



Interstage Big Data Complex Event Processing Server V1.0.0

User's Guide

Linux(64)

J2UL-1665-01ENZ0(00)
October 2012

Preface

Purpose

This manual provides an overview of the features of Interstage Big Data Complex Event Processing Server (hereafter, referred to as "this product"). It also describes the operations required during installation and application development, and the operation and maintenance of this product.

Intended Readers

This manual is intended for users who are considering installing, operating, and developing applications that use the Complex Event Processing feature of this product.

Structure of this Document

This document is structured as follows:

[Chapter 1 Overview](#)

Provides an overview of this product.

[Chapter 2 Features Provided](#)

Describes the features provided by this product.

[Chapter 3 System Configuration and Design](#)

Describes system design for installing this product, such as what kind of information system to build by installing the product, and how to design the operation form.

[Chapter 4 Installation and Setup](#)

Describes the software requirements and resources required to install this product, as well as how to install and uninstall it.

Also explains the setup of this product (that is, how to create an environment for building the system).

[Chapter 5 Development](#)

Describes how to develop applications to run on this product.

[Chapter 6 Operation and Maintenance](#)

Describes how to operate and manage built systems, and how to manage this product.

[Chapter 7 Reliable System Operation](#)

Describes how to operate this product using a reliable configuration.

[Chapter 8 Command Reference](#)

Describes the commands of this product.

[Chapter 9 Definition File Reference](#)

Describes the definition files handled by this product.

[Glossary](#)

Explains the terminology used for this product.

Trademarks

- Adobe, Adobe Reader, and Flash are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.
- Linux is a registered trademark of Linus Torvalds.
- Red Hat, RPM, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

- Microsoft, Windows, MS, MS-DOS, Windows XP, Windows Server, Windows Vista, Windows 7, Excel, and Internet Explorer are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.
- Interstage, PRIMECLUSTER, ServerView, Symfoware, and Systemwalker are registered trademarks of Fujitsu Limited.
- Other company names and product names used in this document are trademarks or registered trademarks of their respective owners.

Note that registration symbols (TM or R) are not appended to system names or product names in this manual.

Export Restriction

This document might contain special technology designated under Foreign Exchange and Foreign Trade Laws. If so, when this document is to be exported or provided overseas, it is necessary to check Foreign Exchange and Foreign Trade Laws, and take the necessary procedure according to these laws.

Copyright

Copyright 2012 FUJITSU LIMITED

October 2012: First edition

Contents

Chapter 1 Overview.....	1
1.1 What is Interstage Big Data Complex Event Processing Server?.....	1
1.2 Product Features.....	2
1.2.1 High-performance CEP Engine.....	2
1.2.2 Simple Rule Description.....	3
1.2.3 Simple Collaboration with External Systems.....	4
1.2.3.1 Distributed Cache Collaboration (XTP Collaboration).....	4
1.2.3.2 Hadoop Collaboration.....	4
1.3 Overview of Features Provided.....	5
1.3.1 Features of the CEP Engine.....	6
1.3.2 Features for Development and Operating Environments.....	7
1.3.3 Features for Status Monitoring.....	7
1.4 What is Complex Event Processing?.....	7
1.4.1 Complex Event Processing.....	7
Chapter 2 Features Provided.....	9
2.1 Input Adapter.....	9
2.1.1 SOAP Adapter.....	10
2.1.2 HTTP Adapter.....	10
2.1.3 Socket Adapter.....	10
2.2 Logging.....	10
2.3 High-speed Filter.....	11
2.3.1 Filter Rules.....	11
2.3.2 Master Data.....	12
2.4 Complex Event Processing.....	12
2.4.1 Features of Complex Event Processing.....	12
2.5 External Data Access.....	13
2.5.1 XTP Collaboration.....	13
2.6 Output Adapter.....	14
2.6.1 SOAP Listener.....	15
2.6.2 Logging Listener.....	15
2.6.3 Debug Log Listener.....	16
2.7 Operation Commands.....	16
2.8 Engine Log.....	16
2.9 Resource Log.....	16
2.10 Cluster Service.....	17
Chapter 3 System Configuration and Design.....	18
3.1 System Configuration.....	18
3.2 Designing the System Configuration.....	18
3.2.1 Designing the System Configuration.....	19
3.2.2 Aspects of Designing the CEP Server.....	19
3.2.2.1 Overall Design.....	19
3.2.2.2 Designing the Input Adapter.....	19
3.2.2.3 Designing the High-speed Filter.....	19
3.2.2.4 Designing Complex Event Processing.....	20
3.2.3 Aspects of Designing a Hadoop System for Collaboration.....	20
3.2.4 Aspects of Designing an XTP Server Node for Collaboration.....	20
3.3 Designing System Resources.....	21
3.3.1 Estimating Memory Usage.....	21
3.3.1.1 Amount of Memory when Using High-speed Filter Rules.....	21
3.3.1.2 Amount of Memory when Master Data is used by the High-speed Filter.....	22
3.3.2 Estimating Disk Usage.....	22
Chapter 4 Installation and Setup.....	25

4.1 Installation Overview.....	25
4.1.1 Installation Methods.....	25
4.1.2 Installed Packages.....	25
4.2 Installation Requirements.....	26
4.2.1 Hardware Environment.....	26
4.2.2 Software Environment.....	26
4.2.2.1 Required Operating System.....	26
4.2.2.2 Mandatory Patch.....	28
4.2.2.3 Mutually Exclusive Software.....	28
4.2.3 Resources Required at Installation.....	28
4.2.4 Resources Required at Operation.....	28
4.3 Installation.....	29
4.3.1 Pre-installation Procedure.....	29
4.3.1.1 Setting /etc/hosts.....	29
4.3.1.2 Checking the Port Numbers to be Used.....	29
4.3.1.3 Checking Free Disk Capacity.....	30
4.3.1.4 Creating the Engine Execution User and Group.....	31
4.3.1.5 Checking Kernel Parameters.....	31
4.3.1.6 Checking Resource Limitations.....	32
4.3.2 Installation Procedure.....	32
4.3.2.1 Attended Installation.....	33
4.3.2.2 Unattended Installation.....	34
4.3.3 Post-installation Procedure.....	34
4.3.3.1 Setting Environment Variables.....	34
4.3.3.2 Applying Updates.....	35
4.3.4 If an Error Occurs during Installation.....	35
4.4 Setup.....	35
4.4.1 Setup Overview.....	35
4.4.2 Setup of Hadoop Collaboration.....	36
4.4.3 Setup of XTP Collaboration.....	36
4.4.4 Setup of the CEP Engine.....	36
4.4.4.1 Status Immediately after Installation.....	37
4.4.4.2 Changing CEP Engine Settings.....	37
4.4.4.3 Creating a New CEP Engine.....	41
4.5 Canceling Setup.....	43
4.5.1 Deleting a CEP Engine.....	43
4.5.2 Canceling XTP Collaboration.....	46
4.5.3 Canceling Hadoop Collaboration.....	46
4.6 Uninstallation.....	47
4.6.1 Pre-uninstallation Procedure.....	48
4.6.1.1 Stopping Event Sending.....	48
4.6.1.2 Backing up User Assets.....	48
4.6.1.3 Stopping the CEP Service.....	48
4.6.1.4 Deleting Updates.....	48
4.6.2 Uninstallation Procedure.....	48
4.6.2.1 Attended Uninstallation.....	48
4.6.2.2 Unattended Uninstallation.....	49
4.6.3 Post-uninstallation Procedure.....	49
4.6.3.1 Uninstalling FJSVod.....	49
4.6.3.2 Uninstalling FJSVmee64 and FJSVsclr64.....	50
4.6.3.3 Removing Environment Variables.....	50
4.6.3.4 Engine Execution User Specified at Installation.....	50
4.6.4 If an Error Occurs during Uninstallation.....	50
Chapter 5 Development.....	52
5.1 Overview of Complex Event Processing of this Product.....	52
5.2 List of Development Assets.....	52

5.3 Task Overview.....	54
5.4 Design.....	55
5.4.1 Association between the Development Asset ID and Definition Information.....	55
5.4.2 Designing an Event Type Definition.....	55
5.4.2.1 Features of Input Events.....	56
5.4.2.2 Recording and Analyzing Events.....	56
5.4.3 Designing a Rule Definition.....	56
5.4.3.1 High-speed Filter Processing.....	57
5.4.3.1.1 Extraction process.....	57
5.4.3.1.2 Extraction process using master data matching.....	58
5.4.3.1.3 Join processing with master data.....	59
5.4.3.1.4 Weighting processing of text.....	60
5.4.3.2 Complex Event Processing.....	61
5.4.3.3 SOAP Listener.....	66
5.4.3.4 Logging Listener.....	66
5.4.4 Designing a Master Definition.....	67
5.4.5 Designing an Event Type Definition (Filtered).....	68
5.4.6 Designing XTP Collaboration.....	68
5.4.6.1 Considerations when Using XTP Collaboration.....	68
5.4.6.2 Using an XTP Cache.....	68
5.4.7 Designing a SOAP Listener Definition.....	70
5.4.8 Event Data (for Testing).....	71
5.4.9 Master Data (for the High-speed Filter).....	71
5.4.9.1 Format of Master Data.....	71
5.4.10 XTP Cache.....	72
5.4.10.1 XTP Cache Compatible Formats.....	72
5.4.11 Designing an Event Sender Application.....	73
5.4.12 Designing a User-developed Web Service.....	73
5.4.13 Designing an Event Log Analysis Application.....	74
5.4.13.1 Output Destination and File Format of an Event Log.....	74
5.4.14 Designing a Cache Access Application for Update.....	76
5.5 Development.....	77
5.5.1 Developing a Definition File.....	77
5.5.1.1 Creating an Event Type Definition File.....	77
5.5.1.2 Creating a Rule Definition File.....	78
5.5.1.2.1 Debug log listener.....	78
5.5.1.3 Creating a Master Definition File.....	79
5.5.1.4 Creating a SOAP Listener Definition File.....	80
5.5.2 Preparing Data.....	80
5.5.2.1 Preparing Event Data (for Testing).....	80
5.5.2.2 Preparing Master Data (for the High-speed Filter).....	80
5.5.2.3 Preparing Data to be Stored in an XTP Cache.....	81
5.5.3 Developing a Collaboration Application.....	81
5.5.3.1 Developing an Event Sender Application.....	81
5.5.3.2 Developing a User-developed Web Service.....	81
5.5.3.2.1 Web service development procedure.....	81
5.5.3.3 Developing an Event Log Analysis Application.....	84
5.5.3.4 Developing a Cache Access Application for Update.....	84
5.6 Deploying Development Assets.....	84
5.6.1 Deploying Definition Information.....	84
5.6.2 Providing Data.....	84
5.6.3 Deploying a Collaboration Application.....	85
5.7 Integration Test.....	86
5.7.1 Integration Test Flow.....	86
5.7.2 Checking an Engine Log.....	87
5.7.3 Starting.....	89
5.7.3.1 Checking the Status of a User-developed Web Service.....	89

5.7.3.2 Starting the CEP Engine.....	89
5.7.3.3 Checking for Syntax Errors in Filter Rules.....	90
5.7.3.4 Checking for Syntax Errors in Complex Event Processing Rules.....	90
5.7.4 Integration Test.....	91
5.7.4.1 Sending Event Data for Testing.....	91
5.7.4.2 Checking the Operation of Filter Rules.....	91
5.7.4.3 Checking the Operation of Complex Event Processing Rules.....	92
5.7.4.4 Checking the Operation of a User-developed Web Service.....	93
5.7.4.5 Checking the Event Log.....	93
5.7.4.6 Checking the Operation of an Event Log Analysis Application.....	93
5.7.5 Stopping.....	94
5.7.5.1 Stopping an Event Sender Application.....	94
5.7.5.2 Stopping the CEP Engine.....	94
5.7.6 Correcting Development Assets.....	94
5.8 Undeploying Development Assets.....	94
5.8.1 Undeploying Definition Information.....	94
5.8.2 Undeploying a Collaboration Application.....	95
5.8.3 Deleting Data.....	95
5.9 Sample Application.....	96
5.9.1 Overview of the Sample Application.....	96
5.9.2 Structure of the Sample.....	96
5.9.3 Events.....	97
5.9.3.1 Location Information Events.....	97
5.9.3.2 Coupon Events.....	97
5.9.3.3 Filtered Location Information Events.....	97
5.9.3.4 Filtered Coupon Events.....	98
5.9.4 Master Information.....	98
5.9.4.1 Member Information Master.....	98
5.9.4.2 Store Information Master.....	98
5.9.5 Rule Definition.....	98
5.9.5.1 Filter Rules (IF-THEN Format).....	98
5.9.5.2 Complex Event Processing Rules (SQL Format).....	99
5.9.6 Event Sender Sample Program.....	100
5.9.7 Directory Structure.....	102
5.9.8 Execution.....	103
5.9.8.1 Deploying Development Assets.....	103
5.9.8.2 Starting the CEP Engine.....	104
5.9.8.3 Sending Events and Checking the Results.....	104
5.9.8.4 Stopping the CEP Engine.....	105
5.9.8.5 Undeploying Development Assets.....	105
Chapter 6 Operation and Maintenance.....	106
6.1 Operating the CEP Server.....	106
6.1.1 Starting the Collaboration System.....	106
6.1.2 Starting the XTP Service.....	107
6.1.3 Starting the CEP Service.....	107
6.1.4 Deploying and Undeploying Definition Information.....	107
6.1.4.1 Deploying Definition Information.....	107
6.1.4.2 Checking Deployed Definition Information.....	108
6.1.4.3 Updating Deployed Definition Information.....	108
6.1.4.4 Undeploying Definition Information.....	109
6.1.5 Starting the CEP Engine.....	109
6.1.6 Typical Operation Tasks.....	110
6.1.6.1 Displaying the Operation Status of the CEP Service.....	110
6.1.6.2 Displaying the Operation Status of the CEP Engine.....	110
6.1.6.3 Monitoring Abnormalities Using Logs.....	111
6.1.6.4 Checking the Resource Usage of the CEP Engine.....	112

6.1.7 Stopping the CEP Engine.....	116
6.1.8 Stopping the CEP Service.....	116
6.1.9 Stopping the XTP Service.....	116
6.1.10 Stopping the Collaboration System.....	117
6.2 Security.....	117
6.2.1 Operation Model.....	117
6.2.2 Prerequisite Knowledge for Designing Security.....	117
6.2.3 Designing Security for this Product.....	119
6.3 Maintenance.....	119
6.3.1 Collecting Data for Investigation when a Problem Occurs.....	119
6.3.2 Backup and Restore.....	119
6.3.2.1 Backup Procedure.....	120
6.3.2.2 Restore Procedure.....	121
6.3.3 Applying Updates.....	121
6.3.4 Tuning.....	122
6.3.4.1 Tuning JVM Options.....	122
6.3.4.2 Tuning File Descriptors.....	125
6.3.4.3 Tuning Trace Logs.....	126
Chapter 7 Reliable System Operation.....	128
7.1 Overview of Reliable System Operation.....	128
7.2 Cluster Service Configuration.....	128
7.3 Building a Cluster Service Environment.....	128
7.4 Operating a Cluster Service.....	130
Chapter 8 Command Reference.....	132
8.1 cepcollectinfo.....	132
8.2 cepconfigeng.....	133
8.3 cepdeployrsc.....	136
8.4 cepdispeng.....	138
8.5 cepdisperv.....	141
8.6 cepgetjvmopt.....	145
8.7 cepgetrsc.....	146
8.8 cepsetjvmopt.....	148
8.9 cepstarteng.....	150
8.10 cepstartserv.....	151
8.11 cepstopeng.....	153
8.12 cepstopserv.....	154
8.13 cepundeployrsc.....	156
Chapter 9 Definition File Reference.....	159
9.1 Defining a CEP Engine.....	159
9.1.1 Engine Configuration File.....	159
9.2 Defining Development Assets.....	161
9.2.1 Event Type Definition File.....	162
9.2.2 Rule Definition File.....	166
9.2.3 Master Definition File.....	168
9.2.4 SOAP Listener Definition File.....	170
9.3 Setting up for Installation.....	171
9.3.1 Installation File.....	171
9.4 Characters Allowed in Item, Tag and Attribute Names.....	172
9.4.1 For High-Speed Filter Rules and Master Definitions.....	172
9.4.2 For Complex Event Processing Rules.....	172
9.5 CSV Format Supported.....	173
Glossary.....	174

Chapter 1 Overview

This chapter provides an overview of the features provided by this product.

1.1 What is Interstage Big Data Complex Event Processing Server?

Interstage Big Data Complex Event Processing Server (hereafter, referred to as "this product") is software that analyzes and assesses massive volumes of event data in real time.

In recent years, there has been a growing demand from companies wanting to use ever-changing event data generated in massive volumes, such as location information sent from smart phones and machines' operation logs, in order to leverage their business activities.

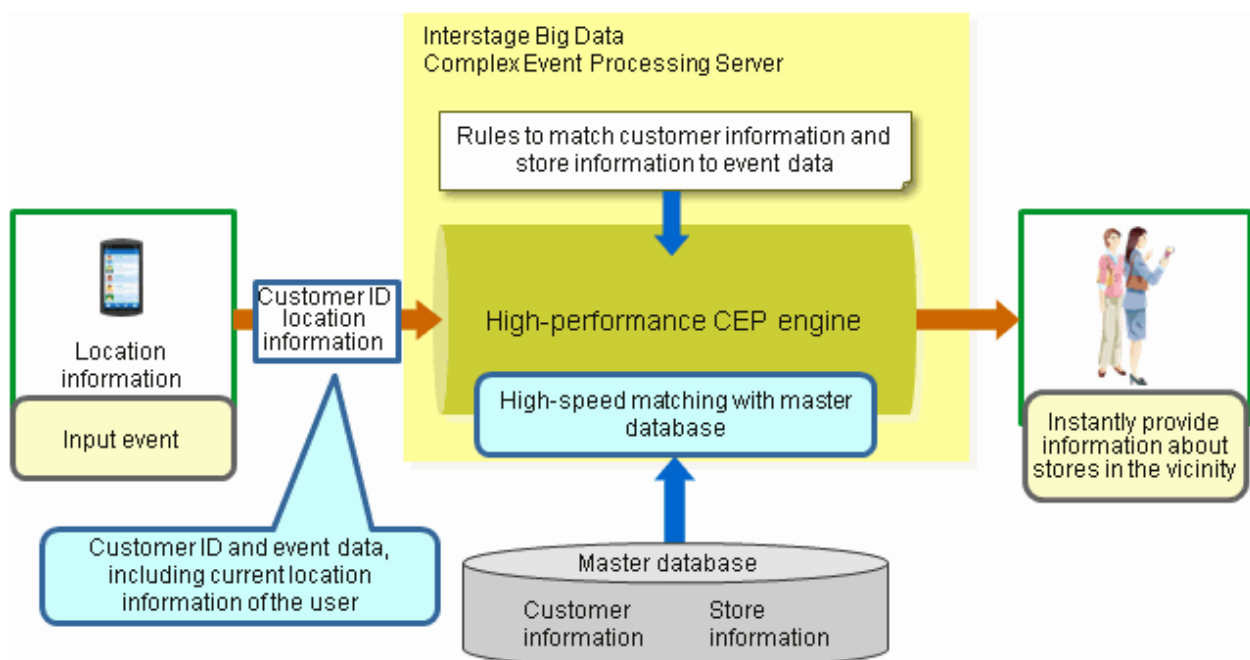
The need to process these kinds of event data in real time has drawn attention to the CEP (Complex Event Processing) technique, which analyzes and assesses massive volumes of data with faster response times than ever before.

The inclusion of a high-performance complex event processing engine (hereafter, referred to as the "high-performance CEP engine"), which integrates a unique high-speed filter processing technique with the complex event processing technique so suitable for processing massive volumes of event data and provides the enhanced processing performance and convenience it needs to support real-time use of massive volumes of event data in corporate systems.

Some scenarios for using this product are described below.

Provide real-time services by utilizing location information

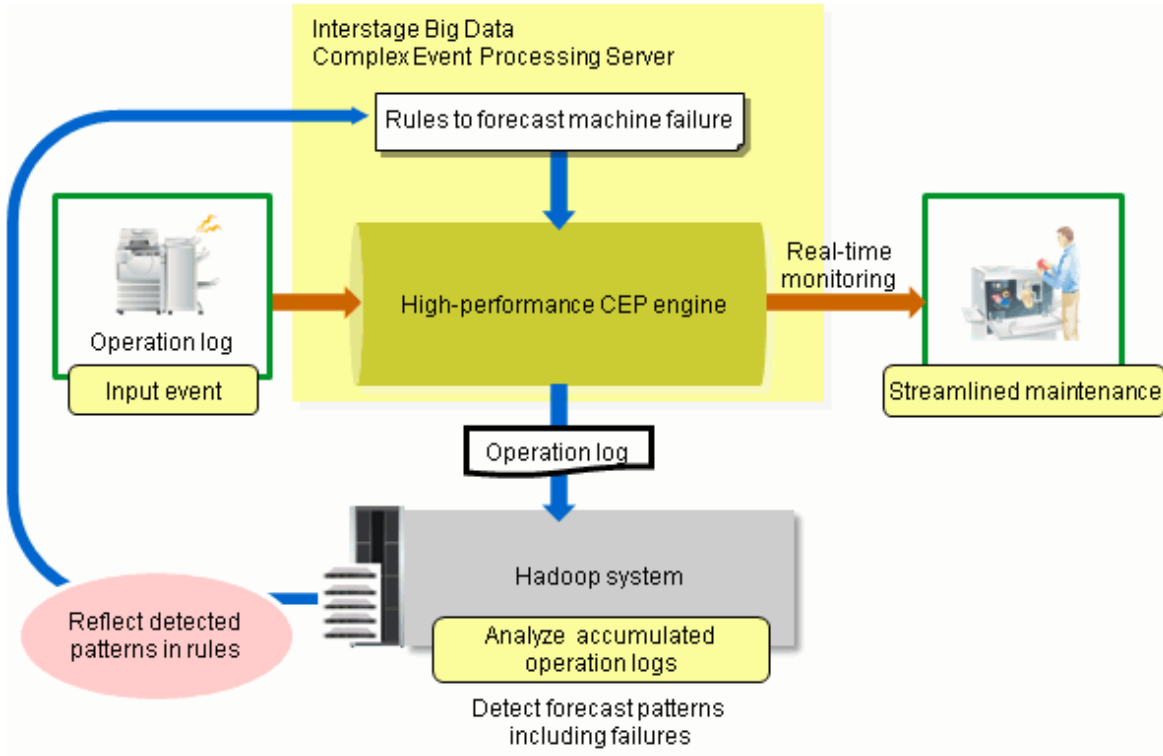
This product allows high-speed matching of real-time customer location information with information registered in the master data, such as customer information and store information. This allows companies to instantly provide services to suit the attributes of customers, such as "provide store coupons to people visiting the vicinity of a store, in real time".



Improve service by monitoring the operation status of sold products

This product allows real-time monitoring the fault prediction of hardware sold to customers, by collecting the operation logs of hardware. This enhances machine availability by allowing preventive maintenance to be performed, which previously may have been impossible in periodic maintenance due to cost or other factors.

The operation logs collected by this product can also be accumulated and analyzed in a Hadoop system, which allows the detection of more refined prediction patterns. Reflecting these patterns in the complex event processing rules allows the implementation of more efficient maintenance services.



1.2 Product Features

This section explains the features of this product, which include the following:

- [High-performance CEP Engine](#)
- [Simple Rule Description](#)
- [Simple Collaboration with External Systems](#)

1.2.1 High-performance CEP Engine

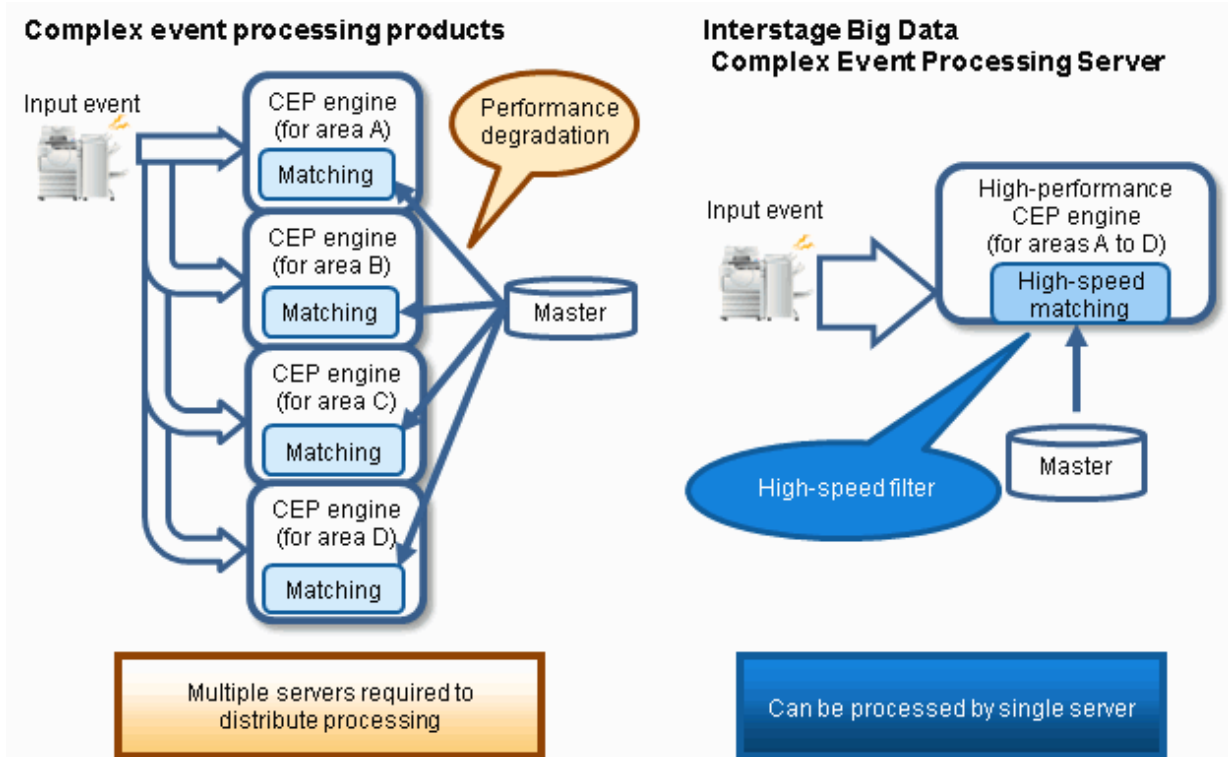
The inclusion of the high-performance CEP engine allows massive volumes of event processing by one server.

In conventional complex event processing products, performance dramatically declines if massive volumes of data are accessed in external master data. To overcome that and ensure performance, multiple servers are provided to distribute processing.

This product uses the unique technique of high-speed filters to allow high-speed matching of input events with master data.

This results in substantially improved performance compared with the conventional complex event processing engine, and allows the required number of events to be processed by one server, with no decline in performance.

Figure 1.1 Comparison when massive volumes of events are processed



1.2.2 Simple Rule Description

Rules must be set in advance in order to execute complex event processing.

This product uses the two description formats below in order to allow rule definitions to be created in a flexible way, according to their purpose.

SQL-type format

This format is based on the database query language SQL.

It is suitable for rules with complex conditional branching or event matching processes.

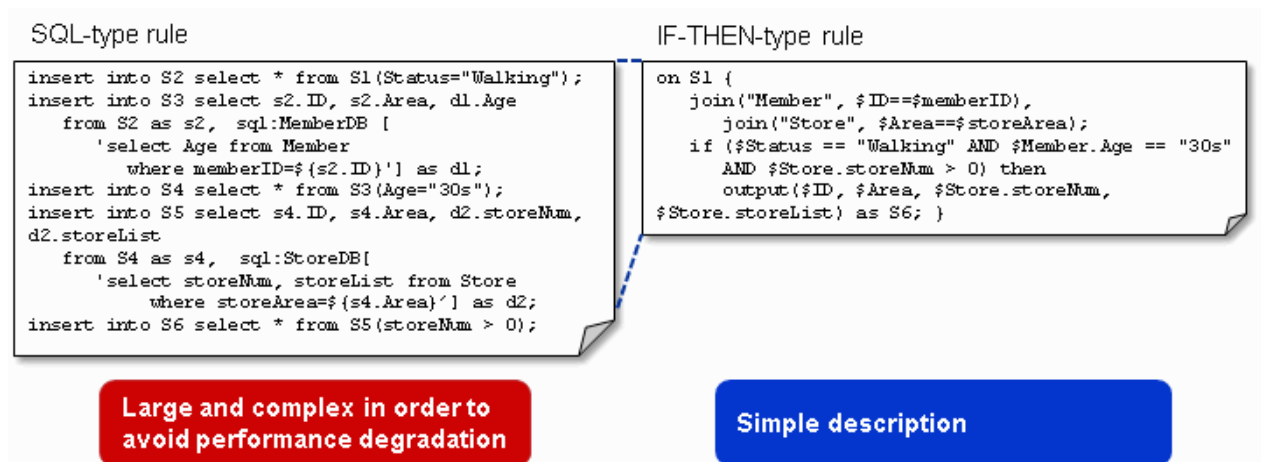
IF-THEN-type format

This format uses the "IF (condition) THEN (process)" structure.

It is suitable for rules that are simpler and easier to understand than SQL-type ones - for example, rules that describe processes such as event filtering, or matching and joining events with master data.

This format also allows simple description of processes that, in conventional SQL-type format, tend to become large and complex, since they need to avoid performance degradation that might be caused by costly join operations.

Figure 1.2 Comparison between an SQL-type rule and an IF-THEN-type rule describing a join with master data



1.2.3 Simple Collaboration with External Systems

1.2.3.1 Distributed Cache Collaboration (XTP Collaboration)

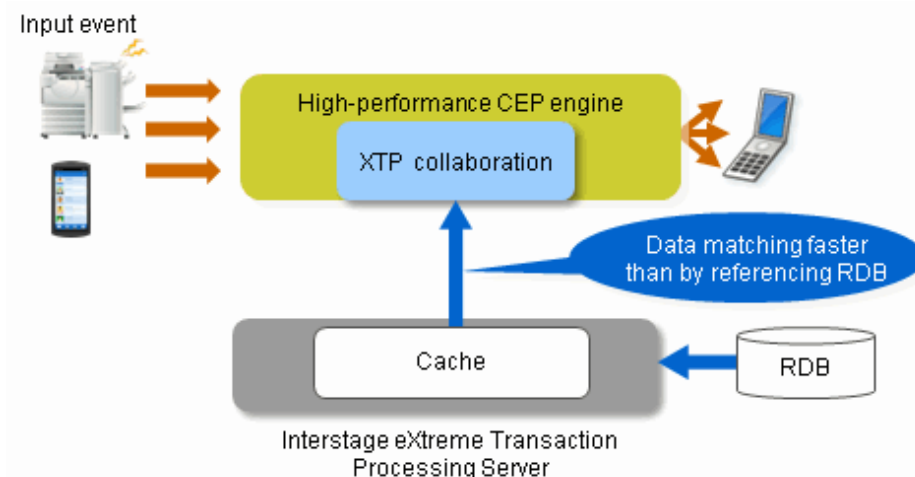
This product can access distributed cache data stored in the extreme transaction processing software Interstage eXtreme Transaction Processing Server (hereafter, referred to as "XTP" - see information below).

This allows external data to be referenced more rapidly than a relational database (RDB).

Information

Interstage eXtreme Transaction Processing Server is an extreme transaction processing software designed to rapidly access massive volumes of a variety of application data and to enhance the scalability and reliability of applications.

Figure 1.3 XTP collaboration



1.2.3.2 Hadoop Collaboration

Collaboration with a Hadoop system allows input events and events output by complex event processing to be accumulated in a Hadoop server.

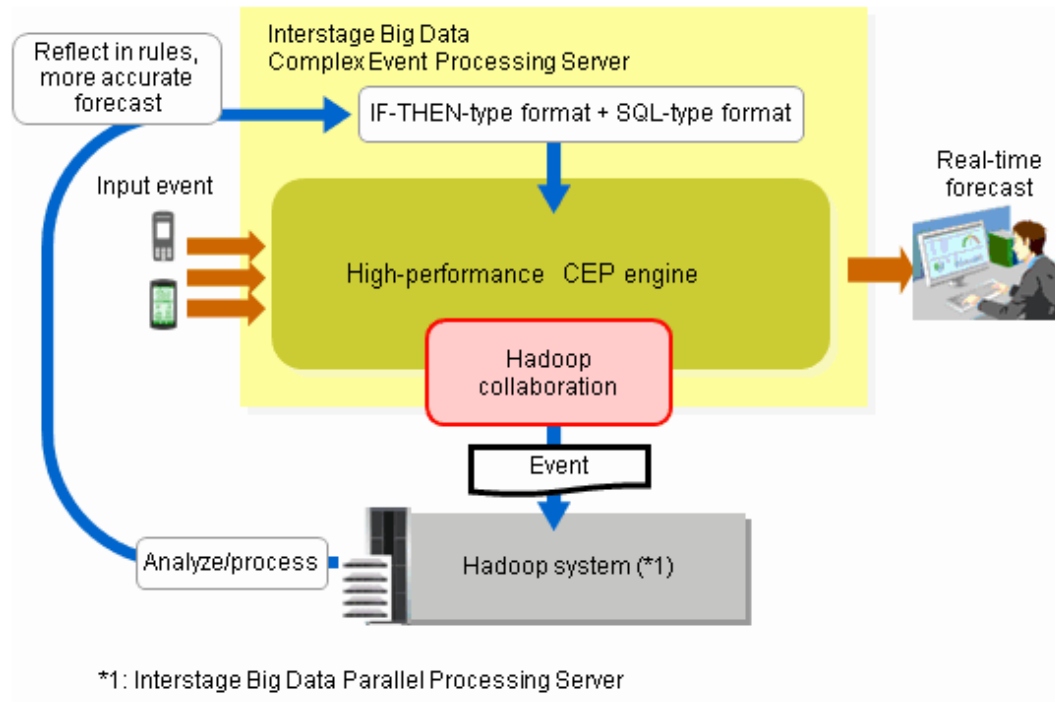
Analysis and processing of the accumulated events by the Hadoop system allows long-term trend analysis of the events, and this can be harnessed for purposes that include business improvement and developing more accurate rules.

The Big Data analysis utilization software known as Interstage Big Data Parallel Processing Server (hereafter, referred to as "BDPP" - see information below) can be used in a Hadoop system.

Information

Based on Apache Hadoop, Interstage Big Data Parallel Processing Server is a Big Data support software on corporate systems that integrates Fujitsu proprietary techniques to further improve processing performance and reliability.

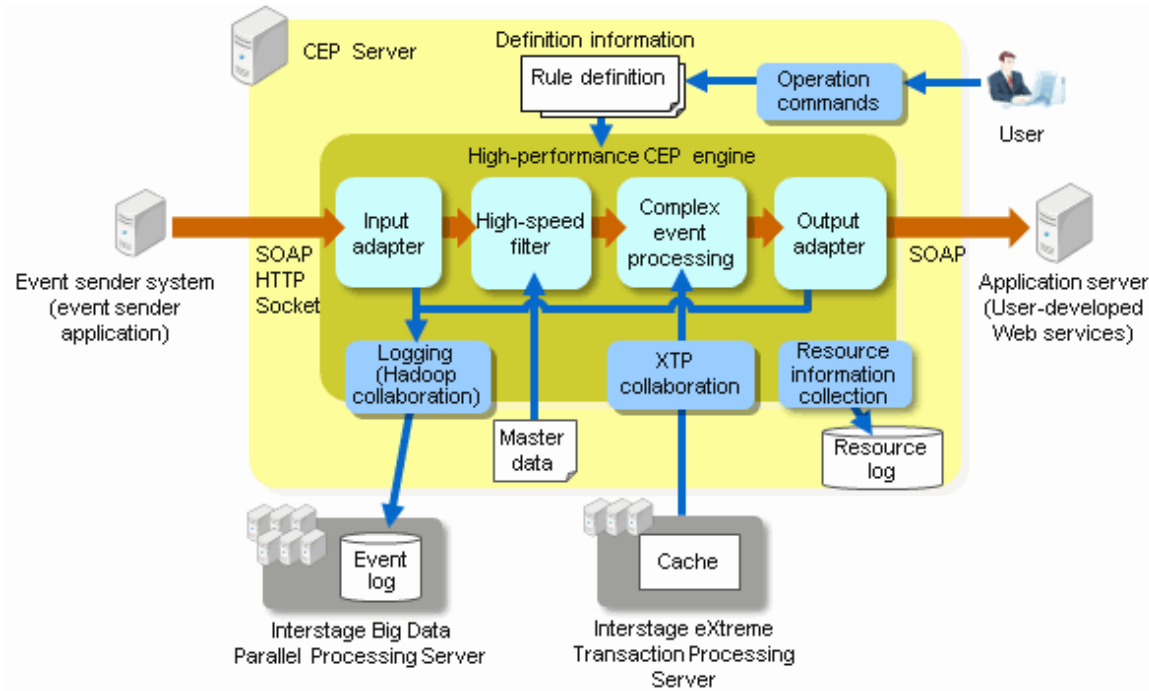
Figure 1.4 Hadoop collaboration



1.3 Overview of Features Provided

This section provides an overview of the features provided by this product (refer to "[Chapter 2 Features Provided](#)" for information on each one).

Figure 1.5 List of features provided



1.3.1 Features of the CEP Engine

Input adapter

Provides three communication methods: SOAP, HTTP, and Socket.

It can connect with input sources such as sensors, smartphones, and SOA systems.

Logging (Hadoop collaboration)

Allows input events and events output by complex event processing to be recorded.

When developing rules, the output destination of logging can be a Hadoop system (Hadoop collaboration), or a CEP Server.

Outputting a log to a Hadoop system allows events to be analyzed in it, and this can be harnessed in monitoring event trends and adding or modifying rules.

High-speed filter

Allows input events to be processed rapidly with the pre-registration of a "rule definition" that describes how to filter input events and perform high-speed matching of events with master data.

Complex event processing

Allows continuously generated events to be analyzed and assessed in real time with the pre-registration of a "rule definition" that describes detection patterns for input events.

External data access

Allows external data to be referenced from complex event processing.

XTP collaboration

Allows cache data stored in XTP to be referenced as external data.

Output adapter

Uses the SOAP listener to allow events output by complex event processing to be sent to a user-developed Web service.

There are also the logging listener (which can be used to store the events output by complex event processing in a Hadoop system), and the debug log listener (which can be used to output the results of complex event processing to a log, for debugging purposes).

1.3.2 Features for Development and Operating Environments

Operation commands

This product provides commands that perform various operations on the CEP engine.

1.3.3 Features for Status Monitoring

Resource log output

This feature logs the CEP engine's resource usage, such as the amount of memory used and the number of input-output events.

The resource log can be used to harness this information in operation and maintenance applications for this product, such as tuning.

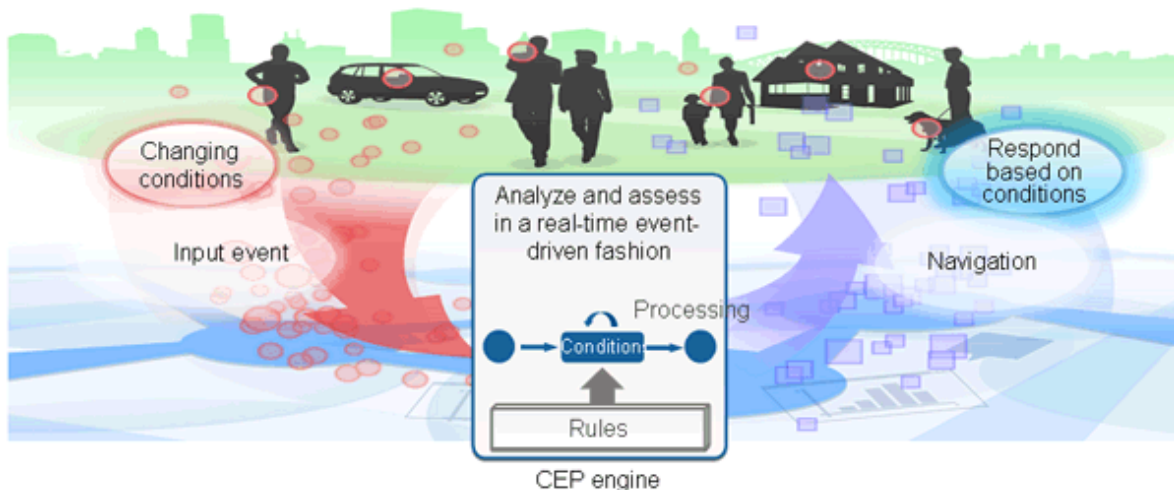
1.4 What is Complex Event Processing?

This section explains complex event processing.

1.4.1 Complex Event Processing

Complex event processing is a technique that analyzes and assesses continuously sent massive volumes of events rapidly and in real time, according to pre-defined rules.

Figure 1.6 Overview of the Complex Event Processing



Complex event processing has the following features:

Real-time processing

In complex event processing, data input from outside is processed in memory as it is, so massive volumes of input data can be processed more rapidly.

As a result, a response can be returned immediately even in ever-changing conditions.

Generally speaking, the throughput and latency (see information below) that can be achieved using complex event processing are as follows:

- Throughput based on rules: Several tens of thousands to several million events/second
- Processing latency: Several microseconds to several milliseconds

Information

In complex event processing, latency is the time elapsed between an event input and its output.

No program development required

The kinds of input events and processes to be processed by complex event processing are described in rules.

No particular program development is required.

By simply changing a rule, you can change the processing content (in order to change the event pattern to be detected, for example).

Chapter 2 Features Provided

This chapter explains the product features.

The features are as follows:

- Main features
 - [Input adapter](#)
 - [Logging](#)
 - [High-speed filter](#)
 - [Complex event processing](#)
 - [External data access](#)
 - [Output adapter](#)
- Operation features
 - [Operation commands](#)
 - [Engine log](#)
 - [Resource log](#)
 - [Cluster service](#)

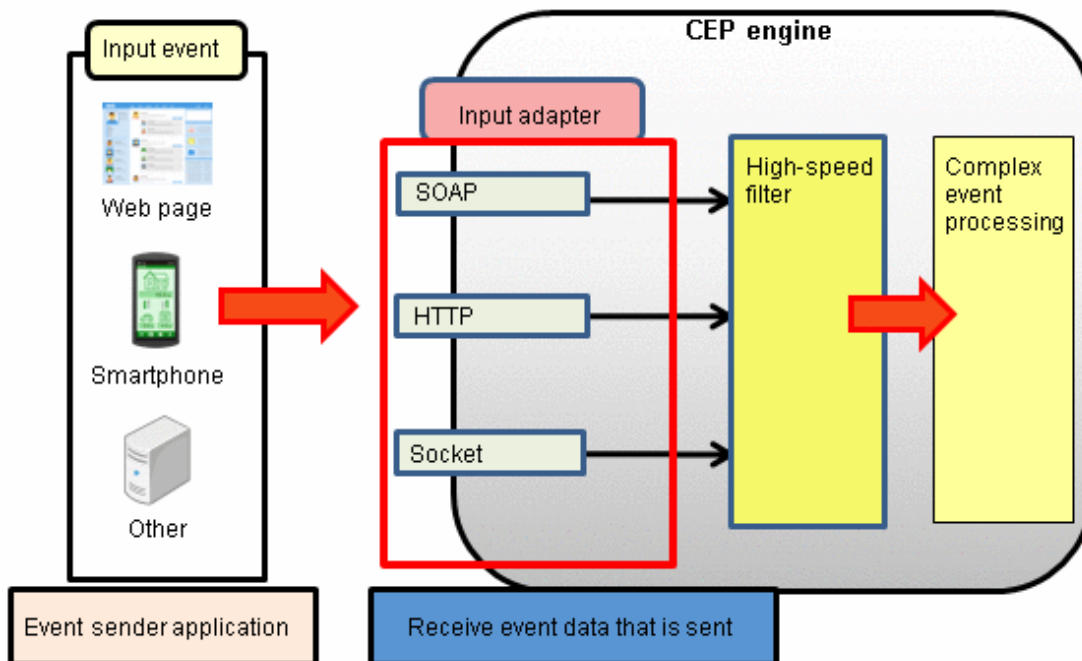
2.1 Input Adapter

The input adapter is a feature that receives events sent from input sources such as sensors, smartphones, and SOA systems.

This product provides the following three input adapters:

- [SOAP adapter](#)
- [HTTP adapter](#)
- [Socket adapter](#)

Figure 2.1 Example of the input adapter



2.1.1 SOAP Adapter

Feature details

The SOAP adapter allows SOAP messages to be received using SOAP communication.

The SOAP adapter extracts event data in XML or CSV format from the received SOAP messages and passes it to the high-speed filter.

Usage scenario

The SOAP adapter is used if event data is to be received from a system that allows SOAP communication, such as an SOA system.

Refer to Chapter 3, "Input Adapter Reference" in the *Developer's Reference* for examples of SOAP messages and for samples of event sender applications.

2.1.2 HTTP Adapter

Feature details

The HTTP adapter allows HTTP requests to be received using HTTP communication.

The HTTP adapter extracts event data in XML or CSV format from the received HTTP requests and passes it to the high-speed filter.

It allows communication that is more lightweight than SOAP communication.

Usage scenario

The HTTP adapter is used if the event sender application is a Web application, or for smartphones.

Refer to Chapter 3, "Input Adapter Reference" in the *Developer's Reference* for information on setting request headers and also for samples of event sender applications.

2.1.3 Socket Adapter

Feature details

The socket adapter allows massive volumes of events to be received at high speed using a communication protocol unique to this product.

The socket adapter passes the received events to high-speed filter processing.

Usage scenario

The socket adapter is used if massive volumes of events need to be processed using high throughput.

Refer to Chapter 3, "Input Adapter Reference" in the *Developer's Reference* for protocol details and also for samples of event sender applications.

2.2 Logging

Logging is a feature that records, as a log, events received from the input adapter prior to high-speed filter processing or events that are the output results of complex event processing.

An Interstage Big Data Parallel Processing Server (Hadoop collaboration) or a CEP Server can be selected as the destination for recording events.

Recording in an Interstage Big Data Parallel Processing Server (Hadoop collaboration)

This selection records events in an event log to be generated in the Interstage Big Data Parallel Processing Server (hereafter, referred to as "BDPP").

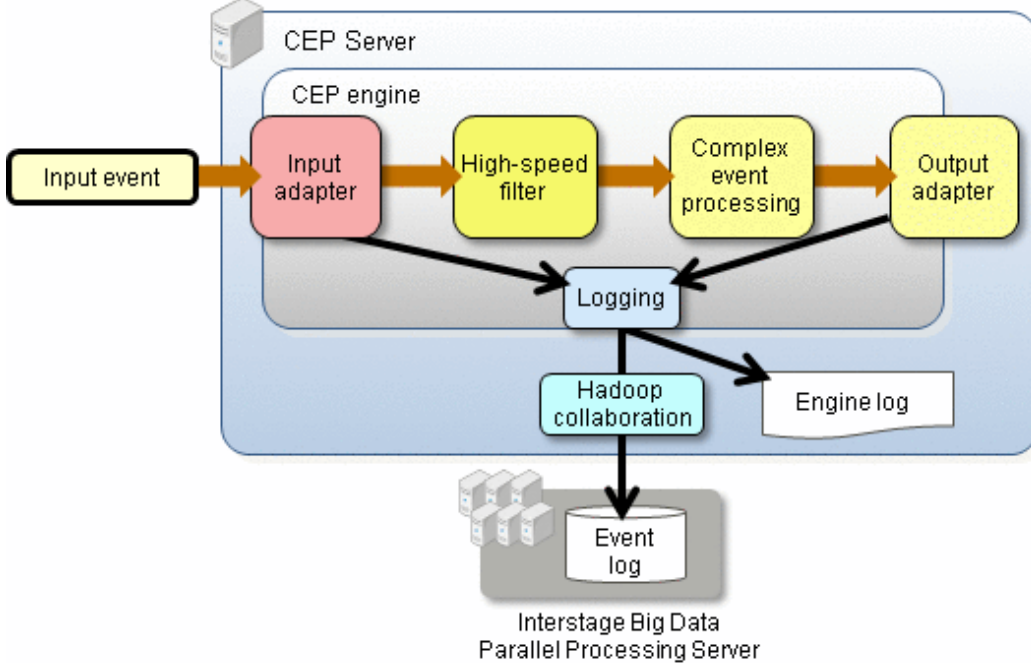
Accumulation, analysis, and processing of the event log can be performed in the BDPP.

The analysis results of the event log can be used for improving the accuracy of the complex event processing rules.

Recording in a CEP Server

This section records events in the [engine log](#) of a CEP Server. The engine log can be used for purposes such as checking event reception and checking rule operation when rules are being developed.

Figure 2.2 Logging



Select the destination for recording events according to your objective.

Refer to "[5.4.13.1 Output Destination and File Format of an Event Log](#)" for information such as the output destination of the event log.

2.3 High-speed Filter

The high-speed filter is a feature that allows the "extraction process of input events" and "join processing of input events with master data" to be performed at high speed.

Events output by the high-speed filter become input events for complex event processing.

This section explains the filter rules used by the high-speed filter as well as master data (files in which data from the master database is stored in CSV format).

2.3.1 Filter Rules

Filter rules define how the "extraction process of events" and "join processing of events with master data" are to be performed.

Input events passed from the input adapter to the high-speed filter are processed based on the filter rules that have been defined.

Typical processes that are performed using filter rules are as follows:

- Extraction process of input events
- Extraction process using master data matching
- Join processing with master data
- Weighting processing of text

Filter rules are described using IF-THEN-type format.

Refer to "[5.4.3.1 High-speed Filter Processing](#)" for information on the typical processes that are performed using filter rules.

Refer to Chapter 2, "Filter Rule Language Reference" in the *Developer's Reference* for information on filter rules.

2.3.2 Master Data

Master data refers to the files in which data from the master database is stored in CSV format.

Master data is used as "join targets" or "extraction conditions of input events" that have been passed to the high-speed filter.

Refer to "[5.4.9 Master Data \(for the High-speed Filter\)](#)" for information on master data.



Master data is loaded in CEP engine memory when the CEP engine starts. Refer to "[3.3.1 Estimating Memory Usage](#)" for information on the required memory usage.

2.4 Complex Event Processing

Complex event processing is a feature that allows any event to be detected and allows "join processing between events" by using pre-registered event processing rules.

This section explains the features in complex event processing.

2.4.1 Features of Complex Event Processing

Complex event processing receives events processed by the high-speed filter and then processes them according to complex event processing rules.

It then passes the events that fall under the rules defined in complex event processing to the output adapter.

Refer to "[2.6 Output Adapter](#)" for information on the output adapter.

Complex event processing rules are described using the SQL-type format (complex event processing language).

Refer to Chapter 1, "Complex Event Processing Language Reference" in the *Developer's Reference* for information on complex event processing language.

The items below are features of Complex Event Processing in this product.

Matching between events

With matching between events, matching is performed between multiple items of event data input from the high-speed filter.

External data access

With external data access, XTP collaboration is used to allow external data to be referenced.

Complex event processing allows external data to be referenced just by defining rules.

Refer to "[2.5.1 XTP Collaboration](#)" for information on the XTP collaboration feature.

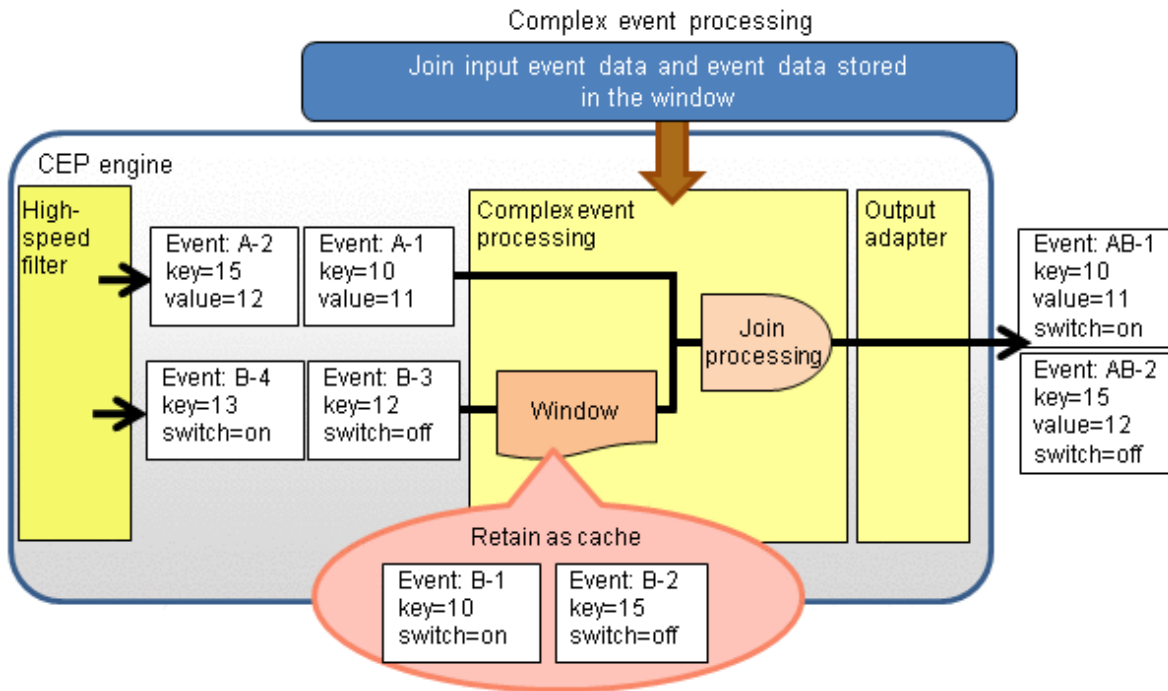
Join processing of cache data and input events

With complex event processing, input events can be retained in the memory cache (feature available in the window).

Join processing of the retained cache data with other input events is also possible.

The example below is an operation image of join processing.

Figure 2.3 Example of join processing of cache data and input events using a window



 Note

- Memory usage must be estimated according to the cache to be retained. Refer to "3.3.1 Estimating Memory Usage" for information on the estimation method.
- If the CEP engine is stopped, the cache retained in memory will be deleted.

2.5 External Data Access

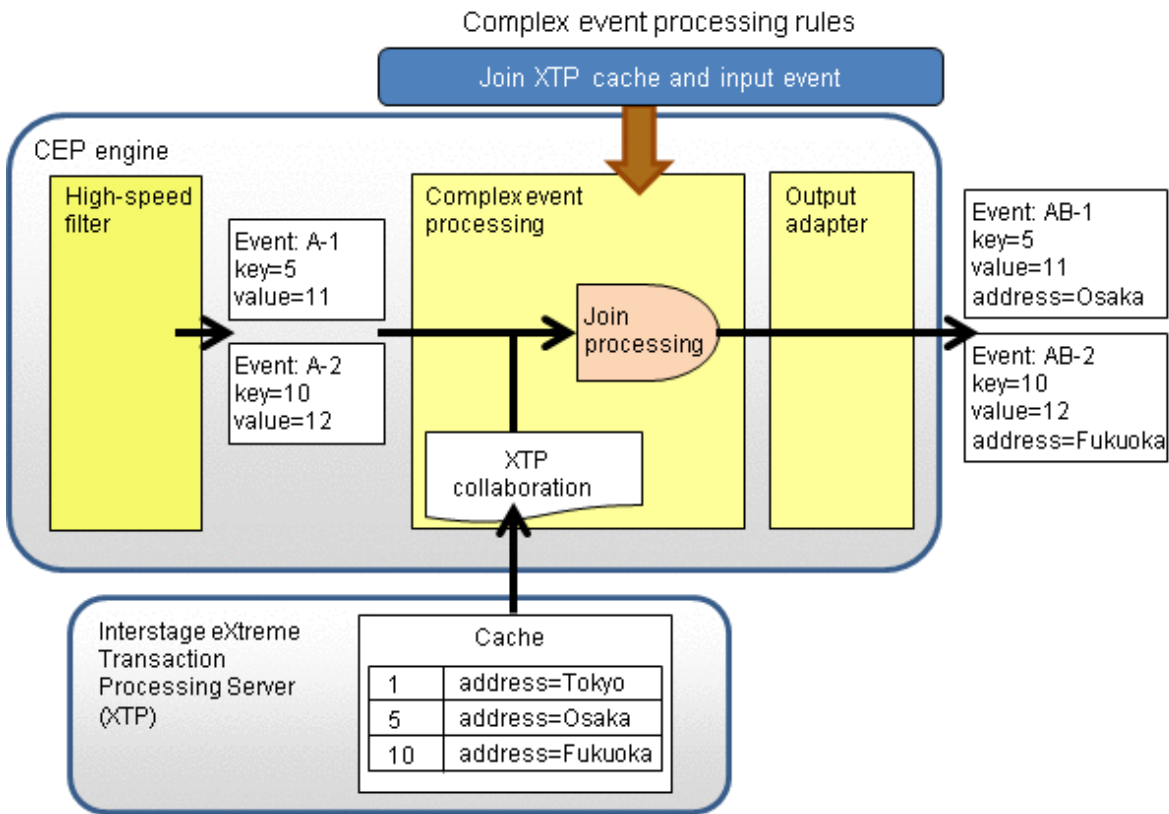
Complex Event Processing allows external data to be referenced.

2.5.1 XTP Collaboration

A cache stored in an external Interstage eXtreme Transaction Processing Server (hereafter, referred to as "XTP") can be referenced from the complex event processing of this product.

In the example below, a cache stored in an XTP is referenced, and the referenced data is joined with input events and then passed to the output adapter.

Figure 2.4 Example of XTP collaboration



Refer to "4.4.3 Setup of XTP Collaboration" for information on how to set up an XTP.

Refer to "5.4.6.2 Using an XTP Cache" for information on how to reference an XTP cache.

Note

In XTP collaboration, only referencing a cache stored in XTP is performed. Updating or deleting a cache from a CEP engine cannot be performed.

2.6 Output Adapter

The output adapter is a feature that outputs the processing results of complex event processing rules externally.

This product provides the three output adapters shown below, according to objective of the user:

Figure 2.5 Output adapter

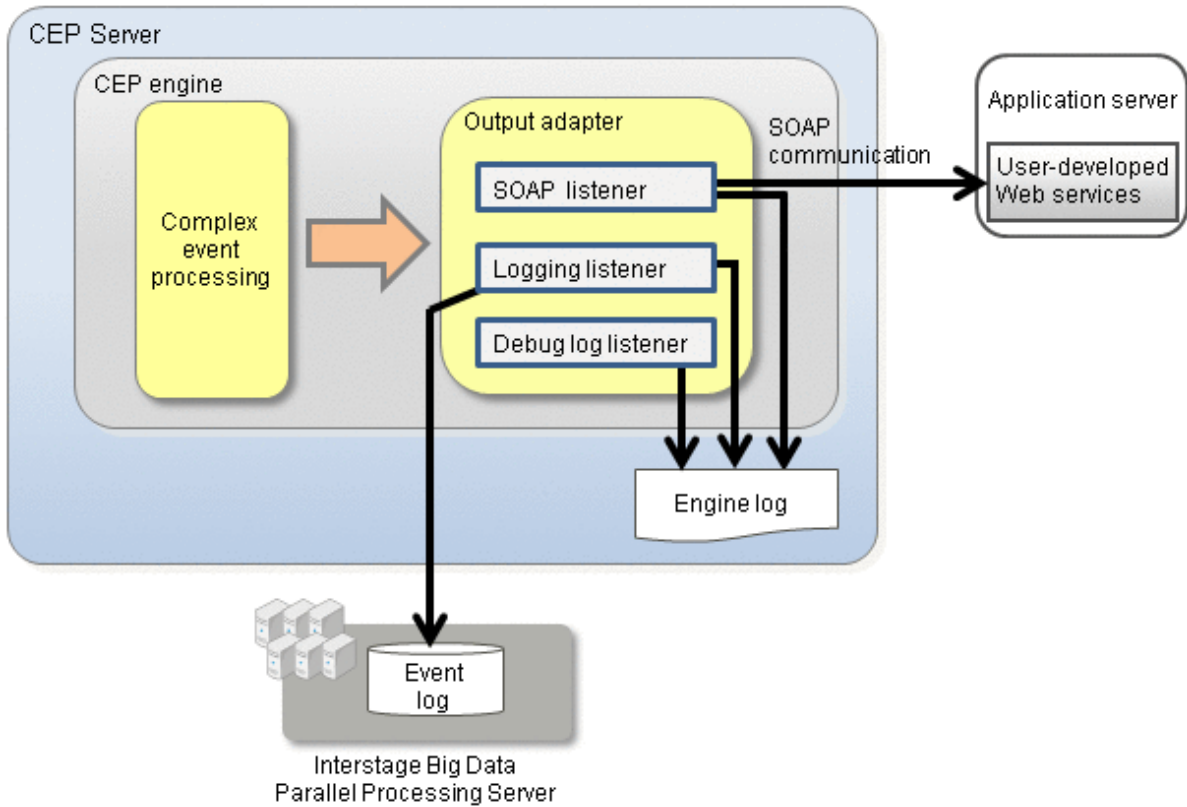


Table 2.1 Overview of the output adapter

Output adapter	Details of feature	Output destination
SOAP listener	Notifies the results of complex event processing (events) to a user-developed Web service, using SOAP communication	<ul style="list-style-type: none"> - User-developed Web service - Engine log (for send records)
Logging listener	Logs the results of complex event processing	<ul style="list-style-type: none"> - Interstage Big Data Parallel Processing Server - Engine log
Debug log listener	Outputs debug information on complex event processing rules	<ul style="list-style-type: none"> - Engine log

2.6.1 SOAP Listener

The SOAP listener sends the results of complex event processing to a user-developed Web service, using SOAP communication.

It also leaves send records in the engine log.

Use the SOAP listener if you want to use the processing results of complex event processing in an external application.

Refer to "[5.4.3.3 SOAP Listener](#)" for information on how to use the SOAP listener.

Also refer to "[5.4.7 Designing a SOAP Listener Definition](#)".

2.6.2 Logging Listener

The logging listener logs the results of complex event processing in the log storage area.

Use the logging listener if you want to analyze the results of complex event processing using a Hadoop system (Interstage Big Data Parallel Processing Server).

Refer to "[5.4.3.4 Logging Listener](#)" for information on how to use the logging listener.

2.6.3 Debug Log Listener

The debug log listener outputs debug information on complex event processing rules to the engine log.

Use the debug log listener if you want to check the operation of complex event processing rules.

Refer to "[5.5.1.2.1 Debug log listener](#)" for information on how to use the debug log listener.



The debug log listener may cause performance to decline, so use it only for development.

2.7 Operation Commands

The operational features of this product are provided using commands.

A list of the commands provided is shown below.

Refer to "[Chapter 8 Command Reference](#)" for information on the commands.

Type	Command name	Command overview
Configuration	cepconfigeng	Creates or deletes a CEP engine
Development	cepdeployrsc	Deploys a development asset
	cepgetrsc	References development assets
	cepundeployrsc	Undeploys a development asset
Operation	cepdispeng	Displays the status of a CEP engine
	cepdisperv	Displays the status of the CEP service (*1)
	cepstarteng	Starts a CEP engine
	cepstartserv	Starts the CEP service (*1)
	cepstopeng	Stops a CEP engine
	cepstopserv	Stops the CEP service (*1)
Maintenance and tuning	cepgetjvmopt	References JVM options
	cepsetjvmopt	Sets JVM options
Troubleshooting	cepcollectinfo	Collects data for investigation in batch

*1: Service that manages the CEP engines on the CEP Server

2.8 Engine Log

Information on errors that have occurred in a CEP engine, such as insufficient memory and logging failures, is output to the engine log.

If some abnormality is detected in a CEP engine, this log can be analyzed in order to identify the cause.

Debug information for complex event processing rules that have been described is also output to the engine log, if the debug log listener has been set in the complex event processing rules.

The engine log is also used by logging as an event recording destination.

Refer to "[6.1.6.3 Monitoring Abnormalities Using Logs](#)" for information such as the output destination of the engine log.

2.9 Resource Log

Resource information for a CEP engine that is collected on a regular basis is output to the resource log.

Analyzing this log allows resource use conditions to be monitored and tuned.

Refer to "[6.1.6.4 Checking the Resource Usage of the CEP Engine](#)" for information such as the output items and output destination of the resource log.

2.10 Cluster Service

The Cluster Service of this product can be used to build a reliable system using PRIMECLUSTER, in order to prevent a long-term suspension of business due to a hardware fault on the CEP Server.

Refer to "[7.1 Overview of Reliable System Operation](#)" for details.

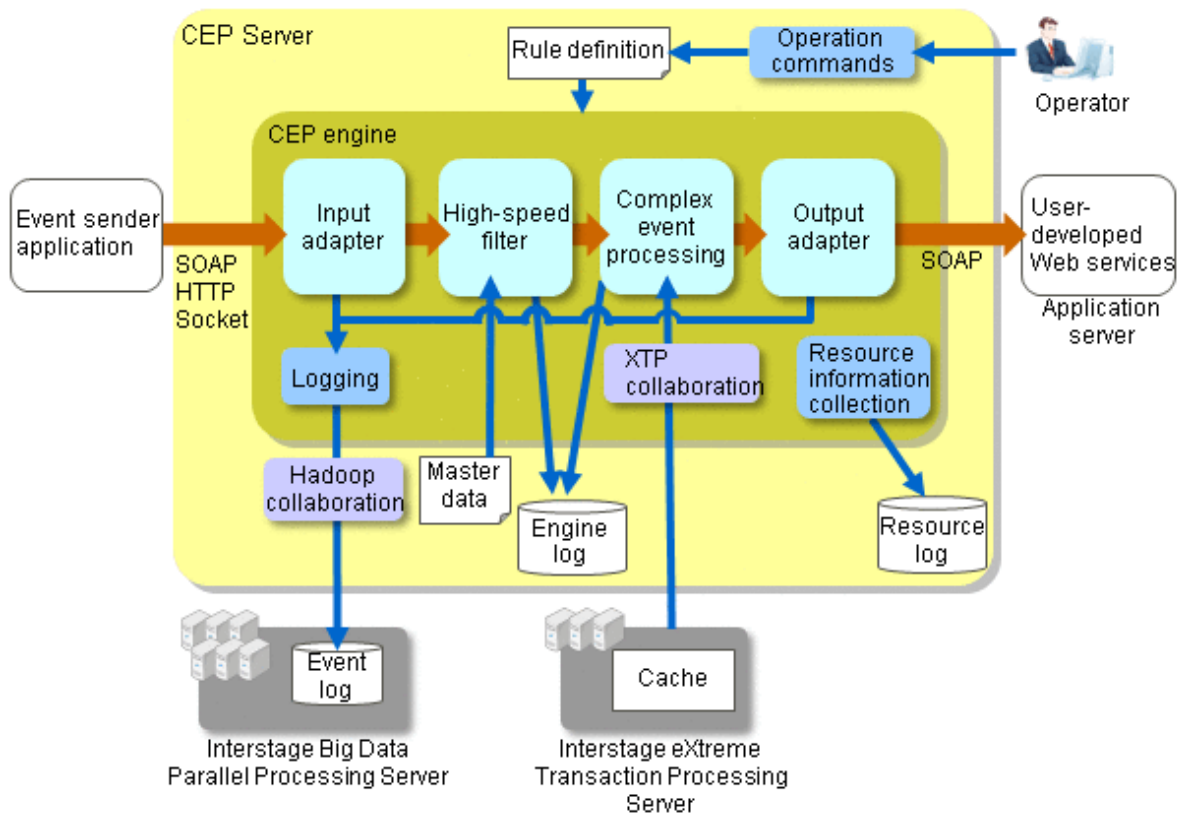
Chapter 3 System Configuration and Design

This chapter explains the system configuration and design of this product.

3.1 System Configuration

A configuration diagram of this product and related products is shown below.

Refer to "[Chapter 2 Features Provided](#)" for information on each feature.



Collaboration features

This product allows collaboration with the following products:

- Product that can be used in Hadoop collaboration
 - Interstage Big Data Parallel Processing Server V1.0.0
- Product that can be used in XTP collaboration
 - Interstage eXtreme Transaction Processing Server V1.0.0

If the collaboration features mentioned above are to be used, a server with the product installed must be provided for each product, in addition to the CEP Server.

Refer to the manual of each product for information on designing the server configuration and also for the installation procedure required for each product.

Refer to "[4.4 Setup](#)" for information on the setup of a CEP Server for collaboration.

3.2 Designing the System Configuration

This section explains designing the system configuration.

3.2.1 Designing the System Configuration

Perform the design tasks in the sections below after identifying which features and collaboration systems will be required, according to the business model and the purpose for which this product is to be used.

Refer to "[3.1 System Configuration](#)" and "[Chapter 2 Features Provided](#)" for information on the system configuration and for details on each feature.

Refer to "[7.2 Cluster Service Configuration](#)" if this product is to be operated using a reliable configuration.

3.2.2 Aspects of Designing the CEP Server

After deciding on the features to be used for each business application, estimate the CEP Server configuration. To achieve a design that meets the estimates requirements use the following features:

- [Overall design](#)
- [Designing the input adapter](#)
- [Designing the high-speed filter](#)
- [Designing complex event processing](#)

The following sections explain the considerations required for the design of each feature. The items considered here are used in "[3.3 Designing System Resources](#)".

3.2.2.1 Overall Design

The main consideration for the overall design of the CEP Server is shown below.

- Number of CEP engines to operate on the CEP Server

3.2.2.2 Designing the Input Adapter

The input adapter design considerations are as follows:

- Number of input event type
- Input event details
 - Data received (per unit of time)
 - Maximum data size
 - Average data size
 - Use of logging
 - Number of items

3.2.2.3 Designing the High-speed Filter

The high-speed filter processing design considerations are as follows:

- Number of high-speed filter statements (IF-THEN statements)
- High-speed filter statement details (IF-THEN statement)
 - Search conditions to be specified in the high-speed filter statement
 - Master data to be used
 - High-speed filter processing event types (those where input events and item content are different)
- High-speed filter processing event details
 - Number of occurrences (per unit of time)
 - Average data size

- Number of items
- Number of master data
- Master data details
 - Number of records
 - Average data size of each item
 - File size

3.2.2.4 Designing Complex Event Processing

The complex event processing design considerations are as follows:

- Use of XTP collaboration
- Number of XTP cache
- Number of user-developed Web service (called from complex event processing rules)
- Number of data type (events or processing results of rules) for complex event processing rules logging
- Details of data to be logged
 - Number of occurrences (per unit of time)
 - Average data size

3.2.3 Aspects of Designing a Hadoop System for Collaboration

The Hadoop system design considerations, when using Hadoop collaboration to perform logging, are listed below.

Based on this information, design storage areas for the required event logs in the Hadoop system. Refer to the Interstage Big Data Parallel Processing Server manuals for information on the design of the Hadoop system.

- Number of data type to be logged (events or processing results of complex event processing rules)
- Details of data to be logged
 - Number of occurrences (per unit of time)
 - Average data size
 - Accumulation period

3.2.4 Aspects of Designing an XTP Server Node for Collaboration

The XTP Server Node design considerations, if XTP collaboration is to be performed, are listed below.

Based on this information, design an XTP Server Node. Refer to the Interstage eXtreme Transaction Processing Server manual for information on the design of the XTP Server Node.

- Maximum number of application communication paths for the server
 - Calculate this by multiplying the number of CEP engines that are to perform XTP collaboration by 2.
- XTP cache collaboration details
 - Number of entries
 - Average data size of entries

3.3 Designing System Resources

3.3.1 Estimating Memory Usage

The formula for calculating the amount of memory to be used by the CEP Server is shown below.

$$\text{Estimated amount of memory required} = A + (B + C + D + E + F + G) \times \text{numberOfCEPengines} \times 1.2$$

Table 3.1 Explanation of items in memory estimation expression

Item	Explanation	Memory usage
<i>A</i>	Base memory amount	2.7 GB
<i>B</i>	Amount of memory when using high-speed filter rules	Refer to "3.3.1.1 Amount of Memory when Using High-speed Filter Rules".
<i>C</i>	Amount of memory when master data is used by the high-speed filter	Refer to "3.3.1.2 Amount of Memory when Master Data is used by the High-speed Filter".
<i>D</i>	Amount of memory when rules are used in complex event processing	The estimation formula is shown below. (MB) $\text{numberOfRuleDefinitions} \times 31 \text{ MB}$
<i>E</i>	Amount of memory required in XTP collaboration	The estimation formula is shown below. (MB) $171 \text{ MB} + 3 \text{ MB} \times \text{numberOfCPUcoresInCEPserver}$
<i>F</i>	Amount of memory when an event type definition is used in complex event processing	The estimation formula is shown below. (MB) $(\text{numberOfInputEventTypes} + \text{numberOfEventTypesThatHaveUndergoneHigh-speedFilterRuleProcessing}) \times 37 \text{ MB}$
<i>G</i>	Amount of memory when a SOAP listener definition is used in complex event processing	The estimation formula is shown below. (MB) $\text{numberOfUser-developedWebServices} \times 2.5 \text{ MB}$

3.3.1.1 Amount of Memory when Using High-speed Filter Rules

The formula for calculating the amount of memory (bytes) when high-speed filter rules are used is shown below.

$$\begin{aligned} \text{Amount of memory when using high-speed filter rules} \\ = \text{Total amount of each } IF-THENstatementMemoryRequirement \end{aligned}$$

$$IF-THENstatementMemoryRequirement = 272 \times 1024 \times 1024 + 960 \times 1024 \times R + 16 \times L + 8 \times a$$

Variable	Meaning	Unit
<i>R</i>	Number of output items to be described in the output expressions of high-speed filter rules	Items
<i>L</i>	Maximum data size of input events	Bytes
<i>a</i>	Area to be used in partial character, character range, and numeric range search (*1)	Bytes

*1:

Use the following formula to calculate this if search is to be executed with a partial character, character range, or numeric range specification:

$$\text{numberOfPartialChars, charRange, or numericRange} \times \text{numberOfKeywordChars} \times 2,048 \text{ Bytes}$$

Use the following formula to calculate this if search is to be executed with a combination of partial character, character range, and numeric range specifications:

$$\text{numberOfPartialChars} \times \text{charRange} \times \text{numericRange} \times \text{numberOfKeywordChars} \times 2,048 \text{ Bytes}$$

- For "number of partial characters", specify the number of parts delimited by a vertical bar (|).
For example, if the search keyword "Sm(ith|ythe|ithy)John" is specified, the number of partial characters will be 3.
- For "character range", specify how many are in the range of the ASCII character code values separated by a hyphen (-).
For example, if the search keyword "class[A-C]" is specified, that range will be 0x41 (A), 0x42 (B), and 0x43 (C), so the character range will be 3.
- For "numeric range", specify how many are in the range of numeric 1 and numeric 2, separated by a comma (.).
For example, if the search keyword "alcohol[9,11]" is specified, that range will be 9, 10, and 11, so the numeric range will be 3.

3.3.1.2 Amount of Memory when Master Data is used by the High-speed Filter

If master data is to be used, the amount of memory capacity shown below will be required in addition to what would be normally required.

```

Memory usage when using master data = Total amount of each IF-THENstatementsMemoryRequirement
IF-THENstatementsMemoryRequirement
    = outputItemMemoryRequirement + joinRelationalExpressionMemoryRequirement
outputItemMemoryRequirement = N x (B + 60)
(If a numeric-type or string-type perfect match is specified)
joinRelationalExpressionMemoryRequirement = N x (216 + A)
(If partial match of a string is specified)
joinRelationalExpressionMemoryRequirement = a x N x (2 x A - logN or 1, whichever is larger) x 144

```

The meaning of each variable is shown below.

Variable	Meaning	Unit
<i>N</i>	Number of master data records	Records
<i>A</i>	Average data size of master items specified in search expressions or in join-relational expressions in join expressions (*1)	Bytes
<i>B</i>	Average data size of master items specified in the output items of output expressions (*1)	Bytes
<i>a</i>	Join key coefficient (*2)	0 < <i>a</i> < 1

*1: If the "val" function is specified in a join-relational expression and an output item, the data size will be 16.

*2: This depends on the content of the master data specified in the join conditions in the high-speed filter rules (see the table below).

Table 3.2 Join key coefficient

Content of master data	Join key coefficient
If values vary widely in the second half of the key Example: (000001, 000002, 000012, 000125, etc.)	0.4
If values vary widely in the first half of the key Example: (100-001, 210-001, 321-001, etc.)	0.6
If values vary widely throughout the key Example: (123456, 234512, 912384, etc.)	0.8

3.3.2 Estimating Disk Usage

The formula for calculating the disk usage required by the CEP Server is shown below.

Disk usage required = (A + B + C + D + E + F + G + H) x 1.2

Table 3.3 Explanation of items in disk usage estimation expression

Item	Explanation	Directory	Estimated disk usage
A	Base disk usage	The directories are as follows: /opt /var/opt /etc/opt	The estimated disk usage is as follows: /opt: 860 MB /var/opt: 53 MB /etc/opt: 20 MB
B	Event log (Before the high-speed filter is used)	Refer to "5.4.13.1 Output Destination and File Format of an Event Log".	(*1)
C	Event log (After complex event processing)	Refer to "5.4.13.1 Output Destination and File Format of an Event Log".	(*2)
D	Resource log	High-speed filter: /var/opt/FJSVcep/cep/flt/logs/ ResourceLog/ <i>CEPengineName</i> Complex event processing: /var/opt/FJSVcep/cep/cep/logs/ ResourceLog/ <i>CEPengineName</i>	The estimation expression is shown below. (MB) $2 \text{ MB} \times \text{numberOfCEPengines}$
E	Engine log	High-speed filter: /var/opt/FJSVcep/cep/flt/logs/EngineLog/ <i>CEPengineName</i> Complex event processing: /var/opt/FJSVcep/cep/cep/logs/EngineLog/ <i>CEPengineName</i>	The estimation expression is shown below. (MB) $200 \text{ MB} \times \text{numberOfCEPengines}$
F	Master data	This will be the path specified in "dataFile" in the master definition file. Refer to "9.2.3 Master Definition File" for details.	This will be the master data file size.
G	Maintenance log of the high-speed filter	/var/opt/FJSVisjee/nodeagents/ijna/ <i>CEPengineName</i> _Flt_Ins/current	The estimation expression is shown below. (MB) $(\text{totalNumberOfFilterStatementsDescribedInHigh-speedFilterRules} \times 1.6 + 0.2) \times 4 \text{ MB}$
H	Other maintenance logs (except for the above logs)	Logs in the following directories: /var/opt/FJSVisjee /var/opt/FJSVjs2su /var/opt/FJSVcep /var/opt/FJSVihs /var/opt/FJSVjs5 /var/opt/FJSVj2ee /var/opt/FJSVisjmx /var/opt/FJSVisas /var/opt/FJSVod	The estimation expression is shown below. (MB) $1301 \text{ MB} + 618 \text{ MB} \times \text{numberOfCEPengines}$

*1: Calculate the events to be logged by the input adapter. The formula is shown below.

Disk usage of event log (B) (KB) = Total amount of each *diskUsageOfEventsToBeLogged*

diskUsageOfEventsToBeLogged (KB)

= *numberOfEventsReceivedPerSecond* x *averageDataSize* (KB) x *eventLogAccumulationPeriod* (Seconds) x 1.2

*2: Calculate the data (events or processing results of rules) to be logged by complex event processing. The formula is shown below.

Disk usage of event log (C) (KB) = Total amount of each *diskUsageOfDataToBeLogged*

diskUsageOfDataToBeLogged (KB)

= *numberOfDataOccurrencesPerSecond* x *averageDataSize* (KB) x *eventLogAccumulationPeriod* (Seconds) x 1.2

Chapter 4 Installation and Setup

This chapter explains how to install, set up and uninstall this product.

Refer to "[7.3 Building a Cluster Service Environment](#)" before operating this product in a high-availability (failover cluster) environment.

4.1 Installation Overview

This section provides an overview of the installation of this product.

4.1.1 Installation Methods

This product is installed using shell scripts. The following installation methods are available:

Attended installation

Use 'attended installation' to execute the installation according to your specific requirements such as the engine execution user name.

Unattended installation

Use 'unattended installation' to execute the installation according to a setup file specified when the Installer starts, with no querying from the Installer.

4.1.2 Installed Packages

Below is a list of the packages installed by this product.

Table 4.1 List of packages

Type	Package name
Basic features of this product	FJSVcep
	FJSVisco
	FJSVisas
	FJSVtdis
	FJSVextp
	FJSVjdk6
	FJSVisjee
	FJSVsclr
	FJSVsmee
	FJSVisscs
	FJSVxmlpc
	FJSVjdk5
	FJSVtd
	FJSVod
	FJSVjs2su
	FJSVihs
	FJSVisjmx
	FJSVejb
FJSVj2ee	
FJSVjs5	

Type	Package name
	FJSVisgui

Note

If packages provided by this product are installed or uninstalled directly, for example by using rpm, they will not operate normally. Unless directed to do otherwise by Fujitsu technical support, always use shell scripts to install and uninstall.

4.2 Installation Requirements

This section explains the resources required for installation.

4.2.1 Hardware Environment

The hardware below is required in order to use this product:

- PRIMERGY RX series or PRIMERGY TX series

This product also requires an environment with sufficient memory available.

Refer to "[Chapter 3 System Configuration and Design](#)" for information on estimating the memory size.

4.2.2 Software Environment

The software below is required in order to use this product.

4.2.2.1 Required Operating System

Either of the operating systems below is required.

Operating system name	Remarks
Red Hat Enterprise Linux 5 (for Intel64)	This operating system supports operation using version 5.3 or later.
Red Hat Enterprise Linux 6 (for Intel64)	This operating system supports operation using RHSA-2010:0842 (kernel-2.6.32-71.7.1.el6) or later. Version 6.1 or later has RHSA-2010:0842 applied.

Note

This product is guaranteed to operate in an environment where the SELinux function is disabled.

Information

If this product is to operate on Red Hat Enterprise Linux 6 (for Intel64), use the packages below in addition to the packages installed with the minimum operating system option.

Package	Architecture
alsa-lib	x86_64
cloog-ppl	x86_64
compat-libtermcap	i686
compat-readline5	i686

Package	Architecture
cpp	x86_64
file	x86_64
gcc	x86_64
gcc-c++	x86_64
gdb	x86_64
glibc	i686
glibc-devel	x86_64
glibc-headers	x86_64
kernel-headers	x86_64
libICE	x86_64
libSM	x86_64
libX11	x86_64
libX11-common	noarch
libXau	x86_64
libXext	x86_64
libXi	x86_64
libXp	x86_64
libXt	x86_64
libXtst	x86_64
libgomp	x86_64
libstdc++-devel	x86_64
libtool-ltdl	x86_64
libxcb	x86_64
make	x86_64
mpfr	x86_64
ncurses-libs	i686
nss-softokn-freebl	i686
perl	x86_64
perl-Module-Pluggable	x86_64
perl-Pod-Escapes	x86_64
perl-Pod-Simple	x86_64
perl-libs	x86_64
perl-version	x86_64
ppl	x86_64
strace	x86_64
tcsh	x86_64
unixODBC	x86_64
zlib	i686

4.2.2.2 Mandatory Patch

The patch below must be installed in advance.

Item No.	Operating system	Patch ID and batch update	Remarks
1	Red Hat Enterprise Linux 6 (for Intel64)	RHBA-2011:0321-1	

4.2.2.3 Mutually Exclusive Software

Do not install the software or packages below on the same system as this product.

Product name	Version	Remarks
Interstage Big Data Complex Event Processing Server	All versions	No more than one instance can be installed on the same operating system.
Interstage Application Server	All versions	(*1)
Interstage Web Server	All versions	(*1)
Interstage Business Application Server	All versions	(*1)
Interstage Service Integrator	All versions	(*1)
Systemwalker Availability View	All versions	(*1)
Systemwalker Centric Manager	All versions	The Job Server can be installed on the same system.
Interstage Big Data Parallel Processing Server	All versions	The Development Server can be installed on the same system.

*1: In Red Hat Enterprise Linux 5 (for Intel64) or Red Hat Enterprise Linux 6 (for Intel64), the product cannot be installed on the same system even if the product supports operation in 32-bit mode.

4.2.3 Resources Required at Installation

The resources below are required for installation of this product.

Disk capacity

The disk capacity below is required for installation of this product. If necessary, extend the size of the relevant file system.

Directory	Required disk capacity
/opt	860MB
/etc/opt	20MB
/var/opt	53MB

Memory capacity

The memory capacity below is required for installation of this product. If necessary, extend the amount of memory installed.

- 2.7GB

4.2.4 Resources Required at Operation

To operate this product, estimate the following resources and allocate the required capacities:

- Memory
- Disk

Refer to "[Chapter 3 System Configuration and Design](#)" for information on estimating resources.

4.3 Installation

This section explains how to install this product.

4.3.1 Pre-installation Procedure

This section explains the tasks required before installing this product.

- [Setting /etc/hosts](#)
- [Checking the Port Numbers to be Used](#)
- [Checking Free Disk Capacity](#)
- [Creating the Engine Execution User and Group](#)
- [Checking Kernel Parameters](#)
- [Checking Resource Limitations](#)

4.3.1.1 Setting /etc/hosts

Configure the /etc/hosts file so that network name resolution is enabled for the CEP Server host name (*1).

When registering the CEP Server host name as "127.0.0.1" (CEP Server loopback address), check the following points:

- When setting the host name of the CEP Server for "127.0.0.1", always first describe the IP address setting used when gaining access from outside the CEP Server.
- Alternatively, do not set the host name of the CEP Server for "127.0.0.1".

*1: The HOSTNAME parameter setting in the /etc/sysconfig/network file.



Example

Example of setting /etc/hosts

Below is an example of setting the host name "cepsv1" for "127.0.0.1". In this example, the IP address setting used when gaining access from outside the CEP Server is "10.10.10.10".

```
10.10.10.10 cepsv1
127.0.0.1 cepsv1 localhost.localdomain localhost
```

4.3.1.2 Checking the Port Numbers to be Used

Check that the port numbers to be used by this product are available for use. The port numbers used by this product are shown below. Use firewall or operating system settings to ensure the relevant port can be used.

Port number	Description
80	Port number used by the input adapter (SOAP adapter or HTTP adapter) to receive input events.
81	Port number used internally by this product.
102	Access from outside the CEP Server is not required.
389	
636	
2000	
2465	
3279	

Port number	Description
4433	
5432	
6666	
8002	
8009	
8080	
8686	
8909	
8919	
9700	
10550	
10555	
12000	
12001	
12200	
12210	
12220	
12230	
13000	
23600	
23601	
23602	
23700-23710	
28080	
28090-28100	
28686-28696	
<i>anyPort</i>	<p>Port number used by the input adapter known as the socket adapter to receive input events.</p> <p>Use this by setting one unused port number between 5001 and 32767 for each CEP engine, and by setting a maximum of five for the entire CEP Server.</p> <p>Setting the port numbers to be used is done in CEP engine setup after installation.</p> <p>Port number 9600 is set in the CEP engine created at initial setup.</p>

4.3.1.3 Checking Free Disk Capacity

Check that the disks have sufficient free capacity. Refer to "[4.2.3 Resources Required at Installation](#)" for information on the required disk capacity.

If there is a shortage of free disk capacity, extend the size.

4.3.1.4 Creating the Engine Execution User and Group

Create a specific user and group to execute the CEP engine. Each of the processes of the CEP engine run with the user and group permissions created here.

Example

Operation example of creating the engine execution user and group using the user name "isbdcep" and the group name "isbdcep"

```
$ su -<ENTER>
# /usr/sbin/groupadd isbdcep<ENTER>
# /usr/sbin/useradd -g isbdcep isbdcep<ENTER>
```

Note

- The user creation method depends upon the management policy of the system. Always check with the system administrator.
- Up to 8 characters can be specified for the user name and the group name.

4.3.1.5 Checking Kernel Parameters

Kernel parameters must be tuned in advance when operating this product.

Edit `/etc/sysctl.conf` to change the values of the target kernel parameters to suitable values, according to the parameter "type".

- If the type is "Maximum":

If the value already set (initial value or previously set value) is greater than the value shown in the table, it need not be changed. If it is smaller than the value in the table, change it to the value in the table.

- If the type is "Additional"

Add the value shown in the table to the value that is already set (initial value or previously set value). Check the system maximum values before adding this value and, if adding that value would exceed the system maximum value, set the system maximum value.

The current kernel parameters can be verified using `/sbin/sysctl -a`.

After making the changes, execute `/sbin/sysctl -p /etc/sysctl.conf` or reboot the OS.

Refer to the documentation for the operating system for information on how to change the kernel parameters.

Below are the kernel parameters to be set.

Shared memory

Parameter	Description	Value to be set	Type
kernel.shmmax	Maximum size in shared memory	57413492	Maximum
kernel.shmmni	Maximum number of shared memory segments	41	Additional

Semaphore

For semaphore settings, set the values for each parameter in the following format:

```
kernel.sem = SEMMSL SEMMNS SEMOPM SEMMNI
```

Parameter	Description	Value to be set	Type
<i>SEMMSL</i>	Maximum number of semaphores for each semaphore identifier	512	Maximum

Parameter	Description	Value to be set	Type
<i>SEMMNS</i>	Number of semaphores for the system as a whole	5763	Additional
<i>SEMOPM</i>	Maximum number of operators for each semaphore call	50	Maximum
<i>SEMMNI</i>	Number of semaphore operators for the system as a whole	1143	Additional

Message queue

Parameter	Description	Value to be set	Type
kernel.msgmax	Maximum message size	16384	Maximum
kernel.msgmnb	Maximum number of bytes of messages in the message queue	32768	Maximum
kernel.msgmni	Maximum number of message queue IDs	526	Additional

4.3.1.6 Checking Resource Limitations

When operating this product, adjust the user limitations for the number of processes (threads) that can be executed.

Edit `/etc/security/limits.conf` and change the number of processes (threads) that the user can execute to an appropriate value.

For this product, set the number of processes (threads) the engine execution user can execute to "2048" or higher. Add "2048" to the already specified value (the initial or previously specified value).

The number of processes (threads) that the engine execution user can execute can be checked with the following command as a superuser.

```
# /bin/su -c 'ulimit -u' engineExecutionUser <ENTER>
```

Reboot the OS after changing the value.

Refer to the OS manual for information on how to change the value.



Example

Example for setting `/etc/security/limits.conf`

This is an example for specifying the number of processes (threads) an engine execution user can execute. In this example, the value is set by adding "2048" to the default value of "1024".

```
isbdcep          soft   nproc           3072
```

4.3.2 Installation Procedure

From the following two types, select the most suitable installation method. If installing using multi-user mode, check that the operations of other users will not affect the installation:

- [Attended installation](#)
- [Unattended installation](#)



Note

This product cannot be installed in an environment where it has already been installed.

4.3.2.1 Attended Installation

1. Log in as a superuser.

```
$ su -<ENTER>
```

2. Load the installation DVD-ROM and execute the installation shell script (install.sh) stored on the DVD-ROM from any directory.

```
# mount /dev/deviceFileName DVD-ROMmountDir <ENTER>
# DVD-ROMmountDir/install.sh <ENTER>
```

3. The product name will be displayed as shown below.

```
+-----+
| Interstage Big Data Complex Event Processing Server V1.0.0 |
|                                                               |
|                               Copyright 2012 FUJITSU LIMITED |
+-----+
```

4. In the interactive process shown below, which continues from the display above, specify the engine execution user and the name of the group to which the engine execution user belongs, which were created in advance.

Only a user or group that has already been created can be specified. You can stop installation by typing "q" and pressing the Enter key.

Below is an input example where "isbdcep" is specified as the user name and "isbdcep" is specified as the group name.

```
Please specify the engine execution user and group.

CEP engine processes will run as the specified user and group.
It differs who performs operational commands.

Please enter the user name of the engine execution user [q]: isbdcep<ENTER>

Please enter the group name of the engine execution user [q]: isbdcep<ENTER>
```

5. Then specify the engine name to be created at initial setup. You can stop installation by typing "q" and pressing the Enter key.

```
Please enter the initial engine's name (default: CepEngine) [q]: CepEngine<ENTER>
```

6. The installation details will be displayed.

In *installPackages*, all of the packages to be installed will be displayed delimited by spaces.

```
Engine execution user (group):
  specifiedEngineExecutionUserName (specifiedEngineExecutionUserGroupName)
Initial engine's name:
  specifiedEngineName
Installation packages:
  installPackages
```

7. Check the content and, to start installing the packages, type "y" and press the Enter key.

To stop installation, type "q" and press the Enter key.

```
Do you want to proceed with the installation? [y,q]:y<ENTER>
```

8. When installation has completed normally, the message below will be output.

```
The installation processing completed successfully.
```

9. After installation has completed normally, reboot the OS. Perform the [post-installation procedure](#) to continue.

```
# shutdown -r now <ENTER>
```

4.3.2.2 Unattended Installation

1. Create an installation file. Create the file after designing and checking the parameters to be specified in advance. Refer to "9.3.1 Installation File" for information on the file.

Example

Below is an example of the file.

In this example, "isbdcep" is specified as the engine execution user, "isbdcep" is specified as the group to which the engine execution user belongs, and "CepEngine" is specified as the name of the CEP engine to be created at initial setup.

```
BDCEP_USER_NAME=isbdcep
BDCEP_GROUP_NAME=isbdcep
BDCEP_INITIAL_ENGINE_NAME=CepEngine
```

Point

A sample of the installation file is stored in "/samples/bdcep.conf" on the DVD-ROM of this product. To work efficiently, copy the sample to the work directory of the system and then edit the parameters.

2. Log in as a superuser.

```
$ su -<ENTER>
```

3. Load the installation DVD-ROM and, from any directory, execute the installation shell script (install.sh) stored on the DVD-ROM, with the "-s *installFilePath*" option added. After execution, installation of this product will start.

```
# mount /dev/deviceFileName DVD-ROMmountDir <ENTER>
# DVD-ROMmountDir/install.sh -s installFilePath <ENTER>
```

Note

The *installFilePath* specification is mandatory.

If there is no file in the specified *installFilePath* or if the read fails, an error message will be displayed and installation will stop.

4. When installation has completed normally, the message below will be output.

```
The installation processing completed successfully.
```

5. After installation has completed normally, reboot the OS.

```
# shutdown -r now <ENTER>
```

4.3.3 Post-installation Procedure

This section explains the tasks after installation.

4.3.3.1 Setting Environment Variables

Add the path below to the `PATH` environment variable of the CEP Server users.

```
/opt/FJSVcep/bin
```

Point

Creating a file with the content below in the `/etc/profile.d` directory of the CEP Server will allow the `PATH` environment variable to be set uniformly for the CEP Server users.

File: `/etc/profile.d/FJSVcep.sh`

```
# Interstage Big Data Complex Event Processing Server V1.0.0
export PATH=/opt/FJSVcep/bin:${PATH}
```

File: `/etc/profile.d/FJSVcep.csh`

```
# Interstage Big Data Complex Event Processing Server V1.0.0
setenv PATH /opt/FJSVcep/bin:${PATH}
```

4.3.3.2 Applying Updates

Refer to "[6.3.3 Applying Updates](#)" to apply product and software (included with this product) updates.

After applying the updates, perform setup for this product. Refer to "[4.4 Setup](#)" for details.

4.3.4 If an Error Occurs during Installation

This section explains how to respond if a failure occurs during installation.

If an error occurs before package installation

After responding in accordance with the error message, re-execute the `install.sh` shell.

If an error occurs during package installation

Execute `uninstall.sh` and uninstall the packages that have been installed. Then, after responding in accordance with the error message output at `install.sh` execution, re-execute the `install.sh` shell.

Refer to Section 2.1, "Errors during Installation" in *Troubleshooting* for details.

Information

Installer log file

Below is the log file output by the Installer. It records detailed operation statuses in addition to the messages displayed in the windows and can be referred to when an issue arises.

```
/var/tmp/bdcep_install.log
```

When installation is successful, the log file will be stored as follows:

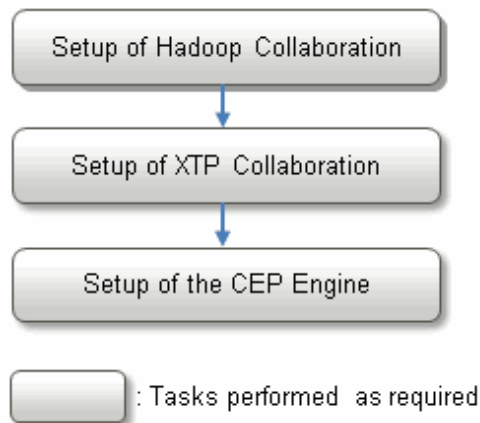
```
/var/opt/FJSVcep/bdcep_install.log
```

4.4 Setup

This section explains setup for this product.

4.4.1 Setup Overview

An overview of the setup of this product is shown below.



4.4.2 Setup of Hadoop Collaboration

With this product, received events and events output from the CEP engine can be logged. The following two destinations can be selected for logging:

- Hadoop system
- CEP Server (*1)

*1: Events are logged in the engine log for the purpose of verifying events received and verifying the results of complex event processing.

To log events in a Hadoop system, setup of Hadoop collaboration is also required.

The Hadoop system that can be linked to is shown below.

- Interstage Big Data Parallel Processing Server V1.0.0 (hereafter, referred to as "BDPP 1.0")

Setup procedure

1. On the CEP Server, install and setup the BDPP 1.0 Development Server. The CEP engine operates as an application for accessing the Hadoop system.
 Refer to Section 3.1.3, "Installing a Development Server" in the *User's Guide* of the BDPP 1.0 manuals for details.
 Refer to E.3, "Adding a Slave Server/Scaling Out" in the *User's Guide* of the BDPP 1.0 manuals also, when you setup a Development Server in addition to the existing Hadoop system.
2. On the Hadoop system, register the engine execution user specified at installation of this product as a Hadoop user.
 Refer to D.6, "User-Specific Hadoop Usage Settings" in the *User's Guide* of the BDPP 1.0 manuals for details.

4.4.3 Setup of XTP Collaboration

Complex Event Processing of this product can collaborate with Interstage eXtreme Transaction Processing Server (hereafter, referred to as "XTP") to rapidly reference the cache on XTP. To perform XTP collaboration, setup of XTP collaboration is also required.

The XTP that can be linked to is shown below.

- Interstage eXtreme Transaction Processing Server V1.0.0 (hereafter, referred to as "XTP 1.0")

Setup procedure

Install and set up the XTP Client Node of XTP 1.0 in the CEP Server. The CEP Server operates as a XTP Client Node, and the CEP engine operates as a cache access application. Refer to the XTP 1.0 manual for details.

4.4.4 Setup of the CEP Engine

This section explains the setup of the CEP engine.

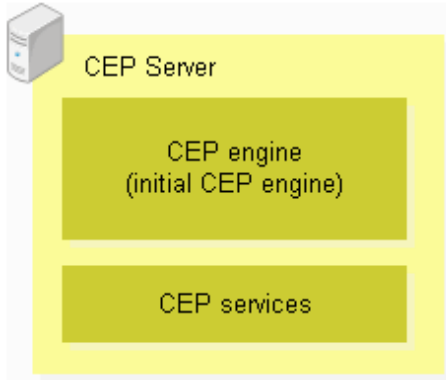
- [Status Immediately after Installation](#)

- [Changing CEP Engine Settings](#)
- [Creating a New CEP Engine](#)

4.4.4.1 Status Immediately after Installation

After installation, the CEP service that manages the CEP engine will start, and one immediately usable CEP engine (hereafter, referred to as the "initial CEP engine") will be created. The CEP engine will have stopped status.

Figure 4.1 Status immediately after installation



This initial CEP engine has the settings shown below. This initial CEP engine can be used to deploy and check the operation of a sample application.

Item	Value to be set
CEP engine name	The name specified at installation. This is "CepEngine" if the specification was omitted at installation.
Logging type	Not set (Logging not used).
Directory name	Same as above
Number of open logging files	Same as above.
Logging cycle time	Same as above.
Socket adapter port	9600
JVM High-speed Filter options (*1)	Maximum value of memory allocation pool: 2048 MB Initial value of memory allocation pool: 512 MB Maximum value of permanent generation area: 192 MB
JVM Complex Event Processing options (*1)	Same as above.

*1: Settings relating to memory used by the CEP engine. Refer to "[6.3.4.1 Tuning JVM Options](#)" for information on each parameter.

Below is the engine configuration file used in initial CEP engine creation.

```
/etc/opt/FJSVcep/cep/sample_eng.xml
```

Refer to "[9.1.1 Engine Configuration File](#)" for information on how to view the content of the engine configuration file.

4.4.4.2 Changing CEP Engine Settings

When using logging in the initially created CEP engine, the CEP engine settings must be changed.

The two methods below are used to change the CEP engine settings, and the items that can be changed will vary depending on the method used.

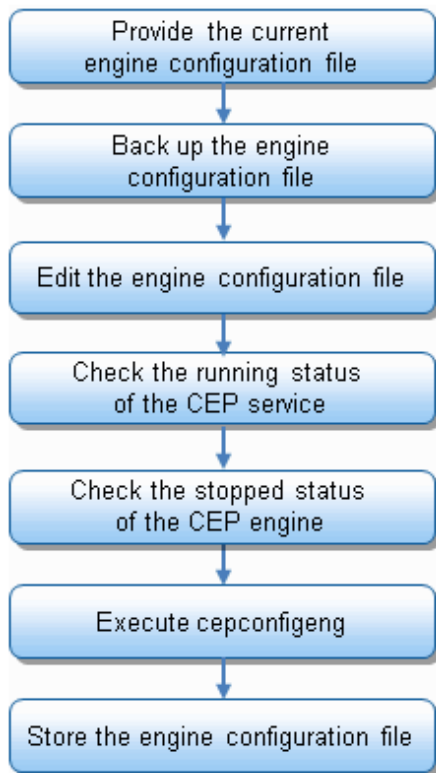
Method for changing settings	Items that can be changed
Using cepconfigeng	- Logging settings - Port number used in socket communication
Using cepsetjvmopt	- Memory used by CEP engine

Changing settings using cepconfigeng is explained below.

Refer to "[6.3.4.1 Tuning JVM Options](#)" for information on changing settings using cepsetjvmopt.

Changing settings using cepconfigeng

Below is a flowchart for changing settings using cepconfigeng.



The procedure for changing CEP engine settings is explained below. For these tasks, log in as a superuser to execute the commands.

1. Prepare the current engine configuration file.

A CEP engine that is not to be changed must also be defined in the engine configuration file to be specified at cepconfigeng command execution. Therefore, define the changes based on the current stored engine configuration file.

The engine configuration file used in the creation of the initial CEP engine is stored in the following location:

```
/etc/opt/FJSVcep/cep/sample_eng.xml
```

Example

Example of preparing the engine configuration file used in the creation of the initial CEP engine

```
# cp /etc/opt/FJSVcep/cep/sample_eng.xml /etc/opt/FJSVcep/Engine.xml<ENTER>
```



Note

If the engine configuration file to be specified at `cepconfigeng` command execution does not include a created CEP engine definition, `cepconfigeng` will delete the CEP engine which is not specified in the engine configuration file.

2. Back up the engine configuration file.

Before editing the current engine configuration file, create a backup of the engine configuration file. Always create a backup to avoid losing definition information erroneously.



Example

Example of backing up the engine configuration file

```
# cp /etc/opt/FJSVcep/Engine.xml /etc/opt/FJSVcep/Engine.bak.xml<ENTER>
```

3. Edit the engine configuration file.

Edit the engine configuration file using a command such as `vi`, and then define the settings for the target CEP engine.

Refer to "9.1.1 Engine Configuration File" for information on the format of the engine configuration file.



Example

Definition example of enabling logging for the initial CEP engine

The definition example below shows "enabling logging of the CEP Server (output destination is the engine log)" for the initial CEP engine immediately after installation.

The "CepEngine" that is displayed in the definition depends upon the CEP engine name specified at installation.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<subSystemConfig xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1">
  <engineConfig id="CepEngine">
    <logging>
      <type>file</type>
    </logging>
    <socketAdapterPort>9600</socketAdapterPort>
  </engineConfig>
</subSystemConfig>
```

4. Check the running status of the CEP service.

The CEP service must be running to execute `cepconfigeng`. Use `cepdisperv` to check the status of the CEP service. Refer to "8.5 `cepdisperv`" for details.

If the CEP service is not running, use `cepstartserv` to start it. Refer to "8.10 `cepstartserv`" for details.



Example

Example of using `cepdisperv` to check the running status of the CEP service

Execute `cepdisperv` to check that the content below is output.

The operating process number will be displayed where "nnnn" is displayed.

Generated CEP engine names are displayed in place of *CepEngine*.

```
# cepdisperv<ENTER>
(...)
Interstage Java EE DAS          started
(...)
Interstage Java EE Node Agent   started
```

```
(...)
CEPAgentIJServerCluster running
CepEngineflt not running
CepEnginecep not running
(...)
Status          : Running
(...)
jsvc (pid nnnn nnnn) is running...
(...)
pg_ctl: server is running (PID: nnnn)
(...)
Command cepdisperserv executed successfully.
```

Information

The CEP service entity

A CEP service is a service made up of multiple processes.

5. Check the stopped status of the engine.

The target CEP engine must be stopped for `cepconfigeng` to change the settings of the CEP engine. Use `cepdispeng` to check the status of the engine. Refer to "8.4 [cepdispeng](#)" for details.

If the CEP engine is running, use `cepstopeng` to stop the CEP engine. Refer to "8.11 [cepstopeng](#)" for details.

Example

Example of using `cepdispeng` to check the stopped status of the initial CEP engine

Execute `cepdispeng` to check the content of the initial CEP engine output, as shown below.

The "CepEngine" that is displayed in the command execution example and output example depends upon the CEP engine name specified at installation.

```
# cepdispeng -e CepEngine<ENTER>
engineId      :CepEngine
port          :9600
status_filter :STOP
status_cep    :STOP
Command cepdispeng executed successfully.
```

6. Execute `cepconfigeng`.

Specify the edited engine configuration file and execute `cepconfigeng`. When the command is executed, confirmation of the change is requested, type "y" to continue execution. Execution of the command can be canceled by typing "n" or "q". Refer to "8.2 [cepconfigeng](#)" for details.

Example

Example of executing `cepconfigeng`

Below is an example of specifying `/etc/opt/FJSVcep/Engine.xml` as the edited engine configuration file.

```
# cepconfigeng -f /etc/opt/FJSVcep/Engine.xml<ENTER>
Are you sure you want to change the CEP Engine configuration? [y,n,q]:y<ENTER>
Command cepconfigeng executed successfully.
```


7. Store the engine configuration file.

The engine configuration file specified in `cepconfigeng` will be required in future for further additions or deletions of CEP engines or for changing settings, so store it in a safe place. Consider creating a backup on external media, as required.



The backup creation method depends upon the management policy of the system. Always check with the system administrator.

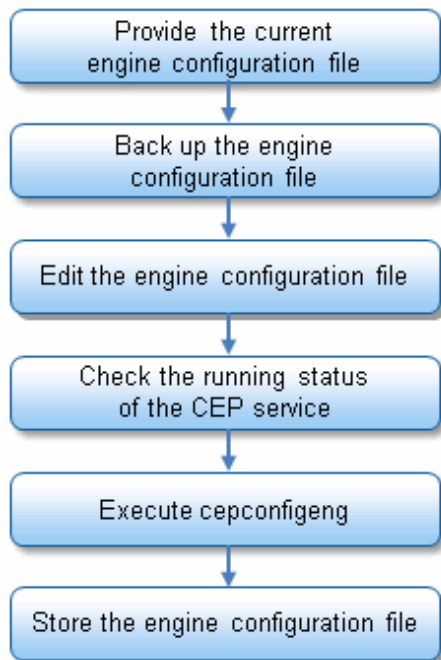
4.4.4.3 Creating a New CEP Engine

In general, only one CEP engine operates, but there may be times when multiple CEP engines need to be provided, such as when a development environment is being divided for multiple development groups.

This product allows a maximum of five CEP engines to be created for one CEP Server. This section explains how to create a new CEP engine using `cepconfigeng`.

Creating a new CEP engine using `cepconfigeng`

Below is a flowchart for creating a new CEP engine using `cepconfigeng`.



The procedure for creating a new CEP engine is explained below. For these tasks, log in as a superuser to execute the commands.

1. Provide the current engine configuration file.

An existing CEP engine must also be defined in the engine configuration file to be specified at `cepconfigeng` command execution. Therefore, define the settings of the CEP engine to be added based on the current stored engine configuration file.



If an operation has been performed previously according to the execution example in "4.4.4.2 Changing CEP Engine Settings", the current engine configuration file will be the following file:

```
/etc/opt/FJSVcep/Engine.xml
```

Note

If the engine configuration file to be specified at `cepconfigeng` command execution does not include an existing CEP engine definition, `cepconfigeng` will delete the CEP engine which is not specified in the engine configuration file.

2. Back up the engine configuration file.

Before editing the current engine configuration file, create a backup of the engine configuration file. Always create a backup to avoid losing definition information erroneously.

Example

Example of backing up the engine configuration file

```
# cp /etc/opt/FJSVcep/Engine.xml /etc/opt/FJSVcep/Engine.bak.xml<ENTER>
```

3. Edit the engine configuration file.

Edit the engine configuration file using a command such as `vi`, and then define the settings for the new CEP engine to be added.

Use a CEP engine name that does not duplicate the name of another CEP engine. Refer to "[9.1.1 Engine Configuration File](#)" for information on the format of the engine configuration file.

Example

Example of creating a new CEP engine called "NewCepEngine"

Create a new CEP engine with the settings below, in addition to the initial CEP engine changed according to the example of executing changes to the settings of a CEP engine in "[4.4.4.2 Changing CEP Engine Settings](#)"

Item	Value to be set	Description
CEP engine name	NewCepEngine	Use a CEP engine name that is unique.
Logging type	Logging is not used, so do not set a value.	
Directory name	Same as above	
Number of open logging files	Same as above.	
Logging cycle time	Same as above.	
Socket adapter port	Socket communication is not performed, so do not set a value.	If not set, only a SOAP adapter or HTTP adapter can be used.

Below is an example of defining an engine configuration file.

The part from "`<engineConfig id="NewCepEngine">`" to "`</engineConfig>`" will become the definition of the new CEP engine to be added.

The existing CEP engine definition will remain without change. The "CepEngine" that is displayed in the definition depends upon the CEP engine name specified at installation.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<subSystemConfig xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1">
  <engineConfig id="CepEngine">
    <logging>
      <type>file</type>
    </logging>
    <socketAdapterPort>9600</socketAdapterPort>
  </engineConfig>
```

```
<engineConfig id="NewCepEngine">
</engineConfig>
</subSystemConfig>
```

4. Check the running status of the CEP service.

The CEP service must be running to execute `cepconfigeng`. Use `cepdisserv` to check the status of the CEP service. Refer to "8.5 [cepdisserv](#)" for details.

5. Execute `cepconfigeng`.

Specify the edited engine configuration file and execute `cepconfigeng`. When the command is executed, confirmation of the change is requested, so type "y" to continue execution. Execution of the command can be canceled by typing "n" or "q". Refer to "8.2 [cepconfigeng](#)" for details.

Example

Example of executing `cepconfigeng`

Below is an example of specifying `/etc/opt/FJSVcep/Engine.xml` as the edited engine configuration file.

```
# cepconfigeng -f /etc/opt/FJSVcep/Engine.xml<ENTER>
Are you sure you want to change the CEP Engine configuration? [y,n,q]:y<ENTER>
Command cepconfigeng executed successfully.
```

6. Store the engine configuration file.

The engine configuration file specified in `cepconfigeng` will be required in future for further additions or deletions of CEP engines or for changing settings, so store it in a safe place. Consider creating a backup on external media, as required.

Note

The backup creation method depends upon the management policy of the system. Always check with the system administrator.

4.5 Canceling Setup

This section explains how to cancel the setup of the features set up in "4.4 Setup".

- [Deleting a CEP Engine](#)
- [Canceling XTP Collaboration](#)
- [Canceling Hadoop Collaboration](#)

4.5.1 Deleting a CEP Engine

This section explains how to delete a CEP engine.

Deleting a CEP engine that is no longer required allows the system resources such as memory and disk that were used by that CEP engine to be used for other purposes. This section explains how to delete a CEP engine using `cepconfigeng`.

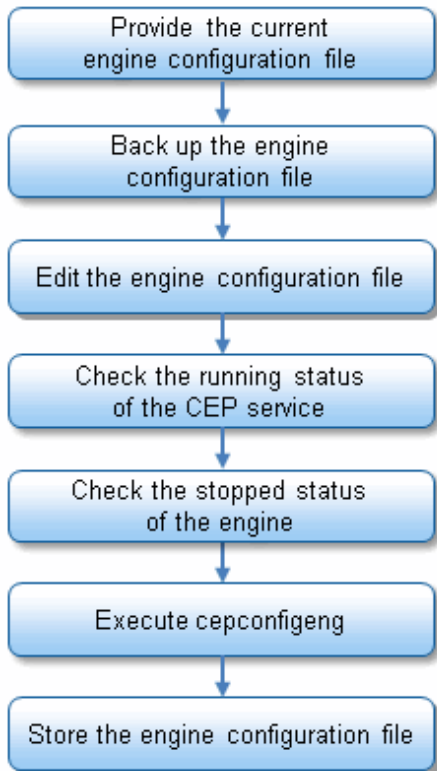
Note

When a CEP engine is deleted, definition information deployed to the CEP engine (such as rule definitions) will also be deleted. The engine log and resource log of the CEP engine will be deleted as well. Take a backup as required. Refer to "6.3.2 [Backup and Restore](#)" for details.

Deleting a CEP engine using cepconfigeng

Below is a flowchart for deleting a CEP engine using cepconfigeng.

Figure 4.2 Flowchart for deleting a CEP engine



The procedure for deleting a CEP engine is explained below. For these tasks, log in as a superuser to execute the commands.

1. Provide the current engine configuration file.

An existing CEP engine must also be defined in the engine configuration file to be specified at cepconfigeng command execution. Therefore, define the settings of the CEP engine to be deleted based on the current stored engine configuration file.

Point

If an operation has been performed previously according to the execution example in "4.4.4.3 Creating a New CEP Engine", the current engine configuration file will be the following file:

```
/etc/opt/FJSVcep/Engine.xml
```

Note

If the engine configuration file to be specified in cepconfigeng does not include an existing CEP engine definition, cepconfigeng will delete the CEP engine which is not specified in the engine configuration file.

2. Back up the engine configuration file.

Before editing the current engine configuration file, create a backup of the engine configuration file. Always create a backup to avoid losing definition information erroneously.

Example

Example of backing up the engine configuration file

```
# cp /etc/opt/FJSVcep/Engine.xml /etc/opt/FJSVcep/Engine.bak.xml<ENTER>
```

3. Edit the engine configuration file.

Edit the engine configuration file using a command such as vi, and then comment out or delete the definition for the CEP engine to be deleted.

Refer to ["9.1.1 Engine Configuration File"](#) for information on the format of the engine configuration file.

Example

Definition example of deleting a CEP engine called "NewCepEngine"

In this example, the CEP engine called "NewCepEngine" that was created according to the execution example in ["4.4.4.3 Creating a New CEP Engine"](#) will be deleted.

Below is a definition example of the engine configuration file.

Rather than simply deleting the definition of the relevant part, the part from "<engineConfig id="NewCepEngine">" to "</engineConfig>" is commented out. Commenting out allows you to thoroughly check the deletion range while you work. The "CepEngine" that is displayed in the definition depends upon the CEP engine name specified at installation.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<subSystemConfig xmlns="urn:xmlns-fujitsu-com:csnf:bdcep:v1">
  <engineConfig id="CepEngine">
    <logging>
      <type>file</type>
    </logging>
    <socketAdapterPort>9600</socketAdapterPort>
  </engineConfig>
<!--
  <engineConfig id="NewCepEngine">
    </engineConfig>
-->
</subSystemConfig>
```

4. Check the running status of the CEP service.

The CEP service must be running to execute cepconfigeng. Use cepdisperv to check the status of the CEP service. Refer to ["8.5 cepdisperv"](#) for details.

5. Check the stopped status of the engine.

The target CEP engine must be stopped for cepconfigeng to be used to delete the CEP engine. Use cepdispeng to check the status of the engine. Refer to ["8.4 cepdispeng"](#) for details.

6. Execute cepconfigeng.

Specify the edited engine configuration file and execute cepconfigeng. When the command is executed, confirmation of the change is requested, so type "y" to continue execution. Execution of the command can be canceled by typing "n" or "q". Refer to ["8.2 cepconfigeng"](#) for details.

Example

Example of executing cepconfigeng

Below is an example of specifying "/etc/opt/FJSVcep/Engine.xml" as the edited engine configuration file.

```
# cepconfigeng -f /etc/opt/FJSVcep/Engine.xml<ENTER>
Are you sure you want to change the CEP Engine configuration? [y,n,q]:y<ENTER>
Command cepconfigeng executed successfully.
```

7. Store the engine configuration file.

The engine configuration file specified in `cepconfigeng` will be required in future for further additions or deletions of CEP engines or for changing settings, so store it in a safe place. Consider creating a backup on external media, as required.

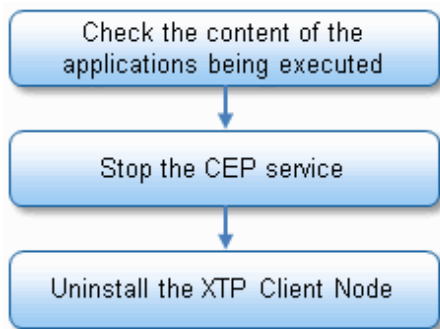
Note

The backup creation method depends upon the management policy of the system. Always check with the system administrator.

4.5.2 Canceling XTP Collaboration

Below is a flowchart for canceling XTP collaboration.

Figure 4.3 Flowchart for canceling XTP collaboration



1. Check the content of the applications being executed.

Make enquiries to the application maintenance staff to check that the applications being executed do not require XTP collaboration.

Information

If an application is to use XTP collaboration to reference the XTP cache, use the `@VDW` annotation in the complex event processing rule definition. Refer to "[5.4.6.2 Using an XTP Cache](#)" for details.

2. Stop the CEP service.

Stop the CEP service that is running. The stop method is shown below. Log in as a superuser to execute the command.

```
# cepstopserv<ENTER>
```

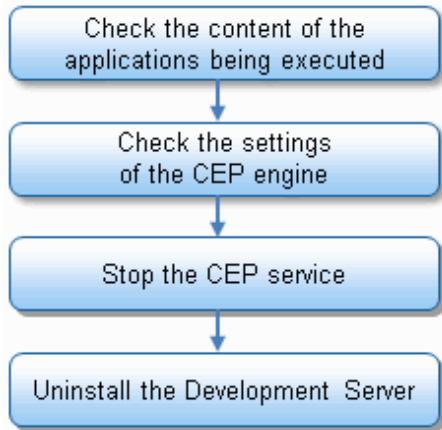
3. Uninstall the XTP Client Node.

Uninstall the XTP Client Node. Refer to the Interstage eXtreme Transaction Processing Server manual for information on how to uninstall it.

4.5.3 Canceling Hadoop Collaboration

Below is a flowchart for canceling Hadoop collaboration.

Figure 4.4 Flowchart for canceling Hadoop collaboration



1. Check the content of the applications being executed.

Make enquiries to the application maintenance staff to check that the applications being executed do not perform Hadoop collaboration. If there is an application that is to perform Hadoop collaboration, consider modifying the application.

 **Information**

An application uses the two methods below to use Hadoop collaboration. Refer to "[Chapter 5 Development](#)" for details.

- Specifying "true" in the useLogging item in the event type definition
- Using the "@LoggingListener" annotation in the complex event processing rule definition

2. Check the settings of the CEP engine.

Check the content of the engine configuration file used in the settings of the existing CEP engine to ensure the existing CEP engine is not specifying a Hadoop collaboration setting. If there is a CEP engine performing Hadoop collaboration, consider changing the settings of the CEP engine.

 **Information**

If an operation has been performed previously according to "[4.4.4 Setup of the CEP Engine](#)", the current engine configuration file will be the following file:

```
/etc/opt/FJSVcep/Engine.xml
```

To perform Hadoop collaboration, "bdpp" will be specified in the "type" element under the "logging" element. Refer to "[9.1.1 Engine Configuration File](#)" for information on the format of the engine configuration file.

3. Stop the CEP service.

Stop the CEP service that is running. The stop method is shown below. Log in as a superuser to execute the command.

```
# cepstopserv<ENTER>
```

4. Uninstall the Development Server.

Uninstall the Development Server. Refer to the Interstage Big Data Parallel Processing Server manual for information on how to uninstall it.

4.6 Uninstallation

This section explains how to uninstall this product.

4.6.1 Pre-uninstallation Procedure

This section explains the tasks required before uninstalling this product.

4.6.1.1 Stopping Event Sending

Stop event sending to the CEP Server.

4.6.1.2 Backing up User Assets

Back up user assets. Refer to "[6.3.2 Backup and Restore](#)" for information on how to back up user assets such as definition information.

4.6.1.3 Stopping the CEP Service

Stop the CEP service that is running. The stop method is shown below. Log in as a superuser to execute the command.

```
# cepstopserv<ENTER>
```

4.6.1.4 Deleting Updates

If the following updates of this product and of the software bundled with this product have been applied in UpdateSite format, delete the updates:

- Interstage Big Data Complex Event Processing Server
- Interstage Application Server Enterprise Edition (64 bit)



See

.....
Refer to the Help in "UpdateAdvisor (middleware)" and the update information files of each update for details.
.....

4.6.2 Uninstallation Procedure

This section explains the uninstallation procedure.

From the following two types, select the uninstallation method that best suits the method of use. If uninstallation using multi-user mode, check that the operations of other users will not affect the uninstallation:

Attended uninstallation

Use 'attended uninstallation' to execute the uninstallation with querying to check that uninstallation is to be executed and then executes uninstallation.

Unattended uninstallation

Use 'unattended uninstallation' to execute the uninstallation with no querying from the Uninstaller.



Note

.....
With uninstallation, the packages installed by this product will be uninstalled. Specific installed packages and features cannot be selected for uninstallation. Some packages will be uninstalled manually.

In addition, the resources under the installation destination directory will be deleted at uninstallation. Always save necessary resources before uninstallation. Refer to "[6.3.2 Backup and Restore](#)" for information on backup.
.....

4.6.2.1 Attended Uninstallation

1. Log in as a superuser.

```
$ su -<ENTER>
```


2. Load the installation DVD-ROM and execute the uninstallation shell script (uninstall.sh shell) stored on the DVD-ROM from any directory.

```
# mount /dev/deviceFileName DVD-ROMmountDir <ENTER>
# DVD-ROMmountDir/uninstall.sh <ENTER>
```

3. The product name will be displayed as shown below.

```
+-----+
| Interstage Big Data Complex Event Processing Server V1.0.0 |
|                                                              |
|                               Copyright 2012 FUJITSU LIMITED |
+-----+
```

4. Continuing from the display above, the uninstallation content will be displayed as shown below. In *uninstallPackages*, all of the packages to be uninstalled will be displayed delimited by spaces.

```
Uninstallation packages:
  uninstallPackages
```

5. Check the content and, to start uninstalling the packages, type "y" and press the Enter key. To stop uninstallation, type "q" and press the Enter key. After typing "y", uninstallation of this product will start. Uninstallation can be stopped by typing "q".

```
All files/directories under the installation destination directory will be deleted.
Please back up the required resources before uninstalling.

Do you want to proceed with the uninstallation? [y,q]:
```

6. When uninstallation has completed normally, the message below will be output. Perform the [post-uninstallation procedure](#) to continue.

```
The uninstallation processing completed successfully.
```

4.6.2.2 Unattended Uninstallation

1. Log in as a superuser.

```
$ su -<ENTER>
```

2. Load the installation DVD-ROM and, from any directory, execute the uninstallation shell script (uninstall.sh shell) stored on the DVD-ROM, with the -s option added.

```
# mount /dev/deviceFileName DVD-ROMmountDir <ENTER>
# DVD-ROMmountDir/uninstall.sh -s<ENTER>
```

3. When uninstallation has completed normally, the message below will be output.

```
The uninstallation processing completed successfully.
```

4.6.3 Post-uninstallation Procedure

This section explains the tasks after uninstallation.

4.6.3.1 Uninstalling FJSVod

The uninstaller does not uninstall the FJSVod package when the following product is installed on CEP Server.

- Systemwalker Centric Manager (Management Server)

Below are manual uninstallation steps for the FJSVod package.

1. Check product usage

Check if the following product is installed. If it is installed, do not uninstall the FJSVod package.

- Systemwalker Centric Manager (Management Server)

2. Uninstall the FJSVod package

Execute rpm as a superuser to uninstall the FJSVod package.

```
# /bin/rpm -e --nodeps FJSVod <ENTER>
```

4.6.3.2 Uninstalling FJSVsmee64 and FJSVsclr64

The uninstaller does not uninstall the FJSVsmee64 and FJSVsclr64 packages, because these packages may be included in Fujitsu non-Interstage products such as Systemwalker Centric Manager.

If you want to uninstall these packages manually, execute following commands as a superuser.

```
# /bin/rpm -e --nodeps FJSVsmee64 <ENTER>
# /bin/rpm -e --nodeps FJSVsclr64 <ENTER>
```

4.6.3.3 Removing Environment Variables

Delete the path below from the PATH environment variable of the CEP Server users.

```
/opt/FJSVcep/bin
```

Point

If the following files were created in the /etc/profile.d directory in "4.3.3.1 Setting Environment Variables", delete the created files:

- /etc/profile.d/FJSVcep.sh
- /etc/profile.d/FJSVcep.csh

4.6.3.4 Engine Execution User Specified at Installation

The engine execution user and the group to which the engine execution user belongs that are specified at installation will not be deleted at uninstallation. If the user is not required, delete the user.

Example

Below is an example of manually deleting the execution user called "isbdcep".

```
$ su -<ENTER>
# userdel isbdcep<ENTER>
```

Note

The user and group deletion method depends upon the management policy of the system. Always check with the system administrator.

4.6.4 If an Error Occurs during Uninstallation

If execution of the uninstaller fails, after responding in accordance with the message, re-execute the uninstall.sh shell.

Refer to Section 2.2, "Errors during Uninstallation" in *Troubleshooting* for details.



Information

Log file of the uninstaller

Below is the log file output by the uninstaller. It records detailed operation statuses in addition to the messages displayed in the windows and is referred to when trouble occurs.

```
/var/tmp/bdcep_uninstall.log
```

When uninstallation is successful, the log file will be deleted.

Chapter 5 Development

This chapter explains event processing using this product, from design to development and testing.

This chapter also explains the sample application.

5.1 Overview of Complex Event Processing of this Product

The event processing of this product is broadly divided into the following two types:

- Event processing that uses logging to record input events in an external Hadoop system.
- Event processing that applies a series of rules to input events and either logs the processing results or sends them to an external Web service. This type of event processing also performs processes such as referencing external data when applying the rules.

Deploying the following definition information allows the system to perform these series of operations. The sections that follow explain the details of this definition information:

- Event type definition (mandatory)
- Rule definition (optional)
- SOAP listener definition (optional)
- Master definition (optional)

In addition to this definition information, using this product to utilize the processing results of events will also require development assets, such as applications for analyzing logged events or processing results and Web services that operate after receiving the processing results of events.

5.2 List of Development Assets

The table below lists the development assets of this product. Development assets are broadly divided into definition information, data, and collaboration applications. Definition information is deployed in the CEP engine for execution. Each of the others is provided (deployed) in its corresponding server.

Table 5.1 List of development assets

Development asset type	Development asset	Explanation	Deployment destination
Definition information	Event type definition	Define the format of the events the input adapter is to receive. Also define the log storage area specification and whether or not complex event processing is to be used. The events the input adapter is to receive can be in XML or CSV format.	CEP engine
	Rule definition	Define high-speed filter processing rules and complex event processing rules. Describe these using filter rule language (IF-THEN format) and complex event processing rule language (SQL format), respectively. The processing results of complex event processing rules can be sent to an external application using SOAP communication, and can also be accumulated on disks using logging.	
	Event type definition (filtered events)	Define this if the items that make up the events will vary according to factors such as join and extraction processing of high-speed filter processing. Filtered events will be in CSV format.	

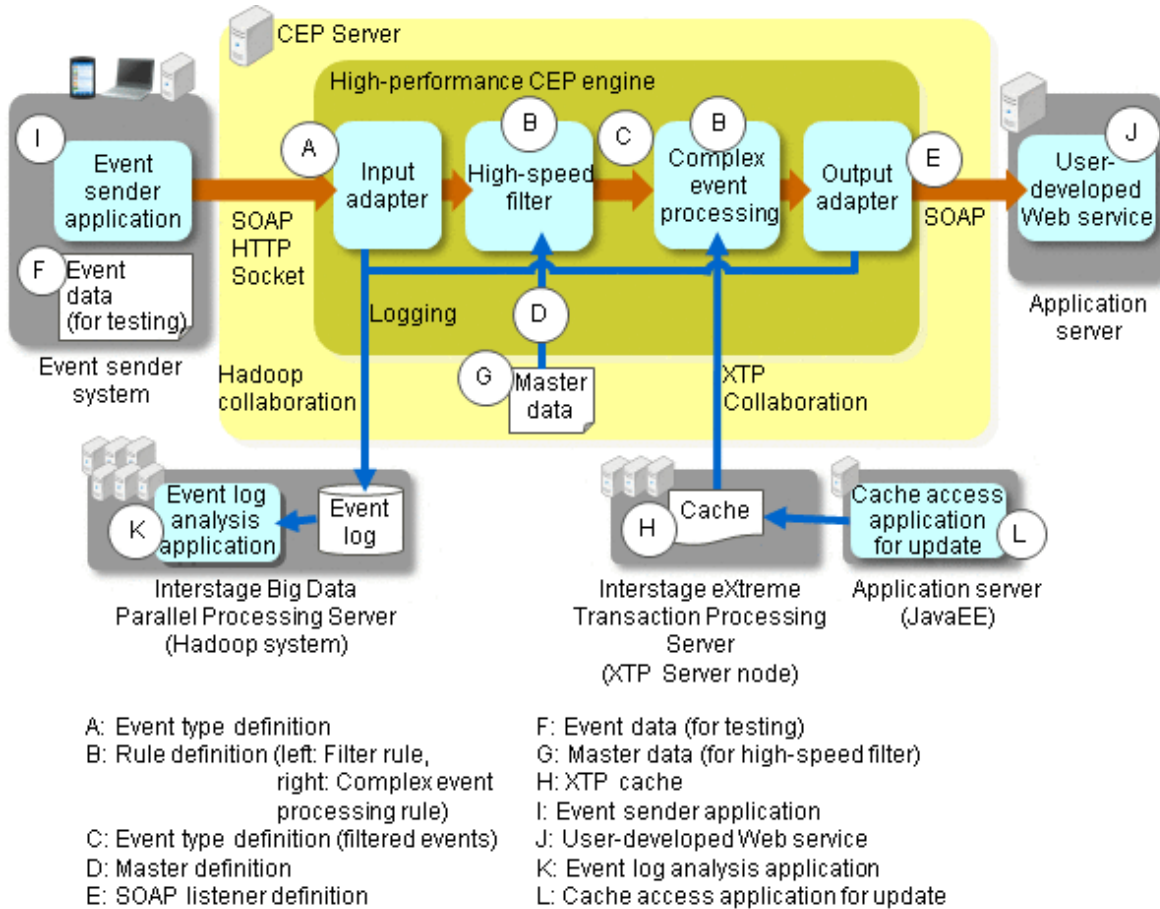
Development asset type	Development asset	Explanation	Deployment destination
	Master definition	Create this if master referencing is to be performed in high-speed filter processing. This is the definition of the master data to be referenced.	
	SOAP listener definition	Define the interface of the user-developed Web service to be used as the send destination when the processing results of complex event processing rules are to be sent to an external application using SOAP.	
Data	Event data (for testing)	This is event data to be sent to the CEP Server to check the operation of definition information.	Event sender system
	Master data (for the high-speed filter)	This is required separately if master referencing is to be performed in high-speed filter processing. Provide this in CSV file format on the CEP Server.	CEP Server
	XTP (*1) cache	This is required separately if XTP collaboration is to be performed. Provide this on the collaborating XTP Server node. A cache access application for update is required separately in order to store or update the data in an XTP cache.	XTP Server node
Collaboration application	Event sender application	This is an application that sends events to the CEP engine. To use SOAP for sending events, provide a SOAP client application. This is not required if events are to be sent directly to the CEP engine using an existing system as the event sender. If a new event sender application needs developing, it must be done according to the device to be used as the sender. The event sender sample program bundled with the product can be used in the operation testing of rule definitions.	Event sender system
	User-developed Web service	This is a Web service (SOAP application) that receives and controls event data sent by the output adapter. Deploy this in the application server which collaborate with CEP server.	Application server
	Event log analysis application	This is an application for analyzing event logs logged in a Hadoop system. Develop this using the Hadoop Java API. Deploy this in the Hadoop system and then execute it.	Hadoop system
	Cache access application for update	This is an application for initially storing data in the XTP cache as well as for updating the cache content. This is not required if an existing XTP cache is to be used.	Application server (Java EE) (*2)

*1: Interstage eXtreme Transaction Processing Server

*2: This is an application server where the XTP Client is installed and set up. It may sometimes be the same as an application server deployed by a user-developed Web service.

The following figure below shows the deployment destination of each development asset.

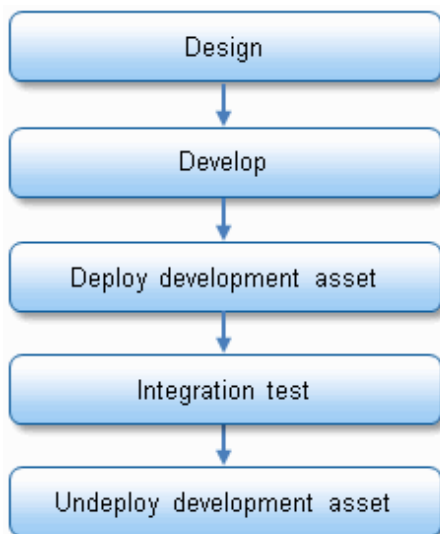
Figure 5.1 Deployment destinations of development assets



5.3 Task Overview

The development flow is shown below.

Figure 5.2 Development flow



Development environment

Design and develop definition information using a developer's own local personal computer, and check the operation by deploying the definition information in a CEP engine on the CEP Server.

For the event data for testing or master data (CSV), use data extracted from an existing database or data generated using a tool such as Excel.

When developing a collaboration application, use a development environment that is suitable for the collaborating system or product.

5.4 Design

This section explains how to design each development asset.

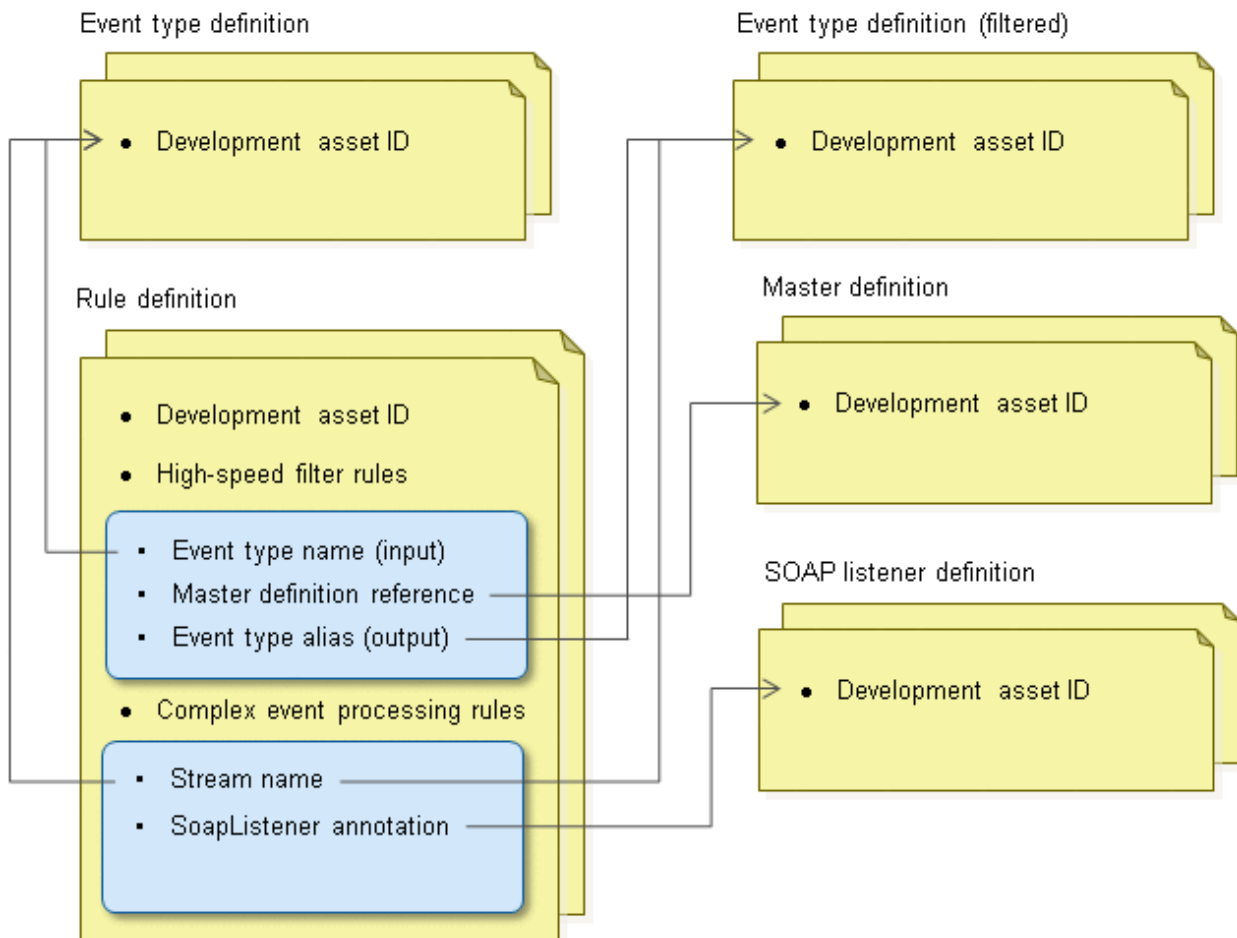
5.4.1 Association between the Development Asset ID and Definition Information

Each item of definition information has a development asset ID used to uniquely identify it as well as to associate it with other items of definition information.

When designing definition information, these development asset IDs must be managed.

Below is an example of the association between each development asset. The development asset IDs and rules described for each development asset must be defined so that they establish this association.

Figure 5.3 Association between development asset IDs and definition information



5.4.2 Designing an Event Type Definition

This section explains the points to consider when designing an event type definition.

5.4.2.1 Features of Input Events

Check details such as the data format and size of input events as well as their frequency of occurrence, as follows:

- If an existing system is to be used as the input event sender, check the event specifications of the existing system.
- Check the requirements relating to the amount of processing of the input events, such as the average input event size and the average number of events received per unit of time (for example, per hour or per second).
- In an event type definition, the two event data formats are XML and CSV format, as follows:
 - Select XML format or CSV format as appropriate.
 - Input events with other data formats such as binary format cannot be directly received, so consider converting the data format between the event sender and the CEP Server.

5.4.2.2 Recording and Analyzing Events

This section explains the points to consider when recording and accumulating events for analysis and other purposes.

To record events, use logging of this product.

Accumulating events

To accumulate the events in the Hadoop system where they are being logged, consider the following points:

- Type and format of the events to be logged
Consider what type of events to log and accumulate.

- Storage destination of the events to be logged

The Hadoop system to be connected is set according to the engine configuration file of the CEP engine.

In the Hadoop system to be connected, check for a path that can be used as a storage destination.

- Event accumulation capacity (free disk capacity required)

The event storage destination will require enough free disk capacity to accumulate massive volumes of events.

Check the average event size, number of recordings per unit of time, and accumulation period in order to consider the disk capacity required for recording and accumulating the events.

If necessary, consider an expansion plan for the system to be used as the event storage destination.

Analyzing accumulated events

An application for analysis must be designed and developed in order to analyze the accumulated events.

Refer to "[5.4.13 Designing an Event Log Analysis Application](#)" for information on designing an application for analysis.

5.4.3 Designing a Rule Definition

A rule definition consists of two types of rules: the filter rules to be used in high-speed filter processing and the complex event processing rules to be used in complex event processing.

This section explains the following items, including considerations when designing rules and the creation procedures to use:

- [High-speed Filter Processing](#)
 - [Considerations when creating filter rules](#)
 - [Processing pattern of filter rules](#)
- [Complex Event Processing](#)
 - [Considerations when creating complex event processing rules](#)
 - [Creation procedure for complex event processing rules](#)
- [SOAP Listener](#)

- [Logging Listener](#)

5.4.3.1 High-speed Filter Processing

Define the filter rules to be used by the high-speed filter.

Considerations when creating filter rules

The items to consider when creating filter rules are as follows:

- Unit of rule creation

Create a filter rule for each event type. Multiple filter rules cannot be defined simultaneously for a certain event type.

- Processing pattern of filter rules

Select a suitable processing pattern from the ones described below, and then create a rule similar to the selected pattern.

Refer to Chapter 2, "Filter Rule Language Reference" in the *Developer's Reference* for information on filter rules.

- Master data

Master data is referenced from within filter rules. Some individual events may contain only limited information such as an ID or code in order to cut down on the volume of event communication. Writing rules to process events is difficult in that situation, so it is possible to use master data that can be referenced using the ID or code as a key. Using master data in this way allows rules to be created more easily.

The master data must be created as files in CSV format by the user beforehand.

Refer to "[5.4.4 Designing a Master Definition](#)" for information on designing master data.

- Memory usage for filter rules

Using filter rules requires a large amount of memory. Refer to "[3.3.1.1 Amount of Memory when Using High-speed Filter Rules](#)" and "[3.3.1.2 Amount of Memory when Master Data is used by the High-speed Filter](#)" for information on the memory required.

Processing pattern of filter rules

By defining filter rules in the high-speed filter, the user can describe event extraction as well as extraction and join processing in combination with master data. The output of the high-speed filter is used directly as the input of complex event processing.

The processes performed by the high-speed filter are generally represented by the following four patterns and their combinations:

- [Extraction process](#)
- [Extraction process using master data matching](#)
- [Join processing with master data](#)
- [Weighting processing of text](#)

5.4.3.1.1 Extraction process

This processing pattern extracts from the input events those events that meet the conditions described in IF-THEN statements in the filter rules.

Consider using this processing pattern if only the events required are to be extracted from massive volumes of events.

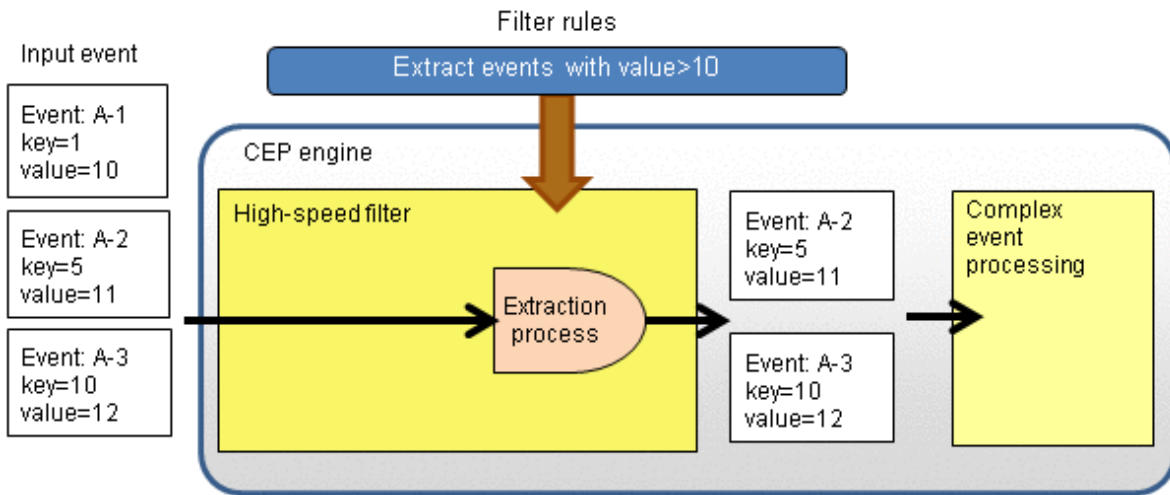


Example

Example of an extraction process

This is an example of extracting events using the content of the events (`value`).

"`value > 10`" is defined as the extraction condition.



Rule to be created

The rule to be created in the example above is as follows:

```
on inputEventTypeID {
  if ($value > 10) then output() as inputEventTypeID;
}
```

- *inputEventTypeID* is the development asset ID of the target event type definition.

Information

Filtering the items in the events to be output

Items in the events to be output can also be filtered. Below is an example of outputting only "key" in the example shown above.

To output using a different format from that of the input events, a corresponding event type definition (filtered) will also be required.

Refer to Section 2.7, "Output Expression Format" in the *Developer's Reference* for details.

```
on inputEventTypeID {
  if ($value > 10) then output($key) as outputEventTypeID;
}
```

- *inputEventTypeID* is the development asset ID of the target event type definition.
- *outputEventTypeID* is the development asset ID of the event type definition (filtered) that represents the results of filtering the item.

5.4.3.1.2 Extraction process using master data matching

This processing pattern matches the relevant entries of master data (CSV files) on the basis of values such as those of the ID or code contained in the input events, and then extracts the events based on the values of the relevant entries.

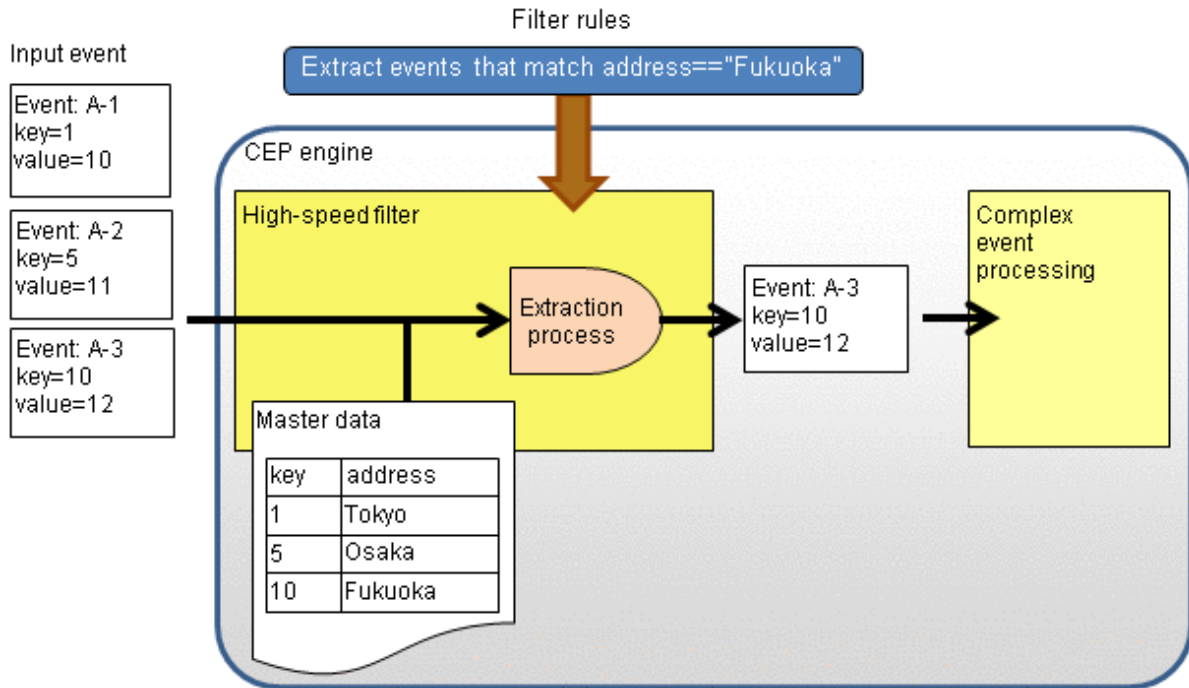
Consider using this processing pattern if only the events required are to be extracted from massive volumes of events but the events themselves do not contain the information required for extraction.

To perform the processing of this pattern, a master definition must also be designed.

Example

Example of the extraction process using master data matching

This is an example of referencing the relevant entries of master data based on the "key" contained in the events, and then extracting the events based on the values of "address" in the entries. "address==Fukuoka" is defined as the extraction condition.



Rule to be created

The rule to be created in the example above is as follows:

```

on inputEventTypeID {
    if (lookup("masterDefinitionID", $key == $key, string($address)) == "Fukuoka") then output() as
inputEventTypeID;
}

```

- *inputEventTypeID* is the development asset ID of the target event type definition.
- *masterDefinitionID* is the development asset ID of the master definition that corresponds to the master data to be referenced.
- To compare "\$address" as a string, use "string(\$address)" to fetch the values.
- The left side of "\$key == \$key" is "key" of the input events and the right side is "key" of the master data.

Note

The master data information will not be assigned to the input events if they are only matched using "lookup". If join processing with master data is required, a join expression must be described. Refer to "5.4.3.1.3 Join processing with master data" below for information on join processing.

5.4.3.1.3 Join processing with master data

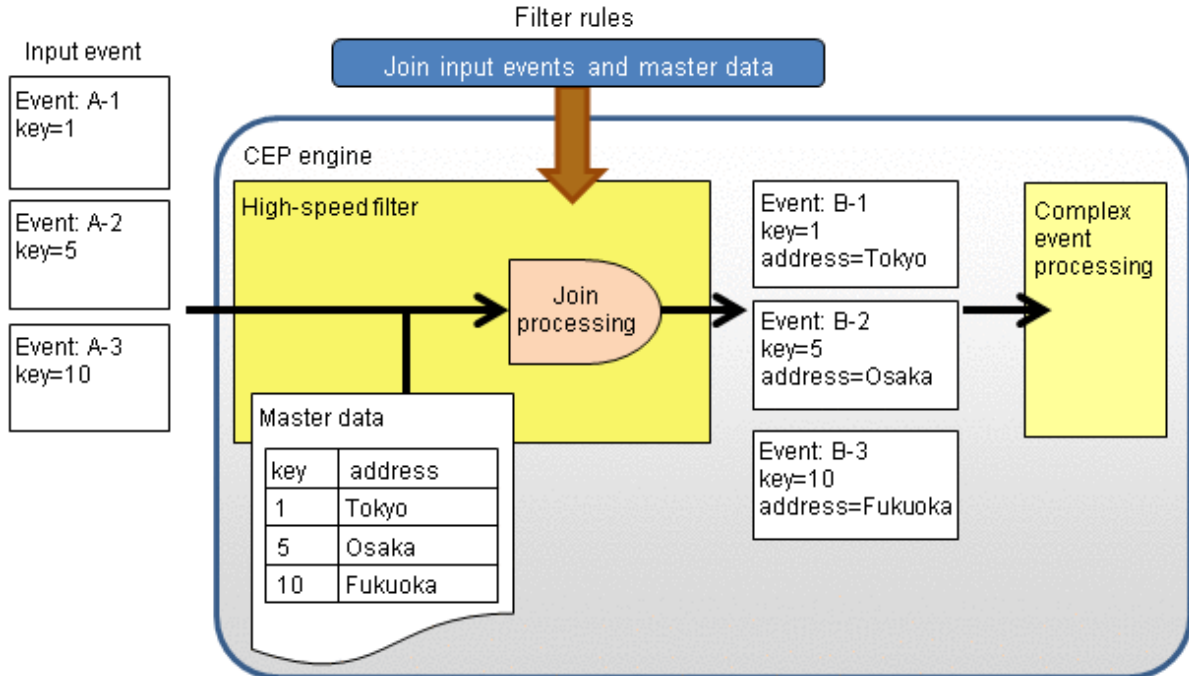
This processing pattern joins input events with master data. Consider using this processing pattern to assign the required data for the next complex event processing.

To pass the results of joining to complex event processing, an event type definition (filtered) corresponding to the join results will also be required.

Example

Example of join processing with master data

This is an example of joining the corresponding master data on the basis of "key" contained in the events, and then assigning "address" to the events.



Rule to be created

The rule to be created in the example above is as follows:

```
on inputEventTypeID {
    join("masterDefinitionID", $key == $key) output($key, "masterDefinitionID".$address) as
outputEventTypeID;
}
```

- *inputEventTypeID* is the development asset ID of the event type definition that is the rule target.
- *masterDefinitionID* is the development asset ID of the master definition that corresponds to the master data to be referenced.
- *outputEventTypeID* is the development asset ID of the event type definition (filtered) that represents the results of joining.

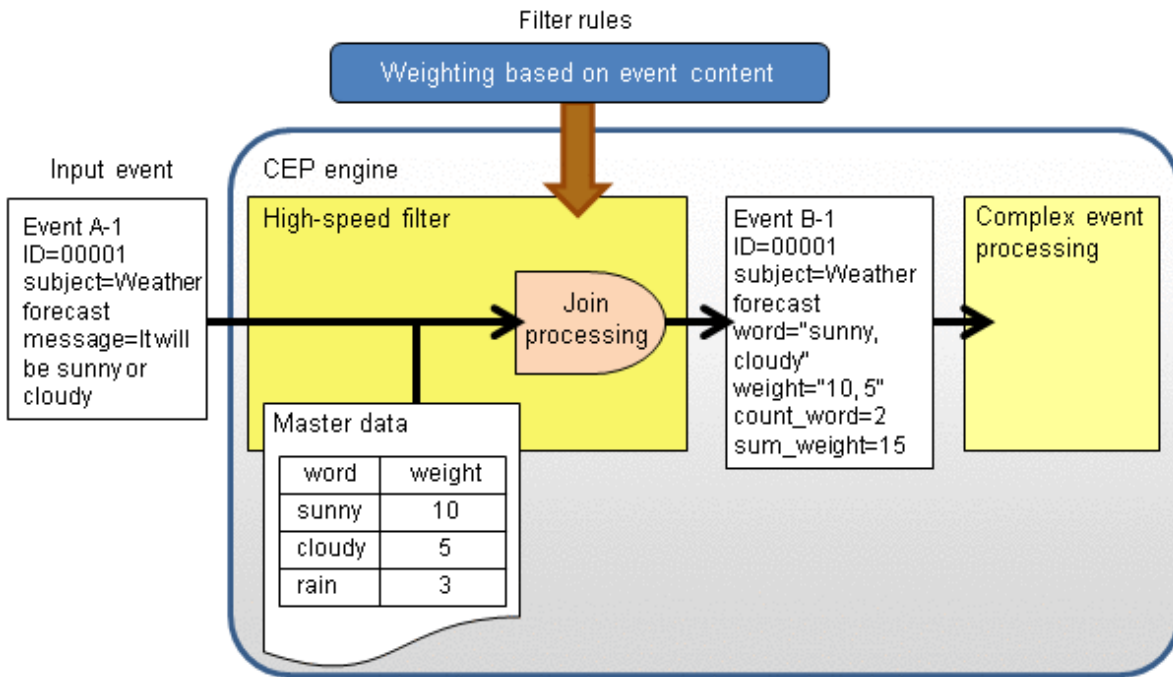
5.4.3.1.4 Weighting processing of text

This processing pattern can weight the text in input events by registering the weight of the keywords in the master data. This in turn allows applications including those that extract only those events with a total weighting that is above a threshold, and those that detect consecutively issued events that are above a threshold.

Example

Example of weighting processing of text

If the text contained in events contains search words that have been defined in the master data, assign the number of search words it contains as well as the total weighting value set for each search word, and then output them.



Rule to be created

The rule to be created in the example above is as follows:

```

on inputEventTypeID {
  join("masterDefinitionID", $message = $word)
  output($ID,
    $subject,
    "masterDefinitionID".$word,
    "masterDefinitionID".$weight,
    lookup_count("masterDefinitionID".$word),
    lookup_sum("masterDefinitionID".$weight)) as outputEventTypeID;
}

```

- *inputEventTypeID* is the development asset ID of the event type definition that is the rule target.
- *masterDefinitionID* is the development asset ID of the master definition that corresponds to the master data to be referenced.
- *outputEventTypeID* is the development asset ID of the event type definition (filtered) that represents the results of joining.

5.4.3.2 Complex Event Processing

Define the rules to be used by complex event processing.

Considerations when creating complex event processing rules

The items to consider when creating complex event processing rules are as follows:

- Unit of rule creation

Create a complex event processing rule for each use application or for each purpose of performing event pattern detection. Decide on a unit of creation in which an event pattern detection described in a certain rule will not affect other rules.

It is also possible to describe multiple processes in one rule definition, but this will create large rule definitions and may lead to reduced maintainability.

For example, if events relating to home electronic equipment are to be processed and the content to be detected varies significantly between domestic appliances and information devices, create the following two rule definitions:

- Rule definition to detect patterns in events relating to domestic appliances

- Rule definition to detect patterns in events relating to information devices
- Referencing external data

Consider whether referencing external data is necessary in event processing. An XTP cache can be used as external data.

Refer to "[5.4.6 Designing XTP Collaboration](#)" for information on referencing an XTP cache.

- Whether processing results are to be sent or logged

The processing results of complex event processing rules can be sent to an external Web service using SOAP but they can also be logged in an event log using logging. Apply the SOAP listener and logging listener, respectively, to the complex event processing rules in these cases.

Refer to "[5.4.3.3 SOAP Listener](#)" for information on how to use the SOAP listener.

Refer to "[5.4.3.4 Logging Listener](#)" for information on how to use the logging listener.

- Rule creation procedure

Design complex event processing rules in stages, without describing statements from the outset, and develop them so that the intended events will be reliably detected. The creation procedure for complex event processing rules is shown below.

Creation procedure for complex event processing rules

This section uses examples to explain the following creation procedure for complex event processing rules:

1. Consider what is to be achieved by using complex event processing.

Decide on what is to be achieved by using complex event processing. If this is unclear at this point, analyzing the collected events can sometimes clarify this.



Example

Example of a rule

"When someone is home, if rain is likely, recommend using the drying feature of the washing machine."

2. Consider the events, external data to be referenced, and output content.

Consider the events, the external data for referencing, and the output content required to create the rule. Also consider using a named window for retaining events in memory.



Example

Examples of events, external data, and output

The rule example above uses the following events, external data, and output content:

- Events
 - TV control event

For the household appliance event data sent from the home gateway of each household, the value of the device category property is to be "Television".
 - Weather forecast event

The forecast is to be represented by the time and weather.
- External data
 - Washing machine model information (whether it has a drying feature)

Whether the washing machine in that household has a drying feature is to be apparent from the home gateway ID.
- Output content
 - Home gateway ID of the household to be given the recommendation
 - Recommendation details

- Named window
 - Named window for weather forecast events
 - The weather forecast events for each time are to be retained for one day.

3. **Refine the processing content.**

Refine the content of the events and of the complex event processing for detecting them.

Here, "refine" is the task of using events and specific information such as complex event processing conditions and external data to make what was previously expressed in everyday language as "what is to be achieved" into a representation closer to the rule to be created.

 **Example**

Example of refining the content of complex event processing

The table below shows the results of refining the "When someone is home, if rain is likely, recommend using the drying feature of the washing machine" rule.

Element of the rule	Refined processing content
"When someone is home"	Determine that someone is home in the household where the TV was controlled (detect a TV control notifying event).
"If rain is likely"	Reference the weather forecast information stored in the named window and, from the time of the weather forecast and the time of the TV control, check whether or not there is a forecast of rain after this time.
"Drying feature of the washing machine"	Obtain product information on the washing machine connected to the home gateway from the XTP cache.
"Recommend"	If there is a forecast of rain and if the washing machine has no drying feature, recommend hanging the clothes inside the house. If there is a forecast of rain and if the washing machine has a drying feature, recommend using it.

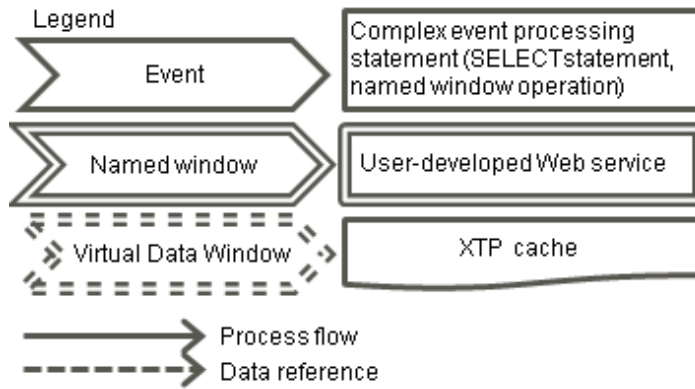
4. **Create an event flowchart.**

After refining the processing content is completed, summarize the event processing flow to create an event flowchart. An event flowchart associates events and their processing content in chart form. Create an event flowchart before describing the processes using complex event processing rule language, as it is useful for checking the processing content to be achieved.

 **Information**

Legend of event flowchart

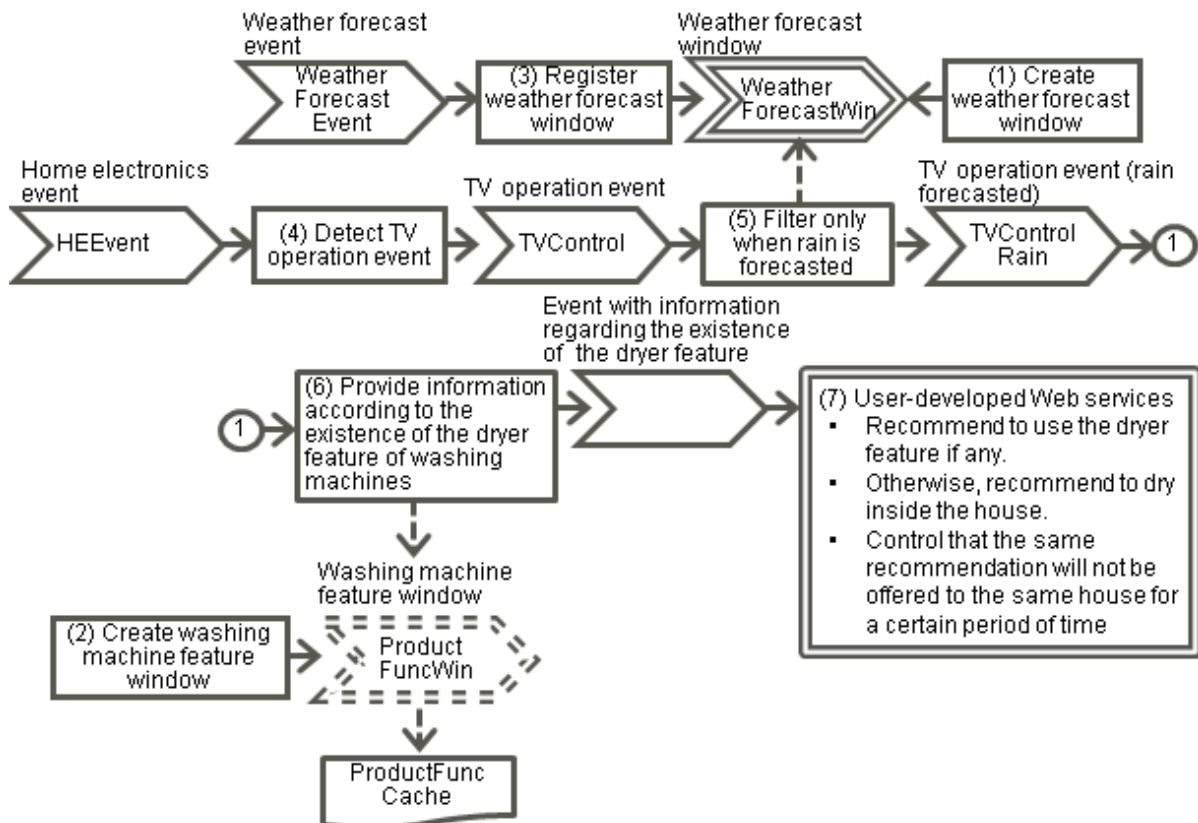
This is the legend for the event flowcharts to be used in this manual.



Example

Example of an event flowchart

This is an event flowchart for "When someone is home, if rain is likely, recommend using the drying feature of the washing machine".



1. Create a named window for retaining weather forecast events.
2. Create a Virtual Data Window (XTP cache) to see whether the washing machine has a drying feature.
3. Store weather forecast events in the weather forecast window.
4. Detect any TV control events from among the household appliance events.
5. Check the weather forecast from after the time that the TV control events occurred and leave only those events with a forecast of rain.

6. Search the washing machine feature window for the households for which a TV control was performed and for which there is a weather forecast of rain, and then add recommendation information according to whether the washing machine has a drying feature and send it to the user-developed Web service.
7. Based on the information received by the user-developed Web service, send a recommendation to the household. Control the user-developed Web service so that the same recommendation is not made to the same household twice within a fixed time period.

.....

5. Create complex event processing rules.

Describe complex event processing rules that correspond to the respective elements in the event flowchart.



Example

.....

Example of complex event processing rules that correspond to the event flowchart

Refer to Chapter 1, "Complex Event Processing Language Reference" in the *Developer's Reference* for information on the meanings of rules.

```

1      // 1. Create a named window for retaining weather forecast events.
2      @Name('WeatherWin')
3      create window WeatherForecastWin.std:unique(FORCASTIME).win:time(1 day)
4         (FORCASTIME long, WEATHER string);
5
6      // 2. Create a Virtual Data Window to check whether the washing machine has a
7      //      drying feature.
8      @VDW(cacheName='ProductFuncCache', keyProperty='gatewayId')
9      create window ProductFuncWin.isxtp:vdw() as (gatewayId string, dryFunc string);
10
11     // 3. Store weather forecast events in the named window.
12     @Name('InputWeather')
13     insert into WeatherForecastWin
14         select weathfore.FORCASTIME as FORCASTIME, weathfore.WEATHER as WEATHER
15            from WeatherForecastEvent as weathfore;
16
17     // 4. Detect TV control events.
18     @Name('ChkTVEvent')
19     insert into TVControl
20         select heevnt.gatewayId as gatewayId, heevnt.updateTime as updateTime
21            from HEEvent as heevnt
22            where heevnt.deviceCategory = 'Television';
23
24     // 5. Check the weather when TV control events occurred (detect forecast of rain).
25     // Use the following expression to evaluate the time zone of weather
26     // forecasts. Assume FORCASTIME is "long" (milliseconds):
27     //   timeOfWeatherForecast
28     //     <= updateTimeOfTVcontrolEvent
29     //     <= timeOfWeatherForecast + 1 hour (3600000 milliseconds)
30     @Name('GetRainyEvent')
31     insert into TVControlRain
32         select tvevnt.gatewayId as gatewayId
33            from TVControl as tvevnt unidirectional, WeatherForecastWin as weather
34            where updateTimeToMillis(tvevnt.updateTime) between weather.FORCASTIME
35               and (weather.FORCASTIME + 3600000)
36               and weather.WEATHER = 'rainy';
37
38     // 6. Check whether the washing machine has a drying feature and send the
39     //      recommendation.
40     // Check whether ProductFuncWin dryFunc property has drying feature ("1": Yes).
41     // Use the SOAP listener to send the output to the application.
42     //   'USE_DRY_FUNC' = Recommendation ID if there is drying feature
43     //   'HANG_LAUNDRY_INSIDE' = Recommendation ID if there is no drying feature

```

```

44      // Use the logging listener to log the output.
45      //   table      : set log storage area.
46      //   properties : set property names output (output results)
47      //               by delimiting with commas.
48      @Name('PutRecommend')
49      @SoapListener('soap-001')
50      @LoggingListener(table='/logsoap',properties='gatewayId,recommendId')
51      select tvevnt.gatewayId as gatewayId,
52             case when product.dryFunc = '1'
53                 then 'USE_DRY_FUNC'
54                 else 'HANG_LAUNDRY_INSIDE'end as recommendId
55      from TVControlRain as tvevnt, ProductFuncWin as product
56      where product.gatewayId = tvevnt.gatewayId;

```

5.4.3.3 SOAP Listener

Use the SOAP listener to send the processing results of complex event processing statements to a user-developed Web service using SOAP.

To use the SOAP listener, assign the "@SoapListener" annotation in front of the complex event processing statement (SELECT statement) for which the processing results are to be sent.

Refer to "[5.4.7 Designing a SOAP Listener Definition](#)" for information on the send destination of processing results and the content of SOAP messages.

Syntax

```

@SoapListener("SOAPListenerDefinitionID")
complexEventProcessingStatement(SELECT statement)

```

SOAP listener definition ID

Specify the development asset ID in the SOAP listener definition that defines information such as the URL that is the send destination of the Web service to which the processing results are to be sent.

Complex event processing statement (SELECT statement)

Specify the complex event processing statement (SELECT statement) for which the processing results are to be sent.

5.4.3.4 Logging Listener

Use the logging listener to log the processing results of complex event processing statements in the event log using logging.

To use the logging listener, assign the "@LoggingListener" annotation in front of the complex event processing statement (SELECT statement) for which the processing results are to be logged.

Syntax

```

@LoggingListener(table="logStorageArea", properties="propertyNameToBeOutput")
complexEventProcessingStatement(SELECT statement)

```

Log storage area

Use an absolute path to specify the path in the Hadoop system in which the event log is logged.

Even if events are to be logged in the engine log of the CEP Server (if "file" is specified in the "type" element of the engine configuration file), specify a virtual path name that begins with a slash (/) (for example, /eventName) to identify the events.



The "table" specification is mandatory. Even if a null value is specified, as in 'table=" "', the CEP engine will start normally but logging will not be performed.

Property name to be output

Of the processing results of the SELECT statement, specify the property name to be logged. Multiple properties can also be specified if delimited using commas (,).

For a property with nested processing results, use periods (.) between the nested properties to join them.

Example

If a "child" property is nested in a "parent" property

If a "child" property is nested in a "parent" property, as shown below, specify the "child" property by describing "parent.child".

```
<root>
  <parent>
    <child>aaa</child>
  </parent>
</root>
```

Note

The "properties" specification is mandatory. Even if a null value is specified, as in 'properties=" "', the CEP engine will start normally but logging will not be performed.

Complex event processing statement (SELECT statement)

Specify the complex event processing statement (SELECT statement) for which the processing results are to be logged using logging.

Note

- The logging to be defined in an event type definition and the logging to be defined in a rule definition (logging listener) can have separate log storage areas as output destinations. Note, however, that logging to a different Hadoop system is impossible.
- If the value of the output property is "null", this is converted to a blank space before being output to the event log.
- If the output property is a numeric item, this undergoes string conversion before being output to the event log.
- If an output property name that does not exist is specified, a blank space is output to the event log.
- If there are double quote marks (") in data, these will be output in duplicate within double quote marks within the data.

Output example of 'aa "bb" aa':

```
"aa " "bb" "aa "
```

5.4.4 Designing a Master Definition

If master data is to be used by high-speed filter rules, design a master definition.

This section explains the items to consider when designing a master definition, as follows:

- Timing of updates to the master data

Updates to the master data will be reflected when the CEP engine is next started.

If data that is being continually updated is to be referenced without stopping the CEP engine, consider using XTP collaboration in complex event processing.

- Amount of memory used by the master data

Master data is loaded in the CEP engine memory when the CEP engine starts. Refer to "[3.3.1.2 Amount of Memory when Master Data is used by the High-speed Filter](#)" for information on the memory required.

5.4.5 Designing an Event Type Definition (Filtered)

If filtered events are to be passed to complex event processing in a format that is different from their original format, due to processes such as join processing with master data using the high-speed filter, they require an event type definition that corresponds to the format used when they are passed to complex event processing.

This section explains the items to consider when designing an event type definition (filtered), as follows:

- Items that make up the filtered events

Check the items of the events output by high-speed filter rules.

The number of items specified in the output expression must match the number of items defined in the CSV column information.

- Event format

The event format that can be specified is CSV format only.

5.4.6 Designing XTP Collaboration

This section explains considerations when performing XTP collaboration, as well as how to use it, as follows:

- [Considerations when Using XTP Collaboration](#)
- [Using an XTP Cache](#)
 - [Creating a Virtual Data Window](#)
 - [Using a Virtual Data Window](#)

5.4.6.1 Considerations when Using XTP Collaboration

This section explains the items to consider when using XTP collaboration for referencing external data in complex event processing, as follows:

- Checking the necessity of XTP collaboration

Unlike when master data is used by the high-speed filter, XTP collaboration allows data that is being continually updated to be referenced (for the high-speed filter, if the CEP engine is running, updates to the data will not be reflected).

If the data to be referenced need not be continually updated, consider using master data in the high-speed filter.

- Structure of the XTP cache

Data in the XTP cache to be referenced by complex event processing rules is managed as entries in Key-Value format. Refer to "[5.4.10 XTP Cache](#)" for information on the Key-Value format used for storing in the cache to be used by complex event processing rules.

- Using an XTP cache

Use the Virtual Data Window feature to use an XTP cache with complex event processing. Refer to "[5.4.6.2 Using an XTP Cache](#)" for information on how to create a Virtual Data Window and how to use a cache via the created Virtual Data Window.

5.4.6.2 Using an XTP Cache

This section explains how to create a Virtual Data Window and how to use the created Virtual Data Window.

Creating a Virtual Data Window

Create a Virtual Data Window (hereafter, referred to as a "VDW") within complex event processing rules in order to reference an XTP cache from the rules.

Specifically, describe this as follows:

Syntax: If an event type ID is used

```
@VDW( cacheName="cacheName", keyProperty="keyPropertyName" )  
create window windowName.isxtp:vdw() as eventTypeID;
```

Syntax: If type information is specified directly

```
@VDW(cacheName="cacheName", keyProperty="keyPropertyName")
create window windowName.isxtp:vdw() as (propertyName type, propertyName type, ...);
```

- Create a VDW using a combination of the "@VDW" annotation and a CREATE WINDOW statement.
To reference the cache entity, the *cacheName* and the *keyPropertyName* for referencing the data must be set. To set these, the "@VDW" annotation must be specified.
- In *cacheName*, specify the name of the cache to be used.
- In *keyPropertyName*, specify the property name for identifying the entry. If an event type ID is to be used, specify the property name of the specified event type. If type information is to be specified directly, any name can be specified. The type of the specified property must be a type that corresponds to the "Key" class in the cache.
- In *windowName*, specify any name. The name specified here will be able to be used to reference from SELECT statements in other complex event processing rules.
- In *eventTypeID*, specify the event specified in the CSV format event type definition. (An XML format event type definition cannot be specified.)
- In *propertyName type*, the property name and type that corresponds to the "java.util.HashMap" key to be set in "Value" in the cache must be specified. The property name and its type specified in *keyPropertyName* must also be specified.
- If the specified *propertyName* has not been set in "java.util.HashMap" to be set in "Value" in the cache, it will be treated as a null by the complex event processing rules.

Example

Description example for creating a Virtual Data Window (VDW)

This is an example of creating a VDW (MarketWindow) to reference an XTP cache (MARKET).

```
@VDW(cacheName="MARKET", keyProperty="code")
create window MarketWindow.isxtp:vdw() as (code string, high int, low int);
```

- "code" is specified as a key property.
- "code (string type)", "high (int type)", and "low (int type)" are defined as the properties.

Using a Virtual Data Window

Use a created Virtual Data Window in the same way as an ordinary window.

Example

Example of using a created Virtual Data Window (VDW)

This is an example of referencing a VDW (MarketWindow) in an event (TickerWindow) to obtain the data of the VDW events (cache entries) that meet the condition.

```
select M.code from MarketWindow as M, TickerWindow as T
where M.code = T.code and ( T.price > M.high or T.price < M.low);
```

Note

Notes on using a Virtual Data Window

The restrictions below apply to a WHERE clause that can be used when accessing the information stored in a Virtual Data Window.

These restrictions apply when records are to be identified from a cache.

If the records have already been identified, the rules can be described in the usual way and the following restrictions will not apply:

- Only key properties specified using the "@VDW" annotation can be specified to identify records.
- Only the "=" comparison operator can be used for key properties that can be specified to identify records.

Description example

After identifying records using "M.code = T.code", compare other items.

```
@VDW(cacheName="MARKET", keyProperty="code")
create window MarketWindow.isxtp:vdw() as (code string, high int, low int);

select M.code from MarketWindow as M, TicketWindow as T
where M.code = T.code and (T.price > M.high or T.price < M.low)
```

The table below shows valid and invalid description examples for the cache above.

No.	Description example	Valid?	Explanation
1	<code>select * from MarketWindow where code = '1111'</code>	Yes	This is acceptable because the "=" operation is used for a key property.
2	<code>select * from MarketWindow where high = 1000</code>	No	This is unacceptable because a WHERE clause is specified for other than a key property.
3	<code>select * from MarketWindow where code > '1111'</code>	No	This is unacceptable because a WHERE clause is only specified for a key property but is performing an operation other than "=".

5.4.7 Designing a SOAP Listener Definition

A SOAP listener definition specifies the interface of the user-developed Web service to which the processing results of complex event processing rules are to be sent. This section explains the points to consider in a SOAP listener definition, as follows:

- Unit of creation

Create a SOAP listener definition for each user-developed Web service. One SOAP listener definition can be used when a generic application is created, and an application can also be separated by pattern matching.

- Development asset ID

Specify a name for the SOAP listener definition that is unique in the CEP engine in which it will be deployed and that suggests the processing content called when rule matching finds a match.

A simple example is where a name such as *powerOn* is used for a process to notify when the power is on.

- Association between a SOAP listener definition and a user-developed Web service

Check that the property value of the corresponding complex event processing rule matches the user-developed Web service to be called, based on the WSDL (interface definition) of the application. Even if an existing Web service is to be used, consider factors such as whether the interface can be used as it is or if it needs to be changed. The aspects that must match are as follows:

- Parameter names

The parameter names to be passed to the user-developed Web service must match the property names selected from within the complex event processing rules.

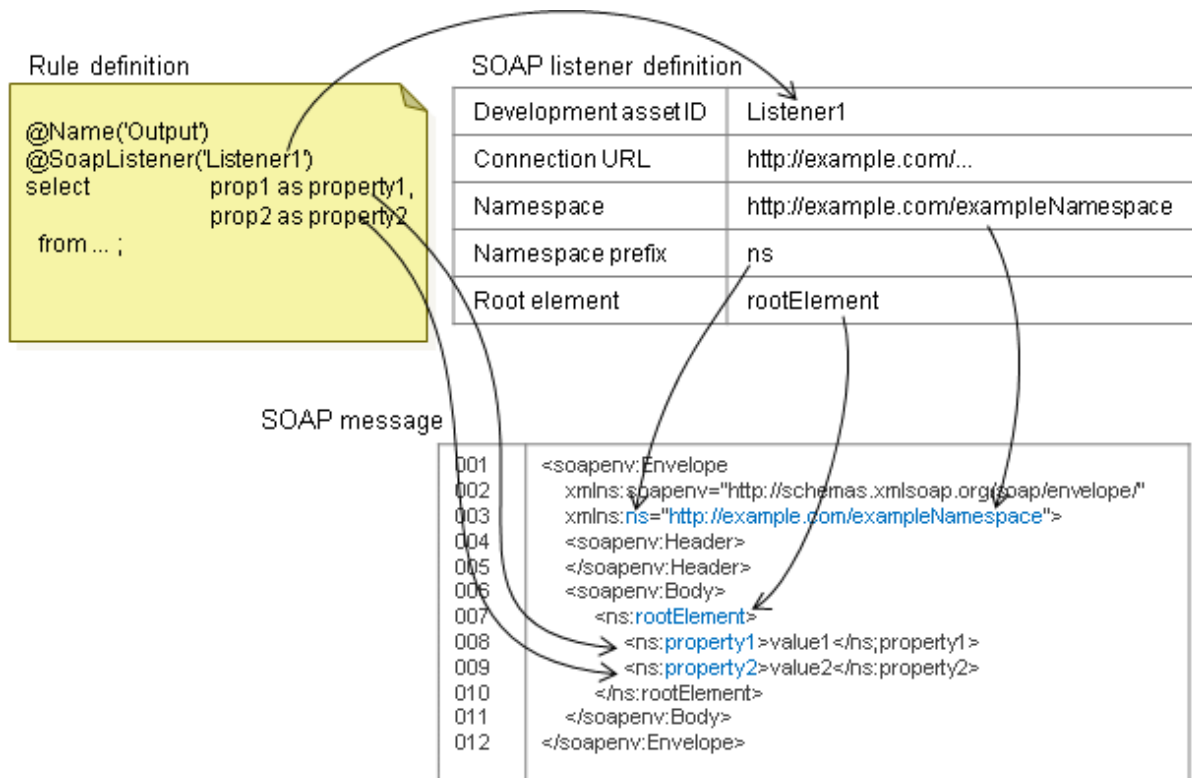
If the property names selected from within the complex event processing rules are to be used as the parameter names to be passed to the user-developed Web service, use them by defining aliases within the rules.

- Parameter types

The parameter types must also match the types used from within the complex event processing rules.

Below is an example of the association between a rule definition and a SOAP listener definition. The SOAP messages to be sent to a user-developed Web service are generated from the rule definition and from the SOAP listener definition associated with it.

Figure 5.4 Example of the association between a rule definition and a listener definition, and the SOAP messages to be sent



5.4.8 Event Data (for Testing)

Consider using event data sent from an event sender application to check the operation of created rules.

Pay attention to the following points when considering this:

- Testing scenario (content of event data required)

Provide event data based on the scenario envisaged for the rules being designed and developed.

Provide data by envisaging that abnormal data as well as normal data will be sent.

- Format of event data

Provide event data using a format that suits the event sender application to be used.

Provide data in CSV format if the event sender sample program included in the samples for this product is to be used as an event sender application for testing.

5.4.9 Master Data (for the High-speed Filter)

Pay attention to the following points when considering the use of master data by the high-speed filter:

- Format of the master data

Refer to "5.4.9.1 Format of Master Data".

5.4.9.1 Format of Master Data

The master data consists of data files in CSV format and a schema file.

Refer to "9.5 CSV Format Supported" for information on the CSV format.

The character code that can be used is UTF-8.

Specify LF or CRLF as the newline code to be described at the end of records.

Schema file

This is a file in which only lines of item names are described.

If a schema file contains information other than item names, an error will occur.

Data file

This is a file in which data is stored. If the item names are described in the first line of the data file, skipping the first line can be set in the master definition. Refer to "[9.2.3 Master Definition File](#)" for details.



Example

Example of a schema file

```
"Kbn", "Number", "Code", "Name", "Value", "Total", "Comment"
```

Example of a data file

```
"01", "1001", "AAA", "BlockA", "1,000", "1,000", "Comment: Memo number 4023"  
"02", "1001", "BBB", "BlockB", "", "1,200", "Comment: Memo number 4023"  
"03", "1002", "CCC", "BlockC", "800", "800", "Comment: Memo number 4023"
```

5.4.10 XTP Cache

Pay attention to the following points when considering using XTP collaboration to reference an external XTP cache from complex event processing rules:

- Initial data of the cache

Initial data must be provided for the cache.

Select a data format according to the specifications of the cache access application for update to be used.

- Format of the XTP cache

An XTP cache must follow a prescribed format.

If an existing cache is to be used, check that it is consistent with "[5.4.10.1 XTP Cache Compatible Formats](#)".

5.4.10.1 XTP Cache Compatible Formats

An XTP cache consists of a "Key" and a "Value". A cache to be referenced by this product must be given the following structure:

Type used for "Key"	Type used for "Value"
java.lang.String	java.util.HashMap<java.lang.String, java.lang.Object>

Each property to be specified in a complex event processing rule corresponds to each "HashMap" element above.

The table below shows the type that are compatible for the value of each "HashMap" element in the cache and the corresponding type in complex event processing rules.

Type used in each "HashMap" element	Corresponding type in complex event processing rules
java.lang.String	string
java.lang.Boolean	bool/boolean
java.lang.Byte	byte
java.lang.Integer	int/integer
java.lang.Long	long
java.lang.Float	float

Type used in each "HashMap" element	Corresponding type in complex event processing rules
java.lang.Double	double

Each "HashMap" to be specified in "Key" and "Value" for an XTP cache must be associated as follows:

- In the "HashMap" key, set the value to be specified in "keyProperty" of the Virtual Data Window.
- In "Key" in the XTP cache, set the same value as that of the "HashMap" key above.

Refer to "[5.4.14 Designing a Cache Access Application for Update](#)" for information on the actual method of use.

5.4.11 Designing an Event Sender Application

Consider which method to use to send events to the CEP engine.

Here, consideration must be given to which communication method to use for the input events and which application to use to send the events, according to that communication method.

Communication method for input events

Consider which communication method to use when the CEP engine is receiving events, according to the characteristics of the input events.

If an existing system that is to be used as the event issuer has an event sending feature, check which communication methods the existing system can use.

With this product, the following three communication methods can be selected:

- SOAP
- HTTP
- Socket

Select a communication method according to the characteristics of the system or device that is to be used as the event sender, and the desired processing performance.

Refer to Chapter 3, "Input Adapter Reference" in the *Developer's Reference* for information on the features of each of these communication methods.

Event sender application

Consider which system or application to use to send events to the CEP Server.

- There are two types of event sender applications: systems that send events during normal business operation and those that send events at any time during rule testing. For the latter, consider using the event sender sample program supplied with the samples of this product, if the volumes of event data are small. Refer to "[5.9.6 Event Sender Sample Program](#)" for details.
- If the system issuing the events can send the events using a communication method supported by this product, check that the CEP engine can receive events using that feature.
- If the event issuer has no event sending feature, or if it does not support a communication method of the input adapter of this product, consider developing an application to send events to the CEP Server (event sender application).
- The design and development of an event sender application must be carried out according to the specifications of the server or device that is to be used as the event sender. Refer to Chapter 3, "Input Adapter Reference" in the *Developer's Reference* for detailed specifications on each communication method and for information on the sample event sender application.

5.4.12 Designing a User-developed Web Service

Refer to the application server manuals for information on designing and developing a user-developed Web service. This section explains the points to consider when designing a user-developed Web service, as follows.

- Integrating similar processes

If multiple recommendations are to be made to an individual, rather than separating Web services by recommendation, group them into one Web service to create one SOAP listener definition, and then use calling parameters to divide up the Web service processes to be executed.

- Separating processes with different responses or targets

If processes have different responses or targets, as with "recommend to a person" and "control a device", the internal logic of the Web service and testing methods will differ, so consider developing such processes as separate applications.

- Execution environment of a Web service with operation checked

This product checks the operation of user-developed Web services in the Web service execution environments below. If executing using another product, perform sufficient connection testing.

Product name	Details
Interstage Application Server	Checked if a Java EE Web service is used.
Apache Axis2	Open source Web service execution framework. The URL is as follows: http://axis.apache.org/axis2/java/core/

If a product other than those above is to be used, check that "Content-Length", "Content-Type", and each value have been set correctly in the HTTP header that will form the response from the user-developed Web service to the CEP engine.

5.4.13 Designing an Event Log Analysis Application

If logging is to be used to accumulate event logs in a Hadoop system, design an application to analyze the content of the accumulated event logs. The application will use the Hadoop API and operate on the Hadoop system.

Refer to the Interstage Big Data Parallel Processing Server (hereafter, referred to as "BDPP") manuals for information on designing and developing applications to operate on a Hadoop system.

The data formats of the event logs to be analyzed by this application are shown below.

5.4.13.1 Output Destination and File Format of an Event Log

Event logs are output to a log storage area specified in the event type definition or in the logging listener in a complex event processing statement. The log storage area that will be the output destination is generated automatically.

If the output destination is a Hadoop system, the details are as follows:

Output destination

The output destination can be changed using the value specified in the `directory` element of the engine configuration file.

If a directory name is specified in the `directory` element, the output destination will be a path made by joining the following values:

- Value set in "`pdfs.fs.local.basedir`" (*1)
- Directory name specified in the engine configuration file
- Log storage area specified in the event type definition or logging listener
- Automatically generated log file name

*1: "`pdfs.fs.local.basedir`" is the Hadoop mount directory. Refer to the BDPP manuals for details.

If a slash (/) only is specified in the `directory` element, the output destination will be a path made by joining the following values:

- Value set in "`pdfs.fs.local.basedir`"
- Log storage area specified in the event type definition or logging listener
- Automatically generated log file name

Example

Example of output destination

The output destination will be `"/mnt/pdfs/hadoop/tmp/logFileName"` for the following conditions:

- If the value set in `"pdfs.fs.local.basedir"` is `"/mnt/pdfs"`; and
- If `"hadoop"` is specified as the directory name in the engine configuration file; and
- If `"/tmp"` is specified as the log storage area specified in the event type definition or logging listener of the complex event processing statement

The output destination will be `"/mnt/pdfs/tmp/logFileName"` for the following conditions:

- If the value set in `"pdfs.fs.local.basedir"` is `"/mnt/pdfs"`; and
- If a slash (/) is specified as the directory name in the engine configuration file; and
- If `"/tmp"` is specified as the log storage area specified in the event type definition or logging listener of the complex event processing statement

Note

If the output destination of the event log is duplicated and the format of the event data is the same, event data of a different event type will be output to the same file. If analysis is to be performed by event type or by output by logging listener, separate the output destinations.

Log file format

The format will be Hadoop SequenceFile (binary file) format.

Log file name

A log file will be automatically generated in the log storage area using the file name shown below.

This file will be renamed with the `".done"` extension in 300 seconds by default.

```
dateTime_VMname_branchNumber
```

- `dateTime`: yyyyMMddHHmmssSSS
- `VMname`: `processID@CEPserverHostName`
- `branchNumber`: 0000000001 to 0000000122

Point

A file with the `".done"` extension will be analyzed by the event log analysis application. Move it to an arbitrary directory to analyze it.

Note

A file with an extension other than `".done"` is a file that is being output, so do not perform an operation on it.

Upper limit of file size

The upper limit of the file size is LONG MAX ($2^{63} - 1$).

Upper limit of number of files

None

Key of SequenceFile

The date and time information (yyyyMMddHHmmss) will be the key. The corresponding Hadoop type (API) is "org.apache.hadoop.io.Text".

The date and time above will be the date and time at which the event data was written. (This may differ from the date and time at which the CEP engine received the events.)

Value of SequenceFile

Input events are output as they are. The corresponding Hadoop type (API) is "org.apache.hadoop.io.BytesWritable".

Compression format of SequenceFile

Record compression

Versions of SequenceFile

6



Information

If outputting to the engine log

Input events are output to the engine log unchanged.

5.4.14 Designing a Cache Access Application for Update

If an XTP cache is to be referenced in complex event processing rules, consideration must be given to an application for storing initial data in the XTP cache. If an existing cache can be used, use an existing cache access application.

Refer to the Interstage eXtreme Transaction Processing Server manuals for information on designing and developing a cache access application.

Refer to "5.4.10.1 XTP Cache Compatible Formats" for information on the format of the cache to be updated.

Each "HashMap" to be specified in "Key" and "Value" for an XTP cache must be associated as follows:

- In the "HashMap" key, set the value to be specified in "keyProperty" of the Virtual Data Window.
- In "Key" in the XTP cache, set the same value as that of the "HashMap" key above.



Example

Example of a Virtual Data Window

```
@VDW(cacheName='CacheA', keyProperty='key')
create window windowName.isxtp:vdw() as (key string, address string);
```

Example of a program to add entries to an XTP cache

```
import javax.cache.*;
import java.util.HashMap;
(...)
CacheManager cacheManager = Caching.getCacheManager();
Cache<String,HashMap> cache = cacheManager.getCache("CacheA"); ... 1.
String keyProperty = "key"; ... 2.
String keyValue = "1"; ... 2.
HashMap value = new HashMap(); ... 3.
value.put(keyProperty, keyValue); ... 3.
value.put("address", "Boston"); ... 3.
cache.putIfAbsent(keyValue, value); ... 4.
(...)
```

1. Obtain the cache to be used by the Virtual Data Window.

2. Create the key property name and its value to be specified in "keyProperty" of the Virtual Data Window. Here, "key" is set as the "keyProperty".
3. Create the "java.util.HashMap" object to be stored as "Value" of the cache. Here, a "java.lang.String" type value is set in the "address" property. Also, ensure that the "keyProperty" created in "2." is set.
4. Add entries to the cache.

5.5 Development

This section explains the development tasks (such as coding) for development assets.

5.5.1 Developing a Definition File

This section explains how to develop a definition file, as follows:

- [Creating an Event Type Definition File](#)
- [Creating a Rule Definition File](#)
 - [Debug log listener](#)
- [Creating a Master Definition File](#)
- [Creating a SOAP Listener Definition File](#)

5.5.1.1 Creating an Event Type Definition File

Create an event type definition file similar to the items of the event type definition designed previously.

Refer to "[9.2.1 Event Type Definition File](#)" for information on the format of an event type definition file.



Example

Example of an event type definition

Below is an example of an event type definition for an event in XML format.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<eventType xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="EVENTTYPE_01">
  <comment>Event type definition 01</comment>
  <type>XML</type>
  <xmlSchema>
    <![CDATA[
      <xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
        xmlns="http://dataaccesscontrol.sspf.fujitsu.com/namespace/xmlmessage"
        targetNamespace="http://dataaccesscontrol.sspf.fujitsu.com/namespace/xmlmessage">
        <xs:element name="messagedata">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="memberID" type="xs:string" />
              <xs:element name="areaID" type="xs:string" />
              <xs:element name="status" type="xs:string" />
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:schema>
    ]]>
  </xmlSchema>
  <root>messagedata</root>
  <useLogging>true</useLogging>
  <loggingTableName>/echonet</loggingTableName>
```

```
<useCep>true</useCep>
</eventType>
```

5.5.1.2 Creating a Rule Definition File

Create a rule definition file similar to the items of the rule definition designed previously.

Refer to "[9.2.2 Rule Definition File](#)" for information on the format of a rule definition file.

Also assign a debug log listener (`@DebugLogListener`) in complex event processing rules, to check operation in an integration test. Refer to "[5.5.1.2.1 Debug log listener](#)" for information on the debug log listener.



Example

Example of a rule definition

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<rule xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="RULE_01">
  <comment>Rule definition 01</comment>
  <filter>
    <![CDATA[
      on EVENTTYPE_01 {
        if ($status == 'Walking') then output() as EVENTTYPE_01;
      }
    ]]>
  </filter>
  <statements>
    <![CDATA[
      @SoapListener('LISTEN_01')
      @DebugLogListener
      select * from EVENTTYPE_01 where areaID = '1010';
    ]]>
  </statements>
</rule>
```

This example describes the following rules:

High-speed filter rule

If the status item contents (string) extracted from the "EVENTTYPE_01" event type input event is "Walking", it is transferred to Complex Event Processing.

Complex event processing rule

This rule notifies the SOAP listener of event data with the "EVENTTYPE_01" event type, and simultaneously outputs debug information to the engine log.

5.5.1.2.1 Debug log listener

Using the debug log listener allows logs for debugging to be output to the engine log when a complex event processing statement is executed.

Specifically, assign the "`@DebugLogListener`" annotation in front of the target complex event processing statement.

Also, using the "`@Name`" annotation to give a name to the target complex event processing statement at the same time will allow the output information of the engine log to be found easily.

Syntax

```
@Name ( " name " )
@DebugLogListener
complexEventProcessingStatement
```

name

This is output at the same time as debug log output and allows the output information to be found easily.

If the same name is given to multiple complex event processing statements, names will be automatically assigned using a format of two hyphens (--) and a numeric will be appended to the end of each name.

complexEventProcessingStatement

This is the complex event processing statement to be the target of debug log output.

Note

The debug log listener may cause a decline in performance, so avoid use during normal business operation.

Information

If the "@Name" annotation is not specified

If the "@Name" annotation is not specified, an automatically assigned unique name such as "0b0562a2-56e7-4cf3-a520-cb1e16ef2992" will be output in place of the value of the "@Name" annotation.

Example of output when the "@Name" annotation is specified (@Name (' EPL ') specified)

```
2012-07-09 19:32:35,495 [DEBUG] EPL:length=1
EPL[0]
```

Example of output when the "@Name" annotation is not specified

```
2012-07-09 19:35:55,244 [DEBUG] 34b1785f-900c-4420-b2bf-ea53aa368b07:length=1
34b1785f-900c-4420-b2bf-ea53aa368b07[0]
```

5.5.1.3 Creating a Master Definition File

Create a master definition file similar to the items of the master definition designed previously.

Refer to "[9.2.3 Master Definition File](#)" for information on the format of a master definition file.

Example

Example of a master definition

This is an example of a master definition (development asset ID: MASTER_01) where the schema file is "/var/tmp/SchemaFile01.csv" and the data file is "/var/tmp/MasterFile01.csv".

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<master xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="MASTER_01">
  <comment>Master definition 01</comment>
  <schemaFile>/var/tmp/SchemaFile01.csv</schemaFile>
  <dataFile>/var/tmp/MasterFile01.csv</dataFile>
  <skipHeader>>false</skipHeader>
</master>
```

5.5.1.4 Creating a SOAP Listener Definition File

Create a SOAP listener definition file similar to the items of the SOAP listener definition designed previously.

Refer to "9.2.4 SOAP Listener Definition File" for information on the format of a SOAP listener definition file.



Example

Example of a SOAP listener definition

This example defines the notification of a message (event) that includes a SOAP body saying that the root element is "cep" in the user-developed Web service with the connection destination URL "http://192.168.11.249/WebServWAR/MyApp1Service".

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<soapListener xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="LISTEN_01">
  <comment>SOAP listener definition 01</comment>
  <url>http://192.168.11.249/WebServWAR/MyApp1Service</url>
  <namespace>http://webservice/</namespace>
  <prefix>ns</prefix>
  <method>cep</method>
</soapListener>
```

5.5.2 Preparing Data

This section explains how to prepare the data to be referenced from rules.

5.5.2.1 Preparing Event Data (for Testing)

Prepare (create) event data to be used in checking the operation of rules.

Use an event data format that suits the event sender application to be used.



Example

This is an example of event data for the event sender sample program supplied with the samples of this product.

/opt/FJSVcep/sample/sample1/event/CouponEvent.csv

```
"STR0001", "CPN0001", "30"
```

5.5.2.2 Preparing Master Data (for the High-speed Filter)

Prepare (create) a schema file and data files of the master data for the high-speed filter.

Refer to "9.2.3 Master Definition File" for information on the schema file and data files.



Example

Example of a schema file

```
"Kbn", "Number", "Code", "Name", "Value", "Total", "Comment"
```

Example of a data file

```
"01", "1001", "AAA", "BlockA", "1,000", "1,000", "Comment: Memo number 4023"
"02", "1001", "BBB", "BlockB", "", "1,200", "Comment: Memo number 4023"
"03", "1002", "CCC", "BlockC", "800", "800", "Comment: Memo number 4023"
```


5.5.2.3 Preparing Data to be Stored in an XTP Cache

Prepare (create) the data that is to be initially stored in the XTP cache to be referenced from complex event processing rules.

If an existing cache is to be used, this data need not be prepared (created).

Select a data format according to the specifications of the cache access application for update to be used.

5.5.3 Developing a Collaboration Application

This section explains how to develop a collaboration application, as follows:

- [Developing an Event Sender Application](#)
- [Developing a User-developed Web Service](#)
- [Developing an Event Log Analysis Application](#)
- [Developing a Cache Access Application for Update](#)

5.5.3.1 Developing an Event Sender Application

Develop an application to send events to the CEP engine.

Develop the application according to the type of input adapter to be used.

Refer to Chapter 3, "Input Adapter Reference" in the *Developer's Reference* for information on sample programs by input adapter type.

5.5.3.2 Developing a User-developed Web Service

Develop a Web service application to be called from the SOAP listener.

5.5.3.2.1 Web service development procedure

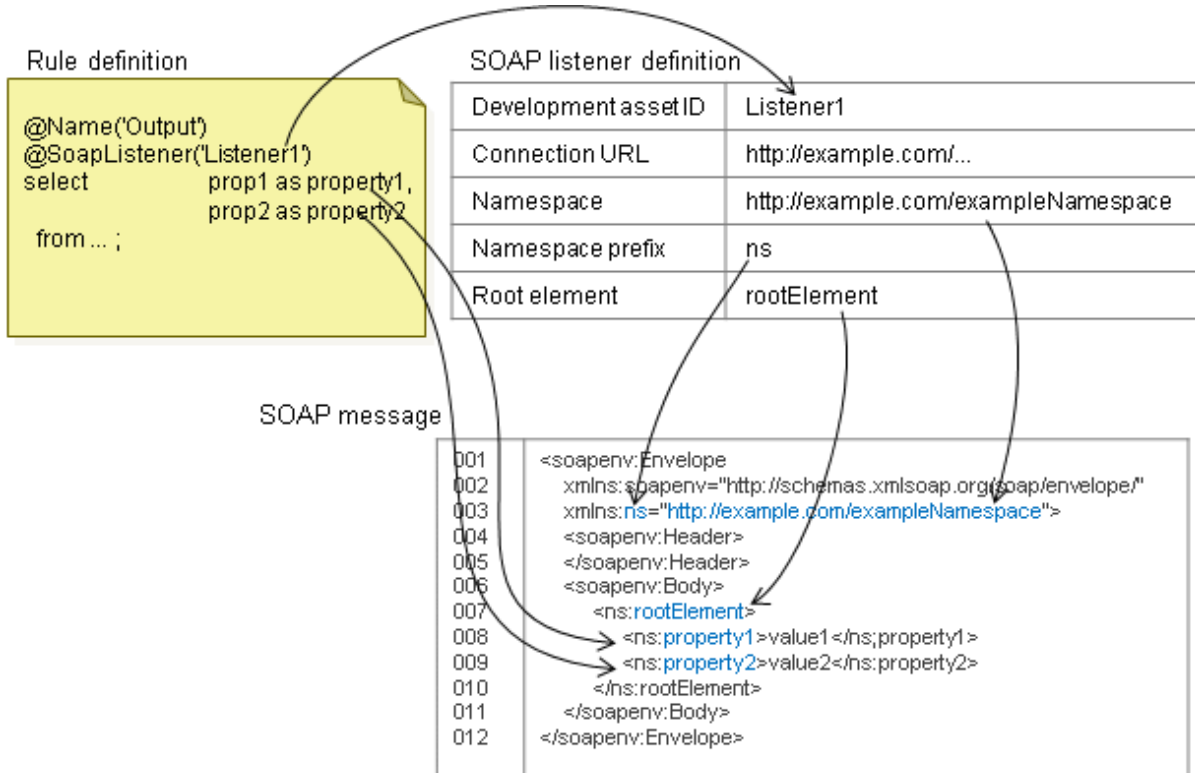
This section explains the Web service development procedure.

1. **Create a WSDL.**

Create a WSDL (interface definition) for a user-developed Web service from the interface information (Web service URL, namespace, prefix, and method) of the Web service to be called, that was defined in the SOAP listener definition, and from the parameters that are specified in the complex event processing statements (SELECT statements).

Below is an example of the association between the complex event processing statement (SELECT statement) for detecting the target events, the SOAP listener definition associated with that rule, and SOAP messages generated from the SOAP listener.

Figure 5.5 Association between a rule definition, listener definition, and SOAP messages to be sent



A sample WSDL for a Web service for receiving these SOAP messages is shown below. This WSDL defines a message receive-only (one-way) Web service. If it is implemented as a Web service that returns a response to the CEP engine (request-response), the CEP engine ignores this response.

Table 5.2 Sample WSDL

001	<?xml version='1.0' encoding='UTF-8' ?>
002	<definitions
003	targetNamespace="http://example.com/exampleNamespace"
004	xmlns="http://schemas.xmlsoap.org/wsd1/"
005	xmlns:tns="http://example.com/exampleNamespace"
006	xmlns:soapbind="http://schemas.xmlsoap.org/wsd1/soap/"
007	xmlns:xs="http://www.w3.org/2001/XMLSchema">
008	
009	<types>
010	<xs:schema elementFormDefault="qualified"
011	targetNamespace="http://example.com/exampleNamespace" >
012	<xs:element name="rootElement">
013	<xs:complexType>
014	<xs:sequence>
015	<xs:element name="property1" type="xs:string" />
016	<xs:element name="property2" type="xs:string" />
017	</xs:sequence>
018	</xs:complexType>
019	</xs:element>
020	</xs:schema>
021	</types>
022	</definitions>
023	
024	<message name="notifyMessage">
025	<part name="body" element="tns:rootElement" />

026	</message>
027	
028	<portType name="eventReceiverPortType">
029	<operation name="notifyOperation">
030	<input message="tns:notifyMessage" />
031	</operation>
032	</portType>
033	
034	<binding name="eventReceiverSOAPBinding" type="tns:eventReceiverPortType">
035	<soapbind:binding transport="http://schemas.xmlsoap.org/soap/http"
036	style="document" />
037	<operation name="notifyOperation">
038	<soapbind:operation soapAction="" />
039	<input>
040	<soapbind:body use="literal" />
041	</input>
042	</operation>
043	</binding>
044	
045	<service name="eventReceiverService">
046	<port name="eventReceiverSOAPPort" binding="tns:eventReceiverSOAPBinding">
047	<soapbind:address location="http://example.com/serviceEndPoint" />
048	</port>
049	</service>
050	</definitions>

Modifying, as follows, the underlined parts according to a rule definition or SOAP listener definition allows the WSDL above to be used to generate a template of the Web service to be created:

- Line numbers 003, 005, and 011

Use the value of the namespace in the SOAP listener definition as the target namespace of the WSDL ("targetNamespace" attribute of the "definitions" element), the declaration of the target namespace prefix ("xmlns:tns" attribute of the "definitions" element), and the target namespace of the XML schema in the WSDL ("targetNamespace" attribute of the "xs:schema" element in the "types" element).

- Line numbers 012 and 025

Use the value of the root element in the SOAP listener definition as the name of the root element of messages to be defined in the XML schema in the WSDL ("name" attribute of the "xs:element" element in "xs:schema" element in the "types" element) and the element defined as a message of the WSDL ("element" attribute of the "part" element in the "message" element).

- Line numbers 015 and 016

Use the output property name of the complex event processing statement (SELECT statement) as the name of the subelement of the root element to be defined in the XML schema in the WSDL ("name" attribute of the "xs:element" element under the root element definition in the "types" element).

Create as many similar lines as the number of properties to be output. In the example above, this element type ("type" attribute of the "xs:element" element) is specifying a string ("xs:string"). Setting this to suit the property type will, depending on the tool used, generate source code according to the type, so type conversion will no longer be required in the program.

- Line number 047

In the final WSDL for service publishing, the actual URL of the service will be entered in the "location" attribute of the "soapbind:address" element under the "service" element, and this will also be the value of the connection URL in the listener definition. However, this connection URL is often unsure at development, so there is no problem with leaving it as in the sample above.

2. Develop the Web service application.

Based on the created WSDL, output a template for the Web service application from the development tool being used, and then add the application logic to it.

5.5.3.3 Developing an Event Log Analysis Application

Develop an application to analyze event logs. Develop the application using the Hadoop API. Refer to the Interstage Big Data Parallel Processing Server manuals for details.

5.5.3.4 Developing a Cache Access Application for Update

Develop a cache access application. Refer to the Interstage eXtreme Transaction Processing Server manuals for details.

5.6 Deploying Development Assets

This section explains how to deploy development assets. The deployment tasks are as follows:

- [Deploying Definition Information](#)
- [Providing Data](#)
- [Deploying a Collaboration Application](#)

5.6.1 Deploying Definition Information

Store the developed definition information in the CEP Server, and then use `cepdeployrsc` to deploy the definition information. Refer to "6.1.4.1 [Deploying Definition Information](#)" for information on how to deploy definition information.

Below is an example of deploying definition information.



Example

Example of deploying definition information

This is an example of command execution when the following definition information (definition file) is deployed in a CEP engine (CepEngine):

```
/application
+-- EVENT01.xml      ... Event type definition      (development asset ID: EVENT01)
+-- RULE01.xml       ... Rule definition           (development asset ID: RULE01)
+-- MASTER01.xml    ... Master definition          (development asset ID: MASTER01)
+-- LISTENER01.xml  ... SOAP listener definition (development asset ID: LISTENER01)
```

The example of command execution is as follows:

```
$ cepdeployrsc eventtype -e CepEngine -f /application/EVENT01.xml<ENTER>
Are you sure you want to deploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
$ cepdeployrsc rule -e CepEngine -f /application/RULE01.xml<ENTER>
(...)
$ cepdeployrsc master -e CepEngine -f /application/MASTER01.xml<ENTER>
(...)
$ cepdeployrsc listener -e CepEngine -f /application/LISTENER01.xml<ENTER>
(...)
```

5.6.2 Providing Data

Provide the data required for rule operation.

Event data (for testing)

Store event data (for testing) in the event sender system.

If an event sender sample program is to be used as the event sender application, store this data in the CEP Server.

Log storage area

Create a directory to be used as the log storage area, and set the write permission for the engine execution user created during installation. This task is usually performed by the system administrator of the CEP Server.

Information

If the directory to be used as the log storage area does not exist, it will be generated automatically at logging.

Use the engine execution user permission to create the directory.

Example

Directory creation

This is an example of creating a directory when the engine execution user and group are "isbdcep".

Log in as a superuser and execute the `hadoop` command to perform these tasks.

If a user-defined directory name is specified in the directory element in the engine configuration file

Create a directory with the specified name in the root directory of the Hadoop file system.

Change the owner of created directory to give write permissions to the engine execution user.

The log storage area to be specified in the event type definition or logging listener will be created automatically.

Below is an example when "hadoop" is specified in the `directory` element.

```
# hadoop fs -mkdir /hadoop <ENTER>
# hadoop fs -chown isbdcep:isbdcep /hadoop <ENTER>
```

If a slash (/) only is specified in the directory element in the engine configuration file

No further action is required if the engine execution user can write to the root directory of the Hadoop file system.

If the engine execution user does not have write permissions, create a directory with the same name as the log storage area to be specified in the event type definition or logging listener.

Change the owner of created directory to give write permissions to the engine execution user.

Below is an example when "/tmp" is specified as the log storage area to be specified in the event type definition or logging listener.

```
# hadoop fs -mkdir /tmp <ENTER>
# hadoop fs -chown isbdcep:isbdcep /tmp <ENTER>
```

Master data

Store a schema file and data files in the path specified in the master definition.

The stored files must have the file read permission set for the engine execution user created during installation.

XTP cache

If a new XTP cache is to be provided, use the cache access application for update that was developed separately to store the initial data in the cache.

Note that the XTP Server node must be set beforehand and a cache to be used as the storage destination must also be created beforehand. Refer to the Interstage eXtreme Transaction Processing Server manuals for information on how to create a cache.

If an existing cache is to be used as it is, no further action required.

5.6.3 Deploying a Collaboration Application

Deploy the developed collaboration application.

Event sender application

Deploy an event sender application in the event sender system. Perform the deployment according to the application deployment method of the event sender system.

If an event sender sample program is to be used as the event sender application, no deployment is required.

User-developed Web service

Deploy a user-developed Web service in an application server. Perform the deployment according to the method in the manual of the application server to be used.

Event log analysis application

Deploy an event log analysis application in a Hadoop system. Perform the deployment according to the method in the Interstage Big Data Parallel Processing Server manual.

Cache access application for update

Deploy a cache access application for update in an application server (Java EE). Perform the deployment according to the method in the Interstage Application Server manual.

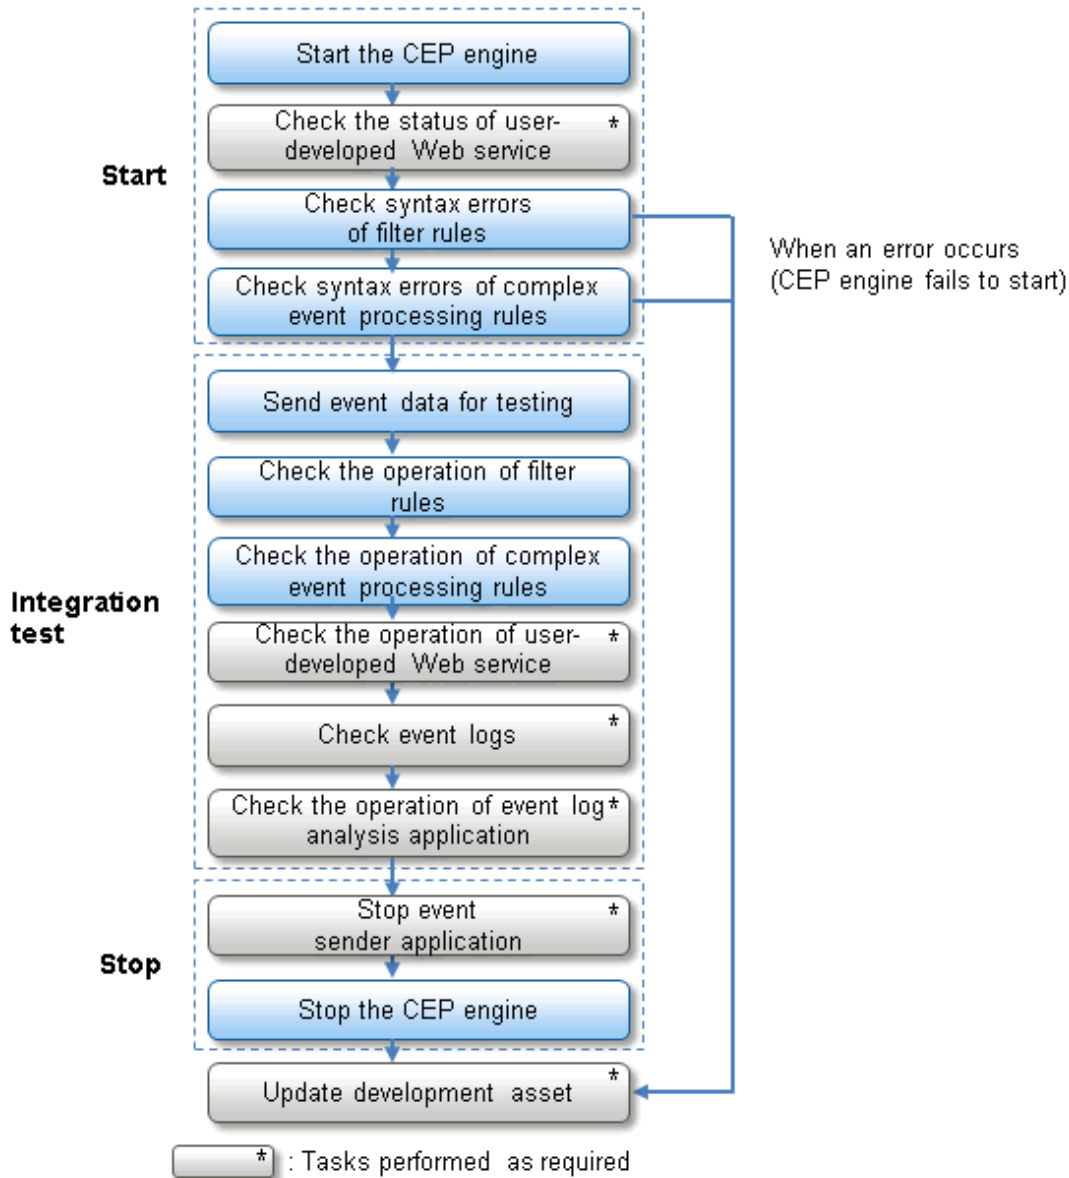
5.7 Integration Test

Send test data to a CEP engine in which development assets have been deployed in order to check operation.

5.7.1 Integration Test Flow

Perform an integration test when the CEP engine is running. Below is the test flow.

Figure 5.6 Integration test flow



5.7.2 Checking an Engine Log

The debug information of a CEP engine and the error messages generated at its start or during its operation will be output to an engine log. Rule errors checked when the CEP engine starts will also be output to the engine log.

One CEP engine outputs two engine logs. One is used for output relating to input adapter and high-speed filter processing, and the other is used for output relating to complex event processing and output adapter processing. The engine logs have no predetermined format.

The output destination of each engine log is shown below.

Engine log of the high-speed filter

```
/var/opt/FJsvcep/cep/flt/logs/EngineLog/CEPengineName/engine.log
```

Engine log of complex event processing

```
/var/opt/FJsvcep/cep/cep/logs/EngineLog/CEPengineName/engine.log
```

If "DebugLogListener" is used in complex event processing rules, processing results will be output to the engine log of complex event processing.



Note

Engine log splitting

If a "DebugLogListener" log of more than 102,400 characters is to be output at once to the engine log, the log will be split and then output. Refer to "Checking when a DebugLogListener log has been split" below for details.

Checking when a DebugLogListener log has been split

If the output results of hits for complex event processing statement conditions exceed 102,400 characters, the output results will be split every 102,400 characters and then output.

At split output, beginning and end identifiers will be added to the split output log. When this happens, only the end identifier will be output in the first output result of the split log, and only the beginning identifier will be output in the last output result.

Beginning identifier

The following content will be output:

```
YYYY-MM-DD hh:mm:ss,sss [DEBUG]
*****CUT_threadID_nanosecond*****
```

End identifier

The following content will be output. *threadID* and *nanosecond* will have the same values as those of the beginning identifier.

```
*****CUT_threadID_nanosecond*****
```



Note

Possibility of log output getting out of sequence

There is a possibility that the processing results of other events being executed simultaneously will interrupt the units into which the log was split. Reference the log according to "Example of output when a log has been split" below.



Example

Example of output when a log has been split

```
2011-12-07 14:01:13,720 [DEBUG] 6e1619bd-a048-4c64-9efb-306e9f2b88d6:length=100000
6e1619bd-a048-4c64-9efb-306e9f2b88d6[0]
    residence      :12: String
    value         :ON: String
    gatewayId     :00000001: String
(...)
6e1619bd-a048-4c64-9efb-306e9f2b88d6[123]
    residence:30: String
    val
*****CUT_94_19131119552293*****           ... 1.

~~~ log of other events ~~~

2011-12-07 14:01:13,720 [DEBUG]           ... 2.
*****CUT_94_19131119552293*****
ue      :ON: String
    gatewayId     :00000001: String
6e1619bd-a048-4c64-9efb-306e9f2b88d6[124]
    residence      :38: String
    value         :ON: String
    gatewayId     :00000001: String
```



```
(...)
6e1619bd-a048-4c64-9efb-306e9f2b88d6[247]
    residence      :30: String
    val
****CUT_94_19131119552293****          ... 3.

~~~log of other events ~~~

2011-12-07 14:01:13,720 [DEBUG]          ... 4.
****CUT_94_19131119552293****
ue      :ON: String
    gatewayId     :00000001: String
6e1619bd-a048-4c64-9efb-306e9f2b88d6[248]
    residence      :38: String
    value         :ON: String
    gatewayId     :00000001: String  ... 5.

~~~log of other events ~~~
```

Explanation of output example:

1. End of the first part of the split output log. Check the identifier (****CUT_94_19131119552293****).
2. Next beginning of the continuing part of the split output log. This is started using the same identifier (****CUT_94_19131119552293****).
3. Next end of the continuing part of the split output log. This is split using the same identifier.
4. Similarly, this is split up to the last part using the same identifier.
5. The end of the last part of the split output log has no identifier.



5.7.3 Starting

Start the deployed user-developed Web service and the CEP engine in which definition information has been deployed.

Also check syntax errors along with starting the CEP engine.

The procedure for the starting tasks is as follows:

1. [Checking the Status of a User-developed Web Service](#)
2. [Starting the CEP Engine](#)
3. [Checking for Syntax Errors in Filter Rules](#)
4. [Checking for Syntax Errors in Complex Event Processing Rules](#)

5.7.3.1 Checking the Status of a User-developed Web Service

If the application server in which a user-developed Web service has been deployed is not running, start the server. After starting the server, check that the deployed user-developed Web service has the status of receiving requests from outside.

Refer to the application server manuals for information on how to start the application server and how to check the status of the Web service.

5.7.3.2 Starting the CEP Engine

After deployment, use cepstarteng to start the CEP engine. Start the CEP engine in which definition information has been deployed. Refer to "8.9 cepstarteng" for details.



Example

Example of executing cepstarteng

```
$ cepstarteng -e CepEngine<ENTER>
Command cepstarteng executed successfully.
```

5.7.3.3 Checking for Syntax Errors in Filter Rules

If the filter rules in the rule definition contain syntax errors, the CEP engine start will fail. If this happens, the cause of the error will be notified to the engine log of the high-speed filter. Correct the syntax error based on the notified error content.

Redeploy the corrected rule definition in the CEP engine. Refer to "[6.1.4.3 Updating Deployed Definition Information](#)" for information on redeploying. After redeploying, repeat the operations from "[5.7.3.2 Starting the CEP Engine](#)".

5.7.3.4 Checking for Syntax Errors in Complex Event Processing Rules

If the complex event processing statements in the rule definition contain syntax errors, the CEP engine start will fail with a "cep20201e" error.

If this happens, the cause of the error will be notified to the engine log of complex event processing using the error message in the "cep20201e" ERRORINFO parameter.

Below is an example of an error message in the ERRORINFO parameter. Correct the syntax error based on this content.

```
Incorrect syntax near xxxxxxxx at line X column Y
```

xxxxxxx: Keyword near the abnormality

X: Number of lines from the beginning of the complex event processing statement in which the error occurred

Y: Number of characters from the beginning of the line of the complex event processing statement in which the error occurred

Redeploy the corrected rule definition in the CEP engine. Refer to "[6.1.4.3 Updating Deployed Definition Information](#)" for information on redeploying. After redeploying, repeat the operations from "[5.7.3.2 Starting the CEP Engine](#)".



Note

Depending on the content of the error message, the location of the error (line and column information) may not be displayed.



Example

Log output example of a syntax error in a complex event processing statement

Target complex event processing statement

```
@Name('EPL3')
@DebugLogListener
select * from FilteredCouponEvent (storeID='STR0001') wherer cast(targetAge,int)>20;
```

Log output (please note that we added newlines in the example below (lines 1 to 6) for readability only - the actual output does not have a newline.)

```
2012-07-15 13:51:18,843 [ERROR] CSPF_CEP: ERROR: cep20201e: EPL module file access
failure. EngineId=CepEngine, FILE=/etc/opt/FJSVcep/resources/CepEngine/rules/SampleRule.
epl, ERRORINFO=com.espertech.esper.client.deploy.DeploymentActionException: Compilation
failed in module url '/etc/opt/FJSVcep/resources/CepEngine/rules/SampleRule.epl' in
expression '@Name('EPL3')@DebugLogListenerselect * from Filt...(115 chars)' : Incorrect
syntax near 'cast' (a reserved keyword) at line 3 column 61, please check the from
clause [@Name('EPL3')]
@DebugLogListener
select * from FilteredCouponEvent (storeID='STR0001') wherer cast(targetAge,int)>20]
```

Explanation of log output:

- ERRORINFO=...

Error information will be output following this ERRORINFO parameter.

- Compilation failed...

This is a broad error classification. "Compilation failed" or "Deployment failed" will be output.

- SampleRule

This is the development asset ID of the rule definition in which the error occurred.

- Incorrect syntax...

This is the error message. It indicates that there is a syntax error near "cast" in line 3, column 61 from the beginning of complex event processing statement in which the error occurred.

- @Name('EPL3')...

This is the complex event processing statement in which the error occurred. In this example, the "where" before "cast" should actually be "where", and the extra "r" has caused a syntax error.

5.7.4 Integration Test

The flow of the integration test tasks is as follows:

1. [Sending Event Data for Testing](#)
2. [Checking the Operation of Filter Rules](#)
3. [Checking the Operation of Complex Event Processing Rules](#)
4. [Checking the Operation of a User-developed Web Service](#)
5. [Checking the Event Log](#)
6. [Checking the Operation of an Event Log Analysis Application](#)

5.7.4.1 Sending Event Data for Testing

Send event data for testing to the CEP engine.

Execute the deployed event sender application or the event sender sample program supplied with the samples.

After sending the events, check the engine log to see if event sending is operating as expected.

Refer to "[5.9.6 Event Sender Sample Program](#)" for information on how to use an event sender sample program.



Example

Sending event data for testing

This is an example of sending the "/application/test.csv" data for testing, which is in CSV format, to a CEP engine (Cepengine) as event data with the event type "EVENT01".

```
$ /opt/FJSVcep/sample/sample1/bin/sendevent.sh EVENT01 /application/test.csv<ENTER>
```

5.7.4.2 Checking the Operation of Filter Rules

Check whether any errors have been notified to the engine log of the high-speed filter, and correct the error based on the notified error content.

In addition to errors, also check whether any content on input events (those with logging enabled in the event type definition) was logged in the engine log, if logging with the "file" logging type is being used.

5.7.4.3 Checking the Operation of Complex Event Processing Rules

Check whether any errors have been output to the engine log of complex event processing, and correct the error based on the notified error content.

In addition to errors, also check the operation of the complex event processing rules by checking the engine log for any content output using the "@DebugLogListener" annotation described in a complex event processing statement.

If the conditions in the complex event processing rule with the "@DebugLogListener" annotation have a hit, the property value of the output event will be output to the log.

Also check whether any content output using the "@LoggingListener" annotation was logged in the engine log, if logging with the "file" logging type is being used.

Note

The following considerations relate to the operation results of complex event processing rules:

- Once checking the operation results of the rules is complete, before deploying the rule definition in the production environment, edit it to delete the "@DebugLogListener" annotation from the rule definition.
- To correct and replace a rule definition during testing, first stop the CEP engine, then redeploy the rule definition, and then restart the CEP engine.
- When the CEP engine is stopped, any data that was in the windows generated by rules before it was stopped will disappear.

Example

Log output example of "DebugLogListener"

Below is an example of engine log output using the "@DebugLogListener" annotation, and will be the execution result of the following complex event processing statement:

Complex event processing statement to be used

```
@Name('EPL3')
@DebugLogListener
@LoggingListener(table='/EPL3', properties='areaID, couponID')
@SoapListener('LISTEN01')
select areaID, targetAge, storeID, couponID
  from FilteredCouponEvent (storeID='STR0001') where cast(targetAge,int)>20;
```

Log output example of "DebugLogListener"

```
2012-07-15 14:21:11,854 [DEBUG] EPL3:length=1      ... 1.
EPL3[0]      ... 2.
      areaID  :2222: String      ... 3.
      targetAge    :30: String
      storeID  :STR0001: String
      couponID   :CPN0001: String
```

Explanation of output example:

1. This displays the number of records that are a hit from the SELECT statement condition. The value specified using the "@Name" annotation appears in the underlined part.
2. This displays the record index. The value specified using the "@Name" annotation appears in the underlined part.
3. This line onwards displays the properties, property values, and property types of the output event.

Information

Example of an engine log output by other output adapters

An engine log output by other output adapters can also be used as debug information because it shows whether listener calling succeeded.

Example of engine log output of the SOAP listener

If the conditions in the complex event processing statement associated with the "@SoapListener" annotation have a hit, the successful listener calling will be output to the engine log. (Please note that we added newlines in the example below for readability only - the actual output does not have a newline.)

```
2012-07-15 14:44:42,556 [DEBUG] id=LISTEN01 destination=http://xxx.xxx.xx.xxx/WebServWAR/MyApp1Service methodName=root<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"><SOAP-ENV:Header/><SOAP-ENV:Body><ns:root xmlns:ns="http://webservice/"><ns:areaID>2222</ns:areaID><ns:targetAge>30</ns:targetAge><ns:storeId>STR0001</ns:storeId><ns:couponID>CPN0001</ns:couponID></ns:root></SOAP-ENV:Body></SOAP-ENV:Envelope>
```

Explanation of output example:

- LISTEN01

This is the development asset ID of the SOAP listener definition that was used.

- http://xxx.xxx.xx.xxx/WebServWAR/MyApp1Service

This is the connection URL. This will be the value specified in the "url" tag in the SOAP listener definition.

- root

This is the root element name. This will be the value specified in the "method" tag in the SOAP listener definition.

- From "root" onwards

This is the SOAP message that was sent.

Example of engine log output of the logging listener

If the conditions in the complex event processing statement associated with the "@LoggingListener" annotation have a hit, the successful logging will be output to the engine log.

```
2012-07-26 00:05:18,692 [DEBUG] Write Log message"2222", "CPN0001"
```

If the engine log was used as the logging destination (by specifying "file" in "type" in the engine configuration file), the processing results of the complex event processing rules will also be output.

```
2012-07-26 00:05:18,692 [DEBUG] TableName:/EPL3; eventdata:"2222", "CPN0001"
2012-07-26 00:05:18,692 [DEBUG] Write Log message"2222", "CPN0001"
```

5.7.4.4 Checking the Operation of a User-developed Web Service

Perform this only if a user-developed Web service has been designed and developed.

Check the operation of a user-developed Web service based on information such as the logs it outputs.

5.7.4.5 Checking the Event Log

If a setting to perform logging using Hadoop collaboration has been specified, check that an event log has been logged in the Hadoop system.

5.7.4.6 Checking the Operation of an Event Log Analysis Application

Perform this only if an event log analysis application has been designed and developed.

Use a recorded event log to check the operation of an event log analysis application.

5.7.5 Stopping

Stop the deployed event sender application and the CEP engine in which definition information has been deployed.

5.7.5.1 Stopping an Event Sender Application

If an event sender application is running and if event sending is continuing, stop the event sender application.

5.7.5.2 Stopping the CEP Engine

Stop the CEP engine. Use `cepstopeng` to stop the CEP engine.

Refer to "8.11 `cepstopeng`" for details.



Example

Example of command execution

This is an example of command execution when stopping a CEP engine (CepEngine).

```
$ cepstopeng -e CepEngine<ENTER>
Command cepstopeng executed successfully.
```

5.7.6 Correcting Development Assets

If the checked operation results have a problem, correct the development asset.

After correcting it, first undeploy the development asset, then redeploy it, and then perform an integration test again.



Point

Updating definition information by overwriting

If only part of the definition information is to be corrected, using the `-o` option of `cepdeployrsc` allows the definition information to be updated by overwriting (and redeployed).

5.8 Undeploying Development Assets

Undeploy any development assets that are no longer required.

5.8.1 Undeploying Definition Information

Execute `cepundeployrsc` to undeploy definition information that is no longer required.

Refer to "8.13 `cepundeployrsc`" for details.

Below is an example of undeploying.



Note

Note on executing `cepundeployrsc`

Definition information cannot be undeployed while the CEP engine is running.

Example

Example of undeploying definition information

This is an example of command execution when the following definition information that has been deployed in a CEP engine (CepEngine) is to be undeployed:

Type of definition information	Development asset ID
Event type definition	EVENT01
Rule definition	RULE01
Master definition	MASTER01
SOAP listener definition	LISTENER01

The command execution example is as follows:

```
$ cepundeployrsc eventtype -e CepEngine -n EVENT01<ENTER>
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
$ cepundeployrsc rule -e CepEngine -n RULE01<ENTER>
(...)
$ cepundeployrsc master -e CepEngine -n MASTER01<ENTER>
(...)
$ cepundeployrsc listener -e CepEngine -n LISTENER01<ENTER>
(...)
```

5.8.2 Undeploying a Collaboration Application

Undeploy any collaboration application that is no longer required.

Event sender application

Undeploy an event sender application according to the application undeployment method of the event sender system.

If an event sender sample program is being used as the event sender application, no undeployment is required.

User-developed Web service

Undeploy a user-developed Web service from the application server. Perform the undeployment according to the method in the application server manual.

Event log analysis application

Undeploy an event log analysis application from the Hadoop system. Perform the undeployment according to the method in the Interstage Big Data Parallel Processing Server manual.

Cache access application for update

Undeploy a cache access application for update from the application server (Java EE). Perform the undeployment according to the method in the Interstage Application Server manual.

5.8.3 Deleting Data

Delete any data that is no longer required.

Event data (for testing)

Delete event data for testing from the event sender system.

If an event sender sample program is being used as the event sender application, delete the event data for testing that was stored on the CEP Server.

Logging storage destination

Delete the directory that was used as the logging storage destination.

If the data in the directory is required, store the data in another location and then delete the directory.

Master data

Delete the master data that was stored in the path specified in the master definition.

XTP cache

Refer to the XTP manuals for information on how to delete an XTP cache.

5.9 Sample Application

This section explains the sample application supplied with this product.

5.9.1 Overview of the Sample Application

This sample is an application for detecting members walking in the vicinity of a registered store and issuing them with coupons that have a time limit.

The application for this service consists of the definition information, which is the part to be processed by the CEP engine, and the master data, which is to be used by the rules. Sample event data and an event sender sample program are also supplied for checking the operation of the application.

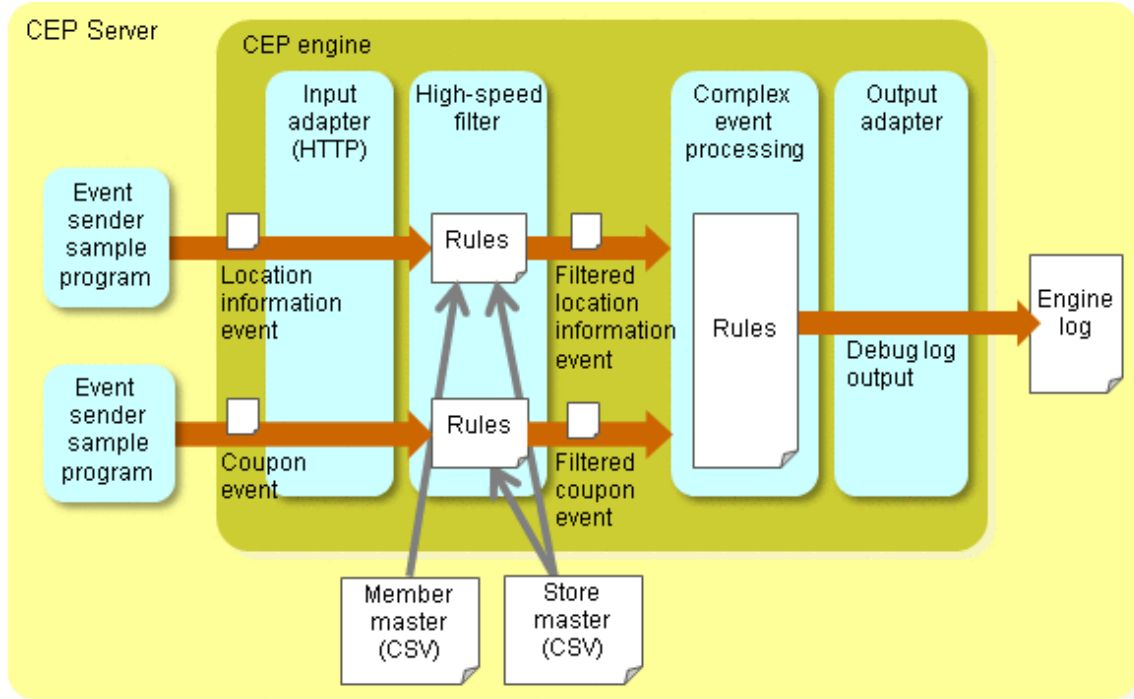
5.9.2 Structure of the Sample

The table below shows the structure of the development assets (and what is provided) in this sample.

Development asset type	Development asset	What is provided
Definition file	Event type definition	Location information event definition Coupon event definition
	Rule definition	Sample rule definition
	Event type definition (filtered events)	Filtered location information event definition Filtered coupon event definition
	Master definition	Member master definition Store master definition
Data	Master data (CSV file)	Member master data Store master data
	Event data	Sample data for location information events Sample data for coupon events
Collaboration application	Event sender application	Event sender sample program (Binary and source code)

Below is an example of how the sample operates.

Figure 5.7 Example of how the sample operates



5.9.3 Events

5.9.3.1 Location Information Events

The table below shows the content of the event type definition for location information events.

Identifier (development asset ID)	Format	Column name
LocationEvent	CSV	memberID
		areaID
		status

5.9.3.2 Coupon Events

The table below shows the content of the event type definition for coupon events.

Identifier (development asset ID)	Format	Column name
CouponEvent	CSV	storeID
		couponID
		targetAge

5.9.3.3 Filtered Location Information Events

The table below shows the content of the event type definition for filtered location information events.

Identifier (development asset ID)	Format	Column name
FilteredLocationEvent	CSV	memberID

Identifier (development asset ID)	Format	Column name
		areaID
		age

5.9.3.4 Filtered Coupon Events

The table below shows the content of the event type definition for filtered coupon events.

Identifier (development asset ID)	Format	Column name
FilteredCouponEvent	CSV	storeID
		couponID
		targetAge
		areaID

5.9.4 Master Information

This section explains the master data to be used by this sample.

5.9.4.1 Member Information Master

The table below shows the content of the member information master definition.

Identifier (development asset ID)	Column name
MemberInfo	ID
	age

5.9.4.2 Store Information Master

The table below shows the content of the store information master definition.

Identifier (development asset ID)	Column name
StoreInfo	ID
	areaID

5.9.5 Rule Definition

The rule definition processes two types of events. It consists of filter rules and complex event processing rules.

5.9.5.1 Filter Rules (IF-THEN Format)

These filter rules are for two types of events.

1	on LocationEvent {
2	if (\$status == "1" AND lookup("StoreInfo", \$areaID == \$areaID) = true()) then
3	join("MemberInfo", \$memberID==\$ID)
4	output(\$memberID, \$areaID, "MemberInfo".\$age) as FilteredLocationEvent;
5	}
6	
7	on CouponEvent {

```

8      join("StoreInfo", $storeID == $ID)
9      output($storeID, $couponID, $targetAge, "StoreInfo".$areaID) as FilteredCouponEvent;
10     }

```

Below is an explanation of each line.

- Line 1

This is a filter rule for a location information event (LocationEvent).

- Line 2

This search expression extracts only the events where a registered store can be found in the area (one with an "areaID" equal to the "areaID" in the "storeInfo" store master) while walking ("status" is "1"). The left side of "\$areaID == \$areaID" is the "areaID" of the input events, and the right side is the "areaID" of the master data.

- Line 3

To assign member master information to the input events, this uses the data of the member master "MemberInfo" to join the events that have equal member IDs.

- Line 4

This outputs the member ID and area ID that are in the input events, as well as the age information of the member master, as a "FilteredLocationEvent".

- Line 7

This is a filter rule for a coupon event (CouponEvent).

- Line 8

This uses the store master "StoreInfo" data to join the events that have equal store IDs.

- Line 9

This outputs the store ID, coupon ID, and target age that are in the input events, as well as the area ID of the store master, as a "FilteredCouponEvent".

5.9.5.2 Complex Event Processing Rules (SQL Format)

These complex event processing rules are for the two types of events processed by the filter rules.

```

1      create schema outEvent (memberID string, storeID string, couponID string);
2
3      create window outEventWin.std:firstunique(memberID,couponID).win:time(3 min) as outEvent;
4
5      @Name('EPL1')
6      insert into outEventWin
7          select loc.memberID as memberID,cpn.storeID as storeID,cpn.couponID as couponID
8             from FilteredLocationEvent as loc unidirectional
9                inner join FilteredCouponEvent.win:time(3 min) as cpn
10                   on loc.areaID=cpn.areaID and loc.age=cpn.targetAge;
11
12     @Name('EPL2')
13     @DebugLogListener
14     select * from outEventWin;
15
16     @Name('filterOut1')
17     @DebugLogListener
18     select * from FilteredLocationEvent;
19
20     @Name('filterOut2')
21     @DebugLogListener
22     select * from FilteredCouponEvent;

```

Below is an explanation of each line.

- Line 1

This defines the format of the events being output (used in the next named window definition).

- Line 3

This is a complex event processing statement that creates a named window to retain the output events for a set time period and to ensure that the same coupon is not sent to the same person twice within that period. This retains the first output events for member ID and coupon ID pairs, one at a time. After the set period (3 minutes in the sample) has elapsed, they are deleted from the window.

- Line 5

This defines the complex event processing statement to perform the main process. Its name is "EPL1".

- Line 6

This inserts the output of the main process in the named window.

- Line 7

This outputs the member IDs of location information events and the store IDs and coupon IDs of coupon events, as the results of the main process.

- Line 8

This is for the "FilteredLocationEvent", which is one source of input for the main process. When "unidirectional" is specified, the reception of this event will be the trigger for operating this process ("FilteredCouponEvent" will not be the trigger).

- Line 9

This is for the "FilteredCouponEvent", which is the other source of input for the main process. This event is retained in memory for a set period only (3 minutes in the sample) and is joined to "FilteredLocationEvent".

- Line 10

The join condition for "FilteredLocationEvent" and "FilteredCouponEvent" is that the "areaIDs" are equal and that "age" and "targetAge" are equal.

- Line 12

This is a complex event processing statement for outputting the processing results of "EPL1" (line 5 onwards). Its name is "EPL2".

- Line 13

By assigning a debug log listener, this outputs the output events to the engine log.

- Line 14

This outputs only the new events registered in the named window. Even if there are more than one of the same event entered in the named window, specifying "std:firstunique" (in "create window") will ensure that only the first event will actually be stored.

- Lines 16 to 18

This is a complex event processing statement for performing debug log output of a filtered location information event (FilteredLocationEvent). This is for checking the filtering process and can safely be deleted.

- Lines 20 to 22

This is a complex event processing statement for performing debug log output of a filtered coupon event (FilteredCouponEvent). This is for checking the filtering process and can safely be deleted.

5.9.6 Event Sender Sample Program

The event sender sample program is an event sender application that sends event data for testing in CSV format to an HTTP adapter.

Execution format

The execution format of the program is as follows:

```
/opt/FJSVcep/sample/sample1/bin/sendevent.sh eventType dataFilePath [ sendDestinationURL ]
```

Arguments

eventType

Specify which event type is to be used to send the data.

Specify the development asset ID of the target event type.

dataFilePath

Specify the path of the CSV format file that contains the event data.

sendDestinationURL

Specify the endpoint address of the HTTP adapter.

If this is omitted, the default is "http://localhost/CepEngineFrontServerService/HttpReceiver".



Information

Endpoint address of the CEP engine

The send destination URL of the CEP engine is as follows:

```
http://CEPengineAddress/CEPengineNameFrontServerService/HttpReceiver
```

The send destination URL of the CEP engine (CepEngine) of the CEP Server ("cephost.example.com") is as follows:

```
http://cephost.example.com/CepEngineFrontServerService/HttpReceiver
```

Source code

Below is the source code of the event sender sample program.

```
/opt/FJSVcep/sample/sample1/src/sample/HttpClient.java
```

```
1 package sample;
2
3 import java.io.BufferedOutputStream;
4 import java.io.BufferedReader;
5 import java.io.FileInputStream;
6 import java.io.InputStreamReader;
7 import java.net.HttpURLConnection;
8 import java.net.URL;
9
10 public class HttpClient {
11
12     static String DEFAULT_URL = "http://localhost/CepEngineFrontServerService/HttpReceiver";
13
14     public static void main(String[] args) {
15         if (args.length != 2 && args.length != 3) {
16             System.err.println("Requires two or three arguments: event-type, file name, [URL]");
17             System.exit(1);
18         }
19         String eventType = args[0];
20         String fileName = args[1];
21         String url = DEFAULT_URL;
22         if (args.length == 3) {
23             url = args[2];
24         }
25
26         try {
27             BufferedReader br0 =
```

```

28         new BufferedReader(new InputStreamReader(new FileInputStream(fileName)));
29
30     String str;
31     while ((str = br0.readLine()) != null) {
32         URL TestURL = new URL(url);
33         HttpURLConnection con = (HttpURLConnection) TestURL.openConnection();
34
35         con.setRequestMethod("POST");
36         con.setRequestProperty("TYPE", "CSV");
37         con.setRequestProperty("EVENT-TYPE-ID", eventType);
38         con.setRequestProperty("Content-Type", "text/plain; charset=utf-8");
39         con.setDoOutput(true);
40
41         BufferedOutputStream bos = new BufferedOutputStream(con.getOutputStream());
42         bos.write(str.getBytes("utf-8"));
43         bos.flush();
44         bos.close();
45
46         BufferedReader br1 = new BufferedReader(new InputStreamReader(con.getInputStream()));
47
48         String line;
49         while ((line = br1.readLine()) != null) {
50             System.out.println(line);
51         }
52         br1.close();
53         con.disconnect();
54     }
55     br0.close();
56 } catch (Exception e) {
57     e.printStackTrace();
58 }
59 }
60 }

```

5.9.7 Directory Structure

This sample is stored in `"/opt/FJSVcep/sample"` using the following structure:

```

sample1
+-- client.jar           Jar file to be used by the event sender sample program
+-- bin                 Directory for shell scripts
|   +-- deployall.sh    Shell script to deploy all sample development assets
|   +-- undeployall.sh  Shell script to undeploy all sample development assets
|   +-- sendevent.sh    Event sender sample program
+-- src                 Directory for source code
|   +-- sample
|       +-- HttpClient.java  Source code of jar to be used by the event sender sample program
+-- master              Directory for master data
|   +-- MemberData.csv    Sample data of member master
|   +-- MemberDataSchema.csv  Schema file of member master
|   +-- StoreData.csv     Sample data of store master
|   +-- StoreDataSchema.csv  Schema file of store master
+-- event               Directory for event data
|   +-- CouponEvent.csv  Sample data of coupon event
|   +-- LocationEvent.csv  Sample data of location information event
+-- resources           Directory for definition files
|   +-- CouponEvent.xml   Event type definition of coupon event
|   +-- FilteredCouponEvent.xml  Event type definition of filtered coupon event
|   +-- LocationEvent.xml  Event type definition of location information event
|   +-- FilteredLocationEvent.xml  Event type definition of filtered location information event
|   +-- MemberInfo.xml    Master definition of member master

```

+-+ StoreInfo.xml	Master definition of store master
+-+ SampleRule.xml	Rule definition

5.9.8 Execution

This section explains how to execute the sample application.

Use general user permissions to execute the commands in the tasks below.

Perform the tasks in the following sequence:

1. [Deploying Development Assets](#)
2. [Starting the CEP Engine](#)
3. [Sending Events and Checking the Results](#)
4. [Stopping the CEP Engine](#)
5. [Undeploying Development Assets](#)

5.9.8.1 Deploying Development Assets

Deploy the development assets when only one CEP engine has been deployed (for example, immediately after the initial setup of this product has completed).

Execute `"/opt/FJSVcep/sample/sample1/bin/deployall.sh"`.

During this process, a number of queries will be made, so press "Enter" for all of them.

```
$ cd /opt/FJSVcep/sample/sample1/bin <ENTER>
$ ./deployall.sh <ENTER>
Are you sure you want to deploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
Are you sure you want to deploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
Are you sure you want to deploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
Are you sure you want to deploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
Are you sure you want to deploy the rule definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
Are you sure you want to deploy the master data definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
Are you sure you want to deploy the master data definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
```

In this shell script, the following commands are executed:

- Deploying event type definitions

```
cepdeployrsc eventtype -f /opt/FJSVcep/sample/sample1/resources/LocationEvent.xml
cepdeployrsc eventtype -f /opt/FJSVcep/sample/sample1/resources/CouponEvent.xml
cepdeployrsc eventtype -f /opt/FJSVcep/sample/sample1/resources/FilteredLocationEvent.xml
cepdeployrsc eventtype -f /opt/FJSVcep/sample/sample1/resources/FilteredCouponEvent.xml
```

- Deploying master definitions

```
cepdeployrsc master -f /opt/FJSVcep/sample/sample1/resources/MemberInfo.xml
cepdeployrsc master -f /opt/FJSVcep/sample/sample1/resources/StoreInfo.xml
```

- Deploying the rule definition

```
cepdeployrsc rule -f /opt/FJSVcep/sample/sample1/resources/SampleRule.xml
```

5.9.8.2 Starting the CEP Engine

Execute the following command to start the CEP engine:

```
$ cepstarteng<ENTER>
```

5.9.8.3 Sending Events and Checking the Results

Use the event sender sample program (`/opt/FJSVcep/sample/sample1/bin/sendevent.sh`) to send the sample data and use the engine log to check the rule operation.

1. Sending location information events and checking the results

```
$ ./sendevent.sh LocationEvent /opt/FJSVcep/sample/sample1/event/LocationEvent.csv<ENTER>
```

Check that "filterOut1" events have been output to the engine log.

Output example

```
2012-07-09 11:11:57,409 [DEBUG] filterOut1:length=1
filterOut1[0]
    areaID :2222: String
    age    :30: String
    memberID      :MEM0008: String

2012-07-09 11:11:57,413 [DEBUG] filterOut1:length=1
filterOut1[0]
    areaID :2222: String
    age    :20: String
    memberID      :MEM0003: String
```

2. Sending coupon events and checking the results

```
$ ./sendevent.sh CouponEvent /opt/FJSVcep/sample/sample1/event/CouponEvent.csv<ENTER>
```

Check that "filterOut2" events have been output to the engine log.

Output example

```
2012-07-09 11:15:38,054 [DEBUG] filterOut2:length=1
filterOut2[0]
    areaID :2222: String
    targetAge :30: String
    storeID :STR0001: String
    couponID :CPN0001: String
```

3. Sending location information events and checking the results

Send location information events within 3 minutes of sending coupon events.

```
$ ./sendevent.sh LocationEvent /opt/FJSVcep/sample/sample1/event/LocationEvent.csv<ENTER>
```

Check that events indicating a coupon issue ("EPL2" event) have been output to the engine log.

Output example

```
2012-07-09 11:17:25,314 [DEBUG] filterOut1:length=1
filterOut1[0]
    areaID :2222: String
    age    :20: String
    memberID      :MEM0003: String

2012-07-09 11:17:25,321 [DEBUG] filterOut1:length=1
filterOut1[0]
```



```

        areaID :2222: String
        age    :30: String
        memberID      :MEM0008: String

2012-07-09 11:17:25,323 [DEBUG] EPL2:length=1
EPL2[0]
        storeID :STR0001: String
        memberID      :MEM0008: String
        couponID    :CPN0001: String

```

4. Sending location information events and checking the results

Send location information events again within 3 minutes of sending coupon events.

```
$ ./sendevent.sh LocationEvent /opt/FJSVcep/sample/sample1/event/LocationEvent.csv<ENTER>
```

Check that events indicating coupon issue ("EPL2" event) have not been output to the engine log.

Output example

```

2012-07-09 11:18:00,691 [DEBUG] filterOut1:length=1
filterOut1[0]
        areaID :2222: String
        age    :20: String
        memberID      :MEM0003: String

2012-07-09 11:18:00,691 [DEBUG] filterOut1:length=1
filterOut1[0]
        areaID :2222: String
        age    :30: String
        memberID      :MEM0008: String

```

5.9.8.4 Stopping the CEP Engine

Execute the following command to stop the CEP engine:

```
$ cepstopeng<ENTER>
```

5.9.8.5 Undeploying Development Assets

Execute `"/opt/FJSVcep/sample/sample1/bin/undeployall.sh"` to undeploy the development assets in batch.

During this process, a number of queries will be made, so press "Enter" for all of them.

```

$ ./undeployall.sh<ENTER>
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
Are you sure you want to undeploy the rule definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
Are you sure you want to undeploy the master data definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.
Are you sure you want to undeploy the master data definition?(default: y) [y,n,q]:<ENTER>
Command cepundeployrsc executed successfully.

```

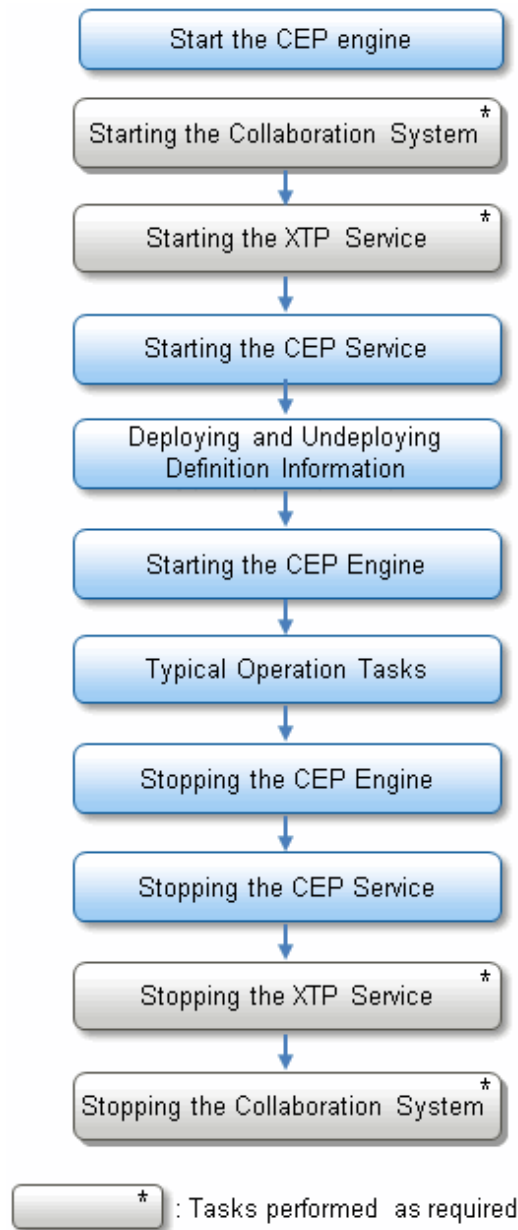
Chapter 6 Operation and Maintenance

This chapter explains how to operate the CEP Server.

If this product is to be operated using a reliable configuration, refer to "7.4 Operating a Cluster Service".

6.1 Operating the CEP Server

This section explains how to operate the CEP Server.



6.1.1 Starting the Collaboration System

If XTP collaboration or Hadoop collaboration is to be performed, first start the relevant collaborating system.

Refer to the manual for the relevant product for information on how to start the system.

6.1.2 Starting the XTP Service

If XTP collaboration is being performed, the CEP Server operates as an XTP Client Node (the CEP engine operates as a cache access application). Therefore, before starting the CEP service, start the XTP service on the CEP Server.

Refer to the *Interstage eXtreme Transaction Processing Server V1.0.0 Setup and Operation Guide* for details.



Do not confuse this XTP service with the XTP service on the XTP Server Node.

6.1.3 Starting the CEP Service

Execute `cepstartserv` as a superuser to start the CEP service.



Automatic start of the CEP service and CEP engine

- The CEP service starts automatically when the operating system starts. The CEP engine that was running the last time the operating system shut down also starts automatically.
- If `cepstopeng` or `cepstopserv` has been used to stop the CEP engine and the operating system is then restarted, the CEP engine will not start automatically when the operating system starts. To start the CEP engine, execute `cepstarteng` again.



Starting the service

When `cepstartserv` is used to start the CEP service.

```
# cepstartserv<ENTER>
(...)
Interstage Java EE DAS          started
(...)
Interstage Java EE Node Agent   started
(...)
UX:ismngconsolestart: INFO: is40041: The service has been activated normally.
(...)
UX:IHS: INFO: ihs01000: The command terminated normally.
(...)
LOG: database system is ready to accept connections
(...)
Starting Tomcat:                 [ OK ]

Command cepstartserv executed successfully.
```

6.1.4 Deploying and Undeploying Definition Information

This section explains how to deploy the definition information stored on the CEP Server, and how to undeploy definition information no longer required.

6.1.4.1 Deploying Definition Information

Execute `cepdeployrsc` to deploy the definition information on the CEP Server.

Refer to "[5.6.1 Deploying Definition Information](#)" for examples of deploying definition information.

A general user permission can be used to execute this command.

If the CEP engine is not stopped, then you must stop it (refer to "6.1.7 Stopping the CEP Engine" for details).

6.1.4.2 Checking Deployed Definition Information

Execute `cepdispeng` and `cepgetrsc` to check deployed definition information

Execute `cepdispeng` with the `-i` option to list the definition information deployed in the CEP engine, and `cepgetrsc` to check details of the deployed definition information.

A general user permission can be used to execute these commands. The definition information can be checked regardless of whether the CEP engine is stopped or running.



Example

Listing the deployed definition information

When listing the definition information deployed to a CEP engine (CepEngine).

```
$ cepdispeng -i -e CepEngine<ENTER>
engineId      :CepEngine
eventtype     :EVENT01
rule          :RULE01
master        :MASTER01
listener      :LISTENER01
Command cepdispeng executed successfully.
```

Listing details of the deployed definition information

When listing the details of an event type definition (EVENT01) deployed to a CEP engine (CepEngine).

```
$ cepgetrsc eventtype -e CepEngine -n EVENT01<ENTER>
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<eventType xmlns="urn:xmlns-fujitsu-com:csf:bdcep:v1" id="EVENT01">
  <comment>Event type definition</comment>
  <type>CSV</type>
  <xmlSchema></xmlSchema>
  <csvColumn>
    <column name="memberID" type="string"/>
    <column name="areaID" type="string"/>
    <column name="status" type="string"/>
  </csvColumn>
  <root></root>
  <useLogging>>false</useLogging>
  <loggingTableName></loggingTableName>
  <useCep>>true</useCep>
</eventType>
Command cepgetrsc executed successfully.
```



Note

- If redirecting the command output of the detailed display (XML format), modify each item accordingly.
- In addition to XML format data, the results of `cepgetrsc` also include output messages. If redirecting the command output to a file, delete unnecessary messages before using the results.

6.1.4.3 Updating Deployed Definition Information

Execute `cepdeployrsc` with the `-o` option to update the deployed definition information.

A general user permission can be used to execute this command.

If the CEP engine is not stopped, then you must stop it (refer to "6.1.7 Stopping the CEP Engine" for details).

Example

Updating the deployed definition information

When updating an event type definition (EVENT01) of a CEP engine (CepEngine) with new details.

The path of the event type definition with the new details (development asset ID: EVENT01) is shown below.

```
/application/EVENT01_new.xml
```

An example of command execution is shown below.

```
$ cepdeployrsc eventtype -o -e CepEngine -f /application/EVENT01_new.xml<ENTER>
Are you sure you want to deploy the event type definition?(default: y) [y,n,q]:<ENTER>
Command cepdeployrsc executed successfully.
```

6.1.4.4 Undeploying Definition Information

Execute `cepundeployrsc` to undeploy definition information no longer required.

Refer to "5.8.1 Undeploying Definition Information" for examples of undeploying definition information.

A general user permission can be used to execute this command. Release the definition information while the CEP engine is stopped.

Refer to "6.1.7 Stopping the CEP Engine" if the CEP engine needs to be stopped.

6.1.5 Starting the CEP Engine

Execute `cepstarteng` to start all CEP engines to be used - starting will be performed one CEP engine at a time.

A general user permission can be used to execute this command.

Example

When the command is executed.

```
$ cepstarteng -e CepEngine<ENTER>
Command cepstarteng executed successfully.
```

Note

Starting the CEP engine will fail in the following cases:

- If there has been no deployment of at least one event type definition.
- If an XML schema specified in an event type definition contains an error.
- If the syntax of a filter rule or complex event processing rule specified in a rule definition contains an error.

Information

How to check the CEP engine name

Execute `cepdispeng` to list the CEP engines that have been set up.

6.1.6 Typical Operation Tasks

This section explains the following operation tasks performed after the CEP engine is started:

- [Displaying the operation status of the CEP service](#)
- [Displaying the operation status of the CEP engine](#)
- [Monitoring abnormalities using logs](#)
- [Checking the resource usage of the CEP engine](#)

6.1.6.1 Displaying the Operation Status of the CEP Service

Execute `cepdisperv` as a superuser to display the status of the CEP service.

This allows you to check that the CEP service is running normally at any time, such as immediately after operation starts or during peak business hours. Execute `cepstartserv` as a superuser to start any service that is not running.



Example

Using `cepdisperv` to check the operation status of the CEP service

Check that the content below is output (*nnnn* indicates the process number, *CepEngine* indicates the created CEP engine name):

```
# cepdisperv<ENTER>
(...)
Interstage Java EE DAS          started
(...)
Interstage Java EE Node Agent   started
(...)
CEPAgentIJSERVERCluster running
CepEngine_flt not running
CepEngine_cep not running
(...)
Status              : Running
(...)
jsvc (pid nnnn nnnn) is running...
(...)
pg_ctl: server is running (PID: nnnn)
(...)
Command cepdisperv executed successfully.
```



Information

The CEP service entity

A CEP service is made up of multiple processes. Refer to "[8.5 cepdisperv](#)" for details.

6.1.6.2 Displaying the Operation Status of the CEP Engine

Execute `cepdispeng` to display the status of the CEP engine.

This allows you to check that the CEP engine is running normally at any time, such as immediately after operation starts or during peak business hours (refer to Chapter 5, "Errors during Operation" in *Troubleshooting* if the CEP engine has ABNORMAL or STOP status when it should be running).

A general user permission can be used to execute this command.

Example

Checking the operation status of the CEP engine

When checking the operation status of a CEP engine (CepEngine) - information such as the number of input and output events is also displayed.

```
$ cepdispeng -e CepEngine<ENTER>
engineId      :CepEngine
(...)
status_filter  :RUN
status_cep    :RUN
inEvent_filter :100
inEvent_cep   :100
outEvent_filter :100
outEvent_cep  :100
(...)
Command cepdispeng executed successfully.
```

6.1.6.3 Monitoring Abnormalities Using Logs

When an abnormality occurs in the CEP service and the CEP engine, a message will be output to the system log and to the engine logs.

The system log and the engine logs are used differently. For example, in addition to the messages reporting abnormalities, the engine logs also contained more detailed information on CEP engine operation, such as records sent using the SOAP listener and debug information. Therefore, the usual abnormality monitoring monitors only the system log and, when an abnormality occurs, references the engine logs to analyze its cause.

When an abnormality occurs, respond on the basis of the messages output to each log. Refer to *Messages* for information on specific responses.

The details of log output are explained as follows:

- Output destinations of the system log and the engine logs
- Message format

Output destinations of the system log and the engine logs

The output destinations of the system log and the engine logs are shown below.

One CEP engine outputs two engine logs - one contains output relating to input adapter and high-speed filter processing, and the other contains output relating to complex event processing and output adapter processing. The engine logs have no predetermined format.

System log

```
/var/log/messages
```

Engine log for the high-speed filter

```
/var/opt/FJSVcep/cep/flt/logs/EngineLog/CEPengineName/engine.log
```

Engine log for complex event processing

```
/var/opt/FJSVcep/cep/cep/logs/EngineLog/CEPengineName/engine.log
```

The engine logs of the high-speed filter and complex event processing undergo rotation as follows:

- File size: 10 MB
- Rotation generations: 9 generations



Engine logs older than 9 generations are automatically deleted. If they are required, periodically store them in a separate location.

Message format

The format of the messages output by this product is shown below.

Refer to Section 1.1, "Message Format" in *Messages* for details.

```
CSPF_CEP: errorType: messageNumber: messageText
```

- Each element is delimited by a colon (:).



Displaying a message

An example message is shown below.

```
CSPF_CEP: ERROR: cep10108e: Event type is not found. EngineId=CepEngine, eventType=EVENT01
```

6.1.6.4 Checking the Resource Usage of the CEP Engine

The resource usage of the CEP engine (such as the heap memory usage of Java VM) is recorded in the resource logs. The logs are output in CSV format and therefore can be analyzed using a tool such as Excel, in order to detect resource excesses or deficiencies in the CEP engine.

The details of the resource logs are explained as follows:

- Output destinations of the resource logs
- Format of the resource log for the high-speed filter
- Format of the resource log for complex event processing
- Collection interval for the resource usage

Output destinations of the resource logs

Two resource logs are output for each CEP engine - one contains output relating to input adapter and high-speed filter processing, and the other contains output relating to complex event processing and output adapter processing.

The output destination of each resource log is shown below.

Resource log for the high-speed filter

```
/var/opt/FJSVcep/cep/flt/logs/ResourceLog/CEPengineName/resource.log
```

Resource log for complex event processing

```
/var/opt/FJSVcep/cep/cep/logs/ResourceLog/CEPengineName/resource.log
```

The resource logs of the high-speed filter and complex event processing undergo rotation as follows:

- Point of rotation: One day's worth (specific point of rotation is set using cron)
- Rotation generations: 13 generations



Resource logs older than 13 generations are automatically deleted. If they are required, periodically store them in a separate location.

Format of the resource log for the high-speed filter

This section explains the format of the resource log for the high-speed filter.

The format is shown below:

```
time, resourceID, jheapNewUsed, jheapNewFree, jheapNewTotal, jheapOldUsed, jheapOldFree,
jheapOldTotal, jheapNewPlusOldTotal, jheapPermUsed, jheapPermFree, jheapPermTotal, VSZ,
numOfInEvents, numOfOutEvents, numOfLogs, numOfRulesDeployed, numOfSocketConnections
```

Please note that we added newlines in the format above for readability only - the actual log does not have a newline.

The table below explains each item:

Item name	Content
<i>time</i>	Output date and time of the resource log. The format is as follows: <div style="border: 1px solid black; padding: 2px; margin: 5px 0;"><code>yyyy-MM-dd HH:mm:ss</code></div> <ul style="list-style-type: none"> - <i>yyyy</i>: Year - <i>MM</i>: Month - <i>dd</i>: Day - <i>HH</i>: Hour - <i>mm</i>: Minute - <i>ss</i>: Second
<i>resourceID</i>	Resource ID. The format is as follows: <div style="border: 1px solid black; padding: 2px; margin: 5px 0;"><code>CSPF_CEP_hostName_engineIDflt</code></div>
<i>jheapNewUsed</i>	Java VM heap memory usage (new generation area) for the high-speed filter (unit: bytes) (*1)
<i>jheapNewFree</i>	Java VM heap free memory (new generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapNewTotal</i>	Java VM heap memory (new generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapOldUsed</i>	Java VM heap memory usage (old generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapOldFree</i>	Java VM heap free memory (old generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapOldTotal</i>	Java VM heap memory (old generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapNewPlusOldTotal</i>	Java VM heap memory (new generation area + old generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapPermUsed</i>	Java VM heap memory usage (permanent generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapPermFree</i>	Java VM heap free memory (permanent generation area) for the high-speed filter. (unit: bytes) (*1)
<i>jheapPermTotal</i>	Java VM heap memory (permanent generation area) for the high-speed filter. (unit: bytes) (*1)
<i>VSZ</i>	Process memory usage for the high-speed filter. (unit: KB) (*1)
<i>numOfInEvents</i>	Number of events input from devices on which events are occurring. (*1)
<i>numOfOutEvents</i>	Number of events sent to the engine on the complex event processing side. (*1)

Item name	Content
<i>numOfLogs</i>	Number of logs in the input adapter. (*1)
<i>numOfRulesDeployed</i>	Number of rules deployed to the high-speed filter (number of high-speed filter statements). (*1)
<i>numOfSocketConnections</i>	Number of simultaneous socket connections. (*1)

*1: If the CEP engine status is not RUN (running normally), an empty string ("") will be output.



Example

Content of the resource log (high-speed filter)

```
2012-08-01 14:00:04,CSPF_CEP_cepsv_CepEngineflt,30130928,9125136,39256064,39224768,
475888192,515112960,554369024,68207736,17292168,85499904,3670240,0,0,0,2,0
```

Please note that we added a newline in the example above for readability only - the actual log does not have a newline.

Format of the resource log for complex event processing

This section explains the format of the resource log for complex event processing.

The format is shown below:

```
time,resourceID,jheapNewUsed,jheapNewFree,jheapNewTotal,jheapOldUsed,jheapOldFree,
jheapOldTotal,jheapNewPlusOldTotal,jheapPermUsed,jheapPermFree,jheapPermTotal,VSZ,
reserve1,reserve2,reserve3,reserve4,reserve5,reserve6,numOfInEvents,numOfOutEvents,
numOfLogs,numOfRulesDeployed,numOfListenersDeployed
```

Please note that we added newlines in the format above for readability only - the actual log does not have a newline.

The table below explains each item:

Item name	Content
<i>time</i>	Output date and time of the resource log. The format is as follows: <pre>yyyy-MM-dd HH:mm:ss</pre> <ul style="list-style-type: none"> - <i>yyyy</i>: Year - <i>MM</i>: Month - <i>dd</i>: Day - <i>HH</i>: Hour - <i>mm</i>: Minute - <i>ss</i>: Second
<i>resourceID</i>	Resource ID. The format is as follows: <pre>CSPF_CEP_hostName_engineID_cep</pre>
<i>jheapNewUsed</i>	Java VM heap memory usage (new generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapNewFree</i>	Java VM heap free memory (new generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapNewTotal</i>	Java VM heap memory (new generation area) for complex event processing. (unit: bytes) (*1)

Item name	Content
<i>jheapOldUsed</i>	Java VM heap memory usage (old generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapOldFree</i>	Java VM heap free memory (old generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapOldTotal</i>	Java VM heap memory (old generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapNewPlusOldTotal</i>	Java VM heap memory (new generation area + old generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapPermUsed</i>	Java VM heap memory usage (permanent generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapPermFree</i>	Java VM heap free memory (permanent generation area) for complex event processing. (unit: bytes) (*1)
<i>jheapPermTotal</i>	Java VM heap memory (permanent generation area) for complex event processing. (unit: bytes) (*1)
<i>VSZ</i>	Process memory usage for complex event processing. (Unit: KB) (*1)
<i>reserve1</i>	Reserved areas. Empty strings ("") will be output for all of these.
<i>reserve2</i>	
<i>reserve3</i>	
<i>reserve4</i>	
<i>reserve5</i>	
<i>reserve6</i>	
<i>numOfInEvents</i>	Number of events input from high-speed filter processing. (*1)
<i>numOfOutEvents</i>	Number of events sent to user-developed Web services. (*1)
<i>numOfLogs</i>	Number of logs in complex event processing. (*1)
<i>numOfRulesDeployed</i>	Number of rules deployed to complex event processing (number of complex event processing statements). (*1)
<i>numOfListenersDeployed</i>	Number of deployed SOAP listener definitions. (*1)

*1: If the CEP engine status is not RUN (running normally), an empty string ("") will be output.



Example

Content of the resource log (complex event processing)

```
2012-08-01 14:00:04,CSPF_CEP_cepsv_CepEngine_cep,14042608,25213456,39256064,41802024,
473310936,515112960,554369024,76981880,8481160,85463040,3385020,,,,,,0,0,0,6,0
```

Please note that we added a newline in the example above for readability only - the actual log does not have a newline.

Collection interval for resource usage

The resource usage is collected in 10-minute intervals and is output to the resource log.

This process is performed by using cron to periodically call an obtain process.

The cron setting is set in the following file. The engine execution user is the user name specified at installation:

```
/var/spool/cron/engineExecutionUser
```

6.1.7 Stopping the CEP Engine

Execute `cepstopeng` to stop the CEP engine.

A general user permission can be used to execute this command.

Example

When a CEP engine (CepEngine) is stopped.

```
$ cepstopeng -e CepEngine<ENTER>
Command cepstopeng executed successfully.
```

Point

When `cepstopserv` is used, all CEP engines running will also be stopped.

6.1.8 Stopping the CEP Service

Execute `cepstopserv` as a superuser to stop the CEP service. Note that the command also stops all CEP engines running.

Example

When the command is executed.

```
# cepstopserv<ENTER>
Stopping Engines.
(...)
Command cepstopeng executed successfully.
(...)
Shutting down Tomcat: [ OK ]
(...)
server stopped
(...)
UX:IHS: INFO: ihs01000: The command terminated normally.
(...)
UX:ismngconsolestop: INFO: is40042: The service has been terminated normally.
(...)
Interstage Java EE Node Agent stopped
(...)
Interstage Java EE DAS stopped
(...)
UX:isstop: INFO: is30160: INTERSTAGE has been terminated normally.

Command cepstopserv executed successfully.
```

Information

When the operating system shuts down, the CEP service stops automatically.

6.1.9 Stopping the XTP Service

If XTP collaboration is being performed, stop the XTP service on the CEP Server.

Refer to the *Interstage eXtreme Transaction Processing Server V1.0.0 Setup and Operation Guide* for details.

Note

Do not to confuse this XTP service with the XTP service on the XTP Server Node.

6.1.10 Stopping the Collaboration System

If XTP collaboration or Hadoop collaboration is being performed, stop the relevant collaborating system after stopping the CEP Server. Refer to the manual for the relevant product for information on how to stop the system.

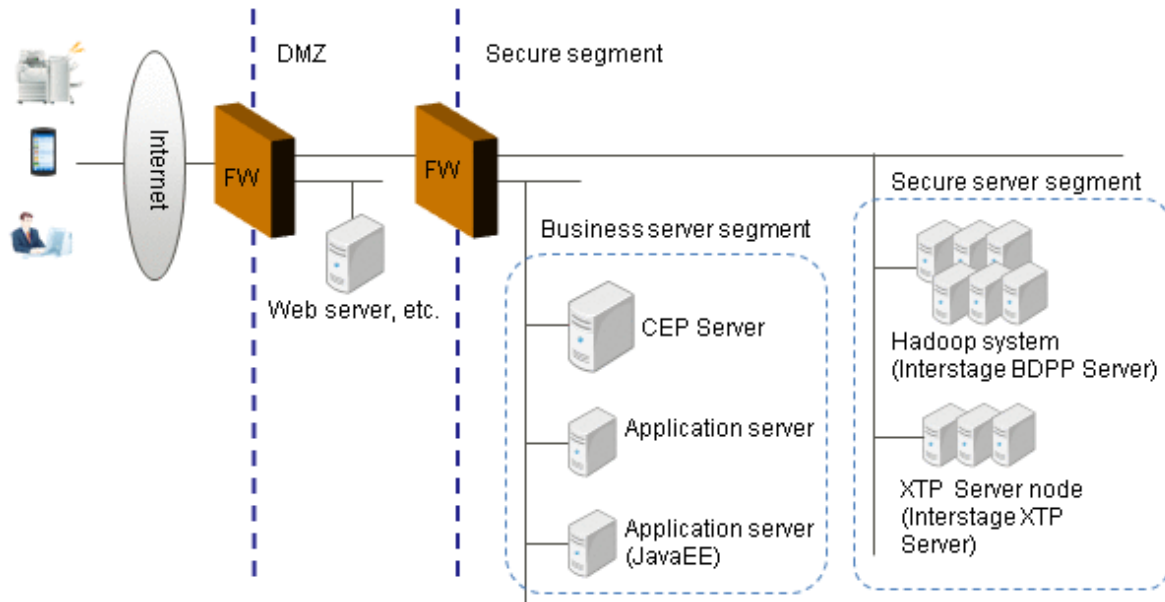
6.2 Security

This section provides information required to operate this product securely.

6.2.1 Operation Model

A typical operation model of this product is shown below:

Figure 6.1 Typical operation model of this product



6.2.2 Prerequisite Knowledge for Designing Security

Take the following elements into consideration when designing, in order to achieve secure system operation using this product:

- Security roles
- Protected resources
- Threats to protected resources and their countermeasures
- Overview of countermeasures for threats to protected resources

Security roles

The table below lists the types of system users that use this product, their security roles, and the corresponding operating system user:

User type	Security role	Operating system user
System administrator	Can perform all operations.	Superuser

User type	Security role	Operating system user
	Can perform operations involving CEP Server, such as starting and stopping the CEP Server or reconfiguring a CEP engine.	
Engine execution user	Can run a CEP engine process.	Engine execution user (create at installation)
Developer	Can deploy definition information to a CEP engine and undeploy it (to check the operation of definition information), as well as start and stop a CEP engine.	General user who can login to the CEP Server

Protected resources

The table below lists the resources to be protected by the CEP Server.

Type	Protected resource	Description
File	Engine configuration file	File used to reconfigure a CEP engine.
	Deployed definition information	Definition information deployed to a CEP engine.
	Master data	CSV files to be referenced by the high-speed filter.
	Resource log	Output file used to investigate the resource usage.
	Engine log	File to which the detailed operation status of a CEP engine is output.
	Data for investigation	File collected to investigate faults.
Network	Event data sent to an input adapter	Packets traveling across the network.
	SOAP messages sent from an output adapter	
	Communication when XTP collaboration is used to remotely access a cache	
	Communication when Hadoop collaboration is used to remotely access a Hadoop system	

Threats to protected resources and their countermeasures

The table below lists the possible threats to protected resources, and their respective countermeasures:

Type of protected resource	Threat	Security countermeasure
File	Tampering or destroying	Setting permissions Authenticating operation permissions for the CEP Server
Network	Sniffing	Placing on a secure segment

Overview of countermeasures for threats to protected resources

The table below provides an overview of each possible security countermeasure:

Security countermeasure	Overview of countermeasure
Setting permissions	Set operating system permissions for files included in the protected resources. Set appropriate permissions to suit the security roles.
Authenticating operation permissions for the CEP Server	Use operating system authentication. Only allow suitable users to login to the operating system.

Security countermeasure	Overview of countermeasure
	<p>This product assumes that users who are given authentication to login to the CEP Server can be trusted with operating a CEP engine and referencing an engine log.</p> <p>In addition, some operations, such as reconfiguring a CEP engine and starting and stopping the CEP service, can only be performed by a superuser.</p>
Placing on a secure segment	To inhibit data sniffing and hacking, place the CEP Server on a secure segment.

6.2.3 Designing Security for this Product

This section explains how to design security for the systems using this product, as follows:

- Authenticating for the CEP Server
- Designing suitable access permissions
- Designing the network

Authenticating for the CEP Server

A superuser of the operating system of the server to which this product is applied can operate the CEP Server.

In addition, authentication is not performed when events are sent to the CEP Server from outside the system. Build a firewall or use event sender business applications to build a system in which authentication is performed when events are sent to the CEP Server.

Designing suitable access permissions

Set suitable access permissions for the files below as a countermeasure to prevent file tampering and destruction.

The table below describes the access permissions to be set for each file.

File	Reference permission	Write permission
Engine configuration file	Superuser	Superuser
Master data	Engine execution user	None
Data for investigation	Superuser	Superuser

The access permissions for files generated by a CEP engine, such as the event log and resource log, will be set automatically.

Designing the network

If the system has been located according to the system configuration supported by this product, a third party will be unable to reference data transmitted over the network.

6.3 Maintenance

This section explains the operations required when maintaining this product.

6.3.1 Collecting Data for Investigation when a Problem Occurs

This product provides a command to collect data for investigation when a problem occurs (refer to Chapter 1, "Collecting Diagnostics Data" in *Troubleshooting* for details).

6.3.2 Backup and Restore

Resources must be backed up periodically, in case the system on the CEP Server fails.

The table below lists the resources to be backed up.

Resource name	Content	When to back up
Engine configuration file	File defining the CEP engine configuration	When settings are changed.
Master data	Master data to be stored in the CEP engine	When storing on the CEP Server. When editing on the CEP Server.
Definition information (definition file)	Definition information to be deployed to the CEP engine	When storing on the CEP Server. When editing on the CEP Server.

- The event log stored on the CEP Server is for checking operation, so it does not need to be backed up or restored.
- If the master data is backed up by the data provider, it does not need be backed up.
- If the definition information is backed up by the definition information developer, it does not need to be backed up.
- Refer to the manual of the relevant collaboration product for information on the backup and restore of resources in the collaboration system.

6.3.2.1 Backup Procedure

Follow the steps below as a superuser to back up the resources:

1. Back up the engine configuration file
2. Back up the master data
3. Back up the definition information (definition file)

This explanation uses the following directory to back up the resources:

```
/backup
```

Back up the engine configuration file

Back up the engine configuration file specified during execution of cepconfigeng.

Example

When the engine configuration file is `"/etc/opt/FJSVcep/Engine.xml"`.

```
# cp -p /etc/opt/FJSVcep/Engine.xml /backup/<ENTER>
```

Back up the master data

Back up the master data stored on the CEP Server.

Example

When the master data is stored in `"/masterdata"`.

```
# cp -rp /masterdata /backup/<ENTER>
```

Back up the definition information (definition file)

Back up the definition information (definition file) stored on the CEP Server.

Example

When the definition file is stored in `"/application"`.


```
# cp -rp /application /backup/<ENTER>
```

6.3.2.2 Restore Procedure

Follow the steps below as a superuser to restore the resources:

1. Restore the engine configuration file
2. Restore the master data
3. Restore the definition information

This section assumes that this product has been installed and set up (refer to "[Chapter 4 Installation and Setup](#)" for details).

It also assumes that the following resource backup directory exists:

```
/backup
```

Restore the engine configuration file

Restore the engine configuration file (refer to "[Chapter 4 Installation and Setup](#)" for information on how to reconfigure a CEP engine using a restored engine configuration file).



Example

When the engine configuration file that is the backup source is in `"/etc/opt/FJSVcep/Engine.xml"` and the backed up engine configuration file is `"/backup/Engine.xml"`.

```
# cp -p /backup/Engine.xml /etc/opt/FJSVcep/<ENTER>
```

Restore the master data

Restore the master data.



Example

When the master data that is the backup source is in `"/masterdata"` and the backed up master data is `"/backup/masterdata"`.

```
# cp -rp /backup/masterdata /<ENTER>
```

Restore the definition information (definition file)

Restore the definition information (definition file) (refer to "[6.1.4.1 Deploying Definition Information](#)" for information on how to redeploy restored definition information).



Example

When the definition information (definition file) that is the backup source is in `"/application"` and the backed up definition information is `"/backup/application"`.

```
# cp -rp /backup/application /<ENTER>
```

6.3.3 Applying Updates

Update information provided by software products including this product can be obtained from "UpdateSite", which is an integrated software update information provision site.

For this product, apply the updates released from the products below:

Product name	Version
Interstage Big Data Complex Event Processing Server	V1.0.0
Interstage Application Server Enterprise Edition (64-bit)	V10.0.0

Follow the steps below as a superuser to apply an update to this product:

1. Stop cron.

```
# /etc/init.d/crond stop<ENTER>
```

2. Execute cepstopserv.

```
# cepstopserv<ENTER>
```

3. Apply the update.

Apply the update according to the update information file.

4. Execute cepstartserv.

```
# cepstartserv<ENTER>
```

5. Start cron.

```
# /etc/init.d/crond start<ENTER>
```



Note

Note on when cron is stopped

This product uses cron to obtain the resource usage, which is configured by this product using the crontab of the engine execution user.

When the CEP service is stopped, and when the process to obtain the resource usage from the cron service is executed, an error may be output using the execution cycle set in the cron service (for obtaining the resource usage, 10-minutes intervals). This error is canceled by starting the CEP service again.

When cron is stopped, both the cron setting for the process to obtain the resource usage and the cron setting for other than the engine execution user stop. Therefore, when applying updates, select a time period when stopping cron will not cause problems.

If the manual is also being updated according to the update information file attached to the update, obtain the latest manual from the following site:

```
http://www.fujitsu.com/support/software/manual/ (as at October 2012)
```

6.3.4 Tuning

This section explains tuning for this product, as follows:

- [Tuning JVM options](#)
- [Tuning file descriptors](#)
- [Tuning trace logs](#)

6.3.4.1 Tuning JVM Options

Two Java VMs (hereafter, referred to as "JVMs") operate within one CEP engine - one performs input adapter and high-speed filter processing, and the other performs complex event processing and output adapter processing.

If analysis of the resource log has revealed an excess or insufficiency of JVM heap memory, tune the JVM options of the CEP engine. This will avoid problems such as decline in performance caused by garbage collections in the CEP engine.

The table below lists the JVM options that can be tuned:

JVM option	Initial value at CEP engine creation
Maximum value of the memory allocation pool	2048MB
Initial value of the memory allocation pool	512MB
Maximum value of the permanent generation area	192MB

Follow the steps below to tune the JVM options:

1. Check the current JVM option settings.
2. Using the resource log, check the heap memory usage.
3. Calculate the heap memory size required - if there is enough memory, there is no need to tune.
4. Change the JVM option settings.
5. Start or restart the CEP engine (to reflect the changed settings).
6. Follow these steps again to check the changed settings.

Check the current JVM option settings

Execute `cepgetjvmopt` to check the JVM option settings of the CEP engine.

A general user permission can be used to execute this command.



Example

Executing `cepgetjvmopt`

When checking the JVM options of complex event processing for a CEP engine (CepEngine).

```
$ cepgetjvmopt cep -e CepEngine<ENTER>
xmxSize           :5120m
xmsSize           :256m
maxPermSize       :96m
Command cepgetjvmopt executed successfully.
```

Note that "m" stands for "MB".

Using the resource log, check the heap memory usage

Analyze the resource log to check the maximum value of the Java VM heap memory usage for the CEP engine. Specifically, check the new generation area, old generation area, and permanent generation area.

The items below relate to JVM heap memory:

- *jheapNewUsed*
- *jheapNewFree*
- *jheapNewTotal*
- *jheapOldUsed*
- *jheapOldFree*
- *jheapOldTotal*
- *jheapNewPlusOldTotal*
- *jheapPermUsed*

- *jheapPermFree*
- *jheapPermTotal*

This applies to the resource log of the high-speed filter and the resource log of complex event processing.

Refer to "[6.1.6.4 Checking the Resource Usage of the CEP Engine](#)" for information on the format of the resource log.

Calculate the heap memory size required

If the resource log has been used to analyze the respective maximum values of the new generation area, old generation area, and permanent generation area, calculate the heap memory size required, as shown below.

If the heap memory size required is less than the setting value, then there is no need to tune.

Calculate the maximum value of the memory allocation pool

$$(maxValOfNewGenerationArea + maxValOfOldGenerationArea + maxValOfPermanentGenerationArea) \times 1.2$$

1.2 is the safety factor - if heap memory usage can vary widely with the time of the year or time period, set a higher safety factor.

Calculate the maximum value of the permanent generation area

$$maxValOfPermanentGenerationArea \times 1.2$$

1.2 is the safety factor - if heap memory usage can vary widely with the time of the year or time period, set a higher safety factor.

Change the JVM option settings

Execute `cepsetjvmopt` as a superuser to change the JVM option settings of the CEP engine.



Example

Executing the `cepsetjvmopt` command

When changing the maximum value of the memory allocation pool in the JVM options for complex event processing of a CEP engine (CepEngine) to 512 MB. The respective defaults are used for the initial value of the memory allocation pool and the maximum value of the permanent generation area.

```
# cepsetjvmopt cep -xmx 512m -e CepEngine<ENTER>
Command cepsetjvmopt executed successfully.
```

Note that "m" stands for "MB".



Note

If an option of `cepsetjvmopt` (-xmx, -xms, or -xxmp) is omitted, the default value for that option will be used.

Start or restart the CEP engine.

Start or restart the CEP engine in order to reflect the changed JVM option settings in the CEP engine.

If the CEP engine is running, it must be stopped (refer to "[6.1.7 Stopping the CEP Engine](#)" for details).

Refer to "[6.1.5 Starting the CEP Engine](#)" for details.

A general user permission can be used to perform this task.



Example

Restarting the CEP engine

When restarting a running CEP engine (CepEngine).

```

$ cepstopeng -e CepEngine<ENTER>
Command cepstopeng executed successfully.
$ cepstarteng -e CepEngine<ENTER>
Command cepstarteng executed successfully.

```

6.3.4.2 Tuning File Descriptors

The number of file descriptors required for operating one CEP engine depends on the number of simultaneous socket connections or high-speed filter statements described in the rule definition filter rules.

If an error occurs due to insufficient filter descriptors, or file descriptors are found to be excessive or inadequate according to resource log analysis, it is necessary to specify the file descriptor upper limit.

The required number of file descriptors for each CEP engine is calculated as follows:

$$\text{numOfFileDescriptorsRequired} = A + (8 \times B) + 371$$

Table 6.1 Placeholders in the file descriptor number tuning formula

Item	Description	Required Number
<i>A</i>	Number of simultaneous socket connections	Use the largest number of assumed simultaneous connections. Refer to the CEP engine operation status or the resource log output of the number of simultaneous connections to check the change in the current or previous number of simultaneous connections. Refer to "8.4 cepdispeng" or "6.1.6.4 Checking the Resource Usage of the CEP Engine" for information on how to reference it.
<i>B</i>	Total number of IF-THEN statements described in the filter rules	Refer to the CEP engine operation status or the resource log output of the number of high-speed filter statements to check the current number of IF-THEN statements. Refer to "8.4 cepdispeng" or "6.1.6.4 Checking the Resource Usage of the CEP Engine" for information on how to reference it.

Note

When the file descriptors are exhausted, operations might no longer be able to continue. Therefore, specify a value with some allowance for the number of simultaneous connections in *A*.

Compare the value calculated above with the actual number of file descriptors that can be used by the engine execution user. If there are multiple CEP engines, compare the largest value of the calculation result from each CEP engine.

Execute the command below as a superuser to display the maximum number of file descriptors that the engine execution user can use:

```
# /bin/su -c 'ulimit -n' engineExecutionUserName <ENTER>
```

If the calculated value is larger, then change `/etc/security/limits.conf` with it and restart the OS (refer to OS documentation for information on how to change the value).

Example

Setting `/etc/security/limits.conf`

This example changes the maximum number of file descriptors that the engine execution user "isbdcep" can use from the default value of 1024 to 2048.

```
isbdcep soft nofile 2048
isbdcep hard nofile 2048
```

6.3.4.3 Tuning Trace Logs

In this product, the input-output information for event sender applications and receipt processing states for event sender application requests are output to trace logs.

By stopping the output of trace logs, it is possible to improve the efficiency of event processing when the HTTP adapter and SOAP adapter are used.



Stopping the output of trace logs will mean that more time will have to be spent to determine the cause of eventual problems. If efficiency is not a problem, it is recommended to output the trace logs.

The flow to stop (or resume) trace log output is as follows (perform these operations as superuser):

- (1) [Stop \(or Resume\) Logging Event Sender Application I/O Information](#)
- (2) [Stop \(or Resume\) Receipt Processing State of Event Sender Application Requests Log Output](#)
- (3) [Restarting the CEP Service](#)
- (4) [Starting the CEP Engine](#)

(1) Stop (or Resume) Logging Event Sender Application I/O Information

Edit `/var/opt/FJSVihs/servers/FJapache/conf/httpd.conf` and set the `IHSTraceLog` setting value to "off".



Setting `/var/opt/FJSVihs/servers/FJapache/conf/httpd.conf`

Before changing

```
IHSTraceLog "|/opt/FJSVihs/bin/ihsrlog -s logs/tracelog 2 5"
```

After changing

```
#IHSTraceLog "|/opt/FJSVihs/bin/ihsrlog -s logs/tracelog 2 5"  
IHSTraceLog off
```

When resuming log output, also edit the `httpd.conf` file and revert the contents to its previous state.

(2) Stop (or Resume) Receipt Processing State of Event Sender Application Requests Log Output

Execute the commands below to stop the receipt processing state log output contents of the event sender application request (please note that we added a backslash ("\") and newline characters in the command line example below for readability - the actual command line does not have a backslash and newline):

```
# /opt/FJSVisjee/bin/asadmin set \  
CEPengineNameflt-config.http-service.property.ISJEELogHttpTraceEnable=false <ENTER>
```

To resume log output, perform the following:

```
# /opt/FJSVisjee/bin/asadmin set \  
CEPengineNameflt-config.http-service.property.ISJEELogHttpTraceEnable=true <ENTER>
```



Executing the command

If the CEP engine name is "CepEngine", perform the following to deter log output (please note that we added a backslash ("\") and newline characters in the example below for readability - the actual command line does not have a backslash and newline).

```
# /opt/FJSVisjee/bin/asadmin set \  
  CepEngine_flt-config.http-service.property.ISJEELogHttpTraceEnable=false <ENTER>  
CepEngine_flt-config.http-service.property.ISJEELogHttpTraceEnable=false
```

(3) Restarting the CEP Service

To reflect your changes, execute `cepstopserv` and then `cepstartserv` to restart the CEP service for the log output.



Example

Restarting the CEP service

```
# cepstopserv <ENTER>  
(...)  
# cepstartserv <ENTER>  
(...)
```

(4) Starting the CEP Engine

Execute `cepstarteng` to start the CEP engine.



Example

Executing `cepstarteng`

```
# cepstarteng -e CepEngine <ENTER>  
Command cepstarteng executed successfully.
```

Chapter 7 Reliable System Operation

This chapter explains how to operate this product (hereafter, referred to as "BDCEP") using a reliable configuration.

7.1 Overview of Reliable System Operation

The Cluster Service of BDCEP can be used to build a reliable system using PRIMECLUSTER, in order to prevent a long-term suspension of business due to hardware faults in the CEP Server.

This feature can be used to allow the events that occur after a cluster switch has completed to be processed by a CEP Server on the standby node.

Knowledge of PRIMECLUSTER is required to use the Cluster Service (refer to the PRIMECLUSTER manuals for details).

The range of application of the Cluster Service is shown below:

Operation form of a cluster service

1:1 standby operation

Cluster products that can be used

- PRIMECLUSTER Enterprise Edition 4.3A10 or higher
- PRIMECLUSTER HA Server 4.3A10 or higher

Features that can be used in a cluster system

All features of this product can be used.



Note

- The statuses of high-speed filter rules and complex event processing rules that exist on the active node prior to or during a cluster switch are not transferred to the CEP Server on the standby node.
- Events that occur during a switch are not received.

7.2 Cluster Service Configuration

To operate BDCEP in a cluster system, the CEP engine must be configured and the development assets (such as definition information and master data) must be deployed so that the respective servers on the active node and standby node will have the same resource configuration.

Set up PRIMECLUSTER so that the operating system also operates on the active node and the standby node using the same IP address, through IP address takeover.

There is no need to deploy BDCEP resources such as the definitions and engine log to a shared disk - instead, use the local disks of the active node and standby node.

7.3 Building a Cluster Service Environment

Follow the steps below to set up the environment to run BDCEP in a cluster system (refer to the PRIMECLUSTER manuals for information on installing and operating PRIMECLUSTER):

1. Install and set up PRIMECLUSTER.

Install PRIMECLUSTER for both the active node and the standby node.

2. Set up IP address takeover.

Set up PRIMECLUSTER so that the active node and standby node can use IP address takeover to switch and then operate using the same IP address.

3. Install and set up BDCEP in the active node.

Activate IP address takeover in the active node, and then install and set up (for example, create a CEP engine) BDCEP, as usual (refer to "[Chapter 4 Installation and Setup](#)" for details).

Note that you must register the active node server information in the master server connection authorization list when performing Hadoop collaboration (refer to E.3, "Adding a Slave Server/Scaling Out" in the *User's Guide* of the Interstage Big Data Parallel Processing Server V1.0.0 for details).

You must also specify the active node server information in the ix-environment.xml <node> tag when performing XTP collaboration (refer to Section 4.6.2, "Tasks for the XTP Client Node to be Extended" in the *Setup and Operation Guide* of the Interstage eXtreme Transaction Processing Server V1.0.0 for details).

4. Deploy development assets to the active node.

Deploy the required development assets, such as rule definitions, to the active node (refer to "[5.6 Deploying Development Assets](#)" for details).

When a status in which events can be processed is reached, stop the CEP engine and the CEP service of the active node.

5. Save the RC procedures of the active node.

Save the start shell script "S99startis" stored in the directories below to any directory for storing backup resources:

- /etc/rc2.d
- /etc/rc3.d
- /etc/rc4.d
- /etc/rc5.d

Save the stop shell script "K00stopis" stored in the directories below to any directory for storing backup resources:

- /etc/rc.d/rc0.d
- /etc/rc.d/rc1.d
- /etc/rc.d/rc6.d

Disable the automatic start of the Interstage Java EE DAS service:

```
# /sbin/chkconfig --del FJSVijdas <ENTER>
```

6. Switch to the standby node.

Use a PRIMECLUSTER cluster switch operation to switch to the standby node.

7. Install and set up BDCEP in the standby node.

With IP address takeover activated in the standby node, install and set up BDCEP in it. Make the CEP engine settings the same as those in the active node (refer to "[Chapter 4 Installation and Setup](#)" for details).

Note that you must register the standby node server information into the master server connection authorization list when performing Hadoop collaboration (refer to E.3, "Adding a Slave Server/Scaling Out" in the *User's Guide* of the Interstage Big Data Parallel Processing Server V1.0.0 for details).

You must also specify the standby node server information in the ix-environment.xml <node> tag when performing XTP collaboration (refer to Section 4.6.2, "Tasks for the XTP Client Node to be Extended" of the *Setup and Operation Guide* of the Interstage eXtreme Transaction Processing Server V1.0.0 manual for details).

8. Deploy development assets to the standby node.

Deploy the same development assets as those in the active node to the standby node (refer to "[5.6 Deploying Development Assets](#)" for details).

There is no need to perform work that duplicates that in step 4, "Deploy development assets to the active node", such as providing data and deploying collaboration applications. Note, however, that if master data is being placed in the local disk of the CEP Server, it must be stored at the standby node in the same way as it was stored at the active node.

When a status in which events can be processed is reached, stop the CEP engine and the CEP service of the standby node.

9. Save the RC procedures of the standby node.

Perform the same operation at the standby node as was performed at the active node.

10. Change and register the Cmdline resource.

Edit the Cmdline resource sample file below and register it in the cluster system, so that the CEP engine will start automatically at a cluster switch:

```
/opt/FJSVcep/HA/sample/SERVICE_BDCEP
```

Edit the sample file as shown below, and then register it in the cluster system using a PRIMECLUSTER operation.

- Change the *Engine-Name* part in the Cmdline resource below to the CEP engine name created at the active node and standby node, and then register it.

```
STARTCMDE='/opt/FJSVcep/bin/cepstarteng -e Engine-Name'
```

- To use multiple CEP engines, describe as many CEP engines as cepstarteng is to start.

Example

Description example when multiple CEP engines are to be used

```
(...)  
STARTCMDE='/opt/FJSVcep/bin/cepstarteng -e CepEngine1'  
STARTCMDE2='/opt/FJSVcep/bin/cepstarteng -e CepEngine2'  
(...)  
start() {  
    ${STARTCMD} > /dev/null 2>&1  
    ${STARTCMDE} > /dev/null 2>&1  
    ${STARTCMDE2} > /dev/null 2>&1  
}  
(...)
```

7.4 Operating a Cluster Service

This section explains some points to consider in cluster service operation.

Starting and stopping

- When the active node is started, the CEP service and CEP engine will start.
- PRIMECLUSTER process monitoring will be performed for the CEP service. Therefore, do not execute cepstopserv unless it is for a cluster switch.

Sending events

- Send events to the takeover IP address.

Changing CEP engines and development assets

- Set up the CEP engine configurations and deploy the development assets (such as definition information and master data) so that the active node and the standby node have the same resource configuration.

Cluster operation and response at an active node fault

- When a hardware fault occurs at the active node and the cluster service detects the abnormality, the active node will be stopped and the CEP engine and CEP service will stop.
- After a switch to the standby node, event sending to the takeover IP address can restart.

- When an abnormality occurs at the active node, take action to resolve the hardware fault that caused the cluster switch, and then take action to allow a switch back in case an abnormality occurs at the standby node.

Chapter 8 Command Reference

This chapter explains the commands that the Interstage Big Data Complex Event Processing Server provides.

Command name	Function	Administrator permission
cepcollectinfo	Collects data for investigation in batch	Yes
cepconfigeng	Configures a CEP engine (*1)	Yes
cepdeployrsc	Deploys a development asset (*1)	No
cepdispeng	Displays the status of a CEP engine	No
cepdispserv	Displays the status of the CEP service	Yes
cepsetjvmopt	Displays JVM options	No
cepgetrsc	Displays development assets	No
cepsetjvmopt	Sets JVM options (*1)	Yes
cepstarteng	Starts a CEP engine (*1)	No
cepstartserv	Starts the CEP service (*2)	Yes
cepstopeng	Stops a CEP engine (*1)	No
cepstopserv	Stops the CEP service (*2)	Yes
cepundeployrsc	Undeploys a development asset (*1)	No

Yes: Must be executed by a superuser

Note

- Commands indicated by (*1) cannot be executed at the same time - if they are (including executing more instances of the same command), commands executed after the first instance will end abnormally.
- Do not execute the commands for starting and stopping the CEP service (indicated by (*2)) at the same time. Also, do not execute other commands (those not indicated by (*2)) while these commands are being executed.

8.1 cepcollectinfo

Name

cepcollectinfo - Data investigation collection

Format

```
cepcollectinfo [path]
```

Function description

This command collects data for investigation for the Interstage Big Data Complex Event Processing Server, in batch.

The compressed file (tar.gz format) below will be generated in the path specified in the command.

path/collect/yyyyMMddHHmmss

yyyyMMddHHmmss indicates the date and time the data was collected.

Arguments

path

Specify the directory for storing the data for investigation.

If this option is omitted, the data will be stored in /tmp.

If the specified path does not exist, the command will return an error without collecting the data for investigation.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



Example

Execution example

When a path is specified:

```
# cepcollectinfo /var/tmp <ENTER>
Compressing collect data is started.
Compressing collect data is finished.
Collecting is finished. (/var/tmp/collect/BDCEP_20120413141134.tar.gz)
Command cepcollectinfo executed successfully.
```

When a path is not specified (data is stored under /tmp):

```
# cepcollectinfo <ENTER>
Compressing collect data is started.
Compressing collect data is finished.
Collecting is finished. (/tmp/collect/BDCEP_20120413141134.tar.gz)
Command cepcollectinfo executed successfully.
```

8.2 cepconfigeng

Name

cepconfigeng - CEP engine configuration

Format

```
cepconfigeng -f XMLfilePath
```

Function description

This command configures a CEP engine.

It compares the content of the specified engine configuration file with that of a current CEP engine that has been configured, and then creates, changes the settings of, or deletes the CEP engine according to the differences between them.

Always save the engine configuration file for the next time the configuration content is to be changed.

When the command is executed, the prompt below is displayed - type "y" and press the Enter key to execute the change. Execution can be canceled by typing "n" or "q" and pressing the Enter key.

```
Are you sure you want to change the CEP Engine configuration? [y,n,q]:
```

Creating a CEP engine

This command creates a new CEP engine.

After the CEP engine is created, it must be started separately.

Changing CEP engine settings

This command changes the settings (Items other than CEP engine name) of a CEP engine that has been configured.

The CEP engine that is to be changed must be stopped beforehand. After this command is executed, the CEP engine must be started separately.

Deleting a CEP engine

This command deletes a CEP engine that has been configured.

The CEP engine that is to be deleted must be stopped beforehand.

If development assets have been deployed to the CEP engine that is to be deleted, the CEP engine can be deleted even if the development assets are not undeployed beforehand.



See

Refer to "8.9 cepstarteng" and "8.11 cepstopeng" for information on starting and stopping a CEP engine.

Arguments

-f XMLfilePath

Specify the engine configuration file that describes the definition content of the CEP engine. We recommend specifying an absolute path to avoid specification error.



See

Refer to "9.1.1 Engine Configuration File" for details.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



Note

- A maximum of five CEP engines can be created. If six or more CEP engines are defined, this command ends abnormally.
- Specify the CEP engine name so as not to be confused with other CEP engine names. If the CEP engine in the CEP configuration file has the same name, this command will end abnormally.
- When creating a new CEP engine, always include in the engine configuration file the configured content of the CEP engine that has already been configured in order to retain that CEP engine. If this content is not included, the CEP engine that has been configured will be deleted.
Similarly, when changing CEP engine settings, always change the target settings only, and be sure to leave the other settings in the engine configuration file.
- If this command ends abnormally due to an error, it is possible that the CEP engine configuration has an inconsistency. To complete changing the CEP engine settings, remove the cause of the error and then execute the command again.

- If a CEP engine is deleted, development assets that have been deployed are undeployed automatically. If this command ends abnormally due to an error during execution, the status of the development asset which was undeployed prior to the error is returned to the one before execution of this command, however there may be a CEP engine configuration inconsistency. To complete the deletion of the CEP engine, remove the cause of the error and then execute the command again.
- If this command is executed while sending an event, it may end abnormally. After stopping all CEP Server event senders, execute the command again.

Example

Example of output at normal end

When CEP engines (CepEngine1 and CepEngine2) are created:

/etc/opt/FJSVcep/Engine.xml:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<subSystemConfig xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1">
  <engineConfig id="CepEngine1">
    <logging>
      <type>bdpp</type>
      <directory>hadoop</directory>
      <loggingMaxOpenFile>6</loggingMaxOpenFile>
      <loggingRotationCycle>300</loggingRotationCycle>
    </logging>
    <socketAdapterPort>9600</socketAdapterPort>
  </engineConfig>
  <engineConfig id="CepEngine2">
    <logging>
      <type>file</type>
    </logging>
    <socketAdapterPort>9601</socketAdapterPort>
  </engineConfig>
</subSystemConfig>
```

Command execution result:

```
# cepconfigeng -f /etc/opt/FJSVcep/Engine.xml <ENTER>
Are you sure you want to change the CEP Engine configuration? [y,n,q]:y <ENTER>
Command cepconfigeng executed successfully.
```

When CepEngine1 is deleted from the CEP engines that have been configured (leaving CepEngine2):

/etc/opt/FJSVcep/Engine.xml:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<subSystemConfig xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1">
<!--
  <engineConfig id="CepEngine1">
    <logging>
      <type>bdpp</type>
      <directory>hadoop</directory>
      <loggingMaxOpenFile>6</loggingMaxOpenFile>
      <loggingRotationCycle>300</loggingRotationCycle>
    </logging>
    <socketAdapterPort>9600</socketAdapterPort>
  </engineConfig>
-->
  <engineConfig id="CepEngine2">
    <logging>
      <type>file</type>
    </logging>
    <socketAdapterPort>9601</socketAdapterPort>
```

```
</engineConfig>
</subSystemConfig>
```

Command execution result:

```
# cepconfigeng -f /etc/opt/FJSVcep/Engine.xml <ENTER>
Are you sure you want to change the CEP Engine configuration? [y,n,q]:y <ENTER>
Command cepconfigeng executed successfully.
```

Example of output at abnormal end

When an attempt is made to delete a CEP engine that is running.

```
# cepconfigeng -f /etc/opt/FJSVcep/Engine.xml <ENTER>
Are you sure you want to change the CEP Engine configuration? [y,n,q]:y <ENTER>
The setup processing failed. EngineId=(CepEngine) Reason=(The job is already running.)
Command cepconfigeng execution failed.
```

8.3 cepdeployrsc

Name

cepdeployrsc - Development asset deployment

Format

```
cepdeployrsc resource [-o] [-e engineName] -f XMLfilePath
```

Function description

This command deploys a development asset to a CEP engine.

Specify the type of development asset (event type definition, rule definition, master definition, or SOAP listener definition) to be deployed.

When the command is executed a prompt is displayed - type "y" and press the Enter key to execute the deployment. Execution can be canceled by typing "n" or "q" and pressing the Enter key.



Example

Example of query (for a SOAP listener definition):

```
Are you sure you want to deploy the SOAP listener definition?(default: y) [y,n,q]:
```

Arguments

resource

Specify the type of development asset.

eventtype

Event type definition

rule

Rule definition

master

Master definition

listener

SOAP listener definition

-o

This option is used to overwrite the definition content.

If this option is omitted, and if a development asset with the same ID as that specified in the definition file has already been deployed, an error message will be output and the deployment will fail.

-e *engineName*

Specify the name of the CEP engine to which the development asset is to be deployed.

If only one CEP engine has been configured, this option can be omitted.

If this option is omitted, and if there are two or more configured CEP engines, an error message will be output and the deployment will fail.

-f *XMLfilePath*

Specify the definition file that describes the definition content of the development asset. We recommend specifying an absolute path.



Refer to "9.2 Defining Development Assets" for information on the definition file.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



- A file storing detailed display results (XML format) output by `cepgetrsc` can also be specified as the definition file (refer to "8.7 `cepgetrsc`" for details).
- If the CEP engine that is the deployment target is running, deployment cannot be performed.



Example of output at normal end

When a SOAP listener definition file (`listenerdeploy.xml`) is deployed to a CEP engine (`CepEngine1`):

```
$ cepdeployrsc listener -o -e CepEngine1 -f /tmp/listenerdeploy.xml <ENTER>
Are you sure you want to deploy the SOAP listener definition?(default: y) [y,n,q]:y <ENTER>
Command cepdeployrsc executed successfully.
```

Example of output at abnormal end

When an attempt is made to deploy a SOAP listener definition during the execution of a command that cannot be executed at the same time:

```
$ cepdeployrsc listener -o -e CepEngine1 -f /tmp/listenerdeploy.xml <ENTER>
Are you sure you want to deploy the SOAP listener definition?(default: y) [y,n,q]:<ENTER>
Processing cannot be performed because another command is executing.
Command cepdeployrsc execution failed.
```

8.4 cepdispeng

Name

cepdispeng - CEP engine information display

Format

```
cepdispeng [-i] [-a | -e engineName]
```

Function description

This command displays the basic information or the operation status of a CEP engine.

Arguments

-i

This option is used to display the basic information of a CEP engine.

If this option is not specified, the operation status of the CEP engine will be displayed.

The table below explains the items displayed:

Basic information

Item name	Content	Displayed?
engineId	CEP engine name	Yes
eventtype	Development asset ID of the deployed event type definition	Conditional
rule	Development asset ID of the deployed rule definition	Conditional
master	Development asset ID of the deployed master definition	Conditional
listener	Development asset ID of the deployed SOAP listener definition	Conditional

Yes: Always displayed

Conditional: Displayed if the development asset has been deployed



If two or more development assets of the same definition are deployed, they are displayed in multiple lines.

Operation status

Item name	Content	Displayed?
engineId	CEP engine name.	Yes

Item name	Content	Displayed?
port	Port number for socket communication (socket adapter port).	Conditional (*1)
socket	Number of simultaneous connections in socket communication.	Conditional (*1)
status_filter	Status of the CEP engine (high-speed filter). RUN : Running normally STARTING : Starting STOP : Stopped STOPPING : Stopping ABNORMAL : Abnormal	Yes
status_cep	Status of the CEP engine (complex event processing). RUN : Running normally STARTING : Starting STOP : Stopped STOPPING : Stopping ABNORMAL : Abnormal	Yes
inEvent_filter	Number of input events for the CEP engine (high-speed filter). Indicates the number of events input to the high-speed filter. Displays a cumulative value from when the CEP engine started.	Conditional (*2)
inEvent_cep	Number of input events for the CEP engine (complex event processing). Indicates the number of events input to complex event processing. Displays a cumulative value from when the CEP engine started.	Conditional (*3)
outEvent_filter	Number of output events for the CEP engine (high-speed filter). Indicates the number of events sent to complex event processing. Displays a cumulative value from when the CEP engine started.	Conditional (*2)
outEvent_cep	Number of output events for the CEP engine (complex event processing). Indicates the number of events sent to a user-developed Web service. Displays a cumulative value from when the CEP engine started.	Conditional (*3)
logging_filter	Number of loggings for the CEP engine (high-speed filter).	Conditional (*2)
logging_cep	Number of loggings for the CEP engine (complex event processing).	Conditional (*3)
heap_filter	Java heap usage and area size for the CEP engine (high-speed filter). Displayed in the following format: Java heap usage (bytes)/area size (bytes)	Conditional (*2)
heap_cep	Java heap usage and area size for the CEP engine (complex event processing).	Conditional (*3)

Item name	Content	Displayed?
	Displayed in the following format: Java heap usage (bytes)/area size (bytes)	
useMemory_filter	Memory usage of CEP engine (high-speed filter) processes (kilobytes).	Conditional (*2)
useMemory_cep	Memory usage of CEP engine (complex event processing) processes (kilobytes).	Conditional (*3)
countRule_filter	Number of rules applied to the CEP engine (high-speed filter).	Conditional (*2)
countRule_cep	Number of rules applied to the CEP engine (complex event processing).	Conditional (*3)
countListener	Number of applied listeners. Indicates the number of SOAP listener definitions applied to the CEP engine (complex event processing).	Conditional (*3)

Yes: Always displayed

Conditional: Displayed in the following cases:

*1: If socket communication is being used (if a socket adapter is being used by the CEP engine)

*2: If the CEP engine (high-speed filter) is running normally (if status_filter is "RUN")

*3: If the CEP engine (complex event processing) is running normally (if status_cep is "RUN")

-a

This option is used to display information on all configured CEP engines.

-e *engineName*

Specify the name of the CEP engine for which information is to be displayed.

If only one CEP engine has been configured, this option can be omitted.

If this option is omitted, and if there are two or more configured CEP engines, an error message will be output and the display will fail.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



Example

Example of output at normal end

When the operation statuses of all configured CEP engines are displayed:

```
$ cepdispeng -a <ENTER>
engineId      :CepEngine1
port          :9600
```

```

socket          :1
status_filter   :RUN
status_cep     :RUN
inEvent_filter  :100
inEvent_cep    :100
outEvent_filter :100
outEvent_cep   :100
logging_filter  :100
logging_cep    :100
heap_filter     :98,725,088 / 532,545,536
heap_cep       :98,725,088 / 532,545,536
useMemory_filter :823,493,339
useMemory_cep  :823,493,339
countRule_filter :1
countRule_cep  :1
countListener  :1

engineId       :CepEngine2
port           :9601
status_filter  :STOP
status_cep    :STOP
Command cepdispeng executed successfully.

```

When the basic information of a CEP engine (CepEngine1) is displayed:

```

$ cepdispeng -i -e CepEngine1 <ENTER>
engineId       :CepEngine1
eventtype     :EVENT_01
rule          :RULE_01
master        :MASTER_01
listener      :LISTEN_01
listener      :LISTEN_02
Command cepdispeng executed successfully.

```

8.5 cepdisperserv

Name

cepdisperserv - CEP service status display

Format

```
cepdisperserv
```

Function description

This command displays the status of the CEP service - it displays the status of each of the multiple services that make up the CEP service entity.

The services that make up the CEP service are as follows:

- Interstage Java EE DAS service
- Interstage Java EE Node Agent service
- IJServer cluster
- Interstage HTTP Server

- Apache Tomcat
- PostgreSQL

The status of each service is explained below.

Interstage Java EE DAS service

Display	Explanation
started	Running status
stopped	Stopped status
starting	Undergoing start processing
stopping	Undergoing stop processing
unknown	Unknown status

Note

"unknown" (unknown status) indicates that the status of the service could not be confirmed.

Some kind of abnormality, such as a service hang-up, is occurring, so wait for a few moments and then execute the command again to confirm the status.

Interstage Java EE Node Agent service

The display is the same as that of the Interstage Java EE DAS service.

IJServer cluster

Display	Explanation
starting	Undergoing start processing
running	Running status
stopping	Undergoing stop processing
not running	Stopped status
partially running	Degraded operation status

Note

When executing this command after cepstopserv, the message below is output, but the IJServer Cluster is stopped, so no action is required:

```
asadmin: ERROR: ISJEE_OM2997: Unable to connect to admin-server at given host:
[localhost] and port: [12001]. Please check if this server is up and running and
that the host and port provided are correct.
asadmin: ERROR: ISJEE_CLI137: Command list-clusters failed.
```

Please note that we added newlines in the example above (line 1 and 2) for readability only - the actual message does not have a newline.

Interstage HTTP Server

Item name	Content
Web Server Name	Web server name.
Status	Status of the Web server. Running: Running Stopped: Stopped
Configuration File	Environment configuration file.
Server Version	Server version of the Interstage HTTP Server.
Current Time	Current date and time.
Start Time	Start date and time.
Daemon Process ID	Process ID of the daemon process.
Listening Port	IP address and port number of the Web server that is to receive connection requests. For communication using SSL protocol, "HTTPS" will be displayed for the port number. An IPv6 address will be displayed within brackets "[" and "]".
requests currently being processed	Number of requests being processed.
idle servers	Number of communication processes (threads) on standby.

Note

If "Status" is "Stopped", from "Configuration File" onwards will not be displayed.

Also, when executing this command after `cepstopserv`, the message below is output, but the Interstage HTTP Server is stopped, so no action is required:

```
UX:IHS: ERROR: ihs81517: The Web Server (Interstage HTTP Server) did not start. [FJapache]
```

Apache Tomcat

Display	Explanation
<code>jsvc(pid <i>nnnn nnnn</i>) is running...</code>	Executing. The IDs of the processes that are running will be displayed in " <i>nnnn</i> ".
<code>jsvc is stopped</code>	Stopping.

PostgreSQL

Display	Explanation
<code>pg_ctl: server is running (PID: <i>nnnn</i>) /.../bin/postgres -D /.../data</code>	Executing. The ID of the process that is running will be displayed in " <i>nnnn</i> ".
<code>pg_ctl: no server running</code>	Stopping.

Arguments

None

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



Example

Example of output at normal end

When all services are running normally:

```
# cepdisperv <ENTER>
Interstage Java EE DAS service status.
ijdasstat
Name                               Status
-----
Interstage Java EE DAS              started

Interstage Java EE Node Agent service status.
ijnastat
Name                               Status
-----
Interstage Java EE Node Agent       started

IJServer Cluster status.
CEPAgentIJServerCluster running
CepEngine_flt not running
CepEngine_cep not running
Command list-clusters executed successfully.

Interstage HTTP Server status.
ihdisp

Web Server Name   : FJapache
Status            : Running
Configuration File: /opt/FJSVihs/servers/FJapache/conf/httpd.conf
Server Version    : FJapache/10.0
Current Time      : Monday, 09-Jul-2012 01:25:17
Start Time       : Sunday, 08-Jul-2012 15:20:31
Daemon Process ID : 2267
Child Process ID  : 2272 2273 2274 2275 2276
Listening Port    : [::]:80
0 requests currently being processed, 5 idle servers

Apache Tomcat status.
/sbin/service FJSVcep-rest status
jsvc (pid 2099 2096) is running...

PostgreSQL status.
su - bdcep_postgres -c "/opt/FJSVcep/postgres/packages/FJSVpgs83/bin/pg_ctl -D /var/opt/FJSVcep/postgres/data status"
```



```
pg_ctl: server is running (PID: 2088)
/opt/FJSVcep/postgres/packages/FJSVpgs83/bin/postgres "-D" "/var/opt/FJSVcep/postgres/data"

Command cepdisperv executed successfully.
```

8.6 cepgetjvmopt

Name

cepgetjvmopt - JVM options display

Format

```
cepgetjvmopt function -e engineName
```

Function description

This command displays the JVM options of a CEP engine.

The JVM options that are displayed are shown below.

Item	Content
xmxSize	Maximum value of memory allocation pool
xmsSize	Initial value of memory allocation pool
maxPermSize	Maximum value of permanent generation area

Arguments

function

Specify the feature for which the JVM options are to be displayed.

filter

High-speed filter

cep

Complex event processing

-e engineName

Specify the name of the CEP engine for which the JVM options are to be displayed.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



Example

Example of output at normal end

When the JVM options of the high-speed filter of a CEP engine (CepEngine1) are displayed:

```
$ cepgetjvmopt filter -e CepEngine1 <ENTER>
xmxSize           :5120m
xmsSize           :256m
maxPermSize       :96m
Command cepgetjvmopt executed successfully.
```

8.7 cepgetrsc

Name

cepgetrsc - Development assets display

Format

```
cepgetrsc resource [-e engineName] [-n resourceID]
```

Function description

This command displays a list of or details of the development assets that have been deployed.

Arguments

resource

Specify the type of development asset.

eventtype

Event type definition

rule

Rule definition

master

Master definition

listener

SOAP listener definition

-e *engineName*

Specify the name of the CEP engine for which the development assets are to be displayed.

If only one CEP engine has been configured, this option can be omitted.

If this option is omitted, and if there are two or more configured CEP engines, an error message will be output and the display will fail.

-n *resourceID*

Specify the development asset ID to be displayed in detail.

When this option is specified, the definition content of the development asset is displayed in XML format.



Refer to "9.2 Defining Development Assets" for information on display results in XML format.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end



- If diverting the results of the detailed display (XML format), modify each item accordingly.
- The results of this command include output messages in addition to XML format data.
If diverting the command results by redirecting them to a file, delete unnecessary messages before using the results.



Example of output at normal end

When the SOAP listener definitions that have been deployed to a CEP engine (CepEngine1) are displayed as a list:

```
$ cepgetrsc listener -e CepEngine1 <ENTER>
LISTEN_01
LISTEN_02
Command cepgetrsc executed successfully.
```

When the event type definitions that have been deployed to a CEP engine (CepEngine1) are displayed in detail:

```
$ cepgetrsc eventtype -e CepEngine1 -n EVENT01 <ENTER>
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<eventType xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="EVENT01">
  <comment>Event type definition</comment>
  <type>CSV</type>
  <xmlSchema></xmlSchema>
  <csvColumn>
    <column name="memberID" type="string"/>
    <column name="areaID" type="string"/>
    <column name="status" type="string"/>
  </csvColumn>
  <root></root>
  <useLogging>>false</useLogging>
  <loggingTableName></loggingTableName>
  <useCep>>true</useCep>
</eventType>
Command cepgetrsc executed successfully.
```

Example of output at abnormal end

When an attempt is made to display an event type definition that has not been deployed to a CEP engine (CepEngine2):

```
$ cepgetrsc eventtype -e CepEngine2 -n EVENT02 <ENTER>
The event type definition does not exist. Id=(EVENT02)
Command cepgetrsc execution failed.
```

The following information will be added to the error message in this example:

Id

Development asset ID that was specified

.....

8.8 cepsetjvmopt

Name

cepsetjvmopt - JVM options set

Format

```
cepsetjvmopt function [-xmx xmxSize] [-xms xmsSize] [-xxmp permSize] -e engineName
```

Function description

This command sets the JVM options of a CEP engine.

Arguments

function

Specify the feature for which the JVM options are to be set.

filter

High-speed filter

cep

Complex event processing

-xmx *xmxSize*

This option is used to set the maximum value of the memory allocation pool.

If this option is omitted, the default (2048m) will be used.

If the value specified here is less than the value specified in the -xms option, an error message will be output when the CEP engine is started, and the start will fail.

The following characters can be specified as units:

To specify KB (kilobytes): "k" or "K"

To specify MB (megabytes): "m" or "M"

If the unit is omitted, the specification will be in bytes. Specify a value greater than 1 MB that is a multiple of 1024.

-xms *xmsSize*

This option is used to set the initial value of the memory allocation pool.

If this option is omitted, the default (512m) will be used.

If the value specified here is less than 2624 KB, an error message will be output when the CEP engine is started, and the start will fail.

The following characters can be specified as units:

To specify KB (kilobytes): "**k**" or "**K**"

To specify MB (megabytes): "**m**" or "**M**"

If the unit is omitted, the specification will be in bytes. Specify a value greater than 1 MB that is a multiple of 1024.

-xxmp permSize

This option is used to set the maximum value of the permanent generation area.

If this option is omitted, the default (192m) will be used.

If the value specified here is less than 20.75 MB, then 20.75 MB will be used.

The following characters can be specified as units:

To specify KB (kilobytes): "**k**" or "**K**"

To specify MB (megabytes): "**m**" or "**M**"

If the unit is omitted, the specification will be in bytes. Specify a value greater than 1 MB that is a multiple of 1024.

-e engineName

Specify the name of the CEP engine for which the JVM options are to be set.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end

Note

This command may be used for the running CEP engine, however, to apply settings, the target CEP engine must be restarted.

Example

Example of output at normal end

When the JVM options of the high-speed filter of a CEP engine (CepEngine1) are set:

```
# cepsetjvmopt filter -mx 5120m -xms 256m -xxmp 96m -e CepEngine1 <ENTER>
Command cepsetjvmopt executed successfully.
```

8.9 cepstarteng

Name

cepstarteng - CEP engine start

Format

```
cepstarteng [-e engineName]
```

Function description

This command starts a CEP engine.

Arguments

-e engineName

Specify the name of the CEP engine to be started.

If only one CEP engine has been configured, this option can be omitted.

If this option is omitted, and if there are two or more configured CEP engines, an error message will be output and the start will fail.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end

Note

- Before this command is executed, cepdeployrc must be used to deploy at least one event type definition. If an event type definition has not been deployed, an error message will be output when the CEP engine is started, and the start will fail.
- This command monitors the start of the CEP engine until it has completed, and the command ends normally if it can confirm that the start has completed.
- If this command is not completed within 180 seconds, force the CEP engine to stop. If the forced stop is then completed within 180 seconds, a message is output stating that the startup processing has timed out and that this command has ended abnormally. If the forced stop is not completed within 180 seconds, a message is output stating that the forced stop has failed and that this command has ended abnormally.

Example

Example of output at normal end

When a CEP engine (CepEngine1) is started:

```
$ cepstarteng -e CepEngine1 <ENTER>
Command cepstarteng executed successfully.
```

Example of output at abnormal end

When an attempt is made to start a CEP engine that does not exist:

```
$ cepstarteng -e CepEngine2 <ENTER>
An incorrect value was entered. Reason=(The engine does not exist.)
Command cepstarteng execution failed.
```

The following information will be added to the error message in this example:

Reason

Cause of the error (reason for the failure to start the CEP engine)

.....

8.10 cepstartserv

Name

cepstartserv - CEP service start

Format

```
cepstartserv
```

Function description

This command starts the CEP service - it starts in batch the multiple services that make up the CEP service entity.

Start processing for services that are already started will be skipped.

The services started by this command are as follows:

- Interstage Java EE DAS service
- Interstage Java EE Node Agent service
- Interstage Management Console
- Interstage HTTP Server
- PostgreSQL
- Apache Tomcat

Arguments

None

End status

The following status codes are returned:

0

Normal end

8

Abnormal end

Note

- If the start of even one of the services that make up the CEP service fails, an error message will be output and the start of the CEP service will fail.
To complete the start of the CEP service, remove the cause of the error and then execute the command again.
- The CEP engine is not started by this command.
- This product uses cron for CEP engine resource usage acquisition. This command does not start cron, it is normally started automatically when OS starts (refer to "6.3.3 Applying Updates" for information on how to start cron manually).

Example

Example of output at normal end

When all services start successfully:

```
# cepstartserv <ENTER>
Starting Interstage Java EE DAS service.
ijdasstart: INFO: ijdas10000: Interstage Java EE DAS service has started.
Name                               Status
-----
Interstage Java EE DAS              started

Starting Interstage Java EE Node Agent service.
ijnastart: INFO: ijnal0000: Interstage Java EE Node Agent service has started.
Name                               Status
-----
Interstage Java EE Node Agent      started

Starting Interstage Management Console.
UX:ismngconsolestart: INFO: is40041: The service has been activated normally.

Starting Interstage HTTP Server.
UX:IHS: INFO: ihs01000: The command terminated normally.

Starting PostgreSQL.
su - bdcep_postgres -c "/opt/FJSVcep/postgres/packages/FJSVpgs83/bin/pg_ctl -D /var/opt/FJSVcep/postgres/data -w start"
waiting for server
to start...LOG:  database system was shut down at 2011-11-21 15:03:21 JST
LOG:  database system is ready to accept connections
LOG:  autovacuum launcher started
done
server started

Starting Apache Tomcat.
/etc/init.d/FJSVcep-rest start
Starting Tomcat:                               [ OK ]

Command cepstartserv executed successfully.
```

Example of output at abnormal end

When the Interstage Java EE DAS service cannot be started:

```
# cepstartserv <ENTER>
Starting Interstage Java EE DAS service.
```



```
ERROR: ijdas10002: Interstage Java EE DAS service cannot be started.  
Command cepstartserv execution failed.
```

8.11 cepstopeng

Name

cepstopeng - CEP engine stop

Format

```
cepstopeng [-e engineName]
```

Function description

This command stops a CEP engine.

Arguments

-e *engineName*

Specify the name of the CEP engine to be stopped.

If only one CEP engine has been configured, this option can be omitted.

If this option is omitted, and if there are two or more configured CEP engines, an error message will be output and the stop will fail.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end

Note

This command monitors the stop of the CEP engine until it has completed, and the command ends normally if it can confirm that the stop has completed.

If the stop has not completed within a predetermined time (within 180 seconds), an error message will be output and the stop will fail with a timeout.

Example

Example of output at normal end

When a CEP engine (CepEngine1) is stopped:

```
$ cepstopeng -e CepEngine1 <ENTER>  
Command cepstopeng executed successfully.
```

Example of output at abnormal end

When stop processing has not completed within a predetermined time (when a timeout has occurred):

```
$ cepstopeng -e CepEngine2 <ENTER>
ERROR: cep30203e: Failed to stop Collection Engine. EngineId=(CepEngine2) JobId=(CepEngine2
-cepjobnet-20120710130118-CepEngine2) Reason=(A timeout occurred.(localhost_FRT))
Command cepstopeng execution failed.
```

Please note that we added a newline in the example above (line 2) for readability only - the actual message does not have a newline.

The following information will be added to the error message in this example:

Engineld

Name of the CEP engine to be stopped

Jobld

The Job ID, in the following format:

```
CEPengineName-cepjobnet-yyyyMMddHHmmss(time of CEP engine start)-CEPengineName
```

Reason

Error details (reason for the failure to stop the CEP engine)

.....

8.12 cepstopserv

Name

cepstopserv - CEP service stop

Format

```
cepstopserv
```

Function description

This command stops the CEP service - it stops in batch the multiple services that make up the CEP service entity.

Stop processing for those services that are already stopped will be skipped.

This command also stops a CEP engine that has been started.

The services stopped by this command are as follows:

- CEP engine
- Apache Tomcat
- PostgreSQL
- Interstage HTTP Server
- Interstage Management Console
- Interstage Java EE Node Agent service
- Interstage Java EE DAS service
- Interstage

Arguments

None

End status

The following status codes are returned:

0

Normal end

8

Abnormal end

Point

Executing this command will stop all CEP engines that are running.

Note

- If the stop of one or more of the services fails, an error message will be output and the stop of the CEP service will fail. To complete the stop of the CEP service, remove the cause of the error and then execute the command again.
- This product uses cron for CEP engine resource usage acquisition. This command does not stop cron (refer to "[6.3.3 Applying Updates](#)" for information on how to stop cron).
- If this command is executed twice in a row, the message below is output, but the CEP engine or Apache Tomcat is stopped, so no action is required:

```
Stopping Engines.
asadmin: ERROR: ISJEE_OM2997: Unable to connect to admin-server at given host:
[localhost] and port: [12001]. Please check if this server is up and running and
that the host and port provided are correct.
asadmin: ERROR: ISJEE_CL1137: Command list-clusters failed.

Stopping Apache Tomcat.
/etc/init.d/FJSVcep-rest stop
Shutting down Tomcat:                               [Fail]
```

Please note that we added newlines in the example above (line 2 and 3) for readability only - the actual message does not have a newline.

Example

Example of output at normal end

When all services stop successfully:

```
# cepstopserv <ENTER>
Stopping Engines.
CEP Engine is not started. EngineId=(CepEngine)
Command cepstopeng executed successfully.

Stopping Apache Tomcat.
/etc/init.d/FJSVcep-rest stop
Shutting down Tomcat:                               [ OK ]

Stopping PostgreSQL.
```

```

su - bdcep_postgres -c "/opt/FJSVcep/postgres/packages/FJSVpgs83/bin/pg_ctl -D /var/opt/FJSVcep/postgres/data stop"
waiting for server to shut down.... done
server stopped

Stopping Interstage HTTP Server.
UX:IHS: INFO: ihs01000: The command terminated normally.

Stopping Interstage Management Console.
UX:ismngconsolestop: INFO: is40042: The service has been terminated normally.

Stopping Interstage Java EE Node Agent service.
ijnastop: INFO: ijnal0001: Interstage Java EE Node Agent service has stopped.
Name                               Status
-----
Interstage Java EE Node Agent      stopped

Stopping Interstage Java EE DAS service.
ijdasstop: INFO: ijdass10001: Interstage Java EE DAS service has stopped.
Name                               Status
-----
Interstage Java EE DAS              stopped

Stopping Interstage.
UX:isstop: INFO: is30160:INTERSTAGE has stopped normally.

Command cepstopserv executed successfully.

```

Example of output at abnormal end

When the Interstage Java EE DAS service cannot be stopped.

```

# cepstopserv <ENTER>
(...)
Stopping Interstage Java EE DAS service.
ERROR: ijdass10003: Interstage Java EE DAS service cannot be stopped.
Command cepstopserv execution failed.

```

8.13 cepundeployrsc

Name

cepundeployrsc - Development asset undeployment

Format

```
cepundeployrsc resource [-e engineName] -n resourceID
```

Function description

This command undeploys a development asset that has been deployed.

Specify the type of development asset (event type definition, rule definition, master definition, or SOAP listener definition) to be undeployed.

When the command is executed, a prompt is displayed - type "y" and press the Enter key, or simply press the Enter key, to execute the undeployment. Execution can be canceled by typing "n" or "q" and pressing the Enter key.

Example

Example of query (for an event type definition):

```
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:
```

Arguments

resource

Specify the type of development asset.

eventtype

Event type definition

rule

Rule definition

master

Master definition

listener

SOAP listener definition

-e engineName

Specify the name of the CEP engine to which the development asset to be undeployed has been deployed.

If only one CEP engine has been configured, this option can be omitted.

If this option is omitted, and if there are two or more configured CEP engines, an error message will be output and the undeployment will fail.

-n resourceID

Specify the development asset ID to be undeployed.

End status

The following status codes are returned:

0

Normal end

8

Abnormal end

Note

If the CEP engine that is the undeployment target is running, undeployment cannot be performed.



Example

Example of output at normal end

When an event type definition that has been deployed to a CEP engine (CepEngine1) is undeployed:

```
$ cepundeployrsc eventtype -e CepEngine1 -n EVENT_01 <ENTER>
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:y <ENTER>
Command cepundeployrsc executed successfully.
```

Example of output at abnormal end

When an attempt is made to undeploy an event type definition that has been deployed to a CEP engine (CepEngine2) that is running:

```
$ cepundeployrsc eventtype -e CepEngine2 -n EVENT_02 <ENTER>
Are you sure you want to undeploy the event type definition?(default: y) [y,n,q]:y <ENTER>
Because the definition is used, it cannot be deleted. Definition=(instream) Id=(EVENT_02-CepEngine2)
Command cepundeployrsc execution failed.
```

The following information will be added to the error message in this example:

Id

Development asset ID and target CEP engine name to be undeployed.

Chapter 9 Definition File Reference

This chapter explains the definition files that the Interstage Big Data Complex Event Processing Server uses.

9.1 Defining a CEP Engine

This section explains the definition file for configuring a CEP engine (engine configuration file).

An engine configuration file for use by the initial CEP engine (/etc/opt/FJSVcep/cep/sample_eng.xml) is created automatically at installation. To change the configuration details, make a copy of this file and then edit the copy.

Creating, deleting or changing the setting of CEP engine is done by executing cepconfigeng with specified engine configuration file.

Point

- The engine configuration file is a text file in XML format.
- The items to be set are specified as XML element and attribute values.
- The character encoding for the engine configuration file is UTF-8.

Note

Creating, deleting, or changing the settings of a CEP engine can also be done by editing the engine configuration file of the initial CEP engine, but to prevent configuration details being lost due to an editing mistake, do not use this method.

9.1.1 Engine Configuration File

This section explains the settings of the engine configuration file.

The engine configuration file is an XML file that has "subSystemConfig" as the root element. The configuration information of one CEP engine is described in an "engineConfig" element. When multiple CEP engines are being created, the configurations of all the CEP engines must be described in one engine configuration file (using consecutive "engineConfig" elements).

Also, if Logging is to be used in the input adapter, describe the "logging" subelement in the "engineConfig" element.

The items to be set in the engine configuration file are shown below:

Element or attribute	Item name	Explanation	Allowed values	Mandatory/ optional
id (attribute of the "engineConfig" element)	CEP engine name	ID that uniquely identifies the CEP engine.	Up to 20 alphanumeric characters, including underscores (_).	Mandatory
type (*1)	Logging type	Log destination. bdpp : Output to a Hadoop system. file : Output to the engine log.	See Explanation.	(*2)
directory (*1)	Directory name	Directory of the Hadoop system that is the log destination. Specify one of the following: - Name of the directory (not its full path). - "/".	Up to 1023 alphanumeric characters, including forward slashes (/). Use only alphanumeric	(*3)

Element or attribute	Item name	Explanation	Allowed values	Mandatory/optional
			characters for subdirectory names	
loggingMaxOpenFile (*1)	Number of open log files	Number of files to be simultaneously opened for the Hadoop system. The default is 6.	1 to 122.	(*4)
loggingRotationCycle (*1)	Logging cycle time	Time from when the event log files are opened until those files are renamed as files that can be analyzed (using the ".done" extension). The default is 300 (seconds).	1 to 2592000. (unit: seconds)	(*4)
socketAdapterPort	Socket adapter port	Reception port number to be used by the socket adapter.	1 to 65535.	(*5)

*1: Subelement of the logging element

*2: Mandatory if logging is to be used

*3: Mandatory if logging is to be used and the logging type is "bdpp"

*4: Optional if logging is to be used and the logging type is "bdpp"

*5: Specified if socket communication is to be used by the input adapter



Point

- The log destination (path) in a Hadoop system is specified in the event type definition. Therefore, if "bdpp" is specified in the logging type, there is no need to specify the log destination (path).
- If "bdpp" is to be specified in the log destination, the Interstage Big Data Parallel Processing Server (hereafter, referred to as "BDPP") must be set up beforehand (refer to "4.4.2 Setup of Hadoop Collaboration" for details).



Note

- Up to five CEP engines ("engineConfig" elements) can be described in the engine configuration file.
- Specify the CEP engine name so as not to be confused with other CEP engine names.
- Specify a port number that is not being used by the system as the socket adapter port.



Example

When two CEP engines (CepEngine1 and CepEngine2) are configured

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<subSystemConfig xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1">
  <engineConfig id="CepEngine1">
    <logging>
      <type>bdpp</type>
      <directory>hadoop</directory>
      <loggingMaxOpenFile>6</loggingMaxOpenFile>
      <loggingRotationCycle>300</loggingRotationCycle>
    </logging>
  </engineConfig>
  <engineConfig id="CepEngine2">
    <logging>
      <type>bdpp</type>
      <directory>hadoop</directory>
      <loggingMaxOpenFile>6</loggingMaxOpenFile>
      <loggingRotationCycle>300</loggingRotationCycle>
    </logging>
  </engineConfig>
</subSystemConfig>
```



```

        <socketAdapterPort>9600</socketAdapterPort>
    </engineConfig>
    <engineConfig id="CepEngine2">
        <logging>
            <type>file</type>
        </logging>
        <socketAdapterPort>9601</socketAdapterPort>
    </engineConfig>
</subSystemConfig>

```

In this example, the following settings are used to configure the CEP engines:

CepEngine1

Uses logging and records events in a Hadoop system (Directory name: `hadoop`). The number of logging files that can be opened is 6, and the logging cycle time is 300 seconds.

Uses the socket adapter listening at port 9600.

CepEngine2

Uses logging and records events in the engine log.

Uses the socket adapter listening at port 9601.

9.2 Defining Development Assets

This section explains the definition files (development assets) for using the various CEP service features.

There are four kinds of development assets, as shown below, and each has an identifier known as a development asset ID, which must be unique per definition type.

[cepdeployrsc](#) deploys a development asset to a CEP engine, and [cepundeployrsc](#) undeploys a development asset.

Definition type	Definition overview	Used by feature				Mandatory/optional
		Input adapter	High-speed filter	Complex event processing	Output adapter	
Event type definition	Defines the data structure of the input event data. Also sets Logging in the input adapter.	Used	Used	Used		(*1)
Rule definition	Describes high-speed filter rules and complex event processing rules.		Used	Used	Used	(*2)
Master definition	Sets the master data used by the High-speed Filter.		Used			(*3)
SOAP listener definition	Sets the send destination information when SOAP communication is to be used to notify user-developed Web service events.				Used	(*4)

*1: At least one development asset must be deployed.

*2: Deploy if the high-speed filter or complex event processing is to be used.

*3: Deploy if master data is to be referenced in the high-speed filter.

*4: Deploy if the SOAP listener is to be used in the output adapter.



- Each definition file is a text file in XML format (except for master data).
- The items to be set are specified as XML element and attribute values.
- The character encoding for each definition file is UTF-8.

9.2.1 Event Type Definition File

This section explains the setup file for describing the event type definition (event type definition file).

The event type definition file is an XML file that has "eventType" as the root element. An event type definition file is created for each event type. The items to be described will vary according to the format of the event data (CSV or XML format).

The items to be set in the event type definition file are shown below:

Element or attribute	Item name	Explanation	Allowed values	Mandatory/optional
id (attribute of "eventType")	Development asset ID	ID unique in the deployment CEP engine. Used as the event type name specified in high-speed filter rules and as the event stream name specified in complex event processing rules.	Up to 39 alphanumeric characters, underscores (_), or hyphens (-). Note: The first character must be a letter.	Mandatory
comment	Comment	Explanation of this definition.	Up to 1,000 characters.	Optional
type	Format	Format of the event data. XML: Events are in XML format. CSV: Events are in CSV format.	See Explanation.	Mandatory
xmlSchema	XML schema	XML schema that represents the structure of the event data. This is described in XML schema language. Therefore, escape must be performed in a way that cannot be interpreted as a markup specification. For simple description, describe using a CDATA section instead of escaping.	Up to 1,048,576 characters.	(*1)
root	Root element	Name of the root element of the event. The root element is one of the elements defined in the XML schema. Specify unique root element names for all event type definitions.	Up to 512 characters.	(*1)

Element or attribute	Item name	Explanation	Allowed values	Mandatory/ optional
csvColumn	CSV column information	Column information that represents the structure of the event data. Describe at least one "column" element (as described below).	Follows the column element (as described below). There is no limit to the number of column elements.	(*2)
useLogging	Whether to use logging	Whether to use logging for pre-processing events of the high-speed filter received from the input adapter. true: Use. false: Do not use.	See Explanation.	Mandatory
loggingTableName	Log storage area	Absolute path of the log storage area for storing events. This is used to store events received by the input adapter. Even if logging an event in an engine log (when the engine structure file type element is set as file), for event identification set the virtual path name (for example, <i>"/eventTypeID"</i>).	Up to 255 characters.	(*3)
useCep	Whether to use complex event processing	Controls whether to use complex event processing. true: Use. false: Do not use.	See Explanation.	Mandatory

*1: Mandatory if "XML" is specified as the format

*2: Mandatory if "CSV" is specified as the format

*3: Mandatory if logging is to be used

If "CSV" is specified as the format, describe the CSV column information shown below in the "csvColumn" element:

Element or attribute	Item name	Explanation	Allowed values	Mandatory/ optional
column	column	Describe as many of these as the number of event (CSV format) columns. The column elements must be specified in the same order as in the event data.	(empty element)	Mandatory
name (attribute of "column")	Item name	Item name of the CSV column. Used as the item name specified in high-speed filter rules and as the property name specified in complex event processing rules.	Refer to "9.4 Characters Allowed in Item, Tag and Attribute Names" .	Mandatory
type	Item type	Data type of the CSV column.	See Explanation.	Mandatory

Element or attribute	Item name	Explanation	Allowed values	Mandatory/ optional
(attribute of "column")		string : String. boolean : Boolean value (true/false). byte : 8-bit signed integer. int : 32-bit signed integer. long : 64-bit signed integer. float : 32-bit float. double : 64-bit double precision float.		

The data types of CSV columns are converted to the appropriate type for complex event processing rules, as shown below:

Data type of CSV column	Data type in complex event processing rules
string	string
boolean	bool/boolean
byte	byte
int	int/integer
long	long
float	float
double	double

Point

If the high-speed filter is being used, the event types of the events passed to complex event processing and of the input events will vary, unless only extraction (filtering) is being used. In this case, a separate event type definition for complex event processing must be deployed.

Note

- Up to 32 event type definitions can be deployed to one CEP engine.
- Regardless of the CEP engine to deploy event type definitions, set a unique root element name for all event type definitions.
- Regarding the useCep attribute (which determines whether to use complex event processing):
Set "false" if only logging is to be used by the input adapter. An investigation of the content of input event data is considered prior to starting operation.
- The column elements must be specified in the same order as in the event data.
- When using cepgetrsc to display details of the development asset, the displayed XML schema uses the same format as used by the CDATA section.

Example

When XML format is selected

Target event data (XML):

```
<?xml version="1.0" encoding="UTF-8"?><messagedata xmlns="http://dataaccesscontrol.sspf.fujitsu.com/namespace/xmlmessage"><memberID>MEM0001</memberID><areaID>1010</areaID><status>1</status></messagedata>
```

Please note that we added newlines in the example above (line 1 and 2) for readability only - the actual data does not have a newline.

Event type definitions:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<eventType xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="EVENTTYPE_01">
  <comment>Event type definition_01</comment>
  <type>XML</type>
  <xmlSchema>
    <![CDATA[
      <xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
        xmlns="http://dataaccesscontrol.sspf.fujitsu.com/namespace/xmlmessage"
        targetNamespace="http://dataaccesscontrol.sspf.fujitsu.com/namespace/xmlmessage">
        <xs:element name="messagedata">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="memberID" type="xs:string" />
              <xs:element name="areaID" type="xs:string" />
              <xs:element name="status" type="xs:string" />
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:schema>
    ]]>
  </xmlSchema>
  <root>messagedata</root>
  <useLogging>true</useLogging>
  <loggingTableName>/echonet</loggingTableName>
  <useCep>true</useCep>
</eventType>
```

When CSV format is selected

Target event data (CSV):

```
"MEM0001", "1010", "1"
```

Event type definitions:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<eventType xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="EVENTTYPE_02">
  <comment>Event type definition_02</comment>
  <type>CSV</type>
  <xmlSchema></xmlSchema>
  <csvColumn>
    <column name="memberID" type="string" />
    <column name="areaID" type="string" />
    <column name="status" type="string" />
  </csvColumn>
  <useLogging>false</useLogging>
  <useCep>true</useCep>
</eventType>
```

In this example, the event types below are defined.

EVENTTYPE_01

Using the input adapter, the event data is output to a log in "/echonet".

This event data is passed to complex event processing (only if it is not being filtered by the high-speed filter).

EVENTTYPE_02

In the input adapter, the event data is not output to a log.

This event data is passed to complex event processing (only if it is not being filtered by the high-speed filter).

9.2.2 Rule Definition File

This section explains the setup file for describing a rule definition (rule definition file).

The rule definition file is an XML file that has "rule" as the root element. The rule definition file describes the rules for the high-speed filter and complex event processing.

The items to be set in the rule definition file are shown below:

Element or attribute	Item name	Explanation	Allowed values	Mandatory/optional
id (attribute of "rule")	Development asset ID	ID unique in the deployment CEP engine.	Up to 39 alphanumeric characters, underscores (_), or hyphens (-). Note: The first character must be a letter.	Mandatory
comment	Comment	Explanation of this definition.	Up to 1,000 characters.	Optional
filter	Filter rule	Description of a high-speed filter rule. Use a method such as a CDATA section if any part of the description can be interpreted as a markup specification.	Up to 1,048,576 characters.	(*1)
statements	Complex event processing rule	Description of a complex event processing rule. Use a method such as a CDATA section if any part of the description can be interpreted as a markup specification.	Up to 1,048,576 characters.	(*2)

*1: Mandatory if a high-speed filter rule is to be described.

*2: Mandatory if a complex event processing rule is to be described.

Point

- Filter rule options (such as "SkipChar", which specifies strings to be excluded as search targets), must be described at the beginning of the filter rule (refer to Chapter 2, "Filter Rule Language Reference" in the *Developer's Reference* for information on filter rule options).

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<rule xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="RULE_02">
  <comment>Rule definition_02</comment>
  <filter>
    <![CDATA[
      @SkipChar("\n")
      @SeparateChar("\t")
      @ANKmix(true)
      @KNJmix(true)

      on EVENTTYPE_02 {
        (...)
      }
    ]]>
  </filter>
</rule>
```

- A huge number of rules can be described in a single rule definition (as long as the number of characters is not exceeded), but when rules become complex, this can negatively affect their ability to be maintained and referenced. When creating rule definitions, consider creating separate rule definitions for the high-speed filter and for complex event processing, and dividing rule definitions into smaller ones that are meaningful and cohesive.

Note

- Up to 32 rule definitions can be deployed to one CEP engine.
- When using `cepgetrsc` to display details of the development asset, the displayed filter rules and complex event processing rules use the same format as used by the CDATA section.

See

- Refer to Chapter 2, "Filter Rule Language Reference" in the *Developer's Reference* for information on how to describe filter rules.
- Refer to Chapter 1, "Complex Event Processing Language Reference" in the *Developer's Reference* for information on how to describe complex event processing rules.

Example

Description example

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<rule xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="RULE_01">
  <comment>Rule definition_01</comment>
  <filter>
    <![CDATA[
      on EVENTTYPE_01 {
        if ($status == 'Walking') then output() as EVENTTYPE_01;
      }
    ]]>
  </filter>
  <statements>
```

```

<![CDATA[
  @SoapListener('LISTEN_01')
  @DebugLogListener
  select * from EVENTTYPE_01 where areaID = '1010';
]]>
</statements>
</rule>

```

In this example, the rules shown below are described.

High-speed filter rule

If the status item contents (string) extracted from the "EVENTTYPE_01" event type input event is "Walking", it is transferred to Complex Event Processing.

Complex event processing rule

This rule notifies the SOAP listener of event data with the "EVENTTYPE_01" event type, and simultaneously outputs debug information to the engine log.



9.2.3 Master Definition File

This section explains the setup file for describing the master definition (master definition file).

The master definition file is an XML file that has "master" as the root element. A master definition file is created for each unit of master data. The master data consists of a schema file and at least one data file, and is matched to a master definition file.

The master data is read when the CEP engine to which the master definition is deployed starts.

The items to be set in the master definition file are shown below:

Element or attribute	Item name	Explanation	Allowed values	Mandatory/optional
id (attribute of "master")	Development asset ID	ID unique in the deployment CEP engine.	Up to 39 alphanumeric characters, underscores (_), or hyphens (-). Note: The first character must be a letter.	Mandatory
comment	Comment	Explanation of this definition.	Up to 1,000 characters.	Optional
schemaFile	Schema file	Full name the schema file.	Up to 1,023 bytes.	Mandatory
dataFile	Data file	Full name of the data file. If there are multiple data files, specify multiple "dataFile" elements. Alternatively, using an absolute path, specify the directory where the data file is located. In this case, all files in the directory will be read as data files.	Up to 1,023 bytes. Note: If specifying a directory, set this so that the directory name plus the file name does not exceed 1,023 bytes.	Mandatory
skipHeader	Skip specification	Whether to skip the first line in the data file when the item name is	See Explanation.	Optional

Element or attribute	Item name	Explanation	Allowed values	Mandatory/ optional
		<p>described in the first line of the data file.</p> <p>true: Assume the first line is not data (ignore).</p> <p>false: Assume the first line is data (default).</p>		

The schema file and data file are text files in CSV format (refer to "9.5 CSV Format Supported" for details). The character encoding for the files is UTF-8.

Schema file

In a schema definition file, one or more item names are described using one record.

Refer to "9.4 Characters Allowed in Item, Tag and Attribute Names" for information on the characters that can be specified in item names.

Data file

In a data file, one line corresponds to one entry of master data.

In each line, describe values (data) that correspond to the schema file items.

If the number of data items is greater than the number of schema information items, the data will be regarded invalid.

Note

- Up to 32 master definitions can be deployed to one CEP engine.
- Up to 100 "dataFile" elements can be specified in one master definition. If this limit is exceeded, locate the data files in any directory and then specify the path of the directory.
- If the number of data file items is less than the number of schema file items, empty strings ("") will be used in the missing items. If the number of data file items is greater, the CEP engine will fail to start.
- If data files are specified using a directory path, ordinary files and symbolic links (with link destinations of ordinary files) located in the directory are read as data files. Aside from subdirectories, do not locate files other than ordinary files and symbolic links (such as named pipes and special files) in the directory.
- Master data is opened in memory when the CEP engine starts. Therefore, memory is consumed in proportion to the master data size.
- If master data is being updated, the CEP engine must be restarted (stopped and then started) in order to reflect the update.
- Read permissions for the engine execution user must be given in the schema file and data files.

Example

- Example of a schema file

/var/tmp/SchemaFile01.csv

```
"Kbn", "Number", "Code", "Name", "Value", "Total", "Biko"
```

- Example of a data file

/var/tmp/MasterFile01.csv

```
"01","1001","AAA","BlockA","1,000","1,000","Comment: Memo number 4023"
"02","1001","BBB","BlockB","","1,200","Comment: Memo number 4023"
"03","1002","CCC","BlockC","800","800","Comment: Memo number 4023"
```

- Example of a master definition file

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<master xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="MASTER_01">
  <comment>Master definition_01</comment>
  <schemaFile>/var/tmp/SchemaFile01.csv</schemaFile>
  <dataFile>/var/tmp/MasterFile01.csv</dataFile>
  <skipHeader>>false</skipHeader>
</master>
```

9.2.4 SOAP Listener Definition File

This section explains the setup file for describing the SOAP listener definition (SOAP listener definition file).

The SOAP listener definition file is an XML file that has "soapListener" as the root element. A SOAP listener definition file is created for each user-developed Web service.

The items to be set in the SOAP listener definition file are shown below:

Element or attribute	Item name	Explanation	Allowed values	Mandatory/ optional
id (attribute of "soapListener")	Development asset ID	ID unique in the deployment CEP engine Used as the output adapter (SOAP listener) in complex event processing rules.	Up to 39 alphanumeric characters, underscores (_), or hyphens (-). Note: The first character must be a letter.	Mandatory
comment	Comment	Explanation of this definition.	Up to 1,000 characters.	Optional
url	Connection URL	Connection URL for the user-developed Web service.	Up to 512 alphanumeric characters and symbols.	Mandatory
nameSpace	Namespace	Namespace of the body of the SOAP message (SOAP body) to be sent to the user-developed Web service.	Up to 512 characters.	Mandatory
prefix	Namespace prefix	Namespace prefix to be used in the body of the SOAP message (SOAP body). When omitted, "ns" is used.	Up to 20 characters.	Optional
method	Root element	Root element name of the body of the SOAP message (SOAP body).	Up to 512 characters.	Mandatory



Note

Up to 32 SOAP listener definitions can be deployed to one CEP engine.



Example

Description example

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<soapListener xmlns="urn:xmlns-fujitsu-com:cspf:bdcep:v1" id="LISTEN_01">
  <comment>SOAP listener definition_01</comment>
  <url>http://192.168.11.249/WebServWAR/MyApp1Service</url>
  <nameSpace>http://webservice/</nameSpace>
  <prefix>ns</prefix>
  <method>cep</method>
</soapListener>
```

This example defines the notification of a message (event) that includes a SOAP body specifying that the root element is "cep" in the user-developed Web service with the connection destination URL "http://192.168.11.249/WebServWAR/MyApp1Service".

9.3 Setting up for Installation

This section explains the setup file to be used at unattended installation of this product (installation file).

9.3.1 Installation File

This section explains the format of and parameters that can be set in the installation file.

Format

Below is a definition example of the installation file - it consists of the line types below:

Parameter definition line

A parameter definition line has the *parameterName=specifiedValue* format, where the value on the right side is assigned to the parameter on the left side. Tabs and spaces at the beginning and end of the line or around the equals sign (=) are ignored.

Blank line

A blank line consists of spaces and tabs only. The entire line is ignored by the Installer.



Example

Below is a definition example of the installation file.

```
BDCEP_USER_NAME=isbdcep
BDCEP_GROUP_NAME=isbdcep
BDCEP_INITIAL_ENGINE_NAME=CepEngine
```

List of parameters that can be set

A list of the parameters that can be set is shown below (all of them are mandatory):

Parameter name	Item name	Explanation	Allowed values	Mandatory/optional
BDCEP_USER_NAME	Engine execution user name	User name of the engine execution user.	Registered user name, up to 8 alphanumeric characters.	Mandatory
BDCEP_GROUP_NAME	Group name	Group name of the group to which the engine execution user belongs.	Registered group name, up to 8 alphanumeric characters.	Mandatory
BDCEP_INITIAL_ENGINE_NAME	Initial CEP engine name	Name of the CEP engine to be created at initial setup.	Up to 20 alphanumeric characters.	Mandatory



Note

Before executing unattended installation, ensure that the user and group specified in the installation file are registered in the system.

9.4 Characters Allowed in Item, Tag and Attribute Names

This section explains the characters that can be specified in item names, tag names, and attribute names.

9.4.1 For High-Speed Filter Rules and Master Definitions

For high-speed filter rules and master definitions, the characters that can be specified in CSV column item names, XML tag names, and attribute names are the single-byte characters shown below as well as multi-byte characters:

!	-	.	0	1	2	3	4	5	6	7	8	9	:
@	A	B	C	D	E	F	G	H	I	J	K	L	M
N	O	P	Q	R	S	T	U	V	W	X	Y	Z	_
^	a	b	c	d	e	f	g	h	i	j	k	l	m
n	o	p	q	r	s	t	u	v	w	x	y	z	



Note

- A hyphen (-) or period (.) cannot be specified at the beginning of a line.
- A space () cannot be specified (the bottom rightmost cell of the table above does not signify a space).

9.4.2 For Complex Event Processing Rules

For complex event processing rules, the characters that can be specified in CSV column item names, XML tag names, and attribute names have the following constraints:

- The first character must be a letter (a to z or A to Z).
- The second and subsequent characters must be letters (a to z or A to Z), digits (0 to 9), or underscores (_).

- The reserved words of complex event processing rules cannot be used, regardless of case.

Refer to Chapter 1, "Complex Event Processing Language Reference" in the *Developer's Reference* for information on the reserved words of complex event processing rules.

Please note that the rules are case-sensitive.

9.5 CSV Format Supported

This section explains the CSV format supported.

Example

CSV description example

```
"01", "1001", "AAA", "BlockA", "1,000", "1,000", "Comment: Memo number 4023"
"02", "1001", "BBB", "BlockB", "", "1,200", "Comment: Memo number 4023"
"03", "1002", "CCC", "BlockC", "800", "800", "Comment: Memo number 4023"
```

- Use commas (,) to delimit the items in the columns in one record.
- Records are separated by newline characters.
- Items can be enclosed by double quote marks (" "). If a comma or newline needs to be specified in an item, enclose the item in double quote marks.
- If a double quote mark (") needs to be specified in the content of an item, enclose the item in double quote marks and then add a double quote mark before the one that is part of the item (that is, specify two consecutive double quote marks).

Example

When items 'A, AA', 'BnewLineBB', and 'C"CC' are included in a record

```
"A, AA", "B
BB", "C"CC"
```

- When data is enclosed by double quote marks, an error will occur if a character other than a comma (,) is used between items. Delete any excess characters around delimiting commas (,).

Example

Incorrect: There are unnecessary spaces (underlined> between items (around the comma).

```
"AAA", "BlockA"
```

Correct: There are no unnecessary spaces between items.

```
"AAA", "BlockA"
```

- The maximum size of one record is 32 MB (megabytes).
- For input events, records do not need to be separated by newlines. Multiple records cannot be handled in one input event.

Glossary

Apache Hadoop

Open-source [Hadoop](#) software developed by the Apache Software Foundation (ASF).

cache access application for update

Application for storing initial data and updating data in a cache in an [XTP Server](#).

CEP (Complex Event Processing)

Refer to "[complex event processing](#)".

CEP engine

Refer to "[high-performance CEP engine](#)".

CEP Server

The server on which this product runs.

CEP service

Service that runs on the [CEP Server](#) and manages the [CEP engine](#).

complex event processing

Technique of rapidly analyzing, assessing and processing in real time the content and status of massive volumes of [event data](#) being sent continuously.

It is used in the [complex event processing](#) feature provided by this product.

complex event processing language

Language for specifying the rules of [complex event processing](#).

It allows the use of SQL format.

complex event processing rule

Rule described using the [complex event processing language](#), consisting of at least one [complex event processing statement](#).

complex event processing statement

Structural unit of rules specified using the [complex event processing language](#).

A complex event processing statement describes processing using statements such as the SELECT statement and the CREATE WINDOW statement.

debug listener

Feature that outputs debug messages to the [engine log](#) when a [complex event processing statement](#) is executed.

engine execution user

User of the operating system when a [CEP engine](#) is running.

engine log

File containing the output of [CEP engine](#) system messages and debug messages defined in the rules.

event data

Time-series data, such as [sensing data](#) or location information about people, and that shows when something occurred.

event log

File containing event data which is recorded by [logging](#).

event log analysis application

Application that analyzes [event logs](#) in the [Hadoop system](#).

event sender application

Application that sends [event data](#) to a [CEP engine](#).

event type definition

Definition information about the format of [event data](#).

The [event data](#) can be specified in XML or CSV format.

filter language

Refer to "[high-speed filter language](#)".

filter rule

Refer to "[high-speed filter rule](#)".

filter statement

Refer to "[high-speed filter statement](#)".

Hadoop

Technique for effectively performing distributed and parallel processing of data accumulated in massive volumes. It basically consists of the [HDFS \(Hadoop Distributed File System\)](#) and the MapReduce which is known as a parallel distributed processing technique.

Refer to "[Apache Hadoop](#)".

Hadoop collaboration

In [logging](#) context, a feature for recording [event data](#) from the [CEP Server](#) directly in the [Hadoop system](#).

Hadoop system

System on which [Hadoop](#) is running.

HDFS (Hadoop Distributed File System)

Distributed file system used by [Hadoop](#).

For the placement of big data, HDFS creates multiple replicas of data blocks and distributes them on nodes called DataNodes, which are then managed by nodes called NameNodes.

high-performance CEP engine

The processing unit in this product.

It allows massive volumes of [event data](#) to be processed rapidly by combining the [high-speed filter](#) with conventional [complex event processing](#).

high-speed filter

Feature that uses a technique unique to Fujitsu to allow high-speed filtering of massive volumes of [event data](#) while matching the [event data](#) with the [master data](#).

high-speed filter language

Language for specifying the rules in the [high-speed filter](#).

It allows the use of IF-THEN format.

high-speed filter rule

Rule described using the [high-speed filter language](#), consisting of at least one [high-speed filter statement](#).

high-speed filter statement

Structural unit of rules specified using the [high-speed filter language](#).

A high-speed filter statement is described for each [event type definition](#).

HTTP adapter

[input adapter](#) that allows [event data](#) to be received using HTTP communication.

It allows more lightweight communication when compared to a [SOAP adapter](#).

initial CEP engine

[CEP engine](#) created at installation.

input adapter

Feature that receives [event data](#) from outside the [CEP Server](#).

logging

Feature that records [event data](#) in a [Hadoop system](#) or in the [engine log](#).

logging listener

Feature that [logging](#) the processing results of a [complex event processing statement](#).

master data

Data used by this product to match [event data](#) in [high-speed filter](#) processing.

It is deployed in the [CEP Server](#) in CSV format, and consists of schema files (where the item names for each column are defined), and data files (where the data rows are stored).

master definition

Definition information of the [master data](#) to be referenced during [high-speed filter](#) processing.

output adapter

Feature that sends the processing results of [complex event processing](#) outside the [CEP Server](#).

resource log

Log containing the resource usage of a [CEP engine](#).

It is stored in CSV format, so it can be analyzed using a tool such as Excel.

rule definition

Processing details for the [high-speed filter](#) ([high-speed filter rule](#)) and [complex event processing](#) ([complex event processing rule](#)) features.

sensing data

Data sent from a variety of sensors.

Sensing data is a type of [event data](#).

SOAP (Simple Object Access Protocol)

Lightweight protocol, which is subject to ongoing standardization work by the World Wide Web Consortium (W3C).

For its communication infrastructure, SOAP uses the Internet standards of Hyper Text Transfer Protocol (HTTP) and (eXtensible Markup Language (XML).

SOAP adapter

[input adapter](#) that allows event data to be received using [SOAP](#).

SOAP listener

Feature that uses [SOAP](#) to send the processing results of a [complex event processing statement](#) to an external system.

SOAP listener definition

Defines information required to perform communication using a [SOAP listener](#).

It defines information such as the URL of the external system that is to be the communication target.

social media

Services and applications that provide open communication generated by diverse content (such as text, voice and video) exchanged and shared on the Internet by a variety of people, in contrast to the traditional information media (such as newspapers and television).

socket adapter

[input adapter](#) that performs communication using a format unique to this product.

Socket adapters are more suitable for receiving massive volumes of event data than a [SOAP adapter](#) or [HTTP adapter](#).

user-developed Web service

Application that uses [SOAP](#) to receive processing results of a [complex event processing statement](#) sent from a [SOAP listener](#).

WSDL (Web Services Description Language)

Interface description language in Web services, which allows the description of information such as the access point in Web services (URL), the protocol to be used (SOAP, HTTP, or MIME), and the message format (XML Schema).

It is used in the development of a [user-developed Web service](#), which is an application for collaborating with this product.

XTP collaboration

In [complex event processing](#) for this product, a feature for referencing the cache of an [XTP Server](#) outside the [CEP Server](#).

XTP collaboration is used to reference frequently updated data from a [complex event processing rule](#).

XTP Server

An XTP Server node in the Interstage eXtreme Transaction Processing Server.
