

SYBASE®

Frequently Asked Questions (FAQ)

Sybase Aleri Streaming Platform 3.1

DOCUMENT ID: DC01297-01-0311-01

LAST REVISED: June, 2010

Copyright © 2010 Sybase, Inc.

All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

DB2, IBM and Websphere are registered trademarks of International Business Machines Corporation.

Eclipse is a trademark of Eclipse Foundation, Inc.

Excel, Internet Explorer, Microsoft, ODBC, SQL Server, Visual C++, and Windows are trademarks or registered trademarks of Microsoft Corp.

Intel is a registered trademark of Intel Corporation.

JDBC, Solaris, Sun and Sun Microsystems are trademarks or registered trademarks of Sun Microsystems or its subsidiaries in the U.S. and other countries.

Kerberos is a trademark of the Massachusetts Institute of Technology.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Netezza is a registered trademark of Netezza Corporation in the United States and/or other countries.

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

SPARC is a registered trademark of SPARC International, Inc. Products bearing SPARC trademarks are based on an architecture developed by Sun Microsystems, Inc.

Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Group Ltd.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian

agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Table of Contents

Introduction	v
1. Related Documents	v
Frequently Asked Questions	1

Introduction

This document contains answers to commonly asked questions about the Sybase Aleri Streaming Platform and associated topics. For more detailed information on specific subjects, see the following list of Related Documents.

1. Related Documents

This guide is part of a set. The following list briefly describes each document in the set.

<i>Product Overview</i>	Introduces the Aleri Streaming Platform and related Aleri products.
<i>Getting Started - the Aleri Studio</i>	Provides the necessary information to start using the Aleri Studio for defining data models.
<i>Release Bulletin</i>	Describes the features, known issues and limitations of the latest Aleri Streaming Platform release.
<i>Installation Guide</i>	Provides instructions for installing and configuring the Streaming Processor and Aleri Studio, which collectively are called the Aleri Streaming Platform.
<i>Authoring Guide</i>	Provides detailed information about creating a data model in the Aleri Studio. Since this is a comprehensive guide, you should read the <i>Introduction to Data Modeling and the Aleri Studio</i> . first.
<i>Authoring Reference</i>	Provides detailed information about creating a data model for the Aleri Streaming Platform.
<i>Guide to Programming Interfaces</i>	<p>Provides instructions and reference information for developers who want to use Aleri programming interfaces to create their own applications to work with the Aleri Streaming Platform.</p> <p>These interfaces include:</p> <ul style="list-style-type: none">• the Publish/Subscribe (Pub/Sub) Application Programming Interface (API) for Java• the Pub/Sub API for C++• the Pub/Sub API for .NET• a proprietary Command & Control interface• an on-demand SQL query interface
<i>Utilities Guide</i>	Collects usage information (similar to UNIX® man pages) for all Aleri Streaming Platform command line tools.
<i>Administrators Guide</i>	Provides instructions for specific administrative tasks related to the Aleri Streaming Platform.
<i>Introduction to Data Modeling and the Aleri Studio</i>	Walks you through the process of building and testing an Aleri data model using the Aleri Studio.
<i>SPLASH Tutorial</i>	Introduces the SPLASH programming language and illustrates its

capabilities through a series of examples.

Frequently Asked Questions

Answers some frequently asked questions about the Aleri Streaming Platform.

Frequently Asked Questions

1. General

Q: What is the Sybase Aleri Streaming Platform?

A: The Sybase Aleri Streaming Platform is complex event processing (CEP) software that performs real-time processing and analysis of incoming data. With the Sybase Aleri Streaming Platform, you can easily build applications for getting timely insight into current conditions or an immediate response to opportunities and threats. The Sybase Aleri Streaming Platform not only has the power and speed to deliver results in milliseconds or less, but also has the flexibility to work with your current data infrastructure.

Q: What is the Aleri Studio?

A: The Aleri Studio is the authoring environment most users initially choose to create, edit and test data models in the Sybase Aleri Streaming Platform. It is an Integrated Development Environment (IDE) based on the Eclipse™ framework. You can create data models using the Aleri Studio to add processing elements called "streams", and connect the streams to define the data flow.

Q: What can I do with the Sybase Aleri Streaming Platform?

A: Any enterprise that needs to filter and analyze vast amounts of event data in real time and act quickly on the information will benefit from the Sybase Aleri Streaming Platform. It's certainly common to see CEP technology being used in financial institutions and markets, where real-time insight and rapid response is critical, and the rate at which new information arrives can be overwhelming. CEP is also being deployed in the front office, where market data represents events, such as executed trades and quotes that have been published. Other industries where CEP is being deployed include telecommunications, network security, e-commerce and retail.

Q: Who can use the Sybase Aleri Streaming Platform?

A: Almost anyone with a little knowledge of data processing, even Excel users, can modify a data model. More programming background is required to create new models, but even here, a basic knowledge of SQL or relational databases often suffices. The Sybase Aleri Streaming Platform makes it possible to create a multi-threaded, 64-bit application that takes full advantage of the latest hardware without knowledge of threading or other systems programming skills.

Q: What is a data model?

A: It's like a program, but written in a simple dataflow language. The Sybase Aleri Streaming Platform runs a data model that contains the event processing logic that will be applied to incoming event data to produce one or streams. Any of the three Sybase authoring environments (Aleri Studio, Aleri SQL, AleriML) can be used to create and edit a data model.

Q: What is a stream?

A: A stream is the fundamental building block of data models, a place where events are received, processed and stored. There are two types of streams: source streams and derived streams. Source streams receive input data from an external source. Every Sybase data model must have at least one source stream, although models routinely have many more. Derived streams always take their input from other streams defined in the data model. Currently, the Sybase Aleri Streaming Platform has nine different types of derived streams.

Q: What kinds of external data sources are supported?

A: The Sybase Aleri Streaming Platform can get data from many common streaming and static data sources: market data feeds (such as Reuters Market Data System, Activ Financial), message buses (MQ Series and JMS), databases (Oracle, Sybase, Microsoft SQL Server, Postgres, Teradata, Netezza, KDB), files, and sockets. You can even get data from other instances of the Sybase Aleri Streaming Platform. Also, you can get data in popular message formats like FIX. And if the list doesn't include what you need, you can build your own connections through our publish APIs, or contact Sybase to build a new one.

Q: On what operating systems is the Sybase Aleri Streaming Platform supported?

A: The following operating systems are supported:

- Red Hat Enterprise Linux AS4 (version .5 or higher) or AS5; 64-bit
- Fedora Core 4, 5, 6; 64-bit (SMP kernel required on multiprocessor machines)
- Solaris 10 SPARC 3/05 or higher
- Solaris 10 x86 3/05 or higher
- Windows XP
- Windows 2003 Server

Q: How many processors does the Sybase Aleri Streaming Platform require?

A: It can run with just one processor, but can take advantage of as many processors as you have. For Windows and Linux, or Solaris with Intel or AMD processors, the clock speed should be at least 300Mhz. For Solaris with SPARC processors, the Sybase Aleri Streaming Platform requires at least two 64-bit Ultra-SPARC II or III (single-core), although Ultra-SPARC IV or IV+ (dual-core) CPUs are recommended.

Q: How much memory and disk does it require?

A: You'll need at least 512MB of RAM on Linux and Solaris and 1GB on Windows, with at least 1GB of disk space on all operating systems. Beyond that, the requirements depend completely on the data model. Some customer data models keep very little data, even though they receive many events (for example, streaming trades data that keeps only a few recent trades for each symbol). Others store all data, or persist data either in the Sybase Aleri Streaming Platform's persistent disk storage or in an outside data warehouse. Sybase offers connectivity to Sybase IQ, Netezza, and

Teradata for bulk storage, as well as to standard relational databases like Sybase ASE, Microsoft SQL Server, Oracle, and IBM DB2.

Q: How fast can it process messages?

A: That depends on the data model. For very small models, the Sybase Aleri Streaming Platform can process over 1,000,000 messages per second. With more realistic models, the Sybase Aleri Streaming Platform can process over 100,000 incoming messages per second, with latencies in the millisecond range. [The STAC Report](#) [<http://www.sybase.com/files/aleri-dunnington-solaris10-orderbooks-rev1.01.pdf>] gives some further statistics.

Q: Is it multi-threaded?

A: Yes, each stream is implemented as a thread, and threads are used other places as well. You don't need to know the details of threading, though, in order to use the Sybase Aleri Streaming Platform; it does the threading for you.

Q: Can I evaluate the Sybase Aleri Streaming Platform for free?

A: Yes, we offer an evaluation license for 180 days. It's the full product too, so you can get a sense of how well it performs.

Q: How do I obtain a license key?

A: Go to [the Sybase website](http://www.sybase.com) [<http://www.sybase.com>] to download the license key for the 180-day evaluation period, or to purchase the full license for the Sybase Aleri Streaming Platform.

2. Data Models and Data

Q: What kind of data can be processed by the Sybase Aleri Streaming Platform?

A: Events come in the form of records of basic values; think of it like a row in Excel or in a database. For instance, you might have a stream of trades, each record having a ticker symbol (string), a number of shares (integer), and a price (a floating-point number). Events can be marked as update or delete to change or remove previously inserted values.

Q: How do I build a data model?

A: Most people initially use the Aleri Studio to build data models, but you can also define a model in the more advanced AleriSQL or AleriML (which is a variation of XML). You can start by either creating a new data model or opening an existing data model. The best step-by-step explanation is found in the tutorial *Introduction to Data Modeling and the Aleri Studio*.

Q:

How much expertise is needed to build a data model?

A: You can build simple and useful models without any programming background. But, to take full advantage of Sybase's feature-rich environment and power tools, you need basic programming skills.

Q: How easy is it to change a data model?

A: The Aleri Studio makes it easy to change the logic of a data model via its point-and-click graphical user interface. Most of the common properties of streams can be changed by simply opening up a stream shape and typing.

Q: How much data can I keep in the Sybase Aleri Streaming Platform?

A: As much as you can store in memory. The Sybase Aleri Streaming Platform works best if all data fits within physical memory; it does not perform as well when spilling over into virtual memory.

Q: How do I get data in?

A: Through connectors, adapters or clients built using the Sybase publish/subscribe APIs.

Q: How do I get data out?

A: The same way as you get data in: through connectors, adapters, or clients built using the Sybase publish/subscribe APIs.

Q: Can I record data and play it back at a later time?

A: Yes. In the Aleri Studio Playback view, you can record in-flowing data to a playback file and play the captured data back into the Sybase Aleri Streaming Platform. There is also a command-line tool that loads data into the Sybase Aleri Streaming Platform from a variety of sources at a specified rate.

Q: How can I build an application using the Sybase Aleri Streaming Platform?

A: To build an application requires identifying the source of data, constructing a data model to transform the data, and sending the data to relevant clients. If you already have data stored in databases or files, it's easy to use the Aleri Studio to "discover" the shape of the input.

Q: Can I keep my data secure?

A: Yes, in three different ways. You can force clients to identify themselves using one of three authentication mechanisms (PAM, RSA, or Kerberos). You can establish groups so that only certain authenticated clients can read data or control the Sybase Aleri Streaming Platform. And you can

encrypt the data passed between clients and the Sybase Aleri Streaming Platform using SSL.

3. Connectors and Adapters

Q: What is a connector?

A: A connector is a part of the Sybase Aleri Streaming Platform for reading or writing data. The Sybase Aleri Streaming Platform contains connectors for databases, files, sockets, and other data sources. The Aleri Studio has intimate knowledge of connectors as well, and in many cases allows for "discovery" of external data sources which permits the importation of data with a simple drag-and-drop.

Q: What is an adapter?

A: An adapter is an external process that publishes data to, or subscribes to data from, the Sybase Aleri Streaming Platform.

Q: What is the difference between a connector and an adapter?

A: Functionally, connectors and adapters serve the same role. Connectors are tightly integrated into the Aleri Studio and are often run inside the Sybase Aleri Streaming Platform process; adapters are not.

Q: What kinds of adapters are supported?

A: Sybase has built adapters for Reuters Market Data System, Activ Financial, Wombat, FIX and other market data sources. The OpenAdaptor software, extended by Sybase, also counts as an adapter.

Q: What kinds of adapters can you build?

A: Almost any. In fact, we use the same publish/subscribe APIs that are shipped with the Sybase Aleri Streaming Platform to build our adapters.

Q: What is the Aleri Dashboard?

A: It is an optional component that makes it easy to build visually rich real-time displays that can include charts, tables and a wide variety of other display components. The Aleri Dashboard, based on technology from SL Corp., comes with a streaming connector that streams live data into the dashboard.

4. Installation

Q: How do I install the Sybase Aleri Streaming Platform?

A: On Windows, the Sybase Aleri Streaming Platform is bundled into an InstallShield package; it's a simple matter of clicking and following the instructions. On UNIX, there is an installation shell script that performs most of the configuration steps.

Q: How many instances can I run on a single machine?

A: As many as you want. Each instance uses at most four TCP/IP ports, so there is a practical limit of 16,384 instances (because the number of TCP/IP ports is limited to 65,536 on all operating systems).

Q: Can I configure to run remotely?

A: Sure. You can easily run the Sybase Aleri Streaming Platform on one machine and send it commands or data from another, since the primary method of communication is via TCP/IP. The Aleri Studio also incorporates means for remote execution and control.

Q: How do I install the Aleri Dashboard?

A: Contact Sybase's Sales representative for help with this installation. They will be happy to walk you through the process.

5. Running a Data Model

Q: How do I start the Sybase Aleri Streaming Platform with a data model?

A: You can start the Sybase Aleri Streaming Platform from a command shell. Alternatively, you can start it from the Aleri Studio by loading an existing data model with the File menu and selecting the green Run button in the toolbar to run the model. This starts the Streaming Processor and loads the active model that is open in the Aleri Studio. Start-up options can also be configured from the drop-down menu available to the right of the Run button on the toolbar.

Q: Can I start the Sybase Aleri Streaming Platform on a remote machine?

A: Yes. In the Aleri Studio, the Start tab lets you start the Streaming Processor either locally or on a remote server depending on the settings in the Configuration tab.

6. Publish/subscribe clients

Q: In what languages can I write clients?

A: The Sybase Aleri Streaming Platform comes with APIs in C++, Java, and .NET (C#). All of these communicate with the Sybase Aleri Streaming Platform over the same rudimentary TCP/IP sockets: a proprietary binary interface for sending and receiving data, and an XMLRPC interface for sending control messages.

Q: What subscription modes are there?

A: You can control how much data a subscribing clients receives in three ways: through pulsed subscriptions that send changes to the underlying streams periodically; with filtered subscriptions that pass along changes that pass a SQL statement; and by lossy subscriptions that drop the connection if the client cannot keep up.

7. SPLASH Language

Q: What is SPLASH?

A: "SPLASH" stands for Streaming Platform LAnguage SHell. It's a small imperative language designed specifically for event processing, with advanced data structures to make programming easier.

Q: What is a FlexStream?

A: A FlexStream is a stream that uses SPLASH to process events. Each input stream is associated with a SPLASH "method" that tells how to process events. FlexStreams can also have SPLASH "timers" that fire periodically.

Q: What kinds of data structures are available in SPLASH?

A: SPLASH has vectors and dictionaries, just like the scripting languages Perl and Python. It also has a powerful "event cache" data structure that makes it easy to aggregate across groups of events.

8. For Further Help

Q: Can I get help building an application?

A: Often it's easier to have an application that you can customize later. Sybase provides pre-sales engineers and consulting services for building models.

Q: What kind of support is available?

A: Sybase offers 24x7 support contracts, as well as less time-critical support.

Q: How can I get more information?

A: You can call Sybase toll free at 1-800-SYBASE6 (792-2735) or 1-925-236-5000. You can also contact Sybase via email: information@sybase.com.