SYBASE®

Product Overview

Sybase Aleri Streaming Platform 3.1

DOCUMENT ID: DC01286-01-0311-01

LAST REVISED: June, 2010

agencies.

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

Table of Contents

About This Guide

1. Purpose

This guide explains how to use the the Sybase Aleri Streaming Platform identifies the major components, and introduces key concepts. It is part of the complete documentation set for the Sybase Aleri Streaming Platform.

2. Organization

This guide includes the following chapters:

Chapter 1, *The Aleri Streaming Platform*	Explains event processing technology, how it is used, and introduces key concepts.
Chapter 2, *Product Components: What's In the Package*	Identifies the components of the Sybase Aleri Streaming Platform and explains how they are integrated within a complete system.
Chapter 3, *Integrating Data Sources and Consumers*	Shows examples of how the Sybase Aleri Streaming Platform can be configured to handle different external data sources.
Chapter 4, *Building a Data Model*	Describes the process for building and deploying an application using the Sybase Aleri Streaming Platform.
Chapter 5, *Features and Benefits*	Describes the Sybase Aleri Streaming Platform's major features.

3. Related Documents

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

Product Overview	Introduces the Aleri Streaming Platform and related Aleri products.
Getting Started - the Aleri Studio	Provides the necessary information to start using the Aleri Studio for defining data models.
Release Bulletin	Describes the features, known issues and limitations of the latest Aleri Streaming Platform release.
Installation Guide	Provides instructions for installing and configuring the Streaming Processor and Aleri Studio, which collectively are called the Aleri Streaming Platform.
Authoring Guide	Provides detailed information about creating a data model in the Aleri Studio. Since this is a comprehensive guide, you should read the *Introduction to Data Modeling and the Aleri Studio*. first.
Authoring Reference	Provides detailed information about creating a data model for the Aleri Streaming Platform.
Guide to Programming Interfaces	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. These interfaces include:

	• 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
Utilities Guide	Collects usage information (similar to UNIX® man pages) for all Aleri Streaming Platform command line tools.
Administrators Guide	Provides instructions for specific administrative tasks related to the Aleri Streaming Platform.
Introduction to Data Modeling and the Aleri Studio	Walks you through the process of building and testing an Aleri data model using the Aleri Studio.
SPLASH Tutorial	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.

Chapter 1. The Aleri Streaming Platform

The Sybase Aleri Streaming Platform lets you quickly build and run applications that can analyze an unlimited amount of incoming data in real time. That means you can get timely insight into current conditions for an immediate response to opportunities and threats, a vital asset in today's volatile business world.

Businesses deal with thousands of events each day, sometimes each second, ranging from the routine of market movements, new orders, purchases and equipment failures, to rare occurrences, such as a natural disaster. If the event generates information that can be processed by a computer system, it can be fed into an event processor to analyze in context of other events and business rules to derive intelligence or initiate a response.

At the core of the software is an unmatched high-performance complex event processing (CEP) engine that can be programmed using Sybase's authoring tools to apply complex logic to streams of incoming data. The Sybase Aleri Streaming Platform not only has the power and speed to deliver results in milliseconds or less, but also the flexibility to work with your current data infrastructure, meaning that you can have it all without the trouble and expense of revamping your existing technology system.

CEP is different from traditional transaction processing technology which process individual transactions because CEP operates on events in the context of other events. Traditional data analysis tools collect event data in a database to analyze as historical data, whereas CEP performs the analysis as soon as each event arrives, continually updating and outputting results. While real-time analysis technology has existed for some time, the Sybase CEP engine represents a new approach that can be used to quickly implement new adaptable applications.

CEP's transformational capabilities can be easily understood with the example of a fictional urban transit system using the technology to improve its customer service. The system has card readers to record when passengers enter a station, which all have sensors to report arrivals and departures of trains. All the event data would be sent in real time to a central event processor that continuously computes the number of people waiting at each station platform and the average times the trains spend in each station. The event processor would then compare passenger numbers and stop times across all stations and train lines to dynamically allocate trains and staff according to demand.

The Sybase Aleri Streaming Platform is a robust, scalable platform designed for speed, efficiency and reliability. Its CEP technology dramatically reduces the time, cost and skill level required to implement and deploy real-time solutions. That means you can focus on the business logic that will increase your profits, without having to design and build an efficient, real-time data infrastructure. It is a multi-threaded 64-bit implementation, which allows it to take full advantage of the latest hardware platforms. But the Sybase solution doesn't require any multi-threaded or event driven programming skills; basic models can be implemented by users without any programming skills.

There are endless opportunities to derive value from event processing technology in any enterprise that needs to filter and analyze vast amounts of event data in real time and act quickly on the information. Here are just a few examples:

Financial Industry

The financial markets are dynamic - they are constantly and rapidly changing. Business rules and processes can't remain static; software needs to adapt to changes in the markets. CEP can help financial firms deploy new technology solutions quickly and in a way that allows solutions to easily evolve and adapt as market conditions change.

In particular, CEP is rapidly being embraced in the capital markets, where real-time insight and rapid response is critical, and the rate at which new information arrives can be overwhelming.

In the front office, market data represents events, such as executed trades and quotes that have been

published. An incoming order is also an event - a decision by a client or portfolio manager to buy or sell a particular security.

Sybase CEP is being used in a variety of ways, including:

- Market data enrichment to synthesize new market data from raw inputs. Combine data from different sources based on rules. Compute insight, such as "best bid and offer" across specific markets.
- Market liquidity analysis to combine order book data from multiple markets to analyze full market liquidity, to determine where, when and how to trade.
- Quantitative trading to monitor the market for pre-defined conditions that represent opportunities, automating the response to capitalize on the opportunity before it disappears.
- Algorithmic execution is the process used for for most common automated trading orders. As more people use this method, algorithms keep changing for those in search of an edge. CEP provides a platform for rapid implementation that can get new algorithms in place quickly.
- Smart order routing with use live market data to intelligently route orders for best results and in compliance with best execution policies.
- Best execution compliance to collect data at the time of each trade to monitor and show compliance with best execution policies.
- Order stream analytics to monitor order flow in real-time for insight into market dynamics and customer behavior; watch for patterns of undesirable trading.
- Real-time portfolio valuation to continuously update positions in real-time using incoming transactions and value positions using live market data, including net gain/loss and exposure calculations. Generate alerts when thresholds or limits are approached or exceeded.
- Real-time risk aggregation to consolidate position and risk information from multiple systems to monitor consolidated exposures in real-time.

Telecommunications

Providing the best customer service is often the deciding factor in the highly competitive telecom industry. One segment where it's particularly critical is the prepaid call authorization market, with numerous competitors fighting for customers. Telecom companies can ill afford a customer service misstep since it's so easy to find another provider. But to get it right, companies need to have complete account information available in real-time so that services can be provided without any problems or delays.

CEP technology gives you the information to do just that. The Sybase Aleri Streaming Platform provides real-time correlation of call-initiated and call-terminated events for this segment. Custom applications have been built to perform the task, but it's difficult and has proven very expensive to maintain. It can also require a great deal of time to adapt the applications to new business models. You can easily create and adapt models for new service plans on the Sybase Aleri Streaming Platform.

Network Security

Network security systems generate thousands of “incidents” every day. Most incidents are harmless, but some combinations can indicate threats, so it's important to detect emerging trends and act on them. An event processor can monitor streams of events such as “incident reports” from network sensors and perform real-time correlation of events. The Sybase Aleri Streaming Platform is also fine-tuned enough to eliminate false-positive alerts, which are costly both in terms of money and time. If a significant pattern is detected, the event processor will alert an administrator or initiate a response. It can also produce high-level statistics on incoming events to monitor trends and generate alerts when

the situation abruptly changes or exceeds defined thresholds.

eCommerce

Knowing your customer is perhaps the most important key to success for any company. Web advertisers and eCommerce firms have an incredible advantage with click-stream technology that records when, where and how customers are surfing the Web. Sybase's CEP software can make click-stream technology far more powerful since it can provide real-time analysis of the data to gain insight into current trends so that companies can respond more quickly. CEP is replacing systems that delivered reports on what users were doing yesterday, with applications that show what users are doing right now. By seeing what is working and what isn't, problems and opportunities can be detected, and adjustments can be made before it's too late, saving time and money.

Retail

Severe pricing pressures exist at every level in this industry, which means retailers have to improve efficiencies and rein in costs to maximize profits. An event processor can help make these goals a reality. CEP engines can be modeled to give retailers the data needed to figure out which products will be bestsellers by showing real-time sales data against inventory, with the ability to factor in the numerous variables that affect demand, such as weather, economic conditions, seasonality and other factors. Once the decision is made on the product lineup, a retailer can then use CEP to track down the most efficient way of getting the product from the warehouse to the shelves. It doesn't stop there, since CEP can give up-to-the-millisecond sales performance evaluations and analysis, so that if a product isn't selling and prospects don't look great based on trends, retailers can make quick adjustments to ensure valuable shelf space isn't wasted.

Chapter 2. Product Components: What's In the Package

This chapter introduces the components that make up the Sybase Aleri Streaming Platform and explains how to integrate other systems and components into a complete application.

2.1. A System Level View

The primary component of the Sybase Aleri Streaming Platform is the *complex event processor* which runs as a server process. The event processor loads a data model that has been defined using one of the Sybase authoring options, and then waits for message(s) to arrive on the source streams. As each event arrives, the processing logic contained in the data model is applied to update one or more *derived streams*. Output messages are continuously generated for the derived streams that have registered subscribers. In addition to real-time streaming output, results can be queried via the *On-Demand SQL* query interface.

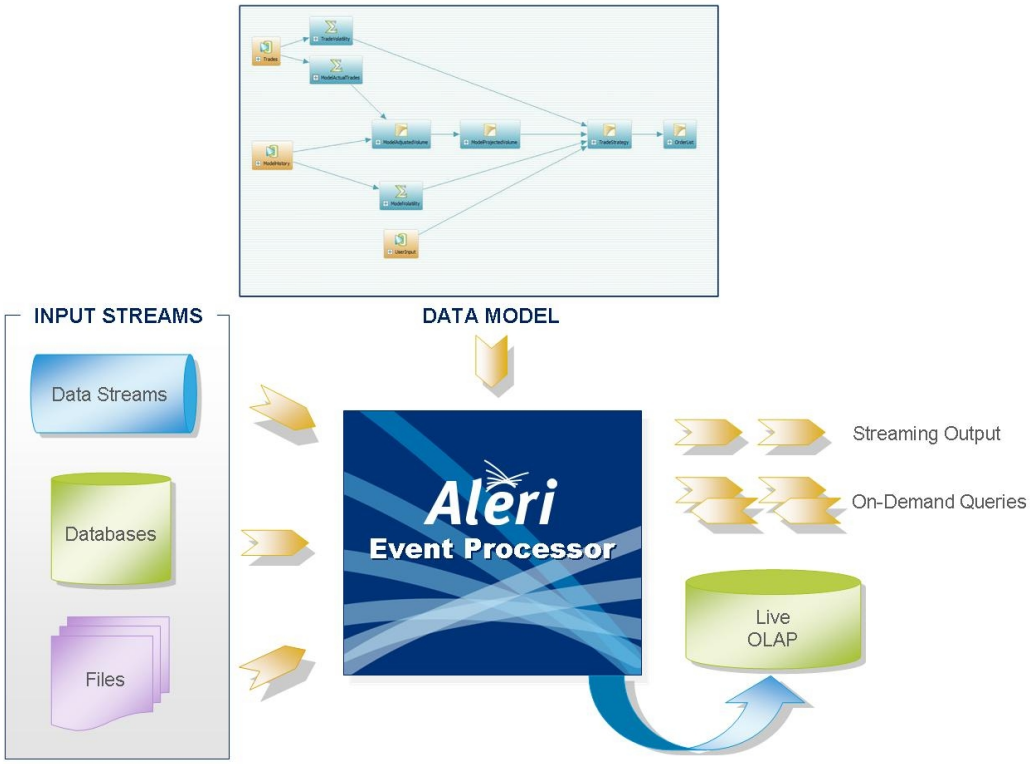

The list below is the complete application environment. Your installation might include some or all of the following, depending on your specific business needs.

Streaming data sources	Provide continuous event data (such as Reuters Market Data System).
Static data sources	Can be databases or files.

Message Bus	Carries messages from data sources to the Sybase Aleri Streaming Platform and delivers streaming output to applications that will use the results (for example, Tibco Rendezvous, International Business Machines Corp., MQ and Java Message Service (JMS).
Dashboard	Can be used to display output of the Sybase Aleri Streaming Platform. Aleri Dashboard tools can be integrated using connectors, custom adapters or an off-the-shelf adapter that publishes in a format compatible with the dashboard application.
Query Tools	Standard query or reporting tools such as Excel, Business Objects and others can be used to query the Sybase Aleri Streaming Platform using the Open Database Connectivity (ODBC), C++ Connectivity API, or Java Database Connectivity (JDBC) interfaces provided with the Sybase Aleri Streaming Platform.
Historical Event Database	When you are capturing a lot of event data and need to build up a permanent event history, the Sybase Aleri Streaming Platform can be integrated with an off-the-shelf database. Any data collected can be retained and periodically moved into the historical database. LiveOLAP accumulates historical event data, which can be queried as the data receives continuously updated feeds.

Connectors and Adapters	Convert incoming messages into the Sybase data format and convert outgoing messages from the Sybase data format into the format expected by the receiving application. These are not needed if the publisher or subscriber uses an Sybase API directly or if a built-in connector is used.
	
	The Sybase Aleri Streaming Platform has a range of built-in connectors that can be used to link the most common general-purpose data sources and destinations. Version 3.1.11 includes connectors for:
	
	• JDBC (one time and periodically polled) which can connect to databases, such as:
	  • Sybase Adaptive Server® (ASE)
	  • Oracle®
	  • SQL Server™ (Microsoft)
	  • Netezza®
	  • kdb+ (Kx Systems)
	  • DB2® (IBM)
	  • Teradata®
	  • PostgreSQL®
	• XML formated messages from/to a file or socket
	• CSV formated messages from/to a file or socket

- Sybase Aleri Streaming Platform to/from Sybase Aleri Streaming Platform
- CSV, XML, Object Array, or FIX messages on JMS
- FIX messages from/to a file or socket
- Mail server (SMTP) (outbound)

The Sybase Aleri Streaming Platform can also be used with external connectors to link with data sources, such as SybaseIQ Output and kdb+tick Streaming Input.

2.2. Adapters

Adapters integrate with specialized data sources and destinations to provide a range of connectivity options beyond what is available with built-in connectors. On the input side, an adapter converts incoming events from the native format to the Sybase message format and passes the message into the Sybase server. On the outbound side it's just the opposite: connectors receive messages from the Sybase server, convert them to the desired format and pass them on.

The list of adapters is constantly growing. Consult Sybase to discuss your specific integration and connectivity needs.

2.3. Authoring Tools

The Sybase Aleri Streaming Platform authoring tools and languages are used to create the data models that contain the business logic applied to incoming event data. The data models define:

- the source (input) streams
- the data schema of each stream (how the information is organized)
- the operators and expressions that are applied to produce derived streams

These derived streams contain the results of the event processing.

The Sybase Aleri Streaming Platform provides a choice of three different authoring environments with the same functionality available in all. Most users choose the Aleri Studio to create, edit and test their models. The Aleri Studio is an Integrated Development Environment (IDE) based on the Eclipse ™framework. It provides both a visual editor and Aleri SQL editor (see below) along with tools for testing and debugging a model.

Aleri Studio It lets you to create *data models* using a visual editor to add source streams and derived streams onto a canvas. The dataflows are defined by simply connecting the source streams and the derived streams that apply the logic of the data model to the input. It also provides tools for testing, debugging and monitoring the streams. The Aleri Studio saves continuous data models as AleriML files.

Aleri SQL Aleri SQL is a subset of ANSI SQL with extensions for real-time stream processing and provides a familiar environment for SQL programmers. It is a text language that can be edited using any text editor.

Aleri SQL editing facilities are also provided within the Aleri Studio. Models built in SQL can be visualized in the Aleri Studio, and then run from the Aleri Studio,

testing and debugging. Aleri SQL is translated to AleriML.

AleriML AleriML is an XML language for defining a Sybase data model. Any text editor or XML editor can be used to create and edit models in AleriML. Models that are created in Aleri SQL or in the Aleri Studio produce a model in AleriML for loading into the Stream Processor.

2.4. Programming Interfaces

Interfaces are supplied to communicate with the Sybase Aleri Streaming Platform. These interfaces provide flexibility for the development of custom adapters, external control and monitoring, and the use of external functions.

Publish/Subscribe API	This interface is used to publish data or load static data into the Sybase Aleri Streaming Platform, and also to subscribe to streams within the Sybase Aleri Streaming Platform. Also known as the “Pub/Sub interface”, it is available in C++, Java and .NET. This API is used to build adapters. It can also be embedded directly in an application that either generates data or consumes data produced by the Sybase Aleri Streaming Platform server.
Command & Control Interface	This interface is used to control or monitor the Sybase server. It is based on the Remote Procedure Call (XML-RPC) standard.
On-Demand SQL Query Interface	This interface is used to run SQL queries against any of the streams in the data model. This interface supports three connection interfaces: ODBC, JDBC or a C++ API. The ODBC/JDBC connections can provide support for off-the-shelf reporting tools that can query an ODBC/JDBC source.
User Defined Function Interface	This interface allows proprietary or off-the-shelf functions to be invoked within an Sybase data model. It allows you to expand the capabilities of the product beyond what is available through the built-in function library.

2.5. AleriRT for Microsoft Excel®

AleriRT is the real-time plug-in for Microsoft Excel on the Sybase Aleri Streaming Platform. It is not packaged with the Sybase Aleri Streaming Platform; your Sybase representative can provide information about downloading it. It includes a Wizard-driven interface for setting up a subscription. Once a subscription is set up, worksheets can be updated in real time with data streamed from the Sybase Aleri Streaming Platform into Excel. You can also use AleriRT to publish data from Excel into the Sybase Aleri Streaming Platform. Data can automatically be injected into the data model whenever a cell in a specified range changes.

2.6. LiveOLAP Server

Online Analytical Processing (OLAP) software integrates with the Sybase Aleri Streaming Platform to collect historical event data in real time, allowing a user to interactively analyze the data across different dimensions. The unique data store allows ad-hoc analysis across any dimension without the pre-aggregation required by traditional OLAP tools. It includes a standard browser-based query interface as well as a Screen SDK (Software Developer's Kit) for building custom browser-based query pages and supports off-the-shelf OLAP query tools.

2.7. Aleri Dashboard

The Aleri Dashboard is an optional component based on real-time display technology from SL Corporation. It is a powerful tool 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 comes with a streaming connector that streams live data into the dashboard.

2.8. Utilities

The Sybase Aleri Streaming Platform includes many utilities, which are fully described in the *Utilities Guide* .

sp	Streaming Processor a server for processing relational stream data
sp_clustermgr	cluster manager for the Sybase Aleri Streaming Platform
p_clustermon	cluster monitor for the Sybase Aleri Streaming Platform
sp_ld	launch cluster daemon for the Sybase Aleri Streaming Platform
sp_monitor	read performance data from a running instance of the Sybase Aleri Streaming Platform and print it in a XML format on standard output.
sp_query	Send an SQL query to the Sybase Aleri Streaming Platform and print results on the screen.
sp_server	Streaming Processor server script
sslwrap	simple TCP service encryption using TLS/SSL
sp_cli	Streaming Processor command line utility
sp_cnc	sample client application for submitting command and control commands to the Sybase Aleri Streaming Platform
sp_convert	Read XML or delimited records from standard input and produce Sybase Aleri Streaming Platform binary records on standard output.
sp_histexport	Historical Data Migration Utility
sp_subscribe	A sample subscription client for the Sybase Aleri Streaming Platform.
sp_upload	A sample record upload client for the Sybase Aleri Streaming Platform.
sp_archive	Archives data from the the Sybase Aleri Streaming Platform. to SybaseIQ either in batch mode or in real-time.
sp_viewer	View the changing state of streams in the Sybase Aleri Streaming Platform.
sp_sql2xml	Compile SQL statements and generate an Sybase Aleri Streaming Platform XML configuration file.
sp_studio	Script for starting the Visual Authoring Environment.
sp_upgrade	Convert old AleriML files to version 3.1.11 of the Sybase Aleri Streaming Platform.

2.9. Security

The Sybase Aleri Streaming Platform includes three mechanisms for protecting sensitive data.

The first level of defense is **authentication**, which validates users via secrets like passwords. The Sybase Aleri Streaming Platform can be started with no authentication checks, with authentication checks provided by Pluggable Authentication Modules (PAM) on Linux and Solaris, RSA public keys or Kerberos. If the authentication check fails — such as the server is configured with password checks, and the user supplies an incorrect password — then the person will not be able to access the server, publish or subscribe to data, or issue commands to control the server process.

A second level of defense is **access control**, that limits access of already authenticated users. For example, you can limit a business user to subscribe to only a handful of streams, or allow a system administrator full access to all streams. Access control is done on Linux and Solaris using the “group” facilities in the operating system, which means that access can be granted or denied for classes of users.

The final level of defense is **encryption**, which protects the data sent over a network. Data sent over the Internet, or even internally within a company, might be sensitive and need to be protected. Individual connections to the Sybase Aleri Streaming Platform can be encrypted via secure sockets (SSL). That makes the data sent over the network unreadable to anyone but the intended user.

These mechanisms, each based on open standards, provide the security needed to use the Sybase Aleri Streaming Platform in an enterprise setting.

Chapter 3. Integrating Data Sources and Consumers

Data Sources

An application that runs on the Sybase Aleri Streaming Platform must have one or more sources publishing data to it. One of these sources generally will be a real-time event stream.

Whether the data source is a real-time stream or static data that gets loaded from a file or database, data enters the Sybase Aleri Streaming Platform in one of three ways:

- streamed into the Sybase server via the Sybase Pub/Sub API (a client that is written using this API)
- via an adapter that uses the API
- via an in-process data connector tightly integrated with the Sybase Aleri Streaming Platform and the Aleri Studio

The API is a relatively simple programming interface, available in C++, Java and .NET versions, that allows an application to connect to the Sybase Aleri Streaming Platform and send records to one or more source streams in the data model. The API also allows the application to query the Sybase Aleri Streaming Platform to find out what streams are available in the data model that's running.

Sybase has connectors to facilitate data absorption from files and sockets, message buses, databases and other data sources. Message bus adapters are available for TIBCO Rendezvous, TIBCO EMS, JMS, and IBM MQ. Database adapters are available for Sybase ASE, Microsoft SQL and Oracle's database. A number of task-specific adapters are also available, including market data adapters for Reuters, ActiveMQ and others. Many connectors support inbound and outbound messages in a variety of formats, including XML, delimited text, SOAP and SWIFT, as well as custom formats requiring regular expressions. A very basic connector and source code that loads XML or delimited text from files is included in the base product.

The choices for connectors and adapters is always expanding so the above list might not include all options. Check with Sybase to see if a specific adapter is available. If not, Sybase can generally build one for you or custom adapters can be developed using one of the Sybase Pub/Sub APIs.

Connect Data Publishers to Source Streams Model

Input data enters the data model running on the Sybase Aleri Streaming Platform via a source stream in the model. Every Sybase data model will have at least one source stream to receive input data. Source streams are passive and accept data from any source that sends it as long as it is correctly formatted. This allows multiple event sources to publish data into the same source stream.

An adapter or application uses the Sybase Pub/Sub API to send data to the Sybase Aleri Streaming Platform server. When the adapter or application hands off an event record (a set of fields) to the API, it specifies the destination stream ID and an operation code (opcode) that indicates how the event record should be processed.

insert	Insert a new record into the stream.
update	Update an existing record into the stream.
upsert	Update an existing record if there is a record with the same key field, otherwise insert this record into the stream.
delete	Delete a record with a specific key field value from the stream.

Data Consumers – Continuous Output

A data consumer works with the results of the event processing performed by the data model running on the Sybase Aleri Streaming Platform. Sybase data subscribers can take several forms:

- An output connector that directly writes streaming data to files, sockets, databases, or message buses (the reverse of the input data source connectors).
- An adapter that passes streaming output directly into another application or publishes it onto a message bus.
- A dashboard that displays real-time output to a user (a custom application or one built with the Aleri Dashboard).
- The LiveOLAP Server, deployed as a consumer of the output of the Sybase Aleri Streaming Platform, which allows users to interactively analyze event data across multiple dimensions.
- AleriRT for Microsoft Excel®, can subscribe to output, pushing it directly into a worksheet where it can feed calculations in the worksheet or be displayed.

A single data model can produce multiple output streams. The data model does not designate a particular stream as an output stream, telling it to publish its results. Rather, it's up to an external client to initiate a subscription to a particular stream. Subscriptions are established using the Sybase Pub/Sub API — the same one used to stream data into the Sybase Aleri Streaming Platform.

When an adapter or client application wishes to subscribe to the output of a particular stream, it tells the Sybase Pub/Sub API to establish a subscription to that stream (by specifying its ID). It can also specify whether it wants to receive all data currently held or just new data as it arrives. The subscription will stay open until the adapter or client application terminates it; the Sybase Aleri Streaming Platform will send all updates in real-time over all open subscriptions. A single application can subscribe to multiple streams and a single stream can have multiple active subscriptions. There are a number of different types of subscriptions, each tuned to a particular need.

Data Consumers - On-Demand Output

You may have applications that need access to live data, but you don't want it continuously delivered. This type of application can issue an SQL query to the Sybase Aleri Streaming Platform in the same way that you would query a database. The Sybase Aleri Streaming Platform includes an ODBC/JDBC provider interface, so that even off-the-shelf tools – any application that can query an ODBC/JDBC source – can query the Sybase Aleri Streaming Platform. Any stream in the data model that retains data can be queried, and while the query returns a snapshot, it runs against live data.

An example of an application that might use this interface would be a financial position application. The trader or risk manager wants the ability to see his or her current positions immediately at any time. Rather than building a display application that continuously receives and caches all updates streaming out of the Sybase Aleri Streaming Platform, it's much easier to use off-the-shelf tools to grab a snapshot of the data when needed.

Chapter 4. Building a Data Model

The Sybase Aleri Streaming Platform runs a *data model* that contains source streams, derived streams and the event processing logic that will be applied to incoming event data to produce one or more output streams or result sets. Any of the three Sybase authoring environments (Aleri Studio, Aleri SQL, AleriML) can be used to create and edit a data model.

Sybase's event processing architecture is based around a relational model with source streams being comparable to tables and derived streams to views. But, think of a data stream as a table that has no beginning and no end. Typically, an Input window defined on the stream specifies the size of the table. For a given stream, you may need to hold a minute's, an hour's, or several days worth of data. It could also be defined in terms of number of records or only operating on the latest record.

Source Streams

A source stream receives input data from a source known as a publisher outside the Sybase Aleri Streaming Platform. Source streams get incoming messages that can be applied as inserts (append), updates, deletes or upserts (update if present, otherwise insert). Streams can also be configured with retention rules that specify how long data must be retained. Other options include streams that accept only inserts and don't retain data.

Derived Streams

There are several types of derived streams that are used to apply different operations to one or more input streams to produce a result stream. Derived streams always take their input from other streams defined in the data model. The following derived stream types are supported on the Sybase Aleri Streaming Platform:

Aggregate	Groups records according to a common value in one or more columns and produces a single summary record for each group. Takes a single input stream.
Filter	Filters records according to a defined filter expression. Only records on the input stream that cause the filter expression to be ‘true’ are included in the output. Takes a single input stream.
Compute	Computes output records from input records. The fields in the output record are computed according to an expression defined for each column. Takes a single input stream.
Extend	Is just like a Compute Stream, but does not require an expression for every column in the output stream. Columns in the input stream are passed through unmodified and any new columns are added according to the column expressions defined for the Extend Stream. Takes a single input stream.
Pattern	Watches for patterns of events and produces one or more output records whenever a defined pattern is detected. Takes any number of input streams. Patterns can be defined for a single stream or across multiple input streams.
Join	Joins two or more streams to produce a single output stream. Just like a database join, output records are produced by matching records from the input streams according to common values in selected columns.
Union	Merges two or more streams into a single stream. Input streams must have the same format (column definitions).
Copy	Produces a copy of the input stream. Primarily used for clustered models to provide a local copy of an input stream. Can also be used to effect a self-join.

Flex	Is a programmable stream. Takes any number of input streams. Events on an input stream invoke a method written in the SPLASH scripting language.

Stores

Every stream is assigned to a “store”. A stream's store holds all data records retained by a stream.

The Sybase Aleri Streaming Platform provides three different types of stores to balance performance against the need to prevent data loss. A single data model can contain multiple stores of the same or different types. The types of stores are:

Log Store	Holds the data in memory and provides disk-based persistence, ensuring that no data is lost when the system is restarted. If the system fails, the stream's full state at the point of failure can be restored.
Memory Store	Retains and holds all the data in memory. Faster than the Log Store, but data is lost if the system is restarted.
Stateless Store	Does not keep a full set of data based on the retention rules, but passes events to downstream nodes. Data cannot be queried in a Stateless Store.

Chapter 5. Features and Benefits

Features of the Sybase Aleri Streaming Platform include:

- Three different authoring environments:
  - The **Aleri Studio** is a visual, interactive development environment based on the Eclipse Framework™.
  - **Aleri SQL** is a subset of standard SQL with extensions for real-time data and can be used with any text editor.
  - **AleriML** uses the Aleri XML schema, providing an option for scripted generation and development of alternate compliers.
- **Ultra-low latency**, on the order of milliseconds, or even microseconds, can be achieved.
- **High throughput** which provides the ability to process hundreds of thousands, or even millions, of messages per second on a single server, and far more in clustered configurations.
- **Scalability** that is multi-threaded to take full advantage of multi-CPU machines and designed with full support for clustering.
- Real-time **streaming output** *and* support for **on-demand queries**.
- Optional **disk-based data persistence**, optimized for throughput and latency, ensuring that critical data is not lost.
- **Security** features include choice of PAM, RSA, or Kerberos authentication, encryption and access control.
- **High availability** with support for hot standby and cold spare configurations.
- Support for **update events** allowing incoming data to be applied as inserts, updates, deletes or upserts, unlike typical event processing that always treats each incoming event as a new item.
- **FlexStream functionality** allows procedural logic to be layered on the relational models, providing incredible versatility for addressing the widest possible range of event processing requirements.
- **Pattern Matching functionality** uses a powerful syntax for matching events and raising alerts.

Chapter 6. Example

The example of a Portfolio Valuation model shows the value of the Sybase Aleri Streaming Platform's CEP technology.

A Portfolio Valuation model calculates the total value of the holdings in real-time. CEP software gives you the edge on creating profits by getting the information needed to make a trade or purchase as soon as a trend appears, which will put you ahead of the crowd during volatile times in the financial industry.

In this model, the Price Feed and Positions are the source or input streams which receive data. Price Feed would be current stock market prices and Positions would be the portfolio's financial instruments. Source streams feed data into selected derived streams that subscribe to it in order to do the assigned processing logic. The information in Price Feeds is subscribed by the VWAP (Volume Weighted Average Price) derived stream, which then sends processed information to the IndividualPositions derived stream, which also receives data from the model's other source stream, Positions. The IndividualPositions derived stream delivers its data to the Value by Book derived stream, which is the value of the Portfolio's book holdings.

Here is how the model looks in the Aleri Studio:

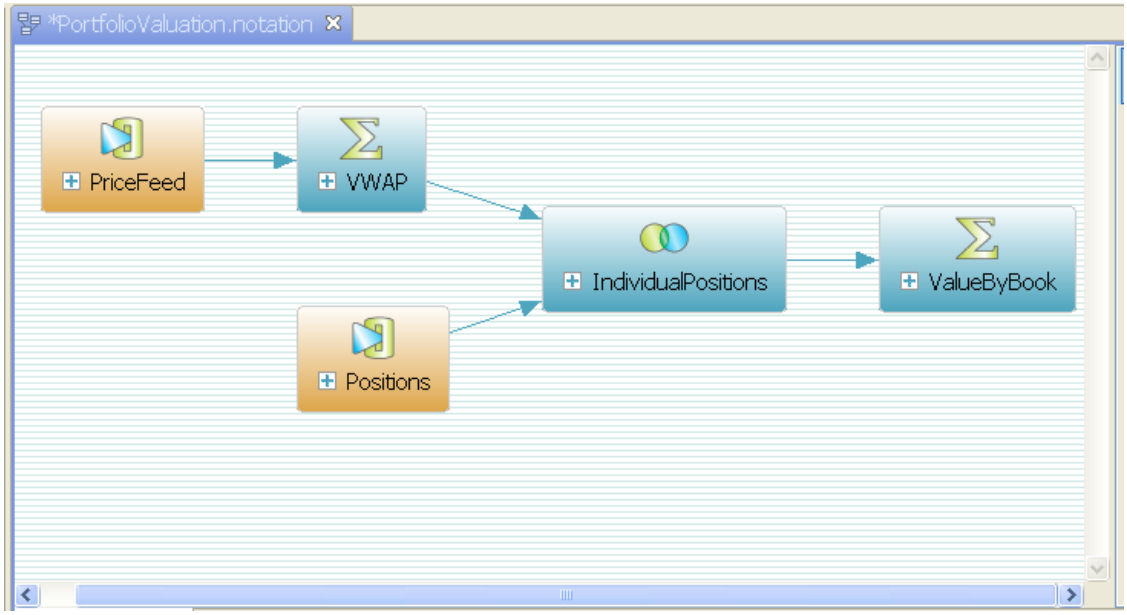

Here are the instructions for building this model in Aleri SQL:

```
*** Create the Store ***/

CREATE STORE store1 MEMSTORE INDEX IS HASH;


/*** Create Table PriceFeed optimized for insert only ***/

CREATE TABLE PriceFeed (

      Id          int32,
      Symbol      string,
      Price       double,
      Shares      int32,
      TradeTime   date,
```

```
      PRIMARY KEY (Id),
      STORE is store1
);



/*** Create the VWAP Materialized View ***/

CREATE MATERIALIZED VIEW VWAP

      PRIMARY KEY (Symbol)

      STORE IS "store1"

AS SELECT

      ticker.Symbol Symbol,

      ticker.Price LastPrice,

      SUM(ticker.Price * ticker.Shares)/ SUM(ticker.Shares) VWAP,

      ticker.TradeTime  LastTime

FROM

      PriceFeed ticker

GROUP BY

      ticker.Symbol;




/*** Create the Book Table ***/

CREATE TABLE Positions (

      BookId                  string,
      Symbol                  string,
      SharesHeld              int32,
      PRIMARY KEY (BookId, Symbol),
      STORE IS store1
);



/*** Create the Individual Positions Join Table ***/

CREATE MATERIALIZED VIEW IndividualPositions

PRIMARY KEY (BookId,Symbol)

STORE IS store1

AS SELECT

      pos.BookId BookId,
```

```
      pos.Symbol Symbol,
      avg.LastPrice * pos.SharesHeld CurrentPosition,
      avg.VWAP * pos.SharesHeld AveragePosition
FROM
      Positions pos LEFT JOIN VWAP avg
            ON (pos.Symbol = avg.Symbol);

/*** Create the ValueByBook Materialized View ***/
CREATE MATERIALIZED VIEW ValueByBook
PRIMARY KEY (BookId)
STORE IS store1
AS SELECT
      ipo.BookId BookId,
      SUM(ipo.CurrentPosition) CurrentPosition,
      SUM(ipo.AveragePosition) AveragePosition
FROM
      IndividualPositions ipo
GROUP BY
      ipo.BookId;
```