

StreamBase High Availability

- Deploy Mission-Critical StreamBase Applications in a Fault Tolerant Configuration

By Richard Tibbetts
Chief Technology Officer, StreamBase Systems

StreamBase High Availability

The StreamBase Event Processing Platform™ provides the ability for customers to deploy mission-critical StreamBase applications in a fault tolerant configuration. High Availability is a broad term, and is used here to mean that the system will continue to operate correctly in conditions such as:

- Hardware failure of a server running StreamBase
- Unexpected software failure in the StreamBase Event Processing Platform™
- Administrative error on a server, such as accidental shutdown
- Failure of network links
- Removing a server from production for maintenance

Correct operation means that the system will continue to correctly process new events and that downstream systems will be insulated from a gap in processing.

Overview

The principal technique for fault tolerance is to deploy a pair of servers, both executing the same application, one running as the primary server and one as the backup. (see Figure A) Data from upstream sources is delivered to both servers. The application running on the primary server is responsible for actual decisions, and for producing the results that are delivered to downstream systems.

The application running on the backup executes all calculations and maintains state, but under normal operation does not deliver results to the downstream system. If the application is non-deterministic, such as making trades based on market timings, then application state can optionally be synchronized from the primary to the backup server. State synchronization assures that the backup server state correctly reflects the processing on the primary and the results already delivered to downstream systems.

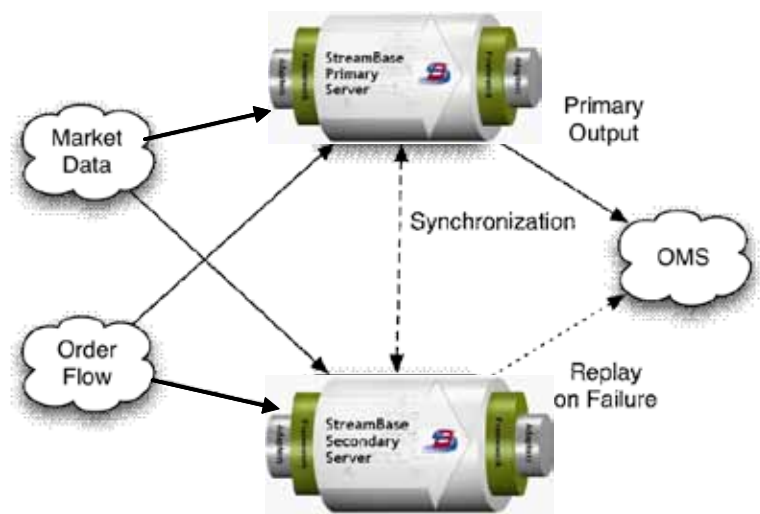


Figure A – StreamBase High Availability

In the event of failure on the primary, the backup server will take over as primary. The backup server uses a

system of heartbeats to monitor the health of the primary. When the primary stops responding to heartbeats, the backup server determines that a failure has occurred. At this point, the backup server stops listening to state updates from the primary, and begins delivering messages to the downstream system, optionally filling in messages that the downstream system missed.

To return to a highly available state, a new backup server can be introduced and synchronized with the now-primary server. Once synchronization is complete, the system is returned to safety.

Components

Several components of the StreamBase Event Processing Platform™ work together to create fault tolerant deployments. (see Figure B) To characterize StreamBase High Availability capabilities in more detail, we describe some of the components below.

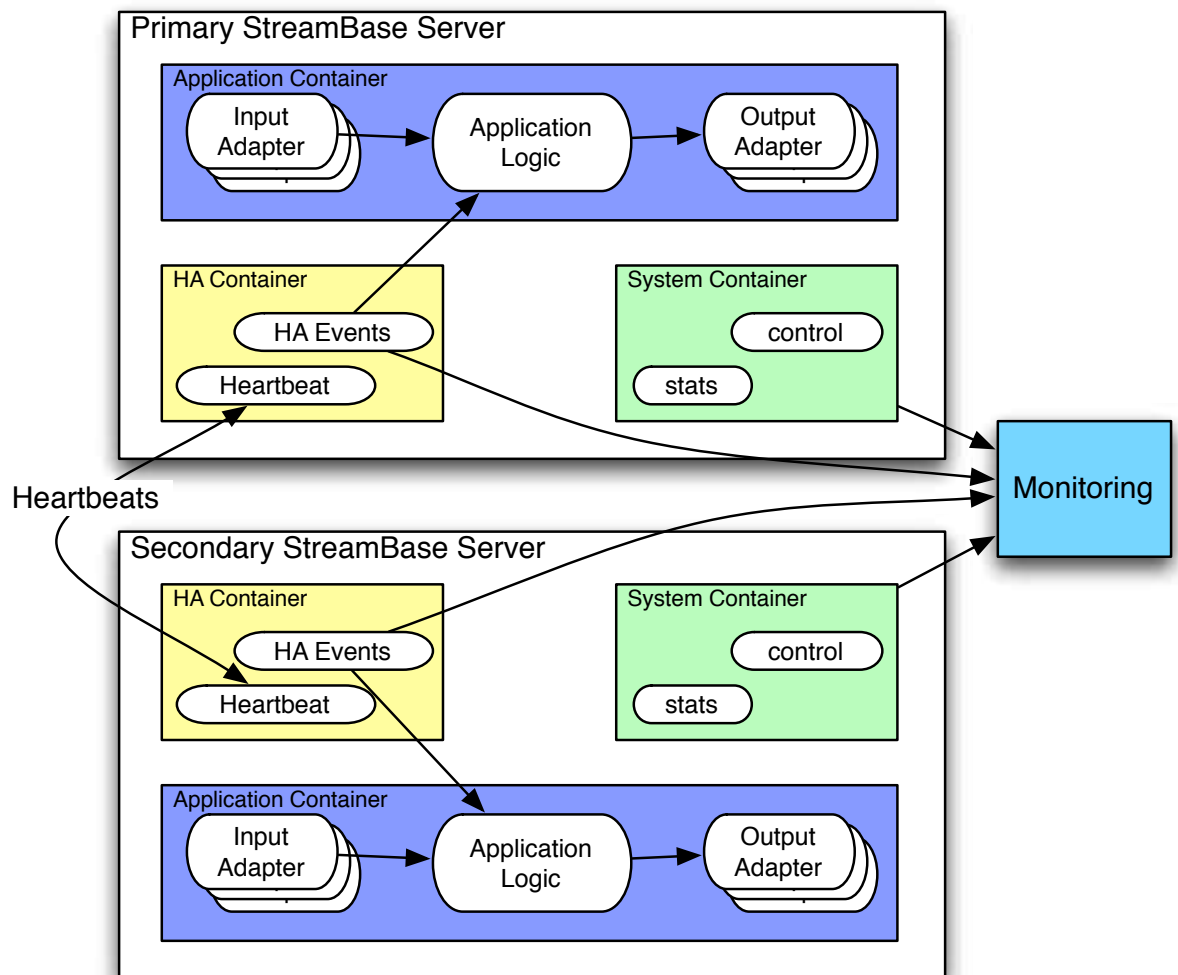


Figure B – Components of StreamBase High Availability

Containers

Each running StreamBase Server can run multiple StreamBase applications, each one of which runs in its own container. A container provides the system services needed by its application, and affords management

functionality, such as starting and stopping applications. The ability to manage containers independently is a key part of StreamBase High Availability.

Leadership

Leadership is a flag that will be set on exactly one of the nodes in the pair. Downstream clients and backup nodes will respect the leader's output messages and state updates over those of other nodes. The leadership flag is maintained by the logic in the HA container.

HA Container

Within a server participating in a highly available server pair, a special container called the HA container will exist. This container executes an HA application, which is responsible for the event processing logic that implements the high availability strategy. This application is specified in StreamBase Event Flow and provided as part of the StreamBase platform. When failure of the primary is detected by the backup, the HA container in the backup decides on the new leader and triggers HA Events which tell applications and clients that a leadership change has occurred. The HA container also publishes events that can be read by clients that want visibility into this recovery process.

Heartbeats

Heartbeats are passed between the HA containers of the two nodes in the pair. These heartbeats come at a configurable interval (generally 10 times per second) and are used to identify a failure in the primary server. When a configurable number of heartbeats are missed (generally 10) the backup server will initiate failover.

Third-Party Components

StreamBase High Availability can be deployed in a pure StreamBase environment, or it can be integrated with additional third-party components.

Highly Available RDBMS

Many StreamBase applications interact with a traditional relational database management system (RDBMS). These systems often provide high availability options, such as Oracle RAC. In a highly available deployment, StreamBase will integrate with the failover strategy of these systems through standard or proprietary APIs, to assure that data is not lost and processing continues.

Storage Area Network (SAN)

Many sites make heavy use of SAN in their high availability or disaster recovery strategies. StreamBase can persist table data to a SAN, and use that data when a backup server takes over. StreamBase can also log to a SAN using an efficient binary representation for recovery or auditing.

Multicast Messaging

Messaging systems such as RMDS and Wombat can use multicast to deliver the same messages to multiple systems simultaneously. A StreamBase primary and backup server can both listen to these messages as a way of assuring that the backup server sees all the messages.

Reliable Messaging

Third-party reliable messaging products, such as Tibco EMS or JMS can be used in place of StreamBase network connections through the use of adapters. StreamBase integrates with the failover strategy of these message queuing products, and will acknowledge the messages at the right stage of event processing.

Hardware Virtualization

StreamBase can easily be deployed on virtualized hardware. In multi-tenant situations sufficient processing power must be available in order to achieve real-time goals, but with StreamBase this is easily accomplished.

Disaster Recovery

Disaster recovery is separate from fault tolerance, because the time scales are different and the faults are larger. In a disaster recovery scenario, an entire operational site is disrupted by a natural disaster or infrastructure failure. Because the upstream and downstream systems feeding StreamBase are likely to be disrupted similarly, and this event is very infrequent, immediate and automatic failover is not generally desirable.

Instead, administrative procedures and operator intervention come into play. (See Figure C) The StreamBase server is generally restarted as part of a larger procedure. Persistent storage, such as a SAN, is used to recover critical application state, such as executed orders. For some applications, the disaster recovery procedure is asymmetric, where the

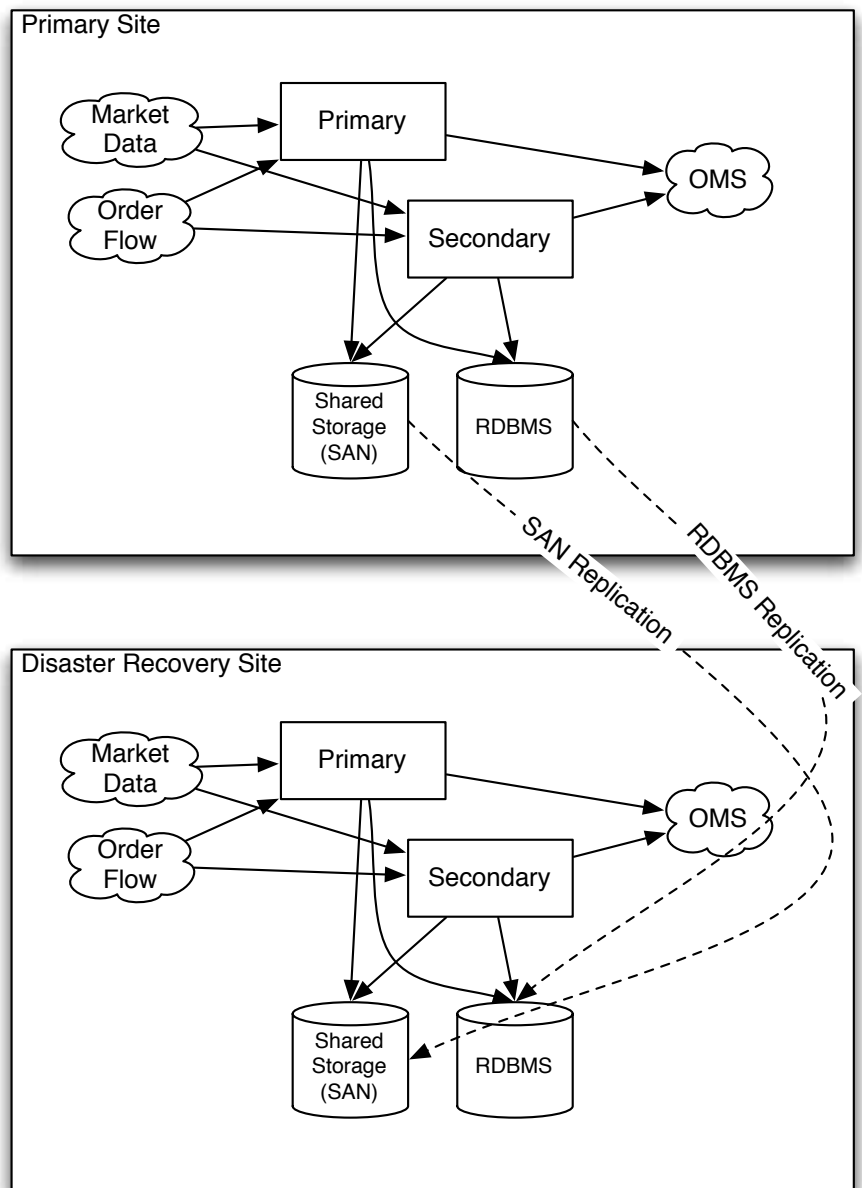


Figure C – Components of StreamBase Disaster Recovery

application running at the recovery site closes out open positions, and then ceases operation. But for many applications, the disaster recovery site achieves full operability and will continue as the primary site for days or weeks. In either case, StreamBase easily fits into existing business continuity planning.

Additional Information

Additional information is available in the following StreamBase documentation topics:

- **Primer on StreamBase Clustering and HA**

This webpage is the primary starting point for understanding clustering and HA in StreamBase. This page provides an overview of the collection of features that can be combined in various ways to support clustered and highly available servers in pursuance of different design goals.

Visit: <http://streambase.com/developers/docs/latest/admin/ha-overview.html>

- **Designing Highly Available Applications**

This webpage describes the application design patterns that HA applications have in common.

Visit: <http://streambase.com/developers/docs/latest/admin/ha-designing.html>

About Richard Tibbetts

Richard Tibbetts is co-founder and Chief Technology Officer at StreamBase Systems. Throughout his career, Tibbetts has concentrated on the design and development of database technologies, while also offering specialization in program analysis and distributed computing.

Tibbetts created the Linear Road Benchmark, the database management industry's guide for comparing the performance of real-time streaming databases with alternative systems, like relational databases. In addition, he has been heavily involved with a number of academic database projects, including Aurora at Brown University, and Medusa at the Massachusetts Institute of Technology (MIT). In 2003, the research from these two projects later led to the founding and commercialization of StreamBase's Stream Processing Engine.

As a co-founder and Chief Technology Officer at StreamBase, Tibbetts is directly involved in architecting the foundation for the next-generation textual language based on Structured Query Language (SQL). The objective of this work is to apply the benefits of SQL for stored data to real-time, transitory data. Tibbetts is also charged with global responsibilities in furthering new StreamBase capabilities through platform and technical innovation.

Tibbetts earned both his bachelor of science in computer science and engineering and his master's of engineering degree at MIT. He completed his graduate research under the direction of Dr. Mike Stonebraker.

About StreamBase

StreamBase Systems, Inc, a leader in high-performance Complex Event Processing (CEP), provides software for rapidly building systems that analyze and act on real-time streaming data for instantaneous decision-making. StreamBase's Event Processing Platform™ combines a rapid application development environment, an ultra low-latency high-throughput event server, and the broadest connectivity to real-time and historical data and leading EMS/ OMS software platforms. Six of the top ten Wall Street investment banks and three of the top five hedge funds use StreamBase to power mission-critical applications to increase revenue, lower costs, and reduce risk. It is also used by government agencies for highly specialized intelligence work. The company is headquartered in Lexington, Massachusetts with European offices in London. For more information, visit www.streambase.com.

Corporate Headquarters
181 Spring Street
Lexington, MA 02421
+1 (866) 787-6227

New York Office
845 Third Avenue, 6th Floor
New York, NY 10022
+1 (866) 787-6227

Virginia Office
11921 Freedom Drive, Suite 550
Reston, VA 20190
+1 (443) 472-0767

European Headquarters
60 Cannon Street
London EC4N 6NP
+44 (0) 20 7002 1095

www.streambase.com