



Complex Event Processing for Government Intelligence and Defense Applications



Contact:
Gary Schrader
Director of Federal Operations
StreamBase Systems, Inc
Phone: 703-608-6958
gary.schrader@streambase.com



Agenda

- About StreamBase
- Need for Real-Time Analysis, Current Challenges
- Real-Time CEP Capabilities
- Example Intelligence and Defense Applications
- Q & A




About StreamBase

StreamBase's software enables agencies to Capture, Filter, Analyze, and Act on real-time event data within milliseconds.



StreamBase is a breakthrough technology for processing and analyzing huge data volumes."

--Dr. Carol Lundquist, Sen. Scientist at Vision Systems and Technology, Inc. (VSTI)



In-Q-Tel is pleased to provide the Intelligence Community with access to StreamBase's platform for building next-generation stream processing applications

--Troy M. Pearsall, Exec. Vice President of Technology Initiatives, In-Q-Tel



About StreamBase: The High Performance CEP Leader

Enable customers to rapidly build systems that analyze and act on real-time streaming data for instantaneous decision-making.

■ Company

- Boston-based: offices in Washington DC, NY, San Mateo, & Europe
- Investments by In-Q-Tel and top-tier venture capital firms

■ Technical Team, Background

- Initial product development at MIT (2001)
- Founded in 2003 by Dr. Mike Stonebraker
- Technical Advisory Board: MIT, Brown, Stanford

■ Broadest Customer/Industry Base:

- Federal Government: Defense and Intelligence agencies
- Financial Services: Majority of top 10 investment banks and hedge funds
- Internet/Web: e-Business, 3-D virtual worlds, massively multi-player online gaming
- Leading ISVs & end-users in anti-money laundering, network/systems monitoring

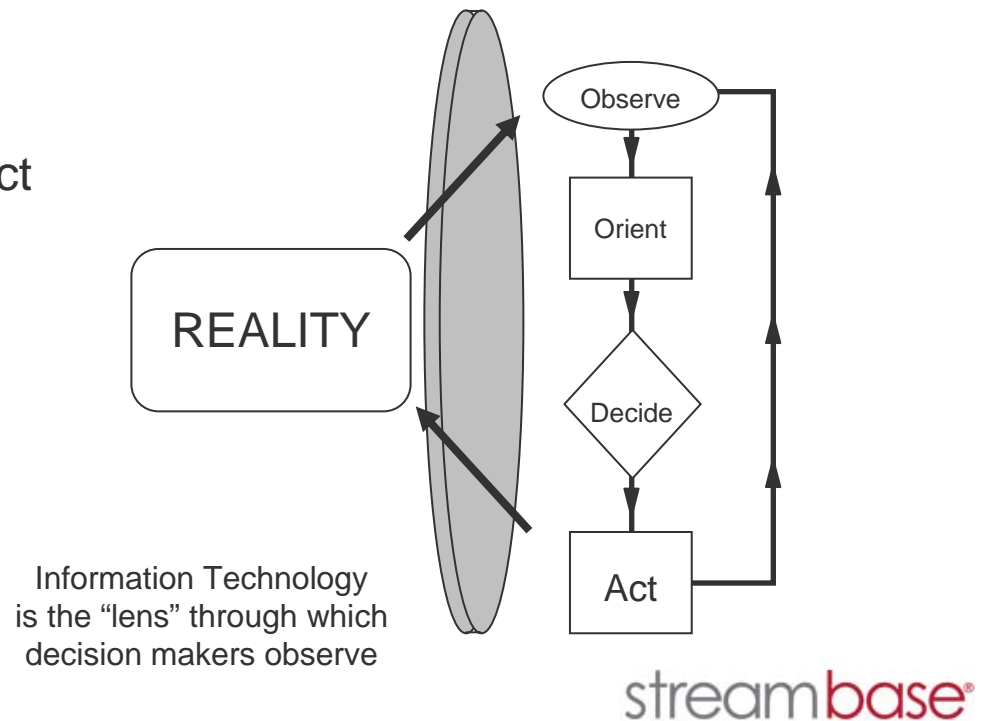


streambase®



Need for Real-Time Analysis


- Exponentially increasing volumes of structured and unstructured data that must be correlated
- Overload: Actionable Intelligence is a function of time
- Decision makers need complete situational awareness
 - Present + Past = context/intent
- OODA Loop
 - Observe, Orient, Decide, and Act





Data Deluge Problem

Disparate Data Sources



Actionable Intelligence




Actions
Alerts
Visualization



Alternatives Are Not Acceptable

■ Traditional database

- Store the data before assessment!
 - Latency
 - What if the data is not important?
- Optimized for business data processing
 - Where you don't trust the app.



■ Custom-coding

- Lots of work (low level paradigm)
- Requires 4* wizard (multi-threading, queue management,)
- Usually not easily changeable
- Maintenance nightmare



streambase®



Data Deluge Problem Solved

Disparate Data Sources



Event Processing Platform



High-performance software for rapidly building systems that analyze and act on real-time streaming data.

Actionable Intelligence



Actions
Alerts
Visualization



Many Applications Must Sense-and-Respond Instantaneously

- **Intelligence & Surveillance**
 - Aggregate data from variety of sources into holistic view
 - Identify threats
 - Dynamically prioritize, analyze, distribute with minimal latency
- **Intrusion Detection & Network Monitoring**
 - Defend against hackers/terrorists
 - Protect information infrastructure from threats
- **Battlefield Command & Control, Visualization**
 - Capture, analyze, and act on real-time sensor data
 - Precision = better, faster decisions



streambase®



Real-Time Analysis Technical Capabilities

- **Complex Event Processing**
 - Connecting and processing “in flight” data across the enterprise
- **Real-Time Analysis**
 - Evaluate the content based on filter and selection profiles, trigger events
 - Perform computational analysis on content and trigger events
- **Analytic Event Correlation**
 - Detect complex conditions based on multiple time-variant lower-level events
- **Response Management**
 - Given one or more event conditions, execute one or more responses
- **Integration with Existing Systems and Technologies**
 - Real-time technologies should augment (not displace) traditional historical analysis and storage systems



High Performance Complex Event Processing


- Use to review very large amounts of data in real time and make processing decisions based on the value or content of the data
 - Real-time alerting
 - Target profiling
 - Filtering
 - Categorizing/Binning
- Can be used for both structured and unstructured data
- Highly scalable and supports very rapid application development
- Runs on standard hardware and operating systems



Capabilities of High Performance Event Processing Platform

Rapid Development

- Graphical StreamSQL Event-Flow or text programming
- Support for full app lifecycle: adapters to processing to output
- Eclipse-based



Extreme Performance

- Throughput: 10K-1M+ msg/sec
- Ultra-low latency
- Scalability: blades, clusters, multi-core CPUs
- Architecture: 64-bit, multi-threaded

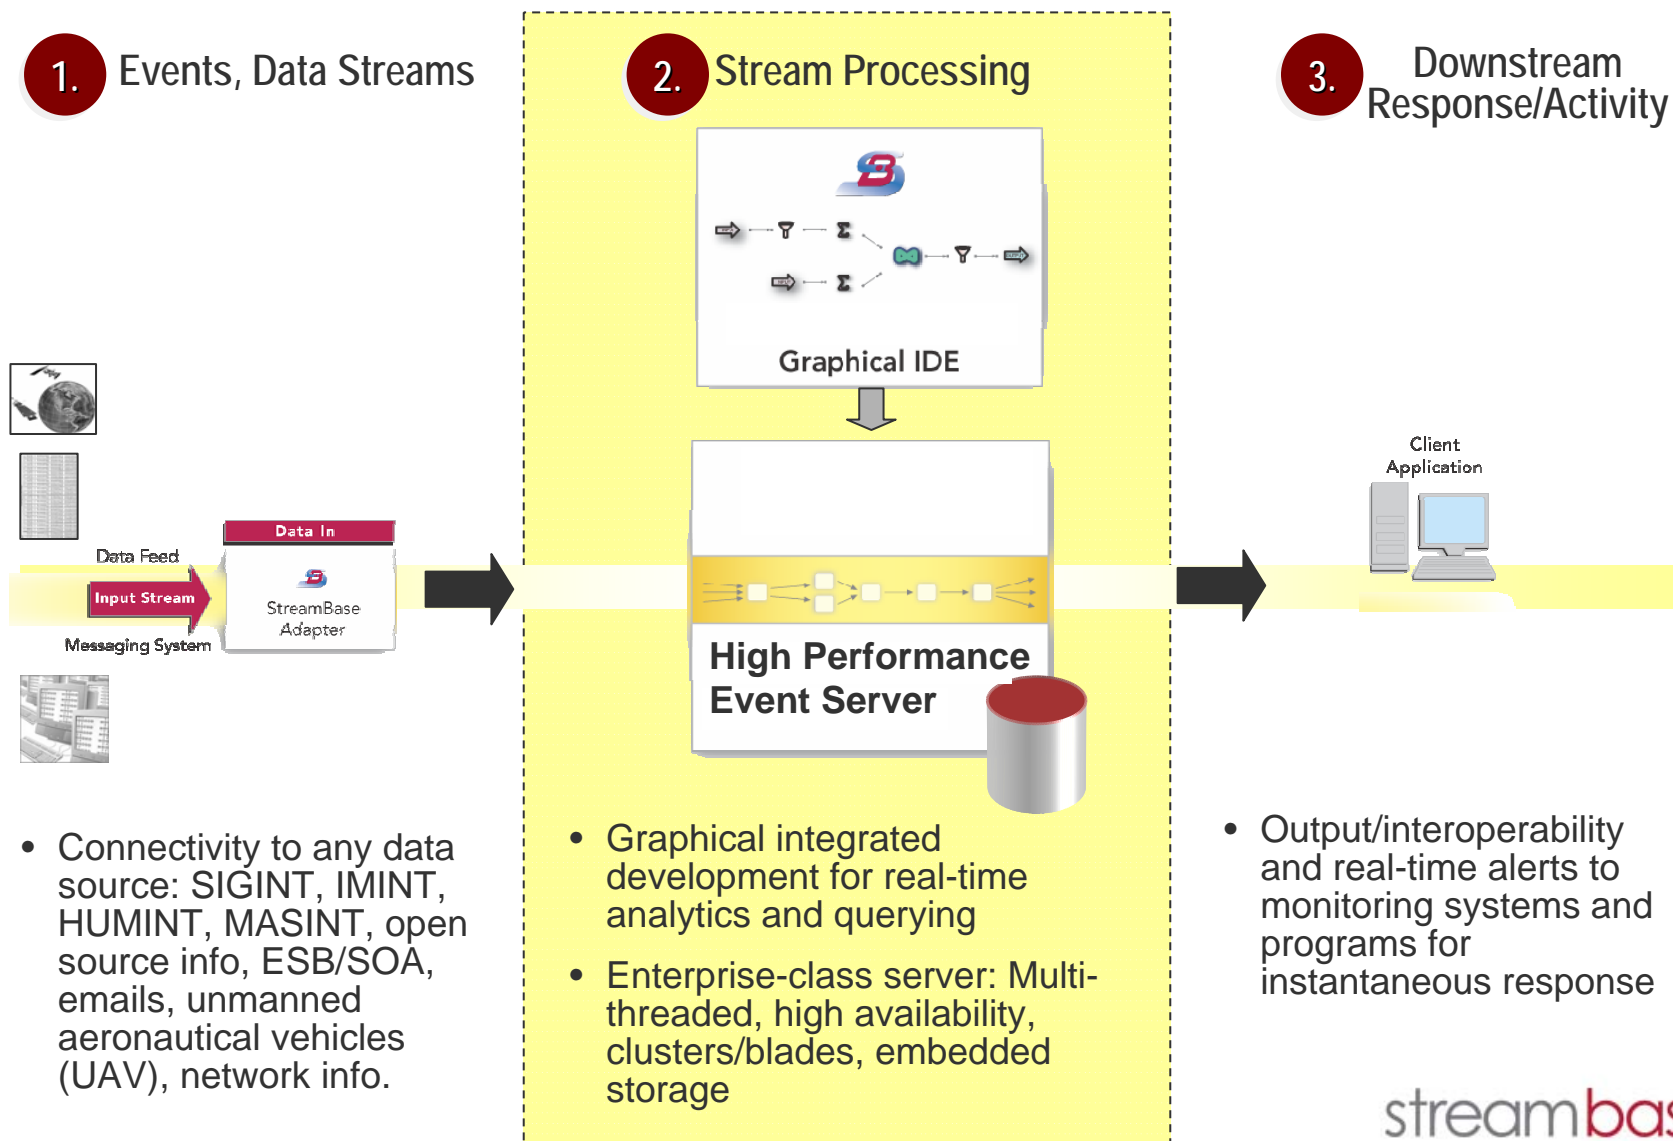
Enterprise Connectivity

- Historical/stored data integration
- Broad set of data and messaging adapters
- Real-time interactive dashboard connectivity
- Rapid adapter development toolkit
- Java, C++, .NET support

streambase®

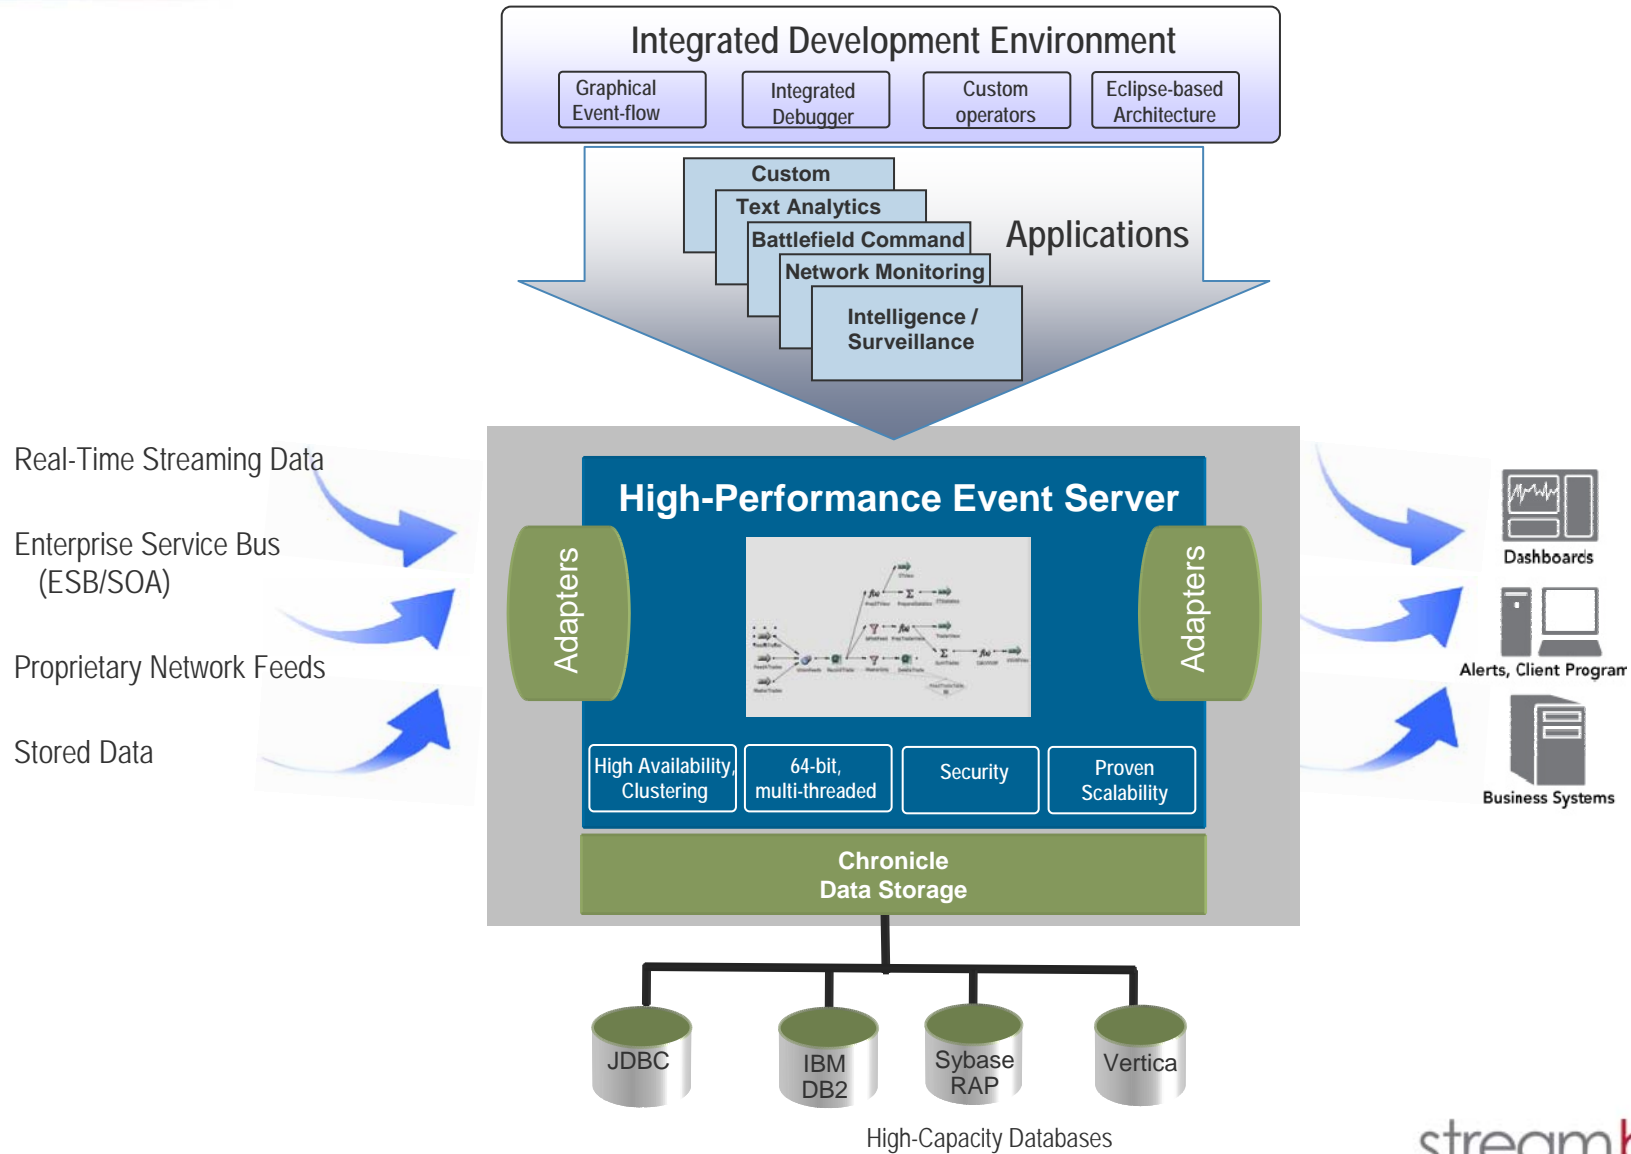


Event Processing Architecture





StreamBase Platform Architecture





Rapid Application Modeling and Development

- Use GUI to quickly model and build apps
 - Confirm requirements met with end-user
 - Graphically build full app, compile, and run on StreamBase Server
- End-to-end development
 - Adapter wizards
 - Stream Record/Playback
 - Feed simulator
 - Debugger, performance monitor
 - Ecosystem of Eclipse plug-ins
 - Version source control, graphical UI development, task management
 - Interfaces with dashboards, visualization tools

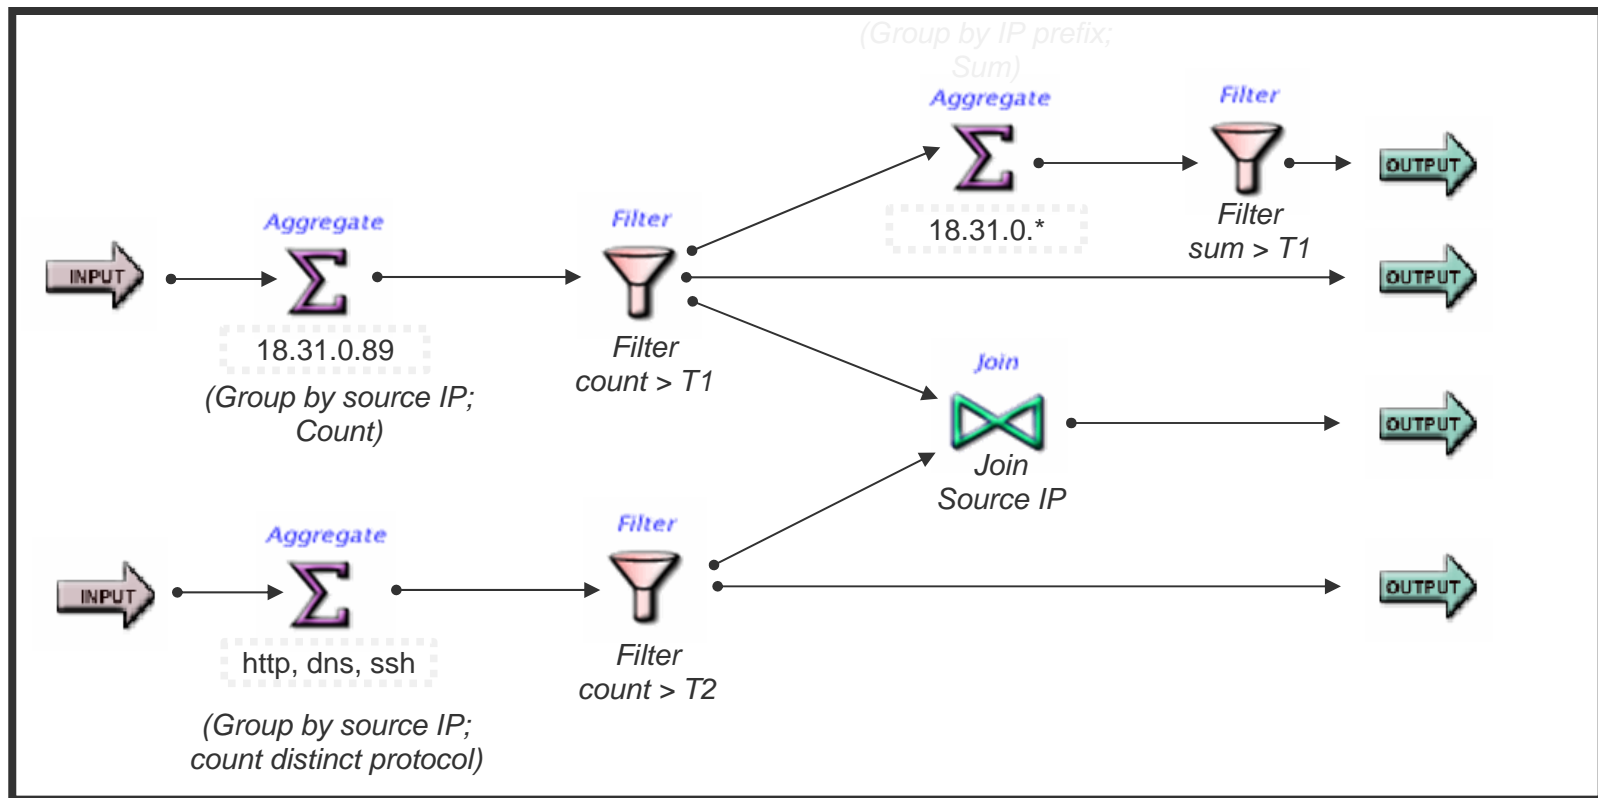
The screenshot displays the StreamBase Studio interface. On the left, a palette contains various operators and data constructs. The main workspace shows a graphical data flow diagram with components like 'InputStream', 'Filter', 'Aggregate', and 'OutputStream'. A text box overlay states: 'Integrated environment for building, testing, and deploying StreamSQL applications: choice of a graphical or text-based application editor'. On the right, a SQL editor window shows the following code:

```
1 CREATE INPUT STREAM InputStream (  
2   SourceIP string(15),  
3   DestIP string(15),  
4   DestPort string(4),  
5   TransactionTime timestamp);  
6  
7 -- Add local time to the stream.  
8 SELECT *, now() AS LocalTime FROM InputStream;  
9  
10 -- Create a one minute aggregate window the  
11 -- hits in that minute.  
12  
13 CREATE STREAM CountryBySourceIP AS  
14   SELECT SourceIP, count(*) AS Count,  
15   firstval(TransactionTime) AS StartTime,  
16   FROM InputStream (SIZE 60 ADVANCE 10 ON  
17   GROUP BY SourceIP);  
18  
19 -- Send an alert when there's > 10 source I  
20 CREATE OUTPUT STREAM AlertTooManySourceIPFI  
21   SELECT * FROM CountryBySourceIP  
22   WHERE Count > 10;  
23  
24 -- Get the source subnet from SourceIP field  
25 -- except for SourceIP.  
26 CREATE STREAM FindSubnetForSource AS  
27   SELECT DestIP AS DestIP, DestPort AS D  
28   TransactionTime AS TransactionTime,  
29   subnet(SourceIP, 0, lastindexof(Source  
30   AS SourceSubnet FROM InputStream;  
31  
32 -- Create 60 second aggregate window that t  
33  
34 CREATE STREAM CountrySourceSubnet AS  
35   SELECT SourceSubnet, count(*) AS Count,  
36   firstval(TransactionTime) AS StartTime,  
37   FROM FindSubnetForSource (SIZE 60 ADV  
38   GROUP BY SourceSubnet);
```

streambase®



Network Intrusion Detection Example



Network intrusion detection query inspired from Autofocus and Snort : detecting most active (clusters of) sources or those that connect over most protocols (10 second window)



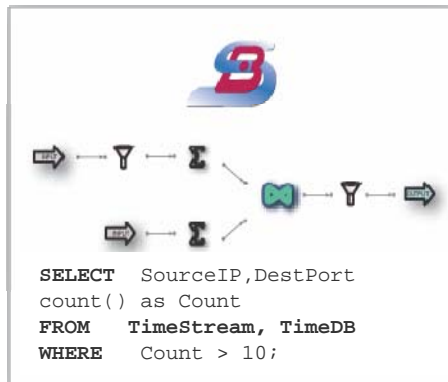
Real-Time Analysis Operators and Constructs

- Time-based or event-based windows
 - To deal with analytics on moving objects
- Stream disorder
 - Late, out-of-order data from networks, intelligence sources
- Integrate application logic and messaging
 - Run an application on multiple machines for scalability and high data rates
 - User-defined functions (Java or C++) and aggregates for secret/proprietary algorithms
- Persistence
 - Trigger real-time alerts in context of historical data
 - Embedded SQL DBMS (for small to large state)
 - JDBC connection (for huge state)
 - Optimized interfaces to high performance databases and data warehouses




StreamBase: Integrated Platform with Historical / Real-time Processing

Common Integrated Development with StreamSQL



StreamBase Real-time Stream Processing





High-volume, IP log data
generated in real-time


[illegible]

Real-time Analysis:
Powerful real-time analysis of high-volume log data (IP address, destination port access)



Comparison with Historical Data:

Integration with stored network intrusion data for better real-time decision in the context of historical network attack information



Real-Time Execution:

Trigger appropriate alert,
initiate lock-down





Example: Locate Terrorist Cells or High-Level Terrorist Target

1 High-volume
real-time data
streams



2

Consolidate
data from
disparate
sources

3

Filter & analyze real-
time data streams for
targets of interest



4

Trigger appropriate real-
time alert or response.



streambase®




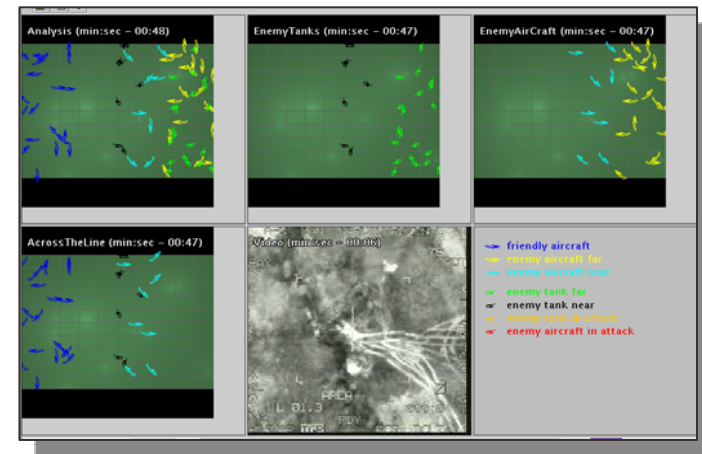
Use Case: Battlefield Monitoring, Visualization with StreamBase

■ Client Scenario

- Government contractor required consolidated reports and data from reconnaissance aircraft of friendly and enemy activity

■ Solution

- Real-time collection, dynamic reprioritization, and delivery of movement data “as it happens”
- Critical alerting established to pinpoint any/every enemy movement



Example of combat military monitoring of friendly and enemy forces in real-time with StreamBase



What People are Saying About StreamBase.



"StreamBase is a breakthrough technology that allows organizations to process and analyze huge data volumes and ensures that critical information doesn't drop on the floor."



StreamBase Systems' Stream Processing Engine is a sophisticated event management product designed for time-sensitive applications, where pattern detection and response must occur in milliseconds."



"We focus on two things - speed and agility. The combination of high performance and the ability to create and modify applications quickly were key factors in selecting StreamBase"



"The Homeland Security Department could use the software to monitor multiple streams of sensory data... The Defense Department could deploy the software to monitor data from vehicles, soldiers' uniforms and other resources, showing positions and conditions of troops and equipment."

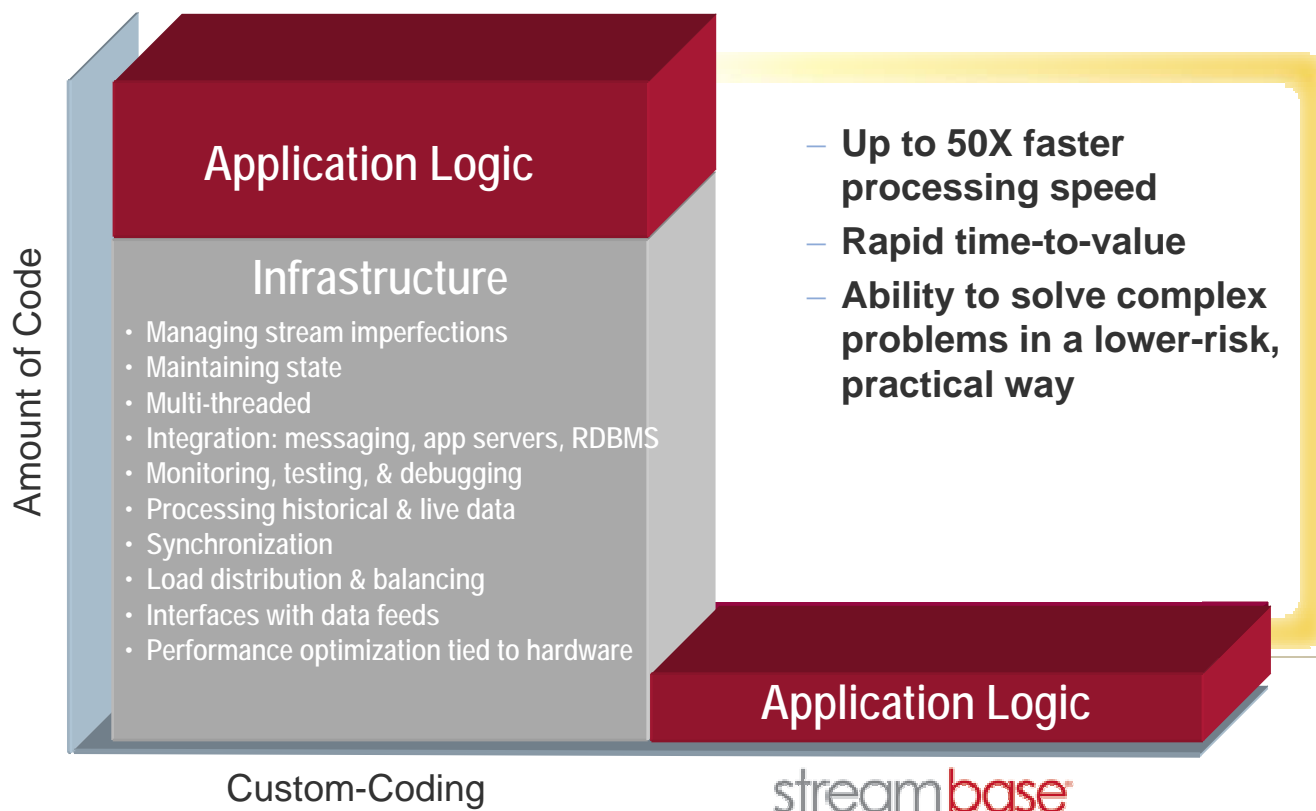


"...used to build real-time applications that previously took months or years to craft by hand. It will also be used to build real-time applications that were not feasible before."





Competitive Edge, Improved Security With StreamBase





The StreamBase Advantage

- **Productivity**
 - 10X faster to build and deploy apps
 - Fastest learning curve
 - Easiest to modify and maintain
- **Performance**
 - 5-10X+ faster performance
 - As much data as you've got
 - Near-zero latency
- **Integration**
 - Flexible interfacing to multiple data streams, messaging systems, apps, and databases
 - Connectivity with market-leading financial services data feeds
 - Rapid time to success with lower project risk
- **Maturity**
 - Advanced enterprise capabilities from security to high availability
 - Deployed in a variety of mission-critical systems across multiple industries



Taking Action

- Joint architectural workshop
- Proof-of-Concept
- Download StreamBase Developer Edition software at www.streambase.com
- Contact StreamBase:

Gary Schrader, Director of Federal Operations
StreamBase Systems, Inc
Phone: 703-608-6958
gary.schrader@streambase.com



streambase®