Altair Panopticon™

Optimize Electronic Trading Performance Across Financial Markets

Using Panopticon with the Confluent Enterprise Kafka Platform

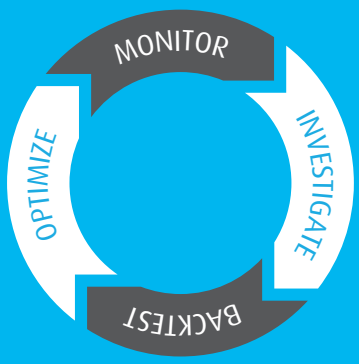

Industry Examples

- **Buy-side:** Monitor and analyze best execution, profitability, risk, and more across brokers and execution venues in real time.
- **Sell-side:** Monitor and analyze best execution, latency, flow, MiFID II and MiFIR compliance, profitability and more in real time.

Improve profitability and ensure regulatory compliance for electronic trading operations

- Compare real-time order and trade flows against historic flows and external market
- Program alerts based on trading anomalies derived from complex calculations
- Fast insight into trading performance and investigation into anomalies
- Optimize performance during trading day; react immediately to market events
- No coding – no developers
- Reduce development time for new analytics dashboards to minutes
- Let traders build their own analytics
- Scale up and out quickly to support fast reactions to spikes in volatility
- Publish new analytical systems to all desks/offices in seconds

Real-Time Transaction Cost Analysis

This dashboard enables traders, quants, and salespeople to monitor best execution and analyze profitability, including transaction costs, in real time. They can respond immediately to market events and threats to customer order flow execution and firm profitability during the trading day.

Risk Sensitivity Monitoring & Analysis

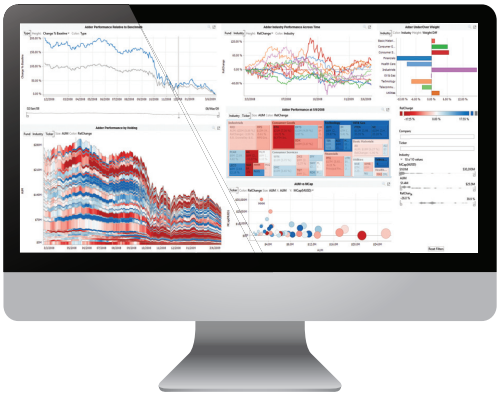

Risk Analysts use visual analytics like this to monitor how much risk the firm is taking on at any given time at the office, desk, and individual trader levels. They can set alerts to notify them when risk sensitivities exceed pre-determined thresholds, and use the visuals and filters to identify the causes of any anomalies in seconds.

Learn more:
altair.com/panopticon

Legacy Architectures Cannot Handle Current Requirements

Capital markets firms continue to use a wide range of systems that cannot support the requirements generated by new regulatation and increasingly competitive markets.

- Ranges from end-of-day reporting only to use of proprietary CEP engines for intra-day monitoring, typically requiring significant custom development
- Applications are written for legacy infrastructures
- Performance analysis typically generated after completion of trading as a batch process
- Heterogenous environments with many software components performing different steps of batch processing
- High availability and replication across office is an afterthought or non-existent
- Development of new capabilities is time consuming and brittle – performed by specialist developers (typically contractors)
- High turnover of development and support staff creates headaches
- Systems often suffer from performance problems, including lost data and failures, during volatility spikes
- Changes to any aspect of system often involve long duration projects and significant re-writing / refactoring
- Changes to analytical processing typically require redeployment of the whole system
- Until the introduction of this platform, moving to real-time analytics from end-of-day batch processes required whole sets of proprietary systems, specialist engineers, and so on.

Schematic – Current (Simplified)

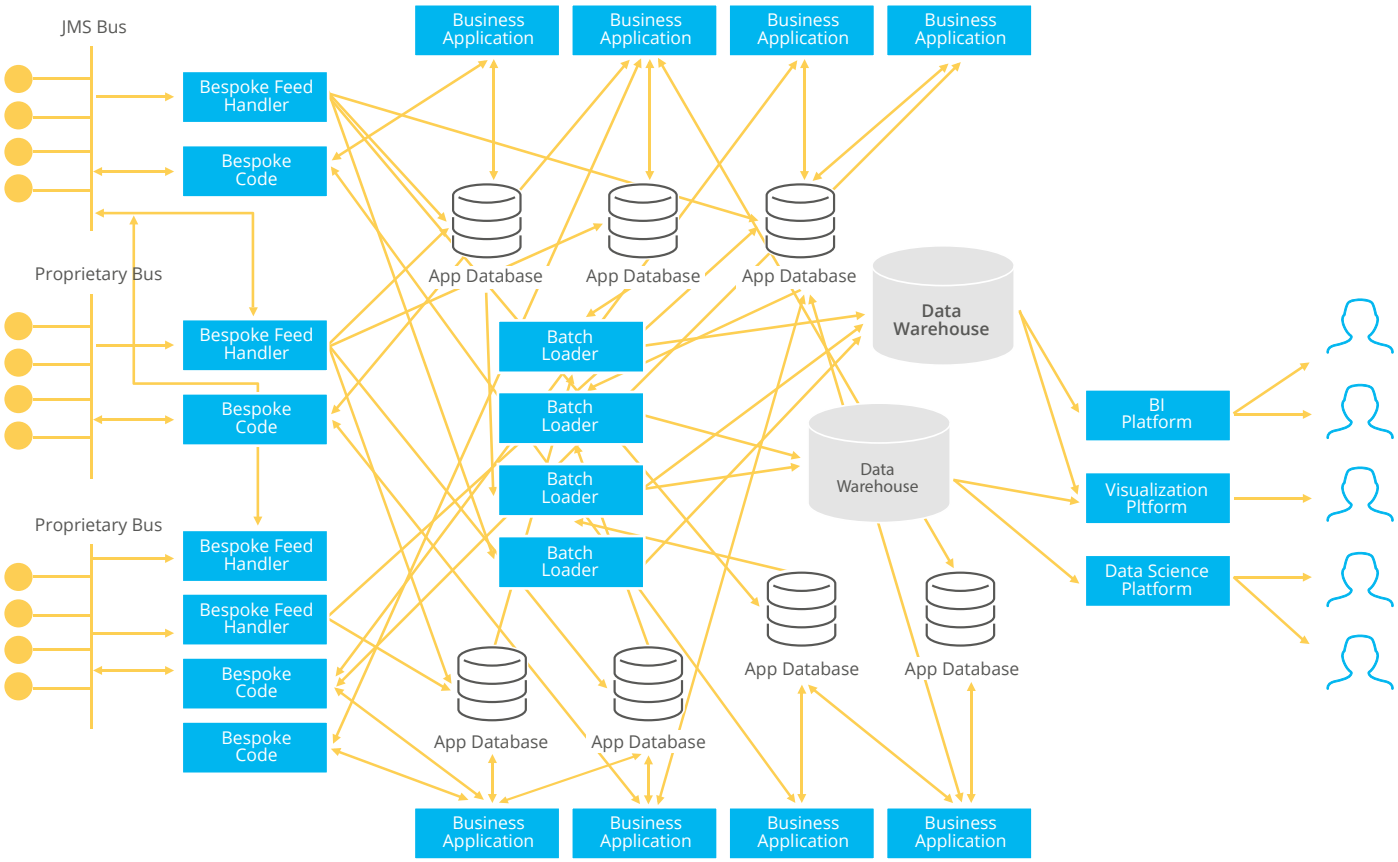

A New Paradigm: Combine Confluent Enterprise with Altair Panopticon to Create A Flexible, Cost-Effective, and Efficient Trading Architecture

- Deploy the Confluent Streaming Platform
- Ingest trading analytics source data streams into the Confluent Platform
- Deploy Panopticon Stream Processing
- Deploy Panopticon Visual Analytics
- Provide Steaming Analytics training to the people working the electronic trading desks
- Traders draw and deploy initial Steam Processing dataflow
- Traders draw and deploy initial Visual Analytics dashboards
- Provide Confluent Enterprise training to DevOps staff
- Scale the solution globally using the Multi Data Centre Replication features of the Confluent Enterprise Platform
- Use the Confluent Enterprise Platform to monitor the solution globally with the Confluent Control Center
- Roll out the Visual Analytics platform to all office locations and disaster recovery sites
- Iterate, expand, and deploy Stream Processing applications
- Iterate, expand, and deploy Visual Analytics dashboards

Target Data Center Architecture

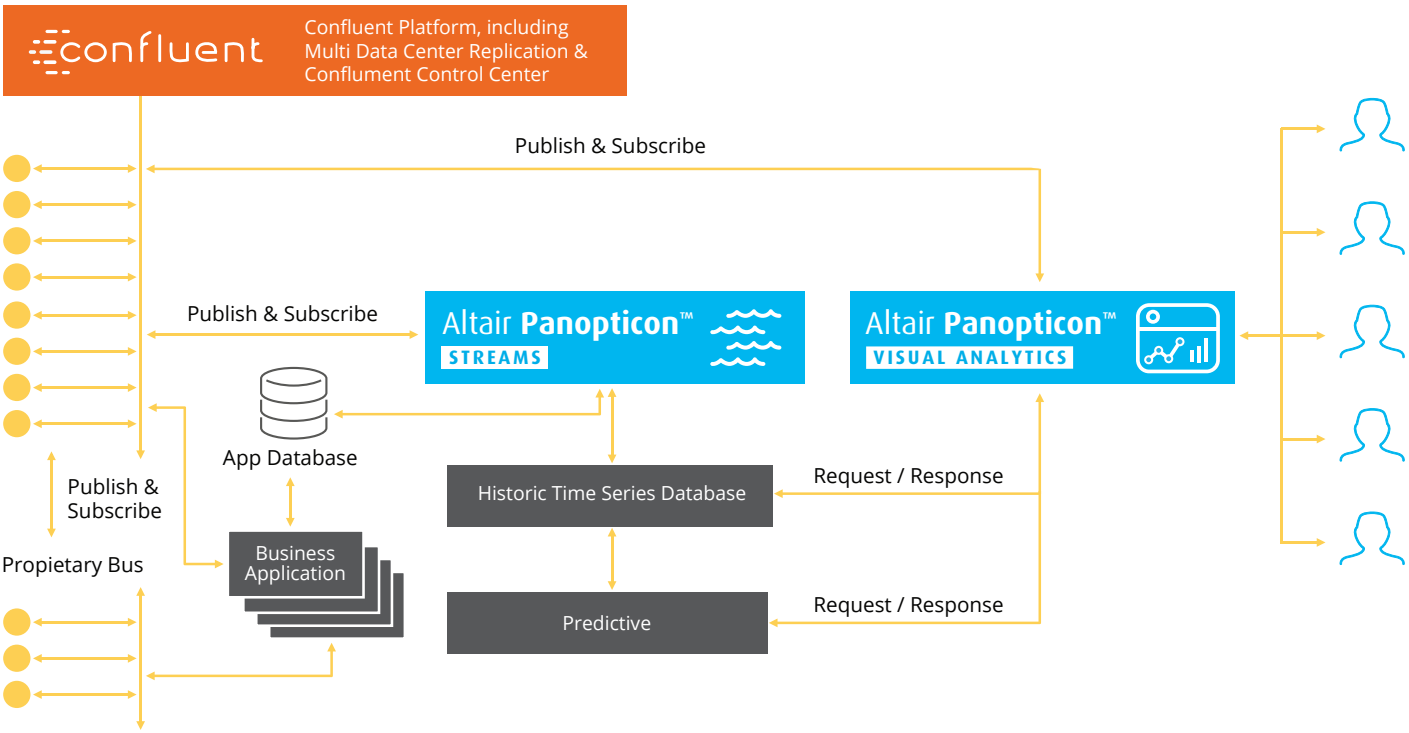

Target Data Center Architecture: A Global View

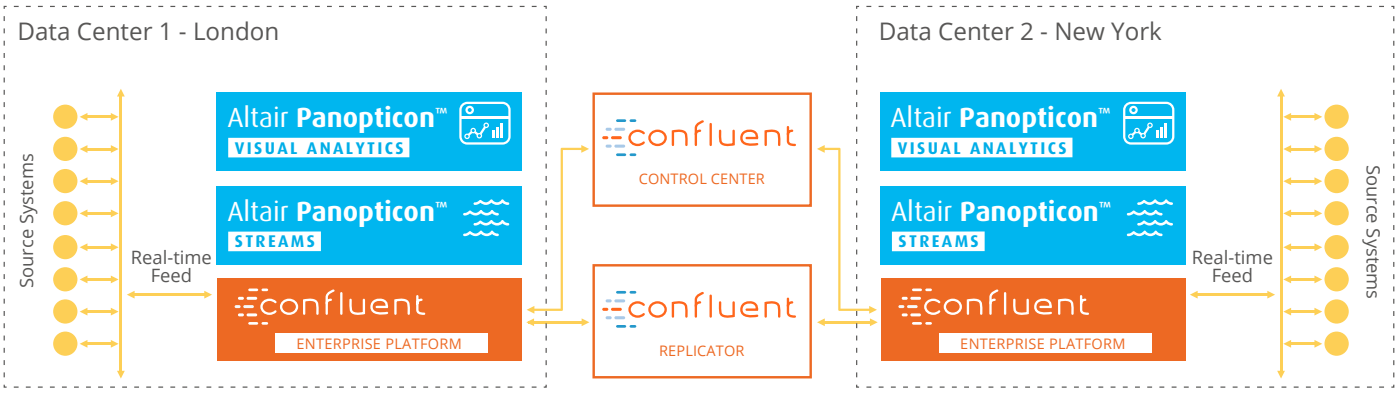

Why Do It Now?

- **Increased demand for real-time visibility into electronic trading activity**
  - Constantly evolving and increasingly strenuous compliance pressure, particularly around best execution (MiFID II, MiFID III, MiFIR, etc)
  - Margin pressure in super-competitive HFT business driving cost cuts, reducing appeal of legacy technologies and large development staffs
  - Customer service based on real-time insights is a competitive advantage
- **Real-time insight into trading activity helps improve profitability, avoids compliance problems, and increases operational efficiency**
- **Improved ability to react quickly to market events and changes in customer behavior to take advantage of profitable opportunities and reduce chances of breaches**
- **Reduced delivery risk: Traders, quants, and compliance officers can build and deploy their own systems, refine them, and then publish them worldwide without cumbersome project plans, time consuming communications with offshore programming and testing teams, and expensive development cycles**

What Next?

1. **Discovery of client's existing infrastructure, data sources, and streams**
2. **Define Proof of Concept (PoC) project, including success criteria**
   a. Determine required Stream Processing and Visual Analytics capabilities
   b. Determine data requirements, including sources and destinations
   c. System integration requirements
   d. Describe specific application scenarios
3. **Implement PoC**
   a. Starts with one week of onsite work, including initial training
   b. Additional three weeks of regular check-ins as customer validates PoC system against success criteria
4. **Purchase**
5. **Implementation and training for first location**
6. **Rollout to global offices**

Final Deployment Size at a Typical Electronic Trading Customer

- **Six offices (London, New York, Chicago, Hong Kong, Tokyo, Sydney)**
- **Three key offices require Disaster Recovery sites (London, New York, Tokyo)**
- **Hundreds of active users in each location**
- **Throughput of 100 million+ messages per day**
- **Schema registry deployment (Master/Slave replication)**
- **Multi-data center Confluent Enterprise Platform deployment (Active/Active between sites)**
- **Global office replication**
- **Panopticon Streams deployment (clustered for high throughput)**
- **Panopticon Visual Analysis deployment (high availability required in each office)**