

ALTAIR® PANOPTICON™: COMPREHENSIVE DATA VISUALIZATION AND STREAMING ANALYTICS

Panopticon provides business users, analysts, traders, and engineers with the monitoring and analysis tools they need to conduct successful, profitable operations while maintaining a close eye on anomalies, trends, clusters, and outliers. They can make insightful, fully informed decisions based on massive amounts of fast-changing data.

Panopticon is enterprise-class software you can deploy in the cloud (public or private) or on-premises. It connects directly to virtually any data source, including big data sources, SQL and NoSQL databases, flat files, and real-time message queues. Users can develop stream processing applications and design sophisticated visual user interfaces without writing a single line of code.



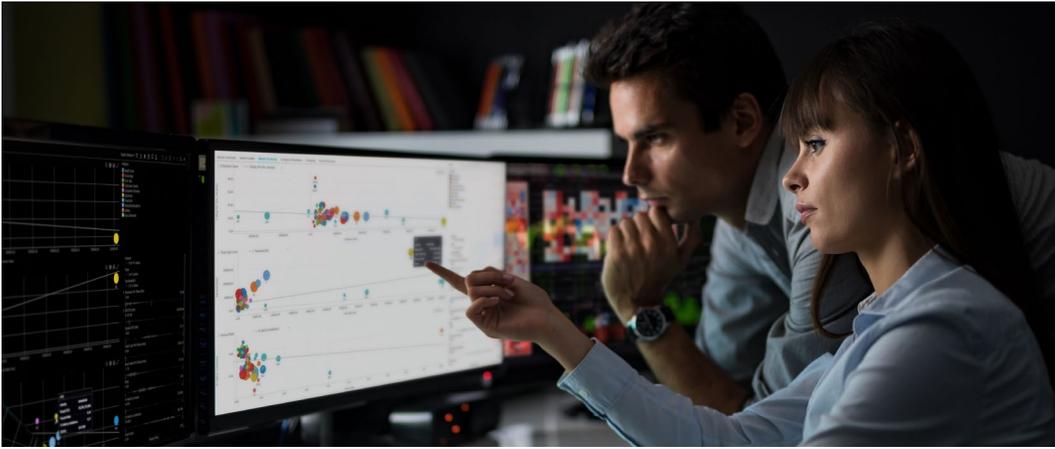
REAL-TIME VISIBILITY INTO YOUR OPERATIONS

Delays in decision-making are costly. Waiting makes it likely you will be late to take advantage of profitable opportunities and avoid threats. The right [data visualizations](#) combined with clean, enriched data help decision-makers focus on the most relevant and important information without losing sight of the big picture.

Panopticon's filtering tools enable users to zoom in and out on the timeline, remove false positives from the screen, and focus on exceptions. Users can solve difficult problems quickly, understand complex relationships in seconds, and identify issues requiring further investigation with just a few clicks.

Panopticon combines:

- Intuitive visual data analysis with high density displays
- Comprehensive filtering and alerting enabling efficient peer comparisons and identifications of correlations, trends, exceptions, and anomalies
- Ability to connect directly to virtually any data source
- Fast UI-based configuration and development with no coding



Panopticon users work in a wide variety of industries including finance, automotive, heavy equipment, and energy.

In addition to data visualization, we offer a stream processing engine that connects directly to real-time streaming and historic time series data sources, including MQTT, Kafka, Solace, and many others. Users can build event processing applications that combine streaming data with historic data, calculate performance metrics using a wide variety of statistical and mathematical functions, aggregate, conflate, compare data sets, and automatically highlight anomalies against user-defined thresholds.

Enterprise Class Data Exploration and Reporting

Understand the structure, interrelationships, causal links, and outliers in your data. Get a comprehensive view of your operations from every angle. Develop and share dashboards and reporting screens securely throughout the enterprise.

Focus on Outliers

Data visualizations help decision-makers focus on the most relevant and important information without losing sight of the big picture.

Leverage All Data Sources

Altair visual analytics software includes native connectors for virtually any data source, from big data repositories and real-time streaming sources to SQL and NoSQL databases and flat files.

Code-free Development

Build analytical dashboards and stream processing applications with a point-and-click user interface. Respond quickly to changing business requirements, reduce risk and cost, and deploy new applications and dashboards in minutes.

Advanced Visual Analysis

Create smart dashboards and generate charts automatically based on data source keywords. Visualize complex datasets using automatic machine learning and predictive models, including “best fit” or “bring-your-own” models, without leaving the user interface.

Visualize Real-time Streaming Data

In addition to standard historical data sources, Panopticon can accept, process, and visualize true real-time streaming data with nanosecond accuracy. Real-time data is a stream of continuous timestamped messages delivered in sequential order — usually transported on a message queue like Kafka or MQTT. In order to make use of such fast-changing data, the system must be able to subscribe to updates as they appear. This is completely different than the traditional database model which requires the data to first be stored and indexed before being “pulled” from the repository into a system for processing. Panopticon’s real-time functions allow you to use data while it is in motion through the server.

Panopticon can handle high velocity/high volume data from multiple streams and make on-the-fly comparisons with historical information stored in high performance in-memory time series databases — down to the nanosecond timestamp if needed. They can rewind and play back data streams at any speed — in real-time, faster than real-time, or slower than real-time.

Panopticon’s capabilities enable an organization to process and visualize large volumes of data quickly enough that decision-makers can react to changing conditions as they happen. For example, they can take immediate action when a dashboard reveals potentially fraudulent activity or other threats. They can jump on opportunities to improve operations or profits as soon as they occur.

Proactive Alerts

Panopticon proactive alerts make it easy to highlight anomalies in trading and/or market activity. You can set up a new alert in a few minutes and specify whether the designated users will receive alerts on their screens, via email, and/or with sounds.

Define proactive alerts using data from streaming data sources, including CEP engines, real-time message queues, and/or historical trading data stored in virtually any kind of database. You can select which users will receive each alert so only the people who are authorized to take corrective action will be notified.

STREAM PROCESSING WITH NO CODING

Panopticon Streams is the stream processing engine that works with Panopticon Visual Analytics software to form the Panopticon Streaming Analytics platform. Streams connects directly to a wide range of streaming and historic sources, including Kafka, and supports these critical functions:

- Real-time data prep: Combines streaming data with historic data
- Calculation engine: Calculates performance metrics based on business needs
- Aggregation engine: Aggregates data as needed
- Alerting engine: Highlights anomalies against user-defined thresholds



You can deploy Panopticon quickly — in days, not months — and new users can be building their own customized dashboards and streaming analytics applications within a few hours.

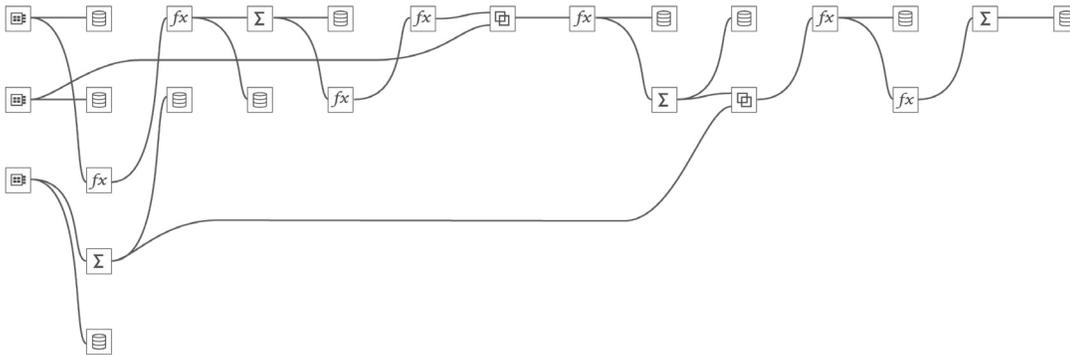
Built on Kafka

The underlying platform for Panopticon Streams is Kafka. It supports all the benefits of Kafka, but without its complexity. Users don't need to know how to write a single line of code in Java, Scala, or even KSQL. Similarly, firms don't need to deploy a proprietary, legacy event processing platform. They can leverage their existing investment in Kafka and get started immediately.

Build Complex Data Flows In a Web Browser

Panopticon Streams allows you to build stream processing applications that:

- Subscribe to streaming data inputs, including Kafka streams and others
- Retrieve from historic and reference data sources
- Join data streams and tables
- Aggregate streams within defined time windows
- Create calculated performance metrics
- Filter streams
- Branch streams
- Union and merge streams
- Pulse output
- Create alerts based on performance metrics against defined thresholds
- Output to Kafka or email, or write to databases including kdb+, InfluxDb, or any SQL database
- Integrate with your infrastructure (LDAP, SSO, AD)



Build event processing applications in minutes using only a web browser and without writing any code. Users who understand the business problems to be solved create data flows using a visual representation of their logic.

LEVERAGE ALL DATA SOURCES

Panopticon lets you extract real value from your existing data infrastructure without implementing new middleware or developing custom code. Panopticon offers native data connectors for a huge variety of sources, including:

- Message buses: MQTT, Solace, ActiveMQ, Google Cloud PubSub, JMS, Kafka, AMQP, RabbitMQ, AMPS
- CEP engines: TIBCO StreamBase, SAP ESP, OneTick, Kx kdb+tickAggregation engine: Aggregates data as needed
- Tick/time series databases: Kx kdb+, OneTick, InfluxDB, Vertica, Timescale DB
- NoSQL: Cassandra, MongoDB, elastic, Spark, Hadoop, Splunk
- Relational databases: Any SQL database, including Oracle®, Microsoft® SQL Server®, MySQL, PostgreSQL, and MonetDB
- Big data: Amazon Redshift, Amazon DynamoDB, Google BigQuery, Hadoop
- Cloud: Google Drive, Microsoft OneDrive, web services
- Flat files: Json, SVG, Text, Excel/CSV, XML
- Python and R



MULTIPLE DEPLOYMENT OPTIONS

Implement Panopticon at the desktop, workgroup, or enterprise level or embed Panopticon into your own applications.

Streaming Analytics in the Cloud

Deploy Panopticon on AWS, Azure, Google Cloud, and Oracle Cloud. It is also fully compatible with containerized cloud implementations using Docker.

Multi-tenant

Clients needing a “zero footprint” implementation can use Panopticon’s multi-tenant capabilities. This supports lower total cost of ownership (TCO), unlimited scalability, and automatic and ongoing updates with no client intervention required without sacrificing configurability and customization options.

On-premises

Many Panopticon clients prefer to deploy the software on-premises in implementations protected by their own firewalls. This gives the IT team complete control over every aspect of their system’s data security, integrations, and reliability.

Embed

Panopticon is designed for embeddability and we have numerous OEM partners and clients who have made Panopticon an integral component in their own applications. Embedding Panopticon provides users with visual analytics capabilities within the context of a business applications. Panopticon making embedding easy with support for single sign-on (SSO), automated embed link generation, and cross-origin resource sharing (CORS).

APPLICATIONS

Our customers use Panopticon to support the continuous optimization of their operations. Business users, engineers, and analysts can connect to data sources, develop stream processing applications, perform on-the-fly data preparation, and build comprehensive dashboards without writing any code. Panopticon makes it easy to spot outliers, trends, clusters, and anomalies in massive amounts of fast-changing data.

Internet of Things (IoT)

Business users, analysts, and engineers can access, process, monitor, and analyze the incredible volume of real-time information streaming in from connected devices, vehicles, equipment, and sensors in the fast-evolving IoT world.

Manufacturing

Panopticon can process real-time sensor data on heat, speed, pressure, and other factors generated by any number of production machinery, make on-the-fly comparisons with historical data, and amplify the accuracy of the machine learning algorithms in remaining useful life (RUL) analysis. It enables manufacturing engineers to identify unusual behaviors or patterns in equipment using sensor data that can prevent small glitches from creating major operational problems. Panopticon can flag potentially serious issues without reducing overall equipment effectiveness (OEE).

Energy

Oil and gas producers, pipeline companies, and electrical grid operators require up-to-the-second views of exactly how systems are performing. In most circumstances, the faster they can respond to an issue, the more efficiently they operate. Specific use cases include oil and gas production analysis, grid utilization monitoring, profitability analysis, capacity analysis, project performance monitoring, and customer service monitoring for call centers.

Telecommunications

Panopticon gives telecom network operators effective ways to correlate and analyze data streaming in from routers, towers, control systems, and third-party systems to make insightful, informed decisions. They can spot small problems before they become major outages, identify maintenance shortfalls, monitor and respond to network and infrastructure issues in real time, and proactively detect and prevent revenue leakage and fraud.

Retail Banking

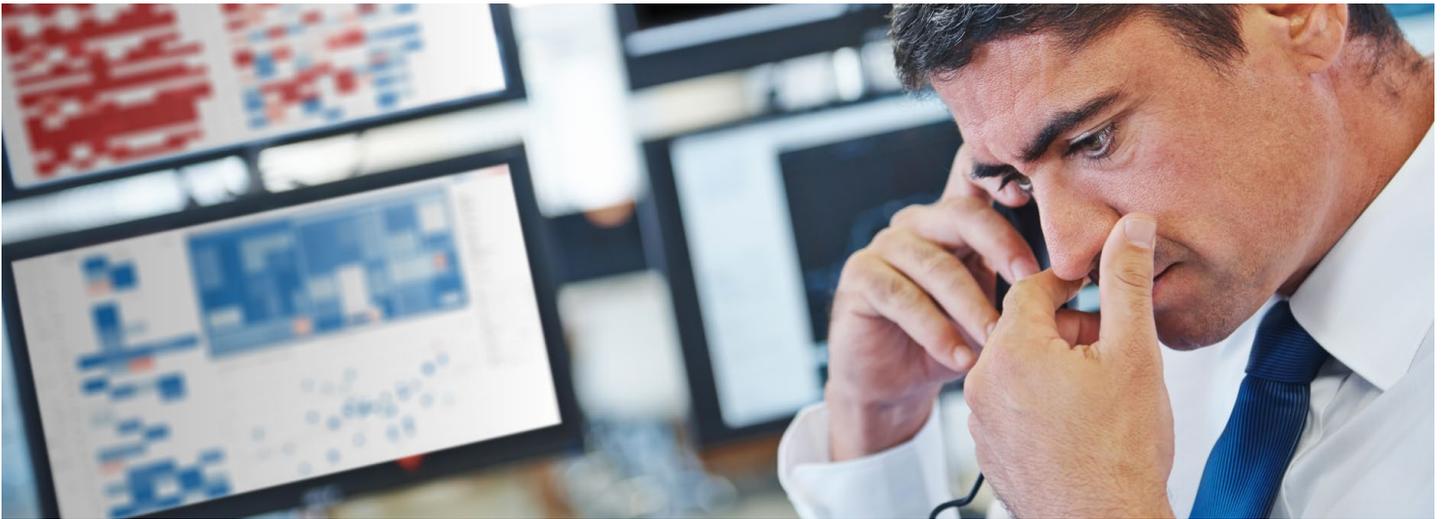
Visualize credit applications, loan performance, risk, exposures, currency and interest rate fluctuations, and asset prices. Examine credit card utilization and customer behavior statistics to identify potential fraud. Improve customer service by aggregating and analyzing telephone, email, and chat interaction data.

Securities Trading

Buy-side and sell-side firms use [Panopticon](#) to monitor and analyze best execution, TCA, latency, client flow, MiFID II and MiFIR compliance, profitability, and more — all in real time. Traders and managers can analyze market conditions and activity to make smart timing and execution decisions and develop insights into the pricing characteristics of all liquidity providers. Panopticon enables them to understand trading behaviors for all clients across all venues.

Panopticon's unique capabilities enable traders, quants, and compliance people to build their own trading analytics systems without writing any code. They can respond to fast-changing market conditions by visualizing real-time orders and RFQs and conduct analysis of historic data to back test trading strategies.

- **Profitability:** Correlate trading costs and risks with revenues by trader, office, asset class, execution venue, or any other dimension. Measure trading effectiveness and see performance trends for all traders and groups. [Monitor trends](#) for RFQ-driven trading activity across all electronic trading channels.
- **Liquidity and Spreads:** [Analyze market microstructure](#) for RFQ-driven markets. Correlate market impact to historic executions, venues, and trading strategies. Compare order and quote data and identify temporary market liquidity imbalances across all venues. Visualize consolidated orders across all available liquidity pools and identify liquidity imbalances.
- **Transaction Cost Analysis (TCA):** [Analyze best execution](#) across fragmented liquidity pools, and through all algo types, books, venues, counterparties/brokers, and instruments.
- **Risk:** [Perform intraday analysis](#) of risk profiles at the office, desk, and trader levels. Analyze liquidity in compliance with BCBS 248 rules and identify value at risk (VaR) and sensitivity trends across all book and instrument hierarchies.
- **Compliance and Surveillance:** [Analyze historic order flow](#) for execution policy breaches, view live execution performance, and compare execution histories to past activities and market data across all liquidity and execution venues in use.
- **Basket Trading:** Analyze portfolio and index basket risk and return analysis, both pre-trade and during basket execution, and identify liquidity spikes, shortfalls, and performance divergences.
- **Trading Fraud:** [Identify cases](#) of potential spoofing, quote stuffing, wash trading, and other fraudulent activity. Playback through a series of trades tick-by-tick to gain a full understanding of exactly what happened.



FIVE EXAMPLE USE CASES

Companies all over the world use Panopticon to improve profitability, reduce the cost of regulatory compliance, decrease operational risks, and identify potential failures before they can impact the business.

Electronic Trading: Client Flow Analysis

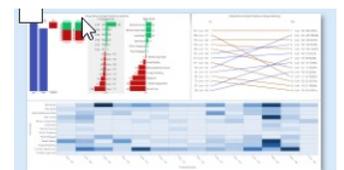
Panopticon gives firms visibility into their customers' trading behaviors. [Portfolio managers](#) can review historic activity and analyze trading patterns for anomalies. They can move away from blotters of RFQs that limit visibility to a comprehensive view of their flow — by book, instrument, client, time period, or any combination of these parameters. Traders can compare today's live cumulative flow to historical averages to identify problem performance, analyze historic customer flow for trends and unusual performance, play back through past activity to identify problem areas for future optimization, and analyze flow across different categories.



Electronic Trading
Client Flow Analysis

Manufacturing: Automated Production Planning

[Automotive manufacturers](#) streamline the assessment of design changes on vehicle weight, risks at all stages, and costs using key performance indicators (KPIs) developed using Panopticon. They integrate, process, and visualize synchronized data from product lifecycle management (PLM), manufacturing resource planning (MRP), enterprise resource planning (ERP), and other databases to develop KPIs that the product management, vehicle integration, weight/mass balance, and cost management teams use to make informed decisions on design changes. This approach enables the teams to examine data from multiple platforms via a single pane of glass with live data.



Manufacturing
Automated Production Planning

Energy: Sensor Data Analysis

Oil wells are equipped with a variety of sensors measuring temperature, torque, friction, pressure, location, depth, flow, and rate of penetration. This data can be streamed into a control room over a real time message bus where engineers can use [Panopticon](#) to monitor that data in real-time. Stream processing algorithms make on-the-fly comparisons between current data with recent trends and historical records and alert users when potential maintenance issues occur before they can bring drilling to a halt.



Energy
Sensor Data Analysis

Manufacturing: Real-Time Equipment Condition Monitoring

Manufacturers can use [Panopticon](#) to generate a real-time view of how all plant equipment is operating. This helps reduce incidences of unplanned downtime and increases overall equipment efficiency with lower costs. In a typical implementation, IIoT sensors and platforms feed real-time streams of data on vibration, load, power consumption, pressure, temperature, production volumes, and other operating parameters to Panopticon over MQTT. In many cases, Altair® Knowledge Studio® is incorporated into the solution to apply advanced machine learning (ML) and artificial intelligence (AI) to the data in order to calculate probability of failure (PoF) and identify machine health trends involving multiple components and/or machines. Proactive alerts notify engineers of threshold breaches and other conditions requiring immediate attention.



Manufacturing
Real-Time Equipment
Condition Monitoring

Logistics: Performance Monitoring for Truck Fleets

This screen is visualizing real-time feeds coming in from a fleet of commercial vehicles. [Data streams](#) containing information on each vehicle's operations, including fuel consumption, time spent idling, actual speed compared to posted, accidents, and route deviations are brought into the system on real-time message buses. The system automatically makes comparisons with historical data and compares the performance of each vehicle and driver with their peers to flag outliers.



Logistics
Performance Monitoring
for Truck Fleets

To learn more, please visit altair.com/panopticon

“The only way to build compelling value-added analytical applications is to bring historical, near real-time and real-time data together. Panopticon, our cloud native visualization technology, is the only technology that enables threading of data at various frequencies to create BI dashboards, real-time dashboards, and control panels.”

Sam Mahalingam, CTO, Altair

Learn More: altair.com/panopticon