

“Panopticon’s OLAP capabilities and data visualizations provide us with a fantastic front end that leverages the capabilities of the entire BI infrastructure.”

— David Fretwell
CEO
ZY Solutions

Effective Monitoring & Analysis of Real-Time Business Data

Panopticon developed its data visualization tools to provide managers, analysts and decision-makers at every level in your organization with the monitoring and analysis tools they need to conduct profitable business. Our platform combines:

- Ability to connect to virtually any data source, including streaming real-time data sources like Active MQ, Apache QPID and Sonic MQ message buses, Complex Event Processing (CEP) engines like Sybase CEP, relational databases like Oracle, DB2, Sybase ASE, or SQL Server, and column-oriented databases like Sybase IQ or Kx kdb+.
- Fast development, deployment and training cycles to ensure quick ROI.

Panopticon is available in three versions:

Developer SDK allows programmers to embed the Panopticon StreamCube™ OLAP data model and data visualizations into their own applications. The SDK supports code-level integration with existing enterprise applications, client/server, or web-enabled, thin-client deployments. The technology is an excellent fit with SOA-enabled architectures. It is capable of receiving streams of data and events, then presenting the information to analysts in a useful graphical format. The system supports Java, .NET and Microsoft Windows Presentation Foundation (WPF) environments.

Panopticon EX is a fully productized system designed for rapid deployment at the workgroup or enterprise level. It incorporates a wide range of data visualizations, from Treemaps and Heatmaps to Scatter Plots and Bullet Graphs, and uses our unique StreamCube™ OLAP data model for on-the-fly data aggregations and slicing and dicing. The platform includes a desktop authoring tool that allows power users to assemble and publish new monitoring and analysis dashboards to the web.

Our unique **Rapid Development Kit** offers a range of features that make it easier than ever to build, publish and embed comprehensive dashboards into your own applications. It utilizes our Panopticon EX Designer tool to allow business users to quickly build their own interactive analytical dashboards. It also uses core components from the Developer SDK to allow you to tightly embed those dashboards into your own software systems. With the RDK, you only need to write a few lines of code to add embedded dashboards to your applications.

In-Memory OLAP Data Model Can Connect To Any Data Source

Panopticon’s unique StreamCube™ data model supports the ability to connect to virtually any data source, including relational databases, real-time streaming data feeds, and the output of CEP engines. It can handle traditional multidimensional analysis for static data (like those specified by the MDX language) and also for dynamic data streams.

Many data models on the market support multidimensional analytical OLAP functionality; however, they are generally not optimized for real-time performance. The Panopticon StreamCube™ is specifically designed to make extremely fast calculations using continuously updated data. The StreamCube™ matches the requirements that live visual interactive displays pose to a data back-end. This involves being able to bind aspects of the data to multiple visuals and associated analysis controls. Most other data models either cannot handle streaming data, or their performance degrades significantly when presented with large streaming feeds of real-time data.

Panopticon’s ability to handle real-time streaming feeds as well as historical information stored in column-oriented and row-oriented databases makes it an excellent fit for organizations that need to make decisions based on fast-changing data sets.

True Real-Time Performance

The idea of “real-time” can be very confusing, since virtually all software companies that do visualizations say they can handle “real-time” requirements. However, what most of them mean is that their software goes out and requests an update from an external data source when the data is needed. For example, a system might make an update request in order to generate a new graph. The system pulls the updated data from the external source and uses it in the display – at the moment when the data was pulled. The data is accurate as of that exact moment, but it becomes out-of-date immediately since no further updates are available until a complete refresh is done.

With Panopticon, external systems like CEP engines and message buses constantly push data into our system – not at a regular interval, but **instantly, as-it-happens and on a tick-by-tick basis.**


Treemaps

Treemaps use our innate ability to comprehend size, color and groupings very quickly – much faster than we can read and interpret reports, tables or diagrams. They enable users to visualize and act fast based on patterns and trends they are able to identify in what otherwise would be an overwhelming amount of data.




Heatmaps

A Heatmap is a special type of color-based data visualization designed to analyze large, flat data volumes using an intuitive graphical display. Heatmaps are good at representing large numbers of data points in ways that would be unwieldy and hard to interpret using traditional tables or charts.



Heat Matrix

A Heat Matrix is similar to a Heatmap or a Treemap in that it displays many different data items and it can represent the value for each item using colors. However, the Heat Matrix has a defined structure in which two attributes define each data item, thus producing a matrix. Each column and row represents a unique attribute, and the point where two variables intersect represents a unique combination of the two attributes.




Powerful Visualizations Designed to Work with Time-Critical Data

The information visualization techniques supported by the Panopticon platform help users see multi-levelled hierarchies, interconnections and details on individual objects in order to solve problems, understand complex relationships or identify areas of concern. They are optimized for handling time-critical data, including data that may be changing very quickly, even in real time. These tools enable complete understanding of the data at a glance by taking into account the latest research on human perception.


Bullet Graphs

Bullet Graphs were designed specifically to display Key Performance Indicators. We provide Bullet Graphs instead of traditional circular gauges or speedometers because they convey more useful information in less screen real estate. With Bullet Graphs, you can understand where you are in terms of KPIs with just a glance.




Line Graphs

Line Graphs are simple, intuitive visualizations that customers often use in Executive Dashboards and presentations. Line Graphs are particularly useful when you have a limited number of series to compare or when you need to examine the performance for a single set of data.




Stack Graphs

Stack Graphs let you visualize quantitative changes to several data sets over time and simultaneously see how each data point contributes to the total. For example, Stack Graphs are a useful way to look at revenue or gross profit figures over time across several product lines. The Percentage Area Graph variant on the Stack Graph is an excellent choice when you need to understand the relative contribution for each data set in the series regardless of the absolute total. You can see trends that you might miss in other visualizations.

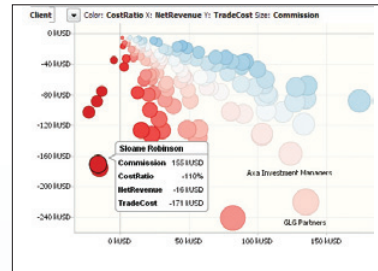


Bar Graphs

Bar Graphs are easy to understand and are a straightforward way to communicate important comparative information. Panopticon's Bar Graphs give you many options to display your data in the ways that make sense. Options include: horizontal or vertical alignment, standard, stacked, or grouped layouts.



Scatter Plots



It is often quite difficult to identify outliers or clustering in very large databases. Reports and standard table displays are time-consuming to interpret. Aggregations of the data — which can make reports and tables easier to understand —


can also mask outliers and make them difficult or impossible to see. Scatter Plots are the right choice when you are looking for outliers, clusters or trends in large data sets.

Geographic Scatter Plots can sometimes reveal surprising and unexpected groupings that would be difficult to identify in traditional reports. They are also excellent presentation tools and can help convey important findings to the public or decision-makers.




Horizon Graphs

Horizon Graphs allow users to view a large number of time series for trends and correlations typically hidden in traditional time series displays. This visualization packs the information in a line graph in 1/6th the space. Users can scan huge amounts of data points across all relevant time series and immediately identify areas of concern that require closer scrutiny.



Pie Charts

The Pie Chart is perhaps the most popular data visualization in the world and is familiar to almost everyone. We support two distinct styles, including Multi-Level Pie Charts, also known as Sunbursts. They allow you to display a hierarchy within the data using multiple levels in the pie.




Surface Plots

Surface Plots make it easy to identify trends and outliers within numeric surfaces. The Surface Plot is made up of a series of points, where each point has:

- X Position
- Y Position
- Color (Which represents the Z Axis).

The Surface Plot supports data sets in which the X and Y data elements have regular and irregular distributions.




Dot Plots

The Dot Plot is an effective alternative to the Bar Graph data visualization in cases where the data contains groupings with similar values. Dot Plots do not use a zero baseline and are less cluttered than Bar Graphs, making them easier to interpret in many cases. They are especially effective as a presentation tool to help people understand comparisons between grouped hierarchies of data.



Needle Graph

Needle Graphs display time-based transactions or occurrence frequencies rather than time-based trends. They are simply a special kind of Bar Graph in which each bar is associated with a particular time point on the axis. They work especially well when combined with a Line Graph visualization. A typical application is to use a Needle Graph to display trading volume for a stock underneath a Line Graph showing price performance data.



Candlestick and OHLC Graphs

Candlestick Graphs are a traditional financial visualization for display time based price distributions.



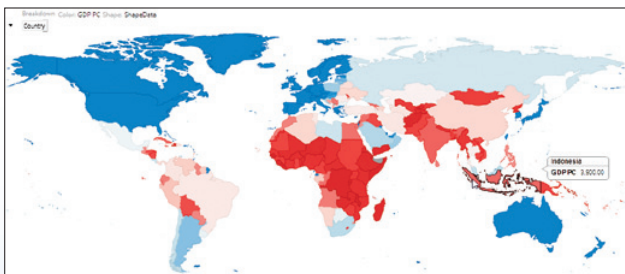
For each time slice, they display:

- Opening Price
- Highest Price
- Lowest Price
- Closing Price



The Candle is filled if the closing price is lower than the open. It is empty if the closing price is higher than the open. The vertical line (or Candle "wick") displays the range of traded prices across the period. OHLC Graphs are similar to Candlestick Graphs in that they also display time-based distributions of price data. As with the Candlestick, the vertical line defines the range of traded prices across the period. OHLC Graphs are different in that the left notch corresponds to the opening price and the right notch displays the closing price.


Shapes and Choropleth Graphs



The Shapes visualization enabled the display of Choropleth Graphs and other visualizations built from SVG Paths. The Shapes visualization can be used to display data where both physical location, and size are important. They clearly show geospatial data correlations and clustering.

Spread Graph

The Spread Graph displays the variance or spread between two time-based data series. This is an excellent visualization to use to compare two correlated series in order to identify important changes in the two series over time.



About Panopticon

Financial institutions all over the world use Panopticon software to help them make insightful, informed and profitable decisions. Panopticon's data visualization software together with its Real-Time Business Intelligence capabilities helps them understand correlations, identify outliers and see trends that are obscured in traditional reports.

Panopticon's clients use the company's data visualization software tools to speed up their business processes, reduce operational and investment risks, detect anomalies (including fraud), and identify opportunities to increase profits and sales. They help people see hidden patterns or spot problems in addition to helping them identify useful opportunities in their data.

Panopticon's platform excels in:

- Transforming data into intuitive academically validated graphical displays that are easy to understand and use.
- Enabling monitoring and analysis of multi-dimensional data with our fast, in-memory OLAP data model.
- Connecting to virtually any data source, including real-time streaming feeds, message buses, the output of Complex Event Processing (CEP) engines, relational databases, and proprietary data formats.

The Panopticon platform is available as an SDK that enables software companies and corporations to embed these capabilities into their own applications and in our Panopticon EX product designed for fast deployment at the enterprise, workgroup or desktop levels. We also offer our unique Rapid Development Kit; the RDK combines the ease-of-use available in EX with the integration capabilities of the Developer SDK.

Contact Us

New York +1 646 912 8444
 San Francisco +1 415 683 7398
 London +44 20 7877 0640
 Stockholm +46 853 480 480
 info@panopticon.com
www.panopticon.com

April 2011

Panopticon for Corporate Applications Brochure
 © 2011 Panopticon Software AB

The information in this document is subject to change without notice and is offered for informational purposes only. All trademarks used herein, whether recognized or not, are the properties of their respective companies.