



AutoPilot Insight:

Getting IT Operational Intelligence
from Logs, Metrics and Transactions

Table of Contents

- 3** Getting Operational Intelligence from Logs, Metrics and Transactions
- 4** Real-Time Analytics for Operational Insight
- 10** Addressing the Real Problem: Too Much Data, Too Little Time
- 15** AutoPilot Insight Features
- 18** Operational Intelligence, IT Ops and the Developer
- 24** Conclusion

IT Operational Intelligence from Logs, Metrics and Transactions

For IT Ops professionals, members of the DevOps group and developers, the ultimate advantage would be to know everything as it happens in their business—and everything that *could* happen. To know all and see all with complete vision is the competitive ideal: operations managers armed with real-time analytics, detecting performance problems before delays arise, the company discovering trends the moment they form. If there is a function in modern technology that offers anything close to this ideal of omniscience, it's providing real-time analytics to prevent problems (or at least their impact) and forensics to resolve the problems you can't prevent.

Such awareness is so difficult to attain, and it is often impossible for companies to know in advance what events need to be analyzed and when that analysis must happen. IT must **store everything** and **analyze everything** or risk missing the most important evidence of operational lags, risks or rising customer trends. For security compliance reasons alone, enterprises are required to maintain good logs, store logs for at least one year, secure these logs—and review them daily.

An important consideration is whether staff is looking at the data with the right degree of resolution. A low-resolution view lets support staff isolate a problem to, say, a specific infrastructure tier and then pass the problem to a specialist—who will start all over again diagnosing the problem with a different tool. While this common approach works, it is time

consuming, expensive and disruptive. A more productive approach would be to use a high-resolution analysis enabling application support to diagnose a problem, providing them with the forensics to immediately begin its resolution.

Real-Time Analytics for Operational Insight

A steadily increasing number of operations staff throughout the IT economy are finding analytics a compelling part of their tactical suites. Nearly half of respondents to a TDWI survey revealed that they are engaged in real-time analytics for operational insight. Roughly 25% of DevOps professionals surveyed by IDC reported a desire for analytics. And the competitive



Figure 1: An example of an operational intelligence dashboard. The data displayed is streaming into AutoPilot Insight from Log4j and Syslog Collectors.

benefits of real-time data are clear. An Aberdeen Group survey revealed that executives who were engaged in real-time streaming analytics averaged nearly 15% improvement in operational cash flow and overall profit. They ported a 6% reduction in operating costs. 70% of them reported that they are able to hold regular operational reviews.

Indeed, analytic platforms designed to deliver operational insights are rising in prominence. Yet for a company to act on trends and problems with precision and immediacy, a real-time, operational intelligence system must allow users to analyze perishable data while it still matters. This is essential in order to achieve the ever-important business goals of keeping customer services levels high, applications running, developers productive, and costs low.

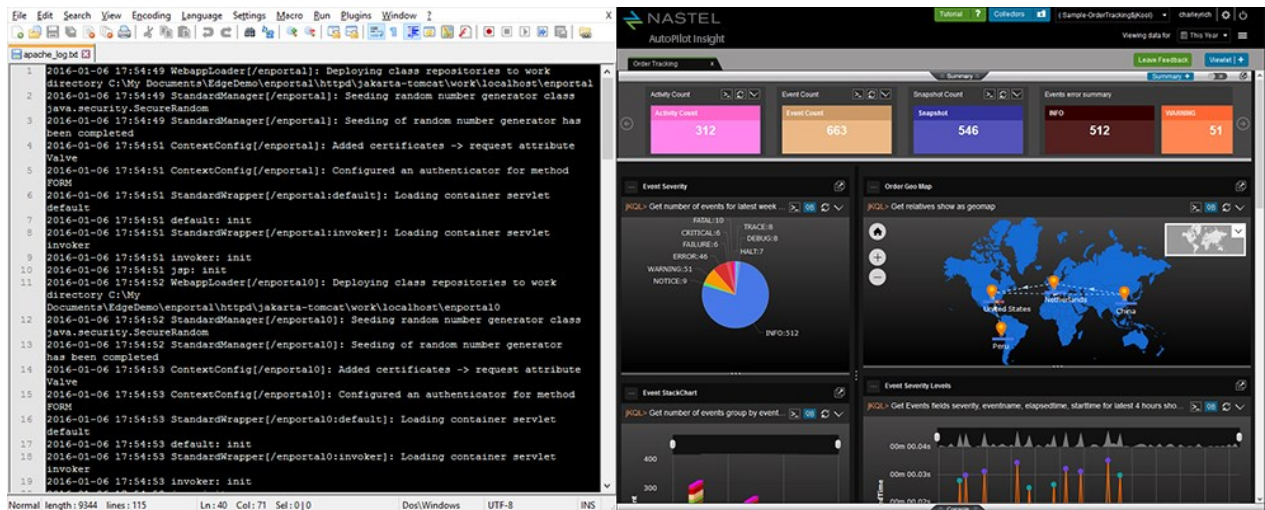


Figure 2: Traditionally, developers are called in to diagnose application behavior by interpreting log data as presented on the above, left. Using AutoPilot Insight's analytics and visualization on the right represents a presentation that Application Support can use.

Transaction information, delivery delays, consumer behavior: the value of this sort of data is ephemeral, its significance at first so high to users and then diminishing rapidly as time passes.

This type of data has an expiration date. As Forrester Research asserted in a 2014 report, “Business won’t wait. That is truer today than ever before because of the white-water flow of data from innumerable real-time data sources. Market data, clickstream, mobile devices, and sensors may contain valuable, but perishable insights. Perishable, because the insights are only valuable if [users] can detect and act on them right now.” For IT Operations these data sources also include Java logs, Syslog data, middleware events, various server and application metrics and transactions.

Too many times, developers are pulled from their critical and customary roles as coders of innovative apps, roles that are essential to the growth of their businesses. They are instead thrust into triage assignments in application support, drafted into diagnosing problems that, as the common perception goes; only they have the skills to address. They lose hours, days, and weeks scanning and annotating logs, tracing transaction and message flows and measuring application performance, their organizations finding no easy way to transfer their skills to less expensive, less critical personnel. Often operational in origin, raw machine data is scattered across multiple logs, servers and event infrastructure tiers.

In such a workflow, it can be very challenging for staff to piece together the story the data has to tell. Yet the right form of real-time, operational intelligence—intelligence that brings all (not some) machine data into a coherent, consumable and configurable visualization—is a superb way to scrutinize a broad range of performance metrics. From Java logs to Syslog, metrics and

Visualization & Reporting Layer

Analytics Processing Layer

Data Connection Layer



Figure 3: AutoPilot Insight architecture is built on open-source technology, built for extreme scalability and designed to ingest data from any source.

transaction records, such intelligence can deliver instant insight users need to keep the operation running, liberating developers to do their true jobs creating apps that drive the business forward.

Professionals at every level in development and operations can take advantage of real-time, operational intelligence to boost immediate insight into their raw, machine data. Seeing performance issues like troublesome trends, patterns and urgent problems as they unfold is crucial—yet awareness is only half of the game. Action and decisiveness are important

capabilities, as well. With real-time analysis highlighting perishable data at its most useful heights, employees are empowered to respond to issues rapidly, to anticipate difficulties, and to make fully informed decisions.

A real-time IT operational analytics (ITOA) product, AutoPilot Insight offers immediate insight, simplified, swift analytics and visualization to IT Operations and DevOps running production, test, and development environments. Users can observe transactions as they occur; teams can consolidate all logs and metrics onto a single configurable web dashboard. With this single pane of glass, staff no longer has to piece together stories from disparate logs or tools. Users

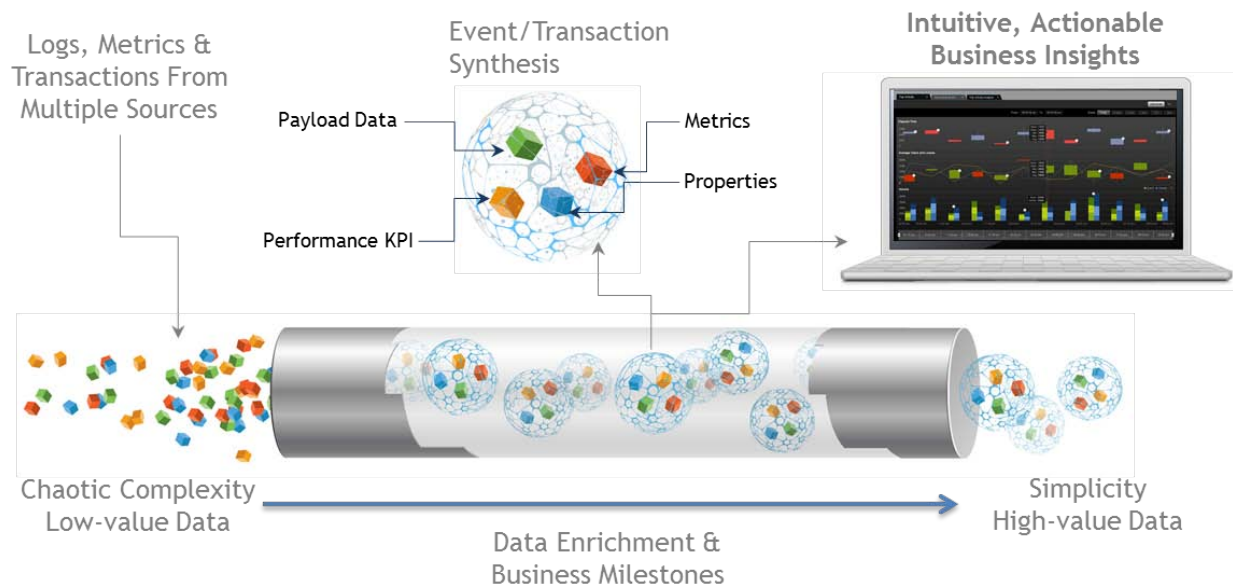


Figure 4: AutoPilot Insight ingests logs, metrics and transactions from multiple sources, enriching them with in-memory metrics and relating them to business objectives. This approach takes complex, low value data and turns it into high value, meaningful situations. Using this “situational analysis”, IT Ops’ productivity increases and the MTTR for problems decreases.

get a detailed record of their transactions stitched together, along with metrics, logs and are then able to see their flow of execution across applications, servers and infrastructure tiers as well as the environment context. This creates a **high resolution, end-to-end image** of your application behavior.

Available on-premises or in a SaaS offering, IT departments can reduce time fixing problems and spend more energy on building business resources, focusing on innovation, as they are freed from providing endless support. Using SaaS, there are no servers, schema or database or storage to manage. In either on-premises or using SaaS versions, AutoPilot Insight can greatly reduce the mean-time-to-repair (MTTR) for application problems through its real-time ability

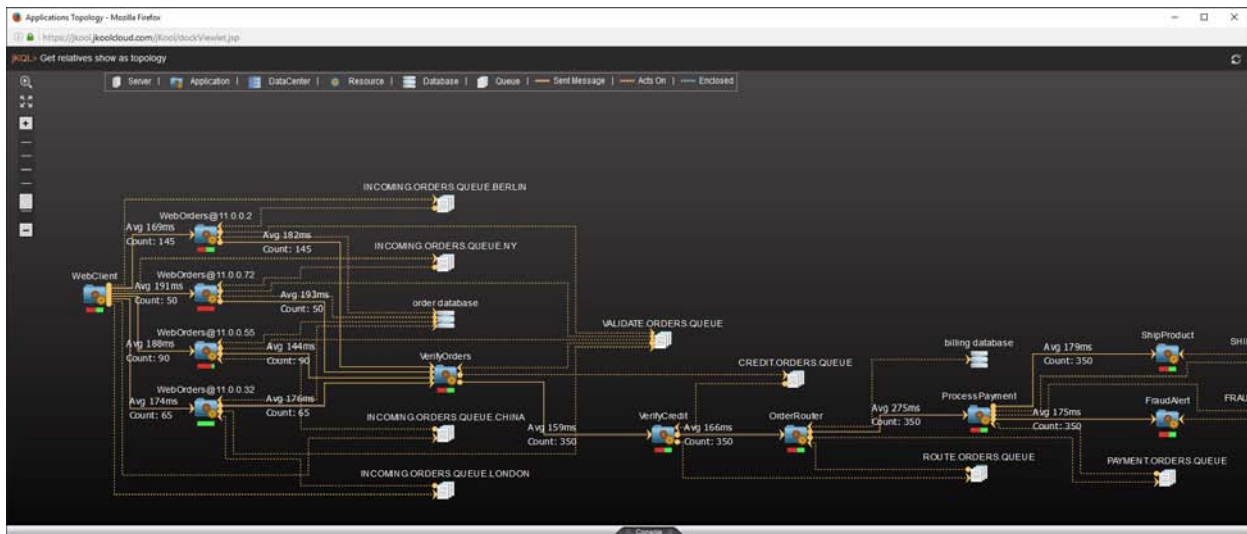


Figure 5: The Topology Viewlet above shows an automatically “stitched” transaction flow. It starts at the left with web browsers, then on to web applications, SQL databases, MQ Queues and more. No manual effort or modeling is required.

to track transactions, provide end-to-end visibility and detect anomalies. As an open source tool, AutoPilot Insight embraces the most compelling technologies in the world today and makes trends and patterns reassuringly clear to users.

Addressing the Real Problem: Too Much Data, Too Little Time

With operational realities generating immense, unorganized levels of data at a relentless pace, many enterprises are unable to keep up with the quantity of data they are accumulating. Institutions collect more data than they can manage, the sheer amount and organization of the data making it difficult to achieve the instant insights staff need for the decision-making

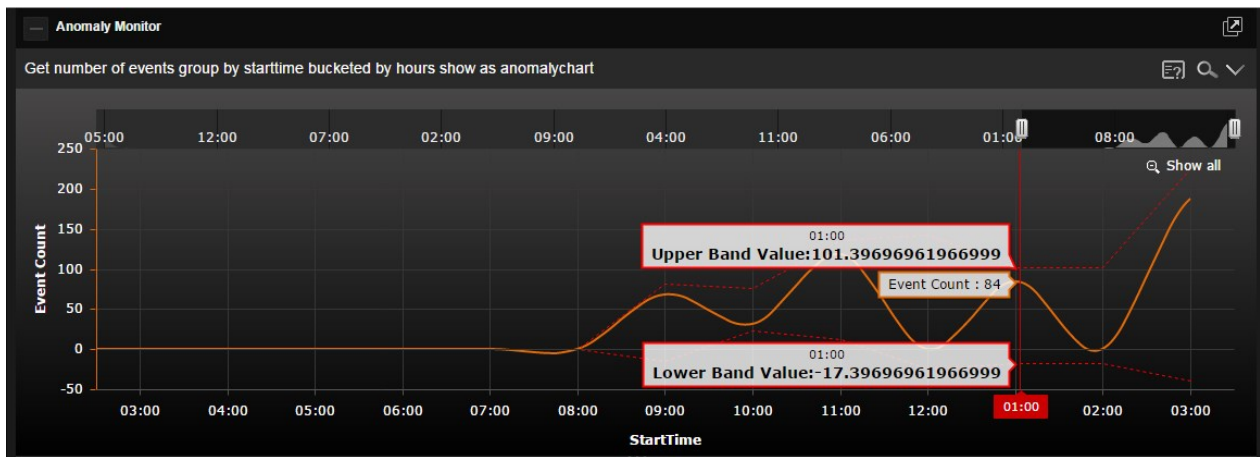


Figure 6: The above image shows how AutoPilot Insight automatically detects anomalies. The Anomaly Chart above uses Bollinger Bands and Exponential Moving Averages (EMA's) to auto-detect anomalies when an event crosses either the lower or upper bands. This is a self-maintaining analysis.

that keeps costs at suitable levels. The competitive advantage that rests in such data is buried in its huge volume and its inconsistent organization. Relational databases are taxed, scalability and performance lagging under the weight. Professionals have no time or means to figure out what the data means fast enough—losing the opportunity to act on trends and patterns before the data’s potency wanes. If organizations can move fast enough, there is opportunity hidden in the data that can provide value.

Support and development of mission-critical applications can be an area of tremendous drain on IT departments. Manual analysis of logs and metrics is costly, slow, and results in errors. Through AutoPilot Insight, users can streamline this responsibility, delivering fast identification of problem areas and performance. Users can address the data with a mere click or two. Raw,



Figure 7: This is an example of a real-time query. Notice the verb "Subscribe" which tells jkQL, the English-like query language to send matching events as soon as they arrive, even before they are saved to disk.

machine data can be presented in a way that enables actionable insight into misbehaving applications and business processes without requiring constant developer involvement. Diagnostics are clear.

AutoPilot Insight permits organizations to leverage less expensive personnel in support who will not need to do heavy lifting to keep the business running. For development organizations,

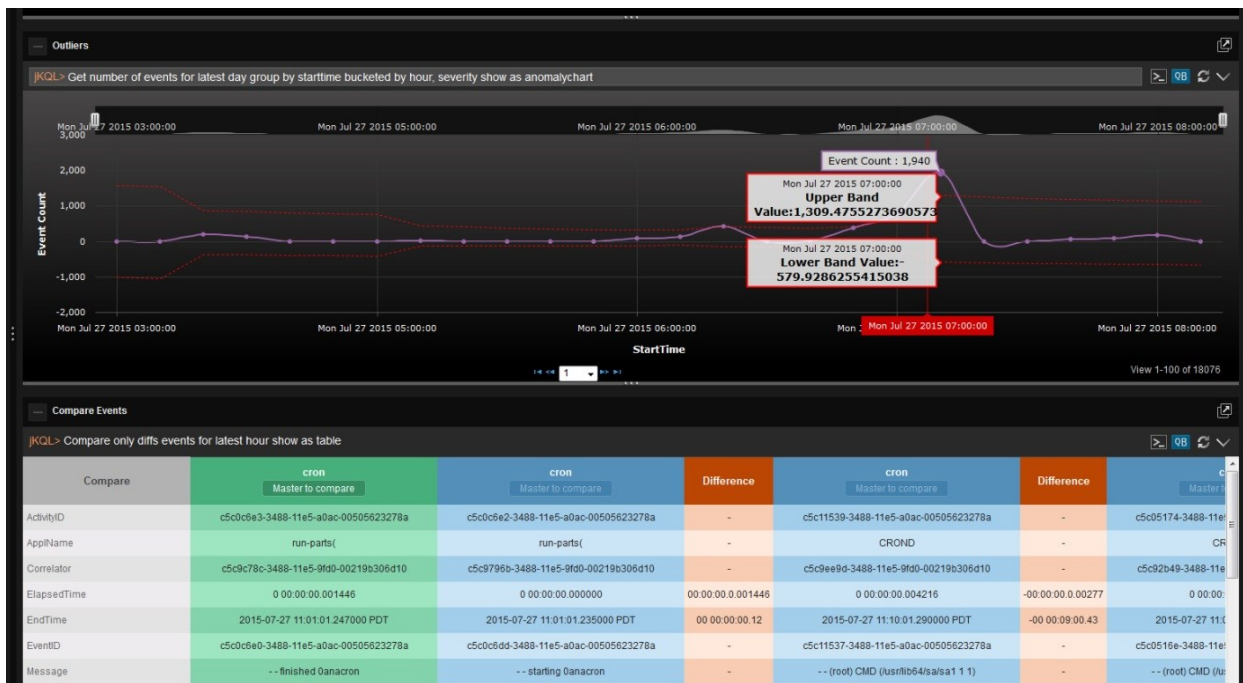


Figure 8: The Viewlet at the top shows an analysis of application events using Bollinger Bands and Exponential Moving Averages (EMAs). The Viewlet on the bottom is a comparison showing the differences between events.

this freedom means higher quality releases, delivered in less time. This can be a big boost in overall productivity for all concerned: IT Ops, Development and DevOps.

Analytics in the moment are useful for IT to achieve aims like loss prevention, projecting when users will run out of resources and rapid diagnosis of problems. These capabilities are valuable to any number of industries in a diversity of contexts. The ability to simulate an application is of great help to enterprise architects, who can use this to improve the effectiveness of pending deployments, as they will be able to visualize transaction dependencies and interactions prior to the actual deployment. They can then adjust their plans based on what they learn from the simulation. AutoPilot Insight’s simulator enables this important benefit.



Figure 9: Real-User Monitoring provides analysis of end-user experience using injection-based, browser instrumentation. The request from the browser is broken up into all its components and measured for latency. The user transaction is stitched to back-end server transactions. The root cause button will take the user to the probable root cause of the latency the end-user is experiencing.

Built to consume time-series machine data, AutoPilot Insight ingests, correlates and visualizes this operational data in memory, in real time. As AutoPilot Insight consumes data, it automatically stores, sequences, correlates, groups, enriches, synchronizes and computes ingested data in-memory and in persistent storage, delivering the instant insight that positions an enterprise to act with speed and confidence.



Figure 10: AutoPilot Insight's multi-panel "stock" chart for analyzing behavior of multiple applications over the same time intervals to detect hidden relationships

AutoPilot Insight Features

From 2012 to 2014, Forrester tracked a 66% increase in firms using streaming analytics. Yet adoption has its challenges, as many of these solutions do not provide out-of-the-box operational intelligence. In essence, such solutions are toolkits from which users must build their own applications. Real-time visualization of streaming data is usually not provided. Other offerings do not yet provide the breadth of functionality—including a useful query language—that is needed to derive instant insight from the data ingested.

Easy to use, AutoPilot Insight was built to provide the predictive analytics to prevent problems and the forensics necessary to diagnose problems fast, when they do occur. AutoPilot Insight offers visualization capabilities to address all elements of Java and other logging and metric sources. The system connects to several logging utilities: Log4j, SLF4J, Logback, Logstash and Syslog. It also ingests data from collectors for various extensions, including JMX, Java EE, StreamGC (Java Garbage Collection), HTTP, HDFS, Apache Flume, Kafka and Spark, JMS, Python, IBM middleware (MQ, IIB and DataPower), end-user experience data, mainframe and more. AutoPilot Insight consolidates and summarizes all logs and metrics, permitting users to detect anomalies and view trends in error rates over time. Self-service capabilities enables users to add metrics via open-source APIs (REST and Java), allowing an automatic track and trace of application transactions. Offering a wide and comprehensive view of the landscape, AutoPilot Insight automatically stitches together transactions and scans problems in operational flow and performance in order to determine problem causality.

The visualization capabilities span:

- **Real-time:** Using in-memory analytics before events are written to disk, AutoPilot Insight can display events and metrics as they arrive.
- **Real-time:** Using in-memory analytics before events are written to disk, AutoPilot Insight can display events and metrics as they arrive.
- **Graphing:** Visualization options include tables, scorecards, histograms, line charts, bar, column charts, pie charts, area charts, summary views and more.
- **Streaming operators:** Comparison of metrics can involve EMAs, Bollinger Bands, Apdex, Abs, Ceiling, Exp, Floor, Ln, Log, Log10, Power, Round Sqrt, count, sum, min, max, avg, first, last, top, bottom latest, earliest, best, worst, largest, smallest, shortest and more.
- **GeoFencing:** AutoPilot Insight's GeoFencing capability can segment events and activities by geolocation, plotting relationships and displays results on a map. AutoPilot Insight automatically converts latitude and longitude to text locales.
- **Comparisons:** jKQL can designate metrics as master data and display other metrics in comparison to main data.
- **Automatic detection:** AutoPilot Insight automatically detects anomalies and displays those using statistical anomaly detection.

With a strong offering of features, AutoPilot Insight grants companies an easy set-up and a level of insight and usability that gets staff acting on trends and problems in short order. For developers, the benefits soothe a constant stress in modern IT departments: the endless pulling of developers into the application support role, a reoccurring situation that drains so much energy and time from their main mission.

AutoPilot Insight's quick-to-use language for analyzing behavior, jKQL, communicates queries with a direct, precise, and logical simplicity. Through jKQL, users initiate ad-hoc analysis of a range of metrics. Using the jKQL language, users can talk directly with their data and get answers. AutoPilot Insight's process of querying is designed to let the user learn things they didn't already know. For example, support working proactively might ask, "Get the worst 10 events fields applname order by elapsedtime desc show as barchart" and receive a barchart representation of the 10 most chatty classes over the last hour. The user can then ask more questions for better insight into the causes and realities of relevant operations, Users can drill down into the data returned; select "analyze" and see a multi-panel chart that shows elapsed time for events, volume of all activities along with user selected business data over time.

Alternatively, a business user can enter, say, "get the worst 50 orders for the last hour". AutoPilot Insight's visualization graphics can display data in a number of different formats, including candlesticks and trend lines. The graphics can then simultaneously show the elapsed time for each order, the value of orders and the overall volume of events during that same time period. The user can view the flow of execution of the order transactions and also isolate them geographically.

jKQL lets users perform analysis using streaming operators. They can compare and highlight differences between activities. Compatible with both browser and tablet formats, jKQL displays results as tables, charts, scorecards, histograms, heat maps, candlesticks, comparison and topology.

Operational Intelligence, IT Ops and the Developer

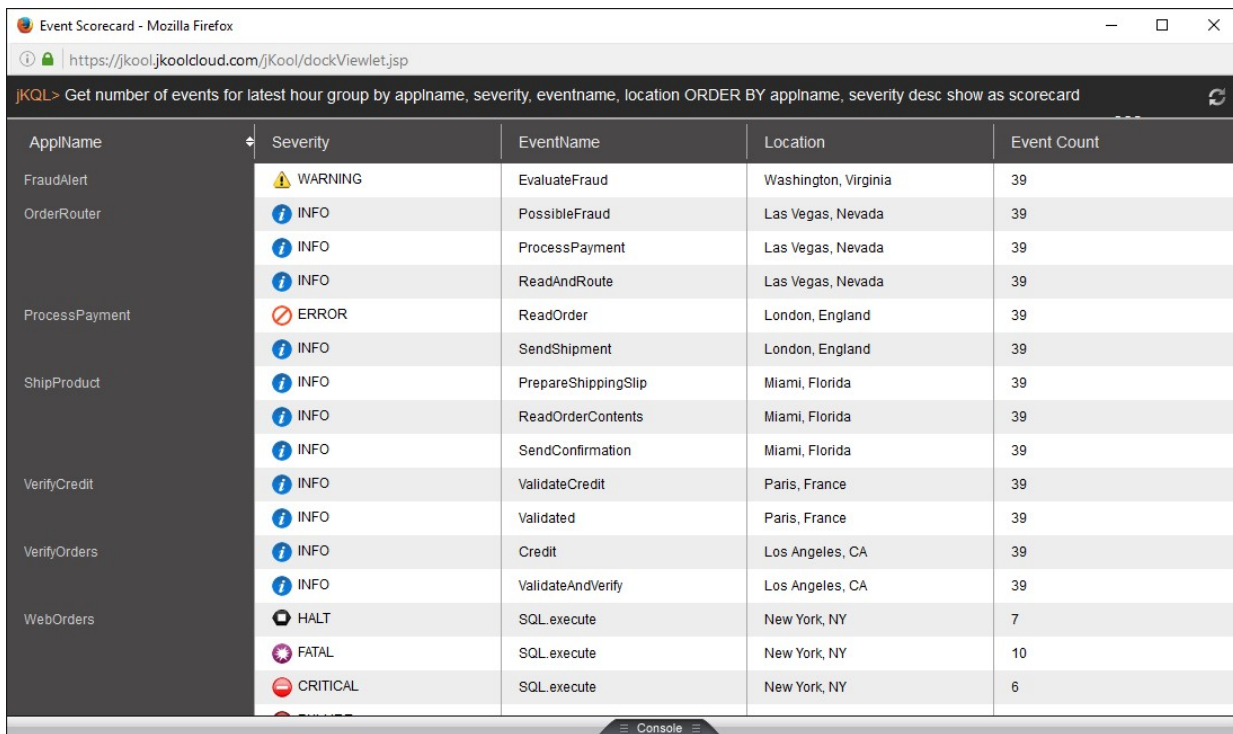
When an application misbehaves, the support group in IT Ops often has to sift through multiple log files in order to trace the path of each application and find out what went wrong. With multi-threaded applications, these logs can appear chaotic, each thread writing a log entry that may have nothing whatsoever to do with the prior entry. Such a lack of coherence makes diagnostics quite challenging and often requires staff using pencils and paper, chalk and string to discover what happened. This is a tedious manual process that easily becomes quite difficult, with each application written to different log files, sometimes on different servers and many times even on differing infrastructure tiers. Multiply this effort by a large number of servers and the situation rapidly becomes unmanageable.

IT Ops can find it difficult to follow the flow of execution, requiring significant involvement by developers, the most critical and expensive employees in the organization. Though the purpose of this diagnostics effort is to understand why performance, the flow of execution as logged may not appear linear. In the jumble of distributed data, causality is very difficult to discern.

Unfortunately, this type of diagnostics occurs all the time in modern IT departments. AutoPilot Insight offers an alternative to this taxing app diagnostic model that is so commonly a default response by companies. By feeding AutoPilot Insight with a range of data collectors, users will find the set up and preparation of the system easy to do.

AutoPilot Insight’s open-source collectors which are downloadable from nastel.com or alternatively from [GitHub](https://github.com). Each collector enables data gathering and delivery to AutoPilot Insight.

Java developers can use the Log4j Collector, an appender that allows developers to add an additional target to the Log4j configuration without any edits to application code. SLF4j, Log4J, Logback or Apache Flume users can consolidate all logs in a single pane of glass; these log visualizations improve diagnostics and root cause identification. A truly comprehensive view of



The screenshot shows a web browser window titled "Event Scorecard - Mozilla Firefox" with the URL "https://jkooljkoolcloud.com/jkool/dockViewlet.jsp". Below the browser window, a jKQL query is displayed: "Get number of events for latest hour group by applname, severity, eventname, location ORDER BY applname, severity desc show as scorecard". The main content is a table with the following data:

AppName	Severity	EventName	Location	Event Count
FraudAlert	WARNING	EvaluateFraud	Washington, Virginia	39
OrderRouter	INFO	PossibleFraud	Las Vegas, Nevada	39
	INFO	ProcessPayment	Las Vegas, Nevada	39
	INFO	ReadAndRoute	Las Vegas, Nevada	39
ProcessPayment	ERROR	ReadOrder	London, England	39
	INFO	SendShipment	London, England	39
ShipProduct	INFO	PrepareShippingSlip	Miami, Florida	39
	INFO	ReadOrderContents	Miami, Florida	39
	INFO	SendConfirmation	Miami, Florida	39
VerifyCredit	INFO	ValidateCredit	Paris, France	39
	INFO	Validated	Paris, France	39
VerifyOrders	INFO	Credit	Los Angeles, CA	39
	INFO	ValidateAndVerify	Los Angeles, CA	39
WebOrders	HALT	SQL.execute	New York, NY	7
	FATAL	SQL.execute	New York, NY	10
	CRITICAL	SQL.execute	New York, NY	6

Figure 11: AutoPilot Insight’s scorecard is a powerful way to “score” your applications. Above it is being used to show the latest hour’s events for an eCommerce book-selling application.

Stitching is a process of connecting events from seemingly unrelated sources, much like putting together a jigsaw puzzle from many different pieces. These events may occur at different times, come from different applications and have different attributes – yet they are related to a specific activity and tell a story about your application: such as payments, claims, searches, orders, etc. The process of stitching combines these events to describe a lifecycle of a composite activity spread across multiple applications, compute analytics and discover topology and flow. The result is high-resolution view of application activities, improved root cause analysis, dramatic reduction in MTTR (mean time to problem resolution) and increase in MTBF (mean time between failures).

the landscape emerges, as users can search log entries across applications and servers, while annotation of log entries bolsters diagnostics.

Other application developers can make use of the simple RESTful API to send their data to AutoPilot Insight.

DevOps or IT Ops can also deploy the StreamJMX Collector to send Java performance metrics to AutoPilot Insight from any relevant Java servers in the enterprise's infrastructure. Downloadable on the AutoPilot Insight site or from GitHub, StreamJMX is a lightweight framework to stream and monitor JMX-borne metrics. IT staff traditionally monitor JVMs through remote JMX monitoring tools, which requires the team to deploy and manage JMX configurations, ports, and SSL. Such remote connectivity is notoriously weak and raises cyber vulnerabilities—further burdening IT with security obligations.

StreamJMX allows developers to stream JMX metrics from JVM out to destinations throughout the network. This functionality allows users to maintain a history of JMX metrics across multiple JVMs. Streaming JMX from within the JVM is often more desirable alternative to remote sampling when monitoring large number of JVMs, improved security (no open

JMX ports), less administrative overhead associated with opening and securing remote JMX ports and connections.

AutoPilot Insight allows for monitoring of JEE Web applications through HTTP Servlet Filter. Users can measure performance, attenuations, and signals; dashboards can display variable



Figure 12: An example of an AutoPilot Insight dashboard with application topology showing transactions spanning Web Apps, IBM DataPower and an ESB

histories of response times. With entire organizations engaging these applications, IT can track end-user behavior and location.

IT Ops professionals using Syslog for system management and security can use AutoPilot Insight to bring together logs from multiple servers on one screen. Viewlets can illustrate the consolidated logs to help with assessing diagnostics and root cause. Users can search Syslog entries across applications, which can be annotated.

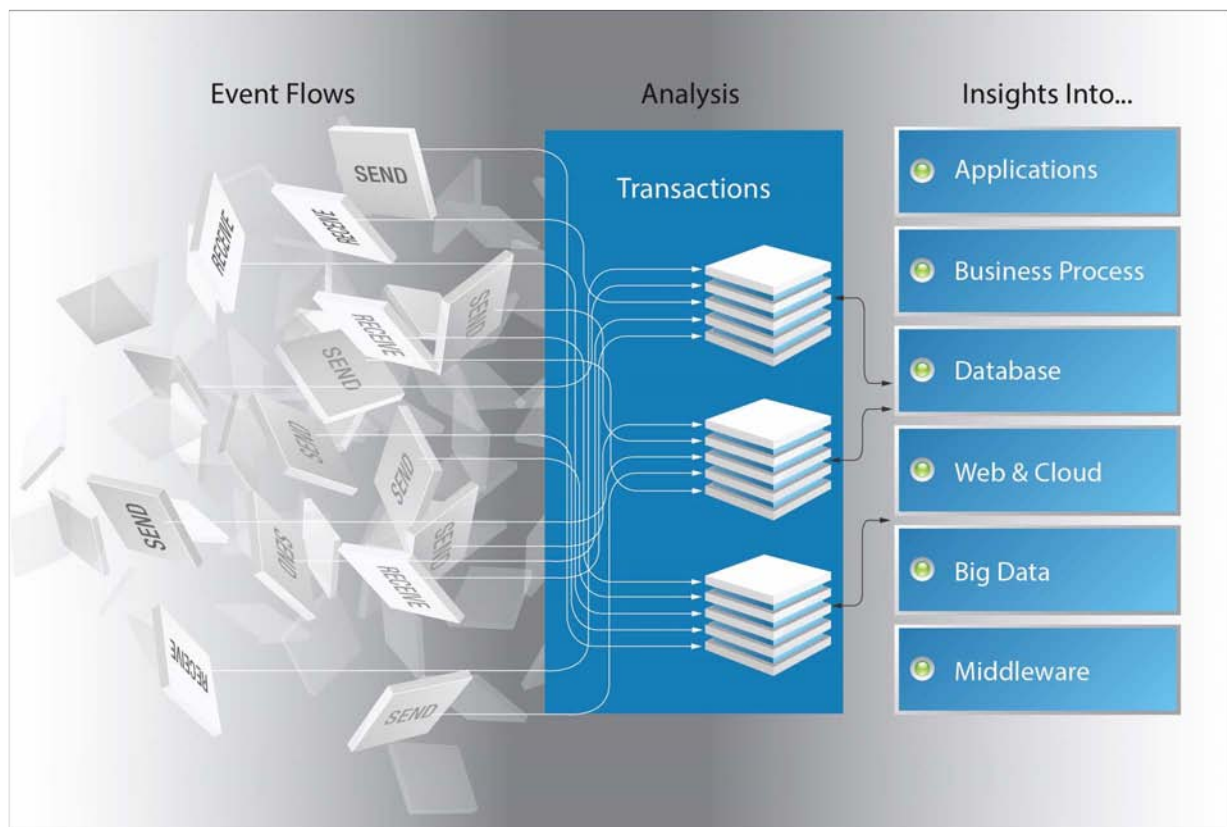


Figure 13: The image above is a conceptualization of AutoPilot Insight ingesting a chaotic stream of events, analyzing and stitching them into an orderly flow of application transactions. This analysis provides IT Ops, DevOps and Development with instant insight into: apps, business processes, database, web, big data and middleware.

With these collectors engaged, AutoPilot Insight automatically sequences, correlates, groups, enriches, synchronizes and computes everything it can about the data as it arrives—immediately providing a rich visualization and analysis to the user.

The collectors come with documentation describing deployment methods and details. Before deploying a collector, users can test their visualization with AutoPilot Insight’s built-in Simulator. This feature generates and displays a randomized version—at selected intervals and volumes—of what users can anticipate in terms of viewing data.

Because AutoPilot Insight’s design prioritizes simple usage, once the data is flowing to it—through the user-friendly jKQL query language and the detail and flexibility of the Viewlet visualizations—less expensive IT staff can serve in support roles for the system.

Users will find jKQL compatible with a host of data collectors. Equally critical to the system’s design is the priority of perishable data, especially machine data that must be captured and understood in the moment. Only then can professionals take the decisive steps that lead to reduced costs and increased profits.

AutoPilot Insight is built on open-source software including the Apache Cassandra NoSQL DB from DataStax; Apache STORM; Apache Spark; and Apache Kafka. AutoPilot Insight’s FatPipes micro-services architecture is designed to orchestrate these open-source technologies into a solution offering extreme scalability and ultimate configurability. Complex event processing delivers faster analytics from in-memory resources. This processing capability goes right to the heart of the challenge—perishable data must be captured, analyzed and visualized in real

time. With its rich set of open-source tools for data ingestion, IT Ops can easily use AutoPilot Insight for their real-time operational intelligence requirements.

With scalability so valuable as operational data levels continue to rise, AutoPilot Insight's Lambda Architecture aims to handle vast volumes of complex, concurrent data streams. The system's compute grid automatically runs queries in parallel to enhance performance and elasticity.

Conclusion

Business is about anticipation and awareness—knowing what is happening, what could happen, and how the company should act to resolve problems and outperform competitors. The pace of operations today can offer tremendous amounts of data to businesses—amounts that can intimidate and overwhelm as much as enlighten. The common approach to data analysis is unable to match this pace. Developers lose too much of their valuable time serving as ad hoc data managers, piecing together logs from disparate applications. By the time they are able to discern patterns in the data, the opportunity to respond to underperforming operations has passed their company by.

AutoPilot Insight provides awareness to such businesses by delivering clear, consumable operational analytics to IT Ops professionals. The product allows users to query data and dive deep into granular details to compare, summarize, and ultimately make swift business decisions.

Through AutoPilot Insight's simplified user design, less expensive IT staff can serve in support roles for the system. Prioritizing extreme scalability and ultimate configurability, AutoPilot Insight's complex event processing delivers faster analytics from in-memory resources. This processing strength goes right to the heart of the challenge—perishable data must be captured in real time.

With extreme scalability, swift visualization and real-time analytics, AutoPilot Insight empowers companies to act with confidence in fast-moving times. With AutoPilot Insight delivering streaming, high-resolution analytics and transaction tracking insights are more readily realized, staff can take more focused action faster, MTTR is reduced, and a crucial, ultimate aim is achieved: more customers are satisfied with much better frequency.