



Add Embedded Analytics to Deliver Groundbreaking Solutions

PUT DATA-DRIVEN INSIGHTS INTO THE HANDS OF USERS AND CUSTOMERS





What's inside:

- 3 How embedded analytics will change your game
- 4 Some of the many advantages of embedded analytics
- 5 How to get started
- 6 Build or buy
- 7 Key platform requirements for your embedded analytics program
- 12 Case study: How PDX revamped with embedded analytics in just three months
- 13 Why PDX chose Snowflake and Looker
- 14 Find out more



How embedded analytics will change your game

TIME TO UPGRADE YOUR APPROACH TO ANALYZING DATA

Today's organizations demand action driven by insights and placed into the hands of employees, partners and customers. Their ultimate goal: Use data to drive better decisions and help build new business models.

Historically, providing key data to internal business users has driven the need for richer data analytics. Business intelligence platforms, paired with data warehouse solutions, emerged to supply that opportunity. That led to dashboards and reports in a fairly standardized way no matter the type of analysis.

But these traditional approaches no longer serve today's organizations. Centralized analytics often lack the context individual groups across the organization need to make informed decisions. A new solution is required, one which speeds decision making, supports more types and greater numbers of users, and offers deeper insights.


Embedded analytics can meet those needs. Unlike traditional business intelligence, which moves data from applications to a separate environment for analysis, embedded analytics delivers live, interactive and contextual analytical insights from inside the transactional business applications.

Some of the many advantages of embedded analytics




Live Data

Reports aren't static but dynamic, based on live data.



Interactive Reporting

Embedded analytics offers more granular views of data and more flexibility over how data is presented versus predetermined data views.



Contextual Insights

Unlike traditional business intelligence, in which reports and dashboards are accessible only within the business intelligence tool, embedded analytics provides insights inside transactional business applications.

For example, PDX Inc. offers a SaaS application to pharmacies and pharmacy chains across the US. Its embedded analytics component enables pharmacies to streamline their operations and provide better patient outcomes and better customer service. PlacIQ is a leading data and technology provider that delivers location data, analytics and insights to leading brands in retail, automotive and global investment companies so they can better understand and serve their customers with timely and targeted products and services.

Organizations use embedded analytics to enhance and differentiate existing products with new capabilities. For those reasons, demand for embedded analytics products continues to soar. The global embedded analytics market is estimated to grow from \$24.46 billion in 2016 to \$46.19 billion by 2021—a Compound Annual Growth Rate (CAGR) of a whopping 13.6%.¹ Simply put, that translates into a big opportunity for businesses to capitalize on this growing technology to enhance their products.

¹ "Embedded Analytics Market by Tool, Business Application, Deployment, Vertical, and Region-Global Forecast to 2021," MarketsandMarkets, August 2016.



How to get started

REQUIREMENTS FOR EMBEDDED ANALYTICS

Traditional business intelligence and analytics platforms aren't going away. Standardized reports are easy to understand, and they're helpful for many professionals. However, they were designed for a different era, one before the emergence of embedded analytics. Similarly, traditional data warehousing platforms were architected in an era when data was scarce, when it had to be painstakingly curated to make sense, and when access to that data was tightly controlled.

Traditional reporting systems are easy to set up, but embedded analytics platforms demand capabilities beyond those of traditional platforms and tools. If you're considering implementing embedded analytics into your applications, here's what's required from an infrastructure standpoint.

FRIENDLY TO DEVELOPMENT AND ITERATION

Embedded analytics is at the core of many data-driven applications. Those applications are in a constant state of iteration and evolution as companies add new features to them. To support development, the technology used to enable embedded analytics must be easy to integrate and manage within the development lifecycle.

TIME TO VALUE

In today's world of software-as-a-service (SaaS) applications, new customers expect near instant access to the application the moment they sign up. As well, customer tolerance for delays and slow performance has vanished. Users accustomed to Gmail, Facebook and Amazon Web Services won't wait for hours for systems to be reconfigured, let alone remain idle for days or weeks for new servers to be deployed. The same goes for an application's embedded analytics.

A CONSISTENT, INTUITIVE VIEW OF DATA

To serve today's broad range of analytics users, all users must have access to the same source of truth. Further, that view of data must be intuitive and concise. Embedded analytics must simplify and represent the underlying data in a way that makes it easy for all types of users to understand.

MINIMAL MANAGEMENT

Embedded analytics handles demands from a constantly fluctuating number of users who may be working at any one time. To handle that scale and flexibility, the technology used to build embedded analytics must enable self-service features, such as intuitive exploration and drill-down capabilities, without requiring manual setup, configuration and management, tuning and monitoring.

LIMITLESS SCALE, GREAT PERFORMANCE

With traditional architectures, the more users you add, the slower everything gets. On the flip side, embedded analytics must automatically support an unpredictable and changing number of users, without any impact to performance. Regardless of how many concurrent users are active, embedded analytics must deliver consistently high performance at any scale.



“We started our proof of concept in December and went to market in March. We had our first paying customer by September.”

JOHN FOSS
VICE PRESIDENT OF
STRATEGIC PRODUCTS,
PDX, INC.

Build or buy

STARTING FROM SCRATCH MIGHT NOT BE IDEAL

So you want to build embedded analytics into your product? The next decision you have to make involves what technology to use.

Traditional data warehousing and business intelligence platforms have formed the backbone of reporting and analytics for decades. Unfortunately, those platforms lack the flexibility, ease of use and scalability needed to support the demands of today's embedded analytics technologies.

Recognizing those limitations, you may consider building a custom solution. That's an option, but building the technology to support embedded analytics can be extremely complicated. Even basic functionality can require months of engineering time, not to mention the resources required later to support, maintain and evolve the solution. This type of development also distracts a business from its core strengths, which typically lie in its data, not in building complex analytics platforms around it.

There is good news. You don't need to compromise on functionality and flexibility by purchasing a legacy platform, and you don't have to build one yourself. Innovative new platforms for data warehousing and analytics have emerged that meet the technical and business demands of embedded analytics. With modern data warehousing and business intelligence platforms, any size organization can dramatically reduce its time to market.

As mentioned above, PDX, Inc., provides SaaS healthcare technology so pharmacies and pharmacy chains can better serve their patients. But PDX wanted to deliver analytics direct to its pharmacy customers, thus the company developed ExploreRX—its embedded analytics product—providing PDX customers with a better picture of the patients they serve. “We started our proof of concept in December and went to market in March. We had our first paying customer by September,” PDX's Vice President of Strategic Products, John Foss said.

Key platform requirements for your embedded analytics program

POWER, SCALABILITY AND EASE OF USE ARE ONLY THE BEGINNING

Let's dig deeper into the technology platform requirements for embedded analytics.

A PLATFORM GENUINELY BUILT FOR THE CLOUD

Only solutions built for the cloud can deliver the elasticity, scalability and flexibility that embedded analytics applications require. Cloud platforms offer the promise of resources on demand, including storage for the ever-increasing volumes of data, and computing capacity to process complex and growing workloads.

These architectural differentiators are paramount to delivering effective embedded analytics to unlimited users with unlimited needs. History proves that any analytics technology creates demand that has far outstripped computing capacity. With cloud-based technology, that decades-old problem disappears.

But it's important to distinguish cloud-enabled legacy systems, which are not fit for embedded analytics, from a true cloud platform. What makes a solution built for the cloud? It must include these characteristics:

- **Automatic deployment and configuration:** Legacy solutions operated in the cloud require manual deployment and configuration of the environment—choosing virtual hardware, installing applications and so on. That's not the case with solutions genuinely built for the cloud, which are automatically deployed by the technology vendor to eliminate delays and errors in configuration.
- **Elasticity:** A solution built for the cloud can add and subtract resources on the fly, without disruption. For a data warehouse, that means being able to immediately scale storage and processing at any time, without taking the system offline and without requiring data redistribution or data reloading. That solution should also independently scale compute from storage, eliminating the need to build complex systems used simply for storing data.



Key platform requirements, cont'd

- **Pay for what you use:** Legacy data warehouse and business intelligence systems require careful upfront capacity planning and large capital expenditures. Many legacy solutions deployed in the cloud still encourage users to make large upfront commitments via fixed-capacity contracts in which costs are independent of usage. With solutions built for the cloud, costs are directly tied to usage. That allows you to start small and only pay for what you need, scaling as your applications grow.
- **Designed for the web:** Legacy solutions designed for client-server architectures or desktop usage are difficult to integrate with today's modern web applications because they rely on heavyweight desktop clients. Solutions built for the cloud can be accessed seamlessly via web browsers and mobile interfaces.

A CONSISTENT VIEW OF DIVERSE DATA

Today's data is not limited to relational data, and it arrives at an unpredictable velocity, often continuously. As a result, embedded analytics requires much more than traditional data warehouse solutions designed for relational data, which arrives at regular intervals. To avoid imposing limitations on both the variety of data and speed of delivery for embedded analytics, a cloud data warehouse must include:

- **Native support for diverse data:** Traditional relational databases were designed for structured data with fixed schemas, requiring data transformation before other types of data can be loaded and exposed to analytics users. However, modern data applications commonly generate and process semi-structured data with flexible, hierarchical schemas (e.g. JSON, Avro, and XML formats). Supporting these applications with embedded analytics requires a solution that can natively load and optimize data in the forms in which the data arrives.

The solution must also expose that data to users through the analytics layer, without requiring cumbersome manual extraction and flattening.

- **Centralized data storage:** Unlike the data silos created by the need to have multiple data marts, such as business intelligence servers and heavyweight desktop clients, embedded analytics needs a centralized repository of data to ensure all users see the same information and insights. A centralized, scalable repository is critical to providing a single version of truth so different users receive consistent insights.
- **Standardized data definitions:** Embedded analytics technology must ensure data definitions are codified in a robust data dictionary or modeling layer to maintain a consistent view of data even as data structures and schemas evolve. Without standardizing data definitions, changes to schemas impose heavy investment from engineering to manually update reports and queries.

“We wanted to scale up our spend as ExploreRx became more successful, and that’s exactly what the Snowflake/Looker solution has allowed us to do.”

JOHN FOSS, VICE PRESIDENT OF STRATEGIC PRODUCTS, PDX, INC.

Key platform requirements, cont'd

HARNESS DATABASE SCALE AND POWER

Traditional business intelligence and analytics tools were built to support a broad range of databases, many of which are unable to deliver the responsiveness and performance that embedded analytics required. To meet performance requirements, those business intelligence tools needed to extract data out of the database or deliver stale data to end users in the form of cached or pre-calculated results.

The myriad layers of management and monitoring required by that approach create overhead and headaches for administering embedded analytics. They require periodic refreshing, rebuilding and recalculation. They also add latency and delays that prevent users from seeing the most up-to-date information.

However, leveraging the power and scalability of a modern data warehouse makes it possible to deliver massive amounts of data to many users at scale. Analytics platforms that intelligently push processing into the data warehouse can rapidly deliver analytics that operate on live data, bypassing the inefficiency of managing optimization layers that live outside the data warehouse.

Scaling also requires leveraging additional storage and compute resources to automatically deliver consistently great performance, when scaling resources up, down or out (concurrency).

Legacy data warehouse technology has limits to concurrency because its architecture forces all concurrent users and queries to share the same resources. That makes it difficult to scale embedded analytics to meet the demands of ever-changing numbers of users and applications. However, a data warehouse with an architecture built for the cloud can automatically scale, providing a consistent experience for users and workloads that rely on embedded analytics.

SIMPLE TECHNOLOGY THAT SUPPORTS SELF-SERVICE

The analytics platform supporting embedded analytics must offer the advanced capabilities that users have come to expect for data-driven applications. For the data warehouse, that means it must have the ability to handle ad hoc and exploratory queries without requiring manual work to reconfigure parameters. For analytics tools, it means they must provide:

- A comprehensive, intuitive way to drill into aggregated metrics to access row-level detail.
- A standardized, central data governance layer built in the same environment that your embedded analytics application uses to prepare and transform that data.
- The ability to create web-native, interactive visualizations that live in-browser, allowing easy sharing with others across the entire data ecosystem.
- The ability to execute queries directly against your database without having to export data excerpts for analysis in other tools.
- Dashboards that you can easily create and edit with drag-and-drop functionality.
- Expandable capabilities via custom visualization plug-ins.

Key platform requirements, cont'd

AUTOMATED MANAGEMENT AND OPTIMIZATION

Embedded analytics also requires technology that eliminates the cumbersome overhead of installing, configuring, monitoring and managing hardware and software. It must provide built-in capabilities that eliminate or automatically handle the laborious tasks that become overwhelming for an application used by large numbers of internal users or SaaS customers.

A solution designed to eliminate these headaches should include:

- **Built-in security:** Properly managed, the security measures provided by a cloud data warehouse can be a much more effective, less expensive option than attempting to manage the security infrastructure yourself. Security is always of paramount importance, but even more so for applications that fall under compliance requirements such as HIPAA and PCI.
- **Instant deployment:** Leveraging the cloud, SaaS solutions for data warehousing and analytics should deploy almost immediately, eliminating bottlenecks to developing new embedded analytics applications.
- **Resiliency and availability:** Embedded analytics should be always available to users, without downtime or disruption. Building and monitoring a solution that incorporates redundancy and automatic recovery from errors is a huge effort that bogs down organizations. Therefore, embedded analytics solutions should have built-in redundancy and failure detection and recovery—all without manual setup, monitoring and management. Further, solutions for embedded analytics must also eliminate planned downtime. In particular, that means eliminating the need to take a system down or place it into maintenance mode for updates and patches.
- **Automatic optimization:** Embedded analytics solutions are impossible to scale and manage when they require manual tuning and optimization. Building and rebuilding indexes, optimizing data partitions, managing job queues, adding query hints—all of these are examples of manual labor associated with legacy data warehousing that can't keep up with the demands of embedded analytics. A data warehouse supporting embedded analytics must eliminate or automate this optimization, automatically delivering optimal performance even in the face of varying workloads and user demands.



“We feel very confident that whatever we run into, we’ll be able to scale the Snowflake/ Looker solution to meet the performance requirements of our pharmacy customers.”

JOHN FOSS, VICE PRESIDENT OF STRATEGIC PRODUCTS, PDX, INC.

SOLID INTEGRATION WITH SOFTWARE DEVELOPMENT TOOLS AND LIFECYCLE

Finally, to support data-driven applications, embedded analytics solutions need to integrate with your software development process and lifecycle. Data-driven applications, largely delivered as SaaS solutions, go through rapid release cycles and constantly evolve. Data warehousing and analytics solutions to support embedded analytics should provide the technical integration, software lifecycle management support and versioning capabilities required by modern application development.

Technology for embedded analytics should be easy for developers to integrate with modern applications.

Capabilities such as the ability for the visualization component to be deployed as an embedded iFrame, or using JavaScript, RESTful APIs, and support for scheduled or programmatic delivery of reports by email or webhook, are examples of features to look for when choosing the right technology.

Technology to support embedded analytics also needs to support the development lifecycle. For one, versioning of data models is necessary to ensure reliable, consistent deployment to production. It’s also important to support the standard software lifecycle, from development to testing to deployment, with capabilities such as instant and easy cloning and promotion of data between environments.

CASE STUDY

How PDX revamped with embedded analytics in just three months

SNOWFLAKE AND LOOKER GIVE USERS FLEXIBILITY AND INSIGHTS INTO DATA

“If you go to a retail pharmacy today and get your prescription filled, chances are you’re being serviced through PDX software,” PDX’s Vice President of Strategic Products, John Foss says. “Our software is used to print labels for the bottles, check for drug interactions and help the store fill the prescription in compliance with state board requirements.” PDX provides more than 10,000 pharmacies and pharmacy chains across the U.S. with vital technology they need to run their business, and helps pharmacies improve patient outcomes and remain competitive.

PDX set a challenge for itself: Build a scalable and flexible analytics service, which could be delivered easily to its pharmacies across the country. After an internal proof-of-concept against various databases and analytics vendors, PDX found a clear winner with the combination of Snowflake and Looker. “Within a week of starting the competition, the Snowflake/Looker team was putting data into the Snowflake cloud-native data warehouse,” Foss said. Within 12 weeks, the product was ready for the market, and PDX’s ExploreRx application went live in March 2016.

The results have been exemplary. “We’re running queries against Snowflake that have over 800 lines of code and are getting results from that kind of complex query in just a minute or two, if not faster,” Foss said.

Customers love the broad array of dashboards and reports they have available but it’s the up-to-the-minute insight that really changes the game. “There are many features that users set up once and then run on a scheduled basis to track sales and margins,” Foss adds.

ExploreRx has been a winning embedded analytics application for both the company and its pharmacy customers. “It’s been a big step forward for us to provide centralized access to data and give our customers the tools they need to change data into actionable information. By any and all measures, ExploreRx is a big success.”

Why PDX chose snowflake and Looker



Ease of use

A single interface with one portal page to log into your accounts, create new databases and warehouses, and manage permissions with a reporting tool that is easy to connect and can produce reports quickly.



Scalability

Snowflake allows PDX to match user growth with the scalability and elasticity of a cloud data warehouse.



Security

Snowflake provides security up to the application layer, with a HIPAA compliant environment. Data is encrypted end-to-end, at rest and in-transit, and strengthened with customer-managed encryption keys, Snowflake encryption keys and user credentials.



Analytics

Snowflake's native integration with analytics visualization tools allows BI users to intelligently push down query processing to Snowflake in ways that take advantage of easy exploration and analysis of structured and semi-structured data (e.g. JSON, Avro, XML, etc).



SAAS Approach

The Snowflake SaaS approach eliminates the distraction of managing infrastructure, tuning knobs, managing indexes or struggling to scale systems.




Performance

Snowflake performed astonishingly better than the competition, right out of the box. Analytics visualization tools could run alongside data loading without contention or performance impact.

ACTION	LEGACY DATABASE (EXISTING)	HADOOP ON CLOUD	SNOWFLAKE
Initial stage load	1800 mins	incomplete	12 mins
Process large file	60 mins	18 mins	8 mins
ETL process	86 mins	Incomplete	12 mins
EDW processing	780 mins	Incomplete	7 mins
Record processing	200k / hr	Incomplete	16MM / hr





Find out more

Modern cloud-built business intelligence platforms and data warehouse solutions provide the right resources and the computing capacity to process and analyze complex workloads. Give your organization access to rich data that will drive better decisions. Find out more at snowflake.net and looker.com.

About Snowflake

Snowflake is the only data warehouse built for the cloud. Snowflake delivers the performance, concurrency and simplicity needed to store and analyze all data available to an organization in one location. Snowflake's technology combines the power of data warehousing, the flexibility of big data platforms, the elasticity of the cloud, and live data sharing at a fraction of the cost of traditional solutions. Snowflake: Your data, no limits. Find out more at snowflake.net.

About Looker

Looker is a complete data platform that offers data analytics and business insights to every department and easily integrates into applications to get data directly into the decision-making process. The company is powering data-driven cultures at more than 800 industry-leading and innovative companies such as Sony, Amazon, The Economist, Kohler, Etsy, Lyft and Kickstarter. The company is headquartered in Santa Cruz, California, with offices in San Francisco, New York, London and Dublin, Ireland. Investors include CapitalG, Geodesic Capital, Goldman Sachs, Kleiner Perkins Caufield & Byers, Meritech Capital Partners, Redpoint Ventures, First Round Capital, Sapphire Ventures, and PivotNorth. For more information, connect with us on LinkedIn, on Twitter, Facebook, G+ and YouTube.