

Embedded Analytics for SaaS Companies

A practical guide for SaaS product leaders evaluating analytics platforms based on real world experience, not vendor hype.





Table of Contents

Introduction **Page 2**

Eight guiding principles for platform evaluation **Page 3**

Four critical areas for platform evaluation

1. Self-service experience **Page 6**

2. Data management capabilities **Page 8**

3. Deployment architecture and flexibility **Page 11**

4. Embedded capabilities **Page 14**

Conclusion

Building a future-proof analytics strategy **Page 17**

Bonus

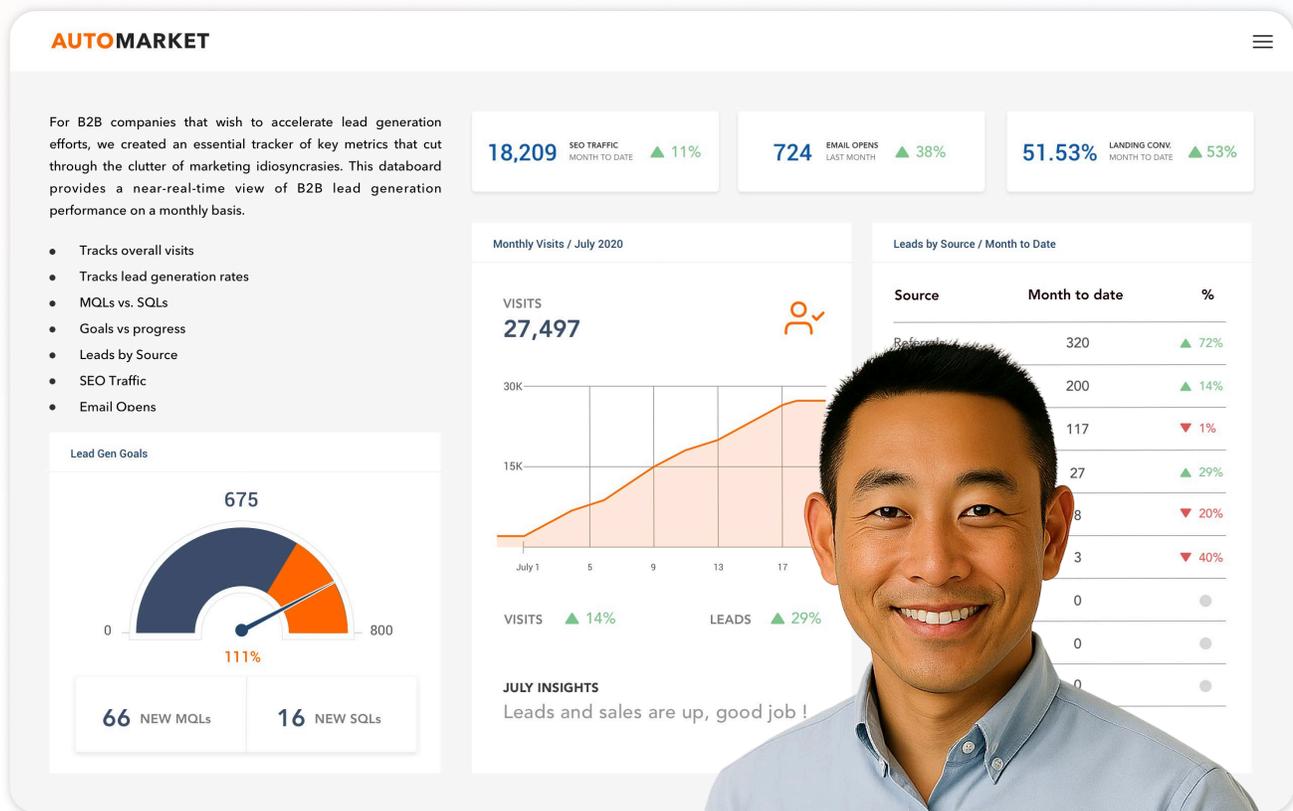
12 Questions to Ask Before Choosing an Analytics Platform **Page 19**

Introduction

The SaaS application market is highly competitive. Vendors face immense pressure to innovate rapidly, while customers increasingly need to justify their return on investment. Embedded analytics play a crucial role in helping SaaS companies retain customers, win new business, and drive new revenue streams.

Many SaaS companies are evaluating platforms for deploying analytics to their customers. This article aims to help product managers, technologists, and business leaders in SaaS companies evaluate analytics platforms based on their real-world requirements.

This document is not written for search engines or by AI. It is based on real-world experience at Qrvey and extensive research of the analytics market.



Eight guiding principles for platform evaluation

Let's start with some guiding principles for your evaluation that can save you time, help you to cover more ground, and avoid spending time on undifferentiated product areas.

1. Multi-tenancy: don't assume feature parity

One of the most significant differences between a SaaS company and an IT department is that the SaaS company deploys to its customers—meaning their analytics capabilities must support multi-tenancy. Some embedded analytics platforms, like Qrvey, are purpose-built for this scenario. However, while many general-purpose BI platforms are marketed to both SaaS companies and IT organizations, their capabilities often differ between single- and multi-tenant deployments.

A common pitfall is assuming that all advertised features will work the same way in a multi-tenant environment. Many platforms offer reduced or altered functionality when deployed across multiple tenants. It's essential to validate how well a platform supports multi-tenancy in practice—not just in theory.

2. Avoid being misled by visual appeal

It's easy to be drawn to visually impressive features—but these rarely determine long-term success. Most analytics platforms offer attractive interfaces and interactive visualizations, which, while important, can also be distracting.

It's reasonable to compare visual experiences and have preferences. But remember: these are not the capabilities that will make or break your deployment. Most SaaS companies will customize the visual layer to match their core application's branding.

This guide focuses instead on the less glamorous—but far more critical—platform capabilities that drive real success in SaaS analytics.

As we move beyond visual appeal, another fast-evolving area that demands attention is artificial intelligence. While AI features are increasingly prominent, their true value lies in how they're integrated and sustained over time.

3. Evaluating AI: prioritize flexibility over novelty

AI is at the heart of a sweeping transformation across the technology landscape, and analytics platforms are no exception. However, the AI features we see in these platforms today are evolving at an extraordinary pace.

Today's cutting-edge features may become obsolete within months. That's why it's critical to evaluate AI in analytics not by the novelty of current features, but by the platform's ability to flexibly integrate new AI capabilities as they emerge—and to progressively surface them as meaningful, customer-facing functionality. Embedding analytics into your SaaS application is a multi-year commitment. In this context, overemphasizing short-lived AI features can be a strategic misstep.

4. Don't overinvest in commonplace features

While embedded analytics platforms boast hundreds of features, not all features are created equal when it comes to driving success in a SaaS context. Many capabilities—like interactive dashboards and standard visualizations—are ubiquitous and offer little differentiation. Spending time evaluating these table stakes is inefficient. Instead, focus your attention on the capabilities that define the frontier of analytics innovation—the ones that separate a platform that merely functions from one that fuels growth. In each section of this guide, we highlight the standard features you can assume are present, so you can concentrate your evaluation on the edge capabilities that truly impact your ability to deliver differentiated, scalable, and customer-centric analytics.

5. Use analytics to differentiate your product

Most SaaS companies provide some form of analytics alongside their core product, either built in-house or powered by third-party platforms. The analytics platform you choose directly shapes your product's differentiation—strong, customer-relevant analytics capabilities can elevate your entire offering above the competition. As you evaluate platforms, prioritize those with distinctive features that will resonate with your users and set your product apart.

6. Plan for evolving requirements

Internal BI deployments typically address well-understood and relatively stable business and technical requirements. In contrast, SaaS companies—even those with decades of experience—often face unpredictable future needs driven by shifting market conditions, competitive dynamics, regulatory changes, product evolution, and corporate events like mergers or acquisitions.

Rather than evaluating analytics platforms based on rigid, predefined requirements, SaaS companies should prioritize flexibility. The right platform must not only meet today's needs, but also adapt to tomorrow's unknowns.

7. Value-added services also matter

Beyond product features, the services a vendor provides can be a powerful accelerator for your analytics strategy. These include onboarding support, solution architecture, advisory, training, and ongoing customer success engagement. Critically, the right partner can also help you shape your

monetization strategy—identifying how analytics can deliver differentiated value to your customers and uncovering new revenue opportunities. As you evaluate platforms, consider not just what the product does, but how the company helps you bring it to market, drive adoption, and turn insights into impact.

8. Prioritize proven experience serving SaaS companies

The capabilities of the platform are only part of the equation—equally important is the company behind it. A vendor’s depth of experience in serving SaaS companies, their understanding of multi-tenant architectures, and their history of successful deployments can be strong indicators of future success. Look for vendors who demonstrate thought leadership, offer proven implementation frameworks, and have a track record of supporting complex, evolving SaaS environments. A knowledgeable partner can help you avoid common pitfalls, accelerate time to value, and adapt as your product and market evolve.

With these principles in mind, let’s turn to the four core areas that define a successful analytics platform for SaaS.



Four critical areas for platform evaluation

The key areas for a platform evaluation are self-service experience, data management, deployment, and embedding.

Self-service experience

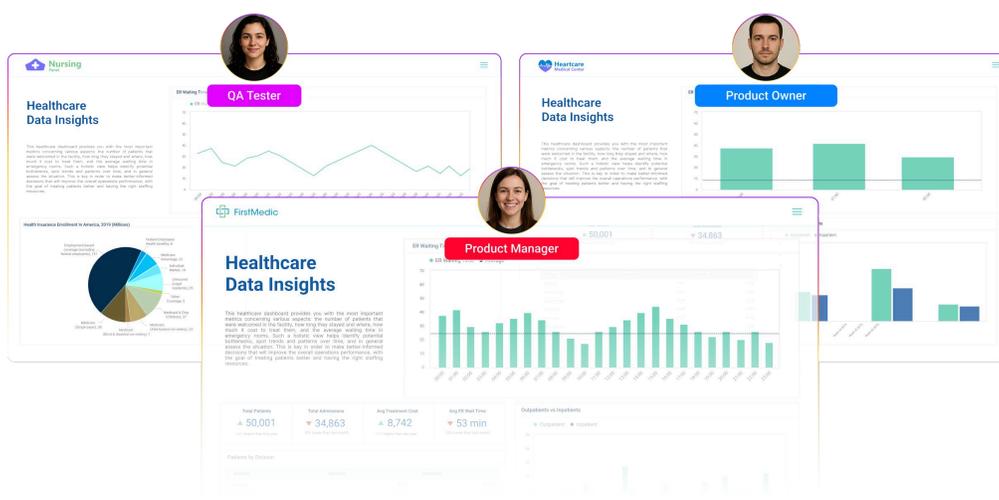
Self-service capabilities—such as building dashboards, customizing visualizations, and modifying datasets—are essential for empowering customers and reducing support overhead. These features let users tailor their analytics, reducing change requests for your development team. This boosts customer satisfaction and streamlines operations. However, as noted earlier, multi-tenancy introduces unique challenges. Features that work well in a single-tenant setup may not scale or function the same way across multiple tenants. Many analytics platforms promote robust self-service capabilities, but these are often optimized for single-tenant environments. In a multi-tenant context, the same features may be limited, unavailable, or require significant customization. When evaluating self-service functionality, it's critical to confirm that these capabilities are not only present but also secure, scalable, and easy to manage in a multi-tenant deployment. Look for platforms that allow tenant-level customization without compromising performance or data isolation.

Key differentiators in self-service

There are two stand-out self-service analytics features that are important to evaluate. The dashboard builder and dataset customization.

Dashboard builder capabilities

The ability for tenant users to independently build dashboards and share them with selected colleagues—without requiring support from your product team—is essential. Without this capability,



customers are likely to become frustrated, and your team may be overwhelmed by requests for custom dashboards.

Dataset customization capabilities

The ability for tenant users to customize datasets—by renaming fields, adding new ones, or creating entirely new datasets from existing ones—is a powerful differentiator. This functionality enables customers to tailor your application to their specific business needs, enhancing both flexibility and user satisfaction.

Standard features found in most platforms

Nearly all analytics platforms provide dashboards with standard visualizations—such as charts, tables, and widgets—and support basic interactivity like filtering and drill-down.

While visual design can typically be customized, this is a baseline expectation for SaaS companies. Since you'll likely tailor the visual experience to align with your brand, there's little value in spending time evaluating out-of-the-box design aesthetics that will ultimately be replaced.

Self-service: questions to ask

- » Can tenant users create dashboards and deploy them to their coworkers autonomously?
- » Can tenant users create their own datasets, with custom calculations, metrics and fields – named as they wish?
- » Are these self-service capabilities available with a commingled data model?

Self-service: platform limitations to watch for

- » SiSense supports self-service dashboard building, but only when deploying one Elasticube per tenant—making it incompatible with commingled data.
- » Looker does not offer a dashboard creation interface in multi-tenant environments.
- » ThoughtSpot lacks an embedded dashboard builder altogether, requiring deployment of the full product to enable dashboard creation.

Self-service: summary

In summary, self-service is a game-changer for SaaS companies—but only if it works reliably across tenants. Prioritize platforms that offer tenant-aware self-service tools that empower your users while keeping your operations lean and scalable.

While self-service empowers users at the front end, the platform's ability to manage data efficiently behind the scenes is just as important. Let's examine how data management influences performance, scalability, and operational complexity.

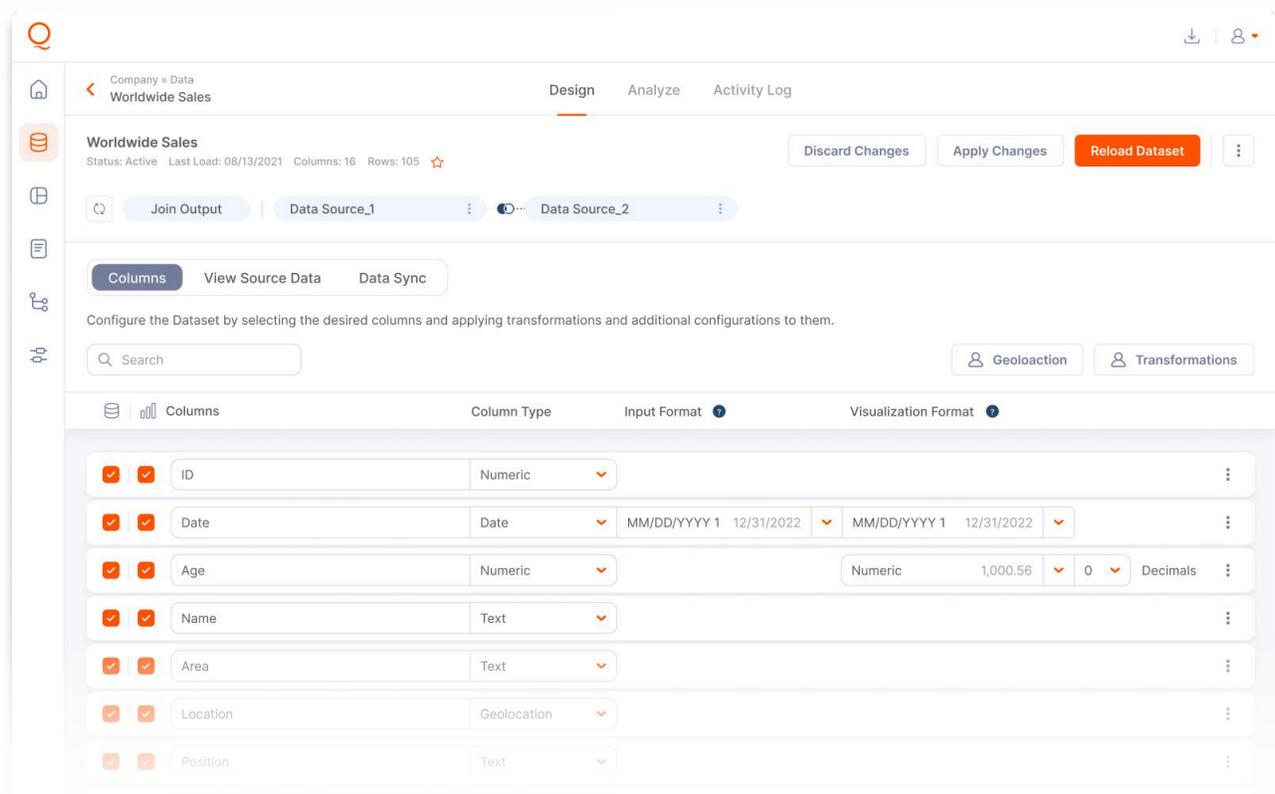
Data Management Capabilities

Data management is a foundational pillar of any analytics platform evaluation—especially for SaaS companies operating in multi-tenant environments. It includes how the platform handles data loading, transformation, storage, and retrieval. These capabilities directly impact performance, scalability, and operational efficiency.

In a multi-tenant context, data architecture becomes even more critical. Poor design can lead to increased costs, limited flexibility, and reduced self-service functionality.

Some platforms lack a dedicated data engine, requiring you to build and maintain your own multi-tenant analytics database—a complex and resource-intensive undertaking. Others rely on proprietary engines that restrict access to vendor-specific tools, limiting your ability to integrate or scale independently.

When assessing data management, look for platforms that support tenant-aware data isolation, efficient query performance, and flexible integration with your existing data infrastructure. Avoid platforms with rigid architectures or heavy customization requirements for multi-tenancy.



Key differentiators in data management

What sets platforms apart is their ability to handle data sources not optimized for analytics—like operational data stores or semi-structured formats. Platforms with built-in transformation and high-performance storage capabilities reduce the need for external tools like Snowflake or BigQuery and the engineering effort required to support them.

Semantic layer capabilities

A semantic layer maps technical table and column names to business-friendly terms and may define metrics, hierarchies, and calculations. It ensures a consistent, intuitive experience for end users and enforces alignment on business definitions.

Built-in analytic database

An analytic database is a built-in data engine that boosts performance and controls costs. For SaaS companies, it's essential that this engine supports both co-mingled and segregated data models and scales efficiently across tenants.

Data transformation capabilities

Data transformation enables raw source data to be reshaped for analytics. When paired with an analytic database, it allows high-performance analysis from virtually any data source.

Why a semantic layer alone isn't enough

Some advocate for using only a semantic layer to manage analytics metadata, avoiding direct data handling in the platform. While this may work in theory—assuming a perfectly tuned, high-performance database—most SaaS companies don't operate in such ideal conditions.

In practice, a built-in data engine offers two key advantages:

Performance: Source systems, especially those backing SaaS applications, often can't support the query load of analytics users. A dedicated engine ensures fast, reliable performance without requiring a major data engineering effort.

Cost: Offloading analytics to external databases can drive up costs. A built-in engine helps contain those expenses while simplifying operations.

Including a data engine is a pragmatic choice that balances flexibility, performance, and cost-efficiency.

Standard data management features

All analytics platforms can successfully access high-performance analytic databases with schemas specifically designed for analytics.

Data management: questions to ask

- » Does your platform include a data engine?
- » Does your platform require a third-party high-performance analytical database like Snowflake?
- » Can your platform deploy with an operational data store (duplicate of production environment) as a source system and still offer secure, high-performance analytics?
- » Does your data engine support both co-mingled and segregated multi-tenant data?
- » Is the data engine standards-based or proprietary?
- » Can data stored in your engine be accessed via third-party applications using a standard interface such as SQL?
- » Can we retain the data engine and data contained therein if we switch to a different analytic platform?
- » Does your platform include data transformation capabilities?

Data management: platform limitations to watch for

- » Some vendors—such as Looker and Omni—take a semantic-layer-only approach, lacking a built-in data engine or transformation capabilities. This makes them better suited for internal deployments backed by high-performance databases like Snowflake or BigQuery.
- » SiSense includes a proprietary data engine (Elasticube), but deploying multi-tenant self-service requires one Elasticube per tenant—an approach that can lead to high operational costs and administrative overhead. You cannot access Elasticube data outside of SiSense, and if you change analytics platforms you will need to redesign and rebuild your data for another engine.
- » ThoughtSpot includes a data engine (Falcon), but it cannot be deployed in a multi-tenant environment.

Data management: summary

In summary, robust data management is essential for delivering fast, reliable analytics at scale. Choose a platform that simplifies multi-tenant data operations, supports open standards, and aligns with your long-term architectural goals.

Strong data management lays the foundation, but how the platform is deployed can make or break your ability to scale and operate efficiently. Deployment architecture is the next critical area to evaluate.

Deployment architecture and flexibility

When evaluating an embedded analytics platform for SaaS companies, deployment capabilities are a critical factor to consider, yet they are often overlooked or skimmed through. Deployment encompasses the platform's support for multi-cloud deployment, the ability to easily migrate content between development, testing, and production environments, and the overall architecture of the deployed software.

Some analytics platforms operate only as SaaS and restrict you to a single cloud provider. This restriction can limit your ability to choose the best cloud environment for your needs and may raise concerns about data residency and security for your customers.

Furthermore, some platforms deploy using outdated technologies like manually managed virtual machines or servers rather than a containerized model, which can hinder performance and scalability. Managing outdated architectures is costly and inefficient.

The ability to easily migrate content between development, testing, and production environments is a critical need for SaaS companies because they deploy analytics as part of a broader application portfolio. Few analytics platforms offer built-in content migration, and it is important to evaluate the ease with which the platform can be integrated with your preferred content deployment approach.

Key differentiators in deployment

Key areas of differentiation include support for multi-cloud deployment, modern deployment technologies, and content migration across environments.

Multi-cloud deployment support

The ability to deploy across public clouds—such as AWS, Azure, or Google Cloud—gives SaaS companies the flexibility to align analytics with their infrastructure strategy, including multi-cloud scenarios.

Modern deployment technologies

Modern deployment methods like serverless and Kubernetes simplify scaling and operations. In contrast, older approaches using virtual machines or manually managed servers increase complexity and overhead.

Managing environments and content migration

Robust support for migrating content between development, staging, and production environments is essential for integrating analytics into existing deployment workflows. Gaps in this area can create friction and slow down release cycles.

Deployment: questions to ask

- » Can the platform be deployed inside my cloud alongside my SaaS application?
- » Do you support multiple clouds (AWS, Azure, Google, other)?
- » Does the platform have built in tools to migrate content between environments like development, testing and production?
- » Does the platform deploy using modern technologies like serverless or Kubernetes containers?
- » Does the platform deployment require manual management of virtual machines or servers?

Deployment: platform limitations to watch for

- » Omni, Luzmo, and QuickSight are SaaS-only.
- » Looker is deployed manually using EC2 or Azure VMs.
- » QuickSight is AWS-only.
- » PowerBI only deploys on Azure.
- » Tableau does not provide VMs or server images, requiring you to manually install onto a Linux server.
- » ThoughtSpot provides only EC2 images.

Deployment: summary

In summary, evaluating the deployment capabilities of an analytics platform involves assessing its support for multi-cloud deployment, ease of content migration between different environments, and the architecture of the deployed software. Be cautious of platforms that are SaaS-only and restricted to a single cloud provider, as well as those that use outdated technologies. Ensuring that the platform you choose offers modern, flexible, and efficient deployment options is essential for maintaining optimal performance and meeting your customers' needs.

Once the platform is deployed, the final step is making analytics accessible and seamless for your users. This is where embedding capabilities come into play.

Before diving into embedding, it's worth addressing a common question: if SaaS companies build SaaS products, why do they often avoid SaaS analytics platforms? The answer reveals deeper concerns about control, security, and customer trust.

Why SaaS companies avoid SaaS analytics platforms

It may seem counterintuitive, but SaaS companies often prefer deployed analytics platforms over SaaS-based ones. The typical benefits of SaaS—like ease of purchase and minimal IT involvement—are less relevant, since deploying and managing cloud software is already a core competency.

More importantly, SaaS companies are entrusted with sensitive customer data. Introducing a third-party analytics provider that stores or processes this data in a separate cloud raises concerns about data residency, privacy, and security—along with potential liability for breaches beyond the company's control.

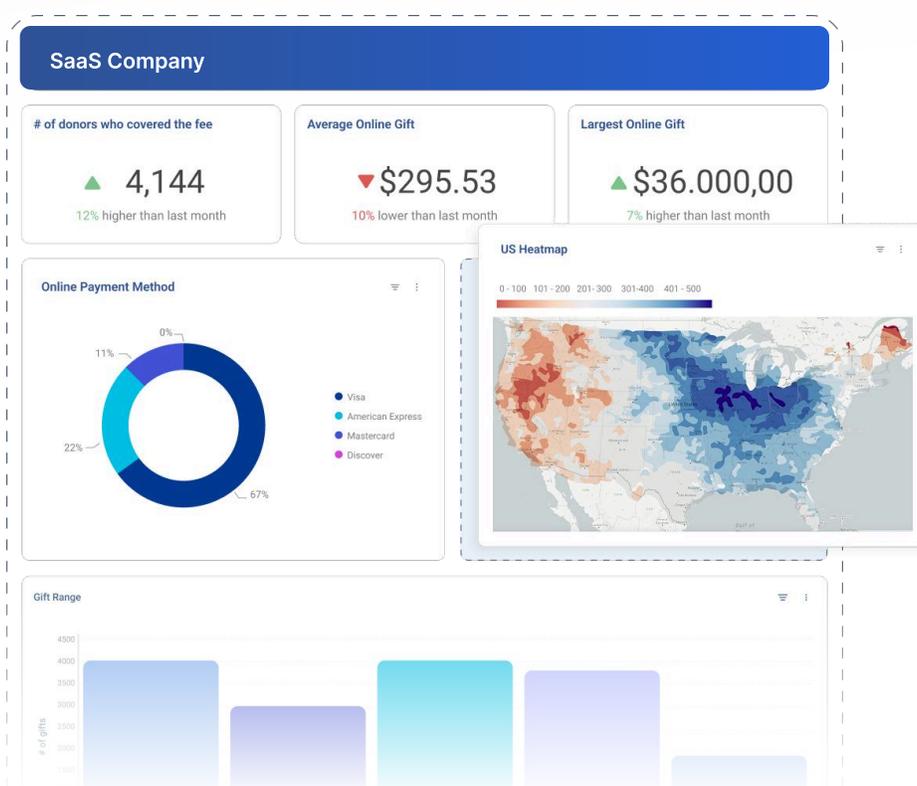


Embedding capabilities

Once you've selected a platform and addressed deployment concerns, the final step is making analytics feel like a seamless part of your product. That's where embedding comes in. Embedding involves the front-end technologies used for integrating components like charts, dashboard explorers, or dashboard builders into your application, as well as the backend platform APIs that facilitate integration with your SaaS application and operations.

Embedding capabilities significantly influence both user experience and operational efficiency. Modern front-end embedding technologies are critical to delivering a seamless, white-labeled analytics experience that feels like a native part of your product. In contrast, outdated methods like iframes often result in a fragmented interface, exposing visible seams between your core application and embedded analytics. These legacy approaches can also introduce security vulnerabilities, posing risks to both SaaS providers and their customers. To avoid these pitfalls, it's essential to choose a platform that offers secure, modern embedding capabilities designed for smooth integration and scalability.

Backend platform APIs are equally critical to your operations. These APIs enable the automation of integration, provisioning, and content migration, which are essential for maintaining efficient and scalable operations. Without reliable and well-documented APIs, integrating the analytics platform with your existing systems can become a complex and time-consuming task.



When evaluating embedding capabilities, keep in mind that this may not be a priority for internal deployments—one reason why some general-purpose BI platforms offer limited support in this area. Internal use cases typically don't require the same level of integration or customization as customer-facing applications, making advanced embedding technologies and APIs less critical.

Key differentiators in embedding

Key differentiators include front-end embedding technologies and the scope of backend APIs.

Front-end embedding technologies

Modern platforms support embedding through web components rather than iframes, and allow embedding of various elements—charts, tables, dashboards, even dashboard builders—directly into your SaaS application.

Back-end API capabilities

While most platforms offer REST APIs, their coverage varies. Comprehensive APIs should expose all major components—users, groups, datasets, semantic layers, dashboards—enabling automation of provisioning, deployment, and integration with external systems.

Embedding: questions to ask

- » Does the platform offer front-end embedding and customization without iframes using modern web components?
- » Is every product capability exposed to tenant users embedded, or do some capabilities require tenant users to operate the analytics platform applications directly?
- » Does the platform include an embedded dashboard builder?
- » Can every dashboard, widget, and visualization be customized in terms of appearance and content (i.e., can colors, fonts, and data all be changed) by the tenant user?

Embedding: platform limitations to watch for

- » Looker and Omni embed using iframes.
- » Looker and ThoughtSpot cannot embed a dashboard builder.
- » Tableau's APIs are limited and inconsistent.
- » Omni backend APIs are limited and immature.
- » Many developers view iframe embedding unfavorably due to security concerns, limited customization, performance issues, responsiveness and browser compatibility

Embedding: summary

In summary, evaluating the embedding capabilities of an analytics platform involves assessing both the front-end technologies for embedding components and the backend platform APIs for integration. These capabilities are crucial for ensuring a positive customer experience and maintaining efficient operations.

With self-service, data management, deployment, and embedding thoroughly evaluated, you're ready to make a strategic decision. But how do you ensure that decision stands the test of time?

Conclusion: Building a future-proof analytics strategy

Choosing the right analytics platform isn't just a technical decision—it's a strategic one that will shape your product, your operations, and your customer experience for years to come.

While it's tempting to chase flashy features or short-term trends, long-term success depends on selecting a platform that is flexible, scalable, and purpose-built for SaaS. The most effective platforms evolve with the AI landscape, support multi-tenant architectures, and empower your team to deliver differentiated value at scale.

Focus your evaluation on the capabilities that matter most: self-service, data management, deployment, and embedding. These are the pillars of a resilient analytics strategy.

Choose wisely—and you won't just meet today's needs. You'll be ready for whatever comes next.

Here's a quick-reference summary of the key areas and considerations covered in this guide.

Quick reference: core evaluation areas

Capability Area	What to Look For	Why It Matters
Self-Service	Tenant-level dashboard building and dataset customization	Empowers users, reduces support load
Data Management	Built-in data engine, transformation, semantic layer	Ensures performance, scalability, and cost control
Deployment	Multi-cloud support, modern deployment tech (e.g., Kubernetes), content migration	Aligns with your infrastructure, simplifies operations
Embedding	Web component-based embedding, full-featured backend APIs	Delivers seamless UX and operational automation

Common pitfalls to avoid

- » Assuming multi-tenancy works the same across platforms
- » Overvaluing visual polish over backend capability
- » Relying on semantic layers without a data engine
- » Choosing SaaS analytics platforms that conflict with your data privacy model

Strategic considerations

- » Flexibility: Can it evolve with your product and market?
- » Differentiation: Does it give your product a competitive edge?
- » Scalability: Will it support growth without re-architecture?
- » Control: Do you retain ownership of data and deployment?

Final takeaway

“ Choose a platform that doesn't just meet today's needs
choose one that's ready for what's next. ”

12 Questions to Ask Before Choosing an Analytics Platform

1. Can tenant users build and share dashboards independently?

Look for true self-service capabilities that reduce support burden and empower users.

2. Can users customize datasets to fit their business needs?

Field renaming, calculated fields, and dataset creation should be intuitive and tenant-aware.

3. Does the platform support secure, scalable multi-tenancy?

Ensure it handles co-mingled and segregated data models without compromising performance or isolation.

4. Is there a built-in data engine and transformation layer?

Avoid relying solely on external databases or semantic layers for performance and cost control.

5. Can the platform evolve with your product and market?

Flexibility is key—requirements will change over time.

6. Does it support modern deployment models (e.g., Kubernetes, serverless)?

Outdated server- or VM-based deployments can limit scalability and increase operational overhead.

7. Can you deploy in your cloud(s) of choice?

Multi-cloud support is essential for compliance, performance, and control.

8. Are content migration and environment management built-in?

Smooth promotion from dev to prod is critical for SaaS workflows.

9. Does it offer seamless embedding with full customization?

Look for web components (not just iframes) and APIs that expose all key functionality.

10. Will it differentiate your product in the eyes of your customers?

Choose a platform that enhances—not just supports—your product's value proposition.

11. Does the vendor offer advisory services to help shape our analytics or monetization strategy?

Look for a partner who can advise on how to turn analytics into a competitive advantage and deliver differentiated value to your customers.

12. What experience does the vendor have working with SaaS companies?

Choose a vendor with a proven track record of success in multi-tenant deployments in production environments.

ABOUT QRVEY

Qrvey is the leader in multi-tenant embedded analytics. With unmatched experience and expertise, Qrvey is the partner of choice for SaaS leaders that want to deliver insight to their customers, agility to their product teams, and growth to their business.

Qrvey offers the most powerful and flexible embedded analytics platform for the SaaS market. With over 20 years of experience, we provide industry-leading guidance and support to ensure our clients achieve their analytics goals. Our deep understanding of multi-tenant SaaS architecture and our comprehensive range of services make us the trusted partner of choice for SaaS leaders.

Learn more at qrvey.com

Embedded analytics designed specifically for SaaS companies.



Insight

- ✓ True self-service with unlimited customization
- ✓ AI-driven insights
- ✓ No-code workflow automation



Agility

- ✓ End-to-end embedded analytics platform
- ✓ Native multi-tenant security
- ✓ Flexible multi-cloud deployment



Growth

- ✓ Flat-rate pricing for scale
- ✓ Unmatched monetization opportunities
- ✓ Embedded Services

Win more business with analytics

See how leading SaaS companies are creating value and differentiation, unlocking new revenue streams, and earning customer loyalty with Qrvey.

[Book a Demo](#)