



Embedded Analytics Bear Traps (and How to Avoid Them)

Embedded Analytics Bear Traps (and How to Avoid Them)

We've all read the articles extolling the benefits of data monetization, of embedding analytics in our products, and making use of all that data sitting idly in our systems. But why haven't more businesses taken this advice? Have they decided they don't want the new revenue streams or the increased user engagement? Or—more likely—have they started their data product journey only to get caught in one of the many “bear traps” that litter the path to data product success?

At Izenda, we've seen our fair share of project teams that have gotten a leg or two caught in one of the data product bear traps and it's always frustrating. It's frustrating because it doesn't have to be this way. As brutal as they are, these bear traps are easily avoided—once you know what to look for. And that's the trick to success: take the journey to building a data product for your users, but avoid the traps that catch so many along the way.

Trap #1: A Product That's a Nightmare to Maintain

Of the three bear traps we'll discuss, this first one is perhaps the most insidious. It's like a trap that you easily spot along the side of a trail, step back to avoid it, only to find your foot caught in another trap that was hidden behind you. It's sneaky like that. And here's why: just like the hidden trap alongside the trail, this trap goes unseen until it's too late to avoid.

The trap gets you when you base your data product on a business intelligence platform that is impossible to maintain. Most teams are on the lookout for the obvious maintenance activities: adding data sources, cleansing data, adding users,



changing metric definitions, etc. but this is a little different.

This trap gets you because the needs of the data product owner are different than those of an enterprise analytics team. These needs are product lifecycle related and often only arise once the product has been launched.

Many analytics platforms focus on one of two areas: the data extraction and cleansing or the data visualization. While these are essential elements of a good data product, alone they are not enough. If you want to make a great data product you need to be able to easily perform core activities such as:

- add new users
- change tenant and user permissions
- connect to new data sources
- blend “master data” with other sources to create derivative analytics
- prepare the data for self-service analytics
- schedule reports and alerts

- roll out updates as you add functionality
- add functionality to customers as they progress from “basic” to “plus” to “pro” analytics tiers
- check usage to make sure customers aren’t exceeding allotted utilization
- audit data updates to ensure that sources are being refreshed properly

Although not as glamorous as data visualization, these are the elements essential to keeping a data product operational over the entire customer lifecycle. And unfortunately, this functionality isn’t present in many of the analytics platforms on the market today.

The builders of these platforms, born in the era of “inside the company, enterprise analytics”, never anticipated the need to push dashboards and other analytics to tens of thousands of customers and to support a full product lifecycle over time. They were built for internal users—not for paying customers likely to express their dissatisfaction by buying elsewhere.

As a result, many of the activities that we as data product owners consider to be “table-stakes” simply aren’t present on these “data product unready” platforms. Want to onboard a new customer to your data product? That requires writing new ETL processes, manually generating a new customer instance, manually adding new users, and manually ensuring that the right users’ permissions are set so they see the right analytics at the right time. Want to roll out version 2.0 of your data product across thousands of your customers, some of whom have slight changes to

the standard product? Sorry, on these platforms you’ll be doing that manually.

The platforms that were designed for internal marketing or sales teams don’t have this essential functionality because the needs of enterprise users and data product users are simply... different.



Unfortunately, it’s extremely easy to fall into this bear trap and pick a platform that’s not designed for product life cycle activities. All these manual tasks don’t seem too terribly time-consuming when you’re in “proof of concept” mode with an analytics vendor. In fact, many of the

vendors will offer to do these activities for you during your assessment phase to shield you from the cumbersome manual tasks required. But extend your gaze out just a year or two past launch.... The tasks that were easy with just one or two beta customers (and vendor resources) become massively more complicated when you’re operating at scale, with 100 or 1,000 customers. If you got caught in the bear trap of selecting an analytics platform that doesn’t have a full product lifecycle management ecosystem as part of its suite of functionality you now have a very big problem.

The good news is, you can avoid this bear trap if you know what to look for when you’re evaluating analytics vendors. In addition to the standard questions about data connectivity and data visualization, ask your vendor about product operations at scale.

Ask things like:

- How many customers (not users—customer instances) does your largest data product support?
- How many resources does your largest customer use to support their data product's daily operations?
- How long does it take the average customer to bring a new client online?
- Can I copy templates of dashboards and other analytics across customer instances?
- Is there some form of version control across ALL my customers' analytics?
- If I have a large customer base and some of the users have a “non-standard” version of the analytics, is it possible to perform mass updates?
- Can updates be automated (without writing a bunch of code or scripts) or would we have to perform these upgrades manually?

When you ask these questions, you'll quickly separate the vendors that have really considered the details of data product management from those that have simply given it cursory thought. Those that live and breathe data products will have answers for these questions (and examples) while those that have simply extended enterprise analytics into the data product arena will have a bit of trouble providing concrete and acceptable answers.



Trap #2: Not Ready for Tomorrow

If the first bear trap—lack of a product lifecycle management ecosystem—was one that snuck up on you from the edge of the trail, this next bear trap is even worse. This one lets you get all the way to the end of the trail, only to realize you've gone the wrong way and you're now boxed in. There's no escape. You're trapped.

This is the trap of a lack of future-proofing, and it seduces you into thinking you're fine by showing you flashy functionality and cool features that

work well for today's needs but will eventually lead to your downfall.

When teams think of future-proofing their analytics product, they often start thinking about the problem of “scaling”...

- Can we store billions of rows of data?
- Can we support tens of thousands of users in a multi-tenant environment?

While these are essential considerations (and should be asked about), most modern analytic platforms handle these types of scaling issues with ease.

At Izenda, when we think about future-proofing for data products, we think a little differently. We think of functionality limitations that will prevent you from offering the data product your customers demand as needs grow over time.

For example, teams evaluating analytic platforms for their data product tend to focus only on today's needs.

- Will it work in the cloud?
- Can I embed analytics in my application?
- Can I show trends and targets?
- Can end users actually create their own reports?

These are important to consider, but it's how you treat the answers that can lead you on to the trail to the box canyon of data product doom.

You need to be thinking not just about today's functionality needs, but tomorrow's as well. You might only need cloud analytics today, but what if a future customer needs its analytics on premise for compliance or other reasons?

You might be fine with embedding an iframe today, but can you afford to lose customers who use mobile devices because it won't support responsive design? And will an iframed solution pass a security audit? What happens when you need customization and you realize the analytics vendor won't give you access to source code for his iframe app?

Is the analytics software consistent with your application's architecture? Will it scale as you increase your customer base? The time to consider the future is today, before selecting a platform, not once you've already invested your resources into a dead end platform.

The key to avoiding this bear trap involves a delicate balancing act. You need to look for an analytics platform that's flexible enough to accommodate both today's and tomorrow's

needs but still provides a rich interface (you don't want to be bogged down writing code for every new capability your customers require.)

Here's how you can avoid finding yourself trapped in an analytics box canyon of no escape:

1. Look for a open and extensible platform.

This requires a detailed architecture review. Those rigid platforms with brittle, bolt-on embedded capabilities might look great, at first. They often come with claims of low cost and speedy implementation times, but they can cause problems as soon as you deviate from their narrowly defined use case.

2. Plan your future needs.

Start by defining your "phase one" user roles and requirements, but don't stop there. Extend your plan to future user roles, future capabilities, and future business models, then make sure that any platform you consider will support them today and tomorrow.

3. Don't fall for the "roadmap trick".

When you ask the tough questions of vendors, you'll often get the same old, stale answer (say it with us...) "it's on the roadmap." And, given an infinite timescale, it probably is on the roadmap. But a theoretical capability tentatively scheduled for possible consideration five years from now isn't going to solve a product need that arises in six months. If they can't do it today or commit to a date, assume they have no plans to deliver the functionality you require and move on in your search.

Trap #3: Operating Blind

We've saved the best (worst) bear trap for last. This bear trap is even more sneaky than the other two because you don't even notice that it's grabbed your ankle until you're ready for the next iteration of your data product. It's the trap of operating blind and it's easy to get caught.

You'd think that, given that they are in the BI and analytics field, most platforms would excel at utilization analytics. You might think that, when using your analytical platform to build a data product, you would have no problem seeing who's using the product, what functionality they're using, what they're not using, and other utilization information that is essential to understanding your product. But unfortunately, this just isn't the case.



With many analytics platforms on the market today, you can build fancy 3D drop-shadow charts with beveled edges and cute animations, but ask the simplest questions about who's using what, when, and how often and you're stuck. Are people using that 3D animated chart? Who knows? You're left operating blind.

The whole point of building a data product is to help deliver information and insights that help users perform their roles—why shouldn't your product team receive similar benefits? Why should your product team have to guess what's being used, what users rely upon, and how the data product is being used? Why should the product team operate blind while building a product designed to help customers... not operate blind? It sounds crazy that most data product teams are left guessing about usage and where they might need to focus future efforts

but with most analytics platforms, this is reality. You can avoid this bear trap by asking tough questions of vendors before you pick your platform. Here are a few that they need to answer:

1. Do you have built-in utilization analytics?
2. Are these analytics more than just logs showing when users logged in?
3. Can you show me the most and least frequently used parts of the system (dashboards, charts, etc.)?
4. Can you show me this information by user role (CEO vs. engineer vs. sales rep)?
5. Can you show me the actions users took, the paths they followed, and where they got stuck within the data product?
6. Are the utilization analytics summarized in such a way that the product team can easily understand where gaps are occurring and where improvement is required?

A vendor that has thought through the needs of a data product team will be able to answer these questions, and not just by saying “well, you can write code to get at that information.”

Demand that your BI and analytics platform deliver the same insights to your product team as you expect to deliver to your end users. Ask vendors to show you exactly how they deliver insights to the product team and how those insights can be used to drive change. A log simply isn't enough. They need to show patterns, trends, and correlations. Without this information, your team is simply guessing at what functionality your users need you to deliver next. As a result, you're increasing the overall risk of data product failure.

Summary

Just because there are bear traps along the route doesn't mean that a journey isn't worthwhile. The journey to delivering a data product can be rich and rewarding, both for your company and for your users. It can open new revenue streams, increase engagement, and reduce churn while helping your users better understand what's happening inside their businesses. The journey is worth taking—just be sure that you know how to recognize the most common traps and how to avoid them.

About Izenda

Izenda is a self-service BI and analytics platform purpose-built for OEM use and embedding. Izenda integrates seamlessly in applications and portals to deliver BI and analytics directly to the people who need it most – users who want to easily analyze, visualize, and share valuable data and insights in real time.

More than 10,000 organizations use Izenda on a daily basis. To find out more, visit www.izenda.com

Other Resources

Which Embedded Analytics Product is Right for You?

This new Eckerson Group report overviews the key issues you need to know before selecting an embedded analytics product.

Embedded BI & Analytics Buyer's Guide

Designed to help independent software vendors and others define and prioritize their requirements for embedded BI, this guide includes questions to ask vendors.

Build vs. Buy: 8 Key Considerations When Adding BI & Analytics to Your Software

Is it better for your team to code business intelligence functionality, or buy a third party solution? This ebook answers eight common questions on the build vs buy debate, including topics such as usability, self-service and responsiveness.