



Bicycle.ai Industry Benchmark

Revenue Leaks: **How eCommerce and Retail Companies Tackle the Silent Margin Killer**

Executive Summary

High-transaction consumer businesses live inside a technical environment that's fraught with danger. Their applications and website are architected for complex consumer interactions, but that same complexity makes it difficult to identify and resolve failures.

Failures matter, because in many cases that technical environment represents a significant source of company revenue. Latency, sluggish API calls, third-party transaction failures, and outages can be temporary issues that over time add up to massive revenue loss.

For years, organizations have relied on business intelligence and analytics platforms to understand performance and diagnose such issues. These tools provide valuable visibility, but they have been fundamentally limited. Data is fragmented across teams and systems, making it difficult to form a unified view of root causes. Analysis often requires manual handoffs between business stakeholders and data teams, slowing investigation and delaying action.

More recently, agentic AI tools have emerged that promise to break through these speed and silo constraints by connecting data, analysis, and action in a single loop. But are they being used correctly by revenue leaders, and are they actually impacting the bottom line?

To understand how organizations are approaching this shift — how they think about revenue-impacting failures, and how they structure teams and toolsets to uncover and resolve revenue leaks — Bicycle.ai conducted a benchmark survey of more than 100 leaders in product and analytics at large retail and eCommerce companies.

We uncovered startling insights:

- Over half of respondents felt they were confident about their ability to identify and resolve revenue leakage in their applications. And yet:
- Only half were able to do so before being flagged by a customer's note or complaint.
- The vast majority took hours to fix the problem, with some even taking up to a day. Furthermore, nearly half of respondents reported that the same issue resurfaced within 30 days.
- The biggest challenges for recovered revenue were integration and data quality, and:
- AI is being used but not comprehensively across all teams and tooling, suggesting the early innings of revenue leaders leveraging AI to recover revenue.

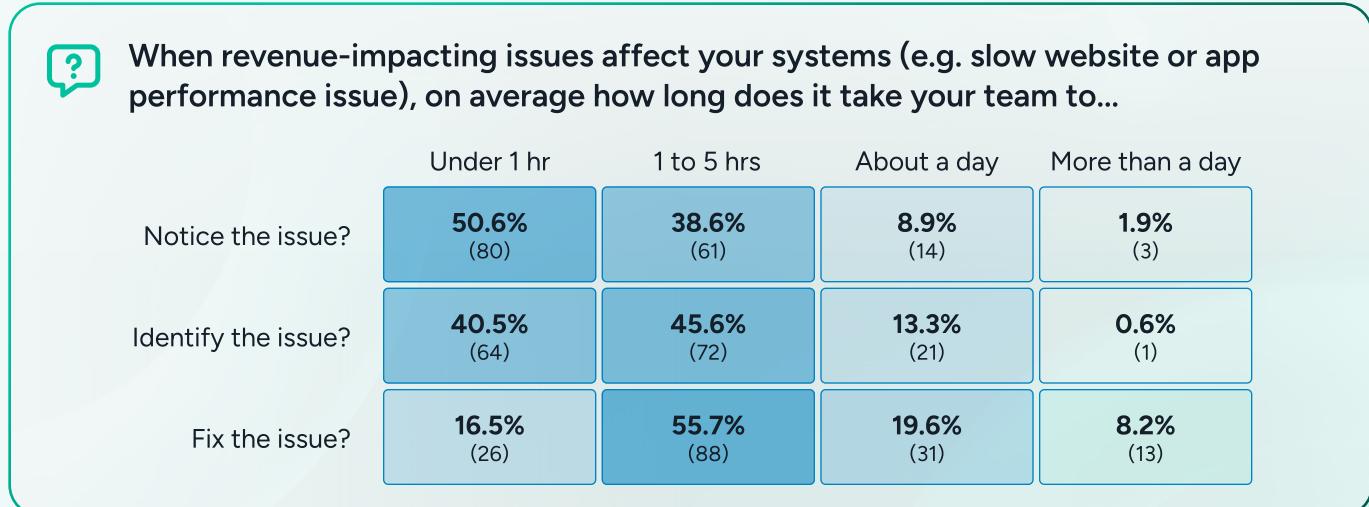
Let's unpack these insights further.

Confidence Game

One of the more surprising insights from the respondents is that over half were confident about their ability to identify and detect most revenue-impacting incidents in real time. 56% felt positive about where they stood in terms of tackling issues affecting revenue.



But findings from the rest of the survey undermine that initial confidence. Only 40% of respondents indicated that they are able to subsequently identify the root cause of problems in that amount of time, and only 16% suggested they were able to fix the problem. The general spread is 1-5 hours, with many respondents taking a day or more to notice, identify and resolve revenue leaks.



From this standpoint, it appears that the initial display of confidence is based on the idea that it should generally take more than an hour – in some cases, much more than an hour – to

identify and resolve problems. Yet it's reasonable to ask whether this should be the status quo in an operational scenario where every millisecond can potentially result in significant revenue loss.

How do revenue leaders find out about problems in the first place? More than half of the time, the alerts come from customers rather than dashboards.



It's safe to say that if awareness of critical issues comes from actual customers, then significant operational challenges remain in terms of how technical problems are triaged and fixed. Customers as the first line of defense invariably mean more inbound inquiries, more code red scrambles, and a far more reactive resolution process.

Furthermore, we can infer a challenge that will become crystal-clear later in the survey: IT and revenue performance data is spread among many systems and teams. If you have multiple teams and multiple sets of tooling focused on recovery, then the reality is that performance data won't come together in a holistic way fast enough to drive outcomes and prevent revenue leaks.

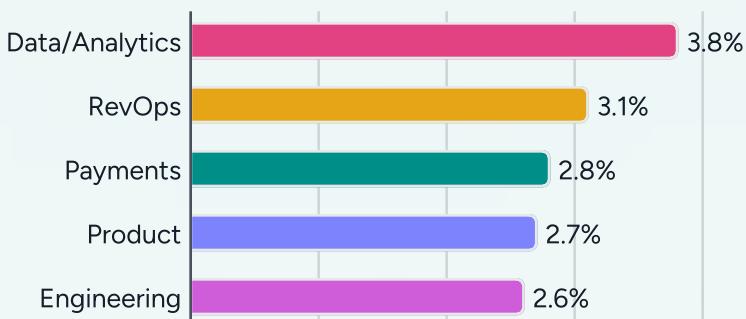
We can investigate this point further by evaluating the tooling and teams that organizations bring to bear on recovery.

The Dream Team is... Everyone

Which teams own recovery rescue? Data teams have the edge, but perhaps the surprising finding here is that they only barely have the edge. Many other teams are similarly on the hook when problems surface.



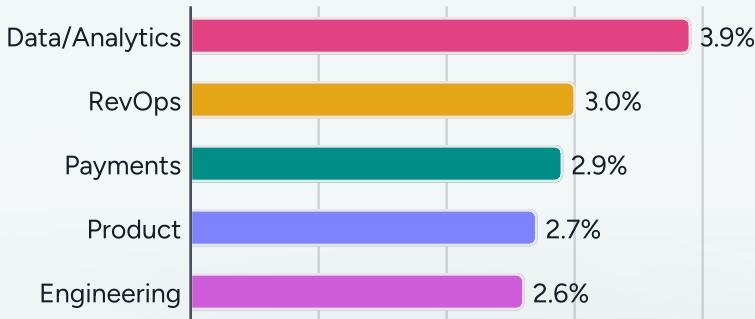
Which team currently owns “revenue rescue” responsibilities in your company?



What teams are being eyed for investment in the future? The answers are nearly identical. The data teams are highly prized, but there's a surprising reliance on other teams as well.



Which team do you plan to invest in more in the future?



The responsibility for revenue recovery seems to be diffused among multiple teams and personas, which may be necessary in the heat of the moment but likely doesn't help reduce mean-time-to-resolution.

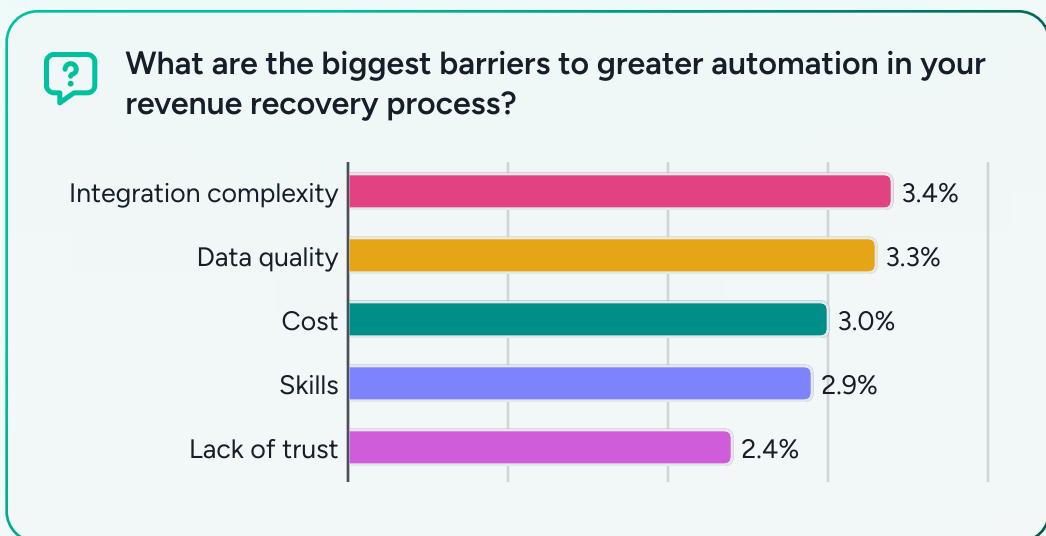
These are interesting numbers, because in their Predictions 2026: Artificial Intelligence report, Forrester indicated that data analytics teams might get less hiring due to the advance of AI. They wrote: “The percent of AI decision-makers who reported that their IT teams planned to hire more analysts this year to support AI projects was 9% less compared to 2024, and the percentage planning to hire more data scientists and data engineers fell by 6% and 5%, respectively.”

But our group of respondents are not following the Forrester drumbeat: they're doubling down

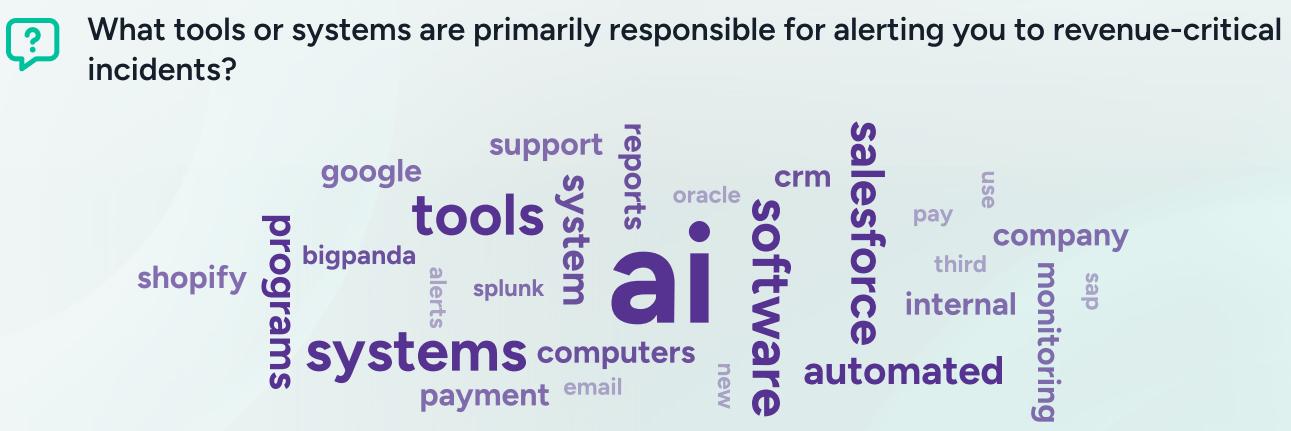
on raw headcount to solve problems. Likely it's because our group are traditional data and analytics leaders, not AI decision-makers. We'll learn more about our group's relationship to AI shortly.

The Tool for the Job

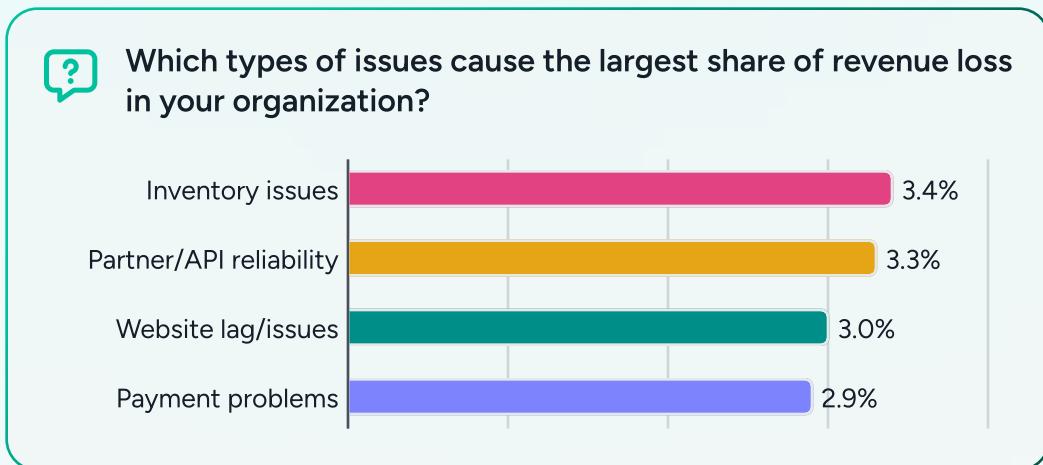
In evaluating the tools being used by these teams, it's important to understand what problems they're trying to solve. The main challenges are integration complexity and data quality.



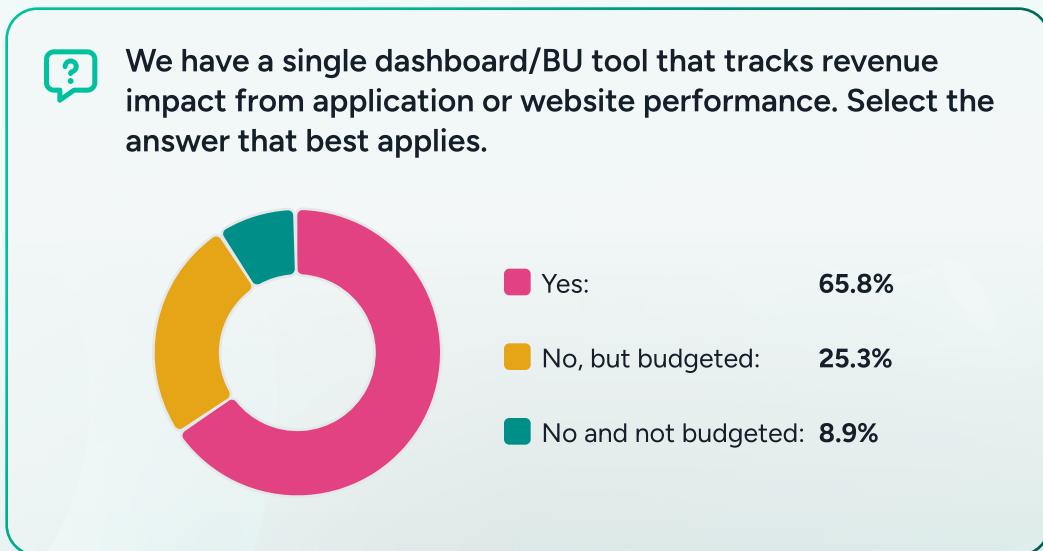
But there's little consistency as to the tools being used, which are a fragmented set of products that range from CRMs to security solutions.



We do see that AI is being leveraged by respondents, which is a promising development that we'll have more to say about in a moment. First, let's take a look at what respondents believe lie at the root cause of revenue issues:



Again, a fairly wide spread among typical problems that most retail and eCommerce companies face. Everyone does seem to have a dashboard to oversee these issues, however:



Is it possible that dashboard heroics are the source of some of the overconfidence evidenced in the report earlier? Respondents are sure that they know when problems occur. But how long does it take to solve them, how often can they do so without customers noticing or being affected, and do they have the right teams and tools for the job?

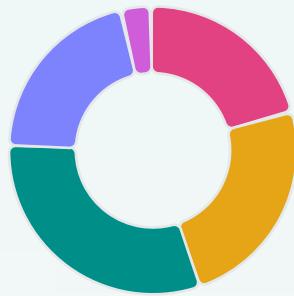
Let's talk about AI for a moment.

AI've Got Questions

The only scalable way to resolve these issues is AI, applied end-to-end across detection, diagnosis, and action. And, as we've already seen, AI is an increasingly important component of the tech stack for revenue recovery teams.



What percentage of incidents are automatically fixed or mitigated by AI-driven workflows?



0:	20.9%
1-10%:	24.1%
11-25%:	31.0%
25-50%:	20.9%
51%+:	3.2%

Remember, respondents are confident they can solve problems, and they have a dashboard to help them surface issues. But why are customers constantly the ones shouting about problems, why do technical problems take more than an hour to solve, and why do they resurface with frightening regularity?

AI could be the answer, but only if integrated throughout the entire environment – and that's integration that seems to be the biggest challenge facing revenue leaders. As a reminder, respondents told us that the biggest challenges for revenue recovery are systems integration and data quality.

Conclusions

High-transaction eCommerce and retail companies operate inside environments where complexity is unavoidable, but prolonged blind spots and slow response times simply can't be tolerated. Customers won't allow them.

Nevertheless, across more than 100 product and analytics leaders, we see a consistent pattern. Teams are confident in their ability to detect and resolve revenue-impacting issues, yet in practice their ability to resolve issues are too slow, hampered by iterations between business and data analyst teams. In fact, issues are most often discovered by customers and take an hour

or longer to fix, and frequently recur within weeks. That delay is unacceptable: every hour translates to significant lost revenue. Dashboard noise and data dispersion also slow down outcomes, and the result is a fragile recovery loop that consumes time, erodes margin, and undermines customer trust.

The root causes are well understood. Integration complexity and data quality challenges are everywhere, and ownership is spread across too many teams. Moreover, tooling is fragmented and often the byproduct of point solutions. Dashboards may create a sense of control, but they do not shorten time-to-action or prevent future crises.

AI represents a clear inflection point, but today it is applied unevenly and too late in the process. When AI is bolted onto isolated tools or limited to post-incident analysis, it cannot overcome the underlying fragmentation of systems, data, and accountability. The opportunity lies not in adding yet another layer of monitoring, but in using AI as connective tissue — continuously correlating signals, identifying revenue risk in real time, and driving coordinated action across the stack.

Consider this: it doesn't even need to be a single team that owns revenue recovery. With the right application of AI, it could be a single person — someone who leverages AI across the entire environment to notice, identify, and resolve issues in seconds.

But doing so requires a mindset shift. It's not just detection and prevention anymore; it's continuous recovery and clear accountability. It's problems that used to take hours or days now being resolved in seconds without a human having to hit the "Red Alert" button. Organizations that make this shift will recover revenue faster and solve problems before customers ever notice.

As revenue leaders stare down the face of 2026, the question is no longer whether revenue leakage exists, or even whether AI can help. The question is who will operationalize it end-to-end —and who will continue to let silent failures compound in the background with an overreliance on dashboards and overly long resolution cycles.

Will we get there? Stay tuned.



Our mission is to turn complex, high-volume transactions into real-time decisions through the power of agentic analytics, driving measurable revenue impact for businesses in retail, payments, and travel.

We're building the next generation of AI powered analytics, moving organizations from signal to action instantly. Our focus is simple: protect revenue and unlock growth for high-volume transaction businesses.

With deep expertise in enterprise data, AI, and real-time systems, we are shaping the future of agentic analytics, where businesses act, not just analyze.

 <https://bicycle.ai>