

Why Most S/4HANA Migrations Miss Their Budget and Timeline, and How to Improve the Odds

SAP transformations have a long, well-documented history of running over budget and behind schedule. With ECC mainstream maintenance ending December 31, 2027, there's less room than ever to absorb a slip, and that is pushing programs toward earlier, evidence-based risk detection, including newer AI-assisted scoring tools.

7 min read | Orpington Technologies Insights

SAP's mainstream maintenance for ECC ends December 31, 2027. That date is fixed, and it is already reshaping the market around it: as more of the installed base moves to S/4HANA, the pool of consultants, implementation partners, and specialized ECC support thins out, a pattern consistent with prior SAP platform transitions. Organizations still running ECC past that point will be working with a shrinking bench and rising support costs, which is the practical argument for migrating on a reasonable timeline rather than waiting.


The harder problem is what happens once a migration is underway. Status reporting in most ERP programs is self-reported, rolls up through several layers of management, and is structurally biased toward optimism until a milestone is missed outright. With a fixed deadline bearing down, a program that slips six months no longer has six months of slack to absorb it. That is driving new interest in predictive risk scoring: a discipline built on tracking program telemetry, defect rates, resource utilization, change velocity, continuously, rather than waiting for a status report to catch up. Much of the current generation of these tools uses AI-assisted modeling to do that scoring, but the real shift is about catching risk earlier, not about the modeling technique itself.

The Track Record That Makes This Necessary

The scale of the problem is well documented, and not by just one study. Horváth's 2025 research found that 65% of S/4HANA transformations experience significant budget deviations: 25% heavily exceeding budget, 40% strongly exceeding it, with only 35% landing on budget or under. Gartner's broader estimate puts the ERP project failure rate, defined as failing to meet original objectives, at 55–75%. Rand Group, an implementation consultancy, reports that 60% of the new clients who come to them arrive after a failed or subpar implementation with a different provider.

Prosci's 2025 research adds a forward-looking data point that explains why risk scoring has become a priority: Gartner projects that more than 70% of recently implemented ERP initiatives will fail to fully meet their original business case by 2027. That figure isn't retrospective. It's a projection about programs that are, in many cases, still underway right now.

How Budgets Actually Land, Once the Program Is Done



Source: Horváth, 200-company S/4HANA transformation study, 2025.

Budget outcomes across a 200-company sample of S/4HANA transformations.

What Predictive Risk Scoring Actually Does

Predictive risk scoring works by treating the program itself as a data source, rather than relying solely on manually updated RAG (red-amber-green) status reports. The discipline draws on the actual telemetry a migration program generates: defect discovery and closure rates in the test-management system, the ratio of open to resolved fit-gap items, resource utilization against plan, the velocity of change requests, and, where available, patterns from comparable historical projects that ended up over budget or behind schedule. Much of the current tooling applies AI-assisted models to weigh those signals, but the underlying method, evidence in place of self-report, predates any particular modeling approach.

The output isn't a single dashboard color. It's typically a risk score by workstream, updated continuously as new data lands, with the specific combination of signals driving the score made visible rather than buried in a black box. A testing workstream where defect discovery is accelerating faster than defect closure is a good example: that's a leading indicator a program often doesn't surface through manual reporting until the test exit criteria are formally missed, weeks later.

The value here has nothing to do with prediction for its own sake. It comes down to lead time. A risk signal that surfaces six weeks before a status report would have caught the same problem gives a program real options: reallocating resources, adjusting scope, resequencing work. None of that is available once a milestone has already been missed.


65%

of S/4HANA transformations experience significant budget deviations. Most of that risk is visible in program data well before it shows up in a status report. (Horváth)

Why This Matters More Than It Might Seem

A 70%+ projected shortfall against original business case objectives is not a minority of poorly run programs pulling down an otherwise healthy average. That's the expected outcome for a majority of ERP initiatives under current practice. Predictive risk scoring won't change that baseline by itself. What it changes is how early a program learns it's tracking toward that outcome, and that's the one variable a steering committee actually controls.

What Prompted the Interest in Predicting Risk Earlier



of recently implemented
ERP initiatives are
projected to fail to fully
meet their original
business case by 2027

Source: Gartner, cited in Prosci, "Why Do ERP Implementations Fail?", 2025.

The Limits Worth Naming

Risk-scoring models are only as good as the program data feeding them. A program with poor data hygiene in its test-management or project-tracking tools will get an unreliable score, no matter how sophisticated the underlying model is. The models are also backward-looking in their training even when the output looks forward: one trained mainly on brownfield conversions won't necessarily transfer cleanly to a greenfield redesign with a very different risk profile. Used well, predictive risk scoring is an early-warning system that buys a steering committee more time to act. It doesn't, on its own, replace the committee actually acting on what it's told.

Where This Leaves You

Orpington Technologies works with clients on both sides of this problem: delivering S/4HANA migrations built to land on budget and on schedule, and providing independent risk visibility, a view of program health that isn't filtered through the team being measured, for programs already underway. AI-assisted scoring is one part of how our Implementation Risk Remediation engagements work, layered alongside SAP delivery experience and a track record of catching problems before they become expensive. With mainstream ECC support ending December 31, 2027 and no upside to finding out late, earlier visibility and proven delivery are worth the investment now. Learn more about Orpington's approach at orpingtontech.com.

Sources & Further Reading

[1] Horváth — Study: SAP S/4HANA Transformations Rarely Go As Planned (2025)

[2] Rand Group — What Percentage of ERP Implementations Fail?

[3] Prosci — Why Do ERP Implementations Fail? (citing Gartner)