



Master Data Governance After Go-Live: Why Clean Data at Cutover Doesn't Stay Clean

Master data governance isn't a task you finish before go-live and file away. It's an operating discipline that has to run for the life of the S/4HANA system, built into the migration program from the start so the organization can hit the December 31, 2027 deadline without inheriting a second cleanup project eighteen months later, with AI-assisted tooling as one aid along the way.

7 min read | Orpington Technologies Insights

SAP's mainstream maintenance for ECC ends December 31, 2027, and the partner and consultant base built around ECC is already thinning as more of the installed base shifts to S/4HANA. Implementation partners, specialist consultants, and third-party tooling vendors have been moving their attention toward S/4 work for a while now, which makes ECC harder to support well even for organizations still running it. That's the pressure behind most migration timelines today, but the deadline only solves half the problem. Getting clean data into S/4HANA at cutover is table stakes; keeping it clean once the project team disbands is the harder, longer-running job, and it's the one most programs plan for the least.

Pre-migration cleansing gets real attention: dedicated data leads, purpose-built tools, executives who can see exactly how many duplicate records remain. None of that infrastructure survives hypercare on its own. Once the system goes live, master data creation goes back to being one task among many for whichever employee happens to be entering a new vendor or customer record that day, and without a governance function already in place, quality drifts back toward where it started within a year or two. That's why governance has to be designed into the migration program itself, with an owner, a process, and a review cadence defined before go-live rather than added afterward once the drift is already visible. AI-assisted monitoring can help that function scale once it exists, but it isn't a substitute for building the function in the first place.

Why Clean Data at Go-Live Doesn't Stay Clean

A migration project puts an unusual amount of concentrated attention on data quality: dedicated data leads, purpose-built cleansing tools, executive visibility into exactly how many duplicate records remain. Almost none of that infrastructure survives past hypercare. Once the system is live, master data creation goes back to being one task among many for whichever employee happens to be entering a new vendor or customer record that day, and most of them don't have the training or the tooling to catch a near-duplicate before it gets saved.

Basis Technologies' modeling of S/4HANA transformations found that projects need more than 75% additional resourcing beyond baseline estimates just to deliver the migration itself. That figure doesn't include the ongoing governance investment needed to protect that work afterward. Organizations that treat data quality as a project deliverable rather than an operating discipline tend to find themselves, eighteen to twenty-four months later, planning another data cleansing initiative to fix a problem their own migration had already solved once.

The Bigger Governance Job Starts After Go-Live

The governance job that starts after go-live is bigger than the one before it, and it needs the same things any operating discipline needs: a defined process for how a new vendor or customer record gets reviewed, a set point in the workflow where that review happens, and someone accountable for closing out flagged cases. AI-assisted matching earns its place at the point of data entry itself. A duplicate-checking model built into the intake workflow can flag a probable duplicate vendor the moment someone tries to create it, instead of waiting for a quarterly audit to catch it months later.

The harder part is drift: records that were correct at go-live but have since gone inconsistent with related records, or fields trending toward the kind of patterns that preceded known data-quality problems in the legacy system. Catching that takes continuous review built into the governance program, not a one-time check. AI-assisted monitoring can run that continuous pass at a scale a manual review team can't match, but interpreting a flagged pattern and deciding what to do about it still belongs to the governance function, not the model.

75%+

additional resourcing beyond baseline estimates is typically required to deliver an S/4HANA transformation. Governance discipline afterward is what protects that investment from eroding. (Basis Technologies)


The Downstream Payoff Is Larger Than It Looks

Master data quality compounds into processes well beyond the systems it directly touches, which is exactly where its value is easiest to underestimate. SAVIC Technologies' benchmarking of AI-assisted demand forecasting found accuracy improvements of 8–15% in environments with average data readiness, climbing to 18–28% in environments with high data readiness. Same underlying forecasting

models. They just perform meaningfully better because the material and sales history data feeding them is cleaner and more consistent.

The pattern repeats in accounts payable. SAVIC's client deployment data shows clean, well-governed vendor master data enabling a 40–50% reduction in invoice processing cycle time and a 65–80% reduction in the manual touches required per standard invoice. Those gains come from automation being able to trust the vendor and payment-term data it's matching against, instead of routing an exception to a human reviewer every time a record looks slightly off.

Forecast Accuracy Gains Track Directly With Data Readiness



Source: SAVIC Technologies, SAP AI deployment benchmarks, 2026. Figures are midpoints of reported 8-15% and 18-28% ranges.

Where Governance Still Needs an Owner

AI-assisted governance tooling is good at flagging: probable duplicates, drifting records, anomalous entries. Letting it auto-resolve those flags unattended is a much worse idea, especially for merge decisions, which are hard to reverse cleanly once downstream transactions have posted against the surviving record. The governance programs getting the best results keep a named data steward (usually embedded in the business, not IT) as the decision-maker on every flagged case. AI's job there is making sure nothing worth reviewing slips through, not replacing the reviewer.

What Clean Vendor Master Data Enables Downstream



Source: SAVIC Technologies, AP automation client deployment data, 2026. Figures are midpoints of reported 40-50% and 65-80% ranges.

Where This Leaves You

A governance model that survives past hypercare needs a named owner, a clear escalation path, and a process built into the migration program from day one, not added after go-live. Orpington Technologies builds that governance function into every S/4HANA migration it leads ahead of the December 31, 2027 deadline, backed by deep SAP delivery experience, a proven cutover methodology, and AI-assisted tooling among the capabilities it brings to bear. Learn more about Orpington's approach at orpingtontech.com.

Sources & Further Reading

[1] SAVIC Technologies — SAP AI ROI: Real Numbers, Enterprise Reality Check, 2026

[2] Basis Technologies — The True State of S/4HANA 2025