


READ THIS BEFORE YOU AUTOMATE ANYTHING

A practitioner's guide to unlocking value
from AI, RPA and ERP implementations

For leaders in Operations, Finance, IT and Transformation



WRITTEN BY

FAMLA AI

Contents

Contents.....	2
Executive summary.....	4
1. The problem.....	5
The promise.....	5
The reality.....	5
The gap.....	6
2. The evidence: four failure patterns.....	6
Pattern 1: Automating noise.....	6
Pattern 2: Accelerating dysfunction.....	7
Pattern 3: Losing the tribal knowledge.....	8
Pattern 4: Buying tools before defining problems.....	8
3. The fix: map, fix, then automate.....	9
Step 1: Map how work actually happens.....	9
Step 2: Fix what is broken or unnecessary.....	10
Step 3: Automate what remains.....	10
The Trust Ladder.....	10
4. The playbook: six steps you can run this quarter.....	11
Step 1: Audit the pipeline.....	11
Step 2: Prioritise by readiness, not by visibility.....	11
Step 3: Map before you build.....	12
Step 4: Build the business case against P&L, not against hours.....	12
Step 5: Run a pre-automation checklist at go-live.....	12
Step 6: Measure relentlessly.....	12
5. Where Famla fits in your automation stack.....	13
What Famla is.....	13
What Famla does not do.....	13
Why this is a problem worth solving.....	13
Why process mining alone is not enough.....	14
The documented outcome.....	14
6. Sources and further reading.....	15
Analyst research.....	15
Public failure case studies referenced.....	15

Executive summary

Most automation programmes do not fail because the technology is wrong. They fail because the processes underneath the technology were never properly understood.

The data is remarkably consistent across every major analyst. McKinsey finds that only around 16 percent of organisations convert automation activity into sustainable performance improvement. Gartner has moved generative AI into the trough of disillusionment, predicts that up to 60 percent of AI projects will be abandoned where the underlying data is not AI-ready, and expects at least 30 percent of generative AI projects to be abandoned after proof of concept. Deloitte's Global Intelligent Automation Survey puts the share of organisations that have scaled their digital workforce at around 3 percent. RAND puts the overall AI project failure rate at roughly 80 percent. MIT Sloan's 2025 research found that around 95 percent of generative AI pilots never reach production. The money is being spent. The value is not showing up.

The root cause is process, not platform. Organisations automate what they think they do, not what they actually do. They speed up broken handoffs. They lose the tribal knowledge that lives in their people. They mistake hours saved for margin gained. They buy tools before they understand the work.

The fix is a sequence, not a new vendor: map how work actually happens, fix what is broken or unnecessary, then automate what remains. This whitepaper sets out the evidence, the sequence, and a practical playbook you can start using this quarter.

1. The problem

The promise

The arithmetic of automation is seductive and has not changed in twenty years. A tool that saves 100 employees five minutes a day returns roughly 2,075 hours to the business each year. Price those hours at fifty dollars and the tool has saved around 103,750 dollars. On paper.

Stack that case across ten processes, then across ten teams, and the numbers become irresistible. This is the business case that gets RPA programmes funded, ERP projects approved, and AI pilots greenlit. It is the case that McKinsey, BCG, and every systems integrator in the world has been selling for a decade. The promise is real. Automation can return time, improve accuracy, and scale capacity.

The problem is not the promise; it is what happens between the slide deck and the P&L.

The reality

Only about 16 percent of organisations achieve sustainable performance improvement from automation, according to McKinsey's tracking across the intelligent automation wave. The headline should stop any sponsor in their tracks. Roughly five out of six automation programmes do not leave a durable footprint on the numbers.

The newer AI data is worse. MIT's 2025 research found that around 95 percent of organisations deploying generative AI saw no measurable return. RAND puts the AI project failure rate at about twice that of non-AI IT projects. Gartner has shifted generative AI into the trough of disillusionment in its latest Hype Cycle, with supply chain pilot-to-production failure rates reportedly above 70 percent. S&P Global found that 42 percent of companies abandoned at least one AI initiative in 2025, up from 17 percent the year before.

ERP, the most expensive of the three categories, produces the most visible failures. Revlon's 2018 SAP S/4HANA rollout disrupted manufacturing at its Oxford facility, generated a reported \$70.3 million quarterly net loss, and triggered shareholder lawsuits. National Grid's SAP implementation, going live days after Hurricane Sandy, ended with roughly \$585 million in clean-up costs and a lawsuit against the integrator. Lidl walked away from an SAP programme after around \$580 million and seven years. MillerCoors sued its integrator for \$100 million. Hertz paid Accenture \$32 million for a website that never went live, and then spent another \$10 million to fix what was delivered.

The gap

None of these failures was caused by a defective tool. RPA works. S/4HANA works. Large language models work. **The gap sits upstream of the technology, in the processes the tools are pointed at.**

The uncomfortable question

You saved time. Great. But did it hit your bottom line? Did it show up on the P&L, or does it only look good in a presentation deck?

When 16 percent is the headline number for sustainable impact, the interesting question is not whether automation can work. The interesting question is what the winning 16 percent do differently. The next section is about that.

2. The evidence: four failure patterns

Across the public failure cases and the analyst research, the same four patterns show up. They are not technology problems. **They are process problems wearing technology clothes.**

Pattern 1: Automating noise

The first failure mode is the simplest. Organisations automate activities without asking whether those activities should exist at all.

A finance team automates the creation of a weekly report. The bot runs flawlessly. Eighteen months later, nobody is reading the report. The team never asked who used the output, or what decision it informed. The automation did exactly what it was designed to do, which was to speed up work that did not need to happen in the first place. Gartner's research lists unclear business value and poor use-case selection at the top of the failure list.

For Finance: What this means for you

- Hours saved translate to margin only if the released capacity is redeployed, headcount is reduced, or volume grows without incremental cost. If none happens, the saved time is slack, not savings.
- Ask for the named P&L line the automation will move, and the owner of that commitment, before the business case is approved.

Pattern 2: Accelerating dysfunction

The second failure mode is more dangerous because the numbers look good for a while. When you automate a broken process, you do not fix the break. You industrialise it.

A loan approval workflow that routed work to the wrong team 20 percent of the time now does so at machine speed. An invoicing handoff that dropped a field in every third case now drops fields thousands of times a day. Revlon's ERP failure is this pattern at scale: consolidation on S/4HANA without a properly validated view of how the business actually worked, locking system mismatches into manufacturing.

BCG's 10-20-70 rule, drawn from hundreds of AI transformations, captures the point. Roughly 10 percent of the effort in a successful AI programme is algorithms. Roughly 20 percent is technology and data. The remaining 70 percent is people and processes. Most organisations invert the ratio and get inverted results.

For IT: What this means for you

- When automation fails, the business blames IT because IT delivered the tool. The root cause is usually upstream: the process was never properly specified.
- Insist on process discovery that captures both the system view and the human view before build begins. This is the single most effective defence against stabilisation pain.

Pattern 3: Losing the tribal knowledge

The third failure mode is the one that surprises executives most, because it is invisible until go-live.

The processes that keep a business running are not the processes that are written down. They live in people's heads: the senior claims adjuster who knows when to escalate, the warehouse supervisor who reads the Monday rota against the forecast, the accounts payable clerk who recognises a vendor anomaly because she has seen it before. This is tribal knowledge, and it is often the difference between the happy path in a process map and what actually gets customers served.

Traditional process mapping workshops fail to capture this knowledge. The senior people are not in the room, the junior people are nervous, and the consultant at the whiteboard is asking for the officially-documented version. Process mining tools downstream then work on event logs that reflect systems, not the human overlays. When the automation arrives, the tribal knowledge does not come with it.

For Operations: What this means for you

- You sit closest to the tribal knowledge and pay the price when it is lost at go-live. You are also best placed to surface it.
- Prioritise processes where the variance between stated practice and actual practice is highest. The bigger the gap, the more the automation will hurt if you do not close it first.
- Treat SOPs as hypotheses to be validated with the people doing the work, not as ground truth. Most SOPs in large organisations are materially out of date within 18 months.

Pattern 4: Buying tools before defining problems

The fourth failure mode is the most expensive because it tends to be sponsored at the top. A board decides the company needs an AI strategy, or an RPA programme, or an ERP upgrade, and picks a vendor before the problem has been defined.

This is how Hertz ended up in a \$42 million dispute over a website that never shipped. It is how Lidl ran an SAP programme for seven years before a new CEO walked away. It is how Birmingham City Council in the UK ended up with an Oracle ERP project whose cost overruns

reportedly ran into the tens of millions of pounds, leaving the council without a functional financial management system for over two years.

Gartner's list of failure drivers for generative AI reads like a pattern check: unclear business value, inadequate risk controls, escalating costs, lack of AI-ready data, and missing executive sponsorship tied to operational outcomes. The through-line is that the tool was chosen before the work was understood.

For Transformation: What this means for you

- Protect the discovery phase. Pressure to start building is constant. The cost of building the wrong thing dwarfs the cost of a few extra weeks of discovery.
- Build your pilot portfolio around processes that will teach the organisation, not just processes that will generate easy wins. One honest hard pilot beats three safe ones.

3. The fix: map, fix, then automate

The pattern that separates the successful 16 percent from the rest is a sequence, not a tool. It has three steps, in order, and skipping any of them is the origin of most failed programmes.

Step 1: Map how work actually happens

Before the tool decision, before the vendor selection, before the use-case workshop, the organisation needs a verified view of how the process actually runs. Not the official version. The real one.

Classical process mining, using tools such as Celonis and SAP Signavio Process Intelligence, is excellent for what it can see: the event logs in the ERP, the CRM, the ticketing system. It is invaluable for bottleneck analysis, conformance checking, and benchmarking. The complement to process mining is process discovery through the people who do the work: stakeholder interviews, gemba walks, and in-app observation that capture the unofficial routing rules, the exceptions handled by hand, the Slack messages that move a case forward, the spreadsheet that bridges two systems.

Process mining and process discovery are a complementary pair

Process mining reads event logs. Process discovery reads people. Neither is sufficient alone. The organisations winning at automation use both, in that order.

Step 2: Fix what is broken or unnecessary

Once the real process is visible, it almost always contains three categories of work. Activity that should exist and is done well. Activity that should exist but is done badly. Activity that should not exist at all. The first category is the automation target. The second is the redesign target. The third is the elimination target.

Skipping the distinction is the classic failure. Automating the third category is noise. Automating the second locks in dysfunction. Only the first category genuinely benefits from automation, and it is usually smaller than the initial business case assumed. The Lean Six Sigma discipline is built around this distinction, which is why the most successful transformation programmes still start with a value stream map and a waste analysis before anyone opens a vendor RFP.

Step 3: Automate what remains

Only at step three does the tool question become the right question. And at step three, it gets easier, because the process is well understood, the benefit case is grounded in actual work, and the stakeholders are aligned on what the automation should do.

The Trust Ladder

For a given process that has survived steps one and two, think of automation as a ladder. Each rung requires more trust in the underlying process and delivers more leverage.

- **Suggest:** AI surfaces insights and recommendations. A human decides and acts. Low risk, low leverage, high learning.
- **Assist:** AI drafts work, fills forms, or pre-populates decisions. A human reviews and approves. Moderate leverage, requires moderate process maturity.
- **Run:** RPA or agentic AI executes end-to-end for defined scenarios. A human handles exceptions. High leverage, requires high process maturity and clear exception routing.
- **Optimise:** AI executes and continuously tunes the process based on outcomes. Humans set objectives and constraints. Highest leverage, requires mature telemetry and governance.

Most organisations try to jump to Run or Optimise on processes that have not earned Suggest. This is the single most common source of the accelerating-dysfunction pattern.

4. The playbook: six steps you can run this quarter

The rest of this section is practical. This is the sequence you can run starting Monday morning, regardless of where your automation programme currently sits.

Step 1: Audit the pipeline

List every active and planned automation initiative. For each, answer four questions: what is the named business outcome; who owns the redeployed capacity; what is the current process maturity; what does go-live look like. Anything that fails on one of the four is a candidate for pause, not acceleration.

Step 2: Prioritise by readiness, not by visibility

Score each candidate against five criteria. Processes that score well on four of the five are automation-ready. Processes that score well on two or fewer are redesign candidates first.

Criterion	Good sign	Warning sign
Process stability	No major changes in the last 12 months; limited change planned.	Being redesigned, or expected to change materially within the year.
Data quality	Inputs structured, consistent, and complete; audit trail exists.	Heavy reliance on unstructured inputs, email attachments, or free text.
Volume	High enough that automation has a measurable denominator.	Sub-scale volume where manual execution remains cheaper.
Exception rate	Exceptions are documented and routed to a named handler.	Exceptions are frequent, undocumented, and handled through tribal knowledge.
Stakeholder alignment	Business owner, IT, and frontline agree on what the process is and should become.	No single owner, competing definitions, or a frontline that was not consulted.

Step 3: Map before you build

For every process that makes the cut, commission a proper process discovery exercise. This is not a two-hour whiteboard session with a consultant. It is structured engagement with the people who actually do the work, validated against system event data, producing a process map that is signed off by the frontline as well as the manager. The length of this phase is the single biggest lever on success downstream.

Step 4: Build the business case against P&L, not against hours

Require every business case to name the P&L line that will move, the magnitude of the move, and the owner of the commitment. Hours saved is an intermediate metric, not a final one. If a business case cannot be tied to a P&L line, it is not ready for investment.

Step 5: Run a pre-automation checklist at go-live

- Has the process been mapped end-to-end, including exceptions, and signed off by the people who do the work?
- Have the handoffs been tested with the downstream team under realistic, not happy-path, conditions?
- Is there a named business owner for each KPI the automation is expected to move?
- Is the rollback plan tested, and is change communication ready for customers as well as staff?

Step 6: Measure relentlessly

The single biggest predictor of whether an automation delivers value is whether someone is measuring it. Set a quarterly benefit realisation review. Compare actual to plan. Shut down what is not working. Double down on what is. Absent this discipline, automation portfolios quietly drift toward the 84 percent that does not deliver sustainable impact.

Vendor red flags

- The implementation plan compresses discovery to a few days or skips it.
- Training and change management are presented as optional.
- The contract includes technical acceptance criteria only, not business outcomes.
- The vendor cannot name specific customers in your industry with measurable outcomes at your scale.

5. Where Famla fits in your automation stack

Most of this whitepaper has been vendor-neutral, because the principles are vendor-neutral. This section is where we tell you what we do and how we fit into the sequence above. We have kept it short on purpose.

What Famla is

Famla is an AI-powered process discovery and mapping platform. We sit at step one of the Map, Fix, Automate sequence. Our AI conducts structured, autonomous stakeholder interviews in multiple languages, ingests the documents and whiteboard images your team already has, and synthesises all of it into verified swimlane process maps, SOPs, and Lean Six Sigma analysis, automatically.

Every element in a Famla map links back to the original interview response or document note that produced it. Nothing is invented. Every claim is verifiable against a real stakeholder input.

What Famla does not do

Famla is not a process mining tool. It does not replace Celonis, SAP Signavio Process Intelligence, ABBYY, KYP.ai, or UiPath Process Mining. Famla is not an RPA platform. It does not replace UiPath, Blue Prism, or Automation Anywhere. Famla is not an ERP. It does not replace SAP, Oracle, or Workday. Famla is not a workflow tool. It does not replace Asana, Monday.com, ServiceNow, or Jira. Famla is not a knowledge base. It does not replace Confluence, SharePoint, or Notion.

Famla is complementary to every one of these. It is the discovery layer that makes all of them work better.

Why this is a problem worth solving

Every process improvement, automation, or digital transformation initiative requires someone to first understand how the process actually works. Before pointing Celonis at a database, someone needs to know which tables are relevant and what the process is supposed to look like. Before automating a workflow with UiPath, someone needs to know what the workflow actually is, including the exceptions and workarounds that never appear in documentation. Before configuring Asana, someone needs to know how work flows so the tool reflects reality rather than assumptions.

This upstream work is consistently the most underinvested and most time-consuming phase of any process initiative. Famla automates it.

Why process mining alone is not enough

Process mining tools like Celonis, ABBYY, KYP.ai, and UiPath Process Mining are powerful, and they share four structural constraints that limit where they can be deployed.

- **They need human guidance before they can start.** Data without context is noise. Someone has to define which tables, which boundaries, and which happy path. Famla produces that baseline in days, without system access or data engineering.
- **They need system access, which takes time to negotiate.** Technical integration and security approvals can kill momentum before a diagnostic even starts. Famla deploys in under five minutes with no connectors, no data extraction, and no IT approvals.
- **They only see what systems record.** Manual tasks, offline handovers, and cross-company flows leave no digital trace. Famla captures what people know, not just what systems record.
- **They talk to systems. Famla talks to people.** Involving frontline teams in discovery builds shared understanding early, surfaces better improvement ideas, and reduces the resistance that top-down data presentations generate.

When Famla precedes a process mining engagement, scope is clearer, data requirements are better defined, and the findings land with the operational context needed to turn them into action, not just dashboards.

The documented outcome

Two benchmarks from our customer base quantify the argument above. Transformation teams using Famla save an average of more than 200 minutes per process mapped compared with manual discovery methods, and they hear roughly 20 percent more stakeholder voices per process than they would in a traditional workshop-based approach. Operations are understood roughly 10 times faster.

The second number matters more than the first. The hours saved are a side benefit. The stakeholder voices heard are the reason the downstream automation does not fail. Tribal knowledge captured in discovery is tribal knowledge that does not get lost at go-live.

Start at famla.com

Anyone can create a free account instantly, with no demo required and no credit card. New users receive free credits with no time restriction. Famla bills on credit volume, not seats, so there is no penalty for involving your whole team. The companies winning at automation are the ones that refuse to automate on top of processes they do not understand. That is the problem Famla was built to solve.

6. Sources and further reading

The data and case studies in this whitepaper draw on publicly available analyst research, academic studies, legal filings, and press coverage. The headline references are listed below.

Analyst research

McKinsey Global Institute and McKinsey Digital on automation productivity, intelligent automation adoption, and the economic potential of generative AI. Deloitte Global Intelligent Automation Survey on scaling patterns, payback periods, and enterprise adoption. Gartner on generative AI abandonment, AI-ready data, agentic AI project failure rates, and the Hype Cycle for supply chain strategy. BCG on the 10-20-70 rule and scaling agentic AI through end-to-end process redesign. RAND Corporation, 2024, on AI project failure rates and root causes. MIT Sloan and Project NANDA, 2025, on pilot-to-production gaps. S&P Global Market Intelligence, 2025, on AI initiative abandonment rates.

Public failure case studies referenced

Revlon SAP S/4HANA (2018 onward), based on SEC filings and Panorama Consulting analyses. National Grid SAP ERP (2012 onward), including the Wipro lawsuit coverage. MillerCoors and HCL Technologies, court filings. Lidl and SAP, public reporting of the seven-year cancellation. Hertz and Accenture, US Southern District of New York filings. Birmingham City Council Oracle ERP, Grant Thornton audit of February 2025. For deeper reading, visit famla.com/blog.

Contents

Contents.....	2
Executive summary.....	4
1. The problem.....	5
The promise.....	5
The reality.....	5
The gap.....	6
2. The evidence: four failure patterns.....	6
Pattern 1: Automating noise.....	6
Pattern 2: Accelerating dysfunction.....	7
Pattern 3: Losing the tribal knowledge.....	8
Pattern 4: Buying tools before defining problems.....	8
3. The fix: map, fix, then automate.....	9
Step 1: Map how work actually happens.....	9
Step 2: Fix what is broken or unnecessary.....	10
Step 3: Automate what remains.....	10
The Trust Ladder.....	10
4. The playbook: six steps you can run this quarter.....	11
Step 1: Audit the pipeline.....	11
Step 2: Prioritise by readiness, not by visibility.....	11
Step 3: Map before you build.....	12
Step 4: Build the business case against P&L, not against hours.....	12
Step 5: Run a pre-automation checklist at go-live.....	12
Step 6: Measure relentlessly.....	12
5. Where Famla fits in your automation stack.....	13
What Famla is.....	13
What Famla does not do.....	13
Why this is a problem worth solving.....	13
Why process mining alone is not enough.....	14
The documented outcome.....	14
6. Sources and further reading.....	15
Analyst research.....	15
Public failure case studies referenced.....	15

Executive summary

Most automation programmes do not fail because the technology is wrong. They fail because the processes underneath the technology were never properly understood.

The data is remarkably consistent across every major analyst. McKinsey finds that only around 16 percent of organisations convert automation activity into sustainable performance improvement. Gartner has moved generative AI into the trough of disillusionment, predicts that up to 60 percent of AI projects will be abandoned where the underlying data is not AI-ready, and expects at least 30 percent of generative AI projects to be abandoned after proof of concept. Deloitte's Global Intelligent Automation Survey puts the share of organisations that have scaled their digital workforce at around 3 percent. RAND puts the overall AI project failure rate at roughly 80 percent. MIT Sloan's 2025 research found that around 95 percent of generative AI pilots never reach production. The money is being spent. The value is not showing up.

The root cause is process, not platform. Organisations automate what they think they do, not what they actually do. They speed up broken handoffs. They lose the tribal knowledge that lives in their people. They mistake hours saved for margin gained. They buy tools before they understand the work.

The fix is a sequence, not a new vendor: map how work actually happens, fix what is broken or unnecessary, then automate what remains. This whitepaper sets out the evidence, the sequence, and a practical playbook you can start using this quarter.

1. The problem

The promise

The arithmetic of automation is seductive and has not changed in twenty years. A tool that saves 100 employees five minutes a day returns roughly 2,075 hours to the business each year. Price those hours at fifty dollars and the tool has saved around 103,750 dollars. On paper.

Stack that case across ten processes, then across ten teams, and the numbers become irresistible. This is the business case that gets RPA programmes funded, ERP projects approved, and AI pilots greenlit. It is the case that McKinsey, BCG, and every systems integrator in the world has been selling for a decade. The promise is real. Automation can return time, improve accuracy, and scale capacity.

The problem is not the promise; it is what happens between the slide deck and the P&L.

The reality

Only about 16 percent of organisations achieve sustainable performance improvement from automation, according to McKinsey's tracking across the intelligent automation wave. The headline should stop any sponsor in their tracks. Roughly five out of six automation programmes do not leave a durable footprint on the numbers.

The newer AI data is worse. MIT's 2025 research found that around 95 percent of organisations deploying generative AI saw no measurable return. RAND puts the AI project failure rate at about twice that of non-AI IT projects. Gartner has shifted generative AI into the trough of disillusionment in its latest Hype Cycle, with supply chain pilot-to-production failure rates reportedly above 70 percent. S&P Global found that 42 percent of companies abandoned at least one AI initiative in 2025, up from 17 percent the year before.

ERP, the most expensive of the three categories, produces the most visible failures. Revlon's 2018 SAP S/4HANA rollout disrupted manufacturing at its Oxford facility, generated a reported \$70.3 million quarterly net loss, and triggered shareholder lawsuits. National Grid's SAP implementation, going live days after Hurricane Sandy, ended with roughly \$585 million in clean-up costs and a lawsuit against the integrator. Lidl walked away from an SAP programme after around \$580 million and seven years. MillerCoors sued its integrator for \$100 million. Hertz paid Accenture \$32 million for a website that never went live, and then spent another \$10 million to fix what was delivered.

The gap

None of these failures was caused by a defective tool. RPA works. S/4HANA works. Large language models work. **The gap sits upstream of the technology, in the processes the tools are pointed at.**

The uncomfortable question

You saved time. Great. But did it hit your bottom line? Did it show up on the P&L, or does it only look good in a presentation deck?

When 16 percent is the headline number for sustainable impact, the interesting question is not whether automation can work. The interesting question is what the winning 16 percent do differently. The next section is about that.

2. The evidence: four failure patterns

Across the public failure cases and the analyst research, the same four patterns show up. They are not technology problems. **They are process problems wearing technology clothes.**

Pattern 1: Automating noise

The first failure mode is the simplest. Organisations automate activities without asking whether those activities should exist at all.

A finance team automates the creation of a weekly report. The bot runs flawlessly. Eighteen months later, nobody is reading the report. The team never asked who used the output, or what decision it informed. The automation did exactly what it was designed to do, which was to speed up work that did not need to happen in the first place. Gartner's research lists unclear business value and poor use-case selection at the top of the failure list.

For Finance: What this means for you

- Hours saved translate to margin only if the released capacity is redeployed, headcount is reduced, or volume grows without incremental cost. If none happens, the saved time is slack, not savings.
- Ask for the named P&L line the automation will move, and the owner of that commitment, before the business case is approved.

Pattern 2: Accelerating dysfunction

The second failure mode is more dangerous because the numbers look good for a while. When you automate a broken process, you do not fix the break. You industrialise it.

A loan approval workflow that routed work to the wrong team 20 percent of the time now does so at machine speed. An invoicing handoff that dropped a field in every third case now drops fields thousands of times a day. Revlon's ERP failure is this pattern at scale: consolidation on S/4HANA without a properly validated view of how the business actually worked, locking system mismatches into manufacturing.

BCG's 10-20-70 rule, drawn from hundreds of AI transformations, captures the point. Roughly 10 percent of the effort in a successful AI programme is algorithms. Roughly 20 percent is technology and data. The remaining 70 percent is people and processes. Most organisations invert the ratio and get inverted results.

For IT: What this means for you

- When automation fails, the business blames IT because IT delivered the tool. The root cause is usually upstream: the process was never properly specified.
- Insist on process discovery that captures both the system view and the human view before build begins. This is the single most effective defence against stabilisation pain.

Pattern 3: Losing the tribal knowledge

The third failure mode is the one that surprises executives most, because it is invisible until go-live.

The processes that keep a business running are not the processes that are written down. They live in people's heads: the senior claims adjuster who knows when to escalate, the warehouse supervisor who reads the Monday rota against the forecast, the accounts payable clerk who recognises a vendor anomaly because she has seen it before. This is tribal knowledge, and it is often the difference between the happy path in a process map and what actually gets customers served.

Traditional process mapping workshops fail to capture this knowledge. The senior people are not in the room, the junior people are nervous, and the consultant at the whiteboard is asking for the officially-documented version. Process mining tools downstream then work on event logs that reflect systems, not the human overlays. When the automation arrives, the tribal knowledge does not come with it.

For Operations: What this means for you

- You sit closest to the tribal knowledge and pay the price when it is lost at go-live. You are also best placed to surface it.
- Prioritise processes where the variance between stated practice and actual practice is highest. The bigger the gap, the more the automation will hurt if you do not close it first.
- Treat SOPs as hypotheses to be validated with the people doing the work, not as ground truth. Most SOPs in large organisations are materially out of date within 18 months.

Pattern 4: Buying tools before defining problems

The fourth failure mode is the most expensive because it tends to be sponsored at the top. A board decides the company needs an AI strategy, or an RPA programme, or an ERP upgrade, and picks a vendor before the problem has been defined.

This is how Hertz ended up in a \$42 million dispute over a website that never shipped. It is how Lidl ran an SAP programme for seven years before a new CEO walked away. It is how Birmingham City Council in the UK ended up with an Oracle ERP project whose cost overruns

reportedly ran into the tens of millions of pounds, leaving the council without a functional financial management system for over two years.

Gartner's list of failure drivers for generative AI reads like a pattern check: unclear business value, inadequate risk controls, escalating costs, lack of AI-ready data, and missing executive sponsorship tied to operational outcomes. The through-line is that the tool was chosen before the work was understood.

For Transformation: What this means for you

- Protect the discovery phase. Pressure to start building is constant. The cost of building the wrong thing dwarfs the cost of a few extra weeks of discovery.
- Build your pilot portfolio around processes that will teach the organisation, not just processes that will generate easy wins. One honest hard pilot beats three safe ones.

3. The fix: map, fix, then automate

The pattern that separates the successful 16 percent from the rest is a sequence, not a tool. It has three steps, in order, and skipping any of them is the origin of most failed programmes.

Step 1: Map how work actually happens

Before the tool decision, before the vendor selection, before the use-case workshop, the organisation needs a verified view of how the process actually runs. Not the official version. The real one.

Classical process mining, using tools such as Celonis and SAP Signavio Process Intelligence, is excellent for what it can see: the event logs in the ERP, the CRM, the ticketing system. It is invaluable for bottleneck analysis, conformance checking, and benchmarking. The complement to process mining is process discovery through the people who do the work: stakeholder interviews, gamba walks, and in-app observation that capture the unofficial routing rules, the exceptions handled by hand, the Slack messages that move a case forward, the spreadsheet that bridges two systems.

Process mining and process discovery are a complementary pair

Process mining reads event logs. Process discovery reads people. Neither is sufficient alone. The organisations winning at automation use both, in that order.

Step 2: Fix what is broken or unnecessary

Once the real process is visible, it almost always contains three categories of work. Activity that should exist and is done well. Activity that should exist but is done badly. Activity that should not exist at all. The first category is the automation target. The second is the redesign target. The third is the elimination target.

Skipping the distinction is the classic failure. Automating the third category is noise. Automating the second locks in dysfunction. Only the first category genuinely benefits from automation, and it is usually smaller than the initial business case assumed. The Lean Six Sigma discipline is built around this distinction, which is why the most successful transformation programmes still start with a value stream map and a waste analysis before anyone opens a vendor RFP.

Step 3: Automate what remains

Only at step three does the tool question become the right question. And at step three, it gets easier, because the process is well understood, the benefit case is grounded in actual work, and the stakeholders are aligned on what the automation should do.

The Trust Ladder

For a given process that has survived steps one and two, think of automation as a ladder. Each rung requires more trust in the underlying process and delivers more leverage.

- **Suggest:** AI surfaces insights and recommendations. A human decides and acts. Low risk, low leverage, high learning.
- **Assist:** AI drafts work, fills forms, or pre-populates decisions. A human reviews and approves. Moderate leverage, requires moderate process maturity.
- **Run:** RPA or agentic AI executes end-to-end for defined scenarios. A human handles exceptions. High leverage, requires high process maturity and clear exception routing.
- **Optimise:** AI executes and continuously tunes the process based on outcomes. Humans set objectives and constraints. Highest leverage, requires mature telemetry and governance.

Most organisations try to jump to Run or Optimise on processes that have not earned Suggest. This is the single most common source of the accelerating-dysfunction pattern.

4. The playbook: six steps you can run this quarter

The rest of this section is practical. This is the sequence you can run starting Monday morning, regardless of where your automation programme currently sits.

Step 1: Audit the pipeline

List every active and planned automation initiative. For each, answer four questions: what is the named business outcome; who owns the redeployed capacity; what is the current process maturity; what does go-live look like. Anything that fails on one of the four is a candidate for pause, not acceleration.

Step 2: Prioritise by readiness, not by visibility

Score each candidate against five criteria. Processes that score well on four of the five are automation-ready. Processes that score well on two or fewer are redesign candidates first.

Criterion	Good sign	Warning sign
Process stability	No major changes in the last 12 months; limited change planned.	Being redesigned, or expected to change materially within the year.
Data quality	Inputs structured, consistent, and complete; audit trail exists.	Heavy reliance on unstructured inputs, email attachments, or free text.
Volume	High enough that automation has a measurable denominator.	Sub-scale volume where manual execution remains cheaper.
Exception rate	Exceptions are documented and routed to a named handler.	Exceptions are frequent, undocumented, and handled through tribal knowledge.
Stakeholder alignment	Business owner, IT, and frontline agree on what the process is and should become.	No single owner, competing definitions, or a frontline that was not consulted.

Step 3: Map before you build

For every process that makes the cut, commission a proper process discovery exercise. This is not a two-hour whiteboard session with a consultant. It is structured engagement with the people who actually do the work, validated against system event data, producing a process map that is signed off by the frontline as well as the manager. The length of this phase is the single biggest lever on success downstream.

Step 4: Build the business case against P&L, not against hours

Require every business case to name the P&L line that will move, the magnitude of the move, and the owner of the commitment. Hours saved is an intermediate metric, not a final one. If a business case cannot be tied to a P&L line, it is not ready for investment.

Step 5: Run a pre-automation checklist at go-live

- Has the process been mapped end-to-end, including exceptions, and signed off by the people who do the work?
- Have the handoffs been tested with the downstream team under realistic, not happy-path, conditions?
- Is there a named business owner for each KPI the automation is expected to move?
- Is the rollback plan tested, and is change communication ready for customers as well as staff?

Step 6: Measure relentlessly

The single biggest predictor of whether an automation delivers value is whether someone is measuring it. Set a quarterly benefit realisation review. Compare actual to plan. Shut down what is not working. Double down on what is. Absent this discipline, automation portfolios quietly drift toward the 84 percent that does not deliver sustainable impact.

Vendor red flags

- The implementation plan compresses discovery to a few days or skips it.
- Training and change management are presented as optional.
- The contract includes technical acceptance criteria only, not business outcomes.
- The vendor cannot name specific customers in your industry with measurable outcomes at your scale.

5. Where Famla fits in your automation stack

Most of this whitepaper has been vendor-neutral, because the principles are vendor-neutral. This section is where we tell you what we do and how we fit into the sequence above. We have kept it short on purpose.

What Famla is

Famla is an AI-powered process discovery and mapping platform. We sit at step one of the Map, Fix, Automate sequence. Our AI conducts structured, autonomous stakeholder interviews in multiple languages, ingests the documents and whiteboard images your team already has, and synthesises all of it into verified swimlane process maps, SOPs, and Lean Six Sigma analysis, automatically.

Every element in a Famla map links back to the original interview response or document note that produced it. Nothing is invented. Every claim is verifiable against a real stakeholder input.

What Famla does not do

Famla is not a process mining tool. It does not replace Celonis, SAP Signavio Process Intelligence, ABBYY, KYP.ai, or UiPath Process Mining. Famla is not an RPA platform. It does not replace UiPath, Blue Prism, or Automation Anywhere. Famla is not an ERP. It does not replace SAP, Oracle, or Workday. Famla is not a workflow tool. It does not replace Asana, Monday.com, ServiceNow, or Jira. Famla is not a knowledge base. It does not replace Confluence, SharePoint, or Notion.

Famla is complementary to every one of these. It is the discovery layer that makes all of them work better.

Why this is a problem worth solving

Every process improvement, automation, or digital transformation initiative requires someone to first understand how the process actually works. Before pointing Celonis at a database, someone needs to know which tables are relevant and what the process is supposed to look like. Before automating a workflow with UiPath, someone needs to know what the workflow actually is, including the exceptions and workarounds that never appear in documentation. Before configuring Asana, someone needs to know how work flows so the tool reflects reality rather than assumptions.

This upstream work is consistently the most underinvested and most time-consuming phase of any process initiative. Famla automates it.

Why process mining alone is not enough

Process mining tools like Celonis, ABBYY, KYP.ai, and UiPath Process Mining are powerful, and they share four structural constraints that limit where they can be deployed.

- **They need human guidance before they can start.** Data without context is noise. Someone has to define which tables, which boundaries, and which happy path. Famla produces that baseline in days, without system access or data engineering.
- **They need system access, which takes time to negotiate.** Technical integration and security approvals can kill momentum before a diagnostic even starts. Famla deploys in under five minutes with no connectors, no data extraction, and no IT approvals.
- **They only see what systems record.** Manual tasks, offline handovers, and cross-company flows leave no digital trace. Famla captures what people know, not just what systems record.
- **They talk to systems. Famla talks to people.** Involving frontline teams in discovery builds shared understanding early, surfaces better improvement ideas, and reduces the resistance that top-down data presentations generate.

When Famla precedes a process mining engagement, scope is clearer, data requirements are better defined, and the findings land with the operational context needed to turn them into action, not just dashboards.

The documented outcome

Two benchmarks from our customer base quantify the argument above. Transformation teams using Famla save an average of more than 200 minutes per process mapped compared with manual discovery methods, and they hear roughly 20 percent more stakeholder voices per process than they would in a traditional workshop-based approach. Operations are understood roughly 10 times faster.

The second number matters more than the first. The hours saved are a side benefit. The stakeholder voices heard are the reason the downstream automation does not fail. Tribal knowledge captured in discovery is tribal knowledge that does not get lost at go-live.

Start at famla.com

- Anyone can create a free account instantly, with no demo required and no credit card.
- **New users receive free credits with no time restriction.**
- Famla bills on credit volume, not seats, so there is no penalty for involving your whole team. In fact, it is highly encouraged to do so.
- The companies winning at automation are the ones that refuse to automate on top of processes they do not understand. That is the problem Famla was built to solve.

6. Sources and further reading

The data and case studies in this whitepaper draw on publicly available analyst research, academic studies, legal filings, and press coverage. The headline references are listed below.

Analyst research

McKinsey Global Institute and McKinsey Digital on automation productivity, intelligent automation adoption, and the economic potential of generative AI. Deloitte Global Intelligent Automation Survey on scaling patterns, payback periods, and enterprise adoption. Gartner on generative AI abandonment, AI-ready data, agentic AI project failure rates, and the Hype Cycle for supply chain strategy. BCG on the 10-20-70 rule and scaling agentic AI through end-to-end process redesign. RAND Corporation, 2024, on AI project failure rates and root causes. MIT Sloan and Project NANDA, 2025, on pilot-to-production gaps. S&P Global Market Intelligence, 2025, on AI initiative abandonment rates.

Public failure case studies referenced

Revlon SAP S/4HANA (2018 onward), based on SEC filings and Panorama Consulting analyses. National Grid SAP ERP (2012 onward), including the Wipro lawsuit coverage. MillerCoors and HCL Technologies, court filings. Lidl and SAP, public reporting of the seven-year cancellation. Hertz and Accenture, US Southern District of New York filings. Birmingham City Council Oracle ERP, Grant Thornton audit of February 2025.

For deeper reading,

Visit famla.com/blog