

Internal Control and Artificial Intelligence

From an army of preparers to a single control ecosystem

«A control that a machine can perform, a machine should perform. What is left for people is the part that genuinely requires judgment, and that part was never the bottleneck.» — Reto H. Wenger

Prepared for Boards of Directors, Audit Committees and Executive Management

GRC: Governance, Risk, Compliance – no frills, just results.

ICS & SOX | Internal Audit | Risk Management | Compliance

Version 1.0 | September 2026 | © LakeHill GRC LLC. All Rights Reserved.

About LakeHill GRC

GRC – Governance, Risk, Compliance. No frills, just results.

Founded in 2024 by Reto H. Wenger on a Hill overlooking a Lake. Bringing two decades of expertise, we deliver GRC solutions focused on efficiency and results. Cutting through the noise straight to the root cause: that is who we are.

«Strong Governance is only as resilient as the organizational, process and IT foundations it rests upon – without a solid structure, even the best frameworks risk collapse.» — Reto H. Wenger

This paper is written for people who sit on the receiving end of an Internal Control System: Boards, Audit Committees and the Executives who pay for it. It is deliberately concrete. Where we give numbers, we say where they come from and what they are worth. Where we see risk, we say so, because a paper on this subject that contains only upside is marketing, not advice.

1. Introduction: the uncomfortable starting point

Most Internal Control Systems are not expensive because they are effective. They are expensive because they are manual.

Walk into any larger organization in the second week of January and count the people working on the Internal Control System. A control preparer pulling screenshots out of the ERP. A control performer signing a checklist. A central ICS function chasing the ones who have not signed. A second-line reviewer sampling what the first line did. Internal Audit re-testing a sample of what the second line already reviewed. External Audit re-testing a sample of what Internal Audit tested. And, somewhere in the last week of the month, an analyst assembling the summary slides for the Audit Committee.

Four groups of people, in sequence, looking at the same fifteen invoices.

That is not a governance problem. It is an industrial engineering problem, and it now has an industrial engineering answer. The purpose of this paper is to describe that answer honestly: what it delivers, what it costs, how long it takes, what it does to your people, and what it does to the assurance model your Board relies on.

What this paper does not claim

Artificial Intelligence does not remove the Board's responsibility for the Internal Control System. That responsibility sits with the Board and cannot be delegated away, whether it arises under the Swiss Code of Obligations, under the corporate governance and audit requirements applying across the European Union and its member states, or under the Sarbanes-Oxley Act and the rules of the US Securities and Exchange Commission. It does not remove the need for judgment, for an external auditor, or for a competent owner of the framework, whatever that framework happens to be called, ICS, SOX, ICFR or otherwise. It does not make a badly designed control good.

What it does is remove the labor. And in most large organizations, the labor is between seventy and ninety percent of what the Internal Control System actually costs.

2. Why AI and the ICS generate real value

Accounting is routine by design. Routine is exactly what a machine does better than a person, and never tires of.

Financial reporting is the most rule-bound, most structured and most repetitive domain in the company. Every transaction sits in a system. Every approval leaves a trace. Every reconciliation follows a logic that somebody once wrote down. This is not a coincidence, but the entire point of double-entry bookkeeping and of an ERP.

That makes it the single best candidate for automation anywhere in the enterprise, and it is faintly absurd that we have spent two decades verifying it with clipboards.

The three ingredients

An AI-enhanced control is not a chatbot with an opinion about your accounts. It has three components, and all three have to be built properly:

- **Connectors.** Read access into the ERP and the surrounding systems, meaning general ledger, sub-ledgers, purchasing, treasury, HR and identity management. Complete, reconciled and monitored. This is plumbing, not intelligence, and it is where most of the real work sits.
- **The control expressed as a trained skill.** You take the control narrative, the same narrative your auditor reads today, and you turn it into an instruction the model executes. Then you train it: against known-good cases, against known-bad cases, against last year's actual exceptions. You tune it. You version it. You document the version. A trained control skill is nothing more exotic than a control design that has been made executable.
- **Evidence that is generated, not collected.** The single largest cost item in a manual ICS is somebody assembling proof after the fact. When the control runs inside the system, the evidence is a by-product of the execution, not a separate exercise.

What the output actually looks like

The deliverable is not «the AI says it is fine». The deliverable is a workpaper: population defined and reconciled, test performed, exceptions listed with the underlying records referenced, judgment calls flagged and routed to a named human, model version and run timestamp recorded.

In our experience that is a better workpaper than most of what is produced by hand today, because it is complete, reproducible and it does not depend on who was on vacation in week 32.

Control activity	How it works today	How it works AI-enhanced
Three-way match / PO compliance	Sample of 25–40 POs, manual tie-out to GR and invoice, screenshots filed	Full population matched continuously; only breaks and tolerance overrides surface for review
Journal entry review	Threshold-based extract, manual scan, judgmental follow-up on a handful	Every line scored against risk criteria (post-date, user, round sums, unusual pairings, weekend entries); reviewer sees a ranked exception list
Access review / Segregation of Duties	Periodic user listing exported, reviewed in a spreadsheet, signed off weeks later	Continuous conflict detection against the SoD ruleset; violations flagged at the moment of provisioning
Bank and account reconciliation	Prepared manually, reviewed manually, evidenced by signature	Auto-reconciled; the control becomes the disposition of the residual items

Control activity	How it works today	How it works AI-enhanced
Master data changes	Sampled after the fact, often quarterly	Every change tested against the approval record in real time
Cut-off and accruals	Analytical review on aggregates, sample of supporting documents	Line-level testing against contract, delivery and payment data across the full period

Note what these examples have in common: not one of them replaces a judgment. Each of them replaces the search for the things that require judgment.

3. A realistic roadmap

Nobody moves from spreadsheets to an ecosystem in a quarter. But nobody needs three years either.

The starting point for most large organizations is what we call a fully staff-managed ICS: a hundred key controls or a good deal more, performed and evidenced by people, coordinated by a central function, monitored by a second line, tested by a third, and re-tested by the external auditor. Everyone in that chain is competent. The chain itself is the problem.

The route out has four stages. The durations below overlap deliberately, because this is not a waterfall.

Stage	Timing	What happens	What the Audit Committee should ask
0 — Honest baseline	Weeks 1–6	Control inventory rationalization. Which controls are duplicative, which are legacy, which are genuinely key. Process and system mapping. A blunt assessment of data quality and ERP fragmentation.	«How many of our key controls would we keep if we designed the framework today?»
1 — Foundations	Months 1–6	Connector architecture. Data completeness reconciliation. Control narratives rewritten to be executable. AI governance policy approved. Custodian roles identified and named.	«Is our data complete, and how do we know?»
2 — Pilot	Months 4–9	Ten to fifteen high-volume, low-judgment controls in one process and one entity. Run in full parallel with the manual control. External Auditor invited to observe.	«What did the dual-run show, and where did the model and the humans disagree?»
3 — Scale	Months 9–18	Extension by process and by entity. Manual duplicates retired only where the dual-run evidence supports it. Exception handling moves into business-as-usual.	«Which manual controls have we retired, and on what evidence?»
4 — Ecosystem	Months 18–30	Monitoring and control converge. Internal Audit shifts to model assurance. Reliance agreement renegotiated. Rotational testing retired.	«What is our residual key-person dependency, and who assures the assurers?»

Two rules that decide the outcome

First: do not start with the tool. Vendor selection is the most visible decision and the least important one. Start with the process and the data. An organization that cannot reconcile its own transaction population will not be rescued by a license.

Second: do not automate a bad control. You will simply obtain a bad control faster, more often, and with markedly better documentation. Stage 0 exists precisely to prevent this.

«Solid foundations are of paramount importance. That was true when we digitized the framework. It is more true, not less, when we teach it to run itself.» — Reto H. Wenger

4. The real savings, once the system is properly trained

The savings are real. They are also not immediate, not uniform, and not where most people expect them.

For the purposes of illustration, take a group with roughly one hundred key controls, fifteen in-scope reporting entities, two ERP instances and an annual ICS cycle that occupies parts of ten months. That is a very ordinary shape for a former SOX filer or a larger Swiss group with a demanding ICS.

The ranges below reflect our experience with control frameworks, control rationalization and GRC IT deployments. They are indicative. They are not audited benchmarks, and no two organizations land in the same place, because outcomes are driven above all by data quality, ERP fragmentation and, frankly, by management's willingness to actually retire the manual control once the machine has proven itself.

Effort category	Indicative reduction	Where the saving comes from
Control preparation and evidence collection	60 – 75%	Evidence is generated by execution rather than assembled afterward
Control performance (first line sign-off)	30 – 50%	The routine controls run themselves; sign-off concentrates on exceptions
Central ICS coordination, chasing, status tracking	50 – 65%	There is nothing to chase; status is a property of the system
Second-line monitoring and review	40 – 60%	Monitoring converges with execution (see Section 11)
Internal Audit control testing	70 – 85%	IA tests five system properties instead of the entire control population
External Audit fees attributable to ICS work	15 – 30%	No archaeology; evidence already exists in the form the auditor needs
Deficiency remediation cycle time	Weeks → days	Exceptions surface on the day they occur, not in the following quarter
Audit Committee reporting preparation	70 – 85%	Reporting is generated from the ecosystem and reviewed, not built by hand

The part nobody puts on the slide

You do not get to zero. You trade a large, variable, people-heavy cost for a smaller, fixed cost in technology, connector maintenance and a handful of expensive custodians. That new cost line is permanent. It is also roughly an order of magnitude smaller than what it replaces and, importantly for a Board, it does not scale with entity count the way headcount does.

Year one is, in our experience, cost-neutral at best. There is a period during which you are paying for the build and still running the manual control in parallel, because retiring it before the evidence supports it would be reckless. Payback typically falls in the twelve-to-twenty-four-month range, and the curve steepens sharply from the second entity onward, because the marginal cost of adding an entity to a trained control skill is close to nothing.

Be skeptical of anyone who promises you savings in the first year. They are either selling a license or they intend to switch off the manual control before it is safe to do so.

5. The transformation of the profiles

The clerical work disappears. The people who understood why the control existed become more valuable, not less.

Let us be direct, because staff will see through anything else: this reduces headcount in ICS-adjacent roles. Any paper that claims otherwise is not being honest with you and, worse, the pretense itself will cost you the trust of the very people whose cooperation the project requires.

But look carefully at which roles disappear. They are the ones nobody wanted in the first place: collecting screenshots, chasing signatures, reformatting evidence into a template, and rebuilding the same Audit Committee slide twelve times a year. We have never met a control preparer who described that work as the reason they joined the company.

The profile that emerges: the custodian

What replaces the preparer is a genuinely demanding role. The custodian of a control skill has to understand the business process, the control objective, the underlying IT, and the behavior of the model, simultaneously. They have to be able to explain, in plain language, to an auditor and to an Audit Committee, why the control flagged what it flagged and why it did not flag anything else.

That is a rare and expensive profile. It is also, for the right person, a far more interesting job than the one it replaces.

Today's profile	Tomorrow's profile	What has to be trained
Control preparer	Control skill custodian	Instruction design, exception triage, reading and challenging model output
ICS coordinator	ICS ecosystem owner	Data lineage, connector governance, version control and release discipline for control skills
Second-line reviewer	Assurance analyst	Reasoning over full populations, drift and bias monitoring, threshold calibration
IA control tester	IA model assurance	Testing the mechanism rather than the transaction; validation methodology; override analytics

Today's profile	Tomorrow's profile	What has to be trained
Head of ICS	Custodian of the framework	AI governance, auditor liaison, board-level explainability

Retention: decide before the project, not after

The people you want are already in your team. In our experience they are, almost without exception, the ones complaining loudest about the manual work, because they understand the process well enough to see how much of their day is waste. Identify them in Stage 0. Tell them what the destination is and what their place in it will be. If you wait until Stage 3, you will be recruiting externally for a profile that the market prices at a substantial premium.

The principle we apply to ICS and Internal Audit staffing holds here with even greater force: these functions are not a terminus, they are a talent development center. A custodian who has spent three years understanding the process, the controls, the systems and the model is exactly the person you want running a finance organization afterward.

6. How light an AI-enhanced annual cycle becomes

The «annual» cycle stops being annual. It becomes continuous, and the annual part becomes a wrap-up.

Today's cycle is a calendar of near-permanent activity: scoping, walkthroughs, design effectiveness testing, quarterly operating effectiveness rounds, an interim push, a year-end push, remediation, and then reporting. Ten months of it, with two months to recover before it starts again.

In an AI-enhanced framework the controls do not have a season. They run with the transaction. Exceptions surface on the day they occur. What is left of the annual cycle is this:

1. Annual scoping and risk refresh. Two to three weeks. Pure human judgment on materiality, entity scope, changes in the business and new systems. Unchanged, and it should be.
2. Design validation of the control skills. Conducted continuously through the year, formalized once annually: has each skill been revalidated against known cases, has any version changed without approval, are the connectors still complete?
3. Operating effectiveness. Already evidenced. Three hundred and sixty-five days, full population, machine-generated. There is nothing left to test, only something to read.
4. Exception disposition. Handled during the year as items arise, not compressed into a January panic. By year-end the population of open items is small and known.
5. Audit Committee reporting. Generated from the ecosystem, reviewed and contextualized by the custodian and, for the first time, reconcilable to the underlying data on request during the meeting itself.

The formal calendar footprint collapses from something like ten months to a handful of weeks of genuine judgment work. What remains is, notably, the part of the job that anyone running an internal control framework, ICS, SOX or any other, would tell you was always the valuable part.

7. External Audit: process, control, system and framework become one ecosystem

The cheapest audit is the one where the evidence already exists in the form the auditor needs.

Consider why testing an Internal Control System is expensive for an external auditor. They must understand your process, understand your systems, locate the control within them, locate the evidence that the control operated, test a sample of it, and reconcile the population back to the financial statements. Six steps, of which five are archaeology.

In a properly constructed ecosystem, the process documentation, the control narrative, the executable skill, the evidence and the tested population are no longer five separate artifacts maintained by five separate people. They are one object. The auditor does not have to reconstruct anything. They inspect what is already there, and they can re-run it.

This is the point at which the phrase «one ecosystem» stops being consultant language and becomes an audit fact: the framework describes the control, the control is the code, the code produces the evidence, and the evidence is the population.

What this is worth, and what it requires

- Fee impact: in the range of 15 to 30 percent of the ICS-attributable portion of the audit fee. Meaningful, but not the headline.
- Disruption impact: this is the headline. The prepared-by-client list stops being a project. Your finance team stops losing six weeks a year to evidence requests.
- Timing impact: interim and year-end testing compress, because there is no sample to accumulate.

One requirement is absolute, and ignoring it is the most common and most expensive mistake we see: engage the external auditor early. Not at the end. An auditor who first encounters your model in the year you intend to rely on it will test it as though it were new, because to them, it is. Invite them into Stage 2. Let them watch the dual-run. Let them shape the evidence format before it is built, not after. The cost of that involvement is a fraction of the cost of a rejected reliance.

8. From sample to population

The sample was never a methodology. It was a concession to the cost of looking.

Twenty-five items. Forty. Sixty for a daily control. Every practitioner knows where those numbers come from, and it is not from risk. They exist because a person can only look at so many things in a working day. Sampling has always been a compromise between what we wanted to know and what we could afford to find out.

A trained model looks at all of them. Every journal line. Every purchase order. Every master data change. Every access grant. Not because it is clever, but because it does not get bored and the marginal cost of the four-millionth record is effectively zero.

Three consequences follow, and they compound:

- Less work, because there is no sample to select, no sampling methodology to defend, and no extrapolation to negotiate.
- Fewer resources, because the effort no longer scales with the number of controls, entities or transactions.

- Higher assurance, because a statement about a population is categorically stronger than an inference from a sample. This is the part that matters to a Board.

It also changes what a deficiency is

Today, a single exception in a sample of forty triggers a familiar and rather unsatisfying conversation: is it isolated, is it systemic, do we extrapolate, what is the projected misstatement, and can we all live with the conclusion. It is an argument about statistics conducted by people who would rather be arguing about controls.

In a full population you do not extrapolate. You know. «Three exceptions in eight hundred and twelve thousand transactions, all in the same cost center, all in March, all by the same user» is a fact, and it is an actionable one. That is a fundamentally better conversation to have with an Audit Committee, and it is a fundamentally better basis for remediation, because the root cause is visible rather than inferred.

A warning worth giving to the Board in advance

The first full-population run will find things. It will surface exceptions that were always there and that sampling never happened to touch. This is not a deterioration of the control environment, but the first honest measurement of it, and the comparison to last year's clean sample-based conclusion is not a like-for-like comparison.

Prepare the Audit Committee for this before the first run, not after. An organization that has not been warned will read its first accurate measurement as a failure, and the instinct in that moment is to blame the instrument.

9. Internal Audit: a genuinely ultra-light testing approach

Internal Audit stops re-performing controls and starts assuring the machine that performs them.

Today, a large part of Internal Audit's control testing effort, whether the framework is called ICS, SOX or anything else, consists of re-testing a sample of a sample: the second line reviewed a selection, and the third line independently selects from the same population to check whether the second line selected well. It is defensible. It is also, viewed from a distance, a remarkable use of some of the most capable people in the organization.

In an AI-enhanced control framework, Internal Audit's object of assurance changes entirely. There are five things to test:

1. Completeness of the data. Does the connector deliver the whole population, reconciled to the general ledger, with no silent filtering?
2. Fidelity of the skill. Does the executable control actually do what the narrative says it does, tested against seeded exceptions and known-bad cases?
3. Integrity of the version. Has any skill, threshold or rule changed without going through approval? Is the change log complete?
4. Disposition of exceptions. Were flagged items actually investigated and resolved, or were they closed?
5. Human override. Who overrode the machine, how often, on what basis, and does the pattern of overrides tell a story?

Five things instead of an entire control population. Four of the five are established once and monitored thereafter. Only the fifth, override analysis, requires recurring judgmental work, and it happens to be the single most informative test in the entire program.

Indicative effort

An annual control testing effort that today occupies an auditor for two to four weeks moves, in our estimation, to roughly two to four days. That is not a marginal efficiency, it is a change of category. It is the difference between an Internal Audit function that spends a meaningful part of its year on compliance testing and one that spends that same time on the risks the Board actually loses sleep over.

And with a reliance agreement, lower still

Reliance on Internal Audit work has always been constrained by the external auditor's confidence in two things: the sampling and the documentation. When there is no sampling, and when the documentation is machine-generated, timestamped, versioned and re-runnable, the traditional objections to reliance substantially fall away.

The prize is that Internal Audit and External Audit test the same five properties, once, on the same evidence, rather than each rebuilding their own view of the entire control population. In a mature ecosystem this is where the largest single block of duplicated assurance effort in the organization simply ceases to exist.

10. The end of the rotational testing model

Rotation exists because you cannot afford to look everywhere every year. When looking everywhere is free, rotation has no purpose.

Rotational scoping, testing entity A this year, B next year and C in year three, is a resource allocation device dressed as a risk methodology. Nobody genuinely believes entity C carries no risk in years one and two. We rotate because we have a fixed number of auditors and a variable number of locations, and we have collectively agreed to call the resulting compromise «risk-based coverage».

A trained control skill runs wherever the data is. Adding the sixteenth entity to scope is a configuration exercise, not a hiring exercise. Anywhere. Anytime. Simultaneously. The marginal cost of universal coverage approaches the marginal cost of the connector.

What follows from that is worth stating plainly to a Board:

- The three-year coverage cycle loses its rationale. Coverage becomes annual and complete by default.
- The annual scope negotiation, always one of the more political conversations in the assurance calendar, largely disappears, because scope is no longer rationed.
- The effort tied to continuous control monitoring disappears with it. Monitoring is no longer an activity that must be resourced. It is a property of the system.

Theoretically, then, rotation ends. In practice we would temper that in one respect, and we would say so to any client: judgmental sampling on the highest-risk, most judgment-dependent controls, such as management estimates, related-party arrangements, going-concern-relevant judgments and unusual transactions, remains a human exercise, and there rotation may still make sense. Rotation dies for routine controls. It does not die for judgment.

11. Monitoring and control become one layer

A control that runs continuously and reports its own exceptions is not a control that needs monitoring. It is the monitoring.

Consider what control performance monitoring actually is. The first line performs a control. The second line then verifies that the first line performed it, on time, completely, and with adequate evidence. An entire discipline has grown up around this, comprising performance dashboards, completion tracking, escalation protocols and reminder emails, and it exists for exactly one reason: we cannot be certain that a person did the thing they said they did.

When the control executes inside the system, on the full population, and writes its own evidence as it goes, that uncertainty is gone. Performance and monitoring are no longer two events separated by a verification step. They are the same event.

The monitoring layer is therefore not abolished by decree. It becomes redundant by design, which is a far more durable way to remove a layer, because it does not require anyone to accept a reduction in assurance.

What remains

Something does remain, and it should not be minimized: monitoring the machine. Drift, threshold calibration, connector health, version integrity, override patterns. This is a small, specialized and high-value activity, it requires genuine expertise, and it belongs unambiguously to the custodian. It is a fraction of the effort of monitoring people, and it is considerably more interesting work.

12. What this means for the Three Lines of Defense

Three lines were an answer to a question about people. Change the actor, and the question changes.

The Three Lines model, for all its merit, is at bottom an organizational answer to a human problem. People who perform work cannot objectively assure their own work. So we add a second line. But the second line is also management, so we add an independent third. The architecture follows from the conflict of interest, not from the control.

When the control is executed by a versioned, logged, auditable system that has no career, no bonus and no reason to prefer one answer over another, the conflict of interest that justified the separation is substantially reduced. Not eliminated, because somebody configures that system, somebody sets its thresholds, and that somebody very much has interests. But reduced, and reduced in precisely the place where the three lines cost the most.

	Three Lines today	In an AI-enhanced ecosystem
First line	Performs the control and produces evidence	Owens the process and disposes of exceptions; the routine execution is machine-performed
Second line	Monitors that the first line performed, reviews evidence, reports status	Converges with the first line, since monitoring is a property of execution (Section 11); residual role is threshold and policy ownership

	Three Lines today	In an AI-enhanced ecosystem
Third line	Independently re-tests a sample of controls across a rotational scope	Assures the ecosystem itself: data completeness, skill fidelity, version integrity, exception disposition, override patterns
Net shape	Three organizational layers, substantial duplication, effort scales with entities and controls	One operating line that performs, evidences and monitors, plus a thin, genuinely independent layer that assures the mechanism

We describe the result as one line plus a very thin independent layer. Not «no lines», because that would be a governance regression and no serious auditor or regulator would accept it. But one operating line that captures performance, evidence and monitoring in a single motion, effortlessly, because the effort has been engineered out rather than reassigned.

A caution for the Board specifically

This is a governance change, not merely an efficiency change, and it must be a conscious, documented, Board-approved one. The failure mode we would warn against is drift: the second line thins out gradually through attrition and hiring freezes, nobody ever decides anything, and two years later the organization discovers that its assurance architecture changed without a decision having been taken.

And one point deserves emphasis precisely because it runs against the grain of the savings argument: the independence of whoever assures the model must be protected with more rigor than before, not less. There are fewer of them. Their conclusions carry more weight. Their access to the Audit Committee, their budget and their freedom to escalate should if anything be strengthened as the rest of the structure lightens.

13. What can go wrong

The risks are real and they are manageable. Pretending they do not exist is the fastest way to lose the Audit Committee.

We would rather present these ourselves than have your auditor present them for us.

Risk	Why it matters	How it is managed
Data completeness	The single largest technical exposure. A control tested on 100% of a 92% complete dataset is worse than a sample, because it feels complete.	Reconciliation of every connector to the general ledger at every run; completeness itself becomes a tested control
Model drift	A skill tuned on last year's data becomes quietly wrong as the business changes.	Scheduled revalidation against seeded and known cases; drift metrics reported to the ICS function and to Internal Audit
Bias in the ruleset	Thresholds encode assumptions. Wrong assumptions produce systematically blind spots rather than random ones.	Independent review of thresholds; periodic challenge of what the model never flags, not only of what it does
Explainability	If the custodian cannot explain in plain language why an item was flagged, the auditor will not rely on it, and neither should the Board.	Explainability is a design requirement from day one, not a feature request later
Change management	An ERP upgrade or a connector change can silently break a control.	Version control, regression testing against a fixed case library, and release approval treated as a key control in its own right
Over-automation	Automating a genuinely judgmental control produces a confident wrong answer, which is worse than no answer.	A deliberate, documented decision on which controls stay human, and the discipline to defend that list
Key-person concentration	Fewer people understand the ecosystem. This exposure is consistently underestimated.	Minimum two custodians per domain, documentation standards, and succession treated as a Board-level topic
AI governance and regulation	Data protection, data residency, vendor dependency and, for groups with EU exposure, the EU AI Act.	A Board-approved AI governance policy covering ownership, approval, logging, override and vendor management, in place before the first control goes live

None of these is a reason not to proceed. Each of them is a reason to proceed in the order set out in Section 3.

14. This is not ambition. This is what we do.

Real-world expertise, not theory. We do not pursue textbook ideals.

Everything described in this paper sits inside work LakeHill GRC already delivers. It is not a new practice area assembled around a trend. It is our existing core service performed with a better instrument.

- We design and implement Internal Control frameworks from scratch, covering scoping, process documentation, risk-control matrices, control design, dry-runs, remediation and implementation. An executable control skill is a control design. The discipline is identical.
- We reshape the IT and non-IT elements of an organization's process structure to give a control framework foundations it can actually stand on, which is precisely the Stage 0 and Stage 1 work described above, and which we were doing long before anyone attached the word «AI» to it.
- We make IT ecosystems compatible with internal control requirements. Connector architecture and data completeness are the modern expression of that same mandate.
- We run pre-audit assessments and audit-readiness testing, where the dual-run discipline in Section 3 is our standard method, applied to a machine rather than a new manual control.
- We establish liaison with Internal and External Audit, including reliance agreements. Section 9 is not a theory for us; negotiating that relationship is an existing service line.
- Our Diagnostics in a Box governance health check already covers AI Governance alongside Access, Security, Continuity, Change and Cyber, because clients started asking, and because it belonged there.

We are not a software vendor and we do not resell licenses. We have no incentive to tell you that a tool will solve a data problem, and we will tell you plainly when your foundations are not ready, because a control framework built on weak foundations, ICS, SOX or any other, does not become stronger by running faster.

«Good governance is not bureaucracy. It is a business enabler. Nothing in this paper changes that. It simply removes the bureaucracy that was never governance in the first place.» — Reto H. Wenger

Where a conversation usually starts

For most organizations the sensible entry point is a short, time-boxed, non-intrusive assessment: what does your control framework actually consist of, how much of it is genuinely automatable today, how complete is your data, and what would the first pilot look like. That is a matter of days, not months, and it produces a decision, not a proposal.

15. Summary and Call to Action

Effective assurance with dramatically less effort is achievable. It is engineering, not magic, and it is earned, not bought.

Summary

1. Artificial Intelligence does not lower assurance in an Internal Control System. Done properly, with full population coverage, generated evidence and versioned skills, it raises it substantially.
2. The savings are real, but they are earned in year two, not year one. Anyone promising otherwise intends to retire the manual control before it is safe.
3. Full-population testing changes the nature of the deficiency conversation from statistical inference to demonstrable fact, and the first honest measurement will be uncomfortable.
4. The profiles change more than the headcount does. Identify and retain the people who understand why the control exists, and do it in Stage 0.
5. Monitoring and control converge. The three lines become one operating line plus a thin, genuinely independent assurance layer.
6. The governance implications require a conscious Board decision. Assurance architecture must never change by attrition.

Call to Action

1. Ask your Head of ICS which of your key controls could run on the full population today. The honest answer is almost always higher than the Board expects.
2. Do not start with a tool. Start with your data and your processes. Licence selection is the most visible decision and the least consequential one.
3. Bring your External Auditor in at the pilot, not at the reliance. Their early involvement costs a fraction of a rejected reliance.
4. Identify your future custodians now and tell them what the destination is. They are already in your team, and they are the ones complaining about the manual work.
5. Dual-run everything. Retire no manual control until the evidence supports retiring it, and document that evidence as though the auditor will read it, because they will.
6. Approve an AI governance policy before the first control goes live. Not after the first exception, and certainly not after the first audit finding.

The organization that gets this right will not end up with a smaller Internal Control System. It will end up with a better one, run by fewer and considerably better people, tested once, and trusted more.

First, foremost and above all: clear, unambiguous and transparent communication at all levels. Always.

THANK YOU FOR YOUR ATTENTION

Reto H. Wenger, Founder and Managing Director

BSc. | CAS | EMBA | Certified Internal Auditor

In the United States:

LakeHill GRC LLC, 7901 4th St. N Ste 300, St. Petersburg, FL 33702, USA

info@lakehillgrc.com | +1 407 456 7353 | lakehillgrc.com

In Switzerland:

LakeHill GRC Reto H. Wenger, Hungerstrasse 51, 8832 Wilen bei Wollerau, Switzerland

reto.h.wenger@lakehillgrc.com | +41 79 613 4990 | lakehillgrc.com/de

Copyright © 2026 LakeHill GRC LLC. All Rights Reserved.

About LakeHill GRC

LakeHill GRC operates independently under Swiss and U.S. legal frameworks. LakeHill GRC is a sole proprietorship based in Wilen bei Wollerau, Switzerland, and also registered as an LLC in the State of Florida, USA. These are two legally distinct entities that do not share client data, operations, or contractual obligations. The only commonality is a shared public-facing website, used solely for informative and promotional purposes. Swiss-based clients are exclusively served by the Swiss entity, subject solely to Swiss law and jurisdiction.

The founder and owner, Reto H. Wenger, is overseeing and providing all core services, with access to a wide network of seasoned professionals, within regulated and non-regulated services, both in the central Switzerland region, the DACH region and globally.

LakeHill GRC does not engage in any activities that fall under regulated professional services. This includes, but is not limited to, statutory financial audits, legal representation or advice, tax advisory services requiring formal authorization or any other services subject to regulatory licensing or oversight under US, Swiss or EU law. Furthermore, LakeHill GRC does not provide any services that involve or require compliance with the Swiss Anti-Money Laundering Act (AMLA), the Financial Institutions Act (FinIA), or the Financial Services Act (FinSA).

Indicative figures, ranges and durations in this paper are drawn from LakeHill GRC's professional experience with internal control frameworks, control rationalization and GRC IT deployments. They are illustrative and are not audited benchmarks, guarantees or assurances of outcome. Any organization's results depend on its data quality, systems landscape, control design and management decisions.

Statement on the Use of Artificial Intelligence in the Preparation of this Document

This whitepaper was prepared with the assistance of an artificial intelligence system. That assistance was confined to language, terminology, grammar and spelling. The subject matter, the structure, the professional analysis, the judgments, the illustrative ranges and every substantive conclusion set out in this document originate from the author, Reto H. Wenger, and reflect his own professional experience and opinion. No section of this document was generated autonomously and adopted unread: the author has reviewed, verified and taken ownership of every statement made here.

This disclosure is made voluntarily and in the spirit of the transparency obligations set out in Article 50 of Regulation (EU) 2024/1689 (the EU Artificial Intelligence Act), which became applicable on 2 August 2026, and of the California AI Transparency Act (Senate Bill 942, as amended by Assembly Bill 853), operative from 2 August 2026, as well as the general prohibition on misleading or deceptive representations under Section 5 of the US Federal Trade Commission Act. LakeHill GRC regards the disclosure of AI assistance in client-facing materials as good professional practice irrespective of whether a specific statutory obligation applies in a given jurisdiction, and applies the same standard to itself that it recommends to its clients.