LakeHill GRC

LAKEHILL GRC
WHITEPAPER

Too Many Controls, Too Little Control

How to cut an overgrown ICS back to the risks that really matter, and make it AI-ready

> *«Every incident got its own new control. Not one of them ever got a new risk assessment.»* — Reto H. Wenger

Prepared for CFOs, Senior Finance Management and Finance Team Leads

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

ICS & SOX | Internal Audit | Risk Management | Compliance

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

1. How a good ICS turns into a bad one

Nobody designs an unmanageable Internal Control System. It grows there, one incident at a time.

Has your ICS grown over time, at times uncontrolled, and is it now in an almost unmanageable state? Far too many controls, many of which are not controls at all, but checkpoints and detailed work instructions? And does this sound familiar: whenever something had gone wrong, a new detective control was written up quickly, so that there was an answer for the auditor or the Audit Committee by the next meeting?

If so, you are in good company. It is the single most common pattern we see in finance organizations of all sizes. The framework was set up years, often decades ago, mostly with a clear logic and a defined scope. Since then, controls were added after every incident, every audit finding, every reorganization and every system migration. Hardly any were ever removed. And needless to say, hardly any were brought up to speed alongside an IT ecosystem that has evolved considerably over those same decades. Many controls still check things the system has long since enforced on its own, or check them in a way that made sense when the ERP was fifteen years younger.

One distinction matters more than any other, because everything else follows from it. A control mitigates a defined risk. It has an owner, a frequency, a clear definition of what is looked at and of what happens when something is found. «Check that the file has been uploaded» is not a control. «Print the report, sign it and file it» is not a control either. These are work instructions. They have their place, in a procedure manual. In your ICS they do nothing but inflate the inventory, consume testing effort and dilute the attention the real controls deserve.

What you end up with

- **Bureaucracy.** Sometimes hundreds of controls, and a finance team that spends a considerable part of its time preparing, signing and chasing evidence instead of doing finance. Whether a control is key or non-key is oftentimes decided by judgment rather than derived from a risk assessment. And the inventory rarely sits anywhere central: more often than not, it lives on a spreadsheet in a folder that nobody really knows of, in a version nobody is quite sure about.
- **Manual detective controls everywhere.** Excel extracts, email approvals and screenshots, performed after the fact on a sample of 25 out of 40,000 transactions. Whatever the sample says, it says very little about the other 39,975. That is not meaningful assurance. It is a ritual with a good conscience.
- **No overview of the risks that matter.** You have a control inventory. You do not have a risk picture. Ask which of your significant risks are mitigated by a preventive control, which by a detective one and which by none at all, and see how long it takes to get an answer.
- **A framework miles away from its origin.** Away from what was originally set up, from how your business works today, and from what your systems are capable of today.

© 2026 LakeHill GRC LLC – No frills, just results.

2. The way out: back to process, systems and real risks

Start where you should always start: with the process, the system and the risk.
Not with the (legacy) control list.

There is a solution, and it is not a new tool. Revisit the underlying processes, your current system landscape and the risks that are real. Do it from the ground up, stream by stream, not line by line through the existing control list. The control list is oftentimes the result of the problem, not the starting point for the solution.

You will most likely discover that the progress IT has made over time changes the picture fundamentally. Today's ERP systems enforce approval workflows and authorization limits, three-way match tolerances, duplicate invoice checks, dual control on master data changes, period locks and segregation of duties through role design. Many of your detective controls compensate for weaknesses your system no longer has. They were written for a system landscape that no longer exists.

This allows you to switch from a pile of detective controls to a much smaller number of preventive controls, keeping only those few detective controls that remain reasonable. In practice, the combination of an effectively rigid segregation of duties, enforced through roles and system access rather than through a policy and a signature, and a small set of well thought-through preventive (and even less detective) controls has the potential to do most of the job.

Then comes the step that requires the most discipline: kill all the risks that are not real risks. The risk that a file is saved in the wrong folder is not a financial reporting risk. Neither is the risk that a report is not printed. What remains after that exercise is compact and comprehensive. For the four core business process streams, the significant risk categories could look like the example below. It is an illustration only: the actual picture depends on your industry, your size and the (regulatory) framework you operate under.

Process stream	The risk categories that matter	Where the system carries the load
Book-to-Report	Unauthorized or inaccurate manual journal entries, including management override; incomplete close and cut-off; judgmental estimates and valuations; consolidation and intercompany	Posting authorizations, workflow approval of manual entries, period locks, system-based reconciliations
Hire-to-Retire	Unauthorized changes to employee and pay master data; fictitious employees; inaccurate payroll calculation and payment; leavers not processed in time	Changes to salary and bank details blocked from payroll release until approved by a second person; no payment to anyone without an active, approved employee record in HR (no record, no pay); automated access removal on exit
Purchase-to-Pay	Fictitious or unauthorized vendors and bank detail changes; payment for goods or services not ordered or received; duplicate payments; segregation of duties	Dual control on vendor master data, three-way match with hard tolerances, duplicate invoice check, role-based SoD

© 2026 LakeHill GRC LLC - No frills, just results.

Process stream	The risk categories that matter	Where the system carries the load
Order-to-Cash	Revenue recognized without performance or in the wrong period; unauthorized prices, discounts and credit notes; billing accuracy and completeness; dunning and collection of receivables	Price master and approval workflows, delivery-based automatic billing, controlled release of billing blocks, system-driven dunning runs, automated cash application

This is deliberately not the full catalogue. The detail, and the value, lies in mapping these categories against your own processes, your own systems and your own history of incidents. That is where the real work is done.

3. Fewer controls. Better controls. Real controls.

Retire what no longer serves you. Keep what genuinely mitigates risk.

Once this is done, you will see, beyond doubt, that the total number of your business process controls diminishes dramatically: oftentimes to fewer than ten per stream. And oftentimes, more than half of those are preventive controls embedded in the ERP/system.

The big advantage of a preventive control is that it always works, and works steadily. It does not take holidays, it does not depend on who is on duty in week 38, and it does not wait for quarter end. It simply blocks what should not happen and reports the exceptions. A detective control is, at best, IT-dependent: a manual review relying on a system-generated report whose completeness and accuracy has to be assured separately. In all cases it uses significantly more resources whilst producing significantly less assurance.

That changes the economics of ICS assurance as well. An automated application control, supported by effective IT general controls, is tested on its configuration and its change history, not on a sample of forty transactions every quarter. Fewer controls, tested less often, with more assurance. That is not a contradiction. That is design.

Retire the screenshot

While you are at it, retire the screenshot. You know the routine: a screenshot from the ERP, pasted into a Word or Excel document, turned into a PDF, stored in a shared folder and emailed to the supervisor for sign-off and, months later, to the auditor. Multiply that by every control, every period and every entity, and a considerable share of your finance team's capacity is tied up in documenting rather than controlling.

In nine out of ten cases, a screenshot proves only that somebody took a picture at a given moment. It does not prove that the population was complete, that the right people approved, or that the risk was mitigated. In rare instances a screenshot does make some sense, for very traditional detective controls. But then ask why that control exists in the first place: either there is no preventive control, or there is one that cannot really be relied upon, or it is legacy bureaucracy that should be questioned and eventually removed. A preventive control needs no screenshot: its evidence is its configuration and its system log. And the remaining detective controls should produce their evidence where they are performed, in the system, not in a folder.

© 2026 LakeHill GRC LLC – No frills, just results.

4. Now your framework is AI-ready

Only a clean framework can be automated. A messy one can only be automated messily.

At this point, and only at this point, the entire framework becomes AI-ready. You now have a reasonable number of controls, conducive to the business, each clearly defined, each tied to a real risk and each based on system data rather than on documents assembled by hand. That matters, because automation and AI can only work with what is structured, defined and available in the system: a clear control logic, complete data and a defined expected outcome. A screenshot in a PDF in a shared folder is none of these.

The remaining resource-intensive standard routine controls, such as account reconciliations, journal entry reviews, access reviews and master data change reviews, can then be turned into trained control skills that run on the full population instead of a sample, generate their own evidence and route only the exceptions to a person who decides. How that works, what it saves and what it requires in terms of data, governance and people is set out in our whitepaper *«Internal Control and Artificial Intelligence – From an army of preparers to a single control ecosystem».*

Avoid the quick cost-saving trap

The temptation is understandable. Automation promises quick savings, and the business case is easy to build. Skipping the steps in between, however, rarely delivers what it promises. Automating a control framework that did not work to your satisfaction produces an automated framework that still does not work to your satisfaction, only faster. Exceptions that nobody can explain accumulate, auditors become cautious about relying on the new set-up, and manual controls are gradually reinstated alongside the tool. In our experience, frameworks automated this way often run into serious difficulty within the first year.

At that point the projected savings are gone, however convincing the business case looked on paper. And your audit bill will most likely wipe out those savings ambitions for years: an auditor who cannot rely on the automated controls has no choice but to fall back on extensive substantive testing. The sequence that works is unspectacular but reliable: rationalize, redesign, embed in the system, and only then automate. Each step pays for itself, and each one makes the next one cheaper.

In short: a lean, risk-based framework built on system-embedded preventive controls is what makes automation and AI worthwhile. Quick-and-dirty, cost-saving driven AI automation on top of an overgrown framework will backfire. Always. And your return on investment moves out by years, because in the meantime your auditor will test the way it was done in the 1980s: substantively, line by line, since he has not placed, and cannot place, any reliance on your automation. And you will pay for it with the highest audit fees ever.

© 2026 LakeHill GRC LLC - No frills, just results.

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 in this paper, such as control counts per stream, 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 drafting the text from the author's detailed outline, and 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.