HOW TO SURVIVE A DATA AUDIT

What Your Auditor is Looking for and what to do before they arrive

What a Data Audit Actually Tests

A data audit tests whether your controls will hold up. Most organisations are not as well-positioned to answer it as they believe.

Over the past decade, insurance firms have invested heavily in the architecture of data governance - data councils, data owners, stewards, policies, frameworks. Additionally, they can point to certain controls running as evidence of operation.

This investment is not wasted. But it has produced organisations that are very good at describing their approach to data quality and how it is working and considerably less certain about *whether the controls giving effect to that approach are actually effective and sufficient.*

Having a framework and having assurance are not the same thing.

An auditor - whether internal, external, or regulatory – will do more than just review your governance documentation and evidence of control operation. That's a given.

They'll want to know how you have designed your data controls framework and why they should accept that design as being effective.

Auditors will not even consider reviewing the operational effectiveness of a control without ensuring that there is an adequate basis for its design.

This guide gives you some of the questions they might ask, where organisations most commonly fall short and what you can do about it before the audit begins.

The Five Questions Auditors Ask

These are not trick questions. They are the five points at which the rigour of your control design becomes visible – or its absence does. In our experience, most organisations struggle with at least three of them.

1. "How did you identify and articulate the risk this control is meant to mitigate?"

This is the first and most important question, because precise risk identification is the foundation of good control design. It is not enough to say there is a risk of data quality failure. If the risk sits in a manual spreadsheet process – say, ensuring formulae are extended to capture new data – the risk that needs articulating is that the formulae are not fully extended, not a generic reference to data quality. A control built against a vague risk statement is a control built against nothing in particular, and it shows the moment an auditor asks what, specifically, it is guarding against.

2. "How did you determine the extent of the inherent risk?"

Without assessing the impact and likelihood of the risk, deciding how rigorous or strong a control needs to be is guesswork. A control that is too light for a high-impact, high-likelihood risk leaves the organisation exposed. One that is too heavy for a low-impact risk wastes effort that could be spent elsewhere. Either way, if the assessment behind that judgement was never written down, it cannot be demonstrated to an auditor – and an undemonstrated rationale is, for audit purposes, the same as no rationale at all.

3. "Is the control narrative clear – and does it say who is accountable?"

A control narrative needs to set out what the control does, why it exists, who is accountable for it, how it is performed, and how its operation is to be evidenced in practice. Leave out the how, and the control will not be operated consistently – everyone will do it slightly differently, if they do it at all. Leave out a single named owner, and the control will quietly die: without one person accountable for it, it becomes someone else's job, and eventually nobody's.

4. "What is the residual risk, and who has signed off that it is within appetite?"

Residual risk is what remains once the control has been applied. It is rarely zero. The questions that matter are whether that residual level has actually been assessed rather than assumed, whether it has been measured against a defined risk appetite, and whose name is against the decision to accept it. A residual risk that has never been sized, or that nobody senior has formally accepted, is not a residual risk assessment. It is a hope.

5. "What compensating controls do you have if the primary control fails?"

No control operates perfectly all of the time. The question is what stands behind it if it does not. This means being able to say what the compensating control actually is, how it was decided that this particular compensating control was the right one, and where in the process it sits relative to the primary control it is backing up. A compensating control that was never deliberately chosen – or that sits in the wrong place to catch the failure it is meant to catch – offers no real protection at all.

The Three Places Frameworks Usually Break

Control frameworks fail in predictable ways. The following three failure modes account for the majority of data audit findings in the insurance firms we have worked with and reviewed.

The Shadow Process Problem

Controls are designed to govern the process as it is documented. They rarely govern the process as it is actually operated.

In most organisations, there is a meaningful gap between the two. Workarounds evolve. Manual interventions become standard practice. Intermediate steps – data extracted to a spreadsheet, manipulated, and re-uploaded – occur outside the systems that your controls are monitoring. These shadow processes are frequently where the most material data quality risks reside, and they will not appear in a system diagram or a process document.

Consider a reinsurance netting process where post-committee adjustments are made in Excel before figures are posted to the ledger. The downstream controls over the ledger entry may be robust. The Excel step is uncontrolled, undocumented, and invisible to the framework. An auditor who asks to see the full data journey from treaty committee to the reserving system will find it.

The Narrative Gap

A control that cannot be operated consistently by anyone other than its original designer is not a control. It is institutional knowledge with a register entry.

The test of a well-designed control narrative is straightforward: could a new joiner, reading it for the first time, execute the control correctly and produce evidence in a form that a second-line reviewer could assess? In most control registers we have reviewed, the answer to that question is no for a significant proportion of controls. The methodology is vague, the evidential requirements are undefined, and the frequency has been set by convention rather than by risk.

This is not a failure of intent. It is a failure of design. And it compounds over time: controls that are loosely defined are operated inconsistently, inconsistently operated controls generate variable evidence, and variable evidence is precisely what an auditor uses to challenge operating effectiveness.

The Inherent/Residual Illusion

Residual risk is frequently assessed as low not because controls are strong, but because nobody has looked hard enough.

There is a structural bias in internally-conducted risk assessments toward crediting controls with more mitigation than they provide. Controls that are technically present receive full credit even when operating consistency is unknown. Inherent risk is graded conservatively to avoid uncomfortable conversations. The result is a residual risk picture that looks reassuring but does not reflect reality.

For firms operating internal models under Solvency II, or producing IFRS 17 disclosures, the stakes here are not abstract. A residual risk assessment that overstates the effectiveness of controls over SCR or IBNR data flows is not simply a documentation problem. It is a misrepresentation of the firm's actual risk position to anyone relying on that assessment – including the board.

A Self-Assessment: Before the Auditor Arrives

Answer these questions honestly. If you find yourself hesitating on more than three, that hesitation is the answer.

1. Can you state, in one precise sentence, the specific risk each key control is designed to mitigate – rather than a generic description such as "data quality risk"?
2. Has the inherent risk behind each control been assessed for impact and likelihood, with that assessment written down rather than held in someone's head?
3. Does every control narrative set out the what, why, who, how, and the evidence to be produced – in enough detail for someone else to operate it?
4. Is there a single named individual, rather than a team or a department, accountable for each control?

5. Has the residual risk for each control been calculated, measured against a stated risk appetite, and formally signed off?
6. Do you know what compensating control exists for each key control, why it was chosen, and where it sits relative to the control it is backing up?
7. Was this reasoning documented at the point of design, or would it need to be reconstructed if someone asked?
8. Has your control design been reviewed by someone who was not involved in building it?

If you found yourself hesitating, that is the answer.

What Good Looks Like

A data control framework that performs under scrutiny is not necessarily more complex than the one you have. In most cases, the difference is not in the number of controls. It is in how precisely they are designed, how honestly they have been assessed, and whether they govern the process that actually exists rather than the one that was last documented.

Organisations that have built their frameworks on this basis share certain characteristics. They can produce control evidence on request without having to reconstruct it. Their residual risk assessments are grounded in observed performance, not assumption. Their data flow documentation captures the shadow processes as well as the formal ones. And when an auditor pulls the thread, the framework holds.

That is not a theoretical standard. It is an achievable one. But it requires the right methodology, applied to the right flows, with the kind of independent challenge that is difficult to sustain from the inside.

The gap between where most firms currently sit and where they need to be is real. The firms that close it before an audit are the ones that never have to explain it afterwards.

The Next Step

I have spent over thirty years in the London insurance market - designing and implementing data control frameworks and auditing CDO functions responsible for them. I know what robust looks like from both sides of the table.

A thirty-minute discovery call will give you an honest, independent assessment of where your framework currently stands: where the most material gaps are likely to lie and what a realistic path to remediation looks like. No deck. No proposal. No sales pitch.

Just a straight answer from someone who has seen what auditors find - and helped organisations make sure they find it first.

Book your discovery call: calendly.com/navin-ahuja-provenancedatarisk/30min

The cost of that conversation is thirty minutes. The cost of not having it tends to be considerably higher.