

THE MISSING SEAT

When Growing Companies Need an
Independent Strategic Adviser



LARRY HARMON

ROWAN PRIME

The Missing Seat

When Growing Companies Need an Independent Strategic Adviser

Larry Harmon | ROWAN Prime

2026 edition | Publication candidate 1.0

Contents

A note to the reader

Chapter 1 The Empty Chair

Chapter 2 The Cost of Not Knowing

Chapter 3 The Resilience Gap

Chapter 4 The AI Governance Crisis

Chapter 5 Intelligence Without Judgment

Chapter 6 The Independent Adviser

Chapter 7 The Economics of Advice

Chapter 8 The DREADNOUGHT Method

Chapter 9 From Fragility to Resilience

Chapter 10 The Executive Advisory Board

Chapter 11 Your First 90 Days

Chapter 12 The Missing Seat

Appendix A Boardroom Blind Spot Worksheet

Appendix B Executive Decision Worksheet

Appendix C Bounded Advisory Engagement Charter

Appendix D Recovery and Authority Exercise

Appendix E Advisory Economics Worksheet

Appendix F Worked Decision Brief

Appendix G Evidence Notes and Reading

A note to the reader

Growing companies need a reliable way to challenge consequential decisions before commitments make correction expensive. Sometimes an experienced leadership team, an effective governing board, or a trusted specialist already supplies that challenge. Sometimes nobody does. This book helps you distinguish those situations and decide what to change.

The missing seat represents a function: informed, independent challenge connected to action and follow-through. A new appointment helps only when somebody can explain the gap, establish a useful scope, and demonstrate work that improves the decision process. Buying a title supplies none of those things automatically.

I argue for a practical advisory relationship focused on AI governance, operating resilience, and executive judgment. I also argue for the right to conclude that your company needs a different solution, or no additional adviser. A book about independent advice ought to survive that conclusion.

The chapters move from the problem to the operating relationship. You will examine how assumptions turn into commitments, how recovery claims differ from demonstrated recovery, and how AI can acquire influence or permission without adequate oversight. You will then consider adviser selection, economics, a proposed DREADNOUGHT review method, and a first ninety-day engagement.

Published frameworks provide selected reference points. Their citations support the specific propositions identified in the text. They do not establish that ROWAN Prime's proposed method has undergone independent validation, that an advisory retainer causes better financial performance, or that every growing business needs to hire me. Practical exercises, templates, and the DREADNOUGHT formulation represent proposed applications, with limits stated where they matter.

Every scenario in these pages serves as a hypothetical exercise. None describes a client engagement, measured client outcome, or disguised testimonial. The financial examples use stated assumptions and arithmetic; they report no achieved return. You should replace their inputs with your own evidence before spending money.

Throughout the book, an executive advisory board means a contracted advisory arrangement that supplies recommendations within an agreed scope. It does not imply appointment to a company's governing board. Actual duties, authority, confidentiality, professional obligations, and liabilities depend on the relevant law, governing documents, agreements, and conduct. Qualified counsel should settle those questions for the particular engagement.

You can read straight through or begin with the fifteen questions in Appendix A. Treat the questions as prompts for evidence and discussion. They do not predict failure, certify compliance, or calculate a scientifically validated risk score.

The central test stays straightforward: can the people responsible for your next consequential decision explain its assumptions, identify its limits, and show what will happen when those assumptions fail?

Chapter 1 The Empty Chair

A meeting can contain every executive and still lack challenge

Imagine a founder-led field-service company preparing to accept a large contract. The sales leader presents revenue. Operations discusses staffing. Finance examines cash requirements. Technology proposes an AI assistant to reduce the administrative workload. Everyone contributes relevant knowledge. The founder hears a credible route to growth.

Now ask who owns the question that crosses all four presentations: what combination of staffing pressure, delayed customer payment, software failure, and management distraction could make the contract destructive rather than profitable?

The hypothetical company may already employ somebody capable of answering. The problem begins when nobody has the time, remit, information, or freedom to do so. Each executive can perform his own work diligently while an important relationship between the parts receives little examination.

An independent adviser should investigate that relationship. He might ask how quickly the business could replace an unavailable subcontractor, how much working capital the payment terms require, and whether the new AI tool can send customer commitments. He should then help the team distinguish a manageable dependency from an unresolved exposure. Raising concern without producing an actionable next step contributes little.

The empty chair does not signify executive incompetence. It signifies an unassigned function. A capable company can accumulate unassigned work as its customers, suppliers, technologies, and obligations change. The question concerns coverage, rather than whether the founder deserves criticism.

Find the function before choosing the person

An adviser search often starts with biography: a prestigious employer, a distinguished title, or access to influential people. Those qualities may help a specific assignment. They can also distract from the assignment that actually needs attention.

Start with the last three consequential decisions your company made. Consider a major hire, a software migration, a customer commitment, a financing decision, or a new operating location. For each, identify who tested the assumptions, who investigated unwelcome evidence, and who checked the result against the original expectation.

Look for work products rather than recollections of a spirited conversation. A written alternative, an amended contract condition, a recovery test, or a decision note can show what changed. A unanimous meeting followed by an attractive presentation cannot establish that adequate challenge occurred.

Your review may reveal strong existing coverage. Perhaps the CFO regularly challenges operating assumptions, the COO tests recovery arrangements, and the governing board reviews the strategic consequences. In that case, a recurring outside appointment may add little. A limited specialist review or a better internal agenda may solve the remaining problem.

Alternatively, you may find that all challenge depends on one person who also owns delivery. That arrangement deserves examination. A person under pressure to launch a project must sometimes evaluate evidence that would delay his own launch. Good character helps; an operating design that permits independent review helps further.

Independence requires a working definition

For this book's advisory model, independence means a disclosed capacity to examine the client's proposal without silently protecting another financial or personal interest. This working definition does not replace formal independence tests that may apply to statutory directors or particular industries.

A paid adviser still has an incentive to retain the client. A specialist who recommends a software platform may receive referral compensation. A founder's longtime friend may understand the company deeply while

hesitating to challenge him. An outside appointment does not remove those incentives. Disclosure, boundaries, scrutiny, and willingness to disagree make them more manageable.

Ask a candidate what would cause him to recommend against retaining him. Ask whether he would propose a cheaper internal solution if the evidence supported it. Ask how he discloses referral income, competing clients, and limitations in his own competence. His answers should translate into actual engagement terms.

The OECD's 2023 corporate governance principles connect board effectiveness with objective judgment and adequate information. Their discussion concerns governing boards and the relevant governance framework; it does not demonstrate the commercial value of this book's advisory offer. The narrower lesson supports examining whether challenge and information actually reach the people making decisions. [S1]

Give challenge a route into the decision

An adviser who learns about a project after the contract signature can offer lessons but may have little room to influence the commitment. An adviser invited early, with no access to the underlying assumptions, faces a different version of the same problem.

Define the decision gates at which advice matters. A technology pilot might require review before customer data enters the system, before external communication begins, and before spending expands. A customer contract might require review when payment terms, service obligations, or dependency concentrations change beyond agreed limits.

These gates should match the company's exposure. Reviewing every routine purchase would consume attention and teach managers to circumvent the process. Reviewing only projects with very large budgets would miss small systems that can access sensitive records or make public commitments. Choose gates based on consequence, authority, and reversibility as well as cost.

The decision owner should receive a concise record: the recommendation, its supporting evidence, the unresolved uncertainty, and the action requested. Record a disagreement faithfully when the owner chooses another route. Do not rewrite the review afterward to create an appearance of unanimity.

A useful first conversation

Before discussing a retainer, take one pending decision and answer five questions. What commitment will the company make? Which assumptions matter most? Who has challenged them? What would change the decision? Who will verify the outcome?

Where an answer already exists, inspect it. Where an answer does not exist, name the person responsible for producing it. That exercise gives you a concrete advisory need rather than a general feeling that successful companies ought to own an impressive boardroom.

The empty chair deserves filling when a material function lacks adequate coverage and a particular person can perform it under workable terms. The furniture carries no strategic value by itself. Your company needs the work.

Chapter 2 The Cost of Not Knowing

Uncertainty needs an owner

Executives rarely receive complete information before a decision deadline. Waiting for certainty can sacrifice an opportunity. Acting without examining uncertainty can create an expensive surprise. The useful question concerns which uncertainty deserves more investigation, which deserves a limit, and which the company can accept.

Consider a hypothetical distributor evaluating a new scheduling platform. The vendor predicts less administrative effort. Managers expect faster dispatch. Nobody has checked whether an export contains enough information to operate during an outage. The company has a benefit hypothesis, an integration plan, and a gap in its recovery knowledge.

Calling the entire project risky adds little. Calling the export untested tells the team what to do next. Ask a dispatcher to work from the export in a controlled exercise, record what he cannot accomplish, and revise the plan. The adviser contributes by converting vague concern into a question that evidence can answer.

Not every unanswered question warrants a test. Some tests cost more than the decision justifies. Others could expose personal information or interrupt service. Establish the decision importance, likely consequence, and safe means of investigation before commissioning more work.

Separate an observation from its interpretation

A report may say that customer demand increased. Its evidence may show more website inquiries during one week. Those observations could reflect seasonality, an advertising change, a competitor's interruption, duplicate submissions, or a genuine demand shift. The decision depends on which interpretation deserves confidence.

A disciplined decision note separates three statements: what the company observed, what management believes the observation means, and what management proposes to do. This separation allows a reviewer to challenge an inference without denying the underlying data.

Take a simple example. The company observed forty inquiries. Management infers enough additional demand to support a new employee. Management proposes a permanent hire. The adviser can ask about inquiry quality, conversion, gross contribution, service capacity, and the hiring alternative. None of those questions accuses management of dishonesty. They identify the reasoning between a count and a commitment.

The same discipline applies to risk. A failed recovery exercise establishes a failure under the tested conditions. It does not establish that every recovery path fails. A successful exercise establishes a result under particular conditions. It does not establish permanent readiness across every outage. Record the conditions with the result.

The cost of information and the cost of delay

More research can improve a decision or become an excuse to avoid one. Establish a stopping rule before beginning. Specify the question, the evidence that would materially change the choice, the investigation budget, and the latest useful decision time.

Suppose a supplier substitution would require a qualification test. If a modest test can expose an incompatibility before a long contract begins, the test may deserve priority. If the supplier already holds adequate verified evidence for the exact use, repeating the work may waste resources. If the test would arrive after the commitment becomes irreversible, its timing needs correction.

The value of information depends on whether the decision maker would actually act differently. A manager who requests another forecast while refusing to consider any different course has asked for reassurance, not decision support.

Ask explicitly: if the result disappoints us, will we reduce scope, delay, select another supplier, or abandon the project? If every answer leads to the same action, identify another reason for the investigation, such as mandatory assurance or operational learning. Do not disguise it as a decision test.

Use expected loss cautiously

Expected-loss arithmetic can clarify assumptions without creating knowledge. Multiply an assumed probability by an assumed consequence and you obtain an implication of those assumptions. You do not obtain a measured exposure simply because the spreadsheet contains a currency symbol.

For a purely hypothetical illustration, assume a disruption probability of ten percent over a defined year and a loss of \$200,000 if that disruption occurs. Their product equals \$20,000. If both inputs rest on weak guesses, the result rests on weak guesses as well. It cannot justify claiming a \$20,000 established loss or a guaranteed prevention opportunity.

Ask who estimated the probability, how the loss would arise, what costs the calculation omits, and whether several losses could occur together. Consider a range. If plausible estimates span orders of magnitude, report that uncertainty prominently instead of presenting the midpoint as an authoritative answer.

Some consequences deserve a different treatment. An unsafe operating condition, a contractual prohibition, or an obligation imposed by applicable law cannot automatically pass review merely because a narrow expected-value calculation looks favorable. Identify the decision constraints before optimizing financial assumptions.

Preserve what the company believed at the time

A decision journal should capture the contemporaneous reasoning. Later memories can drift toward whichever explanation fits the outcome. A good outcome can hide a weak process, while an unfortunate outcome can follow a reasonable decision under uncertainty.

Record the decision owner, the chosen action, the alternatives, the evidence, the major assumptions, and the conditions that would trigger reconsideration. Add a review date. Keep the original entry and append later results rather than silently rewriting history.

When the review arrives, ask two different questions. Did the outcome meet expectations? Did the decision process use the information reasonably available at the time? A fair review distinguishes those questions and looks for specific improvements.

If a market changed after an investment, the lesson may concern detection and adaptation. If management ignored clear contradictory evidence, the lesson may concern decision discipline. If the evidence never reached the decision maker, the lesson may concern reporting. Different causes require different repairs.

Make ignorance smaller before making confidence louder

At your next investment review, identify the three assumptions with the greatest ability to change the result. Assign each an evidence owner, a confidence description, and a practical test or monitoring trigger. Do not create a hundred-item risk register to avoid choosing the few questions that matter.

An independent adviser earns his place by helping management investigate uncertainty that can change action. He should sometimes recommend further research, sometimes a smaller experiment, and sometimes a decision now with explicit limits. His work should make the choice more intelligible, including when he cannot make the future more predictable.

Chapter 3 The Resilience Gap

A recovery promise requires operating evidence

A company can own a continuity document while lacking a demonstrated route back to an essential service. The gap lies between a claim about recovery and evidence of recovery under relevant conditions.

Imagine a hypothetical maintenance provider whose scheduling system becomes unavailable. Managers believe a spreadsheet export will sustain dispatch. During a controlled exercise, the team discovers that the spreadsheet omits access instructions, the contact list lives behind the unavailable login, and only one employee knows the manual process. The document existed. The operating path contained untested dependencies.

The exercise has not proved that the business would fail in every real incident. It has produced specific findings that somebody can repair. That distinction encourages useful testing without turning every defect into a prophecy.

For this book, resilience means the practical capacity to sustain or restore an essential service within the limits leadership has chosen, while adapting when the original plan no longer fits. That working definition covers more than information technology. Its application still requires domain expertise wherever physical safety, regulated activity, or technical engineering enters the picture.

Start with the service that matters

A list of software applications does not tell you which service deserves recovery first. Ask what the customer, employee, or other affected party needs the company to accomplish. Then identify the people, information, equipment, permissions, and external providers required to accomplish it.

Payroll illustrates the point. Restoring a payroll application helps only if the company can obtain the necessary time records, approvals, funding access, and banking arrangements. A technical recovery can succeed while the underlying business service remains unavailable.

Select a small number of essential services and map their dependencies far enough to expose critical assumptions. Give the map an owner and a date. Mark uncertain connections. Update it when operating arrangements change, rather than pretending a diagram remains current because nobody has formally retired it.

NIST Special Publication 800-34 Revision 1 addresses contingency planning for federal information systems. It connects impact analysis, recovery strategies, plans, exercises, and maintenance. That scope supports a useful reference for IT recovery work; it does not certify the broader business method proposed here or establish readiness in a particular company. [S2]

Set a target that somebody can test

Leadership should express recovery requirements in operating terms. A requirement might specify restoring a defined dispatch function within an agreed period using an approved fallback, with stated limitations on missing or outdated information. A test should then record whether the team achieved that requirement under the exercise conditions.

Distinguish the target from the observed result. If management wants restoration within four hours and the exercise takes seven, report both. Calling the target a capability hides the gap the exercise should reveal.

Also distinguish initial availability from usable service. A system that loads successfully may still contain incomplete records, stale instructions, or incorrect permissions. Define what a user must accomplish before declaring recovery complete.

Make the test safe. Use approved simulations, copies, or isolated environments where appropriate. A general adviser should not direct hazardous shutdowns or interfere with production controls beyond his competence and authorization. The technical owner should design the exercise conditions and protect ongoing operations.

Fallbacks can share the original weakness

Two providers may rely on the same communication carrier. Two administrators may depend on one recovery account. A cloud application and its backup may require the same identity service. A second supplier may operate in the same affected region or depend on the same scarce component.

The existence of a second option therefore does not establish meaningful separation. Ask what could disable both options at once. That question can reveal a concentration hidden by a list that counts vendors without examining their dependencies.

Separation also involves people. A trained deputy may still need permission, credentials, access to current instructions, and time to practice. Naming a substitute on paper cannot supply those conditions automatically.

During a tabletop exercise, ask the primary expert to observe while somebody else walks through the recovery steps. Record each point where the team needs undocumented knowledge or an unavailable approval. That controlled approach can identify gaps without pretending to reproduce every pressure of a real event.

Resilience has an opportunity cost

More redundancy consumes money, maintenance effort, and management attention. Some services can tolerate longer interruptions than others. An adviser who insists on maximum protection everywhere has avoided the hard work of prioritization.

Compare alternatives. A secondary platform may cost more than a reliable manual procedure for a low-volume activity. A manual procedure may fail when transaction volume or safety requirements exceed what people can manage. The right choice depends on the service requirement and the evidence from relevant tests.

Record the residual limitation after choosing. Leadership might accept a slower recovery because the alternative costs more than the business can justify. That decision needs a named owner and a reconsideration trigger. An accepted limitation should not disappear from reports merely because someone signed it.

The adviser's contribution

The adviser should connect the recovery question to strategic consequences. Which customer promise relies on the service? Which growth plan increases the dependency? Which acquisition would combine two operations with the same fragile point? Which executive can authorize the improvement?

Technical specialists should establish the technical findings. The adviser can help prevent those findings from becoming an unread annex. His report should explain the operating consequence, the decision required, the alternatives, and the evidence that will demonstrate correction.

Select one essential service this month. Identify its recovery owner, its most consequential dependency, and its last relevant exercise. If the evidence contains gaps, choose a safe test and a proportionate repair. Your company learns more from one honest recovery result than from a polished declaration that nobody has examined.

Chapter 4 The AI Governance Crisis

Capability can expand faster than permission

The executive problem with AI does not begin or end with whether the model gives correct answers. A system may influence a purchase, select a customer response, modify a record, or invoke another tool. Each activity raises a different question about authority, review, and consequences.

Consider a hypothetical assistant introduced to draft supplier emails. At first, an employee reviews every draft. Later, somebody connects the mailbox to save time. A subsequent workflow adds automatic follow-up. The company may have moved from text assistance to external action without a corresponding decision about scope, approval, logging, or withdrawal of permission.

This chapter's use of crisis describes that potential mismatch between capability and accountable control. It does not assert that every company faces an AI emergency or that the book has measured the prevalence of unauthorized systems. The appropriate response starts with the actual use case and its exposure.

Inventory what the system can affect

An AI inventory should explain the workflow, rather than merely listing product names. Record the business purpose, the accountable owner, the data involved, the connected tools, and the actions available. Identify who reviews outputs and who can suspend the workflow.

A note-taking tool processing public webinar material presents different questions from a system processing personnel records. An internal drafting assistant differs from an assistant that can change customer entitlements. The same underlying model can participate in both arrangements while the surrounding permissions create very different consequences.

NIST's AI Risk Management Framework 1.0 organizes AI risk work around Govern, Map, Measure, and Manage. Governance runs across the other functions, and the framework emphasizes context and continuing risk management rather than a universal checklist. The applications proposed in this chapter draw on that general orientation without claiming NIST endorsement or certification. [S3]

Begin with a workflow that already matters to the business. Ask an operator to demonstrate what enters the system, what leaves it, and which actions require human approval. Compare the demonstration with the written description. A difference between those accounts deserves investigation before expanding deployment.

A policy cannot enforce its own instructions

A policy might forbid automatic payments while the connected account still permits them. A prompt might tell an assistant to ask before sending while an integration can dispatch messages without a technical approval gate. The instruction expresses an expectation. The system design determines what can happen when the expectation fails.

For consequential actions, identify the enforcing control. It might involve a restricted service account, an application permission, a transaction limit, a review queue, or a separate approval step. Test the relevant boundary with authorized examples, including attempts that should fail.

Do not infer safety from a model's willingness to repeat the rule. A system can explain its limits accurately in one conversation and still encounter a different state, context, or integration behavior later. Review the actual action path.

The amount of control should fit the exposure. Requiring executive approval for every harmless formatting change would create an unusable process. Allowing an unreviewed workflow to make significant customer commitments because its writing sounds sensible would create a different problem. Define the consequence that triggers stronger review.

Fluent output still requires verification

NIST's Generative AI Profile identifies confabulation, data privacy, information integrity, and human-AI configuration among its risk categories. The last category includes concern about inappropriate reliance on AI.

The profile offers risk-management considerations; it does not supply a guarantee that a particular application will perform reliably. [S4]

An executive should therefore ask what checking method fits the output. A contract summary may require comparison with the actual agreement. A technical procedure may require specialist review. A customer message may require factual and authorization checks. The mere existence of a citation in generated text does not establish that the cited source exists or supports the sentence.

Separate the checking task from stylistic editing. Correct grammar cannot repair an unsupported claim. A calm tone cannot confer authority. A detailed explanation cannot substitute for evidence of the underlying system behavior.

Shadow use calls for investigation and a workable route

Employees may adopt tools outside a formal approval process because the approved route feels slow or fails to address a practical need. That possibility does not excuse uncontrolled use, nor does it establish the motives of any particular employee. Ask what work the tool performs and why the existing process failed to satisfy the need.

Provide a clear route for proposing useful applications. State what information the reviewer needs, who decides, and what lower-risk experimentation can proceed within approved limits. If every request disappears into an indefinite queue, leadership has made a policy while leaving the operating incentive unresolved.

Preserve a distinction between disclosure and punishment. An inventory effort that automatically treats every disclosure as misconduct may impair the information leadership needs. Actual violations still require appropriate handling through established procedures. The advisory task involves helping the company learn what exists without inventing an informal disciplinary system.

The boardroom question

Ask who can authorize an AI workflow, who can change its scope, and who can stop it. Then ask for the most recent evidence that those answers match the actual implementation.

You may find a small, well-bounded use that merits expansion. You may find a consequential system whose owner cannot explain its connected permissions. Those findings call for different actions. A credible adviser should distinguish them and should resist selling the same remedy to both.

AI governance becomes useful when leadership can connect a business objective to an authorized workflow, a practical limit, a review method, and a retained record. The next chapter examines the other half of the problem: even a well-controlled source of information cannot decide what your company ought to value.

Chapter 5 Intelligence Without Judgment

More information does not settle the decision

An executive can receive a detailed research brief and still lack a reason to choose. The brief may describe the market, list competitors, summarize a technology, and attach a persuasive forecast. None of those activities resolves the company's priorities or its willingness to accept a particular consequence.

Imagine a hypothetical company comparing two customers. One offers faster growth with uncertain payment timing. The other offers a smaller, steadier relationship that fits existing capacity. An analytical system can estimate scenarios, subject to its data and assumptions. Leadership must still decide how to weigh cash pressure, operating strain, strategic access, and concentration.

The analysis should make those tradeoffs explicit. If it hides them inside a single ranking, management may adopt somebody else's priorities without noticing. The adviser should ask what the recommendation optimizes and what it sacrifices.

Judgment, as used here, involves deciding how the relevant evidence, objectives, constraints, and uncertainties bear on a particular commitment. It requires accountability for the choice. Access to more material can support that work without replacing it.

Define the question before generating answers

A vague research request invites an attractive but unfocused response. Asking for an AI strategy could produce an inventory of fashionable applications. Asking which administrative activity can safely reduce a specific bottleneck within a defined budget gives the investigation a practical purpose.

Write the decision question in one sentence. Identify the owner, the deadline, the available alternatives, and the constraints. State what the research should help the owner decide. If those items remain unclear, the adviser should help clarify them before commissioning a large report.

The question also establishes the evidence standard. A preliminary exploration can tolerate uncertainty that a final investment decision cannot. A marketing brainstorm can include speculation that a published factual claim cannot. An internal draft and an external commitment deserve different approval arrangements.

When the question changes, record the change. Research designed to compare scheduling tools cannot automatically justify a wider operating transformation. A useful answer to the original question may remain insufficient for the expanded decision.

Examine the chain from source to claim

A research brief should allow a reader to identify the source behind a material claim, the date of that source, and the exact support it provides. A survey of large enterprises may offer context without establishing demand among small industrial service companies. A vendor benchmark may reveal performance under its chosen test without establishing performance in your workflow.

Ask what population, conditions, and measurement produced the result. Then ask whether those conditions resemble the decision in front of you. A weak fit does not necessarily make the source useless. It limits the inference you can reasonably draw.

Do the same with internal records. A sales count may mix renewals with new customers. A support measure may exclude unresolved tickets. A productivity estimate may count time saved by one department while omitting time added elsewhere. The adviser should investigate definitions before praising the trend.

An evidence ledger helps by separating a source's actual finding from the company's proposed application. Use a field for limitations. Preserve disagreement about relevance when reviewers differ. Do not promote a source to proof simply because its publisher carries a recognizable name.

Independent analysis needs more than separate windows

Several AI outputs can repeat the same assumption because they rely on similar source material, prompts, or underlying methods. Agreement across outputs does not by itself establish independent confirmation.

For a consequential review, diversify the basis of scrutiny. Give one reviewer the original records. Ask another to challenge the assumptions. Have a qualified specialist examine the technical issue. Where possible, test the result in a controlled operating setting. Explain the relationship among those checks rather than counting responses as votes.

A human reviewer can also reproduce another person's conclusion without meaningful examination. Independence therefore requires a task design that permits a different finding, adequate access to evidence, and a retained record of reasoning. An instruction to approve by Friday undermines that design when the actual task calls for challenge.

Use dissent constructively. A reviewer should state the disputed proposition, the evidence supporting his objection, the consequence if he proves correct, and a practical route to resolve the question. Disagreement with no identifiable claim or consequence can consume attention without improving judgment.

Match analysis to the decision's reversibility

A small experiment with limited data and no external authority can justify a lighter review than a permanent migration involving customer commitments. The difference concerns the exposure and the ability to recover, rather than whether somebody labels the technology innovative.

Before approving an experiment, define its boundary. Specify the users, information, duration, cost, and permitted actions. Decide what success and failure will look like. Identify the person who can end the experiment and the condition that requires him to do so.

After the experiment, examine the evidence against the original criteria. A successful demonstration may justify another bounded stage. It does not automatically justify a broad rollout. The next stage may introduce new users, new data, or permissions the original test never examined.

The adviser should help leadership recognize these changes in exposure. He should also avoid demanding exhaustive analysis for a harmless reversible decision. Sound judgment includes deciding how much decision process the situation deserves.

Keep the human decision legible

A final decision note should explain why leadership chose the action, what uncertainty remains, and what will cause reconsideration. If an analytical system contributed, record its role without transferring accountability to it.

The phrase the model recommended it cannot settle a question about business authority. Identify the person who accepted the recommendation and the scope of the authority he held. If nobody can do that, the organization has a governance gap regardless of how sophisticated the analysis appears.

In your next executive review, request one page containing the decision, the alternatives, the decisive evidence, the principal uncertainty, and the review trigger. Attach supporting work as needed. This structure does not guarantee a good outcome. It makes the reasoning available for scrutiny before the commitment and for learning afterward.

The missing seat supplies judgment only when the person occupying it connects information to an accountable choice. A larger briefing packet does not perform that function automatically.

Chapter 6 The Independent Adviser

Hire for the assignment you can describe

A useful adviser should fit a defined need, demonstrate relevant competence, and work within clear boundaries. Prestige can contribute to credibility, but it cannot compensate for a mismatch between the candidate's expertise and the company's actual problem.

Suppose a company needs to examine authority in an AI-enabled customer workflow. A former executive with broad commercial experience may contribute strategic questions. A technical specialist may need to investigate permissions, integrations, and system behavior. Counsel may need to examine obligations and contract language. One person may perform several roles, but the engagement should establish the relevant competence rather than assume it.

Ask a candidate to explain what he would personally deliver, what he would delegate, and what requires another professional. Ask for a sample work product with confidential material removed. Examine the reasoning, the evidence, and the treatment of uncertainty. A compelling conversation should lead to a demonstrable method.

Do not require a prospective adviser to disclose another client's information as proof. Respect for confidentiality during selection gives you evidence about how the candidate may handle your own records. A candidate can discuss methods and suitable references without exposing protected material.

Distinguish recurring advice from other services

In this book's proposed arrangement, a strategic adviser reviews recurring decisions within an agreed remit. A project consultant may investigate or implement a defined change. A fractional executive may assume specified operating responsibilities. A governing director occupies a different formal role. Actual arrangements can overlap, so the agreement and conduct matter more than the label.

The comparison should help management allocate work, rather than create a hierarchy of importance. A consultant who implements a necessary control may deliver more immediate value than an adviser who merely recommends it. A fractional operations leader may solve a capacity gap that a monthly review cannot address.

Ask which function the company lacks. Does management need a decision challenged, a problem investigated, a system implemented, or an operation led? If the answer includes several functions, separate their responsibilities, fees, evidence, and acceptance criteria.

Keep the reader instrument and the commercial service distinct. The Boardroom Blind Spot Worksheet contains the fifteen questions in Appendix A. A proposed Boardroom Blind Spot Review would examine agreed evidence about a defined decision or workflow and provide findings, limits, options, and proposed follow-up owners. That separate engagement needs agreed fees, timing, capacity, and terms before anyone promises delivery.

An adviser who also sells implementation should disclose that commercial interest before recommending the work. Independent review of a recommendation can become especially useful when the adviser stands to earn a larger contract from the proposed remedy.

Make scope precise enough to enforce

A scope should identify the decisions and services covered, the expected work products, and the available time. State what the adviser can access, whom he can contact, and which activities require separate approval.

For example, an initial engagement might review one AI-enabled workflow and one operating dependency. Deliverables could include a current-state map, an evidence register, a severity-ranked issue list, and a decision brief. Those outputs give the company something to inspect. A promise to transform resilience gives the reviewer much less to assess.

Define completion. Management should know when a report satisfies the assignment and what happens when access or information proves insufficient. A report that identifies a limitation may still complete a bounded review; it cannot certify a condition the reviewer never examined.

Include an escalation route for material findings and a process for changes to scope. Otherwise, urgency can expand the engagement informally until nobody can explain what authority the adviser actually holds.

Evaluate independence in practice

Ask for disclosure of interests that might affect the assignment. These can include vendor compensation, investment interests, overlapping clients, referral arrangements, and personal relationships. The relevant question concerns the potential effect on the work, along with a proportionate response.

Some interests can receive disclosure and management. Others may require the adviser to step out of a decision or decline the assignment. Have the appropriate decision maker document that determination. An informal assurance that everybody trusts one another does not preserve the reasoning for later review.

Consider renewal incentives too. If an adviser believes every disagreement threatens his contract, the relationship may reward reassurance. Agree at the outset that an evidence-based recommendation to stop a project or end the engagement can count as successful professional work.

Leadership carries a corresponding responsibility. Ask for challenge, listen to its basis, and respond to the actual claim. A founder who requests independent scrutiny and then punishes unwelcome findings has created a contradiction the contract cannot solve alone.

Test the relationship with a bounded assignment

A short initial assignment can reveal working habits before the company assumes a long commitment. Choose a real question with a clear owner, accessible evidence, and a practical decision date.

Observe how the candidate handles incomplete information. Does he identify limits or improvise certainty? Does he ask for evidence before choosing a remedy? Can he distinguish a critical blocker from a lower-priority improvement? Does his output help the decision owner act?

Also observe the company's readiness. Can managers provide the needed records? Do they know who decides? Will they assign an implementation owner? Even a capable adviser can accomplish little when the client lacks capacity or willingness to act on the work.

Do not treat a pilot as a guarantee of future performance. Use it as evidence for a specific next decision about scope and fit. If the assignment reveals a narrower need, narrow the relationship. If it reveals a professional need outside the adviser's competence, route that work appropriately.

Preserve a clean exit

Agree how the engagement ends, how records transfer, how access closes, and what confidentiality obligations continue. Address retention, ownership, and use of work products in terms the parties and their advisers understand.

A company should not need to retain an unsuitable adviser merely to preserve access to its own decision history. An adviser should not carry indefinite operational responsibility after his authority and access end. Counsel should translate the intended relationship into appropriate terms.

Your selection standard can remain direct: choose somebody who can investigate the defined problem, explain his limits, document his reasoning, and help you reach an actionable decision. The best candidate may recommend fewer services than you expected. That possibility belongs in a serious selection process.

Chapter 7 The Economics of Advice

Buy a decision service you can inspect

An advisory fee deserves the same scrutiny as another business expense. Identify the problem, the expected work, the total cost, the alternative, and the means of evaluation. An impressive title does not exempt the engagement from economic reasoning.

Begin with what you can verify. A review can produce an inventory, identify an unresolved permission, test a recovery assumption, or improve a decision record. Management can inspect those outputs. Claims about losses prevented or revenue caused demand more careful attribution.

If an incident does not occur, its absence alone cannot establish that the adviser prevented it. The event might not have occurred under the previous arrangement either. A credible economic account distinguishes observed work, measured operational change, and estimated financial consequence.

The distinction helps both parties. Management can hold the adviser to concrete delivery without pretending to observe an unavailable alternative history. The adviser can explain a useful contribution without claiming responsibility for every favorable event during the contract.

Calculate the full engagement cost

The fee rarely represents the entire cost. Include internal review time, implementation effort, specialist work, software changes, training, and continuing maintenance where those items apply. Consider the cost of disruption while staff learn or operate a new process.

Avoid double counting. If employee time receives a monetary value in the analysis, make clear whether it produces reduced spending, additional usable capacity, or an estimated opportunity value. Time saved does not automatically convert into cash released from the business.

Also include the alternative. Management might perform the review internally, commission a one-time specialist, use an existing board committee, or defer the project. Each alternative has its own cost and limitations. A retainer should compete with the most relevant alternative rather than with doing nothing by default.

When a paid review leads to implementation, price and evaluate the second decision separately. The company should know what the review includes before purchasing it and what further work, if any, remains outside its scope. Completing the book's reader worksheet does not purchase or initiate an advisory engagement.

A transparent break-even illustration

Assume, solely for illustration, an annual advisory fee of \$36,000. Add \$24,000 of implementation expense and \$10,000 of internal time valued as an opportunity cost. The modeled first-year economic cost totals \$70,000. None of these figures describes a quoted ROWAN Prime contract or an achieved client result.

Now assume one relevant disruption would cause \$500,000 of incremental loss. A simplified risk-only model would require a fourteen-percentage-point absolute reduction in its annual probability to offset \$70,000: \$70,000 divided by \$500,000 equals 0.14.

That calculation supplies a demanding question rather than a promise. Can the company credibly estimate such a reduction? Does the baseline probability allow a fourteen-percentage-point reduction? Which control would produce the change? How much of that change could reasonably connect to the advisory work? If the inputs lack support, the economic conclusion lacks support.

The model also omits timing, discounting, tax effects, interactions among incidents, and potential consequences outside the chosen loss estimate. Those omissions may or may not matter for a particular preliminary decision. Record them before presenting the calculation as a financial case.

Use sensitivity to expose the claim

Hold the hypothetical cost at \$70,000. If the relevant loss equals \$250,000, the required absolute probability reduction rises to twenty-eight percentage points. At \$1 million, it falls to seven percentage points. The arithmetic shows how strongly the result depends on the assumed loss.

A decision maker should see these dependencies. A single attractive estimate can hide an argument that collapses under a modest change in assumptions. Test the low-benefit and high-cost cases, not only the preferred case.

If several controls address the same event, avoid crediting each with the entire benefit. If the business already planned an improvement, distinguish the adviser's contribution to timing, design, or verification from the improvement itself. If another project changes the same measure, report that confounding influence.

Use financial analysis proportionately. An exploratory assignment may justify a simple cost cap and concrete deliverables. A major recurring commitment may require a more detailed evaluation. The absence of a perfect causal model does not forbid a purchase; it limits the claims the buyer and seller can honestly make.

Evaluate the relationship with several kinds of evidence

Track delivery against the agreed scope. Did the adviser complete the review? Did his findings identify the relevant evidence? Did the decision owner receive a usable recommendation on time?

Track management action separately. Which recommendations did the company accept? Which did it reject, defer, or supersede? Who owned implementation? An adviser cannot fairly claim success for an unimplemented recommendation, while management should not confuse its own implementation delay with a failure to deliver the agreed analysis.

Finally, track operational results where measurement makes sense. Examples include an observed recovery time under stated test conditions, the number of unresolved privileged accounts, or the completion rate for a defined review process. These measures can show change without automatically proving broad financial causation.

Agree on the baseline and measurement method before the intervention where possible. Preserve unfavorable results. A review process that records only successes cannot help either party learn whether the engagement deserves renewal.

Make renewal a fresh decision

At renewal, ask what material gap remains and why this particular relationship offers a reasonable way to address it. Past effort does not automatically justify future fees. A capable engagement can succeed by reducing the need for recurring outside help.

Possible outcomes include renewal, narrower scope, a periodic review, referral to a specialist, or closure. Write the reasoning. If the adviser argues that every outcome proves the need for more advisory work, challenge the commercial logic.

Advice earns economic credibility when its seller can explain the work, its buyer can inspect delivery, and both can discuss uncertain benefits without manufacturing precision. That standard may sound less dramatic than a promise of extraordinary ROI. It gives leadership a purchase decision it can defend.

Chapter 8 The DREADNOUGHT Method

A proposed discipline for bounded execution

DREADNOUGHT, as developed for this book, organizes a review around five questions: who authorized the action, what bounds it, how people observe it, how an authorized person intervenes, and what evidence survives. This formulation represents a proposed ROWAN Prime method. It does not claim independent validation, a professional certification, or a deployed software capability.

The method starts with a consequential workflow. It can support discussion of AI-enabled processes, customer communications, or other defined operations. Technical implementation and regulated professional work still require appropriately qualified owners. A common review vocabulary cannot erase those requirements.

Apply the method to one workflow before attempting an enterprise-wide declaration. A narrow review produces a clearer account of what the reviewer actually examined. If the scope later expands, preserve that change and revisit the controls that the wider exposure requires.

Established control literature offers technical reference points. NIST Special Publication 800-53 Revision 5 includes least privilege in AC-6, protection of audit information in AU-9, and contingency plan testing in CP-4. These controls address different concerns; they do not constitute the DREADNOUGHT method or validate it. [S6]

A technical owner should select the controls and assessment procedures appropriate to the actual system and applicable requirements. The five questions below organize executive scrutiny. They do not replace a control assessment, establish compliance, or prove that any product implements the controls.

Authorized action

Write the objective and the source of authority. Identify the person or body entitled to grant permission for the specific action. Record what the permission covers and when it ends or requires reconsideration.

A direction to prepare a customer response does not automatically authorize sending it. Approval to test a workflow with artificial records does not automatically authorize using live personal information. Permission to perform a task once does not automatically establish recurring authority.

The review should inspect how the organization represents these distinctions. A task record may preserve the instruction. The application may enforce the permission. A person may approve the final transaction. Check the relationship among those parts instead of assuming that one substitutes for the others.

If the authority remains unclear, identify the responsible decision maker and hold the consequential step until the question receives an appropriate resolution. Continue independent preparation where it falls within the existing scope. This approach prevents an uncertain boundary from unnecessarily stopping unrelated work.

Bounded execution

Specify the operational limits: permitted data, recipients, accounts, spending, duration, transaction size, and tools as applicable. Choose limits that fit the workflow rather than copying a universal list mechanically.

Then identify how each material limit works. A budget line in a document differs from a spending cap enforced by the payment path. A recipient list stored beside a script differs from an allowlist that the dispatch process actually checks. Both documentation and enforcement can matter, but they perform different jobs.

Test an allowed action and a prohibited action under controlled conditions. The allowed action should succeed in the intended way. The prohibited action should stop or route to appropriate review. Record the expected and observed behavior, along with any conditions that limit the result.

An incomplete test deserves an incomplete finding. Do not describe an untested boundary as verified because the design looks reasonable. Conversely, do not describe a design gap as proof that an unauthorized action already occurred.

Continuous observation

Decide what the operator and accountable owner need to see while the workflow runs. Useful signals can include execution failures, unusual volume, changes to permissions, unresolved approvals, or missing evidence. The relevant set depends on the consequence you need to detect.

Continuous observation does not necessarily mean a person watches a screen every second. It means the observation arrangement matches the workflow's timing and exposure. A daily report may suit one activity while leaving an unacceptable gap in another.

Assign responsibility for responding. A dashboard with no owner can display a deteriorating condition without changing it. A notification that nobody receives cannot supply timely intervention. Test the delivery route and the response process, rather than relying solely on the existence of a rule.

Enforceable intervention

Identify who can pause, restrict, or terminate the relevant action. Explain what happens to work already queued or in progress. Distinguish stopping future action from reversing an action that has already affected the outside world.

Revoking a credential may stop one route while leaving another integration active. Pausing a scheduler may leave a dispatch queue intact. A complete intervention test should follow the actual action path far enough to check the intended boundary, under the technical owner's approved procedure.

Plan a safe response to uncertainty. If the system cannot determine whether a consequential authorization has expired, the design may need to hold that action for review. Avoid using that rule to halt unrelated low-impact work whose authority remains clear.

Durable evidence

Retain the instruction, relevant approvals, execution records, exceptions, and test results in an appropriate form. The records should allow an authorized reviewer to reconstruct the material event without exposing information beyond his need to know.

Durability involves more than saving a screenshot. Consider access, retention, integrity, timestamps, version identification, and the connection between the record and the action. A copied summary may omit the exact limitation that mattered.

Proportion matters here too. Logging unnecessary sensitive content can create its own exposure. Determine what evidence the company needs and handle it under an appropriate data and retention policy. Regulated or particularly sensitive records may require specialist review.

From review to disposition

For each of the five questions, record an observation, its evidence, the unresolved limitation, and an owner. Give material findings a disposition: proceed within the tested scope, proceed with explicit conditions, hold a consequential step, or obtain specialist review. These dispositions express the review's operational judgment; they do not confer formal assurance beyond its scope.

Do not average a missing stop mechanism into an acceptable overall score merely because the other controls look strong. A consequential single weakness can require attention on its own terms. Appendix A therefore uses question-level evidence states and decision flags rather than a universal numerical grade.

The method succeeds only if the work connects to an accountable decision. An attractive five-column matrix can still conceal uncertainty. Ask what the company examined, what it learned, and which action changed as a result.

Chapter 9 From Fragility to Resilience

Choose a repair the company can actually complete

A review can generate more recommendations than a company can implement. Without priorities, the result may add a backlog while leaving the most consequential weakness untouched. The adviser should help management select a practical sequence.

Start with a finding that connects to an essential service or a consequential action. Explain the exposure in operating terms. Identify what management can change, what the change will cost, and what evidence will demonstrate completion. Avoid describing every departure from an ideal design as an emergency.

Consider a hypothetical company whose customer communication workflow can dispatch messages from a shared account. Nobody has a complete record of who can change the automation. A broad transformation could involve new software, account redesign, training, and contract changes. An immediate authorized restriction on dispatch permissions may address a material exposure while the technical team evaluates longer-term changes.

The immediate restriction still needs appropriate review. It might interrupt a legitimate customer service or leave queued work in an uncertain state. A bounded repair should account for the operating consequence of the repair itself.

Connect findings to a decision owner

A recommendation addressed to the company can disappear because no individual holds the task. Assign a person who can make the relevant decision or obtain the necessary authorization. Give implementation to a person with suitable competence and capacity. Those roles may belong to different people.

The adviser can recommend priorities and clarify dependencies. He should not silently become the operating owner when the agreement provides only advice. If management wants him to assume implementation responsibility, establish that scope separately and assess competence, authority, access, and cost.

For each accepted recommendation, record the action, owner, target date, completion evidence, and review point. Add a clear state such as proposed, accepted, in progress, awaiting verification, or closed. A task marked complete by its implementer may still need verification against the agreed acceptance criteria.

Rejected and deferred recommendations also deserve a record. Explain the reason and any reconsideration trigger. The absence of action could reflect an informed tradeoff, an unresolved disagreement, a lack of funding, or a forgotten task. Those conditions require different follow-up.

Keep material limitations visible

An executive dashboard should distinguish actual results from targets and plans. A tested recovery time belongs beside its test conditions. An untested target should retain an untested label. A restriction approved for implementation should not appear as an enforced restriction before implementation and verification occur.

NIST's Cybersecurity Framework 2.0 includes a Govern function alongside Identify, Protect, Detect, Respond, and Recover. Its outcome-based structure connects cybersecurity decisions with organizational context, responsibilities, and oversight. This book uses that reference point narrowly; the broader change sequence here represents an original proposed application, rather than a claim of CSF conformity. [S5]

A useful reporting discipline asks what changed since the prior review and what evidence supports the change. If a finding remains open, explain the current obstacle and decision required. Repeating the original warning without updating its state can turn a report into background noise.

Do not reduce serious exceptions to a reassuring percentage. A company could complete many routine tasks while leaving one consequential authority gap unresolved. Report that gap directly. Counts can support management, but their presentation should preserve consequence.

Improve the work without increasing hidden burden

A new control can transfer effort to another part of the business. A mandatory review may reduce inappropriate dispatch while overwhelming the reviewer. A complex approval form may consume time without improving the underlying decision. Examine the operating effect after implementation.

Ask the people who use the process where it creates confusion, delay, duplication, or opportunities for workarounds. Their feedback does not automatically invalidate the control. It helps management determine whether the design fits the real work.

Measure a limited set of outcomes that relate to the control's purpose. For an approval queue, examine whether consequential actions receive the intended review, how long that review takes, and what happens when the reviewer cannot respond. For recovery, examine whether users can complete the defined service within the chosen target under the test conditions.

If the control creates an unacceptable bottleneck, change the design through the appropriate approval route. Do not quietly disable it and preserve a report claiming protection. The improvement process should update the actual system and the record describing it.

Plan for change in people and suppliers

A control that depends entirely on one experienced employee may work until that person changes roles or becomes unavailable. Prepare a practical transfer of knowledge and authority. The substitute needs access to current instructions, a permitted means of action, and an opportunity to demonstrate the task.

Supplier change also deserves review. An application update, contract renewal, ownership change, or integration modification can affect assumptions that previously held. Define which changes trigger re-examination. Avoid claiming that an annual review alone covers every material change between review dates.

The adviser should help connect these triggers to the normal operating process. Purchasing, account administration, project management, and personnel transitions can provide natural review points. A separate governance program that never touches those activities may miss the events that matter.

Retain the prior assessment when a change occurs. It helps reviewers identify what changed and which conclusions may no longer apply. Record the new scope and evidence rather than declaring the entire historical assessment worthless or carrying it forward unchanged.

Practice a decision under pressure

A tabletop exercise can test who notices a problem, who decides, and how the team communicates. Begin with a clearly marked hypothetical event. Keep production systems protected. Give participants the information they would plausibly possess at each stage, then observe how they use their authority.

You might simulate an unavailable supplier during a busy service period. Ask who confirms the disruption, who can authorize a substitute, and who communicates changed commitments to customers. Introduce an uncertain piece of information and observe whether the team labels it appropriately or turns it into fact.

Close with specific findings. Which contact failed? Which authority remained unclear? Which instruction required revision? Which action worked as intended? Assign the repairs and a proportionate retest. A dramatic exercise with no follow-through contributes little operating value.

Define progress without pretending to finish uncertainty

A company cannot remove every dependency or anticipate every event. Progress means addressing selected exposures, improving the evidence supporting management's confidence, and preserving a route to adapt when conditions change.

At the end of an improvement cycle, the adviser should explain what now works under the examined conditions, what remains unverified, and which decision comes next. That account gives leadership a usable picture of the operating position. It also keeps the relationship honest when the work reveals more uncertainty than the original plan anticipated.

Chapter 10 The Executive Advisory Board

A group needs an operating purpose

Adding several advisers can broaden expertise or create another meeting with unclear ownership. The difference depends on the work the group performs and the route from its discussion to management action.

Define the purpose before inviting members. A company might need recurring scrutiny of AI-enabled operations, an external view of growth dependencies, and structured challenge of major investments. Another company might need only a technical specialist for a single project. The second company should not assemble a standing advisory board to solve a temporary need.

In this book's proposed model, the executive advisory board provides recommendations within a contractual remit. The company's authorized leadership retains decision responsibility unless a separate, valid arrangement assigns a particular function differently. Counsel should examine the actual structure and conduct. Neither the name advisory nor a disclaimer alone settles legal obligations.

Write the relationship to the governing board explicitly where one exists. Avoid creating a parallel channel that confuses formal oversight, bypasses required approval, or leaves directors unaware of material issues. The objective should support the company's legitimate governance arrangements.

Select complementary competence

Begin with the decisions likely to arise over the next year. Identify the expertise required to evaluate them. An industrial service business might need operating knowledge, financial discipline, and technical scrutiny of a new system. It may already possess some of that competence internally.

Choose advisers for the missing work. A room filled with similar biographies may produce less useful coverage than a smaller group with genuinely complementary knowledge. Shared industry experience can help participants understand context, while different professional experience can expose assumptions that insiders overlook.

Ask each candidate to identify his competence limits. A specialist should not receive authority over unrelated issues simply because he speaks confidently. A generalist should know when to obtain a specialist rather than expanding his title to fit the question.

Establish an arrangement for occasional expertise. The company may need a lawyer, security engineer, accountant, or safety specialist for a particular matter without requiring every professional to attend every recurring meeting. Scope and confidentiality should cover those contributions appropriately.

Give the chair a practical job

The advisory chair, if the arrangement uses one, should organize the work, preserve useful disagreement, and make the output legible to the decision owner. The role should not confer implied governing authority that the company never intended to grant.

Before a meeting, identify the decisions or questions requiring attention. Send concise materials early enough for review. Ask members to raise important information gaps before the meeting rather than discovering halfway through that nobody can assess the proposal.

During discussion, distinguish a request for information from a recommendation and a recommendation from an authorized decision. Record who owns the next step. Afterward, issue a short decision-support note rather than relying solely on a lengthy transcript.

The chair should also identify when the group lacks a basis to advise. A meeting can properly conclude that the company needs further evidence or specialist review. Pressuring members into a confident consensus can erase the very function the company hired the group to provide.

Build a meeting around one consequential question

An illustrative agenda could begin with the open actions that affect today's question, then review the evidence, examine alternatives, and produce a recommendation with conditions. Reserve time to record dissent and assign follow-up. The meeting duration should follow the complexity of the question and the participants' preparation.

Avoid using the entire session for status presentations that members could read beforehand. Equally, do not force every topic into a compressed format when an important technical finding requires explanation. The chair should protect decision time without making speed the only measure of effectiveness.

For a hypothetical expansion review, the group might examine the cash requirement, the capacity constraint, and the dependency most likely to disrupt delivery. Its output could recommend a bounded pilot, identify unresolved conditions, and name the evidence required before scaling. The output should state which matters the group did not examine.

When a member disagrees, preserve the substance. Record the disputed assumption and the consequence he expects if it fails. Leadership can then decide whether to obtain more evidence, modify the proposal, or accept the uncertainty. A dissent record serves the decision process better than a cosmetic consensus.

Protect information and handle conflicts

Give advisers access appropriate to their assignment. More access does not automatically produce better advice, and unnecessary access increases the burden of protecting information. Establish what members can retain, share, or use in other work.

Set a procedure for disclosures and recusals. A member who holds a relevant supplier interest should disclose it before the recommendation, rather than after procurement. Decide who evaluates the conflict and who records the response. Include new interests that arise during the engagement.

If the company uses AI to summarize discussions, obtain the necessary approval for the tool, the data, and the processing arrangement. Do not assume that permission to attend a meeting automatically authorizes sending its contents to an outside service. Retain the authoritative record and check generated summaries against it.

Confidentiality should not become a pretext for concealing a material limitation from the decision owner. Establish a legitimate route for escalation that respects applicable obligations. Counsel should help resolve situations involving competing duties or legally sensitive information.

Evaluate the board as a service

At a defined review point, inspect the work. Which consequential decisions received useful challenge? Which evidence gaps closed? Which recommendations management accepted or rejected? Did the group identify matters outside its competence and obtain appropriate help?

Ask management about burden as well as benefit. Did meetings consume attention without improving decisions? Did members deliver on time? Did their recommendations account for operating constraints? A technically correct recommendation can still fail to help if it ignores the conditions required for implementation.

Possible responses include changing membership, narrowing the remit, reducing meeting frequency, or ending the arrangement. A board should earn continuation through useful work. Its members should welcome a fair review that permits that outcome.

The executive advisory board offers a structure for recurring challenge. Whether that structure fits your company depends on the gap, the people, and the operating agreement. A charter supplies a starting point; actual conduct determines whether the group performs its intended function.

Chapter 11 Your First 90 Days

Begin with a question that matters

The first ninety days should establish whether the advisory relationship can produce useful work under real conditions. It should not attempt to solve every governance problem or build an elaborate reporting system before understanding the business.

Choose one consequential workflow and one essential service. Identify the management sponsor, the decision owners, the relevant technical or professional owners, and the evidence available. State the intended outputs and the boundaries of the review.

Set a realistic capacity budget. Managers still need to operate the company while supporting the engagement. An ambitious plan that requires interviews, tests, and documentation beyond available time will generate delay or superficial compliance. Agree on priorities and explicitly defer lower-value work.

The schedule below offers a proposed sequence rather than a promised completion standard. Adapt it to the company's size, exposure, obligations, and the time required for qualified review. A serious finding may change the sequence immediately.

Days 1 through 15 establish scope and baseline

Confirm the engagement terms and the authority of each participant. Set access limits and approved communication routes. Record conflicts, professional boundaries, escalation arrangements, and the intended handoff at the end of the period.

Collect a focused baseline. For the selected workflow, document its purpose, inputs, outputs, permissions, approvals, dependencies, and retained evidence. For the selected service, document its minimum operating requirement, recovery owner, fallback, and last relevant test.

Interview the people who actually perform the work. Compare their accounts with policies and technical demonstrations. Where descriptions conflict, preserve the disagreement and identify how to resolve it. Do not choose whichever account sounds most confident.

End this period with a short scope-and-baseline note. It should say what the adviser examined, which material information remains missing, and which question deserves priority. Management should confirm that the review still addresses a useful business need.

Days 16 through 30 test the critical assumptions

Select a small number of tests that could change a consequential decision. The technical or operating owner should approve their design and safeguards. Define expected results before execution. Protect production activity and sensitive information.

For the workflow, a test might examine whether a prohibited action reaches the review queue. For the service, an exercise might examine whether a substitute operator can complete the approved fallback. Neither example justifies performing disruptive or unauthorized actions.

Record the conditions, steps, actual result, and limitations. Classify a test that could not run as incomplete. If the test reveals an unexpected exposure, use the agreed escalation route and document any authorized containment.

By the end of day thirty, leadership should receive a decision brief. It should identify the priority findings, their supporting evidence, the practical alternatives, the proposed implementation owners, and the decisions required. Avoid turning the report into a catalogue of every improvement imaginable.

Days 31 through 60 implement selected changes

Management chooses the changes and assigns accountable owners. The adviser supports clarification and review within his scope. Implementation responsibility should remain explicit, especially when another provider performs the technical work.

Keep a change record. Identify the approved design, the affected workflow, the implementation steps, and the acceptance criteria. Plan how the team will respond if the change produces unacceptable disruption. Record who can authorize reversal or further modification.

Verify the changes against the agreed criteria. An administrator's assurance that he changed a setting may supply part of the evidence. A controlled test of the resulting behavior supplies another part. Match the evidence to the actual claim before closing the finding.

Review operating burden with users. If the control creates confusion or delay, investigate and adjust through the appropriate process. Preserve the protection the company needs while making the work practical. Neither a cumbersome process nor an undocumented workaround deserves automatic acceptance.

At day sixty, compare the observed position with the baseline. State which findings closed, which remain open, and which new facts changed the priority. Distinguish management's chosen residual limitations from issues that simply lost attention.

Days 61 through 90 exercise and decide

Run a proportionate exercise of the revised arrangement. Confirm that the substitute operator, review queue, escalation contact, or stop mechanism works under the selected conditions. Where the company claims improvement, show the comparable evidence and its limitations.

Hold an executive review focused on the next decision. The adviser should explain what the company can now demonstrate, what it still cannot demonstrate, and what ongoing review would accomplish. Management should assess both the outputs and the cost of obtaining them.

Decide whether to renew, narrow, change, or close the engagement. If a recurring role continues, write the next scope around the remaining need. Avoid extending the original assignment automatically because the calendar reached day ninety.

Complete the handoff. Transfer work products under the agreed terms, identify continuing owners, close unnecessary access, and preserve the decision record. If unresolved findings transfer to management, name the owner and review trigger. An engagement should not leave a vague promise of continuing oversight after its authority ends.

Conditions that justify changing the plan

Pause the consequential part of the assignment when authority remains unresolved, required evidence cannot lawfully or safely reach the reviewer, or a material finding requires specialist attention before further action. Continue other authorized work where it remains useful and independent of that boundary.

Reduce scope when the company cannot supply the necessary time or information. Increase technical or professional involvement when the question exceeds the original competence. Reconsider the engagement when management wants endorsement but refuses any meaningful examination.

The adviser should also report when the suspected gap does not appear within the examined scope. A finding that existing arrangements adequately cover the selected question can justify closure or a different assignment. The process should permit that result without treating it as a commercial failure.

What success can honestly mean

After ninety days, the company may possess a clearer account of authority, a tested recovery path, improved decision records, or a set of findings assigned to accountable owners. Describe those results specifically.

Do not call the company resilient in every respect because one exercise succeeded. Do not claim that the adviser prevented an event that never occurred. Do not report implementation when management merely approved a plan.

Success means delivering the agreed work, making the evidence accessible, and helping leadership decide what follows. The scope and the outcome may change as the review produces knowledge. The obligation to report accurately should remain.

Chapter 12 The Missing Seat

Return to the decision in front of you

At the beginning of this book, a hypothetical company prepared to accept a demanding new contract. Its executives understood their own contributions. The unanswered question concerned how those contributions interacted under pressure.

You now have a way to examine that question. Identify the commitment. Distinguish evidence from interpretation. Map the consequential dependencies. Ask who holds authority and how the company can intervene. Consider the expertise needed to challenge the proposal. Evaluate the cost of obtaining that expertise against relevant alternatives.

None of those steps requires an elaborate advisory institution. They require a decision owner willing to make assumptions visible and examine evidence that may change the preferred course.

An adviser can help when he supplies competence, independence, and capacity that the company otherwise lacks. He can also fail when his incentives, access, or methods undermine the work. The business should retain the right to test both possibilities.

Five reasons to consider a bounded engagement

First, management may face a consequential decision outside the team's current competence. A defined outside review can help clarify the issue and identify where specialist work belongs. The adviser should demonstrate relevant capability rather than relying on a general claim to wisdom.

Second, the team may possess the competence but lack the time to investigate a cross-functional exposure. An adviser can supply focused capacity, provided the company still assigns a decision owner and implementation resources.

Third, the preferred project may have strong internal sponsorship and weak independent challenge. A reviewer with access to evidence and permission to dissent can examine the assumptions before commitment. His value depends on actual scrutiny, rather than the ceremonial presence of an outsider.

Fourth, the company may have repeated the same unresolved finding across several reviews. An adviser can help clarify the obstacle, connect it to an accountable decision, and test completion. He cannot substitute for leadership's willingness to assign responsibility and accept the cost of repair.

Fifth, growth may have changed the operating exposure. New customers, new permissions, new suppliers, or a new location can invalidate assumptions that previously held. A bounded review can examine the change without treating every past practice as defective.

Each reason still requires a fit assessment. The same problem may call for an internal change, a specialist, a consultant, or a fractional executive. The book's argument supports investigating the gap; it does not entitle any particular adviser to the contract.

Reasons to wait or choose another route

Do not hire a strategic adviser merely to display a distinguished name. Unless the role includes actual work, the company may purchase appearance while leaving the original gap untouched.

Do not use an advisory appointment to avoid employing a necessary operating leader. If the company needs daily execution and accountable management, a periodic review cannot supply that capacity by implication.

Do not expect a generalist to replace licensed or technically qualified professionals. An adviser should help identify where those professionals belong and coordinate decisions within his remit. He should not absorb their authority through ambiguous marketing.

Wait when management cannot provide the information or attention the assignment needs. A narrower preliminary task may help establish readiness. An open-ended retainer will not manufacture client capacity simply because the contract begins.

Choose another route when existing arrangements already provide adequate coverage. A periodic independent check might suit the need better than a standing relationship. A company that can demonstrate effective challenge deserves acknowledgment of that evidence.

The standard applies to the adviser too

A book advocating evidence should expose its own limits. The proposed DREADNOUGHT method needs field testing, critique, and revision. The fifteen-question worksheet supports structured discussion; it lacks demonstrated predictive validity. The economics depend on assumptions that the buyer must examine. These limits affect what anyone should promise on the strength of the book.

An adviser should welcome scrutiny of his own claims. Ask him which parts of his method derive from published guidance, which represent professional judgment, and which remain hypotheses. Ask how he records corrections and whether he will preserve dissent from another reviewer.

Look for a willingness to say that the evidence does not support a conclusion. That answer should lead to a practical next step where one exists. Permanent uncertainty with no decision route can become as unhelpful as false certainty.

The standard also applies to public positioning. Actual experience, qualifications, product capabilities, and client outcomes should receive accurate descriptions. A developing practice can offer disciplined work without inventing a history it has not earned.

Your next practical step

Select one decision that will matter within the next ninety days. Use Appendix B to write the question and its conditions. Use Appendix A to identify the evidence gaps most relevant to that question. Do not attempt to complete every worksheet before deciding whether the problem deserves attention.

Meet with the people who own the decision and the affected operation. Ask whether they can supply the needed challenge internally. If they can, assign the work and a review date. If they cannot, write a bounded external assignment using Appendix C.

Evaluate a candidate against that assignment. Ask for a method, a work product, a clear statement of limits, and a transparent account of fees and interests. Choose a review point that permits closure as well as continuation.

The decision can end with an adviser, another professional, an internal improvement, or no additional appointment. What matters concerns whether the company's consequential commitments receive adequate challenge and follow-through.

Fill the function

The missing seat has never depended on a particular chair, title, or software package. It depends on whether somebody can ask the necessary question, investigate its answer, and connect the result to an authorized decision.

When that function works, leadership gains an opportunity to change course before a commitment hardens. It also gains a better record of why it chose to proceed when uncertainty remained. Neither benefit guarantees a favorable outcome. Both can make the organization more deliberate about the risks it accepts.

The responsibility stays with the people authorized to lead. Independent advice should strengthen their ability to exercise it. Begin with the next decision. Require evidence. Preserve the right to disagree. Then choose the person or process that can do the missing work.

Appendix A Boardroom Blind Spot Worksheet

Use the questions to find evidence

This fifteen-question worksheet supports an executive conversation about a defined operating scope. It offers an original, provisional discussion instrument. It has not undergone predictive validation, and its categories do not establish a company-wide risk rating. A reviewer should adapt the questions to the company and obtain specialist help where needed.

Choose a scope before answering: a business unit, a proposed investment, a selected workflow, or another clearly defined area. Record the worksheet date, the accountable sponsor, and the people contributing evidence. A conclusion about one workflow should retain that boundary.

For each question, retain the answer, supporting records, date, reviewer, limitation, and follow-up owner. Use one of four evidence states: supported within scope; partially supported; unknown; or outside scope with a reason. Supported means the reviewed evidence adequately addresses the defined question under the stated conditions. It does not imply permanent effectiveness.

Keep decision flags separate from evidence states. A demonstrated defect may require urgent action even though the evidence about it looks strong. An unknown may require investigation without establishing that a defect exists. Possible flags include consequential action held, specialist review needed, improvement planned, and no additional action identified within scope.

Do not add the states into a universal numerical score. A missing stop mechanism cannot safely disappear inside an average. Do not interpret outside scope as a pass. If exclusions remove information necessary for the decision, record the resulting limitation and change the conclusion accordingly.

Use the response lines after each question to identify the actual record or test, rather than merely writing yes. Select one evidence state and add a decision flag when needed. An unknown answer calls for an appropriately scoped investigation; it does not establish that the control failed. A supported answer about a demonstrated defect can still require action.

Governance and authority

Question 1 Who can authorize the consequential decision

Identify the actual decision owner and the source of his authority for the chosen action. Ask how limits and delegations reach the people performing the work. Request the relevant approval record or governing instruction, while handling sensitive material appropriately.

Watch for a difference between the person who sponsors the project and the person entitled to approve it. A sponsor may organize a proposal without holding authority over spending, external commitments, or restricted information. Where that distinction affects the decision, resolve it explicitly.

Follow-up prompt: name the next commitment, identify who can authorize it, and record what evidence will demonstrate the approval before execution.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 2 Which decisions receive independent challenge

Choose a recent consequential decision and identify who examined a material assumption or alternative.

Request a decision note, review comments, or other evidence of substantive challenge. Attendance at a meeting alone does not demonstrate the quality of that examination.

Consider whether the reviewer could reach an unwelcome conclusion without silently protecting an undisclosed interest. If existing internal or governing-board arrangements provide adequate challenge, record that finding. The worksheet should permit the conclusion that no additional adviser appears necessary within the examined scope.

Follow-up prompt: identify the specific decision that lacks adequate scrutiny before proposing a new recurring appointment.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 3 What ends or changes delegated authority

Ask how the company handles expiration, revocation, changes in scope, and completion. Request evidence from the relevant operational path, including queued or in-progress actions where applicable. A reminder that permission ends does not automatically enforce the end.

Where safe and authorized, ask the technical owner to demonstrate an appropriate boundary test. Record the test conditions and limitations. If no test can run, preserve an unverified state rather than claiming either successful enforcement or an actual breach.

Follow-up prompt: establish who owns the authority boundary and what happens when the workflow cannot determine whether permission remains valid.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

AI exposure

Question 4 Which AI workflows use company information

Request a workflow inventory that identifies purpose, owner, information, connected tools, permitted actions, and review arrangements. Start with the chosen scope and distinguish an incomplete inventory from a demonstrated absence of AI use.

Ask operators how they perform the work. Reconcile their accounts with available administrative records and approved technical inspection. Use an appropriate process that respects privacy and does not treat a speculative suspicion as a finding about an individual.

Follow-up prompt: name the material information gap, the appropriate investigator, and the evidence needed to close it.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 5 Which outputs can affect the outside world

Identify whether the system can send, publish, purchase, alter records, or otherwise create an external consequence. Request the applicable permissions and approval arrangement. Distinguish drafting capability from dispatch authority.

Consider the full integration path. A user interface may show a review step while another route permits direct action. A workflow description may omit a recently added connector. Have the qualified owner examine those possibilities within authorized limits.

Follow-up prompt: state the consequential action, the enforcing control, and the latest relevant test of that control.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 6 How does the company check material AI output

Select an actual output relevant to the decision and examine the checking method. Does a reviewer compare factual claims with their sources? Does a qualified person inspect technical or professional content? Does the approval record identify the material version?

Look for a review designed around the output's consequences. A spelling check cannot verify a contractual interpretation. A model-generated confidence statement cannot establish factual support. Review capacity also matters when output volume grows.

Follow-up prompt: choose one output class and define the evidence and competence its approval requires.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Operating resilience**Question 7 Which service must continue or recover first**

Identify the essential business outcome, the accountable owner, the minimum acceptable service, and the consequences of interruption. Request management's priority decision and the basis for it. Do not substitute a list of applications for an understanding of the service.

Compare the service requirement with customer commitments and operating realities. If management has never settled the requirement, ask the appropriate owner to do so before the technical team claims to satisfy it.

Follow-up prompt: write a service-level recovery requirement that an authorized exercise could examine.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 8 What could disable both the primary and fallback route

Map material shared dependencies, including people, identity services, communications, suppliers, facilities, and information. The worksheet should go far enough to investigate the chosen service without pretending to reveal every possible dependency.

Request evidence of meaningful separation where the company relies on redundancy. Two branded services may still depend on a common underlying route. A second operator may still require the unavailable primary expert's credentials or approval.

Follow-up prompt: identify one shared dependency whose failure would undermine the current fallback and choose an appropriate next investigation.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 9 What did the latest recovery exercise demonstrate

Request the date, scope, conditions, participants, expected result, observed result, and unresolved findings. Compare the demonstration with the claim leadership currently makes. A successful narrow test should retain its narrow conclusion.

If no relevant exercise exists, record the absence of evidence and its consequence for confidence. Do not assert that recovery cannot work solely because nobody has tested it. Ask the qualified owner to design a safe, proportionate exercise.

Follow-up prompt: assign an owner, an approved test plan, and an acceptance condition for the selected service.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Executive judgment

Question 10 Which assumptions could change the decision

Ask management to identify the small number of assumptions with the greatest effect on the pending choice. Request their sources, dates, confidence descriptions, and relevant alternatives. A long undifferentiated list may hide the assumptions that matter most.

Identify what result would lead to a different action. If management would proceed regardless, clarify whether the requested analysis serves another purpose. The evidence standard should follow the real decision rather than a ceremonial research exercise.

Follow-up prompt: select one decisive assumption and a practical means of testing or monitoring it.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 11 What would trigger reconsideration

Request the conditions that would cause leadership to reduce scope, pause, change course, or seek additional approval. Identify who observes the condition and who acts. A trigger without a responsible person may remain an unused sentence.

Make the trigger specific enough for the relevant owner to recognize. Avoid false precision where the available data cannot support it. A clearly described event can prove more useful than an unsupported numeric threshold.

Follow-up prompt: connect one reconsideration condition to a named observer, decision owner, and communication route.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 12 Can the company reconstruct its reasoning

Choose a recent decision and request the contemporaneous record of alternatives, evidence, assumptions, authority, and review date. Compare it with later results without rewriting the original entry. Distinguish a weak process from an unfavorable outcome under reasonable uncertainty.

If only recollections remain, record that limitation. A retrospective account can help learning but cannot substitute perfectly for the record available at the time. Begin a proportionate decision journal for future commitments.

Follow-up prompt: adopt a short decision note for the next material commitment and identify where the authoritative version will reside.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Organizational intelligence

Question 13 How do unwelcome operating facts reach leadership

Ask an operator and a decision owner to explain the same escalation path. Request a suitable example or approved exercise showing how a concern travels, who receives it, and what happens next. Handle sensitive reports through the company's appropriate procedures.

Investigate whether the path encourages accurate reporting or rewards the appearance of smooth delivery. Avoid inferring a cultural diagnosis from one anecdote. Look for specific operating features that leadership can change.

Follow-up prompt: identify a material signal that currently lacks a reliable recipient or response owner.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 14 How does a material change reopen a review

Request the events that trigger re-examination of the selected workflow or service. Consider changed permissions, new suppliers, altered information, revised customer promises, and personnel transitions. The relevant triggers depend on the actual scope.

Examine one recent change and ask whether the associated assumptions received review. An annual calendar entry may provide a useful baseline without addressing every intervening change.

Follow-up prompt: connect one material change event to the existing purchasing, project, or personnel process that can detect it.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Question 15 How does the company verify closure

Select a finding marked complete and inspect the acceptance criteria and evidence. Identify who implemented the repair and who checked it. Depending on consequence, a distinct verification step may provide a useful check on implementation claims.

Compare the evidence with the actual closure statement. A purchased tool does not establish an operating control. A revised procedure does not establish operator competence. A completed test does not establish effectiveness outside its conditions.

Follow-up prompt: reopen or qualify any closure claim that exceeds its evidence, then assign the remaining verification work.

Answer and evidence reference: _____

Evidence state: SUPPORTED / PARTIAL / UNKNOWN / OUTSIDE SCOPE

Decision flag and next action owner: _____

Produce a decision from the worksheet

Choose no more than a few findings for the initial executive brief. For each, describe the observation, supporting evidence, potential consequence, and next decision. Label the consequence as observed, inferred, or hypothetical as appropriate.

The disposition may call for an internal repair, a specialist review, a bounded advisory assignment, a later review, or no additional action within the scope. A sale should never serve as the predetermined result of the assessment.

Record exclusions and unanswered questions. Give the decision owner a truthful view of what the exercise established. Keep the full evidence record available to authorized reviewers without flooding the executive brief with every working detail.

Appendix B Executive Decision Worksheet

State the decision

Write one sentence naming the action, the accountable owner, the scope, and the latest useful decision time. Replace a broad ambition such as improve AI with a particular commitment such as approve a limited internal drafting pilot for a defined team and period.

Record the version and date of the note. Identify the people supplying technical, operating, financial, and professional input. Distinguish their advice from the authority to approve the action.

Explain why the decision matters now. Identify the opportunity, the obligation, or the operating problem that creates the need. If no decision actually needs to occur, consider whether the work belongs in research rather than an approval process.

Name the alternatives

Include the proposed action, at least one relevant alternative, and the practical consequence of waiting. The alternative should represent a genuine option, rather than an obviously unattractive construction designed to validate the preferred proposal.

For each, describe its expected benefit, full cost, operating burden, dependencies, and ability to reverse course. Use consistent definitions so leadership can compare the options. Identify where one option requires evidence the others do not.

Explain exclusions. If management rejects an alternative before detailed analysis, record the reason. A legal prohibition, a technical impossibility, and a preference for the incumbent provider involve different reasoning.

Separate evidence and inference

List the material observations with their sources and dates. State the inference management draws from each. Identify the strongest contrary evidence and the uncertainty that remains. A reviewer should have enough information to challenge the reasoning without reconstructing the entire investigation.

Describe source limitations. A vendor demonstration may show only its preferred conditions. A small pilot may omit a difficult operating period. An internal report may contain inconsistent definitions. Preserve those qualifications beside the claim they affect.

Identify the assumptions that could change the choice. For each, state a confidence description and a test or monitoring route where useful. Avoid a percentage confidence estimate unless the method supports that precision.

Define authority and limits

Specify the actions approval would permit. Include relevant limits on data, users, recipients, spending, connected accounts, and duration. Exclude categories that do not apply, with an explanation where their absence might confuse the reviewer.

State which conditions require further approval. Define the stop or pause mechanism and the person entitled to use it. Identify what happens to queued work, retained information, and outside commitments when the activity ends.

Attach the technical or professional review needed for the scope. If such review remains incomplete, explain whether the decision can proceed conditionally or whether a consequential step must wait. The authorized owner should make that determination under the applicable process.

Establish the success and review conditions

State what the company expects to observe and how it will measure the result. Keep targets distinct from achieved results. Identify the baseline, the measurement owner, the test conditions, and the date of review.

Define failure and ambiguity as well as success. A test may fail to produce usable evidence because of incomplete access or an unrelated disruption. Record how the team will handle that condition without relabeling it as a pass.

Specify what would trigger reconsideration before the planned review. Assign responsibility for observing the trigger and reaching the decision owner. Where the relevant owner cannot respond, state the fallback arrangement.

Record the decision and preserve dissent

The final entry should state the authorized action, conditions, approving person, approval time, and material version. Record dissent by describing the disputed proposition and the reason it matters. Avoid reducing an objection to a name beside the word disagreed.

After the review, append the observed result, explanations, and next decision. Keep the original note intact. Where facts require correction, label the correction and its date so another reader can distinguish a contemporaneous belief from later knowledge.

Use the worksheet proportionately. A routine reversible decision may need only a few paragraphs. A consequential investment may require attachments and professional review. The purpose involves making the reasoning accessible, rather than increasing the page count.

Appendix C Bounded Advisory Engagement Charter

Purpose and professional boundary

Use this drafting brief to establish the intended operating relationship before counsel prepares or reviews an agreement. It does not supply a legal instrument or resolve jurisdiction-specific duties. Replace general language with the facts of the proposed engagement.

Name the contracting parties and the authorized signatories. Identify the business question the engagement addresses and the reason existing arrangements need additional support. State the relationship to management and any governing board.

Describe the adviser as providing recommendations within the agreed scope unless the actual agreement expressly and validly assigns another function. Do not rely on the word advisory to settle authority or liability. Have counsel evaluate the planned conduct, access, professional activities, and governing requirements.

Scope and exclusions

Identify the specific workflow, service, decision, or business unit under review. State the period examined, the information available, and the geographic or operational limits. Explain what falls outside the assignment and how an exclusion affects the conclusions.

List the expected work products. Suitable outputs might include an evidence ledger, a current-state map, a findings register, and an executive decision brief. Define the completion standard for each. A limitation can appear in a completed review, but the review cannot claim assurance over facts it never examined.

Separate advisory work from implementation, operational management, formal assurance, and licensed professional services. If the client later needs one of those services, use a defined change process or another engagement with an appropriately qualified provider.

People and responsibilities

Name the executive sponsor, decision owner, operating owner, technical owner, and adviser. Assign only the roles needed for the scope. Explain who provides evidence, who approves tests, who implements repairs, and who verifies completion.

State the expected client participation and time requirements. Identify dependencies on access, interviews, or third-party cooperation. If a dependency fails, require the adviser to explain its effect on scope, timing, cost, or confidence.

Preserve a route for independent challenge. The adviser may raise findings that management dislikes; management may reject recommendations with documented reasons. Neither party should misrepresent the other's position to manufacture consensus.

Authority and external action

List permitted activities, access, contacts, and communications. State what requires additional approval. A permission to review records should not silently expand into permission to contact customers, change systems, or make public statements.

Define any spending or transaction authority explicitly. Explain who may expand or revoke it. Identify start and end conditions and the mechanism for handling work that remains queued or in progress at the boundary.

If AI tools support the engagement, identify approved uses, permitted data, required review, and retained records. Do not grant a service broad access merely because the adviser finds it convenient. The client should understand the relevant processing arrangement before information enters it.

Conflicts and confidentiality

Request disclosure of relevant vendor compensation, referral income, investments, relationships, and overlapping engagements. Assign responsibility for evaluating potential conflicts and deciding whether disclosure, recusal, independent review, or withdrawal provides an appropriate response.

Define information handling, retention, approved sharing, and the return or deletion of records as appropriate. Address ownership and permitted use of work products. Specify how the company preserves access to its own decision history when the engagement ends.

Obtain qualified review where legal privilege, regulated information, professional duties, or competing obligations matter. An operational charter should not promise protections it cannot create.

Fees and acceptance

State the fee structure, payment terms, included work, capacity limits, and treatment of additional work. Identify expenses and any approval needed before incurring them. Avoid open-ended assumptions about unlimited executive access or technical delivery.

Agree on acceptance criteria before work begins. Management should know what it will inspect, how it will identify a material shortfall, and what correction process applies. A commercial promise should match the actual service capacity.

Evaluate benefit claims with the discipline in Chapter 7. A fee proposal should distinguish work products and measured improvements from hypothetical loss reduction. No contract should depend on a client mistaking an illustrative model for achieved performance.

Review and exit

Set the initial review date and the conditions for renewal, narrower scope, specialist referral, or closure. A successful engagement may end when the client gains adequate internal capacity. Renewal should address a remaining need.

Define the handoff, closure of access, treatment of unresolved findings, and continuing obligations. Assign management ownership of any work that transfers. Record the exact end of operational authority rather than leaving it implicit in the final invoice.

Before signature, compare the charter with the actual proposed conduct. The agreement, system permissions, and working habits should tell the same operational story. If they differ, resolve the gap before consequential work begins.

Appendix D Recovery and Authority Exercise

Set a safe and useful scope

This original tabletop exercise combines a service interruption with an authority question. It does not require shutting down production or exposing live records. The responsible operating and technical owners should adapt and approve it for the company.

Use a hypothetical service business that relies on a scheduling platform and an AI-assisted customer message workflow. The exercise asks whether the team can sustain a defined minimum service while preventing unauthorized communication during uncertainty.

Choose an exercise director, a decision owner, relevant operators, an observer, and any necessary specialist. Define the maximum duration, approved information, and stop conditions. Tell participants that the events describe a simulation and identify any actions they must not perform on live systems.

Write the objectives before starting. For example, the team may need to identify the recovery owner, locate an approved fallback, make an authorized customer communication decision, and preserve a useful incident record. Do not grade participants against secret expectations unrelated to those objectives.

Establish the starting position

Give participants a brief description of the selected service, the approved fallback, the communication workflow, and the current authority limits. Provide only information appropriate to their exercise role. Note any assumptions supplied for the scenario rather than verified in the company.

Ask the observer to record timestamps, decisions, information sources, approvals, and unresolved questions. An observation that a participant asked for a record differs from an observation that he obtained and checked it. Preserve that distinction.

If real operating conditions interrupt the exercise, stop or reschedule it. A tabletop should not compete with an actual incident for essential staff attention. Record the interruption without treating it as evidence of exercise failure.

Stage one introduces service interruption

Announce that the scheduling platform cannot currently support the selected function. Ask who confirms the condition and what evidence the team needs before activating the fallback. Do not assume every participant knows the same facts.

Ask the assigned operator to describe how he would perform the minimum service using approved exercise materials. Observe whether he can locate current instructions, contact information, permissions, and the required records. Note each dependency on unavailable people or systems.

The team should identify what it can accomplish and what remains uncertain. If it discovers that the fallback lacks critical information, ask for the next authorized decision. Do not reward a confident declaration of continuity unsupported by the walkthrough.

Stage two introduces uncertain information

Supply a message claiming that the platform provider expects rapid recovery. Identify it as an unverified exercise message. Ask the team how it would confirm the source and how the claim affects its current decision.

Observe whether the group preserves uncertainty or repeats the message as fact. Ask what evidence would justify changing the recovery plan. The exercise should reveal how information influences the decision, rather than testing whether participants can guess the director's preferred answer.

If participants reasonably choose to continue the fallback while investigating the claim, record their reasoning. If they choose another route, ask the same questions about authority, evidence, and consequence. Several defensible responses may exist under the stated conditions.

Stage three introduces a communication request

Present a hypothetical request to send customers an assurance that service will resume within a specified period. Ask who can authorize that commitment and whether the available evidence supports it.

The AI-assisted workflow can draft a message in the scenario, but drafting permission does not settle dispatch authority. Ask the technical owner to explain the actual approval path using approved descriptions or an isolated demonstration. Keep live dispatch disabled throughout the exercise unless a separately authorized test explicitly covers it.

Have the decision owner choose an appropriate communication action within the scenario. Record the approved wording or communication condition, the approving person, and the evidence supporting any factual claim. If the team chooses to withhold an unsupported promise, treat that as a decision to evaluate rather than a failure to communicate.

Stage four changes the authority boundary

Announce that a temporary delegation used in the workflow has reached its stated end. Ask how the team knows which future actions must stop, what happens to queued work, and who can authorize any continuation.

Distinguish a human instruction from a technical enforcement mechanism. A participant may understand the boundary correctly while the system design still needs examination. Record both findings. Do not claim that the tabletop demonstrated technical enforcement when it examined only the explanation.

Ask the team to describe the safe handling of unresolved work. It may need to hold a consequential step while continuing unrelated authorized preparation. Evaluate whether that distinction remains clear under pressure.

Close with repairs and retest conditions

Review each objective against the observations. Identify what worked, what failed, and what the exercise could not establish. Avoid conclusions about actual company-wide readiness from a single simulation.

Select the findings that require action. Assign an owner, a target date, and evidence of completion. A missing contact, an unclear approval, a technical permission gap, and an inadequate fallback each call for a different repair.

Choose a proportionate retest. A corrected contact may need a simple verification. A changed dispatch boundary may need a controlled technical test. A revised operating procedure may need a substitute operator to perform the walkthrough. Close the finding only when the relevant evidence supports the closure statement.

Preserve the exercise record

Retain the scenario version, participants, conditions, observations, decisions, limitations, and follow-up actions under the appropriate information policy. Keep the record factual and respectful. Its purpose involves learning and evidence, rather than manufacturing a flawless performance.

At the next executive review, ask whether the assigned repairs occurred and whether the evidence changed the company's confidence. The exercise produces value when it changes an operating decision or verifies a relevant capability. Its entertainment value does not count as readiness.

Appendix E Advisory Economics Worksheet

Define a model before interpreting its result

Use this worksheet to explore a proposed engagement, not to establish an investment recommendation from generic inputs. Every figure below serves a hypothetical illustration. The model reports neither a ROWAN Prime quote nor a client result. Replace assumptions with evidence appropriate to the actual decision and obtain financial or other professional review where warranted.

Begin with a defined period and event. The simple risk model assumes a single specified event can occur at most once during the chosen year. It uses an assumed probability before intervention, an assumed probability after intervention, and an incremental loss if the event occurs. It does not model repeated incidents, interacting events, or a distribution of loss severities.

Expected loss reduction under those assumptions equals the difference between the two probabilities multiplied by the incremental loss. Net modeled economic benefit subtracts the full economic cost of the intervention. The result describes the model's implications. It does not turn either probability estimate into an observed fact.

The post-intervention probability requires particular scrutiny. Explain which implemented control supposedly changes the event likelihood, what evidence supports that mechanism, and what uncertainty remains. A vendor promise or an adviser's confidence does not establish the size of the change.

Record the complete cost

The illustrative first-year model uses the following inputs. The opportunity value of internal time remains separate from cash spending because reducing time consumption does not necessarily reduce payroll or release cash.

Cost item	Assumed first year amount	Treatment
Advisory fee	\$36,000	Hypothetical cash outlay
Implementation	\$24,000	Hypothetical cash outlay
Internal staff time	\$10,000	Hypothetical opportunity cost
Total cash cost	\$60,000	Fee plus implementation
Total economic cost	\$70,000	Cash plus opportunity cost

Add omitted costs before using the structure in a real review. These could include specialist work, software subscriptions, training, maintenance, transition effort, and disruption. Remove costs that do not apply, with a reason. Avoid counting a software payment in both implementation expense and a separate subscription line for the same period.

For internal time, record the hours, the valuation basis, and what alternative work the company expects to forgo. If the hours would otherwise remain idle, a fully loaded salary allocation may not represent the relevant incremental economic cost. If the engagement displaces valuable work, explain that consequence explicitly.

Examine several assumptions

The following table keeps total economic cost fixed at \$70,000. Probabilities describe annual assumptions for the single defined event. The reduction uses percentage points, not a relative percentage reduction.

Scenario	Baseline probability	After intervention	Assumed loss	Modeled loss reduction	Net economic benefit
Low benefit	5%	4%	\$250,000	\$2,500	minus \$67,500
Middle benefit	15%	10%	\$500,000	\$25,000	minus \$45,000
High benefit	30%	15%	\$1,000,000	\$150,000	\$80,000

The low case uses a one-percentage-point reduction: 0.01 multiplied by \$250,000 equals \$2,500. The middle case uses five percentage points: 0.05 multiplied by \$500,000 equals \$25,000. The high case uses fifteen percentage points: 0.15 multiplied by \$1 million equals \$150,000.

These scenarios have no assigned likelihood. Do not average them into an expected benefit unless a defensible method supplies the weighting. They expose the dependence on assumptions and show that an engagement can fail the modeled economic test.

Against the \$60,000 assumed cash outlay, the middle scenario leaves a \$35,000 modeled shortfall; including the \$10,000 opportunity cost raises the shortfall to \$45,000. The \$25,000 loss-reduction input represents a risk-model estimate, not a cash receipt or an observed saving. Neither calculation describes a client's cash result.

Check feasibility before calculating a break-even claim

Divide the total economic cost by the assumed loss to calculate the required absolute probability reduction in this simplified model. At a \$500,000 loss and \$70,000 cost, the required reduction equals fourteen percentage points.

If the assumed baseline probability equals ten percent, a fourteen-percentage-point reduction cannot occur within a valid probability range. The risk-only argument therefore cannot reach break-even under those inputs, even if the event probability falls to zero. The company might identify other benefits, but it should investigate and value them separately rather than hiding the contradiction.

At a \$250,000 loss, the required reduction equals twenty-eight percentage points. At \$1 million, it equals seven percentage points. State the baseline probability beside the required reduction so the reader can evaluate whether the proposed change even fits the model.

Do not describe a reduction from fifteen percent to ten percent as a five-percent relative reduction. It equals five percentage points, or approximately one third of the baseline probability. Mixing those descriptions can materially distort a sales claim.

Compare another way of doing the work

Assume an alternative specialist assignment costs \$20,000 in total economic terms and, under its own unverified assumptions, changes the event probability by three percentage points for a \$500,000 loss. Its modeled loss reduction equals \$15,000, leaving a modeled net economic result of minus \$5,000.

Compare that with the middle advisory scenario: \$25,000 modeled loss reduction and \$70,000 cost, yielding minus \$45,000. Both examples fail the simplified economic test, but the advisory scenario's larger gross benefit does not make it economically preferable. Cost and the relevant alternative matter.

In a real decision, the alternatives may address different exposures or provide different capacities. Identify those differences rather than forcing dissimilar services into a misleading single measure. If a required safeguard must occur regardless of this calculation, treat the obligation as a constraint and compare compliant alternatives.

Handle operating benefits separately

An intervention may also affect rework, service capacity, decision speed, or customer experience. Specify the measure and baseline before valuing the change. Identify what management actually observed and how much of the observation plausibly connects to the intervention.

For an illustrative time calculation, ten hours saved per month at an assumed \$50 per hour creates \$6,000 of annualized time value: ten multiplied by twelve multiplied by fifty. It does not establish a \$6,000 cash saving. The company needs to explain what the people do with the released capacity and whether a full year's extrapolation fits the observed period.

Do not add this amount to avoided rework if the same hours appear in both measures. Do not count added revenue at its full value when the relevant benefit concerns contribution after incremental costs. Where the attribution remains uncertain, report the measured operating change separately from any financial estimate.

If a ninety-day pilot produces a promising result, identify seasonal effects, workload changes, and other interventions before extrapolating. A limited test can justify further investigation without supporting an annual guarantee.

Keep first year and recurring costs distinct

The first year may contain implementation costs that do not repeat. Later years may contain maintenance, review, and retraining costs that the initial model overlooked. Construct each period deliberately rather than multiplying the first-year result by a desired contract length.

For a hypothetical second year, assume a \$36,000 fee, \$8,000 maintenance expense, and \$4,000 opportunity cost. Total economic cost equals \$48,000. Under the middle scenario's unchanged \$25,000 estimated loss reduction, the modeled result remains minus \$23,000. Removing implementation expense does not automatically establish a positive case.

Longer-term analysis may need discounting, changing exposure, tax treatment, and other financial considerations. An appropriately qualified reviewer should determine what the actual decision requires. The simple worksheet supplies an initial discipline rather than a complete valuation method.

Preserve the assumptions and the commercial decision

Record who supplied each input, its source, the date, the uncertainty, and the reviewer. Label unsupported estimates as assumptions. Where the evidence cannot support a useful numerical estimate, consider a bounded pilot with capped cost and concrete delivery criteria instead.

Management may still choose a service whose financial return remains difficult to quantify. It should explain the decision in those terms. Do not convert an unquantified strategic preference into a numerical ROI claim for marketing convenience.

At renewal, compare the actual delivery, operating observations, full cost, and remaining need. Keep modeled avoided losses distinct from achieved financial results. A transparent account can justify continuation, narrower scope, another provider, or closure. The model should permit all four.

Appendix F Worked Decision Brief

Read the example as a hypothetical exercise

The following brief demonstrates how the book's proposed practices can connect to one decision. The organization, observations, permissions, costs, and results all belong to a hypothetical exercise. They describe no client, existing ROWAN Prime implementation, or achieved business outcome.

The exercise concerns a field-service company's proposal to use an AI tool for internal drafting of customer scheduling messages. Management wants to investigate whether the tool can reduce drafting burden without granting dispatch authority or exposing live customer records during the initial test.

This example stops at an initial decision and defines evidence for a later review. It deliberately contains no invented pilot results. A completed decision process would append actual observations after an authorized test.

Decision and ownership

Proposed action: approve a ten-working-day internal drafting pilot involving two designated employees, an approved set of artificial scheduling records, and a separate review environment. The chief operating officer holds the decision in this exercise, subject to the company's applicable spending and technology approval arrangements.

The operations manager would own the pilot's day-to-day process. The technical owner would verify the tool configuration and access limits. A designated reviewer would check each draft against the approved exercise record. The adviser would examine the decision evidence and produce a review note; he would hold no independent authority to expand the pilot.

The pilot would permit internal drafting only. It would not permit sending messages, using live customer or personnel records, connecting a production mailbox, or making external commitments. Any proposed expansion would require a new decision after the company examines the change in exposure.

Business reason and alternatives

Management reports a drafting burden during schedule changes, but the exercise supplies no verified estimate of its annual financial effect. The pilot would investigate the workflow before management commits to a wider deployment.

Alternative one would improve the existing human message templates. Alternative two would conduct the bounded AI pilot. Alternative three would defer both approaches while collecting a better baseline. A broad autonomous communication rollout would fall outside the immediate decision because management has not examined its permissions or operating consequences.

The reviewer should compare template improvement with the pilot rather than assuming AI provides the only route. Both approaches would require staff time. The simpler option could prove sufficient. The pilot would merit continuation only if actual evidence supports a useful incremental benefit within acceptable limits.

Evidence available and evidence missing

For the exercise, assume management can provide approved message templates and artificial scheduling changes. Those materials would define the factual basis against which the designated reviewer checks each draft. They would not establish performance on live customer situations.

The exercise assumes the technical owner can inspect the pilot environment. The actual pilot should wait for his evidence that the intended access restrictions work. A written statement of the intended restriction would not substitute for examination of the configured environment.

Unknowns would include drafting time, reviewer workload, error patterns, and the practical value of the revised process. Management would not claim those results before testing. Performance with artificial examples would also leave uncertainty about actual operating conditions.

Assumptions that could change the decision

The proposal assumes the company can run the pilot without live confidential information and without production dispatch access. If either assumption fails, the technical owner should stop that consequential step and return the changed proposal to the decision owner.

It also assumes the designated reviewer has sufficient time and competence to inspect each draft. If the review burden prevents meaningful checking, management should reduce the test volume or pause the pilot. A system that produces more text than the reviewer can assess would not meet the intended process.

Finally, the proposal assumes the artificial records represent a useful range of drafting tasks. The operations manager should identify their limitations. The pilot may reveal a narrow result that justifies only another bounded experiment.

Cost and measurement

Management would set a fixed ceiling for pilot expenditure and a staff-time budget before approval. The exercise supplies no quoted software price and no promised saving. A real brief should record the actual approved figures and their source.

The baseline task would ask employees to prepare messages using the existing templates. The pilot task would ask them to prepare comparable messages with AI assistance, followed by the required review. Record drafting and review time separately, along with corrections, rejected drafts, and unresolved factual issues.

Compare like tasks where practical and describe differences where exact comparability proves impossible. Do not count drafting time alone while ignoring review time. Do not convert a small artificial-record experiment into an annual cash-saving claim.

Boundary and intervention evidence

Before the pilot begins, the technical owner would document the approved environment, the accounts involved, and the absence or restriction of dispatch routes. An authorized controlled check would examine the material boundary. Record its expected behavior, observed behavior, and limitations.

The operations manager would hold authority to suspend use within the pilot. The technical owner would manage the relevant technical restriction. The decision owner would decide whether any changed scope could proceed. These roles should appear in the actual working instructions.

At the end of the ten-working-day period, pilot authority would end unless the authorized owner expressly approves a next stage. The team would retain the required evidence, close unnecessary access, and handle records under the agreed information policy.

Proposed disposition

The adviser could recommend conditional approval for the narrow pilot, with no execution until the technical boundary and required approvals receive verification. That recommendation would differ from final permission to start. The decision owner should identify and record the satisfaction of the conditions.

The eventual review could support template improvement, another bounded pilot, a changed tool, or closure. A favorable result on artificial records would not by itself support external dispatch. A poor result would not establish that every AI application lacks value.

This brief gives leadership a way to proceed without pretending that a preliminary proposal already contains its own success evidence. Its practical value comes from connecting the question, owner, limits, test, and next decision in one record.

Appendix G Evidence Notes and Reading

How the references support the book

The references below support selected propositions identified by bracketed source numbers in the chapters. The worksheets, engagement sequence, hypothetical examples, and DREADNOUGHT formulation present original proposed applications. The source organizations have not endorsed those applications or the ROWAN Prime service.

The book does not use these references to claim that an independent strategic adviser causes higher profit, that the questionnaire predicts business failure, or that the proposed fees reflect verified market demand. Those propositions would require different evidence and appropriate study designs.

Source verification for this edition occurred on 16 October 2026 Central time. Publication dates below identify the cited editions. A later revision can change the relevant guidance; readers should inspect the edition and current context before applying it to a consequential decision.

S1 Independent judgment in formal governance

OECD. G20/OECD Principles of Corporate Governance 2023. OECD Publishing, Paris, 2023. Chapter V, especially sections V.E and V.F.

<https://doi.org/10.1787/ed750b30-en>

https://www.oecd.org/en/publications/g20-oecd-principles-of-corporate-governance-2023_ed750b30-en/full-report/component-8.html

The cited discussion addresses governing-board judgment and access to information. Chapter 1 draws a limited analogy when asking whether a company has adequate challenge and information. The reference does not establish the legal role or commercial effectiveness of the contracted advisory arrangement described in this book.

S2 Information system contingency planning

Swanson, Marianne; Bowen, Pauline; Phillips, Amy W.; Gallup, Dean; and Lynes, David. Contingency Planning Guide for Federal Information Systems. NIST Special Publication 800-34 Revision 1, May 2010, including the November 2010 updates. Chapter 3 provides the planning process and impact-analysis discussion.

<https://doi.org/10.6028/NIST.SP.800-34r1>

<https://csrc.nist.gov/pubs/sp/800/34/r1/upd1/final>

The reference supports the specific IT contingency-planning connection in Chapter 3. Broader service mapping, the example exercises, and the advisory operating model remain proposed applications rather than findings attributed to this publication.

S3 AI risk management organization

Tabassi, Elham. Artificial Intelligence Risk Management Framework AI RMF 1.0. NIST AI 100-1, January 2023. Part 2, section 5, AI RMF Core.

<https://doi.org/10.6028/NIST.AI.100-1>

<https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>

Chapter 4 cites the four functions and the cross-cutting role of governance. The cited resource identified a revision in progress at the access date. This manuscript cites the 2023 edition and makes no claim that its terminology represents a newly finalized replacement.

S4 Generative AI risk considerations

Autio, Chloe; Schwartz, Reva; Dunietz, Jesse; Jain, Shomik; Stanley, Martin; Tabassi, Elham; Hall, Patrick; and Roberts, Kamie. Artificial Intelligence Risk Management Framework Generative Artificial Intelligence Profile. NIST AI 600-1, July 2024. Section 2 identifies relevant risk categories.

<https://doi.org/10.6028/NIST.AI.600-1>

<https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

Chapter 4 cites selected risk categories. The book's proposed output-checking examples apply that concern to an executive workflow; they do not assert that the profile validates any particular tool, reviewer, or assurance claim.

S5 Cybersecurity governance outcomes

Pascoe, Cherilyn; Quinn, Stephen; and Scarfone, Karen. The NIST Cybersecurity Framework CSF 2.0. NIST Cybersecurity White Paper 29, February 2024. Section 2 and Appendix A describe the Core and Govern outcomes.

<https://doi.org/10.6028/NIST.CSWP.29>

<https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>

Chapter 9 cites the function structure and the connection with organizational context and oversight. It does not report an assessment of conformity, implementation maturity, or cybersecurity effectiveness for any actual company.

S6 Technical control reference points

Joint Task Force. Security and Privacy Controls for Information Systems and Organizations. NIST Special Publication 800-53 Revision 5, September 2020, including December 2020 updates. Controls AC-6, AU-9, and CP-4.

<https://doi.org/10.6028/NIST.SP.800-53r5>

<https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final>

Chapter 8 uses those control names as limited reference points. NIST lists subsequent minor catalog releases, including Release 5.2.0 in 2025. This edition cites the identified Revision 5 text, not an assertion of the latest implementation baseline. No control mapping or certification follows.

Claims that remain outside the evidence

Whether a company needs additional advice depends on its decisions, existing competence, independent challenge, and capacity to act. This book's framework supports examining that need. It does not establish that every company should purchase an outside advisory engagement.

The DREADNOUGHT method needs independent challenge, field testing, and revision. The worksheet needs evaluation for usability and the quality of decisions it supports. A useful discussion instrument does not automatically qualify as a validated diagnostic product.

Commercial demand, actual service capacity, contracting terms, and offer pricing need their own verification. The hypothetical annual fee and financial scenarios supply arithmetic examples only. The book supplies no verified client outcomes or established board-service record.

Any public description of named software should follow an examination of its actual status and capabilities. A conceptual framework, a local prototype, an implemented internal tool, and a supported production service require different descriptions. This manuscript makes no delivery promise for CACHELINE, DREADNOUGHT RECALL, SHADOW HUNTER, DARAM, OUTCOME INFERENCE ENGINEERING, or BEHAVIOR-BASED SUCCESS.

A disciplined reading practice

When a source appears persuasive, ask which exact proposition it supports and what would make its application misleading. Keep the source's population, context, date, and limitations beside the claim. Use primary material where available and inspect the actual passage rather than relying on a generated summary.

When the book proposes a practice, evaluate it in the same spirit. Ask whether the practice fits the decision, whether its cost appears proportionate, and what evidence would show that it helps. Retain the possibility of a simpler internal solution or no additional intervention. Independent advice deserves independent examination.

WHO CHALLENGES THE DECISION BEFORE THE COMPANY PAYS FOR IT?

Growth adds complexity. AI expands capability. Neither automatically strengthens judgment.

A leadership team can possess experience, information, and confidence—and still overlook the dependency, assumption, or authority gap that changes everything.

The Missing Seat examines the function an independent strategic adviser can fill: asking difficult questions, testing consequential assumptions, and connecting evidence to accountable decisions.

Inside, Larry Harmon explores:

- When outside advice adds value—and when another route makes more sense.
- How to examine AI authority, operating resilience, and hidden dependencies.
- How to evaluate advisory fees without mistaking hypothetical savings for proven results.
- How to structure an advisory board and its first ninety days.

Includes a fifteen-question Boardroom Blind Spot Worksheet, an executive decision worksheet, and a hypothetical worked decision brief.

**You may not need another meeting.
You may need someone willing to challenge its
conclusion.**

REQUIRE EVIDENCE. PRESERVE THE RIGHT TO DISAGREE.

ROWAN PRIME

