oak CONSULT

SECOND EDITION

The Trust Horizon

How Leaders Decide What AI Can Be Trusted to Do



Mark Conway

November 2026

Table of Contents

Foreword 3
Executive summary 4
1 - What a year of adoption tells us 5
2 - Establishing your Trust Horizon 6
3 - Start with purpose and the actual workflow 8
4 - The A.B.C. framework in operation 10
A - AI Assurance 10
B – Built-in Human Controls 11
C - Compliance 12
5 - Test the customer and organisational consequences 13
6 - Establish the economics of trustworthy use 15
7 - A practical first 90 days 18
8 - Prepare for change over the next year 20
9 - The next decision 21
How Oak Consult can help 21
Appendix A - Record AI use and authority 22
Appendix B - Assess one workflow using A.B.C 23
Appendix C - Check whether the benefit arrived 24
Appendix D - Selected regulatory reference points 25
Sources and evidence notes 26

This whitepaper is a fresh treatment of the original Trust Horizon paper. It combines published research with Oak Consult's analysis and practical recommendations. The continuing B2B example and numerical business case are illustrative. Research case studies are attributed to their source. Evidence and regulatory references were checked on 30 October 2026.



Foreword

When I published the first edition of The Trust Horizon, I argued that the question facing business leaders was changing. Understanding what Artificial Intelligence (AI) could do was only the beginning. The harder question was deciding what to trust it to do—and under whose authority.

A year on, that question has become more practical and more pressing. Organisations are finding useful applications, learning where the technology falls short, and discovering how much depends on the way work is organised around it. The opportunity is substantial. Turning it into lasting improvement takes more than access to a capable model.

In preparing this edition, we have applied that same scrutiny to our own thinking. We have revisited the original forecasts, checked the evidence behind the claims, and challenged the assumption that greater autonomy necessarily represents progress. Some ideas have held up well. Others needed revision. That willingness to test and improve is central to the approach we advocate.

We have also worked to make the argument more concrete. Consider a supplier using AI to prepare quotations. Faster drafting could help the team respond sooner and clear a backlog. Testing might also reveal conflicting account terms or unreliable product records. The useful outcome could be to retain AI-assisted drafting, improve the underlying data and keep approval with an experienced colleague. That is a worthwhile improvement, with a clear basis for deciding what comes next.

The customer's experience matters throughout. Does the change make it easier to get the right answer? Does it reduce effort, improve reliability or help someone resolve a problem? These questions give innovation a purpose and provide a more demanding test of success than the amount of work automated.

My ambition for this edition is to help leaders move forward with confidence grounded in evidence. Start with a worthwhile problem. Test the proposed improvement. Learn from what happens. Extend authority when the results justify it. At Oak Consult, we have used AI for several years and seen substantial improvements in its capabilities. It has helped us streamline internal processes, improve how we work and save significant time in producing research, analysis and content.

We have also learned through testing, challenging and refining its outputs. Useful results depend on clear instructions, appropriate boundaries and checks throughout the work. A confident answer can still be wrong, miss important context or fail to serve the intended purpose.

That experience informs this paper. We see considerable value in AI and continue to explore new ways to use it. We also retain responsibility for what we publish, recommend or act upon. The Trust Horizon should move as we learn.

Mark Conway, Oak Consult



Executive summary

AI can release capacity, improve response times and help people complete useful work. Leaders need to decide where those benefits justify the cost and how much authority to give the system.

The 2026 evidence gives reason to start. UK businesses report productivity and process improvements, while enterprise research links stronger trustworthy-AI practices with better reported returns. The same studies reveal uneven adoption, incomplete controls and costs that can erode the gains. Their findings support a practical approach: choose a worthwhile problem, test it in real work and expand when the results justify it. [1–4, 12]

The Trust Horizon is the boundary of AI activity an organisation can justify authorising, given the evidence, controls and accountability available. That boundary can expand as confidence is earned. It can also remain deliberately narrow when a limited application already delivers enough value.

Oak Consult's A.B.C. framework helps leaders make that decision. AI Assurance tests performance. Built-in Human Controls establish ownership, permissions and intervention. Compliance identifies the duties and commitments that apply.

Five questions guide the work:

1. **What business or customer outcome are we trying to improve?**
   Identify a specific problem and establish how you will recognise a better result.
2. **What can the system influence or do on our behalf?**
   Define what it may recommend, change, communicate or commit, and where human approval is required.
3. **What evidence supports that authority in real working conditions?**
   Test normal work, difficult cases and exceptions, including circumstances in which the system should stop or seek help.
4. **Who can intervene and organise recovery?**
   Name the people with the knowledge, access and authority to interrupt the work and put things right.
5. **Is the result worthwhile after the full costs are counted?**
   Include implementation, operation, checking and correction, then establish whether the benefit reaches the business and its customers.

Take a first step and start with one important workflow. Establish its current performance, test the proposed use and decide how the released capacity will be used. Include the customer's experience and the effort of checking and correction.

The result should be a clear decision to continue, improve, restrict, expand or stop. This paper shows how to reach it, using a B2B supplier example, a worked cost assessment and practical review tools.



1 - What a year of adoption tells us

Wider use does not establish operational depth

The UK government's AI Adoption Survey 2026 provides an important baseline. It reports that adoption doubled from 16% to 32%. The research covers private-sector businesses with at least five employees in specified sectors, and its definition of adoption includes piloting. The quantitative fieldwork ran from January to April 2026. The findings describe that population and period. [1]

Among adopters, the average reported proportion of staff using AI changed little, from 30% to 31%. Only 14% rated AI's importance to operations at 8 to 10 on a 0 - 10 scale. Individual benefits can therefore coexist with limited organisational dependence. [1]

A venue in the report's qualitative case studies illustrates the distinction. One operations manager used AI extensively and described useful benefits. A wider rollout provided colleagues with licences after basic training, yet the manager observed limited regular use elsewhere. Access and basic training had not translated into regular use across the team. [1, case study 6]

Leadership teams should ask where AI has changed a repeatable process, what proportion of the relevant work it supports and who experiences the benefit. Those questions are more revealing than an organisation-wide adoption label.

Agentic deployment brings questions of authority

An AI agent is a system configured to pursue a task by selecting and executing actions, potentially using several tools or business systems. Its authority depends on its configuration and permissions. The term alone says little about how independently it operates.

Deloitte's 2026 research surveyed 3,235 business and IT leaders involved in AI programmes across 24 countries. Only 21% reported mature governance for agentic AI, while 74% expected at least moderate agent use by 2027. That expectation highlights the work these organisations need to complete before expanding authority. [2]

Production research offers a useful counterweight to demonstrations of unrestricted autonomy. The current version of Measuring Agents in Production draws on 20 case studies and 86 deployed-systems practitioners across 26 domains. It reports that 68% of the systems examined executed no more than ten steps before human intervention. Deliberate limits were a feature of these deployments. [3]

These studies describe different samples and cannot be combined into a single market estimate. The research cases in this paper illustrate individual experiences. Full methods and limitations appear in the source notes.

Decision question: Where has AI changed the work in your business, and what evidence shows the result?



2 - Establishing your Trust Horizon

The Trust Horizon is the boundary of AI activity an organisation can justify authorising, given the evidence, controls and accountability available. Comfort alone is a weak basis for that boundary. Familiar outputs can look reliable when they contain plausible mistakes. Conversely, an unfamiliar system may perform a narrow task acceptably when its inputs, permissions and consequences are well controlled. Leaders need to connect their confidence to evidence about the particular application.

Authority changes the decision

The evidence and controls should reflect what the system can actually do. The table below shows how the requirements change as its activity reaches further into the business. These activities can overlap within one workflow; each needs appropriate controls.

Activity	Evidence and control to establish
Drafting or summarising	Establish which sources the system should use and what information it may access or disclose. A competent reviewer should check accuracy, important omissions and suitability before the output is used.
Recommending a decision	Check that the recommendation draws on relevant, current evidence and makes material uncertainty visible. Name the person responsible for weighing that evidence and making the decision.
Updating a business record	Limit which records and fields the system may change. Validate proposed updates, retain a history of changes and establish how incorrect entries—and any downstream effects—will be corrected.
Communicating or committing externally	Define approved terms, permitted recipients and limits on prices, delivery promises or other commitments. Require approval or escalation when a case falls outside those boundaries, before anything is sent.
Executing connected actions	Examine the combined effect of actions across systems or agents. Set overall limits, check dependencies and establish how to interrupt the sequence, identify unfinished work and manage commitments already made.

For the supplier, permission to prepare a quotation does not include permission to send it or reserve stock. Each additional action needs its own justification and an explicit authorised boundary.

An application that drafts a quotation and one that sends it to a customer may use similar technology. They create different business exposure. A well-written draft can still rely on an obsolete price, unavailable stock or an ambiguous service promise. Sending it turns those weaknesses into a customer-facing commitment.

Assess the seriousness of a mistake, how quickly it could spread and whether the effect can be reversed. Include people affected by the decision, commercially sensitive information and services on which customers depend. A small transaction repeated thousands of times may create a material aggregate exposure.

Value and dependence are separate questions

Assess value alongside the authority granted to the system and how much the business relies on it. Clear value means a useful outcome demonstrated in practice after the full costs are considered. Exposure increases when the system can make consequential commitments or when its failure would materially disrupt the business.

Demonstrated value	Authority or operational dependence	Leadership response
Clear	**Limited —** The system supports a bounded task. Its permissions are narrow, and colleagues can continue the work without substantial disruption.	**Retain the useful boundary.** Monitor quality, cost and customer outcomes proportionately. Extend its role only where there is a worthwhile purpose and evidence to support the change.
Clear	**Substantial —** The system makes consequential changes or commitments, or the business relies heavily on its availability and outputs.	**Protect the value already established.** Test continuing reliability, permissions, review capacity and recovery. Establish how service will continue during an outage or withdrawal of the tool.
Unclear	**Limited —** Benefits remain unproven, but the activity is contained and the business can readily work without it.	**Run a bounded test or discontinue the use.** Set a baseline, an owner, a spending limit and a review date. Continue when measured benefits justify the effort and cost.
Unclear	**Substantial —** The system has significant authority or has become important to operations, while its overall benefit remains uncertain.	**Prioritise examination and contain exposure.** Restrict unsupported actions, establish a workable fallback and investigate the full costs and outcomes. Manage any withdrawal carefully where customers or staff already depend on it.

This is a discussion aid, not a scoring model or a measure of organisational maturity. The same business may occupy all four positions across different workflows. A useful, limited application can remain the right choice; greater autonomy is not a required destination.

The authorised boundary should be written down. It should state permitted actions, exclusions, the accountable owner and the conditions requiring human approval. Review the boundary when data, suppliers, permissions or customer circumstances change.

Decision question: What may this system do today, and what would justify giving it more authority?



3 - Start with purpose and the actual workflow

Establish a problem worth solving

The starting point is a recurring business or customer problem with an observable consequence. Examples include delays in responding to enquiries, inconsistent handovers or excessive effort preparing reliable information. Establish the current performance before choosing how AI might contribute.

The UK survey found that 65% of businesses neither using nor planning to use AI cited the absence of a clear business case or proven return as a barrier. Start by establishing whether the problem warrants the investment. [1]

Some problems are better addressed through clearer ownership, improved records, simpler software rules or removal of an unnecessary step. Comparing those alternatives helps establish whether AI adds enough value to justify its operating requirements.

A continuing example

Throughout this paper, consider a specialist B2B supplier receiving requests for replacement parts and related support. This is an illustrative scenario, not an Oak Consult client case.

The supplier wants to shorten the time between an enquiry and a usable quotation. A proposed AI workflow reads the request, identifies the customer and equipment, consults account and product information, drafts a quotation and prepares a response.

Later proposals would allow it to update records, send routine quotations and reserve stock.

Before testing the tool, the business must resolve which source governs pricing, how equipment compatibility is established and who can promise a delivery date. It must also identify contract exceptions and information that cannot be shared.

The immediate goal is a dependable draft that a colleague can approve quickly. Automatic sending is a separate decision.

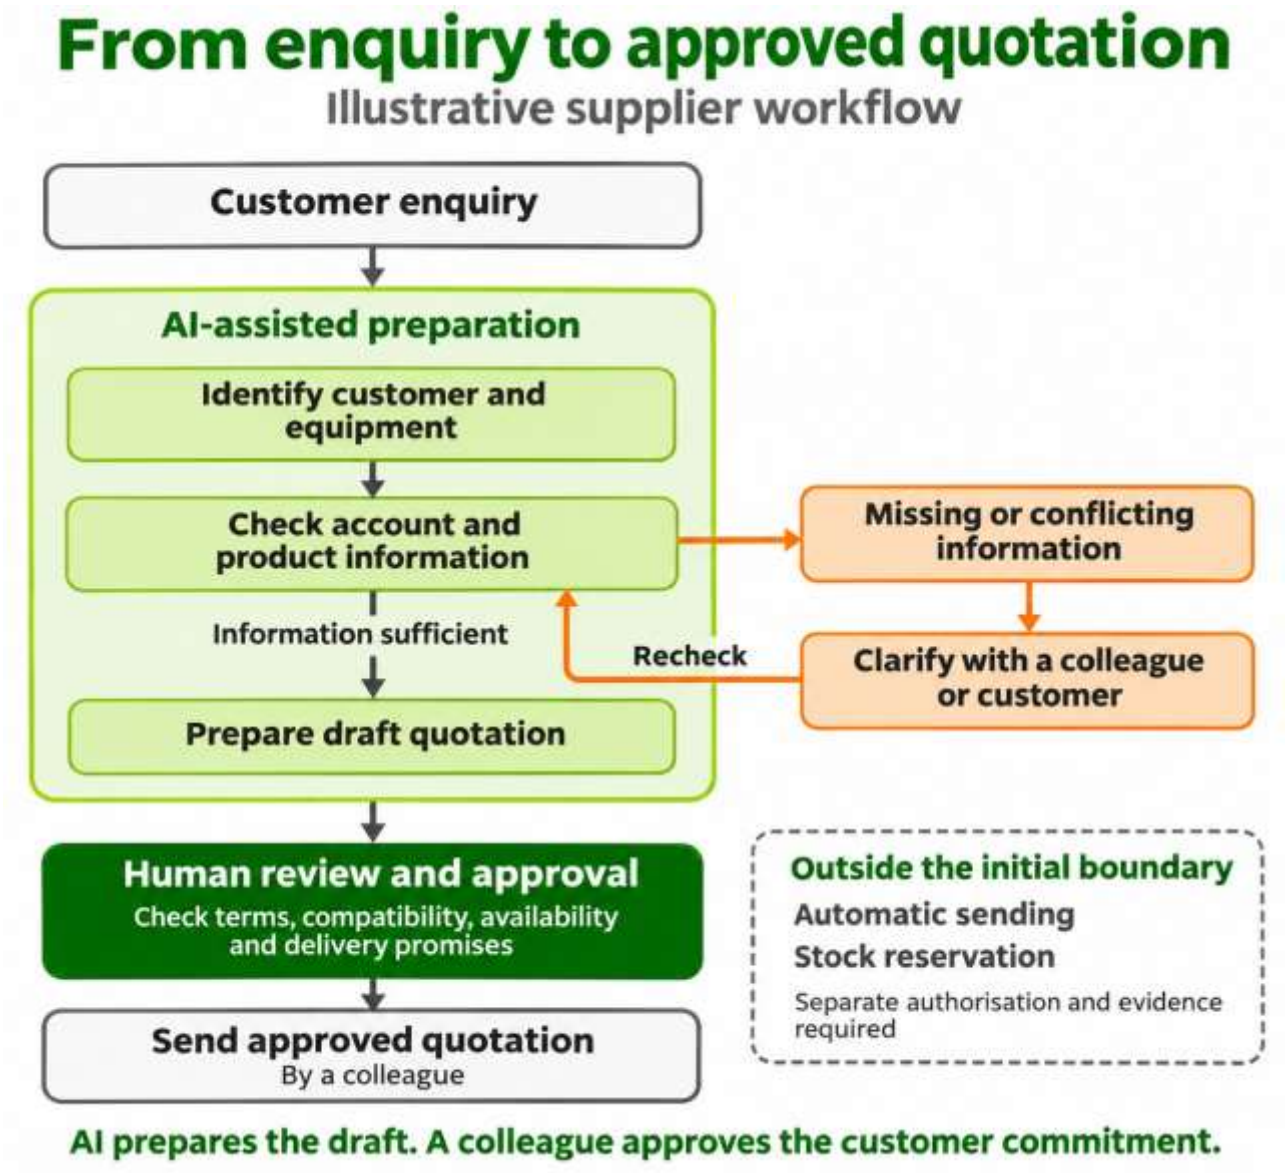

Information and dependencies

Data quality includes meaning and currency. Two records may each be valid yet refer to different customer sites or versions of an agreement. A renewal note may be accurate but incomplete because an unresolved complaint sits elsewhere. The workflow needs rules for conflicting or missing information and a route for clarification.

Document the systems, suppliers and people on which the process depends. Hosted and open-weight options create different responsibilities for security, maintenance, deployment and cost. Access to model weights does not establish data protection, reliability or affordable operation. A hosted product's controls also need examination in the configuration actually purchased and used.

Ask a further commercial question: could this improve delivery of an offer whose market value is falling? In the UK research, a creative agency reported faster development work while its owner described reduced demand for work historically supplied. Its experience shows why the business case must include changing customer demand. [1, case study 4]

Decision question: Are we improving a task that customers still value, using information the business can rely on?



4 - The A.B.C. framework in operation

A.B.C. connects the intended outcome to the evidence and controls needed to authorise the work. Its three elements are AI Assurance, Built-in Human Controls and Compliance.

The framework draws on established concerns in AI risk management and organisational control. NIST's AI Risk Management Framework and ISO/IEC 42001 offer broader reference points. The latter sets requirements for an AI management system. A.B.C. is not a substitute for either, a formal mapping of every requirement or a certification scheme. [5, 6]

A - AI Assurance

Assurance should examine the workflow's performance in its intended context. A model's general benchmark result cannot establish that the application will use the right customer record, interpret a local policy correctly or recover from a failed system connection.

Write conditions that lead to an action. For the supplier, a quotation based on conflicting account terms must be held for clarification. An unauthorised external message should trigger suspension of automatic sending and investigation of affected work. Review effort above the agreed limit should prompt a fresh assessment of the business case. Set these conditions before the trial, then record whether they were met.

Define acceptable outcomes before testing. Include common cases, difficult cases and circumstances in which the system should decline to act. Record the sample used, its limitations, serious failures and what remains untested. Assess output quality and the actions taken, including changes to other systems.

For the illustrative supplier, useful test cases include ambiguous equipment identifiers, conflicting prices, unavailable stock and requests outside an existing agreement. An acceptable response may be a request for clarification or referral to a colleague. A complete-looking answer is not always the right answer.

Continuing evaluation should detect relevant deterioration after release. The evidence may combine sampled checks, customer outcomes, system events and investigation of exceptions. More consequential actions may justify automated monitoring and faster intervention. Monitoring creates information; named people still need to interpret it and act.

Retain records sufficient to investigate material actions: the request, relevant source versions, permission used, action taken, result and any intervention. Logging should be proportionate, protected and subject to appropriate retention. A readable model explanation is not conclusive evidence of why an action occurred. Verify explanations against available sources and observable events.

Security belongs in this examination. Prompt injection occurs when hostile instructions arrive through content a system reads, such as an email or document. NCSC guidance explains why language models do not provide a dependable instruction–data security boundary. Design must therefore consider the consequences of hostile content, including which actions remain possible if the model follows it. [7]

For the supplier, a customer attachment must not acquire the authority to change prices, disclose other accounts or redirect a payment. Restricting tool permissions, separating sensitive actions and independently validating commitments can reduce exposure. Technical specialists should assess the controls and their limitations.

B – Built-in Human Controls

Assign a business owner who can explain the intended outcome and accept responsibility for the authorised activity. Identify the people responsible for technical operation, review, security and recovery. Give those people the time, information and authority to do the work.

Establish what the system may read, change, send and commit. Give agents distinct identities where appropriate and grant only the access required. NIST's 2026 work on software-agent identity and authorisation highlights this developing area. The practical concern is whether an organisation can identify the actor, understand its authority and revoke that authority when necessary. [8]

Approval should occur before an action creates an unacceptable consequence. In the supplier example, staff might approve every quotation initially. Later, a restricted class could proceed without individual approval if evidence supports reliable classification, current prices and enforceable limits. Unusual terms, uncertain compatibility and aggregate exposure could still require review.

Approval by exception depends on detecting exceptions. A system that fails to identify its own uncertainty can send a problematic case down the routine path. Test that routing explicitly, using difficult cases and independent checks rather than relying on confidence labels.

Make human oversight workable

The UK survey found that 65% of adopters reported significant human input or checking and 16% reported some. A further 11% reported a little checking. The survey measures reported involvement; the effectiveness of that checking needs testing in the workflow. [1, figure 27]

A reviewer needs relevant knowledge, access to evidence, sufficient time and authority to challenge the result. They also need to know what to examine. Repeated approval under volume pressure can become habitual. Monitor review workload, observed errors and whether important issues are detected before harm occurs.

Plan for absence and out-of-hours work. Define what stops when a reviewer is unavailable, how urgent work is handled and how the business maintains an acceptable service without the AI. Preserve the skills needed to review and recover; dependence can grow while those skills weaken.

Connected agents require collective controls. A sales agent, stock agent and scheduling agent might each obey a local limit while together promising unavailable capacity. Establish which role owns the combined commitment and how conflicting actions are reconciled. Revoking one agent's access may leave related tasks or messages in progress.

Recovery should distinguish reversible records from irreversible consequences. A database change may be restored; an email, disclosed document or customer decision may require a different response. Rehearse interruption, investigation and customer communication before relying on a stop button.

C - Compliance

Identify applicable obligations by jurisdiction, organisational role, data and use. Distinguish law from contractual commitments, internal policy and voluntary standards. The fact that a supplier describes a product as compliant cannot settle the customer's responsibilities for its deployment.

An obligations record should identify the requirement, accountable owner, supporting evidence and review date. Examine data handling, confidentiality, intellectual property, customer disclosures and any relevant sector duties. Procurement should establish what a supplier will disclose about material changes, incidents and exit arrangements.

For UK uses of personal information, consider current data-protection requirements and safeguards for consequential automated decisions. The ICO's summary of the Data (Use and Access) Act changes explains the importance of information, representations, human intervention and contestability in relevant cases. Specialist assessment should establish how these requirements apply to the workflow. [9]

EU AI Act duties differ across systems and actors, with staged application dates. The UK public sector also has its own transparency arrangements, including the scope and exemptions of the Algorithmic Transparency Recording Standard. Appendix D records selected reference points as at 30 October 2026. Establish the scope for the organisation and use concerned. [10, 11]

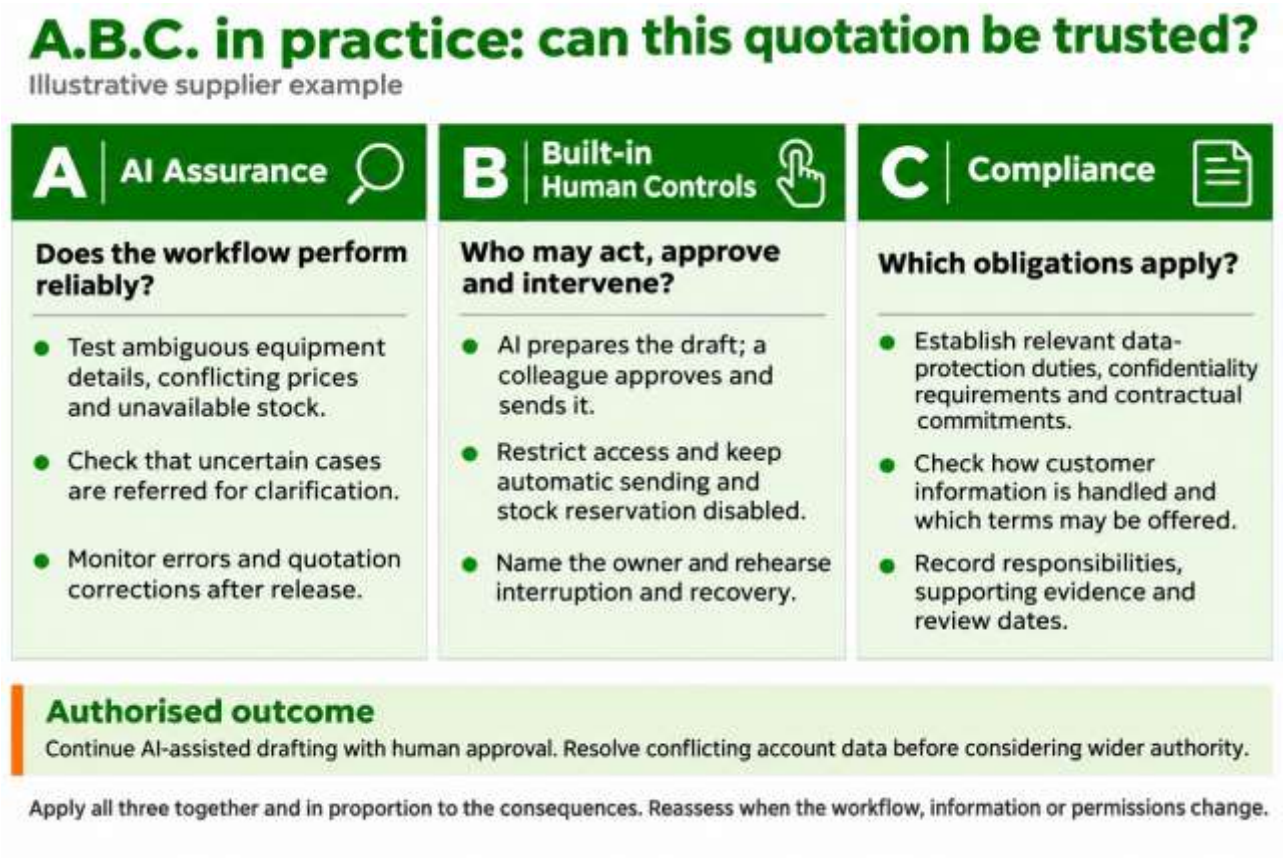

Decision question: Can the owner demonstrate acceptable performance, effective intervention and the evidence required for this use?



5 - Test the customer and organisational consequences

An internal performance improvement deserves examination from the customer's side. Faster response can be valuable, but a quick, incomplete answer may create repeat contact and delay the eventual resolution. A reduction in escalations can reflect fewer problems or a less accessible route to help.

Customer Spectacles is Oak Consult's approach to examining a business through the customer's experience: what they are trying to achieve, what the business asks of them and whether it delivers what matters. Applied to AI, it means following the complete customer task, including what happens when the routine process fails. Ask what the customer had to repeat, check or correct. Could they reach someone with authority when their circumstances required help? Did the faster internal process make their task easier?

Internal measures show part of the result. Examine what happens next: whether the customer gets a usable answer, can obtain help and achieves their intended outcome. Include work passed to other teams or back to the customer.

Internal measure	Evidence to examine alongside it
Faster initial response	Measure the time to an accurate, usable answer or resolution. Check whether the first response moves the customer forward, or leaves them repeating information, correcting mistakes or chasing a complete answer.
Fewer escalations	Establish whether fewer cases need specialist help and whether customers can still reach someone with authority when required. Review difficult cases to check that escalation is working, rather than being delayed, discouraged or missed.
More quotations issued	Check that quotations contain correct terms, compatible products and delivery promises the business can fulfil. Examine whether they lead to suitable, profitable orders, alongside customer queries, corrections and cancellations.
Fewer service tickets	Examine repeat contact across channels, unresolved problems and customers abandoning the process. Establish whether demand has genuinely fallen because issues are resolved, or whether obtaining help has become harder.
Higher staff output	Count the checking, correction and follow-up needed to make outputs usable, including work absorbed by other teams and customers. Assess quality and whether the additional output improves service or creates avoidable work elsewhere.

For the supplier, success means helping customers obtain the right part on dependable terms. Faster replies and more quotations matter when they contribute to that result.

Acceptance is specific to the task

The UK research contains qualitative examples of customers welcoming faster service and others preferring human contact. These business accounts suggest questions to test with customers directly. [1]

A buyer may welcome automatic order confirmation and want a person to handle an urgent replacement affecting production. Test acceptance at the point of consequence. Provide understandable explanations and a workable route to challenge material outcomes where relevant.

Public-service settings require particular attention to access, vulnerability and recourse. A person may have limited choice of provider or limited ability to navigate an automated process. Assess whether the service remains usable when someone lacks the expected documentation or cannot use the preferred channel. The UK business survey does not represent public administration, education or health and social work; its findings cannot establish public-sector readiness. [1]

Examine the work people absorb

Compare management reports with frontline practice. Ask colleagues to show a normal case and a difficult one, including the checking, copying, workarounds and follow-up. A dashboard may record automated completion before another team has made the result usable.

The UK report's fashion-retailer case describes staff largely abandoning an embedded AI analytics assistant because they found its answers inaccurate and the effort of supplying detail unhelpful. The extra work outweighed the convenience staff expected. [1, case study 3]

Where staff intervene, distinguish useful professional judgement from correction caused by avoidable system weakness. Preserve the professional judgement that makes the process dependable.

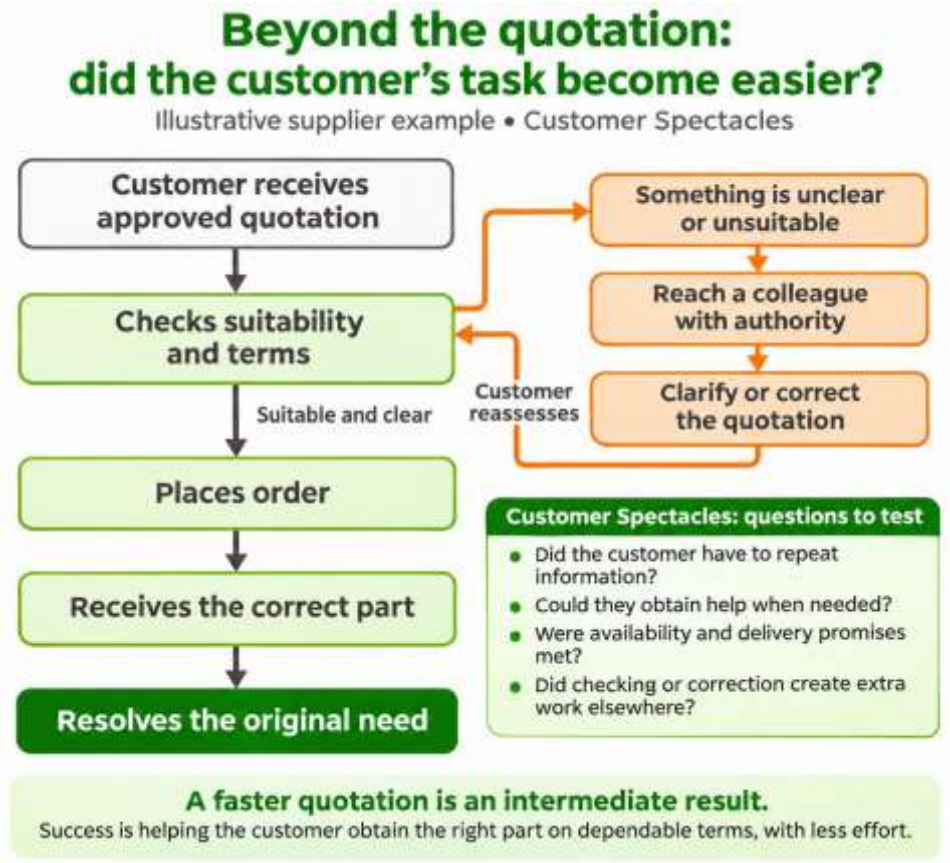

Decision question: Has the customer's task become easier, or has effort simply moved outside the business?



6 - Establish the economics of trustworthy use

What the evidence supports

SAS's account of IDC research reports that organisations with strong trustworthy-AI practices were much more likely to report strong or high returns: 62%, compared with 4% among weaker performers. The study covered 2,699 decision-makers in 28 countries across banking, insurance, life sciences and the public sector. Its trustworthiness assessment included governance, data quality and auditability. [12]

This is a relationship between broader practices and reported outcomes, not fifteen times the monetary return. Governance, management capability and data foundations may all contribute. The findings support investment in those conditions without predicting a particular business's return.

G2's October 2026 research describes a Distrust Tax within its sample of more than 1,000 B2B software buyers. It reports that 45% said oversight and correction consumed more than a quarter of expected productivity gains, while 46% reported higher operating costs than expected. The findings describe reported experience rather than an audited cost study. [4]

These sources justify including control and correction in the business case. Necessary oversight is part of delivering an acceptable outcome. Avoidable rework, unclear responsibilities and excessive manual effort are candidates for improvement. Their presence should prompt investigation before removing controls.

Where the opportunity is clearest

The evidence gives leaders a reason to test useful applications now. In the UK survey, 60% of adopters reported more efficient or flexible processes. The venue case provides a concrete example: its operations manager described faster email drafting and producing marketing materials internally that would previously have required external support. These are reported benefits, rather than audited savings. [1, case study 6]

Oak Consult's assessment is that the easiest starting points to evaluate are recurring tasks with a clear purpose, accessible information and a result that a competent person can check. Drafting routine responses, preparing information and supporting repeatable administrative work can meet those conditions. The gain comes from reducing effort or delay while maintaining quality.

For the supplier, a useful draft could help the sales team respond sooner and spend more time resolving unusual requests. That benefit can be tested while a colleague retains approval. The business does not have to wait for fully autonomous quotation handling to gain value.

The opportunity is strongest when someone owns the outcome and can put the released capacity to work. Keep necessary review in the cost assessment. Test whether the improvement survives normal volumes and difficult cases before extending it.

Follow the benefit into the business

In the UK survey, 71% of adopters selected improved workforce productivity from a benefits list. In a separate question about overall employee productivity, 50% reported an increase. The questions capture different judgements and should not be treated as interchangeable. Meanwhile, 24% reported lower operating costs, 64% no cost impact and 9% higher costs. No cost impact does not exclude benefits in quality, capacity, revenue or employee experience. [1]

The managerial task is to identify what happens to any released capacity. A service team might clear a backlog, reduce overtime or respond more effectively to customers. A commercial team might use the time to qualify enquiries. Each outcome requires an owner and evidence. Saved minutes fragmented across a working day may not be available for a different task without changes to the workflow.

An illustrative cost assessment

Assume a supplier processes 400 routine requests monthly. The existing process takes 20 minutes per request. The proposed workflow takes four minutes of staff handling and five minutes of review per request, with a further 12 hours of monthly exception handling. These are illustrative planning assumptions, not research findings or a forecast for a client.

Monthly measure	Existing process	Proposed workflow
Staff handling time	133.3 hours	26.7 hours
Review time	Included in current handling baseline	33.3 hours
Additional exception handling	Included in current handling baseline	12.0 hours
Total staff time	**133.3 hours**	**72.0 hours**
Capacity potentially released	—	**61.3 hours**

At an assumed loaded staff cost of £30 an hour, the released capacity has a planning value of approximately £1,840 a month. Suppose ongoing tools and assurance cost £700 monthly and an £8,000 setup investment is allocated over 24 months, adding approximately £333 monthly. The indicative capacity value after those costs is about £807 a month.

This is not £807 of cash savings. Salaries may remain unchanged, implementation may consume other staff time and the business may fail to use the released capacity. Validate the baseline, avoid double counting review effort and include any costs omitted from this simplified example. A cash case needs an identifiable avoided expense or reduction in expenditure. A growth case needs evidence that the extra capacity produces profitable demand.

Govern total exposure

A lower unit price does not guarantee a lower workflow cost. Volume, retries, long inputs, additional tools and connected actions can change the total bill. Establish a normal cost per acceptable completed outcome, examine its variation and set budgets and escalation thresholds.

Include shared platform costs and duplicate subscriptions. Examine what happens when demand increases, a difficult case triggers repeated attempts or a supplier changes its charging model. A hard spending limit needs a service plan for work left incomplete.

Define the conditions for continuing before scaling. Performance, customer outcomes, cost and control effectiveness all matter. Confirm that the result remains worthwhile under normal operating support.

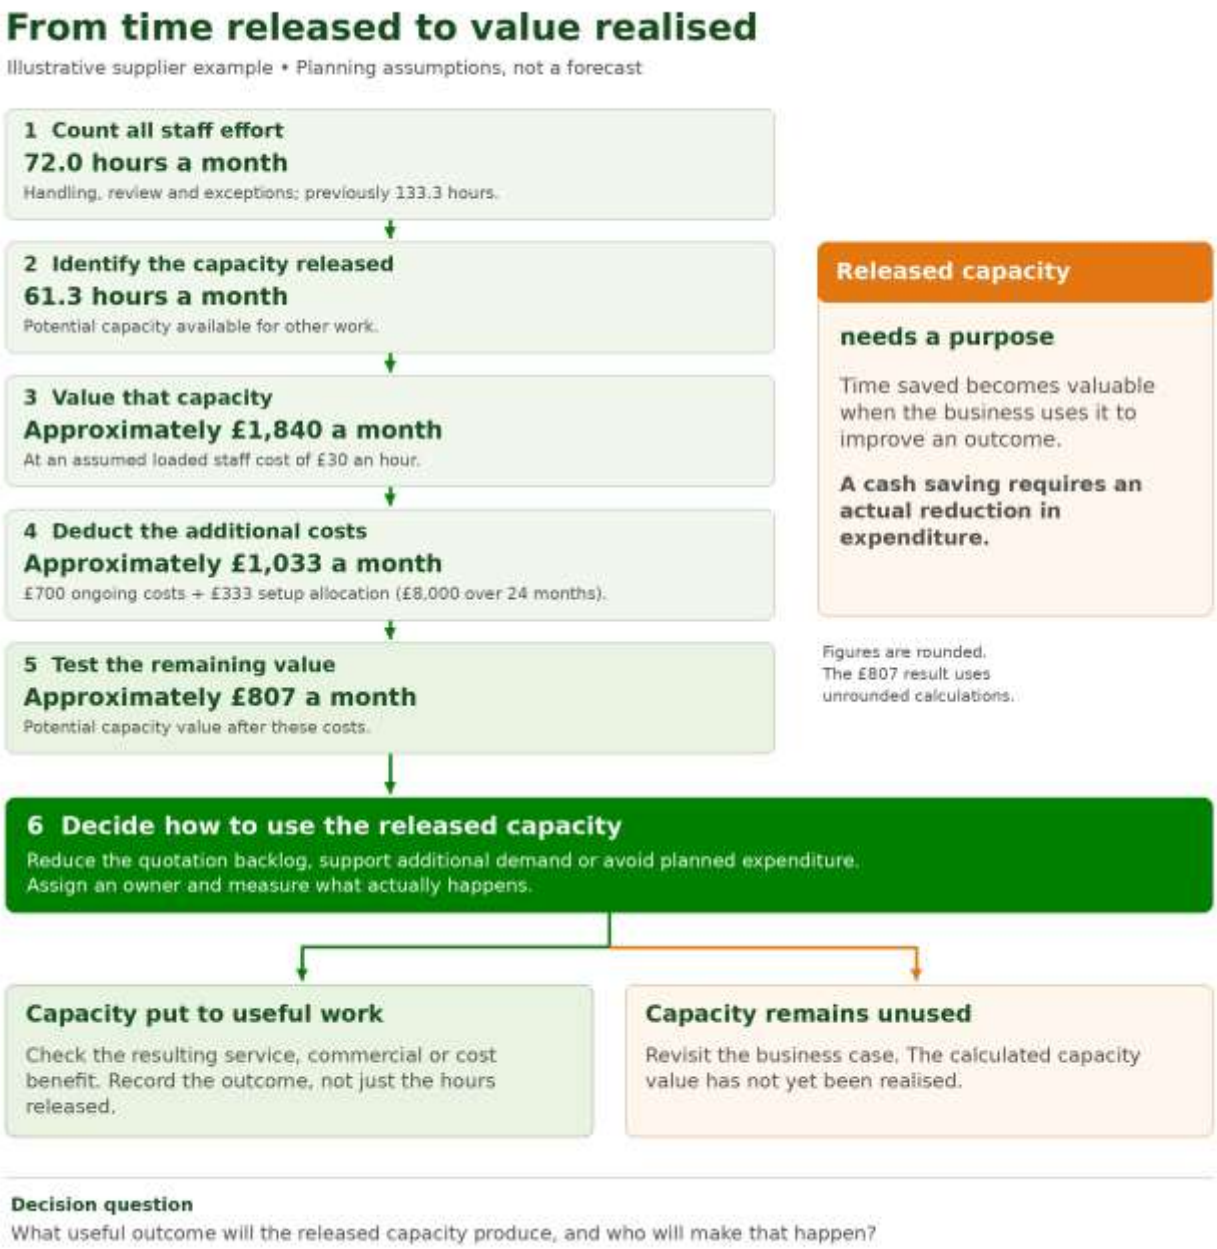

Decision question: What useful outcome will the released capacity produce, and who will make that happen?



7 - A practical first 90 days

Use the first 90 days to establish what works in one workflow and decide the next step. Adapt the pace to its risk, access and dependencies. More extensive assurance or certification may require a longer programme.

Period	Work	Evidence for the next decision
Days 1 to 15	Identify existing use, owners, sensitive information and immediate exposure	Initial register and selection of one important workflow
Days 16 to 30	Observe current work, establish the baseline and define permitted activity	Agreed purpose, authority and test plan
Days 31 to 60	Test representative cases, controls, intervention and cost	Results including failures, exceptions and customer effects
Days 61 to 90	Address findings and review remaining uncertainty	Decision to continue, improve, restrict, expand or stop

Urgent exposure should be addressed when discovered. The sequence is a planning guide, not a reason to leave an unsafe permission active until its scheduled review.

Build a decision brief

The owner should explain the objective, present performance, authorised activity, main findings and material uncertainty. Show the options and a recommendation. Record the person authorising the decision, any conditions and the next review date.

The supplier decision after 90 days

Suppose the trial shows that drafts help colleagues prepare responses, but conflicting account terms make some quotations unreliable. The illustrative decision brief could then read:

- Continue drafting for routine replacement-part enquiries, with a colleague approving each quotation.
- Keep automatic sending and stock reservation disabled.
- Give the commercial operations manager responsibility for resolving conflicting prices and account terms.
- Use released capacity to reduce the enquiry backlog, and measure response time, review effort and quotation corrections.
- Review the results after a further 30 days. Consider automatic sending only for a defined category once data quality and exception routing meet the agreed tests.

The commercial director owns this decision. The cost figures in section 6 remain planning assumptions until actual effort and outcomes are measured. This gives the supplier a useful next step while keeping the expansion decision explicit.

Make the change usable

Before extending the workflow, show affected colleagues what has changed, what remains their responsibility and when to seek help. Practise routine and difficult cases, including what happens when the system is unavailable. Give people a clear route to report problems and assign someone to resolve them.

Where the change affects customers, explain what they need to do differently, what they can expect and how to reach someone with authority. Test those explanations with customers and frontline colleagues. Include their preparation and support needs in the implementation plan and cost assessment.

Give leadership a useful dashboard

Leadership question	Useful evidence	Response when evidence weakens
Do we know what is operating?	Register coverage and checks for unrecorded use	Investigate gaps and assign ownership
Is authority appropriate?	Current permissions and tested boundaries	Restrict unsupported actions
Are results acceptable?	Quality samples, serious failures and customer outcomes	Review affected work and correct causes
Can people intervene?	Recovery exercises, review capacity and escalation outcomes	Restore effective intervention
Is the use worthwhile?	Full cost, realised benefit and trend against baseline	Redesign, limit or discontinue

Agree the sampling method and action thresholds for the workflow. Training completion and policy coverage can help show preparation, but neither establishes that a consequential action is controlled. Leadership reporting should make unresolved decisions visible.

Decision question: What are we authorising next, who owns the conditions and when will we review the result?



8 - Prepare for change over the next year

The next twelve months should be managed through explicit review triggers. A model or platform update, new connector, expanded customer group or changed charging basis can alter an application that previously performed acceptably. Treat a material change as a reason to reassess affected evidence.

Maintain a short record of confirmed requirements and developments to monitor. Separate the regulatory dates in Appendix D from possible supplier capabilities and market forecasts.

Supplier dependence deserves a practical test. Establish how information and action records can be retrieved, what would happen during an outage and whether the workflow can continue or be suspended safely. Examine portability through an exercise rather than a contractual phrase alone.

As agents become more connected, review combined authority, shared data and accumulated commitments. Changes that look small to individual teams can alter the overall exposure. Keep a business owner responsible for the result across the connected process.

Retain the capability to make independent judgements. Staff need opportunities to understand the underlying work, recognise unusual cases and recover service. Role redesign should consider how less experienced colleagues develop those skills when routine tasks are increasingly automated.

Set review frequency according to consequence and rate of change. Every material permission should have an owner, a reason and a review condition.

Decision question: Which change would invalidate today’s evidence, and who would notice it?



9 - The next decision

Choose one workflow where a better result would matter. Establish its current performance and test a specific improvement. Decide what the system may do, who owns the outcome and how the business will use the capacity released.

For the supplier, that means continuing useful drafting, fixing conflicting account data and earning the case for automatic sending. Another business may justify expansion sooner. The value lies in making a clear decision from the evidence and following it through. The Trust Horizon can expand as performance earns confidence. A useful limited application is also a result worth keeping.

Decision question: Which workflow is worth testing first, and who will own the decision?

How Oak Consult can help

Oak CORE provides independent judgement and practical challenge where an important issue needs clearer evidence, ownership or follow-through. An AI-related workflow can be examined through the CORE Check-Up, starting with a specific business question and ending with a recommendation on what should happen next.

The work clarifies the intended outcome, examines operating reality and prioritises changes. We also identify where technical testing, legal interpretation or formal certification requires specialist expertise. To discuss whether a Check-Up would be useful, arrange a Fit Conversation with Mark Conway through Oak CORE.





Appendix A - Record AI use and authority

Complete one record for each material workflow. Group related low-risk uses where sensible, while keeping responsibility and access clear. Store sensitive details securely.

Purpose and ownership

Record the workflow, the business or customer problem, the intended outcome and its accountable owner. Identify the operating team, technical owner and any specialist reviewers.

Systems and information

Record the tools, suppliers, connected systems and relevant configurations. Identify data categories, authoritative sources, access permissions and important retention or contractual conditions. Include AI features embedded in existing software.

Permitted activity

Write down what the system may recommend, change, communicate or commit. State exclusions, approval requirements and aggregate limits. Record which people or customers could be affected and the consequences of an incorrect action.

Evidence and intervention

Identify the evaluation record, quality measures, monitoring arrangements and material unresolved failures. Name the person able to interrupt the workflow, revoke access and organise recovery. Record how incomplete work and affected customers will be handled.

Value and review

Record the baseline, total operating cost, realised outcome and use of released capacity. Include the authorisation date, next review and changes that trigger reassessment.

Use the following questions to test the record:

- Can the owner explain the authority without referring only to a product name?
- Does the evidence cover the cases with the most serious consequences?
- Can the relevant people demonstrate intervention and recovery?
- Is the claimed benefit visible in the business or customer outcome?
- Is the next decision clear if a condition is no longer met?



Appendix B - Assess one workflow using A.B.C.

Use this assessment to prepare a decision. Record evidence and uncertainty in ordinary language; avoid a numerical score that conceals a serious unresolved issue.

AI Assurance

Describe the intended outcome and the cases tested. Record acceptable performance, serious failure types, exclusions and evidence of continuing reliability. Explain how source accuracy, missing information and changes to the system are handled.

Ask which conclusions depend on assumptions, which parts were independently checked and what evidence would change the recommendation.

Built-in Human Controls

Describe the owner, permissions, approval boundaries and escalation process. Identify the reviewers and assess their knowledge, available time and authority. Record the outcome of an intervention or recovery exercise, including work already in progress across connected systems.

Ask what happens when the designated person is absent and whether repeated small actions can exceed an overall limit.

Compliance

Identify applicable duties and customer commitments, who interpreted them and the evidence retained. Record supplier terms and any unresolved specialist questions. Distinguish the organisation’s chosen policy from external requirements.

Recommendation and authorisation

Record one of the following decisions and its reason:

- Continue within the current boundary.
- Improve specified foundations or controls before reconsideration.
- Restrict particular actions, data or customer groups.
- Expand a defined activity because the evidence supports it.
- Stop the use and manage affected work.

Add conditions, responsible people, review date and the evidence required for the next decision. Where investigation is incomplete, identify the missing work and the interim boundary.



Appendix C - Check whether the benefit arrived

Establish the baseline

Choose a representative period and describe volume, case complexity, staff effort, quality, customer outcomes and relevant expenditure. Include work transferred to other teams. Record normal variation and the limitations of the available measures.

Define the proposed improvement

State the expected change and how it will be measured. Identify any assumptions about demand, staffing, supplier performance or customer behaviour. Explain how checking and exception handling will be counted.

Trace released capacity

Name the person responsible for using the time or capacity released. Specify the intended outcome, such as clearing a backlog, improving response, increasing profitable throughput or avoiding external expenditure. Distinguish that outcome from a cash saving.

Count the full cost

Include implementation, integration, training, licences, usage charges, assurance, review, correction and recovery. Allocate shared costs consistently. Test the effect of higher volume, more difficult cases and changes to supplier charges.

Review customer and market effects

Check whether customers complete their task with less effort and acceptable accuracy. Examine repeat contact, complaints and difficult cases. Ask whether the offer remains valuable as customer alternatives and expectations change.

Decide what happens next

Compare the measured outcome with the baseline and agreed conditions. Explain differences and plausible alternative causes. Record whether the use should continue, improve, be restricted, expand or stop.



Appendix D - Selected regulatory reference points

Status checked on 30 October 2026. This is an orientation for leadership review, not a complete legal checklist. Establish scope, organisational role, exemptions and transitional provisions with appropriate advice.

European Union

The Commission's current AI Act timetable distinguishes several stages. Prohibited practices and AI literacy requirements began applying on 2 February 2025. Governance rules and GPAI obligations began on 2 August 2025, subject to transitional provisions. Transparency requirements for certain systems and relevant enforcement arrangements applied from 2 August 2026. The current high-risk dates are 2 December 2027 for Annex III uses and 2 August 2028 for relevant product-embedded systems. [10]

A UK address alone does not determine EU scope. Examine the organisation's role and the deployment concerned. A later high-risk deadline does not suspend duties already applicable.

United Kingdom

UK data protection remains relevant where personal information is processed. The Data (Use and Access) Act reforms broaden circumstances for certain significant automated decisions while retaining safeguards. Consult current ICO guidance and establish the treatment of sensitive information and meaningful human involvement. [9]

In the public sector, apply the Algorithmic Transparency Recording Standard's current mandatory scope and exemptions. It does not impose an identical requirement on every public body or every algorithmic tool. Check applicable sector and service duties separately. [11]

Standards and assurance

ISO/IEC 42001 concerns an AI management system. NIST's AI Risk Management Framework is a voluntary risk-management resource. Use them according to the organisation's needs and obligations; neither a reference to a standard nor use of A.B.C. establishes certification or compliance. [5, 6]

Sources and evidence notes

Sources were checked on 30 October 2026. Numbered references connect material findings to their source. The paper's recommended controls, worksheets and decisions are Oak Consult's analysis. No source is presented as an independent validation of A.B.C. or Oak CORE.

[1] Department for Business, Innovation, Science and Trade. Artificial intelligence adoption research 2026. Published 29 October 2026. The publication page links to the full report and data tables.

Open research

The study used 3,503 telephone interviews between 7 January and 21 April 2026, 50 follow-up interviews and seven site visits. The quantitative scope is private-sector businesses with at least five employees in specified sectors; public administration, education and health/social work were among exclusions. Adoption includes piloting. Results are weighted by size and sector. The reported response rate is 5%, and impacts are self-reported. See methodology and technical annex; figure 27 for oversight; case studies 3, 4 and 6 for the examples discussed. These cases are research participants, not Oak clients.

[2] Deloitte Insights. Business and IT leaders report AI agents are scaling faster than their guardrails. Andy Bayiates, 24 April 2026, drawing on State of AI in the Enterprise 2026.

Open research

Survey of 3,235 leaders involved in organisational AI programmes across 24 countries. The paper uses self-reported governance maturity and expectations, with no inference that the sample represents all UK businesses.

[3] Pan and colleagues. Measuring Agents in Production. arXiv 2512.04123, version 4, revised 4 June 2026.

Open research

Current abstract reports 20 case studies and 86 deployed-systems practitioners. Earlier versions described a wider 306-person survey; this paper uses the current production-focused description. Observed practices do not establish a universal deployment rate.

[4] G2. New G2 Research AI Agents Have Moved From Pilot to Production Now Comes the Distrust Tax. Tim Sanders, 14 October 2026; public findings from Crossing the Agentic Chasm.

Open research

Based on more than 1,000 B2B software buyers and review data. This paper uses the public findings and attributes Distrust Tax to G2. Expected productivity and reported cost experiences are not independently audited outcomes.

[5] NIST. AI Risk Management Framework.

Open research

Reference for voluntary risk-management work; no claim of a complete formal mapping to A.B.C.

[6] ISO. ISO IEC 42001 2023 Artificial intelligence management system. Public standard overview.

Open research

The public overview establishes the standard's subject and purpose. This paper does not claim a clause-level assessment against the full standard.

[7] National Cyber Security Centre. Prompt injection is not SQL injection it may be worse. Dave Chismon, 8 December 2025.

Open research

Used for the security distinction between instructions and untrusted data. Specific implementation controls require technical assessment.

[8] NIST. New Concept Paper on Identity and Authority of Software Agents. 5 February 2026.

Open research

An initiative seeking input, not a claim that agent identity standards or solutions are complete.

[9] Information Commissioner's Office. Data protection summary of changes under the Data Use and Access Act 2025, including automated decision-making; related commencement statement, 5 February 2026.

Open research

Open research

Use current guidance for the specific processing. This paper offers leadership orientation rather than legal interpretation.

[10] European Commission. AI Act and The enforcement framework of the AI Act. Current implementation and enforcement summaries.

Open research

Open research

Dates reflect the current Commission summaries, including changes to high-risk application dates. Check the legal text, scope and transitional arrangements before relying on a deadline.

[11] Government Digital Service. Making the Algorithmic Transparency Recording Standard mandatory across government. 8 May 2025, with links to the scope and exemptions policy.

Open research

[12] SAS. Organisations with trustworthy AI practices are 15 times more likely to see strong ROI per study findings. 1 September 2026; public account of IDC research for the Data and AI Impact Report The New Economics of Trust.

Open research

Survey of 2,699 decision-makers in 28 countries across banking, insurance, life sciences and the public sector. The 62% versus 4% comparison concerns reported strong/high ROI. It is not fifteen times the financial return or evidence that governance alone caused the difference.

Scope of this edition

This is a synthesis of selected secondary evidence with practical analysis. No original survey, client outcome measurement or formal evaluation of A.B.C. was undertaken for this edition. The supplier workflow and cost assessment are explicitly illustrative. Source findings from different populations and periods have not been combined into a single market estimate.

