



Using Raven

A help manual for leaders and investors on understanding how to interpret and put into action the results of a morriganAI Raven Report

Version 1-0, May 1, 2026

How to use this manual

A Raven Report tells you how the work inside an organization actually moves – whose ideas show up in finished documents, where AI is doing the work, and where the business depends more on a single person or a single tool than the org chart would suggest. That picture is what morriganAI calls an organization's AI Identity. This manual is your guide to reading that picture and using it well.

The manual is comprehensive, but you do not need to read it cover to cover. Each section is written so you can read it on its own. The table below shows where to start based on the role you hold and the question you came here to answer.

| If you are a... | Start with these sections | Then read |
|-----------------------------------|--|--|
| CEO or COO | Part 1 (value), Part 3 (how to read findings), Part 4 (limits) | Part 6 (why the data is safe to use) |
| CHRO or HR leader | Part 2 (what is in the report), Part 3 (reading the Contributors table), Part 4 (what needs corroboration) | Part 6 (privacy by design) |
| Banker or private equity investor | Part 1 (snapshot vs. trend value), Part 4 (conclusions Raven cannot support), Part 6 (data handling) | Part 3 (decisions Raven can inform) |
| Anyone who wants the full picture | Read straight through | Use the glossary in the back to look up any term |

A note on language

This manual is written in plain English so it can be shared across an organization. Where we use a technical term, we define it the first time it appears and again in the glossary. We have done our best to keep the writing accurate without making it dense.

Contents

| | |
|---|----|
| How to use this manual | 2 |
| Contents..... | 3 |
| What is AI Identity, and what is Raven looking at? | 6 |
| Part 1 · The value a Raven Report can offer you | 7 |
| 1.1 The problem Raven was built to solve..... | 7 |
| 1.2 The snapshot: one Raven Report at one moment in time..... | 7 |
| 1.3 The trend line: continual reports over time | 8 |
| 1.4 What each role tends to use a Raven Report for..... | 8 |
| For the CEO | 8 |
| For the COO..... | 9 |
| For the CHRO and HR leaders | 9 |
| For the banker..... | 9 |
| For the private equity investor | 9 |
| 1.5 Three case examples of Raven in use | 10 |
| Case one: a COO finding the silent expert | 10 |
| Case two: a CHRO opening a conversation about AI | 10 |
| Case three: a private equity investor pricing operational risk..... | 10 |
| Part 2 · What is in a Raven Report | 12 |
| 2.1 The core sections of a Raven Report | 12 |
| 2.2 The three things every reader should look at first | 13 |
| 2.3 The deeper layers..... | 13 |
| 2.4 The Six Business Phases | 13 |
| Part 3 · How to read the information in a Raven Report..... | 15 |
| 3.1 Reading the Executive Summary and the coverage sentence..... | 15 |
| 3.2 Reading the Contributors table..... | 15 |
| Aggregate Weighted Influence | 15 |
| AI Usage | 16 |
| 3.3 Reading the Phase Analysis..... | 16 |
| Factual versus Conceptual coverage..... | 16 |
| Last Step Similarity, Influence Distribution, and Origin Depth | 17 |
| 3.4 Reading the Key Findings and the Risk Summary | 17 |
| 3.5 Decisions Raven is well-suited to inform | 18 |
| Part 4 · Where to slow down: limits, corroboration, and conclusions Raven cannot support..... | 19 |

| | |
|--|----|
| 4.1 Coverage and what it tells you to do | 19 |
| 4.2 Findings that always deserve corroboration | 19 |
| Anything that names an individual | 19 |
| Findings concentrated in a phase with low coverage..... | 19 |
| High AI Usage on a single contributor..... | 20 |
| Below-threshold or fragmentation findings | 20 |
| Quantified risk numbers in the Risk Summary..... | 20 |
| 4.3 Conclusions Raven cannot support..... | 20 |
| 4.4 Worked examples of corroboration..... | 21 |
| Scenario A: high AI Usage on a senior contributor | 21 |
| Scenario B: process concentration risk in the Do phase..... | 21 |
| Scenario C: a fragmentation finding that turns out to be a non-issue | 22 |
| Scenario D: low coverage on a contributor-specific finding | 22 |
| Scenario E: a leader who looks invisible | 22 |
| 4.5 Connection scope as a source of bias | 23 |
| Part 5 · How a Raven Report is built | 25 |
| 5.1 Where the data comes from, and who decides what Raven can see..... | 25 |
| 5.2 Multiple connections and what each one is for..... | 26 |
| Shared-drive connections | 26 |
| Individual-account connections | 26 |
| 5.3 The customer’s responsibility for connection scope | 27 |
| 5.4 What gets read and what does not..... | 27 |
| 5.5 Pairs and Vectors | 28 |
| 5.6 From Vectors to findings..... | 28 |
| 5.7 How the prose sections of the report get written | 28 |
| 5.8 Coverage and Run Cost | 28 |
| Part 6 · Why it is safe to use Raven..... | 29 |
| 6.1 The starting principle: privacy by design | 29 |
| 6.2 Connection scope: control lives in your identity system | 29 |
| 6.3 The two files that stay on your storage | 30 |
| 6.4 What is stored in morriganAI’s systems and how it is protected | 31 |
| 6.5 Why retaining Raven Vectors is the safe choice – and why it is the useful one..... | 32 |
| 6.6 Avoiding surveillance by design | 32 |
| 6.7 What happens if you end your subscription or disconnect a provider..... | 32 |
| Quick reference cards | 34 |

| | |
|--|----|
| For the CEO and COO | 34 |
| For the CHRO and HR leaders | 34 |
| For the banker and private equity investor | 35 |
| Frequently asked questions | 36 |
| Questions from CEOs and COOs | 36 |
| How long does a Raven Report take to produce? | 36 |
| How often should we run a Raven Report? | 36 |
| What does it cost to run?..... | 36 |
| Questions from CHROs and HR leaders | 36 |
| Will employees know they are being included? | 36 |
| Could a Raven Report be used to discipline an employee? | 36 |
| What about contributors who do important work that does not show up in documents?..... | 37 |
| Questions from bankers and private equity investors | 37 |
| Is the report defensible in a diligence file?..... | 37 |
| Can the same Raven Report be shared with management? | 37 |
| Does the report tell us if AI usage is creating financial liability? | 37 |
| Questions from security and IT teams | 37 |
| What permissions does Raven need? | 37 |
| How do we change what Raven can see, or stop it from seeing something? | 37 |
| Can we connect more than one account? | 38 |
| Where is the data stored? | 38 |
| What happens if morriganAI is breached? | 38 |
| Is the data used to train AI models? | 38 |
| Can we get a copy of all the data on us? | 38 |
| Glossary..... | 39 |
| A closing note..... | 42 |

What is AI Identity, and what is Raven looking at?

Every organization has an identity. It is the way the company actually operates: who knows what, who decides what, how customers are won, how work gets done, how money gets collected, and how the business is governed. That identity is built by the people who work there and the processes they follow, and it is what stakeholders trust when they trust the company.

AI Identity is what that identity becomes once AI is woven into the work done by that organization. AI tools change who is doing what. They change which ideas are persistent and which are disposable. They change how fast new content shows up and how carefully it has been reviewed. Sometimes AI changes an organization's identity in helpful ways. Sometimes it changes the identity in ways nobody noticed and nobody chose.

Raven is the tool morriganAI built to make AI Identity visible. It does this by quietly studying the documents your organization already produces and looking at how ideas flow from older documents to newer ones. From those flows it can see who the steady contributors are, where AI authorship is heavy, where work is concentrated in a small number of people, and where the work seems disconnected. Raven does NOT watch employees. It does not read keystrokes. It looks at finished artifacts – the documents, presentations, and spreadsheets you would already share with a partner – and reasons backward from there.

A short example

Imagine a small services firm where one person on the operations team is the unspoken expert on pricing. Their fingerprints are on every quote that goes out, even when their name is not on the cover page. A Raven Report can see that pattern in the documents themselves and call it out as a process concentration risk – without ever installing anything on a single employee's computer.

Part 1 • The value a Raven Report can offer you

Raven Reports are useful in two distinct ways. The first is the snapshot – a single, point-in-time picture of how the organization is actually operating right now. The second is the trend line – a series of reports run over time that show how that picture is changing. Most leaders need both at different moments, and the value of each is different enough to be worth treating separately.

1.1 The problem Raven was built to solve

AI is changing the inside of organizations faster than the outside is changing. A team can adopt a new AI tool in a week. The org chart, the policies, and the descriptions of how work gets done usually take much longer to catch up. The result is a growing gap between what an organization says about itself and how it actually runs.

That gap matters because it is where operational risk lives. If the company depends on an AI tool that nobody is reviewing, you have a control risk you cannot see in any handbook. If a single contributor's ideas underpin most of the company's recent work, you have a key-person risk that no resume can flag. If two teams are using completely different AI tools and producing disconnected work, you have an integration problem that no productivity dashboard will catch.

Raven was built to close that gap. It gives leaders a way to see the inside of the organization through the artifacts the organization already produces – the documents, presentations, and spreadsheets that travel through cloud storage every day.

1.2 The snapshot: one Raven Report at one moment in time

A single Raven Report is a snapshot. It tells you what the AI Identity of the organization looks like as of the date the analysis was run. That is exactly what you want when you need a credible read on operational reality at a specific moment, especially when something important is about to be decided.

Common situations where the snapshot is the right tool:

- **Due diligence ahead of an acquisition.** A snapshot lets a banker or private equity investor see how concentrated the target's real work is, how heavily AI is being used in each phase of the business, and where documentation is thin.
- **Pre-investment review for a portfolio company.** A board or investor can ask for a snapshot to confirm what management says about how the team is operating before approving a new investment thesis or transformation plan.
- **Transformation planning.** Before launching a redesign, leaders can take a snapshot to know which roles, processes, and tools the current operation actually depends on. That makes the redesign less of a guess.

- **Post-incident review.** After a meaningful operational miss, a snapshot can show whether the incident was a one-off or a symptom of a deeper concentration or governance gap.
- **Onboarding a new executive.** A new CEO, COO, or CHRO can use a snapshot to understand who the real load-bearing contributors are, faster than they could by meeting everyone.

A snapshot has limits. It is one moment in time. It cannot tell you whether the situation is improving, getting worse, or new. For that, you need the trend line.

1.3 The trend line: continual reports over time

When Raven runs on a recurring schedule, each report is compared to the prior one. That comparison is where some of the most useful information lives, because operational risk is rarely about a single moment – it is about the direction things are moving.

A trend line lets you see questions like:

- Is the share of work attributed to AI authorship rising or falling, and in which phases of the business?
- Is the dependence on any one contributor going up or coming down?
- Are people who joined a year ago growing into more central roles when contributing to the work of the organization, or staying on the edge?
- After we set a new AI policy, did AI authorship in the affected phase actually change?
- After we hired a senior leader, are their ideas showing up in the work – or are they still leaning on the same handful of long-tenured contributors?
- Did the integration of two teams after an acquisition actually happen as planned, or are the two teams still producing disconnected work?

The trend line also reveals shifts that no single report could prove. If AI prevalence has been climbing in the same phase for three quarters in a row, that is a meaningful pattern. A single quarter on its own would not tell you the same thing.

1.4 What each role tends to use a Raven Report for

Different leaders bring different questions to the same report. Below is a short summary of how each audience commonly uses Raven, the issues or concerns most central to their role, or the types of insights a Raven Report can offer.

For the CEO

The CEO is usually trying to confirm or challenge a story they have been told. Raven gives the CEO an independent view of how the work actually flows, which contributors are quietly central, and which AI uses might be running ahead of the company's policies. Over time, the trend line shows whether the company's investments – in people, in tools, in process redesign – are landing.

For the COO

The COO uses Raven to find the parts of the operation most exposed to a single point of failure. That can be a person, a tool, or a phase of the business. The Phase Analysis section is especially useful here because it shows how the operation is performing against the six business phases that every organization runs through.

For the CHRO and HR leaders

For HR, Raven is a way to understand the value individuals bring to the work without surveilling them. The Contributors table shows whose ideas keep showing up in finished documents and roughly how much AI those individuals are leaning on. That helps HR identify who to invest in, where cross-training would reduce risk, and where roles are quietly changing because of AI. Importantly, Raven names contributors but does not evaluate them – it is a starting point for a conversation, not a performance review.

For the banker

A banker preparing a deal can use a Raven snapshot to add operational evidence to the diligence file. The report gives a defensible, repeatable view of where the business is concentrated and how heavily AI is being used. That is increasingly relevant when buyers and lenders want to understand whether the company's recent productivity gains are durable or reliant on tools that are not being governed.

For the private equity investor

A PE investor can use a snapshot before close and a trend line after close. Pre-close, the snapshot supports the investment thesis and surfaces operational risks that may need to be priced in or addressed in the first hundred days. Post-close, the trend line is a low-effort way to confirm that operating-partner work is producing visible change and that AI usage is moving in the direction the thesis assumed.

A note on what gets connected and included in a Raven analysis

Later sections discuss how Raven ensures safety by always giving the organization final say over what can be seen by the Raven tool. It is the organization that chooses how much data is visible to Raven based on the connections the organization provides to Raven. This is why it is important to understand that the quality of any Raven Report depends on what the customer chooses to connect. Connecting only storage locations that are broadly shared by an organization's teams gives one picture. Optionally connecting individual leaders' work-account drives can give a richer picture, because many of the most consequential ideas in an organization originate in a leader's personal work area before they appear in a shared document. This is a customer choice, not a requirement, and it is discussed in detail in Part 5.

1.5 Three case examples of Raven in use

The case examples below are short, illustrative scenarios drawn from the kinds of patterns Raven typically surfaces. They are not real customers. Each one shows a different combination of audience and use case so you can see how the report supports the decision behind it.

Case one: a COO finding the silent expert

A regional services firm runs a snapshot at the request of its new COO. The Contributors table shows what the COO suspected and what the org chart did not reveal: a single mid-level operations analyst has the highest Aggregate Weighted Influence in the entire company. Her ideas are showing up across pricing, proposals, and the playbook that operations follows on every job.

No one is doing anything wrong. She is good at her job and willing to share her thinking. The risk is what happens if she leaves. The report does not say to give her a raise – that is a leadership decision – but it does give the COO an unambiguous picture of what would go missing if she did. The COO uses the report to fund a documentation project and pair her with a more junior colleague who can carry her playbook forward.

Case two: a CHRO opening a conversation about AI

A mid-sized professional services firm has heard mixed signals about how AI is being used inside the firm. The CHRO commissions a snapshot. The Contributors table shows three people in the same practice with high AI Usage. The Phase Analysis shows that AI authorship is concentrated in the Tell phase – the public-facing communications work.

The CHRO does not treat the report as a verdict. She invites the three contributors and their manager into a conversation. Two of the three turn out to be careful, fluent users of AI who have learned to draft and revise efficiently and who review their work closely. The third is leaning on AI to produce client-ready material that she has not been reviewing as carefully as the firm would expect. That is a coaching conversation, not a discipline conversation, and the manager is well-equipped to have it because the report opened the door.

Case three: a private equity investor pricing operational risk

A PE firm is in diligence on a software-enabled services business. Management's pitch is that recent productivity gains are durable. The PE firm asks for a Raven snapshot as part of operational diligence.

The snapshot shows two things. First, AI authorship is high in the Do phase – the work where the company delivers on commitments – and the Phase Analysis suggests that work is iterating closely on the most recent prior work without much breadth of input. Second, two contributors hold an unusually large share of Aggregate Weighted Influence in the Agree and Govern phases. The Risk Summary, written with industry context, places these patterns next to known examples of operational disruption when key staff leave or when AI-produced work is not reviewed.

The PE firm does not pull the deal. It does, however, build two things into the term sheet: a transition-services arrangement that retains the two key contributors through the first year, and an operating-partner workplan to introduce a sampling review of AI-produced deliverables. A trend-line subscription is set up to track whether those interventions move the metrics in the right direction over the first eighteen months of the hold.

Common thread across the three cases

In every case, the report did not make the decision. It helped leaders arrive at a better-informed decision. The discipline of the Raven Report is to surface patterns clearly. The discipline of the leader is to bring judgment to those patterns and to act with care, especially when individuals are named.

Part 2 · What is in a Raven Report

A Raven Report is a structured analytical document. Each report is delivered on demand or periodically, as set by the user, and follows the same outline every time. Once you have read one Raven Report, you know how to read the next one. This section walks through what each part contains, what to look at first, and considerations to remember when interpreting the findings of a Raven Report.

2.1 The core sections of a Raven Report

Every Raven Report contains the same sections in the same order. Each section and its purpose is listed below.

| Section | What it is for |
|--------------------|--|
| Executive Summary | A short, plain-language summary of the most important findings, written for a senior reader. Always notes how complete the analysis was. |
| Key Findings | Three to six numbered findings, each one actionable. These are the headlines. |
| Analysis Scope | A paragraph describing what was analyzed: which storage locations, how many artifacts, what time window, and a short description of the organization's industry context. |
| Artifact Summary | Two tables. The first shows the mix of file types analyzed. The second is the Coverage Summary, which tells you how complete the analysis was. |
| Contributors | A table listing every contributor, their AI Usage (low, moderate, or high), and their Aggregate Weighted Influence – a measure of how much their ideas show up in downstream work. |
| Phase Distribution | A table showing how the analyzed documents map to the six business phases (Gather, Tell, Agree, Do, Reward, Govern). |
| Phase Analysis | For each phase, two tables: a user influence table and an artifact characteristics table. This is where you see the operation phase by phase. |
| Risk Summary | For each Key Finding, a longer narrative that places the risk in the context of known industry examples and offers a sense of severity, probability, and remediation time. |

2.2 The three things every reader should look at first

No matter what role you hold, three sections deserve your attention before any others.

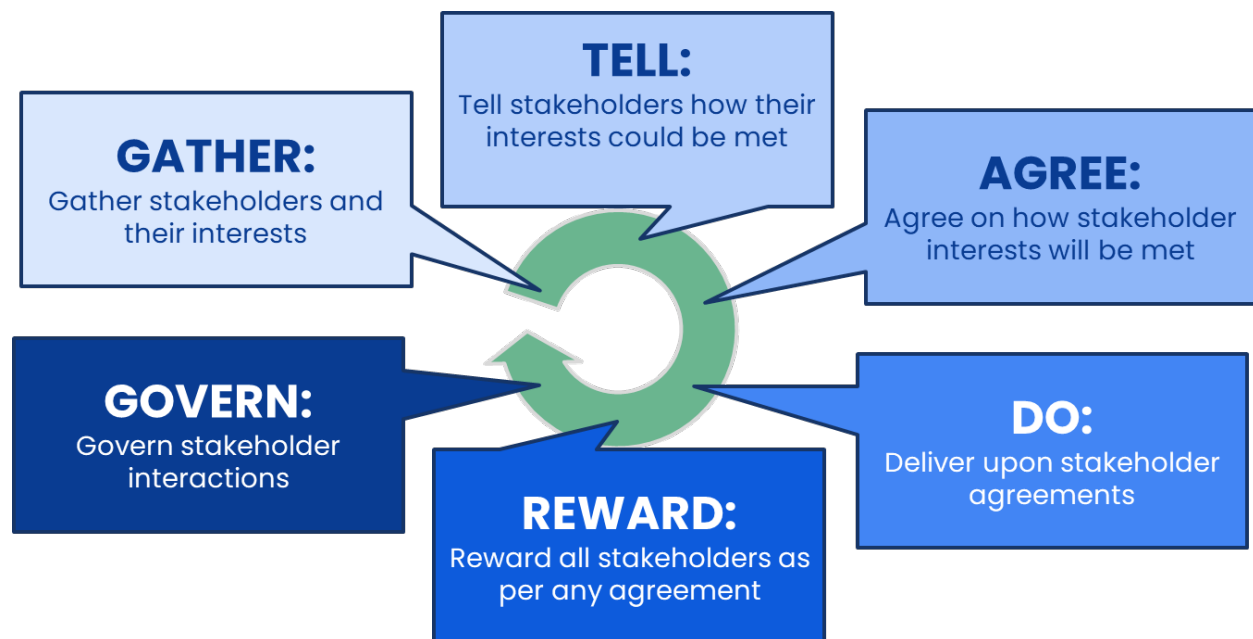
1. The Executive Summary, especially the sentence that tells you what percentage of eligible file pairs were actually analyzed. That sentence sets the confidence level for everything else in the report.
2. The Key Findings. These are the report's headlines, prioritized by significance. Each one is written to be actionable so you can see what corrective step is implied.
3. The Contributors table. This is where the people picture lives. The two columns to focus on are AI Usage (the probability that the work from a specific person is augmented by AI) and Aggregate Weighted Influence (how much that person's ideas show up in finished work).

2.3 The deeper layers

Once you have read the headlines, the rest of the report is there to support and contextualize them. The Phase Distribution and Phase Analysis sections show you the operation phase by phase. The Risk Summary section turns each Key Finding into a longer narrative with industry comparisons.

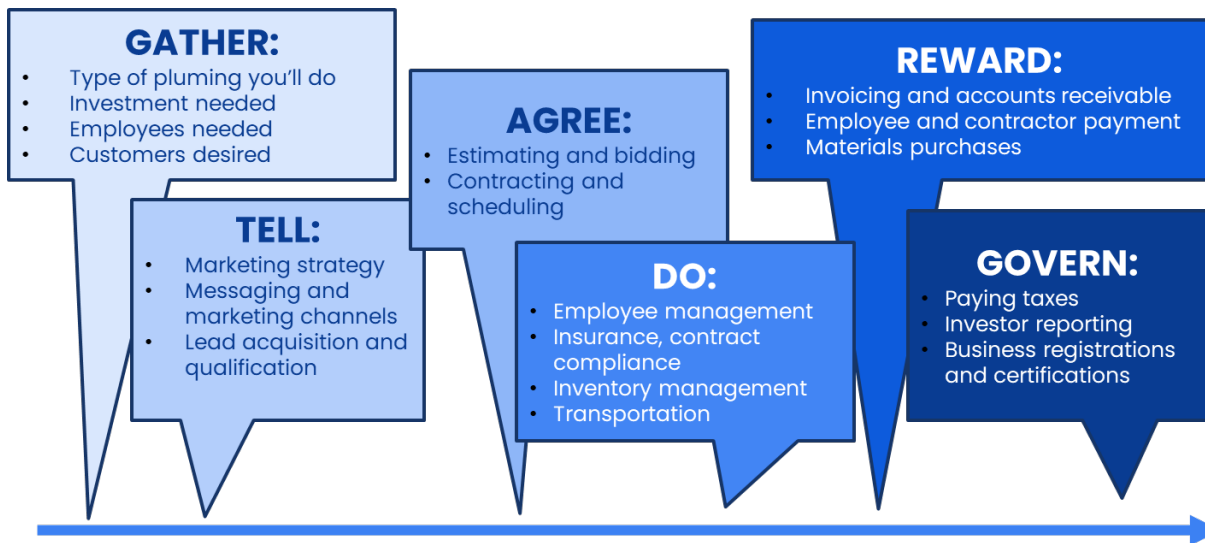
2.4 The Six Business Phases

Raven looks at every document through the lens of six business phases. Every organization, regardless of industry, moves through these six phases as it operates. A document is not assigned to just one phase; it is given a probability of belonging to each, because most documents touch more than one. The phase that scores highest is called the dominant phase, and that is the one used in cases when a document must be classified by only one business phase.



| Phase | Plain-language meaning |
|--------|---|
| Gather | Identifying, researching, and recruiting the people and organizations that have a stake in the outcome – customers, partners, hires, investors. |
| Tell | Communicating who the organization is and what it does – the marketing, the public-facing materials, the pitch. |
| Agree | Reaching a formal agreement – pricing, scenario modeling, proposals, and contracts. |
| Do | Following through on the commitments – design, deployment, delivery, and operations. |
| Reward | Collecting payment, calculating compensation, and distributing the proceeds – billing, payroll, subcontractor payments. |
| Govern | Running the organization itself – board management, security, audits, and the broader responsibilities to the community and regulators. |

As an example of how these phases apply to a traditional business, the following graphic identifies actions associated with each phase when defining the Identity of a plumbing company.



When a Raven Report says something like "AI authorship is high in the Do phase," that is shorthand for "in the work where the company is delivering on its commitments, many of the files contain large amounts of data that is being produced by AI." That is a more useful sentence than "AI usage is up 12%," because it tells you where in the business the change is happening.

Part 3 · How to read the information in a Raven Report

This section explains what each major piece of the report is actually telling you, with examples drawn from the kind of patterns Raven typically surfaces. It also lists the decisions the report is well-suited to inform.

3.1 Reading the Executive Summary and the coverage sentence

Start at the top of the Executive Summary and look for the sentence that mentions pair coverage. It will read something like, "This analysis achieved 84.3% pair coverage." That number tells you what fraction of the eligible document pairs in your time window were actually analyzed in this run, including any analysis that carried over from earlier runs.

Coverage in the high 90s means you are looking at a near-complete picture and you can read the rest of the report with normal confidence. Coverage in the 70s or 80s means the picture is partial, and you should weigh findings about smaller contributors more carefully. If coverage is low enough to materially affect attribution, the Executive Summary will include a callout paragraph telling you so directly. When you see that callout, treat the report as directional rather than precise, and consider running an additional analysis cycle before making consequential decisions.

Why coverage might be less than 100%

Raven uses an AI service to compare document pairs, and each comparison costs a small amount. Most accounts set a Run Cost Limit so a single analysis run cannot exceed a known dollar amount. If the run reaches that limit before all eligible pairs have been analyzed, the report will tell you so plainly. Raising the limit, or running an additional cycle, will fill in the gap.

3.2 Reading the Contributors table

The Contributors table is where the people picture lives. Each row is one contributor – usually a named individual, sometimes a contributor whose identity could not be resolved, in which case they appear as "Unidentified User 1," "Unidentified User 2," and so on. The two columns that matter most are described below.

Aggregate Weighted Influence

This is the headline number for "how much does this person's work show up in the finished output." It is calculated by tracing every document this person was the last modifier of, and seeing how much of their work carried forward into the documents that came after. A person with high Aggregate Weighted Influence is someone whose ideas keep showing up downstream. That is usually a sign of a senior contributor, a subject-matter expert, or a key author.

A handful of people having high Influence is normal in any organization. What deserves attention is when the influence is so concentrated in one person that their absence would noticeably change the work product. That is called process concentration risk and the Key Findings section will usually flag it directly.

AI Usage

This column tells you, directionally, how much of a contributor's recent work appears to be authored with AI assistance. It is shown as low, moderate, or high. The number behind that band is the weighted average of AI authorship across the documents the person last modified, weighted by how informationally rich each document is.

A contributor with low AI Usage is producing work that reads as primarily human-authored. A contributor with moderate AI Usage is using AI as a meaningful part of their workflow. A contributor with high AI Usage is producing work that is heavily AI-authored. None of these labels is, on its own, a problem or a virtue. What matters is whether the level of AI usage is appropriate for the work that contributor is doing and whether that level is being governed.

Important: AI Usage names patterns, not policy violations

Raven cannot tell you whether a person used a particular AI tool, or whether they followed your AI policy. It can only tell you that the documents they last modified read as more or less AI-authored. Treat the AI Usage column as the start of a conversation, not the end of one.

3.3 Reading the Phase Analysis

For each of the six phases, the Phase Analysis section gives you two tables. The first is a user influence table that shows, for each contributor, two percentages: Factual and Conceptual. The second is an artifact characteristics table that shows three percentages: Last Step Similarity, Influence Distribution, and Origin Depth. The phrases sound technical but the meanings are practical.

Factual versus Conceptual coverage

Factual measures how much of a person's specific facts, figures, and details (prices, dates, numbers, names) carried forward. Conceptual measures how much of their underlying logic, framework, or argument carried forward. The two move independently, and the difference between them is informative.

- High Factual and high Conceptual: this person's detailed work is shaping later documents, end to end.
- High Factual and low Conceptual: their numbers are being reused but the surrounding thinking is not. That can be normal (someone owns a price list, for example) or it can mean their work is being copied without being understood.

- Low Factual and high Conceptual: their framework is influencing later work, but the specifics get rewritten each time. That often describes a strategy or design lead.
- Low both: this person is contributing in ways the documents do not capture. They may still be valuable; Raven simply cannot see it.

Last Step Similarity, Influence Distribution, and Origin Depth

These three percentages describe the character of the work in a phase, not any individual person. Read them together rather than in isolation.

- Last Step Similarity tells you how closely a typical new document resembles the most recent older document in the same phase. High numbers mean the work is iterative – people are building on what came right before. Low numbers mean each new document is starting from a clean sheet.
- Influence Distribution tells you whether new work is shaped by many older documents or just a few. High numbers mean broad influence – many sources are feeding the work. Low numbers mean narrow influence – the work is leaning on a small number of sources.
- Origin Depth tells you how far back in time the chain of influence reaches. High numbers mean the phase has long roots – work being done now still carries fingerprints from documents made well in the past. Low numbers mean the phase is operating on recent material.

A phase with high Last Step Similarity, low Influence Distribution, and shallow Origin Depth is a phase running on near-term, narrow-source work. That can be efficient or it can be brittle, depending on the phase. The Risk Summary section will tell you which of these signals the report considered material.

3.4 Reading the Key Findings and the Risk Summary

The Key Findings are the headlines. The Risk Summary expands each headline into a short narrative that places the finding in the context of known industry examples and tries to give you a sense of severity, probability, and how long remediation might take.

The Key Findings are written to be actionable. Every finding implies a corrective step, even if it is not stated as a directive. The most common categories of finding are:

- Process concentration risk – too much of the work depends on one person. The implied action is cross-training, documentation, or hiring.
- AI control risk – AI is producing work that no one is reviewing closely. The implied action is a human-in-the-loop step, a sampling audit, or an updated policy.
- Fragmented document management – teams are producing work that does not connect to other teams' work. The implied action is integration, shared templates, or a shared workflow.
- Low content retention in a critical phase – important details are not being preserved across versions. The implied action is improved documentation standards or version control.

- Inconsistent concept retention – the underlying logic is not carrying forward consistently. The implied action is shared frameworks or training.
- Governance gaps – the Govern phase is not well-supported by the documents. The implied action is more deliberate documentation of board, audit, and oversight work.

3.5 Decisions Raven is well-suited to inform

The list below is not exhaustive, but it covers most of the decisions leaders use Raven Reports to support.

1. Succession and cross-training priorities. Raven shows where the business is most exposed to the loss of a single contributor and where investing in a backup is highest leverage.
2. AI governance and human-in-the-loop design. Phases with high AI authorship and weak governance signals are good candidates for a review step before AI-produced work goes out the door.
3. Documentation and process integration investments. Findings about fragmented work or low retention point directly to where shared templates and process redesign would pay off.
4. Valuation adjustments and diligence questions for an acquirer. A snapshot can support a price discussion or surface questions to ask of management before close.
5. Pre- and post-transformation comparisons. Run a snapshot before a redesign and another after. The trend line shows whether the redesign actually changed how work is done.
6. Onboarding decisions for new senior leaders. A snapshot helps a new executive identify the load-bearing contributors so they can spend their early time well.
7. Board reporting on operational risk. A trend line is a defensible way to show a board that operational concentration and AI control posture are improving (or to admit they are not).

Part 4 · Where to slow down: limits, corroboration, and conclusions Raven cannot support

A Raven Report is most useful when it is read with an honest sense of what it can and cannot prove. This section names the limits explicitly. The discipline this section asks for is not academic – it is what keeps a report from being misused.

4.1 Coverage and what it tells you to do

Coverage is the single most important context for the rest of the report. A coverage of 95% means almost every eligible document pair was analyzed. A coverage of 50% means roughly half of them were not. Coverage below the threshold morriganAI considers material is flagged in the Executive Summary in plain language.

When coverage is low, the aggregate findings (which phase has a lot of AI authorship, where work is concentrated overall) are more reliable than the individual findings (this specific contributor has this specific influence). That is because the unanalyzed pairs are more likely to change the picture for less-influential contributors than for the most-influential ones. If you need confidence in a finding about a specific person – for example, before having an HR conversation, or before pricing a deal on the basis of a single key contributor – and coverage is below the high 80s, run an additional analysis cycle before acting.

4.2 Findings that always deserve corroboration

Some findings should never be acted on by themselves, even when coverage is high. They deserve a second source.

Anything that names an individual

When a finding names a specific person, treat the report as the start of an inquiry, not the verdict. The Contributors table tells you what shows up in the documents. It does not tell you whether the person is performing well, whether they are happy in the role, whether the work attributed to them is theirs alone or part of a team effort, or whether they followed the company's policies. Confirm with managers, conversations, and where appropriate, the contributor themselves.

Findings concentrated in a phase with low coverage

If a Key Finding rests heavily on a phase where many of the eligible pairs were not analyzed, ask for an additional run before treating the finding as settled. The Coverage Summary table will tell you the overall coverage; ask morriganAI for a per-phase breakdown if the finding is consequential.

High AI Usage on a single contributor

If the report shows a person with high AI Usage, that pattern can mean several different things. It can mean the person is a thoughtful power user who has learned to work efficiently with AI. It can mean the person is leaning on AI for work they should be doing themselves. It can mean an AI tool is producing work in their name that they have not closely reviewed. Raven cannot tell you which. Confirm with a conversation before drawing a conclusion.

Below-threshold or fragmentation findings

When the report flags a large number of below-threshold vectors and uses words like "fragmented" or "siloes," confirm that the affected documents are actually meant to be connected. Sometimes documents are below threshold because the teams are siloes and should not be – that is a real problem. Sometimes documents are below threshold because they cover unrelated work that has no reason to be connected – that is not a problem at all. The report cannot always tell which is which, but a manager who knows the work usually can.

Quantified risk numbers in the Risk Summary

The Risk Summary section may include figures like "potential revenue impact of \$500,000 per quarter" or "estimated remediation time of six months." These figures are illustrations drawn from public industry examples, not measurements of your specific business. They are useful for putting a finding in context – they are not Raven's estimate of what your particular outcome will be. Use them to communicate severity, not to set budgets.

4.3 Conclusions Raven cannot support

Some questions feel like they could be answered by a Raven Report but cannot. The list below is the most important boundary in this manual.

1. Raven cannot evaluate an employee's performance. It does not measure productivity, quality of judgment, customer impact, or contribution to outcomes that do not show up in documents. A contributor with a small footprint in the report may be doing some of the most important work in the company.
2. Raven cannot prove that a specific person used a specific AI tool. The AI Usage column reflects the character of the documents, not a record of who logged into what.
3. Raven cannot prove a policy violation. Whether high AI Usage in a particular role is a violation depends on the policy, the role, and the work. That judgment belongs to humans.
4. Raven cannot read content the authorized account did not have access to. If a connection is read-only to one team's storage, work in another team's storage is invisible to it.
5. Raven cannot replace a forensic audit, a legal review, an HR investigation, or a financial audit. It can inform any of those, but it does not stand in for them.

6. Raven cannot value a business on its own. It can inform the operational-risk piece of a valuation conversation. The valuation itself depends on financials, market position, and many other factors.
7. Raven cannot tell you what an organization should be. It describes the AI Identity that is, not the one a leader wants. The decision about whether what the report shows is the right operating model for the business is a leadership decision.

A practical rule of thumb

Use the Raven Report to direct your attention. Use other sources – managers, contributors, financial records, customer outcomes – to confirm what your attention finds. The report’s job is to point. The decision is still yours.

4.4 Worked examples of corroboration

The scenarios below show how careful corroboration can change the meaning of a finding. Each one starts with what the report says, then walks through what a thoughtful leader does next.

Scenario A: high AI Usage on a senior contributor

The Contributors table shows a senior leader with high AI Usage, the only senior leader in that band. A first-pass reading might suggest the leader is offloading work to AI inappropriately.

A careful corroboration looks like this. The CHRO and the leader’s manager sit down with the leader and ask, simply, how the leader is working. They learn that the leader has built a small library of AI prompts that turn rough notes into structured first drafts, which she then revises substantially. Her drafts are good, her output is high, and she is candid about the tools she uses. The corroborated meaning of the finding is that the leader is a model power user, not a problem. The action that follows is to make her practices visible to others on her team.

Scenario B: process concentration risk in the Do phase

A Key Finding flags process concentration risk in the Do phase, attributed to one engineering lead whose Influence in that phase is very high. A first-pass reading might suggest the company is dangerously dependent on one engineer.

A careful corroboration looks like this. The COO checks the Coverage Summary and confirms that coverage in the Do phase is high. Then the COO talks with the engineering manager and learns that the engineering lead does, in fact, hold the deepest knowledge of the platform, and that two recent attempts to cross-train colleagues stalled because of priority shifts. The corroborated meaning of the finding is that the risk

is real and the prior mitigation work was incomplete. The action that follows is to fund the cross-training as a protected commitment, not a side project.

Scenario C: a fragmentation finding that turns out to be a non-issue

A Key Finding flags fragmented document management because a large share of the analyzed pairs were below the information-density threshold. A first-pass reading might suggest siloed teams producing disconnected work.

A careful corroboration looks like this. The CFO asks the operations lead to look at the documents in question. It turns out the affected documents are produced by two teams that genuinely do unrelated work – a customer success team and a back-office reconciliation team. There was never a reason for their documents to share information. The corroborated meaning of the finding is that the fragmentation is structurally appropriate, not a problem. The action that follows is to do nothing about this finding, and to look elsewhere in the report for risks worth acting on.

Scenario D: low coverage on a contributor-specific finding

The report names a specific contributor as central to the Reward phase. The Coverage Summary shows pair coverage at 62%, with the Executive Summary noting that incomplete coverage may materially affect attribution accuracy.

A careful corroboration looks like this. The leader resists the temptation to act on the finding immediately. They raise the Run Cost Limit and run an additional cycle. Coverage rises to 91%. In the new report, the contributor is still central, but a second contributor – who barely showed in the first report – is also significant. The corroborated meaning of the finding is that the Reward phase has two key contributors, not one. The action that follows is to plan resilience around both of them, and to set the Run Cost Limit higher going forward so future reports start from a position of confidence.

Scenario E: a leader who looks invisible

The Contributors table shows the new CEO with surprisingly low Aggregate Weighted Influence – far below what the leader’s actual role in the organization would suggest. A first-pass reading might suggest the CEO is not yet shaping the work.

A careful corroboration starts not with the finding but with the connections. The leader checks which storage locations were connected when the report was run. The connected scope included the company’s broadly shared team drives but did not include the CEO’s personal work-account drive, where many of the CEO’s drafts, memos, and frameworks live before they get cleaned up and shared. Because that personal work area was not in scope, the documents the CEO authored never appeared in the corpus, and any concept that originated with the CEO appears to originate with whoever first put it into a shared document.

The corroborated meaning of the finding is that the report is reflecting what was connected, not what is happening in the organization. The action that follows is to add the CEO's work-account personal drive to the connected scope (with the CEO's consent and on a dedicated connector account) and run another cycle. The next report often shows the CEO's influence rising sharply, sometimes accompanied by a meaningful shift in who else looks central.

A connection-scope rule of thumb

Before treating any influence finding as final, confirm what was connected. A contributor cannot be measured for ideas they wrote in a place Raven could not see. This is the most common cause of surprisingly low influence for senior leaders, and it is fixable by adjusting the connector scope.

What these scenarios share

In each case, the report's job was to direct attention to a real pattern. The leader's job was to bring context that the documents alone could not provide. That partnership is what makes a Raven Report useful in practice.

4.5 Connection scope as a source of bias

Coverage is not the only thing that limits a Raven Report. The other limit is what was connected to Raven in the first place. Raven can only see what the connected accounts have access to. If a folder, a drive, or a leader's personal work area is not in scope, every concept that originated there will appear in the report as if it originated wherever it was first written into a connected location. That is a systematic bias, not a random error, and it is worth understanding before reading any influence finding as a verdict on a person.

The bias has a predictable shape. Senior leaders, founders, and the people who work most closely with them – chiefs of staff, executive assistants, senior advisors – are the contributors most likely to be under-represented when only shared drives are connected. Their early thinking often happens in their personal work areas before it migrates into team folders. The result is a report where their footprint looks smaller than their real influence on the organization.

There are two ways to address this. One is to read every influence finding with awareness of what was connected and to corroborate accordingly, as Scenario E above shows. The other is to expand the connected scope. The mechanics of how connections work, and the customer's control over them, are explained in Part 5; the safety and control implications are explained in Part 6.

When a Raven Report is going to be used to support a high-stakes decision – a deal, a transformation, a senior succession plan – it is worth taking a few minutes before the analysis to confirm that the connected scope reflects how work actually moves in the organization. A report run on the wrong scope will be precise about the wrong picture.

Part 5 • How a Raven Report is built

You do not need to understand the technical details of Raven to use it well, but a basic grasp of how the report is built will help you trust what is in front of you. This section explains the process from beginning to end in plain language. The next section explains why each step is safe.

5.1 Where the data comes from, and who decides what Raven can see

Raven looks at documents that already exist in the customer's cloud storage. The way Raven gets access to those documents is the most important thing to understand about how the system handles data, because the design puts control of scope directly in the customer's hands.

The customer does not give Raven a master key to their cloud storage. Instead, the customer follows a three-step setup that uses the customer's own identity and access controls to decide what Raven can read.

1. The customer creates a new user account inside their own organization. This is an ordinary user account in the customer's own Google Workspace, Microsoft 365, Dropbox, or Box environment. We refer to this account throughout the rest of the manual as the connector account. The customer creates and manages it the same way they create and manage any other user account in their organization.
2. The customer grants the connector account access to whichever folders, shared drives, or sites should be considered by Raven. Access is granted using the customer's normal sharing controls – the same way the customer would invite any colleague to a folder. If a folder is not shared with the connector account, Raven cannot see it.
3. The customer connects Raven to that connector account in the Raven Portal using OAuth, the standard secure sign-in handshake. From that point on, Raven reads only what the connector account has been authorized to read.

The result is that the customer holds the dial. The connector account is an account inside the customer's own system, governed by the customer's own administrators, and visible in the customer's own audit logs. Whatever that account can see, Raven can see. Whatever that account cannot see, Raven cannot see.

A powerful fail-safe

If the customer ever wants to limit Raven's access – to a specific folder, a specific drive, or everything at once – they do not have to call morriganAI, change a setting in the Raven Portal, or trust morriganAI to honor a request. They change the connector account's permissions in their own identity system. The change takes effect immediately, in the customer's own infrastructure, under

the customer's own audit trail. This is one of the strongest control mechanisms a customer can have in any third-party data tool.

5.2 Multiple connections and what each one is for

Raven supports any number of connector accounts on a single Raven Portal account. Each connection is independent: the customer can add one, remove one, or change the access of one without affecting the others. Connections are typically used in two complementary ways.

Shared-drive connections

A shared-drive connection is a connector account that has been granted access to broadly shared team folders, department drives, or company-wide spaces. This is usually the first kind of connection a customer creates, and for many organizations it provides most of the picture. It is the simplest setup: one connector account, granted access to the shared spaces leadership wants Raven to consider, connected once.

Individual-account connections

Customers may also choose to connect individual accounts – particularly the work-account personal storage of senior leaders and the people closest to them. This is optional. It is offered as a choice, not a requirement, because some organizations will want to keep the scope narrow and others will want a richer picture of how ideas originate.

The reason this option exists is that many of the most consequential ideas in an organization originate in a leader's personal work area before they appear in a shared document. A CEO drafts a memo in her own work-account drive and shares the cleaned-up version a week later. A chief of staff prepares the board pre-read in his personal folder and only the final version goes to the board portal. An executive assistant maintains a leader's working notes in a place no team member sees. When those personal work areas are not connected, the documents that come out of them appear in the report to originate with whoever first put the idea into a shared location, even if that was several steps downstream of the actual author.

When a customer wants to add an individual-account connection, the same three-step setup applies: the customer (with the leader's consent and cooperation) creates or selects a connector account, grants that account read access to the leader's relevant work-account drive, and connects it in the Portal. The same fail-safe applies: the leader's administrator can adjust or revoke that access at any moment from inside the customer's own system.

A clarification on the word "personal"

When this manual refers to a leader's personal storage, it means the leader's own drive inside their work account – for example, their My Drive in Google Workspace or their personal OneDrive in Microsoft 365. It does not mean the leader's home email or any cloud storage they use outside of the company. Raven only ever connects to storage that is part of the customer's own organization.

5.3 The customer's responsibility for connection scope

The flip side of putting control of scope in the customer's hands is that the customer is responsible for choosing the right scope. The quality of a Raven Report is conditional on the appropriateness of what is connected to it. Connect everything that matters and the report reflects how the organization actually operates. Leave out a meaningful chunk of the work, and the report reflects only what was connected.

For most organizations, deciding on connection scope is a brief conversation among a few leaders. The questions worth asking before the first run are:

- Which shared drives, sites, or folders contain the work this report is supposed to describe?
- Are there leaders whose personal work-account drives contain meaningful early-stage thinking that we want represented in the picture?
- Are there folders we deliberately want to leave out – for example, legal hold material, HR records, or work that is genuinely unrelated to the operating picture we are trying to see?

These choices can be revisited at any time. A customer can start with a narrow scope and expand later. A customer can add an individual leader's connection ahead of a specific event (a board meeting, a transformation review, a diligence cycle) and remove it afterward. A customer can change scope to support a one-time question and roll it back the next day. Because the controls live in the customer's own identity system, those changes do not require any back-and-forth with morriganAI.

5.4 What gets read and what does not

Once Raven has been connected, it reads the text of supported file types: Word documents, PDFs, presentations, spreadsheets, plain text, CSV, and Markdown. If the customer turns on the optional media setting, Raven can also transcribe video and describe images. If a file type is not supported, Raven simply skips it – failed or unsupported files do not stop the analysis.

Raven does not read your email. It does not watch keystrokes. It does not record screens. It does not monitor what people are doing in real time. It looks at the artifacts your organization has already produced and that already live in cloud storage you control, accessed through the connector account you created.

5.5 Pairs and Vectors

Once Raven has the text of a document, it pairs that document with older documents in the same storage that fall within a chosen time window, typically 90 days. For each pair, Raven uses an AI model to compare the two documents and answer a small set of structured questions:

- How much of the older document’s specific facts and figures show up in the newer one? (Content Retention)
- How much of the older document’s underlying ideas and logic show up in the newer one? (Concept Retention)
- How much of the newer document reads as AI-authored?
- For each of the six business phases, what is the probability that this pair belongs to that phase?

The answers are encoded into a fixed-format 135-character string called a Raven Vector. A Vector contains no human readable or independently decipherable content. It is a compact, structured record of a single comparison between two documents.

5.6 From Vectors to findings

Once enough Vectors exist for an organization, Raven aggregates them into the metrics you see in the report. This step happens in morriganAI’s own systems and does not involve any further AI analysis. It walks the chains of ancestry between documents – what we call the Genetic Trail – to compound how much of one person’s work has carried forward through several generations of edits. It calculates the Aggregate Weighted Influence for each contributor, the AI Usage for each contributor, and the phase-by-phase metrics that fill the Phase Analysis section.

5.7 How the prose sections of the report get written

The Key Findings, the Risk Summary, and the Executive Summary are the three sections of the report that read like prose. They are produced by an AI model – the same model that does the document comparisons – but the model is given the structured metrics, not the raw documents. It is asked to produce findings that are actionable, to compare the current findings to prior reports for the same organization where possible, and to put the findings in the context of public industry examples and best practices.

The Executive Summary is generated last, from the text of the Key Findings and the Risk Summary, and is required to include the coverage sentence and a callout when coverage is materially low.

5.8 Coverage and Run Cost

Each pair comparison costs a small amount because it uses an AI service. Most accounts set a Run Cost Limit so a single run cannot exceed a known dollar amount. If the limit is reached before all eligible pairs are analyzed, the report tells you so plainly in the Coverage Summary table. Raising the limit, scheduling more frequent runs, or running an additional cycle on demand are all ways to fill in the gap. Coverage

tends to climb over time on a recurring schedule because each run only needs to analyze the pairs that are new since the last run.

Part 6 · Why it is safe to use Raven

Leaders are right to be careful about any tool that touches their organization's documents. This section explains, in plain language, why morriganAI built Raven the way it did and why each design choice protects the organization, its employees, and its data. Read this section in full if you are evaluating Raven for the first time, or if your security or HR team has questions about how the system handles information.

6.1 The starting principle: privacy by design

morriganAI's core commitment is that the company holds no personally identifiable information and no commercially sensitive identifiers. Every part of the Raven design follows from that commitment. Raven was built to look at artifacts – the documents your organization already produces – and to draw inferences from them. It was deliberately not built to track people or watch workflows.

This is a meaningful contrast with traditional employee-monitoring or comprehensive-data-collection approaches, which collect more data than they need and create large attack surfaces in the process. Raven is the opposite. It collects only what it needs to draw the inferences leaders care about, and it stores that information in a form that cannot be used to identify anyone on its own.

6.2 Connection scope: control lives in your identity system

The strongest safety feature in the Raven design is one that does not appear on any feature list, because it is structural rather than technical: control of what Raven can read lives inside the customer's own identity system. Part 5 explains the mechanics of how this works. This section explains why it matters as a safety feature.

In a traditional vendor relationship, the customer hands the vendor a credential, the vendor uses that credential to read whatever it has been pointed at, and the customer trusts the vendor to honor whatever access agreement is in place. The customer's control over what the vendor sees is mediated by the vendor's software and the vendor's good behavior.

Raven is built differently. The connector account is not a credential issued by morriganAI; it is an ordinary user account inside the customer's own organization, created and administered by the customer, governed by the customer's own access controls, and visible in the customer's own audit logs. Raven sees what that account has been authorized to see and nothing else. The customer's control is direct, not mediated.

Three properties follow from this design that are difficult to achieve any other way:

- Immediate, unilateral access change. If the customer ever wants to limit Raven's access – to a specific folder, a specific drive, or everything at once – they change the connector account's permissions in their own system. The change takes effect immediately. They do not need to call morriganAI, change a setting in the Raven Portal, or trust morriganAI to honor a request.
- Auditability inside the customer's own infrastructure. Every access the connector account makes shows up in the customer's normal audit logs alongside every other user account. Security and compliance teams do not need a separate audit trail; the connector account is visible in the same place they audit everything else.
- No master credential to misuse or lose. There is no special key the customer hands to morriganAI that would, if mishandled, give Raven access to more than the customer intends. The connector account has exactly the access the customer's administrator has granted it, no more.

Why this is the right answer for security and HR teams

The hardest question security and HR teams ask about any third-party tool is "what happens if we change our minds?" With Raven, the answer is that you change your mind in your own system, with your own administrator, and the change takes effect immediately. That is a stronger control posture than most vendor relationships can offer.

There is a corresponding responsibility, which is honest to acknowledge: because the customer chooses what is in scope, the customer is also responsible for choosing well. Sections 4.5 and 5.3 cover that responsibility from the report-quality angle. From the safety angle, the implication is simply that the customer cannot delegate scope decisions to morriganAI – they must be made by someone in the customer's organization who understands what should and should not be analyzed.

6.3 The two files that stay on your storage

Two small files are central to the Raven design, and both of them live exclusively on your own cloud storage. They are never copied to morriganAI's systems.

- `log.raven` maps a coded file identifier back to the actual file ID in your storage. It allows the analysis to recognize which document is which without morriganAI's systems ever holding the readable file ID.
- `userid.raven` maps a coded user identifier back to a real person's name and email. It is the only place the link between a coded identifier and a real human exists.

Both files are written into the root of your connected cloud storage, hidden by default where the storage provider supports it. They belong to you. If you disconnect Raven from a storage provider, those files

remain in your storage under your control. If you delete them, the next Raven run will rebuild them from the document metadata in your storage.

Why this matters

Because morriganAI never stores the readable file IDs or user identities, a breach of morriganAI's systems would not, on its own, expose who any contributor is or which file is which. The only place that link exists is in your storage, behind your access controls.

6.4 What is stored in morriganAI's systems and how it is protected

morriganAI's systems hold the coded vectors, the run metadata, and the generated reports. The protections on each are described below.

| What is stored | How it is protected |
|----------------------------------|--|
| Raven Vectors | Each Vector is a 135-character coded string. It contains coded references for the file IDs and user identity but no readable content, no document text, and no names or emails. Stored in a managed database with access controls. |
| Run metadata | The dates and status of each analysis run, plus high-level statistics. No document content. |
| Generated reports (PDF and Word) | Stored in encrypted cloud storage with public access blocked. Reports are delivered to you through short-lived download links that expire within fifteen minutes. |
| Cached extracted text | When Raven extracts text from a document for analysis, it caches that text temporarily so a second run does not need to re-download the file. The cache is stored in encrypted cloud storage with public access blocked. It can be cleared on request. |
| Connection tokens (OAuth) | Encrypted at rest using strong encryption. Never written to logs. Refreshed automatically before expiry. |

When AI is asked to compare two documents, only the text of those two documents is sent. No account identifiers, no user names, no connection tokens, and nothing about your organization's wider context are sent with the request.

6.5 Why retaining Raven Vectors is the safe choice – and why it is the useful one

A reasonable question to ask is, "Why does morriganAI keep the Vectors at all?" There are two answers, and they reinforce each other.

The first answer is the safety answer. A Raven Vector contains no readable text. It contains coded identifiers for the two documents it compares and the user who last modified the newer document. The link between those codes and any actual file or person exists only in your storage, in the two files described in section 6.3. A Vector on its own – disconnected from those mapping files – cannot be used to identify a document or a person. It is, by design, useless to anyone outside your organization.

The second answer is the value answer. The Vectors are what make the trend line possible. Each new Raven Report builds on the Vectors created in earlier runs. Without retention, every report would be a cold start; the trend in AI prevalence over time would not exist; the change in influence as a person grows into a role would not be visible; the comparison of before and after a transformation would not be possible. Retention is what turns a series of snapshots into a trajectory.

The combined argument

Retaining Vectors is safe because the Vectors do not hold readable content. Retaining Vectors is useful because they are what enable the trend line. The same design that makes them safe makes them valuable.

6.6 Avoiding surveillance by design

Raven was deliberately designed to avoid being a surveillance tool. It does not watch employees. It does not measure typing speed, screen time, or whether a person is at their desk. It looks at finished artifacts. Those artifacts are work the organization has already chosen to produce and store.

This design choice is not just ethical. It is also more accurate. Direct monitoring of work activity tells you whether someone is busy. Looking at the artifacts they produce tells you what their work is contributing to the organization. The second question is the one leaders actually need answered, and it is the question Raven was built to answer.

6.7 What happens if you end your subscription or disconnect a provider

If you disconnect a storage provider, Raven stops indexing it immediately. The `log.raven` and `userid.raven` files remain in your storage where they were created; they are yours, and you can keep them or delete them as you wish.

If you end your subscription, you can request that morriganAI delete your account's Vectors and any data regarding your account. You will lose the trend history if you do this; so if you choose to come back later, the next run will start from a cold state. Your organization's own documents are never affected by ending your subscription, because they were never copied – they always lived in your storage.

Quick reference cards

The next three pages summarize the manual for the three audiences most likely to read just one section before a meeting. Each card is a one-page recap.

For the CEO and COO

Read first: the Executive Summary, the coverage sentence, and the Key Findings. Then look at the Phase Analysis for the phase most central to your business model.

Use Raven to:

- Confirm or challenge the operating story you are being told.
- Find the parts of the operation most exposed to a single person or a single tool.
- Decide where to invest in cross-training, governance, or process integration.
- Track whether transformation work is actually changing how the work is done.

Do not use Raven to:

- Evaluate an individual employee's performance.
- Prove that someone violated an AI policy.
- Replace a financial audit, an HR investigation, or a forensic review.

Always check coverage. If the Executive Summary flags low coverage, treat individual-level findings as directional, not precise. Run another cycle before acting on a finding about a specific person.

For the CHRO and HR leaders

Read first: the Contributors table, the Phase Distribution, and the Key Findings. Then look at the Phase Analysis for any phase where the Key Findings flagged a concern.

Use Raven to:

- Identify the contributors whose ideas are quietly central to the work.
- See where AI is being leaned on and start a conversation about whether that is appropriate for the role.
- Plan cross-training and succession to reduce concentration risk.
- Track whether new senior hires are actually shaping the work or still on the edge of it.

Do not use Raven to:

- Score, rank, or discipline employees.
- Determine pay, promotions, or terminations on its own.
- Conclude that a person is or is not using AI compliantly.

Treat any finding that names an individual as the start of a conversation, not a verdict. Confirm with managers and, where appropriate, the contributor themselves.

When communicating to the workforce about Raven, you can say accurately that morriganAI sees only what a specific, named user account in your own organization has been granted access to read, and that the company can change or revoke that access at any time from inside its own identity system.

For the banker and private equity investor

Read first: the Executive Summary, the Coverage Summary table, the Key Findings, and the Risk Summary.

Use a snapshot to:

- Add operational evidence to the diligence file before close.
- Pressure-test management's story about how the team is operating and how durable recent productivity gains are.
- Surface concentration and AI control risks that may need to be priced in or addressed in the first hundred days.
- Inform the operational-risk piece of a valuation conversation.

Use a trend line to:

- Confirm post-close that operating-partner work is producing visible change.
- Show the investment committee or LPs how AI control posture and concentration risk are moving over the hold.
- Compare a target to comparable assets in your portfolio in operationally meaningful terms.

Do not treat the Risk Summary's dollar figures as the company's actual exposure. They are illustrations drawn from public industry examples, not measurements of the specific business. Use them to communicate severity, not to set a number.

Frequently asked questions

Questions leaders, HR teams, and security teams commonly ask about Raven, answered in plain language. If your question is not covered here, ask morriganAI directly – the design of the system is meant to be explainable.

Questions from CEOs and COOs

How long does a Raven Report take to produce?

A first run typically completes in hours rather than days, depending on how much content is in the connected storage and the Run Cost Limit. Subsequent runs on the same storage are faster because Raven only analyzes what is new or changed since the last run. This is much faster than traditional risk assessments, which often take months.

How often should we run a Raven Report?

For most organizations, a quarterly cadence is a sensible starting point. Quarterly is frequent enough to see meaningful change, infrequent enough that each report is still substantive. Organizations going through active transformation, an integration after acquisition, or a period of rapid AI adoption may want to run monthly. A snapshot before and after a transformation is also a high-value pattern.

What does it cost to run?

Each run uses an AI service to compare document pairs, and each comparison costs a small amount. The cost charged for a run is the cost of the AI service used, with a small markup. The Run Cost Limit lets you cap any single run at a known dollar figure. Subsequent runs are usually less expensive than the first because most pairs have already been analyzed.

Questions from CHROs and HR leaders

Will employees know they are being included?

Raven looks at documents that already exist in the organization's cloud storage. It does not install anything on any employee's computer or device. It does not change how employees use their applications. As a matter of good practice, leadership should let the workforce know that the organization is using a tool to study its own AI Identity. Many organizations name Raven explicitly in their internal AI policy.

Could a Raven Report be used to discipline an employee?

It should not be, and the manual is direct about this. A Raven Report describes patterns in the documents an organization has produced. It does not measure performance, judgment, or compliance with policy. Using a Raven Report on its own as the basis for a disciplinary action would be a misuse of the tool. Findings that name individuals are starting points for conversations, not verdicts.

What about contributors who do important work that does not show up in documents?

Raven cannot see what is not written down. A contributor who is mentoring others, handling difficult customer conversations, or doing thoughtful one-on-one work may have a small footprint in the report and an outsized contribution to the organization. The report's silence about a person is not a judgment of that person.

Questions from bankers and private equity investors

Is the report defensible in a diligence file?

Yes. The report follows the same structure every time, the metrics are defined in a glossary, and the methodology is described in this manual and in morriganAI's documentation. A snapshot can be archived and referenced later. A trend line can be retrieved as a series of dated reports.

Can the same Raven Report be shared with management?

Yes, and many investors do exactly that as a way to align the operating partner and the management team around a shared view of operational risk. Raven was designed to be the kind of report a board can sit with, not a black-box risk score.

Does the report tell us if AI usage is creating financial liability?

It does not tell you that directly. It tells you where AI is being heavily used and where governance signals look weak. The financial and legal implications are judgments that your counsel and your finance team need to make. The report supports that judgment; it does not substitute for it.

Questions from security and IT teams

What permissions does Raven need?

Raven does not receive a master credential. Instead, you create a user account inside your own organization (called the connector account in this manual), grant that account read access to whichever folders or drives you want Raven to consider, and connect Raven to that account in the Raven Portal using OAuth. Raven also writes two small mapping files (log.raven and userid.raven) into the root of the connected storage; this requires a small amount of write access, scoped to the same storage the connector account already has access to. The exact OAuth scopes used for each provider are documented in morriganAI's technical documentation. See Part 5 for the full setup model and Part 6.2 for the safety implications of this design.

How do we change what Raven can see, or stop it from seeing something?

You change the connector account's permissions in your own identity system, the same way you would change any other user account's permissions in your organization. The change takes effect immediately. You do not need to call morriganAI, change a setting in the Raven Portal, or wait for anything on

morriganAI's side to happen. If you want to stop Raven from seeing anything at all, you can disable the connector account or revoke its access entirely; Raven will simply have nothing to read on its next attempt.

Can we connect more than one account?

Yes. Raven supports any number of connector accounts on a single Raven Portal account. Each one is independent, so you can add, remove, or change the access of any single connection without affecting the others. A common pattern is one or more shared-drive connections plus optional individual-account connections for senior leaders whose early-stage work happens in their own work-account drives. Whether to add individual-account connections is a customer choice, discussed in Part 5.2.

Where is the data stored?

Raven Vectors, run metadata, generated reports, and cached extracted text are stored in encrypted cloud storage operated by morriganAI. The mapping files (log.raven and userid.raven) are stored exclusively in your own cloud storage and are never copied to morriganAI's systems.

What happens if morriganAI is breached?

The data morriganAI holds – coded vectors, run metadata, generated reports, cached document text – is meaningful only when paired with the mapping files that live in your storage. A breach of morriganAI's systems would not, on its own, expose who your contributors are or which document is which, because the link between the codes and the readable identities exists only in your storage. The cached extracted text is the most sensitive content morriganAI holds; it is encrypted at rest and access is tightly controlled.

Is the data used to train AI models?

morriganAI's contractual agreement with its AI provider is for inference only, not for training. The text of your documents is sent only at the moment a comparison is being computed and only as the minimum content necessary for that comparison.

Can we get a copy of all the data on us?

Yes. The Portal lets the account owner download every report. morriganAI can also export the run metadata and the vectors associated with the account on request. The mapping files are already in your storage, where you control them.

Glossary

Plain-language definitions of the terms used in this manual and in a Raven Report.

| Term | Plain-language meaning |
|------------------------------|---|
| AI Identity | How an organization's identity – its people, its processes, and the trust stakeholders place in it – is shaped by the AI tools it uses. |
| AI Usage | A probabilistic estimate of how AI-authored a contributor's recent work appears to be, weighted by how informationally rich each document is. |
| Aggregate Weighted Influence | A measure of how much of a contributor's work shows up in the documents that came after, traced through chains of edits. |
| Below-threshold vector | A document comparison where the information overlap between the two documents was small enough that the comparison is excluded from summary findings. |
| Business Phases | The six phases every organization moves through: Gather, Tell, Agree, Do, Reward, Govern. |
| Concept Retention | How much of an older document's underlying ideas, frameworks, and logic carry forward into a newer document. |
| Conceptual Overlap (PC) | Concept Retention reported at the contributor and phase level after compounding through chains of edits. |
| Connection scope | The set of folders, drives, and sites that the connector accounts have been granted access to. Defines what Raven can see. Controlled entirely by the customer in the customer's own identity system. |
| Connector account | A user account created by the customer inside their own organization (Google Workspace, Microsoft 365, Dropbox, or Box) and granted access to whichever folders or drives Raven should consider. The customer connects Raven to this account using OAuth. Whatever this account can see, Raven can see; whatever it cannot see, Raven cannot see. |
| Content Retention | How much of an older document's specific facts, figures, prices, and dates carry forward into a newer document. |
| Contributors | The people whose work shows up in the analyzed documents, listed in the Contributors table of the report. |

| Term | Plain-language meaning |
|-------------------------------|--|
| Coverage | The percentage of eligible document pairs that were actually analyzed in this run, including any analyzed in earlier runs. |
| Descendent | In a Raven comparison, the newer of the two documents being compared. |
| Factual Coverage (PF) | Content Retention reported at the contributor and phase level after compounding through chains of edits. |
| Genetic Trail | The chain of older documents that influenced a newer one, walked back link by link. |
| Influence Distribution | A phase-level metric of whether new work in that phase is shaped by many older documents or just a few. |
| Individual-account connection | An optional connector account that has been granted access to the work-account personal storage of an individual leader (such as a CEO, COO, chief of staff, or executive assistant). Used to capture early-stage ideas that originate in a leader's own work area before being shared more broadly. |
| Information Density | The average of Content Retention and Concept Retention for a single comparison. Comparisons below 10% are flagged as below-threshold. |
| Last Step Similarity | A phase-level metric of how closely a typical new document resembles the most recent older document in the same phase. |
| Linked Identity | When the same person appears as different users across two storage providers, an account owner can link those users into a single canonical identity in the Portal. |
| log.raven | A small file kept in your cloud storage that maps coded file identifiers back to actual file IDs. Never copied to morriganAI's systems. |
| Origin | In a Raven comparison, the older of the two documents being compared. |
| Origin Depth | A phase-level metric of how far back in time the chain of influence reaches for documents in that phase. |
| Pair | A specific older document and a specific newer document being compared. |
| Pair Coverage | Same as Coverage. The percentage of eligible pairs that have been analyzed. |

| Term | Plain-language meaning |
|----------------------------|--|
| Phase | One of the six business phases. Every analyzed document is given a probability for each phase, and the highest is the document's dominant phase. |
| Process Concentration Risk | The risk that too much of an organization's work depends on a single person. |
| AI Control Risk | The risk that AI is producing work that is not being adequately reviewed or governed. |
| Raven Report | The structured PDF (and Word) document delivered at the end of an analysis run. |
| Raven Vector | A 135-character coded string that records a single document comparison. Contains no readable content. |
| Run Cost Limit | An optional dollar limit on a single analysis run, which keeps spending predictable. |
| Shared-drive connection | A connector account that has been granted access to broadly shared team folders, department drives, or company-wide spaces. Usually the first kind of connection a customer creates. |
| Snapshot | A single Raven Report taken at one moment in time. |
| Time Window | The period – usually 90 days – within which a newer document is paired with older documents for comparison. |
| Trend line | A series of Raven Reports run on a recurring schedule, used to see how the AI Identity of the organization is changing. |
| Unidentified User | A contributor whose coded identifier could not be matched to a name in your storage at report time. Shown as "Unidentified User 1," "Unidentified User 2," and so on. |
| userid.raven | A small file kept in your cloud storage that maps coded user identifiers back to real names and emails. The only place that link exists. Never copied to morriganAI's systems. |

A closing note

A Raven Report is a tool, not a verdict. Its job is to give the leaders of an organization an honest, defensible picture of how the work actually moves and how AI is shaping it. The decisions that follow are still leadership decisions, and the people the report describes are still people whose work and judgment cannot be reduced to any number on any page.

Used well, Raven gives leaders something they have not had before: a way to see the AI Identity of their organization clearly enough to act on. Used carefully – with attention to coverage, with corroboration where individuals are named, and with respect for the conclusions the report cannot support – it can become a steady part of how a company governs itself in the AI era.

morriganAI welcomes questions about anything in this manual. The way Raven was designed and the way Raven Reports are written are both meant to be inspectable. If something is unclear, ask.