# SAUCE: Enterprise AI Guide

> Auto-synced from `docs/proposals/SAUCE-ENTERPRISE-AI-GUIDE.md` at 2026-04-13 19:44 UTC. Source of truth is the repo. Re-sync: `node scripts/notion-playbook-sync.cjs --playbook enterprise-ai-guide`

## The Cost Nobody Budgets For

> **CANONICAL.** Status: v2 | Date: 2026-04-13

>

> **Pricing Note:** This guide uses the canonical pricing framework: Quick Win $1,500–$5,000 (one-time); Visibility $149/mo; Structure $497/mo per desk; Managed $1,249/mo. Tiers are capability ceilings — add-ons (enrichment, blueprints, advanced agents) bolt on when needed. Cross-reference: PRODUCT-SERVICE-CATALOG-v2.md, SAUCE-ADDON-STRATEGY.md, SAUCE-ADDON-PRICING-TABLE.md, scripts/configs/client-*-tier.json.

**A practitioner's guide for profitable businesses that don't have time for an ecosystem.**

---

Your best person quits tomorrow. How much walks out the door with them?

Not their salary. Not even their replacement cost. What walks out is the 42% of institutional knowledge that exists nowhere but in their head — not in a document, not in a system, not in anyone else's memory. That number comes from a study of over 1,000 U.S. workers. It is not a metaphor.

The routing decisions. The client relationships. The "ask Sarah, she knows." Gone.

The cost is real and researched. SHRM estimates replacing a key employee costs 50 to 200% of their annual salary — and most of that cost is not recruiting. It is the 5 to 12 months of lost productivity, stalled projects, and re-learning that follows. A single departure of a $65,000 operations manager costs $32,000 to $130,000 in disruption. For a 15-person service business with two departures a year, that is $55,000 to $220,000 — and it never appears on a P&L.

But the departures are just the visible damage. The invisible damage is happening every week your team is intact.

McKinsey found that knowledge workers spend 20% of their workweek — one full day — searching for and gathering internal information. Not creating value. Not serving clients. Searching for things that were already known, already decided, already documented

somewhere no one can find. In a five-person team, that is 25 hours a week of re-derivation. Every week. Forever.

And when someone finally finds the answer, there is a 58% chance they are re-solving a problem that was already solved. That is not a productivity problem. That is a systems problem — and it has a fix.

This guide is about that fix. Not about AI. Not about digital transformation. About the invisible, compounding cost of running a business where institutional memory lives in people's heads instead of systems — and how to reverse it.

Everything else in this guide — the agents, the dashboards, the automation — is downstream of that one fix.

> You built a business that works. We make it work without you.

---

**Who this is for:** Owner-operators of profitable service businesses. Five or more years in. Revenue between $500K and $5M. Using multiple tools that don't talk to each other. Spending 25+ hours a week on work that should be invisible.

**Who this is not for:** People evaluating AI vendors. People who want a technology comparison chart. If you want to understand large language models, read a technical guide. If you want to understand what changes on Monday morning, read this one.

**How to read this guide:**

- **5 minutes:** Read this opening and **How to Start**. That is the decision.
- **20 minutes:** Add the Assess, Build, and Numbers sections. That is the business case.
- **Full reference:** Everything including appendices. That is for after the first call.

---

## Table of Contents

- **Appendix C: Knowledge Capture Pipeline — Technical Reference**
- **Appendix D: Blueprints, Security & Add-Ons**

---

# The Four Layers

## The Framework

Knowledge in a business moves through four layers. Every business has all four. In most businesses, the first two are broken and the last two do not exist.

Here is what that looks like. Your project manager makes a routing decision in a Slack thread on Tuesday. By Thursday, three people have asked about it. Nobody can find the thread. The decision gets re-made — differently. That is knowledge trapped in the first layer, with no path to the other three.

| Layer | Name | What It Is | What Usually Happens |
|---|---|---|---|
| 1 | **Communication** (Sources) | Where knowledge originates — Slack, email, meetings, calls, notes | It stays trapped in the thread. Nobody extracts it. It is gone in 90 days. |
| 2 | **Documentation** (Hubs) | Where knowledge gets structured — canonical stores, process docs, decision logs | Scattered across 5 tools. Outdated. Nobody trusts it. |
| 3 | **Visibility** (Views) | Where humans see knowledge — dashboards, briefs, reports, review queues | Requires hours of manual assembly. Stale by the time it arrives. |
| 4 | **Execution** (Actions) | Where knowledge drives automated work — agents acting on curated context | Does not exist. Humans re-derive context every time they act. |

The fix is building all four layers so they flow into each other: knowledge is captured from where people already communicate, structured into canonical hubs, surfaced in views that arrive without asking, and used by automated systems that execute on curated institutional memory.

This is what we call the **knowledge system** — the pipeline that makes every other system in this guide work. The technical architecture is in **Appendix C**. The principle is simple: capture what your team already knows, structure it so anyone can find it, show it where decisions happen, and let systems act on it.

The three levels we offer — Visibility, Structure, Managed — map directly to how deep into these four layers your business goes.

## The Exit Value Problem Nobody Talks About

Knowledge loss does not just cost you time. It costs you wealth.

Business brokers apply a formal "key-person discount" to businesses where operations depend on specific individuals. That discount ranges from 5 to 25%, reaching 40% or more in severe cases. For sole proprietorships where the owner *is* the business, the discount can approach 100% — meaning the business has no transferable intangible value beyond personal goodwill.

The data across 80,000 businesses assessed by the Value Builder System shows the gap clearly: the average business receives acquisition offers at 3.5x pre-tax profit. Businesses that score high on systematization receive offers at 6.0x to 7.1x — roughly double. For the $500K to $5M segment, reducing owner dependency and systematizing operations increases your sale multiple by 0.5x to 1.5x on the same earnings base, potentially increasing business value by 25 to 60%.

And 80% of businesses listed for sale never close. Owner dependency is consistently cited as one of the top three reasons deals die.

Every system in this guide — every desk, every agent, every dashboard — reduces that dependency. The knowledge system is how the business runs whether you are there or not. That is what buyers pay for.

## Three Prerequisites

Before any of this matters, three things have to be true:

**1. Your data has to be somewhere findable.** A CRM. An accounting system. A spreadsheet. Anything. If your customer history lives in email threads and memory, the first step is not automation. It is getting the data out of people's heads and into a system. Any system.

**2. You have to have at least one workflow that repeats more than once a week.** Automation only works on things that happen again. If every week is different in every way, you need clarity on what your business actually does. But if there are follow-ups that should happen, reports that should get built, data that should move from one place to another — that is where this starts.

**3. Someone has to own the system after it is built.** Not manage it daily. Not babysit it. Own the decision of whether it is working. Someone who reads the Monday brief and says "this is useful" or "this is wrong." The owner is usually you.

If those three things are true, the rest of this guide applies.

## What "AI in Your Business" Actually Means

When you hear "AI," you probably picture a chatbot or a content generator. That is one small corner of what actually runs.

The Four Layers describe *what* needs to happen — capture, structure, surface, act. But *how* those layers run is a stack of five automation types working together. You are already running

some version of this stack whether you know it or not. The question is whether it was designed intentionally or improvised.

Here is a concrete example. The Monday morning brief that arrives in your inbox? A scheduled job fires it off. A program pulls data from your CRM, your accounting tool, and your project tracker. Another step enriches a new lead and queues a follow-up. An AI role synthesizes all of it into a two-page summary. And a planning layer decided what format the brief should take and what signals to highlight. Five moving parts, one output.

Those five parts have names: **Triggers** fire on a schedule or event. **Scripts** execute repeatable jobs. **Workflows** sequence scripts across systems. **Personas** advise within a single task. **Agents** plan and coordinate. Every automated business operation is a combination of those five types. The full taxonomy is in **Appendix A**.

---

*Now you know the framework: four layers, five automation types, three prerequisites. The next step is figuring out where your business actually stands.*

## Where Are You Broken? — The Assess

Step 1 is not "develop your AI strategy." Step 1 is: what is costing you the most time right now? Every other question is downstream of that one.

### The Assess Conversation

We do not start with technology. We start with how the business actually runs today. The conversation is short and the questions are simple:

**Where does institutional knowledge live?** If Sarah quits, what leaves with her? If you are unreachable for two weeks, what breaks? The answer is usually "everything that requires a routing decision." That knowledge is an asset if it is in a system. It is a liability if it is in someone's head. Research shows that 42% of institutional knowledge is unique to the individual employee — undocumented and unshared. This is the question that reveals the real cost.

**Where is time being lost?** Not "where could AI help" — where does the owner spend hours on work that produces nothing? Data entry. Manual follow-ups. Reconciling numbers across tools. Building reports. Answering the same internal questions every week. The typical owner spends 68% of their time working "in" the business — day-to-day operations and firefighting — versus only 32% working "on" the business. That is not a technology problem. That is a design problem.

**Where is data stuck?** You log into five things to answer one question. The data is there. The visibility is not. Your CRM knows one thing. Your accounting tool knows another. Your project tracker has a third piece. Nobody has all three in one place. The average small business uses over 40 SaaS applications and wastes 20 to 30% of that spend on unused or redundant tools.

**Where does the same work happen twice?** Someone enters data in the CRM. Someone else enters the same data in the invoicing tool. A third person copies it into a report. Three people, one fact, three systems. Research shows 58% of knowledge workers spend less than a quarter of their time on genuinely new work — the rest is rework, re-solving problems that were already solved.

These four questions — knowledge, time, data, duplication — map directly to the Four Layers. Where knowledge lives tells us whether Layer 1 (Communication) is being captured. Where time is lost tells us whether Layer 3 (Visibility) exists. Where data is stuck tells us whether Layer 2 (Documentation) is structured. Where work repeats tells us whether Layer 4 (Execution) is automated.

## Mapping Your Business to the Six Systems

The diagnostic maps your answers to one of six operational systems. You do not choose all six — you choose the one that hurts the most.

| System | Diagnostic Question | Common Finding |
|---|---|---|
| **Executive** | Do you own your morning? Or does your inbox? | Owner reads 50 emails before doing strategic work. |
| **Revenue** | Does every lead get followed up? Who checks? | Deals stall in pipeline with no owner. Follow-ups depend on memory. |
| **Client Intelligence** | Where does customer history live? | In someone's memory and four-year-old email threads. |
| **Financial** | What is your cash position right now — without looking? | Owner guesses. Real number is 10 days old. |
| **Operations** | If you were gone for 2 weeks, what would break first? | Everything that requires a routing decision. |
| **Marketing** | Does content go out on schedule without you approving each piece? | It does not. |

Most businesses need one or two systems to start. You do not need all six on day one.

## What the Assess Reveals

The four questions do more than identify pain. They map exactly to the four layers — and they show you where your knowledge pipeline is broken.

| What the Assess Finds | Which Layer Is Broken | What the Knowledge System Fixes |
|---|---|---|
| Knowledge trapped in people's heads, Slack threads, email | **Layer 1 — Communication** is not being captured | Watches your communication channels and automatically captures decisions, research, and insights as they happen — before they disappear |

| What the Assess Finds | Which Layer Is Broken | What the Knowledge System Fixes |
|---|---|---|
| Data scattered across five tools, nothing structured | **Layer 2 — Documentation** has no canonical store | Classifies every capture and stores it in one structured knowledge base — one source of truth that every other system reads from |
| No dashboards, no briefs, owner chases data manually | **Layer 3 — Visibility** does not exist | Syncs approved knowledge to dashboards, briefs, and reports that arrive without anyone asking |
| Same work done twice, follow-ups depend on memory | **Layer 4 — Execution** is manual | Feeds curated context to agents and automations that act on what the business already knows — not on what someone remembered |

The result of the Assess is not a 40-page roadmap. It is a one-page map: here are your six systems, here is which layer is broken in each one, and here is what we fix first.

### The Data Architecture Question

Every business that runs multiple tools has a topology problem: which system is true when they disagree?

Your CRM says a deal is worth $50K. Your accounting tool says $45K. Your sales rep says $60K in Slack. Three numbers, one deal. Which one wins?

If you do not answer that question before you automate, the automation inherits the confusion. The fix is not better AI. The fix is deciding which source of truth wins.

The principle is simple: **one source of truth per data type.** Operations data has one home. Knowledge has another. Documentation has a third. View layers — dashboards, reports, briefs — receive data but never author it.

The critical rule: **nothing writes back upstream.** Data flows in one direction: from where it originates to where it is stored to where it is displayed. If your dashboard and your CRM disagree, the CRM wins. Always. The technical implementation is in **Appendix C**.

---

*The Assess told you what is broken and which layer to fix first. Now you fix the one that hurts the most — and find out what it costs.*

## What We Fix and What It Costs — The Build

### The Quick Win

Every engagement starts here. The fix for the one thing eating the most time. Fixed scope. 2 to 4 weeks. Running when we are done. Not a pilot program. Not a proof of concept.

**$1,500 to $5,000.** Payment: half upfront, half when it is live. 50/50. No games.

There is a reason we do not run pilots. Research across hundreds of enterprise AI deployments found that 95% of AI pilot programs fail to deliver measurable impact. Gartner predicted that at

least 30% of generative AI projects would be abandoned after proof of concept by the end of 2025. The pilot model is where AI goes to die.

The businesses that get ROI are the ones that skip the pilot, commit to a single production workflow, and have it running in 30 to 90 days. MIT research found that mid-market firms convert pilot to production in approximately 90 days — and that purchasing AI tools succeeds at double the rate of building internally. That is our model.

The Quick Win does three things:

1. **Reduces risk.** $1,500–$5K at 50/50 terms, not a $50K commitment to a vendor you have never worked with.
2. **Builds trust.** You see something working before you expand. No slide deck. No roadmap. A running system.
3. **Reveals the real problem.** The discovery that produced the Quick Win scope almost always surfaces the adjacent issues that become the next engagement.

Four patterns cover most Quick Wins:

| What You Describe | What We Build | Typical Investment |
|---|---|---|
| "Data is scattered across five tools" | Connected systems, one view of your customer | $3–5K |
| "We are basically guessing on decisions" | A dashboard or report that answers your weekly questions | $1,500–3K |
| "If Sarah quits, we are in trouble" | The process in a system anyone can run | $3–5K |
| "We used to do it by hand, never built the system" | The manual work, handled automatically | $2–4K |

Every Quick Win addresses one or two of the Four Layers. And here is the important part: **the Quick Win builds the first piece. Visibility keeps it running.**

The dashboard we build during the Quick Win becomes your ongoing dashboard. The automation we set up becomes the first workflow in your managed system. When you move to Visibility ($149/mo), your Quick Win setup cost applies toward the Visibility setup fee. You are not paying twice.

## From Quick Win to Ongoing System

The Quick Win proved it works. Now the question is: do you want it to keep running — and do you want to add to it?

That is where the three levels come in. Each level builds on the one before it, and each one maps to how deep into the Four Layers your business goes.

| | Visibility | Structure | Managed |
|---|---|---|---|
| **Price** | $149/mo | $497/mo per desk | $1,249/mo |
| **Setup** | $1,500 (credited from QW) | $3,000/desk | $5,000 |
| **Desks active** | Executive only (read-only) | 1–2 (you pick, add more as needed) | All 6 included |
| **Agents** | Dashboard + brief | Core agents per desk, advanced as add-ons | All 54 |
| **Users** | 1 | Up to 5 | Unlimited |
| **What it doesn't include** | No automations running on your behalf, no custom integrations | No client-branded deliverables, no dedicated point of contact, no quarterly reviews | Everything is included |
| **Best for** | See your business clearly | Run the desks that hurt most | Step away entirely |

*At 3 desks ($1,491/mo), Managed ($1,249/mo) becomes cheaper and includes all desks + strategic oversight — the natural upgrade path.*

*Add-ons (enrichment, blueprints, advanced agents, support) bolt on at any level — see [Appendix D](#appendix-d-blueprints-security--add-ons).*

## Level 1: See Your Business — Visibility ($149/mo)

The Quick Win built the first piece. Visibility keeps it alive and adds intelligence.

**What shows up:** Your dashboard stays connected. Your Monday briefing arrives without you asking for it. The knowledge system starts running in the background — capturing decisions and knowledge from your communication channels automatically. You see your business clearly for the first time, every week, without chasing anyone for data.

**What changes in the first month:**

- **Week 1:** Kickoff call. We learn your business. We get access to your tools.
- **Week 2:** Your dashboard is live. All your data in one place.
- **Week 3:** Your first Monday briefing arrives. You see what happened last week without asking anyone.
- **Week 4:** Your first workflow is running. Something that used to take hours now just happens.

**When to choose this:** You need visibility before you can act. You are making decisions on stale data. You need to see the business clearly before you can fix it.

## Level 2: Run Your Operations — Structure ($497/mo per desk)

You have visibility. Now you want one domain to run on its own.

**What shows up:** Think of each desk as a position in your business that runs without someone sitting at it. Pick your desk — Revenue, Marketing, Financial, Client Intelligence, or Operations. The Revenue Desk follows up on leads. The Financial Desk watches your cash. The Operations Desk executes recurring work.

Core agents run per desk, 24/7. Advanced agents are available as add-ons ($75–250/mo each) when you need deeper capability. The knowledge system captures knowledge specific to that domain — pipeline decisions in Revenue, cash-flow patterns in Financial, process knowledge in Operations. Every desk runs on the same knowledge pipeline, applied to its own domain.

**What changes in the first month:**

- **Week 1–2:** Kickoff. We learn the domain you picked and connect your tools.
- **Week 3:** Your first desk is live. Three automations are running without your team touching them.
- **Week 4:** You see the first full week of results. The work that used to take 15+ hours is happening on its own.

**When to choose this:** You have visibility (or do not need it) and you know the one domain where manual work is eating you alive. Start with one desk. Add another when the first one reveals the next pain point. Most clients run 1–2 desks at this level — at 3 desks, the Managed level becomes both cheaper and more capable.

## Level 3: Step Away — Managed ($1,249/mo)

**What shows up:**

- All 6 systems active. All 54 agents running.
- Dedicated contact — a person who knows your business, not a ticket queue.
- Custom builds when you need something that does not exist yet.
- Full monitoring. Weekly strategy calls. Quarterly reviews.

The knowledge system runs across every domain. Knowledge compounds — every captured decision, every classified insight feeds back into every desk. The institutional memory of the business improves every agent output, every week.

**What changes in the first month:**

- **Week 1–2:** Full onboarding. We connect every system and map every workflow.
- **Week 3:** All six desks are live. You see dashboards across the entire business.
- **Week 4:** Weekly strategy call starts. You are reviewing results, not managing operations. Problems are getting caught before they reach you.

**When to choose this:** You want to step away from operations entirely. The business should run whether you are there or not.

Annual commitment available at every level ($119/mo Visibility, $397/mo per desk Structure, $999/mo Managed). Month-to-month is always available. You are never locked in.

Not sure which level fits? That is what the first call answers. We ask three questions, look at your current stack, and tell you which level makes sense — or whether to start with just a Quick Win. **malcolm@saucetech.io** or (847) 558-0405.

## Desk Discovery

Once the Quick Win is running and you have chosen a level, the next question is: which desk do you staff first?

Six pain signals map to six desks:

| Pain Signal | Desk | What Gets Built |
|---|---|---|
| "I spend my morning on email before I can think" | Executive | Morning Brief, Decision Logger, Calendar Intelligence |
| "Leads fall through the cracks" | Revenue | Pipeline Tracker, Follow-Up Sequences, Proposal Generator |
| "I don't know what my customers are experiencing" | Client Intelligence | Relationship Context, Churn Signals, Customer Timeline |
| "I don't know my real cash position without pulling a spreadsheet" | Financial | Cash Clarity, Invoice Flow, Variance Alerts |
| "If I'm out for a week, things break" | Operations | Workflow Documentation, Recurring Task Execution, Escalation Routing |
| "Content goes out when someone remembers to post it" | Marketing | Content Calendar, Publishing Workflows, Performance Reporting |

You pick the one that hurts the most. That desk comes online first. The rest follow as the impact becomes clear.

Every desk runs on the same knowledge pipeline — the same capture-classify-store-surface-act flow — applied to its own domain. That is why adding a second desk is faster than the first — the infrastructure is already running.

## Defining Success

ROI is how we justify cost, not how we sell. The prospect should feel the return before we ever calculate it.

We track three types of metrics in every engagement:

**Time-Recovered Metrics (most visible):** Owner hours saved per week (target: 5–20 depending on level). Staff hours saved per week (target: 4–8 per admin per system). Time-to-answer on key business questions (before: hours chasing data; after: seconds via dashboard).

**Operational Health Metrics:** Lead follow-up rate (target: 100% — every lead gets a touchpoint within 24 hours). Cash reporting lag (target: real-time versus weekly or monthly). Process documentation coverage (target: every repeatable process in a runnable system).

**Business Resilience Metrics:** Can the business run for 2 weeks if the owner is unreachable? Are weekly decisions being made with current data or lagged data? And the one nobody talks about until they want to sell: **systems-driven businesses receive acquisition offers at roughly double the multiple of owner-dependent businesses.** Across 80,000 businesses assessed, the gap between average and systematized is 3.5x versus 7.1x pre-tax profit. That gap is the ROI case that matters most to the owner who plans to sell in 5 to 10 years.

| Level | Year 1 Investment | Hours/Week Recovered | Annual Value | Payoff Month |
|---|---|---|---|---|
| Quick Win | $1,500–5K (one-time) | 4–10 | $15–40K | Immediate |
| Visibility | ~$3,300 | 5+ | $15–30K | Month 2 |
| Structure (1 system) | ~$9,000 | 15–20 | $30–50K | Month 3 |
| Managed | ~$20,000 | 20+ | $50–100K+ | Month 4 |

These are conservative estimates based on real engagements. Time savings only — no revenue uplift included.

---

*That covers what each level delivers and what it costs. The next section explains how the system stays accurate and honest under the hood. If you do not care how it works, skip to [What It Looks Like Running](#what-it-looks-like-running--the-manage).*

## Why You Can Trust It

Most guides on AI adoption assume you have an AI review board, a dedicated team, and 13 months. You have none of those things. That is not a limitation — it is why you need a different approach.

Here is the reality: 80% of AI projects fail, according to the RAND Corporation — double the rate of non-AI IT projects. Gartner's most recent survey found that only 28% of AI use cases fully meet ROI expectations, with the majority failing because organizations "expected too much, too fast." BCG surveyed 1,000 CxOs across 59 countries and found that 74% of companies struggle to achieve and scale value from AI.

The businesses that succeed are not the ones with the most sophisticated tools. They are the ones that solved the data problem first, kept a human in the loop where it mattered, and deployed into production immediately instead of running pilots that never graduate.

This section covers how we build systems that stay accurate, stay auditable, and do not make decisions you never approved.

### The Knowledge Capture Pipeline

The knowledge system turns your team's scattered conversations, decisions, and notes into a single organized knowledge base that everyone — and every automated system — draws from. It is how institutional memory goes from trapped in someone's head to structured in a system that anyone can query, any agent can act on, and any dashboard can surface.

## How It Works

A message appears in Slack — someone on your team writes "We decided to go with Vendor B because their SLA includes weekend support."

Here is what happens next:

1. **Capture.** The system reads the message from your communication channel. It wraps it in a standard format so every source — Slack, email, webhook — looks the same to the rest of the pipeline.
2. **Classify.** An AI classifier reads the message and makes two decisions: what kind of knowledge is this (a decision, a piece of research, an insight, a task, or a reference), and how confident am I? It assigns a score between 0 and 1.
3. **Route.** The confidence score determines what happens next:

- **0.85 or above:** The classifier is certain. The capture is filed automatically in your knowledge base. No one needs to review it.
- **0.70 to 0.84:** The classifier is uncertain. The capture goes to a review queue. You get a notification. You have 72 hours to approve or reject it.
- **Below 0.70:** It was not institutional knowledge. Dropped silently. Logged, but not preserved.

1. **Store.** Approved captures go into your central knowledge store — the single source of truth that every agent, every brief, and every dashboard reads from.
2. **Surface.** Approved records sync to your dashboards and reports automatically. Your Monday brief draws from the same store. Your agents draw from the same store. One knowledge base, many views.
3. **Act.** Agents and automations read from the knowledge store to inform their work. The follow-up sequence knows about the vendor decision. The executive brief includes it. The operations docs reflect it.

That is the entire pipeline: capture → classify → route → store → surface → act. The full technical diagram is in **Appendix C**.

## The Governance Model: AI Recommends. Human Decides.

This is not a philosophical commitment. It is an operational rule encoded in every system we build.

Four roles, strict separation:

1. **The classifier** (an AI model) reads and categorizes. It assigns a confidence score. It never approves.
2. **The review agent** (a second AI step) evaluates ambiguous captures and writes advisory notes. It recommends. It never decides.
3. **The human** (you or someone on your team) approves or rejects items in the review queue. Final authority. No exceptions.
4. **The sync agent** (automated software) moves approved records to your dashboards and reports. It never evaluates content.

The human never sees the noise. Only the borderline cases. The AI never makes the final call on anything it is not certain about.

The 72-hour TTL is intentional. If a capture was not important enough to review in 72 hours, it was not institutional knowledge worth preserving. Expired items are logged, not lost — but they do not enter the canonical record.

**Decision tracking:** Any message that contains a formal decision (marked with a DEC-### pattern) is always routed to human review, regardless of confidence score. Decisions are never auto-approved.

**Categories:** Every capture is classified as one of five types — Decision, Research, Insight, Task, or Reference. This keeps your knowledge base structured and searchable.

## What Runs Behind Each Desk

The opening section introduced five automation types: triggers fire, scripts execute, workflows connect, personas advise, agents plan. Here is what that looks like when a desk is actually running.

**A day in the life of the Revenue Desk:**

**6:00 AM** — A scheduled job fires (that is a **trigger**). It checks your CRM for new leads, stalled deals, and follow-ups that are overdue.

**6:02 AM** — A program pulls the data, deduplicates it, and scores each lead based on engagement and firmographics (that is a **script**). No AI involved — just a deterministic pipeline that runs the same way every time.

**6:05 AM** — A high-value lead appeared overnight. The system kicks off a multi-step process: enrich the contact from public sources, draft a follow-up email, and schedule it for 9 AM (that is a **workflow** — multiple steps connected across systems).

**7:00 AM** — Your Monday brief is being assembled. The system loads your pipeline data, your cash position from Financial, and any decisions captured from last week's Slack threads. An AI

role synthesizes all of it into a two-page summary with recommendations (that is a **persona** — loaded with your business context for this one task, then gone).

**Ongoing** — If the system detects a pattern across multiple weeks — say, deals consistently stalling at the proposal stage — it flags the pattern and recommends a process change (that is an **agent** — persistent, cross-session, capable of spotting what a single script cannot).

Five types, working together, invisible to you. You read the brief. You approve the follow-up. The system handles the rest.

## How We Keep It Honest

### The Review Queue Is the Quality Check

Every borderline classification (confidence 0.70–0.84) goes to human review. When you approve or reject a borderline item, that decision becomes ground truth. Over time, the review queue teaches us where the classifier is strong and where it needs improvement.

### Automatic Cost Shutoff

Every classifier call is tracked: how many tokens went in, how many came out, what it cost. The numbers are transparent.

If classification calls exceed 500 per day, or the AI uses more than 100,000 tokens in a single session, or 5 errors happen in a row — the system stops itself. Remaining items go into a holding queue. An alert reaches the operator. No runaway costs. No surprise invoices.

### The Weekly Report

Every Monday, a knowledge base report summarizes: what was captured, what was approved, what expired in the review queue, and what the system cost to run. Transparency as a feature, not an afterthought.

### Your Data, Your Infrastructure

Every piece of your system runs on accounts you own and control. If our relationship ends for any reason, nothing disappears.

Each client has their own database on their own account. Your data never co-mingles with another client's data. We provide the system — the code, the configuration, the management. You provide the infrastructure accounts.

What this means in practice:

- **Your data** lives on accounts you own and pay for.
- **Your agents** are code in your GitHub repository.
- **Your credentials** stay in your vault.

- **Your dashboards** are in your Notion or Google workspace.

We build it. You own it. Walk away anytime — everything comes with you.

> You own everything. Your data lives on systems you control. You can leave anytime and take it with you.

---

*That covers how the system is built and why you can trust it. What follows is what happens after it is running — because the build is only the beginning.*

## What It Looks Like Running — The Manage

Deployment is not the end of the project. It is the beginning of the managed system. The businesses that get compounding value are not the ones that shipped the most agents. They are the ones that kept the systems running accurately for 18 months.

### What the First 90 Days Look Like

Month 1 was setup — the Quick Win delivered, the dashboard went live, the first Monday brief arrived. Here is what happens next.

### Months 2–3: The Expansion

Each month adds one workflow, one automation, one visibility layer. Every addition is measured against the time-recovered metrics established during Assess. Monthly check-in calls at the Visibility level. Bi-weekly strategy calls at Structure. Weekly calls at Managed.

The system is capturing knowledge in the background. Your review queue is active — borderline captures arrive, you approve or reject them, and the system learns from your decisions.

### Month 4+: The Compound Effect

This is where the Four Layers close the loop — and where the knowledge system reaches steady state.

The business is running the systems. The systems are capturing institutional knowledge from every communication channel (Layer 1). That knowledge is being structured into canonical hubs (Layer 2). Dashboards and briefs surface what matters without anyone asking (Layer 3). Agents and automations act on curated context — follow-ups happen because the system knows, not because someone remembered (Layer 4).

And then the loop feeds back: decisions made at Layer 4 are captured at Layer 1. The institutional memory of the business compounds. Every week, every agent output gets better because it draws from a deeper knowledge base. Every brief gets sharper because there is more context to synthesize.

This is the compounding loop. It takes 4 months to feel it. It is still running 3 years later. And it is the reason that knowledge loss — the invisible cost we opened with — stops compounding and starts reversing.

## How We Keep It Running

**Monitoring:** A health check runs every 10 minutes. Critical failures reach the on-call operator before business hours. Your daily Executive Brief includes system health alongside your business metrics.

**Updates:** Your business vocabulary, client names, product terminology — these are injected through your documentation, not baked into the system's rules. When your terminology changes, the context updates automatically.

**Security:** Every credential for every tool lives in a managed vault structure. When someone leaves, their access is revoked in one place, not twelve. Every agent has a strict list of what it can and cannot access. For deeper security coverage, the Security Package adds monthly digital footprint scans, breach detection, and credential hygiene reviews — see **Appendix D**.

**Infrastructure:** Two runtime modes ensure the system stays up. The automatic cost shutoff ensures no runaway costs — if any threshold is crossed, the system stops gracefully and alerts the operator.

**Quality:** The review queue is the quality check. Human decisions become ground truth. The weekly knowledge base report is the quality summary.

## The Advisor Pattern at Scale

As each new system comes online, the governance model extends. New agents are classified by what they can and cannot do before deployment. Approval authority is documented per agent. Confidence thresholds are set per domain — financial data may require a higher threshold than marketing content.

The fundamental rule does not change at scale: **AI recommends. Human decides.** This holds whether you have 3 agents or 54.

## Growing With You

Desk-stacking is the natural growth pattern:

1. Start with one desk (the one with the most acute pain) — $497/mo.
2. Visibility into that desk reveals adjacent problems.
3. Adjacent problems become the next desk — $994/mo for two desks.
4. At three desks ($1,491/mo), the Managed level ($1,249/mo) becomes cheaper and includes all six desks, strategic oversight, and access to premium capabilities. By month 9–12, most Structure-level clients have hit this inflection point.

**Operational Blueprints** extend what your desks already do. Once a desk is producing results, the adjacent need becomes obvious — the Revenue Desk reveals that your contact data is thin, or your client reports need branding, or your monitoring needs real-time alerts. Blueprints are pre-built automation packages that plug into any active desk. They are optional add-ons, never required. Full details and pricing are in **Appendix D**.

For businesses that outgrow the Managed level — custom agent development, enterprise integrations, or multi-location deployments — we build custom engagements scoped to the specific need.

### What You Own When We Leave

This is what separates a managed system from a vendor relationship:

- **Your data** lives on your infrastructure. Your database, your operational tools, your knowledge workspace.
- **Your agents** are in your code repository.
- **Your credentials** are in your vault.
- **Your process documentation** is in your docs.

If Sauce Technologies ceased to exist tomorrow, you would have everything. The business keeps running.

No lock-in. No exit fees. No hostage situations.

---

## The Numbers

### The Hidden Cost Baseline

Before any conversation about investment, here is what your business is already spending:

### Owner Time

| Cost Type | Conservative Estimate | Calculation |
|---|---|---|
| Owner doing admin | 15–20 hrs/wk | At $75–150/hr: $60K–$150K/year in admin work |
| Context switching (5+ tools) | 5–8 hrs/wk | 30 min/day per disconnected tool × 5 tools |
| Manual follow-ups | 3–5 hrs/wk | Checking on things that should be automated |
| Report building | 2–4 hrs/wk | Pulling data from multiple places |
| **Total hidden cost** | **25–37 hrs/wk** | **$100K–$280K/year in owner time** |

Research confirms these ranges. The Alternative Board's Business Pulse Survey found that SMB owners work 49.4 hours per week and spend 68% of that time on day-to-day operations. Time etc found that 36% of an entrepreneur's work week goes to administrative tasks.

ServiceNow reported that executives spend 16 hours per week on manual administrative work — two full working days.

## Staff Time

| Cost Type | Conservative Estimate |
|---|---|
| Data entry between disconnected systems | 4–8 hrs/wk per admin |
| Manual scheduling and coordination | 3–5 hrs/wk |
| Status updates and check-ins | 2–4 hrs/wk |
| Rework from missed handoffs | 2–3 hrs/wk |
| **Total per admin employee** | **11–20 hrs/wk** |

At a loaded cost of $35–55K/year per admin employee, you are losing **$18–28K/year per person** to manual work.

## Knowledge Loss — The Cost Nobody Tracks

| Loss Type | Estimated Cost | Source |
|---|---|---|
| Key-person departure | $50K–$150K per event | SHRM (50–200% of salary); 5–12 months to recover |
| Re-derivation time | 3–5 hrs/wk per knowledge worker | McKinsey (20% of workweek on information search) |
| Client context loss | 5–15% churn contribution | 42% of knowledge is unique to the individual (Panopto/YouGov) |
| Onboarding drag | 2x ramp time | Without structured knowledge base, new hires take 5–6 months to full productivity |

At $75–$150/hr owner time, re-derivation alone costs **$12K–$40K/year**. A single key-person departure can equal the entire first-year investment in the Managed level. These costs never appear on a P&L — but they compound every quarter the knowledge stays trapped.

## Structural Risks

| Risk | Impact |
|---|---|
| Single point of failure (key person) | If they leave: 5–12 months to recover. Cost: $50K–$150K in disruption. |
| No process documentation | Every new hire takes 2x longer to onboard. Organizations with structured onboarding see 82% better retention. |
| Data lives in someone's head | Decisions made on gut, not data. |
| Business cannot sell without the owner | Exit value discount: 5–25%, up to 40% in severe cases. 80% of listed businesses never sell. |

The question is not whether implementation costs money. It is whether your current approach costs more.

## ROI by Level

| Level | Year 1 Investment | Hours/Week Back | Annual Value | Payoff Month | ROI Multiple |
|---|---|---|---|---|---|
| Quick Win | $1,500–5K (one-time) | 4–10 | $15–40K | Immediate | 3–10x |
| Visibility | ~$3,300 | 5+ | $15–30K | Month 2 | 4.5–9x |
| Structure (1 system) | ~$9,000 | 15–20 | $30–50K | Month 3 | 3.3–5.6x |
| Managed | ~$20,000 | 20+ | $50–100K+ | Month 4 | 2.5–5x |

The right comparison is not zero. It is what you are already spending:

| Compared To | What the Numbers Show |
|---|---|
| Status quo | You are spending ~$60K/year in admin time (15 hrs/wk at $75–150/hr). The Visibility level is $3,300/year. |
| Hiring | A full-time ops person costs $50K before they are productive. Structure is $9K in Year 1 and live in 3 weeks. |
| Hiring an AI engineer | Fully loaded cost: $174K–$189K/year. That is 3.5–35% of total revenue for a $500K–$5M business. Managed is $20K. |
| DIY (Zapier, etc.) | If you have spent 80 hours building automations at your hourly rate, that is ~$6K — and you are still maintaining them. |
| AI consultant | SMB AI consulting runs $100–$300/hour. A typical project costs $10K–$50K. We build the fix for $1,500–5K and it keeps running. |
| Fractional CTO | $3K–$15K/month. Fractional COO: $5K–$20K/month. Managed at $1,249 is a fraction of a fraction — and it includes the systems, not just the advice. |
| Enterprise platform | HubSpot Professional runs ~$800–$1,200/month and most SMBs use a fraction of it. Research shows 80% of software features are rarely or never used. We build exactly what you need. |

The self-funding math: research shows 58% of AI-using SMBs save over 20 hours per month. At $75–$150/hour owner rates, that is $1,500–$3,000 per month in recovered time. Visibility at $149/month is a rounding error. Even Managed at $1,249 pays for itself in recovered owner hours alone.

## By Segment

| Segment | Typical Quick Win | Hours Saved | Annual Value | What Drives ROI |
|---|---|---|---|---|
| Healthcare & Wellness | $2–4K | 10–15 hrs/wk | $25–40K | Patient flow is revenue. Every missed inquiry is a lost recurring patient. |
| Food & Beverage | $2–5K | 8–12 hrs/wk | $20–35K | POS + inventory + labor visibility prevents waste. |
| Construction & Trades | $1,500–3K | 6–10 hrs/wk | $18–30K | Job costing visibility is the difference between profit and loss. |
| Financial & Legal | $3–5K | 12–18 hrs/wk | $35–55K | High hourly rates ($200+). Every admin hour is expensive. |
| Manufacturing & Industrial | $3–5K | 10–15 hrs/wk | $30–45K | Key-person dependency is existential risk. |
| Marketing & Professional Services | $1,500–3K | 5–8 hrs/wk | $15–25K | Project profitability is invisible until month-end. |
| Retail & Local Services | $1,500–3K | 5–8 hrs/wk | $15–25K | Customer retention is cheaper than acquisition. |

These are conservative estimates from real engagements. Time savings only. No revenue uplift included.

---

# How to Start

## Three Questions That Tell Us Everything

1. **What is eating the most time right now — yours or your team's?**
2. **Where does the answer to your most important business question live, and how long does it take to get it?**
3. **If you were unreachable for two weeks, what would break?**

If any of those answers made you uncomfortable, the Numbers section quantifies what that discomfort costs. Typically $100K to $280K a year in owner time alone.

If you can answer those three questions, we can tell you in 15 minutes whether we can help and what the first step looks like.

## The First Step

15 minutes. No pitch. Just: "What is eating your time right now?"

We start with an Assess. We show you what we see. You decide if it is worth fixing.

**Malcolm Sullivan** (847) 558-0405 malcolm@saucetech.io

> The business should run whether you are watching or not. That is what we build.

---

# Appendix A: Five-Tier Taxonomy

Full reference for the automation classification system used across all Sauce Technologies engagements.

| Tier | Term | Definition | Key Rule |
|---|---|---|---|
| T1 | **Agent** | AI-driven orchestration with planning, delegation, and cross-session continuity. Maintains state and makes routing decisions. | Agents **plan**. |
| T2 | **Persona** | Claude Code role activated by prompt and tool policy. Exists for a single session. Produces artifacts within a defined scope. | Personas are **ephemeral**. |
| T3 | **Script** | Deterministic Node.js program that runs on a schedule or trigger. May call AI as a tool, but the script controls flow. | Scripts **execute**. |
| T4 | **Workflow** | External automation flow (n8n, Zapier) on a third-party platform. Triggered by schedules or webhooks. | Workflows are **external**. |
| T5 | **Trigger** | CI/CD or cron definition that invokes a Script. The YAML or config that says "when" and "what" — not the thing that runs. | Triggers are **metadata**. |

## Key Distinctions

- **Agents plan; Scripts execute.** An agent can decide what to do next. A script runs a fixed pipeline every time.
- **Personas are ephemeral; Agents are persistent.** A persona is a session with a role prompt. An agent maintains cross-session state and dispatches work.
- **Workflows are external; Scripts are local.** Workflows live in n8n or Zapier. Scripts are `.cjs` files in the repo.
- **Triggers are scheduling metadata, not actors.** A GitHub Actions `.yml` file is a trigger. The script it calls is the actor.
- **Agents are optional.** The minimum viable runtime is scripts plus triggers. Agents add planning capability when complexity warrants it.

---

# Appendix B: Six Systems — Agent Index

**How to use this appendix:** Find the system below that matches your biggest pain point. Each one lists three **Starter Agents** that go live with your desk at the Structure level, plus Advanced Agents available as add-ons. You do not need to read all 54 — start with the Starter Agents for your desk. For which desk maps to which pain signal, see **Desk Discovery**.

**A note on scope:** Starter Agents are deployed in current client engagements. Advanced Agents are defined, scoped, and ready to build — but not all are running in production today. If a specific Advanced Agent matters to your decision, ask us about its deployment status during the Assess.

Each agent solves one specific, repeatable business problem. Every agent output is validated against the client's shared business context before it reaches the operator. For deep-dives on Starter Agents, see the **System Deep-Dive Library**.

## System 1: Executive — "Run the company. Not the inbox."

| ID | Agent | What It Does |
|---|---|---|
| 1.1 | Inbox Triage | Sorts, prioritizes, and drafts responses to executive email |
| 1.2 | Meeting Intelligence | Preps every meeting; captures decisions and action items |
| 1.3 | Delegation Tracker | Monitors delegated tasks; follows up when things go quiet |
| 1.4 | Report Synthesizer | Assembles business data into executive-ready reports |
| 1.5 | Decision Logger | Captures decisions with context, rationale, and owner |
| 1.6 | Calendar Intelligence | Protects focus time; audits calendar; resolves conflicts |
| 1.7 | Stakeholder Communications | Drafts board updates, investor letters, partner comms |
| 1.8 | Strategic Briefing | Daily morning brief with internal signals and market context |
| 1.9 | Travel & Logistics | Builds itineraries; preps for travel without manual coordination |

**Starter Agents (included at Structure):**

- **Executive Brief** — Your daily morning briefing with key signals and priorities
- **Strategic Intel** — Meeting prep, stakeholder updates, and partner communications in one place
- **Board Readiness** — Investor and board materials assembled from your business data

## System 2: Revenue — "Reps close. The system handles everything else."

| ID | Agent | What It Does |
|---|---|---|
| 2.1 | CRM Hygiene | Auto-audits, deduplicates, and enriches CRM records |

| ID | Agent | What It Does |
|---|---|---|
| 2.2 | Prospect Research | Deep-dive intelligence profiles on target accounts |
| 2.3 | Proposal Generator | Drafts customized proposals from deal data and templates |
| 2.4 | Follow-up Sequencer | Manages post-meeting follow-up sequences |
| 2.5 | Pipeline Reporter | Assembles pipeline health reports without manual exports |
| 2.6 | Contract Processor | Routes contracts for signature; tracks status; files copies |
| 2.7 | Sales Call Prep | Comprehensive briefing package 30 min before every call |
| 2.8 | Win/Loss Analyzer | Documents every closed deal with structured analysis |
| 2.9 | Lead Qualification | Scores and qualifies every inbound lead before a rep touches it |

**Starter Agents (included at Structure):**

- **Pipeline Tracker** — Weekly pipeline health delivered without manual CRM exports
- **Deal Escalator** — New leads scored and qualified before a rep touches them
- **Sequence Automator** — Follow-up sequences and contract routing handled automatically

## System 3: Marketing — "Ideas in. Executed output out."

| ID | Agent | What It Does |
|---|---|---|
| 3.1 | Content Drafting | First-draft production: blog, email, landing page, social |
| 3.2 | SEO Research | Keyword intelligence, topic briefs, competitor gap analysis |
| 3.3 | Social Scheduling | Formats, schedules, and manages content across all channels |
| 3.4 | Analytics Reporter | Aggregates multi-platform data into performance reports |
| 3.5 | Email Campaign Manager | End-to-end campaign production from segmentation to report |
| 3.6 | Campaign Coordinator | Orchestrates cross-channel campaigns from brief to launch |
| 3.7 | Brand Consistency | Reviews content for brand guideline compliance |
| 3.8 | Competitive Intelligence | Monitors competitors; delivers positioning insights |
| 3.9 | PR & Reputation Monitor | Tracks brand mentions, media coverage, reputation signals |

**Starter Agents (included at Structure):**

- **Content Calendar** — Drafts created, scheduled, and posted across channels
- **Campaign Launcher** — Cross-channel campaigns coordinated from brief to launch
- **Attribution Tracker** — Multi-platform performance data in one report

## System 4: Financial — "Clean books. Real-time visibility. Zero manual chasing."

| ID | Agent | What It Does |
|---|---|---|
| 4.1 | Invoice Processing | Automates intake, validation, and routing of vendor invoices |
| 4.2 | Expense Review | Reviews expense submissions for policy compliance |
| 4.3 | Cash Flow Monitor | Tracks real-time and projected cash; surfaces shortfalls |
| 4.4 | Collections Follow-Up | Manages full AR collections lifecycle to reduce DSO |
| 4.5 | Financial Report Assembler | Compiles standard financial statements and board reports |
| 4.6 | Budget Tracker | Monitors departmental spend against approved budgets |
| 4.7 | Payroll Prep | Prepares payroll data; validates inputs; coordinates approval |
| 4.8 | Tax Prep Coordinator | Organizes tax documents; coordinates with accountants |
| 4.9 | Audit Support | Assembles audit-ready packages; responds to auditor requests |

**Starter Agents (included at Structure):**

- **Cash Clarity** — Real-time cash position and projected shortfalls surfaced automatically
- **Forecast Lock** — Financial statements and budget tracking assembled without manual pulls
- **Churn Detector** — At-risk recurring revenue identified before it disappears

## System 5: Client Intelligence — "Every account monitored. Every risk surfaced."

| ID | Agent | What It Does |
|---|---|---|
| 5.1 | Health Score Monitor | Tracks client health signals; surfaces at-risk accounts |
| 5.2 | Onboarding Tracker | Coordinates onboarding milestones across all parties |
| 5.3 | QBR Prep | Assembles quarterly business review materials per account |
| 5.4 | NPS & Feedback Processor | Collects, categorizes, and routes customer feedback |
| 5.5 | Renewal Coordinator | Manages renewal timelines; drafts proposals; routes contracts |
| 5.6 | Escalation Manager | Detects, routes, and tracks client escalations to resolution |
| 5.7 | Expansion Intelligence | Identifies upsell and cross-sell opportunities |
| 5.8 | Client Reporting | Delivers regular performance and ROI reports to clients |
| 5.9 | Voice of Customer | Aggregates cross-channel feedback into strategic insights |

**Starter Agents (included at Structure):**

- **Risk Radar** — At-risk client accounts surfaced before they escalate
- **Expansion Spotter** — Upsell and cross-sell opportunities flagged automatically
- **Health Score** — Client health tracked across feedback, engagement, and delivery signals

### System 6: Operations — "Access granted. Tickets routed. Policies current."

| ID | Agent | What It Does |
|---|---|---|
| 6.1 | Recruiting Coordinator | End-to-end recruiting coordination from post to signed offer |
| 6.2 | Employee Onboarding | Full onboarding lifecycle from pre-start through 30-day mark |
| 6.3 | HR Documentation | Maintains handbook, policies, and benefits FAQ |
| 6.4 | IT Ticket Triage | Categorizes, routes, and tracks every IT support request |
| 6.5 | Systems Documentation | Keeps system config docs and integration maps current |
| 6.6 | User Provisioning | Manages system access for every onboarding and offboarding |
| 6.7 | Contracts & Legal | Drafts agreements, summarizes contracts, tracks lifecycle |
| 6.8 | Compliance Tracker | Monitors compliance obligations and certification renewals |
| 6.9 | Internal Systems Monitor | Monitors business-critical systems; alerts on failures |

**Starter Agents (included at Structure):**

- **Onboarding Automator** — New hire onboarding and system access handled end-to-end
- **Process Guardian** — System docs and compliance tracking kept current
- **Ops Pulse** — Internal systems monitored with automatic triage and alerting

---

## Appendix C: Knowledge Capture Pipeline — Technical Reference

### Pipeline Flow

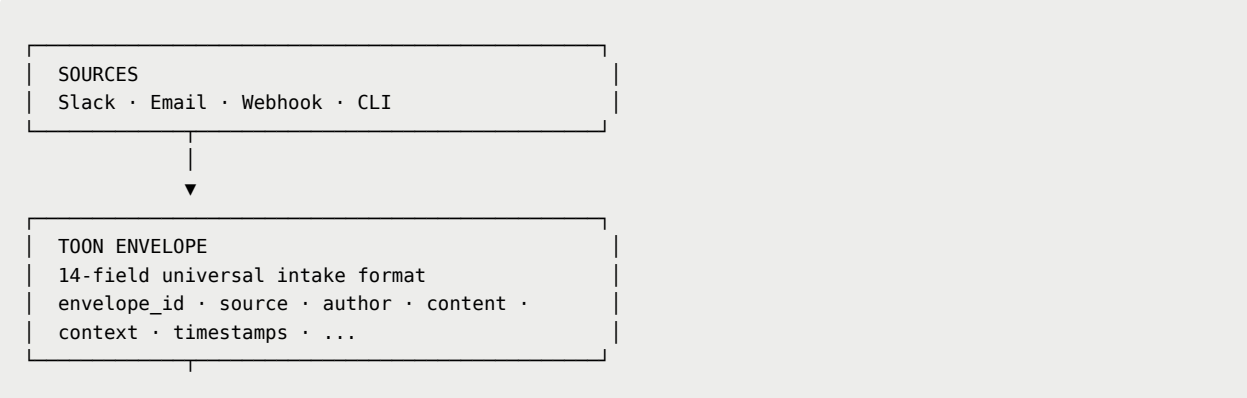

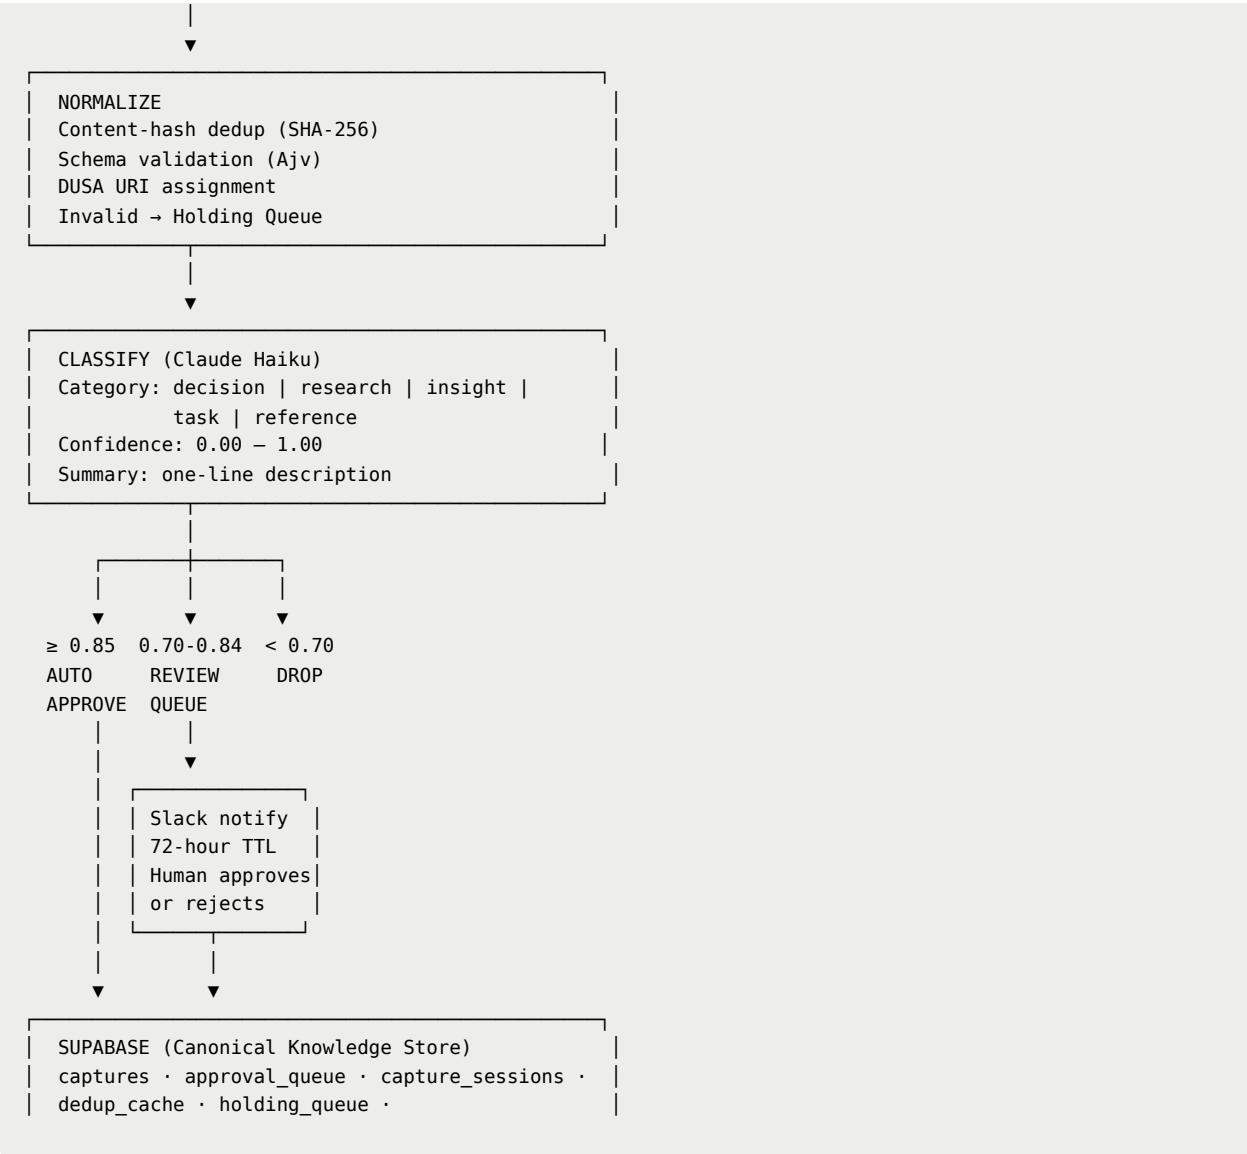


## Confidence Routing Thresholds

| Score | Action | Rationale |
|---|---|---|
| ≥ 0.85 | Auto-approve | Classifier is certain. Human does not need to review. |
| 0.70 – 0.84 | Human review queue | Classifier is uncertain. 72-hour TTL. Slack notification. |
| < 0.70 | Drop silently | Not institutional knowledge. Logged but not preserved. |
| DEC-### pattern | Always human review | Formal decisions are never auto-approved, regardless of confidence. |

For the business rationale behind these thresholds, see the Governance Model in **Why You Can Trust It**.

## Authority Hierarchy

Decisions are escalated based on scope:

```
system < client < team < operator
```

Each level has defined approval authority. A system-level decision (internal tooling choice) can be approved by the operator. A client-level decision (contract terms) requires operator review.

In a client engagement: **system** is an automated decision (tooling, agent configuration). **client** affects a specific customer account. **team** affects internal operations. **operator** is typically the business owner or operations lead — the person who reads the Monday brief and decides what to act on. Sauce Technologies serves as the technical advisor on architecture decisions — the operator (you) is always the final authority on your business.

## Cost Controls

| Control | Threshold | Action When Tripped |
|---|---|---|
| Daily classifier calls | 500/day | Remaining items → holding queue. Telegram alert. |
| Session token limit | 100K tokens/session | Session blocked. Telegram alert. |
| Consecutive errors | 5 errors | Automatic cost shutoff. Remaining items → holding queue. |
| Daily embedding calls | 1,000/day | Embedding paused until next cycle. |

## Session States

```
IDLE → ACTIVE → COMPLETE → IDLE (normal)
ACTIVE → DEGRADED → cleanup → IDLE (partial failure)
ACTIVE → SUSPENDED → retry → ACTIVE (rate limit)
ACTIVE → BLOCKED → retry → ACTIVE (dependency down)
```

Every session tracks: tick count, envelope count, error count, input tokens, output tokens, estimated cost in USD. All persisted to the `capture_sessions` table.

### DUSA — Data Architecture Reference

DUSA (Diatonic Universal System Architecture) is the technical name for the "one source of truth per data type" principle described in **The Data Architecture Question**.

| Data Type | Canonical Source | Role |
|---|---|---|
| Operations (pipeline, contacts, projects, tasks) | Airtable | Operational source of truth. Anything that moves. |
| Knowledge (institutional memory, agent outputs, classifications) | Supabase (Postgres + pgvector) | Knowledge store. Structured data, embeddings, cost tracking. |
| Documentation (process docs, agent code, playbooks) | GitHub | Version-controlled and auditable. |
| Views (dashboards, briefs, reports) | Notion, Google Drive | Read-only projections. Never the source of truth. |

**The critical rule:** Nothing writes back upstream. Data flows in one direction: source → system of record → view layers. View layers receive; they never author.

### BYOD Technical Details

Each client has their own Supabase project on their own account. The knowledge system installs into a `knowledge.*` schema namespace inside the client's existing database. No conflicts with their existing tables. One migration command.

| Client Provides | Sauce Provides |
|---|---|
| Supabase project (their account, their billing) | Schema migration (one command) |
| Slack workspace (bot token, scoped channels) | Agent code and pipeline configuration |
| Notion workspace (integration token) | Sync pipeline |
| Airtable base (access token, base ID) | Sync pipeline |

## Appendix D: Blueprints, Security & Add-Ons

### The Security Package

Security is not a level feature. It is a business requirement. We offer it as an add-on because most businesses want to start with operations. But digital footprint exposure compounds the same way operational debt does — quietly, until it does not.

| Level | Price | What It Covers |
|---|---|---|
| **Standard** | $30/mo | Monthly automated OSINT scan of your digital footprint. Scored risk report. Breach detection alerts. Credential hygiene review. |
| **Plus** | $50/mo | Everything in Standard plus 1Password Business setup and management, quarterly deep audit, and remediation guidance. |

**1Password MSP:** Every credential for every tool lives in a managed vault structure — Shared, Admin, Department. Your team stops reusing passwords. Your admin accounts stop being shared over email. When someone leaves, their access is revoked in one place, not twelve.

**SpiderFoot OSINT:** Monthly automated scans of your business's digital footprint. What domains are exposed? What credentials have appeared in breaches? What does your attack surface look like to someone who is not you? You get a scored report. Critical findings get flagged immediately.

## Operational Blueprints

Blueprints are pre-built automation packages that plug into any active desk. Each one solves a specific operational problem. You add them when the need becomes clear — not before.

Every blueprint comes in two levels: **Standard** (automated pipeline, sensible defaults) and **Plus** (adds custom rules, premium sources, or branded output).

| Blueprint | What It Does For You | Standard | Plus |
|---|---|---|---|
| **Data Aggregation & Enrichment** | Your CRM records get filled in automatically — company info, contact details, social profiles. You never touch the tools. | $35/mo | $65/mo |
| **CRM Segmenting & Scoring** | Your contacts get scored and ranked so your team knows who to call first and who to let go. | $30/mo | $55/mo |
| **Exportable Branded Collaterals** | Reports and proposals generate from your data automatically, in your brand, on schedule. | $25/mo | $50/mo |
| **Augmented Monitoring & Reporting** | Your systems are watched around the clock. You get a weekly digest, and critical changes get flagged immediately. | $40/mo | $75/mo |

The Security Package can be added alongside any blueprint configuration.

**How blueprints fit the Four Layers:**

- **Security Package** strengthens Layer 1 (Communication) — monitoring the external exposure of the business's digital footprint, catching threats where they originate.
- **Data Aggregation & Enrichment** strengthens Layer 2 (Documentation) — better data in the hubs means better outputs everywhere downstream.
- **CRM Segmenting & Scoring** operates at Layer 4 (Execution) — the system acts on scored, segmented data instead of flat contact lists.
- **Exportable Branded Collaterals** enhances Layer 3 (Visibility) — the views that reach your clients look like they came from your team, not a tool.

- **Augmented Monitoring & Reporting** spans Layers 3 and 4 — real-time visibility that triggers automated responses when thresholds are crossed.

Blueprints deploy into your own infrastructure the same way desks do — enrichment data lands in your Airtable, scoring models run against your Supabase, branded collaterals generate from your templates. No co-mingling. No data leaving your accounts.

Blueprints are optional. Most clients start with zero and add one or two after their first desk is running. The expansion is driven by what the desk reveals, not by a sales motion. If you are not sure where to start, Data Aggregation is the most common first add-on — it improves everything downstream.

---

*This guide reflects how Sauce Technologies builds and operates AI systems for real businesses as of April 2026. Pricing is illustrative and subject to change. ROI estimates are conservative, based on real engagements, and reflect time savings only — no revenue uplift included.*

> You built a business that works. We make it work without you.

**Sauce Technologies** Libertyville, IL malcolm@saucetech.io (847) 558-0405