KEYS STAY HOME

DIGITAL SUMMIT MINNEAPOLIS · THU AUG 13 · 1:00 PM

# Give AI the repeat work, not the keys.

MATT SCHWARTZ · FOUNDER, INSPRY + CHECKVIEW

(yes, all of them)

01 · THE GAP

# You automated the content. Not the operations.

| WHERE YOUR TEAM USES AI TODAY | | WHERE THE HOURS ACTUALLY GO | |
|---|---|---|---|
| Content and copy | ✓ | Site audits | MANUAL |
| Ad creative and variants | ✓ | QA and testing | MANUAL |
| Email and social | ✓ | Reporting | MANUAL |
| SEO briefs and research | ✓ | Monitoring and triage | MANUAL |

Why? Because nobody wants to hand an AI the keys to a live site, a customer record, or an invoice.

# The AI you already have, and forgot about

**Writing assistant**
API key living in someone's browser

**Ad platform AI**
who connected this?

**SEO tool agent**
has full account access. why??

**Automation bot**
left behind by an ex-vendor. still running.

**Chat thread**
someone pasted the FTP password in 2024

**Analytics AI**
no log of what it reads

"What can our AI actually touch right now?"

If nobody can answer that in one sentence, you don't have an AI strategy. You have AI sprawl.

# How it started. How it's going.

## HOW IT STARTED

- ✗ 14 browser tabs, 6 logins, one intern
- ✗ Passwords pasted into a chat window
- ✗ Every task re-explained from scratch
- ✗ "Who has access to what?" Silence.

## HOW IT'S GOING

- ✓ One door for every AI request
- ✓ Sign in with your work account. That's it
- ✓ Repeat work saved as named workflows
- ✓ "Who has access to what?" One page.

# One gated door

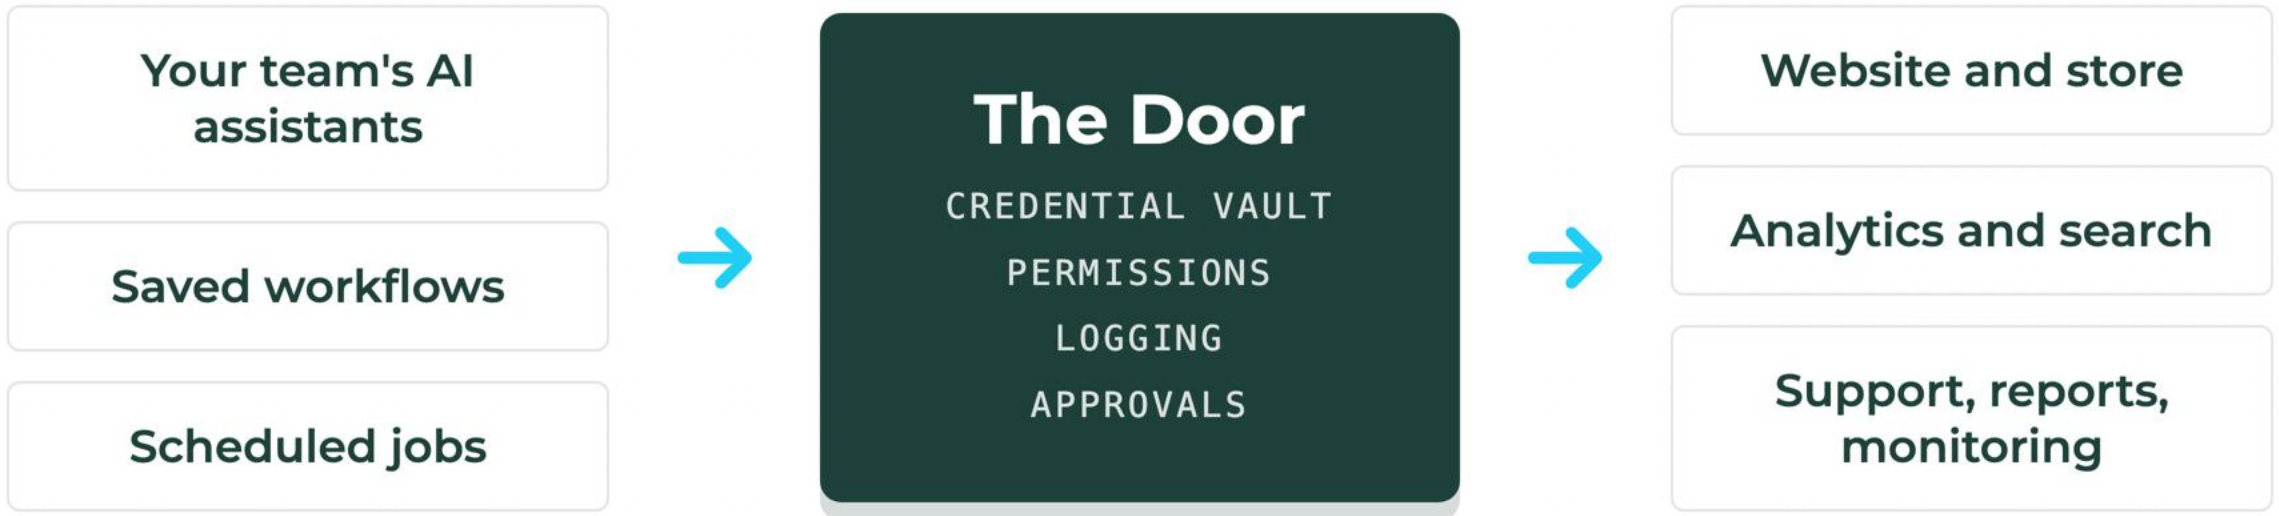


The AI never touches a credential. It carries a badge, not a key: opens what the role allows, logs every swipe, dies the moment you revoke it.

# Permissions you already understand

## READ

### Look, don't touch

Every new connection starts here. Reports, audits, monitoring. Zero risk to anything live.

## EDIT

### Propose, don't publish

The AI drafts the fix or the change. A human approves anything that touches customers.

## ADMIN

### People. Not agents.

Held by humans, granted rarely, logged always, revocable in one place.

If your team understands **editor vs. admin** in your CMS, they already understand AI access control.

# Sort every recurring task into three buckets

1

## Runs on its own

read-only · reversible · boring

- Site and SEO audits
- Reports and dashboards
- Monitoring and health checks

2

## Needs a human yes

anything that changes what's live

- Fixes and deploys
- Settings and config
- Customer-facing content

3

## Stays human

on purpose. forever.

- Client relationships
- Judgment calls
- Strategy and priorities

Then automate the biggest time sink first. Not the loudest request.

# The one-sentence site audit

TEAM CHAT · ANY WEEKDAY

Run the full audit on demo-client.com

performance · security · uptime · plugins · analytics · forms · search · errors

Every system checked, findings ranked by impact, report filed where the team already works.

BEFORE

**[ X hours ]**

by hand · 6+ logins · quarterly, if we were honest

AFTER

**[ Y minutes ]**

one saved workflow · every client site · on a schedule

Read-only, all of it. Bucket 1. No gate needed, no risk taken.

# "This report could have been a cron job."

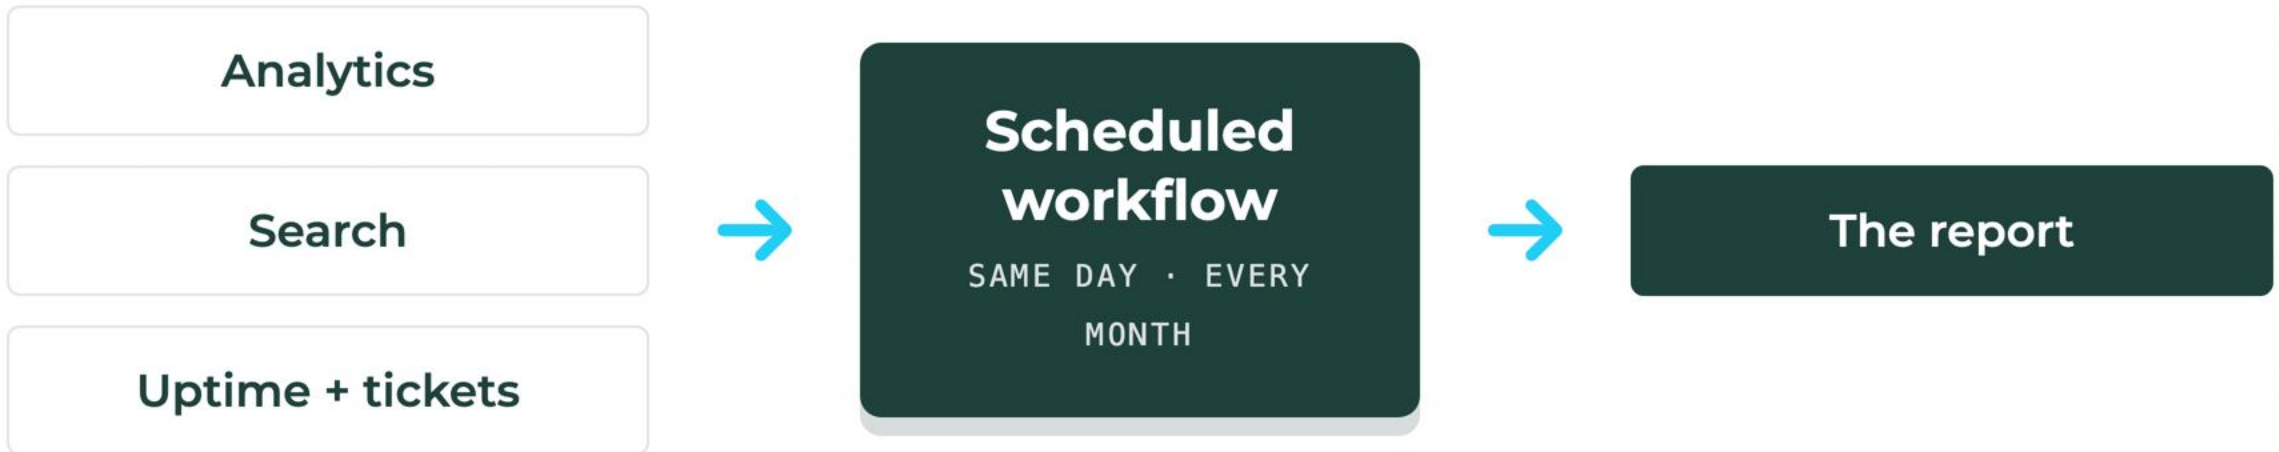


Pulls from the source systems directly. **No copy-paste, no stale exports.**

The human adds the narrative. **Judgment stays in the report.**

09 · THE CATCH

# While everyone slept, a browser with no cookies visited every site we manage

01 · MONTHLY, OVERNIGHT

### Visit like a stranger

Fresh browser, no cookies, never touches the consent banner. Records everything that fires anyway.

02 · THE FINDING

### Trackers before consent

On several sites, ad and analytics tags fired before the visitor agreed to anything. The banner looked fine. The behavior wasn't.

03 · HUMAN DECIDES

### Flagged, not fixed

Ranked by severity on the dashboard. A person confirms each one and decides the fix. Nothing auto-changes.

Those are your marketing tags. A compliant-looking banner and a compliant site are not the same thing, and only one of them shows up in a lawsuit.

# An agent can raise the alarm. Only a human can turn it off.

### It didn't fix it silently

No un-reviewed change ever touches a live site. Even a "safe" one.

### It didn't clear its own alarm

An agent can't decide it was a false positive. That call is human.

A FALSE ALARM COSTS ONE REVIEW. A SILENTLY-CLEARED REAL ONE COSTS AN OUTAGE MID-CAMPAIGN.

# Three things we got wrong (so you don't have to)

## 01 Over-permissioned early

Gave the first connections more access than the job needed, because it was easier.

**FIX**

**Every connection starts read-only. Widen only on evidence.**

## 02 Paid AI to re-think repeat work

Same steps re-figured-out every month. Slow, inconsistent, and the token bill adds up.

**FIX**

**Anything you'll run twice becomes a saved, named workflow.**

## 03 Assumed every vendor gates alike

Some platforms lock everything down. Others hand you full access, and the responsibility.

**FIX**

**Check what your platform blocks for you before you connect it.**

# Your first 30 days

**WEEK 1**

## List and sort

Write down the recurring work. Sort it into the three buckets. That's the whole week.

**WEEK 2**

## Automate ONE

Pick one read-only workflow. A report or an audit. Just one.

**WEEK 3**

## Surface it

Findings go to the dashboard your team already watches. Escalation only.

**WEEK 4**

## Add one gate

Your first approval step. Only then consider write access.

**FOR DIRECTORS** Don't run these systems yourself? This list is the conversation to have with whoever does: your agency, your web team, or IT.

# Yes, it's a real system

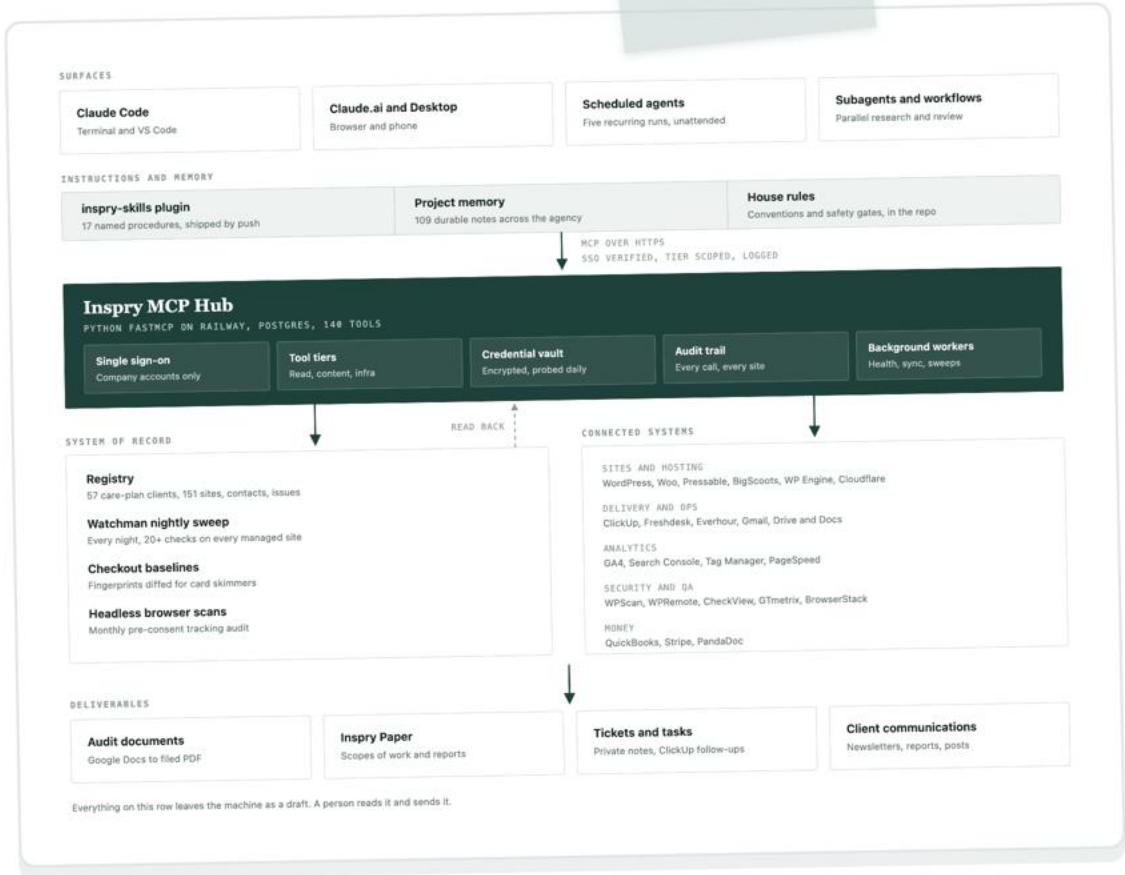


185 sites. 140 tools behind the door. 20+ checks on every site, every night.

The full writeup, with the architecture and the guardrails, is behind the QR on the next slide.

*send this one to your dev team →*

# Repeat work to the AI. Keys stay with you.

ONE GATED DOOR | THREE BUCKETS | ESCALATION ONLY

digitalsummit.inspry.com

The three-bucket worksheet, the 30-day checklist, and the technical deep dive for your dev team.

Find me at the Pressable booth. Bring your most annoying recurring task. We'll figure out which bucket it goes in.