



PSOF V1.0 · PAPERCHASE CONSULTING · PAPERCHASEWEBB, INC.

Governed Systems Operations for the Agentic Enterprise

title: Governed Systems Operations for the Agentic Enterprise subtitle: PaperChase Systems Operations Framework (PSOF) v1.0 ref: PCW-WP-001 version: 1.0 date: 2026-08-01 entity: PaperChaseWebb, Inc. surface: PaperChase Consulting · paperchaseconsulting.com status: public

PaperChase Systems Operations Framework (PSOF) v1.0 White paper · PaperChaseWebb, Inc. · Honolulu, Hawaii · August 2026 REF: PCW-WP-001 · Method SSOT: PCW-PSOF-001

Observe everything. Authorize deliberately. Automate safely. Prove outcomes.

1. EXECUTIVE SUMMARY

Most organizations do not fail for lack of software. They fail because nobody owns the operating layer: the connective tissue between infrastructure, SaaS, data, workflows, automations, and AI agents.

Tools multiply. Incidents stay reactive. AI increases speed without control. Backups exist without restore proof. Vendors renew without scorecards. Agents get tools before services have owners, telemetry, or kill switches.

PaperChase Consulting — a market surface of PaperChaseWebb, Inc. — sells controlled operational outcomes. We design, install, govern, and improve the systems that keep a business available, secure, observable, recoverable, cost-aware, and progressively autonomous.

The method is PSOF v1.0: a proprietary delivery and operating framework with eight control domains, a maturity model from Unknown to Governed Autonomous, production gates, and an AI authority ladder (A0–A4). PSOF is standards-informed. It is not certification theater.

The commercial model is paid to play: diagnose, stand up, maintain — with named packages, written terms, and first-of-month retainers. Entry is a free scorecard (T0) and a fixed Systems Ops Audit (A1 · \$1,000 for standard SMB scope).

2. THE PROBLEM

2.1 SPRAWL WITHOUT OWNERSHIP

Modern stacks mix cloud, SaaS, local systems, integrations, and AI tooling. Without a service catalog and named owners, every incident becomes archaeology. Knowledge lives in people, not records.

2.2 AUTOMATION WITHOUT AUTHORITY

Automation and agents that write to production without intended use, evaluation, least privilege, and human accountability create new classes of risk: tool misuse, identity abuse, cascading failures, and irreversible writes that nobody can roll back.

2.3 METRICS WITHOUT RECOVERY

Dashboards that do not map to user-facing SLIs, and backups that have never been restored, create false confidence. Continuity is proven only by exercise.

2.4 ADVICE WITHOUT INSTALLATION

Generic AI consulting and slide decks do not leave an operating layer. Clients need systems people use: ownership, evidence, gates, and optional paid continuity — not another wishlist.

3. THESIS

The product of serious systems work is not “more AI.” It is a governable operating layer:

- every critical service has a named owner and measurable outcomes;
- every change and agent action has a defined authority class;
- every control has evidence, not assertion;
- recovery is tested, not assumed;
- improvement is driven by real incidents, failed changes, cost, and user impact.

Autonomy is earned. You cannot skip to “governed autonomous” because a demo works. Each maturity level requires proof of the levels below.

4. WHO WE ARE

SURFACE	ROLE	BOUNDARY
PaperChase Consulting (paperchaseconsulting.com)	Public front door: assessment, packages, managed ops	Market and service surface — not a second OS
PaperChaseWebb, Inc.	Contracting entity, client relationship, commercial governance, founder strategy	Named on invoices and agreements
PaperChase Labs	Technical implementation: architecture, integrations, automation, dashboards, docs	Delivery vertical under the same system
AI Install Protocol (AIIP)	Method packaging, training, partner path	Scale after delivery is repeatable
Bloo-Collar	Productized continuity for a specific product line	Never leads consulting pages

Chase Webb is CEO of PaperChaseWebb, Inc., based in Honolulu, Hawaii. Delivery is remote-first. The company is SDVOSB-capable for government-adjacent work where relevant (CAGE on formal government materials).

5. WHAT PSOF IS — AND IS NOT

5.1 PSOF IS

- A proprietary framework for systems operations, AIOps, and agentic enterprise operations.
- A practical crosswalk to recognized service-management, reliability, security, AI-governance, observability, and cost-management practices.
- A repeatable method that turns consulting into assessable controls, implementation gates, operating evidence, and measurable improvement.

5.2 PSOF IS NOT

- An ISO certification, ITIL accreditation, cybersecurity audit opinion, SOC report, or legal-compliance guarantee.
- Permission for an AI model to act freely in production.
- A promise of 24/7 response unless a signed, staffed service plan is priced for 24/7 coverage.
- Another parallel PaperChase operating system or standalone company layer.

Public claims discipline: PSOF is standards-informed or crosswalked to practices. Use “certified,” “compliant,” or “conforms to” only when independent evidence exists.

6. STANDARDS BASIS (INPUTS, NOT CERTIFICATES)

REFERENCE	PSOF USE	GUARDRAIL
ISO/IEC 20000-1:2018	Service-management design	Do not claim ISO certification
ITIL practices	Incident, problem, change, configuration, SLM, suppliers	“Informed by ITIL practices” unless accredited
NIST CSF 2.0	Govern, Identify, Protect, Detect, Respond, Recover	Documented current/target profile + evidence
CIS Controls v8.1	Prioritized technical safeguards	Scope implementation group
ISO/IEC 27001:2022	Risk-based security governance	Do not imply accredited ISMS
NIST SP 800-61r3	Incident response design	Guidance, not certification
NIST AI RMF 1.0 + GenAI Profile	AI governance, map, measure, manage	Intended use, eval, human oversight
ISO/IEC 42001:2023	AI management system patterns	Do not claim ISO 42001 certification
Google SRE	SLIs, SLOs, error budgets	SLOs internal; SLAs contractual
OpenTelemetry	Vendor-neutral telemetry design	Tool choice follows monitoring design
FinOps Framework	Cost allocation and value	Value-linked optimization
OWASP Agentic Security	Threat modeling for agents	Industry guidance

Standards registers are dated and reviewed at least annually (AI RMF revisions are expected).

7. OPERATING PRINCIPLES

- Business service first — Every system, alert, automation, and control maps to a service, owner, user impact, and value metric.
- Evidence over assertion — No control is “implemented” until evidence shows it exists, works, and has an owner.
- Least authority — Humans, services, and agents get minimum access for the approved action and period.
- Observability before autonomy — Visible, baselined, recoverable — then AI production action authority.
- Reversible by default — Previews, staged writes, canaries, backups, defined rollback.
- Human accountability — AI may recommend or execute within policy; a named human remains accountable.
- One source of operational truth — Services, assets, incidents, changes, controls, and decisions link through canonical records.
- Continuity is proven — Backups are not recovery. Restore tests prove continuity.
- Improve from real operations — Incidents, failed changes, overrides, costs, and user feedback set the backlog.

8. EIGHT CONTROL DOMAINS

#	DOMAIN	CORE QUESTION	PRIMARY OUTPUT
01	Service Governance & Value	What service matters, who owns it, what outcome is required?	Catalog, owners, policies, SLOs, value case
02	Assets, Configuration & Architecture	What exists and how does it depend on everything else?	Inventory, configuration, dependency and data-flow maps
03	Observability, Reliability & Capacity	Can we see failure and sustain demand?	Telemetry, dashboards, SLOs, alerts, capacity plan
04	Security, Identity & Risk	Who can do what, to which data and systems?	Risk register, access matrix, security baseline, evidence
05	Incident, Problem & Change	How do we restore, learn, and change production safely?	Runbooks, incident log, RCAs, change records, rollback
06	Automation, AI & Integration	What can be automated, at what authority, with what evaluation?	Workflow specs, agent cards, evals, tool policies, receipts

#	DOMAIN	CORE QUESTION	PRIMARY OUTPUT
07	Continuity & Recovery	How fast and completely can the business recover?	RTO/RPO map, backup policy, restore proof, exercises
08	Cost, Suppliers & Improvement	Is the estate valuable, controlled, and improving?	Cost allocation, supplier scorecards, improvement backlog

Each domain has minimum controls, required evidence, and core measures. Assessments score domains 0–5 and produce a current/target profile with a 90-day control-and-value roadmap.

9. MATURITY MODEL

LEVEL	NAME	MEANING
0	Unknown	Incomplete or undocumented systems and owners
1	Visible	Critical services, assets, owners, baseline metrics identified
2	Controlled	Access, change, incident, backup, and security controls documented and used
3	Reliable	SLOs, actionable monitoring, tested recovery, stable delivery
4	Predictive	Correlation, anomaly detection, proactive problem and cost management
5	Governed autonomous	Bounded agents execute approved reversible actions with eval, audit, budgets, kill switch

Reassess quarterly for managed clients, and after material incidents, acquisitions, migrations, or AI authority changes.

10. AI AUTHORITY LADDER

CLASS	NAME	MEANING
A0	No AI authority	AI not used operationally for this service
A1	Observe / summarize	AI reads allowed data; no production side effects
A2	Recommend	AI proposes actions; human approves all writes
A3	Stage / draft	AI prepares reversible changes; human gates production apply
A4	Execute within policy	AI executes bounded actions with eval, audit, budget, kill switch

Autonomy-control gate: no production write path for agents until intended use, system card, risk class, evaluation set, authority class, monitors, and retirement owner exist.

11. DELIVERY LIFECYCLE AND PRODUCTION GATES

Qualify → Assess → Design → Build → Verify → Transition → Operate → Improve

Hard rules:

- No production write access before design and change-approval gates.
- No autonomous production action before maturity and autonomy-control gates.
- No managed-service transition without documentation, support boundaries, and billing authorization.

Engagement integrity is measured by gate evidence, client communication, scope control, and commercial control — not by demo volume.

12. COMMERCIAL MODEL

Direction (PCW-DIRECTION-001): price to stand up, price to maintain. Prefer many small paid entities over three unpaid almost-deals. CRM and written terms. Paid to play.

PACKAGE LADDER (SUMMARY)

CODE	OFFER	PRICE (PUBLIC)	NOTES
T0	Systems Operations Scorecard	Free	Self-serve qualifier · eight domains
A1	Systems Ops Audit	\$1,000 fixed SMB scope	Maturity scores · top risks · 90-day roadmap · 100% credit to B1/B2 within 14 days of findings
A2	Operator Retainer (Lite)	\$500–\$1,000 / month	\$100/hr · \$500 per 5-hour block · 5–10 hrs · billed first of month
B1	14-Day Controlled Automation Sprint	From \$5,000	One bounded workflow · monitoring · docs · acceptance
B2	30-Day Systems Operations Deployment	From \$10,000	1–3 workflows · control plane · dashboard · runbooks
C1	Managed Systems Operations	\$1,500–\$3,000 / month	Continuity after build · evidence pack · windowed support
C2	Managed Ops Priority	\$3,000–\$7,500 / month	Expanded window · priority queue
D1	Executive Ops Command Center	From \$25,000	Multi-workflow · exec visibility
D2	Enterprise Ops Transformation	From \$50,000 SOW	Multi-service target model

Complex multi-entity or regulated scopes are custom-quoted.

PAYMENT TERMS (STANDARD)

- Audit / assessment: due at checkout or schedule (Net 15 for approved B2B).
- Projects ≤ \$10,000: typically 60% / 30% / 10%.
- Managed ops: month in advance after written authorization.
- Net 15 standard; Net 30 only with approved credit.

13. HOW CLIENTS ENGAGE

1. ■ Free scorecard (T0) — honest self-score across eight domains.
2. ■ Fit call (15 minutes) — confirm scope and constraints.
3. ■ Systems Ops Audit (A1) — paid evidence pack and roadmap.
4. ■ Sprint or Deployment (B1/B2) — bounded build under gates; A1 credit if contracted within 14 days of findings.
5. ■ Retain (A2 or C1/C2) — paid continuity with hours, evidence, and first-of-month billing.

Contracting entity on invoices and agreements: PaperChaseWebb, Inc.

14. WHAT WE REFUSE

- Slideware as the product.
- Unbounded AI in production without authority class and eval.
- Free friendship labor framed as “partnership.”
- Leading consulting with unrelated product brands.
- Claiming certifications we have not earned.
- Managed ops without documentation, access, and billing authorization.

15. CLOSING

The agentic enterprise is not won by the model with the best demo. It is won by the organizations that can see their services, control change, prove recovery, and grant automation only the authority it has earned.

PSOF is how PaperChaseWebb, Inc. delivers that work through PaperChase Consulting.

Start with evidence, not theater.

CONTACT

PaperChaseWebb, Inc. PaperChase Consulting · <https://paperchaseconsulting.com> Honolulu, Hawaii · Remote-first

Chase Webb · CEO mgmt@paperchasewebb.com · (808) 321-9578 <https://paperchasewebb.com>

Book: <https://paperchaseconsulting.com/contact#book> Pay Systems Ops Audit (\$1,000): <https://paperchaseconsulting.com/contact>

DOCUMENT CONTROL

FIELD	VALUE
REF	PCW-WP-001
Title	Governed Systems Operations for the Agentic Enterprise
Method	PCW-PSOF-001 v1.0
Direction	PCW-DIRECTION-001
Offers	PCW-OFFER-001 (prices as of 2026-08-01 site)
Version	1.0
Status	Public white paper
Classification	Public
Owner	Chase Webb, CEO
Entity	PaperChaseWebb, Inc.

This white paper is for information. Engagement terms are defined in written SOW / MSA. PSOF is proprietary method IP of PaperChaseWebb, Inc. Standards references are design inputs, not claims of certification.

PaperChaseWebb Inc. · Honolulu, HI · SDVOSB · CAGE 18N61 · mgmt@paperchasewebb.com · (808) 321-9578

Confidential & Proprietary — © 2026 PaperChaseWebb Inc. All rights reserved. Unauthorized distribution or commercial use is prohibited.

PCW-WP-001 · SYS-001 house format