



Agentic **WorkForce**

AI TEAMMATE GOVERNANCE BRIEF

AI teammates add capacity without adding risk.

How Empower AI defines, levels, gates and authorizes the AI teammates on our federal programs, and why the governance comes before the capability.

October 2026

Empower^{AI}
Elevate the Mission

THE PROBLEM WE SOLVE

The people who know the system best are too busy running it to modernize it.

Critical systems depend on years of institutional knowledge: how the workflows actually operate, where the dependencies live, and which exceptions keep the mission moving. That knowledge sits with the same people responsible for keeping those systems running. They don't have the capacity to run today's mission and teach the modernization team everything they know, so the knowledge baton never gets passed.

Nobody staffs the handoff. We do. Our AI teammates take on defined work alongside your people, creating the capacity to modernize while your experts keep the mission moving. This brief explains how we keep that capacity governed.

01 • DEFINITION

An AI teammate is a governed member of the workforce.

It performs defined work alongside people, within established permissions, oversight and accountability. Not everything that touches AI is a teammate, and being precise about the difference is the first control.



Automation

Deterministic, scripted and API work: identity and access, ticket handling, records movement.

No model in the loop, and we don't call it an agent.



Skill

A capability that does one task. Add a harness and it can become part of a teammate's role.

Not a teammate on its own.



AI teammate

A collection of skills plus a harness, holding a single role. Its own identity, its own queue, and a named supervisor who signs its output.

Not a chatbot, a copilot, or a tool someone opens.



Delivery team

People and AI teammates orchestrated together against a program's work. The teammate joins the team. It does not replace it.

Not a human team with better tools.



Agentic WorkForce. The whole of it: the roster of AI teammates and the delivery teams they sit in. How Empower AI delivers services. **Not a product we sell.**

02 · CLIENT ZERO

Every AI teammate runs inside Empower AI before it runs in your mission.

AI adoption stalls because organizations find out too late how hard it is to make AI work inside their own walls. So every teammate in this brief ran inside Empower AI first, through our own delivery, service management and compliance functions. We found what broke and closed the gaps. Your teammates arrive already de-risked, from an environment that looks and feels like yours.

THE FIVE TESTS EVERY DEPLOYMENT HAS TO PASS

 Data control Your data stays where your policies say it lives.	 Deployment portability It runs in your environment, not only in ours.	 Integration It works inside the systems and workflows you already have.	 Sustainment Someone owns it after go-live: monitoring, retraining, ATO maintenance.	 Post-deployment governance Its level, permissions and supervisor stay visible and changeable.
--	---	---	--	---

No fully autonomous AI teammates run on mission-critical logic; autonomy is reserved for well-defined, low-risk work.

The full account is in our white paper, *From Client Zero to Governed Autonomy*, at empower.ai/insights.

CAPACITY ON THE RECORD

What changes when the teammate joins the team.

~10M

pages of federal health payment records processed, with a person making every determination

1 hr → 10 min

a bulk download in federal immigration case processing

19,676

staff hours returned to the mission on a government-wide enterprise IT program

The measurement period, baseline and population behind each figure are available on request.

03 · THE ROSTER

Eight AI teammates at work today.

Every one has a human supervisor.

Each teammate holds one role on a team of people, under one identity model, one approval path and one supervisor structure. None of them works alone, and none of them runs above the level we publish for it.

AI TEAMMATE	PROGRAM	WHAT IT DOES	LEVEL
 AI Medical Records Reviewer	CMS CERT · air-gapped	Reads the clinical file and bookmarks the evidence behind a payment decision.	L2
 AI Network Engineer	GSA DIGIT	Forecasts the failure, raises the event, diagnoses the cause, proposes the remediation.	L2
 AI Logistics Coordinator	Email-driven workflows	Executes workflows that used to run through email.	L2
 AI Service Desk Technician	DISA JSP	Answers the request before it becomes a ticket.	L2
 AI Recruiter	Staffing at scale	Turns requirements and applicants into an optimized workforce plan in week one.	L1
 AI Data Analyst	GSA DIGIT	Gives plain-language answers across tools and datasets.	L2
 AI Software Developer	USCIS SPEED	Writes, tests and ships code inside the same pipeline and review as the team.	L2
 AI Cloud Engineer	USCIS SPEED	Provisions and hardens cloud, and proposes the change for an engineer to approve.	L2



A named human supervisor signs each teammate's output.

L1 is Output. L2 is Actions and tools. Everything in production on our roster today runs at L1 or L2. The five levels are on the next page.

04 · DELEGATED CONTROL

People decide what AI is allowed to decide.

Five levels. Each one hands the model a different class of decision. We publish the level every teammate runs at, log it, and can change it without redeploying the teammate.

L1	Output · <i>What should I say?</i> The model decides what to generate, and nothing else. No tool is chosen, no branch is taken, no plan is owned.	IN PRODUCTION
L2	Actions and tools · <i>What tool should I use?</i> The model chooses among tools we defined, inside a boundary we drew. It acts on a system only through a door we built and can close.	IN PRODUCTION
L3	Decision points · <i>Which branch should I take?</i> The model chooses between the legs of a workflow a human designed. The path exists before the model runs.	NOT DELEGATED
L4	Process and plan · <i>What plan should I follow?</i> The model owns and revises the plan itself. Predictability is traded away on purpose, so the approval path has to be strongest.	NOT DELEGATED
L5	Initiation and time · <i>When should work begin or resume?</i> The system starts or continues work from triggers or policies, without a person asking.	NOT DELEGATED

Set per role

The level is part of the role definition, agreed with the program before the teammate starts work.

Published and logged

Anyone accountable for the program can see the level, and every change to it is recorded.

Revocable without redeploying

If the work changes or the evidence says so, the level comes down the same day.

Use the lowest viable level. A more agentic system trades predictability and cost for flexibility, so the level is a design decision we justify, not a score to climb. Everything in production on our roster today runs at L1 or L2.

05 · LIFECYCLE

Nothing reaches production without clearing three gates.

No gate is skipped because a schedule is tight. Every teammate on the roster was built this way, here, before it was built anywhere else.

FIVE STAGES, THREE GATES



GATE 1 · AFTER SCREEN

Use-case fit

Problem definition, data and permissions, error tolerance, explainability, sustainability.

GATE 2 · AFTER DEVELOP

Pilot exit

A bounded pilot completed, the agreed test phases run, and an evidence package prepared.

GATE 3 · AFTER SHADOW

Production readiness

Data, tools, ownership and governance, operation, sustainment, and the human-in-the-loop design.

THE GATES STOP THINGS

REJECTED AT GATE 1

Deterministic rules for the CMS CERT use case. Exact-page identification from unstructured OCR text could not be solved without a model, and we said so before building it.

KILLED AT GATE 2

A retrieval-augmented pipeline on a 7B open model. It was piloted, it did not clear pilot exit, and it did not ship.

Model lineage recorded in full. Human review required on high-risk decisions. Audit evidence documented at every gate. Grounded in the NIST AI Risk Management Framework and NIST SP 800-53, and aligned to OMB M-25-21, M-25-22 and M-26-04. Applied at CMS, GSA, DISA, DHA, USCIS and DIU.

06 · AUTHORIZATION

Every path to authorization, mapped before we deploy.

Every AI teammate runs under someone's ATO; the question is whose, not whether. We agree the path in writing with your ISSO before anything goes live.

PATH 1**Inherit the boundary**

TYPICALLY WEEKS

The teammate runs inside a system that already holds an ATO and does not change its risk posture.

PATH 2**Significant change review**

TYPICALLY 1–3 MONTHS

The teammate materially changes the authorized system, and the AO accepts it under the existing ATO. Most agentic deployments land here.

PATH 3**New ATO or FedRAMP**

TYPICALLY 6–24 MONTHS

The teammate stands up its own infrastructure, or it is hosted and access is sold.

We never claim “no ATO required.” What we provide on each path is in the one-page Path-to-ATO Checklist, written for ISSOs and authorizing officials.

CONTROL THAT LASTS

Built for eight teammates and for eight hundred thousand.

We believe federal agencies will soon have as many AI teammates as people. The management model has to hold at both scales, and eight is the only honest time to design it.

Who is doing what

Every teammate carries its own identity, not a borrowed human service account. Attribution per action.

What they may do

Permissions per role and per action, revocable where a person can find them. A level that is set, logged and changeable.

Who supervises them

A named human accountable for a roster, an approval path that scales, and continuous audit.

NO VENDOR LOCK-IN, IN LINE WITH OMB M-25-22**Open by design**

Open architecture, standard APIs and exportable data.

Model-agnostic

Underlying models can change without re-procuring the capability.

Watched in production

Live monitoring against baselines; a breach triggers review and rollback.

Transferable by contract

Run it yourself if you choose. That is a supported outcome.

Let's start a conversation.

questions@empower.ai

EmpowerAI