

The Supervised Digital Workforce

From Shadow AI to Shared Advantage

Ugur Ozer

The Argument in Brief

A complete overview from the book. Original pagination is retained. This sample includes the publication notice and the selected overview only.

Books, purchase links and free working templates

airiskmanagement.ca/books

Publication Notice

Copyright 2026 Ugur Ozer. All rights reserved.

Published by Aeon Digital Publishing.

Sources updated through 2026.

This book provides general management guidance, not legal, accounting, tax, employment, privacy, cybersecurity, investment, or regulatory advice. Verify current duties, evidence, and local requirements for the system and jurisdiction at issue.

Public cases are cited. Scenarios are labeled by type and contain no private outcome claim.

All Leadership Principles are original author text. Historical quotations are used for perspective, not as evidence.

The Argument in Brief

The short version

Three different kinds of AI work often sit behind one label. An employee may prompt a general assistant, interpret the answer, and act manually: **employee tool use**. An embedded system may propose work that a person accepts, changes, or rejects: a **copilot**. A system may pursue a goal across tools, change records, communicate, or transact within limits: a **semi-autonomous agent**.

After this full definition, **agent** means **semi-autonomous agent**. The three modes need different controls because they hold different authority. The executive task is to see the work, decide which tasks should change, grant the least authority needed, and supervise the surrounding work system.

The three mode test

Ask these questions in order:

1. **Who initiates and acts?** If a person chooses the tool, supplies the material, and manually acts, treat it as employee tool use.
2. **Who remains the actor of record?** This is the person who makes or confirms the consequential action and remains accountable for it. If AI is embedded in a workflow but that person accepts or rejects the proposed work before action, treat it as a copilot.
3. **Can the system act across steps or tools?** If it can plan, call tools, change a system, communicate externally, spend, or transact, treat it as a semi-autonomous agent even if the product label says assistant.

Actual access and behavior control the classification. Marketing language does not.

Four axes, four questions

Continuing readers should not treat the series as one maturity ladder.

Axis	Question it answers	It does not establish
Book One adoption or reliance level	How deeply does the company use and depend on AI?	The work mode or permission to operate
Book Two work mode	Who initiates, proposes, decides, and acts?	Proposal or operating status
Book Two proposal or operating status	May this exact use be tested or operated now, and under what conditions?	Formal release or assurance
Book Three formal release status	Has the defined system met the prerequisites for a consequential release decision?	Independent assurance unless that work was separately performed

There is no one-to-one conversion among these axes. Material reliance does not automatically mean agent behavior. Agent behavior does not mean approved operation. Approved operation does not mean formal release or independent assurance.

BOOK 2 / OPERATE

Three work modes require three control postures

Classify the work by actual access and behavior, not the product label.

CONCEPTUAL COMPARISON. NOT A MATURITY LADDER.

INITIATOR / ACTOR	ACCESS / ACTION	CONTROL POSTURE
1 EMPLOYEE TOOL USE PERSON-OPERATED		
Person initiates, interprets, and acts. Employee remains the actor.	Manual copy or upload. No direct authority over company systems.	Approved-use rules Data handling Disclosure route Basic verification Tool inventory
2 COPILOT HUMAN-DIRECTED		
Person directs work and accepts, edits, or rejects the proposal. Person is actor of record.	Scoped workflow context. No consequential action before human decision.	Workflow owner Grounded sources Evaluation and review Version record Outcome measures
3 SEMI-AUTONOMOUS AGENT BOUNDED AUTHORITY		
Person sets the goal, grants limits, and supervises exceptions. Accountability stays named.	Distinct identity and delegated access. May act across tools, change, communicate, or transact within limits.	Agent charter Least privilege Action and confirmation limits Logs and monitoring Stop, revocation, and expiry
AUTHORITY CHANGES BY MODE More authority is a new, reversible decision. Autonomy is not the destination.		

Figure B2.2. Employee tool use is person-operated, a copilot is human-directed, and an agent receives bounded operating authority.

What this figure shows: The left box is employee tool use: an employee pastes routine message text into a public assistant, reviews the rewrite, and sends the final message manually. The middle box is a

copilot: an embedded feature retrieves approved guidance and proposes a request category, and an employee confirms or changes it before routing, remaining the actor of record. The right box is a semi-autonomous agent, shown by the employee-built automation that reads a shared mailbox, drafts supplier replies, and updates a record through a shared credential; the boxes differ by who initiates, who remains actor of record, and whether the system can act across steps or tools.

One proportionality rule

Use the least process that can still support the decision. Lower-consequence, reversible work may need a named owner, a clear information boundary, sampled verification, a correction path, and a dated review. Meaningful system access or moderate consequence adds representative exceptions, current evidence, tested fallback, and explicit operating conditions. External commitment, restricted information, financial authority, rights or safety impact, or hard-to-reverse action keeps decisive authority with an authorized person until identity, permission, evaluation, containment, recovery, and specialist conditions are established.

Each chapter applies this rule to a different decision. Quantities in the 30-day actions are planning prompts, not evidence-based thresholds. Choose enough coverage to represent the work, its variation, and its consequence.

The six decisions

This book asks the executive team to close six decisions:

1. **Visibility:** Where is AI already touching work, information, people, and company systems?
2. **Work design:** Which tasks should be removed, simplified, assisted, redesigned, delegated, or kept human-led?
3. **Relationship:** Should each use remain employee tool use, become a managed copilot, or receive bounded semi-autonomous agent authority?

4. **Operating design:** Who owns the outcome, handles exceptions, maintains knowledge, and preserves critical skills?
5. **Authority:** What may the system read, propose, change, communicate, or spend, and how is that authority revoked?
6. **Supervision:** What evidence, incidents, changes, expiries, and stop decisions will management review?

What to import, and what not to repeat

Keep the AI Investment Receipt as the source record for the approved use. Link its current conditions into operation without rebuilding the business case. Use the Digital Workforce Control File to show what the work system is doing now.

A useful operating record answers questions such as:

- What task is being performed, and in which mode?
- Who is accountable for the outcome?
- Which information and systems are accessible?
- What actions are prohibited?
- Where must a person decide?
- What happens when a case is uncertain or wrong?
- What evidence remains current?
- What triggers narrowing, rollback, retirement, or a new approval?

Your first 30 days

Do not announce an autonomous enterprise. Establish control of the facts.

1. Select a small, representative set of functions where AI use is likely to be material and the team can examine real work properly.
2. Invite disclosure under a published, nonpunitive discovery statement.

3. Interview actual users and map the work they perform, not only the tools they open.
4. Reconcile voluntary disclosure with procurement, vendor, and authorized access signals.
5. Classify each material use as employee tool use, copilot, or semi-autonomous agent.
6. Identify any use whose access or behavior exceeds its label.

30-DAY OPERATING CLOSE

7. Map a manageable candidate set against real baselines and consequences. Five tasks may be enough for a first pass, but the limit should reflect operating variety and consequence.
8. Freeze increases in authority until an owner documents evidence, limits, rollback, and review.

By day 30, management should be able to answer:
Where is AI already acting in our work, what relationship have we chosen, and what would make us change that choice?

DIGITAL WORKFORCE CONTROL FILE

Twelve sections support four operating decisions

One versioned operating record, not 12 separate forms.

CONCEPTUAL FILE ARCHITECTURE. ONE CONTROL FILE.



The 12 sections support four operating decisions while remaining one versioned Digital Workforce Control File.