By Frank Kienle · Version 0.1 · 2026-05-10 · Educational checklist

# Agent Readiness Checklist for Operations Professionals

Before you automate a workflow with AI agents, check whether the process is ready.

By Frank Kienle

Version: 0.1

---

## 0. How to use this checklist

This checklist is for operations, supply-chain, analytics, and process professionals who are being asked a simple but risky question:

> *Can we use AI agents for this workflow?*

Use it before investing time in tools, demos, dashboards, prompt chains, or automation prototypes.

The goal is not to prove that an AI agent should be used. The goal is to find out whether the workflow is ready enough for a small controlled pilot — or whether the real work is still process clarification, data cleanup, governance, and ownership.

Use it to:

- evaluate one candidate use case in 20-30 minutes,
- separate useful agent opportunities from vague AI ideas,
- identify missing process, data, governance, or review conditions,
- prepare a more serious discussion with business, analytics, IT, data, and compliance colleagues,
- decide whether the next step is a pilot, a process/data readiness task, or a clear stop.

What this package is not:

- not legal, financial, compliance, or employer-specific advice,
- not a guarantee of productivity, cost savings, or implementation success,
- not a consulting engagement,
- not a replacement for internal policy, data governance, security review, works council review, or expert judgment,
- not based on employer-confidential, customer-confidential, supplier-confidential, or internally identifiable information.

Recommended route:

```
Pick one workflow
  -> describe the current process
  -> score readiness across seven dimensions
  -> check red flags
  -> choose a next action
  -> document what must stay human-owned
```

---

# 1. The business problem

AI agents are often discussed as if the main question were technical:

- Which model?
- Which framework?
- Which tool?
- Which automation platform?

That is usually the wrong starting point.

In operations, the hard part is not the model. The hard part is knowing which workflows are stable enough, bounded enough, and reviewable enough to let an agent support them without creating operational theatre.

A chatbot can generate an answer. An agent workflow can trigger actions, call tools, collect information, route tasks, update records, and influence decisions. That makes the readiness question more serious.

If the workflow is unclear, the data is unreliable, the owner is missing, or the failure mode is invisible, an AI agent will not fix the problem. It will automate confusion faster.

---

# 2. Why it matters

Agent readiness matters because bad automation does not only waste technical effort. It can create operational risk.

1. Practical consequence: teams build demos that look useful but cannot survive real process exceptions.
2. Human consequence: process owners, planners, analysts, and managers still have to clean up the mistakes.
3. Data/process consequence: missing definitions, unclear handovers, and inconsistent records become harder to see once they are hidden behind an agent interface.
4. Decision consequence: teams may trust an automated workflow before they understand the boundary conditions.

The uncomfortable truth: many workflows are not agent-ready yet. That is not a failure. It is a useful finding.

If the checklist shows weak readiness, the next step is not another tool trial. The next step is to clarify the process, data, responsibility, and review point.

---

# 3. The Agent Readiness Framework

Use seven dimensions. A workflow does not need to be perfect, but weak scores in workflow boundary, data, human review, governance, or failure visibility are serious warning signs.

## 3.1 Overview

| Dimension | Core question | Output |
|---|---|---|
| 1. Business problem | What practical problem should the agent help with? | One clear use-case statement |
| 2. Workflow boundary | Where does the workflow start, stop, and hand over? | Process boundary and trigger |

| | | |
|---|---|---|
| 3. Data readiness | What data, documents, systems, and definitions are needed? | Data/source map |
| 4. Decision and action scope | What may the agent suggest, draft, route, or execute? | Allowed vs forbidden actions |
| 5. Human review | Where must a person inspect, approve, or override? | Review point and accountable role |
| 6. Failure visibility | How will wrong, incomplete, late, or risky outputs be detected? | Failure modes and fallback |
| 7. Governance and confidentiality | Which policies, permissions, and information boundaries apply? | Governance/confidentiality notes |

## 3.2 Dimension details

### 1. Business problem

Bad starting point:

> *We need an AI agent for operations.*

Useful starting point:

> *We want to reduce manual coordination around repeated supplier-delay follow-ups by preparing draft status summaries for human review.*

Reader action:

- Write the use case in one sentence.
- Name the person or team who currently owns the problem.
- State what should improve: speed, consistency, quality of handover, fewer manual checks, better visibility, or better prioritization.

If the business problem cannot be stated without buzzwords, stop and clarify it first.

### 2. Workflow boundary

Agents work better when the workflow is repeated and bounded.

Reader action:

- Define the trigger: what starts the workflow?
- Define the end: what counts as completed?
- List handovers: which teams, systems, or roles are involved?
- Mark exceptions: where does the normal process break?

A workflow with no clear start, no clear end, and many undocumented exceptions is not a good first agent pilot.

### 3. Data readiness

An agent cannot reliably coordinate work if the relevant data is scattered, undefined, stale, or permission-restricted.

Reader action:

- List the required data sources.
- Mark which sources are structured, semi-structured, or free text.

- Check whether terms are defined consistently.
- Check whether the data can legally and practically be used in the agent workflow.

Data readiness does not mean perfect data. It means the minimum information needed for a controlled workflow is available, understandable, and reviewable.

### 4. Decision and action scope

Separate advice, drafts, routing, and execution.

Reader action:

- Decide what the agent may do:
- retrieve information,
- summarize status,
- draft messages,
- recommend priorities,
- route tasks,
- update records,
- trigger actions.
- Decide what the agent must not do.

A sensible first pilot often starts with retrieval, summarization, drafting, or prioritization. Direct execution without review needs a much stronger governance case.

### 5. Human review

Human-in-the-loop is not a decoration. It must be placed where it actually reduces risk.

Reader action:

- Name the accountable role.
- Define what the reviewer checks.
- Define when the reviewer can override or reject.
- Define what happens if no reviewer is available.

If nobody owns the review point, the workflow is not ready.

### 6. Failure visibility

The dangerous failures are not always obvious errors. Sometimes the agent sounds confident while missing context.

Reader action:

- List possible wrong outputs.
- List possible missing inputs.
- List possible timing failures.
- List possible confidentiality or permission failures.
- Define a fallback process.

A useful question:

> *If the agent is wrong, who notices it, how quickly, and what do they do?*

If there is no answer, do not automate yet.

### 7. Governance and confidentiality

Operations workflows often touch supplier, customer, employee, cost, forecast, performance, and commercial information. Treat that seriously.

Reader action:

- Identify confidential, personal, customer, supplier, or commercially sensitive information.
- Check permissions and access boundaries.
- Check whether outputs could expose internal logic or sensitive relationships.
- Check whether the workflow requires compliance, security, legal, data protection, works council, or management review.

Do not use internal employer data, screenshots, dashboards, reports, customer details, supplier details, or identifiable process examples in public demos or knowledge assets.

---

## 4. The checklist

Use this section as the one-page working checklist.

### 4.1 Quick readiness checklist

For one candidate workflow, answer each item with Yes / Partly / No.

| # | Check | Yes | Partly | No | Notes |
|---|---|---|---|---|---|
| 1 | The business problem is specific and owned by a real role/team. | ☐ | ☐ | ☐ | |
| 2 | The workflow is repeated often enough to justify automation support. | ☐ | ☐ | ☐ | |
| 3 | The start, end, inputs, outputs, and handovers are clear. | ☐ | ☐ | ☐ | |
| 4 | Required data/documents/systems are known and accessible for the intended use. | ☐ | ☐ | ☐ | |
| 5 | Key definitions are stable enough for a first controlled pilot. | ☐ | ☐ | ☐ | |
| 6 | The allowed agent actions are explicitly limited. | ☐ | ☐ | ☐ | |
| 7 | Forbidden actions are documented. | ☐ | ☐ | ☐ | |
| 8 | A human review or approval point is defined. | ☐ | ☐ | ☐ | |
| 9 | The accountable business/process owner is named. | ☐ | ☐ | ☐ | |
| 10 | Likely failure modes are visible and testable. | ☐ | ☐ | ☐ | |
| 11 | A fallback/manual process exists if the agent is wrong or unavailable. | ☐ | ☐ | ☐ | |
| 12 | Confidentiality, data protection, permission, and policy boundaries are understood. | ☐ | ☐ | ☐ | |
| 13 | The pilot can be tested safely with low operational impact. | ☐ | ☐ | ☐ | |
| 14 | Success is defined in operational terms, not generic AI excitement. | ☐ | ☐ | ☐ | |

### 4.2 Simple scoring option

Score each dimension:

- 0 = not ready
- 1 = partly ready / needs clarification
- 2 = ready enough for a controlled first evaluation

| Dimension | Score 0-2 | Evidence / notes |
|---|---|---|
| Business problem | | |
| Workflow boundary | | |
| Data readiness | | |
| Decision and action scope | | |
| Human review | | |
| Failure visibility | | |
| Governance and confidentiality | | |
| Total out of 14 | | |

Suggested interpretation:

- 0-5: not agent-ready. Start with process/data clarification.
- 6-9: discussion-ready. Useful candidate, but too many open issues for implementation.
- 10-12: pilot-candidate. Prepare a small controlled prototype with review.
- 13-14: strong readiness signal. Still validate governance, data access, and review before production.

Do not create fake precision. A low score is valuable if it prevents a bad pilot.

### 4.3 Red flags

Hard red flags — stop before building:

- no accountable process owner,
- no clear human review point,
- unclear permission to use the required data,
- confidential or personal data involved without reviewed controls,
- the agent would directly execute operationally relevant actions without approval,
- the use case depends on internal context that cannot be safely exposed to the tool,
- failure would affect customers, suppliers, employees, compliance, safety, finance, or reputation before anyone notices.

Soft red flags — clarify before piloting:

- the workflow has many exceptions and no exception taxonomy,
- data fields mean different things in different systems,
- the current process is already disputed between teams,
- success is described only as “save time” without operational definition,
- people expect the agent to solve a management or ownership problem,
- the pilot is being pushed because AI budget exists, not because the use case is ready.

Frank-style rule:

> *If the process boundary is unclear, AI automation becomes theatre.*

---

## 5. Use-case prioritization

Use this matrix after the readiness score. A workflow can be ready but not worth doing first.

### 5.1 Prioritization questions

| Question | Better first pilot | Risky first pilot |
|---|---|---|
| Frequency | Repeated weekly/daily | Rare, strategic, one-off |
| Risk | Low to moderate, reviewable | High operational/compliance/customer impact |
| Data | Known sources, limited permissions | Sensitive, fragmented, unclear permissions |
| Action scope | Draft, summarize, recommend, route | Execute, approve, negotiate, commit |
| Review | Clear human owner | Nobody owns final check |
| Value | Reduces boring coordination or improves visibility | Vague transformation promise |
| Testability | Can be tested on synthetic/historical/public examples | Needs live sensitive data from day one |

### 5.2 First-pilot sweet spot

Good first candidates often look boring:

- collecting status inputs from known sources,
- preparing draft summaries for review,
- triaging repeated requests,
- checking whether required documents are present,
- routing cases to the right owner,
- comparing a request against a checklist,
- creating a draft explanation that a human approves.

That is not a weakness. Boring coordination is often where agent workflows become useful first.

Avoid starting with:

- supplier or customer negotiation,
- workforce decisions,
- compliance approvals,
- financial commitments,
- production-critical control decisions,
- anything where a wrong action creates immediate external damage.

---

## 6. Blank working template

Copy this block for each candidate workflow.

```
 Candidate workflow:
Business problem in one sentence:
Process owner / accountable role:
Current process summary:
Trigger / start point:
End point / completion condition:
Main inputs:
Main outputs:
Systems / data sources:
Allowed agent actions:
Forbidden agent actions:
Human review point:
Reviewer role:
Likely failure modes:
```

```
Fallback process:
Confidentiality / permission notes:
Governance / compliance review needed:
Readiness score out of 14:
Decision: stop / clarify / discussion-ready / controlled pilot candidate
Next action:
```

---

## 7. Practical example — repeated status coordination

This example is generic and synthetic. It does not describe an employer, customer, supplier, or internal process.

Context:

A planning team receives repeated status updates from different internal sources. The problem is not strategic decision-making. The problem is boring coordination: collecting known inputs, summarizing gaps, and preparing a draft update for a human planner.

Applying the framework:

| Dimension | Example answer |
|---|---|
| Business problem | Planners spend time preparing repeated status summaries instead of reviewing exceptions. |
| Workflow boundary | Trigger: scheduled status review. End: human-approved summary. |
| Data readiness | Known input documents and status fields exist, but definitions need checking. |
| Decision/action scope | Agent may retrieve inputs and draft a summary. It may not approve, commit, or contact external parties. |
| Human review | Planner reviews and edits before anything is shared. |
| Failure visibility | Missing input, stale source, inconsistent status, or wrong summary must be flagged. |
| Governance/confidentiality | Use only approved internal tools/data. Public demos must use synthetic data. |

What this shows:

- A narrow coordination workflow can be a sensible first candidate.
- The agent should draft and flag gaps, not take over ownership.
- Review and fallback matter more than model enthusiasm.

---

## 8. Companion links

Use the trust-loop route from homepage to GitHub to registration.

- Topic page: AI Agents in Operations — https://kienlef.github.io/ai-agents-in-operations/
- GitHub proof artifact: operations_use_case_selection — https://github.com/kienlef/operations_use_case_selection
- YouTube channel until a specific playlist is curated: https://www.youtube.com/@frankkienle7312
- Free checklist registration page: https://kienlef.github.io/register/

Recommended path:

```
Read the AI Agents in Operations topic page
  -> inspect the operations_use_case_selection GitHub material
```

```
  -> use this checklist on one candidate workflow
  -> decide whether the next step is stop, clarify, or controlled pilot
```

## 9. Next step / CTA

If this checklist was useful, continue with the AI Agents in Operations topic page and registration route:

https://kienlef.github.io/ai-agents-in-operations/ https://kienlef.github.io/register/

Use the checklist before investing time in tools, dashboards, prompt chains, or AI-agent prototypes.

A future paid package candidate, `AI for Operations Playbook`, is planned only after the free checklist has been reviewed and the delivery path is tested. It should not be sold or linked as a product before Frank reviews the content, delivery path, and disclaimer.

## 10. Disclaimer

This knowledge package is for educational and professional learning purposes. It does not provide legal, financial, compliance, or employer-specific advice. Adapt the framework to your own context and validate it with your internal policies, data governance, security, compliance, works council, and expert review where relevant.

This is personal educational material by Frank Kienle. Views are personal. Examples are based on public, educational, historical, generic, or synthetic material unless stated otherwise. No employer-confidential, customer-confidential, supplier-confidential, or internally identifiable information is shared.