

FREE LESSON PREVIEW

Module 2 — Delegation: Briefing an Agent Like a Pro

From **AI Agent Command™** — the nine-module course on managing AI agents like infrastructure, by the team that runs them where a wrong answer can unlock a door.

2.1 Bad output is usually a bad briefing

When an agent disappoints, managers blame the model. In our experience running agents in production, **80% of bad output traces back to the delegation, not the model.** Vague goals produce vague work — at machine speed.

The one-liner "Write something for the website about our new service."	The SCOPE briefing "Write a 300-word landing-page section announcing our OSDP migration assessment service. Audience: facility managers who still run Wiegand. Tone: direct, technical-credible, no hype. Must include: the SIA position on Wiegand, our 4-phase method, one CTA to the contact form. Must not include: pricing, competitor names. Output as HTML matching our existing section style."
---	--

Same model. Different manager.

2.2 The Kryos SCOPE briefing method

Every delegation to an agent should cover five elements:

- S Situation**

The context the agent can't guess: who you are, who the audience is, what came before. Agents don't share your head. State the obvious.
- C Constraints**

Hard limits: what it must include, must not include, must not touch, budget/length limits, deadline behavior. Constraints are the part agents violate when unstated — and honor with remarkable reliability when stated.
- O Output**

The exact deliverable: format, length, file type, structure, where it goes. "A markdown file with these four sections" beats "a report."
- P Process**

For multi-step work: the order of operations, which tools to use (or avoid), when to check in, what to do when blocked. This is where you place your checkpoints.

E Examples
One good example is worth 500 words of description. Show a past deliverable you liked, or a sketch. Also show a *bad* example when the failure mode is predictable.

2.3 Sizing the task: the delegation ladder

Match task size to trust level — exactly as you would with a new hire.

Rung	Trust level	Task type	Checkpoint style
1	New agent / new domain	Draft-only; you finalize	Review everything
2	Proven on rung 1	Complete tasks, defined inputs	Review outputs
3	Proven on rung 2	Multi-step projects	Review at milestones
4	Proven + guardrailed	Autonomous runs (scheduled)	Review audit trail + exceptions

Never onboard an agent at rung 4. The teams that get burned are the ones who skip rungs because a demo looked good. Demos always look good.

2.4 The restatement trick

Before letting an agent run a big task, ask it to restate the assignment: *"Before you start, summarize what you're going to do, in what order, and what you will not do."* Ten seconds of reading catches the majority of misunderstandings before they become an afternoon of rework. This is the agent equivalent of "repeat the work order back to me."

CHECKLIST — THE 60-SECOND PRE-DELEGATION CHECK

- [] **Situation:** did I state who/what/why?
- [] **Constraints:** did I list the must / must-nots — especially the "obvious" ones?
- [] **Output:** could a stranger tell if the deliverable matches?
- [] **Process:** are checkpoints placed before anything irreversible?
- [] **Example:** did I show one?
- [] **Ladder:** is this task at this agent's proven rung?

What's in the full course

Core Track — 1 Foundations · 2 Delegation · 3 Oversight & Guardrails · 4 Orchestration · 5 Measurement · 6 Governance

Security Specialization — S1 Agents in Physical Security · S2 Deploying Field Copilots · S3 Life-Safety Guardrails & Audit

Every module ships with lesson content, a field exercise, a checklist you can use tomorrow, and a knowledge check — like the ones above, which are reproduced here in full from the real Module 2.

Join early access at kryossecuritygroup.com/academy — you'll get the launch date, the early-access discount, and the free one-page AI Usage Policy template.