

DAILY ENGLISH READING

9 September 2026 · 30-minute active reading pack for Davin

TODAY'S QUESTION

If AI agents run parts of a clinical trial, who should be allowed to do what?

WHY SELECTED

CDISC TechniCon is happening today, while a recent peer-reviewed ASTRA Perspective proposes a multi-agent Clinical Trial Center. Together they make one issue concrete: clinical automation needs task-by-task autonomy boundaries, not a blanket decision to automate everything.

Difficulty: C1 · multi-agent trials + CDISC + governance + SP autonomy

READING ORDER · 30 MINUTES

Min	Task	Focus
0-3	Preview	Coordinator, teams, shared knowledge, human supervision.
3-14	ASTRA	Task classes, orchestration, blinding, auditability, planned validation.
14-20	TechniCon	Protocol-to-CSR, evidence-backed review, traceability, full lifecycle.
20-25	FDA/EMA	Human-centric, risk-based, standards, context of use, lifecycle.
25-30	Output	Classify six SP activities by autonomy level.

SOURCES · OPEN ACCESS

Main · npj Digital Medicine · 24 Aug 2026

Towards a multi-agent Clinical Trial Center framework to support modern clinical trials

[Open article](#)

Support · CDISC · 9 Sep 2026

2026 CDISC TechniCon

[Open live program](#)

Support · FDA / EMA · Jan 2026

Guiding Principles of Good AI Practice in Drug Development

[Open principles](#)

BRIEF BACKGROUND

ASTRA proposes a central Coordinator Agent plus ten specialized functional teams spanning protocol, data, QA, biostatistics, safety committees, sites, monitoring, patients and laboratories.

Human-led tasks retain judgment and sign-off. System-assisted tasks prepare or reconcile material for review. Bounded agent-executed tasks are routine and rule-governed inside explicit scope, access, escalation and audit limits.

The Coordinator orchestrates dependencies and exceptions but delegates domain-specific execution instead of silently rewriting statistical or safety outputs.

The architecture uses structured communication, a version-controlled knowledge base, and explicit blinded/unblinded boundaries.

Today's CDISC TechniCon demonstrates standards-connected protocol-to-CSR pipelines, AI data review with evidence, protocol-to-ADaM traceability and governed SDTM/ADaM/TFL automation.

ASTRA remains a design framework, not a validated deployment. Retrospective and prospective evidence is still required before autonomy should expand.

KEY VOCABULARY

Term	中文	Meaning / use
multi-agent system	多智能体系统	A system in which multiple specialized agents coordinate through defined roles and interfaces.
Clinical Trial Center	临床试验中心	An operational structure coordinating functions such as protocol, data, quality, safety, statistics, monitoring, and laboratories.
Coordinator Agent	协调智能体	A supervisory agent that allocates work, tracks dependencies, monitors status, and handles exceptions.
human-led task	人工主导任务	Work that requires judgment, ethical discretion, regulatory interpretation, or formal sign-off.
system-assisted task	系统辅助任务	Work prepared or structured by automation but reviewed and decided by a human.
bounded agent-executed task	边界受控的智能体执行任务	Routine, rule-governed work performed autonomously within explicit scope, access, audit, and escalation limits.
task orchestration	任务编排	Coordinating dependent activities across teams, tools, and agents so work occurs in the correct order.
dependency tracking	依赖关系跟踪	Monitoring which tasks, data, approvals, or events must be completed before another task can proceed.
exception handling	异常处理	Detecting unusual states and routing them to the correct recovery or escalation path.
version-controlled knowledge base	版本控制知识库	A controlled repository of protocols, SAPs, schemas, rules, and other authoritative trial knowledge.
structured messaging protocol	结构化消息协议	A predefined message format that carries task identity, status, data pointers, escalation, and approval information.
role boundary	角色边界	A formal limit on what an agent or human is authorized to see, decide, or modify.
blinding boundary	盲态边界	A control that prevents access to information that could compromise trial blinding.
auditability	可审计性	The ability to reconstruct actions, inputs, outputs, versions, approvals, and decisions after the fact.
fallback / escalation path	回退 / 升级路径	A predefined route for handing a problem to another system state or qualified human when automation is insufficient.

USEFUL PHRASES

1. **assign autonomy according to task risk** - *Autonomy should be assigned according to task risk rather than technical possibility.*
2. **keep judgment and formal sign-off with accountable experts** - *The workflow should keep judgment and formal sign-off with accountable experts.*
3. **orchestrate dependencies without modifying domain outputs** - *The coordinator should orchestrate dependencies without modifying domain outputs.*
4. **operate within explicit scope and access boundaries** - *Each agent should operate within explicit scope and access boundaries.*
5. **use a version-controlled source of truth** - *The system should use a version-controlled source of truth.*
6. **make handoffs reproducible and auditable** - *Structured messages make handoffs reproducible and auditable.*
7. **separate blinded and unblinded workflows** - *The architecture must separate blinded and unblinded workflows.*
8. **escalate outputs that fail deterministic validation** - *Outputs that fail deterministic validation should be escalated.*
9. **treat standards as a connected backbone** - *CDISC standards can act as a connected backbone across the clinical-data lifecycle.*
10. **validate the workflow before expanding autonomy** - *The workflow should be validated before expanding autonomy.*

COMPREHENSION QUESTIONS

1. Why does ASTRA keep the Coordinator Agent at the orchestration level?
2. What distinguishes system-assisted work from bounded agent-executed work?
3. Why should blinding be enforced through access control rather than prompts?
4. What does today's TechniCon suggest about the future role of CDISC standards?
5. Why is an attractive architecture still insufficient evidence for production deployment?

RETELLING PROMPTS

30 sec: Three task classes with one example each.

45 sec: Objective -> Coordinator -> specialized team -> evidence -> approval/bounded execution -> audit.

60 sec: Explain how an SP team can automate aggressively without automating final statistical accountability.

5-MINUTE SPEAKING / OUTPUT TASK

Minute 1: Choose six SP activities.

Minutes 2-3: Classify each as human-led, system-assisted or bounded agent-executed.

Minute 4: Add source, validator, permissions, escalation and audit controls.

Minute 5: Defend which autonomy boundary you would refuse to cross.

ONE SENTENCE TO KEEP

The safest multi-agent clinical workflow does not maximize autonomy; it assigns autonomy task by task, keeps authoritative standards and permissions outside the model, and makes every handoff, escalation, and approval reconstructable.