



THE WHOLE SYSTEM

Artificial Organizations

15 visual study sheets


Companion to the

Study Monograph

v1.2 / September 2026

Not a crowd of agents.
An institution that can remember, authorize, and answer for its work.

A condensed map of the book's concepts for returning readers: recover a distinction, choose a mechanism, or locate a question worth testing. These sheets summarize the argument; they do not replace its evidence, assumptions, or detailed examples.



Actor	Institution	Accountability
A runtime participant that observes and acts. Its presence is not an organizational contract.	Durable positions, legitimate acts, obligations and ways to settle what counts as done.	A traceable relation among responsibility, authority, evidence, review and consequences.

🕒 **The central test** Connect every promised outcome to an owner, an authorized action, a check, and an observation of its effects. A convincing narrative supplies none of these automatically.

■ Work & constructive action ■ Authority & caution ■ Evidence & context ■ Critical failure



BEFORE ADDING ANOTHER AGENT

Make organization earn its keep.

Specialization saves expertise.
Every boundary also charges a price.

POTENTIAL GAINS

Specialization
Parallel work
Exception routing

Use scarce expertise only where needed. Preserve stable intermediate results. Put decisions near the knowledge that supports them.

COMPARE

*same
outcome*

same constraints

ORGANIZATIONAL COSTS

Handoffs
Attention
Information loss

Pay for assignment, translation, waiting, monitoring, reconciliation and recovery. Internal instructions are not free.

Decide under limits

Bounded rationality

Evaluate choice against available information, capability, time and resources, not an imaginary omniscient optimizer.

Satisficing

Stop at an explicit acceptability threshold. Reaching it does not prove global optimality or that the threshold is adequate.

Uncertainty absorption

Recipients see an inference instead of its observations. Efficient, but tentative judgment can harden into apparent fact.

Loss-accounted compression

The book's heuristic: preserve paths to evidence, omissions and dissent while reducing reading effort.

Shape the dependency

Near-decomposability

Interactions are stronger within subsystems than between them. Weaker coupling is not independence.

Information need / capacity

Reduce coordination demand through boundaries or slack; increase capacity through systems and lateral relations.

Knowledge hierarchies

Keep routine problems with general workers; send rare exceptions to experts. Task distributions and costs decide the benefit.

Firm boundary

Compare the next activity inside, outside, or in another firm. Include search, bargaining, adaptation and internal administration.

Pooled		Sequential		Reciprocal
Shared resource base	/	One output feeds another	/	Work changes other inputs
<i>Standards</i>		<i>Plans & interfaces</i>		<i>Mutual adjustment</i>

What is standardized? Process, output and skill give different assurances. Supervision and mutual adjustment solve other coordination problems. A title does not certify skill; compliance does not establish an outcome.

Who gets through? Sequential veto screens reject more bad *and good* proposals; independent acceptance channels admit more of both under the compared assumptions. More reviewers is not neutral.



Choose a baseline

Compare with one capable worker and a deterministic workflow. Intermediaries bring commitments and influence, not only capacity: a factory planning model and TVA's institutional relationships answer different questions.

Foundations: Simon; March & Simon; Galbraith; Thompson; Mintzberg; Garicano; Sah & Stiglitz; Coase; Williamson. Models and historical analyses have distinct assumptions.

RESPONSIBILITY SURVIVES REPLACEMENT

The worker changes. The duty remains.

Model the position, its occupant, and its obligations separately.



Four questions before dispatch

1. **Scope.** In which group does this position apply? A role name is not a global credential.
2. **Occupancy.** Who enacted it when the act occurred? Preserve intervals and actor identity.
3. **Compatibility.** May the proposer also be the reviewer? Separate real constraints from bureaucratic habit.
4. **Continuity.** What pending commitments, credentials and attempts remain when the worker changes?

Landmarks, not one compulsory path

A landmark specifies a required state: for example, **recovery evidence accepted**. Several procedures may reach it.

That flexibility changes the route, not the duty, permissible means, or acceptance standard. A new plan must not quietly redefine what counts as success.

⚠ **Replacement is a transition, not a rename.** Test it during unfinished work, not only while idle.

APPROACH	KEEP THIS DISTINCTION	INSPECT
AGR	Agent / group / role	Membership and enactment records
MOISE+	Structure / goals and missions / norms	Role, mission and obligation declarations
OperA	Organizational requirements / participation / interaction	Landmarks and concrete agreements
OMNI	Organizational, normative, ontological <i>concerns</i> versus abstract, concrete, implementation <i>levels</i>	A detailed schema may still leave a policy question unanswered

⊙ **Minimal implementation**

A typed role table may be enough. Use a formal organizational language when its distinctions repay their cost. Stable names without guarded transitions are decorative structure.

Foundations: Ferber & Gutknecht (AGR); Hübner and colleagues (MOISE+); Dignum (OperA); Dignum and colleagues (OMNI). These are related, not equivalent specifications.



CHOOSE THE DEPENDENCY BEFORE THE HEADCOUNT

Six coordination patterns.

Topology describes dependencies, not a ranking of intelligence.

One notation: green circles = S1-S3 specialists; rectangles = coordinating functions or state; blue arrows = information; amber arrows = assignment or selection.

Pipeline

01




Fixed precedence

What must precede what? Each stage consumes an earlier artifact. Define input meaning, preconditions and stage acceptance.

Risk: errors and stale assumptions travel downstream.

Blackboard

03



Selection policy chooses next contribution

What contribution is useful now? Specialists respond to evolving partial solutions.

State is not the scheduler; consistency still costs work.

Peer deliberation

05




Reasons and objections; both directions

Which reasons survive scrutiny? Compare evidence, objections and values; state who closes the decision.

Conversation and agreement do not prove independence.

Supervision

02



Assignments down / summaries up

Who allocates and escalates? Route assignments and exceptions through a bounded authority relation.

Risk: attention bottlenecks and lossy summaries.

Contract Net

04




Requests / proposals; ring marks award

Who should take the work? Allocate among candidates with relevant capability, availability or cost.

An award is not completion or evidence of satisfaction.

Orchestrator / workers

06



How do partial results integrate? Define parallel dependencies, merge criteria and recovery after interruption.

Replanning does not transfer authority automatically.

Pattern ≠ substrate

Tuple space: discover by content. **Actors:** address participants. **Channels:** constrain communication. **Durable workflow:** retain progress. **MCP / A2A:** cross interface boundaries. **Transactions / rights:** protect shared state.

No one-to-one pairing: a pipeline can use actors and durable execution; tuple storage can support a blackboard but does not supply its selection policy.

One handoff

S1 produces, S2 checks, S3 integrates. Changing transport can preserve this precedence; switching to opportunistic selection changes it.

Source families: Smith (1980); Erman et al. (1980); Horling & Lesser (2004); Linda, actor and process models. S1-S3 retain identity; hubs denote coordinating functions, not necessarily extra agents.

THE SMALL OPERATION AND THE LARGER PROMISE

Claiming is not completion.

Coordinate the invariant, not every observation.

LINDA / RINDA

Associative shared space


Match content, not a named recipient.

out / write: publish.
rd / read: observe and retain.
in / take: consume one match.

```
space = Rinda::TupleSpace.new
space.write(["take-test", 1])
space.take(["take-test", nil])
```

Ruby; first require rinda/tuplespace. Returns ["take-test", 1]; nil matches any value.

One take excludes a second taker of that tuple.
It does not prevent duplicate publication, resurrect a stopped consumer, or verify the result.



Different coordination contracts

CSP / HOARE 1978	PI-CALCULUS / NAME PASSING	REO / SYNC CHANNEL	
<pre>X :: *[c:character; west?c -> east!c]</pre> Historical COPY (§3.1): * repeats; ? receives, ! sends. Each handoff waits for its partner. Not modern compiler input.	<pre>$\bar{a}\langle b \rangle.P \mid a(x).Q$ $\rightarrow P \mid Q\{b/x\}$</pre> Send name <i>b</i> on <i>a</i>; receive as <i>x</i>. means parallel; substitute <i>b</i> for free <i>x</i>, avoiding capture. Editorial synchronous reduction; mobility changes connectivity.	<pre>write(a,v) and take(b,v) complete together</pre> Same value <i>v</i>, synchronized ends: a waiting sink blocks completion. Explanatory channel contract, not configuration syntax; a buffer changes the contract.	
Duplicate publication Two tuples can admit two consumers. Deduplicate by operation identity where the effect requires it.	Crash after claim Consumed work can remain unfinished. Define durable attempts, recovery and reassignment.	Lost acknowledgment No reply is not no effect. Inspect or reconcile external state before retrying.	Obsolete owner Timeouts do not stop workers. The effect boundary must reject stale ownership or authority.

DECIDE

Can concurrent operations violate the requirement?

No: let independent evidence collection proceed. **Yes:** use agreement, admission control or appropriately allocated rights at the smallest necessary boundary.

CRDTs require compatible merge semantics. Choose locks, transactions or workflows by invariant, failure model and recovery; exclusive ownership and scarce resources may still need coordination. A lock service is not an outcome verifier.

Test interruption

Stop a worker before and after the external effect. Inspect work identity, attempt state, authority and observed effect separately. Prefer an established substrate whose failure contract is understood.

PERMISSION TO ANALYZE IS NOT PERMISSION TO ACT

Delegate rights. Not just tasks.

Authority belongs at the boundary where the effect becomes real.

01

 **Acquire**

Which observations and records may the worker access?

02

 **Analyze**

Which transformations and interpretations may it perform?

03

 **Select**

Who chooses among consequential alternatives?

04

 **Implement**

Which external effects may it cause, under what guard?

Different stages may have different human/machine allocations. A single autonomy score hides the point where consequences change.

 GATE A ACCESS Can the actor reach or invoke the interface? Technical capability	 GATE B POWER Can this act create a recognized institutional effect? Constitutive rules	 GATE C PERMISSION Is this act allowed in this context? Applicable policy
--	---	---


 **Not three synonyms.** An actor may reach a system, create an effective act, and still violate a rule. Credentials are not a complete authorization decision.

A delegation envelope

- + Objective + acceptance criterion
- + Debtor + accountable recipient
- + Actions, data and resource scope
- + Policy version + capability binding
- + Deadline + authority expiry
- + Evidence + review triggers
- + Fallback + no-response default
- + Revocation + unfinished-work handling

Expiry is not physical stopping

A paused or disconnected actor can resume after its authority expires. A timestamp in the record is not enough.



A check followed by an unguarded remote call leaves a race. Use the receiver's enforceable contract; a generation number in a log is not fencing.

 **Revocation test**

Pause an authorized attempt. Revoke or supersede its authority. Resume it. The important observation is whether the effect-owning boundary refuses the obsolete act, not whether a dashboard says "expired."



ESCALATION IS A SCARCE-RESOURCE DECISION

A quiet dashboard can still be wrong.

8.3%

of alerts are real in this illustrative rare-event example

90% sensitivity.
90% specificity.
1% prevalence.

Headline detector accuracy does not describe the operator's queue.

Positive predictive value depends on the base rate as well as the detector.



Distributed observation, finite human attention. NASA, 1969, KSC-69P-3253. Not the source of this alert example. Image credit.

1,000 hypothetical cases

10 actual events

90% sensitivity

9 true alerts

1 missed event

990 non-events

90% specificity

99 false alerts

891 correct silences

Queue: 9 true + 99 false = 108 alerts. PPV = 9 / 108 ≈ 8.3%. Counts are illustrative arithmetic, not deployment results.

Interrupt for preventable loss

Notify when the expected loss avoided by informing the recipient exceeds the communication and attention cost.

Consequence

What harm can still be prevented?

Recipient knowledge

Does this person already know?

Opportunity

Can they act in time, with useful authority?

Low model confidence alone does not answer these questions.

Neither does a fixed alert threshold.

Two forms of reliance

Compliance

Responding when automation signals a problem. Measure unnecessary action on false alerts.

Reliance

Depending on silence. Measure harmful misses when the system does not alert.

Supervisory capacity

Depends on neglect tolerance and intervention time. Bursts, context switching and simultaneous exceptions defeat simple averages.

Track the queue, not just the classifier

Prevalence PPV Response time Sentinel-found misses Downstream harm



Before scaling

Measure review time and the loss of ignored or late interventions. Adding parallel workers does not add parallel human attention. “No alerts” is an observation, not evidence that checks ran successfully.

Tambe (1997), STEAM communication; Bliss et al. (1995); Meyer (2001, 2004); Lee & See (2004); Crandall et al. (2005). Human-factors findings need contextual transfer to LLM systems.

RECOGNITION DOES NOT IMPLY APPROVAL

What counts as an act?

What is allowed is a separate question.

Observed event

Actor submits a signed record

→ constitutive question

Does this count as an official act?

Check actor capacity, context, required form and the applicable institutional rule.

→ regulative question

Was the act permitted, required or prohibited?


Check scope, conditions, exceptions and the rule's authority. An effective act may still be forbidden.

Regimentation

Prevent a prohibited transition at a boundary the system actually controls. A guard cannot retroactively undo an external effect.

Enforcement

Detect and respond to a violation. Record, notify, sanction or repair where legitimate. A violation log is not prevention.



TRANSITION	MEANING TO RETAIN	QUESTION AT THE BOUNDARY
Create / activate	An undertaking is established or its condition becomes true.	Which parties, terms and authority made it valid?
Satisfy / discharge	The required performance is settled under its criterion.	Which evidence establishes satisfaction?
Cancel / release	Stopping and a creditor's release are different acts.	Who may do it? Do notification or compensation duties remain?
Violate / repair	A requirement was not met; a response may now be owed.	What remedy is authorized, feasible and independently checked?

Conflicting rules need more than recency

Priority, specificity, time and authority can matter, but "newest wins" is not valid when a junior actor cannot supersede the governing rule. Make disagreements inspectable; do not hide them inside a prompt.



Four independent checks

Technical access → institutional power → permission → evidence of the result. Author the relationship among them. None is a synonym for another.



DO NOT COLLAPSE THE RECORDS

Finished. Accepted. *Effective?*

A stopped run, a reviewed submission, and a successful outcome are three different observations.

01 / Commitment	02 / Execution attempt	03 / Submission	04 / Acceptance	05 / Observed outcome
The duty is accepted	Work is attempted	A result is offered	A defined check is met	Consequences are measured
Record debtor, creditor, scope, deadline and criterion.	Record actor, authority, input, progress and effects.	Preserve the artifact and its evidence packet.	A designated reviewer applies the criterion.	Over a stated horizon, effects may confirm success, reveal harm or remain unknown.

These are distinct records, not an unconditional progression. Review can require rework; an accepted submission can still have adverse downstream effects.

Silence has states

Checked, fresh, within limits

Positive health evidence

Checked, uncertain or degraded

Result exists; adequacy is unresolved

Not checked / check failed

Absence of evidence is visible

Evidence stale / not applicable

Time and scope constrain interpretation

What is still in flight?

An empty queue can coexist with active producers, consumed but unfinished tasks, delayed messages, or lost work.

Completion detection must account for producers, attempts, expected results and recovery, not only the absence of a matching tuple.

⚠ A good rationale is a witness statement. It does not replace an action trace or independent acceptance evidence.

GITLAB / PUBLIC INCIDENT

Service restored and **all changes recovered** are different claims. Recovery ownership, artifact dependencies, service checks and observed data loss belong in different records.



Ask “which done?”

Before closing work, name the state being asserted, the criterion, the reviewer and the evidence. State what remains unobserved. Never let a status label answer all five questions.

Book synthesis, Chapters 9 and 14; GitLab's 2017 postmortem; W3C PROV; source and uncertainty traditions. The five-record separation is an operational design synthesis, not a universal standard.



A SOURCE LINK IS THE BEGINNING OF REVIEW

Preserve the path from evidence to claim.

LINEAGE

Where did it come from?

PROV: entities, activities, agents, derivations and attribution. Recorded history does not prove truth or permission.

INFERENCE


Why does it follow?

Toulmin: data support a claim through a warrant; backing supports the warrant; qualifiers and rebuttals bound it.

JUDGMENT

How much support, and for what?

Represent uncertainty, source quality, task-specific reputation and competing explanations separately.



Ignorance ≠ disbelief

<i>b</i> belief	<i>d</i> disbelief	<i>u</i> uncertainty
--------------------	-----------------------	-------------------------

$$b + d + u = 1$$

projected probability = $b + a \cdot u$

In a binomial subjective opinion, **a** is the base rate. Complete uncertainty ($u = 1$) differs from complete disbelief ($d = 1$), even though both have $b = 0$.

Practical equivalent: “not checked” and “checked and invalid” must remain distinguishable even without numerical calculus.

Reputation has a task

Accurate extraction does not establish sound legal judgment, honest reporting or permission to act.

- + Which activity and conditions?
- + Whose experience or social evidence?
- + How much, how recent, how independent?
- + What would reverse the judgment?

A global trust score can erase the very distinction needed for delegation.

Compare hypotheses with discriminating evidence

	H1	H2	H3
Observation A	fits	fits	fits
Observation B	fits	conflicts	unknown

This schematic illustrates ACH reasoning, not a scored dataset. A fits-everything observation discriminates little. Assess quality and dependence before counting inconsistencies.



Review packet

Keep the claim, observations, source version, inference, uncertainty, strongest objection and review condition together. A reproducible derivation can still be wrong.

W3C PROV (2013); Toulmin (1958); Jøsang (2001, 2016); Sabater & Sierra (2001, 2005); Heuer (1999). Calibration compares numerical forecasts with observed frequencies, not rhetorical confidence.



PRESERVE REASONS THAT CAN CHANGE A DECISION


Agreement is cheap. Independence is not.

A role called “critic” is not evidence of independent judgment.

Attacks, not support arrows

Arguments a-f; inspected Tweety graph.
Set notation: attacks R; grounded set G (green).

$$R = \{(a,b), (b,c), (c,d), (d,e), (b,f), (e,f)\}$$
$$G = \{a, c, e\}$$



Semantics

Specify which sets of arguments count as acceptable. **Skeptical** asks about all relevant extensions; **credulous** asks about at least one.

Values

A defeat can depend on an audience's priorities. Make those priorities explicit; a formal solver cannot authorize the organization's values.

Truth

Acceptance in a graph does not prove a premise. Keep sources, warrants, uncertainty and independent checks outside mere voting.

What makes dissent useful?

Contrary evidence

Preserve the observation or argument, not just a theatrical opposing persona.

Authentic versus assigned

Human studies distinguish genuine minority positions from assigned advocacy; transfer requires testing.

Consequential criticism

Can an objection change the decision, halt an act, or require an additional check?

Three agreement traps

Sycophancy

Agreement with a user's belief can be rewarded even when misleading.

Evaluator self-preference

A judge may favor outputs resembling its own. A second voice is not automatically a second standard.

Shared errors

Common data, scaffolds and assumptions can correlate failures. Different names or model families prove little.

Close with a record

Question Options Criteria Evidence Dissent Owner Review trigger



Spend debate wisely

Use deliberation when distinct evidence or values can matter. Compare with a strong single worker and objective checks. End in a decision with an owner and authorized follow-through, not an indefinitely growing transcript.

Dung (1995); Bench-Capon (2003); Nemeth, Brown & Rogers (2001); Sharma et al. (2023); Panickssery et al. (2024). Diagram's accepted set is formal status, not a reliability claim about real-world facts.

RETRIEVE WHAT GOVERNS, NOT JUST WHAT RESEMBLES

Validity before similarity.



This is the book's recordkeeping rule, not a theorem about every retrieval system. A high retrieval score cannot establish that a rule governed an act.

ILLUSTRATIVE HISTORY

Two clocks, two questions

Monday

Tuesday

Wednesday

Thursday

Effective rule

New policy applies from Monday

Local record

Earlier version still held

Update received

What should have governed? Consult effective time. **What did the actor know?** Consult recorded information at the action time. Preserve both histories; do not rewrite the past.

RETAIN	QUESTION IT ANSWERS	CONTROL NEEDED
Policy	What governs this action?	Authority, scope, version, effective time
Precedent / rationale	Why was this tradeoff chosen?	Question, options, criteria, reasons, applicability
Episode / observation	What happened in a particular attempt?	Provenance, context and evidence status
Procedure / skill	How is this work performed?	Version, prerequisites, tested behavior
Transient context	What is useful in the current task?	Expiry, relevance and privacy limits

An archive is not the whole memory

Organizational knowledge persists in people, routines, culture, structures, settings and external archives.

Transactive memory is knowing who knows what. An expertise directory and a document search answer different questions.

These retention locations are not administrative levels such as individual, department and company.



Files, people, routines. TVA central files, 1936; not evidence of policy validity. TVA Web Team, CC BY 2.0, unchanged. Source & credit.

Append facts; revise judgments

An immutable observation can remain recorded while its interpretation changes. Current policy, available budget and final acceptance can depend on newly received facts.

CALM: monotonic semantics relate to coordination-free distributed consistency under assumptions. This does not remove communication or agreement from every decision.




Remember to forget

Retain commitments and evidence under applicable rules; decay transient context and quarantine superseded instructions. A plausible reflection is not evidence of improvement. Test old-but-active policy against newer-but-invalid discussion.

Walsh & Ungson (1991); Wegner (1987); MacLean et al. (1991); Hellerstein & Alvaro (2019). The two-clock example is illustrative, not a legal determination.

THE PROTOCOL CARRIES IT; THE INSTITUTION GOVERNS IT

Interfaces are not institutions.



MCP

Host / client ↔ tool and context server
Tool names, structured requests, results

Conceptual responsibility boundaries, not a mandatory product stack. A2A 1.0 leaves in-task authorization scope, validity and revocation to implementations. Version the protocol you actually use.

KQML / name the communicative act

```
(ask-one :content (PRICE IBM ?price)
:receiver stock-server :language LPROLOG
:ontology NYSE-TICKS)
```

ask-one requests one answer; ?price is the query variable.
Language declares syntax; ontology declares vocabulary. Neither guarantees shared interpretation or permits a trade.
Finin et al. (1994), §3; original query, whitespace normalized.
Inspected, not executed.

REACT
Reasoning with actions and observations; not institutional authority or independent acceptance.

METAGPT
Role-oriented processes and artifacts; role names alone do not enforce duties.

REFLEXION
Feedback across attempts; not proof that feedback is correct or outcomes improve.

MAGENTIC-ONE
Orchestration, progress tracking and replanning; not continuity of obligations or authority after plan changes.

Keep durable state outside context

Record actor, policy, revision, operation identity, attempt and evidence. Context can be replaced; institutional records must remain inspectable.

Inspect, don't infer

For a product: record version, configuration, starting state, operation, observed result, evidence and limitation. Separate advertised, available and demonstrated capability. Recover useful historical contracts, not necessarily old implementations.

Chapters 12 and 15: Finin et al. (1994); FIPA ACL (2002); Singh; MCP 2025; A2A 1.0 (2026); Yao et al.; Shinn et al.; Wu et al.; Hong et al.; Fournery et al.; Park et al. This is a mechanism comparison, not a current-product ranking.

A2A

Agent ↔ independently exposed agent
Discovery, messages, tasks, artifacts

Speech acts / saying can be doing

Separate content (what is said), force (the communicative act), and effect (what happens).

Request	"Please send the report."
Promise	"I undertake to send it."
Inform	"I sent the report."

Editorial examples: a request is not an accepted duty; a promise is not performance; an assertion is not verified evidence.
FIPA-ACL likewise distinguishes request, agree, refuse, inform and failure. KQML is not a FIPA product.

AUTOGEN
Programmable agent interaction; round-robin speaking is not capability-based allocation.

GENERATIVE AGENTS
Memory, reflection and planning in simulation; believable behavior is not institutional readiness.

Bind control to the real effect

A well-formed request can still be unauthorized. Check admission, stale authority, idempotency and uncertain-effect recovery where they can be enforced.

ONE GOVERNED OUTCOME BEFORE A WHOLE COMPANY

Build the complete loop. Then widen it.

DEFINE	BIND	EXERCISE
1. Outcome. Name the result and how it will be observed.	4. Responsibility. Specify roles, duties and enactment.	8. Failure. Interrupt, duplicate, delay and revoke in isolation.
2. Baseline. Try a capable single worker or deterministic process.	5. Authority. Scope data, actions, budgets and expiry.	9. Review. Distinguish submission, acceptance and observed outcome.
3. Acceptance. Define the check and who applies it.	6. Coordination. Match dependencies to the smallest sufficient mechanism.	10. Expand. Increase scope only when representative results justify it.
	7. Evidence. Preserve observations, reasons and revisions.	

Condensed editorial construction sequence from Chapters 13-14. A passing local test does not prove open-ended autonomy.

Choose the case for the question

PUBLIC INCIDENT

Mars Climate Orbiter

Boundary: exchanged data without shared unit meaning and adequate checks.

Transfer: verify semantics at handoff. Passing a physical test does not validate navigation data.

PUBLIC DECISION

Air Canada

Boundary: a chatbot's public representation and the company's responsibility.

Transfer: verify the official rule and authority behind a promise. Fluent answers do not settle accountability.

PUBLIC INCIDENT

Knight Capital

Boundary: deployment and trading controls before consequential effects.

Transfer: test admission, monitoring and containment. Good output elsewhere does not excuse an unguarded action.

PUBLIC INCIDENT

GitLab recovery

Boundary: recovery ownership, distinct copies, dependent restores and checks.

Transfer: separate availability from data completeness and describe observed loss.

RESEARCH FAMILY

Collective judgment

Boundary: evidence, aggregation rules, disagreement and correlated errors.

Transfer: preserve contrary reasons and specify closure. A majority is not independent corroboration.

LOCAL LIBRARY SPECIMEN

Rinda work claiming

Boundary: matching, observation, consuming claims and duplicates.

Transfer: distinguish local atomicity from durable ownership, recovery and acceptance.

One result; several non-substitutable checks

Intelligible handoff Authorized act Completed attempt Accepted result Observed consequence

 *Keep the limits*

Documented incidents, authors' thought experiments, formal models, benchmark simulations and local checks supply different evidence. No single case demonstrates everything an artificial organization needs.

Cases and exact source records: book Chapters 2 and 14; NASA/JPL (1999), Moffatt v. Air Canada (2024), SEC Knight Capital order (2013), GitLab postmortem (2017). Collective judgment and Rinda have different evidence status.



FOR THE NEXT TASK, EXPERIMENT OR REVIEW

Test the distinction.

BEFORE DISPATCH

Whose act is this?

Outcome → responsible role → current actor → authority → acceptance criterion.

- Credentials are not permission.
- Scope is not a role name.
- A plan is not a commitment lifecycle.

Authority / 06

BEFORE CLOSURE

Which claim is supported?

Submission → independent check → acceptance → observed effects.

- A trace is not a proof.
- Consensus is not independence.
- Silence is not verified health.

Closure / 09

WHEN WORK FAILS

Which boundary broke?

Identity → ownership → external effect → acknowledgment → recovery.

- A lost reply is not a failed effect.
- A replacement is not a completed handover.
- Retry is a new attempt under current authority.

Concurrency / 05

BEFORE MORE AUTONOMY

What evidence changes the limit?

Representative task → capable baseline → adverse cases → measured consequences.

- Recency is not validity.
- A success example is not a reliability estimate.
- A benchmark is not a whole organization.

Memory / 12

Four open research demands

Governable specifications: can conversational authoring preserve unambiguous duties and revision rules? **Affordable oversight:** what review burden and miss cost appear under realistic bursts? **Institutional memory:** can stale-but-similar records be rejected without losing useful precedent? **Long-horizon reliability:** does performance survive interruptions, changed conditions and external effects?

The literature does not establish reliable open-ended company autonomy, universal supervision capacity, or one best memory architecture. State what evidence would disconfirm your preferred approach.

References & return routes

Primary companion source: [Artificial Organizations: A Study Monograph](#), v1.2, research through 16 September 2026. Curated and written by AI under the direction of Dr. Paulo Salem. Every sheet's footer points to the relevant chapters and final PDF pages.

Organization: Simon, *Administrative Behavior* (1947); Coase, "The Nature of the Firm" (1937); Garicano, "Hierarchies and the Organization of Knowledge in Production" (2000).

Coordination: Smith, "The Contract Net Protocol" (1980); Gelernter, "Generative Communication in Linda" (1985); Hoare, "Communicating Sequential Processes" (1978).

Authority: Parasuraman et al., "A Model for Types and Levels of Human Interaction with Automation" (2000); Burrows, "The Chubby Lock Service" (2006).

Evidence: W3C, *PROV-O* (2013); Toulmin, *The Uses of Argument* (1958); Heuer, *Psychology of Intelligence Analysis* (1999).

Decision: Dung, "On the Acceptability of Arguments..." (1995); MacLean et al., "Questions, Options, and Criteria" (1991).

Memory: Walsh & Ungson, "Organizational Memory" (1991); Wegner, "Transactive Memory" (1987). [Full citations, links and additional sources.](#)
Icons: Lucide 0.468.0, ISC license.



Return to the evidence

This guide condenses all 17 chapter Core points and 46 Key concepts. It is an editorial synthesis, not a new experiment, systematic review or safety certification. Qualifications that decide an application must be checked in the full discussion.