

The ACE Economy

Status: public working report Author: Njål Gaute Solland Date: 2026-08-27

1. Core proposition

The model begins with the completion required by reality, not with humans, jobs or agents. The primary production unit is the **Governed Completion Unit (GCU)**: one accepted outcome completed under current authority, constraints, governed execution and sufficient evidence.

Humans, agents, models, tools, machines and external services are candidate production resources. **Attentive Cognitive Evaluation (ACE)** remains the scarce human control input: qualified, contextual and authorized attention applied at the moment evaluation, authorization or escalation actually requires it.

The organization that survives is the one that produces the required portfolio of valid GCUs with the best admissible combination of quality, safety, time and total resource cost. Human organization size and disciplinary composition are derived from that demand; they are not starting variables.

The fundamental economic shift is:

Labour economy -> Knowledge economy -> Attention economy -> ACE economy

In the ACE economy, intelligence is abundant, but legitimate evaluation remains scarce.

2. Primary output unit: Governed Completion Unit

A **Governed Completion Unit (GCU)** is one completed outcome that:

1. satisfies its declared acceptance criteria;
2. was admissible under current policy and fresh authority at execution time;
3. traversed the required governed effect path;
4. produced sufficient execution and outcome evidence;
5. was observed as an outcome, not merely attempted or generated.

A draft, recommendation, blocked action, denied action, uncleared escalation, direct machine effect or unobserved result is zero GCU.

GCUs are comparable only inside the same declared outcome and consequence class.

An **H-GCU** has a human primary execution path. An **M-GCU** has a machine primary execution path. An M-GCU may consume ACE for authorization or escalation without becoming an H-GCU. Both must satisfy the same completion contract.

Cost per expected GCU = total governed capacity cost per attempt / expected valid GCU per attempt

GCU contribution margin = price per GCU - cost per expected GCU

2.1 Human control input: ACE-second

An **ACE-second** is one second in which a human:

1. is paying relevant attention;
2. has sufficient context;
3. applies competent judgment;
4. holds valid authority for the matter;
5. accepts or carries responsibility for the consequence.

Ordinary paid time is not automatically ACE. Reading, waiting, status reporting, searching, formatting and duplicated approval may consume attention without producing ACE.

ACE quality

Let:

- A_t = attentive presence;
- C = context completeness;
- J = judgment competence;
- M = mandate validity;
- R = responsibility binding.

Each factor ranges from 0 to 1.

ACE quality = A_t × C × J × M × R

Effective ACE-seconds = clock seconds × ACE quality

The factors are multiplicative because a decision with no attention, context, competence, mandate or responsibility is not fully valid ACE.

3. Cost of ACE

Let:

- C_h = fully loaded human cost per paid second;
- p_a = share of time with real attention;
- p_c = probability of sufficient context;

- p_j = probability relevant judgment is required and available;
- p_m = probability valid mandate is present;
- p_r = probability responsibility is clearly bound.

Then:

$$\text{Cost per effective ACE-second} = C_h / (p_a \times p_c \times p_j \times p_m \times p_r)$$

This explains why senior decision capacity is far more expensive than salary divided by work hours. Most organizations destroy ACE through interruption, reconstruction of context, unnecessary escalation, duplicate approval and unclear ownership.

4. Complementary machine unit: ADS

An **AI Deliberation Second (ADS)** is one second of model, tool or agent activity used to reduce uncertainty, produce options, assemble evidence or prepare action.

ADS includes:

- retrieval;
- synthesis;
- comparison;
- simulation;
- anomaly detection;
- policy checking;
- evidence assembly;
- drafting;
- execution preparation.

ADS and ACE are complementary inputs. GCU is the governed completion output.

ADS produces cognitive supply.

ACE legitimizes consequential choice.

5. The new production function

Traditional simplified production:

$$\text{Value} = \text{labour} \times \text{capital}$$

Knowledge economy:

$$\text{Value} = \text{labour} \times \text{knowledge} \times \text{capital}$$

GCU economy:

$$\text{Governed production} = \text{valid H-GCU} + \text{valid M-GCU, grouped by declared outcome class}$$

$$\text{Realized value} = \text{governed completions} \times \text{outcome value} - \text{total governed capacity cost}$$

Where:

- synthetic capacity includes models, agents, tools, compute and reusable workflows;
- ACE supplies focused human evaluation and legitimate authority;
- execution integrity ensures identity, mandate, policy, context, evidence and enforcement;
- learning retention determines whether each outcome improves the next one.

A more operational form is:

$$RV = O \times Q \times L \times E$$

Where:

- O = output capacity;
- Q = evaluation quality;
- L = legitimacy factor;
- E = execution integrity.

If legitimacy or execution integrity approaches zero, enterprise value collapses even when output is excellent.

6. The ACE conversion chain

The canonical production chain is:

Required outcome -> GCU contract -> Candidate configurations -> Admissibility -> Production -> ACE when required -> Governed execution -> Outcome -> Receipt -> Learning

ACE is concentrated at the points where ambiguity, trade-offs, legitimacy and consequence must be resolved.

The strategic objective is not to eliminate ACE. It is to eliminate waste around ACE and reserve it for the decisions where it produces the highest marginal value.

7. The true AI ROI

Traditional AI ROI is usually calculated as labour cost removed plus throughput gained. This misses rework, governance, false actions, incident risk, review burden, lost learning and authority defects.

The proposed primary measure is **ACE-Adjusted Return on Intelligence (AARI)**.

Let:

- V = gross economic value created;
- L_a = expected losses avoided;
- C_r = economic value of ACE capacity released;
- K = reusable knowledge created;

- S = speed value from earlier valid outcomes;
- C_ai = model, compute, data and tool cost;
- C_g = governance, verification and control cost;
- C_ace = ACE cost consumed;
- C_rw = rework and correction cost;
- C_i = expected incident and remediation cost;
- q = quality factor;
- m = mandate-validity factor;
- e = execution-integrity factor;
- t = trust and adoption factor.

Then:

$$AARI = [(V + L_a + C_r + K + S) \times q \times m \times e \times t] - (C_{ai} + C_g + C_{ace} + C_{rw} + C_i) / (C_{ai} + C_g + C_{ace})$$

This is the actual economic return from converting machine intelligence into legitimate enterprise action.

8. ACE Yield: the new productivity metric

The primary production KPIs are valid GCU by class, time to GCU, cost per expected GCU, contribution margin per GCU and expected loss per GCU.

ACE Yield remains the human-control efficiency metric:

$$\text{ACE Yield} = \text{risk-adjusted realized value} / \text{effective ACE-seconds consumed}$$

At enterprise level:

$$\text{Enterprise ACE Yield} = \text{total risk-adjusted realized value} / \text{total effective ACE-seconds}$$

At organization-design level, gross human capacity must be separated from the scarce attention consumed merely to keep the human organization integrated:

$$\text{Net ACE capacity} = \text{gross effective ACE-seconds} - \text{coordination ACE-seconds}$$

This replaces crude measures such as revenue per employee when most preparatory work is produced by models, agents, software, external platforms and customers. Human headcount is not productive scale when additional people consume more ACE through coordination than they add through judgment, authority, context or relationships.

High ACE Yield means the organization uses scarce human evaluation where it has the highest marginal value.

9. ACE Leverage

ACE Leverage measures how much useful synthetic work is compressed into each effective ACE-second.

Raw ACE Leverage = useful ADS / effective ACE-seconds

A safer measure is:

Effective ACE Leverage = useful ADS × quality × context completeness × evidence coverage / effective ACE-seconds

High raw leverage with low quality is not leverage. It is cognitive dumping.

10. ACE Utilization

ACE Utilization measures the share of scarce human attention spent on work that genuinely requires human evaluation.

ACE Utilization = effective ACE-seconds on human-required matters / total senior attention seconds

Low ACE Utilization indicates that expensive people are spending time on search, formatting, status transfer, duplicate review, routine checking and coordination.

The target is not 100 percent. Reflection, relationships, learning and recovery are productive preconditions for future ACE.

11. Authorization Compression

Authorization Compression measures how much preparation is condensed into the smallest sufficient human decision surface.

Authorization Compression = preparation seconds avoided / ACE-seconds at clearance

The goal is to present the decision owner with:

- the precise action;
- the relevant context;
- the available alternatives;
- the consequence range;
- the evidence;
- the mandate boundary;
- the unresolved uncertainty.

A good system does not ask the human to repeat machine work. It asks the human to perform the irreducibly human evaluation.

12. Decision surface area

The **decision surface area** is the amount of information a human must inspect to make a legitimate decision.

Decision Surface Efficiency = sufficient decision information / total information presented

Poor systems increase review burden by presenting logs, summaries, dashboards and generated prose without identifying the actual choice.

Strong systems reduce the surface to the smallest sufficient set without hiding material uncertainty.

13. ACE capital

ACE is not only a flow. Organizations also accumulate or destroy **ACE capital**.

ACE capital consists of:

- domain experience;
- judgment patterns;
- institutional memory;
- trusted relationships;
- authority structures;
- escalation knowledge;
- verified decision history;
- reusable evaluation criteria;
- calibrated intuition;
- tested ability to take responsibility.

ACE capital appreciates when decisions are documented, outcomes are observed and learning is retained.

It depreciates when senior capability leaves, context is fragmented, junior pathways collapse, decisions are outsourced without learning, or automation hides the reasons behind outcomes.

ACE capital balance

Closing ACE capital = opening ACE capital + learning + mentoring + documented outcome feedback - attrition - fragmentation - unobserved automation - judgment atrophy

14. The ACE balance sheet

An AI-native enterprise should track five forms of capital:

1. **Financial capital** — cash and funding capacity.
2. **Synthetic capital** — models, agents, compute, tools and reusable workflows.
3. **Context capital** — data, provenance, organizational memory and current state.
4. **ACE capital** — human evaluation, authority, accountability and judgment.
5. **Trust capital** — stakeholder willingness to accept decisions and outcomes.

The enterprise survives only when synthetic capital does not outrun context, ACE and trust capital.

15. The survival equation

A practical survival condition is:

Survival capacity = cash runway × ACE Yield × adaptation speed × trust retention × execution integrity / concentration risk

A firm fails in the ACE economy when one or more of the following occurs:

- machine output grows faster than valid evaluation capacity;
- human review becomes a hidden bottleneck;
- decisions are made without authority clarity;
- quality debt grows faster than learning;
- key context is locked in people or vendors;
- customer trust falls after automated errors;
- compute and tool costs erase apparent labour savings;
- junior development collapses and future ACE capital is not replenished.

16. The ACE operating model

Every consequential workflow should be decomposed into four categories:

1. **Production** — generate, retrieve, calculate, transform.
2. **Evaluation** — determine quality, relevance, uncertainty and trade-offs.
3. **Authorization** — decide whether the action is legitimate and within mandate.
4. **Execution** — create real-world consequence.

The failure of most AI transformations is that they automate production but leave evaluation, authorization and execution boundaries undefined.

17. The new firm

The basic production unit is not a human, job or team. It is a GCU contract plus the admissible resource configuration that completes it.

For every required GCU class, the organization compares candidate human-first and machine-first paths on:

- acceptance quality and completion probability;
- current authority and consequence limits;
- no-direct-consequence-path requirement conformance;
- evidence and outcome observability;
- queue, context, execution, coordination, verification and ACE response time;
- price, full resource cost, rework and expected loss;

- availability, sickness, leave, uptime and dependency reserve.

The fastest or lowest-cost path is selected only after all quality, authority, safety and evidence gates pass.

The earlier **ACE cell** becomes a derived organizational form: the bounded humans and synthetic resources required to support a portfolio of GCUs. It is not the primary production unit.

17.1 Derived human coordination constraint

Human capacity enters only after the GCU portfolio has selected its admissible resource configurations. The resulting human disciplines, ACE response coverage, independent duties and availability reserves determine the required human composition.

Human communication is not external overhead. It consumes scarce attention through context reconstruction, alignment, negotiation, handoffs, conflict resolution and authority clarification. Human coordination is therefore a downstream GCU production cost.

Let H be the number of humans inside one integrated coordination core. The upper-bound number of possible bilateral human relationships is:

$$\text{potential_coordination_edges} = H \times (H - 1) / 2$$

This gives 45 possible edges at 10 humans, 435 at 30 and 1,225 at 50. These are possible relationships, not observed message counts; actual load depends on work coupling, network topology and shared context. The equation identifies the coordination surface that may need to be maintained.

An observed coordination edge is a human pair that must directly exchange material context, align interdependent work, resolve authority or negotiate a consequential trade-off during the measurement window. Edge frequency, duration and consequence must be measured separately; edge count alone is not a burden score.

coordination ACE-seconds are scarce human attention seconds consumed solely to maintain integration among people rather than to perform outcome-specific evaluation, authorization, learning or relationship work.

$$\text{Net ACE capacity} = \text{gross effective ACE-seconds} - \text{coordination ACE-seconds}$$

$$\text{ACE Coordination Ratio} = \text{coordination ACE-seconds} / \text{total scarce human attention seconds}$$

A human should be added to an integrated core only while the marginal valid outcome value and ACE capacity they add exceed the marginal communication, context, authority and integration load they create.

After GCU demand has derived the necessary human composition, the following coordination regimes may be tested as a secondary design hypothesis:

- 1-10 humans: small ACE core;
- 11-30 humans: medium ACE core;

- 31–50 humans: large ACE core;
- more than 50 humans: federation of ACE cores, not a larger integrated coordination unit.

The claim concerns the human coordination core. It does not count agents, models, bounded external capability, customers or physically distributed operators who do not all belong to one integrated decision network. A legal entity, industrial system or market presence may remain very large while its governing human structure is a federation of bounded cores.

The 50-person boundary is a falsifiable design hypothesis, not a universal law.

The external 2022 SME technology-implantation review by Reascos and Carvalho supports the mechanism but not the threshold: across 75 reviewed articles, organizational factors were more influential than the technological approach. Soares' review in the same volume identifies team communication as the most reported Scrum challenge. Neither study tests agent-native firms or proves the proposed size bands.

Source: Pereira, Bianchi and Rocha (eds.), Digital Technologies and Transformation in Business, Industry and Organizations (2022), DOI 10.1007/978-3-031-07626-8.

18. New organizational hierarchy

The future hierarchy is not primarily seniority by headcount. It is authority over larger consequence domains.

1. Workflow operator.
2. Workflow evaluator.
3. Workflow owner.
4. Outcome owner.
5. Portfolio owner for human and synthetic capacity.
6. Mandate architect.
7. Institutional consequence owner.

Compensation should increasingly track consequence scope, ACE quality, reusable system creation and verified outcomes rather than volume of personal output.

19. The new wage model

Wages in exposed knowledge work will split into four components:

Compensation = base capability + ACE premium + authority premium + outcome premium

- **Base capability** pays for general skill.
- **ACE premium** pays for scarce judgment under uncertainty.
- **Authority premium** pays for valid decision mandate and accountability.
- **Outcome premium** pays for risk-adjusted realized results.

Routine production receives less premium as synthetic supply expands.

People who can define objectives, detect failure, evaluate ambiguity, own consequences and improve systems receive more.

20. The new cost model

The relevant cost is not payroll. It is **Total Governed Capacity Cost (TGCC)**.

TGCC = human labour + ACE cost + compute + models + data + tools + integration + governance + external capacity + rework + incidents + concentration risk + knowledge depreciation

A workflow is economically improved only when TGCC per valid outcome falls without weakening trust, learning or resilience.

21. The new unit economics

For each validated outcome:

Contribution per outcome = realized customer value - TGCC - expected failure loss

ACE contribution margin = contribution per outcome / ACE-seconds consumed

Synthetic efficiency = valid outcomes / total ADS

Governance efficiency = prevented or modified invalid actions / governance cost

Learning return = future cost or error reduction generated by the current outcome

22. The new P&L

An AI-native P&L should distinguish:

Revenue and value

- revenue from validated outcomes;
- avoided loss;
- speed value;
- retained customer trust;
- reusable knowledge created.

Direct capacity cost

- human production cost;
- ACE cost;
- model and compute cost;

- tool and data cost;
- external capacity.

Control cost

- evaluation;
- verification;
- authorization;
- execution governance;
- audit and evidence.

Quality cost

- rework;
- false positives and false negatives;
- incidents;
- compensation and remediation;
- trust damage.

Capital depreciation

- lost institutional knowledge;
- vendor lock-in;
- model dependence;
- judgment atrophy;
- junior pipeline erosion.

23. The new board dashboard

A board should see:

- Enterprise ACE Yield;
- AARI;
- time to validated outcome;
- ACE Utilization;
- ACE Coordination Ratio;
- human coordination core size and federation map;
- validated outcome value per observed coordination edge;
- authorization compression;
- rework burden;
- evidence coverage;
- exception density;
- governed autonomy rate;
- incidents per million consequential actions;
- learning latency;
- context concentration;

- model and vendor concentration;
- tested human takeover;
- ACE capital replenishment rate.

Headcount reduction is not a primary success metric.

24. The ACE capital replenishment rate

Because junior production tasks are automated early, firms risk consuming ACE capital without replenishing it.

ACE replenishment rate = new effective ACE capacity created / ACE capacity lost

A rate below 1 means the organization is liquidating its future judgment base.

New apprenticeship must include:

- real bounded outcomes;
- visible evaluation;
- access to senior reasoning;
- correction and escalation practice;
- gradual mandate expansion;
- responsibility for improving workflows;
- outcome feedback.

25. The ACE debt model

Organizations accumulate **ACE debt** when they scale automation without scaling evaluation, authority clarity and learning.

ACE debt includes:

- unresolved exceptions;
- undocumented judgment;
- unclear mandate;
- review backlog;
- missing evidence;
- overdependence on key people;
- untested takeover;
- automated decisions no one understands;
- rules that no longer match reality.

ACE debt ratio = unresolved judgment and authority obligations / available ACE capacity

High ACE debt creates sudden operational failure even when day-to-day throughput appears strong.

26. The intelligence inflation thesis

As machine intelligence becomes cheaper, the supply of plausible outputs increases faster than the supply of legitimate evaluation.

This creates **intelligence inflation**:

- more reports, proposals, analyses and options;
- lower average cost of production;
- lower signal-to-noise ratio;
- higher verification demand;
- greater value of trusted judgment.

The price of generic output falls.

The price of trusted evaluation rises.

27. The ACE scarcity law

ACE Scarcity Law:

As the marginal cost of machine intelligence approaches zero, the marginal value of legitimate human evaluation rises, provided consequential action remains bounded by human institutions.

This is not because humans outperform machines at every cognitive task. It is because responsibility, authority, trust and legitimacy do not scale at the same rate as inference.

28. The authorization bottleneck law

Authorization Bottleneck Law:

When production capacity grows faster than legitimate authorization capacity, enterprise throughput becomes constrained by ACE rather than compute.

Adding more agents then increases queues, review burden and risk instead of value.

29. The context conservation law

Context Conservation Law:

Synthetic capacity cannot create reliable enterprise value from context that the organization has failed to capture, preserve and bind to current purpose.

Automation without context does not eliminate human work. It moves it into correction, exception handling and incident response.

30. The execution legitimacy law

Execution Legitimacy Law:

Technically possible action is not equivalent to authorized consequence.

Enterprise value exists only when an action is both operationally correct and institutionally legitimate.

31. The learning compounding law

Learning Compounding Law:

The durable return from an AI-assisted outcome is not only the value of that outcome, but the reduction in cost, uncertainty and error for all future similar outcomes.

This is why evidence, receipts, evaluation sets and reusable workflow improvements are economic assets.

32. Survival strategy for individuals

A person survives and advances by moving upward from production toward evaluation, system ownership and consequence ownership.

The individual strategy is:

1. stop competing primarily on first-draft production;
2. become excellent at framing goals and constraints;
3. learn to evaluate machine output;
4. build domain-specific context;
5. understand authority and consequence;
6. own a workflow, not only tasks;
7. document outcomes and improvements;
8. build reusable systems;
9. cultivate trust and relationships;
10. increase the economic value of each ACE-second.

The future premium goes to people who can answer:

- What are we actually trying to achieve?
- What evidence is sufficient?
- What could go wrong?
- Who is allowed to decide?
- What must remain human?
- What should be stopped?
- What did the outcome teach us?

33. Derived small ACE core: 1-10 humans

A small ACE core should exploit minimal coordination distance and high ownership clarity.

It should:

- organize around a few critical value streams;
- use responsible generalists with broad consequence ownership;
- keep purpose, evaluation, authority and outcome close together;
- standardize repeatable workflows before adding people;
- use agents and bounded external capability to expand production and specialist reach;
- keep one clear escalation route for consequential actions;
- preserve customer relationships, quality and irreversible choices with humans;
- measure cash flow, validated value and coordination ACE per human;
- avoid vendor, principal and single-expert concentration.

Its advantage is not small revenue. It is the shortest possible distance between purpose, context, judgment, authority and outcome.

34. Derived medium ACE core: 11-30 humans

A medium ACE core adds capability breadth without accepting a managerial hierarchy as the default coordination mechanism.

It should:

- form explicit ACE cells around outcomes, not functions around task categories;
- use shared identity, context, evidence and agent infrastructure;
- assign outcome and mandate owners inside each value stream;
- make cross-cell interfaces explicit and machine-readable where possible;
- eliminate status-transfer and duplicate approval work;
- measure ACE Yield and ACE Coordination Ratio by value stream;
- acquire intermittent specialist capability on demand;
- retain institutional context while avoiding permanent coordination layers.

Its main risk is reproducing functional silos and human handoffs inside an otherwise agentic operating model.

35. Derived large ACE core and the end of large as a primary unit

A 31-50-human core is the largest candidate coherent ACE unit. It should exist only when the added domain breadth, relationships, physical supervision or reserved authority cannot be provided by synthetic capacity, bounded external capability or a federation of smaller cores.

Above 50 humans, the default design move is not another management layer. It is federation:

- bounded ACE cores with explicit outcomes and mandates;
- a common identity, authority and execution-governance boundary;
- explicit cross-core contracts and evidence interfaces;
- shared context only where required for the joint consequence;
- portfolio-level capital and exposure allocation;
- no organization-wide all-to-all communication assumption;
- independent evidence, debt, risk and recovery functions where separation is required.

Large therefore disappears as a single integrated human coordination category. It may remain as a legal entity, capital pool, brand, physical operation, regulated envelope or network of ACE cores. Economic scale comes from synthetic capacity and federation, not from continuously enlarging one human coordination graph.

Physical operations, care, construction, logistics and statutory staffing may require far more than 50 people. Those workers are not automatically one integrated ACE core. The relevant boundary is the group that must continuously share context, negotiate priorities, resolve authority and coordinate consequential decisions.

Falsification conditions

The proposed size bands must be modified or rejected if matched real-world evidence shows one or more of the following:

- integrated cores above 50 repeatedly produce higher risk-adjusted valid outcome value per total scarce human attention second after full coordination cost;
- observed coordination ACE does not rise materially with core size or work coupling;
- federation creates more context loss, rework, latency, incident risk or duplicated capability than it removes;
- critical trust, tacit knowledge, physical supervision or legal duties cannot be bounded across cells;
- synthetic and external capacity cannot replace enough permanent internal capability to prevent capability failure.

Until tested, 50 is a candidate design ceiling, not a universal optimum.

36. The new role of management

Management shifts from information transfer to four functions:

1. purpose design;
2. mandate allocation;
3. workflow architecture;
4. exception and consequence ownership.

A manager who only summarizes status loses value.

A manager who converts uncertainty into clear purpose, bounded authority and reusable operating systems gains value.

37. The new role of finance

Finance must move from labour-centric budgeting to governed capacity accounting.

Finance should price:

- one ACE-hour by role and domain;
- one validated outcome;
- one agent workflow run;
- one exception;
- one unit of rework;
- one governance decision;
- one avoided incident;
- one unit of reusable knowledge;
- one concentration risk exposure.

Capital allocation should favor systems that increase ACE Yield and learning retention, not merely token volume or automation count.

38. The new role of HR

HR becomes an architect of capability and authority rather than a custodian of static job descriptions.

HR should manage:

- capability portfolios;
- mandate progression;
- ACE capital replenishment;
- human-agent workflow design;
- evaluation competence;
- apprenticeship pathways;
- authority clarity;
- responsibility boundaries;
- judgment health and overload.

The relevant unit is no longer only the role. It is the person's demonstrated capacity to own increasingly consequential outcomes.

39. The new role of IT and security

IT provides synthetic and context capacity.

Security must move beyond access control toward consequence control.

Identity and credentials prove technical ability.

Execution governance must prove that a specific action is allowed now, for this purpose, under this mandate, with this evidence and consequence profile.

40. execution-boundary check in the ACE economy

execution-boundary check is the control point where abundant synthetic capacity meets scarce legitimate ACE.

Its economic function is to:

- protect ACE from low-value review;
- prepare the smallest sufficient decision surface;
- verify identity, authority, context and evidence;
- route only material exceptions to humans;
- prevent technically valid but illegitimate actions;
- record clearance and execution evidence;
- turn outcomes into reusable learning;
- increase ACE Yield and AARI.

Evaluation assesses evidence.

Authorization remains distinct from evaluation.

Enforcement applies the resulting decision at the declared consequence boundary.

Execution and outcome evidence are recorded separately.

41. The ACE maturity model

Level 0 — Unmeasured AI use

Individuals use models. No workflow ownership, cost model or evidence.

Level 1 — Assisted production

AI accelerates drafts and analysis. Humans repeat most checking manually.

Level 2 — Evaluated workflows

Defined quality criteria, human review and measured outcomes.

Level 3 — Bounded delegation

Agents operate within explicit mandates, limits and escalation rules.

Level 4 — Governed execution

Consequential actions are checked at commit time and produce verifiable receipts.

Level 5 — ACE-native enterprise

ACE is measured, allocated and replenished as the primary scarce production factor. Synthetic capacity, authority and learning are managed as one operating system.

42. The 90-day conversion plan

Days 1-30: Measure reality

- select five consequential value streams;
- separate production, evaluation, authorization and execution;
- measure human attention, review, rework, exceptions and outcome time;
- estimate effective ACE-seconds;
- calculate current ACE Yield;
- identify context and authority gaps.

Days 31-60: Redesign the workflow

- remove duplicate handoffs;
- automate bounded preparation;
- define the smallest sufficient decision surface;
- create evaluation criteria;
- specify mandate and escalation rules;
- introduce evidence capture.

Days 61-90: Govern and compare

- introduce execution checks;
- measure new ACE Yield and AARI;
- compare rework, cycle time, incidents and customer value;
- convert successful patterns into reusable workflows;
- reinvest part of the released capacity in ACE capital and learning.

43. Canonical metrics

Primary completion metrics:

- Governed Completion Unit by declared class;
- H-GCU and M-GCU;
- valid GCU per period;
- time to GCU;

- price per GCU;
- total cost per expected GCU;
- contribution margin per GCU;
- expected loss per GCU;
- ACE-seconds per GCU;
- timely ACE coverage for authorization and escalation;
- blocked direct effects and GCUs prevented by no-direct-consequence-path requirement.

Supporting capacity metrics:

1. ACE-second.
2. Effective ACE-seconds.
3. Cost per ACE-second.
4. ACE Yield.
5. Enterprise ACE Yield.
6. ACE Leverage.
7. ACE Utilization.
8. Authorization Compression.
9. Decision Surface Efficiency.
10. ACE capital balance.
11. ACE replenishment rate.
12. ACE debt ratio.
13. Total Governed Capacity Cost.
14. ACE contribution margin.
15. ACE-Adjusted Return on Intelligence.
16. Time to validated outcome.
17. Governed autonomy rate.
18. Evidence coverage.
19. Exception density.
20. Learning latency.
21. Human coordination core size.
22. Coordination ACE-seconds.
23. ACE Coordination Ratio.
24. Validated outcome value per observed coordination edge.

44. Canonical statements

AI makes intelligence abundant. GCU makes governed completion measurable. ACE remains the scarce human control input.

The real ROI of AI is the risk-adjusted value created per valid GCU after total resource cost, expected loss and ACE consumption.

The future firm is not a collection of jobs. It is a governed system that selects the best admissible human and machine configuration for each required completion.

The organization survives by increasing ACE Yield faster than complexity, risk and trust loss grow.

Human coordination, availability and ACE response become measurable GCU production costs only after the task determines that humans are required.

Beyond the coherent human coordination limit, scale should come from synthetic capacity or federation, not a larger integrated core.

H-GCU and M-GCU share one completion contract but have different time, availability and cost functions.

Machine output with a direct consequence path is not production. It is inadmissible and yields zero GCU.

ACE is qualified attention at the right time: routine bounded M-GCU consumes none, while ambiguity and material consequence must escalate before effect.

Production can be delegated. Consequence must remain governed.

45. Final thesis

The defining production problem of the AI era is not how many humans or agents an organization should contain.

It is which governed completion reality requires, and which admissible combination of humans, machines, models, tools and external capacity can complete it best, safest, fastest and at the lowest total resource cost.

GCU is the common output unit. H-GCU and M-GCU expose different prices, costs, timing, availability and failure modes against the same completion contract.

ACE remains scarce where qualified human attention must arrive with the right context, competence and authority at the exact authorization or escalation point. no-direct-consequence-path requirement makes that control economically possible: routine bounded machine work may scale without ACE, while ambiguity and material consequence remain unable to produce an effect until valid clearance exists.

Human headcount, disciplinary composition and federation boundaries are therefore derived outputs of the required GCU portfolio.

The winners will be those who maximize valid GCU production while minimizing total governed capacity cost, expected loss, latency and unnecessary ACE consumption.