

From Jobs to Governed Completion Units

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

Primary unit

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

1. satisfies its declared acceptance criteria;
2. was admissible under current policy and 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.

An attempt, draft, recommendation, blocked action, denied action, uncleared escalation or unobserved effect is zero GCU.

GCU's are only comparable inside the same declared outcome and consequence class. A low-risk information completion and an irreversible material completion are not interchangeable units.

Task first

The model starts with the required completion, never with humans or headcount:

completion contract -> candidate resource configurations -> admissibility -> execution -> outcome -> receipt -> GCU

The completion contract declares the outcome class, acceptance thresholds, consequence limits, authority, evidence, maximum latency and the optimization objective after the safety gates pass.

Humans, agents, models, tools, machines and external services are candidate inputs. Organization size and disciplinary composition are derived outputs.

H-GCU and M-GCU

H-GCU is a GCU whose primary execution path is human.

M-GCU is a GCU whose primary execution path is machine. An M-GCU may consume human ACE for authorization or escalation without becoming an H-GCU.

Both are measured against the same completion contract. They differ in resource mix, availability, latency, cost, coordination, failure modes and ACE intensity.

Supporting AI does not turn an H-GCU into an M-GCU. Human escalation does not turn an M-GCU into an H-GCU. The producer label follows the primary execution path; every supporting resource remains visible in the receipt and cost model.

Price and total cost

Price per GCU and cost per expected GCU must be recorded separately.

H-GCU cost includes paid human time, coordination, context loading, sickness and leave reserve, training and turnover, tools, governance, evidence, rework and expected incident loss.

M-GCU cost includes compute, energy, models, agents, tools, APIs, data, integration, downtime reserve, governance, evidence, ACE escalation, retry, rework and expected incident loss.

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

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

Cheap generation is irrelevant if the action cannot become a governed completion.

Time and availability

Expected completion time = queue + context load + execution + coordination + verification + evidence + expected ACE response

H-GCU availability includes scheduled coverage, sickness and leave.

M-GCU availability includes machine uptime and external dependency uptime. When an M-GCU may escalate, timely qualified ACE availability is measured separately.

Expected GCU per period = attempts per period × expected valid GCU per attempt

This allows human and machine production paths to be compared without pretending that they have the same resource scale.

ACE inside the GCU model

ACE is not the starting unit. It is a scarce input used where qualified human attention is required at the right moment.

Canonical ACE measures are:

- probability that ACE is required per attempt;
- ACE-seconds when required;
- probability that qualified, contextual and authorized ACE arrives before the deadline;
- ACE response latency;
- ACE-seconds per valid GCU;
- GCU blocked because timely ACE was unavailable.

Routine machine completions should consume no ACE when current bounded authority and deterministic admissibility are sufficient. Ambiguity, material exceptions, stale authority or high consequence produce ESCALATE and reserve the effect until valid ACE clears, modifies or denies it.

Governed consequence path

Machine-produced consequential effects must pass the declared execution boundary. Attempts that bypass the required boundary are inadmissible and count as zero completed outcomes.

Derived organization

For a portfolio of required GCU classes, select the fastest or lowest-cost candidate configuration only after quality, authority, consequence, evidence and expected-loss gates pass.

Then aggregate the selected configurations to derive required human ACE capacity, human disciplines, independent duties, synthetic capacity, coordination load, availability reserves and federation boundaries.

Human headcount is therefore an output of governed completion demand, not an input to the production model.

Current executable proof

Run:

```
npm run gcu:simulate
```

The synthetic example compares an H-GCU, a governed M-GCU with ACE escalation, and a direct M-GCU. The direct path is rejected before optimization. The safe optimizer selects only among configurations that satisfy the completion contract.

Synthetic parameters demonstrate the contract and invariants. They do not establish real prices, costs, availability, task composition or optimal organization size.