

GCU Industry Portfolio Simulation

Status: public working report Author: Njål Gaute Solland Canonical base:
63b4807f34c8611b461396b24c76414b6acebe6a
Date: 2026-08-27

Claim under test

Industry organization is derived from the portfolio of governed completions it must produce. The same GCU contract and governance invariants apply to every industry; only task mix, consequence, physicality, resource use and cost assumptions change.

Initial industry profiles are synthetic hypotheses. They are not empirical estimates, production policy, authorization or claims about optimal headcount.

Simulation boundary

The simulator:

1. defines required GCU demand per task class and period;
2. rejects configurations that fail quality, authority, evidence, loss, latency or governed-effect requirements;
3. selects the declared safe objective within each task class;
4. aggregates cost, latency, ACE, human work, machine work and disciplinary composition;
5. exposes unmet demand and rejected paths.

The simulator must never convert an industry prior into runtime authority. Real conclusions require external research and observed operating data to calibrate the profiles.

Falsification

The thesis weakens when a supposedly superior configuration cannot satisfy the same completion contract, hides human work, exceeds available ACE, omits physical or independent duties, or loses its advantage under credible parameter ranges.

Human output comes after GCU selection

Selected configurations produce two distinct human-capacity outputs:

- `governing_core`: authority, ACE, control, expert accountability and other roles whose mutual coordination constrains the integrated organization;
- `operational_pool`: scalable relationship or physical execution capacity that can be organized in cells, sites, shifts or external delivery networks.

The exact potential-pair count is applied only to the minimum governing core. Applying $n(n-1)/2$ to an entire nursing, retail, construction or field workforce would incorrectly assume all-to-all coordination.

Duty floors remain visible even when average work is below one FTE. Examples are accountable authority, independent control, safety responsibility and maintenance capability. A machine-produced portfolio can therefore require a non-zero multidisciplinary human core.

Executable synthetic baseline

Run all profiles:

```
npm run gcu:industry
```

Run one profile or change demand:

```
npm run gcu:industry -- --industry healthcare --scale 2  
npm run gcu:industry -- --list
```

The built-in profiles use six common task archetypes: routine information, expert analysis, consequential decision, relationship completion, physical execution and creative design. Every industry changes their weights, disciplines and physical-automation prior, while the GCU gates remain identical.

The first normalized run uses 10,000 required GCU in a 30-day period. These are the deterministic outputs of the synthetic v0.1 assumptions:

Industry	M-GCU share	Governing core seats	Operational seats	Total human seats	Core band
Software and SaaS	78%	15	7	22	11-30
Professional services	68%	15	12	27	11-30
Banking and finance	83%	11	7	18	11-30

Industry	M-GCU share	Governing core seats	Operational seats	Total human seats	Core band
Healthcare	60%	8	37	45	1-10
Manufacturing	90%	13	4	17	11-30
Construction	40%	13	59	72	11-30
Logistics and transport	90%	7	6	13	1-10
Retail and hospitality	43%	7	45	52	1-10
Energy and utilities	55%	14	45	59	11-30
Education	57%	14	22	36	11-30

In this baseline, total human seats exceed 50 in construction, retail/hospitality and energy/utilities, but the governing core remains between 7 and 15. That is a generated hypothesis caused by the declared task and capacity assumptions. It is not evidence that real industries have those automation shares, seat counts or core sizes.

Calibration chain

An industry profile becomes calibrated only when every material parameter has a traceable chain:

industry -> occupation mix -> task frequency -> completion contract -> candidate performance -> governed outcome evidence

The initial source classes are:

- O*NET tasks and work activities for detailed occupational task statements and work-activity structure;
- ESCO occupation-skill matrices and Eurostat occupation-by-industry data for European occupation and industry mapping;
- Statistics Norway classifications and earnings data for Norwegian SN2007 industry, STYRK-2008 occupation and labor cost inputs;
- the ILO refined task-level exposure index as an exposure prior, never as proof of valid completion;
- the Anthropic Economic Index as observed AI-use evidence, never as outcome quality or safe automation evidence;
- GDPval as a multi-occupation capability benchmark, with production latency, reliability, authority and integration measured separately;
- causal or field evidence such as Generative AI at Work, Navigating the Jagged Technological Frontier and the METR developer productivity RCT to calibrate only the task populations they actually study.

Exposure, observed use, benchmark capability and causal productivity are separate evidence classes. None may be substituted for governed completion probability.

Proof and disproof conditions

The executable proof is narrow:

- a direct consequence-bearing machine path is rejected before optimization;
- missing capability or unnamed ACE produces an inadmissible candidate or unmet demand;
- human and machine work, ACE, duty floors and discipline composition remain visible;
- demand scaling changes variable capacity while independent-duty floors remain fixed;
- governing core and operational labor are computed separately.

The organizational thesis is weakened if calibrated profiles repeatedly require integrated governing cores above 50, if operational pools cannot be decomposed without losing completion quality, or if machine candidates lose their advantage after real retry, integration, incident, ACE and evidence costs are included.