

The Denied Column: Military AI for Degraded and Denied Operations

By: [Ben Van Roo](#), CEO and Co-Founder

At Project Convergence Capstone 6 this month you could stand in one tent and watch the entire history of command and control run at the same time. On one table, acetate over a map sheet and a grease pencil. On another, a cloud-hosted picture that assumes the network holds. At the edge, in the kit we brought forward, agents working real tasks against hard limits on compute, model size, and power.

Nobody staged that as a demonstration. It is just what a division looks like mid-transition. But the span, grease pencil to agent-coded autonomy, is the most interesting design problem in the field right now, and almost nobody is treating it as one.


These are the author's observations from the field, not official lessons learned from Project Convergence.

Task Ambiguity Sets the Compute Bill

Every antenna a team lights up is a targeting cue

Three things ran side by side. Some processes stayed on acetate, deliberately, because a grease pencil keeps working when everything else stops, and more honestly, because it is what has been used for decades. Some processes ran on cloud-dependent systems, fast and rich when the link held and blank when it did not. Legion Intelligence was there to push agentic AI into the denied and disconnected cases the cloud does not reach, and inside our deployment the activities were finer grained. A status rollup a small model handled without complaint. A planning problem that wanted more reasoning than the form factor could give, and had to be shaped, bounded, or handed back. Some GPUs that were state of the art just a few years ago did not work well for some use cases.

From acetate to edge in one tent is one of the more interesting areas of AI and defense work right now. The constraints are real, they are legible, and nobody has drawn the axis they sit on. So here it is.



Acetate is not off this chart; it sits at the floor, the flat-variance case that performs the same everywhere because its only dependency is a trained soldier. One line on the chart matters more than the rungs. To the left of it, tasks survive a blackout because the compute they need fits at the edge. To the right, they do not, unless somebody collapsed their ambiguity in advance.

That line also sits at a different height for everyone. A corps has feeds, staff, reach-back, and power to spare. A team on an objective has none of it, and every antenna it lights up is a targeting cue. Which is why the loudest arguments here talk past each other. Put AI everywhere is a corps headquarters talking. AI will not survive contact is a team talking. Both are right where they live and wrong one level away.

Why We Climbed the Ladder

Climbing is not a temptation to be resisted. It is where the tempo is.

[Parasuraman, Sheridan, and Wickens laid the framework out in 2000](#) for flight decks, and it still fits. Automation applies separately to information acquisition, analysis, decision selection, and action, and each can be automated to a different degree. The question was never whether to automate but which function, to what level. Sensors now feed collection managers directly and models triage the take. Models draft summaries and run contradiction checks a watch floor would have caught at hour six or never.

Models wargame courses of action faster than a cell can whiteboard one. In our own fielded work we have measured order of magnitude reductions in production timelines.

The pitch that carried this into the building was flexibility: throw the data at a large model, ask it anything, escape the rigid application logic the last generation required. That pitch is largely true. It is also, quietly, a cloud pitch, assuming the transport layer is up and the compute behind it is elastic. The architecture review shows the operating picture, the model endpoint, the data sync, and the reach-back as four line items with four owners. In the spectrum they are one thing on one trunk. A formation that believed it held four independent capabilities finds out under jamming that it held one. Reaching for the cloud is also an emission, and a command post that needs continuous reach-back in order to think is a command post with a signature.

The Adversary Sets the Ceiling

Denial is a choice someone makes, not weather.

Most writing on degraded operations treats connectivity loss as an environmental condition, something that happens to a formation the way weather does. That framing is comfortable and wrong. Denial is an adversary decision, taken at a time selected to maximize our disadvantage. He decides where the ceiling falls and when it drops. So he contests the link precisely when tempo matters, at the transition, the passage of lines, the moment the plan has to change, which makes our worst performance and his best opportunity the same event by construction. He also has an interest in our climbing higher in garrison, because the higher we climb the further we fall when he drops the ceiling. Uncontested connectivity during competition is not a gift. It is the setup.

Least discussed, the attack surface moves. Once a model is the analyst of record, the cheapest attack is no longer against the link at all but against what the model reads. Jamming produces a blank screen and the staff knows it is blank. Poisoned reporting produces a confident answer and the staff knows nothing. Of the two the second is far harder, and the military literature on it has barely started.

The obvious objection is that we are building our way out of this. Proliferated LEO, mesh networking, and the next generation of command and control are all aimed at exactly this problem, and they will raise the floor. They will not remove the adversary's choice of timing, they will not resolve the tension between reach-back and emissions control, and they do nothing about compute, which is a separate constraint from transport. A resilient link still does not put a frontier model in a vehicle. The ceiling rises. It does not disappear.

The Wrong Statistic

The adversary does not sample from your average.

Nearly every evaluation of AI in the command post reports a mean. Average cycle time down, average throughput up. Those numbers are real, and they are the wrong statistic, because the adversary does not sample from your average. Take a planning task a manual staff does in one unit of time, and a campaign where transport is contested twenty percent of the time.

POSTURE	CONNECTED (80%)	DENIED (20%)	CAMPAIGN AVG
Manual staff, no AI	1.0x	1.0x	1.0x
Cloud dependent AI staff	10.0x	0.4x	8.1x
Cloud AI, atrophied staff	10.0x	0.2x	8.0x
Edge harness, bounded	6.0x	3.5x	5.5x

Illustrative, not measured. Placeholders for values every program office could estimate and none currently does.

The campaign averages look close, and that closeness is the trap. Averaging assumes the denied fraction lands randomly across tasks of equal consequence. It does not. That fraction is adversary selected, concentrated on the decisions that matter most, so the column that decides the war is the middle one. The architecture that loses on the brochure wins the campaign.

It falls below manual rather than degrading onto it for the reason [Bainbridge gave in 1983](#). Automate the routine work and the operator stops practicing it, so the takeover arrives exactly when his skills are lowest. Aviation spent two decades and several hundred fatalities learning this. The staff has no check ride for planning.

Moltke Got There First

A nineteenth century staff problem and a 2026 agent orchestration problem have the same answer.

The one variable left to us is where a task sits on the ladder, and that turns out to be movable. A frontier model planning while a small, fast model executed matched the

quality of frontier models throughout, at a fraction of the cost: hundreds of dollars of cheap workers against thousands to tens of thousands of frontier workers, same result. Few moments in a hard task need frontier intelligence, and once a planner has collapsed the ambiguity into explicit instruction, a much weaker model can carry it.

The best part is that this is not new. It is commander's intent with a specified scheme of maneuver, and [mission command exists because Moltke's staff worked out that orders have to survive losing contact with the man who wrote them](#). What is new is that the artifact now has to be executable by a machine at the far end, which raises the specification bar and hands the planning cell a hard new product to own.

The Column Nobody Has Filled In

Capability is a property of the model, the harness, and the context together.

Moving a task down the ladder assumes you know which rung it started on, and for military work nobody does. Both frontier labs have already built the instrument. [OpenAI's GDPval](#) took 1,320 real deliverables from professionals averaging fourteen years of experience across 44 occupations and graded model output blind against expert output. The best model at release won or tied on roughly half, and a stronger model beat that within weeks, so the ladder shifts every time a model ships. The structure is what lasts: the same task runs across different models and scaffolds, so reasoning effort, task context, and scaffolding can each be scored separately. [Anthropic's Economic Index](#) adds the field measurement, and the tasks that automate cleanly are the bounded ones while the tasks that stay in human hands are where judgment still lives.

None of this exists for military staff work. No benchmark scores a warning order, a collection deck, or a sustainment forecast, and none runs those tasks under connected, degraded, and denied conditions, so every claim in this debate, mine included, rests on vendor demonstration. What Project Convergence measured about the network, a standing body should measure about the tasks: which rung each staff product sits on, refreshed every time a model ships, graded under denial, because a rung assignment that holds on fiber and collapses under jamming is worse than none.

Automation Authority

Permit, not enable. The gap between them is where the risk pools.

The question enthusiasm and skepticism both dodge is which decisions we would permit a machine to make, at which rung, at which echelon, under which conditions. The technical frontier of what can be automated moves faster than the doctrinal frontier

of what should be. The instinct is to treat this as unprecedented, and that instinct understates what the profession already holds. Weapons control status moves between free, tight, and hold as conditions change. Fire support coordination measures bound where effects go without further clearance. [DoD Directive 3000.09](#) requires that authority be specified before employment. Nobody has applied that apparatus to cognitive work.

DECISION TYPE	CONNECTED	DEGRADED	DENIED
Report production and rollups	Automate	Automate	Automate
Collection tasking	Automate	Supervised	Supervised, pre-bounded
COA generation and wargaming	Supervised	Supervised	Supervised, pre-bounded
Target nomination	Supervised	Supervised	Human, pre-approved criteria
Weapons release	Human	Human	Human

A straw man for battalion and below, offered as a position rather than an observation. The denied column is where the argument should happen. Fill that column in now, in garrison, with time to argue about it. Leave it blank and it still gets answered, by whoever is standing in the TOC when the screens go dark, under the worst conditions for judgment the profession knows.

Price the Bill Before Doctrine Hardens

The window where these trades are still visible is narrow

The compression is going to be bought, because the mean gains are too large to refuse and refusing them cedes tempo in every fight where the network holds, which is most days of most operations. The argument here is narrower: buy it with the bill priced.

Stand up a standing capability-evaluation body first, because everything else depends on it. Score which model and which harness completes which staff task, refresh it every time a model ships, run it on real products under connected, degraded, and denied conditions, and report the denied column separately. Project Convergence proved the Army will mount an experiment of enormous scale to test a network. The same seriousness applied to the tasks that ride the network is a rounding error by comparison.

Write a PACE plan for cognition. Every signal officer maintains primary, alternate, contingency, and emergency for communications, briefed before crossing the line of departure, and there is no equivalent for thinking. Specify, for each major staff process, the primary means, the alternate when reach-back thins, and the emergency version on a map board.

Separate planning from execution deliberately, collapsing ambiguity where compute is ample and carrying forward explicit, bounded, executable instruction. And certify manual proficiency on a schedule, because untested capabilities are assumed present until the day they are needed.

Then fill in the denied column. Take a position, publish it, and let it be argued with. Extending weapons control status and engagement criteria to cognitive work is an afternoon of argument among people who already speak this language, and it is a better afternoon than the alternative. We are in the narrow window where these trades are visible, the tools are new enough to shape, and the doctrine has not hardened around a bad answer. That is the fun of it. The acetate is still on the table at Project Convergence for a reason, and the open question is everything we are willing to run above it, and how far it falls when the ceiling drops.

None of these questions are hypothetical for Legion. We field agents under exactly these constraints today, from cloud to the disconnected edge, and the sooner the profession argues the denied column, the better every architecture, ours included, will have to be.

References

Lin, W. Agent Swarms and the New Model Economics. Cursor, July 2026.

U.S. Army, Project Convergence Capstone 6, Fort Irwin, July 2026.

Continue the series

Follow Legion Intelligence on LinkedIn: [linkedin.com/company/legion-intel](https://www.linkedin.com/company/legion-intel)

Explore all Command Papers: [legionintel.com/command-papers](https://www.legionintel.com/command-papers)

Ready to talk: <https://www.legionintel.com/contact>