

AETHERMOORE RESEARCH PACKET

Drone Autonomy Training Patterns for SCBE Agent Training

Mapping public autonomy-program patterns onto coding-agent gates, simulations, and promotion tiers.

autonomy

swarms

agents

simulation

training

assurance

Abstract. Government autonomy work tends to progress through simulation, test ranges, trust calibration, runtime assurance, live demonstrations, and repair loops. That structure maps cleanly onto SCBE agent training and receipt-backed promotion.

Claim boundary. Use as an analogy and training-loop design pattern for software agents, not as a claim of drone operation or certification.

Status: Public-pattern synthesis; not a physical drone capability claim.

Research grade: Research Brief

Review model: Public program pattern synthesis with software-agent mapping.

[PDF](#) | [Source note](#)

Source Text

Drone Autonomy Training Patterns for SCBE Agent Training

Date: 2026-05-02

Purpose: capture public, official patterns from DARPA, NASA, Army, and Air Force autonomy work that can inform SCBE's coding-agent and GeoSeal training loops without claiming physical drone capability.

Executive Takeaway

Government drone and autonomy training does not look like one big model run. It looks like a staged progression:

1. high-fidelity simulation,
2. instrumented test ranges,
3. human trust calibration,
4. runtime assurance,
5. live or field-like demonstrations,
6. continuous repair from failures.

That maps cleanly to the SCBE "boss fight" model: run the agent against a frozen gate, record failure mechanics, convert failures into analog repair experience, retry, and only promote when the same gate passes.

Source Patterns

DARPA OFFSET: swarm tactics as a development ecosystem

Source: <https://www.darpa.mil/research/programs/offensive-swarm-enabled-tactics>

DARPA OFFSET frames swarm autonomy around hundreds of small unmanned aircraft or ground systems, human-swarm teaming, fast tactic generation, evaluation, and integration into field operations. The useful pattern for SCBE is not "more drones"; it is the open tactics ecosystem:

- generate candidate tactics,
- evaluate effectiveness,
- preserve the best tactics,
- expose humans to swarm state through a command interface.

SCBE mapping: agent plans should become reusable tactics with score histories, not one-off chat outputs.

DARPA REMA: adapter interface plus mission autonomy

Source: <https://www.darpa.mil/research/programs/rema-rapid-experimental-missionized-autonomy>

REMA focuses on adding autonomy subsystems to existing commercial and stock military drones, with two technical areas: a drone-autonomy adapter interface and mission-specific autonomy software. DARPA explicitly describes repeated "accelerating spirals" of development.

SCBE mapping: GeoSeal should treat tools like adapter surfaces. The coding agent does not need to own every tool internally; it needs a stable adapter interface, mission-specific policies, and repeatable spiral upgrades.

DARPA ACE and AIR: hierarchy, trust, and sim-to-live progression

Sources:

- <https://www.darpa.mil/research/programs/air-combat-evolution>
- <https://www.darpa.mil/news/2024/ace-ai-aerospace>
- <https://www.darpa.mil/research/programs/artificial-intelligence-reinforcements>

ACE separates higher-level human cognitive functions from lower-level autonomous maneuver/tactics. It also names trust calibration as a core challenge. AIR extends this into multi-ship beyond-visual-range missions, emphasizing uncertainty, open-world adaptation, human feedback, modeling, simulation, and live missions.

SCBE mapping:

- Human/user sets mission intent and constraints.
- GeoSeal handles lower-level tool routing, budget checks, and fallback.
- The training system logs where trust should increase or decrease.
- Promotion requires both task success and trust evidence.

DARPA Assured Autonomy: continual assurance for learning-enabled systems

Source: <https://www.darpa.mil/research/programs/assured-autonomy>

DARPA Assured Autonomy targets continual assurance of learning-enabled cyber-physical systems. The key pattern is that assurance is not only design-time certification; it is monitored, updated, and evaluated during operation as the system and environment evolve.

SCBE mapping: a coding agent should not be trusted because it once passed a test. It should carry live assurance state: last gate, failure class, allowed tools, denied tools, and current promotion tier.

NASA simulation and multi-UAV visualization

Sources:

- <https://www.nasa.gov/reference/jsc-simulation-modeling/>
- <https://software.nasa.gov/software/LAR-19641-1>

NASA simulation uses high-fidelity simulation hosts, virtual computers running flight software, partner interfaces, training devices, and human-in-the-loop simulation. NASA WebGS supports multi-UAV test design, automated simulation execution, repeatable/adjustable flight planning, live flights interacting with simulated vehicles, remote monitoring, and multi-user interaction.

SCBE mapping:

- Treat evals as repeatable simulation worlds.
- Keep live systems and simulated systems able to interact through the same message format.
- Make test runs replayable and inspectable by humans.

NASA Armstrong: AI-enabled UAS and collision/landing safety

Source: <https://www.nasa.gov/centers-and-facilities/armstrong/autonomous-systems/>

NASA Armstrong work includes AI-enabled small UAS for disaster/crash-site search, object detection and geo-tagging, improved ground collision avoidance, and autonomous emergency landing. Important training pattern: safety features are paired with maps, reachable sets, real-time updates, and vehicle-specific models.

SCBE mapping:

- Every agent run should track reachable states: what actions are allowed next, what rollback paths exist, and what terrain/map of the repo it is using.
- GeoSeal should expose minimap-like state for code: changed files, tests run, blocked tools, current lane, and fallback route.

Army UAS simulation and swarm test ranges

Sources:

- <https://www.army.mil/article/187989>
- https://www.army.mil/article/236381/army_researchers_find_new_ways_to_test_swarming_drones
- <https://www.tradoc.army.mil/2025/08/19/u-s-army-aviation-center-of-excellence-launches-unmanned-advanced-lethality-course-to-equip-soldiers-for-future-warfare/>

The Army uses accredited simulation systems for UAS crew readiness and keeps simulator software aligned with tactical system upgrades. Army swarm testing also emphasizes large outdoor instrumented ranges, camera navigation, RF communication, heterogeneous swarms, ground/aerial interactions, and human-agent teaming. The newer UAS training course pattern is classroom plus simulator hours before live flight.

SCBE mapping:

- Use simulator-first agent training before granting broader filesystem or network powers.
- Record "simulator hours" as EXP: number of successful gated tasks, failure repairs, and no-regression repeats.
- Do not let a model jump to live actions until it clears the simulation lane.

Air Force Skyborg: open systems, modeling/simulation, and runtime assurance

Source: <https://www.af.mil/News/Features/Article/1796930/skyborg-program-seeks-industry-input-for-artificial-intelligence-initiative/>

Skyborg emphasizes low-cost attritable unmanned aircraft, open systems architecture, modeling and simulation, modularity, autonomous teaming, runtime assurance, and building trust as the system evolves. The Air Force source explicitly connects gaming AI progress to the question of moving AI safely into real flight.

SCBE mapping:

- Use game-like worlds as useful training environments, but require runtime assurance before giving real powers.
- Make tool interfaces modular so small models can learn one subsystem at a time.
- Track trust as a progressive score, not a binary claim.

SCBE Implementation Pattern

EXP Unit

An EXP unit is a verified micro-skill record:

```
{
  "task_id": "stage6_unseen_hex_trace",
  "failure_kind": "byte_hex_compute_trace",
  "missing_required_marker": "compute",
  "repair_action": "analog_repair_rows",
  "gate_retested": true,
  "same_gate_improved": true,
  "heldout_leak": false
}
```

Training Ladder

1. Simulator Lane: frozen prompts, fake filesystem, no live side effects.
2. Instrumented Range: real repo read access, dry-run tool calls, full telemetry.
3. Supervised Live Lane: small patches only, human approval required.
4. Trusted Routine Lane: repeated task family with bounded write scope.
5. Mission Lane: paired agents, switchboard arbitration, full runbook and rollback.

Boss Loop

```
train candidate
run frozen gate
extract failure mechanics
generate analog repair data
apply constrained decoding or preference repair if SFT saturates
retry same gate
promote only after pass
```

What To Build Next

- Add a GeoSeal "simulation world" runner for coding tasks that exposes a minimap: files, tests, tools, budget, failure class, and rollback route.
- Extend boss-retry-plan into a repair-row generator that writes analog SFT/DPO rows from failure kinds without copying held-out prompt text.
- Add a paired-agent flight rule: one agent proposes actions, one agent audits route/budget/tool permission before execution.
- Score EXP as cumulative evidence across runs: pass rate, no-regression count, repair lift, tool discipline, and context continuity.

Generated from research/drone_autonomy_training_patterns_2026-05-02.md by
scripts/research/build_research_library.py.