TheaterForge LLC
CAPABILITY STATEMENT

Rev. 2026-10-15 · v09
West Jordan, Utah
leland@theaterforge.com · theaterforge.com
435-272-9294

Deterministic multi-domain simulation and trade-study software — and the reproducible foundation AI/ML needs to be trained and tested against.

COMPANY DATA

LEGAL NAME
TheaterForge LLC

ENTITY
Utah LLC · No. 14733139-0160

ESTABLISHED
2 September 2026

BUSINESS SIZE
Small business

UEI
FTQ7GVNLKDX5

CAGE CODE
257P5

SAM.GOV
Active

TECH READINESS
TRL 4

PLACE OF PERFORMANCE
United States

NAICS

541715
R&D in the Physical, Engineering, and Life Sciences (primary)
541511
Custom Computer Programming Services
541512
Computer Systems Design Services
541330
Engineering Services

PSC

AC12
R&D Defense System: Applied Research
DA01
IT and Telecom: Business Application Software Development

POINT OF CONTACT

Leland Doss
Sole Member and Manager / Principal Investigator
leland@theaterforge.com · 435-272-9294

CORE COMPETENCIES

Deterministic simulation engineering. Multi-domain theater models — space, air, maritime, ground — executing as reproducible experimental instruments rather than visualizations. Identical inputs produce byte-identical outputs, verified by hashing.

Automated trade-space exploration. Design-of-experiments harness screening large parameter spaces to identify driving factors: full factorial, saturated two-level screening, and Latin hypercube designs, with main-effects analysis and pass/fail verdicts against stated requirements.

AI/ML training and evaluation environments. PLANNED The engine is designed to be wrapped as a Gym-compatible reinforcement-learning environment. Agents act at decision points such as sensor tasking; rewards come from the engine's existing measures; training episodes are drawn across the architecture trade space by Latin hypercube sampling. AI-driven tasking can then be trained and scored against fixed doctrine on identical, reproducible terms.

Browser-native scientific computing. Physics and analysis execute client-side: no install, no server-side compute, no runtime license, no administrative rights on the host workstation.

Analyst-facing decision support. 3D geospatial and 2D tactical visualization, scenario authoring, and self-contained exportable reporting.

DIFFERENTIATORS

Zero-footprint deployment. Runs in a standard browser on the machine the analyst already has — no installation, no backend to accredit — lowering the barrier to fielding analysis inside restricted enclaves.

Offline by default. Ships with its network gate closed and runs fully offline: no telemetry, no scenario upload, no external analysis service. Optional basemap imagery is the only network-capable feature and can be disabled or served locally.

Findings a nominal run reports as success. In a layered air and missile defense study, a raid of eight was fully intercepted with two tenths of a second of worst-case margin — passing the leak-rate requirement and failing the timing requirement in the same run. In a multi-orbit missile-warning study, operator decision delay dominated every architecture parameter by roughly six to one. Both are committed, hashed, and regenerable by a third party. Both are results about the model, not any fielded architecture: parameter values are notional and the sensing models are not yet validated.

Safe to learn against. A measure whose event never occurs returns null, never a sentinel value, and requirement checks reject null — so an undetected threat can never score as comfortable margin, and an agent trained against the engine cannot learn to win by letting threats through.

Provable reproducibility. Every probabilistic outcome is keyed to the identity of the individual event, not a global sequence, so adding an element does not perturb unrelated outcomes — differences between design points are attributable to the design, not to noise.

Screening scale, verified. Resolves all 30 main effects of a 30-parameter trade space in 32 runs, with two-factor interactions aliased by construction and the aliasing structure reported; the design is verified by automated test to be saturated, pairwise orthogonal, balanced, and to recover a known effect. Test output is committed and hashed.

Throughput. Approximately 1,000–3,000× real time on a single laptop core, scenario-dependent.

Engineering discipline evidenced, not asserted. 379 automated tests across 36 suites at tag v0.3.0, static type checking, generated software bill of materials, continuous integration.

Honest verification posture. The system is verified; orbital propagation including J2 secular perturbation has been cross-checked against NASA GMAT, with the residual disagreement documented; no component is claimed as validated. Byte-identical reproduction is claimed for the same code on the same runtime. Capability and limitation are stated together.

PAST PERFORMANCE AND PRINCIPAL EXPERIENCE

TheaterForge LLC has no federal prime or subcontract performance to date. The following is the relevant experience of the principal.

TheaterForge simulation engine — sole architect and developer. Multi-domain theater engine with integrated trade-study harness; ballistic missile defense, layered air and missile defense, multi-orbit missile warning, and naval saturation scenario classes.

Sole developer for his department at a defense services firm, 2023–2026. Simulation and data-visualization software for a U.S. Space Force customer: multi-domain sensor-data ingestion and a satellite-data visualization platform used by defense decision-makers.

Applied AI and reinforcement learning. Reinforcement-learning environment built and shipped (Gymnasium, Stable Baselines3, DQN, curriculum learning, reward shaping); AgentSentinel, an AI-agent monitoring and observability platform.

Security and data posture. No classified data has been used in or is present in the system. Performance figures are compiled from unclassified open-source references with modeling judgment and structured as swappable configuration. TheaterForge LLC seeks SBIR/STTR opportunities, AI/ML and modeling-and-simulation roles on SBIR/STTR and prime teams, and prototype development under OTA or similar instruments. **Approved for public release:** contains no classified, controlled unclassified, or export-controlled technical information.