



Defense as a Service (DAAS)

RE-ARCHITECTING THE DEFENSE INDUSTRIAL BASE FOR CONTINUOUS COMPETITION

Chad Steelberg
Founder & Chief Executive Officer, Tiberius Aerospace

A Research Paper
Prepared for policy makers, defense institutions, and industry leaders

Tiberius Aerospace
Lakestar

January 2026

The views expressed in this paper are those of the author alone and do not necessarily reflect the views of Tiberius Aerospace, Lakestar, or any affiliated organizations.
This paper is intended to contribute to informed discussion on defense industrial policy and does not constitute official policy guidance.

ABSTRACT

Western defense acquisition systems were designed for an era defined by slow technological change, episodic conflict, and industrial mobilization led by a small number of vertically integrated prime contractors. That era has ended. Contemporary conflict is characterized by near-peer competition, rapid adversarial adaptation, compressed innovation cycles, and increasingly constrained public budgets. Under these conditions, the traditional program-centric defense model—built around single-buyer procurement, large upfront Research, Development, Test, and Evaluation (RDTE) investments, and long production tails—has become structurally misaligned with battlefield realities.

This paper introduces Defense as a Service (DaaS) as a market architecture for defense capability development and sustainment. DaaS does not seek to reform acquisition regulations, export controls, or sovereign authorities. Instead, it assumes these features are durable characteristics of mature nation-states and restructures the economic and industrial incentives operating beneath them. By distributing RDTE across multiple allied Ministries of Defense, integrating private capital to absorb early technical and timing risk, licensing rather than transferring intellectual property, and funding continuous innovation through annual program participation, DaaS converts defense procurement from episodic, milestone-driven programs into a continuously competitive, market-disciplined ecosystem.

The DaaS model produces several structural effects: materially lower switching costs by design; competition at the mission level rather than the program level; sustained innovation pressure through continuous integration and deployment; federated manufacturing that improves both wartime surge capacity and peacetime resilience; and transparent market signals that discipline both suppliers and investors. Importantly, DaaS preserves sovereign control over operational use and export decisions while strengthening economic and industrial ties across free-market alliances.

This paper develops the theoretical foundations of DaaS, situates it within the historical evolution of the defense industrial base, quantifies its advantages relative to legacy models, anticipates institutional critiques, and examines current implementations and boundary conditions. The analysis argues that Defense as a Service represents not a novel contracting construct, but a necessary structural evolution if democratic states are to sustain military and industrial advantage in an era of continuous competition.



TABLE OF CONTENTS

Executive Summary

Key Definitions

Part I — The Problem

1. Introduction: The End of Program-Centric Defense

- Structural—not managerial—failure of legacy acquisition
- Why incremental reform repeatedly fails
- Framing DaaS as inevitability, not ideology

2. Historical Evolution: Consolidation, Monopsony, and Incentive Drift

- Post-WWII industrial diversity
- Consolidation into a small set of primes
- Cost-plus dynamics and program lock-in
- Competition moving from lifecycle to award

[Figure 1: Historical consolidation of defense primes vs. industrial diversity]

3. Why Reform Fails: Bureaucracy and Export Controls as Fixed Variables

- Acquisition bureaucracy as a stable equilibrium
- Export controls (ITAR and equivalents) as durable political artifacts
- Why reform-centric strategies are structurally fragile

4. What Changed: Threat Tempo, Countermeasure Cycles, and Industrial Fragility

- Near-peer competition and adaptation speed
- Collapse of RDTE timelines
- Battlefield feedback as a continuous signal

[Figure 2: Weapon effectiveness decay vs. adversary adaptation rate]

5. The Industrial Base in Peace and War: Surge Failure and Downturn Collapse

- Inability to surge production
- Program-captive manufacturing fragility
- Cost of industrial reconstitution

[Table 1: Legacy vs. DaaS industrial resilience comparison]

Part II — The DaaS Market Architecture

6. DaaS First Principles: From Products to Continuous Services

- Capabilities as living systems
- Continuous RDTE as a permanent function
- Markets beneath bureaucracy

7. The Role of Private Capital: Breaking the Monopsony

- Capital as a second buyer of development risk
- Dilution as discipline
- Capital-enabled leapfrogging
- Political legitimacy and sovereign safeguards

[Figure 3: Capital flow comparison — legacy vs. DaaS]

8. Distributed RDTE and TRL-Based Entry

- Coalition-funded innovation
- TRL-based pricing of uncertainty
- Early vs. late entry incentives
- Pre-commitment market signaling

[Table 2: TRL stages vs. risk, pricing, and MOD exposure]

9. Switching Costs as an Arithmetic Constraint

- Structural reduction of switching costs
- Switching Cost $\approx$ MOD Share of RDTE $\times$ Time-in-Program
- Contrast with legacy sunk-cost lock-in

10. Rights, Exit, and Rejoin Mechanics

- Stockpile retention and use rights
- Maintenance and sustainment after exit
- Upgrade and version access
- Rejoin pricing logic
- Wartime access and emergency clauses

Part III — Governance, Certification, and Security

11. Coalition Interface Governance vs. Coalition Engineering Governance

- Interface standards vs. co-design
- Market-driven convergence
- Avoiding committee paralysis

12. Certification, QA, and Version Control Under Continuous Change

- Safety case ownership
- Version taxonomy (patch / minor / major)
- Backward compatibility policies
- Audit cadence and cost allocation

[Figure 4: Versioning and certification pipeline]

13. Security and IP Protection in a Controlled-Modularity Model

- Core vs. peripheral components
- Preventing IP leakage and cloning
- Cryptographic attestation and provenance
- Export-controlled configurations

14. Liability and Accountability in Federated Manufacturing

- Prime vs. manufacturer vs. MOD liability
- Warranty and defect regimes
- Recall authority
- Failure investigation governance

Part IV — Competition and Performance

15. Competition at the Mission Level: The Venn Diagram Model

- Partial mission overlap
- Portfolio allocation by MODs
- Contestability without duplication

[Figure 5: Mission overlap Venn diagrams]

16. Cost-Effective Lethality Score (CELS) as the Clearing Metric

- Definition and rationale
- Environment and adversary adaptation
- Manufacturing rate and sustainment effects

[Table 3: Example mission-level CELS comparison]

17. From Requirements to Outcomes

- Arbitrary requirements vs. non-negotiable constraints
- Requirements capture and vendor shaping
- Outcome-driven competition

18. Continuous RDTE, CICD Culture, and Capital-Enabled Leapfrogging

- Baseline innovation via annual participation
- Burst-mode overinvestment
- Manufacturing share as the scoreboard

[Figure 6: Milestone vs. continuous innovation cadence]

Part V — Manufacturing and Industrial Resilience

19. Federated Manufacturing and Licensing Economics

- Unbundling design and production
- Cost-plus royalty logic
- Margin stability with falling unit cost
- Territorial exclusivity: when it helps vs. hurts

[Figure 7: Federated manufacturing economics]

20. Manufacturer Incentives and Capacity Investment

- Capex without guaranteed volume
- Demand smoothing across programs
- Qualification pipelines and amortization

[Table 4: Manufacturer investment logic — legacy vs. DaaS]

21. Wartime Surge and Peacetime Survival

- Scaling without renegotiation
- Maintaining innovation during downturns
- Deterrence as an option structure

[Figure 8: Downturn survival curves — legacy vs. DaaS]

22. Export Controls and Compliance in a Distributed Model

- Program-level compliance
- Certification as a market gate
- Controlled openness

Part VI — Evidence, Implementation, and Limits

23. Empirical Analogs: Where This Structure Already Works

- Pharmaceutical staged-risk financing
- Semiconductor ecosystems
- Aero-engine risk-sharing
- Subscription software models

[Figure 9: DaaS-to-industry analog mapping]

24. Current Implementation: Tiberius Aerospace, Sceptre, and Invictus (Analytical)

- What is structurally new
- What must be measured
- Validation metrics

[Table 5: DaaS program validation metrics]

25. Investor Instruments, Time Horizons, and Governance Limits

- Equity vs. royalty structures
- Defense time horizons vs. venture logic
- Preventing investor capture of roadmaps

[Figure 10: Capital stack options for DaaS programs]

26. Boundaries and Failure Modes

- No mission overlap monopoly reversion
- Standards fragmentation
- Weak certification governance
- Capital withdrawal scenarios
- Modularity mismanagement

27. Strategic Implications for Alliances and Industrial Policy

- Alliances as markets
- Trade as a security multiplier
- Industrial capacity without dirigisme

28. Conclusion: From Programs to Markets



EXECUTIVE SUMMARY

Western defense acquisition faces a structural mismatch between how weapons are developed and how modern wars are fought. Traditional programs concentrate RDTE risk within a single nation, rely on infrequent competition, and impose high switching costs that lock governments into long-term commitments regardless of battlefield relevance. These dynamics were tolerable in an era of slow technological change and uncontested superiority. They are increasingly untenable today.

Defense as a Service (DaaS) addresses this mismatch by reorganizing defense innovation as a market rather than a sequence of isolated national programs. Under DaaS, multiple allied Ministries of Defense subscribe to a shared capability development effort, entering at different stages of technological maturity and paying entry prices that scale with reduced risk. Annual participation fees fund continuous RDTE, certification, and versioning, ensuring that systems evolve in response to adversary adaptation rather than stagnating after initial fielding.


Private capital plays a defined and constrained role in this architecture. It absorbs early technical and timing risk, breaking the monopoly of the single sovereign buyer without undermining sovereignty. Dilution, rather than cost growth, becomes the penalty for underperformance, creating powerful incentives for DaaS primes to attract allied participation early and to deliver measurable improvements in cost-effective lethality. Capital thus functions as a disciplining mechanism rather than a driver of rent-seeking behavior.

Competition in DaaS occurs at the mission level, not the program level. Different systems may overlap partially across mission sets, forming a portfolio of alternatives evaluated using a Cost-Effective Lethality Score (CELS) that integrates cost, effectiveness, environment, and adversary adaptation. This structure avoids wasteful duplication while preserving contestability where it matters. Switching costs are not reduced by policy or goodwill; they are mathematically bounded by distributed RDTE and time-in-program, typically well below ten percent of legacy program exposure.

Manufacturing under DaaS is federated through licensed, certified producers selected by participating MODs. This separation of design and production lowers unit costs through competition, improves surge capacity during conflict, and allows manufacturers to survive peacetime downturns by pivoting across programs and industries. Deterrence becomes an option structure rather than a sunk cost.

DaaS is not universal. It is least suited to singular strategic deterrents or inherently national systems with no plausible allied overlap. It is strongest where lethality is continuously contested, iteration speed determines relevance, coalition demand exists, and performance can be measured empirically.

The central conclusion of this paper is straightforward: Defense as a Service does not require dismantling existing institutions. It changes the incentive structure beneath them. In doing so, it offers a credible path to sustained innovation, industrial resilience, and allied advantage in an era defined by continuous competition rather than episodic procurement.



KEY DEFINITIONS

TO ENSURE PRECISION AND AVOID AMBIGUITY, THE FOLLOWING TERMS ARE USED THROUGHOUT THIS PAPER WITH SPECIFIC MEANINGS:

Defense as a Service (DaaS)

A market architecture for defense capability development in which RDTE is distributed across multiple allied governments, private capital absorbs early risk, intellectual property is licensed rather than transferred, and continuous innovation is funded through recurring program participation.

DaaS Prime

The entity that owns the core intellectual property of a DaaS program, defines system architecture and interfaces, governs certification and versioning, and licenses manufacturing and program participation to subscribers.

Subscriber / Participating MOD

A Ministry of Defense that joins a DaaS program at a defined Technology Readiness Level, pays an entry fee and annual participation fee, and receives rights to procure, manufacture, and operate the capability under agreed terms.

Research, Development, Test, and Evaluation (RDTE)

The set of activities required to conceive, design, prototype, test, certify, and mature a defense capability. In DaaS, RDTE is continuous rather than episodic.

Technology Readiness Level (TRL)

A standardized measure of technological maturity. In DaaS, TRL determines entry pricing, risk allocation, and rights associated with program participation.

Cost-Effective Lethality Score (CELS)

A comparative metric used to evaluate mission-level alternatives, incorporating weapon cost, operational effectiveness, environmental conditions, and adversary adaptation.

Continuous Integration / Continuous Deployment (CI/CD)

An innovation discipline in which improvements are developed, tested, certified, and released on an ongoing basis rather than through infrequent major upgrades.

Federated Manufacturing

A production model in which multiple certified manufacturers are licensed to produce a system, competing on cost and capacity while adhering to a common design and certification baseline.

Switching Cost



The economic exposure incurred by a subscriber when exiting a program, defined in DaaS as a function of the subscriber’s share of distributed RDTE multiplied by time-in-program.

Coalition Interface Governance

A governance approach in which participating governments influence interface standards and interoperability constraints without collectively designing or specifying the system itself.

PART I_

LAT 34.0522° N
LON 118.2437°



THE PROBLEM



PART I — THE PROBLEM

1. INTRODUCTION: THE END OF PROGRAM-CENTRIC DEFENSE

For most of the modern era, the defense enterprise has behaved as though military advantage is something a state can purchase in discrete units: a program is initiated, requirements are gathered, RDT&E is funded, a weapon system is fielded, and a production run follows. The system is then sustained and incrementally upgraded until it is replaced by the next program. This “program-centric” paradigm is not accidental. It emerged from an environment in which threats were comparatively legible, technology advanced in predictable increments, and the industrial base could be mobilized under centralized control when needed.

That environment no longer exists.

Today, the defining characteristic of military competition is continuous contestation. The battlefield is no longer a proving ground that validates a design; it is a feedback system that punishes static capabilities. A weapon is not “fielded” once. It is challenged continuously—by countermeasures, by deception, by electronic warfare, by new sensors, by industrial scaling, and by adversaries who learn fast because they must. In such a world, a state does not “buy” advantage; it sustains it through persistent iteration across hardware, software, supply chain, and tactics.

The traditional defense model was never optimized for that kind of competition. Its incentive structure favors milestone completion, budget justification, and contractual stability—features that are valuable in governance but corrosive to speed. It concentrates RDT&E risk in a single buyer, creates high switching costs that eliminate meaningful competition after award, and reinforces a cycle in which failure is recognized late and corrected slowly. The result is a structural mismatch: the defense enterprise runs on the clock of programs, while war runs on the clock of adaptation.

This paper argues that the mismatch is not primarily procedural. It cannot be solved by asking the existing system to “move faster” while preserving the same underlying architecture. The program-centric model is a stable equilibrium: it is coherent with the political economy of mature states, with the realities of security classification, and with the bureaucratic need for accountability. That is precisely why it persists—even when it underperforms. The constraints are real and durable. The error is assuming that the only solution is to fight those constraints through reform.

Defense as a Service (DaaS) is instead a structural response: a market architecture that operates beneath the bureaucracy, not against it. It treats bureaucracy and export controls as fixed variables, then rearranges incentives so that continuous competition becomes the natural state of defense capability development. DaaS is not a slogan for faster acquisition; it is a redesign of how risk is financed, how IP is controlled, how manufacturing is organized, and how performance is measured and rewarded.

Two points are central at the outset.





First, this is not an argument that the legacy system is irrational. It produced extraordinary outcomes under earlier conditions. The argument is that it is now misaligned with the dominant drivers of modern advantage: iteration speed, industrial resilience, and measurable cost-effective lethality under contested conditions.

Second, DaaS is not a call to abolish primes, discard certification, or deregulate defense. It is a model that preserves sovereign authority and hard safety constraints while reintroducing market discipline where it matters: in how capability is improved and sustained over time. It is a shift from programs to markets, and from episodic procurement to continuous competition.

This Part I establishes the problem with precision. It traces how the legacy defense industrial base evolved into a consolidated, monopsonistic equilibrium; why decades of reform have not changed its core behavior; how threat tempo and countermeasure dynamics have altered the “physics” of advantage; and why the industrial base has become fragile both in surge and in peace. These are not rhetorical claims. They are the boundary conditions that make a new architecture necessary.

2. HISTORICAL EVOLUTION: CONSOLIDATION, MONOPSONY, AND INCENTIVE DRIFT

The defense industrial base is often described as though it were a natural market. It is not. It is a market shaped by concentrated demand, regulatory barriers, export controls, classification, and a procurement apparatus that is itself a political institution. When viewed through that lens, consolidation is not surprising; it is the expected outcome.

In the decades following World War II, defense production benefited from a broad manufacturing substrate. Many companies could plausibly enter subdomains of aerospace, electronics, propulsion, and munitions because civilian and military manufacturing overlapped substantially. A wide supplier base existed not because defense was inherently competitive, but because the broader industrial economy was.

Over time, several forces drove the defense sector toward a different equilibrium.

Demand concentration. In most countries, defense demand is effectively a monopsony: one principal buyer, occasionally mediated by parliaments, ministries, and agencies, but ultimately centralized. Even where procurement is decentralized, the buyer is structurally singular. A market with one dominant customer does not behave like commercial markets. It rewards relationships, incumbency, and risk management.

Regulatory and compliance barriers. Defense is not a domain where new entrants can quickly “try, fail, and iterate” without friction. Security classification, export controls, and compliance regimes increase the fixed cost of participation. Over time, these fixed costs bias the market toward incumbents with internal compliance capacity and established contracting relationships.



Scale economics and program size. As systems became more complex and program budgets grew, the ability to carry overhead, bid large competitions, and survive long pre-revenue periods became decisive. This further narrowed the set of firms capable of acting as primes.


Buyer preference for stability. Ministries of Defense and Departments of Defense rationally prefer suppliers that appear stable and reliable. In a world where procurement errors can become national scandals, stability is often valued above innovation. Over time, this preference becomes self-reinforcing: stable primes win, primes grow, primes become “too important,” and the buyer becomes more dependent on them.

The result is a defense industrial base characterized by a small number of prime contractors and a layered supply chain. In theory, primes compete. In practice, competition is episodic and constrained because the buyer’s switching cost after award is extraordinarily high. Once a program is awarded, the state often becomes effectively locked into that supplier for a decade or more, because the program’s data rights, sustainment infrastructure, training pipeline, and certification basis become intertwined with the contractor’s proprietary systems.

This is the first major structural defect: competition occurs at the beginning, not throughout the lifecycle. Yet modern military advantage depends on lifecycle performance and iterative improvement. A model that competes once and then stops competing is structurally mismatched with a world of continuous contestation.

In parallel, incentive structures drifted.

In program-centric procurement, the contractor is primarily rewarded for meeting stated requirements and passing milestones. Where cost-plus contracting is used, the contractor may be insulated from cost discipline. Even in fixed-price regimes, the program’s complexity and the buyer’s aversion to cancellation often create renegotiation dynamics that reduce



the contractor’s downside. The most important effect is not any particular contract type; it is the aggregate fact that underperformance is often recognized late, after sunk costs and political commitments make exit difficult.

This creates a predictable behavioral pattern:

- Early optimism: requirements are ambitious, timelines are aspirational, budgets are politically palatable.
- Mid-program reality: technical debt accumulates, integration challenges appear, supply chain constraints emerge.
- Late recognition: cost and schedule overrun become unavoidable; cancellation becomes politically unthinkable; the program continues because the alternative is admitting failure.

Under these conditions, the defense enterprise can produce extraordinary individual systems and still be structurally inefficient at delivering advantage at scale. That is precisely what is happening.

[Figure 1: Consolidation timeline and competition decline]
A timeline chart showing post-WWII industrial diversity evolving into a consolidated prime landscape. Overlay (a) number of primes, (b) share of top 14 primes in procurement spend, and (c) average time from program start to fielding. The figure visually supports the claim that competition has shifted from continuous to episodic and that development cycles lengthened as consolidation increased.

3. WHY REFORM FAILS: BUREAUCRACY AND EXPORT CONTROLS AS FIXED VARIABLES

Every generation of defense leaders eventually reaches the same conclusion: acquisition is too slow, too expensive, and too rigid. And every generation attempts reform. Yet the system's core behavior persists. This persistence is often interpreted as incompetence or corruption. That interpretation is convenient—and wrong. The more accurate explanation is that defense procurement is embedded in the logic of the state.

Bureaucracy is not an accident. It is how states create accountability, allocate resources, manage risk, and enforce legitimacy. The defense domain amplifies these needs because it involves lethal force, national security, and public funds. Reforms can change processes at the margin, but they rarely alter the underlying equilibrium because the equilibrium is maintained by rational forces:

Accountability constraints. In defense, the costs of visible failure are high. Programs are therefore designed to minimize personal and political risk. This leads to reliance on incumbents, extensive documentation, and conservative decision-making.

Information asymmetry. Defense programs often involve complex technologies where the buyer must rely heavily on contractor expertise. This creates inherent asymmetry. Bureaucracies respond with process: reviews, gates, requirements, and oversight. Process reduces some risks but increases cycle time.

Sovereign control. States do not behave like firms. They must satisfy legal requirements, political legitimacy, alliance considerations, and strategic signaling. A program is not only a capability; it is a political artifact.

Export control regimes. Export controls, such as ITAR and analogous allied frameworks, have legitimate security purposes. They also have political and economic functions: controlling the diffusion of sensitive capability, managing alliances, protecting domestic industrial interests, and signaling restraint. These regimes are not easily relaxed because their constraints reflect not only security logic but also domestic politics. Even when reform occurs, it tends to be narrow and exceptional.

These forces create a stark implication: any model that depends on sweeping acquisition reform or export control liberalization is strategically fragile. It may be intellectually appealing, but it will not scale reliably across administrations, parliaments, and crises.

DaaS therefore begins from a different premise: treat bureaucracy and export controls as fixed variables and design the market beneath them accordingly. This is not fatalism; it is engineering. If the constraints are durable, then an effective innovation model must be compatible with them.

The key move is to stop trying to force bureaucracies to behave like markets and instead create market conditions that function within bureaucratic envelopes. That means:

- preserving sovereign authority over operational use and export;
- preserving hard safety and certification constraints;
- preserving compliance requirements;
- while changing how risk is financed, how competition persists, and how suppliers are disciplined.

This shift—from reforming the state to redesigning the market—is the conceptual foundation for the rest of the paper.



4. WHAT CHANGED: THREAT TEMPO, COUNTERMEASURE CYCLES, AND INDUSTRIAL FRAGILITY

If the legacy system has been imperfect for decades, why is this moment different? The answer is that the underlying “physics” of advantage has changed. The most decisive shift is not any single technology. It is tempo.

In earlier eras, the dominant pattern was: build a new generation of capability, field it, and enjoy an advantage for years. Adversaries responded, but response cycles were slower, and the advantage persisted long enough to justify the program model.

In modern conflict, particularly in contested environments with pervasive sensors and electronic warfare, advantages decay rapidly. Systems face countermeasure cycles that occur on operational timescales. The contest shifts from “who has the better platform” to “who can adapt faster.”

This change has several structural consequences:

Effectiveness is not durable. The value of a system is no longer well modeled as a one-time performance attribute. It is a function of continuous iteration. A weapon that is excellent today may be mediocre in months if countermeasures emerge.

Battlefield feedback matters. Performance cannot be fully predicted in the lab. The battlefield becomes a learning system. Organizations that cannot ingest and operationalize battlefield feedback quickly will fall behind.

Innovation must be continuous. A development model with an RDTE spike followed by a long sustainment tail is mismatched. The engineering team cannot disappear after initial fielding; it must remain engaged because the environment continues to change.

Supply chain becomes a weapon. Industrial capacity and logistics speed become strategic determinants. If one side can manufacture and iterate faster, it can compensate for tactical disadvantage. Conversely, an advanced design is irrelevant if it cannot be produced at volume.

These changes are frequently discussed in terms of drones, software, and electronic warfare. But the deeper point is architectural: the defense enterprise must behave like a continuously learning system. That requires a financial and industrial model that supports persistent iteration under hard constraints.

[Figure 2: Weapon effectiveness decay vs. adversary adaptation rate]
A graph showing weapon effectiveness on the y-axis versus time on the x-axis. The legacy model assumes a slow decay curve; the modern contested environment shows steep decay driven by countermeasures. Overlay shows “refresh events” (new versions) and illustrates why long upgrade cycles fail to keep pace.

5. THE INDUSTRIAL BASE IN PEACE AND WAR: SURGE FAILURE AND DOWNTURN COLLAPSE

Industrial performance under stress is the ultimate test of a defense ecosystem. Modern conditions have exposed fragility in both directions: inability to surge during conflict and inability to maintain capacity during peace.

Surge failure. When conflict escalates, demand for munitions and platforms spikes sharply. In a resilient industrial base, surge is achieved through redundancy, flexible capacity, and rapid scaling across multiple producers. In the current model, surge is often constrained by single-source production lines, long lead times, limited tooling, and a thin supplier layer. Programs optimized for efficiency in peacetime become brittle in war.

Downturn collapse. When conflicts subside, defense budgets contract. In the legacy model, this contraction often collapses program-specific manufacturing ecosystems. Suppliers exit; specialized labor disperses; tooling is mothballed; and production knowledge decays. The next crisis then demands expensive reconstitution. The system behaves like a boom-bust commodity cycle, but the commodity is national security capacity.

The combination creates a perverse dynamic: the defense industrial base is too rigid to surge quickly and too fragile to remain healthy between surges. This is the opposite of what deterrence requires. Deterrence is not only about possessing capability; it is about maintaining credible capacity to scale and adapt when threatened.

The legacy model concentrates manufacturing in program-captive structures. When demand shifts, these structures cannot easily pivot. Civilian markets solve this through diversified demand and flexible production networks. Defense markets do not naturally have those properties. DaaS attempts to create them intentionally through federated manufacturing and multi-program participation—topics developed later in the paper. For now, the critical insight is that industrial fragility is not incidental. It is a consequence of a model that treats procurement as episodic and manufacturing as program-specific rather than market-based.


[Table 1: Legacy vs. DaaS resilience comparison] A table comparing legacy procurement and DaaS across surge capacity (number of qualified producers, lead time flexibility, single-point failures), downturn survivability (ability of manufacturers to pivot, continuity of engineering cadence), and reconstitution cost. Intended to show why DaaS improves resilience by design through federated production and subscription-funded continuous RDTE.

Part I Summary: The Problem Statement in One Sentence

The program-centric defense model is a stable equilibrium optimized for governance and procurement accountability, but it is structurally misaligned with modern warfare’s dominant driver of advantage—continuous adaptation—while simultaneously producing an industrial base that is brittle in wartime surge and fragile in peacetime contraction. That is the problem DaaS is designed to solve.

PART II

LAT 34.0522° N
LON 118.2437°



THE DAAS MARKET ARCHITECTURE

200
180
160
140
120
100
80



PART II — THE DAAS MARKET ARCHITECTURE

6. DAAS FIRST PRINCIPLES: FROM PRODUCTS TO CONTINUOUS SERVICES

Defense acquisition has long been organized as if a capability were a finished object: a platform, a missile, a radar, a shell—fielded at some moment in time and then sustained until replacement. This mental model is deeply embedded in policy language (“programs of record”), budgeting constructs (discrete appropriations and line items), and contracting practice (a build-deliver-sustain sequence). It is also increasingly misaligned with how advantage actually behaves under modern contestation.

DaaS begins by replacing the “finished product” assumption with a different first principle:

A defense capability is not an object; it is a continuously contested service state.

A capability only exists to the extent it remains effective against evolving threats within a dynamic environment. That effectiveness is a function of continuous iteration across design, software, manufacturing, and integration. In practice, the “weapon” is the combination of physical system, firmware/software, training doctrine, logistics, and the manufacturing and certification system that sustains the ability to replace and improve it. Under contestation, each of these elements must evolve, or the capability decays.

DaaS translates that reality into an industrial architecture. It does not require ministries to abandon their acquisition frameworks. It does not require states to privatize sovereignty. It does not require export controls to loosen. It simply rearranges incentive structures so that continuous competition becomes the natural mode of operation.

Four principles define the architecture.

Principle 1: Continuous RDTE is normal, not exceptional. Legacy models treat RDTE as a phase; DaaS treats it as a permanent function. There is no “engineers disbanded after fielding” moment. The engineering organization is maintained and incentivized to iterate indefinitely because iteration is the source of advantage.



Principle 2: Distributed risk replaces single-nation risk concentration.

The key failure of program-centric acquisition is not that it spends money; it concentrates risk in a single sovereign. DaaS distributes RDTE across a coalition of allied MODs and introduces private capital as a bounded co-bearer of early risk. This allows more programs to be attempted, more quickly, with lower downside per participant, while still maintaining hard gates for safety and certification.

Principle 3: Licensing replaces transfer; access replaces ownership. In DaaS, the prime retains ownership of core intellectual property (IP) and licenses it to program participants and certified manufacturers. This is not merely a commercial preference; it is a strategic device.

Government-owned bespoke artifacts often become dead ends—hard to iterate, hard to commercialize, hard to maintain. Licensing creates a persistent incentive for the DaaS prime to improve the product because its revenue depends on continued participation and adoption, not on one-time delivery.

Principle 4: Competition must persist after award. Traditional defense competition often occurs once—at the contract award—then collapses. DaaS forces competition to persist throughout the lifecycle, not necessarily through duplicate programs, but through mission-level overlap, measurable outcomes (CELS), and credible exit and reallocation options for MODs.

These principles yield a crucial implication: DaaS is not primarily a contracting instrument. It is a market structure that changes how capital flows, how risk is distributed, how systems evolve, and how suppliers are disciplined.

A common objection at this point is that defense is inherently too different from consumer technology to adopt service logic. The objection is partially correct: defense has hard constraints—safety, lethality, export control, classified subsystems, and sovereign authority. But it is precisely because those constraints exist that the design must be structural rather than rhetorical. DaaS does not attempt to make defense “like Silicon Valley.” It makes defense economically rational under sovereign constraints by creating a system where the prime must compete continuously, not merely comply episodically.

7. PRIVATE CAPITAL’S STRUCTURAL ROLE: BREAKING THE MONOPSONY

Any serious attempt to re-architect defense must confront the central market fact of the industry: demand concentration. In most defense markets, the buyer is structurally singular at the national level and highly constrained in how it can procure and export. This creates a monopsonistic equilibrium. In a monopsony, suppliers optimize to satisfy the buyer’s processes, not to win continuous competition. The buyer, in turn, optimizes for political risk reduction, not for market responsiveness. The equilibrium is stable and self-reinforcing.

DaaS introduces private capital not as decoration, but as a structural lever: a second buyer of development risk.

7.1 Capital as buyer of timing risk

In the legacy model, early-stage RDTE is funded almost entirely by the state. If the state does not fund it, it does not happen. This creates a bottleneck: a program’s existence depends on public budget cycles, political timing, and bureaucratic appetite for risk.

DaaS creates an alternative: private capital can fund early technical and timing risk while coalition participation builds. The purpose is not to privatize defense; it is to avoid a world where the only path to innovation is a single sovereign’s willingness to bear 100% of early uncertainty.

Private capital in this architecture buys timing optionality: it accelerates the start of development and provides runway for iterative testing before broad adoption. It is bounded and disciplined by the fact that success requires government participation and procurement volume. It cannot “force” adoption; it can only price risk.

This is precisely why DaaS requires careful definition of capital’s role. The capital is not meant to govern operational use, exports, or national policy. It is meant to absorb the early-phase uncertainty that the public procurement system is structurally bad at absorbing quickly.

7.2 Dilution as discipline

A subtle but critical consequence follows: because private capital is dilutive to the prime, it functions as an execution penalty when government adoption lags. In the legacy system, cost growth and schedule slip can be survivable, even profitable, because the state is locked in and the program becomes politically hard to cancel. In DaaS, failure to attract subscribers early forces the prime to raise more capital on worse terms.

This reverses a core incentive.

- In the legacy model, slipping can be tolerated because exit is painful for the buyer.
- In DaaS, slipping is punished because dilution grows and market credibility decays.

Dilution therefore replaces cost growth as the primary disciplining mechanism. This is not a moral claim; it is a structural one. It creates a durable incentive for primes to (a) build coalition-relevant capabilities with broad demand and (b) deliver measurable progress quickly enough to convert demand into early subscriptions.

7.3 Capital-enabled leapfrogging (preview)

Private capital also enables asymmetric overinvestment—what we will later call “burst-mode innovation.” A DaaS prime that sees a chance to leapfrog a competitor can raise capital and temporarily increase RDTE spend, accepting dilution in exchange for performance discontinuity and market share capture. This creates a dynamic analogous to competitive R&D investment in commercial markets, but bounded by defense realities (certification, safety cases, and sovereign control).

The key insight is that DaaS does not require primes to remain funding-constrained to annual government cycles. It permits them to accelerate when acceleration is strategically valuable—while imposing cost (dilution) if they do so poorly.

7.4 Political legitimacy and sovereign safeguards

Critics will argue that introducing private capital invites profit-seeking behavior into national security. The critique is natural and must be addressed directly. The answer is not that profit is always good. The answer is that profit already exists throughout defense; the question is whether profit is aligned with outcomes or aligned with inertia.



DaaS constrains private capital in three ways:

1. Sovereign authority remains absolute over operational use, targeting, doctrine, and export decisions. Capital does not touch these levers.
2. Capital is rewarded only by adoption: it is exposed to market demand and therefore disciplined by coalition behavior, not by political lock-in.
3. Program governance remains bounded by certification and interface standards. Capital cannot substitute for safety cases and cannot bypass compliance.

The political argument becomes stronger when framed properly: private capital is not being invited into the kill chain. It is being used to finance early-stage technical risk that the public system is structurally slow to finance, and it is doing so under constraints that preserve sovereignty.

[Figure 3: Capital flow in legacy vs DaaS]

A schematic showing capital sources over time. Legacy: government funds nearly all early RDTE, then production, then sustainment; private capital is minimal. DaaS: early risk partly funded by private capital, gradually replaced by coalition RDTE contributions as MODs subscribe; ongoing annual participation funds continuous RDTE; primes may add “burst” capital for leapfrogging. Highlights dilution pressure as discipline.

8. DISTRIBUTED RDTE AND TRL-BASED ENTRY: PRICING RISK TRANSPARENTLY

DaaS replaces the all-or-nothing funding structure of legacy programs with a staged, coalition-based approach. The core move is simple: share RDTE across multiple allied MODs, allow entry at different times, and price entry as a function of risk.

This architecture reflects an economic truth: the value of joining a program is a function of uncertainty. Early-stage technologies carry high risk; late-stage technologies carry lower risk. In most markets, that difference is priced. In defense acquisition, it is often not. Traditional programs frequently force participants to commit early or not at all, creating perverse incentives for political alignment rather than rational risk-taking.

DaaS formalizes the opposite:

- Early entrants pay less but accept more uncertainty.
- Late entrants pay more but enjoy validated performance.

This creates a rational trade space for ministries. It also generates coalition dynamics that improve outcomes.

8.1 Shared RDTE as coalition-scale innovation

When RDTE is distributed across multiple MODs, three things change simultaneously:

- 1. More programs become feasible. Individual nations can participate in more innovation efforts because their per-program exposure is lower.
- 2. Market signals become legible. If many MODs join early, the program has coalition demand; if they do not, that is a signal of weak product-market fit.
- 3. The prime must optimize for coalition needs, not national idiosyncrasies. In legacy procurement, requirements can become a battleground for niche features that advantage a vendor. In DaaS, niche features reduce the addressable market. The prime is incentivized to focus on common denominators that maximize adoption.

This is a deeper form of interoperability: not a bureaucratically imposed standard, but a market-driven convergence on capabilities that multiple allies will sustain.

8.2 TRL-based pricing as a discipline mechanism

TRL-based entry pricing introduces a new discipline: risk cannot be hidden. In legacy programs, risk is often disguised in schedule optimism, requirements ambiguity, and cost estimates. In DaaS, risk is priced explicitly through entry tiers. Ministries choose when to enter based on their risk tolerance and confidence in execution. This pricing mechanism also improves fairness. Early adopters who bear uncertainty pay less. Late adopters pay more for reduced uncertainty. The system reduces “free riding” without requiring punitive contracts.

8.3 Pre-commitment market signals

The most important effect is often overlooked: DaaS shifts the critical decision point forward. In legacy procurement, program viability is often determined after billions are spent and political commitments are made. In DaaS, program viability is determined earlier by pre-commitment signals:

- how many ministries subscribe;
- at what TRL they enter;
- what level of annual participation they sustain;
- how manufacturing share evolves among certified producers.

This creates an environment where programs that fail to attract demand can die early, cheaply, and visibly—before becoming institutional burdens.

[Table 2: TRL entry tiers vs risk/rights]
A table mapping TRL bands (0–3, 4–6, 7–8, 9) to typical characteristics: technical uncertainty, entry pricing logic (conceptual, not numeric), rights (manufacturing license timing, upgrade access), infl uence on roadmap, expected time-to-fi eld, and risk-sharing proportion across MODs vs capital.

9. SWITCHING COSTS AS ARITHMETIC: WHY EXIT IS CREDIBLE

In the legacy defense model, switching costs are high because they are compounded: sunk RDTE, bespoke integration, proprietary sustainment, and political embarrassment. The high switching cost is not merely financial; it is institutional. Once a nation commits to a program, exiting often means admitting failure, losing capability, and writing off sunk investment. This destroys contestability and creates supplier power.

DaaS reverses the logic by designing switching costs as a bounded function of participation rather than as an emergent hostage situation.

The central relationship is:

Switching Cost ≈ MOD Share of RDTE × Time-in-Program

Because RDTE is distributed across multiple MODs and partially underwritten by private capital, the MOD share is structurally small. Because participation is staged and annualized, time-in-program can be managed. A ministry that concludes a program is not delivering value can exit at defined intervals and lose only what it has contributed to date, not a decade of sunk cost.

This is not merely a financial engineering trick. It is the enabling condition for a functioning defense market. In a market, suppliers compete because buyers can switch. If switching is impossible, competition is ceremonial. DaaS restores switching as a credible option by forcing the architecture to keep exit cheap and legible.

This is where the often-misunderstood relationship between IP ownership and sovereignty becomes clear. In DaaS, the prime owns IP and licenses it. On the surface, critics interpret this as reducing sovereignty. In practice, it can increase sovereignty by ensuring that the ministry never funds an orphaned national artifact it cannot sustain. Under DaaS, the ministry buys access to a living capability ecosystem. If it exits, it keeps what it has purchased to date, but loses future rights. This trade is precisely what makes switching credible: the ministry is not forced to choose between total loss and total dependency. It retains inventory and operational capability while re-allocating future spend.

This logic depends on architecture:

DaaS reduces switching costs not by goodwill, but by architectural constraints: IP licensing rather than government ownership, certified multi-source manufacturing, stable interfaces, and stockpile retention. The result is a competitive market in which a Ministry can reallocate marginal spend across programs without writing off the entire sunk cost base.

Switching costs are therefore not “mitigated”; they are designed.

10. RIGHTS, EXIT, AND REJOIN MECHANICS

If DaaS is to function as a durable market architecture, it must translate the abstract idea of low switching costs into explicit rights and obligations. This is where most conceptual models fail: they describe exit as if it were a cultural preference rather than a legally and operationally governed fact.

The DaaS framework must answer, with clarity, what a subscriber retains, what it loses, and what it can regain.

10.1 Stockpile retention, spares, and maintenance rights

A foundational rule is stockpile retention: a MOD retains everything purchased or manufactured while in the program. This includes inventory and associated documentation required for safe use. Exiting a program does not cause weapons already fielded to disappear. It changes what happens next.

The distinction is:

- Inventory and use rights remain for items acquired during participation.
- Future manufacturing and procurement rights cease upon exit.
- Upgrades and new versions cease upon exit unless negotiated otherwise.
- Maintenance rights for existing inventory persist under defined maintenance agreements, bounded by the last certified version and compliance with sustainment protocols.

This structure avoids the political toxicity of “losing access to our own weapons” while preserving the prime’s incentive to maintain a living subscription ecosystem.

10.2 Upgrade rights, software keys, and version access

Modern systems increasingly require versioned software and cryptographic controls. DaaS must therefore define rights to:

- software updates and patches;
- cryptographic keys and authentication tokens;
- interface specifications and compatibility baselines.

The principle is simple: when a ministry exits, it keeps the versioned capability it has, but it does not receive future improvements. This is analogous to a service model in other domains, but bounded by defense realities.

This separation is important because it creates a strong incentive to remain in the program while leaving the exit option credible. Ministries that want to keep pace with countermeasures maintain subscriptions. Ministries that choose to exit accept that their capability will gradually become less competitive.

10.3 Rejoin pricing logic and settlement rules

Rejoining must not be punitive, but it must not enable free riding. A rational structure is that rejoin requires paying the difference between the program’s current maturity tier and the tier at which the ministry exited, plus any compliance and certification catch-up costs necessary to ensure safe manufacturing and integration. The logic is not punishment; it is fairness. A rejoiner is gaining access to the benefit of RDTE that occurred while it was absent.

This becomes a powerful stabilizer: ministries can exit without catastrophic loss, but they cannot exit and re-enter repeatedly without cost. The program maintains coherence.

10.4 Wartime access assurance and emergency clauses

The most delicate question is wartime: what happens if a ministry exits and later faces urgent need? Here the architecture must balance sovereignty, alliance politics, and market discipline.

The paper’s position is not to prescribe a single universal clause, but to identify the necessary design principles:

- Sovereign emergency authority must be recognized in policy and alliance practice.
- Program rules must remain credible to preserve market discipline; otherwise exit becomes meaningless.
- Coalition arrangements may include emergency provisions for humanitarian or existential scenarios, negotiated outside normal commercial terms, to avoid perverse incentives and political blowback.

The essential point is that wartime exceptions are not a failure of DaaS; they are a reality of state behavior. The architecture must acknowledge them without allowing them to erase the discipline that makes DaaS work in the first place.

Part II Summary: The DaaS Architecture in One Paragraph

DaaS is a market architecture that restores continuous competition to defense by distributing RDTE across allied MODs, using private capital to absorb early risk and enforce discipline via dilution, licensing IP rather than transferring it, and converting procurement into a subscription-funded continuous innovation system. Switching costs become arithmetically bounded rather than politically catastrophic, enabling credible exit. Entry is priced by technological maturity (TRL), creating transparent risk markets for governments. Rights and obligations are explicit—stockpile retention and maintenance continuity paired with loss of future manufacturing and upgrade rights upon exit—so that the system remains both politically viable and economically disciplined.

PART III_

LAT 34.0522° N
LON 118.2437°

GOVERNANCE CERTIFICATION AND SECURITY



PART III — GOVERNANCE, CERTIFICATION, AND SECURITY

11. COALITION INTERFACE GOVERNANCE VS COALITION ENGINEERING GOVERNANCE

A market architecture that invites multiple sovereign customers, multiple manufacturers, and continuous iteration will collapse into chaos unless it is governed correctly. The temptation—especially among governments—is to respond to complexity with central planning: committees, harmonized requirements, and bureaucratic control over design choices. That instinct is understandable. It is also the path back to the very system DaaS is designed to transcend.

DaaS therefore rests on a governance distinction that must be explicit:

DaaS uses coalition interface governance, not coalition engineering governance.

This is not semantic. It is the difference between a system that can evolve under contestation and one that freezes under committee.

11.1 Coalition engineering governance: why it fails

Coalition engineering governance is the traditional multi-national failure mode: many stakeholders attempt to co-design a single platform. Each ministry brings legitimate national constraints—doctrine, terrain, politics, industrial offsets, local procurement rules. The resulting consensus tends to maximize inclusivity and minimize controversy. Requirements accumulate, exceptions proliferate, and the program becomes a negotiated artifact rather than a competitive weapon. Schedule and cost expand because every change must be coordinated across all parties. Innovation cadence collapses because engineering must wait for political agreement.

In effect, coalition engineering governance recreates the “requirements market” that already exists in legacy procurement—only larger. It is a recipe for delay, gold-plating, and brittle interoperability.



11.2 Coalition interface governance: what it is

Coalition interface governance is fundamentally different. It does not attempt to design the weapon together. It governs the rules of compatibility and the conditions of certification under which independent suppliers compete.

The coalition governs:

- Interface standards (mechanical, electrical, software APIs) at defined layers
- Interoperability envelopes (what must remain compatible across versions)
- Backward compatibility policies (how long interfaces must be supported)
- Certification baselines and audit requirements (what is required to be “approved”)
- Information security rules for interfaces and data flows
- Change control procedures for interface evolution (how standards move forward without fragmentation)

The coalition does not govern:

- the prime’s design decisions inside the interface envelope
- the prime’s internal roadmap priorities (beyond subscription and market signals)
- the prime’s engineering management
- detailed feature requirements that become political bargaining chips

This structure is how modern ecosystems scale. It is the logic behind open standards that allow competitive markets: USB, PCIe, TCP/IP, container standards in shipping, and in defense itself, the best of NATO standardization when it is used for interoperability rather than co-design. The system remains coherent because the coalition controls the contract between components, not the components themselves.

11.3 Why DaaS primes accept interface governance

A skeptic might ask why a DaaS prime would accept coalition constraint at all. The answer is economic: in a multi-buyer market, interoperability is a demand multiplier. Standards that allow coalition members to integrate systems and train together increase adoption. A prime that refuses necessary interface convergence may win one national customer but lose coalition scale. The market disciplines the prime in a way that single-buyer procurement does not.

This is one of the structural advantages of DaaS: the prime is not negotiating against one buyer under a captive relationship; it is competing for coalition adoption. Coalitions can therefore enforce interface governance through purchasing power rather than through central planning.

11.4 How interface governance avoids “format wars”

A legitimate concern is standards fragmentation: multiple primes might create incompatible ecosystems, forcing ministries into “format wars.” The DaaS answer is not to pretend this cannot happen; it is to specify the conditions under which it does not dominate.

- Interoperability requirements should attach to interfaces that matter for coalition operations (e.g., targeting data, fire control messaging, safety cases, logistics compatibility), not to internal design.

- The coalition can converge on standards by spend allocation: ministries can preferentially subscribe to programs that adopt the coalition baseline.
- Leapfrogging remains possible: a prime can break ranks with a discontinuously superior solution, but only if the performance gain justifies the interoperability cost. This is how real markets evolve; it is also how stagnation is prevented.

The key is that the system must allow both convergence and disruption. Over-convergence freezes innovation; over-disruption fragments coalitions. Interface governance is the mechanism that balances both.

12. CERTIFICATION, QA, AND VERSION CONTROL UNDER CONTINUOUS CHANGE

No defense capability can be treated like consumer software in one crucial respect: failure is not merely inconvenient. Failure can be lethal, politically catastrophic, and strategically destabilizing. A DaaS system that cannot maintain rigorous certification under continuous iteration is not merely flawed; it is unacceptable.

The central challenge is that DaaS introduces a new tension: continuous change versus safety assurance. Legacy systems reduce this tension by changing rarely. DaaS reduces it by changing continuously—but with disciplined processes that make continuous certification possible.

This section defines the certification architecture required for DaaS to be credible.

12.1 Safety case ownership

Every certified defense system rests on an explicit or implicit safety case: a structured argument supported by evidence that the system performs safely and predictably under defined conditions. In DaaS, the safety case must be treated as a living artifact.

The DaaS prime must own:

- the master safety case for the system baseline
- the definition of test evidence required for certification
- the authority to certify versions and revoke certification
- the audit trail that links changes to test evidence and deployment decisions

Participating ministries may demand access to safety case evidence and may impose additional sovereign constraints. But safety case stewardship cannot be diluted across committees without collapsing accountability. A single accountable steward is necessary.

This is an important inversion: in legacy programs, governments often demand ownership of artifacts as a form of control. In DaaS, control is achieved not through ownership, but through enforceable transparency and auditability, paired with the ability to redirect spend if discipline fails.

12.2 Change control and backward compatibility

Continuous RDTE requires version control, but defense requires something more: version legitimacy. DaaS must specify what constitutes a version change:

- Patch releases: bug fixes, non-functional changes, minor performance improvements within a certified envelope
- Minor versions: changes that affect performance or interfaces but remain backward compatible under defined rules
- Major versions: changes that alter interfaces, require new certification evidence, or change operational assumptions

This taxonomy is not bureaucratic ornament. It is how ministries can maintain operational confidence while allowing rapid iteration. It is also how manufacturers can invest in tooling and QA without fearing chaotic design churn.

Backward compatibility policies must be explicit:

- How long must old interfaces be supported?
- Under what conditions can an interface be deprecated?
- What is the minimum “support window” for prior versions in the field?

These policies create the temporal stability necessary for coalition operations and supply chain planning.

12.3 Audit cadence, test regimes, and who pays

The most common failure of modular markets is quality drift: suppliers optimize to win on price, and quality erodes unless audit and certification are strong. DaaS anticipates this by embedding audit cadence into the model.

A credible DaaS certification system must include:

- Qualification testing for new modules and manufacturers
- Periodic requalification to ensure processes remain within spec
- Random audit testing of production runs
- Digital thread traceability linking design, manufacturing parameters, and test results
- Red-team testing for adversarial environments (EW, spoofing, deception) where relevant

Who pays? The paper’s principle is that the cost of trust must be allocated where it creates discipline:

- The prime funds baseline certification infrastructure through annual participation revenue.
- Manufacturers fund module qualification and ongoing compliance as a cost of market access.
- Ministries pay for sovereign-specific requirements that go beyond coalition baseline (if they insist on unique constraints).

This allocation preserves the core market: manufacturers compete on production efficiency and quality, primes compete on architecture and iteration, and ministries pay for sovereignty where sovereignty diverges.

[Figure 4: Versioning and certification pipeline] A flow diagram showing baseline continuous RDTE, code/design check-in, simulation and verification, test gating, certification authority sign-off, release to certified manufacturers, production QA checks, and field telemetry feedback loops. Shows how “patch/minor/major” versions map to test intensity and certification steps.

13. SECURITY AND IP PROTECTION IN A CONTROLLED-MODULARITY MODEL

A market structure that relies on modularity and federated manufacturing raises a legitimate concern: does openness inevitably mean leakage? The fear is that licensing IP and publishing interfaces will accelerate copying, erosion of advantage, and proliferation of sensitive know-how.

This concern is serious. It is also often misframed. The question is not whether information can leak in theory; it is whether a system can be designed such that the value remains defensible while still benefiting from competitive innovation.

DaaS answers this by adopting a principle that is both technically and economically grounded:

Modularity is discretionary; openness is strategic; the prime controls the standard.

This is the defense analog of how successful commercial ecosystems behave. Apple is not open everywhere. It is open where the market accelerates innovation and closed where integrity, security, and differentiation must be preserved.

13.1 Defining “core” versus “peripheral” in defense

To govern modularity intelligently, the system must define what layers are “core” and which are “peripheral.” In defense, “core” is not a brand preference; it is where safety, security, and mission assurance concentrate. A plausible classification:

Core layers (typically controlled tightly):

- certification baselines and safety case logic
- versioning system and change control policy
- critical guidance architecture and mission software baseline (where applicable)
- cryptographic key management and secure boot chains
- sensitive seeker processing and counter-countermeasure logic (case-dependent)
- export-controlled configurations and compliance gates
- core interface standards (the “constitution” of the ecosystem)

Peripheral layers (often commoditized and competed):

- batteries, power modules (within safety envelope)
- mechanical subassemblies and manufacturing process improvements
- IMUs/GPS modules where supply competition improves cost and performance (case-dependent)
- payload variants and fuzing modules (within legal/ethical constraints)
- non-sensitive sensors and auxiliary comms modules (case-dependent)
- tooling innovations that reduce cost without altering certified behavior

The key is that “peripheral” does not mean unimportant. It means that competition can be safely harnessed there because the interfaces and certification rules constrain outcomes.

This classification is dynamic. As technology matures, a component can move from “core” to “peripheral.” As threats evolve, a formerly peripheral layer can become core. The prime’s ability to reclassify components is itself a competitive dimension.

13.2 Preventing IP leakage and cloning

Licensing and distributed manufacturing raise the risk of cloning. DaaS must therefore rely on layered defenses rather than naive trust. Practical defenses include:

- Partitioned disclosure: manufacturers receive only what they need to build their licensed components, not full system design.
- Black-box modules: for sensitive subsystems, the prime can supply sealed certified units while keeping interfaces open.
- Cryptographic attestation: software-enabled systems can verify that components are authentic and certified before enabling full functionality.
- Digital watermarking and provenance: design files and production artifacts carry provenance signatures enabling traceability and legal enforcement.
- Manufacturing certification as gatekeeper: access to licensing can be revoked; cloned production becomes harder to monetize in legitimate markets without certification.

The point is not that leakage becomes impossible. The point is that the system makes legitimate adoption and coalition-scale demand contingent on certification and provenance—so cloning yields a degraded market outcome for the cloner.

13.3 Handling sensitive subsystems and export configurations

Export controls create a technical reality: some configurations cannot be shared broadly. DaaS must treat export compliance as a configuration problem, not as an afterthought.

A DaaS prime can maintain:

1. a coalition baseline configuration eligible for broad allied adoption
2. restricted configurations for sensitive subsystems, available only to cleared partners
3. modular “exportable cores” with optional add-ons that remain sovereign-controlled

This avoids the false binary of “open to all” versus “closed and stagnant.” It also aligns with the model’s broader thesis: export controls are fixed constraints; the architecture must be compatible with them.



14. LIABILITY AND ACCOUNTABILITY IN FEDERATED MANUFACTURING

Federated manufacturing is essential to DaaS, but it creates an accountability challenge: when something fails, who is responsible? Legacy models hide this question because production is concentrated. DaaS must answer it explicitly or risk institutional rejection.

The correct framing is that DaaS does not eliminate liability; it clarifies it through role separation.

14.1 Liability allocation across prime, manufacturer, and MOD use

A defensible liability model maps to distinct responsibilities:

- DaaS prime liability: design integrity, interface definitions, certification criteria, version releases, and system-level safety case.
- Manufacturer liability: production quality, process control, adherence to certified build instructions, and traceability of materials and workmanship.
- MOD liability: operational employment within defined use conditions, maintenance compliance where required, and adherence to handling/storage protocols.

The distinction matters because it preserves the market: manufacturers can compete on production without inheriting design liability they cannot control; primes remain accountable for architecture and certification; governments retain sovereign responsibility for use.

14.2 Warranty, defect regimes, and recall authority

To operationalize the above, DaaS must include:

- warranty regimes tied to certified manufacturing processes
- defect investigation protocols that preserve evidence chains
- recall authority vested in certification governance: the ability to revoke a version or a manufacturer's certification if defects are systemic
- field telemetry and reporting obligations to detect failure patterns early

The recall authority is crucial: without it, certification is symbolic. With it, the prime and coalition can sustain trust in a multi-manufacturer ecosystem.

14.3 Failure investigation governance

When failures occur, attribution must be fast, credible, and politically defensible. DaaS therefore requires pre-defined failure investigation governance that is neither solely controlled by the prime nor bogged down by committee.

A robust structure can include:

- a standing joint safety board with representation from major subscribers
- technical forensics procedures controlled by the certification steward
- transparency obligations that give ministries confidence the prime is not self-investigating in bad faith
- controlled disclosure rules to protect sensitive IP while preserving accountability


This is where the seriousness of the model is demonstrated: DaaS is not “faster because it is looser.” It is faster because it organizes discipline differently—continuous certification, explicit accountability, and market incentives that punish drift.

Part III Summary: Why Governance is the Make-or-Break Layer


DaaS is only credible if it preserves what defense institutions must preserve: safety, security, interoperability, and accountability. It achieves this not by reverting to central planning, but by separating governance layers: coalition governance of interfaces and certification baselines, prime governance of system design and versioning, manufacturer governance of production quality, and sovereign governance of operational use. This architecture enables continuous iteration without surrendering trust—the essential condition for a market-based defense ecosystem to scale.

PART IV_

LAT 34.0522° N
LON 118.2437°



COMPETITION & PERFORMANCE



PART IV — COMPETITION AND PERFORMANCE

15. COMPETITION AT THE MISSION LEVEL: THE VENN DIAGRAM MODEL

The most common misunderstanding of Defense as a Service is to assume it recreates a traditional procurement competition, merely with new contract language. That misunderstanding leads to the wrong metric of success: “How many primes compete for the same program?” In commercial markets, serious competition rarely takes the form of duplicated products with identical scope. It takes the form of overlapping coverage, specialization, and substitution at the margin. Defense is no different—except that the legacy system has historically suppressed those dynamics.

DaaS restores competition by moving it from the program level to the mission level, and from winner-take-all awards to portfolio allocation.

This shift matters because ministries do not ultimately care about programs; they care about mission outcomes: range, survivability, kill probability, cost-per-effect, logistics burden, and the ability to adapt under countermeasures. Programs are merely a means to those ends. When procurement is organized around programs, the system tends to optimize for contractual closure: a defined requirement set, a single vendor, a long sustainment tail. When procurement is organized around missions, the system can optimize for continuous improvement.

The DaaS competitive landscape is therefore best modeled as a Venn diagram.

- A broad program may cover missions A, B, and C (wide applicability, acceptable performance across many scenarios).
- A second program may be narrower but superior for mission A (optimized, lower cost-effective lethality).
- A third program may overlap on missions B and C with different tradeoffs (different environment assumptions, different countermeasure resistance profile).

The result is not duplication. It is contestability—the ability of a ministry to redirect marginal spend toward better mission-level outcomes.

This matters for two reasons.



First, it avoids the wasteful requirement that there be “two primes per program.” That requirement would either be impossible (because of capital intensity) or counterproductive (because it would encourage redundant cloning). DaaS instead preserves competition where it matters: at the mission boundary where ministries allocate budgets.

Second, it makes switching realistic without forcing ministries to abandon existing inventories. A ministry can keep a broad-coverage system for most needs while allocating incremental procurement and R&D participation toward a mission-optimized alternative for the specific mission where performance matters most. This is exactly how healthy technology portfolios behave in the commercial world: general-purpose platforms coexist with specialized devices and services that win on specific outcomes.

This leads to a precise and disciplined competitive condition for DaaS:

- DaaS does not require identical program duplication.
- DaaS does require mission-level overlap sufficient to preserve credible alternatives.

Where a mission category collapses to one provider, monopoly dynamics re-emerge regardless of contract structure. But the threshold for contestability is lower than critics assume. A ministry does not need a full substitute for every program. It needs substitutes for the mission outcomes that drive marginal spending decisions.

[Figure 5: Mission overlap Venn diagrams] A set of Venn diagrams illustrating three DaaS programs with overlapping mission sets. Each overlap region shows relative CELS values (conceptual), demonstrating how competition exists without duplicated programs. A second panel shows how marginal spend shifts toward the prime with better CELS for a specific mission while retaining broad coverage elsewhere.

This structure reframes competition for policymakers. The question is not “How do we select a single best program?” It becomes “How do we ensure mission categories remain contestable so that suppliers must continuously earn demand?” That is the essence of market discipline.

It also addresses a predictable prime-industry critique: “This will fragment the force.” The response is that fragmentation is a function of unmanaged interfaces, not of mission-level competition. Under DaaS, coalition interface governance preserves interoperability, while mission-level overlap preserves contestability. The two are complements.

Finally, the Venn model creates a new type of strategic behavior for ministries: instead of betting their force modernization on one massive program, they can manage a portfolio. This does not increase complexity for complexity’s sake. It allows ministries to convert uncertainty into optionality: invest early in more than one approach, let performance signals emerge, then allocate procurement volume accordingly.

16. CELS AS THE CLEARING METRIC

If mission-level competition is the market structure, Cost-Effective Lethality Score (CELS) is the clearing metric that makes the market legible.

Traditional procurement tends to treat “cost” and “effectiveness” as separate debates. Cost is handled through budgets and unit price negotiations; effectiveness is handled through requirements, test reports, and doctrinal claims. This separation is increasingly untenable because modern warfare collapses the distinction. A weapon is not “effective” in the abstract. It is effective only in an environment, against a

threat, with countermeasures in play, at a given time in the adaptation cycle. Likewise, cost is not merely unit price; it includes logistics burden, operational risk, training complexity, and replacement rate under contested conditions.

CELS is a disciplined attempt to unify these realities into a decision framework.

At minimum, a CELS construct must incorporate:

- Weapon cost: procurement cost and total cost of ownership per usable effect
- Target value and type: the set of targets the weapon can credibly service
- Probability of effect: kill or neutralization probability under realistic conditions
- Environment: EW density, weather, terrain, sensor saturation, deception
- Adversary adaptation: the rate at which countermeasures degrade performance
- Operational constraints: launch platform availability, basing, safety constraints
- Manufacturing and sustainment rate: ability to replace and scale under demand

The purpose is not to reduce war to a single number. The purpose is to allow rational comparison across mission alternatives, especially when multiple DaaS primes compete on overlapping mission sets. Without a clearing metric, mission-level competition becomes political and subjective. With a metric, competition becomes disciplined.

CELS also forces an uncomfortable but necessary truth into procurement: a low-cost weapon that works poorly is expensive, and a high-cost weapon that works reliably may be cheap. The relevant variable is cost per successful effect in the environment that matters. This is why CELS is a stronger organizing concept than “unit price.”

In a DaaS market, CELS becomes the mechanism by which:

- primes compete to reduce cost per effect, not merely to meet requirements;
- ministries allocate marginal spend toward the best mission-level outcome;
- manufacturing share reflects real-world performance rather than contractual rhetoric.

[Table 3: Example mission-level CELS comparison] A table comparing three mission alternatives for a mission category (e.g., counter-UAS, long-range fires, or coastal strike). Columns include unit cost, estimated probability of effect under baseline EW conditions, degradation rate under countermeasures, logistics burden, production scalability, and resulting conceptual CELS ranking. The table illustrates that “cheaper” unit price does not always mean lower cost per effect.

A predictable critique is that CELS can be gamed. That critique is correct—if the metric is naive. The response is not to abandon quantification; it is to anchor the metric in real operational feedback. DaaS strengthens this anchoring because it keeps engineering organizations continuously engaged and because it creates a market mechanism: primes that optimize for a distorted metric lose manufacturing share when real performance diverges. In effect, the market punishes metric gaming if the scoreboard is tied to operational reality.

CELS therefore plays a role analogous to performance-per-watt in computing or dollars-per-throughput in cloud services: imperfect, but indispensable for disciplined competition. It is a tool for governance, not a substitute for judgment.

17. FROM REQUIREMENTS TO OUTCOMES: AVOIDING ‘ARBITRARY REQUIREMENTS’ WHILE PRESERVING NON-NEGOTIABLE CONSTRAINTS

One of the deepest dysfunctions in legacy procurement is that “requirements” become both the language of engineering and the currency of competition. This creates two distortions:

1. Requirements are often shaped by vendor strategy—each prime seeks to embed advantages that it alone can satisfy.
2. Requirements become frozen early, then defended as if they were truth, even as environments change.

The result is that the acquisition system becomes good at building what was specified and poor at building what remains effective.

DaaS does not abolish requirements. It reclassifies them.

A robust DaaS system distinguishes between:

- Arbitrary requirements: features or specifications that are not directly tied to mission outcomes, often artifacts of procurement politics or over-optimization for edge cases. These requirements tend to be rigid, expensive, and slow.
- Non-negotiable constraints: safety, certification baselines, interoperability envelopes, compliance requirements, and operational boundary conditions. These are hard gates that protect sovereignty and reliability.

This distinction is central to why DaaS can move faster without becoming reckless. Outcomes (measured through CELS and operational performance) replace arbitrary requirements as the primary competitive dimension. Non-negotiable constraints remain fixed and are enforced more rigorously through continuous certification.

In practice, the difference is visible in how change is handled.

- In legacy systems, changing a requirement often triggers program upheaval because the program is not built for continuous evolution.
- In DaaS, change is expected, but it must occur through versioned releases governed by the certification and interface rules described in Part III.

This resolves a common critique: “Outcome-driven systems will become unsafe and chaotic.” DaaS does the opposite. It reduces chaos by shrinking the domain in which requirements committees fight. Safety and interface constraints are stable; competition happens inside that envelope on mission outcomes.

There is also a strategic benefit: outcomes are harder to counterfeit than compliance. A prime can claim compliance; it cannot sustainably fake mission-level cost-effective lethality under contested conditions. If it tries, manufacturing orders and subscriber retention will eventually reflect the truth.

This section also addresses the revolving-door influence critique without needing to moralize. In a market where performance is visible and switching is credible, soft influence loses power because it cannot permanently protect an inferior outcome. Influence can help win an initial decision; it cannot hold market share in a system that re-competes continuously on outcome metrics. That is a structural anti-corruption effect: not because humans become better, but because incentives change.

18. CONTINUOUS RDTE, CICD CULTURE, AND CAPITAL-ENABLED LEAPFROGGING

The legacy model is built around discrete phases: RDTE spikes, production ramps, sustainment tails. This structure is not merely financial; it shapes culture. Engineering becomes episodic. Teams surge to deliver milestones and then disperse. Innovation becomes a project, not a discipline.

DaaS replaces episodic engineering with an always-on innovation engine. The mechanism is simple: annual program participation funds continuous RDTE, certification, and versioning. This creates a permanent capability factory rather than a project team.

But DaaS goes further. It introduces two modes of innovation that together create continuous competition:

1. Baseline innovation cadence funded by recurring participation.
2. Burst-mode innovation enabled by capital when leapfrogging is strategically valuable.

18.1 Baseline innovation via annual participation

Annual fees change the economics of engineering. They provide predictable funding that supports a stable team, continuous test infrastructure, and an iterative release cadence. In a DaaS system, a prime that stops innovating does not merely miss an upgrade opportunity; it begins losing relevance immediately. Continuous RDTE becomes the cost of staying in the market.

This is where “CI/CD” becomes more than a software metaphor. It is a cultural requirement: the ability to integrate changes rapidly, test them continuously, certify them efficiently, and deploy them as new versions without destabilizing the ecosystem. Under DaaS, primes that lack this discipline do not merely underperform; they become non-competitive.



18.2 Burst-mode overinvestment and dilution trade-offs

Baseline cadence is necessary but not sufficient. In real competition, sometimes the rational move is to accelerate dramatically—to invest above the baseline to leapfrog a competitor or respond to a new threat.

DaaS allows this because the prime owns its roadmap and IP and can raise private capital to fund acceleration. This introduces a clear trade:

- Spend more now, accept dilution, deliver a discontinuous improvement, capture market share.
- Or conserve equity, iterate slowly, and risk losing manufacturing orders and subscribers to faster competitors.

This is the economic logic that creates real competitive pressure. It is also the logic that legacy procurement suppresses, because suppliers are constrained by government budget cycles and are rewarded for program persistence rather than for leapfrogging performance.

18.3 Manufacturing share as scoreboard

The decisive disciplinary mechanism in DaaS is not rhetoric or compliance. It is manufacturing share.

Manufacturing orders are the closest thing defense has to a real-time market vote. In a federated manufacturing ecosystem, ministries choose where to allocate procurement volume among certified suppliers and competing primes. If a prime's CELS improves meaningfully, manufacturing orders shift toward it. If a prime stagnates, orders slip away. Subscriber decisions follow these signals.

This creates a brutal but healthy consequence: primes that underinvest in continuous innovation lose volume, then lose relevance, then lose subscribers. This is what makes DaaS a market architecture rather than a rebranded procurement framework.

[Figure 6: Milestone vs continuous innovation cadence] Two curves showing “innovation output” over time. Legacy: spikes at milestone gates with long troughs; DaaS: steady baseline cadence with occasional burst spikes. Overlay shows effectiveness decay and how continuous cadence prevents capability obsolescence.

A critic might say this resembles “commercialization” of war. The more accurate framing is that it resembles competition in survival-critical systems. In modern warfare, the cost of slow iteration is not inconvenience; it is strategic failure. DaaS simply makes the innovation cadence commensurate with the threat cadence, while preserving non-negotiable safety and sovereign control.

Part IV Summary: Competition as a Continuous Condition

DaaS restores competition to defense not by forcing duplicate programs, but by creating mission-level overlap and a portfolio allocation model disciplined by CELS. Outcomes replace arbitrary requirements as the primary competitive dimension, while non-negotiable safety and interoperability constraints remain hard gates. Continuous RDTE funded by annual participation sustains baseline innovation, while capital-enabled burst investment allows leapfrogging under pressure. In this architecture, manufacturing share becomes the real-time scoreboard that converts performance into market consequence—ensuring that suppliers compete continuously, not just at award.

PART V

LAT 34.0522° N
LON 118.2437°

MANUFACTURING & INDUSTRIAL RESILIENCE

200
180
160
140
120
100
80

PART V — MANUFACTURING AND INDUSTRIAL RESILIENCE

19. FEDERATED MANUFACTURING AND LICENSING ECONOMICS

If DaaS is the market architecture for continuous competition, federated manufacturing is the industrial mechanism that makes the architecture real. It is also where the model becomes legible to ministries and industry alike: the economics of cost, volume, resilience, and sovereignty live here.

Legacy defense procurement often bundles three functions into one prime: (1) system architecture and IP creation, (2) program-level integration and certification, and (3) manufacturing and sustainment. This bundling is historically understandable—especially when production required unique facilities and integration knowledge was scarce. But bundling produces a predictable failure mode: the buyer becomes captive not only to a design, but to a production monopoly and a sustainment monopoly. Competition ends at award. Costs drift. Surge is limited to the throughput of a single line. And in peacetime, specialized manufacturing capacity becomes brittle because it exists to serve a narrow program.

DaaS deliberately unbundles these functions.

- The DaaS prime remains responsible for architecture, IP, interface governance, versioning, and certification.
- Manufacturing is licensed to multiple certified producers competing on cost, quality, and throughput.
- Sustainment is funded through annual participation and bounded maintenance contracts for fielded inventory.

This unbundling creates a new economic shape. It allows governments to preserve sovereign control over where production happens while forcing production to behave like a market: efficient, scalable, and resilient.



19.1 Cost-plus royalty logic and margin stability

A central design choice in DaaS is the unit economics of procurement. If the prime earns primarily by selling high-priced units, it will naturally resist cost reductions that lower price per unit. If the prime earns primarily through IP rent extracted from captive sustainment, it will resist interoperability and multi-source manufacturing. Both incentives are incompatible with a resilient coalition market.

DaaS instead favors a simple and powerful unit economic structure:

- Manufacturing cost is competed down by multiple producers.
- The DaaS prime earns an IP royalty (often expressed as a percentage or fixed margin over manufacturing cost) and receives annual participation revenue for continuous RDTE, certification, and version stewardship.

This structure has three important properties:

1. The manufacturing price can fall without destroying prime incentives. The prime’s margin can remain stable (as a percentage), and the prime’s total return is driven by scale and adoption rather than by keeping unit costs high.
2. The coalition benefits directly from manufacturing competition. If a certified manufacturer drives down cost through process innovation, that reduction flows to ministries as lower procurement cost. The prime is not harmed by efficiency; it is rewarded through higher volume and continued subscription.
3. The model disciplines both primes and manufacturers. Manufacturers compete on cost and quality; primes compete on CELS improvement and innovation cadence. Each is judged on what they can control.

This is an inversion of the legacy system, where contractors are often rewarded for controlling scarce production lines and navigating bureaucratic requirements. Under DaaS, the path to value is performance and scale.

A predictable critique is that this resembles commoditization of defense manufacturing and could erode strategic capacity. The critique misses the point: DaaS does not commoditize manufacturing to weaken national industry; it federates manufacturing to strengthen it. A nation’s industrial base becomes more resilient when it participates in multiple programs, competes for production volume, and can pivot in downturns. The coalition becomes stronger when manufacturing is not concentrated in one geography or one firm.

19.2 Exclusive territory rights: when they help vs. hurt

Territorial manufacturing exclusivity is a contentious tool. Used poorly, it recreates mini-monopolies and undermines competition. Used well, it can accelerate capacity build-out and sustain regional resilience. The question is not whether exclusivity is good or bad in the abstract; it is under what conditions it increases the coalition’s total capacity and lowers systemic risk.

Exclusivity can be beneficial when:

- it justifies capital investment in capacity in regions that otherwise would not invest;
- it reduces time-to-scale in a strategically exposed geography;
- it creates clear responsibility for readiness and surge in a defined territory;
- it is time-bounded or performance-conditioned (e.g., must meet cost, quality, and throughput targets).

Exclusivity is harmful when:

- it becomes permanent regardless of performance;
- it blocks alternative producers from entering even when they are cheaper or better;
- it creates rent extraction rather than capacity creation;
- it reduces redundancy in a region that needs multiple sources.

A robust DaaS architecture treats exclusivity as a conditional instrument, not as a default entitlement. It is compatible with market discipline only if it can be revoked for non-performance and if it is structured to increase net capacity rather than merely allocate market share.

[Figure 7: Manufacturing federation economics] A schematic illustrating the separation of value streams: (a) annual participation fees funding continuous RDTE/certification, (b) manufacturing cost competed across producers, (c) IP royalty to the prime. A second panel compares legacy bundled pricing vs DaaS unbundled pricing and shows how cost reductions propagate to buyers.

20. Manufacturer Incentives and Capacity Investment

Federated manufacturing will not materialize because it is morally desirable. It will materialize only if manufacturers can rationally invest in capacity and know-how. This is where DaaS must answer the industrialist’s question: “Why should I spend capital to build this?”

In traditional defense programs, manufacturers invest because the contract provides volume certainty or because the state subsidizes the line. In commercial markets, manufacturers invest because demand is diversified and products serve multiple customers. Defense is structurally neither—unless the architecture creates those properties.

DaaS creates manufacturer investment incentives through three mechanisms: multi-program demand smoothing, certification-driven market access, and credible volume competition.

20.1 Capex justification without guaranteed volumes

Manufacturers rarely require absolute volume guarantees; they require a credible pathway to recover investment. Under DaaS, that pathway is provided by:

- access to multiple subscribing MODs as potential customers;
- transparent pricing structures that allow cost leadership to win volume;
- stable interface and certification baselines that reduce retooling uncertainty.

The key shift is from “guaranteed volumes from one buyer” to “contestable volumes from many buyers.” This resembles commercial aerospace supply more than traditional defense: manufacturers compete for share, and share can be defended through cost, quality, and delivery performance.

DaaS therefore changes the manufacturer’s risk profile: instead of being hostage to a single program’s political fate, a manufacturer’s capacity can serve multiple programs and potentially multiple sectors.

20.2 Demand smoothing via multi-program participation

The volatility of defense demand is a major cause of industrial fragility. When a program ends or budgets contract, specialized suppliers often collapse. DaaS reduces this volatility by encouraging manufacturers to participate across multiple DaaS programs and across multiple mission categories. This is where the mission-level competition model interacts with manufacturing resilience. If manufacturing capacity is tied to one program, it is brittle. If it is tied to a portfolio of programs with overlapping demand cycles, it becomes more stable. The manufacturer’s facility becomes a strategic asset for the coalition, not a single-program liability.

20.3 Qualification pipelines and amortization logic

Certification is often seen as friction. In DaaS, certification becomes a market gate that protects quality and creates value for those who invest in it. The qualification pipeline must therefore be clear, repeatable, and time-bounded. Manufacturers will tolerate high standards if the standards are stable and the rewards are access to a durable market.

Amortization logic also matters. Manufacturers invest not only in machines but in process knowledge, workforce training, and quality systems. The DaaS model must allow these investments to be amortized across:

- long participation horizons (multi-year licensing windows),
- multiple customers (coalition adoption),
- and potentially adjacent products (pivot capacity).

This is one reason federated manufacturing increases resilience: it turns specialized defense capacity into adaptable industrial capability rather than captive infrastructure.

[Table 4: Manufacturer investment logic under legacy vs DaaS] A comparison table listing investment drivers (volume certainty, qualification cost, demand volatility, retooling risk, amortization horizon) and how each differs between legacy single-prime manufacturing and DaaS federated manufacturing. Shows why DaaS improves manufacturer willingness to invest when governance is stable.

21. SURGE AND DOWNTURN: DETERRENCE AS AN OPTION STRUCTURE

Industrial resilience is not a slogan; it is an economic property. It is the ability to preserve readiness across two opposite conditions: sudden demand spikes and prolonged demand declines. Traditional procurement models fail both tests because they treat industrial capacity as a program artifact rather than as a market asset.

DaaS reframes industrial capacity as an option structure: ministries maintain access to innovation and manufacturing pathways even when procurement volume fluctuates. This is a profound strategic shift.

21.1 Wartime surge: scaling without renegotiating the state

In wartime, the limiting factor is rarely design; it is throughput. Legacy systems often require renegotiation, reprogramming, and political bargaining to surge production. Single-source lines become bottlenecks. Specialized suppliers become single points of failure. Long lead items become strategic chokepoints.

DaaS improves surge in three ways:

1. Multiple qualified producers already exist. Surge is a matter of reallocating orders, not creating new lines from scratch.
2. Design authority is centralized, production is distributed. The prime can certify and version quickly while production scales across multiple sites.
3. Manufacturing competition creates latent capacity. Producers invest to capture share; that investment becomes surge-ready capacity.

The state still controls procurement decisions, but it is no longer forced to treat surge as a legislative emergency. It can treat surge as industrial allocation.

21.2 Peacetime downturn: sustaining innovation without sustaining volume

The more subtle advantage is downturn resilience. In peacetime contractions, traditional defense firms often respond by cutting engineering, reducing testing cadence, and hollowing out supply chains—because the funding structure collapses when procurement volume collapses. Programs die; knowledge dissipates; reconstitution becomes expensive.

DaaS breaks this cycle by separating the innovation engine from procurement volume.

- Annual participation funds continuous RDTE, certification, and versioning even if procurement slows dramatically.
- Manufacturers can pivot capacity to other programs or adjacent industries because they are not captive to a single prime’s line.
- Ministries preserve access rights and maintain the option to scale procurement quickly when threats return.

This creates a strategy that legacy procurement cannot easily execute: reduce spending without losing capability trajectory. A ministry can cut procurement volume sharply (even by 50–75%) yet preserve a competitive innovation cadence. When the next crisis comes, it does not restart from zero; it accelerates from an already-warm system.

This is what it means to turn deterrence into an option structure. The state is not buying inventory alone; it is buying the ability to generate and sustain inventory and improvements under changing conditions.

[Figure 8: Downturn survival curves—legacy vs DaaS] A graph of industrial capacity and innovation cadence over time during budget contraction. Legacy model shows sharp drop in engineering capacity and supplier exits; DaaS model shows stable baseline innovation via participation fees and flexible manufacturing capacity pivot, with faster re-acceleration when demand returns.



22. EXPORT CONTROLS AND COMPLIANCE IN A DISTRIBUTED MODEL

A distributed manufacturing model that spans allied borders will immediately encounter export control concerns. If DaaS required wholesale reform of export controls, it would not be a viable architecture. DaaS assumes export control regimes persist and designs around them.

The key reframing is to treat export control as a configuration and governance problem, not as a reason to centralize everything.

22.1 Program-level compliance vs transaction-level friction

Traditional systems often treat each export transaction as an isolated legal event, creating friction and delay. DaaS can shift part of this burden to program-level structures:

- define exportable baseline configurations that are broadly licensable;
- define restricted configurations with controlled access for sensitive subsystems;
- structure licensing and manufacturing rights to align with approved partner sets;
- maintain auditable provenance and certification to support compliance enforcement.

This does not remove export controls. It makes compliance repeatable and scalable.

22.2 Compliance as a market gate

In a federated manufacturing ecosystem, compliance enforcement must be credible or the system becomes strategically dangerous. DaaS addresses this by binding market access to certification and licensing. Manufacturers that cannot maintain compliance lose certification and therefore lose market participation. This is a stronger discipline mechanism than relying purely on post hoc enforcement.

The practical implication: openness is discretionary.

The prime decides how open a component is. This is not arbitrary; it is strategic. Some components will remain tightly controlled because they are export sensitive or because they constitute the core of competitive advantage. Others can be opened to competitive supply because they are commoditizable and benefit from market innovation.

This is a direct application of the “controlled modularity” principle established earlier: openness is not ideological; it is a business and security decision.

Export control critics may argue that distributed manufacturing increases the surface area for leakage. That is true if governance is weak. It is not true if governance is designed correctly: partitioned disclosures, certified processes, cryptographic attestation, and provenance enforcement can create a system that is both scalable and controlled.


DaaS does not eliminate export controls. It reduces the extent to which export controls force the entire industrial base into the slowest possible operating mode.

Part V Summary: Manufacturing as the Physical Enforcement of Market Discipline


DaaS becomes real when manufacturing is unbundled from design and placed into a licensed, certified, competitive ecosystem. Cost-plus royalty economics align incentives so that manufacturing efficiency benefits ministries without undermining prime incentives. Capacity investment becomes rational because demand is diversified across coalition customers and multiple programs. Surge becomes allocation rather than renegotiation; downturn becomes survivable because the innovation engine is funded through participation rather than volume. Export controls remain, but the architecture makes compliance scalable through program-level governance and controlled modularity.

PART VI_

LAT 34.0522° N
LON 118.2437°



EVIDENCE IMPLEMENTATION & LIMITS



PART VI — EVIDENCE, IMPLEMENTATION, AND LIMITS

23. EMPIRICAL ANALOGS: WHERE THIS STRUCTURE ALREADY WORKS

DaaS can appear unfamiliar to defense professionals because it breaks the customary bundling of RDTE, manufacturing, and sustainment under a single long-duration program. Yet the underlying structure—distributed risk, staged commitment, modular standards, continuous iteration, and multi-source production—is not novel in industrial history. It is common in domains where (a) the cost of early uncertainty is high, (b) the value of being late is real, and (c) the market cannot tolerate single-point failures.

Four analog families matter because they validate specific DaaS mechanisms, not because defense is “the same” as these industries.

23.1 Pharmaceutical development: staged risk, priced uncertainty

Modern drug development is a canonical example of priced uncertainty. A molecule begins with high technical risk; it becomes more investable as evidence accumulates through phases. Capital and partners join at different times at different prices; early capital accepts higher risk for lower price, later capital pays more for reduced risk. This is not a moral preference; it is how a complex system finances uncertainty without asking one actor to carry the entire burden.

DaaS mirrors this structure at the coalition level. TRL-based entry pricing is not an invention; it is a defense translation of staged-risk financing. The strategic value is not merely financial. It is informational: early joiners signal confidence, late joiners price proof, and the market allocates resources toward what demonstrates traction.

23.2 Semiconductor ecosystems: core control with peripheral competition

Semiconductor supply chains illustrate the separation of architectural control from manufacturing competition. Foundries, IP core providers, design houses, packaging specialists, and tool vendors co-evolve under interface constraints and rigorous qualification regimes. No single actor controls the whole stack, yet the system remains coherent because standards and qualification gates define legitimate participation.



This maps directly to DaaS’s “controlled modularity” principle. Critical interfaces and certification envelopes are governed, but competition is encouraged wherever competitive suppliers can improve cost and performance. The ecosystem’s speed is not achieved by lowering standards; it is achieved by allowing many actors to innovate inside a disciplined envelope.

23.3 Aero-engine risk-sharing partnerships: distributed RDTE without distributed design authority

Commercial jet engines provide another relevant analog: development is capital-intensive; certification is unforgiving; and global supply chains are inevitable. In response, the industry often uses risk-sharing partnerships where multiple firms fund development and build major modules, while a lead integrator retains system architecture and certification authority.

This is almost a direct analog for DaaS primes and federated manufacturing. What matters is the governance: distributed funding and distributed production can coexist with centralized certification and version authority. That governance structure is precisely what DaaS formalizes for defense.

23.4 Subscription software and “always-on” engineering

In software-defined domains, a product is not “complete” at release; it is continuously updated to remain secure and competitive. Subscription revenue funds the permanent engineering organization required to maintain and improve the service. This is not a business trend; it is a technical necessity driven by continuous vulnerability discovery and evolving user requirements.

DaaS uses subscription logic not because defense is “software,” but because contested environments impose a similar necessity: advantage decays without continuous updates. The defense translation is not “ship fast and break things.” It is “ship continuously and certify continuously,” preserving safety and reliability while sustaining iteration.

[Figure 9: Analog mapping matrix] A matrix mapping DaaS mechanisms to analog industries: staged risk financing ↔ pharma; controlled modularity ↔ semiconductors; risk-sharing development ↔ aero engines; subscription-funded continuous improvement ↔ software. Each cell briefly states what transfers and what does not.

The purpose of these analogs is not to claim equivalence. It is to establish plausibility: the structures DaaS relies upon are proven ways to finance, govern, and scale complex systems under uncertainty. Defense’s uniqueness lies in sovereignty, compliance, and lethality—not in the basic economics of uncertainty and competition.

24. CURRENT IMPLEMENTATION: TIBERIUS AEROSPACE, SCEPTRE, AND INVICTUS (ANALYTICAL)

A theory is not credible until it can be shown to function in contact with institutions. DaaS is already being implemented in practice, most explicitly by Tiberius Aerospace, founded by Chad Steelberg, with programs such as Sceptre and Invictus. This section treats those programs as case material for the model’s mechanics, not as promotion.

The relevant question is not “Are these specific systems good?” The relevant questions are:

- Does the DaaS structure materially change incentives?
- Does it enable faster iteration without sacrificing certification discipline?
- Does it create credible multi-buyer participation and bounded switching costs?
- Does it produce a manufacturing market rather than a captive line?

24.1 What is genuinely new (structurally)

Several aspects differentiate a DaaS program from a traditional program even when the weapon class appears familiar:

1. Licensing-based customer relationship The prime retains IP and licenses it to subscribers for manufacturing and procurement rights. This makes the prime economically dependent on continued customer participation rather than on a one-time program award.
2. Annual participation funding continuous engineering A DaaS program funds “always-on” RDTE, certification, and versioning as a permanent function, rather than an initial spike followed by dormancy.
3. Explicit modularity as a strategic choice Openness is discretionary. The prime can commoditize peripheral components where competitive suppliers improve cost and performance, while protecting core layers required for safety, security, and differentiation.
4. Market discipline through manufacturing share The procurement decision is not a single award; it is ongoing allocation. Manufacturing volumes become the scoreboard that reveals whether performance improvements translate into adoption.
5. These are the mechanisms that matter in evaluating any DaaS implementation.

24.2 How Sceptre and Invictus illustrate the model

Without making promotional claims about specific performance characteristics, Sceptre and Invictus serve as useful examples of DaaS applicability: both occupy mission spaces where (a) lethality is continuously contested, (b) iterative improvements can materially change battlefield outcomes, and (c) coalition demand plausibly exists.

- Sceptre illustrates DaaS in a munition domain where manufacturing can be federated, where versioning can occur as improvements to guidance, seekers, payload options, or countermeasure resilience, and where cost-effective lethality is measurable in operational terms.
- Invictus illustrates DaaS in a missile domain where modularity decisions become strategically important: certain subsystems may be opened to competitive supply (e.g., power, manufacturing process improvements), while other subsystems remain tightly controlled for security and certification reasons.

The purpose of these programs in the thesis is to demonstrate that DaaS is not purely theoretical. The structure can be instantiated in real programs where the market pressures of continuous competition and multi-buyer participation can operate.



24.3 What must be measured to validate a DaaS program

A DaaS model should be judged by measurable signals that reflect the architecture’s claims:

Iteration cadence (rate of certified version releases)

- Time-to-certify change (how quickly improvements can be fielded safely)
- Subscriber growth and churn (whether the coalition market behaves like a market)
- Manufacturing competition (number of certified producers; evidence of cost reduction)
- Cost-effective lethality trendline (whether CELS improves over time under contestation)
- Resilience indicators (ability to maintain innovation cadence during procurement slowdowns)

[Table 5: DaaS program validation metrics] A table listing measurable indicators (cadence, time-to-certify, churn, manufacturing cost curves, number of certified producers, defect rates, CELS trend) and describing what “good” looks like in each. This table serves as a practical evaluation tool for MODs and partners.

This section’s embedded critique response is straightforward: traditional primes will argue they already iterate. The correct answer is that iteration exists today, but it is not structurally funded or structurally disciplined by multi-buyer markets. DaaS is not a claim that engineers elsewhere are incapable. It is a claim that the economic architecture must be different for iteration speed to become the system’s default output rather than its exception.

25. INVESTOR INSTRUMENTS, TIME HORIZONS, AND GOVERNANCE LIMITS

Private capital only strengthens a defense innovation system if it is aligned with sovereign priorities and constrained from distorting outcomes. This requires explicit choices about what investors buy, how they are paid, how long they wait, and what they cannot influence.

25.1 Equity vs. royalty: what capital is actually purchasing

In a DaaS context, capital can plausibly enter via:

- Equity in the DaaS prime (standard venture financing)
- Royalty-linked instruments tied to IP licensing revenue
- Structured project finance for capacity and certification infrastructure
- Hybrid structures combining equity upside with capped returns for certain tranches

The choice matters because defense adoption can be slow even when iteration cycles are fast. Pure venture equity expects outsized growth, which can distort strategy toward “largest possible TAM” even when mission fit is narrower. Royalty-linked structures can be better aligned for defense because they reward adoption and volume without requiring exponential growth narratives.

The paper’s core point is not that one instrument is superior. It is that the capital instrument must match the defense time horizon. A DaaS market can create faster technical iteration, but sovereign procurement still imposes adoption friction.

25.2 The time horizon mismatch: venture logic vs. defense reality

A persistent concern is that venture capital seeks rapid exits, while defense programs require long trust-building cycles. DaaS mitigates this mismatch by creating earlier and clearer demand signals (subscriber entry, churn, manufacturing share). However, it does not eliminate the mismatch.

Therefore, a robust DaaS ecosystem tends to attract a specific profile of capital:

- patient, technically literate capital that understands certification and procurement
- capital that values durable cash flows (subscriptions and royalties)
- capital willing to fund staged risk in exchange for staged adoption

This is closer to infrastructure or industrial technology capital than to consumer internet venture capital. The capital stack must reflect that reality if DaaS is to scale responsibly.

25.3 Governance limits: preventing investor capture of the roadmap

The most legitimate fear is not that investors will control operations; it is that they will distort the roadmap toward capital optics rather than mission outcomes. DaaS must constrain this explicitly through governance boundaries:

- investors cannot override certification discipline to ship faster
- investors cannot force openness of export-sensitive modules
- investors cannot convert a coalition market into a monopoly strategy by enclosure
- investor influence is bounded by the ultimate market discipline: ministries can redirect spend and exit, forcing the prime to prioritize CELS and operational relevance

In other words, investor governance must be constrained by architecture. DaaS’s claim is that the market disciplines both primes and capital because adoption and volume are the only sources of durable return.

[Figure 10: Capital stack options for DaaS programs] A layered diagram showing early-stage risk capital, growth capital, and manufacturing capacity finance, with arrows indicating which risks each tranche absorbs (timing risk, technical risk, scaling risk). The figure illustrates how different instruments align with different program phases without implying any single “correct” stack.

26. BOUNDARIES AND FAILURE MODES: WHEN DAAS REVERTS TO MONOPOLY

A serious architecture must define its failure conditions. DaaS is not magic. It is a market structure. Market structures fail when contestability collapses, governance drifts, or incentives reconverge on rent extraction.

The most important boundary condition, derived directly from the competitive model, is:

If mission overlap disappears, monopoly dynamics re-emerge regardless of contract form.

But that is only one failure mode. A robust DaaS framework must acknowledge several:

26.1 No mission overlap: contestability collapse

If a mission category admits only one credible supplier, ministries lose the ability to reallocate marginal spend. Exit becomes politically harder, pricing discipline weakens, and the DaaS prime can gradually re-enclose the ecosystem. The system reverts toward a traditional prime dynamic even if contracts remain “subscription-like.”

This is why ministries must actively cultivate portfolio overlap at the mission level. DaaS is not self-stabilizing if ministries allow single-provider dominance to emerge in critical mission categories.

26.2 Standards fragmentation: coalition drift into incompatible ecosystems

If interface governance fails, multiple incompatible baselines can emerge, forcing ministries into format wars. This makes switching harder and fractures coalition interoperability.

The defense remedy is not central planning. It is strong coalition interface governance: stable standards at layers that matter, and purchasing-power enforcement of convergence.

26.3 Weak certification governance: speed without trust

If continuous iteration is not paired with continuous certification, DaaS becomes a credibility failure. Ministries will revert to slow, centralized procurement because they cannot accept distributed production without trust. DaaS lives or dies on the quality of its certification and audit discipline.

26.4 Capital withdrawal: funding cliff without coalition depth

If private capital is required indefinitely because ministries do not subscribe, the program becomes fragile. DaaS works when capital bridges early uncertainty and coalition subscriptions replace it. If coalition depth does not materialize, the program becomes a venture project rather than a market structure.

26.5 Modularity mismanagement: “open” without strategic control

If the prime opens too much of the core, IP leakage and security risk grow. If the prime opens too little, supplier innovation stagnates and the system becomes closed. DaaS requires disciplined, component-by-component modularity decisions, revised over time as maturity changes.

These failure modes are not reasons to reject DaaS. They are the conditions that define where DaaS must be governed tightly and where ministries must behave rationally as portfolio managers rather than as single-program patrons.

27. Strategic Implications for Alliances and Industrial Policy

DaaS is not merely an industrial model; it is a strategic alignment mechanism. It strengthens alliances in a way traditional procurement often undermines: by making industrial cooperation economically rational rather than politically negotiated.

27.1 Alliances become markets, not merely treaties

In traditional systems, allied interoperability is often pursued through joint programs that become politically complex and technically slow. DaaS offers an alternative: allied nations converge on shared baselines through subscription, licensing, and manufacturing participation. Cooperation becomes continuous and economically reinforced.

This has two implications:

- Shared RDTE creates shared stakes. Nations that co-fund a capability have aligned incentives to sustain it and improve it.
- Trade becomes a security multiplier. Licensed manufacturing and competitive production strengthen economic ties and increase coalition resilience.

27.2 Deterrence as a function of adaptation capacity

Modern deterrence depends not only on inventory counts but on the credible ability to adapt under contestation. DaaS improves deterrence by preserving an always-on innovation engine and a distributed manufacturing network that can surge.

This reframes strategic posture: the coalition is not only buying weapons; it is buying the ability to remain ahead of countermeasures.



27.3 Industrial policy without industrial dirigisme

DaaS provides a market-based mechanism for industrial capacity without requiring states to centrally plan production. By licensing and certifying multiple producers, it creates incentives for private firms to invest in capacity, skills, and processes. Industrial policy becomes a matter of structuring the market rather than subsidizing a single national champion.

This is strategically important for smaller allies: DaaS can allow them to build sovereign manufacturing capability as part of a coalition ecosystem, rather than relying entirely on external supply chains.

28. Conclusion: From Programs to Markets

The defense acquisition problem is often framed as a process problem: too much bureaucracy, too many rules, too slow a cycle. That framing is incomplete. The deeper issue is architectural: the program-centric model concentrates risk, suppresses continuous competition, and creates industrial structures that are brittle in wartime and fragile in peacetime.

Defense as a Service is a structural response. It does not ask states to abandon sovereign authority or dismantle compliance regimes. It assumes bureaucracy and export controls persist and reorganizes the market beneath them.

DaaS achieves its effects through a coherent set of mechanisms:

- shared RDTE across multiple allied ministries;
- TRL-based entry pricing that makes risk explicit and fair;
- private capital as a bounded buyer of early uncertainty, disciplined by dilution;
- IP licensing that preserves prime incentives to iterate continuously;
- annual participation funding that keeps engineering, certification, and versioning permanently staffed;
- mission-level competition evaluated through CELS;
- federated manufacturing that drives down costs, increases redundancy, and improves surge capacity;
- explicit governance of interfaces, certification, security, and liability.

The core claim is not that DaaS is universally superior. The core claim is that under modern contestation, continuous competition is not optional—and the legacy architecture cannot sustain it reliably. DaaS offers a credible market structure that can.

Its limits are clear. Where mission overlap does not exist, contestability collapses. Where certification governance weakens, trust fails. Where coalition adoption does not materialize, capital cannot bridge indefinitely. These are not flaws; they are boundary conditions. A serious strategy is one that states its boundary conditions plainly.

The decisive question for democratic states is no longer whether to modernize procurement. It is whether to build an industrial structure capable of sustaining advantage under continuous contestation without relying on improbable institutional reform. DaaS is one such structure: a shift from programs to markets, and from episodic procurement to continuous competition.

BIBLIOGRAPHY

Atlantic Council. Commission on Defense Innovation Adoption: Final Report. Atlantic Council, Jan. 2024.

Department of Defense (U.S.). State of Competition within the Defense Industrial Base. Office of the Under Secretary of Defense for Acquisition and Sustainment, Feb. 2022.

Government Accountability Office (U.S.). Weapon Systems Annual Assessment: Programs Are Not Consistently Implementing Practices That Can Help Accelerate Acquisitions. GAO-23-106059, June 2023.

Government Accountability Office (U.S.). Other Transaction Agreements: Improved Contracting Data Helps Assess Effectiveness. GAO-25-107546, Sept. 2025.

International Traffic in Arms Regulations (ITAR). Exemption for Defense Trade and Cooperation Among the United States, the United Kingdom, and Australia (AUKUS). U.S. Department of State, Directorate of Defense Trade Controls, Dec. 2025.

Reuters. “Weapons Makers Have ‘Conned’ U.S. Military into Buying Expensive Equipment, Army Secretary Says.” Reuters, 14 Nov. 2025.

Associated Press. “EU Plans to Boost Defense Industry and Move Away from U.S. Dependency in Face of Russian Threats.” Associated Press, Mar. 2024.

U.S. Department of State. International Traffic in Arms Regulations: Exemption for Defense Trade and Cooperation Among Australia, the United Kingdom, and the United States. Federal Register, 30 Dec. 2025.

U.S. Department of Defense. “Contract Policy: Competition.” Office of the Under Secretary of Defense for Acquisition and Sustainment, Defense Pricing and Contracting, accessed Jan. 2026.