AI Data Centers: The New Center of Gravity

Engineering the Infrastructure the United States Cannot Afford to Lose

Dr. Ron Ross

CEO, RONROSSECURE, LLC.
ISTS Fellow, Dartmouth College

Abstract

Artificial intelligence is rapidly becoming a strategic source of national power. As frontier AI increasingly supports intelligence, command and control, cyber operations, logistics, autonomous systems, scientific discovery, weapons development, economic competitiveness, and critical infrastructure, the facilities that provide the computational foundation for those capabilities acquire strategic significance of their own. This article argues that, in a major conflict involving the United States and a technologically sophisticated adversary, concentrations of advanced AI compute are likely to become a new strategic center of gravity or, at minimum, a critical requirement on which a broader national center of gravity depends. Drawing on U.S. military center-of-gravity doctrine, Sun Tzu's strategic emphasis on understanding strength, weakness, position, and dependency, Engineering for Compromise (EfC), RAND's Secure Inference Data Center (SIDC) concept, and NIST SP 800-160, Volume 1, Revision 1, the article proposes treating critical AI infrastructure as an integrated national-security system encompassing computation, facilities, power, cooling, communications, personnel, supply chains, maintenance, recovery, and external dependencies. The NIST security design principles should therefore be applied recursively from the cyber and cyber-physical architecture into the facility and supporting infrastructure. The objective is not an invulnerable data center. It is a national AI capability engineered so that cyberattack, physical attack, insider action, infrastructure failure, or compromise of individual facilities does not automatically become strategic mission loss.

1. Introduction

Military power has always depended upon more than weapons. It depends upon the industrial, logistical, informational, technological, and organizational capabilities that allow a nation to understand an adversary, make decisions, mobilize resources, develop weapons, coordinate forces, sustain operations, and adapt faster than its opponent. Artificial intelligence increasingly touches all of these functions. Advanced AI can support intelligence analysis, targeting, logistics, cyber operations, autonomous systems, modeling and simulation, communications, weapons design, scientific research, manufacturing, decision support, and the operation of critical infrastructure.

Artificial intelligence is becoming a strategic source of national power, economic competitiveness, military advantage, scientific discovery, and critical-infrastructure capability. As that dependence grows, nation-state adversaries have increasing incentive not only to steal critical AI assets, but to corrupt, deny, manipulate, or destroy the systems and infrastructure that produce them. The broader strategic conclusion follows: if advanced AI becomes increasingly important to the generation and employment of national power, the infrastructure that produces and sustains AI capability becomes increasingly important to national power as well.

Models are digital assets, but frontier AI capability is physically instantiated. Compute, power, cooling, communications, people, and supply chains are inseparable parts of the mission.

Frontier AI does not exist independently of physical reality. Models require processors; processors require servers; servers require facilities; and facilities require immense quantities of electrical power, cooling, communications, physical protection, maintenance, replacement components, and specialized personnel. The strategically relevant object is therefore not merely a model, software stack, or data center. It is a distributed system of systems whose digital and physical elements jointly produce AI capability.

This article makes the case that such infrastructure is likely to become a new *center of gravity* for the United States in a future high-end conflict. The claim should be understood with doctrinal precision. Joint doctrine defines a center of gravity as a source of power that provides moral or physical strength, freedom of action, or will to act, and ties center-of-gravity analysis to objectives, critical capabilities, critical requirements, and critical vulnerabilities [1]. Centers of gravity can change with objectives and can differ across strategic and operational levels. The argument here is therefore not that every AI data center becomes a center of gravity. It is that sufficiently concentrated frontier AI infrastructure may become a strategic source of power, and that the facilities and dependencies supporting that capability may become critical requirements whose disruption produces disproportionate national effects.

2. The Center-of-Gravity Test for AI Infrastructure

Center-of-gravity analysis asks planners to identify the source of power that makes an objective achievable and then determine the critical capabilities, requirements, and vulnerabilities associated with that source [1]. Applied to a technologically advanced nation, this framework provides a disciplined way to distinguish ordinary digital infrastructure from strategically consequential AI capacity.

An AI capability begins to resemble a strategic center of gravity when its loss would materially reduce national freedom of action, military effectiveness, industrial adaptation, intelligence advantage, or the ability to regenerate other forms of power. That threshold will depend on the conflict, mission, and degree of national dependence on AI. As AI becomes more deeply embedded in operational planning, cyber defense, autonomous systems, intelligence fusion, logistics, scientific discovery, and critical infrastructure, the threshold is likely to be crossed by at least some concentrations of advanced compute.

The analytical structure can be stated in military terms. The critical capability is the ability to generate and employ trusted AI at the speed and scale required by the mission. Critical requirements include advanced processors, model and data assets, electrical power, cooling, communications, facilities, skilled personnel, maintenance, supply chains, and trusted recovery. Critical vulnerabilities are those components or dependencies whose disruption can significantly reduce the critical capability. Once framed this way, the strategic problem shifts from protecting individual servers to protecting the causal chain that produces national AI power.

The data center may be the visible fortress. The real target is the dependency architecture that makes the fortress useful.

---

3. Sun Tzu and the Strategy of Dependencies

Sun Tzu's *The Art of War* is not a technical manual for modern infrastructure defense, but its enduring strategic logic is directly relevant. Its treatment of planning, indirect strategy, strengths and vulnerabilities, maneuver, adaptation, terrain, and intelligence repeatedly emphasizes understanding the adversary, avoiding costly contests against strength when advantage can be gained by exploiting dependencies and vulnerabilities, and shaping conditions before decisive action [2]. The Art of War is invoked illustratively, as a widely cited articulation of strategic logic, rather than as a technical or doctrinal source. The implication

for AI infrastructure is straightforward: ***a capable adversary should not be expected to attack the strongest part of the system simply because that is where the most visible asset resides***.

A hardened AI computing facility may therefore be strategically weak if it depends upon an inadequately protected substation, common cooling plant, telecommunications corridor, shared identity infrastructure, maintenance vendor, fuel distribution network, or supply chain. Two data centers that appear redundant may share the same regional grid, fiber path, cloud control plane, administrative authority, or specialized replacement source. A cyber boundary may be formally verified while a maintenance pathway, facility controller, or emergency procedure creates an unmodeled bypass. A sophisticated adversary will search for precisely these asymmetries because indirect attacks on critical requirements can produce the same strategic effect as direct attacks on the protected asset.

The defensive corollary is equally important. Protection must be based on understanding the complete *system of dependencies* rather than the organizational chart of the entities responsible for them. Cybersecurity organizations, physical-security organizations, utilities, facilities teams, telecommunications carriers, equipment manufacturers, logistics providers, and government agencies may each perform their assigned function competently while the integrated AI mission remains exposed to a *common-mode failure*. Strategy demands that the defender analyze the system the way the adversary will attack it: as one interconnected source of capability.

4. From Secure Data Center to Secure National Capability

RAND's Secure Inference Data Center (SIDC) proposal provides an important architectural starting point. The SIDC is a compact, purpose-built cyber-physical facility designed to protect strategically important AI models and inference data against an exceptionally capable nation-state adversary [3]. The RAND approach is notable because it does not begin with a catalog of controls. It begins with unacceptable losses, derives hazards and security constraints, and carries those constraints through requirements, architecture, formally analyzed boundary mechanisms, hardware realization, testing, and adversarial evaluation [3].

The SIDC can be understood as a digital Fort Knox: a purpose-built facility organized around extraordinary protection of critical AI assets [3]. That analogy is powerful, but a strategic center-of-gravity argument requires extending it. Fort Knox is principally an asset-protection metaphor. A wartime AI infrastructure must do something more demanding: preserve sufficient trustworthy AI capability while an adversary is actively attempting to degrade, corrupt, isolate, or destroy the systems and dependencies that generate it.

The design objective therefore changes from protecting a particular data center to preserving the AI mission. A highly secure facility whose loss causes strategic mission failure is still a brittle national architecture. Conversely, an architecture that can lose a facility, a power source, a communications path, or selected supporting elements without losing the mission may be strategically more secure even if no single site is individually invulnerable.

5. Engineering for Compromise at Infrastructure Scale

Engineering for Compromise (EfC) provides the complementary strategy for this problem [4].[1] EfC begins from the premise that mission-critical systems cannot credibly depend on every component, account, service, person, supplier, and supporting dependency remaining trustworthy secure[2] throughout the

[1] EfC is presented as the author's own proposed framework, introduced in a companion analysis, rather than as an established or independently validated standard; it should not be read as settled doctrine on par with NIST SP 800-160 or RAND's SIDC research.

[2] The term *trustworthy secure* is used in this article for consistency with NIST SP 800-160, Volume 1.

system life cycle. It deliberately separates component compromise, system compromise, and unacceptable loss, and uses architecture to place barriers between them [4].

The same reasoning should be applied to national AI infrastructure. Server compromise should not automatically become cluster compromise. Cluster compromise should not automatically become facility compromise. Facility compromise or physical destruction should not automatically become regional AI-capacity loss. Regional loss should not automatically become national mission loss. The objective is to break the causal chain before local adversity becomes strategic catastrophe.

EfC defines a *Compromise Containment Boundary* as a system-level boundary intended to limit the information, authority, control, functionality, and influence available through compromised elements. Importantly, such a boundary can combine cyber, physical, procedural, human, and organizational constraints [4]. That idea can be generalized into an infrastructure compromise containment boundary: an architectural boundary intended to ensure that loss or adversarial control of one facility, utility, support function, administrative domain, supplier, or geographic node cannot automatically propagate into loss of the critical AI mission.

The objective is not invulnerability. It is containment: ensuring that compromise or destruction stops well short of the mission.

---

6. Extending the Design Principles Beyond the Cyber-Physical Boundary

NIST SP 800-160, Volume 1, Revision 1 defines 30 principles for trustworthy secure design and explicitly treats systems security engineering as a transdisciplinary activity spanning software, hardware, information technology, physical elements, people, and other system dependencies [5]. The principles are therefore not limited to computer networks. For critical AI infrastructure, they should be applied recursively at three nested levels: inside the computing architecture, across the physical facility, and across the external infrastructure on which the facility depends.

Domain Separation

Domain Separation should partition not only digital trust realms but also physical and infrastructural failure domains. Within a facility, computing halls, management systems, facility controls, maintenance systems, recovery infrastructure, power-control systems, and cooling controls should be separated according to their protection needs. Across facilities, geographic placement and infrastructure sourcing should prevent one regional event or shared dependency from defeating multiple supposedly independent sites. NIST SP 800-160 provides the systems-security-engineering basis for treating security across software, hardware, physical elements, people, and supporting dependencies [5], while RAND demonstrates the value of carrying security constraints through the integrated cyber-physical architecture [3]. The next step is to engineer the external dependencies supporting those functions with the same discipline.

Least Sharing, Least Privilege, and Least Functionality

Least Sharing exposes the security cost of common infrastructure. Shared substations, telecommunications routes, management planes, fuel suppliers, cooling resources, identity services, maintenance contractors, or logistics hubs may create propagation paths and common-mode failure. *Least Privilege* requires that no operator, contractor, controller, credential, or service possess more authority than its mission function requires. *Least Functionality* removes unnecessary services, interfaces, maintenance paths, and latent capabilities that could become useful to an adversary after compromise [4][5]. For national infrastructure

involving AI, these principles may justify deliberate inefficiency, duplicated or independently operated resources whose security value lies precisely in not being shared.

Mediated Access and Distributed Privilege

Mediated Access requires necessary crossings to be explicit, authenticated, authorized, bounded, observable, failure-aware, and non-bypassable [4][5]. At infrastructure scale, crossings include people, equipment, replacement hardware, removable media, software and firmware updates, fuel, water, power, cooling, physical materials, and maintenance activity as well as digital traffic. *Distributed Privilege* then ensures that no single compromised person, credential, control system, or organization holds unilateral authority for selected catastrophic actions. Physical entry to critical zones, changes to facility-control logic, emergency overrides, model-transfer approvals, destructive actions, cryptographic-root changes, and recovery operations may all require independent authorities or threshold mechanisms.

Continuous Protection, Protective Failure, and Protective Recovery

Continuous Protection requires security-relevant protection to remain effective for as long as it is required. *Protective Failure* recognizes that a trustworthy system may need to degrade, isolate, suspend, or shut down selected functions when continued operation would move the mission closer to unacceptable loss. *Protective Recovery* requires restoration from an authenticated trustworthy basis rather than merely a rapid return to capacity [4][5]. For an AI data center, recovery therefore encompasses hardware identity, firmware, models, keys, reference data, physical boundaries, facility controllers, power systems, cooling, monitoring, communications, personnel authority, and the trusted supply chain used to reconstitute the capability.

Redundancy, Diversity, and Compositional Trustworthiness

Redundancy must be evaluated for independence rather than counted. Two generators supplied by the same fuel infrastructure, two fiber paths in the same conduit, two facilities dependent on the same regional grid, or two recovery sites administered by the same identity plane may be redundant in quantity but not in failure behavior. *Diversity* can reduce common-mode failure when it is deliberately designed and its additional complexity is controlled. *Compositional Trustworthiness* requires evidence that the aggregate of computing, facility, human, and infrastructure elements preserves the *system-level security claim*; the trustworthiness of each part in isolation is not sufficient [5].

Table 1. Summary of NIST SP 800-160 design principles applied at infrastructure scale

NIST SP 800-160 Principle	Infrastructure-Scale Application
Domain Separation	Partition computing halls, facility controls, maintenance, recovery, and power/cooling systems by protection need within a facility; separate geographic and infrastructure sourcing across facilities so one regional event or shared dependency cannot defeat multiple sites.
Least Sharing **Least Privilege** **Least Functionality**	Avoid shared substations, telecom routes, management planes, fuel suppliers, cooling resources, identity services, or maintenance contractors that create common-mode paths; limit operator, contractor, and credential authority to mission need; remove unneeded interfaces and maintenance paths.

NIST SP 800-160 Principle	Infrastructure-Scale Application
Mediated Access **Distributed Privilege**	Require every crossing — personnel, hardware, firmware, fuel, water, power, cooling, physical materials, maintenance — to be explicit, authenticated, authorized, and non-bypassable; require independent authorities or thresholds for physical entry, facility-control changes, emergency overrides, and destructive or recovery actions.
Continuous Protection **Protective Failure** **Protective Recovery**	Keep protection effective for as long as required; allow selected functions to degrade, isolate, or shut down rather than moving the mission toward unacceptable loss; restore from an authenticated trustworthy basis covering hardware identity, firmware, models, keys, and the supply chain — not just a fast return to capacity.
Redundancy **Diversity** **Compositional Trustworthiness**	Evaluate redundancy for independence, not count — shared fuel, conduit, grid, or identity plane can make duplicate assets fail together; use deliberate diversity to reduce common-mode failure; require evidence that the aggregate of computing, facility, human, and infrastructure elements preserves the system-level security claim.

7. A National AI Resilience Architecture

Once AI infrastructure is treated as a strategic source of power, architecture should be organized around mission survival rather than facility survival. The key question is not whether a particular data center remains online. It is how much trusted AI capability must remain available, where it must be available, what dependencies it must retain, and how much adversary success the national architecture can absorb before unacceptable mission loss becomes reachable.

A national AI *resilience architecture* should therefore address at least the following engineering obligations:

- Geographic and infrastructure separation sufficient to prevent one regional event, utility failure, or physical attack from disabling multiple critical AI nodes.
- Independent power, communications, cooling, control, administration, and recovery paths commensurate with the consequences of common-mode failure.
- Predetermined protective states and degraded operating modes that preserve the most critical AI functions when full-capacity operation is unsafe or unavailable.
- Trusted reconstitution capability, including reserve compute, validated hardware and firmware, protected model and data baselines, independent credentials, and tested procedures for restoring service.
- Continuously maintained assurance evidence demonstrating that claimed separation, independence, recovery, and loss margins remain valid as the mission, threat, configuration, suppliers, and infrastructure change.

These obligations convert strategic resilience from an aspiration into an *engineering* problem. They also expose tradeoffs that conventional efficiency metrics can obscure. Consolidation lowers cost but can increase common-mode risk. Shared services simplify administration but can increase compromise reach. Highly optimized utilization reduces idle capacity but can eliminate the reserve needed for wartime reconstitution. The relevant optimization criterion for a strategic AI capability is therefore not peacetime efficiency alone. It is *mission effectiveness* under adversity.

8. Compromise Margin Becomes Strategic Margin

EfC introduces the concept of *compromise margin*: the amount of adversary success, loss of trust, or component compromise that can occur before an unacceptable-loss state becomes reachable [4]. For national AI infrastructure, that idea can be elevated into a *strategic margin*. The questions become concrete: How many facilities can be lost? How much compute can be unavailable? Which regions can be isolated? How many operators can be compromised? Which suppliers can fail? How long can power or communications be disrupted? Which model replicas can be corrupted? How rapidly can trusted capacity be reconstituted?

These are not merely *continuity-of-operations* questions. They are security architecture questions because the answers determine whether an adversary can convert a successful local attack into a national strategic effect. The architecture should make those thresholds explicit and should identify the independent barriers that prevent each transition toward unacceptable loss.

This also changes how redundancy is measured. A nation that possesses ten major AI facilities but depends on one semiconductor supply chain, one class of accelerator, a handful of electrical regions, common software orchestration, shared cloud control planes, or common maintenance authorities may have far less strategic margin than raw data-center counts imply. Conversely, deliberate diversity and distributed reserve capacity may provide a larger strategic margin even when total nominal compute is smaller.

This argument does not resolve, and should acknowledge, a real economic tension. Frontier AI compute has concentrated rather than dispersed in recent years. This is driven by capital intensity, power availability, chip supply, and the efficiency gains of large co-located clusters. Deliberate dispersion, independent infrastructure, and reserve capacity of the kind recommended here carry real capital and operating costs and may run against the market incentives currently shaping where and how frontier compute is built. A credible national resilience architecture must therefore treat *strategic margin* as a cost to be deliberately funded and justified, not an emergent property of ordinary infrastructure investment, and should identify which portion of national AI capacity warrants this premium rather than assuming it applies uniformly.

9. The Adversary Will Attack the System, Not the Organizational Chart

One of the greatest barriers to infrastructure-scale protection is institutional fragmentation. Cybersecurity teams protect computing systems; physical-security organizations protect buildings; utilities provide power; facilities organizations manage cooling; manufacturers provide processors; telecommunications companies provide connectivity; contractors perform maintenance; supply-chain organizations acquire replacements; government agencies protect critical infrastructure; and military organizations plan for wartime continuity. Each organization may satisfy its own requirements while no organization owns the complete security argument for the national AI capability.

EfC explicitly warns that containment can be defeated by acquisition choices, supply chains, maintenance practices, human authority, enabling systems, shared services, and dependencies outside the nominal operational architecture [4]. For national-security systems, the assumed adversary may combine cyber, physical, intelligence, insider, and supply-chain operations over long periods [4]. A defense model partitioned by organizational responsibility therefore gives the adversary seams to exploit.

The unit of analysis must match the unit of attack. For strategically consequential AI, the system of interest should include every dependency whose failure or malicious use can invalidate the mission-level security claim. This is consistent with NIST SP 800-160's foundational systems perspective: *security is an emergent property of the system* and must be engineered through the complete life cycle, not added as an operational overlay [5].

10. A Loss-Driven National Security Engineering Method

The appropriate methodology begins not with a control catalog or a list of facilities to harden, but with unacceptable national losses. EfC expresses the chain as:

> ***Mission -> Unacceptable Loss -> Hazardous System Condition -> Security Constraint -> Architecture -> Mechanism -> Evidence***

Applied to critical AI infrastructure, that chain provides a repeatable national-security engineering method.

- **Mission:** preserve strategically essential U.S. AI capability during conflict and sustained national-level adversity.
- **Unacceptable Loss:** loss, corruption, compromise, or prolonged denial of AI capabilities required for critical national or military missions.
- **Hazardous Conditions:** excessive concentration of compute; common-mode infrastructure dependencies; loss of power, cooling, communications, or trusted recovery; adversarial control of facility systems; corrupted model or data baselines; or inability to replace failed components.
- **Security Constraints:** no single facility, infrastructure dependency, administrative authority, supplier, geographic region, or selected compromise event may be sufficient to produce the identified national loss.
- **Architecture and Mechanisms:** geographically and administratively separated AI nodes, independent infrastructure, bounded interfaces, distributed privilege, diversity, reserve capacity, protective failure, anomaly detection, and authenticated recovery.
- **Evidence:** modeling, infrastructure-dependency analysis, fault injection, physical and cyber testing, adversarial exercises, red teaming, recovery demonstrations, continuously maintained assurance cases, and supply-chain assessments.

This methodology should be applied across the NIST system life cycle. *Mission analysis* establishes the losses that justify extraordinary protection. *Requirements* define survival and degradation obligations. *Architecture* allocates trust, failure isolation, and independence. *Design* and *Implementation* realize those properties. Integration determines whether individually trustworthy secure subsystems create new common dependencies when combined. *Verification* and *Validation* generate evidence. *Transition* ensures the fielded facility and its external environment preserve the assumptions used during analysis. *Operation* monitors the invariants. *Maintenance* must prevent temporary access, replacement equipment, vendor tools, and emergency procedures from becoming a parallel attack architecture. *Disposal* must preserve the same security intent [4][5].

11. From Digital Fort Knox to a Resilient AI Arsenal

The SIDC demonstrates how extraordinary protection can be concentrated around exceptionally valuable AI assets [3]. The center-of-gravity problem suggests the next architectural evolution: a resilient national AI arsenal composed of highly protected facilities whose collective architecture and design assume that some sites, dependencies, and components may be attacked successfully. Such an architecture and design would not abandon hardened facilities. Some AI assets may warrant SIDC-like isolation, physical protection, formal analysis, and minimality. But the national strategy would avoid making any one facility, administrative domain, region, utility, or recovery mechanism indispensable to the mission. The design objective would be to preserve sufficient trusted capacity through dispersion, independence, compartmentalization, reserve capability, and authenticated reconstitution.

Do not design the national AI infrastructure around the assumption that every fortress survives. Design it so missions survive even when some fortresses do not.

This is the infrastructure-scale expression of *Engineering for Compromise*. It also reflects the logic of center-of-gravity analysis: **protect the source of power by understanding its critical requirements and reducing the vulnerabilities through which an adversary could obtain disproportionate strategic effect**.

12. Conclusion

The strategic importance of artificial intelligence is changing what the United States must treat as critical infrastructure. If frontier AI becomes deeply embedded in intelligence, military operations, cyber defense, scientific discovery, economic competitiveness, industrial adaptation, and national decision-making, then the computational infrastructure that generates those capabilities can no longer be treated as ordinary commercial infrastructure.

Center-of-gravity doctrine provides the strategic lens. Sun Tzu provides the enduring logic of attacking weakness, dependency, and advantage rather than simply striking strength. RAND demonstrates that extraordinary protection for critical AI assets can be engineered through a purpose-built cyber-physical architecture. NIST SP 800-160 provides the principles for structuring separation, privilege, sharing, mediation, failure, recovery, redundancy, trustworthiness, and system control. *Engineering for Compromise* supplies the operational premise that ties them together: do not assume that every system element remains trustworthy or survives; engineer the system so that the mission *can*.

For America's most consequential AI capabilities, those ideas must now move beyond processors, software, networks, and the traditional cyber-physical boundary. They must encompass the building, electrical power, cooling, communications, people, maintenance, logistics, suppliers, geographic distribution, recovery systems, and every other dependency whose loss could become mission loss. The future AI data center should therefore be conceived not merely as a trustworthy secure computing facility but as one node in a resilient national capability designed to continue generating strategic power while under attack.

These are not abstract design goals; they are near-term engineering obligations for those who already own a piece of this problem, including the agencies responsible for classified and defense-relevant compute, the entities that designate and oversee critical infrastructure, and the facility/system operators and suppliers whose consolidated footprints already concentrate much of the nation's frontier AI capacity. Each should be able to state, in the terms of Section 8, how much of its *compromise margin* has already been spent.

The decisive question is not whether an adversary can attack, compromise, or destroy part of the AI infrastructure. It is how much of that infrastructure can be lost, compromised, or destroyed before the national mission fails.

If advanced AI becomes a center of gravity in future conflict, the strength of the United States will depend not only on possessing superior models and superior compute, but on whether the nation has engineered the infrastructure behind those capabilities with sufficient separation, independence, resilience, and assurance that *local* defeat cannot be converted into *strategic* defeat.

Acknowledgments

The author gratefully acknowledges the support provided by Dartmouth College's Institute for Security, Technology, and Society (ISTS). The resources provided by the ISTS contributed to the research and development of this paper. The views and conclusions expressed herein are those of the author and do not necessarily represent the views of Dartmouth College.

References

[1] Joint Chiefs of Staff, Joint Publication 5-0, Joint Planning, 1 December 2020.

[2] Sun Tzu, The Art of War, translated by Lionel Giles, 1910. Project Gutenberg, eBook No. 132.

[3] S. F. Comer et al., Secure Inference Data Centers: A Vertically Integrated Strategy for Security Engineering, RAND Corporation, RR-A4827-1, 2026. https://www.rand.org/pubs/research_reports/RRA4827-1.html

[4] R. Ross, "Engineering for Compromise: Designing Systems That Remain Trustworthy Secure Under Adversity," August 2026. https://www.linkedin.com/posts/ronrossecure_engineering-for-compromise-activity-7494928506523414528-vtMr?utm_source=share&utm_medium=member_desktop&rcm=ACoAACOnE_0B4t1GmIXMRt2blqCmFsyrRooSDRc

[5] R. Ross, M. Winstead, and M. McEvilley, *Engineering Trustworthy Secure Systems*, NIST Special Publication 800-160, Volume 1, Revision 1, National Institute of Standards and Technology, November 2022. https://csrc.nist.gov/pubs/sp/800/160/v1/r1/final