Adequate Security: An Engineering Answer to an Impossible Question

Dr. Ron Ross

CEO, RONROSSECURE, LLC
ISTS Fellow, Dartmouth College

Abstract

Security discussions often begin with threats, controls, products, frameworks, or compliance requirements. But the more fundamental engineering question is simpler: How much security does a system actually need? NIST Special Publication 800-160, Volume 1 provides a disciplined answer through its concept of an adequately secure system. Adequate security is neither perfect security nor the maximum security technically possible. It is an evidence-based judgment that the system achieves at least a minimum tolerable level of security and is as secure as reasonably practicable when security is balanced with the system's other cost, schedule, technical-performance, operational, and mission constraints [1]. This article develops that concept as a practical engineering framework. It explains why security is inherently contextual, why acceptable security cannot be reduced to a universal control count, why secure and insecure system states matter, and why the answer must evolve as the system, mission, threat environment, and available engineering options change. The central proposition is that the right amount of security is the amount that can be justified by stakeholder consequences, system context, engineering trade space, and evidence—not by intuition, compliance alone, or the pursuit of an unattainable absolute.

The Question Behind Every Security Decision

Every organization eventually confronts the same question, whether it asks it explicitly or not: How much security is enough? The question appears whenever leaders decide how much to spend, engineers decide how much protection to build into an architecture, operators decide what residual exposure they can tolerate, or mission owners decide whether a system is trustworthy enough to use.

The usual answers are unsatisfying. "More security" is not a requirement. "Best practice" does not establish that a particular system is sufficiently protected. Compliance can demonstrate that prescribed obligations were addressed, but it does not by itself establish that the resulting system is secure enough for its mission. And perfect security is not an engineering objective because complex systems operate under uncertainty, constraints, and adversity.

NIST SP 800-160, Volume 1 reframes the problem around the concept of an *adequately secure* system [1]. That framing is important because it replaces the search for an absolute level of security with a reasoned judgment about what security performance is necessary, tolerable, practicable, and demonstrable for a specific system and its stakeholders. The goal is not maximum security, but security that is sufficient for the consequences at stake—and defensible by evidence.

***"This whole economic boom in cybersecurity seems largely to be a consequence of poor engineering."* [2]**

Security Is an Ideal — Adequate Security Is an Engineering Judgment

NIST begins with an idealized conception of a secure system. Such a system must continue to deliver its required capability despite intentional and unintentional adversity; enforce constraints so that only desired

behaviors and outcomes occur; and enforce rules so that only authorized human-to-machine and machine-to-machine interactions and operations are allowed [1]. These three characteristics establish the direction of travel.

The difficulty is that real systems cannot normally achieve the ideal with certainty. NIST explicitly recognizes a gap between the ideal secure system and the security performance that a real system can dependably achieve. Security is asymmetric: insecurity can be observed, but observing a system does not allow us to prove, in an absolute sense, that an arbitrary system is secure. Uncertainty therefore remains [1].

That gap is where engineering judgment operates. The question becomes not whether a system is perfectly secure, but whether the security performance that has been achieved is adequate for the system's purpose, consequences, operational environment, and constraints.

The Two Tests of Adequate Security

NIST establishes two complementary conditions. A system judged adequately secure must meet minimum tolerable levels of security and must also be as secure as reasonably practicable (ASARP) [1]. Neither condition is sufficient by itself.

Minimum Tolerable Security

The first test establishes a floor. Below the minimum tolerable level, the system is considered insecure. The floor may be informed by experience, analysis, or both [1]. In engineering terms, this means that some residual conditions simply cannot be accepted, regardless of convenience, cost savings, schedule pressure, or the fact that other controls are present.

The minimum tolerable level should therefore be connected to stakeholder protection needs and the consequences of asset loss. A system whose failure can produce catastrophic mission loss cannot rationally use the same minimum threshold as a system whose failure produces minor inconvenience.[1] The threshold is not a generic maturity score. It is a boundary between security performance that stakeholders can tolerate and performance that they cannot.

As Secure as Reasonably Practicable

Meeting the floor is only the beginning. NIST defines ASARP as the point at which an incremental security improvement would (1) require a disproportionate deterioration in other system objectives such as cost, schedule, or technical performance; (2) violate system constraints; or (3) require an unacceptable operational concession [1].

This is a powerful idea because it rejects two extremes. The first is under-engineering: stopping as soon as the minimum threshold is crossed. The second is security maximalism: continuing to add protection regardless of what it does to mission effectiveness, usability, availability, performance, affordability, or schedule. Adequate security occupies the defensible region between those extremes.

Minimum tolerability establishes the floor. ASARP asks whether stopping there can actually be justified.

---

[1] This was the original intent behind the development of FIPS 199 [3]. The level of security must be linked to the *mission impact* (i.e., the *adverse consequences* when something in the system goes wrong).

Security Is a Trade—But Not Just Another Trade

A system is built to deliver *capability*, not merely to be secure. Security therefore exists within a broader engineering trade space that includes cost, schedule, technical performance, mission effectiveness, safety, reliability, maintainability, and other stakeholder objectives. NIST states that the judgment of adequate security requires an evidence-based determination that security performance has been optimized against the other performance objectives and constraints [1].

That does not mean security can always be traded away. The minimum tolerable level constrains the trade space. Once a proposed trade pushes security below that boundary, the system has crossed from a difficult engineering compromise into an unacceptable condition. Above the boundary, however, engineers can evaluate whether additional protection produces sufficient benefit to justify its effect on other objectives. ASARP should not be interpreted as permission to exchange security casually for cost or schedule—but as a reminder that an engineered system must satisfy multiple objectives simultaneously. The final adequacy judgment is about where, within those constraints, the system can justifiably operate.

There Is No Universal Answer

The most consequential implication in NIST SP 800-160 is that adequate security is context-dependent. NIST states that the scope of conditions relevant to security and the acceptable level of security are specific to stakeholder needs [1]. The phrase “How much security do I need?” therefore cannot be answered without first asking: Security for what system, protecting which assets, against which unacceptable consequences, under what forms of adversity, for which mission, in which states and operating environments, and subject to what constraints?

A corporate collaboration platform, an industrial control system, a medical device, a weapons system, an AI data center, and a national-security command-and-control capability may all require security, but *adequacy* means something different in each case. The consequences of failure differ, the adversities differ, the allowable degraded modes differ, the acceptable uncertainty differs, and the evidence needed to support trustworthiness and confidence differs.

This is why security requirements should emerge from mission and stakeholder protection needs rather than from a predetermined inventory of security controls or technologies. A control can be useful without being sufficient. A system can satisfy a checklist and still fail the adequacy test if the resulting security performance does not match the adverse consequences and the stakeholder expectations.

Adequate Security Is About System Behavior

NIST SP 800-160 moves the discussion from static protection to system behavior. Systems operate in states and transition between states. Those states may reflect normal operation, startup, shutdown, recovery, training, maintenance, peacetime, wartime, environmental conditions, or other operational circumstances. From a security perspective, a state or transition may be secure, insecure, or indeterminate [1].

The ideal is straightforward: begin in a secure state and ensure that every state transition ends in the same or another secure state, with the transition itself also secure [1]. Real systems complicate that ideal. Faults occur, components fail, attackers succeed, and operators make mistakes. The system may enter a state whose security is uncertain or known to be unacceptable. Adequacy therefore depends not only on preventing bad states but on recognizing and controlling movement among states.

Protective Failure and Protective Recovery

The state-transition model leads directly to two design principles emphasized in NIST SP 800-160: *Protective Failure* and *Protective Recovery*. Protective Failure requires the ability to detect that the system is in an insecure state (or is transitioning toward one) and to respond so that the failure does not propagate. In

some circumstances, that response may include a secure halt followed by protected recovery into a reconstituted, reconfigured, or alternative secure state [1].

Protective Recovery addresses the complementary problem: how to move securely from an insecure state to a secure state, or at least to a less insecure state. Recovery can include degraded capability or operating modes that constrain further transitions while preserving what the mission most needs [1].

This changes the meaning of "enough security." A system may be adequately secure even though some attacks or failures can succeed, provided the architecture can detect, contain, respond, recover, and continue at an acceptable level without allowing those events to produce unacceptable loss. Conversely, a system with extensive preventive controls may be inadequate if one successful failure moves it directly into catastrophic loss.

Adequate security does not require that nothing bad ever happens. It requires that the system's behavior under adversity remains within defensible bounds.

Adequacy Is Dynamic, Not Permanent

NIST makes another point that is easy to overlook: minimum tolerable security and the meaning of ASARP need not remain fixed throughout the system life cycle. Operational experience, lessons learned, changes in technology, and new information (including threat intelligence) can justify raising the bar [1].

A security decision that was reasonable at deployment may become indefensible after the system is in an operational state. A new attack technique may expose an assumption, a lower-cost protection mechanism may become available, a system may be assigned a more critical mission, interconnections may expand, asset value may increase, or an operating environment may become more hostile. Any of these changes can alter the balance between *tolerability* and *practicability*.

Adequate security should therefore be treated as a dynamic, maintained *engineering judgment*, not a static certificate awarded once. The evidence supporting the judgment must evolve with the system and its environment.

Evidence Is What Turns Opinion into an Engineering Judgment

NIST explicitly characterizes adequate security as an evidence-based determination [1]. That qualifier matters. Without evidence, statements such as "the risk is acceptable," "the system is secure enough," or "additional security is too expensive" are assertions, not engineering conclusions.

A defensible adequacy argument should show (1) what stakeholders need protected; (2) what consequences are unacceptable; (3) what minimum security performance must be achieved; (4) what forms of adversity and system states were considered; (5) what constraints shaped the solution; (6) what additional security options were evaluated; (7) why further improvement is or is not reasonably practicable; and (8) what evidence demonstrates that the implemented system behaves as claimed.

The evidence may include analysis, modeling, test results, verification, validation, operational data, fault injection, penetration testing, red-team results, recovery demonstrations, architecture analysis, and other forms of assurance appropriate to the system. The *rigor* should be commensurate with the consequences of failure.

A Practical Adequate-Security Decision Framework

The concept of adequate security can be translated into a practical sequence for engineering and governance decisions. The sequence is not presented in NIST SP 800-160 as a separate checklist; it is a synthesis of the adequacy logic into an operational decision framework:

1. **Define the required capability.** State what the system must continue to accomplish, including under relevant adversity.
2. **Define unacceptable consequences.** Identify the losses stakeholders cannot tolerate and the assets and mission outcomes associated with them.
3. **Establish the minimum tolerable security level.** Define the security-performance floor below which the system is considered insecure.
4. **Identify relevant states and transitions.** Determine normal, degraded, maintenance, recovery, emergency, and other states that matter to the security judgment.
5. **Engineer protection and recovery behavior.** Use constraints, architecture, mechanisms, Protective Failure, and Protective Recovery to keep behavior within acceptable bounds.
6. **Explore the trade space.** Evaluate feasible additional security improvements against cost, schedule, technical performance, operational needs, and system constraints.
7. **Determine ASARP.** Identify the point at which further incremental security improvement would impose disproportionate degradation or unacceptable concessions.
8. **Substantiate the judgment.** Produce evidence sufficient to justify that the minimum tolerable level is met and the ASARP conclusion is credible.
9. **Revisit the decision.** Reassess adequacy as the mission, system, threats, technology, evidence, and operating environment change.

What Adequate Security Is Not

The concept becomes clearer when contrasted with several common substitutes. Adequate security is not zero risk, because uncertainty and residual adversity remain. It is not maximum security, because systems must satisfy other objectives and constraints. It is not merely compliance, because conformance to a set of prescribed requirements does not automatically establish that stakeholder protection needs are satisfied. It is not a security control count, because controls differ in relevance, effectiveness, integration, cost and assurance. And it is not a one-time system authorization decision, because the basis for adequacy can change over the life of the system.

Most importantly, adequate security is not synonymous with “whatever level of risk management accepts.” The judgment must be constrained by a minimum tolerable security floor and supported by evidence that the system has been made *as secure as reasonably practicable*. Acceptance without that engineering basis can merely institutionalize exposure.

Why This Matters More in the Modern Threat Space

The concept of adequate security becomes more important as systems become more interconnected, software-defined, autonomous, cyber-physical, and dependent on complex supply chains and shared infrastructure. Frontier AI, cloud platforms, operational technology, critical infrastructure, and national-security systems can concentrate enormous capability and consequence in system architectures whose dependencies extend far beyond a traditional cybersecurity boundary.

In such environments, asking whether individual components have enough controls is the wrong level of abstraction. The adequacy question belongs at the system level: Can the complete system continue to deliver the required capability, enforce the necessary constraints, permit only authorized interactions, and recover from relevant adversity with a level of confidence commensurate with the consequences?

From "How Much Security?" to "How Much Confidence?"

There is a deeper and more pervasive question beneath the title of this article. Once stakeholders define the security performance they need, they must decide how much confidence they require that the system will actually deliver it. A low-consequence (i.e., low-impact) system may justify modest evidence. A system whose failure could produce severe or catastrophic consequences (i.e., high-impact) demands substantially stronger assurance.

This is where the concept of adequate security connects naturally to the broader NIST SP 800-160, Volume 1 treatment of *trustworthiness* and *assurance*. Security performance and *confidence* in that performance are related but distinct. A design can claim excellent security properties, but if the evidence supporting those claims is weak, the system may not be sufficiently trustworthy for the mission. The adequacy decision must therefore consider both what the system is expected to do and how convincingly that expectation has been substantiated. The consequences of being wrong should therefore determine not only how much security is engineered into the system, but how much *rigor* is required to substantiate that security.

Conclusion

How much security do I really need? A system needs enough security to meet a minimum tolerable level determined by its stakeholders and consequences, and it should then be made as secure as reasonably practicable within the realities of the engineering trade space [1]. Moreover, the answer is necessarily system-specific. It depends on the capability being protected, the losses that matter, the adversity the system must withstand, the states and transitions through which it operates, the constraints under which it is engineered, and the evidence available to justify confidence. It also changes over time.

The right amount of security is not the most security we can imagine. It is the level of security we can justify—technically, operationally, economically, and through evidence—against the consequences we cannot afford.

The practical implication is straightforward. Stop asking whether the organization has implemented enough security controls. Instead, ask whether (1) the security performance of the system is above the minimum tolerable threshold, (2) whether further improvement remains reasonably practicable, and (3) whether the evidence is strong enough to justify that judgment. That is a question engineering can answer—and one that every consequential system should be required to answer.

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] R. Ross, M. McEvilley, and M. Winstead, *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

[2] C. E. Landwehr, "We Need a Building Code for Building Code," Communications of the ACM, Vol. 58, No. 2, pp. 24–26, February 2015. doi:10.1145/2700341.

[3] National Institute of Standards and Technology (2004) *Standards for Security Categorization of Federal Information and Information Systems*. (U.S. Department of Commerce, Washington, D.C.), Federal Information Processing Standards Publication (FIPS) 199. https://csrc.nist.gov/pubs/fips/199/final