Defining Protect Surfaces in Federal Zero Trust Implementation

Don Yeske

National Defense University

College of Information and Cyberspace

FINAL VERSION as of 27 March 2026

This paper or presentation is my own work. Any assistance I received in its preparation is acknowledged within the paper or presentation, in accordance with academic practice. If I used data, ideas, words, diagrams, pictures, or other information from any source, I have cited the sources fully and completely in footnotes and bibliography entries. This includes sources that I have quoted or that I have paraphrased. Furthermore, I certify that this paper or presentation was prepared by me specifically for this class and has not been submitted, in whole or in part, to any other class in this University or elsewhere, or used for any purpose other than satisfying the requirements of this class, except that I am allowed to submit the paper or presentation to a professional publication, peer reviewed journal, or professional conference. This is not a draft, and is submitted for grading to satisfy in part the requirements for this course. In typing my name following the word 'Signature', I intend that this certification will have the same authority and authenticity as a document executed with my hand-written signature.

Signature DONALD E YESKE

Disclaimer: This paper represents the views and opinions of the author. This paper does not represent official NDU, DoD, USG policy or position.

Acknowledgements

In the interest of full transparency, the author wishes to acknowledge that much of this paper is drawn from personal experience, which includes leading zero trust architecture efforts for the Department of the Navy (2021 through 2023) and the Department of Homeland Security (2023 through 2025). In this capacity, the author reviewed several documents referenced in this paper as they were being developed, worked with many people whose work is cited in this paper, and contributed directly to the development of zero trust knowledge and best practice across the Federal government. Much of this paper can be fairly read as criticism of this work, and it is that. However, where mistakes have been made, they have been made in the process of discovery. Progress in any field requires making mistakes, embracing them, and learning from them. Everyone involved in zero trust implementation from across the Federal government, the author included, has proceeded always with the best of intentions, learning a great deal along the way. Accordingly, this paper should be read not as indictment, but as incitement to continue progress.

Beyond this, the author wishes to acknowledge the contributions of multiple colleagues, academic advisors, and faculty at National Defense University in reviewing, challenging, and improving the ideas put forward in this paper, whose critical analysis has proven invaluable. Ideas can only ever become as good as the people willing to actively challenge them.

Finally, the author wishes to acknowledge the patience and support of family through the yearlong process of writing this paper in evenings and weekends—time that would otherwise have been spent with them. Whether the result is worth the investment is for others to judge, but the result would not have been possible without their investment of time and space in which to research, write, and revise this report.

Defining Protect Surfaces in Federal Zero Trust Implementation

In May 2021, Section 3 of Executive Order 14028, *Improving the Nation's Cybersecurity*, kicked off the Federal government's efforts to implement zero trust architecture (Biden 2021). This order remains in force, and the current administration reaffirmed its commitment to zero trust in *President Trump's Cyber Strategy for America* less than a month ago (Trump 2026, 5). While the last two administrations agree on little else, both describe implementing zero trust architecture across government as both strategically necessary and fundamentally transformative. As stated by the former, "[i]ncremental improvements will not give us the security we need; instead, the Federal Government needs to make bold changes and significant investments in order to defend the vital institutions that underpin the American way of life" (Biden 2021, 1).

Zero trust architecture is necessary and transformative because it corrects a fundamental flaw in the cybersecurity approach that preceded it. Prior to zero trust, cybersecurity focused mainly on the defense of organizational perimeters, such as network boundaries. Actors outside these perimeters were considered untrustworthy and attempts to access organizational networks were treated with high degrees of scrutiny; once actors gained access to the network, they were generally trusted, and their actions were subject to far less scrutiny. The creator of zero trust architecture, John Kindervag, first described the problem with this approach in his seminal paper "No More Chewy Centers: Introducing the Zero Trust Model of Information Security," positing that networks could no longer provide the basis of trust because network breaches had become inevitable, and bad actors could come from within or outside networks, and therefore all access attempts, within and outside of networks, must be treated as equally suspect (Kindervag 2010).

Unfortunately, Federal directives and guidelines underpinning zero trust architecture implementation broadly suffer from the same fundamental flaw because of a critical omission.

Directed actions to implement zero trust have focused on the attack surface—the organizational perimeter—whether by default in not stating a different scope, or explicitly. At the same time, authorities have not directed agencies to define protect surfaces: specific elements within their perimeters that zero trust architectures must protect. The tasks given to Federal agencies are consistent with zero trust architecture principles, but for lack of defined protect surfaces, have not had specific implementation targets. By seeking to protect everything, whether by default or because they were explicitly directed to protect everything, Federal agencies have effectively focused zero trust architecture on protecting nothing. The result is large amounts of activity and investment with limited tangible progress in protecting any defined set of mission outcomes.

This paper begins by defining the protect surface in the context of zero trust architecture. Relying on this definition, this paper critically examines the directives that have been issued to government agencies, noting whether and to what extent agencies have been directed to identify protect surfaces and measure their defenses. Next, this paper reviews publicly reported results of zero trust architecture implementation across the Federal government to date, showing that broad improvements have happened, but little progress has been made in defining, and intentionally designing and implementing, zero trust architectures to protect specific mission outcomes. Finally, this paper proposes actions that the Federal government could direct to better focus, accelerate, and more meaningfully measure zero trust architecture implementation.

Protect Surfaces in the Context of Zero Trust Architecture

The term ‘protect surface’ occurs rarely in Federal zero trust architecture references. Accordingly, one might assume that this concept is not a part of Federal zero trust architecture. This would not be accurate, but to show that, one must first have an objective definition in mind. What follows is a definition of ‘protect surface’ that applies whether the term is used or not.

Zero Trust Architecture

National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207 defines zero trust architecture for the Federal government. It begins with this premise:

> A typical enterprise's infrastructure has grown increasingly complex. A single enterprise may operate several internal networks, remote offices with their own local infrastructure, remote and/or mobile individuals, and cloud services. This complexity has outstripped legacy methods of perimeter-based network security as there is no single, easily identified perimeter for the enterprise. Perimeter-based network security has also been shown to be insufficient since once attackers breach the perimeter, further lateral movement is unhindered.
>
> This complex enterprise has led to the development of a new model for cybersecurity known as "zero trust" … Zero trust security models assume that an attacker is present in the environment and that an enterprise-owned environment is no different—or no more trustworthy—than any nonenterprise-owned [sic] environment. In this new paradigm, an enterprise must assume no implicit trust and continually analyze and evaluate the risks to its assets and business functions and then enact protections to mitigate these risks. In zero trust, these protections usually involve minimizing access to resources (such as data and compute resources and applications/services) to only those subjects and assets identified as needing access as well as continually authenticating and authorizing the identity and security posture of each access request. (Rose, et al. 2020, 1)

Conceptually, zero trust architecture differs from traditional cybersecurity architecture in that zero trust assumes a breach is inevitable, and therefore, network enclave perimeters cannot form the basis of trust in an actor. On this point, there is no disagreement or ambiguity.

The remainder of NIST SP 800-207 is a technical explanation of how organizations accomplish the goals set forth in the opening premise, including what elements comprise a zero trust architecture and how these elements work together, high level architecture patterns that organizations might follow, and practical guidance regarding how to implement zero trust architecture in both green-field (new) and brown-field (existing) technology environments.

The Attack Surface

The NIST Computer Security Resource Center glossary defines the term 'attack surface' as "[t]he set of points on the boundary of a system, a system component, or an environment where an attacker can try to enter, cause an effect on, or extract data from, that system, component, or environment" (NIST n.d.). There are two important things to note about this definition in the context of zero trust architecture.

First: The attack surface of an organizational network or environment is the same perimeter that zero trust architecture defines as being fundamentally and practically indefensible. These things are one and the same. Zero trust is based on the premise that organizations cannot adequately defend attack surfaces, and as such, a new approach is needed. Organizations acting in keeping with zero trust principles "assume that an attacker is present in the environment" (Rose, et al. 2020, 1) and all further actions and decisions proceed from that assumption.

Second: The attack surface is not easily quantifiable. The opening premise of the Federal zero trust reference architecture directly states that "there is no single, easily identified perimeter for the enterprise" (Rose, et al. 2020, 1). This is true for the reasons stated there—i.e., that organizations are complex—but it is also true because the attack surface is a fast-moving target. As Kindervag observes: "The attack surface is like the universe—it's constantly expanding" (Beddell 2024). According to cybersecurity firm Darktrace, reasons why the attack surface constantly grows and changes rapidly are numerous and varied, including such causes as continuous software development and delivery processes, cloud adoption and cloud migrations, risks from the information technology (IT) supply chain (also a central topic of EO 14028), frequent software patches and updates, rapidly evolving attackers and exploit patterns, and changes to remote work practices resulting from the COVID-19 pandemic (Thiele 2021).

Protect Surfaces

In zero trust architecture, a protect surface is the logical inverse of the attack surface. Instead of describing every way in which an environment might be susceptible to attack, a protect surface describes "the area or portion of an organization's technology environment that the Zero Trust policy implementation protects" (Cloud Security Alliance 2024, 9).

A 2022 report to the President on *Zero Trust and Trusted Identity Management* from the National Security Telecommunications Advisory Committee (NSTAC)—one of only a few government documents to define and use this term—cites the protect surface as a key concept, and defines the relationship between protect surfaces and zero trust environments by stating that "[e]ach zero trust environment will have multiple protect surfaces" and "[e]ach protect surface contains a single data, applications, assets, and services (DAAS) element" (NSTAC 2022, 6).

The same report further describes defining protect surfaces as the first step in an industry-developed five-step process for zero trust implementation, observing that "[t]he scope of zero trust can be large and all-encompassing, so breaking the process into smaller and more manageable components is important" (NSTAC 2022, 7). The report illustrates this process as follows:

Figure 1: Five-Step Process for Zero Trust Implementation (NSTAC 2022, 7)

In summary, this process calls for organizations to define protect surfaces, and then to build out their zero trust architectures one protect surface at a time. Each protect surface is mapped, added to the existing zero trust architecture, and then added to the zero trust policy already in place, with access policy determined according to which entities must have access the protect surface, under what conditions. Over time, zero trust architectures grow and mature by incorporating additional protect surfaces (NSTAC 2022, 7-8).

Cloud Security Alliance, a global nonprofit cybersecurity credentialing organization, fleshes out these concepts further by prescribing an order in which protect surfaces should be added to an organization's zero trust architecture. This guidance is illustrated as follows:

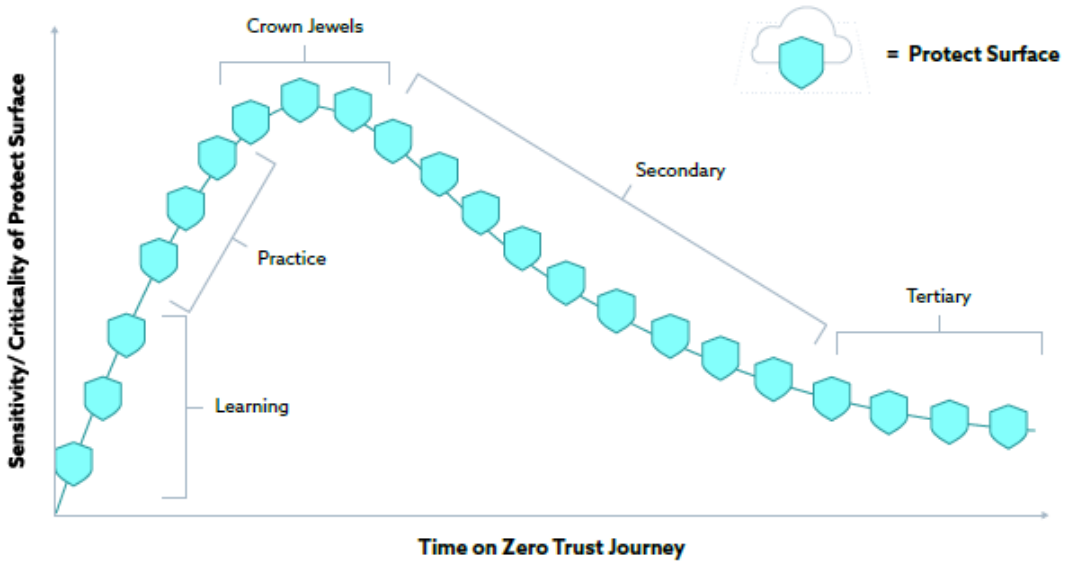

Figure 2: Iterative Execution of the 5-Step Process (Cloud Security Alliance 2024, 13)

Industry best practice, as outlined by Cloud Security Alliance, is for organizations to start building their zero trust architectures initially by incorporating low sensitivity, low criticality protect surfaces as a learning exercise. As the organization's zero trust architecture matures, higher-complexity, higher-criticality protect surfaces should be added. As stated in their report:

> Starting with simpler Protect Surfaces can allow valuable insights to be gained safely and then applied to more complex and higher-risk Protect Surfaces. Especially during the initial phases of their Zero-Trust journey, organizations are advised to implement one Protect Surface at a time. This iterative approach ensures that lessons learned from each implementation can be applied systematically to subsequent Protect Surfaces, facilitating a more informed and effective overall implementation. (Cloud Security Alliance 2024, 14)

NIST SP 800-207 (the Federal reference architecture for zero trust) does not contain the term 'protect surface', but it does provide the same guidance—including the same basic process and the same guidance for how to implement it. Specifically, the final section of the main body of the document provides guidance for migrating to zero trust architecture, including steps for introducing zero trust architecture to a perimeter-based architected network. That guidance calls for organizations to undertake "a survey of assets, subjects, data flows, and workflows" (Rose, et al. 2020, 37) and then to "identify and rank the business processes, data flows, and their relation in the missions of the agency" in order to identify what the document terms "ZTA Candidates" (Rose, et al. 2020, 39). At this point, organizations are given the following guidelines:

> After the asset or workflow is identified, identify all upstream resources (e.g., ID management systems, databases, micro-services), downstream resources (e.g., logging, security monitoring), and entities (e.g., subjects, service accounts) that are used or affected by the workflow. This may influence the candidate choice as a first migration to ZTA. An application/service used by an identified subset of enterprise subjects (e.g., a purchasing system) may be preferred over one that is vital to the entire subject base of the enterprise (e.g., email). (Rose, et al. 2020, 39)

After putting 'ZTA Candidates' in order according to this guidance, organizations are advised to build zero trust policies based on access patterns, design or select their zero trust architecture solutions using what they've learned, implement them in the prescribed manner and order, and add each candidate solution to enterprise policy enforcement and monitoring incrementally (Rose, et al. 2020, 39-41).

This guidance from NIST SP 800-207 constitutes the same process in the same order (identifying protect surfaces, mapping transaction flows, building the architecture, building the policy, and monitoring and maintaining the architecture) and includes the same guidelines for protect surface prioritization and sequencing (starting with narrower and simpler protect surfaces before moving to those that affect the entire enterprise). Federal zero trust reference architecture and industry-provided guidelines and best practices are functionally identical in these regards, even though the word 'protect surface' rarely appears in Federal references on zero trust and does not appear at all in NIST SP 800-207.

In summary, an organization's protect surfaces act as the targets for zero trust architecture implementation. Protect surfaces provide a way to break up the organization's zero trust journey into logical and meaningful steps, a basis for designing the architecture, and a way to expand it incrementally, enabling organizations to make progress while controlling both costs and risks. Protect surfaces are the objects of zero trust architecture protection. Conversely, attack surfaces cannot be easily identified or defended, and this is why zero trust architecture is necessary.

Zero Trust Directives and What They Direct

With these definitions and observations in mind, we can demonstrate that Federal agencies have been directed to do some things that both ignore, and in many cases contradict, fundamental zero trust architecture principles. Specifically, directives that mandate actions targeting the entire enterprise or whole environments at once (i.e., the attack surface) run contrary to a basic zero trust principle: That the attack surface cannot be perfectly defended. Further, failing to direct agencies to identify protect surfaces robs them of the ability to make and measure progress incrementally, while increasing costs and risks and decreasing the effectiveness of zero trust architectures.

Parallel Lines of Directive Authority over Federal IT

The *Federal Information Security Modernization Act*, 44 U.S.C. § 3552-3553, establishes two parallel lines of directive authority over Federal IT systems and networks.

- Directive authority over most Federal department and agency information systems rests with The Director of the Cybersecurity and Infrastructure Security Agency (CISA); the departments and agencies governed thusly are commonly called Federal Civilian Executive Branch (FCEB) agencies. These systems are also subject to general IT-related policies, instructions, and guidelines issued by the Office of Management and Budget (OMB), typically transmitted in OMB memoranda.
- The National Manager for National Security Systems (NSS) holds directive authority over NSS, which the law defines (summarizing here) as those systems involved in defense, intelligence, and classified information processing, excluding systems used only for routine administrative functions. NSS policies and instructions are issued by the Committee on NSS (CNSS), the Department of Defense (DoD) (for defense systems), and the Director of National Intelligence (for intelligence systems).

The following table summarizes the principal directives regarding zero trust issued to both FCEB agencies and to the NSS community, focused specifically on defense IT systems.

Table 1: Summary of Primary Zero Trust Directives

Zero Trust Directive Type	**For FCEB Agency Systems**	**For Defense IT Systems**
Executive orders	EO 14028	National Security Memorandum (NSM) 8
Reference zero trust architectures	NIST SP 800-207	DoD Zero Trust Reference Architecture
Implementing directives	OMB M-22-09 & M-24-14	DoD Zero Trust Strategy
Directed measures of zero trust	CISA Zero Trust Maturity Model	DoD Zero Trust Implementation Roadmap

FCEB Agency Zero Trust References and Directives

As already mentioned, EO 14028, which was issued in May 2021, provided the proverbial starting gun for Federal zero trust architecture implementation efforts. EO 14028 was issued in the aftermath of several high-profile cybersecurity incidents, and as such it touches on several topics relevant at the time (other notable topics include cloud security, software supply chain security, vulnerability disclosure, and cybersecurity incident reporting and data sharing). Section 3 of the order provides the meat of directed actions regarding zero trust, and specifically directs FCEB agencies to develop and begin reporting to OMB on plans and budgets necessary to implement zero trust architecture within 60 days of issuance (Biden 2021, Sec 3(b)). Shortly after directing zero trust architecture actions, the order also directs the same agencies to adopt multi-factor authentication (MFA) and encryption of data at rest and in transit, although this is in a different section of the order than zero trust actions (Biden 2021, Sec 3(d)).

NIST SP 800-207, *Zero Trust Architecture*, already existed when EO 14028 was issued, having been released almost a year prior in August 2020 (Rose, et al. 2020). As previously noted, this document recommends that organizations already architected for boundary defense and seeking to move to a zero trust architecture follow an approach that is substantively identical to the one suggested by Cloud Security Alliance that is otherwise considered commercial best practice—beginning with an inventory of protect surfaces, implementing them incrementally, although this document doesn't use that term (Rose, et al. 2020, Sec 7.3, pp 37-41). Technically, at this point, FCEB agencies might have and arguably should have done exactly that—although in their defense, this guidance appears in the last five pages of the document and is easily missed.

OMB would have received initial plans and budgets from FCEB agencies in July or August of 2021. One month later, in early September 2021, CISA released the first iteration of its

Zero Trust Maturity Model (ZTMM), which established a five-pillar reference model for zero trust activity, and set initial measures of maturity for agencies to follow in each of the five pillars. The zero trust maturity model was later revised in response to public comment to the current 2.0 version in April 2023 (CISA 2023).

Presumably in response to initial agency plans submitted six months prior, and aligning with the CISA ZTMM issued three months prior, in January 2022, OMB issued OMB Memorandum M-22-09, *Moving the U.S. Government Toward Zero Trust Cybersecurity Principles* (Young 2022). This document, for the first time, issued discrete and measurable technical tasks to FCEB agencies to implement zero trust, translating the guidelines in EO 14028 and in NIST SP 800-207, and the measurement framework established by the ZTMM, into specific technical actions for agencies to take—with many of the deliverables due by the end of Fiscal Year (FY) 2024 (Young 2022, 4).

Analysis of FCEB Agency Directives

The high-level direction given in EO 14028 is in no way contrary to the zero trust principles outlined in the previous section, and the reference architecture in NIST SP 800-207 suggests a course of action that, as already mentioned, aligns well with accepted best practices. Unfortunately, the measurement framework established by the CISA ZTMM and the specific tasking directed by OMB M-22-09 do not implement that course of action. Instead, they explicitly direct agencies to implement zero trust architecture across the entire agency at once and fail to direct identification or use of protect surfaces. Specifically, there are three problems.

First: All measures of maturity in the CISA ZTMM are explicitly scoped to the entire agency. Consider, for example, the Initial measure of maturity for the Authentication function,

the first function in the first pillar (Identity), which reads: “Agency authenticates identity using MFA, which may include passwords as one factor and requires validation of multiple entity attributes (e.g., locale or activity)” (CISA 2023, 13). On face value, this measurement applies to the entire agency—not a specific system, or network, or ‘ZTA candidate’, but everything at once. To achieve Initial maturity for Authentication, an agency must be doing what is described here, everywhere. The same pattern holds true for every measurement in the document.

Second: The tasks given to FCEB agencies in OMB M-22-09 are similarly agency-wide in scope, aligned with the CISA ZTMM. For example, M-22-09 directs that “[a]gencies must remove password policies that require special characters and regular password rotation from all systems” within one year of issuance (Young 2022, 28). This explicitly applies to all systems within the scope of a given agency. This pattern also applies to every task in the document.

Finally: Nowhere in the CISA ZTMM or in OMB M-22-09 are FCEB agencies directed to identify protect surfaces, or anything that could reasonably be interpreted as protect surfaces. Some might argue that there are explicit tasks in both documents requiring inventory of things like assets and data, and that this is the same as directing the identification of protect surfaces. However, such an argument would be flawed, given that these tasks are also agency-wide in scope, and at no point are agencies directed to use the results of those tasks to scope or focus other work described and mandated in both documents.

It should be noted that later OMB guidelines did begin to prioritize activities directed in M-22-09. Specifically, OMB Memorandum M-24-14 directed FCEB agencies to prioritize “systems that cannot deploy modern cybersecurity controls” and directed that, in addition to reporting “implementation of zero trust on all information systems,” reports should also describe

the same things for "all high value assets and high impact systems" (Young and Coker, Jr. 2024, 2). However, the scope of directed actions has not narrowed; they still apply to all systems.

FCEB Agency Zero Trust Results

CISA issued a report to Congress on Zero Trust Architecture Implementation in January 2025 (CISA 2025). As of this writing, this report contains the most recent publicly available data. The report lists general progress, identifying high level results as well as common challenges, without identifying any specific agency. Essentially, this report provides trends and analysis.

Broadly, agencies have made progress against the goals established in OMB M-22-09, but none of these goals is reported to be achieved. For example, the following graphic describes implementation of phishing-resistant MFA (the example agency-wide requirement listed earlier):

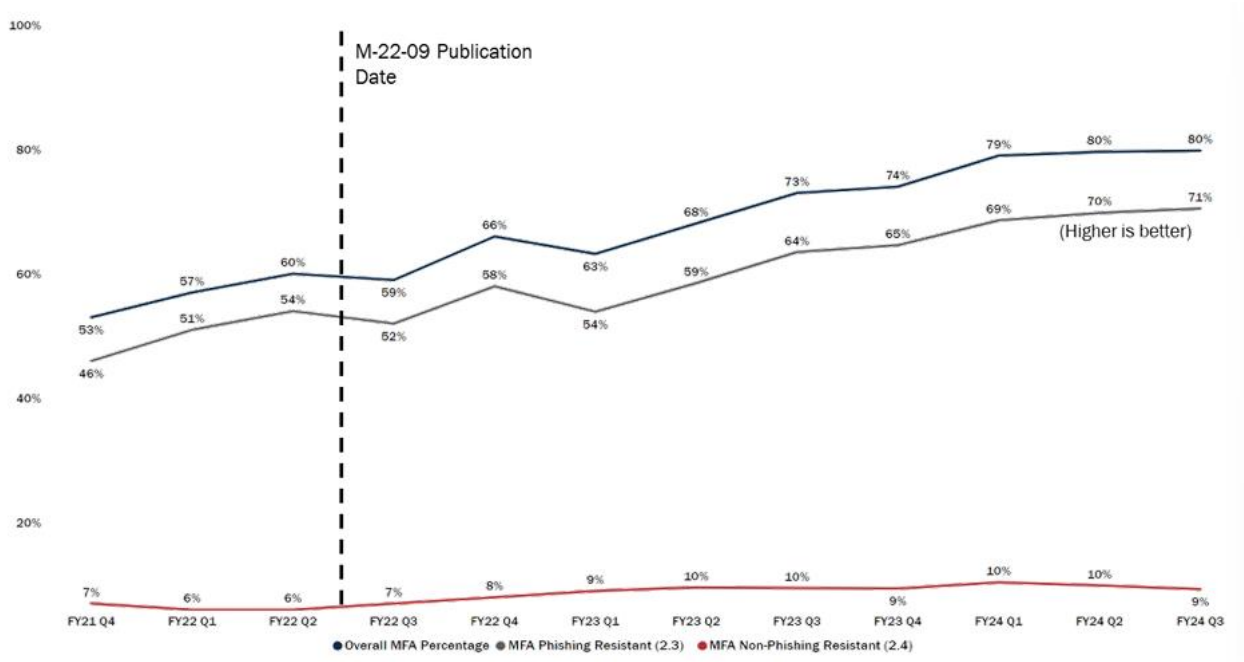

Figure 3: MFA and Phishing-Resistant MFA Implementation (CISA 2025, 7)

This result is typical of what CISA reports overall: Broad-based progress has been made, but none of the directives given in OMB M-22-09 yet have been achieved. Agencies report a

wide variety of challenges, including many one-off scenarios that prevent universal compliance, such as vendors that cannot add necessary software features or the unavailability of products to satisfy requirements under all conditions. Every pillar reports some version of this result.

These results only underscore the condition that gives rise to zero trust in the first place: That it is not possible to perfectly protect the attack surface, and therefore a different approach, zero trust, is required. Zero trust assumes that there will be vulnerabilities—gaps in coverage, cracks in the organization's outer shell—and that attackers will breach organizational perimeters. These results bear that out. FCEB agencies have been directed to close all avenues of attack, albeit only for a few specific kinds of attacks—and they cannot comply. No amount of time, resources, or pressure will change that outcome. A different approach is required.

Defense IT Zero Trust References and Directives

Defense-related IT systems and networks, as outlined previously, fall under a different line of authority, and accordingly, they respond to a different set of directives. These directives contain different, but broadly similar, inherent flaws, and based on the limited information publicly available, defense IT systems have met with similar results to the FCEB agencies.

While parts of EO 14028 apply to the NSS community (including defense IT systems), Section 3, which directed zero trust implementation actions for FCEB agencies, did not immediately apply to DoD or the IC. Instead, Section 9 directed DoD and the IC to "adopt National Security Systems requirements that are equivalent to or exceed the cybersecurity requirements set forth in this order that are otherwise not applicable to National Security Systems" (Biden 2021, Sec. 9(a)). The order that met this requirement is National Security Memorandum (NSM)-8, *Memorandum on Improving the Cybersecurity of National Security,*

Department of Defense, and Intelligence Community Systems (Biden 2022). In substance, the part of NSM-8 that implements zero trust architecture requirements from the original executive order directs the head of each department or agency owning or operating NSS to develop plans to implement zero trust architecture, incorporating NIST SP 800-207 guidance (as appropriate), as well as any relevant CNSS instructions, directives, and policies, and to report the results to the National Manager and to CNSS.

From here, DoD and IC directives diverge somewhat, and in the interest of efficiency, and considering the public availability of the relevant materials, this paper will focus on those directives relevant to defense IT systems.

The *DoD Zero Trust Reference Architecture*, authored by the Defense Information Systems Agency (DISA) and the National Security Agency (NSA), builds substantially on the technical framework in NIST SP 800-207. It establishes a seven-pillar model for zero trust, consisting of the five pillars established by CISA for the FCEB agencies, and adding two of CISA's 'cross-cutting functions' – Automation and Orchestration, along with Visibility and Analytics – as additional pillars (DISA and NSA 2022, 21-23).

The *DoD Zero Trust Strategy* serves as the initial tasking document for DoD zero trust implementation. The strategy establishes the timeline for DoD elements to implement zero trust, with a target completion by end of FY 2027, and establishes three broad Courses of Action (COAs) that DoD elements might take to implement zero trust (DoD 2022, 4-5). Released at the same time, and updated in January 2025, the *DoD Zero Trust Execution Roadmap* outlines the specific activities required to be undertaken by DoD elements, most of which will implement COA 1, which the strategy describes as a brown field approach (i.e., adapting from an existing technology baseline) (DoD 2025). Taken together, these documents provide the measurable tasks

that define zero trust implementation for DoD entities. DoD tasking identifies two specific zero trust implementation targets: A "Target Level" for zero trust compliance is achieved when an organization completes 91 specified activities (out of a total of 152 defined activities); all other activities are defined as "Advanced Level" (DoD 2025).

Analysis of Defense IT Directives

Within the DoD Zero Trust Reference Architecture, the term 'protect surface' is used exactly once across 104 pages of dense technical material (DISA and NSA 2022, 82). Unfortunately, while this single use is contextually consistent with the accepted meaning of that term, the remainder of the document neither defines it, nor employs the concept of the protect surface by any other name. Perhaps ironically, the next two pages of the document present a view of transition from the as-is architecture to the target state—this is exactly where a discussion of identifying protect surfaces and implementing them incrementally should appear, but instead, the architecture describes what must change, without describing any pacing or scoping mechanism (DISA and NSA 2022, 83-84). The document deeply describes zero trust technical concepts, and it does describe the need for incremental technology changes while adopting zero trust—but it never describes what might define an increment.

Unfortunately, an exhaustive review of the strategy and implementation roadmap finds that they have the same basic errors of omission as the DoD Zero Trust Reference Architecture. The term 'protect surface' does not occur, and more importantly, the concept is absent. DoD elements are directed to exhaustively identify all data, assets, applications, and service elements, and to control access to resources at a granular level based on context—but there is no direction in either document regarding scoping, prioritization, or sequencing of these activities, and at

times these activities are explicitly applied to all instances of a specific type of resource. The intention of data tagging activities, for example, is to explicitly tag and granularly control access to **all data**, leveraging increasing degrees of automation over time (DoD 2025, ID# 4.3.1 - 4.3.5). While a few tasks suggest a phased approach, there is no obvious answer to how the activity might be phased. All requirements apply equally to everything, and though activities are logically sequenced (e.g., predecessor and dependent activities are identified), within any given activity, there is no apparent way to group or sequence the objects to which the activity applies.

To consider a counterpoint, some might argue that this is intentional on the part of the DoD and is done because the framework of prescribed capabilities and activities can be applied to a single system just as readily as it can be applied to an entire organization, and to a new environment or an existing one. Such an argument would be flawed, though, given that the framework is most commonly being applied to enterprise-level IT environments—and, within enterprise environments, DoD would need some mechanism to prioritize among multiple objects to which each prescribed action may be applied, because without such a mechanism, all actions apply to everything at once. Thus, the lack of a protect surface or any similar framing mechanism in DoD's strategy and execution roadmap hinders execution at all but the very smallest levels at which this framework might be applied (such as a single system).

Defense IT Reported Results

While there is no official report that publicly accounts for DoD's progress toward its zero trust architecture goals, the limited information published by DoD outlets is not encouraging. The most recent such report was published in an article on the Army.mil website in February 2025, and quotes Col Gary Kipe, then chief of staff of DoD's Zero Trust Portfolio Management

Office (PMO), speaking at the Zero Trust Summit in Washington on February 19th of that year, as stating that: "[According to] the latest analytical review of the implementation plans that we have received back from across the enterprise, we're doing well, but we're not anywhere close to being done." The article goes on to report that according to the PMO, "current data shows that across all 58 components, 14% of target level Zero Trust activities have been completed across DoD" (Lopez 2025). After 50% of the planned schedule, DoD had achieved 14% of the goal.

These results underscore a need for zero trust architecture implementation targets to exist at some level other than the entire enterprise. Considering that DoD is the single largest entity in the Federal government by most measures, it is especially important to identify what, other than everything, must be protected specifically—and especially challenging to accomplish any, let alone all, of the outcomes that DoD has defined for itself. Much like the FCEB agencies, defense IT systems appear to be making progress, but the basic principle remains: That the attack surface is vast, unknowable, and inherently indefensible, and therefore a new approach is needed.

Addressing the Dichotomy in Federal Zero Trust Directives

We have shown that the goals in zero trust directives across the Federal government, for FCEB agencies and for defense IT systems, are broadly unachievable because they are scoped to all things at once, targeting the entire attack surface, whether by default or explicitly. Further, we have shown that no investment of time, money, or other resources will solve this basic problem. Not only is this evident through logical analysis of the tasks, but it is borne out by the results. This dynamic cannot change without resolving the inherent dichotomy in the tasks themselves.

At the same time, the work done over the last several years has produced results. Considering Figure 3, FCEB agencies have made progress; even if 100% is not achievable, fewer systems are using weaker authentication protocols, and more systems are using multifactor

authentication, including phishing-resistant mechanisms. Similarly, 14% Target Level zero trust maturity among DoD components is laudable progress, even if 100% is also unachievable here.

With these things in mind, an actionable strategy to accelerate zero trust architecture implementation across the Federal government should take maximum advantage of the progress and the institutional investments made to date—to include preserving as much of the existing tasking frameworks as possible for both FCEB agencies and defense IT systems.

A Strategic Adjustment

Both the FCEB agencies and defense IT systems have zero trust implementation strategies. However, both strategies suffer from the same fundamental flaw: The actions they direct are either not scoped specifically, or they are scoped far too broadly. At the same time, apart from a section at the end of NIST SP 800-207, which never uses this term, the principal references upon which the Federal government relies for zero trust implementation do not define the term or concept of the protect surface. Apart from these problems of scope and omission, though, the actions directed on all sides are rational, logical, and sound; if the same actions were taken incrementally, one protect surface at a time, the associated goals would become achievable and measurable on the same basis: One protect surface at a time.

The fact that these are common problems means that there is an opportunity to solve them in a common way. If the Federal government developed a shared, government-wide definition of the zero trust protect surface, based on accepted industry-developed best practices, that definition could supplement existing strategies, plans, and task frameworks, closing the same logical gap in both FCEB and defense IT directives without replacing, and perhaps without significantly altering, anything else.

Given a shared Federal definition of the protect surface, the task would then fall to implementing departments and agencies to adjust their existing plans, budgets, and reports on zero trust implementation to be built and measured against this new definition. For example, instead of planning and budgeting to implement every task defined in OMB M-22-09 and OMB M-24-14 across entire agencies, FCEB agencies would identify protect surfaces, sequence them, and then report the same activities, based on the same references, with the target being to complete all of those activities for some defined set of protect surfaces within a given year. Likewise, DoD components could shift from budgeting and reporting progress toward achieving Target Level zero trust for entire enterprise environments, and instead define protect surfaces, sequence them, and then budget, plan, and report progress toward achieving Target Level zero trust activities for a defined set of protect surfaces each year.

Making this adjustment means accepting that not everything will be protected at the same time and to the same extent. Some may point this out and argue that the proposed approach is unacceptable because it means abandoning the idea of implementing zero trust architecture in support of at least some systems and environments. However, such arguments would be flawed. It is true that segmenting zero trust architecture implementation efforts by protect surface would inherently mean that some protect surfaces are prioritized over others, leaving some things more poorly protected than others. However, the status quo alternative is far worse—not as compared with the original plan (to protect everything equally), but as compared with actual results (which show that Federal agencies are broadly failing to protect anything adequately, *because they are trying to protect everything equally*). Choosing to define protect surfaces and using them to sequence and prioritize zero trust architecture implementation offers a much better chance of defending some of the defined protect surfaces adequately through zero trust implementation.

Anticipated Implementation Difficulties

Two aspects of the proposed strategic adjustment are likely to be difficult to implement. This section anticipates these particularly hard problems and offers advice to implementers of the proposed strategic adjustment to help mitigate and resolve them.

First: Defining the protect surface in a way that is both broadly accepted as a definition and useful in the manner intended is likely to become a significant challenge. This is not because the definition itself is inherently complex, but because the resulting list of items will become a means of controlling budget, and as such, defining protect surfaces will become a proxy means of capturing budget for systems, networks, or IT services designated as protect surfaces. To mitigate this challenge, carefully and consistently remind people that there is no logical limitation to the number of protect surfaces that may exist in an organization or environment, and that ultimately, every protect surface will be implemented—this just won't happen all at the same time. It may also be helpful to point out that implementing zero trust architecture against a protect surface creates schedule pressure and a reporting requirement while implementation proceeds—and those system, network, or IT service owners whose programs aren't designated as protect surfaces in process of implementation (yet) do not have those requirements. There will still be winners and losers, but that is unavoidable, given the shift in paradigm.

Second: Making the proposed strategic adjustment will necessitate some level of cultural adjustment. Zero trust implementation is not new; it has been happening now for years, and an entire culture exists across Federal departments and agencies that is steeped in the language and ideas that define zero trust implementation at each department and agency. To mitigate this challenge, it is important to change as little as possible of the language and ideas surrounding zero trust implementation uniquely at each department or agency. For example, within DoD, it

may be necessary to make small adjustments to language when applying existing capabilities and activities in the DoD Zero Trust Implementation Roadmap. If the same task identifiers can be used to refer to a protect surface as were used to refer to a whole enterprise or agency environment—and this should be possible—not only does this preserve existing knowledge, but it also reduces the cultural friction that would otherwise result from changing a plan that has been in process now for years in a fundamental way. By preserving as much as possible of the language and culture that is now in place, departments and agencies can minimize friction.

Conclusion

This paper has introduced an industry standard definition of the protect surface concept. Using this definition, this paper has critically examined foundational Federal zero trust directives, standards, and guidelines to show how these documents mostly fail to define the protect surface concept and completely fail to provide adequate direction to government agencies on the definition and use of protect surfaces to prioritize, sequence, and incrementally implement zero trust architecture. Further, by examining reported results, this paper has conclusively shown that the status quo is not achieving stated goals and argued that the status quo is unlikely to produce the desired results, regardless of any further investments of time or additional resources. Finally, this paper proposes solutions in the form of a targeted, strategic adjustment to existing Federal zero trust architecture directives and guidance, described how to make this adjustment, and suggested ways of mitigating hard problems likely to surface because of the adjustment.

This analysis has demonstrated that Federal zero trust implementation references mostly omit not only the term, but far more importantly the essential concept of the protect surface. This omission renders Federal zero trust strategy unachievable by removing a clear, viable definition

by which organizations can measure zero trust outcomes, resulting in something that is not strategy at all, but a series of tactics that never add up to results.

As famously observed in The Art of War: “Strategy without tactics is the slowest route to victory. Tactics without strategy is the noise before defeat” (Tzu 1971). The Federal government now has two logical choices: Either continue to make lots of noise or start making real progress.

Bibliography

Beddell, Charlie. 2024. *John Kindervag's 3 Zero Trust Truths for Government Agencies.* https://www.illumio.com/blog/john-kindervag-zero-trust-government-agencies.

Biden, Joseph R. 2021. "Executive Order 14028: Improving the Nation's Cybersecurity." *Federal Register.* May 12. https://www.federalregister.gov/d/2021-10460.

—. 2022. "Memorandum on Improving the Cybersecurity of National Security, Department of Defense, and Intelligence Community Systems, National Security Memorandum/NSM-8." *govinfo.gov.* January 19. https://www.govinfo.gov/content/pkg/DCPD-202200025/pdf/DCPD-202200025.pdf.

CISA. 2025. "Zero Trust Architecture Implementation Fiscal Year 2024 Report to Congress." *CISA.gov.* January 29. https://www.dhs.gov/sites/default/files/2025-04/2025_0129_cisa_zero_trust_architecture_implementation.pdf.

—. 2023. "Zero Trust Maturity Model version 2.0." *CISA.gov.* April. https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf.

Cloud Security Alliance. 2024. "Defining the Zero Trust Protect Surface." *Zero Trust Research Working Group.* March 5. https://cloudsecurityalliance.org/artifacts/defining-the-zero-trust-protect-surface.

Cuffari, Joseph V. Ph.D. 2025. "Cybersecurity System Review of a Selected High Value Asset at CISA." *DHS Office of the Inspector General (oig.dhs.gov).* January 15. https://www.oig.dhs.gov/sites/default/files/assets/2025-01/OIG-25-08-Jan25.pdf.

Cybersecurity and Infrastructure Security Agency. 2018. "Binding Operational Directive 18-02: Securing High Value Assets." *CISA.gov.* May 7. https://www.cisa.gov/news-events/directives/bod-18-02-securing-high-value-assets.

Department of Homeland Security. 2025. "DHS Zero Trust Capability Framework." June.

DISA and NSA. 2022. "Department of Defense (DoD) Zero Trust Reference Architecture Version 2.0." *Department of Defense CIO (dodcio.defense.gov).* July. https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v2.0(U)_Sep22.pdf.

DoD. 2025. "DoD Zero Trust Execution Roadmap (COAs 1-3), v1.1." *DoD CIO (dodcio.defense.gov).* January 22. https://dodcio.defense.gov/Portals/0/Documents/Library/ZT-CapabilitiesActivities.pdf.

—. 2022. "DoD Zero Trust Strategy." *DoD CIO (dodcio.defense.gov).* October 21. https://dodcio.defense.gov/Portals/0/Documents/Library/DoD-ZTStrategy.pdf.

GAO. 2018. *GAO-19-128: Weapons System Cybersecurity: DoD Just Beginning to Grapple with Scale of Vulnerabilities.* October. https://www.gao.gov/assets/gao-19-128.pdf.

Kindervag, John. 2010. "No More Chewy Centers: Introducing The Zero Trust Model Of Information Security." Forrester Research, Inc. September 14 (updated September 17). https://media.paloaltonetworks.com/documents/Forrester-No-More-Chewy-Centers.pdf.

Lopez, C. Todd. 2025. "'Zero Trust' architecture could prevent adversary data theft, protect warfighters." *DoD News (army.mil).* February 27. https://www.army.mil/article/283397/zero_trust_architecture_could_prevent_adversary_data_theft_protect_warfighters.

Mulvaney, Mick. 2018. "OMB Memorandum M-19-03: Strengthening the Cybersecurity of Federal Agencies by enhancing the High Value Asset Program." *whitehouse.gov.* December 10. https://www.whitehouse.gov/wp-content/uploads/2018/12/M-19-03.pdf.

NIST. n.d. *Computer Security Resource Center (CSRC) Glossary, attack surface.* Accessed February 27, 2026. https://csrc.nist.gov/glossary/term/attack_surface.

NIST NCCoE. 2025. "NIST SP 1800-35: Implementing a Zero Trust Architecture: High Level Document." June. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1800-35.pdf.

NSTAC. 2022. "NSTAC Report to the President: Zero Trust and Trusted Identity Management." *CISA.gov.* February 23. https://www.cisa.gov/sites/default/files/publications/NSTAC%20Report%20to%20the%20President%20on%20Zero%20Trust%20and%20Trusted%20Identity%20Management.pdf.

Rose, Scott, Oliver Borchert, Stu Mitchell, and Sean Connelly. 2020. "NIST SP 800-207: Zero Trust Architecture." August. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf.

Thiele, Vincent. 2021. "What is the attack surface of an organization?" *Darktrace.com.* June 2. https://www.darktrace.com/blog/understanding-your-organizations-attack-surface-and-why-it-poses-a-risk.

Trump, Donald J. 2026. "President Trump's Cyber Strategy for America." *whitehouse.gov.* March. Accessed March 22, 2026. https://www.whitehouse.gov/wp-content/uploads/2026/03/president-trumps-cyber-strategy-for-america.pdf.

Tzu, Sun. 1971. *The Art of War.* Translated by Samuel B. Griffith. Oxford: Oxford University Press.

Young, Shalanda D. 2022. "OMB Memorandum M-22-09: Moving the U.S. Government Toward Zero Trust Cybersecurity Principles." January 26. https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf.

Young, Shalanda D., and Harry Coker, Jr. 2024. "OMB Memorandum M-24-14: Administration Cybersecurity Priorities for the FY 2026 Budget." July 10. https://www.whitehouse.gov/wp-content/uploads/2024/07/FY26-Cybersecurity-Priorities-Memo_Signed.pdf.