All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

Don’t trust Zero Trust – The Whole of Government Approach

It looks like there is a coordinated effort in the US government to promote so-called “zero trust architecture”. I wrote my last article on the NIST ridiculousness in this arena, and since posting it, I have found that there are several other Federal attempts to promote the so-called architecture both within and outside of government. The NSA attempt was ridiculous on its face. 7 Pages, one (20% full) with references, a header page, and a few bullet lists that aren’t really worth bothering with. But my good friends on the Internet identified some others to me,. And one in particular is worth writing about. It’s the DISA report (prepared jointly with the NSA) titled “*Department of Defense (DOD) Zero Trust Reference Architecture - Version 1.0 - February 2021 - Prepared by the Joint Defense Information Systems - Agency (DISA) and National Security Agency (NSA) Zero Trust Engineering Team*”

Version 1.0 – what do you expect?

Consider that most implementations don’t get usable till about version 2.something. So don’t be disappointing by this early attempt. As a government document, the actual content doesn’t start at all til page 7 (document ‘page 1’) of the 170 pages, and the “Vocabulary” section starts on page 98. It’s good that they define their terms and this is well expected from a government document (unclassified version).

Most of the document is a set of tables, so naturally there is a table of tables… but hey, I love this kind of stuff. It is clear that they have and apply a document standard and that is a good thing. I has a revision history, approval details, and a list of figures. It references “NIST SP 800-207 Zero Trust Architecture, August 2020”, which I covered in some depth in my previous article, quoting “Zero Trust assumes there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the Internet) or based on asset ownership (enterprise or personally owned).”

Purpose of ZT

The identified purpose is:

- “*Zero Trust (ZT) is a cybersecurity strategy and framework that embeds security throughout the architecture to prevent malicious personas from accessing our most critical assets. It provides zones for visibility and information technology (IT) mechanisms positioned throughout the architecture to secure, manage and monitor every device, user, application, and network transaction occurring at the perimeter and/or within a network enclave. Zero Trust is an enterprise consideration and is written from the perspective of cybersecurity. The foundational tenet of the Zero Trust Model is that no actor, system, network, or service operating outside or within the security perimeter is trusted. Instead, we must verify anything and everything attempting to establish access. It is a dramatic paradigm shift in philosophy of how we secure our infrastructure, networks, and data, from verify once at the perimeter to continual verification of each user, device, application, and transaction.*”

And here the problems begin. The purpose is to "*prevent malicious personas from accessing our most critical assets*". The way this is attempted is "*embeds security throughout the architecture*". By conflating the actual purpose with the means of achievement they confound themselves by limiting their perspective. Suppose there are other ways or that embedding security throughout an architecture does not get the job done? This initial assumption is a problem we will see repeated throughout.

The next problem is that "zero trust" is identified as:

- "*Zero Trust is an enterprise consideration …*"
  - So it's really not an architecture at all….
- "*The foundational tenet of the Zero Trust Model is that no actor, system, network, or service operating outside or within the security perimeter is trusted.*"
  - It's a "model", not an "architecture".
  - "Actor, System, Network, Service" are the things that are not "trusted". So here we go to the Glossary of Terms (P96) and find the definitions for these terms.. right? Nope. Actor is not there, System is not there, Network is not there, and Service is not there. So the basic things we are controlling are not defined. Bad start…
    - Note that thee is a section called "Services Definitions" (Table 7) but it does not define what is a Service", it does list a set of service names and their definitions. So if this is what they meant by services, they have both defined and (perhaps over-) constrained that particular thing they are not trusting.
    - I should point out that one of the Services they do not trust is apparently "S1.2.2 Identification and Authentication Services", which includes Identity, Attribute, Credential, Authentication, and Directory Management. So they do not trust the mechanisms that they use to make the decisions about the access they are prevented from maliciously accessing… until they trust them? We will see how they propose to do this… or not.
  - There is "*the security perimeter*" which is both ridiculous and has not been the case in DoD systems I am aware of since the 1970s at least. There are in fact multiple perimeters, a layered approach, and other things, and this has certainly been true since the creation of the early DoD networks (DISN is the earliest one I recall at the moment).
- "*Instead, we must verify anything and everything attempting to establish access.*" But the architecture only verifies "Actor, System, Network, and Services", and since they are not defined, and certainly don't constitute everything and anything, they have doomed themselves to failure by starting on the ridiculous and false premise that they will then deep-end (and depend) on throughout.
- "*It is a dramatic paradigm shift in philosophy*" So it's a philosophical thing… The thing about philosophy, is it depends generally on faith without provable basis or logical analysis. I believe they may be doing the former and are clearly not doing the latter.

From the get-go, they are trying to do something they have not apparently really thought through. Or perhaps they have and this is merely the public face of what really underlies it.

Scope

- "*The scope of the DOD Zero Trust Architecture (ZTA) effort is specifically to determine capabilities and integrations that can be used to successfully advance the Department of Defense Information Network (DODIN) into an interoperable Zero Trust end state.*"

So really, the scope of the effort is not to define ZT, but rather to figure out how to do it, which is to say, they don't yet know how to do it, and there is the (very real) possibility that it in fact cannot be done.

The scope of this document (mine) is then to save hundreds of billions of dollars and hundreds of thousands of person years of effort by identifying the limitations that they face now and helping them redirect their efforts to something worthwhile.

- To be clear, I think that excessive and poorly defined trust models are in use today, and that improving them would be a great step forward for cybersecurity.
- **The first model to abandon is the notion that you cannot trust anything**, because **you have to trust some things to some extent for some purposes**. It is the nature of bootstrapping trust that it starts with an assumption and tests the assumption along with the derivations from it over time.

Here's what's missing

> Trust is widely understood to mean the extent to which you allow that you can be harmed from another.

The concept of Zero Trust is fundamentally and fatally flawed, because without trust, we can go nowhere. The question is not whether we trust, but rather:

> "What and who should we trust for what purposes and to what extents, and how does this change over time?"

More simply, what are we willing to lose by each thing we attempt?

This is a fundamental of risk management that most in cybersecurity fail to understand.

> **<u>We take risks to gain rewards.</u>**
>
> We trust to gain the benefits of trust.
>
> No risks, no rewards! → No trust, no benefits!

But back to the DoD

After some various language with no apparent real implication, we see (P4) (my interpretation indented):

- "*Vulnerabilities exposed by data breaches inside and outside DOD demonstrate the need for a new and more robust cybersecurity model that facilitates well informed risk-based decisions."*
  - In other words, we are scared
- *"Zero Trust is a cybersecurity strategy and framework that embeds security principles throughout the Information Enterprise (IE) to prevent, detect, respond, and recover from malicious cyber activities."*

    - We used a list of things that are important (we think) that we think cover the space (you forgot deter, interdict, and adapt by the way).
- *“This security model eliminates the idea of trusted or untrusted networks, devices, personas, or processes, and shifts to multi-attribute-based confidence levels that enable authentication and authorization policies based on the concept of least privileged access. “*
    - A breakthrough in marketing technology! We go from trusting things to a complex mathematical computation based on a set of poorly defined “attributes” an undefined set of no calculation provided “confidence levels”, to “enable” (cool – we are now all able) the concept (not realization) of “least privilege” based on (technical security) policies (not actual policies in the sense that executives understand them) regarding what you have to prove to be able to do what.
        - In other words, BS “Attributes” BS “least privilege” BS “policies” BS.
            - Yes – Attribute Based Access Control (ABAC) is an interesting approach and useful in many circumstances.
            - Yes – less privilege is a good idea (we cannot do least in any case).
            - Yes – Policies should be in place.
- *“Implementing Zero Trust requires designing a simpler and more efficient architecture without impeding operations to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services viewed as compromised.*”
    - However, in order to do this “simpler more efficient” thing we need lots of complex and inefficient and uncertain things that we don’t yet have (as we will see later).
    - And by the way, it only applies to “*systems and services viewed as compromised*”
        - But the ZTA is based on the assumption that everything is compromised!
            - Lions and Tigers and Circularity – oh my!
- “*Zero Trust focuses on protecting critical data and resources, not just the traditional network or perimeter security. “*
    - OK – Step 1 – identify what’s “critical”
        - Critical “data” – not defined in the document. Best do that...
        - Critical “resources” – not defined in the document. Best do that...
    - Ignore the “tradition” of “just” “network or perimeter security”.
        - OK – For clarity, tradition is something passed from generation to generation over time, and since you are likely 18 before you work for DoD, and your parents were likely 18 give or take when you were born, 36 years ago, (1985) we didn’t have much of an Internet or the Web at all, and the term Computer Virus was only 2 year old. So somehow, I don’t think Tradition is really the word we are looking for here in any case.

- And oh yeah. Cybersecurity 36 years ago was NOT just about “network or perimeter” In fact, it was largely about things like **change control**, which is the thing we are really **missing today** and **the reason for** things like the **Solar Winds attack being so successful**.
- The problem is not that “traditional” approached failed.
  - It’s that the community failed to apply what their elders figured out.
    - It was egotism and not listening to what other folks told you.

- *“Zero Trust implements continuous multi-factor authentication, micro-segmentation, encryption, endpoint security, analytics, and robust auditing to DAAS seven pillars to deliver cyber resiliency. “*
  - Thing 1: ZT does none of these things today. It’s just false.
  - Thing 2: Continuous multi-factor authentication – extremely high overhead and no apparent actual increased protective value. Also, not what they actually are proposing doing as far as I can tell. Also note that technically speaking, digital systems are, by their nature, discontinuous. They don’t normally do much of anything in a continuous fashion except at the device physics level.
  - Thing 3: Auditing, robust or otherwise, doesn’t prevent anything, being retrospective, but it can indeed add to resiliency by allowing subsequent detection and adaptation. I think the term detection might be what they meant here...
- *As the Department evolves to become a more agile, more mobile, cloud-instantiated workforce, collaborating with multiple federal and non-governmental organizations (NGO) entities for a variety of missions, a hardened perimeter defense can no longer suffice as an effective means of enterprise security.”*
  - It never could and never did. Start with the Great Wall of China, which never did it either. But the purpose of perimeter defense is not and never way to stop everything. It brings an economy of scale and reduction in complexity to the defender while increasing attacker workload and narrowing paths of attack.
  - So to be clear, the DoD has decided to stop doing what they never should have done in the first place, and never actually did… Good job declaring victory over your own stupidity. You have now replaced it with new stupidity.
- *“In a world of increasingly sophisticated threats, a Zero Trust framework reduces the attack surface, reduces risk, and ensures that if a device, network, or user/credential is compromised, the damage is quickly contained and remediated.*”
  - None of this has been shown to be true at all. It’s just wishful thinking. How do we know all this ZT stuff won’t in fact make us MORE susceptible to denial of services? And without services, how will we operate at all? Why do we think it reduces risk? What do you call risk anyway? Do you mean risk of leaking information? Is that the mission of the DoD? To keep everything secret? Or are boats intended to go out onto the water? And how does it ensure any such thing? It doesn’t. Not today, and it hasn’t been show to be able to do this for the future.
  - In other words, self aggrandizing BS.

- *“State-funded hackers are well trained, well-resourced, and persistent. The use of new tactics, techniques, and procedures combined with more invasive malware can enable motivated malicious personas to move with previously unseen speed and accuracy. Any new security capability must be resilient to evolving threats and effectively reduce threat vectors, internal and external.”*
  - We’re afraid! You should be. So do what we say and you don’t have to be afraid.
    - In fact, almost none of this is new in any way, and almost all of it has been true for at least 20 years, maybe more like 40 year.
    - And the “any new …” BS is … just BS. Get rid of “new” and it’s just as true.
  - Buy now! Before it’s too late (and you realize you are buying another shiny object that doesn’t actually save you)
- “*Zero Trust end-user capabilities improve visibility, control, and risk analysis of infrastructure, application and data usage. This provides a secure environment for mission execution.*”
  - BS – it’s simply not true.
    - ZT “end user capabilities”? What’s that? Not defined here…
    - How do ZT end user capabilities improve “risk analysis of infrastructure”?
    - How does this provide a “secure environment” for anything?
    - Doesn’t “mission execution” also require availability? ZTA reduces this...
- “*Enabling Zero Trust capabilities address the following high-level goals:*

  Modernize Information Enterprise to Address Gaps and Seams.

  Simplify Security Architecture.

  Produce Consistent Policy.

  Optimize Data Management Operations

  Provide Dynamic Credentialing and Authorization”
  - One out of 5 possibly true… This looks like a case of having a list of objectives somewhere and writing them down, then trying to declare that you met them as justification for your claims. For clarity:
    - ‘’’Gaps and Seems” is really a claim that DoD cyber-security is underfunded. But the problem is, the very agencies that are making the claim have enormous resources they don’t spend on this. They should redirect their own funds for better security, and I can help them do that… but they don’t listen.
    - “Simplify” is patently ridiculous. They offer a great deal of added complexity.
    - “Produce Consistent Policy”? Insane. They depend on policy from others, and the technical security policy they propose is far more complex than the previous ones, involving more things that can go wrong, and more people involved. Nothing in ZTA assures consistency or policy.

- “Optimize… operations”? This has nothing to do with ZTA. But indeed ZTA puts more requirements on those elements being more precisely defined to a very high level of granularity. Which also dramatically increases complexity, which is more likely to make operations less effective and efficient. But of course we won’t know until they do this for 10 years… always claiming they are getting closer (to the moon by each shovel of dirt you put on the mountain).
- “Provide dynamic credentialing and prioritization” – indeed. 1 out of 5

In summary, give us piles of cash and we will spend it. Trust us…

- But this is about Zero Trust!
    - **Don’t trust them!**
        - ***Make them verify these things BEFORE you allow them access (to $s).***

The structure of ZTA in DOD

The overarching structure of ZTA currently aimed toward is described in their Figure 1:

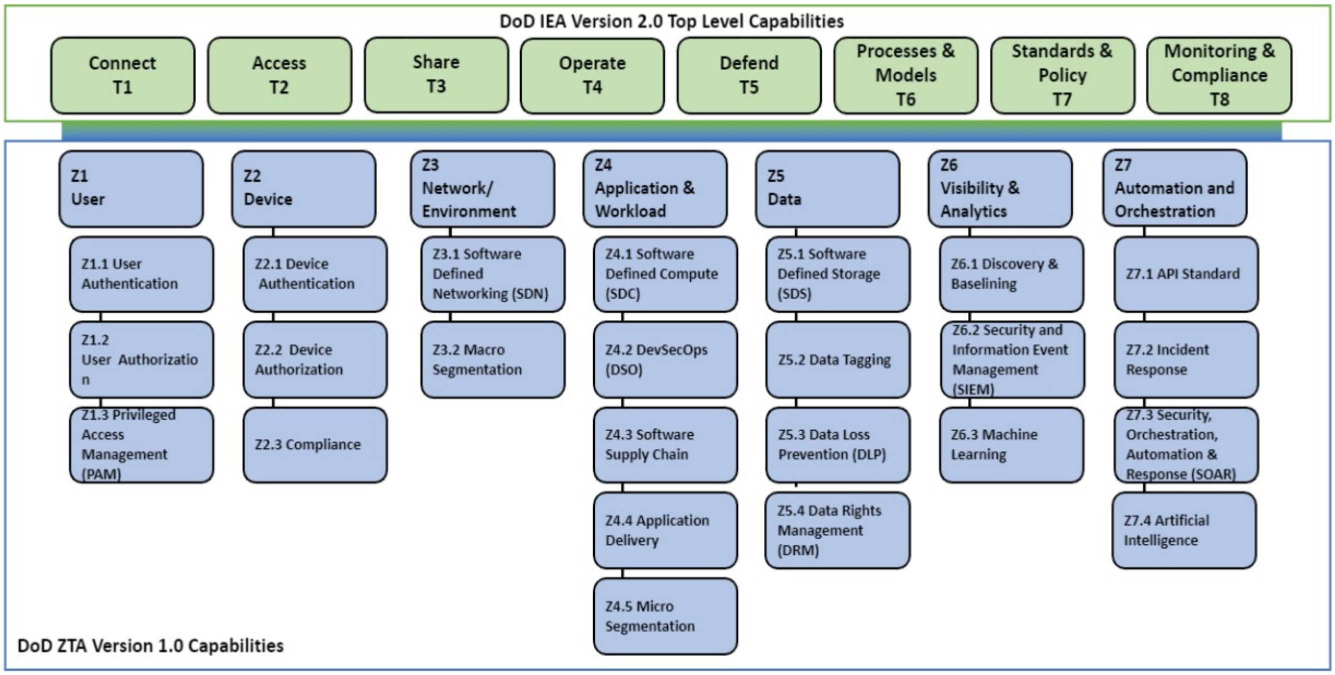
DoD IEA Version 2.0 Top Level Capabilities
Connect T1
Access T2
Share T3
Operate T4
Defend T5
Processes & Models T6
Standards & Policy T7
Monitoring & Compliance T8
Z1 User
Z1.1 User Authentication
Z1.2 User Authorizatio n
Z1.3 Privileged Access Management (PAM)
Z2 Device
Z2.1 Device Authentication
Z2.2 Device Authorization
Z2.3 Compliance
Z3 Network/ Environment
Z3.1 Software Defined Networking (SDN)
Z3.2 Macro Segmentation
Z4 Application & Workload
Z4.1 Software Defined Compute (SDC)
Z4.2 DevSecOps (DSO)
Z4.3 Software Supply Chain
Z4.4 Application Delivery
Z4.5 Micro Segmentation
Z5 Data
Z5.1 Software Defined Storage (SDS)
Z5.2 Data Tagging
Z5.3 Data Loss Prevention (DLP)
Z5.4 Data Rights Management (DRM)
Z6 Visibility & Analytics
Z6.1 Discovery & Baselining
Z6.2 Security and Information Event Management (SIEM)
Z6.3 Machine Learning
Z7 Automation and Orchestration
Z7.1 API Standard
Z7.2 Incident Response
Z7.3 Security, Orchestration, Automation & Response (SOAR)
Z7.4 Artificial Intelligence
DoD ZTA Version 1.0 Capabilities

Figure 1: Capabilities Taxonomy (CV-2)

“*The CV-2 capabilities are centered around the Zero Trust pillars which are User, Device, Network/Environment, Application & Workload, Data, Visibility & Analytics and Automation & Orchestration*”

Of course these are not the “Actor, System, Network, Service” described above, which is more than a bit confusing. User is not the same as Actor because Devices may also be Actors, but are also possibly Systems or parts of Systems, possibly performing Services. Applications may also presumably be Services, although as a Data Centric approach, Data would seem to be the thing we are controlling Access to...

"*The Zero Trust RA High Level Concept provides an operational view on how security measures would be implemented within the architecture.*

> As the Concept of Operations, would be is just right. The term Might may be more appropriate here, since this is not really finished today, but...

"Non-Person Entity (NPE) identity and user identity are tracked independently allowing for separate paths of validating confidence levels across enforcement points."

- It is unclear what these non-person entities are, of course, but I think it's reasonable to identify some of them as any thing – computer, system, etc. Any source of sequences of bits appearing at an interface. But here we start to see a problem. The sources of sequences of bits appearing at an interface are the things we seem to be addressing, but attribution of sequences of bits to sources is problematic, and in fact the underlying source of the problem this architecture is intended to address. Regardless of the method we use, other than physical (read physics) methods, there is no way to definitively establish the source of sequences of bits. So the question is, always has been, and likely will still be for a long time: "How definitive do we want/need to be?".
- The basic concept of using (or allowing for which is trivial) "separate paths of validating…" is an old and well tried and reasonable and prudent practice. But in the Internet era the whole point is economies of scale via shared infrastructure. If we need a second independent and different and never cross-connected network, that will be quite expensive and likely fail to remain separate regardless. Cryptographic separation is commonly used as the surrogate for actual separation, but this inherently limits surety, and the weight we are already putting on these cryptographic systems is far higher than what they can actually sustain.
- The idea that multiple enforcement points somehow get you out of the quandary is also problematic, in that this introduces more opportunities for failure, and more sets of rules that are either shared (leaving risk aggregation at the rule generation point) or independently derived, which tends to lead to inconsistencies in the rules, which is what the architecture is intended to avoid (according to the claims identified earlier).
- And finally we have the "confidence levels", a good idea, but no clue of how you gain these confidences (which by the way is another word for trust). So Zero Trust requires non-zero confidence. I have high confidence that confidence is closely linked to trust.

"Authentication and authorization activities will occur at numerous but focused points throughout the enterprise to include clients, proxies, applications and data. At each enforcement point, logs are sent to the SIEM and analytics are performed to develop a confidence level. Confidence levels of the device and user are independently developed and then aggregated where appropriate for policy enforcement. If the non-person entity or user has a confidence score above a measured threshold, then they are authorized to view the requested data. Data is protected along the way by Data Loss Prevention (DLP) which also feeds the SIEM to ensure the data is being used properly."

- Ouch! It was sounding good there for a while, but then they got to the part about the first "but". So it's not ubiquitous and continuous at all. It's at a substantial number (not yet defined and no formula to calculate how many) of focused points including (without limit?) "clients, proxies, applications, and data".

  - Just a minor point, **data** cannot authenticate or authorize anything. It's just data – sequences of bits. They are stored or transmitted or used, but of themselves are merely physical representations of conceptual decisions. They cannot authenticate or authorize. But data (sequences of bits) can be authenticated (sometimes) with other things (redundancy, consistency checks, physical location, source, etc.) as to some properties, and can be authorized for use, transmission or storage. But to do this for each sequence of bits (for example every byte) would require enormous overhead and of course all the overhead would itself be in these bits (in most cases) which themselves would also have to be authenticated and authorized, and "It's turtles all the way down."
  - Applications are things identified in the document, so I will leave that alone for now.
  - Proxies? Do they authenticate and authorize, or are they authorized and authenticated, or both? I think they meant the proxies to be mechanisms for authorization and authentication, but of course they themselves will also have to be authenticated and authorized… continuously… So how do we do this? Recursively of course. More Turtles.
  - And then we have clients, the still not clear what they are part. And which do they do? And how? And the Turtles? Unanswered questions...
- The log thing and all of that has the same problems, but I will move on for space.
- If the confidence score is above threshold… OK – so this goes back to the long-time approach of things like security levels. Clearances and classifications – the standard approach for along time. But in essence, this is a trust model, not a zero trust model. IT means we decide to trust things at some threshold for everything requiring that level. But this is not the multidimensional mechanism described in other ZT models where there are multiple elements of a 'tuple' that all have to be met. And of course if I get a high score in Geometry, it makes up for failing English, and I still graduate… right?
- Also note that this is all so they can "view" the"data". Nothing about use or use control, alteration or integrity, transparency requirements, accountability, custody… only "view".
- They there is "Data Loss Prevention" (DLP) which is about as good as detecting computer viruses with scanners, and the idea is that this is supposed to limit use! But use is not limited just by limiting access. Use control is a completely separate issue now conflated as well. And DLP as a solution for use control is nuts. Flow control possibly… I think they just had a DLP thing they wanted to include so it went here.

"The following bullets provide additional detail on the decision points, components, and capabilities that are depicted within the OV-1. The capabilities identified below are representative of an end-state Zero Trust implementation. Controlling access to resources based on the risk of the user and devices is the baseline requirement for Zero Trust and is possible without implementation of all identified capabilities."

- They then go and list a bunch of other technologies with possible uses. Essentially all of these are already existing technologies they already have and use that they shoe-horn in so they can keep all of that stuff and call it an architecture.

This document looks more and more like a justification for what they already do than a new approach to anything. Give us more money so we can keep screwing it up. In essence it is a long list of technologies that form the trust, but they all have to be trusted. So it's the opposite of Zero Trust – it's trust a whole lot of other things to play their part in a complex myriad of things which in combination are supposed to make thing simpler. I should add that essentially all of these things are themselves low surety mechanisms with large numbers of false positives and negatives. All of this means that the establishment of trust will almost always depend on a huge group process with some sort of threshold for acceptability. Which is to say, in essence, crowd sourcing trust. Sort of like how the rumor thing works in the Internet. The most popular rumor wins!

It's not all bad – it's just not ZT

I have set a limit of about 10 pages for my review, and in truth, there is so much that's a problem here, that I have spent more than a page per page just identifying the foolishness.

But that's not to say it's all bad. That's not the problem. The problem is it's:

NOT ZERO TRUST

Now I don't mean to put myself out as an authority on what is or is not zero trust. I think it is a ridiculous term and I am ridiculing it for that reason. But that's not the problem here. The problem here is that essentially every component violates the very principals the architecture is intended to create and enforce. It step by step says: "DO THIS – DON'T DO THAT" and then proceeds to "DO THAT" and "NOT DO THIS".

It just looks like a document written to appease a philosophy that's popular and forces that constrain what can be funded or done, while saying fund us to keep doing what we want to do. Which is to say, it's a document part of an influence operation. But whose influence operation?

- Is this the result of Russian or Chinese influence to work the US toward long-term bad security architecture?
- Is it just part of the normal approach that seems to work in the US today – lie to get what you want and hope they don't care (i.e. fake it till you make it)?
- Is it part of a vast conspiracy that the (name your group to dislike) is using to subvert all that's good and right?

There is an old saying. "Never attribute to malice what can be explained by stupidity."

I will leave it to you to decide.

Conclusions

"Don't trust the Zero Trust propaganda misnomer. It's hyperbole to cover a set of pre-existing methods in a so-far inconsistent, poorly thought, expensive, and time consuming approach."[1]

Hmmm.

Not sure I have anything to add to that.

[1] This is the conclusion from my last article on the NIST ZTA document.