Compendium of Cybersecurity and Infrastructure Security Agency (CISA) Technology Evaluations

Publication: December 2023
Cybersecurity and Infrastructure Security Agency

The Cybersecurity and Infrastructure Security Agency (CISA) is identifying technologies that support the CISA Cybersecurity Strategic Plan 2024-2026. Technologies are assessed to determine if they create risk for CISA and its stakeholders; improve the efficiency and effectiveness of the CISA organization; or improve the cybersecurity of CISA stakeholders.

The Compendium provides context for the reader to understand which CISA focus areas may be impacted by the technologies (negatively or positively). The document is organized to provide the reader high level descriptions of the technologies investigated in FY 22-23, and the assessments used to develop the impacts. The report includes descriptions of 19 technologies listed in Table 1. Figure 1 and Table 2 summarize the outcomes of the technology analyses.

The technologies align with/support all three goals of the Cybersecurity Strategic Plan (2024-2026). Given the aspirational, future focus of the technologies investigated the alignment with Goal 1 (Address immediate threats) is limited to 5 of the 19 technologies. The alignment with Goals 2 and 3 is significant, with 18 of 19 technologies supporting all 3 Goal 2 objectives and one or more Goal 3 objectives. In all cases the technologies provide a supporting capability to the techniques and processes needed to produce the measures of effectiveness evidence for each Objective of each Goal in the Strategic Plan. The following diagram illustrates the alignment of the technologies to the goals.

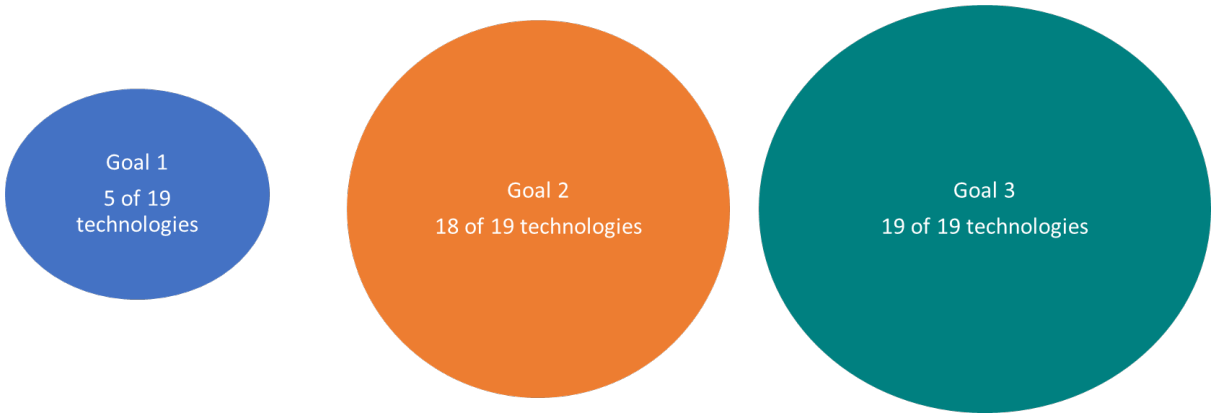

Technology Alignment with CISA Cybersecurity Strategic Plan

The assessments of technology impact urgency and importance, and alignment with CISA strategy are based on expert judgement and qualitative analysis of publicly available information. The information used to assess each technology includes technical research papers, periodicals, media reports, as well as published CISA organization priorities and focus areas. Five technologies were assessed as high urgency and high importance. The results of the assessments are detailed in section 2.

Our hope is that the Compendium becomes a tool for all of CISA to quickly reference technology descriptions and application to our priorities. CISA intends to regularly publish updates to this compendium and refine its presentation based on program needs and user feedback. We encourage all readers to provide feedback to help us improve the value of the report and broaden its readership.

Table of Contents

Contents

List of Figures

List of Tables

1 Introduction

CISA leads the national effort to understand, manage, and reduce risk to the cyber and physical infrastructure that Americans rely on every hour of every day, including the security of critical infrastructure (CI). CISA accomplishes this mission by providing security services to CI providers and establishing guidelines, recommendations, directives, and standards to encourage more secure systems. To do an effective job, CISA needs to continually assess new technologies and respond to those that create new security risks.

This document provides a compendium of the technologies the Cybersecurity and Infrastructure Security Agency (CISA) has assessed in FY22-23. The results of these assessments have been used to inform the development of new policies and knowledge transfer across government.

This compendium provides the reader insights into the topics that CISA has recently studied and provides a summary to the technology, its relevance to CISA, and the findings of research for each topic.

CISA intends to regularly publish updates to this compendium and refine its presentation based on program needs and user feedback. Please provide feedback on this report and requests for further information on the topics in this compendium directly to CISA.

1.1 Topics

Table 1 lists the nineteen topics covered in this compendium. For each topic, there is a standardized summary of the results in the following section of the document.

Table 1: Technolgies of Interest to CISA in 2023

Topic	Description
Web3 and Blockchains	Web3 and blockchains have recently become very popular. What new cyber risks do these technologies introduce to unsuspecting new users?
Large Language Model	LLMs such as ChatGPT have rapidly emerged. These systems have the potential to transform how the Internet is used and effect the ways in which millions do their jobs, but what are the security risks especially for government uses?
LLM Prompt Engineering	The art of interacting with large language model through careful structuring of a sequence of online queries has emerged as a new skill in high demand. What guidance should CISA provide to those who practice prompt engineering?
LLMs Translate C++ to Rust	Rust is a preferred memory safe programming language. Feasibility into the effectiveness of using LLMs to automatically translate existing code to Rust
Post Quantum Cryptographic (PQC) and PQC Transition	Quantum computers continue to gain power and are expected to eventually enable adversaries to crack public-key encryption used for most communications. Assessment of the risk and approaches to transition to quantum computer resistant algorithms.
Quantum Key Distribution (QKD)	Several schemes have begun to mature that enable endpoints to securely generate shared keys that use quantum physics to improve security of the process. Would this have any benefit to CISA?
Smart Manufacturing / Industry 4.0	Review of the current state of modernization of manufacturing control systems. How does this trend impact CISA, CISA stakeholders, and the threatscape?

Topic	Description
Privacy Enhancing Technologies	Available privacy enhancing technologies have been effectively applied in specific use cases outside CISA. Do these technologies and use cases inform benefits to CISA missions?
Anonymous Information Sharing	CISA needs information from CI operators about cyber incidents. What benefits might PETs have in increasing participating in information sharing programs?
SATCOM Cybersecurity	The recent conflict in Ukraine has reminded us of how SATCOM systems can be attacked. There is a renewed interest in increasing the cybersecurity of SATCOM services. What is being done and what else needs to be done?
ICS Virtualization	Assessment of the risks and benefits of virtualizing industrial control systems and moving processing into cloud computing services.
AI for ZTA	Most enterprises including the federal government are moving to Zero Trust Architecture (ZTA) to improve the security of their systems. Where within the ZTA might artificial intelligence technology be of benefit in implementing ZTA?
ZTA Technology Status	ZTA implementation is underway and continues to adapt. What is the current state of ZTA implementation at a broad scale and where is more R&D needed?
CPS-Resiliency	The Colonial Pipeline incident is a reminder that resiliency is a critical quality in the design of our complex CI systems of systems. What should the federal government be doing in this space to encourage more resiliency?
Synthetic Data	Synthetic data is increasingly being used to generate training data sets for machine learning models and to test complex systems. Can new synthetic data generation technologies benefit CISA missions?
Contract Optimization	CISA enters hundreds of contracts for products and services on an annual basis. How can automated tools be used to improve the process and ensure the government gets what it needs at a fair price?
ML Drift Detection	A known problem with machine learning systems is avoiding degradation of performance as streaming data changes over time. How can CISA R&D further approaches to detect and correct for these degradations?
Software Understanding	Tools have emerged that have advanced capability to analyze source code to characterize systems and identify opportunities for improvements. What benefits would software understanding tools have to CISA?
Digital Twin	Digital Twin technology is used to simulate physical systems and maintain their updated operational state through sensor feedback. Digital Twins are useful for predicting when maintenance may be needed or to test changes before implementing them in the live system. How can digital twins be used to improve cybersecurity of critical infrastructure systems?

1.2 Assessments of Urgency and Importance

Figure 1 below illustrates the relative urgency and importance of each topic to CISA at the time this compendium was produced. These values are subject to change over time.

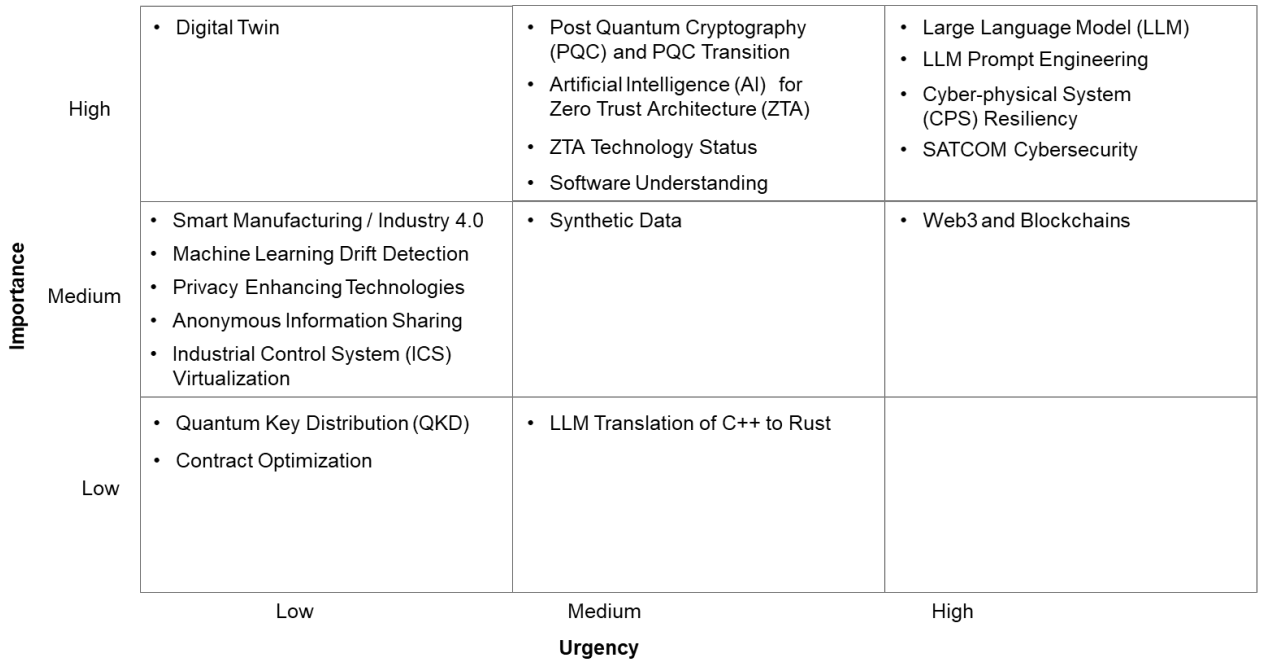

Figure 1: Urgency and Importance of Technologies to CISA

The following definitions were used in the assignment values in the diagram above for each topic.

Urgency

- **High** – Topic potentially relevant to the cybersecurity community's interests in the next 3 – 6 months.
- **Medium** – Topic potentially relevant to the cybersecurity community's interests in the next 6 - 12 months.
- **Low** – Topic potentially relevant to the cybersecurity community's interests in the future (beyond 12 months).

Importance

- **High** – Technologies where CISA action may have significant impactful improvement within the CISA mission space (e.g., improve analysis, increase information sharing, increase CI resilience, reduce cybersecurity incidents, reduce cybersecurity risk, improve NS/EP communications).
- **Medium** – Technologies where CISA may have moderate impactful improvement within the CISA mission space.
- **Low** – Technologies where CISA action may have minimal impactful improvement within the CISA mission space.

1.3 Relevance to CISA Strategic Focus Areas

The following six areas of evaluations in FY23 support the CISA Strategic Plan 2023-2025:

- **Secure by Design/Secure by Default** – "Secure-by-Design" means technology vendors integrate security principles into product requirements, as well as design and development processes, to minimize vulnerabilities malicious actors can exploit. Software manufacturers should perform risk assessments to identify and enumerate prevalent cyber threats to critical systems, and then include protections in product blueprints that account for the evolving cyber threat landscape. "Secure-by-Default" means products are resilient against prevalent exploitation techniques out of the box without configuration changes. Including, a secure configuration should be the out-of-the-box baseline, the complexity of security configuration should not be a customer problem, and manufacturers of products that are "Secure-by-Default" do not charge extra for implementing additional security configurations.[1]
- **AI Security** – Technologies that increase the cybersecurity of AI/ML systems or use AI/ML systems to enhance cybersecurity functions.
- **ZTA** - Technologies that increase the level of ZTA maturity in an enterprise or fill a gap in ZTA capabilities. "Zero trust (ZT) is the term for an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources. A zero trust architecture (ZTA) uses zero trust principles to plan industrial and enterprise infrastructure and workflows."[2]
- **Threat Analysis –** which includes the following two areas:
    - **Cyber Analytics and Platform Capabilities (CAP-C) -** Cyber Analytics and Platform Capabilities (CAP-C): An effort supporting CISA with R&D to advance malware analysis through automation and to develop technology for active defense techniques.[3]
    - **Cyber Analytics for Machine Learning (CAP-M)** - CISA Advanced Analytics Platform for Machine Learning (CAP-M): A joint effort with CISA to develop an R&D environment for advanced analytics.[3]
- **Vulnerability Management –** which includes the following three areas:
    - **Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA)** – The Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA). Enactment of CIRCIA marks an important milestone in improving America's cybersecurity by, among other things, requiring the Cybersecurity and Infrastructure Security Agency (CISA) to develop and implement regulations requiring covered entities to report to CISA covered cyber incidents and ransom payments. These reports will allow CISA to rapidly deploy resources and render assistance to victims suffering attacks, analyze incoming reporting

[1] CISA. (2023, April 13). Shifting the Balance of Cybersecurity Risk: Principles and Approaches for Security-by-Design and - Default

[2] NIST. (2020). Zero Trust Architecture. https://csrc.nist.gov/pubs/sp/800/207/final

[3] DHS. (2023, July 2). Cybersecurity / Information Analysis R&D. Retrieved from https://www.dhs.gov/science-and-technology/cybersecurity-information-analysis-rd.

across sectors to spot trends, and quickly share that information with network defenders to warn other potential victims.[4]

- **Vulnerabilities Equities Process (VEP) -** The Vulnerabilities Equities Process (VEP) is a process used by the U.S. federal government to determine on a case-by-case basis how it should treat zero-day computer security vulnerabilities; whether to disclose them to the public to help improve general computer security, or to keep them secret for offensive use against the government's adversaries.[5]

- **Coordinated Vulnerability Disclosure (CVD) Process -** CISA's CVD program coordinates the remediation and public disclosure of newly identified cybersecurity vulnerabilities in products and services with the affected vendor(s).[6]

- **Software Bill of Materials (SBOM) / Supply Chain Risk Management (SCRM) -** CISA is leading efforts to offer some common guidance and structure for the large and growing global SBOM community.[7]

Table 2 shows the mapping of each topic included in this report to each of the six focus areas. An "**X**" indicates that the topic supports the focus area.

[4]CISA. (2022, July 21). Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA) Fact Sheet. Retrieved from https://www.cisa.gov/sites/default/files/publications/CIRCIA_07.21.2022_Factsheet_FINAL_508%20c.pdf.

[5] White House. (2017, November 15). Vulnerabilities Equities Policy and Process. Retrieved from https://trumpwhitehouse.archives.gov/sites/whitehouse.gov/files/images/External%20-%20Unclassified%20VEP%20Charter%20FINAL.PDF#:~:text=The%20Vulnerabilities%20Equities%20Process%20%28VEP%29%20balances%20whether%20to,such%20as%20intelligence%20collection%2C%20military%20operations%2C%20and%2For%20counterintelligence.

[6] CISA. (n.d.). Coordinated Vulnerability Disclosure Process. Retrieved from https://www.cisa.gov/coordinated-vulnerability-disclosure-process.

[7] CISA. (2023, April 21). CISA Releases Two SBOM Document. Retrieved from https://www.cisa.gov/news-events/alerts/2023/04/21/cisa-releases-two-sbom-documents.

Table 2: Relevance to CISA Strategic Focus Areas

Topic	Secure by Design / Secure by Default	AI Security	ZTA	Threat Analysis	Vulnerability Management	SBOM / SCRM
Web3 and Blockchains	X					X
Large Language Model	X			X		
LLM Prompt Engineering	X			X		
LLMs Translate C++ to Rust	X					
Post Quantum Cryptographic (PQC) and PQC Transition	X		X			
Quantum Key Distribution (QKD)	X					
Smart Manufacturing / Industry 4.0	X	X				
Privacy Enhancing Technologies		X	X	X	X	
Anonymous Information Sharing				X	X	
SATCOM Cybersecurity	X		X			
ICS Virtualization	X					
AI for ZTA	X	X	X			
ZTA Technology Status	X		X			
CPS-Resiliency	X					
Synthetic Data				X		
Contract Optimization						
ML Drift Detection		X		X		
Software Understanding	X	X			X	X
Digital Twin	X			X		

2 Technology Topics

This section summarizes the results of research and analysis of technology topics that were relevant to CISA in FY22-23. The following section organization is used for each topic:

- **Description** – Provides a brief description of the technology.
- **Importance to CISA** – Explains why this technology matters to CISA in terms of the risks or potential benefits to CISA operating divisions; Critical Infrastructure (CI); Federal, State, Local, Tribal, and Territorial (FLSTT); and .gov and how it aligns to CISA strategic focus areas.
- **Details** – Further describe the technology, such as its architecture, key components, and how it works.
- **Findings** – Summarizes the cybersecurity risks and benefits.

2.1 Web3 and Blockchains

2.1.1 Description

“Web3” is the collection of technologies that support distributed applications such as blockchains and cryptocurrencies, which is separate and distinct from "Web 3.0" which is evolution of the "Semantic Web" that has the goal of making the information on the web more machine readable and accessible. They are mainly designed to establish trust between participants and allow them to conduct transactions while avoiding dependencies on centralized intermediaries such as banks and social media platforms. Some forecasts indicate that these technologies will dominate the method by which most things are accomplished on the internet in the future and have the potential to disrupt many industries such as those using online financial transactions, health care record management, and supply chain tracking.

2.1.2 Importance to CISA

There is much hype around Web3 technologies, and they are already in use primarily for cryptocurrency investment. However, these systems are not immune from cyber-attacks, and there have been many instances of enormous cyberthefts that have resulted in the loss of billions of dollars of value from cryptocurrency exchanges. In addition, blockchains are increasingly being considered for use in mission critical applications such as supply chain management, and a host of other information exchange applications. CISA stakeholders require that CISA understand these technologies and their risks to proactively provide guidance and recommendations to assist users of these technologies reduce the likelihood and impacts of potential cybersecurity attacks.

2.1.3 Details

The technology underpinning the evolution of distributed applications is the blockchain. Blockchains can be either permissionless or permissioned. Permissionless (or public) blockchains run on the internet and feature an architecture where all the processing and approval of transactions are open to all participants. Permissioned (or private) blockchains have controlled participation that are only open to approved participants. They are generally used within an enterprise and feature more efficiency and security controls than are generally available in permissionless blockchains. Table 3 below provides a comparison of the features of permissionless and permissioned blockchains.

Table 3: Permissionless vs Permissioned Blockchains

Permissionless (Public) Blockchains	Permissioned (Private) Blockchains
• **Anyone can participate** • **No central control or governance** • **Designed so anyone can read and validate the contents of the ledger** • **Examples are Bitcoin and Ethereum** • **Some claim permissionless networks are the only true Web3 technology**	• Controlled by an enterprise or consortium (e.g., government entity, industry group) leading to more centralized control • Only authorized users allowed to participate • Additional access control enables greater privacy features • Considered to be a more viable approach for coordinating data exchanges among cooperating entities (e.g., health records, supply chains, titles, deeds, driver licenses)

Web3 is a term used to describe permissionless blockchain-based applications that run on the internet such as cryptocurrency, distributed autonomous organizations, and non-fungible tokens. Web3 architecture contains several components each with their own level of cybersecurity concerns. Figure 2 provides a simplified diagram that shows the various components of the Web3 architecture and which components have cybersecurity concerns.

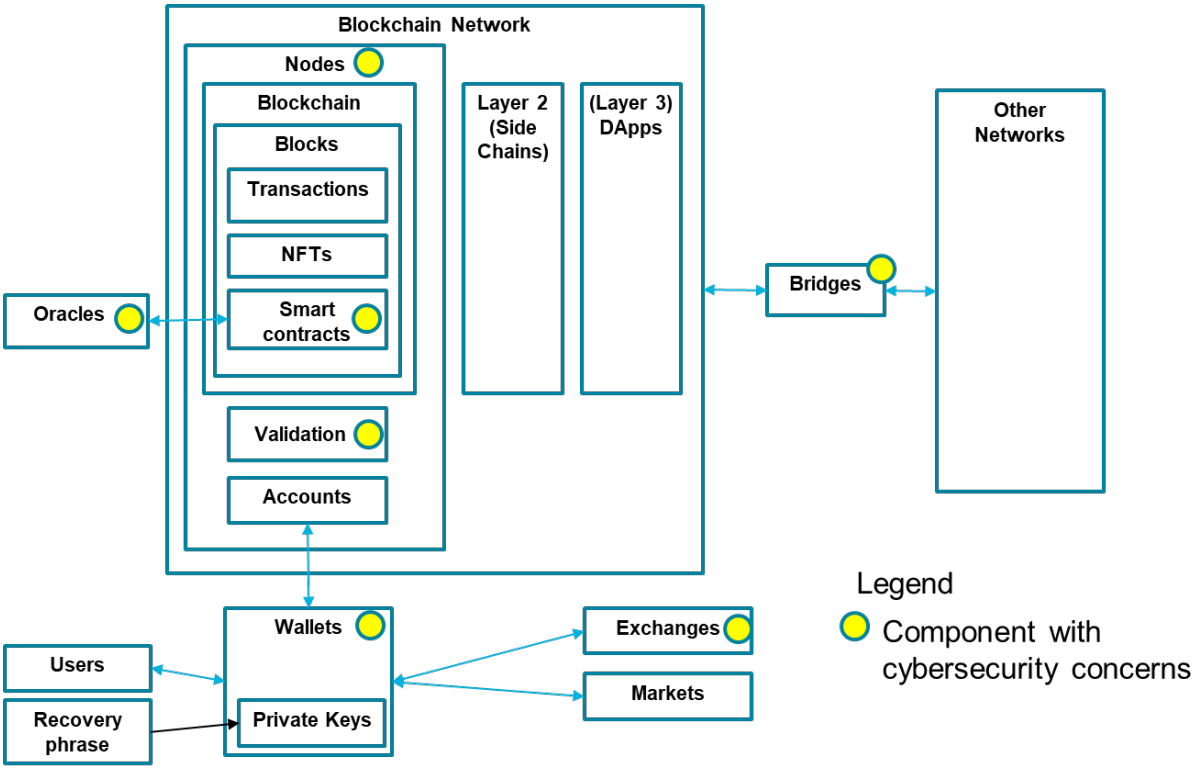

Figure 2: Web3 Architecture

Table 4 provides a high-level description of the key components and associated cybersecurity concerns.

Table 4: Web3 Key components

Component	Description	Cybersecurity risks
Nodes	A server that runs the network code, stores the ledger, and performs the functions of the blockchain that include validation of new blocks using the agreed upon consensus method (e.g., Proof of Work (PoW), Proof of Stake (PoS)), and synchronization with other nodes	Nodes are vulnerable to traditional forms of cyber-attack; however, the ledger or blockchain within the node is considered immune from attacks via encryption and consensus
Smart contracts	A program on the blockchain that represents an agreement between two or more participants that are like regular contracts except they execute automatically and cannot be changed once recorded	Smart contracts are code that can be attacked in many ways. A compromised smart contract can cause significant damage especially if it controls a significant number of transactions
Oracles	An interface between the blockchain and the external world that are used by	Oracles can be attacked through traditional attack methods. If an oracle is

	smart contracts to receive or send data off network (e.g., stock values or IoT sensors or effectors)	compromised, it can be used to trigger or change the effects of a smart contract
Validation	The method used to verify transactions and blocks on the blockchain such as • PoW – energy intensive method used in Bitcoin • PoS – energy efficient method used in Ethereum	Blockchains can be attacked if a user controls a large percentage of the miners or validators on the network. • 51% attack for POW • "Staking" attack for POS (capitalization attack)
Bridges	A system that connects two or more networks or blockchains together and provide a means for transactions to flow from one blockchain to another (e.g., Bitcoin to Ethereum)	Bridges often hold significant crypto balances on all the networks they bridge, which makes them valuable targets for adversaries. These are often attacked and have led to the loss of billions of dollars each year.
Wallets	A container for private keys and interfaces to blockchains, exchanges, and markets	Wallets are applications that can be attacked through traditional attack methods. Wallets contain private keys, which an adversary can use to access and control all assets for the associated accounts
Exchanges	A business (or smart contract) that provides a means to swap one currency for another (e.g., dollars to bitcoin)	Exchanges can be hacked or mismanaged leading to billions of dollars of loss. There are many examples such as Mount Gox

2.1.4 Findings

Web3 is susceptible to commonly understood attack vectors and is subject to the same vulnerabilities found in traditional enterprise IT implementations necessitating CISA develop expertise and operational capabilities to support stakeholders who rely on this technology. Common vulnerabilities and attack patterns remain valid for blockchain based web, which include:

- Known security issues in code; e.g., CVE, CWE
- Poorly designed or implemented architecture
- Poor operations and development practices; e.g., misconfigurations
- Inexperience of developers and operators
- Process and operations gaps and deviation from best practices

The environment is difficult to secure because it is a massive and complex space.

- Every chain is its own operating space and infrastructure that must be secured
- Every smart contract has its own unique programming language

Exploitation effort and time is low with high payoff, leading to:

- FTX collapse - $415M worth of cryptocurrency hacked from exchange accounts.
- Lazarus Group responsible for Ronin Network Attack, largest crypto heist at $624M.

Blockchains have limitations that make them less ideal for mission critical and large-scale applications.

- They do not communicate with each other
- They are slow and have scaling challenges.

These limitations lead to developer work arounds including sidechains, oracles, and cross-chain bridges.

Blockchains are inherently secure and resilient, but it should be noted that current Web3 implementations generally use traditional PKI encryption technologies that will need to be upgraded in the future with new quantum resistant algorithms to remain secure.

The apps and the bridging of tokens from one app to the other create exploitable vulnerabilities. The top crypto hacks are the result of exploiting insecure infrastructure used for sidechains and cross-chain bridges e.g., https://rekt.news/leaderboard/

Oracles exist external to the blockchain. They are subject to common enterprise network, application, and infrastructure attack patterns, and compromise can affect the proper execution of smart contracts.

Wallets are cryptographically secure; however, users are subject to phishing and social engineering attacks leading to asset loss.

Poorly implemented identity management, decentralized ID, and credential verification can lead to a Sybil Attack.

51% and PoS attacks are a concern for smaller networks, the risk diminishes as the networks get larger. These attacks are achieved through common techniques that establish access and execution.

Web3 continues to grow and evolve; the user base and financial interest is already enormous.

The future is uncertain and is being driven by free market forces and technical innovations. It is like the internet boom of the late 1990s and early 2000s.

Permissioned blockchains may have applications for critical infrastructure and some government functions, but it is not clear there are many advantages over other technologies. Some likely applications that will have an impact on critical infrastructure security include central bank digital currencies (CBDCs), energy grid transactions, and digital identities.

CISA should consider primarily focusing on monitoring Web3 due to the instability and rapidly changing environment.

CISA should consider investing in R&D to help CI operators and FSLTT to better understand the useability of permissioned blockchains for use cases such as supply chain.

CISA should consider issuing best practices and recommended security controls for implementing and securing blockchain solutions to include lessons learned from successful exploits.

2.2 Large Language Model

2.2.1 Description

A Large Language Model (LLM) is a type of AI natural language processing trained on massive amounts of unlabeled data, built on a neural network architecture with a vast number of parameters, and using deep learning to understand language data. The LLM is then finetuned for its intended purpose. For example, OpenAI's products were finetuned as follows: ChatGPT for conversational dialog, DALL-E for images, and Codex for coding.

LLMs generate responses to queries based on a probabilistic word matching determined by its algorithms. OpenAI's Generative Pre-Trained Transformer (GPT) is a neural network machine learning model. At its release, the third generation GPT (GPT-3) was the largest trained language model with over 175 billion parameters (the weights and biases of the layers within the model.) GPT-4, released in mid-March, is multimodal with visual processing to accept images in addition to text.

2.2.2 Importance to CISA

LLMs, like GPT-3 or 4, have a wide range of applicable uses. Within the Federal Government administrative uses could be generating first drafts of speeches, report outlines, summarizing articles, contract reviews, and public chatbots. Security operation uses could include cybersecurity incident reports, best security practices, network configuration suggestions, and code review or development. These use cases demonstrate the potential immediate benefits of LLMs. However, use of LLMs also pose several concerns and risks.

First and foremost, LLM responses may lack accuracy. This characteristic is inherent in their probabilistic design, but can also be caused by out-of-date training, or inappropriate training data (query is out of scope). LLMs can also be biased if their training data was biased, or the underlying algorithms are biased. LLMs may also raise privacy concerns. LLM queries expose sensitive information, especially relationships among agencies and technologies, if stored or used to train the model for future queries.

There are other ethical considerations when using LLMs such as possible legal exposure due to use of unlicensed copyrighted material in their training data. The inner workings of how LLMs provide query responses is not visible to the user. This lack of transparency can cause misgivings and lack of trust in the LLM output. Then there is the unknown ability of adversaries to perform attacks on the system, including the potential for malicious actors to manipulate the model's output (e.g., poisoning and contamination).

2.2.3 Details

LLMs hold the promise of large benefits. They may save time and money, for example, by reducing the workload of customer service representatives, improving response times, and enhancing customer satisfaction. With GPT-3, and ChatGPT in particular, OpenAI showed that an LLM trained on enough data can solve natural language processing tasks that it has never previously encountered. The GPT-3 model is a general solution for many downstream jobs without the need for fine-tuning or additional layers on top of the language model.

Even given these potential benefits, LLMs have limitations that must be mitigated: they cannot handle complex tasks; understand the nuances of human language and emotions; or provide 100% accurate responses. Also, there are potential ethical, privacy, and legal concerns on the use of LLMs.

Many companies are building LLMs, and numerous LLMs are under internal company development without a public announcement. An example of LLM's exponential growth since their introduction is presented in Figure 3.[8] More research is needed to better understand LLM's limitations and impacts to not only enterprise security but also to national security.

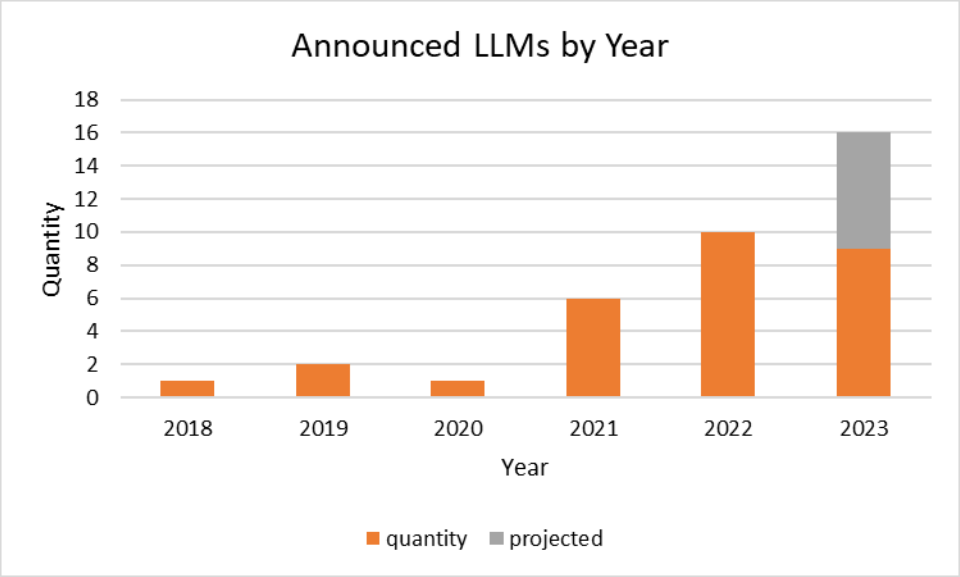

Figure 3: Rapid Development of LLM

2.2.4 Findings

CISA should consider the development of recommendations to guide FCEB use of LLMs for internal and public service use cases. LLM use recommendations should consider a number of items:

1. Education of the reader on how LLMs are experimental, that there is no guarantee of accuracy, privacy, or un-biased responses.
2. Publicly available LLMs should be carefully limited or not used. If used, human validation of all responses for accuracy and current status is needed.
3. Contractual arrangements with LLM owners should be considered to protect sensitive data and intellectual property.
4. Government developed or owned LLMs may overcome some risks with privacy, bias, and ethical concerns. However, human validation of all responses for accuracy and current status would still be needed. Cost-benefit analysis is needed given the high-cost of training an LLM.
5. Upfront planning on CONOPS is needed. Different use cases and user bases will present different risk tolerances and assurance levels. One overarching use policy may not be feasible.

[8] Wikipedia. (2023). Large Language Models. List of large language models released from 2018 through July 2023. Retrieved from https://en.wikipedia.org/w/index.php?title=Large_language_model&oldid=1166453082

6. Need to understand the model being used for applicability. Many government use cases will go beyond the general parameters of public demos (e.g., ChatGPT and Bing Chat)

In addition to stakeholder recommendations, CISA needs to monitor and research the increasing risks that LLMs may pose to national security. LLMs can be a valuable tool for adversary use through improved social engineering; faster development of malware and cyberattacks; and better content for mis, dis, mal-information. There is also the potential to use LLMs for analyzing large volumes of text data, such as emails, chat logs, and social media posts, to gain insights into the activities and plans of foreign governments, organizations, or individuals. LLMs thus become a valuable espionage tool.

2.3 LLM Prompt Engineering

2.3.1 Description

Prompt Engineering (PE) is the process of developing and optimizing prompts for more efficient communication with an LLM. Prompts refers to the queries that are used to interact with an LLM. Utilizing PE techniques when creating prompts can greatly increase both the quality and accuracy of the LLM's output.

2.3.2 Importance to CISA

PE is an evolving field of study. Recent studies have shown that PE can improve the accuracy and context of the LLM response. However, PE does not remove the risk of LLMs generating inaccurate information or hallucinating. Therefore, CISA should not consider relaxing LLM use recommendations and guidance when PE is used by it or its stakeholders.

2.3.3 Details

LLMs have the potential to provide massive benefits. By automating menial tasks (such as first draft generation), they can increase workplace efficiency by allowing employees to focus on more skilled tasks. Compared to traditional AI models which are each only capable of performing a limited range of tasks, LLMs can solve a much wider range. PE has been shown to dramatically increase the performance of LLMs when solving a variety of tasks. This improvement is most easily shown in logic-based tasks such as solving a riddle, solving a math problem, or writing code. However, it also has merit in tasks such as summarizing a document or explaining a topic. These findings, combined with further research, may create benefits if LLMs are integrated into government processes.

2.3.4 Findings

PE can have an impact on AI security. Poorly engineered prompts that return bad information, reduce the value of LLMs. LLM users with PE training and experience can achieve superior results compared to using basic prompts. This improvement is demonstrated in the following graphics which show the results of various PE techniques in solving several problems: Figure 4,[9] Figure 5,[9] and Figure 6.[10]

[9] Yao et al. (2023). Tree of Thoughts: Deliberate Problem Solving with Large Language Models. arXiv:2305.10601

[10] Long. (2023). Large Language Model Guided Tree-of-Thought. arXiv:2305.08291

Method	Success Rate (%)		
	Letter	Word	Game
IO	38.7	14	0
CoT	40.6	15.6	1
ToT (ours)	**78**	**60**	**20**
+best state	82.4	67.5	35
-prune	65.4	41.5	5

Figure 4 Crosswords

Method	Success
IO prompt	7.3%
CoT prompt	4.0%
CoT-SC (k=100)	9.0%
ToT (ours) (b=1)	45%
ToT (ours) (b=5)	**74%**
IO + Refine (k=10)	27%

Figure 5: Game of 24

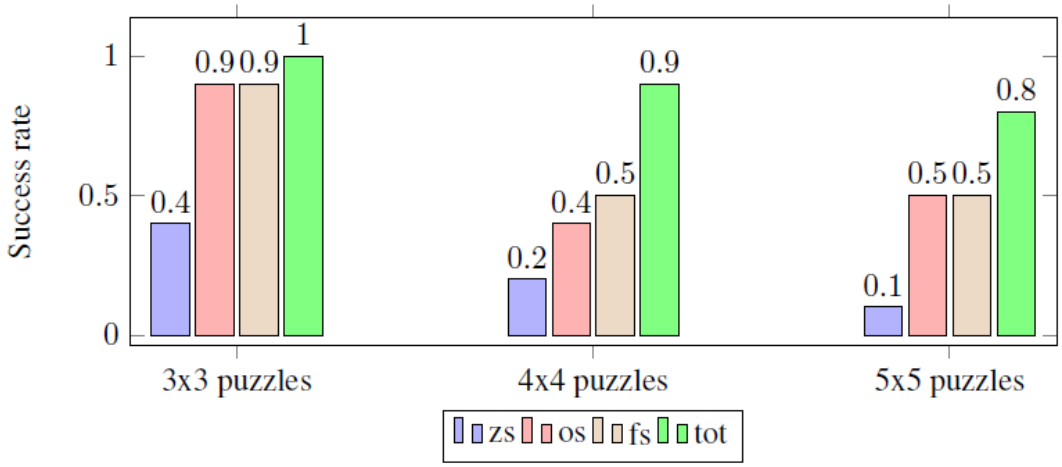

Figure 6: Sudoku

In study 1: IO refers to a basic prompt, CoT to chain-of-thought, CoT SC to chain-of-thought with self-consistency, and ToT to tree-of-thoughts. The images illustrate more advance prompting techniques yield significantly higher performance.

In study 2: zs refers to zero-shot or no user provided example, os to one-shot or one example, fs to few-shot or several examples, and tot to tree-of-thoughts. Once again, by providing more examples or using advanced techniques such as tree-of-thoughts, significant success rate improvements are possible.

2.4 LLMs Translate C++ to Rust

2.4.1 Description

LLMs are powerful tools. They have the promise of providing numerous benefits such as generating text-based products, automating conversational task, creating simplified summaries, mocking up websites, and performing software development functions. This section looks at the specific use of LLMs as the next iteration of tools for automating the translation of software from one programming language to another.

2.4.2 Importance to CISA

Software translation tools are attractive since they can decrease the cost of porting software from one language or environment to another. These types of tools can save time and expedite developers learning new programming languages. They can standardize development by creating consistency across large code bases and augment code by developing documentation and comments. LLMs are no different in this regard. By being easier to use, LLMs could also facilitate an explosion of new code – large number of lines of code in a short period of time.

However, LLMs are not infallible. The translations and suggestions generated by LLMs should be treated as aids rather than absolute truth because LLMs may introduce translation inaccuracies and new vulnerabilities due to translator characteristics noted in section 2.4.3. It is important to validate and verify the translated code manually, especially for critical or high-security systems. Relying solely on LLMs without human expertise can introduce errors, leading to incorrect functionality or vulnerabilities in the translated code.

2.4.3 Details

A generalized view of using LLMs in translating one programing language to another follows.

LLMs, including OpenAI's Codex, lack domain-specific knowledge or explicit programming language rules. Nevertheless, they can still assist in code generation and translation by their ability to infer patterns and generate syntactically correct code. This approach is the same method that LLMs use to process any written and visual inputs.

LLMs are not compilers or programming language specific editors with inherent, embedded language rules. LLMs are predictive algorithms based on tokenized patterns, syntax, and grammar or rules learned during the LLM training and fine tuning. The lack of embedded rules is an important distinction from legacy translation tools.

Nonetheless, LLMs, including ChatGPT and Codex, can be leveraged to automate aspects of programming language translation:

1. LLMs can generate code examples and snippets in different languages. Developers can request examples or specific translations from LLMs to understand how certain constructs can be implemented in other languages. This feature is particularly valuable for programmers who need guidance in a particular language's patterns and practices.

2. LLMs can assist in identifying potential issues or incompatibilities when translating from one language to another. By analyzing the code and generating suggestions, LLMs can help developers make informed decisions during the translation process. They can highlight language-specific patterns and suggest alternative translations.

3. LLMs can generate initial translations. While these translations may not be perfect, they provide a starting point that developers can refine and adjust. This approach may save time and effort, allowing developers to focus on more critical aspects of the codebase.

2.4.4 Findings

Translating specifically from C++ to Rust highlights LLM benefits and drawbacks in code development.

Rust is gaining a lot of attention due to its performance capabilities and ownership model that guarantees memory and thread safety. Since Rust is a newer programming language, it is assumed that most development shops lack deep Rust skill sets. Augmenting skills is where the use of LLMs can increase productivity and efficiency when coding in, or translating to, Rust.

Github's Copilot is an LLM editing assistant, that can virtually look over the shoulder of the developer. Copilot is an add-on to Github's coding editor. It assists developers as they type. Copilot automates processes saving developers time and effort. Copilot can also automate code generation through integrated development environment platforms such as Visual Studio Code, Visual Studio, and JetBrains. It makes suggestions across dozens of programming languages, including C++ and Rust.

While LLMs offer the advantages mentioned above, there are challenges and limitations to consider when using them for C++ to Rust translation:

- C++ and Rust have different semantic models and language features. LLMs may struggle to capture the precise functionality of C++ code and thus generate non-equivalent Rust code. Additionally, syntactic differences between the languages can lead to direct translations that do not fully capture the functionality of the original code or follow Rust's idiomatic style. This discrepancy in direct translations is a problem when using LLMs to translate more than snippets of code.

- Producing idiomatic Rust code from C++ requires careful review, refactoring, and adjustment. Developers must ensure that the translated code takes full advantage of Rust's features and conventions to achieve optimal readability and maintainability. While LLMs may assist in translating snippets of code, often C++ applications need to be redesigned or rearchitected to make proper use of Rust's features and strengths.

- C++ and Rust have complex features. Edge cases, pieces of code that use complex C++ constructs, can pose challenges for LLMs. Advanced C++ features, such as templates and macros, may not have direct equivalents in Rust. LLMs will still attempt the translation, but these situations will require manual intervention. LLMs will also struggle with nuanced scenarios or language-specific intricacies, leading to inaccurate or incomplete translations. Only through extensive testing and manual code review will these cases be identified and corrected.

- LLMs prioritize generating code that produces the desired output but may not account for efficiency and performance considerations. Translated code might lack optimization or not fully leverage Rust's memory safety guarantees. Thorough testing and manual inspection by experienced developers are essential to ensure correctness and performance.

- With the current state of the practice, human expertise plays a vital role in the translation process. Experienced programmers should review and adjust the translated code, ensuring it aligns with the new language's proper style, correctness, and performance requirements. Manual intervention helps overcome LLM limitations and ensures the translated code meets high-quality standards for safety, security, and functionality.

LLM use for everyday tasks is still an emerging field. Though LLMs have promise of better assisting or someday automating programming language translations, more research in these types of use cases is needed.

2.5 Post Quantum Cryptography (PQC) and PQC Transition

2.5.1 Description

Post Quantum Cryptography (PQC): There has been substantial development on quantum computers, which are described by the National Institute of Standards and Technology (NIST) as "machines that exploit quantum phenomena to solve mathematical problems that are difficult or intractable for conventional computers." The integrity of our current public key cryptography relies on the infeasibility of conventional computers to break the current encryption of data at rest and in transit, digital signatures and user, device, and application authentication. A cryptography relevant quantum computer (CRQC) represents a national security threat because it will have the processing capability to break the public-key cryptosystems currently used and the PKI capabilities built into zero trust architectures. PQC aims to develop cryptographic systems that are secure against both quantum and classical computers but will still be compatible with our communications and network infrastructures.

A CRQC large enough to decrypt currently encrypted information is most likely years away from development; estimates range from 2026 to 2041, with 2030 to 2035 being estimated as an inflection point for quantum technologies to start to significantly impact communication and information technology. Despite that timeline, there is evidence that adversaries are currently capturing encrypted data for later decryption, once CRQC's are available. Realizing this situation, OMB issued M-23-02 directing agencies to comply with National Security Memorandum (NSM-10), *Promoting United States Leadership in Quantum Computing While Mitigating Risk to Vulnerable Cryptographic Systems* (May 4, 2022).

PQC Transition: NIST is expected to finalize the standards for PQC algorithms and tools to resist CRQC by the end of 2024, with products anticipated within 12 months of the publication of NIST's PQC standards. Agencies cannot transition to a PQC environment until NIST finishes development and standardization of PQC algorithms and compliant commercial products are available. However, they can develop transition plans, new policies, processes, and testing methods in anticipation of a PQC transition requirement.

2.5.2 Importance to CISA

As the deadline for NIST to release PQC standards is quickly approaching, CISA needs to play a central role in PQC transition across .gov, FSLTT, and CI sectors, which includes a role coordinating with COTS vendors for status and progress. CISA has formal roles around PQC and PQC transition as defined in OMB Memorandum M-23-02. CISA will also have responsibility for supporting the DHS PQC transition as it is described in DHS Policy Directive 140-15 and subsequent DHS CIO Memorandum.

2.5.3 Details

Post Quantum Computing: The RSA public key cryptosystem is the most widely used encryption system. An RSA Public Key is created using two large prime numbers (which are kept private), plus an auxiliary number. Anyone can encrypt a message using the Public Key, but the two large prime numbers (Private Keys) are needed to decrypt the message. RSA public key security relies on the fact that it is easy to multiply two large prime numbers to create a single large number (e.g., the Public Key used for encryption), but it is computationally difficult to reduce a large number into its prime factors (e.g., the Private Key needed for decryption). It is estimated that it may take a conventional computer up to 1

billion years of processing an RSA-2048-bit Public Key to factor it into its prime numbers (source: Dr. Krysta Svore, Microsoft Research). Current Public Keys encryption systems utilize 2048 or 3072 bits.

A CRQC computer will eventually be developed that will be capable of breaking current public key cryptography in minutes or hours. The threat to RSA Public Key posed by a CRQC comes from Shor's algorithm, which is a quantum computer algorithm developed by American mathematician, Peter Shor. On a sufficiently large CRQC, Shor's algorithm can exploit quantum parallelism and constructive interference to factor RSA Public Keys back into their prime factors (Private Keys), thus breaking the encryption, as shown in Figure 7.[11]

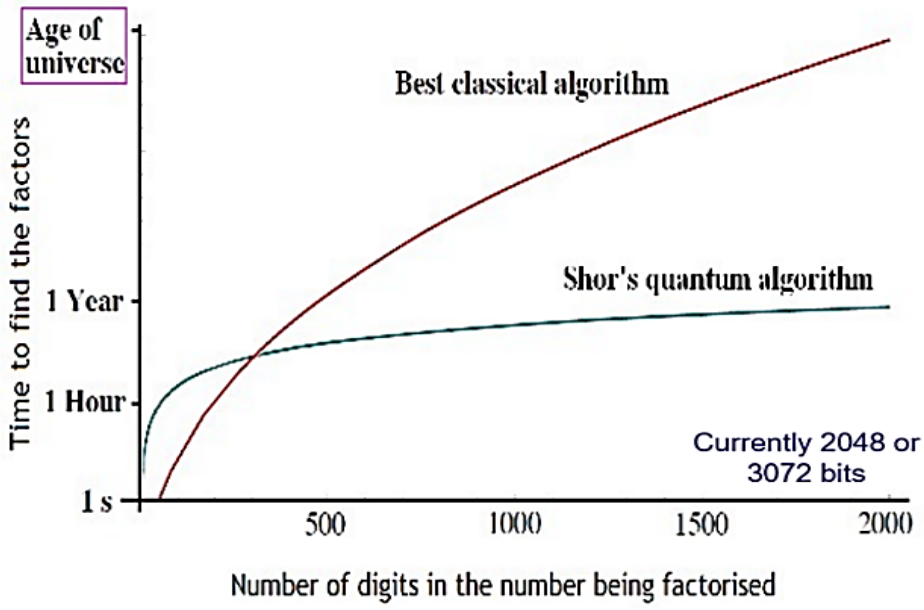

Figure 7: CRQC Using Shor's Algorithm Breaks Current Public Key Encryption in Usable Time

As a matter of implementation, RSA is a relatively slow way to encrypt data because of the math-intensive processing involved with each encrypted message that is sent/received. Practically, RSA is used between end points to transmit symmetric keys which are then used for bulk-encryption since symmetric keys are mathematically easier, and thus faster to use for encryption/decryption. In this scenario, a future CRQC-enabled attacker would intercept the RSA-encrypted symmetric-key exchanges, use Shor's algorithm to decrypt the symmetric-keys, and then decrypt the symmetric-key encrypted dataflow between the two end points.

In 2016, as a response to this threat to our Public Key Infrastructure (PKI), NIST started investigating new standard(s) for encryption methods that will not be susceptible to being broken by Shor's quantum algorithm on a CRQC. This NIST project is named PQC Standardization. NIST expects to publish the initial PQC standards by the end of 2024. It is estimated that PQC compliant commercial products could start being delivered within 12 months of the standard being published. In 2021, NIST National Cybersecurity Center of Excellence (NCCoE) initiated a Migration to Post Quantum Cryptography project which is

[11] Sihare, S, & Nath, V. (2017, February). Analysis of Quantum Algorithms with Classical Systems Counterpart. I.J. Information Engineering and Electronic Business, 2017, 2, 20-2. Retrieved from https://www.mecs-press.org/ijieeb/ijieeb-v9-n2/IJIEEB-V9-N2-3.pdf

developing white papers, playbooks, and demonstrable implementations for organizations to use to implement PQC technology. NIST is also considering a “hybrid mode” as an interim stage to reaching full PQC-only PKI.

PQC Transition: DHS has partnered with NIST to address PQC transition issues as part of DHS CISA’s specific roles to assist the Federal government, *writ large*, with PQC transitioning per the United States Office of Management and Budget (OMB) Memorandum M-23-02 as shown in Figure 8.

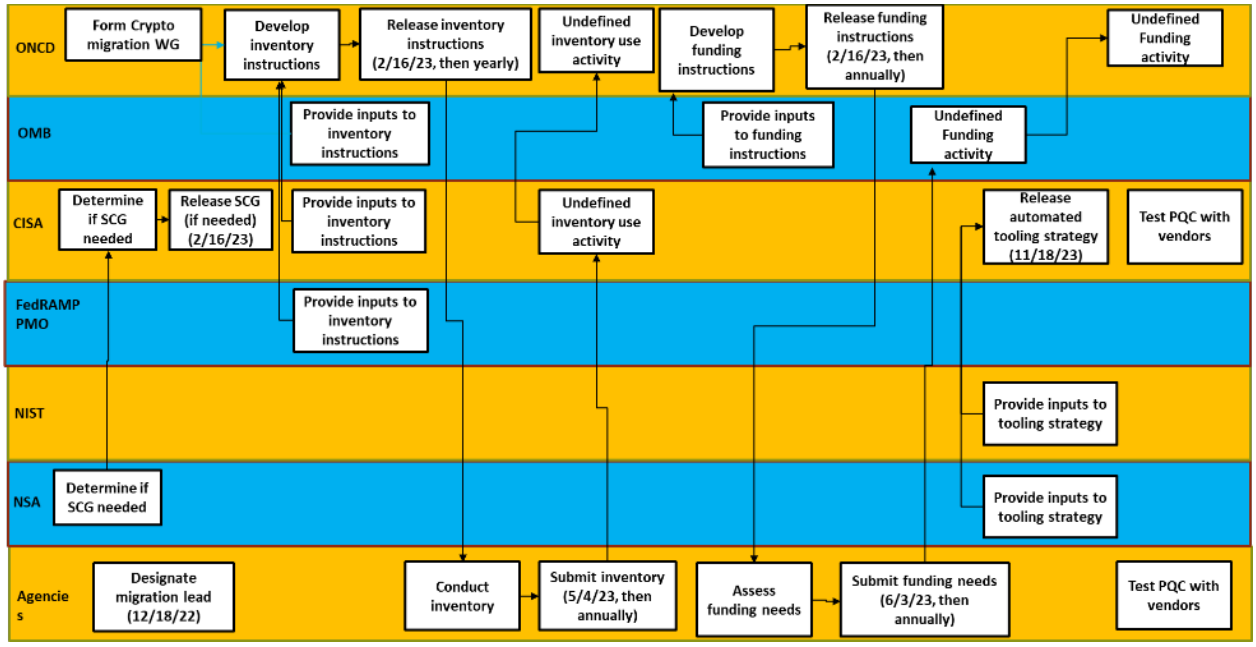

Figure 8: PQC Swim-Lane Activities Based Upon M-23-02

CISA needs to play a central role in PQC transition across .gov, FSLTT, and CI sectors, as well as coordinating with COTS vendors for status and progress. Per OMB Memorandum M-23-03, CISA has formal responsibilities around PQC and PQC Transition. The following table describes potential activities CISA could perform to support their M-23-03 roles:

Table 5: CISA Roles in PQC Transition

CISA Role in PQC Transition per OMB Memorandum M-23-02	Potential Related CISA Support Activity
Participate in Office of the National Cyber Director (ONCD) Crypto Migration working group	CISA has the opportunity to take a leadership role in the transition process in order to influence the overall approach being used to make it more effective. Through the working group (WG), CISA can gain insight into what other agencies are doing to encourage collaboration and feedback on CISA developments related to PQC transition
Work with Vendors to Test PQC Products and Solutions	CISA can leverage its relationship with NIST to participate with NIST/NCCoE work with PQC vendors. CISA can interface with testing through participation with the PQC consortium, Real World Crypto (RWC) Symposium, which is organized by the International Association for Cryptologic Research (IACR). https://rwc.iacr.org/. CISA should develop DHS agency specific use cases and impediments implementations for PQC technology, including transition / implementation scenarios that are specific to DHS Agencies, which it presents to vendors through the NIST/NCCoE and the PQC consortium.
Work with NIST on "tooling strategy"	CISA can work with NIST to provide agencies the tools needed to implement steps in the PQC transition, including tools to: scan and inventory crypto components within systems and enterprise infrastructure; prioritize and plan PQC transition activities; identify cryptographic elements; collect and consolidate PQC survey data; merge and analyze data about agency systems to be used for prioritization and planning; track and report PQC transition status to OMB and ONCD as required
Provide inputs to ONCD on survey instructions	CISA will provide technical support to ONCD to ensure the data needed to support PQC transition activities is requested via an annual worksheet. To improve the survey process, CISA should provide a means and instructions for agencies to report directly into an online system rather than via worksheets (e.g., Cyberscope). The online system should have capabilities that support collection of richer data about system architecture and data stores – may be available in the process to report an HVA, or integrated with Endpoint Detection and Response (EDR) to provide data

Additional CISA responsibilities that are derived from the responsibilities assigned in M-23-02 include:

Table 6: Derived CISA Roles in PQC Transition

Derived CISA Role in PQC Transition from OMB Memorandum M-23-02	Potential Related CISA Support Activity
Provide general guidance for PQC transition	CISA's role is to educate and provide helpful tips to enable agencies to understand the needs and formulate plans to transition to PQC. This includes leveraging existing guidance on PQC transition planning steps and PQC prioritization approaches
Track and report PQC transition status across the federal enterprise	CISA will provide OMB, ONCD, and the White House an authoritative source of PQC transition status across the entire federal enterprise. This can be accomplished via automation for agencies using Federal Information Security Modernization Act (FISMA) reports to Continuous Diagnostics and Mitigation (CDM) or via survey spreadsheets for those systems not having CDM
Provide specific transition support to agencies	Agency specific support can include development of prioritization criteria and algorithm for high-value asset (HVA) and non-HVA systems. CISA can utilize the annual survey data and other data sources to provide recommendations to agencies on what components to prioritize and how to sequence transitions. Furthermore, CISA might consider establishing a center of excellence and building a team of subject matter experts (SME) who can be dispatched to and consulted by agencies that do not have sufficient in-house expertise
Review of data already being collected and its capability to support analysis	CISA needs to ensure collected data supports PQC transition processes and decision making. To do so, CISA should review PQC Inventory data and HVA List data to determine sufficiency for analysis. CISA can use the FISMA ID and HVA ID to correlate data on systems and determine whether new data should be added to list. In order to prioritize PQC algorithm implementation, CISA should ensure that PQC transition priorities are reflected in HVA scores, modifying scoring for roots of trust. CISA should work with HVA Program Management office (PMO) to determine whether emphasis should be renewed by OMB and Office of Science and Technology Policy (OSTP) to generate, update, and maintain a valid HVA list

2.5.4 Findings

CRQC poses a significant security risk to U.S. national interests. While development of a CRQC large enough to break current PKI encryption is years away, adversaries are stock piling encrypted traffic in anticipation of eventually being able to decrypt the data. NIST is developing PQC algorithms that will mitigate the risk posed by CRQC and Shor's algorithm to decrypt our national secrets, with release anticipated in late 2024. While it is impossible to gauge the impact of the (eventual) data breach (involving dated-by-still-classified information), DHS CISA needs to assist the PQC transition process in partnership with NIST and other organizations so that agencies are prepared to implement PQC encryption technology once it is available.

2.6 Quantum Key Distribution

2.6.1 Description

Quantum key distribution (QKD) is a method of sharing a cryptographic key using properties founded in quantum physics to exchange keys in a way that is provable and guarantees security. A quantum computer is not needed to use QKD. QKD does rely on some of the same equipment used in quantum computers for generating and detecting photons and qubits. QKD does not distribute keys, information, data, or video; but rather QKD is used to generate symmetric encryption keys on two endpoints (QKD is only used to generate keys between two end-points – it cannot create shared keys among multiple endpoints). QKD does not use PKI to distribute or generate symmetric encryption keys. As such, it is not susceptible to the same security concerns as current PKI. QKD utilizes quantum mechanics properties to ensure that the generated key has not been observed or altered in transit.

The QKD can occur via fiber optic cable, over-the-air, or satellite nodes. The latter two methods using line-of-sight laser transmissions. Occasionally, QKD is mistermed quantum cryptography only because it is the best-known use case of a quantum cryptographic task. However, the key being shared or distributed is generated using a traditional computer algorithm, such as Advanced Encryption Standard.

DARPA may have been one of the first organizations to establish a working QKD network. That network was operated from 2004 to 2007 with 10 nodes in the greater Boston metropolitan area. Many of the practical limitations mentioned below were likely identified during this trial period.

2.6.2 Importance to CISA

As the use of quantum computers becomes more frequent and presents a threat to current communications infrastructure, QKD may become more relevant to CISA. A resilient defense-in-depth strategy may be needed prior to, and during, industry's transition to post quantum cryptography.

Of concern to CISA, China is the world leader in QKD both in technology advances and standards influence. In early 2021 they announced a network for QKD covering thousands of miles or kilometers. It links four quantum networks in cities in eastern China (Beijing, Hefei, Jinan, and Shanghai) with a remote location in the far west of the country. The system comprises a 2,000 km fiber optic link between the eastern cities and a satellite link spanning 2,600 km between two observatories – one east of Beijing and the other near China's border with Kazakhstan. China has also established a satellite-based quantum node that used QKD to connect China and Austria.

China and the European Union have made heavy investments in QKD while the U.S. has not emphasized QKD to date outside of the Department of Energy (Los Alamos National Labs). The NSA currently advises against the use of QKD. They view post quantum cryptography as a more cost effective and easily maintained solution than QKD. However, as China and the EU continue work and research in hardware development, protocol implementation, standards, and communication technology, QKD solutions may mature to become more important to CISA.

2.6.3 Details

QKD requires a quantum channel and a standard data communications channel between two locations as shown in Figure 9.

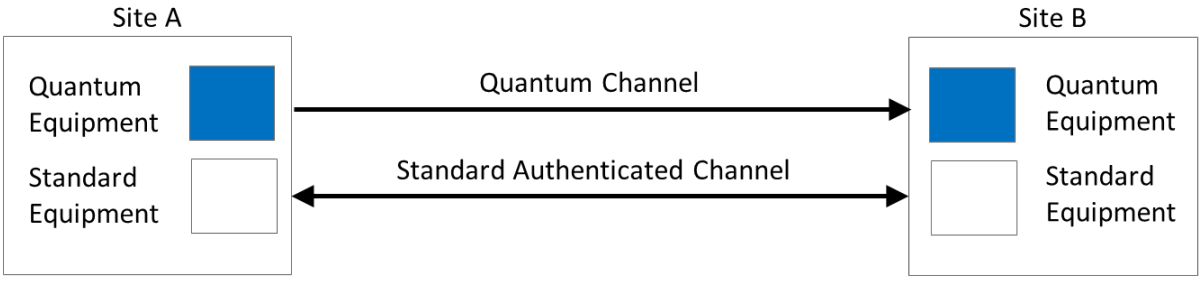

Figure 9: QKD Generates Keys Over Quantum Channel; Sends Encrypted Data Over Standard Channel

An attractive property of QKD is the ability of the two communicating parties to easily detect an eavesdropper. A fundamental property of quantum mechanics is that measuring a quantum system disturbs the system. A third party in trying to eavesdrop on the key must in some way measure the system. This measurement enables detection of the eavesdropping, and thus, makes the distribution of symmetric keys very secure.

In this manner, the security of the key is guaranteed by quantum physics instead of the traditional computational difficulty of algorithms. This type of security is still theoretical. More research is needed to determine if bypassing or masking measurements can be introduced to steal the cryptographic keys. In fact, if the quantum hardware is not built properly there could be a vulnerability in the system potentially allowing for the encryption key to be copied.

There is no one single QKD implementation architecture. Current implementations utilize different:

- Schemes for creating quantum information
- Protocols for generation of symmetric keys
- Standards for hardware fabrication
- Methods for error correction

QKD schemes, protocols, standards, and methods cannot be mixed and matched. Both parties involved in the key exchange must be using the same technology stack.

There are significant drawbacks to QKD in its current form, including:

Authenticated Communications Channel – QKD, by itself, does not have a method to authenticate the transmission source. QKD relies on an already authenticated communication channel. To authenticate the communication channel, a classical key exchange must take place. Assuming the already exchanged keys have proper strength, there is sufficient security without the need for QKD.

Denial of Service - As stated above, an eavesdropper can be easily detected thus terminating the key exchange. If an eavesdropper's intent is to disrupt communication rather than to steal communication, a denial-of-service attack becomes easier to implement. Any communication channels relying only on QKD could be brought to a standstill.

Significant QKD Infrastructure Cost – The attraction of a secure and private means of communication that does not rely on asymmetric encryption has attracted worldwide attention and investment. However, QKD equipment is expensive to obtain, patch, and maintain. Given the current cost and size of quantum hardware, QKD is only used to distribute the key and not to transmit any message data. That is, once the key is exchanged, the subsequent encrypted data is sent over a standard communication channel. This process may help to protect today's communications from an intercept-and-store adversary for future decryption using a CRQC; however, for the time being, QKD cost and complexity limit its applicability.

Physical Limitations of Current Infrastructure/Hardware - Besides imperfections in photon detectors or QKD protocol implementation hardware, there are the limitations of current fiber optic cable infrastructure. Fiber optic cables have a limited distance they can carry a photon. For typical commercial fiber, these distances are in the hundreds of kilometers although, as noted, China has been able to extend this range to about 2,000 km and incorporate satellite links. Despite these advancements, widespread use of QKD may be difficult to achieve due to the limited distance when using current fiber optic cable infrastructure, and inherent line-of-sight distribution mechanics.

2.6.4 Findings

QKD is a highly secure method of exchanging encryption keys, by-passing many of the man-in-the-middle key-exchange exploits, including the ones posed to our current PKI encryption system by the impending development of a CRQC using Shor's algorithm. However, a QKD system is expensive and has limited range at the moment resulting in limited applicability. QKD systems are also vulnerable to misconfigurations and exploits on the non-quantum portion of the system. Therefore, QKD systems still require extensive security controls to ensure that the encryption keys being generated are not compromised.

2.7 Smart Manufacturing/Industry 4.0 Cybersecurity Concerns

2.7.1 Description

Smart Manufacturing is a set of principles, techniques, and technologies that improve the profitability, monitoring, and control of the manufacturing processes beyond the current state-of-the-practice manufacturing, Industry 3.0. This combination of principles, techniques, and technologies collect and analyze data from the "top floor to the shop floor,"[12] by integrating newly available data from new sensors, and new and updated machines that create and test the products being manufactured. The data from the "shop floor" is combined with quality control, marketing, 3rd party suppliers (supply chain members), sales, and other resource data, such as energy consumption, to provide the manufacturer with continuous visibility into the state of the manufacturing process as well as the overall company performance (revenue projections, customer demand, resource usage efficiency, and product manufacturing flexibility). This topic was investigated to gain understanding of potential cybersecurity concerns associated with the transition from Industry 3.0 to Smart Manufacturing/Industry 4.0.

[12] Forbes. (2023). From The Top Floor to The Shop Floor, Planning Needs to Be Aligned. Author: Steve Banker. Retrieved from https://www.forbes.com/sites/stevebanker/2023/04/27/from-the-top-floor-to-the-shop-floor-planning-needs-to-be-aligned/?sh=4aa683e87d71

2.7.2 Importance to CISA

The Critical Manufacturing Infrastructure Sector of the nation is adopting Smart Manufacturing technologies to support the goals stated above. The principles, technologies, and techniques used to achieve Smart Manufacturing change the cybersecurity attack surface. CISA must understand the changing attack surface to support manufacturers as they look to CISA for security recommendations to address threats to their operation. The attack surface changes include: 1) artificial intelligence and machine learning analytics used for decision support and automated responses to process or machine changes, 2) increased number of networked devices, 3) data in transit among the network devices, 4) business IT to manufacturing operational technology communication bridges, 5) integrations with 3rd party suppliers (upstream and downstream supply chain participants), and 6) a lack of skilled cybersecurity and risk practitioners experienced with Industry 4.0 technologies.

The impact of cyber-attacks on Smart Manufacturing implementations may be obvious, or subtle. Obvious impacts may include reduced production capacity. Obvious impacts may also include shutting down a plant, shutting down a manufacturing line in a plant, or changing a process in a way that the end product is unusable or fails quality testing. Subtle impacts may include product or manufacturing process changes difficult for the quality management or process control systems to detect such as quality changes that reduce the useful life of a product, cause a product to fail under specific (untested) conditions, or insert (difficult to detect) malicious software functionality.

2.7.3 Details

Smart Manufacturing is "an optimized connected manufacturing facility, which can facilitate launching new products depending on market dynamics; is scalable enough to meet demand variation for existing products; is able to produce finished goods at least cost; has smart machines, sensors and robots which are seamlessly integrated with information system architecture to enable high level of automation in transaction processing; and has real time analytics that helps in minimizing downtime and improving efficiency. … A Smart factory creates an ecosystem where there is a strong collaboration between all the key players; e.g., suppliers, operations, Information Technology (IT), planning, sales & marketing, and customers. It creates a single platform where multiple business functions such as procurement, planning, manufacturing, sales & distribution, finance, and accounting teams work together to meet overall corporate objectives."[13]

Conversely, Industry 3.0 manufacturing uses siloed automation in manufacturing. For example, robots operate as independent automated process tools, replacing humans for various repetitive tasks. Each robot is designed and programmed to perform a single task at an assembly station. Industry 4.0 extends the concept of automation to integrate all the automated processes used to manufacture a product or set of products. An Industry 4.0 robot may perform multiple tasks/operations, at a single assembly station (multiple welds, or drill hole and insert fastener), or perform a task that changes depending on the process plan during its step of the process (e.g., manufacturing multiple products in serial on a single line). The specific task for each Industry 4.0 robot at its assembly station is orchestrated by a central process control system. The extension also includes increased monitoring of the processes.

[13] IEEE. (2018). Practical Guide to Smart Factory Transition Using IoT, Big Data and Edge Analytics. Retrieved from https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8478188

Extending automation and increased monitoring (via sensors) enables the operator to analyze and model the entire manufacturing process. The manufacturer analyzes all the processes within the product manufacturing lifecycle as a system of systems. This approach enables modeling and feedback for process changes in near real time as well as new flexible/customized product manufacturing. The modeling and feedback may be implemented using a Digital Twin of the manufacturing processes. An example of flexible manufacturing is a Volkswagen plant using Smart Manufacturing principles, techniques, and technologies to produce multiple car models in the same plant on the same manufacturing line. Each car is manufactured to a unique specification programmed into each process along the product manufacturing life cycle.[13]

Smart Manufacturing or Industry 4.0 is the revolution to enable manufacturers to control and optimize each operation in a manufacturing process by collecting data from sensors and machines along the entire manufacturing process. The data feeds support analysis tools: monitoring quality, analyzing resource utilization, supporting marketing plans, predicting machine maintenance needs, predicting profitability, and adjusting product mix. The data provides information needed by designers, process managers, quality control management, as well as business intelligence/analysis tools used by business managers for compliance reporting and corporate leadership decision support. The Figure 10[12] depicts the typical set of Smart Manufacturing processes.

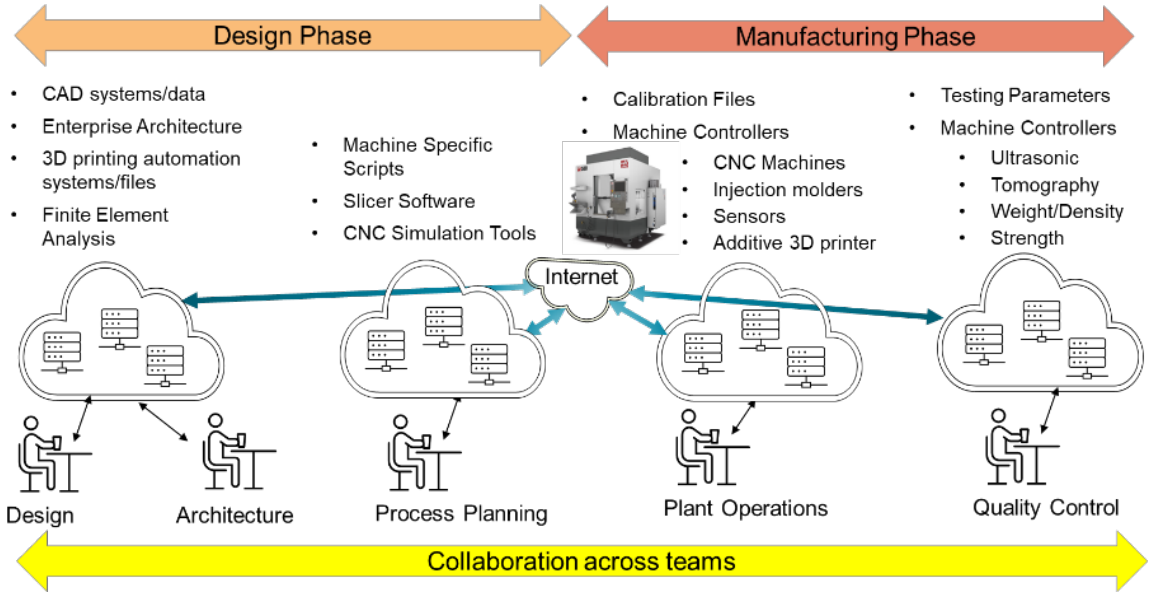

Figure 10 Smart Manufacturing Processes[14]

2.7.4 Findings

Smart Manufacturing requires manufacturers to add technologies that generate new data sets, include flexible manufacturing capabilities, increase data sharing among business units, and increase data analytics. These changes increase the cybersecurity attack surface, and change the cybersecurity needs

[14] IEEE. (2021). Proceedings of the IEEE: Survey of Cybersecurity of Digital Manufacturing. Mahesh et al. Retrieved from https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9247

of manufacturers. CISA should maintain situational awareness of these changes and develop services and guidance to support manufacturers as they transition into Smart Manufacturing.

2.8 Privacy Enhancing Technologies

2.8.1 Description

PET covers a broad range of capabilities that permit the exploitation of data while preserving anonymity and confidentiality. Some of the most important are summarized in Table 7.

Table 7: PET Summary

PET	Description	Notes
Homomorphic Encryption (HE)	HE allows encrypted data to be processed while it remains encrypted, preserving confidentiality in cloud computing and other vulnerable environments. Current implementations of HE imposes a high computation burden, but research on hardware-based solutions show promise.	High computational cost hinders wide adoption. Research initiatives (including DARPA) in progress.
Secure Multiparty Computation (SMPC)	SMPC allows multiple parties to cooperate on joint computations without sharing the contributed data to others. Several implementations have proven worthwhile, but their development can be labor intensive.	Optimized special purpose implementations exist in social sciences, finance, and other fields. Current research focuses on developing turn-key solutions.
Federated Learning (FL)	FL is a machine learning development technique that trains a model on data that is distributed across multiple devices, without the need to share the data itself. Typically, a trusted central server aggregates the inputs and updates the global model, but recent work on Peer-to-Peer FL is enabling the elimination of that requirement	Like MPC, FL has multiple implementations. Google's text-prediction capability is one example. Also, like MPC, FL requires specialized expertise to develop and deploy, Issues around poisoning the model and regulatory uncertainties are also concerns.
Differential Privacy (DP)	DP is a data aggregation method that adds randomized "noise" to the data, allowing for a quantification of privacy risk. It is currently employed by industry and government, including the U.S. Census Bureau. There are no clear best practices or standards on the proper tradeoff between accuracy and privacy.	The census bureau has used DP to mitigate the use of its data to infer Personally Identifiable Information (PII). Current research focuses on improved efficiency, privacy guarantees and identifying and mitigating attacks.
Trusted Execution Environment (TEE)	TEEs are (typically) cloud-based enclaves that shield data and processing from unauthorized users, including cloud administrators	A TEE prohibits execution of any code outside that environment. The confidential computing threat model aims at removing or reducing the ability for a

PET	Description	Notes
		Cloud Service Provider (CSP) and other actors in the tenant's domain to access code and data while being executed.

2.8.2 Importance to CISA

PETs are already in use in a variety of contexts, while continuing research strives to improve the efficiency and usability of various techniques. Any context in which useful research or collaboration is impeded by the need to protect personal information or confidentiality is a potential beneficiary of PET. For example:

Healthcare: PETs can be used to protect the privacy of patients' medical records while still allowing researchers and clinicians to collaborate. For example, the company Care.Trials has launched a block-chain based network for clinical trials, which relies on zero-knowledge proofs, one of the PET building blocks depicted in Figure 11. [15]

Cyber Threat Intelligence: Information of interest to Security Operation Centers (SOCs) and other defenders can be shared without identifying victims or other irrelevant information. CISA oversees a capability called Automated Information Sharing (AIS) that provides real-time exchange of machine-readable threat indicators among participants. The anonymity of the contributors is promised by CISA. Whether the anonymity provided by PET would enhance participation is an open question.

Census Data: has used Differential Privacy to balance the need to collect and report data with the statutory obligation to protect respondent confidentiality.

2.8.3 Details

The term PET covers a range of capabilities that are suitable for differing use cases and depend on different technologies. Of the five technologies described in Table 7, two stand out as having greater maturity and general usefulness: SMPC and TEE.

Secure Multi Party Computation (SMPC)

A working SMPC application consists of multiple interrelating modules built of primitives that execute in the environments of the participants, which adds computation and communication overhead. A continuing goal of SMPC research is to minimize overhead, and to develop tools that can streamline the compilation of primitives into functional applications.

Current research in SMPC is focused on turning the established theoretical models into practical applications. The building blocks noted in Figure 11 form the basis of implemented primitives that specify the sharing or exchange of information under a specified security model. Research on cloud

[15] Adams, Josh; August 2023, Clinical Trials Increasingly Adopt Blockchain and Zero-Knowledge Proofs; https://beincrypto.com/clinical-trials-blockchain-zero-knowledge-proofs

based SMPC makes use of cloud resources to reduce the overhead associated with the inter-party communication and computation.

Based on perceived threats, implementation must also consider the potential for malicious behavior among participants – either intentional or due to compromise by a malicious actor. Current efforts in this are investigating the tradeoffs between robust security, degrees of trust, and efficiency.

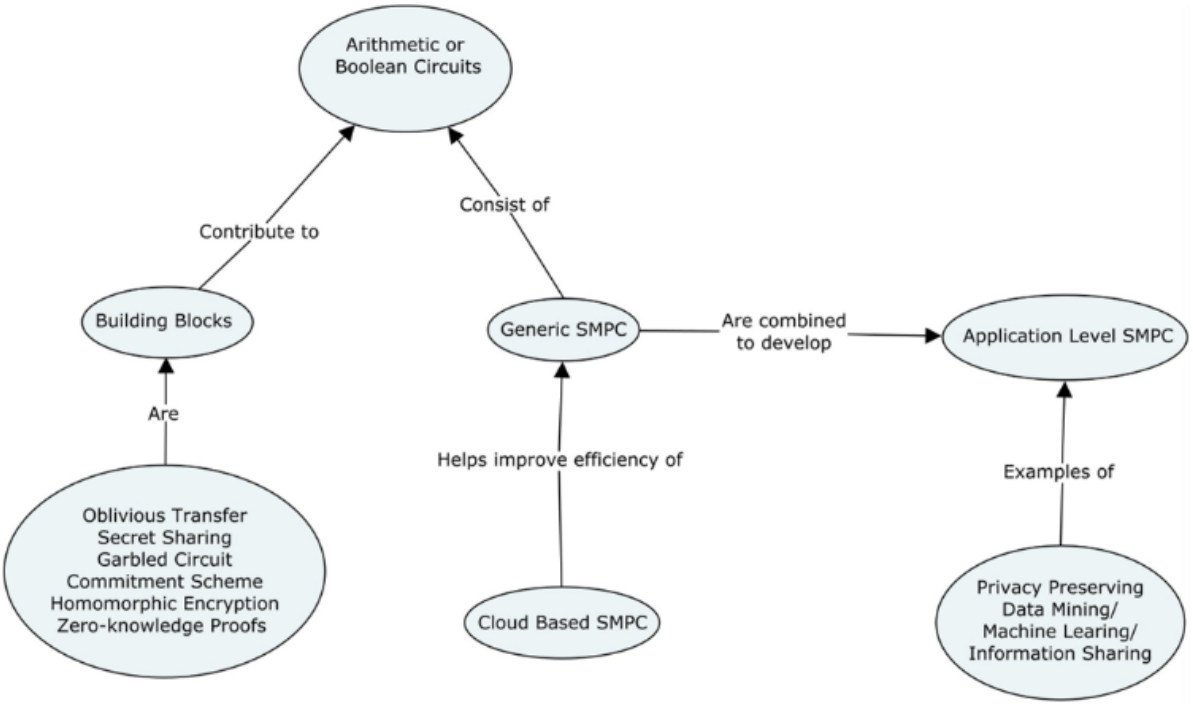

Figure 11: SMPC Components

Trusted Execution Environments (TEE)

TEE technologies provide confidential computing capabilities in the cloud. Vendor implementations differ, but they are typically achieved through a combination of hardware and software mechanisms. For example, Azure Software Guard Extension (SGX) enclaves rely on Intel hardware to isolate the host environment from the enclave. Amazon Web Service (AWS) Nitro Enclaves use cryptographic techniques and a reduced attack surface. (AWS enclaves cannot be accessed by the processes, applications, or users (root or admin) of the parent instance.)

Regardless of implementation, the use cases associated with this technology are comparable to those related to SMPC. As shown in Figure 12,[16] an enclave acts as a trusted third party accepting confidential data, processing a function based on those inputs, and publishing the results without disclosing the individual contributions or even the identities of the participants.

[16] Use AWS Nitro Enclaves to perform computation of multiple sensitive datasets, Sheila Busser, June 2022 retrieved from https://aws.amazon.com/blogs/compute/leveraging-aws-nitro-enclaves-to-perform-computation-of-multiple-sensitive-datasets/

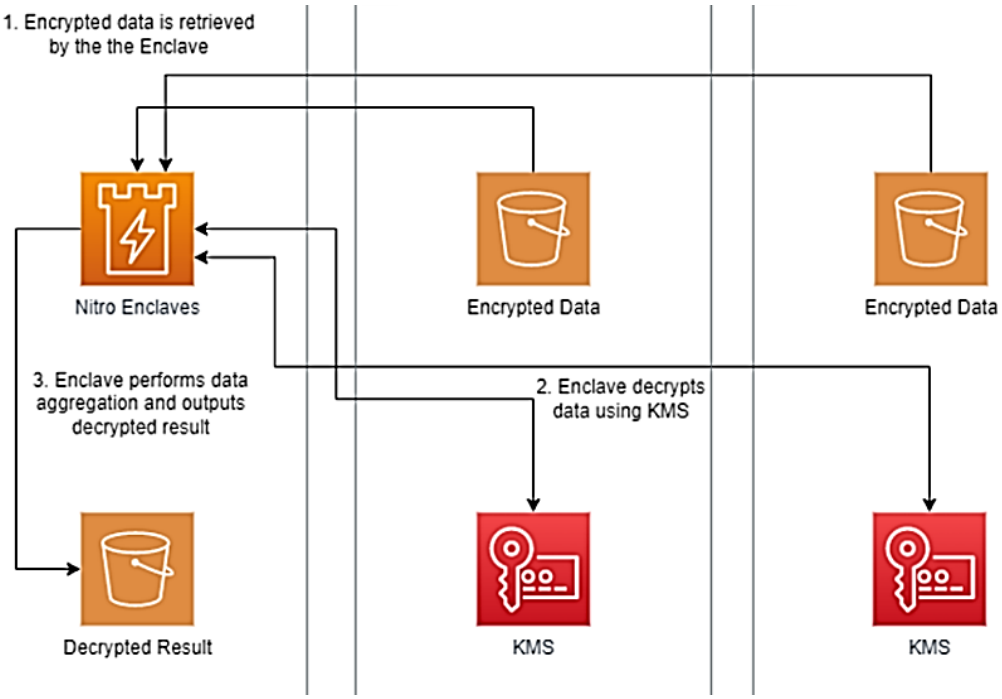

Figure 12: Enclave Based Bidding

Federated Learning (FL)

Federated learning allows multiple collaborators to train ML models without revealing the training data sets to other participants. Any initiative that includes training with proprietary or private data is a potential beneficiary of this capability. For example, federated learning can be used to:

- Train medical imaging models that are used to diagnose diseases.
- Train fraud detection models used by financial institutions to identify suspicious transactions.
- Improve manufacturing by training models that optimize production process, detect defects, and predict equipment failure.

Since reliance on a trusted central server to aggregate and redistribute model updates may be a stumbling block for some potential participants, research on peer-to-peer approaches is noteworthy. In the peer-to-peer approach, the edge devices communicate with each other directly to exchange their local model updates. The devices can either use a fully connected topology or a decentralized topology to exchange model updates. In a fully connected topology, each device communicates with all other devices to exchange their local model updates. In a decentralized topology, each device communicates with a subset of other devices to exchange their local model updates.

Other areas of research include various algorithm optimizations and reduction of communications overhead.

Zero Knowledge Proof

In a zero-knowledge proof (ZPK), one party (the prover) proves to another party (the verifier) that a given statement is true, without conveying anything else.

ZPKs have been deployed in various applications, mostly in the context of cryptocurrencies and blockchain technology. They are also identified as one of the "building blocks" of SMPC's in Figure 11.

However some more unusual use cases have been contemplated.[17] There are potential solutions in those circumstances that entail what is sometimes referred to as "over disclosure". Everyone is acquainted with the requirement to prove one's legal right to purchase liquor, typically by presentation of an identification card that includes the presenter's birthday. This exchange is so common that one forgets how much unneeded information it discloses: at the very least the presenter's actual date of birth and exact age.

ZKPs can protect individuals' digital information by allowing gatekeepers to verify access eligibility without disclosing personal information, reducing the risks associated with identity theft and data breaches.

A more abstruse use case includes the use of ZKP for algorithm verification. Certain government criteria, like those that trigger an IRS audit, are encapsulated in software. While there is a need to maintain the secrecy of the algorithm to avoid aiding its circumvention by would be tax evaders, there is also a legitimate need to assure the public that these algorithms behave predicably based on their inputs and do not unfairly target individuals in an insidious fashion.

While ZKP's have not been deployed for such purposes, they may help solve "verification dilemmas" where individuals and companies must reveal sensitive information for tasks like accessing websites or loans and deal negotiations, potentially compromising privacy, security, and competitiveness.

2.8.4 Findings

PETs are an emerging, but largely untapped technology. While research questions exploring how to minimize overhead and how to quantify levels of trust continue, many deployed applications are achieving beneficial results on a continuing basis. High computational cost will hinder wide adoption of HE for the foreseeable future, but significant research initiatives (including Defense Advanced Research Projects Agency (DARPA)) are in progress. Cloud-based execution environments have the potential to provide more turn-key and scalable offerings that may hasten adoption.

Table 7 summarizes the security risks or benefits of the various privacy enhancing technologies.

2.9 Anonymous Information Sharing

2.9.1 Description

Anonymous Broadcast (AB) is a PET that allows a sender to broadcast a message to a group of recipients without revealing the identity of the sender or the recipients. Cryptographic methods eliminate the requirement for a trusted third party. Some implementations of this idea employ obfuscation techniques, like The Onion Router (TOR) network, but a technique sometimes called "threshold broadcast" adds significant privacy assurance using encryption techniques. AB encrypts a message in

[17] Kenneth A. Bamberger et al. (2021, February 18). Berkley Technology Law Journal. Verification Dilemmas in Law and the Promise of Zero-Knowledge Proofs. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3781082

such a way that the original plain text can be recovered only if at least a specified number (threshold) of the receivers cooperate.

2.9.2 Importance to CISA

This capability has a variety of potential scenarios. In fact, any context in which the disclosure of information considered valuable to the public welfare is impeded by fear of negative consequences accruing to a would-be discloser represents a potential use case. This particular inquiry was prompted by consideration of the benefits related to the sharing of cyber related information:

- Incident reporting
- Vulnerability disclosure
- Threat intelligence sharing

The widespread dissemination of this sort of information would benefit SOCs and malware analysts, as well as a variety of investigators and law enforcement personnel.

Some impediments to information sharing are not addressed by AB. For example, some organizations may be dissuaded by the effort required to establish the capability, since it includes expenses associated with infrastructure, training, policy, and legal concerns.

2.9.3 Details

The underlying concepts for anonymous broadcasts were developed in 1988.[18] Since then, researchers have endeavored to improve on the idea by enhancing efficiencies and minimizing security exposures. Limiting factors are the number of participants and file sizes. One proposed implementation uses "dummy" broadcasts of null files to obscure the source of the sender. These broadcasts are verified and aggregated by some number of participating servers, two are shown in Figure 13[19], and then shared with the participants.

[18] David Chaum. (1988). The Dining Cryptographer's Problem: Unconditional Sender and Recipient Untraceability. Journal of Cryptology, 1(1):65–75, 1988.

[19] Z. Newman, S. Servan-Schreiber, and S. Devadas. (2022). Spectrum: High-bandwidth Anonymous Broadcast, usenix, 2022. Retrieved from https://www.usenix.org/conference/nsdi22/presentation/newman

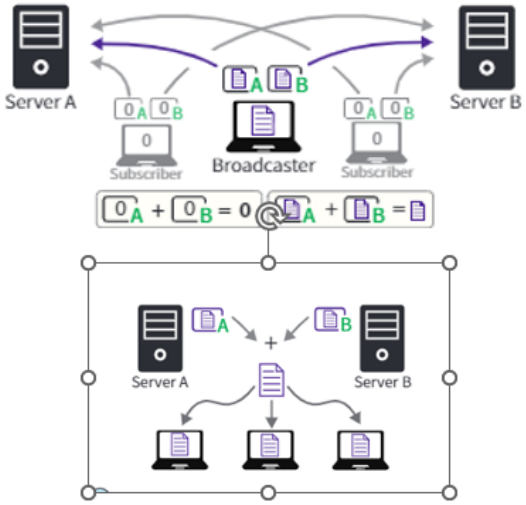

- **Subscribers send shares of null files as cover**
- **Broadcasters send secret shares of real files**
- **Servers verify broadcast permission and aggregate valid shares**
- **Servers reveal their aggregated secret shares to recover the broadcast**

No adversary observing the network and or even controlling a subset of servers and clients can distinguish between an honest subscriber and an honest broadcaster.

Figure 13: Anonymous Broadcast

2.9.4 Findings

It is well established that the widespread and timely dissemination of threat intelligence and vulnerability information can help protect information systems. The premise of AB is that such useful distributions will be encouraged if senders can broadcast a message without revealing their identity; thereby avoiding exposure to any liability or other damages.

However, it is not certain that anonymity concerns are the primary obstacle inhibiting such information sharing. If a trusted third party that collects, collates, and disseminates information is trusted to conceal the identities of contributors, then the implementation of a cryptographically protected infrastructure may be superfluous. (A trusted third party is the model used by CISA's AIS program).

To be genuinely useful the information must be sufficiently specific to motivate specific mitigations from defenders. For example, the CPS employed by elements of CI are often relatively specialized, and subject to attack by highly motivated and skilled adversaries. In such a context, threat intelligence that is both detailed and trustworthy could have high value. CISA should seek to understand the barriers that hinder adoption of this capability (see Section 2.9.2) and initiate appropriate action.

2.10 Satellite Communications Technology (SATCOM) Cybersecurity

2.10.1 Description

Satellite communications systems were first deployed in the late 1950s. Securing SATCOM operations, control, and communications systems is important to ensure operations and communications services are uninterrupted.

2.10.2 Importance to CISA

SATCOM falls within the Communications Critical Infrastructure Sector. The National Risk Management Center is leading the CI Partner Advisory Council (CIPAC) on Space Systems Critical Infrastructure. CIPAC activities include outreach to SATCOM operators to ensure they are aware of CISA's available resources and tools.

2.10.3 Details

Three types of satellite constellations support communications: geosynchronous, medium earth orbit (MEO), and low earth orbit (LEO). Two communication architectures are used to support communications among end users, described in Figure 14, Traditional, and Figure 15, Mesh.

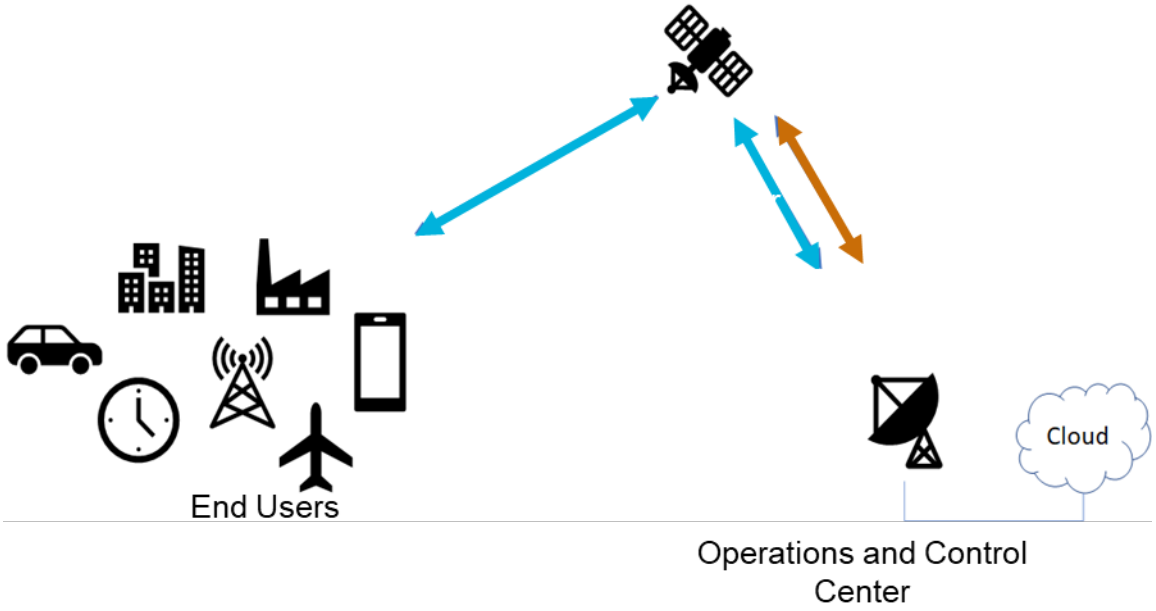

Figure 14: Traditional SATCOM Architecture

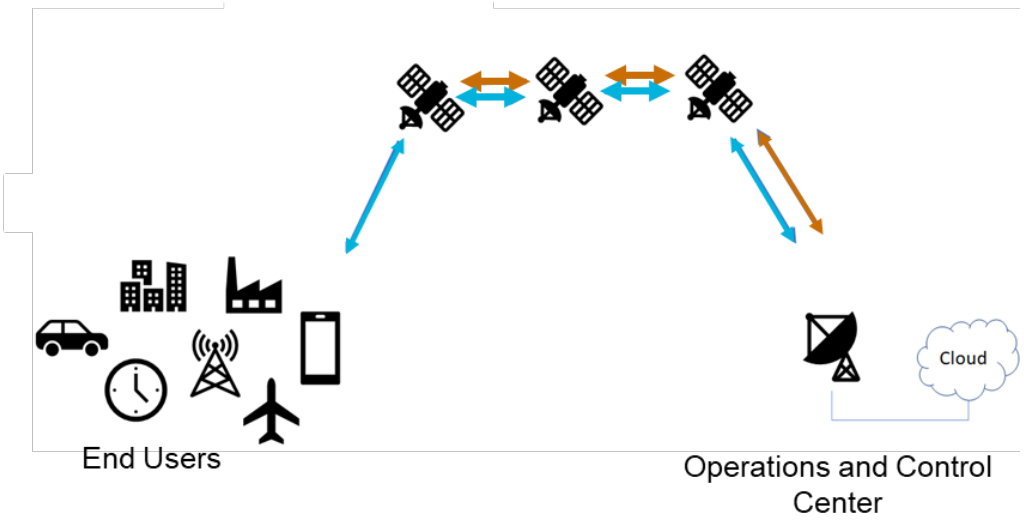

Figure 15: Mesh SATCOM Network Architecture

Commercial SATCOM system cybersecurity is not currently required by regulation. Recent evidence of SATCOM service disruption via cyber attack has raised the level of interest in improving the security of the SATCOM infrastructure. As satellite operations and control systems integrated IP-based communications are replacing non-routable point to point communications protocols, vulnerabilities

similar to IT systems need to be mitigated. The recent cyberattack on a satellite system serving Europe is an example of an adversary successfully utilizing a common IT vulnerability.

Regulations - Commercial and government SATCOM satellite, payload, space vehicle physical movement for launch, orbit insertion, and de-orbit/re-entry is regulated by the FAA. Cybersecurity in the commercial SATCOM industry is not regulated. SATCOM operators recognize the need for security for the space vehicle tracking, telemetry, and control (TT&C) links used to control the satellite payloads and position control. TT&C security controls are not publicly available for cybersecurity and intellectual property protection reasons.

Government resources - In 2022 NASA published a Security Threats Against Space Missions report.[20] The report includes information regarding threats to space missions and includes counter threat options.

In 2018 Department of Defense, National Air and Space Intelligence Center published "Competing in Space."[21] This report describes the various applications and use cases for space-based assets. It also includes descriptions of the challenges created by foreign space assets and describes current and evolving threats to space-based systems, including SATCOM.

In 2019 the Space Information Sharing and Analysis Center (Space ISAC) was formed to collect and share all-threats security information applicable to both public and private space systems. The ISAC is focused on three threat areas: supply chain, business systems, and missions.[22]

2.10.4 Findings

SATCOM operators need assistance to improve the cybersecurity of the systems that operate and control the services they offer, including TT&C for their space assets. CISA is currently working with SATCOM operators and should continue to work with SATCOM operators to develop recommendations and guidance to improve their cybersecurity. CISA currently provides cybersecurity guidance to satellite network providers through advisories such as Strengthening Cybersecurity of SATCOM Network Providers and Customers, Alert CodeAA22-076A.[23]

2.11 ICS Virtualization

2.11.1 Description

Virtualization of Industrial Control Systems (ICS) is the practice of moving from dedicated Operational Technology (OT) computing devices to virtual machines (VM) running in a shared hosting environment. Virtualization has been ongoing in the IT world for decades and has become a preferred approach for implementing systems. Now, it is being used in the OT world driven primarily by cost and agility factors that have been demonstrated in the IT world. This section investigates new risks that may be introduced

---

[20] NASA. (2022). Security Threats Against Space Missions. Retrieved from https://public.ccsds.org/Pubs/350x1g3.pdf

[21] DoD. (2018). Competing in Space. Retrieved from https://media.defense.gov/2019/Jan/16/2002080386/-1/-1/1/190115-F-NV711-0002.PDF

[22] Space ISAC. (2020). Brochure. Retrieved from https://s-isac.org/wp-content/uploads/2020/08/SISAC_8x11_Email-rev1.5.pdf

[23] CISA. (2022). Strengthening Cybersecurity of SATCOM Network Providers and Customers: Alert CodeAA22-076A https://www.cisa.gov/news-events/cybersecurity-advisories/aa22-076a

by this technology when implemented in OT systems that are often critical to the reliability and resiliency of CI.

2.11.2 Importance to CISA

CISA develops guidance and standards for CI which often depends on ICS. Virtualization significantly changes the architecture of ICS systems, which introduces a new set of security vulnerabilities and reliability concerns. Therefore, CISA needs to understand the trend towards ICS virtualization so that adjustments can be made to existing ICS security recommendations, guidance, and standards.

2.11.3 Assessment

ICS virtualization is the replacement of dedicated ICS computing devices (e.g., Programmable Logic Controllers (PLCs), Supervisory Control and Data Acquisition (SCADA), Human Machine Interface (HMI) systems, and Historians) with shared computing services running on-premises or in a cloud. Virtualization is used extensively in the IT world and has advantages and disadvantages included in Table 8.

Table 8: ICS Virtualization Advantages and Disadvantages

Advantages	Disadvantages
• More efficient use of hardware as a single server can support numerous VMs and resources can be dynamically allocated as needed when workloads vary. • Software and data for virtualized devices is kept at, and managed from, central locations. • It is easier to access, monitor, and control configurations of all instances using hypervisor management consoles. • Easier to update and restore, since workloads are containerized and can be quickly restored from backup or base images.	• Increased dependence on the speed and availability of the network connection needed to access the VM and move data in and out. • Not all systems can be readily converted to run in a virtual machine if they use proprietary or unusual operating systems not supported by typical VM hypervisors. • Virtual machine images are suspectable to same forms of attack as any other system. There is no inherent extra security provided by typical hypervisors.

The following diagram shows layers of the Purdue reference architecture[24] that is often used to characterize ICS systems and identifies which layers are, and are not, candidates for virtualization. The lower layers (levels 0 and 1) or those that deal with real time physical process control that are critical to safety are typically not considered as viable to virtualization or cloud hosting given strict timing or safety concerns. These systems often have dedicated real-time networks and are isolated from remote access and higher-level computers through firewalls and data diodes.

[24] D. Garton. (2019, November 12). Purdue Model Framework for Industrial Control Systems & Cybersecurity Segmentation. Retrieved from https://www.energy.gov/sites/default/files/2022-10/Infra_Topic_Paper_4-14_FINAL.pdf.

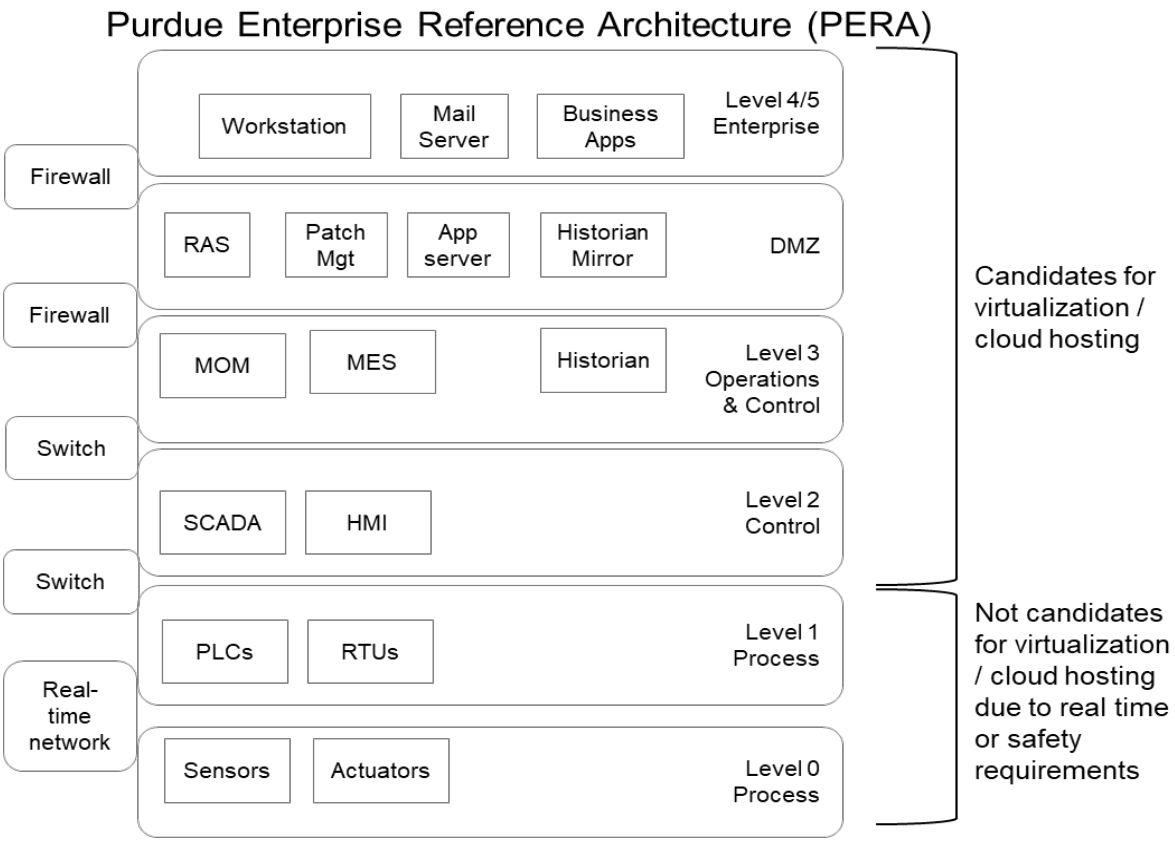

Figure 16: Purdue Enterprise Reference Model (PERA)

A virtualized ICS architecture is depicted in Figure 17.

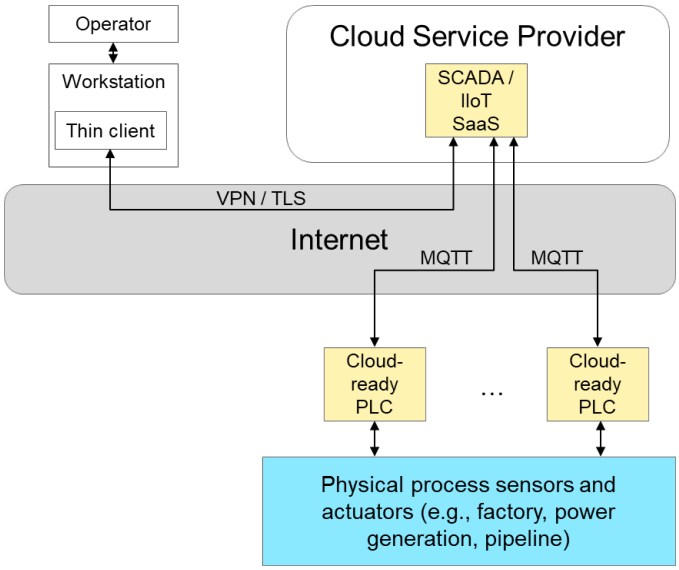

Figure 17: Virtualized ICS Architecture

Note that the internet provides connections between operators, CSPs, and the PLCs used to monitor and control physical systems. The communication ports at the operator workstations, the services running in

a CSP, and the cloud-ready PLCs are potential points where adversaries can attack. One could argue that this architecture is far more vulnerable and difficult to defend than a traditional on-premises ICS architecture where critical systems are all protected by multiple layers of firewalls. However, these traditional architectures have also been compromised especially via remote operator workstations and misconfigurations or omissions of basic security controls due to lack of experience or resources within the plants. Cloud-ready PLCs and CSP-based Software as a Service (SaaS) typically are designed to be secure by default, so they may actually have a security advantage over on-premises one-off designs.

2.11.4 Findings

The following bullets summarize the key findings of research into the trend of ICS virtualization with respect to CISA's concerns:

- ICS virtualization is a trend driven primarily by potential cost and agility benefits.
    - More efficient use of computing resources (less cost)
    - Reduced overall security risks for collections of servers and workstations
    - Greater scalability and flexibility
    - Ease of update
    - Rapid restoration after disaster or attack
- There are many approaches to virtualization, which means every implementation can be unique and highly customized. This uniqueness makes generating policy/regulation based on standard implementations more challenging.
- Virtualized ICS products and services are readily available and have been implemented in operational systems – though examples of implementations for complex ICS have not been identified.
- A cloud-based virtualized ICS approach introduces risks that need to be carefully managed to maintain security.
    - Dependency on external internet connection availability and security
    - Dependency on external cloud service availability and security
    - Greater number of devices exposed to the internet presenting a larger attack surface to adversaries.

2.12 ZTA Technology Status

2.12.1 Description

The Executive Order on Improving the Nation's Cybersecurity[31]31 requires all federal agencies to implement ZTA, which has increased interest in the availability of products and the status of implementation capabilities for ZTA.

The key concept of ZTA is that instead of using perimeter defenses to protect a flat enterprise network, every single access request to sensitive data is checked and connected only to those resources that are permitted by centrally controlled access policies. ZTA is implemented across eight IT functional areas: User or Identity; Device; Network/Environment; Application and Workload; Data; Visibility and Analytics; Automation and Orchestration; and Governance. A description for each functional area is presented in Table 9.

Table 9: ZTA Requirements in Eight IT Functional Areas

IT Functional Area	Description
User or Identity	ZTA requires the ability to continuously authenticate users using multi-factor (MFA), authorize access using role-based access control (RBAC) and attribute-based access control (ABAC) policies, and continuously monitor them while connected to the network.
Device	ZTA requires the identification, authentication, authorization, inventory, isolation, and control of all devices in the ZTA.
Network / Environment	ZTA requires the network to be micro-segmented so that connections between devices and workloads are individually provisioned per each request, continuously monitored, and terminated as soon as they are no longer needed or if traffic deviates from that expected for the type of access requested.
Application and Workload	ZTA requires securing and properly managing the application layer, computing containers, and virtual machines in both on-premises and cloud computing environments.
Data	ZTA requires all data to be inventoried, categorized, classified, and tagged for the purpose of access control, and mechanisms be put in place to prevent and detect exfiltration.
Visibility and Analytics	ZTA requires visibility across all elements to continuously monitor performance and detect anomalous behavior to make dynamic changes.
Automation and Orchestration	ZTA requires the automation of processes to rapidly implement policies across the enterprise.
Governance	ZTA requires governance processes to create and manage policies for device compliance, application access, network communications, and visibility and alerting.

2.12.2 Importance to CISA

The rapid growth and adoption of the Internet of Things (IoT), edge computing, and remote and hybrid work solutions has challenged the ability of traditional perimeter-based security architectures to protect

enterprise assets and CI. Recent cyber incidents[25,26] have highlighted the broad challenges of ensuring effective cybersecurity across the federal government, necessitating a shift to a data-centric approach to enterprise architecture, with fine-grained security controls and access control policies across users, systems, data, and assets that change over time.

CISA has a key role in developing and articulating recommendations and guidance regarding ZTA implementation. CSA has most recently released an update to the CISA ZTA Maturity Model[27] that provides a framework for agencies to use in assessing their level of maturity within each functional area of ZTA. As the Federal government and critical infrastructure seek to implement optimal zero trust implementations, it is important for CISA to understand capability gaps that may hinder organizations from achieving their needed ZTA capabilities and identifying, funding, and/or monitoring R&D and application of new technologies to help improve ZTA implementation.

2.12.3 Details

For each of the capabilities within the CISA Zero Trust Maturity Mode,[27] a survey of open-source literature was conducted to examine the current state of industry, academia, and Federal, State, Local, Tribal, and Territorial (FSLTT) ZTA enabling capabilities. The purpose of this analysis was to identify potential technical capability gaps that may require R&D and/or standardization to address the identified gaps.

ZTA is an architectural strategy that must be tailored to meet individual enterprise needs; therefore, the following facets have not been considered, as these facets should largely be unique to a particular enterprise and not broadly applicable to any or all enterprise(s):

- Notional constructs or use cases,
- Current strategies, implementation plans, or current and planned policy,
- Capability gaps resulting from lack of funding to acquire commercially available products, or
- Capability gaps resulting from lack of skilled staff.

ZTA capabilities were assumed to be fully implementable if common, stable, and commercially available technologies exist. Technical gaps identifying either a need for R&D and/or standardization are presented in Tables 11 – 15, and the capability status key is presented in Table 10.

Table 10: ZTA Capability Status Key

Technology Solutions Available	Standardization Needed	Limited Technology Solutions Available	No Technology Solutions Available

[25] CISA. (2020, December 13). Emergency Directive 21-01- Mitigate SolarWinds Orion Code Compromise. Retrieved from https://www.cisa.gov/emergency-directive-21-01.

[26] CISA. (2021, March 3). Emergency Directive 21-02 - Mitigate Microsoft Exchange On-Premises Product Vulnerabilities. Retrieved from https://www.cisa.gov/emergency-directive-21-02.

[27] CISA. (2023, April). Zero Trust Maturity Model. Retrieved from https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf.

Table 11: ZTA Identity Pillar Status Summary

Technology Pillar	Function	Method
Identity	Authentication	Password
		Biometric
		2FA
		Context-Aware User Authentication
		Continuous Authentication
		Device Authentication
	Access Control	Identity-Based Access Control (IBAC)
		Role-Based Access Control (RBAC)
		Attribute-Based Access Control (ABAC)
		Risk-Based Access Control (RbAC)
		Capability-Based Access Control (CBAC)
	Identity Stores	Self-Managed, On-prem
		Self-managed and Hosted
		Full integration across partners and environments
	Risk Assessments for Identity Risk	Manual methods and static rules
		Limited automation and dynamic rules
		Real-time continuous analysis and dynamic rules
	Access Management	Permanent access w/ periodic review
		Expiring access with automated review
		Need-based and Session-based
		Automated authorization as needed with minimal required access granted to individual actions and resource needs.

Table 12: ZTA Device Pillar Status Summary

Pillar	Function	Method
Device	Device State Monitoring	Network-based Indicators
		Host-Based Indicators*
	Compliance & Policy Enforcement	Configuration Management
		Device Threat Protection*
	Visibility and Analytics	Real-time Monitoring and Centralized Reporting (e.g. SIEM\SOAR)*

Table 13: ZTA Network Pillar Status Summary

Pillar	Function	Method
Network	Network Segmentation	Micro-segmentation
		Transport-Level Access Control
		Label-based Access Control
		DPI-Based Access Control
		API-Aware Access Control
	Software Defined Networking (SDN)	Software Defined Perimeters (SDP)
	Network Traffic Management	Manually Implemented Static Rules and Configurations
		Application Profiles with Distinct Traffic Management
		Dynamic Network Rules and Configurations

Table 14: ZTA Applications and Workload Pillar

Pillar	Function	Method
Applications and Workloads	Software Defined Compute (SDC)	Cloud-ready Virtualized or Containerized Applications and Resources
	DevSecOps	Integrating Security from Design to Delivery
	Software and Supply Chain	SBOM
	Application Delivery	Standardized, granular access to resources (not network) with continuous assessment and logging
	Application Security Testing	Static
		Dynamic
		Integrated throughout SDC
	Application Accessibility	Private Network
		Secure Public Network
		Open Public Network

Table 15: ZTA Data Pillar Status Suminmary

Pillar	Function	Method
Data	Data Tagging	Metadata
		Other Data Encoding i.e. security markings, access rights and handling, enterprise data header
	Encryption	Secure Multi-Party (SMPC)
		Lightweight Encryption
		Lightweight Mutual Authorization
		Homomorphic Encryption
	Categorization	Ad Hoc\Manual
		Automated
	Access	Data Loss Prevention (DLP)
		Software Defined Storage (SDS)
		Data Rights Management (DRM)

2.12.4 Findings

For an optimal level of ZTA implementation as defined in the CISA ZTA Maturity Model, ZTA requires new approaches and algorithms to process additional factors and policies.

There is no one size fits all approach to ZTA; therefore, it is extremely challenging to assess and articulate the present state of ZTA capabilities due to the wide choice of alternatives enterprises may select. After an examination of the current state of ZTA, the following gaps were identified that when addressed, will improve the maturity of ZTA implementation and making an Optimal state of maturity feasible:

- No complete ZTA solution is currently available; achieving ZTA objectives requires standardization that allows for integration of commercially available heterogenous technologies.
- In order to achieve Context Aware User Authentication, mechanisms must be enhanced or developed that can leverage rich contextual information but are also widely useable across all types of devices. Most context-aware authentication mechanisms rely on location information that is concatenated with some other information such as device time or proximity. However, some of the sensors that are deployed for capturing the contextual information (e.g., accelerometer or gyroscope) may not be available on some devices, and thus may not be usable across all types of devices. An open research direction would be to identify suitable alternatives to achieving context aware user authentication that in addition to location, or in lieu of location data, can be utilized with a series of contextual questions that the user must answer to calculate a risk score and authenticates the user.
- In addition to entry point authentication, continuous user authentication is an active area of research for academia and industry. Most continuous authentication mechanisms rely on user's behavioral biometric features associated with typing, tapping, and gait patterns. The problem with this approach is that they rely on sensors embedded in the devices. For continuous authentication across numerous devices (e.g., mobile-phone, laptops, tablets), the reliance upon

sensors is an issue, as not all the devices possess resembling embedded sensors. Besides this, behavioral biometrics are also dependent upon the situation in which they are captured. Development of a mechanism that can work across a wide variety of devices and situations is an open research question that will enable Continuous Authentication.

- Development and standardization are needed to define a viable device identity schema that supports the following properties: unique device identification, tamper resistance and un-clonability, adaptive authentication and access control, end-to-end encryption, and scalability.
- Access control models (e.g., capability-based access control) rely on public-key cryptosystems and digital signatures, warranting changes to utilize a PQC system. NIST has initiated a process to solicit, evaluate, and standardize one or more quantum-resistant public-key cryptographic algorithms.[28]
- Effective data tagging remains elusive to many enterprises and is a necessary capability to control access at the data resource level. Industry has identified numerous best practices and commercially available tools, despite this, many enterprises have not implemented foundational data governance programs. Additionally, many of the enterprise data tagging solutions are costly introducing funding and staffing challenges.
    - Data tagging requires establishment of enterprise policy and execution of well-defined data strategy, as well as investment in skilled staff and commercially available tools.
    - Data tagging can support Identity, Credential, and Access Management (ICAM) and ABAC access control models for fine-grained resource control.
- The symmetric encryption and hash utilized by Lightweight Encryption in IoT and OT devices require larger key sizes and hash lengths to maintain the necessary levels of security.
    - These changes may be problematic for resource constrained devices with limited computational power and memory.
    - Lightweight encryption may not be viable for legacy hardware.
    - Research efforts are ongoing to design a lightweight cryptosystem using quantum permutation pads (QPP). It is not yet known if this approach will protect against known attack vectors or if it can be used with legacy hardware.
    - NIST has selected a group of cryptographic algorithms called Ascon that will be published as NIST's lightweight cryptography standard later in 2023.[29]

[28] NIST. (2023. October 03, 2023). Post-Quantum Cryptography. Retrieved from https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization

[29] NIST. (2023, February 7). NIST Selects 'Lightweight Cryptography Algorithms to Protect Small Devices. Retrieved from https://www.nist.gov/news-events/news/2023/02/nist-selects-lightweight-cryptography-algorithms-protect-small-devices#:~:text=%E2%80%9CPost%2Dquantum%20encryption%20is%20primarily,not%20include%20all%20of%20them.

- Several areas exist where ZTA would benefit from the development of industry standards (e.g., common ontology for ZTA access control attributes (e.g., user and device)) to enable:
    - Consistent access enforcement and promote interoperability of ZTA products.
    - Federation across organizations with varying policies.
    - Interoperability and integration to enable ZTA principles within OT and IoT systems.

2.13 AI for ZTA

2.13.1 Description

Zero Trust Architecture is often traced back to an article published by Forrester in 2010.[30] In 2021, an Executive Order on Improving the Nation's Cybersecurity[31] requires all federal agencies to implement ZTA, which has supported an increase in plans and product offerings in this area.

The key concept is that instead of using perimeter defenses to protect a flat enterprise network, every single access request to sensitive data is checked and connected only to those resources that are permitted by centrally controlled access policies. Table 9 describes what functions are needed to implement ZTA in eight IT functional areas.

2.13.2 Importance to CISA

CISA has a key role in developing and articulating policies about ZTA implementation. They have most recently released an update to the CISA ZTA Maturity model[32] that provides a framework for agencies to use in assessing their level of maturity within each functional area of ZTA. CISA is also involved in supporting the R&D and application of new technologies to help improve ZTA implementation.

2.13.3 Details

NIST has published a ZTA Reference Architecture[33] that has gained a lot of traction within federal agencies. Figure 18 shows the logical architecture for ZTA defined by NIST.

[30] Forrester. by John Kindervag. (2010, November 5). Build Security Into Your Network's DNA: The Zero Trust Network Architecture). Retrieved from https://www.actiac.org/system/files/Forrester_zero_trust_DNA.pdf

[31] White House. (2021, May12). Executive Order on Improving the Nation's Cybersecurity. Retrieved from https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/.

[32] CISA. (2023, April). Zero Trust Maturity Model. Retrieved from https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf.

[33] S. Rose, O. Borchert, S. Mitchell and S. Connelly. (2020, August). Zero Trust Architecture. Retrieved from https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf.

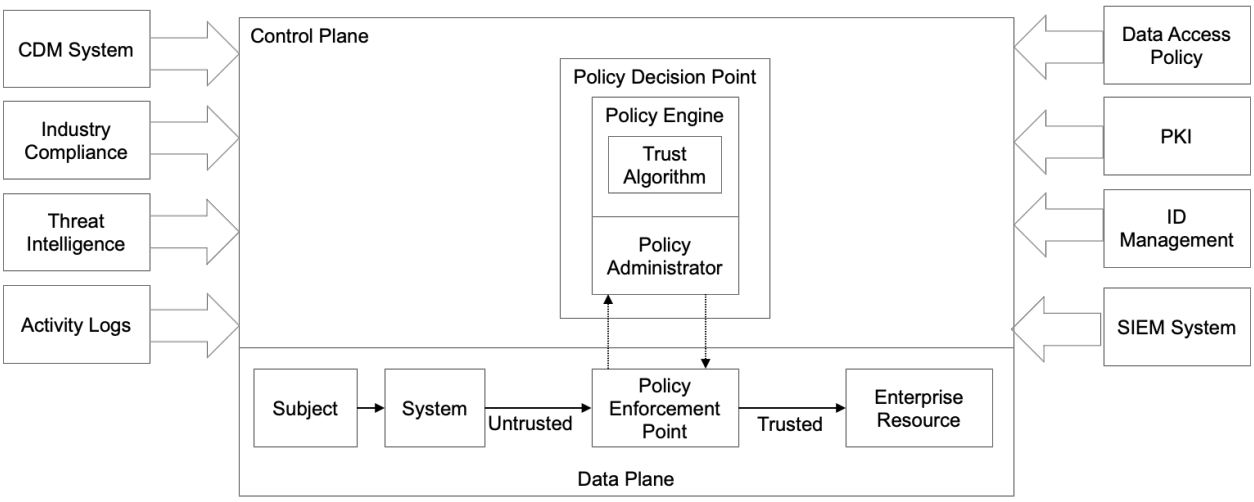

Figure 18: Zero Trust Logical Components

The key components of a ZTA are:

- **Policy Decision Point (PDP)** – Makes decisions to grant or deny access and to establish and terminate connections. The PDP has several subcomponents:
    - **Policy Engine (PE)** – A subcomponent of the PDP that makes decisions to grant or deny access to a resource.
    - **Trust Algorithm** – The analysis process used within the policy engine to grant or deny access. The trust algorithm uses inputs from many sources, such as access request, subject database, asset database, resource requirements, and threat intelligence. Trust algorithms may be implemented in a trust engine that performs the computations, and these trust algorithms are often implemented using the AI/ML techniques described later in this report.
    - **Policy Administrator (PA)** – A subcomponent of the PDP that issues commands for the establishment or shut down of connections to policy enforcement points (PEPs).
- **Policy Enforcement Point (PEP)** – Establishes, monitors, and terminates connections between subjects and resources. The PEP has several subcomponents that may be used independently or in combination:
    - **Agent** – A subcomponent of the PEP that resides on the subject's device.
    - **Gateway** – A subcomponent of the PEP, typically implemented adjacent to the resources that control access to resources.
    - **Gateway Portal** – A subcomponent of the PEP implemented as a central hub through which subjects connect to resources.
- **Resource** – Data item or application to which access must be managed and controlled. Resources exist within trusted enclaves of the enterprise otherwise known as trust zones. Each resource should have a unique set of classification tags that are used for access control. Therefore, it may be necessary to reorganize data items with different tags into separately addressable resources for the purpose of fine-grained access control. For example, a file store that contains files with different sets of classification tags may have to be separated into multiple stores to enable fine-grained access control.

- **Subject** – A user or system that requests access to a resource.

2.13.4 Findings

Realization of ZTA increases the size and complexity of the access control problem by orders of magnitude over traditional perimeter defense models for several reasons:

1. Authentication and authorization must be done for each fine-grained resource access.
2. Additional factors are included into the analysis of each access request such as time, location, and device status.
3. Each connection is continuous monitoring and controlled.

For an optimal level of implementation as defined in the CISA ZTA Maturity model, ZTA requires new approaches and algorithms to process additional factors and policies. AI/ML is a technology that may provide a means to implement these processes. After detailed examination of the current state of AI/ML and the need for ZTA functions, it was determined that AI/ML would be useful to improve maturity in at least the following three areas of ZTA:

- **Identity** – User behavior risk assessment – Developing profiles of user behavior and assigning a risk score that the PDP can use to accept or reject new access requests.
- **Network/Environment** – Threat protection and filtering – Developing profiles of network traffic flows and providing alerts to the PDP when traffic flows vary from expected behavior to use in filtering or isolating session traffic to contain or limit potential damage.
- **Data/Inventory Management** – Scanning, classifying, and tagging all data in the enterprise that needs fine-grained access control.

Figure 19 shows a high-level sequence of transactions that occur among the core components of a ZTA model and where AI/ML functionality could be embedded within the components.

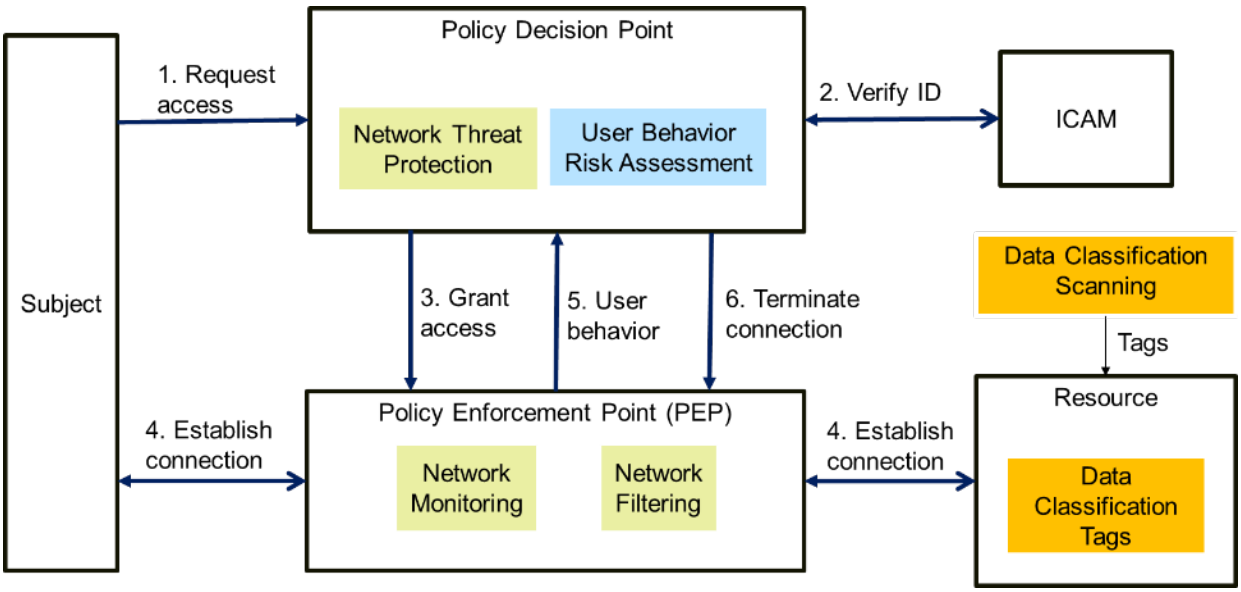

Figure 19: Zero Trust Logical Components with Supporting AI/ML functions

Organizations seeking to improve the maturity of their ZTA implementation might do well to consider integrating one of the many commercial products or services that use AI/ML technology to improve their performance in these areas.

2.14 CPS-Resiliency

2.14.1 Description

Recent high-profile attacks on CI, both domestically and abroad, has amplified the urgency to increase the resiliency of the cyber-physical controls that are integral to many industries and sectors. In response, the President's Council of Advisors on Science and Technology (PCAST) has started a Working Group (WG) on CPS Resilience to "advance existing ideas and efforts as well as to develop new approaches to this problem". They have identified six focus areas (below) and asked for feedback from experts in academia, industry, and government, including DHS. The focus areas of the PCAST Working Group for CPS resiliency are:

- Recovery and survivability in the face of attacks and events.
- Approaches to assure continuity of operations in degraded states.
- Mechanisms to measure and assess modularity and limitations of scope or costliness of failures.
- Incentives to balance efficiency which can reduce resilience vs. the investment needed to maintain sufficient resilience.
- Out-of-band or systems-independent means of assuring physical control in the event of digital failures.
- Methodologies and standards to encourage resilient systems design and adoption.

Federal agencies, including DHS, and industry groups have been developing CPS Resiliency guidelines, standards, and methodologies for years. These activities, however, are not well coordinated nor

orchestrated and they use a wide variety of interpretations of the term “resiliency.” Often the focus of the CPS resiliency activity is centered on availability, reliability, and/or safety vice resiliency. Also, lacking National CPS resiliency outcomes for CI, these CPS resiliency efforts focus on bottoms-up improvements to CPS components in the hopes that these targeted improvements will lead to overall CI resiliency with no goal or objective by which to measure progress.

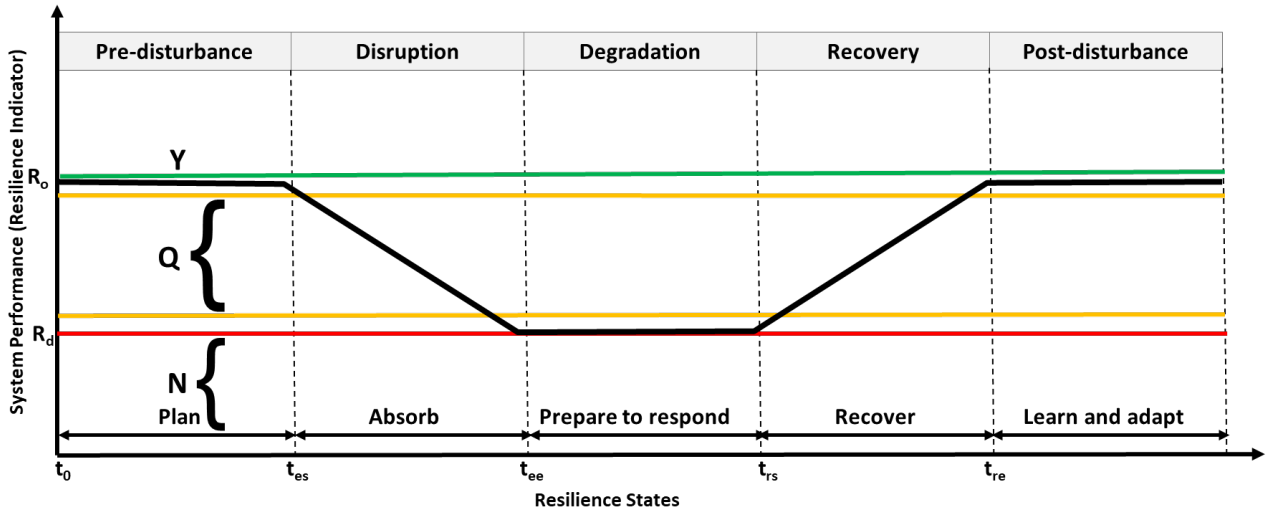

Department of Defense’s Mission Capability Readiness Definitions

Rating	Definition
Y	Unit can accomplish task to established standards and conditions
Q	Unit can accomplish all or most of the task to standard under most conditions. The specific standards, conditions, and shortfalls impacting the unit’s task must be clearly detailed in the mission essential assessment.
N	Unit is unable to accomplish the task to prescribed standard and condition at this time.

https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9255805

Figure 20: Mapping IEEE Resilience States to DoD Mission Readiness to Fully Describe "Resiliency"

MITRE has developed a framework that DHS CISA can utilize to drive CPS resiliency development to meet National CI interests. Incorporating the six PCAST WG CPS focus areas within the framework, National CI Outcomes/Objectives drive Federal policies and standards which in turn can be used to guide/prioritize/harmonize Federal investment in development of new CPS resiliency measures, methodologies, and tools.

2.14.2 Importance to CISA

Protecting critical national infrastructure, much of which relies upon CPS equipment, is core to the mission of DHS. DHS has numerous on-going activities around CI sector resilience and CPS resilience. A framework is necessary in order for CISA to drive a whole-of-government approach to strengthening our critical national infrastructure in accordance with national CI objectives.

2.14.3 Details

The six PCAST CPS WG Focus Areas are presented without any overarching guidance as to how they are to be used, nor to what end they are being implemented. Figure 21 shows a framework whereby National CI Resilience outcomes and objectives are realized via Federal levers (policy, standards, regulation) and enabled/matured by a Federal investment strategy to build prioritized CPS resiliency measures/methodologies/tools.

With the framework, DHS CISA guidance to the PCAST CPS Resiliency WG can focus on a narrative of building metrics and tools to gauge CPS Resilience against high-level Resilience Outcomes and Objectives, then targeting shortfalls within vulnerable "weakest link" CI through updated standards, creation of incentivizing policies, and enforcement mechanisms.

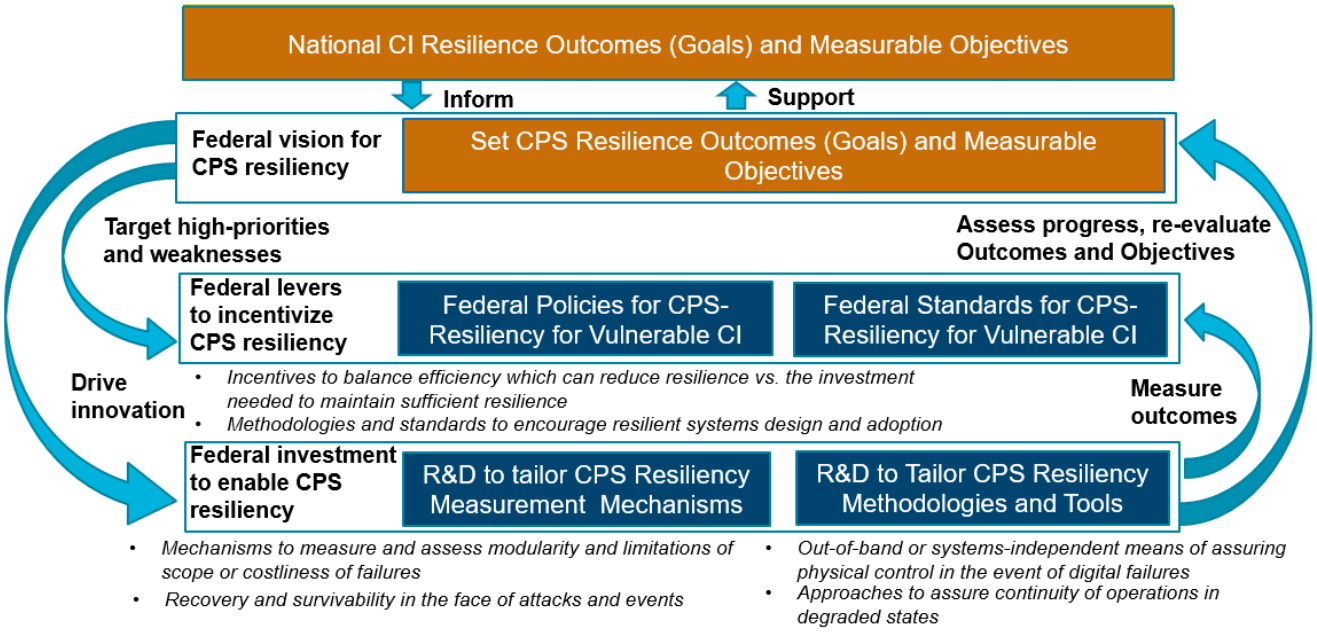

Figure 21: Framework to Drive National CPS Resiliency Activity that Incorporates PCAST CPS WG Goals

The following Table 16 outlines goals and approaches CISA can pursue to use the framework to build CPS-resiliency capacity within CIs in accordance with National objectives:

Table 16: Framework Goals and Recommended CISA Approaches

Framework Goal	Recommended CISA Approach
Create National Critical Infrastructure Resilience Objectives/Outcomes (e.g., target CI resilience levels such as allowable unplanned downtime, per capita downtime) to derive Industry/Component CPS Resiliency Objectives/Outcomes	Create National Standards in conjunction with Federal stakeholders by applying existing DHS CISA work on Critical Infrastructure with NIST approaches. Present Relevant CISA Activity:

Framework Goal	Recommended CISA Approach
	• DHS CISA Resilience Services Branch (RSB) & Infrastructure Resilience Planning Framework (IRPF) • DHS CISA Regional Resiliency Assessment Program (RRAP) NIST Approaches: • NIST Special Publication 1190GB-9: Summarizing Resilience Goals using Performance Goals Tables • NIST-IR 8406 methodology can be extended to create CPS resiliency outcomes (goals) for industry-specific components based upon National Critical Infrastructure priorities
Encourage Federal policies that change the cost/benefit trade-offs of implementing greater CPS resiliency within targeted CI capabilities	DHS CISA can promote the development of Federal policies for CPS Resiliency within the PCAST WG through collaboration with the following initiatives: • DHS CISA Resilient Investment Planning and Development Working Group (RIPDWG) • OSHA Process Safety Management Guidelines for Compliance • Department of Energy (DoE) and their National Cyber-Informed Engineering (CIE) strategy • Work directly with the associated SRMA to develop appropriate policies
Coordinate Federal CPS Resiliency standards development by coordinating across the many different Federal agency and industry groups currently providing CPS and CPS Resilience guidance	Encourage the PCAST WG to harmonize Federal standards for CPS Resilience by working with appropriate SRMAs across the many different Federal and Industry groups creating Resiliency, CPS, and CPS Resilience guidance for vulnerable critical infrastructure
Develop better CPS resiliency metrics to quantify existing shortcomings in CPS-resiliency within Critical Infrastructure	Resiliency metrics are often multi-disciplinary. R&D is needed to focus metrics on the specific contributions of CPS as implemented within targeted CI environments. Metrics should

Framework Goal	Recommended CISA Approach
Develop better CPS resiliency metrics to measure Federal CPS-resiliency Outcomes/Objectives to enforce Federal CPS resiliency policies	be developed that can be used to measure National CI/CPS-resiliency objectives.
Coordinate Federal R&D investments for resiliency metrics/measures that are CPS-specific	CISA should encourage whole-of-government approaches to funding new and existing R&D efforts in academia, government and commercial labs that are specifically targeted at the contributions of CPS as implemented in targeted CI environments
Coordinate Federal R&D investments for resiliency tools/methods that model CPS-specific resiliency techniques	

2.14.4 Findings

The lack of a coordinated, purposeful approach to increasing CPS Resiliency leaves our National CI at risk to cyber-attacks. CPS Resiliency goals will remain ambiguous, with resources potentially being applied to CPS cyber-solutions that are sub-optimal, until the Federal government, led by CISA, can describe CI resiliency outcomes and objectives. Without set goals to focus activity, Federal policy, standards, and R&D investment cannot be synchronized to make meaningful progress. The result of this lack of direction is that DHS will not have necessary tools to strategically increase resilience within vulnerable National CI.

2.15 Synthetic Data

2.15.1 Description

In January 2023, CISA S&T launched the CISA Advanced Analytics Platform for Machine Learning (CAP-M) project to create a multi-cloud research environment ("ecosystem") in which to experiment with analytics on various cyber data sources. The CAP-M environment will include machine learning. The goal of the project is to counter cyberthreats and defend infrastructure from cyberattacks using actual data. CISA is interested in being able to share anonymized cybersecurity data sets from CAP-M with cybersecurity vendors for testing and training purposes. However, DHS cannot risk exposing sensitive cybersecurity data elements or patterns that would reveal capabilities and methods. CISA is investigating Synthetic Data, which is a privacy-preserving technology that creates new data that has been artificially created by computer algorithms. That is, the Synthetic Data software could build anonymized data sets that resemble CAP-M real-world data but does not reveal any sensitive information. These Synthetic Data sets can then be safely shared with cybersecurity vendors for training and testing purposes.

2.15.2 Importance to CISA

Developing a procedure to safely share anonymized real-world CAP-M cybersecurity data with cybersecurity vendors for testing/training purposes is essential for building tools that adequately meet

CISA's cybersecurity challenges. Releasing non-anonymized CAP-M data sets poses a great risk of exposing DHS CISA methods and procedures.

2.15.3 Details

Synthetic Data technology was developed to create the massive data sets required to train and test neural networks and AI/ML systems, but without the problems caused by ingesting mass quantities of real internet data (e.g., personal identifiable information (PII), copyrighted material, bias). Synthetic data can be fully synthetic (containing no original data) or partially synthetic (containing some original data).

The CAP-M project wants to use Synthetic Data for testing/training/validation of cybersecurity vendor capabilities against anonymized CAP-M data. Synthetic data would maintain the statistical properties, distribution patterns, and entity relationships present in the original CAP-M data while removing any direct identifiers, such as PII, data sources, or destinations. This anonymized data will allow CISA to share CAP-M data for purposes such as research, analysis, collaboration, or cybersecurity vendor testing/training.

With the anonymized Synthetic Data, vendors and CAP-M will be able to evaluate the performance of vendor models under different re-created conditions that mimic observed-but-anonymized CAP-M scenarios. In this way, CAP-M can assess the robustness and accuracy of various vendor models. By introducing changes to the data (e.g., a conflicting data entry) into the synthetic dataset, CAP-M will be able to evaluate how the model handles exceptions or detects purposely altered data, revealing vulnerabilities or weaknesses in the vendor model.

Another example of how synthetic data can be used is with simulated network traffic data that includes observed attack patterns and behaviors. With this synthetic data, Intrusion Detection System (IDS) vendors can teach their algorithms to recognize and respond to different types of real-world cyber attacks as seen in the original CAP-M data set. In this way, the synthetic dataset helps vendor re-create a real-world attack scenario without compromising the confidentiality of actual DHS network systems. CAP-M could use its synthetic data capability to widely share threat intelligence without compromising sensitive information.

Synthetic data has many advantages for CISA over real CAP-M data, such as:

- Prevents exposing sensitive data
- Ability to inject more variety into the dataset
- Fills the incomplete, inconsistent, or missing data gaps
- Eliminates the governance burden associated with sharing access to sensitive data (internal and external)
- Enables more efficient use of public SaaS (cloud migration for services) since there is reduced need to protect the data on the part of the SaaS provider
- Avoids real data retention policies
- Cheaper to generate large data sets

When considering the use of synthetic data for testing/training cybersecurity products, however, CISA must consider the relevance of synthetic data's disadvantages such as:

- Cannot handle complicated data sets with a large number of variables
    - Synthetic data set may not properly represent real-world conditions
    - This condition will lead to false insights and erroneous decision-making
- Does not eliminate bias, one of the biggest problems with using data in general
- May lead to "Giraffing," the generic name for the presence of objects where those objects do not exist or overrepresentation of portions of the data. Leads to the creation of bias in data generation.
- AI methods are good at interpolation within a data set, but not so good at extrapolation to new data
- Despite anonymization attempts, it may still be possible to link synthetic data to real people or DHS systems/components
- Vendors have more experience in protecting PII or Protected Health Information (PHI); not business intelligence, trade secrets, or Protected Critical Infrastructure Information (PCII)
- Synthetic data may not maintain currency, as it is a snapshot in time and may diverge from real-world trends

In summary, the CAP-M Synthetic data generation capability must be able to meet the desired outcomes (e.g., anonymizing data, creating new representative data sets) so that the synthetic dataset provided to vendors for testing and training adequately matches the real-world state without compromising sensitive information.

2.15.4 Findings

Synthetic data can provide CAP-M with a valuable tool for furthering cybersecurity research, analysis, collaboration, and vendor testing/training. This tool will increase CISA's ability to perform its cybersecurity mission. However, creating representative and sufficiently anonymized synthetic data will require significant effort. New policies and procedures for creating, testing, and validating synthetic data for release will be required.

2.16 Contract Optimization

2.16.1 Description

Contract optimization tools and services have been in use in commercial industry for several years. Their purpose is to make the creation, solicitation, and management of contracts less labor intensive and more effective. They often use online systems and may employ databases of contract clauses and natural language processing capabilities to create contracts, review contracts, and automate product acceptance and compliance checking.

2.16.2 Importance to CISA

CISA awards and maintains dozens of contracts each year to acquire technologies and services related to cybersecurity functions and R&D activities. They have an obligation to ensure they manage their procurements efficiently and ensure they get exactly the services and products they need from contracts to execute their mission. Contract optimization tools and services may provide them a way to improve the effectiveness of the procurement process.

2.16.3 Details

There are two types of technologies that are often considered part of contract optimization.

- **eSourcing** – using web-based systems to collect and compare information about several suppliers to help buyer select a preferred provider. The most prevalent tool used in the federal government for this is SAM.gov, which is the "The System for Award Management that offers a single login to navigate the federal award lifecycle."[34] DHS uses the Procurement Request Information System for Management (PRISM),[3] which is a software product that provides full procurement lifecycle support including all phases from advanced acquisition planning through contract closeout. There are also many commercial tools and services that perform similar functions. Both SAM and PRISM are effective tools that help CISA automate its procurement processes, but there are still opportunities for improvement.
- **Contract Optimization** – improving the efficiency of creation, analysis, execution, and maintenance of formal agreements. One model identifies the following five stages of contract optimization implementation.[35]
    - **Stage 1** – All contracts in electronic repository
    - **Stage 2** – Analysis and reports of contracts – metadata
    - **Stage 3** – Automatic contract creation
    - **Stage 4** – Obligation management and risk mitigation
    - **Stage 5** – Integrate transaction/billing systems to contractual terms and conditions (T&Cs)

There are a few contract optimization tools available in the marketplace, and more are being developed. The capability they offer that is not already in SAM and PRISM are AI features to review contracts and identify opportunities to improve or eliminate errors or ambiguities in critical contract clauses. Figure 22 shows the stages of a typical procurement cycle where contract optimization tools may be useful.

[34] General Services Administration. (n.d.) Sam.Gov. Retrieved from https://sam.gov/content/home Sam.gov.

[35] CIO-Wiki. (n.d.). Retrieved from https://cio-wiki.org/wiki/Contract_Optimization

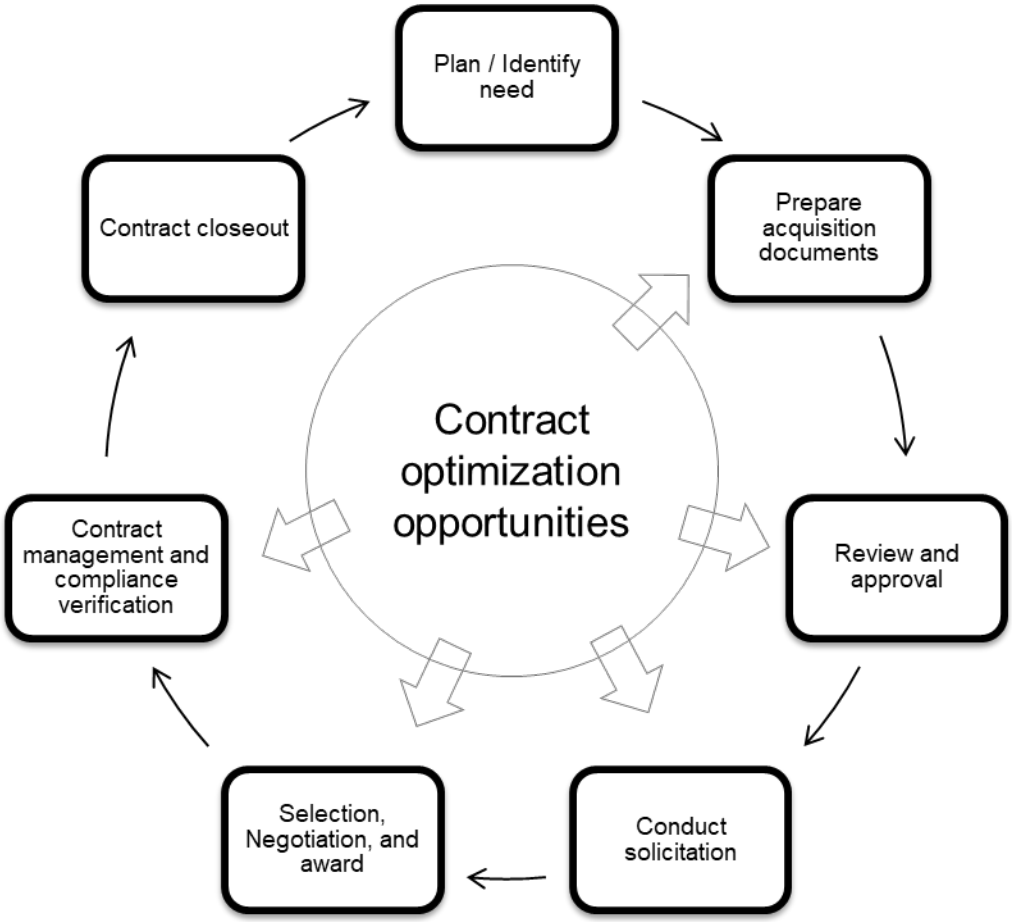

Figure 22: Procurement Lifecycle and Contract Optimization Opportunities

2.16.4 Findings

Contract optimization for CISA needs to fit within the standard acquisition processes required by DHS. The following areas appear to be near term opportunities for automation of acquisition processes within DHS and CISA:

- **Checking Information Technology Acquisition Review (ITAR) compliance** – This process is required and time consuming. It requires the manual review of contracts and supporting documentation prior to completing procurement actions to obtain Chief Information Officer (CIO) approval. There is an opportunity to automate this process as many of the compliance checks are rote.
- **Legal sufficiency review** – This process requires legal staff to review contracts to ensure compliance with applicable statutes, regulations, and procedures. It is a more complicated review process than ITAR compliance checking but may also be an opportunity for automation.

- **Acquisition document creation automation** – It is a tedious and error prone process to ensure all related acquisition documents are consistent. There is an opportunity for tools that can help write acquisition documents and manage changes to them that maintain consistency.

2.17 ML Drift Detection

2.17.1 Description

Model drift detection is a critical function in maintaining the performance and reliability of ML models deployed in dynamic, evolving environments. As models operate within changing data environments, the underlying patterns and relationships in the data can shift over time, leading to a degradation in prediction accuracy. Model drift detection involves continuously monitoring the model's predictions and comparing them to the expected outcomes based on historical data. This process often requires the establishment of a baseline performance metric, against which the model's current performance is measured. Various statistical techniques and data analysis methods are employed to detect significant deviations from the established baseline, signaling the potential for model drift.

The importance of model drift detection lies in its ability to ensure that ML models remain effective and trustworthy over extended periods. In scenarios where accurate predictions are crucial, model drift detection helps organizations identify when a model's performance is compromised due to changing data dynamics. By promptly recognizing and addressing model drift, organizations can take corrective actions, such as updating the training data, retraining the model, or adjusting its parameters to maintain consistent and reliable performance in the face of evolving real-world conditions. Models that drift too far to be corrected should be sunsetted and replaced with a different model.

2.17.2 Importance to CISA

Data and model drift can pose significant challenges to ML systems in production. By understanding the causes and effects of drift and implementing effective drift monitoring practices, organizations can ensure that their ML models remain accurate and reliable over time.

Monitoring the performance of models; using a drift detection model; and regularly retraining on updated data are just a few of the best practices organizations can follow to mitigate the risks of drift. By being proactive about drift monitoring, organizations can ensure that ML systems continue to deliver organizational and mission value.

Monitoring ML models for drift is just one aspect of a broader field called ML operations (MLOps). Understanding MLOps concepts is essential for any data scientist, engineer, or leader to take ML models from a local notebook to a functioning model in production.

CISA should consider an MLOps strategy that encompasses model and data drift detection. Due to the variety of model types in use at the agency, maintaining a centralized and consolidated approach to model and drift detection is challenging. CISA could benefit from investment in proof-of-concept evaluations and prototyping to help identify commercial products that could meet division-specific needs for model and data drift detection.

2.17.3 Details

Model drift is the decay of ML models' predictive power due to changes in real world environments. It can be attributable to a variety of reasons, including changes in the digital environment and ensuing changes in relationship between variables. There are two types of model drift: Concept Drift and Covariate/Data Drift.

Concept Drift occurs when the task that the model was designed to perform changes over time. For example, imagine that a model was trained to detect spam emails based on the content of the email. If the types of spam emails that people receive change significantly, the model may no longer be able to accurately detect spam.

Concept Drift can be further divided into four categories:

- Sudden Drift – occurs when there are sudden changes in the concept of the model.
- Gradual Drift – occurs when there are gradual changes in the concept of time series models.
- Incremental Drift – occurs when there are incremental changes in the concept of time series data.
- Recurring Drift – occurs when model drift re-occurs after a period of time.

Covariate/Data Drift is the change in the distribution of one or more of the independent variables or input variables of the dataset. This means that even though the relationship between feature and target variable remains unchanged, the distribution of the feature itself has changed. When statistical properties of this input data change, the same model which has been built before will not provide unbiased results. This condition leads to inaccurate predictions.

There are three types of Data Drift:

- Covariate Shift - The change of distributions in one or more of the independent variables (input features). This change means that due to some environmental change, even though the relationship between feature X and target Y remains unchanged, the distribution of feature X has changed.
- Prior Probability Shift – Occurs when the distribution of the input variables remains the same, but the distribution of the target variable changes.
- Concept Shift – Occurs when the relationships between the input and output variables change. This change means that the distributions of input variables (such as user demographics, frequency of words) might even remain the same, and instead the focus should be on the changes in the relationship between target variables.

ML model drift significantly reduces the effectiveness of ML models leading to inaccurate inferences and predictions. Implementing robust model drift detection methodologies and capabilities will significantly improve CISA's ability to monitor and proactively retrain ML models in support of their mission.

The Federal enterprise and national CI are becoming increasingly reliant on AI/ML for essential operations. Model and data drift detection capabilities will significantly reduce the risk associated with inaccurate inferences and predictions due to drift. In addition, enhanced visibility will alert engineers early when model performance is potentially impacted.

2.17.4 Findings

A comprehensive model and data drift strategy provides CISA with capabilities for model and data drift detection that will ensure ML models are performing as expected and provide early warning when models begin to drift. A centralized model and data drift capability helps to improve Federal and CI protection through improved detection of failing model performance and early drift detection.

2.18 Software Understanding

2.18.1 Description

Software understanding is an undertaking to discover software behavior by directly analyzing the software artifacts rather than primarily relying on proxy measurements (e.g., documentation, developer attestation, or development processes). While software understanding can be manual, and manual analysis is still done today across many missions; most missions, especially those conducted by CISA, require significant degrees of automation.

Examples of software understanding questions:[36]

- Is there an authentication bypass (e.g., a "backdoor") in this software?
- Could this software encrypt the contents of my archived data?
- Under what conditions might the camera/microphone be turned on?
- Does the control system software have a remote kill switch in it?

To answer mission questions for high-consequence systems and inform risk-based decisions and mitigations, software understanding must perform an evidence-based technical analysis of the software itself and its potential behaviors in order to present that evidence to system domain experts who can judge the risk of the behavior.

2.18.2 Importance to CISA

Software has become ubiquitous, especially throughout the nation's most sensitive and vital mission areas for which CISA is responsible. Software has been integrated into every facet of national security and critical infrastructure missions to the point where mission success depends on the behavior of their software, including third-party software.

Software can often exhibit unexpected behaviors which have been demonstrated to undermine or threaten the missions of the systems relying on such software. These unexpected behaviors are challenging to identify, and more challenging to manage, because of inadequate technical capabilities across U.S. government agencies and missions to analyze the potential behavior of third-party software. This inability to adequately analyze third-party software to answer vital mission questions has the potential to create significant risk for the U.S. government broadly, and CISA specifically.

[36] Note that automated tools cannot determine that these questions are fully answered, especially if adversarial malware is present, but the tools can reduce the gap between software complexity and software understanding depicted in Figure 23.

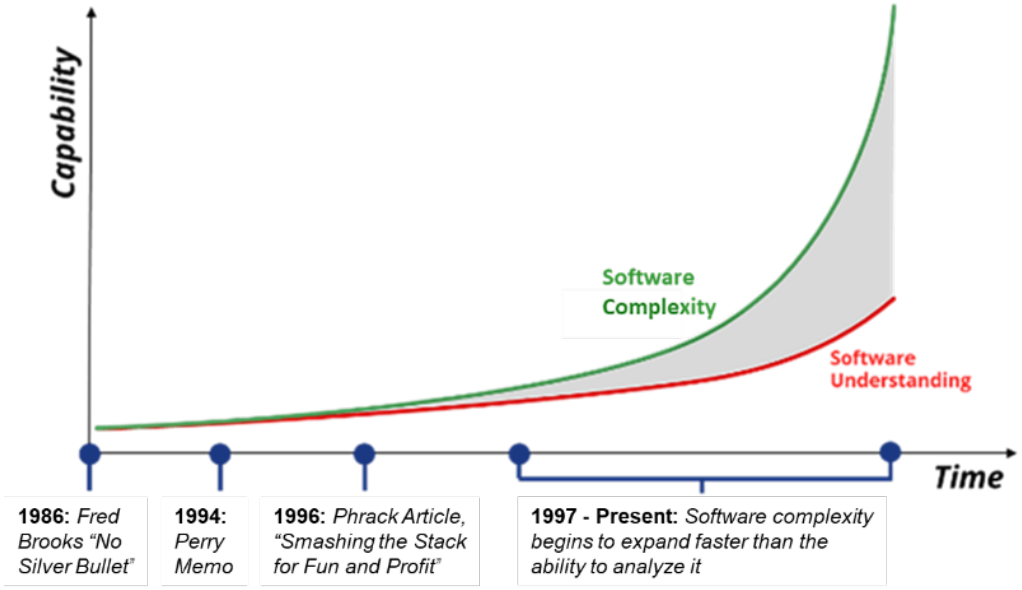

Figure 23: Increasing Gap Between Software Complexity and Software Understanding

The inability to adequately analyze software is partially due to the exponential advancement of software complexity without the simultaneous advancement of software understanding. The more this gap, highlighted in Figure 23, expands, the more it impacts CISA, and other major U.S. government mission owners and operators. The gap will continue to grow exponentially unless a national level software understanding solution is instantiated into policy and sufficient capability is developed.

2.18.3 Details

A national security or critical infrastructure mission owner would ideally rely on a technical evidence package derived from analysis of the software itself to gather information related to a mission question. This type of software analysis would enable the mission owner to make an informed, risk-based decision about the use of the software. This analysis requires a suite of technical capabilities that can handle the variety of software systems used in mission-relevant systems and that are designed to seek evidence related to the mission questions of value to system owners.

Today, software analysis is done either manually, automatically with tools that have limited scalability and limited reusability, or most often, not at all. Instead, mission owners often rely on proxy information (e.g., attestation, country of origin, and technical documentation). The paradigm shift in research required is analogous to car manufacturing where cars in the early 1900s were built bespoke, without interchangeable parts and by small teams of siloed expert mechanics and engineers vice post Ford Model-T where cars are built in a factory with standardized parts. The U.S. government needs to design and build an analogous software understanding factory, depicted in Figure 24. This effort will require an entirely different approach to how Software Understanding research is funded and conducted.

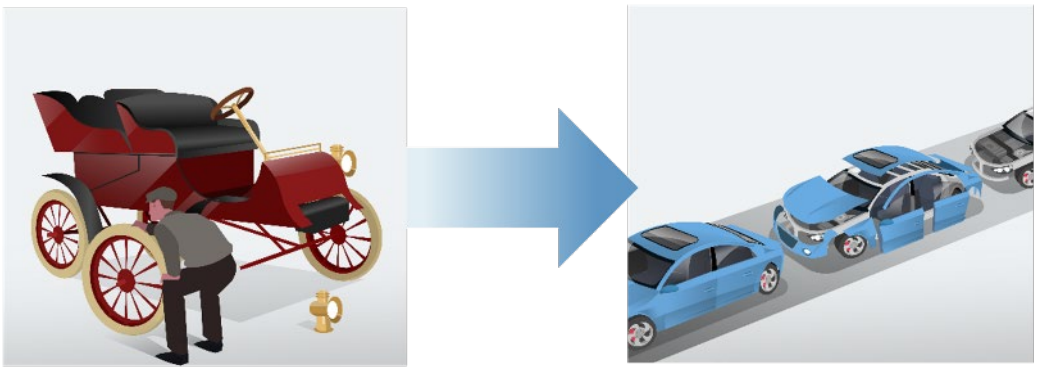

Figure 24: Current Software Development Approach Vice Future Concept

A two-decade legacy approach to investment has enabled software analysis – but not at scale. The ubiquity of software in government and critical infrastructure demands a significant update and unified approach toward developing a national software understanding capability.

2.18.4 Findings

Software understanding can provide valuable information on high-consequence cyber systems. A deeper understanding of the software that is integrated into every aspect of national security and CI will further CISA's ability to successfully perform their mission. However, the current state of limited manual and unscalable automatic analysis is not sufficient to meet the growing needs CISA has to understand the software on the networks they protect. Investments and supporting research into software understanding is needed for CISA to meet these needs and to reduce the gap between increasing software complexity and software understanding.

2.19 Digital Twin

2.19.1 Description

The NIST Internal Report (NISTIR) 8356 *on Considerations for Digital Twin Technology and Emerging Standards* offers the following definition for digital twin: "A digital twin is the electronic representation – the digital representation – of a real-world entity, concept, or notion, either physical or perceived."

The relationship of a Digital Twin to its Physical Twin is depicted in Figure 25.[37]

[37] Image Source: Government Accountability Office (GAO). (2023, February). Publication No. 23-106453

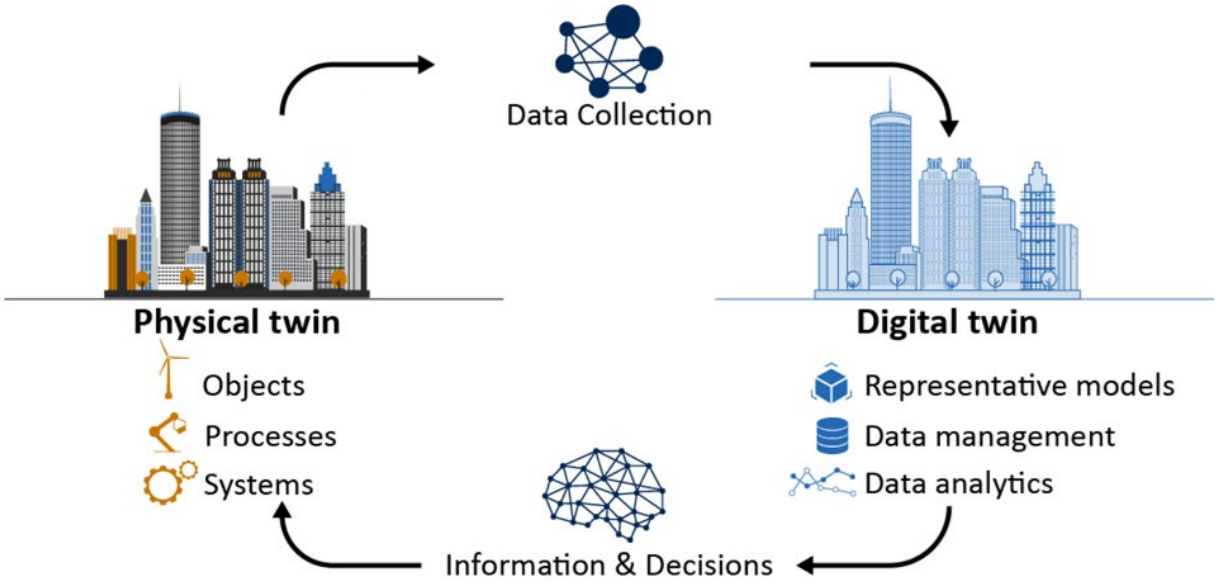

Figure 25: Digital Twin Relationship to Physical Twin

An alternate but parallel definition from the Government Accountability Office's (GAO) Science and Tech Spotlight on "Digital Twins – Virtual Models of People and Objects" is as follows:

> *Digital twins are virtual representations of people or physical objects, processes, or systems, ranging from vehicles to industrial plants to clinical trial patients. These "living" computational models integrate with data from a physical twin, such that any changes made to the physical twin can automatically lead to changes in the digital twin. Digital twins can be used to remotely maintain or monitor the physical twin or predict how it will perform.*

2.19.2 Importance to CISA

Digital twin technology can serve as a key enabler for any mission or activity that could benefit from a portable, digital replica of a real-world environment – whether physical or cyber-physical. Such digital replicas can enable high-fidelity observation, run-time analysis, predictive analysis, and operations within a non-production environment and can lead to downstream optimization of processes and decision-making. Digital twins allow for the execution and analysis of scenarios that may not otherwise be possible in a production environment.

2.19.3 Details

As CISA considers and prioritizes use cases that can benefit from digital twin technologies, it is important to recognize that digital twins can be of varying levels of maturity, complexity, and sophistication. NISTIR 8356 identifies five distinct use case categories, in order from least to most complex/sophisticated:

- Viewing static models
- Executing and viewing dynamic simulation models
- Streaming execution of dynamic simulations

- Real-time monitoring of real-world entities
- Real-time command and control of real-world entities

Some use cases may only benefit from simpler levels of digital twin implementations (e.g., a digital twin built for training and exercises should not have a command-and-control interface with the real-world environment). Other applications, depending on mission requirements and security and privacy restrictions, may be better suited for real-time operations as well as command and control integration. There are a multitude of use cases for digital twins, ranging from production efficiency to safety purposes. They are also becoming the backbone for smart city developments and personalized medicine. Digital twin use cases for simulating potential cyber-attack scenarios are also growing. Siemens is using a digital twin to simulate attacks to understand the impacts to its chemical manufacturing processes and to determine potential courses of action.[38] Similar efforts are underway in other enterprises to bolster security operations center (SOC) capabilities by using digital twins of enterprise networks to run high-performance analytics and perform what-if analysis to identify TTPs, as well as exercising new defensive capabilities.[39] The technology is also useful for identifying extended results of internal or external actions, enabling an organization the opportunity to identify and potentially mitigate negative results before they occur.

Existing standards and frameworks can help CISA better understand the components that may comprise a digital twin system. The standards landscape around digital twins is still evolving and in development. One international and widely-accepted standard that may be able to be extended to the CISA domain is the International Organization for Standardization (ISO) 23247 – Digital Twin Framework for Manufacturing. While focused on manufacturing applications – an area that has commonly used digital twin capabilities for predictive maintenance and cost simulation purposes – the document provides a reference architecture shown in Figure 26[40] on which CISA can build and/or adapt.

[38] Elsby, I. (2019, October 16). Digital Twin Does More Than Designing, Analyzing and Processing; It's the Cyber-Attack Combatant. Retrieved July 01, 2020 from https://news.siemens.co.uk/news/digital-twin-does-more-than-designing-analysing-and-processing-its-the-cyber-attack-combatant

[39] Amy-Vogt, B. (2020, February 26). Q&A: Accenture Creates Cyber Digital Twins to Simulate Potential Attack Scenarios. Silicon Angle. Retrieved July 01, 2020 from https://siliconangle.com/2020/02/26/qa-accenture-creates-cyber-digital-twins-simulate-potential-attack-scenarios-rsac/

[40] Image Source: ISO 23247 – Digital Twin Framework for Manufacturing

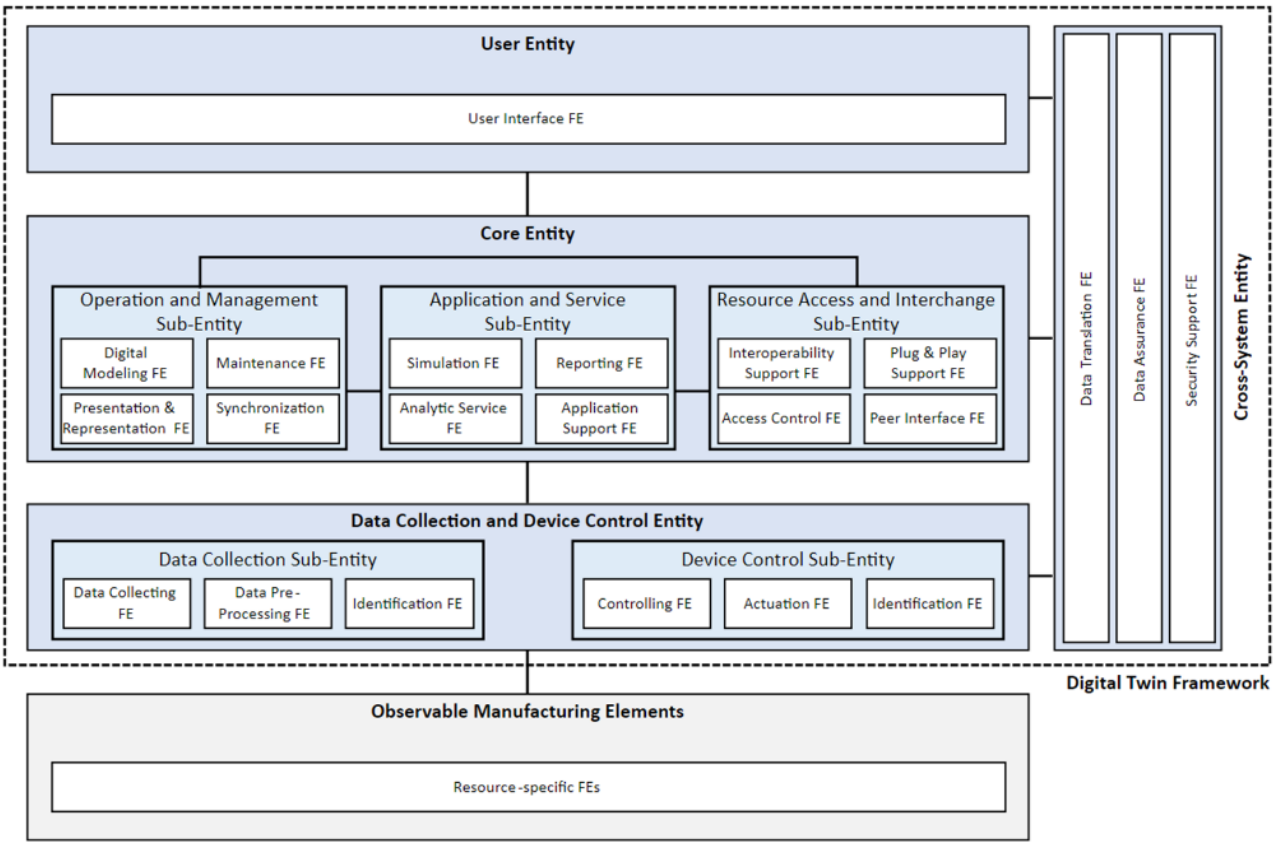

Figure 26: Digital Twin Framework for Manufacturing

This ISO standard consists of four parts, each of which can inform CISA's digital twin planning and implementation processes.

- Part 1 provides general principles and requirements for developing digital twins in manufacturing.
- Part 2 provides a reference architecture.
- Part 3 describes the static and dynamic information attributes necessary to represent physical elements.
- Part 4 presents technical requirements for exchange of information between entities.

There is also a patchwork of existing standards that can apply to digital twin elements (e.g., cybersecurity standards), which could include existing standards, such as, but not limited to, the following:

- IEC 62832 – Digital Factory Framework
- IEEE P2806 – System Architecture of Digital Representation for Physical Objects in Factory Environments
- IPC 2551 – International Standard for Digital Twins

- DIN SPEC AAS – Asset Administration Shell for Industrial Applications

Another pair of key emerging standards that may be of interest to CISA and may help inform potential use cases are the ISO/IEC Approved Work Item (AWI) 30172 on Digital Twin – Use Cases and the ISO/IEC AWI 30173 on Digital Twin – Concepts and Terminology. Both standards are under development.

2.19.4 Findings

Digital twin technology can enhance CISA's ability to reduce risk across cyber, physical, and communications infrastructure as well as CISA's ability to collaborate with key mission partners. Some examples of potential use cases relevant to CISA's mission could include the following:

- Efficient delivery of technical assessment services (e.g., red team assessments, penetration testing, breach, and attack simulation) with no impact or disruption to production environments
- Modeling, simulation, and analysis of critical infrastructure to support understanding of National Critical Function (NCF) inter-dependencies
- Visualization and modeling of environments to enhance analytical capabilities related to threats against physical infrastructure – e.g., active shooter, improvised explosive devices
- Enhanced understanding of CI operating environments and ability to discern differences between benign and malicious anomalies
- Training of mission operators and exercises with mission partners in replicas of real-world environments (versus simulated, representative environments)

Use of digital twin technology does introduce new cybersecurity and operational risks. Some key risk considerations that may be introduced through digital twins include the following:

- Instrumentation and connection of previously unconnected objects for monitoring/modeling purposes (i.e., physical to cyber-physical)
- Centralization of data and control interfaces in digital twin
- Digital twin manipulation (i.e., misrepresentation in the digital twin visualization/presentation to the user or other dependent objects)
- Remote control through digital twin interface
- Digital twin standardization as a vulnerability.

Appendix A: Acronyms

Acronym	Expanded
AAS	Asset Administration Shell
AB	Anonymous Broadcast
ABAC	Attribute-Based Access Control
AI	Artificial Intelligence
AIS	Automated Information Sharing
AWI	Approved Work Item
AWS	Amazon Web Service
CAP-C	Cyber Analytics and Platform Capabilities
CAP-M	CISA Advanced Analytics Platform for Machine Learning
CBDC	Central Bank Digital Currency
CDM	Continuous Diagnostics and Mitigation
CI	Critical Infrastructure
CIE	Cyber-Informed Engineering
CIO	Chief Information Officer
CIPAC	CI Partner Advisory Council
CIRCIA	Cyber Incident Reporting for Critical Infrastructure Act of 2022
CISA	Cybersecurity and Infrastructure Security Agency
CONOPS	Concept of Operations
COTS	Commercial Off-the-Shelf
CPS	Cyber-physical System
CRQC	Cryptography Relevant Quantum Computer
CSP	Cloud Service Provider
CVD	Coordinated Vulnerability Disclosure
CVE	Common Vulnerabilities and Exposures
CWE	Common Weakness Enumeration
DARPA	Defense Advanced Research Projects Agency
DHS	Department of Homeland Security
DIN	Deutsches Institut für Normung
DP	Differential Privacy
EDR	Endpoint Detection and Response
EU	European Union
FAA	Federal Aviation Agency
FCEB	Federal Civilian Executive Branch
FISMA	Federal Information Security Modernization Act of 2014
FL	Federated Learning
FSLTT	Federal, Local, State, Tribal, and Territorial
FTX	FTX Trading, Ltd.
FY	Fiscal Year
GAO	Government Accountability Office
GPS	Global Positioning System

Acronym	Expanded
GPT	Generative Pre-Trained Transformer
HE	Homomorphic Encryption
HMI	Human Machine Interface
HVA	High Value Asset
IACR	International Association for Cryptologic Research
ICAM	Identity, Credential, and Access Management
ICS	Industrial Control System
ID	Identification
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
IP	Internet Protocol
IPC	Institute of Printed Circuits
IR	Interagency or Internal Report
IRPF	Infrastructure Resilience Planning Framework
ISAC	Information Sharing and Analysis Center
ISO	International Organization for Standardization
IT	Information Technology
ITAR	Information Technology Acquisition Review
LEO	Low Earth Orbit
LLM	Large Language Model
MEO	Medium Earth Orbit
MFA	Multi-Factor Authentication
MITRE	The MITRE Corporation
ML	Machine Learning
MLOps	ML Operations
N/A	Not Applicable
NASA	National Aeronautics and Space Administration
NCCoE	National Cybersecurity Center of Excellence
NCF	National Critical Function
NIST	National Institute of Standards and Technology
NISTIR	National Institute of Standards and Technology Internal Report
NS/EP	National Security/Emergency Preparedness
NSA	National Security Agency
NSM	National Security Memorandum
OMB	United States Office of Management and Budget
ONCD	Office of the National Cyber Director
OSHA	The Occupational Safety and Health Administration
OSTP	Office of Science and Technology Policy
OT	Operational Technology
PA	Policy Administrator

Acronym	Expanded
PCAST	President's Council of Advisors on Science and Technology
PCII	Protected Critical Infrastructure Information
PDP	Policy Decision Point
PE	Policy Engine
PE	Prompt Engineering
PEP	Policy Enforcement Point
PET	Privacy Enhancing Technology
PHI	Protected Health Information
PII	Personal Identifiable Information
PKI	Public Key Infrastructure
PLC	Programmable Logic Controller
PMO	Program Management Office
POS	Proof of Stake
POW	Proof of Work
PQC	Post-Quantum Cryptography
PRISM	Procurement Request Information System for Management
QKD	Quantum Key Distribution
QPP	Quantum Permutation Pad
R&D	Research and Development
RBAC	Role-Based Access Control
RIPDWG	Resilient Investment Planning and Development Working Group
RRAP	Regional Resiliency Assessment Program
RSA	Rivest–Shamir–Adleman
RSB	Resilience Services Branch
RWC	Real World Crypto
SaaS	Software as a Service
SAM	The System for Award Management
SATCOM	Satellite Communications
SBOM	Software Bill of Materials
SCADA	Supervisory Control and Data Acquisition
SGX	Software Guard Extension
SME	Subject Matter Expert
SMPC	Secure Multiparty Computation
SOC	Security Operation Centers
SPEC	Specification
SRMA	Sector Risk Management Agency
T&C	Terms and Conditions
TEE	Trusted Execution Environment
TOR	The Onion Router
TT&C	Tracking, Telemetry, and Control
TTP	Tactics, Techniques, and Procedures
U.S.	United States

Acronym	Expanded
VEP	Vulnerabilities Equities Process
VM	Virtual Machine
WG	Working Group
ZKP	Zero Knowledge Proof
ZTA	Zero Trust Architecture