

CANADIAN CENTRE FOR **CYBER SECURITY**

A zero trust approach to security architecture

Management

Foreword

A zero trust approach to security architecture (ITSM.10.008) is an UNCLASSIFIED publication that has been issued under the authority of the Head of the Canadian Centre for Cyber Security (Cyber Centre). For more information, email, or phone our Service Coordination Centre:

Service Coordination Centre

contact@cyber.gc.ca

(613) 949-7048 or 1-833-CYBER-88

Effective date

This publication takes effect on March 15, 2023.

Revision history

Revision	Amendments	Date
1	First release.	March 15, 2023

ISBN 978-0-660-48189-0
CAT D97-4/10-008-2023E-PDF

Overview

This publication provides a description of zero trust (ZT) security concepts and how organizations can benefit from implementing a ZT security architecture to safeguard their assets. It will help organizations understand the importance of transitioning to a ZT security architecture and how they will need to shift their mindset and work in collaboration with all members of their organization to improve their cyber security posture. It provides a description of ZT security concepts and how organizations can benefit from implementing a ZT security framework. We have identified a few of the best practices organizations can follow to help prioritize their efforts when implementing a zero trust architecture (ZTA).

Choosing ZT guidelines or a framework and using trusted expertise and resources are important steps to implementing and continually improving an effective ZT strategy.

The Government of Canada (GC) is developing a ZT security framework that will help GC departments and agencies improve their overall security posture. The GC ZT security framework will align with the pillars in the CISA and NIST references. In the meantime, to help organizations choose which framework or set of guidelines aligns best with their business requirements and network infrastructure, we've provided an overview of the three commonly cited and trusted ZT frameworks/guidelines:

- The National Institute of Standards and Technology (NIST): [special publication 800-207: Zero Trust Architecture](#) [1]
- Cybersecurity and Infrastructure Security Agency (CISA): [Zero Trust Maturity Model](#) [2]
- National Cyber Security Center (NCSC): [Zero trust architecture design principles](#) [3]

Table of contents

1	Introduction	5
1.1	What is zero trust and zero trust architecture?	5
1.2	ZT security approach	6
2	Industry accepted ZT frameworks	7
3	Benefits of a ZT security framework	8
4	Best practices for implementing a ZTA	10
5	Challenges to organizations	14
6	Additional zero trust guidelines	16
6.1	NIST special publication 800-207: Zero Trust Architecture	16
6.2	CISA: Zero Trust Maturity Model.....	17
6.3	NCSC: Zero Trust Architecture Design Principles	20
7	Summary	22
8	Supporting content	23
8.1	List of abbreviations	23
8.2	Glossary.....	24
8.3	References.....	25

List of figures

Figure 1:	CISA's depiction of the foundation of ZT.....	18
-----------	---	----

List of tables

Table 1:	CISA's high-level ZT maturity model.....	20
----------	--	----

1 Introduction

IT networks have grown in size and complexity to meet business needs with new evolved technologies, such as hybrid cloud infrastructures. Cyber threats have kept pace with these changes and often take advantage of the security gaps following these hasty transformations. As users, information, and services are dispersed across various locations, there are no longer defined perimeters around organizations' resources. Organizations can no longer depend on conventional perimeter-based defence to protect their critical systems which makes taking a ZT approach more important now than ever before.

1.1 What is zero trust and zero trust architecture?

A ZTA is an enterprise approach to a system design whose security perspective is based on ZT principles. Its core principle is that inherent trust is never granted by default to any subject.

With a ZTA:

- every interaction initiated between a user and a resource is strongly authenticated and authorized
- access control to resources is made as granular as possible
- access control decisions are based on a **dynamic evaluation** of the trust context for each access request

With ZT the communication between users, systems, and devices is continuously authenticated, authorized, and validated. It's founded on policy-based access controls (PBAC) such as role-based access control (RBAC) and attribute-based access control (ABAC). A ZT architecture enforces access policies based on context such as the user's role, the time of day, geolocation, the device, and the data they are requesting. The level of access that is granted is dynamically adjusted based on the level of trust established with the subject. In short, the more trust that an information system can develop in a subject, the more access that subject can be granted.

A ZTA will also prevent lateral movement within the IT environment. It's important to note that preventing lateral movement is the primary goal of ZT, not the elimination of the legacy boundary defence or bring your own device (BYOD). These are things that may be enabled by ZT but should not be seen as primary reason for doing ZT.

According to NIST [1], an operative definition of zero trust and zero trust architecture is as follows:

Zero trust (ZT) provides a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised. Zero trust architecture (ZTA) is an enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies. Therefore, a zero trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a zero trust architecture plan.

It's recommended that organizations incrementally implement ZT principles and operate, if possible, in a hybrid ZT and perimeter-based model until they fully transition to a ZT model. They should do so while continuing to invest in security

modernization initiatives to improve their cyber security posture. A ZT security model used in conjunction with a risk-based approach to security management, such as [IT security risk management: A lifecycle approach \(ITSG-33\)](#) [4], will enable organizations to be flexible, agile, and adaptable all while providing security assurances of confidentiality, integrity, and availability for business assets.

1.2 ZT security approach

Migrating to ZT is a complex process that requires changes to be made at all levels of an organization. Organization leaders, administrators, stakeholders, and users must fully commit and collaborate to implement new technologies, working practices, and policies to enhance protection and support the development of a ZTA. With ZT, you can improve the security posture of your organization all while providing better use of resources, greater compliance, and resiliency throughout your organization. A shift to a ZT mindset is needed and this means you should:

- assume that the connections to your network infrastructure and resources are always hostile
- assume that all network traffic and requests to access your resources may be malicious
- assume that an adversarial actor will attempt to eavesdrop on communications and data flow
- implement thorough logging and monitoring of access requests, system management, configuration changes and network traffic to measure the integrity and security posture of all assets
- authenticate and explicitly verify and authorize each request for access to the least privilege and ensure that access that has been granted is time-bound
- make access control decisions based on dynamic risk-based policies
- accept that granting access to sensitive resources will increase your risk of cyber threats
- have an incident response and recovery plan in place that will ensure damage control and business continuity

2 Industry accepted ZT frameworks

As federal agencies and organizations move away from relying solely on perimeter defence security concepts to instead adopt a ZT mindset, they are encouraged to reference trusted ZT frameworks and guidelines. CISA, NIST, and NCSC have published ZT guidelines to support agencies and organizations develop and implement a ZT framework that they can tailor to their unique business requirements, mission, and threat landscape.

This section provides a brief description of the industry accepted ZT frameworks. A more detailed overview can be found further in this document.

The GC is developing a ZT security framework that will align with the pillars in the CISA and NIST references. The GC's goal for this shift to ZT is to:

- ensure that the GC continues to support a resilient enterprise digital security ecosystem in which government services are delivered safely and securely
- provide a seamless and enhanced user experience for authorized users
- provide a secure platform that ensures protection of systems and data hosted
- (physical & virtual) within the converged GC's network environment
- offer end-to-end protection of the GC's information, applications, devices, network, hardware, and physical facilities
- develop, adopt, and enforce mature security processes, governance structure, and standards
- ensure the confidentiality, integrity, and availability of the GC's information technology (IT) infrastructure and the government's critical business and customer data

Until the GC ZT framework is published, organizations are encouraged to use the commonly cited ZT framework or set of guidelines from NIST, CISA, and NCSC. Choosing a ZT framework and following the guidance from trusted resources are critical steps to implementing an effective strategy.

NIST developed ZT guidelines recommending agencies design and deploy a ZTA with adherence to seven basic tenets, which we will expand on further in this section.

The Department of Defense (DoD) and the National Security Agency (NSA) guidelines offer a more operational and micro-level approach to ZT than the guidance from NIST. NSA's ZTA is very similar to DoD's ZTA and includes the same seven pillars. These two differ in their focus. DoD's ZTA was developed with a defence-specific mission and requirements in mind and uses the Department of Defense Architecture Framework (DoDAF) while the NSA ZTA was developed for the NSA and defence industrial base organizations.

CISA drafted a ZT maturity model and draws upon the pillar concept from the DoD and NSA ZTA. CISA's ZT model focuses on five distinct pillars supported by overarching capabilities for visibility and analytics, automation and orchestration, and governance. It was created to help all types of federal agencies in the development of their ZTA.

The United Kingdom (UK) NCSC guidance is based on eight principles that represent the main building blocks and architectural considerations needed to develop a ZTA. Every organization will have a somewhat different approach to achieving ZT depending on their business requirements, the technologies they use, and their threat landscape. The NCSC ZT guidance was developed with that in mind and believes that most ZT approaches can be linked to these eight core principles.

3 Benefits of a ZT security framework

A ZT security framework is a holistic approach with the objective of enhancing your organization's security posture to protect its sensitive data and digital assets.

There are several benefits of implementing a ZT security architecture for your organization. Here are a few:

1. Provides greater network and lateral movement protection

We often hear of attacks that involve a compromised user account or device being used as an entry point to an organization's network. With most security controls being on the boundary of organizations' networks, it's difficult to detect an attacker if they successfully breach that first layer of defence. Once in, the attacker will then progress laterally in the network to gain access to credentials or other sensitive information. With a ZT approach, all applications and services can only communicate once they have been authenticated. ZT reduces the risk of lateral movement because all communication, regardless of its origin is considered untrusted until authenticated and specific access has been granted.

In a ZT model, every action a user or device takes is subject to some form of policy decision. This allows the organization to verify every attempt to access data or resources, making it challenging for an attacker to break the system.

2. Provides greater visibility and improved monitoring

The ZT security approach requires organizations to register and maintain compliance on all devices accessing information and resources on the network. It also requires that every user go through a stringent authentication process to gain access to specific resources. This provides visibility on who accesses what resources and for what purpose and can help identify what security measures need to be applied to each resource. ZT requires ongoing monitoring of all activities and communications. This allows organizations to gain full visibility into their network traffic and better detect potential threats and respond to them in a timely manner.

3. Improves incident detection and response

ZT offers new capabilities for incident response as detailed information is available about all suspicious access requests and which user, device, data, and application were involved. With ZT, when an incident is discovered, it can be linked back to specific entities, applications, and data.

4. Improves access control over cloud

Access control and loss of visibility are organizations' greatest worry when moving to the cloud. Although cloud service providers (CSP) offer some security features, the protection of your organization's assets remains a shared responsibility between you and the CSP. You have limited control inside the cloud.

ZT requires that all assets on the cloud be classified so that the right protections and access controls can be selected and implemented. With ZT, you can ensure that everything attempting to connect to your organization's cloud infrastructure is legitimate.

5. Improves data protection

Traditionally, if your network perimeter (such as a firewall) has been breached the attacker can leverage lateral movement to potentially find and steal sensitive information, customer data or intellectual property. This can damage an organization's reputation and potentially lead to legal repercussions. By shifting the focus from perimeter defence to securing individual resources, organizations face reduced risks of data breaches and theft.

Strong authentication and validation of connections under ZT principles makes it possible to ensure data privacy. ZT follows the principle of least privilege, granting only access required by the user to do their work. For each connection request, every entity is assumed to be hostile until both users and devices are authenticated, and permission assessed. The granted access is continually reassessed as the state changes, such as the user's location or the data being accessed.

6. Help support continuous compliance and facilitate auditing

A ZT architecture also helps support continuous compliance with privacy standards and regulations by evaluating and logging every access request. Tracking the user's and application's identity when requesting access, as well as the time and location of when the requests are made, allows for a complete audit trail. As a result, minimal effort is required to comply with audits and to uphold governance.

7. Secures the remote workforce

One of the biggest challenges organizations face today is the rapid shift to a remote working ecosystem. With users working remotely and data being shared, perimeter defences like firewalls are no longer sufficient. Each remote and hybrid worker expands the attack surface and creates new gateways for attackers. With ZT, segmentation of the network and creation of micro-perimeters with stringent identification and validation policies, control is provided as to which user, device, and application have access to the secured zones.

4 Best practices for implementing a ZTA

In this section we describe some of the best practices you can follow to help you implement a ZTA in your organization.

1. Authenticate all connections

Never trust your local network. In traditional architectures, any network connection that originates from inside the perimeter is considered trusted. It's assumed that the user or activity on the network has already been authenticated and authorized. Instead of implicitly trusting network connections that logically originate from your local network, all connections should be properly authenticated before being allowed to connect. At a minimum, authenticate the user and device requesting access. When more stringent security measures are required, add the geolocation, the date, and the time to the authentication requirements. With ZT, you should build trust into the devices, users, and services operating within your network.

2. Implement ZT policies

Creating and implementing policies is one of the most important and labour-intensive steps needed to build a strong ZTA. It requires that organizations truly understand all areas that need to be protected and the level of protection that needs to be applied to secure traffic flow.

When developing and implementing ZT policies, start by answering the following questions:

- Who are the users?
- What do they need to access?
- Why do they need access?
- Where do they need access from?
- Where are the users and endpoints located that are requesting access?
- How is access being granted and approved?

3. Establish a "trust engine"

A trust engine is used to evaluate trust and subsequently allow or deny access based on the collection and assessment of a wide variety of attributes relevant to the security context of an access request. This can include:

- the security state of the devices involved in the request (software versions, patch levels, physical and logical location, date and time, monitoring status, observed access history, etc.)
- the behavioural attributes of the requester (usage patterns, time-of-day, etc.)
- the enterprise-level attributes that represent the current security context (e.g., heightened security state based on monitoring and event indicators)

The centrepiece of ZT is a dynamic "trust engine" that has global visibility across all levels of architecture and that incorporates feeds from key components of operational security. This trust engine will take time and effort to develop.

4. Know your assets and network architecture

When building a ZT architecture, it's important to know your assets and network architecture. Create an inventory of your data, users, devices, and applications being accessed on your network. Know how you will be controlling and managing the access requests and what potential compromises might occur if access is granted. Understanding the value of your organizational data and the risks associated to any compromise of this data is paramount in implementing any significant changes to your architecture and IT environment.

5. Use multi-factor authentication (MFA)

With MFA, two or more different authentication factors are needed to unlock a device or sign-in to an account. MFA uses combinations of the following factors to authenticate a user: something you have, something you know, or something you are. These factors can and should be adjusted depending on the sensitivity of the data and resources being accessed. MFA significantly reduces the chance that attackers can use compromised credentials to access your systems and data. It's an essential prerequisite of ZT.

To better understand how MFA can help secure your devices and accounts refer to our [Secure your accounts and devices with multi-factor authentication](#) [5] publication. If you require assistance in determining a target level of authentication assurance, you can reference our technical guidance [User authentication guidance for information technology systems](#) [6] which complements the Treasury Board of Canada Secretariat (TBS) [Guideline on defining authentication requirements](#) [7].

6. Use encryption for all traffic

ZT requires the encryption of all traffic to reinforce the ZT tenet that all access to resources must be explicitly granted. This is contrary to the traditional perimeter defence model where access is inherently provided. Encryption helps with data loss prevention as files "leaked" or "stolen" are unusable without the ability to decrypt. For this same reason, encrypting helps protect the information end-to-end and helps with data "sniffing".

7. Enforce policy-based access

Organizations need to develop dynamic risk-based policies and make sure they are enforced correctly. These policies are a set of access rules that are assigned to users, data, assets, applications, and services. Resource access permission policies will vary based on the sensitivity of the resource. Least privilege principles should be applied to restrict both visibility and accessibility.

With ZT, the implied trust based on the network location (IP address) is no longer a condition for authentication. Instead, identity-based authentication is used to establish trust and provide access to specific resources at a given time and from a specific location or device.

8. Use privileged access management (PAM) and secure administrative workstation (SAW)

PAM should be used to protect all admin accounts that require elevated privileges. When high-level administrative access is required, it should be granted in a manner referred to as "just-in-time" access. This means temporary access is granted and then revoked as soon as the task is done and access is no longer required. Before anyone is allowed a privileged access session, the request for this session should have been reviewed and granted by a

different user. This prevents someone from self-granting privileged access without proper approval.

PAM solutions can automate this approval process as well as record everything that was done during the privileged access session and save that information.

A SAW, which is a dedicated physical or virtual machine, should be used by administrators strictly to perform administrative tasks. These secured workstations provide added security for IT administrators working with servers and applications to carry out sensitive tasks that pose a higher risk if compromised. The dedicated workstation cannot be used for web browsing, email, and other risky applications.

The PAM works together with secure administrative workstation solutions. Users must log into the SAW through the PAM to access protected accounts. A PAM platform can be leveraged to secure and control all access to privileged accounts, including individualized access permissions, the assigned access time, and the allowable actions.

9. Implement the principle of least privilege, RBAC and ABAC

In a ZT environment, it's important to apply the principle of least privilege in as many areas as possible, especially around privileged access and management of security. ZT is based on the principle that a user should only be given just enough access to allow them to complete a particular task. RBAC can be implemented to enforce least privilege in a codified manner, as it maps user access rights to their role within the organization. Combining RBAC with PAM will further enhance the control of access within your organization.

With ABAC, attributes, at the granular level, are defined to determine which user can access the data. For ABAC, access control policies are based on rules around the characteristics or properties of each requester and every data point is checked to ensure that the attributes match the requester's permissions before being released.

10. Monitor and log devices and services access

Continuously monitor how devices and services are interacting, what is being requested, what activities are performed, and what data is accessed. The continuous collection of log data and the use of security analytics to flag anomalies for investigation will help identify suspicious activity to detect and stop attacks. A properly implemented security information and event management (SIEM) solution will make collecting, correlating, and investigating large amounts of data easier and faster for administrators. SIEM can offer the level of visibility that is needed to ensure that users connected to the network are trustworthy, an essential part of a ZT strategy

11. Manage all devices

Verifying your users is necessary but not sufficient. The principles of ZT also extend to endpoint devices. To ensure that only trusted devices are allowed on your network, start by establishing a unique, traceable identity for each. These identities offer visibility on your network and will expose untrusted devices. The identities you create for devices are necessary to authenticate permissions and access, according to the policies you define. The defined policies should cover device certification, configuration, and compliance. Device certificates can help with inventory control as well as authentication. They should be encrypted and password protected so that they cannot be easily used if leaked.

The operating system and applications that run on the device should be properly configured, securely provisioned, and kept up to date. The access control system should be in place to ensure that all policy controls are in effect before the information/data is accessed. Access should be modified if a security feature is not up to the level required by compliance.

Organizations should consider the use of a trusted platform module (TPM). A TPM, a secure cryptographic integrated circuit (IC), is a hardware-based approach to authenticate a device. It securely stores artifacts, such as passwords, certificates, and encrypted keys, to manage device and user authentication, network access, and data protection. TPMs are used in modern personal computers (PCs), electronic notebooks, mobile phones, and network equipment.

BYOD policies and the use of mobile devices for work have become common in today's modern workforce. With ZT, all devices must be individually authenticated before gaining access to the organization's resources and data. This offers greater control over BYOD environments. It allows IT administrators to set and enforce granular authorization and authentication policies to determine what is required for a device to be allowed to access its resources.

12. Use network segmentation or micro-segmentation

Network segmentation is one of the potential approaches that could be used when implementing a ZTA. It's the practice of creating sub-networks within the overall network to prevent attackers from moving laterally once inside the perimeter. Typically, companies build network segments via virtual local area networks (VLANs) and firewalls, subnets, and security zones.

Micro-segmentation logically divides the data centres and cloud environments into distinct security segments up to the individual workload level. It relies heavily on the use of managed policy enforcement points throughout the network to dynamically control the communication between components based on policy. This is done to protect sensitive data and services from both internal and external threats. It provides layered security and allows for restricted access to assets on a granular level. This ensures that even if an attacker does enter the network, the amount of damage they can cause is limited.

13. Use software-defined perimeter (SDP)

SDP can support many of the ZT concepts and principles. This includes fine-grained, least privilege access control for all access requests, encryption of data in transit, micro segmentation, and so on. It offers an alternative to virtual private networks (VPN) and secure remote access to any application, located anywhere. It's a network boundary that is based on software, not hardware. SDPs are built on an adaptive trust model where access is granted based on user identities and least privilege basis defined by granular ZT policies and principles, not IP addresses. This allows remote users to connect to application without giving them network access which reduces the attack surface.

5 Challenges to organizations

Organizations will face many challenges when transitioning to a ZT security model. With ZT, access authentication, verification, and monitoring need to be controlled at a granular level and most, older technologies are not compatible with this. Organizations must also have a comprehensive understanding of their business requirements as this is imperative to a strong ZT model. As a starting point, organizations should:

- identify their network's most critical and valuable data, assets, applications, and services
- know who their users are, which applications and services they are accessing, their geolocation and how they are connecting

This will allow them to focus their efforts on prioritizing and protecting their resources as part of their ZT implementation journey. For organizations to maintain or improve their security posture, there will be a requirement for more mature and better integrated security across their organization and through all technology layers. Traditional network security mechanisms will still be required when pivoting to a ZT security model and after the implementation of the ZT model to support the transition to smaller trust zones. They will also be used as information sources for your organization's trust engines.

Here are some examples of challenges that organizations may encounter:

- Technicians and administrators across organizations will need to increase their efforts to define and implement detailed attributes of every user and resource to support trust/access decisions.
- Users might become frustrated at having to use multi-factor authentication (MFA) and needing to authenticate themselves more often during their work.
- Devices will require hardware tokens which can be costly and require time to rollout across an organization. There will be an increase in cost for software, hardware, and training.
- Firewalls may be older and not support some of the dynamic functionality required. Phased plans for introducing new equipment must be balanced with cost considerations.
- Technical resources for implementing ZT model may be scarce.

Migration to a ZTA can get messy, as the transition period can get complicated when some systems are ZT compatible while others in the network are not. CISA's draft ZT Maturity Model [2] is one of many roadmaps to support the transition to ZT. The maturity model provides examples of a traditional, advanced, and optimal ZT architecture which will allow organizations to incrementally transition to a ZTA and eventually reach an optimal cyber security posture.

A permanent shift in mindset must be adopted and embraced fully for a ZT solution to work. Implementing ZT requires a holistic integrated effort. A lack of full support throughout the organization, possibly from leadership, administrators, stakeholders, or users, can delay the process and affect productivity.

It can take years to move to a full ZTA. To avoid productivity interruptions, it's recommended that organizations follow an incremental implementation to allow for a smoother transition and more time to adjust to the new security framework, policies, and processes. Organizations need to recognize that ZT is not a passive approach where you build it and leave it. It's a long-term investment that requires time, effort, financial investment, and ongoing maintenance.

As your organization and business grow, it's imperative that your ZT framework evolves simultaneously. For example, access controls need to constantly be updated to ensure the right people have access to specific information. Keeping the permissions accurate and up to date will require ongoing input and is imperative to block access to sensitive information.

Some vendors will claim that their products are the answer to adopting a full ZT security model. Be wary of these vendors. The reality is that there's not a single ZT vendor or solution that can offer all the answers to implementing a ZTA. ZT is more than just a technical solution, it requires a fundamental shift in how security is managed.

6 Additional zero trust guidelines

In the current threat environment, with increasingly sophisticated cyber threats, organizations can no longer depend on conventional perimeter-based defences to protect critical systems and data. As a result, on May 12, 2021, the White House issued [Executive Order \(EO\) 14028 on Improving the Nation's Cybersecurity](#) [8]. It requires US federal government to take action to strengthen national cyber security and achieve certain specific ZT goals by the end of fiscal year 2024. The EO also highlights the need for “comprehensive security monitoring; granular risk-based access controls; and system security automation in a coordinated manner throughout all aspects of the infrastructure in order to focus on protecting data in real time within a dynamic threat environment”. A typical migration plan will assess an agency's current cyber security state and plan for a fully implemented ZTA.

The guidelines described in this section offer a starting point on how to use ZT to strengthen security for every aspect of an IT environment. All these guidelines describe ZT as a strategic approach to cyber security and focus on the core premise of ZT: never trust, always verify. It's the elimination of implicit trust and the continuous validation at every stage of a digital interaction.

6.1 NIST special publication 800-207: Zero Trust Architecture

The NIST ZT guidelines, which were the first to be released in August 2020, introduced a list of ZT basic tenets to be used to build a ZTA. The main objective was to help agencies reduce implicit trust zones and better understand their network infrastructure and how their data, applications and computing systems communicate. It offers a broad set of use cases where ZT can be implemented and how the ZT tenets can be applied to existing federal compliance guidance.

Here are the seven basic tenets recommended by NIST to ensure the success of any ZT security approach. These tenets form the foundation of an architecture that supports the principles of ZT.

1. Every data source and computing service is considered a resource

Any file, data, digital asset, and all types of endpoints that contain company information and communicate with the network must be considered a resource.

2. All communication is secured regardless of the network location

In ZT, the network is always considered hostile. Appropriate security controls should be implemented to protect the confidentiality, integrity, and availability of data in transit. All access requests from assets must meet the same stringent security requirements for authentication regardless of where the requests originated from. The assets may be located on enterprise-owned network infrastructure or any external network. The same security verifications standards must apply to all. Trust can never be implicit.

3. Access to individual enterprise resources is granted on a per-session basis

In an ideal ZT architecture, access to a particular resource (file, data, digital asset) is contingent upon authentication and authorization and is time bound. The access granted should be on the principle of least privilege and is given only to a single resource. Attempts to access other resources will require reauthorization using other

explicit verification rules. Organization administrators will need to identify how to enforce access policies on individual resources and how the multilayers of authentication will occur for each access request.

4. Access is provided based on a dynamic risk-based policy

Access to a resource is not a static concept. An authorized user can still be denied access to a particular resource if there is reason to believe that the access request is suspicious and does not meet policy.

Access to resources is determined by dynamic policy. Those authorized to access the resource must still authenticate themselves and prove they meet the enterprise policy to be granted the session. They must do so by incorporating some combination of:

- state of client identity (including application, service, username, password, etc.)
- asset status (IP address, networks accessed, software version, patch level, geolocation, updates installed)
- other analytics-driven criteria (geolocation, past request patterns)

5. All assets, both internal and external, are continuously monitored and their integrity and security posture are measured

No asset is inherently trusted, every network and device is always vulnerable to attack. Organizations should have a robust monitoring and reporting system so that they are able to continuously evaluate the security posture of the asset when a specific request is received against that asset. Organizations should implement mitigation techniques and apply patches and fixes when needed.

6. All resource authentication and authorization are dynamic and strictly enforced before granting access

Your organization should check explicitly every time that a user attempts to access a resource. Scanning and assessing threats, and continually re-evaluating trust must be an ongoing process. An organization must have identity credentials, access management, and asset management in place.

Implementing MFA along with continuous monitoring is required to ensure that reauthentication and reauthorization occur as defined by security policies.

7. As much information as possible is collected about the current state of assets, network infrastructure, and communications to improve the organization's security posture

Organizations should collect as much information as possible about the current state of the network and communications to help improve the security posture of their overall architecture. Insights provided by this information will allow your organization to continuously learn and improve its security settings and policies to reduce risk and enforce proactive protection.

6.2 CISA: Zero Trust Maturity Model

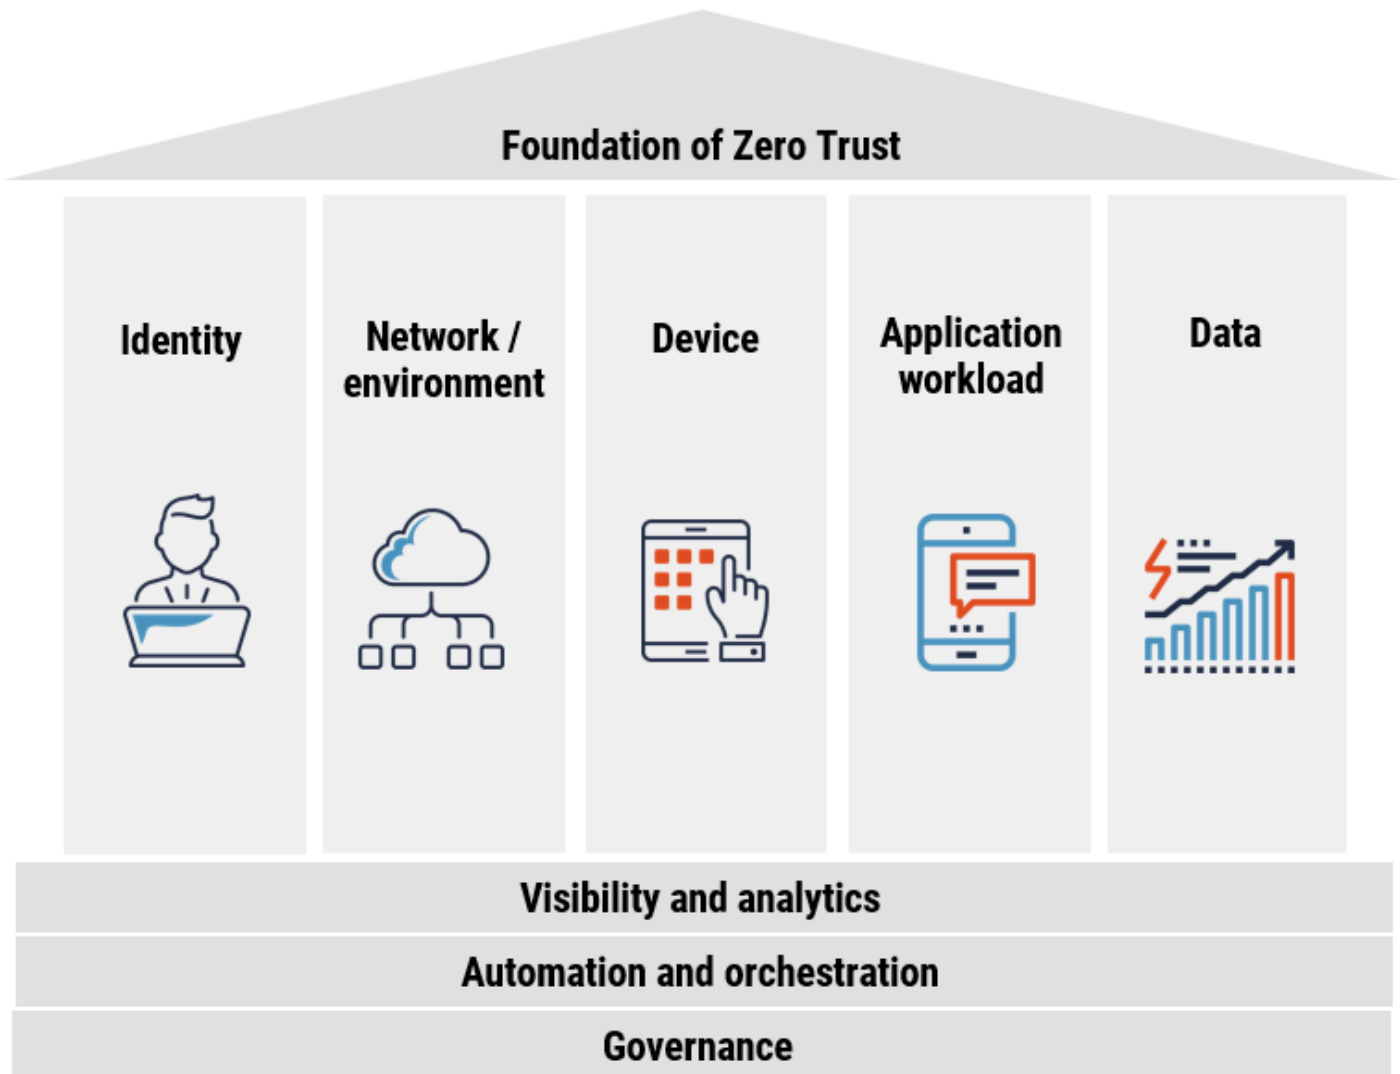
CISA's draft ZT Maturity Model was released in June 2021 in response to the Cybersecurity EO 14028. The model is comprised of five distinct pillars: identity, device, network, application workload, and data. Each pillar also includes general details regarding visibility and analytics, automation and orchestration, and governance. This maturity model is one of many roadmaps to support the transition to ZT. Within each pillar, the maturity model provides examples of a traditional,

advanced, and optimal ZT architecture. Federal agencies' goal should be to incrementally transition to a ZTA and eventually reach cyber security optimization.

Note that CISA will continue to review and modify its ZT maturity model to better align it with the agency's Continuous Diagnostics and Mitigation (CDM) program.

Below is a figure of CISA's depiction of the foundation of ZT and a high-level description of the five pillars.

Figure 1: CISA's depiction of the foundation of ZT



- 1. Identity:** Organizations need to implement technologies that continuously assess identity to then provide or deny access. Identity verification and authentication should be based on the access request and, if access is granted, it should be time bound and limited to the specific resource requested.

2. **Device:** The term “device” includes any hardware asset that can connect to a network, including Internet of things (IoT) devices, mobile phones, laptops, servers, and others. An inventory of all connected devices should be logged, and continuous compliance monitoring and validation of devices’ security posture should be conducted.
 - a. It’s critical to not only ensure the integrity of these devices but also that of the users operating them.
3. **Network environment:** A network environment is an open communication system that ties users together and allows for data sharing. These communication channels, both internal and external, need to be controlled, segmented, and protected based on their unique requirements. Organization’s need to avoid traditional perimeter-defence models, like firewalls, and focus on access control, continuous authentication, encryption, and risk assessment.
4. **Application workload:** An application workload consists of agency systems, computer programs, and services that are executed both on-premises and in the cloud environment. All applications should undergo rigorous empirical testing during the development and deployment process to ensure threat protection. The principle of least privilege (PoLP) should be enforced throughout the application’s lifecycle.
5. **Data:** Data across devices, networks, applications, and cloud must be protected against threats. Inventory of this data should be taken continuously with robust tagging, categorizing, and tracking.

To facilitate transitioning to ZT across the five different pillars, CISA outlined a ZT Maturity Model gradient using three stages, with increasing levels of commitment. CISA provided the following descriptions for each stage to identify the maturity for each ZT technology pillar and to provide consistency across the maturity model:

- **Traditional:** Manual configurations and assignment of attributes, static security policies, pillar-level solutions with coarse dependencies on external systems, least-function established at provisioning, proprietary and inflexible pillars of policy enforcement, manual incident response and mitigation deployment.
- **Advanced:** Some cross-pillar coordination, centralized visibility, centralized identity control, policy enforcement based on cross-pillar inputs and outputs, some incident response to pre-defined mitigations, increased detail in dependencies with external systems, some least-privilege changes based on posture assessments.
- **Optimal:** Fully automated assigning of attributes to assets and resources, dynamic policies based on automated and observed triggers, assets have self-enumerating dependencies for dynamic least-privilege access (within thresholds), alignment with open standards for cross-pillar interoperability, centralized visibility with historian functionality for point-in-time recollection of state.

Table 1 illustrates a high-level view of the ZT Maturity Model across each maturity stages five pillars.

Table 1: CISA's high-level ZT maturity model

Maturity level	Identity	Device	Network/ Environment	Application workload	Data
Traditional	<ul style="list-style-type: none"> • Password or multi-factor authentication (MFA) • Limited risk assessment 	<ul style="list-style-type: none"> • Limited visibility into compliance • Simple inventory 	<ul style="list-style-type: none"> • Large macro-segmentation • Minimal internal or external traffic encryption 	<ul style="list-style-type: none"> • Access based on local authorization • Minimal integration with workflow • Some cloud accessibility 	<ul style="list-style-type: none"> • Not well inventoried • Static control • Unencrypted
Advanced	<ul style="list-style-type: none"> • MFA • Some identity federation with cloud and on-premises systems 	<ul style="list-style-type: none"> • Compliance enforcement employed • Data access depends on device posture on first access 	<ul style="list-style-type: none"> • Defined by ingress/egress micro-perimeters • Basic analytics 	<ul style="list-style-type: none"> • Access based on centralized authentication • Basic integration into application workflow 	<ul style="list-style-type: none"> • Least privilege controls • Data stored in cloud or remote environments are encrypted at rest
Optimal	<ul style="list-style-type: none"> • Continuous validation • Real time machine learning analysis 	<ul style="list-style-type: none"> • Constant device security monitor and validation • Data access depends on real-time risk analytics 	<ul style="list-style-type: none"> • Fully distributed ingress/egress micro-perimeters • Machine learning-based threat protection • All traffic is encrypted 	<ul style="list-style-type: none"> • Access is authorized continuously • Strong integration into application workflow 	<ul style="list-style-type: none"> • Dynamic support • All data is encrypted

6.3 NCSC: Zero trust architecture design principles

This guidance was developed to help IT professionals and cyber security leads design and review a ZT architecture that meets their business requirements. As described by the NCSC, "zero trust is an architectural approach where inherent trust in the network is removed, the network is assumed hostile, and each request is verified based on an access policy".

Here are the eight principles outlines in the UK ZT guidance document.

1. Know your architecture, including users, devices, services, and data

The first step in developing a ZT architecture is to identify all your assets and each component of your architecture including your users, devices, and the services and data that they are accessing. This will allow you to identify where your key resources are, where your data is being stored and their level of sensitivity. Knowing this information will help you develop effective and appropriate access policies that will protect your valuable resources.

2. Know your user, service, and device identities

Identify the identity of all users, services, or devices in your architecture. This is important when deciding who or what will be granted access to data or resources.

3. Assess your user behaviour, devices, and services health

It's imperative to continuously monitor the behaviour of services or devices to identify anomalies. This is an important indicator of their security state. Using the appropriate security tools to measure user behaviour, device and service health is key in a ZT architecture. It helps evaluate the confidence in their trustworthiness.

4. Use policies to authorize requests

Each request for data or services should be authorized against a policy. The power of a ZT architecture comes from the access policies you define.

5. Authenticate and authorize everywhere

Authentication and authorization decisions should consider multiple signals, such as device location, device health, user identity and status to evaluate the risk associated with the access request. You should always assume the network is hostile and want to ensure all connections that access your data or services are authenticated and authorized.

6. Focus your monitoring on users, devices, and services

In a ZT architecture, monitoring is focused on users' behaviour, devices and services and will help you to establish their cyber health. You should know what actions devices, users and services are performing and what data they are accessing. Your monitoring should link back to the policies you have set, verifying they are being enforced as you expect.

7. Don't trust any network, including your own

Don't trust any network between the device and the service it's accessing, including the local network. Communications over a network, to access data or services, should use a secure transport protocol to gain assurance that your traffic is protected in transit and less susceptible to threats.

8. Choose services designed for ZT

Services may not support ZT and thus may require additional resources to integrate and increase support overhead. In these scenarios it may be prudent to consider alternative products and services that have been designed with ZT in mind.

7 Summary

With an ever-changing technology landscape and more sophisticated and persistent cyber threats, transitioning to a ZT security architecture has never been more important. It's important to recognize that the term ZT does not apply to a single product, technology, or architecture layer. Rather, it represents a security architecture and design philosophy whose central tenet is that no subject (application, user, device) in an information system is trusted by default. Trust is reassessed every time a subject requests access to a resource. The degree of access provided is dynamically adjusted based on the level of trust established with the subject.

Migrating to ZT is a complex process that requires changes and efforts to be made at all levels within an organization and a joint commitment from all. Although intricate, this shift will allow organizations to significantly improve their cyber security posture.

This document provided an overview of the ZT concepts and components. It will help organizations better understand what a ZT security model looks like and the associated benefits and challenges. We listed a few recommended best practices and principles to follow when implementing ZT. We provided an overview of three commonly cited ZT frameworks and guidelines to help organizations select the one that is most suitable for their business requirements, network infrastructure, and threat landscape.

8 Supporting content

8.1 List of abbreviations

Term	Definition
ABAC	Attribute-Based Access Control
CDM	Continuous Diagnostics and Mitigation
CISA	Cybersecurity and Infrastructure Security Agency
CSP	Cloud Service Providers
BYOD	Bring Your Own Devices
DoD	Department of Defense
DoDAF	Department of Defense Architecture Framework
EO	Executive Order
GC	Government of Canada
IC	Integrated Circuit
IoT	internet of things
IP	Internet Protocol
IT	Information Technology
MFA	Multi-factor Authentication
NCSC	National Cyber Security Center
NIST	National Institute of Standards and Technology
NSA	National Security Agency
PAM	Privileged Access Management
PAW	Privilege Access Workstation
PBAC	Policy-Based Access Controls
PC	Personal Computer
PoLP	Principle of Least Privilege
RBAC	Role-Based Access Control
SAW	Secure Administrative Workstation
SDP	Software-Defined Perimeter
SIEM	Security Information and Event Management
TBS	Treasury Board Secretariat
TPM	Trusted Platform Module
UK	United Kingdom
US	United States
VLAN	Virtual Local Area Network
VPN	Virtual Private Networks
ZT	Zero Trust
ZTA	Zero Trust Architecture

8.2 Glossary

Term	Definition
Access control	Certifying that only authorized access is given to assets (both physical and electronic). For physical assets, access control may be required for a facility or restricted area (e.g., screening visitors and materials at entry points, escorting visitors). For IT assets, access controls may be required for networks, systems, and information (e.g., restricting users on specific systems, limiting account privileges).
Administrative privileges	The permissions that allow a user to perform certain functions on a system or network, such as installing software and changing configuration settings.
Authentication	A process or measure used to verify a user's identity.
Authorization	Access privileges granted to a user, program, or process.
Availability	The ability for the right people to access the right information or systems when needed. Availability is applied to information assets, software, and hardware (infrastructure and its components). Implied in its definition is that availability includes the protection of assets from unauthorized access and compromise.
Compromise	The intentional or unintentional disclosure of information, which adversely impacts its confidentiality, integrity, or availability.
Confidentiality	The ability to protect sensitive information from being accessed by unauthorized people.
Cloud computing	The use of remote servers hosted on the Internet. Cloud computing allows users to access a shared pool of computing resources (such as networks, servers, applications, or services) on demand and from anywhere. Users access these resources via a computer network instead of storing and maintaining all resources on their local computer.
Cyber attack	The use of electronic means to interrupt, manipulate, destroy, or gain unauthorized access to a computer system, network, or device.
Cyber threats	A threat actor, using the Internet, who takes advantage of a known vulnerability in a product for the purposes of exploiting a network and the information the network carries.
Detection	The monitoring and analyzing of system events in order to identify unauthorized attempts to access system resources.
Encryption	Converting information from one form to another to hide its content and prevent unauthorized access.
Firewall	A security barrier placed between two networks that controls the amount and kinds of traffic that may pass between the two. This protects local system resources from being accessed from the outside.
Integrity	The ability to protect information from being modified or deleted unintentionally or when it's not supposed to be. Integrity helps determine that information is what it claims to be. Integrity also applies to business processes, software application logic, hardware, and personnel.
Intellectual property	Legal rights that result from intellectual activity in the industrial, scientific, literary, and artistic fields. Examples of types of intellectual property include an author's copyright, trademark, and patents.

Term	Definition
Internet of things	The network of everyday web-enabled devices that can connect and exchange information between each other.
IT asset	The components of an information system, including business applications, data, hardware, and software.
Least privilege	The principle of giving an individual only the set of privileges that are essential to performing authorized tasks. This principle limits the damage that can result from the accidental, incorrect, or unauthorized use of an information system.
Malware	Malicious software designed to infiltrate or damage a computer system, without the owner's consent. Common forms of malware include computer viruses, worms, trojans, spyware, and adware.
Multi-factor authentication	Authentication is validated by using a combination of two or more different factors including: something you know (a password), something you have (a physical token), or something you are (a biometric).
Network security zone	A networking environment with a well-defined boundary, a network security zone authority, and a standard level of weakness to network threats. Types of zones are distinguished by security requirements for interfaces, traffic control, data protection, host configuration control, and network configuration control.
Perimeter	The boundary between two network security zones through which traffic is routed.
Security control	A management, operational, or technical high-level security requirement needed for an information system to protect the confidentiality, integrity, and availability of its IT assets. Security controls can be applied by using a variety of security solutions that can include security products, security policies, security practices, and security procedures.

8.3 References

Number	Reference
1	The National Institute of Standards and Technology. Special Publication 800-207: Zero Trust Architecture , August 2020.
2	Cybersecurity and Infrastructure Security Agency. Zero Trust Maturity Model , June 2021
3	National Cyber Security Center: Zero trust architecture design principles , July 2021
4	Canadian Centre for Cyber Security. ITSG-33 IT Security Risk Management: A Lifecycle Approach . December 2014.
5	Canadian Centre for Cyber Security. Secure your accounts and devices with multi-factor authentication (ITSAP.30.030) , June 2020
6	Canadian Centre for Cyber Security. User authentication guidance for information technology systems (ITSP.30.031 v3), April 2018
7	Treasury Board Secretariat. Guideline on defining authentication requirements , November 2012
8	Executive Office of the President. Executive Order (EO) 14028 to Improve the Nation's Cybersecurity , May 2021