Zero Trust for maximum security

A "deny all, allow some" strategy to protect your IT assets across internal networks, clouds and remote working environments

Table of contents

Introduction 2

How implementing Zero Trust protects your enterprise 3

Four key considerations 3

Zero Trust enforcement layers 5

Plan ahead before starting 6

Benefits and challenges by industry 6

Conclusion 7

Introduction

Information security protection has come a long way in the last two decades. It started out with organizations designating their internal network as the trusted zone and everything outside their network as the untrusted zone. They used firewalls to restrict access from all untrusted zones into the trusted zone and gave assets within the trusted zone unfettered access to one another.

Today, in the age of offsite data centers, public cloud environment and the need for remote access from any device and any location, the security perimeter is much more fluid and harder to manage. Security professionals struggle to ensure that access controls remain up to date and relevant to the ever-changing business.

Over the past 5 years, network access controls have started to crack. Adversaries are concentrating on sophisticated exploitation techniques such as those described in the MITRE ATT&CK framework[1], and trusted networks are being infiltrated by attackers that phish employees, use default credentials and exploit unpatched systems.

These vulnerabilities were heightened in the wake of the COVID-19 pandemic, with the sudden shift to remote working and a dramatic rise in bringing your own device and remote access to core systems and sensitive data. Workplace restrictions imposed by governments worldwide have forced most enterprises to rethink traditional approaches to deploying applications and enable new ways of working and collaborating.

The fear is that through this rapid change adversaries will compromise trusted networks and gain unfettered access to the internal environment. Firewalls and endpoint controls are effective to a point, but when a trusted system or an administrative credential has been compromised, most security controls will not be able to discern an attack taking place through these compromised assets.

This is why Zero Trust is so relevant today. Conceived in 2010 by Forrester researcher John Kindervag[2], the model assumes that everything around a network asset is hostile, including network assets from the trusted zone. All access to the network asset is, by default, not trusted. Access is kept to a minimum and permitted only based on certain policies. Ten years later, Zero Trust has massively evolved, and the operating model can be used to protect not only infrastructure assets, but also user credentials and data.

Fortunately, the evolution occurred just as cloud computing is becoming a mainstream platform running organizations' services, and the traditional concept of network access control is completely breaking down. The rising use of serverless technologies as well as containers has fundamentally changed the concept of Zero Trust in a digital environment.

---

[1] MITRE.org, "Enterprise Matrix," https://attack.mitre.org/matrices/enterprise/, 2019.

[2] Forrester.com, "No More Chewy Centers: The Zero Trust Model of Information Security," https://www.forrester.com/report/No+More+Chewy+Centers+The+Zero+Trust+Model+Of+Information+Security/-/ERES56682, March 23, 2016.

Zero Trust is extremely effective in reducing security incidents, as it implements the “deny all, allow some” least-privilege principle even within a trusted environment.

How implementing Zero Trust protects your enterprise

How does Zero Trust architecture help to protect an asset? Traditionally, systems in a trusted zone do not have any restrictions in accessing another system in the same zone, such as employees in the same office who are all connected to the same network. However, if an employee’s PC is compromised, the adversary can make use of it to compromise other employees’ PCs. These other employees will not block such access, as it is assumed that the compromised PC is trusted since they all reside in the same network. If, instead, this network implements a Zero Trust policy that does not allow employee-to-employee PC access, each PC the compromised PC tries to access will ignore the access requests by default.

Zero Trust is extremely effective in reducing security incidents, as it implements the “deny all, allow some” least-privilege principle even within a trusted environment. Such operating behavior may appear to be detrimental to business, but Zero Trust actually improves business operations because it forces stakeholders to be aware of the transactions that need to take place to make a service work. Consequently, stakeholders are more aware of what controls need to be implemented to secure the service as well as to improve ongoing operations.

Four key considerations

Organizations need to be aware of four key activities when implementing Zero Trust in their environment:

1. **Know what access is required to which asset.** The first requirement should be a standard operating behavior applied to all new assets in an environment, including devices, services, accounts and data. For example, a database server can only be accessed via a certain application with specific login credentials through specific network ports; documents stored on a file server can only be accessed by specific users from a specific network. Understanding how the asset is to be accessed is key to the following activities.

2. **Establish a unified identity-centric framework.** The Zero Trust model requires a unified identity-centric framework to be in place to perform the required verification. The identity-centric framework is not only used to authenticate human users, but also to authenticate communication among entities (applications, devices, network services), which is prevalent in an automated environment. Methods of authentication could be through passwords, API keys or public/private keys. This framework is crucial to granting access to verified users or entities, understanding the context of the access and determining the correct policy to apply to the access. The unified aspect of the framework is essential, as many organizations still run a fragmented identity framework in different parts of their business. This may result in inconsistent policies being applied to the same user/entity.

3. **Define a robust policy based on the Kipling method[3].** Even when a user or entity is verified after being properly authenticated, it does not mean that this person/entity will always have complete access to an asset. A strong Zero Trust environment enforces good policies to achieve the least-privilege outcome. The policy should adopt the Kipling method and subject all access to the following six questions:

   - What is the asset being accessed?
   - Who is accessing it?
   - Why is the requestor accessing the asset?
   - How is the requestor accessing the asset?
   - Where is requestor accessing the asset from?
   - When is the requestor accessing the asset?

Even if the user/entity is verified to have the necessary credentials to access the asset, if the policy determines that out of the six questions, the condition for access is not met for even one of the questions, the access will be denied. For example, a finance controller may log into the file server to access sensitive financial reports from the office network, and such access normally will be permitted during office hours.

If the file access occurs during a Saturday night from a remote network connection, the policy will prevent the file from being accessed even if the right credential is provided. If developed correctly, good policies can detect anomalous activities that indicate malicious behavior, and this is how Zero Trust helps to proactively reduce security incidents.

4. **Monitor and optimize all approved access and react quickly to unapproved access.** Lastly, in an effective Zero Trust environment, a security operations center (SOC) must continuously monitor all access, approved or otherwise. Any unapproved access indicates that an anomaly has occurred, and it will require the cybersecurity team to investigate to determine whether there is an active attempt to compromise the environment. For all approved access, it is important to verify that the organization is constantly reviewing policies to ensure that access is always aligned to business requirements.

3 Palo Alto Networks, "All Layers are Not Created Equal," https://blog.paloaltonetworks.com/2019/05/network-layers-not-created-equal/, 2019.

A phased approach allows for a manageable transformation from a legacy “Trust but verify” operation to a “Do not trust and verify” outcome without adversely disrupting other parts of the business.

Zero Trust enforcement layers

From DXC Technology’s perspective, organizations should consider enforcing security best practices for Zero Trust at three different layers:

Infrastructure (network, cloud, containers). The origin of Zero Trust starts at the network level, where it is used to restrict access between endpoints in the same network environment. This method of restricting lateral access is known as microsegmentation and is extremely useful for protecting endpoints against unauthorized access even if the endpoints (such as containers and internet of things devices) do not have built-in security controls to protect themselves against external attacks. Microsegmentation, if done correctly, can help provide an additional level of network access control should the firewalls be compromised or circumvented, or when the use of firewalls may not be feasible (in cloud or virtualized environments).

Identity (users, applications, services). Just because a user or an entity has the privilege of accessing a system or service does not mean they can always have unrestricted access or possess super privileges. Even if an administrator has root privileges to a system, that administrator should not be entrusted to perform privileged operations unless the activity conforms to approved policies. Until then, access to the system being administered should be kept to a minimum for day-to-day operations. In this case, privileged access management (PAM) is a key component for enforcing such least-privilege access.

Data. As organizations embark on their digital transformation and IT modernization journey and move their business services to the cloud, there is less infrastructure to manage, and the concept of restricting network access becomes less relevant. In the cloud context, the focus is to protect access to the data, potentially from any network around the world. Whether the data is managed through a serverless function or through a cloud workload, policies should be enforced to ensure that the data is accessed only under preapproved conditions, and only to the minimum information required for the purpose.

One relevant use case for implementing Zero Trust on data is when an organization needs to perform analytics on its user transactions. Due to regulatory requirements, organizations avoid processing user-identifiable information, as the accidental disclosure of such information can result in a financial and reputational impact. As such, if an organization needs to access a data lake for business analytics, the data should be pseudo-anonymized before it is sent to the analytics team for data processing.

Plan ahead before starting

The value of Zero Trust to help organizations protect their environment is undeniably high, but the journey toward Zero Trust is not straightforward. Organizations need to plan for several considerations before they start working toward a Zero Trust outcome.

When adopting Zero Trust, consider implementing the Zero Trust framework at the design phase. Typically, Zero Trust projects have a higher success rate when implemented in a greenfield environment, as there are no existing processes to disrupt. Retrofitting a Zero Trust framework after the infrastructure, application or service is in production will require careful planning and implementation in stages to minimize disruption.

If developed correctly, good policies can detect anomalous activities that indicate malicious behavior, and this is how Zero Trust helps to proactively reduce security incidents.

For most organizations embarking on Zero Trust for the first time, the biggest challenge is knowing how the service or environment is being accessed and by whom. It will take some time to build an effective policy to enforce Zero Trust, and thus enough leeway must be provisioned to ensure that administrators for Zero Trust projects have time to understand the transactional behavior, build the policies correctly and test the policies sufficiently before enforcing them.

Zero Trust should be implemented in phases rather than as a big bang. A phased approach allows for a manageable transformation from a legacy "Trust but verify" operation to a "Do not trust and verify" outcome without adversely disrupting other parts of the business.

Finally, the biggest impediment to a successful Zero Trust outcome is culture. The concept of treating everything as hostile is a notion that will require all employees to take time to understand and adopt. Zero Trust design is a fundamental change from traditional security architecture where trust among internal users is implicit. In a Zero Trust environment, the trust needs to be explicitly built and verified. It is important that stakeholders fully appreciate the value and challenges of Zero Trust before they start their Zero Trust transformation journey.

Benefits and challenges by industry

While there are common challenges for all organizations in implementing Zero Trust, there are also particular benefits for specific industries.

Healthcare. Most healthcare organizations do not have complex network environments due to budgetary constraints, which prevents implementation of a multizone, multilayer environment commonly prescribed in best practices. Zero Trust works quickly on flat networks — the protection benefits can be realized immediately since there is less complexity to contend with.

Manufacturing, utilities and other industrial environments. Industries dealing with operational technology (OT) security will benefit from a Zero Trust implementation because their operating environment is highly predictable and has very little human intervention. Such transactional behavior allows for the rapid development of robust policies that can be enforced quickly.

Banking and capital markets, insurance, retail. The common theme among these industries is that they all store a lot of personal identifiable information (PII) as part of their business requirements, and thus are subjected to global and local privacy laws. A Zero Trust framework providing bare minimum access to data for processing or analytics purposes helps these industries reduce the likelihood of accidental or intentional data disclosure.

Conclusion

The threat landscape has changed massively in the past two decades, most recently during the spread of COVID-19. Cybercriminals and nation-state groups increased their attacks during the pandemic and are showing no signs of slowing down. With the blurring of assets between on-premises and the cloud in the new digital world, organizations need a different way to protect their data and assets.

Zero Trust is an effective and proven method that can reduce security breaches through its proactive approach of restricting access even for known users and entities. Implementing Zero Trust requires organizations to change the way they design and build their applications and services, and will require stakeholders to understand its benefits and challenges before they embrace it. Once organizations have crossed those hurdles, they will enjoy the benefits of running a secure business for a long time to come.

About the author

TM Ching is chiefly responsible for security thought leadership as well as research and development activities for DXC Security worldwide. He works closely with vendors and professional bodies to identify technological evolutions or disruptions on the horizon, and develops roadmaps for both clients and DXC Security to achieve service readiness to meet the threat landscape changes in the next 12 to 36 months.

Learn more at
dxc.com/security