

The what, why and how of Zero Trust Network Architecture



Table of contents

3

—

4

—

5

—

10

12

—

13

—

18

21

—

22

—

24

—

25

—

26

—

28

—

32

36

—

37

—

39

41

—

42

What is Zero Trust?

Let's take a brief look at what Zero Trust means. Explore the origin and evolution of this term, and the basic tenets of a Zero Trust security model.


The last few years have introduced many changes in the way organizations work and the discourse around IT. Remote work, hybrid workplaces, and cloud solutions have increased in relevance and adoption. But these aren't the only terms to have gained popularity in recent times.

If you work in IT security, or follow IT developments closely, chances are you've come across the term Zero Trust. While the term itself isn't new, it has gained increased prominence in the past few years.

Put simply, the term refers to a security model where nothing and no one is trusted implicitly, and every request for access to an organization's resources is verified—irrespective of where it originated.

In other words: all users and devices are treated as outsiders until their identity and security status is verified.

Before we drill deeper into what Zero Trust is, and what this means for your organization, let's take a quick glimpse at how the term came about.



The origin and evolution of Zero Trust



The first discussion about Zero Trust could be said to have been back in 2004 with the formation of the Jericho Forum™. This forum highlighted the challenges in defining a perimeter for an organization's IT systems. It then defined the trend then known as ["de-perimeterization"](#) and found solutions for it. This culminated in the formation of the [Jericho Forum Commandments](#).

In 2009, John Kindervag, then an analyst at Forrester Research, coined the term "Zero Trust". The original idea of Zero Trust was network-centric. It aimed to use microsegmentation to enforce more granular access controls, and limit lateral movement. Then, in 2014, Google published its BeyondCorp research. In this model, the concept of a privileged corporate network was removed. Instead, access was given solely based on device and user credentials.

Today, the idea of Zero Trust has expanded into a holistic security model that moves beyond just network segmentation or identity based security. It encompasses endpoint security, data security, SIEM, threat intelligence, and more.

Together, these elements come together to form what some might call the Zero Trust ecosystem (or the "Zero Trust eXtended ecosystem" as Forrester calls it).

But what is Zero Trust, really?



Now we know how the term originated and evolved over the years. But what is Zero Trust? Chances are if you search for Zero Trust solutions, you'll come across a plethora of definitions, alongside a bunch of "Zero Trust products". The last of which is a bit misleading.

Zero Trust isn't a product or a service. It's a new strategy for approaching organizational security built on the principle of "Never trust, always verify."

To understand this, let's first explore the current perimeter-focused style of security.

Also known as the castle-and-moat approach, the perimeter-focused model relies on strong firewalls and VPNs to regulate access to company resources. Anyone outside of the network is "the bad guy" and can't access internal resources, while anyone on the inside is "one of the good guys" and has access to, well, almost anything.

The only issue with this is that if an attacker manages to get through the firewall or gains access to a legitimate user's account, they get access to everything. Plus, there is always a risk of insider threats—legitimate users gone rogue.

The Zero Trust approach to security seeks to solve this by focusing defenses on identities, assets, and resources. The goal is to prevent unauthorized access to organizational resources, while making access control as granular as possible.

It does this by enforcing the following:



Explicit verification

Before granting access, each request is explicitly verified, irrespective of where it originated from, based on as much information as can be gathered (device location, device health, user identity, network information, etc.).



The principle of least privilege

Every user only receives the bare minimum level of access required to carry out their jobs.



Just-in-time privilege elevation

Users needing to access sensitive resources (or resources they'd not typically have access to) are given access in a just-in-time manner (when they need



Continuous monitoring

Users' activities are monitored continuously, allowing security teams to detect and act against suspicious behavior right away.



Security automation

Security responses are automated to ensure that action can be taken right away, before serious damage occurs.



Dynamic access

A user's access can be limited or rescinded based on the recommendation from the security and monitoring tools.

This ensures that even if the “bad guys” get inside your network, they won't be able to do much damage.

Infiltrating using a low-privilege account limits their access, and prevents them from moving laterally through the organization's network. On the other hand, logging into a privileged account from an unknown location, at an unusual time, from an untrusted device, and carrying out unusual activities raises the alert and limits or shuts down access.

Another key principle of a Zero Trust security implementation is to “assume breach” (or “limit the blast radius” as some phrase it).

You need to assume that attackers are already inside your network, and work to limit their access and the potential fallout from their actions. This can be accomplished by reducing the size of the implicit trust zone within your network.

We’ll discuss more about this in the following sections.

An abstract view of Zero Trust

In a perimeter-based security model, access to the corporate network, directly or via a VPN, grants users access to most, if not all, internal apps and services. This is what is known as the castle-and-moat model. We could also call it the shellfish model, for those zoologically or gastronomically inclined. If an attacker makes it past the hard outer shell, a tasty treat awaits within.





Figure 1: Perimeter-based "castle-and-moat" network access model with VPN.

In a Zero Trust model, on the other hand, there is no crunchy shell to break through.

If you take an abstract view of Zero Trust access, then a network is comprised of two parts: the untrusted zone, which includes the corporate network, and the trusted zone, where one or more resources reside. In between the two lies the Zero Trust “perimeter” where access is controlled through a policy decision point (PDP) and its corresponding policy enforcement point (PEP).



The PDP is responsible for authenticating a user’s identity, and verifying whether they have the authorization needed to access a specific resource.


The PEP, on the other hand, enables, monitors, and eventually terminates connections between users and enterprise resources.

The Implicit Trust Zone is the area where all entities are trusted to the level of the last PDP/PEP gateway they’ve crossed. Think of it as a movie theater—the ticket checkers are the PDP/PEP, and anyone who makes it past them is

trusted to have bought a ticket to see the show.

Once the audience is in and seated, there's nothing the PDP/PEP can do to further control their access. Thus, to ensure that your access controls are as granular as needed, you need to ensure that the PDP/PEP is as close to the resources—or movie theater—in question. This helps minimize the blast radius if an attacker does manage to get past a specific PDP/PEP.

Zero Trust provides a set of principles that focus on this.



The 7 tenets of Zero Trust

In their special publication on Zero Trust ([NIST SP 800-207](#)), NIST outlines seven basic tenets that should be included in any Zero Trust architecture. While these tenets represent an ideal situation, it's not necessary, or even feasible, to implement them fully for your organization's Zero Trust strategy.

These tenets are:

1. All data sources and computing services are considered as resources:

Be it an IOT device, a SaaS product, or a personal device—if it's within the enterprise network, or can access enterprise owned data or services, it should be counted as a resource.

2. All communication is secured, regardless of network location: Never automatically trust a user or device based on their network location. Whether an access request comes via the enterprise network or from a non-enterprise network, it must meet the same security requirements.

3. Access to individual enterprise resources is granted on a per-session basis: Evaluate the requester's trustworthiness before granting access every single time. Give them the least privileges needed to complete their tasks, and ensure that granting them access to one resource doesn't automatically grant them access to others.

4. Access to resources is determined by dynamic policy: Define your list of resources, your organization members, and who has what level of access to which resources. Take into account context, and vary authentication methods based on the sensitivity of the resource being accessed. The stringency of reauthentication should be proportional to the sensitivity of the resource.

5. Monitor and measure the integrity and security posture of all owned and associated assets: Monitor and evaluate the security posture of all devices accessing your network—this includes known vulnerabilities, patch status, and other cybersecurity risks—before granting it access. Devices that are unmanaged, new to the network, or are managed but are known to have vulnerabilities, should be treated differently from managed devices that are known to be secure.

6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed: Ensure you're continuously checking for threats and reassessing the trust given to users accessing your resources. Use identity and access management solutions to control access and enforce multi-factor authentication (MFA) for access to enterprise resources. Make sure to continuously monitor user activities, reauthenticating and reauthorizing them when needed.

7. Collect as much information as possible about the current state of assets, network infrastructure and communications, and use it to improve security posture: Collect and analyze data about device security postures, network traffic, access requests, etc., to improve your organization's security posture. This data can also provide additional context for the requirements under tenet 4.

A peek under the hood: The components of a Zero Trust ecosystem

In the previous section, we outlined how a Zero Trust network architecture works, and discussed the seven tenets that form the foundation of a Zero Trust implementation.

Now, let's go over the various components of a Zero Trust ecosystem and see what they do.

The logical components of a Zero Trust network access model (as per NIST)


The logical components of a Zero Trust network access model (as per NIST)



A Zero Trust security model is made up of several logical components. These components, whether deployed on-premises or in the cloud, play a vital role in ensuring the security of your network. In fact, some of these components might already be part of your existing security framework.

Before we dive into the details of the figure below, some things to remember:

- Figure 3 is a conceptual framework that shows the basic relationships between these components.
- The model in Figure 3 represents an ideal situation as per NIST SP 800-207, your actual implementation model might not match this 100%.
- All the logical components communicate exclusively over a separate plane called the control plane.
- All application data is communicated over a data plane.
- The PDP from Figure 2 has been broken down into a Policy Engine and Policy Administrator in Figure 3.



Let's start with the three core components:

- **Policy engine (PE):** The brains of this model, the PE handles granting or withholding access to a given resource. It uses organization policies and inputs from other components like the continuous diagnostics and mitigation (CDM) system, SIEM, threat intelligence services, and more to make its decision.
Note: The PE only makes and logs this decision. The execution of its decision falls to the policy administrator.
- **Policy administrator (PA):** The implementer for the PE's decision-maker, the PA is responsible for establishing or severing the connection between identities (users, applications, machines) and the resources they're trying to access. It generates the session-specific authentication tokens and credentials required to access any enterprise resource.

The PA is closely tied to the PE, and depends on its decision to make or break a connection. If a session is authorized and authenticated, the PA commands the PEP to start the session. If the connection request is denied or a previously approval is countermanded, the PA signals the PEP to break the connection.

While this model shows the PA and PE as separate entities, they can also be treated as a single service.

- **Policy enforcement point (PEP):** The PEP is responsible for enabling, monitoring, and terminating connections between identities and the resources they're trying to access. It forwards connection requests to the PA and then creates, denies, or breaks connections with resources as per the PA's commands.

In some implementations, the PEP might comprise of two components: a client-side "agent" and a resource-side "gateway". Others might have just a single portal that acts as a gatekeeper for all communications.

Beyond the PEP is the implicit trust zone wherein lie the organization resources.

As mentioned above, the PE's decisions are informed by data and rules from several other sources. These data sources could be internal or external. They can include:

- **Continuous diagnostics and mitigation (CDM) system:** This system gathers information about the state of all enterprise assets. It also applies updates to asset configurations or software. The CDM provides the PE with information about the asset making the request—such as its OS version, the software it's running, any known vulnerabilities, and more. It is also responsible for identifying and possibly enforcing a subset of policies on non-enterprise devices active within the organization's network.
- **Industry compliance system:** This system ensures that the organization complies with industry-specific regulatory requirements (HIPAA, FISMA, PCI DSS, etc.). It includes any policies and rules that have been created to ensure compliance.


- **Threat intelligence feeds:** These feeds provide the PE with information about newly discovered attacks and vulnerabilities to help it make better decisions. They also include information about new software flaws, malware, and even reports on attacks on other enterprise assets—the lattermost being especially important as the PE will try to deny any compromised assets access to organization resources.
- **Activity logs:** This system, sometimes a part of the SIEM or security orchestration automation and response (SOAR) system, aggregates logs about all activities within the network. This includes network traffic, asset logs, resource accesses, file modifications, and other events. This information provides real-time or near-real-time information on the security posture of an organization's IT assets.
- **Data access policies:** These are the rules and policies that govern access to enterprise resources. They provide information about access privileges for accounts and applications and services, making them the starting point for all access authorization. These policies need to be drafted carefully, keeping organization needs in mind.
- **Enterprise public key infrastructure (PKI):** This system creates and logs certificates issued by the enterprise to resources, systems, and applications. This also includes external PKIs, such as the global certificate authority ecosystem.
- **ID management system:** This system creates, stores, and manages all user accounts and identities. It contains information about any identity in the organization—name, email address, role, assigned assets, access attributes, etc. For the sake of collaboration, this system might sometimes include non-enterprise employees and assets.

- **Security information and event management (SIEM) system:** The SIEM system collects and analyzes logs and other security-centric information to detect threats inside an organization. This can be used to refine organization policies, and warn the security team about possible attacks against enterprise assets.

All of these components together inform the PE's decisions. A misconfiguration in any of these components or the lack of them could hamper the PE's ability to accurately regulate access.

Before we move on, here's an important point to keep in mind: all of the components listed above are logical components. They don't have to be individual systems. For example, your SIEM solution could do the job of the SIEM, activity logs, and threat intelligence feed components.

Similarly, one component might actually be comprised of multiple software or hardware components. For example, the industry compliance system might be comprised of several policies, as well as multiple software tools (like SIEM, endpoint management, and data security tools) and hardware assets.



The Zero Trust eXtended ecosystem



As discussed in the introduction to this e-book, and evidenced by the laundry list of components in the previous section, Zero Trust isn't just about network segmentation or identity and access management. It's a holistic security approach.

To emphasize this, Forrester has defined what it calls the Zero Trust eXtended (ZTX) ecosystem.

A complete implementation of the ZTX ecosystem requires processes and technology capabilities in the following:

1. Data security: Ensuring the confidentiality, integrity and accessibility of data is the goal of any cybersecurity approach. Thus it's no surprise that even the ZTX ecosystem requires organizations to ensure the security of data both at rest and in transit. This element of the ZTX ecosystem differs a little from the data access policies component of the model in Figure 3. Access policies govern who can access the data and resource in question. This element requires organizations to be able to track, classify, manage, and secure their data at all stages.

2. People and identity governance: A key part of a Zero Trust security model is limiting user privileges and securing all user sessions. It's important

to ensure that only the right people get access to the right resources in the right context. This could include monitoring user sessions for any suspicious activity, and terminating them as needed. However, while the ZTX labels this element as “people” other organizations, such as the Identity Defined Security Alliance (IDSA), prefer to talk about identities—which includes both people and machine identities or processes.

3. Devices and endpoint security: The advent of IoT and even policies like bring your own device (BYOD) have expanded an organization’s attack surfaces. Take the example of the casino that was hacked through a smart thermometer. As part of their Zero Trust strategy, security teams need to be able to monitor, manage, and secure every device on their network at all times.

4. Network segmentation and security: The ability to segment and control the network is a key requirement for network segmentation oriented Zero Trust approaches. This can help reduce the risk of lateral movement within an organization’s network and limit attackers’ access in the event of a breach.


5. Workload security: Defined as two separate entities by the IDSA (computing and applications) workloads refer to the entire application stack—from the applications themselves and the task they’re trying to accomplish, to the computing resources associated with them. They are the front-end and back-end systems that help organizations run. As with any other component in an organization’s daily operations, workloads need to be treated as a potential source of threats and secured accordingly.

These elements are also referred to as the five pillars of Zero Trust in the pre-decisional draft of the [Zero Trust Maturity Model](#) from the Cybersecurity and Infrastructure Security Agency (CISA).

In addition to these five pillars, Forrester’s ZTX ecosystem includes two additional components.

1. Visibility and analytics: This element covers tools like SIEM, user entity and behavior analytics (UEBA), security user behavior analytics, and more. It stands to reason that you can't fight what you can't see. Having a visibility and analytics tool like a SIEM and UEBA solution can help organizations gain 360-degree visibility into their network, and detect potential or active threats.

2. Automation and orchestration: Having visibility into the network is of little use if the organization is unable to act on it. Since attackers don't stick to a strict nine-to-five schedule, this means organizations need 24x7 security. Tools with SOAR capabilities can help with this. These tools act as force multipliers for security teams. They reduce grunt work, shorten incidence response times, and help stop attacks.



Now that we've understood the various elements of the Zero Trust security model, let's take a look at various ways of implementing it.

Implementing Zero Trust: 3 approaches, and helping the PE “think”


As discussed in the introduction of this e-book, the earliest definitions of Zero Trust considered network microsegmentation the only way to achieve perimeter-less security. However, as the definition of Zero Trust has evolved over the years, new implementation strategies have emerged.

NIST SP 800-207 outlines three ways in which a Zero Trust architecture can be implemented:

- using enhanced identity governance
- using logical microsegmentation
- using network-based segmentation and software-defined perimeters

Each of these approaches might use different components and have a different primary source for the policy rules that power the PE. However, they all implement all seven tenets of Zero Trust.

Depending on the organization's use case, one approach might be more feasible than the others. Which is not to say the others won't work—they might just be more difficult to implement or require drastic changes in the organization's daily operations.



Zero Trust using enhanced identity governance



In this approach, the identities of people and machines are the key components in policy creation. Here, access policies are based on these identities and their assigned attributes.

Access authorization is based primarily on the access privileges and attributes assigned to the identity raising the request. Other factors, such as the requestor's location, device used, asset status, and other factors, also play a role in this process.

Depending on how these factors play out, the confidence or trust level assigned to a specific access request might change. They could also result in the request being denied altogether.

For example, if a privileged identity requests access to a sensitive resource from a trusted, secure device, from a known location, during their usual active hours, they'll be granted access without any questions.


However, if the same identity attempts to access this resource from a new location using an unmanaged/untrusted device they might be asked to fulfil additional authentication steps or have their access limited or denied.

In this approach, each individual resource or the PEPs protecting these resources need to:

- either have a way to forward requests to the organization's PE for authentication, or
- have a way to authenticate and approve these requests themselves before granting access

Identity-centric approaches work well in situations involving cloud services or platforms which do not allow the use of organization-owned and operated Zero Trust components. Here, the organization can use identities as the means of creating and enforcing access policies, even in the cloud.

It also works well in organizations that use an open network model or have networks which see frequent access from guest users or non-organizational devices. Here network access is open to all. However, organizational resources are only accessible to identities with the right privileges.



Zero Trust using microsegmentation

In this approach, organizations implement Zero Trust access by separating resources or groups of resources on to their own distinct networks, or network segments which are protected by a gateway security component.

This could be accomplished in two ways. The first is using hardware components like smart switches, next-generation firewalls, or other special purpose gateway devices to act as PEPs for each resource or group of resources.

The second is by implementing host-based microsegmentation with the help of software agents and firewalls on endpoint assets. This software-based approach can even be implemented alongside the use of hardware gateways.

This approach requires a strong identity and access management solution to be fully functional. However, it is finally up to the gateway components to act as the PEPs to prevent unauthorized access to resources.

For this approach to work, it's important to ensure that the PEP components:

- are managed by the organization, and
- are able to react and reconfigure as needed to respond to changes in workflows or threats

Zero Trust using network infrastructure and software defined perimeters (SDPs)

In this approach, organizations use their network infrastructure to implement Zero Trust architecture. This implementation could be achieved by using a virtual overlay network on top of your existing network structure. It is sometimes called the SDP approach as it includes concepts from software defined networks.

Just like in the other models, access requests are made via the PEPs which are managed by the PA. Once a request comes in, it is shared with the PA and PE. Then, based on the PE's decision (or a change in its earlier decision), the PA sets up or reconfigures an individual network connection between the requesting asset and the resource.

The PE can use identities, device status and health, location, and other details to inform its decision as to whether the PA should allow this connection.

For this approach to function well, the organization needs to have a strong identity governance solution. It also needs an endpoint management or CDM component in place to ensure the PE doesn't give access to the wrong user or device.

An argument in favor of an identity-based approach

There isn't a right or wrong way to implement Zero Trust security in your organization. However, some approaches reap results quicker than others.


Organizations that focus on the identity and device security components of the Zero Trust have been found to make rapid risk reductions, as detailed in [Forrester's "A Practical Guide To A Zero Trust Implementation"](#).

Implementing identity and access management (IAM) solutions can help organizations:

- **Defend against and mitigate the impact of credential-based attacks:** Credential-based attacks are the most common cyberthreats organizations face today. As per the [Verizon Data Breach Investigations Report 2022](#), nearly 50% of data breaches in 2021 involved credentials. Enforcing MFA adds additional layers of security which can mitigate the risk of these attacks.
- **Centralize your identity governance and access management:** IAM solutions centralize access management and can automate routine identity lifecycle related tasks. Automatic provisioning and deprovisioning of identities can help reduce risks posed by orphaned accounts. They also give IT admins complete insight into all organization identities.

- **Deliver a secure and frictionless experience to your end users:** The use of a mix of on-premises and cloud solutions can result in fragmented identities. It also adds a lot of friction to the end-user experience as they need to remember multiple passwords. This poses a security risk. An IAM solution helps organizations enforce SSO. It can resolve security gaps caused by fragmented identities, and reduce friction in the end-user experience by eliminating the need for multiple passwords.

Identity-governance based Zero Trust models work especially well for organizations using cloud services. If the platform doesn't allow for organization-owned or operated Zero Trust components, identity-based implementations are your best bet.



Some possible Zero Trust deployment models



So far, we've discussed the various logical components of Zero Trust and the elements of the ZTX ecosystem. We've also seen three approaches to implementing Zero Trust in an organization.

However, we still don't have a clue as to how these could look when deployed on ground. So, before we move on to the next section, here's a non-exhaustive list of potential Zero Trust deployment models. Depending on an organization's network set up and requirements, multiple deployment models might be employed for different processes.

1. Device agent and gateway-based deployment

Here the PEP is divided into two components. An agent on organization endpoints, and a gateway on or in front of the resource.




Figure 5: Device agent and gateway-based model.

In this setup, the agent sends a request to the gateway which forwards it to the PA. The PA shares it with the PE. If the PE approves the request, the PA configures a communication channel between the agent and the gateway.

2. Enclave-based deployment

This is similar to the agent and gateway-based model. Only here, the gateway resides on the boundary of a resource enclave. These resources could either serve a single function or be unable to communicate with a gateway, as in the case of legacy systems.




Figure 6: Enclave-based model.

It functions similar to the agent and gateway model. Organizations with legacy applications or on-premises data centers that can't connect to gateways might find this useful.

3. Resource portal-based deployment

In this model, the PEP is a single component instead of being split into an agent and a gateway. It acts as a gateway portal for either a single resource, or a group of related resources.




Figure 7: Resource portal-based model.

It is flexible and allows for BYOD policies as well as collaborative projects between organizations. However, while the lack of an agent makes it easier to set up, it also means the organization has limited visibility and control over the assets accessing a portal.

4. Device application sandboxing

This variation of the agent and gateway model makes use of compartmentalization to ensure security. Here, approved applications are run in a sandbox that is kept separate from other apps and the rest of the asset. The goal is to protect applications and application instances from other applications, or a compromised asset.





Figure 8: Application sandboxing.

Vetted apps can communicate with a PEP to gain access to resources. However, the PEP will refuse connections from other unauthorized apps on the asset.

This model helps where it's not possible to scan assets from vulnerabilities. However, ensuring the security of each sandboxed application on each organization device might require more effort than simply monitoring devices.



Trust algorithms: How a Zero Trust system “thinks”

If the PE is the brain of a Zero Trust model, the trust algorithm is the thought process it follows. It’s how the PE analyzes all the inputs it receives (see Figure 3) and makes its decision to grant or deny access to a resource.

The data the PE receives can be broadly classified into the following categories:

- **Access request:** This is the access request received from an identity. It contains information on the resource that needs to be accessed and the requester. This could include information about the device the request is raised from (OS, patch status, etc.), its location, and more.
- **Subject database:** This helps the PE understand “who” is requesting access to the resource. It contains information about all subjects (people, and machines and processes) and the privileges assigned to them. This could also contain information on their behavior patterns, location, and more.
- **Asset database:** This informs the PE about the observable status of the asset making the request. It contains information about the status of all organization-owned or managed assets (physical and virtual). It informs the PE about the OS, patch status, and other security posture-related information of the asset in question.

- **Resource requirements:** This is the set of policies that determine the minimum privileges and other requirements for accessing the resource. These requirements vary as per the sensitivity of the resource and could include MFA, network location, and other factors.
- **Threat intelligence:** These feeds help inform the PE about threats present on the internet. They often contain information on the latest active malware, attack signatures, and mitigation strategies for various threats. These could also contain information about any suspicious communications seen from the device making the access request.




Figure 9: Trust algorithm inputs.

The trust algorithm determines the weight the PE gives to each of these inputs. This affects its analysis and the final decision it makes. As discussed above, this final decision is then passed on to the PA which takes the necessary action.

Different schools of thought: Variations in trust algorithms

As with Zero Trust access itself, there's no single way to implement a trust algorithm. In general, different organizations will give different weights to various trust algorithm inputs as per their needs.

However, the differences in trust algorithms go deeper than this. The first difference lies in how the various factors are evaluated. Here, trust algorithms can be classified as criteria-based or score-based.

Criteria-based trust algorithms have a set of attributes that must be met for access to be granted. These attributes are evaluated in a binary fashion. If the criteria isn't met, access is denied.

Score-based trust algorithms, on the other hand, calculate a score or confidence level. The organization assigns values and weights for all data sources. The algorithm uses these to calculate the final confidence level. Access is only granted if the confidence level or score is above the threshold value that's been set for the resource.

The second difference in trust algorithms lies in how requests are evaluated in the context of other requests by the same identity or device. Based on this, trust algorithms might be classified as singular or contextual.

Singular trust algorithms, as the name implies, look at each request individually and do not take the requester's history into account. This makes them faster, but could result in attacks going undetected if the attacker's actions are in line with the requester's role. For example, a compromised finance team account accessing payroll information and company finances wouldn't necessarily be considered suspicious.

Contextual trust algorithms take the requester's recent history into account when making their evaluation. Here, the PE stores some information on each identity's recent actions. These algorithms are more likely to detect attackers using compromised accounts if their access pattern is different from the account's usual behavior. Unusual behavior patterns can trigger additional authentication steps or result in limiting or denying access.

These two sets of approaches—how the factors are evaluated, and how each request is evaluated—are independent of each other. It's possible to have any combination of these trust algorithms, like a singular criteria-based algorithm, a contextual score-based algorithm, and so on.

All said and done, a contextual trust algorithm would be the ideal choice. They can be more dynamic in nature, and offer more granular access control. They're also better at detecting insider threats or compromised accounts.

For example, if someone who typically works from morning to evening is suddenly trying to access sensitive files at night, the contextual trust algorithm could ask for more stringent authentications and generate an alert. A singular trust algorithm would not be able to do this.

However, having a contextual algorithm might not be possible given an organization's infrastructure components. In this case, the organization will have no choice but to adopt a singular trust algorithm.

Another potential issue with contextual trust algorithms is usability. An overly cautious algorithm might raise multiple re-authentication requests even for behavior that is in line with an identity's historical trends, leading to usability issues.

In fact, these issues are to be expected in the initial "tuning" phase of any Zero Trust algorithm's deployment. Administrators will need to fine-tune the weights and values of each factor over time to ensure that security policies are enforced without affecting business processes.

Walking the walk: How ManageEngine can help you implement Zero Trust security

As mentioned in the “a peek under the hood” section: Zero Trust is more than just network microsegmentation or identity and access management. It’s a holistic approach towards security.



How ManageEngine approaches Zero Trust security

How ManageEngine approaches Zero Trust security



To achieve Zero Trust security, effectively, you need to have solutions that cover the entire security lifecycle for all five pillars of Zero Trust. The security lifecycle can be defined in terms of the following four stages:

- **Visibility:** Knowing what's happening throughout the network.
- **Deduction:** Detecting changes, incidents, and threats.
- **Response:** Taking action to solve incidents or threats.
- **Resolution:** Strengthening the network to prevent or mitigate the impact of threats.

Combining this with the four pillars of Zero Trust, gives you ManageEngine's take on achieving Zero Trust security:



Have complete visibility into all identities, devices, data, applications accessing your network, and their activities.



Monitor and analyze all entities and activity within your network and flag any suspicious activity or threats.



Respond to threats and risks as appropriate-by forcing additional authorization and

authentication, restricting user privileges, quarantining devices, terminating sessions, and so on.



Resolve security issues by taking preventive action and hardening your security—such as by resetting credentials, patching vulnerabilities, etc.

Besides the five pillars of Zero Trust, you also need to implement security analytics and security automation tools to ensure holistic security coverage.

No matter what approach you choose for implementing Zero Trust in your organization, you'll need to ensure you have each of these elements covered.



How ManageEngine solutions can
help you with your Zero Trust journey



How ManageEngine solutions can help you with your Zero Trust journey



At ManageEngine, we provide a variety of solutions that cover several of these requirements. No matter which approach to Zero Trust you choose, our solutions integrate well with each other and third-party offerings to help you build a holistic security ecosystem.

If you're just starting your Zero Trust journey, our products help you build the technological foundation needed to support your policies and processes.

If you're already on the path to implementing Zero Trust security, our solutions can help fill the gaps in your existing model.

Here's a brief glimpse at our various products and how they can benefit your organization:



Identity security

- Gain complete visibility into all your identities across your on-premises and cloud applications.
- Monitor, record, and remotely terminate privileged sessions as needed.
- Get alerted to suspicious activities within your network, and detect signs of brute-force attacks, account compromise, and more.

- Secure your organization against credential-based attacks with MFA, least privilege, and just-in-time access controls.
- Secure privileged credentials in an encrypted vault, enable secure, passwordless access to privileged accounts, and more.



Device trust

- Monitor and track all endpoints and assets accessing the enterprise network.
- Track each endpoint's health, availability, and security status.
- Remotely lock or quarantine devices, and wipe corporate data from them.
- Detect and patch vulnerabilities across OSs and applications to reduce your attack surface, and quickly restore endpoints using disk images.



Data security

- Locate, identify and classify sensitive files and vulnerable data across your network.
- Monitor all activities and modifications in critical files in real time.
- Prevent unauthorized access and transfer of sensitive data via any medium (email, printers, USB devices, etc.).
- Detect and shutdown ransomware attacks with automated threat response.



Network security

- Monitor your network and servers in real time.
- Detect rogue devices and network misconfigurations.
- Carry out network forensic analysis to detect threats or attacks.
- Backup and restore network configurations for quick disaster recovery.



Security analytics and automation (SIEM and SOAR)

- Analyze network activity to detect and defend against cyberattacks.
- Sync data from threat intelligence services to improve detection and mitigation.
- Detect insider threats and account compromise with UEBA.
- Automate incident response with predefined and custom workflows.

If you're interested in learning more about our products and how they can help you in your Zero Trust journey, please write to tech-expert@manageengine.com. We are happy to help.

A few parting thoughts

Implementing Zero Trust security in an organization is a marathon, not a sprint. IT admins can't simply flick a switch and enable it.

First, it requires the implementation of several technological capabilities. In turn, successfully implementing a component requires time and focused attention to ensure there are no errors and no security gaps in the system.

Secondly, Zero Trust security needs a buy in from everyone, from the C-suite down to the frontline employees. People at every level need to embrace the new organization policies and processes.

Third, IT teams need to ensure they're able to enforce the required access and security policies for all identities in their network at all times.

As mentioned by Joshua Crumbaugh, CTO of Phish Firewall, in the [ManageEngine Insights videocast on Social Engineering and Zero Trust](#), it's important to ensure Zero Trust principles are actively enforced for everyone.

This helps prevent issues like privilege creep, where a user who gets promoted and/or moved across multiple departments, ends up with access privileges for all their past roles. Privilege creep defeats the purpose of least privilege access and makes these users' accounts a juicy target for attackers.

Lastly, remember: there is no one-size-fits-all solution. You need to evaluate your organization's infrastructure and business operations to decide which approach is most feasible, and what roadmap to follow.



References

1. Scott Rose, Oliver Borchert, Stu Mitchell, Sean Connelly (2020), Zero Trust Architecture, NIST Special Publication (SP) 800-207. doi.org/10.6028/NIST.SP.800-207
2. Forrester Research (2021), [The Zero Trust eXtended \(ZTX\) Ecosystem](#).
3. Rory Ward, Betsy Beyer (2014), BeyondCorp: [A New Approach to Enterprise Security](#).
4. Jericho Forum™ (2007), [Jericho Forum Commandments](#).
5. Identity Defined Security Alliance (2019), [The Path to Zero Trust Starts With Identity](#).
6. Cybersecurity and Infrastructure Security Agency (2022), [Applying Zero Trust Principles to Enterprise Mobility](#) (Version: Draft for public comment).
7. Verizon (2022), [2022 Data Breach Investigations Report](#).
8. Forrester Research (2021), [A Practical Guide To A Zero Trust Implementation](#).



ManageEngine is a division of **Zoho Corp.**