Concept Guide

Zero Trust Network Access

DEFINE / DESIGN / DEPLOY / DEMO

Table of Contents

Change Log

Date	Change Description
2022-05-05	Initial release.

What is ZTNA?

ZTNA, by industry standards, is a product or service that protects applications by allowing only trusted entities access to the application. Trust is determined by a trust broker that continually verifies the identity and context of the connecting entity while performing access control based on these factors. Lateral movements within the protected application and network are limited, further reducing the attack surface for compromised hosts.

As you ponder the definition and how this applies to your organization, ask yourself these important questions:

- **Who** are the entities (users, devices, and security posture) that we trust? **What** are the characteristics of entities we do not trust?
- **What** are the applications and resources that we want to protect? What are the levels of trust required to access these applications?
- **Where** are the trusted entities connecting from? **Where** are the protected applications located?
- **How** are the trusted entities connecting to the protected applications?

To visualize, imagine a ZTNA solution delivered by Fortinet:

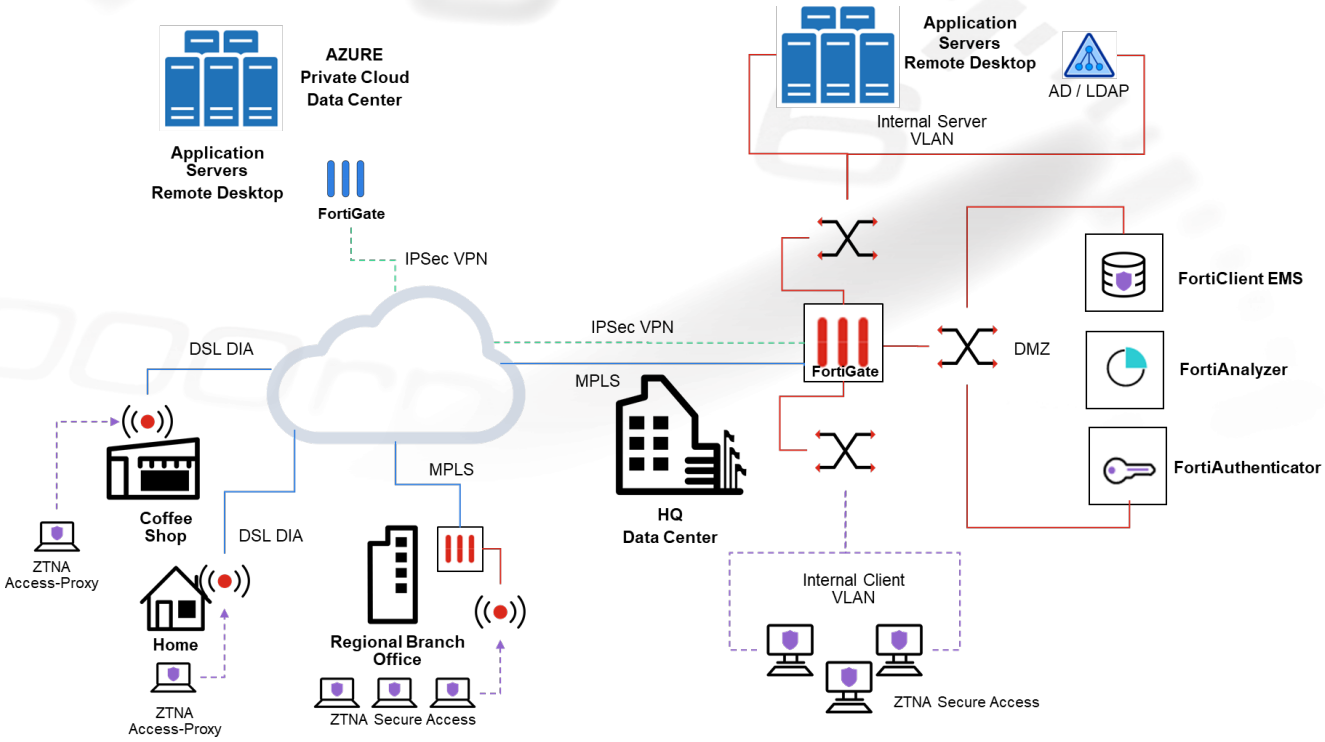

In this example, trusted entities are corporate users and devices accessing from homes, coffee shops, regional offices, and within the corporate headquarters (HQ). They require access to application servers within the internal server VLAN of the corporate HQ and application servers in the cloud. They may be accessing the application servers using remote desktop (RDP), SSH, HTTPS, SMB, or other means.

Traditionally, remote users will have VPN access to each of the protected networks, which requires various remote tunnels to be established. Internal clients may inadvertently be allowed more access than needed, such as access to the DMZ network.

ZTNA assumes no users or devices are trusted, until the trust broker confirms the device and user are who they are, and their security posture is compliant with the company standards. This greatly reduces unauthorized access from an unknown device with stolen credentials, and the chance of a compromised host accessing protected resources. Role-based access control reduces lateral movements and maintains similar access control based on access levels regardless of location.

In this document, we will explore the concept and components, and learn about how Fortinet delivers on each aspect of ZTNA.

Intended audience

Mid-level network and security architects in companies of all sizes and verticals should find this guide helpful.

About this guide

This guide aims to provide a broad overview of Zero Trust Network Access concepts, and introduce products in the Fortinet portfolio that work together to implement a scalable ZTNA solution. Industry standard terminologies are used, with introductions to Fortinet specific terms, concepts, and technologies.

Readers can proceed to the ZTNA Architecture and ZTNA Deployment guides when they are familiar with the concept and terminology, and are ready to explore different designs to use in their environment.

ZTNA concepts

This chapter addresses the following questions:

- Who are the trusted entities?
- What are the protected applications and resources?
- Where is everything located?
- How are protected applications and resources being accessed?

Trusted entities

Zero Trust Network Access (ZTNA) requires a trust broker to verify the identity and context of the connecting entity. Any untrusted devices are prevented from accessing the application. But what is trust in the context of Zero Trust? And **who** is allowed in a Zero Trust access model?

Trust

Trust is not a one-time authorization, but rather a continual evaluation in the ZTNA. Zero Trust is the opposite of implicit trust, which grants trust implicitly based on factors such as the location, device, or user. While trust may be heavily governed by user and device identity, it is not solely based on identity either. A trust algorithm can factor in the location, time, behavior patterns, threat intelligence, security posture, and other measurements dynamically. Trust is a combination of quantifiable factors that are used to evaluate an entity based on a company's requirement.

Trust broker

A trust broker is a generic term for the system that evaluates the entities for its trust context. In the NIST Zero Trust Architecture model, they are referred to as a policy enforcement point, policy engine, and policy administrator. The trust broker uses various inputs that measure trust and processes these inputs using its trust algorithm to produce a policy decision. This decision determines whether a device is trusted and whether it is allowed access to a protected application or resource. In turn, it enforces the appropriate policies on the incoming traffic. The trust broker may be a single system or multiple systems working together.

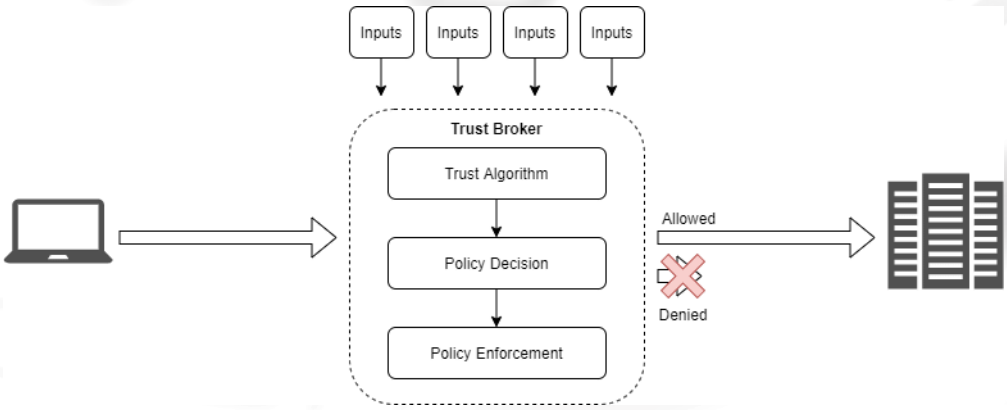

Identity

Identity is usually a heavily-weighted factor in the policy decision, which encompasses both user and device identity. User identity checks verify that the user requesting access is who they are, and determines the access level of the requester. User authentication is typically completed against an authentication server such as Windows Active Directory, an LDAP server, a RADIUS server, or a SAML identity provider. Multi-factor authentication can be used to further secure the authenticating user.

Device identity verifies that the device is who they are. For example, access is only granted to company-issued laptops with the proper client certificate issued by the company's asset management server. In this case, the device's client certificate issued by the enterprise asset management server can be used to verify the device identity.

Security posture

Security posture is a generic term for the overall security attributes that can be evaluated against an entity. This can be the device's OS and software version, known OS and software vulnerabilities, security applications running on the device, and so on. Security posture also produces important inputs for the trust algorithm as it measures how vulnerable a device is to advanced threats and attacks. Devices with outdated OSes, software, and unpatched vulnerabilities can be denied access as they are more susceptible to being compromised.

Enterprise asset management system

While the trust broker accepts various inputs to make a policy decision, the inputs are predominantly measured by external systems. User identity and attributes are handled by an external authentication server, whereas device identity and attributes are often handled by an enterprise asset management system. This system must be able to continually and transparently measure changes to the security posture of assets under management to provide the trust broker with updated information about each device. For example, a detected critical vulnerability may be shared with the trust broker to immediately break connections from the device.

Other environmental factors

Environment factors such as the device's geolocation, time of day, and whether the device is in a public location behind NAT can contribute to the trust algorithm. For example, access from an adversarial geolocation where the company has no employees may cause suspicion. Similarly, access from a trusted

location during non-business hours may equally raise suspicion. Access from an untrusted open WiFi network at a coffee shop might also be prohibited under company policies. These attributes can sometimes play an important role in assessing trust in a device.

In summary, various trust algorithm inputs and measurements are provided by sub-systems such as authentication servers, enterprise asset management systems, and next-generation firewalls (NGFWs).

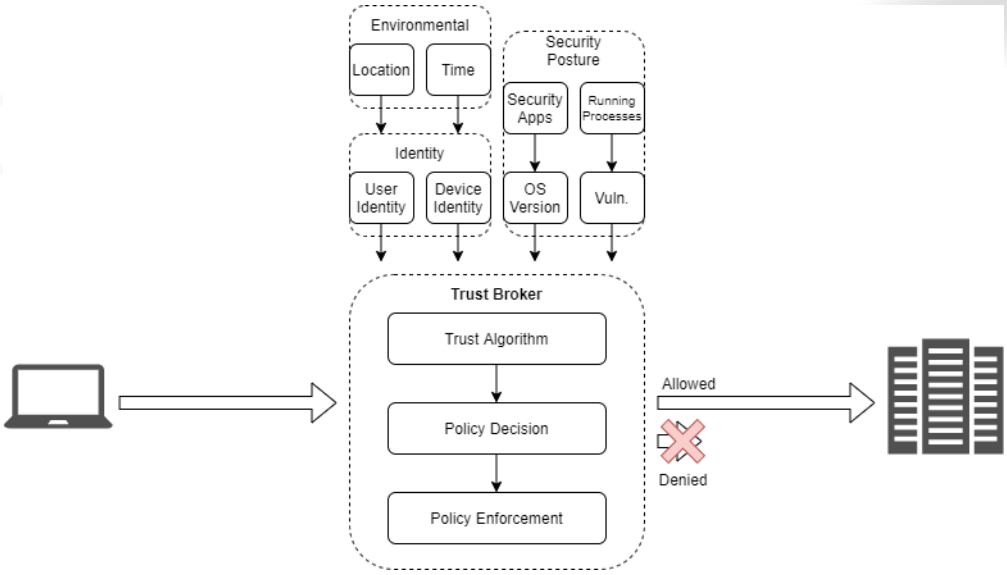

Protected applications and resources

In addition to evaluating **who** is trusted, defining a list of protected applications and resources is just as important. Not every server needs to be exposed to every user. So, mapping who has access to what prevents unnecessary access and lateral movements in the protected network.

Corporate network

In an enterprise network, many resources may be situated within the corporate network within the headquarter or branch offices. These may be broken down further into shared and individual resources and applications, as well as different types of devices. In general, these resources should not be implicitly placed into the same trust zone just because they are located in the same office or within the same department.

Using segmentation and micro-segmentation can help organize the resources into groups that make it more intuitive for network administrators to administer access. For example, when using network access control (NAC) features on a next-generation firewall (NGFW) or intelligent switch, devices can be identified based on their MAC address, vendor type, software version, and more. VoIP phones within a certain department's network, for example, can be grouped into their own micro-segment where remote access is prevented. The PCs inside a department's network may be grouped into a different micro-segments that allow remote access through RDP.

Data center and cloud

Given the distributed nature of today's enterprises, resources are just as often located in data centers and the cloud. Using similar strategies will help separate resources into their own segments for easier administration of role-based access control.

Depending on where the trust broker is located, remote resources in data centers and the cloud should not be accessed directly without passing through the trust broker. Allowing this circumvents the Zero Trust Network Access architecture. For example, if a single trust broker is deployed in the headquarters (HQ), then access to remote resources in the cloud should first pass through the HQ, and then get tunneled to the remote resource through the cloud on-ramp. Similarly, traffic to data centers should also be tunneled through site-to-site VPNs.

Level of trust

Besides understanding that not all resources within a protected location should be implicitly grouped into the same trust zone, it is just as important to clearly define what level of trust is required to access each resource group. Usually most resources' access levels will depend on the authenticated identity of the source, but some resources may require stricter policies in terms of security posture and environmental factors. A good Zero Trust Access solution will incorporate these into its evaluations.

Location

Given that the **who** and **what** are clearly defined, the question of **where** has likely been considered. With the proliferation of mobile workers who work partly from home, public spaces, and in the office using the same laptops and mobile devices, the user and device location are never static. Devices can be compromised outside of the network, causing infection while devices are physically used inside the corporate network. Therefore, clients should never be trusted based on the assumption they are static and stationary. Trust should be evaluated constantly.

Conversely, protected company applications not only reside within a corporate network, but also as cloud applications and applications in remote data centers. When providing access through the trust broker, consider whether the connections between the trust broker to the applications are secure. Also consider if there are routes that users can take without passing through the trust broker that could be exploited.

Accessing protected applications and resources

Understanding **how** a resource and application are accessed allows the trust broker to block any unnecessary and non-essential connection methods that can be exploited by malicious actors. For example, personal desktops that are allowed to be accessed remotely should only require access via RDP. Therefore, the trust broker should only forward RDP traffic, and optionally scan this traffic for anomalies.

Personal desktops should also be restricted from unnecessary lateral access and movements. If a desktop is compromised, it can act as a jump point to infect other devices on the network, or take control of other devices. For example, a compromised desktop that can use SSH to get into other devices in the network might be inadvertently allowed access to corporate firewalls and security systems.

ZTNA components

Conceptually, the definition of Zero Trust Access should not differ too much from vendor to vendor, even if the terminology is different. However, the implementation and architecture of ZTNA can vary widely depending on each vendor's security portfolio.

The Fortinet Security Fabric has a wide array of products, so it delivers on every component of ZTNA described in the previous chapter. This chapter examines each component in detail:

FortiGate as the trust broker

The FortiGate is the heart of the Fortinet Security Fabric, whether it is deployed on the edge, in the cloud, or as an Internal Segmentation Firewall (ISFW). Within the ZTNA architecture, the FortiGate acts as the trust broker to consolidate inputs, process them, make intelligent policy decisions based on the trust algorithm, and enforce the policy.

As a next-generation firewall (NGFW), the FortiGate is able to scan the traffic to detect anomalies, malware, viruses, and the latest threats. Application control allows the FortiGate to granularly identify the application based on traffic patterns to correctly police the application and protocol.

In this example, a ZTNA policy allows traffic from authenticated users in Canada, during business hours, to access desktops using RDP. The security posture check is defined in the ZTNA tags, which require the connecting PCs to have AV enabled and a low vulnerability rating.

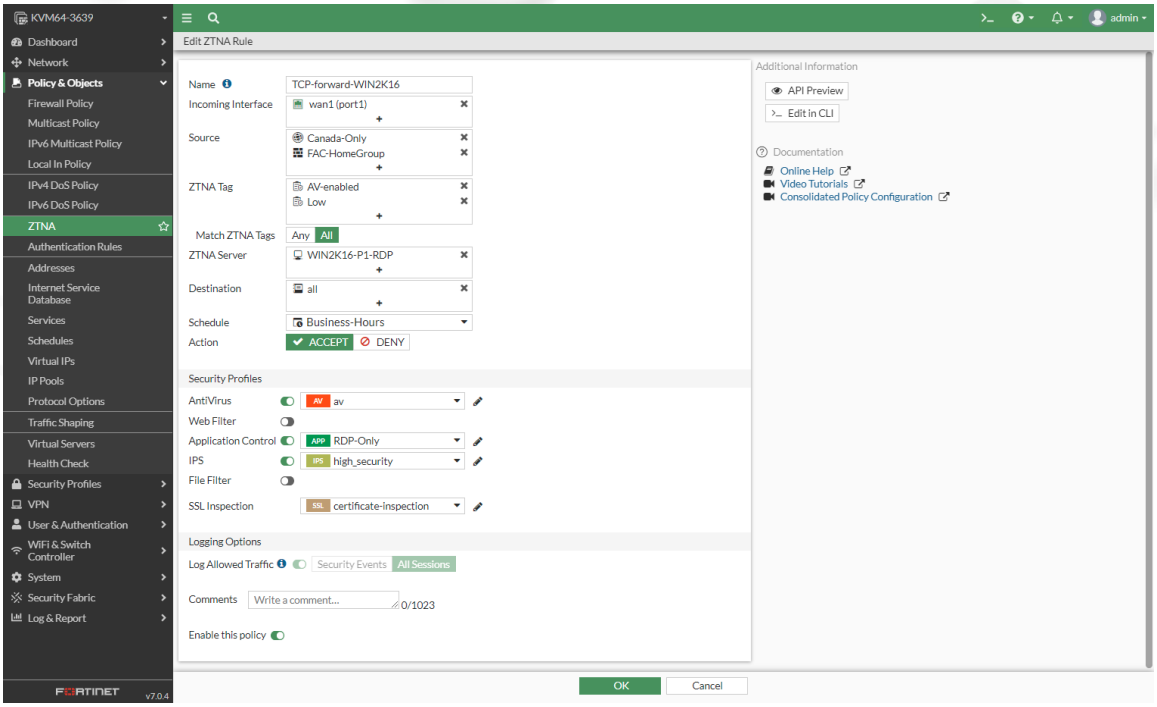

FortiClient EMS

The FortiClient Enterprise Management System (EMS) serves several purposes in the ZTNA architecture:

1. Collect information about managed endpoints used for input in the trust algorithm.
2. Continually monitor managed devices for changes in security posture (such as vulnerability levels and identification of malware) and update the trust broker.
3. Generate and install client certificates on managed endpoints to uniquely identify each endpoint device.
4. Catalog remote servers that can be accessed through ZTNA and install connection rules on the FortiClient endpoint.

Security posture: Zero Trust tags

A registered FortiClient endpoint automatically provides information about the endpoint to the EMS server. Zero Trust tagging rules allow for granular detection of specific attributes based on the connecting device's operating system.

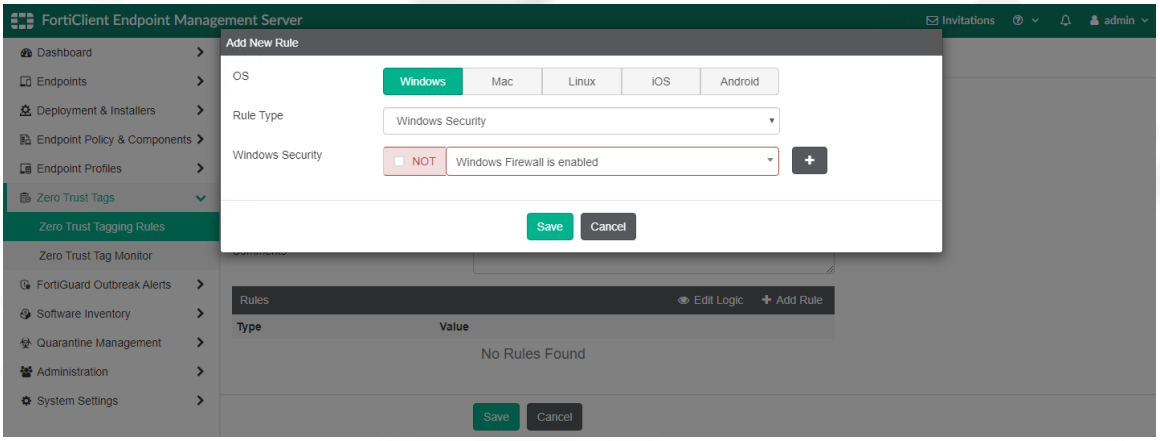

A full list is available in the Endpoint Posture Check Reference guide.

Client certificate

Every FortiClient endpoint that registers to the EMS server is issued a client certificate from EMS's certificate authority. This unique certificate identifies the endpoint when they authenticate against the FortiGate. The FortiGate only trusts devices that are registered to the EMS and hold the client certificate. Whenever a FortiClient endpoint unregisters from the EMS, the client certificate is revoked.

The following example shows an EMS server and the information collected on the registered endpoint.

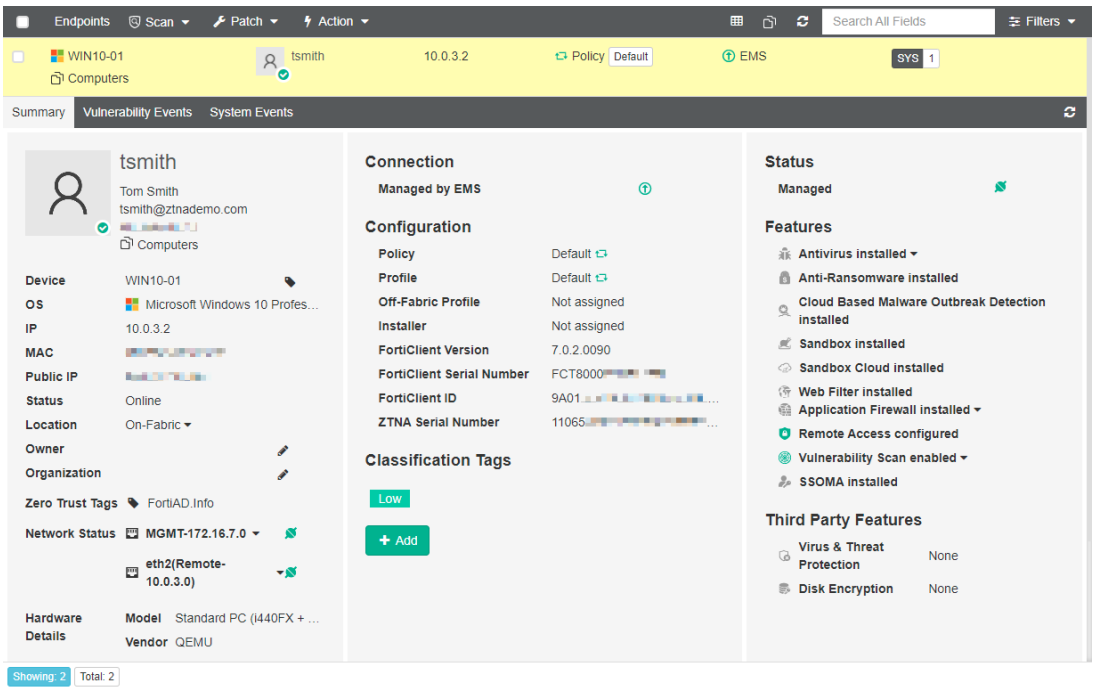

A unique certificate is issued for the registered FortiClient based on its FortiClient ID. A unique *ZTNA Serial Number* is assigned to the certificate.

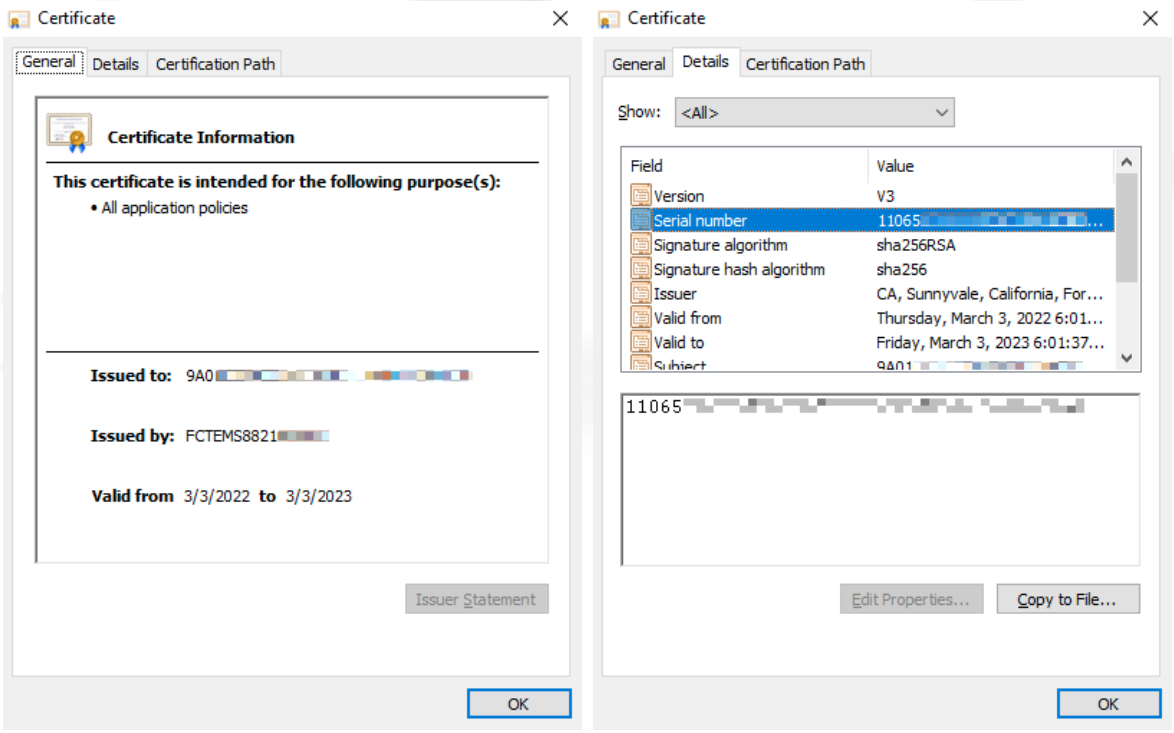

FortiClient endpoint

The FortiClient software has many capabilities to protect the endpoint. In the context of ZTNA, its main roles are to:

1. Register to the EMS server to provide endpoint information to the server.
2. Install the client certificate issued by the EMS to identify itself to the FortiGate.
3. Scan for threats on its system and report back security posture changes to the EMS in real-time.

The majority of the inputs used in the FortiGate's trust algorithm are sourced from FortiClient and pushed to the EMS. The FortiClient endpoint must stay connected to the EMS server in order to connect to resources protected by ZTNA. The information is synchronized automatically, so it is completely transparent to the end user.

Centralized authentication

User identity is taken care of by authenticating against an authentication server. In an enterprise, this may be a Windows Active Directory (AD) server that manages users for a domain. While authenticating directly against the authentication server is very common, when multiple servers are involved, centralizing this within a FortiAuthenticator may be appropriate. A FortiAuthenticator can enhance security further by applying multi-factor authentication using FortiTokens. Centralized authentication allows multiple systems to authenticate against multiple servers seamlessly and efficiently.

For example, an enterprise may have two Windows AD servers working redundantly to provide authentication services to its users; however, it has no multi-factor authentication capabilities. Given that several FortiGates running on the edge and internally as ISFWs require access to the authentication

database and the company wants to enforce multi-factor authentication, a FortiAuthenticator can be put in place to centralize the authentication services.

Secure wireless and switching for NAC micro-segementation

When securing the LAN edge that comprises different wireless and wired device types, it is no longer feasible to segment them simply by location and business groups. Using an Internal Segmentation Firewall (ISFW) is a great way to restrict inter-segment traffic and implement micro-segmentation.

The FortiGate has built-in NAC capabilities that it can apply to managed wireless APs and wired switches. NAC policies can group devices by wildcard MAC addresses, hardware vendor, device family, type, OS, and more. NAC policies can also assign devices to different VLANs, which can then be used in different firewall and ZTNA policies to granularly control the flow of traffic.

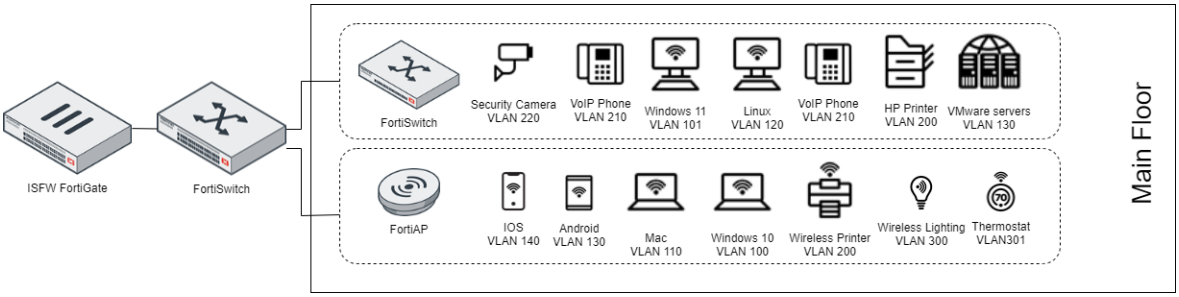

Conclusion

The goal of the Zero Trust Network Access architecture is to increase security and reduce the burden for the security staff when exposing resources and applications for access by organizational users. This can be accomplished by using different products in the Fortinet Security Fabric. The increase in security comes from a seamless integration of device identity, user identity, and posture checking into the trust Fabric. This eases the security burden by reducing the number of point products that manage device certificates, posture checks, and other identity verification technologies.

To learn more about Fortinet's ZTNA solution, you can read more about it in the ZTNA Architecture and ZTNA Deployment guides, or access more content from our resource centers listed in the Appendix.

More information

Feature documentation

- FortiOS 7.0 Administration Guide: Zero Trust Network Access introduction
- FortiClient EMS 7.0 Administration Guide: Zero Trust Tags
- FortiClient 7.0 Administration Guide: ZTNA Connection Rules

ZTNA resources

- Endpoint Posture Check Reference

Resource centers

- ZTNA solution hub
- ZTNA 4-D resources
- Fortinet ZTNA overview

Other 4-D documents

- ZTNA Architecture guide
- ZTNA Deployment guide

External references

- Gartner: Zero Trust Network Access
- NIST: Zero Trust Architecture

www.fortinet.com

01-700-805110-20220505