Maintaining Zero Trust With Federation

William R. Simpson¹ and Kevin E. Foltz²

^{1,2}*Institute for Defense Analyses (IDA), Alexandria, Virginia USA.*

Abstract — Federated activity presents a challenge for enterprises with high-level security architectures. Federation involves information sharing among services and with working partners, coalition partners, first responders, and other organizations. Federation may be unilateral or bilateral with similar or dissimilar information-sharing goals. Strong internal security, including zero trust controls, often do not extend cleanly across enterprise boundaries, potentially leading to insecure shortcuts and workarounds that can become the rule instead of the exception. This paper presents methods for an enterprise to extend its zero trust security policies to include federation partners. It applies to federation partners that support the same security policies with compatible standards and services and to partners that provide a similar but incompatible security framework, a subset of required security services, or no security services. The partner organization may be fully trusted, partially trusted, or untrusted. Even for trusted partners, the services may not meet required security standards. Our solution combines selected partner security services, internal services, derived credentials, delegated authorities, and supplemental services to form the federation security architecture based on zero trust premises to the maximum extent. This paper uses the Zero Trust for Enterprise (ZTE) architecture as the starting point for a secure enterprise and addresses the challenge of extending this model to federate with different types of partners. We review the security approach, the security properties, and several options for an enterprise to maintain the ZTE security properties while enabling federated sharing with other enterprises that have different capabilities and levels of trust.

Keywords—Coalition, Federation, First Responders, Security, Sharing, Trust, Zero Trust.

I. INTRODUCTION

The Zero Trust for Enterprises (ZTE) framework is a way to provide strong security guarantees within an enterprise. It provides rules for secure authentication and authorization that are distributed, scalable, and centrally managed. Entities in the enterprise need only authentication credentials and attributes in data stores to grant access to appropriate applications and services throughout the enterprise. This system works within an enterprise, but often there is a requirement to share resources with different enterprises, which poses a challenge to the ZTE framework.

This work discusses options for federation across enterprises in which one enterprise uses ZTE. The partner

enterprise may use ZTE, a similar security framework, or limited or no security services. The partner enterprise may be trusted, partially trusted, or untrusted. This paper provides options for a ZTE enterprise to enable federated sharing with other enterprises that have different capabilities and levels of trust. This enables more functionality within the ZTE framework and fewer workarounds, which provides stronger and more uniform security measures.

II. BACKGROUND

The existing network defense architecture has failed to stop vulnerabilities and exploits while increasing overall complexity. In response, several approaches have been suggested, including end-to end approaches to network security. Several variants exist: distributed computing approaches [1, 2], endpoint defenses [3, 4], Zero Trust Architecture (ZTA) [5, 6], and ZTE [7, 8].

- Distributed computing is a model in which components of a software system are shared among multiple computers. Even though the components are spread out across multiple computers, they are run as one system. This improves efficiency and performance.
- Endpoint defenses define the requested provider and the requester as the endpoints. They often use endpoint health indicators and request identity information to provide fine-grain access control over network resources.
- ZTA uses the principle of protecting individual resources, such as data and computing, within the enterprise instead of protecting the entire internal network at its borders. Requests coming from the internal network are not inherently trusted, and identities and access credentials must be verified at each resource. ZTA is designed to prevent data breaches and limit internal lateral movement in the enterprise.
- ZTE is a security architecture developed for the U.S. Air Force to overcome the assumptions inherent in fortress defenses. ZTE encompasses all of the above methods and is designed from the ground up to be a

ZTA. The ZTE security approach was formerly named enterprise level security (ELS), which is based upon zero trust principles. A more complete description of the ZTE (ELS) security architecture is provided in section 5.

Although each of these variants are unique in their security approaches, they all share seamless end-to-end encrypted communications in their architectures, as shown in Figure 1. Entities in the enterprise may be active or passive. Passive entities include storage elements, routers, wireless access points, some firewalls, and other entities that do not themselves initiate or respond to web service or web application requests. Active entities are those entities that request or provide services according to ZTE. Active entities include users, applications, and services. This is a basic conflict with current network defense models, which break this connection at one or more internal components, in effect making them active entities.

The figure illustrates end-to-end security's requirement that the front door to the enterprise be passive; however, with the fortress approach, the network defense has evolved into a very active set of entities.

This document follows from our experiences with developing the ZTE architecture. ZTE proceeds from core tenets to application techniques to core requirements.

III. LITERATURE REVIEW FOR CURRENT DEFENSE PACKAGES

Computer network defense is defined as “Actions taken through the use of computer networks to protect, monitor, analyze, detect and respond to unauthorized activity within the enterprise information systems and computer networks.” [9] The current defense package assumes that the threat can be stopped at the front door, as shown in Figure 2. All traffic in the enterprise, both coming and going, are routed through this front door. The front door is often onerous enough that an administrator back door is made available [10] to bypass much of the security check. These back doors are often the target of exploits.

Figure 1: End-to-End Encryption

Figure 2: Fortress-Protected Enterprises

All active entities in ZTE systems have public key infrastructure (PKI) certificates, and their private cryptographic keys are stored in tamper-proof, threat-mitigating storage. Communication between active entities requires bi-lateral, PKI, end-to-end authentication. A verifiable identity claims-based process provides authentication.

The elements involved in implementing network and application defense are numerous and complicated. Functionality is provided by a wide range of appliances. This functionality may be for quality of service to the user or quality of protection to network resources and servers. These appliances are often placed in-line (and many operate at line speeds for all communications coming from or going to the enterprise), and some require access to content to provide their service. The literature is confusing because offerings include multiple services under various titles such as multi-function firewalls or

advanced defense systems. Figure 3 provides a representation of how these appliances come between the user and the application. The most spectacular result of the network packages shown in Figure 3 is that the fortress defense has spectacularly failed—breaches occur almost daily. The appliances in the package do stop the current threats for a short period, but new threats materialize very shortly and once again defeat the fortress approach. Even with detection and mitigation, there is continued threat presence over long periods. The advanced approaches described here assume that the threat is present and in the enterprise at all times. Although this is not true at any given time, it is certainly true at various times during operations.

The number and types of appliances can be quite large. The following is a partial list of functional types as provided in the current literature:

Figure 3: Endpoint

- Header-based scanner/logger [11]
 - Views only unencrypted portion of traffic
 - Synchronous or asynchronous operation
 - Scans for defined behavior and logs traffic
- Content-based scanner/logger [12]
 - Views decrypted transport layer security (TLS) content
 - Synchronous or asynchronous operation
 - Scans for defined behavior and logs traffic/content
- Header-based firewall [13]
 - Views only unencrypted portion of traffic
 - Synchronous operation
 - Scans for and blocks defined behavior
- Content-based firewall—block only [14]
 - Views decrypted TLS content
 - Synchronous operation
 - Scans for defined behavior and blocks (terminates) connection
- Content-based firewall – modify malicious content [15]
 - Views decrypted TLS content

- Synchronous operation
- Scans for defined content; blocks or removes content without blocking the connection
- Access Web accelerator [16]
 - Views decrypted TLS content
 - Synchronous operation
 - Modifies content for performance
- WAN accelerator [17]
 - Views decrypted TLS content
 - Multi-party system
 - Synchronous operation
 - Modifies content representation between parties but no end-to-end modification
- Load balancers [18]
 - Distributes load among destination end points to improve throughput and reduce latency
 - May decrypt content:
 - May combine encrypted flows through a “secure sockets layer (SSL) accelerator”
 - May distribute content by request to different servers based on load
 - These load balancers are active entities
 - May not decrypt content:
 - Using “sticky” or endpoint balances may route all requests from an entity to the same server
 - These load balancers are passive entities

To ensure these appliances operate at line speeds and process all incoming and outgoing communications, they are built from expensive custom hardware, and thus, they are complex with their own set of vulnerabilities and exploits.

IV. SHORTCOMINGS OF THE CURRENT APPROACHES

Each of these appliances offers some functionality and increases the threat exposure. None of these are free from vulnerabilities from a security standpoint, and they do increase the threat surface and the vulnerability space. For example, default passwords or other improperly secured access methods allow an attacker access to any data that the appliance can access.

Most recently, the SolarWinds exploit was accomplished by compromising the supply chain for one of these appliances. For detailed scans, these appliances could decrypt all network traffic to and from a server. With a large number of independent appliances, this represents a significant cost and security risk. Use of any appliance must be balanced by the increased

functionality and the increased vulnerability. The situation is further complicated by vendor offerings of load balancers with firewall capability, “smart” accelerators that scan content, and software-only offerings that will provide most of these functionalities in a modular fashion.

V. THE ZTE ARCHITECTURE

The ZTE architecture has evolved from a fortress approach, in which the threat is assumed to be stopped at the front door, to a distributed security system that eliminates or mitigates many of the primary vulnerability points inherent with that system, as shown in Figure 4.

The basic process of identification involves a two-way contract between two entities initiating a communication. Each entity needs some assurance that the other party is a known entity and specifically the one to whom the communication should be allowed. This is accomplished via the presentation of claims by each party that are verifiable and may be validated. These claims are often in the form of credentials.

Active entities must be named in accordance with enterprise naming rules that assure uniqueness over space and time.

Authentication is implemented by a verifiable identity claims-based process.

Figure 5 illustrates a combination of active and passive entities in a typical request flow. The requester and provider are both active entities using ZTE. To complete the request, the provider calls on a passive data store using non-ZTE methods. Each provider has configuration data for the Enterprise Computing Operations (ECO). ECO is the collection of passive and active entities that comprise the computing system. To mitigate the lower security of the link between provider and data store, the data store interface is often locked down to only communicate with the provider.

ZTE is designed for a high-assurance environment. For ZTE, the primary concerns are the following five security principles:

- Know the Players,
- Maintain Confidentiality,
- Separate Access and Privilege from Identity,
- Maintain Integrity,
- Provide Accountability.

The following sections elaborate on how these are incorporated into the ZTE framework.

Figure 4: Distributed Security Architecture

Figure 5: Communication between Entities

A. Know the Players

All communication in ZTE is conducted between known entities. Requesters know the entity that provides the data, and providers know who is requesting it. Entity identification and authentication is performed using identity credentials, as shown in Figure 6.

Figure 6: Bi-lateral Authentication

In ZTE, the identity credential is an X.509 PKI certificate [19]. PKI certificates are verified and validated. Ownership is verified by a holder-of-key (HOK) check.

B. Maintain Confidentiality

After establishing the identity of the communication partner, the authenticated entities negotiate cryptographic algorithms and keys that enable confidential communication. As shown in Figures 1 and 7, this confidentiality is established through an end-to-end encrypted connection.

Figure 8: Claims-based Authorization

ZTE establishes end-to-end confidentiality using Transport Layer Security (TLS) encryption, in which keys are held only by the endpoints [20].

C. Separate Access and Privilege from Identity

After the connection has been established from requester to provider, the provider knows the requester's identity. In ZTE, this is not always enough to provide access, and an additional authorization credential tied to identity must be provided for access. Figure 8 illustrates the transmission of this access credential.

Figure 7: End-to-End Encryption

Figure 9: TLS Integrity Measures

In ZTE, the access credential is sent using the Security Assertion Markup Language (SAML) [21]. SAML tokens are issued and signed by a trusted Security Token Server (STS) that uses a backend Enterprise Attribute System (EAS) to manage attributes and compute access claims.

In ZTE, data integrity for web requests is implemented by TLS message authentication codes (MACs). The keying material for MACs are exchanged in the initial TLS setup using the Diffie Helman process (see Figure 9). Messages and SAML tokens include XML digital signatures, which are verified and validated [22].

D. Provide Accountability

In a high-assurance environment, accountability is important for allowing rapid identification of problems and tracing them to their source. It also allows better visibility into what is happening, which helps determine how to best allocate the enterprise resources based on who is using which services. Figure 10 illustrates the system of collecting and accumulating information about active entity activity. All active entities must monitor specified activities. For enterprise files, a monitor sweep agent reads, translates, cleans, and submits to a database for recording log records periodically or on demand. The details of this activity are provided in [23] and [24].

Figure 10: Accountability through Centralized Monitoring

VI. FEDERATION TECHNICAL CONSIDERATIONS

The previous section describes some of the technologies and principles of ZTE. This section discusses different ways to extend the principles of ZTE to a federation partner based on the partner's technical capabilities. Options include:

- ZTE federation,
- ZTE-like federation,
- Identity credential federation,
- Weak identity federation,
- Ad hoc federation,
- Person-to-person federation.

The federation options are listed from most to least technologically compatible. Each option on the list also implicitly includes the options lower on the list. For example, a ZTE federation partner provides all the options available to a person-to-person federation, ad hoc federation and weak identity federation partner, as well as identity credential federation and ZTE-like federation. This section discusses the technical means to accomplish federation for each of the options listed.

A. ZTE Federation

ZTE federation is an agreement to accept identity and access claims from another enterprise. Federation is a long-term substantial agreement that is made at the enterprise level. To resolve federation issues, the federation STS relies on the following information:

- Certificates of Federated STSs for validating signatures in SAML tokens and chaining to trusted root CA.
- A set of identity-mapping pairs with the form (Identity1, Identity2), in which Identity1 in a SAML issued by the federated STS is to be mapped to Identity2 in the local enterprise.
- A set of mapping pairs of the form (Claim A, Claim B), where Claim A in a SAML issued by the federated STS is to be mapped to Claim B in the local enterprise.

- Additional attribute mappings associated with claim mappings.

An example of data captured in federation agreements is shown in Table 1. This shows the data for two separate federation agreements. Each web service in the enterprise has a limited number of trusted root CAs for authentication credentials and trusted STS certificates for SAML signatures stored in its trust store. With ZTE federation, a list of trusted partner CAs and STSs is established. Trusted partner CAs for identity credentials are distributed to applications and services. Trusted partner STS credentials are distributed to federation STSs within the enterprise. The federation STS is called by a service when an unknown authorization credential is encountered, and the federation STS checks against known federation partners to validate the credentials, creates a new SAML token with its own signature, and returns this to the application or service for processing.

Table 1 :FEDERATION DATA REQUIREMENTS

Federation Partner 1 Information		
Certificate	Federation Partner 1 certificate and chain to root CA	
Identity Mappings	Identity 1	Identity 2
	Identity A	Identity B
	Identity r	<no change>
	...	...
Claim and Attribute Mappings	Claim A	<null>
	Claim n	Claim z
	Claim y and Attribute q	Claim y and Attribute r
	...	...
Federation Partner 2 Information		
Certificate	Federation Partner 2 certificate and chain to root CA	
Identity Mappings	Identity x	Identity y
	Identity Q	Identity R
	...	...
Claim and Attribute Mappings	Claim n	<no change>
	Claim p with Attributes x, y, z	Claim p with Attribute k
	Claim A	<no change>
	...	...

For identity and claim mappings, the special cases of “null” and “no change” are acceptable in addition to explicit values. “Null” removes the claim or identity, while “no change” leaves the original claim or identity. The claims to be mapped must match claims from sources on both sides. Claims in the federation partner SAML must match the federation agreement exactly. Claims in the re-issued SAML must match claims for the target application or service. Identity and claim mappings are added to the federation store after an amendment to the federation agreement. Revocation of a federation agreement is accomplished by removing the federation partner from the trusted STS data store. When a federation STS recognizes and validates a partner authorization credential, it uses its mapping list to map the received credential into a new credential with possibly different identity, claims, and attributes. This new credential is signed by the federation STS and returned to the requesting application or service. The application or service then processes this new SAML token as though it had received it from a valid requester within the enterprise. Failure to validate an incoming SAML token by the federation STS results in an error message response to the application or service, which leads to an authorization failure at the application or service.

B. ZTE-like Federation

If the federation partner is not using ZTE but does have a way to provide the security functions that ZTE provides, then a more complicated federation agreement may be needed with some translation algorithms. This is similar to ZTE federation, but it accounts for different semantics, formats, and data encodings. The mappings may not be exact, in which case there may be some loss of information in the translation. This may hinder automation and scalability because mismatches in data resolution could lead to either denial when access should be allowed or access when it should be denied. This is generally unavoidable when data is represented in different ways. These differences also apply to tokens and credentials. If the strength of credentials in the partner enterprise is variable, this can weaken the entire authentication process. For example, if the partner provides software authentication credentials for some users instead of hardware credentials for all users, then unless the source of the credential is explicitly provided in the signed portion of the certificate, it is not possible to provide access to hardware-credentialled entities and deny access to software-credentialled entities. The secure choice for an enterprise that requires hardware certificates in such a situation is to deny access even though the entities using hardware credentials should be given access.

For authorization credentials, if the partner uses something other than the approved SAML 2.0 format that is used in ZTE, potential problems exist again because certain guarantees provided by the SAML format—such as validity windows, signatures, or encryption—may be weaker or nonexistent in the partner authorization token. Also, even if the token is equally strong, the local STS needs to understand the partner token format in order to parse, validate, and translate it into a locally comprehensible SAML for the local applications and services. Token translators exist to transform one token format into another, but they address only the format, not the meaning or security properties. When the partner is “ZTE-like,” most of the work is already done, and the process of federating is similar to that with a ZTE partner. The challenge is working out the small differences and making sure that these small differences do not open up large security holes or prevent large groups of entities from receiving proper access. Shortcuts and workarounds may be necessary to provide a seamless experience to the end users; however, the majority of cases should work with standard translation schemes as long as the protocols, standards, and data formats are commonly used and available. In the event that formats and data are incompatible, additional work to make needed conversions at the requesting enterprise may solve the problem. For example, many STSs can be configured to issue tokens according to different protocols and formats, so configuring a partner STS to use the proper SAML format could resolve token format differences. Authentication credentials could be addressed by indicating whether a credential is stored in hardware or software so that receiving entities can decide whether to allow access. This could be accomplished by dividing the issuing CAs into those that issue only hardware credentials and those that issue only software credentials. These solutions are intended as quick fixes, and they do not require any new technology. In some cases, these solutions can promote a ZTE-like partner to a ZTE partner, but in most cases, they remain ZTE-like with more streamlined operations.

C. Identity Credential Federation

A federation partner may provide identity credentials but not authorization credentials, such as an account-based system in which the ID is used to log in. In this situation, a basic function of ZTE is missing, so it is no longer ZTE-like, but the identity credentials can still be used as a starting point for federation.

One solution is a mapping from partner identity to local identity that also includes associated authorization claims. This would be performed at the STS as part of the federation mappings. This is not the typical use for these mappings, and it requires modifications to Table 1 to combine the identity, claim, and attribute mappings.

This method would work if the identity credential is passed for authentication and if the authorization information is exchanged between enterprises and incorporated into the federation agreement. This should generally only be used for small sets of requesters, as these federation mappings are intended to be for claim and attribute equivalences and generic identity transform rules, not explicit per-entity claim and attribute information.

For larger-scale federation, the data owner or some other entity with access to the data can delegate the appropriate claims to the appropriate individuals in the partner enterprise. The delegation framework within ZTE is designed to allow such short-term access to specific individuals. In this case, the data owner, rather than the STS, maintains the mappings through the delegation service. This is a more appropriate place to store this information. However, it still requires manually assigning claims to individuals based on their identities. The basic structure for federated delegation is shown in Figure 11. The dashed lines represent the flow for setting up the delegation. The local delegator uses the special delegation service to assign access claims to the federated partner identity as stored in the authoritative content store for federation partners. The special delegation service is separate from the standard in-enterprise delegation service, and the content store for federation partners is separate from the normal store for in-enterprise entities. This special delegation service uses the claims exposure and editor service to store access claims for the federated identity in the claims repository. The solid lines represent the flow for a federated requester to retrieve a SAML token with delegated claims. The partner interfaces with the STS just as an in-enterprise user does, and the STS then makes a request, just as for any other requester, to retrieve the claims from the claims repository.

Figure 11: Identity credential federation using delegation

A third solution is to use a separate authoritative data store for the partner entities and tie it into the EAS that feeds the SAML creation process. This brings the partner entities into the local enterprise and allows those entities to request local SAML tokens that will work just like those from requesters within the enterprise. Maintenance of the authoritative data store could be performed locally or assigned to the partner enterprise. Local maintenance would again require detailed manual attention, while letting the partner maintain it allows them to manage their own entities. This is the best-integrated solution, as the required attributes are brought into the ZTE system. In all cases, a store maintains information about federated identities. For the mapping method, this is the federation agreement at the federation STS, which contains both identity information and claims and attributes used for access. For the delegation method, it is the store used by the special delegation service, which contains only identities, but no claims because these are delegated as needed. With the separate store, this store contains all identities and their attributes.

D. Weak Identity Federation

If the federation partner does not provide sufficiently strong identity credentials, then identity credential federation is not possible. For example, a federation partner that uses username and password authentication is not compatible with the PKI-based authentication used for ZTE. However, if it is determined that the need to federate outweighs the security risks associated with weak credentials, then there are methods to allow such access while preserving many of the security goals of ZTE. The primary challenge is that ZTE requires PKI authentication to the STS to receive a SAML token and to the application or service to establish a connection, but the federation partner does not have this capability. To allow creation of the SAML, a separate STS can be set up to receive other forms of authentication, such as username and password, Kerberos over TLS, other non-PKI TLS authentication, tokens, single sign on (SSO), or other methods. The STS can use the authenticated identity to generate a SAML token. If the identity format does not match the ZTE standard distinguished name (DN) format, a mapping can be performed algorithmically or based on an enterprise-wide database of identity and derived DN pairs. If the federation partner has an STS that already uses non-PKI authentication, it can be leveraged to create authorization credentials. Application and service endpoints can rely on alternative authentication methods, map the identity to a derived DN if necessary, and then compare this identity to the identity in the SAML token. In the generated SAML token, the authentication method can be indicated, or federation partners with non-PKI authentication at their STSs can be noted in the federation agreement so that receiving applications & services can enforce access accordingly.

For example, standard PKI authentication may provide full access while alternative authentication methods provide limited access that is sufficient for federation purposes. Because there is a clear distinction in authentication method and separate endpoints set up for access, this method provides a way to manage federated access independent of normal access. If a security problem is detected at a federated non-PKI endpoint, that endpoint can be shut down without affecting normal ZTE operations. This process is cumbersome and redundant because it requires new instances of applications and services and potentially standing up a new STS. It may require additional logic to enforce different access based on the authentication method. It also weakens enterprise security by allowing non-ZTE access to applications and services that otherwise require ZTE. The access can be limited, but it may provide an attacker an easier entrance to ZTE-protected data, which could be leveraged through additional attacks to gain full access. This should only be attempted for urgent and infrequent federation agreements with understanding of the risk to all affected applications and services, as well as the risk to the enterprise in general. This is not a ZTE-compliant solution, because it fails to satisfy authentication and accountability requirements. However, it is a way that an enterprise can share all or part of ZTE-protected resources while leveraging existing ZTE infrastructure. As a result, this may be a better solution than other options available.

E. Ad Hoc Federation

If there is no easy way in advance to identify what is to be shared with whom, then an ad hoc approach can address federation requirements. Delegation may be used to explicitly share resources with individuals just as needed. Delegation was mentioned earlier as an option when strong identity credentials are available. It is best suited for when the delegation decisions are distributed. A large federation agreement can be implemented by dividing the resources to be shared and the partner entities to be given access into smaller, manageable groups. A large federation agreement is broken into a large number of locally administered smaller agreements. Central control is maintained by the hierarchical structure of the delegations. This may or may not be possible for the structure of the sharing agreement between enterprises. This approach requires some form of authentication for federated individuals. This is not necessarily a separate solution from those listed earlier, but instead an option for generating access claims for these methods. This is in contrast to explicitly assigning attributes or claims to these identities or mapping federated identities and claims through a federation STS. This option is listed after weak identity federation because of the ad hoc and decentralized nature of the sharing.

With the options listed earlier, a central source determines who can access which resources. In this case, individuals make these decisions rather than the enterprises forming federation agreements. As a result, this is less secure from the enterprise perspective, especially as the scope and scale of sharing increase.

F. Person-to-Person Sharing

If all the previously discussed solutions have technological or other issues that make them infeasible, an individual may make the decision to share information with a federation partner out of band on a person-to-person basis. This type of sharing should be properly authorized, and the person sharing the material should have release authority with guidelines. This could involve emailing documents, placing them on a public website, providing printed copies, or verbally providing the needed information. This goes around ZTE by allowing access without any automated enterprise-level checks of identity or access claims. The last two also lack attribution, because the one sharing and the one receiving are the only entities that know about the transaction. In this case, shared information may be fingerprinted or steganographically imprinted for later forensics in case of leaks. Given the difficulty in stopping it, small amounts of unauthorized sharing are expected throughout the enterprise and with partners. ZTE and the federation options previously listed attempt to provide a feasible option that is better than going completely out of band. Although this method essentially bypasses most or all ZTE security, it is mentioned because it is better than many alternatives that involve giving full electronic access by credential sharing, workarounds, or back doors. This method lacks the formal security of ZTE (it may include accountability provisions on the enterprise side), but it provides the important property that only the needed information is shared because a human with authority to access and release data is making the decision on sharing each item. To help secure this method of sharing information, policies can guide people on which real-world security checks to make related to such sharing. Person-to person sharing has a long history prior to electronic sharing, so this accumulated knowledge can be used to shape policy.

G. Evaluating Options

The options presented earlier describe a spectrum from full ZTE sharing to non-ZTE sharing. It is tempting to say that ZTE is the most secure and person-to-person is the least secure, but in reality scale matters. The security features are reduced from ZTE to person-to-person, but the scale also tends to be reduced. Attacks are easier at the non-ZTE end of the spectrum, but their smaller scale is likely to result in smaller incidents.

At the full ZTE end of the spectrum, the security is high and attacks are difficult, but a successful attack could rapidly escalate to a major incident.

H. Evaluating Options

Figure 12 presents a notional representation of that risk. For efficient sharing, ZTE provides the most secure solution at scale and should be the preferred choice by the enterprise. Other options provide acceptable security at small scales, but they should not be put to widespread use, because the resulting vulnerabilities quickly escalate with scale-up. Figure 9 shows a notional plot of the vulnerability associated with different federation options as a function of scale.

The exact curves vary according to how risk and size are measured, but the key idea is that each federation option will have some maximum scale at which its security is acceptable. At the scale of a few individuals sharing information, the risk remains low for all methods, but as the scale increases, the risk

associated with the approaches diverges, with ZTE offering the only acceptable solution at full enterprise scale.

Despite the scalability limitations, all approaches can be viable options. For example, “Weak ID” federation, although not desirable compared to ZTE, is better than the ad hoc or person-to-person approaches, and making all options available provides a smoother continuum of choices based on resources available. A particular enterprise may wish to limit the options to a subset of these or expand to include more. It may be better to adopt a strict “ZTE or nothing” policy to motivate partners to adopt ZTE, even if it means very limited sharing. For cases in which the focus is on sharing, providing all options helps to provide the best security available for the given situation. To manage overall enterprise risk, it may be necessary to limit the absolute number of federation partners or their constituent members that use each non-ZTE type of federation to maintain the overall vulnerability level below a desired threshold.

Figure 12: Notional comparison of federation security at different scales of operation

VII. FEDERATION TRUST CONSIDERATIONS

The technical considerations are important for federation, but having the technology available does not always mean it is a good option. This section takes a different perspective to federation and considers the trade-offs among federation options based on the level of trust that exists with the federation partner. This is usually important to consider when making information or services available to partners and less important when accessing partner information and services. It is assumed that some sharing is desired, regardless of the level of trust. It is also assumed that sharing should be limited, since this is true even within a single enterprise using ZTE. The challenge is how to use the available technology to build upon the real-world trust relationships to enable as many valid sharing behaviors as possible while preventing behaviors that are not valid or desired. In this section, trust is meant to be the degree to which one enterprise believes that the other enterprise will not abuse technical capabilities to violate the agreed-upon terms of the federation agreement. It does not apply to weak security methods, such as username/password authentication, or to weak cryptographic algorithms, such as RC4 or MD5.

Trust can be a tricky issue. For example, if one enterprise trusts the root certificate authority (CA) of the other enterprise, then PKI authentication can work across the enterprises. Thus, it might be tempting for one enterprise to declare, "Enterprise 1 trusts Enterprise 2's root CA," to make security work. This declaration may sound like, "Enterprise 1 trusts Enterprise 2's root CA and therefore Enterprise 1 will encode that trust in its system to allow both systems to work together." However, stating trust for the purpose of interoperability is really saying, "Enterprise 1 will open up a security vulnerability in its system that Enterprise 2 can exploit and will trust Enterprise 2 not to abuse it." Because such trust opens up vulnerabilities, it is important that statements of trust are truly based on existing or established relationships and not simply put in place for convenience. As trust is usually not an all-or-nothing proposition, any trust relationship must also be accompanied by monitoring and measures to identify and block malicious behavior by a partner, which can be summarized as "trust but verify."

For a trusted partner or a partner for whom there is accountability and a long-standing relationship, trusting a partner CA might be a good way to federate. For a situation in which an enterprise wants to share but really does not trust the partner not to abuse additional privilege, there is no longstanding relationship or accountability, or the damage that can be done is greater than any retribution likely to be enacted, this is a risky choice. The key is to first decide what to actually trust and then

use that as a basis to build the system, not to build a system first and then declare trust where needed to make it work. This section provides options for what to do with different trust levels. Of course, trust can be increased through external policies and agreements as well, so these options do not exist in isolation from the real world. However, it is often easier to build a system based on existing trust than to work to establish trust, as changing the real world to match a technical system is often slower than changing a technical system to match the real world.

A. Full Trust

If one enterprise fully trusts another enterprise, federation can be achieved in any number of ways. This could be the case for two divisions within a larger enterprise, a merger between two companies, two parallel systems within a single enterprise, or a hierarchical arrangement in which one group is sharing with another that contains it. In this case, any option will work, and the best choice is usually the one that enables the most integration and sharing or the one that is simplest to set up. In this case, the degree of ZTE compatibility is likely to be the determining factor for choosing a solution.

B. Infrastructure Trust

If one enterprise trusts the high-level security functions, such as certificate authorities and other enterprise level functions, then the best option is one that leverages this existing trust and limits access based on these trusted services. For example, if both enterprises use ZTE, a federation agreement among STSs is a viable choice, as this leverages trusted root CAs and trusted STSs and attribute stores, which are all high-level security functions. In some cases, only a part of a partner's ZTE infrastructure is trusted or the partner's infrastructure supports most but not all ZTE requirements. For example, if the root CA is trusted but the STSs or attribute stores are not trusted, identity-based federation with delegation or a local attribute store for those identities might be a better option. This leverages the existing trust in the identity credentials while providing the claims through the local STS and internally maintained attributes. This provides the function of the untrusted partner STS or attributes in a trusted way. The lack of trust is similar to a lack of technology, as a function that exists but is not trusted is similar to one that is missing. The complementary situation is when a partner STS is trusted while partner identity credentials are not trusted. This case is not common, because weak authentication to an application often correlates with weak authentication to the STS, which means the SAML token cannot be trusted even if the STS and attribute stores are trusted. The best technical solution when identity services are not trusted may be person-to-person federation, which would limit the scale of sharing to individuals or small groups of known entities.

C. Individual Trust

In some cases, certain individuals of a partner organization are trusted even though the organization's security functions are not. In this situation, identity-based federation may be possible for the specific identities who are trusted. For example, a partner with a smart card credential might be technologically acceptable as the public/private key pair can be recorded and used for authentication without the associated certificate. However, this requires no key escrow or some other method be in use by the partner enterprise to duplicate the private key associated with the credential. This method relies explicitly on the strength of the keys associated with the credential and their storage. This fails to achieve the efficiency and scalability of PK, but allows smaller-scale point-to-point trust to be established using specific public/private key pairs. The certificate is replaced by an internal mapping of public keys to identities. Another option is to provide a local credential to the trusted individual. This brings the individual into the enterprise. This is a well-integrated method, but proper enterprise-level vetting often cannot be completed on such an individual due to the organization he works for. An individual in an enterprise may trust an individual from a partner enterprise, but the enterprise may not trust them sufficiently to issue a credential. A local attribute store can record attributes for the trusted individual to provide proper access credentials through a local STS. This requires the STS's applications and services to check not just issuers of certificates but public keys for specific individual requesters. Revocation checking must also be supported within the enterprise for such individual keys. A simpler option for leveraging this trust in the individual and associated credential is delegation. This allows resource owners to assign which applications and services a partner can access instead of providing access based on the partner's attributes. Automatic access based on attributes provides more scalability and automation than delegation, because with delegation all access must be granted manually. This might be an acceptable tradeoff when limited data is to be shared with a small number of individuals in a secure way. In this case, the STS operates as usual and only the applications and services are modified to trust individual keys.

D. No Trust

In some cases, data owners in an enterprise really do not trust a federation individual or his organization not to abuse their access but they still need to provide some access. As they cannot trust individuals, the option of individual trust does not apply. The problem is that any access granted is likely to be abused, such as copying and exfiltration of data or rendering inoperable or untrustworthy any services or devices to which more generous access is granted.

In this case, the idea is to limit the individual to only what is needed when it is needed and to heavily monitor activity to detect suspicious behavior where limits cannot be strictly enforced through technology. Shared information may be fingerprinted or steganographically imprinted for later forensics in case of leaks. One option is to provide a partner individual with a locked-down device that is preconfigured with a credential that has been delegated appropriate access and to let him use this as his only method of access. In this case, the issuing enterprise tightly controls the hardware, software, credentials, and access. This method limits the information to the display on the locked-down device, which can be remotely monitored. If the device is not returned within the timeframe of the valid access, it can be remotely disabled through device management policy. Monitoring of the device and the data accessed can also reveal potential exfiltration and initiate disabling of the machine and termination of delegation privileges. None of the automated technological solutions are viable because they are designed to facilitate automatic sharing of information. Without trust, no automation is desired. For all trust levels, the idea of "trust but verify" underlies all communication. What changes as trust levels decrease is not just what is shared, but the measures taken to prevent, detect, and limit harm from the federation partner.

VIII. CONCLUSIONS

This paper examines secure enterprise federation options from the perspectives of both technical capabilities and trust. The paper reviews six federation methods:

- 1) ZTE federation,
- 2) ZTE-like federation,
- 3) Identity credential federation,
- 4) Weak identity federation,
- 5) Ad hoc federation,
- 6) Person-to-Person federation.

All solutions are acceptable at small scale, but only those near the top of the list scale to the organizational or enterprise level while maintaining acceptable security. The best choice is generally the one closest to the top of the list, but these top-end options require the most infrastructure for proper operation. When less-secure options are chosen, it is important as part of an overall risk management process for the enterprise to monitor and limit the extent of their aggregate use across the enterprise. This limits the scale of the lower options to maintain the assurance requirements of the enterprise. Technological solutions also must be considered in the context of existing trust. A common but flawed solution is to choose the best technical solution and simply declare any required trust by formal edict.

This masks the underlying trust issues and provides these untrusted partners front-door access to enterprise applications and services. A better approach is to first assess the level of trust in a partner enterprise, its security services, and the individuals involved in the sharing, and then build a federation agreement based on those (and only those) solid building blocks. This research is part of a body of work for high-assurance enterprise computing using web services. Elements of this work include bi-lateral end-to-end authentication using PKI credentials for all person and non-person entities, a separate SAML credential for claims-based authorization, full encryption at the transport layer, and a defined federation process. Many of the elements of this work are described in [25]–[33].

ACKNOWLEDGEMENTS

The authors would like to thank the Institute for Defense Analyses for their unflagging support. The publication of this paper does not indicate endorsement by the Department of Defense or IDA, nor should the contents be construed as reflecting the official position of these organizations.

REFERENCES

- [1] Tanenbaum, Andrew S., and Maarten van Steen, 2002, *Distributed Systems: Principles and Paradigms*. Upper Saddle River, NJ: Pearson Prentice Hall. ISBN 0-13-088893-1.
- [2] Magnoni, L., 2015, "Modern Messaging for Distributed Systems" (sic), *Journal of Physics: Conference Series*. 608 (1): 012038. doi:10.1088/1742-6596/608/1/012038. ISSN 1742-6596.
- [3] Hassan N. A., 2019, "Endpoint Defense Strategies," In *Ransomware Revealed*. Apress, Berkeley, CA. https://doi.org/10.1007/978-1-4842-4255-1_4
- [4] Mark Khai Shean Tan, Sigi Goode & Alex Richardson, 2020, "Understanding Negotiated Anti-Malware Interruption Effects on User Decision Quality in Endpoint Security," *Behaviour & Information Technology*, DOI: 10.1080/0144929X.2020.1734087
- [5] Dayna Eidle, Si Ya Ni, Casimer DeCusatis, Anthony Sager, 2017, "Autonomic Security for Zero Trust Networks," Ubiquitous Computing Electronics and Mobile Communication Conference (UEMCON) 2017 IEEE 8th Annual, pp. 288–293.
- [6] DeCusatis C., P. Liengtiraphan, A. Sager, and M. Pinelli, 2016, "Implementing Zero Trust Cloud Networks with Transport Access Control and First Packet Authentication," 2016 IEEE International Conference on Smart Cloud (SmartCloud), New York, NY, pp. 5–10, DOI: 10.1109/SmartCloud.2016.22.
- [7] Simpson, William R. *Enterprise Level Security – Securing Information Systems in an Uncertain World*. CRC Press, Taylor & Francis Group, Auerbachian Publications, May 2016, 397 pp., ISBN 9781498764452
- [8] Simpson, William R., and Kevin E. Foltz, *Enterprise Level Security 2: Advanced Techniques for Information Technology in an Uncertain World*. CRC Press, Taylor & Francis Group, September 2020, 338 pp., ISBN 9781003080787.
- [9] Science Direct, Editors: Jason Andress, Steve Winterfeld, *Cyber Warfare*, ISBN 9781597496377, 2011, Jason Andress, Steve Winterfeld, Chapter 10 - Computer Network Defense, Pages 179-191, <http://www.sciencedirect.com/science/article/pii/B9781597496377000101>, last accessed on 24 November 2020.
- [10] TechTarget.com. "Backdoor (Computing)," <https://searchsecurity.techtarget.com/definition/back-door>, last accessed 22 November 2019.
- [11] Jack Wallen, September 2011, "Five Free, Dead-Easy IP Traffic Monitoring Tools," *Tech Republic*, <https://www.techrepublic.com/blog/five-apps/five-free-dead-easy-ip-traffic-monitoring-tools/>, last accessed 22 November 2019.
- [12] Moskovitch R, and Y. Elovici., 2008, "Unknown Malicious Code Detection – Practical Issues," In *Proceedings of the 7th European Conference on Warfare and Security* (ECIW '08), Plymouth, UK.
- [13] Begel, A., S. McCanne, and S. L. Graham, 1999, "BPF+: Exploiting Global Data-Flow Optimization in a Generalized Packet Filter Architecture," In *Proceedings of ACM SIGCOMM*, Cambridge, MA, USA, pp. 123–134.
- [14] McDaniel, M., and M. H. Heydari, January 2003, "Content Based File Type Detection Algorithms," In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences, IEEE*, ISBN: 0-7695-1874-5, DOI:10.1109/HICSS.2003.1174905.
- [15] Fisk, Mike, and George Varghese, May 2001, "Fast Content-Based Packet Handling for Intrusion Detection," Los Alamos National Lab Computing Communications and Networking Division, <https://apps.dtic.mil/dtic/tr/fulltext/u2/a406413.pdf>, last accessed 22 November 2019.
- [16] Song, Jian, and Yanchun Zhang, 2007, "Architecture of a Web Accelerator for Wireless Networks," In *Proceedings of the Thirtieth Australasian Conference on Computer Science*, Volume 62 (ACSC '07), Gillian Dobbie (Ed.), Australian Computer Society, Inc., Darlinghurst, Australia, pp. 125–129.
- [17] Shin-ichi Kuribayashi, 2013, "Improving Quality of Service and Reducing Power Consumption with WAN Accelerator in Cloud Computing Environments," *International Journal of Computer Networks & Communications (IJCNC)*, Vol. 5, No.1.
- [18] Afzal, S., and G. Kavitha. 2019, "Load balancing in cloud computing – A hierarchical taxonomical classification." *Journal of Cloud Computing*, 8, No. 22, <https://journalofcloudcomputing.springeropen.com/articles/10.1186/s13677-019-0146-7>
- [19] X.509 Standards, DoDI 8520.2, Public Key Infrastructure (PKI) and Public Key (PK) Enabling, 24 May 2011 also Public Key Cryptography Standard, PKCS #1 v2.2: RSA Cryptography Standard, RSA Laboratories, October 27, 2012
- [20] TLS Family Internet Engineering Task Force (IETF) Standards, Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS) Versions 1.2 and Earlier, 2015-03-09 also RFC 5246 The Transport Layer Security (TLS) Protocol Version 1.2, 2008-08
- [21] Organization for the Advancement of Structured Information Standards (OASIS) Open set of Standards, N. Ragouzis, et al., March 2008, "Security Assertion Markup Language (SAML) V2.0 Technical Overview," OASIS Committee Draft.

International Journal of Emerging Technology and Advanced Engineering

Website: www.ijetae.com (E-ISSN 2250-2459, Scopus Indexed, ISO 9001:2008 Certified Journal, Volume 11, Issue 05, May 2021)

- [22] IFIP Working Group 11.5, William List and Rob Melville, 1994, "Integrity in Information Systems," *Computers and Security*, Volume 13, Issue 4, pp. 295–301, doi:10.1016/0167-4048(94)90018-3
- [23] Simpson, William R., and Coimbatore Chandrasekaran, 2011 "An Agent Based Monitoring System for Web Services," CCCT2010, Volume II, pp. 84–89, Orlando, FL, April 2011
- [24] Simpson, William R., and Coimbatore Chandrasekaran, 2011, "A Multi-Tiered Approach to Enterprise Support Services," 1st International Conference on Design, User Experience, and Usability, part of the 14th International Conference on Human-Computer Interaction (HCI 2011), 10 pp. Orlando, FL, July 2011. Also published in: A. Marcus (Ed.): Design, User Experience, and Usability, Pt I, HCI 2011, LNCS 6769, pp. 388–397, 2011. © Springer-Verlag Berlin Heidelberg.
- [25] Simpson, William R., and Kevin E. Foltz, 2017, "Data Mediation with Enterprise Level Security," In *Proceedings of the 21th World Multi-Conference on Systemics, Cybernetics and Informatics: WMSCI 2017*, Orlando, Florida, July 8–11, 2017, 6 pages.
- [26] Simpson, William R., Kevin E. Foltz, 2017, "Escalation of Access and Privilege with Enterprise Level Security," In *Proceedings of the 22nd International Command and Control Research and Technology Symposium (ICCRTS)*, Los Angeles, CA., September 2017, ISBN: 978-0-9997246-0-6.
- [27] Simpson, William R. and Kevin E. Foltz, 2018, "Electronic Record Key Management for Digital Rights Management," In Sio-Long Ao, et. al. (eds.), *IAENG Transactions on Engineering Sciences, Special Issue of the Association of Engineers Conferences 2016*, Volume II, pp. 475–488, World Scientific Publishing, Singapore, 2018. ISBN 978-981-3230-76-7,
- [28] Simpson, William R., and Kevin E. Foltz, February 2018 "Secure Identity for Enterprises," *IAENG International Journal of Computer Science*, Vol. 45, no. 1, pp. 142–152. ISSN: 1819-656X.
- [29] Simpson, William R. and Kevin E. Foltz, November 2018, "Cloud Security and Scalability," In *Proceedings of the 8th International Conference on Electronics, Communications and Networks (CECNet 2018)*, Volume 1, pp. 27, Bangkok, Thailand.
- [30] Simpson, William R. and Kevin E. Foltz, December 2018, "Insider Threat Metrics in Enterprise Level Security," *IAENG International Journal of Computer Science*, Vol. 45, No. 4, pp. 610–622, ISSN: 1819-656X.
- [31] Simpson W., and K. Foltz K., October 2015, "Maintaining High Assurance in Asynchronous Messaging," In *Lecture Notes in Engineering and Computer Science, Proceedings World Congress on Engineering and Computer Science 2015*, Volume 1, pp. 178–183, Berkeley, CA.
- [32] Simpson, William R., and Kevin E. Foltz, 2018, "Mobile Ad-hoc for Enterprise Level Security," In *Lecture Notes in Engineering and Computer Science: Proceedings of The World Congress on Engineering and Computer Science 2018*, October 23–25, 2018, San Francisco, CA, USA, pp. 172–177.
- [33] Simpson, William R., and Kevin E. Foltz, 2021, "Resolving Network Defense Conflicts with Zero Trust Architectures and Other End-To-End Paradigms," *International Journal of Network Security & Its Applications (IJNSA)* Vol. 13, No. 1.