NIST SPECIAL PUBLICATION 1800-35

Implementing a Zero Trust Architecture:

High-Level Document

Oliver Borchert
Gema Howell
Alper Kerman
Scott Rose
Murugiah Souppaya
National Institute of Standards and Technology

Jason Ajmo
Yemi Fashina
Parisa Grayeli
Joseph Hunt
Jason Hurlburt
Nedu Irrechukwu
Joshua Klosterman
Oksana Slivina
Susan Symington
Allen Tan
The MITRE Corporation

Karen Scarfone
Scarfone Cybersecurity

William Barker
Dakota Consulting

Peter Gallagher
Aaron Palermo
Appgate

Madhu Balaji
Adam Cerini
Rajarshi Das
AWS (Amazon Web Services)

Jacob Barosin
Kyle Black
Scott Gordon
Jerry Haskins
Keith Luck
Dale McKay
Sunjeet Randhawa
Broadcom

Brian Butler
Mike Delaguardia
Matthew Hyatt
Randy Martin
Peter Romness
Cisco

Corey Bonnell
Dean Coclin
DigiCert

Ryan Johnson
Dung Lam
Darwin Tolbert
F5

Tim Jones
Tom May
Forescout

Christopher Altman
Alex Bauer
Marco Genovese
Google Cloud

Andrew Campagna
John Dombroski
Adam Frank
Nalini Kannan
Priti Patil
Harmeet Singh
Mike Spisak
Krishna Yellepeddy
IBM

Nicholas Herrmann
Corey Lund
Farhan Saifudin
Ivanti

Madhu Dodda
Tim LeMaster
Lookout

Ken Durbin
James Elliott
Earl Matthews
David Pricer
Mandiant

Joey Cruz
Tarek Dawoud
Carmichael Patton
Alex Pavlovsky
Brandon Stephenson
Clay Taylor
Microsoft

Bob Lyons
Vinu Panicker
Okta

Peter Bjork
Hans Drolshagen
Omnissa

Imran Bashir
Ali Haider
Nishit Kothari
Sean Morgan
Seetal Patel
Norman Wong
Palo Alto Networks

Zack Austin
Shawn Higgins
Rob Woodworth
PC Matic

Mitchell Lewars
Bryan Rosensteel
Ping Identity

Don Coltrain
Wade Ellery
Deborah McGinn
Radiant Logic

Frank Briguglio
Ryan Tighe
SailPoint

Chris Jensen
Joshua Moll
Tenable

Jason White
Trellix, Public Sector

Joe Brown
Gary Bradt
Zimperium

Jeffrey Adorno
Syed Ali
Bob Smith
Zscaler

June 2025

FINAL

This publication is available free of charge from https://doi.org/10.6028/NIST.SP.1800-35

NIST NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

DISCLAIMER

Certain commercial entities, equipment, products, or materials may be identified by name or company logo or other insignia in order to acknowledge their participation in this collaboration or to describe an experimental procedure or concept adequately. Such identification is not intended to imply special status or relationship with NIST or recommendation or endorsement by NIST or NCCoE; neither is it intended to imply that the entities, equipment, products, or materials are necessarily the best available for the purpose.

National Institute of Standards and Technology Special Publication 1800-35, Natl. Inst. Stand. Technol. Spec. Publ. 1800-35, 55 pages, (June 2025), CODEN: NSPUE2

NIST TECHNICAL SERIES POLICIES

Copyright, Use, and Licensing Statements

NIST Technical Series Publication Identifier Syntax

AUTHOR ORCHID IDS

Oliver Borchert: 0009-0006-1880-0542
Gema Howell: 0000-0002-0428-5045
Alper Kerman: 0009-0000-5880-8369
Scott Rose: 0000-0002-3105-7427
Murugiah Souppaya: 0000-0002-8055-8527
Karen Scarfone: 0000-0001-6334-9486
William Barker: 0000-0002-4113-8861

FEEDBACK

You can view or download the final guide at the NCCoE ZTA project page.

Comments on this publication may be submitted to: nccoe-zta-project@list.nist.gov.

All comments are subject to release under the Freedom of Information Act.

National Cybersecurity Center of Excellence
National Institute of Standards and Technology
100 Bureau Drive
Mailstop 2002
Gaithersburg, MD 20899
Email: nccoe@nist.gov

NATIONAL CYBERSECURITY CENTER OF EXCELLENCE

The National Cybersecurity Center of Excellence (NCCoE), a part of the National Institute of Standards and Technology (NIST), is a collaborative hub where industry organizations, government agencies, and academic institutions work together to address businesses' most pressing cybersecurity issues. This public-private partnership enables the creation of practical cybersecurity solutions for specific industries, as well as for broad, cross-sector technology challenges. Through consortia under Cooperative Research and Development Agreements (CRADAs), including technology partners—from Fortune 50 market leaders to smaller companies specializing in information technology security—the NCCoE applies standards and best practices to develop modular, adaptable example cybersecurity solutions using commercially available technology. The NCCoE documents these example solutions in the NIST Special Publication 1800 series, which maps capabilities to the NIST Cybersecurity Framework and details the steps needed for another entity to re-create the example solution. The NCCoE was established in 2012 by NIST in partnership with the State of Maryland and Montgomery County, Maryland.

To learn more about the NCCoE, visit https://www.nccoe.nist.gov/. To learn more about NIST, visit https://www.nist.gov.

NIST CYBERSECURITY PRACTICE GUIDES

NIST Cybersecurity Practice Guides (Special Publication 1800 series) target specific cybersecurity challenges in the public and private sectors. They are practical, user-friendly guides that facilitate the adoption of standards-based approaches to cybersecurity. They show members of the information security community how to implement example solutions that help them align with relevant standards and best practices and provide users with the materials lists, configuration files, and other information they need to implement a similar approach.

The documents in this series describe example implementations of cybersecurity practices that businesses and other organizations may voluntarily adopt. These documents do not describe regulations or mandatory practices, nor do they carry statutory authority.

ABSTRACT

A zero trust architecture (ZTA) enables secure authorized access to enterprise resources that are distributed across on-premises and multiple cloud environments, while enabling a hybrid workforce and partners to access resources from anywhere, at any time, from any device in support of the organization's mission. This NIST Cybersecurity Practice Guide explains how organizations can implement ZTA consistent with the concepts and principles outlined in NIST Special Publication (SP) 800-207, Zero Trust Architecture. The NCCoE worked with 24 collaborators under Cooperative Research and Development Agreements (CRADAs) to integrate commercially available technology to build 19 ZTA example implementations and demonstrate a number of common use cases. The Guide includes detailed technical information on each example ZTA implementation, providing models that organizations can emulate. The guide also summarizes best practices and lessons learned from the implementations and integrations to make it easier and more cost-effective to implement ZTA. Additionally, this guide includes mappings of ZTA principles and technologies to commonly used security standards and guidelines.

KEYWORDS

enhanced identity governance (EIG); identity, credential, and access management (ICAM); microsegmentation; secure access service edge (SASE); software-defined perimeter (SDP); zero trust; zero trust architecture (ZTA).

ACKNOWLEDGMENTS

We are grateful to the following individuals for their generous contributions of expertise and time.

- Appgate: Jason Garbis, Adam Rose, Jonathan Roy
- AWS (Amazon Web Services): Conrad Fernandes*, Harrison Holstein, Quint Van Deman
- Broadcom: Andrew Babakian*, Genc Domi*, Paul Mancuso, Eric Michael, Dennis Moreau*, Wayne Pauley*, Jacob Rapp*, Lewis Shepherd
- Cisco: Ken Andrews, Robert Bui, Leo Lebel, Tom Oast, Aaron Rodriguez, Kelly Sennett, Steve Vetter, Micah Wilson
- F5: Daniel Cayer, David Clark, Jay Kelley, Darrell Pierson
- Forescout: Yejin Jang*, Neal Lucier*
- Google Cloud: Tim Knudson*
- IBM: Nilesh Atal, Himanshu Gupta, Lakshmeesh Hegde, Sharath Math, Naveen Murthy, Nikhil Shah, Deepa Shetty, Harishkumar Somashekaraiah
- IT Coalition: Aaron Cook, Vahid Esfahani*, Jeff Laclair, Ebadullah Siddiqui*, Musumani Woods*
- Ivanti: Patty Arcano, Jeffery Burton, Jay Dineshkumar
- Lookout: Tyler Croak, Jeff Gilhool, Hashim Khan*
- Microsoft: Thomas Detzner, Ehud Itshaki, Janet Jones, Hemma Prafullchandra*, Enrique Saggese, Sarah Young
- MITRE: Eileen Division*, Spike E. Dog*, Sallie Edwards*, Ayayidjin Gabiam, Jolene Loveless*, Karri Meldorf, Kenneth Sandlin, Lauren Swan, Jessica Walton*
- NIST: Mike Bartock, Julie Chua, Douglas Montgomery, Cherilyn Pascoe, Michael Powell, Kevin Stine
- Okta: Brian Dack, Sean Frazier, Naveed Mirza, Kelsey Nelson, Ron Wilson
- Omnissa: Keith Luck*
- PC Matic: Andy Tuch
- Ping Identity: Ivan Anderson, Aubrey Turner
- Radiant Logic: Bill Baz, Rusty Deaton, John Petrutiu, Lauren Selby
- SailPoint: Peter Amaral, Jim Russell, Esteban Soto
- Tenable: Jeremiah Stallcup
- Zimperium: Dan Butzer, Jim Kovach*, Kern Smith
- Zscaler: Jeremy James, Lisa Lorenzin*, Matt Moulton, Patrick Perry

** Former employee; all work for this publication was done while at that organization*

Special thanks to all who reviewed and provided feedback on this document.

The Technology Collaborators who participated in this project submitted their capabilities in response to a notice in the Federal Register. Respondents with relevant capabilities or product components were invited to sign a Cooperative Research and Development Agreement (CRADA) with NIST, allowing them to participate in a consortium to build this example solution.[1] We worked with:

Technology Collaborators		
Appgate	IBM	PC Matic
AWS	Ivanti	Ping Identity
Broadcom	Lookout	Radiant Logic
Cisco	Mandiant	SailPoint
DigiCert	Microsoft	Tenable
F5	Okta	Trellix
Forescout	Omnissa	Zimperium
Google Cloud	Palo Alto Networks	Zscaler

DOCUMENT CONVENTIONS

The terms “shall” and “shall not” indicate requirements to be followed strictly to conform to the publication and from which no deviation is permitted. The terms “should” and “should not” indicate that among several possibilities, one is recommended as particularly suitable without mentioning or excluding others, or that a certain course of action is preferred but not necessarily required, or that (in the negative form) a certain possibility or course of action is discouraged but not prohibited. The terms “may” and “need not” indicate a course of action permissible within the limits of the publication. The terms “can” and “cannot” indicate a possibility and capability, whether material, physical, or causal.

[1] Note that after the VMware End User Computing division products were implemented at the NCCoE, VMware was acquired by Broadcom, and then the VMware End User Computing Division was divested and reformed under a new entity, Omnissa LLC. Symantec was also previously acquired by Broadcom.

PATENT DISCLOSURE NOTICE

NOTICE: The Information Technology Laboratory (ITL) has requested that holders of patent claims whose use may be required for compliance with the guidance or requirements of this publication disclose such patent claims to ITL. However, holders of patents are not obligated to respond to ITL calls for patents and ITL has not undertaken a patent search in order to identify which, if any, patents may apply to this publication.

As of the date of publication and following call(s) for the identification of patent claims whose use may be required for compliance with the guidance or requirements of this publication, no such patent claims have been identified to ITL.

No representation is made or implied by ITL that licenses are not required to avoid patent infringement in the use of this publication.

Contents

Executive Summary 1

1 Introduction to the Guide 2

1.1 Audience 2

1.2 Scope 2

1.3 How to Use This Guide 2

2 Project Overview 4

2.1 Motivation for the Project 4

2.2 Challenges in Implementing ZTA 4

2.3 Collaborators and Their Contributions 6

3 Architecture and Builds 6

3.1 General ZTA Reference Architecture 7

3.2 EIG Crawl Phase Reference Architecture 9

3.3 EIG Run Phase Reference Architecture 10

3.4 SDP, Microsegmentation, and SASE Reference Architecture 10

3.5 ZTA Laboratory Physical Architecture 11

3.6 Builds Implemented 12

4 Build Implementation Instructions 16

5 General Findings 19

5.1 EIG Crawl Phase Findings 19

5.2 EIG Run Phase Findings 20

5.3 SDP, Microsegmentation, and SASE Phase Findings 22

6 Functional Demonstrations 23

6.1 Demonstration Methodology 23

6.2 Demonstration Use Cases 24

6.2.1 Use Case A: Discovery and Identification 24

6.2.2 Use Case B: Enterprise-ID Access 25

6.2.3 Use Case C: Collaboration: Federated-ID Access 25

6.2.4 Use Case D: Other-ID Access 26

6.2.5 Use Case E: Guest: No-ID Access 26

6.2.6 Use Case F: Confidence Level 27

6.2.7 Use Case G: Service-Service Interaction 28
6.2.8 Use Case H: Data Level Security Scenarios 28
6.3 Functional Demonstration Results 29
6.3.1 Demonstration Result Summaries 29
6.3.2 Demonstration Results in Full 31
7 Risk and Compliance Management 33
7.1 Risks Addressed by the ZTA Reference Architecture 33
7.2 ZTA Security Mappings 35
8 Zero Trust Journey Takeaways 36
8.1 Discover and Inventory the Existing Environment 36
8.2 Formulate Access Policy to Support the Mission and Business Use Cases 37
8.3 Identify Existing Security Capabilities and Technology 38
8.4 Eliminate Gaps in Zero Trust Policy and Processes by Applying a Risk-Based Approach Based on the Value of Data 38
8.5 Implement ZTA Components (People, Process, and Technology) and Incrementally Leverage Deployed Security Solutions 39
8.6 Verify the Implementation to Support Zero Trust Outcomes 39
8.7 Continuously Improve and Evolve Due to Changes in Threat Landscape, Mission, Technology, and Requirements 40
Appendix A List of Acronyms 41
Appendix B References 44
Appendix C Change Log 45

List of Figures

Figure 3-1 General ZTA Reference Architecture 8
Figure 3-2 EIG Crawl Phase Reference Architecture 10
Figure 3-3 Physical Architecture of ZTA Lab 11

List of Tables

Table 4-1 Mapping of Builds to Online Details Regarding Architecture Descriptions and Implementation Instructions 17

Table 6-1 Mapping of Builds to Online Details Regarding Architecture Descriptions and Functional Demonstration Results 31

Executive Summary

A zero trust architecture (ZTA) is an enterprise cybersecurity architecture that is based on zero trust principles, such as those outlined in NIST Special Publication (SP) 800-207, *Zero Trust Architecture* [1], and that is designed to prevent data breaches and limit internal lateral movement. A ZTA can help your organization protect its data and resources no matter where they are located. A ZTA can also enable your workforce, contractors, partners, and other authorized parties to securely access the data and resources they need from anywhere at any time. ZTA implements a risk-based approach to cybersecurity—continuously evaluating and verifying conditions and requests to decide which access requests should be permitted, then ensuring that each access is properly safeguarded commensurate with risk. Because of their effectiveness against both internal and external threats, this architecture is increasingly being adopted, and some organizations are required to use a ZTA.

There is no single approach for each organization to migrate to ZTA. Therefore, the NIST National Cybersecurity Center of Excellence (NCCoE) worked with 24 technology providers to demonstrate practical implementation of ZTA principles from NIST SP 800-207. Together, we have built and implemented 19 example ZTA solutions in lab environments, leveraging the technology from our collaborators. For each of the example ZTAs, we have outlined detailed technical information, including architecture, sample technologies leveraged, specific configurations and integrations of technologies, and use cases and scenarios demonstrated.

We have also created mappings between the example ZTA security capabilities and the NIST Cybersecurity Framework (CSF) versions 1.1 and 2.0 [2][3], NIST SP 800-53r5 [4], and NIST critical software security measures. These mappings were developed to support why and how organizations can implement ZTA.

This guide is intended to help your organization gradually evolve existing environments and technologies into a ZTA over time. It provides practical information that you can use to develop your ZTA roadmap, including models you can emulate and examples of how to best leverage existing technology infrastructure. The lessons we have learned from our demonstrations can benefit your organization by saving time and resources.

By utilizing this guide, your organization can be better positioned to implement a ZTA that achieves the following:

- Supports user access to resources regardless of user location or device (managed or unmanaged)
- Protects sensitive information and other business assets and processes regardless of their location (on-premises or cloud-based)
- Limits breaches by making it harder for attackers to move through an environment and by addressing the insider threat (insiders are not automatically trusted)
- Performs continuous, real-time monitoring, logging, and risk-based assessment and enforcement of corporate policy

This high-level document serves as introductory reading with insight into the project effort. For in-depth details, please refer to the full document in web format.

1 Introduction to the Guide

This paper outlines the guidelines for organizations as they implement zero trust architecture (ZTA). The implementation best practices and lessons learned were identified through a collaborative project at the NCCoE that developed, demonstrated, and documented example ZTAs. The NCCoE and its collaborators have used commercially available technology in lab environments to build 19 interoperable, open standards-based ZTA implementations ("builds") that align with the concepts and principles in NIST SP 800-207, *Zero Trust Architecture* [1]. The implementations include ZTA approaches for enhanced identity governance (EIG), software-defined perimeter (SDP), microsegmentation, and secure access service edge (SASE).

1.1 Audience

The primary audience for this guide is organizations looking to implement ZTA. The document assumes an existing level of cybersecurity knowledge and capabilities to deploy ZTA components and supporting components for data security, endpoint security, identity and access management, and security analytics. The enterprises are also assumed to have critical resources that require protection, some of which are located on-premises and others of which are in the cloud; and a requirement to provide partners, contractors, guests, and employees, both local and remote, with secure access to these critical resources. For a full list of assumptions for this project, see our supplemental Assumptions documentation. This paper is not specific to federal agency audiences.

Readers of this guide should already be familiar with ZTA basics and the topics covered in NIST SP 800-207, *Zero Trust Architecture* [1].

1.2 Scope

The scope of this guide is implementing a ZTA for a conventional, general-purpose enterprise IT infrastructure with support for traditional IT resources such as laptops, desktops, servers, mobile devices, and other systems with credentials. Discovery of resources, assets, communication flows, and other elements is also within scope. The focus is on using the ZTA to protect access to enterprise data, regardless of who initiates the access request (e.g., enterprise employees, partners, contractors, or corporate network guests), from where the access request is initiated (e.g., from the corporate network, a branch office, or the public internet), or where the resources are located (e.g., on-premises or in the cloud).

ZTA for industrial control systems, operational technology (OT) environments, and Internet of Things (IoT) devices are explicitly out of scope for this project. For information on other related NCCoE projects, see [5][6]. Addressing the risk and policy requirements of discovering and classifying data [7] is also out of scope.

1.3 How to Use This Guide

This guide offers two content formats: the "High-Level Document in PDF Format" (this document) and the "Full Document in Web Format." The document in PDF format is meant to serve as an introduction to the project, including a high-level summary of the project goals, ZTA reference architecture, ZTA implementations, and findings. The document in the web format provides in-depth details in terms of

technologies leveraged, their specific integrations and configurations, and the use cases and scenarios demonstrated. The web format document also contains information on the implemented security capabilities and their mappings to the NIST CSF versions 1.1 and 2.0 [2][3], NIST SP 800-53r5 [4], and NIST critical software security measures.

Readers are encouraged to begin by reading the document in PDF format (this document) to gain high-level insight into the project. Readers may then drill down from this document into the deeper sections of the linked online document in web format to access in-depth information as needed. Therefore, this document is organized as follows:

- Section 2 provides an overview of the NCCoE's "Implementing a Zero Trust Architecture" project from the viewpoints of motivation for the project, challenges in implementing ZTA, project execution and implementation approach, as well as collaborating organizations and their contributions to the project.
- Section 3 discusses the reference architectures considered for demonstrating various types of ZTA deployment approaches used across the 19 implementations built. It also lists the technology products, along with out-of-the-box capabilities used in each build. Furthermore, this section provides information regarding the NCCoE lab's physical architecture platform used to implement the builds.
- Section 4 lists 19 example implementations in a table format with relevant columns that identify technology products and capabilities used as "Policy Engines/Policy Decision Points," as well as ZTA deployment approaches used in each implementation. Also, additional table columns provide links to details available in web format with respect to build architecture, technologies used, and flow diagrams, including instructions for each implementation.
- Section 5 explores the noteworthy findings and conclusions recorded throughout the demonstration of each ZTA deployment approach across 19 unique lab implementations.
- Section 6 discusses the essence of functional demonstrations scoped for the project from the viewpoints of demonstration methodology, use cases, and scenarios. It also lists the functional demonstration results for each implementation, both in summary and fully detailed formats.
- Section 7 provides information regarding each build's implemented security capabilities and their mappings to the NIST CSF versions 1.1 and 2.0, NIST SP 800-53r5, and NIST critical software security measures.
- Section 8 concludes this document by sharing a list of takeaways as recommended steps for a zero trust journey, intended for organizations that are considering ZTA adoption for their environments.

ZTA implementers and others seeking detailed information on designing and deploying ZTA solutions are encouraged to read all sections of the guide, as well as utilize the wealth of additional resources linked to throughout those sections.

Cybersecurity professionals, compliance professionals, and others who are primarily concerned with how ZTA solutions relate to the CSF, NIST SP 800-53, and NIST critical software security measures should focus on Section 7 and the resources it links to.

Anyone interested primarily in the lessons learned from the project should focus on the takeaways provided in Section 8.

2 Project Overview

2.1 Motivation for the Project

Protecting enterprise data and resources has become increasingly challenging. Many users need access from anywhere, at any time, from any device to support the organization's mission. Data is created, stored, transmitted, and processed across different organizations' environments, which are distributed across on-premises and multiple clouds to meet ever-evolving business use cases. It is no longer feasible to simply protect data and resources at the perimeter of the enterprise environment or to assume that all users, devices, applications, and services within it can be trusted.

A ZTA enables secure authorized access to assets—machines, applications, and services running on them, and associated data and resources—whether located on-premises or in the cloud, for a hybrid workforce and partners based on an organization's defined access policy. For each access request, ZTA explicitly verifies the context available at access time—this includes both static user profile information or non-person entity information, such as the requester's identity and role; and dynamic information, such as geolocation, the requesting device's health and credentials, the sensitivity of the resource, access pattern anomalies, and whether the request is warranted and in accordance with the organization's business process logic. If the defined policy is met, a secure session is created to protect all information transferred to and from the resource. A real-time, risk-based assessment of resource access and access pattern anomaly detection with continuous policy evaluation is performed to establish and maintain the access. A ZTA can also protect organizations from non-organizational resources that their users and applications may connect to, helping to stop threats originating from outside of the organization's control.

NCCoE has collaborated with ZTA technology providers to build numerous example ZTA solutions and demonstrate their ability to meet the tenets of ZTA described in NIST SP 800-207. The goal of the solutions is to enforce corporate security policy dynamically and in near-real-time to restrict access to authenticated, authorized users, devices, and non-person entities while flexibly supporting a complex set of business outcomes involving both remote and on-premises workforces, use of the cloud, partner collaboration, and support for contractors. The example solutions are designed to demonstrate the ability to protect against and detect attacks and malicious insiders. They showcase the ability of ZTA products to interoperate with existing enterprise and cloud technologies while trying to minimize the impact on end-user experience.

The project can help organizations plan how to evolve their existing enterprise environments to ZTA, starting with an assessment of their current resources, strengths, and weaknesses, and setting milestones along a path of continuous improvement, gradually bringing them closer to achieving the ZTA goals they have prioritized based on risk, cost, resources, and their unique mission. The goal is to enable organizations to thoughtfully apply ZTA controls that best protect their business while enabling them to operate as they need to.

2.2 Challenges in Implementing ZTA

Throughout this project, numerous challenges organizations may face in implementing ZTA have been identified, including the following:

- **Organization buy-in and support, such as:**
    - Perception that ZTA is suited only for large organizations and requires significant investment, rather than understanding that ZTA is a set of guiding principles suitable for organizations of any size
    - Concern that ZTA might negatively impact the operation of the environment or end-user experience
    - Lack of resources to develop necessary policies and a pilot or proof-of-concept implementation needed to inform a transition plan
    - Leveraging existing investments and balancing priorities while making progress toward a ZTA via modernization initiatives
    - Lack of understanding regarding what additional skills and training administrators, security personnel, operators, end users, and policy decision-makers may require
- **Missing foundational pieces, such as:**
    - Lack of adequate asset inventory and management needed to fully understand the business applications, assets, and processes that need to be protected, with no clear understanding of the criticality of these resources
    - Lack of adequate digital definition, management, and tracking of user roles across the organization needed to enforce fine-grained, need-to-know access policy for specific applications and services
    - Lack of visibility of the organization's communications and usage patterns—limited understanding of the transactions that occur between an organization's subjects, assets, applications, and services, and absence of the data necessary to identify these communications and their specific flows
    - Lack of information regarding everything that encompasses the organization's attack surface. Organizations can usually address threats with traditional security tools in the layers that they currently manage and maintain, such as networks and applications, but elements of a ZTA may extend beyond their normal purview.
- **Technical challenges, such as:**
    - Integrating various types of commercially available technologies of varying maturities, assessing capabilities, and identifying technology gaps to build a ZTA
    - Lack of a standardized mechanism to distribute, manage, and enforce security policy, causing organizations to face a fragmented policy environment
    - Lack of common understanding and language of ZTA across the community and within the organization, gauging the organization's ZTA maturity, determining which ZTA approach is most suitable for the business, and developing an implementation plan

There is not a single ZTA that fits all. ZTAs need to be designed and implemented for each organization based on the organization's requirements and risk tolerance, as well as its existing invested technologies and environments. The appropriate logical architecture for a given organization's ZTA will depend on that organization's requirements and technologies.

2.3 Collaborators and Their Contributions

The NCCoE prepared a Federal Register Notice [8] inviting technology providers to provide products and/or expertise to compose example ZTAs. Cooperative Research and Development Agreements (CRADAs) were established with qualified respondents. Collaborators' components have been composed into numerous example implementations (i.e., builds). With 24 collaborators participating in the project, the build teams that were assembled sometimes included vendors that offered overlapping capabilities. We made an effort to showcase capabilities from each vendor when possible. In other cases, we consulted with the collaborators to have them work out a solution.

Each of the technology partners and collaborators participating in the project has provided descriptions of the relevant products and capabilities they bring to this ZTA effort. The descriptions can be found in our supplemental documentation of Collaborators and Their Contributions.

The NCCoE does not certify, validate, or endorse products or services. We demonstrate the capabilities that can be achieved by using participants' contributed technology. Your organization's information security experts should identify the products that will best integrate with your existing tools and IT system infrastructure. Your organization can adopt this solution or one that adheres to these guidelines entirely, or you can use this guide as a starting point for tailoring and implementing parts of a solution.

3 Architecture and Builds

This project began with a clean laboratory environment that we populated with various applications and services that would be expected in a typical enterprise to create several baseline enterprise architectures. Examples include security information and event management systems (SIEMs), vulnerability scanning and assessment tools, security validation tools, and discovery tools.

Next, we used a phased approach to develop example ZTA solutions. This approach was designed to represent how we believe most enterprises will evolve their enterprise architecture toward ZTA, i.e., by starting with their already-existing enterprise environment and gradually adding or adapting capabilities. Our first implementations with minimum viable solutions were EIG deployments because the identity-based controls provided by EIG are foundational components of ZTA. We called this phase of the project the *EIG crawl phase,* which did not include cloud capabilities, and it was followed by the *EIG run phase*, where we added cloud capabilities.

We gradually deployed additional functional components and capabilities to address an increasing number of ZTA requirements and deployed microsegmentation, SDP, and SASE approaches.

Given the importance of discovery to the successful implementation of a ZTA, we initially deployed it to continuously observe the environment and use those observations to audit and validate the documented baseline map on an ongoing basis. Because we had instantiated the baseline environment ourselves, we already had a good initial understanding of it. However, we were able to use the discovery tools to audit and validate what we deployed and provisioned, correlate known data with information reported by the tools, and use the tool outputs to formulate an initial zero trust policy, ultimately ensuring that observed network flows correlate to static policies.

The builds described in this document are examples with the understanding that there is no single approach for migrating to ZTA that is best for all enterprises; ZTA is a set of concepts and principles, not

a set of technical specifications that can be complied with. The objective, instead, is continuous improvement of access control processes and policies in accordance with the principles of ZTA

This section provides information on the project's ZTA builds and the underlying architectures they implemented.

3.1 General ZTA Reference Architecture

Figure 3-1 depicts the high-level logical architecture of a general ZTA reference design. It consists of three types of core components: Policy Engine (PE), Policy Administrator (PA), and Policy Enforcement Point (PEP), as well as several supporting components that assist the policy engine in making its decisions by providing data and policy rules related to areas such as identity, credential, and access management (ICAM); endpoint security; security analytics; data security; and resource protection. Specific capabilities that fall into each of these supporting component categories are discussed in more detail in our supplemental documentation for General ZTA Reference Architecture. The various sets of information, either generated via policy or collected by the supporting components and used as input to ZTA policy decisions, are referred to as policy information points (PIPs). Although the simplicity of the architecture may seem to imply that the supporting components are simple plug-ins that respond in real-time to the PDP, in many cases, the ICAM, endpoint detection and response (EDR)/endpoint protection platform (EPP), security analytics, and data security PIPs will each represent complex infrastructures. Some ZTA logical component functions may be performed by multiple hardware or software components, or a single software component may perform multiple logical functions.

Subjects (human users, devices, applications, servers, and other non-human entities that request information from resources on premises or in the cloud) request and receive access to enterprise resources via the ZTA. Human subjects are authenticated. Non-human subjects are both authenticated and protected by endpoint security. Enterprise resources may be located on-premises or in the cloud.

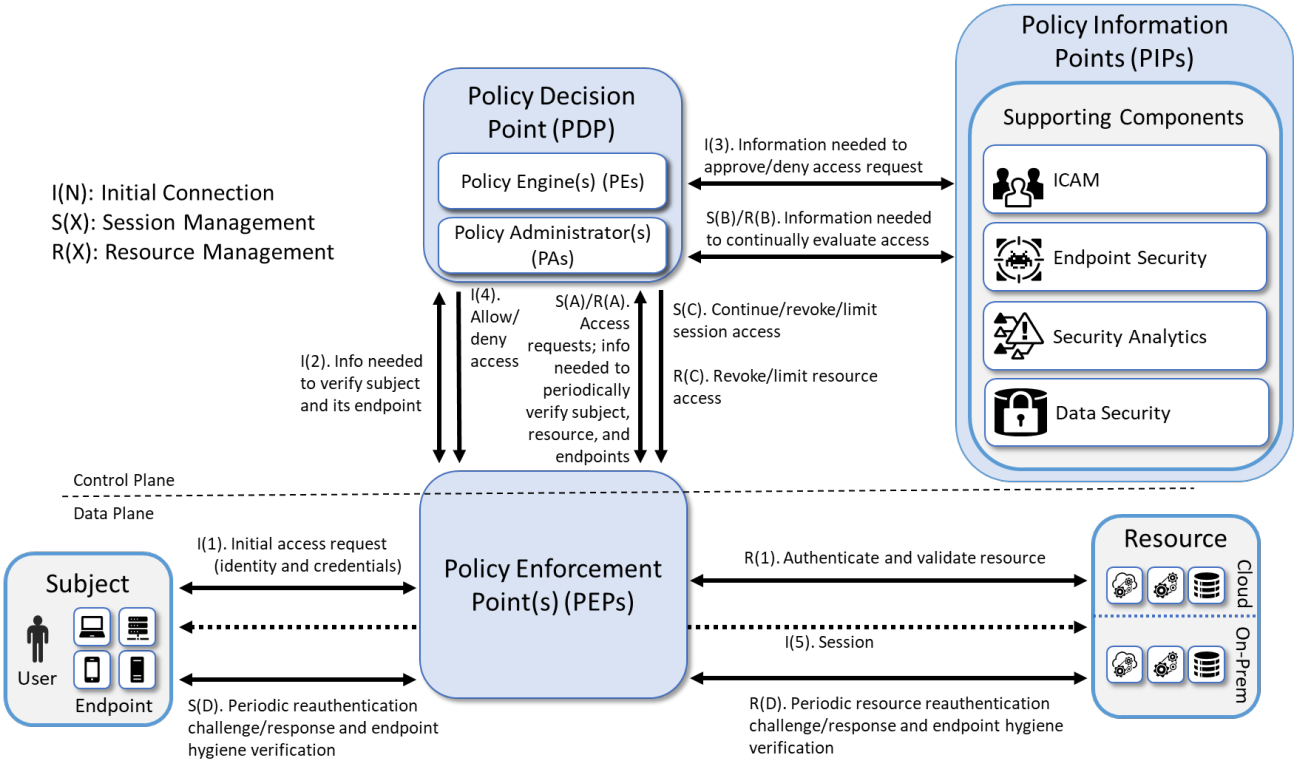

Figure 3-1 General ZTA Reference Architecture

An enterprise ZTA may have numerous PEPs and PDPs. For simplicity, however, Figure 3-1 limits its focus to the interactions involving a single PDP, a single PEP, a single subject, and a single resource. The labeled arrows in Figure 3-1 depict the high-level steps performed in support of the ZTA reference architecture. These steps can be understood in terms of three separate processes:

- **Resource Management—R():** Resource management steps ensure that the resource is authenticated and that its endpoint conforms to enterprise policy. Upon first being brought online, a resource's identity is authenticated, and its endpoint hygiene (i.e., health) is verified. The resource is then connected to the PEP. Once connected to the PEP, access to the resource is granted only through that PEP at the discretion of the PDP. For as long as the resource continues to be online, resource management steps are performed to periodically reauthenticate the resource and verify its endpoint hygiene, thereby continually monitoring its health. These steps are labeled R(1) and R(A) through R(D). Step R(1) occurs first, but the other steps do not necessarily occur in any specific order with respect to each other, which is why they are labeled with letters instead of numbers. Their invocation is determined by enterprise policy. For example, enterprise policy determines how frequently the resource is reauthenticated, what resource-related information the PDP needs to evaluate each access request and when it needs it, and what resource-related changes (environmental, security analytics, etc.) would cause the PDP to decide to revoke or limit access to a particular resource.
- **Session Initiation Steps—I():** Session initiation steps are a sequence of actions that culminate in the establishment of the initial session between a subject and the resource to which it has requested access. These steps are labeled I(1) through I(5), and they occur in sequential order.

- **Session Management Steps—S():** Session management steps describe the actions that enable the PDP to continually evaluate the session once it has been established. These steps begin to be performed after the session has been established, i.e., after Step I(5), and they continue to be invoked periodically for as long as the session remains active. These steps are labeled S(A) through S(D) so that they can be distinguished from each other. However, the letters A through D in the labels are not meant to imply an ordering. The session management steps do not necessarily occur in any specific order with respect to each other. Their invocation is determined by the access requests that are made by the subject in combination with enterprise policy. For example, enterprise policy determines how frequently the subject is reauthenticated, what information the PDP needs to evaluate each access request and when it needs it, and what changes (environmental, security analytics, etc.) would cause the PDP to decide to deny a particular access request or terminate an established session altogether.

Details describing each of the steps in these three processes can be found in our supplemental documentation for ZTA In Operation.

3.2 EIG Crawl Phase Reference Architecture

To support the builds in the EIG crawl phase (the phase without enterprise cloud-based resources), a constrained version of the general ZTA reference architecture depicted in Figure 3-1, called the *EIG Crawl Phase Reference Architecture*, was used. The EIG Crawl Phase Reference Architecture is depicted in Figure 3-2. This architecture included only ICAM, endpoint security, and security analytics components and focused only on protecting resources located on premises. It relied on its ICAM components to provide its PDP functionality, and the only security analytics functionality that it includes is a SIEM. These limitations were intentionally placed on the architecture with the goal of demonstrating the ZTA functionality that an enterprise with legacy ICAM and endpoint protection solutions deployed on premises will be able to support without having to add ZTA-specific capabilities.

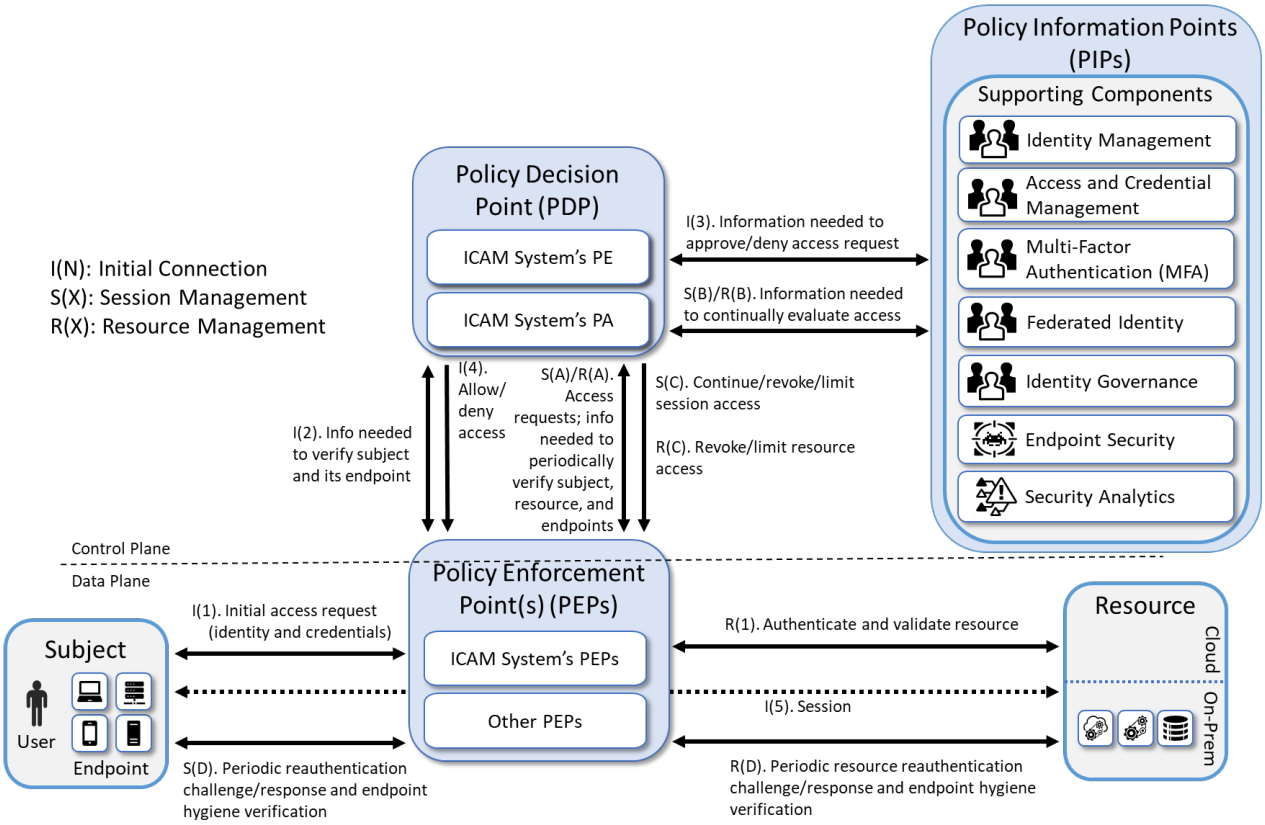

Figure 3-2 EIG Crawl Phase Reference Architecture

3.3 EIG Run Phase Reference Architecture

The EIG run phase, as its name suggests, was built upon the EIG crawl phase architecture. To support the builds in the EIG run phase, some constraints on the EIG crawl phase architecture were lifted. The PDP functionality was no longer required to be provided by the ICAM products used in the build. In addition to protecting access to resources that are located on-premises, the run phase architecture also protects access to some resources that are hosted in the cloud. The EIG run phase also includes a device discovery capability. In addition to monitoring and alerting when new devices are detected, enforcement can be enabled to deny access to devices that are not compliant. The run phase also includes the capability to establish a tunnel between the requesting endpoint and the resource being accessed over which access to the resource can be brokered.

3.4 SDP, Microsegmentation, and SASE Reference Architecture

Unlike the EIG crawl and run phase builds, there are no constraints on the ZTA reference architecture when it is used as the underlying design for a build using the SDP, microsegmentation, SASE deployment approaches, or some combination of these. The SDP and microsegmentation deployment approaches are described in NIST SP 800-207. The microsegmentation approach places one or more resources on unique network segments protected by gateway security components and/or places software agents or firewalls on endpoint assets to implement host-based microsegmentation. The SDP approach involves reconfiguring the network based on access decisions. When implemented at the application layer, this

may be accomplished by establishing a secure channel between a software agent on the endpoint requesting access to the resource and the resource gateway.

SASE delivers converged network and security as a service capability, including Software-Defined Wide Area Network (SD-WAN), Secure Web Gateway (SWG), Cloud Access Security Broker (CASB), Next Generation Firewall (NGFW) and Zero Trust Network Access (ZTNA). SASE supports branch office, remote worker, and on-premises secure access use cases. SASE is primarily delivered as a service and enables zero trust access based on the identity of the device or entity, combined with real-time context and security and compliance policies.

The example solutions implemented as part of the SDP, microsegmentation, and SASE phase also integrated additional supporting components and features to provide an increasingly rich set of ZTA functionalities. The general ZTA reference architecture shown in Figure 3-1, without constraint, is used to support all builds from the SDP, microsegmentation, and SASE phase of this project.

3.5 ZTA Laboratory Physical Architecture

The NCCoE provides virtual machine resources and physical infrastructure for the ZTA laboratory environment. Figure 3-3 depicts the NCCoE ZTA lab. This environment includes four separate enterprise environments, each capable of hosting its own distinct implementation of a ZTA architecture. The enterprises may interoperate as needed by a given use case, and the baseline enterprise environments have the flexibility to support enhancements. The laboratory environment also includes a management virtual local area network (VLAN) on which the following components are installed: Ansible, Terraform, Mandiant Security Validation (MSV) Director, and MSV Protected Theater. These management components support infrastructure as code (IaC) automation and orchestration.

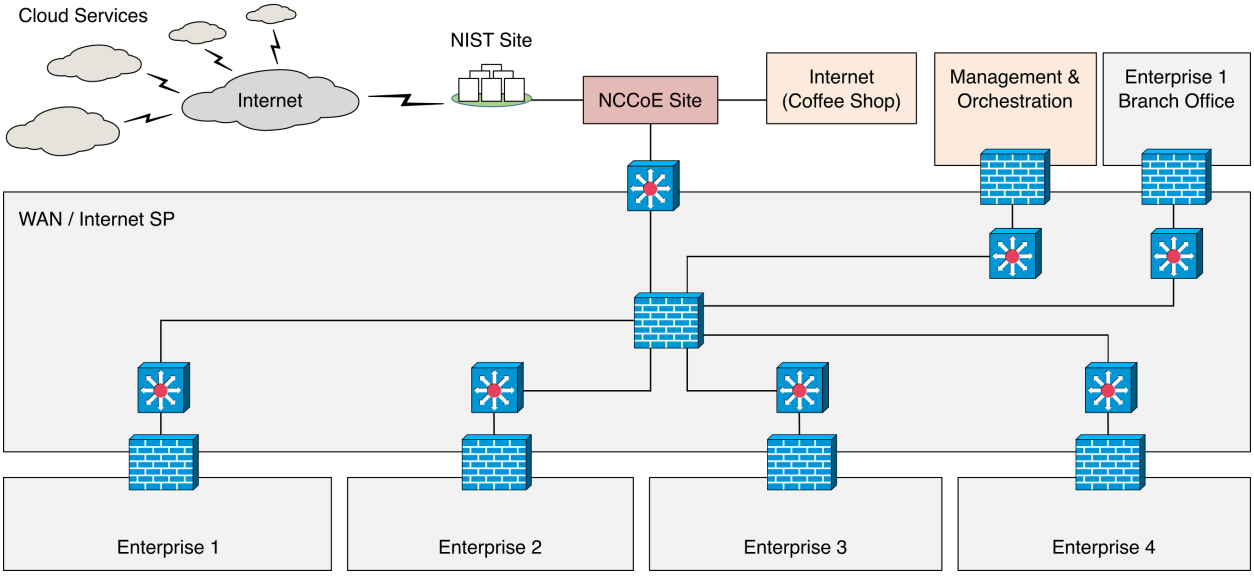

Figure 3-3 Physical Architecture of ZTA Lab

The NCCoE hosts all the collaborators' ZTA-related software for Enterprises 1, 2, 3, and 4. It also provides connectivity from the ZTA lab to the NIST Site, which provides connectivity to the internet and public IP spaces (both IPv4 and IPv6).

Access to and from the ZTA lab is protected by a Palo Alto Networks Next Generation Firewall (PA-5250). (The brick box icons in Figure 3-3 represent firewalls.) In addition to the four independent enterprises (Enterprises 1, 2, 3, and 4) and the management and orchestration domain, the ZTA lab also includes a branch office used only by Enterprise 1, a coffee shop that all enterprises can use, and an emulated wide area network (WAN)/internet service provider. The emulated WAN service provider provides connectivity among all the ZTA laboratory networks, i.e., among all the enterprises, the coffee shop, the branch office, and the management and orchestration domain. Another Palo Alto Networks PA-5250 firewall that is split into separate virtual systems protects the network perimeters of each of the enterprises and the branch office. The emulated WAN service provider also connects the ZTA laboratory network to the NCCoE Site. The ZTA laboratory network has access to cloud services provided by AWS, Azure, IBM Cloud, and Google Cloud, as well as connectivity to SaaS services provided by various collaborators, all of which are available via the internet.

Each enterprise within the NCCoE laboratory environment is protected by a firewall and has both IPv4 and IPv6 (dual stack) configured. Each of the enterprises is equipped with a baseline architecture that is intended to represent the typical environment of an enterprise before a zero trust deployment model is instantiated.

The details of the baseline physical architecture of Enterprise 1, Enterprise 1 branch office, Enterprises 2, 3, and 4, the management and orchestration domain, the coffee shop, and all cloud services, as well as the baseline software and security capabilities running on this physical architecture, are described in our supplemental ZTA Laboratory Physical Architecture documentation.

3.6 Builds Implemented

The following is a list of the builds that have been implemented in the project, organized by build type. Each of these builds instantiates the ZTA architecture in a unique way, depending on the equipment used and the capabilities supported. The products used in each build were based on having out-of-box integration. The details of each build architecture and implementation are shown in Table 4-1.

Note that after the VMware End User Computing Division products were implemented at NCCoE, VMware was acquired by Broadcom, and then the VMware End User Computing Division was divested and reformed under a new entity, Omnissa LLC.

Note that after Enterprise 3's earlier Microsoft builds were completed, the name Azure AD was changed to Entra ID, and the name Defender for Cloud Apps was changed to Defender for Apps.

Note that after Tenable products were implemented at NCCoE, the name Tenable.ad was changed to Tenable Identity Exposure.

EIG Crawl Builds:

- **Enterprise 1 Build 1 (E1B1)** (EIG Crawl; Okta and Ivanti as PEs) uses products from AWS, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zimperium. Certificates from DigiCert are used.

E1B1 components consist of DigiCert CertCentral, IBM Cloud Pak for Security (CP4S), IBM Security QRadar XDR, Ivanti Access Zero Sign-On (ZSO), Ivanti Neurons for Unified Endpoint Management (UEM), Ivanti Sentry, Ivanti Tunnel, Mandiant MSV, Okta Identity Cloud, Okta Verify App, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Tenable.ad, Tenable.io, and Zimperium Mobile Threat Defense (MTD).

- **Enterprise 2 Build 1 (E2B1)** (EIG Crawl; Ping Identity as PE) uses products from Cisco Systems, IBM, Mandiant, Palo Alto Networks, Ping Identity, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.

  E2B1 components consist of Cisco Duo, DigiCert CertCentral, IBM Security QRadar XDR, Mandiant MSV, Palo Alto Networks Next Generation Firewall (NGFW), PingFederate, which is a service in the Ping Identity Software as a Service (SaaS) offering of PingOne, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Tenable.ad, Tenable.io, and Tenable Nessus Network Monitor (NNM).

- **Enterprise 3 Build 1 (E3B1)** (EIG Crawl; Microsoft as PE) uses products from F5, Forescout, Lookout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used.

  E3B1 components consist of DigiCert CertCentral, F5 BIG-IP, Forescout eyeSight, Lookout Mobile Endpoint Security (MES), Mandiant MSV, Microsoft Active Directory (AD), Microsoft Azure AD, Microsoft Defender for Endpoint, Microsoft Endpoint Manager, Microsoft Sentinel, Palo Alto Networks NGFW, PC Matic Pro, Tenable.ad, and Tenable.io.

EIG Run Builds:

- **Enterprise 1 Build 2 (E1B2)** (EIG Run; Zscaler as PE) uses products from AWS, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zscaler. Certificates from DigiCert are also used.

  E1B2 components consist of AWS Infrastructure as a Service (IaaS), DigiCert CertCentral, IBM CP4S, IBM Security QRadar XDR, Mandiant MSV, Okta Identity Cloud, Okta Verify App, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Tenable.ad, Tenable.io, Tenable NNM, Zscaler Admin Portal, Zscaler Application Connector, Zscaler Central Authority, Zscaler Client Connector (ZCC), Zscaler Internet Access (ZIA) Public Service Edges, and Zscaler Private Access (ZPA) Public Service Edges.

- **Enterprise 3 Build 2 (E3B2)** (EIG Run; Microsoft and Forescout as PEs) uses products from F5, Forescout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used.

  E3B2 components consist of DigiCert CertCentral, F5 BIG-IP, Forescout eyeControl, Forescout eyeExtend, Forescout eyeSegment, Forescout eyeSight, Mandiant MSV, Microsoft AD, Microsoft Azure AD, Microsoft Azure AD (Conditional Access), Microsoft Azure AD Identity Protection, Microsoft Azure (IaaS), Microsoft Defender for Cloud, Microsoft Defender for Cloud Apps, Microsoft Defender for Endpoint, Microsoft Intune, Microsoft Office 365 (SaaS), Microsoft Sentinel, Palo Alto Networks NGFW, PC Matic Pro, Tenable.ad, Tenable.io, and Tenable NNM.

- **Enterprise 4 Build 3 (E4B3)** (EIG Run; IBM as PE) uses products from Broadcom (with VMware products), IBM, Mandiant, Palo Alto Networks, and Tenable. Certificates from DigiCert are also used.

E4B3 components consist of DigiCert ONE, IBM CP4S, IBM QRadar XDR, IBM Security Guardium Data Encryption, IBM Security MaaS360 (for both laptops and mobile devices), IBM Security Verify, Mandiant MSV, Palo Alto Networks GlobalProtect Virtual Private Network (VPN), Tenable.ad, Tenable.io, Tenable NNM, and VMware infrastructure.

SDP, Microsegmentation, and SASE Builds:

- **Enterprise 1 Build 3 (E1B3)** (SDP; Zscaler as PE) uses products from AWS, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zscaler. Certificates from DigiCert are also used.

  E1B3 components consist of AWS Infrastructure as a Service (IaaS), DigiCert CertCentral, IBM CP4S, IBM Security QRadar XDR, Mandiant MSV, Okta Identity Cloud, Okta Verify App, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Tenable.ad, Tenable.io, Tenable NNM, Zscaler Admin Portal, Zscaler Application Connector, Zscaler Central Authority, Zscaler Client Connector (ZCC), Zscaler Internet Access (ZIA) Public Service Edges, and Zscaler Private Access (ZPA) Public Service Edges.

- **Enterprise 2 Build 3 (E2B3)** (Microsegmentation; Cisco and Ping Identity as PEs) uses products from Broadcom (with VMware products), Cisco Systems, IBM, Mandiant, Palo Alto Networks, Ping Identity, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.

  E2B3 components consist of Cisco Duo, Cisco Identity Services Engine (ISE), Cisco network devices, Cisco Secure Endpoint (CSE), Cisco Secure Network Analytics (SNA), Cisco Secure Workload, DigiCert CertCentral, IBM Security QRadar XDR, Mandiant MSV, Palo Alto Networks NGFW, Ping Identity PingOne, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Tenable.ad, Tenable.io, Tenable NNM, VMware Workspace ONE UEM and Access.

- **Enterprise 3 Build 3 (E3B3)** (SDP and Microsegmentation; Microsoft and Forescout as PEs) uses products from F5, Forescout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used.

  E3B3 components consist of DigiCert CertCentral, F5 BIG-IP, Forescout eyeControl, Forescout eyeExtend, Forescout eyeSight, Forescout eyeSegment, Mandiant MSV, Microsoft AD, Microsoft Azure AD, Microsoft Azure AD (Conditional Access), Microsoft Azure AD Identity Governance, Microsoft Intune, Microsoft Sentinel, Microsoft Azure App Proxy, Microsoft Defender for Endpoint, Microsoft Azure AD Identity Protection, Microsoft Defender for Identity, Microsoft Defender for Office, Microsoft Entra Permissions Management, Microsoft Defender for Cloud Apps, Microsoft Purview – Data Loss Prevention (DLP), Microsoft Purview Information Protection, Microsoft Purview Information Protection Scanner, Microsoft Intune VPN Tunnel, Microsoft Azure Arc, Microsoft Azure Automanage, Microsoft Intune Privilege Access Workstation, Microsoft Azure Virtual Desktop Windows 365, Microsoft Defender for Cloud, Microsoft Azure (IaaS), Microsoft Office 365 (SaaS), Palo Alto Networks NGFW, PC Matic Pro, Tenable.io, Tenable.ad, and Tenable NNM.

- **Enterprise 1 Build 4 (E1B4)** (SDP; Appgate as PE) uses products from AWS, Appgate, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zimperium. Certificates from DigiCert are also used.

  E1B4 components consist of Appgate SDP Controller, Appgate SDP Gateway, Appgate SDP client, Appgate Portal, AWS IaaS and SaaS, DigiCert CertCentral, IBM CP4S, IBM Security QRadar XDR, Ivanti Neurons for UEM Platform, Mandiant MSV, Okta Identity Cloud, Okta Verify App, Radiant

Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Tenable.ad, Tenable.io, Tenable NNM, and Zimperium MTD.

- **Enterprise 2 Build 4 (E2B4)** (SDP and SASE; Broadcom (with Symantec products) as PE) uses products from Broadcom (with VMware and Symantec products), Google Cloud, IBM, Mandiant, Okta, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.

  E2B4 components consist of Symantec Cloud Secure Web Gateway (Cloud SWG), Symantec Zero Trust Network Access (ZTNA), Symantec Cloud Access Security Broker (CASB), Symantec Endpoint Security Agent, VMware Workspace ONE UEM, Symantec DLP Cloud Detection Service, Symantec ZTNA Connector, Okta Identity Cloud, Okta Verify App, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, IBM Security QRadar XDR, Tenable.io, Tenable.ad, Tenable NNM, Mandiant MSV, Google Cloud, and DigiCert CertCentral.

- **Enterprise 3 Build 4 (E3B4)** (SDP; F5 as PE) uses products from F5, Forescout, Mandiant, Microsoft, Palo Alto Networks, and Tenable. Certificates from DigiCert are also used.

  E3B4 components consist of F5 BIG-IP, F5 NGINX Plus, F5 Access App, Microsoft AD, Microsoft Azure AD, Microsoft Azure AD Identity Governance, Microsoft Intune, Microsoft Sentinel, Tenable.io, Tenable.ad, Tenable NNM, Mandiant MSV, Forescout eyeControl, Forescout eyeExtend, Forescout eyeSight, Forescout eyeSegment, Microsoft Azure (IaaS), and DigiCert CertCentral.

- **Enterprise 4 Build 4 (E4B4)** (SDP; Microsegmentation, and EIG; Broadcom (with VMware products) as PE) uses products from Broadcom (with VMware products), IBM, Mandiant, and Tenable. Certificates from DigiCert are also used.

  E4B4 components consist of VMware Workspace ONE Access, VMware Unified Access Gateway (UAG), VMware NSX-T, VMware Workspace ONE UEM, VMware Workspace ONE MTD, VMware Carbon Black Enterprise EDR, VMware Carbon Black Cloud, VMware vSphere, VMware vCenter, VMware vSAN, IBM Security QRadar XDR, Mandiant MSV, Tenable.io, Tenable.ad, Tenable NNM, and DigiCert ONE.

- **Enterprise 1 Build 5 (E1B5)** (Microsegmentation and SASE; Palo Alto Networks as PE) uses products from AWS, IBM, Mandiant, Okta, Palo Alto Networks, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.

  E1B5 components consist of Palo Alto Networks (PAN) Panorama, PAN Next Generation Firewall (NGFW), PAN Prisma Access, PAN Prisma SASE (Prisma Access & Prisma SD-WAN), PAN Cloud Delivered Security Services (CDSS), PAN Cloud Identity Engine, PAN Global Protect, PAN Strata Cloud Manager, Okta Identity Cloud, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Okta Verify App, IBM Security QRadar XDR, Tenable.io, Tenable.ad, Tenable NNM, Mandiant MSV, DigiCert CertCentral, and AWS IaaS.

- **Enterprise 2 Build 5 (E2B5)** (SDP and SASE; Lookout SSE and Okta Identity Cloud as PEs) uses products from Broadcom (with VMware products), Google Cloud, IBM, Lookout, Mandiant, Okta, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.

  E2B5 components consist of Lookout Security Service Edge (SSE) (includes Secure Private Access [SPA], Secure Cloud Access [SCA], and Secure Internet Access [SIA]), Lookout Secure Private Access Connector, VMware Workspace ONE UEM, Lookout MES, Lookout Client, Okta Identity Cloud, Okta Verify App, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint

IdentityIQ, IBM Security QRadar XDR, Tenable.io, Tenable.ad, Tenable NMM, Mandiant MSV, Google Cloud, Google Workspace, and DigiCert CertCentral.

- **Enterprise 3 Build 5 (E3B5)** (SDP and SASE; Microsoft Entra Conditional Access (formerly called Azure AD Conditional Access), and Microsoft Security Service Edge as PEs) uses products from Mandiant, Microsoft, and Tenable. Certificates from DigiCert are also used.

  E3B5 components consist of Microsoft Entra Conditional Access, Microsoft Security Service Edge (SSE) (which includes Entra Private Access, Entra Internet Access, and Microsoft 365 Access), Microsoft Entra Private Access Connector, Microsoft Entra ID, Microsoft Entra ID Governance, Microsoft Intune, Microsoft Defender for Endpoint, Microsoft Global Secure Access Client, Microsoft Purview DLP, Microsoft Purview Information Protection, Microsoft Purview Information Protection Scanner, Microsoft Entra ID Identity Protection, Microsoft Defender for Identity, Microsoft Defender for Cloud, Microsoft Sentinel, Tenable.io, Tenable.ad, Mandiant Security Validation, Microsoft Azure (IaaS), Microsoft 365 (SaaS), and DigiCert CertCentral.

- **Enterprise 4 Build 5 (E4B5)** (SDP and Microsegmentation; AWS Verified Access and Amazon VPC Lattice as PEs) uses products from AWS, IBM, Mandiant, Okta, and Tenable. Certificates from DigiCert are also used.

  E4B5 components consist of AWS Verified Access, Amazon VPC Lattice, Amazon ECS and AWS Lambda Functions, Okta Identity Cloud, Okta Verify App, IBM Security QRadar XDR, Tenable Cloud Security, Mandiant MSV, DigiCert CertCentral, and AWS IaaS.

- **Enterprise 1 Build 6 (E1B6)** (SDP and Microsegmentation; Ivanti Neurons for Zero Trust Access (nZTA) as PE) uses products from AWS, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.

  E1B6 components consist of Ivanti nZTA, Ivanti nZTA Gateway, Okta Identity Cloud, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Okta Verify App, Ivanti Secure Access Client, IBM Security QRadar XDR, Tenable.io, Tenable.ad, Tenable NNM, Mandiant MSV, DigiCert CertCentral, and AWS IaaS.

- **Enterprise 2 Build 6 (E2B6)** (SASE; Google Chrome Enterprise Premium (CEP) – Access Context Manager as PE) uses products from Google Cloud, IBM, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Omnissa. Certificates from DigiCert are also used.

  E2B6 components consist of Google CEP, Google Application Connector, Omnissa Workspace ONE UEM, Okta Identity Cloud, Okta Verify App, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, IBM Security QRadar XDR, Tenable.io, Tenable.ad, Tenable NNM, Mandiant MSV, Google Cloud (IaaS), Google Workspace (SaaS), and DigiCert CertCentral.

4 Build Implementation Instructions

Table 4-1 identifies the policy engines/policy decision points and types of architecture used in each build. It also links to the online locations where each build architecture is described in detail, as well as the online locations where instructions for implementing each build can be found. These build implementation instructions are designed to enable information technology professionals to replicate all or parts of this project.

To see which build suits your organization, you can first identify which of the ZTA approaches—EIG, SDP, microsegmentation, or SASE—meets your organization's requirements. You can then look at the build

options provided in Table 4-1. Based on your selection of the ZTA approach, you can view the details of the relevant builds by clicking the link in the "Build Architecture, Technologies, and Flow Diagrams" column.

Since most enterprises evolve their enterprise architecture toward ZTA, i.e., by starting with their already-existing enterprise environment and gradually adding or adapting capabilities such as PE, you can start by looking at the builds with the products closest to your existing environment.

Table 4-1 Mapping of Builds to Online Details Regarding Architecture Descriptions and Implementation Instructions

Build	Policy Engines/ Policy Decision Points	ZTA Architecture Instantiated	Links to Online Details: Build Architecture, Technologies, and Flow Diagrams	Links to Online Details: Build Implementation Instructions
E1B1	Okta Identity Cloud Ivanti Access ZSO	EIG Crawl	E1B1 Build Architecture	E1B1 Build Implementation Instructions
E2B1	Ping Identity Ping Federate	EIG Crawl	E2B1 Build Architecture	E2B1 Build Implementation Instructions
E3B1	Azure AD (Conditional Access, later renamed Entra Conditional Access)	EIG Crawl	E3B1 Build Architecture	E3B1 Build Implementation Instructions
E1B2	Zscaler ZPA Central Authority (CA)	EIG Run	E1B2 Build Architecture	E1B2 Build Implementation Instructions
E3B2	Microsoft Azure AD (Conditional Access, later renamed Entra Conditional Access) Microsoft Intune Forescout eyeControl Forescout eyeExtend	EIG Run	E3B2 Build Architecture	E3B2 Build Implementation Instructions
E4B3	IBM Security Verify	EIG Run	E4B3 Build Architecture	E4B3 Build Implementation Instructions
E1B3	Zscaler ZPA Central Authority (CA)	SDP	E1B3 Build Architecture	E1B3 Build Implementation Instructions

Build	Policy Engines/ Policy Decision Points	ZTA Architecture Instantiated	Links to Online Details: Build Architecture, Technologies, and Flow Diagrams	Links to Online Details: Build Implementation Instructions
E2B3	Ping Identity PingFederate Cisco ISE Cisco Secure Workload	Microsegmentation	E2B3 Build Architecture	E2B3 Build Implementation Instructions
E3B3	Microsoft Azure AD (Conditional Access, later renamed Entra Conditional Access) Microsoft Intune Microsoft Sentinel Forescout eyeControl Forescout eyeExtend	SDP and Microsegmentation	E3B3 Build Architecture	E3B3 Build Implementation Instructions
E1B4	Appgate SDP Controller	SDP	E1B4 Build Architecture	E1B4 Build Implementation Instructions
E2B4	Symantec Cloud Secure Web Gateway (Cloud SWG) Symantec ZTNA Symantec Cloud Access Security Broker (CASB)	SDP and SASE	E2B4 Build Architecture	E2B4 Build Implementation Instructions
E3B4	F5 BIG-IP F5 NGINX Plus Forescout eyeControl Forescout eyeExtend	SDP	E3B4 Build Architecture	E3B4 Build Implementation Instructions
E4B4	VMware Workspace ONE Access VMware Unified Access Gateway (UAG) VMware NSX-T	SDP, Microsegmentation, and EIG	E4B4 Build Architecture	E4B4 Build Implementation Instructions
E1B5	PAN NGFW PAN Prisma Access	SASE and Microsegmentation	E1B5 Build Architecture	E1B5 Build Implementation Instructions
E2B5	Lookout SSE Okta Identity Clouds	SDP and SASE	E2B5 Build Architecture	E2B5 Build Implementation Instructions

Build	Policy Engines/ Policy Decision Points	ZTA Architecture Instantiated	Links to Online Details: Build Architecture, Technologies, and Flow Diagrams	Links to Online Details: Build Implementation Instructions
E3B5	Microsoft Entra Conditional Access (formerly called Azure AD Conditional Access) Microsoft Security Service Edge	SDP and SASE	E3B5 Build Architecture	E3B5 Build Implementation Instructions
E4B5	AWS Verified Access and Amazon VPC Lattice	SDP and Microsegmentation	E4B5 Build Architecture	E4B5 Build Implementation Instructions
E1B6	Ivanti Neurons for Zero Trust Access	SDP and Microsegmentation	E1B6 Build Architecture	E1B6 Build Implementation Instructions
E2B6	Google CEP – Access Context Manager	SASE	E2B6 Build Architecture	E2B6 Build Implementation Instructions

5 General Findings

When deploying ZTA, the following capabilities are considered to be fundamental to determining whether a request to access a resource should be granted and, once granted, whether the access session should be permitted to persist:

- Authentication and periodic reauthentication of the requesting user's identity
- Authentication and periodic reauthentication of the requesting endpoint
- Authentication and periodic reauthentication of the endpoint that is hosting the resource being accessed
- Each authentication and reauthentication includes authorization and reauthorization

In addition, the following capabilities are also considered highly desirable:

- Verification and periodic reverification of the requesting endpoint's health
- Verification and periodic reverification of the health of the endpoint that is hosting the resource being accessed

5.1 EIG Crawl Phase Findings

In the EIG crawl phase, we followed two patterns. First, we leveraged our ICAM solutions to also act as PDPs. We discovered that many of the vendor solutions used in the EIG crawl phase do not integrate with each other out-of-the-box in ways that are needed to enable the ICAM solutions to function as PDPs. Typically, network-level PEPs, such as routers, switches, and firewalls, do not integrate directly

with ICAM solutions. However, network-level PEPs that are identity-aware may integrate with ICAM solutions. Also, endpoint protection solutions in general do not typically integrate directly with ICAM solutions. However, some of the endpoint protection solutions considered for use in the builds have out-of-the-box integrations with the mobile device management (MDM)/UEM solutions used, which provide the endpoint protection solutions with an indirect integration with the ICAM solutions.

Second, we used out-of-the-box integrations offered by the solution providers rather than performing custom integrations. These two patterns combined do not support all the desired zero trust capabilities.

Both builds E1B1 and E3B1 were capable of authenticating and reauthenticating requesting users and requesting endpoints and of verifying and periodically reverifying the health of requesting endpoints, and both builds were able to base their access decisions on the results of these actions. Access requests were not granted unless the identities of the requesting user and the requesting endpoint could be authenticated and the health of the requesting endpoint could be validated; however, no check was performed to authenticate the identity or verify the health of the endpoint hosting the resource.

Access sessions that are in progress in both builds are periodically reevaluated by reauthenticating the identities of the requesting user and the requesting endpoint and by verifying the health of the requesting endpoint. If these periodic reauthentications and verifications cannot be performed successfully, the access session will eventually be terminated; however, neither the identity nor the health of the endpoint hosting the resource is verified on an ongoing basis, nor does its identity or health determine whether it is permitted to be accessed.

Neither build E1B1 nor build E3B1 was able to support resource management as envisioned in the ZTA logical architecture depicted in Figure 3-1. These builds do not include any ZTA technologies that perform authentication and reauthentication of resources that host endpoints, nor are these builds capable of verifying or periodically reverifying the health of the endpoints that host resources. In addition, when using both builds E1B1 and E3B1, devices (requesting endpoints and endpoints hosting resources) were initially joined to the network manually. Neither of the two EIG crawl phase builds includes any technologies that provide network-level enforcement of an endpoint's ability to access the network. That is, there is no tool in either build that can keep any endpoint (either one that is hosting a resource or one that is used by a user) from initially joining the network based on its authentication status. The goal is to try to support resource management in future builds as allowed by the technologies used.

5.2 EIG Run Phase Findings

The EIG run phase enabled us to demonstrate additional capabilities over the EIG crawl phase, such as:

- establishment of secure, direct access tunnels from requesting endpoints to private enterprise resources, regardless of whether the resources are located on-premises or in the cloud, driven by policy and enforced by PEPs
- use of connectors that act as proxies for internal, private enterprise resources, enabling resources to be accessed by authenticated, authorized users while ensuring that they are not discoverable by or visible to others

- protection for private enterprise resources hosted in the cloud that enables authenticated, authorized remote users to access those resources directly rather than having to hairpin through the enterprise network
- ability to monitor, inspect, and enforce policy controls on traffic being sent to and from resources in the cloud or on the internet
- discovery of new endpoints on the network and the ability to block newly discovered endpoints that are not compliant with policy

Build E1B2, which uses Zscaler as its PE, PA, and PEP, does not have an EPP because this build does not include any collaborators with EPP solutions that integrate with Zscaler. Zscaler (e.g., the Zscaler client connector) has the capability to enforce policies based on a defined set of endpoint compliance checks to allow or deny user/endpoint access to a resource. However, it does not perform the functions of an EPP solution to protect an endpoint. Zscaler integrates with EPP solutions to receive a more robust set of information about the endpoints in order to make a decision to allow or deny access to a resource. However, in build E1B2, we do not have a collaborator with an EPP solution that can integrate with Zscaler.

Because there is no EPP in E1B2, there is no automatic solution to remediate an issue on the endpoint either.

Build E1B2 also does not have a collaborator with a solution that supports the determination of confidence level/trust scores that can integrate with Zscaler. Due to the absence of a collaborator with this capability, Build E1B2 does not support the calculation of confidence levels/trust scores.

Build E2B1, which uses Ping Identity as its PE and PA and Ping Identity and Cisco Duo as its PEP, does not have an EPP. Cisco Duo provides limited device health information but not the full spectrum that an EPP would provide. Because there is no official EPP in this build, there is no automatic solution to remediate an issue on the endpoint. An EPP for Enterprise 2 was included in a later build phase (E2B3).

When planning a ZTA implementation, organizations should ensure that all of the ZTA core and supporting components that can integrate with each other are selected. This enables having end-to-end ZTA with full functionality.

Build E3B2 currently supports one-way integration between Microsoft Intune and Forescout eyeExtend. If Intune detects an endpoint out of compliance, eyeExtend can become informed of this problem by pulling information from Intune. However, if one of Forescout's discovery tools detects a problem with an endpoint, there is currently no mechanism for this information to be passed from Forescout eyeExtend to Microsoft Intune. Ideally, future integration of these products would allow Forescout eyeExtend to inform Microsoft Intune when it detects a non-Azure AD-connected endpoint that is non-compliant, as this would enable Intune to direct Azure AD to block sign-in from the non-compliant endpoint. Without a mechanism for enabling Forescout eyeExtend to send endpoint compliance information to Microsoft Intune, Azure AD does not have a way of knowing that a non-Azure AD-connected endpoint is not compliant.

5.3 SDP, Microsegmentation, and SASE Phase Findings

More integration of zero trust products from different vendors is needed to support the implementation of ZTAs that are built using components from a variety of vendors. For the most effective zero trust solutions, PDPs should integrate with a variety of security tools and other supporting components that enable the PDP to assess the real-time risk of any given access request.

It is not unusual for a ZTA to have multiple PDPs, each of which may be integrated with one or more different supporting components and/or PEPs. As a result, the policies that the ZTA enforces are not centrally located. Rather, they are configured and managed in association with each of the various PDPs. This makes it challenging to understand, articulate, and manage the ZTA's policies as a comprehensive whole.

In addition, the multiple PDPs that comprise a ZTA do not typically integrate with each other to share information, so they do not have a shared understanding of what users, endpoints, or other subjects may pose risks. For example, one PDP may be aware that an endpoint is non-compliant, whereas this same endpoint compliance information is not available to another PDP. On the other hand, the second PDP may be aware that the endpoint's user may have exhibited suspicious behavior, whereas the first PDP is not. Ideally, when a ZTA has multiple PDPs, it is desirable to have an integrated approach that enables the PDPs to share information so that they can each be more fully informed, share a common, consolidated understanding of risks, and make a decision based on all information available.

The SIEM and/or security orchestration, automation, and response (SOAR) components contain a wealth of information that could prove useful to a PDP as it tries to determine whether any given access request should be allowed or not. Ideally, the SIEM and SOAR should send this information to the PDP in real-time, if possible, to ensure that the PDP's access decisions are fully informed.

Ideally, data security tools should be integrated with the PDP so that the PDP can be made aware of instances in which access requests are denied by the tools that are designed to protect data.

Additionally, risk information and user behavior analytics should be shared with the PDP to potentially improve ZTA security.

Some zero trust SDP solutions for managing endpoints can also manage resources by installing clients onto those resources. However, solutions that are specifically designed to manage resources should be leveraged rather than the zero trust solutions that have the primary purpose of managing endpoints. In some cases, the solutions that manage resources do not have out-of-the-box integration with the PDPs. PDP integration capability should be available in these resource management solutions.

Endpoint compliance is essential for security. It is important to have tools that are capable of detecting when an endpoint is not compliant and ensuring that the endpoint is not permitted to access resources as a result. Furthermore, automatic solutions to remediate noncompliance issues on the endpoint should be deployed when possible, and these should be integrated with the organization's configuration and patch management systems.

6 Functional Demonstrations

This section defines the methodologies we used to demonstrate the capabilities of the project's ZTA builds, summarizes the use cases that were demonstrated, and summarizes the results of performing these use cases with each of the project's builds.

6.1 Demonstration Methodology

We are leveraging two types of demonstration methodologies in this project: manual and automated. Demonstrations that require human interaction (e.g., a user performs MFA) must be performed manually. Demonstrations that do not require human interaction can be performed either manually, automated, or both. It is also possible to perform demonstrations in a hybrid manner in which the early part of a demonstration that requires user authentication is performed manually, followed by an automated portion of the demonstration. This approach can be helpful for demonstrations that are complicated yet nevertheless require human interaction.

We deployed Mandiant MSV throughout the project's laboratory environment to enable us to monitor and verify various security characteristics of the builds. MSV automates a testing program that provides visibility and evidence of how security controls are performing by emulating attackers to safely process advanced cyberattack security content within production environments. It is designed so defenses respond to it as if an attack is taking place within the enterprise. Virtual machines (VMs) that are intended to operate as actors are deployed on each of the subnetworks in each of the enterprises. These actors can be used to initiate various actions for the purpose of verifying that security controls are working to support the objectives of zero trust. We also deployed three VMs that operate as directors, two of which function as applications within Enterprise 1 and Enterprise 3 that are used by those enterprises to monitor and audit their own traffic, and one of which is an overarching director that is located within the management and orchestration domain and used by the project team to demonstrate and audit operations that span multiple enterprises.

This setup enabled the following dual-purpose MSV deployment:

1. **A typical MSV deployment, in which each enterprise deploys MSV as an application within its own enterprise, uses MSV for self-auditing and testing.** Each enterprise deploys a director and multiple actors that function as applications within the enterprise, enabling the enterprise to monitor and test its own enterprise security capabilities, verifying the protections it receives from the ZTA and its ability to operate as expected. In this capacity, MSV is treated just like any other application deployed within that enterprise. The components may be protected by PEPs according to enterprise policies, and directors and actors exchange traffic over the same data communications paths as other enterprise applications. Firewalls and policies within the ZTA must be configured to permit the communications that the MSV components send and receive, including traffic that is sent between actors and the director to control the actions that are performed to test various security controls.
2. **The NCCoE project team, as testers, use MSV to monitor and audit enterprise and inter-enterprise actions.** The project team deploys an overarching director and a management backchannel connecting that director to all actors throughout the laboratory environment. This overarching director is used as a tool to verify the security controls provided by each of the ZTAs

in the various enterprises and to monitor and audit inter-enterprise interactions. In this capacity, MSV does not function as an application deployed or controlled by enterprises; rather, it is a tool used to monitor and audit enterprise and inter-enterprise activity. Communications between the actors and this overarching director occur on a management channel that is separate from the data networks in each of the enterprises. Using a separate backchannel ensures that the tool being used to monitor and verify the various ZTA architectures does not itself impact those architectures. Enabling the overarching MSV director to control the actor VMs via a backchannel requires each of the actor VMs to have two network interface cards (NICs), one for enterprise data and one for MSV tool interoperation. The use of a separate backchannel ensures that enterprise ZTA policies and firewalls don't need to be modified to accommodate the overarching MSV testing by permitting traffic between the overarching director and the actors that would not normally be expected to transit any of the enterprise networks. Such policy and firewall modification would have been undesirable and would, in effect, have amounted to unauthorized channels into the enterprise networks.

An MSV protective theater was also created in the lab. This is a virtualized system that allows destructive actions to be tested without adversely impacting the enterprise deployments themselves. For example, to understand the effects that malware might have on a specific system in one of the enterprises, that system could be imported into the protective theater and infected with malware to test what the destructive effects of the malware might be.

6.2 Demonstration Use Cases

Eight demonstration use cases were defined to exercise the security functionality provided by each of the example solutions that were implemented as part of this project. Each use case consists of one or more scenarios. The use cases and their scenarios are summarized in the following subsections.

More detailed descriptions of each use case and scenario, including their preconditions; demonstration steps; purposes; detailed tables of the various permutations of subject, ID, endpoint, and resource attributes to be exercised; and expected outcomes are available in our supplemental documentation on Functional Demonstrations.

Definitions of terminology used throughout the demonstration scenarios are available in our Demonstration Terminology documentation. The terminology includes identifier, subject, endpoint, and resource types; compliance; authentication status; access levels; user and access profiles; assumptions; and other information that is required to fully describe the demonstration use cases.

6.2.1 Use Case A: Discovery and Identification

Use Case A demonstrates discovery and identification of identifiers, endpoint assets, and data flows. Its scenarios are:

- Scenario A-1: Discovery and authentication of endpoint assets
- Scenario A-2: Reauthentication of identified assets
- Scenario A-3: Discovery of transaction flows

6.2.2 Use Case B: Enterprise-ID Access

Use Case B demonstrates a subject with an ID that is issued and maintained by the enterprise requesting access to a resource. Its scenarios are:

- Scenario B-1: Full/limited resource access using an enterprise endpoint – the subject is granted full, limited, or no access to the requested resource as determined by its authentication status and endpoint compliance status.
- Scenario B-2: Full/limited internet access using an enterprise endpoint – the subject is granted full, limited, or no access to the requested internet domain as determined by enterprise policy.
- Scenario B-3: Stolen credential using an enterprise endpoint – a legitimate user's enterprise ID credential is stolen and is used to request access to an enterprise resource from an enterprise-managed endpoint.
- Scenario B-4: Full/limited resource access using BYOD – a subject using a bring-your-own-device (BYOD) is granted full or limited access to the requested resource as determined by authentication status and enterprise policy.
- Scenario B-5: Full/limited internet access based on ID attributes – the subject is granted full, limited, or no access to the requested internet domain as determined by enterprise ID profiles and enterprise policy.
- Scenario B-6: Stolen credential using BYOD – a legitimate user's enterprise ID credential is stolen and is used to request access to an enterprise resource from a BYOD endpoint.
- Scenario B-7: Just-in-Time Access Privileges – An enterprise provisions access privileges to a resource based on a single business process flow. Temporary privileges are granted to perform a portion of the business process and then revoked when the process is complete.
- Scenario B-8: Enterprise-ID Step-Up Authentication – A subject who already has an active access session with a resource, requests to perform an action on that resource that requires additional authentication checks.

6.2.3 Use Case C: Collaboration: Federated-ID Access

Use Case C demonstrates a subject with a successfully authenticated Federated-ID (i.e., an ID that is issued and maintained by another enterprise in a trusted community of interest) requesting access to a resource. Its scenarios are:

- Scenario C-1: Full resource access using an enterprise endpoint – the subject is granted full access to the requested resource as determined by its endpoint compliance status.
- Scenario C-2: Limited resource access using an enterprise endpoint – the subject is granted limited access to the requested resource as determined by its endpoint compliance status.
- Scenario C-3: Limited internet access using an enterprise endpoint – the subject is granted limited access to internet domains as determined by its endpoint compliance status and enterprise policy.
- Scenario C-4: No internet access using enterprise owned endpoint – the subject is denied all access to internet domains as determined by enterprise policy.

- Scenario C-5: Internet access using BYOD – the subject is granted or denied access to an internet domain as determined by enterprise policy.
- Scenario C-6: Access resources using BYOD – the subject is granted limited access to an enterprise resource as determined by enterprise policy, which dictates that if a subject is using a BYOD, the subject's access to enterprise resources will be limited.
- Scenario C-7: Stolen credential using an enterprise endpoint – a legitimate user's federated ID credential is stolen and is used to request access to an enterprise resource from an enterprise-managed endpoint.
- Scenario C-8: Stolen credential using BYOD – a legitimate user's federated ID credential is stolen and is used to request access to an enterprise resource from a BYOD endpoint.

6.2.4 Use Case D: Other-ID Access

Use Case D demonstrates a subject with an Other-ID (i.e., an ID that is issued and maintained by another enterprise but known or registered by the first enterprise) requesting access to a resource. Its scenarios are:

- Scenario D-1: Full/limited resource access using an enterprise endpoint – the subject is granted full, limited, or no access to the requested resource as determined by its authentication status and endpoint compliance status.
- Scenario D-2: Full/limited internet access using an enterprise endpoint – the subject is granted full, limited, or no access to the requested internet domain as determined by enterprise policy.
- Scenario D-3: Stolen credential using BYOD or enterprise endpoint – a legitimate user's Other-ID credential is stolen and is used to request access to an enterprise resource from either an enterprise-managed endpoint or a BYOD.
- Scenario D-4: Full/limited resource access using BYOD – a subject using a bring-your-own device (BYOD) is granted full or limited access to the requested resource as determined by authentication status and enterprise policy.
- Scenario D-5: Full/limited internet access using BYOD – the subject is granted or denied access to an internet domain as determined by enterprise policy.
- Scenario D-6: Stolen credential using BYOD – a legitimate user's Other-ID credential is stolen and is used to request access to an enterprise resource from a BYOD endpoint.
- Scenario D-7: Just-in-Time Access Privileges – An enterprise provisions access privileges to a resource based on a single business process flow. Temporary privileges are granted to perform a portion of the business process and then revoked when the process is complete.
- Scenario D-8: Other-ID Step-Up Authentication – A subject who already has an active access session with a resource requests to perform an action on that resource that requires additional authentication checks.

6.2.5 Use Case E: Guest: No-ID Access

Use Case E demonstrates a subject that does not have an ID (i.e., a guest on the network) requesting access to a resource. Its scenario is:

- Scenario E-1: Guest requests public internet access – the guest user is permitted to access public internet domains and resources.

6.2.6 Use Case F: Confidence Level

Use Case F demonstrates a subject that has been granted access to a resource and has an active session to the resource. The events listed in the following use cases cause the subject's authorization to access the resource to be re-evaluated:

- Scenario F-1: User reauthentication fails during active session, causing the subject's access to the resource to be terminated.
- Scenario F-2: Requesting endpoint reauthentication fails during active session, causing the subject's access to the resource to be terminated.
- Scenario F-3: Resource reauthentication fails during active session, causing the subject's access to the resource to be terminated.
- Scenario F-4: Compliance fails during active session, causing the subject's access to the resource to be terminated.
- Scenario F-5: Compliance improves between requests – in this case the subject had not been permitted to access a resource due to non-compliance of the requesting endpoint. However, after the endpoint is brought into compliance and access to the resource is requested again, access is granted.
- Scenario F-6: Enterprise-ID Violating Data Use Policy, causing the subject's access to the resource to be terminated.
- Scenario F-7: Other-ID Violating Data Use Policy, causing the subject's access to the resource to be terminated.
- Scenario F-8: Enterprise-ID Violating Internet Use Policy.
- Scenario F-9: Other-ID Violating Internet Use Policy, causing the subject's access to the resource to be terminated.
- Scenario F-10: Enterprise-ID Attempting Unauthorized Access Detection and Response, Access Queries – the enterprise detects a subject's attempt to access an unauthorized resource and responds by revoking access to a resource to which the subject had previously been granted access.
- Scenario F-11: Enterprise-ID Attempting Unauthorized Access Detection and Response, Ongoing Sessions - the enterprise detects a subject's attempt to access an unauthorized resource and responds by terminating the user's active, open access session with a resource.
- Scenario F-12: Other-ID Attempting Unauthorized Access Detection and Response, Access Queries - the enterprise detects a subject's attempt to access an unauthorized resource and responds by revoking access to a resource to which the subject had previously been granted access.
- Scenario F-13: Other-ID Attempting Unauthorized Access Detection and Response, Ongoing Sessions - the enterprise detects a subject's attempt to access an unauthorized resource and responds by terminating the user's active, open access session with a resource.

- Scenario F-14: Enterprise-ID Denied Access Due to Suspicious Endpoint – A subject requests access from an endpoint that had been previously flagged as being suspected of being compromised. The enterprise responds by denying the request and preventing all access requests from the enterprise ID used in this request.
- Scenario F-15: Other-ID Denied Access due to Suspicious Endpoint – A subject requests access from an endpoint that had been previously flagged as being suspected of being compromised. The enterprise responds by denying the request and preventing all access requests from the Other-ID used in this request.
- Scenario F-16: Enterprise-ID Access Terminated Due to Suspicious Endpoint – A subject requests access from an endpoint that had been previously flagged as being suspected of being compromised. The enterprise responds by denying the request and terminating any open access sessions from the Enterprise-ID used in this request.
- Scenario F-17: Other-ID Access Terminated Due to Suspicious Endpoint – A subject requests access from an endpoint that had been previously flagged as being suspected of being compromised. The enterprise responds by denying the request and terminating any open access sessions from the Other-ID used in this request.

6.2.7 Use Case G: Service-Service Interaction

Use Case G demonstrates service-to-service Interactions in which a non-person subject requests access to a resource via API calls. The enterprise can uniquely identify and authenticate both the subject and the resource, and both the subject and the resource are in compliance. Whether or not the access request is granted depends on whether the subject is authorized to access the resource, which depends on enterprise policy. The access request is an API call between two services; the location of the services varies by scenario, as can be seen in the scenarios listed here:

- Scenario G-1: Service Calls Between Resources – both the subject and the resource are located on enterprise-operated infrastructure (on-premises or branch).
- Scenario G-2: Service Calls to Cloud-Based Resources – the subject is located on enterprise-operated infrastructure while the resource is cloud-based.
- Scenario G-3: Service Calls between Cloud-Based Resources – both the subject and the resource are located in the cloud.
- Scenario G-4: Service Calls between Containers – the subject is either in another container in a single container runtime (e.g., Docker), in the same Kubernetes pod, or in a different Kubernetes pod from the requested resource.
- Scenario G-5: Service to Endpoint – an enterprise service attempts to access an enterprise-managed endpoint to perform some action (e.g., maintenance, reconfiguration, etc.).

6.2.8 Use Case H: Data Level Security Scenarios

Use Case H demonstrates data level security scenarios in which a subject requests access to data with different levels of classification. There are at least two different levels of data sensitivity, and a subject who is authorized to access a resource will be authorized either to have full access to the highest level of data or to have limited access to the data (e.g., low/limited/partial access) based on user identity, endpoint type, and other attributes as articulated in the following use cases:

- Scenario H-1: Full/Limited Access to Resource Data Based on Identity Attributes – the subject will be granted full or limited access to different levels of data based on their user identity attributes.
- Scenario H-2: Full/Limited Access to Resource Data Based on Requesting Endpoint – the subject will be granted full or limited access to different levels of data based on whether the requesting endpoint is enterprise-managed or BYOD.
- Scenario H-3: Internet Access restricted when Accessing High Level Data – while a subject has an active access session to a resource storing data with high classification, the enterprise will restrict that subject from accessing public internet resources.
- Scenario H-4: Accessing High Level Data Triggers MFA Challenge – if a subject already as an active access session with a resource and is accessing low-classification data, a request to access high-classification data at that resource will trigger a multi-factor authentication challenge.
- Scenario H-5: Just-in-Time Access to High-Level Data – the enterprise can grant a subject temporary access privileges to high-level data when needed.
- Scenario H-6: Operations Denied When Accessing High Level Data – a subject that is authorized to fully access (e.g., read and write) high classification data when using an enterprise-managed endpoint and located on premises or at a branch office can have their access privileges limited to read-only when using a BYOD or when located remote from enterprise infrastructure.
- Scenario H-7: High Classified Data Has Extra Protection When Stored on Endpoints – when a subject downloads or copies high classification data onto the subject's endpoint, the data is encrypted or has some further protection that requires the subject to pass a challenge before accessing or performing actions on the local copy of the data.

6.3 Functional Demonstration Results

The summary and detailed functional demonstration results are shown in the sections below.

6.3.1 Demonstration Result Summaries

6.3.1.1 EIG Crawl Phase

Three builds were implemented and demonstrated as part of the EIG crawl phase:

- E1B1 (EIG Crawl; Okta and Ivanti as PEs)
- E2B1 (EIG Crawl; Ping Identity as PE)
- E3B1 (EIG Crawl; Microsoft as PE)

The following scenarios were considered out of scope for the EIG Crawl Phase:

- Cloud-based,
- Stolen Credential,
- Just-in-Time Access Privileges,
- Enterprise-ID Step-Up Authentication,
- Federated-ID Access,

- Confidence Level, and
- Service-Service Interactions.

Summaries of the demonstration results for each of these builds can be found in our supplemental EIG Crawl Phase Summary Demonstration Results documentation.

6.3.1.2 EIG Run Phase

Three builds were implemented as part of the EIG run phase:

- E1B2 (EIG Run; Zscaler as PE)
- E3B2 (EIG Run; Microsoft and Forescout as PEs)
- E4B3 (EIG Run; IBM as PE)

The following scenarios were considered out of scope for the EIG Run Phase for builds E1B2 and E3B2:

- Just-in-Time Access Privileges,
- Enterprise-ID Step-Up Authentication,
- Federated-ID Access,
- Confidence Level, and
- Service-Service.

Summaries of the demonstration results for each of these builds can be found in our supplemental EIG Run Phase Summary Demonstration Results documentation.

6.3.1.3 SDP, Microsegmentation, and SASE Phase

Thirteen builds were implemented as part of the SDP, Microsegmentation, and SASE phase:

- E1B3 (SDP; Zscaler as PE)
- E2B3 (Microsegmentation; Cisco and Ping Identity as PEs)
- E3B3 (SDP and Microsegmentation; Microsoft and Forescout as PEs)
- E1B4 (SDP; Appgate as PE)
- E2B4 (SDP and SASE; Broadcom (with Symantec products) as PE)
- E3B4 (SDP; F5 as PE)
- E4B4 (SDP, Microsegmentation and EIG; Broadcom (with VMware products) as PE)
- E1B5 (Microsegmentation and SASE; Palo Alto Networks as PE)
- E2B5 (SDP and SASE; Lookout and Okta as PEs)
- E3B5 (SDP and SASE; Microsoft as PE)
- E4B5 (SDP and Microsegmentation; AWS as PE)
- E1B6 (SDP and Microsegmentation; Ivanti as PE)
- E2B6 (SASE; Google as PE)

All the use cases were in scope. Summaries of the demonstration results for each of these builds can be found in our supplemental SDP, Microsegmentation, and SASE Phase Summary Demonstration Results documentation.

6.3.2 Demonstration Results in Full

Table 6-1 identifies the policy engines/policy decision points and types of architecture used in each build. It also links to the online locations where each build architecture is described in detail, as well as the online locations where the full demonstration results for each build can be found.

Table 6-1 Mapping of Builds to Online Details Regarding Architecture Descriptions and Functional Demonstration Results

Build	Policy Engines/ Policy Decision Points	ZTA Architecture Instantiated	Links to Online Details: Build Architecture, Technologies, and Flow Diagrams	Links to Online Details: Full Demonstration Results
E1B1	Okta Identity Cloud Ivanti Access ZSO	EIG Crawl	E1B1 Build Architecture	E1B1 Full Demonstration Results
E2B1	Ping Identity Ping Federate	EIG Crawl	E2B1 Build Architecture	E2B1 Full Demonstration Results
E3B1	Azure AD (Conditional Access)	EIG Crawl	E3B1 Build Architecture	E3B1 Full Demonstration Results
E1B2	Zscaler ZPA Central Authority (CA)	EIG Run	E1B2 Build Architecture	E1B2 Full Demonstration Results
E3B2	Microsoft Azure AD (Conditional Access) Microsoft Intune Forescout eyeControl Forescout eyeExtend	EIG Run	E3B2 Build Architecture	E3B2 Full Demonstration Results
E4B3	IBM Security Verify	EIG Run	E4B3 Build Architecture	E4B3 Full Demonstration Results
E1B3	Zscaler ZPA Central Authority (CA)	SDP	E1B3 Build Architecture	E1B3 Full Demonstration Results

Build	Policy Engines/ Policy Decision Points	ZTA Architecture Instantiated	Links to Online Details: Build Architecture, Technologies, and Flow Diagrams	Links to Online Details: Full Demonstration Results
E2B3	Ping Identity PingFederate Cisco ISE Cisco Secure Workload	Microsegmentation	E2B3 Build Architecture	E2B3 Full Demonstration Results
E3B3	Microsoft Azure AD (Conditional Access) Microsoft Intune Microsoft Sentinel Forescout eyeControl Forescout eyeExtend	SDP and Microsegmentation	E3B3 Build Architecture	E3B3 Full Demonstration Results
E1B4	Appgate SDP Controller	SDP	E1B4 Build Architecture	E1B4 Full Demonstration Results
E2B4	Symantec Cloud Secure Web Gateway (Cloud SWG) Symantec ZTNA Symantec Cloud Access Security Broker (CASB)	SDP and SASE	E2B4 Build Architecture	E2B4 Full Demonstration Results
E3B4	F5 BIG-IP F5 NGINX Plus Forescout eyeControl Forescout eyeExtend	SDP	E3B4 Build Architecture	E3B4 Full Demonstration Results
E4B4	VMware Workspace ONE Access VMware Unified Access Gateway (UAG) VMware NSX-T	SDP, Microsegmentation, and EIG	E4B4 Build Architecture	E4B4 Full Demonstration Results
E1B5	Palo Alto Networks NGFW Palo Alto Networks Prisma Access	SASE and Microsegmentation	E1B5 Build Architecture	E1B5 Full Demonstration Results
E2B5	Lookout SSE Okta Identity Clouds	SDP and SASE	E2B5 Build Architecture	E2B5 Full Demonstration Results

Build	Policy Engines/ Policy Decision Points	ZTA Architecture Instantiated	Links to Online Details: Build Architecture, Technologies, and Flow Diagrams	Links to Online Details: Full Demonstration Results
E3B5	Microsoft Entra Conditional Access (formerly called Azure AD Conditional Access) Microsoft Security Service Edge	SDP and SASE	E3B5 Build Architecture	E3B5 Full Demonstration Results
E4B5	AWS Verified Access and Amazon VPC Lattice	SDP and Microsegmentation	E4B5 Build Architecture	E4B5 Full Demonstration Results
E1B6	Ivanti Neurons for Zero Trust Access	SDP and Microsegmentation	E1B6 Build Architecture	E1B6 Full Demonstration Results
E2B6	Google CEP – Access Context Manager	SASE	E2B6 Build Architecture	E2B6 Full Demonstration Results

7 Risk and Compliance Management

This section discusses risks addressed by the ZTA reference architecture and provides links to mappings of ZTA security characteristics to CSF Subcategories, NIST SP 800-53 security controls, and NIST critical software security measures. The mappings include both general ZTA logical component capabilities and specific ZTA example implementation vendor technology capabilities.

7.1 Risks Addressed by the ZTA Reference Architecture

Conventional network security has focused on perimeter defense. Historically, most organization resources have been located within and protected by the enterprise's network perimeter, which tended to be large and static. Subjects that are inside the network perimeter are often assumed to be implicitly trusted and are given broad access to the resources within the network perimeter. Attempts to access resources from outside the network perimeter, i.e., from the internet, are often subject to more scrutiny than those originating from within. However, a subject can be compromised regardless of whether it is inside or outside of the network perimeter. Once a subject is compromised, malicious actors—through impersonation and escalation—can gain access to the resources that the subject is authorized to access and move laterally within the network perimeter to access adjacent resources.

By protecting each resource individually and employing extensive identity, authentication, and authorization measures to verify a subject's requirement to access each resource, zero trust can ensure that authorized users, applications, and systems have access to only those resources that they absolutely have a need to access in order to perform their duties, not to a broad set of resources that all happen to be within the network perimeter. This way, if a malicious actor does manage to gain

unauthorized access to one resource, this access will not provide them with any advantage when trying to move laterally to other nearby resources. To compromise those other resources, the attacker would be required to figure out how to circumvent the mechanisms that are protecting those resources individually because it is not possible to reach those resources from nearby compromised resources. In this way, ZTA limits the insider threat because instead of having permission to access all resources within the network perimeter, malicious insiders would only be permitted to access those resources they require to perform their official roles.

In addition, once a subject is granted access to a resource, this access is often permitted to continue for a substantial period of time without being reevaluated based on a defined policy. The access session is often not monitored or subject to behavioral analysis, and the configuration and health of the devices being used to access resources may be subject to initial, but not ongoing, scrutiny. So, if a subject does manage to gain unauthorized access to a resource, the subject often has ample time to exfiltrate or modify valuable information or further compromise the resource and/or use it as a point from which to pivot and attack other corporate resources. ZTA limits these threats by performing continual verification of a subject's identity and authorization to access a resource. It may also perform behavioral analysis and validation of each system's health and configuration, and consider other factors such as day, time, and location of subject and resource. Based on the organization's defined policy, ZTA makes dynamic, ongoing assessments of the risk of each access request in real-time to ensure it poses an acceptable level of risk.

A number of trends, including cloud computing and remote work, have also introduced additional security threats. The growth in cloud computing has meant that enterprises are now storing critical resources (e.g., databases, applications, servers) in the cloud (i.e., outside of the traditional network perimeter) as well as on-premises. As a result, these resources cannot be protected by the network perimeter strategy. A new protection paradigm is needed that focuses on protecting resources individually, no matter where they are located, so that they are not at risk of being subjected to security policies that are not under organization control or not enforced consistently across all enterprise resources. Often the clouds in which resources are hosted are multitenant, meaning that different enterprises have authorized access to their own portions of the cloud infrastructure, with each tenant reliant on the cloud service provider to enforce this separation. If a malicious actor were to figure out how to subvert cloud security and move from one tenant's account to the next, the organization's resources would be at risk. Use of ZTA to protect each resource individually serves as further assurance that the resources will not be accessible to cloud users from other enterprises, nor will they be accessible to users from within the enterprise who do not have a need to access them.

The growth of the remote workforce, as well as collaboration with partners and dependence on contractors are other trends that are also challenging the conventional security paradigm. The subjects requesting authorized access to resources may not necessarily be within the network perimeter. They may be employees working from home or from a coffee shop's public Wi-Fi via the internet, or a partner, contractor, customer, or guest that requires access to some resources but must be restricted from accessing other resources. By relying on strong identity, authentication, and authorization services to determine precisely which resources a subject is authorized to access with respect to their role in or relationship to the organization, ZTA can restrict subjects to accessing only those resources that they have a need to access and ensure that they are not permitted to access any other resources.

While implementing ZTA addresses many risks, it also has limitations. It cannot remove all risks, and the ZTA implementation itself may introduce additional risks that need to be addressed. For more information on the limitations of ZTA, see Section 5 of NIST SP 800-207.

7.2 ZTA Security Mappings

A *mapping* indicates that one concept is related to another concept. This publication introduces mappings for ZTA cybersecurity functions, both those performed by the ZTA reference design's logical components (see Section 3.1) as well as those performed by specific technologies used in the project's builds.

Three categories of ZTA Security Mappings are available in our supplemental documentation:

- Subcategories from *the NIST Cybersecurity Framework 1.1 (CSF 1.1)* [2] and *The NIST Cybersecurity Framework 2.0 (CSF 2.0)* [3]. Note that mapping for CSF 1.1 was done only for the builds that were implemented before CSF 2.0 was finalized. Mapping for CSF 2.0 is done for all builds.
- Security controls from NIST SP 800-53r5 *(Security and Privacy Controls for Information Systems and Organizations)* [4].
- NIST critical software security measures.

These mappings describe how the functions in our ZTA reference design are related to the NIST reference documents within the context of our ZTA reference design. Within each category of mapping, there is both a general mapping from the ZTA reference design logical components to the document being mapped to (i.e., CSF, SP 800-53, or NIST critical software security measures), as well as a set of collaborator-specific mappings from the ZTA technology component capabilities that are included in one or more project builds to the document being mapped to (CSF, SP 800-53, or NIST critical software security measures).

The mappings were developed to support two primary use cases:

1. **Why should organizations implement ZTA?** This use case identifies how implementing ZTA can support an organization with achieving CSF Subcategories, SP 800-53 controls, and NIST critical software security measures. This helps communicate to an organization's senior management that expending resources to implement ZTA can also aid in fulfilling other security requirements.
2. **How can organizations implement ZTA?** This use case identifies how an organization's existing implementations of CSF Subcategories, SP 800-53 controls, and NIST critical software security measures can help support a ZTA implementation. An organization wanting to implement ZTA might first assess its current security capabilities so that it can plan how to add missing capabilities and enhance existing capabilities in order to implement ZTA. Organizations can leverage their existing security investments and prioritize future security technology deployment to address the gaps.

These mappings are intended to be used by any organization that is interested in implementing ZTA or that has begun or completed a ZTA implementation.

The project's mappings use the supportive relationship mapping style defined in Section 4.2 of NIST Internal Report (IR) 8477, *Mapping Relationships Between Documentary Standards, Regulations, Frameworks, and Guidelines: Developing Cybersecurity and Privacy Concept Mappings* [9]. This style uses three relationship types: Supports, Is Supported By, and Equivalent. Each relationship of type Supports or Is Supported By also has a property assigned to it: Example of, Integral to, or Precedes.

8 Zero Trust Journey Takeaways

Based on our experience building example implementations in the lab, we recommend that an organization that wants to deploy and implement zero trust embark on a journey that includes the following steps:

- Discover and Inventory the Existing Environment
- Formulate Access Policy to Support the Mission and Business Use Cases
- Identify Existing Security Capabilities and Technology
- Eliminate Gaps in Zero Trust Policy and Processes by Applying a Risk-Based Approach Based on the Value of Data
- Implement ZTA Components (People, Process, and Technology) and Incrementally Leverage Deployed Security Solutions
- Verify the Implementation to Support Zero Trust Outcomes
- Continuously Improve and Evolve Due to Changes in Threat Landscape, Mission, Technology, and Regulations

8.1 Discover and Inventory the Existing Environment

The first step any organization should take on its zero trust journey is to identify all of its assets by determining what resources it has in its existing environment (hardware, software, applications, data, and services). This may involve deploying tools that monitor traffic to discover what resources are active and being accessed and used. It is necessary to have a complete understanding and inventory of the organization's resources because these are the entities that the ZTA will be designed to protect. If resources are overlooked, it's likely that they won't be appropriately protected by the ZTA. They could be vulnerable to exfiltration, modification, deletion, denial-of-service, or other types of attack. It is imperative that all of the organization's resources, whether on-premises or cloud-based, be identified and inventoried.

Discovery tools that are used to identify organization resources may do so, for example, by monitoring transaction flows and communication patterns. These tools may also be useful in helping the organization identify the business and access rules that are currently being enforced and in identifying access patterns that business operations require. Understanding how resources are accessed, by whom, and in what context will help the organization formulate its access policies. In addition, once the organization has begun deploying a ZTA, continuing to use the discovery tools to observe the environment can be helpful to the organization as it audits and validates the ZTA on an ongoing basis.

8.2 Formulate Access Policy to Support the Mission and Business Use Cases

Once the organization has identified all the resources that it needs to protect and where they are, it may formulate the policies that the ZTA will enforce to specify who is allowed to access each resource and under what conditions. The access policies should be designed to ensure that permissions and authorizations to access each resource conform with the principles of least privilege and separation of duties. Typically, access to each resource will be denied by default, and access policies should be formulated to authorize subjects with the least privileges required in order to perform their assigned task on a resource that they are permitted to access. This requires understanding the types of users that will be accessing resources and their access requirements, work locations, employment arrangements, device types, and ownership models (e.g., BYOD and corporate-owned) because these will all influence policy creation. Access authorizations may be constrained according to the location of the individual requesting access, time of day, or other parameters that can further limit access without interfering with organizational operations. All access policies should be informed by the criticality of the resource being protected.

Initially, an organization may not have a clear sense of what resources each employee needs to access. They may not be aware of which employees are accessing which resources or whether or not such access conforms to the principles of least privilege and separation of duties. Information provided by the tools that were used to discover resources can be useful in this regard. They can monitor access patterns and produce a list of access flows and patterns that are observed. For the remote access example, an organization transitioning from a full device VPN to per-app tunneling could first set up a full device tunnel and observe traffic, then begin enabling only the traffic that is required for the user profile. The organization's security team can then examine this list to determine which access flows should be permitted and then formulate access rules that permit them. Any observed access flows that should not be permitted may be denied by default or explicitly prohibited in the access policy. By basing access policy on observed access patterns, an organization reduces the chances that it will create overly restrictive policies that interfere with its ability to conduct normal operations. By taking into consideration the criticality of the data being protected when formulating the access policy, an organization can help ensure that the protections being provided to a resource are commensurate with its value.

One challenge that organizations may have when formulating policy is that their ZTA may consist of numerous components that each perform policy engine and policy administrator roles. As a result, access policy may not be centralized; rules may be distributed across numerous products, i.e., with some rules configured in an endpoint protection component; some configured in ICAM components; other rules configured in a network security component; and still other rules configured in a data security component or other components. The lack of a single location where all policy rules can be centralized may make it challenging for an organization to maintain an organized, complete, consistent understanding of its access policy. To help manage their access policies, organizations should explicitly keep track of not only what their access rules are but also where each of the rules is configured.

8.3 Identify Existing Security Capabilities and Technology

If an organization is planning to install a ZTA into a greenfield environment, meaning that it will not have any existing IT equipment or security capabilities that it will want to use or accommodate, this step would not be needed. Most organizations embarking on a zero trust journey, however, will not be starting from scratch. Instead, they will have an existing infrastructure and technology systems that already perform security functions. Organizations will typically have at least network firewalls and intrusion detection systems to help provide perimeter security, and identity and credential access management systems that enable them to authenticate users and enforce authorized access based on identity and role. They may have endpoint security systems protecting their laptops and/or mobile devices to provide firewall protections and ensure that they are running required antivirus or other security software. They may have tools for vulnerability and configuration management, log management, and other security-related functions. They also likely have some sort of security operations center.

An organization should identify and inventory its existing security technology components and capabilities to understand what protections they already provide, then determine whether these components should continue to provide these protections as part of the deployed ZTA or should be repurposed. To save money, an organization will want to continue to use or repurpose as much of its existing technology as possible without sacrificing security. Continuing to use existing technology will require the organization to understand what potential zero trust components and products its existing security technology will integrate with. Any additional components that are purchased specifically for deployment in the ZTA should, ideally, integrate with the security technology components that the organization already has and plans to continue to use.

8.4 Eliminate Gaps in Zero Trust Policy and Processes by Applying a Risk-Based Approach Based on the Value of Data

Once an organization has inventories of the resources it needs to protect and the security capabilities it already has, the organization is ready to begin planning its access protection topology, in terms of whether and where its infrastructure will be segmented and at what level of granularity each resource will be protected. The access topology should be designed using a risk-based approach, isolating critical resources in their own trust zones protected by a PEP but permitting multiple lower-value resources to share a trust zone. In designing its access protection topology, the organization will identify which PEP is responsible for protecting each resource as well as what supporting technologies will be involved in providing input to resource access decisions.

Initially, the organization's network may not be well-segmented. In fact, before zero trust is implemented, when the organization is still relying on perimeter-based protections, such a topology can be thought of as the organization protecting all of its resources behind a single PEP, i.e., the perimeter firewall. As the organization implements ZTA, it should segment its infrastructure into smaller parts. Such segmentation will enable it to limit the potential impact of a breach or attack and make it easier to monitor network traffic. In designing its access protection topology, the organization should apply access control enforcement at multiple levels: application, host, and network.

8.5 Implement ZTA Components (People, Process, and Technology) and Incrementally Leverage Deployed Security Solutions

Once an organization has the following, it is ready to begin incrementally implementing ZTA:

- a good understanding of its current environment in terms of the resources it needs to protect and the security capabilities that it already has deployed;
- formulated the access policies that are appropriate to support its mission and business use cases; and
- designed its access protection topology to identify the granularity at which access to various resources will be protected and the supporting technologies that will provide input to the PDP.

Given the importance of discovery to the successful implementation of a ZTA, the organization may begin by deploying tools to continuously monitor the environment, if it has not done so already. The organization can use these observations to audit and validate the ZTA on an ongoing basis.

In addition to discovery tools, the organization should ensure that any other baseline security tools such as SIEMs, vulnerability scanning and assessment tools, and security validation tools are operational and configured to log, scan, assess, and validate the ZTA components that will be deployed. Having security baseline tools in place before the organization begins deploying new ZTA components helps ensure that the ZTA rollout will be well-monitored, enabling the organization to proceed with high confidence that it will understand the security ramifications of the incremental deployment as it proceeds.

Identity, authentication, and authorization are critical to making resource access decisions. Given that making and enforcing access decisions are the two main responsibilities of a ZTA, the organization will want to use its existing or a new ICAM solution as a foundational building block of its initial ZTA implementation. The organization should strongly consider implementing MFA in a risk-based manner for its users. An endpoint protection or similar solution that can assess device health and that integrates with the ICAM solution may also be another foundational component of an initial ZTA deployment. An initial ZTA based on these two main components will be able to use the identity and authorizations of subjects and the health and compliance of requesting endpoints as the basis for making access decisions. Additional supporting components and features can then be deployed to address an increasing number of ZTA requirements. Which types of components are deployed and in what order will depend on the organization's mission and business use cases. If data security is essential, then data security components will be prioritized; if behavior-based anomaly detection is essential, then monitoring and AI-based analytics may be installed. The ZTA can be built incrementally, adding and integrating more supporting components, features, and capabilities to gradually evolve to a more comprehensive ZTA.

8.6 Verify the Implementation to Support Zero Trust Outcomes

The organization should continue to monitor all network traffic in real time for suspicious activity, both to look for known attack signatures and patterns and to apply behavioral analytics to try to detect anomalies or other activity that may be attack indicators. The organization should use deployed discovery and other baseline security tools to audit and validate the access enforcement decision of the ZTA it has provisioned, correlating known data with information reported by the tools. The organization

should perform ongoing verification that the policies that are being enforced, as revealed by the observed network flows, are in fact the policies that the organization has defined. Periodic testing should be performed across a variety of use case scenarios, including those in which the resource is located on-premises and in the cloud, the requesting endpoint is located on-premises and on the internet, the requesting subject is and is not authorized to access the requested resource, the requesting endpoint is and is not managed, and the requesting resource is and is not compliant. In addition, service-to-service requests, both authorized and unauthorized, should also be tested. The use cases selected for testing should reflect those which most closely mirror how the organization's users access the organization's resources on a day-to-day basis. Ideally, the organization can create a suite of tests that it can use to validate the ZTA not only before deploying each new ZTA capability in the incremental rollout process, but also on a periodic basis once the ZTA rollout is considered complete.

8.7 Continuously Improve and Evolve Due to Changes in Threat Landscape, Mission, Technology, and Requirements

Once rolled out, the ZTA must continue to adapt to changing conditions. If technology components used in the ZTA are upgraded or obsoleted by their manufacturer, they should be replaced. If innovative new technologies become available, the organization should consider whether they could be integrated into the existing ZTA to take advantage of new defensive tactics, techniques, and procedures that might improve the organization's security posture. If the organization's security goals change, either as a result of a shifting mission or changes in regulations, the ZTA's policies and the ZTA itself may need to evolve to best address these new goals.

In addition, the ZTA may need to adapt to a changing threat landscape. As new types of adversary attacks become known and prevalent, the ZTA will need to add the threat signatures for these attacks to the list of things it monitors for. Ideally the ZTA will also perform behavior-based monitoring that enables it to detect anomalies that may signal zero-day attacks for which threat signatures are not yet known. Behavior-based monitoring tools provide the ZTA with some degree of agility and readiness with respect to its ability to detect attacks by adversaries who are constantly changing their tactics and techniques. In any case, as the threat landscape changes, the organization's CISO and security team need to continually assess the ZTA's topology, components, and policies to ensure that they are best designed to address newly emerging threats. If the value of one or more of an organization's resources increases substantially, the organization may want to change how that resource is protected by the ZTA, as well as what its access policies are.

As input to this ongoing process of validation and improvement, organizations should continuously monitor their network and other infrastructure and update policies, technologies, and network segmentation topologies to ensure that they remain effective. Creating a ZTA is not a one-time project but an ongoing process. The organization's CISO or other security team members should perform ongoing validation of their ZTA access policies to ensure that they continue to be defined in a manner that supports the organization's mission and business use cases while conforming with the principles of least privilege and separation of duties.

Appendix A List of Acronyms

AD	Active Directory
API	Application Programming Interface
BYOD	Bring Your Own Device
CASB	Cloud Access Security Broker
CISO	Chief Information Security Officer
CRADA	Cooperative Research and Development Agreement
CSF	Cybersecurity Framework
DLP	Data Loss Prevention
E1B1	Enterprise 1 Build 1
E1B2	Enterprise 1 Build 2
E1B3	Enterprise 1 Build 3
E1B4	Enterprise 1 Build 4
E1B5	Enterprise 1 Build 5
E1B6	Enterprise 1 Build 6
E2B1	Enterprise 2 Build 1
E2B3	Enterprise 2 Build 3
E2B4	Enterprise 2 Build 4
E2B5	Enterprise 2 Build 5
E2B6	Enterprise 2 Build 6
E3B1	Enterprise 3 Build 1
E3B2	Enterprise 3 Build 2
E3B3	Enterprise 3 Build 3
E3B4	Enterprise 3 Build 4
E3B5	Enterprise 3 Build 5
E4B3	Enterprise 4 Build 3
E4B4	Enterprise 4 Build 4
E4B5	Enterprise 4 Build 5
EDR	Endpoint Detection and Response
EIG	Enhanced Identity Governance

EPP	Endpoint Protection Platform
IaaS	Infrastructure as a Service
IaC	Infrastructure as Code
ICAM	Identity, Credential, and Access Management
IoT	Internet of Things
IP	Internet Protocol
IR	Internal Report
IT	Information Technology
ITL	Information Technology Laboratory
MDM	Mobile Device Management
MES	Mobile Endpoint Security
MFA	Multifactor Authentication
MTD	Mobile Threat Defense
NCCoE	National Cybersecurity Center of Excellence
NGFW	Next-Generation Firewall
NIC	Network Interface Card
NIST	National Institute of Standards and Technology
OT	Operational Technology
PA	Policy Administrator
PE	Policy Engine
PEP	Policy Enforcement Point
PIP	Policy Information Point
SaaS	Software as a Service
SASE	Secure Access Service Edge
SD-WAN	Software-Defined Wide Area Network
SDP	Software-Defined Perimeter
SIEM	Security Information and Event Management
SOAR	Security Orchestration, Automation, and Response
SP	Special Publication
SWG	Secure Web Gateway

UEM	Unified Endpoint Management
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPN	Virtual Private Network
WAN	Wide Area Network
ZTA	Zero Trust Architecture
ZTNA	Zero Trust Network Access

Appendix B References

[1] Rose S, Borchert O, Mitchell S, Connelly S (2020) Zero Trust Architecture. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-207. https://doi.org/10.6028/NIST.SP.800-207

[2] National Institute of Standards and Technology (2018) Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Cybersecurity White Paper (CSWP) NIST CSWP 6. https://doi.org/10.6028/NIST.CSWP.6

[3] National Institute of Standards and Technology (2024) The NIST Cybersecurity Framework (CSF) 2.0. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Cybersecurity White Paper (CSWP) NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

[4] Joint Task Force Interagency Working Group (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-53, Rev. 5. https://doi.org/10.6028/NIST.SP.800-53r5

[5] National Cybersecurity Center of Excellence, *Internet of Things (IoT)*. Available: https://www.nccoe.nist.gov/iot

[6] National Cybersecurity Center of Excellence, *Manufacturing*. Available: https://www.nccoe.nist.gov/manufacturing

[7] National Cybersecurity Center of Excellence, *Data Classification*. Available: https://www.nccoe.nist.gov/data-classification

[8] "National Cybersecurity Center of Excellence (NCCoE) Zero Trust Cybersecurity: Implementing a Zero Trust Architecture," Federal Register Vol. 85, No. 204, October 21, 2020, pp. 66936-66939. Available: https://www.federalregister.gov/documents/2020/10/21/2020-23292/national-cybersecurity-center-of-excellence-nccoe-zero-trust-cybersecurity-implementing-a-zero-trust.

[9] Scarfone K, Souppaya M, Fagan M (2024) Mapping relationships between documentary standards, regulations, frameworks, and guidelines: developing cybersecurity and privacy concept mappings. (National Institute of Standards and Technology (U.S.), Gaithersburg, MD), NIST Internal Report (IR) NIST IR 8477. https://doi.org/10.6028/NIST.IR.8477

Appendix C Change Log

In June 2025, the public comments received were addressed for the practice guide's final version.

In December 2024, the following changes were made for the practice guide's initial public draft:

- Added builds E2B6 and E4B5

In July 2024, the following changes were made for the practice guide's fourth preliminary draft:

- Introduced a new manner of content delivery in two formats, one we refer to as the "High-Level Document in PDF Format" and the other as the "Full Document in Web Format."
- Added builds E2B4, E3B4, E4B4, E1B5, E2B5, E3B5, and E1B6

In July 2023, the following changes were made for the practice guide's third preliminary draft:

- Added builds E1B3, E2B3, E3B3, E4B3, and E1B4

In December 2022, the following changes were made for the practice guide's second preliminary draft:

- Added builds E2B1, E1B2, and E3B2

In July 2022, the first preliminary draft was created with:

- Created original document including builds E1B1 and E3B1