NIST SPECIAL PUBLICATION 1800-35B

Implementing a Zero Trust Architecture

Volume B:
Approach, Architecture, and Security Characteristics

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

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

Karen Scarfone
Scarfone Cybersecurity
Clifton, VA

Jason Garbis
Peter Gallagher
Appgate
Coral Gables, FL

Adam Cerini
Conrad Fernandes
AWS (Amazon Web Services)
Arlington, VA

Kyle Black
Sunjeet Randhawa
Broadcom Software
San Jose, CA

Peter Romness
Steve Vetter
Cisco
Herndon, VA

Corey Bonnell
Dean Coclin
DigiCert
Lehi, UT

Ryan Johnson
Dung Lam
F5
Seattle, WA

Tim Jones
Tom May
Forescout
San Jose, CA

Tim Knudson
Google Cloud
Mill Valley, CA

Mike Spisak
Harmeet Singh
IBM
Armonk, NY

Corey Lund
Farhan Saifudin
Ivanti
South Jordan, UT

Hashim Khan
Tim LeMaster
Lookout
Reston, VA

Ken Durbin
Earl Matthews
Mandiant
Reston, VA

Clay Taylor
Tarek Dawoud
Microsoft
Redmond, WA

Vinu Panicker
Okta
San Francisco, CA

Sean Morgan
Palo Alto Networks
Santa Clara, CA

Zack Austin
PC Matic
Myrtle Beach, SC

Bryan Rosensteel
Mitchell Lewars
Ping Identity
Denver, CO

Wade Ellery
Deborah McGinn
Radiant Logic
Novato, CA

Frank Briguglio
Ryan Tighe
SailPoint
Austin, TX

Chris Jensen
Joshua Moll
Tenable
Columbia, MD

Jason White
Trellix, Public Sector
Reston, VA

Jacob Rapp
Paul Mancuso
VMware
Palo Alto, CA

Joe Brown
Jim Kovach
Zimperium
Dallas, TX

Bob Smith
Syed Ali
Zscaler
San Jose, CA

December 2022

SECOND PRELIMINARY DRAFT

This publication is available free of charge from
https://www.nccoe.nist.gov/projects/implementing-zero-trust-architecture

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.

While NIST and the NCCoE address goals of improving management of cybersecurity and privacy risk through outreach and application of standards and best practices, it is the stakeholder's responsibility to fully perform a risk assessment to include the current threat, vulnerabilities, likelihood of a compromise, and the impact should the threat be realized before adopting cybersecurity measures such as this recommendation.

National Institute of Standards and Technology Special Publication 1800-35B, Natl. Inst. Stand. Technol. Spec. Publ. 1800-35B, 185 pages, (December 2022), CODEN: NSPUE2

FEEDBACK

You can improve this guide by contributing feedback. As you review and adopt this solution for your own organization, we ask you and your colleagues to share your experience and advice with us.

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

Public comment period: December 21, 2022 through February 6, 2023

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) focuses on protecting data and resources. It 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. Each access request is evaluated by verifying the context available at access time, including criteria such as the requester's identity and role, the requesting device's health and credentials, the sensitivity of the resource, user location, and user behavior consistency. If the enterprise's defined access policy is met, a secure session is created to protect all information transferred to and from the resource. A real-time and continuous policy-driven,

risk-based assessment is performed to establish and maintain the access. In this project, the NCCoE and its collaborators use commercially available technology to build interoperable, open, standards-based ZTA implementations that align to the concepts and principles in NIST Special Publication (SP) 800-207, *Zero Trust Architecture*. This NIST Cybersecurity Practice Guide explains how commercially available technology can be integrated and used to build various ZTAs.

KEYWORDS

enhanced identity governance (EIG); identity, credential, and access management (ICAM); zero trust; zero trust architecture (ZTA).

ACKNOWLEDGMENTS

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

Name	Organization
Quint Van Deman	Amazon Web Services
Aaron Palermo	Appgate
Adam Rose	Appgate
Jonathan Roy	Appgate
Eric Michael	Broadcom Software
Ken Andrews	Cisco
Matthew Hyatt	Cisco
Leo Lebel	Cisco
Tom Oast	Cisco
Aaron Rodriguez	Cisco
Micah Wilson	Cisco
Daniel Cayer	F5

Name	Organization
David Clark	F5
Jay Kelley	F5
Yejin Jang	Forescout
Neal Lucier	Forescout
Andrew Campagna	IBM
Adam Frank	IBM
Nalini Kannan	IBM
Priti Patil	IBM
Nikhil Shah	IBM
Krishna Yellepeddy	IBM
Vahid Esfahani	IT Coalition
Ebadullah Siddiqui	IT Coalition
Musumani Woods	IT Coalition
Tyler Croak	Lookout
Madhu Dodda	Lookout
Jeff Gilhool	Lookout
James Elliott	Mandiant
David Pricer	Mandiant

Name	Organization
Joey Cruz	Microsoft
Janet Jones	Microsoft
Carmichael Patton	Microsoft
Hemma Prafullchandra	Microsoft
Brandon Stephenson	Microsoft
Sarah Young	Microsoft
Eileen Division	MITRE*
Spike Dog	MITRE
Sallie Edwards	MITRE
Ayayidjin Gabiam	MITRE
Jolene Loveless	MITRE
Karri Meldorf	MITRE
Kenneth Sandlin	MITRE
Lauren Swan	MITRE
Jessica Walton	MITRE
Mike Bartock	NIST
Douglas Montgomery	NIST
Kevin Stine	NIST

Name	Organization
Sean Frazier	Okta
Kelsey Nelson	Okta
Shankar Chandrasekhar	Palo Alto Networks
Andrew Keffalas	Palo Alto Networks
Seetal Patel	Palo Alto Networks
Norman Wong	Palo Alto Networks
Shawn Higgins	PC Matic
Andy Tuch	PC Matic
Rob Woodworth	PC Matic
Ivan Anderson	Ping Identity
Bill Baz	Radiant Logic
Don Coltrain	Radiant Logic
Rusty Deaton	Radiant Logic
John Petrutiu	Radiant Logic
Lauren Selby	Radiant Logic
Peter Amaral	SailPoint
Jim Russell	SailPoint
Esteban Soto	SailPoint

Name	Organization
Jeremiah Stallcup	Tenable
Bill Stritzinger	Tenable
Andrew Babakian	VMware
Peter Bjork	VMware
Genc Domi	VMware
Keith Luck	VMware
Dennis Moreau	VMware*
Jeffrey Adorno	Zscaler
Jeremy James	Zscaler
Lisa Lorenzin	Zscaler*
Matt Moulton	Zscaler
Patrick Perry	Zscaler

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

The Technology Partners/Collaborators who have or will participate in this project’s current or upcoming builds 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. We are working with the following list of collaborators.

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

CALL FOR PATENT CLAIMS

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be directly stated in this ITL Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.

ITL may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or

b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL draft publication either:

1. under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or

2. without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to: nccoe-zta-project@list.nist.gov

Contents

1 Summary 1
1.1 Challenge 1
1.2 Solution 2
1.3 Benefits 3
2 How to Use This Guide 4
2.1 Typographic Conventions 6
3 Approach 7
3.1 Audience 8
3.2 Scope 9
3.3 Assumptions 10
3.4 Collaborators and Their Contributions 10
3.4.1 Appgate 11
3.4.2 AWS 12
3.4.3 Broadcom Software 14
3.4.4 Cisco 17
3.4.5 DigiCert 19
3.4.6 F5 20
3.4.7 Forescout 22
3.4.8 Google Cloud 23
3.4.9 IBM 24
3.4.10 Ivanti 26
3.4.11 Lookout 28
3.4.12 Mandiant 28
3.4.13 Microsoft 29
3.4.14 Okta 33
3.4.15 Palo Alto Networks 35
3.4.16 PC Matic 37
3.4.17 Ping Identity 37

3.4.18 Radiant Logic ... 40
3.4.19 SailPoint ... 42
3.4.20 Tenable ... 43
3.4.21 Trellix ... 44
3.4.22 VMware ... 47
3.4.23 Zimperium ... 47
3.4.24 Zscaler ... 48

4 Architecture ... 49

4.1 General ZTA Reference Architecture ... 49
4.1.1 ZTA Core Components ... 50
4.1.2 ZTA Supporting Components ... 51
4.1.3 ZTA in Operation ... 54
4.2 EIG Crawl Phase Reference Architecture ... 59
4.3 EIG Run Phase ... 62
4.4 ZTA Laboratory Physical Architecture ... 62
4.4.1 Enterprise 1 ... 64
4.4.2 Enterprise 1 Branch Office ... 69
4.4.3 Enterprise 2 ... 71
4.4.4 Enterprise 3 ... 71
4.4.5 Enterprise 4 ... 71
4.4.6 Coffee Shop ... 71
4.4.7 Management and Orchestration Domain ... 71
4.4.8 Emulated WAN Service Provider ... 72
4.4.9 Cloud Services ... 72

5 Functional Demonstration ... 77

6 General Findings ... 77

6.1 EIG Crawl Phase Findings ... 77
6.2 EIG Run Phase Findings ... 78

7 Future Build Considerations 80
Appendix A List of Acronyms 81
Appendix B Glossary 87
Appendix C References 89
Appendix D EIG Enterprise 1 Build 1 (E1B1) 90
D.1 Technologies 90
D.2 Build Architecture 94
D.2.1 Logical Architecture 94
D.2.2 ICAM Information Architecture 95
D.2.3 Physical Architecture 107
D.2.4 Message Flow for a Successful Resource Access Request 107
Appendix E EIG Enterprise 2 Build 1 (E2B1) 111
E.1 Technologies 111
E.2 Build Architecture 115
E.2.1 Logical Architecture 115
E.2.2 ICAM Information Architecture 116
E.2.3 Physical Architecture 128
E.2.4 Message Flow for a Successful Resource Access Request 128
Appendix F EIG Enterprise 3 Build 1 (E3B1) 131
F.1 Technologies 131
F.2 Build Architecture 135
F.2.1 Logical Architecture 135
F.2.2 Physical Architecture 136
F.2.3 Message Flows for a Successful Resource Access Request 136
Appendix G EIG Enterprise 4 Build 1 (E4B1) 141
Appendix H EIG Enterprise 1 Build 2 (E1B2) 142
H.1 Technologies 142

H.2 Build Architecture 146
H.2.1 Logical Architecture 146
H.2.2 ICAM Information Architecture 147
H.2.3 Physical Architecture 148
H.2.4 Message Flows for Successful Resource Access Requests 148
Appendix I EIG Enterprise 2 Build 2 (E2B2) 155
Appendix J EIG Enterprise 3 Build 2 (E3B2) 156
J.1 Technologies 156
J.2 Build Architecture 161
J.2.1 Logical Architecture 161
J.2.2 Physical Architecture 162
J.2.3 Message Flows for a Successful Resource Access Request 162

List of Figures

Figure 4-1 General ZTA Reference Architecture 50
Figure 4-2 EIG Crawl Phase Reference Architecture 61
Figure 4-3 Physical Architecture of ZTA Lab 63
Figure 4-4 Physical Architecture of Enterprise 1 65
Figure 4-5 Shared Services Domain of Enterprise 1 67
Figure 4-6 Physical Architecture of the Enterprise 1 Branch Office 70
Figure 4-7 Physical Architecture of the Coffee Shop 71
Figure 4-8 Physical Architecture of the Management and Orchestration Domain 72
Figure 4-9 Physical Architecture of the AWS Infrastructure Used by Enterprise 1 74
Figure 4-10 Physical Architecture of the Azure Infrastructure Used by Enterprise 3 76
Figure D-1 Logical Architecture of E1B1 95
Figure D-2 E1B1 ICAM Information Architecture – Identity Correlation 98
Figure D-3 E1B1 ICAM Information Architecture – New User Onboarding 101
Figure D-4 E1B1 ICAM Information Architecture - User Changes Roles 104

Figure D-5 E1B1 ICAM Information Architecture - User Termination 106
Figure D-6 Successful Access Request Enforced by Okta, Ivanti, and Zimperium Components 108
Figure E-1 Logical Architecture of E2B1 116
Figure E-2 E2B1 ICAM Information Architecture – Identity Correlation 119
Figure E-3 E2B1 ICAM Information Architecture – New User Onboarding 122
Figure E-4 E2B1 ICAM Information Architecture - User Changes Roles 125
Figure E-5 E2B1 ICAM Information Architecture - User Termination 127
Figure E-6 Use Case—E2B1 – Access Enforced by Ping Federate, Cisco Duo, and Radiant Logic 129
Figure F-1 Logical Architecture of E3B1 136
Figure F-2 Use Case—E3B1 – Access Enforced by Azure AD 138
Figure F-3 Use Case—E3B1 – Access Enforced by F5 BIG-IP 139
Figure H-1 Logical Architecture of E1B2 147
Figure H-2 Access to an Internal Resource is Enforced by Zscaler ZPA and Okta Identity Cloud 150
Figure H-3 Access to an Externally-Facing Resource is Enforced by Zscaler ZIA and Okta Identity Cloud 152
Figure J-1 Logical Architecture of E3B2 162
Figure J-2 Use Case— E3B2 – Access to an Internal Resource is Enforced by Azure AD and Azure AD's Application Proxy 165
Figure J-3 Use Case— E3B2 – Access to an Externally-Facing Resource is Enforced by Azure AD and Microsoft Defender for Cloud Apps 167
Figure J-4 Use Case—E3B2 – Forescout Discovers a Non-Compliant Endpoint on the Network and Directs the Palo Alto Firewall to Block it 169

List of Tables

Table 3-1 Technology Partners/Collaborators 11
Table 4-1 Mapping of Builds to Architectures and Appendices 64
Table D-1 E1B1 Products and Technologies 90
Table E-1 E2B1 Products and Technologies 111
Table F-1 E3B1 Products and Technologies 131

Table H-1 E1B2 Products and Technologies 142
Table J-1 E3B2 Products and Technologies 156

1 Summary

1.1 Challenge

Protecting enterprise resources, particularly data, has become increasingly challenging as resources have become distributed across both on-premises environments and multiple clouds. Many users need access from anywhere, at any time, from any device to support the organization's mission. Data is programmatically stored, transmitted, and processed across different boundaries under the control of different organizations to meet ever-evolving business use cases. It is no longer feasible to simply enforce access controls at the perimeter of the enterprise environment and assume that all subjects[1] (e.g., end users, applications, and other non-human entities that request information from resources) within it can be trusted. A zero-trust architecture (ZTA) addresses this challenge by enforcing granular, secure authorized access near the resources, whether located on-premises or in the cloud, for both remote and onsite workforces and partners based on an organization's defined access policy.

Many organizations would like to address these challenges by migrating to a ZTA, but they have been hindered by several factors, which may include:

- 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
- Ever-increasing complexity of communication flows and distributed IT components across the environments on-premises and in the cloud, making them difficult to manage consistently
- 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 awareness regarding everything that encompasses the organization's entire 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

[1] As with NIST Special Publication (SP) 800-207 [1], throughout this document *subject* will be used unless the section relates directly to a human end user, in which case *user* will be used instead of the more generic *subject*.

ZTA may extend beyond their normal purview. False assumptions are often made in understanding the health of a device as well as its exposure to supply chain risks.

- Lack of understanding regarding what interoperability issues may be involved or what additional skills and training administrators, security personnel, operators, end users, and policy decision makers may require; 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
- Integrating various types of commercially available technologies of varying maturities, assessing capabilities, and identifying technology gaps to build a complete ZTA
- Concern that ZTA might negatively impact the operation of the environment or end-user experience
- Lack of a standardized policy to distribute, manage, and enforce security policy, causing organizations to face either a fragmentary policy environment or non-interoperable components
- 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
- 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
- There is not a single ZTA that fits all. ZTAs need to be designed and integrated for each organization based on the organization's requirements and risk tolerance, as well as its existing invested technologies and environments.

1.2 Solution

This project is designed to help address the challenges discussed above by building, demonstrating, and documenting several example ZTAs using products and technologies from a variety of different vendors. The example solutions are designed to provide secure authorized access to individual resources by enforcing enterprise security policy dynamically and in near-real-time. They restrict access to authenticated, authorized users and devices while flexibly supporting a complex set of diverse business use cases. These use cases involve legacy enterprise networks; remote workforces; use of the cloud; use of corporate-provided, bring your own device (BYOD), and guest endpoints; collaboration with partners; guest users; and support for contractors and other authorized third parties. The example solutions are also designed to demonstrate having visibility within the various environments as well as recognizing both internal and external attacks and malicious actors. They showcase the ability of ZTA products to interoperate with legacy enterprise and cloud technologies to protect resources with minimal impact on end-user experience.

The concepts and principles in NIST SP 800-207, *Zero Trust Architecture* are applied to enterprise networks that are composed of pre-established devices and components and that store critical corporate assets and resources both on-premises and in the cloud. For each data access session requested, ZTA verifies the requester's identity, role, and authorization to access the requested assets, the requesting device's health and credentials, and possibly other information. If defined policy is met, ZTA dynamically creates a secure connection to protect all information transferred to and from the accessed resource. ZTA performs real-time, continuous behavioral analysis and risk-based assessment of the access transaction or session.

The example solutions, which are based on reference architectures, are built starting with a baseline designed to resemble a notional existing enterprise environment that is assumed to have an identity store and other security components in place. This enables the project to represent how a typical enterprise is expected to evolve toward ZTA, i.e., by starting with their already-existing legacy enterprise environment and gradually adding capabilities. A limited version of the enhanced identity governance (EIG) deployment approach described in NIST SP 800-207 was implemented first, during what we refer to as the EIG crawl phase of the project. The first iteration of ZTA implementations is based on the EIG approach because EIG is a foundational component of the other deployment approaches utilized in today's hybrid environments. The EIG approach uses the identity of subjects and device health as the main determinants of policy decisions. However, instead of using a separate, dedicated component to serve as a policy decision point (PDP), our crawl phase leveraged the identity, credential, and access management (ICAM) components to serve as the PDP.

After completing the example solutions that were implemented as part of the EIG crawl phase of the project, the EIG run phase was begun. In the EIG run phase, an EIG approach that is not limited to using an ICAM component as the PDP is being implemented. After that, additional supporting components and features will be deployed to address an increasing number of the ZTA requirements, progressing the project toward eventual demonstration of the micro-segmentation and software-defined perimeter deployment options.

1.3 Benefits

The demonstrated approach documented in this practice guide can provide organizations wanting to migrate to ZTA with information and confidence that will help them develop transition plans for integrating ZTA into their own legacy environments, based on the example solutions and using a risk-based approach. Executive Order 14028, *Improving the Nation's Cybersecurity* [2], requires all federal agencies to develop plans to implement ZTA. This practice guide can inform agencies in developing their ZTA implementation plans. When integrated into their enterprise environments, ZTA will enable organizations to:

- **Support teleworkers** by enabling them to securely access corporate resources regardless of their location—on-premises, at home, or on public Wi-Fi at a neighborhood coffee shop.
- **Protect resources and assets** regardless of their location—on-premises or in the cloud.
- **Provision healthy devices from vendors** that can verify that the device is authentic and free of known exploitable vulnerabilities.
- **Limit the insider threat** by rejecting the outdated assumption that any user located within the network boundary should be automatically trusted and by enforcing the principle of least privilege.
- **Limit breaches** by reducing an attacker's ability to move laterally in the network. Access controls can be enforced on an individual resource basis, so an attacker who has access to one resource won't be able to use it as a springboard for reaching other resources.
- **Improve incident detection, response, and recovery** to minimize impact when breaches occur. Limiting breaches reduces the footprint of any compromise and the time to recovery.
- **Protect sensitive corporate data** by using strong encryption both while data is in transit and while it is at rest. Grant subjects' access to a specific resource only after enforcing consistent identification, authentication, and authorization procedures, verifying device health, and performing all other checks specified by enterprise policy.
- **Improve visibility** into which users are accessing which resources, when, how, and from where by monitoring and logging every access request within every access session.
- **Perform dynamic, risk-based assessment** of resource access through continuous reassessment of all access transactions and sessions, gathering information from periodic reauthentication and reauthorization, ongoing device health and posture verification, behavior analysis, ongoing resource health verification, anomaly detection, and other security analytics.

2 How to Use This Guide

This NIST Cybersecurity Practice Guide will help users develop a plan for migrating to ZTA. It demonstrates a standards-based ZTA reference design and provides users with the information they need to replicate one or more standards-based ZTA implementations that align to the concepts and principles in NIST SP 800-207, *Zero Trust Architecture*. This reference design is modular and can be deployed in whole or in part, enabling organizations to incorporate ZTA into their legacy environments gradually, in a process of continuous improvement that brings them closer and closer to achieving the ZTA goals that they have prioritized based on risk, cost, and resources.

NIST is adopting an agile process to publish this content. Each volume is being made available as soon as possible rather than delaying release until all volumes are completed. Work continues on implementing the example solutions and developing other parts of the content. As a second preliminary draft, we will publish at least one additional draft of this volume for public comment before it is finalized.

When complete, this guide will contain five volumes:

- NIST SP 1800-35A: *Executive Summary* – why we wrote this guide, the challenge we address, why it could be important to your organization, and our approach to solving this challenge
- NIST SP 1800-35B: *Approach, Architecture, and Security Characteristics* – what we built and why **(you are here)**
- NIST SP 1800-35C: *How-To Guides* – instructions for building the example implementations, including all the security-relevant details that would allow you to replicate all or parts of this project
- NIST SP 1800-35D: *Functional Demonstrations* – use cases that have been defined to showcase ZTA security capabilities and the results of demonstrating them with each of the example implementations
- NIST SP 1800-35E: *Risk and Compliance Management* – risk analysis and mapping of ZTA security characteristics to cybersecurity standards and recommended practices

Depending on your role in your organization, you might use this guide in different ways:

Business decision makers, including chief security and technology officers, will be interested in the *Executive Summary, NIST SP 1800-35A*, which describes the following topics:

- challenges that enterprises face in migrating to the use of ZTA
- example solution built at the NCCoE
- benefits of adopting the example solution

Technology or security program managers who are concerned with how to identify, understand, assess, and mitigate risk will be interested in this part of the guide, *NIST SP 1800-35B*, which describes what we did and why. Also, Section 3 of *Risk and Compliance Management*, *NIST SP 1800-35E,* will be of particular interest. Section 3, ZTA Reference Architecture Security Mappings, maps logical components of the general ZTA reference design to security characteristics listed in various cybersecurity guidelines and recommended practices documents, including *Framework for Improving Critical Infrastructure Cybersecurity* (NIST Cybersecurity Framework), *Security and Privacy Controls for Information Systems and Organizations* (NIST SP 800-53), and *Security Measures for "EO-Critical Software" Use Under Executive Order (EO) 14028*.

You might share the *Executive Summary, NIST SP 1800-35A*, with your leadership team members to help them understand the importance of migrating toward standards-based ZTA implementations that align to the concepts and principles in NIST SP 800-207, *Zero Trust Architecture*.

IT professionals who want to implement similar solutions will find the whole practice guide useful. You can use the how-to portion of the guide, *NIST SP 1800-35C*, to replicate all or parts of the builds created in our lab. The how-to portion of the guide provides specific product installation, configuration, and integration instructions for implementing the example solution. We do not re-create the product manufacturers' documentation, which is generally widely available. Rather, we show how we incorporated the products together in our environment to create an example solution. Also, you can use *Functional Demonstrations, NIST SP 1800-35D*, which provides the use cases that have been defined to showcase ZTA security capabilities and the results of demonstrating them with each of the example implementations.

This guide assumes that IT professionals have experience implementing security products within the enterprise. While we have used a suite of commercial products to address this challenge, this guide does not endorse these particular products. Your organization can adopt this solution or one that adheres to these guidelines in whole, or you can use this guide as a starting point for tailoring and implementing parts of a ZTA. Your organization's security experts should identify the products that will best integrate with your existing tools and IT system infrastructure. We hope that you will seek products that are congruent with applicable standards and best practices. The example solutions in this guide are not intended to be wholly implemented by most enterprise organizations because each organization's transition to ZT will depend on the organization's risk profile and tolerance, among other factors.

A NIST Cybersecurity Practice Guide does not describe "the" solution, but example solutions. This is a second preliminary draft guide. As the project progresses, this second preliminary draft will be updated, and additional volumes will also be released for comment. We seek feedback on the publication's contents and welcome your input. Comments, suggestions, and success stories will improve subsequent versions of this guide. Please contribute your thoughts to nccoe-zta-project@list.nist.gov.

2.1 Typographic Conventions

The following table presents typographic conventions used in this volume.

Typeface/Symbol	Meaning	Example
Italics	file names and path names; references to documents that are not hyperlinks; new terms; and placeholders	For language use and style guidance, see the *NCCoE Style Guide*.
Bold	names of menus, options, command buttons, and fields	Choose **File** > **Edit**.
`Monospace`	command-line input, onscreen computer output, sample code examples, and status codes	`mkdir`
`Monospace Bold`	command-line user input contrasted with computer output	**`service sshd start`**
blue text	link to other parts of the document, a web URL, or an email address	All publications from NIST's NCCoE are available at https://www.nccoe.nist.gov.

3 Approach

The NCCoE issued an open invitation to technology providers to participate in demonstrating approaches to deploying ZTA in a typical enterprise network environment. The objective was to use commercially available technology to produce example ZTA implementations that manage secure access to corporate resources hosted on-premises or in the cloud while supporting access from anywhere, at any time, using any device.

The NCCoE prepared a Federal Register Notice [3] inviting technology providers to provide products and/or expertise to compose prototype ZTAs. Core components sought included ZTA policy engines, policy administrators, and policy enforcement points. Supporting components supporting data security, endpoint security, identity and access management, and security analytics were also requested. In addition, device and network infrastructure components such as laptops, tablets, and other devices that connect to the enterprise were sought, as were data and compute resources, applications, and services that are hosted and managed on-premises, in the cloud, at the edge, or some combination of these. The NCCoE provided a network infrastructure that was designed to encompass the existing (non-ZTA) network resources that a medium or large enterprise might typically have deployed, and the ZTA core and supporting components and devices were integrated into this.

Cooperative Research and Development Agreements (CRADAs) were established with qualified respondents, and build teams were assembled. The build teams fleshed out the initial architectures, and the collaborators' components have so far been composed into five example implementations (i.e., builds), with several other builds in progress and additional future builds planned. With twenty-four collaborators participating in the project, the build teams that were assembled sometimes included vendors that offer overlapping capabilities. We made an effort to showcase capabilities from each vendor when possible. In other cases, we worked with the collaborators to have them work out a solution. Each build team documented the architecture and design of its build. As each build progressed, its team documented the steps taken to install and configure each component of the build. The teams then conducted functional demonstrations of the builds, including the ability to securely manage access to resources across a set of use cases that were defined to exercise a wide variety of typical enterprise situations. Use cases for the project include the following:

- access by employees, privileged third parties, and guests
- access requested by users who are located at headquarters, a branch office, or teleworking via public Wi-Fi and the internet
- inter-server access
- protection of resources that are located both on-premises and in the cloud
- use of enterprise-managed devices, contractor-managed devices, and personal devices
- access of both corporate resources and publicly available internet services

- the ability to automatically and dynamically calculate fine-grained confidence levels for resource access requests

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. First, we designed and built three implementations of the EIG crawl phase deployment approach using a variety of commercial products. Next, we build two implementations of the EIG run phase deployment approach.

Given the importance of discovery to the successful implementation of a ZTA, as part of the baseline environment we deployed tools that could be run 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 initial ZT policy, ultimately ensuring that observed network flows correlate to static policies.

EIG uses the identity of subjects and device health as the main determinants of policy decisions. Depending on the current state of identity management in the enterprise, deploying EIG solutions is an initial key step that will be leveraged to support the micro-segmentation and software-defined perimeter (SDP) deployment approaches, which will be covered in the later phases of the project. Our strategy is to follow an agile implementation methodology to build everything iteratively and incrementally while adding more capabilities to evolve to a complete ZTA. We started with the minimum viable EIG solution that allowed us to achieve some level of ZTA and then we gradually deploy additional supporting components and features to address an increasing number of the ZTA requirements, progressing the project toward eventual demonstration of more robust micro-segmentation and SDP deployment options.

3.1 Audience

The focus of this project is on medium and large enterprises. Its solution is targeted to address the needs of these enterprises, which are assumed to have a legacy network environment and trained operators and network administrators. These operators and administrators are assumed to have the skills to deploy ZTA components as well as related 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. The reader is assumed to be familiar with NIST SP 800-207, *Zero Trust Architecture*.

3.2 Scope

The scope of this project is initially limited to implementing a ZTA for a conventional, general-purpose enterprise information technology (IT) infrastructure that combines users (including employees, partners, contractors, guests, customers, and non-person entities [NPEs]), devices, and enterprise resources. Resources could be hosted and managed—by the corporation itself or a third-party provider—either on-premises or in the cloud, or some combination of these. There may also be branch or partner offices, teleworkers, and support for fully managed BYOD and non-managed (i.e., guest) device usage. While mobile device management (MDM) is used to support these device types, demonstrating the full spectrum of MDM capabilities is beyond the scope of this project. Initially, support for traditional IT resources such as laptops, desktops, servers, and other systems with credentials is within scope. In future phases, the scope may expand to include ZTA support for Internet of Things (IoT) devices. ZTA support for both IPv4 and IPv6 is in scope, as are the three deployment approaches of EIG, micro-segmentation, and SDP, and both agent and agentless implementations.

It is important to establish the trustworthiness of ZTA component devices to mitigate the possibility that the ZTA will be vulnerable to compromise through the hardware or software supply chain, but discussion of methods for establishing and maintaining the trustworthiness of the underlying hardware and supporting software comprising the ZTA is outside the scope of this document. Also, this document is only concerned with using the ZTA to protect access to enterprise data. Addressing the risk and policy requirements of discovering and classifying the data is out of scope.

This project focuses primarily on various types of user access to enterprise resources sprinkled across a hybrid network environment. More specifically, the focus is on behaviors of enterprise employees, partners, contractors, and guests accessing enterprise resources while connected from the corporate (or enterprise headquarters) network, a branch office, or the public internet. Access requests can occur over both the enterprise-owned part of the infrastructure and the public/non-enterprise-owned part. This requires that all access requests be secure, authorized, and verified before access is enforced, regardless of where the request is initiated or where the resources are located, i.e., whether on-premises or in the cloud. Discovery of resources, assets, communication flows, and other elements is also within scope.

ZTAs for industrial control systems and operational technology (OT) environments are explicitly out of scope for this project. However, the project seeks to provide an approach and security principles for a ZTA that could potentially be extended to OT environments. Any such application of ZTA principles to OT environments would be part of a separate project. Please refer to other related NCCoE projects [4][5][6][7]. The project is not concerned with addressing Federal Risk and Authorization Management Program (FedRAMP) or other federal requirements at this time, although doing so could potentially be a follow-on exercise.

Only implementations of the EIG crawl and EIG run phase deployment approaches are within scope at this time. Builds of more complex ZTAs will be undertaken in later phases of the project.

3.3 Assumptions

This project is guided by the following assumptions:

- NIST SP 800-207, *Zero Trust Architecture* is a definitive source of ZTA concepts and principles.
- Enterprises that want to migrate gradually to an increasing use of ZTA concepts and principles in their network environments may desire to integrate ZTA with their legacy enterprise and cloud systems.
- To prepare for a migration to ZTA, enterprises may want to inventory and prioritize all resources that require protection based on risk. They will also need to define policies that determine under what set of conditions subjects will be given access to each resource based on attributes of both the subject and the resource (e.g., location, type of authentication used, user role), as well as other variables such as day and time.
- Enterprises should use a risk-based approach to set and prioritize milestones for their gradual adoption and integration of ZTA across their enterprise environment.
- There is no single approach for migrating to ZTA that is best for all enterprises.
- There is not necessarily a clear point at which an organization can be said to have achieved a state of “full” or 100% ZTA compliance. Continuous improvement is the objective.
- Devices, applications, and other non-human entities can have different levels of capability:
    - Neither host-based firewalls nor host-based intrusion prevention systems (IPS) are mandatory components; they are, however, capabilities that can be added when a device is capable of supporting them.
    - Some limited functionality devices that are not able to host firewall, IPS, and other capabilities on their own may be associated with services that provide these capabilities for them. In this case, both the device and its supporting services can be considered the subject in the ZTA access interaction.
    - Some devices are bound to users (e.g., desktop, laptop, smartphone); other devices are not bound to users (e.g., servers, applications, services). Both types of devices can be subjects and request access to enterprise resources.
- ZTA components used in any given enterprise solution should be interoperable regardless of their vendor origin.

3.4 Collaborators and Their Contributions

Organizations participating in this project submitted their capabilities in response to an open call in the Federal Register for all sources of relevant security capabilities from academia and industry (vendors and integrators). The following respondents with relevant capabilities or product components (identified as "Technology Partners/Collaborators" herein) signed a CRADA to collaborate with NIST in a consortium to build example ZTA solutions:

Table 3-1 Technology Partners/Collaborators

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

Each of these technology partners and collaborators, as well as the relevant products and capabilities they bring to this ZTA effort, are described in the following subsections. The NCCoE does not certify or validate products or services. We demonstrate the capabilities that can be achieved by using participants' contributed technology.

3.4.1 Appgate

Appgate is the secure access company. It empowers how people work and connect by providing solutions purpose-built on zero trust security principles. This security approach enables fast, simple, and secure connections from any device and location to workloads across any IT infrastructure in cloud, on-premises, and hybrid environments.

3.4.1.1 Appgate SDP

The Appgate SDP solution has been designed with the intent to provide all the critical elements of NIST SP 800-207. The Appgate SDP has a controller that offers policy administrator (PA) and policy engine (PE) functionality and gateways that offer policy enforcement point (PEP) functionality. Appgate SDP natively integrates with components via representational state transfer (REST) application programming interfaces (APIs) and metadata. By providing highly performant, scalable, secure, integrated, and cloaked zero trust access, Appgate SDP is able to ensure that the correct device and user (under the appropriate conditions at that moment in time) are connected. For more information about Appgate SDP, see https://www.appgate.com/zero-trust-network-access/how-it-works.

3.4.2 AWS

AWS provides a platform in the cloud that hosts private and public sector agencies in most countries around the world. AWS offers more than 200 services which include compute, storage, networking, database, analytics, application services, deployment, management, developer, mobile, IoT, artificial intelligence (AI), security, and hybrid and enterprise applications. Additionally, AWS provides several security-related services and features such as Identity and Access Management (IAM), Virtual Private Cloud (VPC), PrivateLink, and Security Hub, allowing AWS customers to build and deliver their services worldwide with a high degree of confidence and assurance. AWS's array of third-party applications provides complementary functionality that further extends the capabilities of the AWS environment. To learn more about security services and compliance on AWS, please visit: https://aws.amazon.com/products/security.

The following subsections briefly list some AWS services relevant to ZTA that are being provided in support of this project, organized by category of service.

3.4.2.1 Identity

IAM: AWS Identity and Access Management (IAM) provides fine-grained access control across all of AWS. With IAM, organizations can specify who can access which services and resources, and under which conditions. With IAM policies, organizations manage permissions to their workforce and systems to ensure least-privilege permissions.

Cognito: Amazon Cognito lets organizations add user sign-up, sign-in, and access control to web and mobile apps quickly and easily. Cognito scales to millions of users and supports sign-in with social identity providers, such as Apple, Facebook, Google, and Amazon, and enterprise identity providers via Security Assertion Markup Language (SAML) 2.0 and OpenID Connect.

3.4.2.2 Network/Network Security

VPC: Amazon Virtual Private Cloud (Amazon VPC) gives organizations full control over their virtual networking environment, including resource placement, connectivity, and security. A couple of key security features found in VPCs are network access control lists (ACLs) that act as firewalls for controlling traffic in and out of subnets, and security groups that act as host-based firewalls for controlling traffic to individual Amazon Elastic Compute Cloud (Amazon EC2) instances.

PrivateLink: AWS PrivateLink provides private connectivity between VPCs, AWS services, and on-premises networks without exposing traffic to the public internet. AWS PrivateLink makes it easy to connect services across different accounts and VPCs to significantly simplify network architecture.

Network Firewall: AWS Network Firewall is a managed service that makes it easy to deploy essential network protections for all of an organization's Amazon VPCs.

Web Application Firewall: AWS WAF is a web application firewall (WAF) that helps protect web applications and APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources.

Route 53: Amazon Route 53 is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to internet applications. Amazon Route 53 is fully compliant with IPv6 as well. With Route 53 Resolver an organization can filter and regulate outbound DNS traffic for its VPC.

3.4.2.3 Compute

EC2: Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers.

ECS: Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that makes it easy to deploy, manage, and scale containerized applications.

EKS: Amazon Elastic Kubernetes Service (Amazon EKS) is a managed container service to run and scale Kubernetes applications in the cloud or on-premises.

3.4.2.4 Storage

EBS: Amazon Elastic Block Store (Amazon EBS) is an easy-to-use, scalable, high-performance block-storage service designed for Amazon EC2.

S3: Amazon Simple Storage Service (Amazon S3) is an object storage service that offers scalability, data availability, security, and performance.

3.4.2.5 Management/Monitoring

Systems Manager: AWS Systems Manager is the operations hub for AWS applications and resources, and it is broken into four core feature groups: Operations Management, Application Management, Change Management, and Node Management.

Security Hub: AWS Security Hub is a cloud security posture management service that performs security best practice checks, aggregates alerts, and enables automated remediation.

CloudWatch: Amazon CloudWatch is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), IT managers, and product owners. CloudWatch provides data and actionable insights to monitor applications, respond to system-wide performance changes, and optimize resource utilization.

CloudTrail: AWS CloudTrail monitors and records account activity across AWS infrastructures, giving organizations control over storage, analysis, and remediation actions.

GuardDuty: Amazon GuardDuty is a threat detection service that continuously monitors AWS accounts and workloads for malicious activity and delivers detailed security findings for visibility and remediation.

Firewall Manager: AWS Firewall Manager is a security management service which allows organizations to centrally configure and manage firewall rules across their accounts and applications in AWS Organizations.

3.4.3 Broadcom Software

Broadcom Software provides business-critical software designed to modernize, optimize, and protect complex hybrid environments. As part of Broadcom Software, the Symantec Enterprise business invests more than 20% of revenue into research and development (R&D), enabling it to innovate across its cybersecurity portfolio and deliver new functionality that delivers both effective zero trust security and an exceptional user experience. With more than 80% of its workforce dedicated to R&D and operations, Broadcom Software's engineering-centered culture supports a comprehensive portfolio of enterprise software, enabling scalability, agility, and security for organizations. For more information, go to https://software.broadcom.com/.

3.4.3.1 Web Security Service with Advanced Malware Analysis

Symantec Web Security Service (WSS), built upon secure web gateway (SWG) technology, is a cloud-delivered network security service that offers protection against advanced threats, provides access control, and safeguards critical business information for secure and compliant use of cloud applications and the web.

3.4.3.2 Web Isolation

Web Isolation enables safe web browsing that protects against malware and phishing threats, even when inadvertently visiting uncategorized and risky websites. Remotely executing web sessions in a secured container stops malware downloads, and read-only browsing defeats phishing attacks. Available as a cloud service or an on-premises virtual appliance, Web Isolation can be standalone or integrated with a proxy or email security solution.

3.4.3.3 CASB with Data Loss Prevention (DLP)

Cloud Access Security Broker (CASB) identifies all cloud apps in use, enforces cloud application management policies, detects and blocks unusual behavior, and integrates with other Symantec solutions, including ProxySG, Data Loss Prevention (DLP), Validation and ID Protection (VIP) Authentication Service, Secure Access Cloud, and Email Security.cloud, to extend network security policies to the cloud. The integration with DLP consistently extends data compliance policies to over 100 Software as a Service (SaaS) cloud apps and automates policy sync with cloud properties. Additional APIs for AWS and Azure also provide visibility and control of the management plane, along with cloud workload assurance for discovering new cloud deployments and monitoring them for critical misconfigurations.

3.4.3.4 Secure Access Cloud

Secure Access Cloud is a cloud-delivered service providing highly secure zero trust network access for enterprise applications deployed in Infrastructure as a Service (IaaS) clouds or on-premises data center environments. This SaaS platform eliminates inbound connections to a network, creates a software-defined perimeter between users and corporate applications, and establishes application-level access. This service avoids the management complexity and security limitations of traditional remote access tools, ensuring that all corporate applications and services are completely cloaked—invisible to attackers targeting applications, firewalls, and virtual private networks (VPNs).

3.4.3.5 Information Centric Analytics (ICA), part of Data Loss Prevention

User and entity behavior analytics is a vital tool to reduce user-based risk. Using it, customers can identify anomalous or suspicious activity to help discover potential insider threats and data exfiltration. It builds behavior profiles of users and entities so high-risk accounts can be investigated. Wider risk context is available when security event telemetry is correlated from many data sources, including DLP, Endpoint Protection, and ProxySG.

3.4.3.6 Symantec Endpoint Security Complete, including Endpoint Detection and Response (EDR) and Mobile Security

Symantec's endpoint security offering delivers protection, detection, and response in a single solution. Symantec Endpoint Security Complete addresses threats along the entire attack chain. It protects all endpoints (workstations, servers, iOS and Android mobile phones and tablets) across all major operating systems, is easy to deploy with a single-agent installation, and provides flexible management options (cloud, on-premises, and hybrid).

3.4.3.7 VIP Authentication Service

VIP is a secure, reliable, and scalable authentication service that provides risk-based and multi-factor authentication (MFA) for all types of users. Risk-based authentication transparently collects data and assesses risk using a variety of attributes such as device identification, geolocation, user behavior, and threat information from the Symantec Global Intelligence Network (GIN). VIP provides MFA using a broad range of authenticators such as push, Short Message Service (SMS) or voice one-time password (OTP), Fast Identity Online (FIDO) Universal 2nd Factor (U2F), and fingerprint biometric. This intelligent, layered security approach prevents inappropriate access and online identity fraud without impacting the user experience. VIP also denies access to compromised devices before they can attempt authentication to the network and tracks advanced and persistent threats. An intuitive credential provisioning portal enables self-service that reduces help desk and administrator costs. An integration with Symantec CloudSOC protects against risky behavior even after application login.

3.4.3.8 VIP Authentication Hub

Authentication Hub is a highly scalable authentication engine that meets zero trust needs by providing phishing-resistant authentication using FIDO2 as well as other multi-factor options, combined with a highly flexible authentication policy model. It includes risk assessment to enable context-sensitive authentication branching. The microservice architecture is built API-first for broad deployment and integration options, and it integrates out of the box with Broadcom's IAM portfolio.

3.4.3.9 Privileged Access Management

Privileged Access Management can minimize the risk of data breaches by continually protecting sensitive administrative credentials, controlling privileged user access, and monitoring and recording privileged user activity.

3.4.3.10 Security Analytics

Security Analytics is an advanced network traffic analysis (NTA) and forensics solution that performs full-packet capture to provide complete network security visibility, anomaly detection, and real-time content inspection for all network traffic to help detect and resolve security incidents more quickly and thoroughly.

3.4.3.11 SiteMinder

While providing the convenience of a single sign-on experience, SiteMinder was built from the ground up using zero trust principles. Every individual resource that is accessed via SiteMinder is only reached once SiteMinder determines if the resource is sufficiently protected, if the user is authenticated, and if the user has authorization to the specific resource. This zero trust approach is applied across all resource access methods (e.g., traditional HTTP, SAML, WS-Federation, OpenID Connect [OIDC], Open Authorization [OAuth]). SiteMinder is deployed in extremely high-performance critical-path business environments. It supports a range of authenticators and in combination with VIP offerings (noted above) provides capabilities to meet the most challenging use cases.

3.4.3.12 Identity Governance and Administration (IGA)

Having a comprehensive ability to manage the lifecycle of user accounts across on-premises and cloud environments is an essential element of a zero trust infrastructure. Symantec IGA delivers comprehensive access governance and management capabilities through an easy-to-use, business-oriented interface. Broad provisioning support for on-premises and cloud apps enables you to automate the granting of new entitlements and removal of unnecessary ones from users throughout the identity life-cycle. Finally, access governance streamlines and simplifies the processes associated with reviewing and approving entitlements, helping ensure a 360 degree view of user entitlements and improving your adherence to zero trust principles.

3.4.4 Cisco

Cisco Systems, or Cisco, delivers collaboration, enterprise, and industrial networking and security solutions. The company's cybersecurity team, Cisco Secure, is one of the largest cloud and network security providers in the world. Cisco's Talos Intelligence Group, the largest commercial threat intelligence team in the world, is comprised of world-class threat researchers, analysts, and engineers, and supported by unrivaled telemetry and sophisticated systems. The group feeds rapid and actionable threat intelligence to Cisco customers, products, and services to help identify new threats quickly and defend against them. Cisco solutions are built to work together and integrate into your environment, using the "network as a sensor" and "network as an enforcer" approach to both make your team more efficient and keep your enterprise secure. Learn more about Cisco at https://www.cisco.com/go/secure.

3.4.4.1 Cisco Secure Access by Duo

Duo is a PE, PA, and PEP for users and their devices. It delivers simple, safe access to all applications — on-premises or in the cloud — for any user, device, or location. It makes it easy to effectively implement and enforce security policies and processes, using strong authentication to reduce the risk of data breaches due to compromised credentials and access from unauthorized devices.

3.4.4.2 Cisco Identity Services Engine (ISE)

Cisco ISE is a network central PDP that includes both the PE and PA to help organizations provide secure access to users, their devices, and the non-user devices in their network environment. It simplifies the delivery of consistent and secure access control to PEPs across wired and wireless multi-vendor networks, as well as remote VPN connections. It controls switches, routers, and other network devices as PEPs, enabling granular control of every connection down to the individual port, delivering a dynamic, granular, and automated approach to policy enforcement that simplifies the delivery of highly secure, micro-segmented network access control. ISE is tightly integrated with and enhances network and security devices, allowing it to transform the network from a simple conduit for data into an intuitive and adaptive security sensor and enforcer that acts to accelerate the time to detection and time to resolution of network threats.

3.4.4.3 Cisco Secure Endpoint (formerly AMP)

Cisco Secure Endpoint addresses the full life cycle of the advanced malware problem before, during, and after an attack. It uses global threat intelligence to strengthen defenses, antivirus to block known malware, and static and dynamic file analysis to detect emerging malware, continuously monitoring file and system activity for emerging threats. When something new is detected, the solution provides a retrospective alert with the full recorded history of the file back to the point of entry, and the rich contextual information needed during a potential breach investigation to both prioritize remediation and create response plans.

As a policy input point, Secure Endpoint delivers deep visibility, context, and control to rapidly detect, contain, and remediate advanced threats if they evade front-line defenses. It can also eliminate malware with a few clicks and provide a cost-effective security solution without affecting operational efficiency.

3.4.4.4 Cisco Firepower Threat Defense (FTD)

Cisco FTD is a threat-focused, next-generation firewall with unified management. It provides advanced threat protection before, during, and after attacks. By delivering comprehensive, unified policy management of firewall functions, application control, threat prevention, and advanced malware protection, from network to endpoint, it increases visibility and security posture while reducing risk.

3.4.4.5 Cisco Secure Network Analytics (formerly Stealthwatch)

Cisco Secure Network Analytics aggregates and analyzes network telemetry — information generated by network devices — to turn the network into a sensor. As a policy input point, it provides enterprise-wide network visibility and applies advanced security analytics to detect and respond to threats in real time. It delivers end-to-end network visibility on-premises, in private clouds, and in public clouds. Secure Network Analytics detects a wide range of network and data center issues ranging from command-and-control (C&C) attacks to ransomware, from distributed denial of service (DDoS) attacks to illicit cryptomining, and from malware to insider threats.

Secure Network Analytics can be deployed on-premises as a hardware appliance or virtual machine (VM), or cloud-delivered as a SaaS solution. It works with the entire Cisco router and switch portfolio as well as a wide variety of other security solutions.

3.4.4.6 Cisco Encrypted Traffic Analytics (ETA)

Cisco ETA helps illuminate the dark corners of encrypted traffic without decryption by using new types of data elements and enhanced NetFlow telemetry independent of protocol details. Cisco ETA can help detect malicious activity in encrypted traffic by applying advanced security analytics. At the same time, the integrity of the encrypted traffic is maintained because there is no need for bulk decryption.

3.4.4.7 Cisco SecureX

Cisco SecureX is an extended detection and response (XDR) cloud-native integrated threat response platform within the Cisco Secure portfolio. Its open, extensible integrations connect to the infrastructure, providing unified visibility and simplicity in one location. It maximizes operational efficiency to secure the network, users and endpoints, cloud edge, and applications. Cisco SecureX radically reduces the dwell time and human-powered tasks involved with detecting, investigating, and remediating threats to counter attacks, or securing access and managing policy to stay compliant. The time savings and better collaboration involved with orchestrating and automating security across SecOps, ITOps, and NetOps teams help advance the security maturity level.

3.4.4.8 Cisco Endpoint Security Analytics (CESA)

Cisco Endpoint Security Analytics (CESA) analyzes endpoint telemetry generated by the Network Visibility Module (NVM), which is built into the Cisco AnyConnect® Secure Mobility Client. CESA feeds Splunk Enterprise software to analyze NVM data provided by endpoints to uncover endpoint-specific security risks and breaches. This data includes information about data loss, unapproved applications and SaaS usage, security evasion, unknown malware, user behavior when not connected to the enterprise, endpoint asset inventory, and destination allowlists and denylists.

3.4.4.9 Cisco AnyConnect Secure Mobility Client

Cisco AnyConnect Secure Mobility Client is a unified endpoint software client compatible with several of today's major enterprise mobility platforms. It helps manage the security risks associated with extended networks. Built on foundational VPN technology, it extends beyond remote-access capabilities to offer user-friendly, network-based security including:

- Simple and context-aware security policy enforcement
- An uninterrupted, intelligent, always-on security connection to remote devices
- Visibility into network and device-user behavior
- Web inspection technology to defend against compromised websites

3.4.4.10 Cisco Network Devices

Cisco network devices do more than move packets on the network; they provide a platform to improve user experience, unify management, automate tasks, analyze activity, and enhance security across the enterprise. In a zero-trust environment, Cisco switches, routers, and other devices provide continuous visibility using the "network as a sensor" to monitor network activity, reporting 100% of NetFlow and other metadata. These devices act as PEPs utilizing a "network as an enforcer" approach to micro-segment network access control to each port and enable dynamic and automated policy enforcement. This policy enforcement simplifies the delivery of highly secure control across environments.

3.4.5 DigiCert

DigiCert is a global provider of digital trust, enabling individuals and businesses to engage online with the confidence that their footprint in the digital world is secure. DigiCert® ONE, the platform for digital trust, provides organizations with centralized visibility and control over a broad range of public and private trust needs, securing websites, enterprise access and communication, software, identity, content, and devices. For more information, visit digicert.com.

3.4.5.1 DigiCert CertCentral TLS Manager

DigiCert CertCentral is used to provision publicly trusted Transport Layer Security (TLS) server authentication certificates. CertCentral relies on DigiCert's publicly trusted root certificates with excellent ubiquity to provide the necessary interoperability with the widest range of third-party products.

3.4.5.2 DigiCert Enterprise PKI Manager

DigiCert Enterprise PKI Manager is a digital certificate management solution for enterprise identity and access public key infrastructure (PKI) use cases. Enterprise PKI Manager simplifies and streamlines certificate lifecycle management for identity and access of users, devices, and applications, supporting a broad array of certificate types with automated workflows, preconfigured templates, multiple enrollment and authentication methods, and a rich ecosystem of integrated technology partners. It is part of the DigiCert family of products delivering digital trust solutions. Enterprise PKI Manager is built on DigiCert ONE's modern, containerized architecture, delivering scalability capable of serving high volumes of certificates, supporting flexible deployment in cloud, on-premises, or hybrid deployment models, and enabling dynamic and rapid intermediate Certificate Authority (ICA) creation to meet the diverse needs of different business groups.

3.4.6 F5

F5 empowers its customers to create, secure, and operate applications that deliver extraordinary digital experiences. Fueled by automation and AI-driven insights, these applications will naturally adapt based on their changing environment—so companies can focus on their core business, boost speed to market, improve operations, and build trust with their customers. By enabling these adaptive applications, F5 with NGINX and F5 Distributed Cloud Services technologies offers a comprehensive suite of solutions for every digital organization.

3.4.6.1 BIG-IP Product Family

The BIG-IP product family provides full proxy security, application intelligence, and scalability for application traffic. As the amount of traffic grows or shrinks, BIG-IP can be adjusted or it can request addition or removal of application servers. It provides rich application traffic programmability to further enhance application security and application traffic steering requirements. In addition, BIG-IP's rich control plane programmability allows for integrations into on-premises orchestration engines, cloud automation/orchestration, and continuous integration/continuous delivery (CI/CD) pipelines, and the ability to deliver application security in a DevSecOps manner. All capabilities can be propagated as common policy throughout the enterprise regardless of whether an organization utilizes F5 hardware or a virtualized on-premises or cloud environment.

BIG-IP modules provide the ability to layer on additional capabilities. The modules being considered for this project are discussed in the subsections below.

3.4.6.1.1 BIG-IP Local Traffic Manager (LTM)

BIG-IP LTM is an enterprise-class load balancer providing granular layer 7 control, Secure Sockets Layer (SSL) offloading, and acceleration capabilities. It allows for massive scaling of traditional and modern apps across the enterprise and provides visibility into TLS-encrypted streams, TLS security enforcement, and Federal Information Processing Standards (FIPS) certified cryptography [8].

3.4.6.1.2 BIG-IP Access Policy Manager (APM)

BIG-IP APM integrates and unifies secure user access to ensure the correct people have the correct access to the correct applications—anytime, anywhere, providing the ability to authenticate users into applications allowing for granular application access control and zero trust capabilities across the application landscape. BIG-IP APM sits in front of applications and APIs to enforce application authentication and access control for each user as part of zero trust.

3.4.6.1.3 BIG-IP Web Application Firewall (WAF)

BIG-IP WAF provides the flexibility to deploy WAF services closer to the apps so they're protected wherever they reside. It has the ability to virtually patch applications for security vulnerabilities such as the latest Common Vulnerabilities and Exposures (CVE) entry without application code changes. It also reduces unwanted application traffic, allowing the application to be more responsive to its intended users while providing complete visibility into the application traffic. WAF provides API security, protecting against web application security concerns. WAF provides secure communication and vetting of traffic to APIs and applications.

3.4.6.2 NGINX Product Family

NGINX is a cloud-native, easy-to-use reverse proxy, load balancer, and API gateway. It integrates advanced monitoring, strengthens security controls, and orchestrates Kubernetes containers.

3.4.6.2.1 NGINX Ingress Controller

NGINX Ingress Controller combines software load balancing with simplified configuration based on standard Kubernetes Ingress resources or custom NGINX Ingress resources to ensure that applications in a Kubernetes cluster are delivered reliably, securely, and at high velocity. It provides security to Kubernetes-based microservices and APIs using API gateway and WAF capabilities. The Ingress Controller protects application and API containers in the Kubernetes environment by enforcing security on all traffic entering the Kubernetes node.

3.4.6.2.2 NGINX Plus

NGINX Plus is an all-in-one load balancer, web server, content cache, WAF, and API gateway. NGINX Plus is built on NGINX Open Source. It is intended to reduce complexity and simplify management by consolidating several capabilities, including reverse proxy and TLS termination, into a single elastic ingress/egress tier. It acts as a webserver to server applications that are secured by the system's zero trust capabilities.

3.4.6.2.3 NGINX Service Mesh

NGINX Service Mesh scales from open-source projects to a fully supported, secure, and scalable enterprise-grade solution. It provides a turnkey service-to-service solution featuring a unified data plane for ingress and egress Kubernetes management in a single configuration. NGINX Service Mesh provides for mutual TLS authentication (mTLS) enforcement, rate limiting, quality of service (QOS), and an API gateway to enforce security at each pod, securing pods from both north/south (N/S) and east/west (E/W) traffic and allowing for zero trust enforcement for all pod traffic.

3.4.7 Forescout

Forescout delivers automated cybersecurity across the digital terrain. It empowers its customers to achieve continuous alignment of their security frameworks with their digital realities, across all asset types – IT, IoT, OT, and Internet of Medical Things (IoMT). Forescout enables organizations to manage cyber risk through automation and data-powered insights.

The Forescout Continuum Platform provides complete asset visibility of connected devices, continuous compliance, network segmentation, network access control, and a strong foundation for zero trust. Forescout customers gain data-powered intelligence to accurately detect risks and quickly remediate cyberthreats without disruption of critical business assets. https://www.forescout.com/company/

3.4.7.1 Forescout eyeSight

Forescout eyeSight delivers comprehensive device visibility across an organization's entire digital terrain – without disrupting critical business processes. It discovers every IP-connected device, auto-classifies it, and assesses its compliance posture and risk the instant the device connects to the network. https://www.forescout.com/products/eyesight/

3.4.7.2 Forescout eyeControl

Forescout eyeControl provides flexible and frictionless network access control for heterogeneous enterprise networks. It enforces and automates zero trust security policies for least-privilege access on all managed and unmanaged assets across an organization's digital terrain. Policy-based controls can continuously enforce asset compliance, proactively reduce attack surfaces, and rapidly respond to incidents. https://www.forescout.com/products/eyecontrol/

3.4.7.3 Forescout eyeSegment

Forescout eyeSegment accelerates zero trust segmentation. It simplifies the design, planning, and deployment of non-disruptive, dynamic segmentation across an organization's digital terrain to reduce attack surface and regulatory risk. https://www.forescout.com/products/eyesegment/

3.4.7.4 Forescout eyeExtend

Forescout eyeExtend automates security workflows across disparate products. It shares device context between the Forescout platform and other IT and security products, automates policy enforcement across disparate tools, and accelerates system-wide response to mitigate risks. https://www.forescout.com/products/eyeextend/

3.4.8 Google Cloud

Google Cloud brings the best of Google's innovative products and services to enable enterprises of all sizes to create new user experiences, transform their operations, and operate more efficiently. Google's mission is to accelerate every organization's ability to digitally transform its business with the best infrastructure, platform, industry solutions, and expertise. Google Cloud helps customers protect their data using the same infrastructure and security services Google uses for its own operations, defending against the toughest threats. Google pioneered the zero trust model at the core of its services and operations, and it enables its customers to do the same with its broad portfolio of solutions. Learn more about Google Cloud at https://cloud.google.com/.

3.4.8.1 BeyondCorp Enterprise (BCE)

BeyondCorp Enterprise (BCE) is a zero trust solution, built on the Google platform and global network, which provides customers with simple and secure access to applications and cloud resources and offers integrated threat and data protection. It leverages the Chrome Browser and the Google Cloud platform (GCP) to protect and proxy traffic from an organization's network. It allows customers to enforce context-aware policies (using factors such as identity, device posturing, and other signal information) to authorize access to SaaS applications and resources hosted on Google Cloud, third-party clouds, or on-premises. This solution is built from Google's own approach of shifting access controls from the network perimeter to individual users and devices, allowing for secure access without the need for a VPN.

BCE key capabilities include:

- **Zero trust access**
    - **Context-aware access proxy (identity-aware proxy):** Globally deployed proxy built on the GCP that leverages identity, device, and contextual information to apply continuous authorization access decisions to applications and VMs in real-time in the GCP, other clouds, or on-premises data centers.
    - **Browser-based application access:** Agentless zero trust access, using Chrome or other browsers, to browser-based apps hosted on the GCP, other clouds (e.g., AWS, Azure), or on-premises data centers.
    - **Legacy client application access (client connector):** Extension that enables zero trust access to non-HTTP, thick-client apps hosted in the GCP, other clouds, or on-premises data centers.

- **Protections**
  - **Data protection:** Built-in Chrome browser capabilities to detect and prevent sensitive data loss, stop pasting of protected content in and out of the browser, prevent accidental and intentional exfiltration of corporate data, and enforce data protection policies across applications.
  - **Threat protection:** Built-in Chrome browser capabilities to filter and block harmful or unauthorized URLs in real-time, identify phishing sites and malicious content in real-time, stop suspicious files and malware transfers, and protect user credentials and passwords.
- **Integrations**
  - **BeyondCorp Alliance ecosystem integrations:** A collection of integrations from BeyondCorp Alliance member partners that enable organizations to share signal information from EDR, MDM, enterprise mobility management (EMM), and other device or ecosystem endpoints to use in access policy decisions. (Members include Broadcom Software, Check Point, Citrix, CrowdStrike, Jamf, Lookout, Netskope, Palo Alto Networks, Tanium, and VMware.)
- **Network connectivity**
  - **On-premises connector:** Private connectivity from Google Cloud to applications outside of Google Cloud (i.e., hosted by other clouds or on-premises data centers.)
  - **VPN interconnect:** Private connectivity via an Interconnect from Google Cloud to applications outside of Google Cloud (i.e., hosted by other clouds or on-premises data centers.)
  - **App connector:** Secure internet-based connectivity from Google Cloud to applications outside of Google Cloud (i.e., hosted by other clouds or on-premises data centers.)
- **Platform**
  - **Google Platform:** Google's public cloud computing services including data management, application development, storage, hybrid & multi-cloud, security, and AI & ML that run on Google infrastructure.
  - **Google Network:** Google's global backbone with 146 edge locations in over 200 countries and territories provides low-latency connections, integrated DDoS protection, elastic scaling, and private transit.

3.4.9 IBM

International Business Machines Corporation (IBM) is an American multinational technology corporation headquartered in Armonk, New York, with operations in over 171 countries. IBM produces and sells computer hardware, middleware, and software, and provides hosting and consulting services in areas ranging from mainframe computers to nanotechnology. IBM is also a major research organization, holding the record for most annual U.S. patents generated by a business (as of 2020) for 28 consecutive years. IBM has a large and diverse portfolio of products and services that range in the categories of cloud computing, AI, commerce, data and analytics, IoT, IT infrastructure, mobile, digital workplace, and cybersecurity.

3.4.9.1 IBM Security Trusteer

IBM Security® Trusteer® solutions help detect fraud, authenticate users, and establish identity trust across a digital user journey. Trusteer uses cloud-based intelligence, AI, and machine learning (ML) to holistically identify new and existing users while improving the overall user experience by reducing the friction created with traditional forms of MFA. Within a ZTA, Trusteer acts as a risk engine that improves the efficacy of policy decisions enforced by various identity and access management solutions.

3.4.9.2 IBM Security QRadar XDR

IBM Security QRadar® XDR suite provides a single unified workflow across an organization's security tools. Built on a unified cross-domain security platform, IBM Cloud Pak® for Security, the open architecture of QRadar XDR suite enables organizations to integrate their EDR, security information and event management (SIEM), network detection and response (NDR), security orchestration, automation, and response (SOAR), and threat intelligence solutions in support of a ZTA.

IBM Security QRadar SIEM helps security teams detect, prioritize, and respond to threats across the enterprise. As an integral part of an organization's XDR and zero trust strategies, it automatically aggregates and analyzes log and flow data from thousands of devices, endpoints, and apps across the network, providing single, prioritized alerts to speed incident analysis and remediation. QRadar SIEM is available for on-premises and cloud environments.

IBM Security QRadar SOAR is designed to help security teams respond to cyberthreats with confidence, automate with intelligence, and collaborate with consistency. It guides a team in resolving incidents by codifying established incident response processes into dynamic playbooks. The open and agnostic platform helps accelerate and orchestrate response by automating actions with intelligence and integrating with other security tools.

IBM Security QRadar XDR Connect is a cloud-native, open XDR solution that saves time by connecting tools, workflows, insights, and people. The solution adapts to a team's skills and needs, whether the user is an analyst looking for streamlined visibility and automated investigations or an experienced threat hunter looking for advanced threat detection. XDR Connect empowers organizations with tools that strengthen their zero trust model and enable them to be more productive.

3.4.9.3 IBM Security Verify

Modernized, modular IBM Security Verify provides deep, AI-powered context for both consumer and workforce identity and access management. It protects users and apps, inside and outside the enterprise, with a low-friction, cloud-native, SaaS approach that leverages the cloud. Verify delivers critical features for supporting a zero trust strategy based on least privilege and continuous verification, including single sign-on (SSO), multi-factor and passwordless authentication, adaptive access, identity lifecycle management, and identity analytics.

3.4.9.4 IBM Security MaaS360

IBM Security MaaS360® with Watson protects devices, apps, content, and data, which allows organizations to rapidly scale their hybrid workforce and BYOD initiatives. IBM Security MaaS360 can help build a zero trust strategy with modern device management. And with Watson, organizations can take advantage of contextual analytics via AI for actionable insights.

3.4.9.5 IBM Security Guardium

IBM Security Guardium® Insights is a data security hub for the modern data source environment. It builds and automates compliance policy enforcement, and streams and centralizes data activity across a multi-cloud ecosystem. It can apply advanced analytics to uncover data risk insights. Guardium Insights can complement and enhance existing Guardium Data Protection deployments or be installed on its own to help solve compliance and cloud data activity monitoring challenges. Built on a unified cross-domain security platform, IBM Cloud Pak for Security, Guardium Insights can deploy and scale in any data environment — as well as integrate and share insights with major security tools such as IBM Security QRadar XDR, Splunk, ServiceNow, and more, in support of a ZTA.

3.4.9.6 IBM Cloud Pak for Security

IBM Cloud Pak for Security is a unified cross-domain security platform that integrates existing security tools to generate insights into threats across hybrid, multi-cloud environments. It provides organizations with the ability to track, manage, and resolve cybersecurity incidents and create response plans that are based on industry standards and best practices.

3.4.10 Ivanti

Ivanti finds, heals, manages, and protects devices regardless of location – automatically. It is an enterprise software company specializing in endpoint management, network security, risk-based vulnerability management, and service and asset management. The Ivanti solution is able to discover, manage, secure, and service all endpoints across the enterprise including corporate/government-owned and BYOD. Ivanti is actively involved with helping to better prepare government and enterprises with cybersecurity and zero trust best practices. Learn more about Ivanti here: https://www.ivanti.com/. The Ivanti solution enables an enterprise to centrally manage/monitor endpoints and trigger adaptive policies to remediate threats, quarantine devices, and maintain compliance.

3.4.10.1 Ivanti Neurons for Unified Endpoint Management (UEM)

Ivanti Neurons for UEM helps enterprises create a secure workspace on any device with apps, configurations, and policies for the user based on their role. Users get easy and secure access to the resources they need for their productivity. For more information, see https://www.ivanti.com/products/ivanti-neurons-for-mdm.

The Ivanti Neurons for UEM platform provides the fundamental visibility and IT controls needed to secure, manage, and monitor any corporate or employee-owned mobile device or desktop that accesses business-critical data. The Neurons for UEM platform allows organizations to secure a vast range of employee and BYOD devices being used within the organization while managing the entire life cycle of the device, including:

- Policy configuration management and enforcement
- Application distribution and management
- Script management and distribution for desktop devices
- Automated device actions
- Continuous access control and MFA
- Threat detection and remediation against device, network application, and phishing attacks

3.4.10.2 Ivanti Sentry

Ivanti Sentry is an in-line intelligent gateway that helps secure access to on-premises resources and provides authentication and authorization to enterprise data. For more information, see https://www.ivanti.com/products/secure-connectivity/sentry.

3.4.10.3 Ivanti Access ZSO

Ivanti Access Zero Sign-On (ZSO) helps identify the user, device, app, network type, and presence of threats. The adaptive access control check is the basis of the zero-trust model. Access provides zero sign-on and security on the cloud and federated enterprise data. The solution is federated with the Okta Identity Cloud to provide continuous authentication and authorization. For more information, see https://www.ivanti.com/products/zero-sign-on.

3.4.10.4 Ivanti Mobile Threat Defense

The combination of cloud and mobile threat defense (MTD) protects data on-device and on-the-network with state-of-the-art encryption and threat monitoring to detect and remediate device, network, app-level, and phishing attacks. For more information, see https://www.ivanti.com/products/mobile-threat-defense.

3.4.11 Lookout

Lookout is a cybersecurity company focused on securing users, devices, and data as users operate in the cloud. The Lookout platform helps organizations consolidate IT security, get complete visibility across all cloud services, and protect sensitive data wherever it goes.

3.4.11.1 Lookout Mobile Endpoint Security (MES)

Lookout MES is a SaaS-based MTD solution that protects devices from threats and risks via the Lookout for Work mobile application. Lookout protects Android and Apple mobile devices from malicious or risky apps, device threats, network threats, and phishing attacks. Lookout attests to the security posture of the mobile device, which is provided to the policy engine to determine access to a resource. The mobile asset is continuously monitored by Lookout for any change to its security posture. Lookout protection can be deployed to managed or unmanaged devices and works on trusted or untrusted networks. Lookout has integrations with productivity and collaboration solutions, as well as unified endpoint management solutions.

3.4.12 Mandiant

Mandiant scales its intelligence and expertise through the Mandiant Advantage SaaS platform to deliver current intelligence, automation of alert investigation, and prioritization and validation of security control products from a variety of vendors. (http://www.mandiant.com/)

3.4.12.1 Mandiant Security Validation (MSV)

Mandiant Security Validation (MSV), continuously informed by Mandiant frontline intelligence on the latest attacker tactics, techniques, and procedures (TTPs), automates a testing program that gives real data on how security controls are performing. This solution provides visibility and evidence on the status of security controls' effectiveness against adversary threats targeting organizations and data to optimize environment against relevant threats. MSV can provide many benefits to an organization (for example, identify limitations in current cybersecurity stack, evaluate proposed cybersecurity tools for an organization, determine overlapping controls, automate assessment actions, and train cybersecurity operators). To support these use cases, MSV emulates 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 across the most critical areas of the enterprise.

Using the natural design of the Security Validation platform, Mandiant is able to support the project in testing and documenting the outcome of one of the key tenets of ZTA, "The enterprise monitors and measures the integrity and security posture of all owned and associated resources." To do this, the software produces quantifiable evidence that shows how people, processes, and technologies perform when specific malicious behaviors are encountered, such as attacks by a specific threat actor or attack vector.

The core Validation components of the MSV platform are:

- The Director - This is the main component of the platform and provides the following functionality:
    - Acts as the Integration point and content manager for the SIEM and other components of the security stack
    - Hosts the Content Library (Actions, Sequences, Evaluations, and Files) used for testing security controls
    - Manages the Actor assignment during testing
    - Aggregates testing results and facilitates report creation
    - Maintains connections with the Mandiant Updater and Content Services, allowing updates to be received automatically for the platform and its content
- Actors (also referred to as flex, Endpoint, and Network Actors) - The components that safely perform tests in production environments. Specifically, use these to verify the configuration and test the effectiveness of network security controls; Windows, Mac, and Linux endpoint controls; and email controls.
- Cloud controls
- Policy compliance

The Director is the component that receives the information from the systems in the environment based on an integration with a SIEM and/or directly with the security appliance itself. Tests are run between Actors and not directly on systems in the environment.

3.4.13 Microsoft

Microsoft Security brings together the capabilities of security, compliance, identity, and management to natively integrate individual layers of protection across clouds, platforms, endpoints, and devices. Microsoft Security helps reduce the risk of data breaches and compliance violations and improve productivity by providing the necessary coverage to enable zero trust. Microsoft's security products give IT leaders the tools to confidently help their organization digitally transform with Microsoft's protection across their entire environment.

3.4.13.1 Azure

Microsoft Azure is Microsoft's public cloud computing platform. It provides a range of cloud services, including compute, analytics, storage, and networking.

3.4.13.2 Azure Active Directory (Azure AD)

Azure AD is an IAM/identity as a service (IDaaS) product from Microsoft that performs ICAM management, authentication (both SSO and MFA), authorization, federation, and governance, and also functions as a PE, PA, and PEP.

3.4.13.3 Microsoft Intune – Device Management

In Intune, devices are managed using an approach that's suitable for the organization. For organization-owned devices, an organization may want full control over the devices, including settings, features, and security. In this approach, devices and users of these devices "enroll" in Intune. Once enrolled, they receive the organization's rules and settings through policies configured in Intune. For example, organizations can set password and PIN requirements, create a VPN connection, set up threat protection, and more.

3.4.13.4 Microsoft Intune – Application Management

Microsoft Intune provides mobile application management (MAM), which is designed to protect organization data at the application level, including custom apps and store apps. App management can be used on organization-owned devices and personal devices. When apps are managed in Intune, administrators can:

- add and assign mobile apps to user groups and devices, including users in specific groups, devices in specific groups, and more;
- configure apps to start or run with specific settings enabled and update existing apps already on the device;
- see reports on which apps are used and track their usage; and
- do a selective wipe by removing only organization data from apps.

3.4.13.5 Microsoft Defender for Endpoint

Microsoft Defender for Endpoint is an enterprise endpoint security platform designed to help enterprise networks prevent, detect, investigate, and respond to advanced threats.

3.4.13.6 Microsoft Sentinel

Microsoft Sentinel is a scalable, cloud-native solution for SIEM. It was previously known as Azure Sentinel.

3.4.13.7 Microsoft Defender for Identity

Microsoft Defender for Identity (formerly Azure Advanced Threat Protection, also known as Azure ATP) is a cloud-based security solution that leverages an organization's on-premises AD signals to identify, detect, and investigate advanced threats, compromised identities, and malicious insider actions directed at the organization. Defender for Identity enables SecOps analysts and security professionals struggling to detect advanced attacks in hybrid environments to:

- monitor users, entity behavior, and activities with learning-based analytics;
- protect user identities and credentials stored in AD;
- identify and investigate suspicious user activities and advanced attacks throughout the kill chain; and
- provide clear incident information on a simple timeline for fast triage.

3.4.13.8 Azure AD Identity Protection

Identity Protection, which is part of Azure AD, is a tool that allows organizations to accomplish three key tasks:

- automate the detection and remediation of identity-based risks;
- investigate risks using data in the portal; and
- export risk detection data to the SIEM.

Identity Protection uses the learnings Microsoft has acquired from its position in organizations with Azure AD, in the consumer space with Microsoft Accounts, and in gaming with Xbox to protect users. Microsoft analyses 6.5 trillion signals per day to identify and protect customers from threats.

The signals generated by and fed to Identity Protection can be further fed into tools like Conditional Access to make access decisions, or fed back to a SIEM tool for further investigation based on an organization's enforced policies.

3.4.13.9 Microsoft Defender for Office 365 (for email)

Microsoft Defender for Office 365 (for email) prevents broad, volume-based, known attacks. It protects email and collaboration from zero-day malware, phishing, and business email compromise. It also adds post-breach investigation, hunting, and response, as well as automation and simulation (for training).

3.4.13.10 Azure App Proxy & Intune VPN Tunnel

Azure Active Directory Application Proxy provides secure remote access and cloud-scale security to an organization's private applications.

Microsoft Tunnel is a VPN gateway solution for Microsoft Intune that runs in a container on Linux and allows access to on-premises resources from iOS/iPadOS and Android Enterprise devices using modern authentication and Conditional Access.

3.4.13.11 Secure Admin Workstation (SAW)

Secure Admin Workstations are limited-use client computers—built on Windows 10—that help protect high-risk environments from security risks such as malware, phishing, and pass-the-hash attacks. They provide secure access to restricted environments.

3.4.13.12 Windows 365 for Enterprise and Azure Virtual Desktop

Windows 365 for Enterprise is a cloud-based service that automatically creates a new type of Windows virtual machine (Cloud PCs) for your end users that provides the productivity, security, and collaboration benefits of Microsoft 365.

Azure Virtual Desktop is a desktop and app virtualization service that runs on the cloud.

For this project, Microsoft 365 for Enterprise and Azure Virtual Desktop can both be used to show how to secure virtual desktop infrastructure (VDI).

3.4.13.13 Microsoft Defender for Cloud

Defender for Cloud is a tool for security posture management and threat protection. It strengthens the security posture of an organization's cloud resources, and with its integrated Microsoft Defender plans, Defender for Cloud protects workloads running in Azure, hybrid, and other cloud platforms. Because it's natively integrated, deployment of Defender for Cloud is easy, providing an organization with simple auto provisioning to secure its resources by default.

3.4.13.14 Microsoft Purview

Microsoft Purview is a unified data governance service that helps organizations manage and govern their on-premises, multi-cloud, and SaaS data. It creates a holistic, up-to-date map of an organization's data landscape with automated data discovery, sensitive data classification, and end-to-end data lineage, enabling data curators to manage and secure the organization's data estate. It also empowers data consumers to find valuable, trustworthy data.

3.4.13.15 Microsoft Defender for Cloud Apps

Microsoft Defender for Cloud Apps is a CASB that supports various deployment modes, including log collection, API connectors, and reverse proxy. It provides rich visibility, control over data travel, and sophisticated analytics to identify and combat cyberthreats across all of an organization's Microsoft and third-party cloud services. Microsoft Defender for Cloud Apps natively integrates with Microsoft solutions and is designed with security professionals in mind. It provides simple deployment, centralized management, and innovative automation capabilities.

3.4.13.16 Microsoft Entra Permissions Management

Microsoft Entra Permissions Management (formerly known as CloudKnox) is a cloud infrastructure entitlement management (CIEM) solution that provides comprehensive visibility into permissions assigned to all identities, for example, overprivileged workload and user identities, actions, and resources across multi-cloud infrastructures in Microsoft Azure, AWS, and GCP.

3.4.14 Okta

Okta is an independent identity provider helping organizations protect the identities of their extended workforces, partners, and customers. With more than 7,000 pre-built integrations to applications and infrastructure providers, Okta provides simple and secure access to people and organizations everywhere, giving them the confidence to reach their full potential. Learn more about Okta here: Okta.com.

3.4.14.1 Okta Identity Cloud

The Okta Identity Cloud is an independent and neutral platform that securely connects the correct people to the correct technologies at the appropriate time. The Okta Identity Cloud includes identity and access management products, integrations, and platform services for extended Workforce Identity and Customer Identity use cases.

The Okta Identity Cloud provides secure user storage, authentication capabilities (primary and MFA) to applications and resources (infrastructure, APIs) regardless of location (on-premises, cloud, or hybrid), as well as automation and orchestration capabilities for identity use cases, such as for automating user on- and off-boarding or for identifying and acting on inactive user accounts. Products used in this project include the following.

3.4.14.1.1 Universal Directory

Okta Universal Directory is a cloud metadirectory that is used as a single source of truth to manage all users (employees, contractors, customers), groups, and devices. These users can be sourced directly within Okta or from any number of sources including AD, Lightweight Directory Access Protocol (LDAP), HR systems, and other SaaS applications.

3.4.14.1.2 Single Sign-On (SSO)

Okta SSO delivers seamless and secure access to all cloud and on-premises apps for end users, centralizing and protecting all user access via Okta's cloud portal.

Okta FastPass, available as a part of Okta SSO, enables passwordless authentication. Organizations can use Okta FastPass to minimize end user friction when accessing corporate resources, while still enforcing Okta's adaptive policy checks.

3.4.14.1.3 Adaptive Multi-Factor Authentication (MFA)

Okta Adaptive MFA uses intelligent policies to enable contextual access management, allowing administrators to set policies based on risk signals native to Okta as well as from third parties, such as device posture from EDR vendors. Okta Adaptive MFA also enables administrators to choose the factor(s) that work best for their organization, balancing security and ease of use with options such as secure authenticator apps, WebAuthn, and biometrics, which many organizations also choose as passwordless options.

3.4.14.1.4 Okta Access Gateway

Okta Access Gateway is an application access proxy that delivers access management (SSO, MFA, and URL authorization) to on-premises apps using legacy on-premises protocols – header-based authentication and Kerberos – without requiring changes in source code. In combination with Okta SSO, it allows users to access cloud and on-premises apps remotely from a single place and delivers the same easy and secure login experience for SaaS and on-premises apps.

3.4.14.1.5 Okta Verify

Okta Verify is a lightweight application that is used both as an authenticator option (e.g., OTP or push, available on macOS, Windows, iOS, and Android) with Okta MFA as well as to register a device to Okta. Registering a device to Okta enables organizations to deliver secure, seamless, passwordless authentication to apps, strong device-level security, and more. Okta Verify is FIPS 140-2 validated. [9]

3.4.14.1.6 Okta Integration Network

The Okta Integration Network serves as a conduit to connect thousands of applications and resources (infrastructure, APIs) to Okta for access management (SSO/MFA) and provisioning (automating on- and off-boarding of user accounts). This integration network makes it easy for administrators to manage and control access for all users behind a single pane of glass, and easy for users to get to the tools they need with a unified access experience.

In addition, the Okta Integration Network also serves as a rich ecosystem to support risk signal sharing for zero trust security. Okta's deep integration with partners in the zero trust ecosystem allows the Okta Identity Cloud to take in risk signals for the purpose of making smarter contextual decisions regarding access. For example, integrations with EMM or EDR solutions allow the Okta IDaaS platform to know the managed state of a device or device risk posture and make decisions regarding access accordingly. Okta can also pass risk signals to third parties such as inline network solutions, which can in turn leverage Okta's risk assessment to limit actions within SaaS apps when risk is high (e.g., read-only). Okta's risk-based approach to access allows for fine-grained control of user friction and provides organizations with a truly zero trust PDP to make just-in-time, contextual-based authentication decisions to any resource, from anywhere.

3.4.15 Palo Alto Networks

Palo Alto Networks is shaping the cloud-centric future with technology designed to transform the way people and organizations operate by using the latest breakthroughs in AI, analytics, automation, and orchestration. By delivering an integrated platform and empowering a growing ecosystem of partners, Palo Alto Networks security technologies enable organizations to apply consistent security controls across clouds, networks, endpoints, and mobile devices.

Their core capabilities include the ability to inspect all traffic, including all applications, threats, and content, and tie that traffic to the user, regardless of location or device type. The user, application, and content—the elements that run your business—become integral components of your enterprise's zero trust security policy.

Towards that end, their Next Generation Firewall (including all hardware-based, VM, and containerized form factors) and Prisma Access have consistent core capabilities fundamental for zero trust policy enforcement—including User-ID, App-ID, and Device-ID.

- *User-ID*™ technology enables organizations to identify users in all locations, no matter their device type or OS. Visibility into application activity—based on users and groups, instead of IP addresses—safely enables applications by aligning usage with business requirements.
- *App-ID*™ technology enables organizations to accurately identify applications in all traffic passing through the network, including applications disguised as authorized traffic, using dynamic ports, or trying to hide under the veil of encryption. App-ID allows organizations to understand and control applications and their functions, such as video streaming versus chat, upload versus download, and screen-sharing versus remote device control.
- *Device-ID*™ technology enables organizations to enforce policy rules based on a device, regardless of changes to its IP address or location. By providing traceability for devices and associating network events with specific devices, Device-ID allows organizations to gain context for how events relate to devices and write policies that are associated with devices, instead of users, locations, or IP addresses, which can change over time.

All NGFW form factors and Prisma Access also include the following cloud-delivered security service (CDSS) capabilities: Advanced Threat Prevention (ATP), Wildfire (WF) malware analysis, Advanced URL Filtering (AURL), and DNS Security (DNS). These capabilities are supported by the GlobalProtect (GP) remote access solution and can all be centrally managed by Panorama.

3.4.15.1 Next-Generation Firewall (NGFW)

The Palo Alto Networks Next-Generation Firewall (NGFW) is an ML-powered network security platform available in physical, virtual, containerized, and cloud-delivered form factors—all managed centrally via Panorama. The Palo Alto Networks NGFWs inspect all traffic, including all applications, threats, and content, and tie that traffic to the user, regardless of location or device type. Built on a single-pass architecture, the Palo Alto Networks NGFW performs full-stack, single-pass inspection of all traffic across all ports, providing complete context around the application, associated content, and user identity to form the basis for zero trust security policy decisions.

Additional NGFWs, including cloud-delivered, software-based VMs (VM-Series), and container-based (CN-Series), are anticipated to be used as part of the micro-segmentation deployment model phase of this project, deployed as PEPs deeper within each enterprise environment. Regardless of form factor, any NGFW or Prisma Access instance can serve as a PEP, enabled by the core (User-ID, Application-ID, Device-ID) technologies described above—helping organizations achieve common zero trust use cases such as data center segmentation, user or application-based segmentation, or cloud transformation.

3.4.15.2 Prisma Access

Prisma Access allows organizations to securely enable remote workforces and branch locations, and will be more extensively demonstrated during the SDP deployment model phase of the project. The cloud-native architecture of Prisma Access is designed to ensure on-demand and elastic scaling of comprehensive networking and security services across a global, high-performance network. Together with Prisma SD-WAN (software-defined wide area network), Prisma Access provides the foundational layer for a complete secure access service edge (SASE) solution that delivers networking and security with a common service delivery model.

Prisma Access combines least-privileged access with deep and ongoing security inspection as well as enterprise DLP to protect all users, devices, apps, and data. Prisma Access fully inspects all application traffic bidirectionally—including TLS-encrypted traffic—on all ports, whether communicating with the internet, the cloud, the data center, or between branches. Additionally, Prisma Access provides more security coverage consolidating multiple point products into a single converged platform that includes Firewall as a Service (FWaaS), Zero Trust Network Access (ZTNA), next-generation CASB, cloud SWG, VPN, and more—all managed through a single console.

Prisma Access connects users and applications with fine-grained access controls, providing behavior-based continuous trust verification after users connect to dramatically reduce the attack surface.

3.4.15.3 Cortex XDR

Cortex XDR is an XDR tool that natively integrates network, endpoint, and cloud data to stop sophisticated attacks. Leveraging behavioral analytics, it identifies unknown and highly evasive threats targeting your environment. ML and AI models uncover threats from multiple sources, including managed and unmanaged devices. Cortex XDR speeds alert triage and incident response by providing a comprehensive picture of each threat and revealing the root cause. By stitching different types of data together and simplifying investigations, Cortex XDR reduces the time and experience required at every stage of security operations, from triage to threat hunting. Native integration with enforcement points lets you respond to threats quickly and apply the knowledge gained from investigations to mitigate future attacks.

Cortex XDR features Identity Analytics, which detects malicious user activities by applying ML and behavioral analytics to users, machines, and entities. Using an analytics engine to examine logs and data, Identity Analytics can understand normal behaviors across your environment and create a baseline so that it can raise alerts when abnormal activity occurs. With this function, suspicious user activity such as stolen or misused credentials, lateral movement, credential harvesting, exfiltration, and brute-force attacks can be detected. This ML-derived insight offers critical identity context specific to each bespoke environment Cortex XDR is deployed into, allowing for higher fidelity alerts to aid organizations in fine tuning access granted to critical assets—an imperative for ZTA.

3.4.16 PC Matic

PC Matic is an endpoint protection solution for enterprises of all sizes, utilizing PC Matic's proactive application allowlisting technology. Through a series of global and local allowlists, PC Matic's software asset management restricts unauthorized programs and processes from accessing resources such as data or services on a network. Unlike traditional application allowlisting products that solely rely on self-made local allowlists, PC Matic operates off both the user's local list and a real-time automated global allowlist consisting of verified files, processes, digital certificates, and scripts. PC Matic eliminates governance issues by granting users the ability to create application, digital certificate, directory, or scripting policies within their local lists. This capability takes immediate effect and can be deployed to individual endpoints, departments, groups, whole organizations, and all agencies and enterprises managed across the account.

3.4.16.1 PC Matic Pro

PC Matic Pro's on-premises endpoint protection provides default-deny protection at the device. PC Matic Pro monitors for any process that attempts to execute and automatically denies access to any unauthorized or known malicious entities. When the unauthorized files and/or processes are denied access, all metadata pertaining to the block is then communicated to the architecture's SIEM for prioritizing and further investigation. This integration provides users with increased visibility over their managed devices and networks. If a block is verified and warranted, the SIEM of choice can utilize the policy engine from either PC Matic or a third-party vendor to create and enforce the exception, granting immediate access to the desired deployment. PC Matic's real-time policy offerings eliminate governance issues, take immediate effect without delay or issue, and provide users with streamlined management across their managed architectures. PC Matic's allow-by-exception approach to prevention enhances the zero-trust model and minimizes the network's attack surface by ensuring only authorized processes are granted privileges to execute and proceed further.

3.4.17 Ping Identity

Ping Identity delivers intelligent identity solutions for the enterprise. Ping enables companies to achieve zero trust identity-defined security and more personalized, streamlined user experiences. The PingOne

Cloud Platform provides customers, workforces, and partners with access to cloud, mobile, SaaS, and on-premises applications across the hybrid enterprise. Over half of the Fortune 100 choose Ping for their identity expertise, open standards, and partnerships with companies including Microsoft and Amazon. Ping Identity provides flexible identity solutions that accelerate digital business initiatives and secure the enterprise through multi-factor authentication, single sign-on, access management, intelligent API security, and directory and data governance capabilities. For more information, please visit https://www.pingidentity.com/.

3.4.17.1 PingFederate

PingFederate is an enterprise federation server that enables user authentication and single sign-on. It is a global authentication authority that allows customers, employees, and partners to access all the applications they need from any device securely. PingFederate easily integrates with applications across the enterprise, third-party authentication sources, diverse user directories, and existing IAM systems, all while supporting current and past versions of identity standards. It will connect everyone to everything.

PingFederate can be deployed within Ping Identity's SaaS offerings, in a customer cloud, as a traditional application, and within air-gapped or network segmented environments.

The deployment architecture of PingFederate eliminates the need to maintain redundant copies of configurations and trust relationships. Supported federation standards include OAuth, OpenID, OpenID Connect, SAML, WS-Federation, WS-Trust, and System for Cross-Domain Identity Management (SCIM).

3.4.17.2 PingOne DaVinci

PingOne DaVinci is a SaaS platform that enables a flexible and adaptive integration framework, allowing you to easily create identity journeys via a drag-and-drop interface. Through DaVinci, administrators can quickly design automated workflows for different identity use cases including authentication, identity proofing, and fraud detection. DaVinci is an open interface with integrations and connections across multiple applications and identity ecosystems.

3.4.17.3 PingOne SSO

PingOne SSO is a SaaS federation platform. Using single sign-on (SSO), users can sign on to all their applications and services with one set of credentials. It gives employees, partners, and customers secure, one-click access from anywhere, on any device, and it reduces the number of separate accounts and passwords they need to manage.

SSO is made possible by a centralized authentication service that all apps (even third-party) can use to confirm a user's identity. Identity standards like SAML, OAuth, and OpenID Connect allow for encrypted tokens to be transmitted securely between the server and the apps to indicate that a user has already been authenticated and has permission to access the additional apps.

3.4.17.4 PingOne Risk

PingOne Risk is a SaaS platform that enables administrators to configure intelligence-based authentication policies by combining the results of multiple risk predictors to calculate a single risk score. Data feeds and inputs roll into set risk predictors. The predictors are assigned different scores and aggregated into a risk policy to determine if a user poses low, medium, or high risk to the organization and what level of authentication will be required. Administrators can create multiple risk policies and apply them in different use cases to meet business requirements.

3.4.17.5 PingOne Verify

PingOne Verify is a SaaS platform that reduces uncertainty during onboarding and prevents fraudulent registration with convenient identity verification. PingOne Verify enables secure user verification based on a government-issued document and real-time face capture (a live selfie). The Verify dashboard summarizes all transactions, which enables you to manage all verifications, exceptions, and rejections within the PingOne platform.

3.4.17.6 PingOne Authorize

PingOne Authorize is a SaaS platform that leverages real-time data to make authorization decisions for access to data, services, APIs, and other resources. Organizations increasingly want to codify their authorization requirements as policies, giving business owners the flexibility to adapt and evolve access control rules over time. Our solution helps organizations accurately control what users can see and do within applications and APIs. With an exploding number of applications, regulations, and access control requirements to manage, abstracting authorization logic to a centralized administrative control plane is the key to enabling scale and consistency.

3.4.17.7 PingID

PingID is a SaaS platform that provides an MFA solution for the workforce and partners that drastically improves organizational security posture in minutes. PingID protects applications accessed via SSO and it integrates seamlessly with Microsoft Azure AD, Active Directory Federation Services (AD FS), and Windows login, Mac login, and SSH applications.

Supported authentication methods include mobile push, email OTP, SMS OTP, TOTP authenticator apps, QR codes, FIDO2-bound biometrics, and security keys.

3.4.17.8 PingAccess

PingAccess is a centralized access security solution with a comprehensive policy engine. It provides secure access to applications and APIs down to the URL level and ensures that only authorized users can access the resources they need. PingAccess allows organizations to protect web apps, APIs, and other resources using rules and other authentication criteria.

PingAccess can be deployed within Ping Identity's SaaS offerings, in a customer cloud, as a traditional application, and within air-gapped or network segmented environments.

3.4.17.9 PingDirectory

PingDirectory is a fast, scalable directory used to store identity and rich profile data. Organizations that need maximum uptime for millions of identities use PingDirectory to securely store and manage sensitive customer, partner, and employee data. PingDirectory acts as a single source of identity truth.

Users get loaded into PingDirectory through import, API connection, manual entry or bidirectional, real-time synchronization from LDAP, RDBMS, JDBC, or SCIM data stores. Both structured and unstructured user data are secured and stored by leveraging encryption, password validators, cryptographic log signing, and more. Out-of-the-box load balancing, rate limiting, and data transformations with an integrated proxy ensure maximum server performance and user data availability at scale during peak usage.

PingDirectory can be deployed within Ping Identity's SaaS offerings, in a customer cloud, as a traditional application, and within air-gapped or network segmented environments.

3.4.18 Radiant Logic

Radiant Logic, the enterprise Identity Data Fabric company, helps organizations combat complexity and improve defenses by making identity data easy to access, manage, use, and protect. With Radiant, it's fast and easy to put identity data to work, creating the identity data foundation of the enterprise where organizations can realize meaningful business value, accelerate innovation, and achieve zero trust. Built to combat identity sprawl, enterprise technical debt, and interoperability issues, the RadiantOne platform connects many disparate identity data sources across legacy and cloud infrastructures, without disruption. It can accelerate the success of initiatives including SSO, M&A integrations, identity governance and administration, hybrid and multi-cloud environments, customer identity and access management, and more with an identity data fabric foundation. Visit http://www.radiantlogic.com/ to learn more.

3.4.18.1 RadiantOne Intelligent Identity Data Platform

The RadiantOne Intelligent Identity Data Platform builds an identity data fabric using federated identity as the foundation for zero trust. It is the single authoritative source for identity data, enabling critical initiatives by making identity data and related context available in real time to consumers regardless of where that data resides. RadiantOne's Intelligent Identity Data Platform uses patented identity unification methods to abstract and enrich identity data from multiple sources, build complete global user profiles, and deliver real-time identity data on-demand to any service or application. Zero trust relies on evaluating a rich and authoritative granular set of attributes in real time against an access policy to determine authorization. RadiantOne provides a single authoritative place for all components of the ZTA to quickly and easily request the exact data they need in the format, structure, schema, and protocol each requires. In order to provide the flexibility and scalability that organizations need, the platform is broken into six distinct modules: Federated Identity Engine; Universal Directory; Global Synchronization; Directory Migration; Insights, Reports & Administration; and Single Sign-On.

3.4.18.1.1 RadiantOne Federated Identity Engine

The Federated Identity Engine abstracts and unifies identity data from all sources (on-premises or cloud-based) to form an identity data fabric that is flexible, scalable, and turns identity data into a reusable resource. The identity data fabric provides a central access point for authoritative identity data to all applications, and encompasses all subjects, users, and objects (employees, contractors, partners, customers, members, non-enterprise employees, devices, NPEs, service accounts, bots, IoT, risk scoring, and data and other assets). RadiantOne gathers, maps, normalizes, and transforms identity data to build a de-duplicated list of users, enriched with all identity attributes to create a single global profile for each user. The Federated Identity Engine is schema-agnostic and standards-based, which allows it to build unlimited and flexible views correlated from all sources of rich and granular identity data, updated in near-real-time, and delivered at speed in the format required by all the consuming applications in the ZTA. These views are stored in a highly scalable, modern big data store kept in near-real-time sync with local identity sources of truth.

3.4.18.1.2 RadiantOne Universal Directory

The RadiantOne Universal Directory provides a modern way of storing and accessing identity information in a highly scalable, fault-tolerant, containerized solution for distributed identity storage. Its highly performant cluster architecture scales easily to hundreds of millions of objects, delivers automation, high availability, and multi-cluster deployments to easily accommodate distributed data centers. Universal Directory is FIPS 140-2 certified for securing data-in-transit and data-at-rest, and it provides detailed audit logs and reports [10]. Universal Directory is accessible by all LDAP, SQL, SCIM, and REST-enabled applications.

3.4.18.1.3 RadiantOne Single Sign On (SSO)

Single Sign On is the gateway between identity stores and applications that support federation standards—SAML, OIDC, WS-Federation—for connecting users with seamless, secure, and uniform access to federated applications. SSO enables a secure federated infrastructure, creating one access point to connect all internal identity and authentication sources for strong authentication. It also provides a self-service portal for managing passwords and user profiles.

3.4.18.1.4 RadiantOne Global Synchronization

Global Synchronization leverages bi-directional connectors to propagate identity data and keep it coherent across enterprise systems in near-real-time, regardless of the location of the underlying identity source data (on-premises, cloud-based, or hybrid). It builds a reliable and highly scalable infrastructure with a transport layer based on message queuing for guaranteed delivery of changes. Global Synchronization reduces complexity and administrative burden, simplifies provisioning and syncing identity centrally, and ensures consistency and accuracy with real-time change detection to underlying identity data attributes.

3.4.19 SailPoint

SailPoint offers identity security technologies that automate the identity lifecycle; manage the integrity of identity attributes; enforce least privilege through dynamic access controls, role-based policies, and separation of duties (SoD); and continuously assess, govern, and respond to access risks using AI and ML. SailPoint Identity Security is the cornerstone of an effective zero trust strategy. Discover more at https://www.sailpoint.com/.

3.4.19.1 IdentityIQ Platform

SailPoint IdentityIQ is an identity and access management software platform custom-built for complex enterprises. It delivers full lifecycle and compliance management for provisioning, access requests, access certifications, and SoD. The platform integrates with SailPoint's extensive library of connectors to intelligently govern access to today's essential business applications. Harnessing the power of AI and ML, SailPoint's AI Services seamlessly automate access, delivering only the required access to the correct identities and technology at the appropriate time.

As an identity governance platform, SailPoint provides organizations with a foundation that enables a compliant and secure infrastructure driven by a zero-trust approach with complete visibility of all access, frictionless automation of processes, and comprehensive integration across hybrid environments. SailPoint connects to enterprise resources to aggregate accounts and correlate with authoritative records to build a foundational identity profile from which all enterprise access is based. Users are granted birthright access based on dynamic attribute evaluation, and additional access for all integrated resources is requested and governed through a centralized SailPoint request portal. The SailPoint governance platform is enriched through its extensible API framework to support integrations with other identity security tools. The IdentityIQ platform contains two components, IdentityIQ Compliance Manager and IdentityIQ Lifecycle Manager.

3.4.19.1.1 IdentityIQ Compliance Manager

IdentityIQ Compliance Manager automates access certifications, policy management, and audit reporting to streamline compliance processes and improve the effectiveness of identity governance.

Access certification ensures least-privileged access by continuously monitoring and removing accounts and entitlements that are no longer needed.

Separation of duties policies enforce business procedures to detect and prevent inappropriate access or actions by proactively scanning for violations.

Audit reporting simplifies the collection the information needed to manage the compliance process and replaces manual searches for data located in various systems around the enterprise through an integrated platform.

3.4.19.1.2 IdentityIQ Lifecycle Manager

IdentityIQ Lifecyle Manager enables an organization to manage changes to access through user-friendly self-service requests and lifecycle events for fast, automated delivery of access to users.

Access requests enable users to request and receive access to enterprise on-premises and SaaS applications and data while ensuring compliance through policy enforcement and elevating reviews for privileged access.

Automated provisioning detects and triggers changes to a user's access based on a user joining, moving within, or leaving an organization. Direct provisioning reduces risk by automatically changing or removing accounts and access in an appropriate manner with automated role and attribute-based access.

3.4.20 Tenable

Tenable®, Inc. is the Cyber Exposure company. Organizations around the globe rely on Tenable to understand and reduce cyber risk. As the creator of Nessus®, Tenable extended its expertise in vulnerabilities to see and secure any digital asset on any computing platform.

3.4.20.1 Tenable.io

Powered by Nessus technology and managed in the cloud, Tenable.io provides comprehensive vulnerability coverage with the ability to predict which security issues to remediate first. Using an advanced asset identification algorithm, Tenable.io can provide accurate information about dynamic assets and vulnerabilities in ever-changing environments. As a cloud-delivered solution, its intuitive dashboard visualizations, comprehensive risk-based prioritization, and seamless integration with third-party solutions help security teams maximize efficiency and scale for greater productivity.

3.4.20.2 Tenable.ad

Tenable.ad is a software solution that helps organizations harden their AD by finding and fixing AD weaknesses and vulnerabilities before attacks happen. Tenable.ad Indicators of Exposure discover and prioritize weaknesses within existing AD domains and reduce exposure by following Tenable.ad step-by-step remediation guidance. Tenable.ad keeps an AD in this hardened state by continuously monitoring and alerting in real time of any new misconfigurations, while Tenable.ad Indicators of Attacks enables detection and response to AD attacks in real time. In addition, Tenable.ad tracks and records all changes to an AD, helping show the link between AD changes and malicious actions. Tenable.ad can send alerts using email or through an existing SIEM solution.

3.4.20.3 Tenable.cs

Tenable.cs is Tenable's cloud security solution to help organizations programmatically detect and fix cloud infrastructure security issues in design, build, and runtime phases of the software development lifecycle (SDLC). Tenable.cs enables organizations to establish guardrails in DevOps processes to prevent unresolved misconfigurations or vulnerabilities in Infrastructure as Code (IaC) from reaching production environments. The product monitors cloud resources deployed in AWS, Azure, and GCP to ensure any runtime changes are compliant with policies, and remediations to address configuration drifts are automatically propagated back to the IaC. Tenable.cs also provides continuous visibility to assess cloud hosts and container images for vulnerabilities whether they're deployed for days or hours, without the need to manage scan schedules, credentials, or agents. All cloud assets—including ephemeral assets—are continuously reassessed as new vulnerability detections are added and as new assets are deployed. This always-on approach allows organizations to spend more time focusing on the highest priority vulnerabilities and less time on managing scans and software.

3.4.21 Trellix

Trellix is redefining the future of cybersecurity. The company's open and native XDR platform helps organizations confronted by today's most advanced threats gain confidence in the protection and resilience of their operations. Trellix's security experts, along with an extensive partner ecosystem, accelerate technology innovation through ML and automation to empower customers. See more at https://trellix.com/. Trellix solutions can play a pivotal role in assisting organizations in meeting their zero trust outcomes through Trellix's extensive portfolio of enforcement points, rapidly growing partner ecosystem, and ability to quickly quantify risk and orchestrate responses.

Trellix offers a comprehensive portfolio of tools that align with zero trust objectives and outcomes. The following subsections discuss the tools from the portfolio currently being included in this NCCoE effort.

3.4.21.1 MVISION Complete Suite

MVISION Complete delivers a comprehensive suite of tools that provide threat and data protection across endpoints, web, and cloud. Individual products included in the MVISION Complete Suite include the following.

3.4.21.1.1 Trellix ePO

Trellix ePolicy Orchestrator (ePO) is a centralized management console for deploying, configuring, and managing Trellix endpoint security solutions including threat prevention, data protection, and EDR. For more information on Trellix ePO, please visit ePolicy Orchestrator | Trellix.

3.4.21.1.2 Trellix Insights

Trellix Insights is a threat intelligence platform integrated with the Trellix solution portfolio that enables customers to gain contextual understanding of active global threat campaigns relevant to their vertical.

Through integrated understanding of compensating controls and detection events, Insights enables organizations to predictively stay ahead of threats, quickly identify campaign activity within their environment, and receive the guidance necessary to proactively defend against campaigns. For more information on Trellix Insights, please visit Trellix Insights | Trellix.

3.4.21.1.3 Trellix Endpoint Security Platform

Trellix Endpoint Security Platform blocks malicious and targeted attacks using traditional and enhanced detection techniques as part of a layered protection strategy. Techniques include generic malware detection, behavioral detection, ML, containment, and enhanced remediation. For more information on Trellix Endpoint Security, please visit Trellix Endpoint Security | Trellix.

3.4.21.1.4 Trellix EDR

Trellix EDR collects and analyzes device trace data using advanced detection techniques in order to surface suspected threats within an enterprise. Trellix EDR empowers security operations teams to gain important context about the environment with true real-time enterprise search capabilities and integrated threat intelligence. Trellix EDR is an asset to resource-starved security operations teams working to keep up with the ever-growing threat landscape by incorporating integrated AI-assisted guided investigations. Guided investigations analyze thousands of artifacts beyond the initial detection event to replicate a traditionally manual playbook process. By automating this process, analysts can reach conclusions faster, reduce time to detection, and accelerate confident response activities. For more information on Trellix EDR, please visit Trellix EDR – Endpoint Detection & Response | Trellix.

3.4.21.1.5 Trellix DLP Endpoint

Trellix DLP Endpoint enables organizations to discover, control, and block access to sensitive data on the endpoint. Trellix DLP Endpoint integrates with identity providers to assign policy based on users' roles and groups, and in a ZTA can adjust data protection policy as user trust changes. Additionally, DLP Endpoint is managed by ePO, and it includes a full case management system for aggregating multiple DLP incidents and identifying malicious insiders. For more information on Trellix DLP Endpoint, please visit DLP Endpoint | Trellix.

3.4.21.1.6 Skyhigh Security SSE Platform

Skyhigh Security, once part of Trellix's foundational company, McAfee Enterprise, has been established as a separate business entity and sister company to Trellix. Skyhigh Security's Security Service Edge (SSE) platform is part of the MVISION Complete Suite, delivered by Skyhigh Security, and offers comprehensive protection for cloud, web, and data protection. Skyhigh Security integrates a CASB platform with strong cloud-hosted web security and data protection controls to deliver a highly secure, highly available platform for protecting hybrid and multi-cloud enterprises. For more information on Skyhigh Security's SSE platform please visit What is SSE? | Security Service Edge | Skyhigh Security.

The MVISION Complete Suite aids in the ability to meet zero trust objectives by delivering device-level protection and alerting, application protection through contextual access controls, user trust through user activity monitoring, data security through comprehensive data protection and discovery, and analytics and intelligence through EDR and Insights.

3.4.21.2 Full Remote Browser Isolation

Remote browser isolation enables organizations to fully contain web applications within a secure container to prevent malware and data leakage and provide complete control over a browser session. The Skyhigh SSE solution out of the box offers remote browser isolation for risky websites to ensure no implicit trust is being granted to web applications prior to trust validation. In some cases, organizations would choose that no implicit trust is ever extended to web traffic, regardless of known reputation. In this scenario, full-time browser isolation is required to meet this objective. The Trellix offering, with sister company Skyhigh Security, includes the ability for full remote browser isolation as an add-on module. For more information on Remote Browser Isolation, see [Remote Browser Isolation | McAfee Products].

3.4.21.3 Helix (XDR)

To achieve zero trust outcomes, it is necessary to have a common platform that applies AI-driven, real-time threat intelligence to data collected from devices and security sensors as a mechanism for surfacing advanced attacks and associated entity risk, and to orchestrate proactive and remediating responses across native and open security tools. Within many zero trust reference architectures, this platform could be considered the dynamic access control plane, or the trust algorithm.

Trellix delivers this capability through Helix. Helix is a cloud-hosted, intelligence-driven platform that collects data from over 600 different sensors and point solutions, analyzes the data against known threats, behaviors, and campaigns using AI and enhanced detection rules, and powers automated and manual responses across Trellix native and third-party policy engines. For more information on Trellix XDR, see [Trellix-Platform | Trellix].

3.4.21.4 CloudVisory

It's no secret that cloud services are now pervasive; many applications have been moved either through SaaS or cloud services development to cloud data centers. This presents new challenges for many organizations as they work to gain better visibility and control over IaaS-hosted cloud applications and the thousands of micro-services that support them. As organizations look to adopt zero trust principles within the cloud, it will become imperative that proper service configuration, IAM roles, cloud network traffic, and workloads are fully evaluated for risk and protected. CloudVisory supports these objectives through:

- CI/CD integration to ensure proper service configuration, and continuous posture assessments to guard against configuration drift
- IAM policy inspection

- intelligent network micro-segmentation
- intra-cloud and cloud-to-cloud network monitoring
- multi-cloud support

For more information on CloudVisory, see ds-cloudvisory.pdf (fireeye.com).

3.4.22 VMware

VMware's content will be included in the next draft version of this practice guide.

3.4.23 Zimperium

Zimperium secures both mobile devices and applications so they can safely and securely access data. Patented on-device ML-based security provides visibility and protection against known and zero-day threats and attacks.

3.4.23.1 Zimperium Mobile Threat Defense

Zimperium Mobile Threat Defense is an advanced MTD solution for enterprises, providing persistent, on-device protection to both corporate-owned and BYOD devices against modern attack vectors. Leveraging Zimperium's patented z9 on-device detection engine, Zimperium MTD detects threats across the kill chain, including device compromise, network, phishing, and application attacks.

Zimperium's MTD provides on-device behavior detection via an on-device agent, even when the device is not connected to a network. Zimperium's MTD begins protecting devices against all primary attack vectors immediately after deployment. The Zimperium zConsole provides a management interface used to configure threat policies, manage device groups/users, and view events and the forensics that are associated with those events.

Zimperium provides critical mobile security data for organizations, with integrations into multiple, concurrent enterprise SIEM/SOAR, UEM, XDR, and IAM platforms. Data is securely shared via REST API, syslog, etc. Zimperium MTD provides comprehensive *device attestation* enabling a complete picture of mobile endpoint security and increased visibility into risks such as jailbreak detections. Zimperium MTD provides continuous protection for mobile devices, providing the risk intelligence and forensic data necessary for security administrators to raise their mobile security confidence. Zimperium integrates mobile threat data into security reporting systems and processes. Using Zimperium's vast integrations ecosystem, mobile device state, security posture, events, etc. are shared, enabling multimodal protections to be automatically deployed, including "conditional access" to sensitive information via MDM/UEMs, SOAR, and IAM, for example. Zimperium MTD protects devices against all primary attack vectors, including via USB, removable storage, and even when the device is not connected to a network.

3.4.24 Zscaler

Zscaler provides secure user access to public-facing sites and on- or off-premises private applications via the Zscaler Zero Trust Exchange, a cloud-delivered security service edge technology. The Zero Trust Exchange helps IT move away from legacy network infrastructure to achieve modern workforce enablement, infrastructure modernization, and security transformation.

Zscaler’s role in the ZTA is to provide full visibility and control of context-based, least-privilege access to internet and SaaS applications as well as private applications in IaaS, PaaS, or internally hosted environments via the Zero Trust Exchange.

3.4.24.1 Zscaler Zero Trust Exchange

Users accessing the internet or a SaaS application can leverage the **Zscaler Internet Access (ZIA)** solution. This solution delivers a comprehensive security stack—including TLS inspection, advanced firewall, SWG, DLP, virus protection, and sandbox capabilities—for end-users, which follows them no matter where they are.

Users accessing private applications either locally or in the cloud can leverage the **Zscaler Private Access (ZPA)** solution, which also provides a virtual PDP+PEP in the cloud.

The **Zscaler Client Connector** brokers access for both ZIA and ZPA, offering lightweight single-agent protection and visibility, as well as optionally gathering telemetry for end-user experience monitoring.

Combining ZIA and ZPA provides a FedRAMP-accredited solution that organizations can integrate into their unique digital ecosystems today. Moreover, since Zscaler is an integral part of any zero trust framework, organizations can leverage Zscaler's cloud service provider, EDR, SIEM/SOAR, and SD-WAN integration partnerships with Microsoft, AWS, Okta, CrowdStrike, and other industry leaders to promote data visibility and access management.

4 Architecture

The project architecture is designed to include the core zero trust logical components as depicted in NIST SP 800-207. In Section 4.1 we present a general ZTA and describe its components and operation. These components may be operated as either on-premises or cloud-based services.

In Section 4.2 we describe a particular version of this general ZTA that we call the *EIG crawl phase* reference architecture. Three of the ZTA builds that are documented in this practice guide are instantiations of this EIG crawl phase reference architecture. This architecture relies mainly on ICAM and endpoint protection platform (EPP) components, does not include any components that are specifically dedicated to providing PE or PA functionality, and is currently limited to protecting on-premises resources.

In Section 4.3 we describe a second version of the general ZTA that we call the *EIG run phase* reference architecture. Two of the ZTA builds that are documented in this practice guide are instantiations of this EIG run phase reference architecture. Like the EIG crawl phase architecture, the EIG run phase architecture bases resource access decisions mainly on information provided by ICAM and EPP components. However, unlike the EIG crawl phase architecture, it may include PA and PE components that are not furnished by the ICAM provider. The EIG run phase architecture also protects both on-premises and cloud resources, and it supports device discovery and the establishment of tunnels between requesting endpoints and resources.

In Section 4.4 we describe the physical architecture of the baseline laboratory environment in which we implemented all of the builds documented in this guide.

Volume B will be updated throughout the project lifecycle as the architecture evolves to include additional functionalities, security capabilities, and ZTA deployment models.

4.1 General ZTA Reference Architecture

Figure 4-1 depicts the high-level logical architecture of a general ZTA reference design independent of deployment models. It consists of three types of core components: PEs, PAs, and PEPs, 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 ICAM, EDR/EPP, security analytics, and data security. Specific capabilities that fall into each of these supporting component categories are discussed in more detail later in this section. 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). Each of the logical components in the reference architecture does not necessarily directly correlate to physical (hardware or software) components. In fact, 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, EDR/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 (devices, end users, applications, servers, and other non-human entities that request information from resources) request and receive access to enterprise resources via the ZTA. Human subjects (i.e., users) are authenticated. Non-human subjects are both authenticated and protected by endpoint security. Enterprise resources may be located on-premises or in the cloud. Existing enterprise subjects and resources are not part of the reference architecture itself; however, any changes required to existing endpoints, such as installing ZTA agents, should be considered part of the reference architecture.

Figure 4-1 General ZTA Reference Architecture

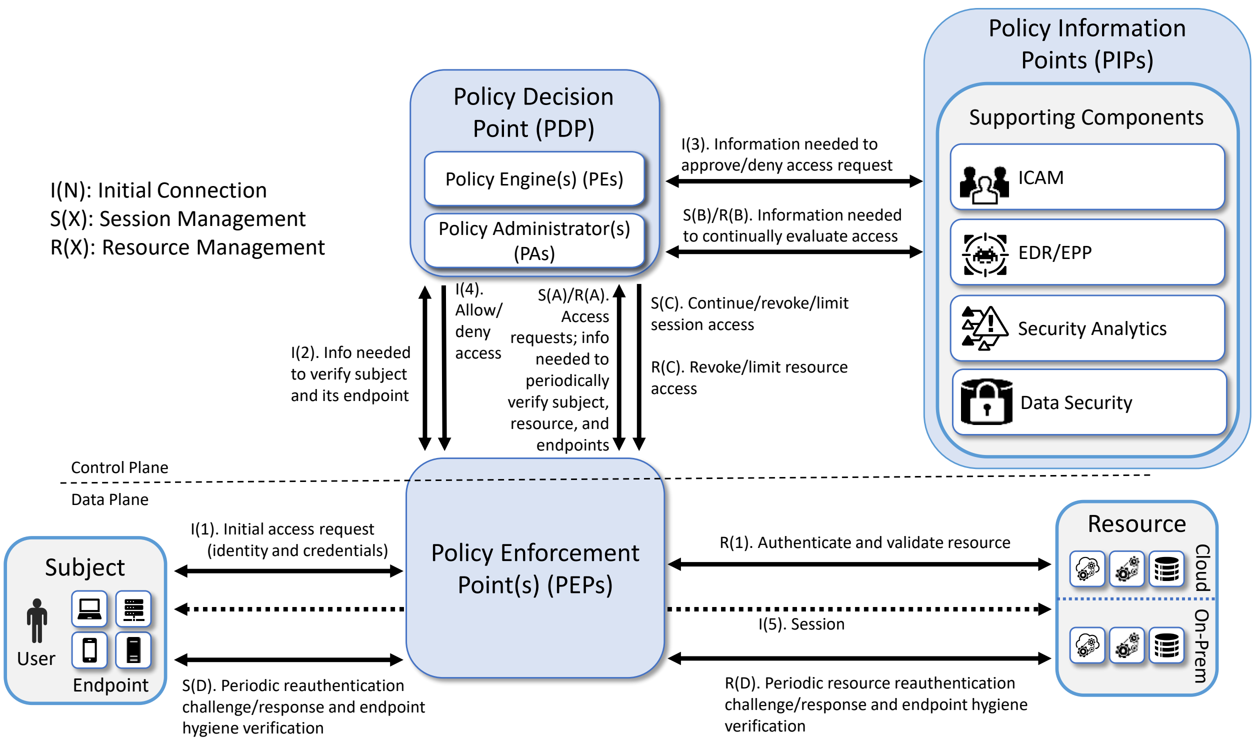

4.1.1 ZTA Core Components

The types of ZTA core components are:

- **Policy Engine (PE):** The PE handles the ultimate decision to grant, deny, or revoke access to a resource for a given subject. The PE calculates the trust scores/confidence levels and ultimate access decisions based on enterprise policy and information from supporting components. The PE executes its trust algorithm to evaluate each resource request it receives. The PE may be a single system or a federation of systems (i.e., a “system of systems”) that covers sectors of the

ZTA. Each PE in the federation would be responsible for its sector based on the overall set of enterprise policies.

- **Policy Administrator (PA):** The PA executes the PE's policy decision by sending commands to the PEP to establish and terminate the communications path between the subject and the resource. It generates any session-specific authentication and authorization token or credential used by the subject to access the enterprise resource.
- **Policy Enforcement Point (PEP):** The PEP guards the trust zone that hosts one or more enterprise resources. It handles enabling, monitoring, and eventually terminating connections between subjects and enterprise resources. It operates based on commands that it receives from the PA.

When combined, the functions of the PE and PA comprise a PDP. The PDP is where the decision as to whether or not to permit a subject to access a resource is made. The PIPs provide various types of telemetry and other information needed for the PDP to make informed access decisions. The PEP is the location at which this access decision is enforced.

Three approaches for how an enterprise can enact a ZTA for workflows can be supported by the architecture represented in Figure 4-1: use of EIG, micro-segmentation, and SDP. If the micro-segmentation approach is used, then when the PEP grants a subject access to a resource, it permits the subject to gain access to the unique network segment on which the resource resides. If the SDP approach is used, then when the PE decides to grant a subject access to a resource, the PA often acts like a network controller by setting up a secure channel between the subject and the resource via the PEP.

4.1.2 ZTA Supporting Components

The various sets of information either generated via policy or collected by the ZTA supporting components and used as input to ZTA policy decisions are referred to as PIPs.

The ZTA supporting components and policy information points are:

- **ICAM:** The ICAM component includes the strategy, technology, and governance for creating, storing, and managing subject (e.g., enterprise user) accounts and identity records and their access to enterprise resources. Aspects of ICAM include:
  - **Identity management** – Creation and management of enterprise user and device accounts, identity records, role information, and access attributes that form the basis of access decisions within an organization to ensure the correct subjects have the appropriate access to the correct resources at the appropriate time. This includes least privilege management, i.e., ensuring that the subject performing the access is given just enough privileges at the time they are needed to complete the task at hand and then removing those privileges to ensure that subjects do not have privileges that are not required. This concept can be characterized as just enough and just in time access rights.

- **Access and credential management** – Use of authentication (e.g., SSO and MFA) to verify subject identity and authorization to manage access to resources. This includes continuous access evaluation, i.e., repeatedly authenticating subjects and verifying their access to resources on an ongoing basis throughout an access session.
- **Federated identity** – The federated identity component aggregates and correlates all attributes relating to an identity or object that is being authorized by a ZTA. It enables users of one domain to securely access data or systems of another domain seamlessly, and without the need for completely redundant user administration. Federated identity encompasses the traditional ICAM data, supports identities that may be part of a larger federated ICAM community, and may include non-enterprise employees. Guidelines for the use of federated identity are discussed in NIST SP 800-63C, *Digital Identity Guidelines* [11].
- **Identity governance** – Use of policy-based centralized automated processes to manage user identity and access control functions (e.g., segregation of duties, role management, logging, access reviews, auditing, analytics, reporting) to ensure compliance with requirements and regulations

- **EDR/EPP:** The endpoint protection component encompasses the strategy, technology, and governance to protect endpoints (e.g., servers, desktops, mobile phones, IoT devices and other non-human devices) and their data from threats and attacks, as well as protect the enterprise from threats from managed and unmanaged devices. In some cases, extended detection and response (XDR) solutions may be used that consolidate multiple EDR/EPP, network monitoring, and other security tools into a unified security solution. Such a unified solution provides automated monitoring, analysis, detection, and remediation for the purpose of improving detection accuracy while simultaneously improving efficiency of security operations and remediation. Some EDR/EPP solutions may depend on EDR/EPP agents being installed on endpoints while other solutions may be agentless. Aspects of endpoint protection include:
  - **Continuous diagnostics and mitigation (CDM) –** Gathering information about enterprise assets and their current state and applying updates to configuration and software components. A CDM system provides information to the policy engine about the asset making the access request. Guidelines for applying patches and updates are discussed in NIST SP 1800-31, Improving Enterprise Patching for General IT Systems: Utilizing Existing Tools and Performing Processes in Better Ways.
  - **Application protection –** Managing and protecting data within an application by enforcing protection policies that apply to the application
  - **Device compliance –** Ensuring that an endpoint contains the hardware, firmware, software, and configurations required by enterprise policy and includes nothing unauthorized by enterprise policy. Guidelines for validating the integrity of computing devices are discussed in NIST SP 1800-34, Validating the Integrity of Computing Devices.

   - **Vulnerability/threat mitigation –** Monitoring endpoint software and configurations to detect known vulnerabilities and, when found, provide alerts that include remediation and mitigation recommendations, if available
   - **Host intrusion protection –** Monitoring an endpoint for suspicious activity that may indicate an attempted intrusion, infection, or other malware; stopping malicious activity on the endpoint, notifying potential victims, logging the suspicious events, and stopping future traffic from suspicious sources
   - **Host firewall –** Preventing the individual endpoint from receiving traffic that is not explicitly permitted, thereby helping to protect the endpoint from receiving malware and other malicious traffic
   - **Malware protection –** Scanning endpoint software for signatures that belong to known malware or using non-signature-based offerings that may use ML or AI to detect malicious code; if detected, disabling the malware, quarantining and repairing infected files if possible, and providing alerts that include any available remediation and mitigation recommendations
   - **Data protection enforcement –** Ensuring that data stored on the device is protected in accordance with enterprise policies
   - **Mobile device management –** Managing and administering mobile devices to ensure that they are secure by provisioning software to the mobile devices in accordance with enterprise security policies to monitor behavior and critical data on the device, thereby protecting the device's applications, data, and content and enabling the device to be tracked, monitored, troubleshooted, and wiped, if necessary
- **Data Security:** The data security component includes the policies that an enterprise needs to secure access to enterprise resources, as well as the means to protect data at rest and in transit. Aspects of data security include:
   - **Data confidentiality –** protecting data from unauthorized disclosure while at rest and in transit
   - **Data integrity –** protecting data from unauthorized modification while at rest and in transit
   - **Data availability –** protecting the ability of authorized users to access data in a timely manner and guarding against unauthorized deletion
   - **Data access policies –** all data access policies and rules needed to secure access to enterprise information and resources
- **Security Analytics:** The security analytics component encompasses all the threat intelligence feeds and traffic/activity monitoring for an IT enterprise. It gathers security and behavior analytics about the current state of enterprise assets and continuously monitors those assets to

actively respond to threats or malicious activity. This information could feed the policy engine to help make dynamic access decisions. Aspects of security analytics include:

  - **SIEM** – Collection and consolidation of security information and security event data from many sources; correlates and analyzes the data to help detect anomalies and recognize potential threats and vulnerabilities; logs the data to adhere to data compliance requirements
  - **Network monitoring and activity logging** – Collection and monitoring of metrics regarding network activity and performance. Collect asset logs, network traffic, resource access actions, privileged tasks, and other events that provide real-time (or near-real-time) feedback on the security posture of enterprise information systems.
  - **Traffic inspection** –Interception, examination, and recording of relevant traffic transmitted on the network. Not all communication may be intercepted and not all intercepted traffic may be subject to the same level of examination (e.g., deep packet inspection, only metadata analysis) depending on policy or capability.
  - **Endpoint monitoring** – The discovery of all IP-connected endpoints and continuous collection, examination, and analysis of software versions, configurations, and other information regarding hosts (devices or VMs) that are connected to the network
  - **Threat intelligence** – Use of information regarding known existing or emerging vulnerabilities, attacks, and other menaces to enterprise operations and assets to inform decisions regarding how to defend against and respond to those threats
  - **User behavior** – Monitoring and analysis of user behavior to detect unusual patterns or anomalies that might indicate an attack
  - **Correlation and analytics** – Use of data analytics and AI to correlate, compare, and analyze all information received from ZTA supporting components (e.g., ICAM, endpoint monitoring, network monitoring, and other related supporting activity) for the purpose of detecting unusual patterns or anomalies that might indicate an attack
  - **SOAR** – Collection and monitoring of alerts from the SIEM and other security systems and execution of predefined incident response workflows to automatically analyze the information and orchestrate the operations required to respond
  - **Security validation** – Continuous validation and measurement of the effectiveness of cybersecurity controls
  - **Firmware assurance** – Continuous monitoring of IT device firmware

4.1.3 ZTA in Operation

Figure 4-1 depicts the general, high-level ZTA reference architecture. If an enterprise has highly distributed systems, it may have many PEPs to protect resources in different locations; it may also have multiple PEPs to support load balancing. For simplicity, Figure 4-1 limits its focus to the interactions involving a single PEP, a single subject, and a single resource. The labeled arrows in Figure 4-1 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 Establishment Steps—I():** Session establishment 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.

The following additional details describe each of the steps in each of the three processes depicted in Figure 4-1:

Resource Management

- **Step R(1). Authenticate and validate resource:** In our model, it is assumed that the resource has already been registered as an authorized resource. Initially, when the resource is brought online, its identity must be authenticated and its endpoint hygiene must be validated to ensure compliance. This authentication and validation could be accomplished by a variety of mechanisms, such as the ICAM and EPP capabilities, the PEP itself, or a connector. The diagram is not concerned with depicting how it is authenticated, just that the authentication and validation are performed.

In some implementations, in order for the resource to communicate with the service provider where the PEP is located, a connector or proxy may need to be installed to enable that connection to the service provider. For example, a database in an existing enterprise may not currently have the capability to interact with a service provider PEP directly. To make this communication possible, a connector, which behaves like a proxy module, may be installed between the resource and the PEP. There are multiple possible types of connectors and ways of connecting. This level of detail (i.e., whether a connector is present and, if so, what type) is not shown in the figure. Authentication and validation of the resource and connection of the resource to the PEP must be completed prior to any users requesting access.

- **Step R(A). Information needed to periodically verify resource and endpoint:** Throughout the lifetime of the session, the PEP will periodically challenge the resource to reauthenticate itself. After doing so, the PEP will provide the PDP with the identity and credentials that the resource provided. Similarly, throughout the lifetime of the session, the PEP will request hygiene information from the resource's endpoint. After obtaining this hygiene information, the PEP will provide it to the PDP. The frequency with which the resource should be issued authentication challenges is determined by enterprise policy, as is the frequency with which the hygiene of its endpoint should be validated.
- **Step R(B). Information needed to continually evaluate access:** Throughout the course of the access session, the PDP requests and receives any resource-related information that it needs to evaluate the resource's ongoing compliance with enterprise policy. This could include information such as authentication information provided by the ICAM system, endpoint hygiene information provided by the EPP, and anomaly detection analysis regarding resource behavior provided by logging and security analytics functionality.
- **Step R(C). Revoke/limit resource access:** The connection between the PEP and the resource may be terminated or reconfigured based on changes to the resource or operating environment that indicate the resource no longer conforms to enterprise policy.
- **Step R(D). Periodic resource reauthentication challenge/response and endpoint hygiene verification:** The resource undergoes continual reauthentication and hygiene checks to ensure that its security posture conforms to enterprise policy. These actions are usually taken by the various systems that may make up the PDP and are performed regardless of any current open sessions. The frequency with which reauthentication and hygiene checks are performed is determined by enterprise policy.

Session Establishment

- **Step I(1)**. **Initial access request (identity and credentials):** The subject interacts with the PEP to request access to the resource and provide its identity and credentials.

- **Step I(2)**. **Information needed to verify subject and its endpoint:** The PEP forwards the subject's identity and credentials to the PE within the PDP.
- **Step I(3)**. **Information needed to approve/deny access request:** The PE requests and receives any additional information that it needs to determine whether it should approve or deny the subject's access request. This includes information provided by the various supporting components of the ZTA. ICAM-related information is used most heavily, i.e., user and endpoint identity, authorization (i.e., subject privileges), federation, and identity governance information; but additional information from other ZTA supporting components, e.g., endpoint compliance, endpoint monitoring, and threat intelligence, may also be relied upon as specified by enterprise policy. The PIPs depicted in Figure 4-1 represent the collection of information required by the PE to decide, in accordance with enterprise policy, whether or not to grant the access request. The PE authenticates the subject, determines what the subject's authorizations are, and evaluates additional information as needed to determine whether to allow or deny the subject access to the requested resource.
- **Step I(4)**. **Allow/deny access:** The PDP informs the PEP whether to allow or deny the subject access to the resource.
- **Step I(5)**. **Session:** Assuming the PDP has decided to allow access, the PEP establishes a session between the subject and the resource through which the subject can access the resource. At the completion of Step I(5), the session is set up and the session management processes begin being performed.

Session Management

Once the session has been established, several session management processes are performed simultaneously on an ongoing basis for the duration of the session. The session management processes depicted in Figure 4-1 include ongoing evaluation of each of the subject's access requests, ongoing continual evaluation of the session, periodic reauthentication of the subject, and periodic verification of the subject's endpoint hygiene. These processes are described below.

Ongoing evaluation of the access requests made by the subject: The steps of this process are depicted by steps S(A), S(B), and S(C) in Figure 4-1.

- **Step S(A). Access requests:** Throughout the course of the access session, the actions that the subject sends to the resource are monitored by the PEP and sent to the PDP for evaluation as to whether the access should continue. When TLS or another form of encryption is used to secure the session between the subject and the resource, it is not possible for a PEP that is situated in the middle of that connection to have visibility into the messages that the subject is sending because they are encrypted. The PEP must have access to the necessary unencrypted traffic needed in order to provide the PDP with the necessary information to make the access decision. The PEP may have full access to monitor the session traffic or may rely on another system (including the resource itself) to monitor the session activity. To enable the access session to be continuously monitored by the PEP, the PEP could be situated adjacent to the subject so it can receive unencrypted requests from the subject and send them to the PDP for monitoring before forwarding them over the encrypted access session to the resource; the PEP could be situated adjacent to the resource so it can decrypt requests it receives from the subject on the access session and send them to the PDP for monitoring before forwarding them to the resource; or the PEP could be located elsewhere and have plaintext requests forwarded to it that it would then send to the PDP for monitoring. Because there are many possible ways the monitoring could be accomplished, Figure 4-1 does not attempt to depict where the access session is terminated with respect to the PEP. It is only meant to convey the fact that the subject's access requests are monitored on an ongoing basis and forwarded to the PDP for evaluation.

- **Step S(B). Information needed to continually evaluate access:** Throughout the course of the access session, the PDP requests and receives any additional information from the PIP that it needs to evaluate the subject's ongoing access to determine whether it should continue. This information is provided by the various ZTA supporting components in the architecture. Examples of such information include subject identity information provided by ICAM functionality, subject endpoint hygiene information provided by endpoint security functionality, and behavioral analysis (e.g., whether the subject has attempted to elevate privileges beyond what is authorized) and anomaly detection information provided by logging and security analytics functionality. Evaluation of the access requests is performed in accordance with enterprise policy.
- **Step S(C). Continue/revoke/limit session access:** If the PDP determines that the access should continue, it will allow the PEP to forward the access request made in step S(A) to the resource. However, if the PDP determines that, in light of the information received from the PIP (e.g., federated identity, endpoint security information, security analytics), the session should be terminated or limited, the PDP may inform the PEP not to forward the action to the resource. Note that in an ideal world, the PEP would wait for the PDP to pass judgement on every request that is made on a session before forwarding each request to the resource. However, in reality, the cost of having the PDP evaluate every individual request in real time may be too great. In most cases the PEP would have a set of rules determining allowed requests and (possibly) a set of policies on when to require reauthentication or additional checks before forwarding requests to the resource.

Ongoing continual evaluation of the session: The steps of this process are depicted by steps S(B) and S(C) in Figure 4-1.

- **Step S(B). Information needed to continually evaluate access:** Throughout the course of the access session, the information in the PIPs is updated by the various ZTA supporting components and made available to the PDP so it can dynamically evaluate whether the session continues to be in accordance with enterprise policy. At any moment, information could become available that causes the session to be non-compliant. For example, threat intelligence information could be received regarding vulnerabilities in the endpoint or software used by the subject, anomalies could be detected in the subject's behavior (e.g., attempts to elevate access), or the subject could fail authentication when trying to access a different resource.

- **Step S(C). Continue/revoke/limit session access:** If the PDP determines that the ongoing access session continues to be compliant, it will permit it to continue. However, if the PDP determines that, based on information available from the PIPs (e.g., endpoint security information, threat intelligence, security analytics), the access session should be limited or revoked, the PDP will direct the PEP to deny some requests that are made on the session or to disconnect the session altogether.

Periodic reauthentication of the subject and periodic verification of the hygiene of the subject endpoint: These are two separate and distinct processes, but they are depicted by the same steps in Figure 4-1, steps S(A), S(D), and S(C), so we will discuss them together:

- **Step S(A). Information needed to periodically verify subject and endpoint:** Throughout the lifetime of the session, the PDP will periodically notify the PEP to challenge the subject to reauthenticate itself. After doing so, the PEP will provide the PDP with the identity and credentials that the subject provided. Similarly, throughout the lifetime of the session, the PDP will periodically notify the PEP to request hygiene information from the subject's endpoint, operating environment, etc. After obtaining this hygiene information, the PEP will provide it to the PDP. The frequency with which the subject should be issued authentication challenges is determined by enterprise policy, as is the frequency with which the hygiene of the subject endpoint should be validated.
- **Step S(D). Periodic reauthentication challenge/response and endpoint hygiene verification:** As directed by the PDP in step S(A), the PEP periodically issues reauthentication challenges to the subject. It also periodically requests and receives endpoint hygiene (software, configuration, etc.) information. The frequency with which each of these types of information is requested is specified by enterprise policy.
- **Step S(C). Continue/revoke/limit session access:** Based on the subject identity and credential information received and/or on the endpoint hygiene information received, the PDP determines whether to permit the access session to continue. If at any time the reauthentication of the subject fails or if the subject's endpoint hygiene cannot be satisfactorily verified (as determined by policy), the PDP will direct the PEP to disconnect or limit the session.

4.2 EIG Crawl Phase Reference Architecture

The reference architecture depicted in Figure 4-1 is intentionally general and is not meant to describe any particular ZTA deployment approach. This project plans to implement all three deployment approaches described in NIST SP 800-207, *Zero Trust Architecture*, beginning with EIG. The EIG approach to developing a ZTA uses the identity of subjects as the key component of policy creation. Access privileges granted to the given subject is the main requirement for resource access. Other factors such as device used, endpoint hygiene and status, and environmental factors may also impact whether and what access is authorized.

Once the EIG approach has been built, additional supporting components and features related to the micro-segmentation and SDP deployment approaches will be added to create a series of subsequent builds that support an increasingly rich set of additional ZTA capabilities, ultimately culminating in the demonstration of a full collection of EIG, micro-segmentation, and SDP-based ZTA functionality.

This section of the practice guide documents the builds that were created in the project's EIG crawl phase. The crawl phase uses what we call an *EIG crawl phase* deployment approach. Figure 4-2 depicts the reference architecture for this approach. The EIG crawl phase reference architecture, as its name suggests, uses a subject's identity and its access privileges as the main determinants for granting resource access, along with the endpoint used and its hygiene status. Hence, as can be seen in Figure 4-2, the reference architecture for this EIG crawl phase build includes ICAM and endpoint protection components. In the area of ICAM, it supports capabilities in all the four main areas of identity management, access and credential management, federated identity, and identity governance.

The labeled steps in Figure 4-2 are the same as those in Figure 4-1. The main difference between the two figures can be found in the set of supporting components that have been included. The EIG crawl phase reference architecture depicted in Figure 4-2 is a constrained form of the general ZTA reference architecture in Figure 4-1. The EIG crawl phase reference architecture relies on the PE and PA capabilities provided by its ICAM components. Also, the only security analytics functionality that it includes is a SIEM. It does not include any additional data security or security analytics functionality. 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 will be able to support without having to add ZTA-specific capabilities.

Figure 4-2 EIG Crawl Phase Reference Architecture

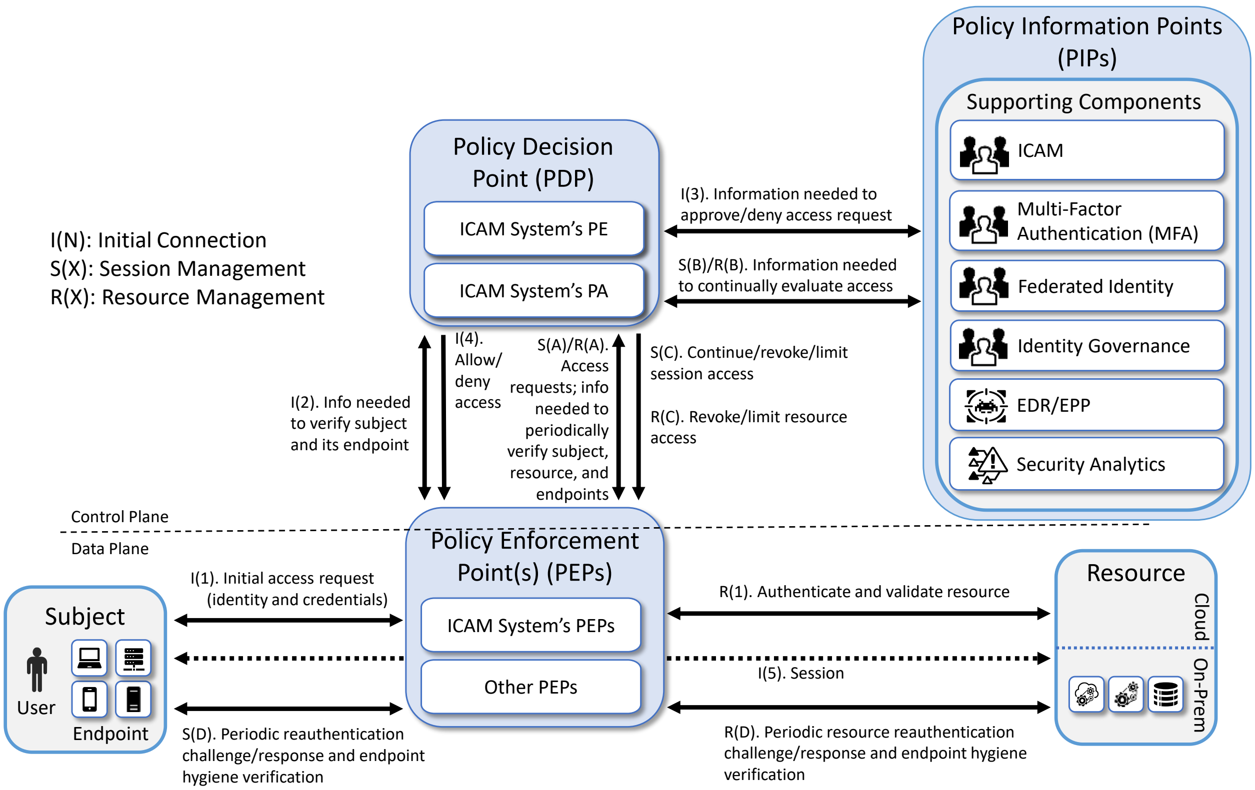

Three EIG crawl phase builds have been implemented. Each of these EIG crawl phase builds instantiates the architecture that is depicted in Figure 4-2 in a unique way, depending on the equipment used and the capabilities supported. The products used in each build were based on having out-of-box integration. Briefly, the three builds are as follows:

- **EIG Enterprise 1 Build 1 (E1B1)** uses products from Amazon Web Services, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zimperium. Certificates from DigiCert are also used.
- **EIG Enterprise 2 Build 1 (E2B1)** uses products from Cisco Systems, IBM, Mandiant, Palo Alto Networks, Ping Identity, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used.
- **EIG Enterprise 3 Build 1 (E3B1)** uses products from F5, Forescout, Lookout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used.

Each of these builds is described in detail in its own appendix (see Appendix D, Appendix E, and Appendix F).

4.3 EIG Run Phase

This section of the practice guide documents the builds that have been created in the project's EIG run phase. The EIG run phase builds upon the EIG crawl phase architecture. The EIG run phase no longer imposes the requirement that the PE and PA components are provided by the ICAM products used in the build. It also adds capabilities to the EIG crawl phase. In addition to protecting access to resources that are located on-premises, the run phase protects access to some resources that are hosted in the cloud. The EIG run phase also includes a device discovery capability, which is performed as part of the baseline. 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.

Two of the builds implemented so far and discussed in the appendices of this document are EIG run phase deployments. Each of these EIG run phase builds is unique, based on the equipment used and the capabilities supported. Briefly, the two builds are as follows:

- **EIG Enterprise 1 Build 2 (E1B2)** uses products from Amazon Web Services, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zscaler. Certificates from DigiCert are also used.
- **EIG Enterprise 3 Build 2 (E3B2)** uses products from F5, Forescout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used.

Each of these builds is described in detail in its own appendix (see Appendix H and Appendix J).

4.4 ZTA Laboratory Physical Architecture

Figure 4-3 depicts the high-level physical architecture of the ZTA laboratory environment, which is located at the NCCoE site. The NCCoE provides VM resources and physical infrastructure for the ZTA lab. It also hosts GitLab, which is used as a DevOps platform that stores Terraform and Ansible configuration information and provides version control for configuration file and change management activities. The NCCoE hosts all the collaborators' ZTA-related software for Enterprises 1, 2, 3, and 4. The NCCoE also provides connectivity from the ZTA lab to the NIST Data Center, which provides connectivity to the internet and public IP spaces (both IPv4 and IPv6).

Access to and from the ZTA lab from within ITOPS is protected by a Palo Alto Networks Next Generation Firewall (PA-5250). (The brick box icons in Figure 4-3 represent firewalls.) The ZTA lab network infrastructure includes four independent enterprises (Enterprises 1, 2, 3, and 4), a branch office used only by Enterprise 1, a coffee shop that all enterprises can use, a management and orchestration domain, and an emulated 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 ITOPS. The ZTA laboratory network has access to cloud services provided by AWS, Azure, 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 ZT deployment model is instantiated.

Figure 4-3 Physical Architecture of ZTA Lab

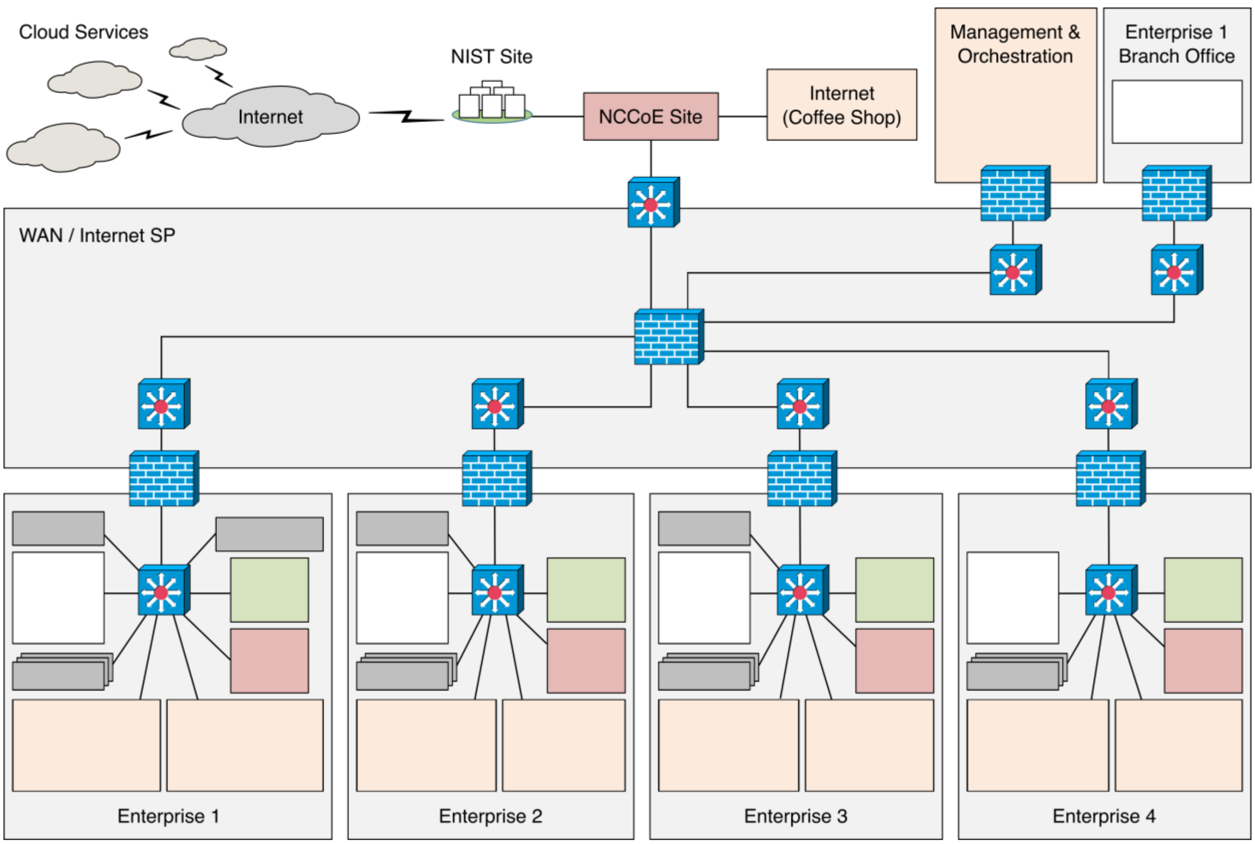

The details of the baseline physical architecture of enterprise 1, enterprise 1 branch office, enterprises 2, 3, and 4, the management and orchestration domain, and the coffee shop, as well as the baseline software running on this physical architecture are described in the subsections below. The details of each of the builds that occupy Enterprises 1, 2, and 3 are provided in the appendices. Table 4-1 maps each build to the appendix where each is described.

Table 4-1 Mapping of Builds to Architectures and Appendices

Build	ZTA Architecture Instantiated	Appendix
E1B1	EIG Crawl	Appendix D
E2B1	EIG Crawl	Appendix E
E3B1	EIG Crawl	Appendix F
E1B2	EIG Run	Appendix H
E3B2	EIG Run	Appendix J

4.4.1 Enterprise 1

Figure 4-4 is a close-up of the high-level physical architecture of Enterprise 1 in the NCCoE laboratory baseline environment. Its components are described in the subsections below.

2301 **Figure 4-4 Physical Architecture of Enterprise 1**

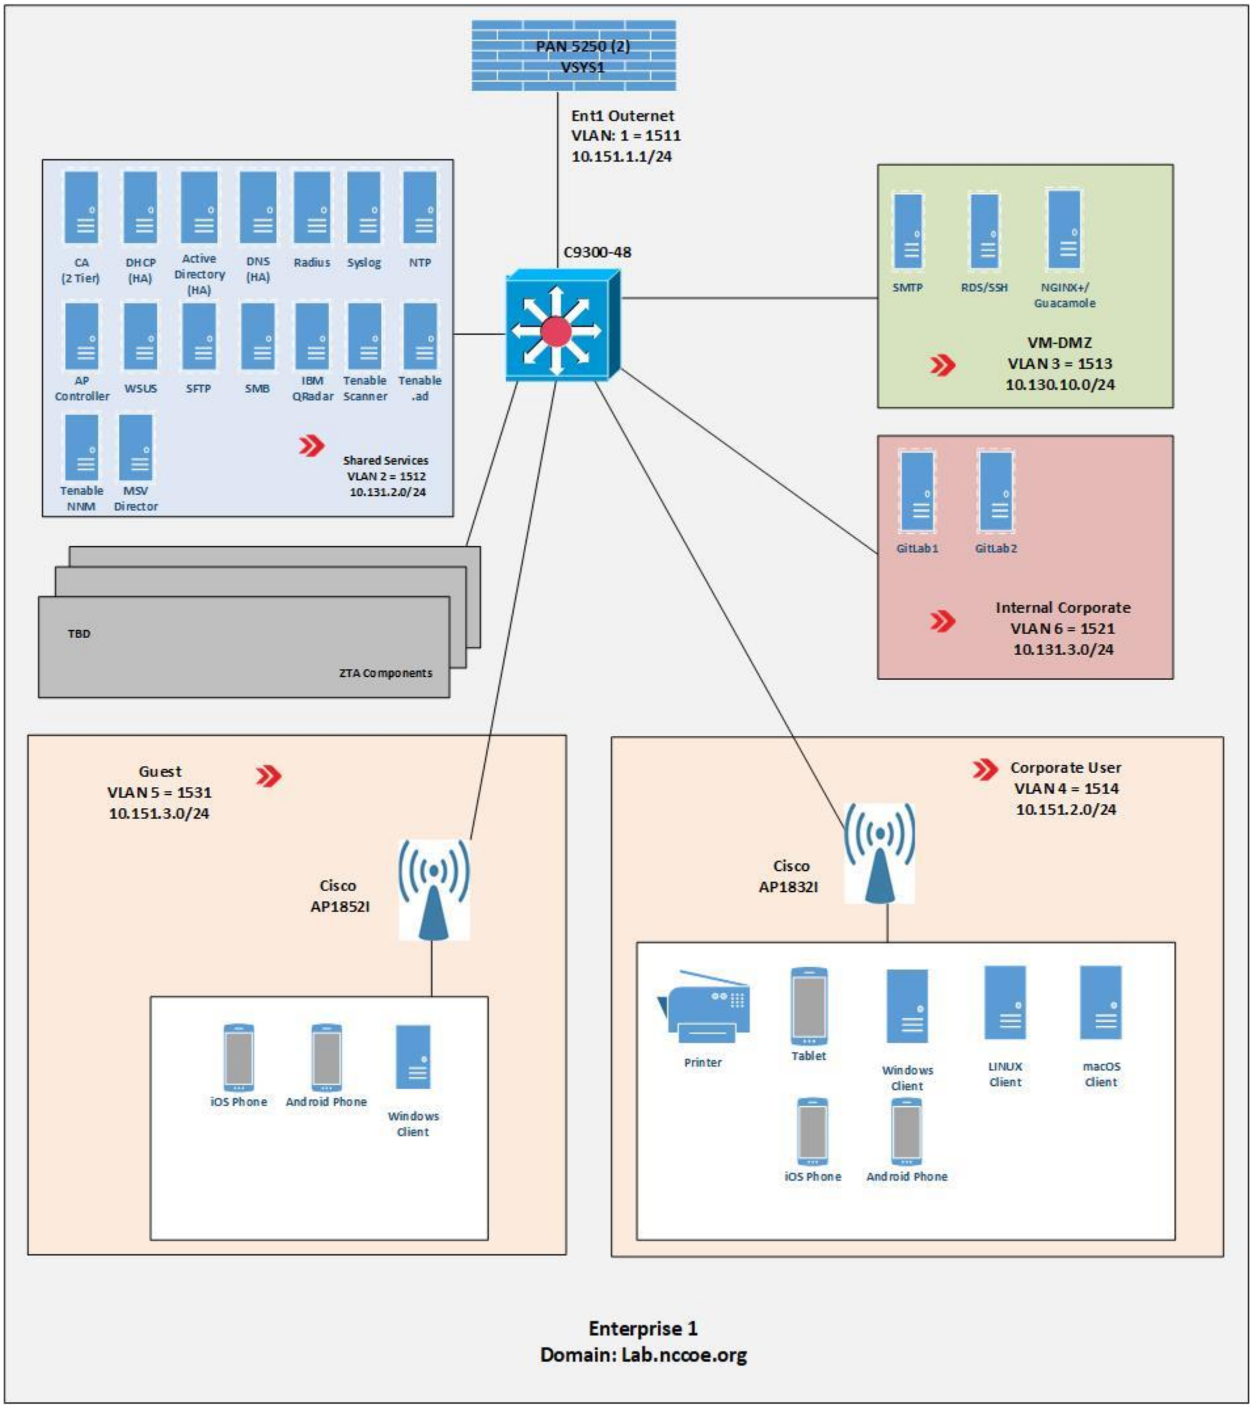

4.4.1.1 Firewall

Enterprise 1, like Enterprise 3, Enterprise 1 Branch Office, and the management and orchestration domain, is protected by a Palo Alto Networks 5250 firewall. This is one physical firewall that provides independent virtual firewalls to protect each of the above domains. Each enterprise is configured with an autonomous ZTA solution set. These virtual firewalls provide firewall and gateway capabilities, support a site-to-site Internet Protocol Security (IPsec) connection between the Enterprise 1 Branch Office and Enterprise 1, provide a remote access VPN (Global Protect) to sites, filter traffic among various internal and external subnets, provide IPv4 and IPv6 routing, and block all inbound traffic unless explicitly allowed, e.g., for communication with cloud resources. These firewalls are integrated with AD to leverage the enterprise user directory store for their respective domains.

4.4.1.2 Switch

Enterprise 1 uses a Cisco C9300 multilayer switch to provide internal network connectivity within the enterprise. It provides layer 2/3 interfaces for each virtual local area network (VLAN) subnetwork with 802.1q trunking. Both IPv4 and IPv6 addresses are assigned. This switch is integrated with the Remote Authentication Dial-In User Service (RADIUS) networking protocol to provide centralized authentication, authorization, and accounting (AAA) management for users requesting access to an Enterprise 1 network service. The switch hosts physical wireless access points and allows connections for their virtual controllers. It also provides wired access for endpoints such as laptops within the lab.

4.4.1.3 ZTA Components Specific to Enterprise 1

Enterprise 1 contains VLANs that pertain specifically to enterprise 1's ZTA build. See Appendix D for a detailed description of the ZTA components used in Enterprise 1 Build 1 (E1B1) and Appendix H for a detailed description of the ZTA components used in Enterprise 1 Build 2 (E1B2).

4.4.1.4 Demilitarized Zone (DMZ) Subnet

Enterprise 1's demilitarized zone (DMZ) is a virtual subnet that separates the rest of the Enterprise 1 network from the internet. The DMZ includes web applications and other services that Enterprise 1 makes available to users on the public internet. For example, the DMZ subnet includes Jump-box Remote Desktop Server (RDS) and Secure Shell (SSH) protocol to provide some collaborators with remote access to Enterprise 1. It also includes applications such as Simple Mail Transfer Protocol (SMTP), NGINX Plus, and Apache Guacamole.

4.4.1.5 Internal Corporate Subnet

The internal corporate subnet is where applications that support Enterprise 1's internal services reside. For example, the internal corporate subnet includes applications such as GitLab.

4.4.1.6 Corporate User Subnet

The corporate user subnet is where users and devices such as mobile devices (iOS and Android), tablets, Windows clients, macOS clients, Linux clients, and printers reside. Some of these devices are connected via wires to the C9300 switch while others are connected via Wi-Fi using the Cisco AP 18321 wireless access point.

4.4.1.7 Guest Subnet

The guest subnet is where guests reside. Guests are users who don't have any sort of network ID and are not authorized to access any enterprise resources. They use their own devices rather than corporate-owned or corporate-managed devices. Devices on the guest subnet include mobile devices, tablets, Windows clients, macOS clients, and Linux clients. The guest subnet allows for BYOD access, with all devices connecting via Wi-Fi using the Cisco AP 18321 wireless access point.

4.4.1.8 Shared Services

A closeup of the shared services domain of Enterprise 1 is depicted in Figure 4-5. The services it includes are discussed in the following subsections.

Figure 4-5 Shared Services Domain of Enterprise 1

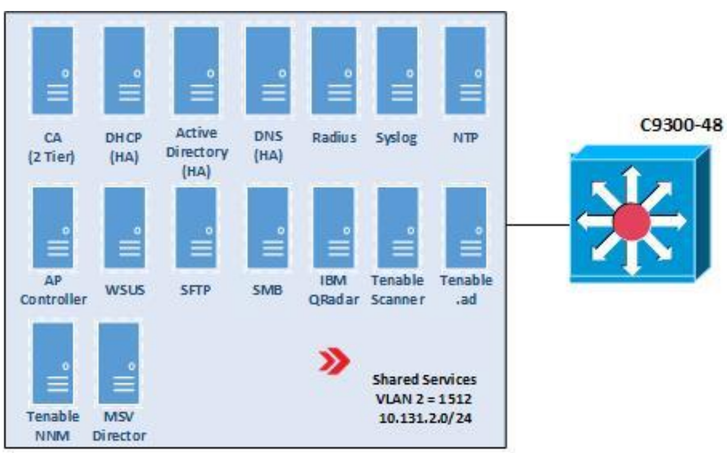

4.4.1.8.1 Certificate Authority (CA)

The CA provides certificate and cryptographic services for the enterprise. It is a Windows 2016 server using AD certificate services. A two-tier CA architecture is used, with an offline CA and an issuing AD-connected CA. The CA automatically issues and reissues certificates via AD group policy, and it can generate and issue certificates to AD domain-connected Windows devices. It issues certificates for both device authentication and web services using TLS.

4.4.1.8.2 Active Directory (AD)

AD provides centralized administration of users, computers, and resources. It runs on Windows 2016 servers and uses multiple domain controllers to ensure high availability and redundancy in hot-hot mode. It also includes a built-in DNS authoritative server and resolver.

4.4.1.8.3 Domain Name Server (DNS)

DNS provides name-to-IP address mappings for internal hosts and answers to DNS queries of external hosts. It runs on a Windows 2016 server and is the authoritative server for the lab.nccoe.org internal domain. Internal DNS services are integrated with AD. DNS servers within ITOps are used as forwarders and to resolve DNS queries from external devices. Two DNS servers are used to ensure high availability and redundancy in hot-hot mode.

4.4.1.8.4 Dynamic Host Configuration Protocol (DHCP)

The Dynamic Host Configuration Protocol (DHCP) allocates and assigns IP address and configuration information to hosts. It runs on a Windows 2016 server and is integrated with AD. Two DHCP servers are used to ensure high availability and redundancy.

4.4.1.8.5 RADIUS

The RADIUS networking protocol is used to provide centralized AAA management services at the switch for users requesting access to Enterprise 1 network services. It runs on a Windows 2016 network policy server (NPS) and is integrated with AD.

4.4.1.8.6 Access Point (AP) Controller

The access point controller manages the enterprise's wireless access points. It runs on a Cisco virtual wireless controller. It manages two APs: models 1852I and 1832I, one for the corporate user subnet and one for the guest subnet.

4.4.1.8.7 SSH File Transfer Protocol (SFTP)

SFTP is used to provide secure file transfer services. It runs on a Windows 2016 server.

4.4.1.8.8 Network Time Protocol (NTP)

NTP provides timing and clock synchronization between systems. It runs on a Windows 2019 server.

4.4.1.8.9 Syslog

Syslog is used to collect logs and diagnostic data. It runs on a Linux Ubuntu 20.04 platform.

4.4.1.8.10 Windows Server Update Service (WSUS)

Windows Server Update Service (WSUS) provides downloads and manages updates and patches for Windows servers. It runs on a Windows 2019 server.

4.4.1.8.11 Server Message Block (SMB)

Server Message Block (SMB) provides Windows file sharing services. It runs on a Windows 2019 server.

4.4.1.8.12 Collaborator Products

The shared services domain of Enterprise 1 also includes some collaborator products that provide shared services for the enterprise. The IBM QRadar, Tenable.ad, Tenable scanner, Tenable NNM, and Mandiant MSV Director are such products.

4.4.1.9 Baseline Applications

The following applications were installed and configured as part of the baseline architecture to represent the types of applications that would be found in a typical brownfield enterprise environment. These applications serve as the enterprise resources to which the ZTA is managing access.

4.4.1.9.1 Guacamole

Apache Guacamole is a remote desktop solution that supports a wide range of protocols such as SSH and Remote Desktop Protocol (RDP).

4.4.1.9.2 GitLab

GitLab is a DevOps tool that allows software developers to develop, test, and operate software in one application. We used GitLab as an enterprise application being accessed by end users.

4.4.1.9.3 NGINX Plus

NGINX Plus is free and open-source software. It is an HTTP server that can also be used as a reverse proxy and a load balancer, among other uses.

4.4.2 Enterprise 1 Branch Office

Figure 4-6 is a closeup of the high-level level physical architecture of the Enterprise 1 Branch Office in the NCCoE laboratory environment. The Enterprise 1 Branch Office has three main components: a firewall, a switch, and a subnet for corporate users.

Figure 4-6 Physical Architecture of the Enterprise 1 Branch Office

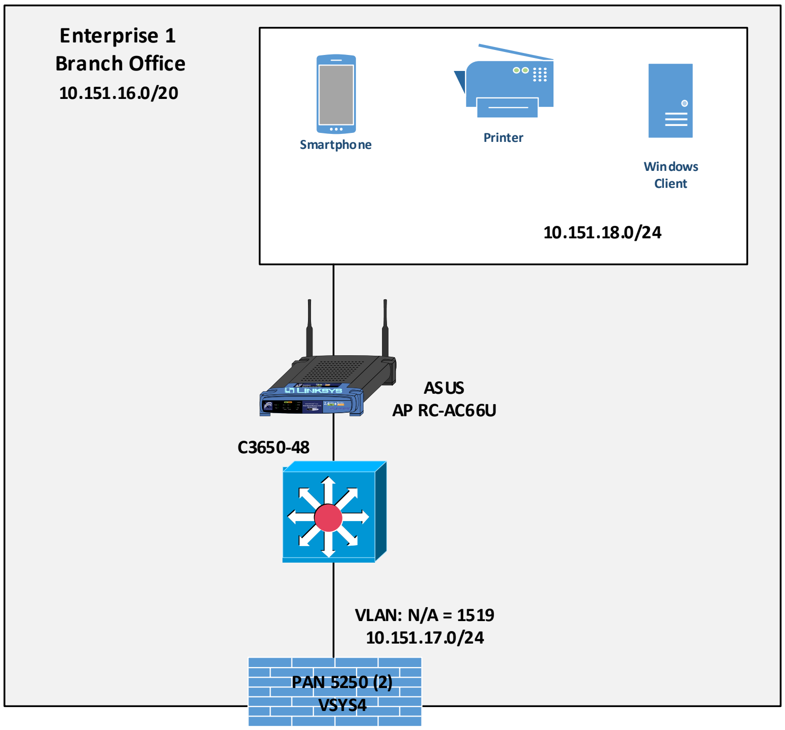

4.4.2.1 Firewall

One of the independent virtual firewalls provided by the Palo Alto Networks 5250 physical firewall is used for the Enterprise 1 Branch Office. It provides firewall and gateway capabilities, connecting the Branch Office to Enterprise 1 via the emulated WAN/internet service provider and supports a site-to-site VPN IPsec connection from the Branch Office to Enterprise 1. This firewall is integrated with the AD of Enterprise 1 so it can leverage Enterprise 1's user directory store.

4.4.2.2 Switch

The Branch Office includes a Cisco C3650 multilayer switch that provides internal network connectivity within the Branch Office. It is integrated with Enterprise 1's AAA (RADIUS) server to leverage Enterprise 1's authentication and authorization services.

4.4.2.3 Corporate Users Subnet

The corporate users subnet at the Branch Office is where users and devices such as mobile devices, tablets, Windows clients, and printers reside. Some of these devices are connected via wires to the Cisco 3650 switch while others are connected via Wi-Fi using an ASUS RC-AC66U wireless access point.

4.4.3 Enterprise 2

The high-level physical architecture of Enterprise 2 is the same as that of Enterprise 1, except Enterprise 2 does not have an associated branch office. The baseline network topology, hardware, and software of Enterprise 2 is configured the same as Enterprise 1's. Enterprise 2 leverages the same setup as Enterprise 1 using the Palo Alto Networks NGFW and Cisco switches. It also includes the same setup and capabilities as Enterprise 1 with respect to its DMZ, internal corporate subnetwork, corporate user subnetwork, guest subnetwork, shared services, and baseline applications. The only differences between Enterprise 2 and Enterprise 1 are with respect to the on-premises and cloud-based ZTA components used in each enterprise. See Appendix E for a detailed description of the ZTA components used in Enterprise 2.

4.4.4 Enterprise 3

The high-level physical architecture of Enterprise 3 is the same as that of Enterprise 2. The only differences between Enterprise 3 and Enterprise 2 are with respect to the on-premises and cloud-based ZTA components used in each enterprise. See Appendix E for a detailed description of the ZTA components used in Enterprise 3 Build 1 (E3B1) and Appendix J for a detailed description of the ZTA components used in Enterprise 3 Build 2 (E3B2).

4.4.5 Enterprise 4

Enterprise 4 is not yet being used in this phase of the project.

4.4.6 Coffee Shop

Figure 4-7 is a closeup of the high-level level physical architecture of the coffee shop in the NCCoE laboratory environment. As shown, the coffee shop provides users and mobile devices (e.g., smartphones and laptops) wireless access to the internet via an ASUS RC-AC66U access point.

Figure 4-7 Physical Architecture of the Coffee Shop

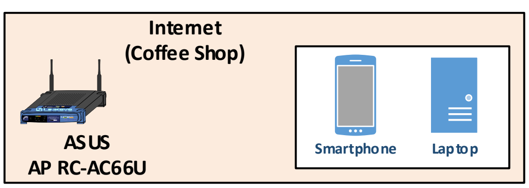

4.4.7 Management and Orchestration Domain

The management and orchestration domain, as depicted in Figure 4-8, includes components that support infrastructure as code (IaC) automation and orchestration across the ZTA lab environment. It includes Terraform, which is used to automate the setup of VMs across the four enterprises, and

2452 Ansible, which automates the setup of VMs and services such as DHCP, DNS, and AD across all four
2453 enterprises. It also hosts the Mandiant MSV Director and the MSV Protected Theater.

2454 **Figure 4-8 Physical Architecture of the Management and Orchestration Domain**

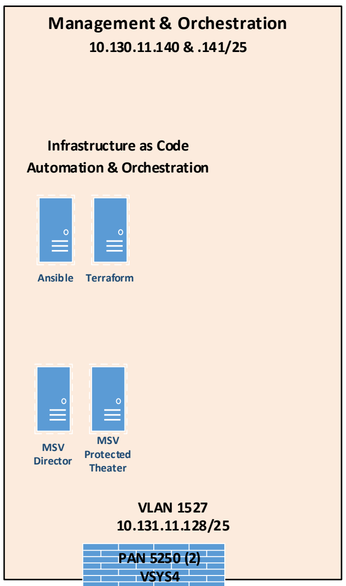

2455 4.4.8 Emulated WAN Service Provider

2456 A subnetwork within the ZTA laboratory network is leveraged to emulate a WAN service provider. The
2457 emulated WAN service provider using a Cisco SG550X switch and a Palo Alto 5250 NGFW provides
2458 connectivity among all the ZTA laboratory network domains, i.e., the enterprises, the coffee shop, the
2459 branch office, and the management and orchestration domain. It also connects the ZTA laboratory
2460 network to ITOPS, which provides connectivity to the internet. Via the internet, the emulated WAN
2461 services provide the ZTA lab network with connectivity to cloud services.

2462 4.4.9 Cloud Services

2463 As mentioned, the NCCoE lab environment has access to various cloud services via the internet. The
2464 cloud services that have been set up during the EIG crawl phase are described in Section 4.4.9.1. Cloud
2465 services will be used as part of the EIG run phase.

2466 *4.4.9.1 IaaS – Amazon Web Services (AWS)*

2467 Figure 4-9 depicts the physical architecture of the AWS infrastructure that has been set up for use by
2468 Enterprise 1. As shown, the NCCoE ZTA lab is connected to AWS via a site-to-site VPN, and work is
2469 underway to set up a direct connection between the NCCoE ZTA lab and AWS as well. Both a production
2470 VPC (labeled Ent 1 Prod VPC) and a management VPC (labeled Ent 1 Mgmt VPC) have been set up within

AWS for Enterprise 1 to use. There is a transit gateway (TGW) for routing traffic between the production and management VPCs, and there is also an NCCoE TGW within AWS. CloudFormation was used to set up the production and management VPC infrastructure within AWS through the NCCoE and Enterprise TGWs. The TGW acts as a hub for routing traffic between production and management VPCs and includes multiple routing tables for secure routing between the VPCs.

2476 Figure 4-9 Physical Architecture of the AWS Infrastructure Used by Enterprise 1

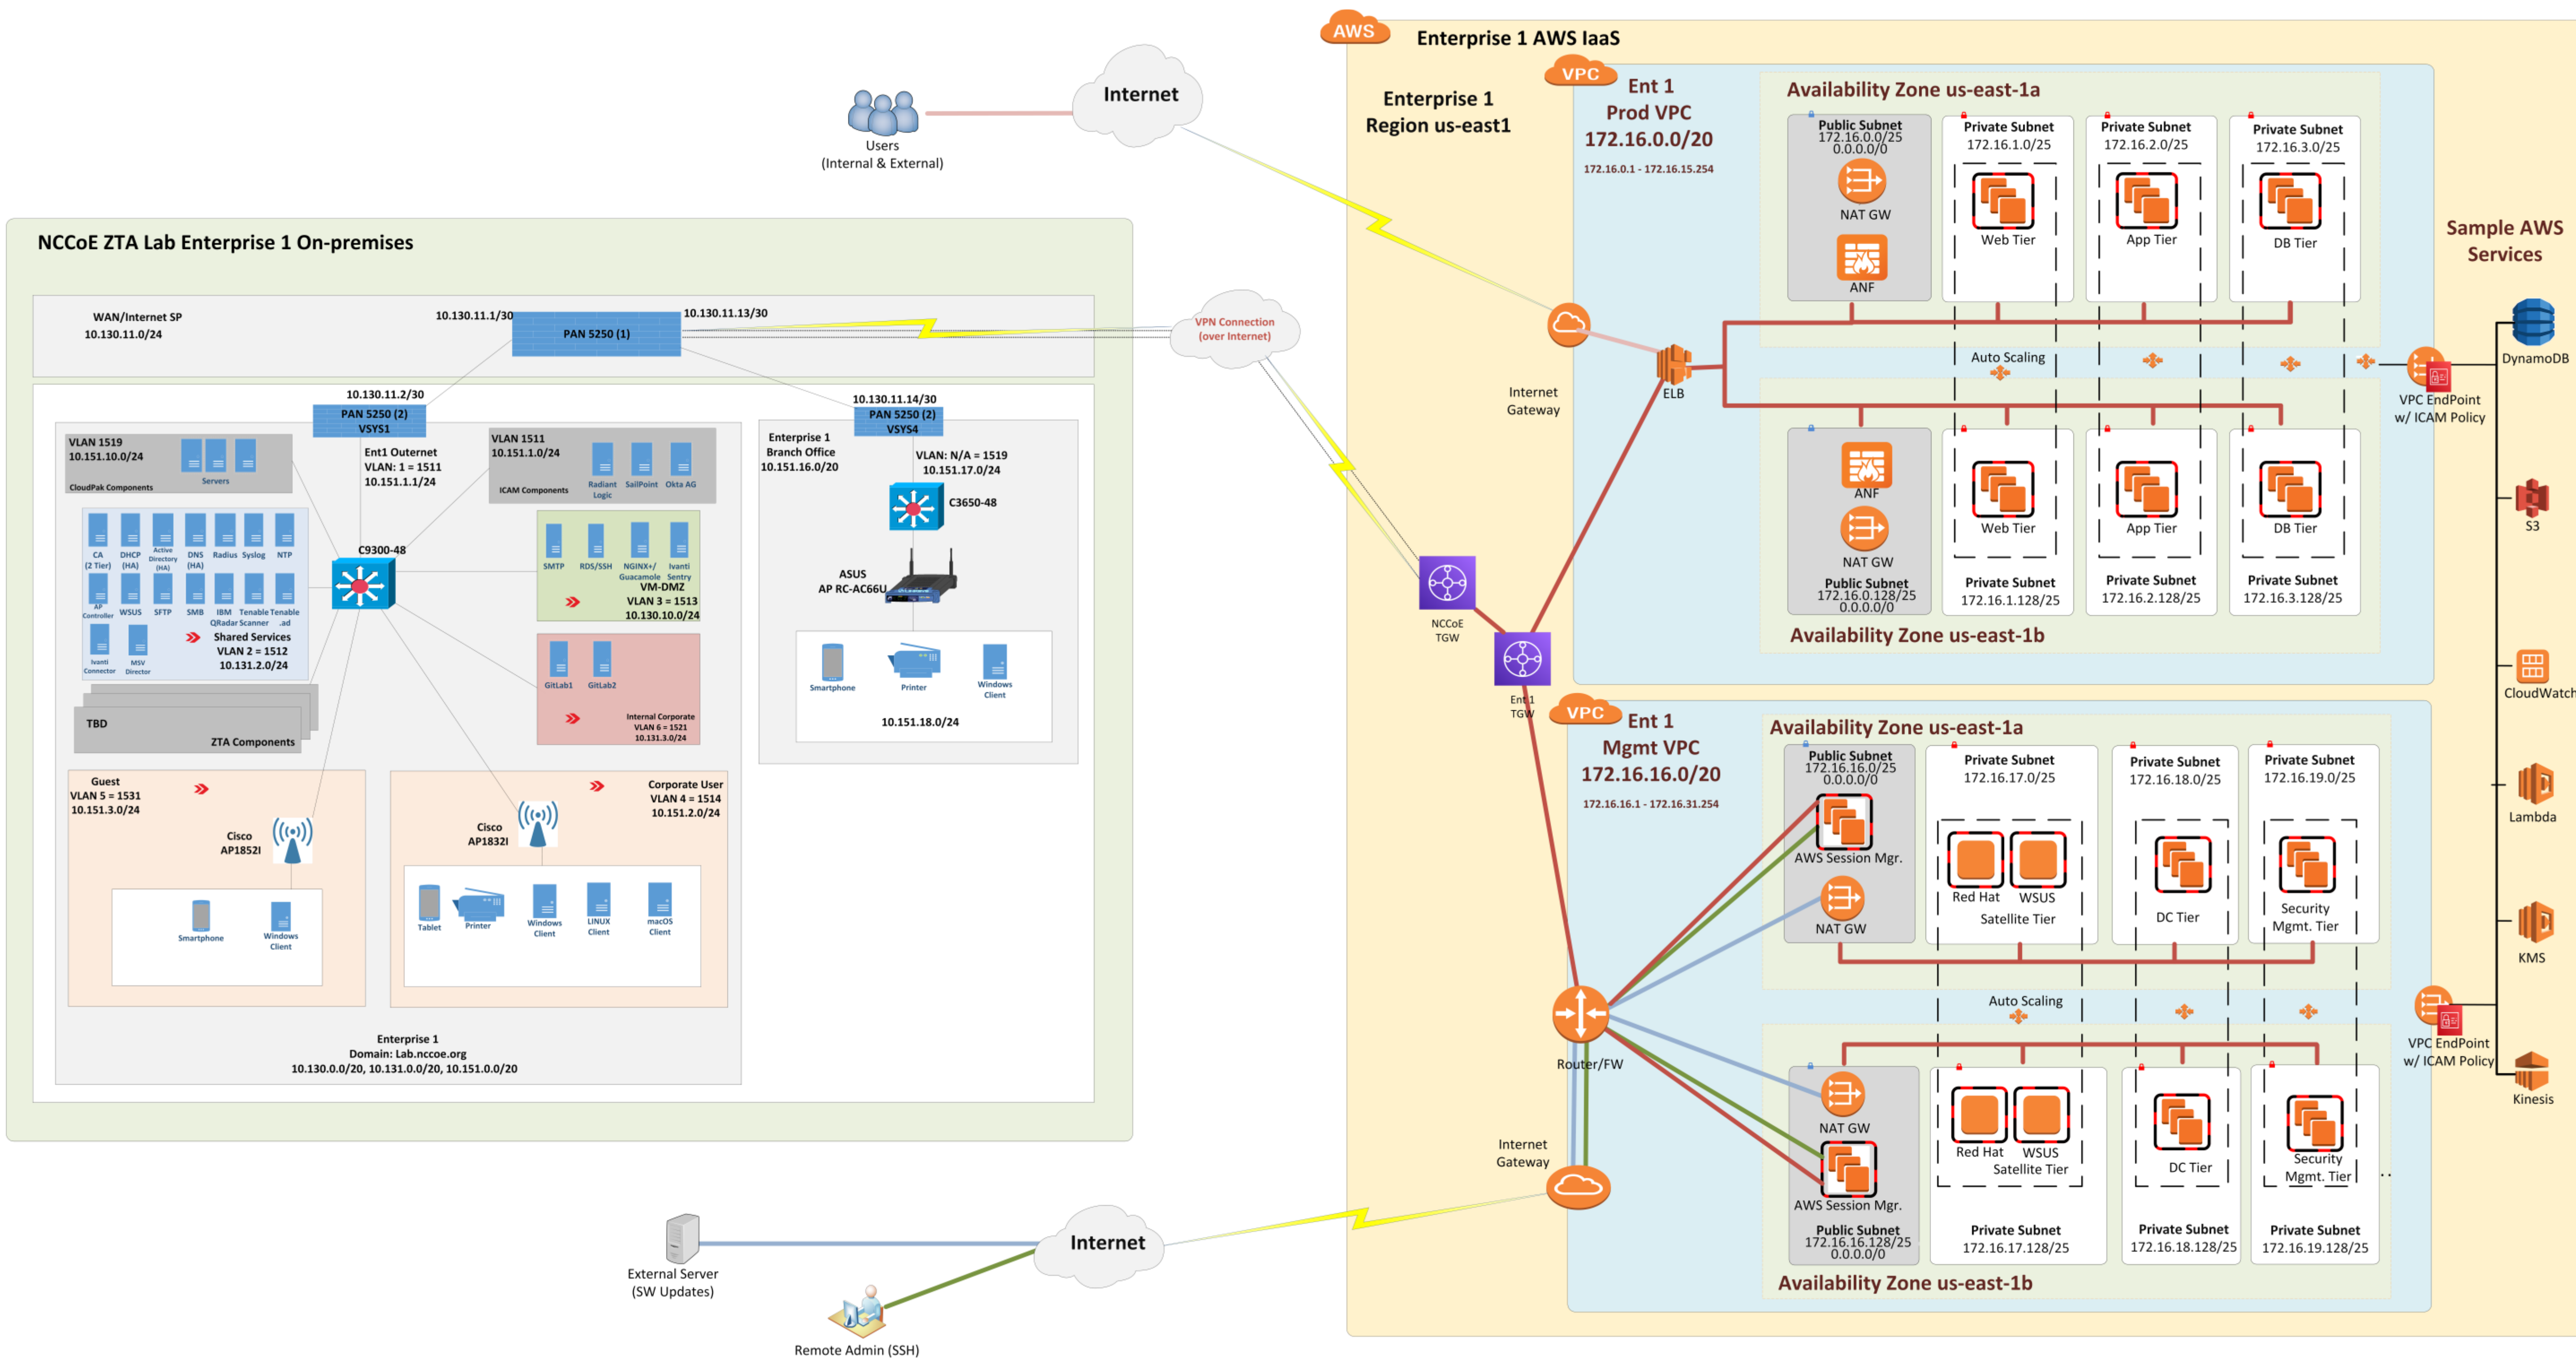

The production VPC has both a public subnetwork and three private subnetworks in each availability zone. The public subnetwork is used for connecting external users to the production VPC. The private subnetworks have EC2s that can host web, application, and database tiers.

The management VPC also has a public subnetwork and three private subnetworks in each availability zone. The public subnetwork is used to support software updates and to enable administrators and other authorized internal staff who are located remotely to SSH into cloud components. The private subnetworks include a satellite tier, domain controller tier, and security management tier.

Each VPC uses two availability zones for redundancy and high availability. Each availability zone uses automatic scaling as needed.

4.4.9.2 IaaS – Google

The NCCoE staff is currently working with its collaborators to set up a cloud environment for Enterprise 2.

4.4.9.3 IaaS – Azure

Figure 4-10 depicts the physical architecture of the Azure IaaS that has been set up for use by Enterprise 3. As shown, the NCCoE ZTA lab is connected to Azure IaaS via a site-to-site VPN. If coming from on-premises through the site-to-site VPN into Azure IaaS, connections go through the hub virtual network before getting to the application virtual networks for both the public-facing and private applications. The hub virtual network consists of the gateway subnet, the firewall subnet, and the bastion subnet. The gateway subnet consists of virtual network gateways in multiple availability zones. The firewall subnet consists of firewalls in multiple availability zones. The bastion subnet consists of Azure Bastion in multiple availability zones.

The public application virtual network consists of a gateway subnet, an application subnet, and utility, web, and data subnets. Each of these subnets is secured by network security group (NSG). The gateway subnet consists of application gateways in multiple availability zones and web application firewall (WAF) Policies. The application subnet hosts the virtual machines and the applications, all of which are secured by application security groups.

The private application virtual network consists of a gateway subnet and an application subnet. Each of these subnets is secured by NSG. The gateway subnet consists of application gateways in multiple availability zones and web application firewall (WAF) policies. The application subnet hosts the virtual machines and the applications, as well as application proxies, all of which are secured by application security groups. The application proxies are meant to be used by remote users connecting to private applications through the internet.

Traffic between subnets is allowed only if the NSGs on the subnets allow it. With the zero trust design, traffic between subnets should not be assumed; it must be explicitly granted.

2511 **Figure 4-10 Physical Architecture of the Azure Infrastructure Used by Enterprise 3**

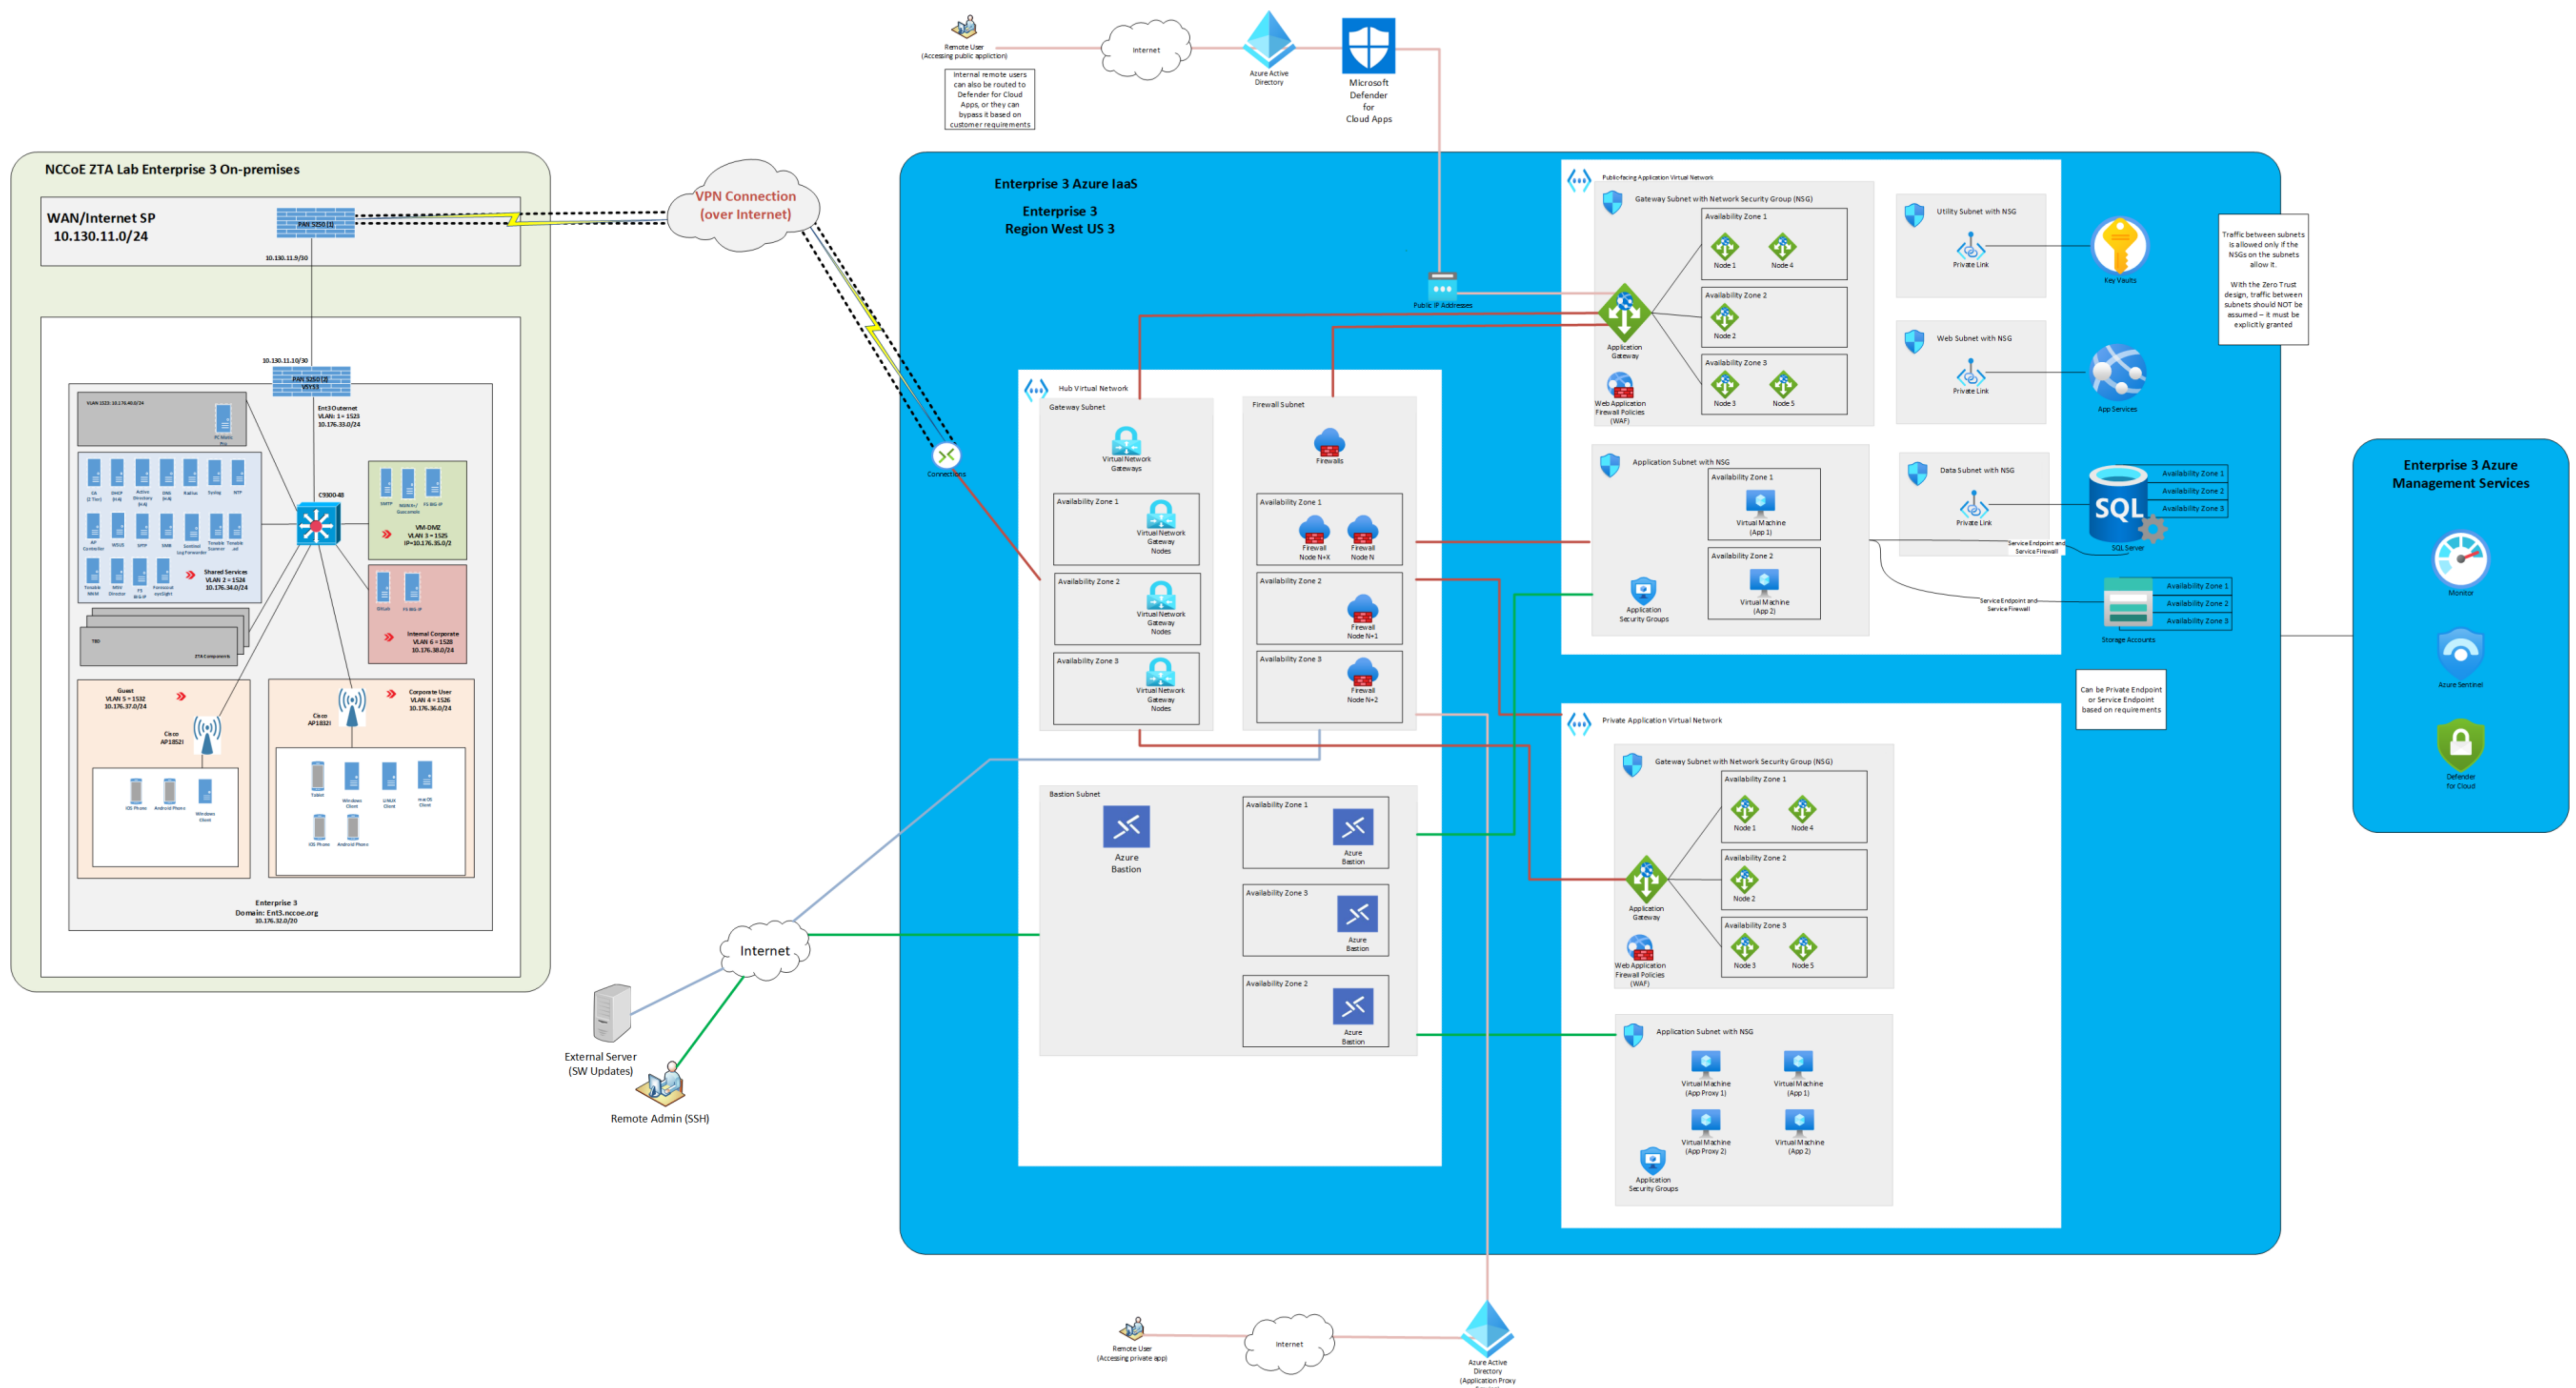

4.4.9.4 SaaS

The project is also using collaborators' ZTA SaaS offerings. The SaaS-based ZTA products used are listed in the appendices describing each build.

5 Functional Demonstration

Functional demonstrations were performed to showcase the security characteristics supported by each ZTA build. These demonstrations show the extent to which the example solutions meet their security objectives under a variety of conditions. NIST SP 1800-35D, *ZTA Functional Demonstrations* will document each of the demonstration scenarios and use cases that have been designed for this ZTA project. The results of the demonstrations that have been conducted on each ZTA build will also be listed in NIST SP 1800-35D.

6 General Findings

When deploying ZTA using the EIG approach, 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

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

6.1 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 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 ZT 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 4-1. 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 include 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.

6.2 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 have 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 capabilities 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 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. The inclusion of an EPP is planned for a later build phase.

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.

7 Future Build Considerations

At this time, three EIG crawl phase builds are complete (E1B1, E2B1, and E3B1). We are skipping the EIG walk phase and have proceeded directly to the run phase. Two EIG run phase builds, Enterprise 1 (E1B2) and Enterprise 3 (E3B2) are also complete. All five of these builds are documented in this guide.

The next phase of the project will focus on the micro-segmentation and SDP deployment models, and a combination of the two. Efforts will be organized into crawl, walk, and run phases that augment the EIG capabilities to support an increasingly rich set of functionalities and additional ZTA capabilities.

Appendix A List of Acronyms

AAA	Authentication, Authorization, and Accounting
ACL	Access Control List
AD	Active Directory
AI	Artificial Intelligence
API	Application Programming Interface
APM	(F5 BIG-IP) Access Policy Manager
ATP	(Microsoft Azure) Advanced Threat Protection, (Palo Alto Networks) Advanced Threat Prevention
AURL	(Palo Alto Networks) Advanced URL Filtering
AWS	Amazon Web Services
BCE	(Google) BeyondCorp Enterprise
BYOD	Bring Your Own Device
C&C	Command-and-Control
CA	Certificate Authority, (Zscaler) Central Authority
CASB	Cloud Access Security Broker
CDM	Continuous Diagnostics and Mitigation
CDSS	Cloud-Delivered Security Service
CESA	Cisco Endpoint Security Analytics
CI/CD	Continuous Integration/Continuous Delivery
CIEM	Cloud Infrastructure Entitlement Management
CISA	Cybersecurity and Infrastructure Security Agency
CRADA	Cooperative Research and Development Agreement
CVE	Common Vulnerabilities and Exposures
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DLP	Data Loss Prevention
DMZ	Demilitarized Zone
DNS	Domain Name System
DTLS	Datagram Transport Layer Security
EBS	(Amazon) Elastic Block Store
EC2	(Amazon) Elastic Compute Cloud

ECS	(Amazon) Elastic Container Service
EDR	Endpoint Detection and Response
EIG	Enhanced Identity Governance
EKS	(Amazon) Elastic Kubernetes Service
EMM	Enterprise Mobility Management
EO	Executive Order
ePO	(Trellix) ePolicy Orchestrator
EPP	Endpoint Protection Platform
ETA	(Cisco) Encrypted Traffic Analytics
E/W	East/West
FedRAMP	Federal Risk and Authorization Management Program
FIDO U2F	Fast Identity Online Universal 2nd Factor
FIPS	Federal Information Processing Standards
FTD	(Cisco) Firepower Threat Defense
FWaaS	Firewall as a Service
GCP	Google Cloud Platform
GDPR	General Data Protection Regulation
GIN	(Symantec) Global Intelligence Network
GP	(Palo Alto Networks) GlobalProtect
HR	Human Resources
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IaaS	Infrastructure as a Service
IaC	Infrastructure as Code
IAM	Identity and Access Management
IBM	International Business Machines Corporation
ICA	Intermediate Certificate Authority
ICAM	Identity, Credential, and Access Management
IDaaS	Identity as a Service
IGA	(Symantec) Identity Governance and Administration
IoMT	Internet of Medical Things
IoT	Internet of Things

IP	Internet Protocol
IPsec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol Version 6
ISE	(Cisco) Identity Services Engine
IT	Information Technology
ITL	Information Technology Lab
ITOps	Information Technologies Operations
KCD	Kerberos Constrained Delegation
LDAP	Lightweight Directory Access Protocol
LTM	(F5 BIG-IP) Local Traffic Manager
MAM	Mobile Application Management
MDM	Mobile Device Management
MES	(Lookout) Mobile Endpoint Security
MFA	Multi-Factor Authentication
ML	Machine Learning
MSV	Mandiant Security Validation
MTD	Mobile Threat Defense
mTLS	Mutual Transport Layer Security
NCCoE	National Cybersecurity Center of Excellence
NDR	Network Detection and Response
NGFW	Next-Generation Firewall
NIST	National Institute of Standards and Technology
NMM	(Tenable) Nessus Network Monitor
NPE	Non-Person Entity
NPS	Network Policy Server
N/S	North/South
NSG	Network Security Group
NTA	Network Traffic Analysis
NTP	Network Time Protocol
NVM	(Cisco) Network Visibility Module
OIDC	OpenID Connect

OMB	Office of Management and Budget
OT	Operational Technology
OTP	One-Time Password
PA	Policy Administrator
PAN	Palo Alto Networks
PDP	Policy Decision Point
PE	Policy Engine
PEP	Policy Enforcement Point
PII	Personally Identifiable Information
PIN	Personal Identification Number
PIP	Policy Information Point
PKI	Public Key Infrastructure
QOS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
R&D	Research and Development
RDP	Remote Desktop Protocol
RDS	Remote Desktop Server
REST	Representational State Transfer
S3	(Amazon) Simple Storage Service
SaaS	Software as a Service
SAML	Security Assertion Markup Language
SASE	Secure Access Service Edge
SAW	(Microsoft) Secure Admin Workstation
SCIM	System for Cross-Domain Identity Management
SDLC	Software Development Lifecycle
SDP	Software-Defined Perimeter
SD-WAN	Software-Defined Wide Area Network
SFTP	SSH File Transfer Protocol
SIEM	Security Information and Event Management
SMB	Server Message Block
SMS	Short Message Service
SMTP	Simple Mail Transfer Protocol

SOAR	Security Orchestration and Response
SoD	Separation of Duties
SP	Special Publication
SQL	Structured Query Language
SRE	Site Reliability Engineer
SSE	(Skyhigh Security) Security Service Edge
SSH	Secure Shell
SSL	Secure Sockets Layer
SSO	Single Sign-On
SWG	Secure Web Gateway
TGW	Transit Gateway
TLS	Transport Layer Security
TTP	Tactics, Techniques, and Procedures
UEM	Unified Endpoint Management
URL	Uniform Resource Locator
USB	Universal Serial Bus
VDI	Virtual Desktop Infrastructure
VIP	(Symantec) Validation and ID Protection
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPC	Virtual Private Cloud
VPN	Virtual Private Network
WAF	Web Application Firewall
WF	(Palo Alto Networks) Wildfire
WSS	(Symantec) Web Security Service
WSUS	(Microsoft) Windows Server Update Service
XDR	Extended Detection and Response
ZCC	Zscaler Client Connector
ZIA	Zscaler Internet Access
ZPA	Zscaler Private Access
ZSO	(Ivanti) Zero Sign-On
ZT	Zero Trust

ZTA	Zero Trust Architecture
ZTNA	Zero Trust Network Access

2619 Appendix B Glossary

Managed Devices	Personal computers, laptops, mobile devices, virtual machines, and infrastructure components require management agents, allowing information technology staff to discover, maintain, and control them. Those with broken or missing agents cannot be seen or managed by agent-based security products. [NIST SP 1800-15 Vol. B]
Policy	Statements, rules, or assertions that specify the correct or expected behavior of an entity. For example, an authorization policy might specify the correct access control rules for a software component. [NIST SP 800-95 and NIST IR 7621 Rev. 1]
Policy Administrator (PA)	An access control mechanism component that executes the PE's policy decision by sending commands to the PEP to establish and terminate the communications path between the subject and the resource.
Policy Decision Point (PDP)	An access control mechanism component that computes access decisions by evaluating the applicable policies. The functions of the PE and PA comprise a PDP. [NIST SP 800-162, adapted]
Policy Enforcement Point (PEP)	An access control mechanism component that enforces access policy decisions in response to a request from a subject requesting access to a protected resource. [NIST SP 800-162, adapted]
Policy Engine (PE)	An access control mechanism component that handles the ultimate decision to grant, deny, or revoke access to a resource for a given subject.
Policy Information Point (PIP)	An access control mechanism component that provides telemetry and other information generated by policy or collected by supporting components that the PDP needs for making policy decisions. [NIST SP 800-162, adapted]
Risk	The net negative impact of the exercise of a vulnerability, considering both the probability and the impact of occurrence. [NIST SP 1800-15 Vol. B]
Security Control	A safeguard or countermeasure prescribed for an information system or an organization designed to protect the confidentiality, integrity, and availability of its information and to meet a set of defined security requirements. [NIST SP 800-53 Rev. 5]
Threat	Any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, or individuals through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service. Also, the potential for a threat-source to successfully

	exploit a particular information system vulnerability. [Federal Information Processing Standards 200]
Vulnerability	Weakness in an information system, system security procedures, internal controls, or implementation that could be exploited or triggered by a threat source. [NIST SP 800-37 Rev. 2]
Zero Trust	A cybersecurity paradigm focused on resource protection and the premise that trust is never granted implicitly but must be continually evaluated. [NIST SP 800-207]
Zero Trust Architecture (ZTA)	An enterprise cybersecurity architecture that is based on zero trust principles and designed to prevent data breaches and limit internal lateral movement. Zero trust architecture is an end-to-end approach to enterprise resource and data security that encompasses identity (person and non-person entities), credentials, access management, operations, endpoints, hosting environments, and the interconnecting infrastructure. [NIST SP 800-207]

Appendix C References

[1] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, *Zero Trust Architecture*, National Institute of Standards and Technology (NIST) Special Publication (SP) 800-207, Gaithersburg, Md., August 2020, 50 pp. Available: https://csrc.nist.gov/publications/detail/sp/800-207/final.

[2] Executive Order no. 14028, *Improving the Nation's Cybersecurity*, Federal Register Vol. 86, No.93, May 17, 2021. Available: https://www.federalregister.gov/documents/2021/05/17/2021-10460/improving-the-nations-cybersecurity.

[3] "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.

[4] https://www.nccoe.nist.gov/iot

[5] https://www.nccoe.nist.gov/manufacturing

[6] https://www.nccoe.nist.gov/energy

[7] https://www.nccoe.nist.gov/healthcare

[8] https://www.f5.com/company/certifications

[9] https://csrc.nist.gov/projects/cryptographic-module-validation-program/Certificate/3344

[10] https://csrc.nist.gov/Projects/cryptographic-module-validation-program/Certificate/3452

[11] P. Grassi, J. Richer, S. Squire, J. Fenton, E. Nadeau, N. Lefkovitz, J. Danker, Y. Choong, K. Greene, and M. Theofanos, *Digital Identity Guidelines Federation and Assertions,* National Institute of Standards and Technology (NIST) Special Publication (SP) 800-63C, Gaithersburg, Md., June 2017, 40 pp. Available: https://www.nist.gov/identity-access-management/nist-special-publication-800-63-digital-identity-guidelines.

Appendix D EIG Enterprise 1 Build 1 (E1B1)

D.1 Technologies

EIG E1B1 uses products from Amazon Web Services, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zimperium. Certificates from DigiCert are also used. For more information on these collaborators and the products and technologies that they contributed to this project overall, see Section 3.4.

E1B1 components consist of Okta Identity Cloud, Ivanti Access ZSO, Ivanti Sentry, Radiant Logic RadiantOne Intelligent Identity Data Platform, SailPoint IdentityIQ, Okta Verify App, Ivanti Neurons for UEM, Zimperium MTD, IBM Security QRadar XDR, Tenable.io, Tenable.ad, IBM Cloud Pak for Security, Mandiant Security Validation (MSV), Ivanti Tunnel, DigiCert CertCentral, and AWS IaaS.

Table D-1 lists all of the technologies used in EIG E1B1. It lists the products used to instantiate each ZTA component and the security function that each component provides.

Table D-1 E1B1 Products and Technologies

Component	Product	Function
PE	Okta Identity Cloud and Ivanti Access ZSO	Decides whether to grant, deny, or revoke access to a resource based on enterprise policy, information from supporting components, and a trust algorithm.
PA	Okta Identity Cloud and Ivanti Access ZSO	Executes the PE’s policy decision by sending commands to a PEP that establishes and shuts down the communication path between subject and resource.
PEP	Ivanti Sentry	Guards the trust zone that hosts one or more enterprise resources; establishes, monitors, and terminates the connection between subject and resource as directed by the PA; forwards requests to and receives commands from the PA.
Identity Management	Okta Identity Cloud	Creates and manages enterprise user and device accounts, identity records, role information, and access attributes that form the basis of access decisions within an organization to ensure the correct subjects have the appropriate access to the correct resources at the appropriate time.
Access & Credential Management	Okta Identity Cloud	Manages access to resources by performing user and device authentication (e.g., SSO and MFA) and using identity, role, and access attributes to determine which access requests are authorized.

Component	Product	Function
Federated Identity	Radiant Logic RadiantOne Intelligent Identity Data Platform	Aggregates and correlates all attributes relating to an identity or object that is being authorized by a ZTA. It enables users of one domain to securely access data or systems of another domain seamlessly, and without the need for completely redundant user administration. Federated identity encompasses the traditional ICAM data, supports identities that may be part of a larger federated ICAM community, and may include non-enterprise employees.
Identity Governance	SailPoint IdentityIQ	Provides policy-based, centralized, automated processes to manage user identity and access control functions (e.g., ensuring segregation of duties, role management, logging, access reviews, analytics, reporting) to ensure compliance with requirements and regulations.
MFA	Okta Verify app	Supports MFA of a user identity by requiring the user to provide not only something they know (e.g., a password), but also something they have (e.g., a token).
UEM/MDM	Ivanti Neurons for Unified Endpoint Management (UEM) Platform	Manages and secures enterprise desktop computers, laptops, and/or mobile devices in accordance with enterprise policy to protect applications and data; ensure device compliance; mitigate and remediate vulnerabilities and threats; monitor for suspicious activity to prevent and detect intrusions; prevent, detect, and disable malware and other malicious or unauthorized traffic; repair infected files when possible; provide alerts and recommend remediation actions; and encrypt data. Pushes enterprise applications and updates to devices, enables users to download enterprise applications that they are authorized to access, remotely deletes all applications and data from devices if needed, tracks user activity on devices, and detects and addresses security issues on the device.

Component	Product	Function
EPP	Zimperium MTD	Detects and stops threats to endpoints through an integrated suite of endpoint protection technologies including antivirus, data encryption, intrusion prevention, EDR, and DLP. May include mechanisms that are designed to protect applications and data; ensure device compliance with policies regarding hardware, firmware, software, and configuration; monitor endpoints for vulnerabilities, suspicious activity, intrusion, infection, and malware; block unauthorized traffic; disable malware and repair infections; manage and administer software and updates; monitor behavior and critical data; and enable endpoints to be tracked, troubleshooted, and wiped, if necessary.
SIEM	IBM Security QRadar XDR	Collects and consolidates security information and security event data from many sources; correlates and analyzes the data to help detect anomalies and recognize potential threats and vulnerabilities; and logs the data to adhere to data compliance requirements.
Vulnerability Scanning and Assessment	Tenable.io and Tenable.ad	Scans and assesses the enterprise infrastructure and resources for security risks, identifies vulnerabilities and misconfigurations, and provides remediation guidance regarding investigating and prioritizing responses to incidents.
Security Integration Platform	IBM Cloud Pak for Security	Integrates the SIEM and other security tools into a single pane of glass to support generation of insights into threats and help track, manage, and resolve cybersecurity incidents. Executes predefined incident response workflows to automatically analyze information and orchestrate the operations required to respond.

Component	Product	Function
Security Validation	Mandiant MSV	Provides visibility and evidence on the status of the security controls' effectiveness in the ZTA. Enables security capabilities of the enterprise to be monitored and verified by continuously validating and measuring the cybersecurity controls; also used to automate the demonstrations that were performed to showcase ZTA capabilities. Deployed throughout the project's laboratory environment to enable monitoring and verification of various security aspects of the builds. 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.
Remote Connectivity	Ivanti Tunnel	Enables authorized remote users to securely access the inside of the enterprise. (Once inside, the ZTA manages the user's access to resources.)
Certificate Management	DigiCert CertCentral TLS Manager	Provides automated capabilities to issue, install, inspect, revoke, renew, and otherwise manage TLS certificates.
Cloud IaaS	AWS - GitLab, WordPress	Provides computing resources, complemented by storage and networking capabilities, hosted by a cloud service provider, offered to customers on demand, and exposed through a GUI and an API.
Cloud SaaS	Digicert CertCentral, Ivanti Access ZSO, Ivanti Neurons for UEM, Okta Identity Cloud, and Tenable.io, and Zimperium MTD	Cloud-based software delivered for use by the enterprise.
Application	GitLab	Example enterprise resource to be protected. (In this build, GitLab is integrated with Okta using SAML, and IBM Security QRadar XDR pulls logs from GitLab.)
Enterprise-Managed Device	Mobile devices (iOS and Android)	Example endpoints to be protected. All enterprise-managed devices are running an Ivanti Neurons for UEM agent and also have the Okta Verify App installed.
BYOD	Mobile devices (iOS and Android)	Example endpoints to be protected.

D.2 Build Architecture

In this section we present the logical architecture of E1B1 relative to how it instantiates the EIG crawl phase reference architecture depicted in Figure 4-2. We also describe E1B1's physical architecture and present message flow diagrams for some of its processes.

D.2.1 Logical Architecture

Figure D-1 depicts the logical architecture of E1B1. Figure D-1 uses numbered arrows to depict the general flow of messages needed for a subject to request access to a resource and have that access request evaluated based on subject identity (both requesting user and requesting endpoint identity), user authorizations, and requesting endpoint health. It also depicts the flow of messages supporting periodic reauthentication of the requesting user and the requesting endpoint and periodic verification of requesting endpoint health, all of which must be performed to continually reevaluate access. The labeled steps in Figure D-1 have the same meanings as they do in Figure 4-1 and Figure 4-2. However, while Figure 4-2 depicts generic EIG crawl phase ZTA components, Figure D-1 includes the specific products that instantiate the architecture of E1B1. Figure D-1 also does not depict any of the resource management steps found in Figure 4-1 and Figure 4-2 because the ZTA technologies deployed in E1B1 do not support the ability to perform authentication and reauthentication of the resource or periodic verification of resource health.

E1B1 was designed with a single ICAM system (Okta Identity Cloud) that serves as the identity, access, and credential manager as well as the ZTA PE and PA. It includes the Ivanti Sentry as its PEP, and it also delegates some PDP responsibilities to Ivanti Access ZSO. Radiant Logic acts as a PIP for the PDP as it responds to inquiries and provides identity information on demand in order for Okta to make near-real-time access decisions. A more detailed depiction of the messages that flow among components to support a user access request can be found in Appendix D.2.4.

Figure D-1 Logical Architecture of E1B1

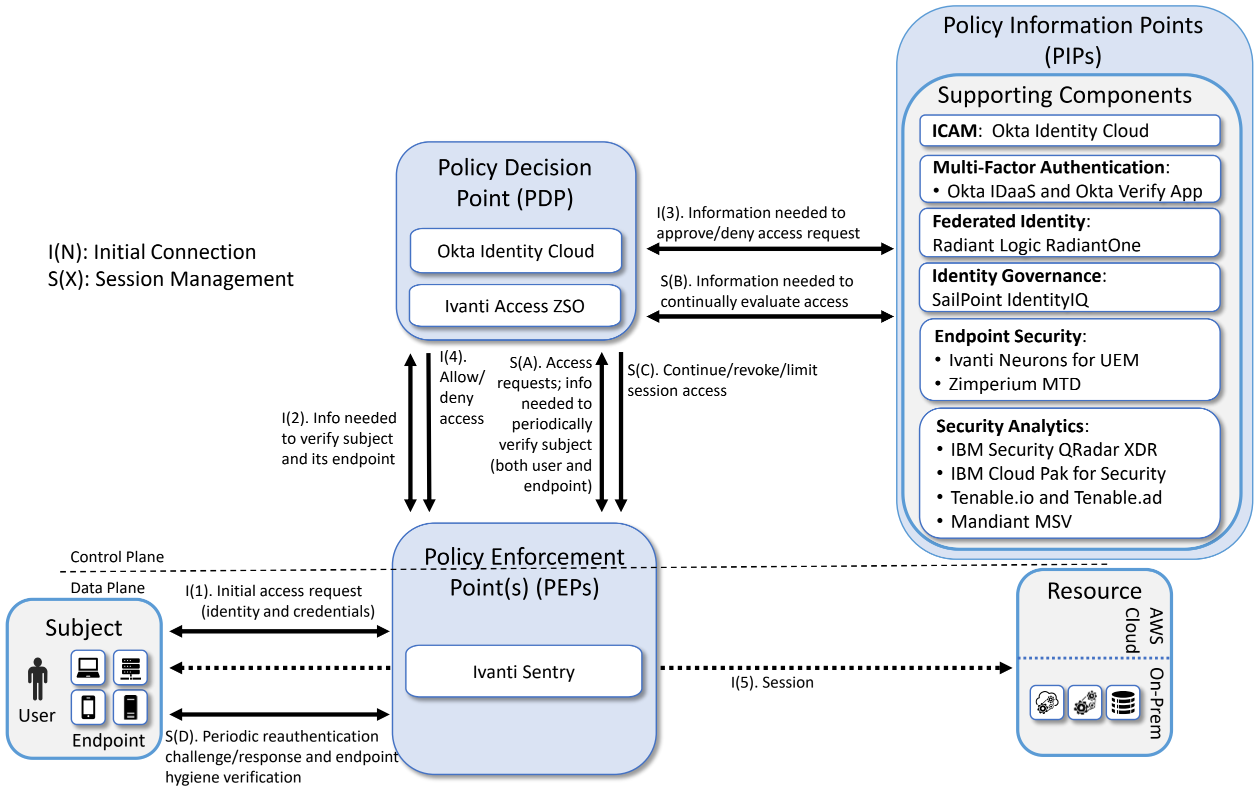

D.2.2 ICAM Information Architecture

How ICAM information is provisioned, distributed, updated, shared, correlated, governed, and used among ZTA components is fundamental to the operation of the ZTA. The ICAM information architecture ensures that when a subject requests access to a resource, the aggregated set of identity information and attributes necessary to identify, authenticate, and authorize the subject is available to be used as a basis on which to make the access decision.

In E1B1, Okta, Radiant Logic, and SailPoint integrate with each other as well as with other components of the ZTA to support the ICAM information architecture. Okta Identity Cloud uses authentication and authorization to manage access to enterprise resources. SailPoint governs and RadiantOne aggregates identity information that is available from many sources within the enterprise. Radiant Logic stores, normalizes, and correlates this aggregation of information and extended attributes and provides appropriate views of the information in response to queries. RadiantOne monitors each source of truth for identity and updates changes in near real-time to ensure that Okta is able to enforce access based on accurate data. SailPoint is responsible for governance of the identity data. It executes automated, policy-based workflows to manage the lifecycle of user identity information and manage user accounts and permissions, ensuring compliance with requirements and regulations. To perform its identity aggregation and correlation functions, Radiant Logic connects to all locations within the enterprise where identity data exists to create a virtualized central identity data repository. SailPoint may also connect directly to sources of identity data or receive additional normalized identity data from Radiant Logic in order to perform its governance functions.

Use of these three components to support the ICAM information architecture in Enterprise 1 is intended to demonstrate how a large enterprise with a complex identity environment might operate—for example, an enterprise with two ADs and multiple sources of identity information, such as HR platforms, the back-end database of a risk-scoring application, a credential management application, a learning management application, on-premises LDAP and databases, etc. Mimicking a large, complex enterprise enables the project to demonstrate the ability to aggregate identity data from many sources and provide identity managers with a rich set of attributes on which to base access policy. By aggregating risk-scoring and training data with more standard identity profile information found in AD, rich user profiles can be created, enabling enterprise managers to formulate and enforce highly granular access policies. Information from any number of the identity and attribute sources can be used to make authentication and authorization decisions. In addition, such aggregation allows identities for users in a partner organization whose identity information is not in the enterprise AD to be made available to the enterprise identity manager, so it has the information required to grant or deny partner user access requests. Policy-based access enforcement is also possible, in which access groups can be dynamically generated based on attribute values.

Although federated identity and identity governance technologies provide automation to ease the burden of aggregating identity information and enforcement of identity governance, they are not required supporting components for implementing a ZTA in situations in which there may only be one or a few sources of identity data.

The subsections below explain the operations of the ICAM information architecture for E1B1 when correlating identity information and when a user joins, changes roles, or leaves the enterprise. The operations depicted support identity correlation, identity management, identity authentication and authorization, and SIEM notification. It is worth noting that both Okta and SailPoint also support additional features that we have not deployed at this time, such as the ability to perform just-in-time provisioning of user accounts and permissions and the ability to remove access permissions or temporarily disable access authorizations from user accounts in response to alerts triggered by suspicious user activity.

D.2.2.1 Identity Correlation

Figure D-2 depicts the ICAM information architecture for E1B1 showing the steps involved in correlating identity information to build a rich global profile that includes not just identity profiles found in AD, but additional profiles and attributes from other platforms as well. The steps are as follows:

1. RadiantOne aggregates, correlates, and normalizes identity information from all sources of identity information in the enterprise. In complex architectures, a ZTA requires an identity data foundation that bridges legacy systems and cloud technologies, and that extends beyond legacy AD domains. In our builds, the identity source used is an example human resources (HR) database that is augmented by extended user profile and attribute information that is representative of information that could come from a variety of identity sources in a large enterprise. A credential management database, an LDAP database, and a learning management application are some examples of such identity sources. These are depicted in the lower left-hand corner of Figure D-2 in the box labeled "Enhanced Identity Data Sources."
2. The correlated identity profiles in RadiantOne are consumed by SailPoint.
3. SailPoint provisions identities into AD. Multiple AD instances may be present in the enterprise, as depicted. However, each of our builds includes only one AD instance.
4. RadiantOne correlates endpoint identities from AD.
5. SailPoint provisions identities into appropriate enterprise resources—e.g., SaaS, IaaS, enterprise applications, and endpoint protection platforms. (This provisioning may occur directly or via Okta.)
6. As the new identities appear in the SaaS, IaaS, enterprise application, endpoint protection, and other components, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization views to reflect the recent changes. Okta will eventually query these authentication and authorization information views in Radiant Logic to determine whether to grant future user access requests.
7. Because Okta is maintaining its own internal identity directory, which is a mirrored version of the information in Radiant Logic, Okta consumes identities from Radiant Logic RadiantOne profiles. However, Okta does not store user password information.
8. RadiantOne correlates identities that it gets from Okta.

The identity correlation lifecycle is an ongoing process that occurs continuously as events that affect user identity information, accounts, and permissions occur, ensuring that the global identity profile is up to date. Example of such events are depicted in the subsections below.

2760 **Figure D-2 E1B1 ICAM Information Architecture – Identity Correlation**

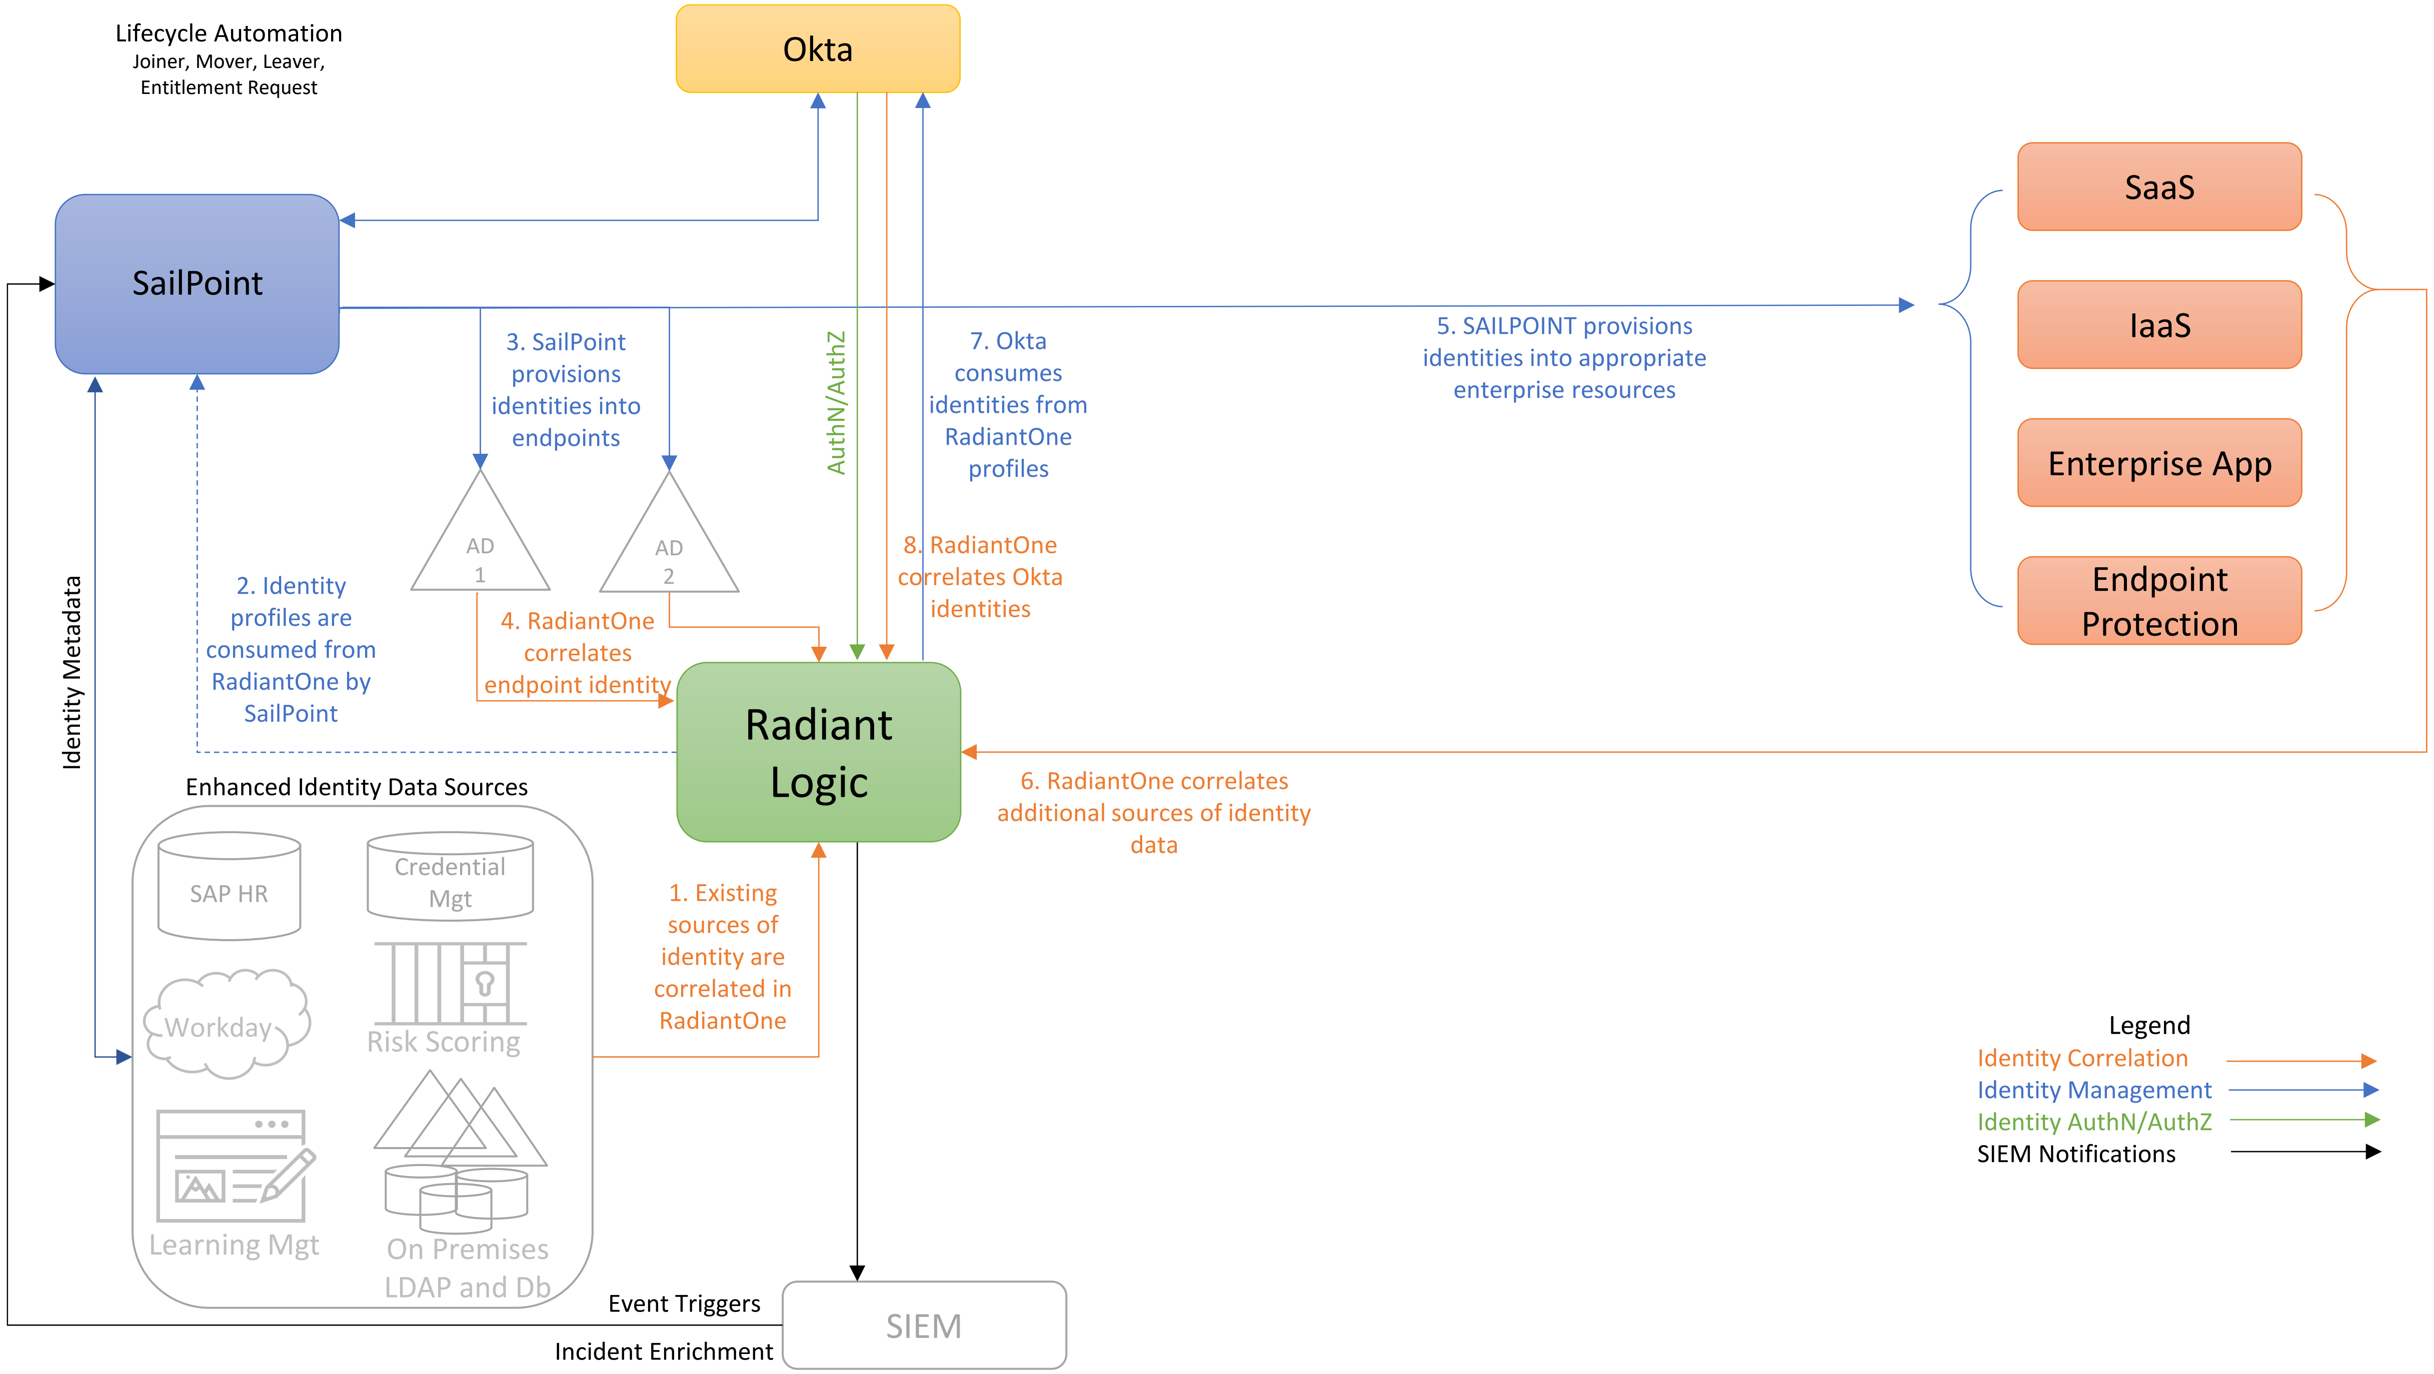

D.2.2.2 User Joins the Enterprise

Figure D-3 depicts the ICAM information architecture for E1B1 showing the steps required to provision a new identity and associated access privileges when a new user is onboarded to the enterprise. The steps are as follows:

1. When a new user joins the enterprise, an authorized HR staff member is assumed to input information into some sort of enterprise employee onboarding and management HR application that will ultimately result in a new, active HR record for the employee appearing in the Radiant Logic human resources record view. In practice, the application that the HR staff member uses will typically store identity records in backend databases like the ones depicted in the lower left-hand corner of Figure D-3 that are in the box labeled "Enhanced Identity Data Sources." As these databases get updated, Radiant Logic is notified, and it responds by collecting the new information and using it to dynamically update its HR view.
2. In the course of performing its governance activities, SailPoint detects the new HR record in Radiant Logic. SailPoint evaluates this new HR record, which triggers a *Joiner* lifecycle event, causing SailPoint to execute a policy-driven workflow that includes steps 3, 4, and 5.
3. SailPoint provisions access permissions to specific enterprise resources for this new user. These access permissions, known as the user's *Birthright Role Access*, are automatically determined according to policy based on factors such as the user's role, type, group memberships, and status. These permissions comprise the access entitlements that the employee has on day 1. SailPoint creates an account for the new user in AD, thereby provisioning appropriate enterprise resource access for the new user. Also (not labeled in the diagram), Radiant Logic then collects and correlates this user information from AD into the global identity profile that it is maintaining.
4. Assuming there are resources for which access is not managed by AD that the new user is authorized to access according to their Birthright Role, SailPoint also provisions access to these resources for the new user by creating new accounts for the user, as appropriate, on SaaS, IaaS, enterprise application, MDM, EPP, and other components. (This provisioning may occur directly or via Okta.)
5. Once the new identity and its access privileges have been provisioned, SailPoint audits the identity and provisioning actions that were just performed.
6. As the new enterprise accounts appear in the SaaS, IaaS, enterprise application, endpoint protection, and other components, Radiant Logic is notified. Radiant Logic collects, correlates and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization (AuthN/AuthZ) views to reflect the recent changes. Okta will eventually query these authentication and authorization

information views in Radiant Logic to determine whether or not to grant future user access requests. (Note that Okta will only query these views in Radiant Logic when a user tries to access a resource; it will not query if there is no action from the user.)

7. In addition, because Okta is maintaining its own internal identity directory, which is a mirrored version of the information in Radiant Logic, Radiant Logic pushes the new account identity information into Okta, thereby synchronizing its extended user profile attribute information with Okta. This provides Okta with additional contextual data regarding users and devices that Radiant Logic has aggregated from all identity sources, beyond the birthright provisioning information that SailPoint provided. Also (not labeled in the diagram), Radiant Logic then collects and correlates identity information from Okta back into the global identity profile that it is maintaining.

2807 **Figure D-3 E1B1 ICAM Information Architecture – New User Onboarding**

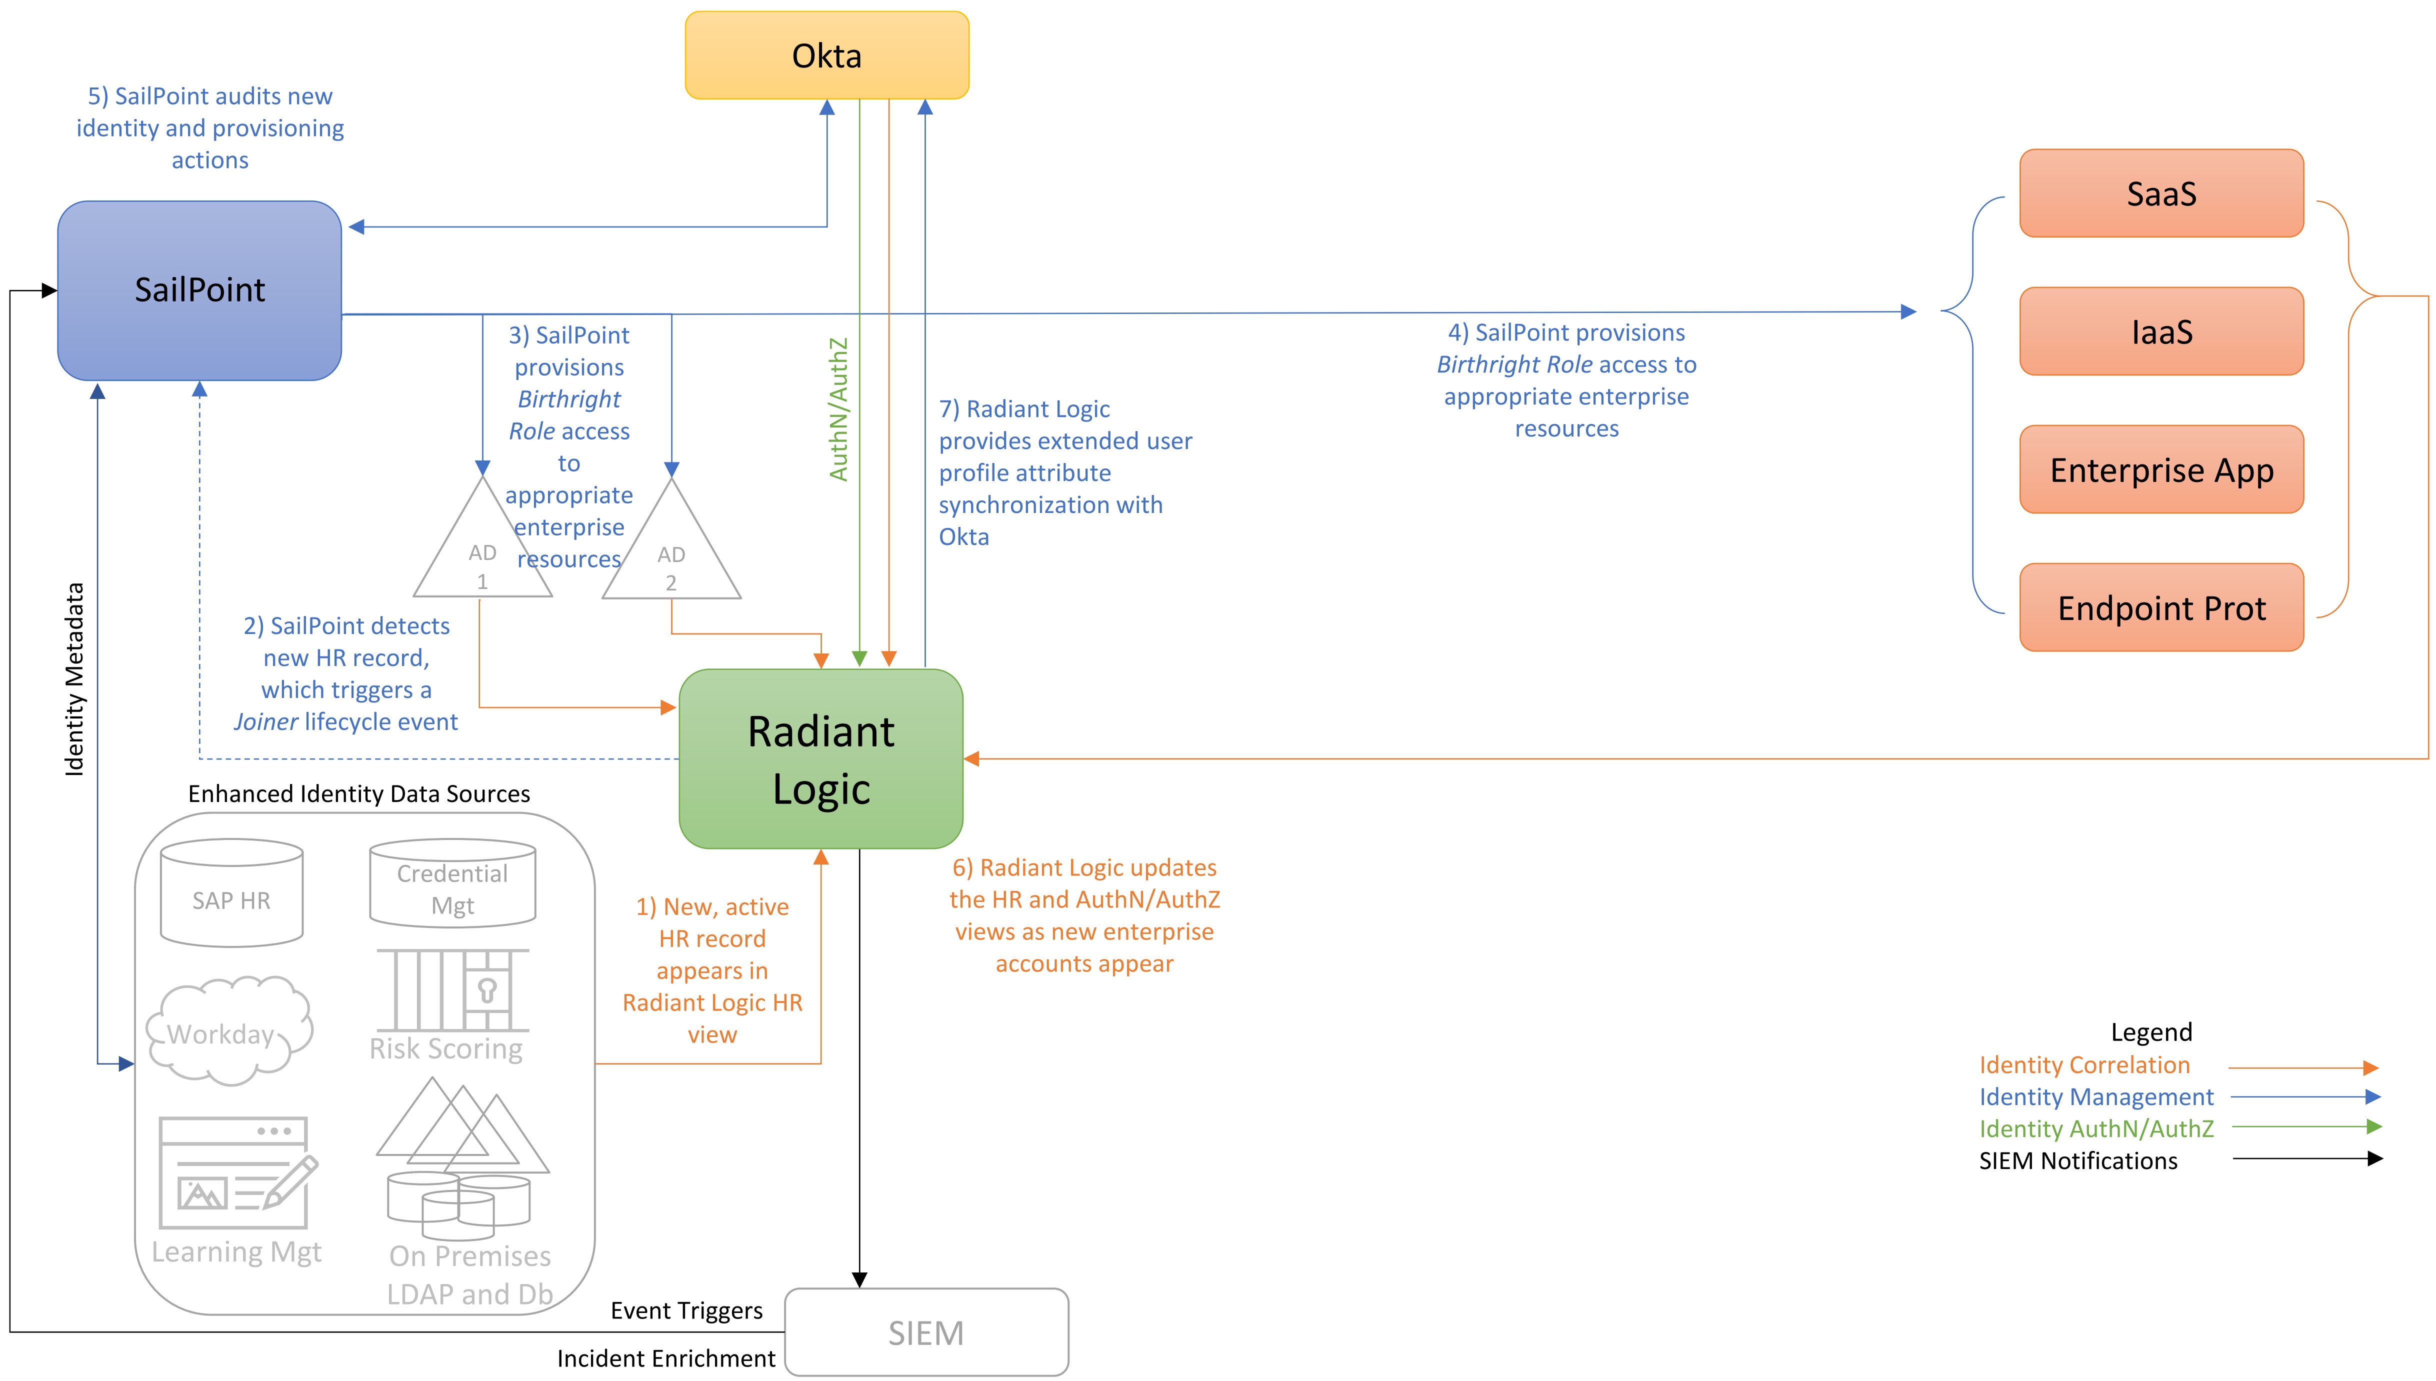

D.2.2.3 User Changes Roles

Figure D-4 depicts the ICAM information architecture for E1B1, showing the steps required to remove some access privileges and add other access privileges for a user in response to that user changing roles within the enterprise. The steps are as follows:

1. When a user changes roles within the enterprise, an authorized HR staff member is assumed to input information into some sort of enterprise employee management application that will result in the Radiant Logic HR record for that user indicating that the user has changed roles.

2. SailPoint detects this updated HR record in Radiant Logic. SailPoint evaluates this updated HR record, which triggers a *Mover* lifecycle event, causing SailPoint to execute a policy-driven workflow that includes steps 3, 4, 5, and 6.

3. SailPoint removes access permissions associated with the user's prior role (but not with the user's new role) from the user's AD account and removes access from other enterprise resources (e.g., SaaS, IaaS, enterprise applications, MDM) that the user had been authorized to access as a result of their prior role, but they are not authorized to access as a result of their new role. Also (not labeled in the diagram), Radiant Logic then collects and correlates any changes that were made to the user's account from AD into the global identity profile that it is maintaining.

4. Assuming there are enterprise resources that the user's new role entitles them to access that are not managed by AD, SailPoint provisions access to these resources for the user by creating new accounts for the user, as appropriate, in SaaS, IaaS, enterprise application, endpoint protection, MDM, and other components. (This provisioning may occur directly or via Okta.)

5. SailPoint generates an access review for management to confirm or revoke the changes that have been made. Such an access review is not strictly necessary. The permission changes could be executed in a fully automated manner, if desired, and specified by policy. However, having an access review provides management with the opportunity to exercise some supervisory discretion to permit the user to temporarily continue to have access to some resources associated with their former role that may still be needed.

6. Once the access review has been completed and any access privilege changes deemed necessary have been performed, SailPoint audits the changes.

7. As the new enterprise accounts appear in the SaaS, IaaS, enterprise application, endpoint protection, and other components, and as existing account access is removed, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization views to reflect the recent changes. Okta will eventually query

these authentication and authorization information views in Radiant Logic to determine whether to grant future user access requests.

8. In addition, because Okta is maintaining its own internal identity directory, which is a mirrored version of the information in Radiant Logic, Radiant Logic pushes the modified account identity information into Okta, thereby synchronizing its user profile attribute information with Okta. Also (not labeled in the diagram), Radiant Logic then collects and correlates identity information from Okta back into the global identity profile that it is maintaining.

2849 **Figure D-4 E1B1 ICAM Information Architecture - User Changes Roles**

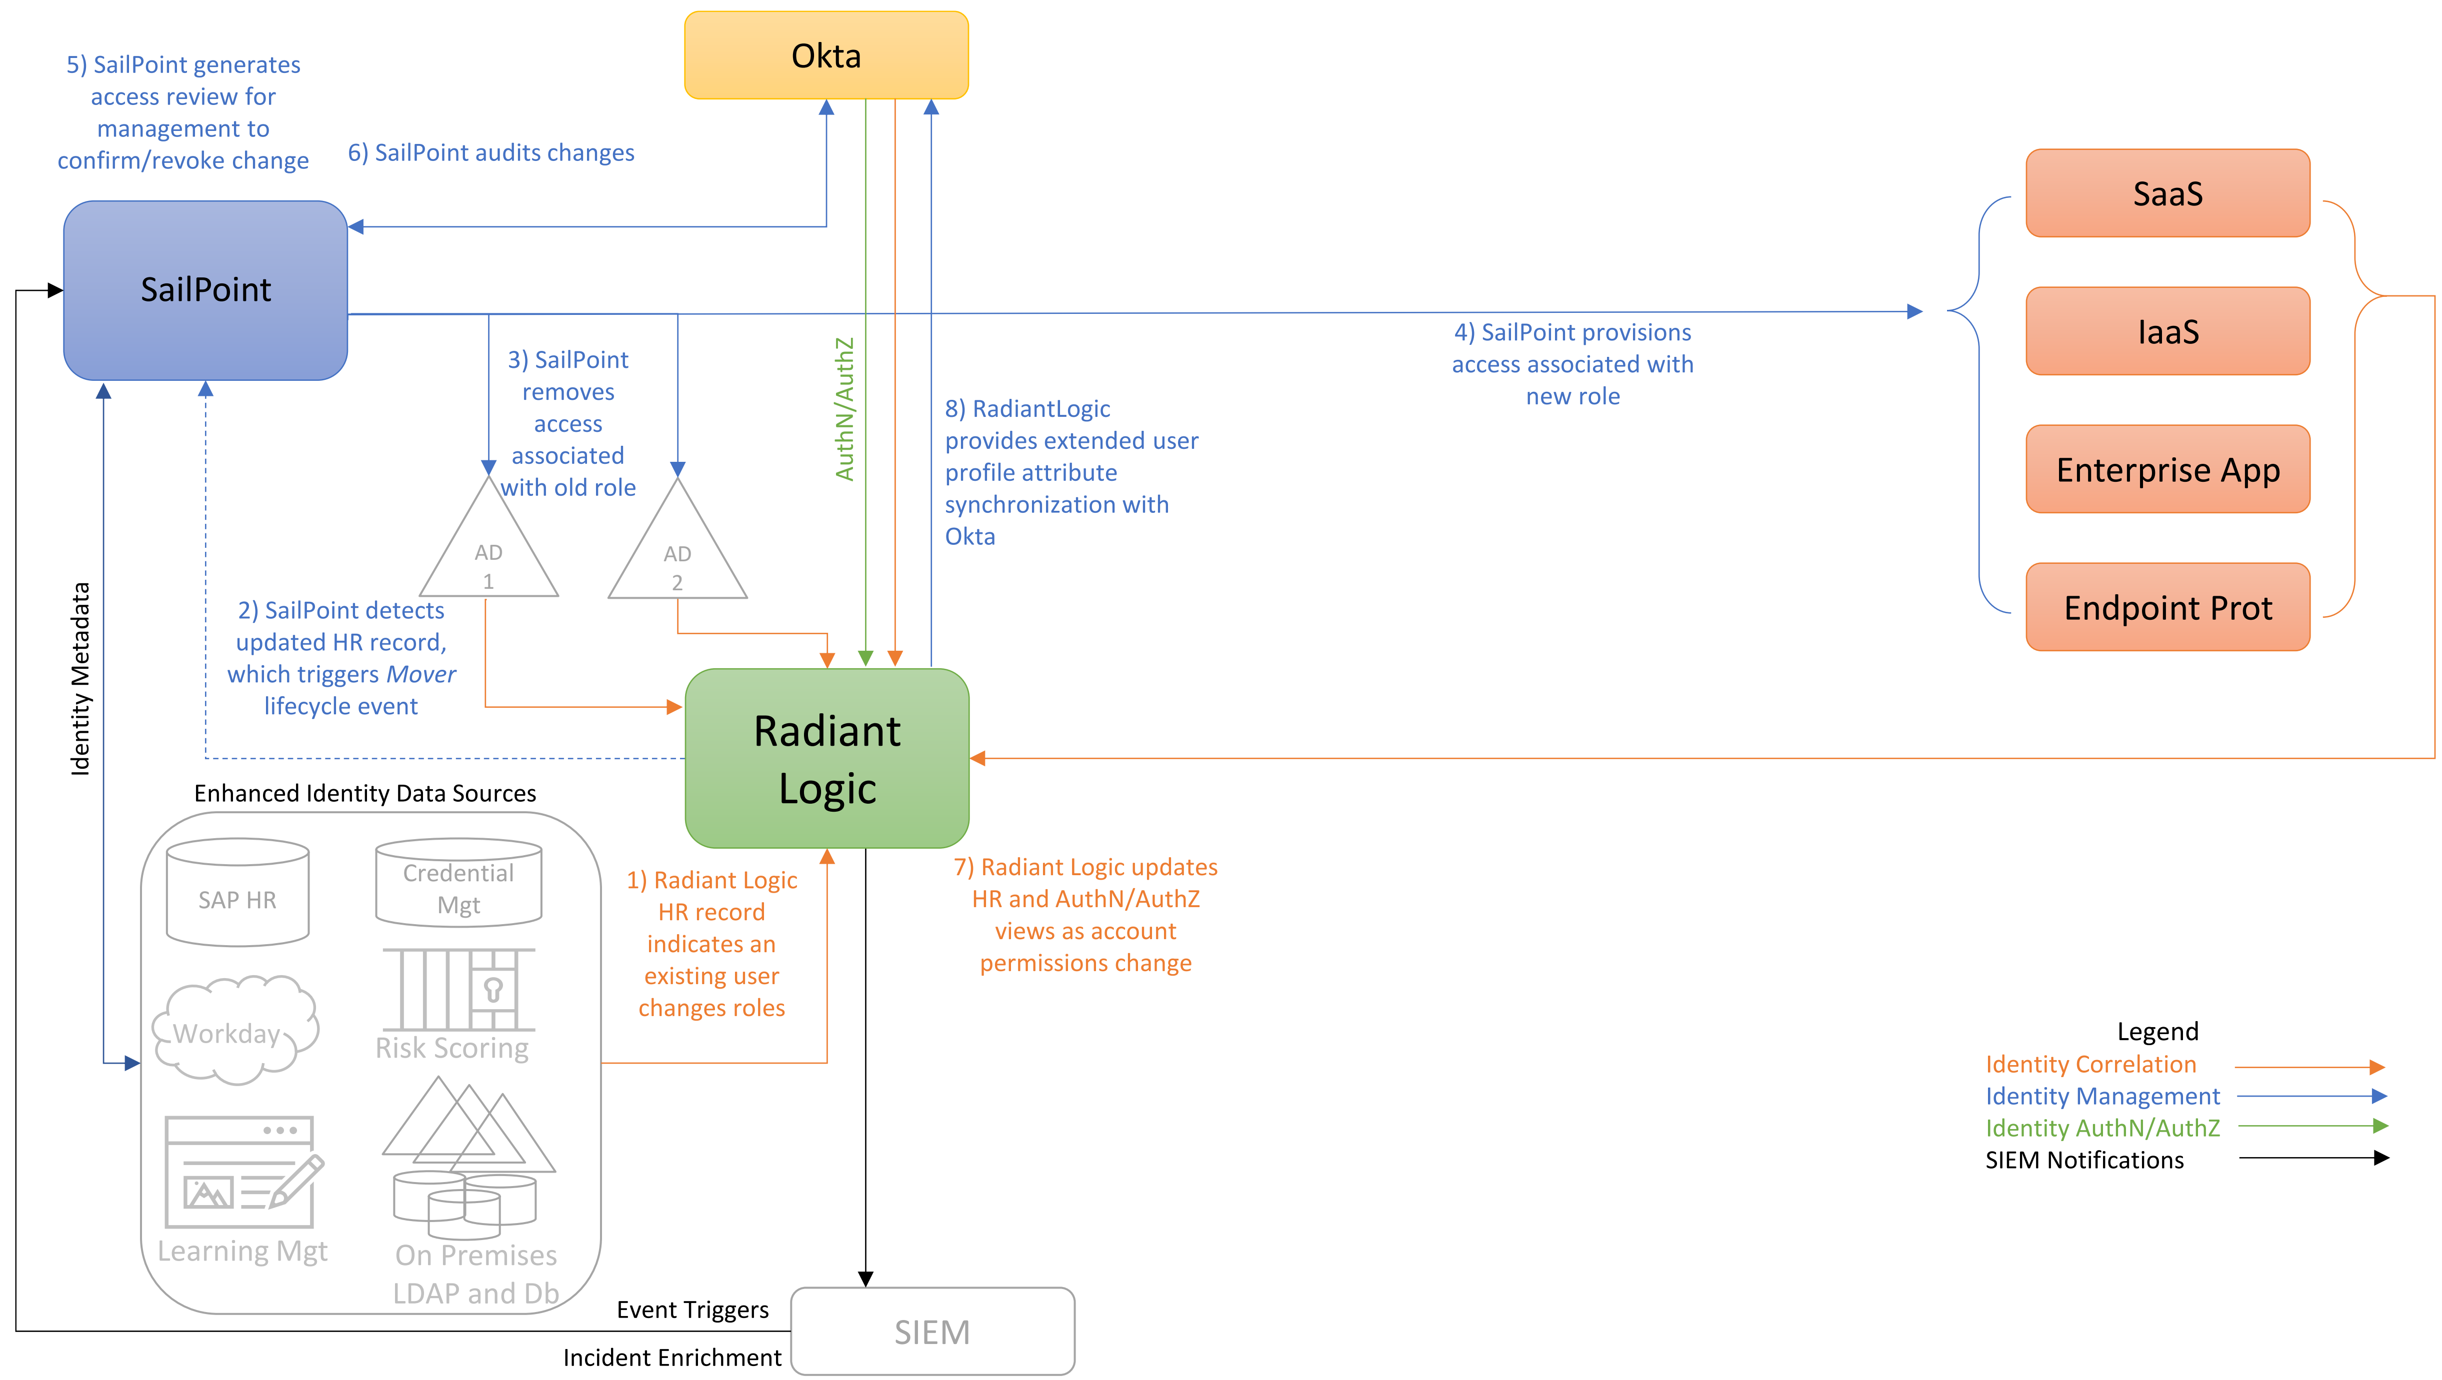

D.2.2.4 User Leaves the Enterprise

Figure D-5 depicts the ICAM information architecture for E1B1 showing the steps required to disable or delete an identity and remove access privileges in response to a user leaving the enterprise. The steps are as follows:

1. When a user's employment is terminated, an authorized HR staff member is assumed to input information into some sort of enterprise employee management application that will result in the Radiant Logic HR record for that user indicating that the user has changed from active to inactive status.

2. SailPoint detects this updated HR record in Radiant Logic. SailPoint evaluates this updated HR record, which triggers a *Leaver* lifecycle event, causing SailPoint to execute a policy-driven workflow that includes steps 3, 4, 5, and 6.

3. SailPoint removes all access permissions associated with the user identity from AD. Also (not labeled in the diagram), Radiant Logic then collects and correlates this user access authorization change from AD into the global identity profile that it is maintaining.

4. SailPoint either disables or deletes all enterprise resource accounts associated with the user identity, as defined by policy, from components such as SaaS, IaaS, enterprise applications, and endpoint protection platforms. (SailPoint may perform these actions directly or via Okta.)

5. SailPoint removes the user identity from all governance groups the identity is in.

6. SailPoint audits the changes made as a result of this user termination.

7. As the enterprise accounts associated with the user's identity are deleted or disabled, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization views to reflect the recent changes. Okta will eventually query these authentication and authorization information views in Radiant Logic to determine whether or not to grant future user access requests.

8. In addition, because Okta is maintaining its own internal identity directory, which is a mirrored version of the information in Radiant Logic, Radiant Logic pushes the modified account identity information into Okta, thereby synchronizing its user profile attribute information with Okta. Also (not labeled in the diagram), Radiant Logic then collects and correlates identity information from Okta back into the global identity profile that it is maintaining.

2880 **Figure D-5 E1B1 ICAM Information Architecture - User Termination**

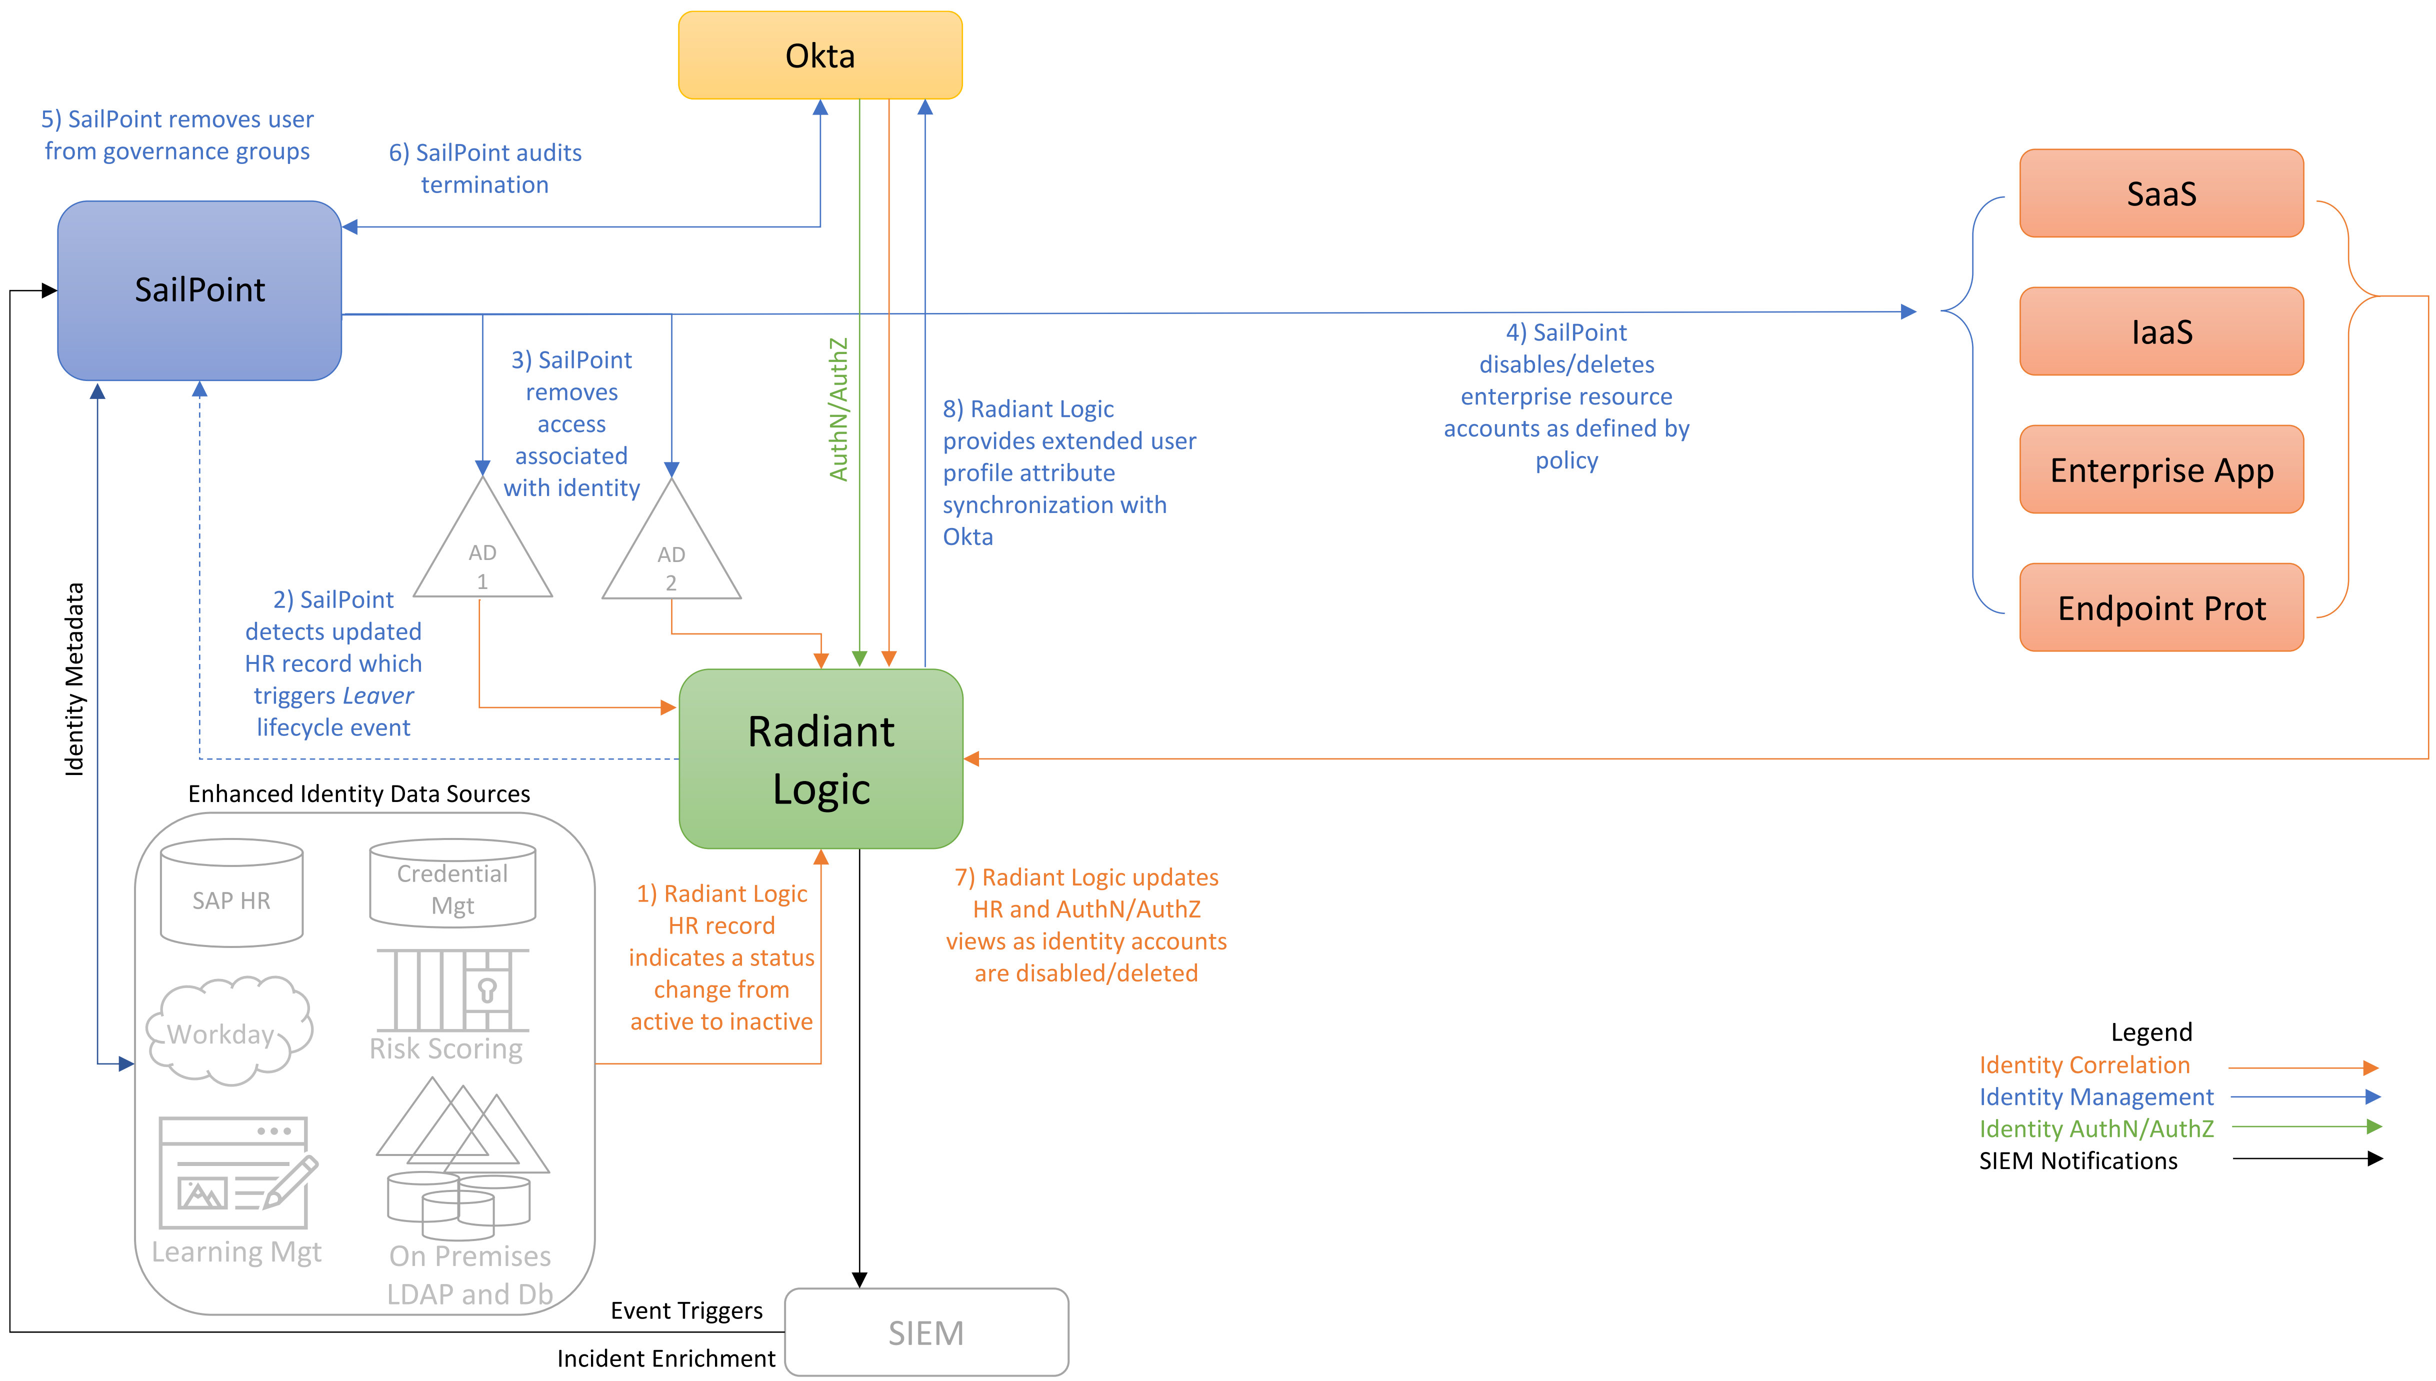

D.2.3 Physical Architecture

Sections 4.4.1 and 4.4.2 describe and depict the physical architecture of the E1B1 headquarters network and the E1B1 branch office network, respectively. In addition to what is represented in Section 4.4, E1B1 has a VLAN on which servers hosting IBM Cloud Pak for Security components reside. It also has MobileIron Connector in its Shared Services VLAN and MobileIron Sentry in its DMZ VLAN.

D.2.4 Message Flow for a Successful Resource Access Request

Figure D-6 shows the high-level message flow for a use case in which a subject who has an enterprise ID, is located on-premises, and is authorized to access an enterprise resource requests and receives access to that resource. In the case depicted in the figure, access to the resource is protected by the Ivanti Sentry gateway, which acts as a PEP; Ivanti Neurons for UEM, which consists of a UEM agent on the endpoint and a cloud component that work together to authenticate the requesting endpoint and determine whether or not it is compliant; Ivanti Access ZSO, which acts as a delegated IdP and consults the Okta Identity Cloud to authenticate the requesting user; and the Okta Verify App, which performs second-factor user authentication.

The message flow depicted in Figure D-6 shows only the messages that are sent in response to the access request. However, the authentication process also relies on the following additional background communications that occur among components on an ongoing basis:

- The Ivanti Neurons for UEM agent periodically synchronizes with Ivanti Neurons for UEM to reauthenticate the requesting endpoint device using a unique certificate that has been provisioned specifically for that device and send Ivanti Neurons for UEM information about device attributes.
- Zimperium periodically sends mobile defense threat information to Ivanti Neurons for UEM.
- Ivanti Neurons for UEM determines device health status based on the above information that it receives from both the Ivanti Neurons for UEM agent and Zimperium.
- Ivanti Neurons for UEM periodically sends device health information to Ivanti Access ZSO.
- Ivanti Neurons for UEM also periodically sends device health information to the Ivanti Sentry gateway.
- Okta periodically synchronizes with Ivanti Neurons for UEM and Ivanti Access ZSO to get the most up-to-date identity information and ensure that the endpoint device is managed by Ivanti Neurons for UEM.

Figure D-6 Successful Access Request Enforced by Okta, Ivanti, and Zimperium Components

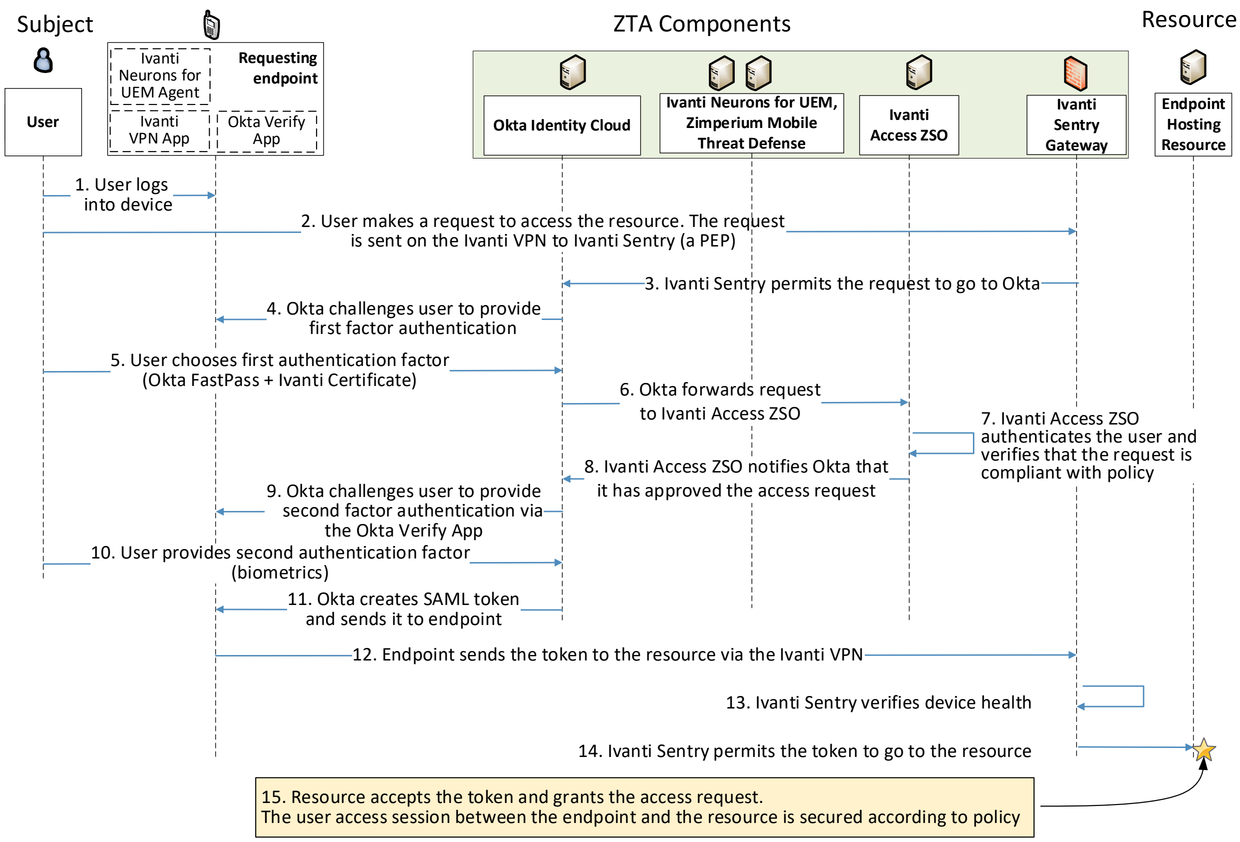

The message flow depicted in Figure D-6 assumes that a VPN between an app on the user's endpoint and the Ivanti Sentry gateway (PEP) has already been set up and connected prior to the user's access request. This VPN connection is established automatically as soon as the device is connected to the network, and it can be configured to be in an "Always On" state. The steps in this message flow, which depicts a successful resource access, are as follows:

1. The user logs into their device and authenticates themselves according to organization policy as configured in Ivanti Neurons for UEM. (This login could be accomplished with a fingerprint ID, face ID, PIN, derived credentials, or any other mechanism that is supported by the device and permitted by organizational policy as configured in the UEM.)

2. The user requests to access a resource. This request is sent on the VPN from the user's endpoint to the Ivanti Sentry gateway, which acts as a PEP.

3. Based on information about the endpoint and user that the Ivanti Sentry gateway has received in the background from Ivanti Neurons for UEM, the Ivanti Sentry gateway determines that,

according to policy, this request is permitted to be sent to Okta, so it allows the access request to proceed to the Okta Identity Cloud component.

4. Okta requests the user to provide authentication information by using Okta FastPass. Okta FastPass allows the user to bypass username and password authentication because Okta trusts that the user properly authenticated when they initially logged into the device in step 1, and Okta knows (from background communications with Ivanti Access ZSO) that Ivanti Neurons for UEM is managing the device.
5. The user provides first-factor authentication information by pressing the Okta FastPass button displayed on the device.
6. Okta forwards the access request information to Ivanti Access ZSO because Okta will rely on and trust Ivanti Access ZSO to perform user authentication and verify the request's attributes to ensure that they conform with policy. In this instance, Ivanti Access will act as a PDP to determine whether the access request should be granted.
7. Ivanti Access authenticates the user using the access request information relayed by Okta. Ivanti Access gets user identities, attributes, and device information from a published certificate that was provisioned uniquely to the device. The certificate contains user information in a Certificate Subject Alternative field. Ivanti Neurons for UEM uses Okta as an identity provider and regularly syncs with Okta to remain up to date. It does not reach back to Okta every time an identity request comes in. Ivanti Access also verifies that the device complies with its conditional access policy. If any policy is being violated, device access is blocked, and a remediation page is presented to the user. Ivanti Access ZSO makes this determination based on information it has been receiving in the background from Ivanti Neurons for UEM and Zimperium.
8. Ivanti Access ZSO notifies Okta that it has approved the access request by signing an authentication token using the Ivanti Access ZSO signing certificate.
9. Okta initiates second-factor authentication using the Okta Verify App. Okta requires the user to present their biometric information to authenticate themselves to the device, and then the Okta Verify App displays a notification on the device informing the user that they must respond (e.g., tap a confirmation button on the display) to prove that they are in possession of the device.
10. The user presents their biometric information and responds to the Okta Verify notification, thereby providing the second authentication factor.
11. Okta creates a SAML assertion and sends it to the requesting endpoint.
12. The requesting endpoint sends the SAML assertion to the resource via the VPN that connects to the Ivanti Sentry gateway.

13. The Ivanti Sentry gateway verifies device health and compliance based on the device information it has been receiving in the background from Ivanti Neurons for UEM.
14. The Ivanti Sentry gateway permits the SAML assertion to proceed to the resource.
15. The resource accepts the assertion and grants the access request. User traffic to and from the resource is secured according to policy (e.g., using TLS or HTTPS).

Note that the message flow depicted in Figure D-6 applies to several of the use cases we are considering. It applies to all cases in which a user with an enterprise ID who can successfully authenticate themselves and who is using an enterprise-owned endpoint requests and receives access to an enterprise resource that they are authorized to access. The message flow is the same regardless of whether the employee is located on-premises at headquarters, on-premises at a branch office, or off-premises at home or elsewhere. It is also the same regardless of whether the resource is located on-premises or in the cloud.

Appendix E EIG Enterprise 2 Build 1 (E2B1)

E.1 Technologies

EIG E2B1 uses products from Cisco Systems, IBM, Mandiant, Palo Alto, Ping Identity, Radiant Logic, SailPoint, and Tenable. Certificates from DigiCert are also used. For more information on these collaborators and the products and technologies that they contributed to this project overall, see Section 3.4.

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

Table E-1 lists all of the technologies used in EIG E2B1. It lists the products used to instantiate each ZTA component and the security function that each component provides.

Table E-1 E2B1 Products and Technologies

Component	Product	Function
PE	Ping Identity PingFederate	Decides whether to grant, deny, or revoke access to a resource based on enterprise policy, information from supporting components, and a trust algorithm.
PA	Ping Identity PingFederate	Executes the PE’s policy decision by sending commands to a PEP that establishes and shuts down the communication path between subject and resource.
PEP	Ping Identity PingFederate Cisco Duo	Guards the trust zone that hosts one or more enterprise resources; establishes, monitors, and terminates the connection between subject and resource as directed by the PA; forwards requests to and receives commands from the PA.
Identity Management	Ping Identity PingFederate	Creates and manages enterprise user and device accounts, identity records, role information, and access attributes that form the basis of access decisions within an organization to ensure the correct subjects have the appropriate access to the correct resources at the appropriate time.
Access & Credential Management	Ping Identity PingFederate	Manages access to resources by performing user and device authentication (e.g., SSO and MFA) and using identity, role, and access attributes to determine which access requests are authorized.

Component	Product	Function
Federated Identity	Radiant Logic RadiantOne Intelligent Identity Data Platform	Aggregates and correlates all attributes relating to an identity or object that is being authorized by a ZTA. It enables users of one domain to securely access data or systems of another domain seamlessly, and without the need for completely redundant user administration. Federated identity encompasses the traditional ICAM data, supports identities that may be part of a larger federated ICAM community, and may include non-enterprise employees.
Identity Governance	SailPoint IdentityIQ	Provides policy-based, centralized, automated processes to manage user identity and access control functions (e.g., ensuring segregation of duties, role management, logging, access reviews, analytics, reporting) to ensure compliance with requirements and regulations.
MFA	Cisco Duo	Supports MFA of a user identity by requiring the user to provide not only something they know (e.g., a password), but also something they have (e.g., a token).
UEM/MDM	None	Manages and secures enterprise desktop computers, laptops, and/or mobile devices in accordance with enterprise policy to protect applications and data; ensure device compliance; mitigate and remediate vulnerabilities and threats; monitor for suspicious activity to prevent and detect intrusions; prevent, detect, and disable malware and other malicious or unauthorized traffic; repair infected files when possible; provide alerts and recommend remediation actions; and encrypt data. Pushes enterprise applications and updates to devices, enables users to download enterprise applications that they are authorized to access, remotely deletes all applications and data from devices if needed, tracks user activity on devices, and detects and addresses security issues on the device.

Component	Product	Function
EPP	None	Detects and stops threats to endpoints through an integrated suite of endpoint protection technologies including antivirus, data encryption, intrusion prevention, EDR, and DLP. May include mechanisms that are designed to protect applications and data; ensure device compliance with policies regarding hardware, firmware, software, and configuration; monitor endpoints for vulnerabilities, suspicious activity, intrusion, infection, and malware; block unauthorized traffic; disable malware and repair infections; manage and administer software and updates; monitor behavior and critical data; and enable endpoints to be tracked, troubleshooted, and wiped, if necessary.
Endpoint Compliance	Cisco Duo	Performs device health checks by validating specific tools or services within the endpoint including antivirus, data encryption, intrusion prevention, EPP, and firewall. If the device does not pass the health check, Duo fails second-factor authentication and denies user access.
SIEM	IBM Security QRadar XDR	Collects and consolidates security information and security event data from many sources; correlates and analyzes the data to help detect anomalies and recognize potential threats and vulnerabilities; and logs the data to adhere to data compliance requirements.
Vulnerability Scanning and Assessment	Tenable.io and Tenable.ad	Scans and assesses the enterprise infrastructure and resources for security risks, identifies vulnerabilities and misconfigurations, and provides remediation guidance regarding investigating and prioritizing responses to incidents.
Network Discovery	Tenable NNM	Discovers, classifies, and assesses the risk posed by devices and users on the network.

Component	Product	Function
Security Validation	Mandiant MSV	Provides visibility and evidence on the status of the security controls' effectiveness in the ZTA. Enables security capabilities of the enterprise to be monitored and verified by continuously validating and measuring the cybersecurity controls; also used to automate the demonstrations that were performed to showcase ZTA capabilities. Deployed throughout the project's laboratory environment to enable monitoring and verification of various security aspects of the builds. 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.
Remote Connectivity	Palo Alto Networks NGFW	Enables authorized remote users to securely access the inside of the enterprise. (Once inside, the ZTA manages the user's access to resources.)
Certificate Management	DigiCert CertCentral TLS Manager	Provides automated capabilities to issue, install, inspect, revoke, renew, and otherwise manage TLS certificates.
Cloud IaaS	None	Provides computing resources, complemented by storage and networking capabilities, hosted by a cloud service provider, offered to customers on demand, and exposed through a GUI and an API.
Cloud SaaS	Cisco Duo, Digicert CertCentral Ping Identity PingOne (PingFederate service), and Tenable.io	Cloud-based software delivered for use by the enterprise.
Application	GitLab	Example enterprise resource to be protected. (In this build, GitLab and WordPress are integrated with Okta using SAML, and IBM Security QRadar XDR pulls logs from GitLab.)
Enterprise-Managed Device	Windows client, macOS client, and mobile devices (iOS and Android)	Example endpoints to be protected. All enterprise-managed devices are running an Ivanti Neurons for UEM agent and also have the Okta Verify App installed.
BYOD	Windows client, macOS client, and mobile devices (iOS and Android)	Example endpoints to be protected.

E.2 Build Architecture

In this section we present the logical architecture of E2B1 relative to how it instantiates the EIG crawl phase reference architecture depicted in Figure 4-2. We also describe E2B1's physical architecture and present message flow diagrams for some of its processes.

E.2.1 Logical Architecture

Figure E-1 depicts the logical architecture of E2B1. The figure uses numbered arrows to depict the general flow of messages needed for a subject to request access to a resource and have that access request evaluated based on subject identity (both requesting user and requesting endpoint identity), user authorizations, and requesting endpoint health. It also depicts the flow of messages supporting periodic reauthentication of the requesting user and the requesting endpoint and periodic verification of requesting endpoint health, all of which must be performed to continually reevaluate access. The labeled steps in Figure E-1 have the same meanings as they do in Figure 4-1 and Figure 4-2. However, Figure E-1 includes the specific products that instantiate the architecture of E2B1. Figure E-1 also does not depict any of the resource management steps found in Figure 4-1 and Figure 4-2 because the ZTA technologies deployed in E2B1 do not support the ability to perform authentication and reauthentication of the resource or periodic verification of resource health.

E2B1 was designed with a single ICAM system (Ping Identity PingFederate) that serves as the identity, access, and credential manager as well as the ZTA PE and PA. PingFederate also serves as its PEP. Radiant Logic acts as a PIP for the PDP as it responds to inquiries and provides user identity and authentication information on demand in order for Ping Identity PingFederate to make near-real-time access decisions. Cisco Duo provides endpoint protection by monitoring the status and configuration of the endpoint to ensure that its health posture continues to conform with enterprise policy. Duo also provides second-factor user authentication. Note that both multifactor authentication and directory services are also available through Ping, but for purposes of this collaborative build, Ping is demonstrating standards-based interoperability by integrating with Cisco Duo for MFA and Radiant Logic RadiantOne for federated identity services. A more detailed depiction of the messages that flow among components to support a user access request can be found in Appendix E.2.4.

Figure E-1 Logical Architecture of E2B1

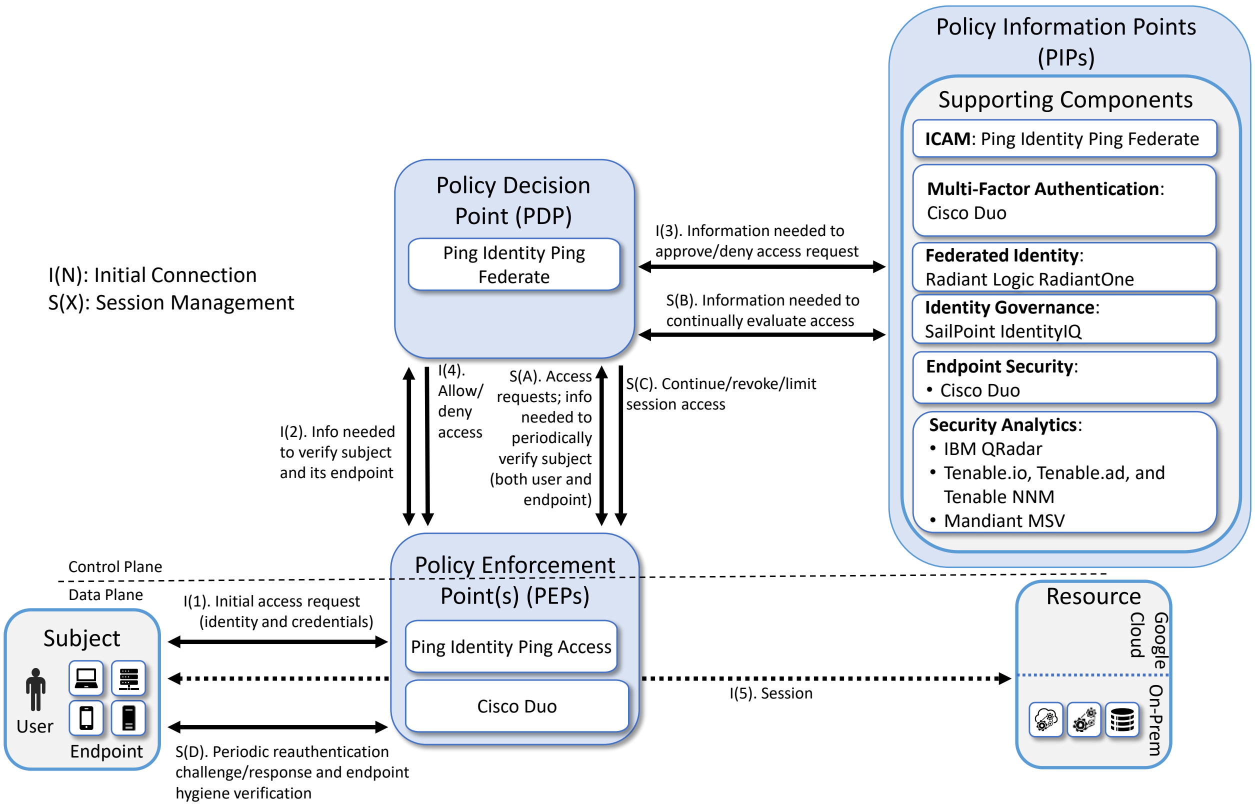

E.2.2 ICAM Information Architecture

How ICAM information is provisioned, distributed, updated, shared, correlated, governed, and used among ZTA components is fundamental to the operation of the ZTA. The ICAM information architecture ensures that when a subject requests access to a resource, the aggregated set of identity information and attributes necessary to identify, authenticate, and authorize the subject is available to be used as a basis on which to make the access decision.

In E2B1, Ping, Radiant Logic, and SailPoint integrate with each other as well as with other components of the ZTA to support the ICAM information architecture. Ping Identity PingFederate uses authentication and authorization to manage access to enterprise resources. SailPoint governs and RadiantOne aggregates identity information that is available from many sources within the enterprise. Radiant One stores, normalizes, and correlates this aggregation of information and extended attributes and provides appropriate views of the information in response to queries. RadiantOne monitors each source of identity truth and updates changes in near real-time to ensure that Ping is able to enforce access based on accurate data. SailPoint is responsible for governance of the identity data. It executes automated, policy-based workflows to manage the lifecycle of user identity information and manage user accounts and permissions, ensuring compliance with requirements and regulations. To perform its identity aggregation and correlation functions, Radiant Logic connects to all locations within the enterprise where identity data exists to create a virtualized central identity data repository. SailPoint may also connect directly to sources of identity data or receive additional normalized identity data from Radiant Logic in order to perform its governance functions.

Use of these three components to support the ICAM information architecture in Enterprise 2 is intended to demonstrate how a large enterprise with a complex identity environment might operate—for example, an enterprise with two ADs and multiple sources of identity information, such as HR platforms, the back-end database of a risk-scoring application, a credential management application, a learning management application, on-premises LDAP and databases, etc. Mimicking a large, complex enterprise enables the project to demonstrate the ability to aggregate identity data from many sources and provide identity managers with a rich set of attributes on which to base access policy. By aggregating risk-scoring and training data with more standard identity profile information found in AD, rich user profiles can be created, enabling enterprise managers to formulate and enforce highly granular access policies. Information from any number of the identity and attribute sources can be used to make authentication and authorization decisions. In addition, such aggregation allows identities for users in a partner organization whose identity information is not in the enterprise AD to be made available to the enterprise identity manager so it has the information required to grant or deny partner user access requests. Policy-based access enforcement is also possible, in which access groups can be dynamically generated based on attribute values.

Although federated identity and identity governance technologies provide automation to ease the burden of aggregating identity information and enforcement of identity governance, they are not required supporting components for implementing a ZTA in situations in which there may only be one or a few sources of identity data.

The subsections below explain the operations of the ICAM information architecture for E2B1 when correlating identity information and when a user joins, changes roles, or leaves the enterprise. The operations depicted support identity correlation, identity management, identity authentication and authorization, and SIEM notification. It is worth noting that both Ping Identity and SailPoint also support additional features that we have not deployed at this time, such as the ability to perform just-in-time provisioning of user accounts and permissions and the ability to remove access permissions or temporarily disable access authorizations from user accounts in response to alerts triggered by suspicious user activity.

E.2.2.1 Identity Correlation

Figure E-2 depicts the ICAM information architecture for E2B1, showing the steps involved in correlating identity information to build a rich global profile that includes not just identity profiles found in AD, but additional profiles and attributes from other platforms as well. The steps are as follows:

1. RadiantOne aggregates, correlates, and normalizes identity information from all sources of identity information in the enterprise. In complex architectures, a ZTA requires an identity data foundation that bridges legacy systems and cloud technologies, and that extends beyond legacy AD domains. In our builds, the identity source used is an example human resources (HR) database that is augmented by extended user profile and attribute information that is representative of information that could come from a variety of identity sources in a large enterprise. A credential management database, an LDAP database, and a learning management application are some examples of such identity sources. These are depicted in the lower left-hand corner of Figure E-2 in the box labeled “Enhanced Identity Data Sources.”
2. The correlated identity profiles in RadiantOne are consumed by SailPoint.
3. SailPoint provisions identities into AD. Multiple AD instances may be present in the enterprise, as depicted. However, each of our builds includes only one AD instance.
4. RadiantOne correlates endpoint identities from AD.
5. SailPoint provisions identities into appropriate enterprise resources—e.g., SaaS, IaaS, enterprise applications, and endpoint protection platforms. (This provisioning may occur directly or via Ping.)
6. As the new identities appear in the SaaS, IaaS, enterprise application, endpoint protection, and other components, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization views to reflect the recent changes. Ping will eventually query these authentication and authorization information views in Radiant Logic to determine whether to grant future user access requests.

   Note that in this architecture, persistent storage of personally identifiable information (PII) is not required within any SaaS service. RadiantOne stores all user identity information, and RadiantOne has been installed on-premises. Ping does not store any user data. When Ping needs user identity data, it looks up this information directly from RadiantOne.

The identity correlation lifecycle is an ongoing process that occurs continuously as events that affect user identity information, accounts, and permissions occur, ensuring that the global identity profile is up to date. Examples of such events are depicted in the subsections below.

3091 **Figure E-2 E2B1 ICAM Information Architecture – Identity Correlation**

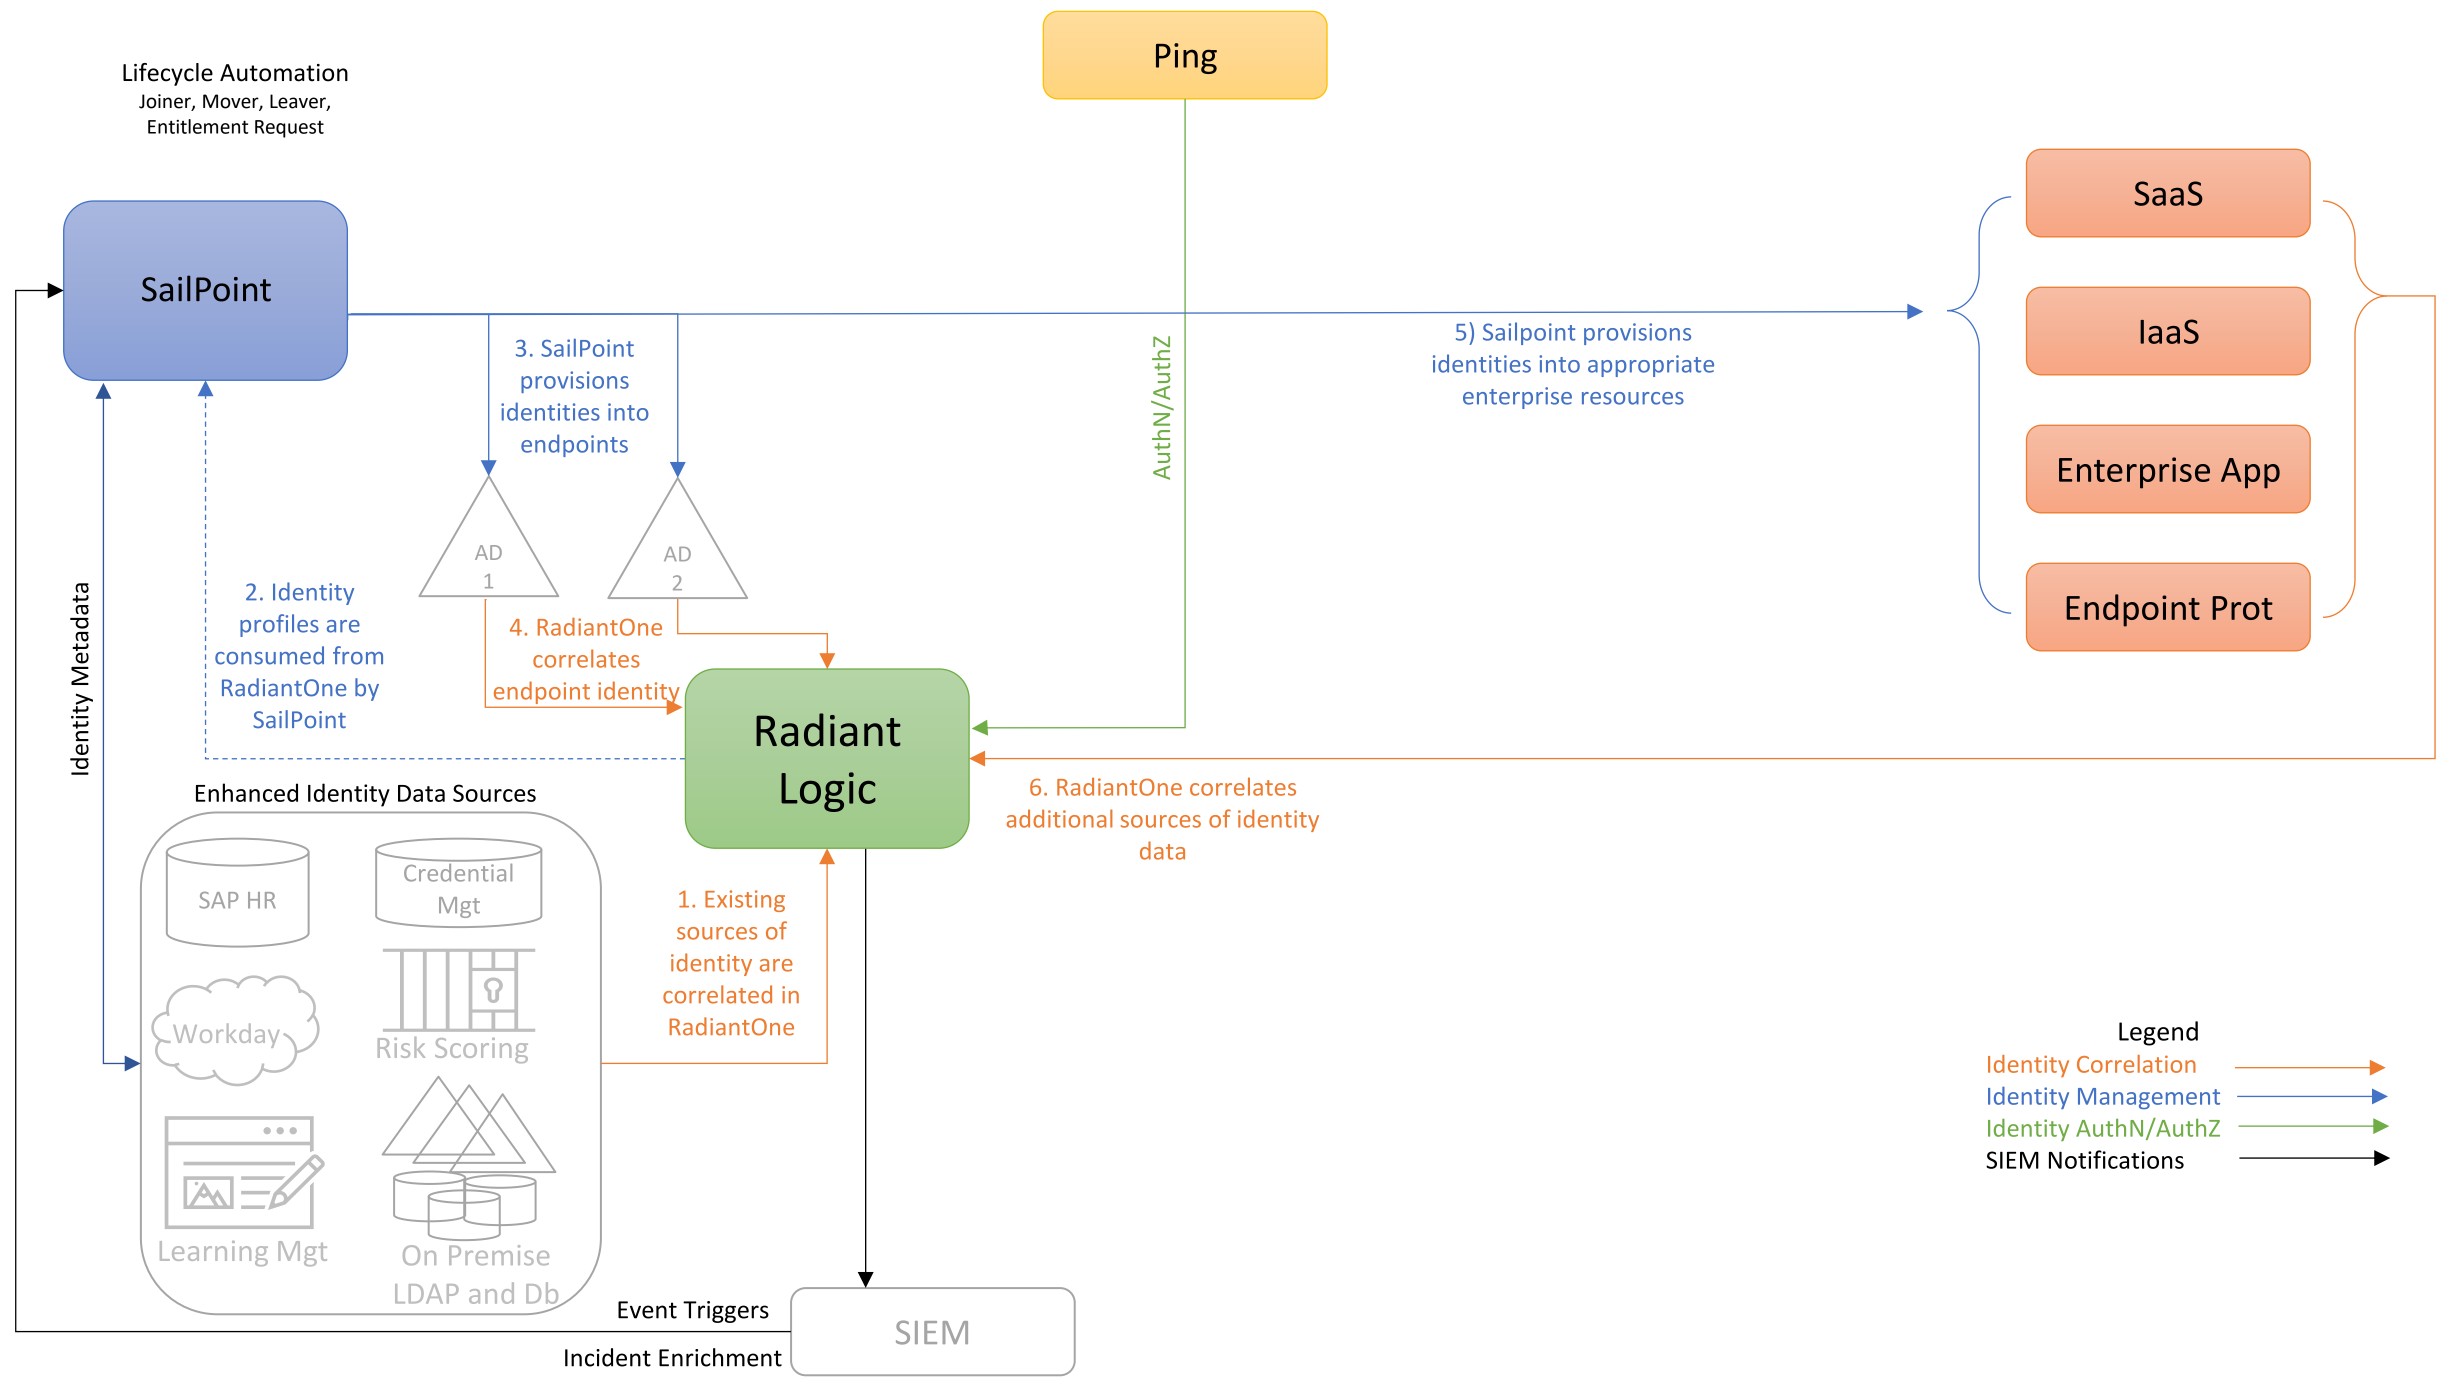

E.2.2.2 User Joins the Enterprise

Figure E-3 depicts the ICAM information architecture for E2B1, showing the steps required to provision a new identity and associated access privileges when a new user is onboarded to the enterprise. The steps are as follows:

1. When a new user joins the enterprise, an authorized HR staff member is assumed to input information into some sort of enterprise employee onboarding and management HR application that will ultimately result in a new, active HR record for the employee appearing in the Radiant Logic human resources record view. In practice, the application that the HR staff member uses will typically store identity records in backend databases like the ones depicted in the lower left-hand corner of Figure D-3 that are in the box labeled "Enhanced Identity Data Sources." As these databases get updated, Radiant Logic is notified, and it responds by collecting the new information and using it to dynamically update its HR view.

2. In the course of performing its governance activities, SailPoint detects the new HR record in Radiant Logic. SailPoint evaluates this new HR record, which triggers a *Joiner* lifecycle event, causing SailPoint to execute a policy-driven workflow that includes steps 3, 4, and 5.

3. SailPoint provisions access permissions to specific enterprise resources for this new user. These access permissions, known as the user's *Birthright Role Access*, are automatically determined according to policy based on factors such as the user's role, type, group memberships, and status. These permissions comprise the access entitlements that the employee has on day 1. SailPoint creates an account for the new user in AD, thereby provisioning appropriate enterprise resource access for the new user. Also (not labeled in the diagram), Radiant Logic then collects and correlates this user information from AD into the global identity profile that it is maintaining.

4. Assuming there are resources for which access is not managed by AD that the new user is authorized to access according to their Birthright Role, SailPoint also provisions access to these resources for the new user by creating new accounts for the user, as appropriate, on SaaS, IaaS, enterprise application, MDM, EPP, and other components. (This provisioning may occur directly or via Ping.)

5. Once the new identity and its access privileges have been provisioned, SailPoint audits the identity and provisioning actions that were just performed.

6. As the new enterprise accounts appear in the SaaS, IaaS, enterprise application, endpoint protection, and other components, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization (AuthN/AuthZ) views to reflect the recent changes. Ping will eventually query these authentication and authorization

information views in Radiant Logic to determine whether or not to grant future user access requests. (Note that Ping will only query these views in Radiant Logic when a user tries to access a resource; it will not query if there is no action from the user. Also, RadiantOne stores all user identity information; Ping does not store any user data. When Ping needs user identity data, it looks up this information directly from RadiantOne.)

3132 **Figure E-3 E2B1 ICAM Information Architecture – New User Onboarding**

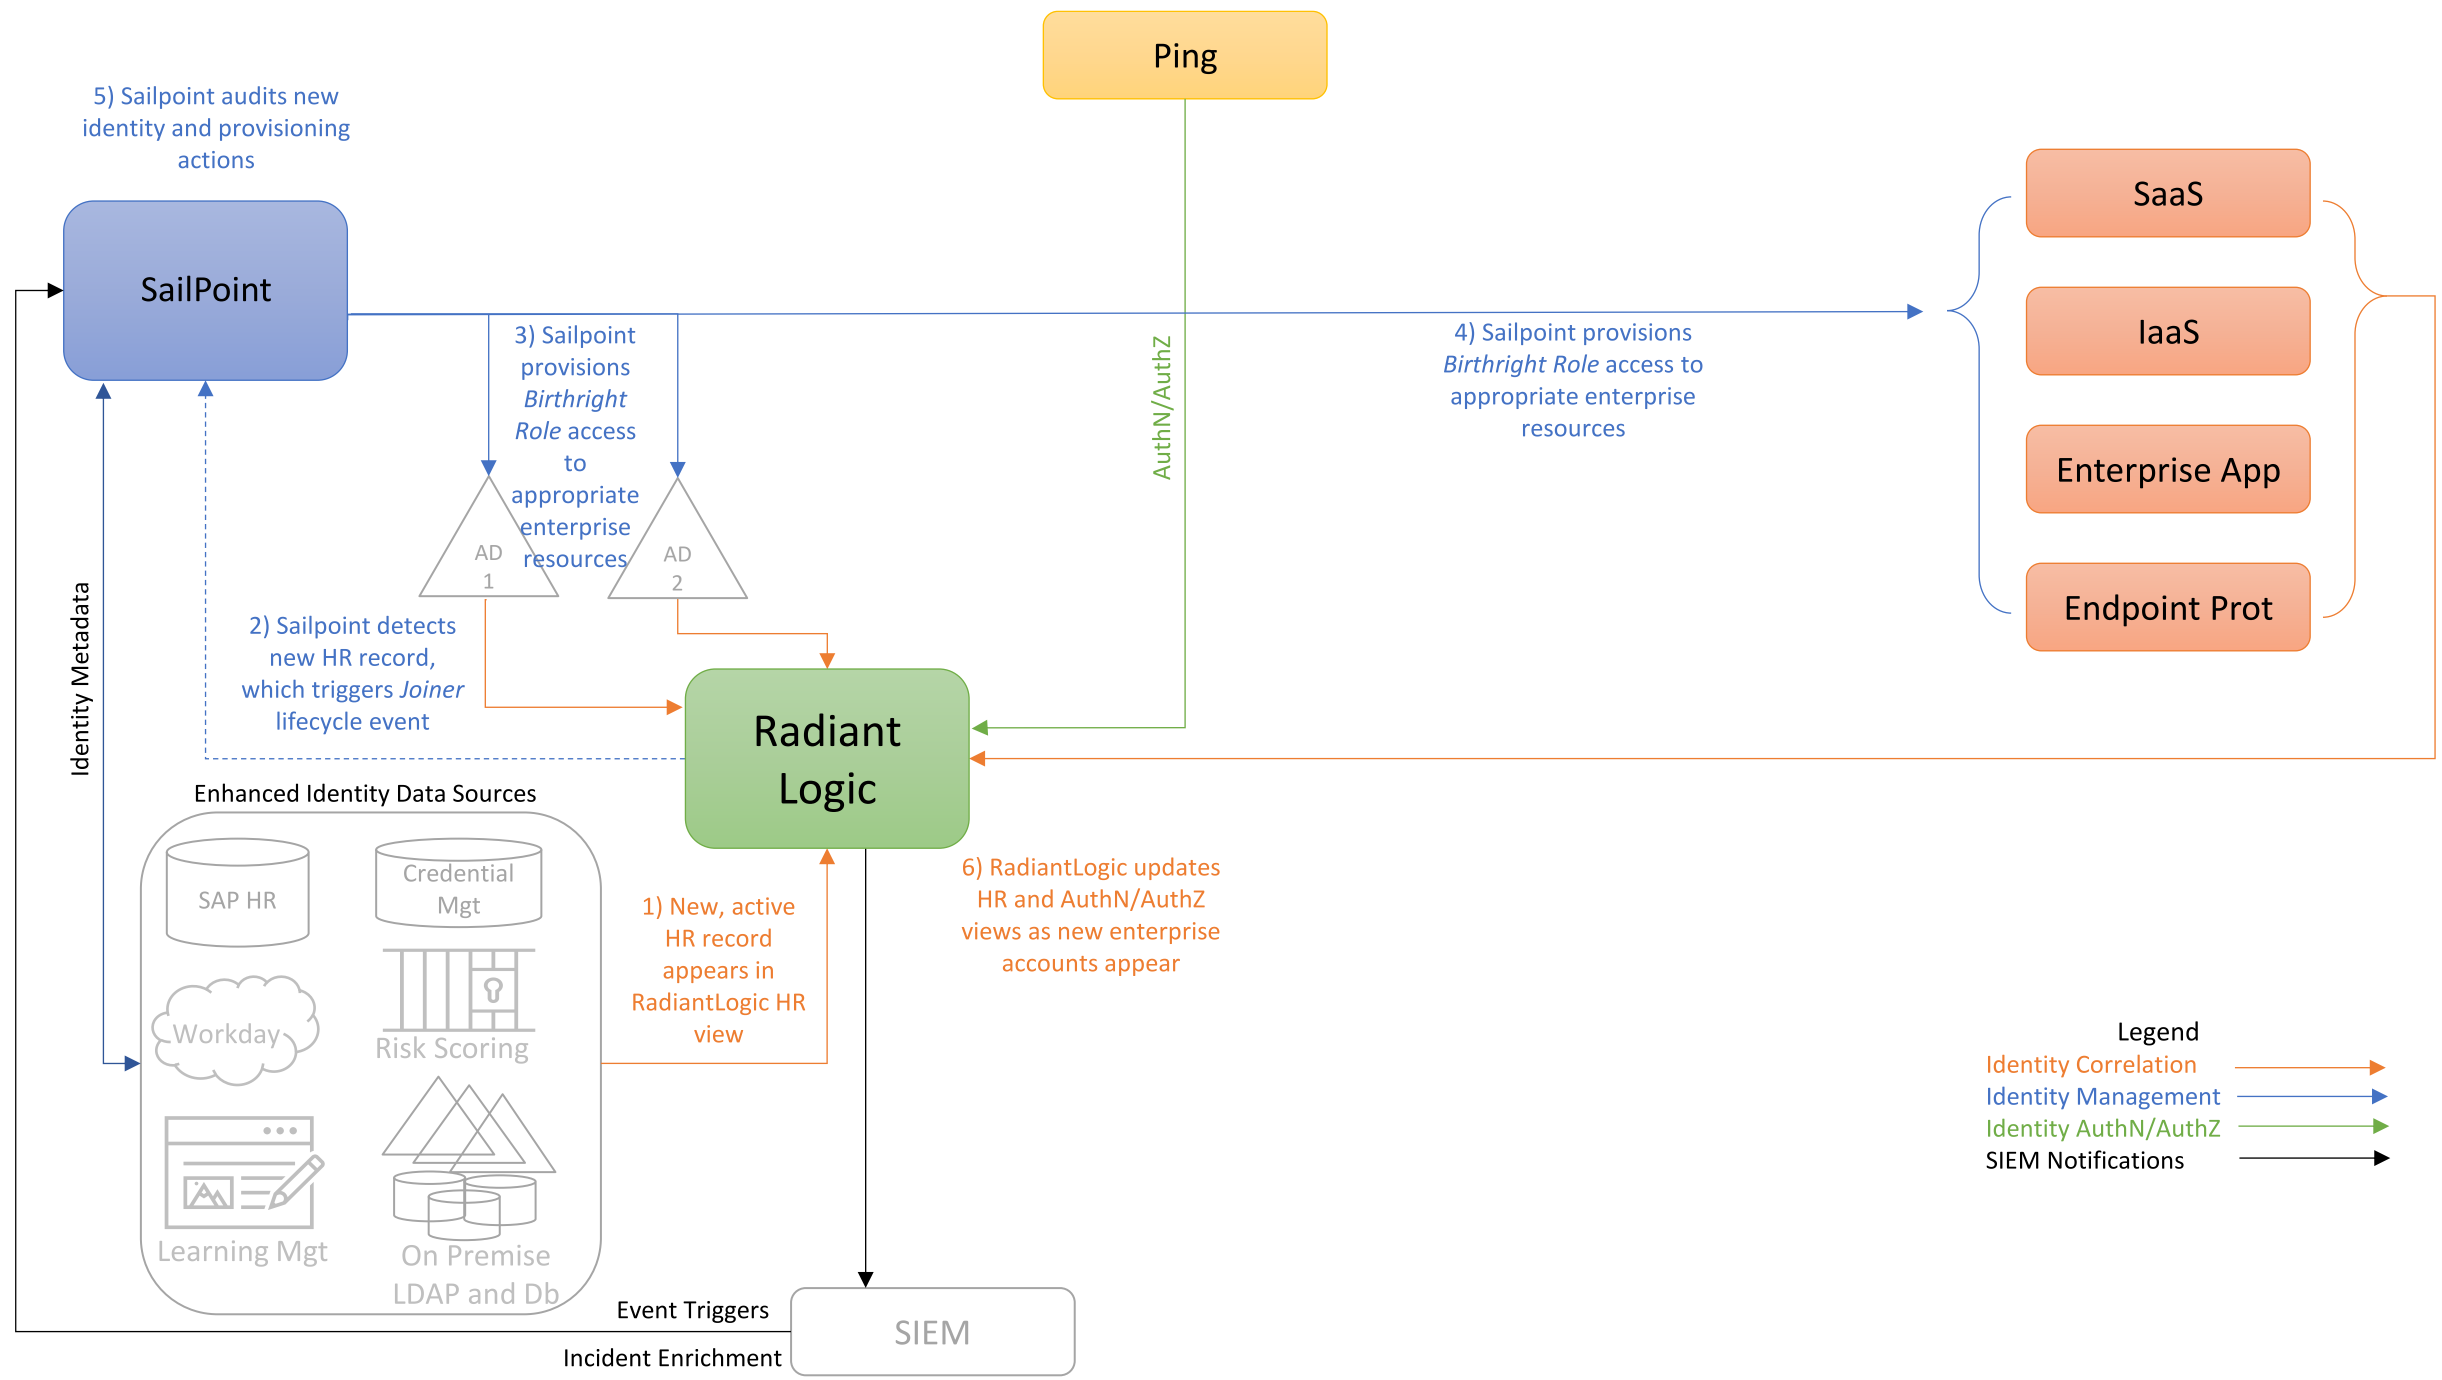

E.2.2.3 User Changes Roles

Figure E-4 depicts the ICAM information architecture for E2B1, showing the steps required to remove some access privileges and add other access privileges for a user in response to that user changing roles within the enterprise. The steps are as follows:

1. When a user changes roles within the enterprise, an authorized HR staff member is assumed to input information into some sort of enterprise employee management application that will result in the Radiant Logic HR record for that user indicating that the user has changed roles.
2. SailPoint detects this updated HR record in Radiant Logic. SailPoint evaluates this updated HR record, which triggers a *Mover* lifecycle event, causing SailPoint to execute a policy-driven workflow that includes steps 3, 4, 5, and 6.
3. SailPoint removes access permissions associated with the user's prior role (but not with the user's new role) from the user's AD account and removes access from other enterprise resources (e.g., SaaS, IaaS, enterprise applications, MDM) that the user had been authorized to access as a result of their prior role but is not authorized to access as a result of their new role. Also (not labeled in the diagram), Radiant Logic then collects and correlates any changes that were made to the user's account from AD into the global identity profile that it is maintaining.
4. Assuming there are enterprise resources that the user's new role entitles them to access that are not managed by AD, SailPoint provisions access to these resources for the user by creating new accounts for the user, as appropriate, in SaaS, IaaS, enterprise application, endpoint protection, MDM, and other components. (This provisioning may occur directly or via Ping.)
5. SailPoint generates an access review for management to confirm or revoke the changes that have been made. Such an access review is not strictly necessary. The permission changes could be executed in a fully automated manner, if desired, and specified by policy. However, having an access review provides management with the opportunity to exercise some supervisory discretion to permit the user to temporarily continue to have access to some resources associated with their former role that may still be needed.
6. Once the access review has been completed and any access privilege changes deemed necessary have been performed, SailPoint audits the changes.
7. As the new enterprise accounts appear in the SaaS, IaaS, enterprise application, endpoint protection, and other components, and as existing account access is removed, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization views to reflect the recent changes. Ping will eventually query these authentication and authorization information views in Radiant Logic to determine whether to grant future user access requests. (RadiantOne stores all user identity information;

Ping does not store any user data. When Ping needs user identity data, it looks up this information directly from RadiantOne.)

3170 **Figure E-4 E2B1 ICAM Information Architecture - User Changes Roles**

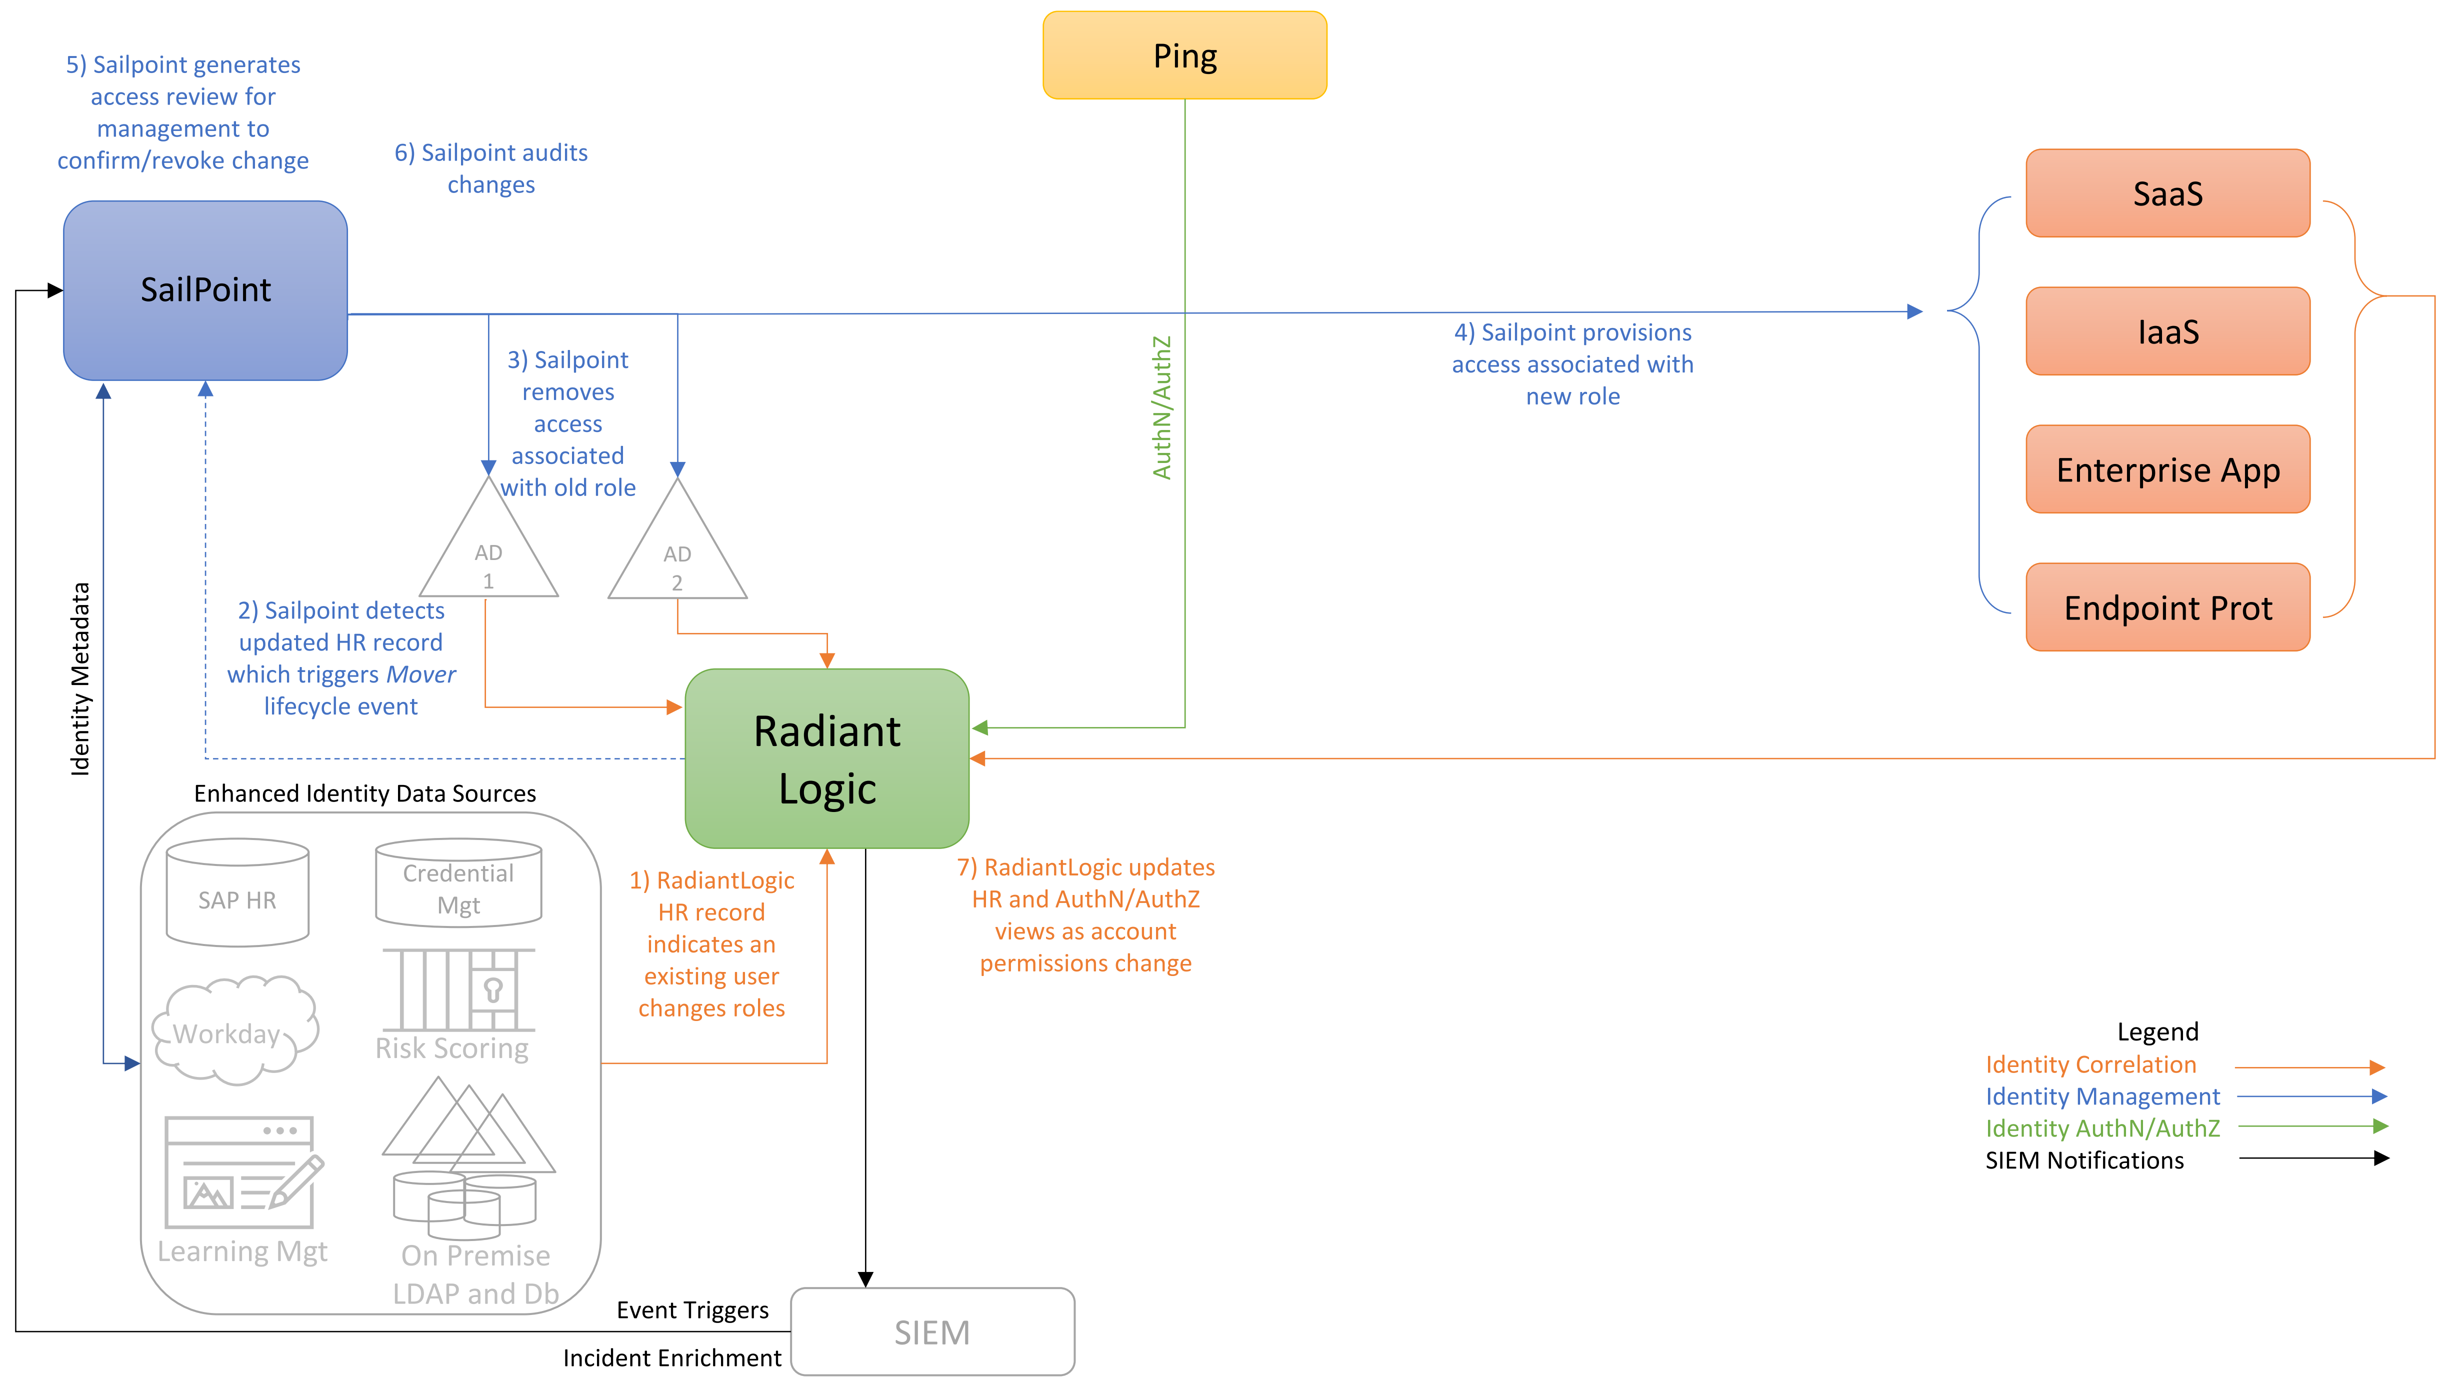

E.2.2.4 User Leaves the Enterprise

Figure E-5 depicts the ICAM information architecture for E2B1, showing the steps required to disable or delete an identity and remove access privileges in response to a user leaving the enterprise. The steps are as follows:

8. When a user's employment is terminated, an authorized HR staff member is assumed to input information into some sort of enterprise employee management application that will result in the Radiant Logic HR record for that user indicating that the user has changed from active to inactive status.
9. SailPoint detects this updated HR record in Radiant Logic. SailPoint evaluates this updated HR record, which triggers a *Leaver* lifecycle event, causing SailPoint to execute a policy-driven workflow that includes steps 3, 4, 5, and 6.
10. SailPoint removes all access permissions associated with the user identity from AD. Also (not labeled in the diagram), Radiant Logic then collects and correlates this user access authorization change from AD into the global identity profile that it is maintaining.
11. SailPoint either disables or deletes all enterprise resource accounts associated with the user identity, as defined by policy, from components such as SaaS, IaaS, enterprise applications, and endpoint protection platforms. (SailPoint may perform these actions directly or via Ping.)
12. SailPoint removes the user identity from all governance groups the identity is in.
13. SailPoint audits the changes made as a result of this user termination.
14. As the enterprise accounts associated with the user's identity are deleted or disabled, Radiant Logic is notified. Radiant Logic collects, correlates, and virtualizes this new identity information and adds it back into the global identity profile that it is maintaining. It also updates its HR, authentication, and authorization views to reflect the recent changes. Ping will eventually query these authentication and authorization information views in Radiant Logic to determine whether or not to grant future user access requests. (RadiantOne stores all user identity information; Ping does not store any user data. When Ping needs user identity data, it looks up this information directly from RadiantOne.)

3198 **Figure E-5 E2B1 ICAM Information Architecture - User Termination**

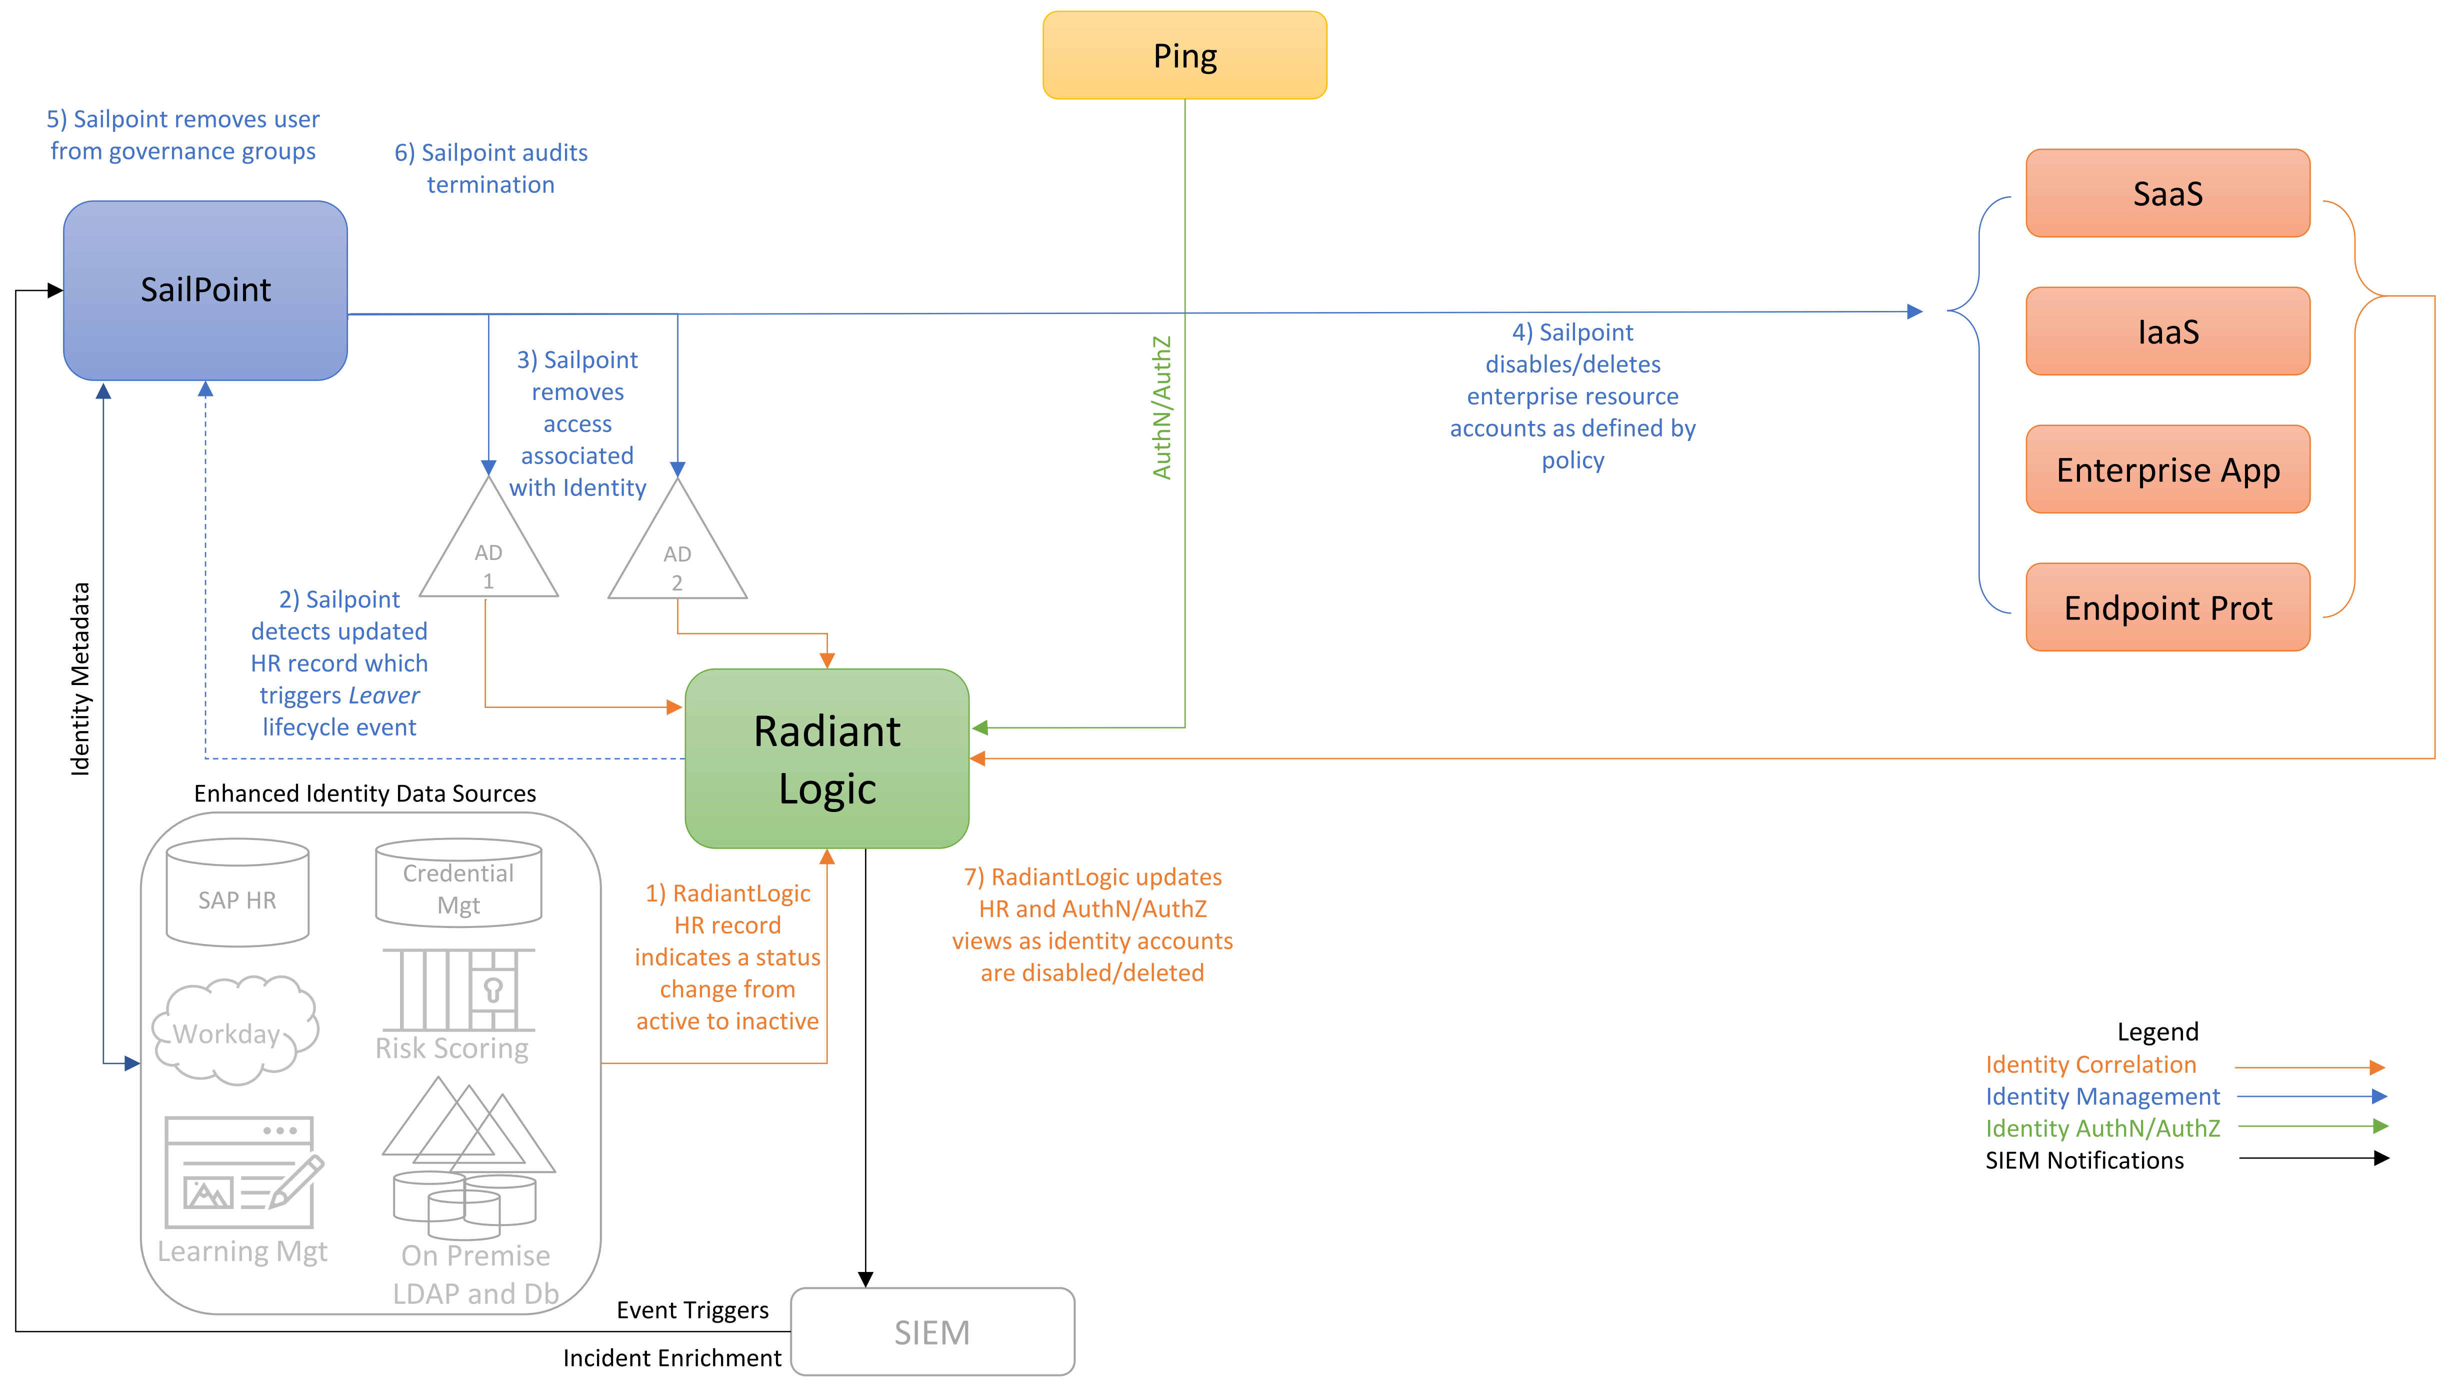

E.2.3 Physical Architecture

Section 4.4.3 describes the physical architecture of the E2B1 network.

E.2.4 Message Flow for a Successful Resource Access Request

Below is depicted the high-level message flow supporting the use case in which a subject who has an enterprise ID, who is located on-premises, and who is authorized to access an enterprise resource, requests and receives access to that resource. In the case depicted here, access to the resource is protected by PingFederate, which acts as a PDP and an identity provider; Cisco Duo, which consists of an agent on the endpoint and a cloud component that work together to perform second-factor user authentication and also to gather device health information to ensure device compliance; and Radiant Logic, which performs credential validation for authentication and provides granular user-relevant attributes and groups for authorization at the request of PingFederate.

The message flow depicted in Figure E-6 shows only the messages that are sent in response to the access request. However, the authentication process also relies on the following additional background communications that occur among components on an ongoing basis:

- The Cisco Duo endpoint agent periodically syncs with the Cisco Duo cloud component to reauthenticate the requesting endpoint device using a unique certificate that has been provisioned specifically for that device and sends the cloud component information about device health (e.g., firewall running, anti-malware software, IoS version).
- Cisco Duo is integrated with PingFederate and periodically sends PingFederate assurance that, based on the device health information collected by Cisco Duo, the device is compliant with configured policy.

Figure E-6 depicts the message flow for the user's request to access the resource.

Figure E-6 Use Case—E2B1 – Access Enforced by Ping Federate, Cisco Duo, and Radiant Logic

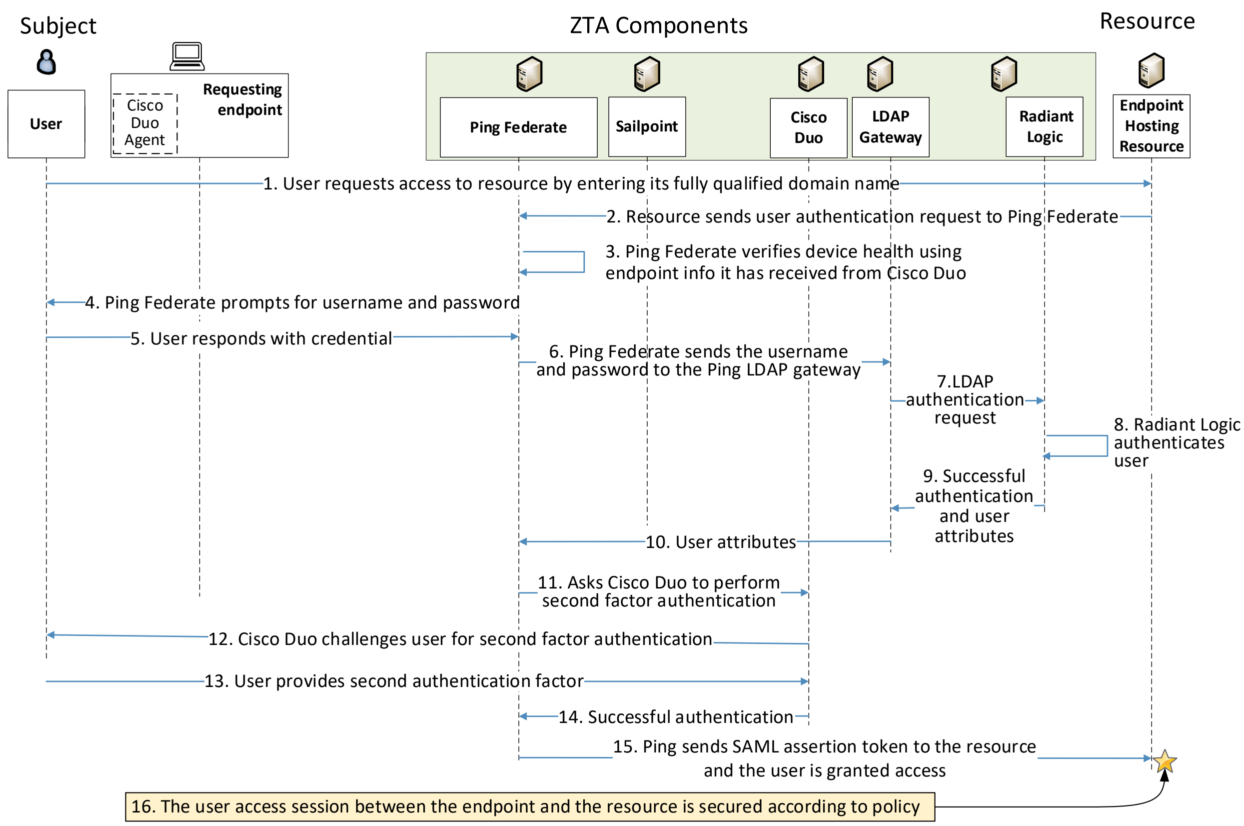

The message flow depicted in Figure E-6 consists of the following steps:

1. A user requests to access a resource by typing the resource's URL into a browser.
2. The resource receives the access request and sends a user authentication request to PingFederate.
3. PingFederate consults the device health information it has received in the background from Cisco Duo verifying that the device has been authenticated and is compliant with policy.
4. PingFederate prompts for username and password.
5. The user responds with username and password.
6. PingFederate sends the user's username and password to the Ping LDAP Gateway to facilitate communication between the cloud-hosted Ping and the on premises Radiant Logic resources.
7. The LDAP gateway forwards the LDAP authentication request to Radiant Logic.

8. Radiant Logic authenticates that the username exists in the master user record and the provided password (credential) is valid based on credentials stored in Radiant Logic or in another source of identity credentials federated by Radiant Logic.
9. Radiant Logic replies to the LDAP gateway with a valid BIND indicating a successful user authentication and all additional user attributes requested by Ping at the time of Authentication
10. The LDAP gateway forwards the response from Radiant Logic to PingFederate with the successful BIND and applicable user's attributes.
11. PingFederate requests Cisco Duo to perform second-factor user authentication.
12. Cisco Duo challenges the user to provide the second authentication factor.
13. The user responds with the second authentication factor.
14. Cisco Duo responds to PingFederate, indicating that the user authenticated successfully.
15. PingFederate sends a SAML assertion token to the resource. The resource accepts the assertion and grants the access request.
16. User traffic to and from the resource is secured according to policy (e.g., using TLS or HTTPS).

Note that the message flow depicted in Figure E-6 applies to several of the use cases we are considering. It applies to all cases in which a user with an enterprise ID who can successfully authenticate themselves and who is using an enterprise-owned endpoint requests and receives access to an enterprise resource that they are authorized to access. The message flow is the same regardless of whether the employee is located on-premises at headquarters, on-premises at a branch office, or off-premises at home or elsewhere. It is also the same regardless of whether the resource is located on-premises or in the cloud.

Appendix F EIG Enterprise 3 Build 1 (E3B1)

F.1 Technologies

EIG E3B1 uses products from F5, Forescout, Lookout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used. For more information on these collaborators and the products and technologies that they contributed to this project overall, see Section 3.4.

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

Table F-1 lists all of the technologies used in E3B1 ZTA. It lists the products used to instantiate each ZTA component and the security function that the component provides.

Table F-1 E3B1 Products and Technologies

Component	Product	Function
PE	Azue AD (Conditional Access)	Decides whether to grant, deny, or revoke access to a resource based on enterprise policy, information from supporting components, and a trust algorithm.
PA	Azure AD (Conditional Access)	Executes the PE's policy decision by sending commands to a PEP that establishes and shuts down the communication path between subject and resource.
PEP	Azure AD (Conditional Access), F5 BIG-IP, and Lookout MES	Guards the trust zone that hosts one or more enterprise resources; establishes, monitors, and terminates the connection between subject and resource as directed by the PA; forwards requests to and receives commands from the PA.
Identity Management	Microsoft AD and Azure AD	Creates and manages enterprise user and device accounts, identity records, role information, and access attributes that form the basis of access decisions within an organization to ensure the correct subjects have the appropriate access to the correct resources at the appropriate time.
Access & Credential Management	Microsoft AD and Azure AD	Manages access to resources by performing user and device authentication (e.g., SSO and MFA) and using identity, role, and access attributes to determine which access requests are authorized.

Component	Product	Function
Federated Identity	Microsoft AD and Azure AD	Aggregates and correlates all attributes relating to an identity or object that is being authorized by a ZTA. It enables users of one domain to securely access data or systems of another domain seamlessly, and without the need for completely redundant user administration. Federated identity encompasses the traditional ICAM data, supports identities that may be part of a larger federated ICAM community, and may include non-enterprise employees.
Identity Governance	Microsoft AD and Azure AD	Provides policy-based, centralized, automated processes to manage user identity and access control functions (e.g., ensuring segregation of duties, role management, logging, access reviews, analytics, reporting) to ensure compliance with requirements and regulations.
MFA	Azure AD (Multifactor Authentication)	Authenticates user identity by requiring the user to provide not only something they know (e.g., a password), but also something they have (e.g., a token).
UEM/MDM	Microsoft Intune	Manages and secures enterprise desktop computers, laptops, and/or mobile devices in accordance with enterprise policy to protect applications and data; ensure device compliance; mitigate and remediate vulnerabilities and threats; monitor for suspicious activity to prevent and detect intrusions; prevent, detect, and disable malware and other malicious or unauthorized traffic; repair infected files when possible; provide alerts and recommend remediation actions; and encrypt data. Pushes enterprise applications and updates to devices, enables users to download enterprise applications that they are authorized to access, remotely deletes all applications and data from devices if needed, tracks user activity on devices, and detects and addresses security issues on the device.

Component	Product	Function
EPP	Microsoft Defender for Endpoint, Lookout MES, PC Matic Pro	Detects and stops threats to endpoints through an integrated suite of endpoint protection technologies including antivirus, data encryption, intrusion prevention, EDR, and DLP. May include mechanisms that are designed to protect applications and data; ensure device compliance with policies regarding hardware, firmware, software, and configuration; monitor endpoints for vulnerabilities, suspicious activity, intrusion, infection, and malware; block unauthorized traffic; disable malware and repair infections; manage and administer software and updates; monitor behavior and critical data; and enable endpoints to be tracked, troubleshooted, and wiped, if necessary.
SIEM	Microsoft Sentinel	Collects and consolidates security information and security event data from many sources; correlates and analyzes the data to help detect anomalies and recognize potential threats and vulnerabilities; and logs the data to adhere to data compliance requirements.
Vulnerability Scanning and Assessment	Tenable.io and Tenable.ad	Scans and assesses the enterprise infrastructure and resources for security risks; identifies vulnerabilities and misconfigurations; and provides remediation guidance regarding investigating and prioritizing responses to incidents.
Security Validation	Mandiant MSV	Provides visibility and evidence on the status of the security controls' effectiveness in the ZTA. Enables security capabilities of the enterprise to be monitored and verified by continuously validating and measuring the cybersecurity controls; also used to automate the demonstrations that were performed to showcase ZTA capabilities. Mandiant MSV is deployed throughout the project's laboratory environment to enable monitoring and verification of various security aspects of the builds. 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.
Network Discovery	Forescout eyeSight	Discovers, classifies, and assesses the risk posed by devices and users on the network.
Remote Connectivity	Palo Alto Networks NGFW	Enables authorized remote users to securely access the inside of the enterprise. (Once inside, the ZTA manages the user's access to resources.)

Component	Product	Function
Certificate Management	DigiCert CertCentral TLS Manager	Provides automated capabilities to issue, install, inspect, revoke, renew, and otherwise manage TLS certificates.
Cloud IaaS	Azure	Provides computing resources, complemented by storage and networking capabilities, hosted by a cloud service provider, offered to customers on demand, and exposed through a GUI and an API.
Cloud SaaS	Digicert CertCentral, Lookout MES, Microsoft Azure AD, Microsoft Defender for Endpoint, Microsoft Intune, Microsoft Office 365, Microsoft Sentinel, and Tenable.io,	Cloud-based software delivered for use by the enterprise.
Application	GitLab	Example enterprise resource to be protected. (In this build, GitLab is integrated directly with Azure AD using SAML, and Microsoft Sentinel pulls logs from GitLab.)
Application	Guacamole	Example enterprise resource to be protected. (In this build, BIG-IP serves as an identity-aware proxy that protects access to Guacamole, and BIG-IP is integrated with Azure AD using SAML. Also, Microsoft Sentinel pulls logs from Guacamole.)
Enterprise-Managed Device	Windows client, macOS client, and mobile devices (iOS and Android)	Example endpoints to be protected. (In this build, all enterprise-managed devices are enrolled into Microsoft Intune.)
BYOD	Windows client, macOS client, and mobile devices (iOS and Android)	Example endpoints to be protected.

F.2 Build Architecture

In this section we present the logical architecture of E3B1 relative to how it instantiates the crawl phase EIG reference architecture depicted in Figure 4-2. We also describe E3B1's physical architecture and present message flow diagrams for some of its processes.

F.2.1 Logical Architecture

Figure F-1 depicts the logical architecture of E3B1. Figure F-1 uses numbered arrows to depict the general flow of messages needed for a subject to request access to a resource and have that access request evaluated based on subject identity (both requesting user and requesting endpoint identity), authorizations, and requesting endpoint health. It also depicts the flow of messages supporting periodic reauthentication of the requesting user and the requesting endpoint and periodic verification of requesting endpoint health, all of which must be performed to continually reevaluate access. The labeled steps in Figure F-1 have the same meanings as they do in Figure 4-1 and Figure 4-2. However, while Figure 4-2 depicts generic crawl phase ZTA components, Figure F-1 includes the specific products that instantiate the architecture of E3B1. Figure F-1 also does not depict any of the resource management steps found in Figure 4-1 and Figure 4-2 because the ZTA technologies deployed in E3B1 do not support the ability to perform authentication and reauthentication of the resource or periodic verification of resource health.

E3B1 was designed with a single ICAM system (Microsoft Azure AD) that serves as identity, access, and credential manager and also serves as the ZTA PE and PA. It includes three PEPs: Microsoft Azure AD, F5 BIG-IP, and Lookout MES. A more detailed depiction of the messages that flow among components to support user access requests in the two different cases when the resource is being protected by the Azure AD PEP versus the F5 BIG-IP PEP can be found in Appendices F.2.3.1 and F.2.3.2.

Figure F-1 Logical Architecture of E3B1

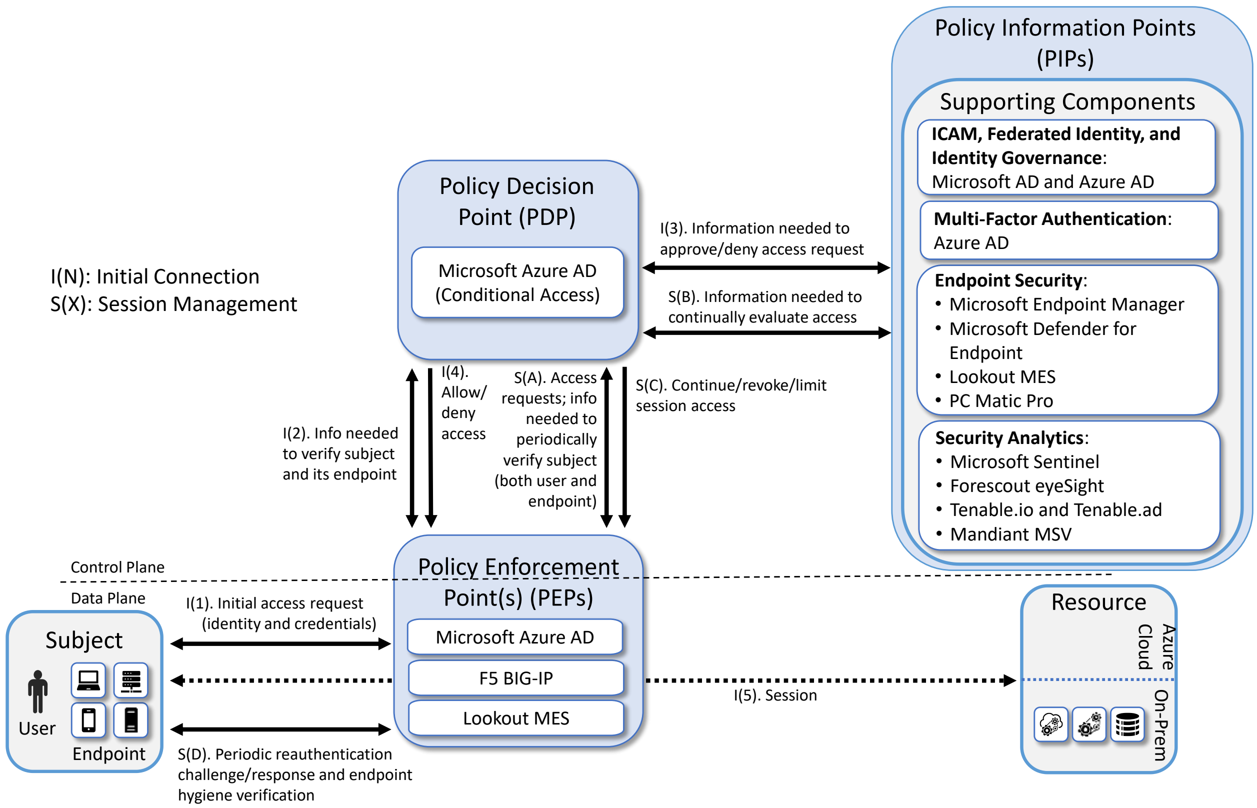

F.2.2 Physical Architecture

Section 4.4.4 describes the physical architecture of the E3B1 network.

F.2.3 Message Flows for a Successful Resource Access Request

This section depicts two high-level message flows, both of which support the use case in which a subject who has an enterprise ID, is located on-premises, and is authorized to access an enterprise resource, requests and receives access to that resource.

The two message flows that are supported by Enterprise 3 for this use case depend on whether the resource being accessed is protected by Azure AD alone (see Appendix F.2.3.1) or by Azure AD in conjunction with the F5 BIG-IP PEP (see Appendix F.2.3.2).

Regardless of which components are being used to protect the resource, all endpoints are enrolled into Microsoft Intune, which is an MDM (and a UEM) that can configure and manage devices and can also retrieve and report on device security settings that can be used to determine compliance, such as whether the device is running a firewall or anti-malware. Non-Windows devices have an MDM agent installed on them to enable them to report compliance information to Microsoft Intune, but Windows devices do not require a separate agent because Windows has built-in agents that are designed to communicate with Intune. Intune-enrolled devices check in with Intune periodically, allowing it to authenticate the requesting endpoint, determine how the endpoint is configured, modify certain configurations, and collect much of the information it needs to determine whether the endpoint is compliant. Intune reports the device compliance information that it collects to Azure AD, which will not permit a device to access any resources unless it is compliant.

For demonstration purposes, one of the criteria that devices are expected to meet to be considered compliant in our example implementation is that they must have antivirus software updated and running. In both scenarios below, some requesting endpoints have Microsoft Defender Antivirus running on them and other requesting endpoints have PC Matic Pro (also antivirus software) running; no endpoints have both turned on. If a device is running Microsoft Defender Antivirus, the Intune MDM can sense this and report it to Azure AD. If a device is running PC Matic Pro, however, the device is configured to notify Windows Security Center that the endpoint has antivirus software installed, and the Security Center provides this information to Azure AD.

The authentication message flows depicted below show only the messages that are sent in response to the access request. However, the authentication process also relies on the following additional background communications that occur among components on an ongoing basis:

- Microsoft AD periodically synchronizes with Azure AD to provide it with the most up-to-date identity information.
- Intune-enrolled devices check in with Intune periodically. Checking in allows Intune to determine how the endpoint is configured and modify certain configurations that have been previously specified. It also allows Intune to report the compliance of the device to Azure AD.
- Microsoft Defender for Endpoint has both a cloud component and built-in sensors that detect threat signals from Windows endpoints. So not only can it tell that a firewall is disabled or antivirus is off, but it can tell when certain malicious signals seen elsewhere have also been observed on your endpoint. It periodically reports this information to its cloud/management component, which uses it for risk determination. This information can be passed off to Intune to include in its compliance determination of an endpoint.
- Microsoft Defender Antivirus (an endpoint agent) periodically syncs with Microsoft Intune and Microsoft Defender for Endpoint.
- Microsoft Intune periodically sends device health information to Azure AD so that it can be sure that the device is managed and compliant.
- PC Matic periodically syncs with Windows Security Center to inform it that that the endpoint has antivirus installed and active.
- Windows Security Center periodically syncs with Azure AD to provide it with endpoint status information, e.g., that endpoints have antivirus installed.

F.2.3.1 Use Case in which Resource Access Is Enforced by Azure AD

Figure F-2 depicts the message flow for the case in which access to the resource is protected by Azure AD (with the Conditional Access feature), which acts as a PDP; and Microsoft AD, which provides identity information.

Figure F-2 Use Case—E3B1 – Access Enforced by Azure AD

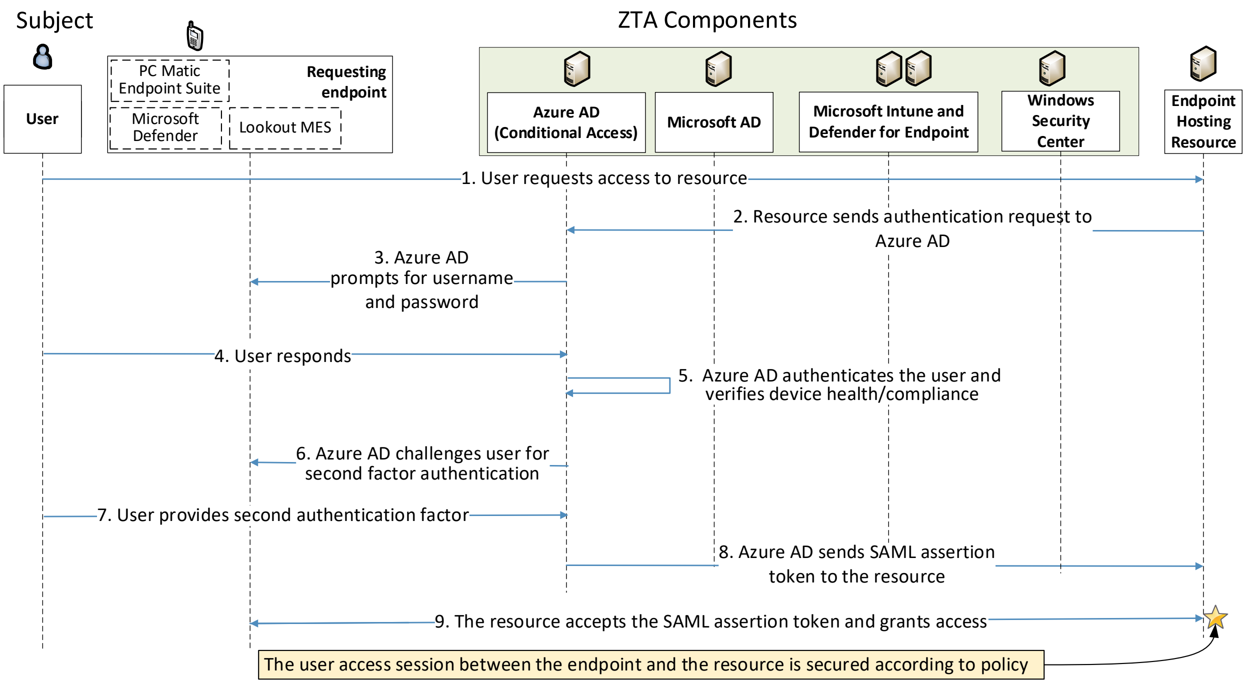

The message flow depicted in Figure F-2 consists of the following steps:

1. A user requests access to a resource.
2. The resource sends the authentication request to Azure AD.
3. Azure AD prompts for username and password.
4. The user responds with username and password.
5. Azure AD authenticates the user. Azure AD consults the information about the device that it has received in the background from Microsoft Intune and Defender for Endpoint to authenticate the device and verify that it is managed and meets compliance requirements. If the device has PC Matic running on it, Azure AD also consults information about the device that it has received in the background from Windows Security Center to verify that the device is running antivirus software.

6. Azure AD challenges the user to provide the second authentication factor.
7. The user responds with the second authentication factor.
8. Azure AD sends a SAML assertion to the resource.
9. The resource accepts the assertion and grants the access request. User traffic to and from the resource is secured according to policy (e.g., using TLS or HTTPS).

F.2.3.2 Use Case in which Resource Access Is Enforced by an F5 BIG-IP PEP

Figure F-3 depicts the message flow for the case in which access to the resource is protected by F5 BIG-IP, which acts as an identity-aware proxy PEP; Microsoft Azure AD, which acts as an ICAM provider and PDP; and Microsoft AD, which provides identity information.

Figure F-3 Use Case—E3B1 – Access Enforced by F5 BIG-IP

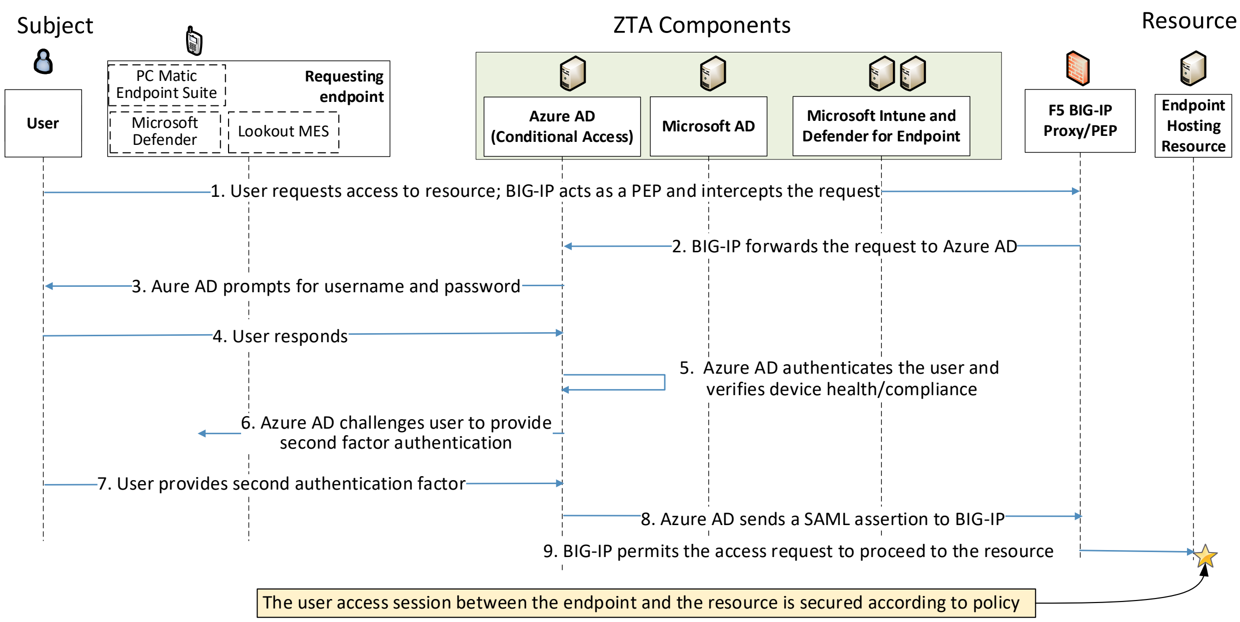

The message flow depicted in Figure F-3 consists of the following steps:

1. A user requests access to a resource.
2. BIG-IP, which is acting as an identity-aware proxy PEP that sits in front of the resource, intercepts and forwards the request to Azure AD.
3. Azure AD prompts for username and password.
4. The user responds with username and password.

5. Azure AD authenticates the user. Azure AD consults the information about the device that it has received in the background from Microsoft Intune and Defender for Endpoint to authenticate the device and verify that it is managed and meets compliance requirements. If the device has PC Matic running on it, Azure AD also consults information about the device that it has received in the background from Windows Security Center to verify that the device is running antivirus software.
6. Azure AD challenges the user to provide the second authentication factor.
7. The user responds with the second authentication factor.
8. Azure AD sends a SAML assertion to BIG-IP which serves as an identity-aware proxy, service provider, and the PEP protecting the resource.
9. BIG-IP accepts the SAML assertion and permits the access request to proceed to the resource. User traffic to and from the resource is secured according to policy (e.g., using TLS or HTTPS).

Appendix G EIG Enterprise 4 Build 1 (E4B1)

This build will be documented in a future version of this publication.

Appendix H EIG Enterprise 1 Build 2 (E1B2)

H.1 Technologies

EIG E1B2 uses products from Amazon Web Services, IBM, Ivanti, Mandiant, Okta, Radiant Logic, SailPoint, Tenable, and Zscaler. Certificates from DigiCert are also used. For more information on these collaborators and the products and technologies that they contributed to this project overall, see Section 3.4.

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

Table H-1 lists all of the technologies used in EIG E1B2. It lists the products used to instantiate each ZTA component and the security function that each component provides.

Table H-1 E1B2 Products and Technologies

Component	Product	Function
PE	Zscaler ZPA Central Authority (CA)	Decides whether to grant, deny, or revoke access to a resource based on enterprise policy, information from supporting components, and a trust algorithm.
PA	Zscaler ZPA Admin Portal and ZPA CA	Executes the PE's policy decision by sending commands to a PEP that establishes and shuts down the communication path between subject and resource.
PEP	Zscaler Public Service Edges	Guards the trust zone that hosts one or more enterprise resources; establishes, monitors, and terminates the connection between subject and resource as directed by the PA; forwards requests to and receives commands from the PA.
Identity Management	Okta Identity Cloud	Creates and manages enterprise user and device accounts, identity records, role information, and access attributes that form the basis of access decisions within an organization to ensure the correct subjects have the appropriate access to the correct resources at the appropriate time.
Access & Credential Management	Okta Identity Cloud	Manages access to resources by performing user and device authentication (e.g., SSO and MFA) and using identity, role, and access attributes to determine which access requests are authorized.

Component	Product	Function
Federated Identity	Radiant Logic RadiantOne Intelligent Identity Data Platform	Aggregates and correlates all attributes relating to an identity or object that is being authorized by a ZTA. It enables users of one domain to securely access data or systems of another domain seamlessly, and without the need for completely redundant user administration. Federated identity encompasses the traditional ICAM data, supports identities that may be part of a larger federated ICAM community, and may include non-enterprise employees.
Identity Governance	SailPoint IdentityIQ	Provides policy-based, centralized, automated processes to manage user identity and access control functions (e.g., ensuring segregation of duties, role management, logging, access reviews, analytics, reporting) to ensure compliance with requirements and regulations.
MFA	Okta Verify app	Supports MFA of a user identity by requiring the user to provide not only something they know (e.g., a password), but also something they have (e.g., a token).
UEM/MDM	Ivanti Neurons for Unified Endpoint Management (UEM) Platform	Manages and secures enterprise desktop computers, laptops, and/or mobile devices in accordance with enterprise policy to protect applications and data; ensure device compliance; mitigate and remediate vulnerabilities and threats; monitor for suspicious activity to prevent and detect intrusions; prevent, detect, and disable malware and other malicious or unauthorized traffic; repair infected files when possible; provide alerts and recommend remediation actions; and encrypt data. Pushes enterprise applications and updates to devices, enables users to download enterprise applications that they are authorized to access, remotely deletes all applications and data from devices if needed, tracks user activity on devices, and detects and addresses security issues on the device.

Component	Product	Function
EPP	None	Detects and stops threats to endpoints through an integrated suite of endpoint protection technologies including antivirus, data encryption, intrusion prevention, EDR, and DLP. May include mechanisms that are designed to protect applications and data; ensure device compliance with policies regarding hardware, firmware, software, and configuration; monitor endpoints for vulnerabilities, suspicious activity, intrusion, infection, and malware; block unauthorized traffic; disable malware and repair infections; manage and administer software and updates; monitor behavior and critical data; and enable endpoints to be tracked, troubleshooted, and wiped, if necessary.
Endpoint Compliance	Zscaler Client Connector	Has capabilities to enforce policies based on a defined set of endpoint compliance checks to allow or deny user/endpoint access to a resource, but does not perform the functions of an EPP solution to automatically remediate an endpoint.
SIEM	IBM Security QRadar XDR	Collects and consolidates security information and security event data from many sources; correlates and analyzes the data to help detect anomalies and recognize potential threats and vulnerabilities; and logs the data to adhere to data compliance requirements.
Vulnerability Scanning and Assessment	Tenable.io and Tenable.ad	Scans and assesses the enterprise infrastructure and resources for security risks, identifies vulnerabilities and misconfigurations, and provides remediation guidance regarding investigating and prioritizing responses to incidents.
Network Discovery	Tenable NNM	Discovers, classifies, and assesses the risk posed by devices and users on the network.
Security Integration Platform	IBM Cloud Pak for Security	Integrates the SIEM and other security tools into a single pane of glass to support generation of insights into threats and help track, manage, and resolve cybersecurity incidents. Executes predefined incident response workflows to automatically analyze information and orchestrate the operations required to respond.

Component	Product	Function
Security Validation	Mandiant MSV	Provides visibility and evidence on the status of the security controls' effectiveness in the ZTA. Enables security capabilities of the enterprise to be monitored and verified by continuously validating and measuring the cybersecurity controls; also used to automate the demonstrations that were performed to showcase ZTA capabilities. Deployed throughout the project's laboratory environment to enable monitoring and verification of various security aspects of the builds. 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.
Remote Connectivity	Zscaler ZPA Zscaler ZIA	ZPA is used to provide remote users' connectivity to on-premises resources. To support remote users' connectivity to resources in IaaS, ZPA is used for private applications and ZIA is used for public-facing applications.
Application Connector	Zscaler Application Connector	Component that is deployed to be the front-end for an internal resource (whether located on-premises or in the cloud) and act as a proxy for it. Requests to access the resource are directed to the connector, which responds by initiating a secure connection to the PEP. A connector enables access to a resource to be controlled without requiring the resource to be visible on the network.
Certificate Management	DigiCert CertCentral TLS Manager	Provides automated capabilities to issue, install, inspect, revoke, renew, and otherwise manage TLS certificates.
Cloud IaaS	AWS - GitLab, WordPress	Provides computing resources, complemented by storage and networking capabilities, hosted by a cloud service provider, offered to customers on demand, and exposed through a GUI and an API. An IPsec tunnel is used to provide a secure connection from the enterprise to the cloud.
Cloud SaaS	Digicert CertCentral, Ivanti Neurons for UEM, Okta Identity Cloud, Tenable.io, Zscaler ZPA, and Zscaler ZIA	Cloud-based software delivered for use by the enterprise.

Component	Product	Function
Application	On-premises - GitLab	Example enterprise resource to be protected. (In this build, GitLab is integrated with Okta using SAML, and IBM Security QRadar XDR pulls logs from GitLab.)
Enterprise-Managed Device	Mobile devices (iOS and Android) and desktops/laptops (Windows and Mac)	Example endpoints to be protected. All enterprise-managed mobile devices are running an Ivanti Neurons for UEM agent and also have the Okta Verify App installed. If Ivanti Neurons for UEM agent is used to push Zscaler Client Connector (ZCC) to the endpoint, that endpoint is considered to be a managed device.
BYOD	Mobile devices (iOS and Android) and desktops/laptops (Windows and Mac)	Example endpoints to be protected.

H.2 Build Architecture

In this section we present the logical architecture of E1B2. We also describe E1B2's physical architecture and present message flow diagrams for some of its processes.

H.2.1 Logical Architecture

Figure H-1 depicts the logical architecture of E1B2. Figure H-1 uses numbered arrows to depict the general flow of messages needed for a subject to request access to a resource and have that access request evaluated based on subject identity (both requesting user and requesting endpoint identity), user authorizations, and requesting endpoint health. It also depicts the flow of messages supporting periodic reauthentication of the requesting user and the requesting endpoint and periodic verification of requesting endpoint health, all of which must be performed to continually reevaluate access. The labeled steps in Figure H-1 have the same meanings as they do in Figure 4-1. However, Figure H-1 includes the specific products that instantiate the architecture of E1B2. Figure H-1 also does not depict any of the resource management steps found in Figure 4-1 because the ZTA technologies deployed in E1B2 do not support the ability to perform authentication and reauthentication of the resource or periodic verification of resource health.

E1B2 was designed with Zscaler components that serve as the PE, PA, and PEP, and Okta Identity Cloud that serves as the identity, access, and credential manager. Radiant Logic acts as a PIP for the PDP as it responds to inquiries and provides identity information on demand in order for Okta to make near-real-time access decisions. A more detailed depiction of the messages that flow among components to support a user access request can be found in Appendix H.2.4.

Figure H-1 Logical Architecture of E1B2

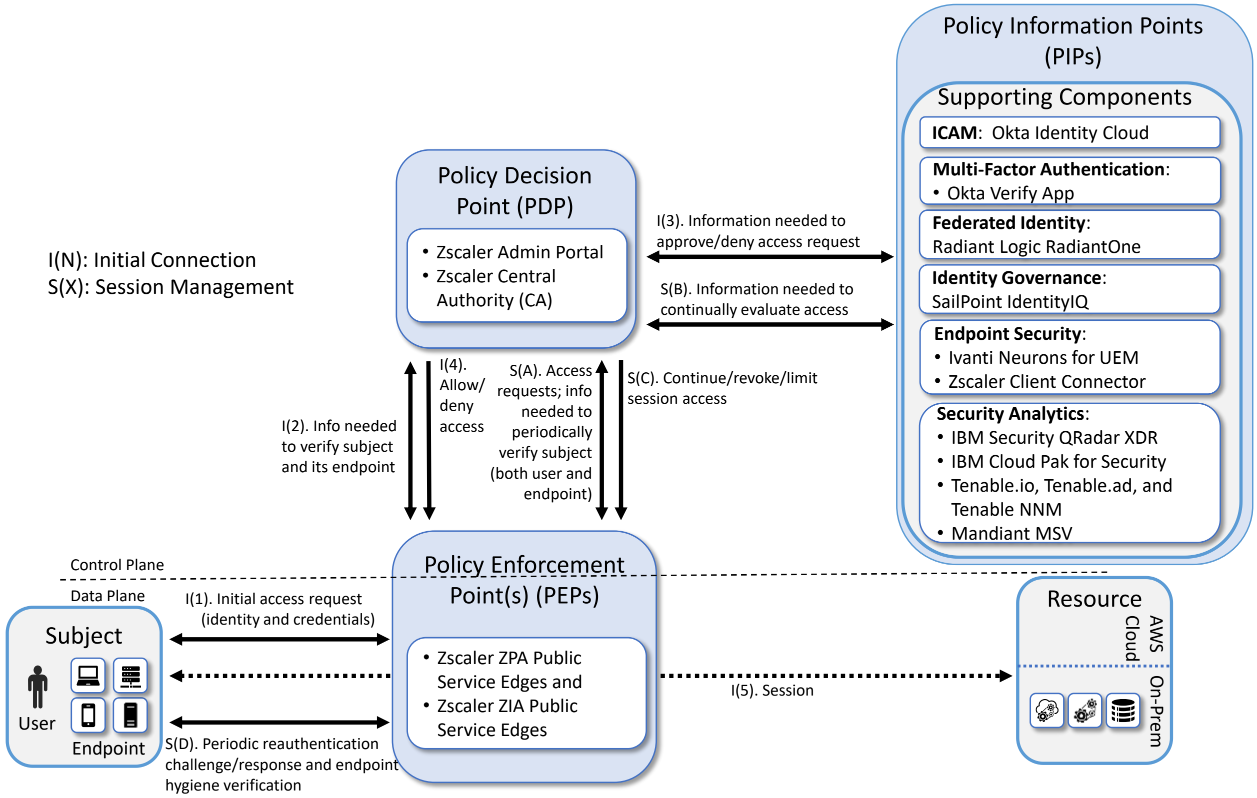

H.2.2 ICAM Information Architecture

How ICAM information is provisioned, distributed, updated, shared, correlated, governed, and used among ZTA components is fundamental to the operation of the ZTA. The ICAM information architecture ensures that when a subject requests access to a resource, the aggregated set of identity information and attributes necessary to identify, authenticate, and authorize the subject is available to be used as a basis on which to make the access decision.

In E1B2, Okta, Radiant Logic, and SailPoint integrate with each other as well as with other components of the ZTA to support the ICAM information architecture. The ways that these components work together to correlate identity information and to support actions such as users joining, changing roles, and leaving the enterprise are the same in E1B2 as they are in E1B1. These interactions are described in Appendix D.2.2.

H.2.3 Physical Architecture

Sections 4.4.1 and 4.4.2 describe and depict the physical architecture of the E1B2 headquarters network and the E1B2 branch office network, respectively. In addition to what is represented in Section 4.4, E1B2 has Zscaler App Connector in the shared services VLAN.

H.2.4 Message Flows for Successful Resource Access Requests

Below are two high-level message flows, both of which support the use case in which a user who has an enterprise ID and who is authorized to access a resource requests and receives access to that resource. The user may be located either on-premises or at a remote location, such as a coffee shop.

In both use cases depicted below, Zscaler platform components are serving as the PDP and PEPs, and Okta Identity Cloud provides a database of users, groups, permissions, and other identity and authorization information that Zscaler consumes. The Zscaler platform and Okta have a SAML federation that provides real-time synchronization of user identity information (to support user authentication) as well as a SCIM federation that provides real-time synchronization of role and group information (to support user authorization). These SAML and SCIM integrations are required because Zscaler relies on Okta to authenticate the identity of users making access requests as well as to help ensure that the user is authorized to access the requested resource.

The Zscaler Central Authority (CA) is the PDP. A Zscaler Client Connector (ZCC) application is assumed to have been installed on the endpoint that the user is using to request access. The ZCC enforces policies that have been configured and applied to the device. When the user requests access to a resource, the ZCC intercepts the request and sends it to either the Zscaler Private Access (ZPA) Service Edge (PEP) or the Zscaler Internet Access (ZIA) Service Edge (PEP). Both the ZPA Service Edge and the ZIA Service Edge perform policy enforcement based on policies that the resource owner is assumed to have already configured. The choice of which PEP to send the request to depends on whether the resource being protected is an internal, private resource (e.g., an enterprise application located on the organization's internal infrastructure--either in an on-premises data center or in the organization's virtual private cloud (VPC) portion of a public cloud infrastructure such as AWS IaaS) or an externally-facing, public resource (e.g., a Microsoft Office 365 application located in a SaaS cloud or a web server on the internet). ZPA is used to broker access to an enterprise's internal resources, while ZIA is used to inspect and secure traffic sent to and from externally facing and public resources.

H.2.4.1 Use Case in which Access to an Internal Resource is Protected Using ZPA

Figure H-2 depicts the message flow for the case in which ZPA acts as the PEP/PDP. In this use case, the resource being accessed is an internal, private resource that does not have a public-facing IP address and may be located either on-premises or in the organization's VPC of AWS IaaS. To support this use case, domains (wildcard or exact) are configured as application segments and context-based access policies must also be configured in the ZPA Administrator Portal (Policy Administrator). ZCC, which is installed on the user's endpoint, validates if a domain accessed is internal based on the Application Segments in the ZPA Administrator Portal. Once ZCC determines the domain is internal, the ZPA Service Edge (PEP) will use the access policies as the basis for deciding whether to broker access to the internal resource. To broker the connection between the ZPA PEP and the internal applications, a ZPA application connector must have been installed near the resource (either on-premises or in the enterprise's VPC in the cloud) and an application segment must have been linked to that connector so that the connector that is near the resources acts as a proxy to the resource(s) on the application segment. ZCC provides a secure, authenticated interface between the endpoint and the ZPA service edge, and the ZPA Application Connector provides a secure, authenticated interface between the resource(s) and the ZPA service edge.

Once the user has logged into the ZCC on his endpoint, all traffic destined for internal resources (e.g., resources within an organization's domain, which may be physically located either on-premises or in a VPC) will be sent to the ZPA PEP in the ZPA cloud that is closest to the user. The ZCC authenticates to the ZPA PEP and then establishes a secure tunnel to it. As a result, user endpoints never connect directly to internal resources. Instead, requests are sent to the ZPA PEP and if they are permitted by ZPA policy (i.e., if the user is authenticated, their access to the resource is authorized, and the requesting endpoint is compliant), then the ZPA PEP brokers access between the user and the application connector for the resource.

Assuming the access request is permitted by policy, another secure tunnel is created between the ZPA PEP and the application connector for the resource. For security reasons, connectors do not accept inbound connections, so the connection that is established between the application connector for the resource and the ZPA PEP is outbound, from the application connector to the ZPA PEP. The ZPA PEP uses the TLS control channel (the reverse TLS tunnel) to signal the application connector to build a data tunnel from the application connector to the ZPA PEP. Then the ZPA PEP stitches together the two TLS tunnels in the cloud, enabling traffic to be exchanged securely between the user endpoint's ZCC and the application connector. If a user connects to multiple resources that are being protected by a single application connector, there will be one TLS/DTLS tunnel created per resource.

When a user requests access to an internal resource, ZCC intercepts DNS lookup queries for these domains and dynamically assigns the domains IP addresses within the 100.64.0.0/16 carrier-grade NAT subnet. Browsers and applications attempting to access the internal resource(s) will route the traffic to the IP addresses set up by ZCC. Due to this, the user accessing the resource never knows the real IP address of the resource, only the address of the temporary IP address assigned by ZCC. The user is not on the network, so connecting to the network via ZPA provides no presumption of access. The only connection that the user's endpoint has is with the ZPA PEP. Logically, the ZPA PEP is positioned between the user endpoint connector and the resource's connector.

All traffic that is sent between a user and an internal resource must be directed through the application connector for that resource. So, for optimal performance, if an enterprise has internal resources in

3502 multiple locations (e.g., both on-premises and in a VPC on AWS), it should deploy application connectors
3503 in each location. Then it should link the respective Application Segment(s) to each location where the
3504 application exists so that the traffic sent from the user to the application can traverse an optimal path
3505 rather than having to be hairpinned through a connector that is not located close to the resource.

3506 **Figure H-2 Access to an Internal Resource is Enforced by Zscaler ZPA and Okta Identity Cloud**

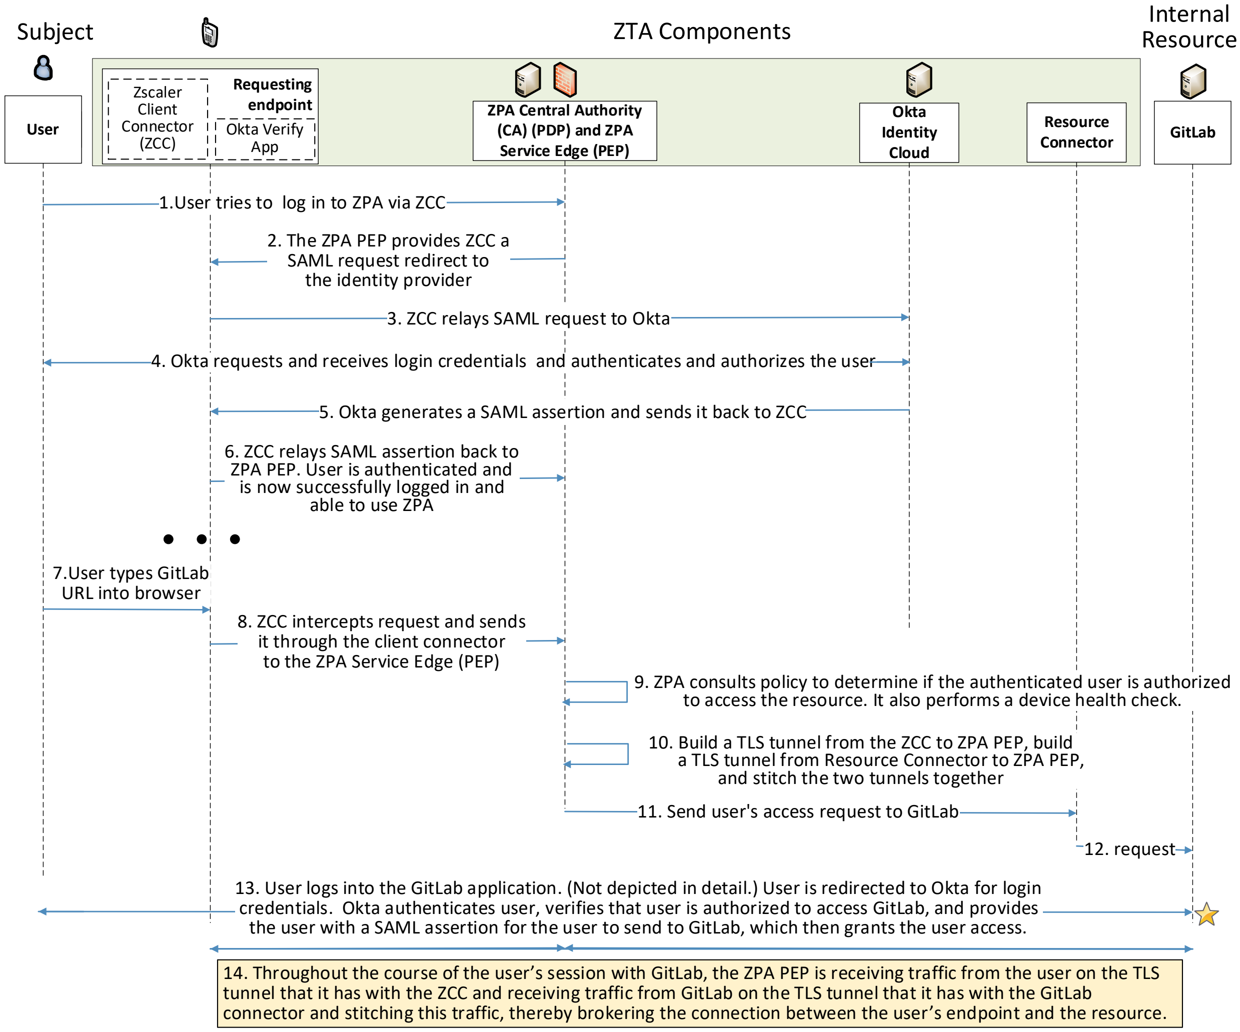

3507 The message flow depicted in Figure H-2 consists of two parts: steps 1-6 depict the high-level message
3508 flow that occurs when a user logs into Zscaler, and steps 7-14 depict the high-level message flow that
3509 occurs when an authenticated user attempts to access an internal resource. The steps are as follows:

3510 1. The user uses the ZCC to try to log into ZPA, and the access request is received at the ZPA PEP.

3511 2. ZPA PEP provides ZCC with a SAML Request redirect to the Identity Provider.

3. The ZCC relays the SAML request to Okta, which is the enterprise's identity provider.
4. Okta requests and receives the user's credentials (and MFA, if configured) and uses these to authenticate the user and ensure that the user is authorized to use ZPA.
5. Okta generates a SAML assertion and sends it back to ZCC.
6. ZCC relays the SAML assertion back to ZPA PEP. The user is authenticated and is now successfully logged in and able to use ZPA.
7. A user requests to access an internal resource by typing the resource URL into their browser.
8. The ZCC intercepts this request, determines if it is an internal resource, and sends it to the ZPA Service Edge (PEP) if it is. (In this use case, the resource is internal.)
9. The ZPA PEP consults access policy to determine if the user is authorized to access the resource. The ZPA PEP performs a device health check to determine if the endpoint requesting access is compliant according to endpoint compliance polices that have been configured in the Zscaler CA (PDP). Information such as device OS version, patch level, anti-virus version, and whether the firewall is running has been collected from the device by the ZCC and provided to ZPA. The ZPA PEP determines if the user is authorized based on username and/or user group.
10. Assuming the user is authorized, the ZPA PEP will broker access to the resource. This is accomplished by building one TLS tunnel from the ZCC to the ZPA PEP and a second TLS tunnel from the resource connector to the ZPA PEP. The ZPA PEP then stitches these two tunnels together in the Zscaler cloud.
11. The ZPA PEP sends the user's original request to access the resource to the resource connector.
12. The resource connector sends the access request to the resource (GitLab).
13. At this point, the user must still complete their login to the GitLab application, so they will select "login via Okta" on the GitLab login screen. The user is then redirected to an Okta screen for login credentials. Okta authenticates the user, verifies that they are authorized to access GitLab, and provides the user with a SAML assertion for the user to send to GitLab. Upon receipt of this SAML assertion, GitLab grants the user access. (These interactions with Okta are not shown in the flow diagram.)
14. Once the user has logged into GitLab, the access session begins. Throughout the course of the user's access session with GitLab, the ZPA PEP brokers the connection between the user's endpoint and the resource. The ZPA PEP receives traffic from the user on the tunnel it has with the ZCC and stitches this traffic to the tunnel it has with the GitLab connector. Similarly, it receives traffic from GitLab on the tunnel it has with the GitLab connector and stitches this traffic to the tunnel it has with the ZCC.

H.2.4.2 Use Case in which Access to an Externally-Facing Resource is Protected Using ZIA

Figure H-3 depicts the message flow for the case in which the ZIA Service Edge acts as the PEP. In this use case, the resource being accessed is externally facing and would typically be located external to the enterprise—e.g., either in a SaaS cloud or on the internet. Once the user has logged into the ZCC on his endpoint, traffic from the user that is destined for external, public resources will be sent to the ZIA Service Edge (PEP) that is closest to the user. A secure TLS tunnel will be established from the ZCC to this ZIA PEP and the traffic destined for this externally facing resource will be forwarded through the tunnel so the ZIA PEP can apply enterprise policies to it.

ZIA PEP is used to determine if access to the resource is permitted at all and, if so, to inspect and secure traffic sent between the requesting endpoint and this external resource. To support this use case, ZIA is typically configured with policies which permit or block access to resources. ZIA can also be configured with traffic inspection policies. The ZIA PEP can inspect all traffic sent between the user and the resource bidirectionally. For example, it can inspect traffic for malware and enforce security, firewall, and web compliance policies (e.g., it may be configured to block PDFs from being sent from the enterprise, or block documents that contain social security numbers). Based on policy, ZIA will either forward the traffic to its destination or drop it. In either case, all traffic is logged and can be reviewed by an administrator.

Unlike ZPA, ZIA does not make use of connectors. The ZIA PEP is used to broker the connection between the user and an externally facing resource. ZIA access policies can be configured based on URLs, URL categories, cloud applications, user location, time, usernames, and/or groups. Providing that the requested resource is permitted based on policy, ZIA enables traffic to be sent directly from the endpoint to the resource (not via a resource connector).

Figure H-3 Access to an Externally-Facing Resource is Enforced by Zscaler ZIA and Okta Identity Cloud

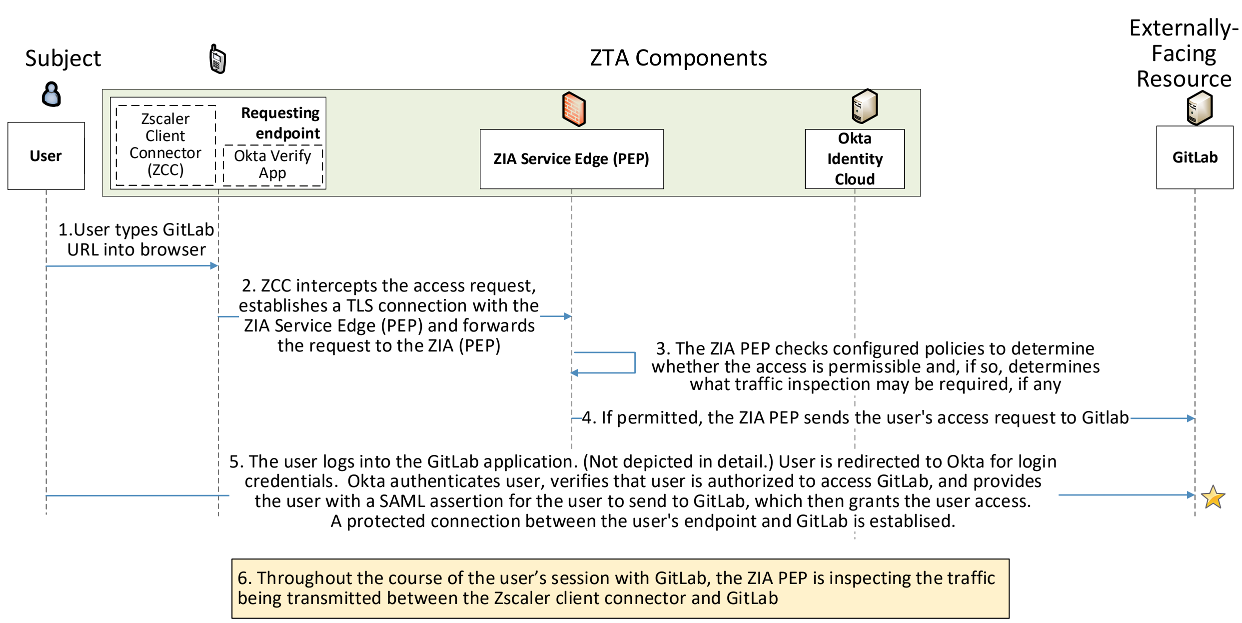

The message flow depicted in Figure H-3 depicts the message flow for the case in which the ZIA Service Edge acts as the PEP. In this use case, the resource being accessed is externally facing and would typically be located external to the enterprise—e.g., either in a SaaS cloud or on the internet. Once the user has logged into the ZCC on his endpoint, traffic from the user that is destined for external, public resources will be sent to the ZIA Service Edge (PEP) that is closest to the user. A secure TLS tunnel will be established from the ZCC to this ZIA PEP and the traffic destined for this externally facing resource will be forwarded through the tunnel so the ZIA PEP can apply enterprise policies to it.

ZIA PEP is used to determine if access to the resource is permitted at all and, if so, to inspect and secure traffic sent between the requesting endpoint and this external resource. To support this use case, ZIA is typically configured with policies which permit or block access to resources. ZIA can also be configured with traffic inspection policies. The ZIA PEP can inspect all traffic sent between the user and the resource bidirectionally. For example, it can inspect traffic for malware and enforce security, firewall, and web compliance policies (e.g., it may be configured to block PDFs from being sent from the enterprise, or block documents that contain social security numbers). Based on policy, ZIA will either forward the traffic to its destination or drop it. In either case, all traffic is logged and can be reviewed by an administrator.

Unlike ZPA, ZIA does not make use of connectors. The ZIA PEP is used to broker the connection between the user and an externally facing resource. ZIA access policies can be configured based on URLs, URL categories, cloud applications, user location, time, usernames, and/or groups. Providing that the requested resource is permitted based on policy, ZIA enables traffic to be sent directly from the endpoint to the resource (not via a resource connector.

Figure H-3 assumes that the user has already logged into ZCC on their endpoint. The message flow consists of the following steps:

1. A user requests access to an externally facing resource (GitLab) by typing the resource URL into their browser.
2. The ZCC intercepts this request, establishes a TLS connection with the ZIA Service Edge (PEP), and forwards the request to the ZIA PEP through this tunnel.
3. ZIA PEP checks configured policies to determine whether the access is permissible and, if permissible, determines what traffic inspection may be required, if any.
4. If permitted, ZIA PEP sends the user's access request to the resource (GitLab)
5. At this point, the user must still complete their login to the GitLab application, so they will select "login via Okta" on the GitLab login screen. The user is then redirected to an Okta screen for login credentials. Okta authenticates the user, verifies that they are authorized to access GitLab, and provides the user with a SAML assertion for the user to send to GitLab. Upon receipt of this SAML assertion, GitLab grants the user access. (These interactions with Okta are not shown in

the flow diagram.) A protected connection between the user’s endpoint and GitLab is established.

6. Throughout the course of the user’s access session with GitLab, the ZIA PEP can inspect the traffic being transmitted between GitLab and the user’s endpoint and either forward or drop the traffic depending upon whether the traffic conforms to the firewall, web, and other security policies that have been defined.

Although ZIA is typically used to protect access to an externally facing resource that is located either in a SaaS cloud or on the internet, NCCoE demonstrated the use of ZIA to protect access to an externally facing resource that is in the NCCoE VPC of AWS IaaS. This resource, GitLab, was placed on a public subnetwork that was segmented from the private subnetwork within that VPC on which internal applications reside. Even though the resource was publicly accessible, access to GitLab was still protected by an identity provider, which in this case is Okta.

Appendix I EIG Enterprise 2 Build 2 (E2B2)

This build will be documented in a future version of this publication.

Appendix J EIG Enterprise 3 Build 2 (E3B2)

J.1 Technologies

EIG E3B2 uses products from F5, Forescout, Mandiant, Microsoft, Palo Alto Networks, PC Matic, and Tenable. Certificates from DigiCert are also used. For more information on these collaborators and the products and technologies that they contributed to this project overall, see Section 3.4.

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

Table J-1 lists all of the technologies used in E3B2 ZTA. It lists the products used to instantiate each ZTA component and the security function that each component provides.

Table J-1 E3B2 Products and Technologies

Component	Product	Function
PE	Microsoft Azure AD (Conditional Access), Microsoft Intune, Forescout eyeControl, and Forescout eyeExtend	Decides whether to grant, deny, or revoke access to a resource based on enterprise policy, information from supporting components, and a trust algorithm.
PA	Microsoft Azure AD (Conditional Access), Microsoft Intune, Forescout eyeControl, and Forescout eyeExtend	Executes the PE's policy decision by sending commands to a PEP that establishes and shuts down the communication path between subject and resource.
PEP	Microsoft Azure AD (Conditional Access), Microsoft Intune, F5 BIG-IP, and Palo Alto Networks Next Generation Firewall (NGFW)	Guards the trust zone that hosts one or more enterprise resources; establishes, monitors, and terminates the connection between subject and resource as directed by the PA; forwards requests to and receives commands from the PA.

Component	Product	Function
Identity Management	Microsoft AD and Azure AD	Creates and manages enterprise user and device accounts, identity records, role information, and access attributes that form the basis of access decisions within an organization to ensure the correct subjects have the appropriate access to the correct resources at the appropriate time.
Access & Credential Management	Microsoft AD and Azure AD	Manages access to resources by performing user and device authentication (e.g., SSO and MFA) and using identity, role, and access attributes to determine which access requests are authorized.
Federated Identity	Microsoft AD and Azure AD	Aggregates and correlates all attributes relating to an identity or object that is being authorized by a ZTA. It enables users of one domain to securely access data or systems of another domain seamlessly, and without the need for completely redundant user administration. Federated identity encompasses the traditional ICAM data, supports identities that may be part of a larger federated ICAM community, and may include non-enterprise employees.
Identity Governance	Microsoft AD and Azure AD Identity Governance	Provides policy-based, centralized, automated processes to manage user identity and access control functions (e.g., ensuring segregation of duties, role management, logging, access reviews, analytics, reporting) to ensure compliance with requirements and regulations.
MFA	Azure AD (Multifactor Authentication)	Authenticates user identity by requiring the user to provide not only something they know (e.g., a password), but also something they have (e.g., a token).

Component	Product	Function
UEM/MDM	Microsoft Intune	Manages and secures enterprise desktop computers, laptops, and/or mobile devices in accordance with enterprise policy to protect applications and data; ensure device compliance; mitigate and remediate vulnerabilities and threats; monitor for suspicious activity to prevent and detect intrusions; prevent, detect, and disable malware and other malicious or unauthorized traffic; repair infected files when possible; provide alerts and recommend remediation actions; and encrypt data. Pushes enterprise applications and updates to devices, enables users to download enterprise applications that they are authorized to access, remotely deletes all applications and data from devices if needed, tracks user activity on devices, and detects and addresses security issues on the device.
EPP	Microsoft Defender for Endpoint, Forescout eyeSight, and PC Matic Pro	Detects and stops threats to endpoints through an integrated suite of endpoint protection technologies including antivirus, data encryption, intrusion prevention, EDR, and DLP. May include mechanisms that are designed to protect applications and data; ensure device compliance with policies regarding hardware, firmware, software, and configuration; monitor endpoints for vulnerabilities, suspicious activity, intrusion, infection, and malware; block unauthorized traffic; disable malware and repair infections; manage and administer software and updates; monitor behavior and critical data; and enable endpoints to be tracked, troubleshooted, and wiped, if necessary.
SIEM	Microsoft Sentinel	Collects and consolidates security information and security event data from many sources; correlates and analyzes the data to help detect anomalies and recognize potential threats and vulnerabilities; and logs the data to adhere to data compliance requirements.
Identity Monitoring	Microsoft Azure AD Identity Protection	Monitors the identity of subjects to detect and send alerts for indicators that user accounts or credentials may be compromised, or to detect sign-in risks for a particular access session.

Component	Product	Function
Vulnerability Scanning and Assessment	Tenable.io, and Tenable.ad	Scans and assesses the enterprise infrastructure and resources for security risks; identifies vulnerabilities and misconfigurations; and provides remediation guidance regarding investigating and prioritizing responses to incidents.
Network Discovery	Forescout eyeSight, and Tenable NNM	Discovers, classifies, and assesses the risk posed by devices and users on the network.
Validation of Control	Forescout eyeSegment	Validates the controls implemented through visibility into network traffic and transaction flows.
Security Validation	Mandiant MSV	Provides visibility and evidence on the status of the security controls' effectiveness in the ZTA. Enable security capabilities of the enterprise to be monitored and verified by continuously validating and measuring the cybersecurity controls; also used to automate the demonstrations that were performed to showcase ZTA capabilities. Mandiant MSV is deployed throughout the project's laboratory environment to enable monitoring and verification of various security aspects of the builds. 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.
Security Analytics and Access Monitoring	Microsoft Defender for Cloud Apps	Monitors cloud resource access sessions for conformance to policy.

Component	Product	Function
Remote Connectivity	Azure AD Application Proxy, Microsoft Defender for Cloud Apps, and Palo Alto NGFW	Palo Alto NGFW is used to provide remote users' connectivity to on-premises resources. Also, two options are available to support remote users' connectivity to resources in IaaS: • The Azure AD Application Proxy can be used to connect directly to private applications, and Microsoft Defender for Cloud Apps can be used to connect to public-facing applications. • Palo Alto NGFW can be used to reach on-premises and then the IPsec tunnel can be used to connect from on-premises to IaaS.
Certificate Management	DigiCert CertCentral TLS Manager	Provides automated capabilities to issue, install, inspect, revoke, renew, and otherwise manage TLS certificates.
Cloud Workload Protection	Microsoft Defender for Cloud	Secures cloud workloads to protect them from known security risks and provides alerts to enable real-time reaction to prevent security events from developing. Monitors traffic to and from cloud and web applications and provides session control to prevents sensitive information from leaving.
Cloud Security Posture Management	Microsoft Defender for Cloud	Continually assesses the security posture of cloud resources.
Cloud IaaS	Azure – GitLab and Wordpress	Provides computing resources, complemented by storage and networking capabilities, hosted by a cloud service provider, offered to customers on demand, and exposed through a GUI and an API.
Cloud SaaS	Digicert CertCentral, Microsoft Azure AD, Microsoft Defender for Endpoint, Microsoft Defender for Cloud, Microsoft Defender for Cloud Apps, Microsoft Identity Governance, Microsoft Intune, Microsoft Office 365, Microsoft Sentinel, and Tenable.io	Cloud-based software delivered for use by the enterprise.

Component	Product	Function
Application	GitLab	Example enterprise resource to be protected. (In this build, GitLab is integrated directly with Azure AD using SAML, and Microsoft Sentinel pulls logs from GitLab.)
Application	Guacamole	Example enterprise resource to be protected. (In this build, BIG-IP serves as an identity-aware proxy that protects access to Guacamole, and BIG-IP is integrated with Azure AD using SAML. Also, Microsoft Sentinel pulls logs from Guacamole.)
Enterprise-Managed Device	Windows client, macOS client, and mobile devices (iOS and Android)	Example endpoints to be protected. (In this build, all enterprise-managed devices are enrolled into Microsoft Intune.)
BYOD	Windows client, macOS client, and mobile devices (iOS and Android)	Example endpoints to be protected.

J.2 Build Architecture

In this section we present the logical architecture of E3B2. We also describe E3B2's physical architecture and present message flow diagrams for some of its processes.

J.2.1 Logical Architecture

Figure J-1 depicts the logical architecture of E3B2. Figure J-1 uses numbered arrows to depict the general flow of messages needed for a subject to request access to a resource and have that access request evaluated based on subject identity (both requesting user and requesting endpoint identity), authorizations, and requesting endpoint health. It also depicts the flow of messages supporting periodic reauthentication of the requesting user and the requesting endpoint and periodic verification of requesting endpoint health, all of which must be performed to continually reevaluate access. The labeled steps in Figure J-1 have the same meanings as they do in Figure 4-1. However, Figure J-1 includes the specific products that instantiate the architecture of E3B2. Figure J-1 also does not depict any of the resource management steps found in Figure 4-1 because the ZTA technologies deployed in E3B2 do not support the ability to perform authentication and reauthentication of the resource or periodic verification of resource health.

E3B2 was designed with Microsoft Azure AD (Conditional Access), Microsoft Intune, Forescout eyesight, and Forescout eyeExtend as the ZTA PEs and PAs, and Microsoft AD and Azure AD providing ICAM support. It includes four PEPs: Microsoft Azure AD (Conditional Access), Microsoft Intune, F5 BIG-IP, and Palo Alto NGFW. A more detailed depiction of the messages that flow among components to support user access requests in the case in which a new endpoint is detected on the network and checked for compliance can be found in Appendix J.2.3.

Figure J-1 Logical Architecture of E3B2

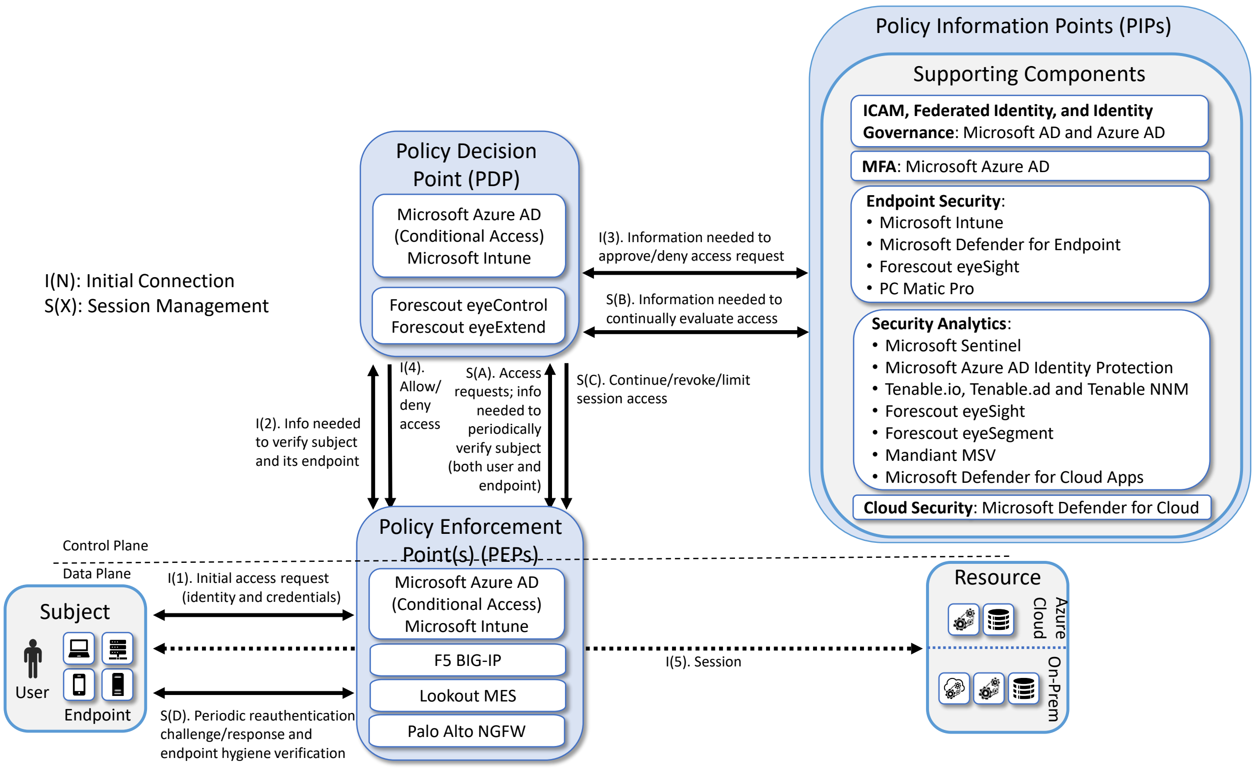

J.2.2 Physical Architecture

Section 4.4.4 describes the physical architecture of the E3B2 network.

J.2.3 Message Flows for a Successful Resource Access Request

The two message flows for E3B1 that are described in Appendix F.2.3 both still apply to E3B2 for cases in which the resource being accessed is located on-premises. Those message flows depict the use cases in which an on-premises resource being accessed is protected by Azure AD alone (see Appendix F.2.3.1), and in which an on-premises resource being accessed is protected by Azure AD in conjunction with the F5 BIG-IP PEP (see Appendix F.2.3.2).

This section depicts three additional high-level message flows. The first two new message flows support the use case in which a user who has an enterprise ID and who is authorized to access a cloud-based resource requests and receives access to that resource. The user may be located on-premises or at a remote location, such as a coffee shop. In the first of these two new use cases, the resource accessed is an internal resource. In the second of these new use cases, the resource is externally facing. The third new message flow presented in this section depicts the use case in which a new endpoint is discovered on the network, found to be non-compliant with enterprise policy, and blocked from accessing all resources.

In both of the cloud-based resource access use cases depicted below, all endpoints are enrolled into Microsoft Intune, which is an MDM that can configure and manage devices, and it can also retrieve and report on device security settings that can be used to determine compliance, such as whether the device is running a firewall or anti-malware. Non-Windows devices have an MDM Agent installed on them to enable them to report compliance information to Microsoft Intune, but Windows devices do not require a separate agent because Windows has built-in agents that are designed to communicate with Intune. Intune-enrolled devices check in with Intune periodically, allowing Intune to authenticate the requesting endpoint, determine how the endpoint is configured, modify certain configurations, and collect much of the information it needs to determine whether or not the endpoint is compliant. Intune reports the device compliance information that it collects to Azure AD, which will not permit a device to access any resources unless it meets configured access policies.

One of the criteria that devices must meet to be considered compliant is that they must have anti-virus software updated and running. Some requesting endpoints have Microsoft Defender Antivirus running on them and other requesting endpoints have PC Matic Pro (also antivirus software) running; no endpoints have both turned on. If a device is running Microsoft Defender Antivirus, the Intune MDM can sense this and report it to Azure AD. If a device is running PC Matic Pro, however, the device is configured to notify Windows Security Center that the endpoint has anti-virus software installed, and the Security Center provides this information to Azure AD.

The authentication message flows depicted below show only the messages that are sent in response to the access request. However, the authentication process also relies on the following additional background communications that occur among components on an ongoing basis:

- Microsoft AD periodically synchronizes with Azure AD to provide it with the most up-to-date identity information.
- Intune-enrolled devices check in with Intune periodically. Checking in allows Intune to determine how the endpoint is configured and modify certain configurations that have been previously specified. It also allows Intune to report the compliance of the device to Azure AD.
- Microsoft Defender for Endpoint has both a cloud component and built-in sensors that detect threats on Windows endpoints. So not only can it tell that a firewall is off or antivirus is off, but it can tell when certain malicious signals seen elsewhere have also been observed on the endpoint. It periodically reports this information to its cloud/management component, which uses it for risk determination. This information can be passed off to Intune to include in its compliance determination of an endpoint.

- Microsoft Defender Antivirus (an endpoint agent) periodically syncs with Microsoft Intune MDM and Microsoft Defender for Endpoint.
- Microsoft Intune periodically sends device health information to Azure AD so that it can be sure that the device is managed and compliant.
- PC Matic periodically syncs with Windows Security Center to inform it that the endpoint has anti-virus installed and active.
- Windows Security Center periodically syncs with Azure AD to provide it with endpoint status information, i.e., that endpoints have anti-virus installed.

J.2.3.1 Use Case in which Access to a Private Cloud Resource is Enforced by Azure AD and Azure AD's Application Proxy

Figure **J-2** depicts the message flow for the use case in which Azure AD's Application Proxy acts as the PEP and Azure AD serves as identity manager. In this use case, the resource being accessed is an internal, private resource that does not have a publicly facing IP address and may be located either on-premises at the owning organization or in a private portion of Azure IaaS or another public cloud that the organization controls. Application Proxy includes both the Application Proxy service, which runs in the cloud as part of Azure AD, and the Application Proxy connector, which is a software agent that runs on a server inside the enterprise's network (either on-premises or in the enterprise's private portion of the cloud) and sits in front of the application being protected to manage communication between the Application Proxy service and the application. The Application Proxy connector uses only outbound HTTPS connections, so there is no need for the enterprise to open inbound ports. The connector can also perform "Kerberos Constrained Delegation (KCD)" in the case of enterprise Kerberos apps, which means that the user authenticating to the cloud can get SSO to Kerberos apps on-premises without re-authentication. For KCD to work, the Application Proxy connector would also need to have a path to an enterprise domain controller.

Figure J-2 Use Case— E3B2 – Access to an Internal Resource is Enforced by Azure AD and Azure AD's Application Proxy

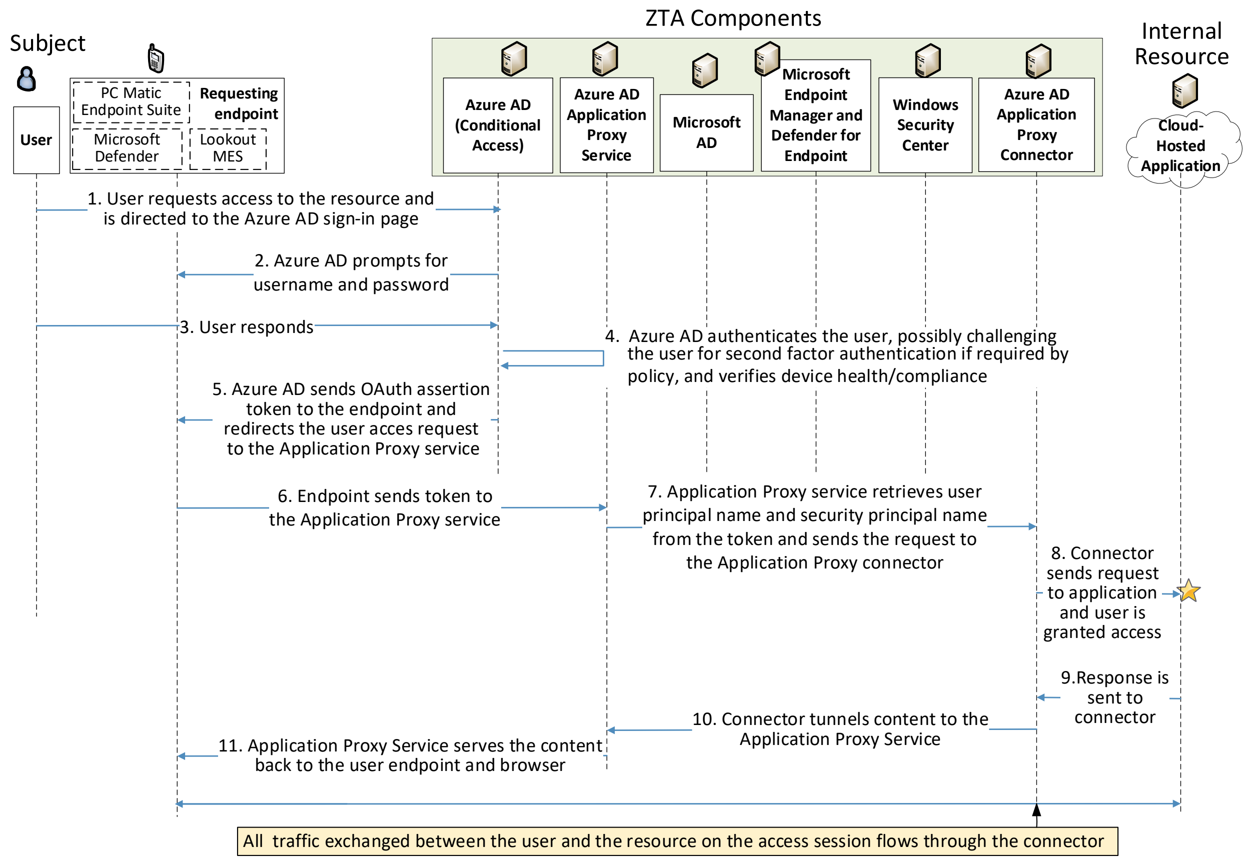

Prior to the flow above, the administrator configures both the Application Proxy connector and the application. This provides the administrator with an internet-facing URL they can give users who are coming off the internet (by default it would be something like app-contoso.msappprox.net, but they can customize the DNS URL with an SSL certificate). The message flow depicted in **Figure J-2** consists of the following steps:

1. A user requests to access an internal resource in the cloud by typing in the external URL provided by the App Proxy service for that resource. This access request is directed to the Microsoft AD sign-in page.
2. Azure AD prompts the user for credentials (e.g., username + password, certificate auth, FIDO2 keys).
3. The user responds with credentials.

4. If required by policy, Azure AD also prompts the user for second-factor authentication. Azure AD Conditional Access can enforce these additional controls (e.g., MFA, device trust, user risk). Azure AD consults the information about the device that it has received in the background from Microsoft Intune and Defender for Endpoint to authenticate the device and verify that it is managed and meets compliance requirements. If the device has PC Matic running on it, Azure AD also consults information about the device that it has received in the background from Windows Security Center to verify that the device is running anti-virus software.
5. Azure AD sends an OAuth token to the user's browser to return to the App Proxy service (SAML can also be configured) and redirects the user access request to Azure AD Application Proxy Service.
6. The endpoint sends the access request and OAuth token to Azure AD Application Proxy Service.
7. The Application Proxy service retrieves the user principal name and security principal name from the token and sends the request to the Application Proxy connector. If KCD was configured (see above), the Proxy Connector reaches out to the domain controller to acquire a Kerberos ticket on behalf of the user identified in the OAuth token for the intended on-premises resource. Alternatively, the Proxy Connector can be configured to inject authentication headers if the application on-premises requests headers. (This KCD-related step is not depicted in the figure because it was not configured in the NCCoE demonstration.)
8. The Application Proxy connector sends the request to the resource (optionally with a Kerberos ticket or headers) and the resource grants the user access.
9. The resource returns content to the Application Proxy connector.
10. The Application Proxy connector tunnels the content to the App Proxy service.
11. The Application Proxy Service serves the content back to the user's end point and browser.

Once the access session is established, all traffic exchanged between the user and the resource flows through the Application Proxy connector.

J.2.3.2 Use Case in which Access to an Externally Facing Cloud Resource is Enforced by Azure AD and Monitored by Microsoft Defender for Cloud Apps

Figure J-3 depicts the message flow for the case in which access to the resource is protected by Azure AD (with the Conditional Access feature), which acts as a PDP; Microsoft AD, which provides identity information, and Microsoft Defender for Cloud Apps, which monitors cloud resource access sessions for conformance to policy. In this use case, the resource being accessed is externally facing, meaning that it has a publicly reachable IP address. Even though the application is externally facing, because the application is in the part of the cloud that is under the organization's control (i.e., configured for SSO with the organization's Identity Provider through SAML or OAuth), it is still protected by the organization's identity provider, Azure AD, which requires the user to authenticate and then verifies that the user is authorized to access the resource and that the resource is compliant before granting access. Once the access session has been established, Microsoft Defender for Cloud Apps monitors all traffic that is exchanged between the user and the resource (see here for a detailed flow explanation). Microsoft Defender for Cloud Apps is therefore able to provide user behavior analytics functionality and prevent harmful or malicious actions within the resource. For example, it can block download of corporate data onto unmanaged devices, or block upload of data onto cloud storage services that contains PII or credit card numbers.

Figure J-3 Use Case— E3B2 – Access to an Externally-Facing Resource is Enforced by Azure AD and Microsoft Defender for Cloud Apps

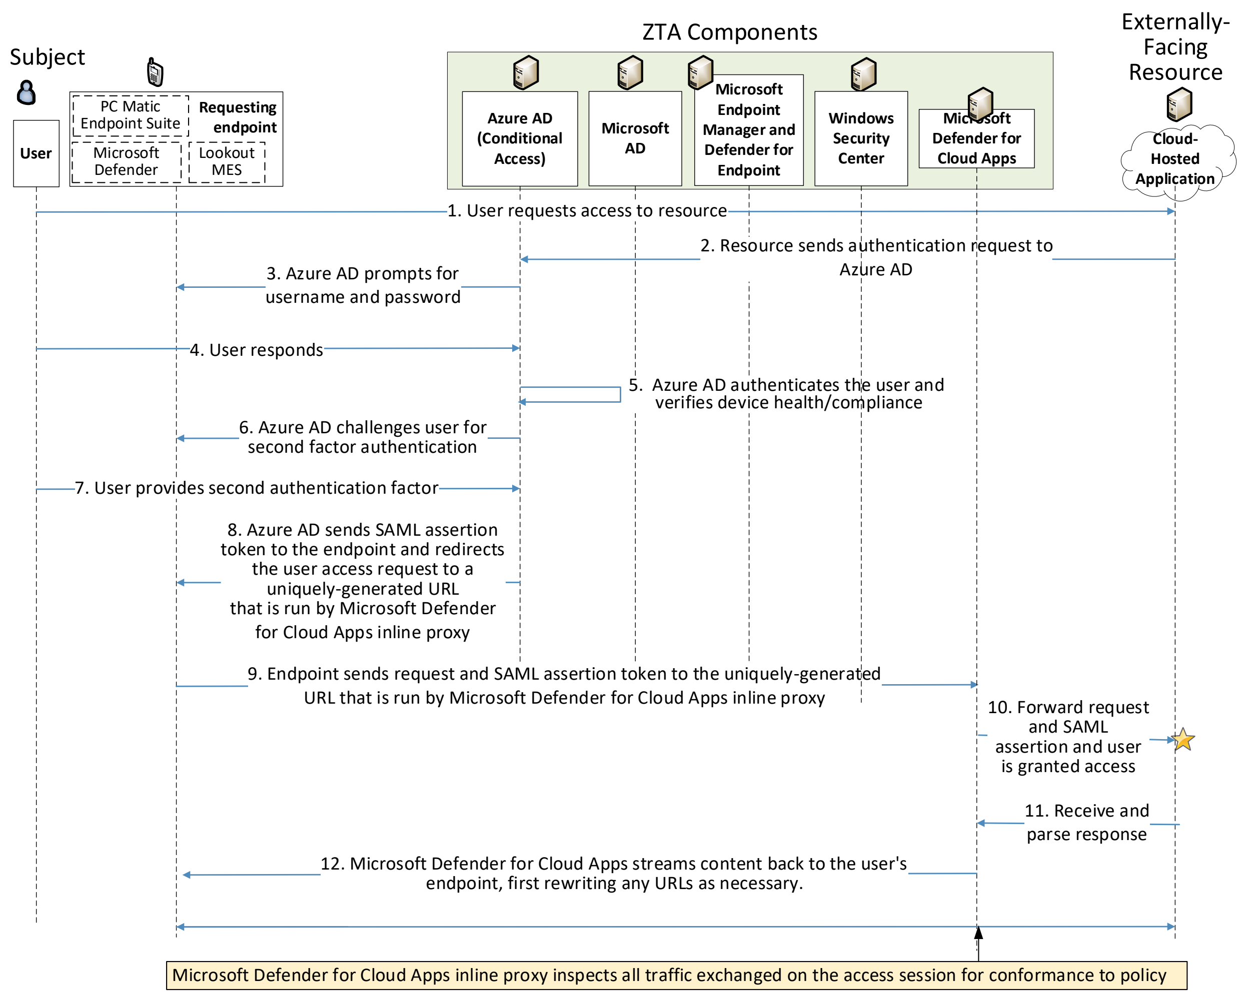

The message flow depicted in Figure J-3 consists of the following steps:

1. A user requests to access an externally facing, cloud-hosted resource, e.g., a SaaS application that has a publicly reachable IP address. For example, app.saas.com.
2. The resource sends the authentication request to Azure AD.

3. Azure AD prompts for credentials.
4. The user responds with credentials.
5. Azure AD authenticates the user. Azure AD consults the information about the device that it has received in the background from Microsoft Intune and Defender for Endpoint to authenticate the device and verify that it is managed and meets compliance requirements. If the device has PC Matic running on it, Azure AD also consults information about the device that it has received in the background from Windows Security Center to verify that the device is running anti-virus software.
6. Azure AD challenges the user to provide the second authentication factor or any other controls.
7. The user responds with the second authentication factor.
8. Azure AD sends a SAML assertion token back to the user's browser/endpoint but does not redirect the user to the resource's original redirect URL configured in the SAML setup (e.g., app.saas.com/saml) and instead redirects the user to a uniquely generated URL that is run by Microsoft Defender for Cloud Apps inline proxy (e.g., app.saml.com.cas.com).
9. The endpoint sends the access request and SAML assertion to Microsoft Defender for Cloud Apps' generated URL.
10. The Microsoft Defender for Cloud Apps inline proxy forwards the request and SAML assertion to the resource's original URL.
11. Microsoft Defender for Cloud Apps receives and parses the response.
12. Before streaming the content back to the user's endpoint, Microsoft Defender for Cloud Apps re-writes any saas.com URLs to be saas.com.cas.com URLs.

The user receives the resulting content from the SaaS app and as they click on any link in the page, they submit their requests back to the Defender for Cloud Apps-generated URL. Defender for Cloud Apps inspects the action and the payload and enforces any DLP or other policies configured. If the action is allowed, Defender for Cloud Apps passes the request on to app.saas.com and, once again, rewrites the URLs of the response before delivery back to the user.

In this manner, for the remainder of the access session, Microsoft Defender for Cloud Apps inline proxy monitors all traffic that is exchanged between the requesting endpoint and the resource endpoint to ensure that is permitted according to enterprise policy. For example, it can inspect the traffic that is sent to and from the cloud for PII or other prohibited content. Microsoft Defender for Cloud Apps inline proxy is integrated with Azure AD Conditional Access, enabling Azure AD to apply its controls to Microsoft Defender for Cloud Apps-governed applications. Furthermore, Defender for Cloud Apps can discover users and endpoints accessing resources, understand and report the risk posture of resources, and identify malicious activity either targeting or sourced from resources, as well as apply DLP policies that mitigate the risk of malicious data exfiltration.

J.2.3.3 Use Case in which a Non-Compliant Endpoint is Discovered on the Network and Blocked from Accessing Resources

Figure J-4 depicts a high-level message flow that supports the use case in which Forescout discovers a non-compliant endpoint on the network and directs the Palo Alto NGFW to block traffic to and from that device.

Figure J-4 Use Case—E3B2 – Forescout Discovers a Non-Compliant Endpoint on the Network and Directs the Palo Alto Firewall to Block it

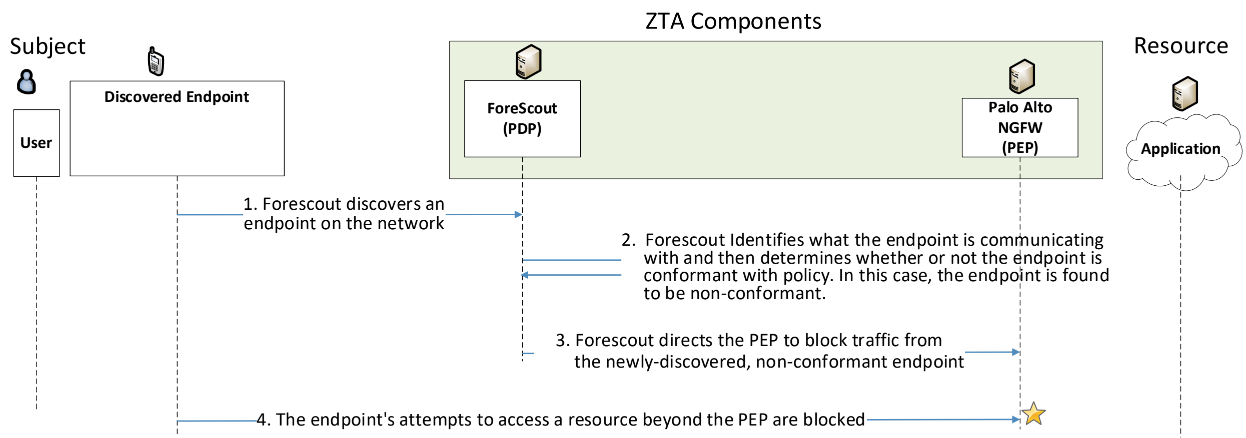

The message flow depicted in Figure J-4 depicts a high-level message flow that supports the use case in which Forescout discovers a non-compliant endpoint on the network and directs the Palo Alto NGFW to block traffic to and from that device.

Figure J-4 consists of the following steps:

1. Forescout discovers a new endpoint on the network.
2. Forescout determines what other resources the endpoint is communicating with and then determines whether or not the endpoint is conformant with policy. (In this use case example, the endpoint is found to be non-conformant.)
3. Forescout direct the Palo Alto NGFW to block traffic to and from this device.
4. When the endpoint attempts to access a resource that is beyond the NGFW, the NGFW blocks the endpoint's traffic.