



1
2
3
4
5
6
7
8
9
10
11
12
13

**NIST Special Publication
NIST SP 800-207A ipd**

**A Zero Trust Architecture Model
for Access Control in Cloud-Native
Applications in Multi-Location
Environments**

Initial Public Draft

Ramaswamy Chandramouli
Zack Butcher

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.800-207A.ipd>

14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36

**NIST Special Publication
NIST SP 800-207A ipd**

A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Location Environments


Initial Public Draft

Ramaswamy Chandramouli
*Computer Security Division
Information Technology Laboratory*

Zack Butcher
Tetrate, Inc.

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.800-207A.ipd>

April 2023



U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

37 Certain commercial equipment, instruments, software, or materials, commercial or non-commercial, are identified in
38 this paper in order to specify the experimental procedure adequately. Such identification does not imply
39 recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or
40 equipment identified are necessarily the best available for the purpose.

41 There may be references in this publication to other publications currently under development by NIST in
42 accordance with its assigned statutory responsibilities. The information in this publication, including concepts and
43 methodologies, may be used by federal agencies even before the completion of such companion publications. Thus,
44 until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain
45 operative. For planning and transition purposes, federal agencies may wish to closely follow the development of
46 these new publications by NIST.

47 Organizations are encouraged to review all draft publications during public comment periods and provide feedback
48 to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at
49 <https://csrc.nist.gov/publications>.

50 **Authority**

51 This publication has been developed by NIST in accordance with its statutory responsibilities under the Federal
52 Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283.
53 NIST is responsible for developing information security standards and guidelines, including minimum requirements
54 for federal information systems, but such standards and guidelines shall not apply to national security systems
55 without the express approval of appropriate federal officials exercising policy authority over such systems. This
56 guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

57
58 Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding
59 on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be
60 interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or
61 any other federal official. This publication may be used by nongovernmental organizations on a voluntary basis and
62 is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

63 **NIST Technical Series Policies**

64 [Copyright, Use, and Licensing Statements](#)
65 [NIST Technical Series Publication Identifier Syntax](#)

66 **Publication History**

67 Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be updated in final publication]

68 **How to Cite this NIST Technical Series Publication:**

69 Chandramouli R, Butcher Z (2023) A Zero-Trust Architecture Model for Access Control in Cloud Native
70 Applications in Multi-Location Environments. (National Institute of Standards and Technology, Gaithersburg, MD),
71 NIST Special Publication (SP) NIST SP 800-207A ipd. <https://doi.org/10.6028/NIST.800-207A.ipd>

72 **Author ORCID iDs**

73 Ramaswamy Chandramouli: 0000-0002-7387-5858

74 **Public Comment Period**

75 April 18, 2023 – June 7, 2023

76 **Submit Comments**

77 sp800-207A-comments@nist.gov

78

79 National Institute of Standards and Technology

80 Attn: Computer Security Division, Information Technology Laboratory

81 100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

82 **All comments are subject to release under the Freedom of Information Act (FOIA).**

83 **Abstract**

84 One of the basic tenets of zero trust is to remove the implicit trust in users, services, and devices
85 based only on their network location, affiliation, and ownership. NIST Special Publication 800-
86 207 has laid out a comprehensive set of zero trust principles and referenced zero trust
87 architectures (ZTA) for turning those concepts into reality. A key paradigm shift in ZTAs is the
88 change in focus from security controls based on segmentation and isolation using network
89 parameters (e.g., IP addresses, subnets, perimeter) to identities. From an application security
90 point of view, this requires authentication and authorization policies based on application and
91 service identities in addition to the underlying network parameters and user identities. This in
92 turn requires a platform that consists of API gateways, sidecar proxies, and application identity
93 infrastructures (e.g., SPIFFE) that can enforce those policies irrespective of the location of the
94 services/applications, whether on-premises or on multiple clouds. The objective of this
95 publication is to provide guidance for realizing an architecture that can enforce granular
96 application-level policies while meeting the runtime requirements of ZTA for multi-cloud and
97 hybrid environments.

98 **Keywords**

99 egress gateway; identity-tier policies; ingress gateway; microservices; multi-cloud; network-tier
100 policies; service mesh; sidecar proxy; SPIFFE; transit gateway; zero trust; zero trust architecture.

101 **Reports on Computer Systems Technology**

102 The Information Technology Laboratory (ITL) at the National Institute of Standards and
103 Technology (NIST) promotes the U.S. economy and public welfare by providing technical
104 leadership for the Nation’s measurement and standards infrastructure. ITL develops tests, test
105 methods, reference data, proof of concept implementations, and technical analyses to advance
106 the development and productive use of information technology. ITL’s responsibilities include the
107 development of management, administrative, technical, and physical standards and guidelines for
108 the cost-effective security and privacy of other than national security-related information in
109 federal information systems. The Special Publication 800-series reports on ITL’s research,
110 guidelines, and outreach efforts in information system security, and its collaborative activities
111 with industry, government, and academic organizations.

112

113 **Call for Patent Claims**

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

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

- 122 a) assurance in the form of a general disclaimer to the effect that such party does not hold
123 and does not currently intend holding any essential patent claim(s); or
- 124 b) assurance that a license to such essential patent claim(s) will be made available to
125 applicants desiring to utilize the license for the purpose of complying with the guidance
126 or requirements in this ITL draft publication either:
 - 127 i. under reasonable terms and conditions that are demonstrably free of any unfair
128 discrimination; or
 - 129 ii. without compensation and under reasonable terms and conditions that are
130 demonstrably free of any unfair discrimination.

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

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

138 Such statements should be addressed to: sp800-207A-comments@nist.gov

139

| | | |
|-----|--|-----------|
| 140 | Table of Contents | |
| 141 | Executive Summary | 1 |
| 142 | 1. Introduction | 2 |
| 143 | 1.1. Background – Zero Trust Principles and Zero Trust Architecture | 2 |
| 144 | 1.2. Relationship to Other NIST Guidance Documents | 3 |
| 145 | 1.3. Scope | 3 |
| 146 | 1.4. Target Audience | 4 |
| 147 | 1.5. Organization of This Document | 4 |
| 148 | 2. The Enterprise Cloud-Native Platform and its Components | 5 |
| 149 | 2.1. Enterprise Infrastructure Layer | 6 |
| 150 | 3. Designing a Policy Framework for ZTA for Cloud-Native Application Environments . | 7 |
| 151 | 3.1. Functional Components of Identity-Based Segmentation Policies for ZTA..... | 8 |
| 152 | 3.2. Shortcomings of Identity-Based Segmentation Policies for Enterprise ZTA | 9 |
| 153 | 3.3. Multi-Tier Policies for Enterprise ZTA | 9 |
| 154 | 4. Implementing Multi-Tier Policies for ZTA for Cloud-Native Application Environments | |
| 155 | 11 | |
| 156 | 4.1. Reference Application Infrastructure Scenario | 11 |
| 157 | 4.2. Role of the Service Mesh in Policy Deployment, Enforcement, and Updates | 12 |
| 158 | 4.3. Policy Deployment for Reference Application Infrastructure..... | 13 |
| 159 | 4.4. Another Application Infrastructure Scenario | 14 |
| 160 | 4.5. Functional Roles of Application Infrastructure Elements in Enforcing Policies | 15 |
| 161 | 4.6. Comparison of Identity-Tier and Network-Tier Policies | 16 |
| 162 | 4.6.1. Approaches for Deployment and the Limitations of Network-Tier Policies | 16 |
| 163 | 4.6.2. Prerequisites for the Deployment of Identity-Tier Policies | 17 |
| 164 | 4.6.3. Advantages of Identity-Tier Policies..... | 18 |
| 165 | 5. Summary and Conclusions | 19 |
| 166 | References | 20 |
| 167 | List of Figures | |
| 168 | Fig. 1. Enterprise infrastructure layer for uniform policy deployment..... | 7 |
| 169 | Fig. 2. Flexibility provided by multi-tier policies | 10 |
| 170 | Fig. 3. Multi-tier Policies for a Hybrid Application Environment | 12 |
| 171 | Fig. 4. An Istio Authorization Policy that allows Service 1 to Service 2 on port 443 but only | |
| 172 | allows it to execute the GET HTTP verb on the “/public” path..... | 14 |
| 173 | Fig. 5. Policy Deployment for a Three-tier Application..... | 15 |
| 174 | | |

175 **Acknowledgments**

176 The author would like to express his thanks to Isabel Van Wyk of NIST for her detailed editorial
177 review of the public comment version as well as the final publication.

178 **Executive Summary**

179 The principles of zero trust, as described in NIST Special Publication (SP) 800-207, have
180 become the guiding markers for developing secure zero trust architecture. A well-established
181 class of applications are cloud-native applications. The generally accepted characterization of a
182 cloud native application includes the following:

- 183 • The application is made up of a set of loosely coupled components called microservices.
184 Each of the microservices can be hosted on different physical or virtual machines (VMs)
185 and even be geographically distributed (e.g., within several facilities that belong to the
186 enterprise, such as the headquarters, branch offices, and in various cloud service provider
187 environments).
- 188 • Any transaction involving the application may also involve one or more inter-service
189 (microservice) calls across the network.
- 190 • A widespread feature (though not necessarily a requirement for cloud-native application)
191 is the presence of a software platform called the service mesh that provides an integrated
192 set of all application services (e.g., services discovery, networking connections,
193 communication resilience, and security services like authentication and authorization).

194 The realization of a zero trust architecture for the above class of cloud-native applications
195 requires a robust policy framework. In order to follow zero trust principles, the constituent
196 polices in the framework should consider the following scenario:

- 197 • There should not be implicit trust in users, services, or devices based exclusively on their
198 network location, affiliation, or ownership. Hence, policy definitions and associated
199 security controls based on the segmentation or isolation of networks using network
200 parameters (e.g., IP addresses, subnets, perimeter) are insufficient. These policies fall
201 under the classification of network-tier policies.
- 202 • To ensure the presence of zero trust principles throughout the entire application, network-
203 tier policies must be augmented with policies that establish trust in the identity of the
204 various participating entities (e.g., users and services) irrespective of the location of the
205 services or applications, whether on-premises or on multiple clouds.

206 This document provides guidance for realizing a zero trust architecture that can enforce granular
207 application-level policies for cloud-native applications. The guidance is anchored in the
208 following:

- 209 • A combination of network-tier and identity-tier policies
- 210 • The components of cloud-native applications that enable the definition and deployment
211 of those policies, such as edge, ingress, sidecar, and egress gateways; the creation,
212 issuance, and maintenance of service identities; the issuance of authentication and
213 authorization tokens that carry user identities in the enterprise application infrastructure
214 that encompasses multi-cloud and hybrid environments

215

216 **1. Introduction**

217 Zero trust (ZT) tenets or principles have been accepted as the guide markers for architecting all
218 applications. There are several reasons why adherence to these tenets is critical for obtaining
219 necessary security assurances, especially for cloud-native applications. The enterprise
220 application environments for this class of applications is highly geographically distributed and
221 span multiple cloud and on-premises environments (e.g., headquarters, enterprise-operated data
222 centers, branch offices, etc.). Further, the user base consists of both remote and on-premises
223 employees. These two features call for establishing trust in all of the data sources and computing
224 services of the enterprise – irrespective of their location – through secure communication and the
225 validation of access policies.

226 Apart from geographic distribution, another common feature of cloud-native applications is the
227 presence of many microservices that are loosely coupled and collectively support business
228 processes through extensive inter-service calls. This is augmented with an integrated
229 infrastructure for providing all application services called the service mesh. These features
230 emphasize the concept of identity for the various components of the application in the form of
231 microservices as well as the users who access them through direct calls or clients (other
232 services). This in turn highlights the critical need for authenticating these identities and for
233 providing legitimate access on a per-session basis through a dynamic policy that takes the current
234 status of the user, service, and requested asset into account.

235 The above requirements can only be met through a comprehensive policy framework. This
236 document provides guidance for developing a policy framework that will form the foundation for
237 realizing a zero trust architecture (ZTA) while incorporating zero trust principles into its design
238 for cloud-native applications. The policy framework should also consist of a comprehensive set
239 of policies that span all critical entities and resources in the application stack, including the
240 network, network devices, users, and services.

241 **1.1. Background – Zero Trust Principles and Zero Trust Architecture**

242 A summary of the zero trust principles and the definition of a zero trust architecture, as described
243 in NIST SP 800-207 [1], are:

- 244 • Zero trust is the term for an evolving set of cybersecurity paradigms that move defenses
245 from static, network-based perimeters to focus on users, assets, and resources. It is a set
246 of security primitives rather than a particular set of technologies. Zero trust assumes that
247 there is no implicit trust granted to assets or user accounts based solely on their physical
248 or network location (i.e., local area networks versus the internet) or on asset ownership
249 (e.g., enterprise or personally owned). Zero trust focuses on protecting resources (e.g.,
250 assets, services, workflows, network accounts) rather than network segments, as the
251 network location is no longer seen as the prime component to the security posture of the
252 resource.
- 253 • A zero trust architecture uses zero trust principles to plan industrial and enterprise
254 infrastructures and workflows.

255 NIST’s guidance on zero trust also contains an abstract definition of zero trust architecture and
256 gives general deployment models and use cases with which zero trust could improve an
257 enterprise’s overall information technology security posture.

258 **1.2. Relationship to Other NIST Guidance Documents**

259 Since the current document provides guidance for the realization of ZTA for cloud-native
260 applications hosted in multiple locations (on-premises and multiple clouds) and the enforcement
261 of ZT principles requires policies that are associated with various security services, it will be
262 useful to refer to the following documents. These documents provide background information for
263 the architecture of a microservices-based application with service mesh as well as guidance for
264 configuring specific security services. The current document expands the reference environment
265 to one where the IT application infrastructure of an enterprise spans multiple premises and
266 multiple cloud provider locations as well as addresses the range of policies that are required for
267 comprehensive security assurance.

- 268 • NIST SP 800-204A, *Building Secure Microservices-based Applications Using Service-*
269 *Mesh Architecture* [2], provides deployment guidance for various security services (e.g.,
270 establishment of secure sessions, security monitoring, etc.) for a microservices-based
271 application using a dedicated infrastructure (i.e., a service mesh) based on service proxies
272 that operate independently of the application code.
- 273 • NIST SP 800-204B [3], *Attribute-based Access Control for Microservices-based*
274 *Applications Using a Service Mesh*, provides deployment guidance for building an
275 authentication and authorization framework within the service mesh that meets the
276 security requirements. This may include establishing (1) zero trust by enabling mutual
277 authentication in communication between any pair of services and (2) a robust access
278 control mechanism based on an access control model (e.g., the attribute-based access
279 control [ABAC] model) that can be used to express a wide set of policies and is scalable
280 in terms of user base, objects (resources), and deployment environment.

281 **1.3. Scope**

282 The scope of this document includes:

- 283 • Identifying the requirements for realizing a ZTA for granular access control in
284 microservices-based application platforms that include a service mesh infrastructure
- 285 • Identifying the infrastructural elements that should be part of the platform in order to
286 configure and implement ZT principles
- 287 • Guidance for deploying a ZTA in the above platform and outlining the security
288 assurances that the deployment can provide

289 **1.4. Target Audience**

290 This guidance is intended for security architects and infrastructure designers in organizations
291 with a hybrid IT environment (consisting of both on-premises and multiple cloud-based
292 applications) with a combination of legacy and microservices-based (i.e., cloud-native)
293 applications with a built-in application services infrastructure, such as a service mesh.

294 **1.5. Organization of This Document**

295 The organization of this document is as follows:

- 296 • Section 2 describes a modern enterprise cloud-native application platform that includes a
297 dedicated infrastructure for providing all application services as well as a management
298 plane when the application spans both on-premises and multiple cloud service provider
299 locations.
- 300 • Section 3 introduces the basic concepts of a policy framework for ZTA for the platform
301 described in the previous section in terms of drivers and design requirements. It also
302 provides an analysis of identity-based policies and introduces the concept of multi-tier
303 policies.
- 304 • Section 4 describes the implementation approach for deploying multi-tier policies for two
305 enterprise application infrastructure scenarios by outlining the roles of the service mesh,
306 the functional components involved, and the advantages of identity-tier policies, which
307 provide service-level segmentation and play a critical role in the security assurance of an
308 application ecosystem to conform to zero trust principles or tenets.
- 309 • Section 5 provides a summary and conclusion.

310

311

312 **2. The Enterprise Cloud-Native Platform and its Components**

313 An enterprise cloud-native platform is increasingly made up of microservices that are
314 implemented as containers and hosted on a container orchestration platform. In addition, it has a
315 dedicated infrastructure layer called a service mesh, which provides a comprehensive set of
316 application services (e.g., network connectivity, network resilience, observability, and security).
317 The application services provided by a service mesh are enabled by the following:

- 318 • A built-in infrastructure for (a) providing service identities, (b) service discovery, and (c)
319 external policy-based authorization engines based on Next Generation Access Control
320 (NGAC), Attribute-based Access Control (ABAC), and Open Policy Agent (OPA)
- 321 • Code for performing network-related functions (e.g., traffic routing) and for ensuring
322 network resiliency through functions such as retries, timeouts, blue-green deployments,
323 and circuit breaking
- 324 • Code for ensuring application integrity and confidentiality through service-to-service and
325 user-to-resource authentications and authorizations

326 More details on the container orchestration platform with an integrated service mesh can be
327 found in [2], and an access control implementation in that platform is described extensively in
328 [3].

329 In the modern enterprise, the platform described above is present in both on-premises data
330 centers and multiple cloud service locations. Assuming that a service mesh instance is deployed
331 for managing a single cluster that consists of the above platforms, there will be multiple clusters
332 spread over multiple on-premises sites and multiple availability zones in different clouds.
333 Consequently, there will be multiple service mesh instances.

334 Each service mesh instance has two main logical components: 1) a control plane that implements
335 the APIs needed to define various configurations and policies that govern access between various
336 microservices in that cluster and 2) a data plane that enforces those policies at runtime. However,
337 a uniform set of policies is also needed to govern access between any pair of microservices or
338 services in the enterprise irrespective of their location or the service mesh instance of which they
339 are a part. This requires a global control plane that can define a uniform set of policies applicable
340 to the entire set of services that operate in the enterprise and disseminate them to the control
341 planes of the individual service mesh instances.

342 It is technically possible to have a single service mesh control plane instance (i.e., single service
343 mesh instance) that manages multiple clusters spanning multiple environments (i.e., on-premises
344 and on clouds). However, this architecture may make the multiple clusters a single failure
345 domain and potentially defeat the very purpose of designing a multi-cluster configuration (i.e.,
346 availability). Thus, running a service mesh control plane instance for each cluster isolates the
347 failure domain and improves availability and scalability. Further, providing the required
348 underlying network connectivity to facilitate every workload (since each workload or application
349 instance has an associated sidecar proxy that forms the data plane) to communicate with a single
350 control plane instance is untenable in most enterprise environments and impossible in many
351 government ones (e.g., air-gapped systems).

352 **2.1. Enterprise Infrastructure Layer**

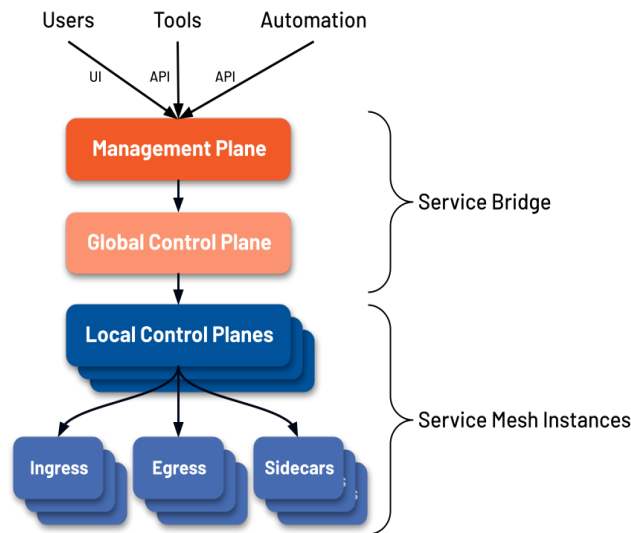
353 The global control plane forms an integral part of the enterprise infrastructure layer. The
354 management plane that contains the various interfaces is hosted within the global control plane.
355 The roles of the global control plane and the management plane are as follows:

- 356 • The global control plane can be leveraged to perform the functions of individual control
357 planes at the enterprise level rather than at the cluster level (e.g., issuing identities to all
358 services in the enterprise by leveraging the enterprise PKI system).
- 359 • The management plane provides the human-computer interfaces (e.g., user interfaces,
360 such as command line interfaces and APIs) that enable enterprise-level systems to work
361 by encoding organizational processes related to the usage of various tools (e.g., policy
362 definition and evaluation tools, telemetry tools, etc.) at the lower layer.

363 In short, the management plane enables the definition and deployment of consistent and uniform
364 policies for all services throughout the enterprise. In addition to the global control plane and
365 management plane, the enterprise infrastructure for a ZTA consists of local control planes
366 (associated with service mesh instances) and a set of various types of proxies that form part of
367 their respective data planes. The proxies act as the policy enforcement points (PEPs) and have
368 three types:

- 369 1. Ingress proxies enforce policies for entering user or service requests from client
370 applications that originate outside of the cluster into any service within the cluster.
- 371 2. Side-car proxies enforce policies between intra-cluster services.
- 372 3. Egress proxies enforce policies for requests that emanate from any service within the
373 cluster to an external application that is outside of the cluster.

374 **Figure 1** shows a schematic diagram of the entire infrastructure layer for uniform (enterprise-
375 wide) policy deployment for realizing a ZTA:



376

377

Fig. 1. Enterprise infrastructure layer for uniform policy deployment

378

3. Designing a Policy Framework for ZTA for Cloud-Native Application Environments

379

380 Based on the set of zero trust principles and some strawman ZTAs provided in [1], the following
381 driver assumptions were formulated for realizing a ZTA for an enterprise cloud-native
382 application environment (i.e., a set of microservices in various clusters with each cluster
383 managed by a service mesh and augmented with an enterprise-level infrastructure that consists of
384 a global control plane and management plane). These driver assumptions are:

385

- 386 • Trust can no longer be based on a network perimeter as perimeters can always be breached.
- 387 • Policies have to be defined based on the assumption that the attacker is already inside of
388 the corporate network.
- 389 • All access decisions have to rely on least-privilege, per-request, and context-based
390 principles and on identities associated with users, services, and devices. This results in a
391 form of runtime isolation for applications, which this document refers to as “identity-
392 based segmentation.”

393

The above driver assumptions provide the design requirements for a ZTA as follows:

394

- 395 • No single component or function is sufficient to implement ZTA. Rather, they must collectively enforce zero trust principles across all applications in the infrastructure.
- 396 • ZTA component functions should be clearly articulated, including their interrelationships
397 and workflows.

- 398 • The enforcement infrastructure that implements the security controls (mainly consisting
399 of PEPs) should satisfy the properties of a security kernel – always invoked (non-
400 bypassable), verifiable, and independent of the application code.
- 401 • The core tenant or primary function of ZTA at runtime is implementing an identity-based
402 segmentation of applications that leverages the enforcement infrastructure.

403 3.1. Functional Components of Identity-Based Segmentation Policies for ZTA

404 The following policy checks should be implemented at runtime through the deployment of
405 identity-tier policies in order to realize identity-based segmentation:

- 406 • ID-SEG-REC-1: Encrypted connection between service endpoints – Service endpoints
407 can be located in different subnets, different availability zones or regions in a cloud
408 provider environment, in different clouds, or on-premises. Wherever they are located,
409 communication between any two should be encrypted to ensure eavesdropping protection
410 and message authenticity.
- 411 • ID-SEG-REC-2: Service authentication – Each service should present a short-lived
412 cryptographically verifiable identity to other services that is authenticated per connection
413 and reauthenticated regularly.

414 Note on the above recommendation: In an ideal situation, services would be authenticated
415 for each service request. Since this is highly disruptive from the point of view of
416 application transaction response, this authentication is accomplished at the connection
417 level via mutual TLS (mTLS) when a service makes an initial connection establishment
418 as part of its inter-service call. This authentication is not performed again in subsequent
419 calls. However, the security of this operation is ensured by not allowing the connections
420 to be very long (usually as long as the TTL of the service’s identity certificate or as short
421 as 15-30 minutes, depending on the configuration).

- 422 • ID-SEG-REC-3: Service to service authorization – Services should leverage runtime
423 service identity (ID-SEG-REC-2) to enforce granular policies and have the capability to
424 call external authorization services if the mesh level proxies are insufficient to enforce
425 dynamic authorization policies.
- 426 • ID-SEG-REC-4: End-user authentication – Since all application requests are triggered by
427 user actions, a robust identity management system is required to assign and maintain user
428 identities and enforce robust protocols with phishing-resistant multi-factor (MFA)
429 authentication. This system should be used to issue a cryptographically verifiable runtime
430 token that represents the user principal to the rest of the infrastructure (e.g., a JSON Web
431 Token [JWT]), and services should authenticate the credential at each hop.

432 Note on the above recommendation: Authenticating the user in session at every hop is
433 impractical at scale. Therefore, NIST recommends using short-lived end-user credentials
434 (e.g., OAuth 2.0 tokens) for external users and exchanging them for a locally
435 authenticatable token, like a JWT, that is authenticated at each hop.

- 436 • ID-SEG-REC-5: End-user to resource authorization – As part of each service access
437 request, the system must ensure that the authenticated end user principal (ID-SEG-REC-
438 4) is authorized to act on the resources designated in the request. This authorization may
439 be performed by the application itself or checked locally (e.g., by checking against a set
440 of claims in a JWT) or externally against an authorization system’s policy decision point.
441 Enforcing end user authorization via the service mesh’s sidecar PEP is particularly
442 effective [3].

443 Context for the application of these policy recommendations and the improved security
444 assurance that emanates from their deployment and enforcement are explained in [2] and [3].

445 **3.2. Shortcomings of Identity-Based Segmentation Policies for Enterprise ZTA**

446 While identity-based segmentation is powerful, purely identity-based policies cannot currently
447 be adopted due to the following scenarios:

- 448 • Identity-based segmentation policies can include access scenarios that cover all origins,
449 such as users, services, and all target resources that consist of services and data.
450 However, enterprise scenarios that involve both on-premises and cloud-based
451 applications require identification of the location of those resources using network
452 parameters. Purely identity-based enforcement should be augmented by other factors
453 (e.g., network location) to evaluate risk when performing context-based authorization.
- 454 • A subset of identity-based segmentation policies (i.e., service identity-based) can be
455 difficult to administer since service identity assignments are often based on specific
456 domains, which makes consistent policy deployment difficult across on-premises
457 systems, cloud-based systems, and different compute runtimes. However, this is
458 mitigated by adopting consistent service names across the infrastructure using the
459 concept of a universal identity domain, as recommended in SM-DR11 of [2].
- 460 • Having network-level policies alone requires high maintenance due to the continuous
461 changes to their location parameters as containers and virtualized workloads are
462 frequently migrated for availability and performance reasons (e.g., migration to different
463 VMs or to a different pod in containerized applications).

464 Network-oriented policies cannot be completely eliminated given current compliance
465 requirements and regulations. However, relaxing requirements at the network level in exchange
466 for introducing more descriptive policy at the identity level could lead to an improved overall
467 security posture compared to network-oriented security alone.

468 **3.3. Multi-Tier Policies for Enterprise ZTA**

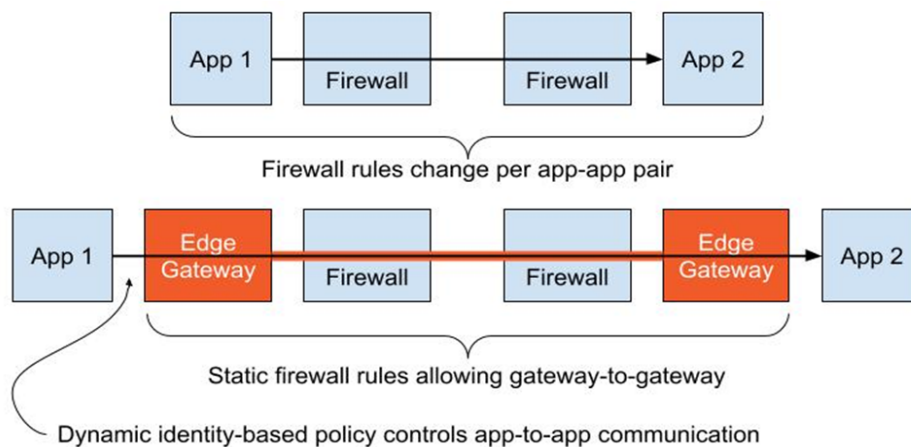
469 A successful enterprise ZTA requires multi-tier policies:

- 470 • Network-tier policies – Allowed communication between enterprise network elements
471 (e.g., firewall rules, which are relatively static)

- Identity-tier policies – Access scope for services and resources based on service and user identities (e.g., dynamic application-to-application communication rules based on identities through a dedicated infrastructure layer, such as user identity provided by an enterprise IAM provider and service identity provided by a standard-based Secure Production Identity Framework for Everyone [SPIFFE] server [4])

Multi-tier policies can be implemented realistically and are non-disruptive to current compliance practices. Other tiers of policy also exist. For example, in the context of the service mesh, there are “application-tier” policies, which apply to the application payload itself. These include coarse-grained WAF rules, fine-grained rules like Spring Cloud Gateway payload validation, and the validation of request semantics via tools like the Open Policy Agent (OPA). Many can even be enforced by a service mesh, but those policies are beyond the scope of this document.

The difficulty with having all network-tier policies is that policies expressed through firewall rules have to be continuously changed, depending on the application pair behind those firewalls. The flexibility in having multi-tier policies is that network-tier policies can be relatively static while identity-tier policies higher up in the stack (e.g., service to service) can be dynamic, as illustrated in Fig. 2.



488

489

Fig. 2. Flexibility provided by multi-tier policies

Implementing identity-tier policies is also a more agile process that allows for new policy capabilities, such as writing policy in terms of identity and application-level action and verb. For example, a network-tier policy would describe the subnets that contain application instances of the client being allowed to call the subnet on a specific port. In contrast, an identity-tier policy would allow the client application identity to communicate with the server application identity via HTTPS on port 443 and execute only the GET method on the /public path. The full range of policies that an enterprise ZTA implemented via a service mesh can enable is outlined in [2] and [3].

Implementing multi-tier policies by relaxing network-tier policies (e.g., by allowing communication across a set of gateways) while introducing identity-tier policies with advanced layer seven controls results in a better overall security posture than either a purely identity-tier or purely network-tier approach.

502 **4. Implementing Multi-Tier Policies for ZTA for Cloud-Native Application** 503 **Environments**

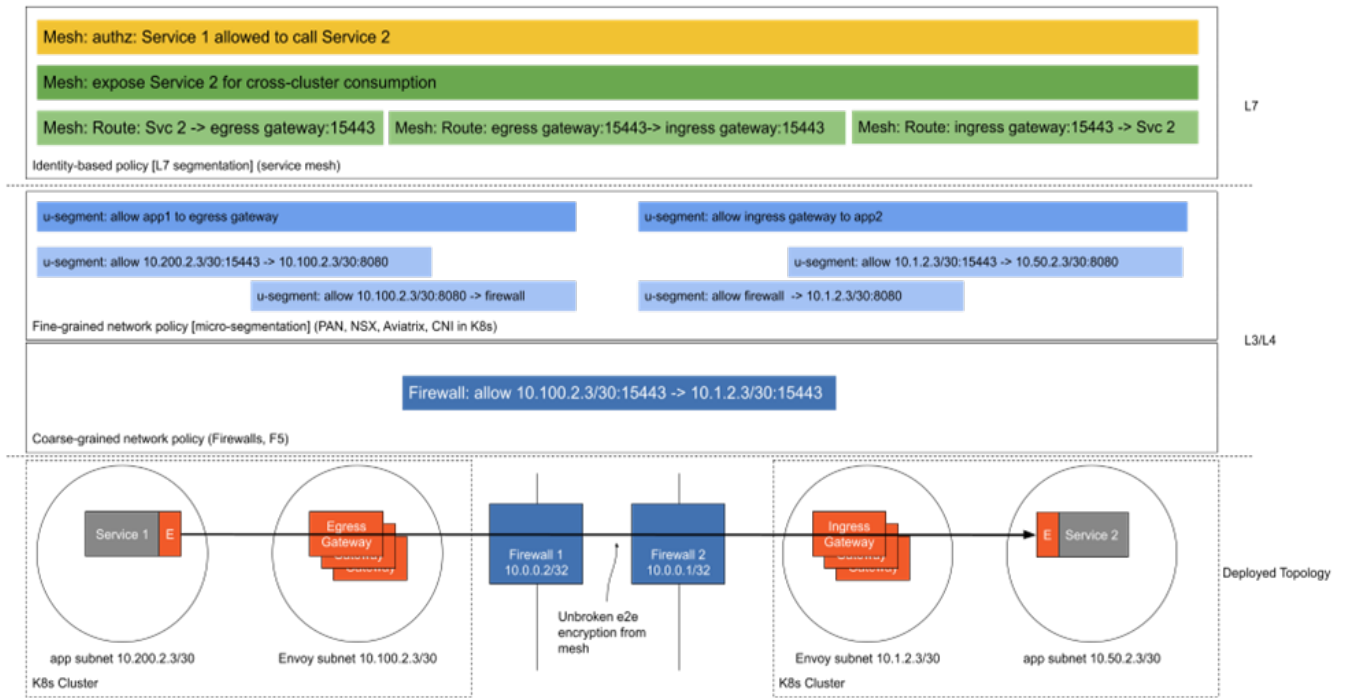
504 This section will consider the implementation of multi-tier policies for realizing an enterprise
505 ZTA using a reference enterprise scenario in which an enterprise hosts microservices
506 applications in several clusters. Each cluster is serviced with a service mesh instance, and
507 clusters are spread out both on-premises and in multiple clouds.

508 Section 4.1 outlines a simple application infrastructure scenario, and Section 4.2 presents a
509 sample set of associated policies that are relevant for that context. Section 4.3 shows how the
510 same set of policies can be defined and deployed for a realistic application infrastructure scenario
511 in which the incoming traffic comes through a DMZ.

512 **4.1. Reference Application Infrastructure Scenario**

513 Consider an application infrastructure of an enterprise where the application topology spans a
514 cloud and on-premises environment. The applications are implemented as microservices with a
515 service mesh instance for each cluster. Hence, a sidecar proxy is associated with each service. At
516 the entry and exit points of each cluster are ingress and egress gateways, respectively. The same
517 data plane (e.g., open-source Envoy) can be used to implement both the sidecar proxy and the
518 transit gateways.

519 Next, consider establishing policies for a scenario that involves two services – Service 1 and
520 Service 2 – that reside in clusters in a cloud and on-premises, respectively. Service 1 in the cloud
521 cluster can interact with services outside of the cluster through an egress gateway. Similarly, all
522 services that attempt to access Service 2 from outside of the cluster have to go through an ingress
523 gateway. All traffic coming out of the cloud has to go through an outbound firewall, and all
524 traffic coming on-premises have to come through an inbound firewall. The paired egress-ingress
525 proxies and the firewall rules that allow them connectivity are collectively referred to as a
526 “transit gateway.” The network location for the two services are each designated by a subnet
527 address. The application topology and policies described so far is shown in Fig. 3.



528

529

Fig. 3. Multi-tier Policies for a Hybrid Application Environment

530 4.2. Role of the Service Mesh in Policy Deployment, Enforcement, and Updates

531 The service mesh has a unique role within the overall policy life cycle activities of policy
532 definition, deployment, enforcement, and update. As already stated, the service mesh is a
533 dedicated infrastructure that provides all application services, including security controls like
534 secure communication and application-level access control. These services are only possible if
535 there are also policies to enforce them during application runtime.

536 Based on the discussion of the control plane in previous sections, it should clear that this
537 component of the service mesh provides access to the interfaces of various policy definition tools
538 through which policies can be defined and updated. Thus, the control plane of the service mesh
539 acts as the *policy administration point*, while the underlying policy tools become the *policy*
540 *decision point*. In addition, the control plane also enables those policies to be distributed to the
541 various proxies described in the previous section. Once distributed, these proxies intercept all
542 traffic in and out of the applications, where it acts as a universal *policy enforcement point*. This
543 allows the service mesh – which centrally manages a fleet of the applications’ proxies – to
544 become the modern cloud-native *security kernel* [3].

545 The proxies – especially the sidecars – can enforce security and traffic policies and generate
546 telemetry data to allow operators to *close the loop* on policy changes by authoring a change,
547 observing its effect on the runtime, and making additional changes as needed in a real-time
548 feedback control loop. In other words, the mesh provides the needed capabilities to implement
549 the runtime controls and achieve a zero trust posture.

550 **4.3. Policy Deployment for Reference Application Infrastructure**

551 Connectivity (between network elements) and access policies (between service instances) are
552 network-tier policies and identity-tier policies, respectively.

553 Consider the following example set of policies that contain a combination of network-tier and
554 identity-tier policies. Network-tier policies can be further categorized into coarse-grained and
555 fine-grained policies.

- 556 • Coarse-grained network-tier policies – These perimeter control policies are informally
557 called firewall rules and are mostly static as they specify:
 - 558 ○ The network location of the egress gateway from which the network edge element
559 at the exit point of a cloud network (e.g., outbound firewall, such as the one at the
560 edge of a cloud) can receive traffic
 - 561 ○ The network location of the ingress gateway to which the incoming traffic that
562 lands at the entry point of the on-premises network edge (e.g., inbound firewall at
563 the entry point to an on-premises network) should be routed:

564 Firewall: allow 10.100.2.3/30 15443 to 10.1.2.3/30:15443

- 565 • Fine-grained network-tier policies – These microsegmentation policies specify the
566 pathways for traffic flowing into and out of the services located within the network
567 subnets at the cloud location or on-premises location.
 - 568 ○ Specify the path on which the outbound traffic from a service or an application
569 (e.g., app 1) can flow. The elements in the path that are specified include the
570 egress gateway at the edge of the cluster and, subsequently, the outbound firewall
571 for the network (cloud network in this example).
 - 572 ○ Specify the path for the inbound traffic into the on-premises network to reach the
573 target application. The elements in the path start from the inbound firewall at the
574 edge of the on-premises network to the ingress gateway in an on-premises cluster
575 to the network subnet where the target service is located.
 - 576 ○ Notably, they specify how traffic can flow “east-west”(i.e., inside of the
577 perimeter). This is in contrast to coarse-grained policies, which specify how
578 traffic can flow “north-south” (i.e., from an external to internal network).
- 579 • Identity-tier policies – These are also called mesh-level policies as they are deployed and
580 enforced at the data plane of the service mesh in the reference platform. In the context of
581 the application infrastructure and the example policies that cover traffic flows from
582 Service 1 to Service 2, these are:

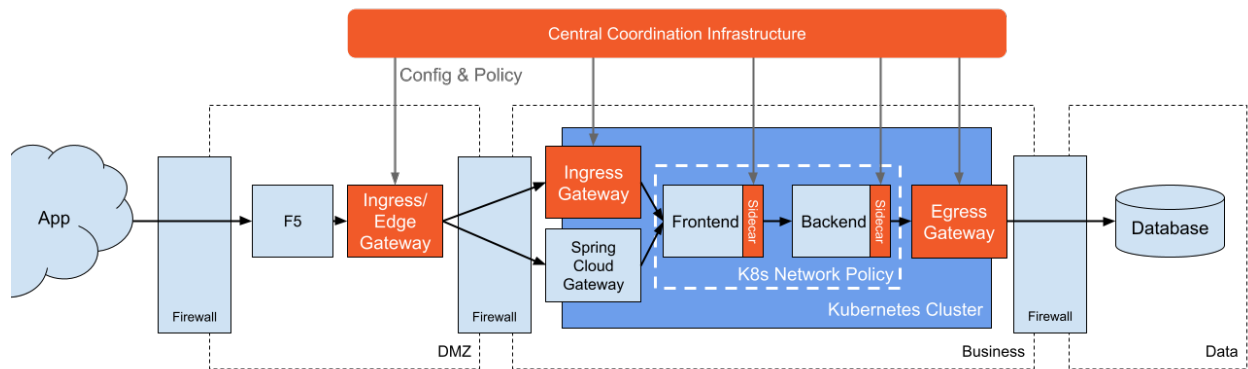
```
selector:  
  matchLabels:  
    app: service-2  
action: ALLOW  
rules:  
- from:  
  - source:  
    principals: ["cluster.local/ns/service-1/sa/service-1"]  
  to:  
  - operation:  
    ports: ["443"]  
    methods: ["GET"]  
    paths: ["/public"]
```

583 **Fig. 4.** An Istio Authorization Policy that allows Service 1 to Service 2 on port 443 but only allows it to
584 execute the GET HTTP verb on the “/public” path

585 This simple example shows some of the advanced capabilities that identity-tier policies can
586 achieve by limiting access based on the application request context. Specifically, this policy
587 limits the application actions to a single HTTP verb on a specific path, but much more
588 sophisticated policy can be implemented as well. See [2] and [3] for detailed overviews.

589 **4.4. Another Application Infrastructure Scenario**

590 Consider another common application scenario in which there is an internal (i.e., within a cluster
591 in the enterprise data center) three-tier application. This application is accessed from outside
592 (through a mobile app or website) through a DMZ. This scenario consists of edge gateways
593 present in the DMZ – an ingress gateway and an egress gateway at the entrance and exit points to
594 and from the data center with firewalls at either side of the gateways. Each of the services that
595 represent the front end and back end (application logic) of the three-tier application have to have
596 a sidecar proxy to enforce policies that pertain to inter-service call requests. This scenario
597 requires the definition and deployment of a combination of network-tier and identity-tier policies
598 that span the various types of gateways and the sidecar proxies. Deploying policies at these
599 multiple locations requires an enterprise-level infrastructure that plays the role of a global control
600 plane, as described in Section 2.1. This is designated as a central coordination infrastructure, as
601 shown in **Fig. 5.**



602

603

Fig. 5. Policy Deployment for a Three-tier Application

604 4.5. Functional Roles of Application Infrastructure Elements in Enforcing 605 Policies

606 This section will show the functionality of each of the application infrastructure elements
607 involved in the policies (e.g., firewalls, gateways, sidecar, transit, and edge proxies) in detail.
608 Since the functionality of firewalls that take part in this context for coarse network-tier policies
609 are well known, this section will focus on the functionality of gateways that take part in fine-
610 grained and identity-tier policies:

- 611 • *Sidecar* – Beside each application instance to intercept all traffic into and out of the
612 application and handles “east-west” internal communication between services in the
613 infrastructure. This is the primary use case of the service mesh.
- 614 • *Ingress gateway* – Controls how applications in the cluster are exposed outside (e.g.,
615 managing what names, certificates, ports, protocols, and application endpoints are served
616 to the world outside of the cluster). Think of this as the service mesh control plane that
617 manages a traditional reverse proxy similar to Spring Cloud Gateway, NGINX, or
618 HAProxy.
- 619 • *Egress gateway* – Controls how applications in the cluster communicate with the outside
620 world. This can be used for traditional egress filtering and logging, like a Squid proxy,
621 but can also implement identity-based policy for what is allowed to call out and perform
622 *credential exchange*, or presenting a set of credentials (e.g., an mTLS certificate for a
623 partner API), on behalf of the application so that the application does not need to handle
624 them (e.g., communicating via mTLS with the partner API). Think of this as a next-
625 generation identity-aware Squid proxy.
- 626 • *Edge gateway* – Accepts external traffic before the ingress gateway and performs fine-
627 grained load balancing across clusters or sites. It is used to terminate external traffic,
628 enable infrastructure-level failover, deploy blue-green clusters, and facilitate ingress-
629 gateway-per-team deployments without requiring each of those teams to have publicly
630 routable ingress gateways. Think of this as a modern software-based local traffic
631 manager, like F5, that can apply policy per-request rather than per-connection.

632 **4.6. Comparison of Identity-Tier and Network-Tier Policies**

633 While network-tier policies are necessary for geographically distributed application
634 infrastructures and for meeting the compliance requirements of regulators, having a combination
635 of network-tier and identity-tier policies allows for some relaxation of the network-tier policies
636 as any unauthorized traffic flow due to an overlooked network element in the path can be
637 addressed through flexible service identity-tier policies. In order to appreciate the need for the
638 coexistence of both policy tiers, it is necessary to know the characteristics of both tiers of
639 policies. This layering of policies, whose strictness can be tuned per organizational needs at each
640 tier, provides agility and operational ease over status quo perimeter-based models while
641 enhancing the overall security posture of the organization.

642 **4.6.1. Approaches for Deployment and the Limitations of Network-Tier Policies**

643 In this approach, applications and service resources with similar security requirements are
644 grouped into a unique segment, and firewall rules are created to block or allow communication
645 with each group or segment [5]. The segments are created using network layer abstractions (e.g.,
646 VLAN IDs or some other tagging approaches), while policies are defined using network address
647 constructs (e.g., IP addresses and ports). Policies apply to subnets (e.g., VLANs) rather than to
648 individual hosts. The assignment of applications to a particular segment can be based on
649 different criteria, such as “all applications with similar security requirements” or “all tiers (web
650 front end, application logic servers, and database servers) associated with a particular application
651 should run in a single segment.”

652 Each segment is protected by gateway devices, such as intelligent switches and routers or next-
653 generation firewalls, which should have the capacity to react and adapt in response to the threats
654 and changes in the application workflows. Segmentation gateways monitor traffic, stop threats,
655 and enforce granular access across east-west traffic (rarely for north-south traffic) within on-
656 premises data centers or cloud regions. The main difficulty with this approach is in mapping the
657 applications’ security requirements-based segments to corresponding network segments. Another
658 difficulty is change management. The mapping between applications and network identities that
659 are being statically maintained has to continuously be kept in sync with the operational scenario
660 in which the application’s network locations are continuously changing due to performance and
661 security.

662 More modern cloud-native deployments that leverage techniques like container network
663 interface-driven network policy are good improvements because they provide identity-tier style
664 policies (i.e., policy in terms of identities and non-network-oriented nouns) while implementing
665 that policy at the network layer (e.g., via eBPF policy or BGP propagation rules). These are a
666 strong upgrade from traditional microsegmentation because they tend to result in finer-grained
667 policies that are easier for the organization to manage over time. However, they typically lack
668 the ability to apply policy per request in the context of the application, which is needed to
669 achieve identity-based segmentation.

670 **4.6.2. Prerequisites for the Deployment of Identity-Tier Policies**

671 Identity-tier policies use contextual, application-driven identifiers (e.g., “order processing front-
672 end service can communicate with inventory back-end service”) instead of network parameters
673 (e.g., “permit calls from 192.168.10.x subnet to 10.0.0.31”). The identifiers assigned to services
674 at runtime are cryptographic identities, which are used for mutual authentication and
675 authorization during each service request and response.

676 Deploying identity-tier policies requires a standardized infrastructure for creating, issuing, and
677 maintaining tamper-proof service identities. Some of the components of this infrastructure are
678 outlined below and are also discussed in [5]:

- 679 • Creation of application identity: The fundamental requirement to enable this is the
680 assignment of a unique identity to each application or service, just like how each user
681 carries a unique identity (e.g., userid). Prior to the era of cloud-based applications,
682 application requests were validated based on the IP subnet or IP address from which they
683 originated. Since ubiquitous access and multi-clouds have eliminated the concept of
684 network perimeters, authentication and authorization based on those parameters are
685 neither feasible nor scalable. Further, the presence of proxies, network address
686 translations, dynamic infrastructures (e.g., migration of applications between VMs), and
687 load balancers make it impossible for the called application to know the IP address of the
688 calling application in order to make authentication or authorization decisions. A unique
689 application identity is required.
- 690 • Establishment of trust in application identity: The created application (workload or
691 service) identity should not be subject to spoofing and should be continuously verifiable.
692 An example of workload identity is a SPIFFE ID [4], which is a string that uniquely and
693 specifically identifies a workload and is encoded as a Uniform Resource Identifier (URI).
694 The SPIFFE ID is carried in a cryptographically verifiable document called a SPIFFE
695 Verifiable Identity Document (SVID). SPIFFE supports multiple SVID formats, but the
696 most commonly used is an X.509 certificate.
- 697 • Discovery of application resources: There should be a robust and secure method for
698 discovering all of the application dependencies consumed over the network (e.g.,
699 services, SaaS endpoints, network appliances, etc.). This capability is enabled through an
700 authenticated service registry.

701 These allowable flows can be based on either (a) the structure of the application (i.e., “the front
702 end of application 1 can call the back end of application 1”) or (b) a legitimate business
703 transaction (e.g., “order processing application can call the shipping application”). Often,
704 organizations do not know all of the allowable service requests in their infrastructure. However,
705 the observability capabilities of the infrastructure (e.g., the metrics provided by the service mesh)
706 can be leveraged to build a view of “requests made today.” From that view, the organization can
707 begin to create fine-grained policies for allowable service requests. Utilizing this observe-and-
708 lock-down methodology builds the organizational processes required to maintain the life cycle of
709 these policies over time.

710 **4.6.3. Advantages of Identity-Tier Policies**

711 Policies based on service and application identities do not use any infrastructure-related variables
712 (e.g., IP addresses, subnets), so they are environment-agnostic and provide the freedom for the
713 services and applications to be migrated to different environments and still maintain the same
714 policies. In other words, there can be a consistent set of policies across cloud providers and on-
715 premises because the policy follows the application rather than the network.

- 716 • Identity-tier policies enable the automated testing of policies. Policies that are
717 independent of infrastructure can be tested by merely exercising the application and
718 observing the outcomes (e.g., trace the sequence of service calls and requests or
719 responses instead of configuring the infrastructure correctly for test runs).
- 720 • Identity-tier policies enable “policy as code” (PaC). With the availability of tools for the
721 declarative specification of policies through PaC, identity-tier policies can be defined and
722 implemented by incorporating the code into automated workflows, such as CI/CD
723 pipelines.
- 724 • Identity-tier policies enable granular (fine-grained) access control by providing visibility
725 into application call sequences/interdependencies and data flows through request-level
726 tracking, which enables the enforcement of security policies for application traffic that is
727 both north-south and east-west, irrespective of the environment (e.g., corporate data
728 center or cloud infrastructure).

729 Additional advantages include:

- 730 • Write once, enforce everywhere – This means that policy can span environments and
731 topologies (i.e., write a policy once, and enforce it everywhere) rather than bespoke
732 policies per environment.
- 733 • Human-readable primitives – The written policies use human-understandable primitives
734 (e.g., “service A can call service B”) rather than network-oriented primitives (e.g.,
735 “10.1.2.3/30 is allowed to call 10.100.2.3/30 on port 8080”). This context is critical since
736 the lack of context for rules is a key reason for the lack of agility around traditional
737 network policy.
- 738 • Contextual intent is codified in a single policy – There is a single policy, not a set of
739 policies that need to be pieced together to understand their intent. A human can read a
740 policy like “the front-end service is allowed to call ‘GET /foo’ on the back-end service”
741 and understand the access that the policy intends to convey even if, for example, the front
742 end is deployed in the cloud and the back end is deployed on-premises. It is significantly
743 harder to read a set of network peering and firewall rules that allow communication
744 across the DMZ for a set of subnets and understand the access that the policy intends to
745 convey. In turn, this means it is harder to write the wrong policy and easier for a human
746 to understand when a policy is incorrect.

747 Identity-tier policies enable only valid network traffic between the various component services of
748 the application due to the mutual authentication and authorization of the service identities, thus
749 enabling the goals of ZTNA to be met.

750 **5. Summary and Conclusions**

751 This document provides guidance for realizing a ZTA for cloud-native application platforms
752 (microservices with a service mesh infrastructure) in the context of an enterprise environment in
753 which applications are hosted in multi-cluster and multi-cloud deployments. A ZTA consists of
754 deployment artifacts that enforce zero trust principles, which is only possible with robust,
755 flexible, scalable, and granular policies that cover all enterprise resources. A policy framework
756 that consists of network-tier and identity-tier policies to meet these goals has been proposed in
757 this document.

758 The artifacts needed for the definition, deployment, and enforcement of these policies have been
759 discussed along with examples of network-tier policies and identity-tier policies. The
760 applicability of these policies in modern enterprise application infrastructures is also illustrated.
761 Finally, the policies that belong to the two tiers are compared in terms of their advantages and
762 limitations, and the critical role of identity-tier policies for realizing a ZTA in the context of
763 modern cloud-native application infrastructures is emphasized.

764

765 **References**

- 766 [1] Rose S, Borchert O, Mitchell S, Connelly S (2020) Zero Trust Architecture. (National
767 Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP)
768 NIST SP 800-207. <https://doi.org/10.6028/NIST.SP.800-207>
- 769 [2] Chandramouli R, Butcher Z (2020) Building Secure Microservices-based Applications
770 Using Service-Mesh Architecture. (National Institute of Standards and Technology,
771 Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-204A.
772 <https://doi.org/10.6028/NIST.SP.800-204A>
- 773 [3] Chandramouli R, Butcher Z, Aradhna C (2021) Attribute-based Access Control for
774 Microservices-based Applications using a Service Mesh. (National Institute of Standards
775 and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-204B.
776 <https://doi.org/10.6028/NIST.SP.800-204B>
- 777 [4] SPIFFE (2022) *SPIFFE Concepts*. Available at [https://spiffe.io/docs/latest/spiffe-](https://spiffe.io/docs/latest/spiffe-about/spiffe-concepts/)
778 [about/spiffe-concepts/](https://spiffe.io/docs/latest/spiffe-about/spiffe-concepts/)
- 779 [5] Chandramouli R (2022) Guide to a Secure Enterprise Network Landscape. (National
780 Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP)
781 NIST SP 800-215. <https://doi.org/10.6028/NIST.SP.800-215>
782