

Withdrawn Draft

Warning Notice

The attached draft document has been withdrawn, and is provided solely for historical purposes. It has been superseded by the document identified below.

Withdrawal Date August 6, 2021

Original Release Date January 26, 2021

Superseding Document

Status Final

Series/Number NIST Special Publication (SP) 800-204B

Title Attribute-based Access Control for Microservices-based Applications using a Service Mesh

Publication Date August 2021

DOI <https://doi.org/10.6028/NIST.SP.800-204B>

CSRC URL <https://csrc.nist.gov/publications/detail/sp/800-204b/final>

Additional Information

2

3 **Attribute-based Access Control for**
4 **Microservices-based Applications**
5 **Using a Service Mesh**

6

7

8

9

Ramaswamy Chandramouli

10

Zack Butcher

11

Aradhna Chetal

12

13

14

15

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-204B-draft>

16

17 **Draft NIST Special Publication 800-204B**

18

19 **Attribute-based Access Control for**
20 **Microservices-based Applications**
21 **Using a Service Mesh**

22

23

24

25

Ramaswamy Chandramouli
Computer Security Division
Information Technology Laboratory

26

27

28

Zack Butcher
Tetrade
San Francisco, CA

29

30

31

32

Aradhna Chetal
TIAA
New York, NY

33

34

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-204B-draft>

35

36

37

January 2021



38

39

40

41

42

43

44

45

U.S. Department of Commerce
Wynn Coggins, Acting Secretary

National Institute of Standards and Technology
James K. Olthoff, Performing the Non-Exclusive Functions and Duties of the Under Secretary of Commerce
for Standards and Technology & Director, National Institute of Standards and Technology

46

Authority

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

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

60 National Institute of Standards and Technology Special Publication 800-204B
61 Natl. Inst. Stand. Technol. Spec. Publ. 800-204B, 37 pages (January 2021)
62 CODEN: NSPUE2

63 This publication is available free of charge from:
64 <https://doi.org/10.6028/NIST.SP.800-204B-draft>

65 Certain commercial entities, equipment, or materials may be identified in this document in order to describe an
66 experimental procedure or concept adequately. Such identification is not intended to imply recommendation or
67 endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best
68 available for the purpose.

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

75 Organizations are encouraged to review all draft publications during public comment periods and provide feedback to
76 NIST. All NIST Computer Security Division publications, other than the ones noted above, are available at
77 <http://csrc.nist.gov/publications>.

78 **Public comment period: *January 27, 2021 through February 24, 2021***

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 Email: sp800-204b-comments@nist.gov

83 All comments are subject to release under the Freedom of Information Act (FOIA).
84

85

Reports on Computer Systems Technology

86 The Information Technology Laboratory (ITL) at the National Institute of Standards and
87 Technology (NIST) promotes the U.S. economy and public welfare by providing technical
88 leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test
89 methods, reference data, proof of concept implementations, and technical analyses to advance the
90 development and productive use of information technology. ITL's responsibilities include the
91 development of management, administrative, technical, and physical standards and guidelines for
92 the cost-effective security and privacy of other than national security-related information in federal
93 information systems.

94

Abstract

95 Deployment architecture in cloud-native applications now consists of loosely coupled
96 components, called microservices, with all application services provided through a dedicated
97 infrastructure, called service mesh, independent of the application code. Two critical security
98 requirements in this architecture are (a) to build the concept of zero trust by enabling mutual
99 authentication in communication between any pair of services and (b) a robust access control
100 mechanism based on an access control such as ABAC that can be used to express a wide set of
101 policies and is scalable in terms of user base, objects (resources), and deployment environment.
102 This document provides deployment guidance for building an authentication and authorization
103 framework within the service mesh that meets these requirements. A reference platform for
104 hosting the microservices-based application and a reference platform for the service mesh are
105 included to illustrate the concepts in the recommendations and provide the context in terms of
106 the components used in real-world deployments.
107

108

Keywords

109 attribute-based access control; authentication policy; authorization policy; CI/CD; DevSecOps;
110 JSON web token; microservices-based application; mutual TLS; next generation access control;
111 policy enforcement point; role-based access control; service mesh; service proxy; zero trust.

112

113

114

Acknowledgments

115 The authors like to express their sincere thanks to Mr. David Ferraiolo of NIST for initiating this
116 effort to provide a targeted deployment guidance in the form of an authentication and
117 authorization framework in a service mesh environment used for protecting microservices-based
118 applications. They also express thanks to Isabel Van Wyk of NIST for her detailed editorial
119 review.

120

Call for Patent Claims

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

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

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

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

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

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

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

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

145 Such statements should be addressed to: sp800-204b-comments@nist.gov

146

147 **Executive Summary**

148 Two significant features of the application environment in emerging cloud-native applications
149 are:

- 150 • Applications have multiple loosely coupled components called microservices that
151 communicate with each other across the network.
- 152 • A dedicated infrastructure called the service mesh provides all services for the application
153 (e.g., authentication, authorization, routing, network resilience, security monitoring),
154 which can be deployed independently of the application code.

155 With the disappearance of a network perimeter because of the need to provide ubiquitous access
156 to applications from multiple remote locations using different types of devices, it is necessary to
157 build the concept of zero trust into the application environment. Further, the cloud-native
158 applications span different domains and, therefore, require increased precision in specifying
159 policy by considering a large set of variables. The service mesh provides a framework for
160 building these and other operational assurances.

161
162 The framework includes:

- 163 • An authenticatable runtime identity for services, the ability to authenticate application
164 (user) credentials, and encryption of communication in transit and between services
- 165 • A Policy Enforcement Point (PEP) that is separately deployable and controllable from the
166 application; the service mesh's side-car proxies
- 167 • Logs and metrics for monitoring policy enforcement

168
169 The service mesh's native feature to authenticate end-user credentials attached to the request
170 (e.g., using a Java Web Token [JWT]) is augmented in many offerings to provide the ability to
171 call external authentication and authorization systems on behalf of the application. The capability
172 to deploy these authentication and authorization systems as services in the mesh also provides
173 operational assurances for encryption in transit, identity, a PEP, authentication, and authorization
174 for end-user identity.

175
176 The objective of this document is to provide deployment guidance for an authentication and
177 authorization framework within a service mesh for microservices-based applications that
178 leverages the features listed above. A reference platform for hosting the microservices-based
179 application and the service mesh is included to illustrate the concepts in the recommendations
180 and provide context in terms of the components used in real-world deployments.
181

182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213

Table of Contents

Executive Summary iv

1 Introduction 1

 1.1 Service Mesh Capabilities..... 1

 1.2 Candidate Applications 3

 1.3 Scope..... 3

 1.4 Target Audience..... 3

 1.5 Relationship to other NIST Guidance Documents 3

 1.6 Organization of this document 4

2 Microservices-based Application and Service Mesh – Reference Platforms ... 5

 2.1 Reference Platform for Hosting a Microservices-based Application 5

 2.1.1 Limitations of Reference Hosting Platform for Security 5

 2.2 Service Mesh Reference Platform – Conceptual Architecture 6

 2.2.1 Service Mesh Functions for Reference Hosting Platform 7

3 Attribute-based Access Control (ABAC) – Background 9

 3.1 ABAC Deployment for Microservices-based Applications Using Service Mesh
 12

4 Authentication and Authorization Policy Configuration in Service Mesh..... 13

 4.1 Hosting Platform Configuration 13

 4.2 Service Mesh Configuration..... 13

 4.3 Higher-level Security Configuration Parameters 14

 4.4 Authentication Policies..... 15

 4.4.1 Specifying Authentication Policies..... 16

 4.4.2 Service-level Authentication 16

 4.4.3 End User Authentication..... 17

 4.5 Authorization Policies..... 17

 4.5.1 Service-level Authorization Policies..... 18

 4.5.2 End-user Level Authorization Policies 18

 4.5.3 Model-based Authorization Policies..... 20

 4.6 Authorization Policy Elements..... 21

 4.6.1 Policy Types..... 21

 4.6.2 Policy Target or Authorization Scope 21

214 4.6.3 Policy Sources 22

215 4.6.4 Policy Operations 22

216 4.6.5 Policy Conditions..... 22

217 **5 ABAC Deployment for Service Mesh..... 24**

218 5.1 Security Assurance for Authorization Framework Enforcement..... 24

219 5.2 Supporting Infrastructure for ABAC Authorization Framework..... 24

220 5.2.1 Service-to-Service Request (SVC-SVC) – Supporting Infrastructure . 24

221 5.2.2 End User + Service-to-Service Request (EU+SVC-SVC) – Supporting

222 Infrastructure 25

223 5.3 Advantages of ABAC Authorization Framework for Service Mesh..... 26

224 5.4 Enforcement Alternatives in Proxies 26

225 **6 Summary and Conclusions 27**

226 **References 28**

227

228 **1 Introduction**

229 Applications based on microservices-based architecture and an application infrastructure based
230 on service mesh that provides various security services through service proxies have emerged as
231 the widespread application environment for cloud-native applications. With the disappearance of
232 the network perimeter due to the need to provide ubiquitous access to these applications from
233 multiple remote locations using different types of devices, it is necessary to build the concept of
234 zero trust [1] into this application environment. Further, the loosely coupled nature of the
235 components of these cloud-native applications (i.e., microservices) facilitates independent
236 design, development, and agile deployment (e.g., CI/CD [2]) of the constituent microservices,
237 enabling paradigms such as DevSecOps [3] need to be used.

238 The security requirements for microservices-based applications are discussed extensively in [4]
239 and summarized here to provide context for this discussion. They are:

- 240 ● Multiple, loosely coupled microservices communicate through network calls, and these
241 communication links must be protected. In the case of monolithic applications, these
242 communications take place through procedure calls.
- 243 ● The entire network is untrusted, and each microservice is untrusted. Therefore, mutual
244 authentication between microservices and secure communication channels between
245 paired microservices through mechanisms such as mutual TLS (mTLS) are required.
- 246 ● The logging data that pertains to each microservice must be consolidated to obtain a
247 security profile in order for forensics, audits, and analytics to assess the overall health of
248 the application.

249 Operating in multiple security domains and multiple clouds, cloud-native applications require a
250 secure authentication and authorization framework. The critical requirements of this framework
251 when implemented within the service mesh are:

- 252 ● The code that is part of this framework is verifiable and non-bypassable (always
253 invoked), thus satisfying the requirements of a security kernel.
- 254 ● The framework should provide authentication and authorization services at multiple
255 levels: service and end-user.
- 256 ● The framework should be able to support a diverse set of authorization policies.

257 The operational assurances required for meeting the above requirements and others are provided
258 by the service mesh. The specific features in service that enable these are given in the next
259 section.

260 **1.1 Service Mesh Capabilities**

261 A service mesh provides a framework for building a set of operational assurances for an
262 organization. That framework includes an authenticatable runtime identity for services, the
263 ability to authenticate application (user) credentials, encryption in transit of communication
264 between services, a Policy Enforcement Point (PEP) separately deployable and controllable from

265 the application (the service mesh’s side-car proxies), and logs and metrics for monitoring policy
266 enforcement. Using these mesh features, a set of controls can be built for all applications that are
267 part of the mesh (e.g., all traffic is encrypted, all traffic to an application goes through the side-
268 car [PEP]). These controls provide a set of operational assurances for applications in an
269 organization deployed in the service mesh.

270 A significant benefit of the service mesh architecture is that the key piece that allows for these
271 controls to be built—the sidecar proxy deployed next to every application—has more security
272 benefits than the traditional approach of building these operational assurances into the
273 application code. First, the life cycle of the sidecar is independent of the application, making it
274 easier to manage across a fleet (e.g., push updates, ensure a consistent version is deployed
275 everywhere). Second, modern implementations (like Istio) allow for dynamic configuration. It is
276 easy to update policies, and updates take effect immediately and without having to redeploy
277 applications. Finally, the mesh’s centralized control allows security teams to build policies that
278 apply to the entire organization so that application developers who build business value are
279 secure by default.

280 A service mesh provides the ability to authenticate end user credentials attached to the request,
281 like a JSON Web Token (JWT). Many service meshes (e.g., Istio) go further and provide the
282 ability for the mesh’s sidecar to call external authentication and authorization systems on behalf
283 of the application. This grants the ability to move request-level policy enforcement out of the
284 application code, trusting instead on the mesh’s assurance that requests that reach the service
285 have been authenticated and authorized for the action that the request is taking. The mesh can
286 even be configured to pass proof of this to the application. This, coupled with the service mesh’s
287 centralized control, means it is possible for a central team to mandate and manage application-
288 level security across the entire organization, delegating to individual application teams only to
289 specify what permissions are required for each applications’ actions.

290 Using the service mesh architecture also means that authentication and authorization systems can
291 be deployed as services in the mesh. Like any other service in the mesh, they benefit from the
292 operational assurances the mesh provides: encryption in transit, identity, a PEP, authentication,
293 and authorization for end user identity. This makes it cheaper to operate an organization’s
294 authentication and authorization systems securely and reliably.

295 In addition to the service mesh features, the capabilities of the access control model play an
296 important role in the authentication and authorization framework. Attribute-based access control
297 (ABAC) has emerged as a promising approach for supporting multiple authorization policies
298 (third requirement above). As per [5], ABAC is defined as “an access control method where
299 subject requests to perform operations on objects are granted or denied based on assigned
300 attributes of the subject, assigned attributes of the object, (optionally) environmental conditions,
301 and a set of policies that are specified in terms of those attributes and conditions.” The main
302 focus of this document is to provide guidance on an authentication and authorization framework,
303 the latter using ABAC to secure microservices-based applications using service mesh.

304 1.2 Candidate Applications

305 The service mesh is most widely used today with containerized applications but can be extended
306 into other environments, such as stateful applications.

307 1.3 Scope

308 This document focuses on providing guidance for building a secure authentication and
309 authorization framework using components of a service mesh for securing services in
310 microservice-based applications. A reference application hosting platform and a reference
311 service mesh platform have been used as examples to illustrate the recommendations in the
312 context of real-world application artifacts (e.g., containers, VMs, etc.). The chosen reference
313 application platform is the open-source Kubernetes, and the chosen reference service mesh
314 platform is Istio. Application infrastructure components in the service mesh that provide other
315 services like network routing, network resilience, and monitoring are outside of the scope of this
316 document.

317 1.4 Target Audience

318 The target audience of the guidance document for developing an authentication and authorization
319 framework for microservices-based applications using the service mesh includes:

- 320 ● Security solutions architects who want to protect the application workloads in microservices-
321 based applications.
- 322 ● Platform architects who want to incorporate a service mesh into the platform offered by their
323 organization to its developers
- 324 ● Developers who want to develop authentication and authorization plug-ins in this application
325 environment

326 1.5 Relationship to other NIST Guidance Documents

327 This guidance document focuses on building an authentication and authorization framework
328 within the service mesh used for securing microservices-based applications. The following
329 publications provide background information for the contents of this document:

- 330 ● Special Publication (SP) 800-204, *Security Strategies for Microservices-based Application*
331 *Systems* [4], discusses the characteristics of microservices-based applications and the overall
332 security requirements and strategies for addressing those requirements.
- 333 ● Special Publication (SP) 800-204A, *Building Secure Microservices-based Applications*
334 *Using Service-Mesh Architecture* [6], provides deployment guidance for various security
335 services (e.g., authentication and authorization, security monitoring, etc.) for a microservices-
336 based application using a dedicated infrastructure (i.e., a service mesh) that uses service
337 proxies that operate independent of the application code.

338 **1.6 Organization of this document**

339 The organization of this document is as follows:

- 340 ● Chapter 2 provides an overview of a microservices-based application, its security
341 requirements, components of a service mesh, and a brief description of the overall
342 architecture of the reference hosting platform and the reference service mesh platform.
343 The latter two are used as examples to illustrate the building blocks involved in the
344 deployment recommendations.
- 345 ● Chapter 3 outlines the advantages of ABAC for the application environment and
346 describes the functional architecture for two of the standard ABAC representations.
- 347 ● Chapter 4 discusses the building blocks of the authentication and authorization
348 framework, the basic configuration that is required in the reference hosting and reference
349 service mesh platform for implementing the framework, and the salient features of the
350 framework.
- 351 ● Chapter 5 provides recommendations regarding deployment of the various use cases
352 pertaining to authorization policies as well as the building blocks (policy components) of
353 these policies.
- 354 ● Chapter 6 provides the summary and conclusions.

355

2 Microservices-based Application and Service Mesh – Reference Platforms

The objective of this document is to offer recommendations for the deployment of an authentication and authorization framework for microservices-based applications within a service mesh that provides the infrastructure for various services, including critical security services. A reference platform for hosting microservices-based applications and the service mesh is included to provide clarity and context for concepts and recommendations in real-world application environments. A brief description of these reference platforms is also provided in terms of their overall architecture and salient building blocks.

2.1 Reference Platform for Hosting a Microservices-based Application

Kubernetes is an orchestration and resource management system widely used for microservices-based applications. In a large application, there will be several microservices, each of which is implemented as a container. Scalable, automated means are required for deployments, operations, upgrading services, and monitoring the health of these containers. The Kubernetes architecture provides the tools to achieve these goals.

To enable application-level, fine-grained access control, it is imperative to have some cluster-level security mechanisms for the clusters that are configured using the hosts of the application components (i.e., microservices). Considering a scenario where the host is a worker node of a Kubernetes platform cluster and the application components are running inside of a container with a pod (i.e., a group of containers) as a deployment artifact, the following cluster-level security measures are required. These measures are defined and enforced through artifacts called pod security policies.

For example, one of the most well-known features of Kubernetes is pod-level *horizontal scaling*. This means that when services receive more traffic, more instances will be generated across machines that grow or shrink on demand. Kubernetes supports auto-vertical scaling on the pod level. Thus, a cluster could be configured to scale the machine on which a pod runs up or down to more accurately fit the anticipated power needs of any microservice. For example, if certain subsets of worker nodes saw spikes in traffic at key times, with the right usage analysis, one could potentially reschedule across machines in order to save costs and optimize performance [7].

Similarly, Kubernetes offers features to monitor the health of the microservices (check the status and readiness). The data to perform these functions is configured in declarative deployment documents, typically as YAML, that describe the port that a pod's containers are listening on. One can specify what to do when services do not start, do not perform as normal, or exit unexpectedly.

2.1.1 Limitations of Reference Hosting Platform for Security

Microservices-based applications require several application infrastructure and security services, such as authentication, authorization, monitoring, logging, auditing, traffic control, caching,

393 secure ingress, service-to-service, and egress communication. Moreover, the following
394 advantages of API architecture are not fully leveraged in the reference platform [8]:

- 395 • A unified way to apply cross-cutting concerns
- 396 • Out of the box plugins to apply cross-cutting concerns quickly
- 397 • A framework for building custom plugins
- 398 • Managing security in a single plane
- 399 • Reduced operation complexity
- 400 • Easy governance of third-party developers and integrators
- 401 • Saving the cost of development and operations

402 By default, communication between Kubernetes containers is insecure, and there is no easy way
403 to enforce TLS between pods since this would result in individually maintaining hundreds of
404 TLS certificates. Pods that communicate do not apply identity and access management between
405 themselves. Though there are tools, such as [Kubernetes Network Policy](#), that can be
406 implemented to act as a firewall between pods, they are a [layer 3](#) solution rather than a [layer](#)
407 [7](#) solution, which is what most modern firewalls are. This means that while one can know the
408 source of traffic, one cannot peek into the data packets to understand what they contain. It does
409 not allow for making vital metadata-driven decisions, such as routing on a new version of a pod
410 based on an HTTP header. There are Kubernetes ingress objects that do provide a reverse proxy
411 based on layer 7, but they do not offer anything more than simple traffic routing. Kubernetes
412 does offer different ways of deploying pods that do some form of [A/B testing](#) or canary
413 deployments, but they are done at the connection level and provide no fine-grained control or
414 fast failback. For example, if a developer wants to deploy a new version of a microservice and
415 pass 10 % of traffic through it, they will have to scale the containers to at least 10—nine for
416 the old version and one for the new version. Further, Kubernetes cannot split the traffic
417 intelligently and instead balances loads between pods in a round-robin fashion. Every
418 Kubernetes container within a pod has separate log, and a custom solution over Kubernetes must
419 be implemented to capture and consolidate them.

420 Although the Kubernetes dashboard offers features like monitoring pods and checking their
421 health, it does not expose metrics that describe how application components interact with each
422 other, how much traffic flows through each of the pods, or what chains of containers make up the
423 application. Since traffic flow cannot be traced through Kubernetes pods out of the box, it is
424 unclear where on the chain the failure for the request occurred.

425 A service mesh addresses these limitations [9]. This document will first consider the service
426 mesh architecture, followed by implementation of service mesh capabilities in the context of the
427 reference platform (Kubernetes).

428 **2.2 Service Mesh Reference Platform – Conceptual Architecture**

429 A service mesh is the network of microservices that make up applications and the interactions
430 between them. It helps to manage microservices-based applications using two major
431 components:

- 432 1. **Data Plane.** This is the component that performs the actual routing or communication of
433 messages between microservices. It also gathers telemetry data, which helps to monitor
434 the health and state of the services. The traffic that flows through the data plane is thus
435 the application-related (business) data.
- 436 2. **Control Plane.** This is the component that provides an API to define policies. This API
437 is often independent of the platform on which the microservices application and, hence,
438 the Service Mesh runs. The control plane also helps the administrator populate the data
439 plane component with configuration that determines how to route traffic. The control
440 plane is the brain of a service mesh. The traffic that flows through the control plane
441 consists of messages of interaction between service mesh components.

442 The control plane may consist of multiple modules, and the distribution of functionality among
443 these modules may be different in different service mesh offerings. However, they all provide the
444 following core functions:

- 445 a. A module that parses the policy rules defined in the control plane and converts them into
446 configuration parameters in the data plane module (i.e., the sidecar proxy). These policies
447 may pertain to various functions, such as authentication and authorization, service
448 discovery, traffic management (including load balancing), intelligent routing, blue-green
449 deployments, canary rollouts, and much more. It may also include configuration
450 parameters related to resiliency in the service mesh, such as timeout, retry, and circuit-
451 breaking capabilities.
- 452 b. A module that provides all of the infrastructure functionality for authentication,
453 authorization, and establishing a secure, encrypted session while two microservices
454 communicate. These functions include user authentication, credential management,
455 digital certificate management, and traffic encryption.

456 2.2.1 Service Mesh Functions for Reference Hosting Platform

457 In order to describe the generic service mesh functions in the context of the reference platform—
458 which, in this case, is Kubernetes—the deployment details of both the microservices application
459 and service mesh components in that platform must be considered. Since authentication and
460 authorization functions are the focus of this document, discussions for those functions on the
461 Kubernetes platform will be confined to the functions in the service mesh.

462 Since the sidecar proxy code implemented as a container is hosted in the same pod as the
463 microservice container, they share the same network namespace and are present in the same node
464 (e.g., VM or a physical machine). Both containers have the same IP address and share the same
465 IP Table rules. That makes the proxy take complete control over the pod and handle all traffic
466 that passes through it [10].

467 Taking the example of establishing a mutual TLS session, the proxy will interact with the
468 module in the control plane of the service mesh to check whether it needs to encrypt traffic
469 through the chain and establish mutual TLS with the backend pod. Enabling this functionality
470 using mutual TLS requires every pod to have a certificate (i.e., a valid credential), and since a

471 good-sized microservice application may be hosted in hundreds of pods, this may involve
472 managing hundreds of short-lived certificates. This in turn requires the service mesh to have a
473 robust identity, access manager, certificate store, and certificate validation. In addition,
474 mechanisms for identifying and authenticating the two communicating pods are required for
475 supporting authentication policies.

476 A service mesh not only provides various application services during runtime but also supports
477 the DevSecOps development and maintenance paradigm. The development team can concentrate
478 their efforts on efficient development paradigms, such as code modularity and structuring,
479 without worrying about the security and management details of their implementation.

480 The service mesh is reference platform-aware and thus automatically injects sidecar containers
481 into the pods. Once the service mesh inserts the sidecar containers, operations and security teams
482 can enforce policies on the traffic and help secure and operate the application. These teams can
483 also configure monitoring of the microservices applications without interfering with the
484 functioning of the applications.

485

3 Attribute-based Access Control (ABAC) – Background

487 Attribute-based access control (ABAC) is an authorization framework or engine that computes
488 decisions for user access requests based on attributes and policies expressed in terms of attributes
489 [5]. The advantages of ABAC for microservices-based applications using service mesh include:

- 490 • Cloud-native applications span different domains and require increased precision in
491 specifying policy by considering a large set of variables. Because of its scalability with
492 respect to attribute and value stores and associated policies, ABAC can meet this
493 requirement.
- 494 • Attributes and their values associated with users, application objects, resources, and
495 environments are independently assigned. Hence, policies based on the attributes do not
496 create a tight subject/object relationship since access decisions are ultimately dependent on
497 dynamic attribute values.
- 498 • Policies are expressed in terms of attributes without prior knowledge of potentially numerous
499 users and resources that are or will be governed under those policies, and users and resources
500 are independently assigned attribute values without knowledge of policy details. This dual
501 feature enables access control decisions to be based on centralized, enterprise-wide policies
502 while also supporting the DevSecOps approach that provides autonomy to each microservice
503 development team to make all decisions regarding their module, including the resource
504 attribute assignments.

505 Due to the features described above, the ABAC authorization framework is a natural fit for the
506 class of cloud-native applications whose design is based on splitting an application into several
507 loosely coupled modules called microservices with each being developed and deployed by
508 independent teams.

509 The ABAC framework has two standardized, representational structures. One uses a platform-
510 neutral text-based language called eXtensible Access Control Markup Language (XACML)
511 Version 3.0, which has been standardized by OASIS. The other is Next Generation Access
512 Control (NGAC), whose data structure and operations have been standardized under INCITS
513 565-2020 [11] – Information technology – Next Generation Access Control. This standardization
514 includes the APIs of functional components (i.e., PEP, PDP, RAP), allowing for the
515 interoperability of these components from different sources. Further, the PEP interface is
516 common for enforcing policies over both application requests and policy administration requests.
517 The biggest advantage of NGAC is the use of linear time algorithms for computing access
518 control decisions and performing policy reviews (i.e., determining the set of resources that a user
519 can access, determining the set of users that can access a resource) [12,13].

520 The functional architectures for these two representational structures are given in Figures 3.1 and
521 3.2, respectively. A brief description of the modules in these two functional architectures is as
522 follows:

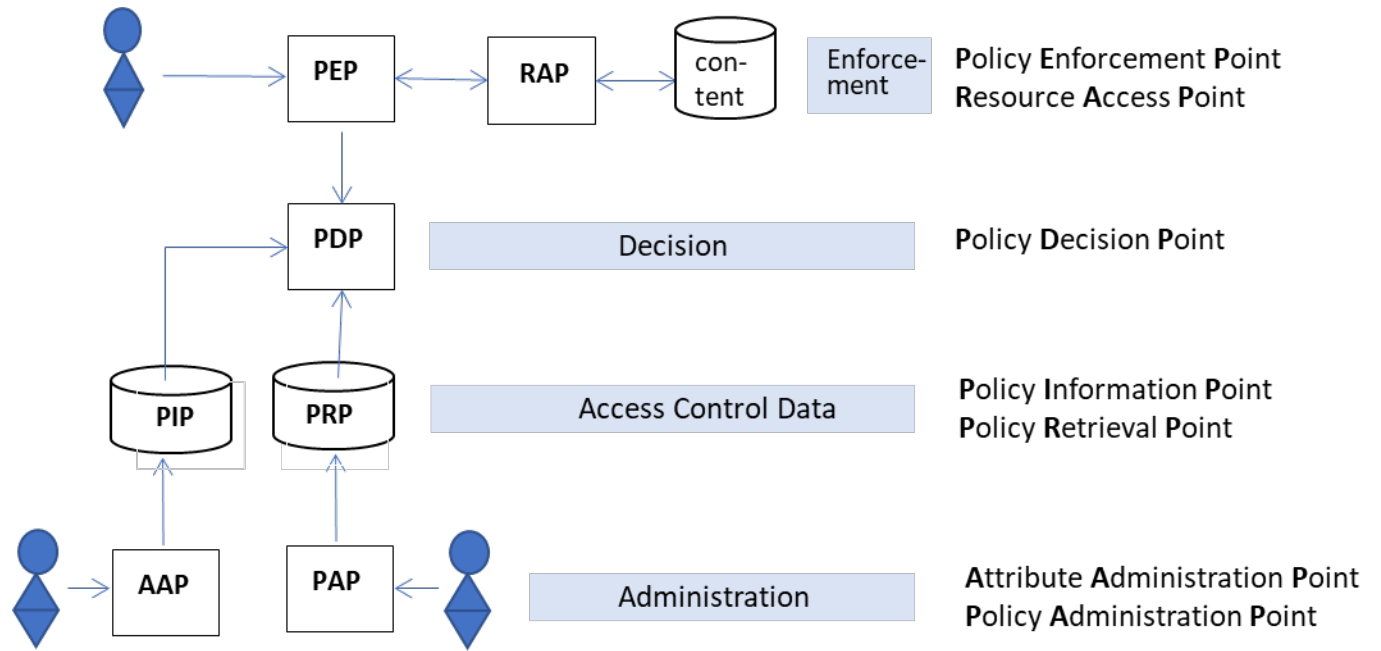
- 523 • Policy Decision Point (PDP) – This is the core module of the ABAC functional

- 524 architecture that computes decisions to permit or deny user access requests for
525 performing actions on resources. Requests are received from and responses sent to a
526 module called Policy Enforcement Point (PEP) in both representations.
- 527 • Policy Enforcement Point (PEP) – This is a module that is part of the application’s
528 platform and tightly integrated with the application. It is designed to intercept all access
529 requests that emanate from the application in both representations.
 - 530 • Policy Information Point (PIP)
 - 531 a. In the XACML representation, this is a module that contains the database of attributes
532 and their associated values for all application-relevant objects or resources. The
533 information here is used to extract the attributes and associated values for users and
534 resources found in the access request to find the applicable target policies in the PRP
535 (described below).
 - 536 b. In the NGAC representation, this is a repository of association relations of the form
537 (u-ai, op-i, o-ai) for a pc-i, where u-ai and o-ai are attribute values associated with a
538 user and object (resource), respectively. op-i denotes a set of allowed operations and
539 pc-i the governing policy classes. To minimize the set of association relations in the
540 authorization database (e.g., having triples to represent every user and every object in
541 the application), containment relations of the form (U < u-ai) are used to show the
542 members of the user group and object group represented in the association relations.
543 In addition, the same set of containment relations are used to denote the applicable
544 policies for each object as well (O < pc-i).
 - 545 • Policy Retrieval Point (PRP) – In the XACML representation, this is the module that is
546 the repository for authorization policies expressed as logical formulas involving
547 predicates on attribute values. The policy representation also contains the target resources
548 that are covered by the policy. The resources requested in the access request are matched
549 to these targets to retrieve the applicable policies by the PDP when computing decisions
550 for those requests. This module is not part of the functional architecture in the NGAC
551 representation.
 - 552 • Attribute Administration Point (AAP) – This is the interface for administering attributes
553 stored in PIP in the XACML representation. This module is not necessary in NGAC
554 representation since its association relations express the access rights on objects
555 instantiated using attribute values.
 - 556 • Policy Administration Point (PAP) – This is the interface for administering policies
557 stored in PRP.

558

559

560




561

562

Figure 3.1 ABAC Functional Architecture based on XACML Representation

563

564



565

566

Figure 3.2 ABAC Functional Architecture based on NGAC Representation

567

568

569 3.1 ABAC Deployment for Microservices-based Applications Using Service Mesh

570 In the context of a microservices-based application using service mesh, an ABAC deployment
571 for access control can take the following forms:

- 572 ● The proxies (e.g., Ingress, sidecar, and Egress) play the role of PEPs since they intercept all
573 requests that emanate from each client, user, service, or external service.
- 574 ● The PEPs can provide either an ALLOW/DENY verdict or a list of allowable objects.
- 575 ● The enforcement function can be provided either natively (using local configuration
576 structures, such as ACLs) or using proxy extensions that call an external authorization server
577 to obtain one or more of the data in the previous bullet.
- 578 ● The assurance mechanisms in the service mesh (e.g., certificate-based authentication, secure
579 session, non-bypassability, execution isolation) can be leveraged to deploy a high assurance
580 authorization framework.

581

582 **4 Authentication and Authorization Policy Configuration in Service Mesh**

583 Fine-grained access control for microservices can be enforced through the configuration of
584 authentication and access control policies. These policies are defined in the control plane of the
585 service mesh, mapped into low-level configurations, and pushed into the sidecar proxies that
586 form the data plane of the service mesh. The configurations enable the proxies to enforce the
587 policies at application runtime (or request time), thus making the proxies act as Policy
588 Enforcement Points (PEPs). As stated in the introduction, the objective of this document is to
589 provide guidance for the deployment of an authentication and authorization framework that is
590 external to the application, agnostic to the platform hosting the application and the service mesh
591 product that implements the application infrastructure. However, Kubernetes is used as the
592 reference application platform and Istio as the service mesh infrastructure platform to provide
593 concrete examples of the concepts and to enable us to make specific recommendations with more
594 clarity and specificity.

595 **4.1 Hosting Platform Configuration**

596 The generic host platform configuration data for microservices-based applications using service
597 mesh that are, at the minimum, needed for authentication and authorization policy configuration
598 are:

- 599 • Metadata, like service name and the sets of instances of that service
- 600 • Runtime data, such as services's protocols and ports
- 601 • Namespaces that provide logical isolation boundaries for sets of services
- 602 • Unique runtime identities for each service

603 In the reference hosting platform Kubernetes, this is realized as:

- 604 • Service resource, which declares a service's name, protocol (e.g. TCP), and ports (e.g.
605 9080)
- 606 • Deployment resource, which declares deployments of pods that implement that service,
607 including metadata such as labels and version
- 608 • Namespace construct and RBAC for managing how users are allowed to publish
609 configuration into namespaces
- 610 • Service Accounts, which are identities unique to each namespace bound to individual
611 services

612 **4.2 Service Mesh Configuration**

613 The installation of any service mesh involves the following components:

- 614 • Ingress Gateway, which is the first point of entry into the microservices-based application.
615 This gateway specification includes names, ports, and routes that the application client
616 must take to access the application.

- 617 ● Egress Gateway for the application to call outside services or applications. Egress
618 gateways are optional since a sidecar proxy can act as an egress proxy for the purposes of
619 policy without deploying an egress gateway.
- 620 ● Injection of sidecar proxies (in the form of containers). The consequence of this is that
621 each of the application's deployments in the platform will now have two containers—the
622 original microservice container plus the mesh's sidecar proxy. These sidecar proxies
623 enforce authentication and authorization policies during application runtime, thus acting
624 as Policy Enforcement Points (PEPs). In addition, proxies should emit metrics and logs to
625 enable continuous monitoring of the system; this can be used to ensure policies are in
626 place and are being enforced.
- 627 ● A Certificate Authority (CA) module is needed to handle certificate requests from sidecar
628 proxies, which need a runtime identity presented as an X.509 certificate. This CA
629 generates, distributes, and manages keys and certificates used by the mesh and enables the
630 mesh to perform automatic certificate rotation. A CRL or OCSP feature is also required to
631 support certificate validation.
- 632 ● A control plane module in the service mesh that monitors configuration data in the hosting
633 platform, encodes policies and distributes those policies in the form of configuration to
634 various proxies in the mesh (e.g., Ingress, sidecar, and Egress).

635 In the context of the reference service mesh platform Istio, to facilitate route specification to the
636 entry service of the application in the Ingress Gateway, a virtual service is defined that specifies
637 the path and hosts making up the virtual service and the first entry service/port to which the
638 gateway must route the incoming request from an application client [14].

639 **SMC-SR-1:** *The signing certificate used by the mesh's CA module should be rooted in the*
640 *organization's existing PKI to allow for auditability, rotation, and revocation.*

641 Some service meshes come with the ability to encrypt traffic using a self-signed certificate; such
642 a certificate should not be used in secure deployments.

643 **SMC-SR-2:** *Communication between the service mesh control plane and the hosting platform's*
644 *configuration server must be authenticated and authorized.*

645 In this reference platform, authentication is typically achieved by the Kubernetes API Server (the
646 configuration server) with simple TLS. Authentication of the client is based on the pod's service
647 account credential. Authorization for the client to receive platform information from the API
648 Server is enforced by Kubernetes RBAC.

649 **4.3 Higher-level Security Configuration Parameters**

650 Since the component microservices of our application are generally implemented as containers,
651 the following higher-level security configuration parameters should be set. In the reference
652 hosting platform Kubernetes, containers are implemented in pods, which contain a microservice
653 container as well as a sidecar container. These higher-level security configurations are set through
654 flags that come under the banner of pod security policies. The recommendations for these flag

655 values are numbered using the acronym HLC-SR-X, where HL stands for higher-level
656 configuration, SR stands for security recommendation, and X is the sequence number. They
657 include but are not limited to the following [5]:
658

659 **HLC-SR-1:** *Containers and applications should not be run as root (thus becoming privileged*
660 *containers).*

661 In Kubernetes, the configuration setting for this is to set the value TRUE for
662 “`MustRunAsNonRoot`” flag.
663

664 **HLC-SR-2:** *Host path volumes should not be used as they create tight coupling between the*
665 *container and the node on which it is hosted, constraining the migration and flexible resource*
666 *scheduling process.*

667 In Kubernetes, the configuration setting for this is to set the value of TRUE to
668 “`readOnlyRootFilesystem`” flag.
669

670 **HLC-SR-3:** *Configure the container file system as read-only by default for all applications,*
671 *overriding only when the underlying application (e.g., database) must write to disk.*
672

673 **HLC-SR-4:** *Explicitly prevent privilege escalation for containers.*

674 In Kubernetes, this is achieved by setting the value FALSE for the
675 “`allowPrivilegeEscalation`” flag.

676 **4.4 Authentication Policies**

677 Authentication policies specify the process for validating identities. The integrity of this process
678 and its strength determines the integrity of the authorization process since the latter depends upon
679 the strength of the authenticated identity. There are two types of identity needed in a
680 microservices-based application:

- 681 • Microservices or workload identity
- 682 • End-user identity

683 Service (microservice) identity is critical for the following reasons:

- 684 • It enables the client to verify that the server to which it is communicating (server identity
685 validated using the certificate it carries) is authorized to run the service. This assurance
686 has to be provided by a secure naming service that maps the server identity to the service
687 identity. In any orchestration platform (including Kubernetes), services can be moved
688 around the nodes (server) for load balancing and service availability reasons. It is the
689 responsibility of the control plane of the service mesh to refresh this mapping information
690 by interacting with the API that contains this configuration information (e.g., through API
691 server in Kubernetes) and convey it to the sidecar proxy in the data plane of the service
692 mesh.
- 693 • The service identity is the basis for the target service to select and enforce applicable
694 authorization policies.

695 4.4.1 Specifying Authentication Policies

696 Associated with these identities are the corresponding authentication processes that the service
697 meshes have to support. They are:

- 698 ● Service-level authentication or peer authentication using service identity
- 699 ● End user authentication or request authentication using end user credentials

700 It is assumed that the reference hosting platform has been configured with the high-level
701 requirements outlined in Section 4.1. It is also assumed that the reference service mesh platform
702 has been installed and configured with the initial requirements outlined in Section 4.2.

703 4.4.2 Service-level Authentication

704 Service-level authentication is the mutual authentication of the communicating services and setup
705 of a secure TLS session. Enabling this requires the capability to define a policy object which
706 should meet the following requirements:

707 **AUN-SR-1:** *A policy object relating to service-level authentication should be defined that requires*
708 *that mTLS be used for communication. The policy object should be expressive enough to be defined*
709 *at various levels (given below) with features for overrides at the lower levels or inheritance of the*
710 *requirement specified at the higher levels. The following are the minimum required levels [6]:*

- 711 a. Global level or the service mesh level
- 712 b. Namespace level
- 713 c. Workload or microservices level – used for applying authentication and authorization
714 policies for a subset of traffic to a subset of resources (e.g., particular microservices, hosts
715 or ports)
- 716 d. Port level, taking into account that certain traffic is designed for communicating through
717 designated ports

718 This form of authentication also requires the assignment of a strong identity to each service and
719 the authenticating of that identity by mapping it to the server identity (where the service is
720 hosted) that digitally signed in a special digital authentication certificate (SPIFFE). To provide
721 assurance that the server whose identity is found in the SPIFFE certificate is the one that is
722 authorized to run the target service, the following requirement (also specified in SP 800-204A) is
723 needed:

724 **AUN-SR-2:** *If the certificate used for mTLS carries server identity, then the service mesh should*
725 *provide a secure naming service that maps the server identity to the microservice name that is*
726 *provided by the secure discovery service or DNS. This requirement is needed to ensure that the*
727 *server is the authorized location for the microservices and to protect against network hijacking.*

728 The information for mapping the server identity to a service is obtained by the control plane of
729 the service mesh by accessing the configuration information from the platform that is hosting the

730 microservices-based application. In Kubernetes, the control plane of the service mesh obtains the
731 mapping information through the API server module of the Kubernetes platform and populates
732 that information in the secure naming service. Thus, the mutual certificate validation not only
733 enables validation of the associated service identities of both the client and target services but
734 also enables creation of a secure mutual TLS (mTLS) session. In Istio, the policy object for this
735 type of authentication is called “peer authentication.”

736 4.4.3 End User Authentication

737 For the mesh to authenticate end user credentials (EUC), the application must participate in some
738 way. Client services that make the request should acquire and attach an appropriate credential to
739 each request (e.g., a JWT) in the request header. End user authentication, or request
740 authentication, is the process of validating the credentials of the end user making a request by
741 extracting from the request’s metadata and authenticating them (locally or against an external
742 server). For example, a common flow at many organizations is to exchange an external EUC, like
743 an Oauth bearer token, at ingress for an internal credential that is encoded within a JSON Web
744 Token (JWT). The JWT can be created by a custom authentication provider or standards-based
745 OpenID Connect provider.

746 **EAUN-SR-1:** *A request authentication policy must, at the minimum, provide the following*
747 *information:*

- 748 ● *Instructions for extracting the credential from the request*
- 749 ● *Instructions for validating the credential*

750 For a JWT, this might include:

- 751 ● Location (header name) of the JWT token that contains the user’s claims
- 752 ● How to extract the subject, claims, and issuers from the JWT
- 753 ● Public keys or the location for the key used for validating the JWT

754 4.5 Authorization Policies

755 Authorization policies, just like their authentication counterparts, can be specified at the service
756 level as well as the end user level. In addition, authorization policies are expressed based on
757 constructs of an access control model and thus may vary based on the nature of the application
758 and enterprise-level directives. Further, the location of the access control data may vary
759 depending on the identity and access management infrastructure in the enterprise. These
760 variations result in the following variables:

- 761 ● Two authorization levels – service level and end user level
- 762 ● Access control model used to express authorization policies
- 763 ● Location of the access control data in a centralized or external authorization server or
764 carried as header data

765 The supported access control in the service mesh uses abstraction to group one or more policy
766 components (described below in Section 4.5.1) for specifying either service-level or end user-
767 level authorization policies. Since microservices-based applications are implemented as APIs
768 (e.g., RESTful API), authorization policy components described using key/value pairs will have
769 attributes pertaining to an API, including the associated network protocols. The types of
770 authorization policies are:

- 771 ● Service-level authorization policies
- 772 ● End user-level authorization policies
- 773 ● Model-based authorization policies

774 4.5.1 Service-level Authorization Policies

775 Service-level authorization policies are defined using a policy object that provides positive or
776 negative permission (authorization) with the following policy components:

- 777 a. The scope of the policy can span all applications at the service mesh level, namespace
778 level, or one or more designated applications (microservice level).
- 779 b. The permissions or operations can be restricted to one or more designated methods of a
780 given service (e.g., an “HTTP GET method on the ‘/details’ path of an application named
781 PRODUCT-CATALOG”) or to designated ports through which an application can be
782 accessed.
- 783 c. Conditions under which access can take place (e.g., possession of a token) are specified.
- 784 d. Sources allowed access are specified at the namespace or a particular service level (in
785 terms of the service’s runtime identity).

786 **AUZ-SR-1:** *A policy object describing service-to-service access should be in place for all*
787 *services in the mesh. At a minimum, these policies should permit access at the namespace level*
788 *(e.g., “services in namespace A can call services in namespace B”).*

789 Ideally, policies should describe the minimum access required for application functionality (e.g.,
790 “service ‘foo’ in namespace A can perform ‘GET /bar’ on service ‘bar’ in namespace B”).

791 4.5.2 End-user Level Authorization Policies

792 Given an authentication policy like Section 4.4.3, a sidecar in the mesh can extract a principal
793 from the request to perform authorization on. Further, the sidecar typically has additional context
794 about the request, including the resource being accessed (e.g., the path in an HTTP/REST API)
795 and the action being taken (e.g., the HTTP verb – GET, PUT, etc. – in the request to that API).
796 This gives the sidecar enough information to act as a policy enforcement point and call a policy
797 decision point.

798 This is the most common case, especially for organizations with traditional IAM systems that
799 exist as an external service, often called by an SDK. To handle this case, a service mesh’s sidecar
800 proxy will typically support calling external services to render an authentication and authorization

801 verdict. For example, the reference implementation Istio supports this via Envoy's (i.e., the
802 sidecar proxy) external authorization service [15].

803 **EUAZ-SR-1:** *When a sidecar communicates with an authentication or authorization system, that*
804 *communication should be secured with the mesh's built-in service-to-service authentication and*
805 *authorization capabilities.*

806 Logs and metrics exported by the sidecar can be used to prove that authentication and
807 authorization was performed by the sidecar on behalf of the application.

808 End user authorization is not applied to the decision endpoint of the external authorization (PDP)
809 service since the service is the principal making the call. It also avoids needing a default policy
810 that allows all users to call the decision endpoint of the PDP. End user authorization should be
811 applied to the PAP and other administrative endpoints of the authorization system, and that can
812 be facilitated by the mesh.

813 However, there is another case that is common enough to address in which an external
814 authorization system is not required. Making a network call to an authorization service for every
815 hop in a service chain can be expensive and cause centralized failures. To mitigate these
816 problems, many organizations will exchange end user credentials at ingress for an internal,
817 trusted, authenticatable credential that conveys not just the user's principal but also that user's
818 capabilities in the system. A JSON Web Token (JWT) is frequently used for this because it is
819 locally authenticatable and conveys the user's principal (the JWT's subject), the issuer of the
820 JWT (issuer), and arbitrary claims that the organization can control (e.g., to use for access
821 control).

822 Performing end user authorization based on a JWT is common enough that it is built directly into
823 Envoy, the sidecar proxy of the reference mesh, Istio. Envoy can be configured with a filter [16]
824 that will process requests in two steps:

- 825 a. JWT token verification involves extracting the token from the request header, verifying
826 whether issuers and audiences are allowed, fetching the public key, and verifying the
827 digital signature on the token.
- 828 b. Match the resources in the request to the claims in the token to determine whether the end
829 user should be allowed access to the requested resources or denied.

830 Envoy's JWT filter act as the PDP, making the access decision entirely locally. This requires that
831 policy documents be small enough to reside on an individual sidecar proxy. Although a full
832 ABAC is ideal for handling resource-level policies, the JWT filter is valuable as a stepping stone
833 from a traditional system that only performs access control on the edge to a zero trust system that
834 performs authentication and authorization at each service.

```

apiVersion: security.istio.io/v1beta1
kind: AuthorizationPolicy
metadata:
  name: backend
  namespace: product
spec:
  action: ALLOW
  rules:
  - from:
    - source:
      principals: ["cluster.local/ns/product/sa/frontend"]
    to:
    - operation:
      methods: ["GET"]
      paths: ["/info*"]
    - operation:
      methods: ["POST"]
      paths: ["/data"]
  when:
  - key: request.auth.claims[iss]
    values: ["accounts.google.com"]

```

Figure 4.1 – An example Istio authorization policy

This allows the front end to call specific methods on the backend only if the request has an EUC attached issued by “accounts.google.com.”

835

836 **EUAC-SR-2:** *All application traffic should carry end-user credentials, and there should be a*
 837 *policy in the mesh enforcing that credentials are present.*

838 We recommend this even if the application is enforcing authentication and authorization
 839 independently of the mesh, because these organization-wide controls allow functionality like
 840 audit to be built on top of the mesh at lower cost to central teams responsible for compliance and
 841 controls.

842 **4.5.3 Model-based Authorization Policies**

843 The service-level authorization policies and a use case of end-user authorization policies that uses
 844 JWT are natively implemented in the proxies. Since these cannot be used for resource-level
 845 authorization policies, we need to support model-based authorization policies as well. As already
 846 alluded to in section 4.5.2, this requires a call from the proxy to an external authorization server
 847 which holds the model-based authorization engine to obtain an access decision.

848 The service principals in these model-based policies are identities (e.g., ServiceAccount) that are
849 provided by the underlying application orchestration platform (e.g., Kubernetes) and is the same
850 that are used by authorization policies natively supported in the proxies. The user principals are
851 usually obtained from the JWT token. The popular access control models in the external
852 authorization servers are either RBAC or ABAC.4.6 Authorization Policy Elements

853 **4.6 Authorization Policy Elements**

854 The authorization policies that can be specified in a service mesh may consist of the following
855 elements:

- 856 ● The policy types – Positive (ALLOW) or Negative (DENY)
- 857 ● The policy target or authorization scope – The namespace, a particular service
858 (application name), and version
- 859 ● The policy sources – Covers the set of authorized services
- 860 ● The policy operations – Specifies the operations on the target resources that are covered
861 under the policy
- 862 ● The policy conditions – The metadata associated with the request that must be met for the
863 application or invocation of the policy

864 **4.6.1 Policy Types**

865 Positive and negative policies are specified in order to set precedence relationships (e.g., DENY
866 overrides, ALLOW, etc.). They are also used for situations that allow one type of policy for all
867 services under a group and to specify exceptions (e.g., have an ALLOW policy for all services in
868 a namespace but a DENY policy for a specified service)

869 **4.6.2 Policy Target or Authorization Scope**

870 This refers to the target resources in terms of a set of services, versions, and the namespaces
871 under which the services are located. The service can be specified in the following ways:

872 Using path: The location of the target resource is specified using paths (e.g., for resources
873 accessed using HTTP or gRPC protocols). The list of paths to be included in the authorization
874 policy scope and paths that need to be excluded can be defined. Both of these sub-elements of the
875 policy target component (i.e., the list of paths to be included and the list of paths to be excluded)
876 are optional.

877 Using host name: In some instances, the target resources are specified using the host sub-element.
878 The list of hosts to be included in the authorization policy scope as well as those hosts that need
879 to be excluded can be defined. Both of these sub-elements of the policy target component (i.e.,
880 list of hosts to be included and the list of hosts to be excluded) are optional.

881 Using network ports: The network port through which the target resource (the service) is accessed
882 is often specified using the port sub-element. The list of ports to be included in the authorization

883 policy scope as well as those ports that need to be excluded can be defined. Both of these sub-
884 elements of the policy target component (i.e., list of ports to be included and the list of ports to be
885 excluded) are optional.

886 **4.6.3 Policy Sources**

887 The policy sources are the set of services that are authorized to operate on the set of resources
888 specified under the policy target (specified using name, path, host name, and ports). The policy
889 sources are usually specified using a service account or name (called principal), all services in a
890 particular logical group (e.g., namespace), or all services that are accessed from a group of
891 network locations (e.g., IP blocks). Both included and excluded principals, namespaces, and IP
892 blocks can be specified in some implementations.

893 **4.6.4 Policy Operations**

894 The set of operations depends on the way the application is implemented. If the application is
895 implemented as a REST API, the following are the common operations (also called HTTP verbs
896 or HTTP methods):

897 POST: This is equivalent to creating a resource.

898 GET: This is equivalent to reading the contents of the resource.

899 PUT: This is equivalent to updating the resource by replacing.

900 PATCH: This is equivalent to updating the resource by modifying.

901 DELETE: This is equivalent to deleting the resource.

902 OPTIONS:

903 HEAD:

904 If the resource is accessed using gRPC instead of a RESTful protocol, there is only one operation
905 or method: “POST.” The authorization policy definition may also have a feature to specify the list
906 of operations (methods) to be excluded. Both policy sub-elements—one to specify the operations
907 to be included in the authorization policy scope and the other to be excluded—are optional.

908 **4.6.5 Policy Conditions**

909 Policy conditions specify the constraints in the form of a key-value pair for the metadata
910 associated with the request. This metadata may cover the following:

911 *Metadata associated with the source:* Some of the metadata (e.g., service account name,
912 namespace, and IP blocks) are specified as part of the policy source specification itself. In
913 addition, it is possible to list IP addresses in CIDR format of the policy sources.

914 *Metadata associated with the request:* In this type of metadata, the parameters or attributes that
915 pertain to a specific request can be specified. These parameters can include an audience that can
916 present the authentication information expressed in the form of a URL (only applicable to HTTP
917 protocol-based requests), a specific end user identifier associated with the audience that can
918 present the authentication credentials, or the claim name that is carried in the token presented by
919 the presenter. In addition, parameters that pertain to the user-agent (e.g., browser name) can also
920 be specified for HTTP protocol-based requests.

921 *Metadata associated with the destination:* The range of allowable IP addresses can be specified in
922 CIDR format as well as the associated list of ports.

923

924

925 **5 ABAC Deployment for Service Mesh**

926 The last chapter introduced three different types of authorization policies including two use cases
927 for end-user level authorization policies. This chapter, we will leverage those architectural choices
928 to describe an ABAC-based authorization framework in the service mesh:

- 929 ● Security assurance for authorization framework enforcement
- 930 ● Supporting infrastructure for authorization requests
- 931 ● Advantages of ABAC Authorization framework for Service Mesh
- 932 ● Enforcement alternatives in Proxies

933 **5.1 Security Assurance for Authorization Framework Enforcement**

934 The authorization policy enforcement mechanism implemented in the service mesh for a
935 microservices-based application must satisfy the three requirements of a reference monitor
936 concept. It must be 1) non-bypassable, 2) protected from modification, and 3) verified and tested
937 to be correct. These three requirements can be ensured by the following:

- 938 ● Every request from a client to the microservices-based application, from one service to
939 another (inter-services call), and from a microservice to an external application is
940 intercepted by the ingress gateway, sidecar proxy, and egress proxy, respectively, and
941 these policy enforcement points (PEPs) are non-bypassable.
- 942 ● The policy enforcement modules are independent executables that are decoupled from the
943 application logic and cannot be modified.
- 944 ● Their outcome can be independently verified and tested through both shadow operations
945 and live production requests.

946 In short, a proxy running in the data plane of the service mesh is the reference monitor with
947 respect to authorization enforcement. The authorization policy engine (e.g., NGAC-based ABAC
948 policy engine) implemented as a container executing either natively in the proxy memory space
949 or callable from a corresponding filter module in the proxy runs as a separate process that does
950 not share any memory space with the calling application. Hence, it satisfies the requirement of a
951 security kernel.

952 **5.2 Supporting Infrastructure for ABAC Authorization Framework**

953 We will now look at the basic building blocks of the supporting infrastructure for service-to-
954 service and end-user+service-to service requests.

955 **5.2.1 Service-to-Service Request (SVC-SVC) – Supporting Infrastructure**

956 The policy object used for authorizing this type of request was described in Section 4.5.1.
957 Service-to-service requests must be authorized based on the identity of the calling and called
958 services. The trusted document that carries the identity of the service is an X.509 certificate

959 issued by one of the control plane components of the service mesh after verifying whether the
960 requested identity is valid for the microservice by consulting an identity registry. The proxy
961 communicates with this control plane component through a local agent, obtains a certificate, and
962 sends it to the proxy, which then performs the certificate validation process on behalf of the
963 calling service or client during each service request. The identity is encoded as URI and carried in
964 a certificate's SAN (subject alternate name) field. It must be mentioned that the certificates that
965 carry service account identities are short-lived certificates (rotated every hour or few hours)
966 rather than the conventional HTTPS TLS terminating certificates whose validity lasts for several
967 months.

968 **5.2.2 End User + Service-to-Service Request (EU+SVC-SVC) – Supporting Infrastructure**

969 The policy object used for authorizing this type of request was described in Section 4.5.2. This
970 request type requires the verification of two identities: the calling user identity and the service
971 identity. As described in the previous section, the service mesh provides the feature to perform
972 authorization based on service identities. Since this is a standard feature, no extra components
973 need to be built in the service mesh infrastructure for this type of authorization. However, when
974 end user identities are introduced for authorization, the authorization framework should be tightly
975 integrated with the following components of the architecture:

- 976 ● The services orchestration control plane for obtaining application object attributes as well
977 as attributes of the registered application users (which includes user credentials), thus
978 playing the role of Policy Information Point (PIP) in ABAC-based authorization
- 979 ● A service mesh control plane for obtaining tokens that encode the claims based on the
980 authorization decision
- 981 ● A service mesh data plane in the service proxy for making calls to the authorization
982 engine (which is just another service), obtaining the authorization decision, enforcing the
983 service-to-service authorization policies, making calls to the service mesh control plane
984 for authorization tokens (e.g., Java Web Tokens [JWT]), and attaching the tokens to the
985 service request.

986 The advantage of an EU+SVC-SVC request processing scheme is that authorizations at a finer
987 level of granularity than the method level can be specified, and conformant claims can be
988 included in the authorization token.

989 A disadvantage is that there is overhead involved in enforcing two layers of authorization—one
990 layer based on policies specified for SVC-SVC requests and a second layer based on EU+SVC-
991 SVC requests. Access control processing logic based on the second layer involves multiple calls
992 by service proxy, such as (a) a call to the authorization engine service to obtain the access
993 decision after obtaining the user attributes (including user credentials) and application object
994 attributes from the orchestration system, (b) obtaining the authorization token from the service
995 mesh control plane based on the access decision, and (c) including the authorization token along
996 with service request.

997 **5.3 Advantages of ABAC Authorization Framework for Service Mesh**

998 We provide here the justification for the various building blocks of the architecture for our
999 authorization framework – the service mesh, the NGAC-based ABAC model etc. We also
1000 highlight the scalability and flexibility of certain components such as proxy APIs, NGAC
1001 authorization engine etc.

- 1002 a. A service mesh is the right architecture for the enforcement of authorization policies since
1003 the components involved are moved out of the application and executed in a space where
1004 they can form a security kernel that can be vetted.
- 1005 b. Both types of authorization requests (i.e., SVC-SVC and EU-SVC-SVC) can be handled
1006 by a runtime infrastructure that involves the coupling of orchestration platform control
1007 plane, service mesh control plane, and mesh data plane to the access control engine.
- 1008 c. The extensible API of the proxy can be used to integrate any authorization engine using
1009 the appropriate type of access control model. ABAC has been found to be one of the most
1010 flexible, scalable access control models because of its ability to incorporate any number
1011 and type of attributes associated with the subject, object, and environment.
- 1012 d. Performance requirements for the authorization engine are met due to the linear time
1013 processing speed of the graph-based, NGAC-based ABAC model.
- 1014 e. The flexibility outlined in (c) can be leveraged to incorporate models for both application
1015 and data protection. Enabling data protection models such as NDAC can be part of the
1016 authorization server.

1017 **5.4 Enforcement Alternatives in Proxies**

1018 Authorization can be enforced through a native structure (e.g., authorization policy) supported in
1019 the particular version of the service mesh or using calls to an external authorization server. The
1020 external authorization server can use any access control model and any representation of policy
1021 expressions (logical rules or acyclic graph representations), but the mediation of a request coming
1022 into the proxy can be performed in the following ways:

- 1023 a. Each request is passed on to the external authorization server through the external
1024 authorization filter in the proxy, and the response from the authorization server is used for
1025 request mediation in the form of ALLOW or DENY.
- 1026 b. Prestored ACLs can be used in the proxy itself, generated by calls to the authorization
1027 server. If the authorization server uses an enterprise-wide access control model, an
1028 administrative API may be needed that will perform the function of mapping the
1029 enterprise resources to resources, users, and groups pertaining to the service served by its
1030 proxy to generate ACLs that are customized for the service.

1031

6 Summary and Conclusions

1033 Deployment guidance has been provided for an ABAC-based authorization framework for
1034 securing microservices-based applications using a service mesh. Background information in
1035 terms of authentication and authorization policies natively supported in proxies of the service
1036 mesh are discussed. For supporting any authentication and authorization framework in the mesh,
1037 the pre-requisites in the form of hosting platform configuration data, the service mesh
1038 configuration and some higher-level security configuration parameters for orchestration of
1039 component microservices (when implemented as containers) are outlined.

1040 The description of the ABAC deployment in the service mesh includes the requirements for
1041 security assurance, supporting infrastructure, advantages of ABAC authorization framework and
1042 enforcement alternatives in proxies.

1043

1044 **References**

- 1045 [1] Rose S, Borchert O, Mitchell S, Connelly S (2020) Zero Trust Architecture. (National
1046 Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP)
1047 800-207. <https://doi.org/10.6028/NIST.SP.800-207>
- 1048 [2] Red Hat (2021) *What is CI/CD?* Available at
1049 [https://www.redhat.com/en/topics/devops/what-is-ci-](https://www.redhat.com/en/topics/devops/what-is-ci-cd#:~:text=CI%2FCD%20is%20a%20method,continuous%20delivery%2C%20and%20c,ontinuous%20deployment)
1050 [cd#:~:text=CI%2FCD%20is%20a%20method,continuous%20delivery%2C%20and%20c](https://www.redhat.com/en/topics/devops/what-is-ci-cd#:~:text=CI%2FCD%20is%20a%20method,continuous%20delivery%2C%20and%20c,ontinuous%20deployment)
1051 [ontinuous%20deployment](https://www.redhat.com/en/topics/devops/what-is-ci-cd#:~:text=CI%2FCD%20is%20a%20method,continuous%20delivery%2C%20and%20c,ontinuous%20deployment)
- 1052 [3] Red Hat (2021) *What is DevSecOps?* Available at
1053 <https://www.redhat.com/en/topics/devops/what-is-devsecops>
- 1054 [4] Chandramouli R (2019) Security Strategies for Microservices-based Application Systems.
1055 (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special
1056 Publication (SP) 800-204. <https://doi.org/10.6028/NIST.SP.800-204>
- 1057 [5] Hu VC, Ferraiolo DF, Chandramouli R, Kuhn DR (2018) Attribute-Based Access Control
1058 (Artech House, Boston USA).
- 1059 [6] Chandramouli R, Butcher Z (2020) Building Secure Microservices-based Applications
1060 using Service-Mesh Architecture. (National Institute of Standards and Technology,
1061 Gaithersburg, MD), NIST Special Publication (SP) 800-204A.
1062 <https://doi.org/10.6028/NIST.SP.800-204A>
- 1063 [7] McEvoy E (2019) *Cordanetes: Combining Corda and Kubernetes* (Medium.com).
1064 Available at <https://medium.com/corda/combining-corda-and-kubernetes-4e2ba54494c7>
- 1065 [8] Ramakani A (2020) *Kong API Gateway – From Zero to Production* (Medium.com).
1066 Available at <https://medium.com/swlh/kong-api-gateway-zero-to-production-5b8431495ee>
- 1067
- 1068 [9] Agarwal G (2020) *How to Manage Microservices on Kubernetes With Istio*
1069 (Medium.com). Available at [https://medium.com/better-programming/how-to-manage-](https://medium.com/better-programming/how-to-manage-microservices-on-kubernetes-with-istio-c25e97a60a59)
1070 [microservices-on-kubernetes-with-istio-c25e97a60a59](https://medium.com/better-programming/how-to-manage-microservices-on-kubernetes-with-istio-c25e97a60a59)
- 1071
- 1072 [10] Agarwal G (2020) *How Istio Works Behind the Scenes on Kubernetes* (Medium.com).
1073 Available at [https://medium.com/better-programming/how-istio-works-behind-the-scenes-](https://medium.com/better-programming/how-istio-works-behind-the-scenes-on-kubernetes-aeb8003f2cb5)
1074 [on-kubernetes-aeb8003f2cb5](https://medium.com/better-programming/how-istio-works-behind-the-scenes-on-kubernetes-aeb8003f2cb5)
- 1075
- 1076 [11] InterNational Committee for Information Technology Standards (2020) *INCITS 565-2020*
1077 *- Information technology - Next Generation Access Control* (INCITS, Washington, DC).
1078 Available at
1079 https://standards.incits.org/apps/group_public/project/details.php?project_id=2328

- 1080 [12] Mell P, Shook J, Harang R, Gavrilas S (2017) Linear Time Algorithms to Restrict Insider
1081 Access using Multi-Policy Access Control Systems. *Journal of Wireless Mobile Networks,*
1082 *Ubiquitous Computing, and Dependable Applications* 8(1):4-25.
1083 <https://doi.org/10.22667/JOWUA.2017.03.31.004>
- 1084 [13] Ferraiolo D, Chandramouli R, Hu V, Kuhn R (2016) A Comparison of Attribute Based
1085 Access Control (ABAC) Standards for Data Service Applications: Extensible Access
1086 Control Markup Language (XACML) and Next Generation Access Control (NGAC)
1087 (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special
1088 Publication (SP) 800-204A. <https://doi.org/10.6028/NIST.SP.800-178>
- 1089 [14] Agarwal G (2020) *Enable Mutual TLS Authentication between your Kubernetes*
1090 *Workloads Using Istio* (Medium.com). Available at [https://medium.com/better-](https://medium.com/better-programming/enable-mutual-tls-authentication-between-your-kubernetes-workloads-using-istio-65338c8adf82)
1091 [programming/enable-mutual-tls-authentication-between-your-kubernetes-workloads-](https://medium.com/better-programming/enable-mutual-tls-authentication-between-your-kubernetes-workloads-using-istio-65338c8adf82)
1092 [using-istio-65338c8adf82](https://medium.com/better-programming/enable-mutual-tls-authentication-between-your-kubernetes-workloads-using-istio-65338c8adf82)
- 1093 [15] Envoy (2021) *External Authorization*. Available at
1094 [https://www.envoyproxy.io/docs/envoy/v1.17.0/intro/arch_overview/security/ext_authz_fi](https://www.envoyproxy.io/docs/envoy/v1.17.0/intro/arch_overview/security/ext_authz_filter#arch-overview-ext-authz)
1095 [lter#arch-overview-ext-authz](https://www.envoyproxy.io/docs/envoy/v1.17.0/intro/arch_overview/security/ext_authz_filter#arch-overview-ext-authz)
- 1096 [16] Envoy (2021) *JWT Authentication*. Available at
1097 [https://www.envoyproxy.io/docs/envoy/v1.17.0/configuration/http/http_filters/jwt_authn](https://www.envoyproxy.io/docs/envoy/v1.17.0/configuration/http/http_filters/jwt_authn_filter#config-http-filters-jwt-authn)
1098 [filter#config-http-filters-jwt-authn](https://www.envoyproxy.io/docs/envoy/v1.17.0/configuration/http/http_filters/jwt_authn_filter#config-http-filters-jwt-authn)
- 1099
- 1100
- 1101
- 1102
- 1103