



NIST Special Publication 800
NIST SP 800-233 ipd

Service Mesh Proxy Models for Cloud-Native Applications

Initial Public Draft

Ramaswamy Chandramouli
Zack Butcher
James Callaghan

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-233.ipd>

NIST Special Publication 800
NIST SP 800-233 ipd

Service Mesh Proxy Models for Cloud-Native Applications

Initial Public Draft

Ramaswamy Chandramouli
*Computer Security Division
Information Technology Laboratory*

Zack Butcher
Tetrade, Inc.

James Callaghan
control-plane.io, Inc.

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-233.ipd>

July 2024



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

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

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

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

Authority

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

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

NIST Technical Series Policies

[Copyright, Use, and Licensing Statements](#)

[NIST Technical Series Publication Identifier Syntax](#)

Publication History

Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be added to final publication.]

How to Cite this NIST Technical Series Publication:

Chandramouli R, Butcher Z, Callaghan J (2024) Service Mesh Proxy Models for Cloud-Native Applications. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-233 ipd. <https://doi.org/10.6028/NIST.SP.800-233.ipd>

Author ORCID iDs

Ramaswamy Chandramouli: 0000-0002-7387-5858

NIST SP 800-233 ipd (Initial Public Draft)
July 2024

Service Mesh Proxy Models for
Cloud-Native Applications

Public Comment Period

July 19, 2024 – September 3, 2024

Submit Comments

sp800-233-comments@nist.gov

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

Additional Information

Additional information about this publication is available at <https://csrc.nist.gov/pubs/sp/800/233/ipd>, including related content, potential updates, and document history.

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

1 **Abstract**

2 The service mesh has become the de-facto application services infrastructure for cloud-native
3 applications. It enables the various runtime functions (network connectivity, access control etc.)
4 of an application through proxies which thus form the data plane of the service mesh.
5 Depending upon the distribution of the network layer functions (L4 & L7) and the granularity of
6 association of the proxies to individual services/computing nodes, different proxy models or
7 data plane architectures have emerged. The purpose of this document is to develop a threat
8 profile for each of the data plane architectures through a detailed threat analysis in order to
9 make recommendations for their applicability (usage) for cloud-native applications with
10 different security risk profiles.

11 **Keywords**

12 proxy model; data plane architecture; service mesh; threat profile; cloud-native application.

13 **Reports on Computer Systems Technology**

14 The Information Technology Laboratory (ITL) at the National Institute of Standards and
15 Technology (NIST) promotes the U.S. economy and public welfare by providing technical
16 leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test
17 methods, reference data, proof of concept implementations, and technical analyses to advance
18 the development and productive use of information technology. ITL's responsibilities include
19 the development of management, administrative, technical, and physical standards and
20 guidelines for the cost-effective security and privacy of other than national security-related
21 information in federal information systems. The Special Publication 800-series reports on ITL's
22 research, guidelines, and outreach efforts in information system security, and its collaborative
23 activities with industry, government, and academic organizations.

24 **Call for Patent Claims**

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

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

- 34 a) assurance in the form of a general disclaimer to the effect that such party does not hold
35 and does not currently intend holding any essential patent claim(s); or
- 36 b) assurance that a license to such essential patent claim(s) will be made available to
37 applicants desiring to utilize the license for the purpose of complying with the guidance
38 or requirements in this ITL draft publication either:
 - 39 i. under reasonable terms and conditions that are demonstrably free of any unfair
40 discrimination; or
 - 41 ii. without compensation and under reasonable terms and conditions that are
42 demonstrably free of any unfair discrimination.

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

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

50 Such statements should be addressed to: sp800-233-comments@nist.gov

51

52	Table of Contents	
53	Executive Summary	1
54	1. Introduction	2
55	1.1. L4 and L7 Functions of Proxies.....	2
56	1.2. Objective & Target Audience	3
57	1.3. Relationship to Other NIST Documents	4
58	1.4. Document Structure.....	4
59	2. Typical Service Mesh Data Plane Capabilities and Associated Proxy Functions	5
60	3. Proxy Models (Data plane Architectures) in Service Mesh Implementations	7
61	3.1. L4 and L7 Proxy per Service Instance – Sidecar Model (DPA-1).....	8
62	3.2. Shared L4 – L7 per Service Model (DPA-2).....	9
63	3.3. Shared L4 and L7 Model (DPA-3)	10
64	3.4. L4 and L7 Part of the Application Model (DPA-4)	11
65	4. Data Plane Architectures Threat Scenarios and Analysis Methodology	12
66	4.1. Threat analysis Methodology.....	13
67	5. Detailed Threat Analysis for Data Plane Architectures	14
68	5.1. Threat Analysis for L4 and L7 Proxy per Service Instance – Sidecar Model (DPA-1)	14
69	5.1.1. Compromised L4 Proxy (TR-1).....	14
70	5.1.2. Compromised Application Container (TR-2)	15
71	5.1.3. Compromise of Business Data (TR-3).....	15
72	5.1.4. Compromised L7 Proxy (TR-4).....	16
73	5.1.5. Compromise of Shared L7 Proxy (TR-5)	16
74	5.1.6. Outdated Client Libraries in Applications (TR-6).....	17
75	5.1.7. Denial of Service (TR-7).....	17
76	5.1.8. Resource Consumption (TR-8)	17
77	5.1.9. Privileged L4 Proxy (TR-9)	18
78	5.1.10. Data Plane (Service Mesh) Bypassed (TR-10)	18
79	5.2. Threat Analysis for Shared L4 – L7 per Service Model (DPA-2)	19
80	5.2.1. Compromised L4 Proxy (TR-1).....	19
81	5.2.2. Compromised Application Container (TR-2)	19
82	5.2.3. Compromise of Business Data (TR-3).....	19
83	5.2.4. Compromised L7 Proxy (TR-4).....	19
84	5.2.5. Compromise of Shared L7 Proxy (TR-5)	20
85	5.2.6. Outdated Client Libraries in Applications (TR-6).....	20

86	5.2.7. Denial of Service (TR-7).....	20
87	5.2.8. Resource Consumption (TR-8)	21
88	5.2.9. Privileged L4 Proxy (TR-9)	21
89	5.2.10. Data Plane (Service Mesh) Bypassed (TR-10)	22
90	5.3. Threat Analysis for Shared L4 and L7 Model (DPA-3).....	22
91	5.3.1. Compromised L4 Proxy (TR-1).....	22
92	5.3.2. Compromised Application Container (TR-2)	22
93	5.3.3. Compromise of Business Data (TR-3).....	23
94	5.3.4. Compromised L7 Proxy (TR-4).....	23
95	5.3.5. Compromise of Shared L7 Proxy (TR-5)	23
96	5.3.6. Outdated Client Libraries in Applications (TR-6).....	24
97	5.3.7. Denial of Service (TR-7).....	24
98	5.3.8. Resource Consumption (TR-8)	24
99	5.3.9. Privileged L4 Proxy (TR-9)	25
100	5.3.10. Data Plane (Service Mesh) Bypassed (TR-10)	25
101	5.4. Threat Analysis for L4 and L7 within Application Model (gRPC proxyless Model (DPA-4))	25
102	5.4.1. Compromised L4 Proxy (TR-1).....	25
103	5.4.2. Compromised Application Container (TR-2)	26
104	5.4.3. Compromise of Business Data (TR-3).....	26
105	5.4.4. Compromised L7 Proxy (TR-4).....	26
106	5.4.5. Compromise of Shared L7 Proxy (TR-5)	27
107	5.4.6. Outdated Client Libraries in Applications (TR-6).....	27
108	5.4.7. Denial of Service (TR-7).....	27
109	5.4.8. Resource Consumption (TR-8)	28
110	5.4.9. Privileged L4 Proxy (TR-9)	28
111	5.4.10. Data Plane (Service Mesh) Bypassed (TR-10)	28
112	6. Recommendations Based on Application Security Risk Profile	29
113	6.1. Cloud-Native Applications with Low Risk Profile	29
114	6.2. Cloud-Native Applications with Medium Risk Profile	30
115	6.3. Cloud-Native Applications with High Risk Profile	30
116	7. Summary and Conclusions	32
117	References.....	33
118		
119		

120 **Acknowledgments**

121 The authors would like to express their thanks to Francesco Beltramini of control-plane.io for
122 participating in our discussions and providing his valuable perspectives.

123 **Executive Summary**

124 Run-time services for Cloud-native applications, consisting of multiple loosely coupled
125 components called microservices, are sometimes provided through a centralized infrastructure
126 called a service mesh. These services include secure communication, service discovery,
127 resiliency, and authorization of application communication. These services are mainly provided
128 through *Proxies* that form the *data plane* of the service mesh, the layer that handles application
129 traffic at runtime and enforces policy.

130 The functions that the proxies provide can be broadly categorized into two groups, based on
131 the OSI model's network layer to which those functions pertain to. These groups are: Layer 4
132 ("L4") and Layer 7 ("L7"). In majority of deployments of service mesh in production
133 environments today, all proxy functions (providing services in both L4 and L7 layers) are packed
134 into a single proxy that is assigned to a single microservice. This service mesh proxy model is
135 called a *sidecar proxy model* since the proxy is not only associated with a single service but is
136 implemented to execute in the same network space as the service.

137 However, performance and resource considerations have led to the exploration of alternate
138 proxy models which involve not only splitting up of L4 and L7 functions into different proxies
139 (instead of a single proxy) but also the association or assignments of these proxies to either a
140 single service or a group of services, thus enabling the proxies to be implemented at different
141 locations - at the granularity of a node rather than at the level of services. Though different
142 models are theoretically possible, we consider only those service mesh proxy models in the
143 data plane implementation of commonly used service mesh offerings, at different stages.

144 We then consider a set of potential/likely threats to various proxy functions. Each of the threats
145 may result in different types of exploits in different proxy models. This variation is due to
146 several factors such as: attack surface (communication patterns to which a particular proxy is
147 exposed), number of clients (services) served and OSI layer functions they provide (e.g., L7
148 functions are more complicated and likely to have more vulnerabilities than L4 functions). The
149 two main contributions of this document are as follows:

- 150 1. The nature of exploits possible for each threat in each of the proxy models are
151 characterized by assigning scores to the impact and likelihood of each of these threats in
152 each of the proxy models or architectural patterns resulting in a threat profile
153 associated with each architectural pattern or proxy model of service mesh.
- 154 2. Each threat profile inherently has a built-in set of security tradeoffs at an architectural
155 level. The implications of these tradeoffs in meeting the requirements associated with
156 security risk profile of different cloud-native applications are analyzed to make a broad
157 set of recommendations towards specific architectural patterns that are appropriate for
158 applications with different security risk profiles.

159

160 **1. Introduction**

161 The service mesh, an application service infrastructure is now an integral part of the overall
162 application infrastructure of cloud-native applications, typically consisting of multiple loosely
163 coupled services or microservices. The infrastructure services or functions provided by a service
164 mesh during application runtime are provided by entities called proxies which constitute the
165 data plane of the service mesh. In addition, the service mesh consists of another architectural
166 component called the control plane which supports the functions of the data plane through
167 interfaces to define configurations, inject software programs and provide security artifacts such
168 as certificates.

169 Based on performance and security assurance data gained over the deployment of service
170 mesh for the last several years, various configurations for proxies are being developed and
171 tested. These configurations are based on the OSI layer functions they provide (see section 1.1
172 L4 and L7 functions of a proxy) and the granularity of association between a proxy and services
173 and go by the name of proxy (implementation) models. Since proxies are the predominant
174 entities of the data plane of a service mesh, these various proxy models are also called data
175 plane architectures.

176 **1.1. L4 and L7 Functions of Proxies**

177 To understand proxy models, there are two aspects we should look at. They are:

178 Proxy Functions: The functions that a service mesh's proxies provide can be broadly categorized
179 into two groups, based on the OSI model's layer [1] to which those functions pertain to. These
180 groups are: Layer 4 ("L4") and Layer 7 ("L7"). The associated proxies are called L4 proxies and L7
181 proxies respectively.

182 Granularity of Association: A proxy can be associated with a single microservice instance, an
183 entire service or it can be deployed to provide functions for a group of services. Depending
184 upon the nature of this association, a proxy may execute within the same network space as the
185 service, or it can execute at the same node where the group of services to which caters to run
186 or in an independent node (dedicated to just proxies where no application services run).

187 The study of proxy functions (the first topic above) in turn requires us to go into fundamentals a
188 little bit and look at what OSI's L4 and L7 layers are, from the network stack point of view and
189 the specific network services provided by those layers.

190 The OSI model [1] is a useful abstraction for thinking about the functions required to serve an
191 application over the network. It describes seven "layers", from the physical wires connecting
192 two machines (Layer 1 – L1 – the physical layer) all the way up to the application itself (Layer 7
193 – L7 – the application layer). When facilitating the communication of cloud-native applications
194 (e.g., two microservices making HTTP/REST calls to each other), we care primarily about layers
195 3, 4, and 7; A brief overview of the functionality of these layers are:

- 196 • Layer 3 (“L3”), the network layer, facilitates baseline connectivity between two
197 workloads or service instances. In nearly all cases, the Internet Protocol (IP) is used as
198 the layer 3 implementation.
- 199 • Layer 4 (“L4”), the transport layer, facilitates the reliable transmission of data between
200 workloads on the network. It also includes capabilities like encryption. TCP and UDP are
201 commonly used L4 implementations, where TLS (*transport layer security* – named after
202 the OSI model) provides encryption.
- 203 • Layer 7 (“L7”), the application layer, which is where protocols like HTTP live – in user
204 applications themselves (e.g., HTTP web servers, SSH servers).

205 With respect to the layers above, in cloud native environments, a service mesh’s proxies are:

- 206 • Are agnostic to L3, so long as microservice instances can communicate at L3 and the
207 proxy can communicate with the mesh’s control plane.
- 208 • At Layer 4 (L4): connection establishment, management, and resiliency (e.g.,
209 connection-level retries); TLS (encryption in transit); application identity, authentication,
210 and authorization; access policy based on network 5-tuple (source IP address and port,
211 destination IP address and port, and transport protocol).
- 212 • At Layer 7 (L7): service discovery, request-level resiliency (e.g., retries, circuit breakers,
213 outlier detection); and application observability.

214 What we have seen so far is one aspect of proxy model or data plane architecture – i.e., proxy
215 functions. The other aspect as we alluded to earlier is the proxies’ granularity of association to
216 services.

217 **1.2. Objective & Target Audience**

218 This document will give a brief overview of the 4 data plane architectures (proxy models) being
219 pursued by a range of service mesh implementations today. It will then develop threat profiles
220 for different proxy models through a detailed threat analysis involving ten types of common
221 threats. These threat profiles will be used to make a set of recommendations regarding their
222 applicability (usage) for cloud-native applications with different security risk profiles. The target
223 audience for these recommendations is:

- 224 • Infrastructure owners and platform/infrastructure engineers (and their team heads)
225 building to build and deploy a secure run-time environment for applications by choosing
226 the right architecture for their environment given the risk factors of the applications
227 they’ll be running (and the resulting security risk profile).
- 228 • Personnel in charge of infrastructure operations to familiarize them with the details of
229 the various building blocks of the proxy models or data plane architectures (and their
230 associated functions and interactions) to troubleshoot in the event of performance
231 (availability) and security issues.

232 **1.3. Relationship to Other NIST Documents**

233 This document can be used as an adjunct to NIST Special Publication (SP) 800-204 series of
234 publications [2,3,4,5], which offer guidance on providing security assurance for cloud-native
235 applications integrated with a service mesh from the following perspectives: strategy,
236 configuration, and development/deployment paradigm. However, this document focuses on
237 the various configurations of the application service infrastructure elements (i.e., proxies) and
238 the resulting architectures (i.e., data plane architecture of the service mesh) that have different
239 security implications for the application that is hosted under each of these configurations.

240 **1.4. Document Structure**

241 This document is organized as follows:

- 242 • Section 2 provides a list of typical capabilities of the data plane of the service mesh
243 under three headings (security, observability and traffic management) and the
244 corresponding L4 and L7 proxy functions implemented under those capabilities.
- 245 • Section 3 provides a brief overview of the four architectural patterns called the proxy
246 models or data plane architectures.
- 247 • Section 4 discusses proxy model threat scenarios and gives a roadmap of the threat
248 analysis methodology adopted in this document for evaluating the threat profile score
249 for the four data plane architectures.
- 250 • Section 5 provides a detailed threat analysis for the four data plane architectures by
251 assigning scores to impact and likelihood factors associated with each threat and using
252 them to arrive at the overall threat score.
- 253 • Section 6 provides the recommendations for the applicability (usage) of each of the 4
254 data plane architectures for cloud-native applications of different security risk profiles
255 based on their security requirements.
- 256 • Section 7 provides the summary and conclusions.

257

258 **2. Typical Service Mesh Data Plane Capabilities and Associated Proxy Functions**

259 Since examining the security tradeoffs of the proxy models (data plane architectures) is part of
 260 our methodology in this document, we have to look at implementations of the various
 261 capabilities (under the umbrella of Security, Observability and Network Traffic Management)
 262 that result as L4 and L7 functions in proxies. To arrive at the totality of proxy functions, we need
 263 to analyze for each capability, which category (L4 vs L7) it falls in to, and the granularity of the
 264 function that it provides at L4 and L7 levels.

265 **Table 1 - Security Capabilities**

Capability	L4 Function(s)	L7 Function(s)
Service-to-service authentication	SPIFFE , via mTLS certs. Control plane issues a short-lived X.509 encoding the pod's service account identity.	N/A—service identity in a service mesh is usually based on TLS only.
Service-to-service authorization	Network-based authorization, plus identity-based policy, e.g.: A can accept inbound calls from only "10.2.0.0/16"; A can call B.	Full policy, e.g.: A can GET /foo on B only with valid end-user credentials containing the READ scope.
End-user authentication	N/A—we can't apply per-user settings.	Local authentication of JWTs, support for remote authentication via OAuth and OIDC flows.
End-user authorization	N/A—see above.	Service-to-service policies can be extended to require end user credentials with specific scopes, issuers, principal, audiences, etc —but it cannot be used for full user-to-resource access control. Full user-to-resource access should be implemented using external authorization.
Mesh proxy's External Authorization API (ext_authz)	Cannot perform any per-request policy; ext_authz API is only configurable for L7 traffic.	Enforce per-request policy with decisions from an external service, e.g., OPA.

266

267

Table 2 - Observability Capabilities

Capability	L4 Function(s)	L7 Function(s)
Logging	Basic network information: network 5-tuple, bytes sent/received, etc.	Full request metadata logging , in addition to basic network information.
Tracing	Not today; possible eventually, with HBONE.	Mesh proxy participates in distributed tracing.
Metrics	TCP only (bytes sent/received, number of packets, etc).	L7 RED metrics: rate of requests, rate of errors, request duration (latency).

268

269

Table 3 – Traffic Management Capabilities

Capability	L4 Function(s)	L7 Function(s)
Load balancing	Connection level only. See TCP traffic shifting task.	Per request, enabling e.g. canary deployments, gRPC traffic, etc. See HTTP traffic shifting task.
Circuit breaking	TCP only.	HTTP settings in addition to TCP.
Outlier detection	On connection establishment/failure.	On request success/failure.
Rate limiting	Rate limit on L4 connection data only, on connection establishment, with global and local rate limiting options.	Rate limit on L7 request metadata, per request.
Timeouts	Connection establishment only (connection keep-alive is configured via circuit breaking settings).	Per request.
Retries	Retry connection establishment	Retry per request failure.
Fault Injection	N/A—fault injection cannot be configured on TCP connections.	Full application and connection level faults (timeouts, delays, specific response codes).
Traffic Mirroring	N/A—HTTP only	Percentage-based mirroring of requests to multiple backends.

270

271 It’s important to note that L7 functions carried out by proxies are much more complex than L4
 272 functions as the latter are carried out in lower layers of OSI stack involving protocols such as IP
 273 and TCP. For example, parsing a TCP stream for L4 functionality requires simply decoding a
 274 fixed set of bytes as integers (the packet header), while handling HTTP requests for L7
 275 functionality requires decoding HTTP headers including complex string parsing and compression
 276 with variable amounts of data. Further, that data dealt with in an L7 function is user-supplied
 277 (i.e., can be controlled by an attacker), while the TCP data at L4 is typically system-supplied as
 278 part of routing a request to your infrastructure – there’s less room to embed malicious data
 279 without breaking the system itself. As one case study, the proxy Envoy is used as the data plane
 280 by several service mesh implementations: historically, majority of Envoy vulnerabilities have
 281 been in L7-function-related code compared to L4-function-related code.

282

283 3. Proxy Models (Data plane Architectures) in Service Mesh Implementations

284 As we had briefly seen before, different data plane architectures or proxy models in service
285 mesh are a consequence of the following parameters.

- 286 • Delineation of L4 and L7 functions
- 287 • Nature of association of a proxy to service instances (1:1 or 1:N)

288 In this section, an overview of the building blocks of different data plane architectures is
289 undertaken to facilitate the threat analysis that follows in section 5. Before we list the different
290 data plane architectures (also called different iterations of service mesh implementations) that
291 have been commonly implemented, it is in order to look at as to why these different
292 architectures were necessitated in the first place. These iterations were driven by the adoption
293 of mesh across a variety of use cases, necessitating different tradeoffs in terms of performance,
294 reliability, and security across a variety of organizations with different application risk profiles.
295 It must be mentioned, however, that in spite of these different operating scenarios, the first
296 model listed here, i.e., the sidecar model, has been the primary predominant method of
297 delivering the capabilities of service mesh for several years.

298 The various alternate data plane architectures, including the one with widespread deployment
299 at present, are:

- 300 • “L4 and L7 per Service Instance” - Side-car Model (DPA-1)
- 301 • “Shared L4 – L7 per Service” (DPA-2) - A shared L4 proxy per node, i.e., shared among all
302 applications that execute on the same physical host, with L7 proxies dedicated per
303 service account or namespace.
- 304 • “Shared L4 and L7” (DPA-3) - A shared L4 *and* L7 proxy per node, i.e., shared among all
305 applications on the same physical host.
- 306 • “L4 and L7 within Application (gRPC proxy-less model)” (DPA-4) – Both L4 and L7
307 functions instead of being implemented in stand-alone proxies are part of the
308 application server itself, e.g., frameworks such as gRPC, Java Spring etc.

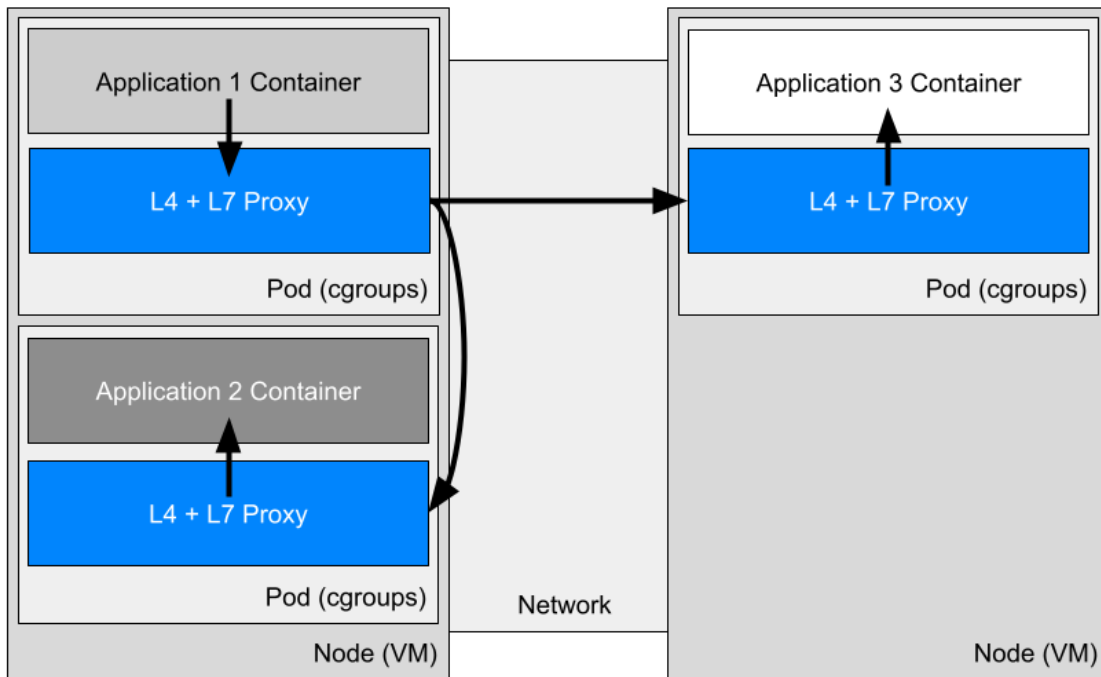
309 It must be mentioned that though the last architectural pattern does not have distinct entities
310 such as proxies, all the service mesh capabilities delivered by proxies are enabled by the
311 frameworks mentioned above.

312

313 **3.1. L4 and L7 Proxy per Service Instance – Sidecar Model (DPA-1)**

314 The first and most common service mesh data plane architecture today dedicates a proxy that
315 has the capability to implements both L4 and L7 functions for each application (service)
316 instance. This is also called a “sidecar model” since the proxy sits beside every instance of every
317 service. The security model here is simple: the proxy holds one identity (for the service it’s
318 deployed beside) and resides in the same trust domain as the application (in Kubernetes, it
319 exists in the same pod; on a VM, it’s deployed in the same VM as the service itself). The service
320 and the proxy communicate with each other through the “local host interface” instead of
321 through a network socket. However, the proxy itself presents a larger attack surface than the
322 service because it implements the complex L7 functions. An example of a data plane
323 architecture is the one that is implemented in the Istio service mesh with an envoy proxy
324 deployed per pod that performs both L4 and L7 functions.

325 A schematic diagram of this architecture is shown in Figure 1.



326

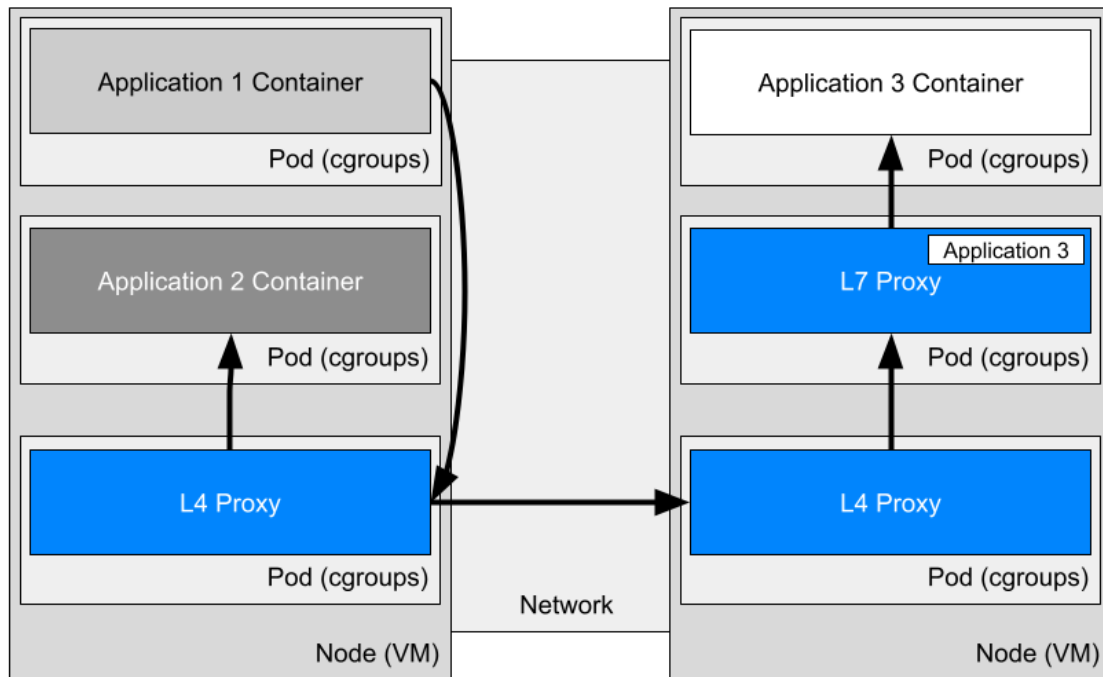
327

Figure 1 – L4 and L7 Proxy per Service Instance (Side Car Model) (DPA-1)

328

329 **3.2. Shared L4 – L7 per Service Model (DPA-2)**

330 In this architecture, there is a shared L4 proxy per node, i.e., shared among all service instances
331 that execute on the same physical host, with L7 proxies dedicated per service account. This is
332 also called “ambient mode”. A variation in this architecture is to dedicate a L7 proxy for an
333 entire namespace. This is not desirable from a security viewpoint based on the same reasons
334 we recommend against shared service account for entire namespace [2] and hence not
335 considered for threat analysis in this document. An example of implementation of this data
336 plane architecture is the Istio Ambient where the per node L4 proxy is called Ztunnel proxy and
337 per service account L7 proxy is called Waypoint proxy [6,7,10,11,13].



338

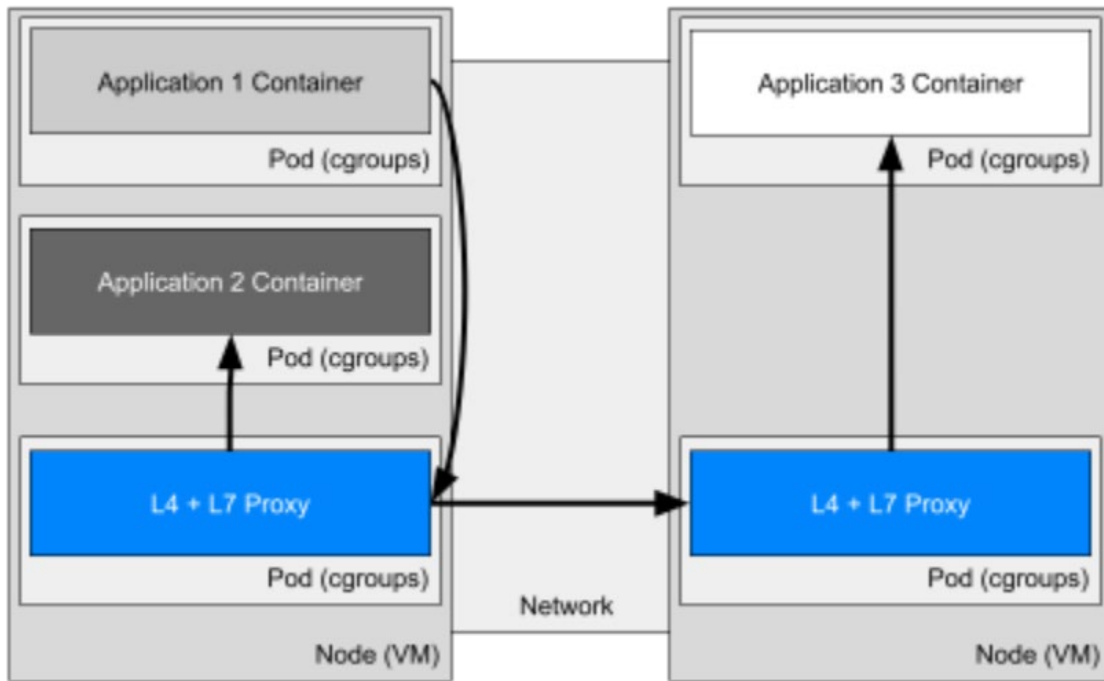
339

Figure 2 – Shared L4 - L7 per Service Model (DPA-2)

340

341 **3.3. Shared L4 and L7 Model (DPA-3)**

342 In this architecture, the L4 and L7 functions are implemented on a per node basis. There is a
343 shared L7 proxy per node, i.e., shared among all service instances that execute on the same
344 physical host and provides L7 functions for all services in that node. However, the L4 functions
345 such as traffic routing can be performed not by proxies but by in-kernel programs (e.g., eBPF
346 programs) or the mesh proxy. An example of this data plane architecture is the Cilium service
347 mesh which deploys the Envoy proxy as L7 proxy based on its CiliumEnvoyConfig specification
348 [8,9,12].



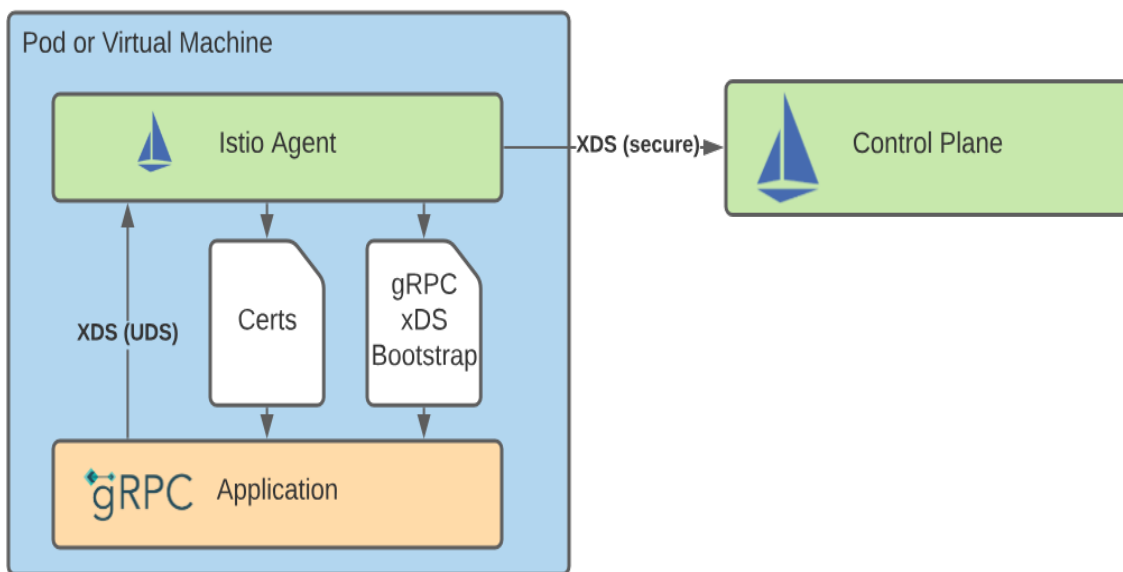
349
350
351

Figure 3 – Shared L4 - L7 Model (DPA-3)

352 **3.4. L4 and L7 Part of the Application Model (DPA-4)**

353 This is a data plane architecture that does not have any proxies. The service mesh control plane
354 dynamically configures proxies using a set of discovery APIs collectively known as xDS APIs. The
355 gRPC client library for applications provides extensive support for the xDS APIs. Leveraging this
356 feature, the service mesh control plane can program L4 and L7 functions into this library in the
357 service container. These gRPC libraries can then provide the L4 and L7 functionality (in general
358 all policy enforcements) to the workloads or service instances to which they are integrated
359 with, thus replicating the exact services which the L4 and L7 proxies provide to those workloads
360 [14].

361 The architecture diagram of the gRPC proxyless data plane architecture for a service mesh is
362 given below:



363

364

Figure 4 – L4 and L7 Part of the Application Model (gRPC proxyless Model) (DPA-4)

365

366 **4. Data Plane Architectures Threat Scenarios and Analysis Methodology**

367 We are studying service mesh deployed in a Kubernetes cluster. Assumption that no human can
368 directly access the cluster, achieved via k8s RBAC; only interaction with the cluster is via CI/CD
369 system controlled declarative configuration in a version-controlled repository with a multi-step
370 approval process to change that configuration (including the approval of each change by at
371 minimum one other human).

372 We start by identifying a variety of access by external threat actors, internal threat actors, and
373 malicious co-tenants.

374 External threat actors include:

- 375 • Compromised workload (application) container, e.g., via a supply chain attack
- 376 • Compromised node L4 proxy or CNI
- 377 • Compromised node L7 proxy
- 378 • Compromised node with limited privileged access, e.g., a container breakout
- 379 • Root compromise of node, e.g., a container breakout chained with exploitation of a
380 privilege escalation vulnerability.
- 381 • Network access to the Kubernetes API server

382 Internal threat actors include:

- 383 • Cluster admins, who have wide-ranging rights to view the cluster and approve changes
384 to the version-controlled repository; they may even have direct access to the
385 Kubernetes cluster, e.g., via a break-glass debugging account – such super-accounts
386 should generate detailed audit records of their usage.
- 387 • Application developers, who can build images and approve configuration that goes into
388 the cluster.
- 389 • Infrastructure engineers, who have permission to deploy and configure the mesh –
390 again, gated by the version-controlled repository’s approval process.
- 391 • Compromised network infrastructure between nodes, e.g., un-encrypted cross-data
392 center communication

393 Finally, malicious co-tenants – in general k8s is not a hard multi-tenant system and we
394 recommend isolating tenants from each other with stronger boundaries. In this context, a
395 malicious co-tenant would fall into one of the internal threat actor personas above.

396 In the context of these threat actors, we introduce the following threats as a minimum set to
397 consider in your environment as they relate to the service mesh:

- 398 1. Compromised L4 proxy
- 399 2. Compromise of the Application Container
- 400 3. Compromise of Business Data

- 401 4. Compromised L7 proxy
- 402 5. Compromise of shared L7 Proxy
- 403 6. Outdated Client Libraries in Applications
- 404 7. Denial of Service
- 405 8. Resource Consumption
- 406 9. Privileged L4 Proxy
- 407 10. Bypassing Traffic Interception

408 In the next section three we will evaluate the impact of these threats on the components of the
409 data plane architecture for each of the four that we have taken up for consideration in this
410 document.

411 **4.1. Threat analysis Methodology**

412 We first identify 10 potential threats to the components that make up the four architectural
413 patterns for the proxy model or data plane architecture. For each threat, we describe how the
414 functionality of each component of the architecture is adversely affected by the threat and
415 then rate the impact and likelihood of their occurrence, justifying each rating. We have chosen
416 three values for ratings - *low*, *medium*, *high*. The values assigned to these ratings are relative to
417 other data plane architectures and are not absolute values based on a metric. For example, the
418 assignment of the rating value “High” for the likelihood parameter for a threat does not imply
419 that the threat is highly likely in all situations; it means that this threat is likeliest to be
420 executable against that architecture *relative to the other architectures under discussion*.

421 For each threat and architecture, we evaluate the *impact (I)* of the exploitation of that threat
422 along with the *likelihood (L)* of that threat being exploited. As we already stated, for both
423 parameters we give a rating of low, medium, and high which we translate to numeric scores 1,
424 2, and 3 respectively. By multiplying these together, $I * L$, we can get a indication of how
425 important that threat is and therefore the necessity to mitigate that threat relative to other
426 architectures under discussion. Summing up the values of this indicator for all 10 potential
427 threats, we obtain an indication of the threat profile for that architectural pattern.

428 For those threats whose impact and likelihood are same irrespective of the architecture – in
429 other words, the threats are agnostic to the architecture, we assign a score of 1 for impact and
430 1 for likelihood due the fact that we stated earlier - these scores are relative scores and not
431 absolute scores.

432 **5. Detailed Threat Analysis for Data Plane Architectures**

433 In this section, we analyze the various potential proxy-functions targeted threats (both for L4 &
434 L7 proxies or the libraries implementing the associated functions), the relevant proxy function
435 that is impacted, the degree of impact, the likelihood of the threat occurring for each of the
436 data plane architectures discussed in sections 3.1 to 3.4.

437 Recapping from Section 4, the 10 threats with their identifiers added that are considered for
438 analysis in this section are:

439 Compromised L4 proxy (TR-1)

440 Compromise of the Application Container (TR-2)

441 Compromise of Business Data (TR-3)

442 Compromised L7 proxy (TR-4)

443 Compromise of shared L7 Proxy (TR-5)

444 Outdated Client Libraries in Applications (TR-6)

445 Denial of Service (TR-7)

446 Resource Consumption (TR-8)

447 Privileged L4 Proxy (TR-9)

448 Data plane (Service Mesh) Bypassed (TR-10)

449

450 The organization of this section is as follows:

451 Section 5.1 will analyze each threat for the “L4 and L7 Proxy per Service Instance – Sidecar
452 Model (DPA-1)” and come up with the overall threat score.

453 Section 5.2 will analyze each threat for the “Shared L4 - L7 per Service Model (DPA-2)” and
454 come up with the overall threat score.

455 Section 5.3 will analyze each threat for the “Shared L4 - L7 Model (DPA-3)” and come up with
456 the overall threat score.

457 Section 5.4 will analyze each threat for the “L4 and L7 Part of the Application Model (gRPC
458 proxyless Model) (DPA-4)” and come up with the overall threat score.

459 **5.1. Threat Analysis for L4 and L7 Proxy per Service Instance – Sidecar Model (DPA-1)**

460 Each of the threats to the data plane of the service mesh is denoted using the mnemonic TR-x
461 where TR stands for threat and x for the threat sequence number.

462

463 **5.1.1. Compromised L4 Proxy (TR-1)**

464

465 **Threat Description:** Compromised L4 proxy (or L4 functions in the case of sidecar proxy with
466 combined L4 and L7 functions) leads to leaked identities for every workload (service) running
467 on the node.

468 **Proxy Function Impacted:** Sidecar proxies negotiate mTLS connections (for communicating with
469 any other service) on behalf of only the single workload it is associated with. In order to
470 compromise key material and identity documents (threat targets) for multiple workloads,
471 multiple proxy (sidecar) instances would need to be compromised.

472 **Impact Score=1:** Because of the nature of impact discussed above (i.e., Single workload / single
473 identity being affected), this threat is assigned the impact score of 1.

474 **Likelihood Score=2:** Code relating to L7 functions is present to be exploited, if it can be
475 triggered. In a pure L4 proxying case it *should* not be triggerable, but this relies on correct
476 configuration from users and the service mesh implementation.

477 5.1.2. **Compromised Application Container (TR-2)**

478 **Threat Description:** Compromised application container (e.g., via a supply chain attack – during
479 development phase) leads to takeover of identity associated with that application.

480 **Proxy Function Impacted:** Proxies run in the same network space (same pod in Kubernetes
481 environment) as the application container, meaning that a compromise of the application
482 container (hosting the service instance) can easily lead to a compromise of any key material
483 (full access to key material pertaining to the identity of the service) possessed by the proxy.

484 **Impact Score=2:** Because of the nature of impact discussed above (i.e., Single workload / single
485 identity being affected), this threat is assigned the impact score of 2. Even though only a single
486 identity is compromised, like TR-1, this has a higher impact score as the application itself must
487 be updated. A compromised proxy can be remediated without requiring the application itself to
488 be updated, so there's a higher chance a central team can successfully remediate a compromise
489 without involving application teams.

490 **Likelihood Score=1:** Same regardless of Architecture.

491 5.1.3. **Compromise of Business Data (TR-3)**

492 **Threat Description:** Compromised identity is used to pivot through the infrastructure, in order
493 to compromise the confidentiality, integrity or availability of business data.

494 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1):** Same regardless of
495 architecture -- this is the fundamental risk of identity-based policy and is why we need to
496 practice the principle of least privilege (PoLP). The telemetry provided by the service mesh
497 (regardless of architecture) is invaluable for understanding communicating in your system and
498 creating accurate access policies (thereby implementing PoLP)

499 5.1.4. **Compromised L7 Proxy (TR-4)**

500 **Threat Description:** Vulnerability in L7 processing stack of the service mesh proxy. As L7
501 processing is inherently more complex, there is a higher probability for vulnerabilities to arise in
502 this part of the stack, as supported by historical CVE data.

503 **Proxy Function Impacted:** No separation between L4 and L7 processing. It can be argued that
504 any exploitable vulnerability in a sidecar proxy can lead to the compromise of all identities in
505 the mesh, however as this would involve more individual proxy instances being compromised, it
506 may be more difficult for an attacker to accomplish this feat undetected.

507 **Impact Score=1:** A single workload is impacted (either leaking credentials, or becoming
508 unavailable due to DoS, depending on the type of L7 attack). The same exploit could be used
509 against all sidecars in the mesh with applications opting in to L7 behavior, resulting in
510 compromise of all identities (Impact 3); in practice this requires many more events than any
511 other architecture, increasing our likelihood of detecting and responding to the event in a
512 timely manner.

513 **Likelihood Score=1:** Full L7 capability is available in the proxy, meaning a relatively large attack
514 surface is exposed; in practice for the service mesh use case, however, it tends to be the HTTP
515 processing that is targeted. If the application is using L7 mesh capabilities, they would be
516 vulnerable to exploit.

517 5.1.5. **Compromise of Shared L7 Proxy (TR-5)**

518 **Threat Description:** Co-tenant exploits L7 traffic processing vulnerability in shared proxy, to
519 affect the confidentiality, integrity or availability of traffic to/from another workload running on
520 the same node.

521 **Proxy Function Impacted:** Because the proxy is dedicated per application, impact on availability
522 is limited to the resource constraints imposed by the scheduling system (e.g. Kubernetes).
523 Confidentiality is impacted the same as if another application itself is compromised -- i.e.
524 containers provide some guarantee, micro-VMs provide a stronger degree of isolation, full
525 blown VMs the strongest.

526 **Impact Score=1:** For noisy neighbors – other L7 proxies on the same host that are compromised
527 – Impact limited by underlying scheduling and resource constraint system (e.g. k8s, VM sizing,
528 etc). Identical across all architectures: for a shared ingress gateway, all services exposed on that
529 gateway would be impacted (Impact 2); for a shared egress gateway, all services utilizing the
530 egress gateway are impacted (Impact 3; typically only a single deployment of egress gateways is
531 used).

532 **Likelihood Score=1:** The sidecar itself is not a shared proxy – by its nature it is dedicated to an
533 individual application. In this case TR-5 refers to both noisy neighbors, other proxies on the
534 same node causing a denial of service, as well as shared ingress or egress gateways. Noisy
535 neighbors are mitigated based on the degree of isolation of the host (container vs micro-VM vs

536 VM). Likelihood of exploiting a shared L7 ingress or egress gateway is the same across all
537 architectures.

538 5.1.6. **Outdated Client Libraries in Applications (TR-6)**

539 **Threat Description:** Client libraries are not updated frequently or consistently across the estate
540 of microservices, leading to potential vulnerabilities and weaknesses that can be exploited.

541 **Proxy Function Impacted:** The proxy's Infrastructure code is decoupled from application code.

542 **Impact Score=1:** The mesh infrastructure is separate from the application itself, therefore it's
543 not impacted by application vulnerabilities directly. Instead, a compromised app would use the
544 (functioning) mesh to hijack the application's identity (see threat on compromised app
545 container, compromised identity). Some application vulnerabilities can be mitigated via policies
546 enforced by the mesh, for example: mesh enforced WAF policy can help mitigate an app
547 vulnerability like Log4j while the organization is patching applications.

548 **Likelihood Score=1:** Same regardless of architecture.

549 5.1.7. **Denial of Service (TR-7)**

550 **Threat Description:** Conventional Denial of Service threat.

551 **Proxy Function Impacted:** Because the proxy is per app instance, a DoS needs to be executed
552 per app. Because the proxy shares resources with the app, a DoS on the mesh data plane
553 directly competes for resources with the app instance itself. The overall blast radius of the DoS
554 is as strong as the underlying isolation mechanism protecting workloads (pods) from each other
555 (VMs, micro-VMs, containers, etc.).

556 **Impact Score=1:** Single instance of a single app

557 **Likelihood Score=1:** L4 and L7 code is able to be exploited; however the attack must be
558 executed across each instance of the target (there's not a central resource to target to achieve
559 a DoS, other than a shared ingress gateway which is identical across all architectures under
560 discussion).

561 5.1.8. **Resource Consumption (TR-8)**

562 **Threat Description:** Overall resource consumption by the data plane of the service mesh
563 infrastructure.

564 **Proxy Function Impacted:** Because sidecars are a separate process and are dedicated per app,
565 they have the worst overall resource consumption:

566 configuration that's identical across all apps must be held by the data plane per app, and can't
567 be shared.

568 static overhead of the sidecar data plane implementation itself (e.g. constant RAM usage,
569 constant CPU overhead, and so on) is duplicated per app instance, and can't be amortized over
570 all apps on the node

571 In part this isolation is what allows sidecars to have lower impact and likelihood across many of
572 the other threats identified here.

573 **Impact Score=3**: Highest resource usage of all options, though good configuration can help
574 mitigate the impact (even then, in well-configured environments sidecars will consume the
575 most resource out of all available options).

576 **Likelihood Score=3**: It is challenging to configure sidecars correctly to minimize configuration
577 and reduce overhead. Some specific implementations do better jobs than others due to
578 engineer tradeoffs (e.g. lazily loading configuration the first time an app needs it, vs eagerly
579 pushing all configuration ahead of use) but overall it's easiest to land in a situation with the
580 most resource utilization with a sidecar architecture.

581 5.1.9. Privileged L4 Proxy (TR-9)

582 **Threat Description**: Service mesh implementation requires L4 component (e.g., deployed as a
583 DaemonSet on a Kubernetes cluster) to run with an overprivileged security context (e.g.,
584 Privileged Pod)

585 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1)**: Same regardless of
586 architecture -- in the per-node case this is usually encapsulated as a container network
587 interface (CNI) provider which runs in a privileged context by default. In the sidecar case,
588 privilege is only needed at startup to establish traffic interception rules; depending on the
589 implementation (e.g., Kubernetes init containers) this can ensure that the privileged user is not
590 run alongside the application but only during initialization. In all cases, typically
591 CAP_NET_ADMIN is the only privilege required for mesh data plane functionality.

592 5.1.10. Data Plane (Service Mesh) Bypassed (TR-10)

593 **Threat Description**: Traffic is sent directly to a workload, bypassing mesh functionality and
594 authorization policies.

595 **Proxy Function Impacted**: Easiest to bypass of all the available models, from app choosing not
596 to use sidecar to container-local bypasses/configurations.

597 **Impact Score=2**: An app is exposed without mesh security controls.

598 **Likelihood Score=2**: Because the proxy runs in user space in the same cgroups as the
599 application, there are a variety of attacks available that are not relevant/applicable to other
600 implementations.

601 **Cumulative Threat Score**: (computed based on the methodology of Section 4.1) = 23

602 5.2. **Threat Analysis for Shared L4 – L7 per Service Model (DPA-2)**

603 5.2.1. **Compromised L4 Proxy (TR-1)**

604 **Threat Description:** Compromised L4 proxy (or L4 functions in the case of sidecar proxy with
605 combined L4 and L7 functions) leads to leaked identities for every workload (service) running
606 on the node.

607 **Proxy Function Impacted:** The L4 proxy has access to all the keys associated with the workloads
608 running on the node.

609 **Impact Score=3:** Identities of all workloads (services) on the node are compromised

610 **Likelihood Score=1:** only code delivering L4 functions is present. This minimal code footprint
611 and functionality presents the lowest attack surface of all options.

612 5.2.2. **Compromised Application Container (TR-2)**

613 **Threat Description:** Compromised application container (e.g., via a supply chain attack – during
614 development phase) leads to takeover of identity associated with that application.

615 **Proxy Function Impacted:** Data plane components are not located in the same pod as workload
616 containers, so a compromised workload does not necessarily lead to the access of keys /
617 secrets.

618 **Impact Score=1:** Single workload / single identity. No direct access to underlying key material.

619 **Likelihood Score=2:** Same regardless of architecture

620 5.2.3. **Compromise of Business Data (TR-3)**

621 **Threat Description:** Threat Description: identity is used to pivot through the infrastructure, in
622 order to compromise the confidentiality, integrity or availability of business data.

623 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1):** Same regardless of
624 architecture -- this is the fundamental risk of identity-based policy and is why we need to
625 practice the principle of least privilege (PoLP). The telemetry provided by the service mesh
626 (regardless of architecture) is invaluable for understanding communicating in your system and
627 creating accurate access policies (thereby implementing PoLP)

628 5.2.4. **Compromised L7 Proxy (TR-4)**

629 **Threat Description:** Vulnerability in L7 processing stack of the service mesh proxy. As L7
630 processing is inherently more complex, there is a higher probability for vulnerabilities to arise in
631 this part of the stack, as supported by historical CVE data.

632

633 **Proxy Function Impacted:** This topology allows 'less complex' L4 capabilities, e.g. mTLS, to be
634 adopted, with L7 processing only occurring if there is a strict requirement for it. Each service
635 account has its own dedicated L7 proxy.

636 **Impact Score=2:** A single set of workloads is impacted (DoS) / single identity leaked. In the
637 event of a DoS, it's much easier to make all workloads unavailable compared to the sidecar
638 model because the mesh's L7 processing is centralized into L7 "middle proxies". We need to
639 DoS this smaller number of middle proxies, vs needing to DoS every instance of the app in the
640 sidecar/library cases.

641 **Likelihood Score=1:** same as sidecar / same argument around potentially impacting all
642 workloads using L7 capabilities – see Section 5.1.4.

643 5.2.5. **Compromise of Shared L7 Proxy (TR-5)**

644 **Threat Description:** Co-tenant exploits L7 traffic processing vulnerability in shared proxy, to
645 affect the confidentiality, integrity or availability of traffic to/from another workload running on
646 the same node.

647 **Proxy Function Impacted:** By limiting the per-node functionality to L4 processing, the attack
648 surface is significantly reduced.

649 **Impact Score=1:** The application workload itself is unaffected, only the proxy – which is a
650 separate deployment. As long as the L7 proxy is not shared with the compromised application,
651 there is no impact.

652 **Likelihood Score=1:** As likely as the previous entry.

653 5.2.6. **Outdated Client Libraries in Applications (TR-6)**

654 **Threat Description:** Client libraries are not updated frequently or consistently across the estate
655 of microservices, leading to potential vulnerabilities and weaknesses that can be exploited.

656 **Proxy Function Impacted:** Infrastructure code decoupled from application code.

657 **Impact Score=1:** Same as the sidecar model, DPA-1 – see 5.1.6.

658 **Likelihood Score=1:** Same regardless of architecture.

659 5.2.7. **Denial of Service (TR-7)**

660 **Threat Description:** Conventional Denial of Service threat.

661 **Proxy Function Impacted:** A DoS executed at L4 has the same impact as the centralized per-
662 node model because the L4 process is centralized per node: all apps on the node are impacted.

663 A DoS executed at L7 impacts all app instances of the target app, since a (set of) dedicated L7
664 proxy(-ies) is deployed per app. The number of proxies implementing L7 functionality is

665 typically (far) less than the number of application instances making them an easier target for
666 DoS than "every instance of the target app".

667 **Impact Score=2**: Every instance of the target app. An L4 DoS would impact all application
668 instances on the target host.

669 **Likelihood Score=2**: The L4 proxy is deployed once per node, so it presents a better target for
670 DoS than DPA-1 or DPA-4; this is mitigated somewhat by the simplified functionality of an L4
671 proxy compared to a combined L4+L7 proxy.

672 The L7 proxy is shared by multiple instances of the same application, it presents an easier DoS
673 target than the application itself. Therefore it is more likely than the sidecar model, DPA-1.

674 5.2.8. Resource Consumption (TR-8)

675 **Threat Description**: Overall resource consumption by the data plane of the service mesh
676 infrastructure.

677 **Proxy Function Impacted**: The shared L4 proxy typically has a much lower memory (RAM)
678 footprint, as well as lower CPU usage overall due to a lower rate of change of config, less config
679 overall, and less responsibility than a combined L4+L7 sidecar proxy, DPA-1. For the service
680 mesh's data plane, L7 processing is typically the dominating CPU cost, followed by encryption.

681 L7 proxies are shared by all instances of the same application, deployed as a few traditional
682 "reverse proxies" per app. This results in much lower resource consumption for L7 processing
683 than the sidecar model (DPA-1). Overall DPA-2 uses more resources than the shared per node
684 model (DPA-3), but substantially less than the sidecar (DPA-1). This is due primarily to reduced
685 overhead -- e.g., an app with 50 instances requires 50 sidecars, but might be served with 5
686 shared L7 proxies (or less).

687 **Impact Score=2**: DPA-3 achieves a good middle ground: lower consumption than sidecar *and*
688 easier to achieve than sidecar (DPA-1); but not as low as all shared (DPA-3) or all in app (DPA-4).

689 **Likelihood Score=1**: Easy to achieve low resource usage.

690 5.2.9. Privileged L4 Proxy (TR-9)

691 **Threat Description**: Service mesh implementation requires L4 component (e.g. deployed as a
692 DaemonSet on a Kubernetes cluster) to run with an overprivileged security context (e.g.
693 Privileged Pod).

694 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1)**: Same regardless of
695 architecture -- in the per-node case this is usually encapsulated as a container network
696 interface (CNI) provider which runs in a privileged context by default. In the sidecar case,
697 privilege is only needed at startup to establish traffic interception rules; depending on the
698 implementation (e.g., Kubernetes init containers) this can ensure that the privileged user is not
699 run alongside the application but only during initialization. In all cases, typically
700 CAP_NET_ADMIN is the only privilege required for mesh data plane functionality.

701 5.2.10. **Data Plane (Service Mesh) Bypassed (TR-10)**

702 **Threat Description:** Traffic is sent directly to a workload, bypassing mesh functionality and
703 authorization policies.

704 **Proxy Function Impacted:** Part of the goal of moving enforcement out of the app context and
705 into a shared context is to use stronger primitives to ensure the non-bypass-ability of the mesh
706 data plane. In general, with a per-node L4 setup, sending traffic to an individual app instance on
707 the node should not be achievable (e.g. similar to [but not necessarily implemented as] a host-
708 level VPN requiring workloads to be part of the VPN overlay to connect).

709 L7 proxies are deployed independently from the applications they represent, which requires
710 special configuration in the mesh to ensure they're routed through, making bypassability easier
711 than other models. Impact of missing L7 policy can be significant. (In other models we rely on 0
712 or 1 things to ensure traffic is directed to the correct policy enforcement point; in this model
713 we rely on 2 things [traffic interception, mesh configuration to route via middle proxies] to
714 ensure traffic is subject to the correct PEPs)

715 **Impact Score=2:** An app is exposed without mesh security controls.

716 **Likelihood Score=2:** L4 controls are by-design built to mitigate this; L7 controls are easier to
717 bypass compared to sidecar model.

718 **Cumulative Threat Score:** (computed based on the methodology of Section 4.1) = 22

719 5.3. **Threat Analysis for Shared L4 and L7 Model (DPA-3)**

720 5.3.1. **Compromised L4 Proxy (TR-1)**

721 **Threat Description:** Compromised L4 proxy (or L4 functions in the case of sidecar proxy with
722 combined L4 and L7 functions) leads to leaked identities for every workload (service) running
723 on the node.

724 **Proxy Function Impacted:** The L4 proxy has access to all the keys associated with the workloads
725 running on the node.

726 **Impact Score=3:** All identities on node

727 **Likelihood Score=3:** L7 code may be enabled for another server (not yours) which can be
728 exploited to affect all apps on the host

729 5.3.2. **Compromised Application Container (TR-2)**

730 **Threat Description:** Compromised application container (e.g., via a supply chain attack – during
731 development phase) leads to takeover of identity associated with that application.

732 **Proxy Function Impacted:** Data plane components are not located in the same pod as workload
733 containers, so a compromised workload does not necessarily lead to the access of keys /
734 secrets.

735 **Impact Score=1**: Single workload / single identity. No direct access to underlying key material.

736 **Likelihood Score=2**: Same regardless of architecture.

737 5.3.3. **Compromise of Business Data (TR-3)**

738 **Threat Description**: Threat Description: identity is used to pivot through the infrastructure, in
739 order to compromise the confidentiality, integrity or availability of business data.

740 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1)**: Same regardless of
741 architecture -- this is the fundamental risk of identity-based policy and is why we need to
742 practice the principle of least privilege (PoLP). The telemetry provided by the service mesh
743 (regardless of architecture) is invaluable for understanding communicating in your system and
744 creating accurate access policies (thereby implementing PoLP)

745 5.3.4. **Compromised L7 Proxy (TR-4)**

746 **Threat Description**: Vulnerability in L7 processing stack of the service mesh proxy. As L7
747 processing is inherently more complex, there is a higher probability for vulnerabilities to arise in
748 this part of the stack, as supported by historical CVE data.

749 **Proxy Function Impacted**: This topology allows 'less complex' L4 capabilities, e.g., mTLS, to be
750 adopted, with L7 processing only occurring if there is a strict requirement for it. Blast radius of a
751 proxy compromise affects all workloads on the node. That means that its failure represents a
752 shared fate outage, and as a shared resource it's susceptible to denial of service attacks.

753 **Impact Score=3**: L7 capability is shared across all applications on the node, so if even a single
754 application's configuration causes the proxy to become susceptible to failure then all
755 applications on the node can be attacked (either a credential leak or denial of service,
756 depending on the attack).

757 **Likelihood Score=2**: For a given app using L7 capabilities, as likely as the sidecar model.
758 However, because workloads that are only doing L4 are susceptible to attack if they share the
759 same node (which under the sidecar model, DPA-1, would have been safe), likelihood is higher.

760 5.3.5. **Compromise of Shared L7 Proxy (TR-5)**

761 **Threat Description**: Co-tenant exploits L7 traffic processing vulnerability in shared proxy, to
762 affect the confidentiality, integrity, or availability of traffic to/from another workload running
763 on the same node.

764 **Proxy Function Impacted**: A single proxy instance does not provide an inherently multi-tenant
765 setup. Hence security concerns arise when combining complex processing rules for L7 traffic
766 from multiple unconstrained tenants in a shared instance. In this configuration, L7 processing of
767 multiple co-tenants' traffic is performed within one process, with no memory protection or
768 isolation benefits that could be gained by containerizing L7 functionality per workload

769 **Impact Score=3**: All workloads on the node are impacted.

770 **Likelihood Score=2**: See section 5.3.4 above – a compromise is as likely as the sidecar model
771 (DPA-1), but applications that would not be susceptible to attack under DPA-1 *are* susceptible
772 under this model, DPA-3.

773 5.3.6. Outdated Client Libraries in Applications (TR-6)

774 **Threat Description**: Client libraries are not updated frequently or consistently across the estate
775 of microservices, leading to potential vulnerabilities and weaknesses that can be exploited.

776 **Proxy Function Impacted**: Infrastructure code decoupled from application code.

777 **Impact Score=1**: Same as the sidecar model, DPA-1 – see 5.1.6.

778 **Likelihood Score=1**: Same regardless of architecture.

779 5.3.7. Denial of Service (TR-7)

780 **Threat Description**: Conventional Denial of Service threat.

781 **Proxy Function Impacted**: Because processing for all app instances on the node is shared, and a
782 single proxy instance is not inherently multi-tenant (provides no controls wrt resource
783 utilization across independent backends and clients), the blast radius of DoS on the mesh data
784 plane is every app on the node.

785 **Impact Score=3**: All workloads on the node.

786 **Likelihood Score=2**: If *any* app configuration triggers exploitable paths in the shared proxy, *all*
787 apps on the node suffer.

788 5.3.8. Resource Consumption (TR-8)

789 **Threat Description**: Overall resource consumption by the data plane of the service mesh
790 infrastructure.

791 **Proxy Function Impacted**: Because *all* functionality is shared at the node level, DPA-3 has the
792 most opportunity for deduplication -- therefore reduction in resource usage. Configuration like
793 service discovery need only be sent a single time to each node, rather than to each and every
794 app instance. Overall this means the lowest rate of change and least data transferred, as well as
795 a lower runtime footprint (RAM, CPU).

796

797 Note some implementations don't fully de-dupe configuration (for a variety of reasons, both
798 due to implementation and as a security measure to provide some degree of isolation), so
799 consume RAM more similarly to a sidecar case than might otherwise appear.

800

801 **Impact Score=1**: lowest overall resource utilization of all available architectures

802 **Likelihood Score=1**: easiest to achieve low resource utilization

803 5.3.9. Privileged L4 Proxy (TR-9)

804 **Threat Description:** Service mesh implementation requires L4 component (e.g. deployed as a
805 Daemon Set on a Kubernetes cluster) to run with an overprivileged security context (e.g.
806 Privileged Pod).

807 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1):** Same regardless of
808 architecture -- in the per-node case this is usually encapsulated as a container network
809 interface (CNI) provider which runs in a privileged context by default. In the sidecar case,
810 privilege is only needed at startup to establish traffic interception rules; depending on the
811 implementation (e.g., Kubernetes init containers) this can ensure that the privileged user is not
812 run alongside the application but only during initialization. In all cases, typically
813 CAP_NET_ADMIN is the only privilege required for mesh data plane functionality.

814 5.3.10. Data Plane (Service Mesh) Bypassed (TR-10)

815 **Threat Description:** Traffic is sent directly to a workload, bypassing mesh functionality and
816 authorization policies.

817 **Proxy Function Impacted:** Part of the goal of moving enforcement out of the app context and
818 into a shared context is to use stronger primitives to ensure the non-bypass-ability of the mesh
819 data plane. In general, with a per-node setup, sending traffic to an individual app instance on
820 the node should not be achievable (e.g., similar to [but not necessarily implemented as] a host-
821 level VPN requiring workloads to be part of the VPN overlay to connect).

822 **Impact Score=3:** All applications on the node are exposed without mesh security controls.

823 **Likelihood Score=1:** By design built to mitigate this kind of bypass

824 **Cumulative Threat Score:** (computed based on the methodology of Section 4.1) = 37

825 5.4. Threat Analysis for L4 and L7 within Application Model (gRPC proxyless Model (DPA- 826 4))

827 5.4.1. Compromised L4 Proxy (TR-1)

828 **Threat Description:** Compromised L4 proxy (or L4 functions in the case of sidecar proxy with
829 combined L4 and L7 functions) leads to leaked identities for every workload (service) running
830 on the node.

831 **Proxy Function Impacted:** mTLS connections are negotiated by the client library inside of the
832 application, with a single identity (the application's). In order to compromise key material and
833 identity documents for multiple workloads, multiple application instances would need to be
834 compromised.

835 **Impact Score=1:** Single workload / single identity.

836 **Likelihood Score=2**: Large surface area if something goes wrong, since we're inside the
837 application's context. Therefore, this is as likely or slightly more likely than DPA-1.

838 5.4.2. **Compromised Application Container (TR-2)**

839 **Threat Description**: Compromised application container (e.g., via a supply chain attack – during
840 development phase) leads to takeover of identity associated with that application.

841 **Proxy Function Impacted**: Compromising the application *is* compromising the mesh in this case;
842 full access to any key material used by the application -- including the mesh identity -- is
843 achievable.

844 **Impact Score=2**: Single workload / single identity. Full access to key material used by that
845 application.

846 **Likelihood Score=2**: Same regardless of architecture.

847 5.4.3. **Compromise of Business Data (TR-3)**

848 **Threat Description**: Threat Description: identity is used to pivot through the infrastructure, in
849 order to compromise the confidentiality, integrity or availability of business data.

850 **Proxy Function Impacted & Impact Score (=1) & Likelihood Score (=1)**: Same regardless of
851 architecture -- this is the fundamental risk of identity-based policy and is why we need to
852 practice the principle of least privilege (PoLP). The telemetry provided by the service mesh
853 (regardless of architecture) is invaluable for understanding communicating in your system and
854 creating accurate access policies (thereby implementing PoLP)

855 5.4.4. **Compromised L7 Proxy (TR-4)**

856 **Threat Description**: Vulnerability in L7 processing stack of the service mesh proxy. As L7
857 processing is inherently more complex, there is a higher probability for vulnerabilities to arise in
858 this part of the stack, as supported by historical CVE data.

859 **Proxy Function Impacted**: Compromising the L7 processing stack results in compromising the
860 entire application, resulting in more risk of compromise beyond runtime identity and DoS for
861 other users.

862 **Impact Score=3**: The application itself is compromised, including non-mesh credentials (e.g.
863 `truncate table users;`) that are not available if only the proxy is compromised.

864 **Likelihood Score=3**: L7 processing code *is* the application, and as a result the surface area is
865 much larger.

866 5.4.5. **Compromise of Shared L7 Proxy (TR-5)**

867 **Threat Description:** Co-tenant exploits L7 traffic processing vulnerability in shared proxy, to
868 affect the confidentiality, integrity or availability of traffic to/from another workload running on
869 the same node.

870 **Proxy Function Impacted:** L7 processing is entirely isolated by whatever mechanisms isolate
871 applications themselves (containers, micro-VMs, VMs, etc). Impact is limited by the strength of
872 that boundary.

873 **Impact Score=1:** See section 5.1.5.

874 **Likelihood Score=1:** As likely as any other application compromise.

875 5.4.6. **Outdated Client Libraries in Applications (TR-6)**

876 **Threat Description:** Client libraries are not updated frequently or consistently across the estate
877 of microservices, leading to potential vulnerabilities and weaknesses that can be exploited.

878 **Proxy Function Impacted:** Infrastructure concerns are embedded within application code.
879 Challenges can arise when enforcing consistency in versions between microservices etc.

880 **Impact Score=3:** The mesh functionality itself is part of the application, therefore bad
881 application updates mean bad mesh updates. This means vulnerabilities stick around for longer.
882 By the same token, since the mesh is part of the app, a vulnerability in the app *is* a vulnerability
883 in the mesh data plane.

884 **Likelihood Score=2:** Depends on frequency of update -- if applications can be updated quickly
885 (i.e., on the order of minutes to hours), likelihood is low. If applications take on the order of
886 weeks to months to update, likelihood is high. In the realm of days-to-update we have a middle
887 ground of risk that's likely acceptable to most organizations. However cross-cutting concerns
888 like mesh data plane, which are critical to the org's overall security posture, should be patched
889 as soon as possible.

890 5.4.7. **Denial of Service (TR-7)**

891 **Threat Description:** Conventional Denial of Service threat.

892 **Proxy Function Impacted:** A DoS of the mesh data plane (L4 or L7) is a DoS of the application
893 itself. In all other respects, it's very similar to the sidecar.

894 **Impact Score=1:** Single instance of single app. See 5.1.7 – an attack could be repeated across all
895 applications.

896 **Likelihood Score=2:** Not just mesh data plane functionality is susceptible to DoS, but application
897 code/functionality itself.

898 5.4.8. Resource Consumption (TR-8)

899 **Threat Description:** Overall resource consumption by the data plane of the service mesh
900 infrastructure.

901 **Proxy Function Impacted:** Because it's built into the app, resources devoted to mesh data plane
902 functionality are *very* low. The only reason resource utilization overall winds up being higher
903 than the shared L4/L7 model (DPA-3) is because some duplication of configuration and
904 processing needs to happen since configuration needs to be pushed to every application
905 instance.

906 **Impact Score=2:** Potentially lower resource usage on a per-app basis than any other model, but
907 likely higher in aggregate because we can't share any resources or configuration across data
908 plane instances.

909 **Likelihood Score=1:** Easy to achieve low resource usage.

910 5.4.9 Privileged L4 Proxy (TR-9)

911 **Threat Description:** Service mesh implementation requires L4 component (e.g., deployed as a
912 Daemon Set on a Kubernetes cluster) to run with an overprivileged security context (e.g.,
913 Privileged Pod).

914 **Proxy Function Impacted & Impact Score (=0) & Likelihood Score (=0):** Mesh data plane
915 functionality runs in the application context without any special privileges -- it's the same as the
916 app itself. No special capabilities or permissions are required to intercept traffic or implement
917 policy enforcement.

918 5.4.10 Data Plane (Service Mesh) Bypassed (TR-10)

919 **Threat Description:** Traffic is sent directly to a workload, bypassing mesh functionality and
920 authorization policies.

921 **Proxy Function Impacted:** App *is* the enforcement point, there is no bypassing.

922 **Impact Score=1:** The app is exposed in a degraded state or without some controls.

923 **Likelihood Score=1:** By nature of RPC frameworks and in-process enforcement, mesh data plane
924 policy should not be bypassable.

925 **Cumulative Threat Score:** (computed based on the methodology of Section 4.1) = 28

926

927 **6. Recommendations Based on Application Security Risk Profile**

928 While the ratings or scores for the impact and likelihood parameters for different threats in
929 different data plane architectures are dictated by the number of service instances affected, the
930 risk profiles associated with applications are determined by the criticality of the entire
931 application with respect to the business process it supports.

932 In arriving at the threat profile for each of the architectural patterns considered in section 5,
933 please recall that we observed that for some threats, the impact and likelihood parameters are
934 the same irrespective of the proxy model or data plane architecture. The ratings assigned to
935 these parameters are as we already stated, are relative ratings, and hence both parameters are
936 assigned the rating 1, resulting in the overall threat rating of 1 for those threats. The threats
937 that come under this category are:

- 938 • Compromise of Business Data (TR-3)
- 939 • Privileged L4 Proxy (TR-9)

940 Hence, we have to ignore the threat ratings of the above listed threats and dwell into the
941 consideration of threat ratings for the other remaining threats. While considering the remaining
942 threats, we must ignore those threats that have no direct security implications but may have
943 performance implications. The only threat that comes under this category is:

- 944 • Resource Consumption (TR-8)

945 **6.1. Cloud-Native Applications with Low Risk Profile**

946 The service mesh capability requirements for this class of application are as follows:

- 947 • LOW-REQ1: Service-to-Service authorization (Service A can call Service B) is based on
948 network location/parameter (e.g., subnet) and authorization at the granularity of the
949 called service method, calling user and per call request are not required.
- 950 • LOW-REQ2: Logging and Metrics need to be captured only at the level of network
951 parameters (e.g., Source/Destination TCP address) and not at the level of per call
952 request.
- 953 • LOW-REQ3: All traffic management capabilities such as load balancing, rate limiting etc.
954 need to be enforced at the network connection level and not at the per call request.

955 Examination of the above capabilities reveals that these essentially involve network
956 transport/network level data, that can be all provided by proxy's L4 functions and hence by L4
957 proxies. Hence the following are recommendations for this class of application.

958 Recommendations:

- 959 1. Since all requirements can be met by L4 proxies or L4 functions built into the libraries,
960 all four data plane architectures can be theoretically used.
- 961 2. Since neither method-level nor per call request handling is required, thus eliminating all
962 L7 functions, data plane architectures that deploy a L7 proxy per service instance (side-

963 car model (DPA-1) expose an unnecessary attack surface. Therefore, either of the two
964 models with a shared L4 proxy (DPA-2 and DPA-3) is recommended. gRPC proxy-less
965 model (DPA-4) is also usable for this class of applications, though it does expose a larger
966 attack surface than DPA-2 or DPA-3.

967 **6.2. Cloud-Native Applications with Medium Risk Profile**

968 The service mesh capability requirements for this class of application are as follows:

- 969 • MEDIUM-REQ1: In addition to Service-to-Service authorization at the level of service, a
970 full authorization policy at the method level (Service A can execute GET on B's Billing
971 method with valid end user credentials containing the READ scope) is required.
- 972 • MEDIUM-REQ2: Logging and Metrics data need to be captured not only at the level of
973 network parameters (e.g., Source/Destination TCP address) but also some metadata
974 such as the called service and method.
- 975 • MEDIUM-REQ3: All traffic management capabilities such as load balancing, rate limiting
976 etc. can be enforced at the network connection level (as in low risk profile case) and not
977 at the per call request or per method level.

978 Examination of the above capabilities reveals that these essentially involve not only network
979 transport/network level data (all L4 functions), but also some L7 functions (not all) such as
980 authenticating user identities not only locally from tokens (e.g., Jason Web Tokens (JWT)) but
981 also remotely using standardized protocols such as OAuth and OIDC. Hence use of L7 proxies
982 with some limited functionality is mandatory. Hence the following are recommendations for
983 this class of application.

984 Recommendations:

- 985 1. Just like for applications with low risk profile, all four data plane architectures can be
986 theoretically used.
- 987 2. Since L7 functions are limited, it is not essential to dedicate a L7 proxy for each service.
988 Hence, data plane architectures that deploy a L7 proxy for each service (side-car model
989 (DPA-1)) may end up consuming more resources than other models for limited
990 additional assurance. On the other hand, as previously discussed, L7 code is where most
991 exploitable vulnerabilities lie. Hence shared L4-L7 model (DPA-3) is not desirable since
992 the shared L7 component introduces risk for all services that share the same physical
993 host. Therefore, the shared L4 -- L7 per service model (DPA-2) is likely the best mix of
994 resource utilization and risk. gRPC proxy-less model (DPA-4) with inclusion of libraries
995 for L4 functions and limited L7 functions is also recommended, with similar risk but even
996 less resource utilization than DPA-2 in most cases.

997 **6.3. Cloud-Native Applications with High Risk Profile**

998 The service mesh capability requirements for this class of application are as follows:

- 999 • HIGH-REQ1: In addition to: (a) Service-to-Service authorization at the level of service
1000 and, (b) a full authorization policy at the method level (Service A can execute GET on B's
1001 Billing method with valid end user credentials containing the READ scope), a full user to
1002 resource level access control is required. The last requirement necessitates the proxy
1003 making an external authorization call for each request.
- 1004 • HIGH-REQ2: Logging and Metrics meta data relating to a request must be captured –
1005 rate of requests, rate of positive outcomes, processing time for each request etc.
- 1006 • HIGH-REQ3: All traffic management capabilities are required at the request level and
1007 should involve application layer parameters in addition to those at the network
1008 connection level.

1009 Examination of the above capabilities reveals that a complete suite of L7 functions is required.

1010 Recommendations:

- 1011 1. Just like for applications with low risk and medium risk profiles, all four data plane
1012 architectures can be theoretically used.
- 1013 2. However, based on the requirements, this class of applications belong to Highly critical
1014 applications, which require a great degree of isolation, where any compromise, if it
1015 occurs should be limited to only one service instances and not multiple service
1016 instances. Hence, data plane architectures that deploy a L7 proxy for each service (side-
1017 car model (DPA-1)) is most applicable. A shared L7 proxy per Service (like DPA-2) can be
1018 an acceptable tradeoff for some organizations, provided they have other mechanisms
1019 for mitigating shared-fate failures of all instances of the service that the shared service
1020 mesh L7 proxy brings (e.g., mitigating a denial-of-service attack via L3 controls outside
1021 the mesh). However, tightly integrating both L4 & L7 functions with the service instance
1022 provides a greater degree of isolation and hence the former data plane architecture
1023 (DPA-1) is highly recommended.

1024

1025 **7. Summary and Conclusions**

1026 Microservices-based applications implemented using containers & VMs and sometimes
1027 spanning on-premises and multiple clouds go by the name of cloud-native applications. In
1028 instances where a centralized service infrastructure is beneficial to the overall security of this
1029 class of applications, this need is met by a service mesh.

1030 Service mesh implementations are characterized by the type of configurations of entities called
1031 proxies which are the engines that enable various capabilities during application runtimes -
1032 such as policy enforcement (including access control), network connectivity (including
1033 establishment of secure sessions), performance monitoring (through collection of data for
1034 computing various metrics) etc. The proxies thus form the data plane of the service mesh, and a
1035 particular configuration of proxies is called a proxy model or a data plane architecture.

1036 The first and still the widely prevalent deployment of the proxy model is the side car model
1037 where a single proxy that provides functions both at the L4 and L7 level is associated with a
1038 service instance. Performance, resource consumption and specific security needs for different
1039 cloud-native applications have led to exploration of alternate proxy models.

1040 In this document, we performed a detailed threat analysis of these alternate proxy models
1041 (including the ones that provide the needed security functions without proxies) by identifying
1042 ten common threats and provided recommendations for their use in cloud-native applications
1043 with different security risk profiles.

1044 **References**

- 1045 [1] Wikipedia (2024) OSI Model. Available at https://en.wikipedia.org/wiki/OSI_model
- 1046 [2] Chandramouli R, Butcher Z (2020) Building Secure Microservices-based Applications Using
1047 Service-Mesh Architecture. (National Institute of Standards and Technology, Gaithersburg,
1048 MD), NIST Special Publication (SP) NIST SP 800-204A.
1049 <https://doi.org/10.6028/NIST.SP.800-204A>
- 1050 [3] Chandramouli R, Butcher Z, Aradhna C (2021) Attribute-based Access Control for
1051 Microservices-based Applications using a Service Mesh. (National Institute of Standards
1052 and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-204B.
1053 <https://doi.org/10.6028/NIST.SP.800-204B>
- 1054 [4] Chandramouli R (2022) Implementation of DevSecOps for a Microservices-based
1055 Application with Service Mesh. (National Institute of Standards and Technology,
1056 Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-204C.
1057 <https://doi.org/10.6028/NIST.SP.800-204C>
- 1058 [5] Chandramouli R, Butcher Z (2023) A Zero Trust Architecture Model for Access Control in
1059 Cloud-Native Applications in Multi-Cloud Environments. (National Institute of Standards
1060 and Technology, Gaithersburg, MD), NIST Special Publication (SP) NIST SP 800-207A.
1061 <https://doi.org/10.6028/NIST.SP.800-207A>
- 1062 [6] Jackson E, Kohavi Y, Pettit J, Posta C (2022) Ambient Mesh Security Deep Dive. (Istio)
1063 Available at <https://istio.io/latest/blog/2022/ambient-security/>
- 1064 [7] Howard J, Jackson EJ, Kohavi Y, Levine I, Pettit J, Sun L (2022) Introducing Ambient Mesh.
1065 (Istio) Available at [https://istio.io/latest/blog/2022/introducing-ambient-mesh/#what-
1066 about-security](https://istio.io/latest/blog/2022/introducing-ambient-mesh/#what-about-security)
- 1067 [8] Turner M (2022) eBPF and Sidecars - Getting the Most Performance and Resiliency out of
1068 the Service Mesh. (Tetrade) Available at [https://tetrade.io/blog/ebpf-and-sidecars-getting-
1069 the-most-performance-and-resiliency-out-of-the-service-mesh/](https://tetrade.io/blog/ebpf-and-sidecars-getting-the-most-performance-and-resiliency-out-of-the-service-mesh/)
- 1070 [9] Graf T (2021) How eBPF will solve Service Mesh - Goodbye Sidecars. (Isovalent) Available
1071 at <https://isovalent.com/blog/post/2021-12-08-ebpf-servicemesh/>
- 1072 [10] Song J (2022) Transparent Traffic Intercepting and Routing in the L4 Network of Istio
1073 Ambient Mesh. (Tetrade) Available at [https://tetrade.io/blog/transparent-traffic-
1074 intercepting-and-routing-in-the-l4-network-of-istio-ambient-mesh/](https://tetrade.io/blog/transparent-traffic-intercepting-and-routing-in-the-l4-network-of-istio-ambient-mesh/)
- 1075 [11] Song J (2022) L7 Traffic Path in Ambient Mesh. (Tetrade) Available at
1076 <https://tetrade.io/blog/l7-traffic-path-in-ambient-mesh/>
- 1077 [12] Cilium (2024) Threat Model — Cilium 1.15.6 documentation. (Cilium) Available at
1078 <https://docs.cilium.io/en/stable/security/threat-model/>
- 1079 [13] Istio (2024) Ambient mode overview: ztunnel. Available at
1080 <https://istio.io/latest/docs/ambient/overview/#ztunnel>
- 1081 [14] Landow S (2021) gRPC Proxyless Service Mesh. (Istio) Available at
1082 <https://istio.io/v1.15/blog/2021/proxyless-grpc/>
1083