

2

3 **Implementation of DevSecOps for a**
4 **Microservices-based Application with**
5 **Service Mesh**

6

7

8

9

Ramaswamy Chandramouli

10

11

12

13

14

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

15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44

Draft NIST Special Publication 800-204C

Implementation of DevSecOps for a Microservices-based Application with Service Mesh

Ramaswamy Chandramouli
*Computer Security Division
Information Technology Laboratory*

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

September 2021



U.S. Department of Commerce
Gina M. Raimondo, 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 Technoogy & Director, National Institute of Standards and Technology*

45

Authority

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

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

59 National Institute of Standards and Technology Special Publication 800-204C
60 Natl. Inst. Stand. Technol. Spec. Publ. 800-204C, 40 pages (September 2021)
61 CODEN: NSPUE2

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

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

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

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

77 **Public comment period: *September 29, 2021 through November 1, 2021***

78 National Institute of Standards and Technology
79 Attn: Computer Security Division, Information Technology Laboratory
80 100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930
81 Email: sp800-204c-comments@nist.gov

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

84

Reports on Computer Systems Technology

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

93

Abstract

94 Cloud-native applications have evolved into a standardized architecture consisting of multiple
95 loosely coupled components called microservices (implemented as containers), supported by
96 code for providing application services called service mesh. Both of these components are hosted
97 on a container orchestration and resource management platform, which is called a reference
98 platform in this document. Due to security, business competitiveness, and its inherent structure
99 (loosely coupled application components), this class of applications needs a different application
100 development, deployment, and runtime paradigm. DevSecOps (consisting of three acronyms for
101 Development, Security, and Operations, respectively) has been found to be a facilitating
102 paradigm for these applications with primitives such as Continuous Integration, Continuous
103 Delivery, and Continuous Deployment (CI/CD) pipelines. These pipelines are workflows for
104 taking the developer's source code through various stages, such as building, testing, packaging,
105 deployment, and operations supported by automated tools with feedback mechanisms. In
106 addition to the application code, and the code for application services (service mesh), the
107 architecture has functional elements for infrastructure (computing, networking, and storage
108 resources), runtime policies (authentication, authorization etc.) and continuous monitoring of the
109 health of the application (Observability), which can be deployed through declarative codes.
110 Thus, separate CI/CD pipelines can be created for all of the five code types.
111 The objective of this document is to provide guidance for the implementation of DevSecOps
112 primitives for a reference platform hosting a cloud-native application with the functional
113 elements described above. The benefits of this approach for high security assurance and for
114 enabling continuous authority to operate (C-ATO) are also discussed.
115

116

Keywords

117 container orchestration and resource management platform; DevSecOps; CI/CD pipelines;
118 infrastructure as code; policy as code; observability as code; GitOps; workflow models; static
119 AST; dynamic AST; interactive AST; SCA.

120

121

Acknowledgments

122 The author would like to express their first thanks to Mr. David Ferraiolo of NIST for initiating
123 this effort to provide a targeted guidance for the implementation of DevSecOps primitives for the
124 development, deployment, and monitoring of services in microservices-based applications with
125 service mesh infrastructure. Sincere thanks to Mr. Nicolas Chaillan, CSO US Air Force, for a
126 detailed and insightful review and feedback. Thanks are also due to Zack Butcher of Tetrade, Inc.
127 for suggesting the title for this document. The author also expresses thanks to Isabel Van Wyk of
128 NIST for her detailed editorial review.

129

Call for Patent Claims

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

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

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

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

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

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

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

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

154 Such statements should be addressed to: sp800-204c-comments@nist.gov

155

156 **Executive Summary**

157 Cloud-native applications have evolved into a standardized architecture consisting of the
158 following components:

- 159 • Multiple loosely coupled components called microservices (implemented as containers)
- 160 • An application services infrastructure called Service Mesh (providing services such as
161 secure communication, authentication, and authorization for users, services, and devices)

162
163 Due to security, business competitiveness, and its inherent structure (loosely coupled application
164 components), this class of applications needs a different application, deployment, and runtime
165 monitoring paradigm – collectively called the software life cycle paradigm. DevSecOps
166 (consisting of three acronyms for Development, Security, and Operations, respectively) is one of
167 the facilitating paradigms for the development, deployment, and operation of these applications
168 with primitives such as Continuous Integration, Continuous Delivery, and Continuous
169 Deployment (CI/CD) pipelines.

170
171 CI/CD pipelines are workflows for taking the developer’s source code through various stages,
172 such as building, functional testing, security scanning for vulnerabilities, packaging, and
173 deployment supported by automated tools with feedback mechanisms. In addition to the
174 application code and the code for providing application services, this architecture may be made
175 up of functional elements for infrastructure, runtime policies (such as zero trust policy), and
176 continuous monitoring of the health of the application, the last three of which can be deployed
177 through declarative codes. Thus, separate CI/CD pipelines can be created for all five code types.
178 The runtime behavior of each of these code types is also described to highlight the roles that they
179 play in the overall execution of the application.

180
181 Though cloud-native applications have a common architectural stack, the platform on which the
182 components of the stack run may vary. The platform is an abstraction layer over a physical (bare
183 metal) or virtualized (e.g., virtual machines, containers) infrastructure. In this document, the
184 chosen platform is a container orchestration and resource management platform (e.g.,
185 Kubernetes). To unambiguously refer to this platform or application environment throughout this
186 document, it is called the *Reference Platform for DevSecOps Primitives*, or simply the *reference*
187 *platform*.

188
189 The objective of this document is to provide guidance for the implementation of DevSecOps
190 primitives for the reference platform. The benefits of this implementation for high security
191 assurance and the use of the artifacts within the pipelines for providing continuous authority to
192 operate (C-ATO) using risk management tools and dashboard metrics are also described.

193
194
195
196
197
198

199
200

Table of Contents

201 **Executive Summary iv**

202 **1 Introduction 1**

203 1.1 Scope..... 1

204 1.2 Related DevSecOps Initiatives 2

205 1.3 Target Audience..... 3

206 1.4 Relationship to Other NIST Guidance Documents..... 3

207 1.5 Organization of this document 4

208 **2 Reference Platform for the Implementation of DevSecOps Primitives 5**

209 2.1 Container Orchestration and Resource Management Platform..... 5

210 2.1.1 Security Limitations of Orchestration Platform..... 5

211 2.2 Service Mesh Software Architecture 6

212 2.2.1 Data Plane..... 6

213 2.2.2 Control Plane 7

214 **3 DevSecOps – Organizational Preparedness, Key Primitives, and**

215 **Implementation..... 9**

216 3.1 Organizational Preparedness for DevSecOps 9

217 3.2 Seamless Evolution from DevOps to DevSecOps 9

218 3.3 DevSecOps – Key Primitives and Implementation Tasks 10

219 3.3.1 Concept of Pipelines and the CI/CD pipeline 10

220 3.3.2 Building Blocks for CI/CD pipelines 12

221 3.3.3 Designing and Executing the CI/CD pipeline..... 13

222 3.3.4 Strategies for Automation..... 15

223 3.3.5 Requirements for Automation Tools in CI/CD Pipelines 16

224 **4 Implementing DevSecOps Primitives for the Reference Platform 17**

225 4.1 Code Types in Reference Platform and Associated CI/CD Pipelines..... 17

226 4.2 CI/CD Pipeline for Application Code and Application Services Code..... 19

227 4.3. CI/CD Pipeline for Infrastructure as Code 19

228 4.3.1 Comparison of Configuration and Infrastructure..... 20

229 4.4 CI/CD Pipeline for Policy as Code..... 20

230 4.5 CI/CD Pipeline for Observability as Code 21

231 4.6 Securing the CI/CD Pipeline..... 21

232 4.7 Workflow Models in CI/CD Pipelines..... 22

233 4.7.1 GitsOps Workflow Model for CI/CD – A Pull-based Model..... 22

234 4.8 Security Testing – Common Requirement for CI/CD Pipelines for All Code

235 Types 23

236 4.8.1 Functional and Coverage Requirements for AST tools 24

237 4.9 Benefits of DevSecOps Primitives to Application Security in the Service Mesh

238 25

239 4.10 Leveraging DevSecOps for Continuous Authorization to Operate (c-ATO)... 26

240 **5 Summary and Conclusion..... 28**

241 **References..... 29**

242

243 1 Introduction

244 Cloud-native applications are made up of multiple loosely coupled components (called
245 microservices implemented as containers), operate in perimeter-less network environments
246 requiring zero trust concepts (on-premises or cloud), and are accessed by users from a diverse set
247 of locations (e.g., campus, home office, etc.). Cloud-native applications do not just refer to
248 applications that run in the cloud. They also refer to the class of applications with design and
249 runtime architectures, such as microservices, and a dedicated infrastructure for providing all
250 application services, such as security. The incorporation of zero trust principles into this class of
251 application provides techniques where access to all protected resources is enforced through
252 identity-based protection, as well as network-based protections (e.g., micro-segmentation) where
253 applicable.

254 Cloud-native applications require agile and secure updates and deployment techniques for
255 business reasons as well as the necessary resilience to respond to cybersecurity events. Hence,
256 they call for a different application development, deployment, and runtime monitoring paradigm
257 (collectively called the software life cycle paradigm) than the ones used for traditional monolithic
258 or multi-tier applications. DevSecOps (Development, Security, and Operations) is a facilitating
259 paradigm for this class of applications since it facilitates agile and secure development, delivery,
260 deployment, and operations through (a) primitives, such as Continuous Integration, Continuous
261 Delivery, and Continuous Deployment (CI/CD) pipelines (explained in section 3); (b) security
262 testing throughout the life cycle; and (c) continuous monitoring during runtime, all of which are
263 supported by automation tools. Thus, DevSecOps has the necessary primitives and other building
264 blocks to meet the design goals of cloud-native applications.

265 1.1 Scope

266 DevSecOps primitives can, in theory, be applied to many application architectures but are best
267 suited to microservices-based ones, which permit agile development paradigms due to the fact
268 that the application is made up of relatively small, loosely coupled modules called microservices.
269 Even within microservices-based architectures, the implementation of DevSecOps primitives can
270 take on different forms, depending on the platform. In this document, the chosen platform is a
271 container orchestration and resource management platform (e.g., Kubernetes). The platform is an
272 abstraction layer over a physical (bare metal) or virtualized (e.g., virtual machines, containers)
273 infrastructure. To unambiguously refer to this platform or application environment throughout
274 this document, it is called the *Reference Platform for DevSecOps Primitives*, or simply the
275 *reference platform*.

276
277 Before describing the implementation of DevSecOps primitives for the reference platform, it is
278 assumed that the following due diligence is applied with respect to deployment of the service
279 mesh component [1]:

- 280 • Secure design patterns for deploying and managing service mesh-based components for
281 infrastructure (e.g., network routing), policy enforcement, and monitoring

- 282 • Tests to prove that these service mesh components work as intended in a variety of scenarios
283 for all aspects of the application, such as ingress, egress, and inside services.

284 The guidance provided for implementation of DevSecOps primitives for the reference platform is
285 agnostic to (a) the tools used in DevSecOps pipelines and (b) the service mesh software, which
286 provides application services, though examples from service mesh offerings, such as Istio, are
287 used to link them to real-world application artifacts (e.g., containers, policy enforcement
288 modules, etc.).

289 The reference platform, along with DevSecOps implementation components, can therefore be
290 looked upon as having the following functional elements or code types:

- 291 1. The application code, which contains the application logic for transactions and database
292 access
- 293 2. The service mesh code, which provides application services such as network routes,
294 network resiliency services (e.g., load balancing, retries etc.), and security services
- 295 3. Infrastructure (consisting of computing, networking, and storage resources) as code
- 296 4. Policy as code for generating rules and configuration parameters for realizing security
297 objectives, such as zero trust through security controls (e.g., authentication,
298 authorization) during runtime
- 299 5. Observability as code for (a) triggering software related to logging, tracing
300 (communication pathways involved in executing application request), and monitoring for
301 recording all transactions and for (b) keeping track of application states during runtime.
302

303 This document covers the implementation of pipelines or workflows associated with all five code
304 types listed above. Out of them, the infrastructure as code, policy as code, and observability as
305 code belong to a special class called declarative code. A declarative code is written in a special
306 purpose language called declarative language that declares the requirements, and an associated
307 tool converts them into artifacts that make up a runtime instance. For example, in the case of
308 infrastructure as code (IaC), the declarative language models the infrastructure as a series of
309 resources. The associated configuration management tool pulls together these resources and
310 generates what are known as *Manifests* that define the final shape and state of the platform
311 (runtime instance) associated with the defined resources. These manifests are stored in servers
312 associated with a configuration management tool and are used by the tool to create compiled
313 configuration instructions for the runtime instance on the designated platform. Manifests are
314 generally encoded in platform-neutral representations (e.g., JSON) and fed to platform resource
315 provisioning agents through REST APIs.
316

317 **1.2 Related DevSecOps Initiatives**

318 **DoD Enterprise DevSecOps Initiative:** The Department of Defense (DOD) Enterprise
319 DevSecOps initiative spans many open-source projects with many state-of-practice ingredients.
320 The artifacts developed under this initiative are used by 12 U.S. federal agencies, five nations,
321 and a substantial number of DoD contractors and vendors. The open-source projects under this
322 initiative include but are not limited to [2]:

- 323 • The Iron Bank, a repository of DOD-vetted hardened container images
- 324 • Platform One, the DevSecOps platform that the department created for software deployments
- 325 from jets to bombers to space to clouds; it is based on the concept of continuous authority to
- 326 operate (cATO), which updated the DOD’s authorization process to accommodate the speed
- 327 and frequency of modern continuous software deployments

328
329 The DoD Enterprise DevSecOps initiative has the following state-of-practice ingredients:

- 330 1. **Abstracted:** To avoid drifts, be agnostic to the environment (e.g., cloud, on-premise,
- 331 classified, disconnected), and prevent lock-ins with cloud or platform layers, CNCF-
- 332 compliant Kubernetes and OCI-compliant containers are leveraged – open-source stacks
- 333 with U.S. observation of code and continuous scanning.
- 334 2. **GitOps/Infrastructure as Code (IaC):** No drift, everything is code (including
- 335 configuration, networking etc.), and the entire stack is automatically instantiated.
- 336 3. **Continuous Integration/Continuous Delivery pipeline (CI/CD):** Fully containerized
- 337 and using Infrastructure as Code (IaC)
- 338 4. **Hardened Containers:** Hardened “Lego blocks” to bring options to development teams
- 339 (one size fits all lead to shadow IT)
- 340 5. **Software Testing:** Mandated high-test coverage
- 341 6. **Baked-in Security:** Mandated static/dynamic code analysis, container security, bill of
- 342 material (supply chain risk), etc.
- 343 7. **Continuous Monitoring:**
 - 344 • **Centralized logging and telemetry**
 - 345 • Automated alerting
 - 346 • **Zero trust**, leveraging service mesh as sidecar (part of SCSS) down to the
 - 347 container level
 - 348 • **Behavior detection** (automated prevention leveraging AI/ML capabilities)
 - 349 • CVE signatures scanning
- 350 8. **Chaos engineering:** Dynamically kills/restarts container with moving target defense

351 1.3 Target Audience

352
353 Since DevSecOps primitives span development (build and test for security, package, delivery,

354 deployment, continuous monitoring) and ensure secure states during runtime, the target audience

355 for the recommendations in this document includes software development, operations, and

356 security teams.

357 1.4 Relationship to Other NIST Guidance Documents

359 Since the Reference Platform is made up of container orchestration and resource management

360 platform and service mesh software, the following publications offer guidance for securing this

361 platform as well as background information for the contents of this document:

- 362 • Special Publication (SP) 800-204, *Security Strategies for Microservices-based Application*
- 363 *Systems* [3], discusses the characteristics of microservices-based applications and the overall

- 364 security requirements and strategies for addressing those requirements.
- 365 ● SP 800-204A, *Building Secure Microservices-based Applications Using Service-Mesh*
- 366 *Architecture* [4], provides deployment guidance for various security services (e.g.,
- 367 establishment of secure sessions, security monitoring, etc.) for a microservices-based
- 368 application using a dedicated infrastructure (i.e., a service mesh) that uses service proxies that
- 369 operate independent of the application code.
- 370 ● SP 800-204B [5], *Attribute-based Access Control for Microservices-based Applications*
- 371 *Using a Service Mesh*, provides deployment guidance for building an authentication and
- 372 authorization framework within the service mesh that meets the security requirements, such
- 373 as (1) zero trust by enabling mutual authentication in communication between any pair of
- 374 services and (2) a robust access control mechanism based on an access control such as
- 375 attribute-based access control (ABAC) that can be used to express a wide set of policies and
- 376 is scalable in terms of user base, objects (resources), and deployment environment.

377 1.5 Organization of this document

378 The organization of the rest of the document is as follows:

- 379 ● Chapter 2 gives a brief description of the reference platform for which guidance for
- 380 implementation of DevSecOps primitives is provided.
- 381 ● Chapter 3 introduces the DevSecOps primitives (i.e., pipelines), the methodology for
- 382 designing and executing the pipelines, and the role that automation plays in the execution.
- 383 ● Chapter 4 covers all facets of pipelines, including (a) Common issues to be addressed for
- 384 all pipelines, (b) descriptions of the pipelines for the four code types in the reference
- 385 platform that are listed in Section 1.1, and (c) the benefit of DevSecOps for security
- 386 assurance for the entire application environment (the reference platform with five code
- 387 types, thus carrying the DevSecOps implementation) during the entire life cycle,
- 388 including the “Continuous Authority to Operate (CAT).”
- 389 ● Chapter 5 provides a summary and conclusion.

390

391 **2 Reference Platform for the Implementation of DevSecOps Primitives**

392 As stated in Section 1.1, the reference platform is a container orchestration and management
393 platform. In modern application environments, the platform is an abstraction layer over a physical
394 (bare metal) or virtualized (e.g., virtual machines, containers) infrastructure. Before the
395 implementation of DevSecOps primitives, the platform simply contains the application code,
396 which contains the application logic and the service mesh code, which in turn provides
397 application services. This section will consider the following:

- 398 • A container orchestration and resource management platform that houses both the
399 application code and most of the service mesh code
- 400 • The service mesh software architecture

401 **2.1 Container Orchestration and Resource Management Platform**

403 Since microservices are implemented as containers, a container orchestration and resource
404 management platform is used for deployment, operations, and upgrading services. Kubernetes is
405 one such open-source container orchestration and resource management platform that is widely in
406 use.

407 A typical orchestration and resource management platform consists of various logical (forming
408 the abstraction layer) and physical artifacts for the deployment of containers. For example, in
409 Kubernetes, containers run inside the smallest unit of deployment called a pod. A pod can
410 theoretically host a group of containers, though usually, only one container runs inside a pod. A
411 group of pods are defined inside what is known as a node, where a node can be either a physical
412 or virtual machine (VM). A group of nodes constitutes a cluster. Usually, multiple instances of a
413 single microservice are needed to distribute the workload to achieve the desired performance
414 level. A cluster is a pool of resources (nodes) that is used as a means to distribute the workload of
415 microservices. One of the techniques used is horizontal scaling, where microservices that are
416 accessed more frequently are allocated more instances or allocated to nodes with higher resources
417 (e.g., CPUs and/or memory).

418 **2.1.1 Security Limitations of Orchestration Platform**

419 Microservices-based applications require several application services, including security services
420 such as authentication and authorization, as well as the generation of metrics for individual pods
421 (monitoring), consolidated logging (to ascertain causes of failures of certain requests), tracing
422 (sequence of service requests within the application), traffic control, caching, secure ingress,
423 service-to-service (east/west traffic), and egress communication.

424 Taking the example of secure communications between Kubernetes containers, these are not
425 natively secure since there is no easy way to enforce TLS between pods as this would require in
426 maintaining hundreds of TLS certificates. Pods that communicate do not apply identity and
427 access management between themselves. Though there are tools that can be implemented to act
428 as a firewall between pods, such as the [Kubernetes Network Policy](#), this is a [layer 3](#) solution

429 rather than a [layer 7](#) solution, which is what most modern firewalls are. This means that while
430 one can know the source of traffic, one cannot peek into the data packets to understand what they
431 contain. Further, it does not allow for making vital metadata-driven decisions, such as routing on
432 a new version of a pod based on an HTTP header. There are Kubernetes ingress objects that
433 provide a reverse proxy based on layer 7, but they do not offer anything more than simple traffic
434 routing. Kubernetes does offer different ways of deploying pods that do some form of [A/B](#)
435 [testing](#) or canary deployments, but they are done at the connection level and provide no fine-
436 grained control or fast failback. For example, if a developer wants to deploy a new version of a
437 microservice and pass 10 % of traffic through it, they will have to scale the containers to at least
438 10 – nine for the old version and one for the new version. Further, Kubernetes cannot split the
439 traffic intelligently and instead balances loads between pods in a round-robin fashion. Every
440 Kubernetes container within a pod has a separate log, and a custom solution over Kubernetes
441 must be implemented to capture and consolidate them.

442 Although the Kubernetes dashboard offers monitoring features on individual pods and their
443 states, it does not expose metrics that describe how application components interact with each
444 other or how much traffic flows through each of the pods. Consolidated logging is required to
445 determine error conditions that cause an application request or transaction to fail. Tracing is
446 required to trace the sequence of containers that are invoked as part of any application request
447 based on the application logic that underlies a transaction. Since traffic flow cannot be traced
448 through Kubernetes pods out of the box, it is unclear where on the chain the failure for the
449 request occurred.

450 This is where the service mesh software can provide the needed application services and much
451 more.

452 **2.2 Service Mesh Software Architecture**

453 Having looked at the various application services required by microservices-based applications,
454 consider the architecture of service mesh software that provides those services. The service mesh
455 software consists of two main components: the data plane and the control plane.

456 **2.2.1 Data Plane**

457 The data plane component performs three different functions:

- 459 1. Secure networking functions
- 460 2. Policy enforcement functions
- 461 3. Observability functions

462 The primary component of the data plane that performs all three functions listed above is called
463 the sidecar proxy. This layer 7 (L7) proxy runs in the same network namespace (which, in this
464 platform, is the same pod) as the microservice for which it performs proxy functions. There is a
465 proxy for every microservice to ensure that a request from a microservice does not bypass its
466 associated proxy and that each proxy is run as a container in the same pod as the application
467 microservice. Both containers have the same IP address and share the same IP Table rules. That

468 makes the proxy take complete control over the pod and handle all traffic that passes through it
469 [6,7].

470 The first category of functions (secure networking) includes all functions related to the actual
471 routing or communication of messages between microservices. The functions that come under
472 this category are service discovery, establishing a secure (TLS) session, establishing network
473 paths and routing rules for each microservice and its associated requests, authenticating each
474 request (from a service or user), and authorizing the request.

475 With the example of establishing a mutual TLS session, the proxy that initiates the
476 communication session will interact with the module in the control plane of the service mesh to
477 check whether it needs to encrypt traffic through the chain and establish mutual TLS with the
478 backend or target pod. Enabling this functionality using mutual TLS requires every pod to have a
479 certificate (i.e., a valid credential). Since a good-sized microservice application (consisting of
480 many microservices) may require hundreds of pods (even without horizontal scaling of individual
481 microservices through multiple instances), this may involve managing hundreds of short-lived
482 certificates. This, in turn, requires each microservice to have a robust identity and the service
483 mesh to have an access manager, a certificate store, and certificate validation capability. In
484 addition, mechanisms for identifying and authenticating the two communicating pods are
485 required for supporting authentication policies.

486 Other kinds of proxies include ingress proxies [8] that intercept the client calls into the first entry
487 point of application (first microservice that is invoked) and egress proxies that handle a
488 microservice's request to application modules residing outside of the platform cluster.

489 The second category of functions that the data plane performs is enforcement of the policies
490 defined in the control plane through configuration parameters in the proxies (policy enforcement
491 service).

492 The third category of functions that service proxies perform sometimes in collaboration with
493 application service containers are to gather telemetry data, which helps to monitor the health and
494 state of the services, transfer logs associated with a service to the log aggregation module in the
495 control plane, and append necessary data to application request headers to facilitate the tracing of
496 all requests associated with a given application transaction. The application response is conveyed
497 by proxies back to its associated calling service in the form of return codes, a description of
498 response, or the retrieved data.

499 **2.2.2 Control Plane**

500 The control plane has several components. While the data plane of the service mesh mainly
501 consists of proxies running as containers within the same pod as application containers, the
502 control plane components run in their own pods, nodes, and associated clusters. The following
503 are the various functions of the control plane [7] (components in the Istio service mesh that
504 perform these functions are given in parentheses):

505 1. Service discovery and configuration of the Envoy sidecar proxies (Pilot)

- 506 2. Automated key and certificate management (Citadel)
- 507 3. API for policy definition and the gathering of telemetry data (Mixer)
- 508 4. Configuration ingestion for service mesh components (Galley)
- 509 5. Management of an inbound connection to the service mesh (Ingress Gateway)
- 510 6. Management of an outbound connection from the service mesh (Egress Gateway)
- 511 7. Inject sidecar proxies into those pods, nodes, or namespaces where application
- 512 microservice containers are hosted (Sidecar Injector)

513 Overall, the control plane helps the administrator populate the data plane component with the
514 configuration data that is generated from the policies resident in the control plane. The policies
515 for function 3 above may include network routing policies, load balancing policies, policies for
516 blue-green deployments, canary rollouts, timeout, retry, and circuit-breaking capabilities. These
517 last three are collectively called by the special name of resiliency capabilities of the networking
518 infrastructure services. Last but not the least are security-related policies (e.g., authentication and
519 authorization policies, TLS establishment policies, etc.). These policy rules are parsed by a
520 module that converts them into configuration parameters for use by executables in data plane
521 proxies that enforce those policies.

522 The service mesh is container orchestration platform-aware, interacts with the API server that
523 provides a window into application services installed in various platform artifacts (e.g., pods,
524 nodes, namespaces), monitors it for new microservices, and automatically injects sidecar
525 containers into the pods containing these new microservices. Once the service mesh inserts the
526 sidecar proxy containers, operations and security teams can enforce policies on the traffic and
527 help secure and operate the application. These teams can also configure the monitoring of
528 microservices applications without interfering with the functioning of the applications.

529 The provisioning of infrastructure, policy generation, and observability services can be automated
530 using declarative code that is part of DevSecOps pipelines. The development team can
531 concentrate their efforts on efficient development paradigms, such as code modularity and
532 structuring, without worrying about the security and management details of their implementation.

533

534 **3 DevSecOps – Organizational Preparedness, Key Primitives, and** 535 **Implementation**

536 DevSecOps incorporates security into the software engineering process early on. It integrates and
537 automates security processes and tooling into all of the development workflow (or pipeline as later
538 explained) in DevOps so that it is seamless and continuous. In other words, it can be looked upon
539 as a combination of the three processes: Development + Security + Operations [9].

540 This section discusses the following aspects of DevSecOps:

- 541 • Organizational preparedness for DevSecOps
- 542 • Seamless evolution from DevOps to DevSecOps (for organizations that already have a
543 DevOps team in place)
- 544 • Fundamental building blocks or key primitives for DevSecOps

545 **3.1 Organizational Preparedness for DevSecOps**

547 DevSecOps is a software development, deployment, and life cycle management methodology that
548 involves a shift from one large release for an entire application or platform to the Continuous
549 Integration, Continuous Delivery, and Continuous Deployment (CI/CD) approach. This shift, in
550 turn, requires changes in the structure of a company's IT department and its workflow. The most
551 pronounced change involves organizing a DevSecOps group that consists of software developers,
552 security specialists, and IT operations experts for each portion of the application (i.e., the
553 microservice). This smaller team not only promotes efficiency and effectiveness in initial agile
554 development and deployment but also in subsequent life cycle management activities, such as
555 monitoring behavior, developing patches, fixing bugs, or scaling the application. This cross-
556 functional team with expertise in three areas forms a critical success factor for introducing
557 DevSecOps in an organization.

558 **3.2 Seamless Evolution from DevOps to DevSecOps**

559 DevOps is an agile, automated development and deployment process that uses primitives called
560 CI/CD pipelines aided by automated tools to take the software from the build phase to
561 deployment phase. These pipelines are nothing but workflows that take the developer's source
562 code through various stages, such as building, testing, packaging, and deployment supported by
563 automated tools with feedback mechanisms. DevOps is primarily a development/deployment
564 pipeline where security assurance is provided through the use of some basic static application
565 security tools (SAST), dynamic security testing tools (DAST), and software composition analysis
566 (SCA) tools. Thus, a DevOps platform only runs during the build and deployment phases of the
567 software life cycle.
568

569 In a DevSecOps platform, security assurance is provided through built-in design features, such as
570 zero trust, during the build and deployment phases in addition to the use of a comprehensive set
571 of security testing tools, such as DAST and SCA tools. Security assurance is also provided during
572 the runtime/operations phase by continuous behavior detection/prevention tools that use

573 sophisticated techniques, such as artificial intelligence (AI) and machine learning (ML).
574 Therefore, a DevSecOps platform not only runs during build and deployment phases but also
575 during runtime/operations phase. DevSecOps pipelines also run in a production environment as
576 opposed to DevOps, which runs in a separate environment for testing/staging since workflows
577 associated with the latter terminate in the deployment phase.

578 Another distinguishing characteristic of a DevSecOps platform is that the security tools (e.g.,
579 SAST, DAST, and SCAs) are so tightly integrated with IDEs and other DevOps tools that they
580 perform application security analysis, such as identifying vulnerabilities and bugs through
581 efficient scanning in the background. This is often invisible to developers without making them
582 call separate APIs for running these tools [10].

583 In summary, DevSecOps has the following distinguishing features as compared to DevOps:

- 584 • Security testing and the robust incorporation of security controls are integral to all
585 pipelines instead of being relegated to a separate task or phase.
- 586 • The DevSecOps platform operates during runtime (in production), providing real-time
587 assurance of security through correction mechanisms and enabling the certification of
588 continuous authority to operate (C-ATO).

590 **3.3 DevSecOps – Key Primitives and Implementation Tasks**

591

592 The key primitives and implementation tasks involved are:

- 593 • Concept of pipelines and the CI/CD pipeline
- 594 • Building blocks for the CI/CD pipeline
- 595 • Designing and executing the CI/CD pipeline
- 596 • Strategies for automation
- 597 • Requirements for automation tools in the CI/CD pipeline

598

599 **3.3.1 Concept of Pipelines and the CI/CD pipeline**

600 DevSecOps is a methodology or framework for agile application development, deployment, and
601 operations for cloud-native applications [11] and is made up of stages just like any other
602 methodology. The sequence and flow of information through the stages is called workflow,
603 where some stages can be executed in parallel while some others have to follow a sequence. Each
604 stage may require the invocation of a unique job to execute the activities in that stage.

605 A unique concept that DevSecOps introduces in the process workflow is the concept of
606 “pipelines” [12]. With pipelines, there is no need to individually write jobs for each stage of the
607 process. Instead, there is only one job that starts from the initial stage, automatically triggers the
608 activities pertaining to other stages (both sequential and parallel), and creates an error-free smart
609 workflow.

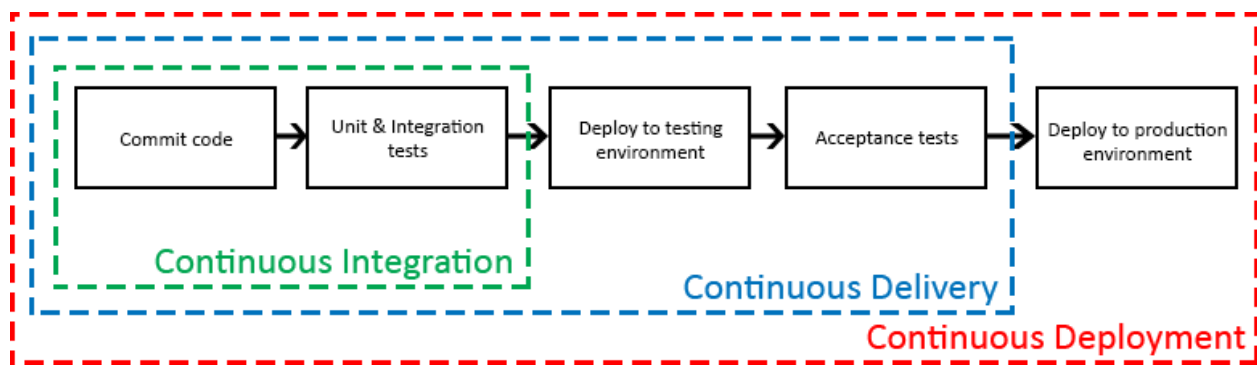
610 The pipeline in DevSecOps is called the CI/CD pipeline based on the overall tasks it

611 accomplishes and the two individual stages it contains. CI stands for the continuous integration
 612 stage. CD can denote either the continuous delivery or continuous deployment stage. Depending
 613 on this latter stage, CI/CD can involve the following tasks:

- 614 • Build, Test, Secure, and Deliver (the tested modified code is delivered to the staging area)
- 615 • Build, Test, Secure, Deliver, and Deploy (the code in the stage area is automatically
 616 deployed)

617
 618 In the former, automation ends at the delivery stage, and the next task of deployment of the
 619 modified application in the hosting platform infrastructure is performed manually. In the latter,
 620 the deployment is also automated. Automation of any stage in the pipeline is enabled by tools that
 621 express the pipeline stage as code.

622 The workflow process for a CI/CD pipeline is depicted in Figure 1 below:



623

624

Figure 1: CI/CD pipeline workflow [13]

625

Figure 3.1. CI/CD pipeline workflow [13]

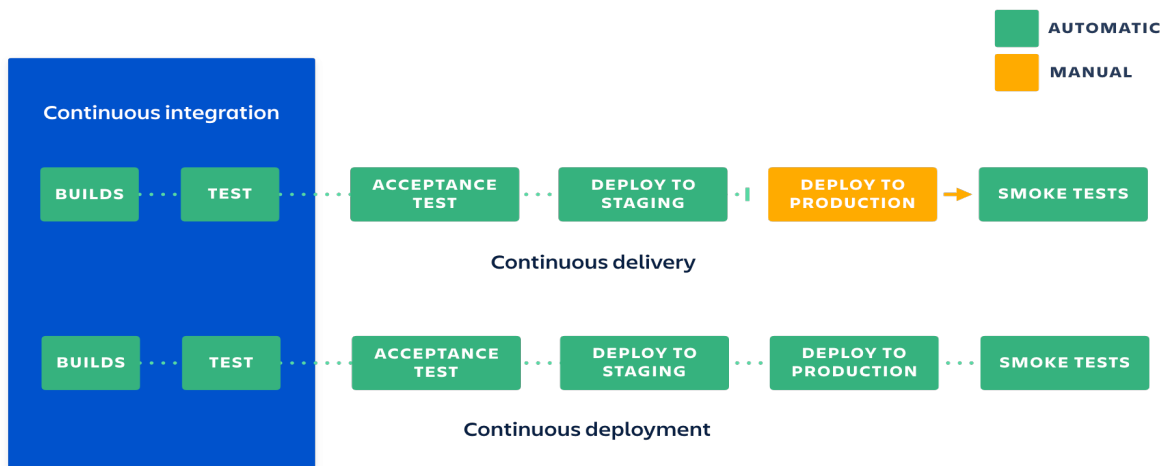
626 The unit and integration tests shown in the diagram use SAST, DAST, and SCA tools described
 627 in Section 3.2. These tests and associated tools are common to DevOps and DevSecOps.

628 Continuous integration involves developers frequently merging code changes into a central
 629 repository where automated builds and tests run. Build is the process of converting the source
 630 code to executable code for the platform on which it is intended to run. In the CI/CD pipeline
 631 software, the developer’s changes are validated by creating a build and running automated tests
 632 against the build. This process avoids the integration challenges that can happen when waiting for
 633 release day to merge changes into the release branch [14].

634 Continuous delivery is the next stage after continuous integration where code changes are
 635 deployed to a testing and/or staging environment after the build stage. Continuous delivery to a
 636 production environment involves the designation of a release frequency – daily, weekly,
 637 fortnightly, or some other period – based on the nature of the software or the market in which the
 638 organization operates. This means that on top of automated testing, there is a scheduled release
 639 process, though the application can be deployed at any time by clicking a button. The deployment
 640 process in continuous delivery is characterized as manual, but tasks such as the migration of code

641 to a production server, the establishment of networking parameters, and the specification of
 642 runtime configuration data may be performed by automated scripts.

643 Continuous deployment is similar to continuous delivery except that releases happen
 644 automatically [14], and changes to code are available to customers immediately after they are
 645 made. The distinction between continuous delivery and continuous deployment is shown in
 646 Figure 2 below.



647

648 **Figure 2 - Distinction between Continuous Delivery and Continuous Deployment [14]**

649 **3.3.2 Building Blocks for CI/CD pipelines**

650 The primary software for defining CI/CD pipeline resources, building the pipelines, and
 651 executing those pipelines is the CI/CD pipeline software. There may be slight variations in the
 652 architecture of this class of software depending on the particular offering. The following is an
 653 overview of the landscape in which CI/CD tools (pipeline software) operate:

- 654 • Some CI/CD tools natively operate on the platform on which the application and the
 655 associated resources needed are hosted (i.e., container orchestration and resource
 656 management platform, such as Kubernetes), while others need to be integrated into the
 657 application hosting platform through its API. An example of a former class of tool is
 658 Tekton, and examples for the latter class include Jenkins, Travis, Bamboo, and CircleCI.
 659 Some advantages of using the CI/CD tools that are native to the application hosting
 660 platform are:
 - 661 (a) It makes it easier to deploy, maintain, and manage the CI/CD tool itself.
 - 662 (b) Every pipeline defined by the CI/CD tool becomes another platform-native
 663 resource and is managed the same way. In fact, all entities required for executing
 664 pipelines, such as Tasks and Pipelines (which then act as blueprints for other
 665 entities, such as TaskRuns and PipelineRuns, respectively), can be created as

666 CustomResourceDefinitions (CRDs) built on top of resources native to the platform
667 [15]. Software with this type of architecture may be used by other CI/CD pipeline
668 software offerings to facilitate faster defining of pipelines.

669 • Some CI/CD tool offerings (e.g., GitLab, Jenkins X) integrate with a Git Repository – one
670 for each application and for each environment (staging/production). All changes in
671 application modules, infrastructure, or configuration are made and stored in these Git
672 repositories. The CI/CD pipeline software is connected to the Git repositories through Git
673 webhooks and gets activated on commits (push workflow model) or pull requests in these
674 repositories.

675 • Some CI/CD tools perform CD functions alone for the native platform (e.g., Jenkins X for
676 Kubernetes platform) or for multiple technology stacks (e.g., Spinnaker for multi-cloud
677 deployment). The difficulty with some in this class of tools is that they may lack native
678 tools for completing the CI functions (e.g., tools to test code, build application images, or
679 push them to registry).

680 3.3.3 Designing and Executing the CI/CD pipeline

681 The purpose of creating the CI/CD pipeline is to enable frequent updates to source code, rebuilds,
682 and the automatic deployment of updated modules into the production environment. The key
683 tasks involved are [16]:

- 684 • Setting up the source code repository: Set up a repository (e.g., GitHub or GitLab) for
685 storing application source code with proper version control.
- 686 • Build process: Configure and execute the build process for generating the executables (for
687 those portions of the code that need to be updated) using an automated code build tool.
- 688 • Securing the process: Ensure that the build is free of static and dynamic vulnerabilities
689 through unit testing with SAST and DAST tools.
- 690 • Creating the deployment environment: Create the environment to deploy the application
691 using an automated deployment tool.
- 692 • Creating the delivery pipeline: Create a pipeline that will automatically build and deploy
693 the application using an automated pipeline tool.
- 694 • Test the code and execute the pipeline: Test the code in the pipeline using an automated
695 testing tool, and execute the pipeline to complete the deployment of updated code.

696 To reiterate, the three primary stages of the CI/CD process are build/test, ship/package, and
697 deploy. The following features transform this into a pipeline:

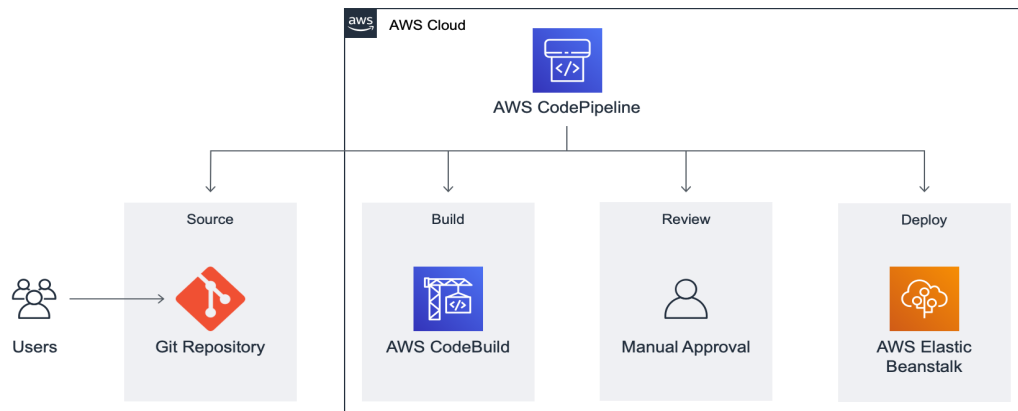
- 698 • When an update is made to the source code for a service, the code changes pushed to the
699 source code repository trigger the code building tool.
- 700 • The code building tool, which is usually an IDE, is often integrated with security testing
701 tools (e.g., static vulnerability analysis tool) to facilitate the generation of secure compiled
702 code artifacts, thus integrating security into the CI pipeline.

- 703 • The generation of compiled code artifacts in code building tools triggers the
704 shipping/package tool, which may be integrated with its own set of tools (e.g., dynamic
705 vulnerability analysis or dynamic penetration testing tools, software composition analysis
706 tools for identifying vulnerabilities in the attached libraries) and also creates the
707 configuration parameters relevant to the deployment environment.
- 708 • The output of the shipping/packaging tool is then automatically fed to the CD tool, which
709 then deploys the package into the desired environment (e.g., staging, production etc.) [17].

710 The workflow of the CI/CD pipeline should not create the impression that there is no human
711 element involved. The following teams provide input into the CI/CD pipeline [18]:

- 712 • Development team incorporates new features into the application
- 713 • Build or integration team creates new CI pipelines and introduces new CI tools and testing
714 tools
- 715 • Security team conducts audits and patching
- 716 • Infrastructure team creates, maintains, and upgrades the infrastructure
- 717 • QA team develops integration test cases
- 718 • Deployment team/release team creates pipelines and packages for various environments
719 (UAT/PreProd/Prod) and performs configuration and provisioning appropriate for these
720 environments

721 Some of the many activities performed by these teams include the customization, update, and
722 enhancement of the tools employed in the CI/CD pipeline (e.g., updating the static vulnerability
723 analysis tool with the latest database of known vulnerabilities etc.). Caution should be exercised
724 during manual operations so that they do not block pipelines. Targets for mean time to production
725 should be set up while also mitigating risks, such as insider threats, through the use of “merge
726 requests” and multiple approvers in merge requests.



727

728

Figure 2 - Examples of Tools Involved in CI/CD Pipeline Tasks

729 This pipeline is designed, maintained, and executed by the post-release team who – in addition to
730 monitoring functions – performs other processes, such as compliance management, backup
731 processes, and asset tracking [19].
732

733 **3.3.4 Strategies for Automation**

734 Compared to other models of software development, which involve a linear progression from
735 coding to release, DevOps uses a forward process with a delivery pipeline (i.e., build/secure,
736 ship/package, and release) and a reverse process with a feedback loop (i.e., plan and monitor) that
737 form a recursive workflow. The role of automation in these activities is to improve this workflow.
738 Continuous integration emphasizes testing automation to ensure that the application is not broken
739 whenever new commits are integrated into the main branch. Automation results in the following
740 benefits:

- 741 • Generation of data regarding software static and runtime flows
- 742 • Reduction of development and deployment times
- 743 • Better utilization of the team’s expertise by assigning routine tasks to the tools

744 The following strategies are recommended [22] for automation so as to facilitate better utilization
745 of organizational resources and derive the greatest benefit in terms of an efficient, secure
746 application environment.

747 Choice of Activities to Automate: For example, the following are productive candidates for the
748 automation of testing activities.

- 749 • Testing of modules whose functions are subject to regulatory compliance (e.g., PCI-DSS,
750 HIPAA, Sarbanes-Oxley)
- 751 • Tasks that are repetitive with moderate to high frequency
- 752 • Testing of modules that perform time-sequenced operations, such as message publishers
753 and message subscribers
- 754 • Testing of workflows (e.g., request tracing) involving transactions that span multiple
755 services
- 756 • Testing of services that are resource-intensive and likely to constrict operations

757 After choosing the candidates for automation based on the above criteria, the usual risk analysis
758 must be applied to choose a subset that provides an optimum return on investment and maximizes
759 desirable security metrics (e.g., defense in depth). Some recommended strategies include:

- 760 • Using the cost-benefit ratio in hours saved per year to prioritize which processes to
761 automate [22]
- 762 • Using key performance indicators (KPI) (e.g., mean time to identify faults or problems,
763 rectify, or recover) as markers to refine the DevSecOps processes [22]
- 764 • Based on the application, applying different weights to infrastructure services (e.g.,
765 authorization and other policies enforcement, monitoring of system states to ensure secure
766 runtime states, network resilience in terms of system availability, latency, etc.) to
767 determine the allocation of resources to DevSecOps processes

768 3.3.5 Requirements for Automation Tools in CI/CD Pipelines

769 The main utility of a performance testing tool is to identify the root cause of an issue, such as
770 application response time or a bottleneck operation. The desirable features of this class of tool
771 are:

- 772 • Easy integration into developers' CI/CD pipelines, which enables this class of tools to be
773 an integral part of the DevSecOps process rather than one managed by a central platform
774 engineering or site reliability engineering team that could become hamper deployments
- 775 • An accessible interface that can be used by all participants involved in the software
776 release workflow and have features that analysts/engineers can use to perform insightful
777 analyses into specific datasets
- 778 • Coverage for a wide variety of back-end infrastructure components in multiple clouds as
779 well as on premises; this abstraction feature is now natively present in container
780 orchestration and resource management platforms (e.g., Kubernetes) or through stacks
781 (e.g., vSphere or OpenStack)

782 The security automation tools for various functions (e.g., static vulnerability analysis, dynamic
783 vulnerability analysis, software composition analysis) used in CI/CD pipelines need to have
784 different interface and alerting/reporting requirements since they have to operate seamlessly
785 depending on the pipeline stage (e.g., build, package, release) at which they are used. These
786 requirements are:

- 787 • Security automation tools should work with integrated development environment (IDE)
788 tools and help developers prioritize and remediate static vulnerabilities. These capabilities
789 are needed to facilitate developer adoption and improve productivity.
- 790 • Security automation tools should be flexible to support specific workflows and provide
791 scaling capabilities for security services.
- 792 • Tools that perform static vulnerability checks at the build phase ensure safe data flows,
793 and those that perform dynamic vulnerability checks ensure safe application states during
794 runtime.

795 **4 Implementing DevSecOps Primitives for the Reference Platform**

796 Various CI/CD pipelines are involved in the reference platform (i.e., microservices-based
797 application with service mesh that provides infrastructure services). Though the reference
798 application is a microservices-based application, the DevSecOps primitives can be applied to
799 monolithic applications as well as applications that are both on-premises and cloud-based (hybrid
800 cloud, single public cloud, and multi-cloud).

801
802 The application architecture consists of other functional elements (in addition to elements for
803 housing executable application code and providing application services) such as for infrastructure
804 resources, runtime policies, and continuous monitoring of the health of the application, which can
805 all be deployed through declarative codes. Hence, separate CI/CD pipelines can be created for all
806 of these code types as well. These five code types will first be considered in the context of the
807 reference platform followed by separate sections that will describe the associated CI/CD
808 pipelines.

- 809 1. Code types in the reference platform and associated CI/CD pipelines (Section 4.1)
- 810 2. CI/CD pipeline for application code and application services code (Section 4.2)
- 811 3. CI/CD pipeline for infrastructure as code (IaC) (Section 4.3)
- 812 4. CI/CD pipeline for policy as code (Section 4.4)
- 813 5. CI/CD pipeline for observability as code (Section 4.5)

814
815 Implementation issues for all CI/CD pipelines irrespective of code types will be addressed in the
816 following sections:

- 817 • Securing the CI/CD pipelines (Section 4.6)
- 818 • Workflow models in the CI/CD pipelines (Section 4.7)
- 819 • Security testing in the CI/CD pipelines (Section 4.8)

820
821 This section will also consider the overall benefits of DevSecOps with a subsection on specific
822 advantages in the context of the reference platform and the ability to leverage DevSecOps for
823 continuous authorization to operate (C-ATO) in Sections 4.9 and 4.10, respectively.

824 **4.1 Code Types in Reference Platform and Associated CI/CD Pipelines**

825 The constituent components of the code types in the reference platform are:

- 826 1. Microservices-based application component (implemented as a series of containers),
827 which contains the application logic (e.g., interacting with data, performing transactions
828 etc.) and the application code
- 829 2. Infrastructure component (containing computer, networking, and storage resources)
830 whose constituents can be provisioned using infrastructure as code
- 831 3. Service mesh component (implemented through a combination of control plane modules
832 and service proxies), which provides application services, enforces policies (e.g.,
833 authentication and authorization), and contains application services code and policy as
834 code
- 835 4. Monitoring component (modules involved in ascertaining the parameters that indicate the
836 health of the application), which performs functions (e.g., log aggregation, the generation

837 of metrics, the generation of displays for dashboard, etc.) and contains observability as
838 code

839 The DevSecOps methodology for developing, testing, delivering, and deploying these five types
840 of codes (i.e., application, application services, infrastructure, policy, and monitoring) requires
841 the corresponding pipelines:

- 842 • CI/CD pipelines for application code and application services code – the former contains
843 the data and application logic for a specific set of business transactions while the latter
844 contains code for all services such as network connections, load balancing, network
845 resilience etc.
- 846 • CI/CD pipeline for infrastructure as code (IaC) – The process of writing the code for
847 provisioning and configuring the steps of infrastructure resources to automate application
848 deployment in a repeatable and consistent manner [23]. This code is written in a
849 declarative language and – when executed – provisions, de-provisions, and configures the
850 infrastructure for the application that is being deployed. This type of code is like any other
851 code found in an application’s microservice except that it provides an infrastructure
852 service (e.g., provisioning a server) rather than a transaction service (e.g., payment
853 processing for an online retail application).
- 854 • CI/CD pipeline for policy as code – Describes many policies, including security policies,
855 as executable modules [24]. One example is the authorization policy, the code for which
856 contains verbs or artifacts specific to the policy (e.g., allow, deny etc.) and to the domain
857 where it applies (e.g., REST API with verbs such as method [GET, PUT, etc.], path, etc.).
858 This code can be written in a special-purpose policy language, such as Rego or WASM,
859 or languages used in regular applications, such as Go. This code may have some overlap
860 with the configuration code of IaC. However, for implementing policies associated with
861 critical security services that are specific to the application domain, a separate policy as
862 code that resides in the policy enforcement points (PEPs) of the reference platform is
863 required.
- 864 • CI/CD pipeline for observability as code – The ability to infer a system’s internal state
865 and provide actionable insights into when and, more importantly, why errors occur within
866 a system. It is a full-stack observability that includes monitoring and analytics, as well as
867 offers key insights into the overall performance of applications and the systems hosting
868 them. In the context of the reference platform, observability as code is the portion of the
869 code that creates agencies in proxies and creates functionality for gathering three types of
870 data (i.e., logs, traces, and telemetry) from microservices applications [25]. This type of
871 code also supplies or transfers data to the external tools (e.g., log aggregation tool that
872 aggregates log data from individual microservices, provides analysis of tracing data for
873 bottleneck services, generates metrics that reflect the application health from telemetry
874 data, etc.). Brief descriptions of the three functions enabled by observability as code are:
 - 875 1. **Logging** captures detailed error messages, as well as debugs logs and stack
876 traces for troubleshooting.

- 877 2. **Tracing** follows application requests as they wind through multiple
878 microservices to complete a transaction in order to identify an issue or
879 performance bottleneck in a distributed or microservices-based ecosystem.
880 3. **Monitoring, or metrics**, gathers [telemetry](#) data from applications and
881 services.

882 The policy and observability code types span the following components of the service mesh.

- 883 • **Proxies (ingress, sidecar, and egress)**: These house encoding policies related to session
884 establishment, routing, authentication, and authorization functions.
- 885 • **Control plane of the service mesh**: This houses code for relaying telemetry information
886 from services captured and sent by proxies to specialized monitoring tools, authentication
887 certificate generation and maintenance, updating policies in the proxies, monitoring
888 overall configuration in the service orchestration platform for generating new proxies, and
889 deleting obsolete proxies associated with discontinued microservices.
- 890 • **External modules**: These house modules that perform specialized functions at the
891 application and enterprise levels (e.g., such as the centralized authorization or entitlement
892 server, the centralized logger, monitoring/alerting server status through dashboards, etc.)
893 and build a comprehensive view of the application status. These modules are called by
894 code from the proxies or the control plane.

895 **4.2 CI/CD Pipeline for Application Code and Application Services Code**

896 Application code and application services code reside in the container orchestration and resource
897 management platform, and the CI/CD software that implements the workflows associated with it
898 usually reside in the same platform. This pipeline should be protected using the steps outlined in
899 Section 4.6, and the application code under the control of this pipeline should be subject to the
900 security testing described in Section 4.8. Additionally, the orchestration platform on which the
901 application resides should itself be protected using a runtime security tool (e.g., Falco) [32] that
902 can read OS kernel logs, container logs, and platform logs in real-time and process them against a
903 threat detection rules engine to alert users of malicious behavior (e.g., creation of a privileged
904 container, reading of a sensitive file by an unauthorized user, etc.). They usually come with a set
905 of default (predefined) rules over which custom rules can be added. Installing them on the
906 platform spins up agents for each node in the cluster, which can monitor the containers running in
907 the various pods of that node. The advantage of this type of tool is that it complements the existing
908 platform's native security measures, such as access control models and pod security policies, that
909 prevent violations of security by actually detecting them when they occur [32].

910 **4.3. CI/CD Pipeline for Infrastructure as Code**

911 Infrastructure as code (IaC) involves codifying all software deployment tasks (allocation of type
912 of servers, such as bare metal, VMs or containers, resource content of servers) and the
913 configuration of these servers and their networks. The software containing this code type is also
914 called a resource manager or deployment manager. In other words, IaC software automates
915 management of the whole IT infrastructure life cycle (provisioning and de-provisioning of
916 resources) and enables a programmable infrastructure. The integration of this software as part of

917 the CI/CD pipeline not only results in agile deployment and maintenance but also in a robust
918 application platform that is secure and meets performance needs.

919
920 The conventional approach to allocating infrastructure for applications consists of initially
921 provisioning compute and networking resources with configuration parameters and ongoing tasks
922 such as patch management (e.g., OS and libraries), establishing conformity to compliance
923 regulations (e.g., data privacy), and making drift correction (where the current configuration no
924 longer provides the intended operational state).

925
926 Infrastructure as code (IaC) is a declarative style of code that encodes computer instructions that
927 encapsulate the parameters necessary to deploy virtual infrastructure on a public cloud service or
928 private data center via a service's management APIs [33]. Depending on the particular IaC tool,
929 this language can either be a scripting language (e.g., Go, JavaScript, Python, TypeScript, etc.) or
930 a proprietary configuration language (e.g., HCL) that may or may not be compatible with
931 standardized languages (e.g., JSON). The basic unit of these instructions is called "configuration"
932 and tells the system how to provision and manage infrastructure (whether that is an individual
933 compute instance or a complete server, such as physical servers or virtual machines), containers,
934 storage, network connections, connection topology, and load balancers. [34]. In some cases, the
935 infrastructure may be short-lived or ephemeral, and the lifespan of the infrastructure (whether
936 immutable or mutable) does not warrant continued configuration management. Provisioning
937 could be tied to individual commits of application code using tools that can connect application
938 code and infrastructure code in way that is logical, expressive, and familiar to development and
939 operations teams, where application code increasingly defines the infrastructure resource
940 requirements for a cloud application [35].

941 942 **4.3.1 Comparison of Configuration and Infrastructure**

943
944 Infrastructure is often confused with configuration [34], which maintains computer systems,
945 software, dependencies, and settings in a desired, consistent state. For example, putting a newly
946 purchased server onto a rack and connecting it to the switches so that it is connected to the existing
947 networks (or launching a new virtual machine and assigning network interfaces to it) belongs to
948 the definition of "infrastructure." In contrast, after the server is launched, installing an HTTP
949 server and configuring it belongs to configuration management. In physical data centers, specific
950 teams purchase servers, install servers, and connect networking cables with the underlying
951 infrastructure in mind.

952 953 **4.4 CI/CD Pipeline for Policy as Code**

954 Policy as code involves codifying all policies and running them as part of the CI/CD pipeline so
955 that they become an integral part of the application runtime. Examples of policy categories
956 include authorization policies, networking policies, and implementation artifact policies (e.g.,
957 container policies). Policy management capabilities in a typical "policy as code software" may
958 come with a set of predefined policy categories and policies and also support the definition of
959 new policy categories and associated policies by providing policy templates [36].

960 Some examples of policy categories and associated policies are given in

961 Table 1 below.

962 **Table 1: Policy Categories and Example Policies**

963

Policy Category	Example Policies
Networking policies and zero trust policies	<ul style="list-style-type: none"> • Blocking designated ports • Designating ingress host names
Implementation artifact policies (e.g., container policies)	<ul style="list-style-type: none"> • Ensuring that containers do not run as root • Blocking privilege escalation for containers
Storage policies	<ul style="list-style-type: none"> • Set persistent volume sizes • Set persistent volume reclaim policies
Access control policies	<ul style="list-style-type: none"> • Ensure that policies cover all data objects • Ensure that there are no policy conflicts
Supply chain policies	<ul style="list-style-type: none"> • Allow only approved container registries • Allow only certified libraries

964 The policies defined in the “policy as code software” may translate into the following in the
 965 application infrastructure:

- 966 • Runtime configuration parameters
- 967 • Policy-enforcing executable (e.g., WASM in service proxies)
- 968 • Triggers for calling an external policy decision module (e.g., calling an external
 969 authorization server for an allow/deny decision based on the evaluation of access control
 970 policies relevant to the current access request)

971 **4.5 CI/CD Pipeline for Observability as Code**

972 Observability as code deploys a monitoring agent in each of the application’s service components
 973 to collect the three types of data (described in Section 4.1), send them to specialized tools that
 974 correlate them, perform analysis, and display the analyzed consolidated data on dashboards to
 975 present an overall application-level picture [20]. An example of such consolidated data are log
 976 patterns, which provide a view of log data that is presented after the log data are filtered using
 977 some criterion (e.g., a service or an event). The data are grouped into clusters based on common
 978 patterns (e.g., based on timestamp or range of IP addresses) for easy interpretation. Unusual
 979 occurrences are identified, and those findings can then be used to steer and accelerate further
 980 investigation [21].

981 **4.6 Securing the CI/CD Pipeline**

982 There are some common implementation issues to be addressed for CI/CD pipelines irrespective
 983 of code type. The first – security – involves securing the artifacts, including hardening the servers

984 and nodes used for generating the build and establishing the authenticity of the build components.
985 Securing the processes involves the assignment of roles for operating the build tasks. Automation
986 tools (e.g., Git Secrets) are available for this purpose. The following security tasks should be
987 performed when securing the CI pipeline [26]:

- 988 • Authentication and validation in code and build repositories
- 989 • Logging all activities associated with code and build updates
- 990 • Sending build reports to developers and stopping everything if a build fails
- 991 • Sending build reports to security and stopping everything if the audit or validation fails

992 Securing the CD pipeline involves the following:

- 993 • Multi-party signing with TUF and in-toto to sign each CD pipeline phase
- 994 • Signing the final artifacts with multiple phase keys to ensure that no one bypasses the
995 pipeline

996 **4.7 Workflow Models in CI/CD Pipelines**

997 The next common issue involves workflow models. All CI/CD pipelines can have two types of
998 workflow models, which depend on the automated tools that are deployed as part of the pipeline.

- 999 1. Push-based model
- 1000 2. Pull-based model

1001
1002 In the CI/CD tools that support the push-based model, changes made in one stage or phase of the
1003 pipeline trigger changes in the subsequent stages and phases. For example, through a series of
1004 encoded scripts, the new builds in the CI system trigger changes to the CD portion of the pipeline
1005 and thus change the deployment infrastructure (e.g., Kubernetes cluster). The security downside
1006 of using the CI system as the basis for change in deployments is the possibility of exposing
1007 credentials outside of the deployment environment in spite of best efforts to secure the CI scripts,
1008 which operate outside of the trusted domain of the deployment infrastructure. Since CD tools
1009 have the keys to production systems, push-based models are rendered insecure. A pull-based
1010 model, which typically uses a GitOps repository for storing the source code and builds, is
1011 therefore highly recommended.

1012 In a pull-based workflow model, an operator that pertains to the deployment environment (e.g.,
1013 Kubernetes Operator, Flux, ArgoCD) pulls new images from inside of the environment as soon as
1014 the operator observes that a new image has been pushed to the registry. The new image is pulled
1015 from the registry, the deployment manifest is automatically updated, and the new image is
1016 deployed in the environment (e.g., cluster). Thus, the convergence of the actual deployment
1017 infrastructure state with the state declaratively described in the Git deployment repository is
1018 achieved. Additionally, the deployment environment credentials (e.g., cluster credentials) are not
1019 exposed outside of the production environment.

1020 **4.7.1 GitsOps Workflow Model for CI/CD – A Pull-based Model**

1021 The GitOps workflow model is an improvement on the CI/CD pipeline (for the delivery portion
1022 of the pipeline), which uses a pull-based workflow model instead of the push-based model

1023 supported by many CI/CD tools. In this model, the CI portion of the pipeline is unchanged since
1024 the CI engine (e.g., Jenkins, GitLab CI) is still used for creating builds for the changed code,
1025 regression testing, and integrating/merging with the main source code in the relevant repositories,
1026 though it is not used to trigger continuous delivery (push updates directly) in the pipeline.
1027 Instead, a separate GitOps operator manages the deployment based on updates to the main (trunk)
1028 code.

1029 An operator (eg. Flux, ArgoCD) is an actor managed by an orchestration platform and can
1030 inherit the cluster's configuration, security, and availability. The use of this actor improves
1031 security because an agent that lives inside of the cluster listens for updates to all code and image
1032 repositories that it is allowed to access and pulls images and configuration updates into the
1033 cluster. The pull approach used by the agent has the following security features:

- 1034 • ONLY carry out operations permitted by authorization policies defined in the
1035 orchestration platform; trust is shared with the cluster and not managed separately
- 1036 • Bind natively to all orchestration platform objects, and know whether operations have
1037 completed or need to be retried

1038 **4.8 Security Testing – Common Requirement for CI/CD Pipelines for All Code Types**

1039 The last common issue is security testing. Whatever the code type is, the CI/CD pipelines of
1040 DevSecOps for microservices-based infrastructure with service mesh should include application
1041 security testing (AST) enabled by either automated tools or offered as a service. These tools
1042 analyze and test applications for security vulnerabilities. According to Gartner, there are four
1043 main AST technologies [27]:

- 1044 1. Static AST (SAST) tools – analyze an application's source, bytecode, or binary code for
1045 security vulnerabilities, typically at the programming and/or testing software life cycle
1046 (SLC) phases. Specifically, this technology involves techniques that look through the
1047 application in a commit and analyze its dependencies [28]. If any dependencies contain
1048 issues or known security vulnerabilities, a commit will be marked as insecure and will not
1049 be allowed to proceed to deployment.
- 1050 2. Dynamic AST (DAST) tools – analyze applications in their dynamic, running state during
1051 testing or operational phases. They simulate attacks against an application (typically web-
1052 enabled applications, services, and APIs), analyze the application's reactions, and
1053 determine whether it is vulnerable. In particular, DAST tools go one step further than
1054 SAST and spin up a copy of the production environment inside the CI job in order to scan
1055 the resulting containers and executables [28]. The dynamic aspect helps the system catch
1056 dependencies that are being loaded at launch time, such as those would not be caught by
1057 SAST.
- 1058 3. Interactive AST (IAST) tools – combine elements of DAST with the instrumentation of
1059 the application under test. They are typically implemented as an agent within the test
1060 runtime environment (e.g., instrumenting the Java Virtual Machine [JVM] or .NET CLR)
1061 that observes operations or identifies and attacks vulnerabilities.

1062 4. Software composition analysis (SCA) tools – are used to identify open-source and third-
1063 party components in use in an application, their known security vulnerabilities, and
1064 typically adversarial license restrictions.

1065 4.8.1 Functional and Coverage Requirements for AST tools

1066 In general, the overall metrics that testing tools (including the specific class of AST tools) should
1067 satisfy are [29]:

- 1068 • Increase the quality of application releases
- 1069 • Integrate with the tools that developers are already using
- 1070 • Be as few test tools as possible but provide the necessary coverage risk
- 1071 • Lower-level unit tests at the API and microservices level should have sufficient visibility
1072 to determine coverage
- 1073 • Higher-level UI/UX and system tests
- 1074 • Have deep code analysis capabilities to detect runtime flaws
- 1075 • Increase the speed at which the releases can be done
- 1076 • Be cost-efficient

1077 The functional requirements for AST tools in particular are that perform the following types of
1078 scans [30]:

- 1079 • **Vulnerability scans:** Probe applications for security weaknesses that could expose them
1080 to attacks
- 1081 • **Container image scans:** Analyze the contents and build process of a container image in
1082 order to detect security issues, vulnerabilities, or deficient practices
- 1083 • **Regulatory/compliance scans:** Assess adherence to specific compliance requirements

1084 The vulnerability scans are to be performed whenever the code in GitHub is revised to ensure that
1085 the current revision does not contain any vulnerable dependencies [31].

1086 The desirable features of AST tools and/or services, along with techniques for behavioral
1087 analysis, are [27]:

- 1088 • Analyze source, byte, or binary code
- 1089 • Observe the behavior of apps to identify coding, design, packaging, deployment, and
1090 runtime conditions that introduce security vulnerabilities

1091 Scanning *application code* for security vulnerabilities and misconfiguration as part of CI/CD
1092 pipeline tasks should involve the following artifacts:

- 1093 • Container images should be scanned for vulnerabilities.
- 1094 • Filesystems should be scanned for both vulnerabilities and misconfigurations.
- 1095 • Git repositories should be scanned for both vulnerabilities and misconfigurations.

1096 Container images include OS packages (e.g., Alpine, UBI, RHEL, CentOS, etc.) and language-
1097 specific packages (e.g., Bundler, Composer, npm, yarn, etc.)

1098 Scanning *Infrastructure-as-code* for security vulnerabilities reduces the operations workload by
 1099 preventing those vulnerabilities from making it to production, although it cannot replace runtime
 1100 security since there will always be the risk of drift. The infrastructure-as-code files can be found
 1101 in the following:

- 1102 • In the container orchestration platform itself to facilitate deployments (e.g., Kubernetes
 1103 YAML infrastructure-as-code files)
- 1104 • In the dedicated infrastructure-as-code files found as part of CI/CD pipeline software
 1105 (e.g., HashiCorp Terraform infrastructure-as-code files, AWS CloudFormation
 1106 infrastructure-as-code files)

1107 Both *application services code*, *policy-as-code*, and *observability-as-code files* can be found in
 1108 the data plane and control plane components of the dedicated application services component
 1109 (e.g., service mesh) and should be scanned for both security vulnerabilities (e.g., information
 1110 leakage in authorization policies) and misconfiguration.

1111 **4.9 Benefits of DevSecOps Primitives to Application Security in the Service Mesh**

1112 The benefits of DevSecOps include [37]:

- 1113 • Streamlined software delivery and deployment processes
- 1114 • Reduction of the number of silos between IT groups
- 1115 • Reduction of attack surfaces by implementing zero trust down to the container level
- 1116 • Introduction of moving target defenses by threatening containers as cattle versus pet and
 1117 going back to an immutable state
- 1118 • Reduction of lateral movement capabilities by leveraging zero trust and continuous
 1119 monitoring with modern behavior prevention capabilities
- 1120 • Increased communications between team members and other stakeholders
- 1121 • Faster incorporation of feedback, resulting in quicker software improvements
- 1122 • Automation of repetitive, manual tasks, leading to increased efficiency
- 1123 • Less downtime and faster time to market
- 1124 • Infrastructure for continuous software development and delivery (as well as admission
 1125 controller and the use of OPA to add guard rails to the runtime)

1126 As already stated, the primary goal of DevSecOps is to have a platform that ensures runtime
 1127 security of microservices-based applications. Runtime security is assured through the following:

- 1128 (a) Validation of every request: Every request from a user or client application (service) is
 1129 authenticated and authorized (using mechanisms such as Open Policy Agent [OPA] or any
 1130 external authorization engine), and admission controllers that are integral parts of the
 1131 platform. While authorization engines provide application domain-specific policy
 1132 enforcement, admission controllers provide platform-specific policy relating to end-point
 1133 objects of a specific platform (e.g., pods, deployment, namespaces). Specifically,
 1134 admission controllers mutate and validate. Mutating admission controllers parse each
 1135 request and make changes to the request (mutate) before forwarding it down the chain. An
 1136 example is setting default values for specifications that are not set by a user in the request
 1137 so as to ensure that workloads running on the cluster are uniform and follow a particular

1138 standard defined by the cluster administrator. Another example is adding a specific
1139 resource limit for the pod (if the resource limits are not set for that pod) and then
1140 forwarding it down the chain (mutate the request by adding this field if it is not present in
1141 the request). By doing so, all of the pods in the cluster will always have a resource limit set
1142 according to a specification unless explicitly stated. Validating admission controllers reject
1143 requests that do not follow a particular specification. For example, none of the pod requests
1144 can have security context set to run as root user [38].

1145 (b) Feedback mechanisms:

- 1146 • Provide feedback loops to the application hosting platform (e.g., a notification to kill a
1147 pod that contains a malicious container)
- 1148 • Provide proactive dynamic security by monitoring application configuration (e.g.,
1149 monitoring new pods/containers introduced into the application and generating and
1150 injecting proxies to take care of their secure communication needs)
- 1151 • Enable several security assertions regarding the application: non-bypassable (policies
1152 always enforced under all usage scenarios), trusted and untrusted portions of the
1153 overall application code, absence of credential and privilege leaks, trusted
1154 communication paths, and secure state transitions
- 1155 • Enable assertions regarding performance parameters (e.g., network resilience
1156 parameters, such as continued operation under failures, redundancy and recoverability
1157 features)

1158 **4.10 Leveraging DevSecOps for Continuous Authorization to Operate (c-ATO)**

1159 In the reference platform, the runtime status or execution state of the entire application system is
1160 due to a combination of executions of infrastructure code (e.g., networking routes for inter-
1161 service communication, resources provisioning code), policy code (e.g., code that specifies
1162 authentication and authorization policies), and session management code (e.g., code that
1163 establishes an mTLS session, code that generates JWT tokens) as revealed by the execution of
1164 observability as code. The observability as code of the service mesh relays the output from the
1165 execution of infrastructure, policy, and session management codes during runtime to various
1166 monitoring tools that generate applicable metrics and log aggregation tools and tracing tools,
1167 which in turn relay their output to a centralized dashboard. The analytics that are integral to the
1168 output of these tools enable system administrators to obtain a comprehensive global view of the
1169 runtime status of the entire application system. It is the runtime performance of a DevSecOps
1170 platform enabled through continuous monitoring with zero trust design features that provides all
1171 of the necessary security assurance for cloud-native applications.

1172 The activities within the DevSecOps pipelines within the service mesh context that enable
1173 continuous ATO are:

- 1174 • Checking for compliant code: Checking whether the updated infrastructure, policy,
1175 session management, and observability codes are compliant with the chosen risk
1176 management framework (e.g., NIST Risk Management Framework) for the enterprise.
1177 This will involve tools with risk assessment features, such as the capability to generate

- 1178 actionable tasks, specify code-level guidance, and test plans for verifying compliance
1179 [31].
- 1180 • Generating a dashboard that displays the runtime status: Using the runtime output of
1181 observability code for generating the dashboard that displays the runtime status of the
1182 overall application
- 1183 Risk assessment tools should provide complete traceability for all of the artifacts displayed in the
1184 dashboard, as well as the reporting capabilities needed for continuous ATO.
- 1185 Dashboard generation tools should enable system administrators to analyze macro-level features,
1186 as well as dynamically change the composition of the artifacts to be displayed based on the
1187 evolving system and consumer needs of the environment in which the application operates.

1188 5 Summary and Conclusion

1189 This document provides a comprehensive guidance for the implementation of DevSecOps
1190 primitives for a reference platform hosting a cloud-native application. It includes an overview of
1191 the reference platform and describes the basic DevSecOps primitives (i.e., CI/CD pipelines), its
1192 building blocks, the design and execution of the pipelines, and the role of automation in the
1193 efficient execution of workflows in CI/CD pipelines.

1194
1195 The architecture of the reference platform, in addition to the application code and the code for
1196 providing application services, consists of functional elements for infrastructure, runtime policies,
1197 and continuous monitoring of the health of the application, the last three of which can be
1198 deployed through declarative codes with separate CI/CD pipelines types. The runtime behaviors
1199 of these codes, the benefits of the implementation for high-assurance security, and the use of the
1200 artifacts within the pipelines for providing a continuous authority to operate (c-ATO) using risk
1201 management tools and dashboard metrics are also described.
1202

1203 **References**

- 1204 [1] Garrison J, Nova K (2017) *Cloud Native Infrastructure* (oreilly.com). Available at
1205 <https://www.oreilly.com/library/view/cloud-native-infrastructure/9781491984291/>
1206
- 1207 [2] Chaillan NM (2020) *DoD Enterprise DevSecOps Initiative and Platform One*. Available
1208 at <https://software.af.mil/dsop/documents>
- 1209 [3] Chandramouli R (2019) Security Strategies for Microservices-based Application Systems.
1210 (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special
1211 Publication (SP) 800-204. <https://doi.org/10.6028/NIST.SP.800-204>
1212
- 1213 [4] Chandramouli R, Butcher Z (2020) Building Secure Microservices-based Applications
1214 Using Service-Mesh Architecture. (National Institute of Standards and Technology,
1215 Gaithersburg, MD), NIST Special Publication (SP) 800-204A.
1216 <https://doi.org/10.6028/NIST.SP.800-204A>
1217
- 1218 [5] Chandramouli R, Butcher Z, Aradhna C (2021) (National Institute of Standards and
1219 Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-204B.
1220 <https://doi.org/10.6028/NIST.SP.800-204B>
1221
- 1222 [6] Agarwal G (2020) *How to Manage Microservices on Kubernetes With Istio*. Available at
1223 [https://medium.com/better-programming/how-to-manage-microservices-on-kubernetes-](https://medium.com/better-programming/how-to-manage-microservices-on-kubernetes-with-istio-c25e97a60a59)
1224 [with-istio-c25e97a60a59](https://medium.com/better-programming/how-to-manage-microservices-on-kubernetes-with-istio-c25e97a60a59)
1225
- 1226 [7] Kubernetes Advocate (2021) *Managing Microservices With Istio Service Mesh* in
1227 Kubernetes. Available at [https://medium.com/avmconsulting-blog/managing-](https://medium.com/avmconsulting-blog/managing-microservices-with-istio-service-mesh-in-kubernetes-36e1fda81757)
1228 [microservices-with-istio-service-mesh-in-kubernetes-36e1fda81757](https://medium.com/avmconsulting-blog/managing-microservices-with-istio-service-mesh-in-kubernetes-36e1fda81757)
1229
- 1230 [8] Ramakani A (2020) *Kong API Gateway – From Zero to Production*. Available at
1231 <https://medium.com/swlh/kong-api-gateway-zero-to-production-5b8431495ee>
1232
- 1233 [9] Srinath K (2020) *DevSecOps – Baking Security into Development Process*. Available at
1234 [https://medium.com/faun/devsecops-baking-security-into-development-process-](https://medium.com/faun/devsecops-baking-security-into-development-process-9579418ad9a7)
1235 [9579418ad9a7](https://medium.com/faun/devsecops-baking-security-into-development-process-9579418ad9a7)
1236
- 1237 [10] Rubinstein D (2020) *AppSec vs. DevSecOps, and what that means for developers*.
1238 Available at [https://sdtimes.com/security/appsec-vs-devsecops-and-what-that-means-for-](https://sdtimes.com/security/appsec-vs-devsecops-and-what-that-means-for-developers/)
1239 [developers/](https://sdtimes.com/security/appsec-vs-devsecops-and-what-that-means-for-developers/)
1240
- 1241 [11] Red Hat (2021) *What is DevSecOps?* Available at
1242 <https://www.redhat.com/en/topics/devops/what-is-devsecops>
1243
- 1244 [12] Red Hat (2021) *What is CI/CD?* Available at

- 1245 [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)
1246 [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)
1247 [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)
1248
- 1249 [13] Ali M (2018) *Continuous Release Practices are evolving, Here is our story*. Available at
1250 [https://medium.com/the-telegraph-engineering/continuous-release-practices-are-evolving-](https://medium.com/the-telegraph-engineering/continuous-release-practices-are-evolving-here-is-our-story-2a4d164e9cac)
1251 [here-is-our-story-2a4d164e9cac](https://medium.com/the-telegraph-engineering/continuous-release-practices-are-evolving-here-is-our-story-2a4d164e9cac)
1252
- 1253 [14] Pittet S (2021) *Continuous integration vs. continuous delivery vs. continuous deployment*.
1254 Available at [https://www.atlassian.com/continuous-delivery/principles/continuous-](https://www.atlassian.com/continuous-delivery/principles/continuous-integration-vs-delivery-vs-deployment)
1255 [integration-vs-delivery-vs-deployment](https://www.atlassian.com/continuous-delivery/principles/continuous-integration-vs-delivery-vs-deployment)
1256
- 1257 [15] Wuestkamp K (2020) *K8s-native Jenkins-X and Tekton Pipelines*. Available at
1258 <https://itnext.io/k8s-native-jenkins-x-and-tekton-pipelines-e2b5a61a1d22>
1259
- 1260 [16] aws.amazon.com (2021) *Create Continuous Delivery Pipeline*. Available at
1261 [https://aws.amazon.com/getting-started/hands-on/create-continuous-delivery-](https://aws.amazon.com/getting-started/hands-on/create-continuous-delivery-pipeline/?trk=gs_card)
1262 [pipeline/?trk=gs_card](https://aws.amazon.com/getting-started/hands-on/create-continuous-delivery-pipeline/?trk=gs_card)
1263
- 1264 [17] aws.amazon.com (2021) *Setting up a CI/CD pipeline by integrating Jenkins with AWS*
1265 *CodeBuild and AWS CodeDeploy*. [https://aws.amazon.com/blogs/devops/setting-up-a-ci-](https://aws.amazon.com/blogs/devops/setting-up-a-ci-cd-pipeline-by-integrating-jenkins-with-aws-codebuild-and-aws-codedeploy/)
1266 [cd-pipeline-by-integrating-jenkins-with-aws-codebuild-and-aws-codedeploy/](https://aws.amazon.com/blogs/devops/setting-up-a-ci-cd-pipeline-by-integrating-jenkins-with-aws-codebuild-and-aws-codedeploy/)
- 1267 [18] mylocaldevstack.pl (2019) *The future of DevOps – Assembly Lines*. Available at
1268 [https://medium.com/@mylocaldevstack/the-future-of-devops-assembly-lines-](https://medium.com/@mylocaldevstack/the-future-of-devops-assembly-lines-40227546d750)
1269 [40227546d750](https://medium.com/@mylocaldevstack/the-future-of-devops-assembly-lines-40227546d750)
- 1270 [19] Socher R (2020) *Best Terraform Tutorial Guides: An Overview*. Available at
1271 <https://faun.pub/best-terraform-tutorial-guides-an-overview-65a6fcee0a24>
- 1272 [20] Pariseau B (2020) *Mendix dumps cluttered DevOps monitoring tools for Datadog*.
1273 Available at [https://searchitoperations.techtarget.com/news/252497697/Mendix-dumps-](https://searchitoperations.techtarget.com/news/252497697/Mendix-dumps-cluttered-DevOps-monitoring-tools-for-Datadog?track=NL-1841&ad=938089&asrc=EM_NLN_151848158&utm_medium=EM&utm_source=NLN&utm_campaign=20210315_Kong+Mesh%27s+built-in+Open+Policy+Agent+simplifies+IT+security+management)
1274 [cluttered-DevOps-monitoring-tools-for-Datadog?track=NL-](https://searchitoperations.techtarget.com/news/252497697/Mendix-dumps-cluttered-DevOps-monitoring-tools-for-Datadog?track=NL-1841&ad=938089&asrc=EM_NLN_151848158&utm_medium=EM&utm_source=NLN&utm_campaign=20210315_Kong+Mesh%27s+built-in+Open+Policy+Agent+simplifies+IT+security+management)
1275 [1841&ad=938089&asrc=EM_NLN_151848158&utm_medium=EM&utm_source=NLN](https://searchitoperations.techtarget.com/news/252497697/Mendix-dumps-cluttered-DevOps-monitoring-tools-for-Datadog?track=NL-1841&ad=938089&asrc=EM_NLN_151848158&utm_medium=EM&utm_source=NLN&utm_campaign=20210315_Kong+Mesh%27s+built-in+Open+Policy+Agent+simplifies+IT+security+management)
1276 [&utm_campaign=20210315_Kong+Mesh%27s+built-](https://searchitoperations.techtarget.com/news/252497697/Mendix-dumps-cluttered-DevOps-monitoring-tools-for-Datadog?track=NL-1841&ad=938089&asrc=EM_NLN_151848158&utm_medium=EM&utm_source=NLN&utm_campaign=20210315_Kong+Mesh%27s+built-in+Open+Policy+Agent+simplifies+IT+security+management)
1277 [in+Open+Policy+Agent+simplifies+IT+security+management](https://searchitoperations.techtarget.com/news/252497697/Mendix-dumps-cluttered-DevOps-monitoring-tools-for-Datadog?track=NL-1841&ad=938089&asrc=EM_NLN_151848158&utm_medium=EM&utm_source=NLN&utm_campaign=20210315_Kong+Mesh%27s+built-in+Open+Policy+Agent+simplifies+IT+security+management)
- 1278 [21] Boutet R (2018) *Log Patterns: Automatically cluster your logs for faster investigation*
1279 (datadoghq.com). Available at <https://www.datadoghq.com/blog/log-patterns/>
- 1280 [22] Soni A (2020) *GitOps: The Next Big Thing for DevOps and Automation!* Available at
1281 [https://medium.com/searce/gitops-the-next-big-thing-for-devops-and-automation-](https://medium.com/searce/gitops-the-next-big-thing-for-devops-and-automation-2a9597e51559)
1282 [2a9597e51559](https://medium.com/searce/gitops-the-next-big-thing-for-devops-and-automation-2a9597e51559)
- 1283 [23] Fallon A (2020) *Understand the role of infrastructure as code in DevOps*
1284 (TechTarget.com). Available at

- 1285 https://searchitoperations.techtarget.com/feature/Understand-the-role-of-infrastructure-as-code-in-DevOps?utm_campaign=20210809_The+next+DevSecOps+challenge%3A+People&utm_medium=EM&utm_source=NLN&track=NL-1841&ad=939963&asrc=EM_NLN_174809933
- 1286
- 1287
- 1288
- 1289
- 1290 [24] Ahmed M (2020) *Introducing Policy As Code: The Open Policy Agent (OPA)*. Available at <https://www.magalix.com/blog/introducing-policy-as-code-the-open-policy-agent-opa>
- 1291
- 1292 [25] Singh P (2021) *Tackle Kubernetes observability with the right metrics*. Available at https://searchitoperations.techtarget.com/tip/Tackle-Kubernetes-observability-with-the-right-metrics?track=NL-1841&ad=938191&asrc=EM_NLN_153034984&utm_medium=EM&utm_source=NLN&utm_campaign=20210322_DevSecOps+leaves+Excel+in+the+dust
- 1293
- 1294
- 1295
- 1296
- 1297
- 1298 [26] Devopscurry (2021) *Securing your CI/CD pipelines with DevSecOps in 2021*. Available at <https://medium.com/devopscurry/securing-your-ci-cd-pipelines-with-devsecops-in-2021-1a6a6e34f2e7>
- 1299
- 1300
- 1301 [27] Gartner (2021) Gardner D, Horvath M, Zumerle D *Magic Quadrant for Application Security Testing*. Available at <https://www.gartner.com/doc/reprints?id=1-262TXQZV&ct=210518&st=sb>
- 1302
- 1303
- 1304 [28] Lewkowicz J (2021) *A guide to automated testing providers*. Available at <https://sdtimes.com/test/a-guide-to-automated-testing-providers/>
- 1305
- 1306 [29] Lewkowicz J (2021) *Automated testing is a must in CI/CE pipelines*. Available at <https://sdtimes.com/test/automated-testing-is-a-must-in-ci-cd-pipelines/>
- 1307
- 1308
- 1309 [30] Angel R (2020) *How to be a CI/CD Engineer*. Available at <https://circleci.com/blog/how-to-be-a-cicd-engineer/>
- 1310
- 1311
- 1312 [31] Alexey K (2020) *Ultimate guide to CI/CD security and DevSecOps*. Available at <https://circleci.com/blog/security-best-practices-for-ci-cd/>
- 1313
- 1314 [32] Agarwal G (2020) *Kubernetes Security With Falco*. Available at <https://betterprogramming.pub/kubernetes-security-with-falco-2eb060d3ae7d>
- 1315
- 1316 [33] Marko K (2021) *Terraform cheat sheet: Notable commands, HCL and more*. Available at https://searchitoperations.techtarget.com/tip/Terraform-cheat-sheet-Notable-commands-HCL-and-more?utm_campaign=20210726_Infrastructure+as+code+still+a+big+security+buzz&utm_medium=EM&utm_source=NLN&track=NL-1841&ad=939808&asrc=EM_NLN_172629823
- 1317
- 1318
- 1319
- 1320
- 1321

- 1322 [34] Guo T (2021) *On DevOps 8: Infrastructure as Code: Introduction, Best Practices and*
1323 *How to choose the Right tool*. Available at [https://medium.com/4th-coffee/on-devops-8-](https://medium.com/4th-coffee/on-devops-8-infrastructure-as-code-introduction-best-practices-and-choosing-the-right-tool-2c8f46d1f34)
1324 [infrastructure-as-code-introduction-best-practices-and-choosing-the-right-tool-](https://medium.com/4th-coffee/on-devops-8-infrastructure-as-code-introduction-best-practices-and-choosing-the-right-tool-2c8f46d1f34)
1325 [2c8f46d1f34](https://medium.com/4th-coffee/on-devops-8-infrastructure-as-code-introduction-best-practices-and-choosing-the-right-tool-2c8f46d1f34)
- 1326 [35] Pulumi.com (2020) *Delivering Cloud Native Infrastructure as Code*. Available at
1327 <https://cdn2.hubspot.net/hubfs/4429525/Content/Pulumi-Delivering-CNI-as-Code.pdf>
1328
- 1329 [36] Chong T (2021) *Product In-Depth: Instantly Secure your Cloud Infrastructure with Policy-as-*
1330 *Code*. Available at [https://www.magalix.com/blog/instantly-secure-your-cloud-](https://www.magalix.com/blog/instantly-secure-your-cloud-infrastructure-with-policy-as-code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-OcHJmlUHG8eFhC5Tfu7MW_Dn3i-g3Ng&utm_content=123795689&utm_source=hs_email)
1331 [infrastructure-with-policy-as-](https://www.magalix.com/blog/instantly-secure-your-cloud-infrastructure-with-policy-as-code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-OcHJmlUHG8eFhC5Tfu7MW_Dn3i-g3Ng&utm_content=123795689&utm_source=hs_email)
1332 [code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-](https://www.magalix.com/blog/instantly-secure-your-cloud-infrastructure-with-policy-as-code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-OcHJmlUHG8eFhC5Tfu7MW_Dn3i-g3Ng&utm_content=123795689&utm_source=hs_email)
1333 [HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-](https://www.magalix.com/blog/instantly-secure-your-cloud-infrastructure-with-policy-as-code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-OcHJmlUHG8eFhC5Tfu7MW_Dn3i-g3Ng&utm_content=123795689&utm_source=hs_email)
1334 [OcHJmlUHG8eFhC5Tfu7MW_Dn3i-](https://www.magalix.com/blog/instantly-secure-your-cloud-infrastructure-with-policy-as-code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-OcHJmlUHG8eFhC5Tfu7MW_Dn3i-g3Ng&utm_content=123795689&utm_source=hs_email)
1335 [g3Ng&utm_content=123795689&utm_source=hs_email](https://www.magalix.com/blog/instantly-secure-your-cloud-infrastructure-with-policy-as-code?utm_medium=email&_hsmi=123795689&_hsenc=p2ANqtz-HbHGALCOrcV_WXjrKE-JgWUWUX9Tg1s9ot5HSox2Ts1rgAhwquoZkBKqZYo-OcHJmlUHG8eFhC5Tfu7MW_Dn3i-g3Ng&utm_content=123795689&utm_source=hs_email)
- 1336 [37] Sheldon R (2021) *Top 30 DevOps interview questions and answers for 2021*. Available at
1337 [https://whatis.techtarget.com/feature/Top-30-DevOps-interview-questions-and-](https://whatis.techtarget.com/feature/Top-30-DevOps-interview-questions-and-answers?utm_campaign=20210712_New+Kubernetes+use+case%253A+Hacking&utm_medium=EM&utm_source=NLN&track=NL-1841&ad=939635&asrc=EM_NLN_170130891)
1338 [answers?utm_campaign=20210712_New+Kubernetes+use+case%253A+Hacking&utm](https://whatis.techtarget.com/feature/Top-30-DevOps-interview-questions-and-answers?utm_campaign=20210712_New+Kubernetes+use+case%253A+Hacking&utm_medium=EM&utm_source=NLN&track=NL-1841&ad=939635&asrc=EM_NLN_170130891)
1339 [medium=EM&utm_source=NLN&track=NL-](https://whatis.techtarget.com/feature/Top-30-DevOps-interview-questions-and-answers?utm_campaign=20210712_New+Kubernetes+use+case%253A+Hacking&utm_medium=EM&utm_source=NLN&track=NL-1841&ad=939635&asrc=EM_NLN_170130891)
1340 [1841&ad=939635&asrc=EM_NLN_170130891](https://whatis.techtarget.com/feature/Top-30-DevOps-interview-questions-and-answers?utm_campaign=20210712_New+Kubernetes+use+case%253A+Hacking&utm_medium=EM&utm_source=NLN&track=NL-1841&ad=939635&asrc=EM_NLN_170130891)
- 1341 [38] Prasad A (2020) *Kubernetes Admission Controllers: Request Interceptors*. Available at
1342 [https://medium.com/cloudlego/kubernetes-admission-controllers-request-interceptors-](https://medium.com/cloudlego/kubernetes-admission-controllers-request-interceptors-47a9b12c5303)
1343 [47a9b12c5303](https://medium.com/cloudlego/kubernetes-admission-controllers-request-interceptors-47a9b12c5303)