



**NIST Internal Report  
NIST IR 8505 ipd**

# **A Data Protection Approach for Cloud-Native Applications**

Initial Public Draft

Ramaswamy Chandramouli  
Wesley Hales

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.IR.8505.ipd>

**NIST Internal Report**  
**NIST IR 8505 ipd**

# **A Data Protection Approach for Cloud-Native Applications**

Initial Public Draft

Ramaswamy Chandramouli  
*Computer Security Division*  
*Information Technology Laboratory*

Wesley Hales  
*Leak Signal Inc.*

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.IR.8505.ipd>

June 2024



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

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

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

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

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

#### **NIST Technical Series Policies**

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

[NIST Technical Series Publication Identifier Syntax](#)

#### **Publication History**

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

#### **How to Cite this NIST Technical Series Publication**

Chandramouli R, Hales W (2024) A Data Protection Approach for Cloud-Native Applications. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) NIST IR 8505 ipd. <https://doi.org/10.6028/NIST.IR.8505.ipd>

#### **Author ORCID iDs**

Ramaswamy Chandramouli: 0000-0002-7387-5858

#### **Public Comment Period**

June 14, 2024 – August 1, 2024

#### **Submit Comments**

[nistir-8505-comments@nist.gov](mailto:nistir-8505-comments@nist.gov)

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

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

1 **Abstract**

2 This document addresses the need for effective data protection strategies in the evolving realm  
3 of cloud-native network architectures, including multi-cloud environments, service mesh  
4 networks, and hybrid infrastructures. By extending foundational data categorization concepts,  
5 it provides a framework for aligning data protection approaches with the unknowns of data in  
6 transit. Specifically, it explores service mesh architecture, leveraging and emphasizing the  
7 capabilities of WebAssembly (WASM) in ensuring robust data protection as sensitive data is  
8 transmitted through east-west and north-south communication paths.

9 **Keywords**

10 data governance; data privacy; data protection; data security; in-transit data categorization;  
11 WASM.

12 **Reports on Computer Systems Technology**

13 The Information Technology Laboratory (ITL) at the National Institute of Standards and  
14 Technology (NIST) promotes the U.S. economy and public welfare by providing technical  
15 leadership for the Nation’s measurement and standards infrastructure. ITL develops tests, test  
16 methods, reference data, proof of concept implementations, and technical analyses to advance  
17 the development and productive use of information technology. ITL’s responsibilities include  
18 the development of management, administrative, technical, and physical standards and  
19 guidelines for the cost-effective security and privacy of other than national security-related  
20 information in federal information systems.

21

## 22 **Call for Patent Claims**

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

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

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

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

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

48 Such statements should be addressed to: [nistir-8505-comments@nist.gov](mailto:nistir-8505-comments@nist.gov)

49

50 **Table of Contents**

51 **1. Introduction .....1**

52 1.1. Existing Approaches to Data Protection and Their Limitations ..... 1

53 1.2. In-Proxy Application for Data Protection ..... 1

54 1.3. Objective and Scope of This Document ..... 2

55 1.4. Organization of This Document ..... 2

56 **2. Web Assembly Background .....4**

57 2.1. Origin..... 4

58 2.2. Progression Into Server-Side Environments ..... 4

59 2.2.1. Development and Deployment Process..... 4

60 2.3. Proxies as WASM Platforms..... 6

61 2.4. Proxy-WASM ..... 7

62 2.4.1. Role of WASM in Different Service Mesh Architectures ..... 7

63 2.5. WASI-HTTP ..... 8

64 2.6. eBPF..... 8

65 **3. Data Protection in Transit.....9**

66 3.1. Data Categorization Techniques ..... 9

67 3.2. Techniques for Data Protection ..... 9

68 3.2.1. Web Traffic Data Protection..... 9

69 3.2.2. API Security..... 10

70 3.2.3. Microsegmentation ..... 10

71 3.2.4. Log Traffic Data Protection..... 11

72 3.2.5. LLM Traffic Data Protection..... 11

73 3.2.6. Credit Card-Related Data Protection..... 12

74 3.2.7. Monitoring Tools to Visualize Sensitive Data Flows..... 12

75 **4. Security Analysis of WASM Modules .....13**

76 4.1. WASM Security Goals and Security Feature Sets..... 13

77 4.1.1. User-Level Security Features ..... 14

78 4.1.2. Security Primitives for Developers ..... 14

79 4.2. Memory Model and Memory Safety..... 14

80 4.3. Execution Model and Control Flow Integrity ..... 15

81 4.4. Security of API Access to OS and Host Resources ..... 16

82 4.5. Protection From Side-Channel Attacks ..... 16

83 4.6. Protection Against Code Injection and Other Attacks ..... 16

84 4.7. Deployment and Operating Security..... 16

85	<b>5. Summary and Conclusions .....</b>	<b>18</b>
86	<b>References.....</b>	<b>19</b>
87	<b>Appendix A. Execution Model for Web Assembly in Browsers .....</b>	<b>20</b>
88	<b>Appendix B. Comparison of Execution Models for Containers and WASM Modules.....</b>	<b>21</b>
89		

## 90 **1. Introduction**

91 In the constantly evolving landscape of cloud-native application architectures, where data  
92 resides in multiple locations (i.e., on-premises and on the cloud), ensuring data security involves  
93 more than simply specifying and granting authorization during service requests. It also involves  
94 a comprehensive strategy to categorize and analyze data access and leakage as data travels  
95 across various protocols (e.g., gRPC, REST-based), especially within ephemeral and scalable  
96 microservices applications. As organizations find themselves governing hundreds to tens of  
97 thousands of services and the inter-service calls between them, a security void has been  
98 identified in observing and protecting sensitive data in transit.

### 99 **1.1. Existing Approaches to Data Protection and Their Limitations**

100 Traditionally, regular expressions (regex) have been widely used for data categorization to  
101 identify patterns that match predefined categories or data classes with the aid of keywords and  
102 validators for enhanced precision. Despite its wide adoption and usage, the approach has  
103 notable limitations. The processing time scales linearly with data volume, making it impractical  
104 for very large datasets. Regex also lacks the capability for logical computations, which are  
105 necessary for complex validations like checksums in credit card numbers. Its effectiveness  
106 heavily relies on the correct proximity to specific keywords, leading to potential false positives  
107 and considerable noise if not managed correctly.

108 Machine learning (ML) offers a promising enhancement to data categorization by learning from  
109 data patterns and improving over time, thus providing a scalable and adaptable solution. ML  
110 algorithms can handle both structured and unstructured data, predict data categories based on  
111 historical data, and adjust to new patterns without explicit reprogramming. This adaptability  
112 significantly reduces the time and computational resources required to manage complex  
113 datasets and is effective for both data at rest and in motion.

114 To address and complement the limitations of traditional data-at-rest inventory, in-transit data  
115 categorization has recently come to light as the next logical step in data protection. Unlike the  
116 former, which only secures stored information, in-transit categorization actively monitors and  
117 secures data as it moves across services and network protocols. This shift to real-time data  
118 analysis within the network brings new observability capabilities, eliminating the need for  
119 traffic mirroring and data duplication.

### 120 **1.2. In-Proxy Application for Data Protection**

121 To address the need for data categorization during travel across services, a relatively new class  
122 of in-proxy application called the WebAssembly program (also called a WASM module) has  
123 been increasingly deployed. A WASM module is a lightweight executable compiled to low-level  
124 bytecode. This bytecode can be:

- 125 (a) Generated from code written in any language using their associated WebAssembly  
126 compilers, including C, C++, and Rust



127 (b) Run using a WASM runtime in an isolated virtual machine (VM) within the proxy, which  
128 allows developers to enhance applications with necessary functionality and run them as  
129 efficiently as native code in the proxies.

130 Over the last few years, the Envoy WASM VM has enabled new types of compute and traffic  
131 processing capabilities and allowed for custom WASM modules to be built and deployed in a  
132 sandboxed and fault-tolerant manner.

133 Additionally, the following features of WebAssembly modules make them particularly effective  
134 for data protection:

- 135 • **Data Discovery and Categorization:** WASM modules can dynamically identify and  
136 categorize data as it traverses the network, ensuring that sensitive information is  
137 recognized and handled appropriately.
- 138 • **Dynamic Data Masking (DDM):** WASM modules can apply DDM techniques to redact or  
139 mask sensitive information in transit, enhancing privacy and security.
- 140 • **User and Entity Behavior Analytics (UEBA):** WASM modules can analyze user and entity  
141 behaviors in real time, detecting anomalies and potential security threats.
- 142 • **Data Loss Prevention (DLP):** WASM modules can enforce DLP policies by monitoring and  
143 controlling data transfers to prevent unauthorized data exfiltration.

### 144 **1.3. Objective and Scope of This Document**

145 All services (e.g., networking, security, monitoring, etc.) for microservices-based applications  
146 are provided by a centralized infrastructure called the service mesh, and the data plane for this  
147 service mesh — which performs all runtime tasks — consists of proxies. This document outlines  
148 a practical framework for effective data protection and highlights the versatile capabilities of  
149 WebAssembly (WASM) within service mesh architectures, multi-cloud environments, and  
150 hybrid (i.e., a combination of on-premises and cloud-based) infrastructures. By focusing on in-  
151 line, network traffic analysis at layers 4–7, organizations can enhance security, streamline  
152 operations, and utilize adaptive data protection measures.

### 153 **1.4. Organization of This Document**

154 This document is organized as follows:

- 155 • Section 2 describes the execution environment for WASM modules in detail, including  
156 the application infrastructure (i.e., service mesh) under which it runs, the specific host  
157 environment (i.e., proxies), the process for generating bytecodes and executables, the  
158 processes for executing the modules using a WASM runtime, and an API (i.e., WASI) for  
159 accessing OS resources of the underlying platform.
- 160 • Section 3 introduces the concept of data categorization and the use of various data  
161 protection techniques (e.g., data masking, redaction, etc.) to ensure the security of data  
162 in different domains or application scenarios using WASM modules, such as web traffic

163 data protection, API Security, microsegmentation, log traffic data protection, LLM traffic  
164 data protection, and integration with monitoring tools for the visualization of sensitive  
165 data flows.

- 166 • Section 4 presents a detailed security analysis of a WASM module by examining its  
167 development, deployment, and execution environment to ensure that the module  
168 satisfies the properties of a security kernel and can provide the necessary security  
169 assurance.
  - 170 • Section 5 provides a summary of the topics covered in this document and discusses how  
171 WASM module functionality must continuously evolve to provide the security assurance  
172 needed to protect against data breaches and exfiltration in the context of increasingly  
173 sophisticated attacks on data.
- 174

## 175 **2. Web Assembly Background**

176 WebAssembly modules are deployed to protect data on microservices-based architectures in  
177 which the entire application (also called a cloud-native application because of its ubiquitous  
178 deployment in cloud and hybrid environments) consists of several distributed, loosely coupled,  
179 and independently scalable components called microservices. All services for this class of  
180 application (e.g., networking, security policies enforcement, state monitoring, configuration of  
181 runtime parameters) are provided by a centralized application-independent service  
182 infrastructure called the service mesh. This service mesh consists of a data plane that is  
183 primarily made up of proxies that house the various service modules. Using the family of APIs  
184 provided by the proxies, relevant service modules (e.g., network path determination) are  
185 implemented using the management/control plane of the service mesh. The WebAssembly is  
186 one such service module ecosystem implemented in the data plane proxies of a service mesh.

### 187 **2.1. Origin**

188 WASM modules originated in browser environments and were designed to run in memory-safe  
189 sandboxes, making them more secure than running client-side JavaScript. The execution model  
190 for running WebAssembly code in browsers is given in Appendix A. In addition to security,  
191 WASM modules have the following advantages [1]:

- 192 • **Performance:** Due to its low-level binary format targeted for modern processors, WASM  
193 modules provide near-native performance. Hence, it is considered the “fourth language”  
194 for the web alongside HTML, CSS, and JavaScript and is designed to enable high-  
195 performance applications in browsers.
- 196 • **Broad support:** It has broad accessibility and is supported in popular browsers, such as  
197 Chrome, Firefox, Edge, and Safari.

### 198 **2.2. Progression Into Server-Side Environments**

199 WASM modules progressed from browser to server environments when Mozilla introduced an  
200 open-source project called the WebAssembly System Interface (WASI) that provided a  
201 framework for WebAssembly apps to access operating system resources [4]. This allowed for  
202 content delivery networks (CDNs) to use WebAssembly to deploy customers’ apps without  
203 giving them access to the underlying CDN infrastructure.

#### 204 **2.2.1. Development and Deployment Process**


205 The emergence of WASM compilers for several languages enabled developers to use their  
206 preferred languages to create server-side applications. Additionally, server-side WASM code  
207 could run inside containers as well as VMs. It is a potential candidate for SaaS-based offerings,  
208 just like VMs and containers. Its portability allows applications to run anywhere, making it an  
209 attractive option for various use cases.

210 The steps involved in developing a WASM module and running it using WASM runtime are [6]:

- 211 • **Source code writing:** Programs are written in languages (e.g., C++, C#, Rust, etc.) that  
212 have target WASM compilers available.
- 213 • **Parsing:** The code is parsed into an abstract syntax tree (AST) structure.
- 214 • **Compiling:** The code in AST structure is then compiled into a WASM module using AOT  
215 or JIT. The WASM module is generated in a binary format that can be executed by  
216 WASM runtime.
- 217 • **WASM runtime loading:** The WASM runtime loads the WASM module (with file name  
218 extension. wasm). If JIT is used, the compilation takes place after loading into WASM  
219 runtime.
- 220 • **Preparation for execution (i.e., instantiation):** The WASM runtime creates an  
221 executable instance from the WASM module by allocating memory, importing functions  
222 and objects, and establishing the execution environment for the module.
- 223 • **Code optimization:** During execution of the byte code, profiling is employed to identify  
224 frequently executed code, and a progressive optimization/re-optimization process takes  
225 place to gradually enhance performance until the code runs efficiently.

226 Figure 1 shows the ability to develop programs in different languages, convert them into WASM  
227 code, and run them under different processor architectures [4]. The execution model for WASM  
228 modules in the server environment and their comparison with the container execution model  
229 are described in Appendix B.

230



231

232

Fig. 1. Generating WASM modules and their execution [4]

### 233 2.3. Proxies as WASM Platforms

234 Proxies are increasingly being used as platforms for executing WASM modules. In cloud-native  
235 and microservices-based applications, proxies mediate inter-service communication. Open-  
236 source projects in proxies, such as Envoy, have extended their filter chain to allow for calling  
237 and executing WASM modules. These WASM modules can enforce policy-based authorizations  
238 or implement network resiliency measures, providing essential security controls for such  
239 applications. Additionally, the capabilities of these modules can be leveraged for data  
240 protection purposes.

241 The advantages of network based WASM modules include:

- 242 1. **Extensibility:** Proxies like Envoy can be extended with WASM modules, allowing  
243 developers to introduce custom logic and functionality without modifying the proxy's  
244 core codebase. This extensibility allows for the seamless integration of new features and  
245 capabilities.
- 246 2. **Security and isolation:** WASM modules run in a sandboxed environment, providing  
247 isolation from the host system and other modules. This isolation enhances security by  
248 preventing unauthorized access to system resources and mitigating the impact of  
249 potential vulnerabilities.
- 250 3. **Portability:** WebAssembly's portability ensures that WASM modules can run  
251 consistently across different proxy implementations and platforms, promoting a write-  
252 once, run-anywhere approach.
- 253 4. **Performance:** WASM modules can potentially offer better performance compared to  
254 the traditional scripting languages used for proxy extensions since they can be compiled  
255 to efficient machine code.
- 256 5. **Policy enforcement and network resiliency:** By executing WASM modules in proxies,  
257 organizations can enforce policies, implement authorization controls, and introduce  
258 network resiliency measures at the proxy level, ensuring consistent and centralized  
259 enforcement across distributed applications.
- 260 6. **Data protection:** WASM modules in proxies can be used to implement data filtering,  
261 transformation, or encryption mechanisms and ensure sensitive data protection as it  
262 flows through the proxy.
- 263 7. **Ecosystem and community:** The growing WebAssembly ecosystem and community  
264 provide libraries, tools, and resources that foster collaboration and accelerate the  
265 development of proxy extensions and data protection solutions.

266 As WASM continues to mature, its role in proxies will expand, enabling proxies to act as robust  
267 platforms for security and application logic execution. This evolution is particularly pertinent to  
268 data protection, which stands as a central theme of contemporary application development.

## 269 **2.4. Proxy-WASM**

270 Envoy Proxy, an open-source edge and service proxy, plays a pivotal role in managing the flow  
271 of traffic between microservices in many service mesh deployments. The collection of  
272 extensible APIs that it provides for various services is designated as xDS. An extension API that  
273 leverages the extensibility of these basic, foundational APIs of Envoy Proxy is the WebAssembly  
274 for Proxies (Proxy-WASM) runtime.

275 Proxy-WASM extends the adaptability of Envoy Proxy by enabling the deployment of  
276 WebAssembly modules within the proxy server. This integration allows for the execution of  
277 custom code directly within the proxy, providing a platform-independent and secure  
278 environment. The modularity of WebAssembly makes it an ideal choice for extending the  
279 functionalities of Envoy Proxy without the need for recompilation or significant changes to the  
280 existing infrastructure.

281 The architecture of Proxy-WASM within Envoy Proxy allows for the seamless integration and  
282 execution of custom logic at various stages of the request-response cycle. For example, a  
283 WASM module can intercept requests, inspect payload data, apply data categorizations, and  
284 redact data before proceeding. This level of granular control enhances the security posture of  
285 microservices architectures while maintaining performance and scalability.

286 Proxy-WASM can be leveraged to implement robust security measures for microservices  
287 communication. WASM modules can perform tasks, such as data categorization and mitigation  
288 directly within the proxy.

### 289 **2.4.1. Role of WASM in Different Service Mesh Architectures**

290 Service mesh architectures have traditionally utilized sidecar proxies, which are implemented as  
291 containers and deployed alongside each service within a Kubernetes pod. These sidecar proxies  
292 manage both inbound and outbound traffic for their respective services, creating an ideal  
293 WASM-based insertion point for in-transit categorization.

294 Additionally, newer architectural patterns (e.g., proxy implementation/deployment models)  
295 recognize that sidecar proxies are excessive because many services do not have L7-level  
296 services. The ambient waypoint proxy pattern seeks to simplify the sidecar model by  
297 centralizing and simplifying traffic management and policy enforcement. In this pattern,  
298 waypoint proxies are deployed at the node level, which provides application services either per  
299 namespace or per service account. They manage all ingress and egress traffic for the services  
300 within their designated scope. In both proxy deployment models, the WebAssembly VM  
301 intercepts and analyzes traffic in the exact same way, providing a transparent deployment for  
302 WASM-based data categorization policies and modules.

303 Outside of traditional Envoy-based service mesh proxies, there are several runtime  
304 environments where WASM modules can be deployed to classify sensitive data in transit. Many  
305 API gateways now support WASM along with commercial content delivery network (CDN)  
306 platforms, such as Fastly's WASM Compute Platform and Cloudflare's WASM Workers.

## 307 **2.5. WASI-HTTP**

308 With its application binary interface (ABI), Proxy-WASM facilitates communication between  
309 WebAssembly modules and host environments, specifically proxies. It has a mature  
310 specification adopted by various proxy servers and traces its origins to efforts within the Envoy  
311 project to extend the capabilities of proxy servers using WebAssembly. Proxy-WASM ensures  
312 that extensions written for one proxy can be reused in others, promoting a write-once, run-  
313 anywhere approach. Proxy-WASM's ABI and event-driven streaming APIs have been  
314 incorporated into several production-level proxies, demonstrating the project's practical  
315 application and influence.

316 In contrast, WASI-HTTP — a WASM-based API — has evolved through iterations to define  
317 interfaces for handling HTTP requests and responses directly within WASM modules. It aims to  
318 provide a minimal and streamlined execution environment for WebAssembly-based HTTP  
319 proxies and is designed to seamlessly integrate with existing web infrastructure, such as service  
320 workers and reverse proxies, without requiring a complex runtime system. WASI-HTTP is  
321 already in production in some environments and supports scalable and dynamic WASM  
322 instance creation in response to web traffic, laying the groundwork for future innovations like  
323 linking HTTP intermediaries through the component model.

324 Both WASI-HTTP and Proxy-WASM are shaping the landscape of WebAssembly in networked  
325 and distributed systems. While WASI-HTTP is allowing for simplified HTTP communication  
326 within WebAssembly applications, Proxy-WASM exemplifies the successful implementation of a  
327 standardized interface across multiple proxy implementations. Their collaborative development  
328 highlights a symbiotic relationship, with WASI-HTTP potentially leveraging Proxy-WASM's ABI to  
329 further enhance the capabilities and reach of WebAssembly in networking scenarios.

## 330 **2.6. eBPF**

331 Using WASM to parse human-readable text in Layers 4–7 offers several advantages over  
332 technologies like eBPF, particularly regarding handling complex application-layer data, such as  
333 HTTP. While eBPF is powerful for data capture and manipulation directly within the kernel, its  
334 use for parsing detailed HTTP traffic can be complex and potentially excessive for some  
335 applications. This complexity stems from the need to handle the intricacies of HTTP within the  
336 kernel — a task that can restrict performance and introduce security concerns if not managed  
337 correctly. Additionally, eBPF imposes numerous restrictions and requires extra effort for data  
338 processing and general-purpose computation.

339 WASM provides a secure, sandboxed environment that is suitable for efficiently executing code  
340 across multiple platforms and parsing application-layer protocols. WASM can be used in user  
341 spaces and server environments, allow easier integration with existing parsing libraries and  
342 tools, reduce complexity, and potentially enhance the reliability of parsing operations. Its  
343 portability and ability to embed in various runtime environments make it a practical choice for  
344 network traffic analysis tasks, including those involving protocols that handle human-readable  
345 text.

### 346 **3. Data Protection in Transit**

347 One of the first and most fundamental tasks in data protection is classifying data to identify the  
348 need for further operations (e.g., sanitization, filtering, etc.).

#### 349 **3.1. Data Categorization Techniques**

350 Data in transit can vary wildly between structured and unstructured formats. For real-time  
351 categorization and protection, care must be taken to formulate the right approach. The  
352 performance of each categorization event is critical to ensuring that minimal latency is added as  
353 the process takes place. By executing WASM modules in proxies, organizations can implement  
354 data categorization and filtering mechanisms at the proxy level. This approach allows for the  
355 identification and protection of sensitive data as it flows between services.

356 Regex and ML models can be used within these WASM modules to detect patterns and classify  
357 data in real time, enabling the implementation of appropriate data protection measures, such  
358 as redaction, encryption, or access control policies. Regex matching can identify complex  
359 patterns for nuanced categorization schemes, and ML tools can detect patterns that signify  
360 categorization attributes. This latter process involves classifying a set of example data and  
361 training one or more models to analyze and classify future data. Though it is potentially the  
362 most effective automatic categorization method, it requires significant setup and management.  
363 The training data sets must be comprehensive to provide ample information for accurate  
364 categorization detection.

365 Unlike other data categorization techniques that operate on data at rest, in-transit  
366 categorization provides the added dimension of time as traffic is analyzed. When combining  
367 data categorization with the time it was accessed or sent, data flows can be visualized and  
368 understood. Once models have been trained on normal data flow patterns, it becomes clear  
369 when a violation in data access has occurred or when an unpermitted data flow has been  
370 established. By leveraging the capabilities of WASM modules in proxies, organizations can gain  
371 visibility into data flows, detect anomalies, and take proactive measures to protect sensitive  
372 data as it moves between services in cloud-native and microservices-based applications.

#### 373 **3.2. Techniques for Data Protection**

374 This section describes the practical uses of the data protection techniques dynamic data  
375 masking (DDM), user and entity behavior analytics (UEBA), and data loss prevention (DLP)  
376 within WASM modules in various application scenarios with a focus on the domain data that  
377 pertains to each application scenario.

##### 378 **3.2.1. Web Traffic Data Protection**

379 In-transit data categorization across web protocols like HTTP/2 and gRPC enable the  
380 observability of data flows between services and clients. By classifying data in motion,  
381 organizations can monitor how sensitive information is accessed by both unauthenticated and



382 authenticated identities. WASM modules can use regex and ML models to identify sensitive  
383 data patterns in HTTP payloads and redact, mask, or block classified data transmissions based  
384 on configured policies. Example applications include:

- 385 • **E-commerce websites:** Monitoring credit card details and personal information during  
386 transactions to ensure that they are properly encrypted and masked, preventing  
387 unauthorized access.
- 388 • **Healthcare applications:** Protecting patient data by detecting and encrypting sensitive  
389 information, such as medical records and personal identifiers before they are  
390 transmitted between systems.
- 391 • **Corporate communications:** Scanning and securing internal emails and messages to  
392 prevent data breaches and ensure compliance with internal data protection policies.

### 393 3.2.2. API Security

394 APIs are critical conduits for sensitive data and are often targeted for attacks. Monitoring data  
395 transmitted to and from APIs is essential for detecting vulnerabilities, such as application-level  
396 DDoS attacks, SQL injection, or data exfiltration. Many API gateways and service meshes  
397 support running WASM modules for enhanced security. These modules can implement  
398 authentication, rate limiting, and payload inspection for API traffic. Example applications  
399 include:

- 400 • **Financial services:** Protecting API endpoints that handle financial transactions by  
401 detecting and blocking SQL injection attempts and unauthorized access attempts.
- 402 • **Social media platforms:** Monitoring data flow through APIs to prevent the exfiltration of  
403 user data and ensure that sensitive information, such as login credentials and personal  
404 messages, are protected.
- 405 • **IoT devices:** Securing data transmitted from IoT devices to backend systems and  
406 detecting anomalies in data patterns that might indicate a security breach.

### 407 3.2.3. Microsegmentation

408 In microsegmentation, in-transit data categorization enhances asset inventory reporting. This  
409 advanced categorization enables organizations to identify and track critical assets and their  
410 data flows to ensure alignment with data protection policies. This granular insight is especially  
411 valuable for assets that handle PII or financial data, bolstering data governance and compliance  
412 efforts.

413 While Kubernetes (K8s) networking policies offer segmentation, managing and testing these  
414 policies can be resource intensive. Traditional network policies rely on static rule sets that  
415 require meticulous configuration and maintenance. Comprehensive testing across dynamic  
416 environments poses operational challenges, and these policies lack granular visibility into data  
417 content, making it difficult to accurately differentiate between legitimate and malicious traffic.

418 In contrast, in-transit data categorization offers a dynamic and granular approach. By analyzing  
419 data flows in real time, organizations gain actionable insights into the content and context of  
420 network traffic. This enables the precise enforcement of security controls based on data  
421 attributes, such as sensitivity levels or compliance requirements. Example applications include:

- 422 • **Financial institutions:** Implementing microsegmentation to protect critical systems that  
423 handle transaction processing to ensure that only authorized services can access  
424 sensitive financial data.
- 425 • **Healthcare providers:** Segregating networks within a hospital to ensure that medical  
426 devices and patient data systems are isolated from less secure administrative networks.
- 427 • **Retail chains:** Using real-time data categorization to manage data flows between point-  
428 of-sale systems and backend inventory systems to prevent unauthorized access to sales  
429 data and customer information.

#### 430 **3.2.4. Log Traffic Data Protection**

431 Regulated organizations often face the challenge of sensitive data leaking into log streams.  
432 Since all log protocols operate at Layer 4 and traverse service proxies within a service mesh,  
433 addressing potential leaks at their source allows organizations to secure data before it disperses  
434 into various storage systems, effectively mitigating the risk of exposure.

435 Example applications include:

- 436 • **Financial services:** Ensuring that transaction logs do not contain unmasked credit card  
437 numbers or personal identification information to prevent accidental leaks.
- 438 • **Healthcare providers:** Protecting patient data in system logs by redacting sensitive  
439 information before it is stored or transmitted to logging systems.
- 440 • **E-commerce platforms:** Monitoring and sanitizing log data to prevent the exposure of  
441 customer order details and personal information.

#### 442 **3.2.5. LLM Traffic Data Protection**

443 Due to their scalability needs, large language models (LLMs) typically operate within service  
444 mesh architectures. Classifying both prompt and response data in transit is crucial for  
445 governance. This enables organizations to maintain visibility over the data flows of deployed  
446 LLMs and ensure compliance with regulatory standards and organizational policies for data  
447 protection.

448 Example applications include:

- 449 • **Customer support systems:** Monitoring interactions between customers and automated  
450 support bots to ensure that sensitive customer data is not inadvertently exposed or  
451 logged.

- 452 • **Content Moderation:** Ensuring that data processed by LLMs for content moderation is  
453 handled in compliance with privacy regulations to protect user information.
- 454 • **Data Analysis Services:** Classifying and securing data used by LLMs in analytics platforms  
455 to prevent unauthorized access to sensitive business insights and customer data.

### 456 **3.2.6. Credit Card-Related Data Protection**

457 WASM modules are also used to protect data related to credit card transactions, as laid out in  
458 PCI DSS 4.0 specifications. This is achieved by incorporating the following functions into WASM  
459 modules:

- 460 • Clearly identify and document all areas in which sensitive data (e.g., cardholder data,  
461 authentication values, encryption keys, etc.) is stored, processed, or transmitted. This  
462 includes databases, servers, applications, and network segments that handle card  
463 holder data.
- 464 • Generate data-flow diagrams or other technical or topological solutions that identify  
465 flows of account data across systems and networks.
- 466 • Identify all data flows for the various stages of payment transactions (e.g., authorization,  
467 capture settlement, chargebacks, and refunds) and acceptance channels (e.g., card  
468 present, card not present, and e-commerce).

### 469 **3.2.7. Monitoring Tools to Visualize Sensitive Data Flows**

470 WASM modules can also be programmed to collect and emit metrics and telemetry data in  
471 various formats to monitoring tools that are used to visualize the flow of sensitive data (e.g.,  
472 Prometheus, Grafana etc.). By examining the normal rate of sensitive data flow over time,  
473 visual indicators, such as spikes, can be used to identify data leakage incidents and  
474 unauthorized data exposures. Subsequent investigations can then ensure compliance with data  
475 protection regulations and reduce the risk of continued data breaches.

476

#### 477 **4. Security Analysis of WASM Modules**

478 To realize the security goals for which the WASM modules are deployed, the whole ecosystem  
479 under which these modules execute must obey the properties of a security kernel:

- 480 1. It is always invoked (i.e., non-bypassable).
- 481 2. It is small and verifiable.

482 Consider the satisfaction of the first property in the context of two proxy implementation  
483 models in a service mesh. In the sidecar proxy model, a proxy is implemented as a container  
484 that coexists with each microservice in the same pod and runs in the same network space as  
485 the service. All traffic coming into and emanating from the microservice must pass through the  
486 proxy and the applications running inside of the proxy. Hence, the WASM module that provides  
487 the data protection function deployed inside the proxy will always be invoked.

488 In the ambient proxy implementation model, the network link to a service or group of services  
489 associated with a namespace has to pass through the node hosting the waypoint proxy serving  
490 that service or group of services for a designated namespace. No direct network paths to the  
491 service or group of services exists. Again, the WASM module provides data protection for  
492 services under the scope of the proxy has to be invoked.

493 To meet the second property of the security kernel (i.e., the security is verifiable), a security  
494 analysis of the entire execution environment for the WASM modules must be performed. The  
495 life cycle of a WASM module begins with a source code in some supported language (e.g., C,  
496 C++, or Rust) that is then compiled using a target compiler (e.g., using LLVM) into a binary byte  
497 code that is run by a runtime module (i.e., WASM runtime). Access to the operating system or  
498 host resources is enabled by calling a module that implements an API called WASM System  
499 Interface (WASI).

500 The security analysis of the WebAssembly ecosystem can be considered in terms of the  
501 following topics:

- 502 1. WASM security goals and security feature sets
- 503 2. Memory model and memory safety
- 504 3. Execution model and control flow integrity
- 505 4. Security of API access to OS/host resources
- 506 5. Protection against side-channel attacks
- 507 6. Protection against injection attacks
- 508 7. Deployment and operating safety

#### 509 **4.1. WASM Security Goals and Security Feature Sets**

510 The WASM security model has two important goals: (1) protect *users* from buggy or malicious  
511 modules, and (2) provide *developers* with useful primitives and mitigations for developing safe  
512 applications within the constraints of (1)[8].

#### 513 **4.1.1. User-Level Security Features**

514 Each WASM module executes within a sandboxed environment that is separated from the host  
515 runtime using fault isolation techniques. This implies that:

- 516 • Applications execute independently and cannot escape the sandbox without going  
517 through appropriate APIs.
- 518 • Applications generally execute deterministically with limited exceptions.

519 Additionally, each module is subject to the security policies of its embedding. Within a web  
520 browser, this includes restrictions on information flow through same-origin policy. On a non-  
521 web platform, this could include the POSIX security model.

#### 522 **4.1.2. Security Primitives for Developers**

523 The design of WebAssembly promotes safe programs by eliminating dangerous features from  
524 its execution semantics while maintaining compatibility with programs written for C/C++.  
525 Modules must declare all accessible functions and their associated types at load time, even  
526 when dynamic linking is used. This allows for the implicit enforcement of control-flow  
527 integrity (CFI) through structured control flow. Since compiled code is immutable and not  
528 observable at runtime, WebAssembly programs are protected from control flow hijacking  
529 attacks.

- 530 • Function calls must specify the index of a target that corresponds to a valid entry in  
531 the function index space or table index space.
- 532 • Indirect function calls are subject to a type of signature check at runtime, and the type  
533 signature of the selected indirect function must match the type signature specified at  
534 the call site.
- 535 • A protected call stack that is invulnerable to buffer overflows in the module heap  
536 ensures safe function returns.
- 537 • Branches must point to valid destinations within the enclosing function.

#### 538 **4.2. Memory Model and Memory Safety**

539 As there are only four primary data types defined by WASM, compilers targeting WASM  
540 implement their own stack in an area called linear memory, which becomes the main memory  
541 of a WASM program. A linear memory is a contiguous, byte-addressable range of memory that  
542 can be considered as an untyped array of bytes. This enables the program to store non-scalar  
543 data and any variable whose address needs to be taken by the module [10]. In addition to linear  
544 memory, there is the code space, execution stack, and runtime data structure [11]. The  
545 execution stack mainly stores local variables, global variables, and return addresses.

546 Compilers that target WASM also create an area for the heap in the linear memory. This area is  
547 reserved at the end of the linear memory so that it can dynamically grow when additional space  
548 is allocated for the linear memory. This linear memory is sandboxed — disjointed from code

549 space, execution stack, and runtime data structure [11] — and prevents WASM modules from  
550 accessing other memory areas. These other memory regions are isolated from the internal  
551 memory of the runtime and are set to zero by default unless otherwise initialized. However,  
552 modules can access the data stored on the execution stack via dedicated instructions. The  
553 actual data address on the execution stack is never shown to the module. A compliant runtime  
554 ensures that the module does not break WASM's memory model [12]. This is done by bounds-  
555 checking access to the linear memory at the region level. If the module accesses the memory  
556 outside of the linear memory, the program traps and prevents modules from accessing data  
557 outside of their allocated memory [11].

558 Another common class of memory safety error involves unsafe pointer usage and undefined  
559 behavior. This includes dereferencing pointers to unallocated memory (e.g., NULL) or freed  
560 memory allocations. In WebAssembly, the semantics of pointers have been eliminated for  
561 function calls and variables with a fixed static scope, allowing references to invalid indexes in  
562 any index space to trigger a validation error at load time or — at worst — a trap at runtime.

563 However, the bounds-checking process is performed at the level of the linear memory, and  
564 modules can access the entire linear memory without restriction. Linear memory is not  
565 protected by standard techniques like stack canaries or guard pages. Therefore, buffer  
566 overflows — which occur when data exceeds the boundaries of an object and accesses adjacent  
567 memory regions — cannot affect local or global variables stored in index space. Data stored in  
568 linear memory can also overwrite adjacent objects since bounds-checking is performed at linear  
569 memory region granularity and is not context-sensitive.

#### 570 **4.3. Execution Model and Control Flow Integrity**

571 WASM code is executed when instantiating a module or when an exported function is invoked  
572 on a given instance [12]. The execution behavior of a WASM module is defined in terms of an  
573 abstract machine that models the program state. This abstract machine includes a stack that  
574 records the operand values and control constructs as well as an abstract store that contains the  
575 global state.

576 WASM primarily achieves control flow integrity through the execution semantics of the  
577 language itself. The definition of the WASM bytecode [12] limits the constructs that are  
578 possible to express. It defines valid code constructs and how control flow may only jump to the  
579 beginning of a valid construct. Arbitrary jumps (e.g., goto statements) are not allowed; only  
580 structured control flow is provided. Consequently, a grammatically valid WASM module can  
581 only jump to the beginning of valid constructs (e.g., conditional constructs or functions) [11].

582 An additional factor contributing to the control flow integrity is the prevention of call  
583 redirection through restrictions on indirect function calls. Restrictions are applied regarding  
584 functions that the module can indirectly call. To indirectly call a function, the module provides a  
585 runtime index to a table. This table holds the signatures of the functions that the module  
586 defines or imports and that can be indirectly called. When an indirect call is made, the runtime  
587 checks that the calling signature and the signature of the called function match. If there is a  
588 type mismatch or an out-of-bounds table access, a trap occurs [11].

#### 589 **4.4. Security of API Access to OS and Host Resources**

590 By default, WASM does not have access to the resources of the host (e.g., file system, network,  
591 system calls). Modules can import externally defined functions provided by the host or other  
592 modules. APIs common to many use cases are currently being standardized in the WASI [13].  
593 The capability-based security model of WASI enables the introduction of a verified secure  
594 runtime system, as shown in [14].

#### 595 **4.5. Protection From Side-Channel Attacks**

596 The WASM language specification [12] clearly states that side-channel attacks are to be  
597 addressed by the runtime. Currently, Wasmtime implements a few forms of Spectre  
598 mitigations. Bounds checks for the runtime index used in indirect calls and some other  
599 instructions are mitigated to ensure that speculation goes to a deterministic place [15].  
600 However, some side-channel attacks can occur, such as timing attacks against modules.

601 In the future, additional protections may be provided by runtimes or the toolchain, such as  
602 code diversification or memory randomization like addressing space layout  
603 randomization (ASLR) or bounded pointers (i.e., “fat” pointers).

#### 604 **4.6. Protection Against Code Injection and Other Attacks**

605 Control-flow integrity and protected call stacks prevent direct code injection attacks. Thus,  
606 common mitigations, such as data execution prevention (DEP) and stack smashing  
607 protection (SSP), are not needed by WASM programs. Nevertheless, other classes of bugs are  
608 not obviated by the semantics of WebAssembly. Although attackers cannot perform direct code  
609 injection attacks, it is possible to hijack the control flow of a module using code reuse attacks  
610 against indirect calls. However, conventional return-oriented programming (ROP) attacks using  
611 short sequences of instructions (i.e., “gadgets”) are not possible in WebAssembly because  
612 control-flow integrity ensures that call targets are valid functions declared at load time.  
613 Likewise, race conditions, such as time-of-check to time-of-use (TOCTOU) vulnerabilities, are  
614 possible in WebAssembly since no execution or scheduling guarantees are provided beyond in-  
615 order execution. Yet another security limitation is that there are no audit tools to track the  
616 changes made by WASM modules.

#### 617 **4.7. Deployment and Operating Security**

618 The security features described so far pertaining to run time security. The following capabilities  
619 relate to the controls that are present for deployment and integrity of operations.

- 620 • The ability to create the WASM filter in the proxy can be controlled through the native  
621 access mechanism in the service mesh (e.g., RBAC).
- 622 • Only calls using HTTP and gRPC protocols are allowed.

- 623
- 624
- 625
- Even for making those calls, only clusters known to the proxy can be used. Similarly, responses coming from clusters already known to the proxy are examined.



## 626 **5. Summary and Conclusions**

627 This document describes how WASM modules can be developed and deployed in service mesh  
628 proxies for the real-time protection of data in transit in cloud-native application architectures.  
629 Various data protection techniques can also be used to protect data in different domains of  
630 various application scenarios. WASM modules can provide telemetry data for monitoring tools  
631 that provide visual images of sensitive data flows. A detailed security analysis of the WASM  
632 module development, deployment, and execution environment can ensure that necessary  
633 security assurances are obtained by running the modules as part of the application  
634 infrastructure environment (e.g., in service mesh proxies).

635 The data categorization and protection techniques built into WASM modules must continuously  
636 evolve to keep pace with increasingly sophisticated attacks on data that result in new forms of  
637 data breaches, data leakages, and other forms of data exfiltration.


638

639 **References**

- 640 [1] Doerrfeld B (2023) Wasm: The Next Generation Beyond Kubernetes? Available at  
641 [https://cloudnativenow.com/features/wasm-the-next-generation-beyond-](https://cloudnativenow.com/features/wasm-the-next-generation-beyond-kubernetes/?utm_medium=email&_hsmt=293085224&_hsenc=p2ANqtz-9qIZw1MWD8-AHNH54OoPyWB7vOLe0uG0KZSle2uH1sbXAD_rmmhyXHMThdOGMdMLUb-w8axL_Gv1L2RM9Nq55L2eCysg&utm_content=293086040&utm_source=hs_email)  
642 [kubernetes/?utm\\_medium=email&\\_hsmt=293085224&\\_hsenc=p2ANqtz-9qIZw1MWD8-](https://cloudnativenow.com/features/wasm-the-next-generation-beyond-kubernetes/?utm_medium=email&_hsmt=293085224&_hsenc=p2ANqtz-9qIZw1MWD8-AHNH54OoPyWB7vOLe0uG0KZSle2uH1sbXAD_rmmhyXHMThdOGMdMLUb-w8axL_Gv1L2RM9Nq55L2eCysg&utm_content=293086040&utm_source=hs_email)  
643 [AHNH54OoPyWB7vOLe0uG0KZSle2uH1sbXAD\\_rmmhyXHMThdOGMdMLUb-](https://cloudnativenow.com/features/wasm-the-next-generation-beyond-kubernetes/?utm_medium=email&_hsmt=293085224&_hsenc=p2ANqtz-9qIZw1MWD8-AHNH54OoPyWB7vOLe0uG0KZSle2uH1sbXAD_rmmhyXHMThdOGMdMLUb-w8axL_Gv1L2RM9Nq55L2eCysg&utm_content=293086040&utm_source=hs_email)  
644 [w8axL\\_Gv1L2RM9Nq55L2eCysg&utm\\_content=293086040&utm\\_source=hs\\_email](https://cloudnativenow.com/features/wasm-the-next-generation-beyond-kubernetes/?utm_medium=email&_hsmt=293085224&_hsenc=p2ANqtz-9qIZw1MWD8-AHNH54OoPyWB7vOLe0uG0KZSle2uH1sbXAD_rmmhyXHMThdOGMdMLUb-w8axL_Gv1L2RM9Nq55L2eCysg&utm_content=293086040&utm_source=hs_email)  
645 [2] Krasnov M (2020) Web Assembly is the End of Internet as we know it. Available at  
646 [https://betterprogramming.pub/webassembly-is-the-end-of-the-internet-as-we-know-](https://betterprogramming.pub/webassembly-is-the-end-of-the-internet-as-we-know-it-9085a49cbc7b)  
647 [it-9085a49cbc7b](https://betterprogramming.pub/webassembly-is-the-end-of-the-internet-as-we-know-it-9085a49cbc7b)  
648 [3] WebAssembly (2024) WebAssembly Concepts. Available at  
649 [https://developer.mozilla.org/en-US/docs/WebAssembly/Concepts#see\\_also](https://developer.mozilla.org/en-US/docs/WebAssembly/Concepts#see_also)  
650 [4] TechTarget (2022) Server-side WebAssembly prepares for takeoff in 2023. Available at  
651 [https://www.techtarget.com/searchitoperations/news/252527414/Server-side-](https://www.techtarget.com/searchitoperations/news/252527414/Server-side-WebAssembly-prepares-for-takeoff-in-2023)  
652 [WebAssembly-prepares-for-takeoff-in-2023](https://www.techtarget.com/searchitoperations/news/252527414/Server-side-WebAssembly-prepares-for-takeoff-in-2023)  
653 [5] Medium (2023) WASM and Kubernetes – A new era of application development.  
654 Available at [https://medium.com/@seifeddineraihi/wasm-and-kubernetes-a-new-era-](https://medium.com/@seifeddineraihi/wasm-and-kubernetes-a-new-era-of-cloud-native-application-deployment-b3c59b39f640)  
655 [of-cloud-native-application-deployment-b3c59b39f640](https://medium.com/@seifeddineraihi/wasm-and-kubernetes-a-new-era-of-cloud-native-application-deployment-b3c59b39f640)  
656 [6] Podobnik TJ (2023) WASM Runtimes Vs Containers: Cold Start Delays (Part 1). Available  
657 at [https://levelup.gitconnected.com/wasm-runtimes-vs-containers-performance-](https://levelup.gitconnected.com/wasm-runtimes-vs-containers-performance-evaluation-part-1-454cada7da0b)  
658 [evaluation-part-1-454cada7da0b](https://levelup.gitconnected.com/wasm-runtimes-vs-containers-performance-evaluation-part-1-454cada7da0b)  
659 [7] ITPro (2024) WASM Today, AI Tomorrow: KubeCon Extends its Reach. Available at  
660 [https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-](https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-tomorrow-kubecon-expands-its-reach?mc=NL_DR_EDT_20240401&cid=NL_DR_EDT_20240401&utm_rid=CPNET000059406774&utm_campaign=57716&utm_medium=email&elq2=a6cba5014e5b49bb9a1fe0c3e0351bd2&sp_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378db1b4436de815)  
661 [tomorrow-kubecon-expands-its-](https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-tomorrow-kubecon-expands-its-reach?mc=NL_DR_EDT_20240401&cid=NL_DR_EDT_20240401&utm_rid=CPNET000059406774&utm_campaign=57716&utm_medium=email&elq2=a6cba5014e5b49bb9a1fe0c3e0351bd2&sp_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378db1b4436de815)  
662 [reach? mc=NL\\_DR\\_EDT\\_20240401&cid=NL\\_DR\\_EDT\\_20240401&utm\\_rid=CPNET000](https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-tomorrow-kubecon-expands-its-reach?mc=NL_DR_EDT_20240401&cid=NL_DR_EDT_20240401&utm_rid=CPNET000059406774&utm_campaign=57716&utm_medium=email&elq2=a6cba5014e5b49bb9a1fe0c3e0351bd2&sp_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378db1b4436de815)  
663 [059406774&utm\\_campaign=57716&utm\\_medium=email&elq2=a6cba5014e5b49bb9a1](https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-tomorrow-kubecon-expands-its-reach?mc=NL_DR_EDT_20240401&cid=NL_DR_EDT_20240401&utm_rid=CPNET000059406774&utm_campaign=57716&utm_medium=email&elq2=a6cba5014e5b49bb9a1fe0c3e0351bd2&sp_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378db1b4436de815)  
664 [fe0c3e0351bd2&sp\\_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378d](https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-tomorrow-kubecon-expands-its-reach?mc=NL_DR_EDT_20240401&cid=NL_DR_EDT_20240401&utm_rid=CPNET000059406774&utm_campaign=57716&utm_medium=email&elq2=a6cba5014e5b49bb9a1fe0c3e0351bd2&sp_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378db1b4436de815)  
665 [b1b4436de815](https://www.itprotoday.com/cloud-computing-and-edge-computing/wasm-today-ai-tomorrow-kubecon-expands-its-reach?mc=NL_DR_EDT_20240401&cid=NL_DR_EDT_20240401&utm_rid=CPNET000059406774&utm_campaign=57716&utm_medium=email&elq2=a6cba5014e5b49bb9a1fe0c3e0351bd2&sp_eh=87aea8874bbd0a1985055c93c957744c11570c6718777eca378db1b4436de815)  
666 [8] Security.md (2018) WebAssembly Security. Available at  
667 <https://github.com/WebAssembly/design/blob/main/Security.md>  
668 [9] Huang W, Paradies M (2021) An Evaluation of WebAssembly and eBPF as Offloading  
669 Mechanisms in the Context of Computational Storage. Available at  
670 [https://marcusparadies.github.io/files/ebpf\\_vs\\_wasm\\_report.pdf](https://marcusparadies.github.io/files/ebpf_vs_wasm_report.pdf)  
671 [10] Daniel Lehmann D, Kinder J, and Pradel M. (2020). Everything Old is New Again: Binary  
672 Security of WebAssembly. In USENIX Security  
673 [11] Haas A., et all (2017). Bringing the web up to speed with WebAssembly. In PLDI.  
674 [12] WebAssembly Community Group (2023). WebAssembly Specification. Draft Release 2.0  
675 (Draft 2023-04-24). Available at <https://webassembly.github.io/spec/>  
676 [13] WebAssembly Community Group (2023). WebAssembly System Interface. Available at  
677 <https://github.com/WebAssembly/WASI>  
678 [14] Johnson E., et all (2023). WaVe: A verifiably secure WebAssembly sandboxing runtime.  
679 In Proceedings of IEEE Symposium on Security and Privacy (SP).  
680 [15] Wasmtime (2023). Security - Wasmtime. Available at  
681 <https://docs.wasmtime.dev/security.html>

682 **Appendix A. Execution Model for Web Assembly in Browsers**

683 WASM runtime originated with browsers that enabled the running of native code (i.e., code  
684 written in low-level languages such as C, C++, Rust, etc.).



685

686

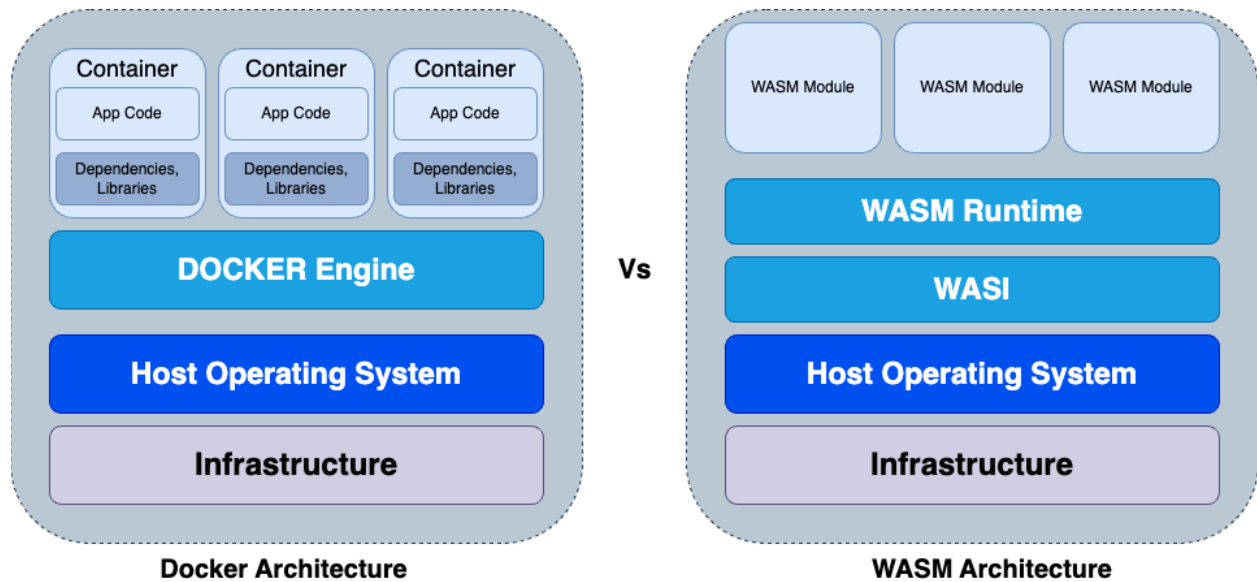
**Fig. 2. WASM Module Development & Execution in Browsers**

687 The WebAssembly program is run through a compiler (also called a WebAssembly target  
688 compiler) that inputs code into an LLVM-compliant language and produces a binary .wasm file.  
689 That file is loaded onto the existing JavaScript code by the JavaScript Interop layer and executed  
690 by the WebAssembly runtime [2]. The .wasm file is a low-level assembly language file in binary  
691 format.

692 The WASM compiler for C, C++, and Rust takes the source code written in those languages and  
693 compiles it into a WASM module. Then the necessary JavaScript “glue” code is generated for  
694 loading and running the module and an HTML document is used to display the results of the  
695 code. The details of this process are explained in [3].

696

697 **Appendix B. Comparison of Execution Models for Containers and WASM Modules**



698

699

**Fig. 3. Comparison of Execution Stack for Containers & WASM Modules**

700 Container images are created by combining the program containing the application logic with  
701 its dependencies (e.g., runtime libraries) in a container runtime (e.g., docker). The container is a  
702 full file system (i.e., utilities, binary), and the generated image should be for a designated OS  
703 kernel and processor architecture (e.g., Intel, Arm, etc.). For example, if a Raspberry Pi OS is  
704 running a docker image, then an image for the C/C++ application based on a Linux image must  
705 be created and compiled for the ARM processor architecture. Otherwise, then container will not  
706 run as expected [5].

707 In contrast, WASM modules and binaries are precompiled C/C++ applications that do not rely  
708 on being coupled with a host OS or system architecture because they do not contain a  
709 precompiled file system or low-level OS primitives. Every directory and system resource is  
710 attached to a WASM module during runtime facilitated by WASI and then run using WASM  
711 runtime. In other words, WASI is used to access all resources under the control of the OS,  
712 essentially decoupling the code from its dependency on the platform architecture.

713