



# ROADRUNNER: Accelerating Data Delivery to WebAssembly-Based Serverless Functions

Cynthia Marcelino

Distributed Systems Group, TU Wien  
Vienna, Austria  
c.marcelino@dsg.tuwien.ac.at

Thomas Pusztai

Distributed Systems Group, TU Wien  
Vienna, Austria  
t.pusztai@dsg.tuwien.ac.at

Stefan Nastic

Distributed Systems Group, TU Wien  
Vienna, Austria  
snastic@dsg.tuwien.ac.at

## ABSTRACT

Serverless computing provides infrastructure management and elastic auto-scaling, therefore reducing operational overhead. By design serverless functions are stateless, which means they typically leverage external remote services to store and exchange data. Transferring data over a network typically involves serialization and deserialization. These operations usually require multiple data copies and transitions between user and kernel space, resulting in overhead from context switching and memory allocation, contributing significantly to increased latency and resource consumption.

To address these issues, we present Roadrunner, a sidecar shim that enables near-zero copy and serialization-free data transfer between WebAssembly-based serverless functions. Roadrunner reduces the multiple copies between user space and kernel space by mapping the function memory and moving the data along a dedicated virtual data hose, bypassing the costly processes of serialization and deserialization. This approach reduces data movement overhead and context switching, achieving near-native latency performance for WebAssembly-based serverless functions. Our experimental results demonstrate that Roadrunner significantly improves the inter-function communication latency from 44% up to 89%, reducing the serialization overhead in 97% of data transfer, and increasing throughput by 69 times compared to state-of-the-art WebAssembly-based serverless functions.

## CCS CONCEPTS

• **Computer systems organization** → **Cloud computing**; • **Applied computing** → **Event-driven architectures**; • **Software and its engineering** → *Middleware*.

## KEYWORDS

Serverless, WebAssembly, Wasm, FaaS, Data transfer, Serialization

### ACM Reference Format:

Cynthia Marcelino, Thomas Pusztai, and Stefan Nastic. 2025. *ROADRUNNER: Accelerating Data Delivery to WebAssembly-Based Serverless Functions*. In *26th International Middleware Conference (MIDDLEWARE '25)*, December 15–19, 2025, Nashville, TN, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3721462.3770777>



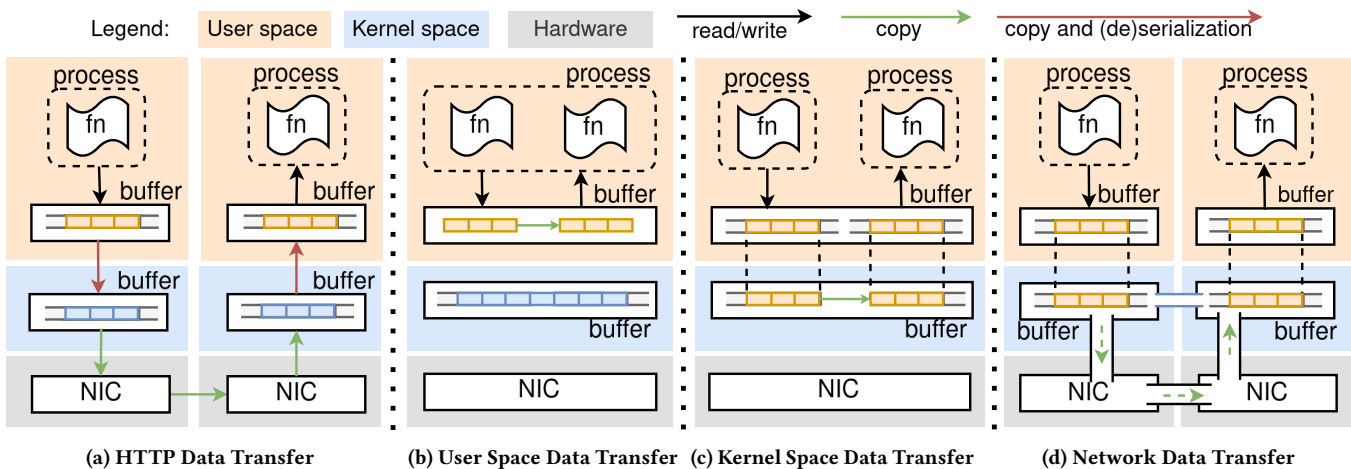
This work is licensed under a Creative Commons Attribution 4.0 International License. *MIDDLEWARE '25*, December 15–19, 2025, Nashville, TN, USA  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1554-9/2025/12.  
<https://doi.org/10.1145/3721462.3770777>

## 1 INTRODUCTION

Serverless functions are short-lived and stateless, which means they often rely on external remote services to store and exchange data [10, 45, 59, 74, 87]. Since serverless functions are short-lived by design, a single function cannot be directly addressed. Therefore, clients rely on the platform ingress and Load Balancers to access the serverless function [4, 18, 26, 46]. WebAssembly (Wasm) has emerged, offering serverless functions with strong isolation, decreased cold starts, and near-native execution speed. Wasm executes small binary modules within a lightweight, memory-safe, and secure sandbox, enabling cross-platform portability while maintaining minimal overhead [11, 21, 22, 50, 58, 75]. Wasm relies on the WebAssembly System Interface (WASI) to interact with the host for essential tasks, such as accessing network interfaces, introducing an additional overhead in wasm-based serverless functions [20, 24, 25, 57]. The WASI dependency for host interactions, including system calls needed for network access, adds extra overhead in serverless functions since these functions typically rely on remote services to retrieve data. These operations lead to multiple context switches and data copies between user and kernel space, increasing function execution time and resource consumption.

In serverless computing, functions typically exchange data via network protocols such as HTTP, which involves serialization of the requested data at the source function and its deserialization at the target (Fig. 1a). Serialization involves converting potentially complex data structures into binary sequences for transfer, followed by deserialization to reconstruct the data structures at the destination. Serialization, along with other low-level operations such as memory allocation, compression, and network stacking, accounts for up to 25% of CPU cycles [8, 30, 65, 69, 89]. During a data transfer, the application in user space prepares the data for the syscall by serializing it into a byte stream. Upon entering kernel space, the syscall copies the byte stream buffers to kernel-space buffers. During this process, the kernel initializes new structures for the transmission, including copying the address, control information, and validating its integrity and permissions. These multiple copies ensure data integrity and security but add overhead due to the repeated copies during context switches [9, 16, 23, 33, 41, 66, 68, 91]. Zero-copy techniques [34, 51, 80] such as RDMA can completely bypass Kernel and CPU, but they require specialized hardware [14, 51, 86]. In Wasm, serialization costs are even higher due to Wasm constraints such as single-threaded execution, which forces the processing of complex tasks to be performed sequentially [24, 53].

Although Wasm offers near-native execution latency, its strict isolation mechanisms introduce additional overhead when interacting with the host through the WASI interface, which can become



**Figure 1: Overview of standard data transfer between Serverless functions and Roadrunner communication models: a) Standard HTTP Data Passing between functions with Serialization/Deserialization, b) User Space Data Transfer, c) Kernel Space Data Transfer between Serverless functions, serialization-free, and d) Serialization-free Network Data Transfer.**

a bottleneck for IO-bound workloads. To address this, we introduce Roadrunner, a novel runtime shim that enables near-zero copy, serialization-free data transfer between functions without requiring specialized hardware such as RDMA-enabled components. Roadrunner establishes a virtual data hose between functions, allowing data to be transferred from memory and bypassing costly serialization and deserialization processes typically required for inter-function communication. Specifically, Roadrunner is tailored for data-intensive edge-cloud scenarios (e.g., ML-based image processing, traffic data analytics), which involve multiple functions that exchange ephemeral data, such as streaming ingestion, frame extraction, data processing, and ML inference. Existing state-of-the-art approaches, such as HTTP communication, either incur high overhead due to WASI-mediated communication or require Wasm runtime modification, compromising Wasm’s isolation guarantees [20, 48, 75]. To address that, Roadrunner introduces novel data-passing abstractions for intra- and inter-function communication that preserve Wasm’s sandboxing by ensuring that only the specific function shim has access to the shared data. In trusted, single-tenant serverless workflows, where functions from the same workflow and tenant are colocated within the same Wasm VM, Roadrunner enables low-latency, near-zero copy communication while maintaining strict isolation. We summarize our contributions as follows:

- *Roadrunner: A novel serverless communication middleware that enables efficient data transfer between Wasm-based serverless functions across user space, kernel space, and the network during runtime. User-space communication (Fig. 1b) operates directly on Wasm linear memory for the Wasm VM with multiple Wasm modules. Kernel-space communication (Fig. 1c) enables co-located functions to exchange data via host mechanisms, skipping network transfer. Network communication (Fig. 1d) establishes lightweight data transfer*

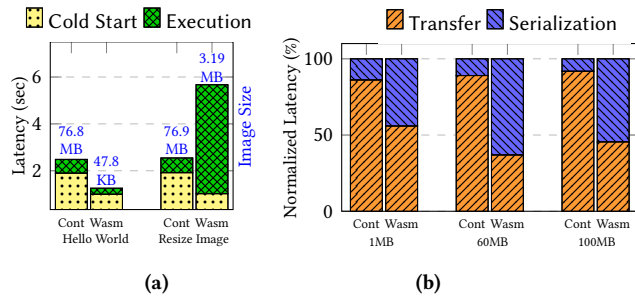
for inter-node functions. This architecture achieves low latency, therefore maximizing the function performance for both local and remote function data exchange.

- *A near-zero copy mechanism for inter-function communication that eliminates multiple data copies between user space and kernel space without requiring specialized hardware such as RDMA-enabled components. Roadrunner establishes a virtual data hose that allows data written to it to prompt the kernel to allocate memory buffers and retain them in its address space. When a read operation occurs, Roadrunner leverages the kernel to reuse the same memory pages for the target function instead of copying the data, thereby avoiding duplication. It reduces intermediate memory copies, minimizes resource overhead, and ensures efficient data transfer between functions. Roadrunner improves the latency by 44% to 89%, increasing the throughput by 69 times compared to the state-of-the-art Wasm runtime.*
- *A serialization-free data transfer mechanism that operates directly on Wasm VM linear memory and manages memory allocation dynamically. Roadrunner allocates memory in the target function to store the incoming data and accesses the memory region specified by the source function using direct memory pointers. The data is transferred by referencing these memory regions without intermediate serialization or deserialization. By directly handling memory allocation and data transfer at the raw binary level, Roadrunner eliminates unnecessary data copies and context switches. Roadrunner reduces the serialization impact by 97% during data transfer between wasm serverless functions.*

## 2 BACKGROUND AND MOTIVATION

### 2.1 WebAssembly in Serverless Computing

*Portability.* Wasm is a portable, low-level binary instruction format designed as a compilation target for high-level languages such



**Figure 2: Cold start and execution latency (a), and normalized I/O breakdown (b) for functions using Docker Containers (Cont) and Wasm. (a) shows Wasm lowers cold starts for both “Hello World” (no WASI) and “Resize Image” (with WASI), though WASI increases execution time. (b) compares transfer and serialization overhead across various input sizes.**

as Rust, C, C++, and Go. Wasm modules are platform-agnostic and can be executed on different architectures (e.g., x86, arm, and aarch64) [44, 61, 92]. Wasm portability contrasts with traditional serverless platforms, which execute functions within OCI bundles (containers) using runtimes such as RunC [29] and Crun [15], which rely on the host kernel and Linux primitives such as cgroups and namespaces for isolation [13, 35, 36, 40, 67], running directly on the host OS. In Wasm, the function code is compiled into small binaries and executed in a secure, sandboxed environment (i.e., Wasm VM) and, therefore, more lightweight compared to Docker images, as they do not require a base image operating system (OS) or extensive initialization provision [20, 25, 61].

**Linear Memory.** Wasm uses a linear memory model, where memory is represented as a contiguous, byte-addressable array. The allocated memory for the Wasm VM can dynamically grow, providing isolated and efficient access for computation between Wasm modules, enabling safe memory management and isolation between the module and host environment [11, 24, 44].

**WASI Overhead.** Wasm follows the deny-by-default principle, which means Wasm binaries do not have access to the host machine unless explicitly requested. To access host functionalities such as the file system and network interfaces, Wasm leverages WASI [20, 21, 58, 75]. Fig. 2a presents the effect of WASI-based host communication on function execution time. Wasm functions exhibit reduced cold start latency, with Wasm binaries of 3.19MB, whereas Docker images have approximately 77 MB. During execution, functions that do not require WASI access achieve lower latency than their Docker container counterparts. In contrast, functions interacting with the host file system through WASI exhibit increased execution times.

## 2.2 Data Passing between Serverless Functions

To exchange data, serverless functions typically rely on external remote services. While using remote services decouples function computation from I/O operations, allowing high scalability, it also introduces additional overhead [19, 49, 70, 73]. The most common approaches to exchange data between serverless functions are: (a) *Remote services* such as object storage [1], KVS [6, 38, 88],

cache [28, 62, 71, 82], distributed storage [5, 52], are commonly used to enable direct communication between serverless functions. However, remote services might increase latency and network overhead. (b) *Direct Communication* such as message queues [2, 13], TCP-punch hole [12], disk storage [52], and local cache [2, 7, 47, 63, 78] allows serverless functions to leverage host mechanisms to exchange data, decreasing latency. Nevertheless, these approaches decrease the function isolation [57, 79]. (c) *Zero-copy communication* such as shared memory [31, 39, 67, 76] enables co-located functions to exchange data from a specific memory region. Nevertheless, it increases development complexity to ensure that only allowed functions can access the data in the shared memory. RDMA [14, 37, 80, 86, 91] enables direct memory access between nodes over the network by skipping CPU and kernel, and directly processing the data by using Network Interface Cards (NICs) to read or write directly to the memory regions specified by functions. However, RDMA usage requires specialized hardware. Moreover, existing approaches such as [2, 31, 52, 75] rely on their own scheduling mechanisms to colocate functions to enable leveraging their data passing optimization, and local Inter Process Communication (IPC) syscalls. Roadrunner optimizes communication regardless of the scheduler’s decisions.

**Wasm-based Serverless Functions.** Typically, serverless functions rely on remote services over protocols such as HTTP, which involves serialization. Although the Wasm design enforces strict isolation, the WASI dependency for host interactions introduces overhead, particularly during data serialization [48, 56, 57]. In Wasm, serialization involves converting complex data structures within the Wasm VM into a linear, standardized format, allocating memory for the serialized output, and copying the data across the Wasm VM boundary, leading to increased execution latency and higher CPU and memory overhead. We further investigate the WASI overhead in Fig. 2b, analyzing the impact of serialization on data transfer in the state-of-the-art serverless Docker container runtime and the Wasm runtime. In our preliminary experiments, we observed that serialization accounts for up to 15% of the execution time of data transfer in Docker runtime, while in Wasm, it contributes up to 60%, significantly impacting the performance of Wasm-based serverless functions. This challenge becomes even more significant for functions that require frequent data access. Therefore, optimizing the transfer of data between functions is crucial to reduce execution time.

## 3 ROADRUNNER DATA ACCESS MODEL & SYSTEM DESIGN


### 3.1 Roadrunner Data Access Model

**Overview.** Roadrunner is designed to optimize Serverless inter-function communication and data passing. Roadrunner supports three communication modes: User space, Kernel Space, and Network, which are integrated with existing serverless platforms through a lightweight shim that runs beside each function. Roadrunner relies on components of serverless platforms that enable the deployment, execution, and observability of a Serverless workflow.

**Wasm-Based Data Access.** Roadrunner enables data access within the Wasm VM by leveraging a combination of direct memory access and controlled memory management functions shown in Fig. 3.

**Table 1: Roadrunner Data Access Function APIs**

Function	Description	Category	Location
void allocate_memory(len)	Allocates a specified amount of linear memory in the Wasm VM	Memory Management	Function
void deallocate_memory(address)	Deallocates a specified amount of linear memory in the Wasm VM	Memory Management	Function
void read_memory_wasm(address, len)	Reads data from a specified memory address and length in the Wasm VM	Data Management	Function
(int,int) locate_memory_region(data)	Returns the memory pointer and length of the specified data	Data Management	Function
void send_to_host(address, len)	Transfers data memory information to the host interface using the address and length	Data Management	Function
void read_memory_host(address, len)	Reads the data from the Wasm VM memory using a specified address and length	Data Management	Shim
void write_memory_host(data[], address)	Write the data into Wasm VM	Data Management	Shim

**Figure 3: Roadrunner data access structured layers for data interactions on the left side and linear memory mapping of Wasm VM for addressable data access on the right side.**

Within the Wasm VM, linear memory is exposed as a contiguous block of memory and accessible through specific offsets to the host. Roadrunner utilizes this memory model to access the data stored directly in the Wasm VM sandbox. When the host needs to access data, it leverages memory pointers provided by the Wasm module to locate specific memory regions without breaking the isolation of the Wasm VM.

**Complex Data Types.** Wasm supports only primitive types (i32, i64, f32, f64), making it highly optimized for low-level operations but requiring additional mechanisms to handle complex data types such as strings. Languages such as C and Rust align well with this model, allowing direct memory management. In contrast, high-level languages such as Java and Python require embedded language runtimes within the Wasm module to manage memory, preventing direct access to linear memory and thus limiting compatibility with Roadrunner [24, 27, 42, 43, 77]. Therefore, the APIs listed in Table 1 are designed for low-level languages to interact with Wasm memory directly and efficiently represent complex data structures such as strings. By directly using pointers to a linear memory region, we can achieve a serialization-free copy. This serialization-free approach is valid under the assumption that all systems use the same endianness, e.g., the little-endian format found in x86 and ARM [3], which dominate edge and cloud environments. Furthermore, Wasm explicitly defines integer sizes (i32, i64), ensuring consistent behavior on 32-bit and 64-bit CPUs and allowing Roadrunner’s pointer-based, serialization-free access to remain valid across different architectures, as long as endianness is consistent.

**Near-Zero Copy Data Exchange.** When data is passed between the shim and the Wasm VM, the shim writes or reads directly from specific memory regions, bypassing traditional intermediary steps such as file I/O or network communication. Roadrunner coordinates memory allocation for the data and data transfer through the Wasm runtime memory interface. The data is written or read directly in

the allocated memory regions. This ensures that data never leaves the Wasm VM unnecessarily, reducing the need for intermediate serialization or external storage. This approach minimizes overhead while maintaining the security and isolation properties intrinsic to Wasm.

**Shared Memory.** While Roadrunner enables shared memory access between Wasm modules, it enforces strict isolation through the shim. Functions never directly access shared memory; instead, all memory operations are mediated by the shim, which manages allocation, access control, and memory mapping. Only functions of the same workflow and tenant are instantiated in the same Wasm VM, which has control over each function’s linear memory space. The shim guarantees that memory visibility is limited to registered functions in the same execution context. To prevent unauthorized access and cross-tenant interference, Roadrunner restricts shim-to-Wasm access to pre-registered memory regions and applies bounds checking before any read or write operation.

## 3.2 Roadrunner System Design

**3.2.1 Lightweight.** Roadrunner uses the compact binary format of Wasm to reduce overhead during function execution. Wasm modules are considerably smaller and more lightweight than traditional container images, such as those used with Docker, which usually require a base image OS and extensive initialization processes, leading to longer cold starts. This design enables quick execution while ensuring compatibility with existing serverless platforms, maintaining low overhead and high performance.

**3.2.2 Function Isolation.** Roadrunner encapsulates each Wasm VM in an OCI-compliant runtime bundle, enabling interoperability with container runtime managers such as containerd. The shim runs as a sidecar alongside each function and manages the Wasm VM lifecycle, including memory configuration, binary loading, and runtime interaction. It handles all function ingress and egress, making data available to the functions via predefined APIs that expose input and output operations, allowing seamless integration with state-of-the-art orchestrators (e.g., Kubernetes) and serverless platforms (e.g., Knative, OpenFaaS, and OpenWhisk) without modifying scheduling, logging, or networking layers.

**3.2.3 Inter-function communication.** Roadrunner improves communication between serverless functions using near-zero copy and serialization-free data transfer mechanisms. It achieves this through direct memory mapping and system-level operations, which reduce the number of context switches between user space and kernel space. Roadrunner supports three modes of data exchange: user

space (Fig. 1b), kernel space (Fig. 1c), and network-based communication (Fig. 1d). These modes allow for optimized data transfers based on the proximity of functions, reducing dependency on remote services and minimizing unnecessary network overhead. By enabling direct communication, Roadrunner decreases latency, supports complex function workflows, and ensures high throughput for applications that require low latency.

**3.2.4 Scalability.** Roadrunner is a shim-based sidecar that lives alongside serverless functions, allowing the container orchestration tool to manage scalability. This design enables Roadrunner to optimize the data transfer while relying on the orchestrator to handle resource allocation, function placement, and horizontal scaling. Thus, Roadrunner seamlessly integrates with these platforms without adding extra complexity to scaling. The sidecar shim approach keeps Roadrunner lightweight and interoperable, allowing the orchestrator to manage increased workloads while maintaining overall system performance efficiently.

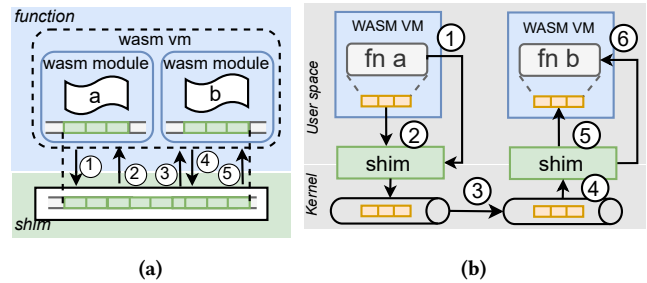
**3.2.5 Function Deployment and Lifecycle.** The Roadrunner shim plays a crucial role in managing the lifecycle of Wasm functions. During initialization, the shim creates a dedicated Wasm VM and configures the Wasm runtime, which includes setting resource limits such as memory. The function binary is then loaded into the isolated memory space of the Wasm VM. To ensure compatibility with standard serverless platforms, the shim packages the Wasm VM as an OCI-compliant bundle. This allows it to be executed as a container by high-level container managers such as containerd, facilitating seamless integration between orchestrators (such as Kubernetes) and Roadrunner. Consequently, Wasm-based functions can utilize the advantages of Roadrunner while using standard container orchestration tools for deployment and execution. This integration simplifies the complexities of Wasm execution and takes advantage of the existing serverless infrastructure.

## 4 ROADRUNNER DATA TRANSFER MECHANISMS

In this section, we introduce the data transfer mechanisms employed by Roadrunner, including user-space, kernel-space, and network-based communication mechanisms.

### 4.1 User Space Data Transfer

In the user-space data transfer between functions, shown in Fig. 4a, functions share the same isolation sandbox, which allows them to be executed within a single process. By sharing the same address space, functions can directly access each other's memory without requiring additional memory copies. This eliminates operations such as duplicating data between different memory regions and context switches between user space and kernel space. As a result, memory access is faster, leading to lower latency and higher throughput in the communication between functions. Roadrunner leverages the shim to create and start the Wasm VM and manage the memory access, which enables read/write data in the function. User-mode communication requires explicit trust; thus users must specify functions within a workflow, which the shim validates using the containerd snapshot. Then, the shim initializes the VM and loads both, *function a* and *function b*, as Wasm modules.




**Figure 4:** a) User-space memory-based data transfer between Wasm serverless functions (modules) using Roadrunner's APIs for allocation, reading, and writing into Wasm linear memory. b) Kernel Space Data transfer where the shim leverages IPC to enable serialization-free data transfer.

When *function a* needs to send data to *function b*, ① *function a* first calls `locate_memory_region(data)` to get the memory pointer and size of the data and forwards them to the shim. In ②, the shim reads it with `read_output(address, len)` via Wasm VM memory APIs. Next, in ③, after reading the data, the shim calls `allocate_memory(len)` to reserve memory in the *function b* for the incoming data. In ④, the shim receives the address of the allocated memory. Finally, in ⑤, the shim calls `write_output(address, len)` in *function b* with the incoming data from *function a*. This process enables efficient and secure data exchange between the two functions within the same Wasm VM, leveraging the shim to abstract memory management and streamline inter-function communication keeping the data movement within the user space.

### 4.2 Kernel Space Data Transfer

In the kernel-space data transfer mechanism, shown in Fig. 4b, functions that are co-located on the same host but in different sandboxes, i.e., each function has its own dedicated shim, and they exchange data via a kernel buffer, avoiding network transfers. Roadrunner enables quick data passing between functions by leveraging host mechanisms that allow data from one sandbox to be transferred directly via a dedicated pipe without serialization and context-switching.

The data transfer process begins in ①, when *function a* invokes the Roadrunner API `locate_memory_region(data)` to determine the memory pointer and size of the data it wants to send. This information is forwarded to the shim managing *function a*. Next in ②, the shim reads the data using `read_output(address, len)` from *function a's* memory. Then, the shim leverages IPC mechanisms to send the data to the shim managing *function b*, in ③. Upon receiving the data in ④, the shim for *function b* calls `allocate_memory(len)` to reserve memory within *function b's* memory space, in ⑤. Finally, the shim writes the data to *function b's* memory using `write_output(address, len)` in ⑥. This kernel-mediated process ensures secure and efficient data transfer between different Wasm VM instances through IPC, leveraging the shim as a bridge between the user-space functions and the kernel.



**Figure 5: Network-based data transfer with near-zero copy mechanisms, mapping user-space memory into kernel buffers to optimize throughput and minimize latency.**

### 4.3 Network Data Transfer

Roadrunner introduces serialization-free mechanisms that enable remote functions to pass data via a dedicated network buffer. Similar to the kernel buffer, Roadrunner creates a virtual data hose between the two hosts where the data in the user space can be directly mapped into this dedicated virtual data hose. Although the application data is located in user space, the Kernel maps it directly from the user space into the virtual data hose, thus avoiding serialization and unnecessary copies. Unlike zero-copy methods such as RDMA, which completely skips the kernel and CPU, the data passing is processed by the kernel, which means the CPU is responsible for the data movement tasks, but the data is mapped directly from the user space memory. Hence, there is no copying between the user and kernel space.

In the network transfer process, the shim in the source Wasm VM first retrieves data from the Wasm VM’s memory space and prepares it for transmission. It begins with ①, where the shim in the source Wasm VM retrieves data from the memory space of the running function using the WasmEdge runtime APIs. In ②, the shim prepares the data and transfers it to the kernel using a syscall for near-zero copy data movement. Next, in ③, the kernel buffers the data and directs it towards the network stack. At this stage, in ④, the kernel passes the data to the network interface card (NIC), which handles the physical transmission over the network in ⑤.

On the target node, the process continues with ⑥, where the NIC receives the data and places it into the kernel buffer. The kernel then prepares this data for further processing. In ⑦, the shim on the target host uses a syscall to move the data from the kernel buffer back to user space with near-zero copy efficiency in ⑧. Finally, the shim invokes the WasmEdge APIs to allocate and write the data into the memory of the target Wasm VM in ⑨, completing the transfer. The data is now available for the function in the target Wasm VM to process. Algorithm 1 provides a pseudo-code representation of these steps, offering an abstraction for implementation.

## 5 ROADRUNNER IMPLEMENTATION

Roadrunner is published as an open-source framework part of the Polaris project. It is implemented in Rust and available on Github<sup>1</sup>. Roadrunner follows OCI standards, which means it can be used in any of the state-of-the-art serverless platforms.

<sup>1</sup><https://github.com/polaris-slo-cloud/roadrunner>

### Algorithm 1 Roadrunner Network Data Transfer

```

1: function FUNCTIONA(input, target)
2:   (address, length) ← LOCATE_MEMORY_REGION(input)
3:   SEND_TO_HOST(address, length)
4: end function
5:
6: function NETWORK_DATA_TRANSFER_SOURCE(target)
7:   data ← READ_MEMORY_HOST(address, length)
8:   vdh ← CREATE_VIRTUAL_DATA_HOSE()
9:   VMSPLICE(vdh, address, length)
10:  socket_target ← CREATE_SOCKET(target)
11:  SPLICE(vdh, socket_target, length)
12:  CLOSE_ALL()
13: end function
14:
15: function FUNCTIONB
16:   ALLOCATE_MEMORY(len)
17:   DEALLOCATE_MEMORY(address)
18:   input ← READ_MEMORY_WASM(address, length)
19: end function
20:
21: function NETWORK_DATA_TRANSFER_TARGET
22:  target_vdh ← CREATE_VIRTUAL_DATA_HOSE()
23:  SPLICE(socket_fd, target_vdh, length)
24:  target_memory_address ← ALLOCATE_MEMORY(length)
25:  VMSPLICE(target_vdh, target_memory_address, length)
26:  WRITE_MEMORY_HOST(data, target_memory_address)
27:  CLOSE_ALL()
28:  DEALLOCATE_MEMORY(target_memory_address)
29: end function

```

For user-space data transfer, we utilize the WasmEdge runtime APIs to directly access the linear memory space of Wasm modules. This allows efficient read and write operations within the same process, avoiding unnecessary data copies and reducing overhead.

For kernel-space data transfer, we employ Unix sockets as IPC mechanism. They provide a reliable and low-latency communication channel between processes, leveraging the kernel to handle data transfer efficiently. We utilize Unix sockets to enable seamless data exchange between different serverless functions while maintaining isolation and security.

For network-based data transfer, we utilize the `splice` [54] and `vmsplice` [55] syscalls. These Linux-specific calls enable zero-copy data transfer between file descriptors, significantly reducing the overhead of data movement across network boundaries. Since `splice` and `vmsplice` bypass the user-space buffer, they allow direct data transfer between kernel buffers and function memory, improving throughput and reducing CPU utilization during network communication.

## 6 EVALUATION

To assess the performance, scalability, and resource efficiency of Roadrunner, we design a series of experiments focusing on data-intensive serverless functions. The evaluation aims to determine

how well Roadrunner optimizes inter-function communication compared to state-of-the-art runtimes. As Roadrunner introduces measurements for intra-node, i.e., user and kernel space data transfer, and inter-node communication, i.e., network data transfer, we present experiments evaluating its performance in both scenarios.

## 6.1 Overview

We perform two experiments to evaluate Roadrunner, using chained serverless workflows consisting of two I/O-bound functions,  $a$  and  $b$ , which exchange serialized strings. The first experiment varies payload sizes (1MB–500MB) to assess communication overhead during sequential execution, while the second evaluates scalability through parallel workflows. These payloads reflect structured data commonly exchanged between serverless functions, with sizes representing edge–cloud serverless workflows [18]. In contrast, the second set assesses scalability through parallel executions. These workflows reflect real-world serverless invocation patterns, such as sequential, fan-out, and fan-in as outlined in [32]. RoadRunner does not aim to provide better performance than serverless state-of-the-art solutions, such as [2, 13], as these solutions operate directly on the host kernel while Roadrunner suffers from Wasm VM IO overhead (shown in Fig. 6 and discussed in Section 6.3). Thus, we compare against RunC (container) as an upper bound for performance, representing the best achievable performance with Wasm. Therefore, we compare Roadrunner with state-of-the-art virtualization runtimes: RunC and WasmEdge. Since Roadrunner aims to optimize data communication, we measure the latency from the moment the source function sends data until the target function receives it.

*Metrics.* Data transfer is crucial to ensure efficient serverless, naturally stateless functions. It is critical to enable seamless communication at scale with minimal latency while avoiding exhausting computing resources. To evaluate the performance, scalability, and resource impact of Roadrunner, we conduct a series of experiments and collect results for the following metrics:

- a) *Latency:* This metric measures the duration from when function  $a$  initiates the data transfer to when function  $b$  has successfully received the message. Latency is recorded in seconds to provide a clear understanding of communication delays.
- b) *Throughput:* This metric evaluates Roadrunner’s ability to handle high workloads without compromising performance. Throughput is measured in requests per second (RPS). For operations completed in less than one second, we extrapolate throughput by calculating the rate of requests over one second. For example, if ten requests are processed in under a second, we estimate the throughput based on their collective execution rate.
- c) *Resource Usage:* For this metric, we monitor the host system’s resource utilization for each function. Specifically, we track CPU usage (in percent) and memory consumption (in MB). We perform fine-grained measurements directly from the cgroup, enabling us to accurately capture the total CPU usage for each sandbox (container), including detailed breakdowns of user space and kernel CPU consumption. This metric allows us to compare the resource overhead introduced by Roadrunner against baseline metrics, ensuring that any increase remains minimal and within acceptable limits.

## 6.2 Experimental Setup

Our experiments focus on evaluating chained serverless functions in a distributed environment. To execute these experiments, we employ RunC [29] as the low-level container runtime, containerd as the high-level container manager, and WasmEdge [83] as the Wasm runtime. We use multiple nodes to facilitate remote data passing between functions. We executed the experiments 10 times and calculated the mean results. We deploy two nodes to enable remote data exchange between functions, emulating an edge–cloud setup. To approximate constrained edge–cloud conditions, we limit the available bandwidth to 100 Mbps using Linux traffic control (tc), and observe a stable round-trip latency of 1 ms between nodes. Each node is a VM with a 4-core 2GHz vCPU, 8GB of RAM, running Ubuntu 22.04 LTS. The underlying server is running an Intel Xeon processor of the Skylake generation. The experimental workflows are implemented in Rust and leverage WasmEdge SDK for WASI-related tasks and tokio [17] for HTTP requests.

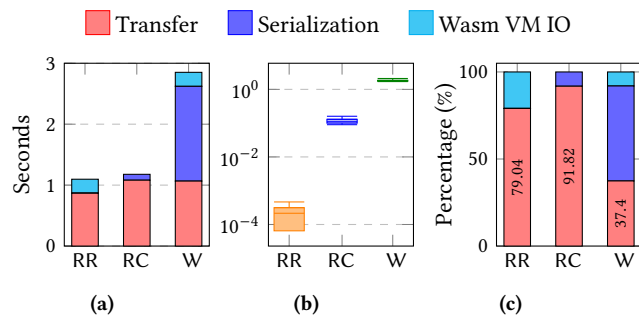
## 6.3 Performance

***Intra-node: Roadrunner (User space) reduces the latency by 44% to 89% compared to WasmEdge and by 10% to 80% compared to RunC. Roadrunner (Kernel space) reduces the latency by 76% to 83% compared to WasmEdge and up to 13% compared to RunC.***

Fig. 7 shows the performance metrics, including Total and Serialization Latency for different payload sizes across three runtimes: Roadrunner, and WasmEdge. Fig. 7a and Fig. 7b represent the total latency and throughput, respectively, including data transfer and data serialization times. Roadrunner (User space) achieves the lowest latency due to near-zero copy transfers, followed by Roadrunner (Kernel space). RunC and WasmEdge exhibit significantly higher latencies. Fig. 7c and Fig. 7d show exclusively the serialization impact on the function data transfer. Roadrunner (User space) and RunC demonstrate negligible latency. Here, it is possible to notice the serialization impact on the WasmEdge function, which is responsible for the high total latency presented in Fig. 7a. Roadrunner (User space) and Roadrunner (Kernel space) significantly reduce both total and serialization latency compared to RunC and WasmEdge, highlighting the impact of near-zero copy and serialization-free transfers on improving function communication performance.

***Inter-node: Roadrunner reduces the total latency by 62% compared to WasmEdge and 7% compared to RunC. For serialization latency, Roadrunner reduces the serialization overhead by 97% compared to WasmEdge and 46% compared to RunC.***

Fig. 6 highlights the gains from Roadrunner compared to RunC and WasmEdge at a fixed size. Fig. 6a shows the total latency contributions from data transfer, serialization, and Wasm VM IO overhead, where Wasm VM IO is the overhead Roadrunner takes to access the data from the host via the data access mechanisms presented in Section 3.1. Although Roadrunner presents gains in the transfer compared to RunC and WasmEdge, it pays a penalty (Wasm IO) to access the data in the Wasm VM. However, due to the near-zero copy and serialization-free improvements, shown in Fig. 6b, Roadrunner offers near-native speed to the Wasm-based functions, reducing the serialization overhead introduced by state-of-the-art



**Figure 6: Breakdown of inter-node transfer latency for a 100MB payload across Roadrunner (RR), RunC (RC), and WasmEdge (W): (a) Detailed analysis of latency components, including serialization, communication transfer, and Wasm IO overhead; (b) Serialization overhead comparison, highlighting Roadrunner’s near-zero copy and serialization-free mechanisms versus HTTP-based communication in RunC and WasmEdge; and (c) Normalized latency distribution, showcasing Roadrunner’s efficiency in reducing serialization impact while maintaining competitive overall latency.**

communication mechanisms such as HTTP. In Fig. 6c, the normalized latency demonstrates that Roadrunner significantly reduces serialization overhead compared to RunC and WasmEdge. As a result, the overall latency in Roadrunner approaches that of RunC, where the majority of the latency is attributed to network transfer. As Roadrunner reduces overall latency, the relative share of Wasm IO overhead increases from 8% in RunC to 20% in Roadrunner.

Fig. 8 shows the varying input size for Roadrunner and baselines. Fig. 8a and Fig. 8b depict total latency and throughput, incorporating both data transfer and serialization overhead. Roadrunner demonstrates a reduction in total latency compared to RunC and WasmEdge, primarily due to its near-zero copy and serialization-free data transfer mechanism. However, the improvement margin is narrower in the inter-node scenario, as network transfer overheads become a dominant factor. Fig. 8c and Fig. 8d focus on serialization overheads and reveal that Roadrunner significantly minimizes serialization latency. Although, in inter-node data transfer, the network contributes to the majority of the latency, Roadrunner offers Wasm-based serverless latency similar to RunC.

## 6.4 Scalability

**Intra-node: Roadrunner (User space) reduces the latency up to 70% compared to RunC and increases the throughput up to 3x; reduces up to 98% compared to WasmEdge and increases the throughput up to 64x. Roadrunner (Kernel space) reduces up to 77% compared to WasmEdge and increases the throughput up to 4x.**

Fig. 9 presents the performance metrics for intra-node fanout scenarios with 10MB function input transfers. Fig. 9a and Fig. 9b present the total latency and throughput across different fanout degrees for Roadrunner (User space), Roadrunner (Kernel space), RunC, and WasmEdge. Roadrunner (User space) demonstrates the lowest total latency due to its efficient near-zero copy mechanism.

Roadrunner (Kernel space) exhibits slightly higher latency than RunC due to the overhead costs imposed by IPC in combination with async libraries to handle the fanout degrees. However, both significantly outperform WasmEdge, which suffers from high serialization costs and network transfer overhead.

Fig. 9c and Fig. 9d isolate the serialization impact, where Roadrunner (User space) and Roadrunner (Kernel space) maintain minimal serialization latency and high throughput, respectively. WasmEdge exhibits the highest serialization latency, which dominates its overall performance. RunC’s serialization latency is higher than Roadrunner but still lower than WasmEdge.

**Inter-node: Roadrunner reduces the latency by up to 65% compared to WasmEdge and increases the throughput up to 2.8x.**

Fig. 10 shows scalability results increasing the fanout degree. Similar as intra-node results, Fig. 10a and Fig. 10b depict the total latency and throughput across varying fanout degrees for Roadrunner, RunC, and WasmEdge. Roadrunner maintains a relatively consistent latency profile across fanout degrees similar to RunC and WasmEdge. Further, it shows a slight increase compared to RunC, indicating the IPC and async library overhead. Roadrunner outperforms WasmEdge, which shows significantly higher latencies due to substantial serialization costs. Fig. 10c further isolates the serialization impact, emphasizing the advantages of Roadrunner near-zero copy and serialization-free mechanism. Meanwhile, Fig. 10d shows Roadrunner consistently outperforming WasmEdge and RunC achieving increased throughput.

## 6.5 Resource Efficiency

**Intra-node: Roadrunner reduces the CPU usage by up to 94% compared to WasmEdge and reduces the RAM usage by up to 50%**

During performance experiments that measure latency and throughput shown in Fig. 7, Roadrunner exhibits significantly lower CPU usage compared to WasmEdge in Fig. 7e. More specifically, the lack of serialization and multiple copies between user and kernel space, Roadrunner decreases the user space CPU usage, shown in Fig. 7f. Fig. 7g shows Kernel Space CPU usage, where Roadrunner keeps consistently close to RunC, presenting a decrease compared to WasmEdge.

During the scalability experiments, in Fig. 9, Roadrunner shows a constant CPU usage even with the fanout degree increase in Fig. 9e. Although Roadrunner handles the data transfer alone in the user space, functions *a* and *b* are embedded in one single Wasm VM, which means it runs in a single process. In this case, the CPU handles both computation and data transfer, leading to an increased CPU burden, especially in the user space CPU usage in Fig. 9f. Nevertheless, Roadrunner (User space) shows consistent and stable usage even for increased fanout degree. On the other hand, Roadrunner (Kernel space) IPC mechanisms result in a decreased user and kernel space CPU usage, shown respectively in Fig. 9f and Fig. 9g, resulting in an overall decrease in CPU usage (Fig. 9e). The trend continues for RAM usage in Fig. 9h where the usage increases with the fanout degree and Roadrunner (Kernel space) outforms RunC and WasmEdge while Roadrunner (User space) shows a steady improvement compared to WasmEdge.

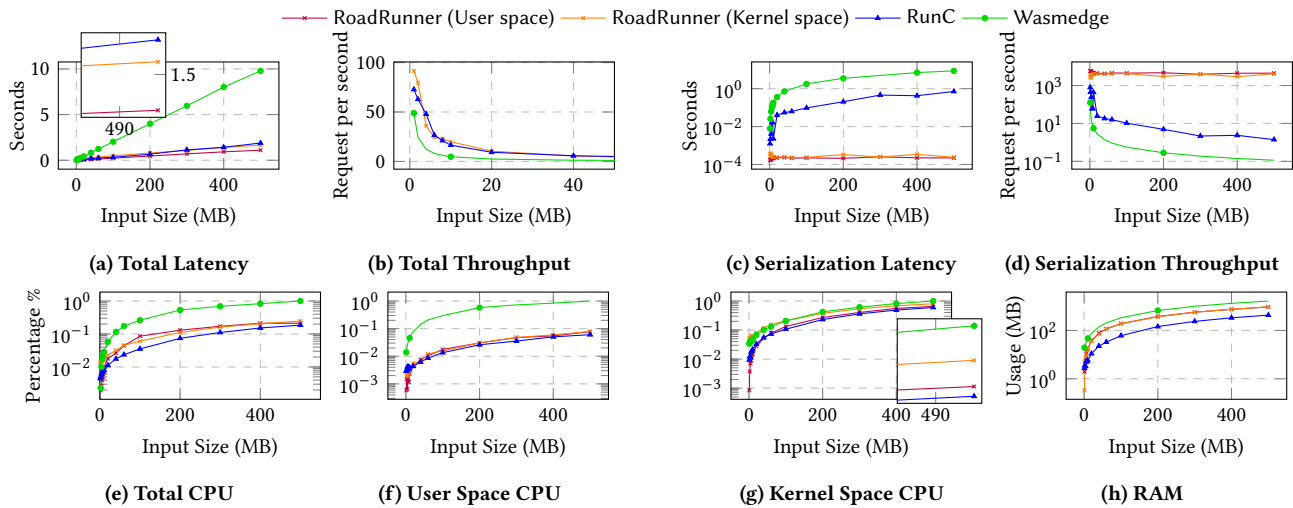


Figure 7: Intra-node overall latency and throughput for varying payload sizes, comparing RoadRunner, RunC (Container) and Wasmedge runtimes. (a) Total latency, including serialization, in seconds. (b) Total throughput, including serialization, in requests per second. (c) Serialization latency, isolating the impact of serialization overhead. (d) Serialization throughput, measuring throughput excluding serialization overhead. (e) Total CPU usage percentage. (f) User space CPU usage percentage. (g) Kernel space CPU usage percentage. (h) Memory usage in MB, illustrating resource consumption for each runtime.

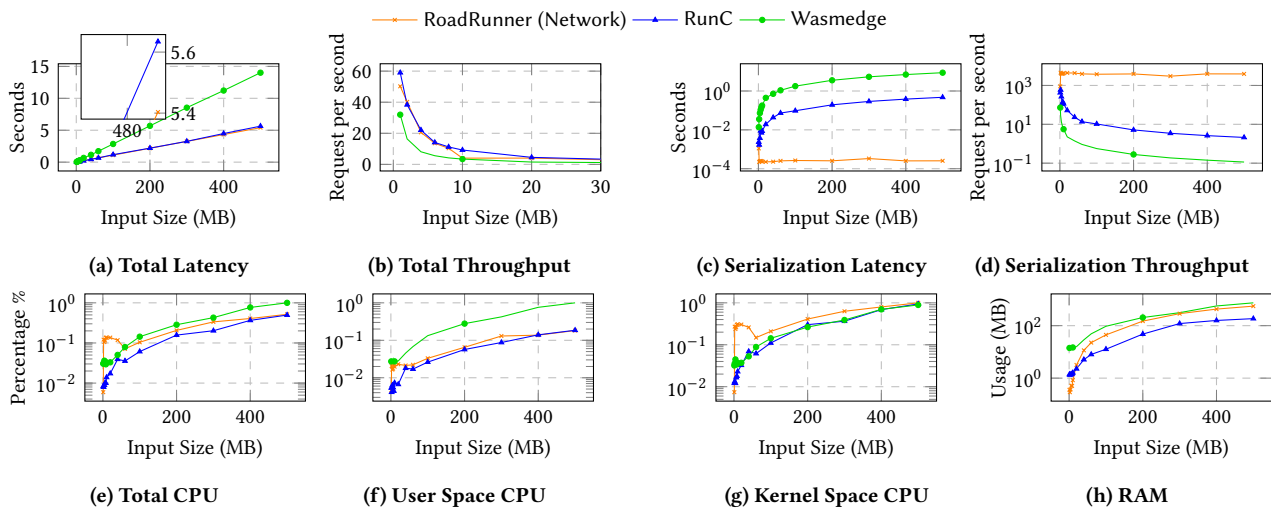


Figure 8: Inter-node overall latency and throughput for varying payload sizes, comparing RoadRunner, RunC (Container) and Wasmedge runtimes. (a) Total latency, including serialization, in seconds. (b) Total throughput, including serialization, in requests per second. (c) Serialization latency, isolating the impact of serialization overhead. (d) Serialization throughput, measuring throughput excluding serialization overhead. (e) Total CPU usage percentage. (f) User space CPU usage percentage. (g) Kernel space CPU usage percentage. (h) Memory usage in MB, illustrating resource consumption for each runtime.

**Inter-node: Roadrunner reduces the CPU usage by up to 85% compared to WasmEdge and reduces the RAM usage by up to 25%**

During performance experiments in Fig. 7, Roadrunner shows a total CPU usage (Fig. 8e) overhead for small payload sizes, which is a consequence of an increased Kernel Space CPU Usage (Fig. 8g), but it stabilizes as the payload size increases, while Roadrunner user

space CPU usage (Fig. 8f) keeps a linear increase similar to RunC and Wasmedge. This means Roadrunner mechanisms have a slight CPU usage overhead for small input sizes, whereas they show decreased CPU usage for larger input sizes. In Fig. 8h, Roadrunner keeps RAM usage low for small payloads and increases linearly, similar to Wasmedge RAM consumption. Roadrunner accepts a higher CPU

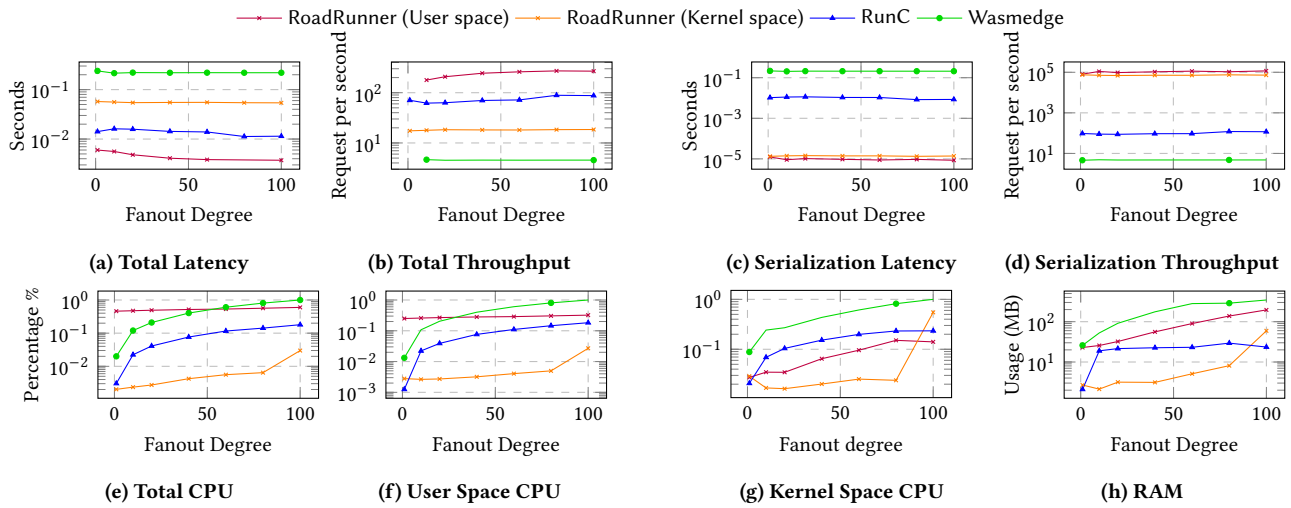


Figure 9: Intra-node Fanout Latency with 10MB function data transfer with varying fanout degrees. (a) Total Latency (b) Total Throughput; (c) Serialization Latency; (d) Serialization Throughput; (e) Total CPU usage; (f) User Space CPU usage; (g) Kernel Space CPU utilization; and (h) RAM usage in (MB).

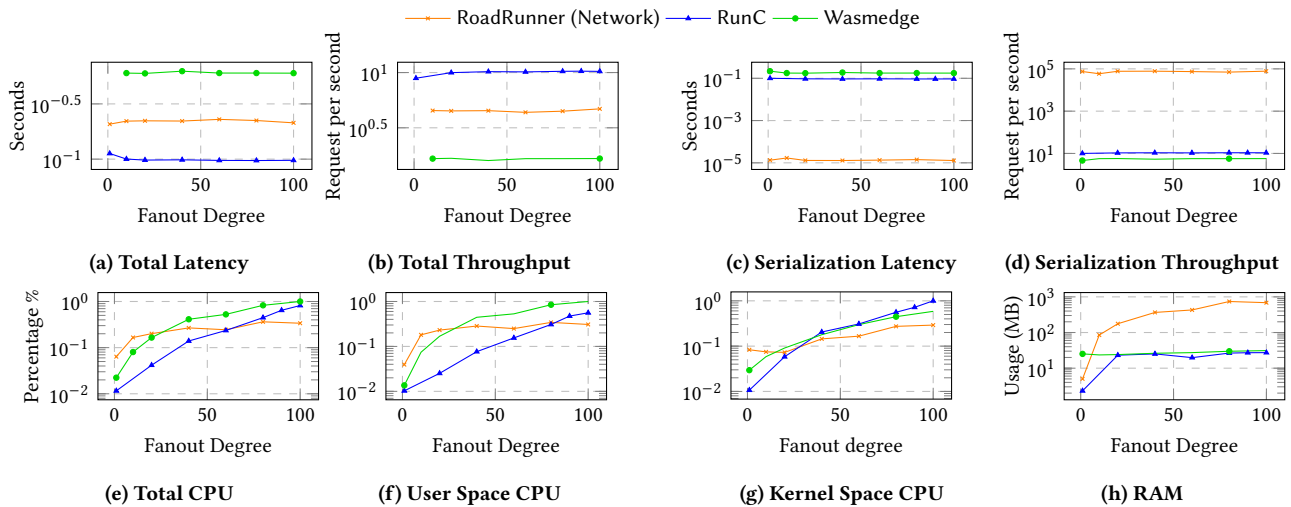


Figure 10: Inter-node Fanout Latency with 10MB function data transfer with varying fanout degrees. (a) Total Latency (b) Total Throughput; (c) Serialization Latency; (d) Serialization Throughput; (e) Total CPU usage; (f) User Space CPU usage; (g) Kernel Space CPU utilization; and (h) RAM usage in (MB).

cost for small payloads to achieve high-speed communication for all payload sizes.

During scalability experiments in Fig. 10, Roadrunner shows a slightly increased CPU usage (Fig. 10e) for small payloads, but then it stabilizes and keeps a slight linear increase over the fanout increase. A similar trend can be noticed in Fig. 8e. Under high load, Roadrunner reports less User CPU usage (Fig. 10f) than Wasmedge and RunC, showing that although the near-zero copy and serialization-free mechanisms might incur slightly higher CPU usage for smaller input sizes, they incur less CPU usage for larger input sizes due to the reduced copies between user space and kernel space. In Fig. 10g,

the increase is linear for Roadrunner and baselines. In Fig. 10h, Roadrunner shows a decreased RAM usage for small input sizes but higher RAM usage as the fanout degree increases. This increase is attributed to the additional memory requirements for handling network buffers, which do not exist for intra-node data transfer since Roadrunner leverages IPC to exchange data.

## 7 DISCUSSION

*Benefits and Trade-Offs.* Each Roadrunner communication mechanism suits the demands of different workflows. Selecting the most appropriate communication depends on balancing the benefits and

trade-offs between communication frequency, data load, latency, scalability, and resource utilization. Therefore, it is crucial to identify when and which communication is most appropriate for a given workflow, ensuring that the selected mode also aligns with the workflow’s defined requirements. User space data transfer ensures low latency and minimal resource contention but is tightly coupled, limiting scalability and introducing fault propagation risks. Kernel space data transfer provides isolation and independent scaling via IPC mechanisms, such as shared memory or UNIX domain sockets, but incurs IPC overhead and increases application complexity. Network transfer across hosts enables high scalability and strong isolation but suffers from increased latency due to network overhead alongside operational costs.

*Interoperability.* Roadrunner’s data transfer model ensures seamless inter-function communication, eliminating the dependency on external storage services for data exchange. This design is not limited to inter-function communication; it can be extended to scenarios involving external services with large input data, reducing the additional overhead associated with serialization, which significantly affects Wasm-based serverless functions. Moreover, Roadrunner remains backward compatible by defaulting to standard WASI communication unless its API is explicitly invoked, allowing unmodified modules to operate normally while enabling near-zero copy, serialization-free communication for those that opt for Roadrunner’s inter-function communication.

*Near-zero Copy Data Transfer.* Roadrunner implements near-zero copy data transfer by using zero-copy syscalls such as `splice` and `vmsplice` to move data directly between the shim and kernel buffers, avoiding redundant copies within the kernel. While this eliminates overhead between user and kernel space, data must still be copied in and out of the Wasm VM’s linear memory due to Wasm’s isolation model. Therefore, we adopted the term “near-zero copy”.

*Security Concerns.* The Wasm VM is designed to ensure strong isolation between the host and the Wasm VM, preventing unauthorized access or manipulation of memory. Although the host accesses the Wasm modules’ linear memory regions, risks such as memory corruption, buffer overflow, and data leakage are avoided due to the strict memory boundaries from Wasm which means, that a specific Wasm VM memory region is only accessible to the specific process that spawns the Wasm VM, i.e., the function shim. Every Wasm module has completely isolated memory regions not accessible to each other. In the event of a boundary violation, the function execution simply fails without affecting other parts of the system. Additionally, the memory access APIs, described in Table 1, are designed to enforce strict access controls, preventing malicious or compromised host processes from exploiting the direct memory interface.

*Threats to Validity.* To the best of our knowledge, Roadrunner is the first to explore serialization-free data transfer for Wasm serverless functions, which introduces certain limitations in the evaluation scope. Consequently, our evaluation focuses on comparing Roadrunner against state-of-the-art solutions used in serverless platforms that rely on traditional communication methods, such as HTTP, which rely on IO serialization. Although our evaluation

presents the most common invocation patterns [32], it is conducted in controlled environments with specific workloads, which may not fully capture the diversity of real-world serverless applications. Additionally, the performance gains observed may vary depending on the characteristics of the underlying hardware, network conditions, and application patterns. Roadrunner reliance on underlying host mechanisms may limit its adoption, however the demonstrated improvements highlights its potential as a foundation for future Wasm-native serverless platforms.

## 8 RELATED WORK

This section provides an overview of existing research that addresses Wasm-based Serverless Data Transfer between Serverless Functions and Sandboxes and Shims, highlighting their contributions, limitations, and how they compare to Roadrunner.

### 8.1 WebAssembly-based Serverless

Faasm [75] proposes a stateful Wasm-based serverless platform. To achieve that Faasm introduces Faaslet, a lightweight sandbox isolation mechanism that relies on cgroups and Wasm isolation to execute serverless functions. Moreover, Faasm relies on Faabric [76] to manage and exchange data between serverless functions. However, Faasm introduces a dedicated serverless platform not compatible with state-of-the-art platforms. Sledge [21] introduces a serverless Wasm runtime designed for edge computing, focusing on data-intensive multi-tenancy. It leverages Wasm sandboxes for lightweight isolation and introduces its own scheduler bypassing the kernel to optimize the execution performance and reduce resource overhead. Nevertheless, Sledge relies on a custom Wasm runtime, which limits compatibility with state-of-the-art serverless and increases development and maintenance complexity. CWASI [57] proposes an OCI-compliant shim that leverages function locality to enable direct inter-function communication: Function Embedding, Local Buffer, and Networked buffer. CWASI Function Embedding enables high-speed data passing, such as shared memory. Local buffer enables inter-function communication via host mechanisms. Nevertheless, CWASI leverages remote storage services for functions placed on different hosts, increasing latency and network overhead. Wow [20] leverages Wasm to provide isolation and portability for heterogeneous environments such as the Edge-Cloud continuum. Wow enables serverless workflows with decreased cold start times. Nevertheless, Wow does not address communication or data transfer, which accounts for a big portion of the serverless function execution time. PSL [81] leverages Wasm’s deny-by-default and isolation mechanisms to execute security and privacy in serverless workloads. PSL relies on a state store to manage state and communication among functions. Nevertheless, it still relies on state-of-the-art communication mechanisms such as HTTP, which incurs serialization and, consequently, latency increase and resource usage overhead. Roadrunner achieves serialization-free communication by leveraging Wasm’s linear memory model to facilitate data exchange between functions. Unlike traditional methods that rely on serialization to encode data for transmission (e.g., via HTTP), Roadrunner maps and transmits memory regions within the function via kernel buffers, minimizing the overhead associated with data

preparation, serialization, and deserialization while maintaining the integrity and isolation of the Wasm sandbox.

## 8.2 Serverless Inter-function Communication

SAND [2] introduces a hierarchical message bus to provide low latency with a global and a local bus. SAND optimizes remote inter-function communication by introducing a global message bus responsible for exchanging data between functions remotely. Moreover, InfiniCache [82] enables in-memory cache communication between serverless functions by providing the proxy address attached to each function. Boxer [85] and FMI [12] explore NAT TCP punch hole. The punch-hole network technique enables direct communication between two clients. In this approach, the rendezvous server controls and exchanges address data between two clients. For example, function A requests a connection to function B to the server S, the server S replies to function A and function B addresses, and forwards function A addresses to function B. Once both functions are aware of each other, they can exchange data directly.

*Zero-copy.* Diffuse [72] presents DSMQueue, a distributed shared-memory queue. DSMQueue leverages RDMA techniques to enable zero-copy data exchange between serverless functions. Nevertheless, it relies on specialized hardware, which might be a bottleneck in resource-constrained environments such as the Edge-Cloud continuum. In [90], authors propose shared memory via co-located function data exchange and a KVS for remote communication. Nightcore [31] places functions from the same workflow in a single host to enable low-latency communication via shared memory. However, nightcore does not address data serialization, which comprises a large portion of the function communication. Wasmer [84] reduces cold-start latency by enabling zero-copy deserialization of precompiled Wasm modules through memory-mapped metadata access. It serializes metadata and compiled code into a cache file that the runtime can later map directly into memory, avoiding parsing, reconstruction, and copying overhead. However, Wasmer does not support zero-copy for runtime inputs, which are still passed via traditional memory copies and host interfaces. While Wasmer focuses on cold starts, RoadRunner eliminates serialization overhead in data exchange between functions during runtime, also enabling direct communication. SharedArrayBuffer [60] enables direct memory sharing between workers, which means it enables memory sharing between workers and threads within the same function, but it does not allow between multiple processes. Roadrunner enables data sharing at the inter-function level. Furthermore, Roadrunner takes a stricter approach to memory isolation than SharedArrayBuffer, all memory sharing in Roadrunner is explicitly mediated by the shim, which controls allocation, visibility, and access semantics.

*Serialization-free.* RMMMap [51] introduces an OS primitive for remote memory map which enables serverless functions to directly access the memory of another function. Nevertheless, RMMMap is designed for RDMA-based network transfer, which requires specialized hardware. PaaS [13] enables functions running as processes, thus co-located functions achieve serialization-free communication by using a shared memory pool for direct object access, allowing zero-copy exchange via pointers in languages such as C, while

languages such as Python still use pickling but with less overhead than traditional IPC.

Roadrunner enables inter-function communication, leveraging host mechanisms to exchange data and decrease serialization overhead. Moreover, Roadrunner introduces near-zero copy techniques that optimize the communication latency and throughput of data exchange between serverless functions.

## 9 CONCLUSION AND FUTURE WORK

In this paper, we introduced Roadrunner, a near-zero copy, and serialization-free data transfer serverless runtime shim designed to optimize inter-function communication in WebAssembly-based serverless functions. Roadrunner leverages Wasm’s linear memory model and syscalls to minimize data transfer overhead between user and kernel space. By eliminating unnecessary serialization and reducing latency, Roadrunner significantly enhances the performance of serverless functions, particularly in data-intensive and latency-sensitive applications.

We evaluated Roadrunner using serverless functions that represent common invocation patterns, such as sequence and parallel functions. Our experiments demonstrated that Roadrunner reduces data transfer latency and improves throughput compared to the data-passing mechanisms of state-of-the-art serverless platforms. The experiments show that Roadrunner improves latency by up to 89% and increases the throughput up to 69x compared to the Wasm baselines, offering close to native speed for IO-bound serverless workloads. These results highlight Roadrunner’s capability to address the key bottlenecks of inter-function communication in serverless computing.

In the future, we plan to expand Roadrunner beyond its current implementation to support a broader range of WebAssembly runtimes, enabling greater portability and integration across diverse serverless platforms. Our future work will focus on developing Roadrunner into a dynamic virtualization runtime that can autonomously select the runtime type, e.g., container and Wasm, and select the most suitable runtime for specific serverless workflows based on workload and environment characteristics. Moreover, we plan to extend Roadrunner with a self-provisioning infrastructure [64] that automatically provides the necessary BaaS for each function. Finally, we aim to introduce function state management and syscall batching, allowing Roadrunner to efficiently handle stateless and stateful serverless functions.

## ACKNOWLEDGMENT

This work is partially funded by the Austrian Research Promotion Agency (FFG) under the project RapidREC (Project No. 903884). This work has received funding from the Austrian Internet Stiftung under the NetIdee project LEO Trek (ID 7442). Partly funded by EU Horizon Europe under GA no. 101070186 (TEADAL) and Ga no. 101192912 (NexaSphere).

## REFERENCES

- [1] 2023. *Amazon S3*. <https://aws.amazon.com/de/s3/>
- [2] Istemi Ekin Akkus, Ruichuan Chen, Ivica Rimac, Manuel Stein, Klaus Satzke, Andre Beck, Paarijaat Aditya, and Volker Hilt. 2018. SAND: Towards High-Performance Serverless Computing. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. USENIX Association, Boston, MA, 923–935. <https://www.usenix.org/conference/atc18/presentation/akkus>

- [3] Arm Limited. 2011. Endianness. <https://developer.arm.com/documentation/den0013/d/Porting/Endianness> Accessed: 2025-05-22.
- [4] Ioana Baldini, Paul Castro, Kerry Chang, Perry Cheng, Stephen Fink, Vatche Ishakian, Nick Mitchell, Vinod Muthusamy, Rodric Rabbah, Aleksander Slominski, et al. 2017. Serverless computing: Current trends and open problems. *Research advances in cloud computing* (2017), 1–20.
- [5] Daniel Barcelona-Pons, Pedro García-López, and Bernard Metzler. 2023. Glider: Serverless Ephemeral Stateful Near-Data Computation. In *Proceedings of the 24th International Middleware Conference* (Bologna, Italy) (*Middleware '23*). Association for Computing Machinery, New York, NY, USA, 247–260. <https://doi.org/10.1145/3590140.3629119>
- [6] Daniel Barcelona-Pons, Marc Sánchez-Artigas, Gerard Paris, Pierre Sutra, and Pedro García-López. 2019. On the FaaS Track: Building Stateful Distributed Applications with Serverless Architectures. In *Proceedings of the 20th International Middleware Conference* (Davis, CA, USA) (*Middleware '19*). Association for Computing Machinery, New York, NY, USA, 41–54. <https://doi.org/10.1145/3361525.3361535>
- [7] Daniel Barcelona-Pons, Pierre Sutra, Marc Sánchez-Artigas, Gerard Paris, and Pedro García-López. 2022. Stateful Serverless Computing with Crucial. *ACM Trans. Softw. Eng. Methodol.* 31, 3, Article 39 (March 2022), 38 pages. <https://doi.org/10.1145/3490386>
- [8] Luiz Barroso, Mike Marty, David Patterson, and Parthasarathy Ranganathan. 2017. Attack of the killer microseconds. *Commun. ACM* 60, 4 (March 2017), 48–54. <https://doi.org/10.1145/3015146>
- [9] Frank Block and Andreas Dewald. 2017. Linux memory forensics: Dissecting the user space process heap. *Digital Investigation* 22 (2017), S66–S75. <https://doi.org/10.1016/j.diin.2017.06.002>
- [10] Paul Castro, Vatche Ishakian, Vinod Muthusamy, and Aleksander Slominski. 2019. The rise of serverless computing. *Commun. ACM* 62, 12 (2019), 44–54.
- [11] Cloud Native Computing Foundation. 2023. *The State of WebAssembly 2023*. Technical Report. Cloud Native Computing Foundation. <https://www.cncf.io/wp-content/uploads/2023/09/The-State-of-WebAssembly-2023.pdf>
- [12] Marcin Copik, Roman Böhringer, Alexandru Calotoiu, and Torsten Hoefler. 2023. FMI: Fast and Cheap Message Passing for Serverless Functions. In *Proceedings of the 37th International Conference on Supercomputing* (Orlando, FL, USA) (*ICS '23*). Association for Computing Machinery, New York, NY, USA, 373–385. <https://doi.org/10.1145/3577193.3593718>
- [13] Marcin Copik, Alexandru Calotoiu, Gyorgy Rethy, Roman Böhringer, Rodrigo Bruno, and Torsten Hoefler. 2024. Process-as-a-Service: Unifying Elastic and Stateful Clouds with Serverless Processes. In *Proceedings of the 2024 ACM Symposium on Cloud Computing (SoCC '24)* (Redmond, WA, USA) (*SoCC '24*). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3698038.3698567>
- [14] Marcin Copik, Konstantin Taranov, Alexandru Calotoiu, and Torsten Hoefler. 2021. rfaas: Rdma-enabled faas platform for serverless high-performance computing. *arXiv preprint arXiv:2106.13859* (2021).
- [15] Crun. 2024. crun: A fast and lightweight fully featured OCI runtime and C library. <https://github.com/containers/crun> Accessed: 2024-12-08.
- [16] IBM Developer. 2017. *Know your TCP system call sequences*. <https://developer.ibm.com/articles/au-tcpsystemcalls/> Accessed: 2025-05-26.
- [17] Tokio Project Developers. 2024. Tokio: An asynchronous runtime for Rust. <https://tokio.rs/>. Accessed: 2024-12-13.
- [18] Simon Eismann, Joel Scheuner, Erwin van Eyk, Maximilian Schwinger, Johannes Grohmann, Nikolas Herbst, Cristina L. Abad, and Alexandru Iosup. 2022. The State of Serverless Applications: Collection, Characterization, and Community Consensus. *IEEE Transactions on Software Engineering* 48, 10 (2022), 4152–4166. <https://doi.org/10.1109/TSE.2021.3113940>
- [19] Simon Eismann, Joel Scheuner, Erwin van Eyk, Maximilian Schwinger, Johannes Grohmann, Nikolas Herbst, Cristina L. Abad, and Alexandru Iosup. 2021. Serverless Applications: Why, When, and How? *IEEE Software* 38, 1 (2021), 32–39. <https://doi.org/10.1109/MS.2020.3023302>
- [20] Philipp Gackstatter, Pantelis A. Frangoudis, and Schahram Dustdar. 2022. Pushing Serverless to the Edge with WebAssembly Runtimes. In *2022 22nd IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. 140–149. <https://doi.org/10.1109/CCGrid54584.2022.00023>
- [21] Phani Kishore Gadepalli, Sean McBride, Gregor Peach, Ludmila Cherkasova, and Gabriel Parmer. 2020. Sledge: A Serverless-First, Light-Weight Wasm Runtime for the Edge. In *Proceedings of the 21st International Middleware Conference* (Delft, Netherlands) (*Middleware '20*). Association for Computing Machinery, New York, NY, USA, 265–279. <https://doi.org/10.1145/3423211.3425680>
- [22] Rastko Gajanin, Cynthia Marcelino, and Stefan Nastic. 2025. Performance Isolation for Serverless Functions. *IEEE Transactions on Services Computing* (2025), 1–18. <https://doi.org/10.1109/TSC.2025.3619151>
- [23] Mohamad Gebai and Michel R. Dagenais. 2018. Survey and Analysis of Kernel and Userspace Tracers on Linux: Design, Implementation, and Overhead. *ACM Comput. Surv.* 51, 2, Article 26 (March 2018), 33 pages. <https://doi.org/10.1145/3158644>
- [24] Andreas Haas, Andreas Rossberg, Derek L. Schuff, Ben L. Titzer, Michael Holman, Dan Gohman, Luke Wagner, Alon Zakai, and JF Bastien. 2017. Bringing the web up to speed with WebAssembly. *SIGPLAN Not.* 52, 6 (June 2017), 185–200. <https://doi.org/10.1145/3140587.3062363>
- [25] Adam Hall and Umakishore Ramachandran. 2019. An execution model for serverless functions at the edge. In *Proceedings of the International Conference on Internet of Things Design and Implementation* (Montreal, Quebec, Canada) (*IoTDI '19*). Association for Computing Machinery, New York, NY, USA, 225–236. <https://doi.org/10.1145/3302505.3310084>
- [26] Joseph M. Hellerstein, Jose Faleiro, Joseph E. Gonzalez, Johann Schleier-Smith, Vikram Sreekanti, Alexey Tumanov, and Chenggang Wu. 2018. Serverless Computing: One Step Forward, Two Steps Back. arXiv:1812.03651 [cs.DC]
- [27] Aaron Hilbig, Daniel Lehmann, and Michael Pradel. 2021. An Empirical Study of Real-World WebAssembly Binaries: Security, Languages, Use Cases. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (*WWW '21*). Association for Computing Machinery, New York, NY, USA, 2696–2708. <https://doi.org/10.1145/3442381.3450138>
- [28] Zhuo Huang, Hao Fan, Chaoyi Cheng, Song Wu, and Hai Jin. 2023. Duo: Improving Data Sharing of Stateful Serverless Applications by Efficiently Caching Multi-Read Data. In *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 875–885. <https://doi.org/10.1109/IPDPS54959.2023.00092>
- [29] Open Containers Initiative. 2024. runc: CLI tool for spawning and running containers according to the OCI specification. <https://github.com/opencontainers/runc>. Accessed: 2024-12-05.
- [30] Jaeyoung Jang, Sung Jun Jung, Sunmin Jeong, Jun Heo, Hoon Shin, Tae Jun Ham, and Jae W. Lee. 2020. A specialized architecture for object serialization with applications to big data analytics. In *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture* (Virtual Event) (*ISCA '20*). IEEE Press, 322–334. <https://doi.org/10.1109/ISCA45697.2020.00036>
- [31] Zhipeng Jia and Emmett Witchel. 2021. Nightcore: Efficient and Scalable Serverless Computing for Latency-Sensitive, Interactive Microservices. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Virtual, USA) (*ASPLOS '21*). Association for Computing Machinery, New York, NY, USA, 152–166. <https://doi.org/10.1145/3445814.3446701>
- [32] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Yadwadkar, Joseph E. Gonzalez, Raluca Ada Popa, Ion Stoica, and David A. Patterson. 2019. Cloud Programming Simplified: A Berkeley View on Serverless Computing. arXiv:1902.03383 [cs.OS]
- [33] Aditya K Kamath and Simon Peter. 2024. (MC)2: Lazy MemCopy at the Memory Controller. In *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*. 1112–1128. <https://doi.org/10.1109/ISCA59077.2024.00084>
- [34] Sagar Karandikar, Chris Leary, Chris Kennelly, Jerry Zhao, Dinesh Parimi, Borivoje Nikolic, Krste Asanovic, and Parthasarathy Ranganathan. 2021. A Hardware Accelerator for Protocol Buffers. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture* (Virtual Event, Greece) (*MICRO '21*). Association for Computing Machinery, New York, NY, USA, 462–478. <https://doi.org/10.1145/3466752.3480051>
- [35] Michael Kerrisk et al. 2023. cgroups(7) - Linux manual page. <https://man7.org/linux/man-pages/man7/cgroups.7.html>. Accessed: 2024-12-12.
- [36] Michael Kerrisk et al. 2023. namespaces(7) - Linux manual page. <https://man7.org/linux/man-pages/man7/namespaces.7.html>. Accessed: 2024-12-12.
- [37] Daehyeok Kim, Tianlong Yu, Hongqiang Harry Liu, Yibo Zhu, Jitu Padhye, Shachar Raindel, Chuanxiong Guo, Vyas Sekar, and Srinivasan Seshan. 2019. FreeFlow: Software-based Virtual RDMA Networking for Containerized Clouds. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA, 113–126. <https://www.usenix.org/conference/nsdi19/presentation/kim>
- [38] Ana Klimovic, Yawen Wang, Patrick Stuedi, Animesh Trivedi, Jonas Pfefferle, and Christos Kozyrakis. 2018. Pocket: Elastic Ephemeral Storage for Serverless Analytics. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. USENIX Association, Carlsbad, CA, 427–444. <https://www.usenix.org/conference/osdi18/presentation/klimovic>
- [39] Swaroop Kotni, Ajay Nayak, Vinod Ganapathy, and Arkaprava Basu. 2021. Faastlane: Accelerating Function-as-a-Service Workflows. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, 805–820. <https://www.usenix.org/conference/atc21/presentation/kotni>
- [40] Tom Kuchler, Michael Giardino, Timothy Roscoe, and Ana Klimovic. 2023. Function as a Function. In *Proceedings of the 2023 ACM Symposium on Cloud Computing* (Santa Cruz, CA, USA) (*SoCC '23*). Association for Computing Machinery, New York, NY, USA, 81–92. <https://doi.org/10.1145/3620678.3624648>
- [41] Linux Kernel Labs. [n. d.]. Typical Operating System Architecture. <https://linux-kernel-labs.github.io/refs/heads/master/lectures/intro.html#typical-operating-system-architecture>. Accessed: 2024-11-13.
- [42] Daniel Lehmann, Johannes Kinder, and Michael Pradel. 2020. Everything Old is New Again: Binary Security of WebAssembly. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 217–234. <https://www.usenix.org/conference/usenixsecurity20/presentation/lehmann>

- [43] Daniel Lehmann and Michael Pradel. 2019. Wasabi: A Framework for Dynamically Analyzing WebAssembly. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems* (Providence, RI, USA) (ASPLOS '19). Association for Computing Machinery, New York, NY, USA, 1045–1058. <https://doi.org/10.1145/3297858.3304068>
- [44] Tianyu Li, Yingpeng Chen, Donghui Yu, Yuanyuan Zhang, and Bert Lagaisse. 2024. Enhancing Effective Bidirectional Isolation for Function Fusion in Serverless Architectures. In *Proceedings of the 25th International Middleware Conference* (Hong Kong, Hong Kong) (MIDDLEWARE '24). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3652892.3654778>
- [45] Yongkang Li, Yanying Lin, Yang Wang, Kejiang Ye, and Chengzhong Xu. 2022. Serverless computing: state-of-the-art, challenges and opportunities. *IEEE Transactions on Services Computing* 16, 2 (2022), 1522–1539.
- [46] Zijun Li, Linsong Guo, Jiagan Cheng, Quan Chen, Bingsheng He, and Minyi Guo. 2022. The Serverless Computing Survey: A Technical Primer for Design Architecture. *ACM Comput. Surv.* 54, 10s, Article 220 (sep 2022), 34 pages. <https://doi.org/10.1145/3508360>
- [47] Zijun Li, Yushi Liu, Linsong Guo, Quan Chen, Jiagan Cheng, Wenli Zheng, and Minyi Guo. 2022. FaaSFlow: enable efficient workflow execution for function-as-a-service. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (ASPLOS '22). Association for Computing Machinery, New York, NY, USA, 782–796. <https://doi.org/10.1145/3503222.3507717>
- [48] Mugeng Liu, Haiyang Shen, Yixuan Zhang, Hong Mei, and Yun Ma. 2025. WebAssembly for Container Runtime: Are We There Yet? *ACM Trans. Softw. Eng. Methodol.* (Feb. 2025). <https://doi.org/10.1145/3712197>
- [49] Qingyuan Liu, Dong Du, Yubin Xia, Ping Zhang, and Haibo Chen. 2023. The Gap Between Serverless Research and Real-World Systems. In *Proceedings of the 2023 ACM Symposium on Cloud Computing* (Santa Cruz, CA, USA) (SoCC '23). Association for Computing Machinery, New York, NY, USA, 475–485. <https://doi.org/10.1145/3620678.3624785>
- [50] Ju Long, Hung-Ying Tai, Shen-Ta Hsieh, and Michael Juntao Yuan. 2021. A Lightweight Design for Serverless Function as a Service. *IEEE Software* 38, 1 (2021), 75–80. <https://doi.org/10.1109/MS.2020.3028991>
- [51] Fangming Lu, Xingda Wei, Zhuobin Huang, Rong Chen, Minyu Wu, and Haibo Chen. 2024. Serialization/Deserialization-free State Transfer in Serverless Workflows. In *Proceedings of the Nineteenth European Conference on Computer Systems* (Athens, Greece) (EuroSys '24). Association for Computing Machinery, New York, NY, USA, 132–147. <https://doi.org/10.1145/3627703.3629568>
- [52] Ashraf Mahgoub, Karthick Shankar, Subrata Mitra, Ana Klimovic, Somali Chaterji, and Saurabh Bagchi. 2021. SONIC: Application-aware Data Passing for Chained Serverless Applications. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, 285–301. <https://www.usenix.org/conference/atc21/presentation/mahgoub>
- [53] WebAssembly. Man. 2024. WebAssembly Data Types. [https://www.webassemblyman.com/webassembly\\_data\\_types.html](https://www.webassemblyman.com/webassembly_data_types.html) Accessed: 2024-11-13.
- [54] man7.org. 2024. splice – Linux manual page. <https://man7.org/linux/man-pages/man2/splice.2.html>. Accessed: 2024-11-11.
- [55] man7.org. 2024. vmsplice – Linux manual page. <https://man7.org/linux/man-pages/man2/vmsplice.2.html>. Accessed: 2024-11-11.
- [56] Cynthia Marcelino, Noah Krennmaier, Thomas Pusztai, and Stefan Nastic. 2025. Lumos: Performance Characterization of WebAssembly as a Serverless Runtime in the Edge-Cloud Continuum. In *Proceedings of the 14th International Conference on the Internet of Things (IoT '25)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3770501.3770515>
- [57] Cynthia Marcelino and Stefan Nastic. 2023. CWASI: A WebAssembly Runtime Shim for Inter-function Communication in the Serverless Edge-Cloud Continuum. In *The Eighth ACM/IEEE Symposium on Edge Computing (SEC 2023)*. <https://doi.org/10.1145/3583740.3626611>
- [58] Cynthia Marcelino, Jack Shahhoud, and Stefan Nastic. 2024. GoldFish: Serverless Actors with Short-Term Memory State for the Edge-Cloud Continuum. In *Proceedings of the 14th International Conference on the Internet of Things (Oulu, Finland) (IoT '24)*. ACM, New York, NY, USA.
- [59] Garrett McGrath and Paul R Brenner. 2017. Serverless computing: Design, implementation, and performance. In *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*. IEEE, 405–410.
- [60] MDN contributors. 2025. "SharedArrayBuffer – JavaScript | MDN". [https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global\\_Objects/SharedArrayBuffer](https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/SharedArrayBuffer). Accessed October 22, 2025.
- [61] J ames M en etrey, Marcelo Pasin, Pascal Felber, and Valerio Schiavoni. 2022. WebAssembly as a Common Layer for the Cloud-Edge Continuum. In *Proceedings of the 2nd Workshop on Flexible Resource and Application Management on the Edge* (Minneapolis, MN, USA) (FRAME '22). Association for Computing Machinery, New York, NY, USA, 3–8. <https://doi.org/10.1145/3526059.3533618>
- [62] Alex Merenstein, Vasily Tarasov, Ali Anwar, Scott Guthridge, and Erez Zadok. 2023. F3: Serving Files Efficiently in Serverless Computing. In *Proceedings of the 16th ACM International Conference on Systems and Storage* (Haifa, Israel) (SYSTOR '23). Association for Computing Machinery, New York, NY, USA, 8–21. <https://doi.org/10.1145/3579370.3594771>
- [63] Djob Mvondo, Mathieu Bacou, Kevin Nguetchouang, Lucien Ngale, St ephane Pouget, Josiane Kouam, Renaud Lachaize, Jinho Hwang, Tim Wood, Daniel Hagimont, No el De Palma, Bernab e Batchakui, and Alain Tchana. 2021. OFC: An Opportunistic Caching System for FaaS Platforms. In *Proceedings of the Sixteenth European Conference on Computer Systems* (Online Event, United Kingdom) (EuroSys '21). Association for Computing Machinery, New York, NY, USA, 228–244. <https://doi.org/10.1145/3447786.3456239>
- [64] Stefan Nastic. 2024. Self-Provisioning Infrastructures for the Next Generation Serverless Computing. *SN Computer Science* 5, 6 (2024), 678 – 693. <https://doi.org/10.1007/s42979-024-03022-w>
- [65] Khanh Nguyen, Lu Fang, Christian Navasca, Guoqing Xu, Brian Demsky, and Shan Lu. 2018. Skyway: Connecting Managed Heaps in Distributed Big Data Systems (ASPLOS '18). Association for Computing Machinery, New York, NY, USA, 56–69. <https://doi.org/10.1145/3173162.3173200>
- [66] On Time RTOS-32. [n. d.]. *Function revc*. <https://www.on-time.com/rtos-32-docs/rtip-32/reference-manual/socket-api/recv.htm> Accessed: 2025-05-26.
- [67] Federico Parola, Shixiong Qi, Anvaya B. Narappa, K. K. Ramakrishnan, and Fulvio Risso. 2024. SURE: Secure Unikernels Make Serverless Computing Rapid and Efficient. In *Proceedings of the ACM Symposium on Cloud Computing (SoCC '24)* (Redmond, WA, USA). Association for Computing Machinery. <https://doi.org/10.1145/3698038.3698558>
- [68] Deepti Raghavan, Philip Levis, Matei Zaharia, and Irene Zhang. 2021. Breakfast of champions: towards zero-copy serialization with NIC scatter-gather. In *Proceedings of the Workshop on Hot Topics in Operating Systems* (Ann Arbor, Michigan) (HotOS '21). Association for Computing Machinery, New York, NY, USA, 199–205. <https://doi.org/10.1145/3458336.3465287>
- [69] Deepti Raghavan, Shreya Ravi, Gina Yuan, Pratiksha Thaker, Sanjari Srivastava, Micah Murray, Pedro Henrique Penna, Amy Ousterhout, Philip Levis, Matei Zaharia, and Irene Zhang. 2023. Cornflakes: Zero-Copy Serialization for Microsecond-Scale Networking. In *Proceedings of the 29th Symposium on Operating Systems Principles* (Koblenz, Germany) (SOSP '23). Association for Computing Machinery, New York, NY, USA, 200–215. <https://doi.org/10.1145/3600066.3613137>
- [70] Philipp Raith, Stefan Nastic, and Schahram Dustdar. 2023. Serverless Edge Computing—Where We Are and What Lies Ahead. *IEEE Internet Computing* 27, 3 (2023), 50–64. <https://doi.org/10.1109/MIC.2023.3260939>
- [71] Francisco Romero, Gohar Irfan Chaudhry,  nigo Goiri, Pragna Gopa, Paul Batum, Neeraja J. Yadwadkar, Rodrigo Fonseca, Christos Kozyrakis, and Ricardo Bianchini. 2021. FaaST: A Transparent Auto-Scaling Cache for Serverless Applications. In *Proceedings of the ACM Symposium on Cloud Computing* (Seattle, WA, USA) (SoCC '21). Association for Computing Machinery, New York, NY, USA, 122–137. <https://doi.org/10.1145/3472883.3486974>
- [72] Andrea Sabbioni, Lorenzo Rosa, Armir Bujari, Luca Foschini, and Antonio Corradi. 2022. DIFFUSE: A Distributed and decentralized platform enabling Function composition in Serverless Environments. *Computer Networks* 210 (2022), 108993. <https://doi.org/10.1016/j.comnet.2022.108993>
- [73] Johann Schleier-Smith, Vikram Sreekanti, Anurag Khandelwal, Joao Carreira, Neeraja J. Yadwadkar, Raluca Ada Popa, Joseph E. Gonzalez, Ion Stoica, and David A. Patterson. 2021. What Serverless Computing is and Should Become: The next Phase of Cloud Computing. *Commun. ACM* 64, 5 (apr 2021), 76–84. <https://doi.org/10.1145/3406011>
- [74] Mohammad Shahrad, Rodrigo Fonseca, Inigo Goiri, Gohar Chaudhry, Paul Batum, Jason Cooke, Eduardo Laureano, Colby Tresness, Mark Russinovich, and Ricardo Bianchini. 2020. Serverless in the Wild: Characterizing and Optimizing the Serverless Workload at a Large Cloud Provider. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*. USENIX Association, 205–218. <https://www.usenix.org/conference/atc20/presentation/shahrad>
- [75] Simon Shillaker and Peter Pietzuch. 2020. Faasm: Lightweight Isolation for Efficient Stateful Serverless Computing. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*. USENIX Association, 419–433. <https://www.usenix.org/conference/atc20/presentation/shillaker>
- [76] Simon Shillaker, Carlos Segarra, Eleftheria Mappoura, Mayeul Fournial, Lluis Vilanova, and Peter Pietzuch. 2023. Faabric: Fine-Grained Distribution of Scientific Workloads in the Cloud. *arXiv preprint arXiv:2302.11358* (2023).
- [77] Brian Stetten. 2021. *WebAssembly: The Definitive Guide*. O'Reilly Media, Inc. <https://www.oreilly.com/library/view/webassembly-the-definitive/9781492089841/>
- [78] Vikram Sreekanti, Chenggang Wu, Xiayue Charles Lin, Johann Schleier-Smith, Jose M. Faleiro, Joseph E. Gonzalez, Joseph M. Hellerstein, and Alexey Tumanov. 2020. Cloudburst: Stateful Functions-as-a-Service. *Proc. VLDB Endow.* 13 (2020), 2438–2452.
- [79] Qiang Su, Chuanwen Wang, Zhixiong Niu, Ran Shu, Peng Cheng, Yongqiang Xiong, Dongsu Han, Chun Jason Xue, and Hong Xu. 2022. PipeDevice: A Hardware-Software Co-Design Approach to Intra-Host Container Communication. In *Proceedings of the 18th International Conference on Emerging Networking Experiments and Technologies* (CoNEXT '22). Association for Computing Machinery, New York, NY, USA, 126–139. <https://doi.org/10.1145/3555050.3569118>

- [80] Konstantin Taranov, Rodrigo Bruno, Gustavo Alonso, and Torsten Hoefer. 2021. Naos: Serialization-free RDMA networking in Java. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*. USENIX Association, 1–14. <https://www.usenix.org/conference/atc21/presentation/taranov>
- [81] Alex Thomas, Shubham Mishra, Eric Chen, and John D. Kubiatowicz. 2024. *Lightweight, Secure and Stateful Serverless Computing with PSL*. Technical Report UCB/ECS-2024-191. EECs Department, University of California, Berkeley. <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2024/ECS-2024-191.html>
- [82] Ao Wang, Jingyuan Zhang, Xiaolong Ma, Ali Anwar, Lukas Rupprecht, Dimitrios Skourtis, Vasily Tarasov, Feng Yan, and Yue Cheng. 2020. InfiniCache: Exploiting Ephemeral Serverless Functions to Build a Cost-Effective Memory Cache. In *18th USENIX Conference on File and Storage Technologies (FAST 20)*. USENIX Association, Santa Clara, CA, 267–281. <https://www.usenix.org/conference/fast20/presentation/wang-ao>
- [83] WasmEdge Project. 2024. WasmEdge - High-Performance WebAssembly Runtime for Cloud Native, Edge, and Decentralized Applications. <https://wasmedge.org>. Accessed: 2024-11-11.
- [84] Wasmer. 2023. Improving With Zero-Copy Deserialization. <https://wasmer.io/posts/improving-with-zero-copy-deserialization>.
- [85] Mike Wawrzoniak, Ingo Müller, Rodrigo Fraga Barcelos Paulus Bruno, and Gustavo Alonso. 2021. Boxer: Data analytics on network-enabled serverless platforms. In *11th Annual Conference on Innovative Data Systems Research (CIDR 2021)*.
- [86] Xingda Wei, Fangming Lu, Tianxia Wang, Jinyu Gu, Yuhang Yang, Rong Chen, and Haibo Chen. 2023. No Provisioned Concurrency: Fast {RDMA-codedigned} Remote Fork for Serverless Computing. In *17th USENIX Symposium on Operating Systems Design and Implementation (OSDI 23)*. 497–517.
- [87] Jinfeng Wen, Zhenpeng Chen, Xin Jin, and Xuanzhe Liu. 2023. Rise of the Planet of Serverless Computing: A Systematic Review. *ACM Trans. Softw. Eng. Methodol.* 32, 5, Article 131 (July 2023), 61 pages. <https://doi.org/10.1145/3579643>
- [88] Chenggang Wu, Jose M. Faleiro, Yihan Lin, and Joseph M. Hellerstein. 2018. Anna: A KVS for Any Scale. *2018 IEEE 34th International Conference on Data Engineering (ICDE)* (2018), 401–412.
- [89] Mingyu Wu, Shuaiwei Wang, Haibo Chen, and Binyu Zang. 2022. Zero-Change Object Transmission for Distributed Big Data Analytics. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. USENIX Association, Carlsbad, CA, 137–150. <https://www.usenix.org/conference/atc22/presentation/wu>
- [90] Minchen Yu, Tingjia Cao, Wei Wang, and Ruichuan Chen. 2023. Following the Data, Not the Function: Rethinking Function Orchestration in Serverless Computing. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1489–1504. <https://www.usenix.org/conference/nsdi23/presentation/you>
- [91] Tianfan Zhang, Huaping Zhou, Chengyuan Huang, Chen Tian, Wei Zhang, Xiaoliang Wang, Yi Wang, Ahmed M. Abdelmoniem, Matthew Tan, Wanchun Dou, and Guihai Chen. 2023. Achieving Zero-copy Serialization for Datacenter RPC. In *2023 IEEE International Performance, Computing, and Communications Conference (IPCCC)*. 304–312. <https://doi.org/10.1109/IPCCC59175.2023.10253859>
- [92] Yusheng Zheng, Tong Yu, Yiwei Yang, and Andrew Quinn. 2024. Wasm-bpf: Streamlining eBPF Deployment in Cloud Environments with WebAssembly. arXiv:2408.04856 [cs.OS] <https://arxiv.org/abs/2408.04856>