

ASPLOS 2023 – Vancouver, Canada – 29th March

Lightning Talk Video

Exit-Less, Isolated, and Shared Access for Virtual Machines

Kenichi Yasukata, Hajime Tazaki, Pierre-Louis Aublin



Internet Initiative Japan

Background

Background

- VMs are one of the most popular computing environments

Background

- VMs are one of the most popular computing environments
- Memory isolation is an important feature of VMs

Background

- VMs are one of the most popular computing environments
- Memory isolation is an important feature of VMs
- But, VMs sometimes need to access the same memory region


Background

- VMs are one of the most popular computing environments
- Memory isolation is an important feature of VMs
- But, VMs sometimes need to access the same memory region
- For example, to share DMA-capable I/O devices

Problem Statement


Problem Statement

Direct-mapping




Problem Statement

Direct-mapping



Problem Statement


Direct-mapping



Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None

Problem Statement


Direct-mapping




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None

Problem Statement

Direct-mapping




Host-interposition




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None

Problem Statement

Direct-mapping




Host-interposition




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None

Problem Statement

Direct-mapping



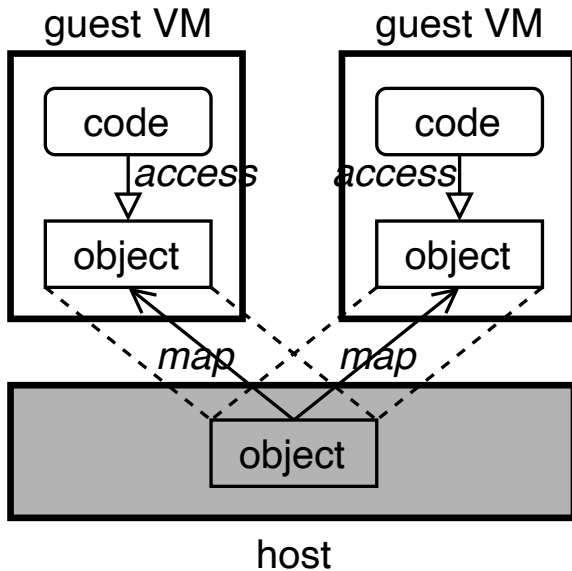
Host-interposition




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None

Problem Statement

Direct-mapping




Host-interposition




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None

Problem Statement

Direct-mapping




Host-interposition




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None
Host-interposition	Isolated	High

Problem Statement

Direct-mapping




Host-interposition




Description	Shared access	Isolation overhead
Direct-mapping	No isolation	None
Host-interposition	Isolated	High

Problem Statement

Direct-mapping



Host-interposition



Description

Direct-mapping

Host-interposition

Shared access

Isolation overhead

No isolation



None



Isolated




High



This Work

ELISA (Exit-Less, Isolated, and Shared Access)




Description	Shared access	Isolation overhead
Direct-mapping	No isolation 😭	None 😊
Host-interposition	Isolated 😊	High 😭
ELISA (this work)	Isolated 😊	Low 😊

This Work

- Isolation
 - Extended Page Table (EPT) separation

ELISA (Exit-Less, Isolated, and Shared Access)

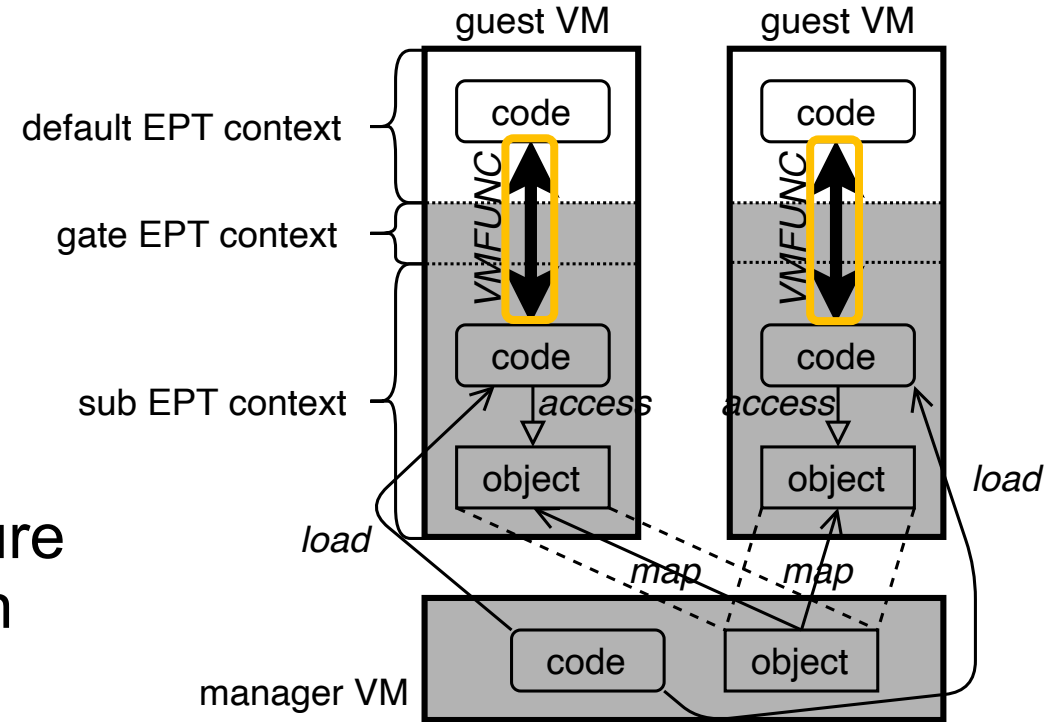


Description	Shared access	Isolation overhead
Direct-mapping	No isolation 😭	None 😊
Host-interposition	Isolated 😊	High 😭
ELISA (this work)	Isolated 😊	Low 😊

This Work

- Isolation
 - Extended Page Table (EPT) separation
- Low overhead
 - Using EPT Pointer (EPTP) switching feature of VMFUNC, a lightweight CPU instruction

ELISA (Exit-Less, Isolated, and Shared Access)




Description	Shared access	Isolation overhead
Direct-mapping	No isolation 😭	None 😊
Host-interposition	Isolated 😊	High 😭
ELISA (this work)	Isolated 😊	Low 😊

This Work

- Isolation
 - Extended Page Table (EPT) separation
- Low overhead
 - Using EPT Pointer (EPTP) switching feature of VMFUNC, a lightweight CPU instruction

ELISA (Exit-Less, Isolated, and Shared Access)



Description	Shared access	Isolation overhead
Direct-mapping	No isolation 😭	None 😊
Host-interposition	Isolated 😊	High 😭
ELISA (this work)	Isolated 😊	Low 😊

Results

- The context switch overhead of ELISA is 3.5 times lighter than that of VMCALL-oriented host-interposition

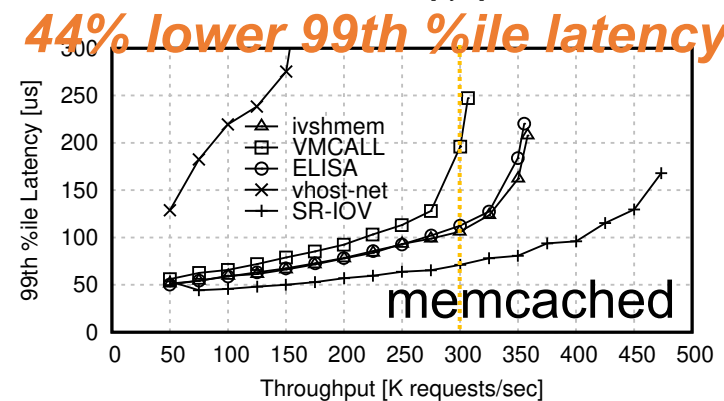

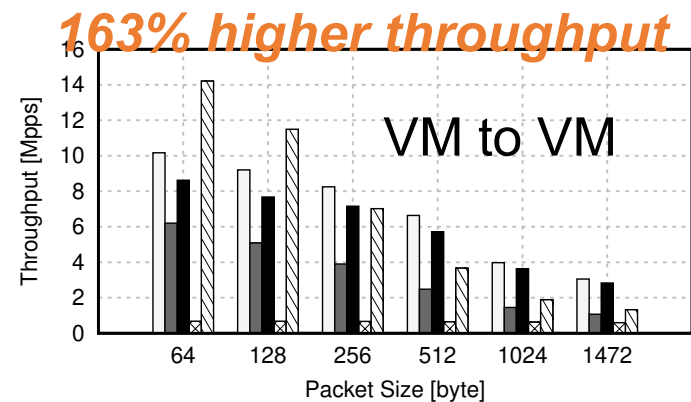
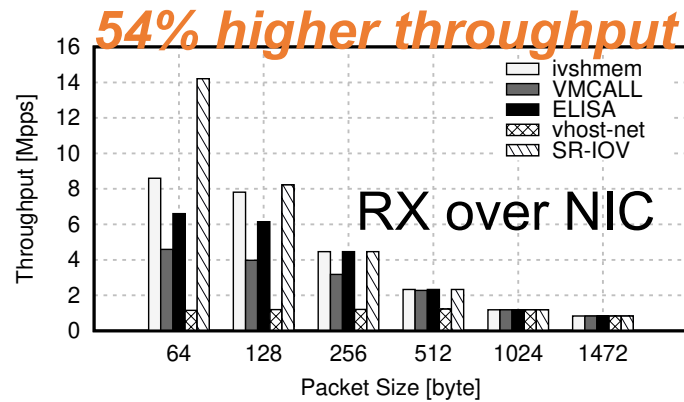
Description	Time [ns]
ELISA	196
VMCALL	699

Context Round-trip Time


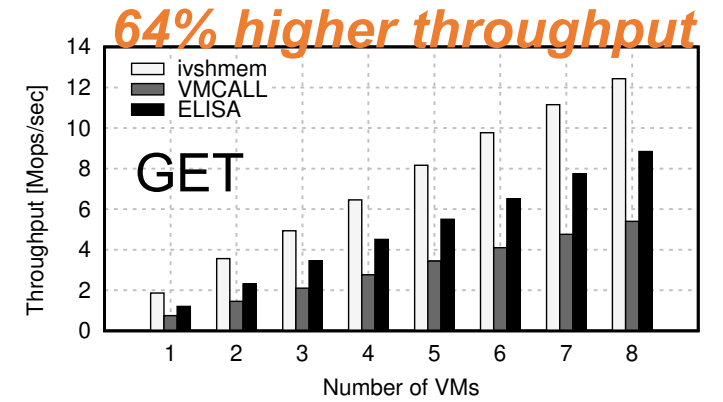
Results

- Replacement from VMCALL to ELISA speeds up applications

VM Networking System



In-memory Key-Value Store



More Information

- Project web page : <https://github.com/yasukata/ELISA>
 - Paper
 - Slides
 - Source code
 - Commentary

