

VAMOS: Virtualization Aware Middleware

Abel Gordon¹ Muli Ben-Yehuda^{1,2} Dennis Filimonov² Mahor Dahan²

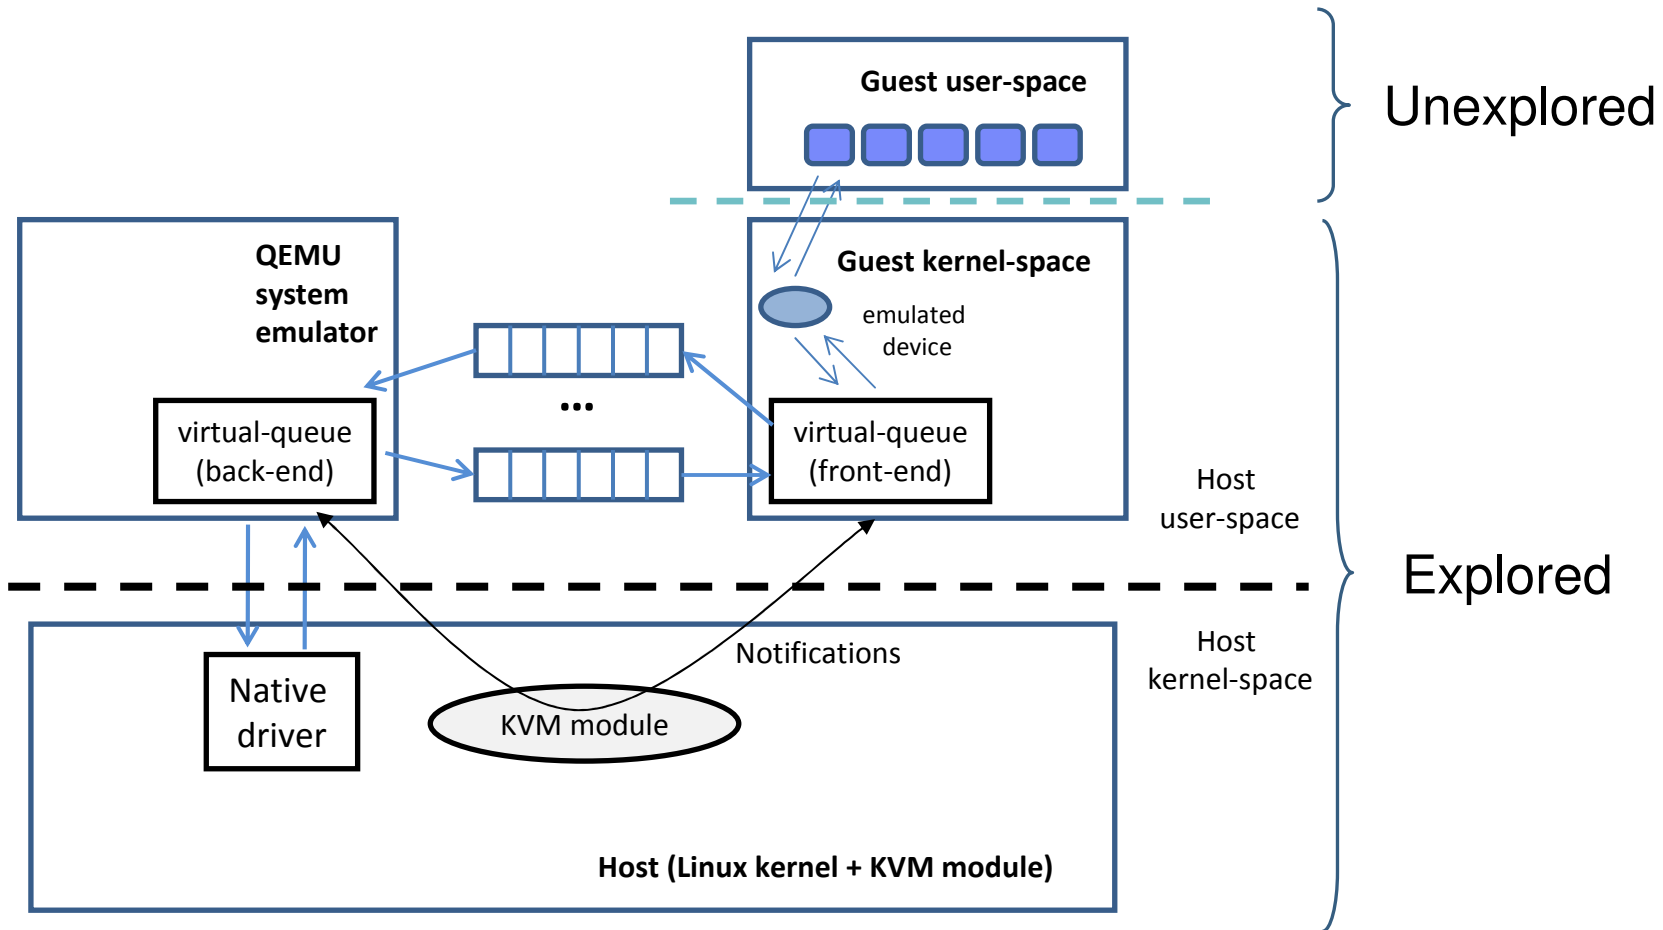
¹ IBM Research – Haifa

² Technion – Israel Institute of Technology

VAMOS: Virtualization Aware Middleware

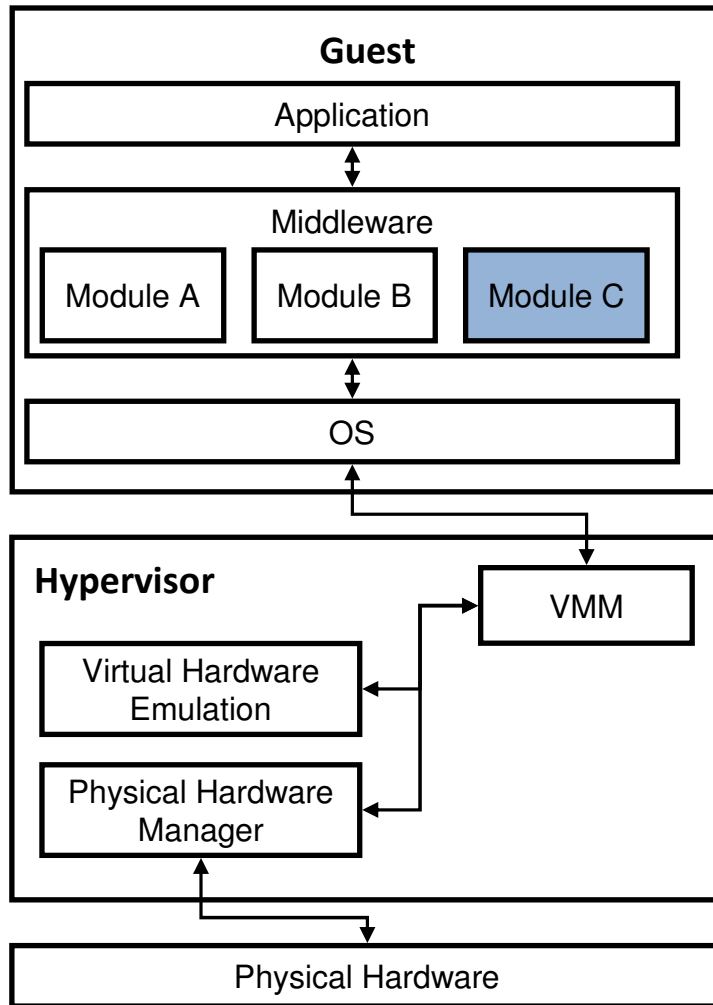
- Virtualization overhead is still high due to the transitions between the guest and the hypervisor
- Motivation: potential optimizations at the application layer have been ignored
 - Software is still built based on models that apply to non-virtual systems
 - Applications are not being adapted for the underlying “virtual” platform
 - No cooperation between the application and the hypervisor
 - Adapting the middleware such as Databases, Web Servers, Application Servers to virtualized platforms we can indirectly adapt many applications and regain lost performance

I/O virtualization with KVM, a long way down....

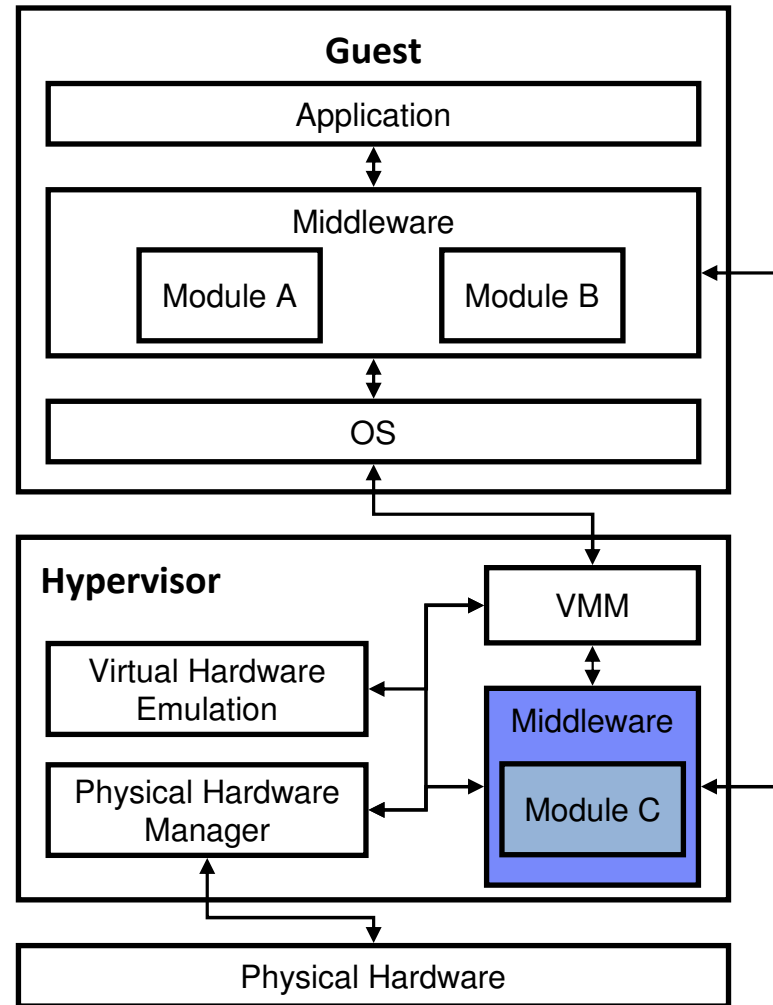


VAMOS Goals

- Reduce virtualization overhead by adding virtualization awareness to the middleware
- Avoid changes in the guest operating system
- Re-use/re-factor existent code by exploiting modularity
- Avoid software re-write/re-design

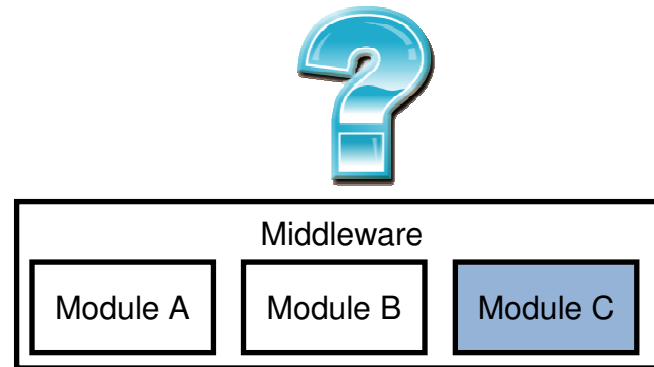


Traditional Architecture



VAMOS Architecture

Middleware Adaptation



1. Modules which interact directly with system resources, generating many transitions to the hypervisor context and requiring emulation of virtual hardware
2. Modules which can be easily re-factored into a client side running in the guest and a server side running in the host
3. Modules which do not share state with other components, avoiding data sharing and synchronization between the guest and the hypervisor.
4. Modules that do not require persisted state at the hypervisor level (do not require special handling for live-migration, checkpoint/restore)

VAMOS Requirements

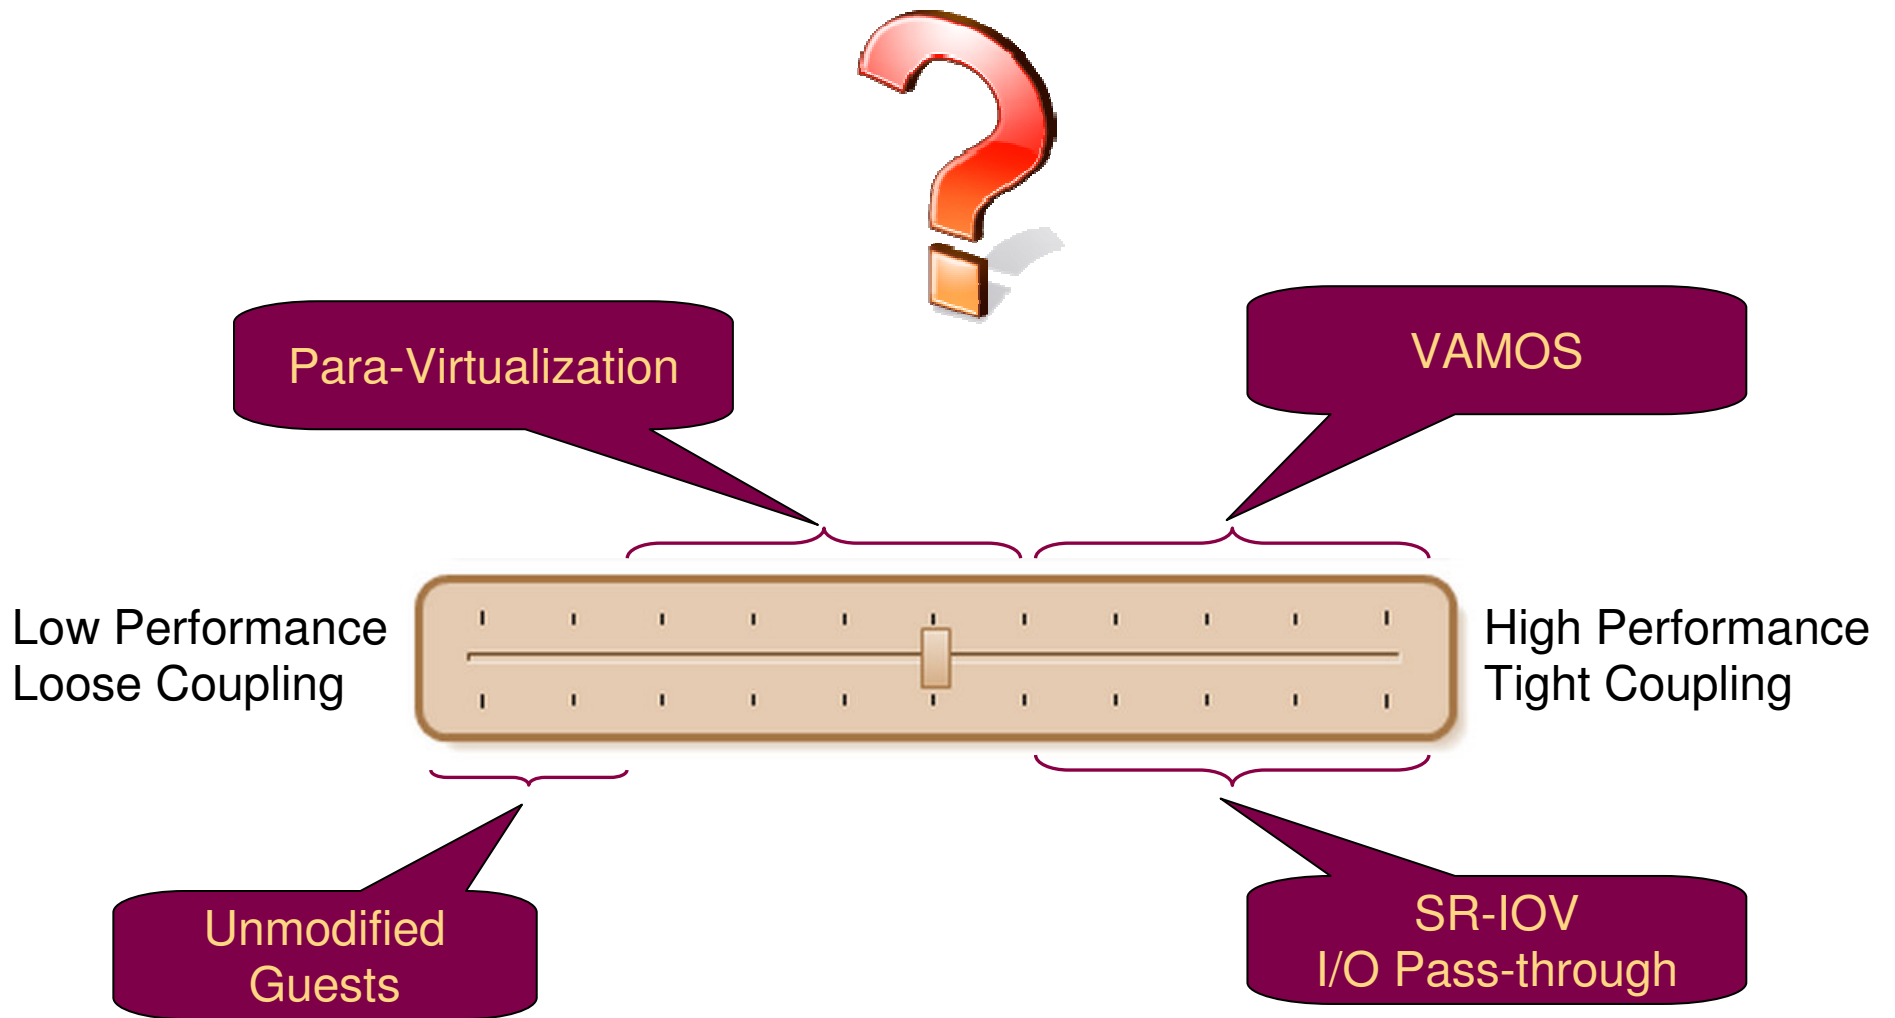
1. A runtime environment:

- Isolated
 - At the hypervisor level
 - Executing middleware code
 - With access to physical resources, such as network devices or disk drives
- Most hypervisors have a general purpose OS

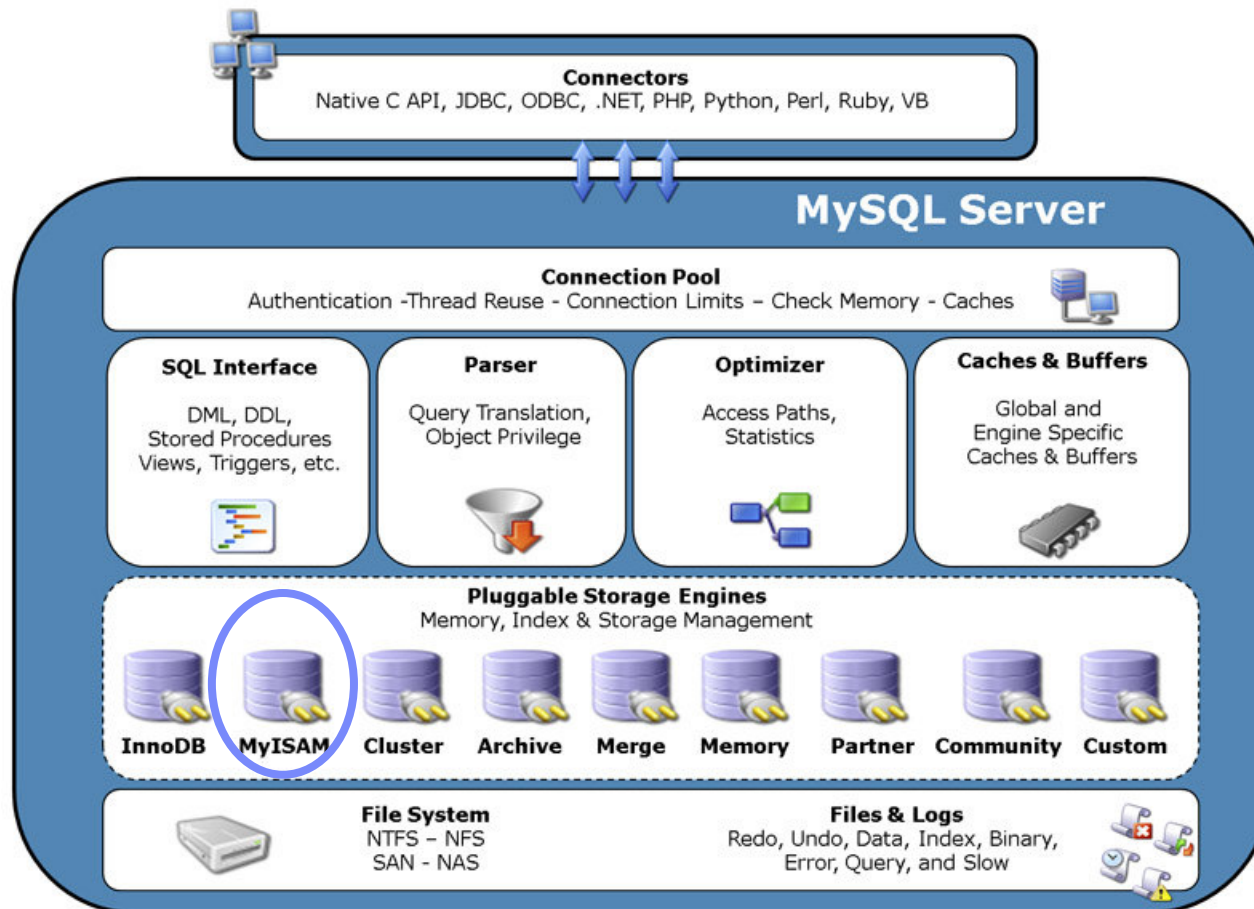
2. A communication channel between:

- the middleware running in the guest
 - the middleware running in the host
 - the hypervisor
- Para-virtualization channels are commonly used

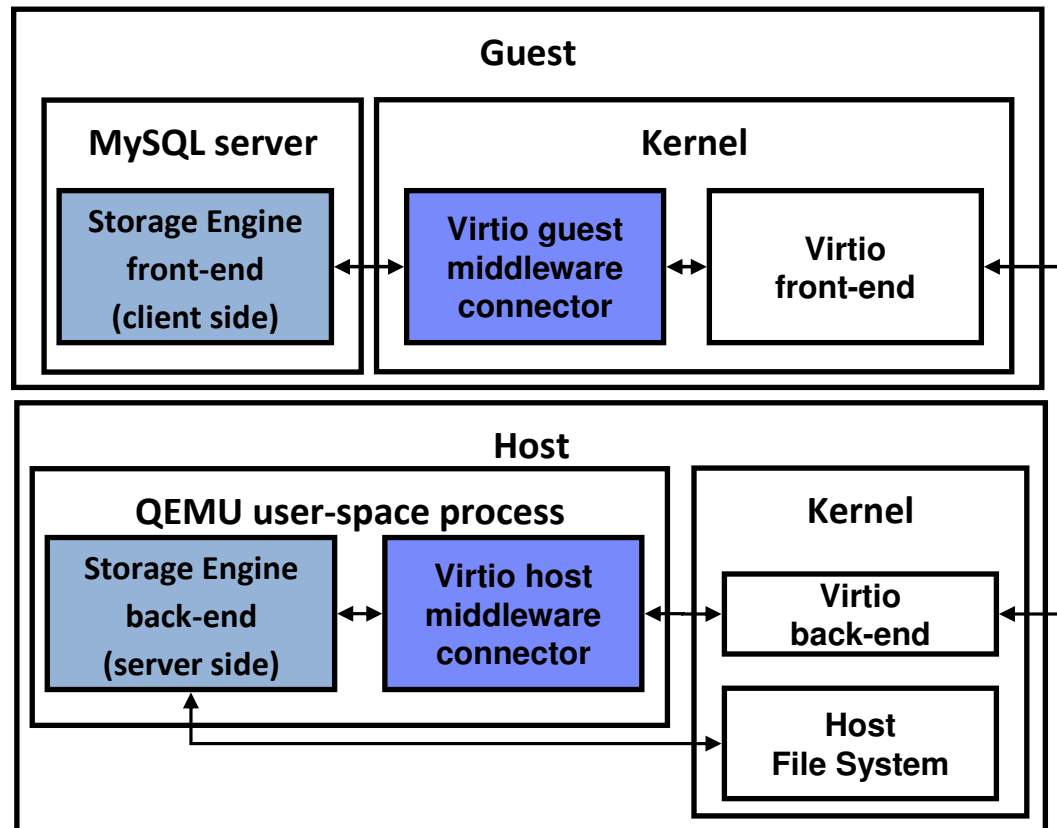
Looking for the right balance




MySQL Software Architecture



VAMOS for MySQL



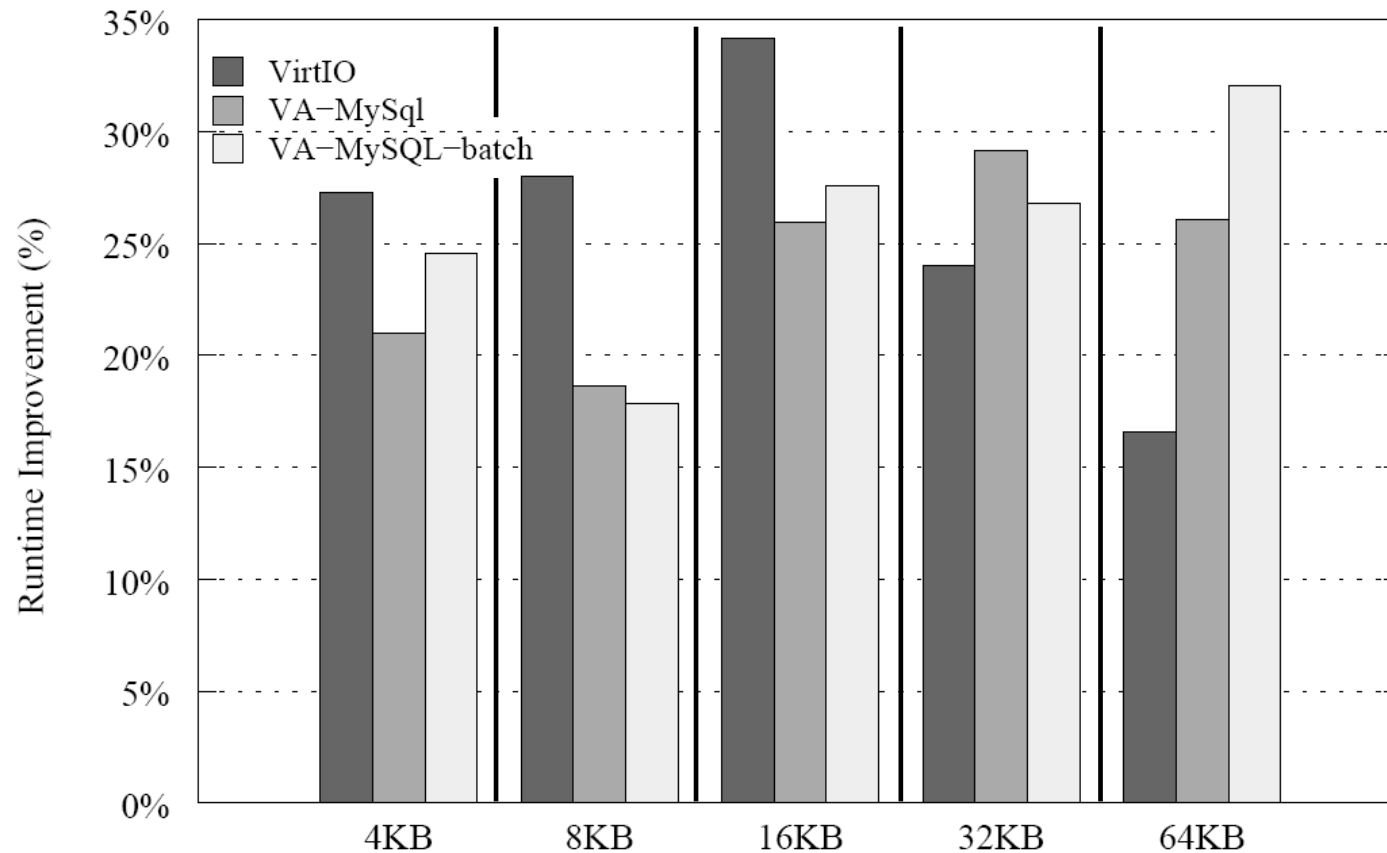
Guest/Host – Cycles Distribution



Experimental setup for different row sizes

- Guest cycles: still the same
- Host cycles: significantly reduced

Runtime Improvement



Tradeoff between amount of data and number of switches:

- VAMOS: number of switches depends on the request type
- Virtio: number of switches depends on the amount of data

Related Work

- Virtual Interface
 - Xen [Barham03]
 - HPC [Gavrilovska08]
 - Virtio [Russell08]

- OS Interface
 - VirtFS [Jujjuri10]
 - Libra [Ammons07]

- Hardware Interface
 - SR-IOV [Dong08, Liu10]

- VAMOS takes virtualization awareness up into userspace (Middleware)

Conclusions & Future Work

- Virtualization overhead is still high due to the transitions between the guest and the hypervisor
- Running part of the middleware at the hypervisor level, VAMOS reduces the overall number of guest/hypervisor switches and improves I/O performance
- Exploiting existing modular designs and abstraction layers, middleware can be adapted to run at the hypervisor level with **modest cost**
- VAMOS presents a new design point to be considered in the [transparency vs. performance] trade-off spectrum
- Next Steps:
 - Apply VAMOS to other middleware
 - Explore additional areas such as memory over-commit
 - Analyze feasibility of building a common infrastructure shared across different middleware
 - Improve middleware isolation and security
 - Guest/Host communication optimizations
 - What can we do if we re-think the middleware from scratch ?

Questions ?