

# Who Decides Migration?

## A Migration Lock Mechanism for Virtual Machines

---

**Xiaolin Wang, Yingwei Luo**

Dept. of Computer Science and Technology

Peking University

SVM'13, 18 October 2013, Zurich

# Agenda

---

- **Why migration lock?**
- **Design of migration lock mechanism**
- **Implementation**
- **How to use migration lock?**
- **Conclusion**

# Why migration lock?

---

- Migration is an powerful management feature**
  - Migrate virtual machines to difference hosts help to**
    - Improve resource utilization and save energy**
    - Achieve load balance and better QoS**
    - React in advance of failure or maintenance**
    - .....**

# Why migration lock?

---

## □ Defects of virtual machine migration

- May slow down the performance of VM during migration

- Network and disk I/O

- Memory access

- Destination host may lack of some hardware resources

- SR-IVO NIC

- General Purpose GPU

- Co-running with different VMs on the same host

# Why migration lock?

---

- **Scenarios in which migration will cause significant performance degradation or even system failure**
  - **The application is running some critical code and does not want to be affected by migration of the virtual machine;**
  - **The application is performing some critical operation, migration will fail the operation;**
  - **Some special optimizations, such as optimizing communication between virtual machines sharing the physical host, requires that several virtual machines run on the same physical machine at the same time;**
  - **Dependence on special hardware resource**

# Why migration lock?

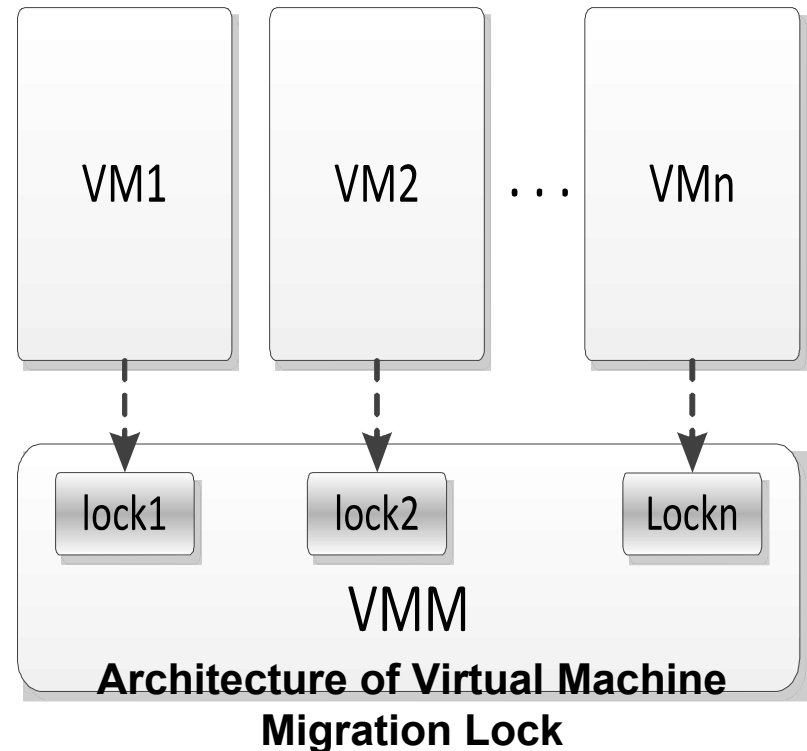
---

## □ Migration lock

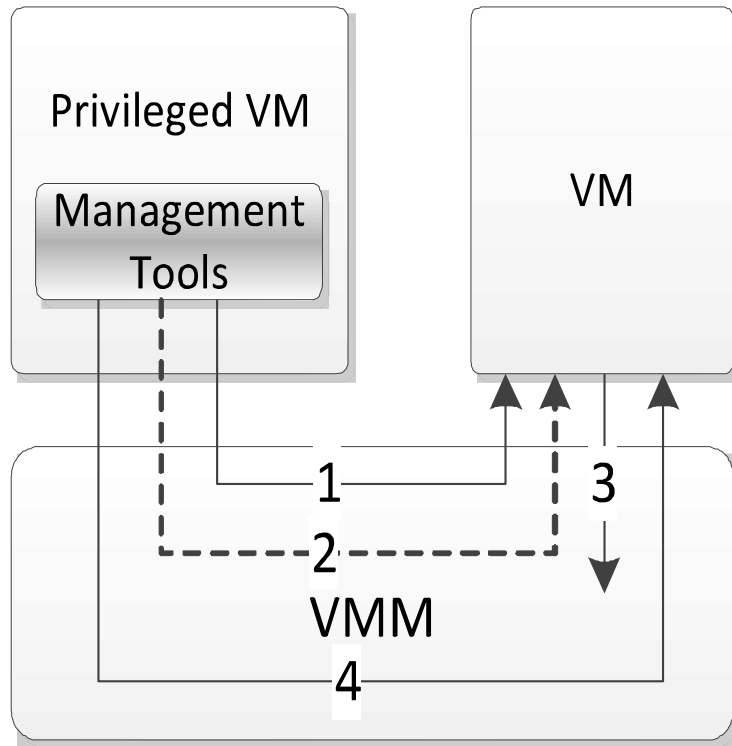
- Enable the VMM learn about whether a VM can be migrated.
- Enable applications in a VM tell the VMM do not migrate the enclosing VM.

# Design of migration lock mechanism

- A status flag added to the virtual machine.
  - VMM must check the flag before migrating a virtual machine.
- Static migration lock
  - Set the lock in configuration file or in the command line of creating the VM.
- Dynamic migration lock
  - Set or clear the lock while the VM is running.



# Design of migration lock mechanism



—————> Successful Operation  
-----> Unsuccessful Operation

- 1: Create a virtual machine and set it cannot be migrated.
- 2: Fail to migrate the virtual machine.
- 3: Set the virtual machine to be migratory during its running period.
- 4: Migrate the virtual machine successfully.

## Set Migration Lock Dynamically



# Implementation

---

- Implemented the migration locking mechanism on Xen-3.31, which involves change in:
  - Xen hypervisor
  - Guest OS of DomU
  - Application in DomU
  - Xen tools

# Implementation in Xen hypervisor

---

- The the structure *domain*
  - `bool_t is_migratory`
  - Setting or clearing migration lock reflect finally on it
  - Checking it before invoking migration
- A new hypercall
  - Hypercall number:  
`__HYPERVISOR_set_migraton_flag`
  - Hypercall handler function:  
`do_set_migration_flag`
  - With parrameter 1 for setting and 0 for clearing

# Implementation in Guest OS of Domain U

---

- To transfer the command of setting migration lock from the application level to the Xen hypervisor
  
- Different implementation for
  - Para-virtualized guest OS – Linux
    - System call
  
  - Full-virtualized guest OS – Windows
    - A special device driver

# Implementation in Guest OS of Domain U

---

## □ Para-virtualized guest OS – Linux

- 64bit Linux, kernel version 2.6.18
- System call with number 235
- Handler function - *sys\_set\_migration\_flag*
  - Which will call *\_\_HYPERVISOR\_set\_migraton\_flag*

# Implementation in Guest OS of Domain U

---

## □ Full-virtualized guest OS – Windows

### ■ In the special device driver

- Firstly, obtain the number of hypercall pages of Xen Hypervisor via CPUID instruction
- Then, allocate a block memory with the same size as the hypercall pages for setting new hypercall pages for Windows Operating System by using WRMSR instruction.
- Finally, invoke the hypercall from the handling function of *MldDDKDeviceIOControl*

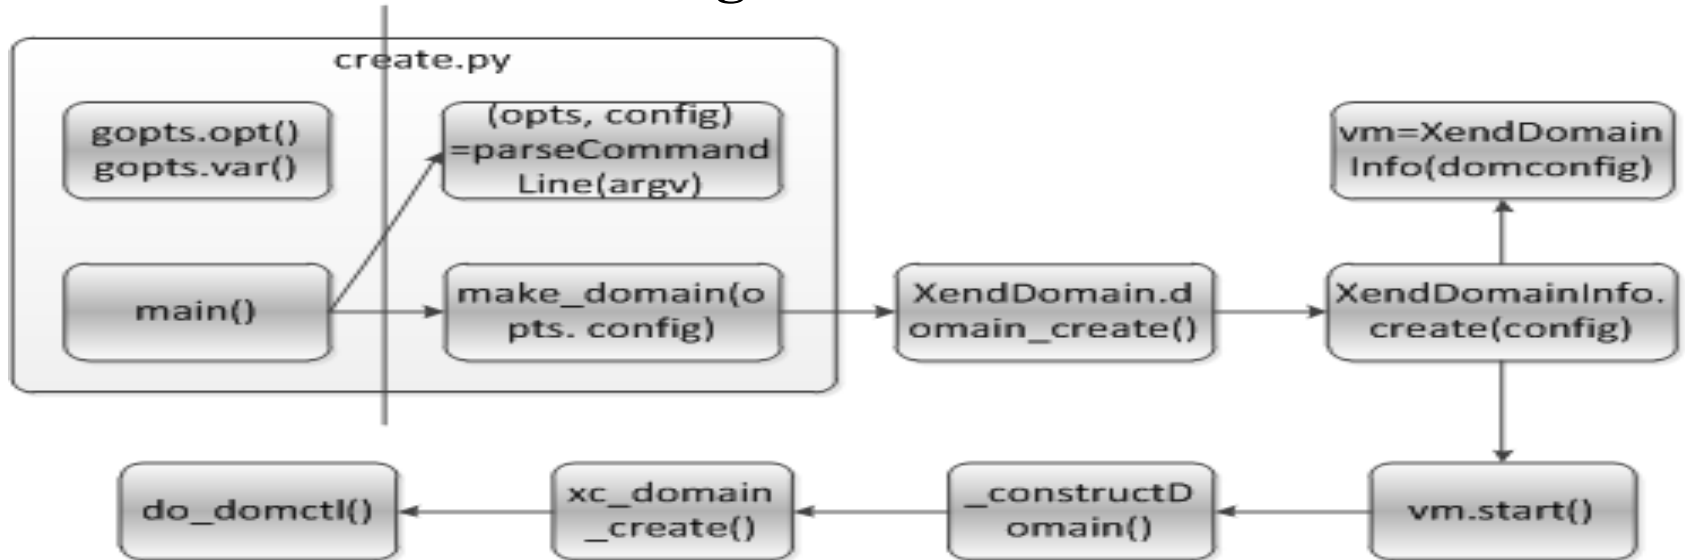
# Implementation in Application in Domain U

---

- An application tools (*set\_migration\_flag*) for users in Guest OS to set migration lock
  - On Para-virtualized guest OS – Linux
    - The tool invoke the system call to set migration lock
  - On Full-virtualized guest OS – Windows
    - The tool invoke the I/O control operation to set migration lock
  - CMD: *set\_migration\_flag flag (flag =1 or 0)*
- Users can also call invoke the system call on Linux or invoke the I/O control operation on Windows in their applications to set migration lock.

# Implementation in Xen tools

## In the Process of Creating Domain U



**CMD:**

*`xm create config_file -migration=yes|no`*

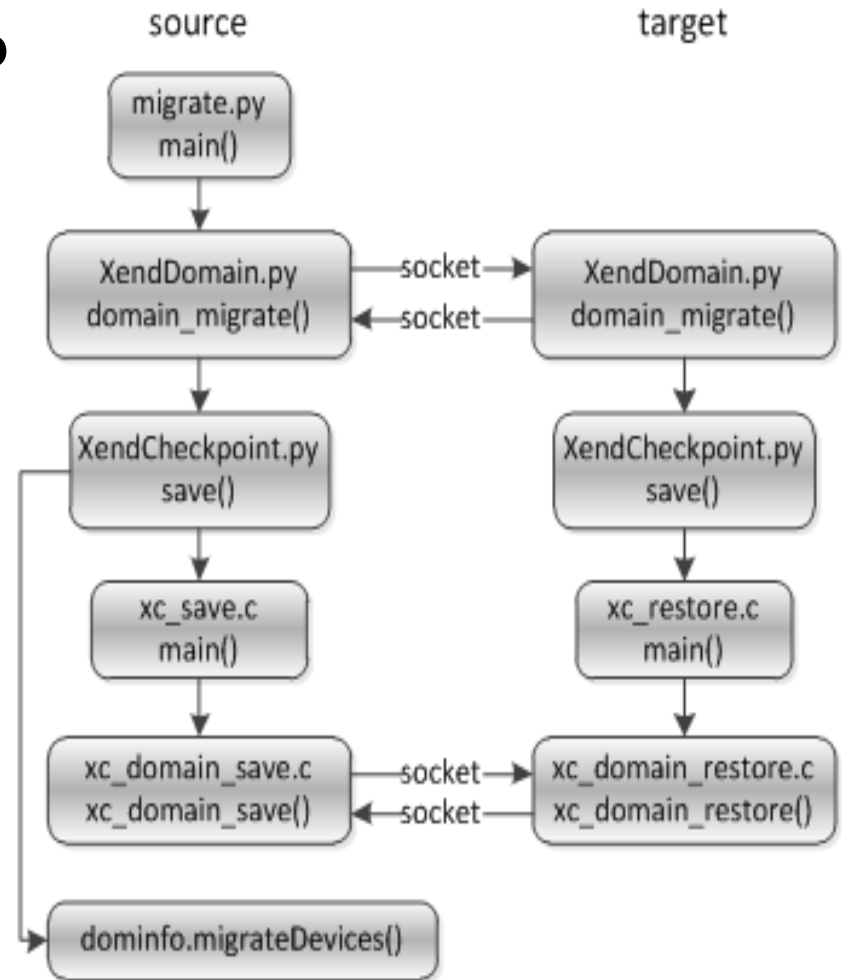
**Configuration file:**

*`migration = yes|no`*

# Implementation in Xen tools

In the Process of Migrating Do  
Checking the lock in two  
functions:

`domain_migrate`  
`xc_domain_save`





# How to use migration lock?

---

- **With migration locking mechanism**
  - **Virtual machine may prevent itself from migration on it's own initiative.**
  - **Current migration strategies should take migration lock into consideration.**
- **Two problems to be solved:**
  - **How to prevent a virtual machine from using migration lock arbitrarily?**
  - **How to make a migration schedule without knowing in advance whether or not a virtual machine can be migrated?**

# Ideas about how to use migration lock

---

## □ To avoid using migration lock arbitrarily

- **Specify in the SLA whether or not a virtual machine may set migration lock.**

Only allowed virtual machines may set migration lock; or some forbidden virtual machines are disable to set migration lock.

- **Tax on setting migration lock according to historical working load.**

To force user to use migration lock properly. To punish those selfish users.

- **Define a constant time slice as the maximum time that a migration lock may persist.**

Prevent a virtual machine hold the migration lock too long. Recover from application failures without releasing migration lock.

- **Provide *trylock* mechanism on migration lock for application.**

Application may have opportunity to avoid doing critical actions or binding to unpopular hardware resources when migration lock is unavailable. (for example, a virtual machine currently in migrating).

- **Atomic and consistent operation to set and check migration lock.**

Avoid data racing, consistent view of migration lock in both VMM and VM.

# Ideas about how to use migration lock

---

## □ To make good migration schedules with migration lock

### ■ Fail and retry strategies

Make migration schedules only on those virtual machines that can be migrated.

When a scheduled VM becomes locked, discard old schedules and try make new ones.

### ■ Probabilities model

With historical information data mining, make migration schedules considering the probabilities of a VM changing from unlocked to locked and vice versa.

### ■ Wait and retry strategies

With the help of *maximum lock slice*, wait for the lock to be released when a scheduled VM is locked when try to migrate it.

# Conclusion

---

- ❑ Scenarios in which migration will cause significant performance degradation or even system failure
- ❑ Migration Lock mechanism can work.
- ❑ Do we need Migration Lock?
- ❑ How to use migration lock?
- ❑ How migration management work together with migration lock?

---

# Thanks!

## Q & A