



PERFORMANCE ANALYSIS OF CLOUD COMPUTING PLATFORMS

Yves Junior BATIONO
December 2016

École Polytechnique de Montréal
Laboratoire **DORSAL**

OUTLINE


- .INTRODUCTION
- .RESEARCH OBJECTIVES
- .METHODOLOGY
- .COMPUTING SERVICE DIAGNOSIS
- .NETWORKING SERVICE DIAGNOSIS
- .STORAGE SERVICE DIAGNOSIS
- .MULTI LAYER ANALYSIS
- .FUTURE WORK



INTRODUCTION

- Complexity of cloud services
- Consumers experiment some services latencies
- Where to start troubleshooting ?
- Complete view of the cloud environment
- Correlate information from different nodes





Objectives

➤ Multi layer Analysis of cloud infrastructure

Services and virtual resources (application, virtualization, kernel layer)

➤ Show OpenStack service efficiency

show interaction and service bottleneck

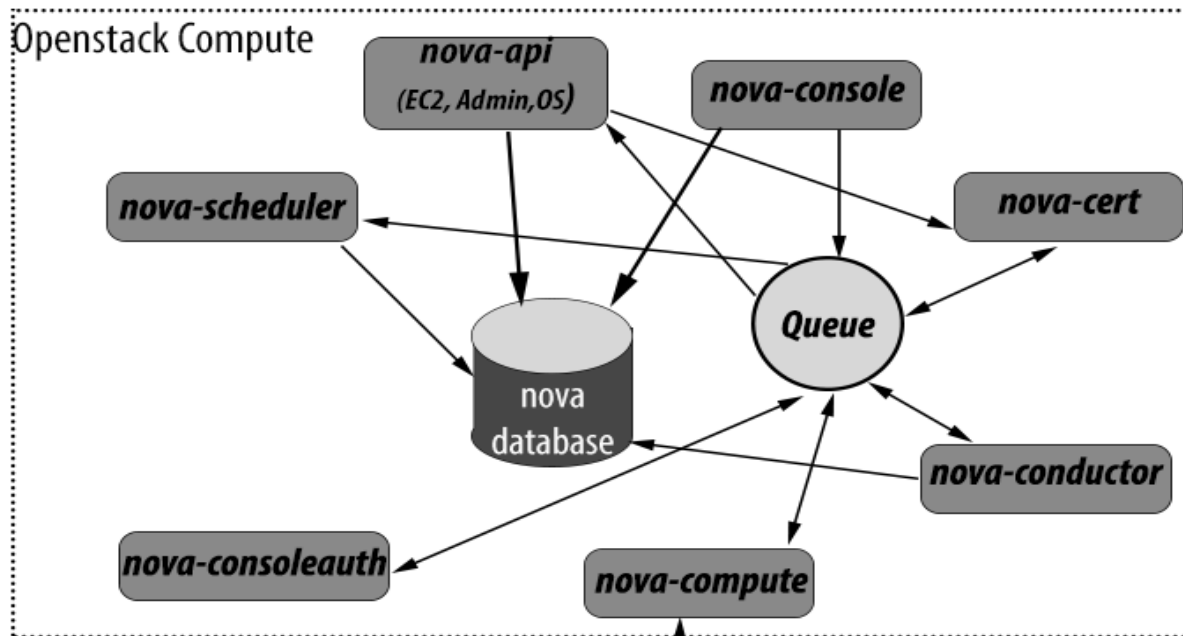
➤ Correlate cloud nodes information

Understand execution failures



➤ 1st Method

Survey OpenStack services interaction through RabbitMQ
This method is **not efficient**.



➤ 2nd Method

OpenStack services instrumentation
OpenStack services tracing with LTTng



TYPE OF SERVICES

Cloud Computing refers to three kind of services

- **Computing services**

Provide virtual machine to consumers

- **Storage services**

Persistent storage for virtual machines

- **Networking services**

Deploy application for virtual machines communication

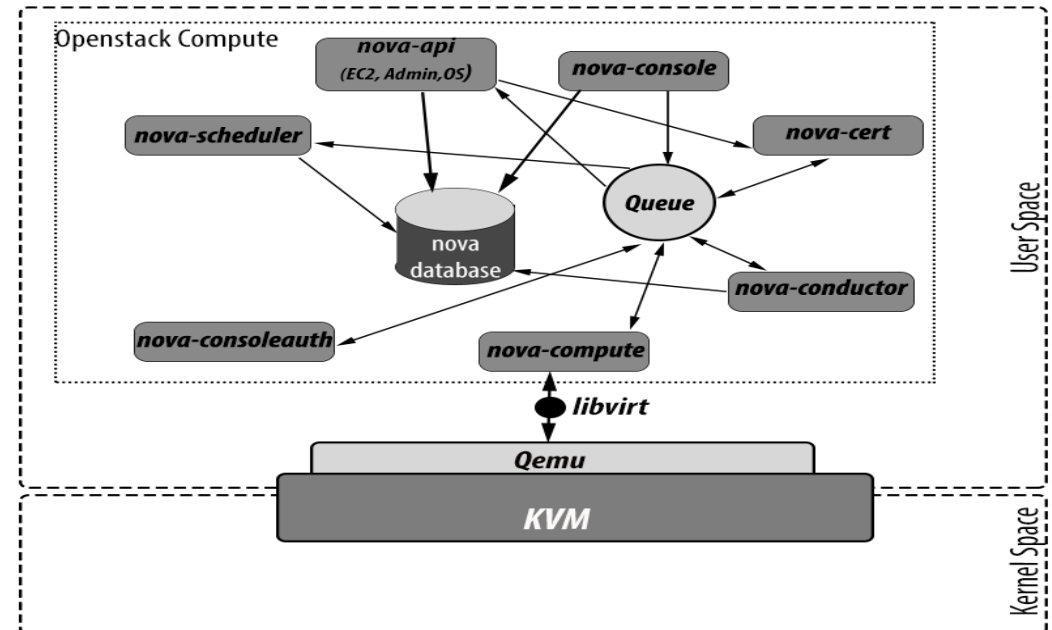


COMPUTING SERVICE DIAGNOSIS

- Provide virtual machines
- shared computer processing resources

Computing Service Layers

Layers	Features
Application	Nova
Virtualization	QEMU/KVM
Kernel	Host Kernel



Application Layer Diagnosis

OpenStack Nova Instrumentation

- Instrumentation is based on Nova logging activities
- LTTng-UST with python binding to collect logs
- Write log in JSON format to retrieve instance attributes

```
Log.Trace(trace_nova(event_type, instance, message))
```

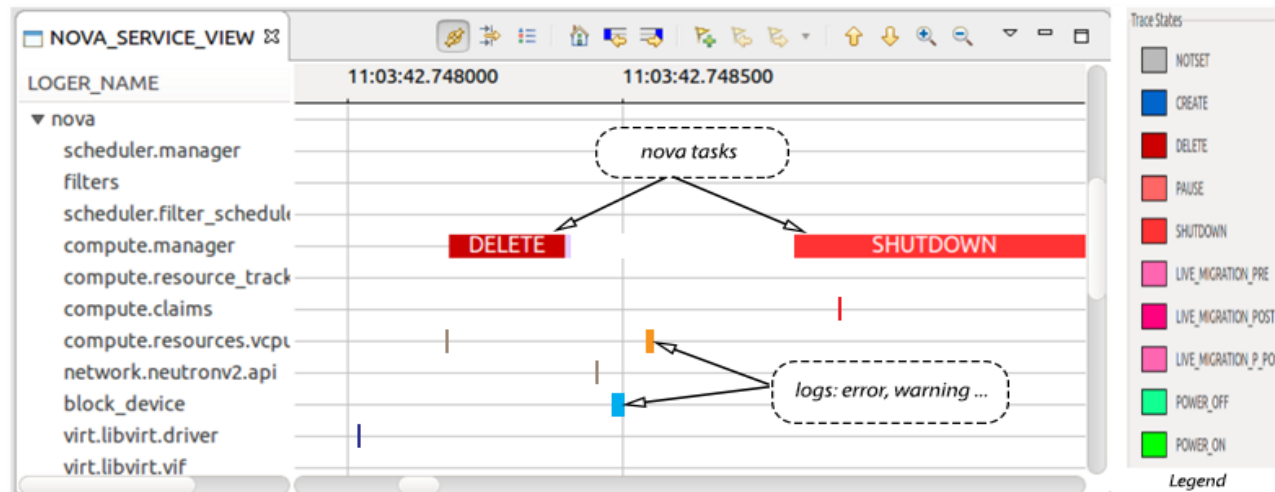
message output in JSON

```
{
  "EVENT": {
    "event_type": event_type
    "vm_name": instance.name
    "state": instance.state
    ....
    "host": instance.hostname
    "request_id": instance.context.req_id
    "msg": message
  }
}
```



Application Layer Diagnosis

OpenStack Nova Trace analysis

- VM state investigation
- Service performance analysis
- Survey request flow



Nova Service View



Virtual Machine state View

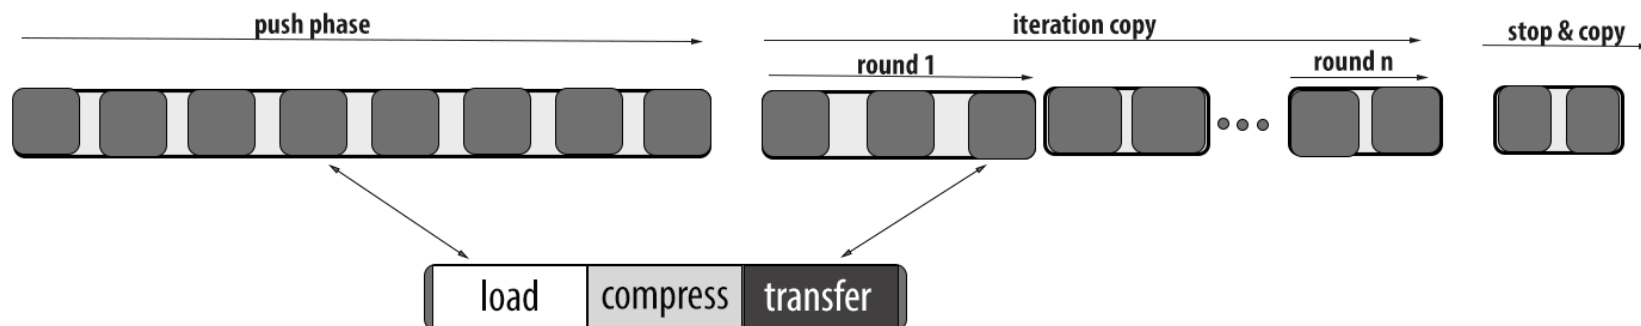
Virtualization Layer Diagnosis

➤ QEMU Instrumentation

Already instrumented!

➤ QEMU Trace analysis

Live migration stage: precopy, iteration copy, stop and copy.

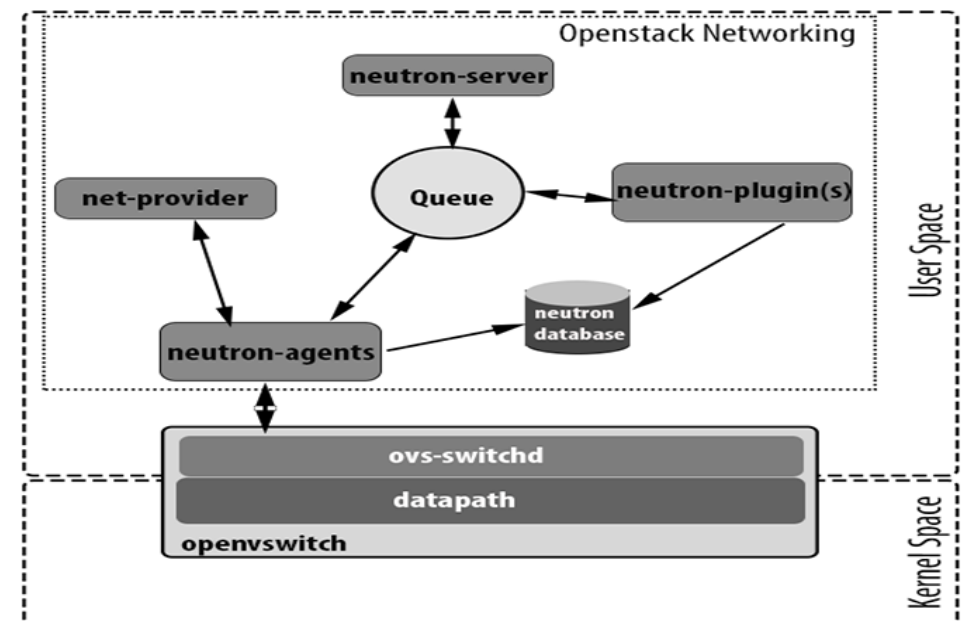


NETWORKING SERVICE DIAGNOSIS

- Virtual machines communication
- Responsible of virtual network, routers... creation

Networking Service Layers

Layers	Features
Application	Neutron
Virtualization	Open vSwitch
Kernel	Host Kernel



Application Layer Diagnosis

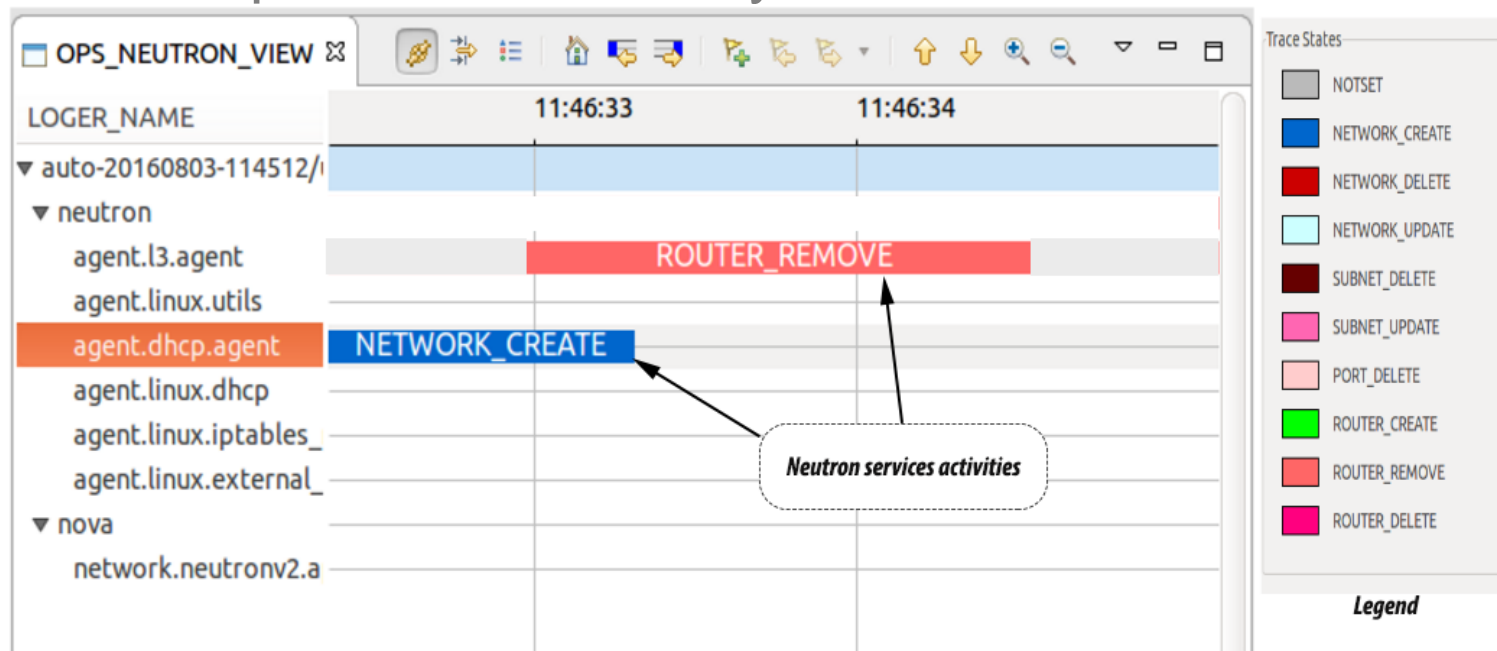
➤ OpenStack Neutron Instrumentation

Instrumented with as Nova with Log in JSON format

➤ OpenStack Neutron Trace Analysis

Survey request flow

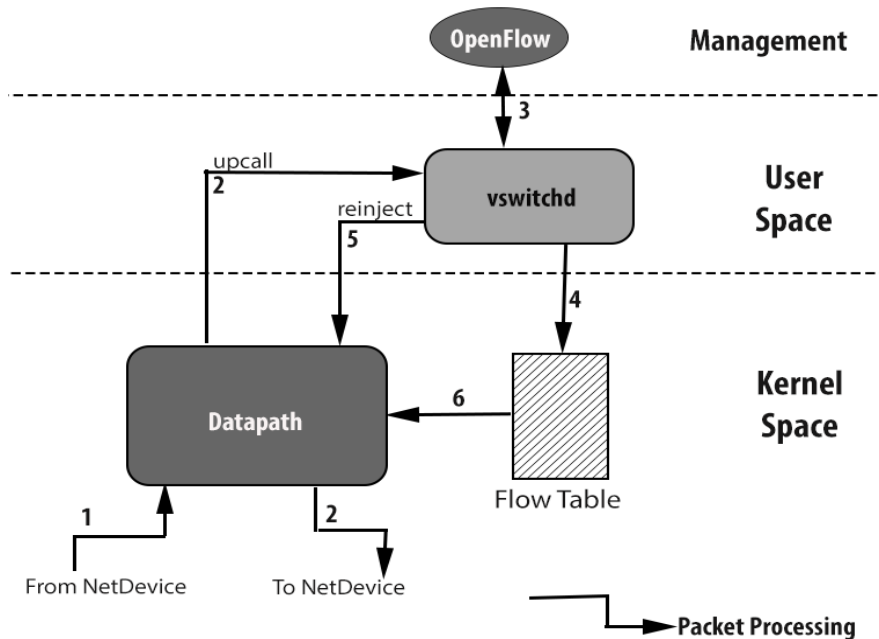
Service performance analysis



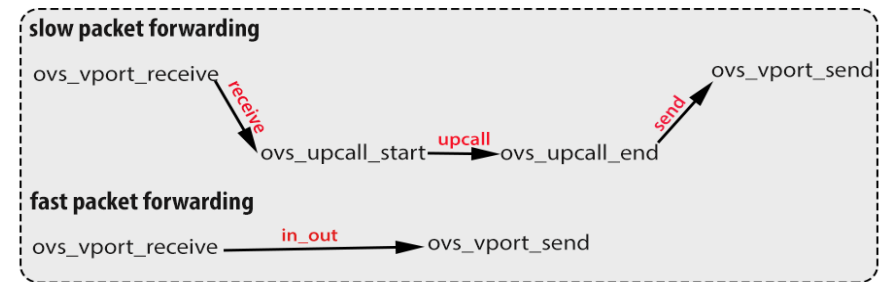
Virtualization Layer Diagnosis

Open vSwitch Instrumentation

- Dynamic instrumentation based on kprobe
- Use LTTng addons to instrument datapath functions
- Tagging Packet tail room



Open vSwitch Packet Processing

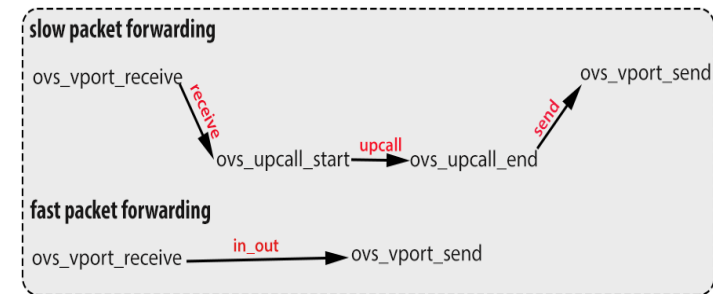


Packet forwarding events


Virtualization Layer Diagnosis

Open vSwitch Trace Analysis

- Packet switching performance
 - Packet latency through OvS components
 - Frequency of network topology configuration
- OvS components interaction
 - Datapath/ovs-vswitchd communication



Packet forwarding events




Open vSwitch View

Storage Service Diagnosis

- Virtualize storage devices
- Provide storage resources to consumers
- Persistent storage to virtual machines

Storage Service Layers

Layers	Features
Application	Cinder
Virtualization	-
Kernel	Host Kernel



Application Layer Diagnosis

➤ OpenStack Cinder Instrumentation

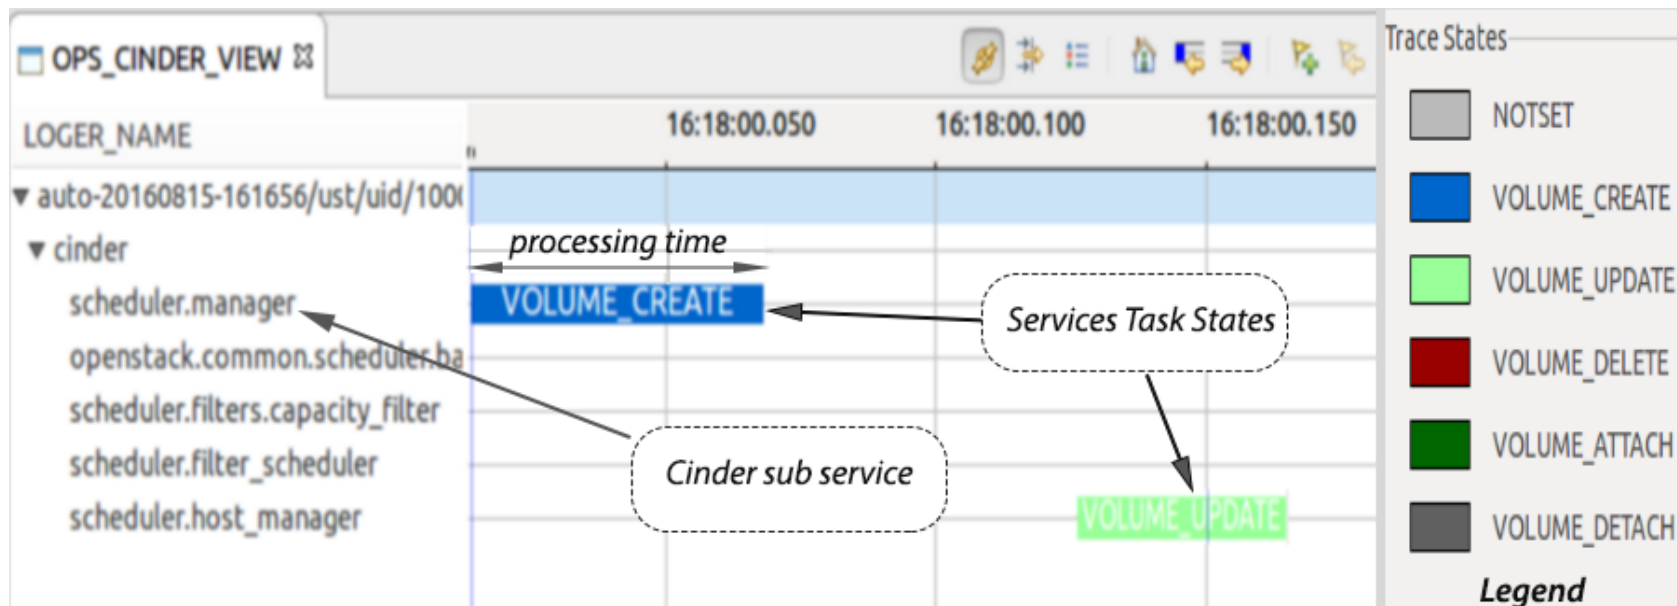
Instrumented with as Nova with Log in JSON format

➤ OpenStack Cinder Trace Analysis

Survey request flow

Performance of Cinder operations


Link virtual volume to physical disk



Cinder View

MULTI-LAYER ANALYSIS

- Link Instances data from the three layers
- For each operation from the Application layer, find sub-operation in the Virtualization and the Kernel layer



Live Migration

Normal case: live migration succeed

vm-29631 live migration


Stages of VM migration (precopy, iteration and stop©)

Downtime=stop©


Round and sub-round for each stage

Dirty page curve converge to 0=migration succeed.

Qemu State View



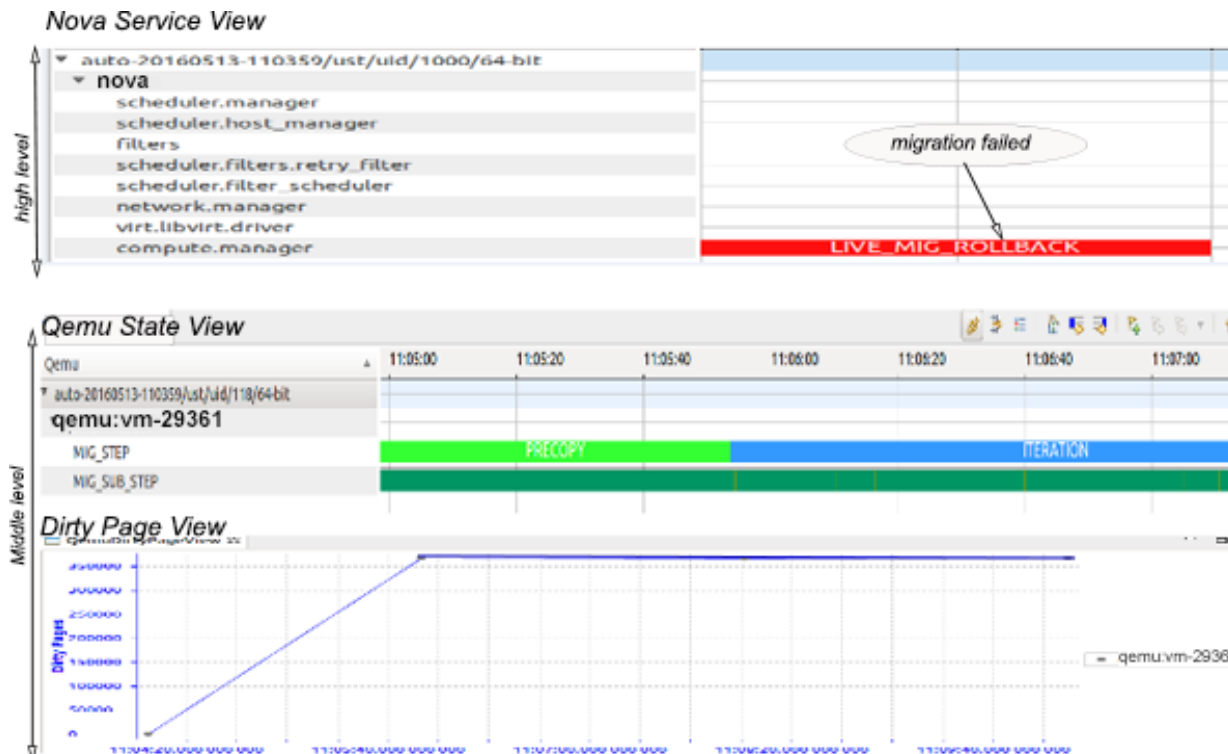
Dirty Page View



Live Migration

Abnormal case: live migration fail

- Live migration rollback; stop© stage is missing
- Dirty page view does not converge to 0
- vm-29361 has a high workload
- Qemu cannot transfer quickly the memory



FUTURE WORK

- OpenStack service diagnosis
associate Nova, Neutron, Cinder interfaces in a single utility
- OvS diagnosis
use packet source and destination address to locate VM



QUESTIONS

yves-junior.bationo@polymtl.ca

