

RED HAT :: CHICAGO :: 2009

SUMMIT

FOLLOW US:

[TWITTER.COM/REDHATSUMMIT](https://twitter.com/redhatsummit)

TWEET ABOUT US:

ADD #SUMMIT AND/OR #JBOSSWORLD TO THE END
OF YOUR EVENT-RELATED TWEET

presented by



RED HAT :: CHICAGO :: 2009

SUMMIT

Building and Leveraging Compute Clouds with Red Hat Enterprise MRG

Bryan Che, Red Hat Product Manager

Shawn Lower, Red Hat GPS Consultant

Mike Santangelo, Red Hat GPS Consultant

Chad Tindel, Red Hat Solutions Architect

presented by



Common Cloud Questions

Cloud computing is a hot topic, but many people have important questions and challenges they need addressed before they can adopt cloud:


How do I build an internal cloud?

How do I avoid lock-in to a single cloud?


How do I mix, match, and blend different cloud resources—including internal and external clouds?

How do I manage a variety of applications and groups with different SLAs, priorities, and resource requirements across clouds?

How do I manage and track cloud resources?



Red Hat Enterprise MRG provides the Answers



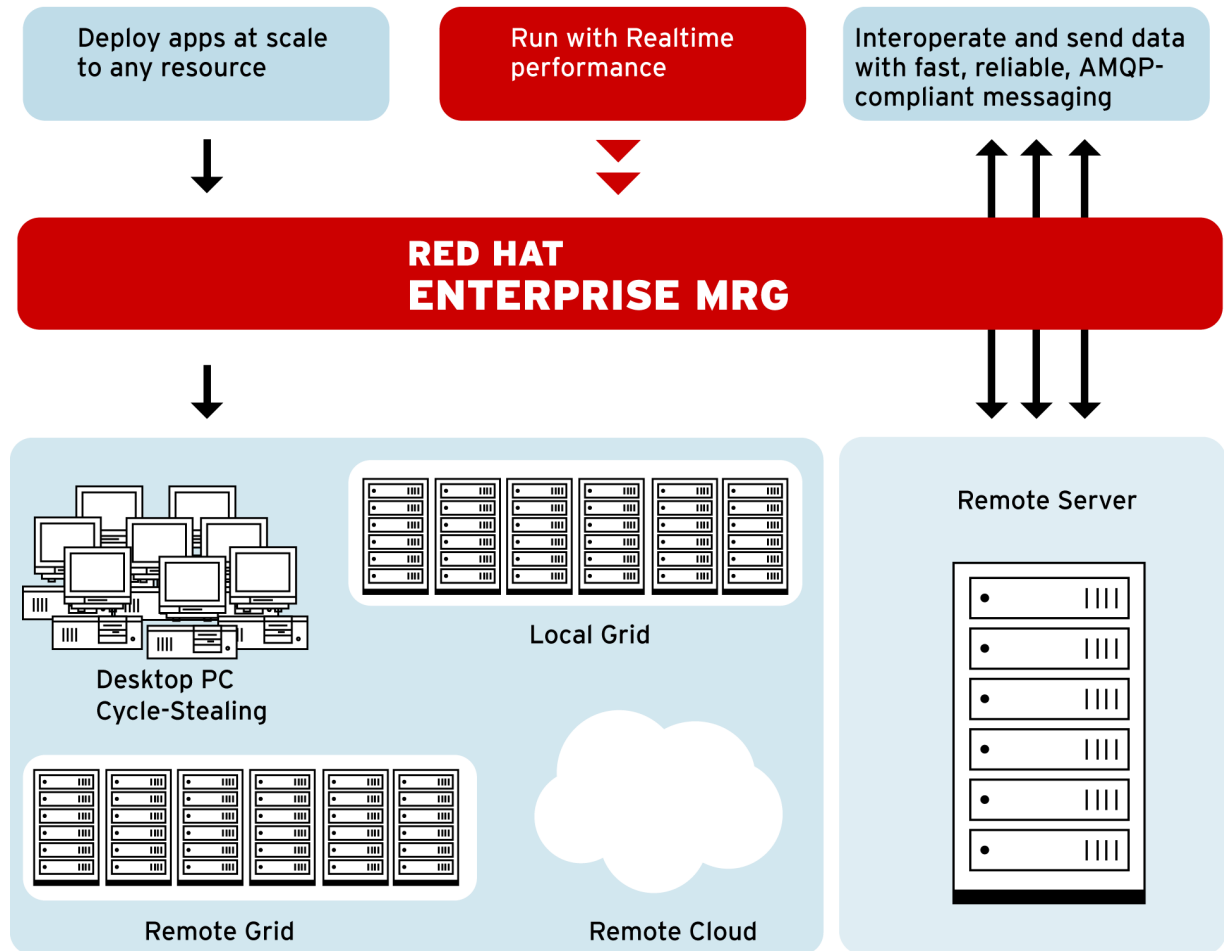
About Red Hat Enterprise MRG

Integrated platform for high performance distributed computing

High speed, interoperable, open standard AMQP **Messaging**

Deterministic, low-latency **Realtime** kernel

High performance & throughput computing **Grid** scheduler for distributed workloads and Utility/Cloud computing



MRG Grid is Based on Condor

MRG Grid is based on the Condor Project created and hosted by the University of Wisconsin, Madison

Condor has a >20-year history and runs many of the largest grids in the world

Red Hat and the University of Wisconsin have signed a strategic partnership around Condor:

University of Wisconsin makes Condor source code available under OSI-approved open source license

Red Hat & University of Wisconsin jointly fund and staff Condor development on-campus at the University of Wisconsin

Red Hat and the University of Wisconsin's partnership will:

Add enhanced enterprise features, management, and supportability to Condor and MRG Grid

Add High Throughput Computing capabilities to Linux



Building Clouds with MRG

Scalable Virtualization

Schedule VMs directly as jobs via libvirt

Provision VMs via Red Hat Enterprise Virtualization

Inject jobs into VMs

Resource Accounting

Track resources via Condor's resource accounting

SLA's

Apply priorities and policies

Apply security – Authentication (e.g. SSL, ...), Integrity, Encryption

Powerful Policies

VMs – run multiple concurrent instances, start on Black Friday or semi-monthly, re-run after fault

Machines – only run VMs from owner's group between 9 and 5, everyone else has a low priority shot from 5 to 9

Global – control limiters (e.g. NFS mount users, licenses),

Various Cloud Services

IaaS clouds: run all workloads as VMs

PaaS clouds: leverage job scheduling with VM scheduling

Aggregating and Bridging Clouds with MRG

MRG includes the ability to schedule jobs and applications to multiple clouds, based on policy

- MRG has the ability to send VMs to other resource managers

- MRG becomes the unified interface to many types of resources – internal VM resources and multiple external clouds

- MRG's life-cycle management, accounting and policy benefits still available

Use cases include

- Manage overflow/spillover

- Access to specialized resource managers

- Transformation between VM types/systems

- Allow a single app/stack to bridge multiple clouds

MRG Cloud Aggregation Architecture

Schedd accepts jobs over SOAP, AMQP, CLI

GAHP: Grid ASCII Helper Protocol

An adapter to an external resource manager

Exist for many batch systems

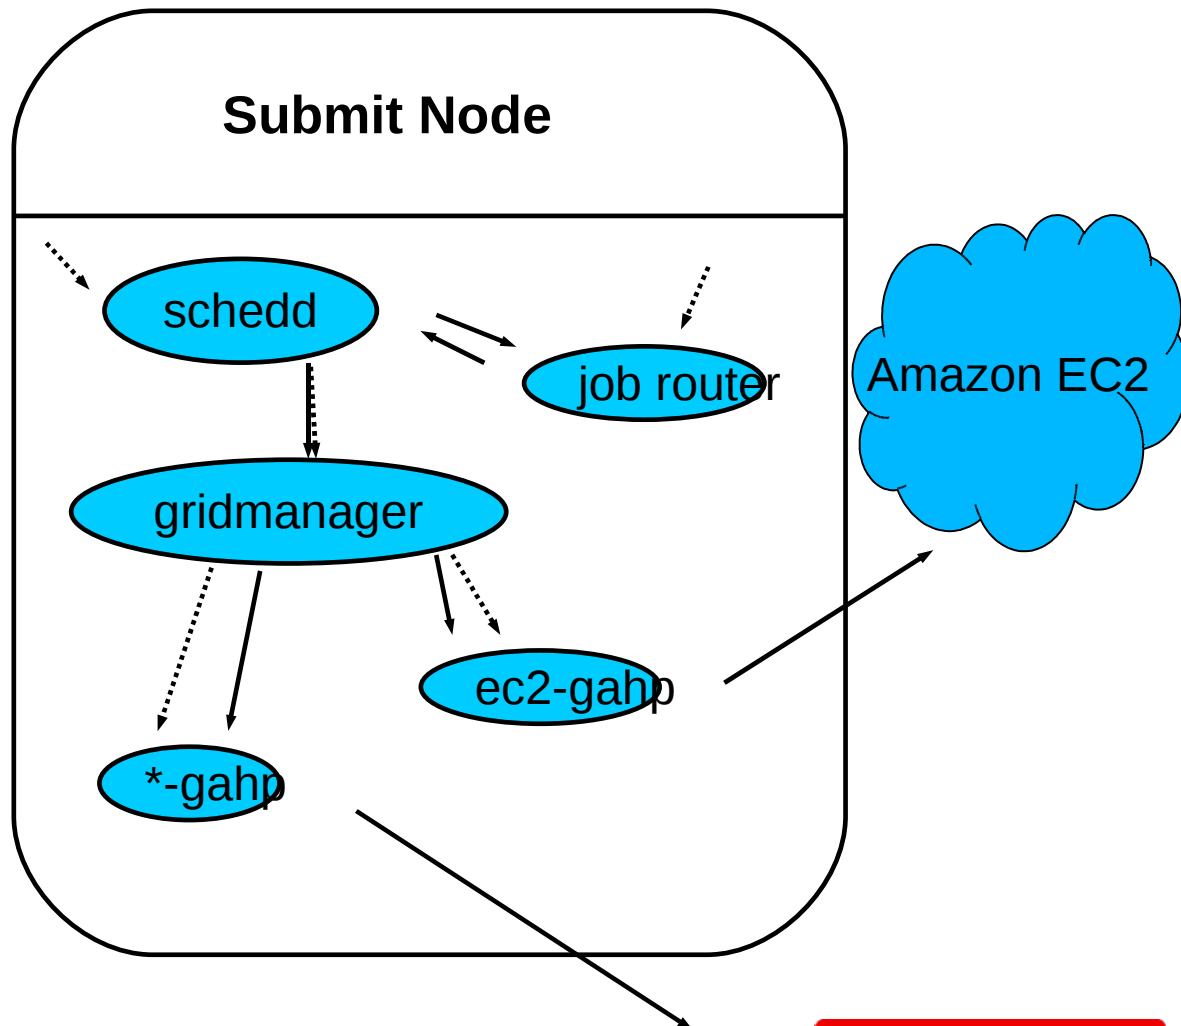
Exists for EC2-like resource managers

Extensible to new resource managers

Job Router transforms types, e.g. stack to VM to EC2 AMI

.....> = Process Spawned

————> = Communication



Condor, Globus, etc

Task: Render “Elephants Dream”

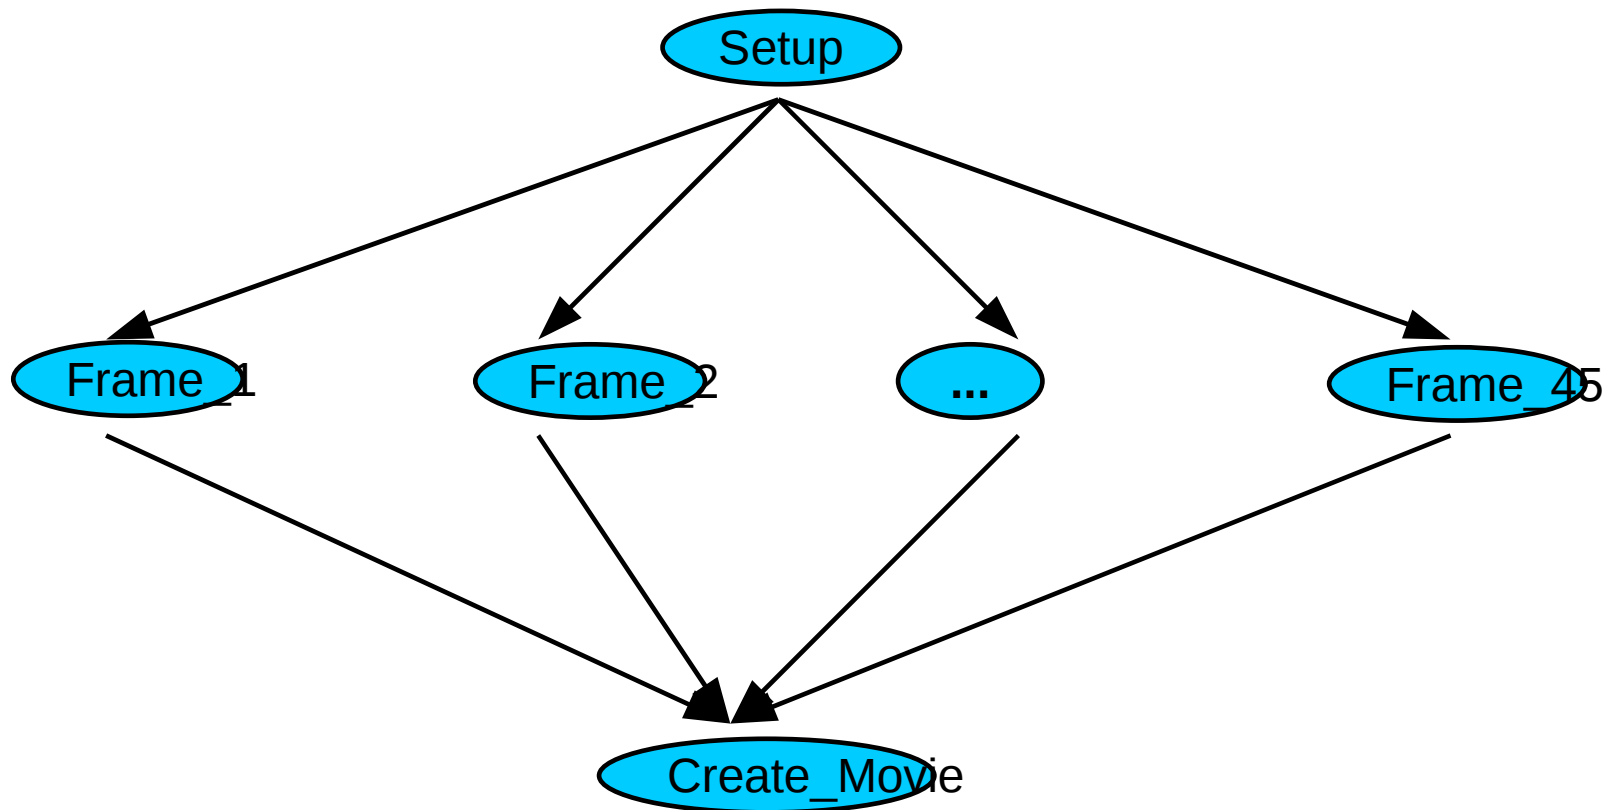
Freely Available, Open Content

Created with Blender, open
source rendering tool

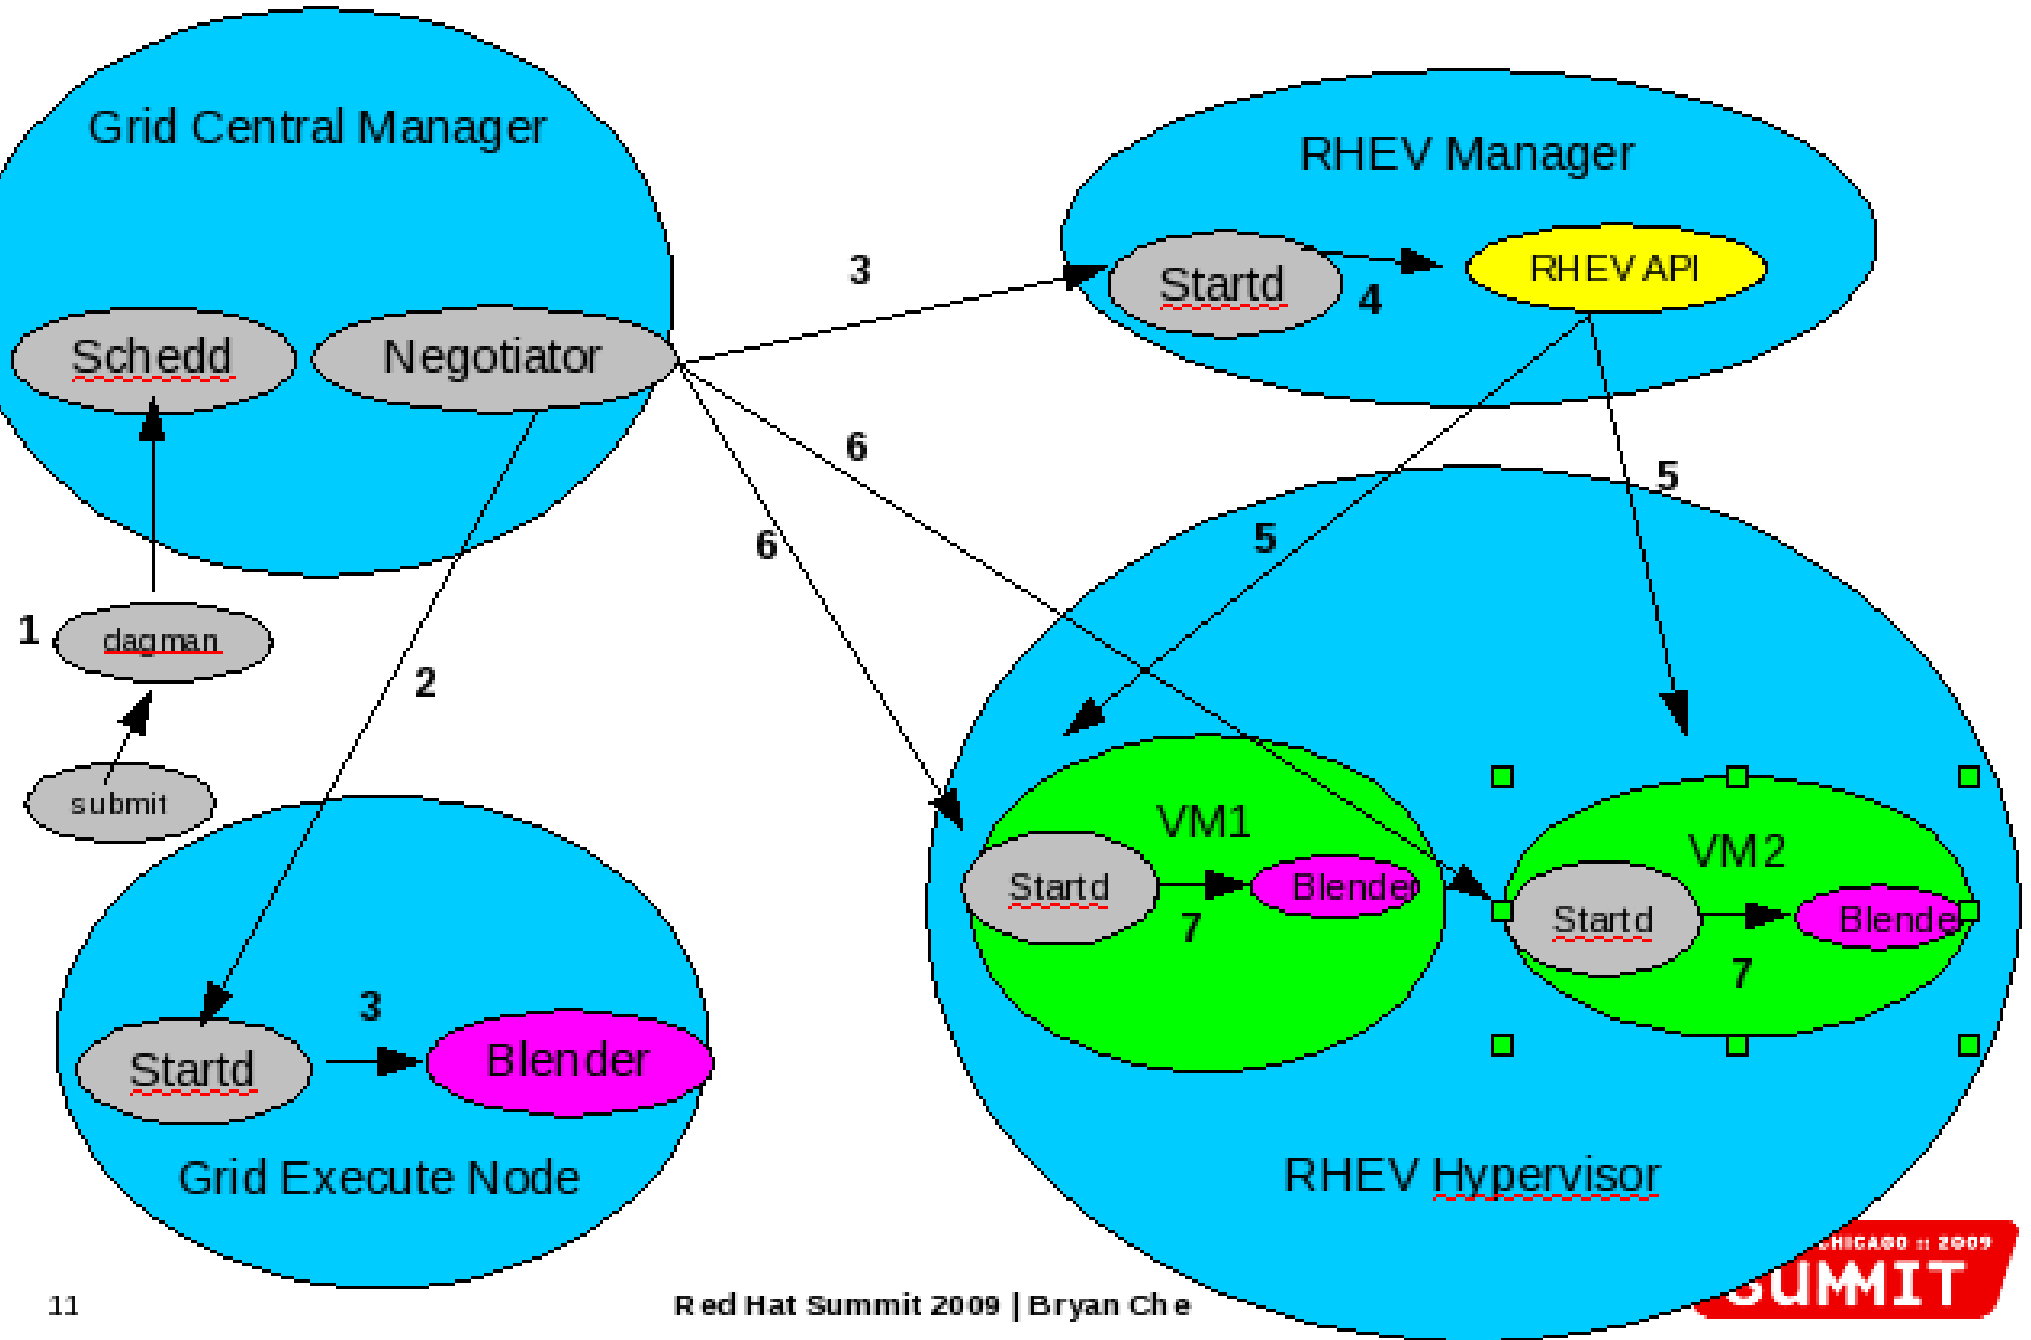
We will render a 45-frame
sequence, roughly 2 seconds
of the movie



Render DAG Definition



MRG-RHEV Architecture



Frame Rendering Job Submit File

Universe = vanilla

Executable = /usr/bin/blender

**Arguments = -b /mnt/render/source/production/07_emo_flips_out/07_01.blend \
-F PNG -o /mnt/render/completed/frames/frame_###.png -f 1**

Log = /home/griduser/elephant/frame1.log

Output = /home/griduser/elephant/frame1.out

Error = /home/griduser/elephant/frame1.err

transfer_executable = false

Requirements = Arch =?= "INTEL"

QUEUE

Movie Creation Job Submit File

Universe = vanilla

Executable = /usr/bin/mencoder

**Arguments = mf:///mnt/render/completed/frames/*.png \
-mf fps=5:type=png -ovc lavc -oac copy \
-o /mnt/render/completed/movies/output.avi**

Log = /home/griduser/elephant/create_movie.log

Output = /home/griduser/elephant/create_movie.out

Error = /home/griduser/elephant/create_movie.err

Requirements = Arch =?= "INTEL"

transfer_executable = false

QUEUE

Render DAG Definition File

JOB Setup render_setup.sub

JOB Frame_1 frame1.sub

JOB Frame_2 frame2.sub

.....

JOB Frame_45 frame45.sub

JOB Create_Movie create_movie.sub

PARENT Setup CHILD Frame_1 Frame_2 ... Frame_45

PARENT Frame_1 Frame_2 ... Frame_45 CHILD Create_Movie

RHEV API

```
Get-VmPools
```

```
Get-Vm [-VmId]
```

```
Start-Vm [-VmId] <Guid> [-DisableHardwareAcceleration] \  
        [-RunAndPause] [-AcpiDisable] [-BootDevice <String>]
```

```
Suspend-VM [-VmId]
```

```
Shutdown-Vm [-VmId] [-WaitBeforeShutdown]
```

```
Stop-Vm [-VmId]
```

Extend Grid Execution into RHEV Compute Cloud

Spin up a remote pool

Using a power shell “glue” script
to integrate MRG Grid with
RHEV

```
$vmPool = Get-VmPools | Where-Object {$_.Name -eq $poolName}  
$poolId = $vmPool.VmPoolId  
$vms = Get-Vm -PoolId $poolId | Where-Object {$_.Status -eq "Down"}  
      | select-object VmId  
  
foreach ($VmId in $vms) {  
    Start-Vm -VmId $VmId.VmId;  
}
```

MRG Cloud Roadmap

MRG already ships cloud capabilities, including:

- Scheduling VMs

- Scheduling to Amazon EC2

- Workflow management

- Policy management

Upcoming MRG features for cloud include:

- RHEV-M integration via QMF for provisioning VMs

- Updated management tools

 - (Demo included working preview of upcoming MRG 1.2 UI)

- Template policies and workflows for cloud

- Integration with other Red Hat cloud technologies (DeltaCloud, Hail, etc)

Additional Information

- www.redhat.com/mrg
- www.youtube.com/watch?v=oSm7Ff8kKjk



QUESTIONS?

**TELL US WHAT YOU THINK:
[REDHAT.COM/SUMMIT-SURVEY](https://redhat.com/summit-survey)**