



Innovation in the Community

Carl Trieloff

Senior Consulting Software Engineer

Director of Engineering for MRG

Background

Red Hat is investing in a set of defining technologies, actively being developed in conjunction with partners. These key projects include among others...

- Grid/ Distributed computing (Partnership with UW around Condor)
- Real-time (provides pre-emption and latency determinism)
- Virtualization (provides scale, VM, and ability to provide environment with Job)
- Distributed computing / Messaging (built on AMQP) www.redhat.com/mrg
incubator.apache.org/qpid www.amqp.org
- Stateless (ability to execute from read only image/device) et.redhat.com
- FreeIPA (Identity, Policy Audit) www.freeIPA.org
- just to name a few

The combination of these types technologies in a standard Linux distribution will provide a new generation of compute platforms.

Today's trends, tomorrow's platform

It is always interesting to see where the innovation is happening in the market, as these trends give us clues of what tomorrow will look like, and what next will be commodity..

- Pay for compute in a cloud (Amazon EC2 , Google, SUN, Red Hat - EC2, etc)
 - The desire to pay for work executed - SAS model today does not scale well outside custom / Open Source application
 - Many vendors have not figured out how to make money in this model yet
- Virtualization - run anywhere (Red Hat, rPath, Xen, WMWare)
- Large community projects
 - OLPC - it has been proved that you can build a \$150 computer
 - GNU, Eclipse.org, Apache Foundation, Fedora, etc
- Free SAS applications - google maps, google xyz...
- Software development philosophies are changing
- **Management = Adoption**

The way people are consuming software is changing.. what is the perceived value of software is changing - estimated 300 Million lines of code in open source

So what is MY point??

Right time to be building community software

- EDU understood this years ago, corporate is waking up 20 years later
- Invest back into the projects you use
- Over time the community will develop everything you try to do proprietary
- Ask the question - what is best for the community?

Don't get caught with your pants down!

- I believe we should all be spending more time thinking about how we want to consume software
 - Note that the answer to this changes very much for every use-case
 - Am I saying everything will be in the cloud - NO.
 - Doing it outside the community always fails - RT examples

Working together we can provide the best compute platform.

- UW-Madison, Red Hat, community, etc are doing this around Condor
- Hope you are too...

OK - enough philosophy ... what is happening in the out there??

Sleepwalker - codename

You might not know, but Condor is in Fedora, just yum install it to get the Condor 7.0.x series

Bringing the Community on-line

- Sleepwalker is a project to bring x Million Fedora users on-line - donate resources based on preferences to help the project
- Why?
 - Community understands benefits of helping each other
 - Put resources on-line for building, custom spins, testing, any project that helps the Fedora community, or community Fedora decides to help
 - Is this not just a MASSIVE version of a cloud?
- **What does this mean for Condor**
 - **How do we manage 3.5+ Million machines?**
 - **How do we determine how much resources where used?**
 - **How do we prevent miss-use?**
 - **How do we get good utilization?**
 - **How do we handle the scale?**

Does this not sound like the same issues every project has, it is about harnessing the community...

EC2 – and the cloud

Few notes – but attend Mike's presentation:

- Clearly communities and foundations are not going to give free time to commerce, they might for good will and research
- This joint work between UW, Red Hat, and Amazon on EC2 integration provides the same cloud benefits for commercial users
 - Grab Matt, or one of the UW guys if you want all the technical details.
 - www.redhat.com/mrg to play with it
 - libvirt integration (bare metal to appliance – same image)

My to cents:

➔ We see an increasing number of companies working with us on these aspects of MRG for:

- ➔ Extra capacity
- ➔ Cost per calculation (cents per hour model) – Condor on EC2 is not interesting – what is, is dynamically putting work onto EC2
- ➔ Startups, build my business in a VM, provision my business in the cloud

Fast scheduling – Low latency schedule

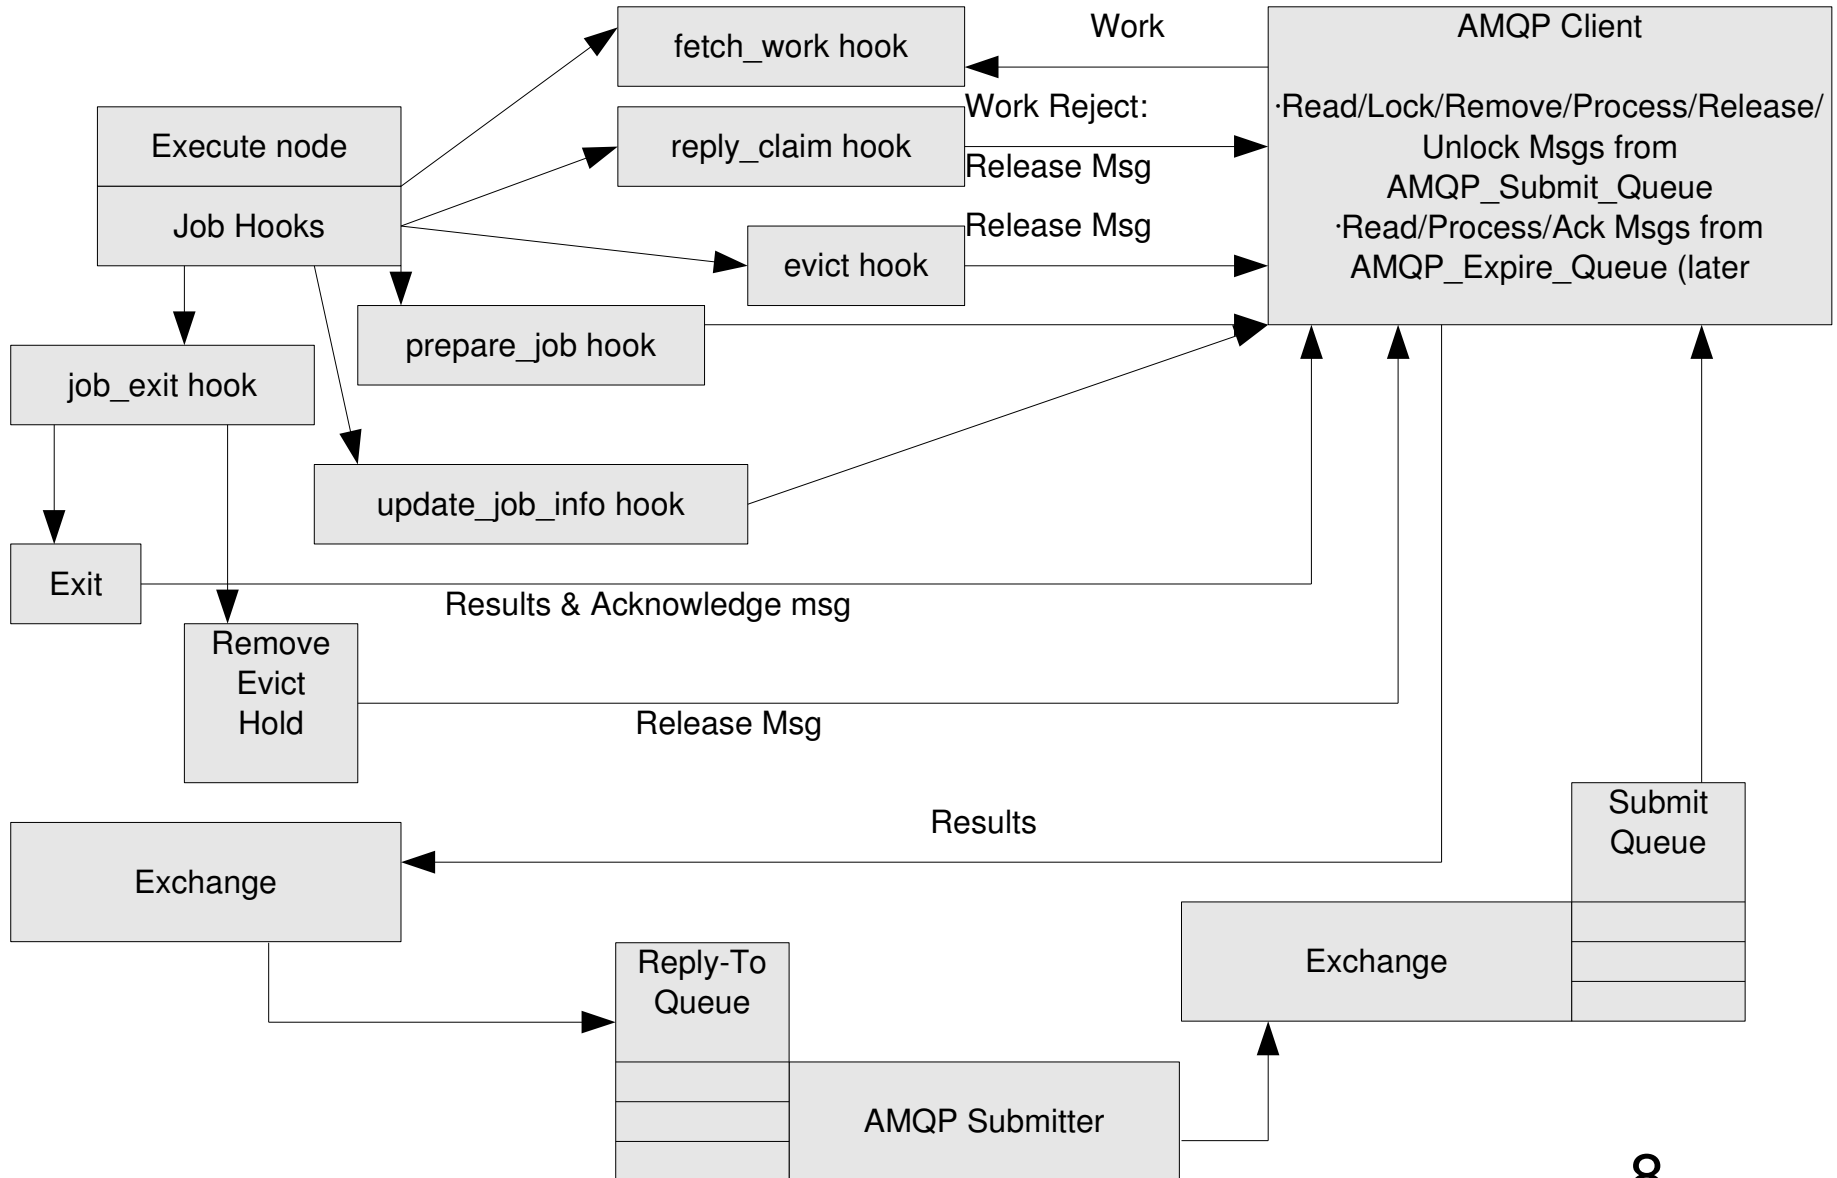
This uses the hooks

--- “That Derek Presented” ---

yes we are using it !!

- By-passes Condor traditional scheduling
- Integration with AMQP – and MRG
 - Allow for work-fetch jobs
 - Support Messaging based paradigm for job submission
- So you can submit through any messaging fount-end,
 - i.e. Use the C# client to submit / get job data from excel.
 - Integrated with any other application (java, C++, Python, Ruby, etc)

What does it look like?



- Working with UW so that it is not yet another custom solution
- Is it just the data from the collector?
 - Instrumentation
 - Configuration
 - Control
- * ... It's all about schema...*

Sec	Depth Range	Enq Rate	Deq Rate
10.01	0..4	3.10	3.20
10.01	0..1	0.30	0.30
10.14	0..3	1.38	1.18
9.86	0..40	4.16	4.16
9.86	0..0	0.00	0.00
9.86	0..1	0.10	0.10
9.86	0..0	0.00	0.00
9.86	0..1	0.10	0.10
9.86	0..6068	34148.37	34148.37
10.00	0..24	2.40	2.40

Commands:

- list - Print summary of existing objects
- list <className> - Print list of objects of the specified class
- list <className> all - Print contents of all objects of the specified class
- list <className> active - Print contents of all non-deleted objects of the specified class
- list <list-of-IDs> - Print contents of one or more objects
- list <className> <list-of-IDs> - Print contents of one or more objects of the specified class
- list is space-separated, ranges may be specified (i.e. 1004-1005)
- call <ID> <methodName> [<args>] - Invoke a method on an object
- schema - Print summary of object classes
- schema <className> - Print details of an object class
- set time-format short - Select short timestamp format (default)
- set time-format long - Select long timestamp format
- quit or ^D - Exit the program

qpid: schema system

Schema for class 'qpid.system':

Element	Type	Unit	Access	Notes	Description
sysId	uuid		ReadCreate	index	
osName	short-string		ReadOnly		Operating System
nodeName	short-string		ReadOnly		Node Name
release	short-string		ReadOnly		
version	short-string		ReadOnly		
machine	short-string		ReadOnly		

qpid: list 1003

Object of type qpid.system: (last sample time: 20:49:53)

Type	Element	1003
config	sysId	d34bcff4-977f-4f7d-9d4e-2d6eda292486
config	osName	Linux
config	nodeName	localhost.localdomain
config	release	2.6.24.4-64.fc8
config	version	#1 SMP Sat Mar 29 09:15:49 EDT 2008
config	machine	x86_64

The screenshot shows the MRG Management interface in Mozilla Firefox. The browser address bar shows the URL: http://localhost:45672/index.html?frame=main.broker.queue;main.m=broker;main.broker.m=queue;main.broker.id=1;main.broker.queue.id=6. The page title is "Queue" and the URL is "http://localhost:45672/index.html?frame=main.broker.queue;main.m=broker;main.broker.m=queue;main.broker.id=1;main.broker.queue.id=6". The page content includes a "Queue" section with the name "mgmt-03ae8943-473f-4353-a2f0-c096d874a7b0". Below this, there are "Statistics" tabs for "Exchange Bindings (2)", "Details", "General", "Durability", and "Transactions". The "General" tab is selected, showing "Input/Output" statistics and "General" statistics. The "Input/Output" statistics table is as follows:

Statistic	Value	Per Second
Msgs. Enqueued	85	0.40
Msgs. Dequeued	83	0.30
Bytes Enqueued	12570	64.80
Bytes Dequeued	12326	50.60
Msgs. Unacked	0	0
Min. Msg. Latency	0	0
Max. Msg. Latency	0	0
Avg. Msg. Latency	0	0

The "General" statistics table is as follows:

Statistic	Value	Per Second
Consumers	1	0
Bindings	2	0
Message Depth	2	0.10
Byte Depth	244	14.20

There are three line graphs on the right side of the page: "Messages Enqueued and Dequeued", "Queue Depth", and "Consumers". Each graph has a time range selector (10 minutes, 1 hour, 1 day) and a legend. The "Messages Enqueued and Dequeued" graph shows two lines: "Msgs. Enqueued / sec" (red) and "Msgs. Dequeued / sec" (blue). The "Queue Depth" graph shows a single line for "Message Depth" (red). The "Consumers" graph shows a single line for "Consumers" (red). The status bar at the bottom of the browser shows "Done" and "<No Referer>".

How are we working together.

We work along side of the Condor team

- Always do what is best for the Condor project – community first!
- Have both on-campus resources and remote, and yes we are hiring

What areas have we working on so far with UW

- Build systems
- Libvirt (native Linux virtualization support – via gap)
- EC2 – more presentations on that
- Using hooks / Low Latency scheduling
- Management

Condor team also get access to make changes in Linux upstream through collaboration

- Expect more happening in this area in future.

Some questions answered

Will Condor on RHEL / Fedora / Linux have an advantage?

- Yes, tighter integration with the OS and working with both Red Hat and other upstream communities should provide benefits.

Is Red Hat going to provide technical support for Condor? At what level?
At what cost?

- Yes, www.redhat.com/mrg for more info

Is Condor going to be packaged as part of the Red Hat distribution?

- Yes, it is in Fedora and will be supported on RHEL 4.6, 5.1 or newer version

What is going to be the relationship with regards to the source code? Is there going to be a fork in the code, aka a UW Condor and a Red Hat Condor?

- Red Hat will co-develop with the Condor team. Work will be performed upstream, and all patches submitted back to UW.

Resources

Contact information

- Carl Trieloff, MRG - cctrielloff@redhat.com
- Matthew Farrellee, Grid team lead - mfarrellee@redhat.com

If you want to join the team (**yes I am hiring**)

- E-mail one of the above contacts.
- Full time and interns this summer! -- only apply if you are good.

Related Projects

- AMQP - www.amqp.org
- Qpid - incubator.apache.org/qpid
- Libvirt - www.libvirt.org
- Fedora - www.fedoraproject.org
- Red Hat MRG - www.redhat.com/mrg

