



HOW-TO

Scaling Magnum and Kubernetes: 2 million requests per second

The CERN OpenStack cloud team walks through how you can upgrade your Mitaka requests.



Superuser
July 29, 2016

We previously described in this [blog post](#) how we deployed [OpenStack Magnum](#) in the CERN cloud. It is available as a pre-production service and we're steadily moving towards full production mode as a standard part of the [CERN IT](#) service offerings to give [Containers-as-a-Service](#)



CONTAINERS-AS-A-SERVICE.

As part of this effort, we've started testing the upgrade procedures, the latest being to the final [Mitaka](#) release. If you're here to see some fancy load tests, keep reading below, but some interesting details on the upgrade:

- We build our own RPMs to include a few patches from post-Mitaka upstream (the most important being the trustee user to support lifecycle operations on the bays) and some CERN customizations (removal of neutron Load-Balancing-as-a-Service and Floating IP which we don't yet have, adding the CERN Certificate Authority, ...). [Check here](#) for the patches and build procedure
- We build our Fedora Atomic 23 image to get more recent versions of Docker and Kubernetes (1.10 and 1.2 respectively), plus support for an internal distributed filesystem called [Cern-VM-File System](#). We do use the upstream disk-imagebuilder procedure with a few additional elements [available here](#)

While discussing how we could further test the service, we thought of this [Kubernetes blog post](#), achieving 1 million requests per second against a service running on a Kubernetes cluster. We thought we could probably do the same. Requirements included:


- Kubernetes 1.2, which our recent upgrade offered
- available resources to deploy the cluster, and luckily we were installing a new batch of a few hundred physical nodes which could be used for a day or two

So along with the upgrade, Bertrand and Mathieu got to work to test this setup and we quickly got it up and running.

Quick summary of the setup:

- 1 Kubernetes bay
- 1 master node, 16 cores (not really needed but why not)
- 200 minions, 4 cores each

In total there are 800 cores, which matches the cluster used in the original test. How did our test go?



Breaking down the Austin Summit for OpenStack operators

Tom Fifield, community

April 23, 2016  Tom Fifield



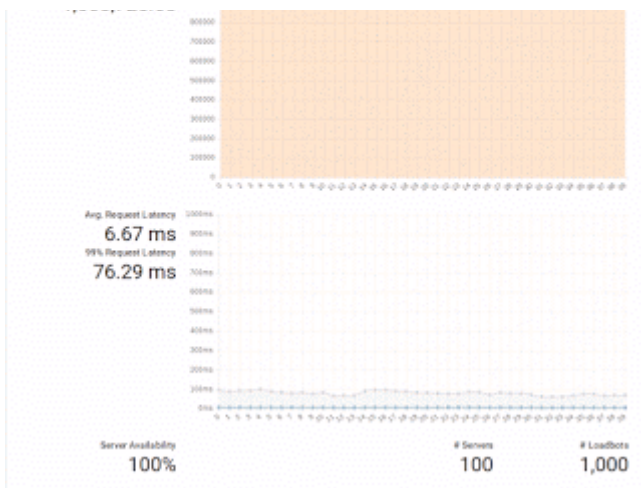
Superuser weekend reading

Keeping OpenStack from getting lost in translation, how Docker plus OpenStack equals cloud powerhouse and why it's

May 8, 2015  Superuser



Helpful Gerrit Queries (Gerrit




2.8 Edition)

A rundown of some handy Gerrit queries: Labels, Labels With Users, Labels

May 10, 2014  Sean Dague

We ended up trying a bit more and doubled the number to 2 million requests per second :)



User Group Spotlight: Greece

July marks OpenStack's 4th birthday and we're celebrating this year by spotlighting user groups from around the world.


July 9, 2014  Brittany Solano

We learned a few things on the way:

- set **Heat's** `max_resources_per_stack` to something big. Magnum stacks create a lot of these, and with bays of hundreds of nodes the value gets high enough that unlimited (-1) is tempting and we have it like that now. It leaves the option for people to deploy a stack with so many resources that Heat could break, so we'll investigate what the best value is
- while creating and deleting many large bays, Heat shows errors like 'TimeoutError: QueuePool limit of size ... overflow ... reached' which we've seen in the past for other OpenStack services. We'll contribute the patch to fix it upstream if not there yet
- latency values get high even before the 1 million barrier, we'll check further the demo code and our setup (using local disk, in this case a solid state drive instead of the default volume attachment in Magnum should help)



Myth-busting the cloud: Five common OpenStack questions answered



Superuser Awards: your vote counts

- Heat timeout and retrial configuration values need to be tuned to manage very large stacks. We're still not sure what are the best values, but will update the post once we have them
- Magnum shows 'Too many files opened' errors, we also have a fix to contribute for this one
- Nova, Cinder (bay nodes use a volume), Keystone and all other OpenStack services scaled beautifully, our cloud usually has a rate of ~150 VMs created and deleted per hour, here's the plot for the test period, we eventually tried bays up to 1000 nodes



Five years in: Charting the OpenStack galaxy



OpenStack User Survey Insights: November 2014



OpenStack + Kubernetes = More choice and flexibility for developers



Take these OpenStack Infrastructure tools and run!



OpenStack crowns "Liberty" its next release name



Getting my hands dirty on an OpenStack lab



Vote now for OpenStack Tokyo Summit presentations



Romancing the devs: another reason Walmart Labs went open source



Superuser weekend reading



What Do OpenStack Operators Do All Day?

And what's next?

- **Larger bays:** at the end of these tests we deployed a few bigger bays with 300, 500 and 1,000 nodes. And in just a couple weeks there will be a new batch of physical nodes arriving, so we plan to upgrade Heat to Mitaka and build on the recent upstream work (by Spyros together with Ton and Winnie from IBM) adding Magnum scenarios to Rally to run additional scale tests and see where it breaks
- **Bay lifecycle:** we stopped at launching a large number of requests in a bay, next we would like to perform bay operations (update of number of nodes, node replacement) and see which issues (if any) we find in Magnum
- **New features:** lots of upstream work going on, so we'll do regular Magnum upgrades (cinder support, improved bay monitoring, support for some additional internal systems at CERN)

And there's also Swarm and Mesos, we plan on testing those soon as well. And kubernetes [updated their test](#), so stay tuned...

Acknowledgements


- Bertrand Noel, Mathieu Velten and Spyros Trigazis from CERN IT, for the work upstream and integrating Magnum at CERN, and on getting these demos running
- [Rackspace](#) for their support within the [CERN Openlab](#) on running

containers at scale

- [Indigo Datacloud](#) building a platform as a service for e-science in Europe
- Kubernetes for an awesome tool and the nice demo
- All in the CERN OpenStack Cloud team, for a great service (especially Davide Michelino and Belmiro Moreira for all the work integrating Neutron at CERN)
- The upstream Magnum team, for building what is now looking like a great service, and we look forward for what's coming next (bay drivers, bare metal support, and much more)
- Tim, Arne and Jan for letting us use the new hardware for a few days



Tips for getting a travel grant to the next OpenStack Summit



How We Scaled OpenStack to Launch 168,000 Cloud Instances

This post first appeared on the [OpenStack in Production blog](#). Superuser is always interested in community content, email: editor@superuser.org.

Cover Photo // CC BY NC



FACEBOOK



TWITTER



LINKEDIN



GOOGLE+

Comments

Community



Login

Recommend

Share

Sort by Best

Start the discussion...

Be the first to comment.

ALSO ON SUPERUSER

[Getting Hyper-V and OpenStack set up quickly](#)

1 comment • a year ago



Roger Nelan — Hello, assuming that the process

[Why large data centers and telcos are fast-tracking](#)

1 comment • 6 months ago



David H Deans — Very insightful. Now. the difficult

...assuming that the process
above is complete, how

...integration with the current
part will be the execution of

Fighting cloud vendor lock-in, manga style

1 comment • 9 months ago •



Sergiu Strat — this is pure awesomess 😊

And the Superuser Award goes to...

1 comment • 3 months ago •



David H Deans — The AT&T digital transformation, Domain

[Subscribe](#) [Add Disqus to your site](#) [Add Disqus](#) [Add](#)

SUPERUSER

- [About](#)
- [Awards](#)
- [Contact us](#)
- [Subscribe](#)
- [User Survey](#)

COMMUNITY

- [User Groups](#)
- [Events](#)
- [Jobs](#)
- [Contribute](#)

USER RESOURCES

- [OpenStack Marketplace](#)
- [OpenStack Training Courses](#)
- [OpenStack Documentation](#)
- [OpenStack Operations Guide](#)
- [Application Developer Resources](#)
- [OpenStack Security Guide](#)
- [OpenStack Questions](#)



OpenStack Superuser is a publication built to chronicle the work of superusers, and their many accomplishments personally, professionally, and organizationally. The emphasis is on a blend of original journalism and user-generated content, ranging from technical to business-level issues with feature stories, case studies, tips and videos for OpenStack cloud architects and administrators.