



## INTRODUCTION

OpenStack® is young, but it's come a long way. Having been formed in a collaboration between Rackspace and NASA in 2010, it's now the de facto standard for enterprise private clouds.

However, that's not to say that implementing OpenStack is easy. If you ask OpenStack operators about their day-to-day activities, you'll quickly hear stories about persistent struggles to manage, monitor and maintain their clouds. Those stories cause concern for organizations like yours that are considering OpenStack for private or hybrid clouds.

One of the largest challenges is that the mental models, runbooks and tooling many organizations have used in traditional virtualization are inadequate for operating a cloud IaaS platform. Many considering OpenStack fear the limits of their mental models, runbooks and tooling will result in OpenStack becoming just another thing to wake them up at 2 a.m. on Sunday morning. Nobody wants that outcome, and we're here to help. Just as application architectures evolved for the cloud era, at Rackspace, we have developed a new model for operating OpenStack you can leverage for your own benefit.

## EXPLORING CLOUD FLEET MANAGEMENT MENTAL MODELS

Shortly after we created OpenStack, we deployed it in our public cloud. That fateful day was also the beginning of our evolution to Cloud Fleet Management. Today, our public cloud has grown to tens of thousands of servers across six different regions, with hundreds of thousands of VMs running at any given second. Over the years, as we searched for adequate tooling to manage a cloud of this size, we quickly realized it wasn't just the tooling which missed the mark, it was also the limitations of how we thought about the cloud – our mental model for management – that didn't align well with new capabilities.

The shift from traditional to cloud infrastructure introduces us to the concept of resiliency. In traditional environments, we balance slow provisioning with resiliency across infrastructure and vertically scaled applications. In cloud environments, we balance fast provisioning with resiliency in the horizontal scale of applications. OpenStack inherently accommodates the speed of provisioning, but it is the shift of resiliency to the application layer which necessitates a new Cloud Fleet Management mental model.

This mental model includes some interesting angles that must be evaluated:

- If applications are built for horizontal scale, what is the impact of moving a virtual machine from one host to another?
- If applications are built to be resilient to infrastructure failure, how painful is a host reboot to users?
- If you are operating at cloud scale, how do you scale your operations team?

In answering these questions, Rackspace came up with the following basic tenets of Cloud Fleet Management:

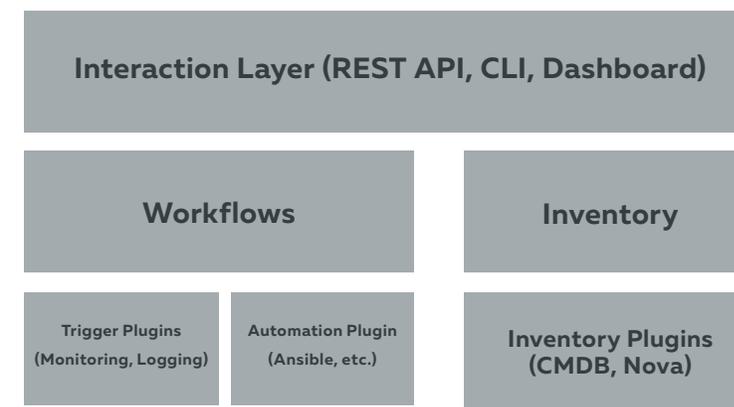
- Invest in the ability to migrate workloads between hosts
- Build tooling to determine if a host is showing any sign of error
- Automate everything for cloud scale

## EXPLORING CLOUD FLEET MANAGEMENT TOOLING

As our public cloud grows, Rackspace's manifestation of our Cloud Fleet Management mental model comes in the form of our tools. We started with a foundational service we call Inventory Management – think of it as a configuration management database (CMDB) for the cloud age. It provides an integrated layer to highlight the state of our cloud, with critical information such as network hierarchy, host status and the correlation between hosts and VMs.

We add to that an auditing and remediation workflow service that plugs into automation systems (e.g., Ansible, Chef, etc.) to collect information or respond to events, consult the inventory and respond appropriately. These auditing and remediation tasks are triggered by monitoring alerts or operator triggers and make up an intelligent runbook – a repository of known symptoms and prescriptive remediation that is constantly evolving. We package this all together in an operator UI with a dedicated central dashboard designed entirely for those tasks that require some degree of human interaction.

## FLEET MANAGEMENT ARCHITECTURE DIAGRAM



This new tooling allows us to manage the day-to-day operations of a multi-region cloud consisting of nearly tens of thousands of compute hosts with only a small team. On average, there are almost 900 hosts per member of the operations team. On any given day, our monitoring system could notice a known issue with a RabbitMQ queue, then trigger an automated remediation to gracefully shut down, drain, reset and return the queue to service, without a Rackspace operator ever being involved.

Cloud Fleet Management might notice a host with an error log known to be a precursor to host failure. While auditing the host, it notices that the host is not at the appropriate OS patch level. The system would then pull the host from Nova's compute pool, migrate all the VMs off the host, reboot the host at the appropriate OS patch level, and return the host to Nova's compute pool. These are just some of the many examples where, via Cloud Fleet Management, Rackspace operates OpenStack at massive cloud scale.

## CONCLUSION

We believe this same mental model and tooling would be valuable to all operators of OpenStack. That's why Rackspace is developing an open-sourced version of Cloud Fleet Management with project Craton. Think of it as Rackspace Cloud Fleet Management for organizations like yours. Craton is a rewrite of our internal tooling developed in the open with other operators who recognize a need for Cloud Fleet Management at scale.

This rewrite delivers a modular design that can connect with a variety of inventory, monitoring, logging and workflow services. Beyond just tooling, we hope this community of operators will evolve to include the sharing of known faults and remediation, giving new OpenStack deployments access to the shared experience of the global OpenStack community. Imagine the possibilities – you’re deploying your first OpenStack cloud without the knowledge that comes from operating OpenStack for years. Ideally, we hope to get to the point where you’ll have access to a living runbook of all previous OpenStack operators, providing you with a powerful tool befitting of the open-sourced community.

### **Conclusion**

Rackspace has considerable experience operating OpenStack. We have run the world’s largest OpenStack public cloud since 2012 and have deployed more OpenStack private clouds than anyone else. With Cloud Fleet Management, we’re driving additional innovation aimed at making your OpenStack experience more positive and efficient, while supporting our goal to deliver the best OpenStack experience in the world. To uncover what we can do for you to improve and enhance your OpenStack cloud capabilities, sign up for our free OpenStack strategy session at <http://go.rackspace.com/OpenStackExperts>.

## ABOUT RACKSPACE

Rackspace (NYSE: RAX), the #1 managed cloud company, helps businesses tap the power of cloud computing without the complexity and cost of managing it on their own. Rackspace engineers deliver specialized expertise, easy-to-use tools, and Fanatical Support® for leading technologies developed by AWS, Google, Microsoft, OpenStack, VMware and others. The company serves customers in 120 countries, including more than half of the FORTUNE 100. Rackspace was named a leader in the 2015 Gartner Magic Quadrant for Cloud-Enabled Managed Hosting, and has been honored by Fortune, Forbes, and others as one of the best companies to work for.

Learn more at [www.rackspace.com](http://www.rackspace.com) or call us at **1-800-961-2888**.

© 2016 Rackspace US, Inc.

This whitepaper is provided "AS IS" and is a general introduction to the service described. You should not rely solely on this whitepaper to decide whether to purchase the service. Features, benefits and/or pricing presented depend on system configuration and are subject to change without notice. Rackspace disclaims any representation, express or implied warranties, including any implied warranty of merchantability, fitness for a particular purpose, and non-infringement, or other commitment regarding its services except for those expressly stated in a Rackspace services agreement. This document is a general guide and is not legal advice, or an instruction manual. Your implementation of the measures described may not result in your compliance with law or other standard. This document may include examples of solutions that include non-Rackspace products or services. Except as expressly stated in its services agreements, Rackspace does not support, and disclaims all legal responsibility for, third party products and services. Unless otherwise agreed in a Rackspace service agreement, you must work directly with third parties to obtain their products and services and related support under separate legal terms between you and the third party.

Rackspace cannot guarantee the accuracy of any information presented after the date of publication.

Rackspace®, Fanatical Support® and other Rackspace marks are service marks or registered services of Rackspace US, Inc. and are registered in the United States and other countries. Other Rackspace or third party trademarks, service marks, images, products and brands remain the sole property of their respective holders and do not imply endorsement or sponsorship.

