

INTRODUCTION

Large IT organizations, telcos and service providers building OpenStack® private clouds have very specific needs. Once organizations resolve issues of stability, scalability and reliability, resource accountability becomes a chief concern.

The basic problem is that different business units and projects are all consuming the same compute, network and storage resources, but those business units and projects aren't accounted for individually — the cost of maintaining the environment is considered in aggregate. As a consequence, you might feel compelled to control costs by limiting resource consumption on the very cloud that was sold to business units as, "high-capacity, fast-provisioning and just like AWS."

A healthier alternative is to extend the costs of infrastructure directly to consumers through a showback system. This allows reporting, spend justification and business models to be based on the infrastructure required to power the relevant applications; otherwise, cloud utilization will be analogous to a black-box cost center.

Recently, Rackspace worked with one of the largest companies in the world to explore the best ways to deliver full-fledged resource accountability within OpenStack. This paper is a summary of our efforts.

EXPLORING THE CHALLENGES

To date, OpenStack's perspective on showback is a bit underdeveloped, with the majority of showback providers being third-party applications. This is a good thing, with the caveat that the appliances mining data are only as good as the inputs they can fetch from OpenStack. And out of the box, OpenStack does not ship the granular level of detail necessary to price cloud resources. Several projects exist to address this limitation. For your purposes, Telemetry is likely to be the most useful.

OpenStack's Telemetry project provides a data source and hooks to capture events and utilization at a high level of detail. Hundreds of measurements can be enabled, down to the number of times an image is launched and the network bandwidth used by a particular Swift container.

While typically used for monitoring and alerting on services, Telemetry's time series records are the perfect base for a showback system. There are some

downsides, primarily the cost of storing data for long periods of time. Even on a smaller cloud, Telemetry quickly reaches millions of records — so query and access time slows to a crawl. Thus the data needs to be routinely mined and aggregated into rollup reports or archived elsewhere, as is quite common in the metering world.

But meter data is only part of the billing puzzle. The other part is resource cost, also known as rating. To address this need, you should consider a Rating-as-a-Service project called CloudKitty. It will eventually be able to store and retrieve meters from Ceilometer, apply pricing algorithms, and display that billing information in a variety of formats.

However, it is still not entirely feature-ready, and its reliance on Telemetry for a data backend means that the data problems previously mentioned need to be addressed prior to production readiness. While there are plans to address the limitations, the fixes aren't likely to be immediate. Thus any showback system for use today needs to deliver:

1. A means to roll up or archive telemetry data
2. Easy access for third-party tooling
3. An approach that's small and future-proof — the goal is to minimize platform changing when Telemetry and CloudKitty come to maturity and are fully integrated into OpenStack

THE RACKSPACE ANSWER

Rackspace has put together a lightweight project to fill this niche, and the rest of this article will cover its design, implementation and best practices for approaching showback in a private cloud or provider cloud. Our aim is to help you overcome some of the limitations in OpenStack without added complexity or cost.

The OpenStack usage-summary tool (<https://github.com/rackerlabs/openstack-usage-report>) provides three small command-line tools that generate CSV billing reports and calculate licensing costs, and a brief reporting tool to break up costs by tag. This tool suite is not intended to be a heavyweight reporting system — another specialist third-party showback provider is almost certain to offer that. Thus the summary and licensing reporting are minimal regarding output and are mainly shown as a way of detailing best practices for ensuring showback readiness.

The useful component here creates an EC2-compatible billing feed that other applications are already comfortable parsing. This lets us bring the rich cloud billing ecosystem to OpenStack and unlocks more metrics for highly customizable billing.

The main way to generate the usage report is through the usage report CLI tool. We chose to implement the billing data in a CSV billing report that is very similar to the EC2 billing feed because it is so frequently adopted with third-party billing data visualization companies. The file format is customizable and can easily be tailored to individual needs. The fields generated in the report are all defined in YAML. Any of the meters in Telemetry can be plugged into this file with rates applied. Here's an example for applying \$.20 per vCPU hour:

```
- meter_name: vcpus

  billing_entity: OpenStack

  line_item_type: Usage

  product_code: Compute

  usage_type: Vcpu Hours

  operation: Runinstances

  currency_code: USD

  item_rate: 0.2

  description: cpu Hours

  product_name: Compute
```

The usage-report tool has a variety of date ranges it can pull from:

- Month to date
- Today (midnight to now)
- Last hour from now
- Any two date ranges

The readme file on the GitHub link contains more information on setting it up and generating reports.

RECOMMENDATIONS AND SUGGESTIONS

The data in the CSV file is an aggregate of all resource usage during that time frame, but it is not broken down by the user. There are quite a few ways to make OpenStack mirror an organization, but our recommendation is that you divide business units by keystone domain. Keystone domains are abstract separations of users and projects. Thus you can have a marketing domain with its own users and projects that remains separate from an internal IT domain. Each application's resources – including instances, volumes, images and Swift containers – should be tagged in OpenStack with an exclusive application identifier so that it's possible to see the exact cost of an application to a particular business unit.

The usage summary CLI tool can provide a simple overview of a cloud's usage via the usage CSV file generated in the previous step. This tool turns a CSV billing file into a chart of spending broken down by keystone domain (business unit) and application ID. Thus every business unit can get a list of every application running under its domain, every resource utilized by an application, and the total cost of every application.

Clever users may realize that OpenStack will let them create any resource without a particular tag by default, enabling them to fly below the radar of cost oversight. We recommend running search and destroy tasks that find cloud resources, like instances and snapshots created (that are not assigned to an application ID), and alerting someone to their existence or deleting them. Additionally, it may make sense to implement a process for obtaining a valid application ID from an approver so that all applications can be overseen by the business unit funding them.

Software licensing is yet another matter to consider. OpenStack has a component known as the Glance metadata catalog. This is a tool that can be used to specify which software is embedded within a Glance image. When images are spun up, Telemetry captures the initial metadata from the image and records it once. The tags are not visible on the running instance, but can be queried from Telemetry data. Thus, software that is registered on an instance is immutable and cannot be removed unless the instance is destroyed, so you can't remove the license reporting from running SQL server instances.

The usage licensing CLI tool follows the same model as the usage-report tool. There is a YAML definition file that specifies licensing costs. Currently, you can choose to charge by the hour or by the total installation count, though additional

licensing models can easily be implemented as modules to the tool. Take a look at this example of a definition file, and you'll see how to reference custom modules in the YAML config:

columns:

- name: Resource Id
func: resource_id
- name: Project Id
func: project_id
- name: Meter Name
func: meter_name
- name: user:environment
func: metadata:environment
- name: image:SQLServer Edition
func: image_metadata:sw_database_sqlserver_edition
- name: image:SQLServer Version
func: image_metadata:sw_database_sqlserver_version

Now in the same file, you can reference the name you set with the price you want. For the SQL server example noted above, the following definition will calculate the total number of SQL server installations during the usage window and apply the price per running instance of \$15 for the enterprise edition and \$8 for the standard edition:

- type: SQLServerCount
project_id_field: Project Id
- costs:

enterprise:

'2016': 15

standard:

'2016': 8

CONCLUSION

This tool and the prototyped summary/license tools show how one of our customers leveraged our innovations to approach showback within their OpenStack private cloud. Showback will continue to evolve in the OpenStack ecosystem as projects mature and vendors provide more advanced functionality. We are looking toward implementing the CSV output as a CloudKitty writer to produce EC2-compatible billing feeds that any showback software provider can read and visualize. For real deployments, an application vendor specializing in showback, like Taligent, can be brought in for deep reporting and cost analysis while implementing these best practices.

Rackspace strives to deliver the best OpenStack experience in the world, with industry-leading reliability, unmatched scalability, innovation-driven agility and a superior overall approach – all of which lead to better outcomes for your organization. To learn more about what we can do to enhance and simplify your OpenStack cloud, 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 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.

