EMBOTICS®

Cloud Manager's Guide

# Taking Control of Public Cloud Consumption

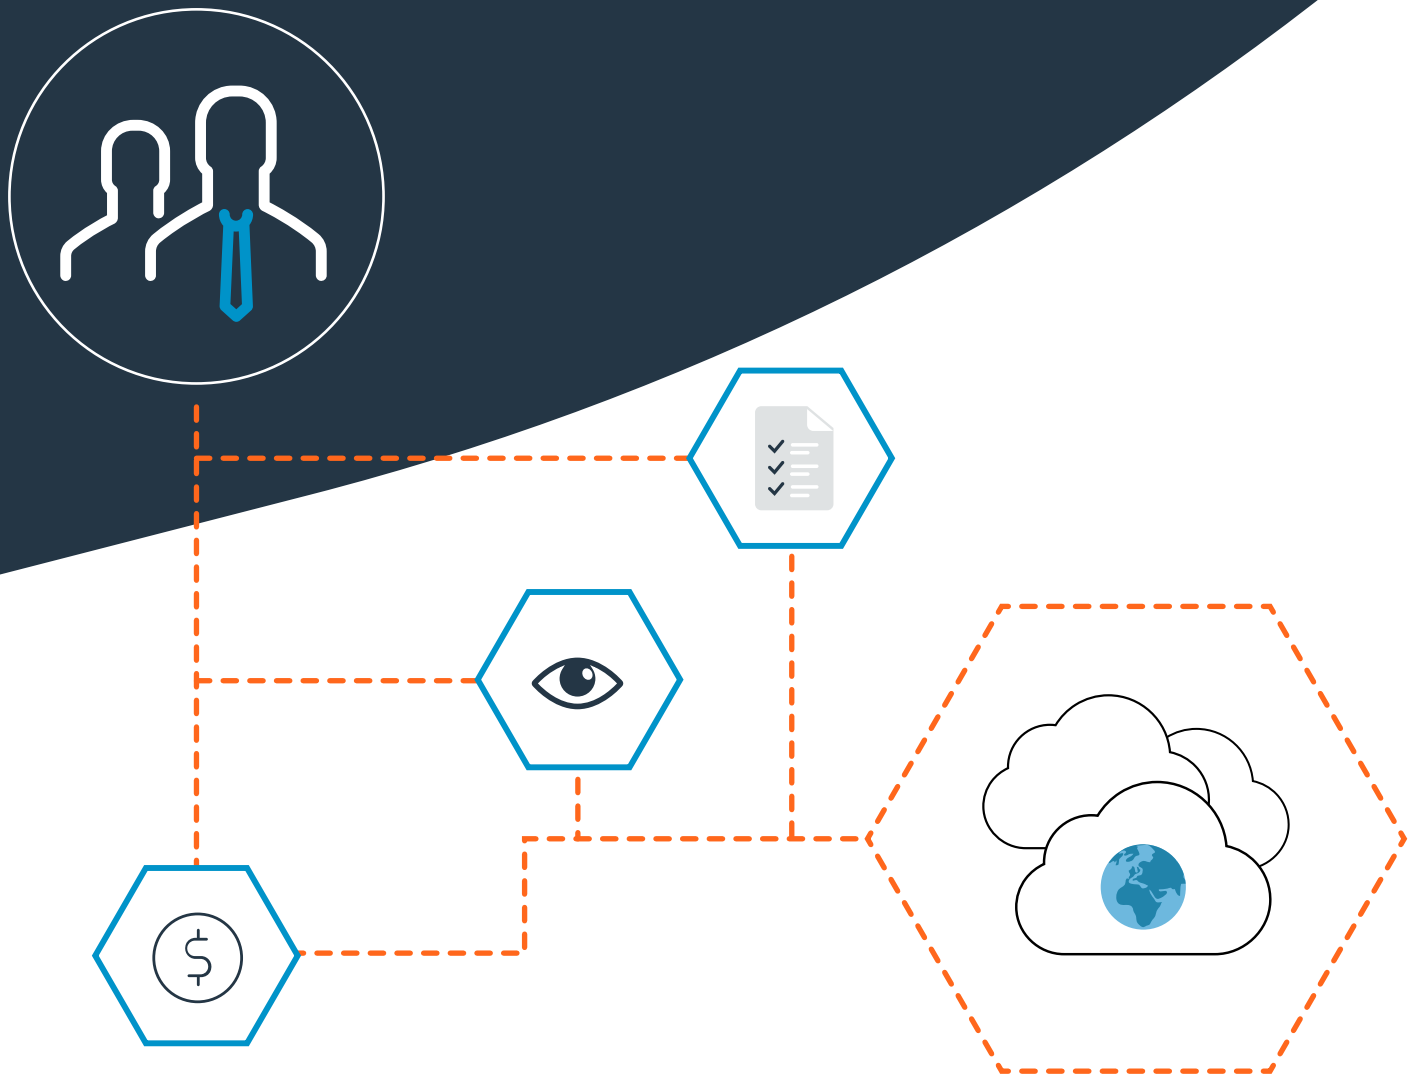

# Table of Contents

## Ghost-wrestling

Getting control of public cloud consumption can be like wrestling a ghost: it's hard to pin down something nebulous.

But if you're in charge of reducing costs without hampering innovation, you've got to pin this ghost down.

To wrestle public cloud consumption into submission, you need visibility, standardization and cost reduction.

A modern cloud management platform (CMP) provides the cloud governance tools you need to control public cloud consumption.

How you get control depends on how your organization consumes cloud resources right now.

There are two main consumption models:

**Cloud Brokered:** Cloud resources are governed by the IT department and consumed through a portal. One method for this model is self-service provisioning.

**Cloud Direct:** Users consume cloud resources directly through APIs, cloud-native tools and public cloud consoles.

A CMP provides automation and governance for both Cloud Brokered and Cloud Direct environments using AWS, Azure and Google Cloud, as well as Kubernetes running in any of these clouds or on-premise.

This eBook will take a look at the tools a CMP gives you for gaining visibility, implementing standards and reducing costs, while making your team more efficient and productive.

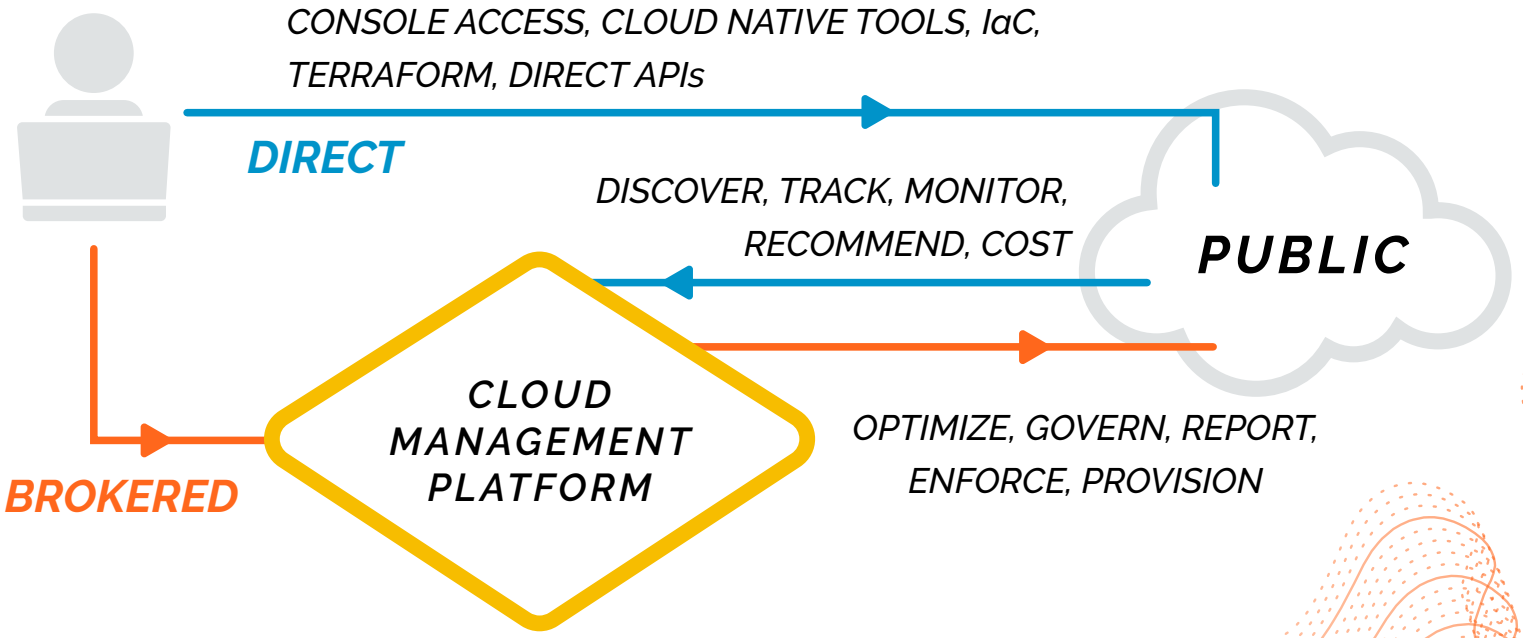

# Gain visibility of public cloud usage and costs

The first step in getting control of your public cloud consumption is to see it clearly.

## Visibility of IT assets

A CMP can provide a unified view of your public cloud infrastructure. Cloud operators can achieve this type of visibility through cloud management architecture designed to be public cloud independent.

Gaining greater visibility into your public cloud environment gives your IT department critical insight into the value and efficiency of your cloud assets. And combined with private cloud visibility, you can achieve a complete hybrid cloud picture.

## Cost visibility

Integrated cloud expense management capabilities provide public cloud cost visibility, cloud cost comparisons and actionable savings recommendations.

Using built-in cost calculation and simple-to-use wizard-based cost models that you can customize for both private and public infrastructure, you can present IT chargeback or showback information to your users.

With a CMP that supports multi-tenant cost configuration, you can present costs for different users, customers or parts of your infrastructure. You can achieve visibility into all cloud costs, run what-if scenarios that compare workload costs on different parts of your infrastructure, and provide out-of-the-box billing reports for the consumers of your IT services.

Public cloud cost visibility allows your IT department to demonstrate the true value of the services they provide and makes business units accountable for their cloud consumption.

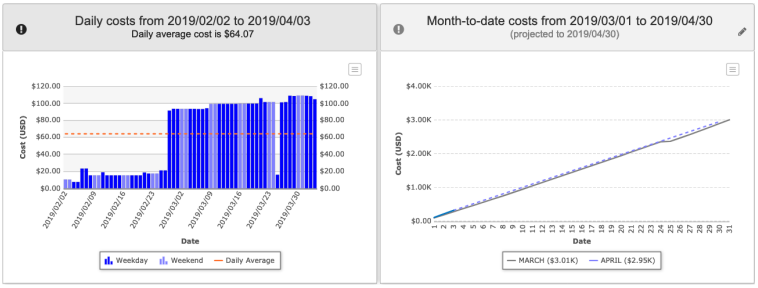


Daily Costs & Projected Monthly Costs –
From Embotics CMP Cost Analytics Dashboard

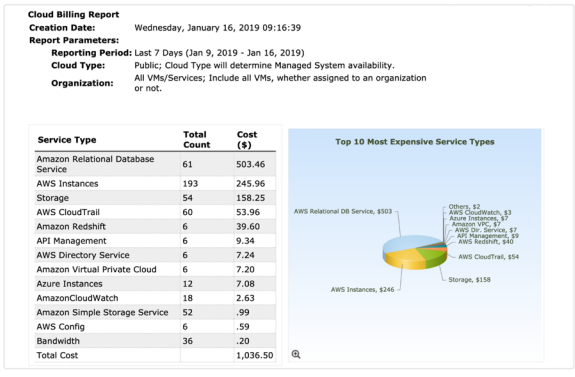

**Cloud Billing Report**
**Creation Date:** Wednesday, January 16, 2019 09:16:39
**Report Parameters:**
**Reporting Period:** Last 7 Days (Jan 9, 2019 - Jan 16, 2019)
**Cloud Type:** Public; Cloud Type will determine Managed System availability.
**Organization:** All VMs/Services; Include all VMs, whether assigned to an organization or not.

| Service Type | Total Count | Cost ($) |
|---|---|---|
| Amazon Relational Database Service | 61 | 503.46 |
| AWS Instances | 193 | 245.96 |
| Storage | 54 | 158.25 |
| AWS CloudTrail | 60 | 53.96 |
| Amazon Redshift | 6 | 39.60 |
| API Management | 6 | 9.34 |
| AWS Directory Service | 6 | 7.24 |
| Amazon Virtual Private Cloud | 6 | 7.20 |
| Azure Instances | 12 | 7.08 |
| AmazonCloudWatch | 18 | 2.63 |
| Amazon Simple Storage Service | 52 | .99 |
| AWS Config | 6 | .59 |
| Bandwidth | 36 | .20 |
| Total Cost | | 1,036.50 |



Get cost visibility through the use
of a cloud billing report

# Standardize your public cloud environment

Now that you can see the ghost clearly, introducing standards allows you to grab hold of it.

Without standards, you're constantly reacting to changes in your environment.

You need to control the change without impeding the pace of development – in fact, a set of solid standards can enhance productivity.

In the Cloud Brokered model, you can enforce governance up-front, so that resources are automatically compliant when provisioned.

In the Cloud Direct model, enforcement of standards must be reactive but can still be quite effective.

## Cloud tagging

Through cloud tag synchronization and a tag Compliance Policy, you can certify that all your instances follow best practices. You can make policy enforcement as strict or as lenient as you like.

In the Cloud Direct model, you can report on those instances not following your corporate tagging rules. With Cloud Brokered consumption, you can be more restrictive – you can ensure that instances are always powered off when not in use, and you can even customize workflows to bring offending instances into compliance.

Ensuring that all members of your organization, from developers to IT, follow corporate standards is a critical part of proper cloud governance.

## Tenant ownership

Public clouds use multi-tenancy to share IT resources securely among multiple applications and tenants (such as businesses units) that use the cloud.

Using CMP technology, you can assign tenant ownership via policy automation if end users are consuming cloud resources directly.

If your organization consumes resources brokered through a CMP, tenants can assign ownership directly through self-service provisioning.

Automatic ownership assignment not only provides enhanced security and easier overall cloud management, but also enables greater agility and scale, allowing for new capabilities to be rolled out by the business units as they require, but in a properly controlled fashion.

## Intelligent placement

With intelligent placement technology, you can ensure that application workloads are automatically deployed to optimal locations, based on criteria you define.

IT departments tap into intelligent placement capabilities through a multi-cloud service catalog, which considers cost, workload business attributes and per-destination/tenant quota.

These features are all part of Embotics' out-of-the-box policy automation capabilities.

## Limited set of instance types

Controlling the instance types your developers can use keeps the number of instance types down to a manageable minimum, allowing you to gain economies of scale and take advantage of public cloud bulk usage discounts for Reserved Instances (RIs).

When using a CMP for self-service provisioning automation, you can control the instance types your developers can use. Specific instance types can be configured for each service catalog entry, providing extensive flexibility across your various service offerings while still enforcing instance type consistency and governance.

## Service catalog and provisioning automation

A modern CMP's end-user service catalog can include single- and multi-tier apps, multi-cloud services, ARM templates and AWS CloudFormation templates.

You can use the built-in orchestration capabilities to automate the approval process, provisioning and decommissioning. Automation reduces provisioning time, manages the lifecycle of virtual assets and eliminates sprawl.

A service catalog allows users to deploy workloads without administrative intervention, streamlining the interface between IT and end users. Users get access to cloud resources more quickly, and your organization benefits from improved productivity, decreased administrative workloads and reduced costs.

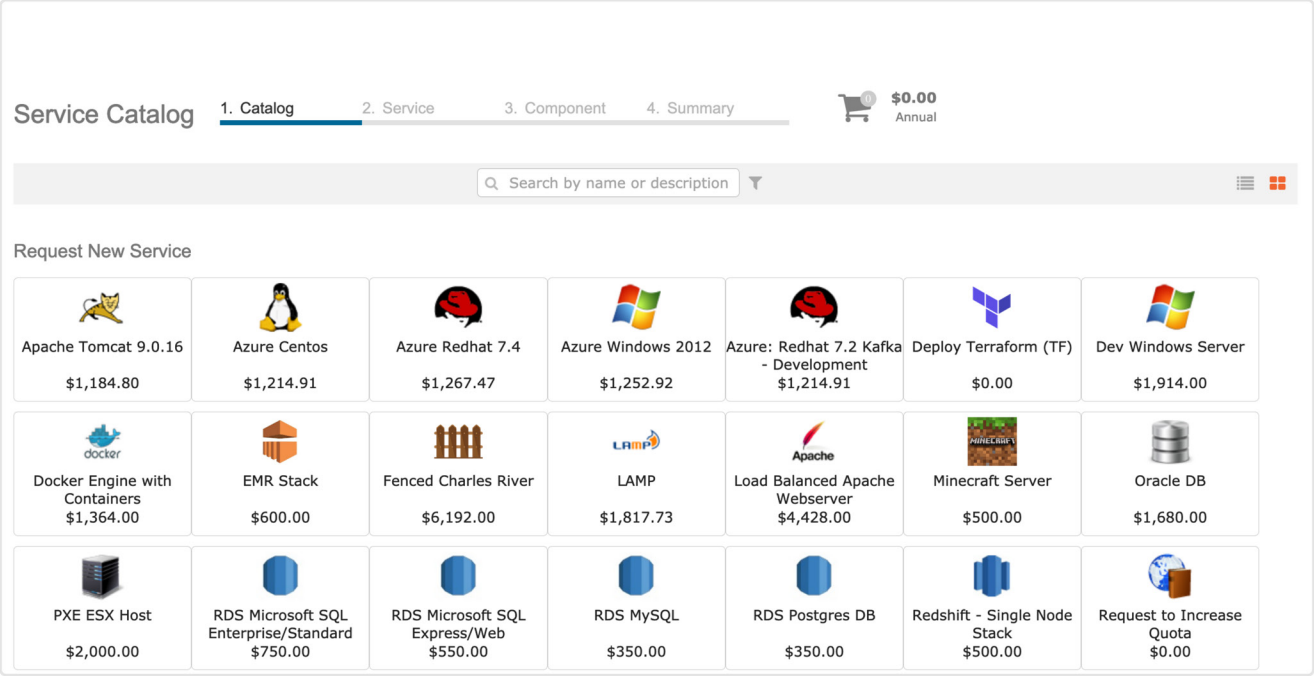


Self-Service users select from a catalog of standardized applications

## Automated decommissioning

Many public cloud application workloads are deployed by engineering teams for short-term use, but then languish and continue to run, in some cases indefinitely, because there's no lifecycle monitoring or enforcement.

A CMP's policy engine automates application lifecycle management, ensuring that services are decommissioned when no longer needed.

Flexible notification and expiry extension options ensure that users are well informed and can keep resources active for as long as they need to, without introducing sprawl.

With cloud management, you can fully automate decommissioning, targeting only public cloud infrastructure specific to R&D (such as accounts, regions or VPCs).

Automated decommissioning allows engineering groups to spin up instances exactly as they do today, without impeding engineering velocity.

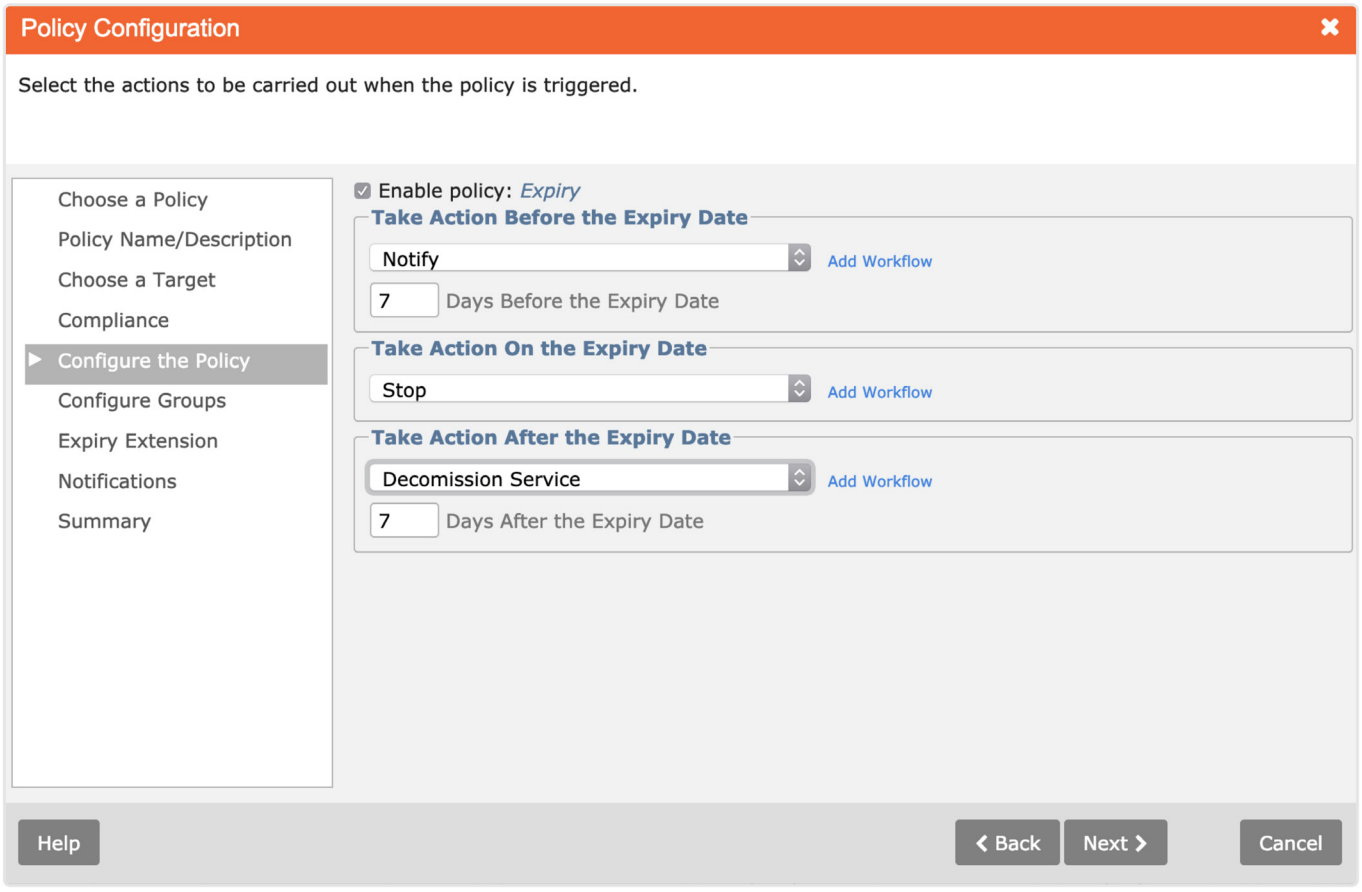


Automatically decommission expired workloads

# Reduce public cloud costs

Once you've got a clear view of your public cloud usage and costs, and you've gained some control through standards, it's time to apply cost-saving measures.

## Instance rightsizing

Rightsizing ensures that all instances in your public cloud environment are allocated the correct resources for their workload.

With a CMP, you can predict and identify constraining resources and provide rightsizing recommendations to both administrators and end users.

In cloud direct environments, rightsizing recommendations can be used by the IT team for reporting, while in the cloud brokered scenario, workflow orchestration can automatically rightsize in a fully governed fashion.

Rightsizing recommendations allow end users to decide when an instance is over-provisioned as well as prevent capacity and performance bottlenecks.

With access to performance metrics and cost data, rightsizing helps your users work with you—not against you—creating a happier, more efficient IT environment for all. And since the majority of instances are over-provisioned, significant cost savings is typically the result.

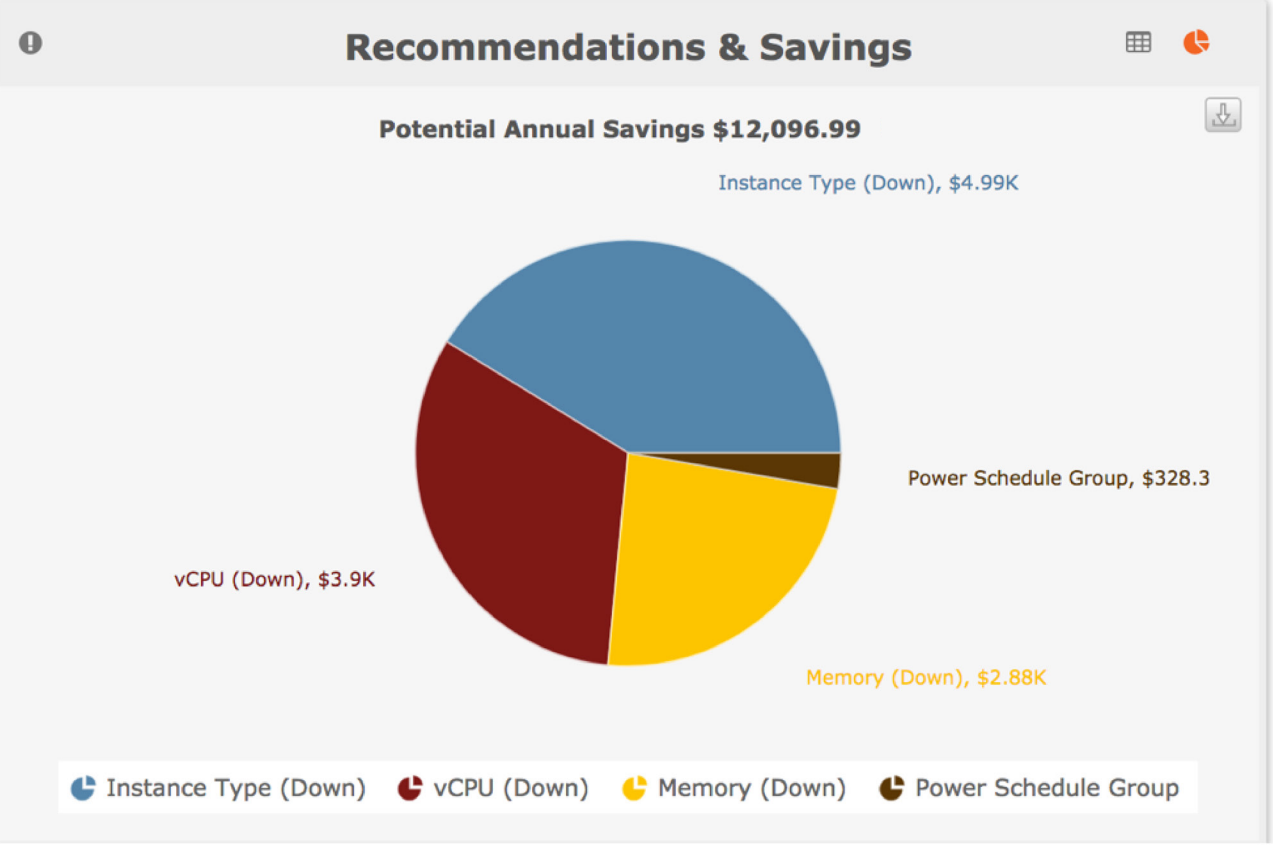


Rightsizing recommendations can cut costs and boost workload efficiency

## Power scheduling

Many types of dev/test workloads don't need to be running 24/7, but it's easy to forget to power them down.

And you pay by the hour for the running instances you host in a public cloud environment, so the less time they're running, the better.

Policy automation supports scheduled workload shutdown and restart. A CMP can also make power schedule recommendations based on tags (such as dev, test and QA).

Power scheduling provides control while allowing engineering teams to use public cloud resources as needed. You can save 60-70% of the workload cost through power scheduling.

## Reserved Instance bulk usage discounts

Bulk usage discounts can have a significant impact on your public cloud bill.

AWS Reserved Instance (RI) purchase recommendations allow you to reserve computing capacity in exchange for significantly discounted hourly rates (up to 75%) compared to on-demand instance pricing.

But determining how to optimize your RI spending can be a challenging task. This is where a Reserved Instance Planner can be crucial for generating the analytics you need and making recommendations on the best return for your RI purchases.

## Cloud and tenant quotas

Cloud and tenant quotas allow you to assign available resources to your consumer groups based on their business requirements.

You can set compute resource or cost limits for each provisioning destination or each tenant. A self-service quota view allows tenant groups and individual group members to police their own consumption. If end-users can provision to more than one location or cloud, quota information is also used as a factor in intelligent placement decisions.

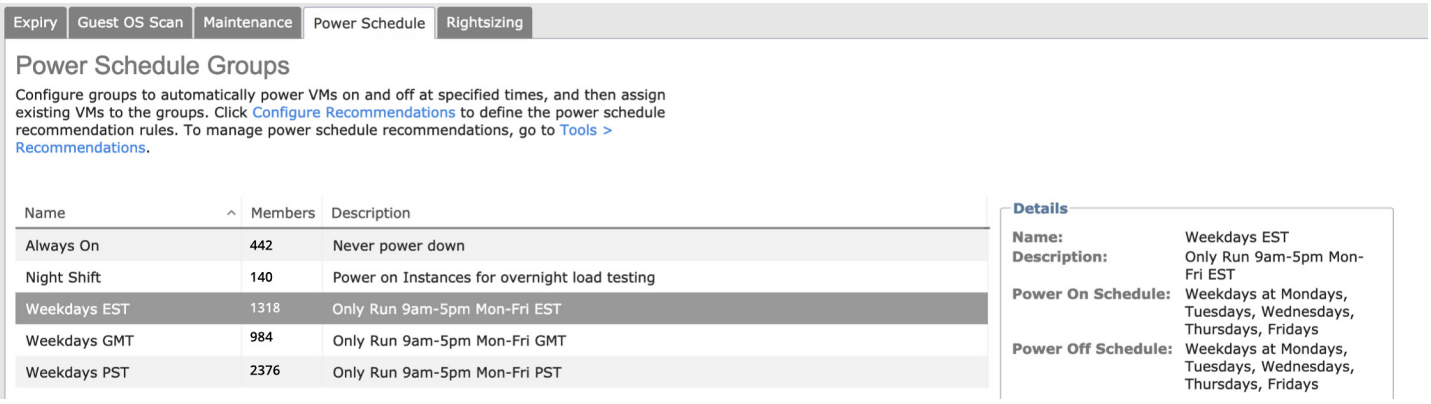


Configure power schedules based on workload type

## Rinse and repeat

Cost optimization is an ongoing, iterative process.

Once you've rightsized, set power schedules and taken advantage of bulk usage discounts, it's critical to monitor your cloud platform infrastructure and your bill, and continually optimize.

## Summary: Controlling your Public Cloud Environment

The following table summarizes vCommander's public cloud capabilities for organizations accessing and consuming cloud resources directly or brokered through a cloud management platform.

| Category | Cloud Direct | Cloud Brokered |
|---|---|---|
| **Visibility** | | |
| IT Visibility | Single pane of glass visibility over multi-cloud environments | |
| Cost Visibility | Showback/chargeback | |
| **Standards** | | |
| Cloud Tagging | Monitored | Enforced |
| Tenant Ownership | Policy Automation | Self-Service |
| Intelligent Placement | - | Policy Automation |
| Limited Instance Types | - | Enforced |
| Service Catalog / Provisioning Automation | - | Self-Service |
| **Cost Reduction** | | |
| Bulk Usage Discounts | Recommendations | |
| Power Scheduling | Policy Automation | |
| Automated Decommisioning | Policy Automation | |
| Instance Rightsizing | Reporting | Self-Service |
| Cloud & Tenant Quotas | - | Enforced |

EMBOTICS®

## About Embotics

Embotics Cloud Management Platform is built from the ground up on a common architecture to deliver an all-in-one solution that supports multiple hypervisors and multiple clouds. Its flagship product is used by organizations including Nordstrom, NASA, Informatica, HBO, NTT Data, and more.

For more information:

www.embotics.com

info@embotics.com

+1 877 599 0494