MindPoint GROUP℠

Elastic | Efficient | SECURE

Solution Template for Cloud Migration

Cloud Cost Tracking: Accounts

PRESENTED BY:

Keith Rhea
Technical Lead, Cloud Architecture and Security Engineering

MindPoint Group, LLC
1330 Braddock Place, Suite 600, Alexandria VA 22314
(o) 703.636.2033 | **(f)** 866.761.7457 | **www.mindpointgroup.com**

TABLE OF CONTENTS

INTRODUCTION

Any organization embarking on a cloud deployment needs to understand the importance of the foundational issues, account setup and resource tagging. These concepts will have significant impacts on a number of areas including resource inventory, security, and cost tracking and allocation. While resource inventory and security may be at the forefront of concerns at the outset, cost tracking and allocation typically is not. We started writing this paper based on real-world experience supporting customers attempting to mature their services deployed to cloud so it is meant to be targeted at providing base information that can help develop a successful cloud deployment plan.

Most people understand that there will be changes in this regard as they adopt cloud. The buying model for IT services inherent in cloud should allow for organizations to save costs (and that is often the siren song leading to cloud adoption), but it can also allow for poorly tracked service provisioning and usage to spiral out of control. Detailed cost tracking in a manner that is as close to real-time as possible is critical to keep costs under control so that IT can stick to budget and demonstrate cost-savings.

This paper is one in a three-part series that will cover the basics you need to know about account setup and resource tagging with the goal of being able to get to a position where you are able to track costs in a sufficient manner. Since we based the series largely on real experience, it will provide as much generally-applicable information on these topics as possible, but it does dive into more detail based on a few specific assumptions - mainly that AWS is the main cloud platform in use and that an existing Splunk installation can be used for developing dashboards and visualizations.

BACKGROUND

Security is often not overlooked in a cloud adoption. In fact, many people over the last several years have remained so skeptical that even when the topic is treated with ample care, skeptics remain unable to be convinced that cloud can often be as secure (or insecure) as their on-premise environments. However, while dealing with the foundational issues of engineering, operations, and security in the initial phases of a deployment, every organization will eventually have to address the issue of tracking costs. Cloud adoption (while not limited to the following examples), could take on any of the following forms: The organization is moving a small DEV workload to the cloud as a proof of concept;

- The organization is going “all-in” and migrating a majority of its infrastructure to cloud; or
- The organization maintains a multi-tenant platform delivering core services to customers using cloud, and therefore must be able to track costs incurred in a way that supports allocating those costs properly.

For any of these common patterns of adoption, there will usually be an end state need to be able to track resource usage and the associated costs incurred in a granular way. This is necessary because whether your IT organization recoups those costs from business units through a chargeback model, simply receives a budget allocation to pay for enterprise IT consumption, or even provides multi-tenant services to external customers as the core business function, understanding how, when, and why costs are incurred will be critical to success in the long term.

For organizations that charge customers for services that are delivered through infrastructure deployed to cloud there can be additional pressure to make sure that this cost tracking is in place. After all, in order to be able to charge the customers for their service usage, the organization must clearly understand how much it costs them to provide those services to each customer. Further, some services may be able to be directly tied and allocated to individual customers while others may be pooled resources that are shared among customers. These challenges make it vital to focus on implementing a good set of processes and powerful tooling to track those costs.

This whitepaper attempts to dig into some of the considerations regarding structuring and managing accounts as well as what impacts those options will have on security and the ability to aggregate cost data. Understanding these fundamentals on tracking resource usage and costs are important to help develop a plan that will lead to successful adoption and migration.

ACCOUNT STRATEGY

With any cloud service provider, there are a number of options for managing accounts. Often organizations embark on adoption in an ad hoc manner using a single account. However, there is nothing preventing the use of multiple accounts and typically it is possible to manage these in a hierarchical manner. Understanding the impacts of single versus multiple accounts is important to developing a plan that will support whatever goals for cost management or even cost recovery your organization might have.

SINGLE V. MULTIPLE ACCOUNTS

Many organizations initiate their cloud environments with a single account. Maybe it is a structured proof of concept or maybe there are a handful of IT administrators that were tasked with “planning” the organization’s cloud migration. No matter the path of initiating your first account, there are some important questions you need to answer before proceeding further. Retro-fitting a design strategy can be a daunting effort and can be easily avoided by properly planning your efforts up front.

In the context of AWS, as with other cloud service providers, accounts provide a line of demarcation. Regardless of the services used, the costs for those services will inherently be associated with the account they are provisioned under. Therefore, in the most basic sense this can be a key tool to help aggregate costs. An organization can create and manage a number of accounts. Since AWS will aggregate the data regarding services used and associated costs with each account ID, this provides a simple way to then provide reporting on costs by account ID.

An example of how this may be used is within the context of an enterprise IT organization providing support to multiple business units, it might make sense to create an account designated for Sales, an account for Manufacturing, and so on. Then it becomes straight forward to track and associate costs with these business units. For an organization building sellable services the usage of multiple accounts can be a method for maintaining separation of cost management between the infrastructures used for different products.

However, it is common for any organization attempting to break out of the cycle of siloed costs and an inability to gain efficiencies in IT spending to attempt to pool resources across customers internal or external. Using multiple accounts does not run counter to that approach, but can be used where it makes sense to draw clear lines of delineation for segregating costs in a discrete manner.

Table 1 provides a comparison of a single versus multiple account strategy.

	Single Account	Multiple Accounts
Pros	1. Limits the amount of management overhead 2. Only requires managing one set of users 3. Requires development of smaller number of security policies 4. All charges are visible on single bill	1. High level of security isolation both in terms of security/users and resources 2. Allocation of discrete costs can be done by account 3. Simplifies tracking resources
Cons	1. Can increase complexity of tagging structure to make up for lack of inherent account demarcations a. For security controls b. For resource usage tracking/billing 2. Can increase complexity of security policies, again, to compensate for account boundaries 3. Complex network policies to segregate applications and operating environments 4. Requires more complexity to implement isolation of users, resources, & data	1. Additional overhead to setup & manage accounts 2. Increases complexity of user and role management across accounts/environments

Table 1: Pros and Cons of Single Versus Multiple Accounts

While each approach has its place, it is almost always advisable to use multiple accounts in all but the simplest environments. So, the question is not so much yes or no, but how many are appropriate and how do you structure them?

SECURITY IMPACTS

Security can be a compelling reason to want to create more accounts. Even if you have a single group of users, multiple accounts may provide a straight forward way to implement security policies depending on the environment. For a single organization, there may be a development area where any number of people can provision and modify resources while the production area is more strictly controlled. By creating two different accounts in which to run these resources, it is simple to define the permission differences in security policies.

This is one of the core tradeoffs between a single account and multiple accounts- with a single account there are generally fewer things (users, security policies, etc) to deal with, but the way in which they must be managed becomes increasingly complex. At this point it is worth mentioning that any organization moving to cloud needs to make sure they have a mature configuration management process and tooling that supports automation. The reason is that if you can automate the process of implementing and enforcing security policies, then the perceived extra overhead burden can be prevented in reality. In fact, it will be more straight forward to create security policy templates and automate their implementation when they are less complex.

COST TRACKING IMPACTS

Outside of security, being able to effectively track billing and resource usage is the most important reason for choosing a multi-account strategy. As we will get into more in the next installment, tagging resources is a valuable tool to be able to track costs, and the same rules apply here in terms of complexity. Tagging will necessarily be more complex when there is a single (or fewer) accounts used. As a simple example, consider again an organization where there are customers that have development and production environments. In a strategy where you provision accounts for each environment that each customer has you would have the following as an example:

- Customer 1: DEV Account
- Customer 1: PROD Account
- Customer 2: DEV Account
- Customer 2: PROD Account

Without using tagging, we can tell which costs were incurred by Customer 2 simply by adding the costs incurred under the two accounts assigned to them. Using a single account, you would need to implement the following tags to achieve the same granularity of tracking costs:

- Customer
- Operating environment

Another complication is that while resources are tied to accounts with 100% assurance, tags are not. That is to say, when a resource is provisioned it is always done within the context of an account. There is never any ambiguity about which account owns a resource. Unless you put strong controls in place, you cannot say the same about tagging resources. By default, it is purely an optional exercise. Also, consider that there are some resources which do not support tagging.

CORNERSTONE ACCOUNTS

While we've hopefully provided some guidance to consider regarding setting up accounts, if you do plan to use multiple accounts, there are a few "cornerstone" accounts that almost every organization will want to implement as part of their overall plan. Those include:

- A master billing account;
- A shared services account; and
- A security account.

These accounts are designed to segment the core services, both general IT services and security-specific services, that the IT organization managing the entire cloud deployment will be providing to all customers as well as an account used to aggregate payment for all costs of linked accounts. This can include things like running a central user directory system; DNS; performance and service monitoring tooling; and project management or development tools for the “shared services account.” For the “security account” it is where the security monitoring services such as log aggregation; vulnerability management; SIEM; incident management; or other similar tools are run.

ACCOUNT SETUP

While using multiple accounts is often beneficial for a number of reasons, it does mean you are responsible for managing more accounts. Automation should be the answer to this problem. Whether you use Ansible, Chef, Puppet, Terraform, Cloud Formation, or a combination of the above, automation is the key to keeping all your accounts and resources uniform and secure. You want to ensure the process is repeatable, efficient, and secure to allow you to quickly provision access for new customers while ensuring that there is little to no manual processing so that provisioning is done consistently and reliably.

Once the process has been defined, it is essential to define who within your organization has privileges to provision new accounts and they understand the process. It is extremely easy for organizations to generate new accounts and drift from the process that has been established when it is implemented manually and there is little or no training regarding the process. This leads to lack of security, lack of visibility, and cost inefficiencies. If you put in all the effort to develop the framework, you want to make sure everyone in your organization knows what it is and how to use it. More than that, you want to make sure that automation reduces the chance for errors.

With the process of creating, connecting, and securing new accounts completed, you now need to develop a strategy of how the “customer” accounts will be organized. There are a few pre-requisite questions you should be asking yourself:

- Who are your customers? Are they internal, or are you building a shared service to sell as a core function of your business?
- What is the current footprint of my IT operating environment?
- What business unit segregation strategy best fits my organization?
- What factors should trigger the need for a new account?
- Do business units or external customers require segregation of workloads?
- Is it possible to project how many accounts are needed based on these initial questions, and if so are there any limits in place with the cloud service provider that would conflict with this?
- How is security managed currently? Is it centralized or are there distributed teams?

Figure 1 provides an example of how accounts might be structured at an IT organization providing services to internal business units that has created accounts for the following:

- A consolidated billing account;
- A shared services account;

- An account for a centralized security team; and
- Individual accounts for DEV, STAGE, and PROD environments for each customer.

Each of the customer accounts is integrated with the resources deployed in the shared services and security accounts as needed.

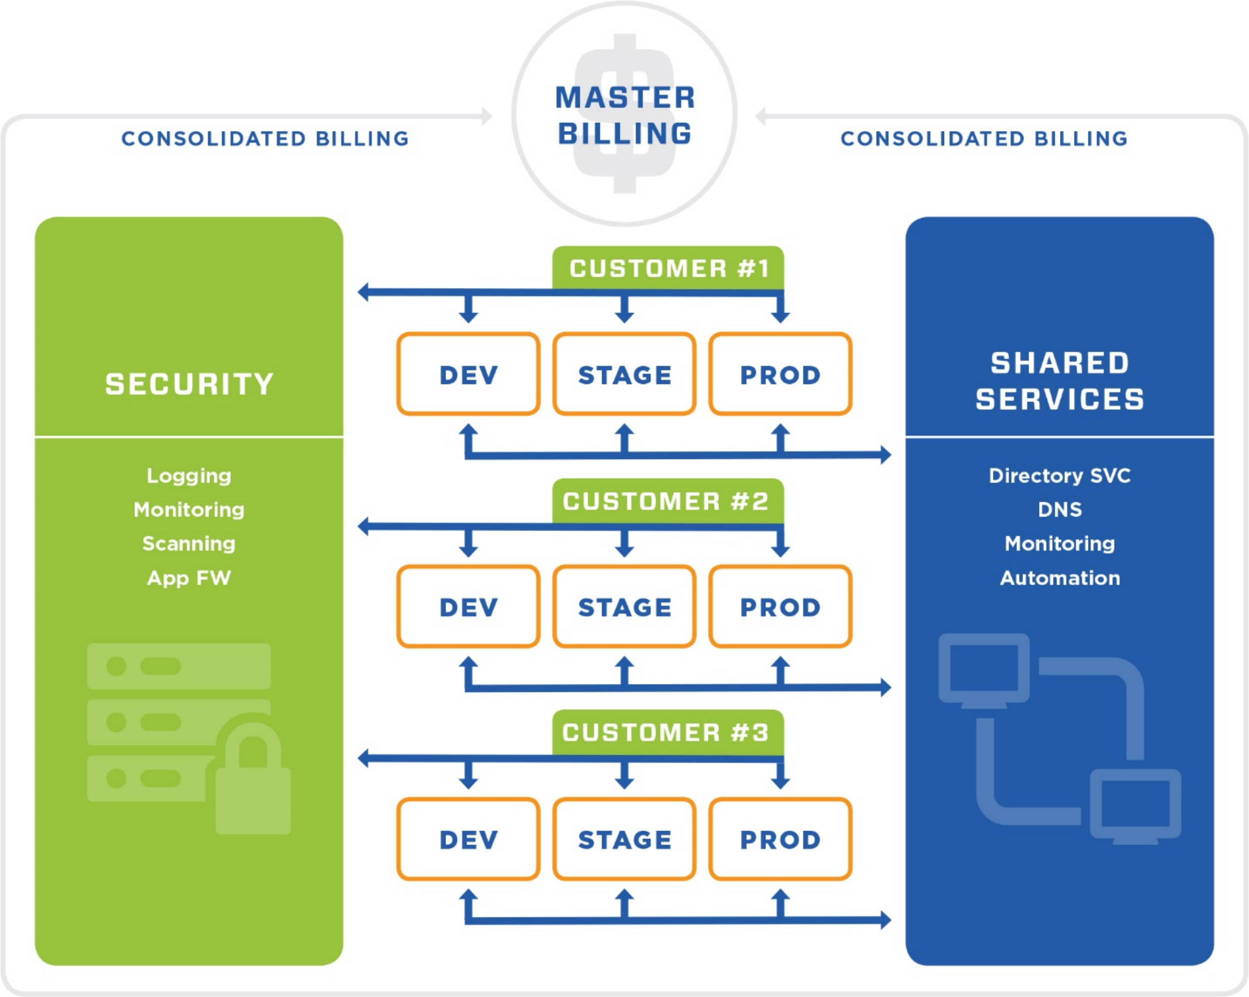

Figure 1 - Multiple Account Master Billing

OPTIMIZATION AND EFFICIENCY

Whether you are a third-party cloud broker or an organization with many independent groups utilizing cloud services, hopefully the design principles in this whitepaper have guided you down the right path to implementing a successful account strategy for your cloud environment. By developing a blueprint, you can make sure your organization can scale efficiently into the cloud using a sustainable model. Also, effectively implementing this model will allow you to implement granular security controls and track resource utilization for everything in your environment in a detailed manner.

Having this ability to analyze cost accrual in a detailed way will enable you to realize the financial benefit of cloud computing. Whether you use a tool like Splunk as we will introduce later or another tool to ingest your billing and usage data, it will be imperative that you understand some of the basic metrics that will present themselves.

- Cost per service;
- Cost per application;
- Cost per customer; and
- Cost per operating environment.

These metrics will provide you with the detailed information to make informed decisions regarding spending, and to more easily find opportunities to reduce waste. This data will also allow your engineering team to build more efficient designs.

ABOUT MINDPOINT GROUP

MindPoint Group is a cybersecurity consulting firm providing innovative solutions that include:

- Cloud Security
- FedRAMP 3PAO Services
- Proactive Security
- Security Architecture & Engineering
- Security Operations
- Governance, Risk & Compliance
- Managed Security Services

Our secure cloud solutions assist organizations in modernizing their legacy IT infrastructure(s) and transferring it securely to a virtualized, elastic, and efficient cloud infrastructure built on Amazon Web Services (AWS) and Microsoft Azure. We successfully apply our breadth and depth in cybersecurity to supporting clients like the National Aeronautics and Space Administration (NASA) where we are helping one of the first and largest cloud brokers in the Federal Government deploy a secure hosting solution to migrate the largest web presence in the Federal Government to the cloud. Very few businesses, large or small have designed and operated a cloud solution at this level for large organizations. Our groundbreaking cloud security solutions have resulted in multiple NASA awards.

ABOUT THE AUTHOR

Keith Rhea is a Senior Engineer with MindPoint Group. He specializes in Network Security, Cloud Security, and Cloud Optimization services for government and commercial sector clients. Keith has been providing Security and Engineering in Information Technology and Information Security since 2005. Keith is an AWS Certified Solution Architect with a Bachelor of Science degree in Computer Information Sciences from Washburn University.

LEARN MORE

For additional information about our cybersecurity services, please visit our website and social media:

- mindpointgroup.com
- LinkedIn
- Twitter
- GitHub
- Glassdoor
- Facebook

To learn more about MindPoint Group's Cloud Services, please email Keith and the rest of the cloud team at cloudsec@mindpointgroup.com