WHITE PAPER

Tagging Strategies:

Importance and Utility of Tagging Cloud Resources

DivvyCloud

What is Tagging?

TAGGING IS THE PROCESS OF APPLYING A SMALL, SEMANTICALLY SIGNIFICANT PIECE OF INFORMATION TO AN OBJECT. TAGGING STARTED TO BECOME COMMON PRACTICE IN TECHNOLOGY AND INTERNET PLATFORMS WITH THE INTERNET 2.0 BOOM OF SOCIAL MEDIA SITES THAT USED, FOR INSTANCE, TAGS ON A RESTAURANT REVIEW SITE TO IDENTIFY A TYPE OF CUISINE, A STYLE OF RESTAURANT, ATTRIBUTES OF THE RESTAURANT'S AMBIENCE, FAVORITE DISHES OR MORE. SIMILAR CONCEPTS WERE USED ON SITES FOR TAGGING PHOTOS, SAYINGS, MESSAGES OR REALLY ANY KIND OF SOCIAL MEDIA POST.

In the case of software-defined "cloud" infrastructure, the "objects" are cloud resources, whether virtual machines (instances), storage (volumes or object storage containers), networks (virtual private cloud subnet segments), security rules/access lists, user accounts or other cloud services.

Examining the cloud context in more detail, "tagging" typically refers to applying metadata tags to various elements of cloud infrastructure. Common examples are:

Applying tags to virtual machines to describe the operating system and applications that are running on each server.

Applying tags to virtual network segments in order to define the purpose of that network.

Applying tags to billable resources to identify cost ownership and facilitate chargeback.

Why Use Tagging?

In traditional IT infrastructure or data centers, the same tagging goals would often have been accomplished by putting physical labels on servers, network switches, or even racks and cages.

With the cloud, customers have no physical access to the infrastructure or data center environments, and therefore cannot see or touch their servers.

Tagging is crucial to help teams understand what the organization is running when cloud infrastructure gets beyond any modest scope. While the organization still has only a handful of virtual machines, running in a relatively steady state, it's easy to read a list of machines, log into each individually when needed, and keep the environments up-to-date and running properly.

However, when environments grow beyond a certain size, and/or are changing rapidly, the software-defined nature of the cloud makes it hard to keep track of resource across regions, cloud accounts and different cloud

platforms (AWS, GCP, Azure, OpenStack, etc.). Tags make it possible for a system administrator to find a given set of servers (virtual machines), or a network or other resource for maintenance or management purposes. Developers can then be sure to deploy applications only to appropriate cloud accounts, instance types and storage media tagged with the right application ID. And security professionals can quickly and easily drill down into any affected network, based on tags identified in log files by suspicious traffic patterns.

Tagging is so important for cloud environments that the recognized leader in the space, Amazon Web Services, recently increased their number of allowed tags to 50 for almost all cloud infrastructure resources on their platform.

Flying Blind Without Tags

MULTIPLE CONVERSATIONS WITH ENTERPRISE CLOUD CUSTOMERS CONFIRM THAT VERY OFTEN CLOUD COSTS CANNOT BE ALLOCATED CORRECTLY, AND WORKLOADS AREN'T PROPERLY UNDERSTOOD BECAUSE OF LACK OF TAGGING. THE FOLLOWING IS A SNAPSHOT OF INSTANCES ACROSS MULTIPLE CLOUD ACCOUNTS FOR A LARGE ENTERPRISE.

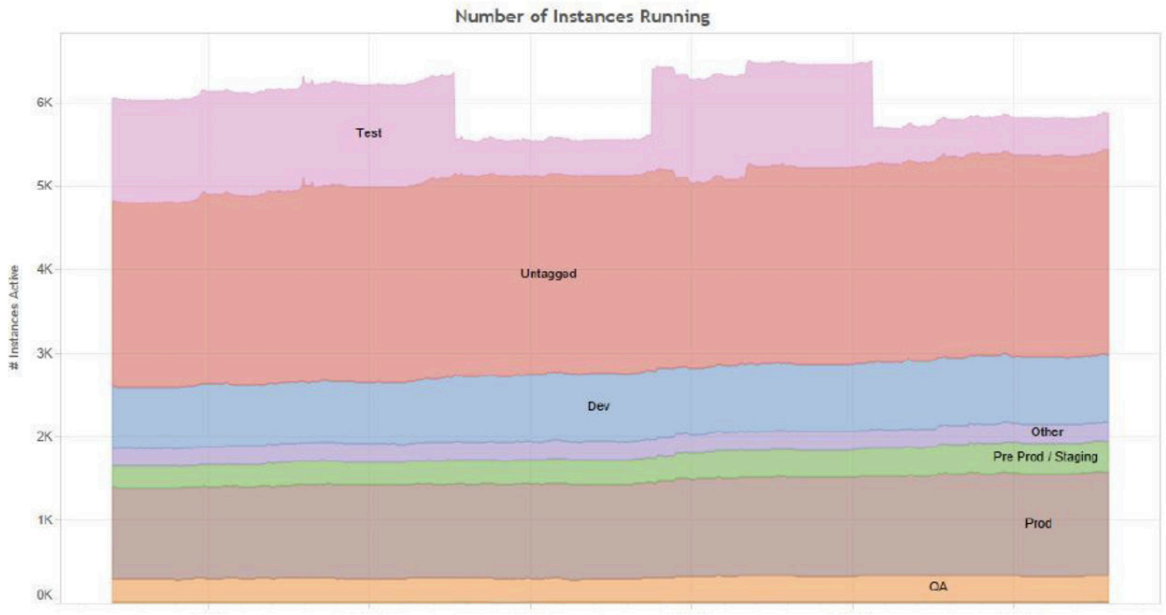

Some key observations from this utilization pattern:

- Overall utilization is increasing across all identified workloads
- Only test instances are stopped on any regular basis
- Within test instances, only a portion are stopped regularly
- **The largest group of virtual machines is the untagged group!**

With that being the case, there's no way for the customer to know whether these machines need to be running all the time. Also, there's no hint as to what security controls or compliance standards need to be verified for all of these servers.

Tags Enable Cloud Automation

A good tagging strategy provides a foundation for automating management of cloud infrastructure. The DivvyCloud platform harvests tag values along with other information about cloud resources. Based on those tags, valuable automation can be enabled to reduce costs, improve security and ensure compliance with industry or enterprise best practices. For example, an organization can ensure that instances or storage volumes with a HIPAA tag are located in the United States and properly encrypted.

DivvyCloud Bots can define non-compliance and trigger the appropriate automated actions based on the operational context provided by tags. Tags can also ensure that actions are NOT taken on specific cloud resources. For example, an overall policy might be for instances to be turned off at night and on weekends to significantly reduce monthly cloud bills. However, any instances with a "production" or "highly critical" tag is exempt from this automated process.

Lastly, Bots can help enforce appropriate tagging strategies. A very common use case is for every cloud instance to have a valid cost code tag. Bots can monitor the cloud environment for any new instances or changed instances, identify missing tags, notify the instance/application owner of the problem, and schedule the instance for termination in 24 hours if the cost center tag is not updated appropriately.

Basic Foundational Tags

This table outlines a number of tags we recommend customer implement as part of their initial tagging strategy. Then you can expand your strategy to include specific tags for your organization's use cases and specific requirements.

Tag	Scenario	Example Values
Environment	identify resources and set policy/workflow by stage of development	Dev, QA, Test, Prod
Resource Owner	for notification, visibility and user-based policies	MarketingAdmin, SalesEngineeringTeam
Cost Center / Accounting Code	for chargeback and budget purposes	#78925
Criticality	understand whether a machine is highly critical or not (useful for opt-out of automations)	P1, P2, P3
Application ID	identify all resources associated with a workload	Widget-2
Encryption	know whether encryption is required or not for data at rest	Required, Recommended, Optional, NoEncryption
Compliance	know whether a resource is subject to compliance/ audit requirements and trigger compliance automation	HIPAA-1, PCI
Schedule	know whether a machine is meant to be turned off on a regular schedule	6pm-6am, weekends off, business hours only, 10pm+
Age Limit	for environments or resources that are meant to be temporary	DeleteMe-12hrs, DeleteMe-30days
Operating System	and OS version, for management, patching and vulnerability assessment	Ubuntu-17-04
Application	and version(s), for management, patching and vulnerability assessment	Portal-123

Summary

Tagging is critical. With the remote, virtualized and dynamic characteristics of cloud platforms, tagging is the only way to effectively manage cloud infrastructure at scale. Tags allow teams of all types (operations, IT, security and development) to know what's what, and act appropriately on the right resources. Tags allow the organization to enable automation that maintains visibility, reduces cloud bills, and ensures security across your cloud deployments.

Please feel free to reach out to us to learn more about cloud management strategies related to tagging, automation and scaling of cloud operations.

DivvyCloud
Founded 2013

About DivvyCloud

DivvyCloud software enables organizations to achieve their cloud computing goals by simplifying and automating security, compliance and cost optimization of public and private cloud infrastructure. Using DivvyCloud, customers can leverage programmatic Bots to identify and remediate common cloud problems in real time. DivvyCloud was founded by seasoned technologists who understand first hand what is necessary to succeed in today's fast-changing, multi-cloud world. For more information, visit: divvycloud.com.

WRITERS
Peter Scott
Jeremy Snyder

DESIGN EDITORS
Tim Clise

CONTACT
DivvyCloud,
1400 Key Blvd,
Level A, Suite #6,
Arlington, VA,
22209, USA,
+1 (571) 290-5077
info@divvycloud.com

DIVVYCLOUD

1400 Key Blvd,
Level A, Suite #6
Arlington, VA 22209
divvycloud.com