

VAST View™

IT and Operations: Cloud Best Practices



POWERED BY **CloudHealth**

Table of Contents

Introduction to VAST View™ Best Practices for IT & Operations Professionals.....	3
Administrative Best Practices.....	3
Managing Cost.....	9
Managing Usage.....	28
Managing Performance with Rightsizing	44
Managing Migrations from the Data Center to the Cloud	47
Managing Security Risks.....	49
Executive Reporting Best Practices	51
Key Performance Indicators and Metrics	52
Frequency of Reports and Policies	61

Introduction to VAST View™ IT and Operations Cloud Best Practices

VAST View provides a consolidated view across your dynamic and fast-changing cloud environment. Infrastructure & Operations professionals responsible for the cloud need to optimize their resources to drive successful business outcomes.¹ Cloud optimization hinges on analyzing and managing data efficiently.

VAST View Technologies IT & Operations best practices look at reporting, analysis, and metrics used by IT & Operations professionals, across industries, on a recurring schedule. Under each task below, VAST View references reporting that you can use to provide key stakeholders (e.g. Executives and Cost Center Owners) visibility into their cost and usage trends, as well as the ability to automatically identify and remediate inefficient resources, and security vulnerabilities.

Administrative Best Practices

VAST View IT & Operations: Cloud Best Practices focuses on managing and optimizing cloud spend and usage through the use of **Perspectives** to create views by Project and Environment. If your business reports on spend and usage in other ways, you need to create a Perspective based on this method, as well. These require Perspectives to be built by account or asset metadata to increase visibility in reporting.

VAST View recommends that each budget owner maintains Perspectives and reporting for their various teams, along with using a role called a “Cloud Steward.” An individual in this role is responsible cross-functionally for the ongoing optimization and governance of the cloud infrastructure, serving as a liaison between various IT/Operations teams and Finance. The Cloud Steward:

- **Defines and manages automated Policies**
- **Analyzes cost, usage, performance, and security across environments**

¹ *Align Your Cloud Infrastructure to Your Business Strategy in Four Steps...or Fail*, Michael Warrilow, Bob Gill, 22 September 2017

Along with a Cloud Steward role, additional roles and permissions in VAST View should be configured and managed through SSO when possible. This automatically reflects internal organizational changes with the appropriate permissions for reporting and access in VAST View without additional manual efforts by other teams. VAST View supports multiple SSO configurations and role permissions, found here.

Businesses are more successful when there is one budget owner from each cost center or Engineering team, along with the Cloud Steward, involved in reporting and optimization efforts. Those users and the Cloud Steward should have VAST View Admin permissions, which allow them to manage users, account access to AWS, Azure, GCP and Data Center, as well as view all account spend and usage in the product. More limited permissions for custom or default roles can be configured in cloud management platform using the instructions here. This includes using VAST View Organizations to limit the accounts visible to users.

All accounts must be **configured** in VAST View before configuring reporting and visibility into cost, usage, performance and security. This enables VAST View to pull metrics and additional metadata about resources. While cloud management platform collects some resource metrics with the initial configuration of AWS accounts, additional metrics, like Memory and Disk, can be gathered:

- **When integrating your Datadog or New Relic accounts with VAST View**
- **By pushing metrics to VAST View with the Metrics API**
- **By installing the VAST View Agent**

You can also integrate configuration management tools, like **Chef**, with VAST View to utilize tags and metadata from those tools, as well as improve configuration as a result of the best practice reporting.

Azure Service Principals provide VAST View with CPU, Memory and Disk I/O metrics once the Service Principal is configured in VAST View. To enable these in Azure, turn on the VM Diagnostic feature for all VMs in the Azure portal. Instructions from Azure can be found here. To obtain Disk Usage metrics, you can install the VAST View Agent on the VMs for reporting.

VAST View's Data Center capabilities product can gather metrics from the VAST View Agent or integrate with VMware. To set up your Data Center account without VMware, instructions can be found [here](#). To integrate your VMware accounts with VAST View, instructions for configuring the aggregator can be found [here](#).

Many best practice recommendations in the following VAST View **Policies** have the option for automation to reduce manual work when optimizing your environment. To enable automation and action options in reporting and Policies, you must update permissions in AWS and Azure following these instructions ([here for AWS](#)) and ([here for Azure](#)).

Defining a Tagging Strategy

Reporting and analysis is much easier with clean data that follows a set strategy for being allocated to various business groups. Tags are the basis for Perspectives and requires an initial strategy to report and govern against. To create a tagging strategy, multiple business leaders need to agree on ways data should be reported on, analyzed and governed. For example, Finance may want to look at spend by cost center for month-end reporting, Product Management may care about spend per customer or project to track margins, while DevOps and Engineering may focus on environments or microservices to monitor their various teams and what they're utilizing.

Developing a Tagging Design Guide enables engineers to reference tags and their value structures when spinning up new infrastructure. This document should include:

- **A set of 5-10 tags that all areas of the business need for reporting and analysis. Common tags used by customers can be found [here](#).**
- **A structure for the tag values, considering case sensitivity, prefixes, and formats like email addresses using '@' symbols**
- **A list or dictionary of approved values that is consistently and regularly updated**

For services where either AWS or VAST View doesn't support tags, you can use VAST View Reallocation Rules to split the spend for the services across Perspective groups. Another option is to allocate all or a specified percentage to a Perspective group. To implement these rules, see the documentation [here](#) to learn how they work and options for allocating spend without tags.

After defining a tagging strategy, enforcing these tags going forward helps keep reporting clean and efficient, as well as limits the effort of retroactively cleaning up data in the future. Configuration tools are a great option for implementing tagging structures to ensure that tags are set with appropriate values from the initial launch. To monitor tags, use VAST View Policies to send alerts for assets that are missing the tags in your tagging schema, as well as use regular expressions to monitor the structure of the values. The following Policies can be done holistically for all accounts if managed and monitored by a centralized group, or by account to send to the respective account owners.

Amazon Assets Missing Tags

The following VAST View Policy monitors your Amazon Assets for a specific set of tags. This should be evaluated daily after the Perspective has been allocated as much as possible to reduce the noise in the Policy alerts. By creating a block in Policies for each tag that you want to monitor, the Policies will evaluate each block individually.

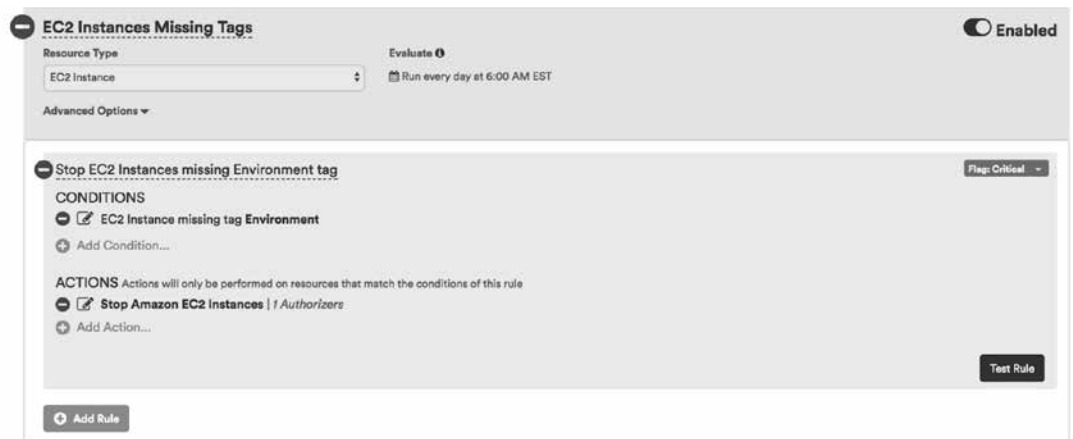


Conditions for Amazon Asset Tagging Policy:

- Resource type = Amazon Assets
- Evaluate Interval = Daily
- Conditions = Missing tag and choose a tag from your tagging strategy that should be governed and enforced
- Actions = Email Cloud Steward. If not governing centrally, this can be filtered within the condition to look at specific accounts and sent to those account owners

Specific Resources Missing Tags

By choosing specific resource types instead of looking at all Amazon Assets, additional conditions and actions are available to be combined with the missing tags conditions. The best practice would be to use automation when possible to manage the violating assets daily. Utilize the following policy conditions to create an automated approach to governing your strategy.




Conditions for Resource Tagging Policy:

- Resource Type = EC2 Instances
- Evaluate Interval = Daily
- Conditions = Missing tag and choose a tag from your tagging strategy that should be governed and enforced
- Actions = Stop Amazon EC2 Instances (either automatically or with an approval chain) although other automated actions, like Launching Lambda Functions to tag the instance or Deleting EC2 Instances is also an effective way to govern the tagging strategy but may be more severe than stopping the instance

Managing Cost

Best Practice Cost Reporting

Using Perspectives to allocate spend based on the environment, service, project or customer, and/or cost center enables the business to track spend broken down to analyze and make recommendations for optimizations. It's recommended review the **Cost History** report on a monthly and daily basis to monitor the changes in spend. For trend analysis, subscribe to monthly view on a monthly basis; Subscribe to the daily view on a daily basis to catch anomalies before they incur cost for too long.

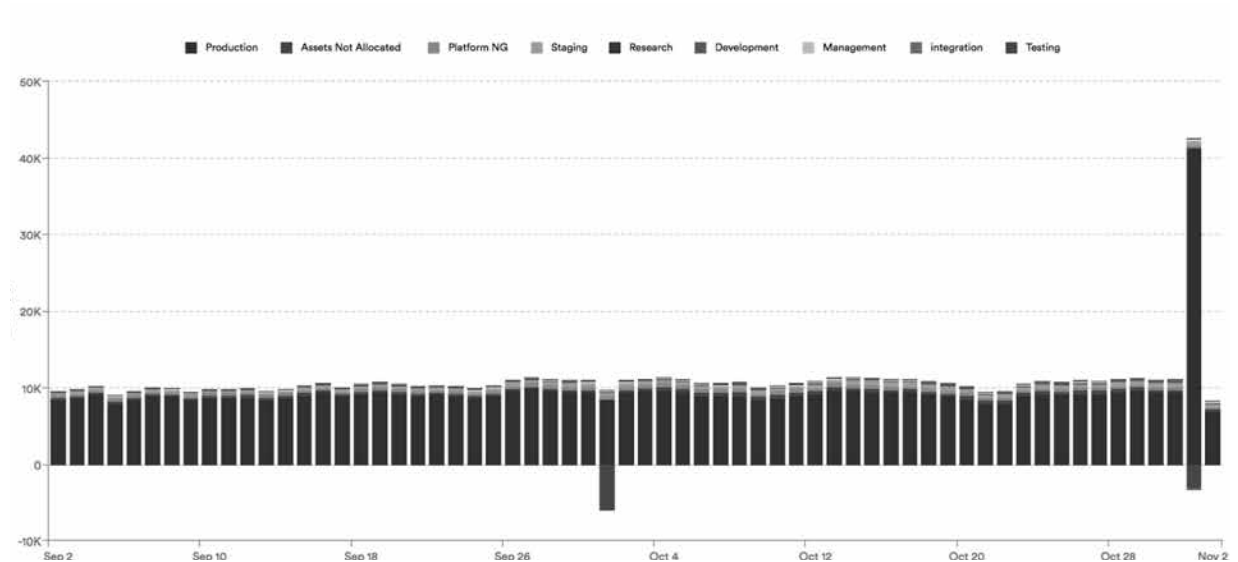


Conditions for Monthly Perspective Cost Trend Report:

- Interval = Monthly
- X-Axis = Months
- Y-Axis = Cost
- Category = Choose Perspective to report on
- Filters = All completed months

Similar reports are available for Azure Cost History and Data Center Cost History.

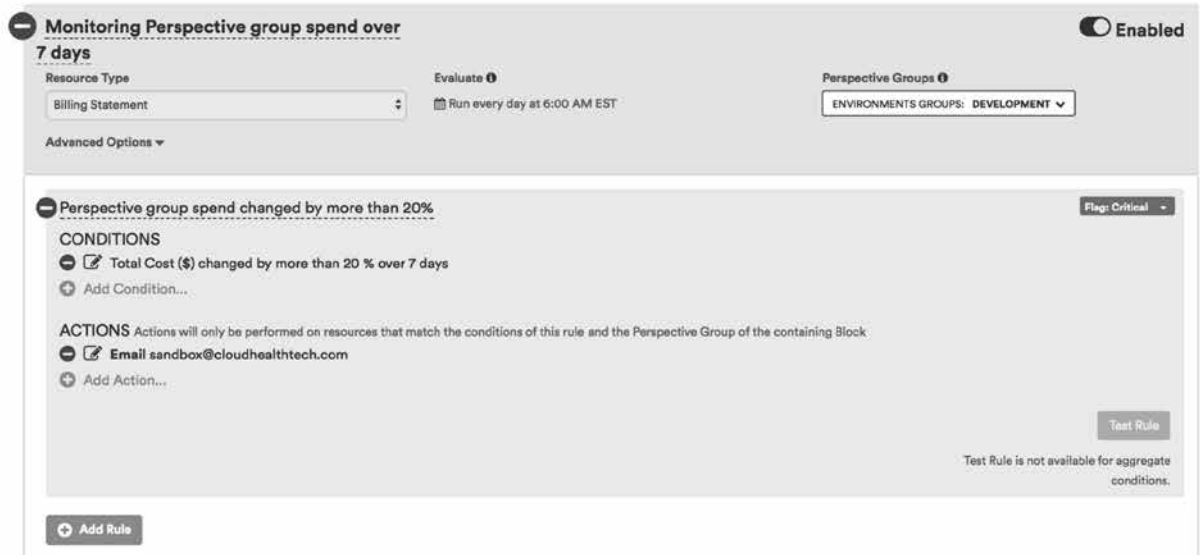
Monitoring daily cost to track the trends at a more granular time interval enables the business to identify and act quickly on abnormal changes to ensure spend stays in line with expectations.



Conditions for Daily Perspective Cost Trend Report:

- Interval = Daily
- X-Axis = Days
- Y-Axis = Cost
- Category = Choose Perspective to report on
- Filters = Filter out last day from the Days filter as the last day's data may be incomplete. If credits are distorting the reporting views, use the Charge Type filter to exclude certain credits

Monitoring changes in Perspective groups with Policies, to send alerts to the owner of an environment, microservice, or cost center allows for less manual monitoring and alerts the appropriate contacts. The following Policy monitors changes in spend over a week and should have a baseline percentage set based on historical trends. The Policy should be created for each Perspective group and the email contact should change according to who owns that group.



The screenshot shows the configuration for a Monitoring Perspective group spend over. The rule is titled "Monitoring Perspective group spend over" and is currently "Enabled". The evaluation interval is set to "7 days". The resource type is "Billing Statement", and the evaluation frequency is "Run every day at 6:00 AM EST". The perspective groups are set to "ENVIRONMENTS GROUPS: DEVELOPMENT".

Conditions:

- Total Cost (\$) changed by more than 20 % over 7 days

Actions:

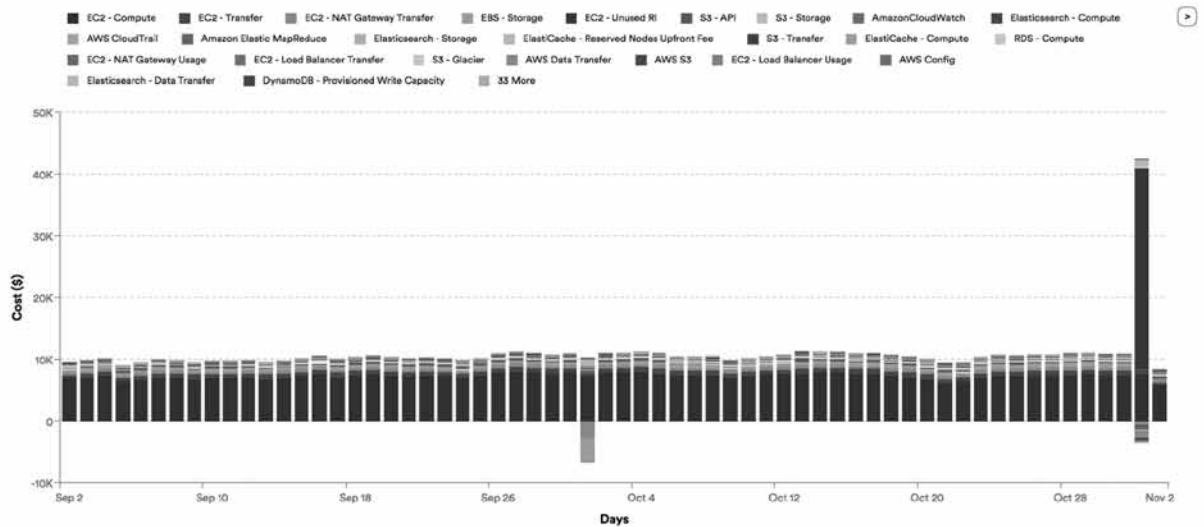
- Email sandbox@cloudhealthtech.com

Buttons for "Add Rule", "Test Rule", and "Add Condition..." are visible. A note at the bottom right states: "Test Rule is not available for aggregate conditions."

Conditions for Perspective Group Cost Monitoring Policy:

- Resource Type = Billing Statement
- Evaluate Interval = Daily
- Conditions = The percentage increase should be determined by the business, although 10-20% is common in these alerts.
- Filters = Select a Perspective at the top of the Policy, which enables a group to be selected in the Policy Block
- Actions = Email Perspective Group Owner or Cloud Steward


Monitoring spend through the lens of Service Items enables the business to pinpoint the cause of variances in spend from month to month. You can add more files to customize this for each cost center owner, business function owner, etc. so that they can also view and monitor their own changes in usage. It's recommended to look at Service Items daily to catch unusual changes before they grow and become an issue.



Conditions for Service Item Cost Trend:

- Interval = Daily
- X-Axis = Days
- Y-Axis = Cost
- Category = Service Items
- Filters = Filter out last day from the Days filter, as the last day's data may be incomplete. If Cost Center Owners are monitoring their own spend, they can filter to a specific Perspective group

While it's important to track variances and changes in cost reporting, comparing actual spend to budgets set by Finance holds teams accountable for their spend, no matter the changes occurring. The VAST View **Budget versus Actual** report enables Cost Center Owners to use budgets set in VAST View by Finance to reports on actual spend and variances to the budget. Cost Center Owners have the option for budgets to rollover from previous months to see variance on a year-to-date level, as well. This report should be subscribed to the Cost Center Owners on a monthly basis.



Conditions for Budget vs Actual Cost Report:

- Cost Center
- Date range = From earliest month available to the last month of the year
- Choose Rollover in the chart options to give visibility into cumulative budget vs actual metrics


Similar monitoring can be done with Policies. The example alerts Cost Center Owners for only their cost center or team's budget. This policy should run daily, based on projected spend and filtered for the cost center belonging to that individual.




Conditions for Budget Policy:

- Resource Type = Billing Statement
- Evaluate Interval = Daily
- Conditions = Monthly Projected Cost and the percentage increase should be determined by the business, although 110% (10% over the budgeted amount) is common in these alerts
- Filters = Select a Perspective at the top of the Policy which will allow a group to be selected in the Policy Block for individual budgets
- Actions = Email Budget Owner or Cloud Steward if not specifying a Perspective group

To track the business's EA commitment, use the **Azure Burndown Cost** report. The EA committed amount is entered in the account setup page when configuring a new **Azure Enrollment**. This enables the business to enter multiple enrollments with varying start dates and track against the total enrollment value.



You can see in the default reporting view that the EA balance trends downwards each month as it is used up. However, a view for executive reporting should track the total spend compared to the commitment. This should be reviewed monthly for month-end reporting.



Conditions for Azure Burndown Report:

- X-Axis = Months
- Y-Axis = Balance, Total Cost
- Category = Enrollments

Managing Cost with Reservations

Reservations enable you to make an upfront monetary commitment in return for a discount on your compute cost. By having customers commit to the usage of specific infrastructure, cloud providers can better manage their capacity and therefore pass these savings to their customers. For more information on what reservations are and how to leverage these for future optimizations, see our Help Center article for AWS reservations [here](#), and Azure reservations [here](#). The best practice is to cover at least 70% of your infrastructure with reservations.

Purchasing AWS Reservations

The VAST View **EC2 RI Optimizer** evaluates your AWS instances to determine if you can benefit from the use of Reserved Instances. During the evaluation, the Optimizer analyzes your previous usage and identifies the optimum purchase for a given budget.

VAST View uses a constraint-based algorithm that analyzes hourly on-demand usage over a time period. The algorithm compares options until it finds the most efficient option for a reservation purchase that fits within your budget constraints. The aim is to provide the best recommendations for your specified constraints such as budget, purchase date, and filters.

Once finalized, quotes can be saved and sent to business owners for approval before purchasing. Purchases can be made through VAST View, using approval processes set by the business, or can be exported and sent to an AWS rep to purchase on your behalf.

OPTIONS ▾

Time Period To Analyze: Last 30 Days | Preferred Term: 1 Year | Target Coverage: 100% | Set Your Budget (Max: \$573,852): \$0 to \$573,853

EXTENDED OPTIONS ▾

Estimated Purchase Date: 03/20/2018
 2 reservations expire by this date
 Only recommend expiring reservations

Billing Account: AWS CHT Sandbox | Purchase Account: AWS CHT Sandbox

Offering Types: All Upfront, Partial Upfront, No Upfront

Choose Lower Prepay Res Type if Savings < 3%

Reserve Capacity | Use Convertible

Filters: **Choose Filter** ▾

Run Optimizer

Recommendation Summary

INSTANCES	RESERVATIONS TO BUY	COVERAGE	UPFRONT COST	BREAK EVEN	SAVINGS	MONTHLY SAVINGS	ANNUAL SAVINGS
1,483	386	95%	\$573,692	6.3 mo	15%	\$42,651	\$511,813

Recommendations by Location


LOCATION	INSTANCE TYPE	ACCOUNT	OS	TENANCY	INSTANCES	INSTANCES (NORMALIZED)	CURRENT	BUY	UPFRONT	COVERAGE	CURRENT	PROJECTED MONTHLY	CURRENT	PROJECTED ANNUAL		
us-east-1					1,483	15587.5	1042	386	\$573,692	95%	\$285,471	\$242,820	\$42,651	\$3,425,652	\$2,913,839	\$511,813
	c3 size flexible	AWS CHT Sandbox	Linux/UNIX	default	157	2380	151	3	\$4,564	98%	\$27,423	\$27,051	\$361	\$329,073	\$324,737	\$4,335
	c4 size flexible	AWS CHT Sandbox	Linux/UNIX	default	40	1208	3	33	\$134,821	96%	\$21,238	\$13,641	\$7,597	\$254,852	\$163,694	\$91,159
	i3 size flexible	AWS CHT Sandbox	Linux/UNIX	default	362	4392	188	135	\$229,066	85%	\$110,592	\$93,656	\$16,936	\$1,327,106	\$1,123,877	\$203,228
	m3 size flexible	AWS CHT Sandbox	Linux/UNIX	default	388	1910	354	31	\$30,722	99%	\$33,936	\$31,654	\$2,282	\$407,236	\$379,851	\$27,385
	m4 size flexible	AWS CHT Sandbox	Linux/UNIX	default	149	1180	72	74	\$37,736	98%	\$17,180	\$12,981	\$4,199	\$206,166	\$155,777	\$50,389
	m5 size flexible	AWS CHT Sandbox	Linux/UNIX	default	47	380	0	46	\$46,169	98%	\$6,547	\$4,047	\$2,600	\$79,764	\$48,565	\$31,200
	r4 size flexible	AWS CHT Sandbox	Linux/UNIX	default	317	4064	265	53	\$88,074	99%	\$67,319	\$58,899	\$8,420	\$807,831	\$706,793	\$101,038
	t2 size flexible	AWS CHT Sandbox	Linux/UNIX	default	23	73.5	9	11	\$2,550	87%	\$1,135	\$879	\$257	\$13,623	\$10,544	\$3,079
Total					1,483	15587.5	1042	386	\$573,692	95%	\$285,471	\$242,820	\$42,651	\$3,425,652	\$2,913,839	\$511,813

Conditions for EC2 RI Optimizer:

- Choose Time Period to Analyze = Last 30 days
- Preferred Time = 1 year; however, 3 year RIs give you larger discounts if your organization's forecasted usage is more stable.
- Target Coverage = 100%. VAST View recommends up to 100% while only recommending the reservations that are cost-beneficial to reserve. The best practice is to be at least 70% reserved.
- Budget = Max budget provides optimal RI recommendations; adjust this to the amount of upfront capital your business can provide, or to see different scenarios at different budget points.
- Estimated Purchase Date = Setting this in the future will take into account RIs that expire between now and that date. Set this to the date that you would like to purchase for to ensure you're including all active RIs at the time. If you set the Purchase Date to a future date, VAST View will alert you if there are reservations expiring by that purchase date and include them in the recommendation. There is also an option below the Purchase Date to recommend only the expiring RIs if you do not wish to purchase additional RIs.
- Offering Types = Best practice is to allow the RI Optimizer to choose the most optimal RI type. Choose the type that accommodates your business's upfront payment ability if certain types are not feasible, but, in most cases, Partial Upfront RIs provide very similar benefits to All Upfront RIs.
- Reservation Types = There are reservation type options to the right of the Offering Types to select for analysis. By default, VAST View will recommend regional, size-flexible reservations.

Tracking Reservation Savings

Obtaining approval for upfront spend for Reserved Instances can be difficult to get from the business in some cases. By tracking the actual savings achieved from RIs, future purchases can be justified easier and past purchases can be proven successful to the business. The **Cost Savings** report is a good visual to present as new RI purchases are made to prove the value of the upfront budget approved to obtain these reservations.



Conditions for EC2 RI Savings Report:

- Interval = Monthly
- X-Axis = Months
- Y-Axis = EC2 Savings
- Category = Cost Category Items

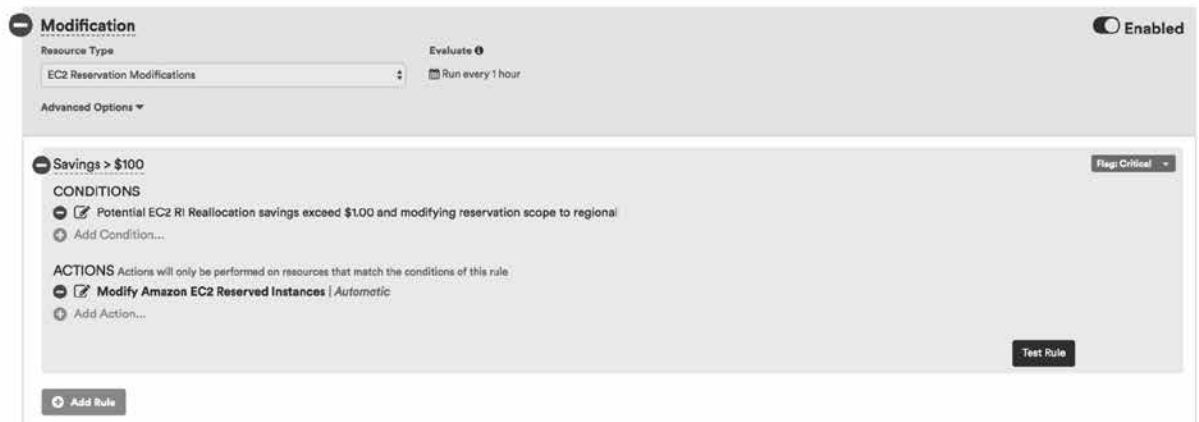
EC2 RI Modifications

The **EC2 RI Modifier** analyzes your Reserved Instances as well as your instance usage, and makes cost-saving modification recommendations to consider. Creating an hourly automated policy, like below, will reduce the need to manually modify these through the report.

Once reservation purchases are completed, monitoring and managing these RIs on a daily basis is important for recognizing the full value of the large investment made. VAST View Policies and reporting recommendations can alert DevOps and Perspective Group Owners when reservations are underutilized as well as when reservation purchases are expiring and need to be repurchased. Policies can then take actions on some of these underutilized reservations on your behalf.

VAST View gives you the ability to modify coverage for all instances tied to your reservation, or just a subset, in one or more of the following ways:

- Switch availability zones within the same region
- Change between EC2-VPC and EC2-Classic
- Change the instance type within the same instance family



Conditions for EC2 Standard RI Modification Policy:

- Resource type = EC2 Reservation Modifications
- Evaluate interval = 1 hour
- Conditions = Savings exceeds \$1 for modifying reservations scope to regional. Note that moving standard AZ reservations to a regional reservation is best practice, but can be changed to maintain scope by editing the policy condition if reserving capacity is necessary for your organization.
- Actions = Modify Automatically

For Convertible Reserved Instances, AWS allows these reservations to be exchanged from one instance type to another instance type within their guidelines listed here. This allows the reservation purchase to be more flexible and align with the business’s infrastructure changes over time.

The **RI Analyzer** is a great report for seeing how size flexible reservations are applied to instance usage now or throughout the current month. This analysis should be done on a weekly basis at minimum, while, if possible, daily analysis would reduce the value of unused RIs even further.

		1811.0 # Instances	1623.7 # Reservations	73.8% Ris (% of Total Hours)	\$90,476 Reservation Savings					
Usage Discount										
FILTERS <input type="checkbox"/> Edit Columns... Analysis Period: Now Account Analysis By: Billing Family										
Billing Account	Recommendations	Location	Instance Type	Operating System	Tenancy	# Instances	# Reserved Instances	Effective Cost Per Month	Reservation Savings Per Month	Potential Savings Per Month ¹
AWS CHT Sandbox	Purchase 37 Partial Upfront reservations	us-east-1e	i3.2xlarge	Linux/UNIX	default	37.0	0.0	\$16,845	\$0	\$5,886
AWS CHT Sandbox	Purchase 36 Partial Upfront reservations	us-east-1d	i3.2xlarge	Linux/UNIX	default	36.0	0.0	\$16,390	\$0	\$5,727
AWS CHT Sandbox	Purchase 27 Partial Upfront reservations	us-east-1b	i3.2xlarge	Linux/UNIX	default	32.0	4.5	\$14,059	\$510	\$4,295

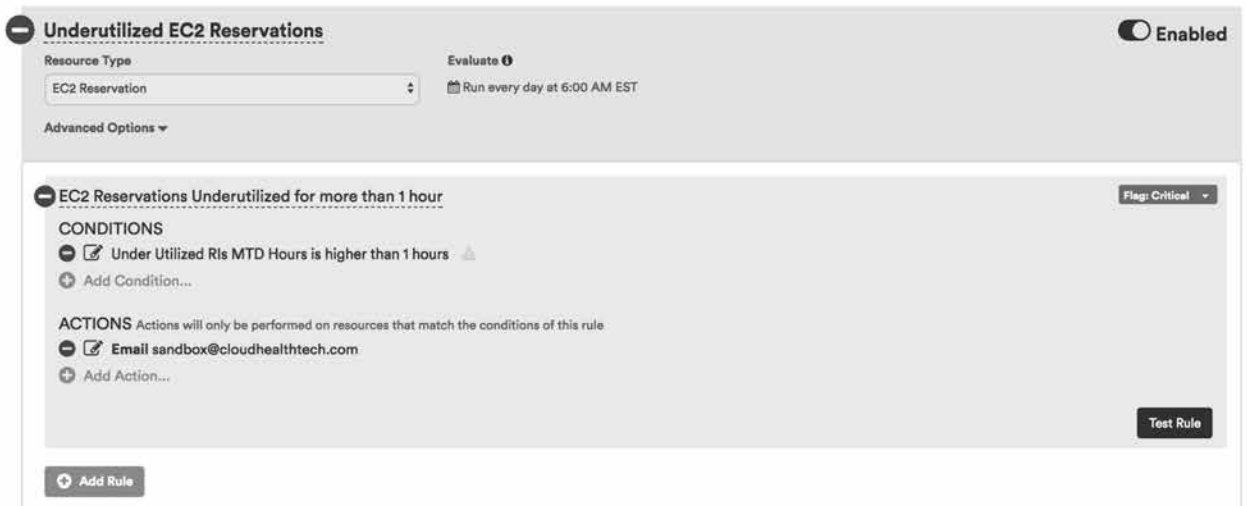
To best way to analyze this report for Convertible RI exchange opportunities is to look at the “Now” analysis period. VAST View will provide recommendations when reservations are underutilized to signify when to sell, modify or convert those reservations. Export this report and filter for the lines that have this recommendation.

Billing Account	Recommendations	Location	Instance Type	Operating System	Tenancy	# Instances	# Reserved Instances	Effective Cost Per Month	Reservation Savings Per Month
AWS CHT Sandbox	Sell, modify or convert 65 reservations	us-east-1	r3.xlarge	Linux/UNIX	default		65.0	\$9,944	-\$9,944

As seen in the previous example, the r3.xlarge Linux line has 65 reservations, or 65 reservations x 8 credits per reservation, while having 0 instance usage to apply these reservations to. This results in 520 reservation credits that can be converted to other instance types. Calculate the credits available to be converted in the export for all remaining lines to come up with a total number of available usage. To determine where to move these reservation credits to, filter the export for all lines that say “purchase X reservations.” Take a look at the instance families represented and determine a number of credits per family that are needed. From there, you can use the AWS console to convert these reservations to the appropriate family type.

Underutilized EC2 RIs

Ensuring that Reservations are utilized every hour helps maximize the benefit of your purchases. Having Policies alert when Reservations are unused for more than 1 hour gives you the opportunity to review changes that can be made in your infrastructure to better utilize these RIs, or identify opportunities for you to sell, modify or convert your existing Reservations.



The screenshot displays the configuration for a CloudWatch policy titled "Underutilized EC2 Reservations". The policy is currently "Enabled". The "Resource Type" is set to "EC2 Reservation" and the "Evaluate" frequency is "Run every day at 6:00 AM EST". The policy name is "EC2 Reservations Underutilized for more than 1 hour" and it is marked as "Flag: Critical".

CONDITIONS

- Under Utilized RIs MTD Hours is higher than 1 hours
- Add Condition...

ACTIONS Actions will only be performed on resources that match the conditions of this rule


- Email sandbox@cloudhealthtech.com
- Add Action...

Buttons for "Add Rule", "Test Rule", and "Add Condition..." are visible.

Conditions for Underutilized EC2 Reservation Policy:

- Resource type = EC2 Reservation
- Evaluate interval = Daily
- Conditions = Underutilized RIs MTD Hours are higher than 1 hour
- Actions = Email Cloud Steward

In addition to a policy, tracking the trend of unused RIs month-over-month gives the business visibility into how the RI strategy is performing and where adjustments need to be made to better utilize the reservations. A cost history report showing the unused RI amount trended over the past 12 months is helpful for visualizing Partial and No Upfront monthly reservation fees. However, to track all unused RIs, utilizing the **EC2 instance usage** report to visualize which instance families are not using all of the purchased RI credits will be more encompassing. The following report uses Normalized Factor Units (NFUs) which AWS describes here.



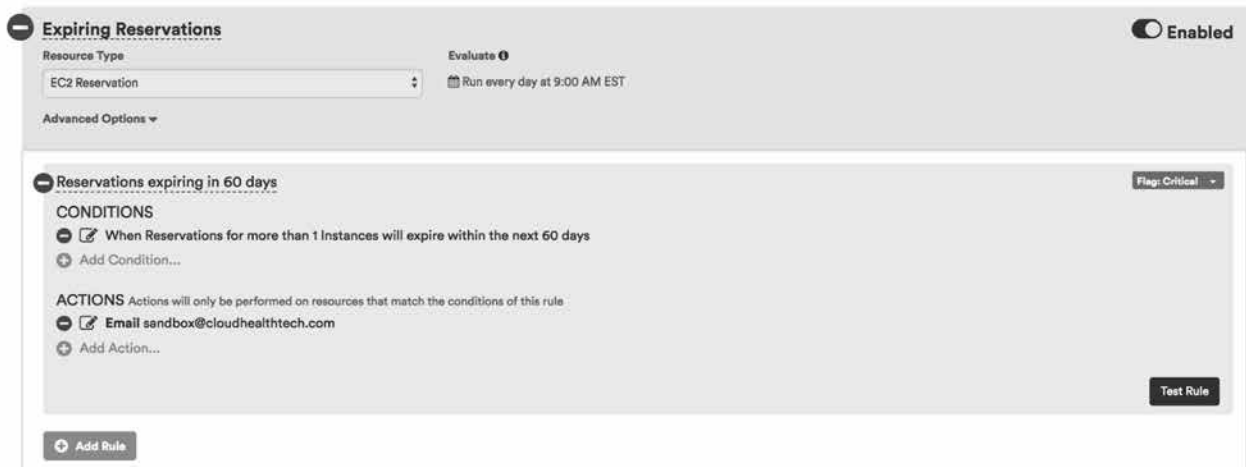
Conditions for EC2 NFU Report:

- Interval = Monthly
- X-Axis = Region
- Y-Axis = Normalized Reservation Surplus (NFU)
- Category = Instance Type Families
- Filters = Months = Previous Month, Operating System = Filter for each OS to align with the reservation purchase attributes

Any positive number of units indicates that there are too many reserved units purchased for that instance family and that region. The report should be filtered to each operating system to align with how the reservations were purchased. Any negative number of units indicates that there are not enough reservation units to cover the usage in that instance family and region and more reservation credits could be purchased.

Upcoming EC2 RI Purchases Expiring

Being aware of previous RIs expiring within 30, 60 or 90 days enables DevOps and Finance to budget for another cash outflow, and begin discussions with Cost Center Owners to determine what Reservations need to be repurchased. A policy to alert the Cloud Steward when reservations are about to expire reduces the need to track these in spreadsheets.



The screenshot shows the AWS IAM Policy console configuration for a policy named "Expiring Reservations". The policy is currently "Enabled".

- Resource Type:** EC2 Reservation
- Evaluate:** Run every day at 9:00 AM EST
- Advanced Options:**
 - Reservations expiring in 60 days:**
 - CONDITIONS:** When Reservations for more than 1 Instances will expire within the next 60 days
 - ACTIONS:** Email sandbox@cloudhealthtech.com

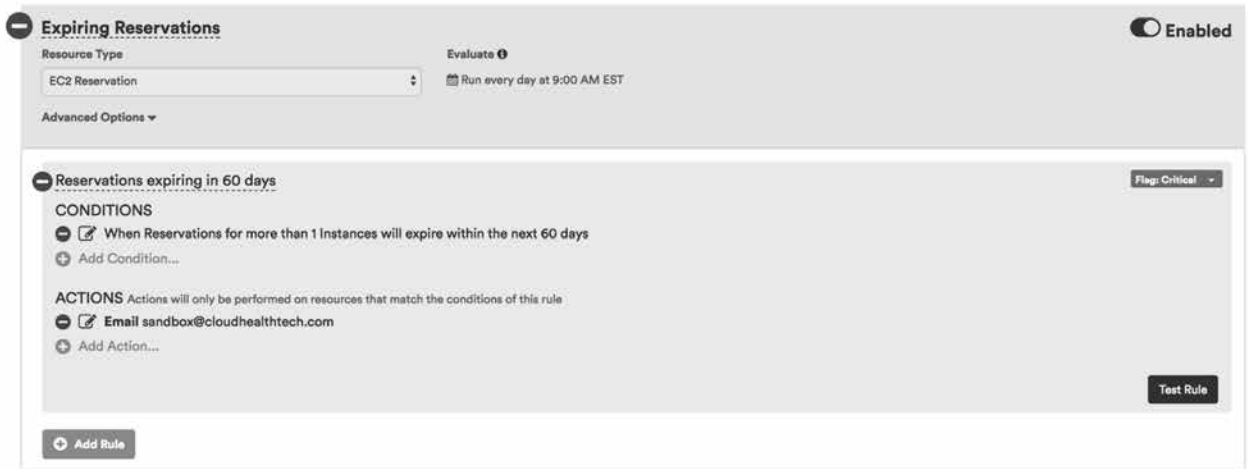
Buttons for "Add Rule", "Test Rule", and "Add Condition..." are visible.

Conditions for Expiring Reservation Policy:

- Resource type = EC2 Reservation
- Evaluate interval = Daily
- Conditions = When Reservations for more than 1 instance expire in the next 30, 60, and 90 days (3 different Policies or blocks)
- Actions = Email Cloud Steward

Managing RDS Spend with Reservations

Similar to EC2 Reservations, VAST View has an **RDS RI Optimizer** that can analyze the on demand usage in your environment to determine the most optimal RI purchases.



Save \$28,832 per month (\$345,981 per year) by purchasing 2 All Upfront and 99 Partial Upfront reservations for a total upfront price of \$129,380

The Effective Cost Savings including the amortized upfront price will be \$18,050 per month (\$216,601 per year)

The projected reservation rate is 66.9% and the payback period is 4.5 months

Conditions for RDS RI Optimizer:

- Choose Time Period to Analyze = Last 30 days
- Preferred Time = 1 year but 3 year RIs give you larger discounts if your organization's forecasted usage is stable.
- Budget = Max budget gives optimal RI recommendations, but adjust this to the amount of upfront capital your business can provide.
- Estimated Purchase Date = Setting this in the future will take into account RIs that expire between now and that date. Set this to the date that you would like to purchase for to ensure you're including all active RIs at the time.
- Offering Types = Best practice is to allow the RI Optimizer to choose the most optimal RI type. Choose the type that accommodates your business's upfront payment ability if certain types are not feasible.

Purchasing Azure Reservations

The VAST View **VM RI Optimizer** evaluates your Azure VM usage to determine if you can benefit from the use of RIs. During the evaluation, the Optimizer analyzes your previous usage and identifies the optimum purchase. By setting a target RI coverage, VAST View will find the most optimal purchase within that coverage constraint. Additional filters can be selected to narrow recommendations down to specific subscriptions, machine series or regions.

OPTIONS ▾

Time Period To Analyze: Last 30 Days ▾ |
 Preferred Term: 1 Year ▾ |
 Target Coverage: 100 % |
 Set Your Budget (Max: \$2,050)

EXTENDED OPTIONS ▾

Estimated Purchase Date: 03/02/2018 |
 Reservation Scope: Shared Subscription Single Subscription |
 Enrollment: cht-ea |
 Subscription: All

Filters

Recommendation Summary

RESERVATIONS TO BUY	COVERAGE	UPFRONT COST	BREAK EVEN	SAVINGS	MONTHLY SAVINGS	ANNUAL SAVINGS
4	99%	\$1,543	5.2 mo	52%	\$168	\$2,012

COLUMNS: (10) ▾ RECOMMENDATIONS BY: LOCATION ▾

Recommendations by Location

Purchasing reservations under an Enrollment Agreement (EA) with shared scope will still require choosing a subscription in Azure portal. You can choose any subscription.

LOCATION	VIRTUAL MACHINE SIZE	ENROLLMENT	ELIGIBLE SUBSCRIPTIONS	CURRENT		UPFRONT	COVERAGE	PROJECTED MONTHLY		PROJECTED ANNUAL			
				BUY	SAVINGS			CURRENT	SAVINGS	CURRENT	SAVINGS		
East US	Standard DS1 v2	cht-ea	Engineering1, Free Trial(Converted to EA) *, Redshirt-azure, Scott's Subscription *	1	2	\$540	99%	\$127	\$69	↓ \$57	\$1,519	\$833	↓ \$687
				1	2	\$540	99%	\$127	\$69	↓ \$57	\$1,519	\$833	↓ \$687
East US 2	Standard F1s	cht-ea	Engineering1, Free Trial(Converted to EA) *, Redshirt-azure, Scott's Subscription *	0	1	\$325	100%	\$71	\$27	↓ \$44	\$849	\$325	↓ \$524
				0	1	\$325	100%	\$71	\$27	↓ \$44	\$849	\$325	↓ \$524
West India	Standard DS2 v2	cht-ea	Engineering1, Free Trial(Converted to EA) *, Redshirt-azure, Scott's Subscription *	0	1	\$678	100%	\$123	\$57	↓ \$67	\$1,479	\$678	↓ \$801
				0	1	\$678	100%	\$123	\$57	↓ \$67	\$1,479	\$678	↓ \$801
Total				1	4	\$1,543	99%	\$321	\$153	↓ \$168	\$3,848	\$1,836	↓ \$2,012

Conditions for Azure VM Reservations:


- Choose Time Period to Analyze = Last 30 days
- Preferred Time = 1 year but 3 year RIs give you larger discounts if your organization's forecasted usage is stable
- Target Coverage = VAST View recommends being at least 70% reserved. While 100% doesn't offer a lot of flexibility for changing your infrastructure, it provides the largest savings across your VM spend and VAST View will recommend an optimal purchase up to 100% coverage.
- Budget = Max budget gives optimal RI recommendations, but adjust this to the amount of upfront capital your business can provide.
- Estimated Purchase Date = Setting this in the future will take into account RIs that expire between now and that date. Set this to the date that you would like to purchase for to ensure you're including all active RIs at the time.
- Reservation Scope = The scope determines how many subscriptions can leverage the reservation benefit and how the reserved VM instance is applied to those subscriptions. Shared Subscription scope is recommended for better RI utilization and should be used if you want reservations to be able to "float" across your entire Azure account. Single Subscription scope should be used if you want to reserved compute capacity for the selected subscription in the filters.

Managing Usage

Monitoring Compliant Cloud Services and Regions

Implementing and enforcing a governance strategy prevents unexpected expenses due to choosing more expensive regions, or using cloud services that are not allowed by the business. Viewing spend by region across all accounts is a quick check for the business to see where outlying services may be used. This should be subscribed to on a weekly basis.

The following **Cost History** report is looking at the Project Perspective for the current month by region. Creating this report by Project view allows the business a visual check based on what regions projects should be built in.




Conditions for Cloud Service Compliance Report:

- Interval = Monthly
- X-Axis = Perspective of Business's choice
- Y-Axis = Cost
- Category = Service Items
- Filters = Month filtered to Current Month

Similar analysis should be done for any Azure and Data Center accounts using the **Azure Cost History** report and **Data Center Cost History** report.

Additionally, monitoring the cloud services used by teams helps the business plan for future spend and service item growth. When services are used outside of what the business has planned, unexpected spend becomes difficult to manage and predict.



Conditions for Cloud Service Compliance Report:


- Interval = Monthly
- X-Axis = Perspective of Business’s choice
- Y-Axis = Cost
- Category = Service Items
- Filters = Month filtered to Current Month

Governing Approved Instance Types

Spend doesn't always tell the full story of what's happening in an environment, microservice, or cost center. By monitoring usage within Perspective groups and by Service Items, it can be determined if incorrect instance types are being used, if particular services related to a project or customer have changed abnormally, or if particular environments are growing while other environments are offsetting those changes by decreasing, therefore leading to overall spend appearing steady. Use the following usage reports both daily and monthly to identify short and long-term trends and changes. These reports are individual to each service and should be monitored separately.

To monitor approved EC2 instance types, an **EC2 Instance Usage** report shows you the daily trend of the quantity of each instance type used. Additional views can be created by filtering to certain Perspective groups to audit what teams are using, or filter to specific instance types that aren't allowed by the business and categorize on Perspectives to determine who is using incorrect instance types.



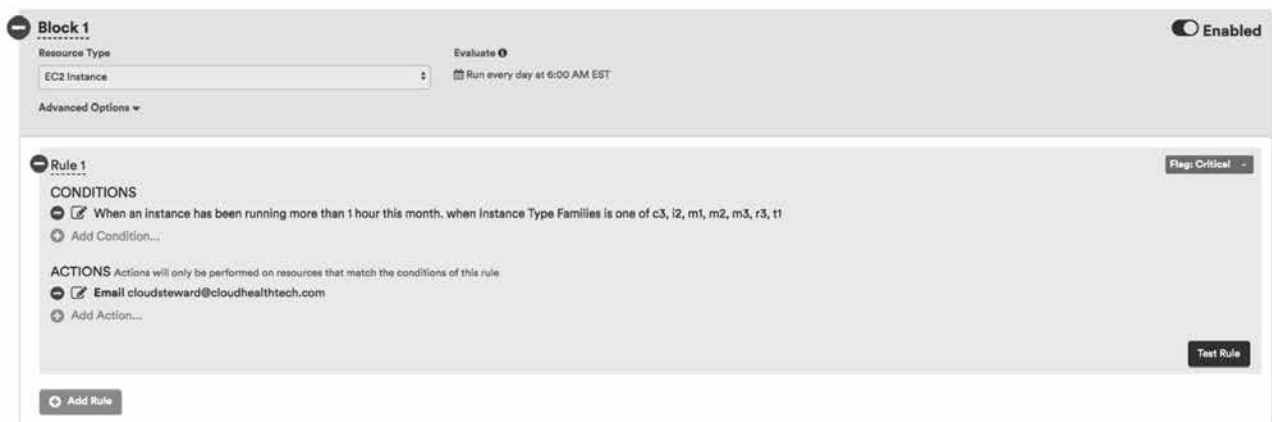


Conditions for EC2 Instance Compliance Report:

- Interval = Daily (Repeat conditions for Monthly intervals for overall trends as well)
- X-Axis = Days
- Y-Axis = # Instances (Note that this is a calculation of the number of instance hours used/number of instance hours available for the time period)
- Category = Instance Types
- Filters = Filter out last day from the Days filter as the last day's data may be incomplete. Another filter should be added for Perspective groups if sending directly to Perspective owners.

Use **Azure VM Usage** report to track Azure Virtual Machine types.


In addition to tracking variations through reporting, Policies provide alerts to teams based on conditions set to monitor usage for specific instance types or families. The Policy below should be run daily to alert users to stop instances before they incur a large amount of spend. This is best used for old generation instance types that should no longer be considered when spinning up new infrastructure.



Conditions for Instance Type Compliance Policy:

- Resource type = EC2 Instance
- Evaluate interval = Daily
- Conditions = When an instance has been running for more than 1 hour this month, when Instance Type Families is not one of the approved instance types
- Actions = Email Cloud Steward

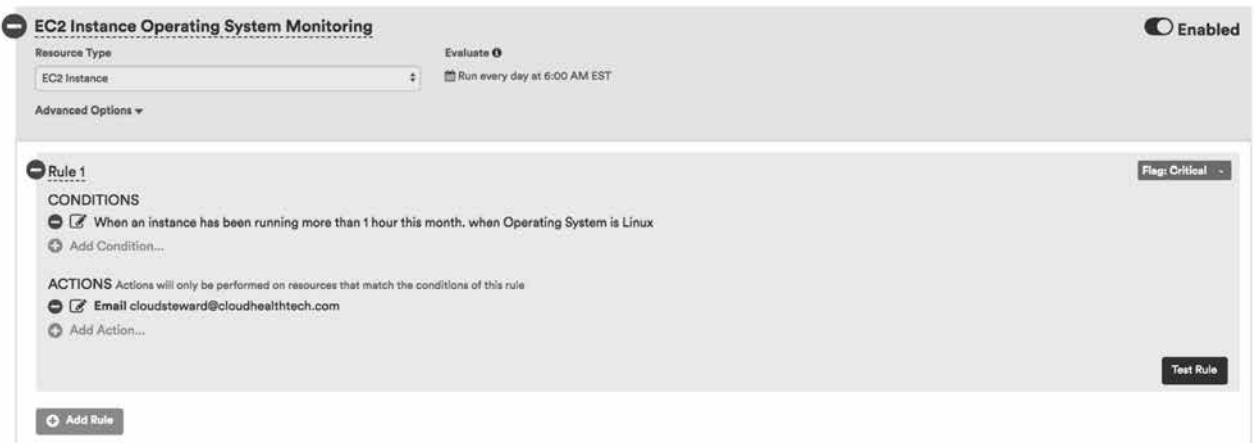
It is important to monitor the Operating System of each instance or VM, since the price per server can change depending on the OS or license used. Similar to the instance usage reports and Policies discussed, you can change the category to view by OS instead of by instance type. Viewing this report on a daily basis enables the business to immediately catch and correct instances running the wrong OS.



Conditions for Operating System Compliance Report:

- Interval = Daily
- X-Axis = Days
- Y-Axis = # Instances (Note that this is a calculation of the number of instance hours used/number of instance hours available for the time period)
- Category = Operating System
- Filters = Filter out last day from the Days filter as the last day's data may be incomplete. Another filter should be added for Perspective groups if sending directly to Perspective owners.

A Policy is best for being notified of these types of errors quickly. The Policy shown monitors instance usage on an hourly basis and notifies the Cloud Steward if any instances are running a non-approved operating system.



Conditions for Operating System Compliance Policy:

- Resource type = EC2 Instance
- Evaluate interval = Daily
- Conditions = When an instance has been running for more than 1 hour this month, when Instance Type Families is not one of the approved instance types
- Actions = Email Cloud Steward


Monitoring Unused Infrastructure

It's important to review the top 5 services based on spend within your accounts to determine which services will have the largest impact when optimized. The **Cost Summary Pulse** report is helpful for showing the top Perspective groups contributing to spend, as well as the top services being used. Common underutilized resources among customers tend to involve EC2 instances, RDS instances, EBS Volumes, and S3 Buckets.

The following are a few ways to optimize your environment based on commonly underutilized service items.

Managing Unattached Volumes

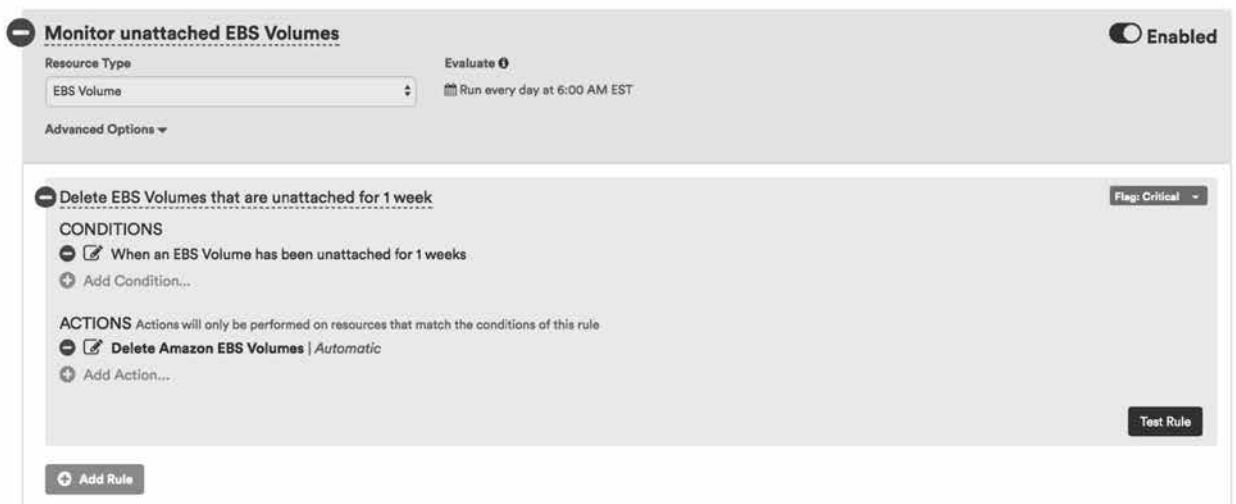
Monitoring EBS Volume usage is important to ensure that volumes are being deleted when the EC2 Instance is deleted. When “Available” volumes in the **EBS Volume** report continues to grow, this is an indication that the volumes have been unnecessarily left and should be deleted. It’s recommended to track this usage daily so that volume spend doesn’t increase due to unused volumes.



Conditions for EBS Volume Status Report:

- Interval = Daily
- X-Axis = Days
- Y-Axis = Storage
- Category = EBS Volume Status
- Filters = Filter out last day from the Days filter as the last day’s data may be incomplete. Another filter should be added for Perspective groups if sending directly to Perspective owners.

The EBS Volume report is helpful for tracking the trend of unattached volumes, but a policy should also be created to automatically delete these volumes after 7 days. This gives engineers enough time to reattach the volumes if they need to be reused, and deletes any volumes that have not been reattached. Instructions for fully automating, or setting up approval processes for deletion, can be found here.



The screenshot shows the AWS IAM console interface for creating a policy. The policy name is "Monitor unattached EBS Volumes" and it is currently "Enabled". The resource type is set to "EBS Volume". The evaluation is set to "Run every day at 6:00 AM EST". The policy is configured with the following conditions and actions:

- CONDITIONS:**
 - When an EBS Volume has been unattached for 1 weeks
 - Add Condition...
- ACTIONS:**
 - Delete Amazon EBS Volumes | Automatic
 - Add Action...

Buttons for "Add Rule", "Test Rule", and "Flag: Critical" are also visible.

Conditions for Unattached EBS Volume Policy:

- Resource type = EBS Volume
- Evaluate Interval = Daily
- Conditions = When EBS Volumes have been unattached for more than 1 week (This can be filtered to particular environments or accounts that are safer to delete volumes in, while excluding those that volumes should be looked at before deleting)
- Actions = Delete Amazon EBS Volumes Automatically, or send email to Perspective Group Owner if filtering by a Perspective

Managing Aging EBS Snapshots

EBS Snapshots are a point-in-time view of an EBS volume, typically used for backups or baselines for new volumes. In most cases, these snapshots don't need to stay active longer than 6 months, and, once a new snapshot has been taken, the old snapshot should be deleted. Monitoring these snapshots through Policies takes the manual work out of determining how long snapshots have been active and can have actions taken to delete these snapshots automatically.



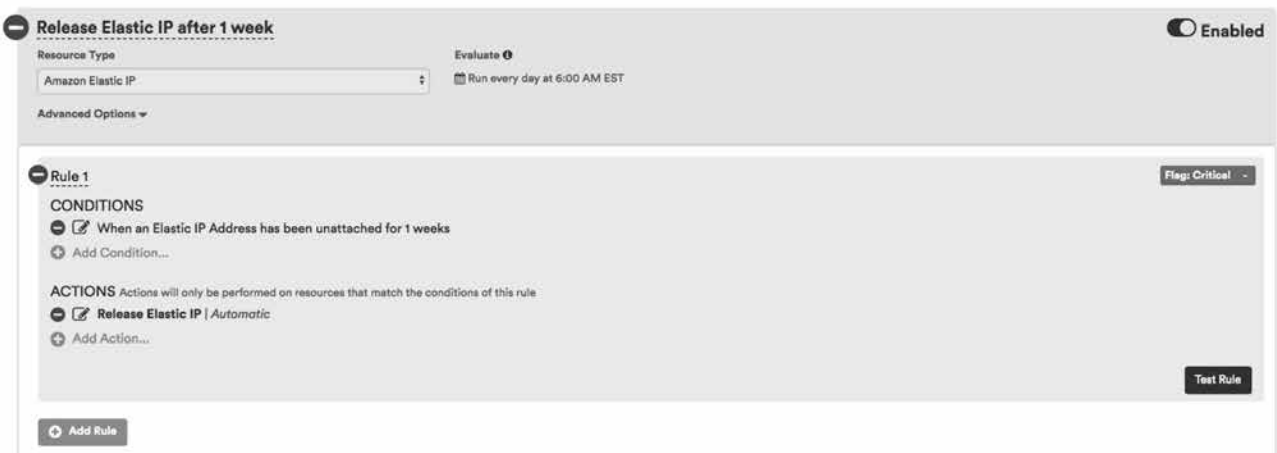
Conditions for Aging EC2 Snapshot Policy:

- Resource type = EC2 Snapshot
- Evaluate Interval = Daily
- Conditions = When EBS Snapshot is older than 6 months. This should be set to a shorter aging policy for environments or accounts that take daily snapshots in order to reduce spend as much as possible.
- Actions = Delete Amazon Snapshots Automatically, or send email to Perspective Group Owner if filtering by a Perspective

You can also review a tabular report to determine when a snapshot was first created. The report shows all snapshots active across your accounts. It can be sorted, or exported, to analyze the dates further. From the **Snapshots Assets** report, you can delete snapshots individually, or in bulk, if permissions from the Administrative section above were enabled for snapshot deletion. Use the Edit Columns button at the top to add the Create Date column. This identifies when the snapshot was created. You can add other columns should to help \the business identify the snapshot's purpose.

Managing Unattached Elastic IPs

AWS only charges for Elastic IPs (EIPs) when they are no longer attached to an instance. If EIPs are whitelisted by your business, the following policy should not be considered. If EIPs are not whitelisted by your business, the following policy should run on a daily basis to clean up any unattached EIPs.



Conditions for Unattached Amazon Elastic IP Policy:

- Resource type = Amazon Elastic IP
- Evaluate Interval = Daily
- Conditions = When an Elastic IP Address has been unattached for 1 week
- Actions = Release Elastic IPs Automatically, or send email to Perspective Group Owner if filtering by a Perspective

Managing Underutilized Elastic Load Balancers (ELBs)

Load Balancers are a critical component of web architecture, but often show an asymmetric skewed distribution of usage with well over 50% of load balancers doing little to no work. The **Load Balancer Metrics** report provides many column options to analyze various metrics when determining whether a load balancer is efficient or not. You should perform the following analysis monthly to determine which load balancers should be terminated.

Name	Avg Healthy Hosts	Max Healthy Hosts	Sum Requests
api-vpc	6.82	7.0	120352866.0
cht-api-vpc	14.87	15.0	426110.0
webserver-app-vpc	9.92	10.0	906185.0

Conditions for Load Balancer Metrics Report:

- Filters - Range Filters = Monthly, Choose previous complete month
- Columns = Name, Avg Healthy Hosts, Max Healthy Hosts, Sum Requests (By default, Month is included as the last column)

To analyze these 3 columns, sort the Sum Requests column from smallest to greatest to identify those with very few, or possible no, requests. Those with no, or very few, requests are good candidates to terminate. Similarly, you can sort the Max Healthy Hosts column to identify those with 1 or no hosts. With only 1 or no hosts, there is nothing to balance and, therefore, a load balancer isn't needed.

Optimizing Underutilized EBS Volumes

There are three important metrics when analyzing whether an EBS volume is utilized efficiently or not.

- **The speed, or latency of the disk, determines how long it takes the volume to respond to a query.**
- **IOPS measures how many requests per second are made to the volume.**
- **Throughput determines how quickly large files can be read from, or written to, a volume.**

With these metrics, you can focus on matching the optimal EBS volume based on the latency, IOPS, and throughput required. EBS volumes are charged per GB per month. It's worth noting that io1 is more expensive, both in the cost per GB, and in the cost of adding additional I/O, with a slightly higher SLA for sustained I/O. Gp2 has the the same seek time as io1, and you'll notice that gp2 is four times the cost of sc1 per GB. These nuances are vital to understand how to optimize the EBS spend and which EBS volume type is the correct one.

Configuring equivalent baseline IOPS in gp2

In most cases it's not economical to use io1 unless a volume has absolutely critical I/O needs, or has sustained volume above 10K IOPS. Gp2 can deliver the same IOPS for a lot less. GP2 provisions baseline IOPS at a rate of 3 per GB up to a limit of 10K. To figure out the baseline IOPS for a gp2 volume, simply multiply the GB by 3. For example, 400 GB gp2 will deliver a baseline of 1200 IOPS volume at much lower cost than a 400 GB io1 drive with 1000 PIOPS. The volume will be able to burst to higher levels until it's 'Burst Balance' is exhausted and the performance reverts to baseline, but it will provide the baseline performance for as long as needed.

Focus first on the Consumed Read/Write Ops columns to identify candidates to move to gp2 volumes. This combines reads and writes, and compensates for larger reads and writes. I/O operations that are larger than 256K are counted in 256K capacity units. For example, a 1024K I/O would count as 4 consumed IOPS.

Finding candidates for migrating to magnetic volumes

Volumes with high throughput are good candidates for st1, which has a high sustained baseline MB/s of I/O, with the ability to burst above baseline. To evaluate a transition to sc1 or st1, it is recommended to run an **EBS Volume Metrics** reports for a particular month, and filter by average I/O levels.

Range Filters

Report Period
Monthly ▾

Month
2018-02-01 ▾

Smart Filters

Account
All ▾

Availability Zone
All ▾

Current
Active ▾

Status
All ▾

Type
All ▾

Utilization
 All
 Low - less than 10 IOPS
 Medium - 10-85 IOPS
 High - more than 85 IOPS
 Health Check ▾

Tag Keys
All ▾

Tag Values
▾

Account Name	Volume Name	Instance Id	Type	Size (GB)	PIOPS	Avg Read Ops (/s)	Max Read Ops (/s)	Avg Write Ops (/s)	Max Write Ops (/s) ▾
CloudHealth Mgmt			standard	100 GB	4.7	275.1	61.4	3,134.8	
AWS CHT Sandbox			gp2	60 GB	9.0	1,239.0	31.8	3,044.1	
AWS CHT Sandbox			gp2	60 GB	5.3	1,640.5	17.9	3,004.9	

Conditions for EBS Volume Metrics Report:

- Filters - Range Filters = Monthly, Choose previous complete month
- Filters - Smart Filters = Utilization = Medium – 10-85 IOPS
- Columns = Volume Name, Volume ID, Instance ID, Type, Size (GB), PIOPS, Avg Read Ops (/s), Max Read Ops (/s), Avg Write Ops (/s), Max Write Ops (/s) - (Month will be included by default as the last column)

Finding candidates to migrate to st1

When looking for good candidates to migrate to s1, start by sorting volumes by max read and write MB/H to find your high throughput volumes. Sc1 volumes can sustain 250 mb/s, or 90,000 MB/h. St1 adds headroom to 500 MB/S baseline with a burst capability as described here.

Account Name	Volume Name	Volume Id	Instance Id	Type	Size (GB)	PIOPS	Avg Read (MB/h)	Max Read (MB/h)	Avg Write (MB/h)	Max Write (MB/h)
AWS CHT Sandbox				gp2	48 GB	6,980.2	121,786.1	4,925.6	95,504.8	
AWS CHT Sandbox				gp2	48 GB	10,209.1	108,018.9	8,823.7	87,639.3	
AWS CHT Sandbox				gp2	48 GB	10,330.7	103,523.3	8,332.6	77,954.1	

Conditions for EBS Volume Metrics Report:

- Filters - Range Filters = Monthly, Choose previous complete month
- Columns = Volume Name, Volume ID, Instance ID, Type, Size (GB), PIOPS, Avg Read (MB/h), Max Read (MB/h), Avg Write (MB/h), Max Write (MB/h) -(Month will be included by default as the last column)

Identifying Idle EBS Volumes

To find idle EBS volumes, run the same EBS volume metrics for the previous month, and add columns for sum read ops and sum write ops. Sort by the sum of read or write ops and the idle volumes will sort to the top. Those volumes with no read or write ops are candidates for termination.

Account Name	Volume Name	Volume Id	Instance Id	Type	Size (GB)	PIOPS	Sum Read Ops	Sum Write Ops	Avg Read (MB/h)	Avg Write (MB/h)
AWS CHT Sandbox				gp2	22 GB	0.0	0.0	0.0	0.0	
AWS CHT Sandbox				gp2	50 GB	0.0	0.0	0.0	0.0	
AWS CHT Sandbox				gp2	50 GB	0.0	0.0	0.0	0.0	

Conditions for EBS Volume Metrics Report:

- Filters - Range Filters = Monthly, Choose previous complete month
- Columns = Volume Name, Volume ID, Instance ID, Type, Size (GB), PIOPS, Sum Read Ops, Sum Write Ops, Avg Read (MB/h), Avg Write (MB/h) -(Month will be included by default as the last column)

Rightsizing Underutilized RDS Instances

Amazon's Relational Database Service (RDS) is a cost-efficient service for automating tasks, like hardware provisioning and backups. **The RDS Instance Metrics** report provides multiple metrics column options to analyze when determining the efficiency of an RDS instance. Metrics should be collected when a new instance is first spun up and running a typical workload. Various time intervals should be collected, from one hour, to one month, with average, maximum, and minimum values during these various time intervals. These metrics help determine a baseline for ongoing monitoring of performance. Analyze the following metrics monthly to identify underutilized RDS instances.

API Name	Zone Name	DB Name	Min Freeable Memory	Max Freeable Memory	Max CPU Utilization
db.r4.2xlarge	us-east-1e	cloudpercept	15137390592	22347112448	100.0
db.t2.xlarge	us-east-1d	cloudpercept	14985752576	15121809408	0.75

Conditions for RDS Metrics Report:

- Filters - Range Filters = Monthly, Choose previous complete month
- Columns = API Name, Zone Name, DB Name, Min Freeable Memory, Max Freeable Memory, and Max CPU Utilization (Month will be included by default as the last column)

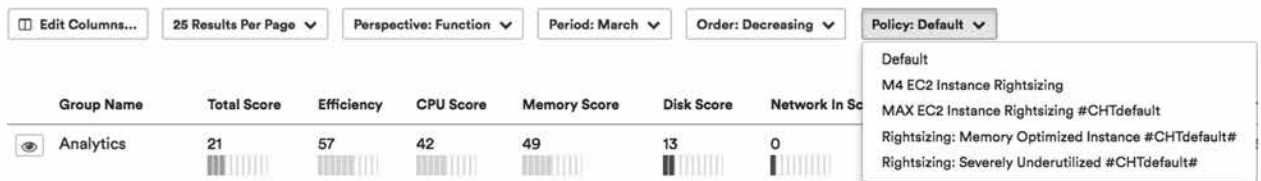
When analyzing these columns, there may be use cases that require additional columns to be added. One way to determine if an RDS instance is underutilized is if Min Freeable Memory is more than 50% of Max Freeable Memory. It's recommended to compare the Max Freeable Memory in the report to the baselines gathered previously. If there are no baselines set, Max Freeable Memory as a percent of total memory should remain between 50% and 75%. Any instance consistently going over 75% of the allocated memory may indicate that the instance is close to being undersized.

Similarly, compare the Max CPU utilization to the baselines gathered at the launch of the instance. Low CPU utilization may be a sign of an oversized instance but should be analyzed in conjunction with memory to determine if the next instance size would be adequate for the workloads running on the instance.

Managing Performance with Rightsizing

Monitoring tools, like Datadog and New Relic, are great resources for tracking and alerting in real-time when performance spikes or drops suddenly. VAST View integrates with these tools to track performance over a period of time so that trends can be analyzed for future capacity planning. For example, smaller instances can be provisioned in the beginning to reduce wasted cost and effort required by teams to rightsize in the future. If certain instances are consistently under or overutilized for particular projects or tasks, VAST View can identify these and provide recommendations for downsizing or terminating instances.

In addition to the default recommendations, **custom rightsizing policies** can be created by instance family or function to determine more accurate thresholds and weights for CPU, Memory, Disk and Network analysis. When changing the percentage thresholds for underutilized instances, the scoring and battery meters will change accordingly. However, customizing the thresholds from average metrics to maximum or minimum metrics will affect the recommendations provided. Any custom policies will be available as options in the Policy dropdown at the top of the rightsizing reporting.



Group Name	Total Score	Efficiency	CPU Score	Memory Score	Disk Score	Network In Score
Analytics	21	57	42	49	13	0

Rightsizing AWS EC2 Instances

Using Perspectives to identify certain environments or microservices enables the business to focus on the groups that can be rightsized to smaller instance types for underutilized EC2 instances. VAST View **EC2 Perspective Rightsizing** report analyzes CPU, Memory, Disk and Network In/Out for each EC2 instance and aggregates this information per Perspective group. More information on how these reports are calculated can be found [here](#).

Group Name	Total Score	Efficiency	CPU Score	Memory Score	Disk Score	Network In Score	Network Out Score	# of Instances	Projected Cost	Recommendation Savings
Production	21	55	25	63	13	0	0	964	\$121,529.48	\$28,154.12
Management	12	45	9	39	12	0	0	9	\$1,697.53	\$789.86
Staging	9	23	1	35	12	0	0	1	\$382.03	\$194.18
Platform NG	46	65	50	87	7	0	0	48	\$9,716.92	\$0.00

Once you determine the Groups that are candidates to be rightsized, click on the eye icon to see a filtered view of the EC2 instances that roll under that Group. Subscribe the correlating Cost Center Owners and System Admins to this filtered view on a weekly basis to identify opportunities for additional optimizations. It's important to discuss the recommendations internally while incorporating business context behind the purpose of certain instances, focusing on those instances that have the largest savings available. Those instances with a red battery meter are considered underutilized, while green is considered well-utilized.

Total Score	Efficiency	CPU Score	Memory Score	Disk Score	Network In Score	Network Out Score	Zone Name	Instance Name	Recommendation	Projected Cost	Recommendation Savings	Actions
22	26	40	61	9	<1	<1	us-east-1a	selenium-grid	Downgrade to m3.xlarge	\$382.15	\$194.18	Actions
11	28	5	24	27	<1	<1	us-east-1a	jenkins-master	Downgrade to c3.xlarge	\$302.38	\$153.30	Actions
2	31	<1	2	6	<1	<1	us-east-1a	mgmt-consul	Downgrade to r3.large	\$239.54	\$121.91	Actions
2	31	<1	5	6	<1	<1	us-east-1b	mgmt-consul	Downgrade to r3.large	\$239.42	\$121.91	Actions
2	31	<1	2	6	<1	<1	us-east-1e	mgmt-consul	Downgrade to r3.large	\$239.42	\$121.91	Actions
20	42	38	42	18	<1	<1	us-east-1a	sensu-prod	Downgrade to c3.large	\$172.45	\$76.65	Actions
19	68	4	75	16	<1	<1	us-east-1a	pritu1-1	No recommended change	\$53.84	\$0.00	Actions
19	77	2	75	15	<1	<1	us-east-1b	pritu1-2	No recommended change	\$34.16	\$0.00	Actions
19	77	2	75	15	<1	<1	us-east-1d	pritu1-3	No recommended change	\$34.17	\$0.00	Actions

Rightsizing Azure VMs

As with rightsizing AWS instances in VAST View, Perspectives help identify certain environments or microservices that can be rightsized to smaller VM types for underutilized VMs in Azure. The **Virtual Machine Perspective Rightsizing** report analyzes CPU, Memory and Disk I/O for each VM and aggregates this information per Perspective group. More information on how recommendations are made can be found here.

Once underutilized VMs that are candidates to be resized have been identified, you can set up actions in VAST View to resize selected VMs in the report. Options include automatically resizing the VM to the recommended size, or resizing the VM to a custom size using the custom actions setup previously.

Total Score	CPU Score	Memory Score	Disk Score	Vm Name	Size Name	Region Name	Projected Compute Cost	Recommendation Savings	Recommendation	Actions
25	4	46	No Data	demoecomm	Standard_A2	East US	\$40.19	\$20.12	Downgrade to Standard_A1	Actions -
13	3	30	7	cent1	Standard_DS1_v2	East US	\$46.29	\$7.87	D	<ul style="list-style-type: none"> Create JIRA Issue Stop & Deallocate Azure Virtual Machines Resize Azure Virtual Machines Restart Azure Virtual Machines Custom Resize Action
40	40	No Data	No Data	agenttest-win	Basic_A1	East US	\$0.00	\$0.00	N	

Managing Migrations from the Data Center to the Cloud

Organizations looking to migrate to the cloud face big decisions — such as which workloads they want to migrate, and what approach they should take. Primary considerations when starting are how to determine which workloads are suitable for migration and understanding the TCO of migrating those workloads to the cloud. A recent study from IDG found that the number one area where IT professionals needed help selling cloud to internal stakeholders was around TCO, which also happens to be the top investment driver for migrating to the cloud.

The VAST View Data Center product continuously monitors the performance and configuration of physical and virtual servers running in one or more data centers. The platform then analyzes this data and provides recommendations for moving a workload or a subset of infrastructure running the workload. Recommendations are made on asset types, region, reservations, and associated projected costs. This enables you to compare the TCO of running workloads on-premises with the public cloud, and make intelligent migration decisions. More information on configuring a migration assessment can be found [here](#).

Recommendations by Data Center

DATA CENTER	SERVER	CPU	MEM (GB)	DISK (GB)	CONFIGURATION	INSTANCE	STORAGE (GB)	ON DEMAND	UPFRONT	MONTHLY RI
CHT vSphere Boston		48	55.75	1037.79	\$12,241.97		1138	\$1,777.23	\$12,965.00	\$1,080.51
	cht-catalunya-linux-01	1	0.25	16.0	\$240.76	m3.medium	17	\$48.88	\$353.00	\$29.42
	CHT-Customer-Aggregator	2	2.0	16.0	\$405.43	m3.medium	17	\$48.88	\$353.00	\$29.42
	cht-daytona-linux-01	1	0.25	16.0	\$240.76	m3.medium	17	\$48.88	\$353.00	\$29.42

To create **custom pricing** for your data center servers for assessments and reporting, VAST View allows you to assign pricing based on configuration or tags to apply pricing to all similar server types. This provides visibility to the business to base spend on the server cost, building cost, heating and cooling, and more.

1. Select account(s). DATA CENTER ACCOUNT: CHT VSPHERE BOSTON

2. Price assets per: INTERVAL: DAY

3. Price assets by: Configuration Tag

4. Price each configuration.

Configuration Name:

Per Core: \$ 1.33

Per GB of Memory: \$ 2.33

Per GB of Storage: \$ 0.00

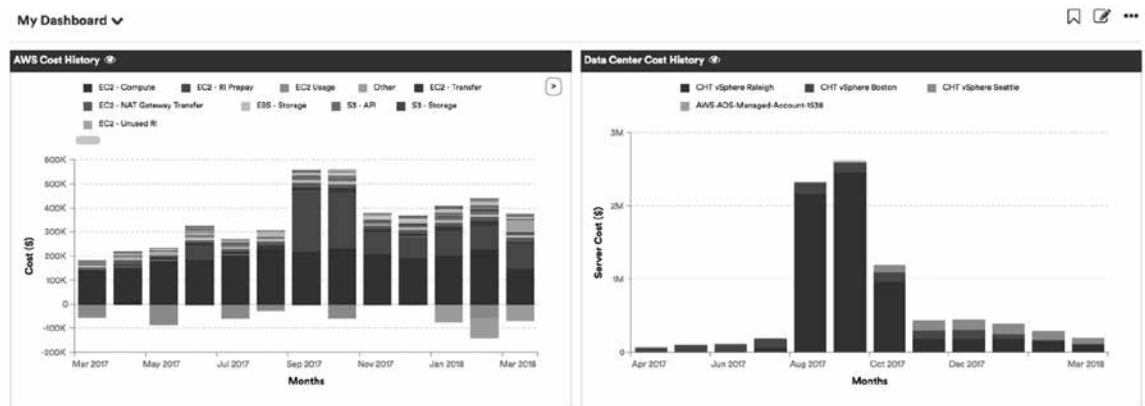
Windows License: \$ 4.50

Linux License: \$ 1.00

fixed: \$ 5.00

Calculating pricing will not affect existing pricing.

You can create a dashboard in VAST View to track the progress of your migration from the Data Center to the Cloud. For example, it can show spend in the Data Center going down, while the Cloud spend trends upwards as servers are migrated to new instances. Save the preferred view for **Cost History** reports for Data Center and any Cloud providers so that they can be added to **custom dashboards**. Drag and drop the graph widgets to your new dashboard to choose the saved Cost History reports.



Managing Security Risks

Monitoring security risks is important to identify risks early and remediating security issues quickly. A subscription to the **Security Recommendations** report should be set up to delivery to DevOps and the business’s Security team on a daily basis to monitor these risks. The underlying **Policies** that populate the recommendations for AWS Best Practice Policies and CIS Foundations Policies should be set to run hourly for important security risks to the business so that alerts can be sent directly to the DevOps teams and Security teams quickly.

Severity ^	Policy Block Name	Summary	Updated At
▲	Root Account API Access	1 Amazon Account has Root Account API Access	2017-11-03 05:27:50 UTC
▲	Root Account MFA Access	2 Amazon Accounts have Root Account MFA Disabled	2017-11-03 05:27:51 UTC
▲	CloudTrail Enabled	2 Amazon Accounts do not have CloudTrail enabled for all regions	2017-11-03 10:12:55 UTC
▲	CloudTrail CloudWatch Logs	8 Amazon CloudTrail Trails are not integrated with CloudWatch Logs	2017-11-03 10:12:55 UTC
▲	CloudTrail Bucket Access Logs	2 Amazon CloudTrail Trail's S3 buckets are not configured to create access logs	2017-11-03 10:12:57 UTC
▲	EC2 Security Group Ingress ICMP	18 Amazon Security Group rules have an open icmp ingress source range allowing all traffic	2017-11-03 05:27:59 UTC
▲	IAM User Password Policy	2 Amazon Accounts have a password policy that does not adhere to best practices	2017-11-03 05:27:59 UTC
▲	IAM User has access key(s) that needs rotation	109 Amazon Users have access keys that have not been rotated in 90 days	2017-11-03 05:27:53 UTC
▲	EC2 Security Group has large ingress port range	20 Amazon Security Group rules have a port range greater than 50	2017-11-03 05:27:56 UTC
▲	CloudTrail Log Encryption	11 Amazon CloudTrail Trails have logs that are not encrypted at rest	2017-11-03 10:12:57 UTC
▲	EC2 Security Group Egress Range	122 Amazon Security Group rules have an open egress source range allowing all traffic	2017-11-03 05:27:59 UTC
▲	IAM User with no Group	1 Amazon User has No assigned IAM Group	2017-11-03 05:27:51 UTC
▲	IAM User has access key(s) that are not being used	41 Amazon Users have access keys that have not been used in 60 days	2017-11-03 05:27:53 UTC
▲	IAM Role Changed	9 Amazon IAM Roles have changed or been created in the last 7 days	2017-11-03 05:27:56 UTC
▲	IAM Group Changed	1 Amazon Group has changed or been created in the last 7 days	2017-11-03 05:27:56 UTC
▲	IAM Policy Changed	26 Amazon IAM Policies have changed or been created in the last 7 days	2017-11-03 05:27:56 UTC
▲	IAM User Changed	2 Amazon Users have changed or been created in the last 7 days	2017-11-03 05:27:57 UTC
▲	IAM User Inactive	12 Amazon Users have not logged in for 60 days	2017-11-03 05:27:57 UTC
▲	EC2 Security Group Changed	14 Amazon Security Groups have changed or been created in the last 7 days	2017-11-03 05:27:57 UTC
⊙	CloudTrail Log Validation	No violations at the present time	2017-11-03 10:12:55 UTC
⊙	EC2 Security Group CIDR Rules	No violations at the present time	2017-11-03 05:27:57 UTC
⊙	EC2 Security Group Ingress 0.0.0.0/0	No violations at the present time	2017-11-03 05:27:58 UTC
⊙	RDS Security Group CIDR Rules	No violations at the present time	2017-11-03 10:12:55 UTC



Further customization of the Policies should be considered to align with your business’s requirements, such as password Policies or access key rotation Policies. For any security checks that are particularly sensitive to your business, actions should be added to the Policies, in addition to being subscribed to the overall report.

Edit Condition for Amazon Account

Choose a Topic: Security | Data Type: Per-resource | Choose a Measure: Password Policy

Build a Condition

When any Amazon Accounts have a password policy that do not adhere to best practices (Zeroes mean a check is disabled)

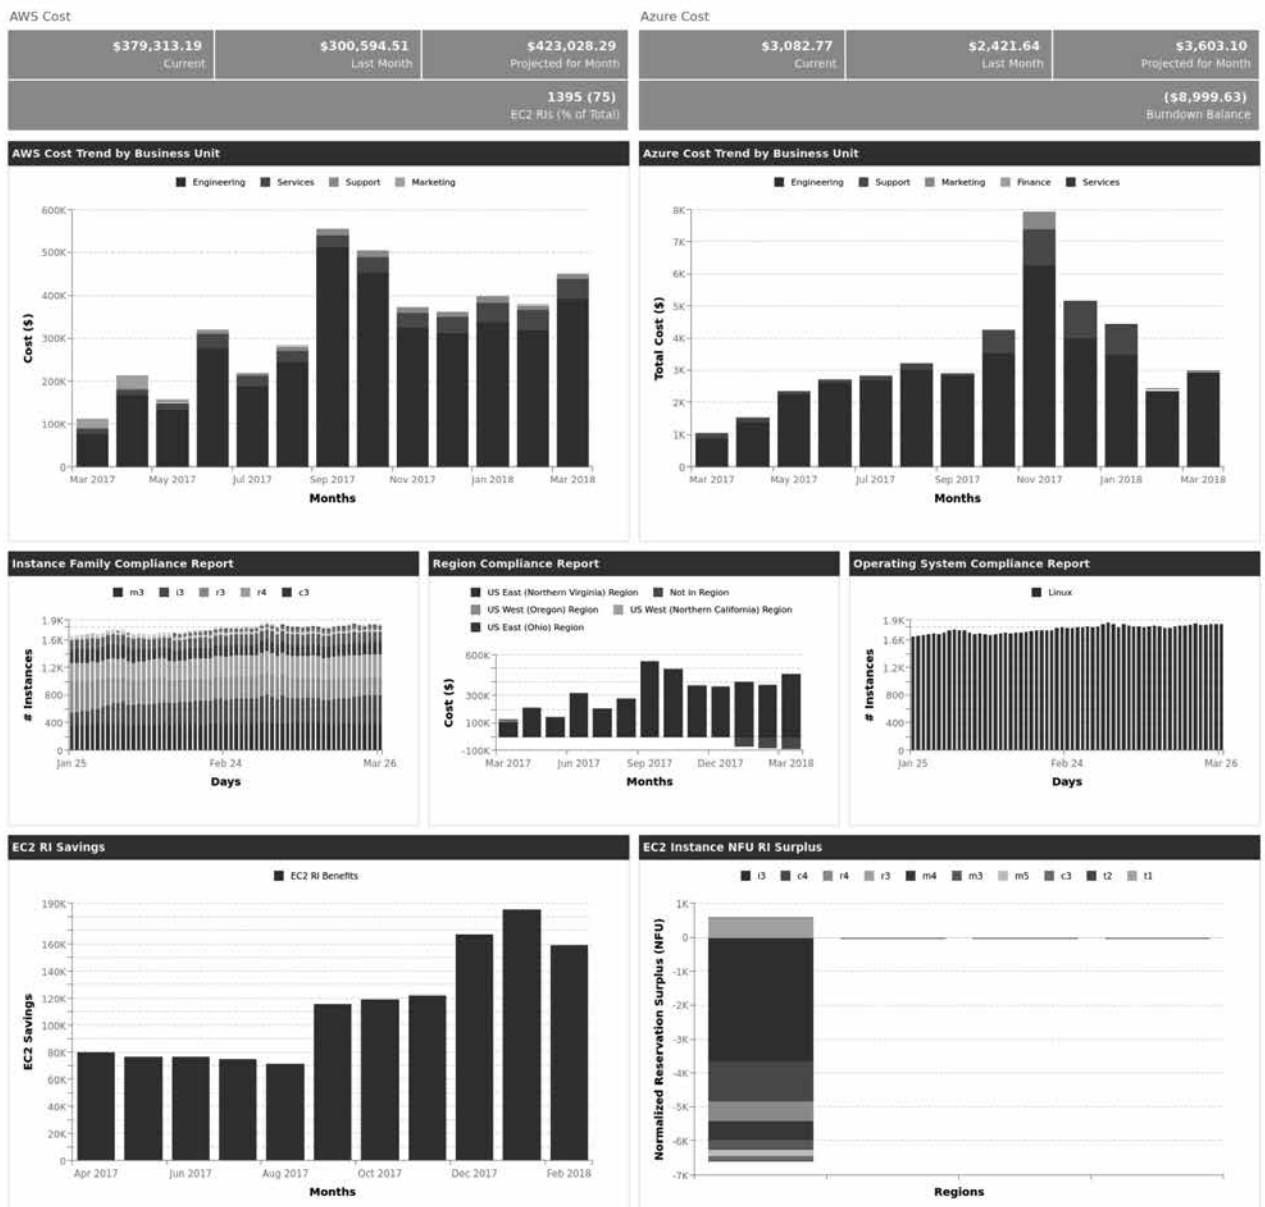
Minimum password length:	14 characters
Max password age:	90 days
Remember and prevent past passwords:	24 passwords
Allow users to change password:	Yes
Require uppercase characters:	Yes
Require lowercase characters:	Yes
Require numbers:	Yes
Require symbols:	Yes
Hard expire passwords:	No

Another aspect of securing your infrastructure is determining which instances have ports open to the public. The **Health Check Pulse** report lists EC2 instances with all ports open to the public, as well as instances with unencrypted public port exposure. The results are hyperlinked and will bring you to a filtered view based on those checks. It's recommended to subscribe to the Health Check Pulse report on a weekly basis to identify additional optimizations, as well as security risks.

76 Running Instances With All Ports Open to the Public	
TCP - 1 Instance	High
ICMP - 75 Instances	Low
390 Running Instances With Unencrypted Public Port Exposure [†]	
FTP (21) - 52 Instances	High

Executive Reporting Best Practices

Dashboards allow the Cloud Steward to create executive and cost center views to provide visibility to various teams. By grouping the best practice interactive charts seen in the above sections into one central dashboard, users can be provided these reports in one view when they log in, or be subscribed to it on a weekly or monthly basis. It's recommended to make a dashboard per role, and a dashboard per cost center, based on the descriptions and recommendations in the reporting sections above



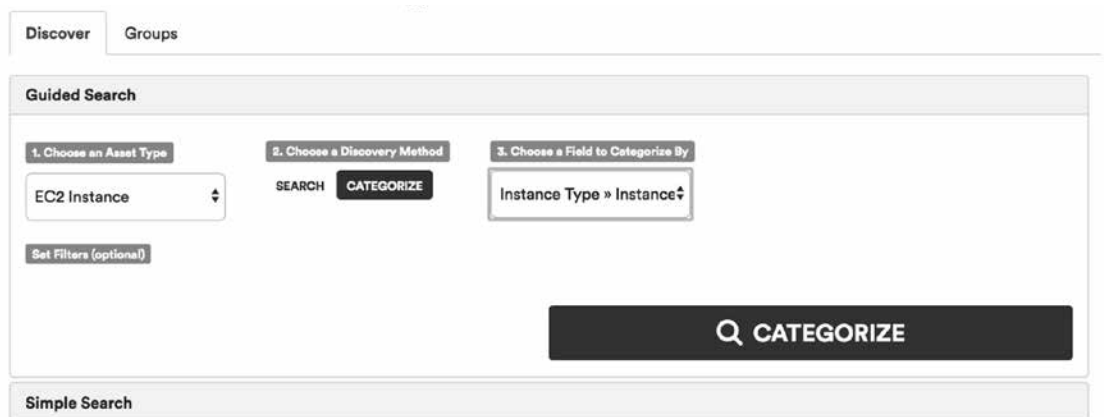
Key Performance Indicators and Metrics

Cost per EC2 Instance

Most customers find EC2 to be the largest component of their spend. Tracking this spend by instance hour allows the business to ensure they are moving in the right direction with reservations, while considering this by instance family to limit the number of variables impacting their spend. This requires looking at instance hours by family, as well as the spend per instance family. To create this analysis, you need to have a Perspective by instance family. Use the following instructions to set this up. Then combine the outputs from the following instance usage report and amortized cost report to mimic the following calculations.

Perspective by Instance Family

Set up an Instance Family Perspective using the asset type “EC2 Instance” and categorize by Instance Type > Instance Family. Add all results as a dynamic group to reduce maintenance when new types are used.




The screenshot shows the 'Guided Search' interface. At the top, there are tabs for 'Discover' and 'Groups'. Below this, the 'Guided Search' section contains three numbered steps:

- 1. Choose an Asset Type:** A dropdown menu is set to 'EC2 Instance'.
- 2. Choose a Discovery Method:** Two buttons are visible: 'SEARCH' and 'CATEGORIZE'.
- 3. Choose a Field to Categorize By:** A dropdown menu is set to 'Instance Type » Instance'.

Below these steps is a 'Set Filters (optional)' section. At the bottom right of the 'Guided Search' area is a large black button with a magnifying glass icon and the text 'CATEGORIZE'. Below the 'Guided Search' section is a 'Simple Search' section.

Instance Usage Report

Create an **instance usage** report with the following conditions below to look at the number of instance hours by instance family.




Instance Type Families	2017-10	2017-09	2017-08	2017-07	2017-06	2017-05	2017-04	2017-03	2017-02	2017-01	2016-12	2016-11	2016-10
r4	519.00	12,382.47	11,695.00	8,929.00	8,235.00	1,593.00	720.00	722.00					
r3	32,414.56	504,362.38	451,795.65	436,452.00	420,770.00	433,168.00	342,990.00	316,143.00	264,598.00	286,232.00	261,278.00	219,162.00	187,012.00
c3	8,325.00	130,857.00	134,804.90	139,882.00	124,214.00	127,214.00	123,763.00	118,530.00	101,240.00	109,322.00	105,255.00	101,871.00	99,892.00
t2	877.69	14,547.00	15,174.00	24,079.00	29,962.00	32,816.00	22,155.00	15,594.00	7,831.00	7,644.00	7,455.00	6,517.00	3,738.00
m4	2,575.00	34,560.95	36,388.01	14,326.00	5,607.00	1,701.00	216.00						22.00
i3	3,875.00	22,625.00	8,489.00	4,696.00	6,288.00	4,085.00	724.00	725.00					
m3	14,939.09	236,131.70	236,731.97	227,476.00	215,562.00	229,740.00	275,134.00	266,409.00	218,928.00	238,771.00	247,480.00	233,046.00	245,593.00
c4	635.61	8,022.00	8,409.00	7,355.00	5,506.00	5,548.00	1,620.00			3.00			
t1	47.00	721.00	341.00		519.00	2,232.00	2,160.00	2,232.00	2,016.00	2,233.00	2,278.00	2,160.00	1,752.00
m1	45.00	720.00	744.00	744.00	720.00	744.00	935.00	1,488.00	769.00	744.00	744.00	720.00	762.00
d2	45.00	720.00	744.00	456.00									
i2		72.00											
Total	64,299.94	965,731.50	905,316.53	864,395.00	817,383.00	838,841.00	770,417.00	721,843.00	595,382.00	644,949.00	624,490.00	563,476.00	538,771.00

Conditions for Instance Usage Report:

- Interval = Monthly
- X-Axis = Months
- Y-Axis = # Instances Hours
- Category = Instance Type Families

Amortized Cost Report

Create an **amortized cost** report by instance family using the conditions shown.



Conditions for Amortized Cost Report:

- Interval = Monthly
- X-Axis = Month
- Y-Axis = Cost, Amortized Cost
- Category = Instance Type Families Perspective
- Filters = Months -All complete months, AWS Services -All EC2 services under “Amazon Elastic Compute Cloud -Direct” (EC2 -Compute, EC2 -Transfer, EC2 -EBS Surcharge, EC2 -Load Balancer Transfer, EC2-Load Balancer Usage, EC2 -NAT Gateway Transfer, EC2 -NAT Gateway Usage)

Combining Exports for Analysis

Export both reports to combine in excel. The amortization and cost tabs in the amortized cost report should be summed into one value per month and instance family to then be compared to the number of instance hours running in that month per instance family. Shown is an example of what the end analysis would look like. A downward trend in spend per instance hour would indicate that reservations are being effectively used to drive cost down, while noting that many other attributes can affect the cost, like region, operating system and size of instance type.

Cost per Instance	2017-10	2017-09	2017-08	2017-07
r3	\$ 37.47	\$ 39.81	\$ 10.69	\$ 6.62
c3	\$ 1.02	\$ 1.06	\$ 1.08	\$ 0.94
m3	\$ 0.10	\$ 0.09	\$ 0.09	\$ 0.09
r4	\$ 1.01	\$ 0.40	\$ 0.36	\$ 0.27
m4	\$ 0.19	\$ 0.20	\$ 0.19	\$ 0.19
t2	\$ 0.03	\$ 0.03	\$ 0.03	\$ 0.03
r3	\$ 37.47	\$ 39.81	\$ 10.69	\$ 6.62

Underutilized EC2 Percentage

Trending how utilized EC2 instances are over time gives the business the information needed for future capacity planning, data for improving utilization in certain environments or microservices, and allows changes to be made to downsize instance types that will lead to decreasing EC2 instance spend over time.

Use the **EC2 Rightsizing** report to identify instances that are underutilized. Add columns to the report for visibility into their associated Perspective groups.

	Total Score	CPU Score	Memory Score	Disk Score	Network In Score	Network Out Score	Zone Name	Instance Id	API Name	Environments	Recommendation	Projected Cost	Recommendation Savings
	6	3	No Data	No Data	11	4	us-east-1e	i-0dbedc5d0a773166f	c4.xlarge	Production	Downgrade to c4.large	\$652.23	\$508.08
	6	3	No Data	No Data	12	4	us-east-1e	i-0ccta18ad53b4974f	c4.xlarge	Production	Downgrade to c4.large	\$652.23	\$508.08
	10	5	No Data	No Data	19	6	us-east-1e	i-08c82ee6960bf501e	c4.xlarge	Production	Downgrade to c4.large	\$652.46	\$507.57
	10	5	No Data	No Data	19	6	us-east-1e	i-06339d310e964a8b8	c4.xlarge	Production	Downgrade to c4.large	\$651.55	\$506.75
	6	3	No Data	No Data	11	4	us-east-1e	i-06d94d23ba06b40c6	c4.xlarge	Production	Downgrade to c4.large	\$644.73	\$499.94

Conditions for EC2 Rightsizing Report:

- Period = Prior Month
- Columns = Listed in the image above

Once exported, filter the report for any instances with a Total Score of less than 68. Then use the count formula in Excel to count the number of instances that have a score of less than 68. Use this in your calculation for the number of instances underutilized. VAST View determines this threshold using its **Scoring Mechanism**. You can use custom rightsizing Policies to change the score and you can customize the threshold to align with one that is approved by your business. Combine this underutilized instance count with the total number of instances ran in the previous month using an **EC2 Instance Usage** report.




Because the Rightsizing report displays two months of detailed data rather than a trend of underutilized instance counts, this calculation should be saved each month locally and combined with previous calculations to observe the trends and make adjustments to capacity planning as a result.

Conditions for Instance Usage Report:

- Interval = Monthly
- X-Axis = Month
- Y-Axis = # Instances
- Category = None
- Filters = Previous complete month

Reservation Type Breakdown

Another metric to analyze when utilizing reservations is the number of Reserved Units (NFUs) split between RIs, Spot and On Demand. The RI and Spot NFUs should account for at least 70% of the NFUs utilized in a month. Below is the **EC2 Instance Usage** report categorized by Reservation Type and should be analyzed on a monthly basis.




Conditions for Instance Usage Report:

- Interval = Monthly
- X-Axis = Months
- Y-Axis = Normalized Usage (NFU)
- Category = Reservation Type

Ratio of Prod vs. Non-Prod Resources


You should analyze the split of resources between Production and Non-Production environments on a weekly basis to ensure that the spend of Non-Production resources doesn't continue to outgrow Production at a much faster rate than Production spend grows. Non-Production spend should be less than 25% of total infrastructure spend. Using the **Cost History** report, a trend of the breakdown of environments can be exported to convert into percentages. If the environment detail is not necessary in analyzing your business's breakdown, a simpler Perspective can be created by merging the Non-Production environments into one group.



Conditions for Cost History Report:

- Interval = Weekly
- X-Axis = Weeks
- Y-Axis = Cost (\$)
- Category = Environment Perspective
- Filters = AWS Services = All services excluding prepay services

To track this breakdown on a monthly basis, you can select just one month in the Cost History report filters, and change the report to a Pie chart. This can be used for executive reporting purposes and should be compared to previous months' breakdowns.

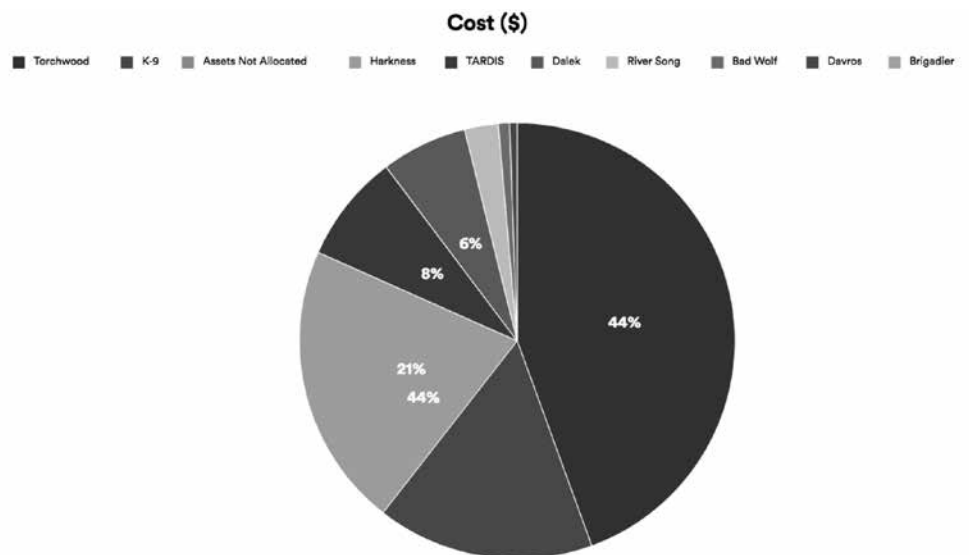


Conditions for Cost History Report:

- Interval = Monthly
- X-Axis = Environment Perspective
- Y-Axis = Cost (\$)
- Category = None
- Chart Type = Pie
- Filters = AWS Services = All services excluding prepay services, Month = Previous Month

Tagging Compliance

With the recommendations previously discussed to help with tagging compliance, it's important to track that the changes to tags are actually being made. Perspectives are a great way to track this as any infrastructure with incorrect tags will not be allocated to a Perspective group. To track the tagging compliance at a point in time, use a **Cost History** report to view one month's spend by a Perspective. Because a resource's historical spend will follow the resource to a Perspective group, a historical trended cost report will not accurately show the spend for untagged, or unallocated, resources in prior months. The percentage of resources in Assets not Allocated should be as low as possible, if not zero. If certain accounts don't need to follow a tagging structure for certain tag keys then they should be excluded from the report using filters to ensure only actionable resources are included in the analysis.



Conditions for Cost History Report:

- Interval = Monthly
- X-Axis = Project Perspective
- Y-Axis = Cost (\$)
- Category = None
- Chart Type = Pie
- Filters = Month = Previous Month, Accounts = Only Accounts included in Tagging Policy

Frequency of Reports and Policies

STAKEHOLDERS				
	Cloud Steward	Tech Ops/Sus Admin	Cost Center Owner	VP of IT/Operation
Monthly	<ul style="list-style-type: none"> Actual Spend by Group Budget vs Actual Spend Azure Burndown Trend RI Optimizer (EC2, RDS, VM) RI Savings EC2 NFU RI Usage Report Compliant Instances/VMs Underutilized ELBs Underutilized EBS Volumes Rightsizing RDS Instances Rightsizing EC2 Instances Rightsizing Azure VMs Data Center Migrations 	<ul style="list-style-type: none"> Compliant Instances/VMs Underutilized ELBs Underutilized EBS Volumes Rightsizing RDS Instances Rightsizing EC2 Instances Rightsizing Azure VMs Data Center Migrations 	<ul style="list-style-type: none"> Actual Spend by Group Budget vs Actual Spend RI Optimizer (EC2, RDS, VM) RI Savings EC2 NFU RI Usage Report Compliant Instances/VMs Underutilized ELBs Underutilized EBS Volumes Rightsizing RDS Instances Rightsizing EC2 Instances Rightsizing Azure VMs Data Center Migrations 	<ul style="list-style-type: none"> Actual Spend by Group Budget vs Actual Spend Azure Burndown Trend Data Center Migrations Health Check Pulse Report
Weekly	<ul style="list-style-type: none"> Convertible RI Analysis Compliant Regions Report Compliant Services Report Compliant OS Report 	<ul style="list-style-type: none"> Compliant Regions Report Compliant Services Report Compliant OS Report 	<ul style="list-style-type: none"> Convertible RI Analysis Compliant Regions Report Compliant Services Report Compliant OS Report 	<ul style="list-style-type: none"> Executive Dashboard
Daily	<ul style="list-style-type: none"> Resources Missing Tags Spend Increase Monitoring Projected Spend vs Budget Auto RI Modification Underutilized RI Monitoring Expiring RIs Daily Spend by Group Compliant Instances/VM EBS Volume Status Security Recommendations Health Check Pulse 	<ul style="list-style-type: none"> Resources Missing Tags Compliant Instances/VMS Compliant OS Unattached EBS Volume Aging Snapshots Unattached EIPs Compliant Instances/VMs EBS Volume Status Security Recommendations Health Check Pulse 	<ul style="list-style-type: none"> Spend Increase Monitoring Projected Spend vs Budget Compliant Instances/VMS Compliant OS Unattached EBS Volume Aging Snapshots Unattached EIPs Daily Spend by Group Compliant Instances/VMs EBS Volume Status Security Recommendations Health Check Pulse 	

- Policy
- Report

CONCLUSION

It's important to remember that these best practices are not meant to be one-time activities, but ongoing processes. Because of the dynamic and ever changing nature of the cloud, cost optimization activities should ideally take place continuously.

Learn more about how you can help you automate the continuous optimization by visiting www.vastITservices.com

POWERED BY **CloudHealth**

VAST View™ is a trademark of VAST IT Services.

