OpenCost

Kubernetes Cloud Cost Monitoring with OpenCost & Optimization Strategies

Matt Ray

Senior Community Manager for OpenCost

mattray@kubecost.com

SCaLE 20x

Who am I?

Matt Ray

- Senior Community Manager for OpenCost at Kubecost.
- Co-host of the Software Defined Talk podcast for 8+ years.
- Living in Sydney, Australia for 6+ years after relocating from Austin, Texas.
- Active in Open Source for much, much longer.
- mattray@kubecost.com
- https://www.linkedin.com/in/mhray/
- @mattray on GitHub, Mastodon, too many Slacks

The complexity of operating Kubernetes efficiently is real

Technical

COMPLEXITY

Higher-level abstractions, more shared resources, and increasingly dynamic

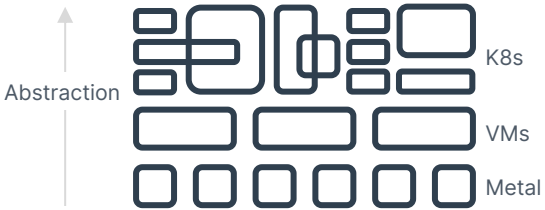

Behavioral

COMPLEXITY

Decentralized releases means any engineer can increase spend quickly

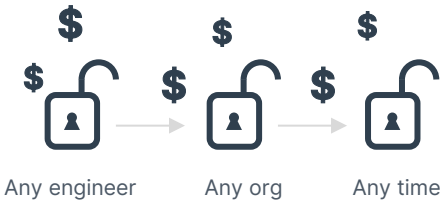

OpenCost

Open source Kubernetes cost monitoring

Specification and Implementation

Cloud Native Computing Foundation Sandbox Project

FinOps Certified Solution

- https://opencost.io
- https://github.com/opencost
- https://www.cncf.io/projects/opencost/

OpenCost Specification

Created by a community of Kubernetes practitioners

https://github.com/opencost/opencost/blob/develop/spec/opencost-specv01.md

- Adobe
- Armory
- AWS
- D2IQ
- Google Cloud
- Kubecost
- Mindcurv
- New Relic
- Pixie
- Red Hat
- SUSE

OpenCost Specification

https://github.com/opencost/opencost/blob/develop/spec/opencost-specv01.md

Kubernetes workloads are often transient and vary in the resources they consume.

How do we measure who is responsible for what and how much?

- Management fees
- Expenses from nodes
- Persistent volumes
- Attached disks
- Load balancers
- Network ingress/egress

Total Cluster Costs

`Total Cluster Costs = Cluster Asset Costs + Cluster Overhead Costs`

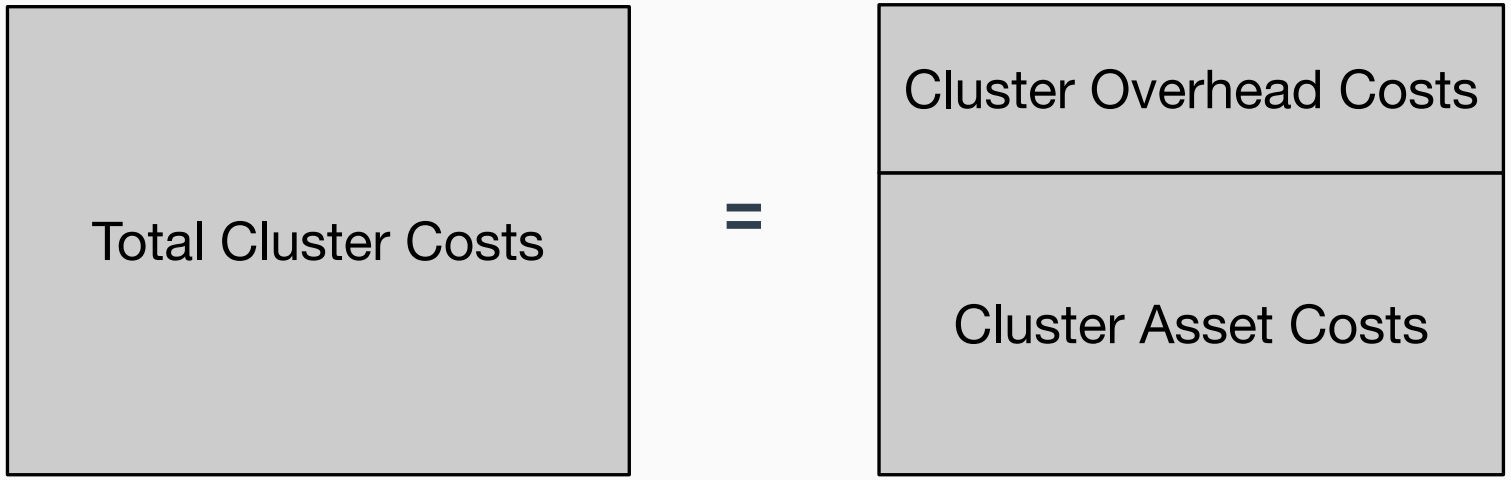

Cluster Asset Costs

`Cluster Asset Costs = Resource Allocation Costs + Resource Usage Costs`

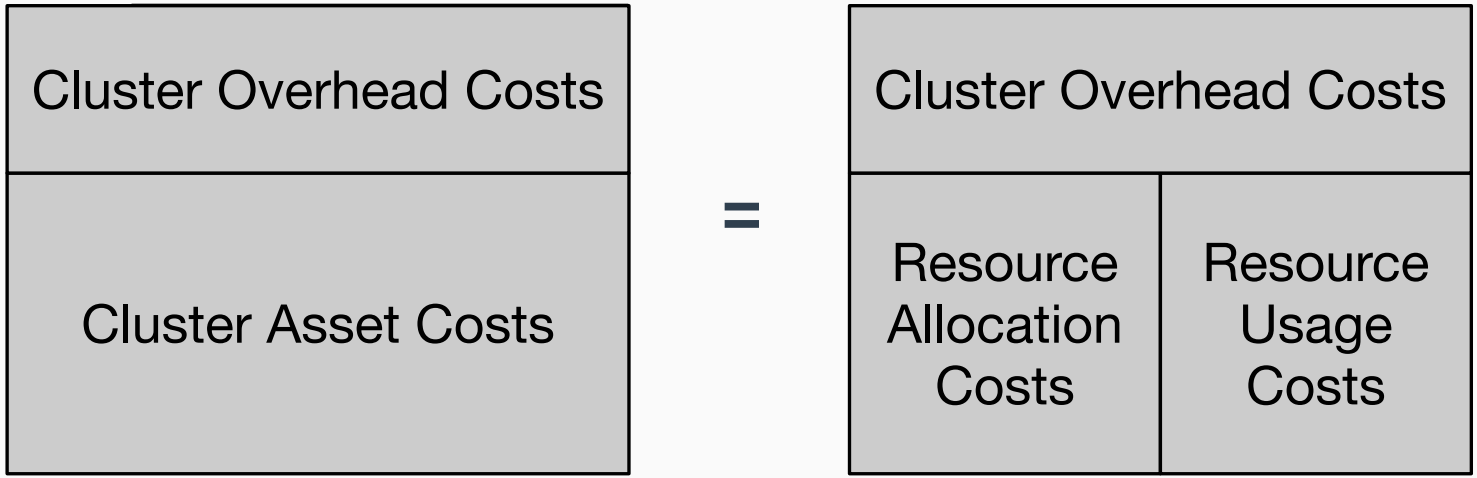

Cluster Asset Costs: Node Costs

`Cluster Asset Costs = Resource Allocation Costs + Resource Usage Costs`

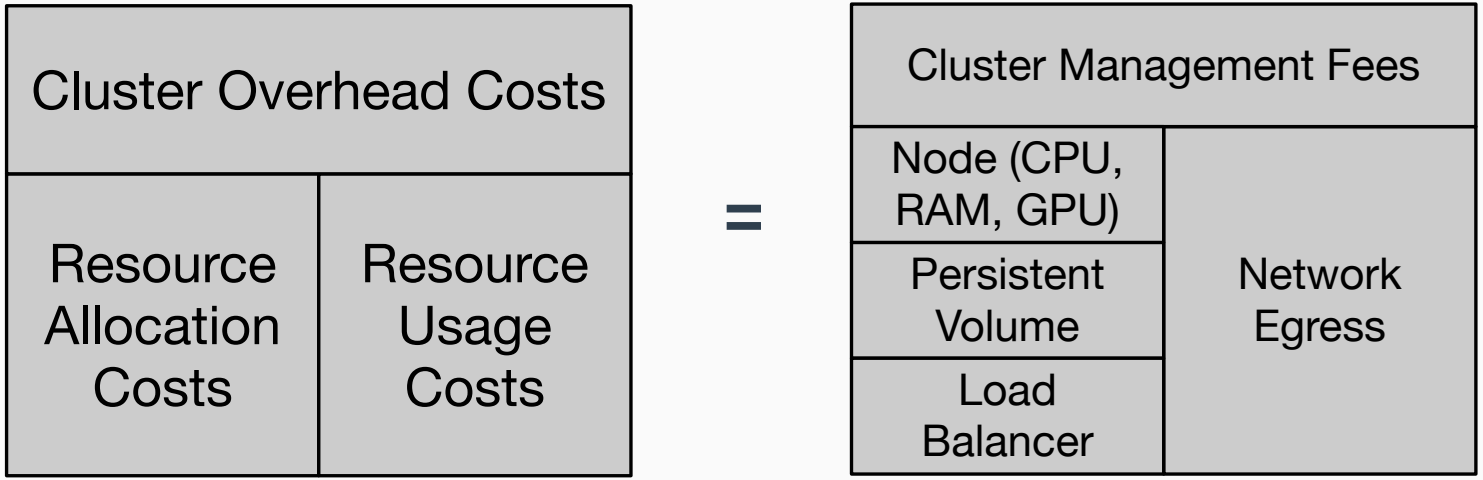

We’ve got the cost of our Kubernetes assets

Now let's distribute them across Workloads

Workload Costs

Inside the Kubernetes Cluster

What is Measured

- CPU
- Memory
- GPU
- Storage Volume
- Load Balancer

Aggregations

- Container
- Pod
- Deployment
- StatefulSet
- Job
- Controller Name
- Controller Kind
- Label
- Annotation
- Namespace
- Node
- Cluster

Workload Costs + Cluster Idle Costs

`Total Cluster Costs = Workloads + Cluster Idle Costs + Cluster Overhead Costs`

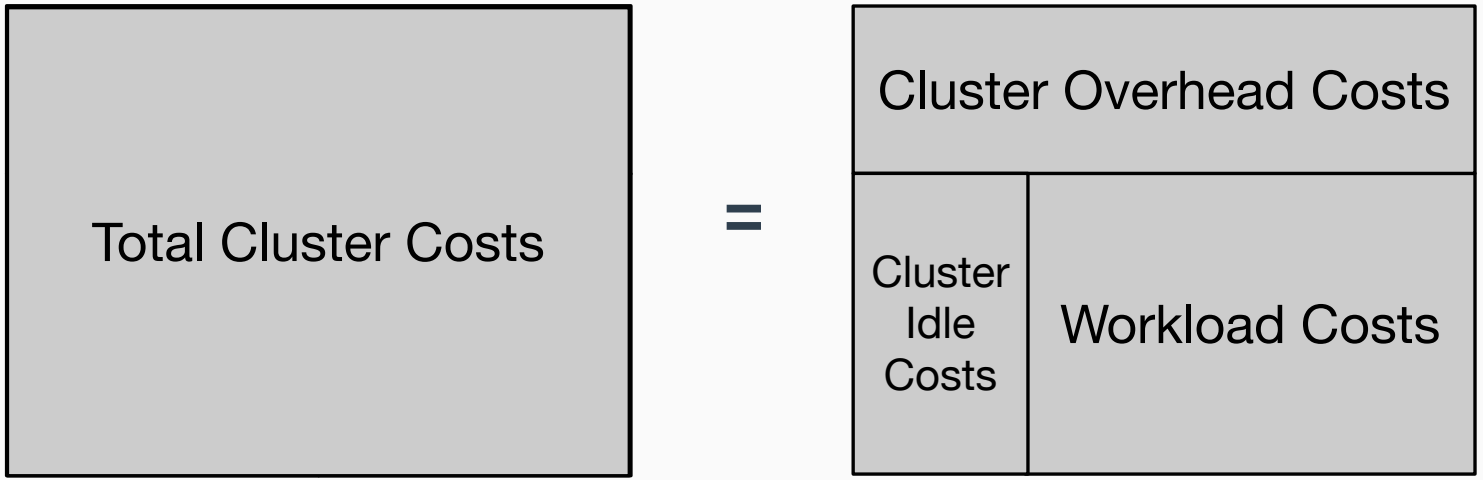

Workload Costs + Cluster Idle Costs

`Total Cluster Costs = Workloads + Cluster Idle Costs + Cluster Overhead Costs`

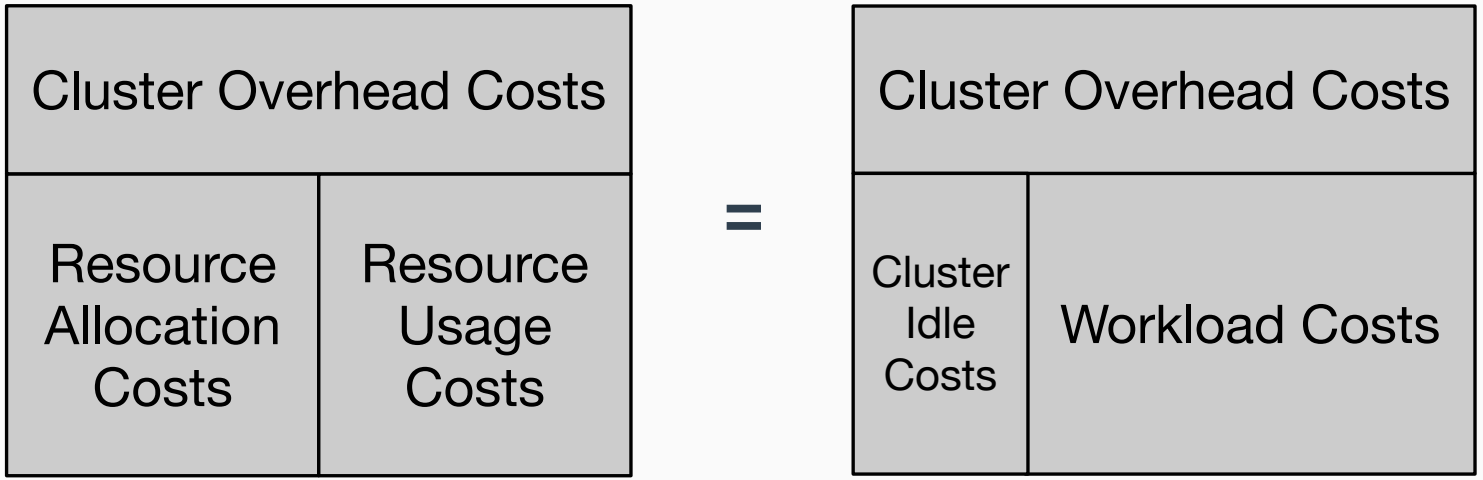

Workload Costs + Cluster Idle Costs

`Total Cluster Costs = Workloads + Cluster Idle Costs + Cluster Overhead Costs`

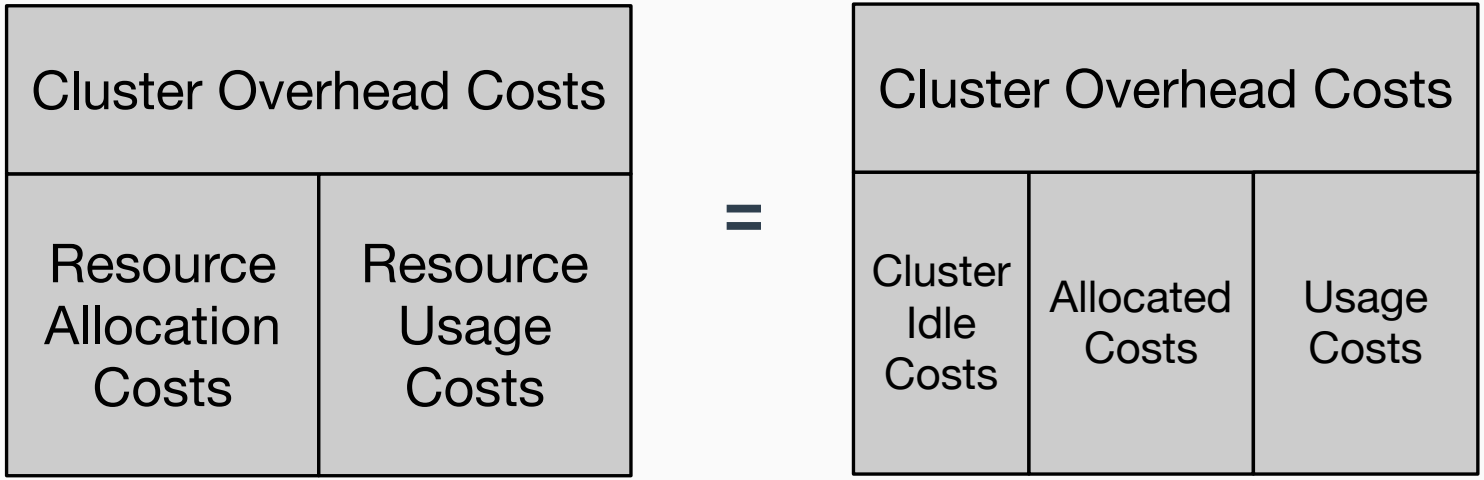

Cluster Idle Cost

`Cluster Idle Cost = Cluster Asset Costs - Workload Costs`

Idle costs can be calculated at the Asset/Resource and at the Workload level.

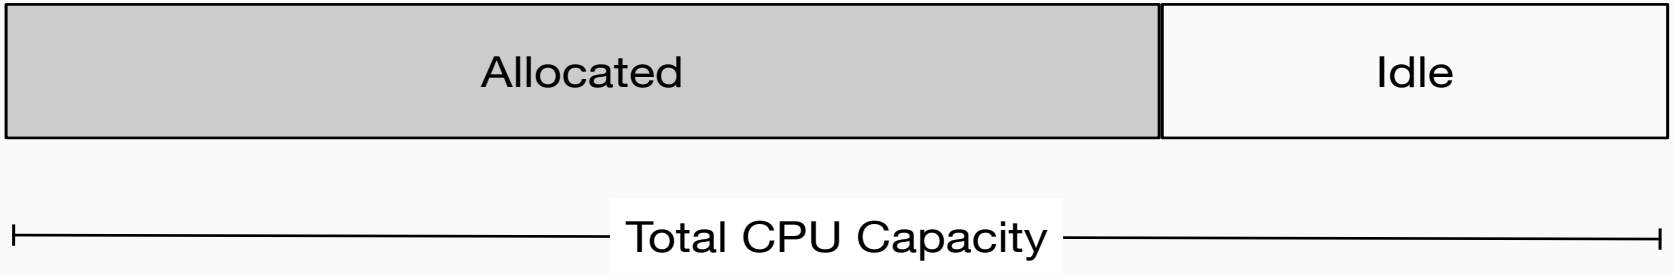

I understand what we're measuring

How do I get at the metrics?

OpenCost Architecture

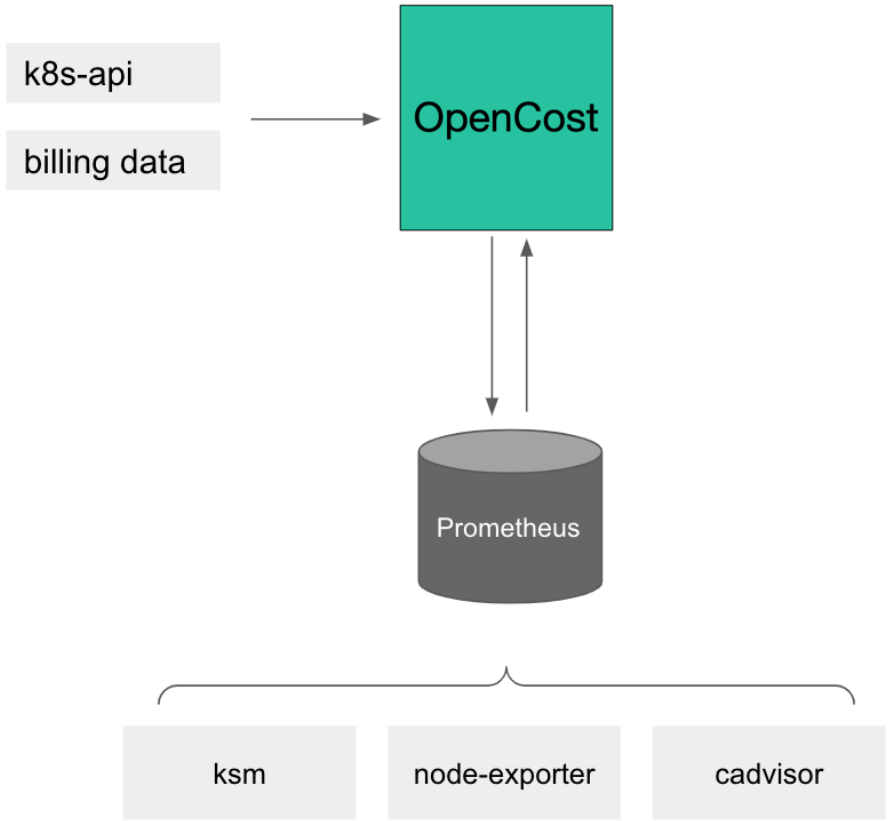

Deploying OpenCost

https://www.opencost.io/docs/install

Prometheus

- Prom community Helm chart the default

OpenCost Manifest

- `kubectl apply --namespace opencost -f https://raw.githubusercontent.com/opencost/opencost/develop/kubernetes/opencost.yaml`

OpenCost Helm Chart

- https://github.com/opencost/opencost-helm-chart/
- Configurable settings (Prometheus, namespaces, etc.)

Accessing OpenCost

- **API**
- **Web UI**
- **kubectl cost**
- **Prometheus**

```
+---------+---------------+---------------------+-----------------+
| CLUSTER | NAMESPACE     | MONTHLY RATE (ALL) | COST EFFICIENCY |
+---------+---------------+---------------------+-----------------+
|         | opencost      |           18.295200 |        0.231010 |
|         | prometheus    |           17.992800 |        0.000000 |
|         | kube-system   |           11.383200 |        0.033410 |
+---------+---------------+---------------------+-----------------+
| SUMMED  |               |           47.671200 |                 |
+---------+---------------+---------------------+-----------------+
```

OpenCost

Last 7 days by controller daily

25 January 2023 through now by Controller

Date Range: Last 7 days | Breakdown: Controller | Resolution: Daily

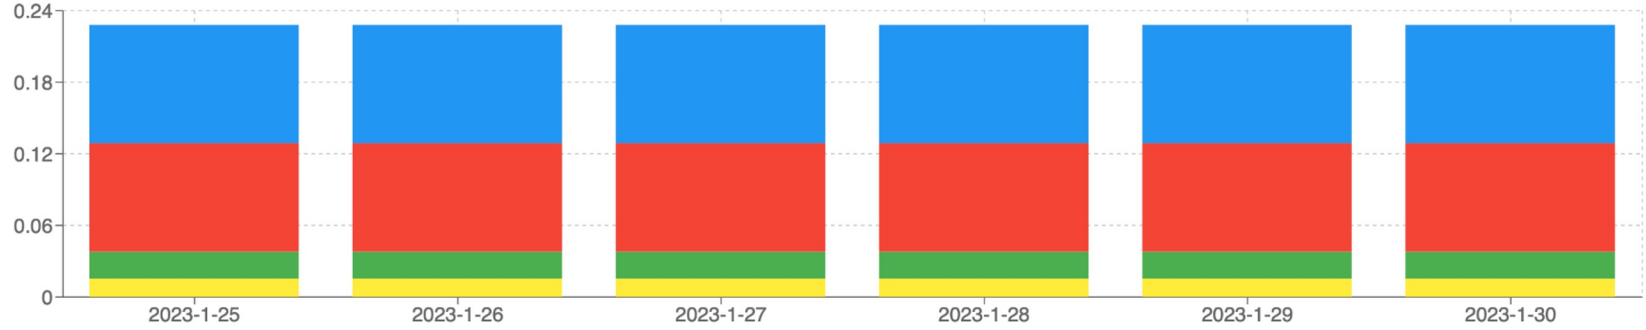

Name	CPU	RAM	PV	Efficiency	↓ Total cost
Totals	$1.28	$0.09	$0.00	15.4%	$1.37
deployment:coredns	$0.55	$0.05	$0.00	3.2%	$0.60
daemonset:kube-proxy	$0.55	$0.00	$0.00	0.3%	$0.55
daemonset:aws-node	$0.14	$0.00	$0.00	12.8%	$0.14
deployment:opencost	$0.05	$0.04	$0.00	18.3%	$0.09

What's the Future of OpenCost?

What's the Future of OpenCost?

What do you want it to be?

Near-Term Roadmap

- External Asset Costs
- Backstage integration
- More Clouds
- More Documentation
- More Integrations

Get Involved with OpenCost

https://www.opencost.io

Slack

- https://slack.cncf.io/ #opencost

GitHub

- https://github.com/opencost/opencost
- https://github.com/opencost/opencost-helm-chart
- https://github.com/opencost/opencost-website

OpenCost Working Group

- https://bit.ly/opencost-calendar
- https://bit.ly/opencost-meeting

LinkedIn

- https://www.linkedin.com/showcase/opencost/

Kubernetes Optimization Strategies

We've got the numbers, now what do we do?

FinOps

https://finops.org

The FinOps Foundation provides guidance on cloud financial management through best practices, education, and standards.

Establish a FinOps practice within your organization.

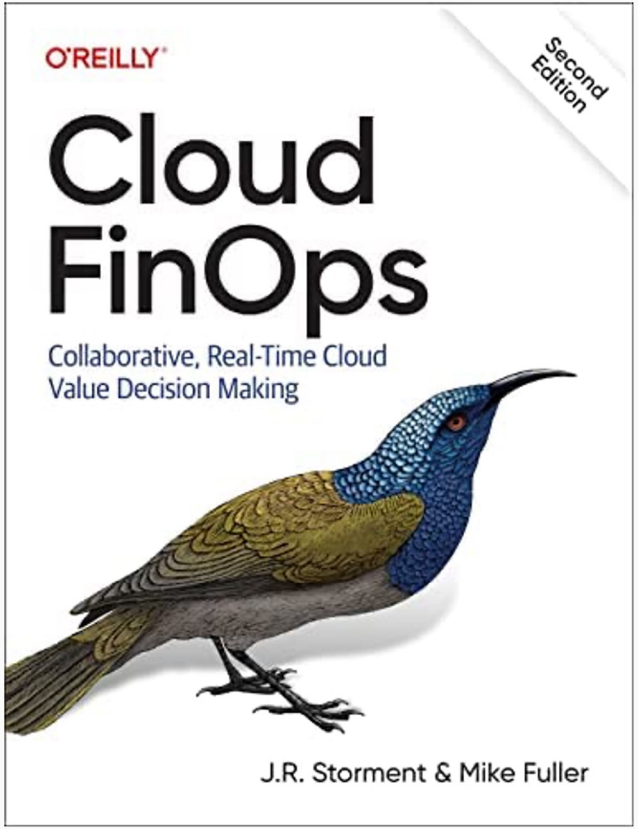

FinOps
Framework

FinOps is an evolving **cloud financial management discipline** and **cultural practice** that enables organizations to get **maximum business value** by helping engineering, finance & business teams to collaborate on data-driven spending decision

Maturity

Principles

- Teams need to collaborate
- Everyone takes ownership for their cloud usage
- A centralized team drives FinOps
- Reports should be accessible and timely
- Decisions are driven by business value of cloud
- Take advantage of the variable cost model of the cloud

Personas

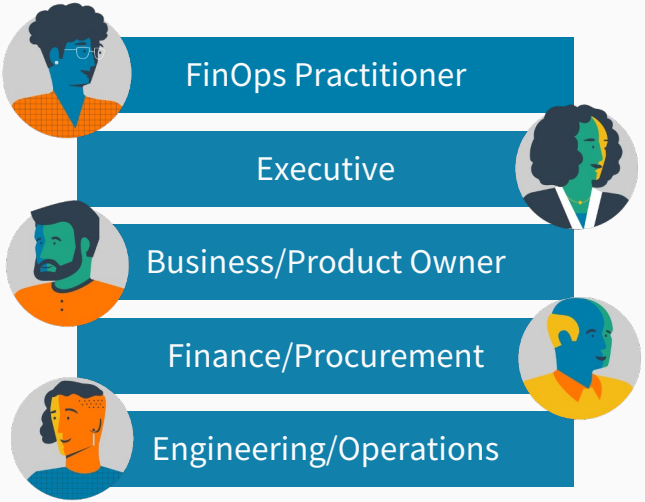

Phases

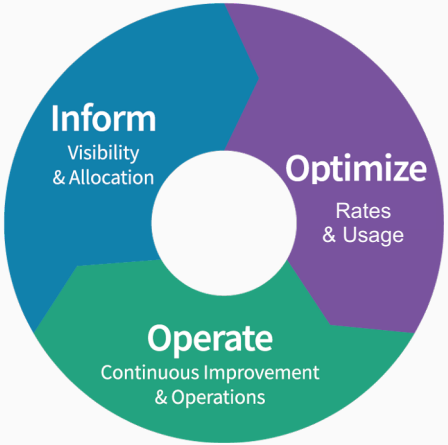

Domains

Understanding Cloud Usage and Cost	Performance Tracking & Benchmarking	Real-Time Decision Making	Cloud Rate Optimization	Cloud Usage Optimization	Organizational Alignment

Optimization Strategy

Start at the top

Efficiencies and cost savings compound

Coordinate your savings across the org

This is an iterative process

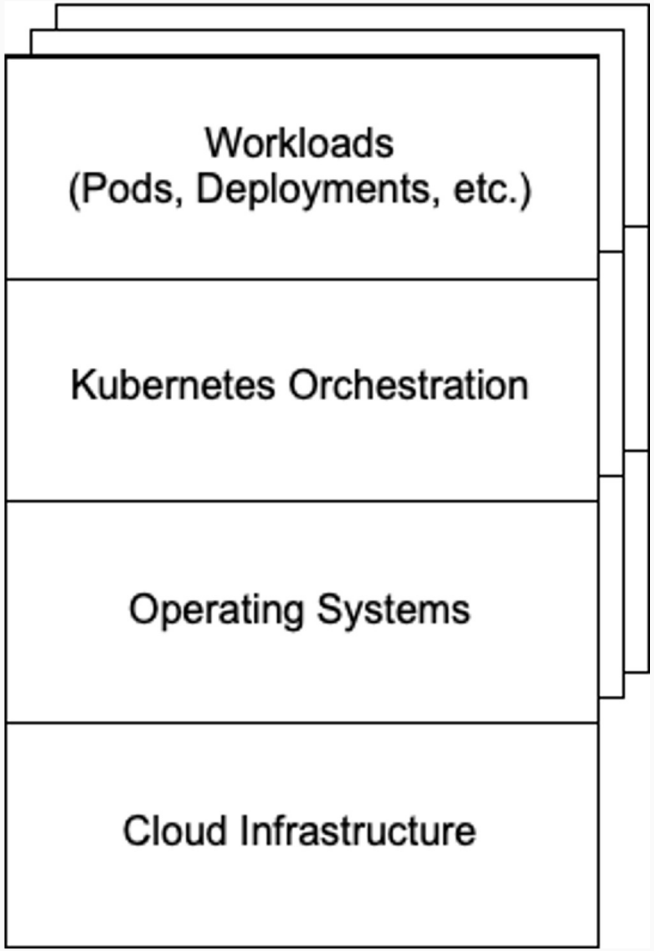

Workload Strategies

Applications Running on Kubernetes

Abandoned Workloads

- Deleting abandoned pods, controllers, or even entire namespaces

Right Sizing Containers

- Updating pod manifests to reflect observed usage
- Providing `requests` and possibly LimitRanges for default resource allocations
- Always provide CPU requests, probably do not use CPU limits
- Always use memory requests and make limits equivalent

Managing Unclaimed Volumes

- Delete volumes that are unused by any pods or move them to a cheaper storage tier

Kubernetes Strategies

Cluster Configurations

Right Sizing Cluster Nodes

- Adjust the number and size of your cluster's nodes to stop overspending on unused capacity
- AMD CPUs may be less expensive than Intel for some workloads

Underutilized Nodes

- Adjust the number and type of your cluster's nodes to stop overspending on unused capacity
- Check CPU, memory, storage class, and network requirements

Managing Unclaimed Volumes

- Delete volumes that are unused by any pods or move them to a cheaper storage tier

Operating System Optimizations

Under the Kubernetes Clusters

Delete Unassigned Resources

- Disks and IP addresses that are not being used by any clusters may continue to incur charges

Resize Local Disks

- Resize local disks with low utilization

Switch to Arm architecture

- Arm CPUs are generally less expensive than Intel across cloud providers for similar performance

Cloud Infrastructure Optimizations

Cloud FinOps

Reserved Instances

- Consider purchasing reserved instances based on historical resource usage patterns

Spot Instances

- Identify workloads ready for spot (preemptible) nodes and resize your cluster to realize the savings of migrating workloads to spot

Savings Plans

- Talk to your cloud vendor about all your options

This is why you have a FinOps team.

Optimization Strategy

Start at the top

Efficiencies and cost savings compound

Coordinate your savings across the org

This is an iterative process

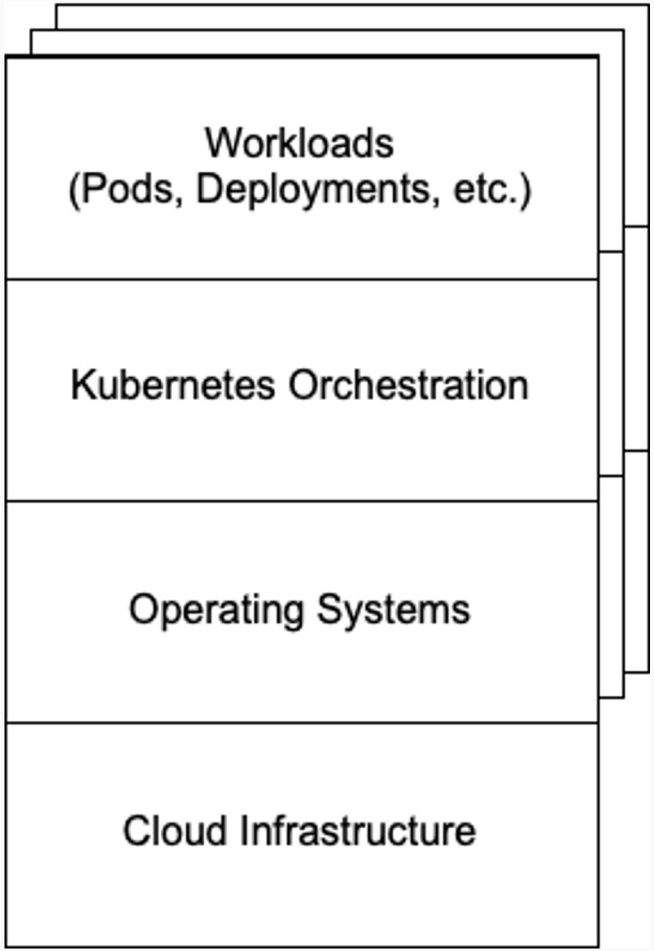

Thanks!

opencost@kubecost.com

OpenCost