Technomics
Better Decisions Faster

Shining Rays of Light & Savings on Cloud Portfolios: An Important Advance!

ICEAA 2023
San Antonio

Tuesday, May 16th, 2023

Reference to any specific product or entity does not constitute an endorsement or recommendation by the authors. The views expressed by the authors are their own and do not imply an endorsement of them or any entity they represent. Views and opinions expressed by Amazon Web Services employees are those of the employees and do not necessarily reflect the view of their employer, the Government, or any of its officials.

Abstract

- Existing cloud cost estimating and pricing tools have at least two significant drawbacks:
  1. They require many assumptions, therefore a detailed understanding of current/future architectures, and
  2. They do not provide insights required to manage costs.
- Technomics developed a groundbreaking alternative that will revolutionize how cloud costs are estimated and managed.
- This paper describes our parametric-based toolset, which enables cloud lifecycle cost estimation, cost reduction and efficiency analysis.

- Track: IT & Cloud Computing (IT06)

Who We Are

RJ Krempasky

ICEAA PCE/A

ADC Engineer at Amazon Web Services.

Performed cloud cost monitoring to assist customers with maintaining dev/prod environments and operations of mission critical systems within executable cost allocations.

Developed cloud cost models, designed/configured databases, and created data visualization tools.

BA in History from Penn State University and MS in Data Analytics from Western Governors University.

Ken Rhodes

ICEAA CCE/A

Project Manager and Employee Owner at Technomics, Inc.

Over fifteen years of experience performing cost analysis and acquisition decision support for DoD customers.

Develops life-cycle cost estimates, cost / price assessments, and data visualization products for software and IT programs.

BS in Industrial & Systems Engineering and MS in Systems Engineering from Virginia Tech.

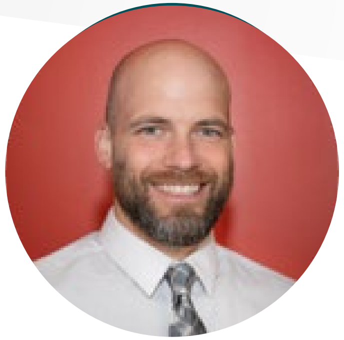

Alex Wekluk

ICEAA CCE/A

Project Manager and Employee Owner at Technomics, Inc.

Over seventeen years of experience performing cost, risk, and technical analyses for the DoD and the IC.

USPTO patent for Marine Corps weapon design work and earned the IC Meritorious Unit Citation for exemplary performance identifying cost-reduction measures.

BS in Mechanical Engineering from Virginia Tech and MA in Economics from George Mason University.

Agenda

Check out our Long-Form Research Paper!

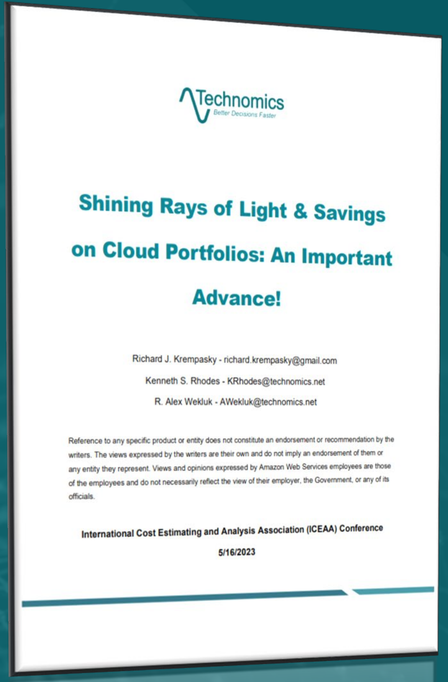
Technomics
Better Decisions Faster

Shining Rays of Light & Savings on Cloud Portfolios: An Important Advance!

Richard J. Krempasky - richard.krempasky@gmail.com
Kenneth S. Rhodes - KRhodes@technomics.net
R. Alex Wekluk - AWekluk@technomics.net

Reference to any specific product or entity does not constitute an endorsement or recommendation by the writers. The views expressed by the writers are their own and do not imply an endorsement of them or any entity they represent. Views and opinions expressed by Amazon Web Services employees are those of the employees and do not necessarily reflect the view of their employer, the Government, or any of its officials.

International Cost Estimating and Analysis Association (ICEAA) Conference
5/16/2023

- Cloud Framework
- MUSCLE Overview, Data Regression & CERs
- MUSCLE Inputs and Outputs
- FLECS Overview, Cloud Service Data
- FLECS Inputs and Outputs
- Conclusion

Problem Statement

Survey feedback from the cloud user community consistently highlights two top cloud adoption challenges facing organizations[1]:

1. Getting engineers to take action on cost optimization
   - Actively monitoring cloud costs to measure performance and reduce waste
   - Better understanding of costs and the relationship to requirements
   - Integration of tools and data to better manage/understand the cost drivers and improve forecasting

2. Accurate forecasting of spend
   - Cost estimators need a data-driven method for estimating cloud costs – based on data, not SME judgement
   - Available early in the software development process, possibly before technical design is available
   - Directly tied to requirements

[1] The State of FinOps Report 2022. FinOps Foundation, 2022, https://data.finops.org/. Accessed 16 August 2022.

Alternatives: Which Estimating Method is Best?

- The best approach to cloud cost estimates depends on when the estimate is needed and **data available**.

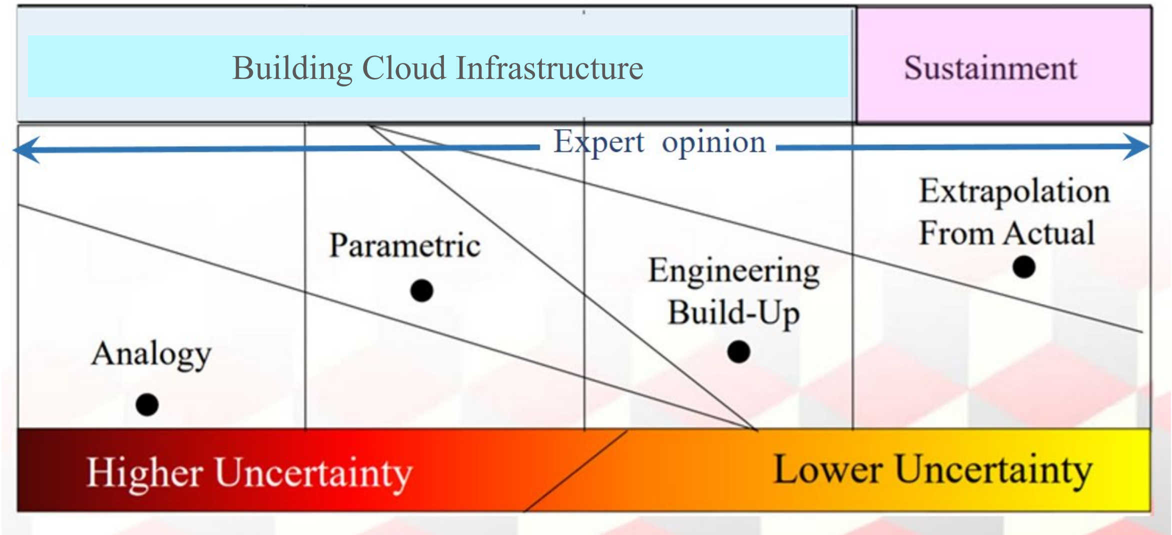

Cost Estimating Body of Knowledge – Software (CEBoK-S), ICEAA, 2021.

Lesson learned, methods improved throughout the lifecycle

- Early-cycle and Mid-cycle estimates, where detailed technical definition is not available need another method such as *Parametric*
  - Requires a standard sizing metric related to Key System Attributes or Key Performance Measures
  - **MUSCLE - Multi-Service Cloud Estimating Model**
- Cloud environment set-up and initial implementation
  - *Engineering Build-up* requires detailed system hardware and cloud requirements inputs
  - Well-Architected Framework
- During Operation, estimates generated late in the life cycle are most defensible with *Extrapolation from Actuals*
  - Requires analysis of real-time cost and consumption measures
  - **FLECS - Framework for Lifecycle Execution of Cloud Systems**

Technomics

Framework

- MUSCLE
  - Produces long-term estimates based on programmatic requirements (ex. # of Galleries, duration of custody)
  - Flexible model accommodates requirements changes over time
  - Given historical data on analogous systems, model accommodates wide-range of program applications
- FLECS
  - Cost optimization (i.e., cost management) uncovers possible opportunities to reduce costs by scaling resources, eliminating unused resources, using automation, and improving processes
  - Efficiencies gained by analyzing service and cost limits, volumes, usage, and efficiency goals can result in common policies and governance that optimize the environment

Lifecycle Cloud Efficiency Modeling

FLECS monitoring helps maximize use of resources

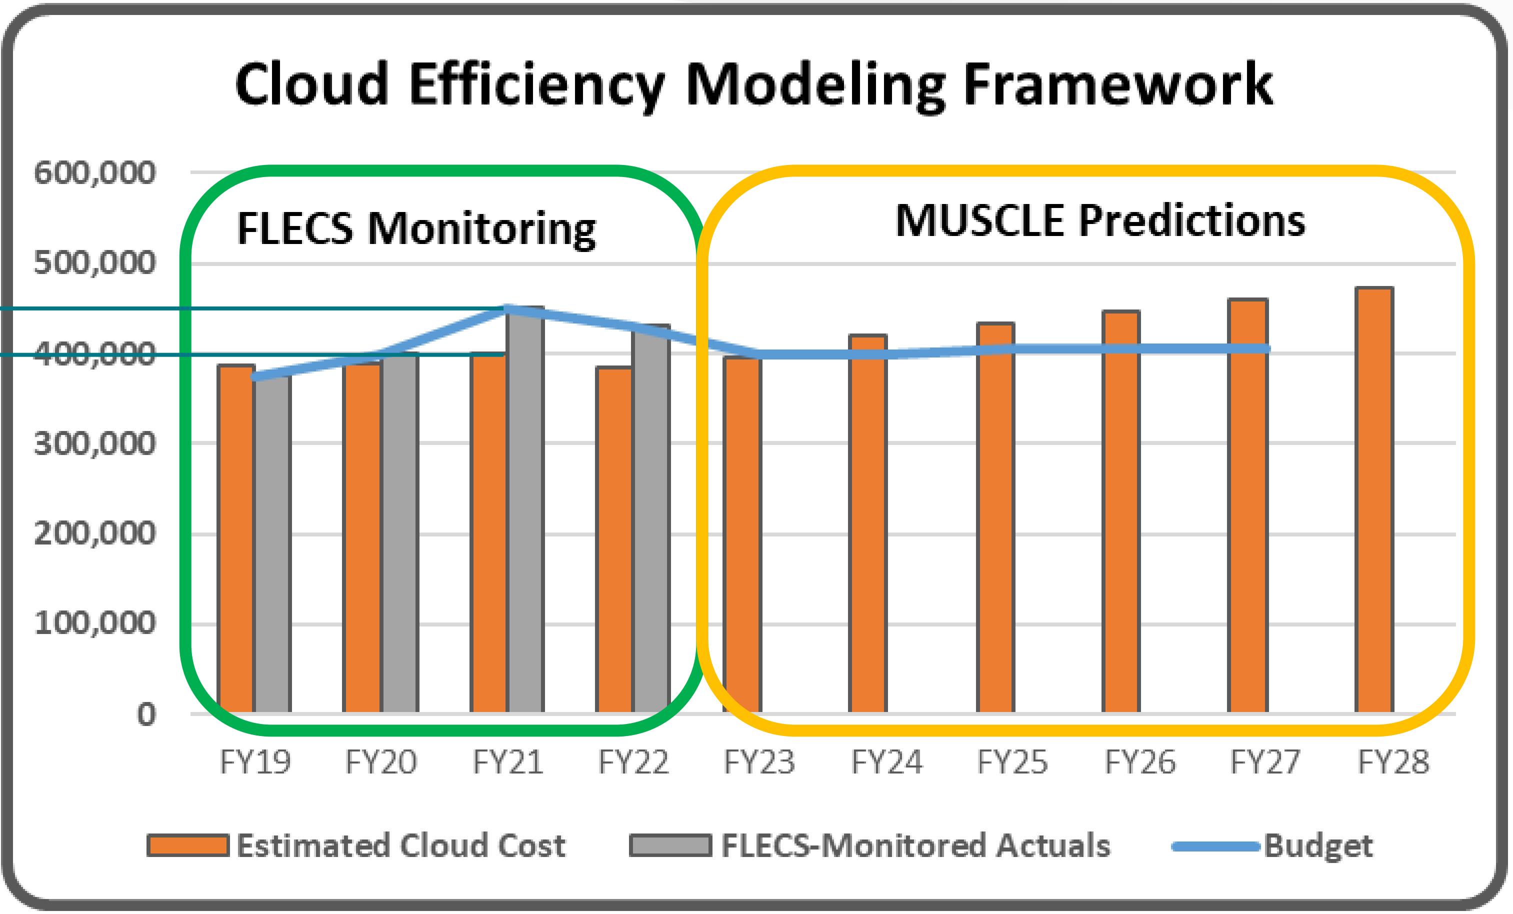

Budgets informed by:

- Agencies extrapolating current costs
- Oversight (Independent Cost Estimates)

MUSCLE Forecasts include changes due to:

- Pricing fluctuation
- Requirements changes
- Developer, user and data growth

Using a combination of realistic cost predictions and active monitoring, cloud usage can maximize mission effectiveness

Technomics
Better Decisions Faster

MUSCLE

Multi-Service Cloud Estimating Model

A modular and scalable cloud estimating model allowing infinite cost scenarios, trades and technical baselines to be examined in one place.

MUSCLE Description

Model informs budgetary and programmatic decisions with robust modular cloud service cost estimates and scenarios

Estimates:

- Cloud instances based on technical **or programmatic** requirements (see Slide 14)
- Inflation using government-approved rates
- Phasing over time
- Calibration to historical costs or pricing quotes
- Cost improvement due to technology maturation
- Increase in cloud services based on usage

Potential Cost Scenarios:

- Predictive pricing
- Vary schedules by instance or end item
- Accuracy in cloud pricing relationship through multi-variate statistical analysis
- Manage entire portfolio of programs
- Model current and future cloud costs with broader participation from carriers
- Pareto analysis (vary technical solutions)
- Cost As an Independent Variable (CAIV)

Model provides a linkage between customer technical and programmatic requirements to cloud pricing through regression analysis

Model Schematic

Cost Data Analysis

Instance Type	Memory	vCPUS	Storage	Network	Hourly Price	Monthly
c6gn.16xlarge	128 GiB	64	EBS only	100 Gigabit	2.7648	2018.304
c6g.medium	2 GiB	1	EBS only	Up to 10 Gigabit	0.034	24.82
m6i.2xlarge	32 GiB	8	EBS only	Up to 12500 Megabit	0.384	280.32
c5d.9xlarge	72 GiB	36	1 x 900 NVMe SSD	10 Gigabit	1.728	1261.44

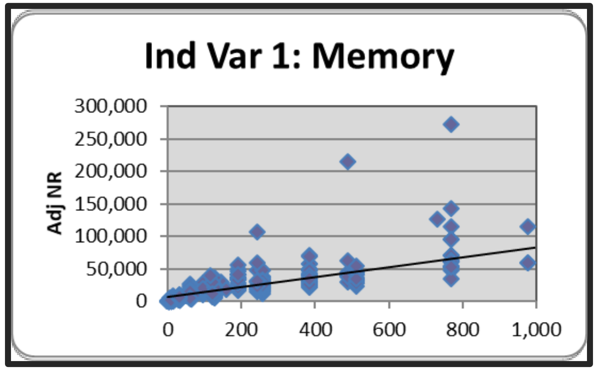

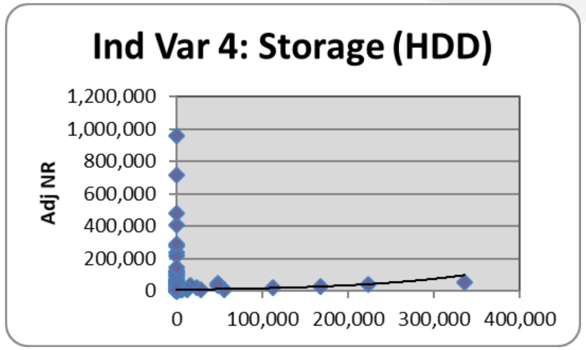

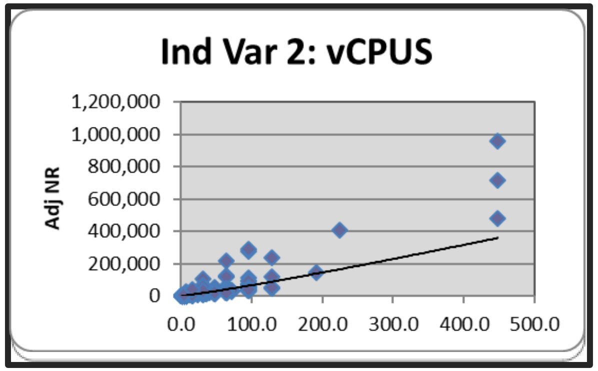

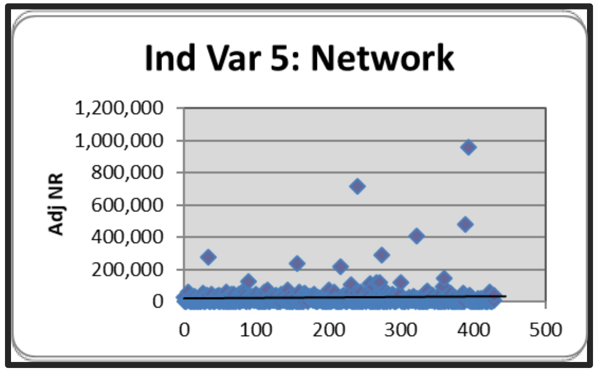

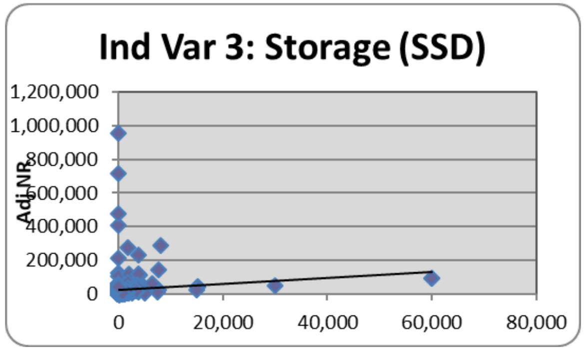

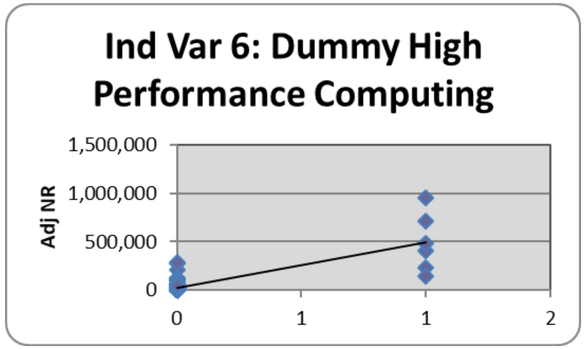

- **Regressed:**
  - 430 data points for Compute
  - 135 data points for Database
- **Graphs showing:**
  - X-axes: Amount of independent technical variable
  - Y-axes: Cost per year
- **Sample of Variables Tested:**
  - Memory, vCPUs, Storage (EBS and S3), Network Speed, SQL (Y/N), Utilization %, Provider, Access
- **Sample of Variables Not Tested:**
  - Data center location, Database type, 1-yr vs 3-yr commitment, Region, Processor (RDS), Architecture
  - Some variables were not used because they represent too-specific a technical solution

Team is actively collecting service provider pricing data and developing new methods

Regressions for CER Library

- Utilized a **Method Development Template** run regressions
- Team ran over **200 regressions** using different combinations of variables
  - vCPUs
  - Storage (SDD and HDD)
  - Network
  - Memory
  - Dummy variables tested, including for subjective 'High-Performance Computing'
  - Operating system
- vCPUs had to be binned (low, medium, high)
- Desire for low number of variables in regression
- Regression statistics with 3 Variables (Memory, vCPUs, Network)
  - **Multiple CERs had:**
  - **Adjusted R^2** > 75%
  - **SPE** < 30%

Tool runs with person <u>in the loop</u> looking for *causation* not machine learning ID'ing *correlations*

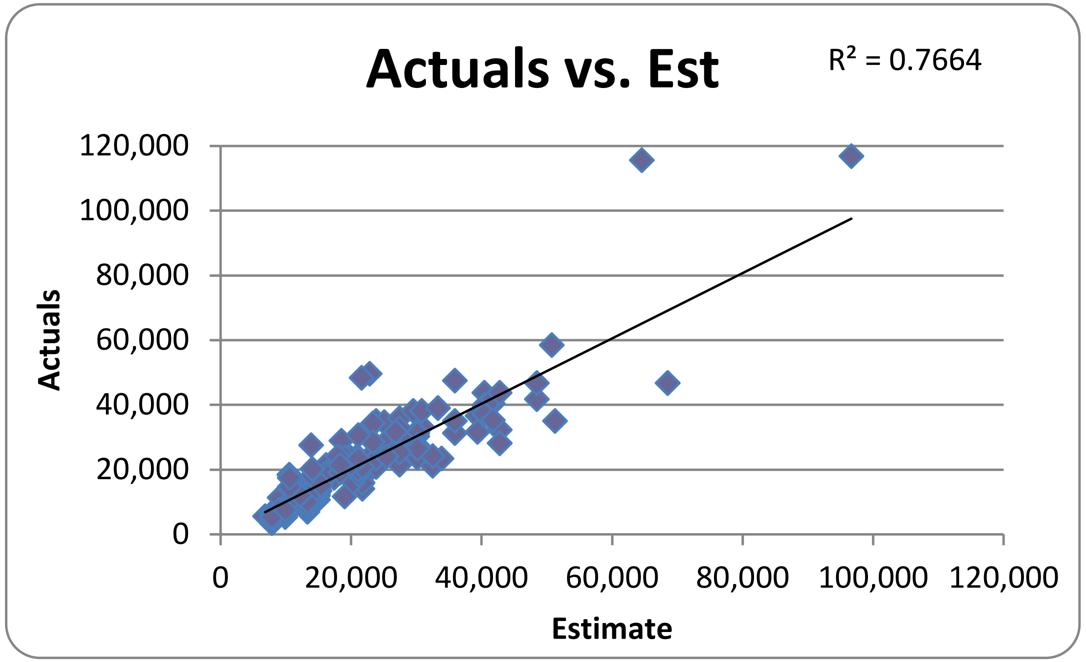

MUSCLE Sample

			Cost Estimating Relationship (CER)	Instance X-Check		Schedule		Independent Var 1		Independent Var 2		Independent Var 3	
Item	Cost Type	Instance Qty	CER Default	Instance	Yrly Cost	Start Date	End Date	Var 1	Input	Var 2	Input	Var 3	Input
Production Controllers	Compute	1	EC2 General	c5.4xlarge	5,957	10/1/2020	9/30/2025	Memory (GB/vCPU)	2	vCPUs	16	Network Speed (Gbps)	10
Production Templates	Compute	4	EC2 General	c5.2xlarge	2,978	10/1/2020	9/30/2031	Memory (GB/vCPU)	2	vCPUs	8	Network Speed (Gbps)	10
Testing Controllers	Compute	3	EC2 General	c5.4xlarge	5,957	10/1/2020	9/30/2031	Memory (GB/vCPU)	2	vCPUs	16	Network Speed (Gbps)	10
Testing Templates	Compute	1	EC2 General	c5.2xlarge	2,978	10/1/2020	9/30/2031	Memory (GB/vCPU)	2	vCPUs	8	Network Speed (Gbps)	10

- Delivered multiple cost estimates using current model
  - Translated quantity of image galleries created to vCPU/instances requirement
- Readily configurable CER Library allows for many forms with any desired independent variables
  - CERs are based on cloud technical requirements and ***programmatic variable***
  - Multi-variate approach translated cloud requirements into unique program-specific variables for mission needs
  - Examples of programmatic requirements that drive cloud computing costs:
    1. Number and type of images
    2. Data rates and time access from communication links
    3. Number of flights entering the country
    4. Number of banking transactions
    5. Privacy/Access (Public, Private, Hybrid)
    6. Data Export Requirements

Model Outputs

Note: Costs and Schedules are Notional

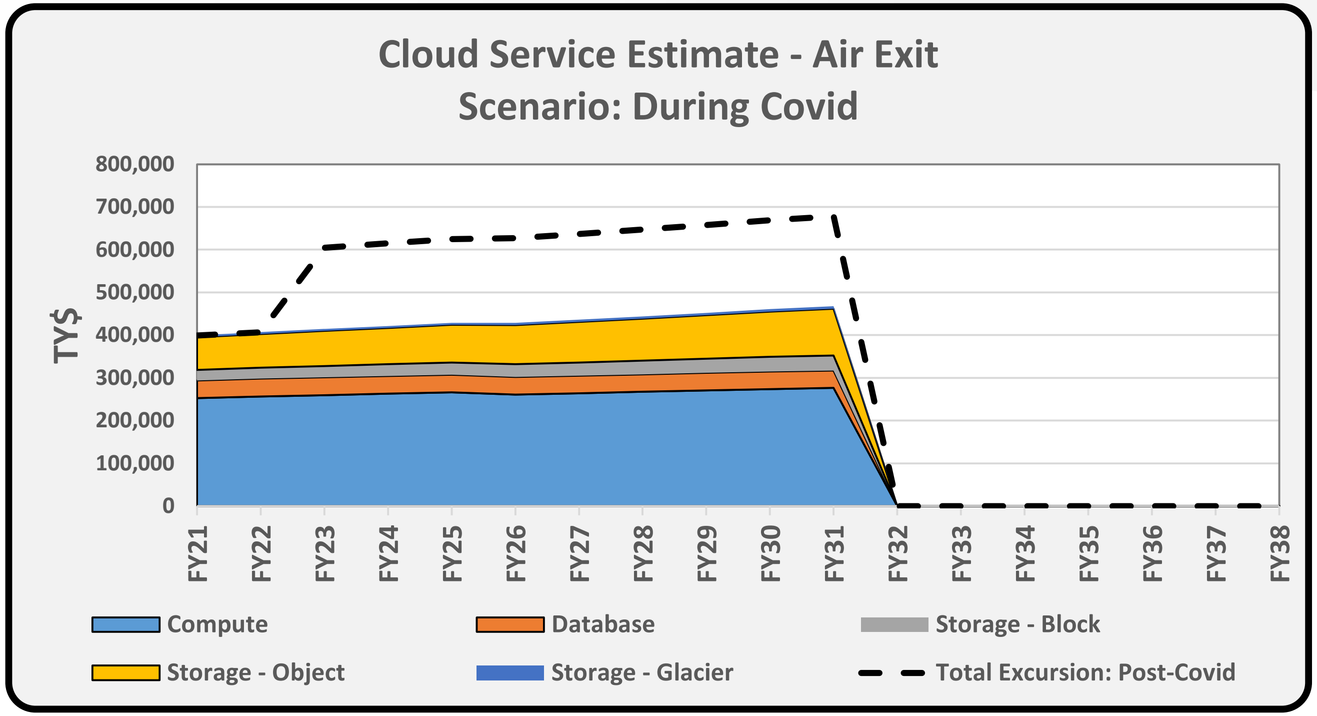

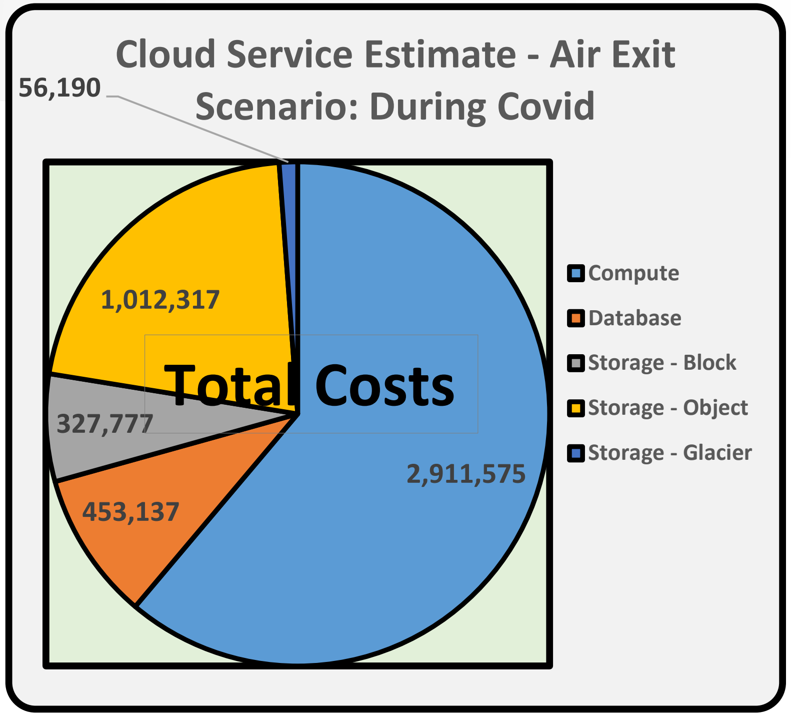

- Sample MUSCLE output from cost dashboard
- Outputs reconfigurable to conform to sponsor standards or mission needs
- **Parametric approach to model minimizes error bounds relative to engineering buildup**

Team and model translate cloud technical requirements into programmatic, enhancing decision-making information available to leadership

FLECS

Framework for Lifecycle Execution of Cloud Systems

Integrated set of tools that cohesively optimizes and manages cloud environments and solutions for our government sponsors

FLECS Description

Leverage an integrated (or multi-app) toolset to cohesively optimize and manage otherwise disparate cloud environments and solutions

Operational and Cost Metrics:

- Cloud technologies heavily instrumented leading to surplus of data
- Multitude of methods used to gather necessary data
- Identification of program/customer metrics pulled from cloud providers categorizing utilization into proper context
- Categorizing and mapping of billing data to align with utilization metrics

Analysis, Architectural & Efficiency Insights:

- Analysis of all identified metrics to monitor cloud usage
- Identification of efficiencies within environment to optimize both end-user service and cost
- Administering Architectural Reviews across entire portfolio
- Continual evaluation of available cloud technologies and impacts on current programs

FLECS enables the continual lifecycle execution analysis as programs live and evolve on the cloud

Operational and Cost Metrics Data Collection

- The modern cloud environment pushes procurement to the edge of development requiring a new paradigm for data collection and analysis
- Each vendor has multiple ways to gather account data (Programmatic API's, Dashboards, Downloadable Metrics)
- Understanding each vendor's environment enables the gathering of complex operational and cost metrics
- Costs associated with each account can be paired with various metrics to bring forth program specific insights

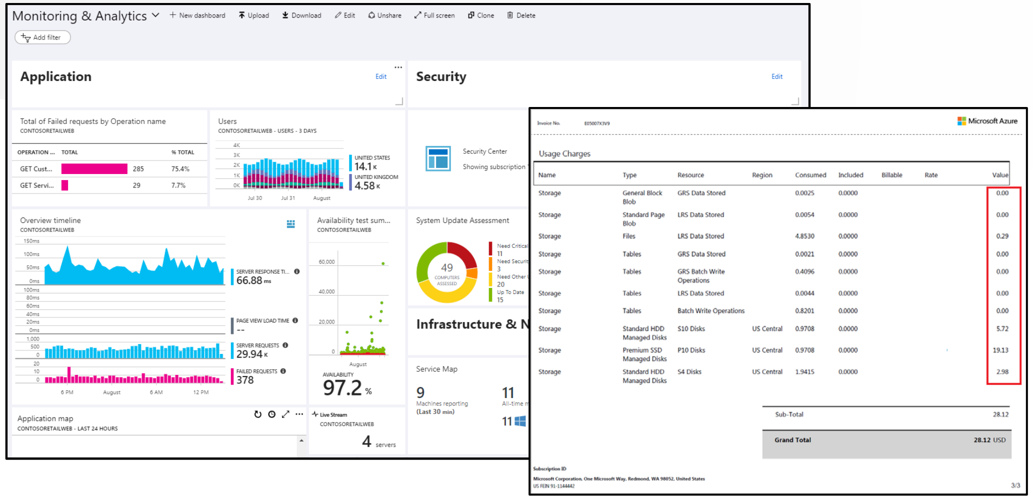

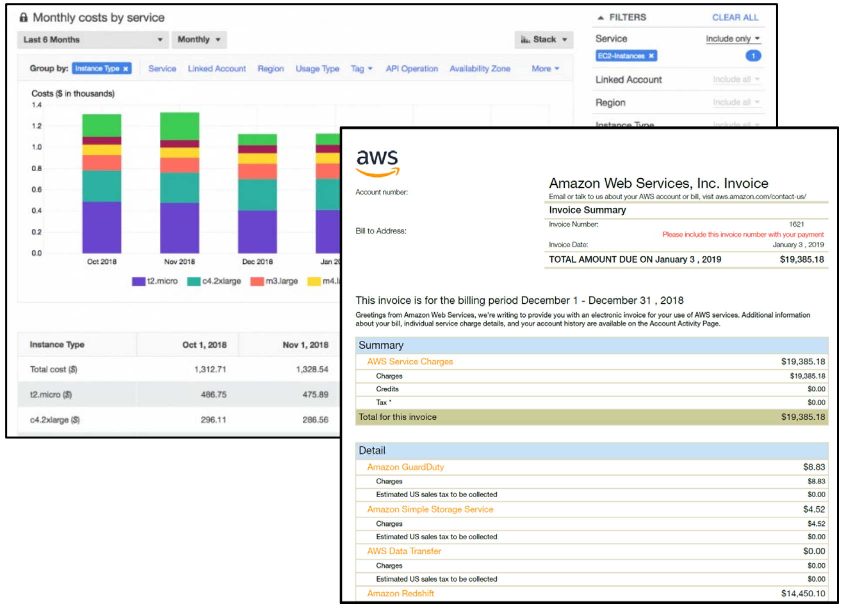

Generating and identifying key data enables lifecycle management in the cloud

Cloud Data Analysis

- Analysis of cloud metrics leads to mission and cost efficiencies
  - CPU Utilization and RAM usage measurements can properly size Virtual Machines to accommodate appropriate workflow for the mission
  - Storage access metrics aid in determining data lifecycle polices, resulting in significant storage cost savings
  - Network latency measurement analysis informs improvements in user experience on deployed apps
- Combinations of metrics can generate specific insights into daily operations custom-built to individual programs (cost per user, number services delivered, user experience)
- Metrics gathered and analyzed can be imported into MUSCLE along with informing architecture reviews to aid in cloud optimizations & cost management

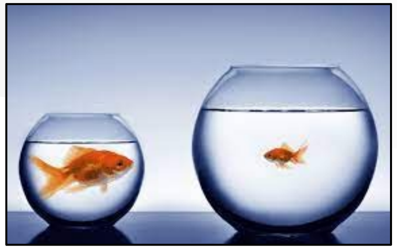

Right-sizing cloud services is critical for optimizing mission within budget constraints

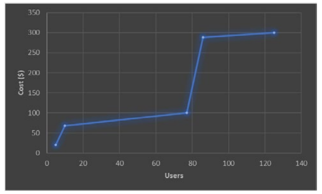

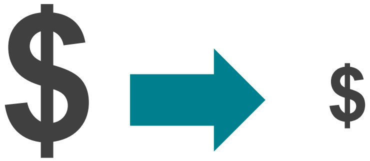

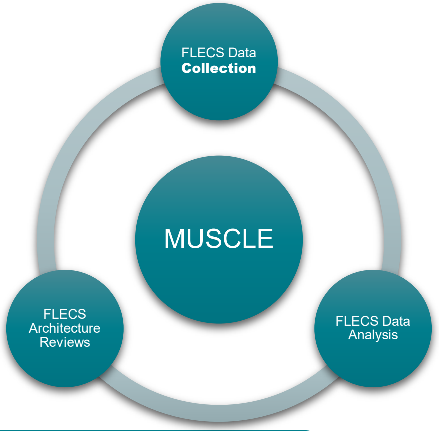

Ingesting disparate automatically-generated datasets combined with expert analysis and visualization results in significant cost savings and increased mission execution

Cloud Architecture Reviews

- Understanding the underlying architecture is crucial to the cloud resource management
- Every action on the cloud utilizes resources that are charged directly to the customer
- Establishing a cadence to evaluate each program within the codified best practices of each vendor yields improved costs and increases delivery of products to the end-user
- Not utilized as an audit but a collaborative experience between cloud providers, developers, and government customers
- Evaluation of new deployed technologies and practices to enable continued evolution of programs

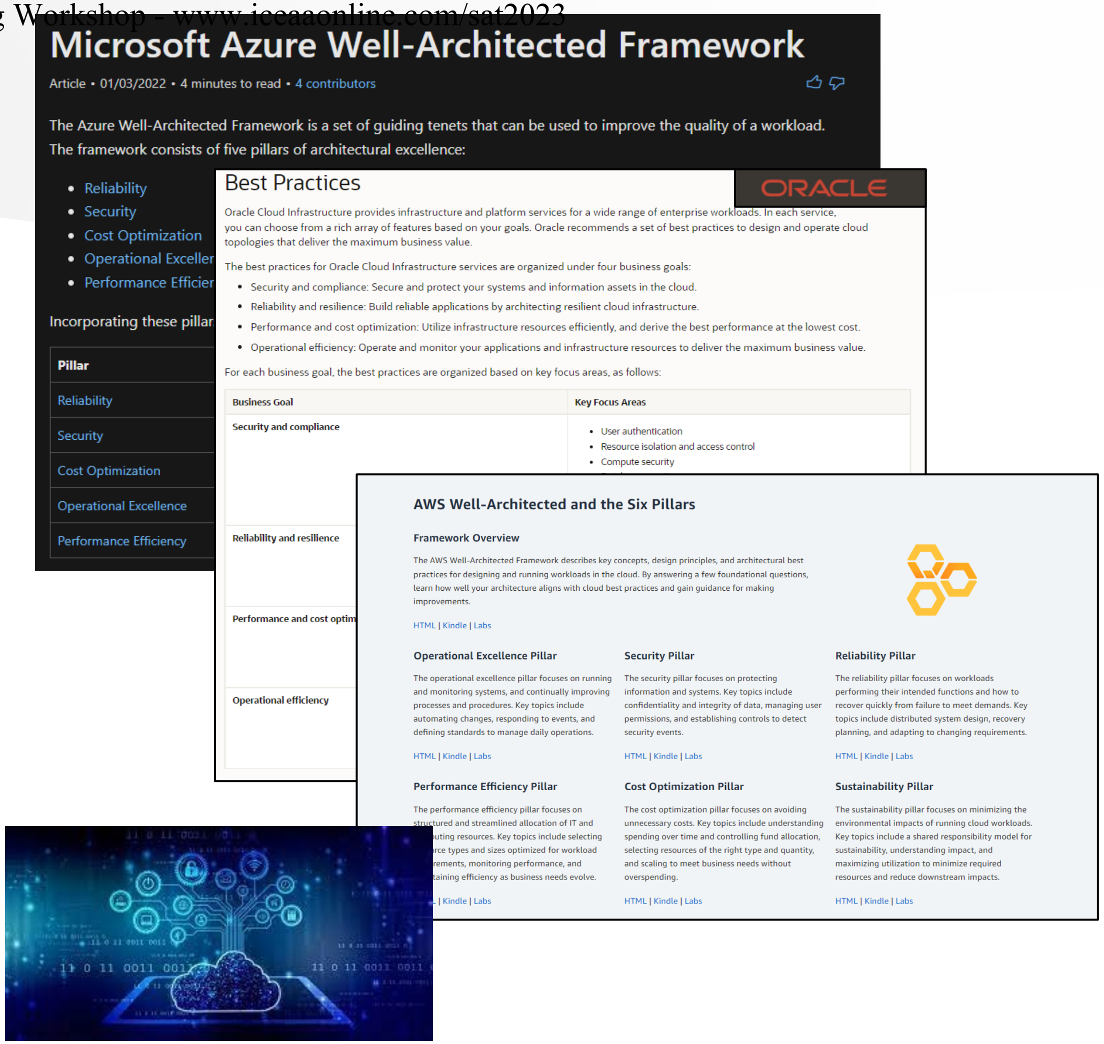

Constant evaluation and best-practice adoption maximizes resource efficiency

FLECS Takeaways

- Currently FLECS is deployed on a government site assisting in the management of over 40 separate cloud accounts hosting mission critical applications
- Data from cloud providers ingested using series of python scripts to output useful datasets
- Statistical insights are generated on the resultant data using a suite of tools
- Then assembled dashboards using the analysis to present insights to decision makers
- Analysis supporting the usage of reserved instances as well as enabling the usage of modern architecture techniques results in significant savings across the cloud accounts
  - Utilizing various API calls pulls data from the cloud accounts and highlights underutilized assets and cloud cost waste focus areas for best practice implementation

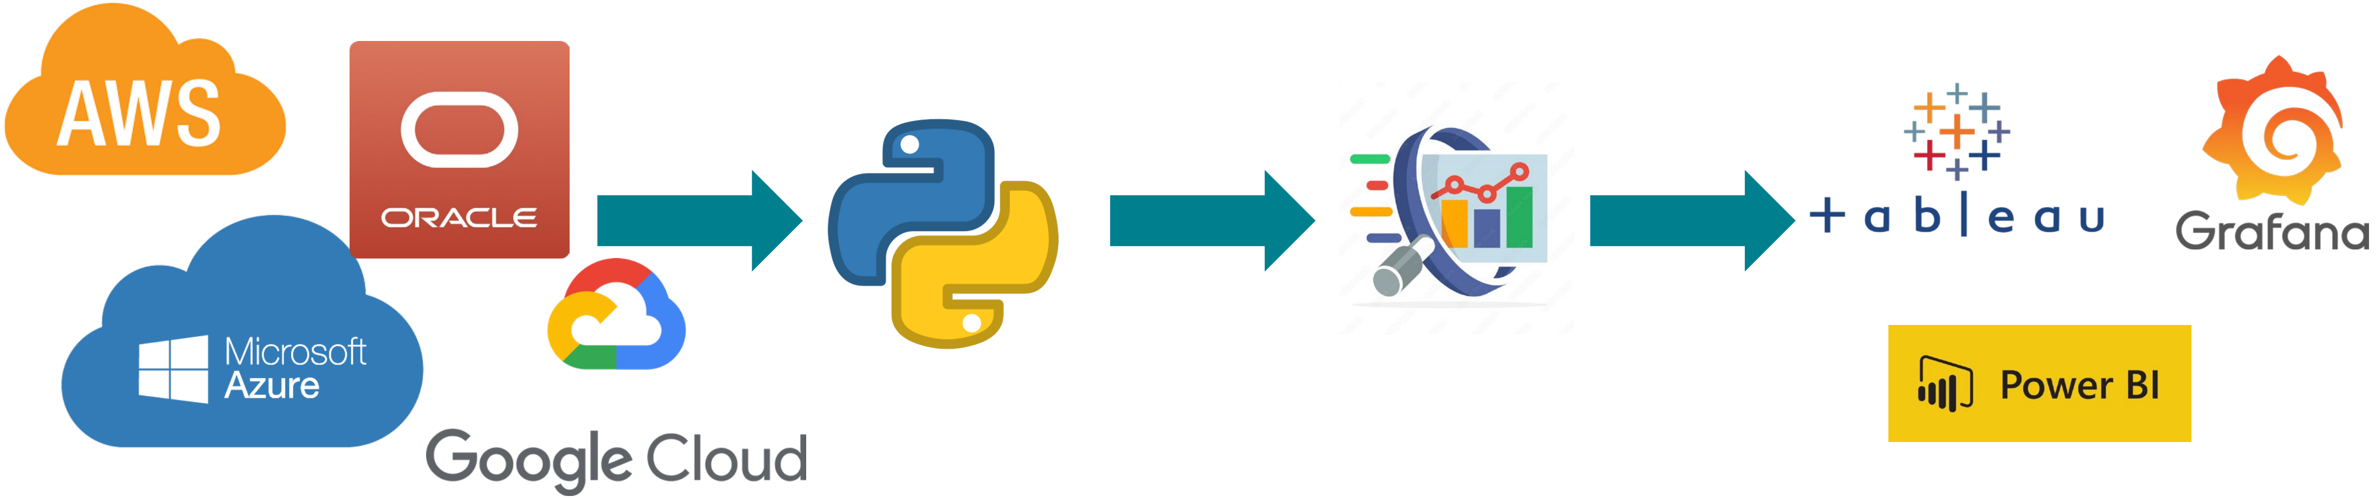

Lessons Learned

- Our objective for sharing the ideas behind the tools is to improve the broader community's ability to maximize their business value by leveraging the principles of active cloud account monitoring and waste reduction
  - Balance understanding of the financial and technical aspects
  - Managing cloud costs is critical for successful digital transformation
  - Cost optimization training, resources and management tools are readily available
  - Take comprehensive review across all accounts and services to fully understand the lifecycle costs
  - Scrutinize spending and leverage variable pricing models in cloud

Next Steps in Cloud Monitoring and Optimization

- Continue building Cloud Pricing database
  - Additional providers, Government vs Commercial, domain specific, etc.
  - Assess price changes over time
- Develop additional CERs for wider-range of programmatic requirements
  - Images per day, # of users, etc.
- Refine and expand methods for translating program requirements to Cloud cost estimates via continuous feedback from FLECS to MUSCLE

Understand cloud purchases and identify areas for cost savings while improving mission performance

RJ Krempasky
ADC Engineer
571-212-8584 (cell)
Richard.krempasky@gmail.com

Ken Rhodes
Employee Owner
571-366-1430 (office)
krhodes@technomics.net

Alex Wekluk
Employee Owner
703-647-2262 (office)
awekluk@technomics.net

Thank you!

Technomics
Better Decisions Faster