



Updated August 2019

Cisco Workload Optimization Manager

Workload optimization and automation for multicloud environments

Speaker name

Speaker title

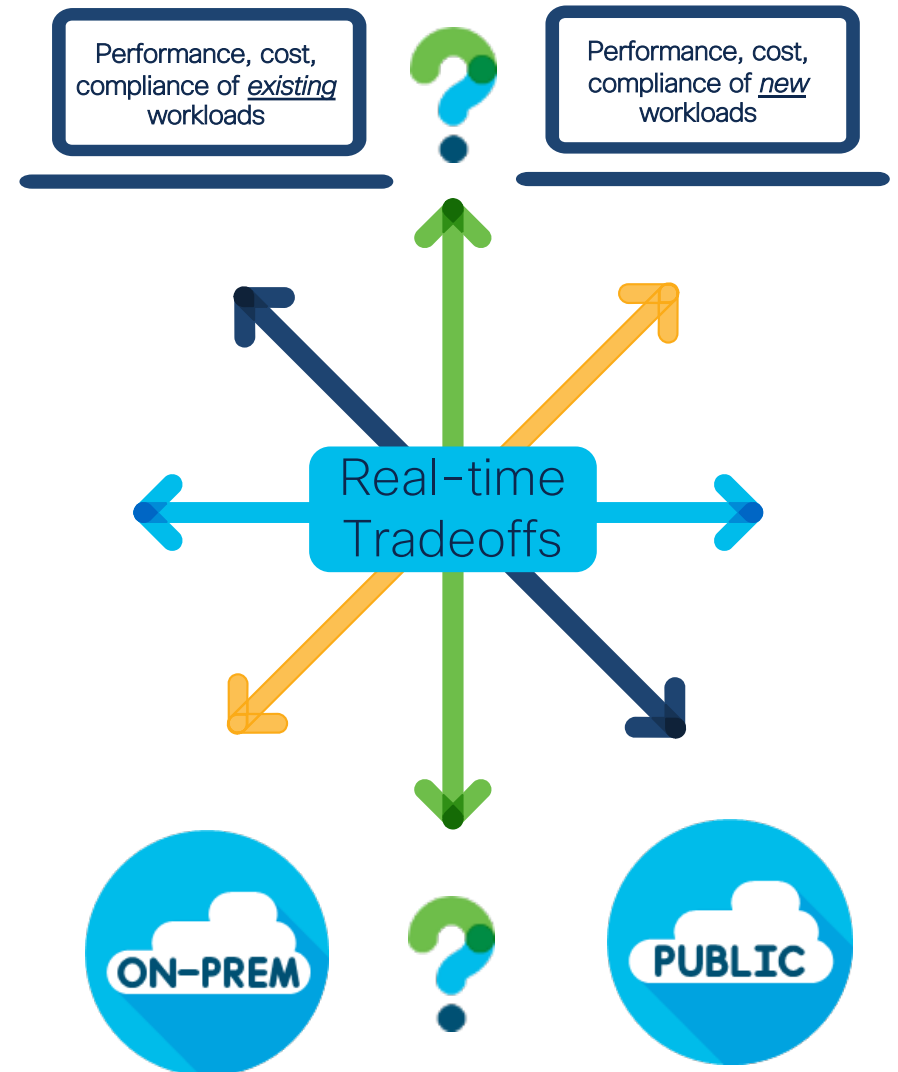
Agenda

- What is Cisco Workload Optimization Manager?
- Customer benefits
- How it works
- Integrations with broader Cisco Multicloud portfolio
- Cisco IT success story
- Appendix: Complementary capabilities in VMware environments

Workload assurance and optimization is a complex problem

- Workload deployments will continue to increase in volume and frequency.
- Multi-cloud environments expand workload placement options, increasing cost overruns and compliance risks.
- Tradeoffs must be made in real time to keep pace with the business.

Decision automation is required



Simultaneously assuring workload performance and compliance while maximizing utilization



The Desired State

CCWOM: Driving Toward the Desired State



Always solving for, simultaneously:

1. Better Performance
2. Increased Compliance
3. Lower Costs

What is Cisco Workload Optimization Manager?

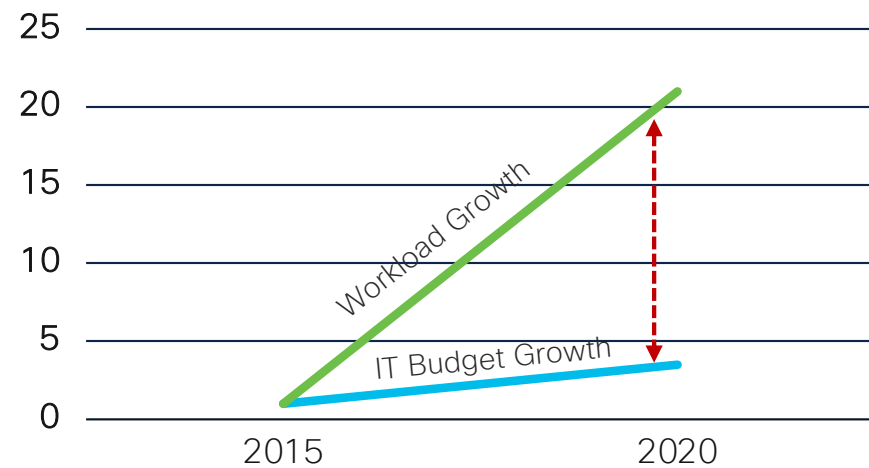
A decision engine for hybrid cloud environments



- Software that **continuously analyzes workload consumption**, costs, and compliance constraints and **automatically allocates resources in real-time, on-premises and in the cloud.**
- It **assures workload performance** by giving workloads the resources they need when they need them.

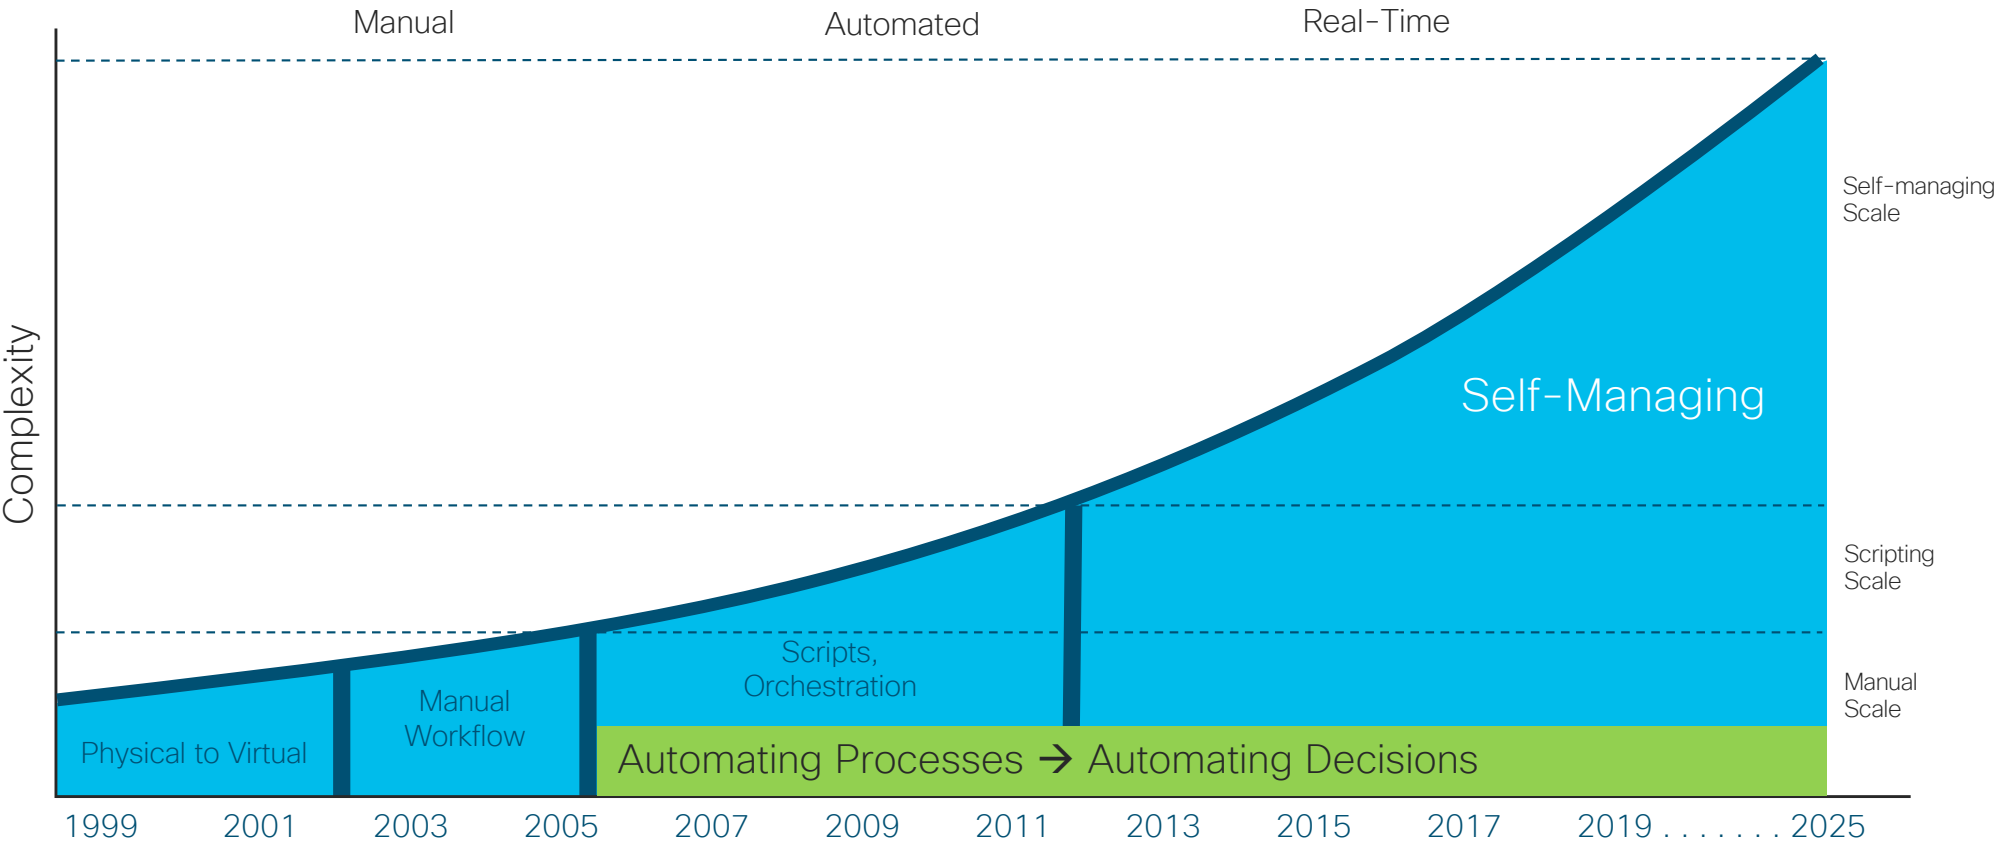
Performance is Critical, but Budgets are Flat

- Outages are labor intensive
- Workloads grow but hardware spend doesn't
- IT fights to keep head above water



How do you assure performance?

Evolution of Automation




© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Partner Confidential

Process Automation versus Decision Automation

Process

- Problems typically addressed after alerting—**reactive**
- Labor intensive
- More data = more noise




Decision

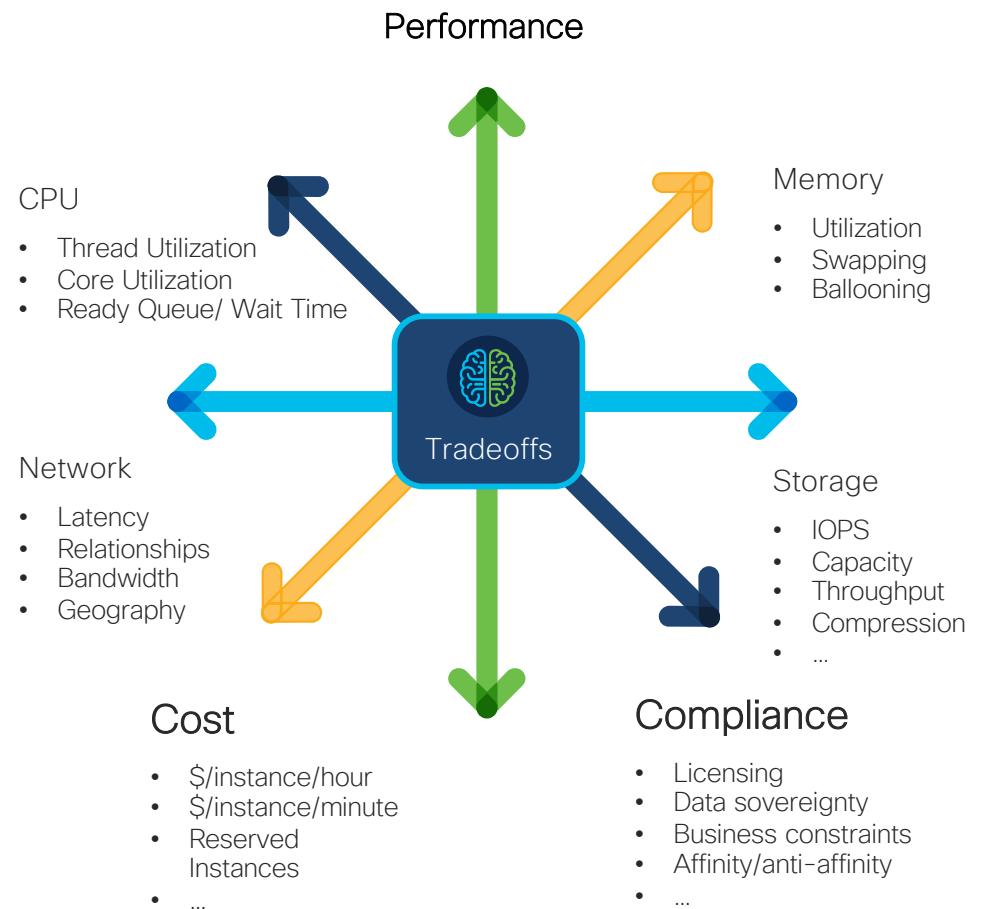
- Problems typically addressed before alerting—**preventative**
- Little to no human intervention
- More data = better decisions



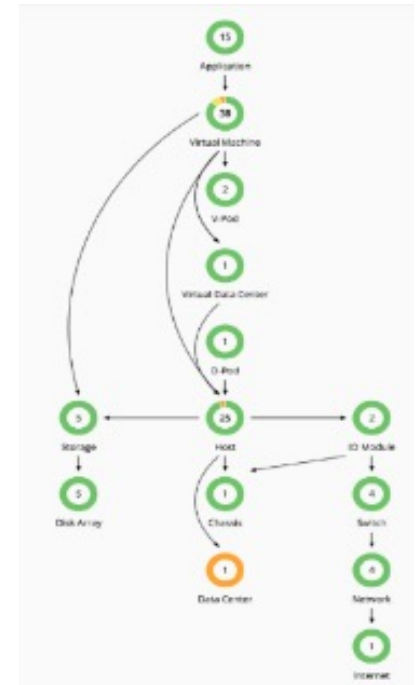
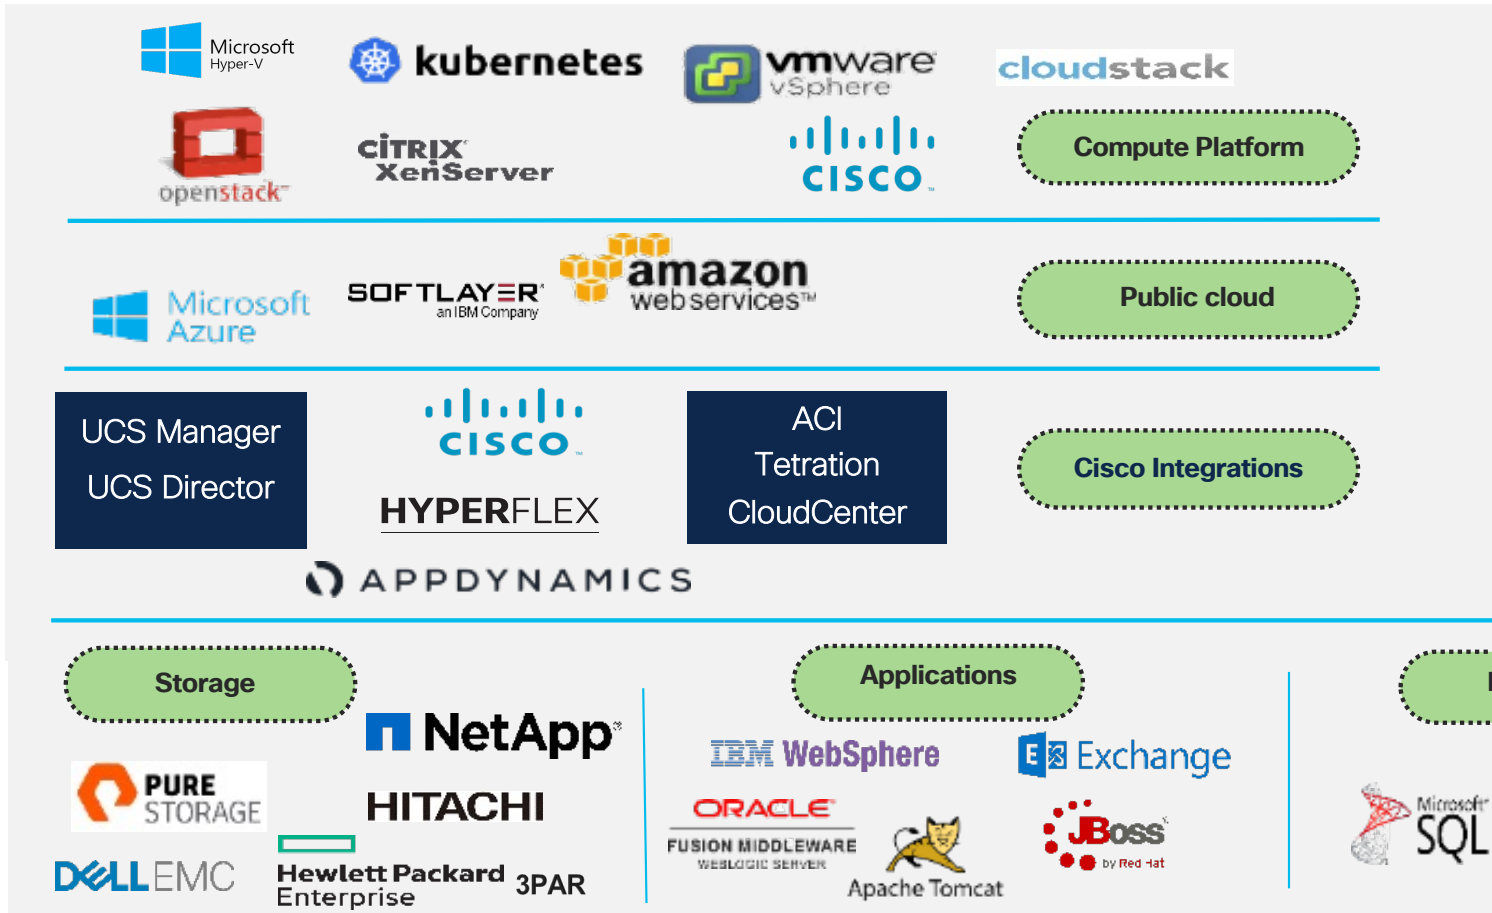
Dynamic Adjustment

- Navigate multiple tradeoffs
- Operate in real time
- Self-managing

- 
- Initial Workload Placement
 - Increase Resources
 - Decrease Resources
 - Move Workload
 - Retire Resources



Cisco Workload Optimization Manager Integrations with broad ecosystem



Cisco Workload Optimization Manager Integrations

Delivering value across the stack and into the cloud

Application-Aware Infrastructure

Drive better optimization through the infrastructure with [AppDynamics](#) metrics.

Self-Managing Container Platforms

Accelerate cloud native projects with production-scale [Kubernetes](#), [OpenShift](#) & [Cloud Foundry](#).

Multicloud Deployment

Deploy workloads with [Cisco Cloud Center](#), optimized for performance, cost, & compliance with Cisco Workload Optimization Manager.

Cloud Elasticity On-Prem

Safely maximize cloud elasticity in [Cisco HyperFlex](#) & [UCS](#) environments.

Super Cluster Optimization

Extend the hypervisor platform and maximize virtualization and [Cisco Hyperflex](#) investments.

Multicloud Dynamic Optimization

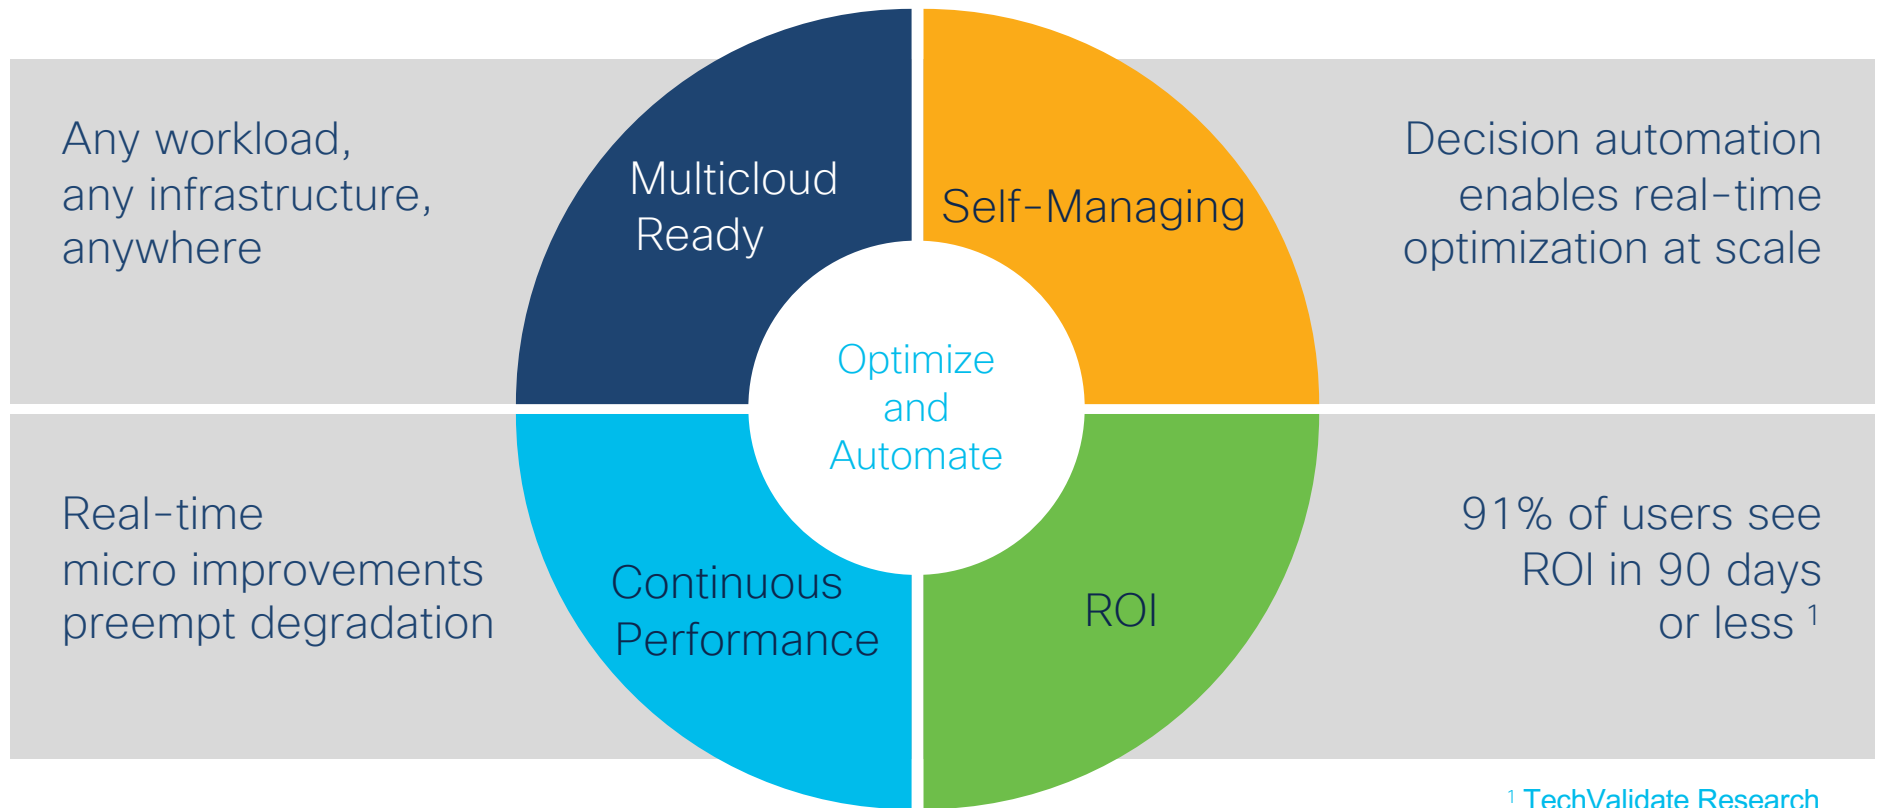
Optimize performance, cost, & compliance in the data center or public cloud ([AWS](#) and [Azure](#)) with one platform.

Network-Aware Optimization

Reduce latency by dynamically localizing “chatty” workloads with [Tetration Analytics](#)

Cisco Workload Optimization Manager

Ensure continuous application performance



¹ [TechValidate Research](#)

Proven Business Outcomes

37%

Better application response time.¹

87%

Increase utilization by 20% or more.²

55%

Reduce user generated tickets by 20% or more.³

91%

See ROI within 90 days or less.⁴


1. [Principled Technologies Report](#)
2. [TechValidate Research](#)
3. [TechValidate Research](#)
4. [TechValidate Research](#)[TechValidate Research](#)

- Performance
- Efficiency
- Compliance

Let's Get Started

- 30-minute install through VM & single OVA file
- Improvement actions appear within 1 hour
- Agentless

What to Expect



Continuous Optimization

Real-time actions drive continuous performance, efficiency, and compliance.



Capacity Management

Quickly & accurately model what-if scenarios: workload growth, add/remove hardware, cloud costs



Compliance & Business Policies

Easy custom policies ensure CWOM actions abide by business and compliance requirements.

When application workloads get the resources they need, when they need them



Applications
continuously
perform



Resources
are efficiently
utilized



Cloud
spend is on
budget



Policies
are
followed




IT
can
innovate

How does CCWOM work?

(How do you automate decisions?)

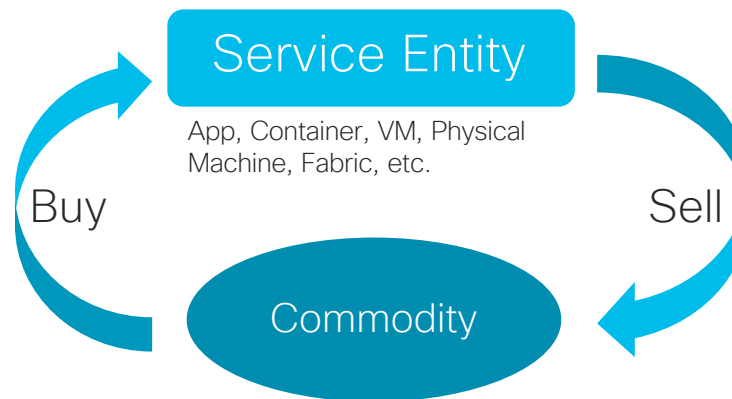
- Abstraction
- Analysis
- Automation

Abstraction: The Supply Chain Market

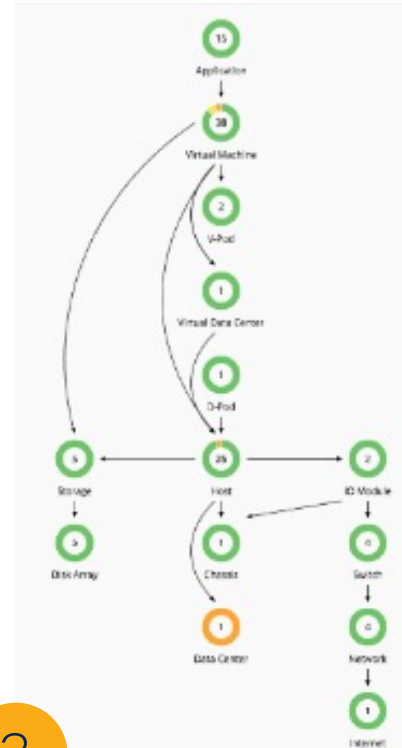


1 Everything in the data center is abstracted into a supply chain market.

2 Services entities shop for the best overall price for every commodity (resource) they need to perform.

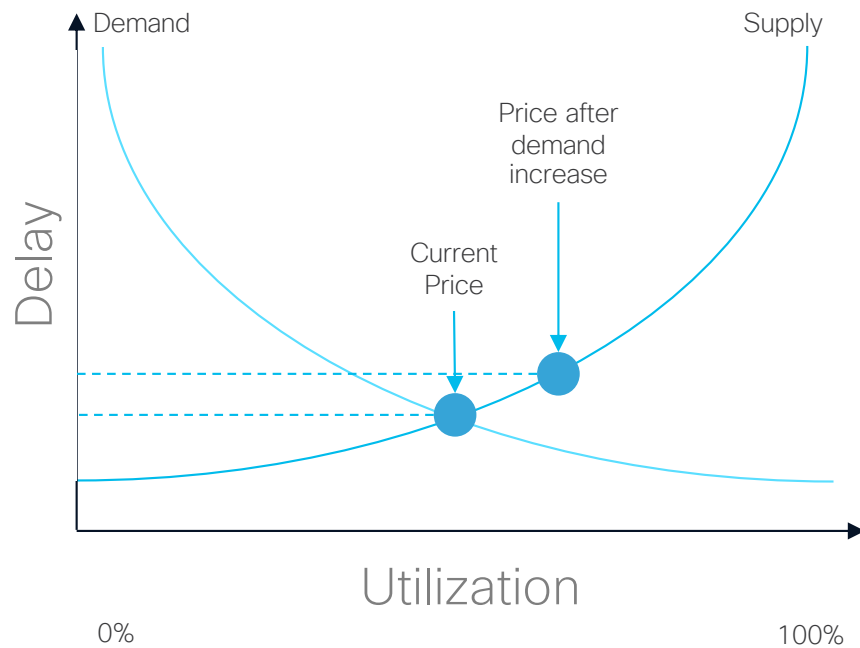


CPU, Memory, Flow, IO, Storage, IOps, etc.



3 Within 1 hr. you see these relationships in the CWOM UI.

Analysis: Economic Supply, Demand, and Price



- Utilization (demand/supply) determines price.
- Workloads/service entities make scaling, placement, and capacity decisions based on *all* the resources they need.

Automation: Real-time Action

Continuous Optimization

Real-time actions drive continuous health:

- Placement
- Sizing
- Provisioning

Capacity Management

Quickly & accurately model what-if scenarios:

- Workload growth
- Add/remove hardware
- Cloud costs

Automation: An Action, Examined

What is the logic behind an action?

Automation: Why Move?

Move VM to Host
Student8IMC200

PERFORMANCE ASSURANCE

Move Virtual Machine 'Student8IMC200' from Host 'm4-esx4.ednvt.cisco.com' to Host 'm4-esx1.ednvt.cisco.com', to Improve workload distribution

Student8IMC200
VIRTUAL MACHINE

| | | |
|----------------------------|-----------------------|---------------------------|
| 0.2 % | 0 % | 2 % |
| 1.9 TB | 81.9 Gb/s | 488.6 GB |
| STORAGE AMOUNT | NET THROUGHPUT | MEMORY ALLOCATION |
| 0.6 % → 0.6 % | 0 % | 1 % |
| 8.2 GHz | 127.9 GB | 0.8 THz |
| VIRTUAL CPU | BALLOONING | CPU PROVISIONED |
| 0 % | 0 % | 0 % |
| 5,000.0 IOPS | 164.6 GHz | 167.6 Gb/s |
| STORAGE ACCESS | CPU ALLOCATION | IO THROUGHPUT |
| 0 % | 1 → 0 | 0 % |
| 100 msec | PRODUCER | 40 Mq/s |
| STORAGE LATENCY | SWAPPING | |
| 2.5 % | 7.4 % | 0.8 % |
| 3.9 TB | 127.9 GB | 1.2 TB |
| STORAGE PROVISIONED | MEMORY | MEMORY PROVISIONED |
| 0.2 % | 2.6 % → 2.6 % | 0 % |
| 20 sec | 12 GB → 11 GB | 91.8 GHz |

Pending Actions

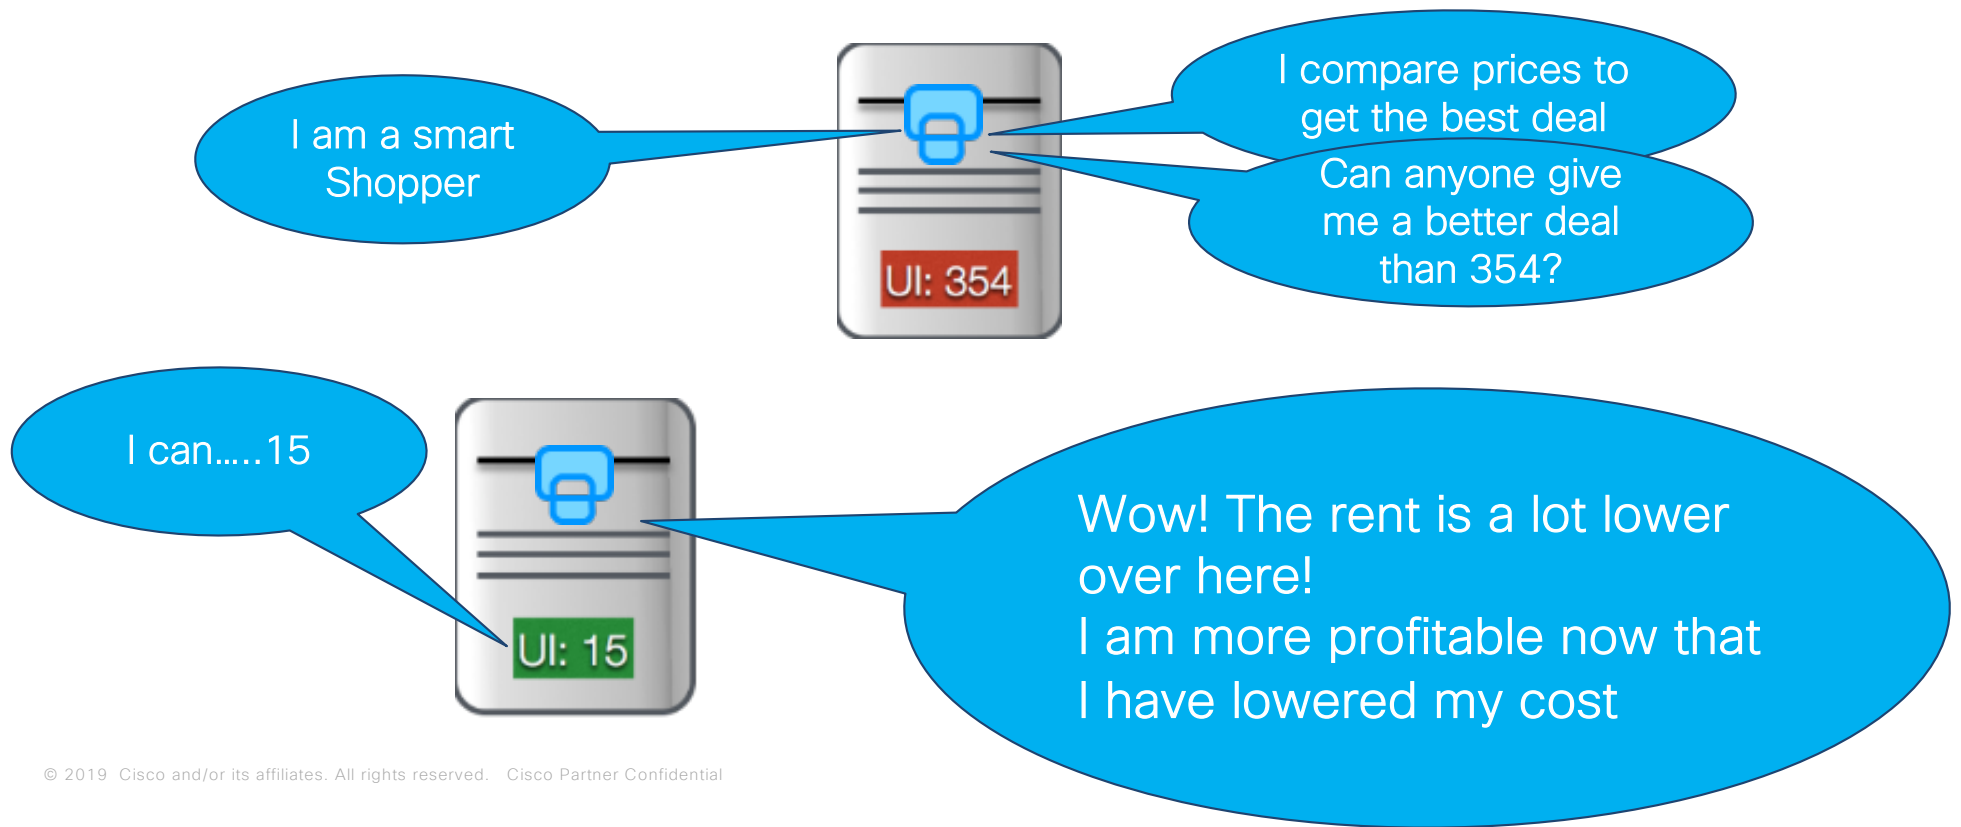
- Move VM to Host
Student8IMC200
- Move VM to Host
Student14JMP7VNX
- Move VM to Host
Student34JMP7VNX
- Move VM to Host
Student2JMP7VNX

Pending Actions

| | | | |
|-----------------------|--|--|----------------------------------|
| PERFORMANCE ASSURANCE | \$36,350 one time Estimated Savings | \$400 one time Estimated Cost | 11 → Placement Actions |
| PERFORMANCE ASSURANCE | | | |
| PERFORMANCE ASSURANCE | 59 Scaling Actions | | 1 Start Actions |
| PERFORMANCE ASSURANCE | | | |

Show all 71 Actions

Entities in Cisco CWOM are Smart Shoppers



Entity Behavior

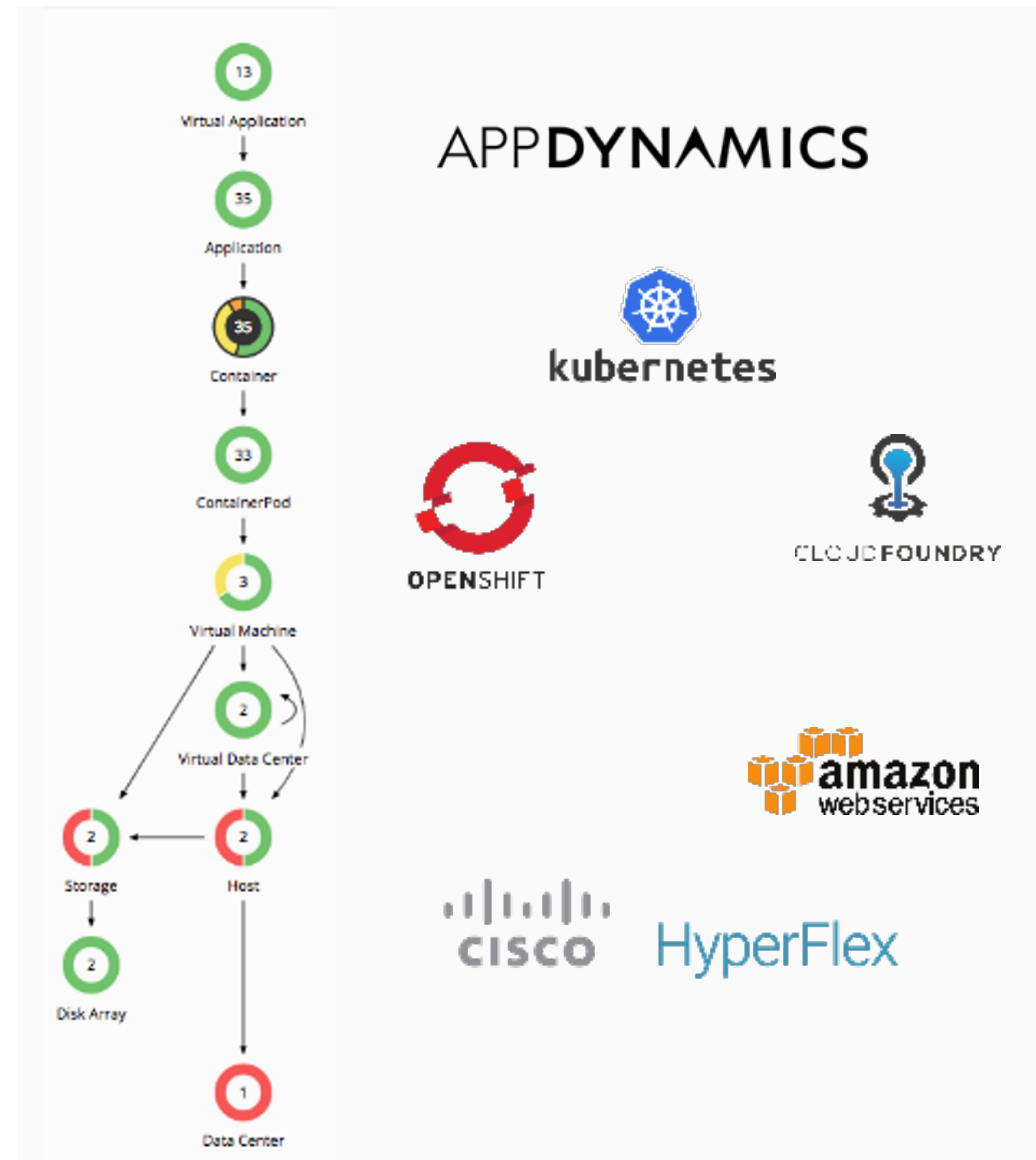
- Entities in CWOM behave according to few simple heuristics
 - Shop around for resources, buy (move to) at the lowest price
 - Evaluate my ROI (Revenue - Expenses) and act accordingly
 - At Equilibrium – hold
 - Very profitable – expand my business
 - Unprofitable – contract my business
- An entity does not know (and does not need to know) the whole supply chain, just the current prices offered by all eligible providers

Basket of Goods

- An entity buys a basket of goods from a provider
 - Example: A VM buys CPU, Mem, ReadyQueue, Ballooning, Swapping, and many other commodities from PMs
- The price charged by the provider is a function of the utilization of every commodity in the basket
 - Higher Utilization → Higher Price
 - Because of the way pricing works, the most utilized resource has the biggest effect on the overall price

CCWOM 2.2 Highlights

- *True* application-aware infrastructure
- Dynamic optimization on HCI—super clusters!
- Self-managing container platforms
- New public cloud savings with continuous optimization of AWS RIs



Cloud Customers Tell Us They are Struggling to Balance Elasticity and RI Savings

Some Customers try to be “Elastic” with People and Spreadsheets



Some Customers Consume the Cloud in “Static” Mode

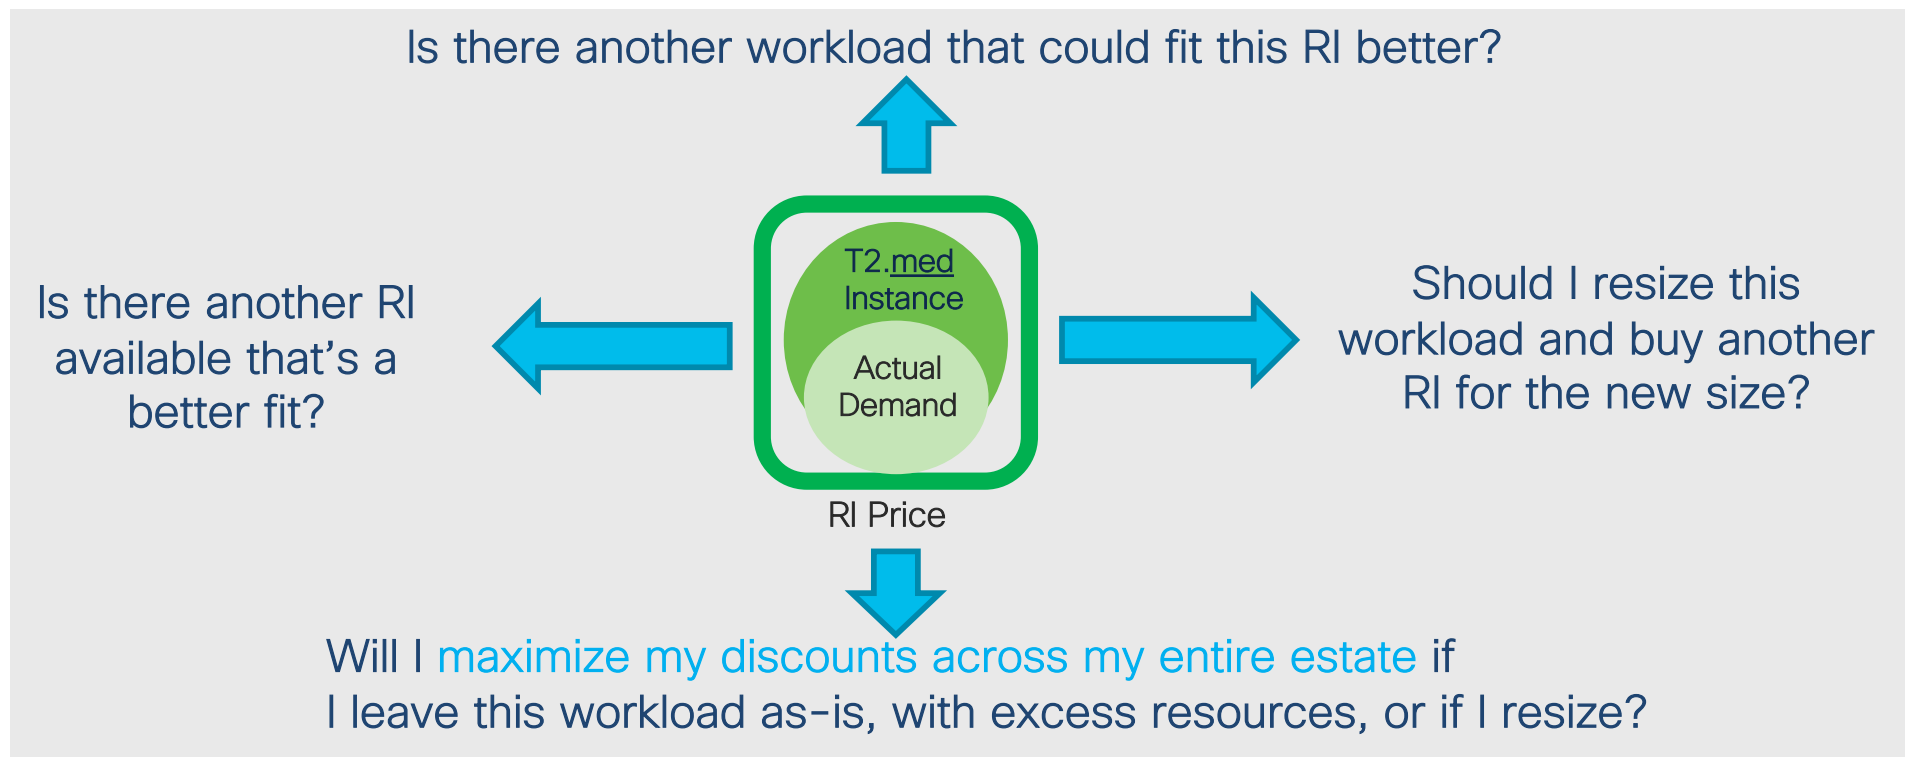
application demand



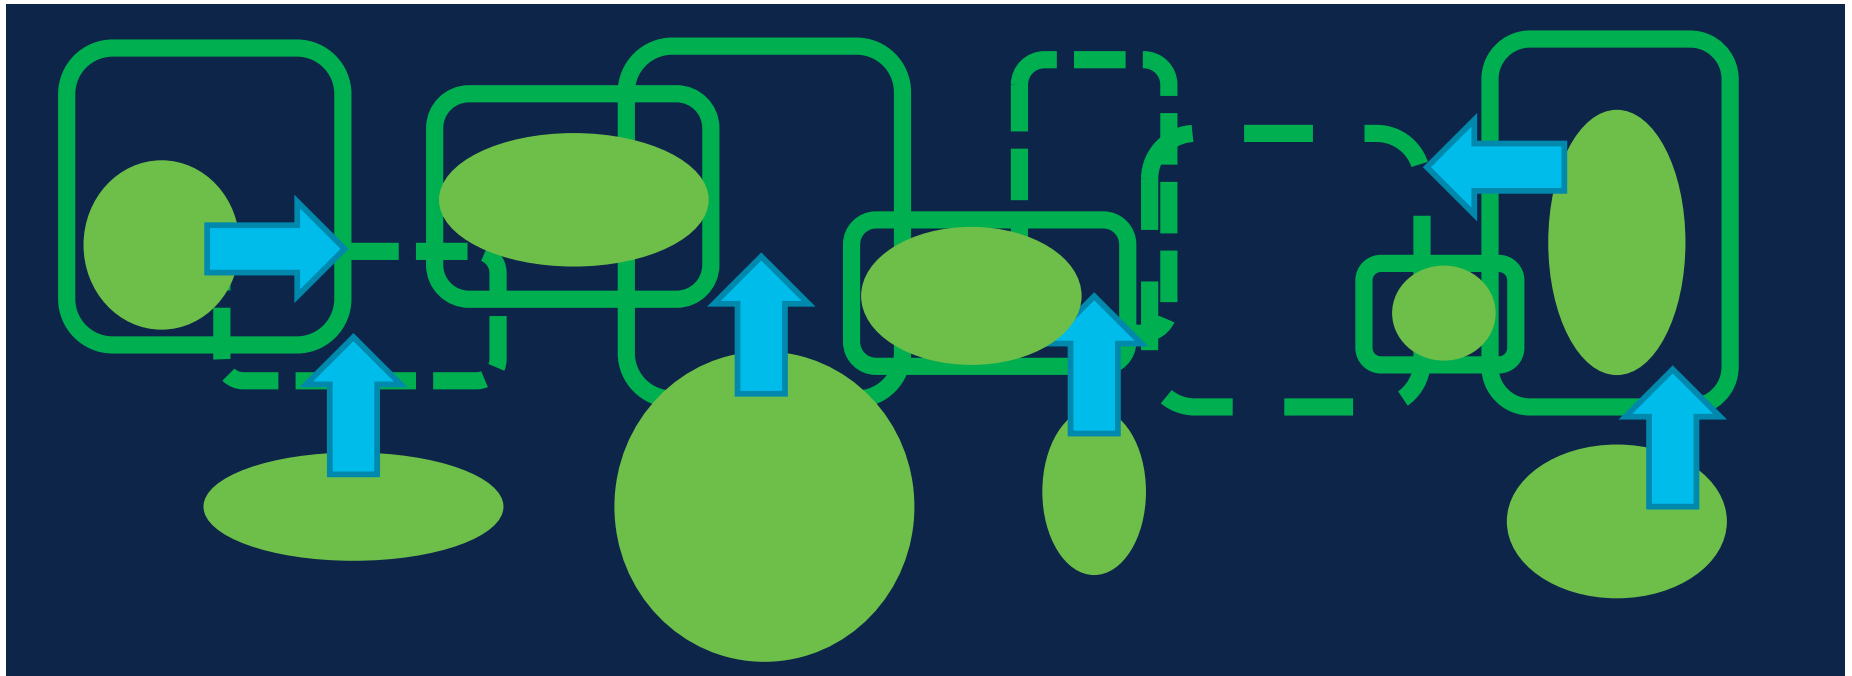
Only Purchase RIs for Steady State Applications

Correct sizing and coordinating with RIs is incredibly complex for even one instance

Example: AWS Instance is Over-Provisioned



Manually Driving the Maximum Cost Savings from RIs – at Scale – Simply isn't Possible



Our customers report fluctuating workload demand that they are unable to optimally match to existing – or future – RI purchases

New in Cisco CWOM 2.2

Continuous optimization of AWS RI consumption – in concert with resizing automation

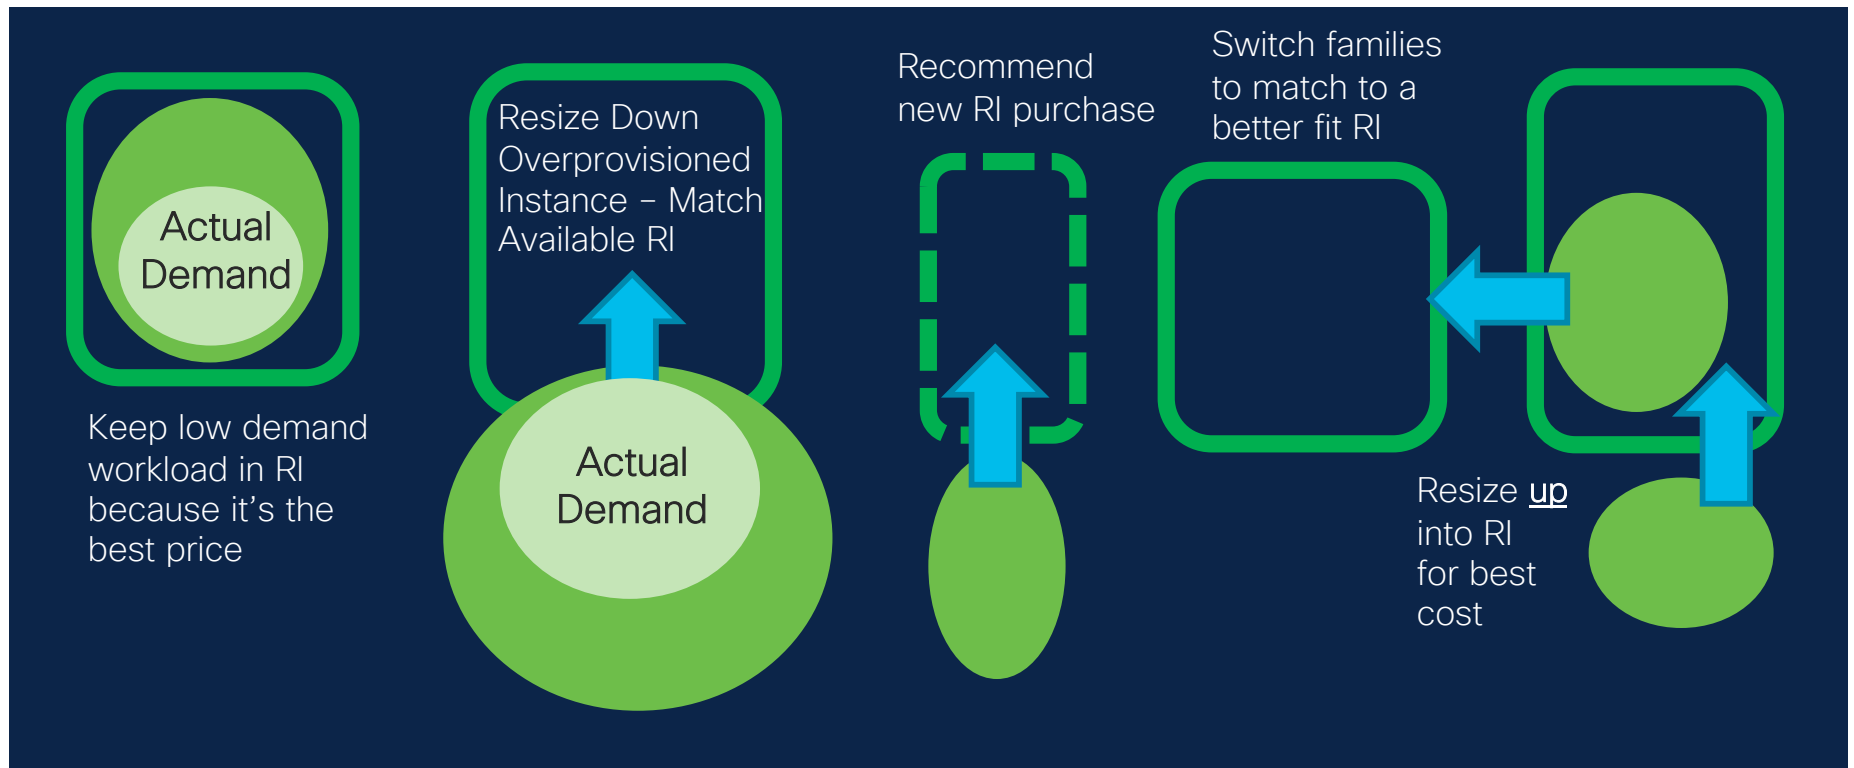
The Cisco CWOM Decision Engine uses one “brain” to holistically:

1. Look at what RIs are available
2. Look at the workloads and their demands
3. Decide and action “best fit at best cost” for **all workloads**
 - Resize up or down
 - Maximize use of available RIs
 - Recommend new RI purchases

Result: Continuously Maximizes RI Savings Across the Entire AWS Estate

New Cisco CWOM 2.2

Example: Cisco CWOM decisions that are coordinated across the estate to maximize AWS savings



Container Projects Underway and Growing Rapidly

Management challenges, resource contention, cost overruns are right behind

CCWOM now optimizes and automates container platforms



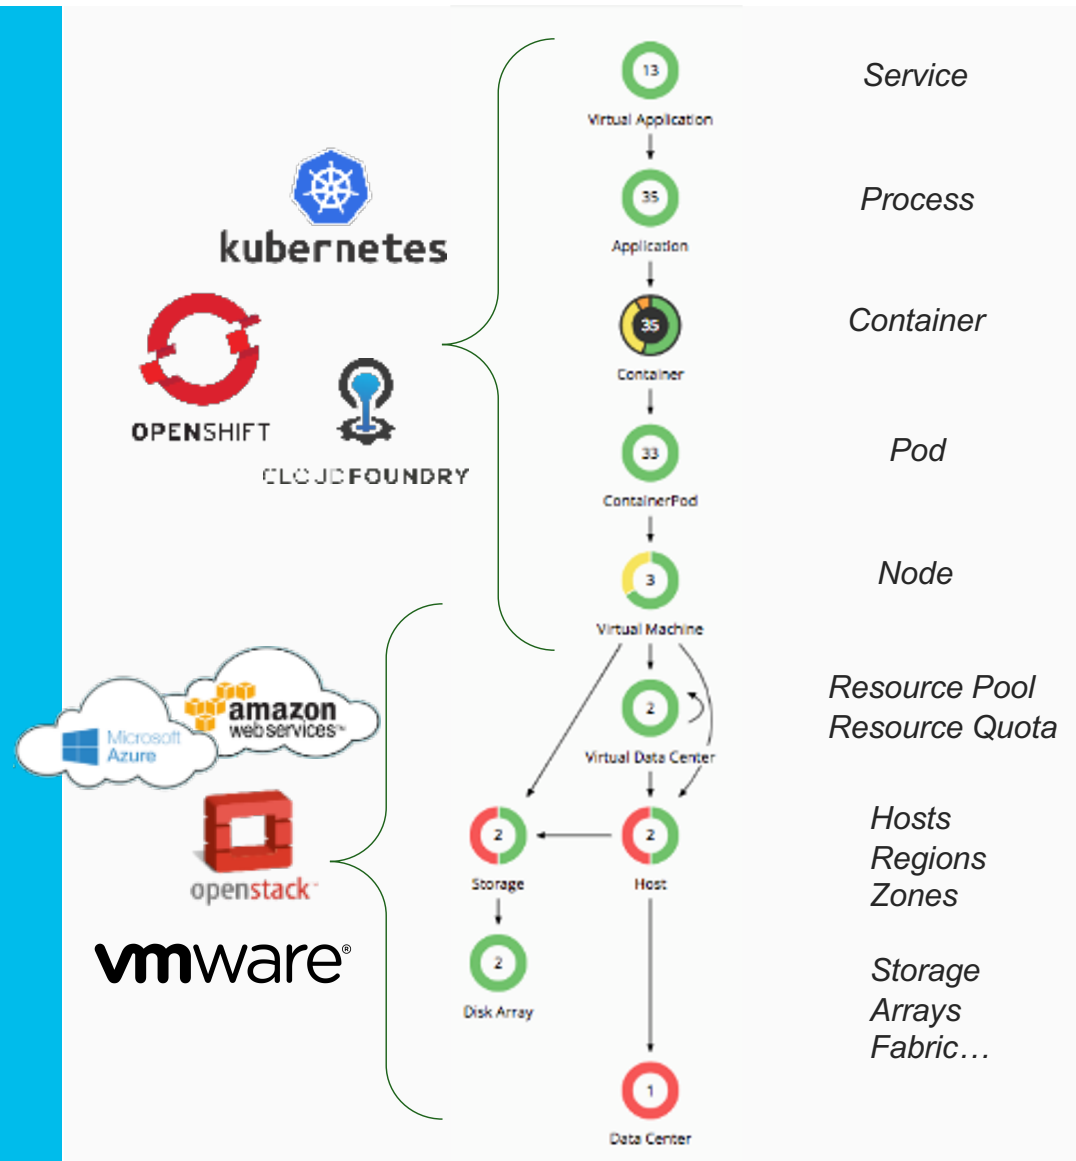
Self-Managing Container Platforms

Self-managing workloads optimize container platforms so IT organizations can scale and accelerate cloud native strategies.


Results:

- ✓ Minimal human intervention – no thresholds to set!
- ✓ Automated rescheduling of pods assures performance
- ✓ Intelligent cluster scaling ensures elastic infrastructure
- ✓ Full-stack control unites DevOps and Infrastructure

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Partner Confidential




The Need for Continuous Placement: Pod Rescheduling





Avoid Noisy Neighbor

- Workload that always peaks together



Avoid Performance Degradation

- CPU starvation due to node cpu congestion



Avoid long pending pod

- Resource Fragmentation

General 2.x Platform Enhancements

CCWOM 2.0 introduces an enhanced decision engine which will provide the following benefits:

- Greater speed in Real-Time Operations and Decisions
- Significantly faster run-time for environment modeling, simulations via plans
- CPU Capacity support: now understands differences between different generations of CPU architectures for more accurate planning and recommendations

Industry Standard SPEC CPU Benchmark

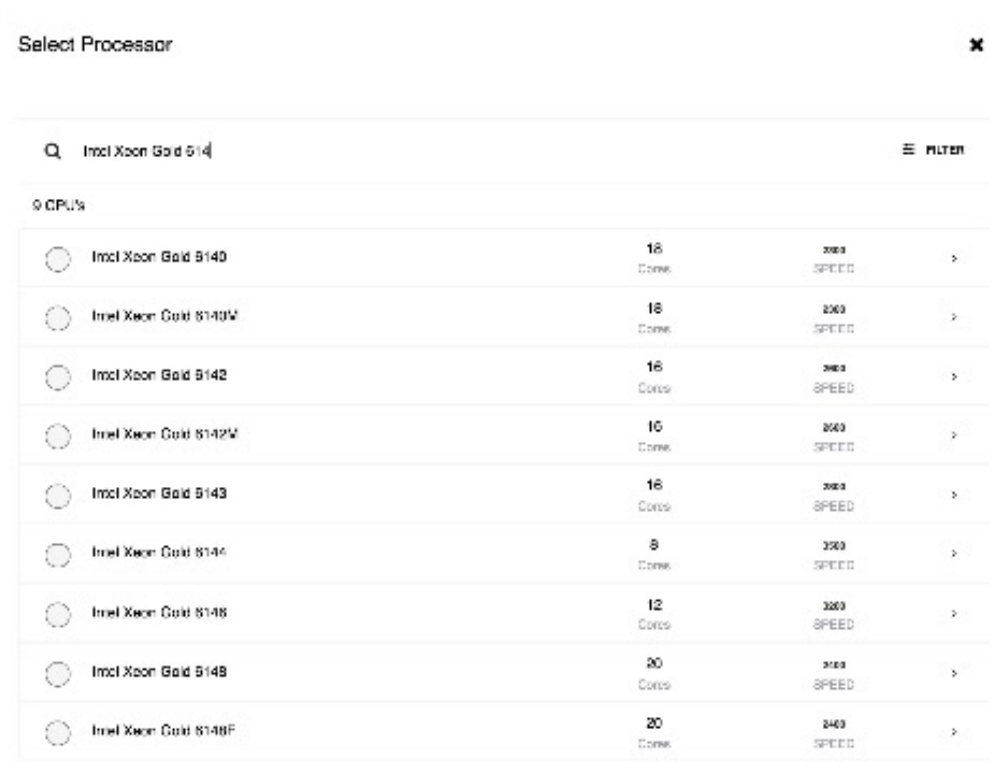

CCWOM 2.0 understands the differences between CPU generations/architectures

Vendors publish metrics

CCWOM gets periodic updates

| | B | C | E | F | H | I | J |
|-----|----------------------------|--|--------|------------------|-----------------------|---------------|----------|
| | Hardware Vendor | System | # CPUs | # Cores Per Chip | Processor | Processor MHz | Baseline |
| 826 | Dell Inc. | PowerEdge R730 (Intel Xeon X5687, 3.80 GHz) | 2 | 4 | Intel Xeon X5687 | 3500 | 45.1 |
| 827 | Dell Inc. | PowerEdge T330 (Intel Xeon E5-1270 v5, 3.60 GHz) | 1 | 4 | Intel Xeon E5-1270 v5 | 3500 | 72.7 |
| 828 | Dell Inc. | PowerEdge T330 (Intel Xeon E5-1270 v5, 3.60 GHz) | 1 | 4 | Intel Xeon E5-1270 v5 | 3500 | 69.9 |
| 829 | Dell Inc. | PowerEdge T630 (Intel Xeon X5687, 3.80 GHz) | 2 | 4 | Intel Xeon X5687 | 3500 | 45.1 |
| 830 | Dell Inc. | PowerEdge T710 (Intel Xeon X5687, 3.80 GHz) | 2 | 4 | Intel Xeon X5687 | 3500 | 45.1 |
| 831 | Hewlett-Packard Enterprise | ProLiant ML130 Gen9 (Intel Xeon E5-1270 v5) | 1 | 4 | Intel Xeon E5-1270 v5 | 3500 | 73.2 |
| 832 | Dell Systems | Dell UCS B230 M3 (Intel Xeon E5-2647 v2, 3.50 GHz) | 2 | 4 | Intel Xeon E5-2647 v2 | 3500 | 51 |
| 833 | Cisco Systems | Cisco UCS B230 M3 (Intel Xeon E5-2648 v2, 3.50 GHz) | 2 | 6 | Intel Xeon E5-2648 v2 | 3500 | 57.5 |
| 834 | Dell Systems | Dell UCS B230 M3 (Intel Xeon E5-2647 v2, 3.50 GHz) | 2 | 4 | Intel Xeon E5-2647 v2 | 3500 | 61.2 |
| 835 | Dell Systems | Dell UCS B230 M3 (Intel Xeon E5-2647 v2, 3.50 GHz) | 2 | 4 | Intel Xeon E5-2647 v2 | 3500 | 63.2 |
| 836 | Cisco Systems | Cisco UCS C220 M3 (Intel Xeon E5-2648 v2, 3.50 GHz) | 2 | 6 | Intel Xeon E5-2648 v2 | 3500 | 57.4 |
| 837 | Cisco Systems | Cisco UCS C220 M4 (Intel Xeon E5-2637 v3) | 2 | 4 | Intel Xeon E5-2637 v3 | 3500 | 61.1 |
| 838 | Cisco Systems | Cisco UCS C220 M4 (Intel Xeon E5-2637 v4) | 2 | 4 | Intel Xeon E5-2637 v4 | 3500 | 65.1 |
| 839 | Dell Systems | Dell UCS C240 M3 (Intel Xeon E5-2647 v2, 3.50 GHz) | 2 | 4 | Intel Xeon E5-2647 v2 | 3500 | 52.2 |
| 840 | Cisco Systems | Cisco UCS C240 M3 M3 (Intel Xeon E5-2648 v2, 3.50 GHz) | 2 | 6 | Intel Xeon E5-2648 v2 | 3500 | 57.6 |
| 841 | Cisco Systems | Cisco UCS C240 M3 (Intel Xeon E5-2637 v3 @ 3.20 GHz) | 2 | 4 | Intel Xeon E5-2637 v3 | 3500 | 60.2 |
| 842 | Dell Systems | Dell UCS C240 M3 (Intel Xeon E5-2647 v2, 3.50 GHz) | 2 | 4 | Intel Xeon E5-2647 v2 | 3500 | 64.5 |
| 843 | Dell Inc. | Dell Precision R6400 (Intel X5270, 3.50 GHz) | 2 | 2 | Intel Xeon X5270 | 3500 | 22.4 |
| 844 | Dell Inc. | Dell Precision R7400 (Intel Xeon X1270, 3.50 GHz) | 2 | 2 | Intel Xeon X1270 | 3500 | 26.3 |
| 845 | Dell Inc. | Dell Precision T3610 (Intel Xeon E5-1650 v2, 3.50 GHz) | 1 | 6 | Intel Xeon E5-1650 v2 | 3500 | 52.2 |
| 846 | Dell Inc. | PowerEdge 1950 II (Intel Xeon X5270, 3.50 GHz) | 2 | 2 | Intel Xeon X5270 | 3500 | 26.5 |
| 847 | Dell Inc. | PowerEdge 2950 II (Intel Xeon X1270, 3.50 GHz) | 2 | 2 | Intel Xeon X1270 | 3500 | 26.5 |
| 848 | Dell Inc. | PowerEdge 2950 II (Intel Xeon X1270, 3.50 GHz) | 2 | 2 | Intel Xeon X1270 | 3500 | 26.5 |
| 849 | Dell Inc. | PowerEdge P680 (Intel Xeon E5-2687 v4, 3.50 GHz) | 2 | 4 | Intel Xeon E5-2687 v4 | 3500 | 66 |
| 850 | Dell Inc. | PowerEdge M620 (Intel Xeon X1270, 3.50 GHz) | 2 | 2 | Intel Xeon X1270 | 3500 | 26.5 |

How to enable CPU Capacity



- To add CPU model to existing template:
 - Settings → Templates
 - Edit Template
 - Toggle “Select from Catalog”
 - Choose specific CPU
- Run any plan with this template



CWOM & UCS Manager

Optimize workloads across UCS estates with integrated knowledge of fabric, IO modules, power & cooling.

Context

- IT must move quickly to keep up with business needs.
- Programmable infrastructure increased IT agility.
- But, what about the workload?

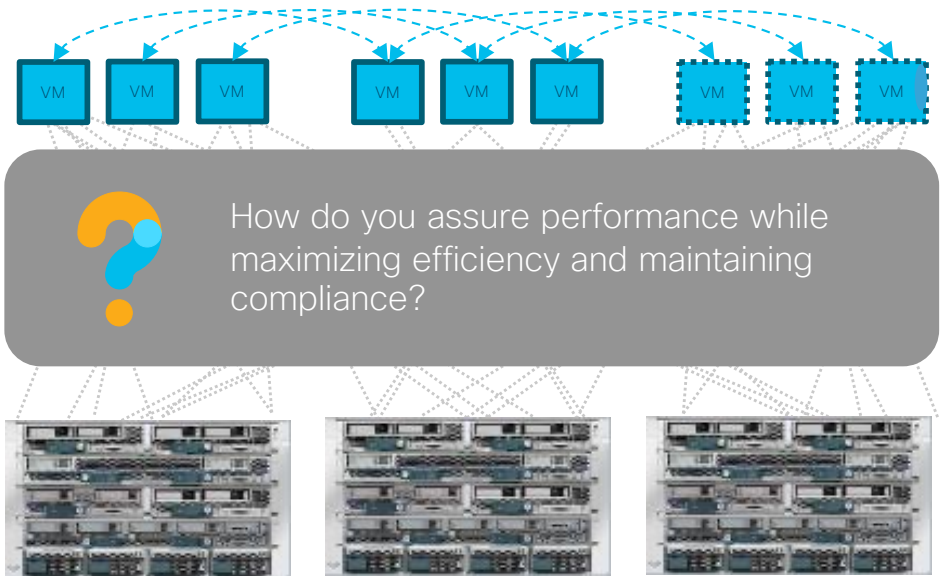


IT agility =

simplified infrastructure management
+
workload performance

Problem

Application consumption has implications from the virtualization layer to the physical layer, including fabric, IO modules, power & cooling.



What actions does it automate?

Documented benefits of CWOM automated on UCS...



- 33% increase in orders per minute
- 24% better response time

© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Partner Confidential

Enhanced VM intelligence when...

- Provisioning additional resources (vMem, vCPU)
- Moving Virtual Machine
- Moving Virtual Machine Storage
- Reconfiguring Storage
- Reconfiguring Virtual Machine

Physical Machines

- Start Physical Machine
- Provision Physical Machine
- Suspend Physical Machine

Chassis

- Provision New Chassis (recommend only)

Fabric Interconnect

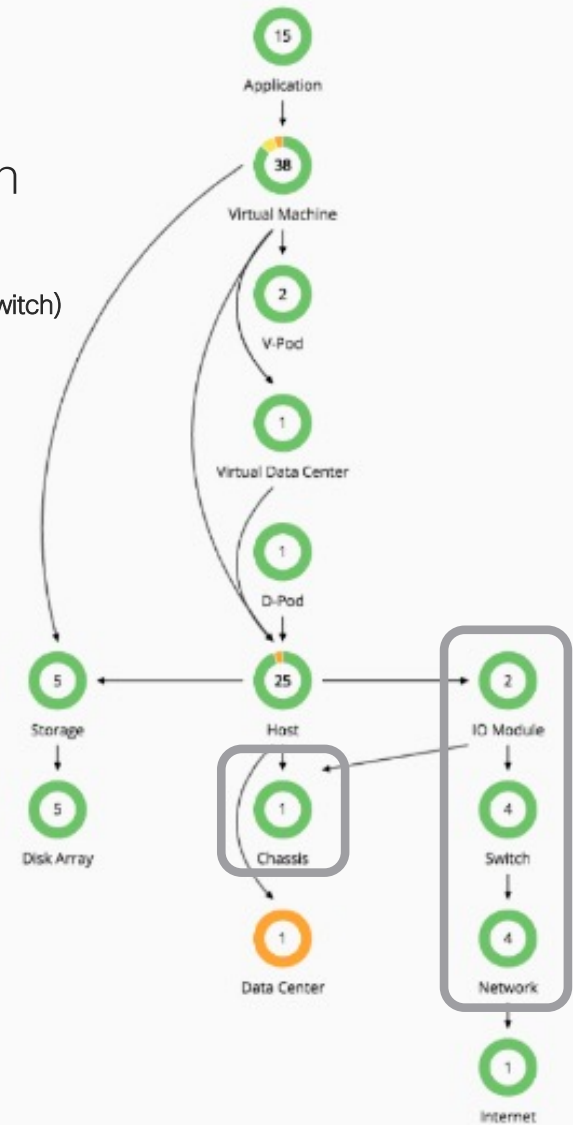
- Add Port to Port Channel (recommend only)
- Remove Port from Port Channel (recommend only)
- Add Port (recommend only)

How does CWOM make its decisions?

1

Extend the Supply Chain

- IO Module
- Fabric Interconnect (Switch)
- Domain (Network)
- Chassis



How does CWOM make its decisions? (cont'd)

2

Buy, sell, or hold commodities



Real-time Actions

Chassis

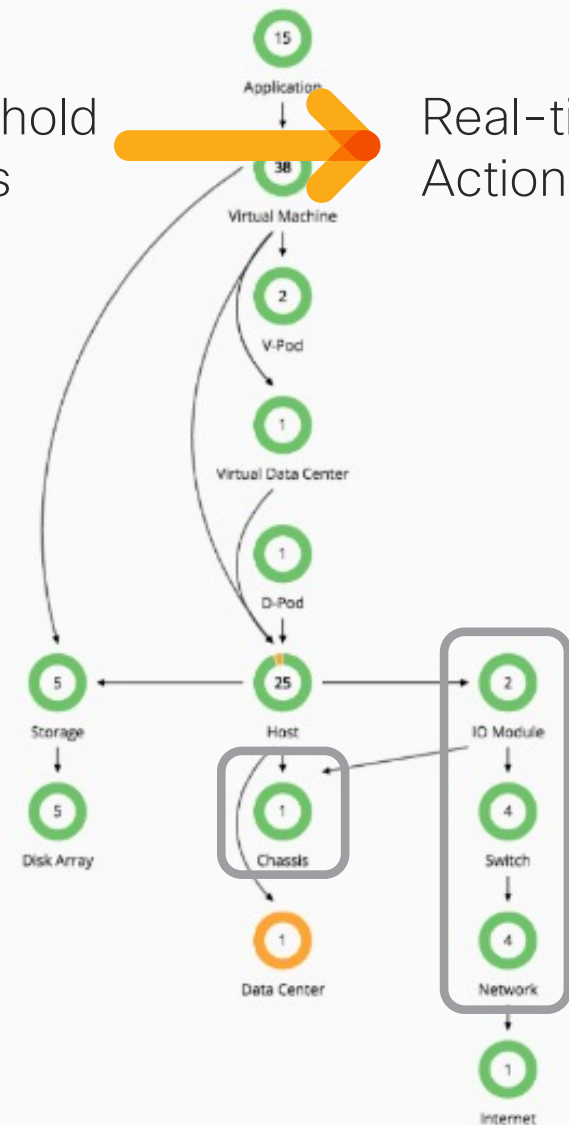
- Space
- Power
- Cooling

IO Module

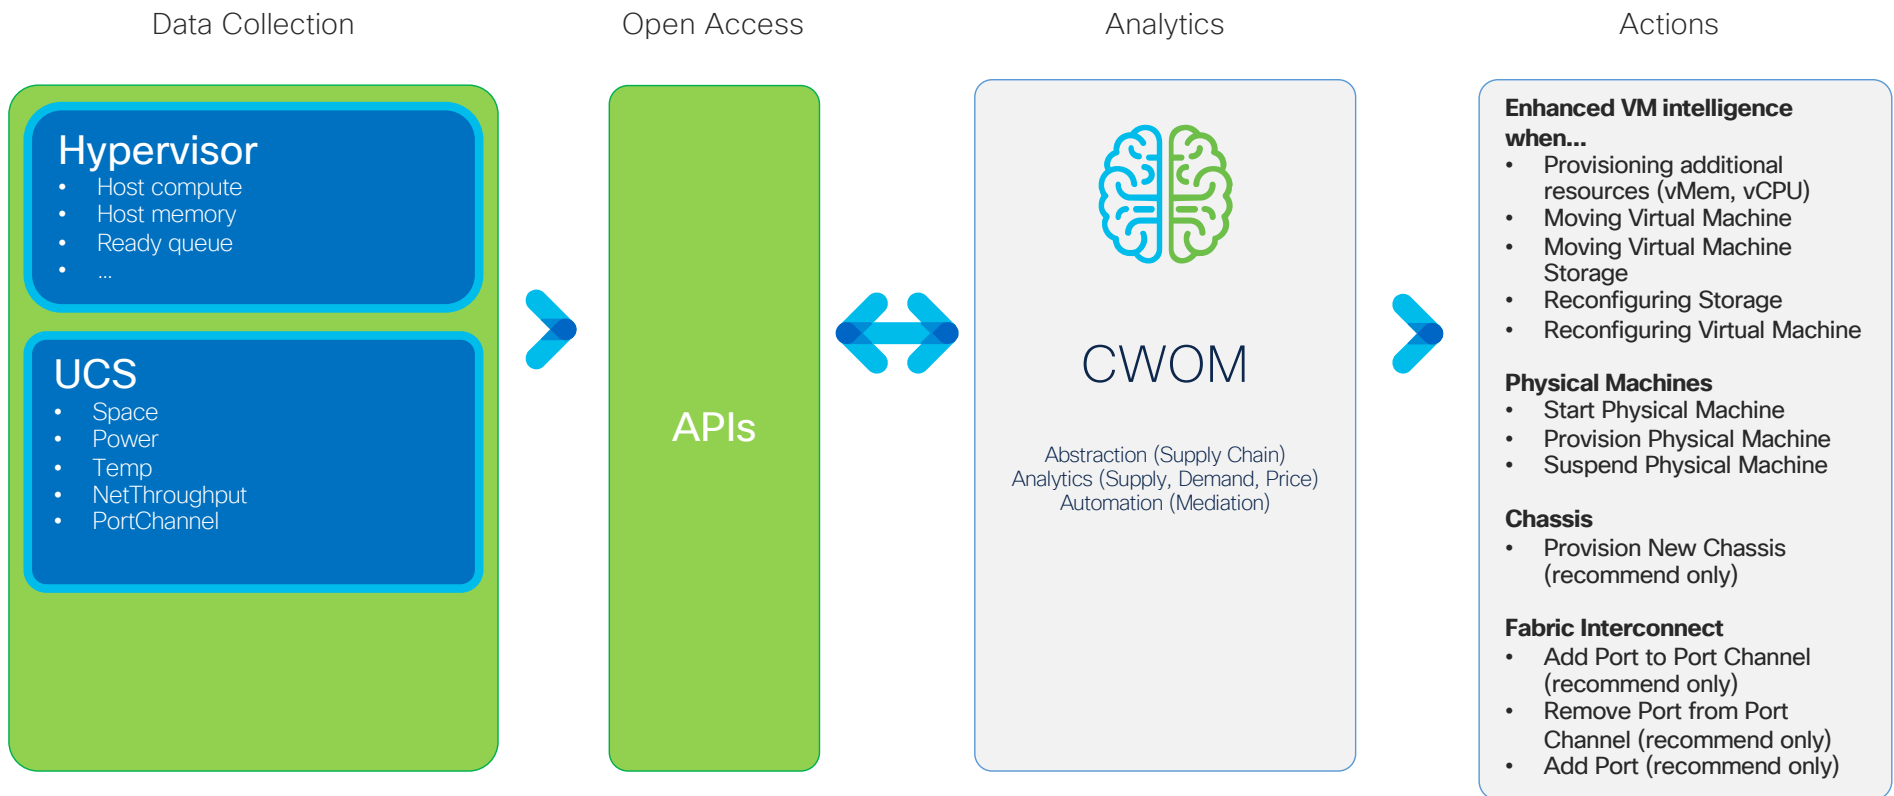
- NetThroughput
- PortChannel

Fabric Interconnect

- NetThroughput



How does it work together?





CWOM & UCS Director

End-to-end automation of the virtual, network, storage, and physical stack based on real-time workload demand.

Problem

Elastic real-time infrastructure requires two types of automation.

Process
Automation



Decision
Automation

What to do

When to do it? Why?



UCS Director



CWOM

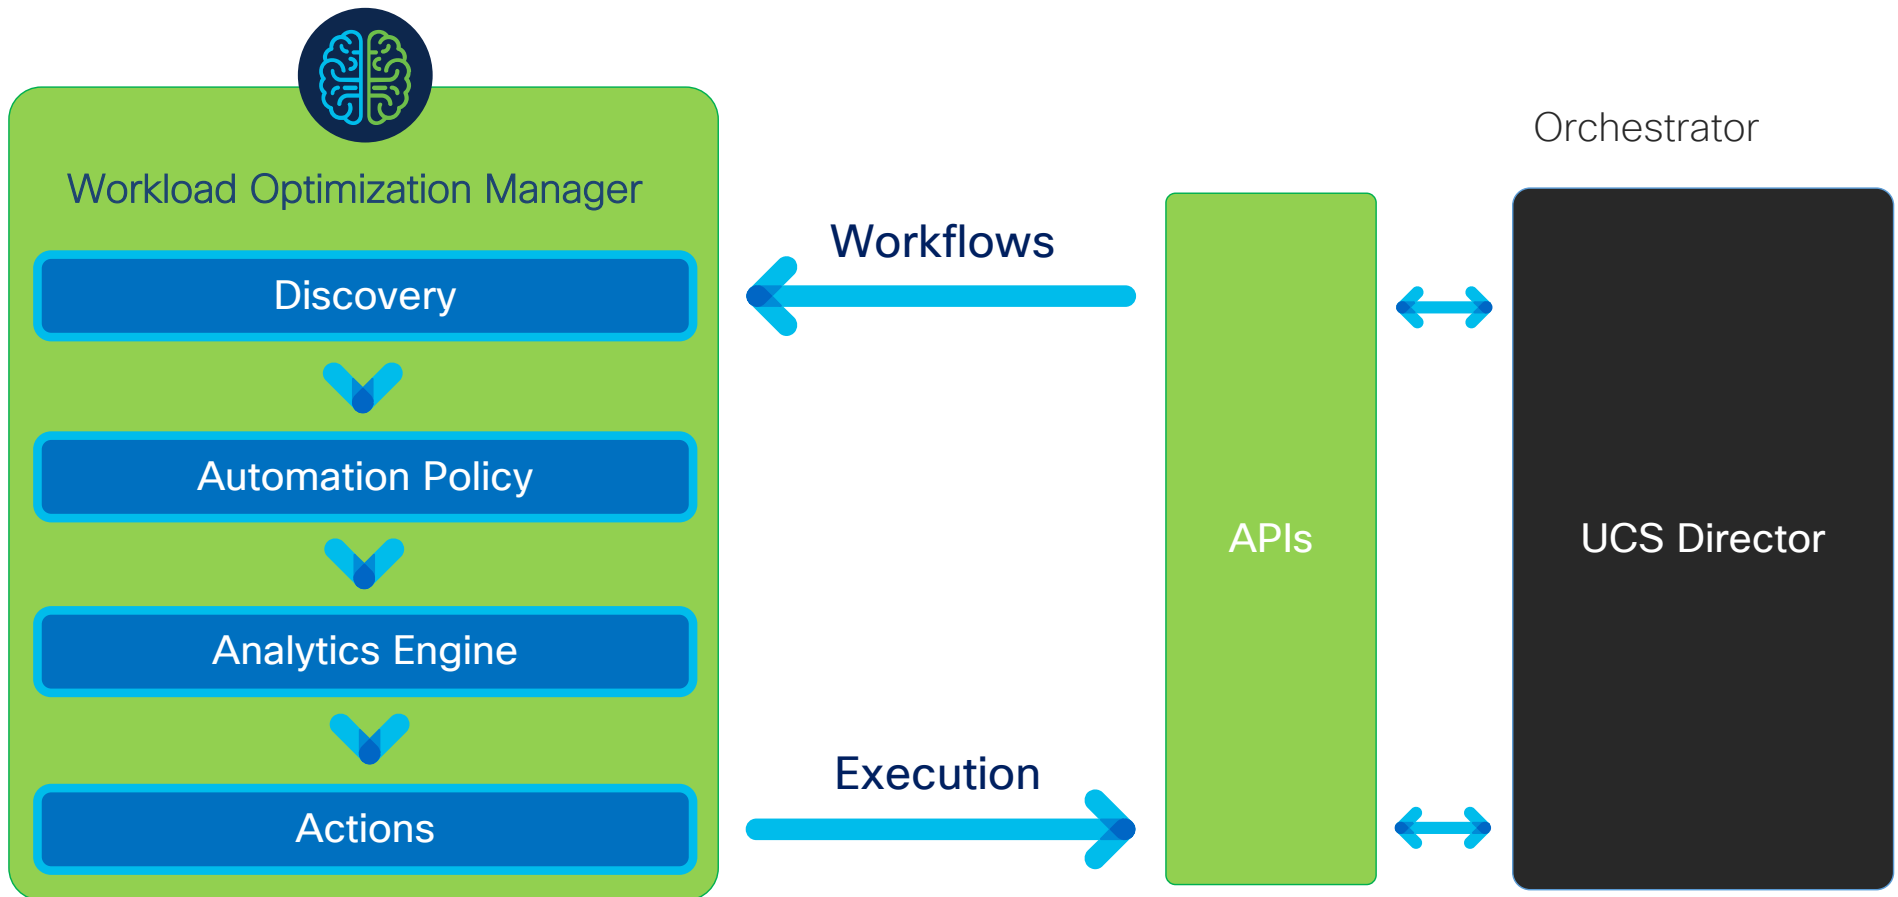
What actions does
it automate?

End-to-end automation...

- Provision new host
- Provision new storage
- Decommission host
- Decommission storage
- Resize storage

...based on real-time
workload demand.

How does it work together?






CWOM & Tetration

Network-aware optimization for distributed applications

Context



New application architectures increase east-west traffic and risk network congestion.



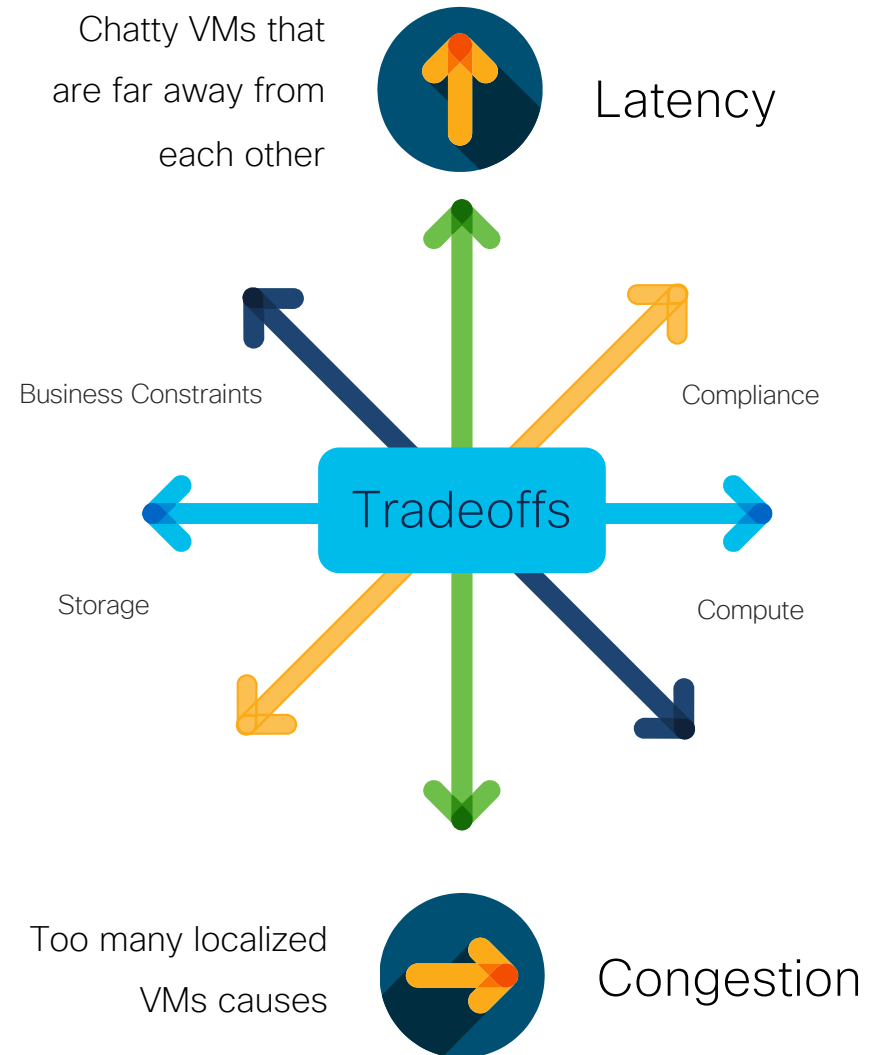
By 2020, 86% of data center traffic will be east-west...



Network-aware optimization reduces latency.

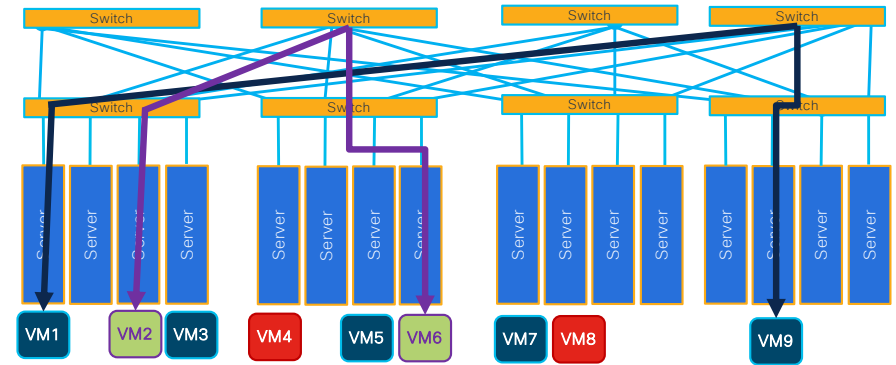
Problem

- East-west traffic growth
- More dynamic demand on networks

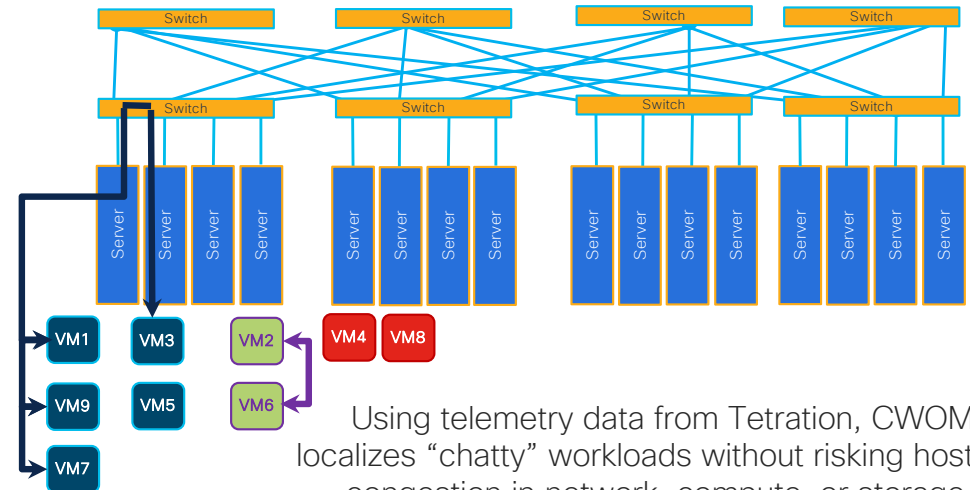


What is being optimized?

CWOM reduces latency by dynamically localizing “chatty” workloads without risking host congestion and while adhering to customer’s business requirements.



Chatty workloads across distances = latency.



Using telemetry data from Tetration, CWOM localizes “chatty” workloads without risking host congestion in network, compute, or storage.

How does CWOM make its decisions?

1

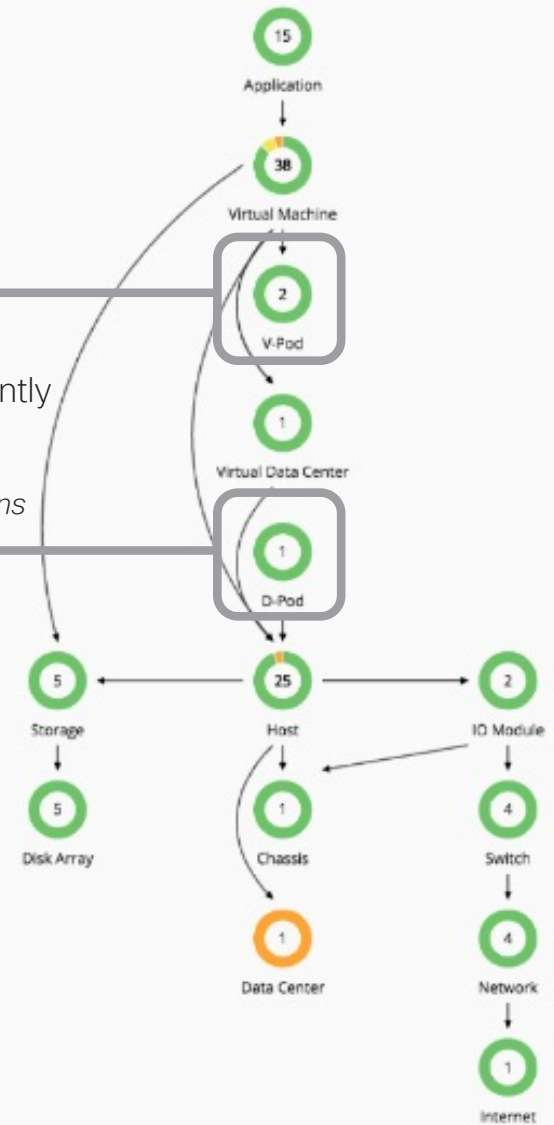
Extend the Supply Chain

V-Pod
Group of VMs that frequently communicates

Example: multi-tier applications

D-Pod
Group of Physical Resources that are closely located

Example: Servers in the same TOR Switch or UCS Domain



How does CWOM make its decisions?

2

Add Flow as a Commodity

Level 0: Intra-host
Flow that represents chatty
VMs in the same host



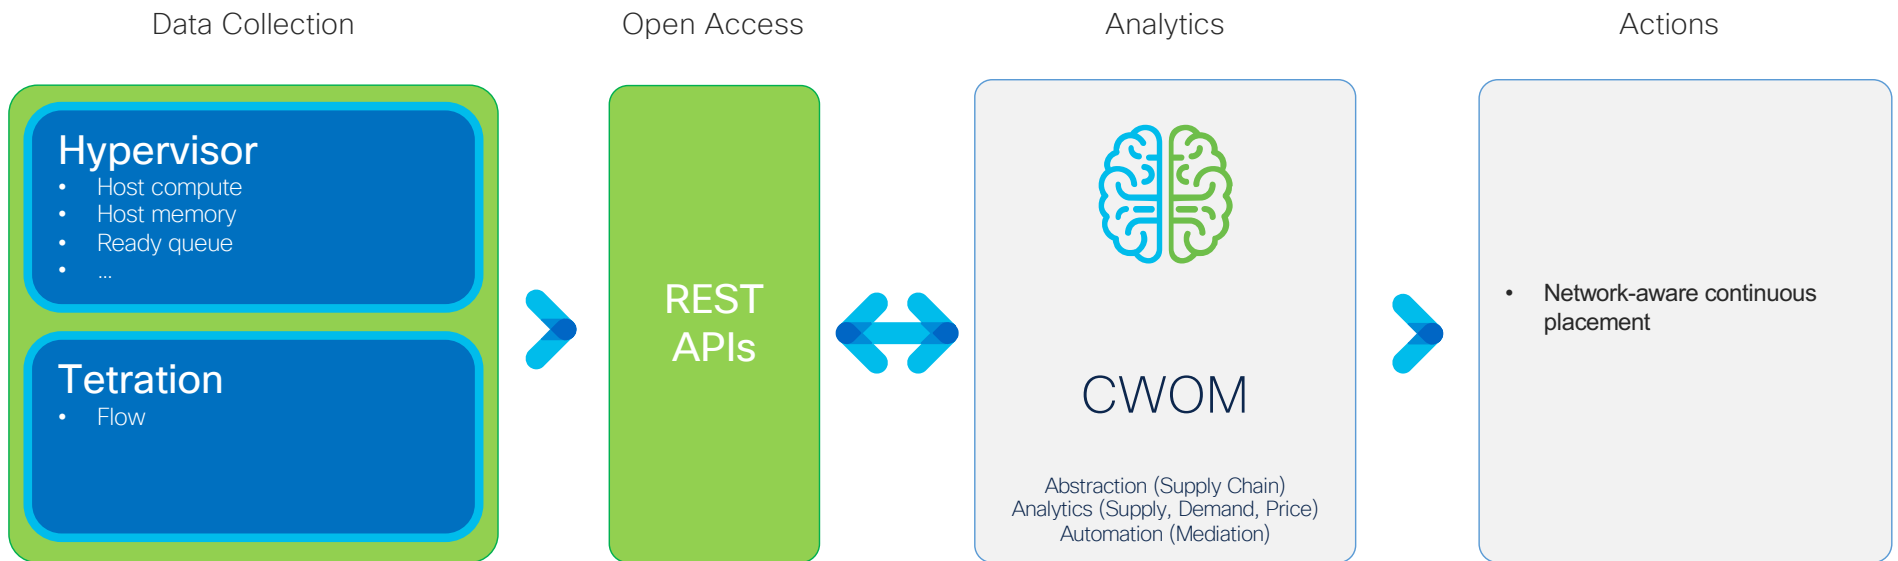
Level 1: Intra-D-Pod
Flow that represents chatty
VMs in the same D-Pod



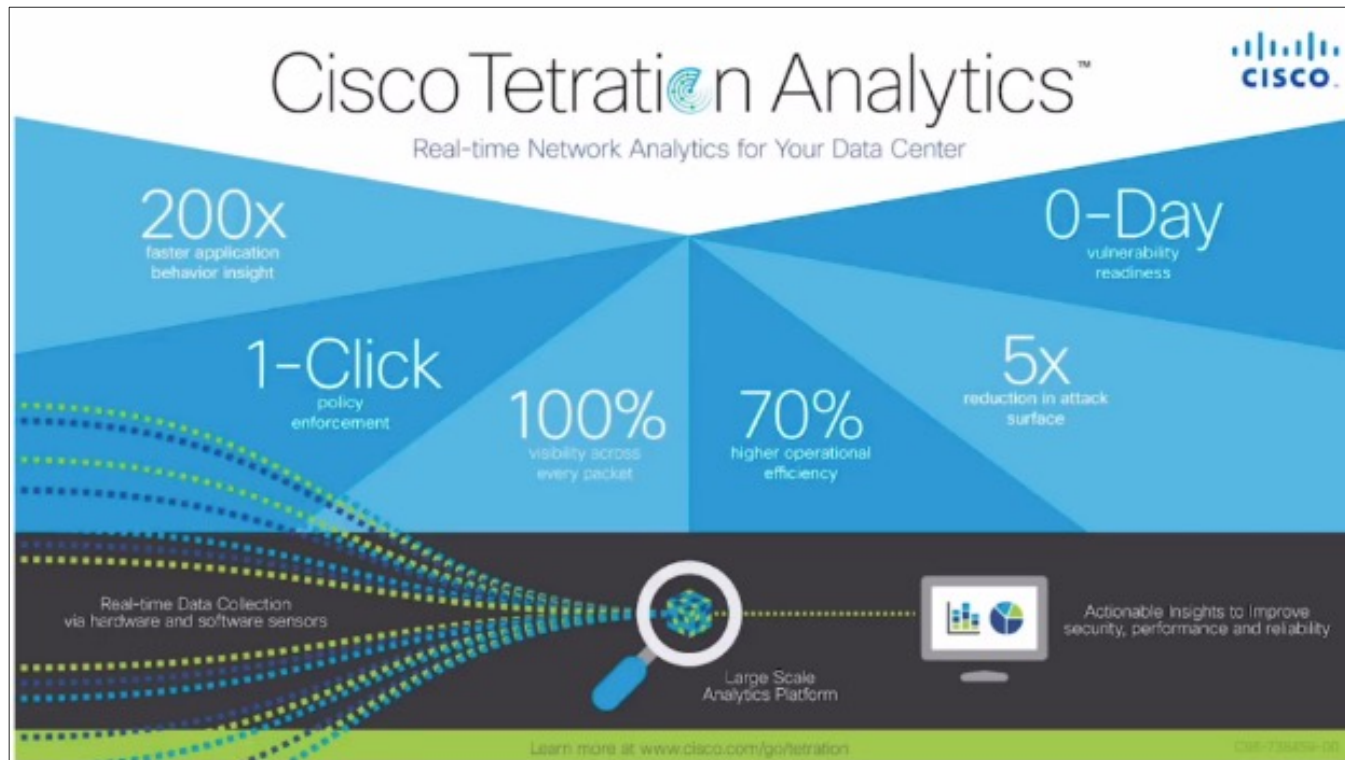
Level 2: Cross-D-Pod
Flow that represents chatty
VMs across D-Pods



How does it work together?



Demo: Tetration and CWOM



[Click here](#)



CWOM & AppDynamics

End-to-end performance optimization, from applications to infrastructure.

Context

Applications **are** the business.

62%

of users have **increasing expectations** of how well digital services should perform.



The *busiest bank branch* is the mobile app on the 7:15 a.m. train



Pearson

Education publishing is no longer about text books



U B E R

The world's largest *taxi company* owns no vehicles



A technology business that happens to be in the betting industry.

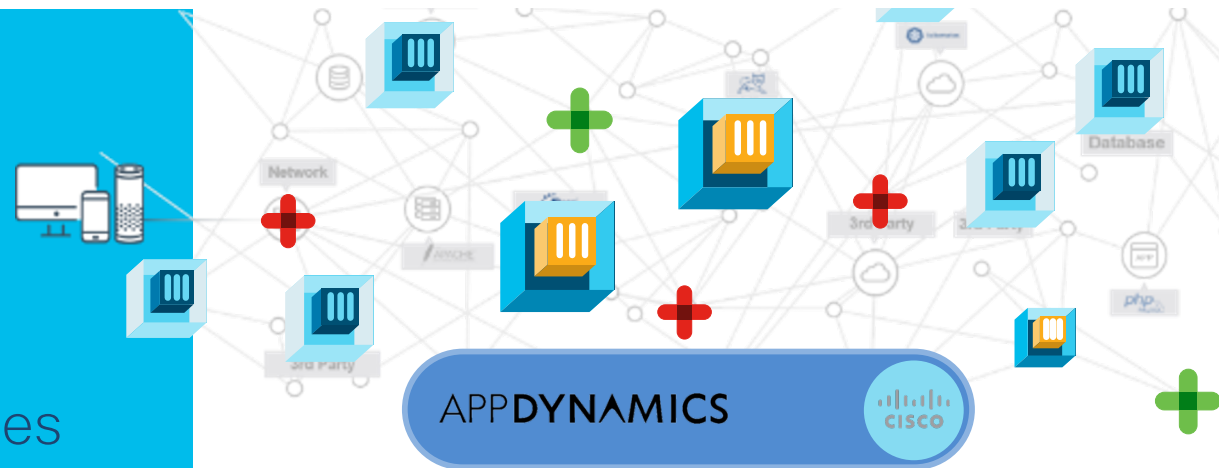
Application performance relies on two interdependent things...

1


Apps are written and architected well.

2

Apps get the resources they need when they need them.

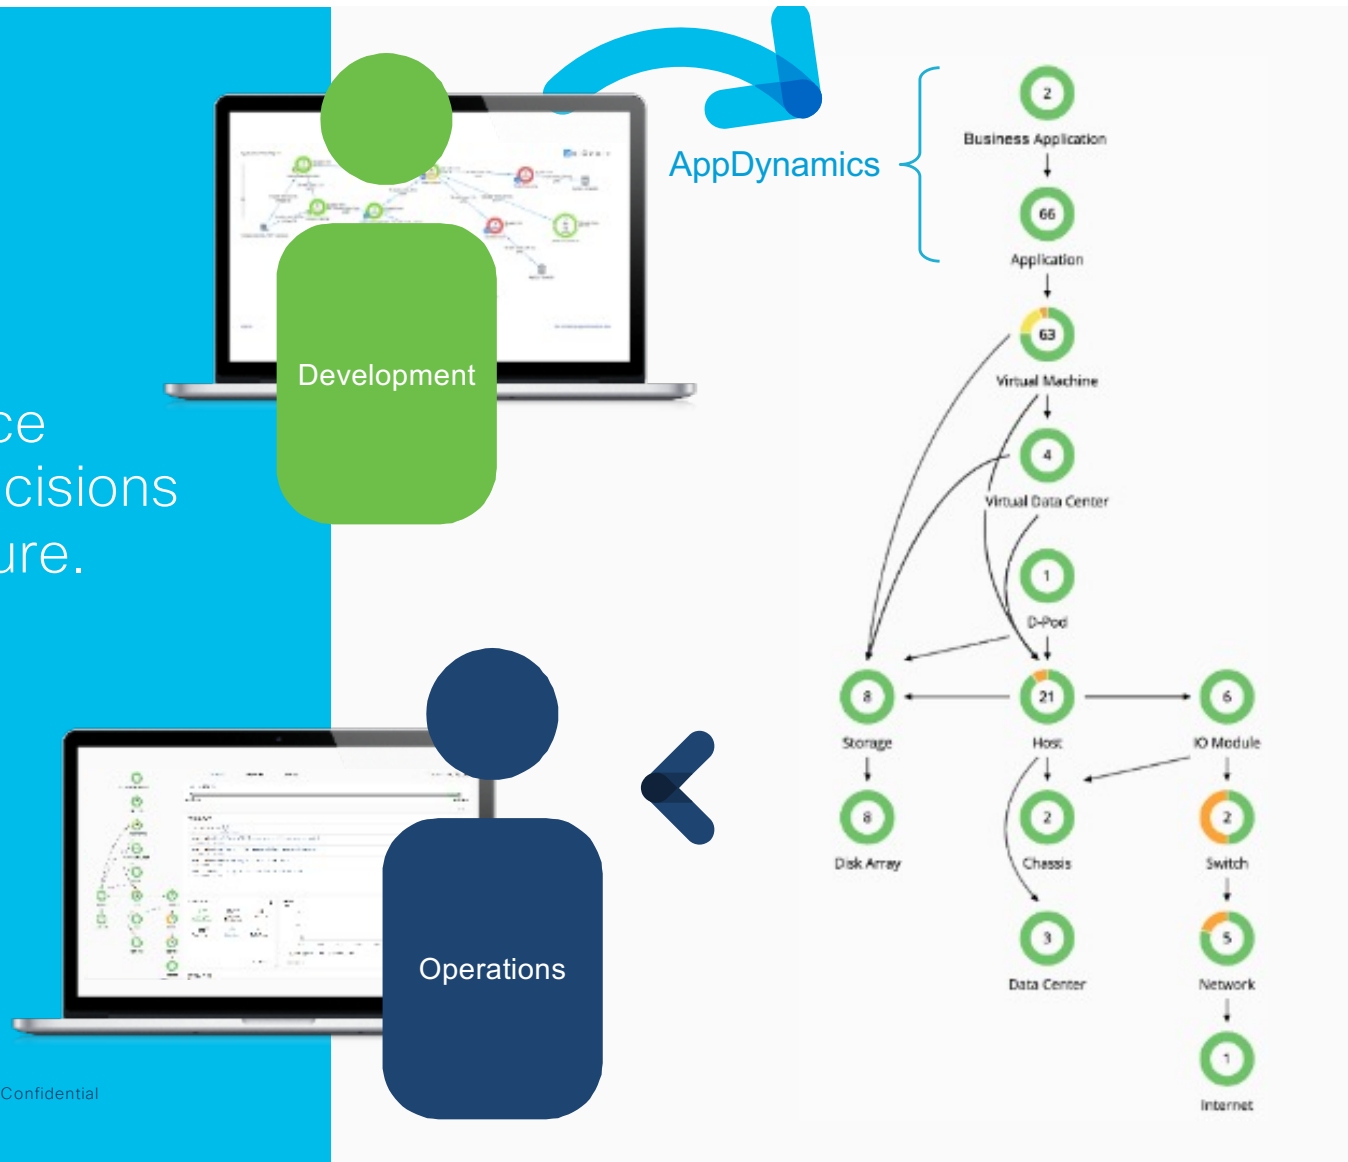


Workload Optimization
Powered by Turbonomic



What's being optimized?

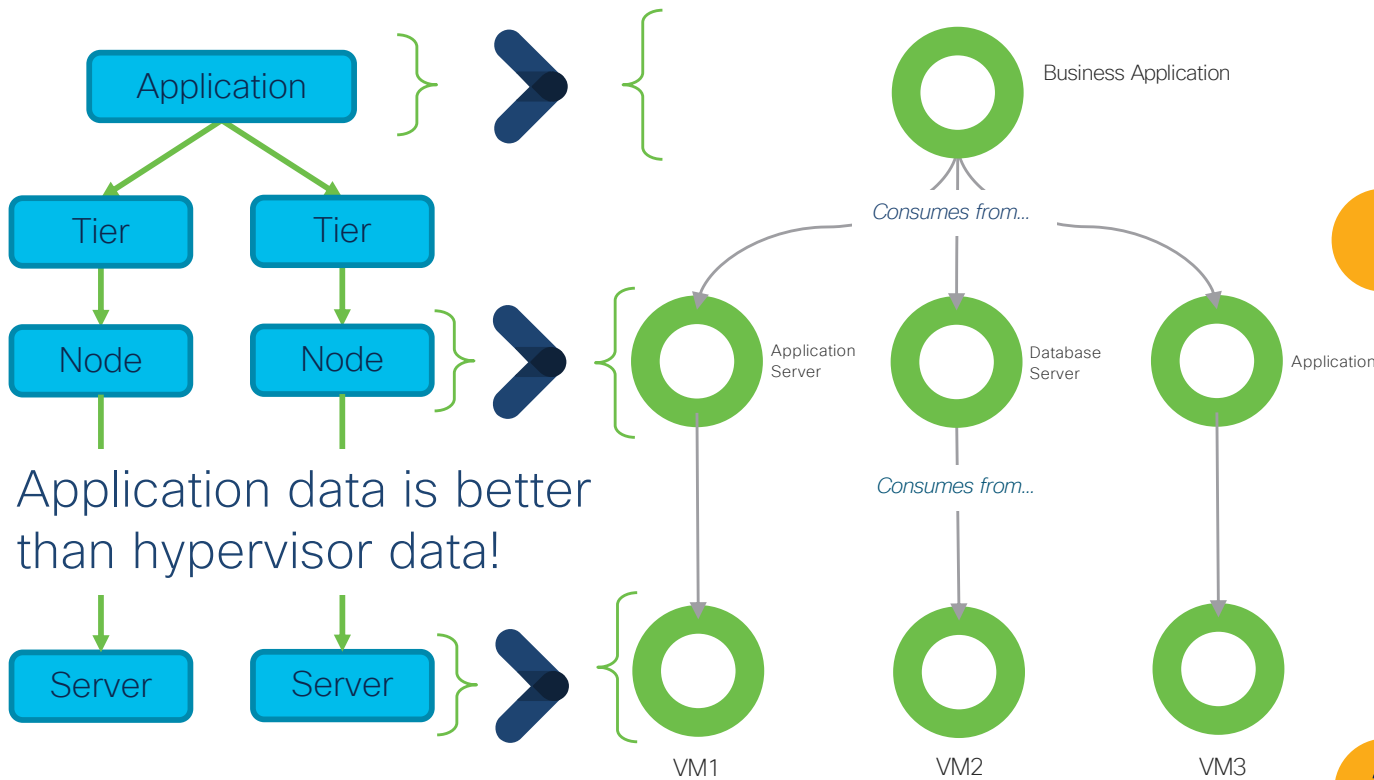
Application performance metrics drive better decisions through the infrastructure.



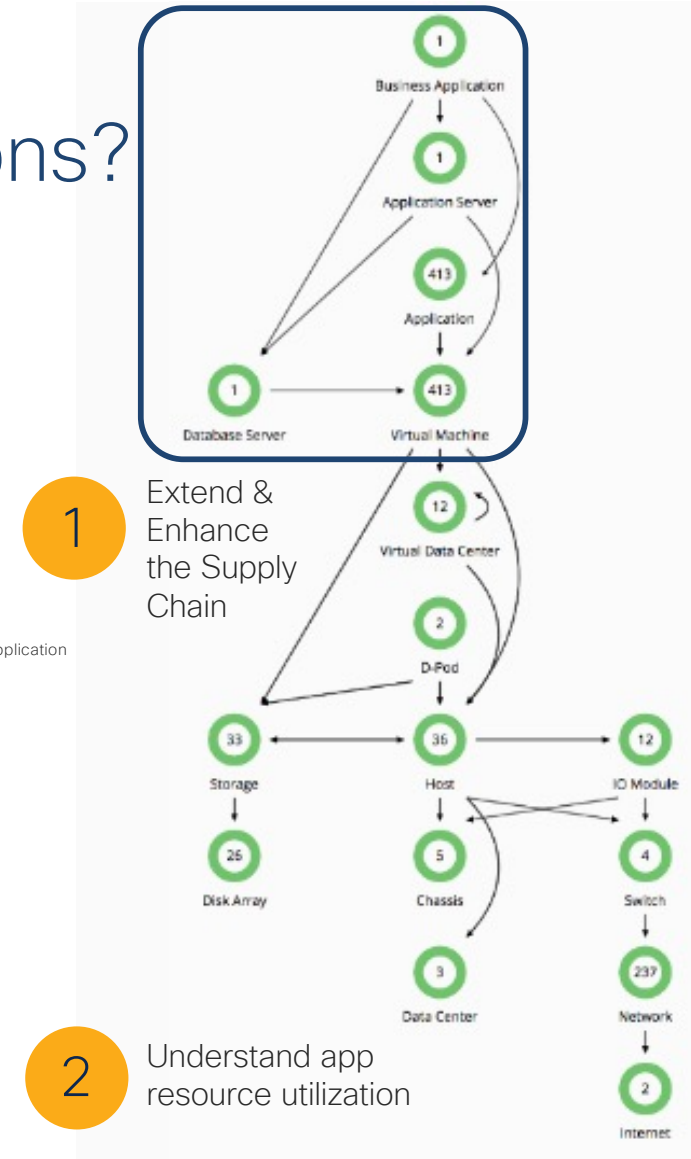
How does CWOM make its decisions?

AppDynamics terminology:

CWOM supply chain:



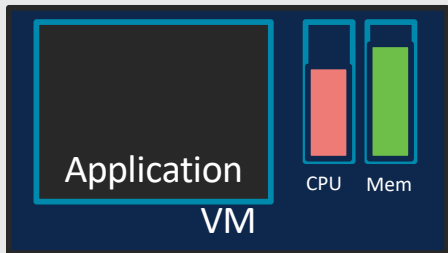
Application data is better than hypervisor data!



Cisco CWOM Integration with AppDynamics

Application-Level Understanding Enables Trustworthy Decisions, Automation

Traditional Infrastructure Monitoring



Application is a black box

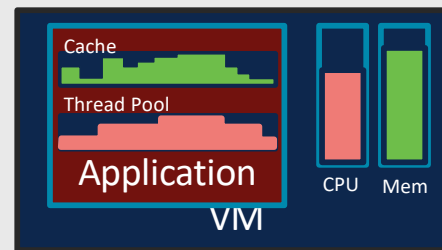


Components are a black box




No concept of application performance


Application Aware Infrastructure



Understands application resource utilization



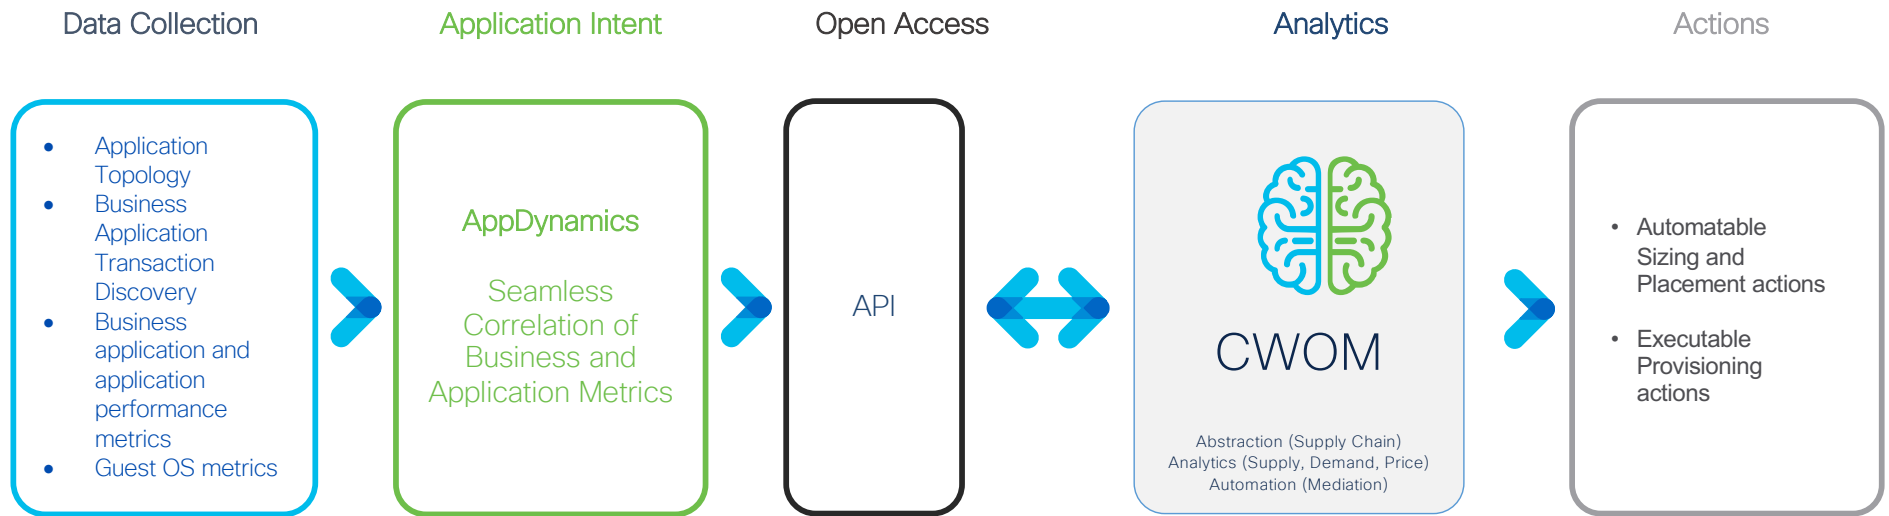
Understands application composition



Understands application performance

Architectural Overview

Application Metrics drive **sizing**, **placement**, and **capacity** actions across the stack.



Benefits



Assure App Performance

Eliminate applications performance risk due to infrastructure, without over-provisioning.



Show IT's Business Value

Resource decisions in the infrastructure are directly tied to the performance of business critical applications.



Bridge the App-Infra Gap

Full-stack automation elevates teams, full-stack visibility provides a common understanding of app dependencies.



Accelerate App Migration

De-risk app migration projects with holistic understanding of app topology and the data center stack.

Demo Video: AppDynamics & CWOM

- Just Turbo-AppD Demo Video: <https://vimeo.com/238472559>

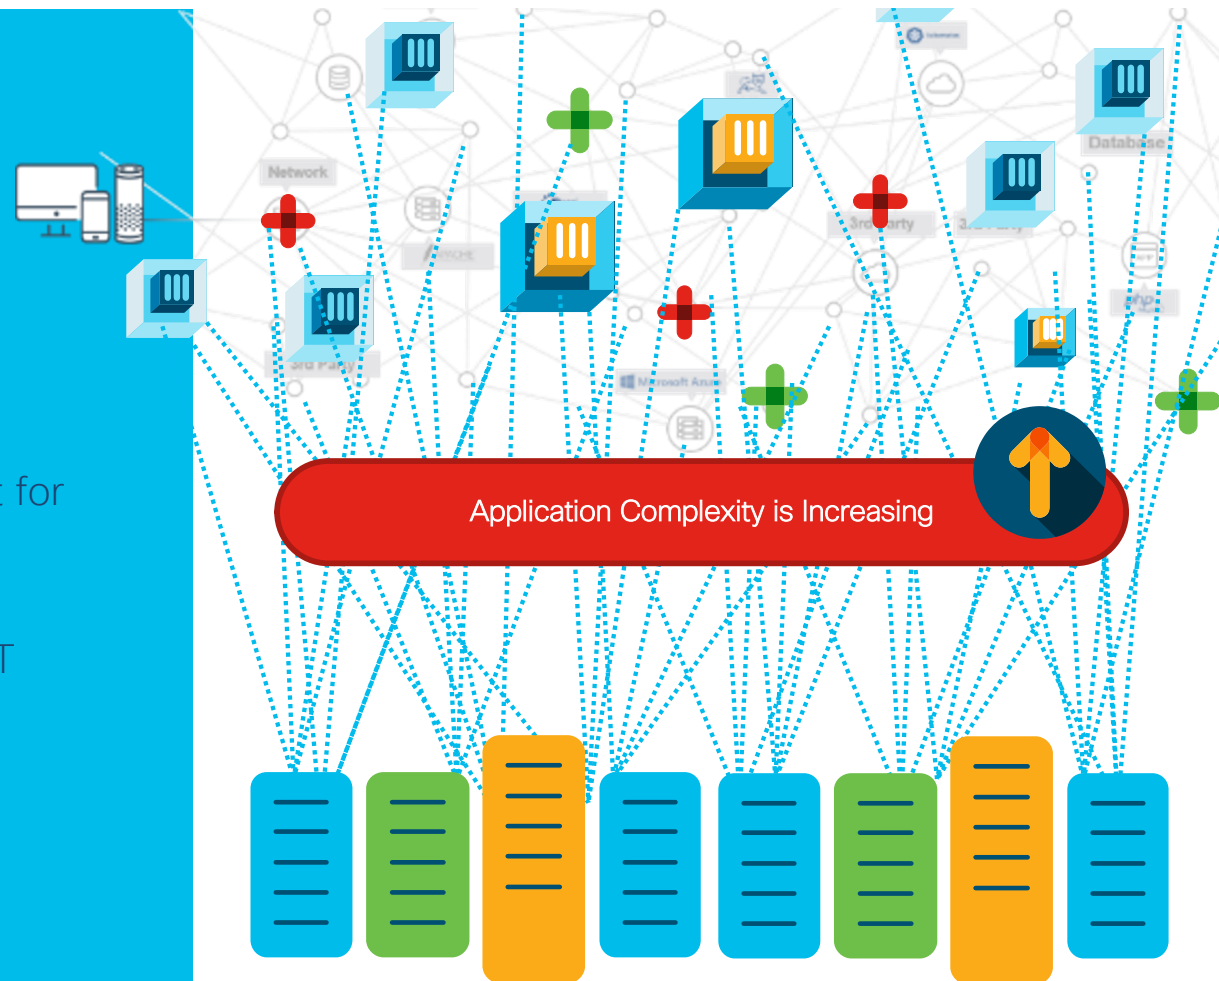


CWOM & HyperFlex

Achieving Agility & Predictable Performance with Self-Managing Hyper-Converged Systems

Context

- Application complexity is increasing...
- Cloud native & microservices are great for developers, but challenge IT agility.
- Hyper-converged is about simplifying infrastructure management to enable IT agility, but....



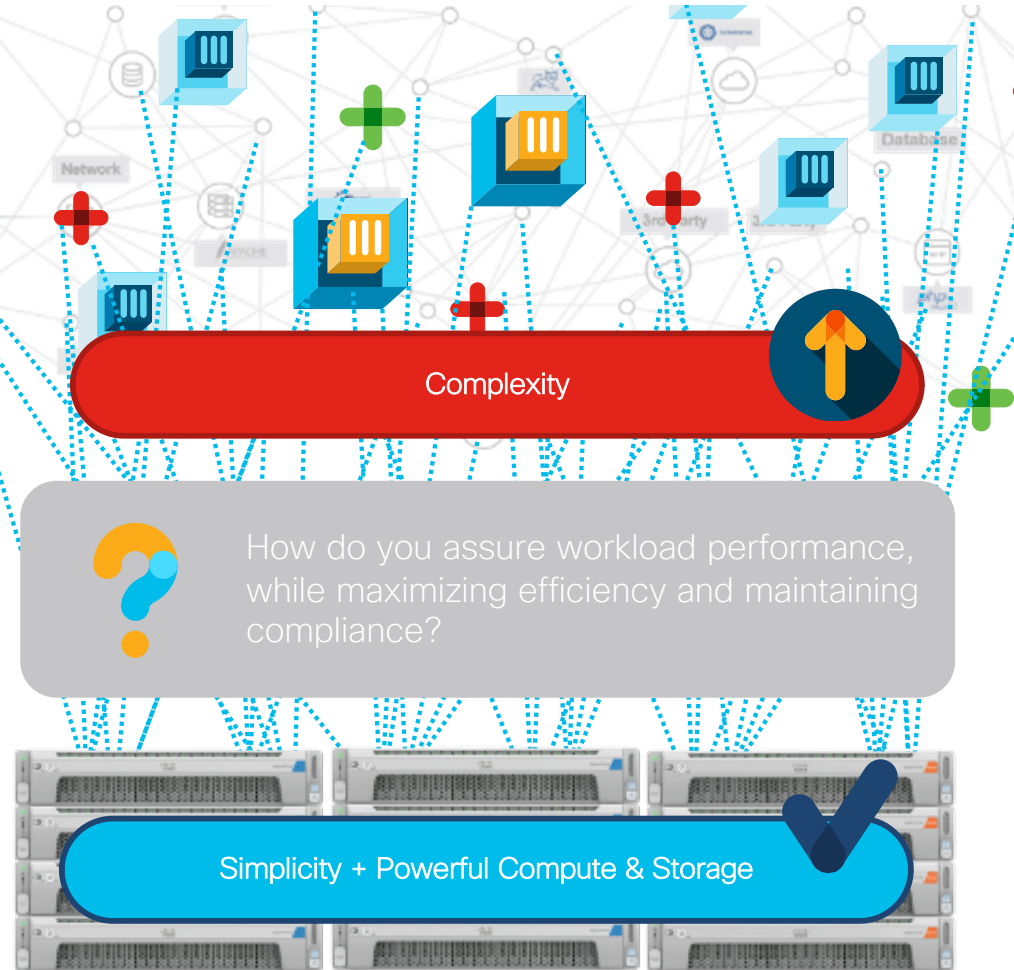
Problem

Agility requires two things:

Simplifying infrastructure management ✓

+

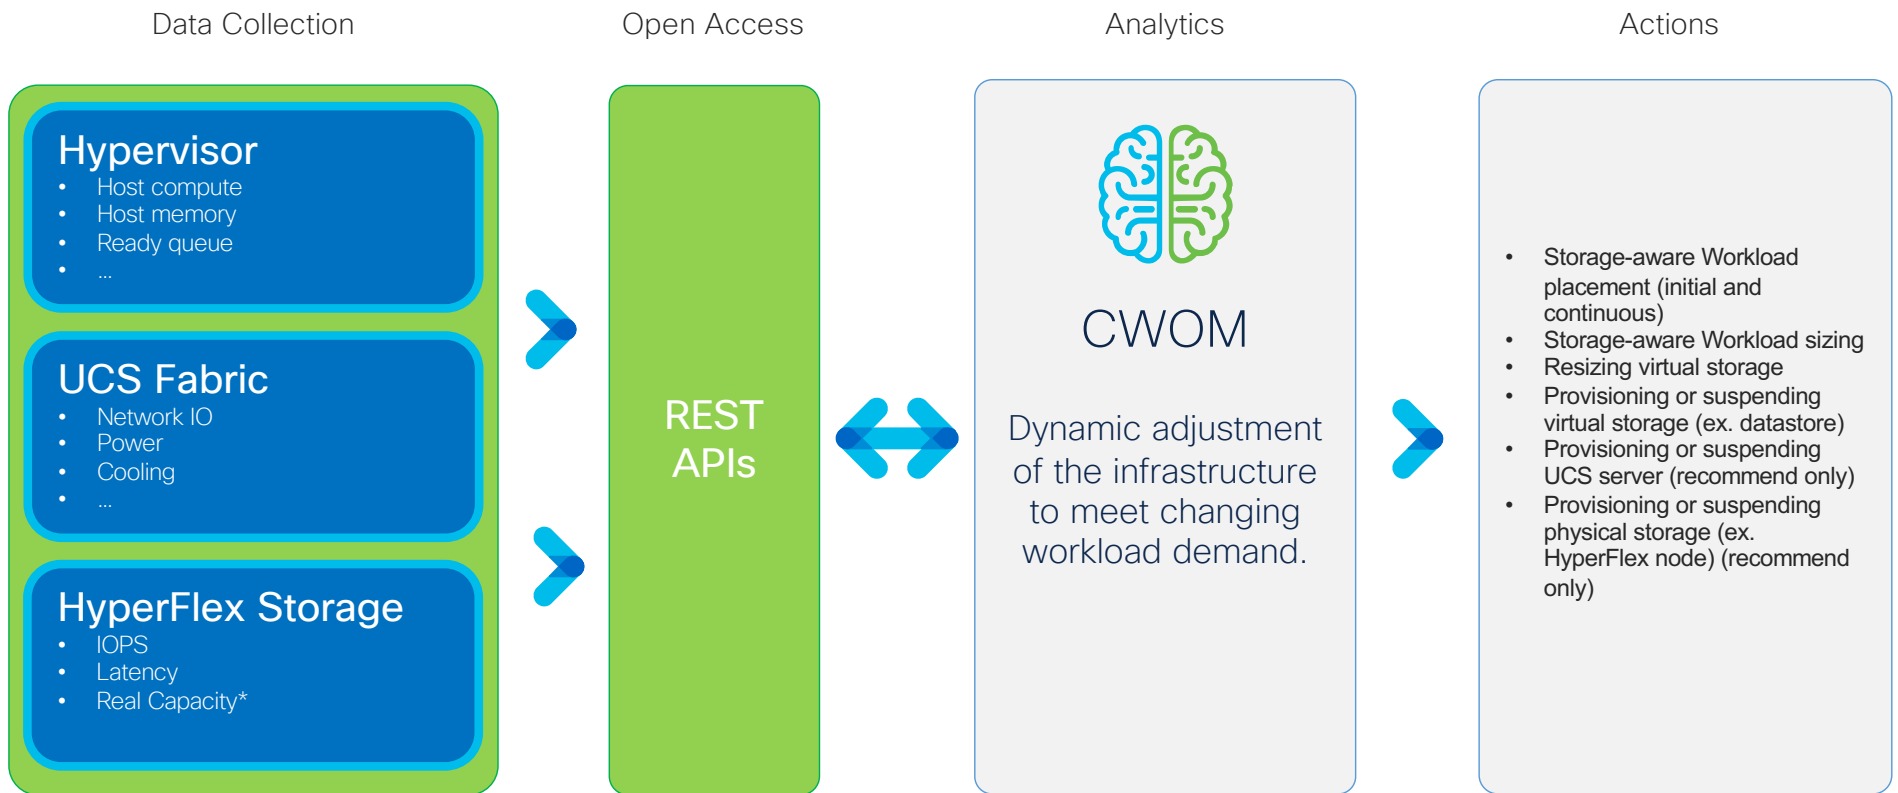
Performance ?



What actions does it automate?

- Storage-aware workload placement (initial and continuous)
- Storage-aware workload sizing
- Resizing virtual storage
- Provisioning or suspending virtual storage (ex. datastore)
- Provisioning or suspending HyperFlex compute-only nodes (recommend only)
- Provisioning or suspending physical storage (ex. HyperFlex node) (recommend only)

How does it work together?



Key Use Cases for CWOM & HyperFlex

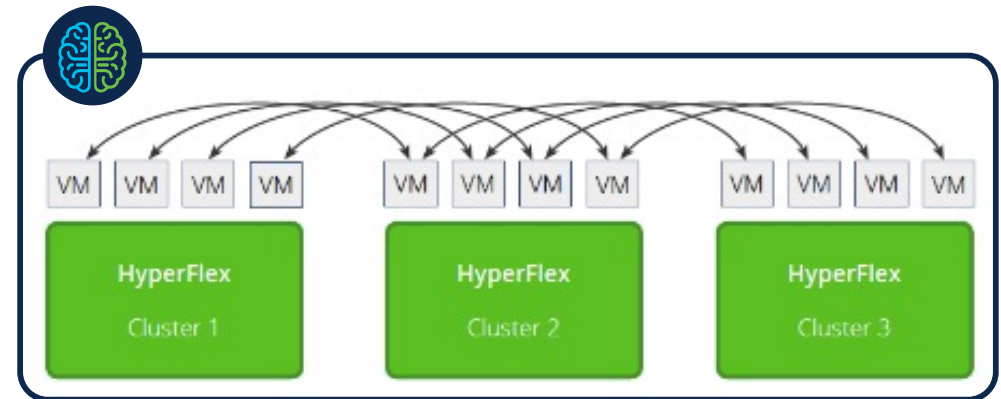
1. Super-Cluster Optimization
2. Intelligent Independent Scaling of Compute & Storage
3. Modernize at the Pace of Your Business



Self-managing HyperFlex systems continuously deliver predictable performance so IT can maximize agility.

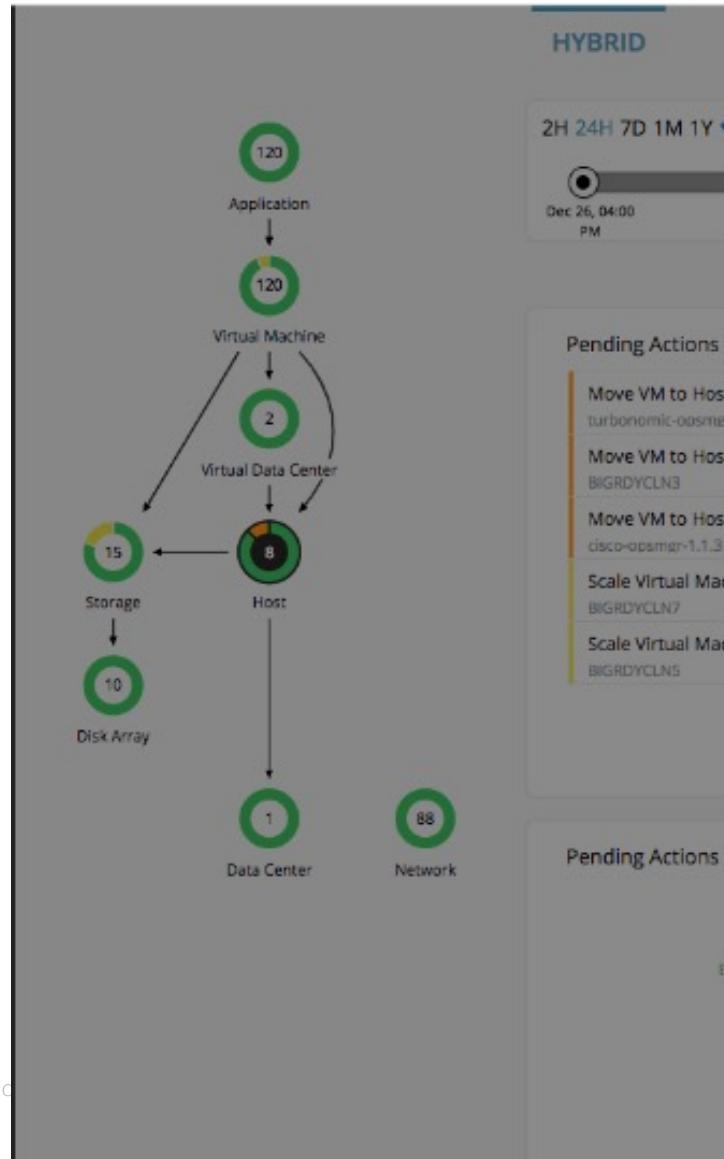
Use Case

Super-Cluster Optimization



- CWOM continuously analyzes workload consumption across HyperFlex clusters to determine exactly when, where, and how to move and/or resize existing workloads.
- Analysis accounts for different versions of HyperFlex (hybrid or all-flash) so there's no need for storage tiering.

CWOM provides a specific action to **move the virtual machine** 'cisco-opsmgr-1.1.3' from host '10.0.210.54' to host '10.0.210.52' to ensure workload performance.



23 Pending Actions

Move VM to Host PERFORMANCE ASSURANC
 cisco-opsmgr-1.1.3

Move Virtual Machine 'cisco-opsmgr-1.1.3' from Host '10.0.210.54' to Host '10.0.210.52', to improve workload distribution

| cisco-opsmgr-1.1.3 | | |
|---|---------------------|-------------------|
| VIRTUAL MACHINE | | |
| 1.42 → 1 | 1.7 % | 0 % |
| RISK INDEX | 598.3 GHz | 38.1 Mb/s |
| | CPU PROVISIONED | SWAPPING |
| 3.9 % | 0.4 % | 0.1 % |
| 383.9 GB | 3.7 TB | 20 sec |
| MEMORY | MEMORY PROVISIONED | Q4VCPU |
| 0 % | 0 % | 1.1 % |
| 5 TB | 78.1 Gb/s | 1.4 TB |
| STORAGE AMOUNT | NET THROUGHPUT | MEMORY ALLOCATION |
| 2.5 % | 0 % | 0 % |
| 10 GHz | 383.9 GB | 228,800.0 IOPS |
| VIRTUAL CPU | BALLOONING | STORAGE ACCESS |
| 0.1 % | 0 % | 0.8 % |
| 213.5 GHz | 3.9 Gb/s | 100 msec |
| CPU ALLOCATION | IO THROUGHPUT | STORAGE LATENCY |
| 1 → 0 | 0 % | 16.1 % |
| PRODUCES | 10 TB | 16 GB |
| | STORAGE PROVISIONED | VIRTUAL MEMORY |
| 0.4 % | 14.1 % | 1 |
| 59.8 GHz | 183.9 GB | HOST |
| CPU | VIRTUAL STORAGE | |

CWOM provides a specific action to **move the virtual machine 'turbo61hx' and storage 'FNTOP'** between clusters—from host '10.0.210.53' to host 'm4-esx1.cdnivt.cisco.com'—to ensure compliance with business policies.



34 Pending Actions

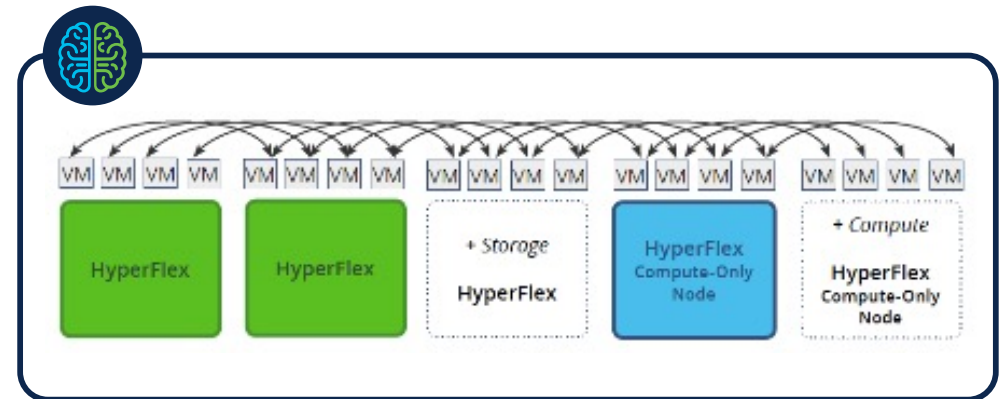
Move VM to Host turbo61hx COMPLIANCE

Move Virtual Machine 'turbo61hx' from Host '10.0.210.53' to Host 'm4-esx1.cdnivt.cisco.com' and Storage 'FNTOP', to ensure compliance with Workload Placement::Place Policy/m4-esx1.cdnivt.cisco.com

| turbo61hx | | |
|-----------------|--------------------|---------------------|
| VIRTUAL MACHINE | | |
| 1.2 → 1.18 | 1.7 % | 0 % |
| RISK INDEX | 598.3 GHz | 38.1 Mb/s |
| 4.1 % | 0.4 % | 0 % |
| CPU PROVISIONED | SWAPPING | |
| 383.9 GB | 3.7 TB | 20 sec |
| MEMORY | MEMORY PROVISIONED | Q4VCPU |
| 2.9 % | 0 % | 1.1 % |
| 5 TB | 78.1 Gb/s | 1.4 TB |
| STORAGE AMOUNT | NET THROUGHPUT | MEMORY ALLOCATION |
| 1 % → 1 % | 0 % | 0 % |
| 10 GHz | 383.9 GB | 228,800.0 IOPS |
| VIRTUAL CPU | BALLOONING | STORAGE ACCESS |
| 0 % | 0 % | 1.3 % |
| 213.5 GHz | 3.9 Gb/s | 100 msec |
| CPU ALLOCATION | IO THROUGHPUT | STORAGE LATENCY |
| 1 → 0 | 0 % | 1.5 % |
| PRODUCES | 4.9 Mbytes → | 10 TB |
| | 9.8 Mbytes | STORAGE PROVISIONED |
| | FLOW | |
| 8.8 % → 8.1 % | 0.2 % | 8.3 % |

Use Case

Intelligent Independent Scaling of Compute and Storage



- HyperFlex gives customers the flexibility to scale compute and storage independently based on their unique requirements.
- CWOM scales these resources based on real-time workload consumption, providing intelligent elasticity in HyperFlex clusters that preempts performance degradation.

Below, CWOM provides the action to provision a new disk array similar to 'CCHX' (a HyperFlex cluster) in order to assure performance.

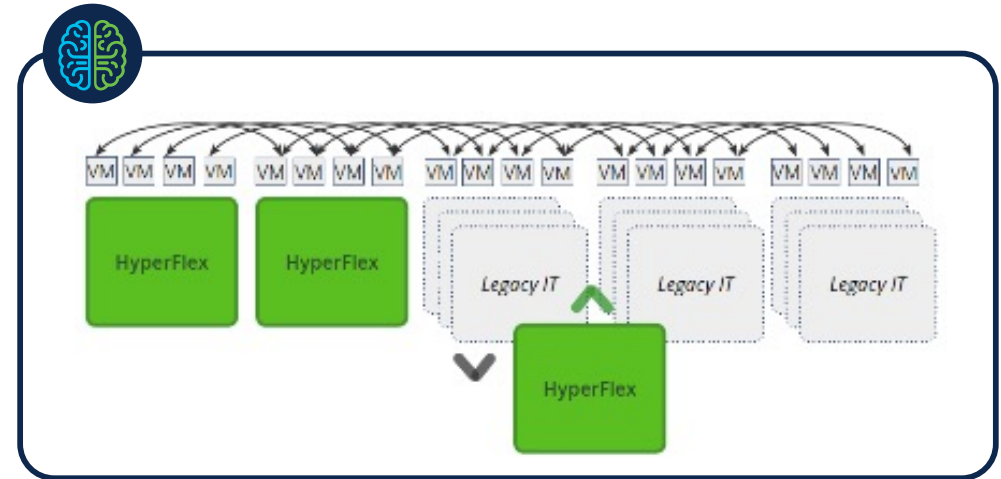
The screenshot displays the CWOM interface for a HyperFlex cluster. On the left, a dependency graph shows the following components and their relationships:

- Application (25) depends on Virtual Machine (25).
- Virtual Machine (25) depends on Virtual Data Center (1).
- Virtual Data Center (1) depends on Host (4).
- Host (4) depends on Storage (6).
- Storage (6) depends on Disk Array (1).

The main interface includes a navigation bar with 'HOME / CCHX', 'DAG', and 'CCHX' icons. A 'PLAN' button is visible in the top right. The main content area is titled 'OVERVIEW DETAILS POLICY ACTIONS (1)'. A timeline shows the current time as 'Dec 14, 06:00 PM' and a green bar indicates the state 'After Actions'. Below this, a 'Pending Actions' section lists 'Provision Disk Array Similar to CCHX' with a 'PERFORMANCE ASSURANCE' tag and a 'Show all' link. The bottom section contains three panels: 'Health' (green bar), 'Entity information' (ID: 48972875266470886445014959464179659105), and 'Related Group Information'.

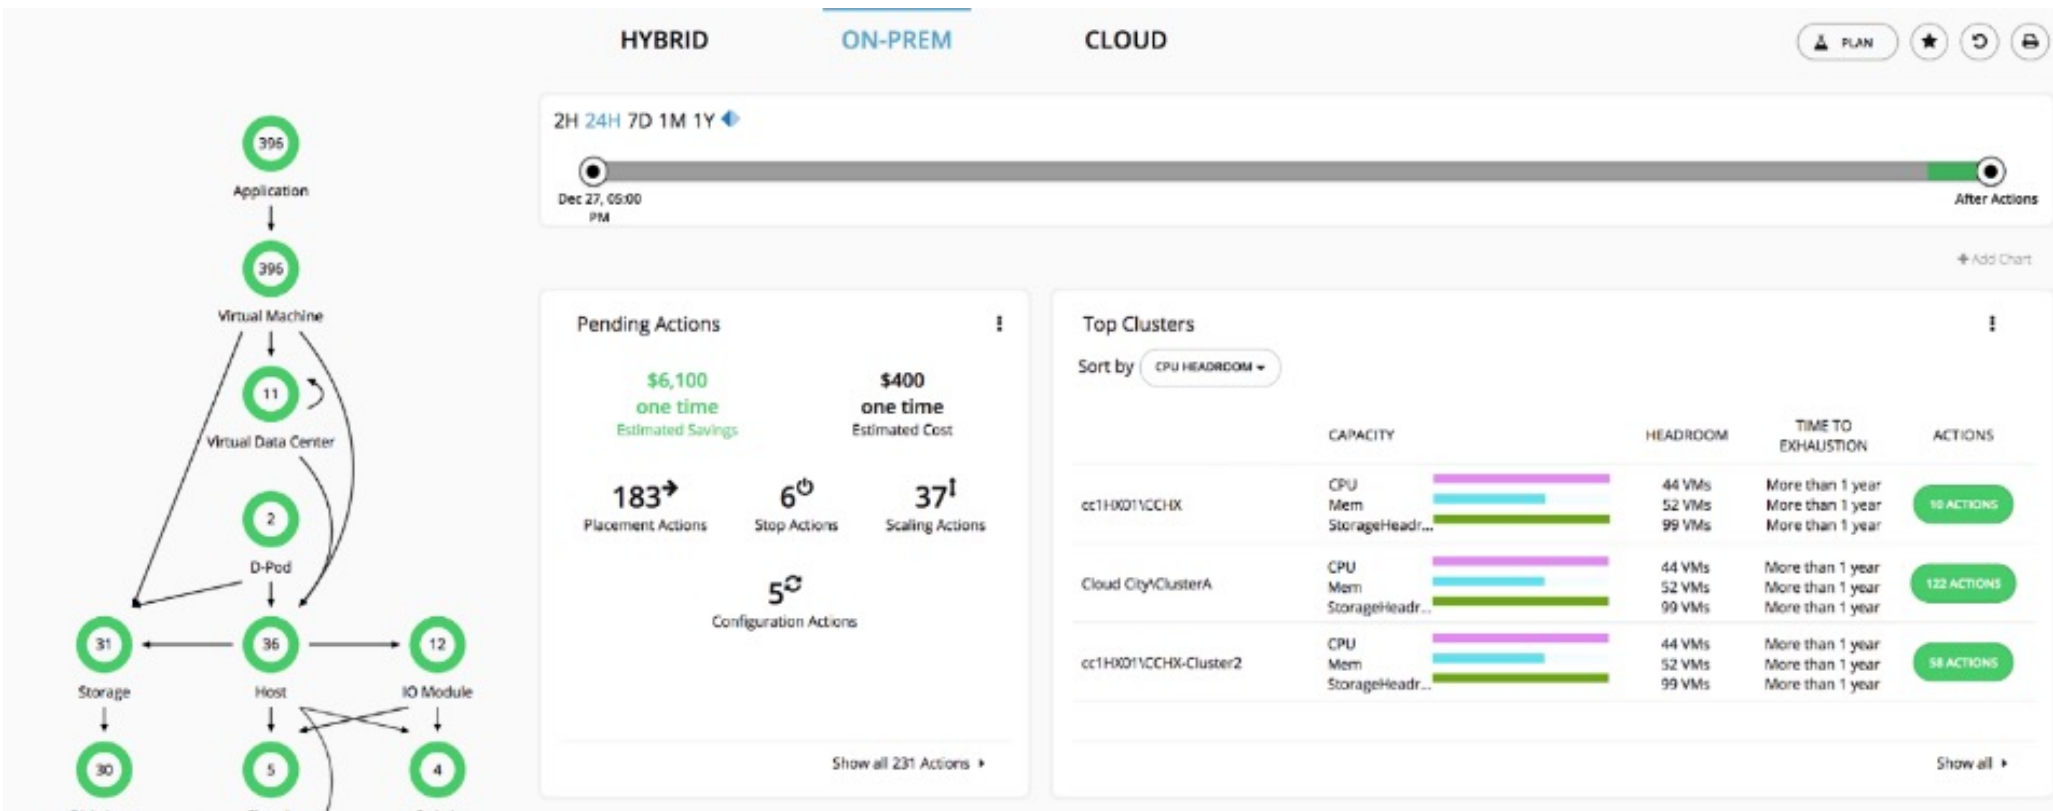
Use Case

Modernize at the Pace of Your Business

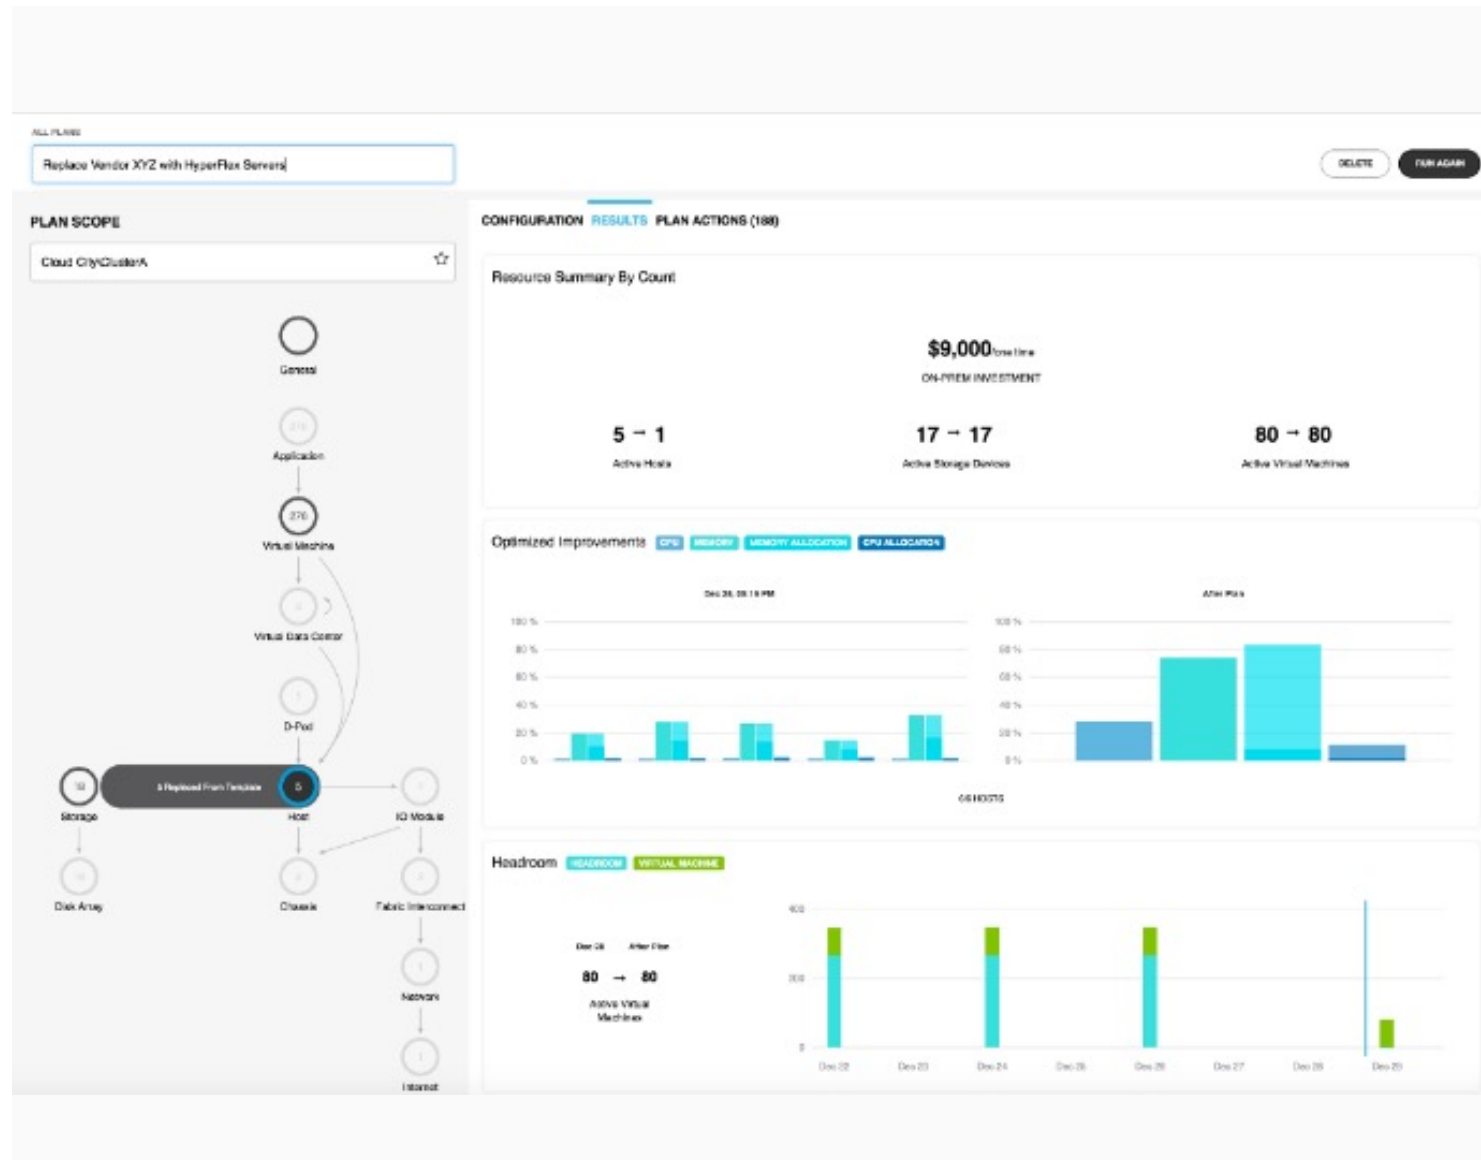


- HyperFlex Systems and CWOM seamlessly integrate with the data center you have today.
- Quickly model “what if” scenarios as your business grows and as you retire legacy infrastructure.

Below, CWOM analyzes workload growth trends to determine when new hardware will need to be purchased. Note the 'Top Clusters' highlighted and their 'Time to Exhaustion,' ensuring that you always have the hardware you need to support growth.



Below, CWOM plan results show that workload demand, in this case 276 virtual machines (see supply chain on left), on 5 legacy hosts can be supported with 1 HyperFlex host for a one-time investment of \$9,000. Customers get specific before and after views in order to make fully informed decisions as they modernize their infrastructure.



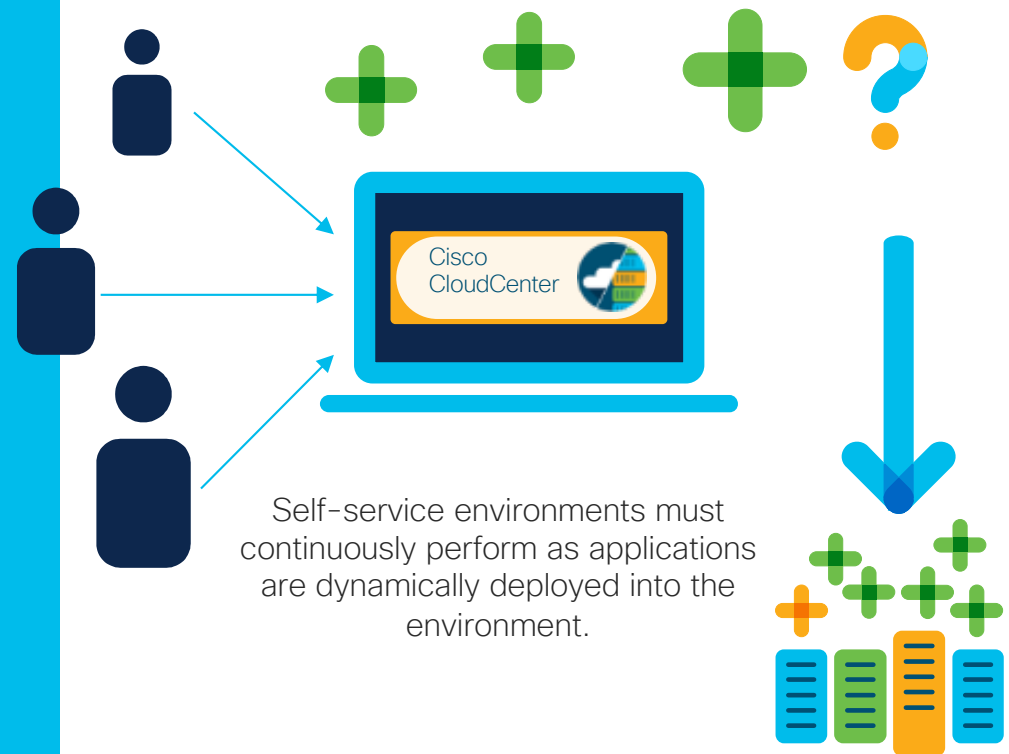


CWOM & CloudCenter

Assure Performance, Minimize Cost, and Maintain Compliance in Dynamic Self-Service Environments

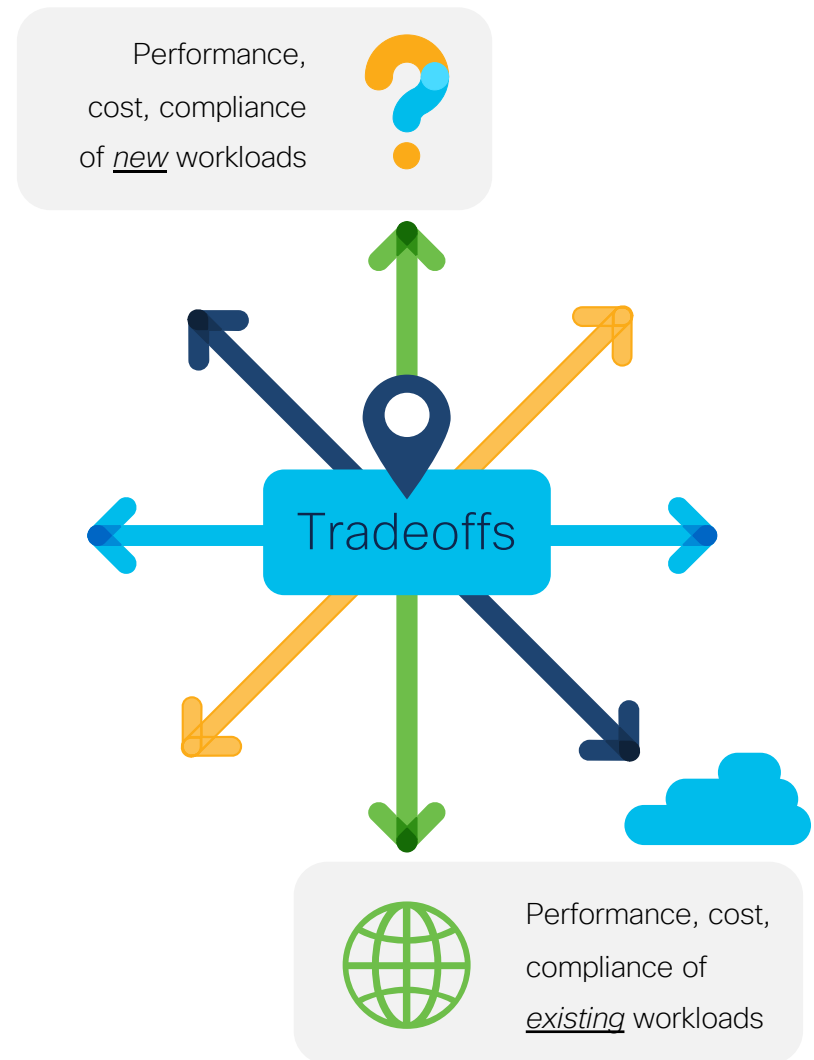
Context

- By 2019 60% of workloads will run in cloud, whether on-premises or off.
- Workload deployments will increase in volume and frequency.
- 98% of organizations say a single hour of downtime costs over \$100,000.



Problem

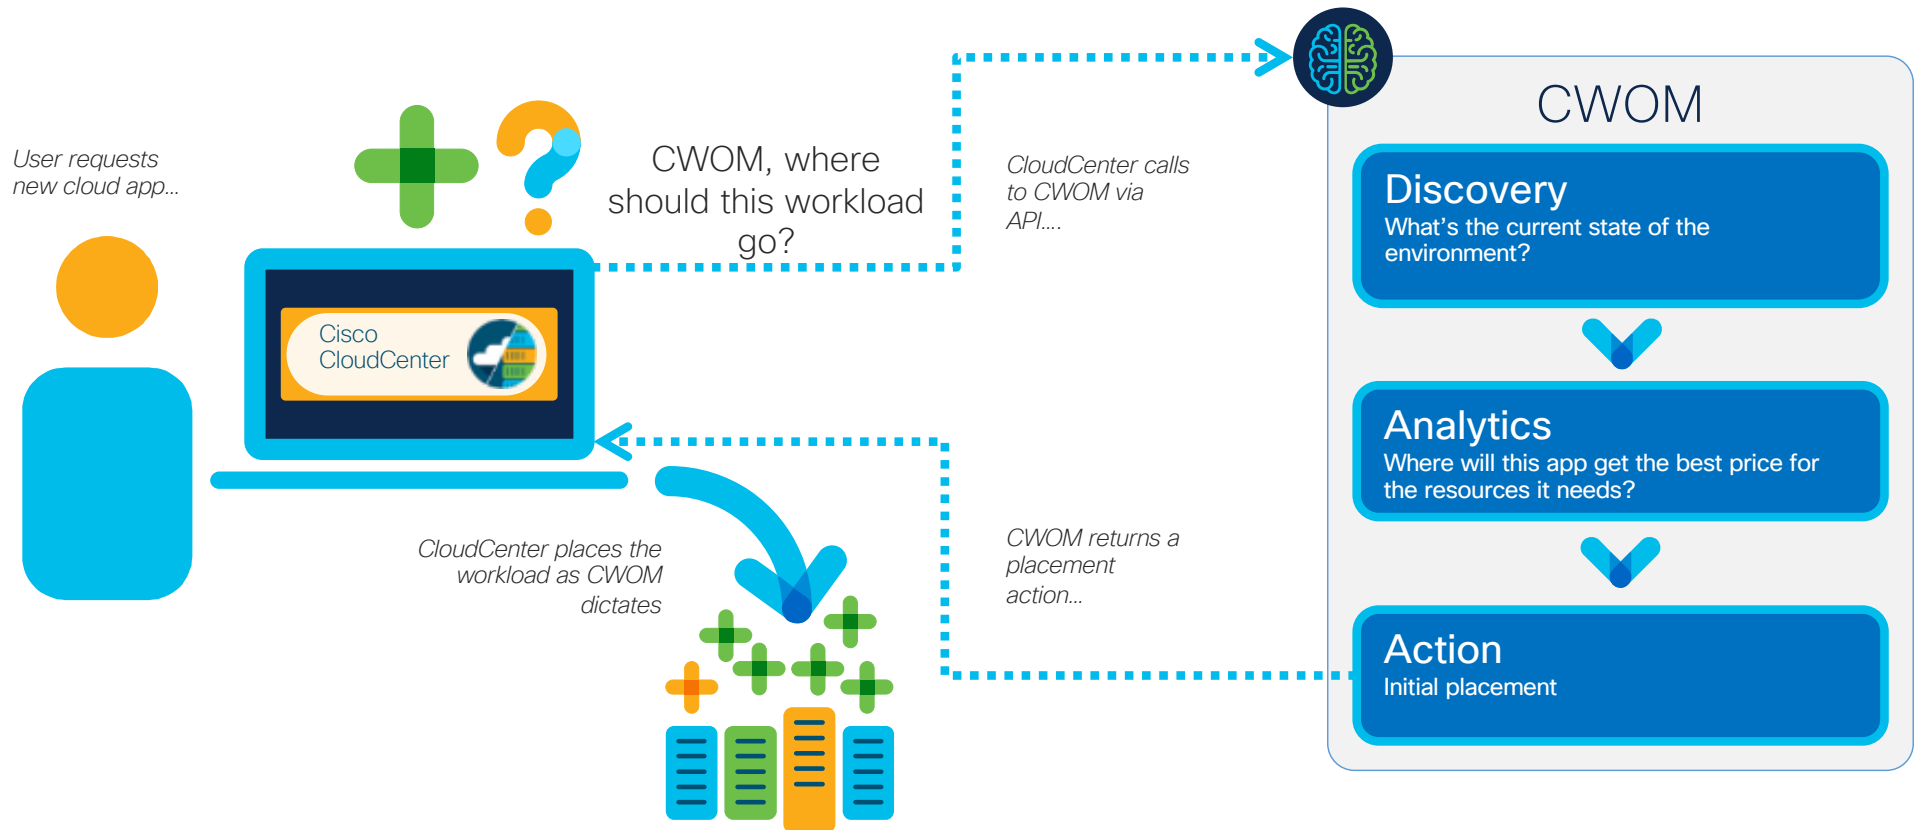
- Self-service = dynamic change.
- Multi-cloud estates expand workload placement options, increasing cost overruns and compliance risks.
- Tradeoffs must be made in real time to keep pace with the business.



What actions does it optimize?

- Initial workload placement on-premises
- Initial workload placement across clouds*

How does it work together?



Demo Video: AppDynamics & CloudCenter with CWOM

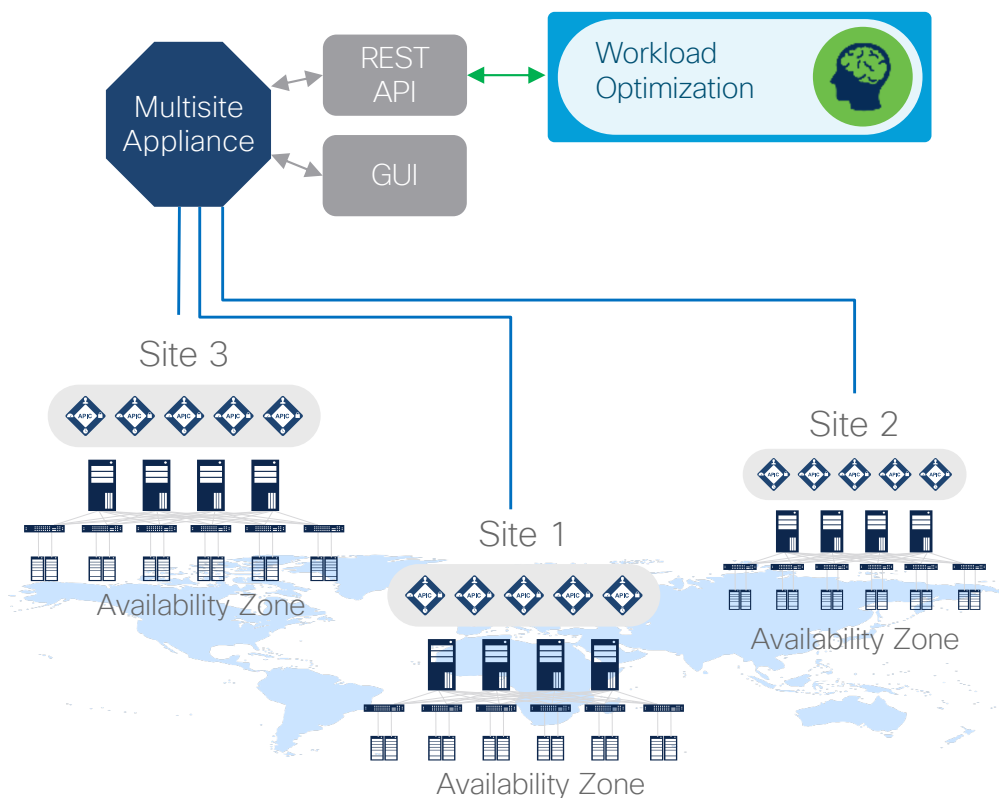


[Click here](#)

ACI

Multi-site mobility and disaster recovery with performance assurance and policy enforcement.

Real-time workload placement that satisfies resource, physical, and policy requirements across multi-site architecture.



Success Stories

Cisco IT ...takes 2,500 CWOM actions / day to mitigate risks



1,514,100

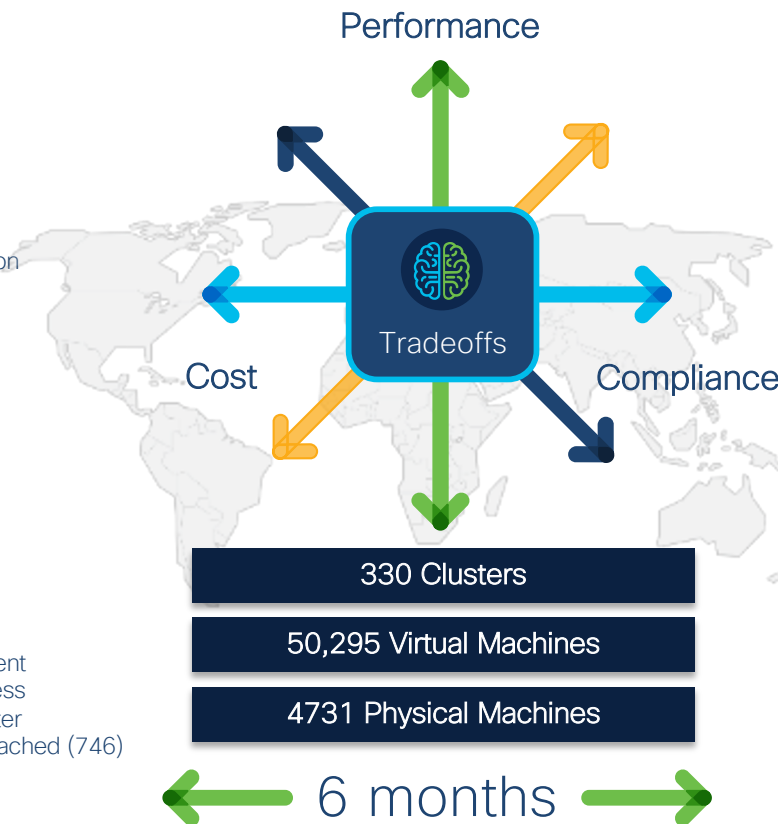
Performance Risks / Automated:

- CPU
 - 299,376 CPU Congestion
 - 20,772 ReadyQueue Congestion
 - 398,226 CPUProvisioned Congestion
- MEM
 - 765,832 Mem Congestion
 - 3,812 MemProvisioned Congestion
- Throughput
 - 6,26 NetThroughput Congestion
 - 165 IOThroughput Congestion

6,028

Compliance Risks / Automated:

- 323 Affinity/Anti-affinity
- 4,959 Configuration Drift
 - 2,380 Violation of Storage Placement
 - 856 Misconfigured Storage Access
 - 1,723 Misconfigured Storage Cluster
- ESXi Management Agent could not be reached (746)



REAL-TIME Pending Resize Actions:

- 1,607 vCPU Decreases due to Ready Queue
- 9,272 vCPU Decreases due to under-utilization
- 410 vCPU Increases for CPU Congestion
- 13,827 vMem Decreases for under-utilization

Driving Efficiency / Automated:

- \$2.8M Yearly Opex
- \$17M Capex Savings
- 130 TB Allocated Memory
- 111 TB Consumed Memory
- 7 TB Allocated vMEM Reservation
- 32,610 Allocated vCPU's
- 3,403 Consumed CPU GHz (50,911 Total)
- 612,000 MHz of Allocated vCPU Reservation

Cisco CWOM & VMware vRealize Suite

