

Technical Validation

# Cisco Intersight Workload Optimizer

Continually Optimizing Performance, Utilization, and Compliance

By Kerry Dolan, Senior IT Validation Analyst

July 2020

This ESG Technical Validation was commissioned by Cisco and is distributed under license from ESG.

## Contents

Introduction .....	3
Background .....	3
Cisco Intersight Workload Optimizer .....	4
Supply Chain Economic Principles .....	6
ESG Technical Validation.....	6
Deployment and Operations .....	6
ESG Testing .....	6
Public Cloud Integration .....	7
Containers.....	8
Integration with AppDynamics .....	9
Cisco HyperFlex Integration .....	10
ServiceNow Integration .....	10
Healthcare Customer Saves Millions with Workload Optimizer .....	11
Cloud-based Healthcare Technology Company Saves by Reclaiming Resources .....	12
The Bigger Truth .....	14

### ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

## Introduction

This report documents hands-on testing of Cisco Intersight Workload Optimizer. This software analyzes application demand, resource consumption, costs, and compliance requirements and provides proactive management to continuously ensure application performance and optimize infrastructure resource utilization, while minimizing costs and ensuring compliance.

## Background

Digital transformation is a key focus of organizations across every industry. Organizations depend on business applications to transform their product, customer, marketing, sales, and logistics operations to deliver high value and competitive differentiation while keeping costs down. According to ESG research, the most cited objective for digital transformation initiatives is to become more operationally efficient; others include providing better and more differentiated customer experiences and developing data-centric and innovative products and services (see Figure 1).<sup>1</sup>

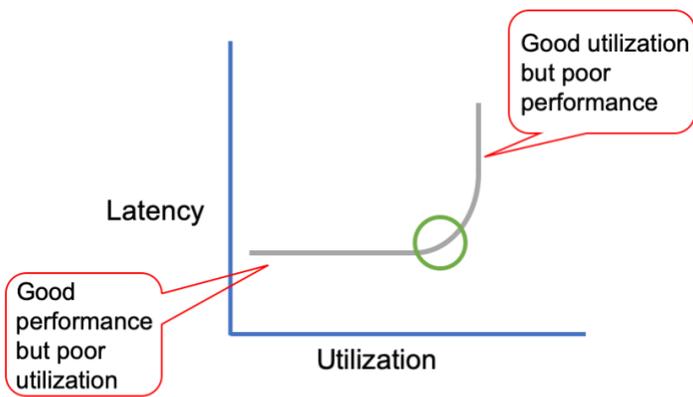
**Figure 1. Objectives for Digital Transformation**



Source: Enterprise Strategy Group

However, the complexity and scale of today’s IT landscape makes this a daunting task, particularly with distributed applications, VMs, cloud services, containers, and microservices administered by multiple monitoring, orchestration, and management solutions. Getting the most performance out of applications is key, but how can IT accomplish that with thousands of applications and microservices on premises and in the cloud—and tens of thousands of VMs, cloud instances and containers—and still keep costs down? How can IT teams design infrastructure deployments to produce the application performance they need, not just today, but every day? How can they assure performance while ensuring high utilization of resources and compliance?

<sup>1</sup> Source: ESG Master Survey Results, [2020 Technology Spending Intentions Survey](#), January 2020.



**Goal: Increase utilization and cost efficiency but stop before you impact application latency.**

Managing infrastructure is a game of tradeoffs. For each application, IT can overprovision to get higher performance and provide a positive customer experience, but this often results in poor utilization that is wasteful, expensive, and leaves resources idle. Or, IT can reduce costs by underprovisioning—resulting in high utilization, but poor performance and user experience. Many organizations overprovision to prevent users from complaining, which can be a costly choice.

The simplified graph at left shows the optimal mix of utilization and performance: increasing utilization and cost efficiency but stopping before you impact application latency. This seems simple enough, **but there are as many graphs as there are metrics**: a complete picture must combine

utilization, latency, IOPS, memory, CPU, network bandwidth, VM memory, vCPU, *ad infinitum*. Add compliance requirements to this, and scale to thousands of workloads, servers, VMs, containers, etc., and **it is clearly an impossible task for humans to manage with spreadsheets**. IT may be able to optimize a single silo, but manually finding the optimal intersections of all resources is impossible.

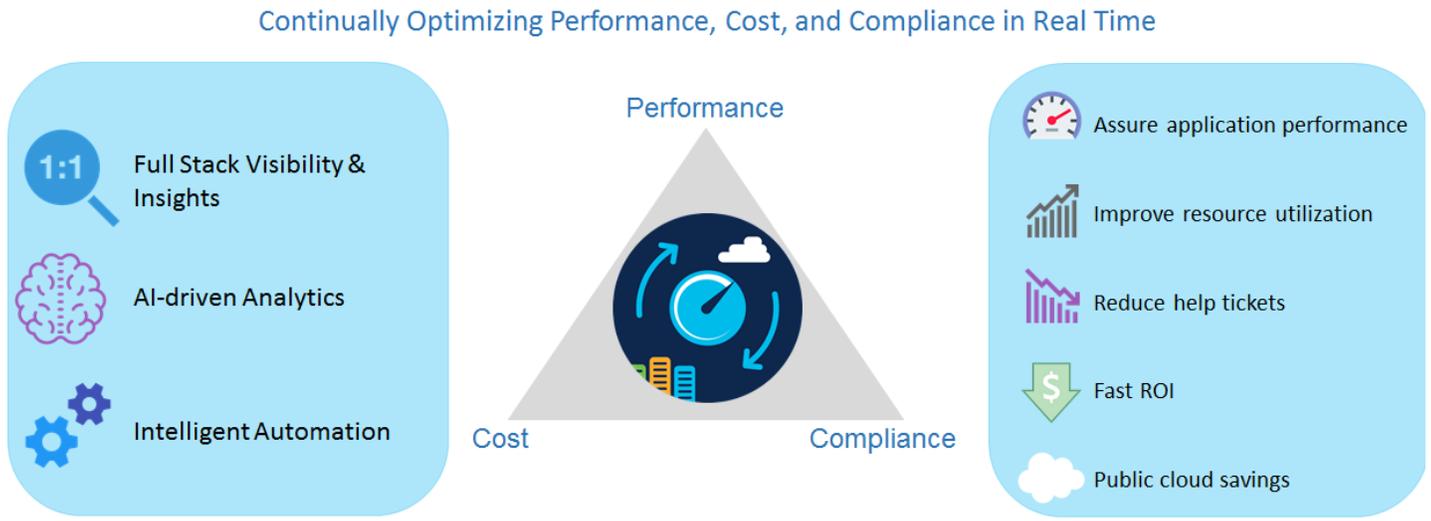
In this complex environment, IT ends up operating blind, taking guesses and often erring on the side that they believe will cause the fewest help desk calls. Organizations need better visibility and insights into the infrastructure and interdependencies underlying their applications, and a systematic, scalable way to implement changes to assure performance and maximize utilization.

### Cisco Intersight Workload Optimizer

Cisco Intersight Workload Optimizer is intelligent application resource management software that delivers visibility and provides insights into your infrastructure—its realities and interdependencies—along with AI-powered, real-time analytics providing resource actions, and automation to implement these actions. Application performance is the number one priority; Workload Optimizer was designed to get applications the resources they need when they need them to ensure performance, and then to increase resource utilization to optimal levels.

Workload Optimizer drives infrastructure toward the desired state of high performance, maximum utilization, and full compliance on an ongoing basis by continually reviewing infrastructure details. This information is fed into a real-time decision engine that provides decisions on where, when, and how to run workloads. Workload Optimizer provides visibility and insight throughout the infrastructure stack—including on-premises and multi-cloud applications, VMs, databases, storage, compute platforms, and containers.

**Figure 2. Cisco Intersight Workload Optimizer**



*Source: Enterprise Strategy Group*

Workload Optimizer is easy to install and is agentless—data is gathered through API calls. Once installed, Workload Optimizer immediately creates a topology map of the on-premises, cloud, and hybrid environments showing infrastructure resources and interdependences between them. Armed with a complete view of the environment, Workload Optimizer then offers decisions for workload placement, resizing, and capacity utilization, and can automate actions based on that information to support ongoing health and performance of applications and workloads.

Features include:

- Designed for virtualized on-premises and public cloud infrastructures; supports multiple hypervisors, containers, and cloud platforms.
- Integrated with Cisco ecosystem including UCS, HyperFlex, Tetration, AppDynamics, and ACI.
- Integration with AppDynamics and other application performance management solutions delivers the application insights and interdependencies, not just the infrastructure view.
- Integration with ServiceNow, Ansible, and other automation/orchestration platforms.
- Fast, easy deployment—Workload Optimizer provides actionable insights within an hour once you add targets.

### Supply Chain Economic Principles

The key intellectual property of Workload Optimizer is an internal decision engine that uses economic principles based on supply and demand of the commodities at every layer of the environment: from virtual data centers and applications to VMs, servers, CPU, network components, containers, storage, disks, controllers, I/O modules, etc. Every resource has commodities to “sell” and to “buy.” For example, a VM might “buy” CPU and memory from the host and “sell” those virtualized resources to applications. The “price” of every commodity is determined internally based on the current need and availability—demand and supply—and the component in need takes the resources from the lowest priced seller. For example, if current workloads are IOPS/performance hungry, CPU and storage will be more valuable; but if a workload starts doing a lot of ingest and analytics, then bandwidth becomes more valuable. Thousands of transactions happen in a kind of “auction” every few minutes, with the need, availability, and prices fluctuating, and transactions occurring among commodities across every part of the infrastructure to achieve the desired state.



This results in thousands of decisions (such as start/buy, stop, delete, place, scale, and configure) in categories (such as performance or efficiency) that can be viewed and either manually or automatically implemented. Examples include scaling up VM memory, re-provisioning storage, or moving a workload to another node. For on-premises actions, the costs are configurable and often reported as a percentage of potential savings (e.g., save 30% of your vCPU). For cloud workloads, decisions include real currency savings, as Workload Optimizer can provide decisions such as moving to a lower cost template or moving a workload to a reserved instance, with savings based on regional pricing.

### ESG Technical Validation

ESG viewed Workload Optimizer demos that included initial startup and ongoing management of VMs and containers; other areas of focus included integration with AppDynamics, ServiceNow, and Cisco HyperFlex. It should be noted that Workload Optimizer has many ways to view information with various dashboards and filtering mechanisms; only a few are used in this report.

### Deployment and Operations

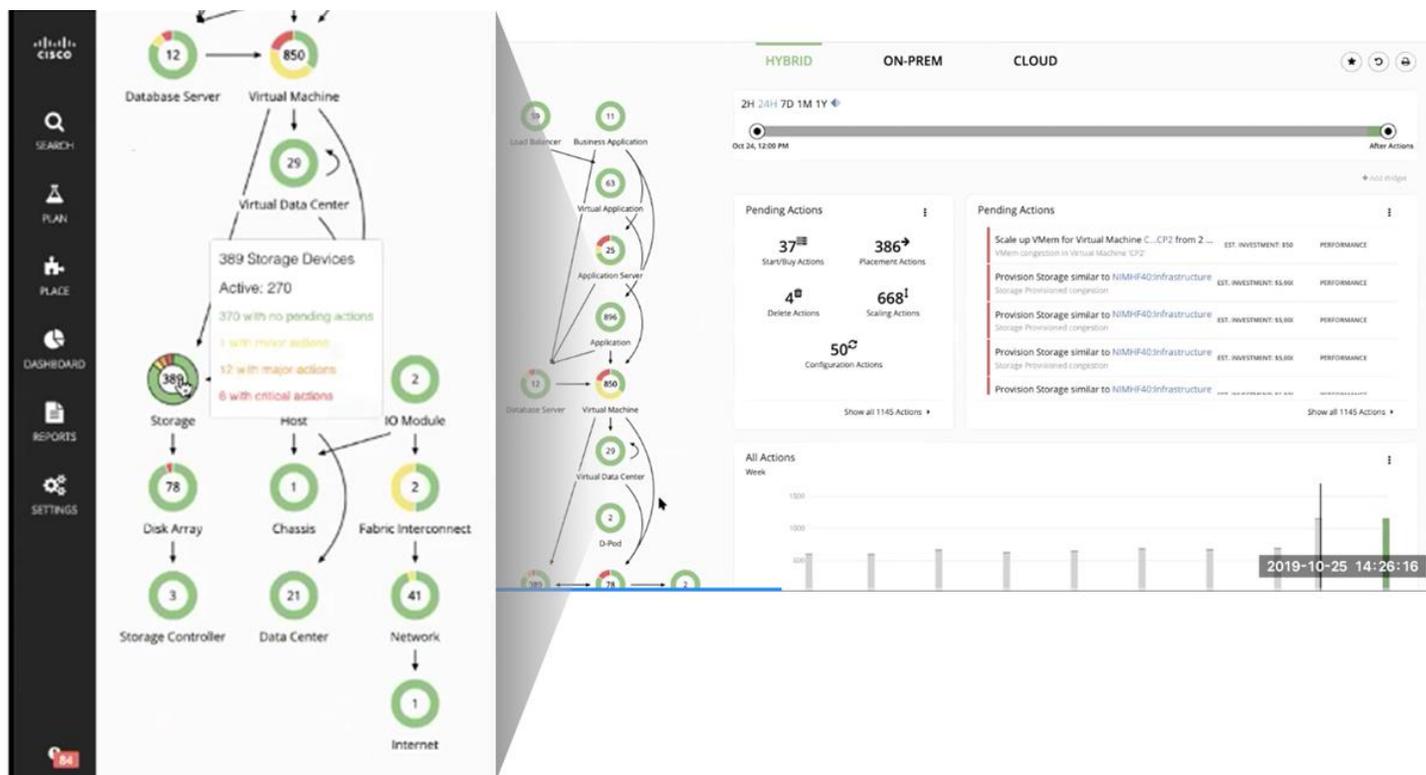
#### ESG Testing

ESG started by logging in and entering the license key; the next step was to select targets for Workload Optimizer access. The typical first targets are the hypervisor and cloud management, followed by other categories such as servers, storage, databases, hyperconverged infrastructure, orchestrators, fabric, and network. In minutes, Workload Optimizer began collecting information about the environment and created a topology map showing specific entities, their relationships to one another, and color-coded health status (see Figure 3).

Note that IT can view any resource—on-premises, cloud, or hybrid—by timeline. On the left of Figure 3, Workload Optimizer shows the topology map, and IT can click on any item—from high-level business applications down to storage components such as I/O modules—and drill down into its components to view status and decisions. Markers on the map are color-coded to indicate the severity of potential actions. On the right, Workload Optimizer lists potential actions that

can be taken by category, such as start/buy, place, delete, scale, and configure; pending actions can also be viewed by color-coded severity, with notations about estimated costs and type of impact, such as performance or efficiency. Scrolling down the view shows actions executed during the week with dates, times, and whether the action succeeded. IT can also view by cluster and see CPU, memory, and storage headroom. There are also graphs available that show the current status and projected improvements.

Figure 3. Workload Optimizer View of Infrastructure



Source: Enterprise Strategy Group

Most pending actions listed can be executed by checking its box and clicking **Apply**; this creates an API call to the appropriate target, such as to vCenter to scale down vCPU for an overprovisioned VM, recouping resources. Actions can be nondisruptive if supported by the underlying target (e.g., hypervisor, storage, etc.). Cisco reports that many customers begin by reviewing the actions and manually implementing them but switch to automated implementation after they see the benefits.

### Public Cloud Integration

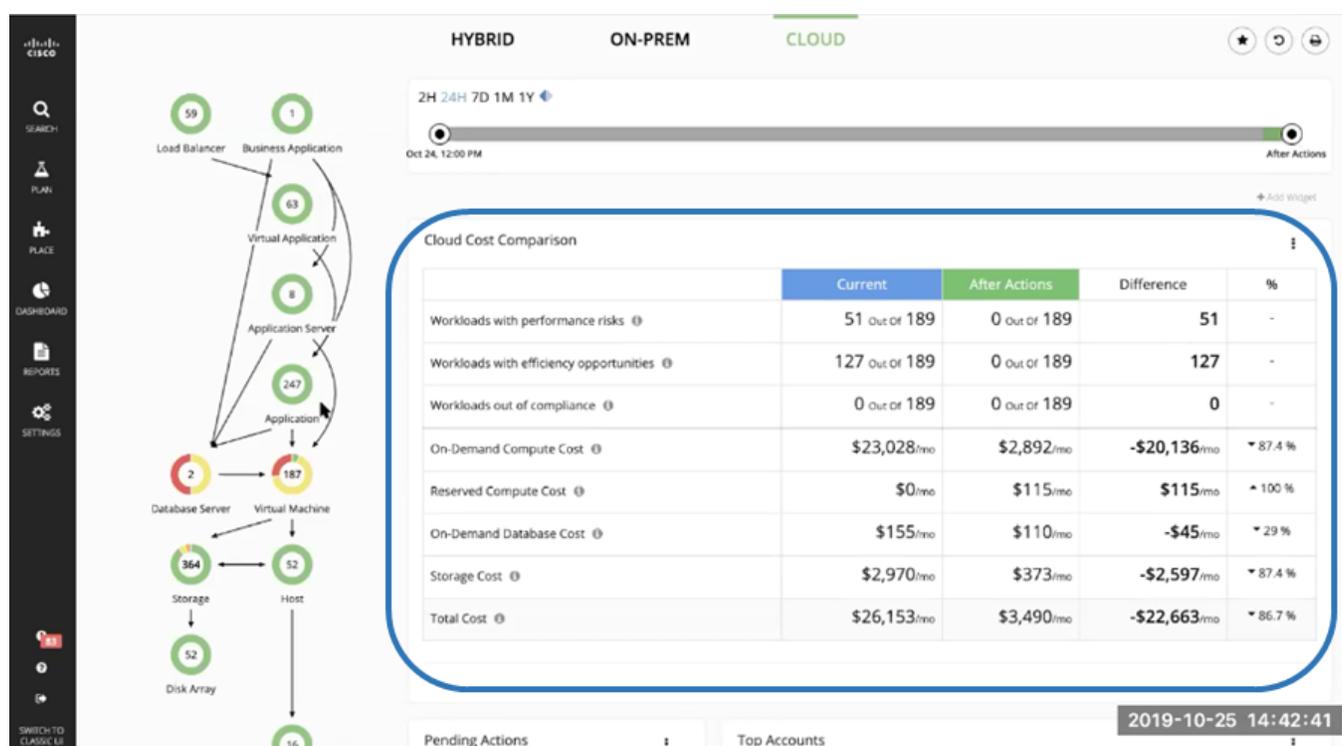
Workload Optimizer performs similar activities with cloud resources (including AWS and Azure, and Google). In the cloud, workloads are run on templates for on-demand instances or prepaid reserved instances (RIs) configured with CPU, memory, storage, etc., for a certain cost. Each cloud provider’s catalog changes continually, which is difficult for users to keep up with. Workload Optimizer maintains updated catalogs for all cloud providers.

A significant challenge with cloud workloads is keeping up with each application’s needs and optimizing utilization. As workload usage shifts over time, it is extremely difficult to know which is the best instance for the money today, next quarter, and next year. Organizations may purchase an RI and run a workload on it that they later abandon, leaving a resource paid for and unused. Workload Optimizer can take into account performance and compliance requirements,

select the cloud instance that offers a lowest cost taking these considerations into account, and move workloads automatically.

Figure 4 shows the cloud dashboard of the demo environment. On the right side, it shows a cost comparison for each category between the deployment as currently configured and how it could be after Workload Optimizer executes actions. Workload Optimizer selected various instance changes—eliminating some and purchasing others—based on driving the applications to the desired state and provided actions to implement them. In this case, workloads with performance risks would drop from 51 to 0, and workloads that could be more efficient would drop from 129 to 0. It also shows specific cost changes for compute instances, databases, and storage; on-demand cloud computing costs would drop by \$20K/month, reserved instance costs would increase by \$115/month, on-demand database costs would drop by \$45/month, and storage would drop by \$2.5K/month, for a total savings of more than \$22K/month, or 86% of current costs.

**Figure 4. Cloud Resource Cost Savings**



Source: Enterprise Strategy Group

## Containers

Organizations are increasingly deploying applications using containers, and Kubernetes has become the *de facto* standard for container cluster management. The Scheduler within Kubernetes orchestrates container workloads by placing “pods”—groups of containers—on infrastructure servers with available resources. However, it is not application-centric; it deploys pods wherever there is available space, and until the pod is expired, it does not return to check on that pod’s utilization or performance. As a result, it does not recognize when resources become constrained and application performance is impacted.

This is where Workload Optimizer can make a significant difference—not replacing the Kubernetes Scheduler, but optimizing it. Workload Optimizer can identify performance and utilization details and make changes. For example, when two pods encounter peak usage at the same time, Workload Optimizer can move one pod to a different node to avoid

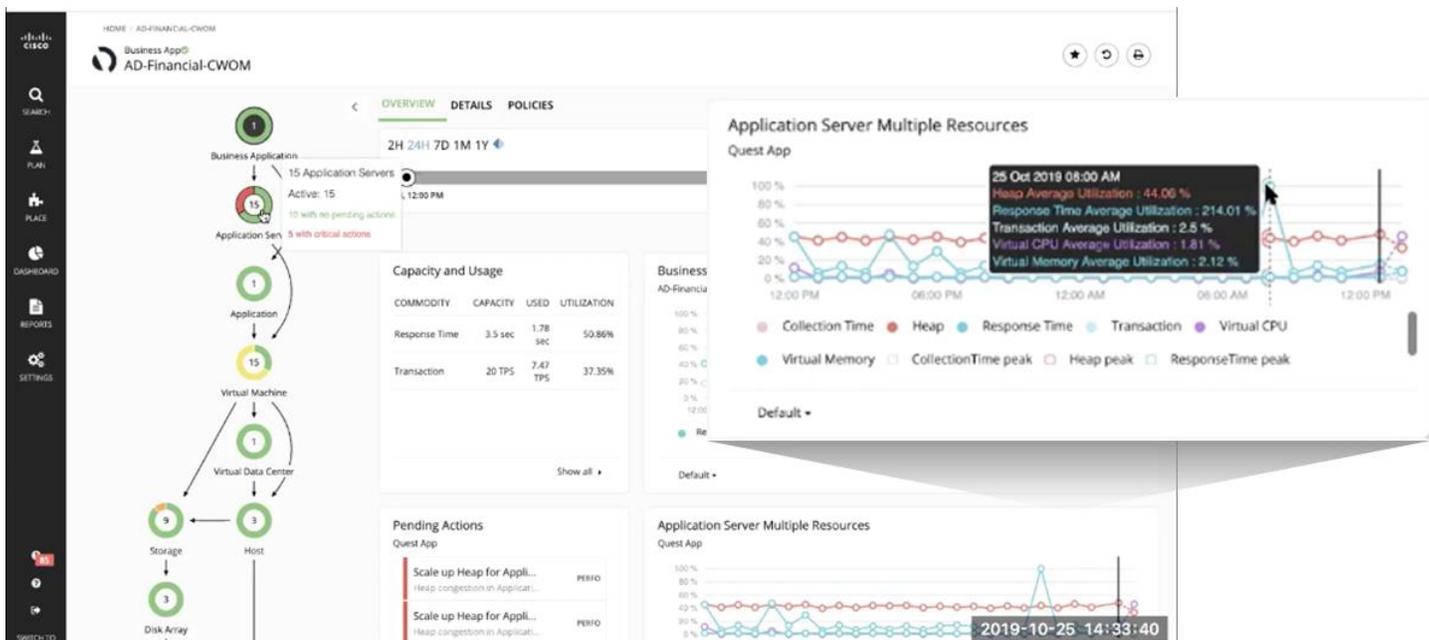
“noisy neighbor” performance impacts. In another example, if IT wants to deploy a new pod but both Node A and Node B lack sufficient CPU or memory for it, the Kubernetes Scheduler would not deploy the new pod. In contrast, Workload Optimizer can identify the resource contention, move a workload from Node A to Node B, and then place the new workload on Node A. When you factor in the scale of thousands of containers for hundreds of workloads, Workload Optimizer can make a significant difference.

### Integration with AppDynamics

Workload Optimizer is integrated with AppDynamics, an application performance management and analytics platform that can be added as a target for on-premises or SaaS applications.

On the topology map, we clicked on the business application *AD-Financial-CWOM* that appeared partially red, indicating that there were critical actions. We drilled down to see the actions, which included scaling up the heap memory resource in the JVM. This information is available to Workload Optimizer because of the AppDynamics integration; Workload Optimizer can ingest information and view the AppDynamics application-based view of the infrastructure. Since Workload Optimizer is getting metrics directly from AppDynamics, it can show a graph of the response time utilization (see Figure 5) and safely recommend/implement actions. For example, it won’t recommend sizing the VM memory below the heap size of the running JVM within. The Workload Optimizer and AppDynamics integration is bi-directional and it provides closed-loop optimization. Workload Optimizer infrastructure resourcing actions assure that applications adhere to Service Level Objectives (SLOs), with real-time validation by AppDynamics.

**Figure 5. AppDynamics Integration**



Source: Enterprise Strategy Group

Without AppDynamics, Workload Optimizer can still get information from API calls to application and database servers, but the logical interdependencies among application components would not be visible. Since AppDynamics has already instrumented the servers, it can provide Workload Optimizer insight from the business application perspective immediately. IT can see clearly when a “noisy neighbor” workload may impact other workloads and take action to correct that in advance. Integration with AppDynamics makes Workload Optimizer smarter by connecting the application’s view of resources with the infrastructure view, connecting the actions back to the intended business outcomes and the end-user application experience.

This application-specific information is especially helpful for planning and creating “what-if” scenarios. Want to see what to expect after a big launch? What happens if IT moves some workloads and retires old ones? How will this impact each application and your environment as a whole? This is where application-level detail is extremely helpful, rather than piece-by-piece components.

### Cisco HyperFlex Integration

Workload Optimizer is integrated with Cisco’s HyperFlex hyperconverged infrastructure. This enables Workload Optimizer to connect directly to vCenter, Cisco UCS Manager, and the HyperFlex storage controller that manages the file system. Workload Optimizer can understand the different performance capabilities of hybrid and all-flash arrays, recommend actions such as storage vMotion, or move VMs based on knowledge of the system’s CPU, memory, and storage capacity.

In addition, it uses the Supercluster concept to streamline vCenter usage and minimize sprawl. Typically, IT organizations build a vCenter cluster and when they reach a defined threshold, they stop adding workloads and deploy another chunk of infrastructure for a new cluster. That old cluster may become inefficient or underutilized, but IT rarely returns to review its usage and improve management; they just create new silos continually. Workload Optimizer can logically combine physically separate vCenter clusters using Groups and manage them through placement policies. IT can merge two vCenter clusters and treat them as a single entity, providing decisions to optimize their resources. For example, Workload Optimizer can provide actions and automate placing workloads with high performance demands on an all-flash cluster and other workloads on a hybrid cluster.

### ServiceNow Integration

With ServiceNow integration, Workload Optimizer can track Workload Optimizer decisions and actions with an audit trail for compliance troubleshooting and can make Workload Optimizer decisions part of a customer’s ServiceNow review and approval workflow. When creating a ServiceNow ticket for tasks such as adding workloads, ServiceNow can reach out to Workload Optimizer to discover the optimal location based on performance and compliance needs, and make recommendations that application owners can review, approve, and permit Workload Optimizer to implement.



### Why This Matters

The desired state of optimal performance, utilization, and compliance cannot be achieved using a series of point tools. Looking at any subset of infrastructure resources in isolation cannot achieve modern IT goals; the scale is simply too vast for humans with spreadsheets and point tools. But abstraction can simplify the management of complex, heterogenous environments. Using the basic laws of economics and a robust analysis engine, Workload Optimizer can make decisions—and take actions—to ensure that CPU, memory, disk I/O, latency, cloud template placement, network resources, VM placement, etc., are used in ways that assure application performance with maximum utilization, and within rules of compliance.

Cisco’s own implementation of Workload Optimizer has paid dividends. One of the company’s data scientists commented,

*“[Workload Optimizer] enables elasticity for us by continuously sweeping the environment to identify what clusters are running hot, what clusters are running cold, what VMs are actually using all of what they have today and may need more. And simultaneously, what VMs are sitting there idle but have memory allocated to them sitting unused. You really see the enablement of elasticity—not necessarily resizing the VM, whether smaller or larger, but actually increasing the performance. We saw this in our RTP1 data center—we saw that enabling elasticity across about half of our overall infrastructure and downsizing about a quarter of it actually reduced our overall resource contention in the data center by 80%.”*

## Healthcare Customer Saves Millions with Workload Optimizer

ESG spoke about Workload Optimizer with the Senior Manager of Servers and Virtualization for a customer in the healthcare industry. This organization includes multiple hospitals, 30,000 employees, and 500 IT staff. A team of 12 is responsible for the 450 servers and 3,600 VMware virtual machines that support the organization's critical applications. The key applications are the EPIC electronic health record software and SAP for business and finance. Other applications include clinical software for various departments such as radiation oncology and ancillary lab applications; business intelligence; and Microsoft SQL Server. Servers are located in two data centers located within 50 miles of each other.

The organization initially sought a solution to help them evaluate their resource usage. They had an inkling that they could save money by improving resource utilization but lacked the visibility to clearly understand their usage. After a successful Proof of Concept (PoC), they rolled Workload Optimizer out to the production environment for moves, migrations, and placement decisions, and have seen significantly fewer help tickets due to application performance issues since then.

*“Workload Optimizer is saving us more than \$2M in licensing costs.”*

--Senior Manager of Servers and Virtualization,  
Workload Optimizer healthcare customer

This organization has seen substantial results with Workload Optimizer in production for a little more than a year.

- Licensing cost savings of \$2M.** A change by Microsoft in the SQL Server licensing model, from per-CPU to per-core, resulted in unexpected high costs last year. However, an option to consolidate SQL Server VMs into a single cluster would enable licensing by host cores instead. This organization is using Workload Optimizer to move all SQL server VMs by tagging them for automated placement on that cluster, saving more than \$2M in licensing costs this year, as well as the IT staff time it would have taken to move these.
- Cost analysis for moving to cloud.** When Windows 2008 was going out of support, Microsoft offered a cost advantage for moving those workloads to the cloud. This organization had Workload Optimizer review all the Windows 2008 servers in the virtual environment and determine what the cost would be to move them to Azure instead of paying for extended support. Workload Optimizer provided the cost analysis they needed to make a decision about how to proceed. The manager commented, *“This gave us real confidence. It’s a little scary to use the cloud because you’re not sure what will happen, and what the cost will be. Workload Optimizer did the analysis, and based on our application history, it told us where to reduce CPU and memory for potential savings in the cloud too.”*
- Storage Migration.** After purchasing a new storage array, the organization wanted to move VMs from the old array to the new one. They used Workload Optimizer to migrate the workloads by turning on storage automation. Instead of manually moving hundreds of VMs, they let Workload Optimizer take care of it in the background. The manager expected this task to take three weeks if done manually, while Workload Optimizer could do it in less than a week. In addition to savings on IT staff time, they were able to save the cost of running the two arrays in parallel for only about a third of the time expected.
- Increased VM density, \$1.5M hardware cost savings.** This organization has almost doubled VM density, from about 10 VMs per host to 15-20 VMs per host. Commented this manager, *“In the first year alone, we were able to avoid about \$1.5M in hardware. If we were doing business as usual, we would have ended up buying more hardware than necessary.”*
- Audit log.** One unexpected benefit of Workload Optimizer has come from periodically reviewing the audit log, where Workload Optimizer tracks all changes, moves, expansions, etc. This manager says that when Workload Optimizer has

not been able to move a VM for some reason, he has been able to find errors such as incorrect network connections and resolve them quickly. *“I am discovering inconsistencies in my environment that I wasn’t aware of before—all these other tools aren’t telling me that there’s something wrong,”* he commented.

For this organization, the next phase of Workload Optimizer usage will include having Workload Optimizer suggest what VMs to add CPU or memory to. *“We’ll be able to tell application owners that we did them a favor, adding resources before they even knew they might have a performance issue,”* said the manager. While their change control process requires them to manually make these changes, the manager said that with Workload Optimizer, “manual” is just a matter of clicking “apply” versus doing all the work yourself, so it is simple.

*“There is a transformation happening, [with] IT moving from a cost center to a business partner. With Workload Optimizer, we can help the business save money.”*

--Senior Manager of Servers and Virtualization, Workload Optimizer healthcare customer

## Cloud-based Healthcare Technology Company Saves by Reclaiming Resources

ESG also spoke with a Senior Technology Manager and a Lead Engineer of a cloud-based healthcare technology organization about their Workload Optimizer success. This \$500M organization employs 2,500 people and helps customers in all areas of healthcare.

*“We have saved about \$200K in Windows licensing, and more than \$500K in SQL licensing, just based on efficiency right sizing. And, if you’re reclaiming hosts, you also reduce Linux and ESXi licensing, so there is additional cost avoidance there.”*

--Senior Technology Manager, Workload Optimizer healthcare customer

This organization’s key objective when they purchased was to improve infrastructure efficiency on-premises and in the cloud. The IT staff had tried to manually identify resources for reclamation, with a team of seven reviewing inventory and validating servers to support various applications. This was difficult and consumed significant staff time, so they searched for a technology solution. After looking at several potential solutions, they selected Workload Optimizer because it offered a solution to support both on-premises and cloud resources, and because it supported their Pivotal Cloud Foundry resources. In this organization, IT manages infrastructure to support about 4,500 VMware VMs. The primary applications are Microsoft

SQL databases, and big data solutions including Hadoop, Greenplum, and MongoDB. They have both Azure and AWS public cloud presences.

The managers we interviewed said the Workload Optimizer implementation was simple, and documentation was good. They also took advantage of professional services, which they said were extremely helpful to ensuring that they got more than they expected from Workload Optimizer. In addition to Pivotal, Azure, and AWS, Workload Optimizer targets include their Cisco UCS Managers, AppDynamics, and numerous Dell EMC storage solutions.

They are extremely pleased with the results so far. The organization has been able to:

- Achieve efficiency objectives. According to the Senior Manager, *“In a six-month period we rightsized or scaled down about 650 VMs, allowing us to reclaim about 1,000 virtual CPU cores and 6,000 GB of virtual memory. We also identified idle workloads, and were able to decommission 25+ additional machines, saving an additional 100 vCPU and close to 300 GB of virtual memory.”*

- Increase cluster efficiency by moving VMs between hosts, improving VM host density from 10:1 to 15:1. Said the Senior Manager, *“This also allowed us to reclaim infrastructure, put it into available pools, and reallocate that based on demand to support client expansions, new products, and new business. This has prolonged capital investment, and delayed additional compute procurement for about 12 months.”*
- Reduce licensing costs, in particular for Microsoft Windows and SQL. *“We have saved about \$200K in Windows licensing, and more than \$500K in SQL licensing, just based on efficiency right sizing. And, if you’re reclaiming hosts, you also reduce Linux and ESXi licensing, so there is additional cost avoidance there.”*
- Gain cloud resource efficiency. The cloud team has also seen benefits from right sizing, decommissioning of idle workloads, and reclaiming orphaned storage, reducing costs.
- Leverage the planning feature. This organization has benefited from the opportunity to preview changes such as adding or removing hosts from clusters, identifying the cost of moving applications to public cloud resources, and resizing VMs when going from on-premises to cloud.
- Troubleshoot VM and application performance. Said the lead engineer, *“The [Workload Optimizer] dashboards give more information and are easier to use than vCenter diagrams and such, and we can go back and look for trends or to see if there have been issues.”*
- Leverage integration with AppDynamics for several applications. AppDynamics provides additional data that increases visibility and insight into servers and applications and helps to identify bottlenecks that can then be targeted for remediation. Commented the Senior Manager, *“[With AppDynamics] We’re getting better visibility because we can see from the application perspective, not just the infrastructure view.”*
- Repurpose to other business challenges five of the seven staff members who had been trying to manually find areas to improve efficiency. Now, a team of two is responsible for using Workload Optimizer for both efficiency and performance objectives.

*“Workload Optimizer has provided tangible benefits in terms of infrastructure efficiency. But there have been other benefits that we didn’t expect: reducing and adjusting our workforce, better planning, and better application performance.”*

--Senior Technology Manager, Workload Optimizer healthcare customer

## The Bigger Truth

Digital transformation has become a key priority for many organizations as they strive to leverage data to improve every aspect of their businesses. While these initiatives can significantly improve operational efficiency, customer experience, and business innovation, they also contribute to the complexity of today's IT landscape—and any time there is complexity, there is cost. When ESG asked survey respondents how complex their organizations' IT environments were relative to two years ago, 64% reported that it was more or significantly more complex, and only 3% reported that it was less complex.<sup>2</sup>

Organizations with hundreds of distributed applications built on tens of thousands of individual infrastructure components struggle to deliver the performance required and keep costs down. A primary reason is that they lack full visibility and insights. As a result, most manage by making tradeoffs between performance and resource utilization, based on the limited knowledge that point management solutions can provide.

Cisco Intersight Workload Optimizer can provide that full view not only across individual devices, but also by application and interdependencies, for on-premises and cloud workloads. Its AI-powered analytics engine provides and can automate actions that drive organizations to the highest possible performance with the optimal utilization while ensuring compliance. These actions can result in assured application performance, higher resource utilization, fewer help desk calls, and lower costs.

ESG validated the ease of deployment and operations. It was easy to select targets to include in the topology map and to view the relationships between components. The suggested actions and their impacts showed performance risks, improved utilization, and demonstrated the abundant savings that organizations could achieve with their on-premises and cloud workloads—savings that would be virtually impossible for humans armed with spreadsheets to even identify. In addition, integration with Kubernetes orchestration, AppDynamics and ServiceNow platforms, public cloud services (AWS, Azure), and the Cisco portfolio demonstrate the expansive use cases for Workload Optimizer.

Many of today's IT environments are too large, distributed, and complex for IT administrators to manage to their fullest—it's time for software to manage all this complexity with full stack visibility and insights to prevent application performance issues and proactively optimize your infrastructure. Your mileage may vary based on the specific challenges of your infrastructure deployment. But if you are looking for a powerful solution that can help you assure application performance, and optimize utilization and compliance levels, ESG recommends evaluating Workload Optimizer.

*“We put all these tools in place—we understand what we have and how we can scale and grow it, and then suddenly we don't need to spend \$17M, and more importantly, we don't even need to spend the \$3M to rent RTP7. So, the first year it was going to be the \$17M savings, and subsequently in perpetuity \$2.8M savings a year. It saved us 54TB of memory and 4800 virtual CPU cores. So, we could add, after that reboot, another 2200 servers to the environments for free.”* -Senior Director, Cisco Cloud Orchestration & Platform Services, speaking about Cisco's own implementation of Workload Optimizer.

<sup>2</sup> Source: ESG Master Survey Results, [2020 Technology Spending Intentions Survey](#), January 2020.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



**Enterprise Strategy Group** is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

