The Smart Approach to Cloud Workload Placement Decisions
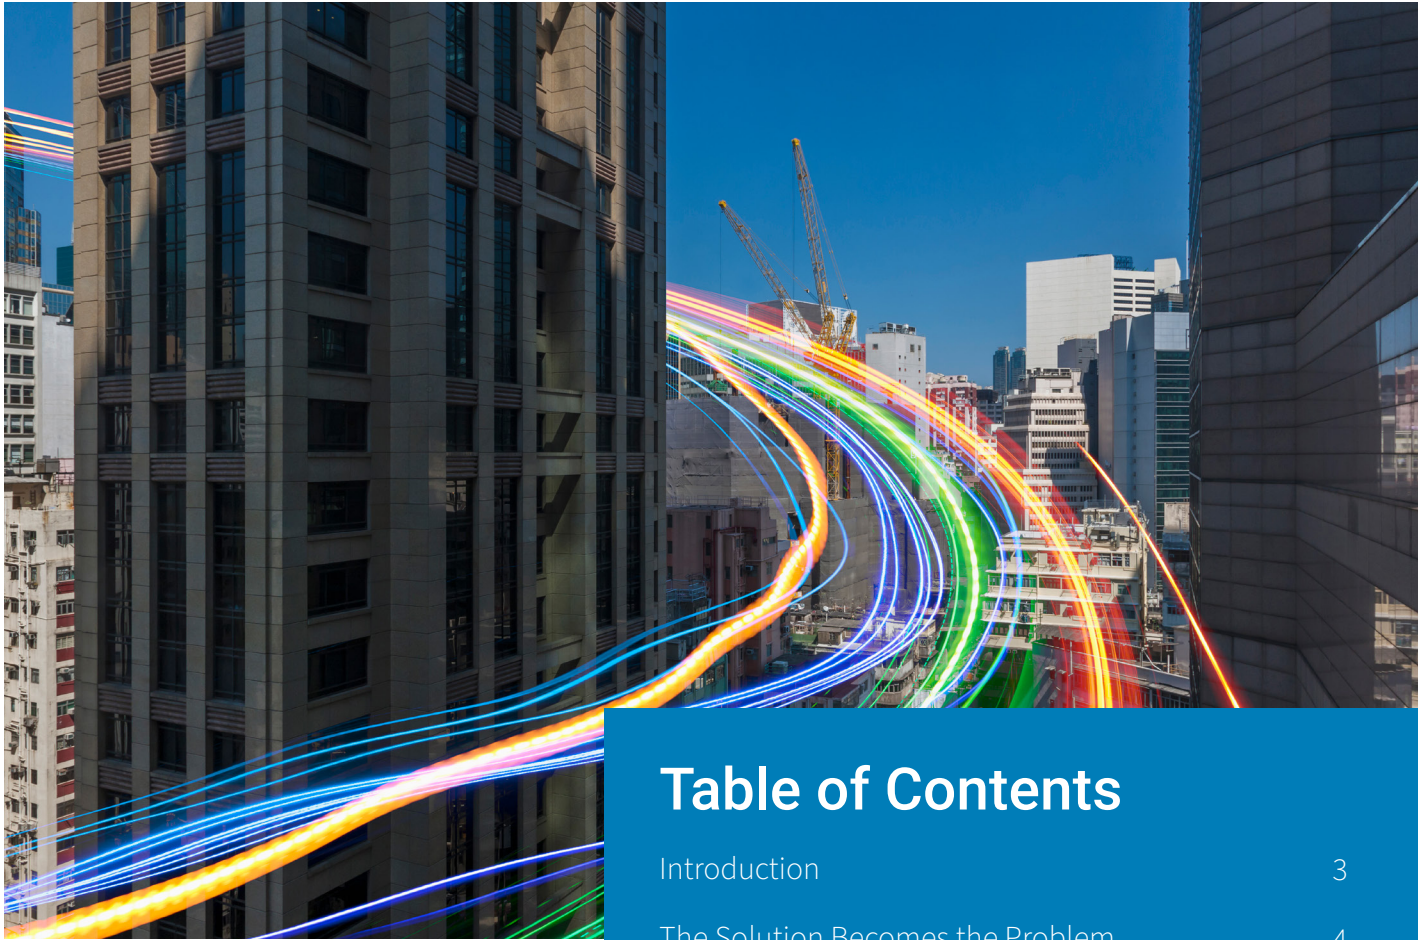
Table of Contents

DELLTechnologies
Dell Expert Network

Introduction

Public cloud spending is expected to reach $304.9 billion in 2021, an 18 percent increase from the previous year, according to Gartner. For a preponderance of businesses, moving as many workloads to the cloud as possible has become a go-to strategy. The cloud-first approach is popular because it allows companies to accelerate adoption of new technologies to gain new capabilities and improve competitiveness. Instead of making capital investments in on-premise IT infrastructure, businesses embrace cloud migration as a way to control IT spending.

While the cloud is a sound strategy in many cases, it isn't suitable to every workload. After moving to the cloud, some businesses find that costs are substantially higher than they hoped and performance falls short of expectations because of application incompatibilities with cloud environments. Unfortunately, these realizations occur after workloads are already in the cloud, creating new challenges that must be addressed.

Workload repatriation – that is, the process of shifting workloads back on premise – is often the answer to a company's cloud woes. But before that ever becomes necessary, companies can avoid ill-advised cloud migration decisions with proper planning. And that involves going through a process of workload identification and discovery to calculate the costs and resources needed to properly run those workloads.

Attempting to perform the process manually is time-consuming, inefficient and error-prone. But there is a way around it – by leveraging a workload placement decision platform that automates discovery and cost analysis. Managed solution providers (MSP) working with clients on cloud placement or repatriation can take advantage of such a platform to steer them in the right direction and avoid the common pitfalls of cloud migration. In so doing, MSPs add value and strengthen customer trust by showing they are working in the client's best interests.

The Solution Becomes the Problem

Cloud-first mandates typically are driven by budgetary considerations and the growing popularity of the cloud. A move to the cloud carries the promise of increased flexibility, scalability and the acceleration of go-to-market strategies. These benefits are achievable in the right context, but if not planned properly, cloud migration creates more problems than it solves.

A common scenario involves virtualization, which companies use to avoid server sprawl. Instead of procuring more and more servers to scale IT operations, companies minimize hardware costs by adopting a virtualization strategy. Sometimes the result mirrors server sprawl, except that instead of physical hardware, the sprawl involves virtual machines (VM). So organizations start moving their VMs to the cloud.

However, without a compatibility check, the VM may not perform as well as expected once in the cloud. Legacy applications that weren't developed for the cloud perform poorly because they should be rewritten or replaced by cloud-native services. In some cases, actually running a VM in the cloud turns out far more costly than if a business keeps it in house.
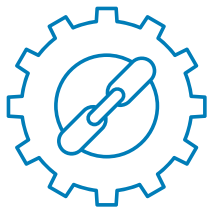
VM overprovisioning is a common occurrence, resulting in unused capacity and higher costs. Most VMs are oversized in a private cloud, where there is no cost penalty, but public clouds are different. Companies often assume it's a simple matter of a lift and shift, but they need to right-size their VMs for the public cloud. Many don't because they don't know how to approach right-sizing or where to get help to do it. MSPs can help them by leveraging a workload placement decision platform for right-sizing and compatibility checks.
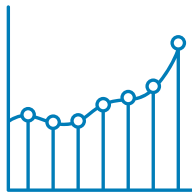
Calculating Cloud Costs

When discussing workload cloud placement or repatriation, cloud decisions often are based on what businesses see as the biggest priority – cost. Budget-driven decisions are at the heart of countless ill-advised migrations. But while budgetary considerations are important, organizations also need to take into account other factors, such as identifying which workloads qualify as the "heavy hitters" that consume the most resources.

When it comes to cost, public cloud providers such as Amazon, Google and Microsoft offer online cloud cost calculators to help users project the costs of using their platforms – usually in one-year increments. The problem is the calculators aren't as easy to use as they may appear, typically requiring hours of work in scripting APIs or manually entering values. Pricing is calculated against a matrix of configurations in a mindboggling number of variances in CPU, memory and capacity.
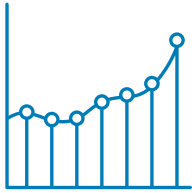
This level of complexity, while an attempt to provide users with all the information they need to make smart decisions, ends up having the opposite effect. All the work involved in entering values and modeling configurations discourages users, be they IT professionals or department managers, from completing the process. Instead of making a smart decision by availing themselves of all relevant data, in many cases they make a determination based on an incomplete picture.

This is especially problematic with "shadow IT" situations, when department managers bypass the IT department in adopting a cloud service to fill a business need. Lacking experience in IT-related decisions, these managers can get in over their heads because of increased complexity and cost overruns. In many of these cases, repatriation is the best solution.
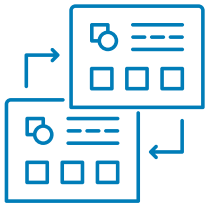
Workload Repatriation

In cases where workload repatriation becomes necessary, businesses need to avoid repeating the mistakes they make with ill-advised moves to the cloud. And this translates to proper planning. With repatriation, the process of discovery, resource allocation and workload placement is just as critical as when a company is deciding whether to transfer workloads to the public cloud.

The process is essentially reversed, in that now businesses, or MSPs working on their behalf, have to perform discovery in the cloud. This entails identifying each cloud VM and ascertaining its function, level of criticality, resource consumption, and dependencies on other resources. Businesses also have to analyze costs for each VM and calculate the projected savings associated with workload repatriation.
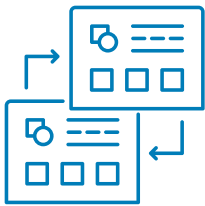
If approached manually, the process is time-consuming, complex and error-prone – in much the same way it would be when planning a cloud migration. The same issues of locating and identifying specific workloads, and their relationship with other resources, are bound to pop up. This is why a manual approach is inadvisable. Rather, businesses should use data-driven discovery with a workload placement decision platform, which accelerates discovery and takes guessing out of the process.

The Live Optics Approach

In searching for a workload placement decision platform that drives value for end customers, MSPs should leverage Dell EMC's Live Optics. Available for free online, Live Optics is a vendor and platform-agnostic agentless service that provides discovery and analysis. It can be used for either determining which workloads belong in the cloud or identifying workloads already in the cloud that should be moved back in house.

Live Optics allows users to collect, visualize and share data about different operating systems, hardware and virtual environments. It features three components that perform specific functions: collector, project and web portal. The collector captures data from servers, VMs and operating systems, and sends it to secure servers. The project component organizes the data, providing users with numeric, tabular and graphical statistics about the environment. An online analytics engine measures and analyzes configuration and performance data to create scenarios to organize workloads and improve resource usage. Projects are viewable in a web portal, from which all the data collected and analyzed can be shared with other stakeholders in formats such as PDF and PowerPoint.

Live Optics offers AWS, Azure, and Google Cloud profiling features to help determine costs and cloud suitability before moving a workload to the cloud. When performing a price analysis of cloud workloads, the tool lets users add up the costs of all VMs based on a cloud provider's one-year discount pricing. Live Optics looks at average VM specifications to calculate monthly costs for each cloud provider. Thanks to the use of public APIs, profiling a public cloud is a speedy process, taking only 30 seconds to a few minutes to collect data.

This type of analysis is key to making a well-informed workload placement decisions, preventing cloud migrations for workloads that are best kept on premise. But even when a suitable workload lands in a cloud, all is not lost. Businesses can move them back through repatriation and run them on Dell EMC on-premise environments.

Live Optics is a valuable tool for MSPs by accelerating IT requirement discovery and workload placement with point-and-click functionality. Manually, the process can take weeks or months. While it is available free of charge, the platform gives MSPs an opportunity to monetize planning and workload placement services – or to perform assessments free of charge that generate more business.

The platform automates data collection and analysis, requiring little effort by MSPs to gather insights about their customers' environments to cost-effectively prepare workload placement plans. MSPs should consider adding Live Optics to their toolbox as a profit-generating tool that adds value and strengthens the IT trusted advisor role.
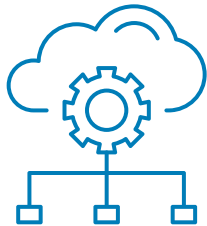
Conclusion

Cloud migration offers many advantages to businesses, including flexibility and scalability, but it is not the suitable to all workloads. Without proper planning and cost analysis, a migration can lead to performance issues and cost overruns.

Such a scenario is avoidable by leveraging Dell EMC Live Optics, a workload placement decision platform that automates the discovery process to match the right workloads to the cloud. It also makes it easy to calculate the costs of repatriating cloud workloads that fall short of expectations. The platform is available free of charge to MSPs, allowing them to generate new revenue from services built around Live Optics data to accomplish these goals:

- Help customers right-size and determine cloud compatibility to avoid overpaying and regretting the move
- Help identify cloud workloads that might be better suited to running in a private environment

These services address a critical need and allow providers to solidify their role as IT trusted advisors and strengthen customer loyalty. For more information about Live Optics or the Dell Expert Network, reach out to www.dell.com/expertnetwork.