THE
BLUEPRINT
FOR CLOUD
SUCCESS

ABSTRACT

In this paper, we will highlight the importance of a multi-cloud strategy which includes a right mix of on-premises as well as off-premises investments based upon application and business requirements. This paper will outline the various approaches to cloud migration including the pro's and con's such as: "lift and shift", "resize and rehost" vs "refactor, replatform, and rewrite". It will also delve into the topic of technical debt for which organizations have accumulated as a result of virtualization (Debt is accumulated when the software upgrade cycle is not in sync with hardware upgrade cycle) and how it can impact cloud adoption. This paper will outline use cases that are suitable for public cloud as well as the trend of repatriating workloads from the public cloud when applications are better served by private or hybrid deployments. Finally, the paper will disclose best practices for driving successful cloud adoption leveraging both existing virtualization as well as the emerging container strategies for enabling a true multi-cloud operating environment. This will include real-world statistics based upon work with organizations in SMB, mid-market and large enterprises.

CONTENTS

INTRODUCTION

Spending on cloud computing in 2018 was over $200 Billion USD and is expected to grow by approximately 20% in 2019.[1] With cloud now approaching 20% of overall global IT spend, the trend is clear, organizations of all sizes are embracing cloud computing. Cloud computing provides a transparent and consistent operating model which is viewed as superior to traditional IT operating environments. Cloud computing promises agility and scalability not available in traditional IT operating environments. Cloud computing also comes in various forms: private, public, and hybrid.

A LOOK BACK AT 2018: THE RISE AND FALL OF THE CLOUD FIRST STRATEGY

One of the more interesting developments in 2018 was the rise and fall of the cloud first strategy. A cloud first strategy is often viewed as an edict that all application workloads should be migrated whole-sale to the public cloud. The popularity of this approach was beneficial in that it accelerated cloud adoption across the industry, however, it has not always proven to deliver the value organizations expected. In the last couple years, cloud technology has dramatically matured, and as result, there are many lessons that can be learned to help organizations avoid the pitfalls of yesterday.

Although the public cloud increases application agility, time to value, and offers an easier way to interact with infrastructure, not all apps are best suited for cloud deployment. A variety of factors including user demographics, service level, security concerns, technical dependencies and financial drivers impact the business case when determining if an application is suitable for cloud deployment. After carefully weighing these factors, organizations can confidently deploy workloads in the right place and capitalize on each cloud's unique benefits.

Our customer engagements reveal that a lift and shift cloud first approach can take focus away from the real business problem the organization is trying to solve. Instead of mandating a cloud first strategy, you should focus on the business outcome, i.e. on optimizing your spend and lowering your risk. It is often the case that a multi-cloud strategy best accomplishes the business goals.

The majority of surveyed organizations reported unexpected costs when migrating all of their applications to public cloud. [2]

THE RISE AND FALL OF THE CLOUD FIRST STRATEGY, CONT'D

To avoid the issues commonly associated with a misinformed cloud first strategy, organizations should analyze their current IT portfolio to determine which applications will have near-term and/or long-term benefits from hybrid cloud deployment. Application workloads should be prioritized for treatment and migration to a cloud based platform based upon the level of effort as well as expected benefits. Each application will fit somewhere in the continuum of resize, rehost, replatform, refactor, retire, rewrite or retain.

APPLICATION CONTINUUM

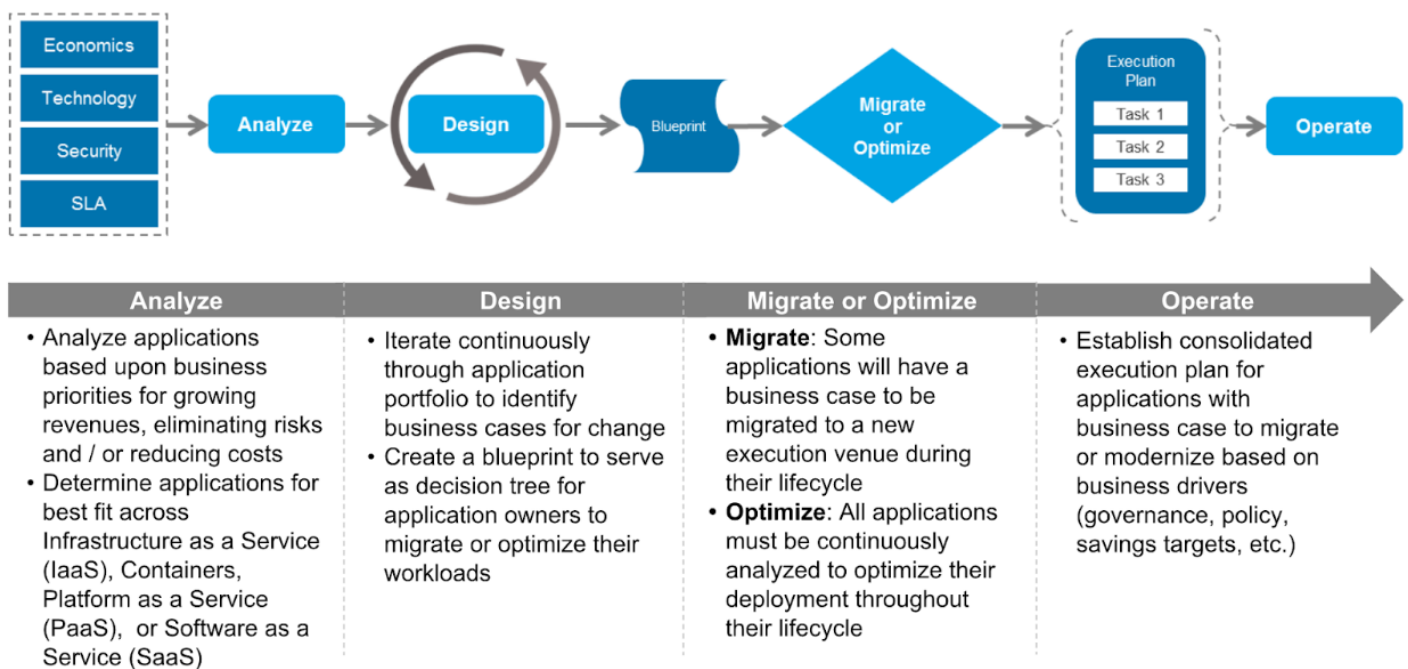| Application Treatment	Application Treatment Outcome
Resize	Reduce or add capacity based upon workload utilization needs
Rehost	Migrate the workload to a new Infrastructure as a Service (IaaS) or Platform as a Service (PaaS) execution venue
Replatform	Upgrade legacy hardware and operating system environments to enable cloud deployment
Refactor	Repackage or re-architect the application workload to take advantage of Containers and Platform as a Service (PaaS)
Retire	Eliminate application workloads no longer in use or better suited for Software as a Service (SaaS)
Rewrite	Invest in the development of a cloud native application
Retain	Leave application workload in current configuration and in its current execution venue

Furthermore, a repeatable process should be implemented. Intelligent workload placement requires an assessment of an organization's tolerance for change and it aligns to an organization's priorities to grow revenues, improve profitability and lower risks. Without having a repeatable process for making workload placement decisions organizations are putting their applications at risk.

"While cloud native application development is an integral component of an organization's digital transformation, the reality is that not all applications have a business case to be rewritten. A successful DevOps practice needs to embrace agile approaches to evolving applications over time as well. This includes resizing, replatforming, refactoring, and rehosting applications to make them suitable for cloud."

- Ethan Michaud - Managing Director, DevOps - Moody's

CLOUD STRATEGY: APPLICATION PORTFOLIO REMEDIATION WORKFLOW

The following workflow is to be followed to deliver continuous optimization

Analyze	Design	Migrate or Optimize	Operate
• Analyze applications based upon business priorities for growing revenues, eliminating risks and / or reducing costs • Determine applications for best fit across Infrastructure as a Service (IaaS), Containers, Platform as a Service (PaaS), or Software as a Service (SaaS)	• Iterate continuously through application portfolio to identify business cases for change • Create a blueprint to serve as decision tree for application owners to migrate or optimize their workloads	• **Migrate**: Some applications will have a business case to be migrated to a new execution venue during their lifecycle • **Optimize**: All applications must be continuously analyzed to optimize their deployment throughout their lifecycle	• Establish consolidated execution plan for applications with business case to migrate or modernize based on business drivers (governance, policy, savings targets, etc.)

REPATRIATION IS A THING

As many organizations have rushed to "lift and shift" their workloads to the public cloud, they have begun to discover that while there are many benefits to the public cloud, the economics are not always better than a private cloud or hybrid cloud approach. The transparency of public cloud pricing has created the illusion that the public cloud is less expensive compared to a well researched and planned hybrid cloud strategy.

REPATRIATION IS A THING, CONT,D

In practice, the public cloud is less expensive for certain use cases. The economics of public cloud make great sense for ephemeral workloads (e.g. development, testing, and disaster recovery), data archival and small compute configurations. That said, our data shows that public cloud can also be up to two times the cost of the unit economics available in a data center. This is especially true for workloads that require constant access and continuous use (e.g. production) with high levels of security and service level requirements.

82% of organizations surveyed who leveraged a "lift and shift" approach to public cloud migration have experienced increased costs.[2]

The economics of cloud depends on a business' scale and IT purchasing power. As the purchasing power of an organization increases the economics of operating data centers becomes much more competitive with the unit economics of public cloud. For very large organizations, the unit economics of the data center can be more attractive than the comparable costs in public cloud.

As organizations have internalized the true cost to operate their workloads in the public cloud, some of our customers have decided to repatriate public cloud workloads back into their data centers. Repatriation is the scenario where a workload is created in or migrated to cloud only to be migrated back to the organization's data center.

Unexpected costs are not the only drivers for repatriation in the market. Service levels and security gaps are also driving repatriation. This is not an issue with the capabilities of public clouds per se, the problem is organizational maturity not public cloud capability. Business units and IT organizations are underestimating the people and process changes necessary to operate cloud environments effectively. Today organizations lack technical and business operational maturity in the public cloud environment with the internal IT skill sets that are trained for in house data center deployment and management. Over time organizations will get better at delivering high levels of service and ensuring secure operations in the cloud.

Repatriation is not always a negative outcome. In fact, repatriation can be an expected outcome in a healthy cloud computing operation. One of the greatest benefits of the public cloud is the speed with which new products and services can be prototyped. Advanced cloud operators leverage the public cloud as an inexpensive way to design and deliver new services quickly. This gives a "time to market" advantage as well as limits invested capital in ideas that have yet to be proven viable.

REPATRIATION IS A THING, CONT,D

New services that prove to be viable in the market can continue expanding in public cloud to scale or relocate to private alternatives. The unit economics, service level and security capabilities of mature data centers are still very compelling. The decision to move often comes down to the cost of relocation and a comparison of relevant features between venues.

TECHNICAL DEBT IS HOLDING YOU BACK

Over the last 10 years, virtualization has disrupted the way organizations deploy and manage applications in the on-premises data centers. The benefits were clear: higher utilization of IT capacity resulted in a lower total cost of ownership. As such, the vast majority of organizations operate the majority of their applications in a virtualized environment.

One unforeseen consequence of the extremely successful adoption of virtualization in the data center is the impact of virtualization on the software lifecycle. Prior to the widespread use of virtualization technologies, organizations would upgrade their hardware and software as an integrated process. Post virtualization, it was no longer necessary to upgrade your software when you upgraded your hardware. This has resulted in many organizations accumulating significant technical debt as software updates were no longer a mandatory component of refreshing the underlying hardware.

Delaying software updates creates significant risk for an organization. As software approaches end of life, the software manufacturer offers fewer security patches and updates. Furthermore, when software is unsupported, the software manufacturer stops providing security patches and updates all together. Therefore, a continuous process of software upgrades is essential for minimizing the potential for and impact from cyber security threats.

The majority of companies interviewed reported 60-70% of their operating systems were either running an unsupported or soon to be unsupported version. [2]

Virtualization provides greater levels of reliability and availability by abstracting the software layer from the hardware layer in a technology deployment. As hardware ages it becomes less reliable.

While leveraging virtualization best practices do help organizations avoid the risks associated with a hardware failure becoming a business outage, delaying hardware upgrades does put an organization at risk.

TECHNICAL DEBT IS HOLDING YOU BACK, CONT'D

Refreshing hardware regularly ensures an organization benefits from the security, performance, and reliability enhancements that are continuously being introduced into the market by hardware manufacturers.

The current generation Intel® Xeon® Scalable processors optimized for both public and private cloud deliver up to four times VM density performance and lower infrastructure TCO by up to 60%, when upgrading from a 5 to 6 year old server. [3]
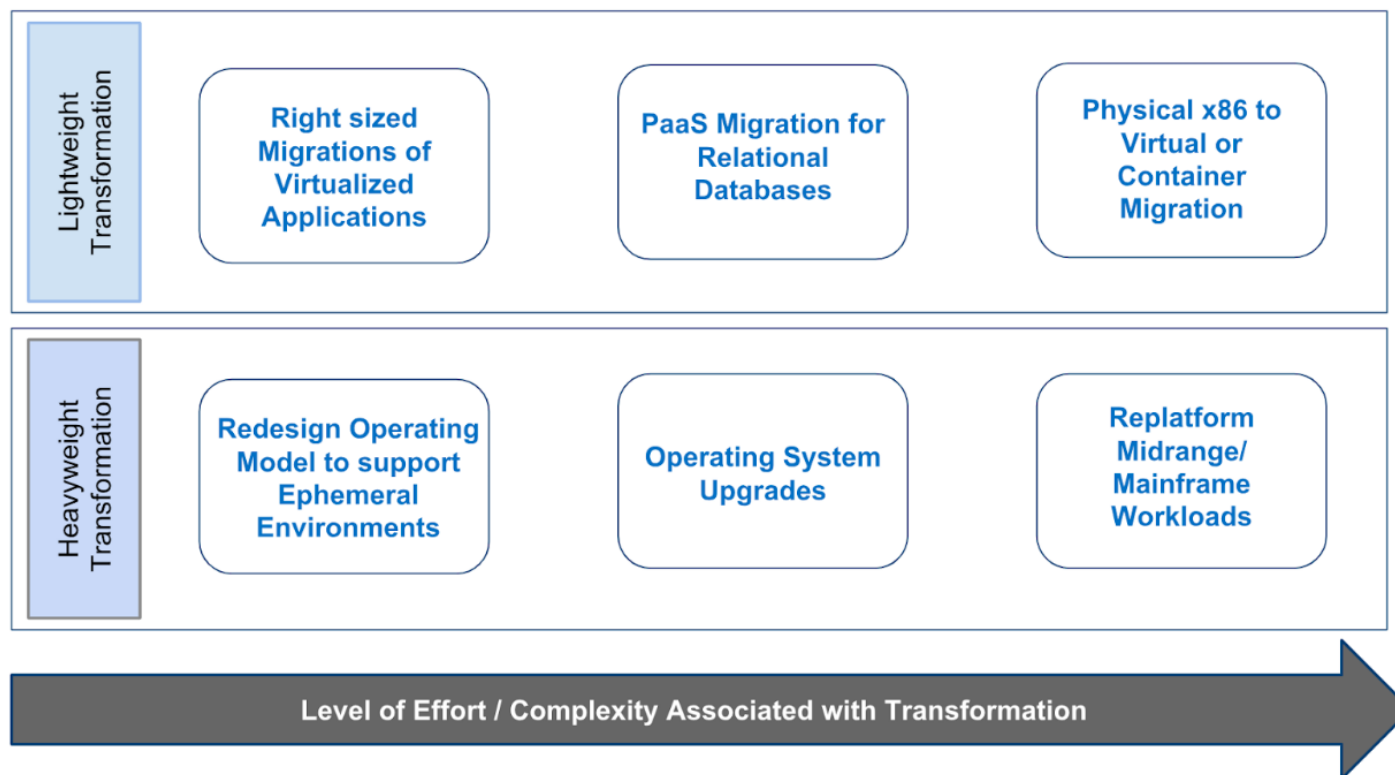
Virtualization is not the only source of technical debt. Many organizations are also running significant mission and business critical workloads on legacy physical x86 and non x86 platforms. These legacy platforms tend to be more expensive when compared to cloud alternatives but are also extremely reliable and highly performant. The catalyst for change in these environments is more often business risk than cost.

In addition, mainframe and midrange systems implemented in the 1980s and 1990s were delivered on proprietary hardware architectures and platform specific software. The lack of open standards and modern software frameworks make it difficult to maintain and support these systems as the world has moved to open and extensible application architectures. Furthermore, the greatest risk comes in the form of a lack of available skilled resources to operate these environments moving into the future. As such, most organizations must find ways to lessen their dependence on these legacy platforms and must proactively eliminate this technical debt from their business operation.

Paying down your technical debt requires a well thought out strategy that balances quick wins with long-term investments. Technical debt has been accumulated over decades of technology adoption and evolution. This debt will not be eliminated in a calendar quarter and it could take considerable time and planning to eliminate. In addition, the acknowledgement that technical debt will be an ongoing issue is critical to a business' long-term success. A debt reduction plan needs to be a continuous process. If organizations are not proactive, the debt they eliminate from their traditional IT operation will reappear in their cloud environment.

Modern software runs best on modern hardware. Waiting too long to refresh hardware or avoiding regular upgrades to software adds risk to the security and performance of an organization's workloads while also delivering suboptimal total cost of ownership. It is important to understand the benefits of modernizing hardware and software together in both private and public cloud.

REDUCING TECHNICAL DEBT THROUGH TRANSFORMATION

Lightweight Transformation

- Right sized Migrations of Virtualized Applications
- PaaS Migration for Relational Databases
- Physical x86 to Virtual or Container Migration

Heavyweight Transformation

- Redesign Operating Model to support Ephemeral Environments
- Operating System Upgrades
- Replatform Midrange/ Mainframe Workloads

Level of Effort / Complexity Associated with Transformation →

LOOKING FORWARD: DATA DRIVEN DECISIONING

Moving forward, more organizations will acknowledge that the complexity and variability associated with making cloud decisions requires a data driven decision making process. Historical approaches to technology procurement and operations are being disrupted by the proliferation of infrastructure, platform and software-as-a-service (SaaS) alternatives available in the market. There is now an array of options to consider ranging from modernizing the data center to hosting applications in co-location or regional data center providers, to leveraging public cloud services. Each venue presents its own advantages and disadvantages with regards to agility, economics, service level, and security. Organizations will need to invest in defining their own blueprint for multi-cloud operations which includes both governance of supplier use as well as policies that guide desired business outcomes.

DATA DRIVEN DECISIONING, CONT'D

Organizations will need to develop policies and procedures that govern the use of supplier offerings. One of the greatest advantages of utilizing the public cloud is the ability to quickly procure, provision, and operate technology workloads. This agility also presents a risk to the business. The consumerization of IT service procurement and operations makes it easy for a line of business to launch a new service without considering the impacts associated with supplier proliferation, compliance, audit, and long-term total cost of ownership. With a well-documented policy governing cloud usage, a business can provide the agility desired without adding risk to the business operation.

"As adoption of public cloud accelerates, organizations must dedicate resources to govern public cloud consumption to ensure a compliant and secure operating environment. Organizations should also integrate workload placement and code execution venue decisioning into their DevOps pipelines and continuously analyze their cloud environments to rapidly identify and remediate security risks."

- Charles Gautreaux - CTO, Cloud Services - AIG

A strategy must be built to continuously evaluate opportunities to achieve an agile, secure, and highly performing IT operating environment. Many organizations are beginning to discover that it is not uncommon for services and offerings to vary significantly across different data center locations and public cloud instances for the same service. As such, organizations need to adequately execute performance testing of their applications in the venue where they plan to deploy to ensure quality. In addition, certain security and technical capabilities are only available in certain geographic regions. Therefore, to be successful in making sound workload placement decisions you must be aware of the variability of service level, security and technical capabilities across provider data centers and across providers.

The most successful organizations implement a continuous process of discovery and analysis to ensure maximum value for their technology spend. Organizations need to invest in automation and tools that will help to drive a consistent experience across public and private cloud data center operations.

DELIVERING CONTINUOUS OPTIMIZATION OF TECHNOLOGY SPEND

The following workflow is to be followed to deliver continuous optimization
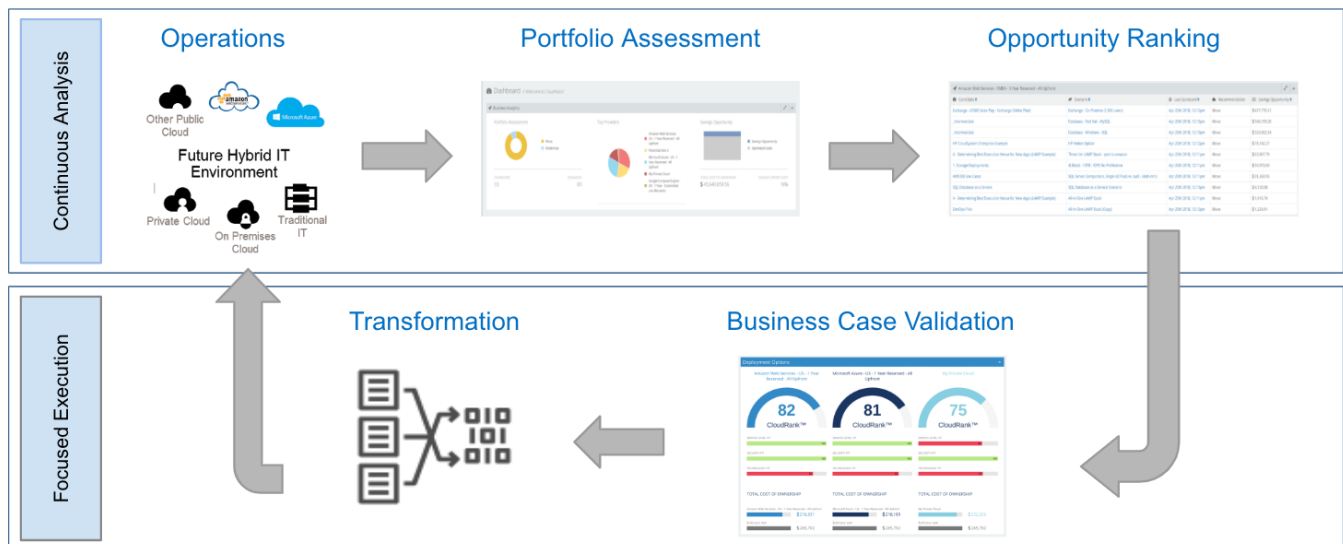
BALANCING THE IT BUDGET

Transparent pricing continues to be a driver for cloud growth. The consumption-based pricing provided by the public cloud is an alternative to traditional IT pricing. Consumption based pricing allows for greater transparency into the unit economics of the IT spend. Cloud pricing also eliminates the need for upfront capital investment, making it more attractive when compared against traditional on-premises cost models.

Developers and line of business application owners alike prefer the transparency of public cloud but are also finding it difficult to predict their ongoing costs due to the frequency with which public cloud pricing changes.

In 2018, the top 3 public cloud vendors by revenue (Amazon, Azure and Google) changed their pricing 150+ times, increased the number of infrastructure and platform options to choose from by 475.96%, introduced multiple new licensing models, opened 22 new data center regions and added support for more than a dozen new security certifications. [4]

BALANCING THE IT BUDGET, CONT'D

The public cloud market moves fast. As such, it has proved challenging for organizations to predict their future public cloud spend as well as manage their current spend. While the pricing of the public cloud may be transparent, the variability and complexity associated with consumption-based models available in the market are confusing to the consumer long term. To deal with these challenges, organizations must implement spend management and continuous spend optimization strategies to ensure the business receives the best value for their IT spend.

The economics of traditional IT lacks transparency and, while it is not nearly as variable or complex as the public cloud, it is also expensive when compared to cloud alternatives. That said, there is tremendous potential to deliver both transparency as well as improved economics in the data center.

First, organizations must begin adopting a software-defined data center architecture powered by hyper-converged infrastructures to drive down the costs to operate application infrastructure in the data center. Hyper-convergence eliminates unnecessary costs associated with integrating and operating disparate storage, networking and compute environments.

Second, organizations should begin to embrace pay as you go pricing on-premises. Many of the hardware and software OEMs now offer this model as an alternative to large upfront capital investments traditionally associated with a technology refresh. Embracing consumption based pricing in the data center will showcase the lower cost unit economics available in the data center when purchasing or leasing technology at scale.

MULTI-CLOUD IS THE FUTURE

As organizations reduce their technical debt, improve the economics and cost transparency in the data center, and establish best practices for making workload placement and optimization decisions they will uncover their right mix of on-premises and off-premises technology investments. The data shows that most organizations will operate in a multi-cloud environment leveraging the best that private and public cloud can afford.

To succeed with a multi-cloud strategy, organizations will need to invest in packaging, automation and tooling that provides continuous optimization through intelligent workload placement and decouples applications from the infrastructure that supports them. Another significant trend heading into 2019 is the adoption of containerization as a multi-cloud enabling technology. Containers are providing an abstraction layer on top of the multi-cloud infrastructure that enables consistent and reliable execution regardless of the venue from which the containers run. This technology is the next innovation in a line of evolution that decouples the application from the data center where it is executed. Containers present the opportunity to make an application portable and greatly increase workload mobility across data centers and clouds.

A sound multi-cloud strategy also ensures the right competitive motivation in the market. While empowering your business with agility and choice, a business with a multi-cloud strategy ensures there is enough competition in the market to ensure that suppliers continue to innovate and price their offerings competitively.

When implemented correctly, a multi-cloud strategy will deliver a consistent experience regardless of where a data center is located and who operates that data center.

CONCLUSION

Cloud computing is an operating model, not a target state. The path to a successful cloud operation starts with an honest assessment of an organization's readiness to embrace the changes required to receive the benefits cloud can afford. This journey may take several months or years depending on the size and complexity of the organization and their existing investments in technology. Once correctly implemented, a multi-cloud strategy will improve the unit economics and agility available in the data center while also driving intelligent use of public cloud services.

Here are actionable recommendations to ensure cloud computing success:

1. Focus on assessing application workloads and their dependencies not on infrastructure. Each application should have a treatment plan that includes some combination of recommendations to resize, rehost, refactor, replatform, rewrite, retire and retain.

2. Implement a workflow to continuously assess and optimize technology investments. The workflow should provide a decisioning framework for both new workload placement as well as optimization of existing deployments

3. Execute a program to assess and eliminate technical debt. The program should prioritize both lightweight and heavyweight transformation targets to deliver long term reduction in risk and improvement in total cost of ownership.

4. Establish governance around public cloud usage to ensure a compliant and secure operating environment. This includes an investment in tooling and automation to integrate workload placement and code execution venue decisioning into DevOps pipelines.

5. Invest in packaging and automation solutions that decouple applications from the underlying infrastructure that supports them. To realize the full benefits of multi-cloud, application workloads must be portable across execution venues.

ABOUT CLOUDGENERA

CloudGenera provides vendor agnostic information technology analysis that allows organizations to choose the technology solution that best fits their needs. Headquartered in Charlotte, North Carolina, CloudGenera is one of the fastest growing technology companies in the United States and is internationally recognized as a thought leader in cloud transformation. The company's highly visual software automates cloud decisioning for a clear view of available private, hybrid and public cloud solutions, eliminating costly and error-prone manual evaluation methods and speeding enterprise time to digital transformation. With CloudGenera, some of the world's largest organizations have transformed their technology operations to realize the business benefits of cloud. For more information, please visit *www.cloudgenera.com*.

CITED PAGE

1

Gartner calculates 2018 public cloud spend to be $175.8 Billion USD. Wikibon calculates 2018 private cloud spend to be 32.0 Billion USD.

2

CloudGenera surveyed public, private and government organizations, analyzing 100,000 plus application environments in 2018 to determine current or future state costs

3

Performance results are based on testing as of June 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.   For more complete information visit *www.intel.com/benchmarks*. Configuration details:

Per node 4X higher integer throughput performance: estimate based on SPECrate*2017_int_base on Intel internal platforms as of June 2018: 1x node, 2x Intel® Xeon® Processor E5-2690, 128GB total memory, 16 slots / 8 GB/ 1600MT/s DDR3 RDIMM, Benchmark: SPEC CPU2017 V1.2, Compiler: Intel® Compiler IC17 update 2, Optimized libraries / versions: IC18.0_20170901, Other Software: MicroQuill SMART HEAP. uCode: 713, OS: Red Hat Enterprise Linux* 7.4, Kernel: 3.10.0-693.11.6.el7.x86_64 x86_64, Score 65.5 vs. 1x Node, 2x Intel® Xeon® Platinum 8180 Processor, 384GB total memory, 12 slots / 32 GB / 2666 MT/s DDR4, Benchmark software: SPEC CPU® 2017, Compiler: Intel® Compiler IC18 OEM, Optimized libraries: AVX512, ucode:0x043, Red Hat Enterprise Linux* 7.4, 3.10.0-693.11.6.el7.x86_64, Score: 281.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction. Configuration details: Up to 60% TCO savings with Intel® Xeon® Scalable processor compared to 5-year old system. Example based on estimates as of June 2018 of equivalent rack performance over 4-year operation on integer throughput workload (estimate based on SPECrate*2017_int_base on Intel internal platforms) running VMware vSphere Enterprise Plus on Red Hat Enterprise Linux Server and comparing 20 installed 2-socket servers with Intel® Xeon® processor E5-2690 (formerly "Sandy Bridge-EP") at a total cost of $737,460 [Per server cost $36.8K: acquisition=12,5K, infrastructure and utility=4.5K, os & software=10.2K, maintenance=9.7K] vs. 5 new Intel® Xeon® Platinum 8180 (Skylake) at a total cost of $294,540 [Per server cost $58.9K: acquisition=12,5K, infrastructure and utility=10.1K, os & software=10.1K, maintenance=9.7K].

Assumptions based on *https://xeonprocessoradvisor.intel.com*, assumptions as of June 6, 2018.

4

2018 statistics referenced in this paper were compiled and analyzed by CloudGenera CloudResearch. CloudGenera CloudResearch continually analyzes the qualitative and quantitative capabilities of 400+ cloud service provider catalogs worldwide.