

PUBLIC CLOUD MIGRATION WITH KEMP APPLICATION DELIVERY FABRIC



INTRODUCTION

The IT industry has seen its share of technological revolutions, from PCs and the World Wide Web to the Internet of Things and everything in between.

Today, we are on the brink of another disruptive paradigm shift as an increasing number of businesses are starting to transition to an on demand digital business model to maintain a competitive edge. While cultural and operational adjustments are needed to make this happen, the melding of the appropriate emerging technologies is just as important. A key supporting technology pillar that enables this transformation is cloud computing. The only way to be as agile as possible and respond to scale and intelligent data requirements in near real-time is to leverage an infrastructure built for it. By their very design, public clouds such as Microsoft Azure, Amazon Web Services (AWS) and Google Cloud Platform satisfy these requirements in an on demand fashion.

While here are technical motivators for transitioning applications and infrastructure to the cloud, such as resource elasticity challenges and difficulty to procure and implement new systems, cloud migration as an enterprise agreed approach is almost always a business decision. The objective benefits of the organization vary but often center around:

- **CAPEX Reduction**
Unlike the install-and-own capital model of on premise infrastructure, the cloud offers pay-as-you-go subscriptions that shift up-front expenses from the customer to the service provider.
- **OPEX Reduction**
Economies of scale and efficiency allow cloud providers to offer a variety of attractive pricing models for organizations running large numbers of applications.

Additionally, cloud allows resource costs associated with operation and management of the underlying infrastructure for services to be shifted to the vendor.

- **Business Growth**
Planning for increased customers or employees typically means additional strain on infrastructure. It is almost always more economically sound to leverage cloud to support this than leasing or building physical data center space.
- **Business Continuity**
Enterprise approved CIO-led cloud initiatives often start as a way to satisfy business continuity requirements related to disaster recovery and resilience for critical applications.
- **Agility**
Organizations in the cloud can deploy new services in minutes and hours as opposed to days and weeks as well as scale existing applications on demand in order to respond to requirements customer demands more rapidly

Whatever the motivating factor might be, the shift to the public cloud is undeniable. According to Gartner research, Amazon Web Services and Azure combined yearly revenue has reached ~\$14B. In 2016, RightScale, a SaaS-based cloud computing management provider, conducted the largest survey to date on the use of the cloud infrastructure. The results revealed that cloud adoption increased from 63% to 77%, and 95% of organizations surveyed are running applications or experimenting with infrastructure-as-a-service. According to Gartner analyst Daryl Plummer, “adoption of the cloud is rising rapidly, with no sign that it’s going back”. As was the case with virtualization, despite perceived blockers, a catalyst has already been triggered that can be reverted. A recent Gartner survey backs up Plummer, with 90% of respondents saying they are conducting some form of cloud computing and 78% saying they plan to increase cloud spending through 2017. An IDC survey also recently noted that “savvy CIOs now see the cloud as being an extension of their sourcing strategies.”

CLOUD OPTIONS

When evaluating cloud, enterprises have a number of options to choose from in terms of the model they will leverage for their approach. These are categorized into 3 main branches:

- Software-as-a-Service (SaaS)
- Platform-as-a-Service (PaaS)
- Infrastructure-as-a-Service (IaaS)

SaaS is a delivery model where software is consumed on a subscription basis from a provider that centrally hosts the infrastructure. A common example of this would be Salesforce. In order to move legacy applications to this model, data transformation and consulting is normally required, especially when custom in-house applications are involved. The benefits

here is that the provider in this case is responsible for SLA and all infrastructure and core code management.

PaaS is a type of cloud computing where a platform is provided by a cloud provider for development, operation and management of applications without having to manage the actual underlying infrastructure. An example would be Amazon DynamoDB which provides extremely low latency NoSQL database services for mobile, web, gaming and IoT workloads. Another example would be Azure App services which provides a framework for building and running .NET, NodeJS, PHP, Python and java apps. In both cases, the responsibility for management, maintenance and configuration of the underlying infrastructure is on the cloud provider, in these examples, Microsoft or Amazon. The benefits of this approach is agility for launching new applications and services without the challenges associated with infrastructure management. By the nature of its design, scaling on demand also becomes frictionless. Economies of scale also generally make this approach financially attractive for new application development.

IaaS provides virtualized computing resources accessed over the web. Azure Virtual Machines and AWS EC2 VMs would both be examples of IaaS. This model still abstracts the customer from the underlying physical infrastructure but does require management of virtual resources. The benefits are greater control, reduction in capital investments and fairly simple approach to expand as you grow. For many organizations today, this is the model used for transitioning existing applications into the cloud.

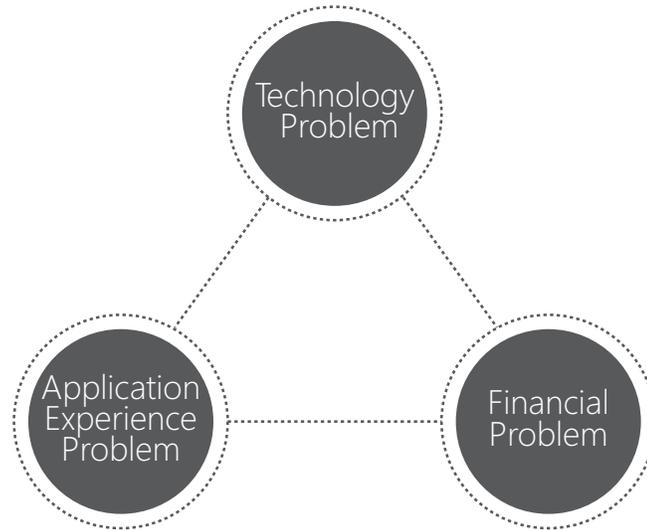
While there are a number of cloud providers offering infrastructure-as-a-Service (IaaS) such as Google, IBM, VMware and others, Microsoft Azure and Amazon Web Services (AWS) dominate the space. The majority of others have specific value for specific use cases but presently lack broad applicability. Still, even with Azure and Amazon's maturity and feature depth, they still fail to meet the requirements for the easy delivery of all traditional applications that an organization may be looking to transition. This is a challenge since by and large, enterprises today are still in lift and shift mode for their workloads to IaaS as opposed to full re-architecture for PaaS.

ENTERPRISE CHALLENGES

While cloud clearly offers modern enterprises many benefits and the technology has come a long way, a transition to it is not without its challenges. These range across technology alignment, application cloud readiness, visibility, lack of in-house expertise and data sovereignty. And while cloud is almost always more cost-effective at a macro level, when vendor capabilities are required that are not provided by the chosen cloud provider, the price of getting into the cloud can significantly grow. Overall, the various challenges can be placed into three main problem groups

Technology Challenges

The prospect of migrating tens, hundreds or thousands of applications with limited service interruption is daunting at best. In addition to the application itself, data needs to be migrated, there may be new plug-ins that require integration for environmental compatibility and the infrastructure layer in the cloud may not offer the same features and capabilities that were instrumented on premise. With the application delivery controller (ADC) typically



serving as the publishing point for applications and providing critical functions for availability, ensuring that these requirements are met in the cloud platform is a key consideration as well. At this point, native ADC capabilities in the majority of public cloud environments are not suited for all traditional, multi-tier enterprise applications. And because the operating concepts for on premise ADCs are significantly different from

how they operate in the cloud, additional issues can arise after the migration has been completed.

Financial Challenges

When native cloud capabilities do not meet the needs of the applications being migrated, costs associated with deployment or migration of 3rd party systems can be high. A way to address this is to re-architect the application to match the capabilities provide by the target cloud environment. This can be costly in terms of project management with estimates up to \$250,000. Time to market is also likely lengthened for an application migration project when this becomes necessary.

Application-experience Challenges

Customers expect the same quality of experience and reliability from applications they consume regardless of where the application is located or how it's connected – they shouldn't have to care. However, application performance and response times are less predictable in the cloud, more complex to design and more challenging to troubleshoot when issues do arise. For traditional applications migrating to the cloud this can be even more challenging and may require capabilities not offered natively in the cloud environment without the introduction of 3rd parties.

CIO CONCERNS

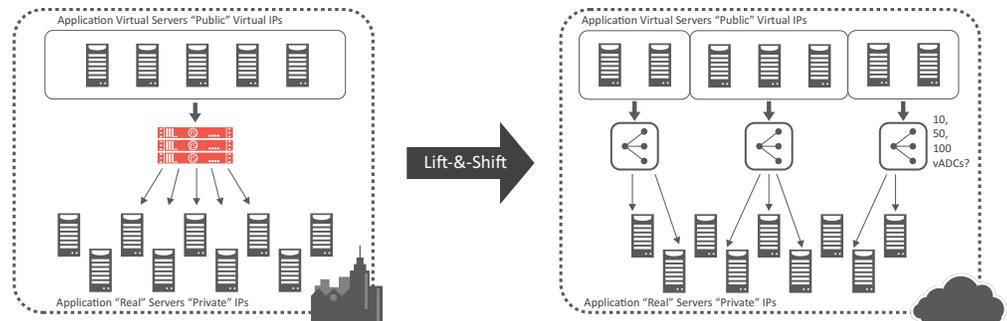
In addition to the enterprise challenges noted, planning a move to the cloud also raises concerns among CIOs, who realize that migration of existing business-critical packaged applications is not as easy as flipping a switch. This is often exacerbated by the concerns over security, vendor lock-in, liability and visibility. As an example, regardless of whether or not applications end up in cloud through a sanctioned change management process or because of a rogue business unit, it's ultimately the CIO's problem if something goes wrong. Fortunately, 3rd party vendors and cloud providers themselves have been responsive to these challenges and developed offerings that help to address these. As an example, a variety of cloud management platform vendors have developed to help with governance and policy enforcement across a heterogeneous cloud infrastructure. While data privacy is still sometimes seen as a blocker to migration, cloud providers have built their business around a commitment to protecting data and keeping it private. Cost models for running apps in the cloud also often provide a long term advantage. With connected devices and IoT initiatives increasing the attack surface for organizations, CIOs are also finding an advantage with cloud providers who have a depth of security expertise and are often at the forefront of identifying and mitigating the latest emerging vulnerabilities.

A recent survey by Gartner noted that approximately 70% of CIOs surveyed will be jettisoning on premise data centers in favor of the cloud over the next two years. The primary reason given for this shift is that most traditional IT service providers aren't delivering public cloud-related services or products that are suited for digital business transformation.

APPLICATION MIGRATION CHALLENGES

For the majority of workloads an application delivery controller(ADC) provides load balancing functions to ensure high availability and serve as the publishing endpoint for customer access. Most traditional workloads have very specific requirements when it comes to the functionality of an ADC ranging from intelligent session persistence, context-based traffic steering and application-level awareness. Most public clouds offer some type of native application delivery and load balancing functionality. However, these are not always sufficient for legacy applications. The reason is that the capabilities native to Azure and AWS have been designed with cloud native applications in mind which typically have different requirements. As an example, most cloud native applications are stateless meaning that session information is replicated within the application and that a client does not have to maintain a contiguous connection with a single endpoint for a good UX. So, if customer Joe accesses a cloud-hosted app and connects to Server 1, then gets disconnected, Joe can access the application again regardless of the server to which he reconnects to without having a negative user experience. Challenges can arise with packaged applications that were designed with legacy principles that assumed things like statefulness and application layer proxying. In this case an interruption in service will require the application to reconnect to the same server in order to provide users with the same experience they enjoyed prior to the interruption. An application delivery controller is normally used to meet this requirement.

When application delivery-related capabilities not offered by a cloud provider are needed, one approach is to simply “bring-your-own-ADC” by leveraging available 3rd party offerings built for these environments. This ensures that the same capabilities used by the application on premise is present when the application becomes operational in the cloud provider environment. However, this presents another set of challenges that are more financial than technical in nature.



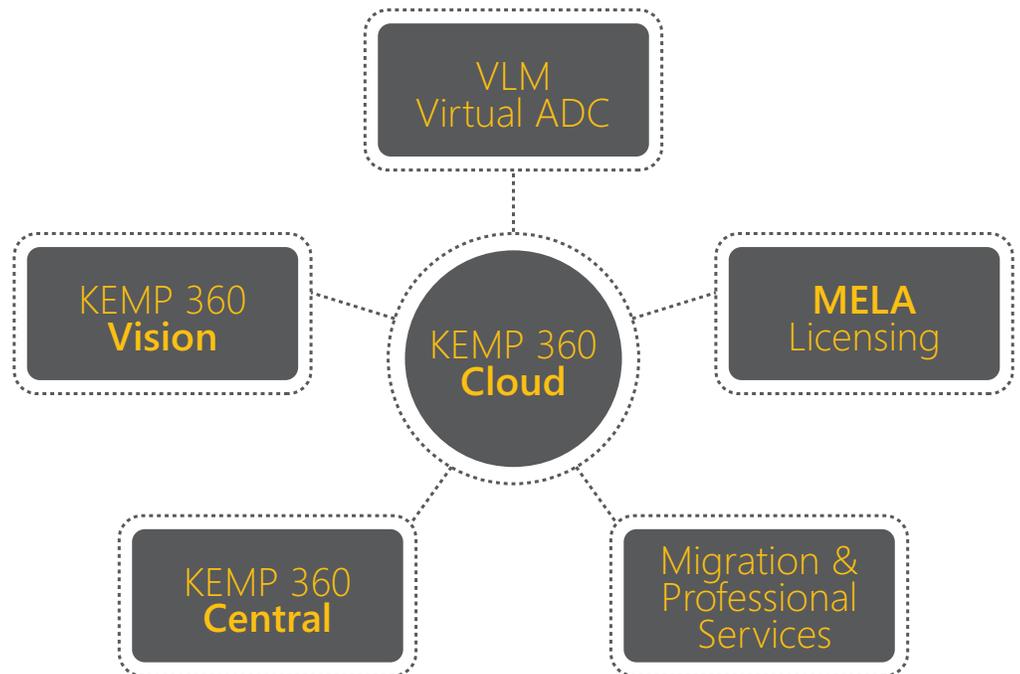
A core driver for moving to the cloud is the flexibility it offers in terms of infrastructure planning and upfront costs. The traditional application infrastructure planning approach requires an enterprise architect to predict what maximum usage will be for the applications in their environment over some period of time, normally ~3 years, and purchase products and services based on that. This often results in periods of over and under subscription often referred to as peaks and troughs. When adoption or utilization is not at the planned levels, the environment has excess capacity which equates to funds that could have been allocated to impactful business projects. When utilization exceeds what was originally envisaged, user experience issues ensue and expensive rip and replace of infrastructure must be performed. Public clouds like Azure and AWS solve this challenge by providing a scale-on-demand environment with pay-as-you-go utility based licensing. If though, you need to bring existing technology from your data center into the cloud such as load balancing, you’re presented with an inconsistency since the consumption models likely don’t match, again requiring you to predict maximum usage over a pre-specified term in order to control costs. While an increasing number of vendors do have cloud marketplace offerings, doing the math will reveal that the 1-3 year TCO is significantly higher than a traditional perpetual license. It is also likely that the number of individual instances required when moving many applications to the cloud will be higher given the differences in network architecture and public IP address allocation. When high availability pairing or clustering is accounted for as well, the costs can be exponential, creating an impasse for customers looking to complete large scale application migration projects. For these reasons, a simple vendor model that is consistent with cloud consumption principles and provides cost control is needed.

APPLICATION DELIVERY FABRIC

KEMP Technologies is a leading provider of application delivery solutions for enterprises, is in the top 3 most deployed ADCs in the industry, top 10 most deployed ISV solution in the Azure Marketplace and was the only vendor in the Visionary category of the 2015 Gartner Magic Quadrant for Application Delivery Controllers.

Supporting customers who are on a cloud journey is a key objective of the company and reflected in their being the first ADC commercially available for Microsoft Azure, continued cloud focused product integration and ongoing collaboration with a global network of cloud system integrators to deliver joint customer value.

KEMP’s Application delivery fabric is a subscription based product and services suite focused on supporting the migration of applications to a cloud operating model. It combines flexible billing models and a simplified bundling of essential application focused services to help make migration projects as simple as possible.



KEMP’s software focus enables delivery of the same high performance and functionality across all cloud and virtualized environments without a dependency on specialized hardware. Hybrid and multi-cloud deployments are easily implemented using integrated VPN and global load balancing (GSLB) for resiliency. KEMP’s application delivery fabric enables customers to deploy KEMP’s application delivery and management technology without concerns over predicting the total capacity required in terms of instance count. This provides greater flexibility and allows for application delivery elements to be consumed in an on-demand cloud-centric model. Where a mix of cloud and on premises platforms is in use, the application delivery fabric integrates each to offer a consolidated view of resource consumption and performance as well as a single point to perform operational tasks such as

software upgrades and configuration backups. Migration of application delivery services between platforms is greatly simplified through a simple drag-and-drop console while configurations and policies are consistently applied across all resources regardless of location enabling greater control.

Insight on current and historical usage is the basis for accurate migration planning. In order to architect a cloud infrastructure for existing applications, it's important to have a clear understanding of the requirements for those applications. With KEMP's application delivery fabric, existing on-premises usage can be collated from both KEMP and 3rd party application delivery instances to allow for correct resource allocation in the cloud. Once a cloud migration is completed, the collation of performance metrics continues and provide the analytics for optimal ongoing performance.

KEMP also offer a unique service where the application delivery infrastructure is integrated with the KEMP customer support to combine early detection and notification of issues along with incident response and management automation. With early detection, the KEMP team can use their expertise and institutional knowledge to diagnose the issue and reach out to the customer with actionable insight to mitigate and remediate the issue. Operational costs are significantly lowered as 24/7 coverage and access to the subject matter experts reduces the need to expand in-house expertise and headcount.

A subscription to KEMP's application delivery fabric affords a customer the benefit of getting the technologies support needed support a seamless migration of existing application delivery capabilities to the public cloud. This is all provided via flexible licensing and consumption that prevents customers from needing to go through tedious à la carte SKU procurement.

As has been highlighted, resource deployment and consumption in the cloud is dynamic in nature and the pricing models that are focused on ADC capacity and number of instances do not meet the demands of application delivery in the cloud. With KEMP's Metered Enterprise Licensing (MELA), KEMP's application delivery fabric recasts the licensing model to deliver unlimited instances of uncapped capacity across all delivery platforms within an organization. For billing purposes, usage across instances and platforms is aggregated on a monthly basis with built in overage terms to accommodate unexpected usage spikes. Since a metered licensing approach is leveraged, there also is flexibility in terms of responding to usage growth and shrinkage without negative financial implications.

WHY KEMP? WHY NOW

At KEMP, we are cloud application delivery specialists – it's our core business and the only thing we do. Unlike "born in cloud" providers, we have the experience of 40,000 actual deployments across hundreds of applications. Unlike "industry gorillas", our goal is to help you deploy your apps where it best suits your business, not where it's most financially

beneficial for us and our stakeholders. Additionally, we are the only provider that has a vision to deliver the application experience management that your enterprise needs.

In today's highly competitive global marketplace, companies must make progress towards transitioning to a digital on-demand business model if they want to be the disruptors as opposed to the disrupted. From a technology perspective, public cloud will continue to play a growing part in making this possible. Organizations now need to align with technology partners that have a clear cloud strategy or replace them with those that have an approach that will better enable them for success.

The question enterprises need to ask today no longer is *if* they should adopt a cloud strategy and seek to transition application services but *how* they will do it and *who* they will choose as technology advisors to help them on the path to greater agility, increased scale and increased capabilities to deliver intelligent and responsive on demand services for their customers.

AUTHORS

Peter Melerud – CMO, KEMP Technologies

Jason Dover – Director, Product Line Management, KEMP Technologies

REFERENCES

https://en.wikipedia.org/wiki/Cloud_computing

<https://support.rackspace.com/white-paper/moving-your-infrastructure-to-the-cloud-how-to-maximize-benefits-and-avoid-pitfalls/>

<http://www.gartner.com/newsroom/id/2685515>

<http://www.gizmoxts.com/transposition/the-need/item/28-the-cost-of-rewriting-applications>