Deloitte.

Cloud and infrastructure
How much PaaS can you really use?

Contents

About the series

Overview

In order to enable new business and preserve existing value, global Information Technology (IT) executives should address the growing dichotomy between the agility companies want, and the stability they need. This dichotomy is becoming exacerbated as legacy data centers become farther and farther removed from the cloud and from mobile end users. Meanwhile, delivering data to devices and taking advantage of platforms often means giving up some control over operating systems, hardware, and data centers. It's 9 a.m. on Monday. Do you know where your data is?

Introduction

As IT executives look to provide value from their IT portfolios, they are balancing a mix of emerging, current, and legacy technologies. With the world of technology ever evolving at pace, the gulf between emerging and legacy continues to widen. The consumer market drives advances in end user devices and end user expectations. Service vendors invest in cloud, virtualization, and orchestration, while manufacturers attempt to deliver more compute horsepower at lower cost. Behind the tablets, clouds, and chips, there still sits a data center and a room full of legacy infrastructure waiting for refresh.

This widening gap between end user devices, data mobility, cloud services, and back office legacy systems can challenge the IT executive to manage and maintain technology in a complex array of delivery capabilities. From mobile apps to mainframe MIPS, and from in-house servers to sourced vendor services, managing this broad range requires a view on how much can change by when, an appropriate operating model, and a balanced perspective on what should be developed and controlled, and what needs to be monitored and governed.

To help equip the IT executive in forming those views and making those judgements we present points of view on key trends and topics.

Background

Many IT executives are always on the lookout for ways to standardize software and IT infrastructure, and to simplify development, support, and maintenance processes. At the same time they consider where there may be competitive advantage, and how best to balance the emerging, current, and aged technologies available. This is particularly true with the recent growth of areas such as analytics, mobility and social media, and the Internet of Things, which could each be addressed via platform services.

Maturing "platforms" are rising above the hypervisor, providing enhanced automation and developer self-service, programming interfaces, and more integrated middleware and management capabilities. These emerging services abstract applications from infrastructure, eliminate the need to buy, host, and operate compute and storage, and offer the potential to simplify and standardize in one go.

But it's unlikely that a cloud platform-as-a-service is the best answer for all use cases, or even whether

an organization wants too many apps in one basket. Even if you want to, your current applications might need significant rework in order to be cloud and platform ready.

So you're potentially left with a mixed bag of cloud, internally optimized, and legacy capabilities—each requiring a different management approach.

The Past: Technology Stacks

In the past, enterprise Chief Information Officers (CIOs) and Chief Technology Officers (CTOs) largely built and owned their technology stack. They would buy or develop business application software that utilized a given IT infrastructure. And they would buy, support, and manage that IT infrastructure. Their operating model would enable application software to be designed, deployed, maintained, and refreshed using a specific stack. And they would deploy, maintain, and refresh that infrastructure themselves. However managing this breadth and depth of the supply chain and its complexity could hinder growth and innovation in an organization.

The Present: IaaS Model

Today, sourcing and governing an Infrastructure-as-a-Service (IaaS) is more common. An integrated IT infrastructure stack is procured, based on a set of standards, and is presented to consumers with some level of automated provisioning, likely with some virtualization and a utility consumption model. This might be provided on or off premise, and in a private or public service. Public cloud IaaS now enables organizations and individuals to leverage computing power without the need to worry about buying or housing IT infrastructure. In an IaaS model,

application development is still oriented to take advantage of a specified infrastructure and there is likely still a need to configure and deploy the business software and infrastructure layer to work with each other.

The Near Future: PaaS

A much hyped but still maturing service is Platform as a Service (PaaS). A "platform" in this sense builds on the IaaS model by providing not only the compute and OS layers, but the server and runtime management and additional automation and orchestration. These provide application software developers with the means to deploy and manage their software themselves with application relevant service levels and characteristics. By having the middleware, operating system, and hardware abstracted from the application software, the organization needs less infrastructure operations, and the application developer spends less time deploying and provisioning, and more time developing valuable software.

PaaS adoption stats
The PaaS market is estimated to reach $7bn by 2018 implying a 5 year CAGR of almost 23 percent[1]

Overview of PaaS

Market drivers

Many organizations today are creating a generation of custom applications that are leveraging the potential of cloud capabilities. They often attempt to deal with the pressures of creating great business value by providing increased flexibility and speed, while managing greater uncertainty. They use approaches to validate, learn, and pilot new capabilities, while balancing risk in new development so they may fail fast, learn from the effort, and move on. All the while, they are dealing with unpredictable requirements and capacity needs.

Trends

In response, SaaS and IaaS providers alike are moving into PaaS to help provide enterprise customers the agility and flexibility they require. For example, AWS and Azure are evolving from IaaS with value added point solutions to a turnkey PaaS offering which is still emerging. Many organizations are trending towards turn-key private PaaS offerings (e.g. Pivotal Cloud Foundry, Apprenda). While the market and products are still maturing many firms wait for the winners to emerge. The market is evolving quickly and we expect significant growth in the coming 2–3 years (15–30 percent growth by category).

Service offering

Like other cloud services, a PaaS can be provided on or off premise, and in a private or public manner. Differing delivery models offer variety in terms of agility and control—where public hosted or managed can provide greater agility, private hosted or managed provides more control.

A PaaS provider may offer abstraction and shared services via a container-based approach, a hypervisor and virtual machines, or both. There is debate about full machine virtualization, or OS level virtualization provided by containers. But as much focus as there is now on containers, from Docker to cluster managers like Kubernetes, there is a place for both. Google for example runs all of their services in containers, and that includes their VMs[2]. Service providers are likely to provide both capabilities.

But by now the IT executive might be less interested in the mechanics of how a PaaS provider manages their underlying compute layer and more interested in the features, characteristics, and price points of the service.

PaaS options should be considered for the level and richness of the abstraction and automation layer. A few questions in your checkbox could be:

What limitations or requirements are there in terms of deploying an application?

Can a developer or other appropriate staff truly "click and deploy" software?

Are the service levels clear and unambiguous?

Are there straightforward but appropriate options in terms of availability, latency, response time, user concurrence, RPO & RTO?

Are there controls over data encryption, key management, and data at rest and in transit?

Are patches, upgrades, or refresh to the middleware, operating system, and hardware transparent or self-service?

Do you need anyone at all to support or maintain anything beneath the application layer?

Considerations for PaaS

Enterprise value
An enterprise value-driven approach is an option to enable PaaS adoption within the IT service operating model. Understand industry trends, their impact on the organization, and how they could impact the operating model. Look to articulate business and organization objectives and goals that will drive positive transformation, and then develop a business case and benefits tracking mechanism to validate benefit realization. Develop a value-driven enterprise capability architecture. These steps enable the organization to look at any PaaS potential being informed by critical business capabilities, and to understand any opportunity to rationalize the applications portfolio aligned to those capabilities. PaaS solutions can then be selected based on the expected benefits vs your overall management of risk.

Feature maturity
To realize PaaS's full potential, consumers should determine if the service itself is fit for its intended purpose.

It can be useful to have differentiated service levels and price points that can be used for application workloads of differing criticalities and service windows. Variability in performance and disaster recovery characteristics, dev/test vs. production, and the associated unit price points ensures a broader spectrum of candidate applications and increased cost variability:

- How mature are the service providers, commercial models, features, and functions?
- What technologies are being used within the service, what is one's view on their quality and features, and do these align with one's needs?
- Are there any proprietary components or features that could lock one in or make moving workloads difficult?
- Are the services standardized across the PaaS offering?
- Is the PaaS offering portable between vendors?

Understanding the gap
PaaS services include items like development tools, middleware, a runtime environment, and self-service options. But there may be entry requirements to satisfy if your workloads are to operate in the PaaS. Application interfaces, support and maintenance approaches, and data management aspects could be defined in a PaaS standard. Understand that the approaches, tools, languages, and technologies are things that can be used, are appropriate, and don't create unnecessary vendor and/or proprietary technology lock-ins. And of course, consider the future development and application needs to meet the rapidly evolving demands on the business.

A gap analysis will be required to take these items into account and compare the enterprise with the supplier in terms of technology roadmap, and application, data, and security architectures.

Key PaaS Figures

Thinking beyond strategic IT

"More and more business leaders are recognizing its [the cloud's] profound implications for how enterprises can make money, differentiate, and compete. Business leaders of all stripes—Finance, Sales & Marketing, Product Development and more—are becoming increasingly focused on the business value cloud provides. Over the next three years, cloud's strategic importance to business users is expected to double from 34 percent to 72 percent, even surpassing their IT counterparts at 58 percent." [3]

Trends in the suppliers

• The PaaS Market is a highly fragmented and crowded market, however over 50% of the vendors are niche players with only a few dozen clients

• Almost 73 percent of the firms are clients of the top 5 providers and about 53 percent are clients of the top 2

• The fast growth rate of the PaaS industry and its strategic relevance to IT is what is attracting investments from both new players as well as established vendors [5]

Trends in the Industry

• Application PaaS holds the highest segment of PaaS market share (35 percent) followed by ADLM PaaS (16.3 percent)

• Mature markets such as Western Europe, U.S. and Japan are leading the PaaS adoption

• Vendor innovation and demands of the user are driving the availability of leading edge capabilities in PaaS which is its advantage over software on-premises

• Private PaaS is seeing the highest adoption rates [4]

Trends in buyers

CIOs primarily buy cloud services for an increase in business agility and speed of business. The following were cited as the primary reason to purchase cloud:

• Scalability & Agility: 50 percent

• Cost Reduction: 14 percent

• Innovation: 13 percent

• Quality & Reliability: 10 percent

• Other: 7 percent

• Financial considerations: 5 percent [6]

Based on the gap analysis, understand what changes may be needed that enable existing applications to migrate and run on PaaS. This could be anything from programming language and SDK support, to data architectures and APIs. Estimate the effort and time needed to make those changes and the effort to migrate and manage the process.

And look for improvement opportunities from the gaps. If a target PaaS provides easier provisioning, have a plan to optimise the need for those skills and processes accordingly.

Technology service providers generally require customers to maintain currency with supported and maintained releases. Where are the PaaS providers going with their particular technology and release roadmap? Understand what will be needed to keep up, or what could change over time in order to maintain software currency with changing PaaS standards.

Dealing with another roadmap

Along that same vein, the CIO will likely be used to laying out a technology roadmap, not being a consumer of someone else's. A public or hybrid PaaS model challenges the CIO to select the available vendor architectures and engineering standards as opposed to defining or developing them in-house. And it changes the model in which you sustain compliance with vendor services over time. The semblance of control over architecting and engineering the middleware and OS/hardware stack decreases, and by definition the infrastructure operations is outsourced.

In the event an organization has a mixed portfolio of PaaS, IaaS, and in-house or legacy technology, its operating model will need to cater for all those lifecycles. A firm will likely need to apply varying levels of architecture and engineering standards to different things. And it can require a combination of people that develop for the PaaS services and self-service their workloads, and for retained services that can manage the full support cycle (develop, deploy, provision, support, maintain, and change applications and IT infrastructure).

These aspects again highlight the need to look at iterative approaches (validate, learn and pilot) and to mitigate risks in learning to fail fast, frequently, and small. Your own roadmap could be based on an agile approach, given that a more standard SDLC may not succeed in adopting nascent PaaS models.

Consider diversification

Be flexible and dynamic with your choice of service providers. Have a target ratio of services to put into PaaS, and with a single provider. Depending on size, geographic spread and corporate risk profile, understand if a single or multi-vendor approach is more appropriate. Ensure that applications are fungible between platforms and understand how workloads can be moved between providers or to future providers.

Public cloud services spending forecast (Gartner, 2015)[7]

Billions of dollars — 2013: $131, 2014: $154, 2015: $175, 2016: $203, 2017: $237, 2018: $274, 2019: $312
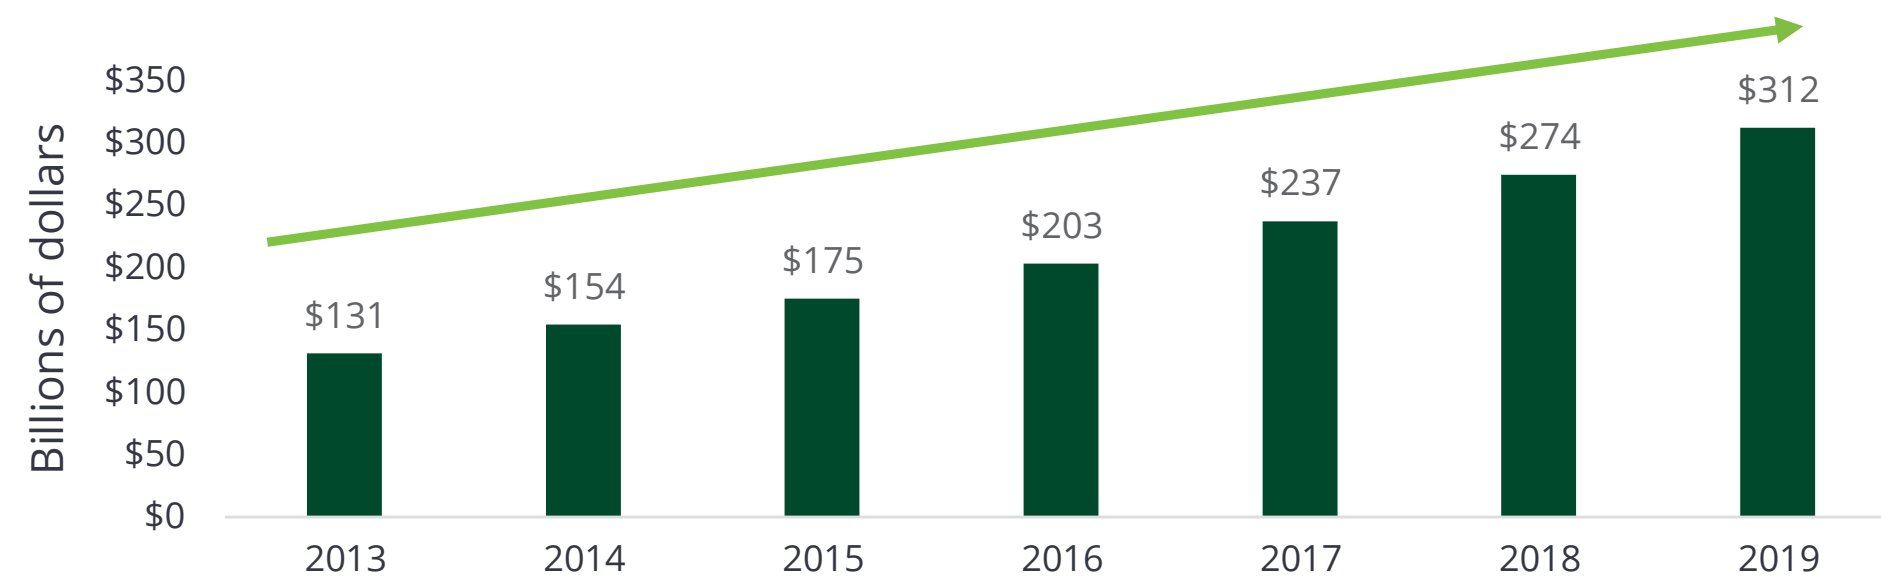
Existing portfolio mix

Looking at PaaS, the IT executives should consider what balance their IT portfolio should have. Do you get competitive advantage from the application functionality, from the technical ability of the IT infrastructure layer, or both together? There may be things that should be highly bespoke and customized, some other things that leverage common elements or technologies, and things that are highly standardized using commodity technologies.

While there may be a real advantage and business case to having bespoke applications on specific technology stacks, these may not be appropriate for a public/shared off premise PaaS model. For example, specialized low-latency workloads and/or those that use specific hardware devices. In addition there may be parts of the landscape that are on legacy technology not readily transferable to a PaaS model—needing too much time and money to worry about a PaaS target state before some intermediate refresh step.

Applications on wintel x86 architectures, perhaps already using virtualization, cross-platform applications already on web and/or Java technology, and those with a stateless data approach are the more obvious parts of the landscape to be early PaaS migrators.

For those things that may not be PaaS targets, you'll consider some other choices to make the most of your IT infrastructure—for example driving up CPU utilization, thin provisioning your storage, or using Virtual Machines or Solaris Zones.

And while it would be nice to think that all your workloads could be serviced in a one-server-to-many-app ratio, the likelihood is that you still have some larger workloads where the ratio is inverted to one app to many servers. So virtualization is not the panacea for all things, as compute grids and standalone high performance servers do have their place. A sample framework to consider is as below, where you map out your infrastructure to the below matrix and take the necessary action.

The forward planning analysis will be "what advantages can one get from PaaS, and what is the ratio of things one can get in a given timeframe?"
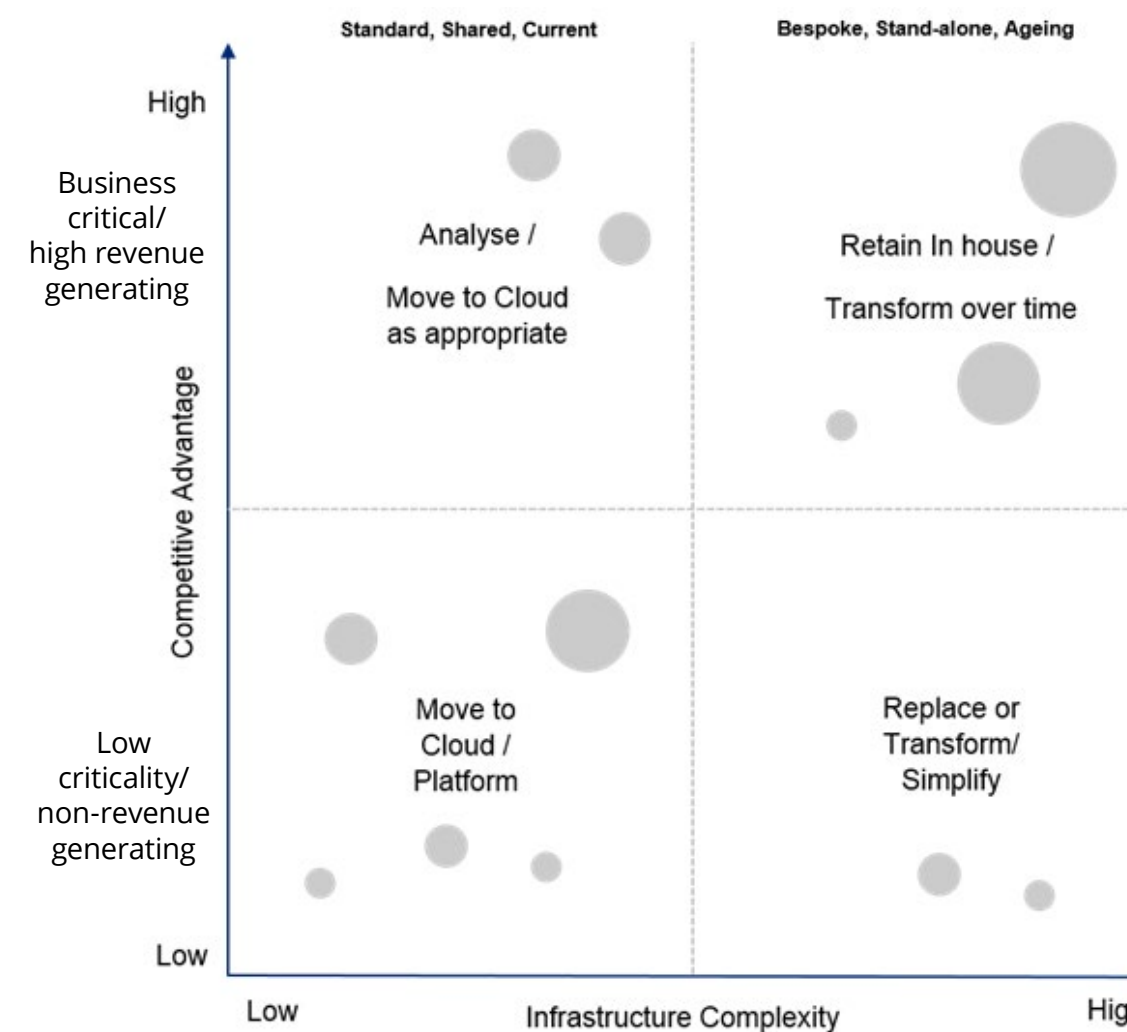
Show me the money

A simple PaaS equation for your business case is:
[(Benefits of PaaS/Proportion of estate that can get on PaaS) − (Cost of entry to PaaS + Cost of exit of existing/legacy + Cost of ongoing maintenance)]

To sort out the more detailed financial calculations, the finance team will need to understand:

· the ratio of applications to put into PaaS

· the potential benefits of optimising the application development and infrastructure operations costs

· the mix of prices and services being targeted

· the effort and time needed to change and migrate

· the ongoing cost of that portfolio mix over time

· the book value, depreciation schedule, and exit strategy regarding assets that may no longer be required

A cost and spend optimization approach can be to align the migration of applications off of assets at the point in time those assets fully depreciate.

PaaS budget share
US CIOs expect their external IT budgets to increase to 4.8 percent and European CIOs expect a growth to 3.7 percent in 2014. Cloud and data analysis projects expected to get the maximum share of the budget.[8]

Assuming that not everything gets to PaaS at once, or even over time, there is a mix of financial scenarios to model and work with. Have a view on what compute, network, and datacentre assets, sizes, and locations are maintained over time, and how they will be used for retained applications and technology. And have a view on the business appetite to book a write-off, or if an asset sale or sale and leaseback model could be considered.

Traditional IT models own IT infrastructure assets and licenses and use them over time. Whether or not an asset is being used 24 hours a day does not affect the capital spent or its depreciation schedule. PaaS service models may offer pricing based on resource units consumed over time. Understand the resource units and how they are defined, measured, and priced.

There may be different service levels with different price points. Have a target for how much of the environment should be in which different service tier. There may be discount tiers for volume consumption, or financial floors or ceilings as commitments. Have a control and governance regime over who, how, and when they can request any new items. And be sure to model the number and frequency of changes to applications that may attract charges.

Many service providers use a number of strategies to differentiate themselves, and to attract and maintain clients. Understanding the provider business model helps the CIO to be aware of motivations and things to consider. The provider may use proprietary technology that could be an advantage, or a hindrance when you want to move

somewhere else. They may require you to keep only on current or n-1 versions, necessitating a regular refresh cycle that could cost. And they might be planning on chargeable changes and growth in unit consumption—betting that you change and add more than you might expect.

There are still rules and risks

If your business operates internationally, then you are likely subject to a mix of different policies, rules, and regulations from different jurisdictions. Information security and personal data protection will be particular focus areas, and you may also find legal, compliance, operational risk, and business continuity interested.

Understand which policies, rules, and regulations are relevant in each jurisdiction. An appropriate level of due diligence should be exercised to ensure the PaaS provider is compliant with internal needs, and the relevant legal and regulatory environments. And understand if the service provider maintains its own policies and if and how it checks and maintains compliance against those requirements.

Maintaining a clear understanding of where business data actually resides and how it flows across borders and between legal entities is a complex imperative. Depending on your legal and regulatory regime, you may be required when using a sourced service to know and evidence where a piece of personal data resides, and who could access it. In a shared cross platform cloud service this raises some challenges. Have an understanding of what the options may

be in terms of where data is stored and processed. Inform yourself on who should, and who could, possibly access your data from which locations. And be aware of the different legal entities of your customers and the primary and sub-contracted vendors in the service delivery chain.

Moving from the virtual to the physical, and depending on the level of scrutiny you may require, have a perspective on the service provider physical and technical information security. Review the buildings, locations, shared and dedicated infrastructure components, and how secure they are.

Where there are shared infrastructure components, understand in which way they are shared, what this means for you, and how that is controlled and managed. You may also look at the access rights, controls, and segregation of duties related to how the services are provide and supported. Who can access which components for what purpose under what authority and control? Who has access to instances of things hosting my applications, and who has access to my data? What are the controls and procedures over who can access what? In using a PaaS at any scale in your enterprise it is important to exercise appropriate governance and control. Monitor and review performance standards, incidents and outages, policy and regulatory compliance, and be sure to verify consumption and the bill. Depending on the size and risk profile of the endeavour, consider a dedicated governance and control function to look at these and other aspects.

One method of computing your risk would be to categorize risks, assign them values, calculate risk

scores, and then assess the investment required for the treatment of those risks. A go/no-go decision can be made, as we compare the scores of the residual risks to that of the risk treatment.

Sample risk treatment model[9]

Legend: ■ Category residual risk ■ Risk treatment
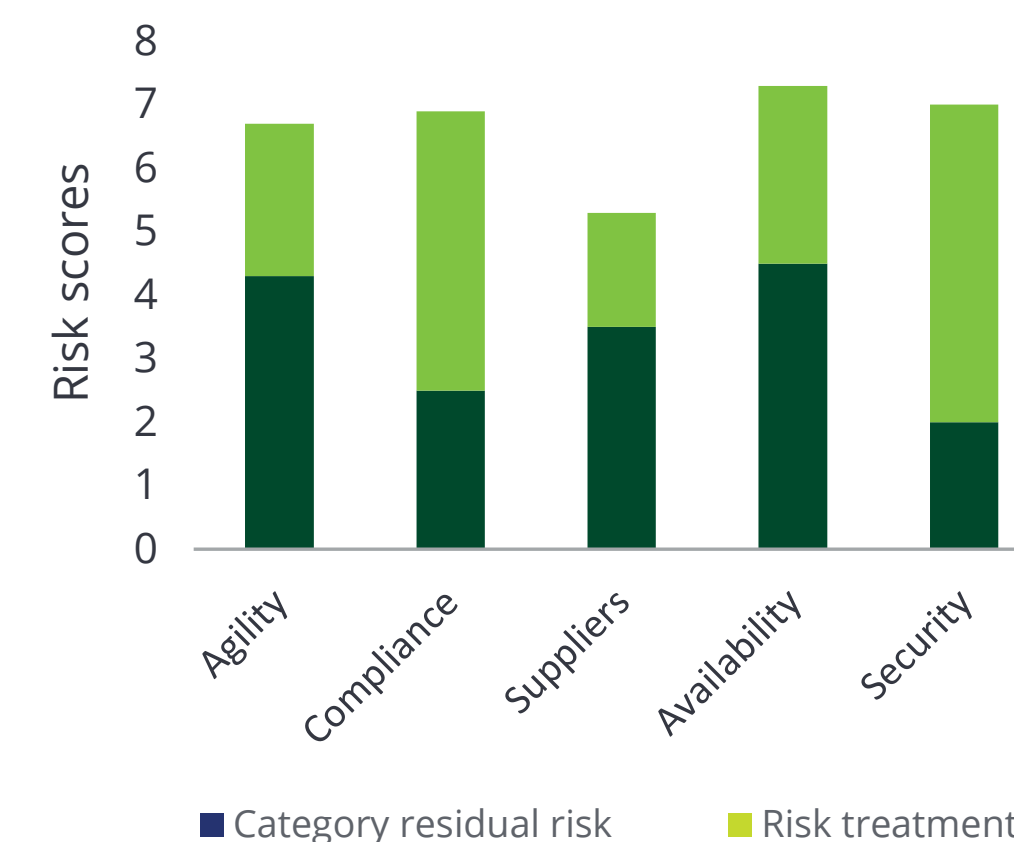
People and operating model

Having determined how much PaaS can be consumed, and what the transformation and migration might look like over time, determine the right balance of people, skills, and locations for the target operating model. Review which skills and management capabilities are needed, which to increase or decrease, and how to fulfil those needs.

First and foremost the organizational culture should become one that is more inquisitive and that thrives on change. Some aspects to consider in terms of a people strategy include having development skills for

leveraging emerging platform services, scripts, and languages—and maintaining development skills in current and legacy technologies. Ensure that there are skills in house to review, verify, and influence supplier technology and architecture choices, as well as those needed to service internally run and managed systems. From an IT infrastructure perspective, vendor governance, client and demand management skills may need to be supplemented, while the datacentre and systems operations workforce is adjusted as needed. Consider the value of where IT resources are located and why, and adjust the location strategy and footprint accordingly. And be sure to have the right people and processes in place to manage any transition and transformation execution plans that might be required.

Cloud employment trends
Employment in cloud topped 18 million globally in 2014, with China leading the recruitment space. Cloud related jobs are expected to grow by 26 percent in 2015 [10]

To the right, we look at some of these same items and a few more as consideration for a target operating model:
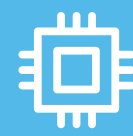
Have the capability to develop, manage, and support both retained legacy and new platform oriented services.

Have an architecture and model that can move quickly and flexibly with the changing technology, vendor and service landscape.

Ensure appropriate levels of vendor management and governance—be an intelligent commercial buyer and manager/governor of services and service levels.

Have someone focused on consuming and paying less (the service provider wants you to consume and pay more).

Understand and control application interfaces, data storage, residency, and flow between applications, databases, internal and external source, vendors, and services.

Be able to support a lengthy and complex transition and transformation.

Understand and govern both the sourced service and the technology stack. Don't replicate the services of the service provider.

Ensure there is demand management— monitor requests and consumption of services that cost money.

Understand and manage the financial impact of consumption and service over time.

In any scenario where you are have an element of sourced services in your environment, make sure "the implicit becomes explicit". Everything needs to be clear and documented when dealing with a service provider.

Conclusion

We have presented a point of view where the maturity and service coverage of PaaS can be considered as part of your portfolio, but it's likely not quite right for all of your needs. And we've covered some details on what to look out for and suggestions on mitigations in terms of the service, cost, and risk dimensions of managing your technology organization. It's likely that for some time you continue to balance a portfolio of emerging, current, and legacy technology. Which means you need to have a view on what is happening beneath the services layer inside the IT infrastructure. Two important and expensive items are in providing and hosting your compute capability. The next article in this series considers developments in compute power and data centre approaches, and questions if some key historical trends have broken down.

For more information

Ranjit Bawa
Principal
Deloitte Consulting LLP
rbawa@deloitte.com

Stephen Swartz
Managing Director
Deloitte Consulting LLP
stswartz@deloitte.com

Nitin Tandon
Principal
Deloitte Consulting LLP
ntandon@deloitte.com

Contributor

Richard Pone
rpone@deloitte.com

Endnotes

[1] Mynewsdesk. (2013, February 18). *Global PaaS market: $7 billion industry by 2018*. Retrieved from http://www.mynewsdesk.com/uk/pressreleases/global-paas-market-7-billion-industry-by-2018-837920

[2] Wilkes, J. (2015). *Running billions of containers at Google with Borg*. Google Cloud Platform.

[3] Firstforcloud. (n.d.). *Cloud Computing—A Paradigm Shift Within the World of IT*. Retrieved from http://www.firstforcloud.com/index.php/about-us/cloud-computing/

[4] Natis, Y. V., Pezzini, M., Tapadinhas, J., Petri, G., Guttridge, K., Thoo, E., & Murphy, J. (2014). *Predicts 2015: Private, Public and Hybrid: PaaS Advances*. Gartner.

[5] Pezzini, M., Natis, Y. V., Malinverno, P., Iijima, K., Thompson, J., & Thoo, E. (2014). *Magic Quadrant for Enterprise Integration Platform as a Service*. Gartner.

[6] Biscotti, F., Natis, Y. V., Pezzini, M., Malinverno, P., Thompson, J., Cantara, M., & Murphy, J. (2014). *Market Trends: Platform as a Service, Worldwide, 2013-2018, 2Q14 Update*. Gartner.

[7] Nag, S., & Lai. (2015). *Forecast: Public Cloud Services, Worldwide, 2013-2019, 3Q15 Update*. Gartner.

[8] Perez, J. C. (2014, January 14). *CIOs project healthy IT budget increases for 2014*. Retrieved from http://www.computerworld.com/article/2487695/it-management/cios-project-healthy-it-budget-increases-for-2014.html

[9] Proctor, P. E., Plummer, D. E., & Heiser, J. (2015). *A Public Cloud Risk Model: Accepting Cloud Risk Is OK, Ignoring Cloud Risk Is Tragic*.

[10] Anderson, C., & Gantz, J. F. (2012). *Climate Change: Cloud's Impact on IT Organizations and Staffing*.