RBA

A Digital and Technology Consultancy

Gaining a Competitive Advantage With Platform-as-a-Service (PaaS)

Gaining a Competitive Advantage With Platform-as-a-Service (PaaS)

The benefits of the cloud are widely known by IT professionals, and to realize those benefits most corporate IT groups have made at least some efforts to move their applications and services to the cloud. However, the most impactful benefits for companies in moving applications to the cloud aren't technical in nature; rather, it's that a cloud-based platform enables new business opportunities—opportunities to drive more revenue, enable better collaboration, and streamline business processes to reduce costs.

But moving systems and applications to the cloud is rarely a straightforward process and is hindered by confusion over terminology and options for migration. This paper seeks to demystify the options and make a case for why Microsoft Azure is the right cloud platform for most companies, and why RBA is the right partner to help companies make the transition.

Defining "The Cloud"

There are a lot of misconceptions about what "the cloud" really is. For many, it simply means that an application, server or database is hosted at an off-site data center and accessed over the Internet, as opposed to residing on hardware on a company's premises. There are many reasons to host in the cloud, ranging from cost-efficiency, reduced maintenance, and better scalability. In some cases, the infrastructure is owned and managed by the company, while in others it is owned or managed by the cloud operator.

In this model, called **Infrastructure-as-a-Service (IaaS)**, a company's systems are essentially "lifted" to the cloud as-is and the physical and network layers are provided as a service. For example, a company could take an application server that once existed on the company's internal network and place it in the cloud and the company will have no awareness or concern as to the type of hardware, the manufacturer of the disk drives, or the type of network it is running on. As far as the company is concerned, it's still a similar instance running the same application and more capacity can be easily added by requesting more instances.

Another cloud model is **Software-as-a-Service (SaaS)**. In this model, software applications are written specifically for the cloud to provide extreme efficiency. For instance, Microsoft Office can be installed on an individual computer or local network, or companies can use Microsoft Office 365, the cloud-based SaaS version of the same family of applications. Both versions of Office provide the same applications and functionality, but each are architected differently. In addition to being a technical model, SaaS is also a business model for applications; access to an application is typically priced on a monthly or per-user basis.

What makes SaaS very approachable and economically feasible is a cloud model called **Platform-as-a-Service (PaaS)**. PaaS is about decomposing workloads and processes into discrete components, and then mixing-and-matching the PaaS capabilities to enable those components to build, run and manage powerful cloud-based applications.

It's important to clarify what the cloud isn't. It's not a place. It's not a single dedicated computer running your and only your application or service. Rather, it's a shared set of resources spread across a number of physical systems that collectively execute that application or service. This understanding is important because it is, for many, a paradigm shift.

Critical Advantages of PaaS

There are several interrelated advantages of PaaS over traditional methods of building applications.

PaaS provides extreme scalability

It's important to move away from the perception that those individual components—which are technically known as services—reside on a specific piece of hardware. It's not as if, for example, a server running one service has to communicate with a server running another service in the same data center. PaaS is a virtual model in which one physical server may be running multiple services—or just one.

What's important is that PaaS removes concerns about physical server scalability. The services scale—not the hardware. PaaS is a collection of services—not a collection of hardware components. The catalog of services available in a PaaS platform like Microsoft Azure enable all kinds of capabilities—making applications accessible to mobile devices, Big Data, even machine learning. Microsoft is continuously releasing new services that make it progressively easier to add compelling new functionality.

Suppose you're building an application that will need to store large amounts of data for tens of thousands of your employees. Off the bat, you might assume you'll need a farm of expensive SQL database servers. But in a PaaS model, you don't need to worry about that. You could deploy a small number of less expensive servers running a "NoSQL" database, which in a PaaS environment cost-efficiently scales "sideways" as your needs grow.

PaaS is on-demand infrastructure. As load increases from more users or more activity per user, additional infrastructure resources—more memory or processing power, more database or file storage, more network connections—can be "turned on" immediately to accommodate the increase. Traditionally, software developers would attempt to forecast the maximum demand condition and deploy infrastructure to accommodate that maximum. This resulted in overbuilding. And, it meant that if the maximum were ever exceeded, additional infrastructure would need to be stood up—but that would take time, and would require a transition plan for the system, potentially requiring down time and/or code changes.

PaaS circumvents these issues by focusing on the smaller building blocks, which can be scaled more quickly and smoothly than big blocks. If your application needs more memory to meet performance requirements, you would no longer

buy an entirely new server—you'd increase the number and/or size of instances that would provide the additional capacity, with the click of a mouse.

PaaS is flexible

A PaaS platform like Microsoft Azure is open—it's agnostic about what it runs and what it communicates with. An application using a database could be written in a number of different languages, using a variety of tools, and your choice of operating system. Microsoft Azure, for example, works with Microsoft Windows or Linux, and provides relational and NoSQL databases provided by numerous vendors (Microsoft, Oracle, Hadoop – just to name a few). Your choice of data storage is made completely independently and in fact, you can swap one language or tool for another in building the application without needing to modify or impact the data storage. Likewise, services can be continually added, removed, or modified as needed to accommodate shifting business requirements, and done so in a way that doesn't necessitate re-architecting the entire application or worrying about hardware impacts.

Another example of flexibility is that some infrastructure resources can be blended in what are called **hybrid cloud solutions**. As an example, a company might have a database that contains highly sensitive information the company is reluctant to move off-premise. But, employees need to be able to access that data remotely through a mobile application. A company could build a cloud-based application using PaaS services that accesses that database. In short, not every component needs to reside in the cloud. In fact, an on-premises application might run five concurrent workloads and two of them could be split off and run in a PaaS model. From a user perspective, this would all be seamless.

PaaS is cost-efficient

Services can be thought of as building blocks. And with Microsoft Azure, companies pay only for the building blocks they need to build the applications they want. If you need raw storage, for instance, you can buy just that; you don't need to buy all the tools and other things that come with a traditional instance of a relational database engine like SQL Server or Oracle.

Equally important is that PaaS enables cost savings in other ways. For example, suppose you have $1 million allocated to hardware upgrades for your on-premises infrastructure. You could instead invest some of that in redesigning applications for the cloud with a more nominal amount, say a few thousand per month, for the cloud infrastructure. In that approach, the up-front costs are significantly lower, and your company can take the savings and invest it in product development.

PaaS affords business agility

The modular nature of services means that applications can be built and deployed quickly and cost-effectively. This enables businesses to launch and test not just new applications but entirely new ways of doing business. The biggest risk historically has actually been success. There are numerous examples of new products or services that quickly achieve incredible success that cannot be sustained due to insufficient resources or support to sustain that success. The initial success is quickly countered with dissatisfied customers

or wannabe customers that lose trust and confidence that cannot be recovered.

Earlier, we discussed an application that would support tens of thousands of employees. The exact same application that would support 100,000 users would support 100. So, a company could build an application and test it with those 100 users. If it works well for those 100, it could be assumed that it will work equally well for 1,000, or 100,000.

This means that companies can try low-cost, low- or no-risk proof-of-concept experiments in building applications for employees, partners, or end customers. It empowers companies to be much bolder about what they can try to bring to market, how they can influence their business ecosystem, and the value they can deliver to their shareholders.

Consider the payment processing space. It's a hyper-competitive, fast-growing market that's fueled with innovation. It also requires significant computing power. Historically in this setting, any company attempting to enter this space with a novel solution would face significant risk in the time and money necessary to build its own IT infrastructure from scratch. Today, it could instead build a proof-of-concept in cloud infrastructure and test it in the marketplace. If it a failure, it can be scrapped and the total investment lost would be fairly modest. If the results are mixed, adjustments can be made quickly. If it's a runaway success, it can be scaled quickly.

Or consider a company that wants to establish a joint customer-focused program with another partner. This requires quite a bit of data communication and integration, which makes it ideal for the cloud platform. Suppose that the program is a resounding success and now the company wants to expand the program to other partners. If the application infrastructure was designed correctly from the outset, scaling the application to include other partners will be a straightforward and cost-efficient undertaking.

Companies can even run parallel experiments to see which yields better results. Perhaps a company is building a customer-facing application and debating the merits of two different business models for the application, which in turn dictates different sets of functionality. A company could build out both options, determine which one customers prefer, and then scale up that one while shutting down the other.

Drivers for Cloud Migration

There are many reasons a company might begin to consider leveraging the cloud for existing applications and systems. A common one is that an application's existing infrastructure is nearing its end of life (EOL). Perhaps the servers it currently runs on are now being too taxed to run efficiently and need to be upgraded. Or the server might be fine, but its warranty has expired. But, rather than buy new hardware, it might be a good time to evaluate whether the application should be rebuilt in the cloud. The same is true of servers running older operating systems. Or, an application that might be getting outdated; if it needs to be rewritten, anyway, the option to write it with modular, reusable components deployed in a PaaS model should be considered.

Another set of "trigger events" might be business changes. One company might have

merged with another and similar applications should be merged. Business models evolve and new partnerships or departments might result in new expectations for applications that were never anticipated during their design. The biggest trigger to move to the cloud is the desire to launch an entirely new capability or service. Any customer-facing initiative should strongly consider leveraging the cloud platform to maximize the value and flexibility of what you can quickly bring to market today and down the road.

Debunking Objections to Cloud Migration

There are several common obstacles to embracing the opportunity of a PaaS model. One is the entrenched paradigms. This manifests in a few ways with statements such as:

- "We've always done it this way. Why change it now?"
- "We've invested a lot in our current on-premises infrastructure. We want to preserve that investment."
- "It might not be perfect but it ain't broke, so let's keep it."
- "It's my job to manage this equipment. If we move to the cloud I won't have a job."

Of course, the counter to all of these objections is that the cloud opens up new possibilities. What a company did or spent in the past is just history. The important question is what the company will do now and in the future to set it itself up for greater business success.
Another common concern to moving applications to the cloud are perceived risks to mission-critical infrastructure. For instance, "What if the data center loses power or Internet connectivity?" or "Can't hackers more easily steal sensitive business data if it's in the cloud?"

These kinds of concerns are usually misplaced. In reality, it's far more likely that a company's on-premises infrastructure will be subject to power loss or a loss of Internet connectivity than a Tier 4 data center would. And if on-premises servers fail, there may be no backup; in a PaaS environment, the servers can reliably failover to a working one.

And all data that is in any way connected to external systems is subject to potential loss or theft, regardless of physical location. The true level of risk is at worst identical, or at best favors a PaaS provider that invests heavily in protecting its infrastructure. Most corporate IT teams, despite their best intentions, simply can't keep up with combatting the latest methods of data intrusion.

Approaches to Cloud Migration

Starting with a SaaS platform can be the safest approach. It's already built and tested, and generally supported by another company that invests in its continued development. The functionality of some SaaS platforms can be extended with customization or integration to other systems. If that platform can meet at least most of your needs, it's worth considering doing that rather than deploying a new system with PaaS or IaaS.

"Lift and shift" is rarely, if ever, the ideal option. If a company has an application designed for on-premises hosting and wants

to move it to the cloud, it could purchase cloud-based virtual machines that matched the hardware it runs on today and copy disk images to those machines. But this approach could incur additional costs in remediation work which might mitigate any cost savings but provide no performance or functional benefits. None of the advantages of having a "true" cloud-based application would be realized.

Not everything needs to be moved to the cloud, or moved at one time. As discussed earlier, hybrid cloud models mean that certain parts of your infrastructure can remain on-premises or hosted elsewhere in an IaaS environment, while other parts can be deployed in PaaS environments. By moving only pieces, or by moving pieces one at a time in a well-planned and orchestrated migration process, technical and business risks and disruptions can be minimized.

The Migration Process

The first step in the migration planning process is to determine if a portion of the application's workloads will be moved to the cloud (a hybrid model), or the entire application.

In many cases, a hybrid model should be employed, at least initially. In that case, the first step is to identify all the workloads an application has and segregate them into components. Some components can (and should) be moved to the cloud to achieve better performance, more flexibility, add new functionality or realize other benefits, while others should stay where they are. One reason a portion of an application might stay on premise might be its inability to efficiently work in a distributed model. For instance, a transactional process might require sub- second response times from a server or database, and the existing application might not provide a reliable data access connection over Internet Protocols (IP) that are used to connect in the cloud. Another reason is that your connection to the cloud might not be reliable enough to hit that requirement every time. The ability to meet user experience expectations and Service Level Agreements (SLA) are an important consideration that comes into play here.

To also help determine the scope of the migration, it's important to consider interdependencies. Sometimes moving one component can adversely impact another, such as by breaking a reporting process that involves data stored in multiple databases. Sometimes a user is routinely does something unusual that you're not even aware of, and moving components to the cloud might disrupt that undocumented process.

Once you know what should be moved to the cloud, the design phase begins. Here, there are myriad decisions that need to be made about how best to deploy the components in the cloud. Some of the areas that need to be considered include performance, functionality, user experience, the product roadmap, and integration points.

As an example, Web servers are one of the easiest components to move to the cloud. Few changes are required, and anything changed can be rolled back relatively easily, if necessary. This might be a good candidate to migrate first, especially when there are minimal interdependencies on other systems. A second step might be to redesign the way

the Web server communicates with databases (the services layer) so that the size and number of the databases available to it becomes more easily scalable—a critical advantage of PaaS. A third step might be to actually move the data from, say, a single large SQL database to cloud-based, non-relational databases for faster execution.

Your current infrastructure should remain where it is while the cloud components are designed, deployed, and thoroughly tested. Part of the design process should include a transition/release plan to switch off the old pieces and stand up the new ones (or the entire application if a hybrid model will not be used).

The ease with which you can move over the Web and the services really comes down to how well-designed and current the existing system is. Migrating older systems with large relational databases will be trickier than if, for example, you have implemented a robust service-oriented architecture that allows you to move components and layers with minimal impact to the rest of the system.

Why RBA is the Right Partner

Suffice it to say, planning a migration to the cloud is not a trivial task, especially if you are seeking to maximize its value. And before any components are migrated, applications that currently exist on-premises should be evaluated for improvement. The cloud opens up tremendous opportunities for making applications better—simply re-creating what you already have represents a missed opportunity for improvement. RBA can help with every aspect of this process.

It starts with strategy, not technology. RBA will learn about your business objectives and work with you to map out a plan of how to capitalize on the benefits of the cloud to achieve those objectives. All cloud options are considered (SaaS, PaaS, IaaS, hybrid), and we have in-depth experience with each. We have been successfully implementing Microsoft Azure solutions since the platform was introduced so our team consists of some of the initial thought leaders in cloud-based design and development.

Our team has hands-on experience implementing cloud solutions for clients using a proven methodology. We've seen what works—and what doesn't. We've overcome roadblocks. We can bring that expertise to bear on your project.

Ultimately, the cloud is about delivering business value through the deployment of services and experiences that delight your employees, customers, partners and other stakeholders. Partnering with RBA ensures you'll get there in a way that minimizes risk and expense and maximizes your success.

RBA
A Digital and Technology Consultancy

RBA is a digital and technology consultancy with roots in strategy, design and technology. Our team of specialists help progressive companies deliver modern digital experiences backed by proven technology engineering.

Minneapolis, Minnesota
294 Grove Lane E,
Suite 100
Wayzata, MN 55391
Phone: 952.404.2676

Dallas, Texas
1431 Greenway Dr.,
Suite 1000
Irving, TX 75038
Phone: 972.573.0995

Denver, Colorado
4582 South Ulster Street Parkway,
Suite 1350
Denver, CO 80237
Phone: 720.235.3420

www.rbaconsulting.com