White Paper

Considering Cloud Pricing Models and Microsoft Azure

Sponsored by: Microsoft Corp.

Al Gillen Larry Carvalho
April 2018

IDC OPINION

The industry has concluded that cloud computing offers significant benefits of cost and agility compared with on-premises computing environments. In IDC's conversations with customers, CIOs increasingly bring up the challenge of having good visibility to cloud pricing. List prices offer a good starting point when it comes to evaluating one or more public cloud solutions. However, customers must consider their needs in a more holistic manner and should think of cloud computing solutions as a broad portfolio of technologies. As such, customers not only should consider price as a key factor but also should consider the wider range of technologies they will consume and how smooth of a transition they can make from existing on-premises installations to a mix of on-premises and cloud as they navigate what is likely to be a years-long transition to cloud computing.

Microsoft's portfolio has grown dramatically in recent years, and today the company offers solutions that generally match those of its competitors on price for commodity services. However, leveraging customers' installed base of Windows Server licenses makes it possible to offer customers a hybrid cloud solution that provides significant cost advantages.

In summary, price is an important consideration, and a cloud service provider without a compelling story on base costs will lose ground to competitors in many market segments. We believe Microsoft is well positioned to compete in this dimension. However, most customers need to take a larger view and consider cost as one dimension and consider the value-added services they would consume, the full range of cloud capabilities, and the ability to operate a hybrid environment for a years-long transition while leveraging existing investments throughout this transition period.

SITUATION OVERVIEW

In yesterday's world of on-premises computing, customers typically heavily overprovisioned servers to support theoretical maximum workload deployments, leading to an overall utilization rate that would typically be well below 20%. While virtualization software helped customers improve utilization rates, even in a virtualized environment, few customers were able to achieve utilization that exceeded 40–50%, in part because of widely variable consumption based on the time of day, month, and year and the need to have servers on hand for peak demand periods. The infrastructure capital expenditure associated with datacenter investment remained relatively static, regardless of how heavily it was utilized.

In contrast, modern cloud deployment solutions offer much greater granularity, empowering customers to select the exact capacity they need for specific time periods and dial deployment and consumption — and the associated costs — up and down as business needs change over time. Customers are faced

with two important decisions: the consumption model they will choose to utilize and the economics of their architectural choices.

That sounds simple enough, but it's surprisingly complicated to make deployment decisions because of the incredible variety offered by vendors. Let's first consider the choices available to customers.

Cloud Economics and Pricing Models

Multiple cloud pricing models are available in the market today. Within each category of offerings detailed in this section, most cloud service providers offer a multiplex of offerings that consider the currency of the underlying CPU, the number of virtual cores or virtual CPUs, the amount of RAM and storage required, and the associated payment terms. Customers must first select the general approach they want to use and then drill down deeply into the exact configuration that will meet their requirements.

For most customers today, a mix of infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS), and software-as-a-service (SaaS) resources will serve their needs. For large enterprises with a substantial investment in traditional datacenters, the need for IaaS and PaaS solutions helps drive their economic decisions. This IDC White Paper focuses primarily on IaaS consumption, but customers desiring PaaS resources will find compelling pricing in that dimension as well, with similar pricing models in use for container-based and other PaaS environments.

Pricing models interesting to clients include the following:

Introductory Offers

- **No-cost trial instances.** Most cloud providers offer customers and prospective customers limited-length, limited-capacity free access to cloud resources, so customers can deploy an experimental application without any up-front financial commitment. This makes it easier for developers or administrators to try out a cloud environment without engaging management or needing to secure a purchase order. No-cost trials represent a great on-ramp for customers to gain some hands-on experience without making any specific commitments to a given cloud service provider.
- **Development/test instances.** Development/test instances used by developers generally include more resources and longer use periods than no-cost trial instances. These developer instances sometimes are a low- or no-cost benefit of a larger overall relationship with a given cloud vendor. In some cases, cloud providers offer accompanying benefits such as discounts (trial or longer term) on related products or free access to other software, development tools, and other cloud resources and services. We see the use of low- or no-cost test instances as providing a fair return on investment versus making a longer-term commitment for resources through an operational subscription. Taking this approach enhances the ability of an organization to try multiple initiatives.

Ad Hoc Consumption

- **Pay as you go.** Pay as you go is no-commitment billing for use of cloud computing services. Users typically pay only for the resources they consume, but the rate they are paying per time unit tends to be substantially higher than that of other subscription types, such as reserved instances. Pay as you go is a cost-effective strategy that requires no up-front investment. The use of pay as you go has application for most workload types such as spontaneous projects, short-lived deployments, and certain long-term workloads.

Committed Consumption

- **Reserved instances.** A reserved instance is a commitment for a specific period for a given workload or service. The rate per time unit tends to be considerably lower than that of pay-as-you-go services, but a reserved instance is expected to be billed for 24 x 7. It is effectively an analog to having a server running around the clock in a customer's datacenter. In cases where customers have a steady state application that needs to be "always on," such as a database, this can be the best approach, although an application need not be operated around the clock to be more cost effective than pay as you go. Where customers benefit versus an on-premises deployment is they can size the reserved instance appropriately and rely on pay as you go or other solutions to handle known peak demand periods.
- **Dedicated host.** Customers can subscribe to a dedicated server (host) that is theirs to use for a predetermined period of hours or perpetually as a reserved resource. The benefit of using a dedicated host is that customers having software licenses that are tied to a server can consume that resource because the licenses will be deployed on a known host with a known configuration (number of processors, number of cores, memory, etc.) that may meet the software licensing terms and conditions. This approach make sense from a software licensing perspective but loses some of the prime benefits that cloud computing brings to customers in terms of agility and ability to purchase only the computing resources they need. Customers have dedicated use of that server, meaning if they are underutilizing the server, there might be no way to rebalance that workload with the capacity it requires.

Managing Capacity Shortfalls and Excess Capacity

- **Ad hoc capacity boosting.** Most cloud service providers offer some form of on-demand capacity boosts and use different terminology to describe this solution. In some cases, there are offers that include virtual machines (VMs) that are, by definition, lower capacity and able to burst to additional VMs when required. Other cloud vendors offer other unique approaches to acquiring and disposing of ad hoc capacity. Solutions include those that can allow workloads to be quiesced by the cloud vendor based on capacity availability.
- **Returning excess capacity.** The opposite scenario from ad hoc capacity boosting is the release of unused reserved instance capacity. As with ad hoc capacity boosting, there is no standard approach to releasing capacity. Some clouds offer a guaranteed return value, while others require users to sell capacity back to another customer. In both boosting and returning of ad hoc capacity, it behooves customers to consider how likely they will be to need to acquire or dispose of additional resources.
- **Batch computing.** Some cloud providers offer access to spare capacity that can exist in each region, for a significantly reduced price, for off-hours computing needs. Optimizing consumption of this spare capacity can significantly reduce costs, especially where there is less urgency to complete a job. Applications where reduction of cost is important and applications that can tolerate nodes being terminated fit well here. Some high-performance compute tasks could fit in this category.

Bring Your Own License

Another dimension that is inconsistent in the industry is the ability to bring your own license to cloud. The combination of independent software vendors and cloud service providers may impact the options that are available. If they are one and the same, portability options may be offered that are not available for a multivendor scenario. In some cases, customers are required to reserve an entire server to be able to reuse an existing license for a cloud deployment. In other cases, there may be options to migrate a license from on-premises to a cloud deployment scenario, allowing customers to acquire cloud computing resources at a lower cost.

CONSIDERING MICROSOFT AZURE

Microsoft Azure is one of the leading tier 1 cloud platforms available in the market, allowing customers to build and deploy solutions on a worldwide basis. Microsoft Azure has another benefit in that it is the only public cloud environment that has an operational analog in Azure Stack, which is available for customers to deploy on-premises, making it the only consistent hybrid cloud in the market.

The Microsoft Azure cloud offers customers a broad variety of services, ranging from common IaaS VMs, multiple open source products including Linux, PaaS environments, support for containers, functions, and a variety of data services. Microsoft has made tight integration with its Visual Studio and Visual Studio Team Services products a priority for the Azure cloud to ensure development and deployment experiences are smooth and seamless for developers and DevOps practitioners across a range of operating systems, programming languages, frameworks, databases, and devices. And like other hyperscale cloud providers, Microsoft is aggressively adding new machine learning/deep learning, artificial intelligence, and IoT services, and the company continues to develop other emerging services.

For professional developers looking to rapidly create enterprise back-end services to connect to any platform or device, Microsoft delivers an enterprise-ready, fully integrated, event-based, and serverless compute experience. Developers benefit from an integrated service with security, broad programming language support, built-in DevOps support, and an ecosystem of prebuilt triggers and bindings. Microsoft is able to deliver this experience on a global scale to IT customers because of its experience building apps and services supporting its own compute, search, and software products and services.

Differentiations for the Microsoft Azure Cloud

Like other cloud service providers, Microsoft has a broad portfolio of services that matches its key competitors "one for one" in many dimensions. But the company also offers unique differentiation, which at times is related to the company's extensive software portfolio and deep relationships with enterprise customers, and other benefits that are more directly related to the company's willingness to push the envelope to undercut the competition with a stronger offer. While pricing is a complicated topic that includes a multitude of dimensions that must be considered, IDC research has found that Microsoft's pricing models are, at face value, competitive with alternatives. But for most enterprise customers, there can be deeper cost savings that come from their broader relationship with Microsoft. Customers new to Microsoft can also access some of these same benefits, including start-ups with a limited budget whose goal is to build immediate solutions with no on-premises investments.

- **Azure Hybrid Benefit.** Azure Hybrid Benefit helps customers get more value from their Windows Server and SQL Server licenses. Customers that have Software Assurance (or subscription licenses) on Windows Server products can take advantage of Azure Hybrid Benefit. Depending on the specific edition in use, customers can convert or reuse their licenses to run Windows Server virtual machines in Azure but pay the lower base-Linux virtual machine rates. Each 2-processor license or each set of 16-core licenses is entitled to two instances of up to 8 cores or one instance of up to 16 cores. The Azure Hybrid Benefit for Standard Edition licenses can be used only once either on-premises or in Azure. Datacenter Edition benefits allow simultaneous use on-premises and in Azure. Microsoft recently announced the Microsoft Hybrid Benefit for SQL Server offering, which gives customers the option of using Microsoft Azure Data Services such as Azure SQL Database as well as SQL Server for Virtual Machines, depending on the type of on-premises SQL Server license the customer holds. The Azure Hybrid Benefit is unique to Microsoft because Microsoft's Software Assurance contracts are directly related to Windows Server and other enterprise products

used by Microsoft's largest and most loyal clients. This benefit has the potential to lower the subscription costs for cloud instances because customers don't have to procure new enterprise licenses.

- **Flexible Azure Reserved VM Instances.** Microsoft offers reserved instances on the Azure cloud, as do other cloud service providers. Microsoft offers several differentiations with its reserved instance offering, including making it possible for customers to combine Azure Hybrid Benefit with Azure Reserved Instances. This allows customers to obtain the base computing costs for those instances, if they have Software Assurance-covered licenses on-premises that they wish to relocate to the Azure cloud for deployment. Comparing similarly configured VMs between Microsoft and competitors shows similar pricing. It is possible to be slightly above or slightly below Microsoft pricing for generic Linux VMs or for branded Linux VMs. However, when one considers the use of Azure Hybrid Benefit for Microsoft Windows and SQL Server licenses, the Azure cloud can offer significant cost savings versus other clouds. Microsoft offers customers the option to exchange or cancel reserved virtual machines at any time during that period and receive a prorated refund less an early termination fee. The ability to convert reserved instances can play an important role in pricing determination; competitive solutions don't provide that same level of flexibility.

  Cost savings are additive, as presented in Figure 1. Customers that opt to use Azure Hybrid Benefit either alone or in conjunction with Azure Reserved Instances on a three-year term can enjoy substantial savings. Customer cost reductions of up to 40% can be realized from using Azure Hybrid Benefit when placed on a pay-as-you-go VM. When Azure Hybrid Benefit and Azure Reserved Instances are combined, customers can realize up to 82% cost reductions compared with straight pay-as-you-go scenarios.

FIGURE 1

Cost Comparison of Azure Pay as You Go Versus Azure Hybrid Benefit and Azure Reserved Instances Combined with Azure Hybrid Benefit

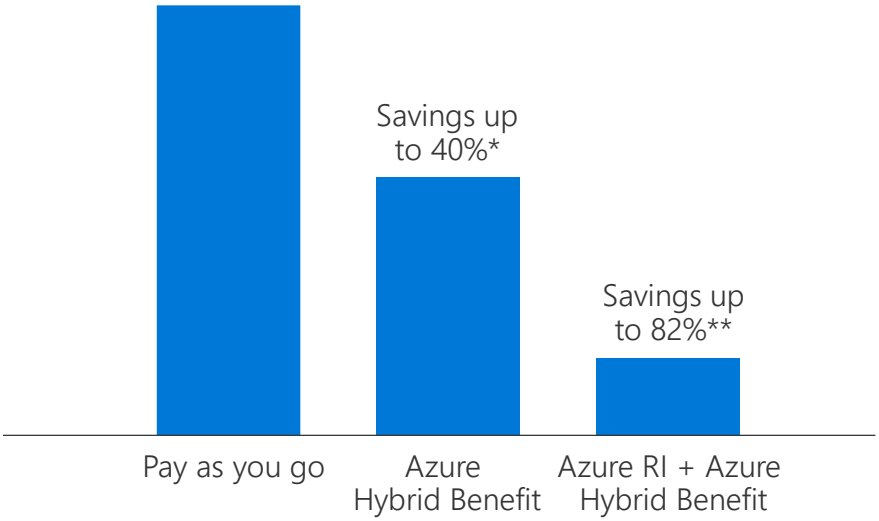

* Savings versus pay as you go

** Based on Dv2 three-year RI with Azure Hybrid Benefit

Source: Microsoft, 2018

- **Free accounts and new customer credit.** Like other cloud service providers, Microsoft offers customers free access to try the Azure cloud. Microsoft offers new customers a $200 credit, which can be used on any Azure product for 30 days, to get them started. Customers can build their solutions with free access to Azure's most popular products for 12 months. This kind of pricing is often called a freemium service and encourages developers to showcase prototypes to management without any investment. Try-before-you-buy offers abound in the IT industry, especially among cloud providers, and have a variety of terms and conditions. Microsoft's integration with its development tool portfolio and its support for Windows, SQL Server, and Active Directory and other products serves as a differentiation for Windows shops considering a move to public cloud. For non-Microsoft-centric shops, Microsoft's support for a broad range of open source software technologies means that customers are likely to have support for their preferred environment in the Microsoft Azure cloud, and the $200 credit for 30 days without restrictions gives these customers additional flexibility.
- **Global scale.** The Microsoft Azure cloud is announced for 50 regions worldwide and is available in 42 regions now, with major installations in North America, Western Europe, Australia, India, China, and Asia. Azure Reserved Instances can be deployed to any region or moved from region to region as needed. Microsoft's worldwide datacenter reach was extensive due to internal utilization of technologies such as Office 365, Bing, and Windows Update, giving the company experience in managing a broad, widely distributed worldwide network. The growth of the Azure cloud added to that network. By comparison, major cloud competitors have fewer regions and zones than Microsoft.
- **Dev/test pricing benefits.** Microsoft offers discounted pricing for development and testing work that applies to development conducted in Azure DevTest Labs and development work conducted using other tools. Most customers today are investing in development and testing of new applications, and they often conduct that work using clouds because of the robust availability of test and simulation resources without requiring a large investment. Microsoft offers an integrated development and test environment through Azure DevTest Labs. Like its competitors, Microsoft offers a variety of ways to consume the Azure cloud for development and test work. Low-cost pay-as-you-go options are open to all customers, and monthly credits are available to Visual Studio customers. Enterprise Agreement customers using Visual Studio can consume VMs at a preestablished discounted rate, and there may also be prepaid Azure monetary commitment funds available as part of those Enterprise Agreements.
- **Cloudyn included.** In 2017, Microsoft acquired Cloudyn, which offered software to help customers track their spend across multiple cloud environments. Now called Azure Cost Management by Cloudyn, this service allows customers to monitor usage, cumulative costs, and comparisons between costs incurred versus budgeted totals. While adopting new cloud technologies, customers often deploy resources in a wasteful manner. Cloudyn enables users to identify inefficiencies in the architecture. This tool provides visibility into spending on Azure services and can also combine that with spending on other cloud providers including AWS and Google. Most cloud vendors offer some form of an economic analysis tool, but in most cases, competitive offerings are focused on that vendor's environment and a comparison versus different deployment options. Microsoft's Cloudyn extends that analysis to look at multiple cloud environments and is free to Microsoft Azure customers.

CHALLENGES/OPPORTUNITIES

With technical change come both challenges and opportunities for customers. Most customers will review the cloud computing portfolio of more than one major vendor as they decide on a given provider for a given workload. Customers need to pay attention to the following areas:

- **Architecture.** Customers evaluating cloud environments must consider the benefits and the challenges they will see from each vendor's architectural model. Challenges may include lack of a hybrid cloud component to ease a transition and pricing models that penalize customers that have made investments in on-premises software licenses. Benefits may include tighter architectural integration, ability to extend existing business relationships at a lower cost, and interconnection of existing infrastructure with cloud solutions.
- **Pricing models.** Although cloud computing has been available for a decade, it remains a relatively nascent industry, where each vendor is defining its own value proposition, and there is still variation on the core services and contract terms incorporated into a service offering. As a result, it is difficult to create a precise service-versus-service comparison across cloud vendors. Further, for most customers, a subscription will require basic compute resources (either IaaS or PaaS) and will include value-added services such as database services, management or security services, or possibly consumption of PaaS extensions such as artificial intelligence, machine learning, IoT, anomaly detection, and identity and access management services. Those services are usually priced in addition to basic PaaS or IaaS services and likely have varying degrees of compatibility with existing customer solutions.
- **On-premises/cloud mix.** Most enterprise customers today will have a mix of on-premises and off-premises computing needs. Determining the right combination will help define costs and subscription needs for off-premises offerings. Further, organizations should recognize that these are likely to be targets rather than firm delineation points. The ability to adjust the mix of on-premises/off-premises computing provides considerable flexibility for an organization's agility needs. Finally, customers should look for a solution that offers them a relatively consistent experience with on-premises and off-premises computing resources, which helps reduce the friction associated with dialing up or down either compute resource as needed.
- **Software licensing.** For better or for worse, software licensing costs continue to be a key consideration when looking at cloud computing. In some cases, customers may find their existing licenses are tied to on-premises computing deployments and offer no cloud deployment rights. In other cases, the pricing model may be unattractive for deployment on a given cloud environment but may be more favorable on another cloud computing environment. Finally, having the ability or flexibility to deploy infrastructure software, databases, and applications on-premises or off-premises without any penalties is a desirable scenario for most customers and may result in lower opex costs over the longer term.

CONCLUSION

As customer adoption of cloud resources to meet their IT computing needs increases, cost is becoming an increasingly important factor when choosing cloud vendors. The availability of emerging cloud technologies such as containers and serverless computing also affects pricing and is becoming a differentiating aspect for vendors and for customers. For many organizations, the challenge today is no longer about whether to embrace cloud or not; rather, it is about how to get the most economic benefit from their shift to a cloud computing consumption model.

By offering a variety of pricing options, such as the reserved instances and hybrid models discussed in this white paper, Microsoft is well positioned to serve the needs of a majority of customers.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com