Tech, Media, and Telecom Practice

Winning in the SMB Cloud:

Charting a path to success

July 2011

Zoe Diamadi
Abhijit Dubey
Darren Pleasance
Ashish Vora

Contents

Introduction

In the past five years, it has been difficult to escape the hype around cloud computing. These technologies have been proclaimed by many as forces that will "reshape IT"[4] and "power innovation,"[5] while others describe them as a passing fad and "just another form of outsourcing."[6] As is often the case with potential disruptions in the high-tech industry, the real promise of cloud computing lies somewhere between these breathless pronouncements.

Virtually all of the largest technology players (e.g, IBM, HP, Microsoft, Cisco, Oracle, EMC, and others) are making significant cloud investments as are a number of smaller and newer entrants including Amazon.com, Google, VMWare, and Salesforce.com. Add to this large telecommunications companies and service providers (e.g., AT&T, Verizon, Comcast, TimeWarner, and others) and a long and diverse group of other participants (e.g., hosting providers, cloud solution providers, resellers, and other players) and it quickly becomes clear that the cloud landscape is not only increasingly crowded, but will certainly be dynamic and evolving as key players jockey for position in the coming years.

Small businesses (companies with fewer than 100 employees) and midsize players (companies with 100 to 1,000 employees) stand to gain much from the promise of cloud-computing technologies. Cloud computing offers SMBs access to reliable and scalable infrastructure resources (for example, computing and storage), configurable platforms that allow for integration between the business and vendors or customers, and rich application functionalities that can be paid for on an ongoing basis. Consequently, cloud computing offers SMBs the opportunity to enhance or improve IT capabilities in a way they previously could not.

Given the large investments being made by established technology companies, is there room for smaller and more varied providers of cloud services to SMBs? Or will cloud computing become a winner-takes-all marketplace with a handful of dominant companies? What should smaller cloud players do to compete and win?

In this white paper, we will examine opportunities for SMB cloud service providers and explain how they can seek to do more than merely survive in the marketplace—we will explore how they can be active, successful participants. In the first section, "Follow the growth," we describe the dynamics that make the SMB cloud services space so compelling. In the second section, "Size is not destiny," we use a combination of historical examples and economic theory to explain why the SMB cloud services market will remain fragmented. In the last section, "Playing to win," we share strategies and tactics that are being successfully employed by SMB cloud service providers in the marketplace today.

4 http://www.windowsitpro.com/article/news2/gartner-cloud-computing-is-reshaping-it.aspx
5 http://www.huffingtonpost.com/ping-li/cloud-computing-is-poweri_b_570422.html
6 http://www.computerworlduk.com/news/cloud-computing/3261721/cloud-computing-is-just-outsourcing-says-information-security-forum/

Follow the growth

As mentioned in the introduction, breathless pronouncements about cloud computing have been many and varied in the past several years. But what exactly is cloud computing? Why do consumers of technology—particularly SMBs—care about it? And what should providers looking to enter the cloud technology space understand about the market as they make decisions on how to compete and where to play?

This section explains the basics of cloud computing and details marketplace dynamics that demonstrate the potential for cloud-based services, especially in the SMB space. Participants entering this market must understand the growth of cloud-computing technologies and then pursue strategies targeting that growth.

Definitions: what is cloud computing?

Since the term "cloud computing" is used repeatedly, and different organizations often use conflicting definitions, it is important to develop a clear and consistent definition and an explanation of the associated concepts. McKinsey relies on the National Institute for Standards and Technology (NIST) definition of cloud computing and its associated concepts, described below.

- **Cloud computing**. This is a new computing paradigm in which dynamically scalable and multitenant resources are provided "as a service."

This definition includes two essential elements:

- **Dynamically scalable and multitenant**. One of the central features that differentiates cloud computing from other forms of computing is the idea that computing resources (for example, processing power or storage) can be scaled up and down elastically according to customer demand. For providers of cloud-computing services, this elastic provisioning capability can be provided economically by pooling the resources across a number of customers (that is, multitenancy) to help reduce variability in demand patterns.

- **Resources provided as a service**. Cloud computing abstracts the underlying technologies and operations from the user, providing access to various kinds of computing resources as services, rather than products. A key feature of this service-based approach to computing resources is the idea that consumers of the resources are charged for consumption on a variable, pay-as-you-go basis, similar to the way consumers of electricity or gas pay for their utilities. Service models for computing resources can deliver infrastructure, platforms, or software.

The three service models merit closer attention:

- **Infrastructure as a service (IaaS)**. In this service model, the key computing resources that are provided are computing/processing capabilities, storage, and network infrastructure. Users of IaaS offerings can deploy and run arbitrary software, including operating systems and applications. The user does not manage or control the underlying cloud infrastructure, but has control over systems that are run on top of the infrastructure (e.g., operating systems, applications). Examples of IaaS include Amazon's EC2 and S3 services, XDrive, and offerings by players such as Savvis and Rackspace.

- **Platform as a service (PaaS)**. In this service model, companies provide computing platforms that enable application design, development, testing, or deployment as services to end consumers. Users of PaaS offerings can use specific programming languages or frameworks and tools to deploy applications onto the Cloud infrastructure. PaaS offerings often act as integration platforms between companies and their vendors or customers. Examples of such offerings include Microsoft's Windows Azure and Salesforce.com's Force.com platform.

- **Software as a service (SaaS)**. In this model, software applications are delivered as services. The underlying responsibilities of software installation and maintenance rest with the SaaS vendor, not the end user. Numerous SaaS applications have emerged in the marketplace, most notably customer-relationship-management (CRM) applications from Salesforce.com and personal-productivity application suites such as Google Docs, Microsoft Office 365, and Zoho.

Any discussion of cloud computing would be incomplete without a description of the various deployment models that can be used to deliver cloud-based services. We refer to the two primary deployment models as "public cloud" and "private cloud." While there are some variants that blur the lines, such as "hybrid clouds", there are a few basic distinctions:

- **Public cloud**. In this deployment model, the cloud infrastructure is available to the general public and is owned by a large provider of cloud services. A company's information stored in the public cloud would by definition reside off premise and be collocated with other stores of information (for other companies or for individual consumers).

- **Private cloud**. By contrast, private cloud services consist of a cloud infrastructure that is operated solely for a single enterprise or organization. The infrastructure may be managed by the organization itself or by a third party, and the information can exist either on or off premise.

- **Hybrid cloud**. This deployment model combines public and private clouds together, allowing for some workloads to be processed on public cloud infrastructure, while others are run in private clouds. While the public and private clouds remain distinct from one another, technology fosters the ability for these clouds to interoperate with one another.

Market characteristics

MARKET SIZE

While estimates of the overall market size vary considerably, the consensus is that cloud computing is growing rapidly, measured both in actual dollars and in share of worldwide IT spending. McKinsey's estimates of cloud computing's size show the combined market for public and private cloud services growing from about $11 billion in 2010 to between $65 billion and $85 billion by 2015. To understand the SMB opportunity, the more relevant segment is the portion of the overall cloud-computing market related to public cloud, since the predominant portion of private cloud spend will come from large enterprises.

Exhibit 1 shows that the public cloud portion[4] of the overall market is expected to grow from $9 billion in 2010 to between $40 billion and $50 billion by 2015. In this time frame, the lion's share of growth will be driven by SMB consumption of public cloud services; SMBs will make up around 80 percent of public cloud spending in 2010 and 65 percent in 2015.[5] Public cloud spending alone will grow from about 2 percent of IT spending in 2010 to more than 6 percent by 2015; it is expected to reach 10 percent by 2020.

EXHIBIT 1

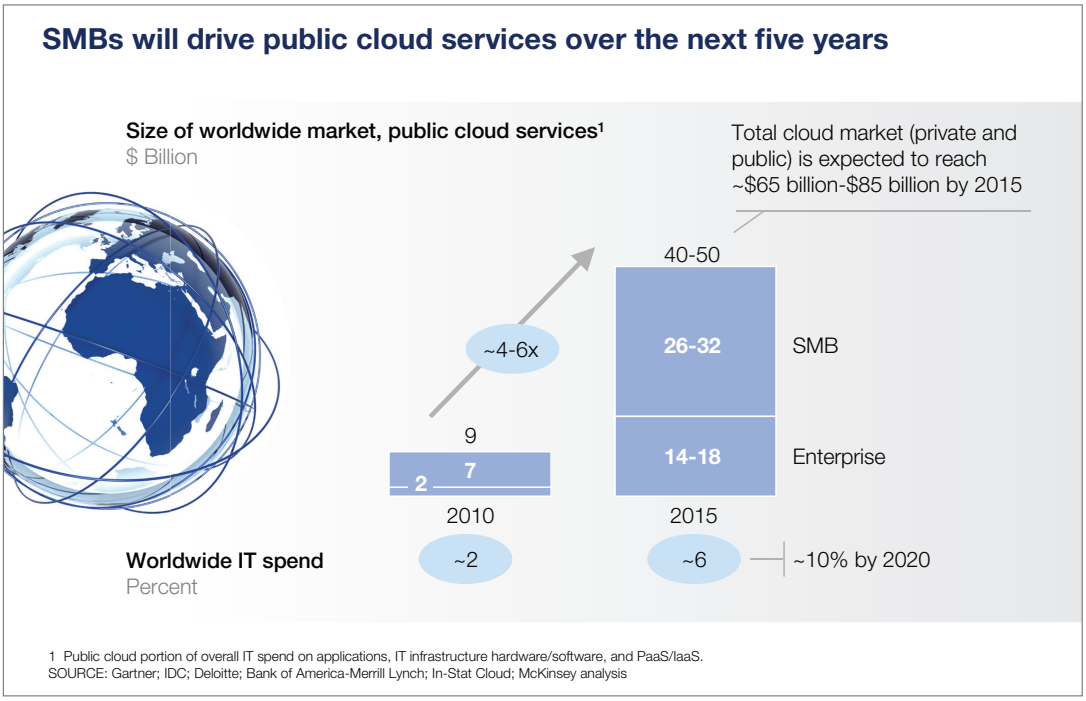

As illustrated above, the overall cloud market should grow rapidly. Our research shows that growth will be uneven across technologies; this presents opportunities for cloud players to seek out segments where they can most effectively compete. The data in Exhibit 2 demonstrate that public cloud migration will be heavily driven by customer-facing and employee-productivity applications such as CRM, e-commerce/Web, and collaboration applications. Other applications, such as data warehousing, enterprise resource planning (ERP), operations, and engineering management, will likely either migrate to private clouds or remain on premises in noncloud versions.

4 This includes the public cloud portion of overall IT spending on applications, IT infrastructure hardware and software, and PaaS/IaaS (Source: Gartner; IDC; Bank of America–Merrill Lynch; Deloitte; McKinsey analysis).
5 The 2010 public cloud spend split was 80 percent SMB and 20 percent enterprise; this assumes 2015 cloud IT spending reaches 12 percent for SMBs and 6 percent for enterprises (Source: In-Stat 2010; McKinsey analysis).

EXHIBIT 2

Growth won't be even–companies will migrate to the cloud in some areas more than others

CIOs surveyed on application migration
Percent expecting migration to public/hybrid cloud by application

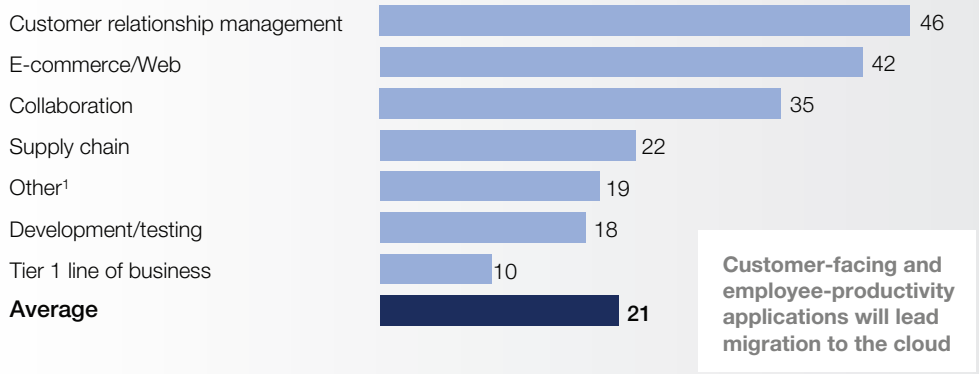

1 Includes weighted average for content management, other lines of business, data warehouse, operations, enterprise resource planning, engineering, and IT applications.

SOURCE: McKinsey CIO/CTO Cloud Online Survey; Gartner; McKinsey analysis

Exhibit 3 demonstrates that the growth potential for cloud computing exists in all geographies, though the vast majority of investment in cloud computing over the next several years will be driven by developed markets in Western Europe and North America.

EXHIBIT 3

Potential exists in all regions, but developed markets will drive most investment in the near term

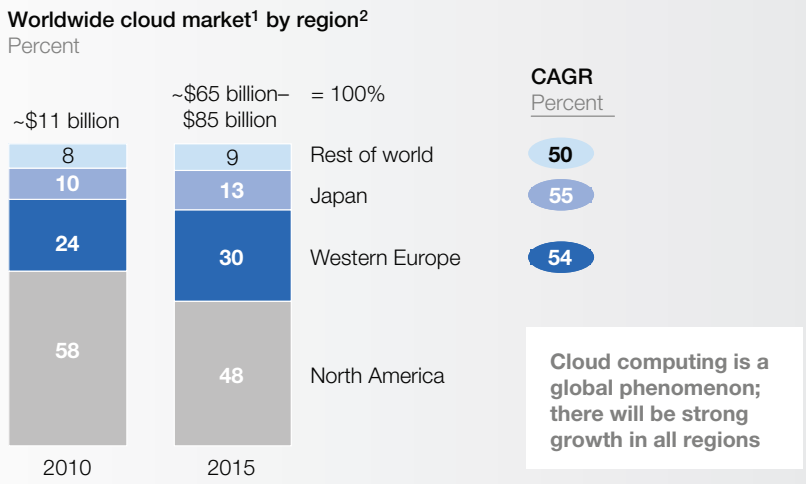

1 Cloud portion of overall IT spend on applications, IT infrastructure hardware/software, and PaaS/IaaS (includes both public cloud and private cloud markets).
2 This figure should be regarded as indicative; it is based on Gartner's top-down estimates given the early stage of adoption.
SOURCE: McKinsey CIO/CTO Cloud Online Survey; Gartner; McKinsey analysis

WHY THE CLOUD MATTERS TO SMBs

Estimates of market size clearly show that SMBs will drive public cloud adoption in the years to come—but why? What, exactly, does cloud computing offer to these businesses that is so compelling?

The benefits of cloud computing for SMBs can be understood from two perspectives. First, from a capability standpoint, cloud computing allows SMBs to gain access to technologies (infrastructure, platforms, and software) that they might otherwise need sophisticated IT support to obtain. For example, few small businesses have the knowledge and expertise to perform ongoing management of dedicated servers and storage in order to run commercially available ERP or CRM packages. By leveraging public cloud-based IaaS and SaaS solutions, these companies can stop worrying about the details of installing and running infrastructure or sophisticated software packages and instead simply contract for these offerings as services. Additionally, by making use of the cloud, SMBs can gain access to service levels (especially with respect to reliability and performance) that are much higher than they would be able to provide with a typical on-premise installation. In short, leveraging cloud solutions in a service-oriented manner allows SMBs to focus their efforts on the most important aspects of running their businesses, and it allows technology to be simply one of many actors that support the company.

Second, we can build on the argument related to capabilities and translate it into a financial argument. For many SMBs, the ability to outsource large portions (or all) of the company's IT needs to the cloud leads to significant cost savings.

EXHIBIT 4

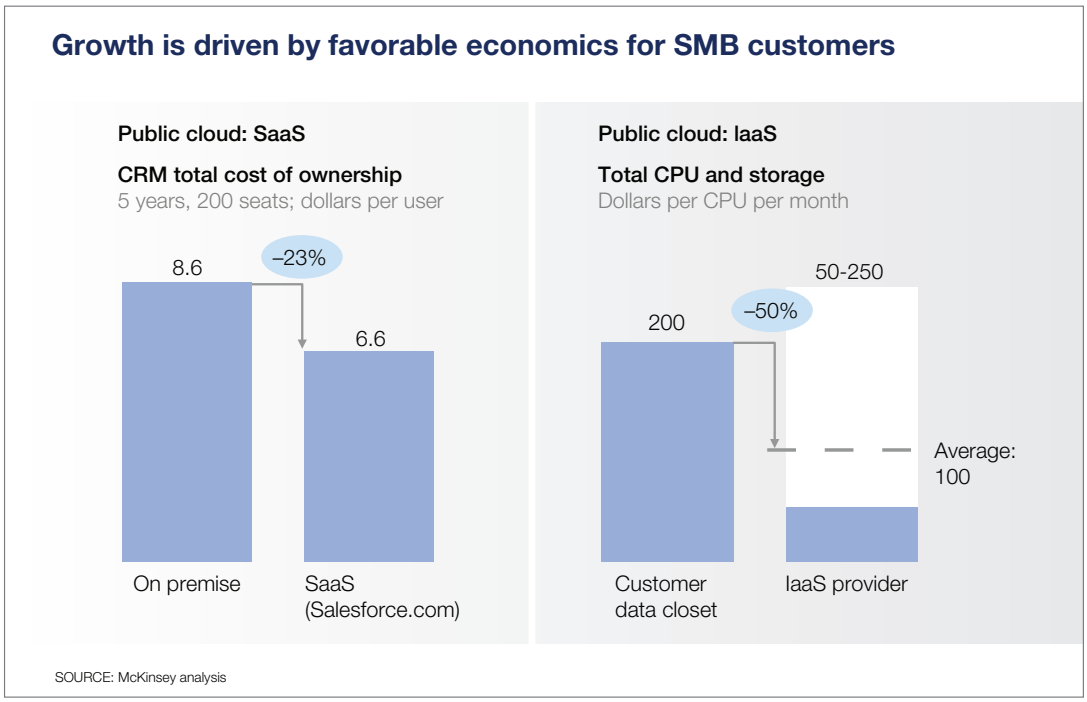

As shown in Exhibit 4, SMBs can save 20 to 25 percent by using SaaS-based CRM solutions in place of traditional on-premise applications. Cost savings to SMBs for moving infrastructure to the cloud can be well over 50 percent.

Size is not destiny

The analysis of the market opportunity in the previous section demonstrates cloud computing's growing importance, particularly for SMBs. But this growth raises its own set of questions: Who is driving technological innovation in cloud computing? What dynamics are shaping the industry? Will the promise of cloud computing only be realized by a handful of technology providers that make sustained and substantial investments, or will it remain highly fragmented, much like it is today?

In this section, we answer these questions by examining the competitive landscape in the cloud technology space and drawing on examples from other industries and economic theory to arrive at a view of the likely evolution of the Cloud industry.

Competitive dynamics

CLOUD TECHNOLOGY LANDSCAPE

The landscape of cloud players is diverse and has become more so as technologies have evolved. In Exhibit 5, we look specifically at some public cloud infrastructure players, and segment the different types of players by the amount of infrastructure investments they have made.

EXHIBIT 5

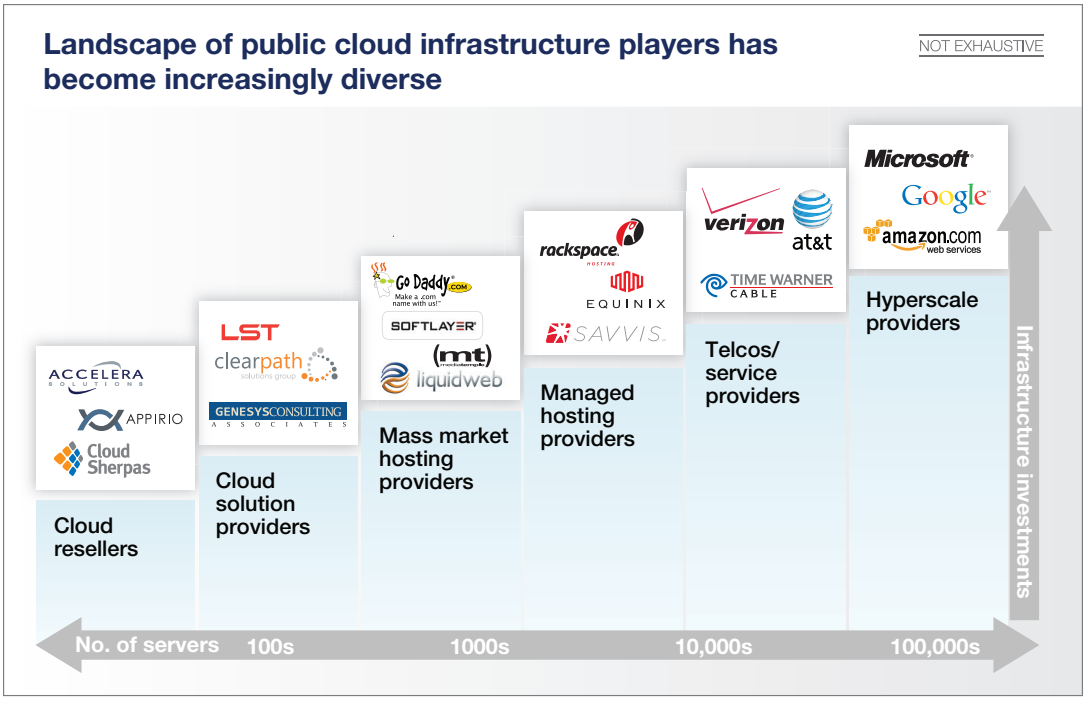

While the lines between these players are blurring, certain characteristics define each segment. The left side of the diagram shows that players that invest less in infrastructure, such as cloud resellers and cloud solution providers, have succeeded using a variety of business models. Some of these players act largely as consultants providing a variety of web services (e.g., web design, IT consulting) while others act as value-added resellers (VARs) for technology offerings from much larger players, such as Cisco, HP, IBM,

and Microsoft. These VARs have often built small-scale data centers and sell compelling microvertical applications that target specific SMB customers as well as providing implementation services to help migrate existing customer systems to the cloud.

Moving to the right on the diagram, the next major segment is mass-market hosting providers that offer services such as Web hosting and domain registration. In this segment, some players (particularly larger ones) build their own infrastructure capacity instead of renting it out. The key advantages of this approach include the ability to more actively manage infrastructure components, provide higher levels of service to end customers, and run operations in a more cost-effective manner.

The next segment—managed hosting providers—includes companies that have sophisticated infrastructure offerings that they lease out to others through collocation or dedicated-hosting arrangements. These providers have invested large sums in expanding their physical assets such as data centers, dedicated and virtual server offerings, and storage configurations.

The last two segments, large telecommunications providers and hyperscale providers, represent companies that have made significant investments in building out cloud infrastructure capabilities. Companies such as Verizon, with its early 2011 acquisition of Terremark, or Amazon, with its market-leading offerings of storage and computing power in the cloud (S3 and EC2), are pushing the envelope on scale in the cloud-computing marketplace through the aggressive acquisition of capacity and greenfield buildouts. Exhibit 6 provides a glimpse of the scale of the investments that a few hyperscale providers have made in the past decade in data centers, R&D, and acquisitions related to the cloud-computing space.

EXHIBIT 6

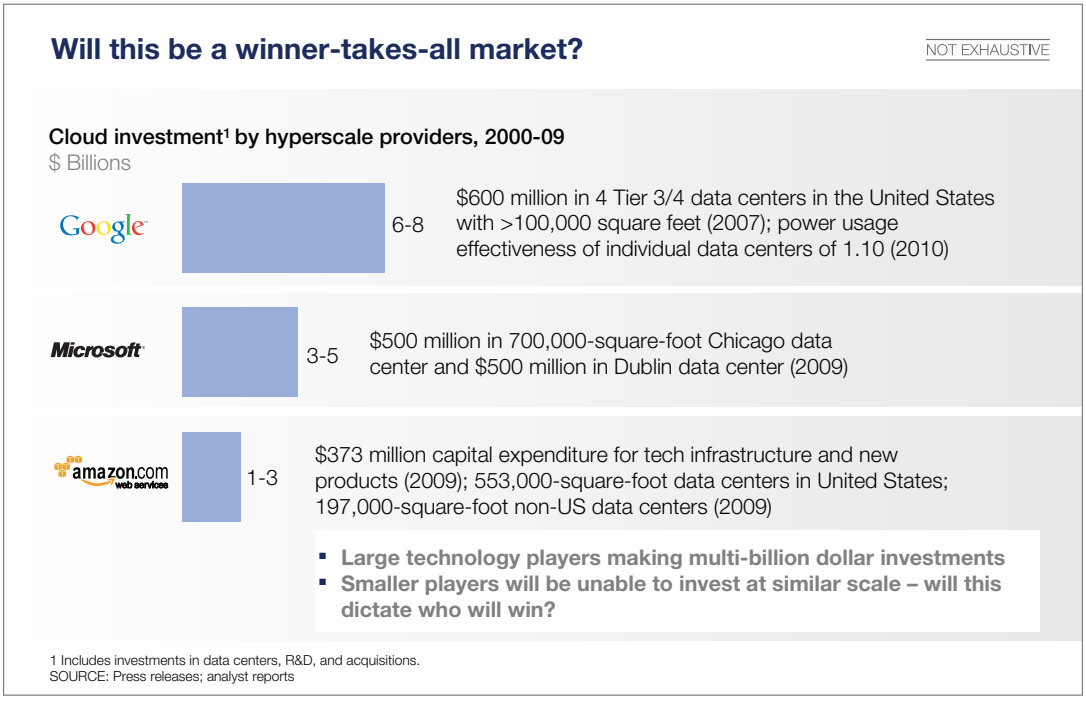

ECONOMICS OF CONSOLIDATION AND FRAGMENTATION

Given the huge bets that these companies are placing, it is natural to wonder whether this market will turn into the exclusive domain of just a few large players. After all, the last time the technology industry experienced this kind of booming investment, the winners ended up taking it all: Google has dominated Internet search, and Facebook has come to define social media.

Is this the future that awaits the cloud?

To answer this question, it is helpful to look at an analogous example from the retail industry. The retail landscape parallels the diversity found in the cloud landscape, with a wide range of players of varying sizes and offerings as well as companies that play across multiple categories versus those that are single-format. And as in the cloud market (particularly for infrastructure and platform players), there are some definite advantages to scale in the retail industry. Purchasing power, sophisticated distribution and logistics, and reduced levels of variability are all hallmarks of scale-players in retail. And finally, the retail industry has its own examples of players making large strategic bets to try to change competitive dynamics –Wal-Mart's IT–driven supply chain innovations are a perfect example.

Interestingly, in spite of the benefits of scale and the actions of large cross-category players, consolidation has not taken hold in the retail industry to the extent one would expect, as shown in Exhibit 7. While consolidation certainly took place—the market share of the top five cross-category retailers in the United States doubled between 1991 and 2008—robust competition remains, with room for many types of special purpose retailers. The rest of the industry—which accounts for over 80 percent of industry revenues—remains heavily fragmented, with almost one million retail establishments.

EXHIBIT 7

Market development of industries with similar characteristics (such as retail) suggests a fragmented future for Cloud services

While some consolidation has occurred in the retail industry...

Market share of retailers[1]
Percent of industry sales

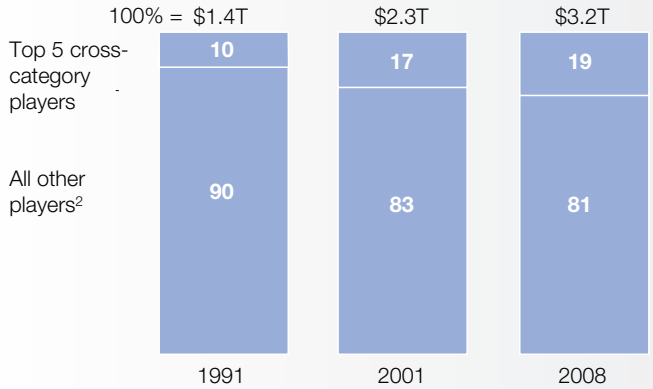

...fragmentation persists in the overall marketplace

- ~1 million retail establishments in 2008 representing >80% of retail industry revenues
- Fragmentation primarily driven by diversity of customer needs and desires

1 Retailer definitions based on US Census retailer categories (excludes motor vehicle/parts dealers and food services/drinking places
2 All other players include both cross-category players as well as dingle-format retailers
SOURCE: US Census, Stores magazine, McKinsey analysis

Analysis of other industries such as banking reveal a similar pattern of persistent fragmentation, in spite of a handful of large and powerful players. Why has fragmentation remained in the retail and banking industries, and what can it tell us about what could happen in the cloud-computing market? The answer rests in understanding the economics behind the reasons companies make the decisions they make—in particular, why they choose to either build out their own capabilities (perhaps by consolidating other firms) or buy these capabilities from the marketplace on an ongoing, contractual basis.

The basic economic principle at play is the notion of transaction costs, an idea often attributed to economist Ronald Coase. Transaction costs are any costs that are incurred as a result of participating in a marketplace—that is, any costs that exceed the cost of the product itself. For example, when an individual goes to the ice-cream store to purchase a scoop, the transaction costs include the costs of getting in a car, driving to the shop, parking, standing in line, and so on. Coase used the theory of transaction costs to explain why corporations arise in the first place—without transaction costs, it would generally make sense for all activities to take place freely through market-driven mechanisms. However, because transaction costs exist, companies sometimes find it cheaper to do things themselves rather than bear the transaction costs of interacting in a market. As Exhibit 8 shows, companies make decisions by examining the internal and external transaction costs of conducting an activity. When internal transaction costs are less than external transaction costs, companies will choose to perform an activity themselves (that is, build). At some point, however, the costs of carrying out an activity internally will be far greater than the cost of the activity provided by an outside entity, and consequently the company will contract out for services (that is, buy).

EXHIBIT 8

Principles of transaction economics

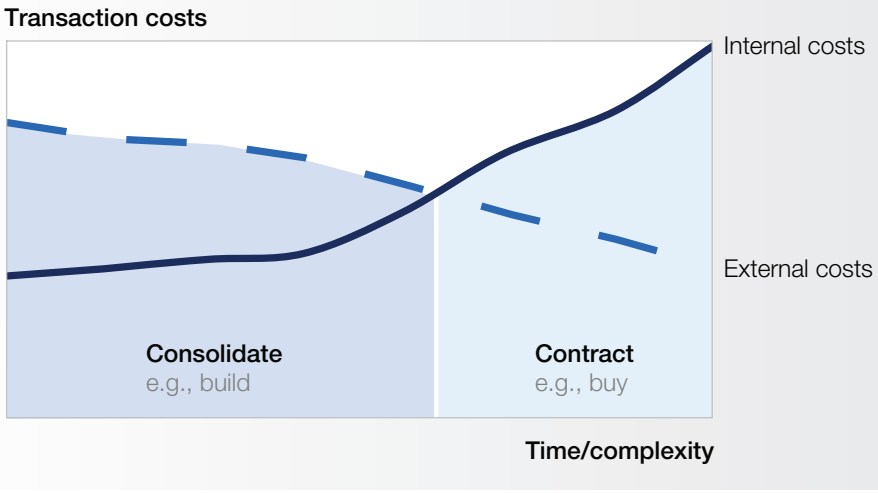

SOURCE: Institute for Research in Economics and Business Administration (SNF); McKinsey analysis

How does this relate to the questions of consolidation and fragmentation? There are three primary reasons fragmentation persists in a marketplace: disintegration, industry standardization, and customization.

- **Disintegration**. This is a simple extension of the transaction-cost theory. When a firm that has consolidated numerous players begins crossing into the area where internal transaction costs exceed external transaction costs, the firm will experience pressure to disintegrate some of their activities through spinoffs or other vehicles. For example, in the late 1990s, several US automakers went through waves of disintegration when they spun off their parts suppliers. The parts divisions at these automakers were facing increasingly non-competitive cost structures, and as a result, imposing higher internal transaction costs on the automakers than if the parts were obtained from other parts suppliers. In an effort to simplify the cost structure and consequently reduce the transaction costs of obtaining car parts, General Motors spun off its parts division into Delphi, and Ford separated out its Visteon division. Both Delphi and Visteon remain independent entities today.

- **Industry standardization**. In industries with heavy knowledge-transfer requirements (for example, those where proprietary knowledge can be a source of competitive differentiation), standardization can drive fragmentation by decreasing external transaction costs and lowering barriers to entry for new market participants. Illustrations of this type of fragmentary pressure abound in the technology space—for example, the standardization of Internet telephony protocols, such as voice over Internet protocol (VOIP), has led to significant competition in the telecommunications industry. New companies, such as Vonage and Skype, and established companies from related industries, such as Cisco, have all been able to establish footholds in the formerly concentrated telecommunications market.

- **Demand-side diversity**. As companies grow larger and more complex, it becomes more difficult (and more expensive) to provide customized offerings to different segments of customers. In other words, the internal transaction costs of trying to serve multiple market niches simultaneously may be much higher than the external transaction costs for other players to serve these niches. We can see this clearly from our retail analogy. While many think of Wal-Mart as retail's dominant force, the company maintains a share of only about 13 percent of the heavily fragmented space. Why? Because Wal-Mart cannot be all things to all people—it provides an offering to customers that seek strong value; customers that are looking for high-fashion or luxury merchandise will naturally turn to retailers better suited to provide those goods.

  In order to address demand-side diversity, companies typically differentiate their products or services across dimensions that appeal to their customers. Examples of these dimensions include the level of service or support, the trust facilitated by local relationships, the level of customization, the cost effectiveness, or the reliability of the product or service.

All three of these factors will play a part in keeping the cloud-computing landscape fragmented. However, issues related to the diversity of demand are particularly relevant for SMB cloud services.

WHY FRAGMENTATION WILL PERSIST IN THE SMB CLOUD SERVICES SPACE

In 2010, McKinsey asked SMBs in the United States to describe the way their companies manage IT. The results, shown in Exhibit 9, provide insight into the needs of SMBs and demonstrate how these needs vary in the segment. As the chart shows, many SMBs surveyed, particularly the smaller ones, manage IT in a self-service fashion, using tools and information from the Internet and offline publications to guide their IT-related decisions. For many of these self-service SMBs, offerings from hyperscale providers such as Google, Amazon, and Microsoft will be the most attractive way to begin utilizing cloud services.

However, a number of SMBs either leverage third-party partners to actively guide their IT decisions, or they've

created internal IT functions which in most cases will also rely on the support provided by local third-party advisers.

SMBs described the reasons they prefer smaller and more localized IT providers across a number of dimensions, which speak to the diversity of demand that exists in the SMB space:

EXHIBIT 9

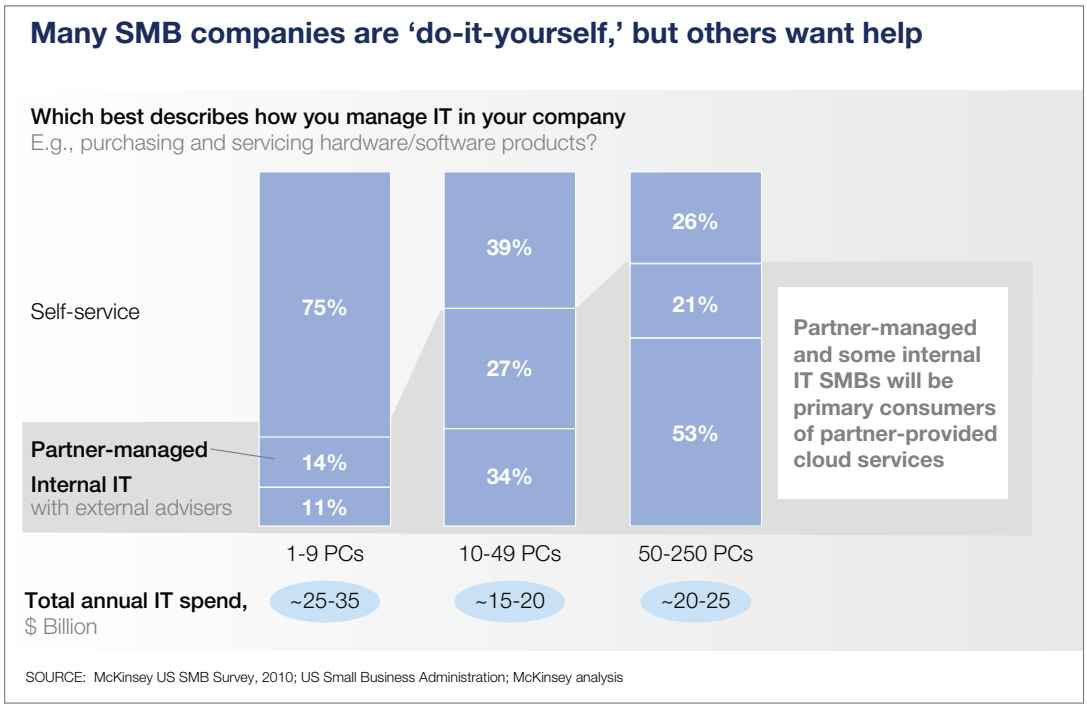

- **Service levels and support**. "My [IT] consultants are a family business – I can always get a hold of them. With larger companies, it takes longer to get someone to solve a problem." – Small business

- **Local relationships**. "I like to shake someone's hand and be able to look them in the eye." – Small business; "Vendors usually have eight seconds to convince me that I should pay attention to them, and I listen a lot more to those I have a good relationship with, than those I don't know." – Medium business

- **Flexibility/customization**. "[Big companies] are arrogant. They tell you what benefits them and not what would benefit you. We want to work with companies that make products that meet our needs particularly around compliance and the nature of our industry" – Medium business

- **Effectiveness**. "Sure cost matters, but it's not the most important thing. Quality matters a lot too. If I buy something that fails, or that I can't get fixed quickly, then my whole business is at risk, and that's a lot more important to me than saving a dollar here and there." – Small business

- **Reliability**. "Frankly, if I'm thinking about IT, then it probably means something has gone wrong and I'm not happy. IT is necessary, but should be behind the scenes, reliable, and painless. I want to work with a provider who makes IT as carefree as possible." – Medium business

SMB customers represent a compelling opportunity for many smaller cloud service providers; these organizations can use their strengths in agility and responsiveness to cater to SMBs in an area that larger providers may struggle with.

CAN BIGGER PLAYERS OPERATE MORE EFFICIENTLY?

Given the dynamics of managing demand-side diversity and the importance of localized support, it is reasonable to think that fragmentation of providers to SMBs will remain high for the foreseeable future. But can smaller providers offer their services in a cost-effective manner? Won't hyperscale providers be able to achieve levels of efficiency that are beyond the reach of smaller players?

The drivers of operational efficiency

If we examine the ways that large and small providers bring cloud services to the market, we can see several areas where operational efficiency is important. One area of particular importance is data center operations, the costs of which are borne by both large and small providers alike, and which contribute in a significant way to the breadth and quality of a cloud offering. The major drivers of operational efficiency in data center operations are:

- **Efficient procurement of large areas of land.** Cloud service providers need to obtain land on which to situate data center facilities, while taking into account geographic advantages and risks, interconnectedness with existing facilities and networks, and considering how the needs of the facility may evolve over time (e.g., space for future build-out).

- **Low-cost fabrication of facilities on this land.** Facility construction costs for data centers encompass a wide range of costs, from the physical shell of the building itself, to the cooling, power generation and distribution, and basic rack infrastructure (racks, cabling, raised floor) equipment that is required in every data center.

- **Ability to cheaply acquire volumes of low-priced power.** Data centers consume significant quantities of power, and as such an important focus of operational efficiency improvements is on achieving cost-advantages in procuring and consuming power. This can happen in a number of ways, from taking advantage of geographical advantages of a particular site such as access to natural cooling sources (water, wind) to negotiating volume discounts with utility providers.

- **Improvements in overall energy efficiency.** On top of simply finding cheaper power, providers can also pursue innovations that improve the overall energy efficiency of their operations, leading to better power usage effectiveness (PUE) ratings of the facility.

- **Negotiation of discounts on computing equipment.** Providers can find several ways to reduce the costs of computing equipment, from negotiated discounts based on volume purchase of standardized platforms or development of custom, "white-box" solutions optimized for cost-savings from both a construction and operations perspective.

- **Increasing productivity of labor resources.** Labor costs for operating data centers can be managed through the use of automation technologies that improve the ratio of equipment managed to employees. These automation technologies are often made feasible as a result of the investments in standardization on computing equipment described above.

Does scale matter in terms of these operational efficiencies?

While the operational efficiency drivers outlined above are discrete levers, the most effective providers are able to simultaneously drive efficiency across several (or all) of them. Doing so requires providers to put concerted focus on developing innovative new ways of running and managing their data centers.

Interestingly, based on conversations with both large and small providers, we do not believe that the focus on simultaneous innovation is causally driven by the scale of the provider. However, the clearest demonstrations of focused innovation around operational efficiency come from hyperscale players such as Google, Amazon and Microsoft, possibly because their consistent focus from the outset has enabled them to grow to their current size more easily than their less-focused peers.

In spite of the advances that have been made in operating data centers more efficiently, these innovations and best practices are not the exclusive purview of hyperscale providers; indeed, from our research, there are many instances of smaller providers managing costs in a creative manner to help operate with a lower cost profile:

- **Capturing vendor discounts on OEM hardware and software.** Small players have become quite effective at negotiating discounts of 10 to 30 percent (or more) that approach the deals that OEMs are making with hyperscale providers, particularly by driving purchase of more standardized systems. In fact, there is growing anecdotal evidence that in their rush to gain market share, OEMs have been selling to large hyperscale providers at, or even below, cost. If this is true, the increasingly small price advantages on OEM equipment that hyperscale providers have negotiated are bound to disappear soon.

- **Taking creative approaches to facilities and energy.** Many smaller cloud providers have begun performing sophisticated analysis around build-versus-buy decisions, and in some cases they are discovering that it is cheaper to build and own data-center facilities rather than to contract out for them. Innovations in modular data center design are allowing for more cost-effective construction of facilities. Better availability of next-generation cooling technologies and increased awareness of the importance of managing energy costs has allowed a number of these players to achieve PUE ratings of 1.2 or lower, rivalling their bigger peers' energy efficiency. Finally, small players are increasing the sophistication of their negotiations with utility providers and achieving lower costs for procuring energy.

- **Maximizing the useful life of equipment.** Some SMB cloud providers explained novel strategies for stretching the useful life of their infrastructure investments by creatively redeploying machines (for example, repurposing old PCs as low-end servers).

- **Making operations lean and nimble.** Finally, a number of cloud providers explained how the increased levels of standardization in their operations allowed for heavy use of automation and other tools to manage their operations. As a result, they were operating much more efficiently than they had in the past and driving administrative support ratios to levels well beyond the typical 80 to 100 units per FTE seen in many small operations. Indeed, several players were easily exceeding 500 to 1,000 units supported per FTE – one small provider indicated a support ratio approaching 3,000 units per FTE – indicating levels of automation that had traditionally only been achieved by hyperscale providers.

While these creative strategies do not fully eliminate the economic advantages that hyperscale providers enjoy, they do help narrow the gap significantly. In the end, for customers of these services, what matters is total cost of ownership (TCO). As shown in Exhibit 10, in a customer's typical TCO for a CRM application, infrastructure costs (ostensibly where hyperscale providers see the biggest economies of scale) make up a little more than 10 percent of overall TCO. In other words, hyperscale cost advantages of 40-50 percent only translate to a 4-5 percent cost advantage from an end-customer perspective. Many other costs, such as data migration, customization, and training, are the types of services most often provided by smaller, more localized providers.

EXHIBIT 10

Infrastructure costs represent a small fraction of overall information technology spend and application total cost of ownership

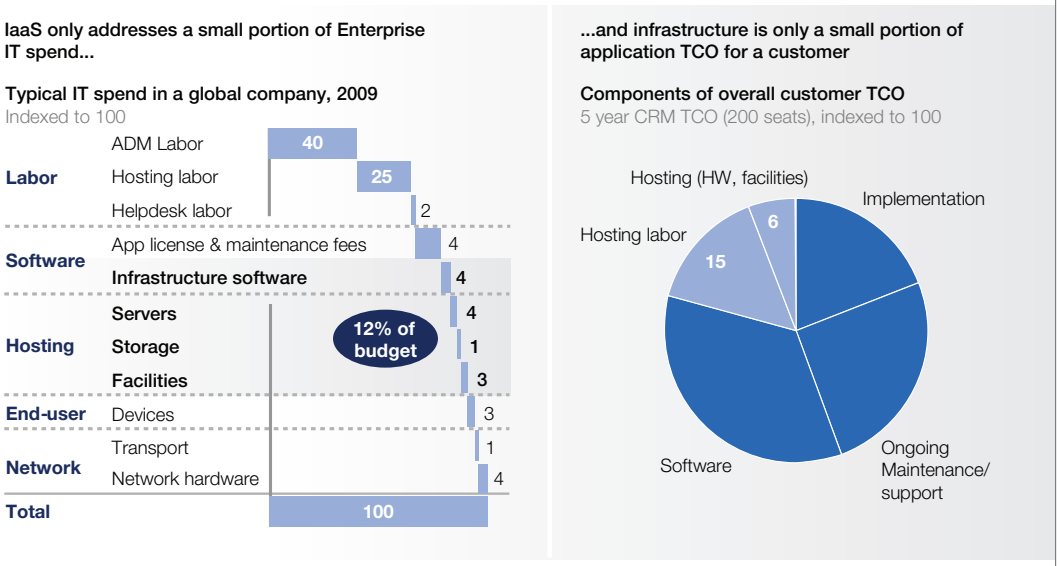

SOURCE: Gartner, McKinsey IT metrics

WHAT THE FUTURE WILL LOOK LIKE

Based on both historical analogs and economic principles, we believe that the cloud services market, particularly for SMBs, will remain fragmented over the next 5 to 10 years. While some consolidation will certainly take place, in the end, the market will consist of both asset-intensive players (such as the hyperscale providers) and asset-light players (such as VARs and solution providers) and variants in between. Success in this market will be defined by the availability of high-quality self-service interfaces that allow smaller companies to manage IT more effectively, as well as localized relationships with IT solution providers with related expertise.

Playing to win

If the cloud market increases rapidly in size yet remains significantly fragmented, what does this mean for participants in the diverse ecosystem of cloud computing? What will it take to succeed in this environment?

In our discussions with a variety of cloud solution providers and mass-market hosting providers, we were struck by the fact that players were making healthy profits with many different business models. Exhibit 11 demonstrates that operating margins of 15 to 20 percent or more can be achieved, and many of the players we spoke with managed various aspects of their businesses to sustain this level of profitability.

EXHIBIT 11

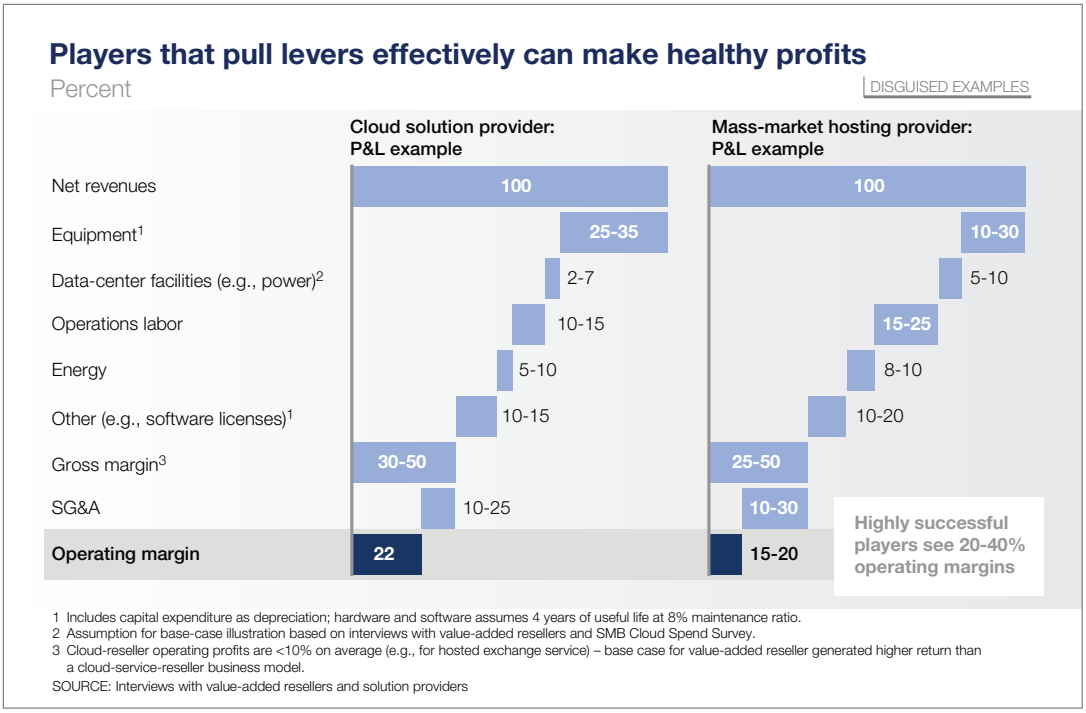

In this final section, we lay out four key principles that will help companies chart a path to success in the cloud services market. Although these lessons and best practices are drawn from our discussions with smaller cloud providers, the principles apply to both smaller players and larger players, albeit in different ways.

Key economic levers that will drive success

Our discussions with cloud solutions providers and mass-market hosting providers identified four key economic levers these players used to manage their businesses:

1. Increasing average revenues per user (ARPU) through cross-selling and up-selling
2. Driving loyalty and stickiness
3. Pursuing operational excellence
4. Minimizing customer acquisition costs

Let's look at each of these in more detail.

INCREASING ARPU THROUGH CROSS-SELLING AND UP-SELLING

A critical metric in tracking the health of cloud services organizations is ARPU. The best companies understood the importance of increasing ARPU, since the ability to drive higher revenues from existing customers often allowed them to spend little to pursue new customers. The key to increasing ARPU effectively is to understand customer needs and tailor offerings that speak to what customers will want. Players that do so can reap huge rewards, as shown in Exhibit 13. This chart is based on survey responses from about 500 SMBs; it shows the average monthly spending for various services and the percentage of the sample SMBs that adopted the services. Interestingly, less than 3 percent of SMBs adopted four or more of these services. For providers that are able to capture customers using all of these services, the payoff in ARPU could be dramatic. In other words, driving customers to more and varied services has a significant impact on the overall revenue opportunity.

EXHIBIT 12

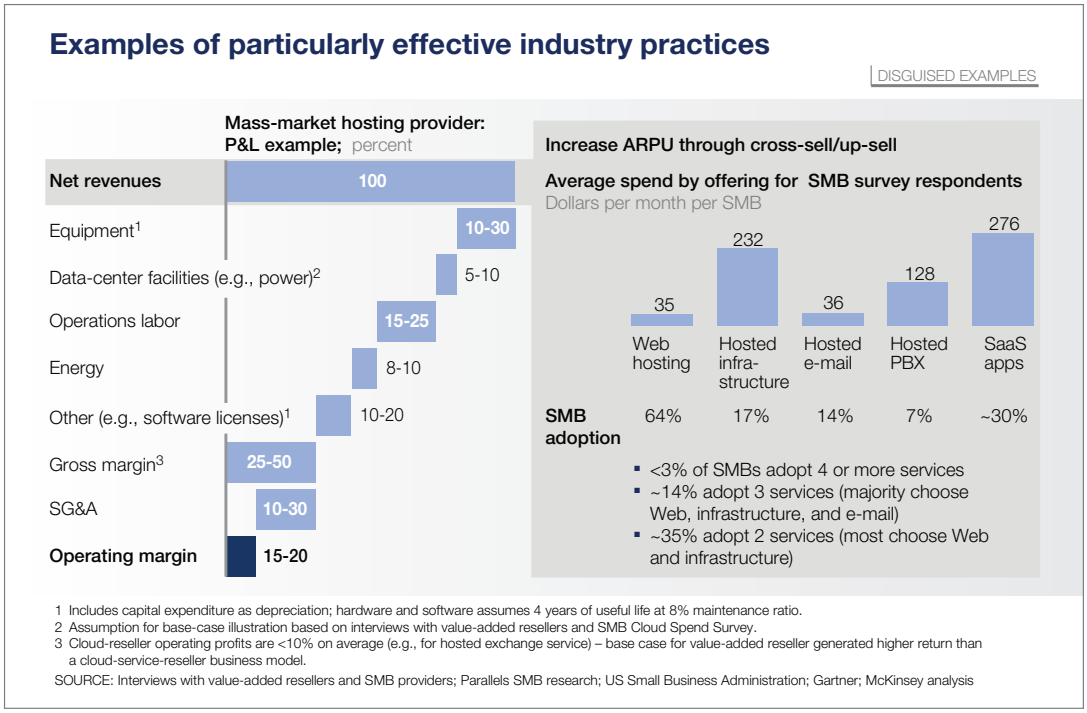

This notion is best exemplified in a quote from a mass-market hosting provider: "We're moving up the food chain and trying to eliminate people who only do part of the stack." The implications are clear—cloud service providers must think about adding more offerings. If multiple services are provided but customers are not buying them, the company must figure out if the services should be improved or if sales efforts need to be revamped. Service providers' competitors are already assessing this—and these players should consider whether they can afford not to.

DRIVING LOYALTY AND STICKINESS

The second major economic lever that successful providers pull is driving loyalty and retaining customers by actively managing churn. In our research, we learned that there is a dramatic difference in churn rates for different providers. The best-performing providers have churn rates that are 60 percent lower than those of the worst-performing providers. For a typical provider, with economics similar to those of players we spoke with, reducing churn by 1 percent can translate into a 3 to 5 percent increase in profitability.

Clearly, keeping customers is important—but how can they be retained? The best performers we spoke with did a number of things well when it came to managing churn. First, these companies adopted a disciplined approach to customer life-cycle management (CLM) that focused on understanding customer lifetime value and identifying behavioral cues (for example, purchase decisions and support calls) that might help predict churn.

Second, many companies adopted a proactive approach to preventing churn. As one mass-market hosting provider remarked, "If [our support desk does] proactive outreach to our hosted exchange customers, we improve retention by 80 percent."

Third, successful players drove customers toward stickier services. For example, a number of companies explained how providing backup services increased stickiness (as well as ARPU). A cloud solutions provider described how customization could improve stickiness, noting that of his competitors, "the smart people were focused on personalization."

Fourth, strong churn-management programs involved designing "save desks" that act as a line of last resort when a customer is about to leave. Individuals on these desks were given greater flexibility to try to either solve customer concerns or offer preferential pricing as a way of enticing customers to stay.

Finally, a small number of players described efforts they were taking to identify "serial churners" before acquiring them as customers; in some cases, these companies chose not to do business with them.

PURSUING OPERATIONAL EXCELLENCE

As described in the section on operational efficiencies, one of the key success factors for a number of SMB cloud providers has been to take novel approaches to cost management. For example, regarding the cost of equipment, many companies found creative ways to reduce their hardware-acquisition costs. One mass-market hosting provider described how they "owned [their] own network" to efficiently connect multiple data-center locations. A cloud solutions provider talked about redeploying old PCs for low-end server usage and taking advantage of OEM end-of-quarter discounts to source hardware cheaply. Other companies described their abilities to get significant volume discounts on OEM equipment purchases designed for cloud services as a result of their large dedicated-hosting businesses. Another hosting provider we spoke with had developed sourcing strategies with white-box manufacturers in China to source servers for €150 each.

Companies were even more creative when it came to managing facility and energy costs. For smaller providers, it is not uncommon to leverage leased facility space from some of the managed hosting providers—but we also found many examples of small solution and service providers that were cost-effectively running

their own facilities. One solution provider in North America had built two 6,000-square-foot data centers that were collocated with the company's office space; in the winter, the waste heat from the data center was used to warm the office space. We spoke with a provider that relocated its data center from one European country to another to take advantage of significantly lower electricity costs and the environmental benefits of being near a riverbank as well as the opportunity to occupy an old tobacco factory at attractive rates and repurpose it as a data center. The access to naturally cool water allowed the data-center facility to be run with a PUE rating significantly lower than the level Google considers "best practice."[6] A cloud solutions provider in North America described negotiating heavily with the utility provider in a state where it was the only large data-center facility; it realized 50 percent savings on energy costs as a result.

Companies also described extensively using automation technologies to reduce the amount of labor necessary to run and manage their facilities. As one mass-market hosting provider put it, "Automation, standardization, and cross-training our employees is critical for us." Another spoke of the competitive advantages of automation, saying "Other organizations that depend more on people than we do just can't compete with our ability to layer on new services." A third described the extent to which automation solutions influenced product decisions, stating "We will not offer anything unless it's automated."

MINIMIZING CUSTOMER ACQUISITION COSTS

Finally, successful companies we spoke with minimized the costs of acquiring new customers. Once again, they employed many approaches to do so. The simplest approach, and one that many companies followed, was to start with existing (non-cloud) customers and offering "risk-free" 30-60 day trials to migrate them to cloud offerings. Typically, these companies would start with relatively basic offerings such as web hosting, and then migrate up through hosted email, remote data storage and archiving, and ultimately to various line of business applications. Customer acquisition costs for this path were extremely low.

For companies that were not in a position to leverage their existing user base, establishing partnerships with other players in the cloud landscape was often an effective approach. For example, one solutions provider described working with independent software vendors (ISVs) that provided software for microvertical markets such as specific areas in the health-care sector. The solutions provider would bundle the ISV's software solution into a cloud-based offering with additional services and then leverage the ISV's established connections to the microvertical.

Still other providers relied on mergers and acquisitions to "buy" customers. These companies spent the majority of their customer acquisition budgets on purchasing other companies that fit the customer demographics they were targeting.

6 http://www.google.com/corporate/datacenter/efficiency-measurements.html

Applying the economic levers to larger players

While the prior sections described some best practices observed among SMB cloud service providers, the key economic levers are just as applicable to larger players in the cloud landscape. Hyperscale providers and telecom companies must think carefully about increasing ARPU, managing retention, operating efficiently, and minimizing customer acquisition costs.

In managing ARPU, the primary driver of success for larger players is bringing higher-quality self-service interfaces to market. A large segment of SMB customers has the appetite to manage IT independently, but this group needs compelling offerings that make this process simple and manageable; such players will otherwise seek external support from smaller, more local providers.

To better manage churn, large providers should adopt many of the same practices used by successful smaller players—essentially building out world-class CRM capabilities.

A number of the larger providers, particularly the managed hosting providers, should aggressively invest in lower-cost, next-generation infrastructure to remain competitive on costs. As one Asian hoster told us, “I can outcompete the telcos any day because I’m fully automated; I can rapidly provision a computing environment for a customer and then deliver a variety of cloud services with a single bill; and I can do this far faster and for lower cost than the local telco. They’re running on technology they built many years ago, and it’s hard for them to get stuff done.”

Finally, since many SMBs will remain difficult to reach without localized channel outreach, it behooves larger players to provide clear and compelling value propositions to smaller players in the channel such as VARs. Many of the VARs and local cloud resellers we spoke with questioned the economics of reselling cloud services from hyperscale providers, relative to other options they had, and asserted that the story had to be “a lot clearer and more compelling” than the hyperscale providers were currently providing. Without such a value proposition, many large companies will fail to make inroads in large segments of the market that demand higher levels of service and localized support.

Conclusion

In summary, cloud computing is a large and growing market that is particularly attractive for SMB customers and that offers significant potential for players of all sizes. The cloud landscape is rich and robust; because of this richness, size will not be destiny, and the market will remain fragmented with opportunity for many. To stay competitive, companies must focus on managing the top and bottom line, using traditional and creative means to do both.

About the Authors

Zoe Diamadi is an Associate Principal and Abhijit Dubey is a Principal in McKinsey's San Francisco office. Darren Pleasance is a Principal in the Silicon Valley office. Ashish Vora is an Engagement Manager in the Chicago office.

Acknowledgements

The authors wish to thank Birger Steen, Terra Terwilliger, Jack Zubarev and their McKinsey colleagues William Forrest, Yuval Harry, Loralei Osborn, Brent Smolinski, and Kara Sprague for their substantial contributions to this paper.

Tech, Media, and Telecom Practice
July 2011
Designed by US Design Center

www.mckinsey.com/en/Client_Service/High_Tech.aspx