



Cloud Management & Consumption

Ending the public vs. private debate?

SEPTEMBER 2017

COMMISSIONED BY





About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

© 2017 451 Research, LLC and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with 451 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. 451 Research disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although 451 Research may discuss legal issues related to the information technology business, 451 Research does not provide legal advice or services and their research should not be construed or used as such.

451 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

NEW YORK

1411 Broadway
New York NY 10018
+1 212 505 3030

SAN FRANCISCO

140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON

Paxton House
(Ground floor)
30, Artillery Lane
London, E1 7LS, UK
P +44 (0) 207 426 1050

BOSTON

75-101 Federal Street
5th Floor
Boston, MA 02110
Phone: +1 617.598.7200
Fax: +1 617.357.7495

EXECUTIVE SUMMARY

Cloud service providers and their customers are becoming increasingly savvy when it comes to making choices about infrastructure. While IT spending continues to move to off-premises models, the overnight shift to public cloud has not materialized. End users are realizing the value of hosted private cloud, hybrid and other models – identifying the potential gaps in deployment requirements. As interest shifts from service construction to service consumption and delivery management, end-user demand points to the growth of spending on bundled offerings and managed services – the so-called ‘everything as a service’ (XaaS). Channel partners and service providers will benefit from cloud management tools and platforms that provide customization and automation for their clients while offering total cost of ownership and other benefits. This, in turn, helps address gaps and potential opportunities such as security, compliance, data protection/sovereignty, performance, control, etc.

KEY FINDINGS

- As cloud delivery models mature, end-user interest is shifting toward the consumption of cloud services, rather than infrastructure construction.
- Best execution venue (BEV) is increasingly relevant to workload placement and service delivery in order for users to control where they want their workloads and how their clouds are delivered.
- Total cost of ownership (TCO) has become a critical component of opex measurement for cloud. The scale of delivery and utilization rates affect the choice of public, private, proprietary and open source platforms.
- As infrastructure becomes increasingly ‘invisible’ to end users, operational managed services will be further bundled with standard offerings as interest in management tools progresses.
- The service provider and IT distribution firms that make up the cloud ecosystem stand to gain most from the management of cloud service consumption. Tools, automation and service migration are key differentiators.

Platforms and Buyers

In any discussion about cloud, it is important to provide a baseline definition, given the many different opinions about how to describe it. 451 Research defines the three main cloud delivery methods used by service providers as follows:

- **Hosted private cloud.** Infrastructure deployed with a hosting provider but not shared with other customers. Hosted private cloud is configured for resource pooling, automation and orchestration and may also include self-service, catalogs, metering and chargeback services.
- **Public cloud.** Multi-tenant infrastructure shared with other customers. Public cloud is configured for resource pooling, automation and orchestration, but may also include self-service, catalogs, metering and chargeback services.
- **Hybrid cloud.** A delivery model that provides coordination and interoperability between two or more distinct cloud entities to deliver seamless business functions.

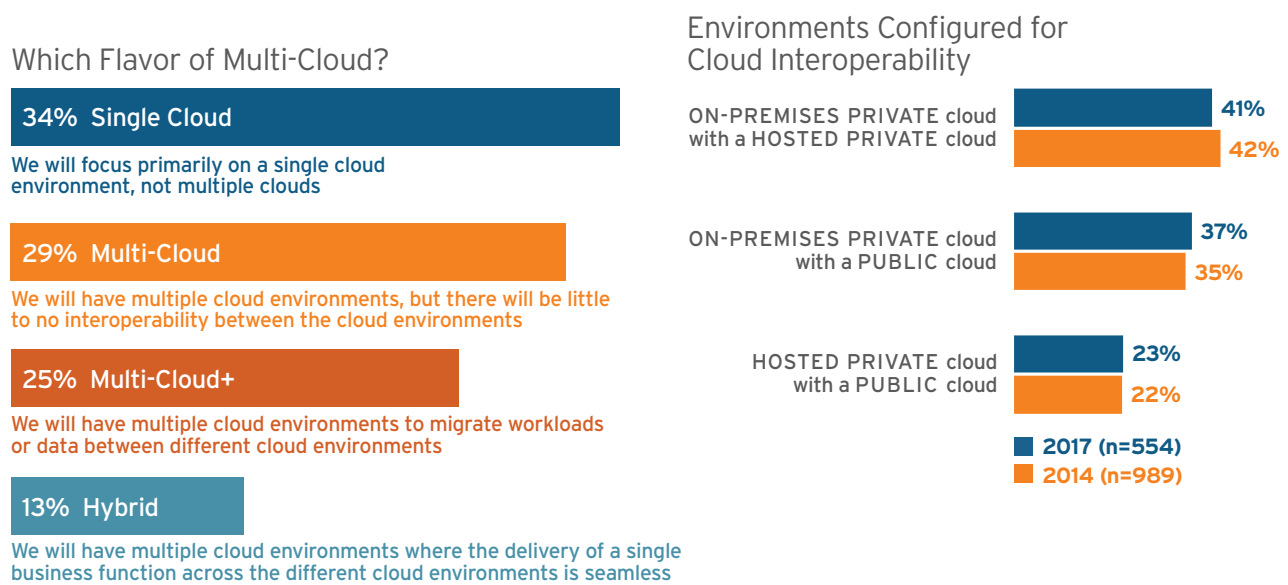
There is a fourth category – multi-cloud – that is not a specific delivery method but, rather, describes how companies use a combination of the three delivery models described above.

BEST EXECUTION VENUE AND MULTI-CLOUD

While many vendors (and some service providers) will try to convince organizations that one cloud delivery model is suitable for all applications and services, the experience of most end users is that this is not the case: for each workload, there is a ‘venue’ best suited to its deployment – the ‘best execution venue’ principle.

451 Research’s Voice of The Enterprise (VoTE) Cloud Transformation research shows that while hybrid cloud is often the eventual objective for many digital transformation projects, a single business function using interoperable clouds is, in fact, some way off for most companies. Figure 1 shows that, at the same time, while 34% of survey respondents plan to focus on a single cloud environment, the majority – 54% of firms surveyed – will use a multi-cloud infrastructure. This will entail a varying level of interoperability between clouds, as well as the migration of workloads.

Figure 1: We increasingly live in a multi-cloud world



Source: 451 Research Voice of the Enterprise Cloud Transformation, 2016

There are several drivers for the adoption of multi-cloud, but it ultimately reflects some of the key trends in infrastructure technology and use. We have already seen how the virtualization of server, storage and networking functions has led to a disaggregation of technology from IT infrastructure, and the convergence of on- and off-premises colocation, cloud and hosting services has led to the creation of new forms of managed services.

WORKLOAD CONSIDERATIONS

Different workloads have different characteristics, and the best platform for a particular component of a workload depends on a range of business requirements and technical criteria – including those related to performance and resiliency, network, elasticity, data privacy, infrastructure management and integration. Some applications require seasonal or ‘point in time’ scalability, while others have predictable and consistent needs.

451 Research Market Monitor, Cloud Computing shows that the market is in transition, moving from Phase 1, on-demand infrastructure, to phase 2, concerned with the management of cloud environments. As this transition progresses, consumers of cloud services will undertake greater enterprise adoption of more complex cloud workloads. These workloads will be product and mission-critical in scope and will require a greater use of management tools and vertical offerings.

Multi-cloud and multi-workload considerations, in turn, enhance the need for further cloud management considerations, and point to the requirement for tools, platforms and business models around which end users and their hosting partners can provide services for diverse infrastructures.

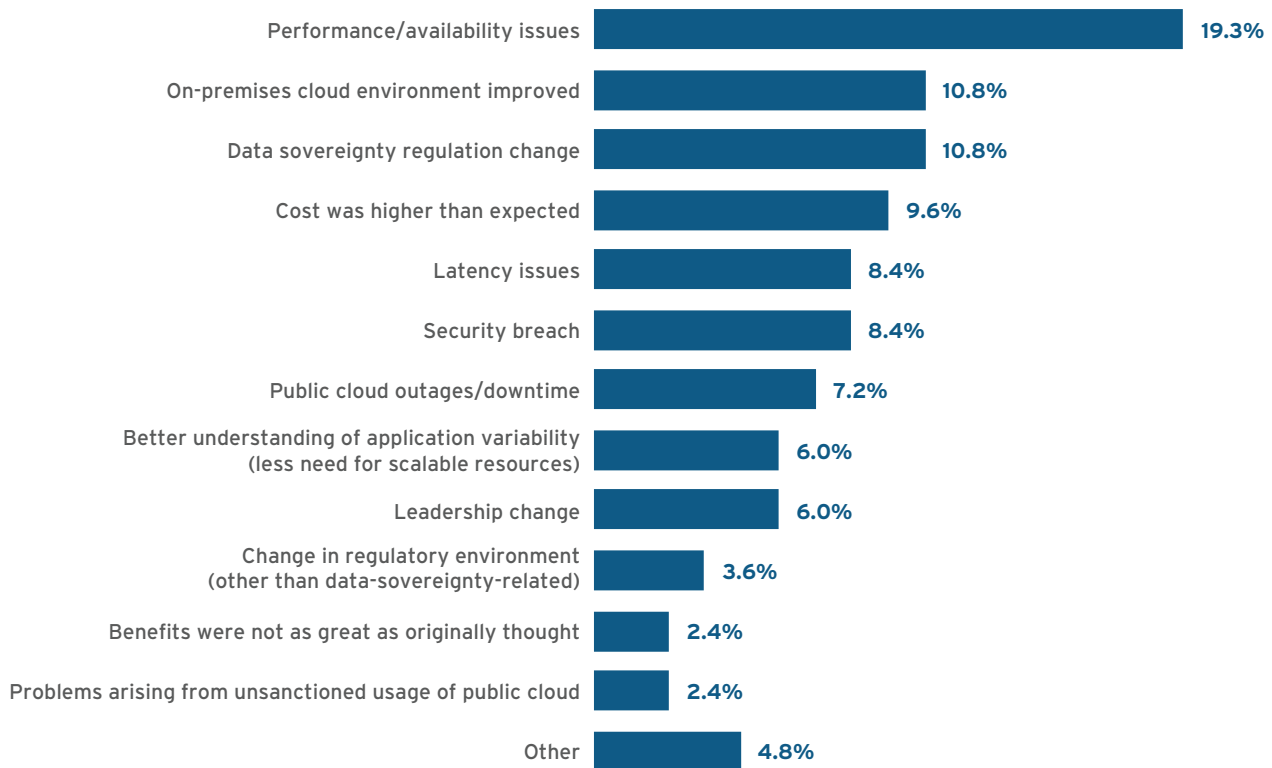
WORKLOAD MIGRATION

451 Research, Voice of the Enterprise: Cloud Transformation, Organizational Dynamics 2017, has shown that over the course of the previous 12 months, 32% of survey respondents had migrated applications or data that were primarily part of a public cloud environment to an on-premises private cloud or hosted private cloud environment.

Figure 2 shows that – aside from improvements in on-premises cloud – performance, price, latency, security and data sovereignty were the key drivers in the decision to migrate out of public cloud.

Figure 2: Drivers for workload migration from public cloud to private or non-cloud environments

(percentage of sample, n = 83)



It is clear from this chart that providers of public cloud infrastructure need to better address their go-to-market approach, SLA expectations, data-sovereignty position, ecosystem relationships and other aspects. Aside from management changes and 'shadow' IT, many of the above topics can be addressed via other cloud deployment models – an opportunity for the service provider community.

Investment and Return

BUILD VS. CONSUME

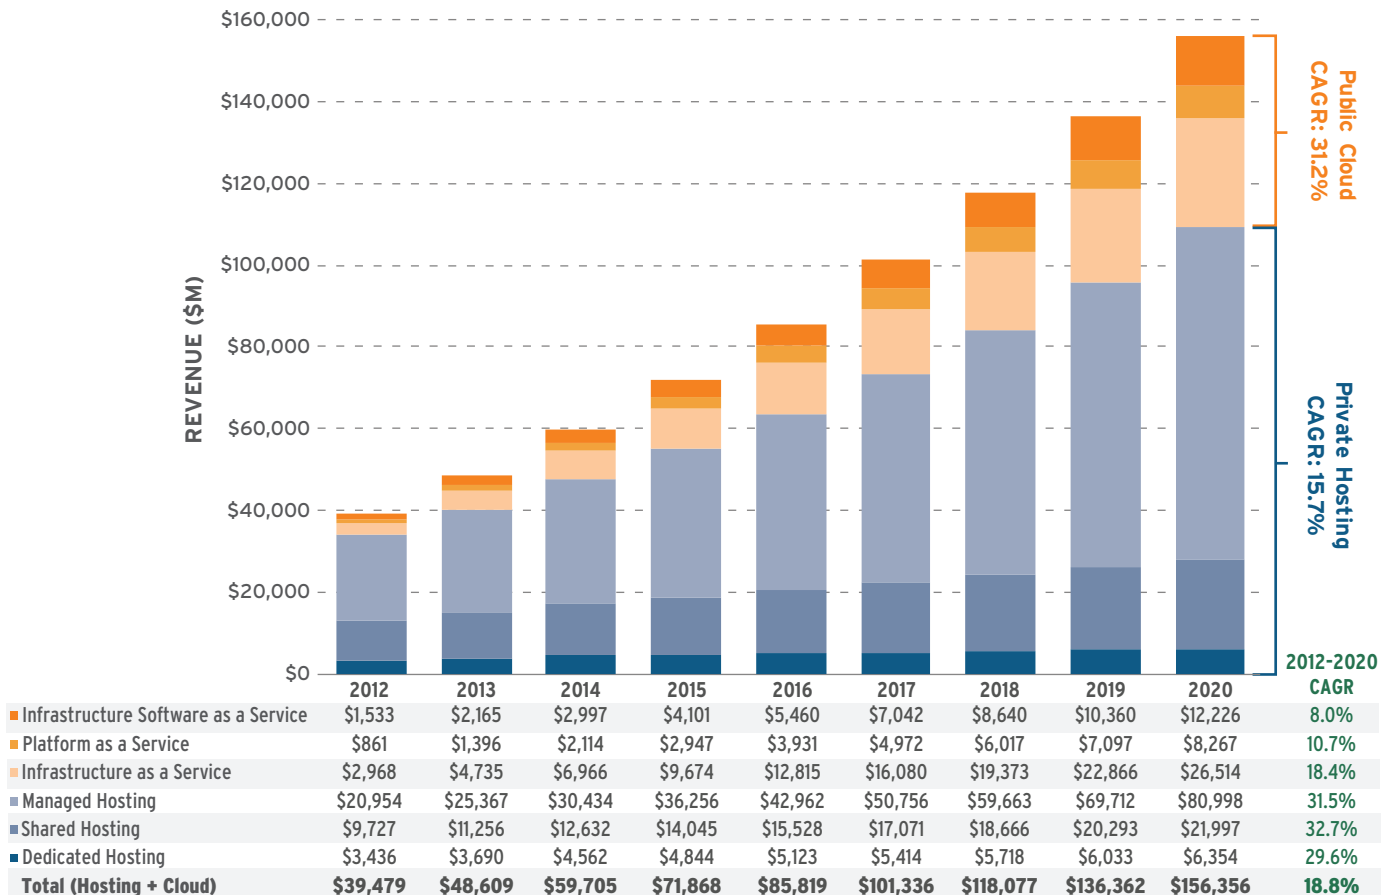
In the datacenter sector, the 'build vs. buy' debate has continued to prompt discussion among service providers and enterprise buyers of the benefits of infrastructure ownership vs. rental or colocation. The shift from infrastructure asset-heavy to asset-light can be seen in the moves by some service providers – such as Verizon's sale of its datacenter estate – to focus on more strategic areas further up the cloud stack, e.g., collaboration, connectivity and managed services.

Clearly, there are pros and cons to infrastructure disinvestment and reliance on third-party management. Those that own and manage datacenters, network connectivity, storage facilities, etc., emphasize the importance to their business model of end-to-end service capabilities, integration options for legacy infrastructure, hybrid IT development and other aspects. For regional firms delivering a large portfolio of traditional hosting services, this is a standard model under which to operate.

At the same time, data from 451 Research's Cloud Computing Market Monitor service shows that the services portion of the global hosting and managed services market is growing faster than either applications or infrastructure, at a CAGR of 18.8% through 2020. Where cloud is concerned, this growth indicates that demand is increasing in those areas where management of cloud consumption – as opposed to cloud construction – is core.

Figure 3 provides some perspective as to how this services growth is reshaping opportunities in the overall hosting and cloud space. While more traditional forms of hosting – dedicated and shared – become increasingly commodified and show much lower rates of growth, managed hosting (in all its forms) continues to take the majority of revenue. This segment is the main driver of growth for the overall hosting and cloud services sector as IaaS, PaaS and infrastructure software as a service continue to increase their share of infrastructure spending.

Figure 3: Worldwide hosting and cloud services opportunity 2012-2020



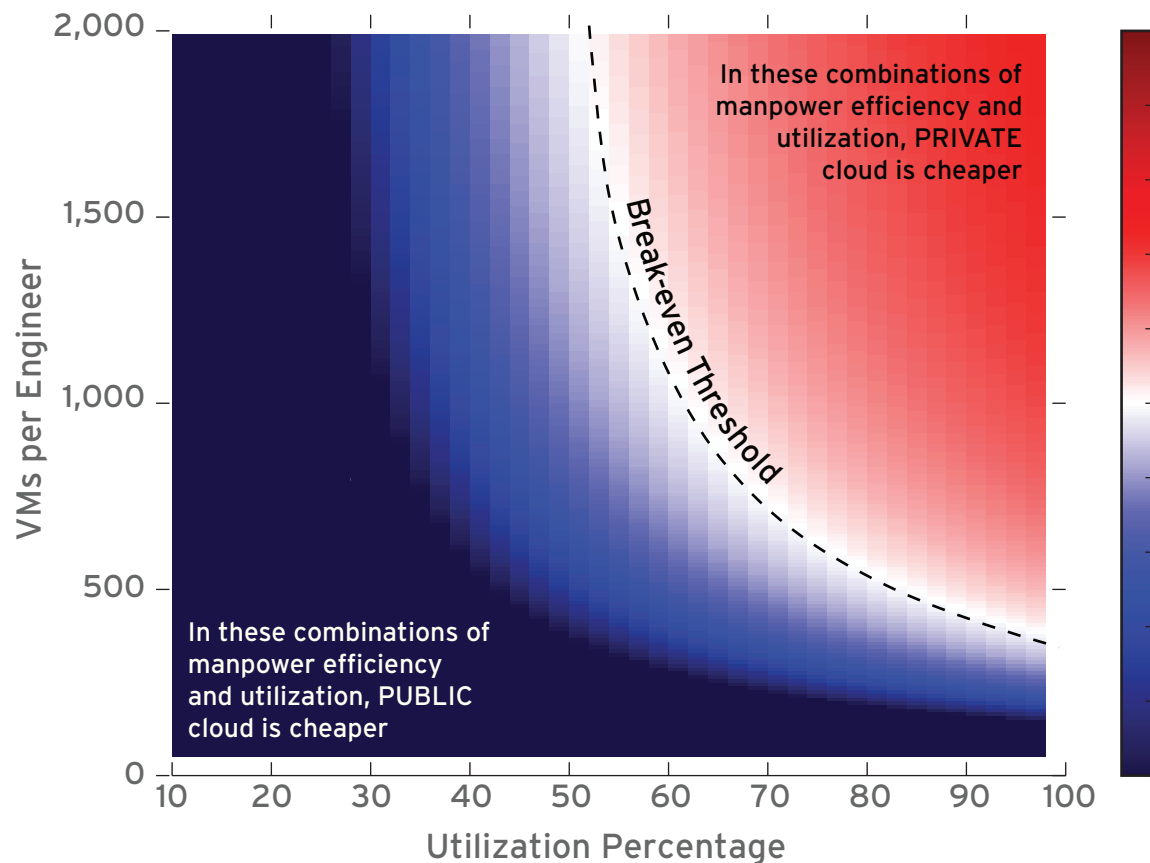
Feedback from a variety of service providers also indicates that while many services are not delivered as 'off the shelf' offerings (in the way that most public cloud offerings are), this growth is not primarily due to a corresponding increase in customized service delivery. The increasing use of software-based automation, service integration, multi-cloud management tools and white-label components enables faster go-to-market strategies, standardization and more consistent delivery.

HYPERSCALE CLOUD AND TCO

The three leading hyperscale cloud providers – Amazon Web Services, Microsoft and Google – have created global infrastructures for large-scale deployment, but not everyone chooses to use a hyperscale provider's infrastructure or can self-manage compute and storage resources internally. As we mentioned above, skills, technical resources, data-protection requirements, security concerns and regulatory aspects come into play.

There is also the question of TCO. As cloud usage has become more prevalent and the experience of running workloads on public, private, open source and other platforms has increased, users have been able to gain insights into what service utilization is costing them in the long term. The 451 Research Cloud Price Index shows that platform choice and utilization rates play an important part in TCO. Figure 4 highlights the TCO of supporting public cloud vs. private cloud based on the number of VMs that can be sufficiently managed by a single appropriately qualified engineer. The chart illustrates that to achieve a price point of \$0.15 per VM, a commercial platform only needs an engineer to be able to manage 170 VMs, on average.

Figure 4: TCO depends on utilization and manpower efficiency rates



Source: 451 Research Cloud Price Index 2016

Service providers offering managed private cloud should provide services to help increase utilization, which may include scheduling or flexible contracts to allow end users to buy capacity on demand.

Enterprise firms looking to implement their own private cloud with a lower unit price than public cloud need to achieve both labor efficiency and utilization, and keep utilization high. Tools, processes and automation need to be in place. It is possible to achieve cost benefits over public cloud by being rigorous in the approach to management – and this is where providers of tools and platforms can help.

Stable workloads are the most likely to achieve high utilization, and variable ones are least likely. 451 Research recommends that this should be factored into end-user plans, including the potential hosting of multiple applications on the same private cloud. By giving these applications different priorities, it is possible to keep the cloud highly utilized while allowing some apps to be terminated ad hoc, thereby freeing up capacity for more pressing requirements.

PORTFOLIO MANAGEMENT, BUSINESS MODELS AND INVISIBLE INFRASTRUCTURE

The rise of XaaS has shown that there is a demand for hosted cloud services, but not just by the hyperscale providers. Whether inside a large enterprise or via a third party, in order to provide something 'as a service,' the hardware it runs on needs to be maintained, servers need to be cooled, and failed drives need to be swapped out, etc. Whether an internal/external service provider or distribution partner, the retail model for cloud consumption at point of sale can be used in today's multi-cloud world as described above. Using BEV, the role of service intermediary for integration, connectivity and deployment has already started to emerge, with XaaS facilitating workload placement and hybrid integration.

And XaaS for hybrid IT is not something that has emerged purely as an alternative to self-service public cloud. According to 451 Research's VotE: Hosting and Cloud Managed Services study, managed services are becoming a larger portion of overall hosting and cloud services spending, expected by respondents to increase from 10% to 13.3% of the total amount over the next two years. Hosted infrastructure spending is also increasingly being bundled with operational managed services and security services; 69.5% of respondents indicate their next hosted or cloud infrastructure engagement would be managed, rather than unmanaged or self-managed.

This bundling of services and the increase in third-party operational management mirrors the virtualization of compute, storage and networking hardware: technology is becoming increasingly 'invisible' to the end user. This does not mean that physical infrastructure is disappearing – far from it – but consumers of cloud services are not required to be aware of CPU speed, RAM, network capacity, single-tenancy/multi-tenancy, etc. The concept of 'public' or 'private' is, therefore, less relevant as a discussion point. When talking with end-user organizations, providers of cloud services can increasingly focus on aspects of service provisioning that reflect business efficiencies – availability, automation, tools for managing consumption, billing, etc.

Channels and Ecosystems

Any significant move to cloud delivery will be dictated by demand from the service consumer. The service providers and IT distribution channel partners that are growing today and will thrive going forward are those that listen to what their customers are asking for now (whatever that may be) while balancing this with emerging and innovative services, regardless of platform.

SERVICE PROVIDER AND IT DISTRIBUTION CHANNEL DYNAMICS

As infrastructure becomes increasingly invisible to the users of cloud services and the technology behind it less relevant to consumption, differentiation becomes more important to those supplying cloud services. The management of legacy and third-party environments – whether cloud or non-cloud, on-premises or off-premises – will become more workload- and vertical-specific in order to be competitive and meet customer demand. Figure 4 illustrates the shift in workload placement that will take place over the next two years and the increase in service usage. Much of the responsibility for enabling this shift will come from the distributor in its emerging channel role as a builder of digital services to facilitate consumption through value-added resellers, systems/service integrators and others.

The ability to provide virtual IT management – with the emphasis on high-quality, user-focused services – is a sweet spot for the cloud service provider. At the same time, those channel partners that have customers at many different stages of cloud transformation – from standardization through virtualization and automation – also stand to gain. The most successful ones will be able to satisfy a variety of diverse needs while running a profitable business, realizing immediate capex benefits through establishing simple workflows and process integrations incorporating repeatable, out-of-the-box components. It is reasonable to expect that most (but not all) end users will expect a certain level of customization when progressing beyond basic cloud services. Over time, additional layers of complexity can be added by the service provider where required. Taking a turnkey approach combined with automation is an efficient way of controlling the cost of service provisioning. Providers should look to recruit workflow and process specialists – particularly those with artificial intelligence, machine learning and operational integration experience – as opposed to hiring from within IT teams.

Standardization and automation promise to give managed service providers greater efficiency and productivity, higher quality of service, scalability, increased compliance and greater profit margins.

HIGH-VALUE SERVICES - THE MANAGED OPPORTUNITY

451 Research defines managed services as those delivered by a service provider in connection with the operation of hosted infrastructure and application services, including those operated by the provider of managed services and those sourced from third parties. Examples of managed services where we see growth opportunities include end-to-end application management, infrastructure/application monitoring and alerting, premium 24/7 support services, disaster recovery, and back-up and recovery.

Having identified growth potential, it is also important to look at the 'business of cloud.' Developing sustainable run-rate revenue, maintaining commercial partnerships that deliver results and an understanding of organizational change are all essential requirements. For service providers, we have already outlined why service-level differentiation is important, but there is a balance that needs to be maintained: not enough differentiation, and competition will be broad and erosive; too much, and niche markets may be too small to sustain more than a 'hobby' business. At the same time, time to market is critical, particularly for new service lines; vertical-specific offerings may offset some of this, but not much.

There is no 'one way' to provide high-value cloud services, in the same way as, for example, no single public cloud is appropriate for all workload types. However, 451 Research VotE Managed Hosting & Cloud Services has shown that, by providing services focused on refactoring, migration, moving data and adopting new applications, providers can demonstrate an additional level of expertise in both the workload in question and the hosted environment. Hosting and cloud service vendors can and should strategize around these requirements, either by focusing on attracting the workloads that best fit their capabilities to support them or by developing the capabilities to support the workloads they intend to target.

For those enterprise companies in the process of becoming an internal service provider for their stakeholder communities, the challenges in providing cloud services are often not related to commercial or technical aspects. While budgets, cost of migration and integration difficulties remain perennial topics, cloud deployment is often delayed due to human aspects – resistance to change, fear of redundancy, lack of skills, etc. This, in turn, can inhibit business growth and slow the implementation of well-funded and strategic initiatives.

Service providers again have a role to play in this scenario, not only by helping enterprise IT departments 'join the dots,' bridge skills gaps and provide training, but also by reducing fear and uncertainty through non-disruptive means: baked-in automation, service analytics and self-healing, as well as more advanced aspects such as machine learning and artificial intelligence.

Conclusion

Cloud workload rates are increasing, but it is not a one-way street toward a single method of deployment. Because of the experience end users have gained from running workloads in clouds, cost and BEV, some clouds are being repatriated, while multi-cloud has become the 'new norm.' For end users – the consumers of cloud services – it should not matter where workloads are located or how they are delivered, as long as the business function of that application or service is fulfilled. Infrastructure is more critical than ever to make this happen, yet is increasingly virtualized, integrated and converged, making it more 'invisible' to consumers. The distinction between private and public should be immaterial and of little concern: awareness of cloud construction becomes much less important than the process of cloud management and consumption.

For those delivering cloud services, delivery and management of multi-cloud are creating opportunities further up the stack for value-add services. At the same time, expectations of availability are higher than ever, creating a requirement for significant levels of automation, rapid provisioning and services that address complex enterprise needs while retaining high-touch delivery. At this level, choices need to be made about cloud delivery, but again, there is no 'one size fits all' for the cloud.

For the enterprise user, individual workload requirements, long-term TCO and integration levels matter more. For the service provider, achieving this level of service provisioning is not easy, but the opportunity exists for ecosystem partnerships with technology vendors that offer 'best of breed' management tools, deployment platforms, delivery capabilities and service differentiation.



The following information is provided by Ormuko. 451 Research has not validated any claims and does not endorse Ormuko or its offerings.

ENDING THE PUBLIC VS PRIVATE CLOUD DEBATE

Managed Service Providers Can Offer 'The Best of Both Worlds'

Enterprises today are focusing on cloud consumption rather than cloud construction; what matters most to enterprise cloud consumers is meeting business objectives, not engaging in a debate about which IT infrastructure best supports their workloads and applications. In fact, most organizations are realizing that the public cloud versus private cloud debate is over; a hybrid or multi-cloud approach—where the platform of choice, or “Best Execution Venue,” is dictated by the needs of each workload at any given moment—is the optimal path forward.

This is excellent news for Managed Service Providers who are able to provide Best Execution Venues and exceptional service while making infrastructure invisible to their customers.

HERE'S HOW ORMUKO CAN HELP

Ormuko's cloud solution—Ormuko Stack—offers service providers, independent software vendors (ISVs) and resellers true workload portability in an open, vendor-neutral, hybrid cloud environment.

Ormuko Stack is a full-featured, turnkey cloud platform that works seamlessly across public and private environments and addresses performance, security and data sovereignty requirements, enabling businesses around the world to choose where and how their applications are deployed and data stored.

With Ormuko Stack, VARs, SPs and ISVs can offer their enterprise customers a public cloud experience—easy, convenient and self-service—inside their own datacenters. Ormuko Stack also allows enterprises to burst seamlessly from on-premises to off-premises managed clouds when workloads demand it, without having to worry about complex integrations. Ormuko Stack offers this because the architectures of the on-premises cloud and the off-premises cloud are fully compatible.

Fully managed and updated by Ormuko, the platform is engineered around RESTful APIs and achieves operational excellence through artificial intelligence (AI), automation and machine learning, thereby ensuring scalability, efficiency and cost optimization.

BENEFITS OF THE ORMUCO SOLUTION

- **Intelligent, Cost-Effective Cloud, Rapidly Deployed:** AI and machine learning deliver self-healing systems with better economics, improved agility, streamlined operations, and enhanced customer experiences. Self-healing and end-to-end automation deliver quick and efficient cloud deployments with dramatic reductions in time to market.
- **Open Source Core & Powerful Economics:** Ormuco Stack leverages open source technologies and its operational expertise to deliver continuous innovation and significant cost savings.
- **Simplicity:** Ormuco Stack spares customers the hassle of complex integrations, complicated upgrades, and tedious lifecycle management. Customers are free to focus on consumption, not construction.
- **Service:** Customers benefit from Ormuco's depth of experience and customer service commitment. Ormuco installs, configures, manages, operates, upgrades, and supports the Ormuco Stack platform so that VARs, SPs and enterprises can focus on value-added services.

For more information, visit www.ormuco.com or e-mail us at sales@ormuco.com