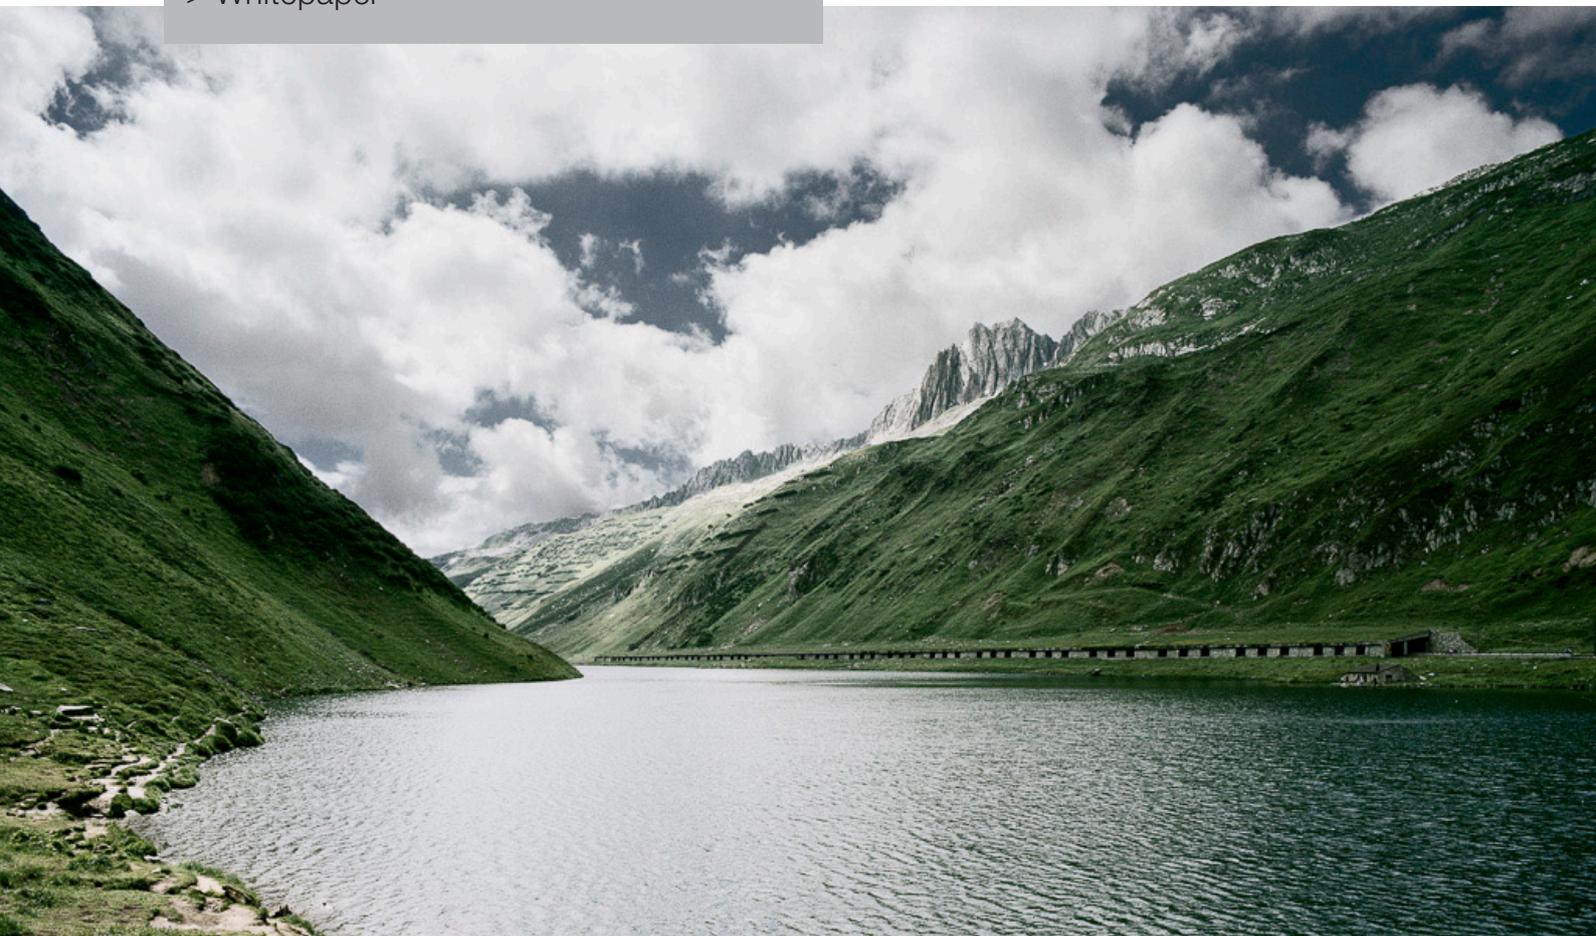


Hybrid Cloud

The way to a successful hybrid cloud strategy

> Whitepaper



Hybrid Cloud

Introduction

We already use many different cloud services in our daily lives these days. Even our non-technically-inclined friends and acquaintances share their files via Dropbox or similar applications. Some of us listen to music using streaming services such as Spotify or Simfy. Many of us use Google Mail or Outlook.com for their e-mails and are connected with social networks via Facebook, Twitter, Instagram and Pinterest.

Now even a strongly growing number of companies are also using cloud-based IT services. Amazon Web Services (AWS) boasts an annual turnover of almost 4 billion dollars, and is still growing at a breathtaking rate. Google Cloud Platform and Microsoft Azure compete in the same area and offer similar functionality. Certain providers have put highly solution-specific services on the market. Salesforce.com, for instance, rapidly became the number one CRM system of choice for small, mid-sized and large enterprises.

Cloud services are highly diverse and we could ponder whether it makes sense to think of them as a unit. Yet, cloud services do share a whole series of similarities. The most important feature is resource pooling. By pooling resources of many different customers, users and applications, users are able to achieve an economy of scale that leads directly to better efficiency. These resource pools have to be accessible over an extensive network and, for the costs to be properly allocated, the precise use of the services must be logged in detail. This paves the way towards greater elasticity, flexibility and agility of the services.

There are of course a number of challenges to take into consideration. A company's security profile, for instance, changes significantly when using shared infrastructure and publically accessible networks. The distribution of applications and data to multiple providers can also add to the complexity. By outsourcing parts of the infrastructure, the company or user often has to re-establish adapted processes and introduce new organizational structures.

None of these points, however, is an unsolvable problem. They all simply need to be accounted for in a careful plan. At the heart of this plan is the data center. It is the basis of all cloud services. It is essential to choose a data center that suits the targeted IT or cloud architecture.

Need

It makes sense to tackle this plan by first posing the question: Why opt for the cloud? Today's need for cloud computing services stems directly from the current information technology requirements. Above all else, these IT requirements follow three main tenets: better efficiency, rapid elasticity and focus on core business.

Many see lowered costs as the main advantage of cloud computing. Yet, it is important to recognize that a public cloud is not always cheaper than an in-house solution. To reveal the true benefit, one has to take a closer look at the utilization of the individual applications.

Utilization curve 1

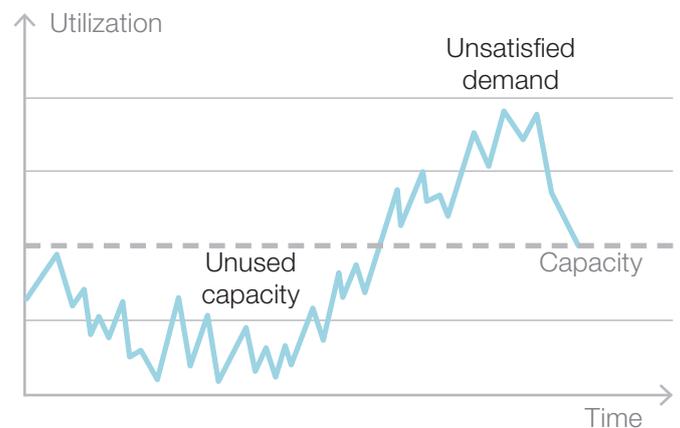


Figure 1

When applications have a highly variable need for resources (memory, processor, or network) then the result is double inefficiency (Figure 1). Non-virtualized environments these days often run at a server utilization of less than 20%. That means there are necessarily periods in which the full capacity is not used, but the full costs still have to be paid.

On the other hand, there are some times when the capacity falls short of the application demand. This puts limits on services that could be profitable or business-critical. While it is possible to minimize the frequency of these bottlenecks by forever increasing the capacity, this simultaneously increases the inefficiency during times of overcapacity.

In a cloud, it is no longer necessary to overprovision resources just to cope with peaks in demand (Figure 1). You pay only for the capacity you have actually used. That allows you to reduce the costs and even allocate them more easily than continuing to account for them internally.

A cloud-based infrastructure further expands the company architecture with additional flexibility and agility. It becomes far easier, for example, to introduce new services or to take down old applications when they are no longer needed. It is not necessary to acquire or cheaply dispose of the corresponding hardware. As alluded to at earlier, one can scale a service upwards or downwards without delay.

Hybrid Cloud

With Internet services, above all, there are cases where the resource requirements will increase tenfold from one day to the next, and then just as suddenly decrease again. The elasticity of the cloud allows the company to adapt its resources exactly to the demand.

Its flexibility also allows more rapid introduction of new applications. A common criticism managers have of their IT department is that they often have to wait months to introduce new services, and thereby lose a valuable edge in the market. With an automatically provisioned cloud, the typical lead time for acquiring the necessary equipment can be reduced to a few minutes since the resources can be provisioned at request.

A globally replicated cloud facilitates access from any location using any equipment at any time, and contributes to the flexibility and productivity of the users. This advantage becomes even clearer when it is necessary to integrate business processes with those of your suppliers, partners and customers. A simplified firewall configuration makes it easier to allow fine-grained access to services without compromising the security of your sensitive data. The fact that some of the IT services can be outsourced to a cloud provider reduces the effort and administration that an IT department would otherwise necessitate. The scope of these duties covers user provisioning, application management and troubleshooting. By automating these services, experts can be made free for activities that support the company's core business.

Models

Cloud services are often divided into two dimensions. The Cloud Service Model distinguishes between various technical types of cloud services. The Cloud Deployment Model relates to the commercial agreements between service provider and client, where these also have technical implications.

Cloud service model

Cloud services are often divided into three categories or layers. The bottom layer offers infrastructure services – such as Amazon Web Services, Google Compute Engine or Openstack. The middle layer offers platform services – such as Google App Engine, Amazon Elastic Beanstalk, Windows Azure or Cloud Foundry. The top layer offers software services – such as Google Apps, Microsoft Office 365 or Salesforce.com.

All three layers have one principle of cloud computing in common: resource pooling (Figure 2). Infrastructure services abstract the hardware – and thereby allow the bundling of physical infrastructure. Platform services also share the underlying operating system. And software services benefit from common application instances. Most cloud applications are based on an IaaS platform, since it is highly suitable for putting legacy applications into the cloud.

Utilization curve 2

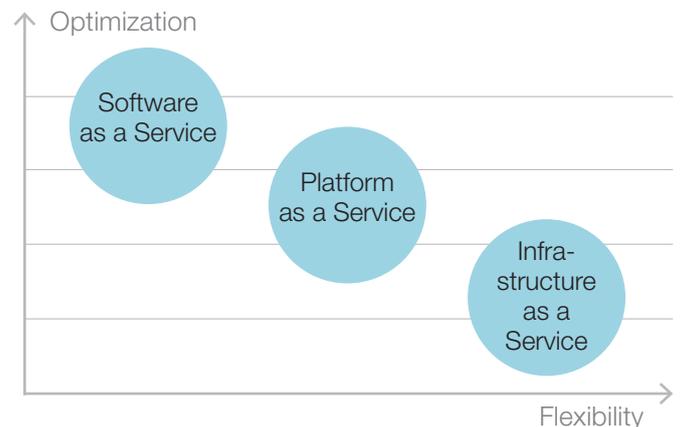


Figure 2

Because it only shares very few resources, however, it only offers limited advantages, whereby only short-term improvements can be achieved.

Cloud deployment model

There is yet another distribution of cloud services, based on ownership and operation. Before we take a closer look at this division, however, we should examine another trend, which already began before cloud computing, namely the attempt to centralize computing power. Large companies had already expanded their computing space with the transition from mainframe to client server computing, in order to merge as much capacity as possible. This resulted in the creation of large data centers that brought certain economies of scale with them in terms of power, cooling and physical security.

Smaller companies could not keep up with this on their own. Soon, however, there were colocation service providers on the market who were equipped with space, power and network connections, and bundled additional hardware, software and management along with it. Many large enterprises thereupon made use of these services for various reasons. Some preferred not to run their own data centers; others only needed minimal capacity at remote locations, or only for a short time.

This is where cloud computing goes a step further, in that it offers standardized computing services. Depending on who provides these services, we can distinguish between a public, private or hybrid cloud. Public cloud services generally benefit from high economies of scale, and thereby offer better efficiency and the illusion of unlimited scalability, where private implementations cannot keep up easily. For this reason, many experts argue that the only legitimate clouds are the public clouds.

Hybrid Cloud

In reality, however, the established IT providers have already been working for many years on developing a convergent infrastructure, and many of these aims and technologies are very similar to a public cloud. Indeed, they often reflect a public cloud with the sole distinction that the clients are internal departments and the providers the IT department.

Just like the public clouds, this convergent infrastructure improves resource efficiency. This is achieved by bundling applications and users from various specialist departments. They employ a high-performance network and require fine-grained reporting for their internal accounting. They also endeavor to improve operating efficiency using the same technologies, such as automation and self-service. Sourcing efficiency represents the last step in introducing cloud computing. It provides the flexibility to allocate services and resources from several internal and external providers without changing the enterprise architecture. The only way to achieve this flexibility is to align all systems strictly to the principles of service orientation and service management.

As soon as the companies have fully decoupled their sourcing, they are flexible with respect to using the required services. That means, they can continue to receive the applications from IT, but also have the option of swapping to an external service provider if that provider is more efficient or more reliable. Do not forget, however, that this independence can work in both directions. If IT modularizes your services, then the organization also has the flexibility to offer certain functions on the external market, and so to monetize the IT investments in a way that was not possible before.

If an organization embarks on a journey into the hybrid cloud, then the task is not only to select an increasing number of publically offered services. The bigger challenge is to integrate them dynamically, with both internal and external services.

To make this clear, we shall look at some of the integration options for hybrid cloud computing. For a first step into the hybrid cloud, all that is needed are services that run in independent silos, without any interaction between them. For example, an organization could employ Microsoft Exchange internally and Salesforce.com as its public CRM tool. But the goal must be to ultimately achieve an exchange between the two applications.

The next step would be to integrate them where it makes the most sense. This could be connecting Salesforce.com with the internal Active Directory, in order to allow single sign-on or use the service to send e-mails and enter appointments. Of course, this integration requires careful planning in order to guarantee compatibility of all components and to protect sensitive data.

The last step is to support a dynamic load distribution. That means the internal and external services have to be homogeneous and interface-compatible. One reason for taking this approach would be to allow disaster recovery. Upon a total failure of the internal data center, the company could push its jobs into the cloud and start up its services there again.

Another ambitious goal would be cloud bursting, which is a good way to optimize costs and flexibility. If the organization is in a position to shift its workloads in real time, then, during normal times, it can have its services run internally, where the costs are lower. Should it come to peaks loads, or should the service grow faster than originally expected, then the company can outsource the incremental load that exceeds its internal capacity.

Implementation

Cloud solutions bring many advantages, yet it is equally important to consider the challenges.

Security

The most common concerns when it comes to the cloud are security and risk.

- Client capability increases the risk of data loss upon access by another client
- Provision of the services over the Internet creates a larger target in terms of data connections
- Lack of clarity in international law and national legislation make it difficult to comply with statutory requirements
- Virtualization hinders the monitoring of resources and can lead to uncontrolled proliferation in the server landscape
- Outsourcing reduces both control and transparency with regard to services and data
- Lack of standardization can lead to unwanted commitment to the provider

A common obstacle to adopting cloud computing is the fact that the service provider hosts sensitive data – possibly even in a multitenant environment. The client must therefore decide whether he can trust that the provider will not inadvertently (or even intentionally) compromise the information.

Furthermore, the data must also be transmitted between the cloud and the company. If this connection runs through the public Internet, then there is a great risk that competitors or hackers could intercept and thereby gain unauthorized access to sensitive information.

Data security is of critical importance not only as a competitive advantage. Especially in Germany, there are very strict laws that prescribe how and where personal data may be stored. Equally, many authorities and trade organizations demand a strict degree of security of financial data (e.g. credit cards).

It is therefore highly important to be mindful of the location of a cloud computing center. In the cloud, especially, it is sometimes difficult to discern exactly where data are being stored. Yet, their location is what determines which laws regulate access and under what circumstances the data may be forwarded to third parties and what security measures must protect this data.

Hybrid Cloud

The monitoring of network, server and storage also becomes much more difficult given the additional virtualization layer, since it is no longer possible to recognize all systems. From this, there is also a greater risk that unauthorized virtual machines could be running, or legitimate systems are not properly shut down and continue to create costs over an indeterminate period of time.

The public cloud is essentially an outsourcing service that brings all its advantages and disadvantages with it. I.e. the company gives control and even a degree of visibility to the service providers, and is only indirectly able to enforce its own requirements.

Yet, should it ever be necessary to change provider, then a fundamental problem often materializes. Since the interfaces are typically different from one provider to another, there is a considerable effort and expense involved in making such a changeover.

Integration

There are numerous technical challenges that an enterprise must address if it is to introduce a cloud solution. It is no easy task to integrate the various networks in a reliable, scalable and secure fashion. There may be uncontrollable sources of latency or bottlenecks in data transfer. It can also be difficult to manage the encryption keys needed to protect all channels of communication.

Another common challenge is integration at the application level across organizational and technical boundaries. Again, there are security precautions to consider, but there are also more general problems of interoperability and standardization of interfaces. Furthermore, it can be difficult to maintain data integrity when critical activities need to be handled asynchronously.

Management

The client must consider how he can organize operational processes such as incident management, change management and a service catalogue in a multi-vendor, multi-provider environment. To reduce this complexity, it is often helpful to consolidate the management of different systems over a uniform console.

The cloud management platforms give the administrator an overview of the entire system landscape and centralized management. Even cross-cloud actions can be performed to shift resources between providers. The technical implementation of these continually growing requirements, however, requires careful planning and in most cases extensive configuration.

Organization

Ultimately, the shifting of functions into a public cloud means the outsourcing of services. This outsourcing usually has repercussions for the employees. It doesn't necessarily mean that employees previously responsible for performing these services become redundant. But it does mean that their duties need to be reviewed in order to harmonize them with the new operating model. There

are many changes that can happen in this scenario:

- **Market-oriented architecture and design** – Architecture and design functions play an important role in a complex cloud environment. They require a deep understanding of market trends and of the services and architectures present on the market.
- **Shorter decision cycles** – A key benefit of cloud computing is that it reduces time-to-market. To realize this benefit, however, internal decision-making processes need to be accelerated.
- **Closer integration of operations and finance** – Cloud computing offers greater elasticity with a pay-per-use model. However, this also requires integration of operational and financial processes so that the organization can harness the benefits of elasticity without losing control of its finances.
- **Increased role of provider management** – As the organization becomes more dependent on the cloud provider, it will need more mature processes for managing provider contracts and changes to contracts. The organization must become much more agile and flexible given the rapidly changing market landscape (new offers and new providers can rapidly change the landscape).
- **Mature incident management capability** – IT needs clear responsibilities, workflows and escalation paths to perform incident management and other ITIL processes.
- **Well-defined SLAs/OLAs** – Effective management of providers requires well-defined service-level agreements (SLAs) and operational-level agreements (OLAs) with the most important key indicators that effectively and reliably measure the quality of services.
- **Continuous development** – Since the organization's IT capabilities depend increasingly on using the cloud effectively, and at the same time the cloud technologies and services are changing very quickly, the employees must continuously develop their skills to ensure they stay up to speed with market trends and best practices. The IT organization must become more agile in order to respond to changing business expectations and market conditions.

Implications

Transformation from a conventional infrastructure to a cloud-based architecture is no trivial task. Nevertheless, it is possible to meet the new challenges listed under paragraph four. It is possible, for example, to encrypt sensitive data, to employ tools combining networking, application and management elements, and to introduce new governance processes and policies. These tasks are made considerably easier if a hybrid cloud is planned carefully from the outset. I.e. it must be decided early on what data and applications are suitable for a private cloud or a public cloud. It must also be considered what in-house

Hybrid Cloud

requirements influence the choice of cloud provider. An essential point is the data center in which the services are to be hosted.

The new role of the data center, taking e-shelter as an example data center service provider

The change in IT technologies and new approaches such as cloud computing, big data and mobile computing, driven by the unrelenting forces of standardization and competition, influence the importance of the modern data center.

The data center is developing into a center of gravity for an agile and flexible IT structure that can dynamically follow the adaptations to the requirements of modern IT. With the increase in IaaS, PaaS and SaaS providers, the data centers are now in direct competition with external and internal IT service providers in the enterprises. Yet, they can also complement one another very well given their different focuses. The enterprise's in-house IT must realize, however, that the relationship between the enterprise units and the data center is changing, and that they must face up to these changes. The general trend among enterprise data centers not to operate entirely on their own, but to outsource to data center service providers, is strongly supported by this technological change.

Since the importance of the data center, its location and its services is continuing to rise as part of the IT value added chain, it is very important to select suitable data centers.

The criteria for this are many and varied. One first filter for larger clients is capacity. It only makes sense to consider a data center if it offers ample space, can deliver the required power and has a sufficiently powerful network connection. In order to guarantee high availability, these should be redundant and boast automated switching processes.

Security is a complex topic that no enterprise can afford to ignore. Good physical protection, strict access control and certified security processes are crucial factors to be verified. It is also indispensable to verify that the data is kept in compliance with the law. This topic is of utmost importance in Central Europe especially, since the laws on the storage and geographic location of personal data are very strictly formulated by international comparison.

Enterprises can achieve the best economic benefit from a hybrid cloud. In the hybrid cloud approach, IT resources are intelligently distributed taking into account the various requirements of business applications. Figure 3 shows an example for the distribution of resources.

The basic concept of a hybrid cloud allows the combination of on premise applications (installed locally) with cloud applications (private or public) by failsafe and high performance connections.

“Home to the Cloud”: e-shelter offers secure, direct access to all providers of cloud solutions

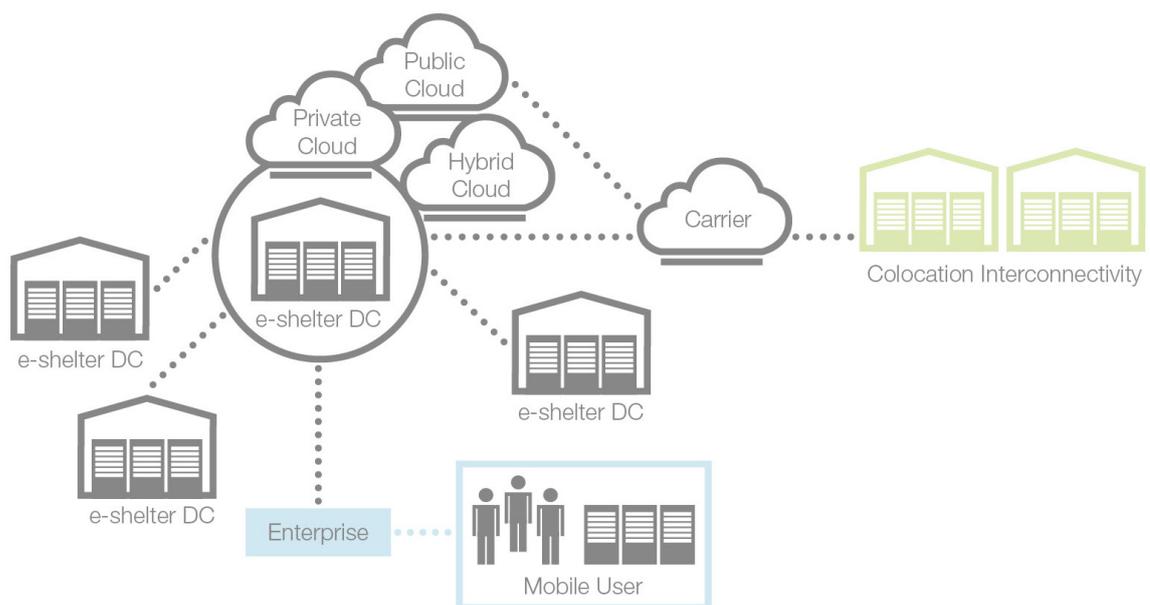


Figure 3

Hybrid Cloud

Access to cloud services and Internet nodes at the e-shelter location



Figure 4

The ideal case is when some of the cloud service providers are located in the same data center, and therefore a connection can be offered with minimal latency, highest bandwidth and the best service levels. If this is not the case, then a broad range of stable and secure connections of the data center locations and reliable access into further cloud services (direct connect), backed by SLAs, ensure the dynamic use of the hybrid cloud.

The colocation provider e-shelter names its cloud ecosystem “Home to the Cloud”. The concept stands out for its high agility and flexibility. e-shelter is able to provide colocation spaces at its locations for enterprises at short notice. Furthermore, the large footprint of e-shelter gives rise to a broad range of

- Public cloud providers
- Private cloud providers
- SaaS providers
- Carrier POPs
- Cloud direct connects
- Internet exchange connections
- Cloud consulting companies

who offer their services at the e-shelter locations, or with whom e-shelter cooperates closely. The short paths between the public cloud, private cloud and colocation areas alone yield many advantages in terms of latency, security and costs.

That makes e-shelter an attractive location not only for enterprises to operate part of their IT infrastructure or utilize cloud services.

The e-shelter cloud eco-system is also very interesting for the cloud providers themselves, since the providers’ requirements for agility and flexibility, service and short implementation times can also be met. Furthermore, the location becomes more attractive as a “market place” to customers or enterprises, given the diversity of providers.

Conclusion

Awareness is increasing that there is no escaping the cloud. Better efficiency and flexibility are indispensable in a highly competitive market climate.

The products and services of all vendors are each adapting in turn to the new technologies. And the fact that suppliers, clients and partners have switched to cloud-based delivery models forces the transition to compatible interfaces that work on the same basis. Nevertheless, companies still want to use traditional IT as well. Existing applications and infrastructure cannot be replaced in the short term. They often still fulfil their purpose very well, and produce only minimal extra costs during their lifespans. A transition means also radical rethinking of security, control and management, which can only be undertaken with careful planning and a certain lead time.

A hybrid cloud is the combination of private and public cloud services that maintains the advantages of both solutions. If set up

Hybrid Cloud

correctly, it allows dynamic elasticity and cost-based pricing of the public cloud for irregular and unforeseeable services. It can even continue to use the existing systems and to maintain traditional security and management policies.

The biggest challenge is its complexity. The topic is extensive and involves not only the sum of all private and public cloud technologies. It is also necessary to integrate the entire solution and to factor in higher latency times, unreliable connections and heterogeneous security and maintenance measures.

For every core problem, there is a corresponding solution, but depending on how deeply the internal IT department can and will deal with the issues, it is often sensible to employ a consultant who knows the market and the technologies and who already has experience in implementing cloud computing.

The cloud will bring about profound changes. One can therefore expect it will take several years until all aspects of the cloud are implemented and its maximum value is exploited. Yet, precisely because of this long introduction phase, it is recommendable to start as soon as possible.

You need to consider now whether cloud services make sense for your enterprise in principle. A comprehensive, long-term cost-benefit analysis can complement the initial analysis in order to prove whether or not the cloud comes into question for your enterprise.

You can then assert the cloud focus on all new investments and all new applications. In order to exploit the advantages of standardization, for example, you should define architecture and security standards, which you can then harmonize with those of the cloud service provider in your service level agreement. You must also not neglect the necessity to equip your employees with the necessary skills to successfully operate and manage the cloud. The cloud offers many new advantages and at the same time brings new challenges with it. In times of change, it is often the case that the market changes in several dimensions simultaneously, where there will always be winners and losers. Companies that recognize these changes in time have the best opportunities to exploit the advantages by satisfying their customers, increasing their profits and improving their position relative to the competition.

About Accenture

Accenture is a global management consulting, technology services and outsourcing service provider with around 319,000 employees serving clients in more than 120 countries. As a partner for major business transformations, Accenture brings unparalleled project experience, comprehensive capabilities across all industries and business functions, and knowledge from qualified analyses of the world's most successful companies into a collaborative partnership with its customers. Accenture generated net revenues of USD 30 billion for the fiscal year ended 31 August 2014. Its web address is www.accenture.de.

About e-shelter

High-availability data centers are the backbone of the digital economy and the Internet. Since 2000 e-shelter designs, builds and operates high-availability data centers, whose infrastructure guarantees the highest standards of physical security and operational reliability. e-shelter operates approximately 90,000 sqm of data center space, on eight sites. 60,000 sqm space alone are located at e-shelter's main site in Frankfurt, Germany, making it Europe's largest single data center site. Additional sites are located in Berlin, Frankfurt, Hamburg, Munich as well as in Zurich, and Vienna. Among e-shelter's clients are financial services companies, telecoms operators and IT service providers as well as cloud service providers. As a part of NTT Communications Corporation e-shelter provides access to a network of 140 data centers around the world.

Our own security personnel of e-shelter security guarantees the security of our data centers and develops safety concepts according to individual requirements.

With our extensive experience in operations of data centers we are a sought-after contact especially when it's about individual solutions for complex projects and high power density. Because of the size of our data centers we offer the necessary flexibility for hybrid IT solutions and direct access to cloud providers in particular.