intug

INTUG Paper

The Chain of Trust

A Cloud Service Provider Evaluation Guide

June 2013

Content

About INTUG
INTUG is an international association of business users of telecommunications, bringing together national and multinational user associations throughout the world. With members and contacts in all five continents it has a global presence.

1. INTRODUCTION

Selecting the right Cloud Service Provider (CSP) to trust with your applications and data can be a difficult process. The aim of this publication is to guide you in the selection of a CSP. It needs to be stressed this guide is not a one-size-fits-all exercise and as such the conclusions derived from it are highly dependent from the context in which it is used.

This guide is meant to help in the evaluation of CSPs. It provides no guidance in the assessment of whether or not a service is suitable to be moved to the cloud, nor is it intended to help you in building the business case or the risk assessment for moving applications and data to the cloud.

As is the case with many other industries, the CSP landscape is very diverse; there are lots of different solutions available, each with its specific advantages and disadvantages. The real challenge is to find the offering best suited to your specific needs. Compare this with buying a car: a 2-seat Ferrari is a very nice sports car but it is probably not the most practical choice for bringing the kids to school. This brings us to a number of important questions you should ask yourself when assessing a possible CSP solution:

- Does the service proposed offer a proper match for your business needs?
- Does the service pricing fit with the budget stated in the business case?
- Have all risks and the costs to mitigate those risks been addressed?

Keeping these questions in mind when evaluating cloud services will help you in choosing the proper solution for your needs.

Allow us to be your guide for cloud solutions, by highlighting a number of questions to ask that can help you with evaluating CSP offerings.

The evaluation process proposed in this guide is based upon the five essential characteristics of cloud services as defined by the National Institute of Standards and Technology (NIST[1]):

- On-demand self-service
- Broad network access
- Resource pooling
- Rapid elasticity
- Measured service

In addition to assessing the degree of cloud-readiness of a CSP based on these five essential characteristics a number of additional questions to ask your potential CSP will be provided at the end of this document.

We will only briefly touch upon each question in this guide, and you should make sure your assessment takes into account both your context and expectations.

[1] Founded in 1901, NIST is a non-regulatory federal agency within the U.S. Department of Commerce. NIST's mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve the quality of life. Source: the NIST web site

2. NIST CLOUD CHARACTERISTICS

It is essential to check if the CSP solution proposed is compatible with the characteristics of a cloud-computing environment as defined by the NIST.

2.1 On-demand self-service

> "A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider."
> (Source: NIST)

Q: What is the provisioning time for adding additional resources?

A: Some platforms support automated provisioning and have a lot of spare capacity. In such a case, the provisioning of additional resources should take only minutes. Other environments require manual intervention or even the manual addition of physical resources such as RAM or storage capacity before the provisioning can be performed. In such a case, it might take a few days before the requested resources are made available.

Q: Is auto-provisioning available?

A: Suppose your web server experiences a massive load increase during the evening following a new product announcement. In such a scenario, you might want additional servers to be provisioned automatically to make sure your website response time remains acceptable. Once your web server's load decreases, resources should be removed automatically to prevent you from having to pay for resources not used or needed.

Q: Is self-provisioning possible?

A: Depending on your company's knowledge and experience, you might want to provision resources yourself using interactive tools instead of having to call the CSP whenever you want to change your cloud-computing environment. Try to get a detailed description of the configuration tools available and verify the availability of management APIs allowing you to integrate cloud management in your existing IT management framework.

2.2 Broad network access

> "Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations)."
> (Source: NIST)

Q: What connectivity options are available?

A: In certain scenarios, a standard Internet connection might be what you need but in other cases you might need a Layer 2 connection to your CSP.
There are a lot of different connectivity options available, so make sure you know what is offered and at what cost before you sign a contract.

Q: Does the CSP offer redundant network connectivity?

A: If you are moving business critical applications to the cloud the availability of these applications is of extreme importance. Therefore, all cloud computing components necessary to support these applications should support at least N+1 redundancy, including the network connectivity components. Of course, CSP network redundancy will not help if your site loses its non-redundant network connectivity!

Q: What is the minimum guaranteed bandwidth?

A: Sad but true: bandwidth overbooking happens a lot, based on the idea that not all customers will be using the maximum available bandwidth simultaneously. Ask your CSP what the minimum guaranteed bandwidth is under full load conditions. Remember that you need to be able to send and receive network traffic, so make sure to get the minimum guaranteed bandwidth for both incoming and outgoing traffic.

2.3 Resource Pooling

> "The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth."
> (Source: NIST)

Q: Is the CSP's infrastructure sufficiently redundant for supporting your applications and data?

A: Resource pooling refers to the grouping together of resources (assets, equipment, personnel, effort, etc.) for the purposes of maximizing advantage and/or minimizing risk to the users. By pooling resources, which involves virtualization of typical IT stacks server, storage and networking (but also on the level of datacentre power and cooling), users benefit from lower individual investments, since resources are shared. Although shared infrastructures have huge benefits, potential issues on the shared components will impact all users of the shared environment. A thorough analysis of the infrastructure is recommended to identify potential single points of failure.
You may opt for "private" instances (private clouds) for specific needs or for specific reasons. In terms of resource pooling, bigger suppliers tend to have the benefit of being able to provide shared support services with round-the-clock service. Which do you prefer: round-the-clock access to a support service with potentially less expertise, or relying on a single support engineer who is *on duty* during off-peak hours?

Q: With what and how many other customers are we sharing the cloud computing resources, and what is the projected impact of these customers?

A: Remember that cloud computing is about sharing resources with a number of, in general, unknown outside parties. A thorough risk assessment is needed, and you should be informed about the measures and techniques used by the CSP to guarantee logical separation of data between the different tenants sharing the cloud-computing resources. This risk assessment should, of course, also take into account the nature of the data you intend to move into the cloud. You might want the CSP to provide a configuration that minimizes the sharing of confidential resources such as the disks containing your organisation's sensitive data.

Q: Are backups included in the price plan? Is there a "no data loss" guarantee?

A: Quite often backups are a billable option, but some CSPs offer a "no data loss" guarantee as a basic service. A "no data loss" guarantee is usually implemented using data replication between multiple datacentres and is to be considered as essential. Some CSPs offer "no data loss" guarantees without stating a specific Recovery Time Objective (RTO) and/or Recovery Point Objective (RPO). Make sure you fully understand the possible consequences of this sort of guarantee.

Q: Is it possible to define backup policies and to initiate restores?

A: How much control do you have for implementing a backup/restore policy adapted to your business needs? Make sure you have detailed information available about the backup and restore policies and procedures. Find out what possibilities you have concerning the testing of your backup and restore policies.

Q: How many datacentres does the CSP operate and what are their geographic locations?

A: For continuity and availability reasons, a CSP should operate at least two datacentres. It is also important to know where these datacentres are located, in order to assess their chances of surviving a catastrophe and to check compliance with data storage legislation relevant to your business.

Q: Does the CSP hold the required certificates and does it allow auditing by its customers and/or third party organisations?

A: You will want to do business with a CSP you can trust. Certification and auditing provides the foundation for trust. CSPs need to be able to show that they can live up to the promises they are making.

Certification and auditing can cover many control objectives:

- CSP organisation, planning, governance and risk management
- Documented policies and procedures
- Service change management
- Event management
- Logical security
- Change management
- Data integrity
- Physical and environmental security
- Service level agreements
- Client reporting, billing and satisfaction
- Financial health

Q: What channels are available for communicating with the CSP service organisation?

A: You may be interested in evaluating the way your CSP communicates with you. For critical events concerning your applications and data, you will probably want to communicate with your CSP via telephone using the local language. For less urgent matters email communication should be sufficient.

2.4 Rapid Elasticity

> "Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time."
> (Source: NIST)

Q: What are the upper limits on the expansion of computer resources such as CPU cores, RAM, IO bandwidth and storage?

A: It is important to find out what the maximum limits are in terms of computing resources. It would be unpleasant to find out after signing a contract that the CSP's environment cannot support your application because of one or more limitation on the resource level. It is equally important to find out how a resource reconfiguration will impact the active services. It might be necessary to reboot servers to make the expanded resource available to your services.

Q: Does the system scale resources automatically or do you need to rescale resources manually using, for example, a web application?

A: Is the scaling of resources automated or not? You should be able at least to set alerts on resource thresholds to allow you to monitor resource consumption and, if needed, to reconfigure the resources assigned to your services. Some systems will assign additional resources automatically when needed, but you still need to be informed. Also, make sure the scaling mechanism covers both the addition and removal of resources.

Q: Are there any resource reconfigurations that require a system reboot?

A: Most CSPs advertise scalability, but often they forget to mention if a server reboot is necessary after reconfiguring the assigned resources. Depending on the server's role, a reboot might be something you want to plan carefully if it is unavoidable.

Q: Does the CSP offer price protection?

A: Unfortunately, some of the bad licensing practices found in the on-premises enterprise software world have been carried over in the cloud-computing world. Thus it is important to check if the CSP allows for:

- Usage-level alignment (up or down) based on customer needs
- Application of, for example, monthly "rollover" usage to address peak usage
- Long-term price protection

2.5 Measured Service

> "Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service."
> (Source: NIST)

Q: What billing model is used to support the "pay-per-use" principle?

A: Some cloud-computing services are charged per instance-hour consumed, from the time an instance is launched until it is terminated or stopped. This billing model might not suit your business case, so you will have to find out if an alternative billing model is available to you. Some CSPs offer prepayment schemes linked to price reductions while others offer mixed "fixed + variable" billing models. It is up to you to decide what model best fits your requirements.

Q: What real-time resource monitoring tools are available?

A: As already mentioned, when discussing scalability it is rather important to be able to check resource consumption, preferably in real-time. Automatic scalability is an excellent cloud feature but proper systems management practice dictates that you should monitor your systems in an effort to detect and thus be able to react to anomalies such as an unusual increase of memory usage on your server. If you do not monitor your cloud infrastructure, your CSP bill at the end of the month might be an unpleasant surprise. Do not forget to check the availability of application level monitoring tools.

Q: How do you determine if a Service Level Agreement (SLA) is being met?

A: Most CSPs offer an SLA on the availability of the cloud service. Typically this availability would be between 99,95% and 99,99%. It is important to know how the effective availability of the cloud service is calculated. If calculated on a yearly basis, a 50-minute downtime event still meets the 99,99% availability SLA. If effective availability is calculated on a monthly basis, a 5-minute service interruption breaches that very same SLA. Make sure you understand the SLA details and the other general contract conditions. Negotiate if the standard SLAs do not fit your needs, and make sure your contract clearly states the penalties (if any) for not meeting the SLA.

3. ADDITIONAL QUESTIONS

The answers to the above questions will give you a good idea if you are dealing with a true CSP or instead with a traditional hosting provider. They can also help you to identify "cloud-washed" offerings, which are traditional solutions dressed up with a cloud marketing label.

In addition to the questions already mentioned, you need to ask the following questions to verify the suitability of the service proposed by the CSP.

Q: Does the CSP meet general and industry-specific compliance and security standards?

A: Since your organisation remains accountable to regulators, business partners, customers and employees, you should not consider using a particular

CSP unless it meets the compliance and security standards applicable to your industry. Remember that outsourcing an application and its data to a CSP does not include a transfer of the responsibilities that go with it. Pay particular attention to the respect of the applicable privacy rules and regulations. You may want to check compliance with the EU Data Protection Directive (European Commission Data Protection web site) and the EU Proposed Directive on Network and Information Security (Frequently Asked Questions on the Proposed Directive on Network and Information Security).

Q: Does the CSP offer a service catalogue with predefined service templates?

A: Service catalogues list standard products and predefined cloud services. These nice-to-have catalogues give you flexibility and speed up deployment, saving valuable time for an often small, additional cost.

Q: What is the minimum contract term?

A: Although cloud services are sold based on "pay-per-use", most CSPs offer contracts with a minimum term of 12 to 36 months. In general longer-term contracts tend to be somewhat cheaper than short-term contracts. Just make sure to check that the contract term meets your expectations.

Q: What is the contract renewal policy?

A: Do not assume your contract will end automatically. Long-term cloud contracts usually include an automatic renewal policy, unless you terminate your contract as specified in the contract.

Q: What application and management APIs are made available?

A: APIs are essential if you want to integrate cloud services with other services and applications. Some CSPs offer them, others don't. Make sure to check the details concerning the availability and the maintenance of the APIs. Take care to check the availability of management APIs if you want to integrate the management of the cloud services within your existing management framework.

Q: What happens to my data when a provider goes out of business or when my contract ends?

A: In general, CSPs should have documentation on how to migrate your data into their cloud and will help you to do so. Unfortunately, few CSPs document how to get your data out of their cloud. In general, the data export process will not be very different from the data import process but it is better to check. In some cases, you might have to go through a complex and costly migration to reconvert your data before importing it into a new processing environment. Regardless of the operational procedure, the CSP is legally bound to give you back your data. Make sure there is no doubt over who owns the data. The worst case scenario is, of course, that your CSP goes bankrupt. In that case, you will have to negotiate with the CSP's trustee/curator responsible for the liquidation of the remaining CSP assets.

Q: Is it possible to migrate my cloud resources to another CSP or environment and are there additional expenses associated with this?

A: This particular question deals with moving full environments including virtual servers, networking and such. How easy is it to migrate a full environment to an alternative CSP or to a local infrastructure? In certain scenarios additional expenses will be invoiced because additional bandwidth and/or support is needed to support the move. Some contracts also specify the applicability of a cancellation fee.

4. CONCLUSION

Cloud is not so much about a new technology as it is a new business model needing a different approach from all parties involved, since it impacts the complete chain linking the manufacturers, vendors, distributors, integrators and customers. All parties involved are facing a transition period and some will be faster and more successful in adapting than others.

Solutions are becoming less transparent, as the focus is now more on performance than features. Consequently there is a shift of responsibility from the end user to the supplier. A CSP that wants to flourish in cloud computing will have to be seen as a trustworthy partner by his customers and prospects. CSP customers are relying more and more heavily on their CSP, and security issues are still seen by many companies and organisations as the major hurdle to overcome in moving services to the cloud. This principle of trust is applicable to the complete chain leading from the hardware/software manufacturer to the CSP customer.

This BELTUG guide is a first assistance for identifying true cloud offerings. Each of the topics mentioned is documented extensively in many specialised publications, but in this document we deliberately decided not to go into too much detail, which could easily overwhelm a cloud novice. It should also help you to choose and trust the proper offering by asking the right questions. After all, information and communication are key in avoiding disappointment.