2- Illusions

We introduce in this chapter common approaches which many companies believe to be sufficient in order to protect their critical data and applications from discontinuity and permanent loss. We will show how each approach provides the illusion of resilience yet does not provide a valid disaster recovery plan.

Illusion 0: Laissez Faire

Cloud Computing providers are supposed by contract to handle backups, service continuity and disaster recovery. Some companies believes that this contractual obligation is sufficient.

However, the notion of "force majeure" defined in most contracts does not cover the Disaster Case No 4 and would probably not cover any of the Disaster Case triggered by a large scale power outage or natural disaster. Moreover, any Cloud provider might fail to its contractual obligations and, depending on the contract, would not be forced to compensate much for this failure.

For example, we are aware of some cases of users who experienced unexpected data losses with a large Web email provider. Since there was no backup and no access to access logs, there was no way to prove that data was lost. After this experience, they decided to implement backup of their emails and no longer trust the reliability of Cloud providers.

Laissez Faire is thus not an option.

Illusion 1: VM Backup

Backup of Virtual Machine (VM) images from one Infrastructure as a Service (IaaS) provider to another is becoming increasingly popular. As long as the content of the image is consistent, such backup can be useful to migrate from one Infrastructure as a Service (IaaS) to another one and recover from the unavailability of the first Infrastructure as a Service (IaaS) provider.

However backup of virtual images does not provide a solution to most disaster cases of Chapter 1. In particular, backup of an inconsistent Virtual Machine (VM) image is an inconsistent backup. It can not be trusted for disaster recovery. Virtual Machine (VM) images can become inconsistent because of storage bugs or because of application bugs or because of system corruption.

The case of system corruption can happen for example whenever a virtual image has been initially prepared with an operating system which is long to install and is not very reliable. If during the installation procedure, the host operating system has to be restarted for some reason, the guest operating system may not be able to flush a few storage blocks. The resulting image may contain incorrect blocks either related to the application or the operating system. The consequence of those incorrect block can be immediate or may need to wait months before triggering a permanent bug which prevent the application to restart.

This example demonstrate the limitations of VM Backups for disaster recovery and why it is not sufficient.

Illusion 2: SAN Replication

Many private Clouds rely on a combination of server virtualization and storage virtualization. Storage virtualization is often implemented by disk bays also known as Storage Area Network. Many companies believe that disk bays are a perfect tool capable through snapshots and replication on multiple sites to protect data from any disaster. The use of replicated disk bays creates an atmosphere of confidence which leads IT operation teams to put less effort in monitoring daily backups.

We have seen in Chapter 1 that intentional destruction, replication bugs, storage bugs and malware can affect disk bays. Many companies already lost critical data due to the illusion that disk bays are sufficient to prevent any data loss. We even know one case in which the snapshot of a disk bay took so much space that it prevented the synchronization of storage blocks between the main memory of virtual servers and the physical disks. Many hours of production data were lost due to the use of snapshots for backup.

No disaster recovery plan should ever be based on any form of trust of the replication technology of disk bays.

Illusion 3: Untested Data Backup

Most companies implement automated backup of their critical applications. They use very sophisticated backup tools with beautiful user interfaces and expensive price, which gives the illusion reliability and resilience.

A daily job copies data to a dedicate storage area on a disk bay, to a removable disk or to a tape. Once the automated backup has been setup, it is very often forgotten until the day an incident happens. And it is only once the incident happens that the IT team discovers that the backup job was not launched for a while or that the data which is present in the backup is not sufficient to recover the application.

Overall, no matter the backup tool, doing data backup without testing the possibility to reconstruct the whole system from the data backup and from the application original files on a fresh operating system provides no guaranty for disaster recovery.

Illusion 4: Centralized Vaults

In order to pass certifications imposed by government authorities, some companies purchase a "Vaulting Services" to companies specialized in disaster recovery. Those vaults include servers, disk bays and everything needed to rebuild the IT infrastructure of a company in case of destruction of its data centers, for example by natural disaster.

Once the disaster recovery vault is purchased and paid, the certification is obtained by claiming that replication between storage area networks is set. It is then forgotten, not tested and not monitored.

The high price tag of disaster recovery vaults, the long list of video monitoring features and an impressive list of customer references creates the illusion of a trustable disaster recovery plan. The absence of regular tests of disaster recovery combined with storage replication technologies with significant risks of bug means that in most cases, disaster recovery will fail.

Reliance on a single, famous, global provider is the consequence of typical purchasing policies which try to reduce costs by aggregating requirements in a single tender. A big company also gives the illusion of trust and resilience. However, reliance on a single provider creates many risks.

The case of wikileaks demonstrates the consequences of relying on a payment system under the authority of a single country: once wikileaks was under scrutiny of US authorities, no payment though Visa, Mastercard or SWIFT was accepted any longer, even in Europe. SWIFT European representation had to implement restrictions defined by US government.

Same phenomenon can happen with Cloud. A company suspected by US justice to detain certain information on the Cloud or to participate to illegal activities might see its servers or virtual machines seized, just like it happened to Megaupload. The Wikileaks case shows here that it was quite useful to use multiple providers subject to different legislations to implement service continuity. Various small hosting providers in Europe which were not subject to US legislations directly or indirectly ensured business continuity.

Cases which involve seizure can happen more often than it could seem: suspicion of intellectual property theft, suspicion of corruption, suspicion of money laundering, suspicion of illegal government subsidies, suspicion of illegal export, suspicion of anti-competitive practice, etc. Any large company with activities in many countries and multiple product lines could fall under any of the above suspicions, which are often triggered by a competitor. It would thus be very unwise to rely on a single supplier or on suppliers under the control of a single legislation.

Relying on a single provider also poses risks in relation with social and political instability, which could happen up once in a century. The risk is low, but high enough to be taken into account by selecting multiple providers in multiple countries. In particular, a single supplier is more subject to government pressures or blackmail than multiple providers.