1- Cloud Disasters

We introduce in this chapter scenarii which can lead the the interruption of service of public or private clouds, some of which lead to the permanent loss of data. We include in public clouds both infrastructure as a service (ex. Amazon EC2), software as a service (ex. Salesforce), platform as a service providers (ex. SlapOS). We mainly include in private clouds server virtualization (ex. VMWare) and storage virtualization (ex. EMC).

Unless there was public information available, we did not provide detailed references to the cases bellow in order to protect the identity of the companies or countries which have experienced a disaster. However, each case was certified by a group of 2 independent scholar and 2 industrials.

Disaster Case 1: international simultaneous destruction

We are aware of at least two cases in which both data centers of a government were intentionally destroyed the same day, as well as all tape backups. The technique used for destruction combined fire and robbery. The consequences of this destruction was the loss of critical government database which are required to exercise sovereignty.

Although this type of incident is uncommon, it happens from time to time in countries subject to civil war. It could also be the consequence of terrorism or political instability.

The case of terrorism should be considered in particular for data centers with a certification, such as payment card industry compliant (PCI) data centers which are not that many in the world. A targeted attack on such data centers based for example on the use of graphite bombs in the air conditioning system could be planned with few financial resources and have an economic impact much stronger than September 11 attacks.

Political trouble with riots and sabotage are a risk to consider with the growth of unemployment rate. Countries with more than 15% employment rate tend to experience mass protests and at very rare occasions, violence against data centers could happen in a way similar to the destruction of train infrastructure which happened in the United States in the 19th century.

Disaster Case 2: replication bug

Most Cloud infrastructure include a data replication system of some kind: synchronous, asynchronous, real-time, scheduled, etc. Data replication is used to migrate data from one data center to another, or from one server to another, either to optimize performance or to implement high availability in case of failure of one data center due - for example - to natural disaster. As any software, the data replication system has bugs [*between 20 to 40 defects per KLOC*, Watts S. Humphrey, Software Engineering Institute, Bugs or Defects?] with a probability of occurrence. Bugs of the data replication system can thus sometimes lead to service discontinuity or data losses each time Cloud service is migrated from one data center to another or from one server to another.

In 2011, Amazon EC2 experienced a 4 days service outage. This outage happened after one of Amazon data centers became unavailable due to the scale of the power disruption, a large number of EBS servers lost power. Data replication started at once between that data center and all other data centers. This sudden start of replication between all data centers caused bandwidth congestion which eventually lead to 4 days of unavailability, without data losses.

Replication bugs are sometimes caused by human errors rather than software errors. A financial institution used a tool to synchronize two disk bays during the upgrade of its data center. The options passed to the synchronization software did not include consistency checks. For some reason, the synchronization tool had to be launched multiple times in order to copy data. However, not all data was flushed out to disks between each launch. Some block were thus not copied from one device to the other. Critical financial transactions were thus lost. A government agency experienced even worse incident when it started the replication between two disk bays in both ways at the same time. Suddenly, filesystems on both disk bays became corrupted and impossible to mount on servers. The government agency then searched for backups made by the automated backup tool which was installed on disk bays and found that no backup had been made for a couple of months due to some exception. Critical government accounting data was permanently lost.

Disaster Case 3: storage bug

Disk bays which are also known as Storage Area Network (SAN) are an essential component of many Infrastructure as a Service (IaaS) services. They consist of grouping hundred or thousands of hard disks within a single device with advanced software to implement synchronization, redundancy and caching between disks. By grouping disks and using fast networks to interconnect them, it is possible to provide high performance, easy to manage and supposedly resilient storage to servers. However, if the advanced software contains bugs, the use of disk bays can generate disasters.

The same risk applies actually to any complex software based storage technology, including distributed block storage, distributed databases and distributed filesystems. Implementing a reliable distributed storage requires at least 10 years of research and development focused on scalability and quality assurance. The nature of research and development for distributed storage is so complex that it can only be conducted by very few individuals in the world. For example, the

development of Oracle BTRFS filesystem heavily relies on single individual, Chris Mason, who has left Oracle for Fusion-IO in june 2012.

It is thus not surprising that, even with the best people and the best technologies, distributed storage itself experiences bugs which lead to permanent data loss. In August 2011, Amazon suffered an incident experienced by an error in the EBS software that cleans up unused [EBS] snapshots that incorrectly thought some of the blocks were no longer being used and deleted them. This incident caused permanent data loss to many European customers. Another incident which caused thousands of company to lose their data permanently happened in Japan during summer 2012. Firtserver, a company of Yahoo Japan group, upgraded its disk bays with a software patch required to fix security bug. Both bays were upgraded at the same time. The software patch which was supposed to fix a security bug was itself buggy. The upgraded software suddenly started to erase data contained on the disk bays, causing permanent losses to thousands of Japanese companies.

Disaster Case 4: malware deployment

In July 2012, all Drupal based web sites had to be upgraded suddenly after the discovery of a bug which cause security leaks. This bug cause numerous trouble to numerous governments including leaks of sensitive information and service discontinuity during the upgrade. However it did not cause significant data loss.

Let us imagine now two scenarii with worse consequences.

The development of cybercrime in modern societies has lead to the development of cyberpolice technologies and business, that is of software which helps police and governments to take control remotely of desktop PCs, servers and application servers. A few companies in the world develop software which are sold to most police departments worldwide under the term "legal enforcement" or "IT intrusion". Some of those software can be used to take control and erase all data on a phone, on a desktop PC or on a Linux server. Under the wrong hands, such software could used to alter the configuration of application servers which power a multitenant Platform as a Service (PaaS) providers, causing data losses to all tenants of the Platform as a Service (PaaS).

This first example demonstrates the risk of commonality posed by multitenant Cloud platforms. Single deployment of malware can cause damage to thousands of businesses. This risk of commonality also exists whenever large number of developers rely on a the same binary distribution of a software.

Let us imagine for example that developers of a Cloud Computing platform working in Central Europe for a US based operating system vendor are in charge of the block storage subsystem. A certified version of the block storage subsystem is compiled, packaged and published daily on the internal server of the software publisher. A cybercrime organization could tap the connection between the Central Europe team and the distribution server located United States. After a few weeks, binary package signatures could be cracked or stolen. Then, a forged version of the of binary package which includes a timebomb could be signed and uploaded transparently to the internal servers. A binary package which includes a timebomb would then be present in all so-called certified versions of the Cloud Computing platform distributed worldwide with no possibility for third parties to recompile it from clean sources. The day the timebomb expires, all block storage subsystems start erasing all data.

The scenario we have just described is realistic in the sense that the most technologies to implement it are available and the profits which a criminal organization could make from a timebomb bribery are huge. The risk of occurrence can thus not be neglected.

Disaster Case 5: large scale electrical outage

In 2011, Azure service experienced a service outage of 48 hours due tothunderstorms which affected electrical power supply. In theory, such power outage should not affect the availability of Cloud services since Cloud providers can rely on data centers locate on multiple continents. However, certain customers require their data to be located in certain countries for legal reasons. This is sometimes the case of governments, of military industry, health industry, etc. Most Cloud providers only have a few data center per country or continent, which creates a significant risk for some of their customers.

Business which require their data to be located in a single country should thus take into account in their disaster recovery plan the probability of regional or national power outage. This type of outage has been experienced in 2012 in the USA at New-York and Ashburn, in Germany, in Japan and in France in 2011. Batteries and power generators of Data Centers could be insufficient to power the data centers more than a day or two. After that time, data and applications can not be migrated to another continent due to strict policies in terms of geographic location.

Case 6: natural disaster

In 2012, Sandy storm lead to the outage of the Web site of Gawker, Gizmodo, Buzzfeed and Mediatewhich was hosted in the data center of East Coast. The use of Cloud services based on data centers located in multiple regions or continents should protect from this kind of event, besides bugs related to the replication system.

However, natural disasters and fire in particular are still a cause of permanent data loss or service discontinuity for businesses which rely on a private Cloud infrastructure hosted in a single data center. This is probably still the case of many

companies and government worldwide.