Disaster Recovery in the Cloud Era… Part 1

Disaster Recovery… it's one of those words that conjures up all sorts of emotional responses; from "we don't need that" through, "I want it, but cannot afford it", to "how could I afford NOT to have it"…

Whatever the response, the vast majority of companies would understand and appreciate the value of being able to recover services should a disaster strike. It is the degree of financial value placed on this ability to recover that determines the willingness to invest in technology that offers business and/or IT continuity.

In a survey completed by Forrester research, enterprises were asked whether they actually have a disaster recovery plan in place. Interestingly, a large proportion of respondents have no formal recovery plan in place.
What is interesting with this statistic is that whilst many businesses talk about the impact to them of a DR event, very few actually take any preventative action. This statistic coupled with another statistic; where one in four companies that suffer a DR event will never reopen their doors; means many businesses are taking an interesting gamble.

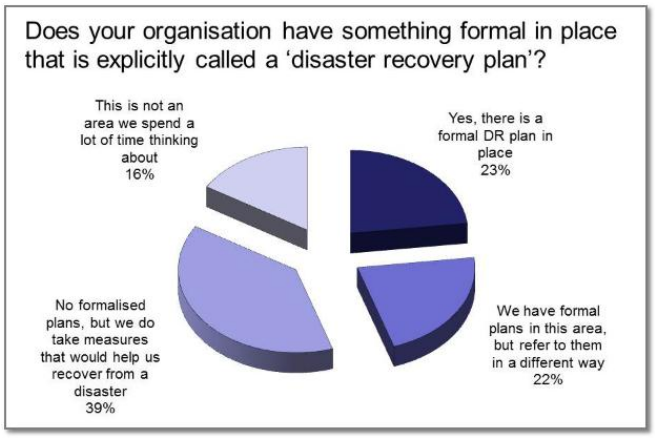

Before we get further into this article, first of all we need to understand the definition of the word "disaster" in a business context; this can simply be described as any non-planned event that disrupts the business' ability to complete its primary function; so when we talk about disaster recovery, that would imply the ability to restore (recover) the business' ability to perform its primary function.

The ability to restore a business' core function encompasses much much more than simply having another copy of your data somewhere else; in fact, disaster recovery is much larger than IT… A true and comprehensive DR plan (also known as a Business Continuity Plan; BCP) needs to encompass the people, their access, the facilities, any plant or equipment, and IT. Only by having a BCP plan will the business have complete assurance of being able to recover from a disaster.

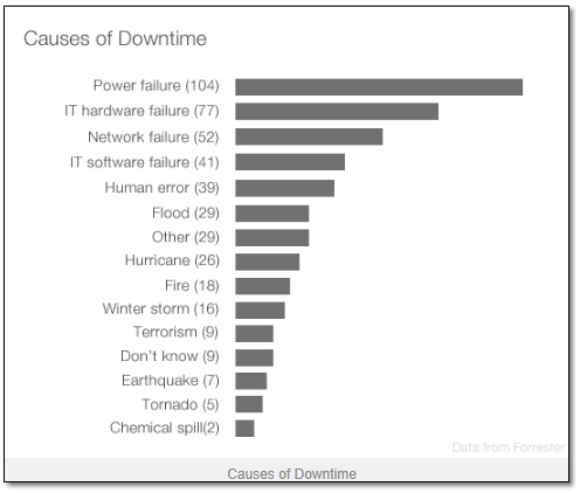

Whilst IT can, and should, assist with the formation of a BCP plan, the primary focus of IT would be to develop and maintain an accurate IT/CP (IT Continuity Plan). The scope of an IT/CP is to provide assurance to the business that any critical IT services required to deliver business functions would be available should disaster strike. In the same survey completed by Forrester, causes of existing downtime events were captured, and dispelling many myths, 68% of events were directly caused by IT technology and only 32% by natural disasters/outside influences. Given the vast majority of issues are caused by IT; IT has a responsibility to take preventative actions to limit the business impact of IT failures.

Bringing all this DR talk into context; just how likely is a DR event to occur, and is it simply "scare-mongering" from people trying to sell DR solutions? Well, once again the Forrester survey provides us with that information, with 73% of respondents reporting that in the last 5 years they had never called a DR event; but inversely, 27% had actually need to rely on their DR process. So, is 27% of customers considered scare-mongering? That is up to each business to make a calculated risk decision upon; however given the impact of NOT having a DR plan in place, it is likely that should a cost effective DR solution be available, that almost every business would seriously consider adopting it.

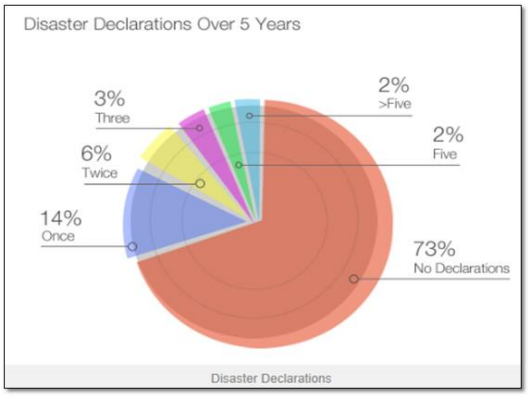

Based on local customer feedback, the main reason why organisations do not have a DR plan in place is purely related to the capital cost of implementing and maintaining DR. Traditional DR solutions are complex to set up. They require a secondary site, expensive WAN circuits, dedicated infrastructure, and hardware-based replication to move data to the secondary site. As highlighted previously, all of this investment is only to cover the relatively "small" chance of a DR event actually occurring. As a result, DR is relegated to being something that only the largest of organisations can afford, and then even within them; DR is only offered to the most critical of IT resources, and very rarely is the entire production landscape protected. Even in highly regulated industries like Banking and Telecommunications, it is common to see subsets of production servers protected with DR.

But disaster recovery is a natural fit for the evolving cloud computing model, and public cloud service providers such as IndonesianCloud have DR-as-a-Service offerings that aim to make disaster recovery broadly accessible for all companies by providing cost-efficient, automated and simple disaster protection.

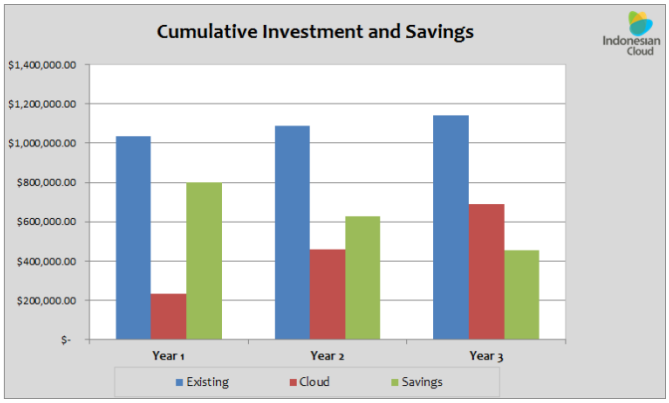

Cloud service providers have the ability to "share" expensive infrastructure across many customers, and by virtue of this sharing, enables the cost of standby DR infrastructure to be divided across a much wider pool of customers. This all results in reliable and comprehensive Cloud DR solutions for a price far below that of traditional self-built / self-hosted models. Generally, the cost over 3 years would be near 50% of the legacy model. It is not unusual to see companies with 5% of their production servers protected using legacy self-built models being able to leverage Cloud DR to increase their DR coverage to 100% for the same costs as they were previously spending for 5% coverage over a 3 year period. This can only be good for businesses as 100% protection is absolute assurance of recoverability in the event of a disaster.

Disaster recovery solutions, when delivered from the cloud, change the dynamic of how the service is both funded, and delivered to the business. Cloud based solutions are subscription based, which means 100% operational expense (OpEx), and generally they come with few (if any) long term contractual obligations. Cloud solutions also remove the need for expensive dedicated WAN connections, in favour of virtual private networks (VPN) delivered over standard internet, and already provide service from a remote datacentre, removing the need to either build your own, or undertake physical collocation/hosting contracts . Cloud based DR can also be setup and operational in days, not weeks, and due to the subscription model, does not require substantial up-front investment.

One of the biggest positive changes available by moving to Cloud based DR is the ready access to complete DR test simulations; it is quite common to hear of enterprises that complete DR exercises

annually, and generally, even that DR test is not an accurate representation of a true DR failure. The problem with annual DR tests is that infrastructure and systems change, and they change quickly. An annual test might miss critical infrastructure components that may have been introduced, removed, or simply updated. For recovery assurance, a more regular DR testing regime is recommended, at least quarterly. Using traditional DR, not only is the build and maintenance expensive, but generally so is the recovery testing; tradition DR is very manual, it requires trained people to perform the recovery, and this recovery is normally completed out of business hours. The vast majority of companies executing DR tests organise to have vendors on "standby" and mandate that all IT staff are present to help deal with issues that might arise. In a real disaster, those same vendors and trained IT staff may not be available; they might be sick, or just otherwise unavailable. With Cloud DR, the infrastructure platform is available for testing at any time, it resides on an isolated network, and generally is underpinned by some sort of virtualization technology; this all combines to allow the Cloud provider to offer a semi-automated, and non-intrusive DR test/recovery process, with common tasks completed by the DR system rather than people. This provides two benefits; first, it allows the business to recover without reliance on anyone else, and second, it reduces the time and therefore cost to recover (or test recovery).

It is only by having a cost effective, automated, and reliable DR recovery process does the business have complete confidence and assurance from IT that IT/CP process has been developed, and the required preventative measures taken to ensure uptime.

In part two of this three part series, we will discuss data protection and networking requirements for seamless and transparent DR; and in part three we will discuss DR Service activation and re-protection (failback).

For more information, contact the author: Neil.Cresswell@IndonesianCloud.com