

BCAP	July/August 2010	Circulation: 12,120	Page: 72
------	------------------	---------------------	----------



Business continuity

Dealing with DISASTER

Volcanoes make headlines but business continuity is not just about dealing with natural disasters. It's about sweating the small stuff

BY TIM MENDHAM

Pandemics? Volcanoes? Earthquakes? Tsunamis? Are these the stuff of business continuity? Gartner has issued several papers covering major disasters such as the Iceland volcano eruption and its impact on business travel, admitting that “few, if any, businesses plan for a volcanic ash disruption scenario”, which is probably the understatement of the year.

In contrast, Stephen Hopkins, head of security practice for BT global services, says that most disasters are small scale — “death by a thousand cuts” — rather than a huge catastrophic event.

The most common threats are localised, he says, such as losing telecommunications or losing power. But many people underestimate the window of disruption. It could be one hour, it could be one week, or even months. For many people, their recovery vision is too short term, he says. And the vision might also not take into account all the implications of business continuity, which go well beyond getting the email system back up and running to include business, legal and regulatory requirements. Then there are brand management issues that can impact an organisation well beyond the extent of the disaster itself.

And to a certain extent, this sense of limited vision is borne out by Gartner predictions on strategic

planning issued at the end of last year. According to the organisation's business continuity management (BCM) research community:

- By the end of 2012, fewer than 10 per cent of enterprises will have received external certification for their BCM and IT disaster recovery programs; those that do are either regulated to do so or will be mandated to do so by their supply chain partners.
 - By 2014, only 50 per cent of in-house transaction document printing and mailing operations will have alternative BCM providers in place.
 - By the end of 2015, 15 per cent of global 2000 enterprises will have transformed their BCM programs into a cross-enterprise business operations strategy and planning function.
 - By the end of 2015, one-third of work-at-home programs will be verified and ‘tuned’ for BCM readiness when teleworkers log in for their work shifts.
- At the same time as the lack of forward-looking vision are the requirements for increasingly short response time objectives (RTO). Further, Gartner research shows that RTOs are shrinking: In a 2010 report, 63 per cent of survey respondents said the RTOs for their mission-critical business processes were less than 24 hours.

Nobody knows when something will go wrong, so risk mitigation is what BCP is all about

“With such short RTOs,” the report says, “it is imperative that BCM plans are current and easily available during a crisis.” The problem is that, in some circumstances, 24 hours is not a short RTO, it’s a disaster in its own right. In some instances, a response within minutes, if not seconds, is a concrete requirement. What do you do then when the power goes off — or the volcano erupts?

Business continuity

Scott Henderson, CIO with mining software and consulting company Runge, says disaster recovery is generally not well handled. “There has been heightened awareness over the last few years ... but many non-IT people only have a cursory understanding of the issues.

“It’s more than just a consultant’s analysis and, while IT is a central component in business continuity planning [BCP], particularly as we are all so reliant on IT as a corporate asset, it is not just an IT problem. Nobody knows when something will go wrong, so risk mitigation is what BCP is all about.”

And mitigating risk is what Henderson’s company did.

Operating since 1977, Runge moved into a new head office building in Brisbane two years ago with a new data centre. “We looked at our core systems, and asked ‘what’s the minimum amount of time we could do without a particular operation?’ We committed to less than one hour downtime for everything.”

Email and finance were the most critical operations, “anything else we could probably do without for 24 hours, but we still stick to the one hour policy”.

The company now has 15 servers in 14 locations around the world, with a replicated disaster recovery site in Sydney and global latency of under one second. “We use multiple layers — no BCP is fail-safe. We have local ▶

The law, the brand and security



Business continuity plans are not a business differentiator, says Stephen Hopkins, head of security practice at BT Global Services, Asia Pacific. “Everyone has to do it.”

But in some cases it goes further than a good and sensible thing to do; it is a requirement, a compliance issue that you are legally, contractually or morally obliged to perform.

“For some organisations, like essential and financial services, there could be legal and regulatory requirements on top of best practice.”

The legal obligations also extend to the supply chain and customers. Partners and associates both upstream and down expect you to be available to do business, and do it effectively and ethically. “Just because you’re going through a disaster, you can’t start exposing customer records,” Hopkins says.

Dean Redman, Australia New Zealand country manager for IT security hardware and service supplier Soricwall, points out that there are also brand protection issues. “You need to ask: how much do you protect yourself? What does it do to my brand and company reputation if I have a disaster, or my plans are revealed? Be careful what you reveal about your organisation and your plans, and who to. Social media raise many issues of security and revealing too much about your business continuity planning. Likewise, be aware of phishing to gain background information. All of this can lead to denials of service or worse, cyber terrorism or blackmail.”

back-up archiving for desktops, laptops et cetera — it all goes to the servers. We use virtual servers for the entire operation. We do quarterly tests on a round robin basis by taking down one server at a time. We have a rooftop generator. There is an expectation that IT will be up 100 per cent. We don't leave it to chance.

"We went with Double-Take Availability and Backup as a one stop shop ... and also one throat to choke approach. It works with a lot of different applications, and that's worked for us; we've achieved 99.6 per cent up-time generally, and 99.8 per cent for services."

He adds there is a borderline of being so preventative that it almost tips over into being an efficiency issue. "You need to take a balance of how much to prevent and how much to recover."

We look for
simplification and
end-user experience as
the ultimate objective

Service providers

Hopkins says that there is an important decision to make in this process — how to run the disaster recovery and business continuity facilities more effectively; can you use or partner with service providers?

Andrew Pritchett, CIO of intellectual property law firm Griffith Hack, agrees that there are benefits in finding the right technology partner, particularly in a specific area of the process. They need to fit well with the organisation's objectives and procedures, however. Sometimes that technology solution can be a problem in its own right.

"We had undertaken analysis to determine whether to outsource or insource. We'd selected a supplier who simply offered rack space and rented equipment for use in a disaster. As it turned out, the restore time objective was under-estimated, and the restore point objective became days instead of hours or minutes, and then the recovered infrastructure still didn't work as required," says Pritchett.

There were some contractual and technical limitations including a suite of problems with the solution, both technical and commercial. For example, the IP firm had small WAN links that were under contract and it was considered cost prohibitive to change the speed.

"We made a concerted effort to rectify the in-place solution, to simplify and reduce cost. We installed on-site rack and equipment, and used Riverbed WAN optimisation as the backbone for the system for data compression and

de-duplication. We are still in the process of making our DR site live, but we've already significantly cut down on recovery point objective and RTO — the latter has been cut down to a tenth of previous performance.

"We look for simplification and end-user experience as the ultimate objective in everything we do. It is imperative that we achieve this."

The new system, working on the Riverbed backbone, has helped in various ways, he says. "We now have five sites, including the DR. We've been able to decommission 30 servers nationally and, as a result, we've achieved ROI on the technology in our first year."

Pritchett says there's also been an improved end-user experience overall and this has taken pressure off the IT team. "We've radically improved performance for log-in times, file sharing, and web surfing. We've now reduced back-up costs, tape costs, operator costs, and the number of servers. Interstate offices' internet use has gone through the roof. And, of course, the partners and staff around the country, who previously tolerated a very slow system — virtually dial-up — have now adjusted to the new system and consider the current performance as the 'norm'." And, though Pritchett doesn't say it, once the staff get used to a new operating environment, they might be wanting even more; such being the downside of successfully improving performance.

Of course, there is financial and technical commitment for a system that, hopefully, you'll never use. And you need to convince management of the need to make this commitment.

While Henderson admits there's an issue with pitching the need for apparently 'excess' technology, he doesn't think it's that hard to convince people of the need for disaster recovery systems. "There's a tangible outcome and you have to present that." But once it's in place, he says there is a tendency for some to want to use the facility while it's not required. "Sometimes people misunderstand. They ask: 'Can we use these back-up systems for something else while there's no disaster?' And the answer is no!"

CIOs have to stress the reason for this, and that is the insurance approach: The old adage of an ounce of prevention is worth a pound of cure; the price of making good after a disaster is a lot more than the cost of buying a system in the first place.

"There is a difference between cost and price," Henderson says.

Taking that one step further, Pritchett maintains disaster recovery needn't be a cost centre. "If you do it right, you can get performance benefits, improved facilities and services, and cost savings." Not to mention a little peace of mind.

And next time a volcano erupts, hopefully you will have the plan in place to deal with it. **CIO**

Tips for business continuity

10 best practices for creating and maintaining effective business continuity management plans

Gartner, February 2010

- 1 Executive management commitment is required for the business continuity management (BCM) program
- 2 Business units must develop their own BCM plans
- 3 BCM plans must follow a standard process and formality and not be done on an ad hoc basis
- 4 BCM plans must be developed to cover a longer outage time frame
- 5 BCM plans must be regularly exercised
- 6 Develop a structured framework of BCM plans
- 7 Keep BCM plans relevant to their purpose
- 8 Provide relevant information in BCM plans to facilitate recovery within defined recovery time frames
- 9 Establish a central repository and administration process for BCM plan maintenance
- 10 Use automation to mature BCM plan management

Disaster recovery checkpoints

Scott Henderson, CIO, Runge

- Use multiple layers of protection
- Understand how people use IT
- Business continuity is not just a process, procedure or structural thing — it involves people
- Don't put your business continuity plan on the shelf. You need active management rather than an expensive paper weight
- What was right five years ago is not necessarily OK today — things change
- The IT department needs to be easy to do business with
- Reference all areas, physical and software — it might not just be a system crash, it could be a building collapse — be aware of security in all circumstances

Mark Deguara, Emerson Network Power

- Is your disaster recovery location susceptible to the same issue that has caused your requirement for it?
- Has your disaster recovery location got the capability of supporting your IT requirements for an extend period of time?
- Do you test your disaster recovery procedure regularly — both in terms of the critical infrastructure and from an IT perspective?
- Ensure your staff are familiar with the disaster recovery procedures and, if it is an offsite location, make sure they know how to get there and what equipment is there
- If you have best practices in your main facilities, ensure you implement best practices in your disaster recovery facility. If and when it is needed, it better work (eg cold aisle/hot aisle configuration, redundancy as required, battery autonomy to support the load, precision air conditioning, and so on)
- Monitor, monitor, monitor — as you need to be able to know the status of the DR site as much as you need to know the status of your main site
- Consider managed power rails so critical items can be powered up remotely as well as being monitored