8. Contingency Planning and Management

Learning objectives of this chapter: Business disruptions, business continuity planning (BCP) and business continuity management (BCM), crisis management, disaster recovery planning and management, incident handling, organisational and performance considerations.

Businesses and governments across the globe depend heavily on Information Technology to manage various processes and deliver services. It is thus paramount that the services relying on the underlying IT systems continue to operate or are minimally disrupted even in the event of major disruptions to the underlying IT and physical infrastructure -for instance, triggered by a major security incident or a natural disaster, and so on. This may be achieved through a gamut of mechanisms, applied individually or in conjunction - for instance, switching the operations to an alternate site, or having backups to recover necessary data and information to resume services, have mechanisms in place to carry out some vital services manually while the IT infrastructure is not operational and so on. Contingency planning and management is then essentially coming up with and deploying the plans, processes, and technical solutions to effectuate such resilience, maintaining them over time and operationalising them whenever needed.

Contingency planning: What is it? Why do it?

For most part of this course, we have delved into aspects of operational security — identifying information assets, classifying the nature of sensitivities involved to determine corresponding security objectives — and accordingly determining the necessary controls to be put in place.

Contingency plan complements operational security mechanisms, and is devised for outcomes and circumstances beyond the expected norm. The ultimate objective being business continuity, by ensuring in the short term that core/critical operations of the business continue (or can be resumed within a very short span of time), notwithstanding if the unexpected events render useless the underlying services that these core services rely on, and in the longer run, to be able to restore all operations of the business back to normal.

Contingency planning in fact permeates beyond IT security, and is necessary for other aspects for an organization's operations as well. For instance, a government ideally needs to be able to operate even in case of a nuclear attack or other catastrophic event. In the context of an organization's reliance on IT infrastructure, contingency planning is necessitated by a multitude of factors. Several of these are in fact valid reasons to have contingency plan for non-IT aspects of an organization as well.

Trivia: In the United States, a designated survivor (or designated successor) is a member of the United States Cabinet who is appointed to be at a physically distant, secure, and undisclosed location when the President and the country's other top leaders (e.g., Vice President and Cabinet members) are gathered at a single location, such as during State of the Union addresses and presidential inaugurations. This is intended to maintain continuity of government in the event of a catastrophic occurrence which kills many officials in the presidential line of succession. Were such an event to occur, killing both the President and Vice President, the surviving official highest in the line, possibly the designated survivor, would become the Acting President of the United States under the Presidential Succession Act. (excerpt from Wikipedia)

Foremost, as has been iterated on multiple occasions, there is no perfectly secure system. Particularly, failures in complex systems is an eventuality, irrespective of how much so ever care is taken in operating and testing the system, or the degree of redundancy built into the system. Thus, there is a need to plan for being able to operate vital and critical services notwithstanding the failure of the underlying system that enables the functioning on a day to day basis.

Market and economic forces, such as the emphasis on getting a product's basic functionality and usability right, and getting the product to market on time (A 'Ship it on Tuesday and get it right by version 3' attitude) also back-burns security considerations.

Unexpected circumstances also emerge as a consequence of a series of unfortunate events. For instance, a Tsunami leading to equipment failures and eventual nuclear meltdown as happened in Fukushima, Japan in March 2011.

In principle, a robust contingency plan thus needs to be effective notwithstanding the nature of event and consequent outage of the normal operations. In practice, one's contingency plans are often based on what one imagines as the worst case scenario, and thus, the plans need to be revisited and redrawn with one's changing understanding or perception of worst case scenario. Case in point is an interesting observation from Hong Kong Monetary Authority (HKMA)'s [1] Business Continuity Planning manual. "Efforts put into, and experiences gained from the preparation of, the Y2K contingency plan were obviously not sufficient to cope with 9/11. Y2K was a known event and was essentially a software problem. Also, it did not raise the issue of destruction of people and property." Essentially, redundant mechanisms planned as contingency to carry out core functions of a financial institution (e.g., manually for a transitory

[1] Hong Kong Monetary Authority (HKMA). Supervisory policy manual - business continuity planning. URL http://www.hkma.gov.hk/media/eng/doc/key-functions/banking-stability/supervisory-policy-manual/TM-G-2.pdf

period) in the event of a known software issue is fundamentally different in nature than physical destruction of the whole infrastructure, and so on. In fact, the events of 9/11 forced a major rethink on how companies view their IT contingency plan. To quote Keith Payne, IT security officer at Javitch, Block & Rathbone, a law firm from a PC world article[2]: "Before 9/11, our customers did not heavily evaluate the possibility that the entire firm could cease to exist with no pre-indicators." This thought process can be seen in action across the board, and well encapsulated in the following HKMA's statement[3]: "plan on the basis that they may have to cope with the complete destruction of buildings in which key offices or installations are located (rather than just denial of access for a period) and the loss of key personnel (including senior management)".

The culmination of such a thought process is reflected at multiple granularities: starting from how data backup is being done by many organizations, for instance, moving away from periodic (say, weekly) transfer of back-up tapes to an alternate storage location which is often inadequate for critical data, to real-time online back-up solutions, often enabled by new technology such as cloud services, to the provisioning of alternate physical site ready to be moved into and operate from at any time with no or little notice, are some elements of a robust contingency plan.

And yet, as the maxim goes, one size does not fit all. The nature of trigger events would dictate the nature of effective contingency mechanisms. For instance, for a flu pandemic like SARS, where a whole team may be confined at home under quarantine, even if they are actually healthy to work, the disruption does not come from destruction of infrastructure, and depends on whether the services can be kept operational while employees work remotely, and so on.

Even as the nature of threats (or our perception) thereof is in a continuous flux, the technology landscape is also dynamic, and emerging technologies allow us to revisit the possible solutions. For instance, cloud services facilitate new back-up solutions. Amazon's web service deploy geographically distributed data centers, with multiple availability zones within individual geographic regions, whereby data centers utilise independent resources (e.g., independent power supply, etc.), so realise distribution and fault tolerance at an unprecedented scale. Many organizations have embraced and codified the usage of pesonal computing devices for work related activities under bring your own devices (BYOD) mechanisms, or allow remore working using virtual private networks, and so on. These provide further capabilities when dealing with contingencies. It is however to be noted that each of these new opportunities also bring along new attack surfaces that need to be holistically protection

[2] http://www.pcworld.com/article/239693/911_attacks_changed_the_way_companies_view_it.html

[3] http://www.hkma.gov.hk/eng/key-information/guidelines-and-circulars/circulars/2002/20020131a.shtml

against while devising operational security of the systems.

Beyond an organization's own self-interests in planning for contingency, and technical considerations, another key driver is an organization's regulatory obligations. Depending on the criticality of the operations of an organization, for instance, a financial or health institution, or a federal entity, the organization's functioning may have wide-scale implications on the society at large, and thus may fall under the purview of an appropriate regulatory agency. The above anecdotes from Hong Kong Monetary Authority provide a ready example. Likewise, federal organizations in the United States need to be FISMA (Federal Information Security Management Act) compliant, and to that end NIST (National Institute of Standards and Technology) has formulated several documents, for instance, NIST special publications [4] addressing issues relevant for contingency planning and business continuity. In addition to the explicit legal provisions requiring an organization to implement proper business continuity and disaster recovery mechanisms, there could be implicit drivers as well. For instance, in the United States, all publicly traded companies need to comply with the auditory requirements of Sarbanes-Oxley Act of 2002 (SOX), which then in turn requires preservation of all data essential for a proper audit.

Shared purpose, multiple responses

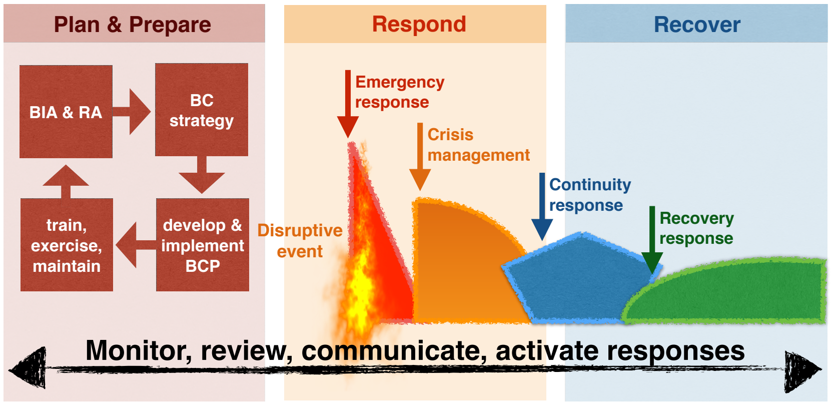

Figure 1: This figure has been adapted from Business Continuity Management Framework 2014 -18 document by Queensland Government, which in turn had adapted the visualisation from Marsh & McLennan Companies

As mentioned above, contingency planning needs to take into account the nature of event, to determine suitable response. Suitable response also needs to be determined based on the specific objectives that one wishes to achieve with a response. For instance, during a natural disaster, the priority could be to act in a timely manner to save lives, and contain the effect on the physical infrastructure to minimize data loss. However, after a certain time window, the priority could shift to not so much the life saving activities, but

[4] Swanson et al., 2010; and Blank and Gallagher, 2013

disaster recovery related activities instead. This exemplifies that as part of contingency planning, one needs to chart out varied responses (or a combination of several responses), as shown in Figure 1: e.g., emergency response, crisis management, business and disaster recovery, etc., depending on the immediate objectives in mind, but they all share the ultimate purpose of business continuity. We next elaborate the characteristics of some important response variants and their underlying rationale.

Emergency management: Emergency is something time critical, which needs quick response to reduce damage/losses of people?s life, physical or information assets. The purpose of emergency management is to save human lives, and try to contain or stabilize the situation, reducing loss to physical or IT infrastructure, and carry out damage assessment. For example, terror attack, or a fire on the physical premises of an organization will require the activation of emergency responses.

Crisis management: Crisis is a situation with potential knock-on and long term adversarial effects, affecting an organization's reputation, stock prices, stock market, etc. Crisis management is strategic in nature, and needs to take into consideration (operational) policy issues, manage communication with various stake holders (employees, customers, suppliers, etc.) through various channels including media and social media, liaise with necessary external entities, coordinate service recovery. The December 2013 Target data breach[5] is a relevant example. The incident tarnished Target's reputation, and damaged customer confidence leading to loss of business, fall in stock prices, etc. Two years down the line, with heavy shake-up in the management, and a lot of painful restructuring and layoffs later, Target appears to have weathered the crisis reasonably.[6]

Continuity & Business recovery: Given resource and time constraints, it may not be possible to restore normalcy immediately after an event, or to restore all functionalities simultaneously. It is thus necessary to determine the impact of an event on different aspects of a business, and identify the criticality of different activities, and accordingly carry out a phased recovery of business critical processes and services. This includes identifying and continuing the very critical services of the business, even as the recovery process to restore full normalcy is gradually activated by carrying out disaster recovery activities such as recovery of data and restoration of infrastructure and services.

[5] http://krebsonsecurity.com/2013/12/sources-target-investigating-data-breach/

[6] Target stocks were traded in NYSE at as much as above 65$ in November 2013, but had dropped to around 56$ in mid-February 2014 as details and implications of the data breach were getting exposed, however the stock price had risen up to 74$ by end of November 2014. When the fluctuations of stock prices is considered across a longer span of time, the crisis triggered by the data breach does not appear to have had any longterm ill effects on Target. Ironically, FireEye, the company behind the malware detection software used at Target, which traded at around 37 dollars in November 2013, fuelled by the positive publicity of how it had actually detected the attacks (but the red flags were ignored by Target's IT security team) had risen beyond 85 dollars by end of February 2014, and yet, by December 2014, it traded at less than 30$!

The show must go on: How to do business continuity planning?

Even an innocuous incident such as a water leakage or small fire in a part of a building may render the whole building temporarily unusable. The emergency response in this case will be to carry out evacuation to ensure there is minimal loss of life. If business is halted for an extended period of time as a consequence, it may lead to enormous financial loss, as well as loss of reputation and future business opportunities. Consequently a business may need to temporarily relocate its employees to operate from alternate sites within a short span of time. This may need operational support in many manner, such as a mechanism to divert calls to the right temporary offices that the personnel operate from, that they have the necessary equipments to work with, etc. Eventually, the business should be able to locate to a permanent site (the old one restored, or a new one), and resume operations at full capacity. The resilience in the system should be built in a manner so as to cope with different degrees of disruptions.

The mantra behind the design of Amazon's Dynamo[7] database, originally used in the backend of their e-commerce site, and eventually also provided as part of their cloud service offering is surmised in the following quote from the system's designers "customers should be able to view and add items to their shopping cart even if disks are failing, network routes are flapping, or data centers are being destroyed by tornados". This design principle aptly captures the essence of business continuity — notwithstanding disrupting events — the core functions of the business must go on. In fact, resilience is ingrained in Amazon's infrastructure design, and other major cloud services also adopt similar techniques. For instance, with the elastic compute cloud (EC2) offering, Amazon introduced redundancy at multiple granularities across data centers - spread across geographic regions, and furthermore, within each region, deploying data centers in isolated locations, relying on independent physical infrastructure, and called availability zones. End users using a service like EC2 can choose to configure their deployments to take advantage of different degrees of resiliency.

Most businesses have neither the economic wherewithal nor the technical expertise to build all the IT solutions and infrastructure in-house to meet their business continuity (BC) and disaster recovery (DR) needs. Consequently, many businesses may rely on managed services for business continuity and disaster recovery. With the proliferation of cloud services, businesses are increasingly relying on managed disaster recovery services. Given the vital role of BC/DR services, be it carried out in-house or is out-sourced, there

[7] DeCandia et al., 2007

has been several national and international standards and guidelines. Prominent ones include ISO/IEC published guidelines (ISO/IEC 24762:2008) for provisioning of information and communications technology disaster recovery (ICT DR) services [8], which itself was in part derived from the Singapore standard (SS507:2008) for information and communications technology disaster recovery services. The has since been withdrawn, and ISO 22301:2012 [9] on business continuity management systems and ISO/IEC 27031:2011 [10] guidelines for information and communication technology readiness for business continuity are some other more recent relevant documents, as are the NIST special publications [11] mentioned previously. Instead of getting into the specific details and differences across these myriad of guidelines and standards, we will instead next summarize the key ideas derived from across these.

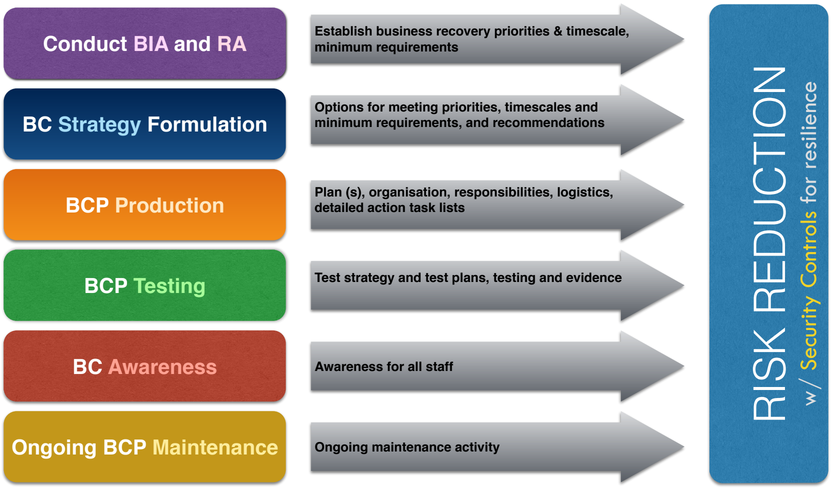
Conduct BIA and RA
Establish business recovery priorities & timescale, minimum requirements
BC Strategy Formulation
Options for meeting priorities, timescales and minimum requirements, and recommendations
BCP Production
Plan (s), organisation, responsibilities, logistics, detailed action task lists
BCP Testing
Test strategy and test plans, testing and evidence
BC Awareness
Awareness for all staff
Ongoing BCP Maintenance
Ongoing maintenance activity
RISK REDUCTION
w/ Security Controls for resilience

Figure 2: Business continuity planning approach from ISO/IEC 24762:2008. The following abbreviations have been used in this figure — Business Impact Analysis (BIA), Risk Assessment (RA), Business Continuity (BC), Business Continuity Plan (BCP).

Business continuity plans and implementation involve costs, so essentially BC/DR can be viewed a within the ambit of risk management. Each organization needs to determine the worst case scenarios that it wishes and can afford to prepare for. ISO/IEC 24762:2008 outlines a number of discrete but dependent stages, as depicted in Figure 2, showing how business continuity planning is interconnected with risk reduction. The first five stages, namely (i) Business Impact Analysis (BIA) and Risk Assessment (RA), (ii) Business continuity (BC) strategy formulation, (iii) Business continuity plan (BCP) production, (iv) Business continuity plan testing and (v) ensuring staff awareness regarding business continuity are executed sequentially, while the sixth stage is a process to review and revise the existing plan, to be carried our periodically or after major changes that may effect the efficacy of the existing plan (e.g., change in the organization's infrastructure or nature or business, new legislations coming into effect, etc.). The seventh stage, that of risk reduction by

[8] ISO/IEC, 2008

[9] ISO, 2012

[10] ISO, 2011

[11] Swanson et al., 2010; and Blank and Gallagher, 2013

applying necessary security controls, is carried out in parallel, based on the outcomes of the other six stages.

Business Impact Analysis

Business impact analysis (BIA) is the process of analysing the consequences on an organization due to the interruption of (some of) its activities that support its business activities, products and services. Since there could be interdependencies across functionalities, a rigorous BIA ought to identify such connections and correlations, to determine the criticality of individual functions and resources. Based on the extent of effect individual elements in the whole system may have, one then needs to carry out risk assessment (how likely they are going to happen, what will be the extent of exposure, etc.) and determine the overall impact — be it quantitative (e.g., directly attributable financial loss) or qualitative (e.g., loss of reputation) in nature. The impact is often dependent on the degree and duration of an outage, which needs to be taken into account while considering response options. The BIA along with risk assessment thus help determine the maximum tolerable downtime (or minimal acceptable level of service), and accordingly the recovery time objective (RTO) and recovery point objective (RPO), which in turn then determine the necessary response strategy and resources. The BIA and risk assessment exercise is naturally also coupled with risk management and risk reduction activities, but it should be noted that these are still two distinct activities. For instance, an organization may decide to accept some risks in not investing in any controls to mitigate it - for instance, because the cost of controls may be too high, or the likelihood of the event may be too less. Nevertheless, the event actually may still happen, and business impact analysis needs to determine the consequences if that is the case, and how to deal with the consequence. In that sense, risk management and contingency planning are opposite sides of the coin — the former exploring mechanisms to mitigate risks, the later looking into how to respond if the risk actually comes to pass.

The BIA process can thus be seen as comprising of the following logical steps:

- Determine Critical Business Processes, Services, and Products: Those that must be restored immediately after a disruption to ensure the affected organization's ability to protect its assets, meet its critical needs, and satisfy mandatory regulations and requirements.

- Identify activities that support provision of critical business pro-

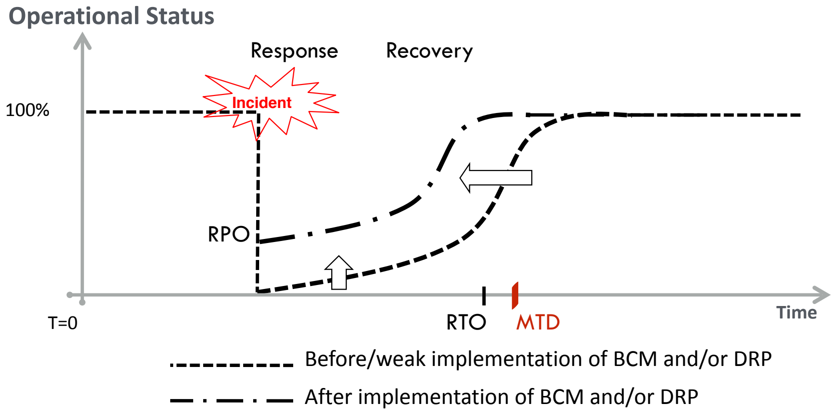

Figure 3: Business continuity management (BCM) and disaster recovery planning (DRP) practices focus on shortening period of disruption (RTO) within a maximum tolerable downtime (MTD), and reducing the impact of an incident (RPO) by risk mitigation and recovery planning.

cesses, services, and products: While the critical processes or services and products for the business are often obvious to the corresponding users, the underlying resources enabling these may be overlooked. e.g., documentation, back-up (and periodically checking that restoration works), operating systems, database management systems and data center tools, etc.

- Assess impacts over time of not performing these activities: e.g., Loss of life, damage to physical assets, denial or disruption of critical technology services, etc.
- Set prioritized timeframes for resuming these activities: e.g., Minimum acceptable level; Maximum tolerable downtime.
- Identify dependencies and supporting resources for these activities, including all third parties dependencies: e.g., if the data center of the back-up service provider is confined in the same geographic region as the business itself, that may be affected by a natural catastrophe such as earth quake that affects the business itself.

The key considerations to determine the recovery requirements include customer impact analysis, system impact analysis, policy and regulatory compliance requirements and third party dependencies, which together dictate the *Maximum Tolerable Downtime* (MTD) a.k.a. Maximum Tolerable Period of Disruption (MTPD), i.e., the period of time after which an organization's viability will be irrevocably threatened if delivery of a particular product or service cannot be resumed. Ascertaining dependencies across activities, along with prioritization based on their criticality aide the linking of activities with similar recovery requirements to form a timeline based recovery action plan. Accordingly, a *Recovery Time Objective* (RTO), i.e., the duration of time, from the point of disruption, within which an activity should be restored, is determined per activity. The organization also needs to classify the criticality of different data, and determine corresponding *Recovery Point Objective* (RPO). It essentially gives systems designers a limit to work to - e.g., if restoring data till up to last week is deemed sufficient, an offline tape based back-up solution, where the tapes are moved off-site on a weekly basis would suffice, while if data up-till the last hour is required to be retrieved, it will necessitate a fundamentally different system design.

The Plan-Do-Check-Act Cycle for Business Continuity Management

ISO 22301:2012[12] adapts the management method of Plan-Do-Check-Act (PDCA) in the context of Business Continuity Management Systems (BCMS), as depicted in Figure 4. Quoting from the standard itself, the purpose of the four stages can be summarized as follows:

Plan: Establish business continuity policy, objectives, targets, controls, processes and procedures relevant to improving business continuity in order to deliver results that align with the organization?s overall policies and objectives.

Do: Implement and operate the business continuity policy, controls, processes and procedures.

Check: Monitor and review performance against business continuity policy and objectives, report the results to management for review, and determine and authorize actions for remediation and improvement.

Act: Maintain and improve the BCMS by taking corrective action, based on the results of management review and reappraising the scope of the BCMS and business continuity policy and objectives.

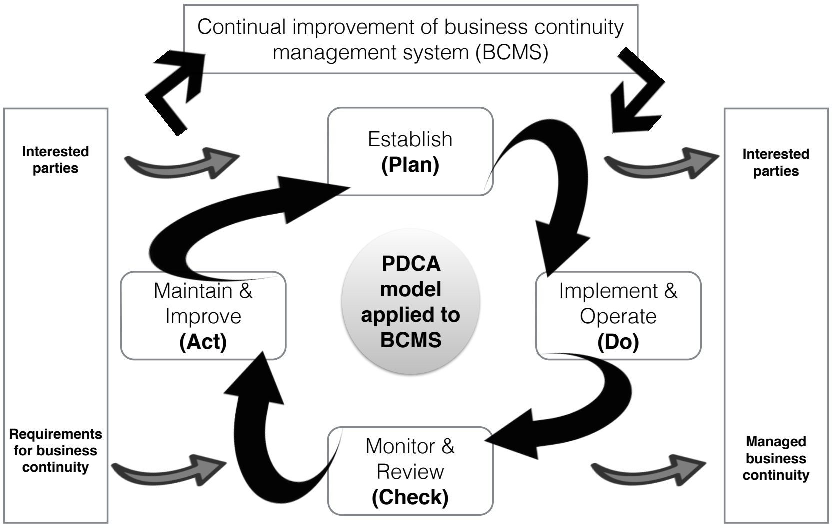
Continual improvement of business continuity management system (BCMS)
Interested parties
Requirements for business continuity
Establish (Plan)
Maintain & Improve (Act)
PDCA model applied to BCMS
Implement & Operate (Do)
Monitor & Review (Check)
Interested parties
Managed business continuity

Figure 4: This Plan-Do-Check-Act (PDCA) cycle for Business Continuity Management Systems (BCMS) figure has been adapted from ISO 22301:2012 standard. Plan-Do-Check-Act (PDCA) is a general purpose iterative four-step management method used in business for the control & continuous improvement of processes and products.

[12] ISO, 2012

The rationale behind the PDCA approach for BCMS is to capture and adapt to a changing landscape (or our understanding) of worst case scenarios. For instance, a contingency plan under the assumption of destruction of a building, but assuming that there will be some survivors who can then operate from an alternate site may not be suitable for a fundamentally different event such as a pandemic flu — where there is no physical harm to the buildings and IT infrastructure, and thus provisions to operate from an alternate site is not the solution, since a whole team involved in orchestrating a specific business activity may have to be quarantined in their respective homes (even if they are in fact physically fit to work).

This example also highlights another interesting subtlety to business continuity and recovery planning, namely that, it needs to be scenario driven. The type of disaster (earthquake, fire, pandemic, haze, etc.) would determine the point(s) of impact as well as the depth and scope of impact, and the availability of resources to respond with. For instance, back-up servers located in the same building may not be robust against a fire that destroys the building, while even an offsite storage facility in the same city may also not be robust against a natural disaster affecting a whole geographic region. And yet, if there is a major breakdown of the communication channels (e.g., ISP failure), then, a local back-up mechanism may come handy.

Since the business continuity management system (BCMS) is itself realized using Information and communication technology (ICT), an integral part of evaluating the efficacy (check) of the BCMS processes involves ascertaining ICT readiness for business continuity (IRBC). ISO 27031:2010[13] deals with the issue of IRBC and articulates on how to conduct tests and exercises with progressive level of meticulousness, in order to validate the BCM/DRP arrangements with increasing degree of confidence. To quote ISO 27031:2010 "The organization should exercise all elements of the ICT service recovery as appropriate to its size, complexity and business continuity management scope. The exercising should not focus solely on service recovery and resumption, but should include the reliability of the resilience capability, system monitoring and alert management. Finally, the organization should exercise at component level through to full location-based system testing in order to achieve high levels of confidence and resilience." Accordingly, the testing spans review of and familiarization with the existing documentation (desktop process review), test of individual components and integrated end-to-end services (recovery simulation), and eventually full-scale (operational) simulation by switching ICT and end-to-end services between primary and secondary.

[13] ISO, 2011

Further BCM/DRP Considerations

So far we have articulated the main concepts of business continuity management (BCM) and disaster recovery planning (DRP). We wrap up the topic by discussing some practical considerations and specific elaborative examples.

Organizational considerations

An organization needs to clearly identify and enable a team with well defined roles for business continuity planning. The actual business recovery in case the business contingency plan has to be activated may be carried out by a somewhat ad-hoc team, depending in part on the available human resources a posteriori.

The planning is typically done while the organization is operating in its normal mode. The business continuity manager would interact with the business units in order to understand their operations, and carry out business impact analysis, and liaise with support functions for people, logistics, infrastructural facilities, technology related issues, and report to the upper management with the findings, strategic inputs, budgetary approval, etc. The executive management however is unlikely to be involved in the day-to-day activities of the business continuity planning team.

In case there is an incident, a first verification (triage) that there is genuinely an incident has to be made by whoever encounters it, and this person needs to notify the operational team (notification) that is most directly responsible for dealing with or affected by the incident. Once the team confirms the incident (escalation), it needs to carry out damage assessment, and determine the severity and decide whether it can be contained by applying the normal operational reactive controls. However if the severity is beyond a specific threshold, or continues to grow, then the executive management team needs to be notified, who would have to make a decision on whether to declare a disaster, and accordingly activate the response plan and mobilize a response team. This whole plan activation process is depicted in Figure 5. The response team may be ad-hoc in nature, depending on the available personnel. Unlike the planning phase, the executive team is typically more directly involved in the recovery activities, simply because that becomes the primary activity for the transient period.

It should also be noted that in practice, some of the logical steps of the plan activation may disappear or collapse into a single step, depending on the nature of incident. For instance, a fire or terror attack in a building would likely trigger emergency responses without ex-

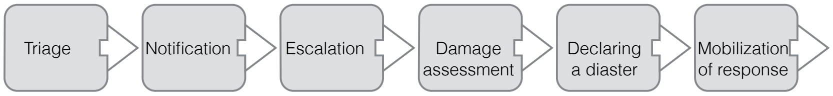

Figure 5: Plan activation for business/disaster recovery and other (emergency/crisis) responses.

plicit intervention and declaration from the CEO of a company. The planning phase should also clearly articulate a chain of succession for key roles, so that if some of the top executives are also adversely affected by the incident, there is no ambiguity on who ought to fill up the roles for the transient period, and with what authority and responsibilities. This is crucial, given that declaration of the disaster itself is an executive decision, as is the execution of the recovery activities. The identified persons for specific roles need also be competent for the corresponding tasks, even if that is not what they do during the normal operations.

Communication

Communication plays vital role in all aspects of business continuity and disaster recovery activities. To start with, all the processes need to be well documented, and personnel need to be made aware of what is expected of them in case of an incident. For instance, everyone needs to know the evacuation plan, alternate site where to report after evacuation, or what other actions to take, etc.

An effective incident response structure is also essential. This includes detection and notification channels for timely dissemination of information regarding incidents to various stake holders — employees, customers, suppliers, etc.

Alternate sites

The organization should provision for an alternate head quarter from where the recovery team can coordinate and execute the recovery activities. This is in addition to (though, it may be co-located) having a back-up site, where the organization can relocate after a disaster such as fire, flood, terror threat, etc., to continue its business operations.

Broadly, there are three categories of back-up sites: cold, warm and hot. A cold site is the cheapest to provision, since it does not come equipped with the hardware infrastructure (nor does it have any backed up data, or working software, etc.), but in case of a disaster, it will require the additional time to set-up all the capabilities before resuming operations. In contrast, a hot site is essentially a duplicate of the original site, with fully functional IT-infrastructure, which is synchronized (almost) in real time with the primary site of operations, so that the organization may relocate anytime, and resume normal operations with minimal disruptions. However, the capacity of the hot site may or not match that of the primary site, and it may be able to handle only a lower volume of business. Naturally, a hot site is the most expensive option, and is opted by a limited set of organizations, such as financial institutions, government agencies, etc. Warm sites look for a compromise between the two extremes and may have partial infrastructure or data in place to start with.

(Third party) dependencies

An organization should determine if its business continuity plan is incompatible with the third party service provider's business continuity plan (for instance, in terms of mismatched RTO).

Service level agreements notwithstanding, there are additional risks of correlated failures at different levels of granularities. If the third party is vulnerable to the same threats that an organization is trying to alleviate risks by outsourcing a service, then it may not be useful. If multiple organizations all utilize the same third party service for some crucial functionalities, then the failure of this common third party service may lead to disruption of functions for all the organizations. For certain kinds of organizations, for instance, financial institutions, disruption of their services can have wide-scale societal repercussions. For this later scenario, it is desirable to avoid such correlated failures by instead using independent third party services.

At the same time, business units performing similar tasks (both within an organization, or across organizations) could make arrangements for mutual support or shared services, to restore operations. In fact, even without explicit arrangements, organizations may be able to leverage on redundancies to determine the criticality of the specific functions. For instance, financial institutions may leverage on the fact that depositors can carry out deposits with other organizations operating in the same geography (assuming that to be the case), and thus determine that the deposit function is relatively less critical [14] than others like clearing and settlement.

Technological aide

Even as business continuity planning and management appears to be a daunting and expensive exercise, technological advancements in several areas make it increasingly more tractable.

Sophisticated business continuity management tools[15] have evolved over time, which can help carry out risk assessment for availability, BIA from loss of people, IT, facilities, suppliers, business

[14] FSB, 2013

[15] Witty and Morency, 2014

process & IT dependency mapping, workflow management, analytics for understanding effectiveness, risk reduction & cost, etc.

There are multiple avenues to amortize resources — for instance, by re-purposing existing infrastructure, pooling resources through virtualization, or adopting cloud based solutions which remove the need for capital investment or a large specialized IT team (for businesses, for whom IT is not the core competence), bare-metal restore which allows having cost-effective warm sites practical, to name a few.

Concluding remarks

Contingency planning and management is key to business continuity, and complements risk assessment and management. Not surprisingly, there is no perfect plan, nor one specific solution that fits all scenarios, but there are several standards & guidelines, as well as software tools to facilitate the process, which needs to be regularly tested, updated, communicated to stake holders, exercised and tested again, and so on and so forth, in a continuous PDCA (plan, do, check, act) cycle. There are different responses, such as emergency, crisis, business recovery and disaster recovery, but with shared purposes. The nature and the specificities of the response depends on the specifics of a scenario, and thus the contingency planning needs to span multiple worst case scenarios. Even if an organization decides to accept certain risks, the contingency planning exercise, and in particular, the business impact analysis, needs to take those risks also into account. Last but not the least, for successfully managing business continuity, it is paramount that there is support and commitment which is both top-down and bottom-up.

Bibliography

Europol. EU Organized Crime Threat Assessment. OCTA 2011, 2011a.

Europol. Major international network of payment card fraudsters dismantled, 2011b. URL https://www.europol.europa.eu/content/press/major-international-network-payment-card-fraudsters-dismantled-1001.

David Wolman. *The End of Money: Counterfeiters, Preachers, Techies, Dreamers–and the Coming Cashless Society*. Da Capo Press, 2012.

Christopher Hadnagy. *Social Engineering: The Art of Human Hacking*. Wiley Publishing, 2011.

Charles P. Pfleeger and Shari Lawrence Pfleeger. *Security in Computing*. Pearson, 2007.

Rebecca Boyle. Hackers could access pacemakers from a distance and deliver deadly shocks, 2012. URL http://www.popsci.com/technology/article/2012-10/hacker-attackers-could-reverse-pacemakers-distance-delivering-deadly-shocks.

Alan Grau. Hackers invade hospital networks through insecure medical equipment. *IEEE Spectrum*, 2015.

John D. Howard and Thomas A. Longstaff. A common language for computer security incidents. Sandia report, 1998.

ISO. Information processing systems – open systems interconnection – basic reference model – part 2: Security architecture. Technical Report ISO 7498-2:1989, International Organization for Standardization (ISO), 1989.

Gary Stoneburner, Clark Hayden, , and Alexis Feringa. Engineering principles for information technology security (a baseline for achieving security), revision a. Technical Report NIST Special Publication 800-27 Rev A, National Institute of Standards and Technology, 2004.

Bob Bakely, Ellen McDermott, and Dan Geer. Information security is information risk management. In *Workshop on New security paradigms*, 2001.

James Bayne. An overview of threat and risk assessment. Technical report, SANS Institute, 2002.

Kelley Dempsey, Nirali Shah Chawla, Arnold Johnson, Ronald Johnston, Alicia Clay Jones, Angela Orebaugh, Matthew Scholl, and Kevin Stine. Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations. Technical Report Special Publication 800-137, National Institute of Standards and Technology, 2011.

Ding Tan. Quantitative risk analysis step-by-step. Technical report, SANS Institute, 2002.

Hong Kong Monetary Authority (HKMA). Supervisory policy manual - business continuity planning. URL `http://www.hkma.gov.hk/media/eng/doc/key-functions/banking-stability/supervisory-policy-manual/TM-G-2.pdf`.

Marianne Swanson, Pauline Bowen, Amy Wohl Phillips, Dean Gallup, and David Lynes. Contingency Planning Guide for Federal Information Systems. Technical Report Special Publication 800-34, National Institute of Standards and Technology, 2010.

Rebecca M. Blank and Patrick D. Gallagher. Security and Privacy Controls for Federal Information Systems and Organizations. Technical Report Special Publication 800-53, National Institute of Standards and Technology, 2013.

Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. Dynamo: Amazon's Highly Available Key-value Store,. In *Twenty-first ACM SIGOPS Symposium on Operating Systems Principles (SOSP)*, 2007.

ISO/IEC. Guidelines for information and communications technology disaster recovery services. Technical Report ISO/IEC 24762:2008, ISO/IEC, 2008.

ISO. Business continuity management systems — requirements. Technical Report ISO 22301:2012, International Organization for Standardization (ISO), 2012.

ISO. Guidelines for information and communication technology readiness for business continuity. Technical Report ISO/IEC

27031:2011, International Organization for Standardization (ISO), 2011.

FSB. Recovery and resolution planning for systemically important financial institutions: Guidance on identification of critical functions and critical shared services. Technical report, Financial Stability Board, 2013.

Roberta J. Witty and John P Morency. Magic quadrant for business continuity management planning software. Technical report, Gartner Inc., 2014.