Key Considerations for Business Resiliency

John P. Pironti, CISA, CISM, CGEIT, CISSP, ISSAP, ISSMP, is the president and founder of IP Architects LLC. He has designed and implemented enterprisewide electronic business solutions, information security programs, business resiliency capabilities, and threat and vulnerability management solutions for key customers in a range of industries, including financial services, government, hospitality, energy, aerospace and information technology on a global scale. He is also a published author and writer, frequently quoted and interviewed, and a frequent speaker on electronic business and security topics at domestic and international industry conferences.

Business resiliency is the maturation and amalgamation of the individual processes of crisis management, incident response, business continuance and disaster recovery into one succinct set of processes and capabilities that work collectively, instead of independently. This combination allows organizations to have minimal disruption in the event of a business-impacting incident that affects the entire organization, instead of focusing on incidents that involve specific information infrastructure areas. When evaluating these capabilities, it is important to understand that they are only as effective as the proactive planning and considerations that go into their development. Too often, planning accounts for only the most obvious considerations and does not incorporate crucial and essential considerations that have a greater effect on the business.

CRISIS MANAGEMENT

The crisis management capability represents the umbrella under which all other business resiliency capabilities fall. This capability includes the decision-making element of the business resiliency program, often known as command and control. Command and control involves the key elements essential to the initial and ongoing management of any business-impacting event or incident. Typically, it is comprised of the organization's senior leadership, but it should also include key stakeholders throughout the organization.

When developing a crisis management capability, it is important to identify specific scenarios in which the organization establishes predetermined action plans and then formulates a generic action plan for all other scenarios. Common crisis management scenarios establish action plans for a logical disruption of business activities, a physical disruption of business activities, negative media attention, employee safety and soundness, and financial distress or insolvency. Each of these scenarios has specific requirements and actions along with generic requirements and actions that are universal among all scenarios. The universal elements in all scenarios typically include communication capabilities (internal and external), legal considerations, financial capabilities, facilities, and personnel.

Command and Control

A command and control capability involves the people, process, procedures and facilities required to identify, analyze and react appropriately to both predefined and general incidents that can affect the organization's ability to conduct itself in a business-as-usual manner. The first 72 hours of an incident are usually the most critical and require the most advanced and prescriptive planning.

Leadership Identification and Availability

When developing a command and control capability, it is important to identify the key leaders and stakeholders in the organization and document their normal business responsibilities and any business-critical organization information of which they may be custodians. This information generally includes key documentation, organizational responsibilities, financial and signature authority, contacts with outside organizations and customers, and legal information. A common and critical mistake made by many organizations is to assume that the senior leadership will be available and capable of making decisions for the business during a business-impacting incident. In an effective business resiliency program, all factions must have a contingency plan for all critical elements, including key staff availability. The key leaders of the organization should identify multiple tiers of delegation-of-authority (a minimum of two backups, geographically separated, when possible) to act on their behalf in the event that they are unable to participate in crisis management activity.

Communication Plan

The most crucial and most often underdeveloped component in a command and control capability is an effective communication plan. In a crisis, interested parties (internal and external) focus on the organization and have an insatiable need for information. If a clear and concise communication plan is not in place, an organization runs the risk of creating an

environment of misinformation resulting from fragmented data and assumptions. Senior leadership should develop and approve an initial communication plan to accommodate zero-hour communications (communications at the time of incident identification and declaration) in advance of any incident. Typical zero-hour communication includes language that states that the organization has encountered a business-impacting event and is investigating the situation. It also includes guidelines to determine when to provide future updates to interested parties and how the interested parties can receive these updates.

A regular stream of updates is critical to any crisis communication strategy. These updates should be communicated on multiple platforms, including web sites, e-mail, voice mail broadcasts and, when possible, in-person briefings. During the first 72 hours, updates should be on a regular basis, typically four-hour increments. If possible, during the initial hours of the crisis period, it is suggested that updates be provided each hour, but quickly reduced to a longer window of time to allow for crisis remediation activities to take place. Even when there is nothing to report, the issuance of an update conveys that the organization is addressing the situation and more information will be available at the next scheduled update, or as warranted. Otherwise, interested individuals may begin to lose confidence in the organization's ability to remediate the incident.

Communications need to be consistent for both internal and external audiences during a crisis. Consistency reduces miscommunication and minimizes the socialization of misinformation. The use of consistent language, terms and speakers (if possible) is preferred to allow the audience to receive accurate messaging. Press releases to media and posted on web sites are generally acceptable for initial communications, but often there is a need for interactive discussion within the first hours of an incident to allow for questions and feedback.

It is also important to ensure that the method of communication is consistent. To accomplish this, a leading practice is to use web sites associated with the organization and provide a telephone number for clients to call with inquiries. When establishing the telephone number, it is important to use one that is separate from the organization's business-as-usual call center, to avoid an influx of inquiries that would impede normal business call traffic. Industry-leading practices in this case are to contract a call center organization in advance, provide scripts for zero-hour communications and provide them with scripted updates as the incident response activities continue to operate.

Internal communication elements such as logistic information for staff involved in recovery efforts and confidential information should utilize the same information infrastructure capabilities as external communication capabilities. The only difference should be the addition of identification and authentication elements to ensure that only authorized individuals are receiving the information and to track who has received the information for audit and accountability purposes. This capability will also allow general staff to know when and where they can return to work once the recovery efforts allow.

Another critical consideration is the engagement of outside assistance to help in the crisis. Many organizations are confident in their ability to deal with any external communication requirement, but, in reality, most do not have experience with media inquiries. An external organization (such as a public relations firm) introduces an unbiased and unaffected perspective to crisis communication, ensuring that only appropriate information is released and transmitted through the most effective channels. When organizations do not properly vet their external communications, they may release potentially harmful information in an emotionally charged environment, due to the stress of actively remediating the incident. It is important to contract the outside organization in advance of a crisis and establish strategies in advance.

Legal Considerations

During a crisis, the senior leadership of an organization may not be available to execute legal agreements or make decisions. It is important to implement delegations of authority for all identified contingency staff prior to the occurrence of a crisis. This allows these individuals to execute binding legal agreements on behalf of the organization. It also enables expanded signature authority to enter into agreements with higher-than-normal signing authority for emergency considerations.

Important components of a command and control tool kit are the declaration of incident and completion of incident legal documents. During an incident, these documents allow the initiation and legal recognition of the delegation of authority and expanded signature and financial authority by external organizations. Upon completion of the incident or when no longer necessary, the documents remove these capabilities. Otherwise, the individuals granted the expanded signature authority and responsibilities may not be recognized by external parties, which can delay or prevent critical agreements being executed. The individuals who declare the incident sign these documents and internal or external legal counsel countersign the documents.

Just as in a communication capability, an organization should retain external legal counsel with experience in crisis management to ensure availability during an incident. It is important to confirm that appropriate remediation steps are taking place to reassure interested parties (internal and external) and to allow for the availability of an unbiased perspective when making legal decisions.

INCIDENT RESPONSE

Incident response represents the second tier of an organization's response (the first tier is operational response) and, typically, the first organized and focused approach (see **figure 1**). By invoking the incident response capability, an organization has concluded that a situation is no longer an event that requires further investigation. The situation is now an incident that has the potential to impact business operations. This is an important distinction because an incident is very different from an event (an event is something that has occurred within the information infrastructure or operating environment of the organization that has the potential to cause a disruption and warrants further investigation). The declaration of an incident and activation of the incident response function represents a decision of the organization to identify, analyze and remediate the incident until it no longer has an impact on the organization's business. This is important both internally and externally to the organization. If an organization declares an incident and chooses not to take these actions, internal stakeholders and external examiners, customers and litigators may regard them as negligent.

One of the most important considerations when declaring an incident and invoking the incident response capability is determining whether the analysis and recovery will be performed in an operational or forensic manner. Both of these eventually result in return-to-normal business operations, but they do differ. An operational response has the goal of resolving the identified issue as quickly as possible, while causing minimal disruption to business operations. A forensic response focuses on the preservation and integrity of evidence while identifying and rectifying the business-disrupting incident.

After the incident response leader determines the style of incident response, root cause analysis and the execution of the remediation plans occur. As important as it is to determine when and how to declare the event, an aspect of incident response commonly overlooked is an understanding of when the incident has completed and when the organization should discontinue remediation activities. In the case of physical incidents (e.g., failed equipment, physical security breach), this is more apparent. Logical incidents can be much harder to measure because they can involve elements that are not immediately apparent (such as multiphase virus attacks that use diversion techniques to implant code into systems while initial remediation is taking place).

The current leading practice to determine the effectiveness of incident remediation in an operational response is to use key performance indicators (KPIs) or measurements that align to business operation effectiveness and the continued efforts of the incident response team. If the organization's business operations are able to function appropriately within the range of "minimal impact" to "no impact" after an identified incident, the incident response team can return the remediation activities to the operations organization, which can finalize the return-to-normal actions for the business operations. The incident response team may still assist the operations team with return-to-normal activities, but they no longer need to lead, organize and manage these efforts. The KPIs should have established thresholds that define the risk and operational effectiveness limits that the organization is willing to accept and, in turn, activate or deactivate its incident response capabilities.

Figure 1—Incident Escalation Matrix

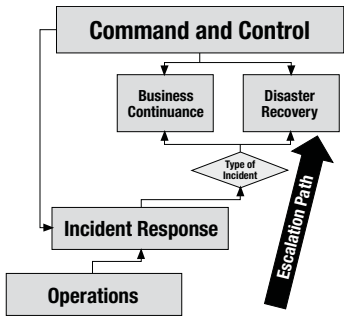

In the case of a forensic response, the same KPIs measure the incident response requirements, but a legal opinion is advantageous prior to deactivating the incident response capabilities. This ensures proper substantiation of all legal, chain of custody, and evidence preservation and collection considerations.

BUSINESS CONTINUANCE

Business continuance focuses on an organization's ability to continue to operate effectively in the event of a business-debilitating incident. The first step in developing a business

continuance capability is to identify the business processes important to the organization, the level of capability required to meet minimum effectiveness requirements and how the business can sustain business operations no matter what the disruption.

Historically, the best way to determine the business processes most important to an organization was to map the revenue streams on which the organization depended. While this is still a valid technique, current business conditions also introduce other factors that need consideration, such as regulatory and industry compliance requirements, contractual arrangements, and customer expectations.

One of the most overlooked, but important, business continuance concerns is the organization's impact on its partners and vendors. Today, most organizations provide products and services that support the success of its customers' business activities. Many organizations have contractual requirements with these organizations in regard to availability of services and capabilities. These agreements often take the form of service level agreements (SLAs) that have associated financial and potential legal consequences if the organization cannot meet the SLA requirements. Therefore, a business process that seemingly represents a medium priority to the organization from a revenue perspective may, in fact, be a high priority in a business continuance situation, due to the financial and legal ramifications that an organization may face if the process is unavailable.

Once these situations are identified, it is important to establish secondary capabilities that can be leveraged during a business-impacting crisis. This can include arrangements with partner, or even competitive, organizations that provide similar services. These arrangements should be made in advance of a crisis and should be reciprocal in nature.

Business Impact Analysis

It is important to develop a business impact analysis that enumerates the impact of a loss of some or all business process capabilities. Many organizations conduct this analysis through a series of questionnaires in which they ask business process stakeholders about the potential impact if their process were to become partially or completely unavailable. This method tends to identify obvious business impacts, but often fails to enumerate all of the business process elements or the impact of their loss.

One of the most effective ways to approach a business impact analysis is to first conduct a business process mapping activity. Business process mapping provides a visual depiction of all of the process elements and associated dependencies for a particular business process (**figure 2**). This information is particularly useful when utilizing the business continuance plan. The visual depictions enumerate interdependencies and key elements within business processes on which a plan can focus its immediate attention.

After the organization completes the business process mapping, it identifies an inventory of the information infrastructure and data elements that support the business process. This includes the people, processes, procedures, technical infrastructure and data used in the business process. This provides a complete understanding of the dependencies

Figure 2—Sample Business Process Map

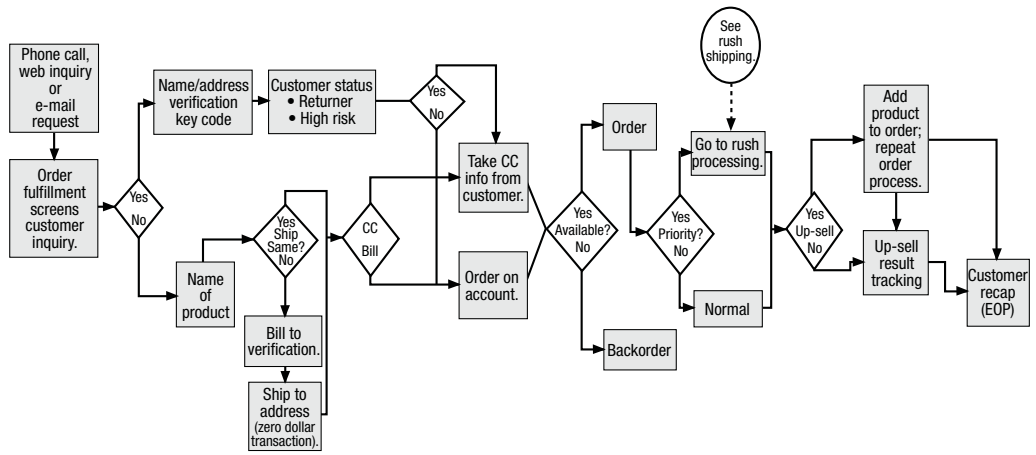

(resiliency capabilities inherent to the business process based on its design), key components and weaknesses that exist in its current implementation.

When conducting a business impact analysis, it is important to consider the effect of a partial and a whole loss of the business process. Many business impact analysis activities assume that a complete loss of a business process is less likely to occur than a partial loss (i.e., the loss of key infrastructure or personnel elements, data flows, or vendors). By using the output of the business process mapping activity, an organization can perform a business impact analysis by developing scenarios that assume the loss of key elements as well as the entire process. This ensures that the implemented preventive controls and recovery efforts are business-appropriate and most effective.

One of the key considerations when performing a business process analysis is to define the recovery point objectives (RPOs) and recovery time objectives (RTOs) associated with the analyzed business process. This allows an organization to define the minimal capabilities required for the business process to be effective and valuable. These minimal capability metrics drive the appropriate level of effort and the investment in a particular business process during a continuance effort. These metrics also help an organization understand when it can reduce its focus on recovery efforts as well as when it is no longer in the organization's best interests to continue the effort in the event of an unsuccessful recovery.

Competency Models and Staff Availability

An often neglected element in a business continuance plan is the ability for key staff to be available, willing to function and capable of carrying out their duties in a crisis. A business continuance plan should not assume that key staff members will work at levels beyond their typical capabilities and do everything possible to remediate a crisis. In many crises, when organizations initiate the business continuance plan, due to stress, the staff works at a depreciated level of competency.

To combat this situation, it is important to establish competency models for the roles identified as necessary for the business continuance plan to operate. These competency models are similar to job descriptions, but focus on specific activities that an individual will be required to perform during a business continuity situation. Competency models should include the specific skills, personnel profiles, knowledge and competencies that are required to perform the tasks outlined in the business continuity plan.

Once an organization establishes the competency models, it is important to create a sourcing strategy that can accommodate the defined requirements. Often, an organization finds that it has many skilled workers who have the appropriate skills to carry out their business continuity plans, but the workers are not currently in positions that would make this obvious. The most effective approach to identifying these individuals is to work with the human resources department to develop a skill inventory database of all employees and then map their skills to the individual business continuity plans.

Once the organization identifies the individuals who have the necessary skills, it is important to ask the individuals if they would be willing to take part in a business continuity effort and are capable of operating in a crisis. Some individuals may not feel comfortable with this role, and it is important to identify them in advance.

It is also important to recognize that the identified staff may not be available when the organization enacts the plan or the plan may require more staff than are currently in place. In this case, it is important to establish relationships with staffing organizations that specialize in providing staff with the required skills. The most effective way to do this is to communicate the competency models for the plans to these staffing organizations and implement a retainer contract that specifies that the staffing organizations keep an adequate inventory of available staff or, at minimum, have a contact database of possible candidates that map directly to the required competencies.

Financial Planning and Reserves

Business continuance plans, when enacted, can represent a significant financial cost to the organization. One of the key considerations that an organization often overlooks is the ability to continue to fund the plans until the organization has recovered from the incident. In many cases, the plans assume that the staff will continue to be paid and that the organization will cover expenses associated with the plan. Unfortunately, many organizations today do not have adequate financial reserves to pay all staff identified in the plan for an extended period, especially if the incident affects the organization's ability to generate revenue or its accounts receivable and payable functions. There is a high likelihood of failure if the organization does not compensate the required staff or pay the associated expenses for the plan.

It is important to work with the finance team to ensure the establishment of adequate reserves to meet the financial considerations of the business continuity plan and to understand how to access the funds if they are required. In many cases,

organizations achieve this through insurance vehicles, but often these capabilities do not provide immediate reimbursement during the time of the crisis. To counteract this, the plan should have an estimate of cost associated with its activities, detailed with timelines of when funds will be required. This allows the finance organization to create a reserve for immediate funding needs and then use insurance vehicles for future funding based on the reimbursement arrangements in the insurance policies that are held for this purpose.

Unavailable Workforce

A recent consideration organizations have added to many business continuance plans is unavailable workforce or workforce quarantine (often associated with pandemic preparation). Many plans prepare for this scenario by establishing remote terminal service capabilities or issuing mobile devices, such as smart phones and laptops, to their workforce. While these capabilities provide the computing capabilities to allow a remote workforce to continue to work, they do not account for the potential lack of available local service provider bandwidth. If an unavailable workforce or workforce quarantine situation affects a region, many more people will be utilizing their remote and home network capabilities during peak working hours. They will also be utilizing more bandwidth than normal, especially if they are utilizing terminal services or application virtualization solutions.

The most effective way to counteract this situation is to develop alternative operating capabilities that require a minimal amount of bandwidth to operate effectively. These capabilities may include text input screens instead of enhanced graphical user interfaces, minimal file transfer or batch file transfers in off-hour windows (such as overnight periods), and the establishment of time-sharing schedules that have a minimum number of users accessing the computing infrastructure at any given time.

DISASTER RECOVERY

Disaster recovery is typically associated with facilities and IT requirements required to recover from a business-debilitating event. Many organizations today are effective in developing secondary physical infrastructure and IT capabilities through mirrored facilities, data replication capabilities, and environments to counteract natural disasters and physical disruptions of their business. Unfortunately, many organizations do not appropriately consider disruptions of a logical nature when developing these capabilities.

Recover Remote or Recover in Place

One of the key considerations in a disaster recovery situation is to decide whether to recover the technical environment in place or revert to remote facilities. The best way to determine the most appropriate option for recovery is to assess the situation based on the recovery point and time objectives for the affected business processes. When possible, it is always preferable to recover in place, because this typically least affects the business operation and cost and is less disruptive to the organization during the return-to-normal operations process.

Overlooked Threat Scenarios

A significant, but often overlooked, threat is one by adversaries who are interested in causing a business disruption by compromising servers or computing capabilities attached to the alternate facility's replication capabilities of an organization. In this scenario, an adversary compromises a system and installs malicious code onto the target system in an inconspicuous way. Instead of activating the code immediately, the adversary waits for an extended period of time (typically three months) before utilizing the code. This amount of time ensures the replication of the malicious code throughout all of the mirror facilities and backup and archive solutions.

When the adversary has a reasonable belief that the code has replicated, they enact it to cause the business disruption. In a typical disaster recovery situation, the organization uses the mirror facility to counteract the attack. Unfortunately, one of the first things most disaster recovery plans require is the reestablishment of networking and redirection of the Domain Name System (DNS) to the alternative site. Even if the organization has not disclosed the physical location of the alternative facilities, the adversary is able to activate the replicated code and continue with the attack as soon as the DNS has propagated.

This situation, as well as other file-level logical attacks, can be counteracted though simple countermeasures such as file integrity checks prior to backup or replication. One method is to use cryptographic hash algorithms to create a library of hash outputs of master production files (such as system files). These files should not change while in development and prior to deployment into the production environment without a change control process. Prior to backing up or replicating data in the production environment, the data in this environment should be hashed using the same cryptographic algorithm and the output compared to the master output that was created earlier. This quickly identifies the modification or corruption

of any of the data to prevent backups or replications, and detects any files that exist in the environment but should not.

Access and Availability of Facilities

Another overlooked scenario in disaster recovery situations is the access and availability to facilities in the event of a disaster situation. Many organizations prepare their primary facilities to be able to function in a loss-of-service situation, such as a power failure or network interruption, by installing fuel-based electrical generators and separate path network connectivity. This allows them to recover in place instead of implementing a remote recovery, which is typically a more expensive and business-disrupting activity. These organizations typically contract with local service providers to provide fuel refill and network repair capabilities. These capabilities work well in the event of a localized incident that affects the organization facility or local facilities, but do not work well in regional incidents. In the case of regional incidents, local or federal law enforcement can choose to declare a state of emergency and prevent these service providers from providing services to replenish fuel supplies or repair telecommunications and network infrastructure.

An equally challenging issue when a state of emergency is declared by local or federal authorities is the ability of staff required to be part of the recovery efforts to either access the facilities or exit the primary facility to access the backup facilities. If the staff required for the recovery is unavailable to reach the secondary facility or leave the primary facility, the organization may not meet the RTOs and RPOs of the plan. In this case, it is important to have competency models developed for all required staff in the recovery efforts.

Availability of secondary facilities can also be a challenge if any organization is contracting with a third party for its disaster recovery data center and networking capabilities. These organizations typically follow a business model where they are prepared to assist in the recovery of a single or small number of clients at the same time. Often, from a facilities perspective, these organizations are not prepared to handle a large number of clients simultaneously.

Even though an organization may have sufficient physical space, computing capabilities and cooling, it often underestimates the need for network bandwidth to the Internet in a disaster. During disasters, many organizations find that they utilize high levels of network bandwidth for activities such as data synchronization, remote workers accessing the environments, and customers and the public attempting to access the computing infrastructure. If the organization uses a shared facility, it be important that there be adequate bandwidth.

Backup of the Backup Facilities

Another consideration is the establishment of secondary recovery facilities if the primary facilities are in use or not available at the time of need. The secondary recovery facilities should be located a great distance from the normal and primary recovery facilities, to minimize the potential involvement of the secondary facilities in the incident. The organization should make the same considerations, as in the primary facilities, in the architecture, design, implementation and operation of the secondary recovery facilities. The secondary facilities typically can be kept in a limited state of readiness (if the business impact, threat and vulnerability analysis allows for this) compared to the primary recovery facilities, to contain costs. This can include remote management, minimal staffing and less frequent data synchronization. It is important to ensure that the organization synchronizes the site with the primary backup facilities on a regular basis, based on the risk tolerance of the organization. It should not be kept in a dark state for a period beyond the recovery point. The organization may not be able to meet time objectives due to preparation activities required to bring the facilities into a fully operational state prior to usage, unless this activity is built into the disaster recovery plan.

TABLE-TOP VS. ACTUAL TESTS

Many organizations test their business continuance and disaster recovery capabilities annually or semiannually. Many of these tests take place using table-top exercises designed to simulate the use of these capabilities without actually enacting them. Unfortunately, many organizations that have had to use their business continuance or disaster recovery capabilities have found that they have failures because they did not perform actual tests. Actual tests have the benefits of identifying weaknesses in the plans, information infrastructure that support the plans and employee readiness. If an organization chooses to utilize actual tests instead of table-top exercises, it is important to do this unannounced, during a time that the test can cause minimal business disruption (e.g., weekends, evenings). If the tests are scheduled, they may not provide an accurate simulation of a crisis, because, due to the notice, the individuals expected to be part of the plans may prepare themselves in advance and ready the capabilities for which they are responsible.

Whether an organization uses table-top or actual tests, the most important activity is a postmortem exercise to identify areas of improvement for the plans and capabilities. These tests typically enumerate areas for improvement to address in the plans. No plan will ever be able to incorporate every situation or threat scenario, but the more the plan is tested, the more it will be effective when enacted in an actual situation.

RETURN-TO-NORMAL CONSIDERATIONS

Return-to-normal procedures and activities are often the most overlooked portion of a business resiliency capability. Crises tend to drive heroic activities and extensive cooperation within organizations to resolve these situations. Organizations often quickly and appropriately perform the use of command and control, business continuance and disaster recovery capabilities. Often, this heightened level of awareness support begins to diminish, if an organization employs these capabilities for extended periods, and the organization often adapts to working in these modes as "business as usual" in relatively short periods. To be successful, it is as important to detail and test the return-to-normal activities as it is to detail the immediate recovery activities for an organization.

Return-to-normal considerations should use the same methods as other elements of the plan, including business impact analysis and the use of RPOs and RTOs.

HETEROGENEOUS APPROACH

One common mistake made in the development of business resiliency capabilities is the development of individual capabilities that are independent of the others. For example, many organizations design disaster recovery plans for complete technical recovery, but do not appropriately account for the logical elements included in the business continuity plans that are connected to the same processes. In many scenarios, an organization will use multiple components of their business resiliency capabilities to recover from and/or remediate a business-impacting event. Neglecting to cooperatively develop capabilities and consider the scenarios heterogeneously will impede the successful resumption of normal business operations with minimal cost and operational impact. Wherever possible, the business resiliency capabilities should share capabilities, methods and procedures to ensure consistency with minimal cost and confusion (**figure 3**).

Figure 3—Interdependency of Capabilities

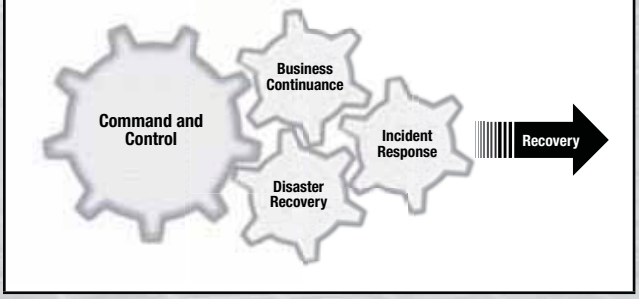

One of the most effective ways to understand where the capabilities are lacking these considerations is to develop test cases that include the utilization of all business resiliency functions during the same incident. These test cases identify gaps where technical and business-logic concepts have not been developed to work together heterogeneously. An example of this kind of test scenario would be a wide-scale disruption of services that includes loss of facilities and disclosure of sensitive information (such as health record information about individuals) as a result of an incident to a public forum (such as a web site) in the middle of the night of a widely observed holiday (such as Christmas).

CONCLUSION

Advanced planning is crucial. There will always be elements of business resiliency capabilities that do not work as planned or scenarios not considered. If the organization develops its capabilities proactively and appropriately, though, it can quickly adapt and ensure minimal disruption at minimum cost.