# DISASTER RECOVERY AND BUSINESS CONTINUITY:

## MORE IMPORTANT THAN EVER IN TODAY'S RISK CLIMATE

Blass, Gerry
President & CEO, ComplyAssistant
gerry@complyassistant.com
Ferrari, Mark
Vice President and CISO, BluePrint HIT
mark.ferrari@blueprinthit.com

# Contents

Disaster Recovery and Business Continuity: More Important than Ever in Today's Risk Climate

# Introduction

Few industries incorporate prevention of and recovery from disaster at the forefront of their day-to-day activities than those in the field of Information Systems. Further, in few industries is it more important to assure the rapid return to systems' pre-disaster capabilities than it is in the field of healthcare.
With the reliance on Electronic Health Records (EHRs), healthcare organizations have low tolerances for system downtimes; they must plan for recovery of business operations as well as the foundational IT systems and the data required to treat patients.

This white paper presents at a high level the integration points of Incident Response, Business Continuity, and Disaster Recovery. Key elements of these plans will be presented to assure a full understanding of the scope of each plan and to help healthcare organizations become more resilient to disruptions and adverse events.

Business Continuity and Incident Response plans must not only account for traditional threats such a fires and floods, but also a new breed of cyber threats such as ransomware and other malicious code. Hackers and other cybercriminals are taking advantage of security weaknesses within healthcare organizations infrastructure. They infiltrate the defenses of the network, encrypt the data and charge the healthcare organization a premium for the key to unlock and retrieve their patient data. Even more insidious are the recent attacks, such as NotPetya, where hackers were not interested in ransom, but rather the destruction of data and disruption to the compromised organization.

In preparing for any type of contingency, three plans are critical to the continuity of operations and restoration of capabilities:

1. Business Continuity
2. Incident Response
3. Disaster Recovery

## Key Distinctions

Business Continuity and Disaster Recovery are concepts that are often used interchangeably in many organizations. There can be a tendency for organizations to revert solely to Disaster Recovery as a first step in responding to an incident, often in lieu of a dedicated Incident Response Plan or Process. Definitions of each of these plans and the synthesized employment of them during a contingency can be lacking in many organizations.

There are many components of planning for, managing, and recovering from contingencies; these processes and plans are complementary and transition to and from one another. However, each type of plan has a very specific focus:

- Business Continuity Plans exist to ensure that critical business functions and processes are sustained during and following periods of degradation.

- Disaster Recovery Plans are information system-focused and are designed to ensure the restoration of target system, infrastructure, or other components as soon as possible following a contingency event.
- Incident Response Plans (also called "Cyber Incident Response Plans") are designed for Information Security personnel to identify, mitigate, and recover from malicious computer events or incidents (NIST, 2012).

The illustration below graphically represents a simple but effective means to conceptualize the employment of each of the three plans being discussed and to highlight the differences among them.

Figure 1

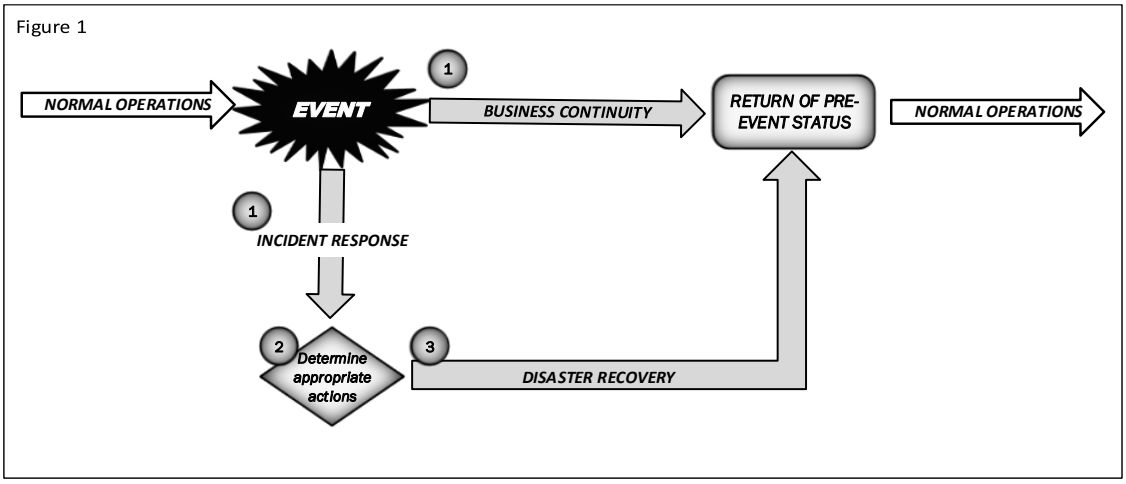


As Figure 1 shows, Disaster Recovery is not an initial step in response to an event.
The immediate steps employed following an event are the employment of Business Continuity measures and the activation of an Incident Response. Consider a clinical system outage as an example:

1. The first steps employed are to continue the delivery of patient care in whatever form that takes (i.e. via Business Continuity measures) while simultaneously performing a degree of Incident Response activity to assess the situation, gathering as much information as possible.
2. Only once the situation or event is understood would a decision point be reached as part of the Incident Response regarding next appropriate steps.
3. That decision point may at that time involve the determination to employ Disaster Recovery activities, among other security response activities.

Overall, the goal is to employ the most effective Incident Response and Disaster Recovery Plans to affect the shortest possible period in which Business Continuity Plans must be utilized. However, Business Continuity Plans must be prepared to account for both short and long-term degrades.

## The Importance of a Business Impact Analysis

When developing effective Disaster Recovery and Business Continuity Plans, the first step should be to conduct a thorough Business Impact Analysis (BIA). NIST Special Publication 300-34, "Contingency Planning Guide for Federal Institutions," identifies three main outcomes of performing a BIA:

1. The correlation of systems with the critical business processes that they support or affect
2. The consequences to the business of a disruption to a system's availability
3. The determination (or confirmation) of priorities for Business Continuity and Disaster Recovery Plans (NIST, 2010)

When performing a BIA, it is important to first identify the correct stakeholders for each system and facilitate a structured discussion/discovery session. This must include personnel from both Information technology and the business side. Information gathering sessions are most effective when both technology and business representatives are in attendance; even the most tenured staff invariably learn something new about their systems during BIA discovery sessions. Key elements to include in these discussions are listed below:

General

- A brief description of the system or application's function
- What critical processes does the system affect? (specific clinical workflows, billing, payroll, etc.)
- Classification of data contained in or processed by the system
- What is the Maximum Tolerable Downtime (MTD) of the system?
- Where is the system hosted? Describe the environment
- What are the dependencies among critical systems?

Contingency status

- Do documented downtime/outage procedures exist for this system?
- How often is the system backed up?
  - *This informs the system's Recovery Point Objective (RPO)*
- Personnel needed to restore the system
- How long does it take to restore the system and its data?
  - *This informs the system's Recovery Time Objective (RTO)*
- Once the system is restored, is there additional work that must be performed by the business prior to resuming system use?
  - *This informs the system's Work Recovery Time (WRT)*

Thorough discovery and analysis of the elements will yield the information necessary to establish system criticality tiers, identify recovery priorities, and understand resource requirements.

In evaluating the metrics of RTO, WRT and MTD, a simple equation must be kept in mind: the time needed to recover a system (RTO) ***plus*** any additional work necessary prior to the resumption of system use (WRT) must not exceed the maximum tolerable downtime (MTD). If the BIA reveals that MTD is exceeded, senior leadership must evaluate compensating controls or additional resources to ensure system availability is returned in a shorter amount of time and within the MTD.

## Business Continuity Planning

As mentioned earlier in this paper, Business Continuity and Disaster Recovery are not synonymous; organizations often erroneously present these two elements in the same plan. Business Continuity involves maintaining business operations during a period in which information systems, personnel, or facilities may be degraded or unavailable. For healthcare providers, the use of downtime plans is the quintessential example of Business Continuity activity. For smaller Business Associates, Business Continuity activity may take the form of staff members working from home if a main office/facility is unavailable (due to fire, flood, etc.). Business Continuity measures are initiated and remain underway as Incident Response activities and, ultimately, Disaster Recovery are pursued.

Specific guidance in constructing a Business Continuity Plan (BCP) can be found in the HITRUST Common Security Framework. Examples of critical elements of a BCP are:

- Assurance of the safety of personnel and the protection of information assets and organizational property
- Identification of critical business functions/processes
- Understanding of the risk(s) the organization faces in terms of likelihood and impact in time, including an identification and prioritization of critical business processes
- Understanding of the impact that interruptions caused by information security incidents are likely to have on the business
- Implementation of preventive and detective controls for critical assets that support key business functions
- Identification of financial, organizational, technical, and environmental resources needed to address information security requirements during times of degradation.
- Testing and updating of the BCP at least annually;
- Ensuring that management of business continuity is incorporated in the organization's processes and structure
- Assigning responsibility for the business continuity management process at an appropriate level within the organization (HITRUST Alliance, 2017)

Once documented, BCPs should be communicated to all appropriate resources within the organization and copies should be stored in areas accessible to staff, but insulated from potential threats. Examples of this may be as simple as a hard copy being stored in a fireproof container or, more effectively, electronic copies being stored in a secure, web-based file sharing solution that assures high availability access for authorized users and devices

## Incident Response Planning

An Incident Response Plan (IRP) is a very specific set of documented processes and procedures to be followed when responding to an information security event. A good IRP enables a consistent and comprehensive response to incidents, which is critical in minimizing loss and disruption of services.

NIST SP 800-61 Revision 1 provides excellent foundational guidance for developing a specific organization's IRP, which should include the following elements or phases:

1. Preparation (documented plans, data gathering forms, controls to mitigate effects)
2. Detection and Analysis
3. Containment, Eradication, and Recovery
4. Post-Incident Activity (Lessons Learned)

NIST 800-61 focuses primarily on responding to attacks or malicious incidents that affect information systems and data. When developing an Information Security IRP, it is important to consider all potential threats to the confidentiality, integrity, and availability of information systems and data, which would have to include both inadvertent, environmental, and malicious events (NIST, 2012)

In its Comprehensive Preparedness Guide for Developing and Maintaining Emergency Operations Plans, the Federal Emergency Management Agency (FEMA) advocates what is termed an "all hazards approach" to planning for contingency incidents. To directly quote this publication:

*"While the causes of emergencies can vary greatly, many of the effects do not. Planners can address common operational functions in their basic plans instead of having unique plans for every type of hazard or threat."*

In its Comprehensive Preparedness Guide, the Federal Emergency Management Agency advocates the adoption of an "all hazards approach" when planning for the management of contingency situations (FEMA, 2010). This approach is well advised when developing plans for Information Security IRPs. The goal is an IRP that accounts for not only the wide variety of known threats, but also for unknown and emerging threats. With respect to Information Security, we know these as "zero-day threats." Whatever the specific attack or source of degradation, the effects are the same: loss of availability of availability personnel, facilities, or systems, and compromises to the confidentiality and integrity of systems and data.

## Disaster Recovery Planning

Disaster recovery planning is the process used by an organization to define and document the detailed processes required to recover after a disaster. While the BCP is used to recover business processes, the Disaster Recovery Plan (DRP) and the Information System Contingency Plan(ISCP) are the two primary plans used to recover IT systems.

### Disaster Recovery Plan

The DRP is designed to orchestrate the recovery of IT systems at an alternate facility after a major disruptive event (e.g. fire, hurricane, flooding) has negatively impacted operations at a healthcare organizations primary site. The DRP provides an enterprise view of recovery sequencing for essential IT services that reflects the business needs of the healthcare organization, system interdependencies, and competing priorities of all stakeholders. It establishes comprehensive procedures to recover critical IT services quickly and effectively following a disaster or extended critical disruption.

### Information System Contingency Plan

An ISCP is an application or system-focused plan containing established, detailed procedures designed to assess and recover an application or system at either the primary site or an alternate site following.an incident that threatens to exceed the RTO/RPO for that application or system. An ISCP may be activated independent from other plans or as part of a larger recovery effort coordinated with a DRP, COOP, and/or BCP. A healthcare organization may have multiple ISCPs to recover individual applications or systems, but there is only one DRP per location. If multiple recovery sites exist, a DRP with recovery instructions specific to each site would be . Table 1 provides a comparison of the DRP and ISCP.

*Table 1. DRP vs. ISCP Comparison*

| DRP | ISCP |
|---|---|
| Orchestrates the recovery of data center operations in a different location | Orchestrates the recovery of a single information system |
| Focused on an individual site | Usually focused on an individual application |
| Executed to move site processing capabilities from a primary site to a recovery location | Executed to restore functionality to an application or system at the primary site (e.g. incident) or at another site (e.g. disaster) |
| Contains site-level procedures to accomplish transition to another site | Contains detailed (e.g. keystroke level) recovery procedures for the individual application |
| Can activate one or more ISCPs for recovery of individual systems | May be activated independently from other plans (e.g. incident) or as part of a larger recovery effort (e.g. disaster) |

## Backup and Recovery Planning

The ability to care for patients is directly impacted by the ability to access medical records and other personal information. If an EHR or other critical system crashes or is impacted by ransomware one must ask, is there a plan to recover the data? Establishing a data backup and recovery plan is required by the HIPAA Security Rule and important for every healthcare organization. The HIPAA Security Rule requires the safeguarding of ePHI, but it does not specify

specific technologies. It does require that backup data is protected from threats, environmental hazards and unauthorized access.

Ensuring an accurate inventory of systems and locations of where data is stored is the first step in developing a data backup plan. Questions must be considered such as: Is the data local (e.g. laptops, servers, storage area networks (SANs) or in the cloud? Are the configurations and data for system and network infrastructure accounted for? Data in all its forms must be backed up and tested regularly to enable the recovery of systems and data at the alternate recovery site.

The RTO/RPO of IT systems will drive the backup frequency and mechanism used to recover the data. For example, if an EHR has a 15 min RTO/RPO, then a real-time backup mechanism is required, because a tape backup would not meet the recovery requirements of the system. The RTO/RPO will help in deciding how to best to back up an organization's systems. EHR data must be backed up more frequently than less critical data; cloud-based solutions are becoming an increasingly popular option for business continuity and disaster recovery functions. If using a cloud based EHR, it is important to understand the cloud provider's recovery plans to ensure access to patient records remains possible in the event of a failure or degrade within the cloud provider's environment. Many cloud providers offer solutions where both the primary and secondary servers are in an active state and immediate failover is available if a single server fails. This configuration provides immediate recovery in the event of a primary server failure, but in the event of a compromised server, both servers will more than likely be compromised, making the use of backup data that much more important.

Backing up data on regular intervals not only protects the organization in the event of an outage or adverse physical event, but is also required to enable the capability to restore data in the event of a cyber event. IT system backups (e.g. EHR, Financial) should be completed on a regular basis and include plans for both system-level and data-level backups. System level and data backups will be performed at different time intervals. For example, system backups will usually include the operating system (OS), current patches and the application(s) running on top of the OS. The system level backup should be conducted before and after an update to the OS or at longer intervals if the hardware, operating systems, and applications change less frequently.

Based on the RTO/RPO of the system, data backups (e.g. PII, ePHI) may require a frequent backup solution (e.g. real time, 15 – 30 minute intervals), while data not considered critical can be backed up less frequently. (e.g. hourly, daily, weekly) There is a cost associated with the frequency of backups, so it is important to identify the criticality of all systems and data as part of a BIA. This classification enables the leveraging of the most cost-effective backup strategy that aligns with the ability to meet the RTO/RPO requirements of the system.

## Types of Backup

There are a few ways a health organization can choose to back up their data. A common back up strategy for critical data is called the "3-2-1" rule. This rule involves having at least 3 copies of data (primary data and two copies), in two different formats (e.g. disk, tape) with at least

one of them off-site. (e.g. off-site tape storage, cloud) When recovering data from an alternate site, if one backup dataset contains errors or is corrupt, there remains another backup copy from which restoration can be accomplished.

When backing up data locally (e.g. Tape, Network Attached Storage (NAS), SANs), the three most common types of backup are full, incremental or differential backups. A full backup is a complete backup of the system and is commonly performed weekly.

An incremental backup involves only the backing up of changes to the data since that last full or partial backup. For example, for an incremental backup, a complete backup would be run on Sunday, but on Monday only the changes that have occurred since Sunday are backed up. On Tuesday, only what has changed since Monday is backed up, and the pattern continues. Incremental backups shorten the time required to run daily backups and tend to be one of the lower cost backup strategies. The disadvantage to incremental backups is that it may take longer to recover data as each day must be restored in order. Also, if any of the days have errors or are missing, a full system restoration may not be possible.

A differential backup is like an incremental backup as it only involves backing up changes in data since the last full backup. Again, if we use the same example, we still run a full backup on Sunday night. On Monday night, we only back up the changes that have happened since Sunday, same as the incremental backup. On Tuesday, when we run our backup, we are backing up everything that has changed since the Sunday full backup, so we are backing up Monday and Tuesdays changes. The advantage when using differential backups is an organization must only retain the full backup and the differential with all the changes. This makes differential backup restoration quicker than restoring an incremental backup.

In addition to backups of an organization's data, it is also necessary to perform backups of operating systems and applications that are in use. If running virtual machines in a data center, backups are usually done at the host level rather than on each individual virtual machine.

HIPAA requires storing ePHI backup offsite, but it is a good practice to store all backup data offsite. It is optimal to store backup data in a location that is outside the geographical area or an organization's primary location. In the event of a natural disaster such as a hurricane or flood, the probability of the same event affecting both the primary and backup locations can be greatly reduced. While encryption is identified as "addressable" with respect to HIPAA, a best practice is to ensure ePHI is encrypted both at rest and in transit.

## Information System Recovery Plans

The following sections provide greater detail into the types of information found in a DRP and ISCP, as well as how they are used in a well-documented recovery plan. This is not an all-inclusive list of each section of the DRP or ISCP, but will help when considering the types of information needed when developing these documents. NIST SP 800-34 can provide additional details and templates; below are some of the types of info contained in an ISCP. (NIST, 2010)

## Concept of Operations - ISCP

As mentioned earlier, an ISCP is application specific plan and can be activated independent of the DRP. For instance, if a data center is running several virtual machines hosting various applications, one might include the recovery of the virtual machines as part of a common infrastructure ISCP and then develop an ISCP for each application running on the virtual server. Any number of ISCPs can be incorporated into a recovery plan, but usually only one DRP will provide the overall prioritization of the recovery sequence of the ISCPs. ISCP recovery is implemented in 3 phases; activation and notification, recovery and reconstitution (NIST, 2010).

### Activation and Notification - ISCP

Contact information, along with roles and responsibilities of each person responsible for recovery of the specific system should be included in the ISCP. This includes internal staff, as well as external recovery support, such as system vendors.

### Recovery - ISCP

The ISCP should provide a complete description of the IT system and the system architecture. It should include a complete inventory of components, as well as identifying all system interconnections, dependencies and any associated plans. The ISCP provides a prioritized and sequenced series of activities required to recover the system. Given that the ISCP is system specific, it should have keystroke level recovery instructions included as part of an appendix to the plan. Depending on the recovery event, primary staff may or may not be available to recover the system. Therefore, detailed instructions along with validation and functional system tests greatly improve organizations' ability to recover systems as quickly as possible.

For servers running virtual machines, the ISCP should include plans to recover the hardware, hypervisor and virtual machine configurations, in addition to any applications and data. Depending on who has responsibility for the environment, organizations may have an ISCP to recover the VM environment and the business organization may be responsible to restore the application and data. It is also important to be aware that each virtual machine may have a different RTO/RPO, so recovery procedures must be prioritized to meet their recovery objectives.

If the IT system is not running on a virtual machine, the recovery plan should identify how to recover the hardware, operating system, business application and its data. It is crucial to ensure that recovery documentation is tied to asset and configurations management processes to ensure the most recent changes are reflected in recovery documentation. Steps to retrieve and recover backup data and system installation media should be included in the ISCP.

### Reconstitution - ISCP

Reconstitution takes place following recovery and includes activities for returning information systems to fully operational states. The recovered system is tested to validate system capability

and functionality, recovery activities are completed, and normal system operations are resumed.

## Concept of Operations - DRP

The DRP is the master document that orchestrates the recovery activities for the network and IT systems. The DRP should provide guidance for the phases implemented as part of the recovery. The incorporation of checklists, rather than lengthy narratives, will help track what steps have been implemented. This is especially useful if the recovery effort takes days and staff rotates throughout the event. Like the ISCP, the DRP recovery is implemented in 3 phases; activation and notification, recovery and reconstitution (NIST, 2010).

### Activation and Notification - DRP

The Activation and Notification Phase defines the activities required to activate the DRP and notify supporting recovery personnel. It defines what constitutes a disaster, as well as identifies who has the authority to declare a disaster and activate the plan. Once the plan has been activated, senior management as well as IT staff need to be notified to support the recovery. The DRP should contain the contact info of key personnel, as well as their role in the recovery, so they can be notified. Automated notification-calling tree tools will help to accelerate the notification process. A common mistake made when developing a disaster recovery plan is assuming that the existing IT staff will be able to get to the recovery site to restore systems. In an emergency, as has been seen with recent hurricanes, IT staff may have trouble getting from their homes to the recovery site. Preparing for this type of scenario may include plans to augment staff with vendors or other trusted sources.

A list of external stakeholders requiring notification should also be included in the DRP. Contact and account info for other external stakeholders such as Telecom/Internet providers, Hardware/Software vendors and Insurance agents will also need to be notified. If leveraging an MSP or cloud provider for recovery services, they would also need to be notified that a disaster has been declared.

Finally, a table outlining roles and responsibilities will help to ensure everyone knows what they are accountable for as part of the recovery. The table should include a leadership or decision-making role, a recovery or coordinator role, and other roles as needed to support disaster recovery operations.

### Recovery - DRP

The recovery phase focuses on implementing recovery strategies to restore essential IT healthcare systems and data capabilities at the recovery site. The DRP should identify a prioritized list of applications (e.g. ISCPs) and the activities that need to occur as part of the recovery process. The BIA will have identified the RTO/RPO of the IT systems, but determining criticality must also consider the various dependencies between IT systems and processes. For example, data that is critical for forming the complete picture necessary for clinical decision processes may originate from multiple systems that provide input to the EHR. Focusing solely

on the most important systems and not understanding these critical relationships can put patient care at risk.

Healthcare organizations may use third party services, such as cloud based EHRs, so it is important to identify in the DRP what data applications run locally and which are external to the organization. Acquisitions, mergers and strategic alliances may also create a need for external connectivity to access critical data. Identifying where all applications reside, along with their data, will help to identify what systems and data need to be restored at the recovery site, as well as identify any external connections that need to be restored to third party vendors or business associates.

For systems hosted outside of an organization's data center, it is important to ensure that the information necessary to access that system from the recovery site is recorded and accessible by all relevant staff. An advantage of having a cloud based EHR or other web-based services is the ability to access the data from any location with internet access. Where dedicated, encrypted links to a third-party provider or business associate are employed it is important to ensure access to the information required to re-establish the connection to/from the recovery site (e.g. IP address, password, encryption keys).

### *Recovery Sites*

The recovery site should, optimally, be located outside an organization's geographical area so that each facility/location would not be impacted by the same adverse event (e.g. Hurricane, earthquake). Common configurations for recovery sites are hot, warm, and cold sites.

A hot site, commonly referred to as an active/active architecture, is a duplicate of the existing systems in the primary data center. It usually runs in parallel and data is updated concurrently at both sites. If the primary site fails, the secondary site takes over processing and no data is lost and no downtime is incurred. This is also the costliest of the recovery options, since systems are duplicated and active at both the primary and secondary sites.

A warm site is also a duplicate of the primary data center but rather than having both sites running and updating concurrently, the secondary site is up, but is not processing data. Network infrastructure and servers are in place and, in the event of a disaster, data would be recovered to the warm site to make it fully operational.

A cold site is simply a facility with power and cooling. Everything required to restore systems must be procured and delivered to the site. (e.g. servers, switches, routers, cables, access points) Recovering to a cold site is generally rare given today's dependency on quick access to electronic health data. If a cold site is part of an organization's plan, it is important to plan for recovery of shared infrastructure, which would have to be operational prior to the recovery of essential systems. A shared infrastructure includes all the hardware (e.g. servers, storage, network devices) and software (e.g. AD, WINS, virtualization) required to support applications. Recovery of systems is not possible until the shared infrastructure is operational. The time

required to recover shared infrastructure should be considered when developing the recovery plan.

Rather than recovering systems locally, it is possible to leverage the services of a Disaster Recovery as a Service (DRaaS) cloud provider to recover essential services. In the event of a disaster, servers and storage can be provisioned in the DRaaS cloud service providers environment to restore essential IT services and it obviates the need for maintaining infrastructure at a recovery site. When normal conditions are restored, the infrastructure in the cloud can be dismantled and normal operations can be moved back to the primary location.

### Reconstitution - DRP

As described above, reconstitution takes place following recovery and includes activities for returning information systems to fully operational states. In the context of the DRP, reconstitution means the systems covered by the DRP are back to a fully operational state.

## Preventive Controls

The aim of DRP is to minimize business downtime so that your technology gets back up and running in the shortest possible time. BCP, on the other hand, looks at maintaining the functionality of the business as a whole and covers a much wider scope, including putting in place preventative controls and managing staff and customers.

It's important for everyone on the team to realize that the BCP is the most important corrective control the organization will have, and to use the planning period as an opportunity to shape it. The BCP is more than just corrective controls; the BCP is also about preventive and detective controls. These three elements are described here:

- Preventive—Including controls to identify critical assets and prevent outages
- Detective—Including controls to alert the organization quickly in case of outages or problems
- Corrective—Including controls to restore normal operations as quickly as possible

Traditionally, a disaster recovery system involved cutover or switch-over recovery systems. Such measures would allow an organization to preserve its technology and information, by having a remote disaster recovery location that produced backups on a regular basis. However, this strategy proved to be expensive and time-consuming. Therefore, more affordable and effective cloud-based systems were introduced.

Some of the most common strategies for data protection include:

- Backups made to tape and sent off-site at regular intervals
- Backups made to disk on-site and automatically copied to off-site disk, or made directly to off-site disk

- Replication of data to an off-site location, which overcomes the need to restore the data (only the systems then need to be restored or synchronized), often making use of storage area network (SAN) technology
- Private Cloud solutions that replicate the management data (VMs, Templates and disks) into the storage domains that are part of the private cloud setup. These management data are configured as an xml representation called OVF (Open Virtualization Format), and can be restored once a disaster occurs.
- Hybrid Cloud solutions that replicate both on-site and to off-site data centers. These solutions provide the ability to instantly fail-over to local on-site hardware, but in the event of a physical disaster, servers can be brought up in the cloud data centers as well.
- The use of high availability systems that keep both the data and system replicated off-site, enabling continuous access to systems and data, even after a disaster (often associated with cloud storage).

In many cases, an organization may elect to use an outsourced disaster recovery provider to provide a stand-by site and systems rather than using their own remote facilities, increasingly via cloud computing.
In addition to preparing for the need to recover systems, organizations also implement precautionary measures with the objective of preventing a disaster in the first place. These may include:

- Local mirrors of systems and/or data and use of disk protection technology such as RAID
- Surge protectors — to minimize the effect of power surges on delicate electronic equipment
- Use of an uninterruptible power supply (UPS) and/or backup generator to keep systems going in the event of a power failure
- Fire prevention/mitigation systems such as alarms and fire extinguishers
- Anti-virus software and other security measures
- Water sensors in the data center ceiling and floors
- Plastic tarps that may be unrolled over IT equipment to protect it from water damage

## Plan Testing, Training and Exercises (TT&E)

It is not enough to develop a plan, put it on the shelf and not look at it again until an adverse event occurs. Once a plan is developed, it must be tested to ensure it meets the RTO/RPO's of the systems being recovered.

Once a plan is developed, a first step would be to gather the plans stakeholders together to walk through the plan on paper. This is referred to as a "table-top" exercise and is the least disruptive way to test a plan. During a tabletop exercise, participants walk through the steps as they are written, looking to uncover holes in the process or steps that have been missed. Organizations would then incorporate the lessons learned and update the plan to reflect those changes.

Once a tabletop exercise has been successfully completed, a next step may be to execute a live exercise that tests all or parts of the organization's contingency plans. Live exercises can be costly and disruptive, however, which is why the employment of tabletop exercises is a recommended first step. Live exercises are vitally important and will uncover additional areas for plan improvement that could not have been identified by the tabletop exercise. Each time recovery staff goes through these exercises, they become more and more familiar with the process. As with the ISCP, lessons learned should be incorporated back into the plan. As part of ISCP recovery testing, it is helpful to include staff members who are less familiar with the system to try and follow the recovery steps. This helps to ensure recovery steps are detailed enough so someone less familiar with the system can effectively follow the process. In the event of a natural disaster that impacts an organization's local area, staff may experience issues in accessing the recovery site; reliance on personnel that are less familiar with the systems may be necessary to complete a recovery. The more detailed the plan, the better chance there is for a successful recovery.

### Plan Maintenance

All the plans that make up an IT Continuity response should be reevaluated annually, at a minimum. Staff may leave or change roles, vendors or models of equipment change, configuration of the equipment can change, these all impact the ability to recover systems if recovery documentation does not reflect these changes. Ideally, plans should be incorporated into an overall service management (SM) framework. (E.g. ITIL) A service management framework defines processes such as asset, change and configuration management. If built into an SM process, once a change occurs, it is not only updated in the database tracking the change, but also any documentation dependent on the change (e.g. DRP, ISCP).

## Conclusion

Events such as increased frequency of catastrophic weather events and healthcare-targeted cyber events has highlighted the importance of having plans in place to facilitate business continuity and the rapid recovery of critical data. Reliance on EHRs have created low tolerances for system downtimes, which could negatively impact quality of care and patient safety. The use of plans to address incident response, business continuity and disaster recovery are all key elements in helping healthcare organizations become more resilient to disruptions and adverse events. It is no longer a question of "if" an organization will be impacted, but "when" it will be impacted. Having recovery plans in place that are tested and updated on a regular basis, will not only improve an organization's chance to recover from a catastrophic incident, but will show due diligence in attempting to protect your patients and their data.

## References

United States, National Institute of Standards and Technology (NIST), Department of Commerce. (August 2012) *Special Publication 800-61 Revision 2 - Computer Security Incident Handling Guide*.

United States, National Institute of Standards and Technology (NIST), Department of Commerce. (May 2010) *Special Publication 800-34 Revision 1 – Contingency Planning Guide for Federal Information Systems*.

HITRUST Alliance. (September 2017) *HITRUST CSF Version 9.*

United States, Federal Emergency Management Agency (FEMA), Department of Homeland Security. (November 2010) *Comprehensive Preparedness Guide (CPG) 101 Version 2 – Developing and Maintaining Emergency Operations Plans.*