cloud security alliance℠

*SecaaS Implementation Guidance*

# Category 9 // Business Continuity / Disaster Recovery

September 2012

# Contents



 

 

# Foreword

Cloud Computing represents one of the most significant shifts in information technology many of us are likely to see in our lifetimes. We are reaching the point where computing functions as a utility, promising innovations yet unimagined. The major roadblock to full adoption of Cloud Computing has been concern regarding the security and privacy of information.

Much work has been done regarding the security of the cloud and data within it, but until now, there have been no best practices to follow when developing or assessing security services in an elastic cloud model—a model that scales as client requirements change.

One mission of the Cloud Security Alliance is to provide education on the uses of Cloud Computing to help secure all other forms of computing. To aid both cloud customers and cloud providers, the CSA SecaaS Working Group is providing Implementation Guidance for each category of Security as a Service, as delineated in the CSA's SecaaS Defined Categories of Service. Security as a Service was added, as Domain 14, to version 3 of the CSA Guidance.

Cloud Security Alliance SecaaS Implementation Guidance documents are available at https://cloudsecurityalliance.org/research/working-groups/security-as-a-service/.

We encourage you to download and review all of our flagship research at http://www.cloudsecurityalliance.org.

Best regards,

| | | |
|---|---|---|
| Jerry Archer | Alan Boehme | Dave Cullinane |
| Nils Puhlmann | Paul Kurtz | Jim Reavis |

The Cloud Security Alliance Board of Directors

# Letter from the Co-Chairs

Security as a Service is a specialized area categorized two years ago as growing rapidly and in unbound patterns. Vendors were struggling. Consumers were struggling. Each offering had its own path. We felt it was urgent to address the needs and concerns common to the implementation of Security as a Service in its many forms.

The Defined Categories of Service helped clarify the functionalities expected from each Category. In this series, we hope to better define best practices in the design, development, assessment and implementation of today's offerings.

We want to thank all of the many contributors worldwide who have worked so hard to produce these papers providing guidance for best practices in Cloud Computing Security. Many have been with the Security as a Service Working Group since the beginning; many others joined in this effort. Each has spent countless hours considering, clarifying, writing and/or editing these papers. We hope they help move forward toward those unimagined innovations.

Sincerely,

Kevin Fielder and Cameron Smith
SecaaS Working Group Co-Chairs

# Acknowledgments

Co-Chairs

Kevin Fielder, Canada Life

Contributors

Steve Markey, N Control
JP Morgenthal, EMC

Peer Reviewers

Henry St. Andre, inContact
Michael Roza, Bridgestone

CSA Global Staff

Aaron Alva, Research Intern
Vicki Hahn, Technical Writer/Editor
Luciano JR Santos, Research Director
Kendall Scoboria, Graphic Designer
Evan Scoboria, Webmaster
John Yeoh, Research Analyst

# 1.0 Introduction

Business Continuity and Disaster Recovery (BC/DR) are the contingency plans and measures designed and implemented to ensure operational resiliency in the event of any service interruptions. BC/DR plans have always been important to any business. As IT systems have become more central to all areas of business, the ability to quickly and reliably recover these systems has become critical. This, in conjunction with increased focus on BC/DR from regulatory bodies, has helped ensure that BC/DR planning and testing is much higher on the agenda for most businesses.

Traditional BC/DR plans have involved storage of backup media at off-site facilities. Should BC/DR processes be invoked, media were restored either to the business's existing infrastructure or to the infrastructure supplied by outsourced BC/DR specialist companies.

Cloud centric BC/DR makes use of the cloud's flexibility to minimize cost and maximize benefits. For example, a tenant could make use of lower specification guest machines as replication targets for data and systems, thus minimizing costs. With the ability to quickly ramp up these machines, and even the number of machines (guests), in a BC/DR scenario, the tenant would have the same benefits as hosting the fully specified systems 24/7. This ability to minimize cost, while also providing full performance should DR be invoked, is a key benefit of the cloud's flexibility and resilience.

When using the cloud for operational processes and/or production systems, an organization's BC/DR requirements must be included in their procurement, planning, design, management, and monitoring of their cloud environments and cloud service providers. BC/DR requirements should be embedded in service or operational level objectives.

Cloud providers typically do not offer a complete BC/DR service, which may include office space, monitors, phones, etc. Cloud providers primarily offer availability for servers/systems and storage, and possibly end-user access via virtual desktops. All aspects of the services to be provided in the case of an invocation should be negotiated and defined fully in their service level agreement (SLA).

## 1.1 Intended Audience

This document has been written for system auditors, system engineers, system architects, system implementers, system administrators, project planners, project coordinators, cloud architects, cloud engineers, and cloud administrators of private/public/hybrid/community cloud consumers.

Section 2 provides a high level overview of BC/DR as a service, the requirements addressed and benefits of BC/DR in the cloud. This section will be particularly useful to C-level executives and project managers wanting a better understanding of this category of service.

Sections 3 and 4 cover implementation considerations/concerns and actual architecture, respectively. These sections will primarily be of use to those charged with designing and implementing this service in their environment.

## 1.2 Scope

This document describes and discusses the following aspects of Business Continuity (BC) and Disaster Recovery (DR) in the cloud:

- Business Continuity & Disaster Recovery BC/DR as a service, including categories such as complete Disaster Recovery as a Service (DRaaS), and subsets such as file recovery, backup and archive;
- Storage as a Service including object, volume, or block storage;
- Cold Site, Warm Site, Hot Site backup plans;
- IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service);
- Insurance;
- Business Partner Agents (business associate agreements);
- System Replication (high availability);
- Fail-back to Live Systems;
- Recovery Time Objective (RTO), Recovery Point Objective (RPO);
- Encryption (data at rest [DAR], data in motion [DIM], field level);
- Realm-based Access Control;
- Service-level Agreements (SLA); and
- ISO/IEC 24762:2008, BS25999, ISO 27031, and FINRA Rule 4370.

Regarding the wider CSA research, this guidance maps to the following areas:

- GRC Stack,
- CCM 1.3,
- CSA CloudCERT, and
- Cloud Data Governance.

This guidance also maps to NIST Special Publication 800-53 Revision 3, *Recommended Security Controls for Federal Information Systems and Organizations.*

Out of Scope:

For the context of this document, BC/DR refers specifically to the recovery of IT systems and services. Anything not explicitly listed above is out of scope.

For the wider recovery of the business, a Business Continuity Management (BCM) strategy comes into play, including plans to provide users with working environments and even travel arrangements when necessary.

Telecommunications deserve mention specifically as out of scope, as telecommunications are so critical and also may be overlooked when planning for BC/DR. The Telecom Working Group is a separate CSA research group dedicated entirely to telecoms in the cloud. For further information on telecoms in the cloud and telecom-specific BC/DR, please refer to https://cloudsecurityalliance.org/research/telecom/

# 2.0 Requirements Addressed

The majority of BC/DR considerations and best practices for cloud-based BC/DR services are very similar to those for traditional non-cloud solutions in terms of what is required, and how the business must prepare prior to implementation.

There are numerous business drivers when planning, executing and testing BC/DR solutions either for traditionally hosted systems or those already in the cloud. Cloud-based BC/DR solutions offer:

- Flexible Infrastructure;
- Secure Backup;
- Monitored Operations;
- Third-Party Service Connectivity;
- Replicated Infrastructure Components, and/or Replicated Data (core/critical systems); and
- Data and/or Application Recovery.

It is likely that the CSP can help the customer with various non-technical issues such as addressing legal issues, data privacy concerns, and operational concerns, which can benefit the executive management team on multiple levels.

The chosen combination of these benefits will influence a business's BC/DR strategy.

In addition to the logical components of BC/DR, there also is a need to consider physical locations, such as alternate sites of operation, geographically distributed data centers/infrastructure and relevant jurisdictions, network survivability, and the incorporation of third-party ecosystems in planning and testing. Conflicting requirements may occur when addressing geographic relocation responses to large scale disasters. Data protection requirements may not allow PII (Personally Identifiable Information) data to leave certain legal boundaries. A careful choice of vendor and a clear understanding of where the vendor's data centers are, along with their data movement policies, can mitigate this.

## 2.1 High-Level Benefits of Cloud-Based BC/DR vs. Traditional Approaches

Some of the key benefits of using a reliable, highly available cloud-based service for BC/DR of systems include:

- **Recovering Workloads to the Cloud** – The CSP (Cloud Services Provider) will host all replicated systems in the cloud, making them available should DR be invoked.
- **Effectively Unlimited Scalability** – A key component of any cloud-based service is its elasticity and effective unlimited scalability. This ensures that the Cloud Services Consumer's (CSC) BC/DR systems will scale up as required.

- **Secure and Reliable Infrastructure** – A CSP's business and reputation are based on services that are secure and available. Most CSPs host multiple infrastructures for hundreds or thousands of customers, so they invest in large scale, resilient, redundant and secure systems and facilities. The scale of these infrastructures far exceeds what most individual businesses are able to support, maintain and afford. By spreading these costs across many hundreds of customers, individual customers are able to enjoy their benefits at a fraction of the costs.
- **Pay per Use** – Consumers pay only for the actual use of the service. This translates to lower service costs during times of normal operation, when minimal capacity and performance are required for storage and replication of systems and data to the BC/DR solution. When events invoke a BC/DR response, requisite supporting services ramp up to meet the need.
- **No (Minimal) Systems BC/DR Expertise Required** – There is a significant reduction in the expertise and effort required on the part of the consumer versus traditional BC/DR solutions. The CSP can manage scaling up capacity, public DNS records, and provide guidance for BC/DR best practices. This is especially true when looking at DR in the cloud for cloud-based systems, as the CSP usually will have automated fail-over facilities protecting their systems and infrastructure.

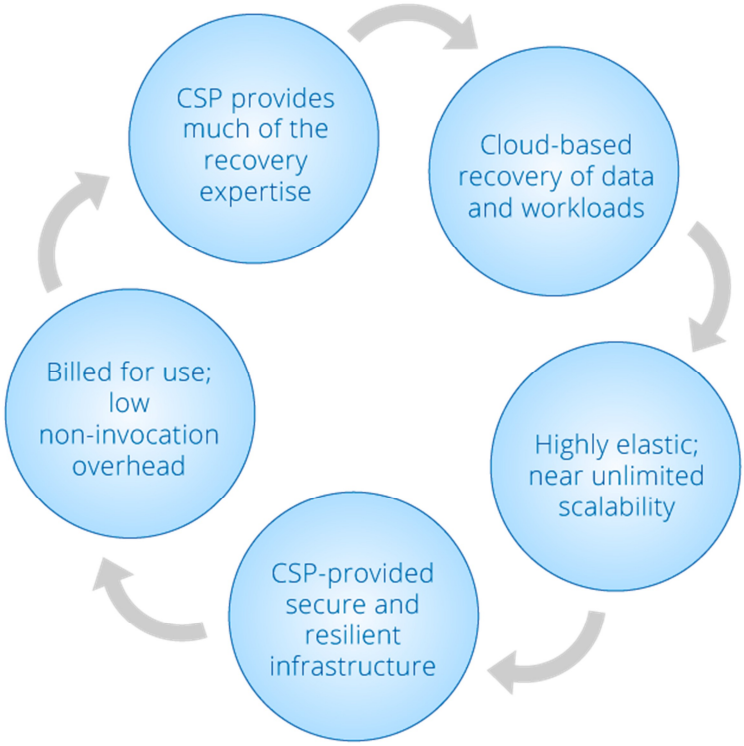


*Figure 1: Benefits of Cloud-Based BC/DR*

# 3.0 Implementation Considerations and Concerns

The considerations and concerns presented in this section should be part of any discussion regarding cloud-based BC/DR.

## 3.1 Considerations

- What is the value of an entity's data to the business and how much will it cost to replace it, if it is lost or stolen?
- After conducting a Risk Assessment/Risk Analysis, what data should go into the cloud?
- How effective is the CSP's own BC/DR Planning?
- How is access to the BC/DR systems and access from the BC/DR systems to third parties managed? If there is an established relationship between the CSP and the BC/DR provider, procedures already may be in place. Do they mesh well with the consumer's needs?
- Data Replication is NOT the same as data archiving or backing up data. Unless snapshots and journaling are used, replication will only replicate whatever is "live" so deletions/errors/corruptions will be replicated to the local Highly Available (HA) instance and the remote BC/DR instance(s). To provide solid DR, backups, whether as snapshots to disk or more traditional backups to tape, also should be implemented.
- Elasticity of the Cloud Provider – Can they provide all the resources if BC/DR is invoked?
- How is DR testing achieved? Does the CSP support DR testing?
- How are physical and virtual workloads accommodated when using cloud for BC/DR?
- How are the DR services accessed if invoked?

### 3.1.1 BC/DR Service Level Agreements (SLAs)

- Determine the best acceptable RTO/RPO (Recovery Time Objective/Recovery Point Objective) for all systems whether Cloud or Non-cloud Apps/Data/Systems.
- Integrate Cloud versus Non-cloud Services into SLAs.
- Establish Performance Requirements.
- Cite Bandwidth Requirements for things like replication to the CSP and client access to hosted services.
- Detail System Requirements for Recovery Workstations.
- Provide a Declaration of DR Procedures and Workflow.
- Specify how failover and failback are handled.

### 3.1.2 Services Already Hosted in the Cloud

Ensuring that services work correctly in performance and scale, as well as functional terms, are critical considerations. Understanding the location of the CSP's DR site(s) also is critical, as it may impact both performance (latency issues) and regulatory compliance, in addition to survivability concerns.

If your services are hosted already with a cloud provider who has multiple data centers in your legal jurisdiction/region, or if there are no regulatory or business reasons why your data cannot move, then setting up BC/DR for your services will, in most cases, be relatively simple and primarily concerned with contractual issues, SLAs and how/when services should fail over. Many of the challenges often faced, such as available bandwidth for replication of data, especially when there is a high rate of data change, will be taken care of by the CSP and their highly scalable infrastructure and services.

### 3.1.3 In-house "Traditional" Non-Cloud Services

Disaster Recovery for traditional, non-cloud hosted systems, when provided as a cloud service, has many of the same issues and concerns as any other DR solution relying on data replication to a remote DR site.

An added challenge when using a CSP over more traditional approaches to DR is that the CSP will only host virtual servers and may place limitations on the Operating Systems it will support. These limitations likely are driven by the underlying hardware placing O/S limitations, even when IaaS services are used. This should be a key consideration when assessing the challenges and effort involved in any project to move BC/DR to the cloud. Considerations regarding physical systems should include:

- Can the system/application be virtualized?
- Is virtualization supported by the vendor?
- Does the system rely on any non-standard physical components (dedicated PCI cards, USB/parallel port dongles, etc.)?
- Are there licensing issues with virtualizing the application?

## 3.2 Concerns

- Jurisdiction of Data
    - Where is your data hosted?
    - Where is your data replicated?
    - Are there any rules governing the movement of your data, such as the EU data privacy laws covering personally identifiable information?
    - Are there legal restrictions on the use of encryption (e.g., requirements that keys be disclosed)?
    - Can the CSP guarantee adherence to any relevant regulations while still guaranteeing resilience and fail over?
- Data Protection/Encryption
    - Is your data encrypted in transit to and from the CSP and between the CSP's sites?
    - Is your data encrypted in transit between guests within the CSP's datacenter(s)?
    - Is your date encrypted at rest?
    - Which algorithms and key lengths are employed?
    - Is the cryptographic implementation trusted (e.g., FIPS 140.2)?
    - Are secure key management processes in place and maintained?
    - Are there adequate data confidentiality, integrity and availability (CIA) protections in place surrounding the virtual or shared servers, including when data is being processed?

- Separation of Duties
    - Are there well documented and agreed roles within the datacenter? (This helps ensure that CSP staff, such as network administrators and DBAs, cannot perform unauthorized acts, either accidentally or intentionally.)
    - Is there secure separation between the various cloud-based services to ensure that those running DR (Disaster Recovery) processes cannot access SIEM (Security Information Event Management) or DLP (Data Loss Prevention) systems, and vice versa?
    - Are policies and procedures in place, documented and up-to-date, to ensure these roles and responsibilities are strictly followed?
- Access Controls
    - Are there strict access controls in place for all customers of the CSP, ensuring that each customer can only access their own systems and data?
    - Are there strict access controls in place, documented, up-to-date and strictly followed for all employees of the CSP?
- Metadata Retention, Separation, and Protection
    - In addition to your actual data and systems hosted in the cloud, are there secure and proper controls surrounding metadata and other data the CSP will capture/log/record during the normal running and hosting of your systems?
- CSP Resilience
    - Does the CSP provide the agreed upon levels of resilience, contained in the SLA, both for systems within the data center and for the datacenters themselves?
    - Are supporting services, such as third-party links to the CSP, including locally hosted CSP infrastructure, as resilient and secure as is required?
- Licensing
    - How is the licensing of any legacy application worded regarding cloud/virtualized deployments? Many legacy applications do not have licensing models that gel well with virtual deployments.
    - What are the vendors' (of your applications and systems) licensing requirements for DR deployments of their solutions?
    - Does the CSP offer any guidance or support regarding licensing concerns or issues (especially if they have existing customers running the same applications), and examples of how licensing concerns or issues have been addressed?

As a general rule, the more current the application/service is, the less likely it is that there will be challenges with virtualizing it or moving it to a cloud-based service. This, however, should not be assumed.

If your business is completely or heavily virtualized already, the move to cloud-based BC/DR will be considerably easier. Many CSPs offer a service where you can directly copy the virtual machine files to the cloud, then spin them up in the cloud and either leave dormant, should DR be invoked, or perform any reconfiguration (IP address, hostname, etc.) to allow them to run and be replicated to, should this be required.

# 4.0 Implementation

This section will provide an overview of various ways in which the cloud can be used to provide BC/DR services to IT systems. Cloud-based services can be used to provide BC/DR to both traditional, on-premise IT systems, and IT systems already based in the cloud.

While the considerations regarding BC/DR process are very similar in both scenarios, the technical questions regarding configuration and replication will have very different responses for services already hosted in the cloud when compared to those hosted in house.

## 4.1 Architecture Overview

### 4.1.1 Services Already Hosted in the Cloud

For services already hosted in the cloud, it is very likely that the CSP will have comprehensive BC/DR capabilities that either come automatically as part of the service or that can easily be added to the service. Many even will include managing tasks like updating public DNS records to ensure services remain available even when automatically migrated to the CSP's DR site.

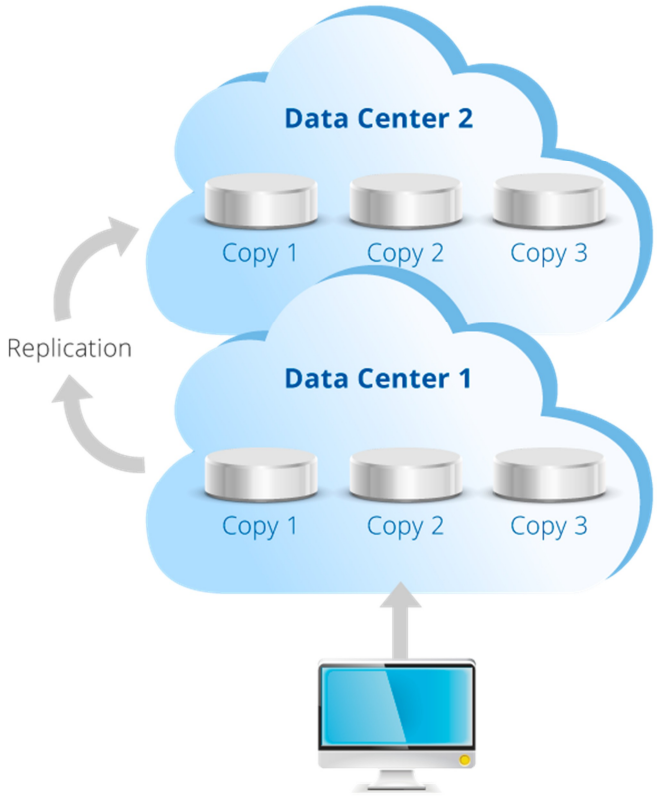


*Figure 2: Cloud to Cloud Disaster Recovery*

Figure 2 shows an example of the resilience to be gained from utilizing a CSP for hosting services and DR. In this instance, the data is replicated across three storage devices at the "live" site (Data Center 1), and also to the DR site, where it is again stored across three devices. Should DR be invoked, the external DNS records will be updated to point to the systems in Data Center 2. Data centers hosted by the same CSP likely will have large interconnects and associated technology, such as compression and packet shaping, that facilitate their ability to replicate large volumes of data.

Various other methods of replication and data transfer may be used by the CSP, ranging from those provided by SAN/storage vendors at an array to array level to much more tailored offerings. For most BC/DR replication requirements, the CSP will have the underlying technology/infrastructure in place.

### 4.1.2 In-House/Traditional Non-Cloud Services

Architecturally, traditional BC/DR will look similar to using a second site you own or a data center in the more traditional hosting/co-location model. Your systems will still be hosted on site and, depending on the model/service chosen, some or all of your data and systems will be replicated to the CSP. In hybrid models, the CSP will host infrastructure at the consumer's site to manage monitoring, replication and fail over. These different scenarios are discussed in the sections below.

Figure 3 shows how traditional systems can be replicated to the cloud for BC/DR, and how they then would re-replicate to the primary facility once the situation is returned to normal.

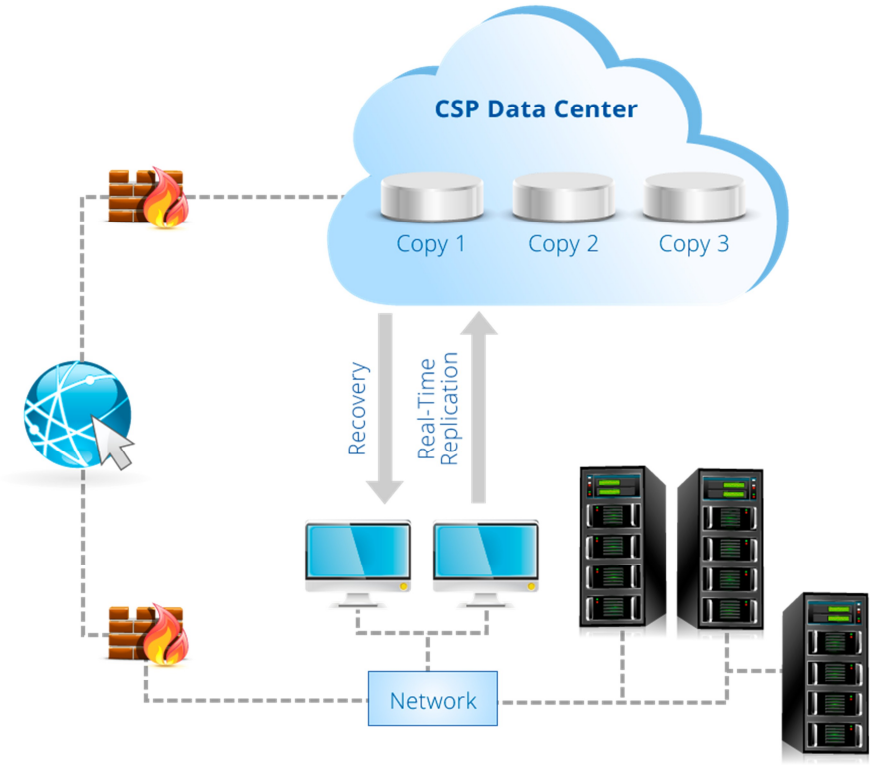


*Figure 3: Traditional System to CSP Replication and Recovery*

## 4.1.3 Types of Service

### 4.1.3.1 Object-Level Storage for Backups

Most of the larger CSPs provide object-level storage via Infrastructure-as-a-Service (IaaS) offerings that can be used for BC/DR. Often, small-to-mid-sized businesses (SMBs) will leverage this model for storing locally executed backups remotely, in lieu of having to use backup tapes, offsite storage vendors, and couriers.

When using the cloud for this method of off-site data storage, consider how systems and services will be restored in the event BC/DR is invoked, as this solution only provides safe offsite storage in a similar manner to manually moving backup tapes off site.

Primary technical concerns include:

- **Replication Management** – A primary difference between replication and backing up is that a backup is static, so you can recover from errors such as corruption and accidental file deletion by reverting to the backed up copy. Replication should employ snapshots at specified intervals to allow for reversion, should files be corrupted or accidentally deleted.
- **Legal and Regulatory Requirements** – Know the legal and regulatory requirements for data retention. Consider whether it remains necessary to perform backups to tape or other long term storage media.
- The manner by which BC/DR will be performed must be considered when using only cloud for "backup as a service," as this effectively only replicates moving tapes off site.

Consider whether traditional BC/DR will be used, and data either accessed from or copied back out of the cloud, or if further cloud-based services will be utilized to provide DR systems and services.

### 4.1.3.2 Block/Volume-Level Storage for Backups

Cloud consumers may use block/volume-level storage for BC/DR purposes when leveraging IaaS compute services. Block level replication is likely more bandwidth efficient than object level replication, so will often require a smaller link to the CSP and therefore be either more feasible, lower cost, or both.

The considerations listed in 4.1.3.1 for object level backups also apply to block level backups.

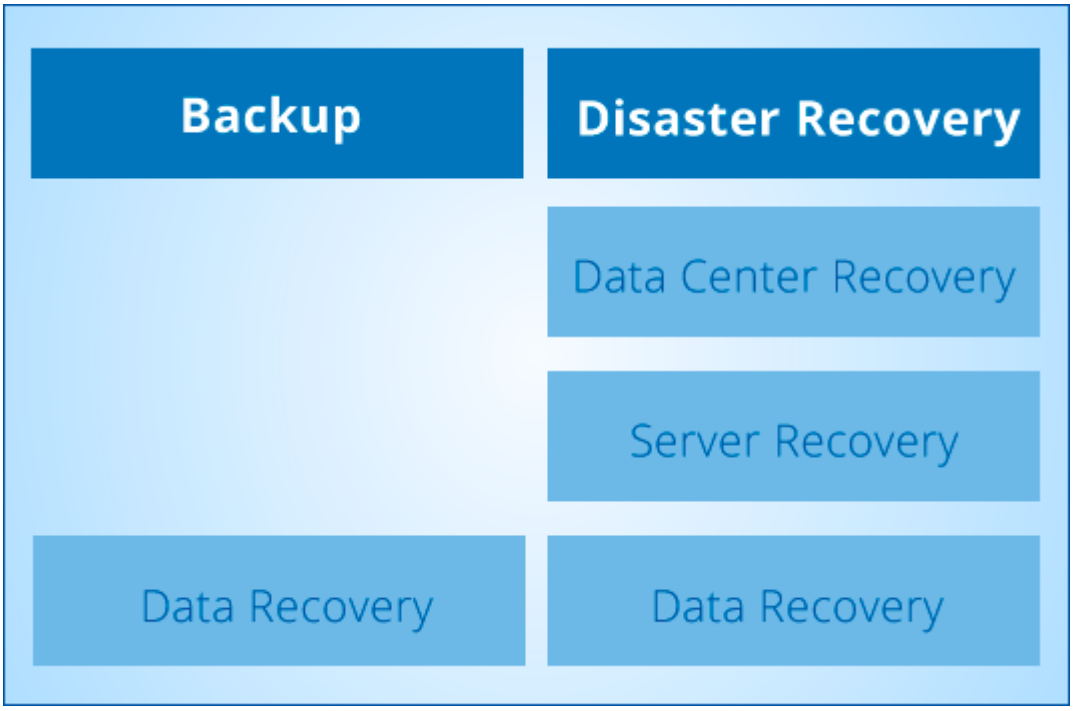


*Figure 4: Highlighting Differences between Backup and Recovery*

The diagram above shows that there is much more involved in recovering working systems than just having a safe copy of the data. This should be a prime consideration when looking at cloud services. Backup to the cloud is just that. If you want the service to include recovery services so systems are live if BC/DR is invoked, you will need to consider further services to provide recovery in addition to backups.

### 4.1.4 PaaS – Replicated DB Instances

When using PaaS or database as a service (DBaaS) for a production environment, it behooves a consumer to leverage a replicated environment via geographic separation.

When replicating databases in the cloud, an obvious concern is cost. However, for a production system of a transactional nature, replication may be more cost effective than engaging in a hot site contract with a continuity provider. Concurrently, a replicated system may provide load balancing, as well.

Databases in the cloud also can be provided via an IaaS style implementation. For example, by installing an instance of the DB software on an IaaS O/S instance, the client may be able to utilize any database software and version, rather than relying on the existing PaaS offerings. In this instance, the client would be responsible for managing replication via technologies such as mirroring or log shipping.

Although this discussion specifies databases in a PaaS context, PaaS services cover a wider range of platforms than just databases. The list continues to grow and covers various development frameworks including PHP, Java and .Net; cloud only software platforms; and a variety of hosting and database services.

### 4.1.5 IaaS – Replicated OS Instances

In a replicated OS scenario, the CSP will provide the underlying infrastructure services only. As with all IaaS solutions, the onus is on the customer to securely protect and manage their servers, systems and data.

If using IaaS for DR from traditional on premise systems, the customer will be responsible for configuring the IaaS systems, setting up replication, updating public DNS records, and ensuring the data and systems are functioning correctly. In short, outside of the underlying infrastructure, the customer is largely responsible for ensuring their DR services function as expected should DR be invoked. Apart from the elasticity and other generic cloud benefits, IaaS is the closest cloud solution to traditional DR solutions in terms of effort required and flexibility.

The benefits of an IaaS solution are that it offers the greatest flexibility in terms of what is run on the CSP systems, with the only limitations coming from the CSP side. Limitations from the CSP tend to be limited to which operating systems/versions are compatible or permitted.

IaaS is often the ideal solution for businesses whose systems are less compatible with PaaS or SaaS offerings. As with all cloud offerings, once the customers systems are hosted with the CSP it is likely that the CSP will offer automated replication of the IaaS systems within the data center and between their data centers, for resilience.

### 4.1.6 PaaS – Application Failover

Platform-as-a-Service (PaaS) environments offer businesses scalable application containers that can operate across a multitude of infrastructures. In contrast to IaaS, PaaS offers greater built in resilience for applications, because it creates a further abstraction away from guest operating system environments. PaaS provides a greater level of homogeneity, but often at the cost of requiring applications to be modified for operation within the specific PaaS environment.

Due to the greater abstraction, PaaS has the added benefit of more easily deploying across cloud service and deployment models. Some IaaS environments force a specific guest operating system architecture upon the user, which makes it difficult to share a single image across multiple service providers. The PaaS abstraction allows a single application image to operate on all these IaaS, thus allowing cloudbursting, on-demand scaling, service provider redundancy, and geographic independence; all which enable greater levels of business continuity and disaster recovery at much lower costs than traditional BC/DR using redundant owned or leased data center environments.

Many PaaS environments can naturally load balance and scale across PaaS node instances that are geographically dispersed. As with any redundant application architecture, load balancing, DNS and identity management services need to be able to failover accordingly. Assuming these services are also highly available, the interoperable nature of PaaS means that the application will be accessible.

### 4.1.7 SaaS – Replicated File-Level Storage for Clients/Hosts Backups

CSPs typically have automated methods to backup and sync data to offsite locations. The only requirement for an administrator or user is to install and configure the solution for the backup jobs; backups then are conducted at specified intervals. These solutions are convenient and cost-effective. However, the need for a consumer to test the backups is extremely important. The backup process should be tested every six months.

### 4.1.8 SaaS – Replicated File-Level Storage for File Shares

Much like replicating file-level storage for client systems, an organization may engage in using a CSP to back up user or organizational file shares via a software client/agent. This solution works best for SMBs, as large enterprises will have a much larger data footprint, thus adding costs and time for the backups to execute.

### 4.1.9 SaaS – Data Export

To negate vendor lock-in, SaaS cloud consumers should validate that a cloud service provider (CSP) provides data export functionality. Consumers may use this functionality to export data to another geographic location and/or CSP, to ensure continuity in the event of an outage. Furthermore, consumers can use this model to ensure that proper and agreed-upon recovery time objectives (RTO) and recovery point objectives (RPO) are met. For portability and interoperability purposes, the consumer should ensure that the SaaS CSP provides a consistent standard for data export, and one that is compatible with the customer's and other cloud vendors' solutions.

Data export processes can be optimized by leveraging open standards such as ASCII, as well as by leveraging a data dictionary. A common data model, data dictionaries or master data management (MDM) will allow the consistent and automated exporting of information between geographically separate systems. This expectation is predicated on the expectation that the CSP's and the consumer's systems have data import capabilities as well.

### 4.1.10 IAM – Leveraging Federated Identities for BC/DR

Federated identity management allows organizations like enterprises and service providers to securely exchange user information across partners, suppliers and customers. By utilizing standards-based methods, identity federation can reduce costly repeated provisioning, security loopholes and user inconvenience, which are often the consequences of rigid, proprietary, tightly-coupled application architectures. Organizations that have deployed federated identity management software remove barriers from logging in, improve collaboration with partners, enhance customer service, accelerate partnerships and alliances, reduce costs associated with integrating to outsourced services, and free themselves from large vendor specific, all-encompassing systems.

Refer to the *CSA SecaaS Category 1: Identity and Access Management Implementation Guidance* for further details on Federated Identities.

### 4.1.11 Recovery as a Service

Recovery as a Service (RaaS) is effectively a catchall term for performing BC/DR in the cloud. It usually refers to a considerably more complete service than the other options mentioned and may include replication to the CSP, alerts of issues, and potentially automated fail over options, including component and whole site failover. This service also often will include error reporting, compression and replication, etc.

RaaS will encompass many options, including replication, compression, encryption, fail over, IaaS/PaaS as required, possibly wrapped up into a single "service." As such, it is potentially more costly than doing cloud BC/DR yourself, but likely considerably easier in terms of ensuring a complete service and getting the most support from the CSP.

RaaS primarily relates to performing BC/DR for traditional on premise systems. Many CSPs offer very extensive BC/DR services for their systems when your primary services are hosted in the cloud.

## 4.2 Guidance and Implementation Steps

### 4.2.1 Replication of Files and Data

Replication, from a BC/DR perspective, copies an entire logical object to the DR site for the purposes of system/application/data recovery. Depending upon the solution used, it may also provide off site backup functionality as well. Replication can be done in multiple ways, including file level synchronization, file copies, byte or block level.

Files are usually replicated or synchronized across different geographic locations for continuity purposes. Note that synchronization is more bandwidth-efficient than complete file replication; the entire file is not copied over, just the modifications. The all or nothing nature of file level replication can lead to higher bandwidth requirements and a loss of efficiency. However this solution can be easier to implement.

Block level replication only replicates block level changes on the storage device and is the most efficient in terms of bandwidth requirements. However, block level replication may not be compatible with systems that need to be live at the DR site, especially if immediate fail over is required. An example would be a database, where it is likely best to rely on in built-in replication technologies, such as mirroring.

In order to minimize bandwidth and off-site storage requirements, consideration and assessment should be undertaken to determine exactly which data must be replicated for DR, rather than just replicating all data.

Note that replication is not a replacement for backing up data, as it will replicate any corruptions/errors/unintended file deletions, etc.

When it comes to replicating data for DR, it is critical to understand:

- The systems/applications for which data is being replicated, including any dependencies (if data in system X is from xx time, then data in system Y must be from xx time also).

- The mode of DR: if systems are being replicated as live systems, or if only data is being replicated, and the DR systems will be "turned on" if DR is invoked.
- What the rate of change in the systems is, and how much bandwidth will be required.
- How near to "real time" the replication must be.
- What technologies are supported by the CSP (e.g., can SAN to SAN type replication be performed or are other software or devices required?).

### 4.2.2 Clustering

Systems are clustered when they are joined in a pool of hardware and software resources that act as one endpoint. Clustering allows for the failover of hardware and/or software within this pool for continued operations. Note that clustered systems are usually constrained to one geographic location, so location-based failures (e.g. power outages/fluctuations) may affect a clustered solution.

Given the virtual nature of cloud-based services, implementing relatively complex cluster solutions within this environment may not be as worthwhile as in the "physical" domain. Hosting the service on multiple IaaS instances, utilizing horizontal scaling and high availability technology, likely provides more resilience than a traditional cluster.

### 4.2.3 Server-Side Encryption

Server-side encryption enables cloud consumers to be sure that information uploaded by end-users is protected once the data is within the walls of the CSP. This solution is a final safeguard for those organizations that require protection for their data at rest (DAR).

Depending upon the offering consumed, the customer will have differing levels of control in this area. For example, in an IaaS solution, the customer likely will be able to use whichever encryption method they want, but in a SaaS solution the encryption level will be defined by the CSP.

### 4.2.4 Client-Side Encryption

For cloud consumers located in heavily regulated jurisdictions, or for those in heavily regulated industries, the use of client-side encryption is recommended to ensure the confidentiality and integrity of your information. This can be done via data at rest (DAR) encryption solutions. For further information on Encryption, please refer to the CSA SecaaS Category 8: Encryption Implementation Guidance.

### 4.2.5 Transport/DIM Encryption

CSPs should support the use of transport layer security (TLS) and secure sockets layer v3 (SSLv3) protocols for secure communications. TLS is recommended, as this is superseding SSL, although SSLv3 is still in wide use and likely acceptable to many businesses. When consumers upload content via electronic data interchange (EDI) methods, the CSP should allow for the use of secure shell (SSH) technologies. FTP is widely considered obsolete.

#### 4.2.5.1 Encryption Considerations

If data is to be processed in the cloud, it will at some point have to become unencrypted in order to be processed, even if it is then re-encrypted for transport or storage.

In order to ensure that the levels of encryption are adequate for the sensitivity of the data and the length of time that the data must remain secure, it is essential to understand the requirements for data handling. Session keys for a web application session that expire after a specified period only need to be protected in transit with relatively low level encryption. Compare this to data backups that may need to be secured for many years.

The value of data to attackers, and how long this value lasts, also should be considered when designing the encryption requirements for different data and components of the solution.

### 4.2.6 ADFS

For those organizations that have implemented a Microsoft Windows platform for their client and server-side architecture, active directory federation services (ADFS) is an option for single sign-on (SSO) mechanisms. ADFS merges the nuances of the lightweight directory access protocol (LDAP) with an alternative to security assertion markup language (SAML). Through the use of AD group policy objects, which enable role-based access controls (RBAC), the organization can have a more comprehensive ability to control authentication and authorization.

When using the cloud for production and or DR services, the company's domain can be extended into the cloud to enable unified access and authentication across both cloud- and non-cloud-hosted services, along with providing DR and resilience for domain-based services. This can also true for the customer's DNS namespace.

### 4.2.7 SAML

SAML (Security Assertion Markup Language) is an XML-based authentication and authorization standard for single sign-on (SSO) purposes. This platform is consumed by a service provider and offered by an identity provider on behalf of a cloud consumer. As SAML is web-based, this solution works through a user's web browser.

### 4.2.8 Transaction Journaling

For both continuity and tracking purposes, it is essential that CSPs use the journaling features available on both appliance and software-based systems. Journaling enables an audit feature that allows systems to be restored to a particular point and/or provides a detective safeguard for auditors from an information assurance perspective.

### 4.2.9 Realm-Based Access Controls

While many organizations leverage role-based access controls for user authorization, others use realm-based controls. The difference is that realm-based authorization is based on domains instead of user roles. For

example, a large organization that has many different domains for their business units may leverage those domains for access controls.

## 4.2.10 BC/DR Economics

As organizations use the cloud and CSP-based services for their BC/DR needs, cost effectiveness will play a key role in vendor and solution selection. As additional players in the market are introduced it is essential that cloud consumers identify the most cost-effective CSPs and RTO/RPO levels.

Most cloud consumers are familiar with dealing with continuity providers on a contractual basis for BC/DR services through cold sites, warm sites or hot sites. However, the cloud has changed that paradigm as consumers have additional options available. To ensure that organizations realize the cost savings they desire for choosing CSP services versus a traditional continuity provider's, it is essential to have the proper monitoring and administration controls in place.

## 4.2.11 BC/DR for Traditional On-Premises Infrastructure to the Cloud

A CSP may offer services that provide DR capabilities for on-site IT systems. The architecture will depend on the RPO (Recovery Point Objectives) and RTO (Recovery Time Objectives) of the cloud customer. The overall principles will be the same as traditional BC/DR, and will consist of:

- Configuring the relevant local systems to replicate data.
- Configuring systems and storage at the CSP to be replicas of the live systems, likely with considerably lower specifications.
- Defining the required specifications at the CSP should DR be invoked. Ensure the CSP has enough elasticity to provide the increases in performance and scale within the required DR time lines.
- Replication of systems and data to the CSP. Note that it is imperative to replicate systems and data, not just data. While it is very possible, and in some cases reasonable, to just replicate data in order to have off site copies, this is analogous to storing tapes off site. Such a plan can be used to support DR, but does not in itself provide DR capabilities. Planning and attention should be paid to the rate of data change, and therefore bandwidth required, between the customer and the CSP in order to replicate systems for DR.
- Working with any third parties to ensure the cloud hosted systems can access the third party, if DR is invoked, as per locally hosted systems.
- Testing access to the systems in the CSP from any relevant locations, including third parties, customers, employees (both from the live site, and any BC sites).
- Testing the actual DR systems themselves to ensure they work as expected and process correctly.
- Performance testing. There often are concerns regarding the performance of virtualized systems, so performance tests should be carried out to confirm to the business that the chosen solution will perform as designed, should DR be invoked.
- Ensuring any legal/licensing concerns are addressed.
- Ensuring all applications/systems are supported in a cloud/virtual environment.

- Potentially designing alternative solutions for applications/systems that will not function in cloud / virtual environments. Work with vendors of these systems on potential solutions to enable virtualized/ cloud-based DR.

## 4.2.12 BC/DR for Cloud Service to Cloud Service

If DR is required when using cloud services to provide the Infrastructure/Platform/Software as a Service for production environments, ensure that services are backed up and/or replicated to another site.

Many CSPs offer BC/DR capabilities for systems they host, via replication to another datacenter either by default or as an extra service, or a combination of the two. Some platforms replicate all data to an additional datacenter by default, but not the whole system in a live state. Maintaining live systems at the DR site may be added as additional services, usually incurring additional fees.

In terms of architecture, DR when hosted in the cloud is considerably simpler than when hosted locally. The CSP will manage replication, bandwidth requirements, storage, networking, DR hosts/systems, etc.

Consider where the DR site is – data movements in and out of certain legal boundaries are controlled by regulations. For example, personally identifiable information held in the EU about EU residents must not leave the EU. Choosing a CSP with multiple datacenters in the same legal boundary may be a requirement.

Depending on user/customer access to the systems, proximity to users or customers may be a concern if large data volumes are regularly transferred.

Latency and other constraints must be considered regarding ensuring users access, access to and from third parties, etc.

When planning any kind of cloud BC/DR service, the customer should also reference CSA SecaaS Category 8: Encryption Implementation Guidance to ensure the data is securely transported and stored.

## 4.2.13 Failover Automation

Cloud consumers need to be cognizant of what failover functionality a CSP's environment offers and how to automate that task. Many IaaS and PaaS providers should have this functionality available at an additional cost. Once a consumer has added automated failover capabilities to their environment it is extremely critical that they occasionally test this functionality to ensure that it works properly.

The failover process will vary considerably depending upon the service model implemented:

- **SAAS** – When utilizing a SAAS-based system, failover likely will be entirely handled by the CSP, including ensuring data is consistent across components of the system, and making requisite external changes, such as updating DNS records, as part of the service offering.
- **PAAS** – The CSP will provide failover for the platform along with applications and data residing on the platform. Areas such as re-pointing other application components and other parts of the system will either be the responsibility of the customer, or would need to be added to the basic contract.

- **IAAS –** An IAAS deployment, while the most flexible option for the customer, is the option where the onus for implementation and management resides with the customer. The CSP will replicate all data and O/S instances, but bringing DR systems online, ensuring consistency of replicated data, etc., will fall to the customer.

## 4.2.14 Final Notes

### 4.2.14.1 Traditional Non-Cloud Systems Recovery – Recovery as a Service

For complete BC/DR for systems/services from a CSP, consider Recovery as a Service (RaaS). RaaS providers typically offer a total solution. The CSP will work with you, potentially even hosting hardware at your local site, and offer fully or partially automated fail over services for requirements from individual systems to the whole site. The CSP also will work with you on the failback process which is often overlooked and is as critical as the initial fail over process in terms of enabling the business to return to normal, once the disaster has been resolved.

From an architectural perspective, RaaS will look similar to a combination of the solutions incorporating replication of systems and data, with additional wrappers for better monitoring, and assistance with fail over, etc.

### 4.2.14.2 Cloud-Based Systems Recovery

For businesses whose systems are already hosted in the cloud, it is recommended that they leverage the inbuilt BC/DR capabilities provided by the CSP. Most CSPs will offer, at a minimum, automated replication and the ability to automatically (or at least quickly) fail over systems, even when using an IaaS-based solution. This option means the customer must contractually specify the resources required at the DR site, define the level of ramp up if BC/DR is invoked, and understand the repercussions of failing over to another location, including its effects on customers and partners. The big benefit of the RaaS solution is that there is little or no actual implementation or architectural work required of the consumer. Requirements of the customer are limited to contractual design, scaling and testing.

# 5.0 References and Useful Links

http://en.wikipedia.org/wiki/Disaster_recovery

http://www.silicon.com/management/cio-insights/2010/09/30/cloud-computing-is-it-ready-for-disaster-recovery-39746406/

http://blogs.forrester.com/rachel_dines/11-08-29-disaster_recovery_meet_the_cloud

http://www.usenix.org/event/hotcloud10/tech/full_papers/Wood.pdf

http://www.finra.org/Industry/Issues/BusinessContinuity/

http://docs.oracle.com/cd/E19857-01/817-1831-10/agsecure.html#wp1028500

http://www-935.ibm.com/services/uk/en/it-services/vsr_whitepaper.pdf

http://www.hknet.com/en/ict-solution/cloud-hosting/virtual-server-hosting/whats-on.html

http://blogs.forrester.com/rachel_dines/12-03-22-cloud_based_disaster_recovery_demystified

http://www.cio.com/article/704036/How_the_Cloud_Democratizes_and_Complicates_Disaster_Recovery_

http://www.informationweek.com/news/services/disaster_recovery/232800112

http://searchdisasterrecovery.techtarget.com/feature/Disaster-recovery-in-the-cloud-explained

http://www.forbes.com/sites/microsoft/2011/05/12/disaster-recovery-in-the-cloud/

 