

Cloud Incident Response (CIR) Framework



The permanent and official location for the Cloud Incident Response Working Group is <https://cloudsecurityalliance.org/research/working-groups/cloud-incident-response/>.

© 2021 Cloud Security Alliance – All Rights Reserved. You may download, store, display on your computer, view, print, and link to the Cloud Security Alliance at <https://cloudsecurityalliance.org> subject to the following: (a) the draft may be used solely for your personal, informational, non-commercial use; (b) the draft may not be modified or altered in any way; (c) the draft may not be redistributed; and (d) the trademark, copyright or other notices may not be removed. You may quote portions of the draft as permitted by the Fair Use provisions of the United States Copyright Act, provided that you attribute the portions to the Cloud Security Alliance.

Acknowledgments

Lead Authors:

Soon Tein Lim
Alex Siow
Ricci Jeong
Michael Roza
Saan Vandendriessche

Key Contributors:

Aristide Bouix
David Chong
David Cowen
Karen Gispanski
Dennis Holstein
Christopher Hughes
Ashish Kurmi
Larry Marks
Abhishek Pradhan
Michael Roza
Ashish Vashishtha

Reviewers:

Oscar Monge España
Nirenj George
Tanner Jamison
Chelsea Joyce
Vani Murthy
Sandeep Singh
Fadi Sodah

CSA Global Staff:

Hing-Yan Lee
Ekta Mishra
Haojie Zhuang
AnnMarie Ulskey (cover design)

Special Thanks:

Bowen Close

About the Cloud Incident Response Working Group (WG)

With today's emerging and rapidly evolving threat landscape, a holistic cloud incident response framework that considers an expansive scope of factors for cloud outages is necessary. The Cloud Incident Response (CIR) Working Group (WG) aims to develop a holistic CIR framework that comprehensively covers fundamental causes of cloud incidents (both security and non-security related) and their handling and mitigation strategies. The aim is to serve as a go-to guide for cloud users to effectively prepare their detailed plan to respond and manage the aftermath of cloud incidents. The CIR is also a transparent and common framework for cloud service providers to share their cloud incident response practices with cloud customers. This framework's development includes imperative factors of cloud incidents such as operational mistakes, infrastructure or system failure, environmental issues, cybersecurity incidents, and malicious acts.

Table of Contents

1. Introduction	7
Purpose	7
Target Audience.....	7
2. Normative References.....	8
3. Definitions.....	9
4. CIR Overview	10
5. CIR Framework	13
5.1 Phase 1: Preparation and Follow-on Review.....	13
5.1.1 Documentation	17
5.2 Phase 2: Detection and Analysis.....	17
5.2.1 Inducement	17
5.2.1.1 Cause of Cloud Incident.....	17
5.2.1.2 Signs of an Incident	18
5.2.1.3 Common Sources of Precursors and Indicators	19
5.2.2 Incident Analysis to Determine Impacts	19
5.2.2.1 Incident Analysis	19
5.2.2.2 Incident Notification.....	20
5.2.2.2.1 Incident Notification Timing	20
5.2.2.3 Incident Impacts.....	21
5.2.3 Evidence Gathering and Handling.....	23
5.3 Phase 3: Containment, Eradication, and Recovery.....	23
5.3.1 Choosing a Containment Strategy.....	25
5.3.2 Eradication and Recovery	25
5.4 Phase 4: Post-Mortem	26
5.4.1 Incident Evaluation	26
5.4.1.1 Incident Evaluation Metrics.....	26
5.4.1.2 Incident Classification	27
5.4.2 Incident Closing Report.....	28
5.4.2.1 Lessons Learned	29
5.4.3 Incident Evidence Retention.....	30

6. Coordination and Information Sharing	31
6.1 Coordination	32
6.1.1 Coordination Relationships	32
6.1.2 Sharing Agreements and Reporting Requirements.....	33
6.2 Information-Sharing Techniques	33
6.3 Granular Information Sharing	33
6.3.1 Business Impact Information.....	34
6.3.2 Technical Information.....	34
6.3.2 Technical Information.....	34
6.4 Table-Top Exercises and Incident Simulations.....	35
7. Summary	36

1. Introduction

In today's connected era, a comprehensive incident response strategy is an integral aspect of any organization aiming to manage and lower its risk profile. Many organizations and enterprises without a solid incident response plan have been rudely awakened after their first cloud incident encounter. Significant downtime can occur for numerous reasons, such as a natural disaster, human error, or cyberattacks. A good incident response plan helps ensure organizations are well-prepared at all times. There are, however, many considerations when it comes to incident response strategies for cloud-based infrastructure and systems, due in part to the nature of its shared responsibility.¹

Incident response frameworks have already been documented in many governmental and industry guidelines, such as the *NIST 800-61r2 Computer Security Incident Handling Guide* or *SANS Institute Information Security Reading Room Incident Handler's Handbook* for traditional on-premises information technology (IT) environments. However, when a cloud computing environment is incorporated, the roles and responsibilities defined in the traditional incident response frameworks must be revised and refined to align with the roles and responsibilities of cloud service providers (CSPs) and cloud service customers (CSCs) for different cloud service models and deployment models.

Purpose

This document seeks to provide a Cloud Incident Response (CIR) framework that serves as a go-to guide for a CSC to effectively prepare for and manage cloud incidents through the entire lifecycle of a disruptive event. It also serves as a transparent and common framework for CSPs to share cloud incident response practices with their CSCs.

Target Audience

The key beneficiaries are CSCs. This framework guides CSCs to figure out their organization's security requirements and thus opt for the appropriate level of incident protection. Through this, CSCs can negotiate with CSPs or select security capabilities that are made-to-measure—providing a clear understanding of the division of security roles and responsibilities.

¹ Cloud Security Alliance, Cloud Incident Response, <https://cloudsecurityalliance.org/research/working-groups/cloud-incident-response/>

2. Normative References

The CIR Framework refers to several industry-accepted standards and frameworks to plan and prepare for cloud incidents, mitigation strategies, and postmortem processes.

- CSA Security Guidance For Critical Areas of Focus In Cloud Computing v4.0
- NIST 800-61r2 Computer Security Incident Handling Guide
- ITSC Technical Reference (TR) 62 – Cloud Outage Incident Response (COIR)
- FedRAMP Incident Communications Procedure
- NIST 800-53 Security and Privacy Controls for Information Systems and Organizations
- SANS Institute Information Security Reading Room Incident Handler's Handbook
- ENISA Cloud Computing Risk Assessment

Figure 1 shows the relationship between the CIR phases and the primary references.

Phase 5.1 Preparation	Phase 5.2 Detection and Analysis	Phase 5.3 Containment, Eradication and Recovery	Phase 5.4 Postmortem
CSA Sec. Guidance v4.0 9.1.2.1 Preparation	CSA Sec. Guidance v4.0 9.1.2.2 Detection and Analysis	CSA Sec. Guidance v4.0 9.1.2.3 Containment, Eradication, and Recovery	CSA Sec. Guidance v4.0 9.1.2.4 Postmortem
NIST 800-61r2 3.1 Preparation	NIST 800-61r2 3.2 Detection and Analysis	NIST 800-61r2 3.3 Containment, Eradication, and Recovery	NIST 800-61r2 3.4 Post-Incident Activity
TR 62 0.1 Cloud Outage Risks	TR 62 4.2 COIR Categories	TR 62 5.2 During Outage: CSC	TR 62 5.3 After Outage: CSC
FedRAMP Incident Comm. Procedure 5.1 Preparation	5.1 Before Cloud Outage: CSC	6.2 During Outage: CSP	6.3 After Outage: CSP
NIST (SP) 800-53 r4 3.1 Selecting Security Control Baselines	6.1 Before Cloud Outage: CSP	FedRAMP Incident Comm. Procedure 5.3 Containment, Eradication, and Recovery	FedRAMP Incident Comm. Procedure Post-Incident Activity
Appendix F-IR IR-1, 1R-2, 1R-3, IR-8	FedRAMP Incident Comm. Procedure 5.2 Detection and Analysis	NIST (SP) 800-53 r4 Appendix F-IR	Incident Handlers Handbook 7 Lessons Learned
Incident Handlers Handbook 2 Preparation	NIST (SP) 800-53 r4 AT-2, 1R-4, IR-6, 1R-7, IR-9, SC-5, SI-4	1R-4, IR-6, IR-7, IR-9	8 Checklist
8 Checklist	Incident Handlers Handbook 3 Identification	Incident Handlers Handbook 4 Containment	
ENISA Cloud Computing Security Risk Assessment Business Continuity Management, page 79	8 Checklist	5 Eradication	
		6 Recovery	
		8 Checklist	

Figure 1: Incident Life Cycle and Normative References

3. Definitions

- Assets: An asset is anything of value to the organization. Assets can be abstract assets (like processes or reputation), virtual assets (data, for instance), physical assets (cables, a piece of equipment), human resources, money, et cetera.²
- Incident: An issue that harms the operation of network and information systems core services.
- Reportable incidents: Incidents deemed to have a significant enough impact that they need to be reported outside the entity according to laws or regulations.
- Incident handling³: The corrective action to address an issue/incidence in violation of security practices and recommended practices.
- Incident response plan: A clear set of instructions that helps an organization prepare, detect, analyze and recover from an incident.
- Incident reporting: The procedure by which the reporting party (cloud provider or cloud operator) shall submit to a national competent authority a report with information on the incident on an ad-hoc basis.
- Impact: A measure of the extent of damage caused by an incident before it can be resolved.
- Root cause: The reason (ultimate root cause) that caused the incident. (A root cause analysis could identify multiple "causes and effects" but will have a single root cause)
- Threat: A threat is any circumstance or event with the potential to cause harm to an information system in the form of destruction, disclosure, adverse modification of data, and/or denial of service.⁴
- Vulnerability: A defect or weakness in a particular system, module, or component that leaves it open to being compromised due to attack, disaster, or other causes.

2 ENISA 2015, Technical Guideline on Threats and Assets, <https://www.enisa.europa.eu/publications/technical-guideline-on-threats-and-assets>

3 NIST.SP800-61r2: Computer Security Incident Handling Guide

4 NIST SP 800-32 under Threat NSTISSI 4009

4. CIR Overview

CIR can be defined as the process designed to manage cyberattacks in a cloud environment and comprises four phases:

- Phase 1: Preparation
- Phase 2: Detection and Analysis
- Phase 3: Containment, Eradication and, Recovery
- Phase 4: Postmortem

There are several key aspects of a CIR system that differentiate it from a non-cloud incident response (IR) system, such as governance, shared responsibility, and visibility.

Governance

Data in the cloud resides in multiple locations, perhaps with different CSPs. Getting the various organizations together to investigate an incident is a significant challenge. It is also resource-draining on large CSPs that have a colossal client pool.

Shared responsibility

Cloud service customers, CSPs, and/or third-party providers all have different roles to ensure cloud security. Generally, customers are responsible for their data and the CSPs for the cloud infrastructure and services they provide. Cloud incident response should always be coordinated across all parties.

The domains of shared responsibilities are also different between the CSPs and CSCs depending on the model of cloud services chosen, such as software-as-a-service (SaaS), platform-as-a-service (PaaS), and infrastructure-as-a-service (IaaS). This idea has to be well understood. For example, in IaaS, managing the operating system (OS) lies with the CSC. Therefore the IR responsibilities for the OS also lie with the CSC.

Responsibility	On-Prem	IaaS	PaaS	SaaS	
Data classification & accountability risks	●	●	●	● - - -	Requires Internal Trust
Client & endpoint risks	●	●	●	● ●	
Identify & access risks	●	●	● ●	● ●	
Application risks	●	●	● ●	● - - -	Requires External Trust
Network risks	●	● ●	●	●	
Host risks	●	● ●	●	●	
Infrastructure risks	●	●	●	●	

 Cloud Provider is responsible
  Cloud Customer is responsible

Figure 2: CSC and CSP Shared Responsibility Risk Matrix⁵

It is essential to discuss—in granular detail—that roles and governance are clear and well-documented in the contract or service-level agreement (SLA) with the CSP. The CSC should not create or settle for any policy that cannot be enforced. Organizations should understand they can never outsource their part of governance or shared responsibilities.

Service Provider Diversity

Organizations should have a consistent and well-defined multi-cloud strategy/framework towards engaging with CSCs, CSPs, and/or third-party cloud providers. Any organization going with an ‘all-in’ strategy with a single CSC, CSP, or third-party cloud provider is indirectly introducing a single point of failure in case of an outage at the service providers’ end.

A single CSP approach to the supply of cloud services may result in a situation where the organization’s business could suffer a sustained outage in case of any failures introduced at the CSC/ CSP over which the organization does not have control. This scenario will impact business operations substantially and raises the possibility of a business continuity plan (BCP) strategy unable to recover—resulting in a systemic CIR event.

When approaching service provider diversity from the CIR perspective, organizations are also encouraged to consider aspects of digital service sovereignty (e.g., data residency, data sovereignty) in their plans.

⁵ Microsoft TechNet 25 October 2019, Shared Responsibilities for Cloud Computing, <https://gallery.technet.microsoft.com/Shared-Responsibilities-81d0ff91>

Visibility

A lack of visibility in the cloud means that incidents that could have been remediated quickly are not addressed immediately and are at risk of further escalation. The cloud can make for a faster, cheaper, and more effective IR if appropriately leveraged. There are already many built-in cloud platform tools, information sources, services, and capabilities provided by CSPs and their partners to significantly enhance detection, reaction, recovery, and forensic abilities. Care must be taken when developing the IR process and documentation when leveraging cloud architectures instead of traditional data center models. The CIR must be proactive and architected to sustain against failure throughout the process.

5. CIR Framework

Incident response and management are considered reactive actions to minimize the damage of incident outbreaks. It is a critical facet of any information security program, as stated in the ninth domain of the *CSA Securitys Guidance for Critical Areas of Focus in Cloud Computing v4.0*.⁶ With appropriate incident response processes and plans defined, CSCs can ensure that detected incidents can be managed and controlled.

Incident response and management frameworks have been developed and documented by many organizations, as stated in chapter 2 of this document. Different frameworks have their objectives and target audiences. This framework has adopted the commonly accepted "Incident Response Lifecycle" described in *CSA Security Guidance for Critical Areas of Focus in Cloud Computing v4.0* and *NIST Computer Security Incident Handling Guide (NIST 800-61rev2 08/2012)*.



5.1 Phase 1: Preparation and Follow-on Review

In the preparation phase, it is necessary to establish the incident response capability so the organization is ready to respond to incidents. In other words, it is vital to know the environment and the "enemy."

When an incident occurs, the objectives of CIR are to achieve the following:

- Provides rapid detection, isolation, and containment
- Minimizes exposure and compromise of personal data, proprietary, and sensitive information
- Minimizes disruption to business and network operations
- Establishes controls for proper retrieval and evidence handling
- Provides incident communication to all affected parties
- Provides accurate reports and useful recommendations
- Protects the organization's reputation and assets
- Educates employees based on lessons learned
- Reviews and improves the CIR Plan based on lessons learned

To understand the organization incident response capability, one of the critical differences between the traditional IR framework and the CIR framework is the presence of the "Shared Responsibility

⁶ Cloud Security Alliance 2017, *Security Guidance for Critical Areas of Focus in Cloud Computing v4*, <https://cloudsecurityalliance.org/artifacts/security-guidance-v4/>

Model” between the CSC and the CSP. In the conventional IR framework, the organization that owns the systems would be solely responsible for the systems. Its computer incident response team (CIRT) should develop its process, procedures, plan, and playbook to handle different types of incidents, including security incidents. Because the organization solely manages the systems, the CIRT commander or lead should coordinate, manage and oversee all the affected systems and collect necessary logs and artifacts from those systems.

However, in a cloud environment, the CSC is not the owner of all the systems. Depending on the adopted service models and their corresponding shared responsibility model, some artifacts and logs are managed by the CSP. When third-party IR providers are engaged, the CIR plan should also include them in the overall process. This juncture presents an appropriate opportunity for organizations to consider vetting any third-party IR vendors to ensure quick access to resources should they be needed in an emergency response situation.

Organizations should familiarize and make full use of their CSPs’ business continuity and disaster recovery capabilities to invoke them in incidents. Thus, it would be necessary for the CSC to understand the IR procedures of the CSP and to align with them through the SLA and contract. To manage and execute this initiative, a CIR plan should include:

1. An analysis of the existing environment, cloud architecture, and responsibility model.
 - a. [CSC] Identify and prepare the inventory, service components, and corresponding service model and deployment model of the cloud services to be used.
 - b. [CSC] Review the compliance requirements (such as data privacy and local regulatory requirements) and extract the compliance requirement, such as the data breach reporting time requirement.
 - c. [CSC] Collect the existing contract and SLA and determine the roles and responsibilities of different parties in the cloud architecture and their obligations according to the shared responsibility model. **A clear description of roles and responsibilities** prevents overlap or overlooking of tasks and unnecessary time wastage in assigning roles during an incident.
 - d. [CSC] Collect the contact methods between different parties (internal team, managed service provider, CSP, or other third parties). **Incident reporting structure** should include contact information such as phone numbers and email addresses.
 - e. [CSC] Collect the incident assistance teams from the CSP. **Incident Assistance teams** include help desks, field-support teams, and other assistance services such as the security operations center or SOC.
 - f. [CSC] Review the CSP's delegated administration privileges, which allows a CSP to access the CSC's tenant with the same level of privilege as a super user. Although a CSP may enforce strong security controls over the CSC's tenant, a threat actor that compromises the CSP may be able to access the CSC's environment. Hence, a CSC should verify if their CSP requires these delegated administration privileges. If the CSP requires delegated administration privileges, a CSC should ensure the CSP has implemented appropriate controls, such as monitoring, to alert the CSC of potential misuse. The CSP should also utilize conditional access policies to restrict access to the CSC's environment where applicable. If the CSP does not require delegated administration privileges, a CSC should ensure the CSP has this privilege removed. Lastly, establish a CIR organization (CSC and CSP)

- g. [CSC] Take the collected contact and identified parties within the organization to formulate the incident response team.
 - h. [CSC] Define the CIR organization structure and appoint the incident response commander, corresponding system owner, technical response lead(s), and technical coordinator(s). Depending on the roles and responsibilities identified in the previous step (analysis of existing environment, cloud architecture, and responsibility model), the CSC may have to assign a technical response lead that can support the implemented systems in the IaaS or arrange a technical coordinator to communicate and collect support indicators of the incident or logs for the following phase. [Phase 2: Detection and Analysis]. Ensure all cloud components are handled with responsible parties.
2. Establish incident handling plan, process, and procedures/playbook (CSC and CSP) for effective and efficient CIR response and remediation.
- a. **Create incident reporting processes and procedures** with contact information such as phone numbers and email addresses.
 - b. **Issue tracking system** to record and track the incident status.
 - c. **Formulate incident response process and procedures, including coordination with third parties and crisis communication**, establish contacts with clear roles and segregated responsibilities, define escalation plans, assign staff, determine procedures, allocate responsibilities formally.
 - d. **Develop a process to update the CIR plan** with any organizational changes.
 - e. **Ensure access to archived lessons learned**, which all team members can access for reference.
 - f. **Subscribe to third-party threat intelligence services** to gain knowledge about current and potential threats.
 - g. **Test CIR plan** through mock incident scenarios as part of staff training; ideally, review and update this plan annually. No matter how well-thought-out the plan is, it will break down if employees are not well-prepared.
 - h. **Develop a continuous training process** for staff on tasks within their scope of responsibilities. This will equip staff with the necessary knowledge to react in demanding times.
 - i. **Define and document the contact list** of the CIR organizations in the CIR plan and procedures. The contact list should be regularly updated.
3. Technical level preparation (CSC and CSP) to proactively watch for indicators of operational deficiencies and malicious activities.
- a. [CSC] The roles and responsibilities should derive from the previous step (analysis of existing environment, cloud architecture, and responsibility model). The CSC should review the entire architecture to determine if any gaps exist within the architecture. The CSC would perform the incident response process through the shared responsibility model in the following phases by the corresponding internal team. The CSC should collect logs and health status from the CSP for those handled by the CSP.
 - b. [CSC] The CSC should compare the logs and health statuses collected from CSP(s) with the list of logs and health statuses that the CSC defined to ensure necessary logs for analysis have been collected. The CSC should also understand the limitations of the logs and artifacts collected from the CSP, especially on the expected logs' availability and retention period.

- c. [CSP] The CSP should **constantly monitor its infrastructure and applications** by proactively scanning system and data center health status and network monitoring.
 - d. [CSC] The CSC should define **preventive measures**, such as having redundancy for critical operations and storage, storage backup, intrusion detection systems and prevention systems, file integrity monitoring systems, antivirus solutions, vulnerability remediation, and firewalls—as well as adopting secure software development life cycle (SDLC) practices.
 - e. [CSC] The CSC should identify the centralized log management and log analysis facilities locations. In many CSC environments, logs are stored in different CSP facilities and cloud servers established by the CSC—and maybe within on-premises servers. Logs should be consolidated for effective incident response and analysis.
 - f. [CSC] Conduct **regular assessments** of vulnerability and risk—including threat detection capabilities—to improve security posture.
 - g. [CSC] **Maintain incident analysis hardware and software** for preserving log files for digital forensics, restoring backups, report writing, etc.
 - h. [CSC] **Establish an automated support** mechanism to request assistance or distribute information. The CSC should also identify and prepare its jumpseat toolkits that would be used for incident response in their cloud environment.
 - i. [CSC] Verify **internal documentation** includes port lists, asset lists, network diagrams, and current network traffic baselines.
 - j. [CSC] A robust **business continuity plan (BCP)** will significantly enhance the organization's operating resilience to manage and recover from incidents. The scope should include services provided by the CSP.
 - k. [CSC] Purchase **cyber-insurance** where available, as it may help mitigate the potential financial impact of a cloud incident.⁷
 - l. [CSC] The CSC should understand the CSP logging schema and how it may differ from on-prem logging. The use of dynamic fields may limit the CSC's ability to query necessary data and create efficient alerts in their security information and event management (SIEM) solution.
 - m. [CSC] The CSC should document logging requirements as CSP products may not possess the capability to support necessary log collection in a centralized SIEM.
4. Communication Channel Preparation (CSC and CSP)
- a. [CSC] Assemble a company team to act as the primary, lone contact for all communications with the CSP.
 - b. [CSC] Develop crisis communication protocol with external parties, such as with the CSP. The **communication methods** should also be ready to reach out to key parties within or beyond the organization to enable smooth communication during an incident
 - c. [CSC] Ensure the team should have an updated contact list of internal and external parties. **The emergency contacts list should include** other IR teams within or beyond the organization, on-call staff information, legal counsel, law enforcers, and other essential incident handler facilities.

⁷ Wikipedia, Cyber insurance, https://en.wikipedia.org/wiki/Cyber_insurance
 AIG, Cyber insurance, <https://www.aig.com/business/insurance/cyber-insurance>
 CHUBB, Cyber Insurance <https://www.chubb.com/sg-en/business/cyber-insurance.html>

The preceding list summarizes the main actions to prepare during the incident response preparation phase. What follows is the list of deliverables resulting from the above actions:

1. Create an IR plan, policy, and procedures.
2. Develop an asset inventory list (including cloud services, servers, account list, implemented security defense mechanisms, expected log files, and facilities).
3. Formulate an incident response role matrix (including the CIRT of the CSC and the other participants' roles from the CSP).
4. Incident response drill test plan and test results.
5. Incident response jumpseat toolbox.

5.1.1 Documentation

Throughout the IR process, the organization shall maintain incident documentation to ensure a systematic record for efficient review of the incident and lessons learned. The organization should manage the following information about an incident record:

1. The current status of the incident ("new," "in progress," "forwarded for investigation," "resolved," etc.).⁸
2. A summary of the incident.
3. Indicators of compromise related to the incident.
4. Other incidents related to the original incident.
5. Actions taken by handlers of this incident.
6. Chain of custody, if applicable.
7. Impact assessments related to the incident.
8. Contact information for other involved parties (e.g., system owners, system administrators).
9. A list of evidence gathered during the incident investigation.
10. Comments from incident handlers.
11. Planned next steps (e.g., rebuild the host, upgrade an application).
12. Restrict access to the incident record to appropriate personnel since it may contain sensitive information with regulatory or compliance implications, IP addresses, exploited vulnerabilities, confidential business information.
13. Retrospective/lessons learned: Document any lessons learned, such as successes, areas of improvement, actions to avoid, and new procedures to improve outcomes.

5.2 Phase 2: Detection and Analysis

5.2.1 Inducement

5.2.1.1 Cause of Cloud Incident

A cloud incident—as defined in this document—is an occurrence that harms the operation of IaaS, PaaS, desktop-as-a-service (DaaS), SaaS, and related services that CSPs provide. Cloud incidents

8 NIST.SP800-61r2, Computer Security Incident Handling Guide

can cause cloud outage (a period when cloud services are unavailable). Cloud incident causes and downtime can fall into one of the following categories:

1. Natural disasters (e.g., flooding, fire)
2. System problems'
 - a. Internal (e.g., software bugs, faulty hardware)
 - b. External (e.g., loss of power supply, telco network connectivity issue)
3. Man-made
 - a. Non-deliberate (e.g., human error)
 - b. Deliberate (e.g., government sanctions, hacker/ DoS attack, ransomware)

5.2.1.2 Signs of an Incident

There is usually a sign before an incident. According to the National Institute of Standards and Technology (NIST) definition, scenarios that constitute a sign include:

- A precursor (a sign that an incident may occur in the future).
- An indicator (a sign that an incident may have occurred or may be occurring now).

	Precursor	Indicator
Natural Disaster	Bad weather forecast	Multiple power interruptions
System Problems	• Lag in response for multiple software services • Web server log entries that show vulnerability scanner usage	• Multiple power interruptions • Noticeable period of fluctuation in power supply • Continuous period of temperature increase in direct current (DC) • Network intrusion detection sensor alerts when buffer overflow attempt occurs against database server
Man-made	• Announcement of new exploit that targets vulnerability of organization's mail server • A threat from a group stating that the group will attack the organization	• Antivirus software alerts when it detects that a host is infected with malware. • A system administrator sees a filename with unusual characters.

Figure 3: Signs of an Incident

5.2.1.3 Common Sources of Precursors and Indicators

The CSP and CSC must have a system or process to detect these signs, which may prevent an actual occurrence. The common sources of precursors and indicators include:

1. Alerts
2. Logs
3. Indicators of compromise (IoC)
4. Industry events
5. Market analysis reports
6. Threat intelligence reports
7. Publicly available information
8. People
9. Social media

It is recommended to have systems in place to collect and analyze these precursors and indicators, ranging from system logs, alerts, SIEM, and a security ops center to an integrated ops center. Ideally, monitor and correlate the various alerts, logs, events, calls, and logs for comprehensive cyber-situation awareness via an integrated ops center. In all cases, the collection and analysis scope must cover the cloud's management plane and not merely the deployed assets.

5.2.2 Incident Analysis to Determine Impacts

5.2.2.1 Incident Analysis

Part of the incident information collection effort is to determine if the issue is a false positive or a false negative.⁹ If the issue is a "false alarm," then the documentation (i.e., ticket) should be updated to document this assessment and close the issue. Each indicator must be evaluated to determine legitimacy.

The following are recommendations for incident analysis:¹⁰

1. Profile networks and systems: Profiling a system—such as baselining—will help better identify when changes occur so they can be better identified.
2. Understand normal behavior: Conducting log reviews should help analysts better notice trends, such as trends over time and abnormal events, that may indicate an incident.
3. Perform event correlation: Evidence of an incident may be captured in several logs containing different data types. A firewall log may have the source IP address used, whereas an application log may include a username.
4. Run packet sniffers to collect additional data: Sometimes, the indicators do not record enough detail to permit the handler to understand what is occurring. If an incident occurs over a network, the fastest way to collect the necessary data may be to have a packet sniffer capture network traffic.

⁹ The SANS Institute, 2011, *Following Incidents into the Cloud*

¹⁰ NIST, *Computer Security Incident Handling Guide*, SP.800-61r2

5. Leverage data analytics to analyze all data sets: One effective strategy to address voluminous indicators is to filter out categories of indicators that tend to be insignificant. Another filtering strategy is to show only the types of indicators that are of the highest significance. However, this approach carries substantial risk because the new malicious activity may not fall into one of the chosen indicator categories. Hence, it would be best if data analytics can be deployed to monitor all collected indicators.

5.2.2.2 Incident Notification

The incident response plan should be organized systematically to minimize impact to business and service operations, and relevant parties should be informed when incidents occur. Incident escalation should be based on the severity of the incident's impact. Because of the high volume of incidents in a highly complex cloud environment, senior management should only be informed of critical and high-impact incidents. The CSP and CSCs should develop and integrate an escalation matrix into the contract and/or SLAs. *Note: The CSP may obligate CSCs to inform the CSP of any significant CSC incidents, as these may pose a threat to the CSP's infrastructure and operations.*

For accurate reporting, the following critical information (5Ws) should be gathered from the incident reporter and, if possible, the affected environment:

1. What happened? Did the user take any actions before and after the incident?
2. Where did the incident happen? Has it been contained, or have any other areas been impacted? What is the confidence level for non-impacted zones?
3. When did it happen?
4. Who discovered it? Who is affected or not affected? How was it discovered?
5. Why did the incident occur? Has the source or "patient zero" been found?

5.2.2.2.1 Incident Notification Timing

Time is of the essence. Though there is a need to resolve the incident quickly, it is equally essential to inform relevant stakeholders promptly to allow them to understand the situation so they can advise or take necessary actions to reduce incident impact. During an incident, crisis communication is an integral part of the crisis management plan covering any incident related to service or business outages, including a cyberattack. Poor incident management can lead to regulatory fines, reputational damage, loss of customer trust, and severe financial loss.

- The initial incident notification should be disseminated to key stakeholders internally and externally within the first two to eight hours to enable horizon scanning at the CSC/CSP/third-party provider.
- A primary informational incident report should be shared with internal stakeholders containing information on at least the first 4Ws within the first four to 48 hours (depending on the incident's impact). When necessary, external stakeholders (CSP/third-party providers) may need to be involved in the investigation and containment. When necessary, external stakeholders may also need to be involved.
- The CSCs/CSPs usually undertake self-reporting within an agreeable timeframe, as per generalized contractual terms and conditions. Organizations may want to undertake a

review if this reporting threshold meets their requirements and aligns with the overall incident management framework.

- Organizations must be aware of regulatory requirements applicable in the country/region/district they operate. For example, the EU's General Data Protection Regulation (GDPR) requires companies to report a breach within 72 hours of becoming aware of the breach (when feasible). This obligation is imposed upon organizations anywhere, so long as they target or collect data related to people in the EU and/or process personal data of EU citizens or residents.

Depending on the escalation workflow, organizations should send notifications expeditiously via the agreed-upon medium (phone call, SMS, email, etc.). Incidents of different severity levels should be escalated to different execution and management parties as agreed to in the CIR plan. If there is a critical impact on business continuity or reputation, organizations should also activate their BCP and/or crisis management plans (CMP).

5.2.2.3 Incident Impacts

An incident impact model must be developed upfront and used by the CSP and CSC to ensure consistency in event assessment, impact, notifications, and actions needed to respond accordingly. The incident priority matrix (also called an "impact and urgency matrix") is derived from impact severity and urgency levels. A quick and proper impact assessment must be performed to determine the damage extent. The following examples include the key impact types that both the CSP and CSC should consider together:

- Business: Scale and level of business criticality
- Financial: Downtime loss or impact to reputation
- Regulatory/legal: Data privacy and contractual terms

Organizations must establish and define proper categorization of the impact severity levels based on their tolerances and appetite for risk. Under the *European Union Agency for Cybersecurity (ENISA) Cloud Security Incident Reporting*, one or more parameters can assess the level of impact. For example: For one day of downtime and a geographic spread of 70 percent, the incident's impact would register "Level 2/Level 1." Upon this determination, users should reference containment guidelines for incidents with "Level 2/Level 1" impact. It is important to note that the given values serve as examples, and the values should be adjusted to reflect organizational nature, priorities, and business objectives.

Urgency levels range from the lowest ("Level 5") to the highest ("Level 1/2") using the following considerations:

- Are the systems or services currently affected critical?
- Are there any workarounds or mitigating actions that can be deployed?
- How many users are affected?
- Can this incident be contained effectively?
- Is this incident spreading slowly or quickly—and affecting other users or systems?
- Other considerations? For example, are there potential legal or regulatory ramifications?

This self-assessment will guide the mobilization of needed resources and determine the extent of actions required to quickly manage and neutralize the incident within the required timeframe. For example, an incident with the highest impact and urgency ("Level 1") would typically result in a "P1" ("Priority 1") designation, which could be a crisis that triggers the organization's crisis management plan (CMP) and escalation to senior management and/or the board.

Organizations should adopt incident classification scales used by several standards and guidelines to help users gauge the severity of impact and/or the relevant importance of cloud services availability to business operations. The following is a set of policies based on the current operational trend of CSPs:

Priority Code = Incident Scale	Incident Impact	Target Response Time	Target Resolution Time
1	Critical	< 5 min With a 24-hour response team	< 1 hour
2	High	< 15 mins during office hours < 2 hours after office hours for an office-hour response team. Otherwise, 4-8 hours depending on site.	< 4 hours
3	Medium	< 15 mins during office hours < 2 hours after office hours for an office-hour response team. Otherwise, 4-8 hours depending on site.	< 8 hours
4	Low	< 15 mins during office hours < 2 hours after office hours for an office-hour response team. Otherwise, 4-8 hours depending on site.	< 24 hours
5	Very Low	No response needed with system auto-filter.	--

Figure 4: Incident response policies.

Organizations may also wish to undertake a business impact analysis (BIA)—or a threat, vulnerability, and risk assessment (TVRA) specific to organizational parameters—and consider purchasing cyber-insurance to mitigate the potential financial impact of a cloud incident.

5.2.3 Evidence Gathering and Handling

Identifying data relevant to the investigation is vital in determining the root cause of the incident and identifying lessons learned to avoid repeated incidents. Identified data can also help support beneficial information-sharing initiatives for leverage to prevent similar incidents.

Note that CSPs may limit log retention periods due to GDPR or other compliance requirements. These limitations must be understood and accounted for in incident response planning as log availability will affect necessary evidence gathering (depending on the cloud service chosen).

Possible locations of relevant data include storage drives attached to virtual instances and the memory space of an instance. By utilizing CSP capabilities, such as for-instance snapshots, CIR teams can obtain snapshots of the virtualized storage drives attached to incidents and utilize them for further analysis and discovery. These snapshots can be mounted to digital forensic investigative resources for scrutiny with widely used forensic analysis tool sets.

Any collected evidence should also undergo hash activity processes. This helps ensure the integrity of the collected information and that the data has not been altered from its original source. This undertaking also helps ensure evidence admissibility regarding potential legal proceedings. Ensure forensic work is performed on a copy of the collected evidence (rather than the original data that has been hashed) for court admissibility.

For cybersecurity incidents, the following steps should occur to identify attacking hosts:

- Validate the attacking host's internet protocol (IP) address/domains/emails/other info
- Research the attacking host through search engines
- Use incident databases
- Monitor possible attacker communication channels
- Create IoC alerts to SIEM or other tools to assist with finding the attacking host

Any collected evidence should utilize a hash activity to ensure the integrity of the collected data. This process can be used to verify that evidence has not been altered from its original source, and helps ensure admissibility for potential legal proceedings.

5.3 Phase 3: Containment, Eradication, and Recovery

Containment: The methods for containing damage when responding to a security incident are unique to the incident and the organization. After identifying an incident correctly, the policy should list the actions to be taken based on the incident type. Containment deals with isolating the infected system.

Note: Depending on the incident and its effects, containment, eradication, and recovery may all be part of the same process.

Containment is essential upon detection of a security incident to prevent further attacker activity and system re-entry. Unchecked activity may overwhelm resources or increase damage. From an attacker's perspective, a typical attack starts with the initial compromise, then establishing a foothold by downloading malware, an escalation of privileges, and then network exploration. Up to this point, the attacker is likely limited to a single machine and cannot exfiltrate data.

Next, the attacker may move laterally and establish persistence by installing different malware on a small number of other machines. This keeps their detection risk low while providing means of re-entering the network should the initial compromise be detected. The attacker is now established in the system and will start to execute their mission.

Upon incident discovery, affected organizations should execute predefined CIR plans (as stipulated in "Phase 1: Preparation"), such as taking systems offline, quarantining systems, and restricting connectivity. It is paramount not to remove the threat by blind deletion, as this destroys forensics evidence needed for CIR plan revisions. Containment provides time to develop a remediation strategy. An essential part of containment is decision-making (e.g., shut down a system, disconnect it from a network, delete API keys, disable username). Such decisions are much easier to make with predetermined strategies and procedures for incident containment. To define and document the strategies and procedures, IR teams should utilize playbooks and runbooks to simplify tasks.

Organizations should define acceptable risks in dealing with incidents and develop strategies accordingly. Containment strategies vary based on incident types. For example, the process to contain an email-borne malware infection is quite different compared to a network-based DDoS attack response. Organizations should create separate containment strategies for each major incident type, with criteria documented clearly to facilitate decision-making.

Criteria for determining the appropriate strategy include:

- Business impact
- Potential resource theft and damage
- Need for evidence preservation
- Service availability (e.g., network connectivity, services provided to external parties)
- Time and resources needed to implement the strategy
- Strategy effectiveness (e.g., partial containment, full containment)
- Containment approach duration, complexity (e.g., an emergency workaround to be removed in four hours vs. temporary workaround to be removed in two weeks vs. permanent solution)
- Resource availability (particularly technical expertise)
- Availability and integrity of backup/copies/snapshots
- Availability of sandbox/honeypots environments

The appropriate containment strategy's ultimate goal is to limit the attacker's movement and prevent further unauthorized access or infection within the shortest possible time while minimizing service disruptions. An appropriate strategy will prevent further damage from happening while preserving forensic evidence necessary for investigation.

5.3.1 Choosing a Containment Strategy

In certain cases, some organizations redirect attackers to a sandbox (a form of containment similar to a honey pot) so they can monitor the attacker's activity (usually to gather additional evidence). The IR team should discuss this strategy with its legal counsel to determine feasibility.

Organizations should not implement alternative methods to monitor attacker activities (other than sandboxing). If an organization detects a system compromise and allows the compromise to continue, the organization may be held liable if the attacker uses the compromised system to attack other systems.

The delayed containment strategy is dangerous because an attacker could escalate unauthorized access or compromise other systems. Another potential issue is that some attacks may cause additional damage after containment. For example, a compromised host may run a malicious process that pings another host periodically. When the incident handler attempts to contain the incident by disconnecting the compromised host from the network, the subsequent pings will fail.

As a result of the failure, the malicious process may overwrite or encrypt all the data on the host's hard drive. Even after a host has been disconnected from the network, handlers must not assume that further damage to the host will be prevented.

5.3.2 Eradication and Recovery

Eradication: The removal of the problem. This includes minimizing loss, information theft, and service disruption—and threat elimination. The eradication step may be necessary to restore operational levels for all affected system(s). The threat, infection, or damage must be removed to return the system(s) to operational levels. This could entail wiping disks clean, deleting compromised code and user accounts.

Recovery: Includes restoring computing services securely and promptly.¹¹ The recovery process repairs a system to its original—or enhanced—condition. This procedure returns it to production by applying patches, rebuilding the system's key files, reinstalling applications, changing passwords, and restoring files from backups.

11 FedRAMP PMO 2017, FedRAMP Incident Communication Procedure, https://www.fedramp.gov/assets/resources/documents/CSP_Incident_Communications_Procedures.pdf

5.4 Phase 4: Post-Mortem

The final phase of the CIR process is a post-mortem. The objective of this pivotal phase is to evaluate how the incident was processed and managed by enterprise and CSP teams with the aim to improve future incident handling procedures. The evaluation is underpinned by reviewing incident data and after-action reports that contain "Lessons Learned."¹² The crucial question to answer: what could have been done better? This feedback should translate into new countermeasures flowing back into Phase 1.

5.4.1 Incident Evaluation

An analysis of incident characteristics may indicate—at a minimum— security weaknesses and threats, cloud configuration weaknesses, and changes in incident trends. This data can be added back as a feedback loop into the risk assessment process, which may lead to the selection and implementation of additional controls, processes, and preventive measures.

An objective postmortem will also help the team use collected information to gauge the overall effectiveness of the CIR process.

Questions may include:

- How did they respond?
- What were their strengths and weaknesses?
- What were their lessons learned?

If incident data is collected and stored properly, it should highlight several measures of success (or at least activities) of the IR team.

5.4.1.1 Incident Evaluation Metrics

Incident data can also be collected to determine if notable trends exist over time. These patterns may reveal more about how the team is doing over a defined duration and if there are improvements (e.g., a decreasing number of incidents) or areas that warrant increased attention (e.g., a spike in security-related incidents). Organizations must typically report such information in regulated industries—especially major incidents—to regulatory bodies and management. The CSCs are expected to collect necessary data in a timely, accurate, and complete manner to meet these requirements.

Data such as flow logs or other traffic logs should be collected to review unauthorized access or suspicious traffic.

¹² FedRAMP PMO 2017, FedRAMP Incident Communication Procedure, https://www.fedramp.gov/assets/resources/documents/CSP_Incident_Communications_Procedures.pdf

Collected incident data should consist of metrics (performance indicators) that capture the following information:

- **Mean time to detect (MTTD):** The average time to discover the security incident. How long did it take from when the incident started until the team became aware of it? This is directly linked to attacker dwell time (the time between attacker infiltration and the detection point).
- **Mean time to acknowledge (MTTA):** The time it takes a security operator to respond to a system alert. While MTTD measures the time before an attacker is noticed, MTTA focuses on measuring a security operator's time responding to the security alert and starting the analysis.
- **Mean time to recovery (MTTR):** The time required to bring a system back into an operating state (linked to phase 3).
- **Mean time to containment (MTTC):** The average time required to detect, respond to, eradicate and recover from an incident. The MTTC can be calculated by adding up the MTTD, MTTA, and MTTR for all in-scope incidents, divided by the number of in-scope incidents. This metric is considered a key metric (key performance indicator, or KPI) as it shows how well the incident response team is organized. An elevated MTTC signals that some subprocesses are not optimal during incident response. A lower MTTC indicates the team is very well-organized.
- Threat metrics, e.g., Gbps or Tbps if a DDoS attack
- Threat actor TTPs (tactics, techniques, and procedures). These include phishing and account manipulation. More examples can be found in MITRE's ATT&CK® Cloud Matrix¹³

5.4.1.2 Incident Classification

Severity and urgency classifications (H/M/L) may change after a postmortem.

High severity incidents that compromise the confidentiality/integrity of personally identifiable information (PII) or personal health information (PHI)—and the availability of services for a significant number of customers—can have a significant financial impact. Examples include:

- Confirmed breach of PII/PHI
- Successful root-level compromise of production systems
- Financial malware
- Denial of service attacks resulting in severe outages

Medium-severity incidents represent attempts (possibly unsuccessful or not-yet-successful) at breaching PII or those with limited availability/financial impact. Examples include:

- A suspected PII breach
- Targeted attempts to compromise production systems
- DoS attacks resulting in limited system degradation or other performance issues

¹³ MITRE ATT&CK® Matrix for Enterprise covering cloud-based techniques
<https://attack.mitre.org/matrices/enterprise/cloud/>

Low severity incidents do not impact PII, availability, or financial impact to the enterprise or customers. Examples include:

- Attempted compromise of non-important systems (e.g., staging/testing instances)
- Incidents involving specific employees
- DoS attacks with no noticeable customer impact

5.4.2 Incident Closing Report

Once the incident has closed, the CIR team that managed the event shall compose a formal after-action report (AAR) using data collected from previous phases and incident evaluation. This task is essential in the postmortem phase and should be performed as soon as possible while lessons are still fresh. If delayed, critical details may be lost or forgotten—potentially making a significant difference in future incident prevention. The CIR team should present the AAR to key stakeholders within two weeks of the incident closure.¹⁴ Appropriate countermeasures must be formulated and validated by (senior) management. The AARs are best created using a formally approved reporting template to ensure that reports consistently meet expected standards.

An incident report should contain the following:

- Date and time of the incident
- Date and time of incident closure
- Scope of the incident
- Name of the person who reported the incident
- Organization and business unit of the affected person
- Incident description
- Affected cloud system(s) and providers/on-premises resources (hardware, software, locations) and respective SLAs
- Business service owner and CSP points of contact (if applicable), including CSP actors involved during incident management (if applicable)
- Incident classification (severity classification)
- Company/customer impact analysis
- Resolution
- Recommendations
- A "lessons learned" section to determine successes and needed improvements to develop an enhanced response to prevent future incidents

While writing the report, consider the following elements:

- Review the incident's timeline and any CIRT and CSP CIRT observations.
- Perform a thorough root-cause analysis supported by the "5 Whys" (or "5Y") technique to identify and review all contributing event factors.

¹⁴ SANS Institute 2021, Incident Handler's Handbook, <https://www.sans.org/reading-room/white-papers/incident/incident-handlers-handbook-33901>

- Prioritize remediation steps to reduce the future likelihood of another event.
- Use the AAR as training material for new team members to communicate how more experienced team members respond to incidents.
- Centralize and index the AAR (per classification levels), and generate a follow-up report for each incident. Reports are valuable references when dealing with similar, future incidents.
- Review communication channels (CSP<->CSC) and update where necessary.
- Review forensic capabilities and determine if any elements are missing from the “cloud jump kit.”¹⁵
- Review data identified in the incident related to exploitable security vulnerabilities, sensitive data details, or other details impacting PII/PHI.
- Review breach notification timelines (e.g., GDPR) and processes.
- For CSCs, review delivered provider support during incident responses and evaluate if contractual adjustments are required to facilitate enhanced provider support.

It is often considered necessary to publish the report to the broader public after reporting the information to top management to facilitate incident information sharing across enterprises. This transparency helps peers better identify and control risks.

5.4.2.1 Lessons Learned

The final step in handling a security incident is determining what was learned. If gaps are identified during the incident response related to personnel, processes, or technology, they must be addressed. The person who closes the event must ensure that a retrospective review is held regarding the security incident—an undertaking referred to as “lessons learned.” Use “lessons learned” to help revise and solidify the CIR plan. Each IR team should proactively evolve to reflect new threats, improved technology, and lessons learned¹⁶—improving future response actions.

Security Guidance: Pay particular attention to data collection limitations and determine how to address the issues moving forward. As cloud data resides in multiple locations (and perhaps with different CSPs), the following considerations present challenges to this phase of the process:

- Challenges related to obtaining and coordinating incident data collection from various third-party providers (internet service providers)
- Resource dependency from third-party providers (potentially due to the size of dependencies from their client pool).

The following suggested questions can help CSCs come up with their own inquiries:

- What part(s) of the service layer had an issue? What was the impact on the affected applications and users?
- How long did the problem last, and at what time?
- Is the problem cause known?
- What was learned that could prevent or mitigate the occurrence of this event?

¹⁵ CSA Security Guidance For Critical Areas of Focus In Cloud Computing v4.0, section 9.1.2

¹⁶ NIST.SP.800-61r2 Computer Security Incident Handling Guide

- What actions should be taken?
- Was there anything suspicious from a security perspective?
- How well (or how quickly) did the provider or broker provide incident support (if applicable)?
- How well-identified were in-scope technologies for forensic evidence gathering?
- Did someone—or did automated monitoring or other scanning systems—detect unauthorized access or suspicious traffic across a remote connection? Were roles and responsibilities clear from the time of the incident throughout the incident's life cycle?
- Did the technologies raise an alert?
- Can previous "lower"-classified incidents be linked to the root causes?

5.4.3 Incident Evidence Retention

All identified evidence collected during "Phase 2 Detection and Analysis" must be retained according to the requirements set forward for the enterprise's applicable legal, regulatory, industry, or contractual obligations. Evidence is kept for the following three (3) purposes:

- Regulatory compliance requirements (i.e., specific levels and granularities of audit logging, alerts generation, activity reporting, and data retention). Data retention may not be a part of standard service agreements impacted by providers.
- Legal: To support a prosecution for compromise of PII/PHI or enterprise systems.
- Risk management: To reflect and reassess new threats tactics and techniques.
- Training: To facilitate better team preparedness for future incidents, incorporate adaptive incident learning.¹⁷

The enterprise forensic model must be capable of facilitating the required evidence retention periods and technologies used. As per prior CSA guidance.¹⁸ The CSC should work with the CSP to evaluate incident handling. The retention of digital forensic evidence in a cloud context must be seen as an integrated model between the CSP and CSC.¹⁹

17 Incident Response Teams – Challenges in Supporting the organizational Security Function, Ahmad, Hadgkiss & Ruighaver 2012; Shedden, Ahmad & Ruighaver 2011 <https://www.sciencedirect.com/science/article/pii/S0167404812000624?via%3Dihub>

18 CSA Security Guidance For Critical Areas of Focus In Cloud Computing v4.0

19 An integrated conceptual digital forensic framework for cloud computing, Martini and Choo <https://www.sciencedirect.com/science/article/abs/pii/S174228761200059X>

6. Coordination and Information Sharing

Addressing the complexities of the shared responsibility model in terms of incident response requires considerable and varied proactive investments by cloud users and CSPs. Effective use of these investments is critical to ensure an efficient and effective CIR. All cloud stakeholders should jointly develop short-term and long-term goals for CIR. Some examples of long-term goals include building/continuously enhancing frameworks to engage affected users to mitigate losses and strategizing business recovery methods.

The communication path between the provider and the users should be established appropriately. Regular updates should be available for any impacted users to mitigate losses and strategize business recovery methods. Effective coordination and communication go beyond just reporting to the customers.

Because of the shared nature of cloud computing, an attack typically affects more than one organization simultaneously. Thus, incident information sharing is mutually beneficial in helping involved organizations guard against the same threats. The CSA runs the Cloud Cyber Incident Sharing Center (CloudCISC)²⁰ that facilitates incident data sharing between participating CSPs.

Coordination with key partners, IR teams in other departments, and law enforcement agencies significantly reinforce CIR capabilities. This communication should be set up from the start—during the planning phase—and maintained throughout the entire CIR process, as necessary.

²⁰ More information on CloudCISC: <https://cloudsecurityalliance.org/research/working-groups/cloudcisc/>

The following infographic exemplifies the various stages an organization transitions through to ensure effective communication in case of a crisis:²¹

Preparedness	Identifying Communication Team	Selecting Communications Channels	Message for Target Audience
CCMP: Incident Management Plan	Chief Marketing Officer	Internal and External emails	Regulator
	Communication Lead		Board of Directors
Maintaining a RACI Matrix or a linear responsibility chart	Advisors	Press Release to Media	Workforce
	Subject Matter Experts	Boardroom Presentation	Third Party
Setting up war room	Company Secretary	Regulatory Reporting	Customers
Cyber Crisis Table Top Exercises		Shareholder's Meeting	Insurer
		IVR Service	Law Enforcement Agency
RACI: Responsible, Accountable, Consulted, and Informed		Notice/Briefing to and via regional office/branch network	Channel Partners
		Website	Creditors
		Social Media	Shareholders
		Customer Support	

Figure 5: Effective crisis communication stages.

6.1 Coordination

6.1.1 Coordination Relationships

All stakeholders should work together to identify their roles and responsibilities during cloud security incidents explicitly. Traditionally, these roles closely tie with their duties in the shared responsibility model. For example:

- a security incident occurring in the platform or service layer for a PaaS or SaaS application should be driven by the CSP;
- a security incident occurring in the application layer for a PaaS application should be driven by the CSC;
- a security incident occurring in the platform layer for an IaaS infrastructure cloud should be driven jointly by the CSC and the CSP to determine if it originated in the CSC's environment or the CSP's environment.

21 REBIT Cyber Crisis Communications Playbook <https://rebit.org.in/playbooks-and-presentation/cyber-crisis-communications>

Usually, all incidents require close cooperation between CSCs and CSPs for effective incident management.

Stakeholders should proactively identify such incident scenarios along with their roles and responsibilities. They should also identify communication channels (e.g., email, video/call conference call details) for use during incidents so that stakeholders know how to share information efficiently.

Stakeholder communications: Communication recommendations should be based on the different first responder possibilities (e.g., CSPs as the first responders vs. cloud users as the first responders).

6.1.2 Sharing Agreements and Reporting Requirements

Once stakeholders identify their roles and responsibilities, it's essential to get these relationships formalized in contract agreements. These agreements should include nondisclosure agreements (NDAs) for all stakeholders so they can share information confidentially (including an enterprise's most sensitive information). Organizations trying to share information with external organizations should consult with their legal departments before initiating coordination efforts. There may be contracts or other agreements that must be put into place before discussions occur.

Organizations should also consider any existing reporting requirements, such as sharing incident information with an information sharing and analysis center (ISAC) or reporting incidents to a higher-level CIRT.

6.2 Information-Sharing Techniques

Cloud stakeholders must have the capability to identify threats and share security information with key stakeholders. Often, stakeholders do not have clear direction regarding the most optimal ways to discover or share critical information about an incident—and to assess their capabilities objectively. Sharing techniques must be evaluated to determine effectiveness to reduce stakeholder burden while ensuring interconnectivity and resiliency. Even the smallest organizations must retain a capacity to share incident information with peers and partners to realize positive outcomes. Organizations should share information throughout the incident response life cycle and not wait until an incident has been fully resolved.

Information sharing is a fundamental element in enabling coordination across organizations.

1. Ad hoc
2. Partially automated
3. Security considerations

6.3 Granular Information Sharing

Organizations must also weigh the benefits of information sharing with the drawbacks of sharing sensitive information. Enterprises should only share necessary data with appropriate parties. Ideally, all stakeholders should have an NDA in place, providing contractual protections for sensitive and proprietary information.

6.3.1 Business Impact Information

Cloud security incidents are both business problems and IT problems. Cloud security incidents may cause a range of negative business impacts, such as financial loss (e.g., service unavailability, loss of compliance certifications resulting in an inability to do business, incident response costs), reputational impact (loss of customer trust), trade secret disclosures, intellectual property theft, sensitive data breaches, or other issues.

Business impact information is only useful for reporting to organizations that are interested in ensuring the mission of the affected enterprise. In many cases, IR teams should avoid sharing business impact information with outside organizations unless there is a clear value proposition or formal reporting requirements. However, in some cases, organizations may be forced to share this information publicly due to regulatory and legal requirements.

Business impact information describes how the incident affects the organization in terms of mission impact, financial impact, etc. At least at a summary level, such information is often reported to higher-level coordinating IR teams to communicate an incident's damage estimate.

Business impact information is only useful for reporting to organizations that have some interest in ensuring the mission of the organization experiencing the incident. In many cases, IR teams should avoid sharing business impact information with outside organizations unless there is a clear value proposition or formal reporting requirements.

6.3.2 Technical Information

Because CSPs cater to many clients, adversaries often use the same weakness to compromise multiple CSP customers. Once a CSC/CSP extracts technical details about an attack or emerging threat, this data can be distributed to enhance defenses against the specific attack.

In today's digital economy, speed and efficiency are essential. The speed at which cybercriminals operate can be worrying for those tasked with defending networks from attacks. The industry must share more security intelligence with industry peers to better protect and adapt to evolving threats. While enterprises gain value from collecting their internal indicators, they may gain additional value from analyzing indicators received from partner organizations and sharing their internal indicators for external analysis and use. If organizations receive external indicator data about an incident they have not seen, they can use that indicator data to identify the incident as it begins. Similarly, organizations may use external indicator data to detect an ongoing incident it was not aware of due to a lack of internal resources to capture the specific indicator data.

Technical indicator data is useful when it allows an organization to identify an actual incident. However, not all indicator data received from external sources will pertain to the organization receiving it. External data may occasionally generate false positives within the receiving organization's network and cause unnecessary resource allocation on nonexistent problems.

While organizations gain value from collecting their internal indicators, they may gain additional value from analyzing indicators from partner organizations and sharing internal indicators for external

analysis and use. If an organization receives external indicator data on an incident they have not seen, they can use that indicator data to identify the incident as it begins to develop. Similarly, an organization may use external indicator data to detect an ongoing incident it was not aware of due to a lack of internal resources. Organizations may also benefit from sharing internal indicator data with external organizations.

Organizations should share as much insight as possible. However, there may be security and liability reasons dictating why organizations may withhold details of an exploited vulnerability.

Technical indicator data is useful when it allows an organization to identify an actual incident. However, not all external source indicator data will pertain to the organization receiving it. In some cases, this external data will generate false positives within the receiving organization's network—causing unnecessary resource allocation for nonexistent problems.

6.3.3 CSP Dashboard

A CSP should offer a self-service customizable dashboard for users that notifies them about incidents so customers are up-to-date. These dashboards are typically used to communicate incidents that impact a great number of customers. The CSPs should also support configuration options to customize cloud alerts and create personalized dashboards to analyze relevant incidents, monitor cloud resource impacts, provide guidance and support, and share details and updates. These dashboards can be designed as the single source of truth concerning cloud resources and should give users more visibility into any issues that may affect them.

6.4 Table-Top Exercises and Incident Simulations

It's difficult for most enterprises to prepare for a security incident with tangible, "real world" experience—outside of a handful of progressive organizations. These realistic exercises introduce harmless (but real) security vulnerabilities and simulate external exploitation to evaluate readiness for these organizations. In such activities, a small organizational team is aware of the exercise. For everyone else, there is no practice. It's an actual security incident.

Therein lies the value in table-top exercises, a pure simulation of an attack scenario, and a security incident preparedness activity. Table-top exercises help organizations consider various security incident scenarios and prepare for potential cyber threats by guiding participants through the process of responding to a simulated incident scenario. The experience provides hands-on training for participants that can then highlight flaws in the IR process.

Any organization should be able to perform table-top exercises (as opposed to introducing bugs in customer environments that require sophisticated technical and operational capabilities). Furthermore, table-top exercises are far less resource-intensive compared to the "real-world" simulations.

Table-top exercises help improve the overall incident response posture and the collective team preparedness and decision-making process when incidents occur. Exercises begin with the IR plan and gauges team performance against it. Since most organizations are unprepared for cloud security incidents, having a well-executed IR plan is critical.

7. Summary

Benjamin Franklin is quoted as saying, “By failing to prepare, you are prepared to fail.”

In many ways, this sentiment hits home for organizations concerning cyberattack threats. Organizations should develop a solid understanding of the incident response process—and its incident response capabilities—to prepare for any potential incidents.

This paper explored the CIR framework and the preparation required to respond to incidents effectively. It serves as a go-to guide for a CSC to prepare for and manage cloud incidents through the entire lifecycle of a disruptive event. It also provides a transparent, common framework for CSPs and CSCs to share cloud incident response practices.

We presented the CIR framework in four phases (plus a final section covering coordination and information sharing).

Preparation addresses the strategies and actions required in advance of a cloud incident. An effective incident response plan includes forming a CIR team (CIRT), strategy planning and preparation, procedures development, technical preparation, and communication plan creation.

Detection and analysis covers the various signs and possible causes of cloud incidents for early detection. To determine the root cause, multiple means are discussed. The speed of early incident notification (and the corresponding resolution timing based on business impacts) is also highlighted for CSP/CSC consideration.

Containment, eradication, and recovery explain the importance of choosing the right strategy to stop the attacker from doing further systems damage while investigations and forensics are undertaken.

The **Postmortem** process identifies gaps in personnel, processes, or technology and translates these into “lessons learned” that must be ingested in the preparation phase. The key objective of this closing phase is to improve future incident handling. To improve an enterprise’s security capabilities, it is critical to review the incident/forensic support of the CSP(s) (if applicable), the available technological tools to support event analysis, the TTPs used by the actor, and to conduct forensic investigations.

The **Coordination and information sharing** section describes how the complexities of threats to the cloud requires stakeholders to coordinate and share security information to mitigate losses.

In conclusion, this framework will help guide CSCs in determining their security requirements and appropriate incident protection levels. Additionally, CSCs can use this guide to negotiate with CSPs and/or third parties to ascertain capabilities and shared responsibilities.