

NIST Special Publication
NIST SP 800-201 ipd

NIST Cloud Computing Forensic Reference Architecture

Initial Public Draft

Martin Herman
Michaela Iorga
Ahsen Michael Salim
Robert H. Jackson
Mark R. Hurst
Ross A. Leo
Anand Kumar Mishra
Nancy M. Landreville
Yien Wang

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-201.ipd>

NIST Special Publication
NIST SP 800-201 ipd

NIST Cloud Computing Forensic
Reference Architecture

Initial Public Draft

Martin Herman*	36	Michaela Iorga	
Information Access Division	37	Computer Security Division	
Information Technology Laboratory	38	Information Technology Laboratory	
Ahsen Michael Salim	39	Robert H. Jackson	
American Data Technology, Inc.	40	Mark R. Hurst	
	41	SphereCom Enterprises, Inc.	
Ross A. Leo		Anand Kumar Mishra	
University of Houston-Clear Lake	42	National Institute of Technology Sikkim	
The CyberSecurity Institute	43		
Nancy M. Landreville	44	Yien Wang	
Graduate School of Cybersecurity	45	Auburn University	
and Information Technology	46		
University of Maryland (GC)	47		
	48		

*Former NIST employee; all work for this publication was done while at NIST.

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-201.ipd>

February 2023



U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at <https://csrc.nist.gov/publications>.

Authority

This publication has been developed by NIST in accordance with its statutory responsibilities under the Federal Information Security Modernization Act (FISMA) of 2014, 44 U.S.C. § 3551 et seq., Public Law (P.L.) 113-283. NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official. This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

NIST Technical Series Policies

[Copyright, Use, and Licensing Statements](#)
[NIST Technical Series Publication Identifier Syntax](#)

Publication History

Approved by the NIST Editorial Review Board on YYYY-MM-DD [will be added upon final publication]

How to Cite this NIST Technical Series Publication

Herman M, Iorga M, Salim AM, Jackson RH, Hurst MR, Leo RA, Mishra AK, Landreville NM, Wang Y (2022) NIST Cloud Computing Forensic Reference Architecture. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-201 ipd. <https://doi.org/10.6028/NIST.SP.800-201.ipd>

Author ORCID iDs

Martin Herman: 0000-0001-9315-6458

Michaela Iorga: 0000-0001-7880-6045

Public Comment Period

February 8, 2022 - March 31, 2023

Submit Comments

sp800-201@nist.gov

National Institute of Standards and Technology

Attn: Computer Security Division, Information Technology Laboratory

100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

All comments are subject to release under the Freedom of Information Act (FOIA).

Abstract

This document summarizes research performed by the members of the NIST Cloud Computing Forensic Science Working Group and presents the NIST Cloud Computing Forensic Reference Architecture (CC FRA, also referred to as FRA for the sake of brevity), whose goal is to provide support for a cloud system's forensic readiness. The CC FRA is meant to help users understand which cloud forensic challenges might exist for an organization's cloud system. It identifies challenges that require at least partial mitigation strategies and how a forensic investigator would apply that to a particular forensic investigation. The CC FRA presented here is both a methodology and an initial implementation. Users are encouraged to customize this initial implementation for their specific situations and needs.

Keywords

civil litigation; criminal investigation; cybersecurity; digital forensics; enterprise architecture; enterprise operations; forensic readiness; incident response.

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, and its collaborative activities with industry, government, and academic organizations.

Call for Patent Claims

This public review includes a call for information on essential patent claims (claims whose use would be required for compliance with the guidance or requirements in this Information Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be directly stated in this ITL Publication or by reference to another publication. This call also includes disclosure, where known, of the existence of pending U.S. or foreign patent applications relating to this ITL draft publication and of any relevant unexpired U.S. or foreign patents.

ITL may require from the patent holder, or a party authorized to make assurances on its behalf, in written or electronic form, either:

- a) assurance in the form of a general disclaimer to the effect that such party does not hold and does not currently intend holding any essential patent claim(s); or
- b) assurance that a license to such essential patent claim(s) will be made available to applicants desiring to utilize the license for the purpose of complying with the guidance or requirements in this ITL draft publication either:
 - i. under reasonable terms and conditions that are demonstrably free of any unfair discrimination; or
 - ii. without compensation and under reasonable terms and conditions that are demonstrably free of any unfair discrimination.

Such assurance shall indicate that the patent holder (or third party authorized to make assurances on its behalf) will include in any documents transferring ownership of patents subject to the assurance, provisions sufficient to ensure that the commitments in the assurance are binding on the transferee, and that the transferee will similarly include appropriate provisions in the event of future transfers with the goal of binding each successor-in-interest.

The assurance shall also indicate that it is intended to be binding on successors-in-interest regardless of whether such provisions are included in the relevant transfer documents.

Such statements should be addressed to: sp800-201@nist.gov

Table of Contents

Executive Summary	1
1. Introduction	2
1.1. The Need for a Cloud-specific Forensic Reference Architecture	2
1.2. The Approach	5
2. Overview of NIST Cloud Forensic Challenges	7
3. Overview of CSA's Enterprise Architecture	9
4. The Forensic Reference Architecture Methodology	11
5. The Forensic Reference Architecture Data	17
6. Conclusion	20
References	21
Appendix A. Acronyms	23
Appendix B. Glossary	25
Appendix C. CSA's Enterprise Architecture	27
Appendix D. NIST's Forensic Reference Architecture Data Set	28

List of Figures

Fig. 1. Forensic Reference Architecture Overlaying Approach	6
Fig. 2. Mapping Flowchart	12
Fig. 3. Excerpt of the Forensic Reference Architecture (Challenges vs. Capabilities Mapping Table).	18
Fig. 4. CSA's Enterprise Architecture (v1.1)	27
Fig. 5. NIST's FRA Data Set	28

Acknowledgments

This report is dedicated to the memory of our colleague, collaborator, and friend, Ernesto F. Rojas of Forensic & Security Services Inc., who passed away unexpectedly.

This publication was developed by the [*NIST Cloud Computing Forensic Science Working Group \(NCC FSWG\)*](#) co-chaired by Dr. Martin Herman and Dr. Michaela Iorga. NIST and the co-chairs wish to gratefully acknowledge and thank the members whose dedicated efforts contributed significantly to researching this topic and to generating the data included in this publication, which constitutes the foundation of NIST's Cloud Computing Forensic Reference Architecture. The authors thank Richard Lee of Citizens Financial Group, who served as our reviewer; Rodrigo Sardinias of Auburn University; Pw Carey of Grapevine Solutions; Richard Austin, formerly of Hewlett-Packard Enterprise; Dr. Kevin Stavinoha, formerly with Cisco Systems; and Otto Scot Reemelin – who was with Ipro Tech during this work – for contributing to the data collection, analysis, and review. The authors would also like to thank their peers who continuously provided feedback during the research and development stages of this document.

194 **Executive Summary**

195 The rapid adoption of cloud computing technology has led to the need to apply digital forensics
196 to this domain. New methodologies are required for the identification, acquisition, preservation,
197 examination, and interpretation of digital evidence in multi-tenant cloud environments that offer
198 rapid provisioning, global elasticity, and broad network accessibility. This is necessary to
199 provide capabilities for incident response, secure internal enterprise operations, and support for
200 the U.S. criminal justice and civil litigation systems.

201 This document presents the NIST Cloud Computing Forensic Reference Architecture (CC FRA,
202 also referred to as FRA for the sake of brevity), whose goal is to provide support for a cloud
203 system's forensic readiness. The CC FRA is meant to help users understand the cloud forensic
204 challenges that might exist for an organization's cloud system. It identifies forensic challenges
205 that require mitigation strategies and how a forensic investigator would apply that to a particular
206 forensic investigation.

207 The CC FRA provides a useful starting point for all cloud forensic stakeholders to analyze the
208 impacts of cloud forensic challenges previously reported by NIST. It does so by considering each
209 cloud forensic challenge in the context of each functional capability presented in the Cloud
210 Security Alliance's Enterprise Architecture.

211 While the CC FRA can be used by any cloud computing practitioner, it is specifically designed
212 to allow cloud system architects, cloud engineers, forensic practitioners, and cloud consumers to
213 ask specific questions related to their cloud computing architectures. The CC FRA is both a
214 methodology and an initial implementation, and users are encouraged to customize this initial
215 implementation for their specific situations and needs.

216

1. Introduction

The [NIST Cloud Computing Forensic Science Working Group \(NCC FSWG\)](#) previously published NIST IR 8006, *NIST Cloud Computing Forensic Science Challenges* [1], which was the result of collaboration between volunteers from the private and public sector. That document highlighted digital forensic challenges triggered by the specific characteristics and business model of public cloud computing services.

The approach to examining digital forensics in the cloud was to first understand cloud computing technology and to identify and elucidate its essential and unique characteristics, which play a significant part in three aspects of operation: normal operations, adverse operations when cloud computing resources are under attack, and operations during criminal exploitation.

The second phase of this approach was a close examination of the challenges that were identified in the previous NIST report. This examination involved analyzing the Cloud Security Alliance's (CSA's) Enterprise Architecture (EA) [2], its various functional capabilities and processes, and the potential impact of each challenge on performing a forensic investigation if a specific functional capability or process were involved in an attack and breach or were used during criminal exploitation. The analysis presumed fictive use case scenarios that would exploit potential weaknesses, vulnerabilities, exposures, or cloud technology for criminal activities. Such elements are of fundamental concern in forensic analysis as they present points that adversaries may seek to exploit or characteristics that can be used by criminals. In either case, there will be evidence of the attack or criminal exploitation for future forensic analysis. The EA is composed of a large number of specific functional capabilities that enable detailed consideration of the effects of each forensic challenge on each of the capabilities.

The third phase of this work has been to examine the nature of each challenge (i.e., whether the challenge is technological or non-technological) to determine its role and impact on the forensic examination process. As each challenge was analyzed, the applicability of techniques or technologies became clearer in terms of how they function and ultimately contribute to the forensic processes of identification, acquisition, preservation, examination, and interpretation of evidence.

This work brings value by clarifying how forensics in the cloud can achieve the same acceptance as forensics in traditional computing models. This document, the associated research, and NIST IR 8006 [1] proactively address the White House Executive Order of May 12, 2021, entitled *Executive Order on Improving the Nation's Cybersecurity* [3], which points out the importance of having forensic-ready information systems, including cloud systems, to improve the Nation's cybersecurity.

1.1. The Need for a Cloud-specific Forensic Reference Architecture

Digital forensics is the application of science and technology to the discovery and examination of digital artifacts within information systems and networks to establish facts and evidence concerning events and conditions that occur within them. Digital forensics is traditionally used for judicial proceedings and regulatory issues but may also be used for other purposes as described below.

Digital forensics continues to evolve in step with computer and information science. As these technologies, their implementations, and their operations have changed, digital forensics has

adapted. The number of scenarios that may require the application of digital forensic techniques have increased along with the complexity of the underlying architectures .

One common scenario involves the detailed investigation of criminal activities. As computers become widely available and develop greater capabilities, criminal elements worldwide have adopted them as tools to manage their endeavors. These include both “traditional” forms of crime (e.g., violent crime, property crime, drug trafficking, human trafficking, white-collar crime) and crimes that occur in cyberspace (e.g., ransomware attacks, data breaches, identity theft, cyber-terrorism, distributed denial of service, illicit cryptocurrency mining, child pornography, and attacks against governments, key corporations, or power grids). Forensic procedures involve locating and analyzing digital traces that can help solve the crime and/or allow for incident response.

Forensic procedures are also used to investigate civil actions, such as divorce proceedings, asset discovery, insurance claims, lawsuits, and similar cases that often require forensic methods to determine the presence, absence, and movement of data and funds.

An example of how forensic techniques are used involves the collection of a laptop computer while apprehending a presumed perpetrator of an illegal act. The suspected act could involve – for instance – financial exploitation of stolen identities, hacking into a hospital’s records management system to implant ransomware, electronic entry of a corporate system in attempted commercial espionage, or penetrating a government or military computer. Similarly, civil actions can require forensic examination, such as discovering financial assets for a divorce proceeding.

In each of these cases, forensics plays an essential role in determining facts; assisting in the analysis, validation, and authentication of data; and enabling documentation of findings to present to a court and attorneys.

The application of forensic methods may also be required for normal business operations. For example, forensic methods may be employed to recover data that, at first, appears to be lost or destroyed on computer drives. During incident response, additional goals of using forensic methods may include mitigating future cyberattacks, preventing system failure, or minimizing data loss.

In the commercial context, the use of forensics in incident response can help determine the root cause of an outage event, such as a component failure, corrupted software, or intentional sabotage. Other scenarios may involve close examination of system configurations, potentially questionable employee data storage and activities, and operational aspects related to compliance matters. In any of these cases, forensic methods may supply insights that are not available through any other means.

For decades, information processing systems have enabled the storage, processing, and transmission of information for public and private organizations and individuals. The maintenance, operations, and protection of these information systems have become paramount concerns since a disruption of sufficient magnitude or specific type could threaten business activities. In addition, the use of these systems in support of criminal activities has been of major concern.

Industry and government have an array of authoritative sources that guide the design, engineering, and operations of information systems. Each of the frameworks listed below can

provide core support for the design, implementation, assessment, monitoring, and operations of information systems:

- NIST Risk Management Framework (RMF) [4] – A focused guide to information system risk management
- ISO 27000 Series [5] – A series of standards dealing with a wide range of information security topics, such as:
 - ISO/IEC 27001 [6] – Information Security Management
 - ISO/IEC 27002 [7] – Information Security Controls
 - ISO/IEC 27018 [8] – Security of Personally Identifiable Information (PII) in the Cloud
 - ISO/IEC 27035 [9] – Incident Response
 - ISO/IEC 27037 [10] – Digital Evidence Collection and Preservation
- IT Infrastructure Library (ITIL) [11] – A service-oriented architecture (SOA)
- Sherwood Applied Business Security Architecture (SABSA) [12]
- The Open Group Architecture Framework (TOGAF) [13] – A general security framework
- Cloud Security Alliance STAR program [14] – A progressive security certification

The focus of each of these frameworks varies but generally facilitates architecting, implementing, and operating secure and resilient information systems. The RMF is focused on security from a risk identification and management perspective. As varied as the ISO 27000 series [5] is, it contains standards that address digital evidence and incident response. Interestingly, however, there is not a readily apparent, in-depth exploration of cloud-system forensics.

The endeavor presented here deals with the matter of forensics performed within a cloud computing environment. The advent of cloud computing has simplified business operations and introduced a level of business agility not previously experienced with traditional or on-premises computing. However, cloud computing has also introduced a range of security and forensics challenges. Enhanced capabilities enjoyed by legitimate businesses and friendly governments are often equally available to opposing nation-states, terrorist groups, and international criminal elements and assets. As a result, targets that were once unassailable by nefarious actors may now be vulnerable to attack or exploitation.

To a great extent, cloud computing runs on virtualization – that is, the creation of processing resources that have hardware as their basis but run as multiplexed programs and are thus functionally multiplied through it. Cloud forensics involves performing analysis on “virtual machines” using techniques that rely on having “real machines” on which to work. In addition, there is the issue of the information obtained. If the “machine” is essentially “unreal,” what does that say about any evidence derived from it? This evidence is therefore different from traditional digital evidence.

Cloud computing has become increasingly pervasive as more entities discover its advantages. These entities include legitimate businesses, governments, and individuals who use SaaS cloud

platforms, as well as criminal and terrorist organizations and opposing nation-states. For legitimate consumers, cloud computing provides capabilities such as:

- More rapid business continuity and disaster recovery
- More effective incident response
- Improved information access, management, and archiving
- Easier and more immediate collaboration between widely separated individuals and groups

This research has adapted solutions that originated in the on-premises data center to the significant differences presented by the cloud.

As important as they are for addressing significant events related to business operations (as described above), forensic methods have at least equal importance when contributing to matters of compliance, legality, and criminal exploitation. Careful treatment has been given to these questions during this research to ensure that the findings do not merely consider technical aspects but also address the broader aspects of their material application. Unquestionably, close examination of these adverse events is required to understand their incipience and progression and – in particular – to ensure that remediation, event reconstruction, and attribution are effectively and credibly realized.

Thus, it has been the specific focus and goal of this effort to research these issues, examine and clarify the forensic challenges, and ultimately formulate and validate the capabilities required to apply accepted forensic techniques and technologies to this unique computing environment. The result is the Cloud Computing Forensic Reference Architecture.

In as much as a security reference architecture is required to incorporate standards and requirements that will inform system actualization and operation with respect to security, applying the forensic reference architecture will likewise inform that system actualization and operation with the capability to more effectively examine, understand, reconstruct, and remediate the variety of system events and disruptions being experienced.

The goal of the CC FRA is to provide support for a cloud system's forensic readiness. It is meant to help the user understand the cloud forensic challenges that might exist for an organization's cloud system. It identifies which forensic challenges require mitigation strategies and how a forensic investigator would apply that to a particular forensic investigation. The CC FRA presented here will likely evolve over time with more use and research.

1.2. The Approach

The CC FRA builds on several foundational layers. We begin with the understanding that this reference architecture addresses forensics in the context of a cloud computing environment. Building upon the fundamental relationship between security, incident response, and forensics, the CC FRA is designed as an overlay to NIST SP 500-299/SP 800-200, *NIST Cloud Computing Security Reference Architecture* (Draft) [15]. This document discusses the Security Reference Architecture (SRA) and leverages the CSA's Enterprise Architecture (EA). Section 3 provides descriptions of the CSA's EA and its use in the SRA, while Section 4 elaborates on the overlay approach employed for the CC FRA.

Figure 1 depicts the overlaying approach in which cloud functional capabilities comprising the EA are analyzed using the NIST cloud computing forensic challenges to identify the functional capabilities' potential for supporting a cloud system's forensic readiness.




Fig. 1. Forensic Reference Architecture Overlaying Approach

The bottom layer in Figure 1 graphically represents the NIST cloud security reference architecture (SRA). The middle layer represents the NIST cloud forensic challenges. The top layer represents the NIST forensic reference architecture (FRA) described in the current document as an overlay (subset) of the graphical representation of the CSA EA – more precisely, the [CSA TCI v1.1](#), which is the initial version of the CSA's EA (see Appendix C).

In Figure 1, the FRA layer leverages the two layers graphically represented beneath it by analyzing each capability of the SRA (these capabilities being derived from the CSA EA) in the context of the challenges documented in NIST IR 8006 [1]. For each challenge, the analysis determines whether the challenge *affects* the capability if implemented in a cloud environment as part of a cloud service or solution. If the challenge *affects* the capability, then the functional capability is considered to have forensic importance, and it is imported to or considered being a capability of the FRA.

2. Overview of NIST Cloud Forensic Challenges

The [NIST Cloud Computing Forensic Science Working Group \(NCC FSWG\)](#) was established to research forensic science challenges and architectures related to the cloud environment. The Working Group surveyed the literature and identified a set of challenges related to cloud computing forensics. These challenges are presented in NIST IR 8006 [1], where each of 62 challenges is described along with potential results of overcoming each challenge. In addition, the document provides a preliminary analysis of these challenges by including 1) the relationship between each challenge and the five essential characteristics of cloud computing, as defined in the NIST cloud computing model [16]; 2) how the challenges correlate to cloud technology; and 3) nine categories to which the challenges belong. The analysis also considers logging data, data in media, and issues associated with time, location, and sensitive data. In addition, the relevance of topics such as rapid elasticity, multi-tenancy, and hypervisor/virtual machine layers is discussed. These 62 challenges support the criminal justice and civil litigation systems, security incident response, and internal enterprise operations.

The nine categories to which the challenges belong are reproduced below (from NIST IR 8006 [1], pp. 8-9):

1. Architecture (e.g., diversity, complexity, provenance, multi-tenancy, data segregation). Architecture challenges in cloud forensics include:
 - a. Dealing with variability in cloud architectures between providers
 - b. Tenant data compartmentalization and isolation during resource provisioning
 - c. Proliferation of systems, locations, and endpoints that can store data
 - d. Accurate and secure provenance for maintaining and preserving chain of custody
2. Data collection (e.g., data integrity, data recovery, data location, imaging). Data collection challenges in cloud forensics include:
 - a. Locating forensic artifacts in large, distributed, and dynamic systems
 - b. Locating and collecting volatile data
 - c. Data collection from virtual machines
 - d. Data integrity in a multi-tenant environment where data is shared among multiple computers in multiple locations and accessible by multiple parties
 - e. Inability to image all of the forensic artifacts in the cloud
 - f. Accessing the data of one tenant without breaching the confidentiality of other tenants
 - g. Recovery of deleted data in a shared and distributed virtual environment
3. Analysis (e.g., correlation, reconstruction, time synchronization, logs, metadata, timelines). Analysis challenges in cloud forensics include:
 - a. Correlation of forensic artifacts across and within cloud providers
 - b. Reconstruction of events from virtual images or storage
 - c. Integrity of metadata
 - d. Timeline analysis of log data, including synchronization of timestamps
4. Anti-forensics (e.g., obfuscation, data hiding, malware). Anti-forensics are a set of techniques used specifically to prevent or mislead forensic analysis. Anti-forensic challenges in cloud forensics include:
 - a. The use of obfuscation, malware, data hiding, or other techniques to compromise the integrity of evidence

- b. Malware may circumvent virtual machine isolation methods
- 5. Incident first responders (e.g., trustworthiness of cloud providers, response time, reconstruction). Incident first responder challenges in cloud forensics include:
 - a. Confidence, competence, and trustworthiness of the cloud providers to act as first responders and perform data collection
 - b. Difficulty in performing initial triage
 - c. Processing a large volume of collected forensic artifacts
- 6. Role management (e.g., data owners, identity management, users, access control). Role management challenges in cloud forensics include:
 - a. Uniquely identifying the owner of an account
 - b. Decoupling between cloud user credentials and physical users
 - c. Ease of anonymity and creating fictitious identities online
 - d. Determining exact ownership of data
 - e. Authentication and access control
- 7. Legal (e.g., jurisdictions, laws, service level agreements, contracts, subpoenas, international cooperation, privacy, ethics). Legal challenges in cloud forensics include:
 - a. Identifying and addressing issues of jurisdictions for legal access to data
 - b. Lack of effective channels for international communication and cooperation during an investigation
 - c. Data acquisition that relies on the cooperation, competence, and trustworthiness of cloud providers
 - d. Missing terms in contracts and service-level agreements
 - e. Issuing subpoenas without knowledge of the physical location of data
- 8. Standards (e.g., standard operating procedures, interoperability, testing, validation). Standards challenges in cloud forensics include:
 - a. Lack of minimum/basic SOPs, practices, and tools
 - b. Lack of interoperability among cloud providers
 - c. Lack of test and validation procedures
- 9. Training (e.g., forensic investigators, cloud providers, qualification, certification). Training challenges in cloud forensics include:
 - a. Misuse of digital forensic training materials that are not applicable to cloud forensics
 - b. Lack of cloud forensic training and expertise for both investigators and instructors
 - c. Limited knowledge about evidence by record-keeping personnel in cloud providers

3. Overview of CSA's Enterprise Architecture

The Cloud Security Alliance's Enterprise Architecture (CSA's EA) [2] is both a methodology and a set of tools that enable security architects, enterprise architects, and risk management professionals to leverage a common set of solutions and controls. These solutions and controls fulfill common requirements that risk managers must assess regarding the operational status of internal IT security and cloud provider controls. These controls are expressed in terms of security capabilities and designed to create a common roadmap to meet the security needs of businesses.

CSA designed the EA understanding that business requirements must guide the architecture. In the case of the Enterprise Architecture, these requirements come from a controls matrix partly driven by regulations such as Sarbanes-Oxley [17] and Gramm-Leach-Bliley [18], standards frameworks such as ISO-27002 [7], the Payment Card Industry Data Security Standards [19], and the IT Audit Frameworks such as COBIT [20], all in the context of cloud service delivery models such as software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS).

From these requirements, a set of security capabilities have been defined and organized according to the following best practice architecture frameworks. The Sherwood Applied Business Security Architecture (SABSA) [12] defines a security model from a business perspective. The Information Technology Infrastructure Library (ITIL) [11] specifies the schema needed to manage a company's IT services, including the security guidelines to manage those services securely. The Jericho Forum [21] designates technical security specifications that arise from the reality of traditional technology environments in the data center and shift to one where solutions span the internet across multiple data centers, some owned by the business and some purely used as outsourced services. Lastly, The Open Group Architecture Framework (TOGAF) [13] provides an enterprise architecture framework and methodology for planning, designing, and governing information architectures, concluding in a common framework to integrate the work of the security architect with the enterprise architecture of an organization.

The CSA EA is reproduced in Appendix C, and the domains covered are:

1. Business Operation Support Services (BOSS) – These functional capabilities are associated with cloud IT services that support an organization's business needs. BOSS embodies the direction of the business and objectives of the cloud consumer. BOSS capabilities cover compliance, data governance, operational risk management, human resources security, security monitoring, internal investigations, and legal services.
2. Information Technology Operation and Support (ITOS) – These functional capabilities are associated with managing the cloud IT services of an organization. ITOS capabilities cover IT operation, service delivery, and service support.
3. Security and Risk Management (S&RM) – These functional capabilities are associated with safeguarding cloud IT assets and detecting, assessing, and monitoring cloud IT risks. S&RM capabilities cover identity and access management, GRC (governance, risk management, and compliance), policies and standards, threat and vulnerability management, and infrastructure and data protection.
4. Presentation Services – These functional capabilities are associated with the end user interacting with a cloud IT solution. The capabilities cover presentation modalities and presentation platforms (including end points, handwriting, and speech recognition).

5. Application Services – These functional capabilities are associated with the development and use of cloud applications provided by an organization. The capabilities cover programming interfaces, security knowledge life cycle, development processes, integration middleware, connectivity and delivery, and abstraction.
6. Information Services – These functional capabilities are associated with the storage and use of cloud information and data. The capabilities cover service delivery, service support, reporting services, information technology operation and support, business operations and support, data governance, user directory services, risk management, and security monitoring.
7. Infrastructure Services – These functional capabilities are associated with core functions that support the cloud IT infrastructure. The capabilities cover facilities, hardware, networks, and virtual environments.

Together, there are 347 functional capabilities within these domains.

As mentioned above, the CSA's EA functional capabilities are leveraged by the NIST Cloud Security Reference Architecture (SRA) [15], which is comprised of a formal model designed as a security overlay to the NIST Cloud Computing Reference Architecture [22] and a methodology for architecting and orchestrating a cloud-based solution. The methodology allows cloud architects to identify the system's functional capabilities. The orchestration employs a risk-based approach that follows the Risk Management Framework (RMF) [4] applied to cloud-based systems.

The SRA's risk-based approach for determining a cloud actor's responsibilities for implementing specific system components supports a clear delineation between the security responsibilities of cloud providers and consumers and a clear understanding of the customer responsibility matrix. Specifically, for each cloud service model, system components are analyzed to identify the level of involvement of each cloud actor when implementing those components.

4. The Forensic Reference Architecture Methodology

The Cloud Computing Forensic Reference Architecture introduced in this document aims to help the user understand the cloud forensic challenges that might exist for an organization's cloud systems. When architecting or orchestrating a new cloud system, cloud architects and cloud security and forensic practitioners are encouraged to use the CC FRA to identify which challenges could impact the system and therefore require at least partial mitigation strategies to minimize the risk incurred during operations by, for example, allowing real-time interventions based on the proactively generated forensic data and to eliminate potential negative impacts on digital forensic investigations if the need arises.

While the FRA can be used by any cloud computing practitioner, it is specifically designed to help the following target audiences by finding answers for specific questions related to their cloud computing architectures:

- **Target Audience #1: Cloud System Architects and Engineers.** This target audience might ask: *"To what extent does the cloud system I'm designing facilitate the use of digital forensics?"* The architectural methodology and initial architecture presented in this paper can help this audience identify where there could be potential challenges for conducting forensics and can allow them to focus on areas of potential concern. System trade-offs can be considered as well (e.g., the more that a system facilitates the use of forensics, the greater the negative operational or economic impacts might be, or the greater the chance that privacy might be impacted negatively).
- **Target Audience #2: Forensic Practitioners.** This target audience might ask: *"What items do I need to be aware of to conduct digital forensics in the cloud environment versus a traditional or on-premises computing environment?"*
- **Target Audience #3: Consumers Who Want to Procure Cloud Services from Providers.** This target audience might ask: *"What forensic questions and issues do I need to consider when discussing what a cloud provider has to offer?"*

The Cloud Computing Forensic Reference Architecture provides a useful starting point for all cloud security and forensic stakeholders to analyze the extent to which the cloud forensic challenges identified in NIST IR 8006 [1] are impacting their systems.

The 62 forensic challenges and 347 functional capabilities described in Section 2 and Section 3, respectively, provide the basis for determining which capabilities are *affected* by each of the challenges. All possible pairs of challenges and capabilities are considered. The capabilities help focus possible mitigation efforts as follows. If a challenge *affects* a capability, there may be mitigation approaches that can be used to perform better forensics with regard to that capability. Such information could prove useful for forensic practitioners, developers, and researchers.

The [NCC FSWG](#) has developed a mapping between functional capabilities and forensic challenges. For each functional capability, the mapping shows all of the forensic challenges that *affect* that capability. This has resulted in a Mapping Table of 347 rows (one for each capability) and 62 columns (one for each challenge). An entry in the table is YES if the associated challenge *affects* the associated capability; otherwise, the entry is NO. (See Figure 3 for an excerpt of this table.)

When the question is asked: *does a forensic challenge affect a functional capability*, it is defined to mean: *if the challenge were overcome, would that make it easier to conduct a cloud forensic investigation on the considered functional capability?* This is the relationship that the mapping between challenges and capabilities is attempting to capture.

To help answer this question, the [NCC FSWG](#) developed a summary for each of the 62 challenges. This summary answers the following question for each specific challenge: *What advantages would be provided to a forensic investigator if this challenge were overcome (or mitigated)?* If these advantages imply that the quality of forensics that can be performed on the functional capability could be improved, then the answer is *YES*, *overcoming the challenge could make it easier to perform a forensic investigation on the capability*. The summaries for the 62 challenges are found in NIST IR 8006 [1], Annex A, Table 1.

The goal was to provide a narrow, precise mapping between challenges and capabilities. A flowchart was developed that was followed to achieve this mapping, as shown in **Fig. 2**.

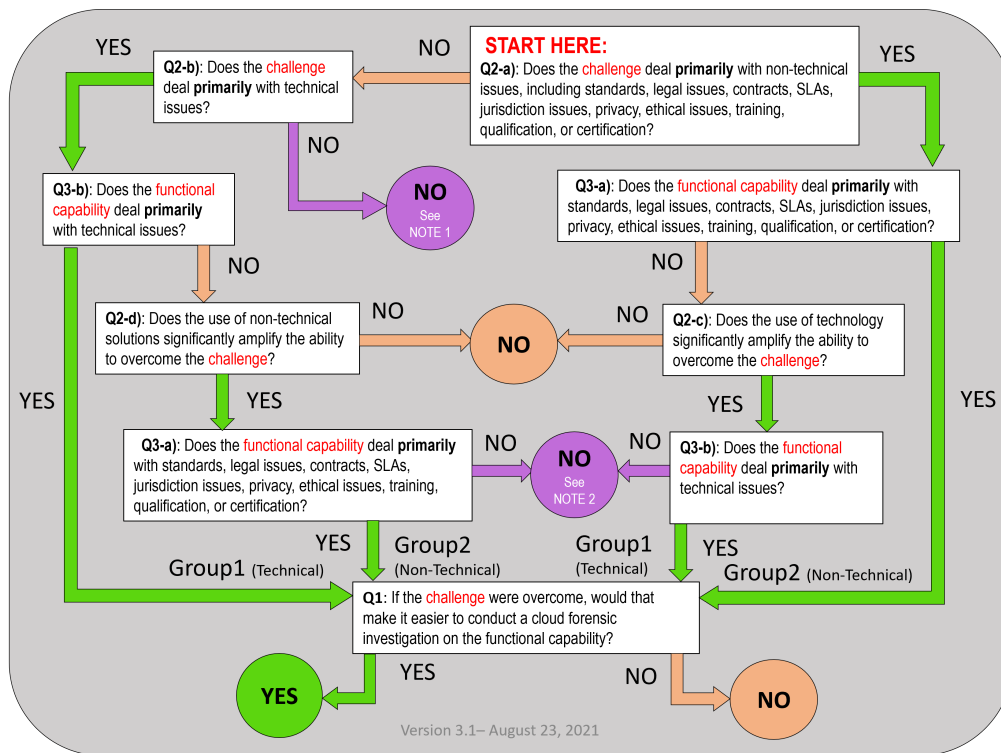


Fig. 2. Mapping Flowchart

The flowchart provides users with a uniform method for determining the applicability of a challenge to a particular capability. In conducting the analysis, the [NCC FSWG](#) placed each cloud forensic challenge into one of two groups: 1) challenges that are primarily technical in nature (e.g., architecture), or 2) challenges that are primarily non-technical in nature (e.g., legal). This led to the creation of questions Q2-a, Q2-b, Q2-c, and Q2-d in the flowchart, which perform the placement into the two groups. If a challenge deals primarily with standards, legal issues, contracts, service-level agreements, jurisdiction issues, privacy, ethical issues, training, qualifications, or certifications, then the challenge is considered non-technical. Otherwise, it is

considered technical. This grouping provides a simple and straightforward method for analyzing the high-level characteristics of each challenge.

Similarly, the [NCC FSWG](#) placed each of the cloud functional capabilities into one of two groups: 1) primarily technical or 2) primarily non-technical in nature. If a capability deals primarily with standards, legal issues, contracts, service-level agreements, jurisdiction issues, privacy, ethical issues, training, qualification, or certification, then the capability is considered non-technical. Otherwise, it is considered technical. This led to the creation of questions Q3-a and Q3-b.

The flowchart attempts to map challenges that are primarily technical only to capabilities that are primarily technical and challenges that are primarily non-technical only to capabilities that are primarily non-technical. This results in a precise and limited mapping. If a challenge and a capability pair are assigned to the same group, the question is asked whether overcoming the challenge makes it easier to conduct forensics on the capability. The answer determines whether the capability is *affected* by the challenge. In summary, if the appropriate grouping is done and overcoming the challenge makes it easier to conduct forensics, then the challenge is considered to *affect* the capability (i.e., the mapping is YES; otherwise, the mapping is NO).

There can, of course, be challenges in one group that affect capabilities in another group, but that does not provide the precise, limited mapping. In such cases, the mapping is considered to be NO.

The following is an example of what is meant by a precise, limited mapping. Suppose the challenge deals with training (e.g., Challenge FC-65: *There is a lack of training materials that educate investigators on cloud computing technology and cloud forensic operating policies and procedures*; see [1], page 52). This is a non-technical challenge. In addition, suppose the capability under consideration is technical. Enhanced training would clearly provide significant benefit to forensic investigators and cloud providers because training is so broadly applicable and would help to perform forensics more easily on most capabilities. However, a cloud forensic architecture in which training *affects* almost every capability is undesirable because then the architecture applies too broadly; most of the capabilities are not *affected* by this challenge in an important way. This makes the architecture less useful because the architecture will have many challenges that *affect* too many capabilities. Rather than this broad mapping of challenges to capabilities, a narrower mapping is preferred. Narrowing the number of capabilities *affected* by the challenge allows the mapping to be more powerful because the challenge can be used as an effective tool of identifying the capabilities that are more likely to be *affected* by the challenge in an important way. The architecture with a narrower mapping is also more practical because the fewer YESs in the mappings, the easier for an investigator to apply the mappings in real-world scenarios.

As described above and shown in Figure 2, if both the challenge and the capability being evaluated deal with the same type of issue (i.e., *technical* or *non-technical*), then the following question is asked: “If the challenge were overcome, would that make it easier to conduct a cloud forensic investigation on the functional capability?” If the answer is “yes,” then the mapping is YES.

However, if the challenge is primarily technical in nature and the capability is non-technical in nature (or vice versa), then an analysis is conducted to determine whether the use of technical or non-technical solutions to implement the capability would significantly enhance the ability of a

forensic investigator to overcome the challenge, as illustrated in questions Q2-c and Q2-d. If the answer to this question is “no,” then no further analysis is required. If the answer to question Q2-c or Q2-d is “yes,” then the analysis will continue to determine: “If the challenge were overcome, would that make it easier to conduct a cloud forensic investigation on the functional capability?”

Using this methodology, it is possible to determine in a well-defined, structured fashion whether it would be easier to conduct a cloud forensic investigation on a functional capability if the forensic challenge were overcome. As a result, the flowchart will help cloud designers, forensic investigators, and other interested parties focus specifically on those functional capabilities that are affected by a specific cloud forensic challenge.

The process of traversing the flowchart involves asking questions about the particular challenge and capability pair that is being analyzed. Starting at the top right of the flowchart (labeled “Q2-a”), each box asks a question about the challenge or the capability. The answer to each question – YES or NO – then leads to either another box with a question or to one of the circles shown in **Table 1**.

Table 1. The meaning of the circles within the flowchart of **Fig. 2**



When following the logical flowchart and answering the guiding questions, if the final answer is a YES marked with a green circle, then the challenge DOES affect the capability.



When following the logical flowchart and answering the guiding questions, if the final answer is a NO marked with an orange circle, then the challenge DOES NOT affect the capability.



When following the logical flowchart and answering the guiding questions, if the final answer is a NO marked with a purple circle, then the challenge DOES NOT affect the capability for reasons explained in NOTE 1 and NOTE 2, below.

To determine whether *the forensic challenge affects the functional capability*, three fundamental types of questions are asked:

1. Question 1 (Q1) – If the challenge were overcome, would that make it easier to conduct a cloud forensic investigation on the functional capability? Note that the term “cloud forensic investigation” means the identification, acquisition, preservation, examination, interpretation, and reporting of potential digital evidence in the cloud. When analyzing Question 1, it is narrowly considered only with regard to the particular functional capability, ignoring all other capabilities as if they do not exist. So, the question really asked is: *If the challenge were overcome, would that make it easier to conduct a cloud forensic investigation on this functional capability only while ignoring other capabilities?*
2. Question 2 (Q2-a, Q2-b, Q2-c, and Q2-d) – These questions relate only to the challenges and not capabilities. The purpose of these questions is to determine whether the challenge deals with technical or non-technical issues and if either technical solutions or non-technical solutions significantly amplify the ability to overcome the challenge.

3. Question 3 (Q3-a and Q3-b) – These questions relate only to the capabilities and not the challenges. The purpose of these questions is to determine whether the capability deals primarily with technical or non-technical issues.

Questions 2 and 3 ask about the issues that a challenge or capability deals with, which are determined as follows. As discussed in Section 2, the [NCC FSWG](#) labeled each of the 62 challenges according to the following nine categories: architecture, data collection, analysis, anti-forensics, incident first responders, role management, legal, standards, and training. The labels for each challenge may be found in [1], Annex A, Table 2, in the columns labeled “Primary Category” and “Related Category.” These categories and the challenge descriptions are used to determine the type of issue each challenge deals with. If the primary issues are standards, legal issues, contracts, service-level agreements, jurisdiction issues, privacy, ethical issues, training, qualification, or certification, then the challenge is considered non-technical. Otherwise, it is considered technical.

Similarly, if a capability deals primarily with standards, legal issues, contracts, service level agreements, jurisdiction issues, privacy, ethical issues, training, qualification, or certification, then the capability is considered non-technical. Otherwise, it is considered technical.

The [NCC FSWG](#) developed consensus answers for all of the questions related to Question 2 and Question 3 in the flowchart. Therefore, when a particular challenge and capability pair was considered, all questions – except for Question 1 – were already answered. This resulted in much more consistent mappings across all challenges and capabilities.

When traversing the flowchart starting at the box labeled “Q2-a,” if a NO node is *not* reached, then the box labeled “Q1” is eventually reached. For any challenge and capability pair, it may lie in one of two groups when Q1 is reached (see Figure 2). As discussed above, Group 1 is the “Technical Group,” and Group 2 is the “Non-technical Group.” They are defined as follows:

• **Group 1** (Technical Group) –

[The *challenge* is technical, **OR** the *challenge* is non-technical but requires technology (at least partially) to overcome the *challenge*.] **AND** [The *functional capability* is technical.]

• **Group 2** (Non-Technical Group) –

[The *challenge* is non-technical, **OR** the *challenge* is technical but requires non-technical solutions (at least partially) to overcome the *challenge*.] **AND** [The *functional capability* is non-technical.]

The reason for these groups – to map technical challenges to technical capabilities and non-technical challenges to non-technical capabilities – was explained above. Once a challenge and capability pair is assigned to the appropriate group, the question of whether overcoming the

challenge makes it easier to conduct forensics on the capability is asked. This determines whether the capability is affected by the challenge. If the grouping is appropriate and overcoming the challenge makes it easier to conduct forensics, then the challenge is considered to affect the capability (i.e., the mapping is YES).

However, suppose a challenge is non-technical but requires technology to overcome the challenge. Examples of non-technical challenges that have both non-technical and technical solutions include the following ([1], Annex A):

- FC-56 (Confidentially and PII) deals with legal/privacy issues (a non-technical challenge). Privacy issues can be resolved with a combination of legal steps (e.g., legislation) and technology steps (privacy-enhancing technologies).
- FC-64 and FC-65 deal with training (non-technical challenges). Training issues can be resolved with better and more widely available training classes, but they can also be resolved with better technology to perform the training.

There are non-technical challenges that require solutions that are non-technical, technical, or a combination of both. If the non-technical challenge requires only a non-technical solution (and the capability is non-technical), it is in Group 2. If it requires only a technical solution (and the capability is technical), it is in Group 1. If it requires both, then it is in Group 1 or Group 2, depending on whether the capability is technical or non-technical.

When a challenge is technical but requires a non-technical solution to overcome the challenge (and the capability is non-technical), then this challenge is in Group 2.

In **Fig. 2**, the two purple circles refer to two notes, as follows:

- NOTE 1: When this circle is reached, the challenge does not fit in either of the two groups. It is neither technical nor non-technical. Fortunately, none of the challenges reach this node as none have this property. This node is included simply for logical completeness of the flowchart, so that every node has both a YES exit path and a NO exit path.
- NOTE 2: When this circle is reached, the capability does not fit in either of the two groups. It is neither technical nor non-technical. There are a few capabilities that reach this node. However, these capabilities do not deal with issues directly related to digital forensics for cloud computing. Instead, they involve controlling physical access to facilities (e.g., using barriers, security patrols, checking physical ID cards, etc.). They also involve mitigating physical threats to facilities, such as installing fire suppression equipment.

The process described in this section, which is employed for the analysis of any pair consisting of a cloud functional capability and a cloud forensic challenge, represents a core component of the CC FRA – the methodology – and can be applied to any set of capability-challenge pairs, either modified from the sets used in this document or adapted from a different architectural framework or empirical data.

5. The Forensic Reference Architecture Data

The data that supplements the CC FRA methodology described in Section 4 represents the result of an analysis performed by the [NCC FSWG](#) members. The methodology was applied to all possible pairings of cloud forensic challenges (62 total challenges) with cloud functional capabilities (347 capabilities). In total, 21,514 challenge-capability pairings were evaluated using the flowchart in Figure 1.

All users of CC FRA data are encouraged to use the data as an initial implementation of the methodology but use their own judgment when employing the CC FRA methodology in the context of their cloud systems and modify or customize NIST's initial dataset for their specific situations and needs.

For example, if the existing capabilities are not appropriate for the user's situation, some or all can be removed, and new ones can be added. Similarly, new challenges appropriate for the user's situation can be added, or those challenges that have been adequately mitigated can be removed. This architectural methodology has the advantage of helping to focus on how challenges can be mitigated because it considers each challenge specifically in the context of affected capabilities.

The results of the NCC FSWG's analysis are summarized in a Mapping Table (MT). An entry in the MT is YES if the associated challenge was identified as *affecting* the paired capability. Otherwise, the entry is NO.

The CC FRA data set provides all interested parties with the responses for every challenge-capability pairing based on the analysis performed by the authors and collaborators of this document. A sample excerpt of the table is displayed in Figure 3. The full CC FRA Mapping Table is available for download (see Appendix D for a partial image and a link for downloading the data).

The CC FRA data has 62 cloud forensic challenges obtained from NISTIR 8006 [1]. In the CC FRA Mapping Table, each cloud forensic challenge is shown across the top row (i.e., Forensic Challenge 1 [FC01], Forensic Challenge 2 [FC02], etc.). In Figure 3, only FC01-FC09 and FC58-FC65 are shown, and the rest of the challenges are hidden for the sake of readability in the figure. See Appendix D for the full Mapping Table.

The CC FRA data has 347 cloud functional capabilities. In the CC FRA Mapping Table, each cloud functional capability is listed on the left column labeled "CAPABILITY" (see Figure 3). The CC FRA data set preserves the grouping of the cloud functional capabilities provided by the CSA EA [2] into "CONTAINERS" and "DOMAINS."

In Figure 3, the first nine capabilities are shown, as are the last nine; the rest are hidden. Each row, therefore, represents a separate capability and includes the following information: the domain of the capability (all of the domains are described in Section 3), the container (the highest-level elements within the architectural diagram in Appendix D¹), the name of the capability, and a description of the capability (not shown in Figure 3 but shown in Appendix D).

¹ The container is a high-level collection of capabilities consisting of related processes and procedures within the domain.

						FC01	FC02	FC03	FC04	FC05	FC06	FC07	FC08	FC09	...HIDDEN...	FC58	FC59	FC60	FC61	FC62	FC63	FC64	FC65
					2a	No	No	Yes	No	No	Yes	No	No	Yes		Yes	No	Yes	No	Yes	Yes	Yes	Yes
	Components descriptions also available on CSA's int https://research.cloudsecurityalliance.org/tci/				2b	Yes	Yes		Yes	Yes		Yes	Yes			Yes	Yes		Yes				
					2c			Yes			Yes			Yes		Yes		No		Yes	No	Yes	No
Index	DOMAIN	CONTAINER	CAPABILITY (process or solution)	3a	3b \ 2d	No	No		No	No		No	No				Yes		Yes				
4	BOSS	Compliance	Intellectual Property	Yes	No	NO*	NO*	NO	NO*	NO*	NO	NO*	NO*	NO		YES	YES	YES	NO	NO	NO	NO	NO
5	BOSS	Data	Handling/ Labeling/	Yes	No	NO*	NO*	NO	NO*	NO*	NO	NO*	NO*	NO		YES	YES	YES	NO	NO	NO	YES	YES
6	BOSS	Data	Clear Desk Policy	Yes	No	NO*	NO*	NO	NO*	NO*	NO	NO*	NO*	NO		NO	NO	NO	NO	NO	NO	NO	NO
7	BOSS	Data	Rules for Information	No	Yes	YES	YES	NO	YES	YES	YES	YES	YES	NO		NO	NO	NO*	NO	NO	NO*	YES	NO*
8	BOSS	Human	Employee Awareness	No	Yes	YES	YES	NO	YES	YES	YES	YES	YES	NO		NO	NO	NO*	NO	NO	NO*	YES	NO*
9	BOSS	Security	Market Threat	No	Yes	YES	YES	NO	YES	YES	YES	YES	YES	NO		NO	NO	NO*	NO	NO	NO*	YES	NO*
10	BOSS	Security	Knowledge Base	No	Yes	YES	YES	NO	YES	YES	YES	YES	YES	NO		NO	NO	NO*	NO	NO	NO	YES	NO*
11	BOSS	Compliance	Audit Planning	Yes	No	NO*	NO*	NO	NO*	NO*	NO	NO*	NO*	NO		YES	YES	YES	NO	NO	NO	NO	NO*
12	BOSS	Compliance	Internal Audits	No	Yes	YES	YES	NO	YES	YES	YES	YES	YES	NO		YES	YES	NO*	NO	NO	NO*	YES	NO*
	...HIDDEN...																						
342	S & RM	Infrastructure	Network	No	Yes	YES	YES	YES	YES	YES	YES	YES	Yes	NO		YES	YES	NO*	YES	YES	NO*	YES	NO*
343	S & RM	Data Protection	Data Lifecycle	No	Yes	YES	YES	YES	YES	YES	YES	YES	YES	NO	NO		YES	YES	NO*	YES	YES	NO*	YES
344	S & RM	Cryptographic	Signature Services	No	Yes	YES	YES	YES	YES	YES	YES	YES	YES	NO	NO		YES	YES	NO*	YES	YES	NO*	YES
345	S & RM	Governance	IT Risk Management	Yes	No	NO*	NO*	NO	NO*	NO*	YES	NO*	NO*	NO		NO	NO	NO	NO	NO	NO	NO	NO
346	S & RM	InfoSec	Risk Portfolio	Yes	No	NO*	NO*	NO	NO*	NO*	YES	NO*	NO*	NO		NO	NO	NO	NO	NO	NO	NO	NO
347	S & RM	Privilege	Authorization Services	No	Yes	YES	YES	YES	YES	YES	YES	YES	YES	NO	NO		YES	YES	NO*	YES	YES	NO*	YES
348	S & RM	Privilege	Authorization Services	No	Yes	YES	YES	YES	YES	YES	YES	YES	YES	NO	NO		YES	YES	NO*	YES	YES	NO*	YES
349	S & RM	Policies and	Information Security	Yes	No	NO*	NO*	NO	NO*	NO*	YES	NO*	NO*	NO		NO	NO	NO	NO	NO	NO	NO	NO
350	S & RM	Privilege	Privilege Usage	No	Yes	YES	YES	YES	YES	YES	YES	YES	YES	NO	NO		YES	YES	NO*	YES	YES	NO*	YES

Fig. 3. Excerpt of the Forensic Reference Architecture (Challenges vs. Capabilities Mapping Table).

The entry in the table that corresponds to a specific row and column (i.e., a specific challenge-capability pair) is either YES or NO based on the result of traversing the mapping flowchart in Figure 2. Traversing the flowchart requires answers to Questions 1 (Q1), 2 (Q2-a, Q2-b, Q2-c, Q2-d), and 3 (Q3-a, Q3-b). As described in Section 4, Q1 must be answered for each individual challenge-capability pair that reaches Q1 when the flowchart is traversed. However, Questions 2 and 3, which relate only to challenges and capabilities separately, can be answered ahead of time, and consensus answers were developed for these by the [NCC FSWG](#). These answers are shown in the table in Figure 3. The second row in the table has the answers for Q2-a, the third row for Q2-b, the fourth row for Q2-c, and the fifth row for Q2-d. The fifth column in the table has the answers for Q3-a and the sixth column for Q3-b.

Each entry in the table is color-coded as follows:

- Orange – A NO is obtained before reaching question Q1 in the flowchart. These entries can be filled in automatically once the answers to questions Q2-a, Q2-b, Q2-c, Q2-d, Q3-a, and Q3-b are entered.
- Red – A NO is obtained as a result of answering Q1.
- Green – A YES is obtained as a result of answering Q1.

Analysis of the correlation between the forensic science challenges and the functional capabilities constitutes the foundation for achieving consistent and repeatable answers to the

questions identified in the CC FRA methodology. Each challenge is further categorized based on its overall *impact* on cloud functional capabilities. This categorization is focused on the overall number of affected capabilities, identifying if only a limited set of capabilities is impacted versus most capabilities composing the cloud ecosystem being impacted. The term *impact* is used to indicate how broadly or narrowly a challenge *affects* the set of functional capabilities. Therefore, the *impact* of each challenge was categorized along a *generic-to-specific* scale as follows (see NIST IR 8006 [1], Annex A, Table 2, column 4):

- *Generic (G)* – A challenge is labeled *generic* if it *affects* most of the capabilities.
- *Specific (S)* – A challenge is labeled *specific* if it *affects* a limited set of capabilities.
- *Quasi (Q)* – A challenge is labeled *quasi* if it falls somewhere between generic and specific.

A *specific* challenge applies narrowly and *affects* only a limited number of capabilities, while a *generic* challenge *affects* a broad set of capabilities. The *specific* challenge *affects* a capability in a direct manner that is determined by the particular issues addressed by the capability. This results in the capability being *affected* in an important and profound way. On the other hand, because the *generic* challenge *affects* most of the capabilities, the *affect* is not tied closely to the issues addressed in each capability, and the capabilities are *affected* in a much less important and profound way. (See Section 4 in which the “precise, limited mapping” is explained.) Thus, a *specific* challenge is more impactful overall than a *generic* one when it comes to conducting a cloud forensic investigation. The *generic-to-specific* label of each challenge is also part of the Forensic Reference Architecture, as shown in Appendix D. The [NCC FSWG](#) developed consensus labels for all of the challenges [1].

6. Conclusion

This document presents the NIST Cloud Computing Forensic Reference Architecture (CC FRA) comprised of:

- a) A methodology for analyzing the functional capabilities of an existing architecture – preferably a security architecture like the Cloud Security Alliance’s (CSA’s) Enterprise Architecture (EA) [2] – through a set of cloud forensic challenges, such as the set identified in NIST IR 8006 [1]
- b) A data set that aggregates the results of the above methodology applied to the CSA’s EA [2] and the NIST IR 8006 [1] set of cloud forensic challenges

The goal of the FRA is to enable the analysis of cloud systems to determine the extent to which a system proactively supports digital forensics. More precisely, the FRA is meant to help users understand how the previously identified cloud forensic challenges might impact an organization’s cloud-based system. When developing a new system or analyzing an existing one, the FRA helps identify those cloud forensic challenges that could affect the system’s capabilities and, therefore, require at least partial mitigation strategies to support a complete forensic investigation. The FRA also identifies how a forensic investigator would apply the mitigation strategies to a particular investigation. While the FRA can be used by any cloud computing practitioner, it is specifically designed to enable cloud system architects, cloud engineers, forensic practitioners, and even cloud consumers to analyze and review their cloud computing architectures for forensic readiness.

The FRA data provided in this document offers an initial implementation of the FRA methodology and a useful starting point for all cloud forensic stakeholders to analyze how the NIST cloud forensic challenges presented in NIST IR 8006 [1] affect each functional capability present in the CSA’s EA [2].

All users are encouraged to customize this initial implementation (shown in Appendix D) for their specific situations and needs. For example, if the existing functional capabilities are not appropriate for the user’s situation, some or all can be removed, and new ones can be added. Similarly, new forensic challenges appropriate for the user’s situation can be added, and challenges that have been adequately mitigated can be removed. The FRA methodology promotes analysis of how cloud forensic challenges affect particular functional capabilities and helps determine whether mitigations are necessary to ensure forensic readiness related to the respective capability. This means that users can replace all cloud forensics challenges or functional capabilities used in the current FRA data set with their own.

The FRA presented here will likely evolve over time, and methods for quantifying impact will be developed to enhance FRA usability.

References

- [1] Herman M, Iorga M, Salim AS, Jackson R, Hurst M, Leo R, Lee R, Landreville N, Mishra AK, Wang Y, Sardinas R (2020). NIST Cloud Computing Forensic Science Challenges. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) 8006. <https://doi.org/10.6028/NIST.IR.8006>
- [2] Cloud Security Alliance Enterprise Architecture. Available at <https://ea.cloudsecurityalliance.org/>
- [3] The White House, Executive Order on Improving the Nation's Cybersecurity, May 12, 2021. Available at <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>
- [4] Joint Task Force (2018). Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-37, Rev. 2. <https://doi.org/10.6028/NIST.SP.800-37r2>
- [5] International Organization for Standardization, ISO 2700 Standards. Available at <https://www.27000.org/index.htm>
- [6] ISO/IEC 27001, Information Technology — Security Techniques — Information Security Management Systems — Requirements, 2013. Available at <https://www.iso.org/standard/54534.html>
- [7] ISO/IEC 27002, Information Security, Cybersecurity and Privacy Protection — Information Security Controls, 2022. Available at <https://www.iso.org/standard/75652.html>
- [8] ISO/IEC 27018, Information Technology — Security Techniques — Code of Practice for Protection of Personally Identifiable Information (PII) in Public Clouds Acting as PII Processors, 2019. Available at <https://www.iso.org/standard/76559.html>
- [9] ISO/IEC 27035-2, Information Technology — Security Techniques — Information Security Incident Management — Part 2: Guidelines to Plan and Prepare for Incident Response, 2016. Available at <https://www.iso.org/standard/62071.html>
- [10] ISO/IEC 27037, Information Technology — Security Techniques — Guidelines for Identification, Collection, Acquisition and Preservation of Digital Evidence, 2012. Available at <https://www.iso.org/standard/44381.html>
- [11] IT Infrastructure Library (ITIL). Available at <https://www.ibm.com/cloud/learn/it-infrastructure-library>
- [12] The SABSA Institute, SABSA Enterprise Security Architecture. Available at <https://sabsa.org/>
- [13] The Open Group, The TOGAF Standard, Version 9.2. Available at <https://www.opengroup.org/togaf>
- [14] Cloud Security Alliance – Security, Trust, Assurance and Risk (STAR). Available at <https://cloudsecurityalliance.org/star>
- [15] NIST Cloud Computing Security Reference Architecture (Draft). (National Institute of Standards and Technology, Gaithersburg, MD). NIST Special Publication (SP) 500-299/800-200. Available at https://github.com/usnistgov/CloudSecurityArchitectureTool-CSAT-v0.1/blob/master/Documents/NIST%20SP%20800-200-SRA_DRAFT_20180414.pdf

- [16] Mell PM, Grance T (2011) The NIST Definition of Cloud Computing. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-145. <https://doi.org/10.6028/NIST.SP.800-145>
- [17] United States Congress, Sarbanes-Oxley Act of 2002, Public Law 107–204, 107th Congress. Available at <https://www.govinfo.gov/content/pkg/PLAW-107publ204/pdf/PLAW-107publ204.pdf>
- [18] Federal Trade Commission, Gramm-Leach-Bliley Act (Financial Services Modernization Act of 1999). Available at <https://www.ftc.gov/tips-advice/business-center/privacy-and-security/gramm-leach-bliley-act>
- [19] PCI Security Standards Council, Payment Card Industry (PCI) Security. Available at https://www.pcisecuritystandards.org/pci_security/
- [20] ISACA, COBIT – Control Objectives for Information Technologies. Available at <https://www.isaca.org/resources/cobit>
- [21] Jericho Forum. Available at https://en.wikipedia.org/wiki/Jericho_Forum (information aggregator) or <https://publications.opengroup.org/catalogsearch/result/?q=jericho+security+reference+architecture>
- [22] Liu F, Tong J, Mao J, Bohn R, Messina J, Badger L, Leaf D (2011). NIST Cloud Computing Reference Architecture. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 500-292. <https://doi.org/10.6028/NIST.SP.500-292>
- [23] SWGDE Digital and Multimedia Evidence (Digital Forensics) as a Forensic Science Discipline, Version 2.0, September 5, 2014. Available at <https://drive.google.com/file/d/1OBux0n7VZQe7HSgObwAtmhz5LgwwX0oY/view>
- [24] ISO/IEC 2382, Information technology - Vocabulary, 2015. Available at <https://www.iso.org/standard/63598.html>
- [25] Scarfone K, Souppaya M, Hoffman P (2011). Guide to Security for Full Virtualization Technologies. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-125. <https://doi.org/10.6028/NIST.SP.800-125>

Appendix A. Acronyms

Selected acronyms and abbreviations used in this paper are defined below.

BOSS

Business Operation Support Services

CC FRA

Cloud Computing Forensic Reference Architecture

COBIT

Control Objectives for Information Technologies

CSA

Cloud Security Alliance

EA

Enterprise Architecture

FC

Forensic Challenge

FISMA

Federal Information Security Modernization Act

FRA

Forensic Reference Architecture

GRC

Governance, Risk management, and Compliance

IaaS

Infrastructure as a Service

ID

Identification

IEC

International Electrotechnical Commission

ISACA

Information Systems Audit and Control Association

ISO

International Organization for Standardization

ITIL

Information Technology Infrastructure Library

ITL

Information Technology Laboratory

ITOS

Information Technology Operation and Support

NCC FSWG

NIST Cloud Computing Forensic Science Working Group

NIST IR

NIST Interagency or Internal Report

985	NIST SP
986	NIST Special Publication
987	OMB
988	Office of Management and Budget
989	PaaS
990	Platform as a Service
991	PCI
992	Payment Card Industry
993	PII
994	Personally Identifiable Information
995	Rev.
996	Revision
997	RMF
998	Risk Management Framework
999	S&RM
1000	Security and Risk Management
1001	SaaS
1002	Software as a Service
1003	SABSA
1004	Sherwood Applied Business Security Architecture
1005	SLA
1006	Service Level Agreement
1007	SOA
1008	Service-Oriented Architecture
1009	SOP
1010	Standard Operating Procedure
1011	SRA
1012	Security Reference Architecture
1013	STAR
1014	Security, Trust, Assurance and Risk
1015	SWGDE
1016	Scientific Working Group on Digital Evidence
1017	TOGAF
1018	The Open Group Architecture Framework

1019 **Appendix B. Glossary**

1020 **challenge**

1021 For this paper, a currently difficult or impossible task that is either unique to cloud computing or exacerbated by it.

1022 **cloud computing**

1023 A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing
1024 resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released
1025 with minimal management effort or service provider interaction. This cloud model is composed of five essential
1026 characteristics, three service models, and four deployment models. [16]

1027 **cloud consumer**

1028 A person or organization that maintains a business relationship with and uses service from cloud providers. [22]

1029 **cloud provider**

1030 The entity (a person or an organization) responsible for making a service available to interested parties. [22,
1031 adapted]

1032 **criminal exploitation**

1033 The exploitation of computing resources by criminals. Criminal activities are planned and/or carried out using these
1034 computing resources.

1035 **digital forensics**

1036 The process used to acquire, preserve, analyze, and report on digital evidence using scientific methods that are
1037 demonstrably reliable, accurate, and repeatable such that it may be used in judicial proceedings. [23, adapted]

1038 **flowchart**

1039 A diagram that shows step-by-step progression through a process using boxes to show the steps and connecting
1040 arrows between the boxes to show their order.

1041 **forensic investigator**

1042 A person who is an expert in acquiring, preserving, analyzing, and presenting digital evidence from computers and
1043 other digital media. This evidence may be related to both computer-based and non-cybercrimes, including security
1044 threats, cyber-attacks, and other illegal activities.

1045 **forensic readiness**

1046 The ability to collect digital evidence effectively and quickly with minimal investigation costs. This involves being
1047 able to define the digital evidence required to reconstruct past computing events of interest.

1048 **functional capability**

1049 Cloud processes or solutions in the Cloud Security Alliance's Enterprise Architecture that cover business operations,
1050 IT operations, security and risk management, presentation services, application services, information services, and
1051 infrastructure services. [2, adapted]

1052 **incident response**

1053 The mitigation of violations of security policies and recommended practices. Addressing and managing the
1054 consequences of a security breach or cyberattack.

1055 **mapping**

1056 An operation that associates each element of a given set with one or more elements of a second set.

1057 **security**

1058 Measures and controls that ensure the confidentiality, integrity, and availability of the information processed and
1059 stored by a computer.

1060 **virtual machine**

1061 A virtual data processing system that appears to be at the exclusive disposal of a particular user but whose functions
1062 are accomplished by sharing the resources of a real data processing system. [24]

1063 **virtualization**

1064 The simulation of the software and/or hardware upon which other software runs. This simulated environment is
1065 called a virtual machine. [25, adapted]

1066 **Appendix C. CSA's Enterprise Architecture**

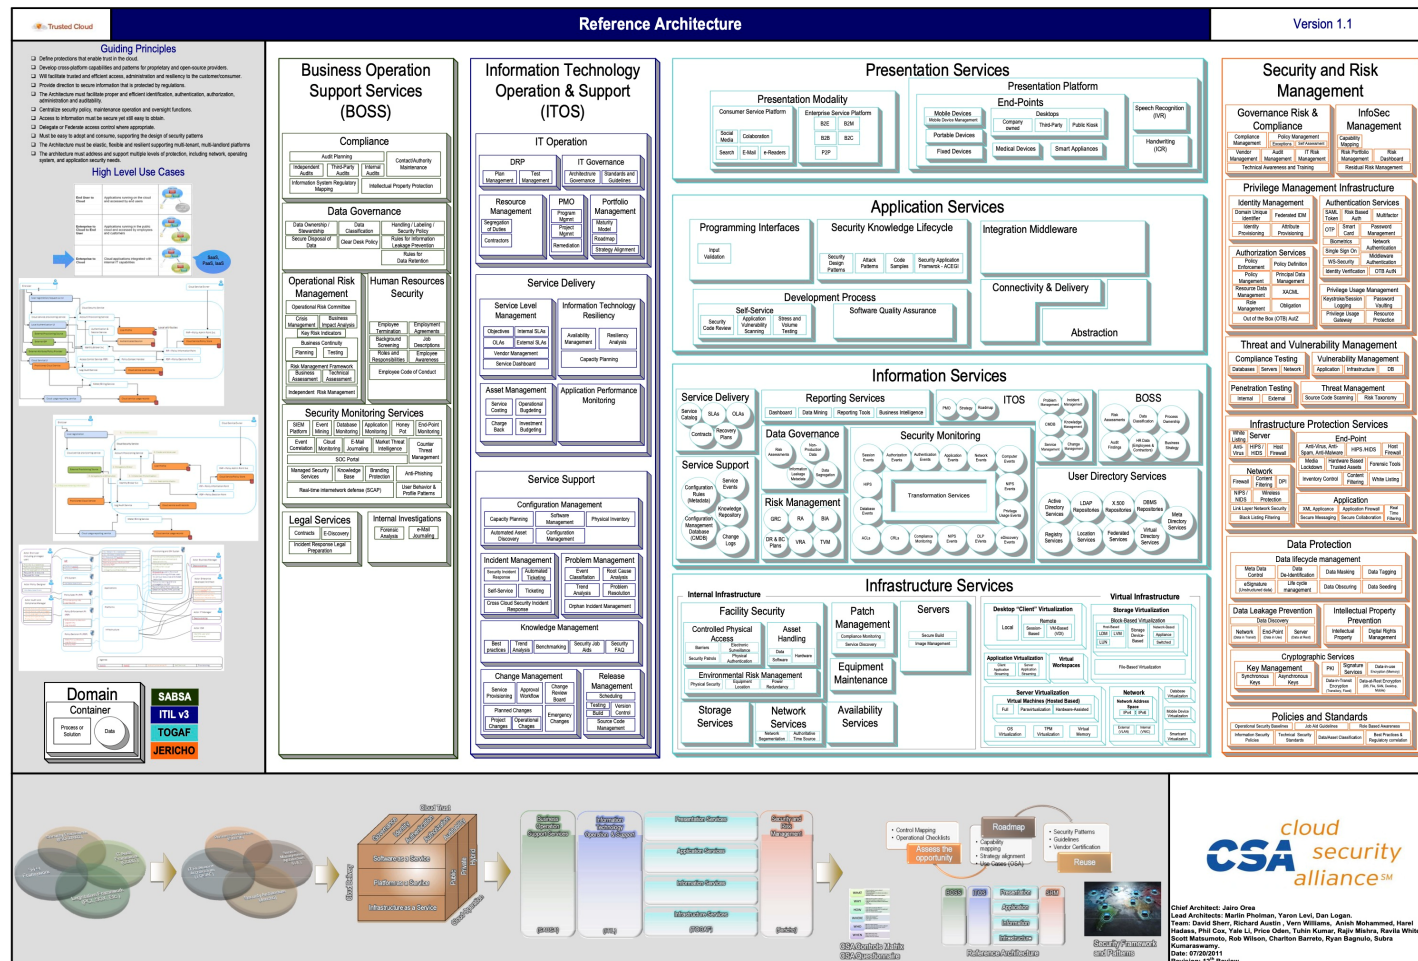


Fig. 4. CSA's Enterprise Architecture (v1.1)

The CSA's Enterprise Architecture v1.1 and v2.0 are available for download as PDF files that can be easily enlarged for further review at NIST's FRA [GitHub repository](#) and the [NCC FSWG website](#).

Section 5 of this document describes how the FRA methodology can be applied to analyze and review the functional capabilities of a cloud system by using a known set of forensic challenges to determine forensic readiness as related to these capabilities. To demonstrate its use, NIST provides an initial implementation of the FRA methodology by generating the FRA data set captured in the [workbook](#) available for download at the [FRA’s GitHub repository](#) or the [NCC FSWG website](#). The [workbook](#) contains the summary of data analyzed by the NIST Cloud Computing Forensic Science Working Group using the FRA methodology that leverages [NIST IR 8006, NIST Cloud Forensic Science Challenges](#), applied to the Cloud Security Alliance’s Enterprise Architecture. The FRA dataset can be found under the “Capabilities vs. Challenges Data” tab of the downloadable [workbook](#).

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	AA	AB	AC	AD	AE	AF	AG	AH	AI	AJ	AK	AL	AM	AN	AO	AP	AQ	AR	AS	AT	AU	AV	AW	AX	AY	AZ	BA	BB	BC	BD	BE	BF	BG	BH	BI	BJ	BK	BL	BM	BN	BO	BP	BQ	BR	BS	BT	BU	BV	BW	BX	BY	BZ	CA	CB	CC	CD	CE	CF	CG	CH	CI	CJ	CK	CL	CM	CN	CO	CP	CQ	CR	CS	CT	CU	CV	CW	CX	CY	CZ	DA	DB	DC	DD	DE	DF	DG	DH	DI	DJ	DK	DL	DM	DN	DO	DP	DQ	DR	DS	DT	DU	DV	DW	DX	DY	DZ	EA	EB	EC	ED	EE	EF	EG	EH	EI	EJ	EK	EL	EM	EN	EO	EP	EQ	ER	ES	ET	EU	EV	EW	EX	EY	EZ	FA	FB	FC	FD	FE	FF	FG	FH	FI	FJ	FK	FL	FM	FN	FO	FP	FQ	FR	FS	FT	FU	FV	FW	FX	FY	FZ	GA	GB	GC	GD	GE	GF	GG	GH	GI	GJ	GK	GL	GM	GN	GO	GP	GQ	GR	GS	GT	GU	GV	GW	GX	GY	GZ	HA	HB	HC	HD	HE	HF	HG	HH	HI	HJ	HK	HL	HM	HN	HO	HP	HQ	HR	HS	HT	HU	HV	HW	HX	HY	HZ	IA	IB	IC	ID	IE	IF	IG	IH	II	IJ	IK	IL	IM	IN	IO	IP	IQ	IR	IS	IT	IU	IV	IW	IX	IY	IZ	JA	JB	JC	JD	JE	JF	JG	JH	JI	JK	JL	JM	JN	JO	JP	JQ	JR	JS	JT	JU	JV	JW	JX	JY	JZ	KA	KB	KC	KD	KE	KF	KG	KH	KI	KJ	KL	KM	KN	KO	KP	KQ	KR	KS	KT	KU	KV	KW	KX	KY	KZ	LA	LB	LC	LD	LE	LF	LG	LH	LI	LJ	LK	LL	LM	LN	LO	LP	LQ	LR	LS	LT	LU	LV	LW	LX	LY	LZ	MA	MB	MC	MD	ME	MF	MG	MH	MI	MJ	MK	ML	MM	MN	MO	MP	MQ	MR	MS	MT	MU	MV	MW	MX	MY	MZ	NA	NB	NC	ND	NE	NF	NG	NH	NI	NJ	NK	NL	NM	NO	NP	NQ	NR	NS	NT	NU	NV	NW	NX	NY	NZ	OA	OB	OC	OD	OE	OF	OG	OH	OI	OJ	OK	OL	OM	ON	OO	OP	OQ	OR	OS	OT	OU	OV	OW	OX	OY	OZ	PA	PB	PC	PD	PE	PF	PG	PH	PI	PJ	PK	PL	PM	PN	PO	PP	PQ	PR	PS	PT	PU	PV	PW	PX	PY	PZ	QA	QB	QC	QD	QE	QF	QG	QH	QI	QJ	QK	QL	QM	QN	QO	QP	QQ	QR	QS	QT	QU	QV	QW	QX	QY	QZ	RA	RB	RC	RD	RE	RF	RG	RH	RI	RJ	RK	RL	RM	RN	RO	RP	RQ	RR	RS	RT	RU	RV	RW	RX	RY	RZ	SA	SB	SC	SD	SE	SF	SG	SH	SI	SJ	SK	SL	SM	SN	SO	SP	SQ	SR	SS	ST	SU	SV	SW	SX	SY	SZ	TA	TB	TC	TD	TE	TF	TG	TH	TI	TJ	TK	TL	TM	TN	TO	TP	TQ	TR	TS	TT	TU	TV	TW	TX	TY	TZ	UA	UB	UC	UD	UE	UF	UG	UH	UI	UJ	UK	UL	UM	UN	UO	UP	UQ	UR	US	UT	UU	UV	UW	UX	UY	UZ	VA	VB	VC	VD	VE	VF	VG	VH	VI	VJ	VK	VL	VM	VN	VO	VP	VQ	VR	VS	VT	VU	VV	VW	VX	VY	VZ	WA	WB	WC	WD	WE	WF	WG	WH	WI	WJ	WK	WL	WM	WN	WO	WP	WQ	WR	WS	WT	WU	WV	WW	WX	WY	WZ	XA	XB	XC	XD	XE	XF	XG	XH	XI	XJ	XK	XL	XM	XN	XO	XP	XQ	XR	XS	XT	XU	XV	XW	XX	XY	XZ	YA	YB	YC	YD	YE	YF	YG	YH	YI	YJ	YK	YL	YM	YN	YO	YP	YQ	YR	YS	YT	YU	YV	YW	YX	YY	YZ	ZA	ZB	ZC	ZD	ZE	ZF	ZG	ZH	ZI	ZJ	ZK	ZL	ZM	ZN	ZO	ZP	ZQ	ZR	ZS	ZT	ZU	ZV	ZW	ZX	ZY	ZZ																																																																																																																																																																																																		
							NIIST DEFINITION OF FORENSICS SCIENCE: "Digital Computing forensic science is the application of scientific principles, technological practices and derived and proven methods to process past cloud computing events through: — identification, — collection, — preservation, — examination and — reporting of digital data for the purposes of facilitating the reconstruction of these events." When considering Q1, consider it narrowly only with regard the particular functional capability, proving all other capabilities as if they don't exist.																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																						