

NIST Special Publication 800-210

General Access Control Guidance for Cloud Systems

Vincent C. Hu
Michaela Iorga
Wei Bao
Ang Li
Qinghua Li
Antonios Gouglidis

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-210>

C O M P U T E R S E C U R I T Y

NIST
National Institute of
Standards and Technology
U.S. Department of Commerce

NIST Special Publication 800-210

General Access Control Guidance for Cloud Systems

Vincent C. Hu
Michaela Iorga
*Computer Security Division
Information Technology Laboratory*

Wei Bao
Ang Li
Qinghua Li
*Department of Computer Science and Computer Engineering
University of Arkansas
Fayetteville, AR*

Antonios Gouglidis
*School of Computing and Communications
Lancaster University
Lancaster, United Kingdom*

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-210>

July 2020



U.S. Department of Commerce
Wilbur L. Ross, Jr., Secretary

National Institute of Standards and Technology
Walter Copan, NIST Director and Under Secretary of Commerce for Standards and Technology

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.

National Institute of Standards and Technology Special Publication 800-210
Natl. Inst. Stand. Technol. Spec. Publ. 800-210, 34 pages (July 2020)
CODEN: NSPUE2

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.SP.800-210>

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 NIST, 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>.

Comments on this publication may be submitted to:

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930
Email: sp800-210-comments@nist.gov

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

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.

Abstract

This document presents cloud access control characteristics and a set of general access control guidance for cloud service models: IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service). Different service delivery models require managing different types of access on offered service components. Such service models can be considered hierarchical, thus the access control guidance of functional components in a lower-level service model are also applicable to the same functional components in a higher-level service model. In general, access control guidance for IaaS is also applicable to PaaS and SaaS, and access control guidance for IaaS and PaaS is also applicable to SaaS. However, each service model has its own focus with regard to access control requirements for its service.

Keywords

access control; access control mechanism; Cloud; cloud systems; policy; authorization ABAC; RBAC.

Acknowledgements

The authors, Vincent C. Hu and Michaela Iorga of the National Institute of Standards and Technology (NIST), Bao Wei, Ang Li, and Qinghua Li of Department of Computer Science and Computer Engineering University of Arkansas, and Antonios Gouglidis of School of Computing and Communications Lancaster University wish to thank Isabel Van Wyk, Jim Foti, and David Ferraiolo (NIST) who reviewed drafts of this document. The authors also gratefully acknowledge and appreciate the comments and contributions made by government agencies, private organizations, and individuals in providing direction and assistance in the development of this document.

Patent Disclosure Notice

NOTICE: The Information Technology Laboratory (ITL) has requested that holders of patent claims whose use may be required for compliance with the guidance or requirements of this publication disclose such patent claims to ITL. However, holders of patents are not obligated to respond to ITL calls for patents and ITL has not undertaken a patent search in order to identify which, if any, patents may apply to this publication.

As of the date of publication and following call(s) for the identification of patent claims whose use may be required for compliance with the guidance or requirements of this publication, no such patent claims have been identified to ITL.

No representation is made or implied by ITL that licenses are not required to avoid patent infringement in the use of this publication.

Executive Summary

Cloud systems have been developed over time and conceptualized through a combination of software, hardware components, and virtualization technologies. Characteristics of the cloud, such as resource pooling, rapid elasticity, and pay-as-you-go services, accelerated its wide adoption by industry, government, and academia. Specifically, cloud systems offer application services, data storage, data management, networking, and computing resources management to consumers over a network (the internet in general). Despite the great advancements of cloud systems, concerns have been raised about the offered level of security and privacy. The importance of these concerns becomes more evident when considering the increasing number of users who have adopted cloud services.

This document presents cloud access control (AC) characteristics and a set of general access control guidance for cloud service models—IaaS (Infrastructure as a Service), PaaS (Platform as a Service), and SaaS (Software as a Service). The main focus is on technical aspects of access control without considering deployment models (e.g., public, private, hybrid clouds etc.), as well as trust and risk management issues, which require different layers of discussions that depend on the security requirements of the business function or the organization of deployment for which the cloud system is implemented. Different service delivery models need to consider managing different types of access on offered service components. Such considerations can be hierarchical, such as how the access control considerations of functional components in a lower-level service model (e.g., networking and storage layers in the IaaS model) are also applicable to the same functional components in a higher-level service model (e.g., networking and storage in PaaS and SaaS models). In general, access control considerations for IaaS are also applicable to PaaS and SaaS, and access control considerations for IaaS and PaaS are also applicable to SaaS. Therefore, AC guidance for IaaS is applicable to PaaS and SaaS, and AC guidance for IaaS and PaaS is also applicable to SaaS. However, each service model has its own focus with regard to access control requirements for its service.

Table of Contents

Executive Summaryiv

1 Introduction 1

 1.1 Purpose 1

 1.2 Scope 1

 1.3 Audience 1

 1.4 Document Structure 1

2 Cloud Access Control Characteristics 3

3 Access Control Guidance for IaaS 8

 3.1 Guidance for Network 8

 3.2 Guidance for Hypervisor 8

 3.3 Guidance for Virtual Machines 9

 3.4 Guidance for APIs 9

 3.5 Recommendations for IaaS Access Control 9

4 Access Control Guidance for PaaS 11

 4.1 Guidance for Memory Data 11

 4.2 Guidance for APIs 11

 4.3 Recommendations for PaaS Access Control 11

5 Access Control Guidance for SaaS 13

 5.1 Guidance for Data Owner’s Control 13

 5.2 Guidance for Confidentiality 13

 5.3 Guidance for Privilege Management 14

 5.4 Guidance for Multiple Replicas of Data 14

 5.5 Guidance for Multi-tenancy 14

 5.6 Guidance for Attribute and Role Management 14

 5.7 Guidance for Policies 15

 5.8 Guidance for APIs 15

 5.9 Recommendations for SaaS Access Control 15

6 Access Control Guidance for Inter- and Intra- Operation 18

7 Conclusions 20

References 21

List of Appendices

Appendix A— Guidance and SP 800-53 Revision 4 Access Control (AC) Family Mapping 25

List of Figures

Figure 1: The general architecture of a cloud system..... 4
Figure 2: The service models of a cloud system..... 5
Figure 3: Accesses controlled by the cloud service provider and the consumer 6
Figure 4: Multiple applications and users of an SaaS provider..... 13
Figure 5: The external collaboration (inter-operation) between different Clouds..... 18
Figure 6: The internal collaboration (intra-operation) within the same cloud..... 19

List of Tables

Table 1: Potential policy rules expressed by Subject, Operation, Object for IaaS AC policy 10
Table 2: Potential policy rules expressed by Subject, Operation, Object for PaaS AC policy..... 12
Table 3: Potential policy rules expressed by Subject, Operation, Object for SaaS AC policy..... 17
Table 4 Mapping the cloud access control guidance to the AC controls listed in NIST SP 800-53, Revision 4..... 25

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.800-210>

1 Introduction

1.1 Purpose

Access control (AC) dictates how subjects (i.e., users and processes) can access objects based on defined AC policies to protect sensitive data and critical computing objects in the cloud systems. Considering the heterogeneity and remote nature of the cloud service models, AC and its general concepts should be revisited. In recent years, many works have focused on AC in cloud systems [23, 25, 26, 27]. However, these are primarily ad hoc solutions targeted at specific cloud applications and do not provide comprehensive views of cloud AC.

This document presents a set of general AC guidance for cloud service models independent from its deployment models because it requires another layer of access control that depends on the security requirements of the business function for which the cloud system is used. As shown in Figure 3, different cloud service models require the management of access to different components of the offered service. Since such cloud service models can be considered hierarchical, the AC considerations of functional components in a lower-level (according to Figure 2) service model (e.g., networking and storage layers in the Infrastructure as a Service (IaaS) model) are also applicable to the same functional components in a higher-level service model (e.g., networking and storage in Platform as a Service (PaaS) and Software as a Service (SaaS) models). In general, AC considerations for IaaS are also applicable to PaaS and SaaS, and AC considerations for IaaS and PaaS are also applicable to SaaS. Thus, AC guidance for IaaS is applicable to PaaS and SaaS, and AC guidance for IaaS and PaaS is also applicable to SaaS. However, each service model has its own focus with regard to AC. For instance, an IaaS provider may put more effort into virtualization control, and in addition to the virtualization control, a SaaS provider needs to consider data security and the privacy of services it provides.

1.2 Scope

This document focuses on providing guidance for access control systems that are applicable to an organization's cloud implementation and security management. It does not prescribe the internal cloud access control standards that an organization may need in their enterprise systems or within a community other than the organization itself.

1.3 Audience

The intended audience for this document is an organizational entity that implements access control solutions for sharing information in cloud systems. This document assumes that readers are familiar with the cloud and access (authorization) control systems and have basic knowledge of operating systems, databases, networking, and security. Given the constantly changing nature of the information technology (IT) industry, readers are strongly encouraged to take advantage of other documents—including those listed in this document—for more current and detailed information.

1.4 Document Structure

The sections and appendix presented in this document are as follows:

- Section 1 states the purpose and scope of access control and cloud systems.
- Section 2 provides an overview of cloud access control characteristics.
- Section 3 discusses guidance for access control systems for IaaS (Infrastructure as a Service).
- Section 4 discusses guidance for access control systems for PaaS (Platform as a Service).
- Section 5 discusses guidance for access control systems for SaaS (Software as a Service).
- Section 6 discusses guidance for access control systems for inter- and intra-cloud operations.
- Section 7 concludes the document with future directions.

2 Cloud Access Control Characteristics

With the support of different service models, cloud systems can provide a wide range of services to its end-users, developers, and system administrators. Cloud systems have been developed over time and conceptualized through a combination of software, hardware components, and virtualization technologies. Characteristics of the cloud, such as resource pooling, rapid elasticity, and pay-as-you-go services, have accelerated its wide adoption by industry, government, and academia. Specifically, cloud systems offer application services, data storage, data management, networking, and computing resources management to consumers¹ over a network (and the internet in general). Examples of popular cloud applications include web-based email services (e.g., Google's Gmail, Microsoft's Office 365 Outlook), data storage (e.g., Google Drive, Microsoft's OneDrive, Dropbox) for end users, and consumer relationship management and business intelligence systems (e.g., Customer Relationship Management (CRM) Cloud, Workday) for business management. Despite the great advancements of cloud systems, concerns have been raised about offered levels of security and privacy. The importance of these concerns becomes more evident when considering the increasing number of users that have adopted cloud services [1].

NIST publications defines cloud computing as “a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” [2,3]. Cloud deployment models (e.g., public cloud, private cloud, community cloud, hybrid cloud, etc.) are configured by the scope of cloud users, services, and resources based on service requirements, they may be deployed privately, hosted on the premises of a cloud consumer or provider's dedicated infrastructure, or hosted publicly by one or more cloud service providers. The system may be configured and used by one consumer or a group of trusted partners or support multi-tenancy and be used publicly by different end users who acquire the service. Depending on the type of cloud deployment model, the cloud may have limited private computing resources or access to large quantities of remotely accessed resources. The different deployment models present a number of trade-offs in how consumers can control their resources as well as the scale, cost, and availability of those resources [4]. As depicted in Figure 1, the architecture of a cloud system is composed, in general, by layers of functions:

- VM (Virtual Machine), including:
 - Applications
 - Application Programming Interface (API)
 - Operating System (OS)
- Hypervisor
- Storage
- Networking
- Hardware

¹ Cloud service **consumers** play various roles in the consumption of the cloud services, e.g. system planners, program managers, technologists. **End-users** are individuals using cloud services as direct clients of a cloud provider, of a cloud consumer leveraging a cloud service, or individuals employed by a cloud consumer. A **user** is in a generic term associated with any entity using the cloud service. Depending on scenario, the user can be referred as either cloud service consumer or end-user where applicable.

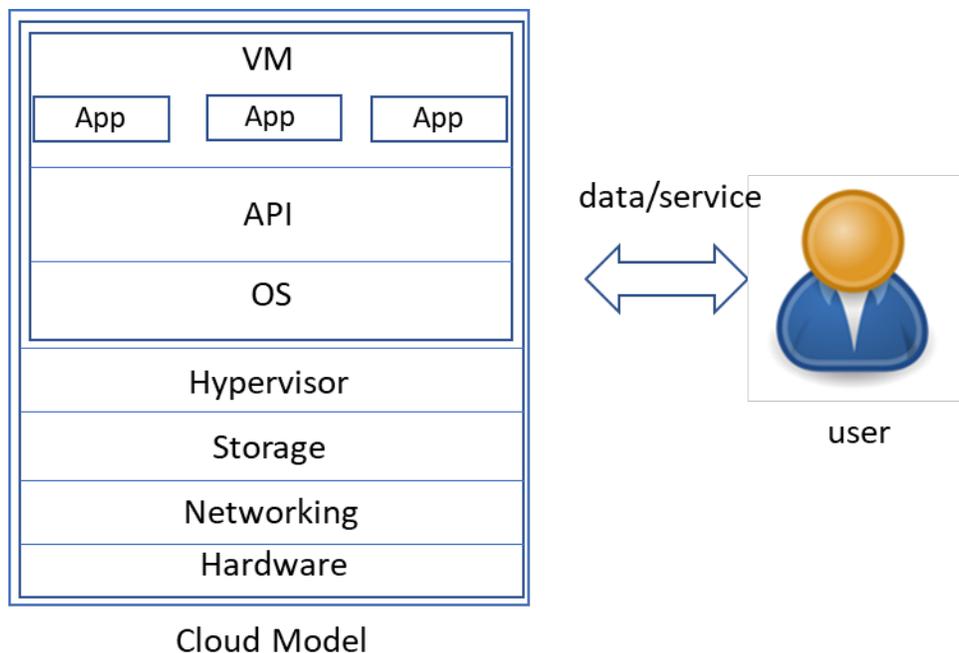


Figure 1: The general architecture of a cloud system

A cloud service can provide access to software applications, such as email or office productivity tools (i.e., the Software as a Service, or SaaS, service model); an environment for consumers to build and operate their own software (i.e., the Platform as a Service, or PaaS, service model), or network access to virtualized computing resources such as processing power and storage (i.e., the Infrastructure as a Service, or IaaS, service model). The different service models have different strengths and are suitable for different consumers and business objectives [4], as illustrated in Figure 2, the arrows show the support relations between models.

A cloud system that deploys the SaaS model can be accessible over a network by an end user utilizing various client devices (e.g., a thin client interface, such as a web browser, for accessing a web-based email application) or via a program with the correct set of interfaces whose execution would enable communication with a cloud application. In the SaaS model, an application user is limited to user-specific application configuration settings and does not manage or control the underlying cloud infrastructure, which typically includes the network, servers, operating systems, storage, or individual applications.

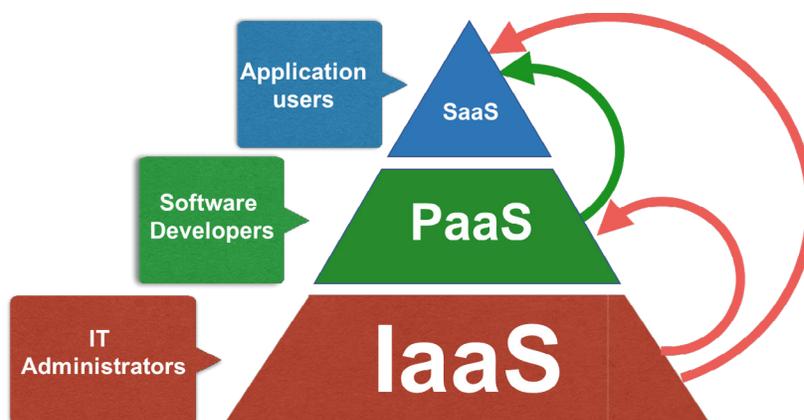


Figure 2: The service models of a cloud system

The PaaS model in a cloud system allows developers to create and deploy applications onto the cloud infrastructure using programming languages, libraries, services, and tools. A software developer does not manage or control the underlying cloud infrastructure but has control over the deployed applications (software) and, possibly, configuration settings for the application-hosting environment.

When analyzing the responsibilities between consumer and cloud service providers for protecting cloud data, it is not always clear-cut, if a IaaS systems provide only the computation resources, or offers also the virtualized storage, and network resources to consumers for deploying and running arbitrary software, including operating systems and applications. The consumer may in turn have control over virtual storage, virtualized network components, and the ability to deploy their own VMs and applications given access provisioned by the cloud service provider.

The shared responsibility of access control needs to be considered in the PaaS and SaaS model [42]; For example software developers might need to access data in systems provided by PaaS for their developmental needs, and internal application users (i.e., users that need to access the application system data) might need to access application system data that is managed by the applications. In general, for PaaS, consumer software developers might share access control responsibilities with cloud service providers; for SaaS, internal application users might share such responsibilities with cloud service providers.

Note that unless there is express prior approval from the consumer, a PaaS or SaaS provider must manage access control with the IaaS provider and the consumer (if it is not also the IaaS provider). If the consumer approves, the provider should inform the consumer of its intention to store the specified data in the IaaS provider, where it will be accessed as well as the extent to which the data can be accessed by the IaaS provider, foreign entities, or authorities. A public consultation and hearing process must then be conducted before a decision is made.

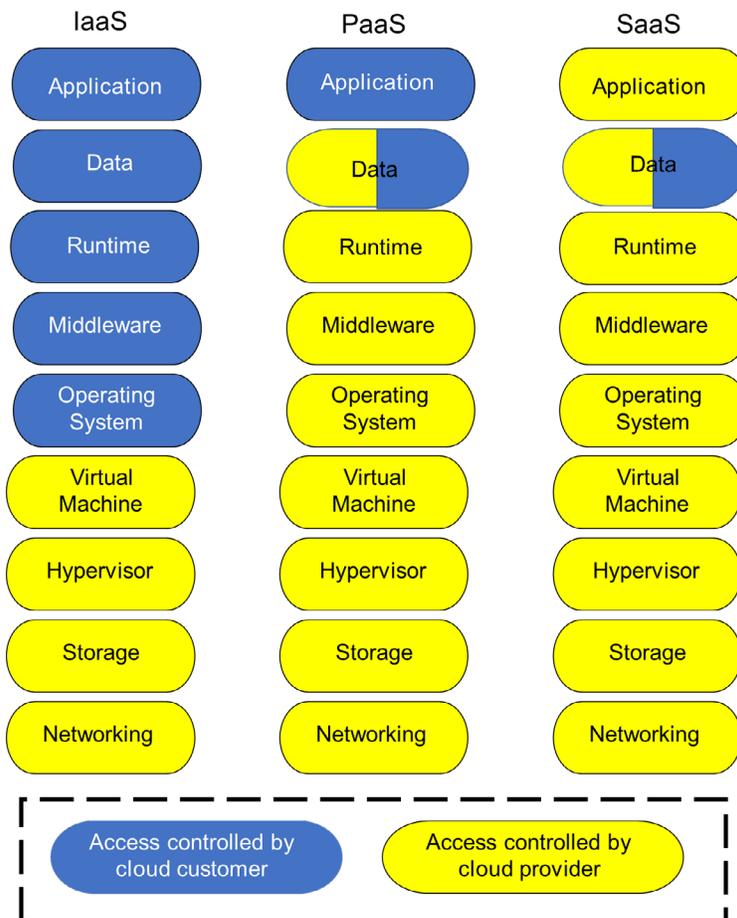


Figure 3: Accesses controlled by the cloud service provider and the consumer

The five essential characteristics that challenge AC system design are summarized as follows [2]:

1. *Broad network access*: Cloud services are available over the network and accessible through standard mechanisms that promote use by heterogeneous thick and thin client platforms (e.g., mobile phones, tablets, laptops, workstations). This raises security concerns with regard to network access. For example, denial of service (DoS) attacks can be launched against a cloud system, rendering its resources unavailable to legitimate users. Thus, AC for network access should be managed.
2. *Resource pooling*: The computing resources of a cloud system (e.g., storage, memory, processing, network bandwidth) are pooled to serve multiple consumers using a multi-tenant model (i.e., a single instance of the software and its supporting infrastructure serves multiple consumers) through different physical and virtual resources, each dynamically assigned and reassigned according to consumer demands. Information may be leaked if the resource allocated to a consumer can be accessed by another co-located consumer or if the allocated resource, such as memory, is not wiped before being reallocated to another consumer. There is also a sense of location independence in that the consumer generally has no control over or knowledge of the exact location of the provided resources. Location may be specified at a higher level of abstraction (e.g., country, state, data center) that brings

security concerns. Therefore, methods for implementing resource pooling while ensuring the isolation of shared resources should be considered in the AC design.

3. *Rapid elasticity*: Cloud services can be elastically provisioned and released—automatically, in some cases—to rapidly scale outward and inward commensurate with demands. To the consumer, services available for provisioning often appear to be unlimited and appropriated in any quantity at any time and are supported by adding new *virtual machines* (VMs) with specified computing resources. A challenge for AC design involves the capability to rapidly verify the security of new VMs and determine whether the newly added VMs are qualified to execute a specific task.
4. *Measured service*: Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, active end user accounts). Resource usage is monitored, controlled, and reported to provide transparency to both the provider and consumer of the utilized service. To maintain resource usage, cloud consumers should be authorized to review but not to modify their own metering data since this could lead to the falsification of payments required for cloud services. Thus, it is reasonable for AC to consider the protection of metering data.
5. *Data sharing*: Sharing information among different organizations is not a trivial task since a cloud system needs to meet the same security requirements of organizations to achieve that. To facilitate data sharing, concepts such as trust of federated identities and AC attributes need to be considered, and building that trust is paramount. In this document, it is assumed that trust and federated identities/attributes are already established, and further discussion on that topic will be considered in another document. Regardless of the service model, consumers are entitled to be responsible for the security of their cloud-based data and, implicitly, of who has access to it [5]. For this reason, data is never controlled by cloud service providers but rather always stays with the cloud consumers. (The exception to this is log data, but consideration should still be given to how privacy and security is affected by such data.) Although a cloud service provider might become the custodian of consumers' data, it should not have access to that data. If a consumer's data is not encrypted, then cloud administrators might be able to read it. In such a case, the consumer's data should be identified (by the provider's access privileges to the data) and red-flagged as accessible by the service provider, and the consumer should be informed immediately.

Guidance for AC system for each cloud service model, as described in Sections 3, 4, and 6 of this document, can be further extended to system requirements by referring to the AC control elements listed in NIST SP 800-53, Revision 4, *Security and Privacy Control for Federal Information Systems and Organizations* [6] based on the operation requirements of the cloud service. Appendix A maps the guidance to the AC control elements listed in the NIST SP 800-53, Revision 4.

3 Access Control Guidance for IaaS

IaaS is the cornerstone of all cloud services that offer computing and storage through a network such as the internet. Through virtualization technology, IaaS enables end users to dynamically allocate computing resources by instantiating new *virtual machines* (VMs) or releasing them based on their requirements. A VM is a software container that behaves like a physical machine with its own operating system (OS) and virtual resources (e.g., CPU, memory, hard disk, etc.). Leasing VMs is more cost-effective than purchasing new physical machines. The virtualization technology is composed of VMs and a *hypervisor*, as shown in Figure 1. VMs are managed by the hypervisor, which controls the flow of data and instructions between the VMs and the physical hardware. On the consumer side, system administrators are usually the major users of IaaS services since IaaS services are flexible to configure resources (e.g., network, data storage).

Cloud virtualization adds additional security management burdens by introducing security controls that arise from combining multiple VMs onto a single physical computer, which can have potential negative impacts if a security compromise occurs. Some cloud systems make it easy to share information among VMs by, for instance, allowing users to create multiple VMs on top of the same hypervisor if multiple VMs are available. However, this convenience can also become an attack vector since data leakage could occur among VMs. Additionally, virtualized environments are transient since they are created and vanish frequently, thereby making the creation and maintenance of necessary security boundaries more complex.

As shown in Figure 3, data in the middleware, data, applications, and OS layers is owned and controlled by the consumer. The IaaS system and the consumer need to ensure that access to the data is not granted to IaaS system administrators or any other IaaS consumers in these layers unless any of them are permitted. IaaS administrators are responsible for access control on the virtual machine, hypervisor, storage, and networking layers and should consider Sections 3.1 to 3.5 below.

3.1 Guidance for Network

The network is shared among IaaS consumers, and it is important to secure the network traffic and the cloud's environment from being exploited by unauthorized consumers. Thus, access control for network boundaries and allowlists for network communications are required and may be applied through, for example, dedicated virtual local area networks (VLANs) leveraging automated access control lists (ACLs). Using the Institute of Electrical and Electronics Engineers (IEEE) 802.1Q VLAN tagging for network traffic with a cloud data center will result in routing only traffic tagged with the server's unique VLAN identifier to or from that server [7].

3.2 Guidance for Hypervisor

A hypervisor plays an important role in the security of the entire virtualized architecture since it manages consumer loads and guest operating systems (OSs),² creates new guest OS images, and controls hardware resources. The security implications of actions like managing guest OS and hardware resources means that access to the hypervisor should be restricted to authorized cloud administrators only. Otherwise, a cloud end user could potentially obtain a VM from the cloud

² An OS that is secondary to the originally installed OS.

service provider and install a malicious guest OS that compromises the hypervisor by gaining unauthorized access to and altering the memory of other VMs [8]. Moreover, an attacker in a VM with lower access rights may be able to escalate their access privilege to a higher level by compromising the hardware resources allocation within the hypervisor [9]. Protecting the hypervisor from unauthorized access is therefore critical to the security of IaaS services.

3.3 Guidance for Virtual Machines

VMs that are created by different end users allow resources to be shared among multiple end users. In such cases, it must be ensured that no application from one VM can directly access other VMs since covert channels [10, 11] may leak information between VMs by accessing shared physical resources (e.g., memory). Similarly, although the ability to copy and paste information between VMs via the clipboard is a convenient feature, such a capability could be made available on other VMs running on the same hypervisor and thus introduce an attack vector (i.e., information can be leaked to other VMs through the clipboard). Organizations should have policies regarding the use of shared clipboards.

Isolation between VMs is necessary to keep VMs running independently of each other, and quotas on VM resource usage should be regulated so that a malicious VM can be prohibited from exhausting computation resources. If a malicious application consumes the majority of computation resources, legitimate applications may not be able to obtain sufficient resources to perform their operations. Moreover, end users might terminate the execution of their tasks before they are finished. The state and data of the current VM would then be saved as a guest OS image, and when the task is resumed, the VM might be migrated from a different hypervisor. In such scenarios, guest OS images must be protected from unauthorized access, tampering, or storage. Furthermore, VMs that are not active may also store sensitive data. Monitoring access to the sensitive data in inactive VMs should be considered.

3.4 Guidance for APIs

There are several popular open-source platforms for deploying an IaaS system [12, 13, 14]. These solution platforms enable APIs to manage access control of VMs, hypervisors, and networks (note that a consumer cannot control hypervisors and networks in a multi-tenant environment unless it is a private cloud). For example, [14] consists of control components, including API, communication, lifecycle, storage, volume, scheduler, network, *API server* for managing AC policies for hypervisors, and *network Controller* for constructing network bridges and firewall AC rules. The lack of monitoring AC within these APIs might result in unenforced or wrongly enforced AC policies by the hypervisors, VMs, and networks. Thus, a service for monitoring the AC APIs in cloud platforms should also be taken into consideration.

3.5 Recommendations for IaaS Access Control

As shown in previous sections, the security of an IaaS cloud system is heavily dependent on the virtualization (hypervisor). One of the most widely adopted solutions for protecting them is a *virtualization management system* [15], which lies between the underlying hardware and the hypervisor. The virtualization management system enforces AC on both hypervisors and VMs in different ways. Virtualization management systems enforce different levels of access on different

users. Some users are given read-only access to the administrative interface of a guest OS; some are allowed to control particular guest OSs; and some are given complete administrative control. There are existing solutions for providing AC for hypervisors and VMs. For example, the approach in [16] secures the hypervisor against control hijacking attacks by protecting its code from unauthorized access and offering isolation of VMs with the flexible security of mandatory access control (MAC). To enforce AC on interoperations, a service level agreement should be designed to include appropriate control to secure external interoperations. Other isolation mechanisms [17, 18] are helpful in ensuring the security of internal interoperations.

Guideline rules for IaaS AC policy that consider the main elements in AC (i.e., subject, object, and operation) are listed in Table 1. While each row indicates a possible AC rule, the AC policy designer should ultimately decide whether the decision in each rule is permitted or denied based on system requirements. For example, if an authorized IaaS end user requires the use of cloud services, a login operation in the hypervisor for the end user should be granted; otherwise, it should be denied.

Table 1: Potential policy rules expressed by Subject, Operation, Object for IaaS AC policy

Subjects	Operations	Objects	Environment Conditions
IaaS end user	Login, Read, Write, Create	Hypervisor	Time, Location, Security impact level etc.
IaaS end user	Read, Write, Create	VMs	Time, Location, Security impact level etc.
VM	Write	Hypervisor	Time, Location, Security impact level etc.
VM	Read, Write	Other VMs within the same host	Time, Location, Security impact level etc.
VM	Read, Write, Create	Guest OS images	Time, Location, Security impact level etc.
VM	Read, Write	Other VMs from different hosts but within the same IaaS provider	Time, Location, Security impact level etc.
VM	Read, Write	Other VMs from different IaaS providers	Time, Location, Security impact level etc.
Hypervisor	Read, Write, Create	Guest OS images	Time, Location, Security impact level etc.
Hypervisor	Read, Write	Hardware resources	Time, Location, Security impact level etc.
Hypervisor	Read, Write, Create	VMs	Time, Location, Security impact level etc.

4 Access Control Guidance for PaaS

PaaS is a platform that provides a framework for developers to create and deploy customized applications. As shown in Figure 3, security assurance considerations include some and all below the data level, and during the application development process lifecycle should be offered by the PaaS provider. The primary focus of AC in the PaaS model is to protect data during runtime, which is managed by middleware and OS. PaaS systems are primarily concerned with developing, deploying and operating customer applications. The security and privacy offered by the PaaS provider protect the applications and data from potential leaks through a covert channel introduced by unsecure shared memory. Therefore, enforcing AC over data during runtime in the PaaS is critical for the security of PaaS services.

The PaaS system administrator is responsible for the access control of runtime, middleware, OS, virtual machine, hypervisor, storage, and networking layers, as described by the guidance in Sections 4.1 to 4.3 below.

4.1 Guidance for Memory Data

The PaaS system permits users to deploy tasks in a provider-controlled middleware and host OS, which may be shared with other PaaS applications. As such, PaaS typically leverages OS-based techniques (e.g., Linux Containers and Docker for isolating applications) [19]. However, numerous existing memory-related attacks can compromise sensitive application-related data by hacking through the shared OS memory in PaaS [20]. Thus, AC for OS memory, such as AC of different processes on top of processor caches [21], should be considered.

4.2 Guidance for APIs

As the PaaS system allows cloud developers to build applications on top of the platform, APIs should control the scope of each user's application such that user data remains inaccessible between different applications. In addition, packaged APIs can be serviced as microservices in a PaaS cloud. A centralized architecture for provisioning and enforcement of access policies governing access to all microservices is required due to the sheer number of services needed for service composition to support real-world business transactions (e.g., consumer order processing and shipping). Since each of the microservices may be implemented in a different language, policy provisioning and computation of access decisions may require the use of an authorization server [22].

4.3 Recommendations for PaaS Access Control

An efficient method should be established for protecting memory data by flushing processor caches during context switches. However, in order to avoid significant performance degradation, only highly sensitive memory data should be flushed.

To handle access control for multiple replicas of data, a method to manage the central AC policy system should be introduced. Thus, once the data within a PaaS provider is duplicated across PaaS providers, any change in the policy should result in an appropriate update to the central AC policy

system. Moreover, the AC policy related to the replicated data in other PaaS providers should be synchronized accordingly based on an AC policy in the central system.

Guideline rules for PaaS AC policy are listed in Table 2 with respect to the three basic elements of AC (i.e., subject, object, and operation). Each row indicates a possible AC rule, but the AC designer should decide whether access should be granted or denied based on the system requirements. For example, if a user of an application needs to access memory data related to their application, permission to read memory data will be granted. However, access to that memory data will be denied to other users.

Table 2: Potential policy rules expressed by Subject, Operation, Object for PaaS AC policy

Subjects	Operations	Objects	Environment Conditions
Application user	Read	Memory data	Time, Location, Security impact level etc.
VM of a hosted application	Read, Write	Other applications' data within the same host	Time, Location, Security impact level etc.
Application developer	Create, Read, Write	Middleware data, memory data	Time, Location, Security impact level etc.
Cloud service provider	Replicate	Application-related data	Time, Location, Security impact level etc.

5 Access Control Guidance for SaaS

In SaaS, a cloud service provider delivers an application as a service to end users through a network such as the internet. Thus, there is no need for users to install and execute applications locally on their own computers. As shown in Figure 4, multiple applications and users can be supported simultaneously by the cloud system to share common resources, including applications and underlying databases.

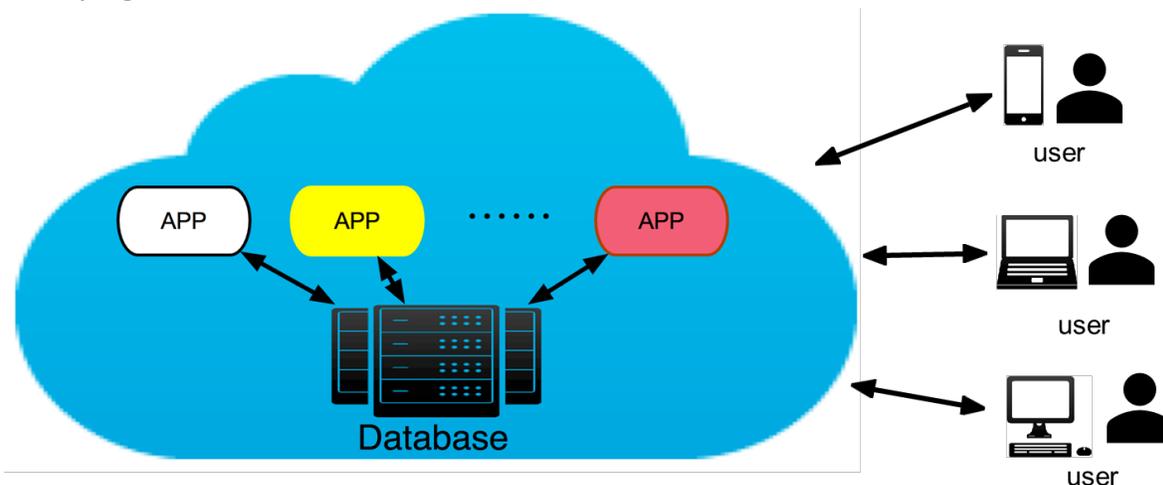


Figure 4: Multiple applications and users of an SaaS provider

If a developer deploys a third-party application, data in that application and other unrelated applications might be stored in the cloud system. End users have to rely on the security and privacy offered by the cloud service provider to protect their data from unauthorized access introduced by those unrelated applications. Note that data managed by the application layer is owned and controlled by the consumer. The SaaS system and consumer need to ensure that access to application data in these layers is not granted to the SaaS system administrator, consumers, or other users unless they are trusted. SaaS administrators are responsible for the access control of all operation layers except for the consumer's application data as shown in Figure 3 and should consider the guidance in Sections 3, 4, and 5.1 to 5.4.

5.1 Guidance for Data Owner's Control

A data provider is the creator or source of application data owned by consumer organizations. Application data is typically stored in the SaaS service provider's database. How a data provider manages access to its data is a challenge. Example questions to be addressed are related to data retention by the provider (e.g., where data is kept and for how long) and whether the provider has any permission to determine access rights to the data it hosts. If a data provider has the capability to determine access rights on data it holds, consideration should be given to ensure that an up-to-date AC policy is always enforced within the SaaS system.

5.2 Guidance for Confidentiality

In the application deployment model, the integrity of sensitive data residing within the data owner's domain must be protected. Protection mechanisms for application data include data

encryption schemes by which data can be encrypted through certain cryptographic primitives, and decryption keys will only be disclosed to authorized users [23]. For such enforcement, attribute-based access control (ABAC) [24] and attribute-based encryption (ABE) schemes can be used to control access to SaaS data [23, 25, 26, 27, 28] since these schemes can use the identity of users through attributes to manage, encrypt, and decrypt application data. However, considering the high volume of data in the SaaS model, the involved encryption and decryption significantly reduce performance. Hence, when encryption is used, consideration should be given to ensure the confidentiality of data while offering good performance.

5.3 Guidance for Privilege Management

In addition to AC enforcement, privilege management involves adding, removing, and changing the privileges of a subject. It is crucial to design a flexible or real-time mechanism for assigning and revoking privileges to maintain the usability of the SaaS service [29].

5.4 Guidance for Multiple Replicas of Data

To maintain high availability, the cloud service provider may replicate data at multiple locations, even across countries. Thus, it is important to make sure that all data replicas are protected under the same AC policy. In other words, the same AC policy for the replicated data object should be populated to all hosts that process the same data. The technology for policy synchronization upon changes must also be considered for inclusion.

5.5 Guidance for Multi-tenancy

The SaaS system introduces additional considerations with regard to the management of access to applications. An immediate necessity is to focus on users' access to applications. The access rights are granted to end users through AC policies based on predefined attributes or roles. This can be specified by attribute-based access control (ABAC) policy models [30, 31], role-based access control [32] (RBAC), and context-based access control [33] (CBAC).

The SaaS model is a typical, multi-tenancy platform that supports multiple end users simultaneously accessing an application with the data of different users' applications residing at the same location. Exploiting vulnerabilities in the application or injecting code into the SaaS system might expose data to other users [34]. Therefore, strategic planning should be given to implementing multi-tenancy while segregating data from different users' applications during the design of an AC system.

5.6 Guidance for Attribute and Role Management

In the SaaS system, attribute and role-based AC management employs policies and predefined roles to manage access rights to applications and underlying databases. The primary challenge of deploying attribute or role-based AC management is reaching an agreement on what types of attributes or roles should be used and what should be considered when designing the AC systems [35]. If the set of considered attributes or roles is too small, flexibility will be reduced. However, if the number of attributes or roles is too large, the complexity of policies will increase.

5.7 Guidance for Policies

SaaS applications provide application-specific access control configurations for different user applications, and in this case, user policies for each application are enforced by the SaaS provider. This configuration does not support collaboration between the SaaS provider and the consumer's access control infrastructure. For example, while large organizations often employ on-premises access control systems for managing their users centrally and efficiently, SaaS applications typically provide organizations with an AC configuration interface for managing AC policies, which forces the AC policies to be stored and evaluated on the SaaS provider's side. This approach might result in disclosing sensitive data required for evaluating the AC policies to the SaaS provider. Therefore, methods for enforcing authorization in the SaaS provider while not disclosing sensitive access control data to the SaaS provider should be considered. Federated authorization [36] is an efficient technique that utilizes a middleware layer to transfer the management of access control policies from the SaaS provider to the consumer side and enforce policies on the SaaS applications without disclosing sensitive data required for evaluating the policies.

5.8 Guidance for APIs

An API in the SaaS model serves as an interface between the cloud server and its users. The API should be designed to protect against both accidental and malicious attempts to circumvent any AC policy. Applications for organizations and third parties often build upon the APIs, which introduce the AC complexity of the new layered API. For example, if the APIs do not require memory access for their tasks, then the AC policy for the APIs should enforce the non-memory access. Additionally, AC policies should be specified to manage the authorization process for web APIs. For example, when APIs connect through SOAP and REST protocols, the AC should control whether to allow end users to interface between Microsoft or non-Microsoft tools and technologies. For authorized API connections through Simple Object Access Protocol (SOAP) and Representational State Transfer (REST) protocols, the AC should grant all related access requested by the protocols. For unauthorized API connections through these protocols, no access or partial access should be granted by the AC.

5.9 Recommendations for SaaS Access Control

With regard to multi-tenancy, authorization may be enforced using a *centralized*, *decentralized*, or *hybrid* authorization system. In a centralized authorization system, the SaaS provider manages a central authorization database for every end user and their accounts [37]. In a decentralized or hybrid authorization system, individual tenants are responsible for all or part of the authorization process. Note that different tenants may require different systems. Considering the attributes or roles of tenants is crucial when selecting the most suitable system. There are many ways to specify attributes or roles, such as in ABAC and RBAC models [31,32]. Attributes or roles must be well-designed and take into account hierarchy relationships when implementing AC policies for different tenants.

Authorization federation [36] is an efficient way to enforce AC policies in the SaaS provider. A generic middleware architecture that incorporates access control requirements from consumers and handles local and remote attributes or roles can be used to extend and shift AC policy management from the SaaS provider to the consumer side. This approach centralizes consumer AC policy

management and lowers the required trust in the SaaS provider. In addition, the AC for VM-supporting federation operations should also be specified (e.g., an end user may create a VM to run different applications). Within the VM of the same host, one application may need to access the application code of other applications to fulfill its task. Unlike the PaaS architecture, where consumers can fully manage the design, testing, and development of the software, SaaS consumers have limited control of the applications hosted in the cloud server.

To achieve the application data owner's control, a security class agreement (SCA) [28] may be of use. SCA is mutually agreed upon by both the data provider of PaaS subscribers and the PaaS service provider and is used for defining the security class of data providers. Multiple replicas of the same data share the same security level as its data provider. This means that given data from a particular data provider, the security class for multiple replicas of the data should be identical. As a result, the host within the PaaS service that is qualified for executing the access request can be determined by referring to the SCA. The data provider can manage access to its data by specifying security classes for the SCA to keep the data provider and the cloud host synchronized in determining the access right of data. For example, in a Bell-LaPadula model [38], assuming a patient's report is written by a doctor with confidential clearance, the report can only be read by a host with the same or higher security clearance. Additionally, when multiple data sources that are not intended to be accessed in the same cloud system are accessed, the privacy of data should not be leaked due to different security classes of these data sources and their data in the SCA. However, due to the high computation complexity of encryption and decryption, cryptographic schemes should be carefully designed to maintain the performance of cloud systems while protecting data confidentiality.

A privilege management infrastructure (PMI) [39] can be employed to dynamically manage assigning and revoking privileges through the use of attributes or role specification certificates in the PaaS model. PMI specifies the privileges for different users and links the privileges with different attribute or role specification certificates, which contain different attribute or role assignments to enforce privilege management.

To handle access control of multiple replicas of data, a method to manage the central AC policy system should be introduced. Thus, once the data within an SaaS provider is duplicated across SaaS providers, any change in the policy should result in an appropriate update to the central AC policy system. Moreover, the AC policy related to the replicated data in other SaaS providers should be synchronized accordingly based on an AC policy in the central system.

Guideline rules for SaaS AC policy are listed in Table 3. The AC designer should decide whether access in each rule is permitted or denied based on the system requirements. For example, during federation operation, VM read/write to other application code within the same host is permitted; otherwise, it is denied.

Table 3: Potential policy rules expressed by Subject, Operation, Object for SaaS AC policy

Subjects	Operations	Objects	Environment Conditions
Application user	Read, Write	Application-related data	Time, Location, Security impact level etc.
Application user	Read	Memory	Time, Location, Security impact level etc.
Application user	Execute	Application	Time, Location, Security impact level etc.
Application user	Read, Write	Application data	Time, Location, Security impact level etc.
Application user	Execute	Application code	Time, Location, Security impact level etc.
VM of a hosted application	Execute	Other application code within the same host	Time, Location, Security impact level etc.

6 Access Control Guidance for Inter- and Intra- Operation

In general, collaboration (i.e., two or more systems that work together as a combined system) in the context of the cloud may lead to a seamless exchange of data and services among various cloud infrastructures. There are two types of collaborations: *inter-operation* and *intra-operation*. Inter-operation refers to the capability of using multiple cloud infrastructures. For example, as shown in Figure 5, a consumer may purchase IaaS services from two different cloud service providers, *Cloud A* and *Cloud B*, and the collaboration between them should be allowed due to data processing requirements.

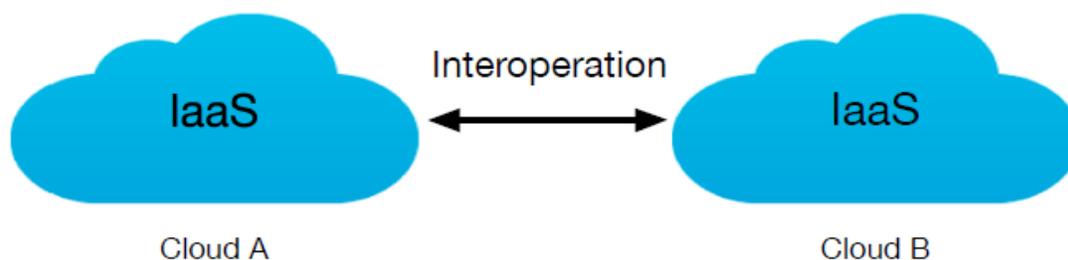


Figure 5: The external collaboration (inter-operation) between different Clouds

Intra-Operation

With regard to intra-operation, two scenarios on intra-operation can be presented as derived from Figure 6. First, a consumer may own multiple VMs in a single cloud host (e.g., *VM A* and *VM B*), and communication among those VMs may be required. Second, a consumer may rent multiple hosts within the same IaaS service, and collaboration among VMs from these different hosts may be required (e.g., an inter-operation between *VM B* and *VM C*).

For intra-operation, the AC policy should enable the operations of VMs for the same consumer to access each as needed during the collaboration period and disable access when the collaboration period ends. There are two primary cases in intra-operation: inter-host case (i.e., VMs from different cloud hosts are operating collaboratively) and intra-host case (i.e., VMs are from the same cloud host and must exchange data and services). Additionally, for some applications, VMs might be distributed in multiple host computers, so the AC policy should cover both intra-host and inter-host cases.

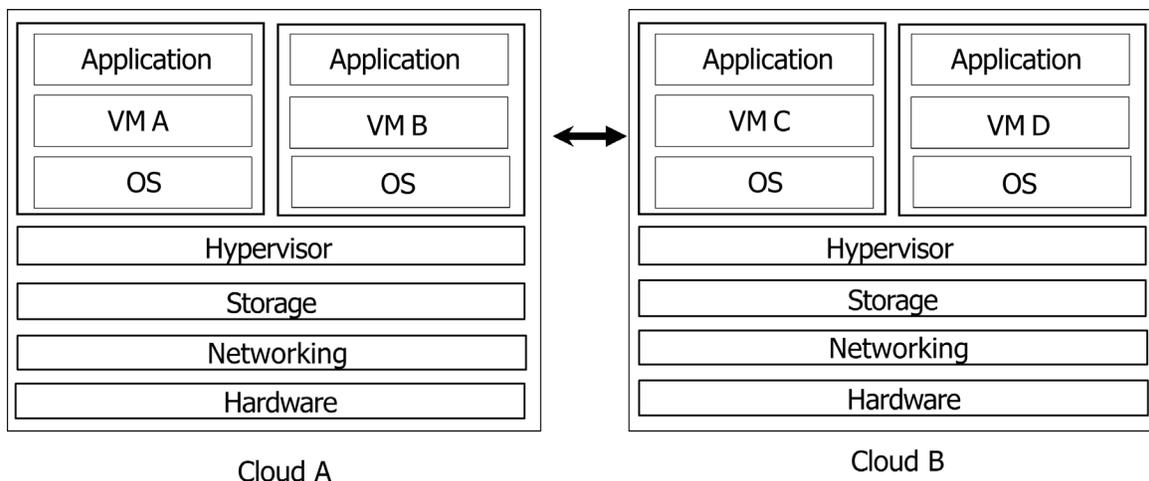


Figure 6: The internal collaboration (intra-operation) within the same cloud

Inter-Operation

There is the possibility that inconsistent management of access elements leads to incorrect access control policy integration for inter-operation. For instance, different cloud service providers using different sets of subject attributes for AC may cause potential conflicts or leak access permissions [40]. Attributes with the same name may result in different privileges when switching providers. Enforcing AC among different cloud service providers without incurring conflicts or blocks of privilege for individual users/VMs is a challenge. This would require examining how to achieve secure inter-operation among the cloud service providers [1], such as in cross hybrid environments. Some cloud AC systems adopt centralized mechanisms to create global AC policies that manage policy integration among different cloud service providers [41]. However, the cloud inter-operation is transient and, thus, inefficient to manage global AC policies as frequent updates for individual cloud AC policies.

7 Conclusions

This document presents an initial step toward understanding access control (AC) challenges in cloud systems by analyzing the AC considerations in all three cloud service delivery models— IaaS, PaaS, and SaaS. Essential characteristics that would affect the cloud’s AC design are also summarized, such as broad network access, resource pooling, rapid elasticity, measured service, and data sharing. Various guidance for AC design of IaaS, PaaS, and SaaS are proposed according to their different characteristics. Recommendations for AC design in different cloud systems are also included to facilitate future implementations. Additionally, potential policy rules are summarized for each cloud system. However, many issues remain open, such as AC management across different devices and platforms, as well as new challenges that have yet to emerge with the wide adoption of the cloud.

References

- [1] Gouglidis A, Mavridis I, Hu VC (2014) Security policy verification for multi-domains in Cloud systems. *International Journal of Information Security* 13(2):97-111. <https://doi.org/10.1007/s10207-013-0205-x>
- [2] 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>
- [3] Liu F, Tong J, Mao J, Bohn R, Messina J, Badger ML, 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>
- [4] Badger ML, Grance T, Patt-Corner R, Voas JM (2012) Cloud Computing Synopsis and Recommendations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-146. <https://doi.org/10.6028/NIST.SP.800-146>.
- [5] Federal Information Security Modernization Act of 2014, Pub. L. 113-283, 128 Stat. 3073. <https://www.govinfo.gov/app/details/PLAW-113publ283>
- [6] Joint Task Force Transformation Initiative (2013) Security and Privacy Controls for Federal Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-53, Rev. 4, Includes updates as of January 22, 2015. <https://doi.org/10.6028/NIST.SP.800-53r4>
- [7] Bartock MJ, Souppaya MP, Scarfone KA, Carroll D, Masten R, Scinta G, Massis P, Prafullchandra H, Malnar J, Singh H, Yeluri R, Shea T, Dalton M, Dukes A, Phoenix C Swarts B (2018) Trust Cloud: Security Practice Guide for VMware Hybrid Cloud Infrastructure as a Service (IaaS) Environments. (National Institute of Standards and Technology, Gaithersburg, MD), Preliminary Draft NIST Special Publication (SP) 1800-19B. Available at <https://www.nccoe.nist.gov/projects/building-blocks/trusted-cloud>
- [8] Szefer J, Lee RB (2011) A case for hardware protection of guest VMs from compromised hypervisors in cloud computing. *2011 31st International Conference on Distributed Computing Systems Workshops (ICDCSW)* (IEEE, Minneapolis, MN), pp 248–252. <https://doi.org/10.1109/ICDCSW.2011.51>
- [9] Krutz RL, Vines RD (2010) *Cloud security: A comprehensive guide to secure cloud computing* (Wiley Publishing, Indianapolis, IN).
- [10] Wu J, Ding L, Wu Y, Min-Allah N, Khan SU, Wang Y (2014) C2detector: a covert channel detection framework in cloud computing. *Security and Communication Networks* 7(3):544–557. <https://doi.org/10.1002/sec.754>

- [11] Rushby J (1992) Noninterference, transitivity, and channel-control security policies. (SRI International, Menlo Park, CA), Technical Report CSL-92-02. Available at <http://www.csl.sri.com/papers/csl-92-2/>
- [12] Change ATC, Foster JL, Hall DK (1987) Nimbus-7 SMMR derived global snow cover parameters. *Annals of Glaciology* 9:39-44. <https://doi.org/10.3189/S0260305500200736>
- [13] Nurmi D, Wolski R, Grzegorzczak C, Obertelli G, Soman S, Youseff L, Zagorodnov D (2009) The Eucalyptus open-source cloud-computing system. *9th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID'09)* (IEEE, Shanghai, China), pp 124-131. <https://doi.org/10.1109/CCGRID.2009.93>
- [14] Sefraoui O, Aissaoui M, Eleuldj M (2012) OpenStack: toward an open-source solution for cloud computing. *International Journal of Computer Applications* 55(3):38-42. <https://doi.org/10.5120/8738-2991>
- [15] Scarfone KA, Souppaya MP, 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>
- [16] Wang Z, Jiang X (2010) Hypersafe: A lightweight approach to provide lifetime hypervisor control-flow integrity. *2010 IEEE Symposium on Security and Privacy (SP)* (IEEE, Berkeley/Oakland, CA), pp 380–395. <https://doi.org/10.1109/SP.2010.30>
- [17] Berger S, Cáceres R, Pendarakis D, Sailer R, Valdez E, Perez R, Schildhauer W, Srinivasan D (2008) TVDc: managing security in the trusted virtual datacenter. *ACM SIGOPS Operating Systems Review* 42(1):40–47. <https://doi.org/10.1145/1341312.1341321>
- [18] Sailer R, Valdez E, Jaeger T, Perez R, Doorn LV, Griffin JL, Berger S (2005) sHype: Secure hypervisor approach to trusted virtualized systems. (IBM Research Division, Yorktown Heights, NY) IBM Research Report RC23511. Available at [https://domino.research.ibm.com/library/cyberdig.nsf/papers/265C8E3A6F95CA8D85256FA1005CBF0F/\\$File/rc23511.pdf](https://domino.research.ibm.com/library/cyberdig.nsf/papers/265C8E3A6F95CA8D85256FA1005CBF0F/$File/rc23511.pdf)
- [19] Zhang Y, Juels A, Reiter MK, Ristenpart T (2014) Cross-tenant Side-channel Attacks in PaaS Clouds. *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (ACM, Scottsdale, AZ), pp 990–1003. <https://doi.org/10.1145/2660267.2660356>
- [20] Osvik DA, Shamir A, Tromer E (2006) Cache attacks and countermeasures: the case of AES. Pointcheval D. (eds) *Topics in Cryptology – CT-RSA 2006*. CT-RSA 2006. Lecture Notes in Computer Science 3860 (Springer, Berlin), pp 1–20. https://doi.org/10.1007/11605805_1
- [21] Tromer E, Osvik DA, Shamir A (2010) Efficient cache attacks on AES, and countermeasures. *Journal of Cryptology* 23(1):37–71. <https://doi.org/10.1007/s00145-009-9049-y>

- [22] Chandramouli R (2019) Security Strategies for Microservices-based Application Systems. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-204. <https://doi.org/10.6028/NIST.SP.800-204>
- [23] Yu S, Wang C, Ren K, Lou W (2010) Achieving secure, scalable, and fine-grained data access control in cloud computing. *INFOCOM, 2010 Proceedings* (IEEE, San Diego, CA), pp 1-9. <https://doi.org/10.1109/INFCOM.2010.5462174>
- [24] Hu VC, Ferraiolo DF, Kuhn DR, Schnitzer A, Sandlin K, Miller R, Scarfone KA (2014) Guide to Attribute Based Access Control (ABAC) Definition and Considerations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-162, Includes updates as of August 02, 2019. <https://doi.org/10.6028/NIST.SP.800-162>
- [25] Sahai A, Waters B (2005) Fuzzy identity-based encryption. *Advances in Cryptology – EUROCRYPT 2005*. Lecture Notes in Computer Science 3494 (Springer, Berlin), pp 457–473. https://doi.org/10.1007/11426639_27
- [26] Nali D, Adams CM, Miri A (2005) Using threshold attribute-based encryption for practical biometric-based access control. *International Journal of Network Security* 1(3):173–182. Available at http://ijns.jalaxy.com.tw/download_paper.jsp?PaperID=IJNS-2005-06-30-2&PaperName=ijns-v1-n3/ijns-2005-v1-n3-p173-182.pdf
- [27] Zhu Y, Hu H, Ahn G-J, Huang D, Wang S (2012) Towards temporal access control in cloud computing. *INFOCOM, 2012 Proceedings* (IEEE, Orlando, FL), pp 2576–2580. <https://doi.org/10.1109/INFCOM.2012.6195656>
- [28] Hu VC, Grance T, Ferraiolo DF, Kuhn DR (2014) An access control scheme for big data processing. *2014 International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)* (IEEE, Miami, FL), pp 1–7. <https://doi.org/10.4108/icst.collaboratecom.2014.257649>
- [29] Hu VC, Scarfone KA (2012) Guidelines for Access Control System Evaluation Metrics. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) 7874. <https://doi.org/10.6028/NIST.IR.7874>
- [30] Vipul G, Pandey O, Sahai A, Waters B (2006) Attribute-based encryption for fine-grained access control of encrypted data. *Proceedings of the 13th ACM Conference on Computer and Communications Security (CCS '06)* (ACM, Alexandria, VA), pp 89-98. <https://doi.org/10.1145/1180405.1180418>
- [31] Hu VC, Kuhn DR, Ferraiolo DF, Voas J (2015) Attribute-based access control. *Computer* 48(2):85-88. <http://doi.org/10.1109/MC.2015.33>
- [32] Sandhu RS, Coyne EJ, Feinstein HL, Youman CE (1996) Role-based access control models. *Computer* 29(2):38-47. <https://doi.org/10.1109/2.485845>

- [33] Rubart J (2005) Context-based access control. *Proceedings of the 2005 Symposia on Metainformatics (MIS '05)*. (ACM, New York, NY), pp 13-18. <https://doi.org/10.1145/1234324.1234337>
- [34] Subashini S, Kavitha V (2011) A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications* 34(1), pp 1–11. <https://doi.org/10.1016/j.jnca.2010.07.006>
- [35] Jin X, Krishnan R, Sandhu R (2012) A unified attribute-based access control model covering DAC, MAC, and RBAC. *Data and Applications Security and Privacy XXVI, DBSec 2012*. Lecture Notes in Computer Science 7371 (Springer, Berlin), pp 41-55. https://doi.org/10.1007/978-3-642-31540-4_4
- [36] Decat M, Lagaisse B, Van Landuyt D, Crispo B, Joosen W (2013) Federated authorization for software-as-a-service applications. *On the Move to Meaningful Internet Systems: OTM 2013 Conferences*. Lecture Notes in Computer Science 8185 (Springer, Berlin), pp 342–359. https://doi.org/10.1007/978-3-642-41030-7_25
- [37] Dimitrios Z, Lekkas D (2012) Addressing cloud computing security issues. *Future Generation Computer Systems* 28(3):583-592. <https://doi.org/10.1016/j.future.2010.12.006>
- [38] McLean J (1985) A comment on the ‘basic security theorem’ of Bell and LaPadula. *Information Processing Letters* 20(2):67-70. [https://doi.org/10.1016/0020-0190\(85\)90065-1](https://doi.org/10.1016/0020-0190(85)90065-1)
- [39] Blobel B, Nordberg R, Davis JM, Pharow P (2006) Modelling privilege management and access control. *International Journal of Medical Informatics* 75(8), pp 597–623. <https://doi.org/10.1016/j.ijmedinf.2005.08.010>
- [40] Bertino E, Federica P, Rodolfo F, Shang N (2009) Privacy-preserving digital identity management for cloud computing. *IEEE Data Engineering Bulletin* 32(1):21-27. Available at <http://sites.computer.org/debull/A09mar/bertino.pdf>
- [41] Catteddu D (2010) Cloud Computing: Benefits, risks and recommendations for information security. *Web Application Security*. Communications in Computer and Information Science 72 (Springer, Berlin), pp 17-17. https://doi.org/10.1007/978-3-642-16120-9_9
- [42] Simorjay F, Tierling E (2019) Shared Responsibility for Cloud Computing. (Microsoft, Redmond, WA), Version 2.0. Available at <https://gallery.technet.microsoft.com/Shared-Responsibilities-81d0ff91/file/225366/1/Shared%20Responsibility%20for%20Cloud%20Computing-2019-10-25.pdf>

Appendix A—Guidance and SP 800-53 Revision 4 Access Control (AC) Family Mapping

The following table maps the cloud access control guidance to the AC controls listed in NIST SP 800-53, Revision 4, *Security and Privacy Controls for Federal Information Systems and Organizations*.

Table 4 Mapping the cloud access control guidance to the AC controls listed in NIST SP 800-53, Revision 4

Guidance	AC Control in 800-53
3.1 Guidance for Network	AC-1, AC-3, AC-4, AC-5, AC-10, AC-17, AC-21, AC-22
3.2 Guidance for Hypervisor	AC-1, AC-3, AC-5, AC-17, AC-21
3.3 Guidance for Virtual Machine	AC-1, AC-3, AC-4, AC-5, AC-11
3.4 Guidance for API	AC-1, AC-3, AC-4, AC-5, AC-11, AC-17, AC-21, AC-22
4.1 Guidance for Memory Data	AC-1, AC-3, AC-4, AC-5, AC-10, AC-11, AC-21
4.2 Guidance for APIs	AC-1, AC-3, AC-4, AC-5, AC-10, AC-11, AC-21
5.1 Guidance for Data Owner's Control	AC-1, AC-3, AC-5
5.2 Guidance for Confidentiality	AC-3, AC-6, AC-21
5.3 Guidance for Privilege Management	AC-2, AC-11, AC-14, AC-22
5.4 Guidance for Multiple Replicas of Data	AC-1, AC-3, AC-4, AC-5, AC-17, AC-21
5.5 Guidance for Multi-tenancy	AC-1, AC-2, AC-3, AC-4, AC-5, AC-10, AC-11, AC-21
5.6 Guidance for Attribute and Role Management	AC-6, AC-1, AC-3
5.7 Guidance for Policies	AC-1, AC-3
5.8 Guidance for APIs	AC-1, AC-2, AC-3, AC-4, AC-5, AC-6, AC-11, AC-14, AC-17, AC-21

AC-1: Access Control Policy and Procedures

AC-2: Account Management

AC-3: Access Enforcement

AC-4: Information Flow Enforcement

AC-5: Separation of Duties

AC-6: Least Privilege

AC-10: Concurrent Session Control

AC-11: Session Lock

AC-14: Permitted Actions without Identification or Authentication

AC-17: Remote Access

AC-21: Collaboration and Information Sharing

AC-22: Publicly Accessible Content