Table of Contents

List of Figures

List of Figures

Chapter 3 Cloud Computing Security Essentials and Architecture

3.1 The 3F Inflection Point in the History of the Internet and Information Systems

The evolution of the Internet can be divided into three generations: in the 70s, the first generation was marked by expensive mainframe computers accessed from terminals; the second generation was born in the late 80s and early 90s, and was identified by the explosion of personal computers with Graphical User Interfaces (GUIs); the first decade of the 21st century brought the third generation, defined by mobile computing, the "Internet of Things" and cloud computing.

In 1997, Professor Ramnath Chellappa of Emory University defined cloud computing for the first time, calling it an important new "*computing paradigm where the boundaries of computing will be determined by economic rationale rather than technical limits alone*."[1] Even though the international IT literature and media have come forward since then with many definitions, models, and architectures for cloud computing, autonomic and utility computing were the foundations of what the community commonly referred to as "cloud computing." In the early 2000s, companies started rapidly adopting this concept upon the realization that cloud computing could benefit both the Providers as well as the Consumers of services. Businesses started delivering computing functionality via the Internet, enterprise-level applications, Web-based retail services, document-sharing capabilities, and fully hosted IT platforms, to mention only a few cloud computing use cases of the 2000s. The latest widespread adoption of virtualization and service-oriented architecture (SOA) has promulgated cloud computing as a fundamental and increasingly important part of any delivery and critical-mission strategy. It enables existing and new products and services to be offered and consumed more efficiently, conveniently, and securely. Not surprisingly, cloud computing became one of the hottest trends in IT, with a unique and complementary set of properties, such as elasticity, resiliency, rapid provisioning, and multi-tenancy.

Information systems are now at a triple or 3-factor inflection point in the IT's evolution (Figure 1). Virtualization of computing infrastructure set the foundation for the **technological** inflection point, providing *ubiquitous*[2] cloud computing that nurtured the evolution of *pervasive*[3] mobility and rapid expansion of the Internet of Things (IoT) or Network of Things (NoT). Cloud computing, mobility, and IoT/NoT are the steering components that induced the business **operations** inflection point, transforming the world from *connected* to *hyper-connected*. Due to its resilience and expandable capacity offered at

[1] "Cloud Computing for Teaching and Learning: Strategies for Design and Implementation", Lee Chao, University of Huston-Victoria, USA

[2] In 1991, Mark Weiser and his colleagues at the Palo Alto Research Centre introduced the terms 'ubiquitous' and 'pervasive' computing, used initially interchangeable to describe how computing was going to change from desktop, personal computing to a more distributed, mobile, and embedded form. Despite being used interchangeably, they do refer to different forms of computing. *Ubiquitous* means "the state of being everywhere," while *pervasive* means to "pass through, to be diffused throughout" (these definitions are taken from the Concise English Dictionary, 1984). In the computing world, ***ubiquitous computing*** describes the underlying framework, the embedded systems, networks, and displays that are invisible and everywhere, allowing us to 'plug-and-play' mobile devices and tools.

[3] ***Pervasive computing***, on the other hand, refers to the distributed set of tools and devices within our environment, through which we access information anytime, anywhere.

reduced cost, cloud computing resources became the target and the source of malicious activities, triggering an evolution among attackers and inducing an inflection in the sophistication and strength of **attacks**, resulting in the exponential increase of cybercrimes.

Figure 1: Information Systems' 3-Factors Inflection Point

3.2 Cloud Computing Definition

The United States' National Institute of Standards and Technology (NIST) was the first standards organization to define cloud computing and identify its main characteristics, deployment, and service models. According to the definition published in NIST Special Publication (SP) 800-145,[4] "*cloud computing is 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*." Enterprises can use these resources to develop, host and run services and applications on demand in a flexible manner anytime, anywhere and on any device. This definition is widely accepted as providing a clear understanding of cloud computing technologies and cloud services and has been submitted as the U.S. contribution for International standardization.

The NIST definition also provides a unifying view of five essential characteristics of cloud services: *on-demand self-service, broad network access, resource pooling, rapid elasticity,* and *measured service*. Furthermore, NIST identifies a simple and unambiguous taxonomy of three "service models" available to cloud Consumers: Infrastructure-as-a-Service (IaaS), Platform-

[4] NIST SP 800-145, *The NIST Definition of Cloud Computing*, September 2011. Available at: http://csrc.nist.gov/publications/PubsSPs.html#800-145.

as-a Service (PaaS), Software-as-a-Service (SaaS); and four "cloud deployment modes": Public, Private, Community and Hybrid. When combined, a service model and deployment model categorize ways to deliver cloud services. NIST SP 800-145 defines the three service models as follows:

1. *Infrastructure as a Service (IaaS) - The capability provided to the Consumer is to provision processing, storage, networks, and other fundamental computing resources where the Consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The Consumer does not manage or control the underlying cloud infrastructure, but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).*
2. *Platform as a Service (PaaS) - The capability provided to the Consumer is to deploy Consumer-created or acquired applications onto the cloud infrastructure that are created using programming languages and tools supported by the Provider. The Consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control over the deployed applications and possibly the application-hosting environment configurations.*
3. *Software as a Service (SaaS) - The capability provided to the Consumer is to use the Provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface, such as a web browser (e.g., web-based email). The Consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application-configuration settings.*

ISO/IEC JTC1 SC38 WG3 and ITU-T also developed a cloud computing taxonomy that is derived from NIST SP 800-145: International Standard ISO/IEC 17788 | Recommendation ITU-T Y.3500 "*Information technology - Cloud computing - Overview and vocabulary.*" [5]

The main concepts of cloud computing and many of the terms are largely interchangeable between the NIST and ISO/IEC standards. However, since NIST's cloud computing definition has been available longer and constitute also the core concept defined by ISO/IEC standard, this book leverages the NIST definition.

Each of the three cloud service models allows the following capabilities:

- **IaaS** allows cloud Consumers to run any operating systems and applications of their choice on the hardware and resource abstraction layers (hypervisors) furnished by the cloud Provider. A Consumer's operating systems and applications can be migrated to the cloud Provider's hardware, potentially replacing a company's data center infrastructure.
- **PaaS** allows Consumers to create their own cloud applications. Basically, the cloud Provider renders a virtualized environment and a set of tools to allow the creation of new web applications. The Cloud Provider also furnishes the hardware, operating

[5] Publicly available at: http://www.itu.int/rec/T-REC-Y.3500/en.

systems, and commonly used system software and applications, such as Database Management System (DBMS), Web Server, etc.

- **SaaS** allows cloud Consumers to run online applications. Off-the-shelf applications are accessed over the Internet. The cloud Provider owns the applications, and the Consumers are authorized to use them in accordance with a Service Agreement signed between parties.

In summary, cloud computing provides a convenient, on-demand way to access a shared pool of configurable resources (e.g., networks, servers, storage, applications, and services), enabling users to develop, host and run services and applications on demand in a flexible manner anytime, anywhere on any device.

3.3 Cloud Computing Reference Architecture

NIST was also the first to define a technology- and implementation-agnostic *Cloud Computing Reference Architecture* (NIST SP 500-292) that identifies the main cloud Actors, their roles, and the main architectural components necessary for managing and providing cloud services (e.g., service deployment, service orchestration, service management, service aggregation, etc.).

Derived from NIST SP 500-292 [6], ISO/IEC JTC1 SC38 WG3 and ITU-T also developed a reference architecture standard: International Standard ISO/IEC 17789 | Recommendation ITU-T Y.3502 *"Information technology - Cloud computing - Reference Architecture"* [7] that describes cloud computing Actors, focusing on *cloud Provider* and *cloud Customer*, while grouping the other cloud Actors in a separate *cloud Partners* category.

Cloud reference architectures and a cloud taxonomy are foundational documents that help cloud computing stakeholders communicate concepts, architecture, or operational and security requirements, to enumerate just a few of their benefits.

The technology-agnostic cloud computing Reference Architecture (RA) introduced by NIST in NIST SP 500-292 is a logical extension of NIST's cloud computing definition. As highlighted earlier, the cloud RA is a generic, high-level conceptual model that facilitates the understanding of cloud computing's operational intricacies. The RA does not represent the system architecture of a specific cloud computing system; instead, it is a tool for describing, discussing, and developing a system-specific architecture using a common framework of reference. This architecture is not tied to any specific vendor products, services, or reference implementations, nor does it provide prescriptive solutions. The RA defines a set of cloud Actors, and their activities, and functions that can be used for orchestrating a cloud

[6] NIST SP 500-292, *NIST Cloud Computing Reference Architecture*, September 2011. Available at: http://www.nist.gov/customcf/get_pdf.cfm?pub_id=909505.
[7] Publicly available at: http://www.itu.int/rec/T-REC-Y.3502/en.

Ecosystem.[8] The cloud computing RA relates to a companion cloud computing taxonomy and contains a set of views and descriptions that are the basis for discussing the characteristics, uses, and standards for cloud computing. The Actor-based model is intended to serve stakeholders by representing the overall view of roles and responsibilities in order to assess and manage the risk by implementing security and privacy controls.

As shown in **Figure 2,** the RA identifies the five major cloud Actors; Consumer, Provider, Broker, Carrier, and Auditor.

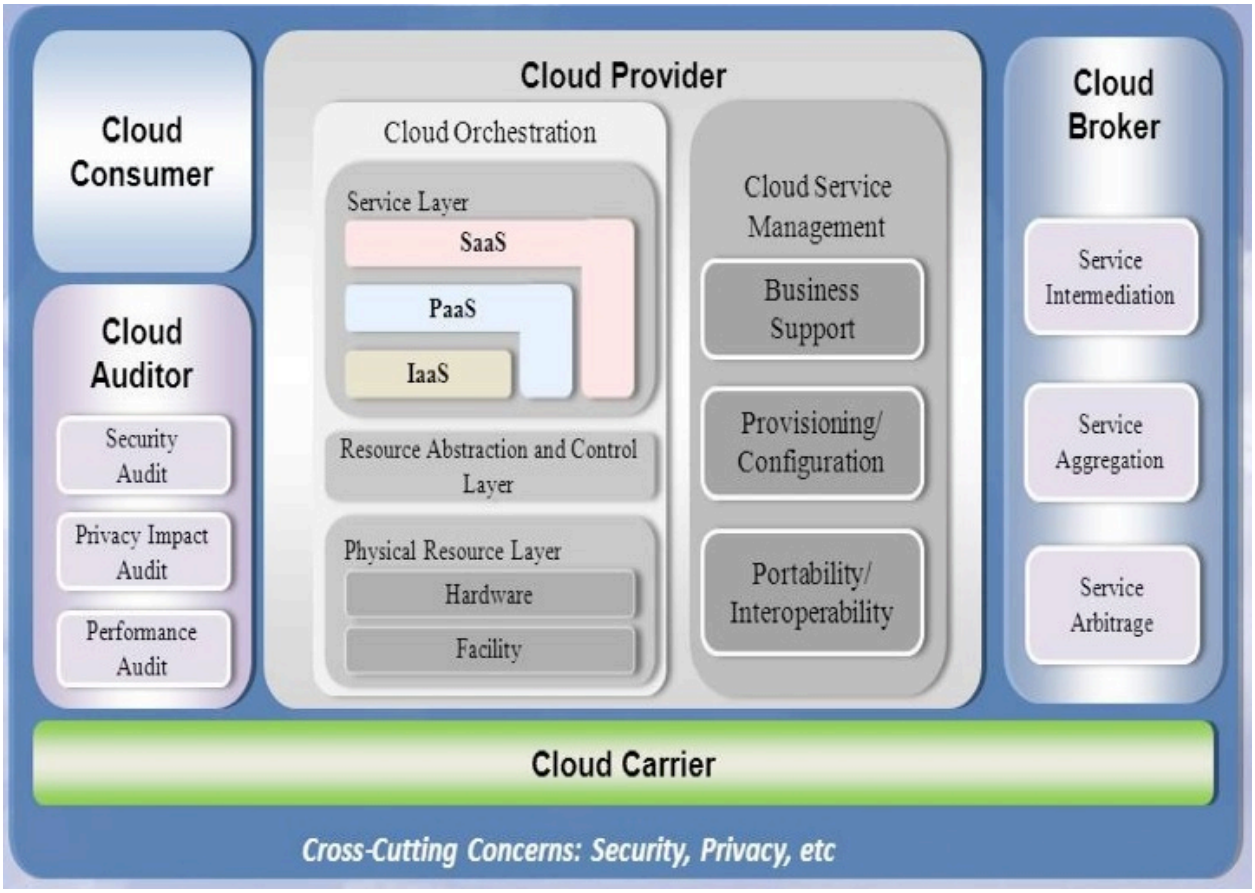

Figure 2: NIST Cloud Computing Security Reference Architecture Approach
(Source: NIST, SP 500-292)

Each cloud Actor defined by the NIST RA is an entity (a person or an organization) that participates in a transaction or process and/or performs tasks in cloud computing. The definitions of the cloud Actors introduced by NIST in SP 500-292 are reproduced below in **Table 1**.

Table 1: Cloud Actor Definitions (Source: NIST, SP 500-292)

[8] "Cloud Ecosystem" is a term used to describe the complex system of interdependent components that work together to enable a cloud-based information system, which can be orchestrated by multiple cloud Actors. Components of one cloud Ecosystem can be shared with other cloud Ecosystems serving different information systems.

Actor	Definition
Cloud Consumer	A person or organization that maintains a business relationship with, and uses service from, *Cloud Providers*.
Cloud Provider	A person, organization, or entity responsible for making a service available to interested parties.
Cloud Auditor	A party that can conduct an independent assessment of cloud services, information system operations, performance and security of the cloud implementation.
Cloud Broker	An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between *Cloud Providers* and *Cloud Consumers.*
Cloud Carrier	An intermediary that provides connectivity and transport of cloud services from *Cloud Providers* to *Cloud Consumers*.

The NIST RA diagram in **Figure 2** also depicts the three service models discussed earlier in **Section 3.2**: IaaS, PaaS, and SaaS in the 'inverted L" representations, highlighting the stackable approach of building cloud service. Additionally, the NIST RA diagram identifies, for each cloud Actor, their general activities in a cloud Ecosystem. This Reference Architecture is intended to facilitate the understanding of the operational intricacies in cloud computing. It does not represent the system architecture of a specific cloud computing system; instead, it is a tool for describing, discussing, and developing a system-specific architecture using a common framework of reference that we plan to leverage in our later discussion of key management issues in a cloud environment.

To enhance the NIST SP 500-292 cloud RA, NIST identified in NIST SP 500-299, *Cloud Security Reference Architecture*, two types of cloud Providers:

1. Primary Provider, and
2. Intermediary Provider;

and two types of cloud Brokers:

1. Business Broker, and
2. Technical Broker.

Figure 3, To enhance the NIST SP 500-292 cloud RA, NIST identified in NIST SP 500-299, *Cloud Security Reference Architecture*, two types of cloud Providers, the key management functions that fall under the Provider's responsibilities might need to be divided among the two Providers, depending on the architectural details of the offered cloud service. From the cloud Consumer's perspective this segregation is not visible.

A Primary Provider offers services hosted on an infrastructure that it owns. It may make these services available to Consumers through a third party (such as a Broker or

Intermediary Provider), but the defining characteristic of a Primary Provider is that it does not obtain the sources of its service offerings from other Providers.

An Intermediary Provider has the capability to interact with other cloud Providers without offering visibility or transparency into the Primary Provider(s). An Intermediary Provider uses services offered by a Primary Provider as invisible components of its own service, which it presents to the customer as an integrated offering. From a security perspective, all security services and components required of a Primary Provider are also required of an Intermediary Provider.

A Business Broker only provides business and relationship services, and does not have any contact with the cloud Consumer's data, operations, or artifacts (e.g., images, volumes, firewalls) in the cloud and, therefore, has no responsibilities in implementing any key management functions, regardless of the cloud architecture. Conversely, a Technical Broker *does* interact with a Consumer's assets; the Technical Broker aggregates services from multiple cloud Providers and adds a layer of technical functionality by addressing single-point-of-entry and interoperability issues.

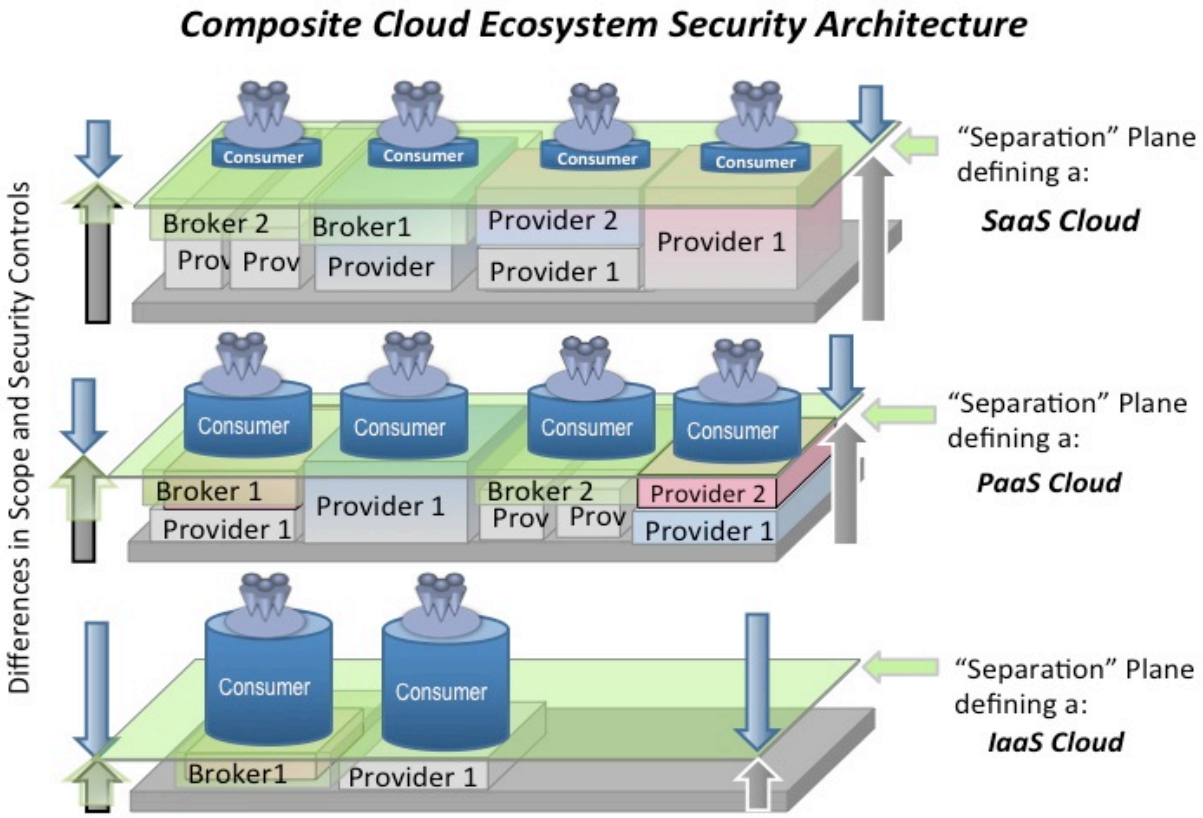

Figure 3: Composite cloud Ecosystem security architecture
(Source: NIST SP 500-299, draft)

There are two key defining features of a cloud Technical Broker that are distinct from an Intermediary Provider:

1. The ability to provide a *single consistent interface* (for business or technical purposes) to multiple differing Providers, and
2. The *transparent visibility* that the Broker allows into who is providing the services in the background – as opposed to Intermediary Providers that do not offer such transparency.

Since the Technical Broker allows for this transparent visibility, the Consumer is aware of which cloud capabilities are implemented by the Technical Broker versus the ones provided by cloud Provider(s) working with the Technical Broker. This case is different from the case in which an Intermediary Provider is involved, since the Intermediary Provider is opaque, and the Consumer is unaware of how the key management functions are divided, when applicable, between the Intermediary Provider and the Primary Provider.

3.4 Cloud Computing Security Essentials

Cloud computing provides enterprises with significant cost savings, both in terms of capital expenses (CAPEX) and operational expenses (OPEX), and allows them to leverage leading-edge technologies to meet their information processing needs. In a cloud environment, security and privacy are a cross-cutting concern for all cloud Actors, since both touch upon all layers of the cloud computing Reference Architecture and impact many parts of a cloud service. Therefore, the security management of the resources associated with cloud services is a critical aspect of cloud computing. In a cloud environment, there are security threats and security requirements that differ for different cloud deployment models, and the necessary mitigations against such threats and cloud Actor responsibilities for implementing security controls depend upon the service model chosen and the service categories elected. Many of the security threats can be mitigated with the application of traditional security processes and mechanisms, while others require cloud-specific solutions. Since each layer of the cloud computing Reference Architecture may have different security vulnerabilities and may be exposed to different threats, the architecture of a cloud-enabled service directly impacts its security posture and the system's key management aspects.

For each service model, **Figure 4** below uses a building-block approach to depict a graphical representation of the cloud Consumer's visibility and accessibility to the various layers of a cloud environment. As the figure shows, in a IaaS service model, the cloud Consumer has high visibility into everything above the API layer, while the cloud Providers implement controls below the API layer (which are usually opaque to Consumers). The cloud Consumer has limited visibility and limited key management control in a PaaS model, since the cloud Provider implements the security functions in all layers below the integration and middleware layer. The cloud Consumer loses visibility and control in a SaaS model, and in general, controls below the presentation layer are opaque to the cloud Consumer, since the cloud Provider implements all security functions.

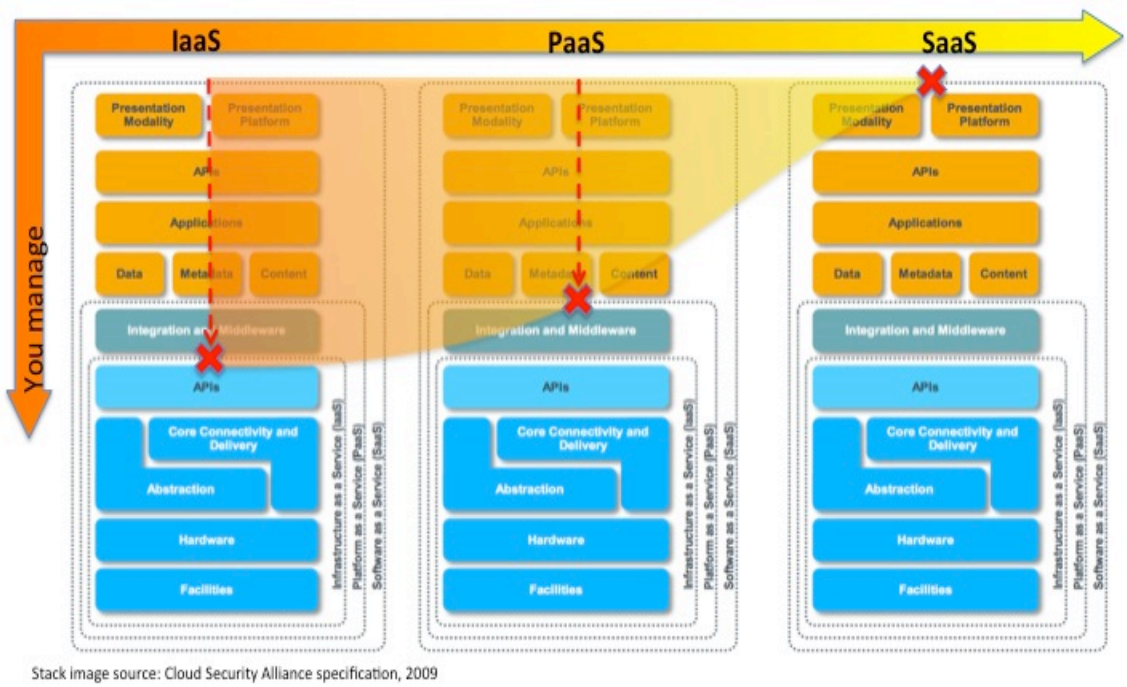

Figure 4: Consumer's level of control
(Source: NIST SP 800-173 (NIST internal working draft))

While all cloud Actors involved in orchestrating a cloud Ecosystem are responsible for addressing operational, security and privacy concerns, cloud Consumers retain the data ownership, and therefore remain fully responsible for:

- properly identifying data's sensitivity,
- assessing the risk from any exposure or misuse of the data and the impact to their business,
- identifying security requirements commensurable with the data sensitivity, and
- approving necessary risk mitigations.

Some of the cloud Consumer's areas of concern are:

- Risk Management
  - Risk Analysis
  - Risk Assessments
  - Vulnerability Assessments
  - Incident Reporting and Response
- Business Continuity
  - Disaster recovery plans
  - Restoration plan incorporating and quantifying the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) for services
- Physical Security
  - Physical and Environmental Security Policy
  - Contingency Plan

    - Emergency Response Plan
    - Facility Layout
    - Security Infrastructure
    - Human Resources
    - Environmental Security
    - Visual inspection of the facility
- User Account Termination Procedures
- Compliance with National and International/Industry Standards on Security
- Transparent view of the security posture of the cloud Providers, Brokers, and Carriers.

Technological advancements have led to cloud computing's emergence as a viable alternative for meeting the technology needs of many organizations. However, for cloud Consumers to take full advantage of cloud computing's economies of scale, flexibility, and overall full potential, Consumers need to address the concerns listed above and quantify the risk associated with the adoption of a cloud-based information system. Since gauging the risk and managing it in a cloud Ecosystem is a complex problem, a separate chapter, Chapter 8, is dedicated to this topic.

Cloud computing security refers to the set of procedures, processes and standards designed to provide information security assurance in a cloud Ecosystem. The massive concentration of specialized resources in a cloud Ecosystem has the potential to provide, on one hand, more robust, scalable and cost-effective defenses. On the other hand, these same specialized resources and the massive concentration of data present an attractive target to attackers.

Cloud computing security addresses both physical and logical security issues across all the different service models of software, platform and infrastructure. It also addresses how these services are delivered in the Public, Private, Hybrid and Community delivery models.

The new economic model facilitated by cloud computing technology has driven substantial technical changes for cloud-based information systems in terms of *scale*, *architecture*, *security*, and *privacy*.

> **Scale**. The commoditization of cloud computing and the organizations' drive towards economic efficiency have led to massive concentrations of hardware resources necessary to provide these services.
>
> **Architecture.** On-demand use of computing resources, the resources abstraction from the underlying hardware, and the multi-tenancy that brings together unrelated individuals or organizations who share hardware and software resources are only a few specific characteristics of this relatively new technology. Massively distributed computing, content storage, and data processing relying only on logical isolation mechanisms to protect it are also characteristics of cloud computing. Global markets for commodities demand edge distribution networks where content is delivered and received as close to customers as possible. This tendency towards global distribution and redundancy provides increased resilience for the cloud-based information

systems while, on the down side, means that the resources are usually managed in bulk, both physically and logically.

Security. The centralization of data and increase in security-focused resources can improve security, but concerns can persist about losing control of certain sensitive data, and the lack of security for stored kernels. Security is often as good as or better than traditional systems, in part because providers are able to devote resources to solving security issues that many customers cannot afford to tackle. However, the complexity of security greatly increases when data is distributed over a wider area or over a greater number of devices, as well as in multi-tenant systems shared by unrelated users. In addition, user access to security audit logs may be difficult or impossible for cloud Providers to grant to cloud Consumers. Private cloud installations are in part motivated by users' desire to retain control over the infrastructure and avoid losing control of information security.

Privacy. Cloud computing possesses privacy concerns because the service providers have access to the data that is stored on their infrastructure. Cloud Providers could accidentally or deliberately alter or even delete information. Many cloud Providers can share information with third parties if necessary without a warrant. The permission is granted in their privacy policy, which users agree to before they start using cloud services. Privacy solutions include policy and legislation as well as end users' choices for how data is stored. Users can encrypt data that is processed or stored within the cloud to prevent unauthorized access.

Since different users are sharing a cloud Provider platform, there may be a possibility existing that information belonging to different customers resides on the same data server. Therefore, information leakage may arise unintentionally when information for one customer is given to another customer. Additionally, hackers are spending substantial time and effort looking for ways to find vulnerabilities in the cloud infrastructure that would allow them to penetrate the cloud. Because data from hundreds or thousands of companies can be stored on large cloud servers, hackers can theoretically gain control of huge stores of information through a single attack of the hypervisor - a process referred to as "hyperjacking."

Another cloud Ecosystem issue is the legal ownership of the data and the responsibilities and privileges of the data owner and data custodian. Because cloud Consumers retain ownership of the data residing in a cloud Ecosystem, they usually keep the security authorization in-house and are responsible for identifying all security requirements pertaining to the cloud Ecosystem's hosting and processing of this data. However, since a cloud Consumer's level of control and management of the cloud Ecosystem's stack is limited by the adopted cloud architecture (see discussion related to Figure 4), cloud Providers and cloud Technical Brokers (when involved) become the data custodians and are responsible for fulfilling all security and privacy requirements identified by the cloud Consumer. It is always recommended that cloud Consumers review the implementation of all the security and

privacy controls and ensure that all the requirements are met before authorizing the use of a cloud-based information system.

3.5 Dividing Operational Responsibilities

Once a cloud Consumer selects the most suitable cloud architecture and identifies the other cloud Actor partners to orchestrate the cloud Ecosystem, all Actors must work together to clearly identify their operational responsibilities. These responsibilities are often split among Actors with the level of responsibility shifting based on the deployment and service models adopted. Ideally, the cloud Consumer should be ultimately responsible for defining the security and privacy controls required to safeguard the data and cloud-based information system. The implementation of many of these controls is often the responsibility of the cloud Providers or cloud Technical Brokers (when involved).

Once the cloud architecture is defined, cloud Actors involved in orchestrating the Ecosystem identify the control interfaces exposed to cloud Consumers. Examples of control interfaces that a cloud Provider and/or Broker can expose include: system, security and application logs; broker APIs for instrumentation; or the Broker's web application for managing cloud Consumer applications. Ultimately, each cloud Actor is responsible for their respective operational tasks as defined in the security authorization for the cloud-based information system.

3.6 Visibility and Trust in the Cloud Ecosystem

Under the cloud computing paradigm, an organization relinquishes direct control over many aspects of security and privacy, and in doing so, confers a high level of trust onto the cloud Provider(s) and the cloud Technical Broker. At the same time, cloud Consumers, as data owners, have a responsibility to protect information and information systems commensurate with the risk and magnitude of the harm resulting from unauthorized access, use, disclosure, disruption, modification, or destruction, regardless of whether the information is collected or maintained by or on behalf of the cloud Consumer. In order to maintain trust in the cloud Ecosystem and properly mitigate risks associated with the cloud-based information systems, cloud Actors need visibility into each other's area.

Transition to cloud computing services entails a transfer of responsibility to implement necessary security and privacy controls to the cloud Providers and cloud Technical Brokers for securing portions of the system on which the cloud Consumer's data and applications operate.

Visibility into the way the cloud Provider operates, including the provisioning of composite services, is a vital ingredient for effective oversight over system security and privacy by a cloud Consumer. To ensure that policy and procedures are being enforced throughout the system life cycle, service agreements should include some means for the organization to gain

visibility into the security and privacy controls and processes employed by the cloud Provider and their performance over time.

Trust is an important concept related to risk management. How cloud Actors approach trust influences their behaviors and their internal and external trust relationships. The reliance on cloud computing services results in the need for trust relationships among cloud Actors. However, building *trustworthiness* requires visibility into Providers' and Technical Brokers' practices and risk/information security decisions to properly gauge the *risk* and estimate the *risk tolerance*. It is important to note that the *level of trust* can vary and the *accepted risk* depends on the established *trust relationship*.

The next section further discusses the importance of building trust and introduces the concept of *trust boundary*. Moreover, Chapter 8 discusses in detail the cloud Consumer's risk management in a cloud Ecosystem

3.7 Boundaries in a Cloud Ecosystem

In a cloud Ecosystem, it is of critical importance for cloud Consumers to establish the clear demarcation of information-system boundaries on all levels in a vendor-neutral manner. Furthermore, it is incumbent upon the cloud Consumer to establish measures to ensure appropriate protection, regardless of vendor, ownership, or service level for the cloud-based information system.

To avoid vendors lock-in and to allow for a vigilant improvement of designed countermeasures, cloud Consumers need not only establish a plan to adopt a cloud-based solution, but also be prepared to transition to alternate cloud Providers or Brokers. Therefore, at each layer and subsystem level, a cloud Consumer needs to identify the security and privacy controls and negotiate which cloud Actor is responsible for the implementation and operation of each control function. Each cloud Actor needs to monitor and manage the service levels and the licensure, and needs to support the integrity and availability of the information system on a boundary-by-boundary basis. Furthermore, if external integrations to the cloud service are providing functionality, data feeds, or services, all strata need to be identified and the information system control boundary established. Also, for the aggregates cloud service, cloud Actors need to establish clear ownership of the methodology to maintain, monitor, and protect the externally provided functionality, the transactions, and the associated data.

The process of establishing information system boundaries and the associated risk management implications remains an organization-wide activity independent of vendor interaction. Cloud Consumers need to carefully negotiate with all Actors participating in the orchestration of the cloud Ecosystem solutions for all of an organization's business requirements, all complex technical considerations with respect to information security and the programmatic costs to the organization.

To build the foundational level of protection for the data and to provide the adequate overall security posture of the cloud-based information system, the inherited security and privacy controls implemented by cloud Providers and cloud Technical Brokers (when participating

in the orchestration) need to be properly assess and monitor at each boundary. To elevate the systems' security posture and protect data commensurable with its sensitivity, cloud Consumers often need to negotiate tailoring of existing controls via parameter selection or via implementation of compensating security and privacy controls. Because data owners retain the responsibility and accountability to ensure that all cloud security controls are managed and tracked on an ongoing basis, it is important to incorporate in the security plans and in the service agreements, clear coordination of and consideration for:

- The selection, implementation, assessment, and monitoring of security controls for cloud-based systems;
- The effects of changes in the cloud service functionality on the overall security posture of the cloud-based information system and on the mission and business processes supported by that system; and
- The effects of changes to the information system on the cloud service and its controls.

Security controls identified by the cloud Consumer and implemented by cloud Actors are documented in the security plan for the holistic information system and assessed for effectiveness during the risk management process (i.e., during the initial authorization of the information system and subsequently during the continuous monitoring process). Cloud security controls are also assessed for effectiveness if additional functionality is added after the information system is authorized to operate.

As owners of the data, cloud Consumers need to take appropriate measures to ensure that changes within any inner boundary of the cloud system do not affect the security posture of the overall system. Additionally, they need to aggregate, at the data level, applications, platforms and infrastructure level, all pertinent information obtained from the cloud Providers and cloud Technical Brokers, and to consolidate the aggregated information, to conduct near real-time monitoring and to perform security impact analyses.

The following sections identify and discuss each logical or physical boundary in the cloud Ecosystem. When architecting a cloud-based information system and orchestrating the supporting cloud Ecosystem, the cloud Consumer starts by categorizing the user's data and the application and identifying the corresponding boundaries. Next, the Consumer needs to identify functional capabilities or components needed to support the application and secure the data, the multiple boundaries corresponding to the service model, the cloud Ecosystem's orchestration, the cloud deployment model, and last, but not least, the trust boundary. In the next sections, we discuss these boundaries.

3.7.1 User-Data Boundary

The core of the cloud Ecosystem is the user-data boundary. This boundary traverses all stackable functional layers of the cloud Ecosystem and contains the cloud Consumers data, which defines the required level of security in all outer layers. The way the user data boundary intersects with the presentation, API, and application boundary requires a clear understanding of the value of the information stored within the user-data perimeter and the corresponding security controls required to instrument said outer functional layers.

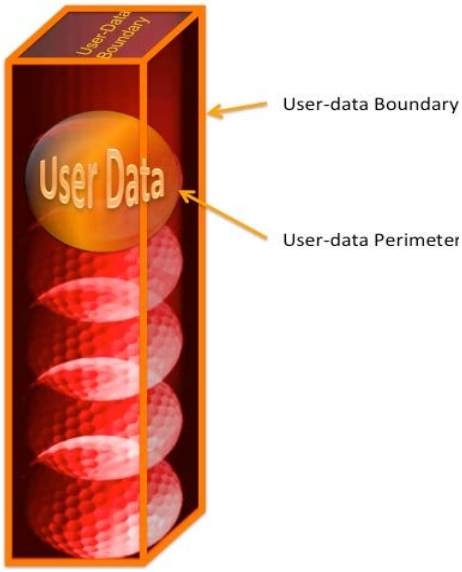

Figure 5: User-Data boundary

As the center of the cloud Ecosystem, the user data boundary (**Figure 5**) contains user data encompassed within the user-data perimeter. The user-data perimeter is the logical containerization of user data as it traverses the cloud Ecosystem between cloud Consumer and across all cloud Actors. As the user data contained within the *user-data perimeter* moves from cloud Provider to cloud Consumer, the *user-data boundary* traverses the presentation, the API, and the application boundary and needs to ensure the security of this information.

A data-centric architecture leveraging a boundary approach warrants that all elements of a cloud Ecosystem are designed and instrumented based on the sensitivity of the cloud Consumers' data.

3.7.2 Service Boundary

Service Boundary is a general concept introduced to identify the service layers acquired by a cloud Consumer or implemented by cloud Actors other than the Consumer.

This generic *service boundary* can be of IaaS, a PaaS, or SaaS type, based on the architectural service layers defined in NIST SP 800-145:

- Software as a Service Boundary
  - Presentation Modality Boundary
  - Presentation Platform Boundary
  - Application Programming Interfaces Boundary
  - Applications Boundary
  - Data Boundary
  - Metadata Boundary
  - Content Boundary
- Platform as a Service Boundary
  - Integration & Middleware Boundary
- Infrastructure as a Service Boundary

- Application Programming Interfaces Boundary
- Core Connectivity and Delivery Boundary
- Abstraction Boundary
- Hardware Boundary
- Facilities Boundary

The following sections discuss key elements of boundary definition and acceptable risk. Because the Consumer's view is provided in these sections, the functionality the Consumer manages is perceived as internal, and to better highlight the data-centric architecture with layers wrapping around user's data, the boundaries defining Consumer's managed layers are referred to as *internal service boundaries.* In contrast, the boundaries defining the layers managed by other cloud Actors (Provider, Technical Broker, etc.) will be referenced as *external service boundaries*. Moreover, due to the similarities in graphical representation between the three types of service boundaries, only a graphical representation for the Platform as a Service boundaries is provided below.

3.7.2.1 IaaS Security Boundaries

NIST SP 800-145 defines Infrastructure as a Service (IaaS) as follows:

> The capability provided to the [cloud Consumer] is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The [cloud Consumer] does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).[9]

There are both *internal* and *external* boundaries, which the cloud Consumer must establish with the cloud Provider to delineate management control and scope of responsibilities.

The *IaaS boundary* divides the cloud Ecosystem at the infrastructure layer exposing as a service the IaaS API, while delineating the layers *external* to Consumers as the interconnected stack that encompasses core connectivity, hardware and facilities.

In a logical way, outside the IaaS boundary lies the *Ecosystem orchestration boundary*, *cloud deployment boundary*, and *trust boundary.* The *internal* and *external* IaaS boundaries require coordination to establish an acceptable level of trust and coordination of security with other cloud Actors.

The Consumer establishes trust within the IaaS boundary in concert with any contracted service Providers. This trust must be established with the IaaS whether the service is provided within the Consumer's control or not. Well-defined boundaries should clearly delineate responsibilities for security, privacy, and quality of services within the *service*

[9] NIST SP 800-145, p. 3.

boundaries. Consumers need to assess the trustworthiness of all interfaces (logical and physical) with other Actors both inside and outside system boundaries.

The cloud deployment model chosen by the cloud Consumer has a direct impact on the trust relationship with the cloud Provider(s). For the IaaS service model, the cloud Consumer assumes a greater level of responsibility than the cloud Provider or other Actors for the service provided.

3.7.2.2 PaaS Security Boundaries

NIST SP 800-145 defined Platform as a Service (PaaS) as follows:

> The capability provided to the [cloud Consumer] is to deploy onto the [cloud Provider] consumer-created or acquired applications created using programming language, libraries, services, and tools supported by the provider. The [cloud Consumer] does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.[10]

The *PaaS boundary* divides the cloud Ecosystem at the platform layer offering an integrated development environment and integration point, while delineating the layers *external* to Consumers as the interconnected stack that bundles network, servers, operation systems, and storage, from the operating environment down to facilities, allowing cloud Consumers to deploy or build their choice of compatible applications.

Similar to the IaaS service boundaries, a PaaS-based Ecosystem has PaaS *internal* and *external boundaries*, which the cloud Consumer establishes with the cloud Provider to delineate management control and scope of responsibilities (see Figure 6 and Figure 7). Providers assume increasing levels of responsibility for implementing and monitoring security.

Figure 7 depicts the PaaS external boundaries consisting of an interconnected stack that links the facility boundaries of the IaaS with the integration boundary of the PaaS. Below the PaaS boundaries lay the API, the connectivity and delivery, the abstraction and control, and the hardware and facilities boundaries. The boundaries that are providing PaaS interfaces require coordination to establish an acceptable level of trust and coordination of security with the cloud Provider.

[10] NIST SP 800-145, p. 2.

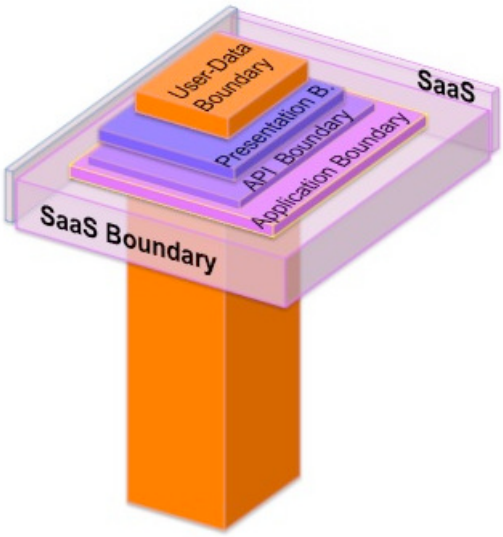

Figure 6: Platform as a Service boundary - Consumer's layers

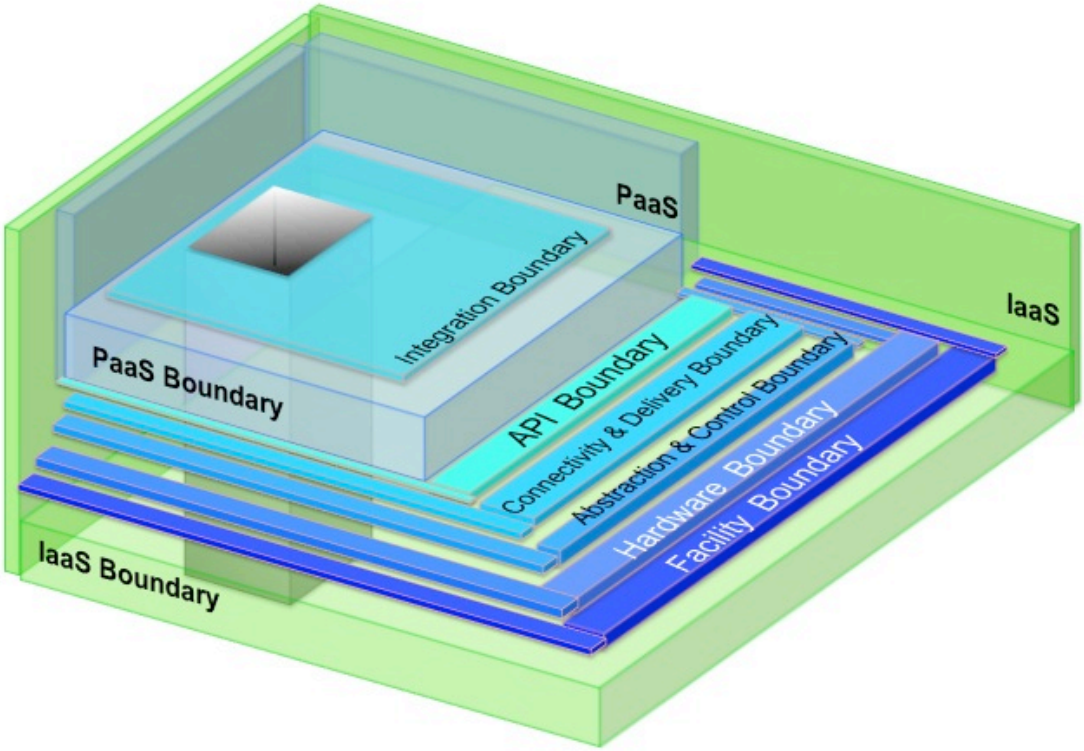

Figure 7: Platform as a Service boundary – Provider's layers

As mentioned previously, cloud Consumers need to assess the risk of using the system and to establish the risk tolerance. To authorize the use of the cloud service once the assessment is complete, Consumers need to establish, in concert with any contracted cloud Actors, a trust relationship with all parties involved in orchestrating the PaaS-based cloud Ecosystem. Cloud Providers, in most cases, assume greater responsibility for security and service coordination than cloud Consumers in a PaaS-based cloud Ecosystem.

3.7.2.3 SaaS Security Boundaries

NIST SP 800-145 defines Software as a Service (SaaS) as follows:

> The capability provided to the [cloud Consumer] is to use the [cloud Provider's] applications running [in a cloud Ecosystem managed by the Provider or Technical Broker]. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. [Cloud Consumers] do not manage or control the underlying cloud [Ecosystem] including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.[11]

Cloud Providers must assume the greatest level of responsibility for meeting all standard compliance requirements and for implementing and monitoring security and privacy controls.

The *SaaS boundary* divides the cloud Ecosystem at the application layer exposing as a service the application and SaaS API(s), while delineating the layers *external* to Consumers as the interconnected stack that encompasses from the applications layer down to facilities.

The *internal SaaS boundaries* consist of an interconnected stack of upper layer boundaries that include the user-data, presentation, API, and application. The SaaS *external boundaries* start at the SaaS layer and build upon PaaS external boundaries. Between the PaaS and SaaS layers lies the integration boundary. The boundaries that expose interfaces at the SaaS layer require operational, security and privacy coordination with the cloud Provider to establish an acceptable level of trust. Trust within a SaaS-based cloud Ecosystem needs to be established by the cloud Consumer in concert with any contracted cloud Actors (Providers, Brokers, etc.). Within the SaaS boundaries, establishing trust is not only more challenging but also is a more critical component since the provider is assuming most and sometimes all of the responsibilities for deploying and operating the service. Since the service is outside the cloud Consumer's physical or logical control, establishing and maintaining trust can only be done through well-defined deployment and orchestration boundaries with enforceable terms and conditions.

Relative to the IaaS and PaaS service models, in a SaaS-based cloud Ecosystem, cloud Providers assume the greatest responsibility for implementing security and privacy controls and coordinating and operating the service. The level of trust within SaaS boundaries and

[11] NIST SP 800-145, p. 2.

between the internal and external SaaS boundaries—for both cloud Consumer and cloud Provider—needs to be the highest attainable, and therefore more restrictive service agreements and SLAs are required, with well-defined penalties and liabilities.

3.7.3 Ecosystem Orchestration Boundary

To minimize business expenses and reduce the cost of cloud services, Providers design cloud solution sets targeting as many potential customers as possible. Such solutions are easier for industry segments to both understand and move workloads within and to the cloud. These pre-packaged solution sets often contain modules of components that are identical, with identical configurations, and that are easily reproducible in various cloud Ecosystems. This chapter defined the cloud Ecosystem as a complex system of interdependent components that work together to enable a cloud-based information system.

It is very important to note that while serving a cloud-based information system, a cloud Ecosystem can be orchestrated by multiple cloud Actors that collaborate to build it. The foundation of the Ecosystem is built by cloud Providers. Cloud Technical Brokers may provide layers of functionality that provide intermediation, aggregation or interoperability. The layers built by Brokers or Intermediate Providers inherit the controls from the lower layers in the stack implemented by Providers. Depending on the service model, a cloud Consumer adds functionality to the cloud Ecosystem, while inheriting security and privacy controls implemented by all other cloud Actors. Often, due to the multi-tenancy nature of cloud computing, components of one cloud Ecosystem are shared with other cloud Ecosystems serving different information systems.

Moreover, with the exception of an on-premise Private cloud, most clouds run in third-party data centers. And, even in an on-premise Private cloud, there are likely to be provisions for "*cloud burst*," into another cloud under extreme conditions. One of the impediments to broader cloud computing adoption is the cloud Consumers' inability to continuously monitor the controls implemented by other cloud Actors or the operation of the components managed by these Actors. By ensuring that all cloud Actors have a clear understanding of their responsibilities and that the cloud Actors properly implement agreed-upon security and privacy controls as identified in the security plans, it is possible for the cloud Actors to define the cloud *ecosystem orchestration boundary* and to properly assess the inherited risk from the use of the particular orchestration for the information system under discussion.

Orchestration of the cloud Ecosystem allows Public, Private, and Hybrid clouds to operate with elasticity, scale, and efficiency. The *Ecosystem orchestration boundary* is identified when the decisions are made to include certain cloud Actors and to define their responsibilities. For example, a cloud Ecosystem may be supported by a single cloud Provider that offers its services to a cloud Consumer. Alternatively, a similar SaaS-based Ecosystem might be architected such that services from multiple cloud Providers are aggregated by a Technical Broker and offered to a cloud Consumer as a SaaS-based information system. In particular cases, cloud Consumers might prefer to gain more control over the cloud Ecosystem and therefore decide to leverage PaaS or IaaS services to build a similar information system by adding the necessary functional layers to the PaaS or IaaS offer, composing a final SaaS-like

solution. Figure 8 graphically depicts the alternatives described above while highlighting the cloud *Ecosystem orchestration boundary.*

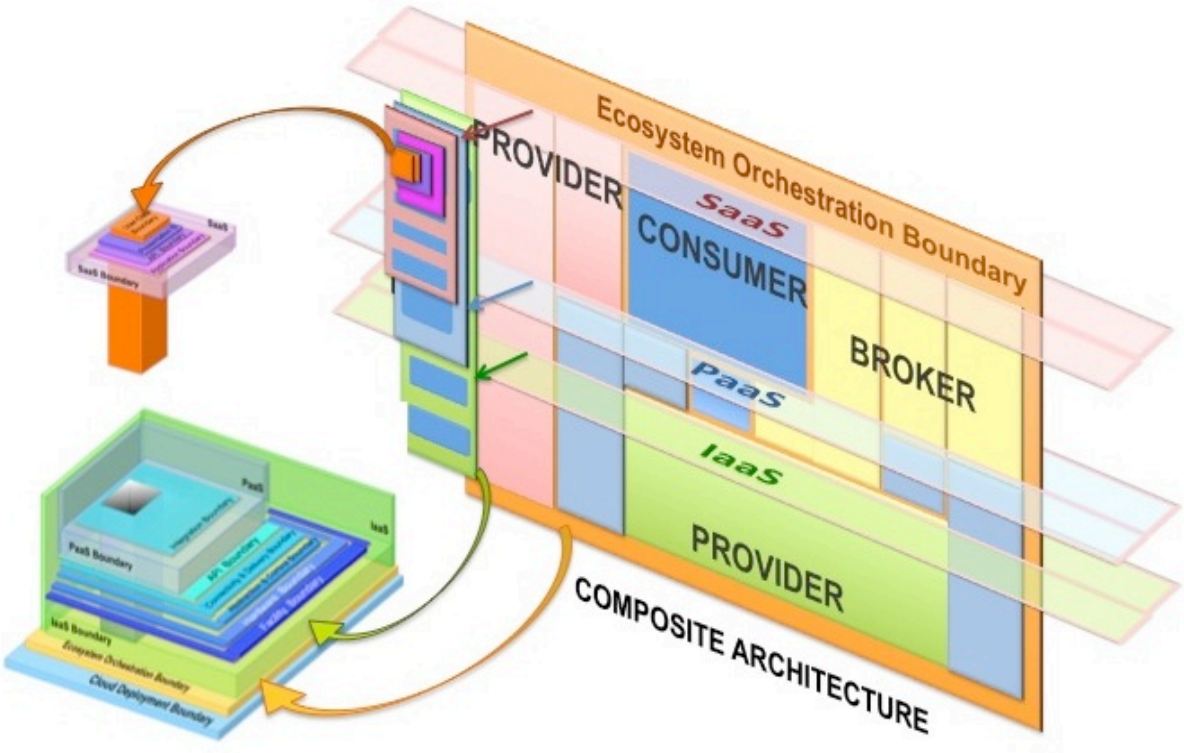

Figure 8: Cloud Ecosystem orchestration boundary

The *Ecosystem orchestration boundary* needs to incorporate automated workflow functionality and management of the cloud Ecosystem's components (e.g. compute, Identity, Credential and Access Management (ICAM). A cloud Actor that orchestrates the cloud Ecosystem needs to ensure that all cloud resources serving an information system and their configuration management capabilities are identified and placed inside the *ecosystem orchestration boundary* for both proper assessment of the inherited risk and adequate continuous monitoring. When identifying the *ecosystem orchestration boundary*, it is important to ensure that all configurable interconnections and interactions among cloud-based and on-premises resources (dependent on the cloud deployment model) are accounted for. Cloud orchestration is complex as it involves accounting for automation of interconnected processes running across heterogeneous systems, potentially in multiple locations. Often processes and transactions may have to cross multiple organizations, systems, networks and boundary-protection devices.

The orchestration function is a high-priority target from a threat perspective. Properly identifying all orchestration components and including them within the cloud *Ecosystem orchestration boundary* to be accounted for and detailed in the information system security plan is critical.

3.7.4 Deployment Boundary

Once the cloud *Ecosystem orchestration boundary* is established, the next logical step is to select the cloud deployment model that best meets cloud Consumer's needs. The four types of cloud deployment models are Private, Public, Hybrid, and Community. A cloud *deployment boundary* is a logical boundary, which provides a common framework for assessing the level of exclusivity the cloud Consumer needs for the cloud-based information system. Often the information system's impact level drives the final decision regarding the cloud deployment model. In Figure 9, the cloud deployment boundary is graphically represented depicting all the elements contained therein, including the ecosystem orchestration boundary, IaaS, PaaS, SaaS and User Data boundaries.

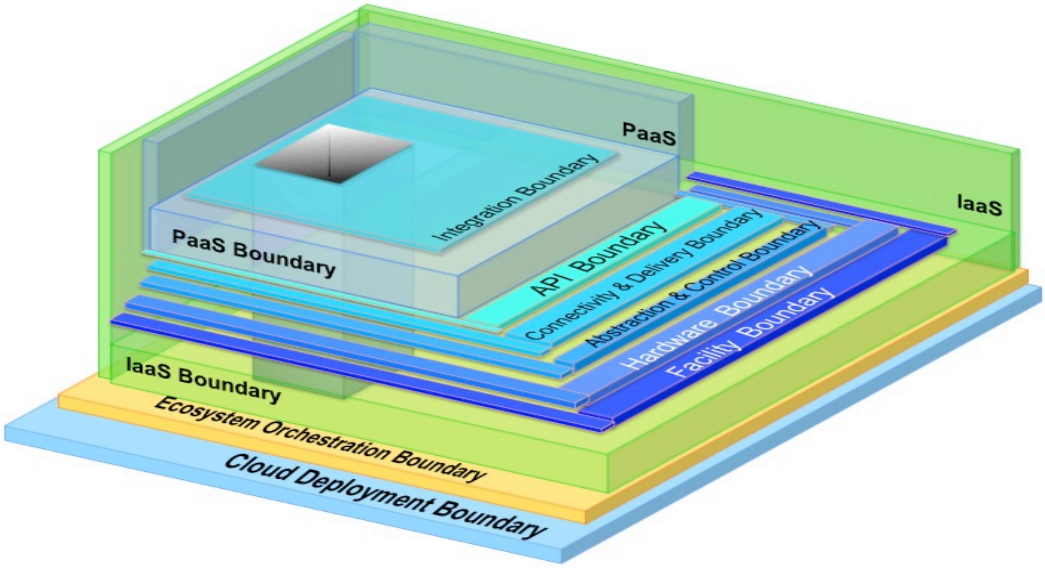

Figure 9: Deployment Boundary with PaaS external layers

The NIST Cloud Computing Reference Architecture (NIST SP 500-292) and NIST Cloud Computing Security Reference Architecture (Draft NIST SP 500-299) documents introduce and discuss these deployment models:

Private – The cloud's infrastructure is operated for the exclusive use of a single owner. The cloud instance could be managed by the owning organization or run by a third party. Private cloud can be on- or off-premises.

Public – The cloud's infrastructure is available for public use alternatively for a large industry group and is owned by an organization selling cloud services.

Community – Provides a cloud instance that has been organized to serve a common purpose or function.

Hybrid – Provides for an integration of multiple cloud models (private, public, community) where those cloud tenants retain uniqueness while forming a single unit. Common ubiquitous protocols are provided to access data for presentation.

3.7.5 Trust Boundary

In order to consume a service, a cloud Provider and a cloud Consumer each has to extend trust beyond their own IT resources, beyond the demarcation service access point between the cloud Consumer and other cloud Actors. A cloud Consumer is responsible for the implementation of the security and privacy controls required on its side, but is dependent on the service(s) implemented by the other cloud Actors. Many of the security and privacy controls implemented by cloud Consumer are inherited from the other cloud Actors. Therefore, a cloud Consumer entrusts the cloud Provider and associated Actors with implementing the security measures necessary to protect the cloud Consumer's data and to fulfill the Service Agreement and the Service Level Agreement, if they exist. Identifying all system components, deciphering the intricacy of this complex Ecosystem, identifying the logical boundary of all trusted components that service the cloud-based information system and constitute the cloud Ecosystem – the *trust boundary*, and ultimately building a trust relationship among cloud Actors is critical for cloud Consumers and for the successful deployment and operations of the cloud-based information system.

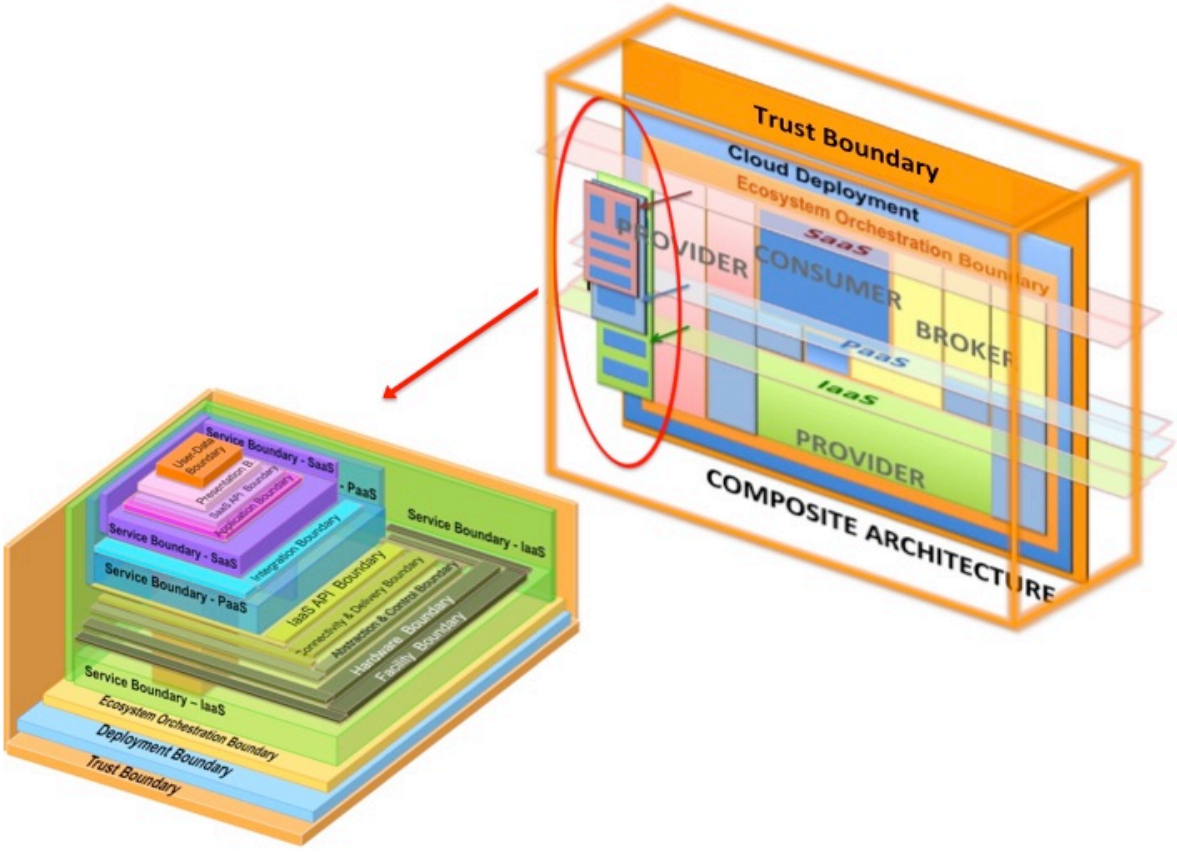

Figure 10: Trust Boundary – Concept explained

A *trust boundary* is the logical perimeter that typically spans beyond physical boundaries to represent the extent to which cloud-based IT resources are trusted within an established cloud Ecosystem (see Figure 10 for a graphical representation of the concept).

This extended *trust boundary* encompasses the resources from all cloud Actors and identifies a logical dynamic border of the cloud-based information system and of the supporting subsystems, viewed from the cloud Consumer's perspective. The *trust boundary* is elastic and adapts to the cloud Ecosystem's dynamic changes triggered by provisioning or decommissioning of the resources, and by data securely traveling or resting.

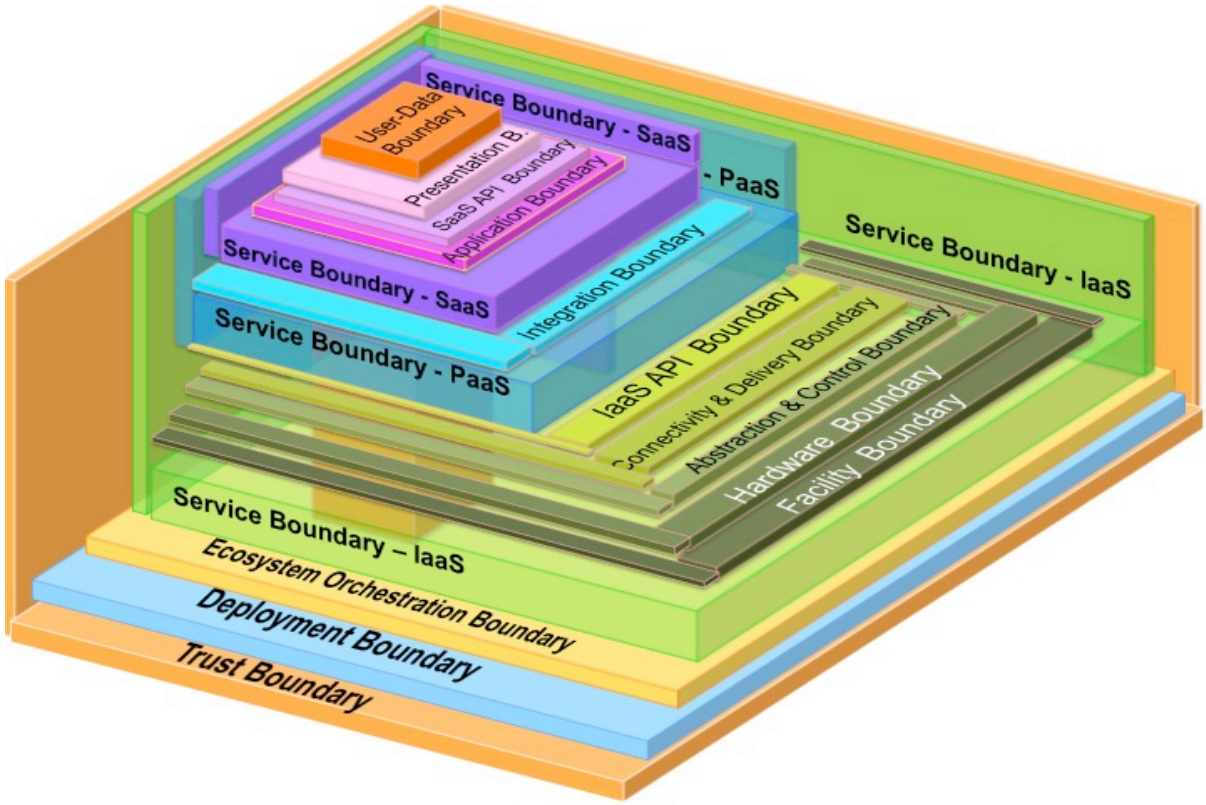

Figure 11: Trust Boundary

To build and maintain trust in the cloud Ecosystem, cloud Consumers need to be able to examine the security controls deployed inside this boundary and to determine the organization's *risk tolerance* to the confidentiality, integrity and availability risks resulting from operating this cloud-based information system. The typical method to establish an agreement with a cloud Provider is via Service Agreement and Service Level Agreements that describe security needs, capabilities, and agreed-upon standards, policies, and methods of trust implementation (including monitoring and auditing).

Figure 11 depicts the trust boundary as the outer-most of the boundaries.

For example, building trust and identifying the *trust boundary* in an IaaS cloud Ecosystem, means establishing the process for creating trusted platforms and aggregating them into

trusted pools of resources at design time. At run time, *trust boundaries* become elastic and dynamically adjust as the multi-tenancy and resource pooling characteristics of the cloud are exhibited. For example, a "burst out" to a cloud from on-premise resources requires that the trust boundary dynamically re-shapes to cover the "burst out" cloud compute infrastructure, and therefore this infrastructure needs to be trusted. At the other end, the users accessing the cloud resources needs to be trusted, so the supporting authentication and access control mechanisms and the networking that connects them to the resources need to be trusted. In this scenario, "trusted" means that the level of assurance has been established and the security posture of the components has been assessed and the residual risk gauged for all aspects of the processing based on the sensitivity of the data at the *user-data boundary*.

At run time, auditing and logging need to support assurance mechanisms that all critical aspects of the *trust boundaries* are present for workload processing and are meeting data confidentiality, integrity, and availability requirements. Continuous monitoring is also required for the status of the security program and serves as a critical part of the risk management process. The organization's overall security architecture and accompanying security policies and controls are monitored to ensure that organization-wide operations remain within an acceptable level of risk, considering any changes that occur.

3.8 Defining your root of trust

Trust is an intransitive relation with a specific hierarchy. What that means is that trust flows down a chain until it reaches the *root of trust*. Cloud implementations have multiple layers of abstraction, from hardware to virtualization to guest operating systems. The security and privacy of the user's data depend on the integrity and *trustworthiness* of the cloud Ecosystem, which depends on the cumulative *trustworthiness* of the layers that could potentially manipulate or compromise data integrity or confidentiality. The *trustworthiness* of each layer relies on the hardware or software secure modules (HSM/SSM) that are inherently trusted and that perform the cryptographic functions engineered to secure the data and the operations of each layer of the cloud stack.

Understanding who owns the root of trust is a foundational element to the architecture of an information system. Roots of trust are not only the underlying anchors for all compute elements that support secure operations of the cloud Ecosystem, but the roots of trust need to be trusted by the cloud Actors in order to assess the integrity and trustworthiness of the cloud Ecosystem, to identify the trust boundary and to build the necessary trust relationship among cloud Actors.

In a data-centric architecture, it is important that the cloud Consumer owns the root of trust as it pertains to the cloud Consumer's user-data and associated user-data boundary. This means that cloud Consumer needs to own the cryptographic keys used by the HSM/SSM that is securing the cloud layers (storage, hypervisors, virtual machines (VM), applications, and user-data at rent, in transit and in memory). The cloud Consumer should own the key used to secure the lowest common denominator of the cloud Ecosystem based on the sensitivity of the data housed therein. Information systems containing non-sensitive data may only need to have the cloud Carrier encrypted. Information systems containing more sensitive

information may require Virtual Machine or Storage Encryption, wherein the cloud Consumer owns the key and the Virtual Machine or Storage is unlocked using a hardware or software encryption appliance. Cloud Access Security Brokers serve to encrypt data in transit and at rest within cloud Providers, ensuring that cloud Consumers' data remains encrypted as it traverses the cloud ecosystem. Defining a root of trust is a critical element of cloud architecture, and should be determined before issuing a security authorization for the information system.

3.9 Managing user authentication and authorization

Understanding and defining user authentication and authorization among cloud Actors is another critical element of cloud architecture. Without knowing who is logging into the cloud-based information system, and who is accessing what data, cloud Actors are not able to protect the data housed by a cloud Ecosystem. Understanding who the users are, what data they are trying to access, where the data is stored and how are users trying to get to this data—these are critical pieces of information that help cloud Consumers determine an appropriate cloud architecture and deployment model.

User authentication is the process of establishing confidence in the *identity* of a user, typically by entry of a valid username and a valid *token* (password, key, and biometrics information) for the purpose of granting access to a particular information system(s) or resources. An authentication server compares the user's authentication *credential(s)* with the database storing all user credentials. A *credential* is an object or data structure that authoritatively binds an *identity* (and optionally, additional attributes) to a *token* possessed and controlled by the user. For example, a *username* and *password* pair is a data structure or a credential. If the provided credential matches the information in the authentication database, the user is granted access to the information system. If the credential does not match, the authentication fails and access is denied.

The type of credential used should be commensurate with the level of assurance defined by the sensitivity of the cloud Consumer's user-data. By leveraging user, data and location, a varying level of credentials can be used if any of the aforementioned variables change. For example, if a cloud user is currently in the United States and normally accesses a cloud information system via a Web browser on their personal computer, they would be prompted to enter their username and password to access said system. In the background, the information system can verify additional information collected from the user's device, such as geo-location, IP address, etc. When the same cloud user travels internationally and accesses the cloud information system via a Web browser on a public computer, the cloud information system's authentication server can identify a different IP address or a different geo-location. As soon as this new information is collected from the user's device, the system can prompt the user to provide additional credentials for a higher level of assurance while validating the identity of the user before granting said access.

User authorization is the process of enforcing policies such as determining what resources or services a user is permitted to access. Typically, *user authorization* occurs within the context of authentication. Once a user is authenticated, they may be authorized to access different components of a cloud information system. Ensuring that *user authorization* is

applied to the lowest common denominator of each element of a cloud Ecosystem is vital to ensuring the security of the data stored within the cloud information system. Granting users more authority than they require can compromise a system. Furthermore, safeguarding user *credentials* to protect against tampering or misuse is critical and needs to be part of the *security policies* employed within the security authorization program of the cloud-based information system.

Enforcing authorization *policies* is critical in a cloud Ecosystem. The enforcement can be instrumented by the user authentication and authorization server. The cloud Ecosystem architecture will dictate which cloud Actor is responsible for managing the server and, authenticating and authorizing users. Effective management of user authentication and authorization is a vital element of a secure cloud information system. Cloud Consumers are required to select the best fitting solution for their cloud-based information system, since the *user authentication* and *authorization* processes [12], policies [13] and procedures [14] are instrumental in protecting their data in a cloud Ecosystem.

In summary, technological advancements have led to ubiquitous cloud computing, which emerge as the most viable alternative for meeting the technology needs of many organizations. However, for cloud Consumers to take full advantage of cloud computing's economies of scale, it is important to build the necessary level of trust and gain visibility into the service in order to fully leverage the cutting-edge technologies embedded into cloud Providers' and cloud Technical Brokers' offers and to provision resources quickly and elastically, in a manner commensurable with the speed and dynamic changes of the business.

[12] Processes are a high level, overall view of the identified tasks.
[13] Policy is a guideline or law that drives the processes and procedures
[14] Procedures are the detailed steps required to perform an activity or a task within a process.

References

NIST Special Publication 500-292, *NIST Cloud Computing Reference Architecture*, September 2011.

NIST Special Publication 500-299, *NIST Cloud Security Reference Architecture* (draft).

NIST Special Publication 800-144, *Guidelines on Security and Privacy in Public Cloud Computing*, December 2011.

NIST Special Publication 800-145, *The NIST Definition of Cloud Computing*, September 2011.

NIST Special Publication 800-146, *Cloud Computing Synopsis and Recommendations,* May 2012.