IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AGENT BASED DATA SECURITY APPROACH FOR HYBRID CLOUD COMPUTING

SHRAVYA SUDHA O
PG Scholar, Dept of CSE
Cambridge Institute of Technology
Bengaluru, India

PANKAJA K
Associate professor, Dept of CSE
Cambridge Institute of Technology
Bengaluru, India

Abstract:

Cloud Computing sees a technical and cultural shift of computing service provision from being provided locally to being provided remotely, and en masse, by third-party service providers. Users have lost control over the protection of their data: No longer is our data kept under our own watchful eyes. In general, an information security approach estimates the risk, where the risk is to occur due to an unusual event, and the associated consequences for cloud organization. Information Security and Risk Management (ISRA) practices vary among cloud organizations and disciplines. There are several approaches to compare existing risk management methods for cloud organizations but their scope is limited considering stereo type criteria, rather than developing an agent based task that considers all aspects of the associated risk. It is the lack of considering all existing renowned risk management frameworks, their proper comparison, and agent techniques that motivates this research. This paper proposes Agent Based Information Security Framework for Hybrid Cloud Computing as an all-inclusive method including cloud related methods to review and compare existing different renowned methods for cloud computing risk issues and by adding new tasks from surveyed methods.

Keywords:
Lambda services , Hybrid cloud , SQS ,SNS SES services ,Dynamo db, Encryption , IP & MAC Address ,SSH(perm file)

1. Introduction:

Traditional Computing Models. Data is a valuable resource for organizations and individuals. Their management is an important task and includes assuring integrity of data. For decades, organizations, Individuals have been using computer hardware such as hard discs, DVDs, CDs, discs, and floppy discs to store their data.

Introduction of database systems enhanced information management and made it more effective [1]. The last decades have made a reality information processing where data can be processed efficiently on large computing and storage platforms accessible via the Internet [2]. Advancements in database systems and networking (including the Internet) enabled the development of new computing models. They include grid computing developed in the early 1990s; as well as utility computing and cloud computing, developed around 2005. Cloud computing (CC) can be defined as a computing model that enables convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management efforts or service provider interactions. Cloud computing involves the provision and use of IT infrastructure, platforms and applications 2 of any kind in the form of services that are electronically available on the Internet. Just a few examples of applications using cloud services include: online file storage, social networking sites, web mail, and online business applications [2]. As enterprises strive to identify new methods for driving their businesses forward, surging demand has shifted to solutions that provide lower-cost solutions for use of computing systems (both in terms of access to computing infrastructure and operating costs).This resulted in the exponential growth of cloud computing, which was found to be more effective than the earlier solutions [3]. As technological advances continue, the reach and influence of cloud computing continue to rise. Even so, when organizations outsource data and business applications to CC providers (who are third parties for them), security and privacy issues emerge as crucial concerns. Virtualization in Cloud Computing. Virtualization is one of the key technologies of cloud computing services, facilities, aggregation of multiple standalone systems into single hardware platform by virtualizing computing resources (e.g.: Network, CPU's, Memory, Storage). Virtualization is enabled by hardware abstraction, which hides the complexity of managing the physical computing platform and simplifies scalability of computing resources. It is implemented via hypervisors. A hypervisor is responsible for isolation of Virtual Machines (VMs), so that

they are prevented from to directly accessing other VMs' virtual disks, memory, or applications on the same host. Virtualization provides scalability and multi-tenancy (the latter occurs when a single instance of a software application serves multiple customers [4]). These two properties are significant characteristics of CC, and facilitate sharing and pooling of resources in order to improve agility, flexibility, reduce costs and enhance business value. The practical aspects of virtualization related to configuration, networking, and sizing of cloud systems are faced with challenges [8, 11] . In cloud virtualization, provisioning is a basic mechanism for allocation of a cloud provider's resources to a customer. When a cloud provider accepts a request from a customer, it must create the appropriate number of virtual machines (VMs) and allocate resources to support them.

The process is conducted in several different ways: advance provisioning, dynamic provisioning, and user self-provisioning [7]. The dynamic provisioning of cloud resources and services faces a number of challenges such as the optimal configuration for VMs, and the limitations in the number and capabilities of CPUs,memories, disks and network bandwidth to be partitioned among the resident VMs [12]. Cloud service providers undertake a significant effort to make the virtualization mechanisms secure by striving to eliminate or at least reduce vulnerabilities, threats and attacks. Classic "CIA" Security Triad. A classic definition of security—in terms of its basic characteristics—specifies it in terms of the CIA triad; the acronym "CIA" stands for confidentiality, integrity and availability--three key requirements for any secure system [16, 17].They are defined as follows: Confidentiality: It is the ability to hide information from those people unauthorized to view it. It is the basis of many security mechanisms protecting not only information but other resources. Integrity: It is the ability to ensure that data is an accurate and unchanged representation of the original information [10]. Availability: It ensures that a resource is readily accessible to the authorized user upon the user's request. This model is applicable across the whole subject of security analysis, from access to a user's Internet history to security of encrypted data across the Internet. [11].

2. Cloud Computing Definitions:

A scientific definition is proposed by the GRIDS Lab at the University of Melbourne:"A Cloud is a type of parallel and distributed system consisting of a collection of interconnected and virtualized computers that are dynamically provisioned and presented as one or more unified computing resources based on service-level agreement established through negotiation between the service provider and consumers."

Berkeley's defines it as:"Cloud Computing refers to both the applications delivered as services over the Internet and the hardware and systems software in the data centers that provide those services (Software as a Service -SaaS).

i. Storage-as-a-Service
ii. Database-as-a-Service
iii. Information-as-a-Service
iv. Process-as-a-Service
v. Application-as-a-Service
vi. Integration-as-a-Service
vii. Security-as-a-Service
viii. Management/Governance-as-a-
ix. Service
x. Testing-as-a-Service

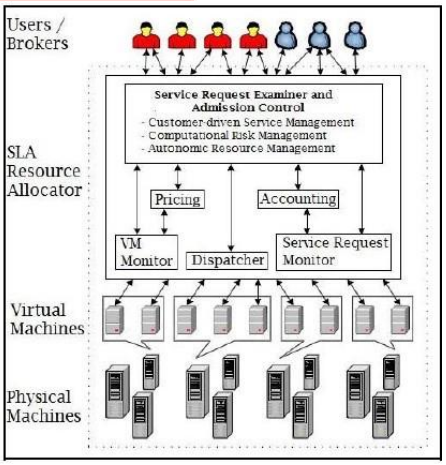

Fig.1. Cloud Computing

3. Hybrid Cloud Computing:

A hybrid cloud is a composition of at least One private cloud and at least one public cloud.

A hybrid cloud is typically offered in one of two ways: a vendor has a private cloud and forms a partnership with a public cloud provider, or a public cloud provider forms a partnership with a vendor that provides private cloud platforms. A hybrid cloud is a cloud environment in which an organization provides and manages some resources in- house and has others provided externally. For example, an organization might use a public cloud service, such as Amazon Simple Storage Service (Amazon S3) for archived data but continue to maintain in- house storage for operational customer data. Ideally, the hybrid approach allows a business to take advantage of the scalability and cost-effectiveness that a public cloud environment offers without exposing mission-critical applications and data to third-party vulnerabilities. This type of hybrid c loud is also referred to as hybrid IT.

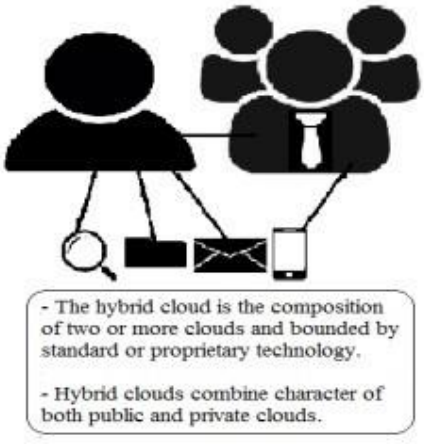

Fig.2. Hybrid Cloud Computing

4. AWS Lambda:

AWS Lambda is an event-driven, server less computing platform provided by Amazon as a part of Amazon Web Services. It is a computing service that runs code in response to events and automatically manages the computing resources required by that code.

AWS Lambda was designed for use cases such as image or object uploads to Amazon S3, updates to Dynamo DB tables, responding to website clicks, or reacting to sensor readings

from an IoT connected device as shown in Fig.3. AWS Lambda can also be used to automatically provision back-end services triggered by custom HTTP requests, and "spin down" such services when not in use, to save resources.

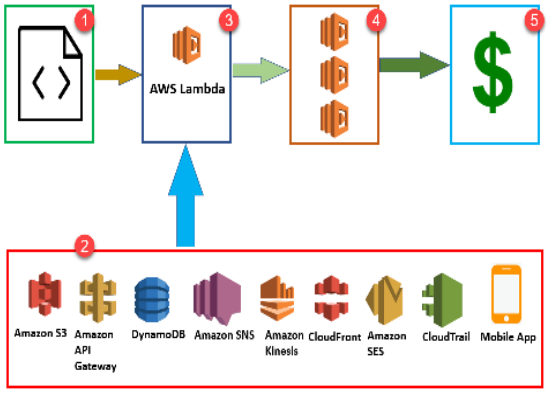

Fig. 3.AWS lambda flow

Step 1: First upload your AWS Lambda code in any language supported by AWS Lambda. Java, Python, Go, and C# are some of the languages that are supported by AWS Lambda function.
Step 2: These are some AWS services which allow you to trigger AWS Lambda.
Step 3: AWS Lambda helps you to upload code and the event details on which it should be triggered.
Step 4: Executes AWS Lambda Code when it is triggered by AWS services:
Step 5: AWS charges only when the AWS lambda code executes, and not otherwise.

5. SQS,SNS,SES:

SQS-Amazon Simple Queue Service (SQS) is a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating message oriented middleware, and empowers developers to focus on differentiating work.

Amazon Simple Notification Service (SNS) is a cloud service for coordinating the delivery of push messages from software applications to subscribing endpoints and clients. All messages published to Amazon SNS are warehoused across several availability zones to prevent loss.

Amazon Simple Email Service (SES) is a flexible and highly-scalable email sending and receiving service. Using SES, you can send emails with any type of correspondence

6. Dynamo DB:

Amazon Dynamo DB is a key-value and document database that delivers single-digit millisecond performance at any scale. It's a fully managed, multi-region, multi-active, durable database with built-in security, backup and restore, and in-memory caching for internet-scale applications. Dynamo DB can handle more than 10 trillion requests per day and can support peaks of more than 20 million requests per second.

The following are the basic Dynamo DB components:

i. **Tables** – Similar to other database systems, Dynamo DB stores data in tables. A *table* is a collection of data. For example, see the example table called *People* that you could use to store personal contact information about friends, family, or anyone else of interest. You could also have a *Cars* table to store information about vehicles that people drive.

ii. **Items** – Each table contains zero or more items. An *item* is a group of attributes that is uniquely identifiable among all of the other items. In a *People* table, each item represents a person. For a *Cars* table, each item represents one vehicle. Items in Dynamo DB are similar in many ways to rows, records, or tuples in other database systems. In Dynamo DB, there is no limit to the number of items you can store in a table.

iii. **Attributes** – Each item is composed of one or more attributes. An attribute is a fundamental data element, something that does not need to be broken down any further.

For example, an item in a *People* table contains attributes, called PersonID, Last Name, First Name, and so on. For a Department table, an item might, have such as DeptID, Name, Manager, and so on. Attributes in Dynamo DB are similar in many ways to fields or columns in other database.

Fig. 4. Person Table with attributes,items

There is a usability issue with the PKI certificate authentication, since non-IT users might have difficulties understanding how to keep public/private keys secret and yet available on a computer for login.

ii. PKIX.509 Methodology:

Fermi-Cloud [18] uses this approach for authentication and authorization - it utilizes PKIX.509 certificates for user identification and authentication. Fermi Cloud is built in OpenNebula1and it develops both X.509 authentication in Sun stone Open Nebula - a Web interface intended for user management – and X.509 authentication via command line interfaces. To avoid the limitations of Open Nebula access control lists that are used for authorization after successful authentication of users, authors integrated an existing local credential mapping service. This solution has also been extended in cloud federations to authorize users across different cloud providers that have established trust relationships through trusted certification authorities. The important functionalities of identity management systems for the success of clouds in relation to consumer satisfaction are discussed by Leonardo et al. In [19].

iii. Trusted Cloud Computing Platform (TCCP):

A trust relationship between a given user and SaaS domains is required so that SaaS user scan access the application and resources that are provided. In a PaaS domain, there is an

7. Literature Survey:

i. Multi factor authentication

There has been numerous research on security and privacy of sensitive data in cloud computing in both industry and academia. Most notably developing protocols and tools for anonymization or encryption of data for confidentiality purposes. To identify the related research on security and privacy of sensitive data in the cloud, we categorized the activities of cloud providers into five main categories: service deployment, resource abstraction, physical resources, service management, security, and privacy Multifactor authentication is a mechanism that requires more than one factor to verify the identity of users to achieve strong authentication.

For example, smart cards and tokens [15] or two-factor authentication [14] through popular services such as Short Message Service (SMS) or direct phone calls are multi-factor authentication mechanisms that effectively protect classified information. We use two-factor authentication using mobile devices and Yubi key tokens because setting up the infrastructure for smart cards, phone calls or SMS based authentication can become costly to maintain. Public Key Infrastructure (PKI) is another strong authentication approach that provides scalable secure communication solutions in open networks based on an asymmetric pair of keys known as public (shared with all the parties) and private(owned only by the user) keys [16], [17].

interceptor that acts as a proxy to accept the user's requests and execute them.The interceptor interacts with the STS, and requests the security token using the WS-Trust specification. Santos et al. [21] extend the Terra [22] design that enables users to verify the integrity of VMs in the cloud. The proposed solution is called the Trusted Cloud Computing Platform (TCCP), and the whole IaaS is considered to be a single system instead of granular hosts in Terra. In this approach, all nodes run a trusted virtual machine monitor to isolate and protect virtual machines. Users are given access to cloud services through the cloud manager component. The External Trusted Entity (ETE) is another component that provides a trust coordinator service in order to keep track of the trusted VMs in a cluster. The vulnerability reports in [23] show 59 vulnerability cases for Xen and 38cases for KVM. Approximately, 50 percent of all vulnerabilities are common regarding confidentiality, integrity, and availability. The

remote management software of Xen contributes to 15.3 percent of the vulnerabilities that demonstrates the increase attack surface by non-essential services. VM management component contains 11.9percent of the vulnerabilities in Xen compared to 5.3 percent in KVM. The lower vulnerability rate in KVM is due to the libvirt toolkit inside the hypervisor, where Xen allocates an entire privileged are in Dom0. Janus Version 2 (J2) [24] uses filtering and sand boxing, and Ostia [25]uses a delegation. Ostia also provides a hybrid interposition architecture, which allows for kernel level enforcement and user policies. However, system call

interposition has many problems. OS semantics are very difficult to replicate correctly. Indirect paths to resources are often overlooked, and there are side effects to denying system calls [26].

iv. Centralized-monitoring solution for cloud applications

Anand [27] present a centralized monitoring solution for cloud applications consisting of monitoring the server, monitors, agents, configuration files and notification components. Redundancy, automatic healing, and multi-level notifications are other benefits of the proposed solution and typical drawbacks of a centralized monitoring system, such as limited scalability, low performance and single point of failure.

8. Proposed system:

In the proposed system, we present a user-centric privacy-preserving cryptographic access control protocol calledK2C(Key To Cloud) that enables end-users to securely store, share, and manage their sensitive data in an untrusted cloud storage anonymously.K2Cisscalable and supports lazy revocation.

It can be easily implemented on top of existing cloud services and APIs – we demonstrate its prototype based on Amazon S3 API.K2Cis realized through our new cryptographic key-updating scheme, referred to as AB-HKU. The main advantage of the AB-HKU scheme its support for efficient delegation and revocation of privileges in hierarchies without requiring complex cryptographic data structures. We analyze the security and performance of our access control protocol, and an open source implementation. Two cryptographic libraries, Hierarchical Identity-based Encryption and Key-Policy Attribute-Based Encryption, developed in this project are useful beyond the specific cloud security problem studied.

We design and implement a scalable and privacy-preserving access control frame-work for existing untrusted cloud services. Our framework supports lazy revocation and access hierarchies. We present a signature scheme for Key-Policy Attribute-Based Encryption. Using our signature scheme, users can prove that they own a key that its policy satisfies with a set of attributes, without revealing their identity or credentials. We provide the first open source implementation of cryptographic libraries for Hierarchical Identity-Based Encryption and Key-Policy Attribute-Based Encryption schemes

We design and implement a new identity-management framework – Web2ID – for sensitive client-side mash up applications.

Web2ID supports identity authentication on the Internet without any centralized trusted party by using the mac and IP address. In addition, user's privacy (in terms of service history) is protected because identity providers and service providers do not communicate directly about the user's requests

9. System Architecture:

The proposed system has seven components as shown in the fig and they are namely ,

i. IP&MAC Authentication
ii. SSH Authentication
iii. Notifier system
iv. Dynamo DB
v. Data Transfer
vi. User Interface
vii. Aws Lambda

All these modules are called as software agent that is they act as software agent in protecting the hybrid cloud from network threats. Here the notifier system, aws lambda, dynamo db will be running in the backend ,dynamo db is used as the database to store the user account details ,but this will be called by aws lambda services when user login into the system and for the notifier system also it is same it is invoked when user login to the system. The explanation for each of them is given below,

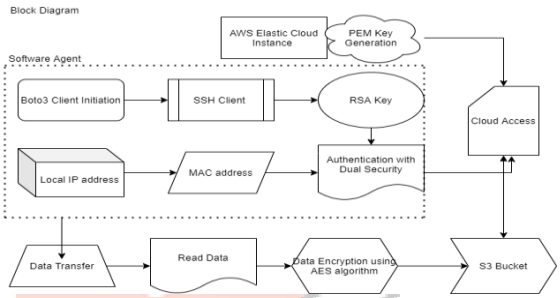

Fig.5: System Architecture of proposed system

i. Ip & mac authentication:

Here we check mac address and ip adress for authentication purpose because even if ip address us traced mac cannot be traced and cannot be recreated.

ii. Ssh authentication :

Here we check the aws credentials (perm file) authentication

iii. Notifier system :

Once the user login every transaction info will be notified through email and phone to the registered user. This is part of authentication as shown in fig5

iv. Data transfer:

Once the authentication is success through agent the files are transferred from cloud to the user and vice versa

v. Dynamo db:

The user account details ,the admin account details , cloud watch logs are maintained here. This is part of authentication as shown in fig5

vi. User interface:

The GUI we use django , since they have crs token by default that protects from cross browser forgery

vii. AWS LAMBDA:

It is a server less component that will be running in aws instances all the services will be running here from UI the agent invoke the services that are inside aws lambda

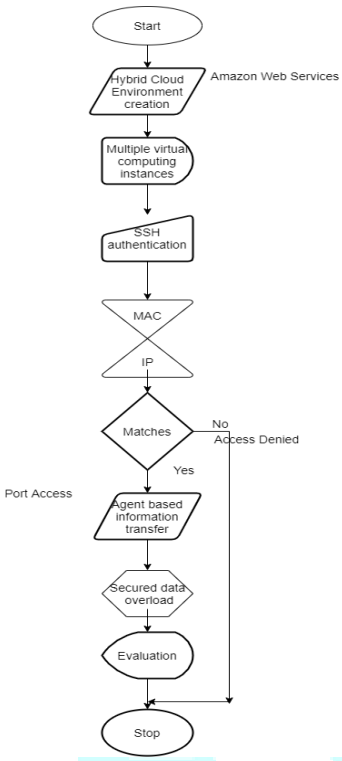

Fig.6. Flow Diagram of proposed system

10. Implementation:

The user needs to register first through offline once after registration the credentials will be stored in the dynamo db. Here the main catalyst that provides security is software agent.

When the user login to the hybrid cloud first the MAC address and IP address of the user system are checked against the credentials stored in db, which is also called stage 1 authentication as shown in Fig.5.

After the Stage1 Authentication is passed, next is stage 2 authentication. In stage 2 we check aws credentials which is also called as perm file whether it is authorized file.

After stage 2 authentications, the agent moved to the lambda trigger, Lambda triggers the handler which in turn invoked the services on the lambda, first it invokes notifier services in the lambda, that notifies about successful authentication to the user registered through phone or sms.

After the notifier services are executed, then agent provides access to the cloud. The files are transferred to the user and cloud through the software agent securely, For upload files, the agent encrypts and uploads the file to the cloud.

For download file the agent downloads the file from cloud and decrypts then stores in to the user local machine directory

Agent provide Admin privilege level as well as user privilege level login.

In the admin privilege login, the admin can monitor the number of the ec2 instances that are accessing the application. The flow of the proposed system is pictorially presented through flow diagram as shown in Fig.6.

11. Conclusion and Future Work:

Cloud computing is an emerging paradigm which its cost-effectiveness and flexibility have given it a tremendous momentum. However, there are many security challenges that, if not addressed well, may impede its fast adoption and growth. This dissertation primarily addresses the problem of sharing, managing and controlling access to sensitive resources and services in an integrated cloud environment. The primary conclusion of our research is that adoption of user-centric security models and shifting certain parts of communication and computation to the client side allows us to provide the cloud consumers with more visibility and control over their resources.

Therefore, using this approach not only the security and privacy concerns of cloud consumers can be addressed more effectively, but also the burden of managing end-users' identities and fine-granular access control will be reduced from cloud service providers

12. References:

[1]. Adam Barth, Collin Jackson, and John C. Mitchell. Robust defenses for cross-site request forgery. In Proceedings of the 15th ACM Conference on Computer and Communications Security, 2008.

[2] Adam Barth, Collin Jackson, and John C. Mitchell. Securing Browser Frame Communication. In Proceedings of the 17th USENIX Security Symposium,

2008.

[3] Mihir Bellare and Phillip Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In Proceedings of the 1st ACM conference on Computer and communications security, CCS '93, pages 62–73. ACM, 1993.

[4] John Bethen court, Amit Sahai, and Brent Waters. Cipher text-policy attribute-based encryption. In Proceedings of the 2007 IEEE Symposium on Security and Privacy, SP '07, pages 321–334, Washington, DC, USA, 2007. IEEE Computer Society.

[5] Abhilasha Bhargav-Spantzel, Anna Cinzia Squicciarini, and Elisa Bertino. Establishing and Protecting Digital Identity in Federation Systems. Journal of Computer Security, 14(3):269–300, 2006.

[6] Marina Blanton. Key Management in Hierarchical Access Control Systems, 2007.PhD Thesis, Purdue University, Aug. 2007.

[7] Marina Blanton, Nelly Fazio, and Keith B. Frikken. Dynamic and Efficient Key Management for Access Hierarchies. In
Proceedings of the ACM Conference on Computer and Communications Security, 2005
.

[8] Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In EUROCRYPT, pages 127–144. Springer-Verlag, 1998.

[9] Piero A. Bonatti and Pierangela Samarati. A Uniform Framework for Regulating Service Access and Information Release on the Web.Journal of Computer Security,10(3):241–272, 2002.

[10] Dan Boneh and Matthew Franklin. Identity-based encryption from the weil pairing.SIAM J. Comput., 32:586–615, March 2003.

[11] Jose M. Alcaraz Calero, Nigel Edwards, Johannes Kirschnick, Lawrence Wilcock, and Mike Wray. Toward a multi-tenancy authorization system for cloud services.IEEE Security and Privacy, 8:48–55, 2010.

[12] Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In ACM Computer and Communication Security 2002. ACM, 2002.

[13] Jan Camenisch, Abhi Shelat, Dieter Sommer, Simone Fischer-H¨ubner, MaritHansen, Henry Krasemann, G. Lacoste, Ronald Leenes, and Jimmy Tseng. Privacy and Identity Management for Everyone. In Proceedings of the 2005 ACM Workshop on Digital Identity Management, pages 20–27, November 2005.

[14] S. Cantor, F. Hirsch, J. Kemp, R. Philpott, E. Maler, J. Hughes, J. Hodges,P. Mishra, and J. Moreh. Security Assertion Markup Language (SAML) V2.0.Version 2.0. OASIS Standards.

[15] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. Bigtable: A distributed storage system for structured data. In Proceedings of the 7th symposium on Operating systems design and implementation - volume 7,pages 205–218, 2006.

[16] Richard Chow, Philippe Golle, Markus Jakobsson, Elaine Shi, Jessica Staddon, Ryusuke Masuoka, and Jesus Molina. Controlling Data in the Cloud: Outsourcing Computation without Outsourcing Control. In Proceedings of the 2009 ACM workshop on Cloud computing security, CCSW '09, pages 85–90, New York, NY,USA, 2009. ACM.

[17] Qian Wang, Cong Wang, Jin Li, Kui Ren, and Wenjing Lou. Enabling public veri-fiability and data dynamics for storage security in cloud computing. In Proceedings of the 14th European conference on Research in computer security, ESORICS'09,pages 355–370, Berlin, Heidelberg,2009.

[18] A. Whitten and J.D. Tygar. Why Johnny can't encrypt: a usability evaluation of PGP 5.0. In8th Usenix security symposium, pages 169–184, 1999.

[19] Huijun Xiong, Xinwen Zhang, Wei Zhu, and Danfeng Yao. CloudSeal: End-to-End Content Protection in Cloud-based storage and delivery services. In Secure comm,2011.

[20] Danfeng Yao, Keith B. Frikken, Mikhail J. Atallah, and Roberto Tamassia. Point-Based Trust: Define How Much Privacy Is Worth. In Proc. Int. Conf. on Information and Communications Security (ICICS), volume 4307 of LNCS, pages 190–209.Springer, 2006.

[21] Shucheng Yu, Cong Wang, Kui Ren, and Wenjing Lou. Achieving secure, scalable,and fine-grained data access control in cloud computing. In Proceedings of the29th conference on Information communications, INFOCOM'10, pages 534–542, Piscataway, NJ, USA, 2010. IEEE Press.

[22] Saman Zarandioon, Danfeng Yao, and Vinod Ganapathy. K2C: Cryptographic Cloud Storage With Lazy
Revocation and Anonymous Access. Technical report,Rutgers University. DCS-tr-688.

[23] Saman Zarandioon, Danfeng Yao, and Vinod Ganapathy. OMOS: A Framework for Secure Communication in Mashup Applications. InACSAC'08: Proceedings of the 24th Annual Computer Security Applications Conference, December 2008.

[24] Paul Stanton, William Yurcik, and Larry Brumbaugh. Protecting multimedia data in storage: A survey of techniques emphasizing encryption. In IS and T/SPIE International Symposium Electronic Imaging / Storage and Retrieval Methods and Applications for

Multimedia, pages 18–29, 2005.

[25] Hassan Takabi, James B.D. Joshi, and Gail-Joon Ahn. Security and Privacy Challenges in Cloud Computing Environments.IEEE Security and Privacy, 8:24–31,2010.

[26] Helen J. Wang, Xiaofeng Fan, Jon Howell, and Collin Jackson. Protection and Communication 2002. USENIX Association

Abstractions for Web Browsers in Mashup OS. In ACM Symposium on Operating Systems Principle (SOSP), pages 1–16. ACM Press, 2007.

[27] Erik Riedel, Mahesh Kallahalla, and Ram Swaminathan. A framework for evaluating storage system security. In Proceedings of the 1st USENIX Conference on File and Storage Technologies, FAST '02, Berkeley, CA, USA,