Confidentiality and Privacy in Cloud Computing using Hybrid Execution Method

B.Jaswanthi, M. NaliniSri

Abstract— Today cloud computing has become ubiquitous and we see everybody lot of data being transferred and being accessed from the cloud. At the same time this phenomenon presents us with a great risk of data theft and privacy issues .Among these privacy is the main reason that many companies and also individuals to some extent are avoiding the cloud, which also needs be addressed. For this purpose we are proposing a new model for enabling 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. This report analyses the challenges posed by cloud computing and the standardization work being done by various standards development organizations (SDOs) to mitigate privacy risks in the cloud, including the role of privacy-enhancing technologies (PETs).And a new execution model for confidentiality and privacy in cloud computing, called the Hybrid Execution model. This model provides a seamless way for an organization to utilize their own infrastructure for sensitive, private data and computation, while integrating public clouds for non-sensitive, public data and computation. We outline how to realize this model in one specific execution environment, Map Reduce over Big table.

I. INTRODUCTION

A considerable amount of cloud computing technology is already being used and developed in various flavors (e.g., private, public, internal, external, and vertical).1 Not all types of cloud computing raise the same privacy and confidentiality risks.2 Some believe that much of the computing activity occurring today entirely on computers owned and controlled locally by users will shift to "the cloud" in the future. Whether this will turn out to be the case is uncertain and not especially important here. This analysis does not support or oppose cloud computing. The continuing development and maturation of cloud computing services is an undeniable reality. The definitional borders of cloud computing are much debated today.For present purposes, *cloud computing* involves the sharing or storage by users of their own information on remote servers owned or operated by others and accessed through the Internet or other connections. Cloud computing services exist in many variations, including data storage sites, video sites, tax preparation, sites, personal health record websites, photography websites, social networking sites, and many more. Any information stored locally on a computer could be stored in a cloud, including email, word processing documents, spreadsheets, videos, health records, photographs, tax or other financial information, business plans, PowerPoint presentations, accounting information, advertising campaigns, sales numbers, appointment calendars, address books, and more. The entire contents of a user's storage device may be stored with a single cloud provider or with many cloud providers. Whenever an individual, a business, a government agency, or other entity shares information in the cloud, privacy or confidentiality questions may arise.

Cloud computing allows large organizations to tap into a virtually infinite pool of resources with the ability to control cost. Cloud providers give this power to their customers by offering a pay-as-you-go pricing model as well as elasticity and availability of computing and storage resources. Due to this benefit, cloud computing has quickly gained popularity in recent years. Yet many organizations have not widely adapted the use of clouds due to the concerns of confidentiality and privacy. For example, a recent survey collected responses from more than 500 IT executives from around the world, and reports that IT executives prefer their existing internal infrastructure (i.e., their private cloud) over a third-party cloud (i.e., a public cloud) due to security threats and lack of control over their data and systems that handle it . For industries such as finance and healthcare, explicit regulations regarding data protection — Payment Card Industry Data Security Standard and Health Insurance Portability and Accountability Act — severely limit the potential use of public clouds. These concerns are well-founded. Researchers have shown that an outside attacker can extract unauthorized information in Amazon EC2. Other researchers discovered a vulnerability that allows user impersonation in Google Apps. Encryption can only provide a limited guarantee, since any computation on encrypted data either involves decrypting the data or has yet to be practical even with fully homomorphism encryption. One line of research to resolve these issues is to make public clouds more secure. However, a public cloud is a shared platform managed by a third-party with potential security risks such as insider attacks and software vulnerabilities.

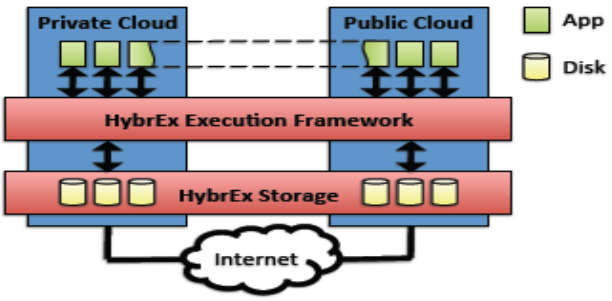

Fig. 1: The Architecture of the Hybrid Execution Model

Manuscript received April 2013.

B.Jaswanthi, Electronics& Computer Engineering, K.L.University, Vijayawada (A.P.) India.
M.Nainisri, Electronics& Computer Engineering, K.L.University, Vijayawada (A.P.) India.

These risks are difficult to eliminate as evidenced by the examples mentioned above as well as the history of security breaches and patches in general — after all, this is precisely the reason why organizations are hesitant to adapt the use of public clouds. Recognizing this difficulty, we argue for an alternative that treats public clouds as an inherently insecure environment instead of trying to make them more secure; we propose an execution model that utilizes public clouds only for safe operations while integrating an organization's private cloud. We refer to this as the Hybrid Execution model. More specifically, our Hybrid Execution model utilizes public clouds only for an organization's non-sensitive data and computation classified as public, i.e., when the organization declares that there is no privacy and confidentiality risk in exporting the data and performing computation on it using public clouds. For the organization's sensitive, private data and computation, the Hybrid Execution model utilizes their private cloud. Moreover, when an application requires access to both the private and public data, the application itself also gets partitioned and runs in both the private and public clouds. Figure 1 depicts the architecture with a Hybrid Execution framework that partitions and runs applications, as well as a Hybrid Execution storage that manages private and public data separately. The main benefit of the Hybrid Execution model is integration with safety, i.e., the ability to add more computing and storage resources from public clouds to a private cloud without the concerns for confidentiality and privacy. By partitioning data and computation, the Hybrid Execution model side-steps the question of trustworthiness of public clouds and

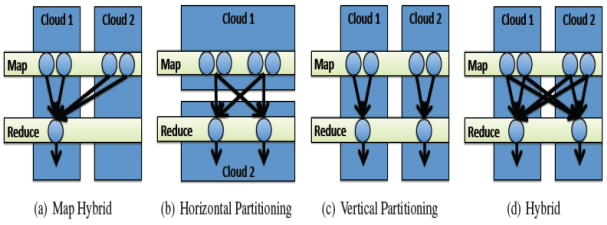

Fig. 2: Execution Categories for HybrEx MapReduce

provides the same level of confidentiality and privacy guarantees as the traditional local computing provides. Businesses are already looking into the integration of private and public clouds mainly for capacity and performance, i.e., to elastically scale out from their private cloud to public clouds [14]; the Hybrid Execution model can give an additional benefit of confidentiality and privacy to these businesses. In order to concretely explore this general direction, we first focus our effort on how to realize the Hybrid Execution model in one specific execution environment, Map-Reduce over Big table, using Hadoop Map Reduce and HBase. We have chosen this execution environment for two reasons. First, Map Reduce is arguably the most popular execution environment in cloud computing. Second, Map Reduces massively parallel nature of execution, combined with the semi-structured data management of Big table, allows clean and well-defined partitioning between private and public clouds. Using partitioning for secure computing is not a new idea. However, realizing it in the Hybrid Execution model and Map Reduce brings its own set of challenges. First, we need to be able to partition data and computation. Second, utilizing both private and public clouds for a single (partitioned) Map Reduce job means running the job over the wide-area Internet. Due to the all-to-all communication pattern in Map Reduce as well as its master-slave architecture, providing reasonable performance over the wide-area can be challenging. Finally, while partitioning provides confidentiality and privacy for private data and computation, it does not ensure integrity of public data and computation. Section 3 outlines these challenges and our research directions. The Hybrid Execution model enables new kinds of applications utilizing both private and public clouds. We discuss this in the context of Map Reduce in the next section.

II. EXECUTION CATEGORIES

Cloud computing has significant implications for the privacy of personal information as well as for the confidentiality of business and governmental information. This document identifies multiple and complex privacy

and confidentiality issues that may be of interest or concern to cloud computing participants. While storage of user data on remote servers is not a new activity, the current emphasis on and expansion of cloud computing warrants a more careful look at the privacy and confidentiality consequences. Map Reduce presents a unique opportunity to realize the Hybrid Execution model since a Map Reduce job is sub-divided into tasks that run massively in parallel. We present four categories (Figure 2) showing how Hybrid Execution Map Reduce enables new kinds of applications that utilize both private and public clouds. These categories highlight both the integration and safety aspects of the Hybrid Execution model Map. Many Map Reduce applications analyze private and public data sets; a number of public data sets are available for domains such as forensic analysis, spam detection, and genome analysis, via well-known repositories, e.g., Amazon AWS Public Data Sets and CDC WONDER. Using these data sets, organizations can analyze both their own data sets and public data sets together for comparison or accuracy improvements. Since these applications process both private and public data sets, it is difficult to execute them in a public cloud without compromising on confidentiality and privacy. Hybrid Execution Map Reduce enables this type of applications to safely utilize a public cloud by executing the Map phase in both private and public clouds while executing the Reduce phase in only one of the clouds. We refer to this category as Map hybrid, depicted in Figure 2(a). We illustrate this with an example called Cloud-Burst , a bioinformatics Map Reduce application. Simply put, it implements an algorithm that compares genome sequences, which serves as a basis for other biological analyses such as comparing a patient's genome to a reference human genome for medical purposes. Cloud-Burst's input consists of (potentially) private data called target genomes (e.g., patients' genomes) and public data called reference genomes (e.g., human reference genomes available from a public repository). Cloud-Burst uses the same algorithm to process both the reference genomes and the target genomes in the Map phase. Thus, Hybrid Execution Map Reduce can safely utilize the public cloud for the reference genomes while utilizing the private cloud for the target genomes during the Map phase.

On the other hand, the Reduce phase of a Cloud-Burst produces (potentially) private outputs, e.g. a genome comparison result for a patient. Thus, we can utilize only the private cloud for the Reduce phase. Horizontal Partitioning There are two popular usage cases for the current public clouds. The first case is long-term archiving of an organization's data, where the organization encrypts their private data before storing it on a public cloud. The second case is exporting and importing data for running periodic Map Reduce jobs in a public cloud. This is in fact a common usage case in Amazon Elastic Map Reduce; an organization first exports its data to a public cloud, runs a Map Reduce job in the public cloud, and (optionally) imports the result back to its own private storage. Since Hybrid Execution Map Reduce seamlessly integrates private and public clouds, it can automate these usage cases by utilizing

different clouds for different phases. We refer to this category as horizontal partitioning, depicted in Figure 2(b). For example, in the long-term archiving case, Hybrid Execution Map Reduce can run Map tasks that encrypt private data in the private cloud, transfer encrypted data to the public cloud via the Shuffle phase, and run Reduce tasks that store the data in the public cloud. Vertical Partitioning In 2007, the New York Times transformed its scanned images of the public domain articles to PDF files using Map Reduce running in Amazon EC2. 1 Each original article comprised of many small TIFF images, and a Map Reduce job "glued" these images together to produce one PDF file per article. However, they only transformed public domain articles; if an organization wants to process private and public documents at the same time, it becomes difficult to do so in a public cloud due to confidentiality and privacy. Hybrid Execution Map Reduce enables this type of applications to safely utilize a public cloud by executing a Map-Reduce job in both private and public clouds while avoiding any inter-cloud shuffling of intermediate data. Although this type of partitioning is technically akin to running two separate jobs in private and public clouds, Hybrid Ex Map Reduce supports this naturally without the overhead of separate management of jobs and data. We refer to this category as vertical partitioning. Figure 2(c) depicts this category. In vertical partitioning, Hybrid Execution Map Reduce workers in the public cloud execute Map and Reduce tasks using public data as the input, shuffle intermediate data among them, and store the result in the public cloud. Workers in the private cloud do the same with private data. In general, Hybrid Execution Map Reduce can run a Map Reduce job this way when the job can process private and public data in isolation. Many applications belong to this category such as separate indexing of private and public webpages of an organization, pattern search (grep), etc. Hybrid Many organizations are looking into ways to integrate public clouds for performance reasons or due to resource limitations in their private cloud. These organizations mainly utilize their private cloud even for storing and processing public data, but occasionally want to scale out to a public cloud. Hybrid Execution Map Reduce can achieve this by utilizing both private and public clouds in all three phases of Map Reduce as shown in Figure 2(d). We refer to this simply as hybrid. Beyond Pure Map Reduce For general applications, the pure Map Reduce programming paradigm is limiting as there are only two primitives to play with. Thus, we are exploring ways to overcome this limitation. The first is an optional phase in between Map and Shuffle called the Sanitize phase. Using this phase, Hybrid Execution Map Reduce could allow programmers to explicitly sanitize intermediate data for more flexible utilization of private and public clouds. We are also exploring the applicability of Hybrid Execution in Pig 2 and Hive 3. Both systems are based on Map-Reduce, but allow more flexibility and generality in programming. We believe that supporting these systems will likely result in wider applicability of Hybrid Execution.

III. RESEARCH CHALLENGES AND DIRECTIONS

The promise to deliver IT as a service is addressed a large range of consumers, from small and medium-sized enterprises (SMEs) and public administrations to end-users. According to industry analysts, the ICT sector is poised for strong growth of cloud services. Users are creating an ever-growing quantity of personal data. IDC predicts that the "digital universe" – the amount of information and content created and stored digitally – will grow from 1.8 zettabytes8 (ZB) in 2011 to over 7 ZB by 2015.

This expanding quantity of personal data will drive demand for cloud services, particularly if cloud computing delivers on the promises of lower costs for customers and the emergence of new business models for providers. Among the main privacy challenges for cloud computing are:

a) Complexity of risk assessment in a cloud environment
b) Emergence of new business models and their implications for consumer privacy
c) Achieving regulatory compliance.

3.1 Complexity of risk assessment

The complexity of cloud services introduces a number of unknown parameters. Service providers and consumers are cautious, respectively, about offering guarantees for compliance-ready services and adopting the services. With service providers promoting a simple way to flow personal data irrespective of national boundaries, a real challenge arises in terms of checking the data processing life cycle and its compliance with legal frameworks.

In a cloud service, there are many questions needing to be addressed in order to determine the risks to information privacy and security:

- Who are the stakeholders involved in the operation?
- What are their roles and responsibilities?
- Where is the data kept?
- How is the data replicated?
- What are the relevant legal rules for data processing?
- How will the service provider meet the expected level of security and privacy?

To address these issues, the Madrid Resolution states that every responsible person shall have transparent policies with regard to the processing of personal data. Stakeholders need to specify requirements for cloud computing that meet the expected level of security and privacy.

In Europe, the European Network and Information Security Agency (ENISA) provides recommendations to facilitate understanding of the shift in the balance of responsibility and accountability for key functions such as governance and control over data and IT operations and compliance with laws and regulations.

We need to overcome three main challenges to implement the Hybrid Execution model — data partitioning, system partitioning, and integrity — and achieve integration with safety. The following outlines these challenges and our research directions.

3.2 Data Partitioning

Since the Hybrid Execution model proposes the use of partitioning for confidentiality and privacy, it raises an immediate question of how to partition data. We address this question by using two labels, the private label and the public label. Hybrid Execution Big table and Hybrid Execution Map Reduce recognize these labels and determine data and computation placement accordingly. This use of labels is inspired by previous information flow control techniques used in programming languages such as Jif and systems such as Asbestos . While these approaches use labels to control and track how information flows among system components, we only need to use labels to determine the placement of data and computation. We assume that organizations have policies to determine data sensitivity; we hypothesize that in many cases, organizations can perform labeling programmatically at the time of importing data into Hybrid Execution Big table by running a Map Reduce job. For example, if an organization has a file-naming convention that indicates the sensitivity of a document, it can run a Map Reduce job that labels data according to file names.

3.3 System Partitioning

Since the Hybrid Execution model utilizes both private and public clouds, any system that implements the Hybrid Execution mode l has to partition its components, i.e., it needs to place some components in the public cloud and others in the private cloud. This naturally raises two sub-questions —

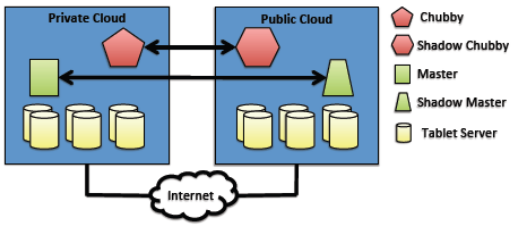

Fig3: The Architecture of Hybrid Ex Big table

i)How to keep public components from accessing unauthorized information (e.g., private data), and ii)how to reduce the wide-area communication overhead if the communications between private and public components are necessary and become a bottleneck. We discuss how we are addressing these issues in Hybrid Execution Big table and Hybrid Execution Map Reduce. Hybrid Execution Big table In order to concretely discuss the two questions of system partitioning for Big table, we briefly introduce its three main components—the master, tablet servers, and Chubby. The master is responsible for the overall management of the system. Tablet servers manage stored data and handle read and write requests. Chubby provides meta data consistency and reliability.

With these components, we can rephrase the two questions of system partitioning as follows. First, when utilizing both private and public clouds, we need to place tablet servers in both clouds as they directly handle data. Thus, we must keep public tablet servers from handling private data. Second, we avoid placing the master and Chubby in the public cloud as we need to trust them for overall correctness. However, the master, tablet servers, and Chubby communicate with each other frequently for correct functioning of the system. Thus, we need to reduce the wide-area communication overhead between public tablet servers, the master, and Chubby. In order to address these issues, we introduce two new components that we refer to as the shadow master and the shadow Chubby as shown in Figure 3. These components are public counterparts to the master and Chubby that run in the public cloud. Each shadow component is a restricted version of its private counterpart, and does not have any access to information regarding private data. Moreover, we only allow tablet servers in the public cloud to communicate with these shadow components. An immediate benefit of this architecture is the ability to avoid the master-slave wide-area communications, as it mostly localizes the communications among system components; the components in the private cloud do not communicate with the components in the public cloud for the most part, and vice versa. Hybrid Execution Map Reduce Since both Big table and Map-Reduce have a similar master-slave architecture, we apply the same

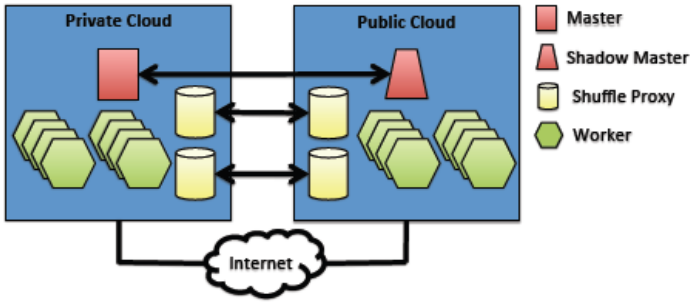

Fig 4:The Architecture of Hybrid Ex Model

general approach to system partitioning we partition the main components of the original Map-Reduce, the master and the workers. The master has its restricted public counterpart, the shadow master, and we place the workers in both private and public clouds as shown . As in Hybrid Execution Big table, this architecture avoids the master-slave wide-area communications. However, Map Reduce has a critical difference from Big table in terms of wide-area communication overhead as some Map Reduce jobs require shuffling of intermediate data over the wide-area as shown in Section 2. Thus, we introduce a new component called shuffle proxies that transfer intermediate data over the wide-area on behalf of the workers. This is different from the original Map Reduce that allows the workers to directly transfer intermediate data among them. Shuffle proxies give us the benefit of having a separate architectural component where we can apply optimization techniques to reduce the wide-area overhead.

We are exploring techniques such as caching, aggregation, compression, and de-duplication of intermediate data.

We recognize that using a public cloud in addition to a private cloud gives more computing power, which can lead to better performance despite the wide-area overhead. For example, in a small-scale experiment involving 5 local machines in Buffalo and 5 Emu lab machines in Utah, the execution time of Hadoop sorting 20GB of input data turns out to be faster over the wide-area (702sec with all 10 machines over the wide-area vs. 937 sec with 5 local machines). However, overhead reduction is still important to benefit a broad range of applications and system configurations.

3.4 Integrity

The last question is how to provide integrity for data and computation in the public cloud, as our basic assumption is that it is not safe to trust the public cloud. To address the question of data integrity, Hybrid Execution Big table keeps the hashes of the public data in the private cloud. Hybrid Execution Big table can verify the integrity either when there is a request for the public data or proactively by sampling. For computation integrity, Hybrid Execution Map Reduce checks the integrity of the results from the public cloud in two modes that provide different levels of fidelity.The first mode is full integrity checking, where the private cloud re-executes every Map and Reduce task that the public cloud has executed. Hybrid Execution Map Reduce provides this mainly as a means to enable auditing at a later time, e.g., when an organization wants to verify the correctness of past computations from the public cloud. Obviously, the overhead of doing the full integrity checking can be costly. Thus, Hybrid Execution Map Reduce provides quick integrity checking, where the private cloud selectively checks the integrity of the results from the public cloud. Hybrid Execution Map Reduce provides this mainly to check the integrity at runtime for probabilistic detection of suspicious activities in the public cloud. For this purpose, we store data items that we call inspection points in the private cloud. Inspection points can be either new synthetic data items that we add to the public data or existing public data items selected randomly for the purpose of verification. For example, for a Map-Reduce job that counts words in a public document, we can either add new unique words to the document or select existing words at random from the document, and store them in the private cloud. We can verify that the result from the public cloud contains the accurate counts of these words by running the same job in the private cloud with the inspection points. The frequency, overhead, and effectiveness of inspection points are our current subjects of investigation.

IV. PRIVACY BY DESIGN

Privacy is an essential human right, enshrined in the Universal Declaration of Human Rights and International Covenant of Political and Civil Rights9. Article 12 of the Universal Declaration of Human Rights states that “No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honor and reputation. Everyone has the right to the protection of the law against such interference or attacks.” In Europe, the Charter of Fundamental Rights of the European Union (2000) became legally binding in European Union law as part of the Lisbon Treaty (in force since December 2009). EU Directive 95/46/EC, and e-privacy and electronic communications Directive 2002/58/EC covering also data retention, are the main legal instruments in Europe covering privacy and the processing of personal data. The recent Madrid Resolution provides international standards for the protection of privacy, but there is as yet no universally binding privacy legislation covering all the countries in the world. In a cloud computing service, privacy becomes more complex. Applying legal frameworks to the cloud is not easy when regimes are not harmonized, depend on the location of data and involve blurred division of responsibilities between stakeholders. In Europe, the 27 Member States have implemented the 1995 EU Directive differently, resulting in difficulties in enforcement. According to the European Commission, a single EU law can do away with the current fragmentation and costly administrative burdens, which could save businesses some €2.3 billion a year10. A recent ENISA report summarizes a number of rules and challenges associated with Directive 95/46/EC in the context of "the cloud computing environment, for which the roles of controller and processor still need to be determined on a case-by-case basis and in relation to the nature of the cloud services". The United States does not have an overarching governmental regulation as is the case in Europe, but follows a sectorial approach with privacy and data protection needs being addressed through a plethora of regulations and laws, including self-regulation. A number of privacy principles are also to be found in other organizations and countries11. The privacy legislations in other countries are also important. Countries in the developing world (e.g on the African continent, India and China) are currently planning the introduction of privacy and data protection laws. The differences in privacy legislation globally can become a trade barrier and prevent innovation.

V. RELATED WORK

In addition to the previous work on cloud security, partitioning, and information flow control discussed in Section 1 and 3, the recent line of work on hybrid clouds, cloud accountability, security and privacy in Map Reduce, and un trusted public clouds is closely related. Notably, Cloud Net proposes a VPN-like network to provide secure and seamless resource integration between private and public clouds. The work on Accountable virtual machines proposes an account ability scheme for virtualized Environments. Since the AVM approach is basically VM logging-and replaying, it is effectively the same as our full integrity checking, potentially with more overhead. Air vat provides security and privacy guarantees in Map Reduce by mandatory access control and differential privacy. Recent works such as Depot and SPORC also assume public clouds as an untrusted environment.

VI. CONCLUSIONS AND FUTURE WORK

Cloud computing is still in its infancy. This is an emerging technology which will bring about innovations in terms of business models and applications.

The widespread penetration of smartphones will be a major factor in driving the adoption of cloud computing. However, cloud computing faces challenges related to privacy and security.

In cloud services, the implementation of PETs will depend on the availability of standards to assess privacy risks and describe means of ensuring data protection compliance. PETs can ensure that breaches of the data protection rules and violations of individuals' rights are not only forbidden and subject to sanctions, but are also a technically daunting undertaking. The embedding of privacy by design features when designing technologies is increasingly supported by regulators and is also being included in the reform of the EU Data Protection Directive.

We are currently implementing Hybrid Execution Map Reduce and Hybrid Execution Big table. Moving forward, we believe that the Hybrid Execution model could be useful for other execution environments such as VM-based hybrid clouds and smart phones interacting with clouds. We plan to explore the feasibility of applying the Hybrid Execution model in these environments.

Cybercriminal activities impacting cloud computing environments − for example, fraud and malicious hacking − are threats that can undermine user confidence in the cloud. Cloud computing providers face multiple, and potentially conflicting, laws concerning disclosure of information. Achieving a better understanding of jurisdictional issues is critical and should be tackled through enhanced dialogue.

1) In this paper, we start from the position that it is fundamentally difficult to secure public clouds, and then outline an execution model called the Hybrid Execution model that uses partitioning of data and computation as a way to provide confidentiality and privacy. We discuss how we can realize this model in one specific execution environment.

REFERENCES

1. A. Armando et al. Formal Analysis of SAML 2.0 Web Browser Single Sign-On: Breaking the SAML-based Single Sign-On for Google Apps. In ACM FMSE, 2008.
2. F. Chang et al. Big table: A Distributed Storage System for Structured Data. In USENIXOSDI, 2006.
3. S. Chong et al. Secure Web Applications via Automatic Partitioning. In ACM SOSP, 2007.
4. J. Dean and S. Ghemawat. Map Reduce: Simplified Data Processing on Large Clusters. In USENIXOSDI, 2004.
5. P. Efstathopoulos et al. Labels and Event Processes in the Asbestos Operating System. In ACMSOSP, 2005.
6. A. J. Feldman et al. SPORC: Group Collaboration using Untreated Cloud Resources. In USENIXOSDI, 2010.
7. C. Gentry. Fully Homomorphism Encryption using ideal lattices. In ACM STOC, 2009.
8. A.Haeberlenetal. Accountable Virtual Machines. In USENIXOSDI, 2010.
9. Health Insurance Portability and Accountability Act of 1996(HIPAA), Public Law 104-191.
10. P. Mahajan et al. Depot: Cloud Storage with Minimal Trust. In USENIX OSDI,2010.
11. A.C. Myers and B. Liskov. A Decentralized Model for information Flow Control. In ACM SOSP, 1997.
12. Survey: Cloud Computing 'No Hype', But Fear of Security and Control Slowing Adoption. http://www.circleid.com/posts/20090226_cloud_computing_hype_security.
13. PCIDSSv2.0.https://www.pcisecuritystandards.org/documents/pci_dss_v2.pdf, 2010.
14. Forecast for 2010: The Rise of Hybrid Clouds. http://gigaom.com/2010/01/01/on-the-rise-of-hybrid-clouds, 2010.
15. T. Risten part et al. Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds. In ACMCCS, 2009.
16. I. Roy et al. Air vat Security and Privacy for Map Reduce. In USENIX NSDI, 2010.
17. N. Santos et al. Towards Trusted Cloud Computing. In USENIX Hot Cloud, 2009.
18. M. C. Schatz. Cloud Burst: Highly Sensitive Read Mapping with Map Reduce. Bioinformatics, 25(11):1363–1369, 2009.
19. P.Sirota. Keynote: Making aHdoop Enterprise Ready with Amazon Elastic Map Reduce. Hadoop Summit, 2010.
20. T. Wood et al. The Case for Enterprise-Ready Virtual Private Clouds. In USENIX Hot Cloud, 2009.
21. S. Zdancewic et al. Untrusted Hosts and Confidentiality: Secure Program Partitioning. In ACMSOSP, 2001.
22. L. Zheng et al. Using Replication and Partitioning to Build Secure Distributed Systems. In IEEE Oakland, 2003.