Report from Dagstuhl Seminar 11492

Secure Architectures in the Cloud

Edited by

Sabrina De Capitani di Vimercati¹, Wolter Pieters², and Christian W. Probst³

- 1 Università degli Studi di Milano, IT, sabrina.decapitani@unimi.it
- 2 Delft University of Technology, NL, w.pieters@tudelft.nl*
- 3 Technical University of Denmark, DK, probst@imm.dtu.dk

— Abstract

This report documents the outcomes of Dagstuhl Seminar 11492 "Secure Architectures in the Cloud". In cloud computing, data storage and processing are offered as services, and data are managed by external providers that reside outside the control of the data owner. The use of such services reduces the burden of the owners in managing their data, and may provide significant cost savings. However, cloud computing introduces new security and privacy concerns. In fact, there is little consensus on how to guarantee the confidentiality, integrity, and availability of data in cloud computing scenarios. Also, it is unclear to what extent parties can be held accountable in case something goes wrong. In this seminar, we searched for architectures, modelling approaches, and mechanisms that can help in providing guarantees for cloud security. We proposed the concept of verification-as-a-service that can guide architectures for verification of cloud architectures and configurations, as well as results of computations. We also proposed architectures for organising customisability of security and privacy for cloud customers.

Seminar 4.–9. December, 2011 – www.dagstuhl.de/11492

1998 ACM Subject Classification K.6.5 Security and Protection

Keywords and phrases attestation, auditing, cloud computing, security architectures, security modelling, verification

Digital Object Identifier 10.4230/DagRep.1.12.11

1 Executive Summary

Wolter Pieters Christian W. Probst Sabrina de Capitani di Vimercati

License <a>
 (c) Creative Commons BY-NC-ND 3.0 Unported license

 Wolter Pieters, Christian W. Probst and Sabrina de Capitani di Vimercati

Introduction

In cloud computing, data storage and processing are offered as a service, and the data resides outside the control of the owner. It is often argued that clouds improve security, as the providers have more security expertise than their (smaller) customers. However, despite

Except where otherwise noted, content of this report is licensed under a Creative Commons BY-NC-ND 3.0 Unported license Secure Architectures in the Cloud, *Dagstuhl Reports*, Vol. 1, Issue 12, pp. 11–26 Editors: Sabrina De Capitani di Vimercati, Wolter Pieters, and Christian W. Probst

^{*} At the time of the seminar, the second editor was affiliated with the University of Twente. His research was supported by the research program Sentinels (http://www.sentinels.nl). Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs.

NAGSTUHL Dagstuhl Reports REPORTS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

12 11492 – Secure Architectures in the Cloud

theoretical breakthroughs in cryptography, there is little consensus on how we can provide architectural solutions guaranteeing that cloud data remains confidential, uncorrupted, and available. Also, it is unclear to what extent parties can be held accountable in case something goes wrong. In seminar 11492 Secure Architectures in the Cloud, we searched for architectures, modelling approaches, and mechanisms that can help in providing guarantees for cloud security. The main question was which cloud-specific security architectures should and could be devised, and how they can be matched to security policies. The seminar was attended by researchers from different academic and industry communities, making it possible to propose integrated solutions and research directions that transcend disciplines. Four main topics have been the subject of this seminar (see also [14]):

- 1. Data protection. Data outside the data owner's control implies that privacy and even integrity can be put at risk. Guaranteeing the privacy and integrity of the data, whether stored in the system or communicated to external parties, becomes a primary requirement, and has raised the attention of both individuals and legislators. Cloud providers have to properly protect the privacy of (possible sensitive) information when storing, processing or sharing it with others [19], and have to adopt adequate access control solutions for enforcing selective access to the data. New approaches have emerged for identifying persons and roles and linking them to access privileges, such as identity-, attribute-, claims-, and data-based access control (e.g., [7, 15]). We discussed challenges of the cloud to the notions of privacy, accountability and user empowerment, their legal, ethical, and architectural implications, and possible solutions.
- 2. Simulating physical constraints in the cloud. In the cloud, we cannot easily enforce where data is stored and how long, and from where it is accessed. Location-based access control aims at limiting access to specific locations, thereby seemingly putting physical limitations back in place [23]. Measures proposed include use of GPS, trusted platform modules (TPMs), but also physically unclonable functions (PUFs) [21]. Also, data could be moved away from attacks [17]. With respect to time, mechanisms have been proposed to assure deletion of data in the cloud [9, 12, 22]. We assessed to which extent these approaches are sufficient to simulate physical constraints, and which architectural solutions are needed to make such forms of assurance possible in practice.
- 3. *Misuse detection*. Many methods have been proposed for intrusion detection, penetration testing and digital forensics. Are these sufficient for cloud environments? The seminar identified necessary adaptations to system and threat models as well as security metrics, to adequately indicate which attacks are possible and which are actually happening, and thereby reduce cybercrime.
- 4. Splitting the clouds. Public clouds, containing data from different parties, are not deemed suitable for particularly sensitive information. This means that decisions will have to be made about which data to put in the cloud and which data not, which security properties to outsource and which not, and how to make sure that the entire system conforms to the security requirements (cf. [4]). The seminar investigated suitable architectures for "splitting the clouds". For example, in "security-as-a-service", not only IT infrastructure is rented, but also the security that is added to it. For authentication this seems to work pretty well, but how far can this concept be stretched to other security properties such as confidentiality and integrity?

Processing encrypted data was discussed in the parallel seminar 11491 Secure Computing in the Cloud. This report covers the results of the seminar on Secure Architectures in the Cloud, abstracts of presentations, and proceedings of the working groups. The topics have been restructured during the seminar, and we will refer back to the topics originally proposed where appropriate. Several follow-up initiatives have been assigned to the participants.

Main Findings

As a general observation, we concluded that clouds require a different kind of architectural decisions than traditional information systems. In complex systems such as clouds, we cannot do lots of things manually anymore. For example, there is usually no way to inspect a cloud for evidence manually after an incident. This means that the architecture needs to allow for automation of such tasks, by providing not only functional services, but also meta-services to perform automated maintenance, recovery, etc. Moreover, the processes that make use of such meta-services need suitable architectures themselves. In particular, the following meta-services are needed:

- Automated policy checking,
- Automated configuration verification,
- Automated incident management,
- Automated auditing, and
- Automated forensics.

These processes could be deployed again in (different) clouds, but then the same security concerns apply to them as well.

In this sense, the cloud paradigm begs the question whether we can do everything as a service (XaaS). The participants came up with many different XaaS concepts. In particular, we proposed the concept of verification-as-a-service, which can refer to both the verification of the results of computations, as well as the verification of the (security) architecture and configuration in place at the cloud provider. The former is well-known in the field of electronic voting systems (cf. [20]); the latter resonates with the practice of security auditing. Verification-as-a-service is the main focus in relation to original topics 3 and 4. Specific challenges relate to the verification of negative properties (something is not the case in the architecture) and verification of the results of randomised algorithms. Also, testing-as-a-service could be employed to test functional and non-functional properties of cloud services.

As an instance of simulating physical constraints (topic 2) in relation to verification-as-aservice, we discussed the possibilities for verifying the location of data in the cloud (cf. [23]). One proposal is to integrate GPS with trusted hardware (such as TPM) to prove locations.

Verification-as-a-service provides a paradigm to organise accountability in the cloud. This could be realised by different techniques, for example by:

- Transparency of architecture/configuration (inspection/attestation),
- Forensics (e.g., watermarking),
- Regulation (precaution) and enforcement,
- Incident response (logging), or
- Creating incentives.

Verifying the *integrity* of data seems to be more intuitive than verifying its *confidentiality*. With integrity, it is possible, for example, to compare two different copies. With confidentiality, one would have to prove that only certain parties possess a copy. It only seems to be possible to falsify this after the fact, when it is indeed discovered that data has been leaked. Even in that case, one would need some kind of watermark to prove who leaked the data, for it might have been the user as well as the provider. How to develop a service that provides such watermarking in relation to confidentiality-as-a-service has been identified as an open problem, in relation to topic 4 (splitting the clouds).

Especially on the user side, accountability can be further enhanced by modifiability, or customisability, which allows the user to adapt services to his or her own policies. This requires negotiation on policies, not only between the user and the initial provider, but 13

14 11492 – Secure Architectures in the Cloud

also between providers within the supply chain (cf. [24]). Customisability is the main topic discussed in relation to original topic 1 (data protection). Again, special services can be set up that allow the user to achieve this for multiple cloud services at the same time, which would amount to modifiability-as-a-service. Such services could be standardised to make sure that they really empower the user, by employing certain privacy policies themselves, and providing an understandable interface (cf. [6, 10]). We would then have achieved "standardised customisability".

We formulated several attacker models that lie behind these proposals. Many standard attacker models are problematic in the cloud. An evil/malicious cloud service provider implies that we cannot solve anything without advanced encryption methods, which are costly or even infeasible in many scenarios. Assuming that computations are performed in the clear, we have to assume that the cloud service provider is *indifferent*, not curious. Thus, we trust the cloud provider on the issue of confidentiality, in the sense that we do not expect the provider to leak or misuse data intentionally. However, the provider may still be a:

- Sloppy provider (makes mistakes),
- Lazy provider (simplifies computations), or
- Greedy provider (reduces security to save money).

The sloppy and lazy provider might compromise the integrity of the result of computations. Verification of results would be a countermeasure here, for example by executing the computations on multiple, independent clouds.

Greedy providers are willing to violate policies for economic reasons, thereby exposing the data to insider or outsider threats. Although we do not assume malice on the side of the provider, we do assume malice on the side of other cloud users, who may or may not have specialised access (e.g., administrators). In relation to the greedy provider, one would want to have some means to verify the architecture in place.

Especially if services have been customised, one would want to have some kind of assurance that there is actually a change in configuration taking place based on the customisation. We proposed the development of a tool suite to support remote measurements of architectural variables, which would include existing proposals (cf. [2, 18, 25]). Care needs to be taken that acquiring such information does not violate customer privacy or company property rights [3]. Also, even if the architecture would be (partly) known, the user would then need meaningful support to choose among different providers (and thereby different architectures). This provides another incentive to develop quantitative models that can indeed calculate overall security risks from system architectures, based on existing qualitative approaches [1, 8, 11, 13, 16]. The user can then compare risks and costs to make decisions [5]. Such decisions could even be made in real-time based on information on the current security situation, leading to what has been called fluid information systems [17].

A remaining question is how to create incentives to invest in cloud security. If there is no immediate impact, investments may lag behind with respect to threat levels. Ironically, you can gain a competitive advantage by making your competitors invest in security. Do we really need big scandals to improve security? In any case, achieving more security by (self-)regulation, whether by law, seals, or otherwise, requires architectures such as proposed here, for it is impossible to impose constraints if they cannot be verified.

In conclusion, this seminar proposed architectures for verifying the results of cloud computations, verifying the configuration of cloud architectures, and supporting customisability of cloud services in terms of security. These were defined in relation to cloud-specific attacker models. Visual representations of the proposed architectures can be found under the results of the working groups. Open problems are defined at the end of this document.

References

- S. Bleikertz, T. Groß, and S. Mödersheim. Automated verification of virtualized infrastructures. In CCSW'10: Proceedings of the 2010 ACM Cloud Computing Security Workshop, pp. 47–58. ACM, 2010.
- 2 S. Bleikertz, M. Schunter, C.W. Probst, D. Pendarakis, and K. Eriksson. Security audits of multi-tier virtual infrastructures in public infrastructure clouds. In CCSW'10: Proceedings of the 2010 ACM Cloud Computing Security Workshop, pp. 93–102. ACM, 2010.
- **3** T.D. Breaux and C.B. Lotrionte. Towards a privacy management framework for distributed cybersecurity in the new data ecology. In *HST'11: Proceedings of the IEEE International Conference on Technologies for Homeland Security*, pp. 6–12. IEEE, November 2011.
- 4 S. Bugiel, S. Nürnberger, A.-R. Sadeghi, and T. Schneider. Twin clouds: Secure cloud computing with low latency. In B. De Decker, J. Lapon, V. Naessens, and A. Uhl, editors, *Communications and Multimedia Security*, vol. 7025 of *Lecture Notes in Computer Science*, pp. 32–44. Springer Berlin / Heidelberg, 2011. DOI: 10.1007/978-3-642-24712-5_3.
- 5 Y. Chen and R. Sion. To cloud or not to cloud?: musings on costs and viability. In Proceedings of the 2nd ACM Symposium on Cloud Computing, SOCC'11, pp. 29:1–29:7, New York, NY, USA, 2011. ACM.
- 6 L. Coles-Kemp and E. Kani-Zabihi. On-line privacy and consent: a dialogue, not a monologue. In *Proceedings of the 2010 workshop on New security paradigms*, NSPW'10, pp. 95–106, New York, NY, USA, 2010. ACM.
- 7 S. De Capitani di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, G. Pelosi, and P. Samarati. Encryption-based policy enforcement for cloud storage. In *ICDCSW'10: IEEE 30th International Conf. on Distributed Computing Systems Workshops*, pp. 42–51. IEEE, 2010.
- 8 T. Dimkov, W. Pieters, and P.H. Hartel. Portunes: representing attack scenarios spanning through the physical, digital and social domain. In Proceedings of the Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security (ARSPA-WITS'10). Revised Selected Papers, Paphos, Cyprus, vol. 6186 of LNCS, pp. 112– 129, Berlin, March 2010. Springer Verlag.
- 9 R. Geambasu, T. Kohno, A.A. Levy, and H.M. Levy. Vanish: increasing data privacy with self-destructing data. In *Proceedings of the 18th conference on USENIX security symposium*, SSYM'09, pp. 299–316, Berkeley, CA, USA, 2009. USENIX Association.
- 10 L. Jedrzejczyk, B.A. Price, A. Bandara, and B. Nuseibeh. "privacy-shake": a haptic interface for managing privacy settings in mobile location sharing applications. In *MobileHCI'10: Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services*, pp. 411–412, New York, NY, USA, September 2010. ACM.
- 11 B. Kordy, S. Mauw, S. Radomirović, and P. Schweitzer. Foundations of attack-defense trees. In Formal Aspects of Security and Trust, 7th International Workshop, FAST 2010, vol. 6561 of LNCS, pp. 80–95. Springer, 2011.
- 12 R. Perlman. The ephemerizer: making data disappear. Technical Report SMLI TR-2005-140, Sun Microsystems, Inc., Mountain View, CA, USA, 2005.
- 13 W. Pieters. Representing humans in system security models: An actor-network approach. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, 2(1):75–92, 2011.
- 14 W. Pieters. Security and privacy in the clouds: a bird's eye view. In S. Gutwirth, Y. Poullet, P. De Hert, and R. Leenes, editors, *Computers, Privacy and Data Protection: an Element of Choice*, pp. 445–457. Springer, Dordrecht, 2011.
- 15 W. Pieters and Q. Tang. Data is key: introducing the data-based access control paradigm. In E. Gudes and J. Vaidya, editors, *Data and Applications Security 2009, Montreal, Canada*, vol. 5645 of *Lecture Notes in Computer Science*, pp. 240–251. Springer, 2009.

16 11492 – Secure Architectures in the Cloud

- 16 C.W. Probst and R.R. Hansen. An extensible analysable system model. *Information* security technical report, 13(4):235–246, 2008.
- 17 C.W. Probst and R.R. Hansen. Fluid information systems. In Proceedings of the 2009 workshop on New security paradigms workshop, pp. 125–132. ACM, 2009.
- 18 C.W. Probst, M.A. Sasse, W. Pieters, T. Dimkov, E. Luysterborg, and M. Arnaud. Privacy penetration testing: How to establish trust in your cloud provider. In S. Gutwirth, R. Leenes, P. De Hert, and Y. Poullet, editors, *European Data Protection: In Good Health?*. Springer, Dordrecht, 2012.
- 19 J. Ruiter and M. Warnier. Privacy regulations for cloud computing: Compliance and implementation in theory and practice. In S. Gutwirth, Y. Poullet, P. De Hert, and R. Leenes, editors, *Computers, Privacy and Data Protection: an Element of Choice*, pp. 361– 376. Springer Netherlands, 2011. DOI: 10.1007/978-94-007-0641-5_17.
- 20 P.Y.A. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia. Prêt à voter: a voterverifiable voting system. *IEEE Transactions on Information Forensics and Security*, 4:662– 673, 2009.
- 21 G.E. Suh and S. Devadas. Physical unclonable functions for device authentication and secret key generation. In *Proceedings of the 44th annual Design Automation Conference*, DAC'07, pp. 9–14, New York, NY, USA. ACM, 2007.
- 22 Q. Tang. Timed-ephemerizer: Make assured data appear and disappear. In F. Martinelli and B. Preneel, editors, *Public Key Infrastructures, Services and Applications*, vol. 6391 of *Lecture Notes in Computer Science*, pp. 195–208. Springer Berlin / Heidelberg, 2010. DOI: 10.1007/978-3-642-16441-5_13.
- 23 A. van Cleeff, W. Pieters, and R. J. Wieringa. Benefits of location-based access control:a literature study. In Proceedings of the 3rd IEEE/ACM International Conference on Cyber, Physical and Social Computing (CPSCom 2010), Hangzhou, China, pp. 739–746, Los Alamitos, CA, November 2010. IEEE Computer Society.
- 24 M. Winslett, T. Yu, K.E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. Negotiating trust in the web. *IEEE Internet Computing*, 6(6):30–37, 2002.
- 25 W. Xu, X. Zhang, H.n Hu, G.-J. Ahn, and J.-P. Seifert. Remote attestation with domainbased integrity model and policy analysis. *IEEE Transactions on Dependable and Secure Computing*, 99(PrePrints), 2011.

2 Table of Contents

Executive Summary	
Wolter Pieters, Christian W. Probst and Sabrina de Capitani di Vimercati	11
Overview of Talks	
Adaptive Information Security for Cloud Services: Relating Security Requirements to Design Arosha Bandara 1	18
Security Assurance in Virtualized Infrastructures Sören Bleikertz	18
Privacy and Security Requirements in the Cloud Travis D. Breaux	19
Versatile Key Management for Secure Cloud Storage Sebastian Graf 1	19
Energy Efficiency in Cloud and Related Security Issues Toni Mastelic 2	20
Privacy and Integrity Issues in Data Mining Outsourcing Anna Monreale 2	20
To Cloud or Not To. Simple Musings on Cloud Viability Radu Sion	21
On Securing Untrusted Clouds with Cryptography Radu Sion	21
BOTCLOUDS – The Future of Cloud-based Botnets? Martijn Warnier	22
Working Groups	
Privacy, Data Protection and User Empowerment	22
Verifying Configurations	22
Verifying Computations	23
Open Problems	24
Panel Discussions	25
Participants	26

3 Overview of Talks

3.1 Adaptive Information Security for Cloud Services: Relating Security Requirements to Design

Arosha Bandara (The Open University – Milton Keynes, GB)

License 🐵 🏵 Creative Commons BY-NC-ND 3.0 Unported license © Arosha Bandara Joint work of Bandara, Arosha; Yu, Yijun; Tun, Than Thien; Nuseibeh, Bashar

Information security involves protecting valuable information assets from possible harm. With the increasing use of cloud computing services, the technical and social contexts in which software applications are expected to operate become increasingly dynamic. As a result, the assets, their values, and attack scenarios can easily change. This increases the challenge of finding out what the information assets are, who their owners are, where in the system vulnerabilities lie, and the extent to which the security requirements need to be enforced. In such an environment, information security has to be highly context-sensitive: software applications must adapt to the changing contexts and respond quickly and appropriately to ensure that the requirements for information security are not violated. We call this notion Adaptive Information Security, and focus on three of its prerequisites in the context of cloud computing: (1) understanding user requirements for cloud applications; (2) traceability between requirements, design and implementation of cloud services; and (3) adaptive design for dynamic contexts.

3.2 Security Assurance in Virtualized Infrastructures

Sören Bleikertz (IBM Research – Zürich, CH)

Cloud computing and virtualized infrastructures are often accompanied by complex configurations and topologies. Dynamic scaling, rapid virtual machine deployment, and open multi-tenant architectures create an environment, in which local misconfiguration can create subtle security risks for the entire infrastructure. This situation calls for automated deployment as well as analysis mechanisms.

We present a platform that combines a static information flow analysis and a virtualization assurance language with state-of-the art verification methods. The system discovers the actual configurations of diverse virtualization environments and unifies them in a graph representation. Using graph traversal, it computes the transitive closure of information flow. The language integrates descriptions of virtualized infrastructures, their transformations, their desired security goals, and evaluation strategies. The different verification tools range from model checking to theorem proving; this allows us to exploit the complementary strengths of methods.

We demonstrate the feasibility of our approach by a real-world case study of a virtualized infrastructure of a global financial institution.

3.3 Privacy and Security Requirements in the Cloud

Travis D. Breaux (Carnegie Mellon University – Pittsburgh, US)

Cloud computing enables organizations to cheaply and quickly obtain computer resources on an as-needed basis, allowing them to more efficiently and effectively provide services to their consumers. Despite the cloud's ubiquitous appearance, however, data provenance within the cloud presents a multi-jurisdictional challenge, as privacy laws and regulations that govern data may be applicable based upon the residence of the entity who owns the data, the type of organization that uses the data, and any intermediaries facilitating the handling of the data, such as cloud service providers. To this end, we are developing a modeling framework for determining jurisdictional applicability in which these entities are respectively designated as data subjects, data users, and data custodians. We foresee organizations using this framework during system design to determine and resolve the complex issues about where to provide services, store and transfer data.

3.4 Versatile Key Management for Secure Cloud Storage

Sebastian Graf (Universität Konstanz, DE)

License (© (©) Creative Commons BY-NC-ND 3.0 Unported license © Sebastian Graf

Storing data on cloud-based infrastructures facilitates infinite scalability and all-time availability. Putting data in the cloud additionally offers a convenient way to share any information with user-defined third-parties. However, storing data on the infrastructure of commercial third party providers, demands trust and confidence. Often simple approaches, like merely encrypting the data by providing encryption keys, which at most consists of a shared secret supporting rudimentary data sharing, do not support evolving sets of accessing clients to common data.

Based on well-established approaches regarding stream-encryption, we propose an adaption for enabling scalable and flexible key management within heterogeneous environments like cloud scenarios. Representing access-rights as a graph, we distinguish between the keys used for encrypting hierarchical data and the encrypted updates on the keys enabling flexible join-/leave-operations of clients. This distinction allows us to utilize the high availability of the cloud as updating mechanism without harming any confidentiality. Our graph-based key management results in a constant adaption of nodes related to the changed key. The updates on the keys generate a constant overhead related to the number of these updated nodes.

The proposed scalable approach utilizes cloud-based infrastructures for confidential data and key sharing in collaborative workflows supporting variable client-sets.

3.5 Energy Efficiency in Cloud and Related Security Issues

Toni Mastelic (TU Wien, AT)

Cloud computing is a promising approach for implementing scalable on- demand computing infrastructure. It includes business aspects like SLAs and customer-provider relationship, as well as organizational issues like scheduling, resource allocation, all the way to a technical details like VM monitoring and application deployment. While energy efficiency is mostly managed on an organizational level, it is realized by actions on the level of clusters, physical machines, VMs or even a single application.

By monitoring customer's applications for a purpose of more efficient scheduling, provider reaches the privacy border. Also, by applying energy efficient measures like time-sharing VMs and running multiple VMs on a single physical machines, provider creates vulnerable environments for customer's applications. Can customer trust provider's measurements; how secure is his application; is customer's privacy being threatened; these are all the questions which cannot be neglected for benefit of energy efficiency, but should certainly be considered.

3.6 Privacy and Integrity Issues in Data Mining Outsourcing

Anna Monreale (University of Pisa, IT)

"Privacy-preserving data mining from outsourced databases," in S. Gutwirth, Y. Poullet, P. D. Hert, R. Leenes (eds.): Computers, Privacy and Data Protection: an Element of Choice, pp. 411–426, Springer Netherlands, 2011.
URL http://dx.doi.org/10.1007/978-94-007-0641-5_19

Spurred by developments such as in cloud computing, there has been considerable recent interest in the paradigm of data mining-as-service. A company (data owner) lacking in expertise or computational resources can outsource its mining needs to a third-party service provider. In this paradigm two problems arise: i) both data and the knowledge extractable from the outsourced database are considered private property of the corporation (data owner), and so there arises serious privacy issues ii) a dishonest service provider may return inaccurate mining results to the data owner, so there arises serious integrity issue of the mining results.

To protect corporate privacy, the data owner has to transform its data and ship it to the server, send mining queries to the server, and recover the true mining results from the extracted knowledge received from the server. To detect security issues, the data owner has to apply an efficient and practical auditing approach that can verify the correctness and the completeness of mining results.

3.7 To Cloud or Not To. Simple Musings on Cloud Viability

Radu Sion (Stony Brook University, US)

License 🛞 🛞 😑 Creative Commons BY-NC-ND 3.0 Unported license Radu Sion Joint work of Sion, Radu; Chen, Yao; Main reference Y. Chen, R. Sion, "To Cloud Or Not To Cloud? Musings On Costs and Viability," in the ACM Symposium on Cloud Computing (SOCC 2011). URL http://www.cs.sunysb.edu/~sion/research/cloudc2010-draft.pdf

In this talk we aim to understand the types of applications for which cloud computing is economically tenable, i.e., for which the cost savings associated with cloud placement outweigh any associated deployment costs.

We discover two scenarios. In an (i) "unified client" scenario, once cloud-hosted, applications are meant to be accessible only to a single cloud customer (or small set of associates). It then becomes important to ensure that the cost savings (mainly computation-related) can offset the often significant client-cloud distance (network costs etc).

In a (ii) "multi-client" setting on the other hand, outsourced applications serve numerous different third parties. We show that then clouds begin to act similarly in nature to contentdistribution networks – by comparison, their better network integration is simply too good to pass on, when compared to locally hosting the applications (and incurring associated network costs).

Ultimately, we hope this work will constitute a first step in an objective evaluation of the technological side of costs of outsourcing and computing in general.

3.8 On Securing Untrusted Clouds with Cryptography

Radu Sion (Stony Brook University, US)

License 🐵 🕲 🔁 Creative Commons BY-NC-ND 3.0 Unported license © Radu Sion

- Joint work of Sion, Radu; Chen, Yao;
- Main reference Y. Chen, R. Sion, "On Securing Untrusted Clouds with Cryptography," in the ACM Workshop on Privacy in the Electronic Society (WPES 2010). ${\tt URL}\ http://www.cs.sunysb.edu/\ sion/research/sion2010 wpes-pcost.pdf$

In a recent interview, Whitfield Diffie argued that "the whole point of cloud computing is economy" and while it is possible in principle for "computation to be done on encrypted data, [...] current techniques would more than undo the economy gained by the outsourcing and show little sign of becoming practical". In this talk we explore whether this is truly the case and quantify just how expensive it is to secure computing in untrusted, potentially curious clouds.

3.9 BOTCLOUDS – The Future of Cloud-based Botnets?

Martijn Warnier (TU Delft, NL)

Many Cloud Service Providers (CSP) offer access to scalable, reliable computing resources following a payas-you-go model. Research into security of the Cloud focuses mainly on protecting legitimate users of Cloud services from attacks by external, malicious users. Little attention is given to prohibit malicious users from using the Cloud to launch attacks, such as those currently done by botnets. These attacks include launching a DDoS attack, sending spam and perpetrating click fraud. This paper discusses the threat of Cloud-based botnets, or botclouds and the need for new techniques to detect them. Two experiments show how simple and cheaply these attacks can be launched from botclouds.

4 Working Groups

4.1 Privacy, Data Protection and User Empowerment

In this working group, the focus was on supporting the user in making meaningful choices with respect to security and privacy of data. To enable such choices, an interface is needed in which the user can obtain the relevant information, and make her choices known. This requires careful selection of the information and choices presented to the user, as well as the possibility to *actually* change decisions on the fly, that is, to move from one cloud service to another whenever the security requirements change. It also requires the propagation of information all the way up the supply chain from infrastructure to user interface, which requires cooperation of the different providers, and therefore standardisation. Important questions are:

- Modifiability: How to reach-back for security/ privacy customisation?
- Diversity / supply chains: How to manage these issues across multiple clouds?
- Scalability: How to make the approach work without overloading the user with information and choices?
- Mobility: How to swap in/out sub-clouds?
- User affordances: If clouds don't enshrine real-world process complexity, what happens?
- Auditing: How to make auditing manageable if all security is customised?

4.2 Verifying Configurations

The objective of verifying configurations is to know that a certain architecture configuration is running on a cloud. Such an architecture consists of both hardware and software. An example application is when law requires an "adequate level of protection", for example for privacy-sensitive data. The high-level research problems are:

- How does the cloud provider itself know what is running on its systems?
- How can one transfer such complicated knowledge to the cloud user?



Figure 1 An architecture for verifying configurations

An additional challenge is that the cloud provider has right to its own privacy: one may not be able to study, or at least not publish, the underlying hardware because it's a business secret. This requires mechanisms to ensure the confidentiality of the architecture, while ascertaining its high-level security properties. Logging may help as a basis, but if the logs are maintained by the cloud provider, trust is still required.

A basic architecture was proposed for verifying architecture configurations in the cloud (Figure 1). In this architecture, measurements performed by a Measurement Provider provide the basis for an assessment by a Configuration Verifier, which is then communicated to the user. Depending on confidentiality requirements on the architecture, different ways of communicating the information can be proposed.

Apart from verification of the configuration, the user would also be interested in verifying relevant data security properties:

- locality of data,
- integrity of data,
- confidentiality of data,
- availability of data,
- deletion of data, and
- non-repudiation of data leakage.

4.3 Verifying Computations

In the working group on verifying computations, the basic assumption was that, in order to assure integrity of results, it is not enough to check the systems. Following up on electronic voting research, there needs to be a way to check that the results of calculations are correct. One simple way to do this would be to perform the calculations in multiple independent clouds, and then use a voting algorithm to determine the correct result. This would provide statistical confidence about the correctness. However, it is not evident how to assure that clouds are independent, as for example to independent SaaS providers may rent infrastructures from the same IaaS provider. Also, if computations are not deterministic, one cannot expect the results to be the same for different instances.



Figure 2 An architecture for verifying computations

A basic architecture was proposed for verifying the results of computations in the cloud (Figure 2). In this architecture, a proof of the result is generated by a Proof Generator service, which can be checked by a Proof Verifier service. The latter may or may not be localised on the user device. If it is not, for example because of the complexity of the task, another (simpler) mechanism is needed to assure the user of the correctness of the verification of the proof.

5 Open Problems

At the end of the seminar, we identified the following open problems and research directions:

- Architectures for trusted computing without trusted hardware,
- Security product lines for the cloud,
- Further refinement of cloud attacker models,
- Architectures to provide accountability for data leakage,
- Corporate social responsibility and transparency in relation to cloud security,
- Use cases for secure cloud architectures,
- Methods for security-related decision support for cloud consumers,
- Specification of the relation between technical measures and regulation,
- Further refinement of the Verification-as-a-Service concept,
- Development of a measurement/reconstruction suite for cloud architecture configurations, and
- Development of methods for cloud forensics.

All of these problems have been assigned to seminar participants as follow-up activities. For details, please contact the seminar organisers.

6 Panel Discussions

Two cloud security related seminars took place simultaneously at Schloss Dagstuhl from December 4 to December 9, 2011. Seminar 11491 Secure Computing in the Cloud focused on the verifiability, auditability and confidentiality of computation and data, while seminar 11492 Secure Architectures in the Cloud discussed architectures for verification of computations and configurations, as well as customisability of cloud security and privacy. The joint panel discussion featured the panellists Radu Sion, Martijn Warnier, and Marianne Winslett (11492), as well as Ari Juels, Ahmad Sadeghi and Nigel Smart (11491).

Topics of the panel discussion included, but were not limited to, the following. The panel discussed the "big question" whether small or medium-sized enterprises are more secure in the cloud or using their own systems. Naturally, no answer was found. Here as well, an estimation of the security of cloud providers compared to the security of local infrastructures is essential. For this purpose, self-regulatory or government-initiated penetration testing agencies were suggested in order to assess different cloud infrastructures in an objective fashion.

We discussed the security consequences of providing complete infrastructure-as-a-service (IaaS) images in an App-Store like fashion for clouds. This raises security concerns both for the users of such images (potentially malicious software pre-installed) and for the providers of the images (full erasure of sensitive, private data from the images). Consequently, automated checks are needed to address these problems – with some technical details still being challenging.

Other topics included the efficient verifiability of outsourced computation in the general setting that the cloud provider is not fully trusted. Moreover, the internet-of-things was also a topic. That includes car-to-X communication as well as device clouds. The latter allows the creation of ad-hoc clouds, e.g. for the purpose of sharing an internet connection with people who are travelling in order to save roaming fees.

Further, we discussed the possibility of buying insurance for the data stored in the cloud. This, however, requires precise definitions of (a) the coverage of the insurance (data loss, leakage or corruption) and (b) how to assess whether such an event has indeed occurred.

Participants

Arosha Bandara The Open University – Milton Keynes, GB Sören Bleikertz IBM Research – Zürich, CH Travis D. Breaux Carnegie Mellon University -Pittsburgh, US Julien Bringer Morpho, SAFRAN Group, FR Sven Bugiel TU Darmstadt, DE Lizzie Coles-Kemp RHUL - London, GB Sabrina De Capitani di Vimercati University of Milan, IT

Trajce Dimkov
University of Twente, NL
Sebastian Graf
Universität Konstanz, DE
Fabio Massacci
University of Trento – Povo, IT
Toni Mastelic
TU Wien, AT
Sjouke Mauw
University of Luxembourg, LU
Anna Monreale
University of Pisa, IT
Sebastian Pape
TU Dortmund, DE
Wolter Pieters
University of Twente, NL

Christian W. Probst Technical University of Denmark – Lyngby, DK Peter Y.A. Ryan University of Luxembourg, LU Matthias Schunter IBM Research – Zürich, CH Radu Sion Stony Brook University, US André Van Cleeff University of Twente, NL Marcel Waldvogel Universität Konstanz, DE Martijn Warnier TU Delft, NL Marianne Winslett Univ. of Illinois – Urbana, US

