

Thunder in the Clouds: Security Challenges and Solutions for Federated Clouds

Karin Bernsmed*, Martin Gilje Jaatun*, Per Håkon Meland* and Astrid Undheim†

*SINTEF ICT, Norway

Email: {karin.bernsmed,martin.g.jaatun,per.h.meland}@sintef.no

†Telenor Research and Future Studies, Norway

Email: astrid.undheim@telenor.com

Abstract—Cloud federation brings together different service providers and their offered services so that many Cloud variants can be tailored to match different sets of customer requirements. To mitigate security risks and convince hesitant customers, security must be an integrated part of the federated Cloud concept. This paper surveys the state of the art in Cloud computing security, identifies unsolved issues related to federated Clouds, discusses possible approaches to deal with the threats and points out directions for further work.

I. INTRODUCTION

Cloud computing services can currently be found almost everywhere, offering all kinds of IT services imaginable in an on-demand and scalable manner. Cloud computing makes it possible for enterprises via the Internet to manage and deliver services, which can be rapidly provisioned and released by customers. *Federated Clouds* (also known as hybrid Clouds or Clouds-of-Clouds) consist of Cloud services that are composed of one or more other Cloud services. Cloud federation (i.e. the forming of federated Clouds) makes it possible for customers to utilize different types of Clouds to fulfill the requirements they find necessary for their different types of applications. Just as public Clouds enable customers to handle peak demands for information storage, processing and transfer, federated Clouds enable individual Cloud providers to cope with unexpected demand variations [1]. In addition, Cloud federation facilitates for smaller providers to enter the market. By offering service components that can be used as building blocks in more complex Cloud services small and medium-sized businesses and new entrants can become Cloud providers.

Cloud federation is not a new concept, but has until recently existed mostly as a vision rather than deployed solutions. This situation is about to change [2]. The federated Cloud concept is expected to bring Cloud computing to a new level, where organizations and businesses can capitalize on the advantage of aggregated capabilities of disparate data centers, including their own [?], [3]. Such services rely on interoperability among different Cloud providers, and allow providers to dynamically partner with each other [?]. However, the perceived lack of security threatens to be a showstopper for the adoption of the

federated Cloud paradigm.

What makes security difficult in federated Clouds? To put it simply, Cloud computing is outsourcing, and outsourcing implies bidirectional trust relationships. In a Cloud, the responsibility for implementing and maintaining efficient security mechanisms will be in the hands of the provider. To alleviate their customers' fear of the Cloud, these providers need to convince them that their data and applications will be properly secured. The federated Cloud concept takes this uncertainty a step further; in a Cloud-of-Clouds, security responsibility will be split amongst the different actors involved. When security is handled by providers' sub-providers, it is virtually impossible for a Cloud customer to require (for example) on-site security auditing or data center inspections at all these providers. In fact, the customers may not even be aware that there is more than one service provider involved in the final service delivery.

The purpose of this paper is to identify the main security threats related to federated Clouds and to discuss possible solutions. The paper is organized as follows. Section II briefly explains the technical basis for federated Cloud service delivery. In Section III we explain the main security challenges associated with federated Clouds. Section IV discusses three possible approaches that may be used as building blocks for trustworthy Cloud federations. In Section V we review industrial approaches and ongoing research projects related to federated Cloud security. Section VI discusses advantages and weaknesses associated with the three possible approaches identified in Section IV. Finally, Section VII concludes the paper and points out directions for further work.

II. FEDERATED CLOUD SERVICE DELIVERY

The federated Cloud paradigm relies on interoperable services that can be composed from two or more smaller services. The underlying enabler for Cloud compositions is inherited from service oriented architectures (SOA) [6], which enables both loosely and tightly coupled functions or services to operate over a network. Here we define the SOA related concepts used in the rest of the paper. A *service* is a means of delivering value to customers. A service represents some function or type of task performed by a provider on behalf of

a customer. Examples of services are a hotel booking service listed in a public registry or a computing service provided by Amazon¹. A *composite service* is an aggregation of multiple sub-services or service components. These more fine-grained services may be atomic services or other compositions. A *service provider (SP)* is an organization supplying services to one or more internal or external customers. An example is Google, who provides e.g., electronic mail services. A *service customer (SC)* (or just customer) is someone who orders/buys services.

Service compositions can be formed either during the design phase or at run-time. Service components are usually selected based on properties such as functionality, QoS parameters and cost. Today this is mainly done in a static fashion, since fully automatic service composition is still very immature. An inherent problem with most service compositions is that customers buying composite services may know very little about both the service providers and the service components themselves. The reason is that in service oriented architectures, a service component has the nature of a black box, which means that its external interface is exposed, but it is very difficult to check what the component actually does internally [7].

Federated Clouds consist of Cloud services that are composed of one or more other Cloud services. In a federated Cloud the customers normally have a single entity, a “home provider”, with whom they have a relationship. The home provider serves as a gateway to other Clouds by providing the customer with a unified view of the different 3rd party Cloud services. The principle is illustrated in Fig. 1, which displays an example of a federated Cloud where the services offered to the customers are based on a combination of public and private Cloud services, accessible through the private Cloud infrastructure [?].




Fig. 1. An example of a federated Cloud, which combines public Cloud services through private Cloud infrastructure (adapted from [?]).

The federated Cloud paradigm is especially interesting for SMEs, which often seek to avoid large investments. For example, an enterprise that needs to offer its employees a Unified Communication (UC) solution, may buy this service from a Cloud provider that bundles its own voice service with

additional services, such as email and presence, from external service providers [8]. Another example of services that will benefit from the federated Cloud paradigm are SAP systems, such as Customer Relationship Management (CRM), which are complex systems that are used for a variety of business applications [3].

Cloud federation has been a vision for a long time but adoption has been slow. One of the main reasons is the lack of trust between the different parties involved in the service composition. How can one ensure security in federated Clouds? Just trusting the nearest Cloud provider will not be enough. The next section explains the main security challenges associated with Cloud computing and points out three emerging security challenges specifically related to federated Clouds.

III. SECURITY CHALLENGES IN FEDERATED CLOUDS

Security in the Cloud has been a hot topic since the Cloud computing concept first emerged. When asked in surveys, potential customers state security as the main barrier for adopting Cloud computing services [9]. This contrasts with the fact is that there are very few new and unique security issues related to Cloud computing; most of them have been investigated and addressed in the traditional system and network security context. For instance, mechanisms for data protection, access control, trust delegation, mitigation of DDoS attacks and code verification are well known and have been (more or less) successfully applied to large-scale systems and distributed software. Much of what is described in the literature as Cloud computing security should therefore rather be paraphrased as “old security transported to the Cloud” [10]. However, some characteristics of Cloud computing are fundamentally new in an outsourcing perspective, such as multi-tenancy and on-demand elasticity, which open up to a new sets of threats we should be aware of.

The National Institute of Standards and Technology (NIST) has put together a set of guidelines on security and privacy in Cloud computing [11] that can be used by any organization that plans to migrate their data, applications or infrastructure to a public Cloud. Other relevant reading on Cloud computing security includes the European Network and Information Security Agency’s (ENISA) Cloud computing risk assessment guide [?] and the Cloud security guidance report by the Cloud Security Alliance (CSA) [13].

Federated Clouds take the security challenges associated with Cloud computing to a new level. In a federated Cloud, new services can be formed by combining service components from public Clouds with on-site services residing in private Clouds, as well as replacing service components during the lifetime of an existing service. We have identified four new problem areas in security that arise from the formation of federated Cloud services: the longer chain of trust, the limited auditability, the risk of malicious service components, and liability and legal issues. While these problems currently are mostly of academic interest, they risk becoming major obstacles as the federated Cloud concept matures and becomes widely adopted.

¹<http://aws.amazon.com/ec2/>

A. Cloud Federation Challenge: Longer Chain of Trust

Most of the security and privacy concerns that Cloud customers experience are similar to those of traditional non-Cloud services. However, outsourcing data and applications to a public Cloud means that more responsibility is placed on the Cloud provider. The customer loses his previous direct control over both the physical and the logical aspects of data and application storage, processing and transferring in the network. The customer thereby delegates installation, configuration and management of the necessary security mechanisms to the provider, which implies an increased level of confidence and trust in the provider.

Trusting one such provider is one thing; however, with the emerging federated Cloud paradigm, where services are composed of several other services from different providers, we have a situation that implies a *chain of transitive trust*. For example, a Cloud customer may have his application delivered as a SaaS from provider A, who has compiled the application by aggregating additional services from provider B and C, Provider C may in turn have his software running on a PaaS and an IaaS delivered by provider D and E, respectively. Even though the customer has faith in his nearest provider, he may not (and should not) implicitly trust other service providers involved in the service composition. Assuring the customer that adequate security mechanisms exist and are correctly implemented throughout the whole chain of providers will therefore be a huge challenge.

B. Cloud Federation Challenge: Limited Auditability

Limited auditability is a problem in Cloud computing in general and for federated Clouds in particular. Cloud providers are generally not willing to open their systems for outside audits, rather they tend to hide their inner workings from their customers. In federated Clouds auditing becomes even more difficult. Even though customers may have some level of transparency with respect to their home Cloud service provider, they have no way to verify that the underlying components involved in a service composition behave as expected.

The insider security threat is known to be difficult to deal with in most organizations, and this threat applies to Cloud computing as well. In the Cloud, the service provider's system administrators often have unlimited access to the customer data and applications, which implies that illegitimate tampering or usage of customer property may easily go undetected. From the customers' perspective, use of federated Cloud services implies an increased risk; control over who can access and alter the customer data will be extremely difficult to achieve.

Also, in contrast to the traditional ICT system context, mutual auditability is a security issue new to Cloud computing [10], since both the customer and the provider may be the source or the target of an attack. In federated Clouds, the service providers face an increased risk, since they may not know or do not have any possibility to control who the data and applications residing in their data centers belong to.

C. Cloud Federation Challenge: Malicious Service Components

Additional security issues may arise from the service composition itself. Federated Clouds will be based on a service-centric environment, where services will be composed according to properties such as price, quality, latency, security, geography, etc. Today this matching is in most cases performed manually; the service composition is adapted to the customer's current requirements and redesigned whenever a change is necessary. The emerging federated Cloud concept will however require services to be composed dynamically in a more autonomous and ad-hoc manner. This opens the door to a wide range of new threats, such as hi-jacking service components, forming compositions based on both legitimate and illegitimate service components, and injecting malicious services into otherwise legitimate service compositions. An example of the latter threat is the service-injection attack described in [7].

D. Cloud Federation Challenge: Liability and Legal Issues

Cloud customers often have much to lose in case of a service failure or downtime. Most existing contractual agreements between Cloud service provider and customer states a guaranteed service uptime (availability) to be compensated by customer credits should the service fail, but security and performance guarantees are almost always lacking. Many Cloud providers try to elude responsibility by limiting their liability to a level that is far below the potential risk faced by the customer [14]. From the customer's point of view, the more complex services associated with federated Clouds will increase this risk.

Liability issues threaten to be a showstopper for federated Clouds, since a composite service may not have a single "owner" of the system [15]. It is not always straightforward to outsource liability, and companies tend to disacknowledge responsibility in case of security failures. There are several recent examples of Cloud service providers blaming their partners when something goes wrong, e.g. the incident where Play.com customer email addresses were stolen, which they maintained occurred outside their domain [16].

Other issues may arise due to the cross-border nature of federated Clouds. It is often unclear what law is applicable to what Cloud service. An illuminating example is the storage and transfer of personal identifiable information (PII). For instance, the strict privacy regulations adopted by the European Union countries (including the recently introduced Data Retention Directive) are not necessarily compatible with the multitude of different privacy legislations in e.g. Asia. To protect European citizens it will be necessary to guarantee that all handling of PII in all the service components comply with the EU regulations, something that may be very difficult in a federated Cloud.

IV. BUILDING BLOCKS FOR TRUSTWORTHY FEDERATED CLOUDS

A federated Cloud business model will be based on a foundation of trust. In order to interact with external services

that are implemented, deployed and run off-premise, Cloud providers and customers need assurance that information exchanged will not be compromised by other parties. No ready-made solutions for trustworthy federated Clouds currently exist, but we have identified three different approaches that may contribute to a solution to achieve this trust. Note that the examples we highlight are not necessarily implemented in Cloud solutions today.

A. The Trusted Computing Approach

The first approach is based on relying on a third party to implement the basis for trust in external Cloud service providers' infrastructure, storage and processing capabilities. Trusted computing is a technical solution that guarantees the confidentiality and integrity of computations. For example, trusted computing makes it possible for the customer to verify that the service provider administrators do not tamper with the customer VMs even though they have the technical means to access them. By incorporating the concept of trusted computing in the Cloud, Cloud services can be secured from unauthorized access, modification or usage.

Trusted computing is fundamentally based on the presence of a trusted platform module (TPM) on the computing hardware [17]. Such TPMs are now routinely shipped with commodity PCs and servers. A TPM can verify integrity of operating systems and applications, often via a bootstrapping procedure where e.g. a signature on the BIOS firmware is first verified, then the BIOS checks the signature on the bootloader, the bootloader checks the signature on the operating system kernel, and so on. A problem with the original TPMs is that they are not developed for the virtualized environments where Cloud services reside. Therefore virtualized TPM specifications (VTPMs) [18] have been developed, which can be implemented as software instances of TPMs for each virtual machine running on a trusted platform.

Santos et.al [19] propose a trusted Cloud computing platform (TCCP) that can be used to protect confidentiality and integrity of computations that are outsourced to IaaS providers. They propose a TCCP design that includes a set of nodes, which are controlled by an external coordinator controlled by a trusted third party. The nodes embed certified TPM chips, which enables them to run customer VMs securely. The nodes are managed by an (untrusted) Cloud manager, which makes the Cloud services available to the users. Also Krauthem [20] proposes to achieve Cloud security through the introduction of TPMs in all datacenter equipment.

B. The Algorithmic Approach

The second approach is to implement the protection mechanisms on the Cloud customer side, ensuring that information is disseminated in a manner that prevents breach of e.g., confidentiality. Many of the current proposals are based on complex cryptographic principles, as will be outlined in the following.

1) *Secure multiparty computation*: Secure multiparty computation is one of the classic exercises in cryptographic protocols. The concept can be illustrated with a toy example, due to Schneier [21]:

If a group of employees want to calculate their average salary without revealing exactly how much they earn to each other, the first person could add a random value to their salary, and pass the sum on to the next person. This person adds their salary, passes the new sum on to the next person, and so on. When all employees have added their salaries, the final sum is returned to the first person, who deducts the initial random value to extract the sum of all salaries. The average salary can then be calculated by dividing the sum by the total number of employees who participated.

This simple example only works if all the participants are honest, i.e., faithfully execute the protocol without lying, and also requires that intermediate messages are kept secret from all but the two communicating parties (i.e., user n_i and n_{i+1}). More serious schemes have been suggested e.g. by Chaum et al. [22], which can guarantee a secure result provided (in the worst case) at least $(2n/3)+3$ of the n participants are honest.

A practical example of using secure multiparty computations is the Sharemind system [23], in which a community of users can perform privacy-preserving calculations on a group of three so-called data miners. Although not specifically designed for Cloud computing, there should be no inherent barriers to porting Sharemind to the Cloud. Note that although each individual data miner does not have to be trusted, confidentiality can be breached if the miners collude.

2) *Fully homomorphic encryption*: Fully homomorphic encryption [24], [25] is an encryption scheme where, e.g., arithmetic operations can be performed on the encrypted data, and return the correct plaintext result when the encrypted result is decrypted. To carry our secure multiparty example further, an alternative way of calculating the average salary would be if each employee encrypted their salary using fully homomorphic encryption with the same key, and sent this to an "untrusted" Cloud processing provider in a secure manner. The Cloud provider could then add up all the encrypted salaries, and divide by the number of participating employees. The result can then be returned to everyone, and be decrypted with the inverse of the key everyone used in the first step.

A less trivial example is provided by Mowbray et al. [26], who describe a scheme for privacy in the Cloud by using obfuscation of sensitive data. Although their scheme is not really homomorphic encryption, the example of an online share portfolio would be applicable here as well: The user would encrypt (in our version of the example, using fully homomorphic encryption) the amounts of her various stock holdings (e.g., 200 000 shares of MSFT, 40 000 bushels of AAPL, etc.) before putting everything in the Cloud. The Cloud provider thus only has access to the encrypted values, but can at any time calculate an encrypted total value of the stock

portfolio with current stock prices. This value can be returned to the user, who can decrypt it to establish her net worth.

3) *Security through data splitting*: Another algorithmic approach, which is not inherently based on cryptography, is to split data across different Cloud providers. An example is the Redundant Array of Independent Net-storages (RAIN) proposed by Jaatun et al. [27], [28]. Their model provides confidentiality control of data stored in the Cloud. In the RAIN approach, data is split into segments and distributed between multiple storage providers, which makes it virtually impossible for a single observer to re-assemble the original data. The main advantage is the preservation of data confidentiality without having to rely on heavy cryptographic operations. The main challenge may be the reliance on an anonymous sender-receiver framework (e.g., TOR [29]), which in turn usually is implemented using fairly computationally-intensive cryptographic mechanisms [30], [31]. Other similar approaches are HAIL [32], which uses RAID-like techniques across storage vendors to ensure high-availability and integrity of data stored in the Clouds, and RACS [33], which is a proxy that transparently spreads the storage load over many providers. All three approaches are ideally suited to a federated Cloud environment, which can easily provide a suitable mixture of independent storage service providers.

C. The Contractual Approach

The third approach to achieve trust in federated Clouds is through the use of contracts. A Service Level Agreement (SLA) is a common way to specify the conditions under which a service is to be delivered. Today, a typical SLA for a Cloud service is specified at a top-level between the customer and the provider, usually limited to availability levels and credits/penalties. Since the SLA is used to explicitly state the obligations of the provider, the implemented security mechanisms, their effectiveness, and the implications of possible mismanagement should be a part of this agreement. This concept is sometimes known as Quality of Protection (QoP), which comprises the ability of a service provider to deliver service according to a set of specific security requirements [8].

1) *Security SLAs*: There have been some projects in the research community looking into various aspects of security in SLAs. Early work on security agreements was performed in 1999 by Henning [34], who already then raised the question whether security can be adequately expressed in an SLA. More recently, Casola et al. [35] proposed a methodology to help evaluate and compare security SLAs. Frankova and Yautsiukhin [36] also recognized the need for security in SLAs. Their approach focused on the process of selecting the optimal service composition based on a set of pre-defined requirements. Chaves et al. [37] argued that security in SLAs has an important role in Cloud computing. They explored security in SLAs, with focus on measurable security metrics combined with a monitoring and controlling architecture [38]. As pointed out by Chaves et al., it is a challenge to define quantifiable security metrics, but they give examples related to password management, frequency of backups and repair/recovery time.

Bernsmed et al. introduced a framework for incorporating security in Cloud SLAs [39]. They presented a set of standard security mechanisms organized according to different types of Cloud services.

2) *Dynamic SLA management*: SLA management today is mainly a static affair where the contract terms are defined by the provider, typically published on a Web page, and intended to be read by humans. With a shift towards a more dynamic service environment – where services change terms, are composed from resources in federated or hybrid Clouds, and where more interactive SLA negotiations take place – the SLA management must become a more automatic process performed by software agents [8]. More detailed SLAs would also be an argument for dynamic management as frequent service updates easily could cause more contract violations. Early work on SLA management for federated environments was performed by Bjoh et al. [40] who developed an architecture to allow SLA monitoring and sharing of selective management information across administrative domain boundaries. They demonstrated its applicability using a prototype implementation that measures the availability, performance and utilization of an email service.

Dynamic SLA management has recently been targeted in sister sciences to Cloud computing, such as Grid and Web services. There are two main specifications that describe SLAs for web services. The first is the Web Service Agreement (WS-Agreement), developed by the Open Grid Forum (OGF) [41]. The second is the Web Service Level Agreement (WSLA) framework [42], which was developed by IBM for SLA monitoring and enforcement in a Service Oriented Architecture (SOA). Patel et al. [43] propose a mechanism for managing Cloud SLAs using the WSLA framework. Their main contribution is the usage of the 3rd party support feature of WSLA to delegate the monitoring and enforcement part of the SLA management to trusted 3rd parties. Relevant work has also been performed by Comuzzi et al. [44] and Theilmann et al. [45], contributing to reference architectures for multilevel SLA management. Unfortunately, none of these approaches include security as an attribute in the SLA.

V. CLOUD FEDERATION STATE-OF-THE-ART

2010 was the year when Cloud computing moved from hype to reality. Today, many businesses are either already using or planning to use Cloud services. The next years are predicted to be all about Cloud federation [2]. The vision of an open federated Cloud ecosystem is evident in both industry and academy, which both have security and interoperability on their current agenda.

A. Industrial Efforts

On the industry side, Intel is leading the Open Data Center Alliance (ODCA)², which at the time of writing has more than 300 members, envisioning an open Cloud federation ecosystem. Their goal is to enable Cloud service delivery

²<http://www.opendatacenteralliance.org/>

and consumption in a more secure, interoperable and efficient manner. The purpose is to allow data to be shared between private and public Clouds, to facilitate automatic movement of applications and resources and to make the Cloud services aware of the client devices.

In addition, several independent actors have appeared on the market. SpotCloud³ was one of the first marketplaces where service providers could offer their services, thus utilizing otherwise unused capacity in Cloud datacenters. ScaleUp Technologies⁴ is a German-based Cloud infrastructure provider that offers support for federation across multiple Cloud providers from within their Cloud management platform. Their software is based on OpenStack⁵, which is an open source IaaS platform originally founded by hosting company RackSpace and NASA. Scaleup claims to be an enabler of marketplaces where Cloud providers can offer their services. A third actor is ScaleXtreme⁶, which provides its customers with the ability to manage IaaS from a number of public Cloud providers.

The ODCA has put security high on their agenda but has not yet stated how they will achieve trustworthy Cloud federations. Neither SpotCloud, ScaleUp nor ScaleXtreme use any of the proposed trust-based approaches mentioned above today, but for example SpotCloud claims that they are watching the demands of the market closely to determine when there will be a business case for prioritizing SLAs.

B. Federated Cloud Research in the EU

Security in Cloud computing is a topic that has gained increased interest in the research community in the last few years. There are several EU funded projects under the seventh framework programme (FP7) that are currently addressing this topic. ANIKETOS [47] focuses on secure and trustworthy service compositions. They work on new technology, methods, tools and security services that support the design-time creation and run-time dynamic behavior of composite services, based on socio-technical aspects as well as basic technical issues. The TClouds project [?] aims to develop a security architecture for federated Cloud infrastructures. TClouds will develop privacy-protecting protocols for transferring information securely between different Cloud service providers. They will also look into developing security standards and building open APIs (application programming interfaces) and secure Cloud management components. Also the OPTIMIS project [?] will deliver a specification and toolkit that enables the generation of federated Clouds, where security is one of the non-functional requirements that will decide the final composition of a service.

In addition, several EU funded projects address the topic of Cloud federation, but focusing on other aspects than security. The RESERVOIR project [3], [49] develops a model and architecture for open federated Cloud computing. The project addresses the limited scalability of a single-provider Cloud, the

lack of interoperability among Cloud providers, and the lack of built-in Business Service Management support in current Cloud offerings. The VISION project [50] is another EU funded project that aims to deliver an architecture and reference implementation of a Cloud-based infrastructure, which is to be built on open standards and new technologies. They will focus on scalable, flexible, and dependable framework to deliver data-intensive storage services.

VI. DISCUSSION

Cloud services are not necessarily stand-alone or independent from other services provided by different vendors; the federated Cloud concept is a prime example of that. Federated Cloud services depend on mechanisms to establish mutual trust and secure interactions. The underlying enabler for federated Clouds is inherited from service oriented architectures (SOAs) and federated Clouds are likely to suffer from similar challenges that are still to be solved with SOAs. The lack of security specifications [51], guarantees of protection [52] and trust requirements in Web service orchestration and choreography are already known challenges, and are likely to emerge in a broader scale as we see more Cloud services and more collaboration and competition between Cloud service providers.

Trusted computing is an approach that can help secure data and applications in federated Clouds. However, it depends on standardization between the various hardware and software platforms in the Cloud. In addition the trusted computing concept seems incompatible with the concept of open source software⁷, on which Cloud federation relies heavily. If TPMs are to be effective in virtualized environments, trust mechanisms need to be built at every layer in the Cloud using both hardware TPMs and virtual machine monitors [53].

The main challenge with solutions such as secure multiparty computation and fully homomorphic encryption is that they remain too demanding with respect to processing effort to be employed for the general case. It is possible that future advances in processing capacity may render these concerns moot, but in the meantime we have to keep searching for approaches that are deployable for the general case without overstepping the bounds of “reasonable processing”. Security through data splitting is a promising alternative for securing federated Cloud storage, but requires an additional management layer in form of a proxy architecture. This is still an immature research field where more research is necessary in order to evaluate the performance and cost implications of the proposed architectures.

To mitigate the security risks associated with the federated Cloud, existing security mechanisms and their effectiveness can be formalized in contracts. The absence of security mechanisms in today’s SLAs combined with the lack of methods for making objective comparisons between different service offerings makes it virtually impossible for Cloud providers

³<http://www.spotCloud.com/>

⁴<http://www.scaleupCloud.com/>

⁵<http://www.openstack.org/>

⁶<http://www.scalextreme.com/>

⁷<http://www.cl.cam.ac.uk/~rja14/tcpa-faq.html>

to offer trustworthy services to their customers when external providers are involved [8]. Organizations that use Cloud services can in some cases have strict security policies on where their data should be stored, and according to Buyya et al. [54], it should be possible to have SLAs that specify the location of Cloud resources. Disappointingly, an investigation by Honeyball [55] showed that it was not possible for EU customers of the Microsoft Azure platform to get any legally binding guarantee where their data would or would not be stored. This also reflects the disadvantaged bargaining position small users find themselves in when negotiating with a huge Cloud provider, often resulting in unilateral agreements presented by the provider in a “take it or leave it” fashion.

Extending an SLA to include specific security mechanisms may not be overly difficult, but the challenge is to define a *security level*, and then verify whether the provider actually delivers on his promise. Future research therefore needs to investigate how security agreements (SLAs) for Cloud computing services can be formed, verified and maintained over time. Security SLAs will not only to increase the trust in the provider, but also facilitate objective comparisons between different service providers on the basis of their security features. Such an approach will also form a basis for composing services from different providers, based on a set of pre-defined security requirements [56].

A trustworthy service composition will be fundamental for both Cloud customers and providers. Developers need to be able to describe security properties as well as safe and secure behavior, from both a technical point of view and the organizational and business perspective. Run-time monitoring and automatic adaptation of services are needed due to an evolving environment of threats and operating conditions. This requires methods and tools for designing, developing, composing and running services, based on usable and efficient security mechanisms. To achieve this, new methods for including threat awareness and threat response capabilities in the software services must be developed.

VII. CONCLUSION

Building secure software services will be a prerequisite for succeeding in the Cloud, where anything may be thought of as services; storage capacity, computing resources, software applications, etc. The federated Cloud provides an environment in which a diverse range of services are offered by a wide range of suppliers, and where collaboration, competition and dynamic changes in the behavior are key concepts. Composing services from different providers in several steps represents a security challenge that may be a showstopper for mainstream adoption of federated Clouds. This paper has identified and discussed some of the most urgent security challenges associated with the Cloud federation concept.

Our further work will be focused on development of security SLAs [56] in federated Cloud scenarios that include a brokerage function.

ACKNOWLEDGMENT

This research has been supported by Telenor through the SINTEF-Telenor research agreement.

REFERENCES

- [1] E. R. Gomes, Q. B. Vo, and R. Kowalczyk, “Pure exchange markets for resource sharing in federated clouds,” *Concurrency and Computation: Practice and Experience*, pp. n/a–n/a, 2010. [Online]. Available: <http://dx.doi.org/10.1002/cpe.1659>
- [2] K. Subramanian. Research Report: Cloud Trends In 2011 And Beyond. [Online]. Available: <http://www.cloudave.com/9587/research-report-cloud-trends-in-2011-and-beyond/>
- [3] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Cáceres, M. Ben-Yehuda, W. Emmerich, and F. Galán, “The RESERVOIR model and architecture for open federated cloud computing,” *IBM J. Res. Dev.*, vol. 53, pp. 535–545, July 2009. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1850659.1850663>
- [4] K. Salmon, “Demystifying the Cloud.” [Online]. Available: <http://www.kurtsalmon.com/publications/demystifying-the-cloud/>
- [5] E. Network and I. S. Agency, “The Future of Cloud Computing. Opportunities for European Cloud Computing Beyond 2010.”
- [6] M. Vouk, “Cloud computing - issues, research and implementations,” in *Information Technology Interfaces, 2008. ITI 2008. 30th International Conference on*, june 2008, pp. 31–40.
- [7] P. H. Meland, “Service injection: A threat to self-managed complex systems,” in *Dependable, Autonomic and Secure Computing, IEEE International Symposium on*. Los Alamitos, CA, USA: IEEE Computer Society, 2011, pp. 1–6.
- [8] K. Bernsmed, M. G. Jaatun, P. H. Meland, and A. Undheim, “Security SLAs for Federated Cloud Services,” in *Proceedings of the Sixth International Conference on Availability, Reliability and Security (AREs 2011)*, 2011.
- [9] (2011) Cloud Security Survey Global Executive Summary. Trend Micro Inc. [Online]. Available: http://us.trendmicro.com/imperia/md/content/us/trendwatch/cloud/global_cloud_survey_exec_summary_final.pdf
- [10] Y. Chen, V. Paxson, and R. H. Katz, “What’s New About Cloud Computing Security?” EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2010-5, Jan 2010. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-5.html>
- [11] W. Jansen and T. Grance., “Guidelines on security and privacy in public cloud computing,” NIST, Draft Special Publication 800-144, Jan 2011.
- [12] European Network and Information Security Agency (ENISA), “Cloud Computing: Benefits, risks and recommendations for information security,” Nov. 2009.
- [13] “Security Guidance for Critical Areas of Focus in Cloud Computing,” 2010. [Online]. Available: <https://cloudsecurityalliance.org/guidance/>
- [14] W. M. H. and D. J. Buller, “Cloud computing: Emerging legal issues for access to data, anywhere, anytime,” *Journal of Internet Law*, July 2010.
- [15] L. M. Vaquero, L. Rodero-Merino, J. Cáceres, and M. Lindner, “A break in the clouds: towards a cloud definition,” *SIGCOMM Comput. Commun. Rev.*, vol. 39, pp. 50–55, December 2008. [Online]. Available: <http://doi.acm.org/10.1145/1496091.1496100>
- [16] E. Doyle, “Who Is Carrying The Can For Cloud Data Security?” March 28th 2011. [Online]. Available: <http://www.eweekeuropa.co.uk/comment/who-is-taking-responsibility-for-cloud-data-security-25023>
- [17] “Trusted computing group - home,” 2011, <http://www.trustedcomputinggroup.org/>.
- [18] S. Berger, R. Cáceres, K. A. Goldman, R. Perez, R. Sailer, and L. van Doorn, “vtpm: virtualizing the trusted platform module,” in *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*. Berkeley, CA, USA: USENIX Association, 2006. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1267336.1267357>
- [19] N. Santos, K. P. Gummadi, and R. Rodrigues, “Towards trusted cloud computing,” in *HOTCLOUD*. USENIX, 2009.
- [20] F. J. Krauthem, “Private virtual infrastructure for cloud computing,” in *Proceedings of the 2009 conference on Hot topics in cloud computing*, ser. HotCloud’09. Berkeley, CA, USA: USENIX Association, 2009, pp. 5–5. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1855533.1855538>
- [21] B. Schneier, *Applied Cryptography*, 2nd ed. John Wiley & Sons, 1996.

- [22] D. Chaum, C. Crépeau, and I. Damgård, "Multiparty unconditionally secure protocols," in *Proceedings of the twentieth annual ACM symposium on Theory of computing*. ACM, 1988, pp. 11–19.
- [23] D. Bogdanov, S. Laur, and J. Willemson, "Sharemind: a framework for fast privacy-preserving computations," Cryptology ePrint Archive, Report 2008/289, 2008, <http://eprint.iacr.org/>.
- [24] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proceedings of the 41st annual ACM symposium on Theory of computing*. ACM, 2009, pp. 169–178.
- [25] N. Smart and F. Vercauteren, "Fully homomorphic encryption with relatively small key and ciphertext sizes," in *Proceedings of Public Key Cryptography–PKC 2010*. Springer, 2010, pp. 420–443.
- [26] M. Mowbray, S. Pearson, and Y. Shen, "Enhancing privacy in cloud computing via policy-based obfuscation," *The Journal of Supercomputing*, pp. 1–25, 2010, 10.1007/s11227-010-0425-z. [Online]. Available: <http://dx.doi.org/10.1007/s11227-010-0425-z>
- [27] M. G. Jaatun, Å. A. Nyre, S. Alapnes, and G. Zhao, "A Farewell to Trust: An Approach to Confidentiality Control in the Cloud," in *Proceedings of the 2nd International Conference on Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology (Wireless Vitae Chennai 2011)*, 2011.
- [28] M. G. Jaatun, G. Zhao, A. Vasilakos, Å. A. Nyre, S. Alapnes, and Y. Tang, "The design of a redundant array of independent net-storages for improved confidentiality in cloud computing," *Journal of Cloud Computing: Advances, Systems and Applications*, vol. 1, no. 1, p. 13, 2012. [Online]. Available: <http://www.journalofcloudcomputing.com/content/1/1/13>
- [29] R. Dingleline, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proceedings of the 13th conference on USENIX Security Symposium–Volume 13*. USENIX Association, 2004, pp. 21–21.
- [30] D. Chaum, "The dining cryptographers problem: Unconditional sender and recipient untraceability," *Journal of Cryptology*, vol. 1, pp. 65–75, 1988, 10.1007/BF00206326. [Online]. Available: <http://dx.doi.org/10.1007/BF00206326>
- [31] D. L. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Commun. ACM*, vol. 24, pp. 84–90, February 1981. [Online]. Available: <http://doi.acm.org/10.1145/358549.358563>
- [32] K. D. Bowers, A. Juels, and A. Oprea, "Hail: a high-availability and integrity layer for cloud storage," in *Proceedings of the 16th ACM conference on Computer and communications security*, ser. CCS '09. New York, NY, USA: ACM, 2009, pp. 187–198. [Online]. Available: <http://doi.acm.org/10.1145/1653662.1653686>
- [33] H. Abu-Libdeh, L. Princehouse, and H. Weatherspoon, "Racs: a case for cloud storage diversity," in *Proceedings of the 1st ACM symposium on Cloud computing*, ser. SoCC '10. New York, NY, USA: ACM, 2010, pp. 229–240. [Online]. Available: <http://doi.acm.org/10.1145/1807128.1807165>
- [34] R. R. Henning, "Security service level agreements: quantifiable security for the enterprise?" in *Proceedings of the 1999 workshop on New security paradigms*, ser. NSPW '99. New York, NY, USA: ACM, 2000, pp. 54–60.
- [35] V. Casola, A. Mazzeo, N. Mazzocca, and M. Rak, "A SLA evaluation methodology in Service Oriented Architectures," in *Quality of Protection*, ser. Advances in Information Security, D. Gollmann, F. Massacci, and A. Yautsiukhin, Eds. Springer US, 2006, vol. 23, pp. 119–130.
- [36] G. Frankova and A. Yautsiukhin, "Service and protection level agreements for business processes," in *Young Researchers Workshop on Service*, 2007.
- [37] S. A. de Chaves, C. B. Westphall, and F. R. Lamin, "SLA Perspective in Security Management for Cloud Computing," in *Proceeding of the 2010 Sixth International Conference on Networking and Services*. IEEE, March 2010, pp. 212–217.
- [38] R. R. Righi, D. L. Kreutz, and C. B. Westphall, "Sec-mon: An architecture for monitoring and controlling security service level agreements," in *XI Workshop on Managing and Operating Networks and Services*, 2006.
- [39] K. Bernsmed, M. G. Jaatun, and A. Undheim, "Security in Service Level Agreements for Cloud Computing," in *Proceedings of the 1st International Conference on Cloud Computing and Services Science (CLOSER 2011)*, 2011.
- [40] P. Bhoj, S. Singhal, and S. Chutani, "SLA management in federated environments," *Comput. Netw.*, vol. 35, pp. 5–24, January 2001. [Online]. Available: <http://portal.acm.org/citation.cfm?id=370802.370805>
- [41] OGF, "Web Services Agreement Specification (WS-Agreement)," Tech. Rep., 2007.
- [42] H. Ludwig, A. Keller, A. Dan, R. P. King, and R. Franck, "Web Service Level Agreement (WSLA) Language Specification," IBM, Tech. Rep., 2003.
- [43] P. Patel, A. Ranabahu, and A. Sheth, "Service Level Agreement in Cloud Computing," in *Proceedings of OOPSLA 2009*, 2009.
- [44] M. Comuzzi, C. Kotsokalis, C. Rathfelder, W. Theilmann, U. Winkler, and G. Zacco, "A framework for multi-level sla management," in *Proceedings of the 2009 international conference on Service-oriented computing*, ser. ICSSOC/ServiceWave'09. Springer-Verlag, 2009, pp. 187–196.
- [45] W. Theilmann, J. Happe, C. Kotsokalis, A. Edmonds, K. Kearney, and J. Lambea, "A Reference Architecture for Multi-Level SLA Management," *Journal of Internet Engineering*, vol. 4, no. 1, pp. 289–298, 2010.
- [46] "Aniketos: Ensuring Trustworthiness and Security in Service Composition," 2011. [Online]. Available: <http://www.aniketos.eu/>
- [47] "The TClouds project," <http://www.tclouds-project.eu/>.
- [48] "The Optimis project," <http://www.optimis-project.eu/>.
- [49] "RESERVOIR: Resources and Services Virtualization without Barriers," <http://www.reservoir-fp7.eu/>.
- [50] "VISION Cloud," <http://www.visioncloud.eu/>.
- [51] D.-H. Xu, Y. Qi, D. Hou, G.-Z. Wang, and Y. Chen, "An improved calculus for secure dynamic services composition," in *Proceedings of the 2008 32nd Annual IEEE International Computer Software and Applications Conference*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 686–691. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1444455.1446065>
- [52] A. Singhal, "Web services security: Challenges and techniques," in *Proceedings of the Eighth IEEE International Workshop on Policies for Distributed Systems and Networks*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 282–. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1263544.1263904>
- [53] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: state-of-the-art and research challenges," *Journal of Internet Services and Applications*, vol. 1, no. 1, pp. 7–18, 2010. [Online]. Available: <http://www.springerlink.com/index/10.1007/s13174-010-0007-6>
- [54] R. Buyya, C. S. Yeo, and S. Venugopal, "Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities," in *High Performance Computing and Communications, 2008. HPCC '08. 10th IEEE International Conference on*, sept. 2008, pp. 5–13.
- [55] J. Honeyball, "The truth about microsoft azure - and where your data will be kept," *PC Pro*, August 2009.
- [56] P. H. Meland, K. Bernsmed, M. G. Jaatun, H. Castejon, and A. Undheim, "Expressing Cloud Security Requirements in Deontic Contract Languages," in *Proceedings of the 1st International Conference on Cloud Computing and Services Science (CLOSER 2012)*, 2012.