

Kent Academic Repository

Full text document (pdf)

Citation for published version

Pawar, Pramod and Sajjad, Ali and Chadwick, David W. and Dimitrakos, Theo (2015) Security-as-a-Service in Multi-cloud and Federated Cloud Environments. In: Trust Management IX. IFIP Advances in Information and Communication Technology . Springer International Publishing, Hamburg, Germany, pp. 251-261. ISBN 978-3-319-18490-6.

DOI

https://doi.org/10.1007/978-3-319-18491-3_21

Link to record in KAR

<http://kar.kent.ac.uk/56874/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Security-as-a-Service in Multi-cloud and Federated Cloud Environments

Pramod S. Pawar¹, Ali Sajjad¹, Theo Dimitrakos^{1,2}, David W Chadwick¹

¹ The University of Kent, Canterbury, Kent CT2 7NZ, United Kingdom

² British Telecommunications, Adastral Park, Ipswich IP5 3RE, United Kingdom

{pramod.s.pawar, ali.sajjad, theo.dimitrakos}@bt.com,
d.w.chadwick@kent.ac.uk

Abstract. The economic benefits of cloud computing are encouraging customers to bring complex applications and data into the cloud. However security remains the biggest barrier in the adoption of cloud, and with the advent of multi-cloud and federated clouds in practice security concerns are for applications and data in the cloud. This paper proposes security as a value added service, provisioned dynamically during deployment and operation management of an application in multi-cloud and federated clouds. This paper specifically considers a data protection and a host & application protection solution that are offered as a SaaS application, to validate the security services in a multi-cloud and federated cloud environment. This paper shares our experiences of validating these security services over a geographically distributed, large scale, multi-cloud and federated cloud infrastructure.

Keywords: application security, data security, cloud security, multi-cloud, federated cloud

1 Introduction

Cloud computing provides flexible and dynamic access to virtualized computing and network resources, however, its complexity, especially in multi-cloud and federated cloud environments gives users cause for concern over the security of services hosted in the cloud[1]. Due to the dynamic nature of clouds, new categories of security threat emerge [2]. Hashizume et al. provide an analysis of security issues in the cloud considering the three service delivery models, SaaS, PaaS and IaaS [3]. The major security concerns for SaaS applications include data location, legislative compliance of its data, security policy of the providers and data protection. This obliges the SaaS provider to have additional security mechanisms beyond what is offered by the PaaS and IaaS providers, in order to protect their applications and data.

The emphasis of this paper is on the security and protection of the SaaS applications and user/consumer data. Due to the dynamic nature of the cloud, significant challenges exist in applying the traditional security solutions such as firewalls, IDS/IPS and data protection, in the cloud environment. This paper proposes security as a value added

service that forms an additional layer of security for hosts and applications in the cloud and is dynamically provisioned. Specifically, it considers a data protection solution and an application protection solution that are offered as a SaaS application, to demonstrate security services in a multi-cloud and federated cloud environment.

Most of today's existing data protection or secure cloud storage services focus on file-level encryption of the user's data, following one of two approaches: either the data is uploaded to the cloud provider and is then encrypted, in which case the keys are managed by the service provider, e.g. Dropbox, Google Drive, Microsoft Sky Drive; or the data is encrypted at the user end and then uploaded to the secure storage service provider and the keys are managed by the user, e.g. BoxCryptor or the Virtual Cloud Drive [4]–[8]. Although these approaches are suitable from the point-of-view of online backup and write-once read-many types of scenarios, they become very unwieldy when the data has to be modified frequently and resides on virtual machines in the cloud. This is further complicated when the user wants to take advantage of the growing inter-cloud usage scenarios. The data protection solution considered in this paper is developed to address these issues, however there is a need to validate the functioning in a multi-cloud use case scenarios.

The host and application protection solution considered in this paper is developed to enable cloud providers to dynamically provision the protection functions to cloud service users while allowing them to have full control over the security of their applications and hosts. As this solution is required to be integrated in the provisioning workflow of the cloud service provider and required to be offered to large numbers of cloud consumers, there is a need to verify the automation workflow since it will automatically deploy the protection components in a multi-cloud environment, on a heterogeneous and scalable infrastructure that can host hundreds of VMs.

The validation of these security services is performed as experiments executed over the Fed4FIRE (Federation for Future Internet Research and Experimentation) infrastructure which is an EU funded, geographically distributed, multi-cloud and federated cloud infrastructure [9]. This paper describes the experiments performed for the dynamic provisioning and automation of services, our experiences and findings and provides feedback about the Fed4FIRE infrastructure. The rest of this paper is organized as follows. Section 2 provides the details of the application protection and data protection solutions used in this paper. Section 3 describes the Fed4FIRE infrastructure used for evaluating the solutions. Section 4 provides the objectives of the experiment and section 5 describes the experiments. Section 6 provides the experimental results. Finally, section 7 provides the concluding remarks and the future work.

2 Overview of the security solution used for experimentation

2.1 Secure Cloud Storage - Data Protection Service

Secure Cloud Storage (SCS) is a cloud security service that provides data protection for public and private clouds and other virtualization platforms. It allows users to protect and control their confidential and sensitive information with a user-friendly file and

volume encryption service that keeps their data private and helps meet their regulatory compliance requirements. Secure Cloud Storage can be deployed as a hosted Software-as-a-Service or as an On-Premise software application, but in either case only the user has access and control of the decryption keys, giving them the freedom to decrypt their data on-demand and in real time. It offers users the capability of applying policy-based key management, a means to validate the identity and integrity of virtual machines requesting the encryption keys, and it can specify where and when the encrypted data can be accessed.

Key technical features of this service include:

(1) Policy-driven Key Management: this (a) uses identity and integrity-based policy enforcement to ensure only authorized virtual machines receive keys and access secure volumes (b) automates key release and virtual machine authorization for rapid operations or requires manual approval for increased security (c) enables the use of policies to determine when and where keys were used;

(2) Advanced Encryption Techniques: this (a) features FIPS 140-2 certification and FIPS approved AES encryption (b) encrypts and decrypts information in real time, so data at rest is always protected (c) applies file and volume encryption to secure all data, metadata, and associated structures without impacting application functionality;

(3) Robust Auditing, Reporting, and Alerting: this (a) logs actions in the management console for audit purposes (b) provides detailed reporting and alerting features with incident-based and interval-based notifications.

Users of the Secure Cloud Storage service can monitor the integrity and protection status of their volumes via a configurable web-based dashboard offered to them over a secure channel via BT's hosted service. Users can define different roles for their cloud administrators, security operation teams and auditors giving them different levels of control visibility and rights to define policies as appropriate. Security administrators with the appropriate rights can define policies for each volume on any virtual machine that has been registered by a secure cloud agent.

2.2 Intelligent Protection - Host & Application Protection Service

Intelligent Protection is a cloud security service that is designed and developed to protect virtual servers and hosted applications on cloud infrastructures [10]. The novelty of this service centers on offering security as a value-added service (multi-tenant security SaaS) while enforcement is delivered via the cloud infrastructure, with minimal integration overhead. It enhances cloud user experience by offering more secure, flexible, automated security management for applications deployed or on-boarded on cloud infrastructures (IaaS) such as BT Compute or other 3rd party equivalents (e.g. Amazon EC2 or V-Cloud enabled IaaS) while placing the users in control of their own security operations through its Security SaaS operations dashboard [11], [12].

Intelligent Protection enables its users to automatically perform the following security functions via an intuitive web interface: (1) Virtual Security Patches (2) Intelligent Intrusion Detection and Prevention (IDS/IPS) (3) Bi-directional stateful firewall (4) Anti-Malware (5) Integrity Monitoring (6) Incident Reporting and Analysis (7) Recommendation Scans.

A user of the intelligent protection service can: analyze their virtual networks, servers, and applications for vulnerabilities; obtain recommendations of missing security patches, and the best security to address the identified vulnerabilities on each system that has been analyzed; continuously monitor for attacks, intrusion, viruses, exploits and any other security incident; monitor performance and scalability; and apply corrective measures updating the security policies accordingly.

Users of the Intelligent Protection service can monitor the health and protection of their virtual machines and servers via a configurable dashboard offered to them over a secure channel via BT's hosted service. Tenants can define different roles for their cloud administrators, security operation teams and auditors giving them different levels of control visibility and rights to define policies as appropriate. Security administrators can define security policies via an intuitive GUI by combining policies from a library or by asking the Intelligent Protection system to analyze a selected environment's vulnerabilities and recommend which security policy rules to apply on that environment.

3 Overview of the Fed4FIRE facilities

. The goal of the EU FP7 Fed4FIRE project is to federate various Internet experimentation facilities to enable an innovative cross-domain experimentation platform, providing researchers with easy access to resources on different facilities. The project currently involves multiple facilities, introducing a diverse set of technologies such as cloud computing, wired and wireless network, software defined networking, Internet of Things and smart cities. Some example facilities are: EPCC BonFIRE(UK), Virtual Wall (EU), PlanetLab Europe (EU), Smart Santander(EU), NORBIT(Australia), KOREN(Korea), Sanford optical access testbed(USA) [9].

The incentive for using the Fed4FIRE facilities stems, on the one hand, from the scale and heterogeneity that the Fed4FIRE facilities bring to the experimentation, and on the other hand from the likelihood of continuity and expansion of research experimentation by using the FIRE facilities that expose open standards management interfaces. Furthermore it enables the verifiability and validity of the protection services over 3rd party Cloud platforms. These characteristics make Fed4FIRE an appropriate environment for evaluating Secure Cloud Storage (SCS) and Intelligent Protection solutions. Furthermore, the geographically distributed infrastructure of Fed4FIRE, which is connected through the Internet, allows the experiments to be run under real-world network behavior that is not present in a simulation or a single public or private cloud. The comprehensive monitoring infrastructure and services of Fed4FIRE enables the scalability and performance of the security solutions to be easily measured. Additionally, Fed4FIRE provides the jFed tool and Rspec. The jFed tool is a Java based framework that provides an integrated view and interface to communicate with all the infrastructures available in Fed4FIRE. Rspec is the resource specification written in the jFed tool that allows provisioning of VMs on the Fed4FIRE infrastructure.

4 Objectives of the experimentation

The evaluation of our data protection service and host & application protection service is targeted on the automation, scalability and performance of the solutions. Accordingly, the following objectives are specified for the data protection service:

(a) to be able to automatically deploy the Secure Cloud Storage solution on a hybrid multi-cloud environment within the Fed4FIRE Cloud facilities,

(b) to provide a method for provisioning encrypted volumes to VMs deployed on the multiple cloud platforms, and

(c) to assess and automate the methods for transferring secure volumes from one VM to another and from one cloud platform to another.

The host and application protection service included the following objectives:

(a) to be able to automatically deploy the Intelligent Protection components in a hybrid environment of public cloud, private cloud and physical servers,

(b) to automatically provision Intelligent Protection agents and protect IT assets on a large number of virtual machines emulating real-world scenarios, and

(c) to develop a blue print of a managed security service to apply virtual patches to proprietary applications deployed on multiple virtual machines.

The following parameters of the data protection service were measured:

(a) the complexity of the workflow required to achieve the deployment and configuration of the SCS service components on multiple cloud platforms,

(b) the time-scales involved in the workflow,

(c) the mix and heterogeneity of the cloud platforms,

(d) the time for the creation of data volumes of different sizes on multiple clouds of Fed4FIRE,

(e) the number of volumes attached to the VMs deployed on these clouds, and

(f) the time for volumes to be encrypted by employing the full disk encryption technique.

The measurable parameters of the application protection services are as follows:

(a) the feasibility and speed of deployment of the agents in the hybrid multi-cloud environment,

(b) the maturity of the automation of workflow required to achieve the agent deployment and configuration on multiple cloud platforms,

(c) the time for deployment, configuration and security management of the agents in up to 50 virtual machines on multiple cloud platforms, and

(d) the performance associated with the security management function such as the automatic security patching for standard applications such as databases, Email servers, Apache web servers, and various Windows and Unix operating systems.

5 Architecture and Experiment Set up

5.1 Data Protection

The architecture of the SCS is composed of three main components:

- (a) Key Management Server (KMS): this is a key server self-hosted or hosted by a trusted Service Provider. It is the only place where the encryption keys are stored persistently.
- (b) Web Console: this is a console that allows SCS administrators to review and approve pending key requests, check device status and integrity, set up policies, manage devices, check reports and logs, and manage user accounts.
- (c) Agent: this is software installed on a virtual machine that communicates with the KMS and performs the actual encryption.

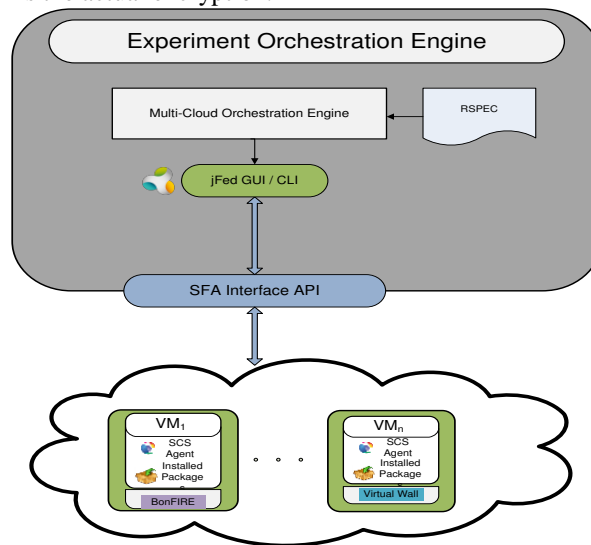


Fig. 1. Detailed Design of CMS

The experiment uses the federated tools and interfaces provided by Fed4FIRE to discover, reserve and provision the required resources. The resources required for this experiment are VMs and persistent storage blocks or volumes that can be attached to those VMs. The experiments use Virtual Wall 1 & Virtual Wall 2 and BonFIRE as a multi-cloud infrastructure. To help with the deployment and provisioning of the SCS agent on the Fed4FIRE infrastructure VMs, Puppet is used as the configuration management service (CMS). This allows contextualization of the VMs to fulfil all the dependencies of the SCS agent and also to deploy the appropriate version and build of the SCS agent on the target VM's. To conduct the experiment in a multi-cloud environment, the jFed tool is used, which allow end-users to provision and manage experiments via a GUI and CLI. The design of the CMS in this context is given in **Fig. 1**. The experiments are specified and orchestrated automatically by constructing a RSPEC file provided to the jFed tool. The Key Management Server and the SCS Web Console were hosted at BT on a dedicated facility and offered as a cloud-based service to the experiment.

5.2 Host & Application Protection

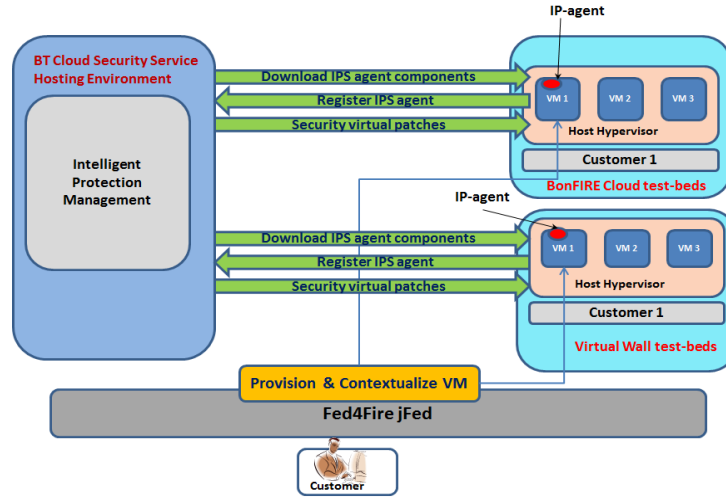


Fig. 2. Intelligent Protection provisioning architecture in Fed4FIRE

The Intelligent Protection architecture comprises of the following to support on-boarding and enablement of application protection elements in the Fed4FIRE infrastructure:

- (a) Intelligent Protection Management Server: This has a web interface to manage the Intelligent Protection (IP) agents that protect the cloud nodes. The components of the server include: a dashboard, Alerts, Events, Computers, Policies and Administration.
- (b) IP-agent: This is agent software that can be either manually deployed or deployed using deployment scripts on the VMs to be protected.
- (c) Deployment Scripts: This is a script, which on execution on a VM, automatically downloads the agent from the server, then installs and registers the agent with the server. This setup requires one or more nodes in the Fed4FIRE infrastructure to have connectivity to the Intelligent Protection Management Server and an RSpec that is configured for automatic provisioning of an agent. An RSpec document is prepared that contains nodes of Virtual Wall 1 & 2 infrastructures and the deployment script for automatic provisioning of the agents.

6 Experiment Results

6.1 SCS Performance experiments

The objective of this experiment is to analyze the provisioning of SCS agents on VMs and assess the data volume encryption overhead on I/O bandwidth and I/O latency. The tests compare file I/O performance of VMs before and after the data volumes are encrypted with the data protection agent using the AES-256 algorithm. The tests are

run on VMs provisioned on the BonFIRE and Virtual Wall infrastructures and benchmarked with the read and write I/O operations. The effects on bandwidth and latency of these I/O operations are measured with the help of the FIO storage benchmarking tool for both the unencrypted and encrypted volumes [13]. The volumes used in the BonFIRE and Virtual Wall infrastructures are 1 GB in size.




Fig. 3. Encryption overhead on I/O Bandwidth




Fig. 4. Encryption overhead on I/O Latency

The FIO tool is configured to use whole volumes in the read/write tests i.e. 1 GB of data is written and read in all the tests. All the volumes are formatted with the ext4 filesystem with 4K block sizes. **Fig. 3** and **Fig. 4** show the encryption overhead on I/O bandwidth and I/O latency respectively. The measurements shown in **Fig. 3** and **Fig. 4** demonstrate that the data protection solution enables data security in clouds while maintaining reasonable performance for typical filesystem and database workloads.

6.2 Scalability experiments

The aim of this experiment is to evaluate the performance of automatic provisioning of protection to the VMs, in a scaled environment. The provisioning of protection includes downloading of the agent from the Intelligent Protection Management Server, installing

the agent on the VMs, registering the agent with the server and updating the agent with the security policies. To perform the scalability evaluation, 50 VMs (40 VMs on Virtual Wall 2 and 10 VMs on Virtual Wall 1) are simultaneously created in a multi-cloud environment and the Intelligent Protection Server is hosted in the BT cloud.




Fig. 5. Agent Deployment Performance in a Scaled Federated cloud environment

VM nodes 1-40 are created on Virtual Wall 2 and VM nodes 41-50 are created on Virtual Wall 1. The time durations of booting the VMs and provisioning the agents on the infrastructure is recorded. The observation in **Fig. 5** shows that the Intelligent Protection agents are deployed and configured in less than 1 minute after the VMs are booted on both infrastructures. Although the VM provisioning performance is different in the two infrastructures, it was observed that the Intelligent Protection agent deployment was consistent even in the scaled environment of 50 VMs. These time scales were obtained for the automatic provisioning workflow without any manual intervention from the user or the administrator of the targeted Cloud platform.

6.3 Feedback to the Fed4FIRE Infrastructure

The various characteristics of the Fed4FIRE infrastructure such as scale, heterogeneity, multi-cloud, and federation are potentially valuable to the experiments and a detailed feedback was provided based on our experience of the experiment. The feedback also included some of the limitations of the existing infrastructure which were the obstacles we experienced while performing the experiments. These were:

- (a) The SCS software can only run on recent versions of Ubuntu, CentOS or Windows operating systems, none of which were readily available and required templates to be provisioned on request.
- (b) The Intelligent Protection Agent is only supported for Windows and unix, but for widely used variants of these operating systems, such as Windows 32/64, Red Hat Enterprise Linux 6.0/5.0, SUSE Linux 11, Ubuntu 12.04/10.04, Amazon Linux, Oracle RedHat Enterprise Linux 6/5, Solaris 9/10/11 and HP-UX 11i v3 IA-64. However the Fed4FIRE infrastructure only supports Ubuntu 12.04/10.04 and no other widely used operating system for which the agent is available.

- (c) The Virtual Wall infrastructures do not provide facilities for creating and attaching data volumes with virtual machines, which prevented the data protection experiment from performing volume encryption and confined it to file encryption.
- (d) Some Fed4FIRE infrastructures do not provide an option to create VM templates.

7 Conclusion and Future Work

The paper describes experimentation with a data protection service and a host & application protection service offered as a value added security service in Fed4FIRE environment and presents the experimental results. The experiments verify the provisioning and automation workflow and the results validate the performance and scalability of these services in the Fed4FIRE which confirms the capability of these services to perform in a large scale multi-cloud and federated cloud environments. The obstacles experienced for the experiments on the Fed4FIRE infrastructure were provided as feedback to the operators. As a future work, we intend to perform more experiments after enhancements are made to the infrastructure.

Acknowledgement

This work has been partially supported by the EU within the 7th Framework Programme under contract ICT-318389 – Fed4FIRE (Federation for FIRE).

References

1. P. Mell and T. Grance, ‘The NIST Definition of Cloud Computing’, <https://nist.gov/publications/nistpubs/800-145/SP800-145.pdf>, Sep. 2011.
2. A. S. Ibrahim, J. H. Hamlyn-harris, and J. Grundy, *Emerging Security Challenges of Cloud Virtual Infrastructure*. 2010.
3. K. Hashizume, D. G. Rosado, E. Fernández-Medina, and E. B. Fernandez, ‘An analysis of security issues for cloud computing’, *J. Internet Serv. Appl.*, vol. 4, no. 1, pp. 1–13, 2013.
4. ‘Dropbox’. [Online]. Available: <https://dropbox.com/>. [Accessed: 13-Mar-2015].
5. ‘Google drive’. [Online]. Available: <https://drive.google.com/>. [Accessed: 13-Mar-2015].
6. ‘Microsoft Sky Drive’. [Online]. Available: <https://skydrive.live.com/>. [Accessed: 13-Mar-2015].
7. ‘Boxcryptor | Encryption for cloud storage | Window, Mac, Android, iOS | boxcryptor.com’. [Online]. Available: <https://www.boxcryptor.com/>. [Accessed: 13-Mar-2015].
8. ‘TERENA/CloudDrive · GitHub’. [Online]. Available: <https://github.com/TERENA/CloudDrive/>. [Accessed: 13-Mar-2015].
9. ‘Fed4Fire’. [Online]. Available: <http://www.fed4fire.eu/>. [Accessed: 13-Mar-2015].
10. J. Daniel, T. Dimitrakos, F. El-Moussa, G. Ducatel, P. Pawar, and A. Sajjad, ‘Seamless Enablement of Intelligent Protection for Enterprise Cloud Applications through Service Store’, 2014, pp. 1021–1026.
11. ‘Amazon Web Services’, 22-May-2013. [Online]. Available: <http://aws.amazon.com/>. [Accessed: 22-May-2013].
12. ‘vCloud’. [Online]. Available: <http://vcloud.vmware.com/>. [Accessed: 13-Mar-2015].
13. ‘Flexible I/O tester · GitHub’. [Online]. Available: <https://github.com/axboe/fio/>. [Accessed: 13-Mar-2015].