D1.3.2
Cloud Computing: Business Impact Analysis

Project number:	257243
Project acronym:	TClouds
Project title:	Trustworthy Clouds - Privacy and Resilience for Internet-scale Critical Infrastructure
Start date of the project:	1st October, 2010
Duration:	36 months
Programme:	FP7 IP

Deliverable type:	Report
Deliverable reference number:	ICT-257243 / D1.3.2 / 1.0
Activity and Work package contributing to deliverable:	Activity 1 / WP 1.3
Due date:	September 2012 – M24
Actual submission date:	26th October, 2012

Responsible organisation:	IBM
Editor:	Christian Cachin
Dissemination level:	Public
Revision:	1.0

Abstract:	This report summarizes the services of a representative selection of IaaS cloud providers. Furthermore, it explores the impact of trust-related aspects on cloud-computing business models, specifically in terms of security and resilience features.
Keywords:	IaaS, data governance, isolation failure, insider fraud, management-interface attacks, secure deletion, data protection, encryption, resilience, compliance, liability, transparency, accountability.

Editor

Christian Cachin (IBM)

Contributors

Christian Cachin, Kristiyan Haralambiev, Elmar Husmann, Nikola Knežević (IBM)

Mina Deng, Ya Liu (PHI)

Disclaimer
This work was partially supported by the European Commission through the FP7-ICT program under project TClouds, number 257243.

Executive Summary

Infrastructure clouds affect business models, processes, and vendor-client relations. This work explores the business impact of engaging in trustworthy cloud computing activities, with special focus on security and resilience.

This report makes two main contributions. The first consists of summarizing and analyzing the services of a representative selection of IaaS cloud providers. A particular focus is put on security features, resilience aspects, and regulatory compliance.

The second contribution explores the impact of trust-related aspects on cloud-computing business models, specifically in terms of security and resilience features. The covered technical features include governance, isolation failure, insider fraud, management-interface attacks, secure deletion, data protection and encryption, (data and computation) resilience, as well as legal features, covering compliance, liability, transparency, and accountability. Problems, technical approaches, and the corresponding business impact are described.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 TClouds — Trustworthy Clouds

TClouds aims to develop *trustworthy* Internet-scale cloud services, providing computing, network, and storage resources over the Internet. Existing cloud computing services are today generally not trusted for running *critical infrastructure*, which may range from business-critical tasks of large companies to mission-critical tasks for the society as a whole. The latter includes water, electricity, fuel, and food supply chains. TClouds focuses on power grids and electricity management and on patient-centric health-care systems as its main applications.

The TClouds project identifies and addresses legal implications and business opportunities of using infrastructure clouds, assesses security, privacy, and resilience aspects of cloud computing and contributes to building a regulatory framework enabling resilient and privacy-enhanced cloud infrastructure.

The main body of work in TClouds defines an architecture and prototype systems for securing infrastructure clouds, by providing security enhancements that can be deployed on top of commodity infrastructure clouds (as a cloud-of-clouds) and by assessing the resilience, privacy, and security extensions of existing clouds.

Furthermore, TClouds provides resilient middleware for adaptive security using a cloud-of-clouds, which is not dependent on any single cloud provider. This feature of the TClouds platform will provide tolerance and adaptability to mitigate security incidents and unstable operating conditions for a range of applications running on a clouds-of-clouds.

1.2 Activity 1 — Legal and Business Foundations for Cross-Border Computing

The scope of Activity 1 is to identify requirements and boundaries for cloud computing. The Activity aims at providing a guidance framework to address both legal requirements and business interests in cross-border infrastructure clouds. Based on the expertise and input from users and stakeholders the activity researches relevant interests, drivers and obstacles for the use of cloud computing services for privacy-sensitive and business-critical applications — with a focus on the implication of cross-border cloud deployment.

Furthermore, an analysis of the European legal framework for data protection and data security identifies the regulatory foundation for cloud computing and leads to an investigation of its privacy impact. The Activity addresses the business impact of cloud computing as well as the accompanying privacy and security concerns. Requirements derived from this tense relationship of business benefits and regulatory boundaries will be mapped to organisational, contractual and technical measures and enablers.

1.3 Workpackage 1.3 — Business Impact of and Business Models for Infrastructure Clouds

WP1.3 addresses the effects of a cloud infrastructure on business models, processes, and vendor-client relations. It places a special focus on security and resilience features in cloud computing offers and addresses the benchmark scenarios of smart lighting and home healthcare.

1.4 Deliverable 1.3.2 — Cloud Computing: Business Impact Analysis

Overview. This report describes the impact of cloud computing models on businesses. It specifically focuses on infrastructure-cloud services, also called Infrastructure-as-a-Service Clouds (IaaS) [MG11]. In IaaS, customers obtain virtual computing resources on the infrastructure level, be it storage devices, network infrastructure, or bare-bones virtual machines. There are no sophisticated interfaces for end-users. Usage is typically billed per amount of the consumed resource.

Many concerns exist around security and resilience of cloud computing. These factors clearly affect the business models of customers. Many excellent studies have been developed so far that address this field; they analyze existing problems, show areas of concern, and identify requirements for future security technology. The most prominent and most widely accepted frameworks of this kind have been produced by industry consortia and government agencies [Clo11, ENI09, JG11a, BGPCV12].

Finally, let us note that this report does try to address questions of business advantages in relation to trustworthiness of clouds (or, more specifically, to make predictions on whether different aspects of cloud's trustworthiness will influence various business to adopt a cloud). This report explores the impact of security and resilience features on cloud-computing business models, listing possible risks, solutions and outlooks on mitigating these risks. The next deliverable (deliverable D1.3.3) in this work package, will offer an analysis on the potential business value and impact of such solutions, based on two demonstration scenarios.

Structure. This report makes two main contributions. In Chapter 2, the services of a representative selection of IaaS cloud providers are analyzed and summarized. A particular focus is put on security features, resilience aspects, and regulatory compliance.

Chapter 3 explores the impact of *trustworthiness* aspects in cloud computing *on* cloud-computing business models, specifically in terms of security and resilience features. The covered technical features include governance, isolation failure, insider fraud, management-interface attacks, secure deletion, data protection and encryption, (data and computation) resilience, as well as legal features, covering compliance, liability, transparency, and accountability.

Deviation from Workplan. This deliverable aligns with the DoW/Annex I, Version 2. Compared to D1.3.2 as foreseen in the original DoW/Annex I, the scope if D1.3.2 is reduced to analyzing economic implications and business models for infrastructure clouds only. The description and analysis of of the potential business value and impact of the solutions developed in the two benchmark scenarios is deferred to D1.3.3 ("Cloud Computing: Business Impact Analysis of the benchmark use cases").

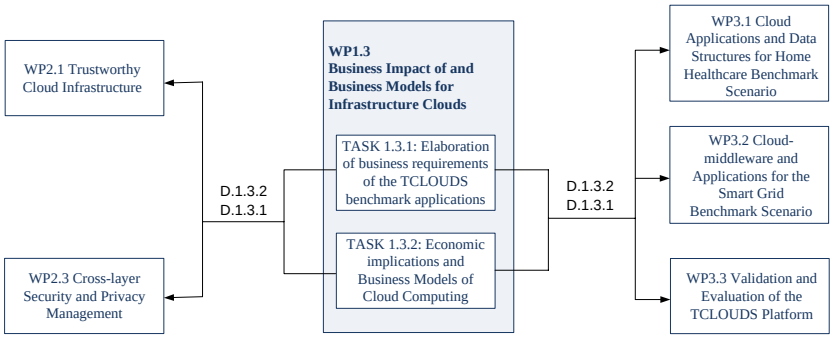

Figure 1.1: Graphical structure of WP1.3 and relations to other workpackages.

Target Audience. This deliverable aims at professionals, business executives, and researchers in the IT industry.

Relation to Other Deliverables. Figure 1.1 illustrates WP1.3 and its relation to other workpackages according to the DoW/Annex I (specifically, this figure reflects the structure in Annex I, Version 2).

This document, deliverable D1.3.2, extends D1.3.1 and elaborates on particular business factors identified in D1.3.1. Special consideration is given to trustworthiness, that is, security- and resilience-enabling technologies. Mechanisms developed in WP2.1 and WP2.3 can play a significant role in building trustworthy clouds, as detailed in Chapter 3 of this report.

Chapter 2

Comparison

Chapter Authors:
Elmar Husmann and Nikola Knežević (IBM)
Mina Deng and Ya Liu (PHI)

2.1 Introduction

In IaaS cloud computing, customers obtain virtual computing resources on the infrastructure level, be it storage devices, network infrastructure, or bare-bones virtual machines. There are no sophisticated interfaces for end-users. Usage is typically billed for units of the consumed resource, per time and/or per volume.

This chapter presents a representative selection of IaaS cloud providers and summarizes the features of the offered services. Apart form the basic service models, the comparison includes also security factors, resilience aspects, and regulatory compliance.

2.2 Provider features

IaaS providers provide the most basic IT services including servers, networking, and storage, on a utility model with optional managed hosting services [MG11]. While there are many benefits of adopting the infrastructure offered by a cloud-service provider (CSP), the applicability of these depends on the nature of the company need. With a growing list of CSPs, various types of infrastructure as a service have made the decision for customers to be very complex. To investigate the current IaaS vendors, the approach of choosing a cloud provider must be based on an analysis of the offered features, functional and non-functional.

In the following sections, key features of some representative cloud providers are compared against each other. The comparison uses a European perspective and, where applicable, investigates features relevant for privacy-sensitive healthcare-related data.

An estimate of response times have been obtained via the Global Provider View of CloudSleuth (`https://cloudsleuth.net/global-provider-view`). This service runs test transactions from Compuware's servers in the Gomez Performance Network and uses the same deployed target application to monitor the response time. Specifically, the response time concerns loading two web pages hosted on the respective cloud provider. One page is full of item descriptions and associated pictures and the other page consists of a single large picture.

2.2.1 Self-hosting oriented providers

Self-hosting oriented vendors are those who focus on provide various specialized cloud infrastructure with options on computing capability, storage and network. Most of them provide public cloud for not only enterprise but individuals.

Amazon Web Services

Highlights

- AWS is the leader in the IaaS market with high market awareness and good reputation.
- The infrastructure options are various from high-performance computing to large data applications.
- Various templates and pre-configured images are available.
- Extensive APIs for developers to customized management of their cloud service.
- AWS has obtained many security compliance related certifications including HIPAA.
- Quick response time in EU, see Figure 2.1.
- Outstanding geographic strategy covering US, Europe, and Asia-Pacific.
- Highly customized monitoring and auditing methods are provided with full API access.
- Storage encryption is provided as a service.

Critical points

- Low SLAs compared to other vendors. SLAs do not cover all the services (no EBS).
- Self-service: No managed network, security, and hosting services are provided.
- As many value added services has separated pricing mode. Pricing strategy or sales strategy is difficult to be understood by normal customers, except for IT experts.

Figure 2.1: Amazon EC2 response-time measurement

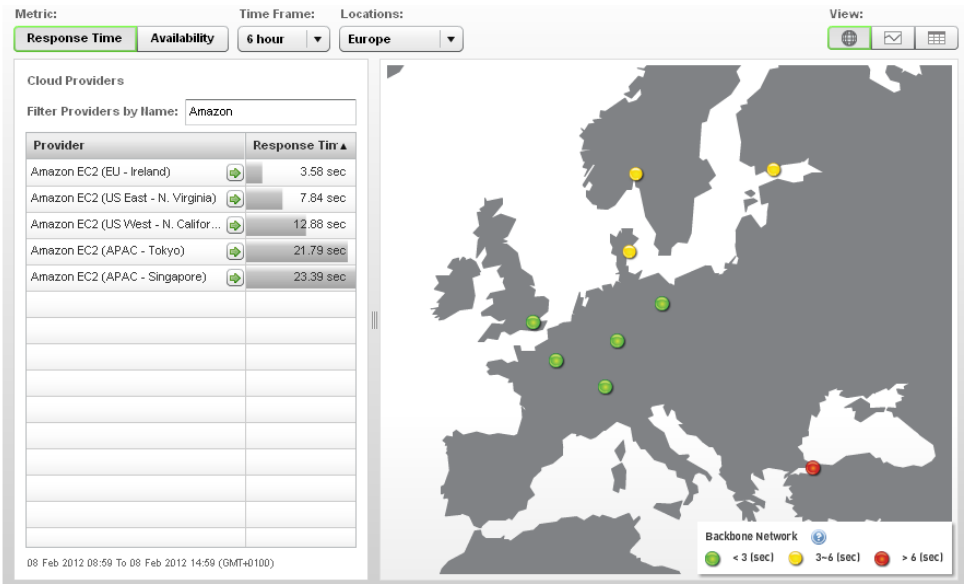

Microsoft Azure

Highlights

- Quick response time in EU, see Figure 2.2.
- Extensive APIs for developers to customized management of their cloud service.
- Easy integration with Microsoft technology like SharePoint, MS SQL, Windows Live.
- Patch management is provided on the platform where applications are running over.
- Outstanding geographic strategy covering US, Europe, and Asia-Pacific.
- Highly customized monitoring and auditing methods are provided with full API access.

Critical points

- Low flexibility of platform choice, only Windows Azure guest operating system is available which is running on top of Windows 2008.
- Low SLAs compared to other vendors.
- No server side encryption is provided.
- Microsoft applied a non-mainstream Hyper-V as Hypervisor; only Windows platform available, may lead to vendor lock-in.

Figure 2.2: Windows Azure response-time measurement

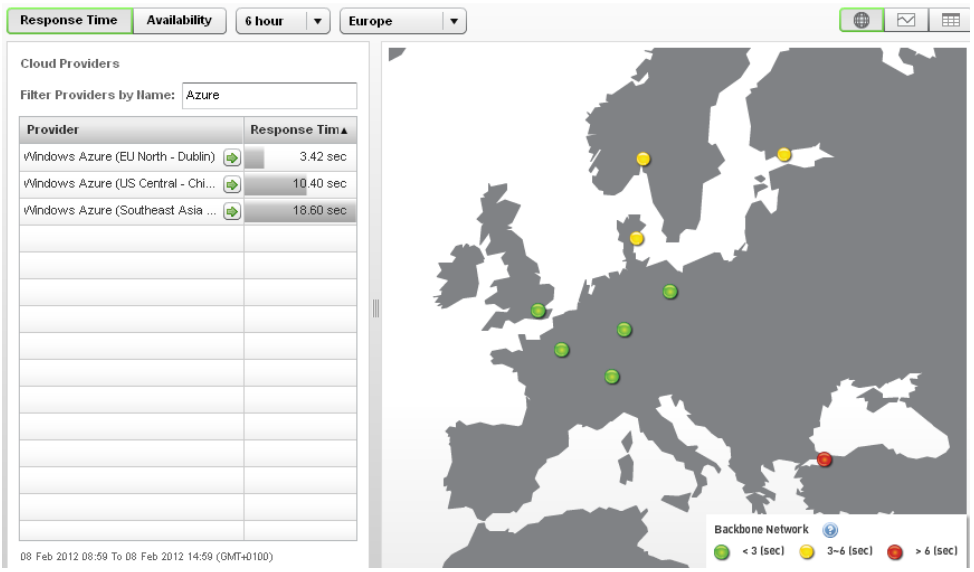

GoGrid

Highlights

- Excellent SLAs for all services.
- Successfully integrated managed hosting with self-service IaaS.
- GoGrid Exchange allows partners to distribute preconfigured server images to the GoGrid community.

Critical points

- Only preconfigured, standard OS can be chosen from the template list offered by GoGrid. No VM import by users.
- Slow response time in EU, which will degrade the user experience and service quality, see Figure 2.3.
- Does not offer any API support for Monitoring. Third-party applications needed.
- No encryption offered as managed security service for storage.
- Narrow geographic strategy (only in US).

Figure 2.3: GoGrid response-time measurement

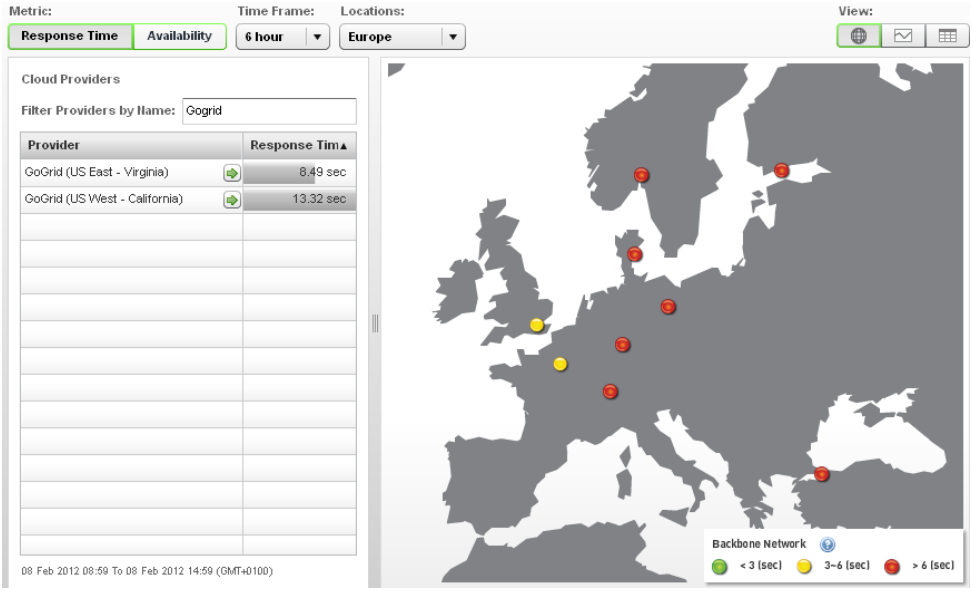

Rackspace

Highlights

- Rackspace supports and develops OpenStack, the open-source cloud application which will increase the third-party application support.
- Excellent SLAs for all services.
- Extensive API access.
- Provides patching management for the platform where customer's applications are running above.

Critical points

- Service from Rackspace is oriented towards hosting use case.
- No managed service for network management and monitoring.
- Only standard OS can be chosen from the template list offered by Rackspace. No VM import by users.

Figure 2.4: Rackspace response-time measurement

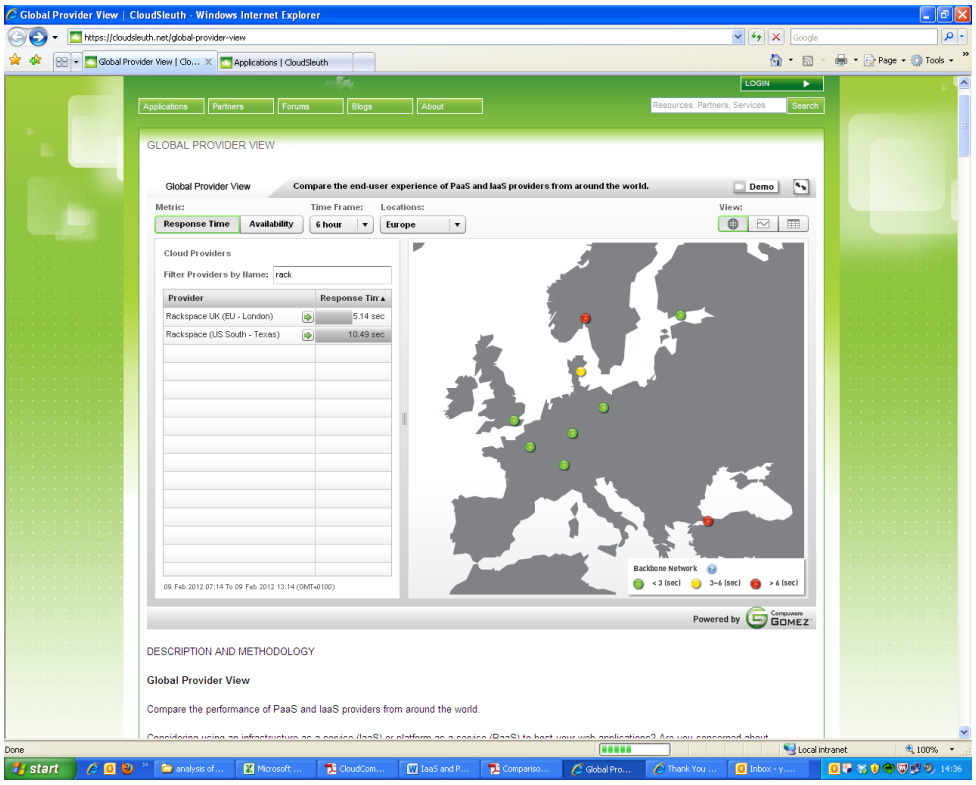

2.2.2 Managed-service-oriented providers

Vendors offering managed services provide managed cloud hosting as their major service including cloud network management, data backup, load balancing monitoring, etc. Their target customers are enterprises.

DataPipe

Highlights

- Outstanding geographic strategy covering the US, Europe, and Asia-Pacific, including China.
- Provides managed cloud service which uses AWS as infrastructure. Easy conjunction with Amazon Web Service.
- Managed Monitoring provides various monitoring objects from computing performance to network health.

Critical points

- Lack of self-hosting service, because of emphasis on AWS.
- Poor technical supports for public cloud developers.
- Only standard OS can be chosen from the template list. No VM import by users.

Carpathia

Highlights

- Emphasis on compliance hosting for special use cases (medical and government), satisfying various industry compliance requirement including HIPAA.

Critical points

- Managed service, not self-hosting, which limits customers choice.
- Lack of information can be found if the user wants to do investigation on its service capability.
- Only standard OS can be chosen from the template list. No VM import by users.

2.2.3 Hybrid-model service providers

The following providers offer balanced self-hosting and managed-service models. Their focus are enterprise customers .

IBM SmartCloud Enterprise.

Highlights

- High market awareness and good reputation for its IT service. Is built on top of IBM's private IaaS technology.
- Various value-added services and side-products including DB2 and Tivoli software can be easily integrated into SmartCloud.
- Various technical documents and security related reports can be easily found for normal users and developers. Full API support is provided.
- Outstanding geographic strategy covering the US, Europe, and Asia-Pacific.
- Has obtained many security compliance related certifications including HIPAA.

- Provides patching management for public and private images.
- Pre-configured Linux images can be imported.

Critical points
- Low SLAs compared to other vendors
- Low pace and slow reaction in developing public IaaS market, which degrade its performance in public IaaS market.

OpSource

Highlights
- Excellent SLAs for network uptime and server availability.
- Emphasis on compatibility and avoiding vendor lock-in, by offering cancellation and data return service.
- Managed hosting covers various fields including backup, disaster recovery, network management, etc.
- Enables user to create its own image to upload, which makes it more compatible for different use cases.
- Has obtained many security compliance related certifications including HIPAA.

Critical points
- Slow response time in EU, which will degrade the user experience and service quality.
- OpSource service is more managed-hosting oriented, other than offering high performance capability infrastructure.

CSC

Highlights
- Outstanding geographic strategy covering US, Europe, and Asia-Pacific, including Australia.
- Pre-configured images can be imported.
- Provides integrated development environment (IDE), and network simulator for developer better experience.
- Snapshot is provided for data restore which is not offered by every VMware based providers.
- Various managed services are provided for customers.

Critical points
- Low SLAs for its service.
- Lack of information introducing its services in technical respect.

AT&T

Highlights
- Provides medical image for medical use case with storage encryption and other managed security service.
- Erasure coding is a feature provided by its cloud storage service for data recovery purpose.

Critical points
- Low SLAs compared to other vendors.

- Only standard OS can be chosen from the template list offered by AT&T. VM import by users is not available.
- Poor third-party integration due to its poor API support.

Savvis

Highlights
- Customized web management portal is offered for better hosting experience.
- Excellent SLAs for all services.
- Pre-configured images can be imported.
- Outstanding geographic strategy covering US, Europe, and Asia-Pacific.
- Various managed services are provided for customers.

Critical points
- Little technical information, e.g. technical documents, can be found online for in-depth investigation.

2.3 Comparison

On the following pages the capabilities and features of the IaaS cloud providers listed before are compared side-by-side.

In order to reflect the most important technical business aspects of those providers, the following dimensions are created as criteria to represent the level of service, flexibility, security, and support offered to customers:

Features: Reflects the computing capability, the flexibility of configuring a server, and the variety of data storage functions provided by a vendor.

Data management: Addresses how stored data is maintained, including the location of data center, the policy of data backup and restore, and data segmentation.

Security and privacy: Security of cloud computing is always a significant aspect needs to be considered when choosing a provider, including network security, identity management, access control and cryptographic support.

Industry-regulation compliance: Compliance- and security-related certifications are important factors for special use cases.

Support: The extent of the available technical supports measures, including documents, API descriptions, and live-support channels; these are essential for future integration and expansion.

Other elements: Further features and relevant details, like the perceived response time for servers (again obtained via the Global Provider View of CloudSleuth), cloud-related security breaches, and sample customers.

Table 2.1: General features overview

	Virt. Arch	VM Import	Standard Platforms	Monitoring approach	Scaling up/down	Load balancing
Amazon	Xen	Yes	UNIX-like: Red Hat Linux, Oracle Linux, SUSE Linux, Amazon Linux, Ubuntu, Fedora, Gentoo Linux, Debian Windows-based: Windows Server 2003/2008	Amazon CloudWatch provides: resource utilization, application performance, and operational health on AWS Resources. It supports Custom Metrics, Alarm, API	Auto scaling on demand	Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances
Windows Azure (IaaS)	Hyper-V	No	Windows-based: Windows Azure guest operating system	Service Dashboard, Event Monitors, User Perspective Monitors, Performance Monitors. Administrative operations are audited, audit trail can be viewed.	Auto scaling on demand	Network load balancing and failover are handled automatically
GoGrid	Xen	No	UNIX-like: CentOS, Red Hat Linux, ubuntu, Debian Windows-based: Windows Server 2003/2008	Needs third-party apps: Cloudkick, Cacti, etc.	Auto scaling on demand	GoGrid includes fully integrated and redundant F5 load balancers
Rackspace Cloud	Xen	No	UNIX-like: CentOS, Ubuntu, Debian, Gentoo, Fedora, Red hat, Arch Linux Windows-based: Windows Server 2003/2008	Needs third-party apps: Cloudkick, etc.	Auto scaling on demand	Customized load balancer
IBM SmartCloud Ent.	VMware	Yes	UNIX-like: Red Hat Linux, SUSE Linux Windows-based: Windows Server 2003/2008	IBM Tivoli® Monitoring provides a user interface that includes graphical views of performance and availability data. Enables problem detection and recovery of potential performance and availability problems of IBM SmartCloud instances	Reserved capacity	IBM Tivoli Service Automation Manager enables the definition of rules to automatically distribute the workload among VMs
OpSource	VMware	Yes	UNIX-like: CentOS, Red Hat Linux, Ubuntu Windows-based: Windows Server 2003/2008	Auditing and reporting on all activities. System monitoring is offered as managed hosting	Auto scaling on demand	Supports load-balancing and port translation across multiple virtual servers
AT&T	VMware	No	UNIX-like: CentOS, Ubuntu, Debian, Fedora, Red hat, FreeBSD Windows-based: Windows Server 2003/2008 AT&T Medical Image	Customer can create customized reporting for hardware and OS monitoring, system health monitoring, applications monitoring	Auto scaling on demand	Local load balancer policy with a virtual IP address to distribute traffic among multiple virtual machines
Savvis	VMware	Yes	Savvis Symphony Dedicated: managed instance Savvis Symphony Open: customized image	VMware HA provides high-availability monitoring and automated failover. The system also performs active monitoring of errors or failures at the physical and virtual server level	Auto scaling on demand	Load balancing is offered as a service
CSC	VMware	Yes	VMware-based platforms Customer can upload any customized image	Hardware, system monitoring and operating logs	Auto scaling on demand	Managed Hosting Network Services provide geographic and local load-balancing services
DataPipe	Xen	No	UNIX-like: CentOS, FreeBSD, Red Hat, ubuntu Windows-based: Windows Server 2003/2008	Monitoring services: Content Checking, Database Monitoring, ICMP - HTTP - FTP - SMTP - SSH, Vital Services, CPU Utilization, Memory Utilization, Disk Storage Capacity, Running Processes or Applications	Auto scaling on demand	High availability load balancing is provided as managed hosting
Carpathia	Xen	No	Xen-based platforms	Managed monitoring service	Auto scaling on demand	
Google App Engine	Not available	No	Java Runtime Environment Python Runtime Environment Go Runtime Environment	GAE monitors the serving status of Java, Python and Go. As well as the APIs status. Monitoring includes Latency, error rate, throughput, CPU. It also provides APIs for customization	Auto scaling on demand	
Windows Azure (PaaS)	Hyper-V	No	.NET framework Node.js Java, php	Service Dashboard, Event Monitors, User Perspective Monitors, Performance Monitors. Microsoft administrative operations are audited. The audit trail can be viewed to determine the history of changes.	Auto scaling on demand	

Table 2.2: Storage features overview

	Integrated DB	Separated DB Service	Cloud Storage
Amazon	Amazon EC2 Relational Database AMIs (MS SQL 2003/2008, MySQL 5, Oracle 10g/11g, IBM DB2, Postgre SQL, SYBASE)	Amazon RDS (MySQL or Oracle) Amazon SimpleDB (non-relational, but scalable) Amazon DynamoDB (non-relational)	Amazon S3 (Simple Storage Service): file storage Amazon EBS (Elastic Block Storage): block storage
Windows Azure (IaaS)		SQL Azure Azure table storage	BLOB (Binary Large Object) storage Windows Azure Drive:block-based
GoGrid	MS SQL 2003/2008		GoGrid Cloud Storage: file storage
Rackspace Cloud	MS SQL 2008		Rackspace Cloud Files: file storage
IBM SmartCloud Ent.	IBM DB2® Enterprise Developer Edition IBM DB2 Express-C		IBM SmartCloud object storage: object-based IBM SmartCloud Archive: file-based
OpSource	MS SQL 2000/2005/2008, MS SharePoint 2010, Oracle 9i, 10g, 11g, MySQL 4.x, 5.x		OpSource Cloud Files: file storage
AT&T	MS SQL 2008, MySQL, Cloudera's Distribution for Hadoop		AT&T Synaptic Storage: file storage
Savvis		Savvis Symphony Database	Savvis' enterprise storage: file storage
CSC DataPipe	Oracle, IMB DB2, MS SQL, MySQL, PostgreSQL, MongoDB, Oracle		SAN (Storage Area Network): block storage NAS (Network Attached Storage): file storage
Carpathia	MS SQL 2008		Managed storage solution: file storage
Google App Engine		The App Engine datastore is a schemaless object datastore, with a query engine and atomic transactions, and provide standard API for Java, Python and Go	Google Cloud Storage : file storage
Windows Azure (PaaS)		SQL Azure Azure table storage	BLOB (Binary Large Object) storage Windows Azure Drive: block-based

Table 2.3: Data handling overview

	Data center locations	Data backup	Disaster recovery	Data segmentation
Amazon	North America -N. California, N. Virginia, Oregon South America - Sao Paulo Europe - Ireland Asia-Pacific - Singapore, Tokyo	Amazon S3 for backup	Data in Amazon S3, Amazon SimpleDB and Amazon EBS is redundantly stored in multiple physical locations to recovery in the event of a natural or man-made disaster	Physically, user can choose multiple geographic regions. Logically, hypervisor-based isolation
Windows Azure (IaaS)	North America - Virginia, Washington Europe - Dublin, Amsterdam, Asia-Pacific - Singapore, Hong Kong, Japan		Windows Azure Blobs, Tables are replicated three times in the same data center against hardware failure at no additional cost.	Hypervisor-based isolation of root VM from guest VMs and guest VMs from one another.
GoGrid	North America -San Francisco, California and Ashburn, Virginia	Data backup can be provided by GoGrid community	Disaster recovery can be provided by GoGrid community	Logically, hypervisor-based isolation
Rackspace Cloud	North America - Dallas, Chicago Europe - UK	Rackspace Server Backup for backup and data restore		
IBM SmartCloud Ent.	North America - Raleigh, Boulder, Toronto Europe - Germany Asia-Pacific - Japan, Singapore	IBM SmartCloud object storage can be used as backup for other storage	Geo-replication for disaster protection	Hypervisor-based isolation
OpSource	North America - Ashburn VA, San Jose CA Europe - London, Paris	OpSource Backup Services	OpSource offers a fully-managed Disaster Recovery option that provides a path to recovery in the event of a natural or man-made disaster affecting the primary data center	Logically, hypervisor-based isolation
AT&T	North America - Atlanta, Annapolis, Md. and the New York/New Jersey Europe - UK, Asia-Pacific - Japan, Hong Kong		Data is protected using erasure coding which is a software-based data protection scheme that allows for data recovery in the event of hardware failures	Hypervisor-based isolation
Savvis	North America (many locations) Europe - UK	Savvis provides scheduled backup	Savvis provides on-demand restore	Hypervisor-based isolation
CSC	Asia-Pacific - Singapore, Japan North America - Chicago, Newark, Chantilly Europe - Copenhagen, UK Asia-Pacific - Sydney	Data backup and restore as an option for customers	Quick restore of large datasets via snapshots	Hypervisor isolation for network adapters
DataPipe	North America - San Jose, New Jersey Europe - London Asia-Pacific - Hong Kong, Shanghai	Datapipes backup and restoration services secure critical data		
Carpathia	North America-Ashburn, Dulles, Phoenix, Los Angeles, Harrisonburg	Managed backup solutions	Replicated data ensures you have a current off-site copy of mission critical information at a fraction of the cost of traditional disaster recovery solutions when application becomes either corrupted or lost, user can restore datastore from backup	Hypervisor-based isolation
Google App Engine	North America - South Carolina, Iowa, Georgia, Oklahoma Europe - Finland, Belgium Asia-Pacific - Hong Kong, Singapore	User can use Datastore Admin tab of the Admin Console to backup entities of selected kinds and when needed restore from a selected backup	Windows Azure Blobs, Tables are replicated three times in the same data center against hardware failure at no additional cost.	Sandbox isolates application in its own secure, reliable environment.
Windows Azure (PaaS)	North America - Virginia, Washington Europe - Dublin, Amsterdam, Asia-Pacific - Singapore, Hong Kong, Japan			Hypervisor-based isolation of root VM from guest VMs and guest VMs from one another.

Table 2.4: Security features overview

	Identity and Access Management	Intrusion detection/prevention	Firewall	Patch management	Data encryption
Amazon	Role-based permissions allow administrator to limit sub-administrators to manage only certain resources	Detection and prevention of DDoS, MITM, IP Spoofing, etc.	Default restrict firewall (80, 443), customized host-based firewall		Amazon S3 supports Server Side Encryption/Decryption and key management
Windows Azure (IaaS)	Azure Management portal provides Role-based access control management	Filtering Routers prevents common attacks, including DDoS.	Customized Firewalls for Azure SQL	Windows Azure handles automatically patching the OS	No server side encryption for cloud storage
GoGrid	Role-based permissions allow administrator to limit sub-administrators to manage only certain resources	Intrusion detection/prevention is provided by GoGrid community	Customizable host-based firewall		No server side encryption for cloud storage
Rackspace Cloud		Dedicated intrusion detection devices which is optional	Dedicated firewall and VPN services to block unauthorized system access	System patching configured by Rackspace to provide ongoing protection from exploits	
IBM SmartCloud Ent.	Role-based permissions allow administrator to limit sub-administrators to manage only certain resources	Firewall and IPS/IDS between guest virtual machines (VMs) and the Internet	Customizable host-based firewall	Public images patched and scanned regularly. Patch servers for private images are also available	
OpSource	Role-based permissions allow administrator to limit sub-administrators to manage only certain resources	Denial of Service (DoS and DDoS) Protection	Customizable host-based firewall		OpSource Cloud File provides In-flight SSL 128 bit and At-rest 256bit AES encryption
AT&T	Access via the AT&T web portal or VMware vCloud™ API which enables role-based access control	AT&T offers Managed Cybersecurity for intrusion detection and prevention	AT&T offers Managed Cybersecurity for firewall set up		Server side encryption service is offered for AT&T Medical Image
Savvis	Access via VMware vCloud™ portal or API which enables role-based access control	Savvis managed network services provide intrusion detection solution: network-based DDoS mitigation	Savvis managed security service provides network-based firewall, web application firewall and etc.		
CSC	Access via VMware vCloud™ portal or API which enables role-based access control		Customized virtual and network perimeter firewalls	Patch management is offered as support option	
DataPipe		Managed intrusion prevention service provide maximum host-level protection via continuous, non-intrusive multi-layer vulnerability monitoring and prevention	Managed firewall services enables all inbound data traffic flows to a firewall (s), filtering traffic based on users' specified requirements	Datapipe utilizes a state-of-the-art patch management solution to continuously test new patches and updates before deploying to servers.	Managed DB encryption, file encryption, DB backup encryption
Carpathia	Virtual desktop solutions based on Citrix XenDesktop allow simplified management	Managed intrusion prevention service	Managed firewall services	Managed patching service	It supports server side encryption for storage
Google App Engine	Role-based access control (viewer, developer, owner). It also provides APIs for authenticating users and sending email using Google Accounts	It provides API for DoS Protection Service Configuration	By using the Google Secure Data Connector (SDC), application can connect to systems behind firewall	Support	Secure connections via HTTPS for URLs; When a request accesses a URL using HTTPS, both the request data and the response data are encrypted by the sender before they are transmitted, and decrypted by the recipient after they are received
Windows Azure (PaaS)	Azure Management portal provides Role-based access control management.	Filtering Routers prevents common attacks, including DDoS.	Customized Firewalls for Azure SQL	Windows Azure handles automatically patching the OS your applications run on top of, removing the need for you to manually coordinate or script this.	No server side encryption for cloud storage.

Table 2.5: Regulatory compliance overview

	Certification
Amazon	PCI DSS Level 1, SOC 1/SSAE 16/ISAE 3402, FISMA Moderate, ISO 27001, ITAR, FIPS 140-2, HIPAA
Windows Azure (IaaS)	Safe Harbor , ISO/IEC 27001, FIPS , PCI DSS, SAS 70 Type II, HIPAA
GoGrid	SAS 70 Type II, PCI DSS, HIPAA
Rackspace Cloud	Safe Harbor , SSAE 16 Type II SOC 1*, PCI DSS
IBM SmartCloud Ent.	California Senate Bill No. 1386, FISMA, Gramm-Leach-Bliley Act compliance solution, HIPAA, PCI DSS, SOX, CADA Assessment and SCADA solutions
OpSource	SAS 70 Type I and Type II , PCI, HIPAA, Salesforce.com AppExchange Certification, Safe Harbor
AT&T	SAS 70, Type II, HIPAA
Savvis	PCI DSS, SSAE 16 (SOC 1) , Nasdaq Requirements
CSC	SAS 70 Type II, HIPAA
DataPipe	HIPAA, SOX and SSAE 16, PCI Level 1
Carpathia	SAS 70 Type II, FISMA, DIACAP, HIPAA
Google App Engine	Safe Harbor, SAS 70 Type II, SSAE 16 Type II, ISAE 3402 Type II
Windows Azure (PaaS)	Safe Harbor, ISO/IEC 27001, FIPS, PCI DSS, SAS 70 Type II, HIPAA

Table 2.6: Service-level agreements and support overview

	Service-level agreements	Technical support for developers
Amazon	Amazon EC2 – 99.95% monthly uptime Amazon S3 – 99.9% monthly uptime	Live technical support Extensive APIs Extensive technical documents support
Windows Azure (IaaS)	99.95% availability for compute 99.9% instance detection, storage, SQL, Access control	Live technical support Extensive APIs Extensive technical documents support
GoGrid	100% for all services	Live technical support Average APIs Basic technical documents support
Rackspace Cloud	100% for all services	Live technical support Extensive APIs Basic technical documents support
IBM SmartCloud Ent.	99.9% availability for IBM SmartCloud Enterprise	Live technical support Extensive APIs Extensive technical documents support
OpSource	100% Network Uptime 100% Server Uptime Guarantee	Live technical support Extensive APIs Average technical documents support
AT&T	99.9% availability for storage, compute,	Live technical support Basic APIs Average technical document supports
Savvis	100% availability for all services	Live technical support Average APIs Basic technical document supports
CSC	99.0-99.95% availability	Live technical support Average APIs Basic technical document supports
DataPipe	100% network uptime	Live technical support Average APIs Basic technical document supports
Carpathia	100% power availability 99.90% uptime	Live technical support Average APIs Basic technical document supports
Google App Engine	100% uptime for all services	
Windows Azure (PaaS)	99.95% availability for compute 99.9% instance detection, storage, SQL, Access control	Live technical support Extensive APIs Extensive technical documents support

Table 2.7: Business factors overview

	Response time EU (s)	Start year	Sample customers (medical domain)	Published security breaches	Vendor lock-in	Value added services
Amazon	3.58 (Ireland)	2006	Nimbus Health, CloudPrime, TC3 Health	Using Amazon cloud in attack; Amazon security bulletins [Ama12]	Because Amazon is industry leader, its interfaces have been copied by others	Industry leader; tools provided by Amazon: CloudWatch, S3, SimpleDB and many more
Windows Azure (IaaS)	3.42 (Dublin)	2008	T-Mobile, HealthVault		Abstraction is a method to avoid vendor lock-in issue in PaaS; depends on the implementation of application itself	MS SharePoint, easy integration of other MS products
GoGrid	8.49 (US-East)	2008	Condé Nast Digital Germany	Credit card leakage to authorized third party	FREE data transfer to and from your cloud servers to Cloud Storage	GoGrid Exchange allows partners to distribute preconfigured GoGrid Server images; CloudLink allows customers to have a direct connection between GoGrid's data centers
Rackspace Cloud	5.22 (London)	2006	TweetPhoto, IRX		Less risk of technology or vendor lock-in by using openstack	Subversion hosting, Attachment fu in Ruby, Cloudvox, Nautilus Cloud Files Plugin, Paperclip-Cloudfiles, Olark Live Website Chat, Vanilla Free Forum Hosting
IBM SmartCloud Ent.			Ottawa Hosiptal, Freie Universität Berlin		Open Java and cross-platform support with no vendor lock-in	IBM Hosted Application Security Management, IBM Hosted Mobile Security Management, IBM SmartCloud Virtualized Server Recovery, IBM SmartCloud Archive
OpSource	8.45 (US-East)	2009	Thermo Fisher, MediServe		Supports on-demand cancellation and return of data	Various managed services can be chosen for network security and server management
AT&T			Buzz, Sample of AT&T ForHealth			Intrusion detection systems (IDS) and assessments, Anti-virus and anti-spyware protection, VeriSign SSL certificates, AT&T Medical Image
Savvis		2008	EasyJet, Starz Entertainment			Savvis Managed Dedicated Load Balancing and SSL Acceleration Service
CSC		2008	Pharmaceutical Company Manages Growth; Hosts SAP Applications with CSC			Industry's only private on-premise cloud billed as a service, trusted and transparent cloud-enabled data centers
DataPipe		2006				Managed hosting, security support
Carpathia			Ascend Healthcare Systems, ChainDrugStore			Community Cloud for developer association
Google App Engine	8.36 (US)	2008			Python users can avoid "vendor lock-in" via the Django-nonrel project by porting their applications from webapp to Django. Data Liberation helps users to move their data in and out of Google products	Google Secure Data Connector, Private gadgets, Google Visualization API, Google Apps APIs, Google web toolkit, IDE support
Windows Azure (PaaS)		2008	T-Mobile, HealthVault		Abstraction is a method to avoid vendor lock-in issue in PaaS; depends on the implementation of application itself	

Chapter 3

Impact of Cloud Computing

Chapter Authors:
Christian Cachin, Kristiyan Haralambiev, Elmar Husmann and Nikola Knežević (IBM)

3.1 Methodology

Cloud computing profoundly affects the business models of the IT industry. Due to the many new available forms of outsourcing IT functions to the cloud, with their numerous advantages, cloud computing as a business model is a success. This chapter describes several factors related to the *trustworthiness* of clouds and how they may impact cloud-business models.

These topics chosen here are based on the *cloud security and privacy* factors and the *cloud legal and compliance* factors as described by TClouds deliverable D1.3.1 (Chap. 3). For each topic, the current situation in respect to cloud computing is analyzed, existing solutions are presented, and prospects are given about the future development of businesses relevant to addressing the concern.

3.2 Loss of governance

3.2.1 Overview

Data governance or *information governance* addresses the control over all activities in a company to collect, store, share, and process information. Data governance imposes formal processes on all aspects of data handling, such as schema definitions, data-retention policies, or security requirements. These processes usually also govern how they subject rules should evolve over time.

Information governance is closely related to the protection of information but more general; broadly speaking it spans *data quality*, the *lifecycle* of information, and the protection of *privacy and security*. In this analysis we focus on the latter aspect.

According to IBM [Soa10, IBM12d] the following are the main tasks for implementing data governance:

Understand and define: Understand where data resides, what domains of information exist, how it is related across the enterprise, define rules for quality and consistency of data, and develop policies and metrics for securing and protecting that data.

Manage and access: Enforce defined quality standards on data, maintain consistent representations, and ensure that only relevant data is used for production.

Consolidate and archive: Continuously monitor data and applications, consolidate or eliminate unnecessary applications, and archive obsolete data, but respect data retention rules even after the application has been retired.

Secure and protect: Protect data across the enterprise — in both production and non-production, both structured and unstructured — from unauthorized use.

Monitor and audit: Ensure information remains protected from authorized and unauthorized users on an ongoing basis, assess vulnerabilities and validate compliance. Report the status to auditors both internally and externally.

The "governance" factor is closely linked to several other factors mentioned in this chapter, especially to confidentiality and compliance. The content of Sections 3.7 and 3.9 should therefore be considered as well.

3.2.2 Relevance

Loss of governance is a *key concern* for the adoption of cloud computing. Since management, storage, and computation functions are outsourced to a third party, information governance becomes very important. This has been recognized widely in the industry [IDC10, SVC10, Clo12a, Sym11, BGPCV12].

Cloud providers often adopt a one-size-fits-all approach, where only standard governance functions are provided. These cannot be tailored to specific customer needs, hence it is much more difficult to implement the governance operations.

Concerning the scope of the problem, data governance is most relevant for *public clouds* and less relevant for *private clouds*. By their nature, public clouds limit access to data and programs for the data owner; if governance mechanisms are not offered as part of the service, then the customer usually cannot exercise its protection, monitoring, and auditing functions. In private clouds, on the other hand, governance mechanisms are more readily obtained because the data owner controls the cloud infrastructure.

3.2.3 Solutions

Governance concerns primarily the procedures carried out by humans and not directly technical solutions embodied in IT products. In this section we discuss solutions of both kinds that are available to businesses today. Organizational solutions usually encompass processes, frameworks, and guidelines; technical solutions typically support these processes.

Organizational solutions. Procedures are maintained in the form of guidelines and auditable compliance frameworks. Many regulations exist in this space and require processes for data provenance, traceability, audits and identity management. Regulation needs to be audited. This has created a market for compliance checking, audits, and certification by external entities.

Harmonization of regulations is underway. However, not many regulations have specifically addressed cloud computing so far.

Technical solutions. Data governance is most difficult when one deals with unstructured information. Many existing products in this space target unstructured information and data storage, but also databases. Much consideration is also given to data provenance tracking.

Families of products supporting information governance are today available from many IT vendors. For instance, the EMC SourceOne™ family, IBM InfoSphere Platform solution portfolio (addressing data warehousing, Information integration, master data management, big data analytics, and lifecycle management), or the Oracle Fusion Governance, Risk, and Compliance (GRC) component of the Oracle Fusion Applications suite.

Some existing products already extend to cloud-specific issues, like record- or field-level encryption for outsourced databases. But the majority of products focuses still on governance inside enterprises without giving special consideration to the cloud model.

3.2.4 Outlook

Future cloud services can go much beyond the currently offered *computing* services, whether they concern infrastructure, platform, or software services [MG11]. There exist many opportunities for cloud-specific governance services to appear, which cover all governance aspects of cloud services, organizational and technical.

On the organizational side, auditable regulations and frameworks for cloud services will appear and create a new market. Certifications for cloud providers will make their offerings more comparable. Especially those governance factors that become business differentiators will depend on new regulation frameworks. Such works are currently being developed, e.g., by the Cloud Security Alliance [Clo12a].

In terms of technical developments, for example, cloud storage services could be complemented with features allowing information discovery, data-retention policies, audits, long-term archiving, storage monitoring, and resilience/business continuity.

Cloud providers may offer data governance services in order to differentiate their products from others, in addition to their standard services. This permits providers to differentiate the value of their basic services and to sell higher-level services for a higher price. Examples of differentiated services exist already, such as Amazon S3's support for server-side encryption of stored data or the choice of data center and resilience group offered by many cloud infrastructure providers. However, many more features of this kind will become available, such as higher availability guarantees, better risk assessment and reporting, and richer security controls.

Today many cloud vendors are readily working on such extensions, which differentiate their offering from the basic "one-size-fits-all" paradigm. Customers would ideally want to obtain the same control over the delivered service as they have with in-house solutions. Achieving that in the cloud-computing model is not possible, however. Higher levels of service and value-added components for governance will be key differentiators in future cloud markets. The low-level cloud commodity offerings will compete alone by price; but only delivering additional value will generate higher revenues.

3.3 Isolation failure

3.3.1 Overview

Multi-tenancy and shared resources are defining characteristics of cloud computing. In such environments, *isolation* refers to logical segregation over multiple layers, rather than physical separation of resources. High degrees of multi-tenancy over a large number of platforms are needed because it:

- ensures envisioned flexibility of on-demand provisioning of reliable services, and

- enables cost benefits and efficiencies due to economies of scale.

To operate within these high scales of consumption, cloud providers have to ensure dynamic, flexible delivery of service and isolation of consumer resources.

Multi-tenancy spans the layers at which cloud services are provided:

IaaS: Tenants share infrastructure resources like hardware, compute servers, network and data storage devices.

PaaS: Tenants share application (container) servers, and, thus, the execution environment.

SaaS: Tenants are sourcing the same application; this, in turn, means that the data of multiple tenants is likely stored in the same database and may even share the same tables.

In case of IaaS, multi-tenancy is achieved through multiplexing the execution of virtual machines from potentially different tenants on the same physical server, isolating each virtual machine from others (hypervisor-level isolation). In case of PaaS, multi-tenancy is achieved through execution of every tenant's application in a separate application container, within server, thus providing execution isolation. In case of SaaS, multi-tenancy is enabled in the application or the database, through authentication and Access Control Lists, that govern the isolation of users among themselves.

Isolation failure denotes the failure of these mechanisms separating storage, memory, network routing, or execution. In addition, isolation failure includes exploitation of reputation of different tenants (for example, when attacker breaks a less secure VMs, escapes isolation and attacks another, more secure VM in the system) [ENI09].

3.3.2 Relevance

As isolation is the main approach in providing multi-tenancy, any failure is detrimental to both customer and cloud provider's business [JG11a].

Isolation failure is most relevant for public clouds, and less relevant for private clouds [ENI09]. As public clouds hold data of various tenants, they are a good target for various malicious entities, that may want to get into the possession of the data of other tenants. Hence, isolation is important in this setting. Private clouds usually belong to the same entity, where tenants have the same goal, and, as such, isolation failure is of low impact. Nevertheless, isolation is necessary for private clouds, as a mean of achieving fair access and predictable performance.

The impact can be a loss of valuable or sensitive data, reputation damage and service interruption for cloud providers and their clients.

3.3.3 Solutions

Isolation relies solely on technical solutions that are part of different IT products. This section lists different solutions available as products, categorized by their applicability (whether IaaS, PaaS or SaaS).

IaaS. IaaS hosting providers have to provide full isolation between different workloads that belong to different tenants. Server virtualization [Gol72] provides a good level of isolation between virtual machines. Solutions currently available on the market, and widely used by different IaaS providers include: Xen [BDF+03] (and its commercial variants Citrix XenServer [Cit10]

and Oracle VM [Ora11b]), Red Hat's KVM [Red08], Microsoft's Hyper-V series of products [Mic07], and VMWare ESX/ESXi family of virtualization solutions [VMW08]. The field of server virtualization experienced a large consolidation in recent years, and many products disappeared. As such, it is not expected to see any new offerings in this field.

In addition to server virtualization, secure isolation includes isolation at the network layer. Network layer in IaaS provides Layer-2 connectivity between different workloads that run over the same infrastructure. Isolation has to be compatible with the isolation in the physical data center to meet customers' expectations and not be a barrier for cloud adoption. In a private cloud setup, isolation is almost as important as it is to IaaS hosting providers, as it is often required to isolate certain workloads and environments from others, such as the systems for finance and human resources departments. Apart from using VLAN approaches, there are several commercial solutions to network virtualization. Open vSwitch [PPK+09] is an open source virtual switch that supports OpenFlow, while Cisco offers their Nexus 1000v [Cis11] virtual switch for VMWare environments. Microsoft has developed a virtualized networking solution for their Hyper-V products, called Hyper-V Extensible Switch [Mic12].

As this is a recent technology, we expect many new products in this area in the near future [Xu06]. Business impact of IaaS isolation lay in better isolation with lower performance overhead, this will increase consolidation.

PaaS. PaaS denotes running multiple applications from different customers inside the same infrastructure and on top of the same platform. Usually, PaaS uses the infrastructure consisting of a number of ordinary servers or virtual machines, typically provided by IaaS providers. Isolation in PaaS means one application cannot interfere with another application running inside the same platform. To achieve this in the context of security, PaaS provider must ensure memory, network and filesystem isolation.

There are basically two approaches to realize isolation. One is to rely on the same mechanism as IaaS providers, and use virtualization. However, this approach requires more resources, and incurs performance drawbacks. The other approach is to leverage existing security mechanisms in the operating system: SE-Linux [LS01], approaches like Linux VServer [lin05], OpenVZ [ope12] (and it's commercial variant Parallels Virtuozzo Containers [Par10]), Solaris Containers [Ora11a], FreeBSD Jails [KW00], or other approaches described as operating system level virtualization could be used to provide application sandboxing.

In the area of PaaS isolation, the proliferation of UNIX like hosting platforms and virtualization allows business to migrate their applications from one cloud to another with little or no cost.

SaaS. According to Chong et al. [CCW06], isolation in SaaS model is achieved through data isolation, either through database isolation (where each tenant has different database), shared database with separate schemas, or shared database with shared schemas. Here, protection and isolation relies on the database access control mechanism, and is thus related to the particular database. Isolation is not a technological problem, but administrative.

3.3.4 Outlook

Emerging research topics on isolation are in the areas of verification and auditing of isolation and other infrastructure properties. For example, Bleikertz et. al. [BSP+10, BGM11] present a system for analyzing the environment from within the VM, and auditing its security properties in a domain specific language.

3.4 Insider fraud

3.4.1 Overview

Insider fraud refers to the malicious or criminal attacks perpetrated upon business or government carried by employees, contractors, or other business partners with legitimate, past or present, access to the organization's network, system, or data [Pon11]. To carry such an attack, the insider intentionally exceeds or misuses that access in a manner that negatively affects the confidentiality, integrity, or availability of the organization's data or resources [CMTS09], often for personal financial gains.

According to a recent *Cost of Data Breach* study by Ponemon institute [Pon11], 31% of all data breaches were caused by malicious insiders. While studies differ on how to measure the cost of insider fraud, two recent billion-loss cases from the financial industry [CJ08, Sal11] shows how costly and reputation damaging such incidents could be.

Insider fraud is among the most significant threats to any institution. While the fundamental nature of this threat is not be changed by cloud computing, there are new exploit possibilities stemming from migration to the cloud. We will focus on these new attack vectors for the remaining of this section; for a more general discussion and studies on insider fraud, the reader could refer to [CMTS09, Pon11, CLM+12] as a starting point.

3.4.2 Relevance

The most often considered cloud insider threat is the malicious administrator [Clo10, Sla11, CN12]. A rogue administrator employed by the cloud provider could make an unauthorized copy of a user's database or modify the stored data. This leads to loss of data confidentiality or integrity, respectively. The usual motivation would be financial gains, though one should also consider the possibility of an insider seeking to harm the cloud provider, due to unresolved previous conflicts, by sabotaging its customers. Unlike a typical enterprise environment, the malicious administrator threat has several levels: applications administrators, system administrators, virtual image administrators, and hosting company administrator. If data protection is not present, administrators at each level also have the capabilities of those "above" them, i.e. listed prior in the above list.

Another insider threat is nefarious use of cloud computing by an employee inside the institution. The most common example is a person copying proprietary or confidential documents to external location before quitting their job. In a typical enterprise environment, this will raise a red flag. But for data stored in the cloud, the institution might not have the same level of monitoring tools. Alternatively, using the high computational power of the cloud, an insider could try to mount a denial-of-service attack against their institution or to crack a colleague's password by brute force, for example, after obtaining a file encrypted with that password.

Lastly, we discuss vulnerabilities exposed by the use of cloud services which can be exploited by skilled employees [CN12]. If the security policies or access control models used within the institution are not fully compatible with the available cloud functionality, this might allow unauthorized access to data or resources. A more sophisticated attack could also make use of the cloud replication lag. In the example given in [CN12], an insider triggers an update on the server replica with the highest update latency seconds before a regular one. This way, the malicious update will be overwritten by the regular update, but nevertheless it will give the insider the possibility to exploit such short discrepancy; for example, to buy goods at much lower price.

3.4.3 Solutions

Countermeasures against malicious administrators vary from relying solely on the cloud provider, through enforcing service level agreements, to protecting all data stored with the provider. Given that it is in the cloud provider's own business interest to protect its administrators from accessing customers' data, one could expect cloud providers to have rigorous privacy policies which are strictly enforced similarly to other IT companies [Zet10].

While trusting the cloud provider might be overall sufficient, in many situations any access to customers' data might be critical. In such cases, protecting all the data might be the only option. There are various solutions for that like the ones described in Section 3.7 as well as novel research ideas [dVFJ+07, IKC09] which further leverage the data protection for managing access control and privacy. It also worth mentioning that once the fully homomorphic encryption [Gen09] transitions from research to practical products, malicious administrators could easily be prevented from accessing customers' data.

The solution proposed in [Clo10] is to enforce strict supply chain management and conduct a comprehensive supplier assessment, to specify human resources requirements as part of legal contracts, to require transparency into overall information security and management practices, and to determine security breach notification policies. Many of these could be achieved through careful management and enforcement of service level agreements according to [CN12]. Recent research ideas [MBS12, ENC+12] could also contribute for the use of service level agreements between cloud providers and institutions.

Regarding possible inside fraud by employees, institutions should resort to current enterprise solutions and migrate them to the cloud. These include specifying and enforcing clearly defined access policies and controls, monitoring and responding to suspicious activities, implementing strict password and account management policies, enforcing separation of duties, considering insider threats in the software development process, exercising extra caution with privileged users and system administrators, developing an insider incident response plan, and many others [CMTS09].

Most technical weakness introduced by the use of cloud computing can be addressed by careful planning of the cloud migration and maintaining of the services. Institutions and cloud providers should have agreements about updating software and responding timely to security incidents. Proper monitoring and logging tools should also be available. Enforcing separation of duties and other fundamental security policies is critical for both parties. Organizations should consider proper authorization methods and access controls for cloud computing.

3.4.4 Outlook

Insider fraud looks like a "cat-and-mouse game" between insiders and their institutions. We expect to see a lot of research and new products towards this problem in the coming years. Ultimately, cloud computing gives an opportunity for novel solutions based on normal use behavior analysis which was previously missing enough real-world data. Combined with other approaches like socio-technical predictive models and identification of indicators of insider threats [CN12], this could limit the extent of insider fraud.

3.5 Compromise of management interfaces

3.5.1 Overview

Management interface is a management control panel or cloud interface, that may be either web-based, or any kind of a remote shell (a SOAP-based API, SSH shell, ...). Through it, a customer manages all the services and resources provided on the cloud, or the cloud provider manages it's resources and client assignments.

3.5.2 Relevance

Compromise of the *cloud management interface* is quite important, as this interface mediates access to a large set of resources on the cloud. It occurs due to the nature of the management interface — it is exposed to the Internet (increased attack surface). It poses a high risk for public clouds and public parts of hybrid clouds, on all levels (IaaS, PaaS, and SaaS). The insecurity may come from both the cloud provider side (through security holes in the interface), or from the customer side (through phishing attacks). Moreover, due to vulnerabilities on the end-point machines (used by operators at the customer site, or cloud operators), the cloud infrastructure could be jeopardized through weak authentication of responses and requests [SHJ+11].

The factors that affect the security of the cloud management interface come from [ENI09]:

- a poor system for authentication, authorization and accounting,
- vulnerabilities in the remote access protocol,
- misconfiguration,
- various vulnerabilities in the system or OS hosting the management interface, or
- poor patch management.

The issue of the cloud management interface is closely related to authentication mechanisms compromise (such mechanisms rely on various protocols, such as SOAP and SAML, and thus are exposed to vulnerabilities in the protocol implementations). Moreover, due to Internet accessibility of clouds, password based authentication attacks (for example, through corporate passwords stealing Trojans) have much larger impact, than on standard machines and services, rendering such authentication mechanisms inefficient. Hence, stronger authentication mechanisms are needed in order to further mitigate risks leading to cloud management interface compromise.

Risk related to cloud management interface could be further mitigated through greater investments in the security on the cloud provider site, or through the education of operators on the customer side, as well as hardening of the end-point machines that are used for administrative tasks on the cloud.

Compromise of cloud management interfaces is ranked among top 10 cloud-computing related risks by ENISA [ENI09].

3.5.3 Solutions

In the previous section, we listed possible factors that attribute to cloud management interface compromise. Two listed factors are particularly tied to the cloud: authentication, authorization and accounting issues, and remote access protocol vulnerabilities; hence, these two factors will be discussed in this section. Furthermore, techniques and solutions could be either *preventive* (active) or *detective* (passive), and we will attempt to classify each of them as such.

Authentication, authorization and accounting factors. Preventive solutions incorporate multi-factor authentication, authentication schemes with stronger security than password-based approaches, and solutions against compromised end-hosts. Some of the existing implementations based on open standards for authentication are OpenID [RR06], OAuth [HL10], Shibboleth [SC05], and Mozilla Persona [Moz12]. Many of these implementations use SAML standard [sam05] for exchanging authentication and authorization data. Another important scheme that reduces the attack surface on the authentication system is *single sign-on*. Some of the well known systems that support this scheme are: myOneLogin by VMWare [VMW12a], Okta [Okt12], Oracle Access Manager [Ora12], Ping Federate [Pin12], Tivoli Identity Manager [IBM12c], and SmartSignin [Sma12]. For multi-factor authentication, well known solutions are RSA SecurID [EMC12] and VASCO's DIGIPASS [VAS12]. A well known secure authentication solution, protecting against compromised end-hosts is ZTIC by IBM [WKH+08].

Detective solutions span access reporting, logging and correlation, as well as security scanning. Such solutions are offered, among others, by: SolarWinds [Sol12], AlertLogic [Ale12], Loggly [Mar11b], LogRhythm [Log11], Splunk [Spl12], while others are listed in Section 3.8.

Remote access protocol factors. Some of the preventive solutions are listed in the previous section, as secure authentication protocols also relate to hardened remote access protocols. Other solutions encompass use of VPN, or other secure tunnelling mechanisms, either through libraries or appliances. As mostly all network equipment providers offer solutions and appliances in this category, we will refrain from listing them. Another class of solutions belongs to web applications security scanners, that help in detecting known (or, in certain cases, unknown) security threats to exposed interfaces. A list of offers in this class of solutions [The11] contains well known Metasploit [May07], IBM's AppScan family of products [IBM12b], HP's WebInspect [HP 12], Hailstorm [Cen12], Sentiel from WhiteHat [Whi12].

Detective solutions span the same set of offerings listed in the previous paragraph.

3.5.4 Outlook

Current solutions for authentication and authorization may have sound theoretical background, however, different implementations may be far from secure, as there is a tradeoff between convenience and security. Hence, most of the research is focused on providing analysis and automated tools (penetration test tools) for assessing and inspecting various protocols and interfaces [SHJ+11, SMS+12]. Also, another area of research that receives a lot of attention is about log analysis, in order to find long-lasting, stealth attacks [BMO11, HZ12].

3.6 Secure data deletion

3.6.1 Overview

Data deletion can be as critical as preventing data loss. The latter is usually realized through replicating the data across several servers at different locations. Snapshots and backups of the servers make the deletion of a simple piece of data even more complicated as their data can stay around long after the deletion operation. Moreover, one should keep in consideration the possibility of recovering the deleted data from a physical device either by an insider with access to the hardware or a party handling retired/disposed devices. While there are several standards

for data sanitation, e.g., the NIST guidelines [KSSL06], cloud providers are not known to implement such measures at present.

Secure data deletion stands for a data deletion process at the end of which is guaranteed that the deleted data cannot be recovered by anybody. As the users have no control over the physical storage of the data, it can be realized through proper policies and procedures implemented by the cloud provider *or* through cryptographic techniques used to protect and access the data.

In principle, the goals of secure deletion and resilience (Section 3.8) for stored data contradict other. This creates an interesting tension for data owners, which extends also directly to cloud business models.

3.6.2 Relevance

Permanently deleting data does not only save space but also reduces liability and legal risk in certain cases. For example, in the U.K., patients can request inaccuracies in their medical records to be deleted under the Data Protection Act. In the U.S., under the Fair Debt Collection Practices Act, everybody is entitled to request deletion of invalid or time-elapsed entries. Similarly, a court order could require some information from a person's file to be deleted. Failing to fully comply with such a request leaves a business with the possibility of litigation.

As discussed in the last paragraph, data sanitation is an important issue when disposing storage devices. If data is kept in clear and devices are not properly sanitized, important information can be leaked. Valli and Woodward [VW08] studied the remnant data on Australian second-hand enterprise-level hard drives and found sensitive information that related to critical infrastructure providers. One can find multiple similar cases from around the world in the media.

To avoid similar problems in cloud computing, one basic solution is server-side data encryption combined with proper deletion handled by cloud providers. Examining Table 2.4, we see that data encryption is not currently supported by major providers. This shows a somewhat popular view that data deletion is a secondary concern or should be handled by the user. While it might be true is many cases, it is worth pointing out that ensuring permanent data deletion upon user's request, which also enables the implementation of retention policies, is critical for allowing business like law firms, tax preparation companies, hospitals, and others to migrate to the cloud.

3.6.3 Solutions

The most prevalent solutions providing secure data deletion rely on encryption to securely store and access data. A file is accessible as long as there is a copy, possibly a backup one in a previous snapshot, and its encryption key is still available. Therefore, a file can be deleted if all snapshots which contain it are deleted by the cloud provider or its encryption key is no longer available.

Nasuni [Nas11], a storage solution provider, allows customers to establish snapshot-retention policies based on the snapshot age or the number of snapshots. Once the last snapshot containing the deleted file is erased, the file is permanently removed. This is not perfect, but is the best one can hope to achieve with simply key-management system, e.g., all data is encrypted with the same user key. SafeNet [Saf12] which builds atop Amaozn AWS provides data protection for the cloud which can be combined with their enterprise key-management solution. The latter allows granular encryption capabilities applied to databases, applications, and individual files. This way, one could encrypt individual sensitive files separately with their own key. To delete

such a file requires simply the revocation of its encryption key. While this resolves the issue with the immediate deletion of a file, it is worth pointing out that the key managing authority is still capable of producing the key. So, theoretically, the file is still accessible if there is a copy of the file and the key-managing authority is forced to reveal the key, e.g., in the case of subpoena. Other solutions based on data protection, which provide secure data deletion within the above-described restrictions, can be found in Section 3.7.

A recent line of work tries to overcome the limitation of key retrievability even by the key-managing authority. This is achieved by cryptographic key operations across a quorum of key managers. Tang et al. [TLLP12] designed and implemented a prototype system atop Amaozn S3 which achieves fine-grained policy-based access control with secure file deletion. Files are associated with policies and each policy has its own encryption key. Once a policy is revoked its key becomes unrecoverable to anyone. Although far from a business solution, this work provides useful insight about possible future products.

3.6.4 Outlook

Current solutions rely on data protection by controlling the access to encrypted data. This can be view as a side result of "one-size-fits-all" paradigm with all users and all files being treated the same way. Cloud providers have the capabilities to delete all copies of a file in all replicas and snapshots, but that would be costly in general. Nevertheless, it might be worthwhile for highly sensitive data. Combined with a proper sanitation of the sensitive data or key management, this will provide fast and reliable deletion. Alternatively, the fine-grained policy-based access control allows secure data deletion. We expect to see this idea into real-world solutions in the near future.

3.7 Data protection

3.7.1 Overview

For public infrastructure clouds, both computing and storage clouds, data resides in a shared environment collocated with data from other tenants. Therefore, *data protection* denotes the means by which access to the data is controlled and how the data is kept secure, when placing sensitive and regulated data into a public cloud. In other words, the aspects of data protection are:

Confidentiality: Maintaining the secrecy of the data that relates to accessibility of information and its isolation[1] from other tenants; and

Data encryption: Cryptography provides a sound way for the protection of data; this makes it "future-proof."

According to a NIST study [JG11a], the *key concerns* of protecting data stored in a cloud arise because of the following elements:

Value concentration: The risk of an intrusion increases because valuable, high-profile data from many clients resides in a common location, which is well-known and accessible

[1] isolation of data is conceptually different from the isolation of execution, discussed in Section 3.3, as data resides on storage. Additionally, data is passive, hence there must be mechanisms of isolating data, even without any execution mechanisms.

to clients. Moreover, a determined attacker may also attack the cloud provider (either through Denial-of-service attacks, or through social engineering) in order to gain access to this valuable data. This concern is important, as the intrusion detection and prevention mechanisms are mostly controlled by the cloud provider (not the organization that owns the data).

Data isolation: As discussed in Section 3.3 (but see Footnote 1), cloud providers share their resources by design. As such, data must be kept away from unauthorized users. There are two means of achieving such goal: access controls mechanisms and encryption. Most of the access control mechanisms are centered around identity, hence, authentication mechanisms become an important issue for cloud computing. As tenants do not have physical control over the data storage, encryption is the only way to ensure that the data is protected, both at rest and in transit. Furthermore, proper encryption ensures that data will remain secure even if someone steals it and tries to decrypt it later on (i.e., future-proofing). As mentioned in Section 3.3, database access is shared in SaaS and PaaS environments, and involve clear tradeoffs between features and isolation. This requires careful evaluation of the suitability of the data management solution for the data involved. Some industries, such as healthcare, have particular requirements that highly influence (and constrain) the choice of database and data organization used in appropriate application. Finally, encryption is only as strong as the processes for managing the encryption keys. Thus, tenants should either store keys privately, or on a cloud, only if all the risks are carefully weighted.

Data sanitation: Data that is deleted from the client's perspective may still reside in some form on the cloud. This particular concern is addressed in Section 3.6.

Data encryption relies on encryption keys, hence, the protection of these keys is quite important topic, as the encryption is as strong as the system for storing and handling the keys. This section focuses on encryption, as well as access control and management aspects of data protection.

3.7.2 Relevance

As mentioned in Section 3.2, protecting the privacy and security of the data within the cloud often means encrypting data. There is a lot of published research and case studies concerning data protection, and, more specifically, encryption. However, implementing encryption in practice is not an easy task. Data protection in terms of encryption is important mostly for public and hybrid clouds, but it has benefits for private clouds, as it offers a long-term protection (if implemented correctly) in case of a breach.

As with all security measures, data protection is complete only if applied on all levels. In case of IaaS, this means that proper isolation must be in place. If virtual machines use any cryptographic primitives, corresponding keys need to be protected against unauthorized access. Moreover, data containing executable code such as disk images need to have integrity protection (to discover tampering) and encrypted (to prevent leakage of sensitive material). For PaaS, data protection is performed through system's access control mechanisms, in order to isolate different execution domains, and through per-tenant encryption of data.

For a whole-system data protection, data must be protected during its delivery to the cloud, while residing on the cloud, and during retrieval.

3.7.3 Solutions

As data protection is an important topic, there are several important offerings and services for both standalone and cloud deployments.

In the area of data access control and management, but also including protection, there are Sensitive Information Management suite, the Privileged Session Management and Privileged Identity Management suites from Cyber-Ark [Cyb12], PowerBroker series of products by BeyondTrust [Bey12] that, among others, handle privilege management processes, and Halo from Cloud Passage [Clo12b], that deals with privilege management and isolation of SaaS solutions.

In the area of data encryption solutions for enterprises, there are many offerings: solutions from Voltage [Vol12], CipherCloud [Cip12], Thales [Tha12], Vormetric [Vor12] (that offers scalable solutions for encryption and key-management on all commercial databases), GreenSQL [Gre12] (that offers simplistic approach to database security for MSSQL, MySQL and PostgreSQL), IBM's InfoSphere Guardium [IBM12a], and Application Security Inc [App12] (offering some specialized products, such is dbProtect, that discovers sensitive data within databases on the corporate network). Like any security offering, these products do not match all demands — one needs to do a thorough research in order to ensure that any fits within business goals for the long term. Thus, the business impact of any product of this category is high, if such product is highly customizable to wide-range of different business needs.

3.7.4 Outlook

Many current solutions deal with providing seamless data encryption and isolation. These solutions encrypt the data while the system is otherwise idle, and decrypt data while in use. The recent research trends point toward computation on encrypted data, thus increasing security, and protecting against malicious or curious cloud providers. Apart from the seminal work on fully homomorphic encryption by Gentry [Gen09], there are much more practical and pragmatic approaches that involve trusted intermediaries, like CryptDB [PRZB11]. Some of these may form the basis of future commercial offerings.

Another research topic relevant to data protection focuses on describing and implementing policies on information-flow of sensitive (and encrypted) data. There are several notable works in this research area, such as Excalibur [SRGS12], Dstar [ZBWM08], Flume [KYB+07], and SilverLine [MRF11].

3.8 Resilience

3.8.1 Overview

In general terms, *resilience* is a property of the system to *provide and maintain* an acceptable level of service in the face of faults (both benign and malicious) and provide *availability* even in critical security situations [SH07]. For the purpose of cloud computing, resilience is related to redundancy and replication mechanisms (collectively described as availability), and integrity protection mechanisms (either due to replication or without).

3.8.2 Relevance

There are two aspects of resilience, when it comes to cloud computing. The first aspect is the resilience implemented by the cloud-service provider, that seeks to have all possible resources

available at any moment. The other aspect is the resilience that tenants could obtain from the cloud, through methods implemented by the cloud provider, or custom methods (such as replication over multiple machines or even multiple providers; see TClouds' activities on cloud-of-clouds middleware).

Resilience (and more specifically, availability) has a big impact on the business, as any (possibly unplanned) downtime directly affects operation. Threats to availability come as denial of service attacks, human factor, equipment outages, and natural disasters, and impact it temporarily or permanently, while a loss can be partial or complete [JG11a].

There are various degrees of control over resilience in the cloud. For a private cloud, since both the provider and the customer are the same entity, the degree of control is the same on all levels. However, in public clouds, customers have very little control [Sav11]. In the PaaS model, customers have more control and the ability to incorporate resiliency into their application (such as stateless design, distributed execution and decoupled logic/paradigm). In both models customers must develop a thorough understanding of the cloud provider's resiliency features, such as data replication, snapshots and location, in order to build resilient applications. Cloud customers have the greatest control over resiliency in the IaaS model. For example, Amazon offers several global regions and availability zones, that customers can leverage to build resilient systems [Var10].

As clouds host large quantities of data, processing it at a large scale, *silent data corruption* is a serious issue for cloud customers [Har07]. Hence, cloud providers need to have mechanisms in place for *data integrity*. On the other hand, customers also need to have *end-point integrity* protection, in order to have an all-round integrity protection of their most valuable asset [MM10].

3.8.3 Solutions

Resilience in the cloud is usually implemented by creating a copy of a VM on another server (and possibly in another, isolated availability zone). Different solutions, some tied to hypervisor provider, like VMware's vSphere [VMw12b], or independent, like Marathon Technologies' EverRun [Mar11a] keep these copies loosely synchronized, and restart the redundant VM as soon a fault occurs on one server and heartbeat signals are disrupted.

Eliminating single points of failure is another method toward increased resilience, and clustered servers generally implement redundant server components (for example, redundant power supplies) along with LAN and SAN connectivity. All major SAN and network providers provide cloud-based high-availability solutions.

Monitoring is another crucial aspect of achieving high-availability, which contributes to increasing resilience. Availability of any application is governed by its stability and the computing resources available. Since the goal of cloud providers is to maximize resource sharing, some virtual machines might run on physical systems with depleted resources. In extreme cases, that may crash the whole server and disable all VMs. All hypervisor providers (see Section 3.3) provide monitoring systems, but there are other offers on the market: Zenoss vCloud Monitoring [Zen12], Zenoss vCloud Monitoring system [Zen12], Compuware's APM [Com11], and vFoglight from Dell [Que10], to name a few; all can track server resource use, report resource shortages and help capacity planning. This area is quite populated with solutions, so enterprises can find the best offer for their particular setup.

Cryptography plays an important role in *integrity protection* of the data. Such tools allow detection of data-tampering, and, as such, enable stronger data resilience. GuardTime and Joyent offer data integrity protection [JG11b] for enterprises through a hierarchy of GuardTime

Gateways, each using keyless signatures to electronically sign any data entering the cloud. This approach ensures that one can prove that data was not tampered with, but also, more importantly, where the data came from and what paths it took through the system. Another solution for data integrity is through Cloud Data Management Interface (CDMI) SNIA standard [Sto11], that enables one CDMI compliant system to query another CDMI compliant system for data object hashes, and verify that two copies of the data are identical. Although there are many other data integrity protection solutions, in order for this section to remain concise, lastly, it is worth mentioning ZFS [BM05] as a mean of providing data integrity. ZFS is a file-system developed from ground-up with data integrity protection, through checksumming, data healing and copy-on-write mechanism. Some big-data solution providers, like EMC in their Greenplum database, use ZFS in their offers [Sta07].

3.8.4 Outlook

Similarly to the outlook on the loss of governance issue (Section 3.2), there are many opportunities for resilience related offerings. Standard computing and storage services could be complemented with features regarding automatic crash-resilient replication, Byzantine Fault Tolerance for both computation and storage, or integrity protection.

For example, many cloud providers have diverse, isolated data centers — Amazon's S3 and EC2 support a choice of data-center region and resilience group. We expect to see more differentiated offers in the future, with higher availability guarantees and fine-grained integrity protection controls.

The state-of-the-art research on resilience in cloud computing focuses on two approaches: developing techniques for assessing resilience (or, dependability), and developing methods for managing and improving the resilience. Both approaches are important for business applications, as they allow monitoring system state and making informed business decisions, as well as lowering the risk of downtime. Kounev et al. [KRB+12] provides an extensive review of the-state-of-the-art research on both topics. Proper assessment techniques will lead to better modeling of cloud performance and availability. As such, we can expect to see future systems that provide even better resource utilization, through tighter packing of virtual machines. These techniques will also improve resilience through better fault isolation, as faults tend to cluster in both temporal and spatial domain.

Besides focusing on assessing and modeling the resilience, there is a recent research effort on the cloud-of-clouds computing paradigm as a path to achieve cloud computing resilience. TClouds makes major contributions along this line through its work on the cloud-of-clouds paradigm, as documented in the corresponding reports. The concept of an Intercloud or a cloud-of-clouds is a natural extension of the cloud concept, through leveraging the availability of multiple, independent clouds. The cloud-of-clouds empowers users to self-organize different cloud offerings and tailor them toward their use cases, while increasing availability and performance.

Finally, there is ongoing research [CFJS12] on automating integrity checks between various tiers, with no changes to the source code. Proliferation of such systems would enable cost-effective development of cloud-based solutions.

3.9 Compliance

Cloud providers have-to or can voluntarily comply with a number of security and privacy policies. This includes those externally set by regulators as well as binding corporate rules, volun-

tary codes of conducts or organisation level standards such as the ISO 27000 family.

Compliance can further include policies that are specific to a customer. Typically, those specific policies need to be negotiated and agreed on purpose in the cloud contract and service agreement, whereas the compliance to external policies is realized and demonstrated through other means.

In the following we decompose this adherence to external policies into two sub aspects: first the *accountability* that builds on responsible (self-controlled), reliable and verifiable actions by the provider. Second, the *transparency* towards external parties (in particular the customers) that is closely linked to this.

3.9.1 Accountability

Overview

Accountability in cloud computing is a relatively broad concept. It describes that a cloud provider adheres to policies — in particular with regard to security and privacy — that are externally established and accepted. While perfect control or enforcement of this adherence is impossible, accountability typically implies [EU 10] responsible, reliable and verifiable activities by the cloud provider. This is the basis for trust.

Table 2.5 in this document gives an overview on the regulatory compliance of leading cloud providers. It is in general possible to audit at a given point in time the compliance to these regulations and their respective policies — and potentially have this certified by an independent external party.

However, this can only mean that the cloud provider demonstrates appropriate processes, procedures, roles, technical requirements etc. . . But the compliance in daily practice — e.g. if procedures are executed in the specified way — is a question of accountability of the cloud provider.

This is further linked to the issue of transparency as it will be described in the next section. In general, cloud providers should ensure appropriate transparency to demonstrate their accountability and increase the level of trust that customers and partners put in them.

Relevance

Accountability addresses several fundamental issues of trusting a cloud provider. In this context it needs to be acknowledged that it is not only the cloud provider that has to be accountable but also the cloud customer — e.g. as data owner towards an end customer. Hence, accountability is the basis for establishing chains of trust between cloud customers and providers as well as in-between collaborating cloud providers. As the ENISA risk study on cloud computing [ENI09] puts it: “Ultimately, you can outsource responsibility but you can’t outsource accountability.”

Solutions

There are different levels of implementing accountability for cloud computing. At the level of the regulators there is interest in a legal framework for accountability [EU 10, EU 12]. One element of this is to include the obligation for accountability as an integral element of policies that shall be adhered to be cloud providers. This obligation would still leave the cloud provider self-responsible for the way how he is observing policies but it would make the existence of appropriate reliable governance mandatory. In the same way, the cloud provider would also be obliged to be able to prove this at any time.

Legally enforcing accountability is an alternative and complementary way of implementing cloud policies and binding rules by making certification mandatory. In fact, external certification has some deficits in flexibility e.g. with regard to matching different firm sizes of providers or to cope with the permanent technical, organizational and service development of the provider. Hence, a combination of accountability and voluntary certification —as now suggested in the European Cloud Strategy [Eur12] is a pragmatic way.

On the level of the provider, the corresponding concept to accountability can mostly be summarized under the term governance . This would imply the establishing of governing roles and procedures but also the use of more proactive methods such as Privacy Impact Assessments (PIAs) [Pea11], automated analysis and policy enforcement.

Outlook

Apart from the implementation of accountability via governance, there is also ongoing research work on a technical representation of the concept of accountability [WABL+08]. In particular this relates to the question how security and privacy policies could be implemented rather than as ex ante (up front) controls through a policy aware infrastructure and by attaching information with adequate policy relevant information. Weitzner et al. [WABL+08] have presented an interesting perspective on this for the general Internet information architecture that can also be of interest for cloud computing.

Further to this, there is also the TClouds perspective on reducing the dependency on trustworthiness — and therefore accountability - of singular cloud providers.

3.9.2 Transparency

Overview

Transparency is a complementary concern to the demand for *accountability*. As it is explained in the previous section, a viable route towards trustworthy cloud computing is to call upon cloud provider self control while establishing widely accepted policies and binding rules within a legal framework of accountability. This in return also reduces the need for mandatory provider certification and external auditing when dealing with regulatory compliance. Hence, lighter options like voluntary certification become viable in combination with accountability and transparency.

With regard to transparency there is always a tradeoff between a necessary and reasonable level of transparency that a provider can offer as well as the goal to not constrain cloud provider innovation capability in how to organize internally for security, privacy and also general service delivery. This even becomes more complex when multiple providers are involved in delivering the cloud service — as in stacked services such as Dropbox.

Some transparency of cloud data processing is necessary in particular to verify that data location is compliant to regulations — e.g. when certain data needs to be kept in a geographic area. On the other hand, the cloud provider must also remain able to realize the dynamic data allocation between its data centers and external provider partners that is needed to perform its services most efficiently.

It should be noted that in addition to voluntary transparency there are also mandatory demands for transparency — e.g. with regard to notify customers of data breaches.

Relevance

It was discussed that accountability and transparency together provide a basis for implementing regulatory compliance in a trustworthy way.

Solutions

There are different means to realize transparency. A first option is allowing *auditability*. Whereby auditability means the openness of the cloud provider for security and privacy audits conducted by its customers. This includes insights into the providers technologies and processes.

Related to this is the option of *regular auditing by an external trusted party* that confirms the positive outcome of the audit via a *certificate*. As already discussed, these auditing options have some constraints and should not be seen as the only possibility towards increased cloud provider transparency.

Also, as already mentioned, *notifications to customers* are a further element in transparency. In particular, data breach notifications are also legally required. However there are further possibilities of notifications such as regarding outage times.

On the more technical side, transparency can also relate to allowing *monitoring*, and *automated validation*.

A connected element to technical transparency is the use of *open cloud standards* and in particular the further integration of monitoring capabilities and meta information in such standards.

Connected to this is the need for increased collaboration between cloud providers, standards organizations and customers in the definition of such standards — as e.g. through initiatives like the OMG Cloud Standards Customer Council (CSCC).

Outlook

From the viewpoint of cloud customers, there is a second aspect to transparency apart from contributing to compliance. Cloud computing has typically taken away control from IT departments and introduced a level of transparency between the customer and the providers (compared to inhouse provisioning of IT). The transparency discussed here can not completely resolve this situation but it can further change the relation between the cloud customer (that currently has mostly to accept that cloud services are executed with little transparency) and the cloud provide (that needs to find a way to provide better transparency while preserving its business flexibility and specific innovation). This can also address resistances against cloud computing.

3.10 Liability

3.10.1 Overview

Liability concerns in cloud computing relate to incidents such as data breach, data loss or disruption of the service and their consequences. Further issues can relate to illegal activities conducted with the help of the cloud or the storage and distribution of illegal or rights protected data with the help of the cloud. In that sense, not only cloud customers are interested in cloud provider liability but also organizations or individuals, for example, that have been targeted by cyber attacks via the cloud, right holders of data in the cloud.

A related concern springs from the fact that many cloud providers collaborate in the delivery of the service with other providers such as Dropbox that builds their storage service on top of Amazon S3. This leads to concerns about a chain of liability and to the question which organization would finally be responsible.

It is also distinguished if a potential liability incident has occurred by force majeure, by external influence out of control of the provider (e.g. through a cyber attack to the cloud) or by factors that are in control of the provider.

A further influence factor is the behavior of the cloud customer or data owner that can itself have had an influence on either the occurrence of the incident itself and also on the consequences. For example, the cloud customer might have ignored or violated ground rules of behavior defined by the cloud provider.

3.10.2 Relevance

Due to the severe business impact that most of the issues named above can have, it is almost impossible to limit the potential damage. Hence, also the potentials for liability claims are basically unlimited.

In a recent study [IDC11] prepared by IDC for the European Commission, liability was cited as a primary blocking factor for the adoption of cloud computing. Also in the European Commission's 2012 communication on Unleashing the Potential of Cloud Computing in Europe [Eur12], there is a concern expressed over "how liability for service failures such as downtime or loss of data will be compensated."

In fact, such liabilities are in general excluded by the terms and conditions of commodity cloud providers or they are strictly limited to a small factor of the volume of the cloud service contract.

3.10.3 Solutions

The issue of liability closely relates to the terms and conditions of the cloud provider and the contract of the provider with its customer as well as inter-provider contracts.

A central solution proposed in the context of the European Cloud Strategy [Eur12] is the establishing of safe and fair European model contract terms and conditions for the supply of cloud services. This shall in particular support consumers and small sized businesses.

Some of the concerns are also addressed in the upcoming Common European Sales Law. So the model terms and conditions shall in particular address further issues in cloud computing that are not addressed by the Common European Sales Law.

With particular regard to the issue of data breach and data protection, the European Commission has further proposed to the cloud industry to establish a code of conduct and to national data protection agencies to agree on European Binding Corporate Rules

3.10.4 Outlook

Liability will remain a central element of concern in cloud computing. The establishing of European model terms and conditions for cloud computing, the Common European Sales Law as well as a code of conduct and European Binding Corporate Rules on data privacy in the cloud, will all together provide for more trust in the use of cloud services.

There are further positions from cloud customers associations, e.g., from the European CIO Association (EuroCIO), that call for "risk sharing" solutions between cloud providers and cloud

customers. Currently, it seems that these are most likely to occur in individual and typically large scale cloud contracts. This in return also implies further obligations on the cloud customer — such as towards conducting due diligence and monitoring risk factors on the own side.

Chapter 4

Conclusion

Cloud computing is an ubiquitous phenomenon, with large adaptation in various fields and growing market penetration. We are currently witnessing that many businesses move to cloud computing or even exist solely on a cloud. Whereas cloud computing seems as an enabler of new computing strategies, it has its own share of risks. In this report, we described some of the risks associated with Infrastructure-as-a-Service cloud computing, their impact on business, and forthcoming security solutions.

Current cloud-computing solutions do not yet offer sophisticated security measures, as they are designed more with functionality in mind than with security and resilience. Many security issues have been recognized so far and have to be addressed in the future. This report lists concerns and solutions around the security of cloud computing and how it affects businesses using cloud computing. This will be a significant development, as it is widely expected in the IT industry that the cloud-security market will grow tremendously in the next 5 to 10 years. This development will start from the situation described in this report and along avenues explained here.

Specifically, cloud-specific security threats, which are intrinsic in the cloud model, will be addressed by novel products and features; moreover, security threats that exist in workloads migrating to the cloud have to be treated *in the cloud*. The most important business enablers for this development consist of forthcoming compliance rules and security standards for cloud computing.

Furthermore, security solutions will increasingly be delivered *from the cloud* — also to IT workloads that are not necessarily running in the cloud; this makes such security products scale much better than current solutions that are applied locally. This move also brings significant changes to the existing IT security market (however, this dimension goes beyond the scope of this report).

Bibliography

[Ale12] AlertLogic. Log manager product overview. Whitepaper, available at http://www.alertlogic.com/wp-content/uploads/datasheets/AL_log-manager_overview.pdf, 2012.

[Ama12] Amazon Web Services. Security bulletins. http://aws.amazon.com/security/security-bulletins/, 2012.

[App12] Application Security Inc. Effective database defense. http://www.appsecinc.com, 2012.

[BDF+03] Paul Barham, Boris Dragovic, Keir Fraser, Steven H, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. Xen and the art of virtualization. In *SOSP*, pages 164–177, 2003.

[Bey12] BeyondTrust. Beyondtrust home page. http://www.beyondtrust.com, 2012.

[BGM11] Sören Bleikertz, Thomas Groß, and Sebastian Mödersheim. Automated verification of virtualized infrastructures. In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, CCSW '11. ACM, 2011.

[BGPCV12] Lee Badger, Tim Grance, Robert Patt-Corner, and Jeff Voas. Cloud computing synopsis and recommendations. NIST special publication 800-146, National Institute of Standards and Technology (NIST), 2012. Available from http://csrc.nist.gov/publications/PubsSPs.html.

[BM05] Jeff Bonwick and Bill Moore. Zfs — the last word in file systems. 2005.

[BMO11] Beth E. Binde, Russ McRee, and Terrence J. OConnor. Assessing outbound traffic to uncover advanced persistent threat. Technical report, 2011.

[BSP+10] Sören Bleikertz, Matthias Schunter, Christian W. Probst, Dimitrios Pendarakis, and Konrad Eriksson. Security audits of multi-tier virtual infrastructures in public infrastructure clouds. In *Proceedings of the 2010 ACM workshop on Cloud computing security workshop*, CCSW '10. ACM, 2010.

[CCW06] Frederick Chong, Gianpaolo Carraro, and Roger Wolter. Multi-tenant data architecture. Microsoft Corporation — Article from http://msdn.microsoft.com/en-us/library/aa479086.aspx, 2006.

[Cen12] Cenzic. Cenzic hailstorm technology dynamic application security testing (DAST). Datasheet, available at http://www.cenzic.com/downloads/datasheets/Cenzic-datasheet-Hailstorm-Technology.pdf, 2012.

[CFJS12] Miguel Correia, Daniel Gómez Ferro, Flavio P. Junqueira, and Marco Serafini. Practical hardening of crash-tolerant systems. In *Proceedings of the 2012 USENIX conference on Annual Technical Conference*, USENIX ATC'12. USENIX Association, 2012.

[Cip12] CipherCloud. Building trust in the cloud. http://www.ciphercloud.com, 2012.

[Cis11] Cisco Corporation. Cisco 1000v virtual switch. White Paper, available from http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns892/ns894/at_a_glance_c45-492852.pdf, 2011.

[Cit10] Citrix. Citrix xenapp/xendesktop and xenserver: Solution guide. Citrix software — White Paper, available from http://www.citrix.com/English/ps2/products/documents_onecat.asp?contentid=683148&cid=White+Papers, 2010.

[CJ08] Nicola Clark and David Jolly. French Bank Says Rogue Trader Lost $7 Billion. The New York Times, published 25/01/2008, available at http://www.nytimes.com/2008/01/25/business/worldbusiness/25bank.html, 2008.

[CLM+12] Adam Cummings, Todd Lewellen, David Mclntire, Andrew P. Moore, and Randall Trzeciak. Insider Threat Study: Illicit Cyber Activity Involving Fraud in the U.S. Financial Services Sector. available at https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf, 2012.

[Clo10] Cloud Security Alliance. Top Threats to Cloud Computing v1.0. available at https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf, 2010.

[Clo11] Cloud Security Alliance. Security guidance for critical areas of focus in cloud computing (V3.0). Available from http://www.cloudsecurityalliance.org/, 2011.

[Clo12a] Cloud Security Alliance. Cloud data governance working group. https://cloudsecurityalliance.org/research/cdg/, 2012.

[Clo12b] CloudPassage. Cloudpassage home page. http://www.cloudpassage.com, 2012.

[CMTS09] Dawn Cappelli, Andrew Moore, Randall Trzeciak, and Timothy J. Shimeall. Common Sense Guide to Prevention and Detection of Insider Threats 3rd Edition Version 3.1. available at http://www.cert.org/archive/pdf/CSG-V3.pdf, 2009.

[CN12] William R Claycomb and Alex Nicoll. Insider Threats to Cloud Computing: Directions for New Research Challenges. available at http://www.cert.org/archive/pdf/CERT_cloud_insiders.pdf, 2012.

[Com11] Compuware. Performance monitoring from the first mile to the last mile. Brochure, available at https://compuware.my.salesforce.com/sfc/p/00000000hdRBsZQIQ2y8QE0sm0_Bs39vSG6Jsb0=, 2011.

[Cyb12] Cyber-Ark. Cyber-ark home page. http://www.cyber-ark.com, 2012.

[dVFJ+07] Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. Over-encryption: Management of access control evolution on outsourced data. In *VLDB*, pages 123–134, 2007.

[EMC12] EMC Corporation. Rsa securid authenticators. Datasheet, available at https://www.rsa.com/products/securid/datasheets/2305_h9061-sid-ds-0212.pdf, 2012.

[ENC+12] Vincent C. Emeakaroha, Marco Aurélio Stelmar Netto, Rodrigo N. Calheiros, Ivona Brandic, Rajkumar Buyya, and César A. F. De Rose. Towards autonomic detection of sla violations in cloud infrastructures. *Future Generation Comp. Syst.*, 28(7):1017–1029, 2012.

[ENI09] Cloud computing — Benefits, risks and recommendations for information security. European Network and Information Security Agency (ENISA), 2009.

[EU 10] EU Article 29 Working Party. Opinion 3/2010 on the principle of accountability. July 2010.

[EU 12] EU Article 29 Data Protection Working Party. Opinion 05/2012 on cloud computing. July 2012.

[Eur12] European Commission. Unleashing the potential of cloud computing in Europe. European Commission Communication COM(2012) 529 Final, http://www.ec.europa.eu/information_society/activities/cloudcomputing/docs/com/com_cloud.pdf, September 2012.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09. ACM, 2009.

[Gol72] R. P. Goldberg. *Architectural Principles for Virtual Computer Systems*. PhD thesis, Harvard University, 1972.

[Gre12] GreenSQL. The database security company. http://www.greensql.com/, 2012.

[Har07] Robin Harris. Data corruption is worse than you know. Blog post, available at http://www.zdnet.com/blog/storage/data-corruption-is-worse-than-you-know/191, 2007.

[HL10] Eran Hammer-Lahav. The OAuth 1.0 Protocol Specification. Technical report, 2010.

[HP 12] HP Corporation. HP WebInspect — test web applications by mimicking real-world attacks. Web page, available at http://www.hpenterprisesecurity.com/products/hp-fortify-software-security-center/hp-webinspect/, 2012.

[HZ12] Ray Hunt and Sherali Zeadally. Network forensics — an analysis of techniques, tools, and trends. *Computer*, 99(PrePrints):1–1, 2012.

[IBM12a] IBM Corporation. Ibm infosphere guardium data encryption. http://public.dhe.ibm.com/common/ssi/ecm/en/imd11887usen/IMD11887USEN.PDF, 2012.

[IBM12b] IBM Corporation. IBM Security AppScan: Application security and risk management. Whitepaper, available at http://public.dhe.ibm.com/common/ssi/ecm/en/rab14001usen/RAB14001USEN.PDF, 2012.

[IBM12c] IBM Corporation. Ibm security identity manager — drive effective identity management and governance across the enterprise. Datasheet, available at http://public.dhe.ibm.com/common/ssi/ecm/en/tid10294usen/TID10294USEN.PDF, 2012.

[IBM12d] IBM Corporation. Secure enterprise data and ensure compliance. IBM Software — White Paper, available from http://www-01.ibm.com/software/data/information-governance/, 2012.

[IDC10] IDC Corporation. Information governance in the cloud. White Paper, available from http://www.emc.com/collateral/analyst-reports/1010-idc-paper.pdf, 2010.

[IDC11] IDC Corporation. Quantitative estimates of the demand for cloud computing in Europe and the likely barriers to take-up. SMART 2011/0045 Open Workshop, proceedings available from http://cordis.europa.eu/fp7/ict/ssai/study-cc-workshop_en.html, 2011.

[IKC09] Wassim Itani, Ayman I. Kayssi, and Ali Chehab. Privacy as a service: Privacy-aware data storage and processing in cloud computing architectures. In *DASC*, pages 711–716, 2009.

[JG11a] Wayne Jansen and Timothy Grance. Guidelines on security and privacy in public cloud computing. NIST special publication 800-144, National Institute of Standards and Technology (NIST), 2011. Available from http://csrc.nist.gov/publications/PubsSPs.html.

[JG11b] Joyent and GuardTime. Achieving integrity in the cloud. White paper, available at http://www.guardtime.com/assets/joyent_guardtime_integrity_in_the_cloud_wp.pdf, 2011.

[KRB+12] Samuel Kounev, Philipp Reinecke, Fabian Brosig, Jeremy T. Bradley, Kaustubh Joshi, Vlastimil Babka, Stephen T. Gilmore, and Anton Stefanek. Providing Dependability and Resilience in the Cloud: Challenges and Opportunities. Springer Verlag, June 2012.

[KSSL06] Richard Kissel, Matthew Scholl, Steven Skolochenko, and Xing Li. Guidelines for media sanitization. NIST special publication 800-88, National Institute of Standards and Technology (NIST), 2006. Available from http://csrc.nist.gov/publications/PubsSPs.html.

[KW00] Poul H. Kamp and Robert N. M. Watson. Jails: Confining the omnipotent root. In *In Proceedings of the 2nd International SANE Conference*, 2000.

[KYB+07] Maxwell Krohn, Alexander Yip, Micah Brodsky, Natan Cliffer, M. Frans Kaashoek, Eddie Kohler, and Robert Morris. Information flow control for standard os abstractions. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, SOSP '07. ACM, 2007.

[lin05] *Virtualization of Linux based computers: the Linux-VServer project*, May 2005.

[Log11] LogRhythm. Solera and logrhythm partner solution brief. Datasheer, available at http://logrhythm.com/Portals/0/pdf/Solera_%20LogRhythm_Solution_Brief.pdf, 2011.

[LS01] Peter Loscocco and Stephen Smalley. Integrating flexible support for security policies into the linux operating system. In *Proceedings of the FREENIX Track: 2001 USENIX Annual Technical Conference*, pages 29–42, Berkeley, CA, USA, 2001. USENIX Association.

[Mar11a] Marathon Technologies Corporation. everrun mx: Powering application availability for the always on world. White Paper, available at http://www.marathontechnologies.com/documents/Meet_the_Needs_of_the_Always_on_World.pdf, 2011.

[Mar11b] Raffael Marty. Cloud application logging for forensics. In *Proceedings of the 2011 ACM Symposium on Applied Computing*, SAC '11. ACM, 2011.

[May07] D. Maynor. *Metasploit Toolkit for Penetration Testing, Exploit Development, and Vulnerability Research*. Elsevier Science, 2007.

[MBS12] Michael Maurer, Ivona Brandic, and Rizos Sakellariou. Self-adaptive and resource-efficient sla enactment for cloud computing infrastructures. In *IEEE CLOUD*, pages 368–375, 2012.

[MG11] Peter Mell and Timothy Grance. The NIST definition of cloud computing. NIST special publication 800-145, National Institute of Standards and Technology (NIST), 2011. Available from http://csrc.nist.gov/publications/PubsSPs.html.

[Mic07] Microsoft Corporation. Windows server 2008 Hyper-V product overview — an early look. White Paper, available from https://www.microsoft.com/en-us/download/details.aspx?id=6415, 2007.

[Mic12] Microsoft Corporation. Hyper-V virtual switch overview. Technical document, available from http://technet.microsoft.com/en-us/library/hh831452.aspx, 2012.

[MM10] Linda Musthaler and Brian Musthaler. Enchanced trust and data integrity in the public cloud. Web article, available at https://www.networkworld.com/newsletters/techexec/2010/101206bestpractices.html, 2010.

[Moz12] Mozilla Corporation. Mozilla persona documentation. Web-based documentation, available at https://developer.mozilla.org/en-US/docs/persona, 2012.

[MRF11] Y. Mundada, A. Ramachandran, and Nick Feamster. Silverline: Data and network isolation for cloud services. *2nd USENIX Workshop on Hot Topics in Cloud Computing*, 2011.

[Nas11] Nasuni Corporation. Nasuni Announces New Snapshot Retention Functionality in Nasuni Filer; Enables Fail-Safe File Deletion in the Cloud. Web page, at http://www.nasuni.com/news/press_releases/11-nasuni_announces_new_snapshot_retention, 2011.

[Okt12] Okta. On-demand identity and access management for the cloud. Datasheet, available at http://www.okta.com/pdf/Okta_Datasheet_Overview.pdf, 2012.

[ope12] Openvz project. Web site at http://wiki.openvz.org/Introduction_to_virtualization, 2012.

[Ora11a] Oracle Corporation. Consolidating applications with oracle solaris containers. White Paper, available from www.oracle.com/technetwork/server-storage/solaris10/documentation/consolidating-apps-163572.pdf, 2011.

[Ora11b] Oracle Corporation. Oracle vm 3: Application-driven virtualization. White Paper, available from http://www.oracle.com/us/technologies/virtualization/ovm3-app-driven-459334.pdf, 2011.

[Ora12] Oracle Corporation. Complete and scalable access management. White paper, available at http://www.oracle.com/technetwork/middleware/id-mgmt/overview/complete-and-scalable-access-mgmt-1697349.pdf?ssSourceSiteId=ocomen, 2012.

[Par10] Parallels. An introduction to os virtualization and pvc. White Paper, available from http://www.parallels.com/products/pvc/learn-whitepapers-more/, 2010.

[Pea11] Siani Pearson. Toward accountability in the cloud. *IEEE Internet Computing*, 15(4):64–69, 2011.

[Pin12] Ping Identity. Cloud identity management — solutions guide. White paper, available at https://www.pingidentity.com/unprotected/upload/Enterprise_Solutions_Guide.pdf, 2012.

[Pon11] Ponemon Institute LLC. The Risk of Insider Fraud: U.S. Study of IT and Business Practitioners. available at http://resources.idgenterprise.com/original/AST-0060002_Ponemon_2011_Insider_Fraud_Survey_Results.pdf, 2011.

[PPK+09] B. Pfaff, J. Pettit, T. Koponen, K. Amidon, M. Casado, and S. Shenker. Extending networking into the virtualization layer. In *Proceedings of Workshop on Hot Topics in Networks (HotNets-VIII)*, 2009.

[PRZB11] Raluca Ada Popa, Catherine M. S. Redfield, Nickolai Zeldovich, and Hari Balakrishnan. Cryptdb: protecting confidentiality with encrypted query processing. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, SOSP '11. ACM, 2011.

[Que10] Quest Software (now Dell). Maximizing virtual machine performance. White Paper, available at http://communities.quest.com/servlet/JiveServlet/downloadBody/8328-102-1-10079/Maximizing%20VM%20Performance%201.2.pdf, 2010.

[Red08] Red Hat. KVM — kernel based virtual machine. White Paper, available from http://www.redhat.com/resourcelibrary/whitepapers/doc-kvm, 2008.

[RR06] David Recordon and Drummond Reed. Openid 2.0: a platform for user-centric identity management. In *Proceedings of the second ACM workshop on Digital identity management*, DIM '06. ACM, 2006.

[Saf12] SafeNet Inc. SafeNet Cloud Security. Web page, at http://www.safenet-inc.com/Products/cloud/SafeNet_Cloud_Security/, 2012.

[Sal11] Matthew Saltmarsh. UBS Blames $2 Billion Loss on Rogue Trader. The New York Times, published 15/09/2011, available at http://dealbook.nytimes.com/2011/09/15/ubs-reports-2-billion-loss-to-rogue-trader, 2011.

[sam05] Security assertion markup language (saml) v2.0, 2005.

[Sav11] Marcia Savage. Cloud availability and resiliency: Planning for failure. Interview, available at http://searchcloudsecurity.techtarget.com/news/2240039009/Cloud-availability-and-resiliency-Planning-for-failure, 2011.

[SC05] Tom Scavo and Scott Cantor. Shibboleth architecture — technical overview. Technical report, 2005.

[SH07] James P.G. Sterbenz and David Hutchison. Resilinets definitions. Web page, available at https://wiki.ittc.ku.edu/resilinets_wiki/index.php/Definitions#Resilience, 2007.

[SHJ+11] Juraj Somorovsky, Mario Heiderich, Meiko Jensen, Jörg Schwenk, Nils Gruschka, and Luigi Lo Iacono. All your clouds are belong to us: security analysis of cloud management interfaces. In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*, CCSW '11. ACM, 2011.

[Sla11] Slavik Markovich. THREE CLOUD-COMPUTING DATA SECURITY RISKS THAT CANT BE OVERLOOKED. available at https://blog.cloudsecurityalliance.org/2011/03/21/, 2011.

[Sma12] SmartSignin. Smart single sign-on. White paper, available at http://www.smartsignin.com/resources.aspx, 2012.

[SMS+12] Juraj Somorovsky, Andreas Mayer, Jörg Schwenk, Marco Kampmann, and Meiko Jensen. On breaking saml: be whoever you want to be. In *Proceedings of the 21st USENIX conference on Security symposium*, Security'12. USENIX Association, 2012.

[Soa10] Sunil Soares. *The IBM Data Governance Unified Process*. MC Press Online, 2010.

[Sol12] SolarWinds. Estimating log generation for security information event and log management. Whitepaper, available at http://content.solarwinds.com/creative/pdf/Whitepapers/estimating_log_generation_white_paper.pdf, 2012.

[Spl12] Splunk. Splunk for cyber threat analysis. White paper, available at http://www.splunk.com/web_assets/pdfs/secure/Splunk_for_Cyber_Threat.pdf, 2012.

[SRGS12] Nuno Santos, Rodrigo Rodrigues, Krishna P. Gummadi, and Stefan Saroiu. Policy-sealed data: a new abstraction for building trusted cloud services. In *Proceedings of the 21st USENIX conference on Security symposium*, Security'12. USENIX Association, 2012.

[Sta07] John Stanik. A conversation with jeff bonwick and bill moore. *Queue*, 5(6), 2007.

[Sto11] Storage Networking Industry Association. Information technology — cloud data management interface (CDMI). Technical position, available at http://snia.org/sites/default/files/CDMI%20v1.0.2.pdf, 2011.

[SVC10] Javier Salido, Patrick Voon, and Doug Cavit. A guide to data governance for privacy, confidentiality, and compliance. White Paper, Microsoft Corp., available from http://www.microsoft.com/, 2010.

[Sym11] Symantec. The secure cloud: Best practices for cloud adoption. White Paper, available from http://www.symantec.com/, 2011.

[Tha12] Thales. Information technology security. http://www.thales-esecurity.com, 2012.

[The11] The Web Application Security Consortium. Web application security scanner list. Web page, available at http://projects.webappsec.org/w/page/13246988/Web%20Application%20Security%20Scanner%20List, 2011.

[TLLP12] Yang Tang, Patrick P. C. Lee, John C. S. Lui, and Radia Perlman. Secure overlay cloud storage with access control and assured deletion. *IEEE Trans. Dependable Sec. Comput.*, 9(6):903–916, 2012.

[Var10] Jinesh Varia. Architecting for the cloud: Best practices. White Paper, available at https://jineshvaria.s3.amazonaws.com/public/cloudbestpractices-jvaria.pdf, 2010.

[VAS12] VASCO. DIGIPASS pack for remote authentication. Web page, available at https://www.vasco.com, 2012.

[VMW08] VMWare Corporation. Architecture of vmware esxi. White Paper, available from http://www.vmware.com/files/pdf/ESXi_architecture.pdf, 2008.

[VMW12a] VMWare. Vmware horizon application manager. Datasheet, available at https://www.vmware.com/files/pdf/horizon/VMware-Horizon-App-Manager-Datasheet.pdf, 2012.

[VMw12b] VMware Corporation. Wonderware system platform 2012 — superior agility with vmware vsphere 5. White Paper, available at https://www.vmware.com/files/pdf/techpaper/vmware-vsphere-wonderware-system-platform-tech-guide.pdf, 2012.

[Vol12] Voltage Security. Voltage security home page. http://www.voltage.com, 2012.

[Vor12] Vormetric. Enterprise encryption and key management simplified. http://www.vormetric.com, 2012.

[VW08] Craig Valli and Andrew Woodward. The 2008 Australian study of remnant data contained on 2nd hand hard disks: the saga continues. 2008.

[WABL+08] Daniel J. Weitzner, Harold Abelson, Tim Berners-Lee, Joan Feigenbaum, James Hendler, and Gerald Jay Sussman. Information accountability. *Communication of the ACM*, 51(6), June 2008.

[Whi12] WhiteHat Security. WhiteHat Security Sentinel Service — website risk management solutions. Datasheet, available at https://www.whitehatsec.com/assets/DS/DS_4pgSentinel040611.pdf, 2012.

[WKH+08] Thomas Weigold, Thorsten Kramp, Reto Hermann, Frank Höring, Peter Buhler, and Michael Baentsch. The Zurich Trusted Information Channel — an efficient defence against man-in-the-middle and malicious software attacks. In *Proceedings of the 1st international conference on Trusted Computing and Trust in Information Technologies: Trusted Computing - Challenges and Applications*, Trust '08, 2008.

[Xu06] Howie Xu. 5 predictions: How virtual networking will change the network industry. CIO — News article from http://www.cio.com/article/645344/5_Predictions_How_Virtual_Networking_Will_Change_the_Network_Industry, 2006.

[ZBWM08] Nickolai Zeldovich, Silas Boyd-Wickizer, and David Mazières. Securing distributed systems with information flow control. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'08. USENIX Association, 2008.

[Zen12] Zenoss Corporation. Zenoss vCloud Monitoring. Web page, at http://www.zenoss.com/solution/vcloud-monitoring, 2012.

[Zet10] Kim Zetter. Ex-Googler Allegedly Spied on User E-Mails, Chats. available at http://www.wired.com/threatlevel/2010/09/google-spy/, 2010.