Critical Services in the Cloud: Understanding Security and Resilience Risks

Thomas Hecht and Paul Smith
AIT Austrian Institute of Technology
Vienna, Austria
Email: {thomas.hecht|paul.smith}@ait.ac.at

Marcus Schöller
NEC Europe Ltd.
Heidelberg, Germany
Email: marcus.schoeller@neclab.eu

***Abstract*—The promise of low costs, adaptation to customer load, and fast service roll-out has made cloud infrastructures a primary choice for many service providers. So far, this has been largely for end-user and enterprise services. Recently, the cloud paradigm is being considered by service providers of critical infrastructures. A prominent example of this is ETSI's Industry Specification Group (ISG) on Network Function Virtualization, which provides guidelines on how to move telecommunications services to the cloud. But other critical infrastructure providers are following closely. Common characteristics of critical infrastructure services, such as network functions, are their high requirements for service dependability and security. In this paper, we present a risk assessment method for assessing the risks of moving critical infrastructure services to the cloud. To achieve this, we have extended a well-established information security risk assessment process and developed an extensive cloud-specific fault and challenge catalogue.**

***Index Terms*—Network functions virtualization, critical infrastructures, risk assessment, security, resilience**

I. Introduction

Because of its many benefits, moving services to the cloud (the so-called *cloudification* of services) is a trend that is unlikely to abate. The benefits of using the cloud can include reduced operating costs and the ability to flexibly scale services in response to demand. Furthermore, due to the use of virtualization, services that are deployed in the cloud can be more resilient to problems such as underlying hardware failures and power outages. To date, the cloud has been primarily used to host enterprise and end-user (consumer) services. However, the operators of *critical infrastructures* are considering using the cloud to implement their high-assurance IT services.

For these high-assurance IT services, there are stringent security and resilience requirements, which are arguably higher than for enterprise and end-user services. In some cases, the failure of IT services that support critical infrastructures could have safety-related implications and thus strict regulatory frameworks for security and dependability have been defined. Such safety implication become obvious for traffic management services, for example, which we consider in this paper. Another trend that can be observed is the cloudification of real-time services that, for example, provide an enterprise telephone system (PBX). Moving beyond the realisation of PBX services in the cloud, ETSI's Industry Specification Group (ISG) on Network Function Virtualization (NFV) is providing guidelines for moving telecommunications services, e.g., a 3GPP evolved packet core (vEPC) or a Broadband Remote Access Router (vBRAS), to the cloud. These systems are generally acknowledged as critical.

Despite the many benefits and drive towards the cloudification of critical infrastructure high-assurance IT services, the security and resilience implications of doing so are arguably not well understood. Previous research has examined the security faults and challenges that are associated with using the cloud [1], and some risk assessment approaches have been proposed [2]. However, a consistent organisation of these cloud-specific faults and challenges, and how they can be applied as part of a risk assessment to a cloud infrastructure that is targeted to supporting critical infrastructure IT services, has not been proposed. We suggest that it is important to address this shortcoming, so that operators can get a clear understanding of the security and resilience implications of cloud usage, before choosing to cloudify their services.

In this paper, we present an overview of an extensive[1] cloud-specific fault and challenge catalogue; entries are organised into categories that relate to different aspects of cloud usage [3]. The catalogue was derived via the systematic analysis of related work and a cloud architectural model that is targeted to support high-assurance services, such as those being cloudified for telecommunications. The catalogue can be used as input to a risk assessment method for determining the security and resilience of cloud usage. Accordingly, we present an extension to a risk assessment method that operators can use to understand the risks that are specifically associated with cloud usage. In order to evaluate the utility of these items, they have been applied to a traffic management system.

To evaluate the risk assessment process, we have decomposed the traffic management system into its high-level processes and modelled the resources that support their correct operation. Subsequently, these resources are associated with the faults from the catalogue; finally the associated challenges are modelled that can activate these faults and consequently lead to a service failure. We have used a well-established information security management tool, called Verinice, to implement these models and to conduct the assessment.

[1]At the time of writing, the catalogue contains over 170 entries, and can be downloaded from https://seccrit.eu

The rest of the paper is organized as follows. First, in Sec. II, we present the state of the art in assessing risks in general and for critical infrastructure services, in particular. Afterwards, we briefly introduce the architectural framework that we have developed in Sec. III, followed by an excerpt from our fault and challenge catalogue in Sec. IV. The extended risk assessment process, which can be used to assess the risks associated with moving critical infrastructure services to the cloud, is presented in Sec. V. Our use-case is described in Sec. VI, along with an evaluation and key insights in Sec. VII.

II. Related Work

The ISO 27000 series of standards relate to information security management, risk management and security controls. They provide a code of practice for information security, which outlines potential controls that may be implemented, and guidance on the development and use of measures and measurement for the assessment of the effectiveness of an implemented information security management system and controls. NIST SP800-53 [4] is similar to the ISO 27000 standards, and lists and classifies security control requirements, from which each security process should extract a baseline. Furthermore, it encompasses the risk management process, specifying all activities for the selection of security controls to their application to an organisation's information systems. Meanwhile, ISO 31000 provides principles, a framework and a process for managing risk, and aims at improving the identification of opportunities and threats, and effectively allocating and using resources for risk treatment. ISO 31010 provides information on risk assessment concepts, processes and the selection of risk assessment techniques.

Behnia *et al.* [5] point out the range of differences between security analysis methodologies, including OCTAVE [6], IS-RAM [7], CORAS [8] and several others. Furthermore, the European Network and Information Security Agency (ENISA) maintains a repository of risk assessment standards, methods and tools from a European perspective [9]. The Operationally Critical Threat, Asset and Vulnerability Evaluation (OCTAVE) method can be used for identifying and managing information security risks. It contains methods, techniques and tools for an *asset-driven* evaluation approach, focusing on security practices and strategic issues, and self-direction. Similarly, Magerit is a risk analysis and management methodology that has been developed in Spain [10]. In a similar manner to OCTAVE, it is driven by an analysis of the assets that are associated with an organisation.

These risk assessment and management approaches can be used to assess the risks to critical infrastructure services. However, they do not support the assessment of cloud-specific risks. For example, when deploying IT services in the cloud, specific threats and vulnerabilities must be evaluated, which are not accounted for systematically in previous work. Furthermore, when critical infrastructure operators are considering deploying their ICT services in the cloud, to the best of our knowledge, previous work has not provided specific guidance regarding how to evaluate the relative cybersecurity risk with respect to an existing deployment. It is these shortcomings that we address in this paper.

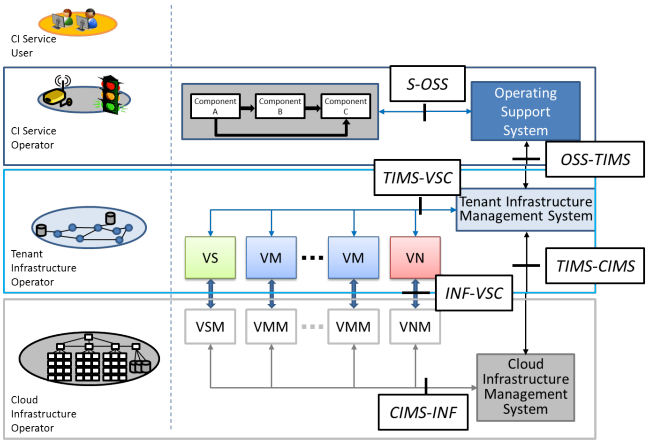

Fig. 1. The management view of our layered architecture model for critical infrastructure service deployment and operation

III. A Cloud Architectural Model for Critical Infrastructure Services

To support the creation of our fault and challenge catalogue, and the execution of the risk assessment process that is summarised in Sec. V, we made use of a cloud architectural model that is targeted at supporting high-assurance critical infrastructure services. The model is described in detail in [11], and includes multiple views on a cloud architecture. A short top-down introduction to the layers of abstraction in the model and the management view is given here, as it is used in the subsequent sections. Amongst other items, Fig. 1 summarises the components that exist in the management view and highlights reference points.

A. Cloud Architectural Model Layers

The cloud architectural model describes four layers of abstraction:

a) User Level: At the top level of the model, the critical infrastructure *service user* remotely accesses the critical infrastructure service. For instance, an urban traffic management operator could observe and control the traffic flow within a city, using web-based interfaces, as well as various distributed sensors that deliver measurement data.

b) Service Level: This level is managed by the critical infrastructure *service provider*. The critical infrastructure service can be composed of several interacting components that collectively provide the actual service. The service provider monitors the service operation and performance at this level. The service components usually either provide the application or the platform that are required at user level while being instantiated on the virtual infrastructure resources.

c) Tenant Infrastructure Level: At this level, a virtual infrastructure that consists of virtual compute, virtual storage, and virtual network resources is provided. Those virtual resources are managed by a *tenant infrastructure operator*. Several tenants are typically hosted within one cloud infrastructure, which is described next. The tenant infrastructure

operator may provide either a pure virtual infrastructure (IaaS), or some basic services as a platform (PaaS) to the service provider.

d) Physical Cloud Infrastructure Level: this level provides the underlying physical compute, storage, and network resources, which are hosted in data centres. These are administered by *cloud infrastructure operators*. This level usually provides virtual resources (IaaS) to its upper level, i.e., the tenant. The virtualization solution provides (a certain degree of) isolation between the different tenants that are multiplexed onto the same physical infrastructure, and thus permits the sharing of resources. For increasing resource efficiency, the cloud infrastructure provider can usually transparently move virtual resources across its physical infrastructure, unnoticed by the tenant.

B. Cloud Service Management

Critical infrastructure services can be managed in our cloud architectural model, depicted in Fig. 1, in the following way. Starting at the service level, the service operator controls their services by using an *Operating Support System (OSS)* component. New services are introduced to the system by providing a description of the service [12] and the software component executables. Both are provided by a service vendor and fed into the OSS via the `S-OSS` interface. Service onboarding is an offline activity, in which the service component executables are put into one or more cloud image repositories; this is outside the scope of our analysis. The OSS is responsible for requesting service instantiation and termination from the *Tenant Infrastructure Management System (TIMS)*. Additionally, the OSS sets operational policies for the TIMS via the `OSS-TIMS` interface.

To instantiate a service, in a first step, the TIMS has to request adequate resources, specifying the necessary constraints like co-location restrictions and redundancy requirements, from one or more *Cloud Infrastructure Management Systems (CIMS)* by passing corresponding resource descriptions to them. Such requests are passed via the `TIMS-CIMS` interface. The CIMS is provisioning the requested virtual resources by slicing its physical resources, i.e., compute, network, and storage resources, via the `CIMS-INF` interface. Second, the TIMS has to execute the lifecycle management of the service components, e.g., reconfiguration of the service chain, fault correlation and management, or service scaling; this is performed via the `TIMS-VSC` interface (VSC: Virtual Service Components). Each of the service components is tied to a virtual resource provided by CIMS using the `INF-VSC` interface.

IV. A Cloud-specific Fault and Challenge Catalogue

To support the analysis of the risks of moving high-assurance services to the cloud, such as a network function or components of a traffic management system, we have developed a *cloud-specific fault and challenge* catalogue. As part of a risk assessment, the items in this catalogue can be examined for their likelihood of occurrence (challenges) and severity (faults) – we summarise how this can be achieved in Sec. V. To create the catalogue, we performed an extensive literature survey, drawing on sources from academia and other areas, such as the Cloud Security Alliance (CSA). Furthermore, we carried out an analysis of different cloud deployment models, as described below.

The resulting catalogue is organised into a number of categories, as shown in Fig. 2, which are primarily based on the five National Institute of Standards and Technology (NIST) essential cloud characteristics: *(i)* on-demand self-service; *(ii)* broad network access; *(iii)* resource pooling; *(iv)* rapid elasticity; and *(v)* measured service [13]. Additionally, we organise items in the catalogue into domains that reflect the importance of *virtualization as a key enabling technology*; the often overlooked *organisational issues* associated with using cloud services [14]; and the provisioning of the *underlying physical infrastructure*. For each of these categories, we have developed a fault and challenge table, which we introduce below.

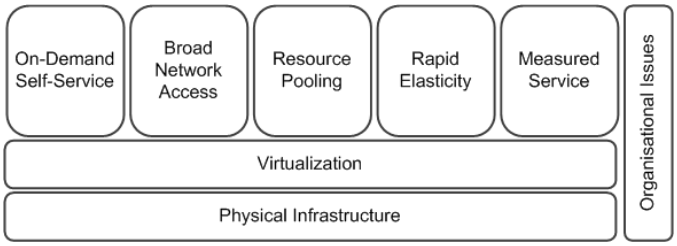

Fig. 2. Categories of fault and challenge domains

A. Failure Modes

As mentioned earlier, we performed an analysis of the failure modes that can be associated with different service deployment scenarios in the cloud. The findings of this analysis contributed to the fault and challenge catalogue. Depending on the type of service deployment, the impact of failure will vary.

The classical "box-model" mode of operation is shown as *Option 1* in Fig. 3. Note that the figure shows a logical view – the service usually consists of a set of components cooperatively providing the specified service (these components are represented by the box that is labelled *Service1*). The simplest approach to virtualizing such a service is to take the existing software, install it into a virtual machine (VM) image and execute it on virtual resources provided by the hypervisor (*Option 2*). The additional software introduced to the system, i.e., VMM and VM, adds new failure modes which did not exist in the box-model. For example, a programming bug of the VMM is triggered by certain VM configurations.

In order to maximise hardware utilization, physical resources are sliced into a set of virtual resources and provided to a multitude of tenants. Thus, several service (component) instances can be hosted on the same physical host (*Option 3*). This adds another set of failure modes to the system, e.g., a potentially negative performance impact of *Service2*

TABLE I
AN EXAMPLE VIRTUALIZATION-RELATED FAULT

Name	*Description*	*Aspect*	*Mapping*
Failure and challenge detection mechanisms not in place	Mechanisms to detect challenges to or failures of CIMS and infrastructure components are not in place.	C-I-A	CIMS, compute, network, storage

TABLE II
EXAMPLE VIRTUALIZATION CHALLENGES THAT ARE RELATED TO THE FAULT DESCRIBED IN TABLE I

No.	*Name*	*Description*	*Aspect*	*Mapping*
VMC-1	Side-channel attack	Cross virtual machine side-channel attack leading to information leakage	C	INF-VSC
VMC-2	Virtual resource intrusion	An attacker breaks into the resources of other tenants	C-I	INF-VSC
VMC-3	VIM intrusion	Specialized attack against the VIM to gain access to tenant information and free virtual resources	C-I	TIMS-CIMS
VMC-4	Uncontrolled request or illegitimate for resources	A malicious tenant or faulty service instance requests an unusually (and damaging) amount of resources	I-A	TIMS-CIMS
VMC-5	VIM session hijacking or riding	Weak authentication or communication channel protection between tenants and the VIM	C-I	TIMS-CIMS

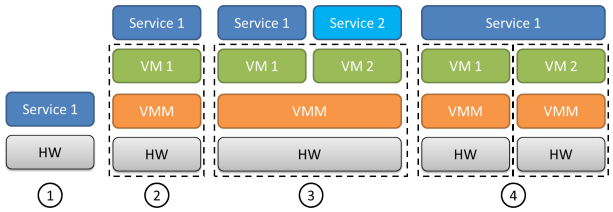

Fig. 3. Deployment options of CI services on a cloud infrastructure

on *Service1*, if resource isolation does not work or an attack against a VM leads to collateral damage.

Finally, a critical infrastructure service can also span across physical and virtual host boundaries. If a service is composed from multiple service components, each component will be deployed into its own virtual machine on the same physical host or on a different physical host on the cloud (*Option 4*). Again, new failure modes are introduced, e.g., simultaneous failures of multiple service components due to a failure of the underlying hardware or service failures due to communication failures between its components. More concretely, the service will fail if any component of the component chain fails and no redundancy has been provided.

B. Faults and challenges of virtualization

To give an indication of the items that are included in the fault and challenge catalogue, we explore the fault that *failure and challenge detection mechanisms are not in place* – see Table I. This fault can be activated by a number of different challenges, leading to an error or even a failure. For example, a type of attack that is introduced by the use of virtualization is VM side-channel attacks (see *VMC-1* in Table II). If this attack is successful, the confidentially of authentication keys cannot be guaranteed any more. In a second step the attacker can attempt a virtual resource intrusion attack (*VMC-2*). This leads to further loss of confidential data, as well as a loss of the integrity of this resource. Alternatively, an attacker can try to execute a man-in-the-middle attack on the management channel between the virtual resource and the management system (*VMC-5*), which again can impact the confidentially and integrity of that system. Finally, an attacker or a faulty software instance can request a large amount of virtual resources (*VMC-4*) resulting in a denial of resources to other tenants. Any of these challenges should be detected by an appropriate mechanism built into the system.

The full list of faults and challenges for this and all other categories can be found in [15]. For each of the faults and challenges, we highlight the primary dependability and security objectives they affect: Availability (A), Confidentiality (C), and Integrity (I). Table I discusses the faults which might exist in the system; for each a name and a description is provided as well as impacted aspect, and the component of our cloud architectural model (see Fig. 1). Moreover, we link each fault to the challenges by which they might be triggered. These challenges are presented in Table II, but with the interface the challenge interacts with the system to trigger the fault.

V. ASSESSING THE RISK OF CLOUD DEPLOYMENT

Building on the challenge and fault catalogue, we have developed a process that critical infrastructure operators can use to determine the risks associated with moving their high-assurance services to the cloud. The process is intended to be realised as an extension to an existing asset-driven risk assessment. As such, its implementation should incur a relatively minor additional overhead. A summary of the risk assessment method is depicted in Fig. 4, which shows two stages. The first stage relates to the implementation of a *base* asset-driven risk assessment, which results in the current, non-cloud related, risks being assessed. In this stage, the assets, including services and data, that are being considered for migration to the cloud are modelled. In the second stage, the risks associated with these assets are considered in the context of a target cloud deployment. To support the practical application of our approach, the steps that are shown in Fig. 4 follow those that are used to realise the process in the open-source Verinice ISMS tool[2]. Despite this, the overall approach

[2]Verinice: http://www.verinice.org

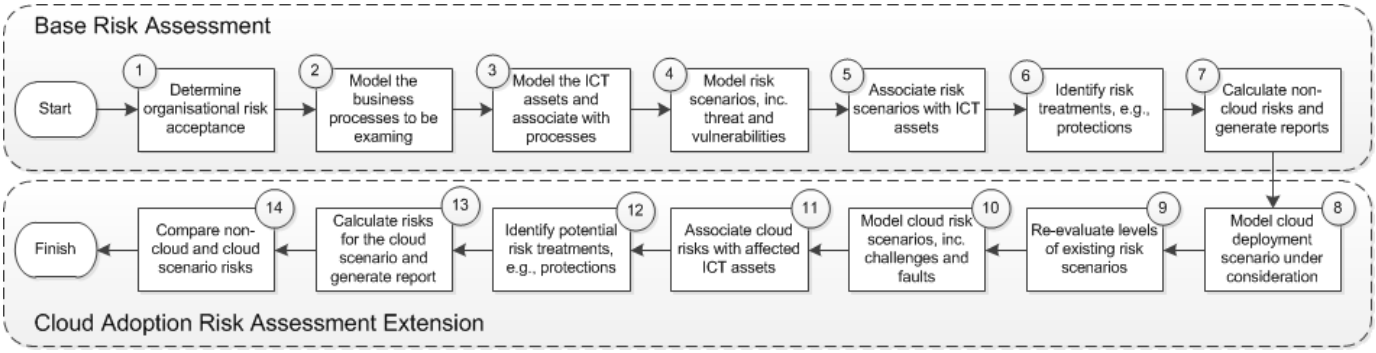

Fig. 4. An extended risk assessment process that can be used to determine the risks associated with moving services to the cloud.

is readily applicable to other well-established asset-driven risk assessment methods, such as Magerit and OCTAVE, for example.

To realise the risk assessment extension that is shown in Fig. 4, initially the cloud deployment under consideration must be modelled. This involves modelling the target cloud infrastructure and its relationship to the previously defined critical infrastructure assets, including the dependencies between them. For the purpose of the analysis of the traffic management system, we modelled the cloud infrastructure using the architectural model that is described in Sec. III. A benefit of this approach is that once a cloud infrastructure has been modelled by a provider, it can be reused to analyse the risks of migrating further critical infrastructure services.

Subsequently, the next stage is to re-evaluate the risks associated with the existing challenge scenarios, which are defined as part of the base assessment. Deploying assets in the cloud can alter the nature of these challenge scenarios in a number of ways, including *(i)* completely negating them, e.g., consider moving assets to a cloud data centre that is not affected by local environmental phenomena such as flooding; *(ii)* reducing risk, e.g., consider the fault-tolerance benefits associated with the use of virtualization technology; and *(iii)* increasing risk – the challenges of implementing authentication schemes in the cloud are well understood [16]. Of course, existing challenge scenarios may remain unaffected by the use of cloud.

The cloud-specific risk scenarios must then be modelled, which include the relevant challenges and faults. For this purpose, the catalogue that is outlined in Sec. IV can be applied. For each of the challenges, the likelihood of occurrence should be estimated and the severity of faults assessed. Subsequently, the risk scenarios need to be associated with the critical infrastructure assets and the model of the cloud infrastructure, as appropriate.

Finally, the risk values are calculated for each of the scenarios; these are determined by considering the aforementioned values (probability and severity), and the (business) impact associated with the compromise of the assets the risks are related to. Using the Verinice tool, this step can be realised automatically. Based on the values from the original base risk assessment and those determined using our process extension, a critical infrastructure operator can better understand the risks associated with deploying their assets in the cloud.

VI. Cloudifying a Traffic Management System

To demonstrate our risk assessment process, we have considered the movement of the components of a traffic management system to the cloud. An overview of the components associated with such a system are depicted in Fig. 5. This represents an abstraction of a more detailed architectural view of the traffic management system that we used for our assessment.

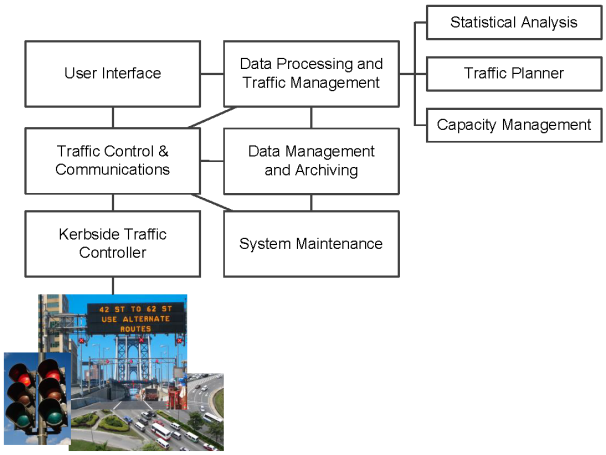

Fig. 5. An overview of the traffic management that was considered for cloudification

At the bottom of the figure, the distributed *kerbside traffic control* component is shown, which controls and monitors traffic lights, information displays and inductive loops. These distributed elements are controlled by a *traffic control and communications* component, which acts as a central point of interaction for the other components in the system. The data that is collected from the kerbside equipment is distributed from this component. Based on this data, a set of *data processing and management* components perform *statistical analysis* functions, create traffic management strategies (using the *traffic planner*), and perform *capacity management* operations, for example. The output from these activities is made available to human operators via a sophisticated *user interface*. Long-term archiving of traffic and system performance data

is supported by a *data management and archiving* component. Building on the archival data, a *system maintenance* component performs data analysis to determine whether the overall performance of the system can be improved. Along with short-term control strategies, these improvements are communicated to the kerbside systems via the traffic control and communications component.

Deployments of the traffic management system have been running for several years in dedicated data centres, which requires that each city using it to install and operate their own hardware infrastructure. Moving components of the system into a virtualized cloud environment is expected to provide several benefits, including centralized service provisioning and management, and simpler application update procedures. All the components depicted in Fig. 5 are being considered for migration to the cloud, except for those distributed at the kerbside. The complexity of the application makes it necessary to deploy the various functional components onto multiple virtual machines with corresponding virtual interconnects.

From a security and resilience perspective, the use case highlights the stringent requirements for service availability and integrity for such a system. For example, in order to manage traffic incidents in a city, the system must be highly available – outages could have a significant impact. The integrity of measurement data and control messages is of the utmost importance, as compromises to these could result in sub-optimal traffic management strategies being proposed (by the *traffic planner* component) and deployed throughout the city. Moreover, the security of archival data is important, as it used for city development plans that incur high costs.

VII. Evaluation

We evaluated the security and resilience risks that are associated with moving the traffic management system described in Sec. VI using our challenge and fault catalogue, and the assessment process that is summarised in Sec. V.

To achieve this, we initially determined the processes the traffic management system implements (Step 2 in Fig. 4). In total, we identified nine high-level processes that relate to implementing traffic control strategies, maintaining the traffic control system itself, and providing the general public with information about the status of traffic. Based on our knowledge of the system, we evaluated the impact of the processes being compromised from a confidentiality, integrity and availability perspective. These values were calibrated with the operator of the traffic management system. Subsequently, we modelled the components and sub-components of the traffic management system, and their dependencies, which are outlined in Fig. 5. These *assets* are then associated with the processes that were previously defined, i.e., assets were linked to the processes they support. As defined in Sec. V, we then defined a number of risk scenarios that relate to the traffic management system assets. The German Federal Office for Information Security (BSI) Baseline Security Catalogue[3] was used as a basis for identifying these risk scenarios. Example scenarios related to environmental aspects, such as fire and flooding, data protection issues, and network and system management. Challenges and faults were associated with the risk scenarios and assessed for the likelihood of occurrence and severity, respectively. The product of these values, and the impact values that are associated with the processes, are used to determine the severity of the risk scenarios.

In a second stage, we modelled the cloud architectural model that is outlined in Fig. 1. This included defining the dependencies between the different components at the various architectural layers. Subsequently, the dependencies between the traffic management system and the cloud architectural model components were defined. For example, there are direct dependencies between the *statistical analysis* component and the *Operating Support System* and *Virtual Machine Manager* components, amongst others. In Fig. 4, this activity relates to Step 8. Subsequently, the risk scenarios that relate to the base risk assessment were reconsidered in the context of the cloud deployment. A number of risks related to in-house *network and service management* were considered no longer relevant, for example. Having reassessed the existing risks, the faults and challenges from our catalogue were applied and assessed, as described in Sec. V.

Inevitably, to some extent, such an assessment will reflect the subjective opinions of the assessors. This activity has been carried out in the context of a research project that includes experts on the traffic management system under scrutiny, and vendors and operators of cloud solutions. To ameliorate this problem, the analyses of the authors were validated and in some cases adjusted in workshops with these experts.

A. Key Findings

Based on the reports generated by the Verinice tool, which was used to implement the risk assessment, we can identify some key findings. It should be noted these findings are based on an assessment that does not consider the implementation of protection measures (see Steps 6 and 12 in Fig. 4), which could reduce risk levels. The primary purpose of the report generated by the Verinice tool is to alert the assessor to the *high risk scenarios*, i.e., those that are considered unacceptable and protection measures must be urgently implemented for.

In our assessment, the largest number of high risk scenarios, when considering the cloud-based deployment, relate to challenges and faults that could affect the *availability* of the traffic management system's processes. In part, this is due to the relatively high impact associated with a lack of availability for such a system. Additionally, this relates to the items in our catalogue primarily affecting this property – for example, a third of the items in the *virtualization* category exclusively impact upon the availability of the assets they affect; much more than confidentiality and integrity.

Fig. 5 shows the number of medium and high risk scenarios for the different risk categories, again considering the cloud deployment. A risk scenario includes a set of challenges and

[3]BSI IT-Grundschutz-Kataloge: https://www.bsi.bund.de/DE/Themen/ITGrundschutz/ITGrundschutzKataloge/itgrundschutzkataloge_node.html

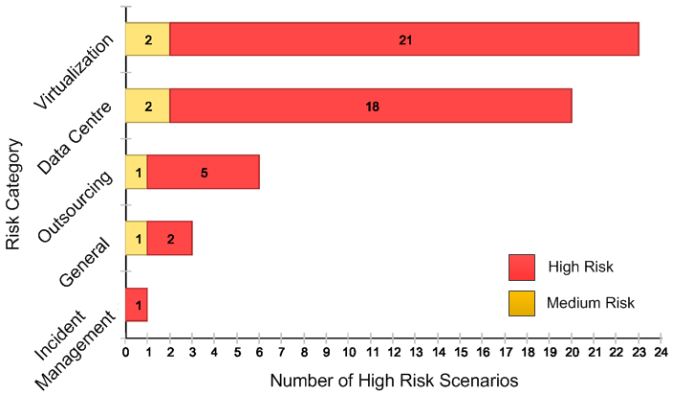

Fig. 6. The number of high and medium risk scenarios for the different risk categories for the cloud deployment

faults that affect an asset. It can be seen in Fig. 5 that the greatest number of high risk scenarios relate to the challenges and faults associated with virtualization, such as those presented in Tables I and II. The next highest category relates to *data centre* risks, such as the *failure of protection measures*, including physical equipment failure. Whilst this category includes fewer high risk scenarios compared to the non-cloud deployment, because of the use of virtualization, it is still prominent. A key challenge for a critical infrastructure service provider is accurately assessing these risks for cloud data centres that are not under their organisational jurisdiction. Finally, a number of risks that relate to *outsourcing*, which are captured as organisational issues in our category, are considered high risk. An example fault in this category relates to *unclear incident management processes* across the organisations involved in the service provisioning, which could result in protracted periods of unavailability of a system.

VIII. Conclusion

In this paper, we have shown how we model critical infrastructure services deployed on cloud infrastructure to assess the operational risk. This is based on an extensive fault and challenge catalogue. As an example, we conducted a risk analysis for a traffic management system to be moved to the cloud. We found that service availability is impacted most by that move and especially the use of virtualization is contributing dominantly to this. This is not surprising as it has also been the working assumption of the ETSI ISG NFV reliability and availability WG but we can put risk values to support this.

Having gained an understanding about the high risk scenarios and the associated categories allows for a better design of effective countermeasures. This is the obvious next step in our project and affects all levels of our architecture. Ideally, failures get detected and repaired locally but in complexed systems that is often not feasible. Thus, entities of the next higher level need to be prepared for the occurrence of these new failure modes to effective mitigate or at least contain the failure. Ultimately, the service consumers need to be educated for interacting with the migrated system and the ways the system might behave; e.g., short service downtimes might indicate the migration of virtual resources but not necessarily indicate a service failure. Hence, the development of security and dependability guidelines for the various operators and service users is key.

Acknowledgement

The research presented in this paper has been funded by the European Commission in the context of the Research Framework Program Seven (FP7) project SECCRIT (Grant Agreement No. 312758).

References

[1] B. Grobauer *et al.*, "Understanding cloud computing vulnerabilities," *IEEE Security Privacy*, vol. 9, no. 2, pp. 50–57, 2011.

[2] M. Theoharidou *et al.*, "In Cloud We Trust: Risk-Assessment-as-a-Service," in *7th IFIP WG 11.11 International Conference on Trust Management*, (Malaga, Spain), pp. 100–110, June 2013.

[3] J. Busby *et al.*, "Deliverable 3.1: Methodology for Risk Assessment and Management," December 2013. https://seccrit.eu/upload/D3-1-Methodology-for-Risk-Assessment-and-Management.pdf.

[4] Joint Task Force Transformation Initiative, "Security and Privacy Controls for Federal Information Systems and Organizations." http://dx.doi.org/10.6028/NIST.SP.800-53r4, April 2013. NIST Special Publication 800-53 Revision 4.

[5] A. Behnia *et al.*, "A Survey of Information Security Risk Analysis Methods," *Smart Computing Review*, vol. 2, February 2012.

[6] Software Engineering Institute Carnegie Mellon, "OCTAVE Information Security Risk Evaluation." http://www.cert.org/octave/, 2013.

[7] B. Karabacak and I. Sogukpinar, "ISRAM: information security risk analysis method," *Computers & Security*, vol. 24, no. 2, pp. 147–159, 2005.

[8] F. den Braber *et al.*, "The CORAS methodology: model-based risk assessment using UML and UP," *UML and the Unified Process*, pp. 332–357, 2003.

[9] ENISA, "Inventory of risk management/risk assessment methods and tools." http://www.enisa.europa.eu/activities/risk-management/current-risk/risk-management-inventory, 2013.

[10] F. L. Crespo *et al.*, *MAGERIT – version 2 Methodology for Information Systems Risk Analysis and Management Book I - The Method*. Ministerio de administraciones públicas, 2006.

[11] M. Schöller *et al.*, "An architectural model for deploying critical infrastructure services in the cloud," in *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on*, vol. 1, pp. 458–466, Dec 2013.

[12] M. Schöller *et al.*, "Resilient deployment of virtual network functions," in *Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT), 2013 5th International Congress on*, pp. 208–214, Sept 2013.

[13] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Tech. Rep. Special Publication 800-145, National Institute of Standards and Technology (NIST), September 2011.

[14] D. Molnar and S. E. Schechter, "Self Hosting vs. Cloud Hosting: Accounting for the Security Impact of Hosting in the Cloud.," in *Workshop on the Economics of Information Security (WEIS)*, (Cambridge, MA, USA), June 2010.

[15] M. Schöller and N. Khan, "ETSI ISG NFV: Resiliency Requirements," 2014.

[16] European Network and Information Security Agency (ENISA), "Cloud Computing: Benefits, Risks, and Recommendations for Information Security," November 2009. https://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment.