SPONSORED BY THE

Federal Ministry
of Education
and Research

ASMONIA

Attack analysis and **S**ecurity concepts
for **MO**bile **N**etwork infrastructures,
supported by collaborative **I**nformation exch**A**nge

Evaluation of the Intercloud Demonstrator

D3.4-1.0

Contributors: Cassidian / EADS Deutschland GmbH

ERNW Enno Rey Netzwerke GmbH

Fraunhofer Research Institution for Applied and Integrated Security (AISEC)

Hochschule Augsburg

Nokia Siemens Networks Management International GmbH

RWTH Aachen

Editor: Mirko Haustein (Cassidian / EADS Deutschland GmbH)

Author(s)	Company	E-mail
Mirko Haustein	Cassidian	mirko.haustein@cassidian.com
Herbert Sighart	Cassidian	herbert.sighart@cassidian.com
Matthias Luft	ERNW	mluft@ernw.de
Bernd Jäger	Nokia Siemens Networks	bernd.jaeger@nsn.com

About the ASMONIA project

Given their inherent complexity, protecting telecommunication networks from attacks requires the implementation of a multitude of technical and organizational controls. Furthermore, to be fully effective these measures call for the collaboration between different administrative domains such as network operators, manufacturers, service providers, government authorities, and users of the services.

ASMONIA is the acronym for the German name* of a research project that aims to improve the resilience, reliability and security of current and future mobile telecommunication networks. For this purpose the ASMONIA consortium made up of several partners from academia and industry performs a number of research tasks, based on the specific expertise of the individual partners. The project running from September 2011 till May 2013 receives funding from the German Federal Ministry of Education and Research (Bundesministerium für Bildung und Forschung, BMBF). Various associated partners further contribute on a voluntary basis.

* The full name is "**A**ngriffsanalyse und **S**chutzkonzepte für **M**Obilfunkbasierte **N**etzinfrastrukturen unterstützt durch kooperativen **I**nformations**A**ustausch" (Attack analysis and security concepts for mobile network infrastructures, supported by collaborative information exchange).

Partners:	EADS Deutschland GmbH / Cassidian
	ERNW Enno Rey Netzwerke GmbH
	Fraunhofer Research Institution for Applied and Integrated Security (AISEC)
	Hochschule Augsburg
	Nokia Siemens Networks Management International GmbH
	RWTH Aachen
Associated Partners:	Federal Agency for Digital Radio of Security Authorities and Organizations (BDBOS)
	Federal Office for Information Security (BSI)
	Deutsche Telecom AG (DTAG)

For more details about the project please visit www.asmonia.de.

Executive Summary

This paper describes the evaluation work of the developed ASMONIA Intercloud demonstrator based on a cloning scenario. The description cloning contains the cloning of a SIP server, which is hosted in a virtual machine, to other available cloud resources to mitigate an overload state resulting from an attack. For this it deals with a design security review where the design of the Intercloud demonstrator has been reviewed and assessed. Elastic systems/cloud computing environments are not a strictly defined term or technology and mostly industry-based origin leads to a lack of standardizations or clear terminology in the domain of Intercloud processes. To get a better understanding of the principles and necessities of cloud networking a design security review had been carried out. Chapter 2 describes the methodology of this review and the applied criteria and the tasks and gives a conclusion of the results. In the second part of this document the results from the Intercloud demonstrator have been reviewed and finally the scenario had been increased in terms of traffic by amount of data and verified using a simulation based validation. The necessary steps and assumptions for the implementation and the validation of an Intercloud demonstrator are described in chapter 3. This chapter includes a description of the results coming out from tests which had been carried out on the demonstrator. The results of this of chapter 3 were the base for the implementation of a simulation model of the demonstrator and in a further step the extension of the scenario in this simulation model. The goal of the simulation was the validation of the principles of Intercloud networking. Chapter 4 gives a description of the simulation model and a discussion of the results coming out from the simulation based on different use cases. A final conclusion of the evaluation work for the developed Intercloud demonstrator is put down in chapter 5 and completes this deliverable.

Table of Contents

1 Introduction

The focus is to evaluate the behavior of the developed Intercloud demonstrator based on a cloning scenario as described in our previous work in [ASMONIA_D3.3] under different overload conditions and to estimate the necessary security criteria. In addition the Intercloud demonstrator shall be validated by using simulations. The availability of network simulation tools will enable the validation of defined scenarios using simulations. The aim is to investigate the effectiveness of cloud systems in supporting servers in overload situations. It is planned to carry out this task at first on a hardware based Intercloud demonstrator. All the necessary tasks for a live cloning process will be applied on this and the test results will be taken for further use in the definition of a simulation based test bed. A scaling process will be carried out to prepare the simulation model on its readiness for further investigations regarding cloud networking with focus on the live cloning process. Specific live cloning scenarios will be described and their behavior tested in the simulation. It is expected that comparable results for a later analysis regarding different behavior in a scenario during applied process modifications will be derived.

2 Design Security Review

In order to ensure that the developed approach to integrate elastic systems into the overall architecture both fulfills defined goals and requirements and provides security benefit, a document-based security review is to be performed. This review uses ASMONIA deliverables and information from live presentations for the analysis against international standards, internal project goals, and best practices.

2.1 Review Methodology

Elastic systems/Cloud computing environments are not a strictly defined term or technology. Its heavily industry-based origin leads to a lack of standardizations or clear terminology. Hence the design of Cloud environments does not follow strict guidelines and must be developed and reviewed individually taking the characteristics, requirements, and objectives of the particular target environment (of evaluation) into account. The following subsections develop review criteria which are assessed and rated using the following possible results:

Result	*Description*
✓ Fulfilled	The design/environment fulfills all requirements and objectives.
~ Partially Fulfilled	The design/environment fulfills most requirements and objectives. No critical requirement violations.
× Not Fulfilled	The design/environment does not fulfill the requirements and objectives. Critical requirement violations.

Table 2-1: Possible Review Results

2.2 Review Criteria

The review criteria can be categorized into functional/architectural aspects and security aspects. The following subsections describe the results of the review and the description of the different review criteria and its results.

2.2.1 Architecture & Functionality

The design of the ASMONIA Intercloud is described in [ASMONIA_D3.2]. Different internal requirements are described in [ASMONIA_D3.1], [ASMONIA_D3.2], and [ASMONIA_D3.3]. As the design of the Cloud services is tightly integrated with the requirement engineering, the analysis as for the implementation of the internal requirements is *fulfilled*.

In order to ensure a holistic review to support both common, internal, and intrinsic (security) goals, additional requirements from international standards and best practices are taken into account. Even though Cloud environments are (due to the development originating from the industry) lacking strict specifications and standards, [NIST_SP800-145] provides a de-facto standard, which is used by most Cloud practitioners, defining and setting the basic parameters of Cloud environments. The design described in [ASMONIA_D3.2] is reviewed as for the compliance with the five basic Cloud characteristics (which can also be mapped to functionality and use cases) described in [NIST_SP800-145]. The results, where applicable, are described in the following subsections and summarized in the following table:

Requirement	*Result*
Broad Network Access	Not applicable.
Rapid Elasticity	✓ Fulfilled
Measured Service	✓ Fulfilled
On-Demand Self-Service	✓ Fulfilled
Resource Pooling	✓ Fulfilled

Table 2-2: Functional Requirements Results

2.2.1.1 Broad Network Access

One of the use cases of the ASMONIA Intercloud is the handling of overload situations: The unavailability/outage/overload of the network connection is one of the key overload situations. Hence, the requirement for broad network access is not fully applicable and is hence not evaluated. However, the general network availability is given by design and assured in normal situations. The handling of network outages increases the relevance of the requirement *Rapid Elasticity.*

2.2.1.2 Rapid Elasticity

The possibility to rapidly deploy and use new instances is one of the key aspects to handle overload situations. The ASMONIA Intercloud is compliant with the requirement *Rapid Elasticity* as it defines all functionality required (e.g. [ASMONIA_D3.2], 5.1.1.6 and 5.1.1.1) and addresses shortcomings of the design (e.g. [ASMONIA_D3.2], 5.1.1.3) using appropriate controls ([ASMONIA_D3.2], 5.1.1.4.1).

2.2.1.3 Measured Service

The Cloud requirement to measure service is appropriately addressed in order to support future billing/accounting requirements/functionality (e.g. [ASMONIA_D3.2], 5.1.1.5).

2.2.1.4 On-Demand Self-Service

As the ASMONIA Intercloud is a “Cloud of Clouds”, this requirement is fulfilled on a per-participant base per design. For the Intercloud design, the requirement is fulfilled as protocols for the service management between the different clouds are defined and support all use cases (e.g. [ASMONIA_D3.2], 4.4.1, 4.4.2, and 4.4.3)

2.2.1.5 Resource Pooling

As the ASMONIA Intercloud is a “Cloud of Clouds”, this requirement is fulfilled on a per-participant base per design. For the Intercloud design, the requirement is fulfilled as the architecture is purely designed to handle overload situations and hence pool resources between the different participants of the Intercloud.

2.2.2 Security

In addition to the functional review, a security review was performed in order to ensure the secure operation of the Intercloud regarding the highly distributed design and the interfaces between participants. The evaluation criteria are developed and consolidated based on [NIST_SP800-144], [CSA_TopThreats], [ENISA_CloudRisks], [ASMONIA_D3.1], and

[ASMONIA_D3.2]. The results, where applicable, are described in the following subsections and summarized in the following table:

Requirement	*Result*
Governance	✓ Fulfilled
Trust	~ Partially Fulfilled
Infrastructure & Design	✓ Fulfilled
Identity and Access Management	✓ Fulfilled
Isolation	~ Partially Fulfilled
Data Protection	✓ Fulfilled
Availability & Resource Management	✓ Fulfilled
Incident Response	Not applicable.
Security of Management Interfaces & APIs	~ Partially Fulfilled
Lock In	Not applicable.
Subpoena and E-Discovery	Not applicable.
Changes in Jurisdiction	Not applicable.
Compliance	✓ Fulfilled

2.2.2.1 Governance

A potential lack of governance as a) users may circumvent corporate governance mechanisms and b) public Cloud environments result in a lack of certain control mechanisms is a central concern for the use of Cloud services.

As “The cloud control layer is solely managed by the cloud service provider and provides functions like pooling of cloud resources, deployment, monitoring of these resources […]” ([D3.2], p17) the Intercloud allows the implementation of governance mechanisms and processes for the internal Cloud infrastructure as well as for the Intercloud functionality, as the interfaces and services are also hosted by the particular participants (refer to e.g. [D3.2], 4.4.1, 4.4.2, and 4.4.3).

2.2.2.2 Trust

Trust between Cloud provider and clients can be established based on different factors, such as transparent documentation of all characteristics and processes of the Cloud environment, consistent handling of (e.g. service or infrastructure) changes, and a symmetric trust relationship with the customer. However as the Intercloud approach results in different trust requirements as a provider-client relationship, the requirement is only partially fulfilled as it is explained in [D3.2], p46:

> *“It remains future work to investigate appropriate approaches for modeling trust and reputation in the ASMONIA collaborative cloud. But these approaches can help heightening the chances for acceptance on the side of the users, i.e. the mobile network providers.”*

Until the future work to design all aspects of Intercloud participant relationships and capabilities is finished, the design is only partially compliant with the requirement.

2.2.2.3 Infrastructure & Design

In order to support an ongoing secure operation, the design must be developed having different security best practices, such as clear interface definition and a reduction of attack surface in mind as well as the availability of communication channels. The overall design is suitable to fulfill these requirements. Certain operational aspects are out of scope for the design and mentioned in Sections 2.2.1.5 and 2.3.

2.2.2.4 Identity and Access Management

The Intercloud approach allows the transfer of the responsibility for the identity and access management to each participant. For the Intercloud-wide IAM, the design includes support for the federation of identities ([ASMONIA_D3.2], 4.4.2.4)

2.2.2.5 Isolation

Cloud environments are often (by design) multi-tenancy environments, serving different parties on the same infrastructure. As the ASMONIA Intercloud is designed to allow the offloading of services to other cloud participants in order to handle overload situations, the isolation of different tenants is a key requirement as well. The isolation must be ensured on

- Memory & CPU
- Network
- Storage
- Management Interface & Functionality

level. Every single entity carrying out tasks related to the mentioned levels of isolation (e.g. hypervisors, network devices, storage nodes, and management applications) must support isolation. Taking recent Cloud vulnerabilities into account as described in [ACM], [Insinuator] and [VMWareSecAdv], the operation and implementation of all components is of particular relevance. Different shortcomings of the design such as

- [ASMONIA_D3.2], 4.4.2.3.1.2.3: “The security of the host in a virtual machine environment needs no special configurations when the setup of the infrastructure is done in a secure way.”: Recent vulnerabilities require different controls such as patch- and vulnerability management or input validation of all input sources and emphasize the need for classical hardening mechanisms (such as the removal of unnecessary virtual devices).
- Requirements for operational security documentation: Even though this is not the primary focus of a design document, it is lacking a reference to the requirement to create operational security documentation (e.g. covering vulnerability management and hardening measures) which must be present and implemented in order to ensure the secure ongoing operation of the Intercloud environment.

lead to partial compliance with the requirement, as they are flagged as work in progress or are not in primary scope of the design document.

2.2.2.6 Data Protection

Legal data protection issues can be addressed following the Intercloud approach and flexible onboarding processes between the participants. Additional, custom agreements can be used to address legal requirements resulting from e.g. the German Bundesdatenschutzgesetz.

2.2.2.7 Availability

The requirement for availability is a core design goal of the Intercloud environment. While the different participants are responsible for the availability of their own systems, the Intercloud design provides functionality for an available overlay Cloud (e.g. by exposing functionality as VM Monitoring [D3.2], 5.1.1.5, Load Balancing [D3.2], 5.1.1.7, and Auto Scaling [D3.2], 5.1.1.6).

2.2.2.8 Security of Management Interfaces & APIs

While the different Intercloud participants are responsible for the security of the management interfaces of their particular Cloud environments, the Intercloud management interfaces need to be designed in a secure way. As the specific secure implementation is out of scope of the design document, the design mentions mechanisms for the authentication and authorizations concepts ([D3.2], 6.5).

2.2.2.9 Compliance

Compliance issues can be addressed following the Intercloud approach and flexible onboarding processes between the participants. Additional, custom agreements about controls to be implemented and security levels to be fulfilled can be met between the participants in order to fulfill different internal and external compliance requirements.

2.3 Conclusion

Except for the statement of the hardening of the virtualization host systems resulting in a minor non-compliance, all non-compliances result from aspects not primarily in scope of the design document. The issues are mentioned in order to ensure that the connected security aspects are regarded and appropriately addressed during the actual implementation of the Intercloud environment.

3 Validation of the Intercloud demonstrator

This section discusses the results of the ASMONIA Intercloud demonstrator as presented in [ASMONIA_D3.3] and build the base for the following investigations using simulation. The key parameters to evaluate the results of the ASMONIA demonstrator are the size of the live-cloned OpenSIPS VM and the bandwidth provided between the two Collaborative Clouds.

- **Size of OpenSIPS VM: ~ 2 Gbyte**
- **Bandwidth between the Collaborative Clouds: ~ 1 Gbps**

Depending on these parameters the behavior and the benefit for larger VM sizes or a larger amount of VMs or both can be roughly estimated.

The Intercloud live cloning can be separated into multiple phases as shown beneath:

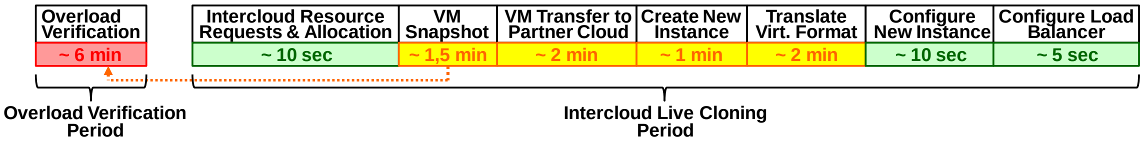

Figure 1: Time period for Intercloud Live Cloning

3.1 Intercloud Demonstrator Phases

Below each phase name you find roughly measured time intervals gathered during implementation and testing of the ASMONIA Intercloud demonstrator. Please be aware that the length of an Intercloud live cloning phase in the figure is not corresponding to the size of the rectangle, instead the size of the time interval is indicated by the bottom time highlighted in traffic light colors: from green like 'negligible' up to red like 'overwhelming'.

The main security objective of the ASMONIA Intercloud demonstrator is to maintain or increase the availability of cloud-based telecommunication infrastructures in case of a Denial of Service (DoS) attack. Relevant for the usefulness of this solution for co-operating telco operators is the period of time until the additional resources are available to mitigate the attack. The Intercloud Live Cloning period as measured with the ASMONIA Intercloud demonstrator is around 7 minutes.

Subsequently the specific phases are discussed in more detail concerning their relevance and potential improvements.

3.1.1 Overload Verification

The Overload Verification Period is not considered for the length of the Intercloud Live Cloning Period. But as the Overload Verification triggers the Intercloud Live Cloning the length of this period is nevertheless important. Therefore the actual implementation and possible improvements shall be discussed.

The Overload Verification phase consists of two sub-phases:

1. Detection of overload state by [Nagios] as outlined in [ASMONIA_D3.3] (~ 1min)
2. Monitoring that the overload state persists for at least 5 minutes → trigger to start the Intercloud live cloning

Overload Detection as implemented per default in the ASMONIA Intercloud demonstrator is by far the largest time interval. While the sub-phase 'Detection of overload state by Nagios' seems not significantly reducible, it is worth to think about possible improvements to deduce the Intercloud Live Cloning trigger.

For the demonstrator a rather simple trigger mechanism with a 5 minutes hysteresis was selected because it was not in the scope of the work to develop a more sophisticated solution. Furthermore the size of the hysteresis is configurable in the Collaboration Cloud dashboard and can therefore be reduced by the user if desired.

With behavior models analyzing also the kind of traffic increase, perhaps combined with time thresholds it seems probable that the trigger time can be significantly reduced. Another possible input for Overload Verification in the context of ASMONIA would be mobile network status indications from the 'Monitoring and Analysis' functionality of WP4. This functionality that permanently monitors the mobile network status in a highly sophisticated manner via distributed sensors could perhaps be able to already detect a malicious network behavior before the overload at the potentially affected telco equipment really occurs.

Without concrete evidence it is assumed that the Overload Verification period could be approximately reduced to a time of ~ 2 minutes.

3.1.2 Intercloud Resource Requests and Allocation

This phase can be decomposed into three sub-phases:

1. Request of additional resources in the Intercloud, initiated by the operator cloud
2. Activation of Intercloud autoscaling functionality to detect appropriate free resources in the Intercloud
3. Activation of Intercloud matchmaking to allocate specific appropriate resources with regard to the Intercloud policies

Derived from the 'basic' implemented ASMONIA demonstrator Intercloud functionality it can be estimated that this phase is with ~ 10 seconds negligible.

3.1.3 VM Snapshot

The snapshot of a 2 Gbyte VM takes around 1.5 minutes. The overall Intercloud Live Cloning period can be reduced if taking of the snapshot starts already in advance during the Overload verification period.

3.1.4 VM Transfer to Partner Cloud

This phase strictly relates to the size of the VM image (2 Gbyte) and the available bandwidth between the Collaborative Clouds (1 Gbps gross) with respect to the parallel background traffic. This phase was averaged with ~ 2 minutes under different background traffic conditions in the Fraunhofer Cloud. Unfortunately the size of the background traffic and therefore the remaining bandwidth for the VM image transfer is not exactly known.

Therefore we try to derive a reasonable transfer time from a theoretical analysis of the network between the Collaborative Clouds, the ACEX network. This network is already during normal conditions (no Intercloud Live Cloning) filled with a certain amount of background traffic, caused by the Intercloud services traffic. During a DOS attack it must be expected that the background traffic further increases due to attack-related Intercloud services activity. Nevertheless it is assumed that the control traffic only takes the minor part of the bandwidth and that at least 50% to 70% of effective bandwidth (assumed 80% of netto bandwidth) is available in case of a 1 Gbps link. With 500 Mbps effective bandwidth (worst case) the transfer of a 2 Gbyte image would take around 30 seconds reducing the 'VM transfer to Partner cloud' to around 0.5 minutes and therefore the total Intercloud Live Cloning period to around 5.5 minutes.

Although a transfer time of 30 seconds appears relatively short, this changes immediately if larger images (e.g. 20 Gbyte or even 200 Gbyte) or multiple images simultaneously have to be transferred. Then this phase can be significantly reduced by bandwidth over-provisioning, e.g. by 10 Gbps instead of 1 Gbps leading to a transfer time of (with 8 Gbps = 80% of effective bandwidth assumed)

- ~ 2 seconds in case of a 2 Gbyte image (instead of ~ 30 seconds for 1Gbps)
- ~ 20 seconds in case of 20 Gbyte image (instead of ~ 5 minutes for 1 Gbps)
- ~ 200 seconds in case of 200 Gbyte image (instead of ~ 50 minutes for 1 Gbps)

While being advantageous for performance, there is in case of a 10 Gbps link presumably a conflict between economics and performance as (potentially long-haul) high-speed optical links between the Collaborative Clouds are expensive and will usually only be under-utilized, supposed that DoS attacks are more an exception than normality.

It can also be seen that the Intercloud control traffic can be absolutely neglected in case of a 10 Gbps link.

Another aspect that affects the available bandwidth between the Collaborative Clouds significantly more than the Intercloud control traffic is the communication traffic between the cloned network element (the OpenSIPS SIP server in case of the ASMONIA Intercloud demonstrator) and its connected neighboring elements (these are the SIP load generators and the OpenSIPS database in case of the ASMONIA Intercloud demonstrator) remaining in the original operator cloud. This traffic must be additionally bi-directionally transmitted over the ACEX network (connecting the Collaborative Clouds) when at least one cloned image gets operational in a foreign Collaborative Cloud.

Although this traffic may be not very high in case of the ASMONIA Intercloud demonstrator, it will get relevant in case of real-life applications. The equivalent of the OpenSIPS server in a mobile core network is for example the CSCF (Call Session Control Function, a SIP server in the IMS network) connected to the HSS (Home Subscriber Server, an IMS subscriber database). The size of a CSCF-VM can be around 30 Gbyte, the subscriber SIP traffic to/from the CSCF is assumed around 25 Mbps per direction and the control traffic between the CSCF and the HSS is assumed around 10 Mbps per direction. In case of a DoS attack around 500% of the original resources must be additionally provided in foreign Collaborative Clouds (rough assumption of the associated ASMONIA member Telekom) to mitigate the attack. That leads to five cloned CSCF images with a total of ~ 175 Mbps (effective bandwidth) additional load in the ACEX network. Already this example shows that 1 Gbps links in the ACEX network will soon get overloaded if more than one network element with a significant VM size and with significant traffic relations to its neighbors is attacked.

Not regarded in the context of the ASMONIA demonstrator is the latency that is additionally inserted by the Collaborative Cloud and the interconnecting ACEX network. Although these latency aspects were neither measured nor estimated, it is expected that they play (besides the bandwidth issue) an important role. Especially in case of a SIP server connected to a user database with its ping-pong-like traffic behavior additional latencies may sum up and lead to an increased connection setup time.

3.1.5 Create New Instance

This time is directly corresponding to the size of the VM image and can be presumably hardly influenced by ASMONIA measures (only by an increases cloud performance which is not in the scope of ASMONIA).

3.1.6 Translate Virtualization Format

Not necessarily but in the most general case the two Collaborative clouds may use different virtualization formats. This doesn't apply for the ASMONIA demonstrator. As a consequence no measured time intervals are available from the ASMONIA demonstrator.

As a good estimation,

The results achieved by [Fallenbeck2011] for the translation of a 2 Gbyte image between VDI/VMDK-format and the Xen-RAW-format were used for a good estimation. The translation time is

- ~ 3 minutes for VDI → RAW
- ~ 1 minute for RAW → VDI

As the kind of translation is not predictable, an average of 2 minutes is assumed. The translation between other virtualization formats is estimated to be comparable.

As a consequence also the Intercloud live cloning demonstrator underlines the necessity of a standard open virtualization format as proposed with OVF superseding the necessity of virtualization format translation.

3.1.7 Configure New Instance & Configure Load Balancer

Configuration of new instances and load balancer (e.g. IP addresses, especially if the loadbalancer is dynamically inserted) is inevitable and with a timer interval of 10 seconds negligible.

3.2 Summary

Summarized it can be stated that the Intercloud Live Cloning of a 2Gbyte VM over real 1 Gbps links as provided by the existing ASMONIA demonstrator takes around 5.5 minutes (with the adjusted VM transfer time and an assumed translation between two different virtualization formats). If considering also the Overload Verification period (default ~ 6 minutes) an overall time interval of ~ 11.5 minutes is needed from start of the attack until the cloned image gets operational.

With further improvements as discussed before (also potentially uneconomical improvements like bandwidth over-provisioning) a reduction of the overall time interval down to around 4.5 to 5 minutes is imaginable without assuming improvements of the existing cloud technology.

How does this scale in case of different (larger) VM image sizes or in case of multiple simultaneously VM Intercloud Live Cloning? As a rough estimation the summarized size of the VM image sizes can be calculated and the effective available bandwidth to the Collaborative Clouds has to be evaluated. Looking at Figure 1 it can be derived

- that the overload verification time doesn't depend on the VM image size and the available bandwidth. Therefore this phase is assumed to stay fairly constant. (But be aware that VM snapshotting parallel to Overload Verification period gets increasingly ineffective if the VM image sizes are significantly larger than 2 Gbyte).
- that the time for the phases 'VM Snapshot', 'VM Transfer to Partner Cloud', 'Create New Instance' and 'Translate Virtualization Formats' behave in the first approximation rather linearly depending on the VM image size and in case of 'VM transfer to Partner Cloud' also depending on the effective available bandwidth.
- that the time for the phases 'Intercloud Resource Requests & Allocation', 'Configure new Instance' and 'Configure Load Balancers(s)' will increase but presumably less

than linearly. Due to the minimal time intervals of these phases their behavior is in the first step negligible.

One important aspect in case of multiple VM images Intercloud Live Cloning is also to consider possible parallelization of tasks. This is for example obvious for the phases ‘Create New Instance’ and ‘Translate Virtualization Formats’ which can be done in parallel, supposed that the processing power of the cloud is performing enough. However, in case of ‘VM Image Transfer to Multiple Partner Clouds’ parallelization may be only possible with a corresponding network architecture. In case of a meshed network or an optical ring potentially a parallel VM image transmission to multiple partner clouds is possible while in case of multiple Collaborative Clouds connected to a central routing instance (star-like network architecture) only a subsequent transmission of VM images is possible.

Final question: Is Intercloud live cloning beneficial for cloud-based cooperating telco operators to mitigate DoS attacks? The answer is ‘Maybe’! In case of long-lasting DoS attacks (over several hours) it is certainly an improvement that after some 10 to 30 minutes the consequences of such attacks may be (significantly?) diminished. On the other hand this method is at the time being no magic cure to prevent users of a specific service to be affected by DoS attacks at all.

Although not in the scope of the ASMONIA research project, the Intercloud Live Cloning could nevertheless be interesting for telco cloud operators to prevent normal overload situations caused for example by time-of-day dependent traffic or by specific events like football matches or concerts and not by DoS attacks. These are by far more predictable and here the disadvantages of the time period for Intercloud live cloning can be largely reduced by planning while maintaining the economic advantages of resource sharing.

4 Simulation

4.1 Use Case Specification

The ASMONIA Intercloud demonstrator as described in [ASMONIA_D3.3] is evaluated in chapter 3 and shows the Live Cloning of a telco VM instance (an OpenSIPS SIP server) from an Operator Cloud to the Collaborative Cloud of another operator. The results of the demonstrator shall be validated by means of simulation and shall furthermore be put on a more general basis by simulating also more complicated scenarios with variations of parameters like bandwidth of the ACEX network or the VM image size.

As an input for the simulation not the OpenSIPS SIP server but a more realistic example from the ASMONIA mobile network is assumed which is shown in Figure 2.

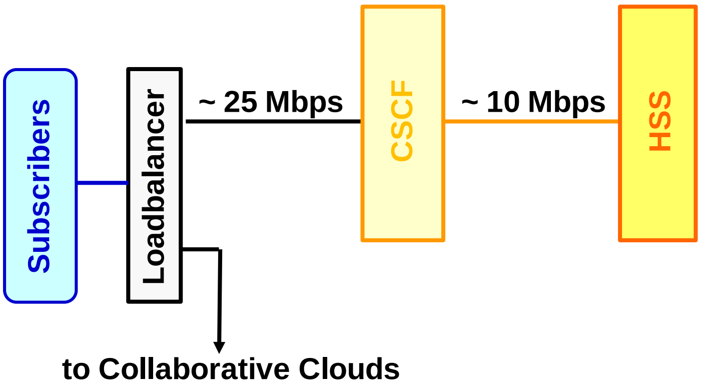

Figure 2: SIP Use Case Scenario

In the example the CSCF (Call Session Control Function, the SIP server of the IMS network) is used instead of the OpenSIPS server and the HSS (Home Subscriber Server, the subscriber database) is used instead of the OpenSIPS database. A real CSCF has a VM image size of around 30 Gbyte and the subscriber traffic to/from the CSCF is ~ 25 Mbps (25 Mbps per direction) while the traffic between the CSCF and the HSS is ~ 10 Mbps (10 Mbps per direction).

Compared to the ASMONIA demonstrator the simulation could be enhanced in the following ways:

- Intercloud Live Cloning to multiple Collaborative Clouds (proposed 3) instead of only one Collaborative Cloud
- Creation of multiple VM clones in the Collaborative Cloud(s) in parallel instead of only one clone

  It is proposed to simulate 5 copies of one specific VM and 3 copies of different VMs. These clones can be distributed in several ways:

  - in one specific foreign Collaborative Cloud
  - in all other Collaborative Clouds
  - in the operator-internal Collaborative Cloud part
- Consideration of different VM sizes ranging from 2 Gbyte to 200 Gbyte

- Consideration of different bandwidth for the ACEX network connecting the Collaborative Clouds ranging up to 10 Gbps
- Consideration of split functionalities after Intercloud VM cloning

  Split functionalities means that in our example the CSCF-VM is transferred to a foreign Collaborative Cloud while the subscribers and the HSS database still reside in the original Operator Cloud. That has the consequence that the subscriber traffic of ~ 25 Mbps per direction and the HSS database traffic of ~ 10 Mbps per direction must be transferred over the ACEX network interconnecting the Collaborative Clouds.
- Consideration of different network architectures concerning parallel transmission (e.g. meshed, star-like, ring networks)
- Consideration of translation between different virtualization formats

  This comprises also the evaluation whether it is advantageous to perform the translation already in the original Operator Cloud or only in the foreign Collaborative Cloud

4.2 Scenario Description

Live cloning is a sub-process in Intercloud operation. The following described principle has been partly implemented in a simulation model:

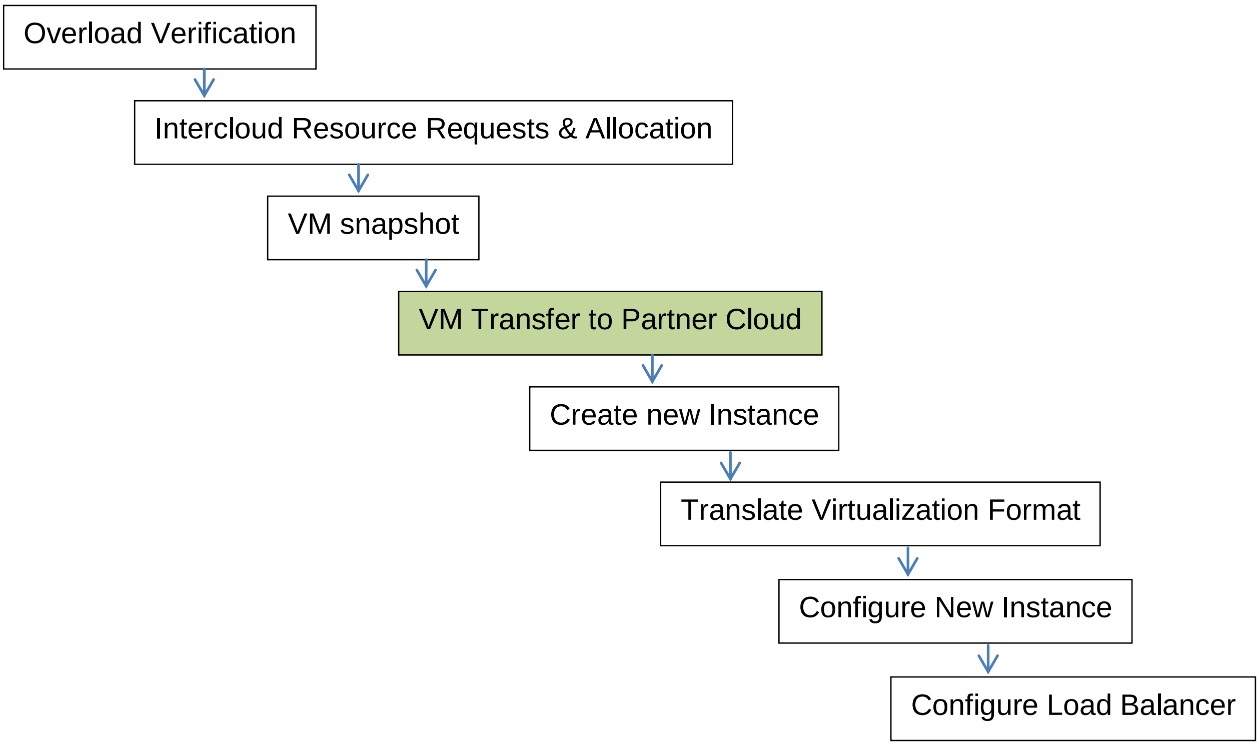
Overload Verification
Intercloud Resource Requests & Allocation
VM snapshot
VM Transfer to Partner Cloud
Create new Instance
Translate Virtualization Format
Configure New Instance
Configure Load Balancer

Figure 3: Intercloud Live Cloning

Goal of this process was to detect overloaded elements (server) in an own network and to establish a redundant server in a 'cloud' on available capacities provided by a cooperating network partner (see Figure 3). Overload verification and the following both steps to create a snapshot (image) of the overloaded server are possible in real test beds. The interesting task is the transfer of this image to the partner cloud during overload conditions in the network structure. The question was how much time this process would take under such conditions.

It was important to define a simplified simulation model of the network with a realistic data flow behavior. [OPNET Modeler] as the modeling and simulation platform provides tools to define the saturation of network links (link utilization) and by this to define additional 'harmful' traffic load, to observe the data flow and to record the behavior of the network elements for analysis purposes.

The scenario should fulfill the following conditions:

- The results of the implementation must be like those of the real experiment
- The scenario shall be expandable up to five cloud elements
- The considered data volume must be equivalent to realistic ones.

The architecture of the simulation model is represented in Figure 4. This model includes the basic components of a mobile network necessary for the considered cloning process. It is assumed that five network providers cooperate in this scenario (networks A, B and C). The basic network represented for provider A is the same as for B to E but not pictured in this scheme. All VoIP management relating data are provided and handled by Database_OperatorCloud_A1. It was assumed that this base had to be accessed during the whole simulation process.

During normal operation the voice communication in 'Operator Legacy Network_A' is controlled via 'Server_OperatorCloud_A1' using 'Database_OperatorCloud_A1'. In case of an overload situation on 'Server_OperatorCloud_A1' the traffic is redirected by ACGW_A (load balancing function) to the Collaborative Cloud represented by the load balancing units 'LB_A' to 'LB_C' and the attached servers (Server_CollabCloud_Ax) as well as the Intercloud connection network routers (ACEX_A/B/C/D/E). The model considers primarily the traffic flow to and within the collaborative cloud taking into account varying data rates depending on the tested scenarios. As basic case the same scenario as applied in the cloud demonstrator was tested in the simulation model for a verification of the adjustments regarding link bandwidth and transferred data volume. In a number of experimental runs different use cases had been tested and analyzed. A detailed representation and discussion of the results is given in chapter 9.

The simulation model includes a number of VoIP participants connected to different provider networks. Each provider owns a SIP server with a limited call handling capacity. A call can only be established if the SIP-server has the capability to handle this call otherwise the call is rejected. To avoid this, a kind of 'load balancing' had to get active and to redirect the call to a backup server for establishing the call from this. This procedure was implemented as described in the following. All known and unswayable processes and delays had been defined by fixed parameters (e.g. overload verification, Intercloud resource request & allocation, VM snapshot) and set in a defined relation to the traffic scenario regarding to the estimated parameters from the real test bed. Depending on the network traffic load (link utilization) the duration for the VM transfer was estimated and extended by the necessary processing time to get a cloned backup server.

4.3 Model Description

Figure 4 shows the model architecture for the simulation model including 5 different provider networks connected in the collaborative cloud.

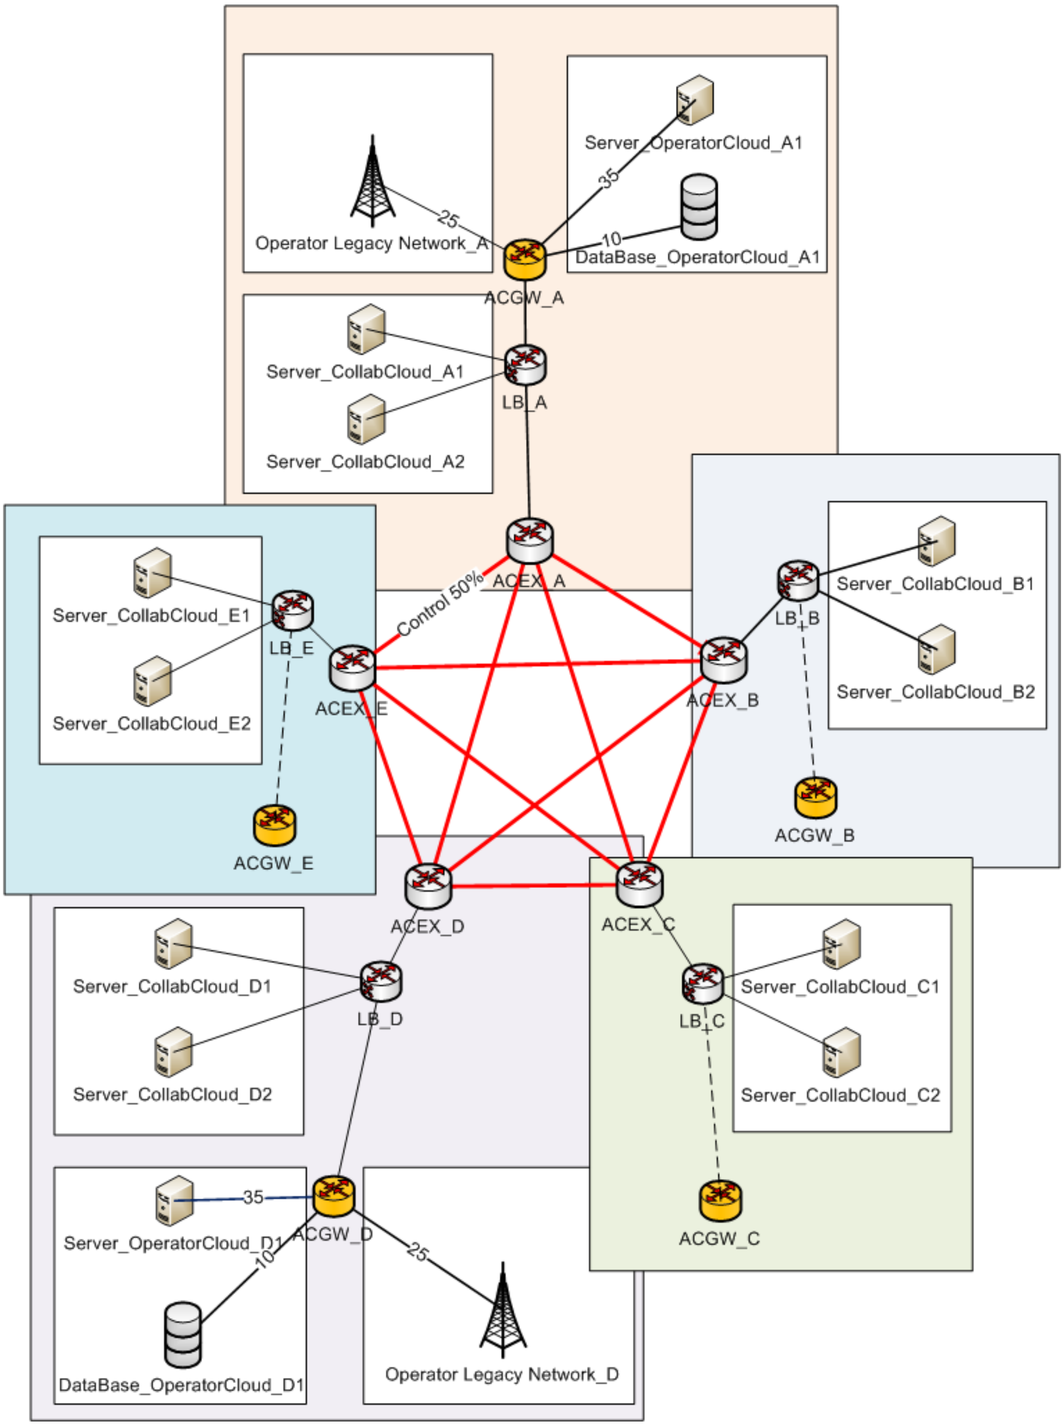

Figure 4: Model architecture

4.3.1 Traffic model

An increasing number of server accesses caused a saturation of this affected server and lack of capability to handle new incoming requests. In order to ensure the continuity of services a load balancing mechanism got active for the redirection of requests during peak demands to stand-by systems located in the Collaborative Cloud. The simulation considered the background traffic in a network (by link saturation) and the traffic necessary to operate and administrate specific services as well as the upcoming traffic load during server image transfers. In the applied traffic model quality of service (QoS) was not considered.

4.3.2 Network Parameters

4.3.2.1 Latency vs. Bandwidth

Although the theoretical peak bandwidth of a network connection is fixed according to the used technology, the actual bandwidth to be obtained varies over time and is affected by high latencies. Excessive latency creates bottlenecks that prevent data from filling the network pipe, thus decreasing effective bandwidth. Latency arises in a router from data processing delays caused e.g. by packet queuing.

4.3.2.2 Packet Size Definition

Network utilization and latency are usually inversely proportional. Smaller packets will be transmitted over the network faster and therefore will have lower latency. However, many smaller packets require greater network overhead (IP headers and Ethernet headers) than fewer larger packets. To achieve the maximum throughput the applied packet size in the simulation model was auto computed based on exchanged packets/sec and bits/sec.

4.4 Simulation Process

4.4.1 Initial Situation

An increasing number of call requests in Operator Legacy Network_A leads to an overload of Server_OperatorCloud_A1 and no further calls could be established. It was assumed that Server_OperatorCloud_A1 is running in a VM (virtual machine). This was considered in the simulated traffic model as exceeding number of SIP-requests represented by a specific data transfer rate of 25 Mbps between Operator Legacy Network_A and Server_OperatorCloud_A1. Furthermore a continuous 10 Mbps data traffic was applied between Database_OperatorCloud_A1 and Server_OperatorCloud_A1 (see Figure 2).

4.4.2 Cloud Preparation

The scenario was that after the detection of the overload situation on Server_OperatorCloud_A1 a snapshot of the VM had been taken. This was considered as delay in starting the image data transfer. It was assumed but not modeled that during this overload situation incoming calls could not be handled and were rejected.

4.5 Model Validation

Our first step was the adjustment of the simulation model to achieve comparable results to the real test bed. The applied link utilization was scaled for an adaption of the transfer time to that estimated in the real test bed. The transmission delay in the present simulation model depended on the link utilization; increasing link utilization caused an increasing transmission delay. This dependency is represented in Figure 5. The measured time of approximately 120 seconds for the transfer of a 2 GB image (as measured for this transfer in the Intercloud demonstrator) was achieved on a link utilization of 85%.

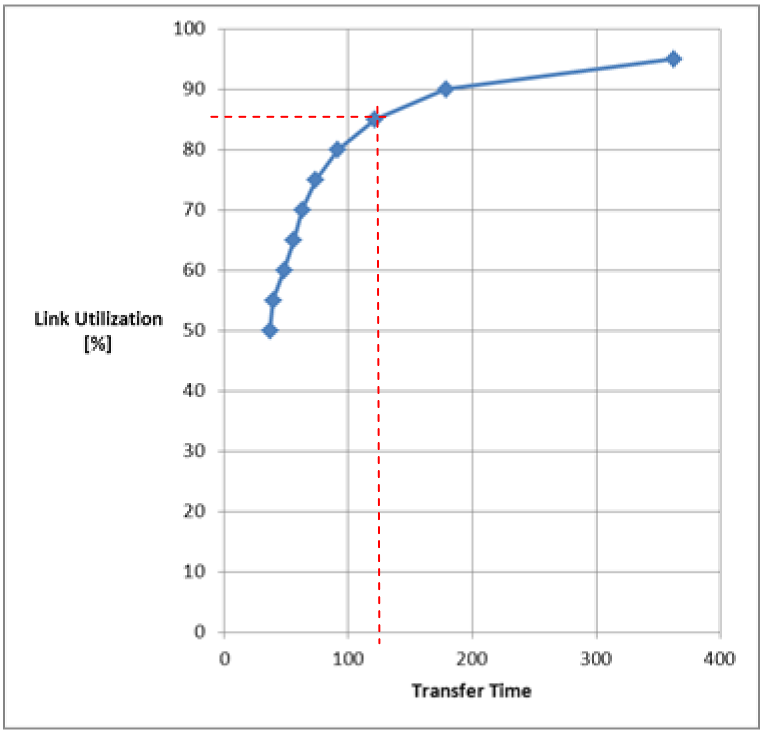

Figure 5: Transfer time vs. Link utilization

4.5.1 VM-Image Transfer

Purpose of live cloning is the punctual supply of an extended infrastructure in the case of a capacity bottleneck. Task of the simulation was the estimation of the backup data (VM snapshot) transfer behavior between several servers respectively from one to several servers during operation. A separate TCP connection for each image transfer was established in the simulation.

The simulation model included five service providers each contributing server capacity to the Collaborative Cloud. It was assumed that each provider procured two backup servers. Figure 7 represents the principle: the available link capacity between the routers ACEX_A, ACEX_B, ACEX_C, ACEX_D and ACEX_E was reduced by applying link saturation to achieve the desired transfer duration as had been estimated in the real test bed (~120 sec) for the image transfer (image size ~ 2 Gbyte).

The traffic profile for the sequential distribution of a 2 GB VM-image is shown in Figure 6. The image data transfer was started after a defined time period of 400 seconds as explained in chapter 3.

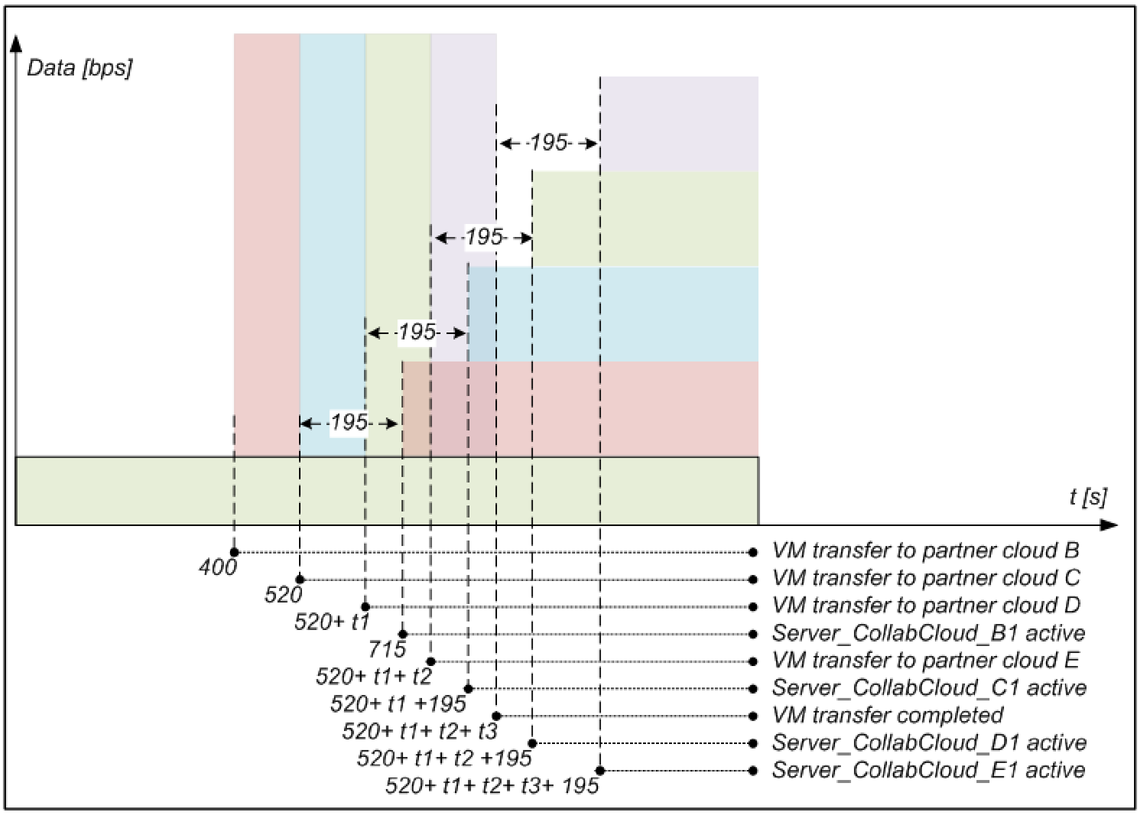

Figure 6: Traffic timing for a 2GB VM-image transfer

The cloud resource was in operation after a preparation time of 195 seconds (Server_CollabCloud_B1 active). This time results from the following summarized Intercloud demonstrator phases (see Figure 1):

- Create New Instance (60 sec)
- Translate Virtualization Format (120 sec)
- Configure New Instance (10 sec)
- Configure Load Balancer (5 sec)

The necessary transfer time (t1, t2 and t3) for each VM-image was determined by the simulation taking into account effects like bandwidth limitations. This caused an increasing delay for the ongoing transfer process to the remaining Collaborative Cloud servers. The reason was additional dataflow between this new SIP server and the load balancer ACGW_A as the initiator of the re-routing process. As shown in Figure 6 the basic traffic load from the initial phase increased by the number of connected servers.

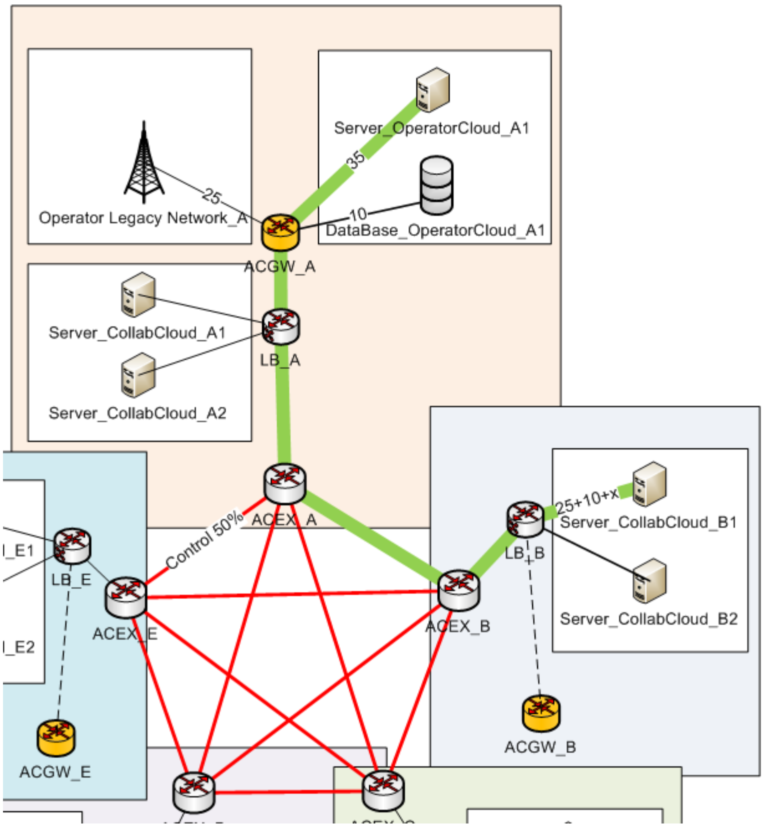

Figure 7: Data transfer delay adjustment

This model enabled the simulation of image-data transfers from operator network A as the origin of the images to all servers placed at the disposal. Two use cases had been considered: first a sequential data-image transfer from the source to multiple destinations and second the parallel transfer to multiple destinations.

The following selection from the recommended use cases from chapter 4.1 was implemented in our simulation model. Following requirements had been selected:

- Intercloud Live Cloning to multiple Collaborative Clouds instead of only one Collaborative Cloud
- Creation of multiple VM clones in the Collaborative Cloud(s) in parallel instead of only one clone
  - in one specific foreign Collaborative Cloud
  - in all other Collaborative Clouds
  - in the operator-internal Collaborative Cloud part
- Consideration of split functionalities after Intercloud VM cloning

  Split functionalities means that in our example the CSCF-VM is transferred to a foreign Collaborative Cloud while the subscribers and the HSS database still reside in the original Operator Cloud. That has the consequence that the subscriber traffic of ~ 25 Mbps per direction and the HSS database traffic of ~ 10 Mbps per direction must be transferred over the ACEX network interconnecting the Collaborative Clouds.
- Consideration of different network architectures concerning parallel transmission (meshed, star-like networks).

Use cases 1 and 2 considered the meshed ACEX network architecture while use cases 3 and 4 considered the star-like ACEX network. Use case 5 considered the concurrent transfer of VM images from different providers and all use cases considered all the split functionalities after Intercloud VM cloning.

4.6 Results

This simulation was carried out and analyzed under normal network conditions and by applying irregular traffic conditions. One of the simulation results was to show the time to readiness of the cloud system. The results are represented and discussed in this chapter.

4.6.1 Use Case 1

In this use case a sequential data transfer of the VM-images was applied. That means, only one image is transferred at the same time. This scenario refers to Figure 4 and considered the transfer of one 2 GB image to four foreign cloud segments (partner clouds b to E also referred to as collaborative clouds). The simulation results showed not the expected increasing delay caused by increasing link load directly. Figure 8 shows the simulation results; the transfer time for the images varied only slightly.

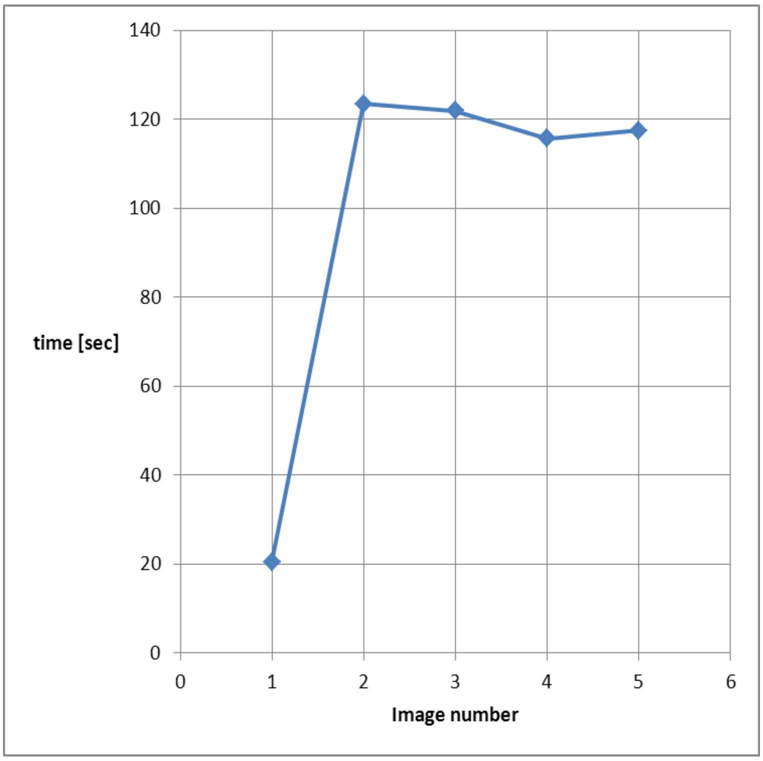

Figure 8: Sequential image transfer time

The reason was the meshed structure in the core network (ACEX_A to ACEX_E) of the network model.

As represented in Figure 6, the image data transfer started after finishing the image preparation process and this transmission was repeated four times. This caused a constant traffic load because the necessary bandwidth for this transfer in a given time of 120 seconds was limited to 150 Mbps in the core network – but not in the connections between the core network and the load balancers of each network provider.

The core net provided a free link capacity of 150 Mbps between each of the routers and by this the traffic load was split over all available connections; the result was sufficient bandwidth during all load situations (link utilization <25% to both directions, see Figure 9).

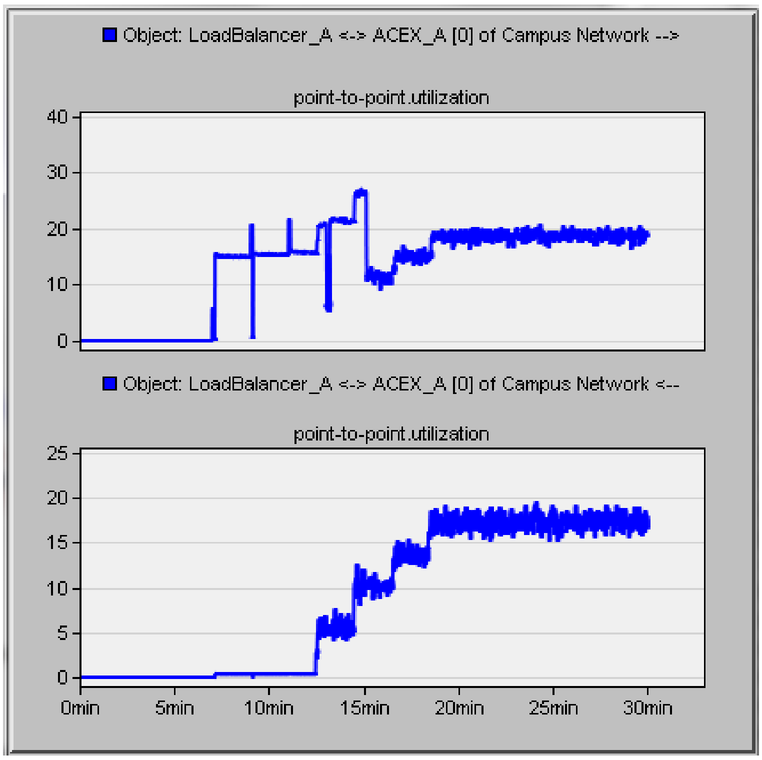

Figure 9: Link utilization for a sequential image transfer

It was shown that increasing traffic load had less influence on the image transfer time. The impact of limited bandwidth was considered in Use Case 3. Principally sequential image replication provided early operational readiness in live cloning in a cloud application.

4.6.2 Use Case 2

Scenario 2 was the same as described in 4.6.1 but the image data transfer was carried out as parallel transfer to four collaborative cloud servers. Due to no significant bandwidth limitations and the low transfer time (Figure 10), compared to use case 1, this seems to be the preferred transfer process (Figure 11).

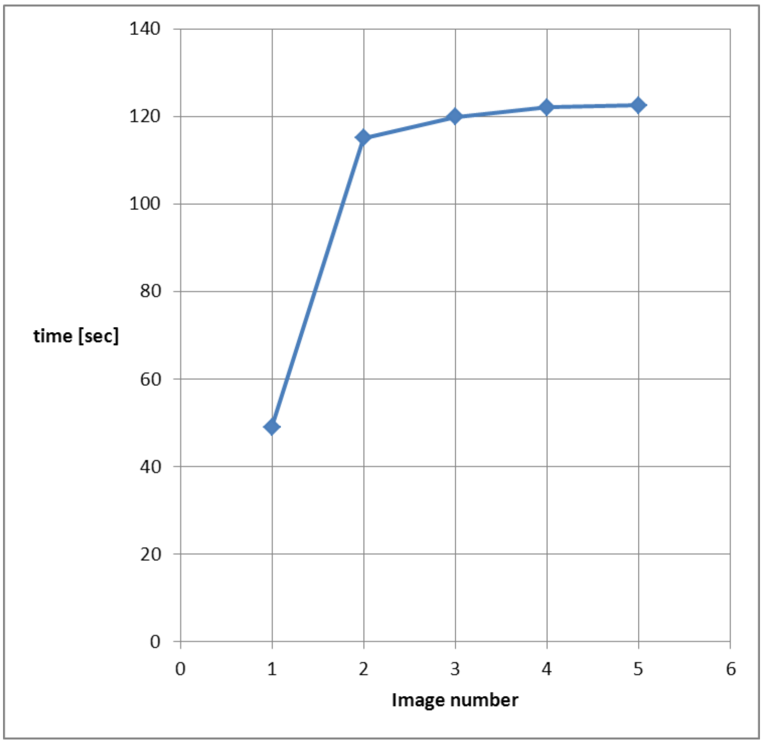

Figure 10: Parallel image transfer time in unlimited bandwidth

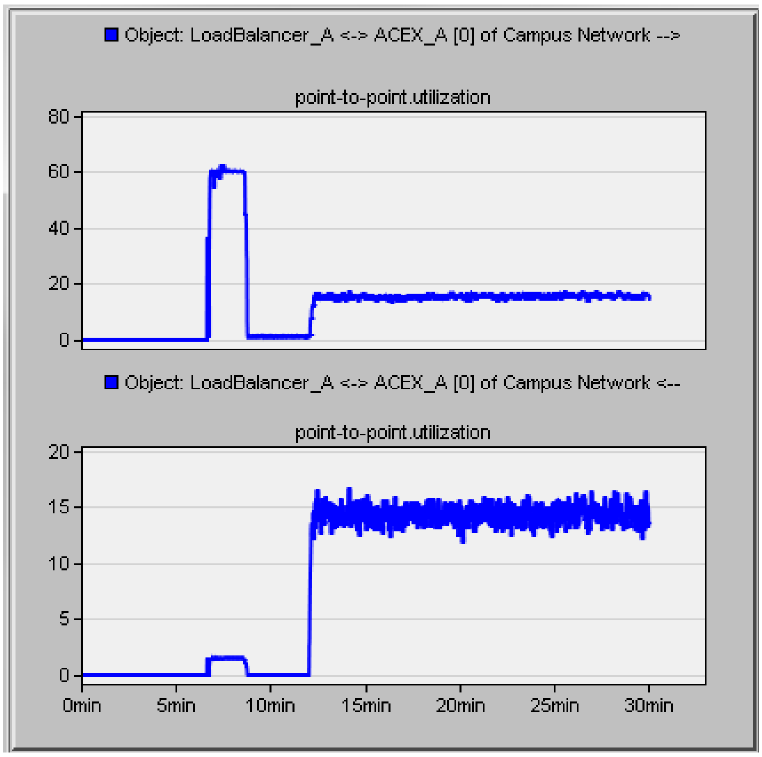

Figure 11: Link utilization for the parallel image transfer

4.6.3 Use Case 3

The impact of limited bandwidth in the ACEX network using a star-like architecture was analyzed in this scenario by limiting the bandwidths to 150 Mbps in all links between the core net (ACEX_x) and the load balancers (LB_x) in the simulation scenario. This bandwidth limitation caused an increasing congestion of the network connection between the initiator of the image transfer and the cloud. Figure 12 shows this effect: the first image was transmitted as expected but following image transfers could not use the whole link capacity. The reason was the starting data communication between this established cloud server and the overloaded server as well as the necessary access of this cloud server to the provider database server.

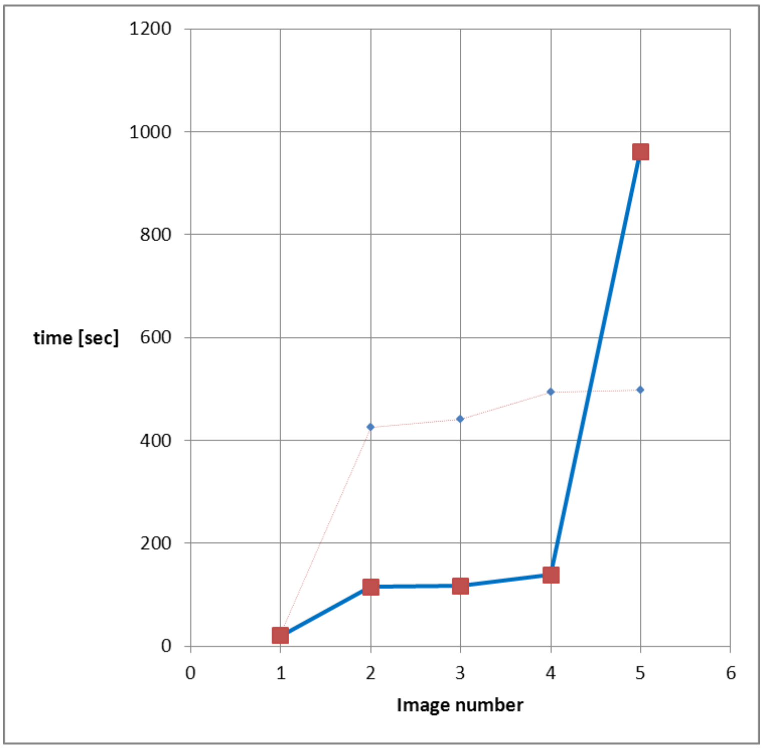

Figure 12: Sequential image transfer duration in bandwidth limited networks

It was visible that each active cloud server caused increasing traffic load on provider connection LB_A – ACEX_A. The impact of this constrained traffic flow on the utilization of this link is represented in Figure 15. The image transfer caused a link utilization of 100% between load balancer LB_A and router ACEX_A (upper trace). The traffic from ACEX_A back to LB_A also represented in Figure 13 (lower trace) shows the increasing communication traffic between the cloud resources and Server_OperatorCloud_A1 (the router under overload).

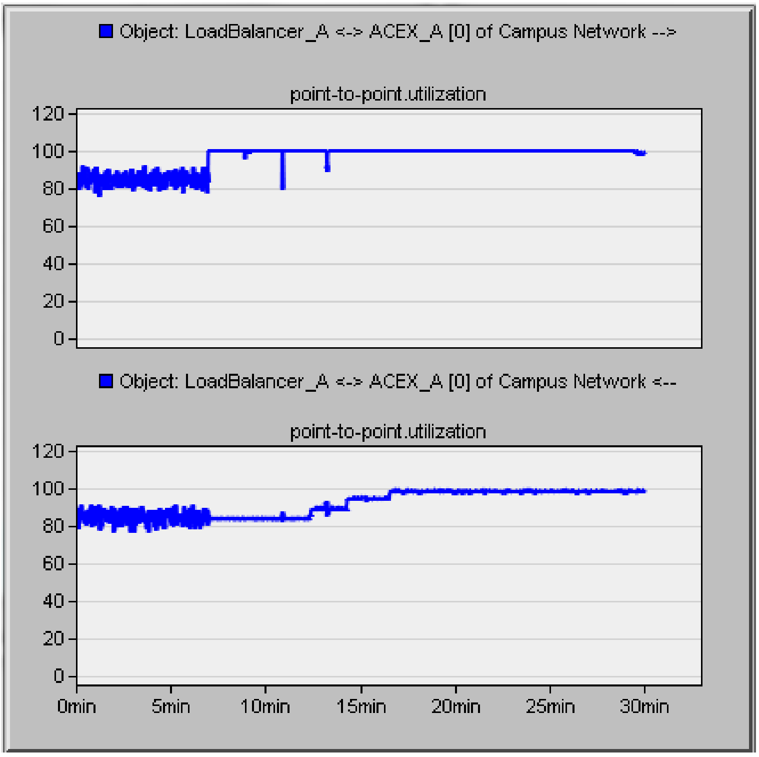

Figure 13: Link utilization for the sequential image transfer in bandwidth limited networks

4.6.4 Use Case 4

Parallel image transfer promised immediate availableness of all cloud resources. As we discovered in the simulation, a parallel transfer provided all resources within a shorter time but the duration to get them in operation was significantly longer (Figure 14, Figure 15).

The first cloud element was located in the own cloud and the bandwidth for the image transfer was sufficient (1 Gbps). In these circumstances the image transfer was finished after approximately 20 seconds. But the transfer to all remaining foreign resources claimed substantially more time. Only at an increasing number of servers the parallel image transmission was advantageous (Figure 15).

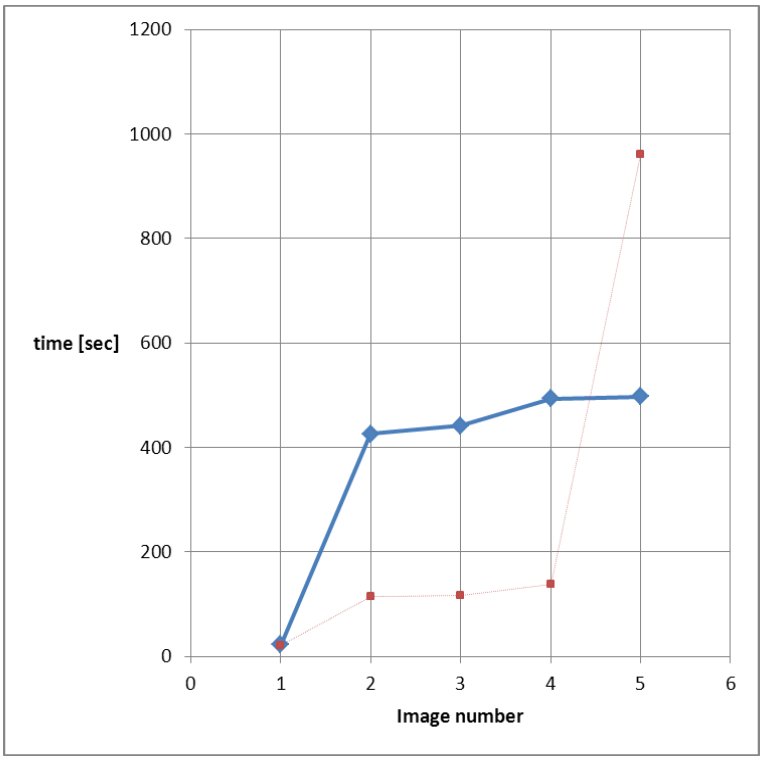

Figure 14: Parallel image transfer duration in bandwidth limited networks

Figure 15 shows that parallel transmission depends very strongly on the available bandwidth. This has to be considered in the resource planning. It is also visible that a parallel transfer has a strong influence on the remaining data traffic. The remaining traffic increases rapidly after finishing the image transfer and the initialization of the cloud servers because of the simultaneous starting of the data traffic between cloud servers and local servers using load balancer functionality.

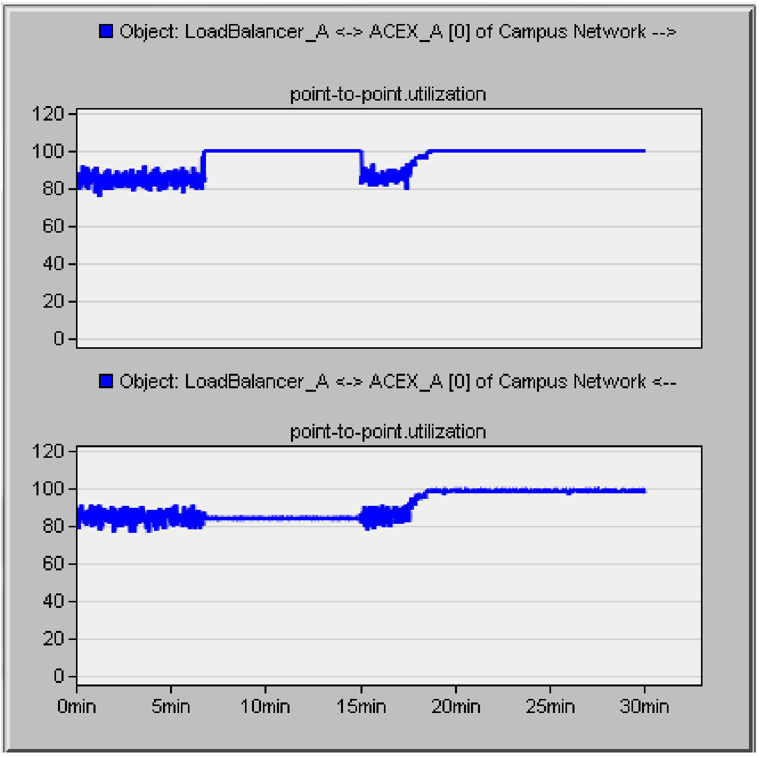

Figure 15: Link utilization for the parallel image transfer in bandwidth limited networks

4.6.5 Use Case 5

Since it was assumed that an overload situation will not be limited to one provider, a ‘mixed scenario’ was simulated in a further step. This scenario considered the sequential transfer of a 2 GB image from Server_OperatorCloud_A1 (see Figure 4) to the own cloud resources (Server_CollabCloud_A1) and to four foreign collaborative clouds (Server_CollabCloud_B1/ C1/D1/E1).

In difference to all previous scenarios the simulation model was extended in provider network D by an Operator Legacy Network_D, a Server_OperatorCloud_D1 and a Database_ OperatorCloud_D1 similar to provider A.

It was assumed that both providers start the cloning process at the same time. Provider D preferred a parallel transfer for the image from Server_OperatorCloud_D1 to three foreign cloud resources (Server_CollabCloud_B2/C2/E2). The operational readiness for both cases is shown in Figure 16. It is visible that network operator A (sequential transfer) gets early an increasing access to the requested cloud resources but the full availability is delayed. Network operator D (parallel transfer) gets access to all resources nearby simultaneously but he has to take into account the delayed availability of them.

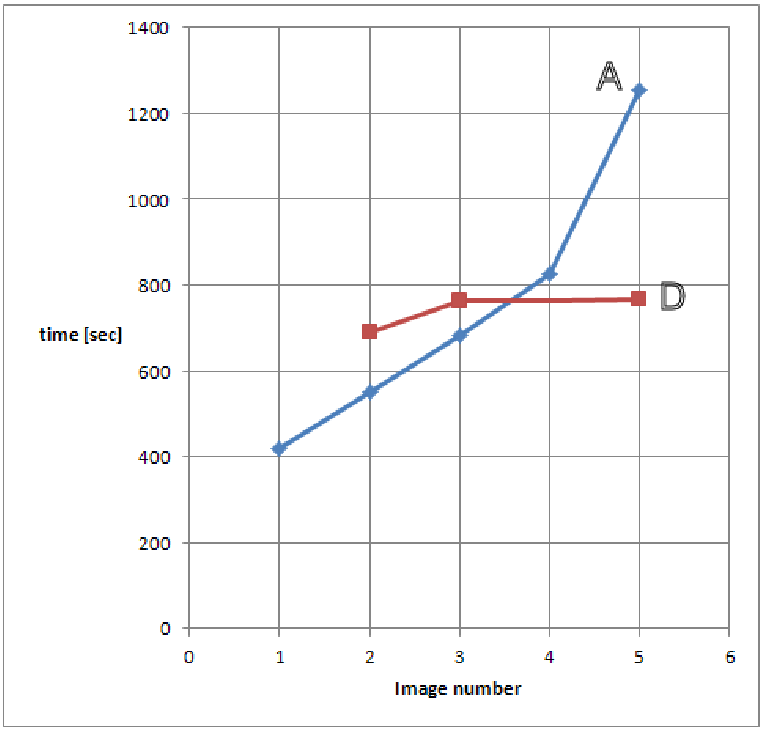

Figure 16: Cloud operational readiness

The link utilization for provider link A is shown in Figure 18 and Figure 19 and represents this behavior for provider link D. It is visible that a sequential transfer gives advantages regarding operational readiness. First resources have been available after a short time and a continuous increasing number of available resources were provided afterwards. This could be helpful to prevent a server from rapidly increasing access rates.

Figure 17 depicts the amount of transfer time for each image for the sequential case (A) as well as the parallel one (D). It is shown that during a sequential VM-image transfer the time will increase because each transferred image will cause additional data traffic between cloud and supported servers and the link utilization also will increase.

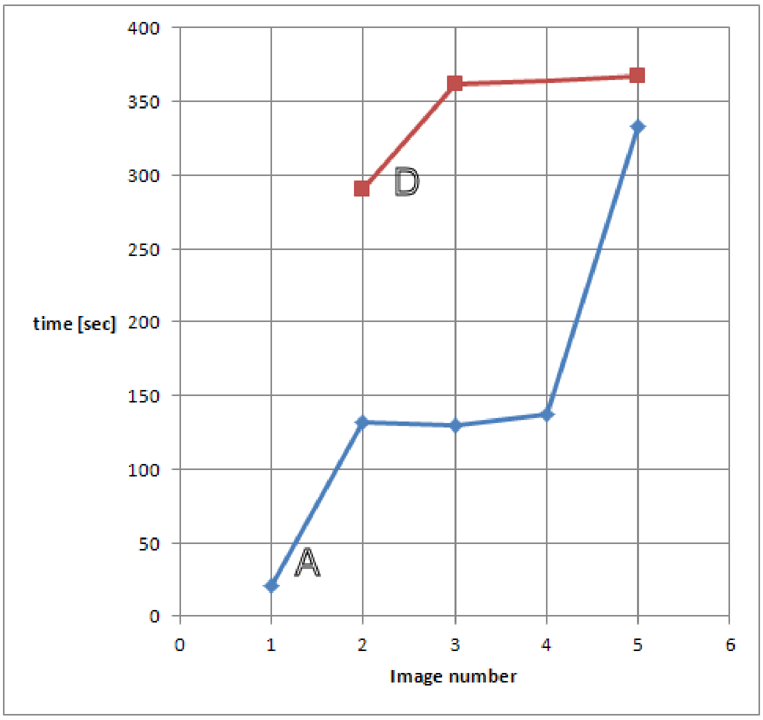

Figure 17 Image transfer time for sequential and parallel transfer

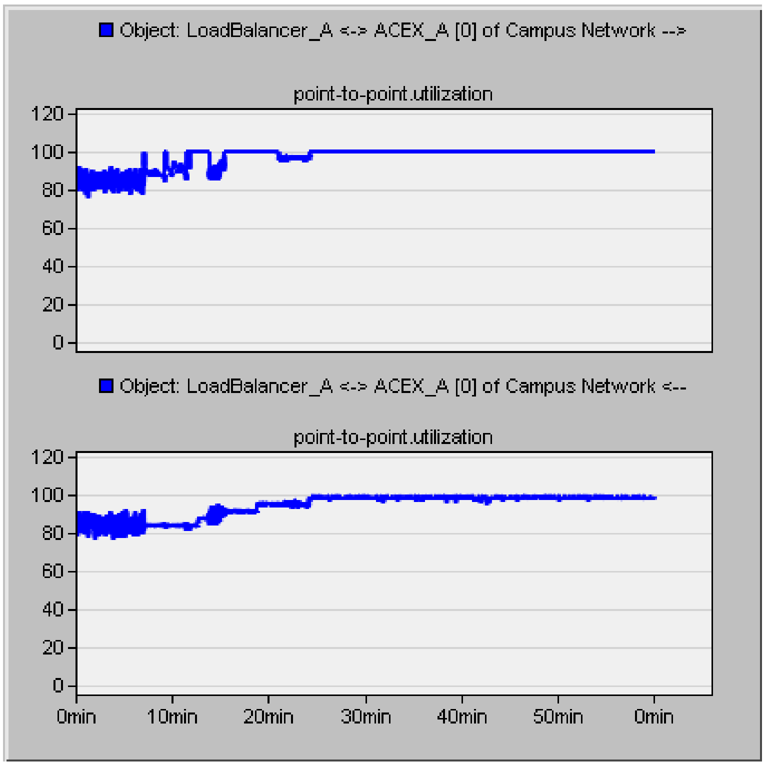

Figure 18: Link utilization provider link A

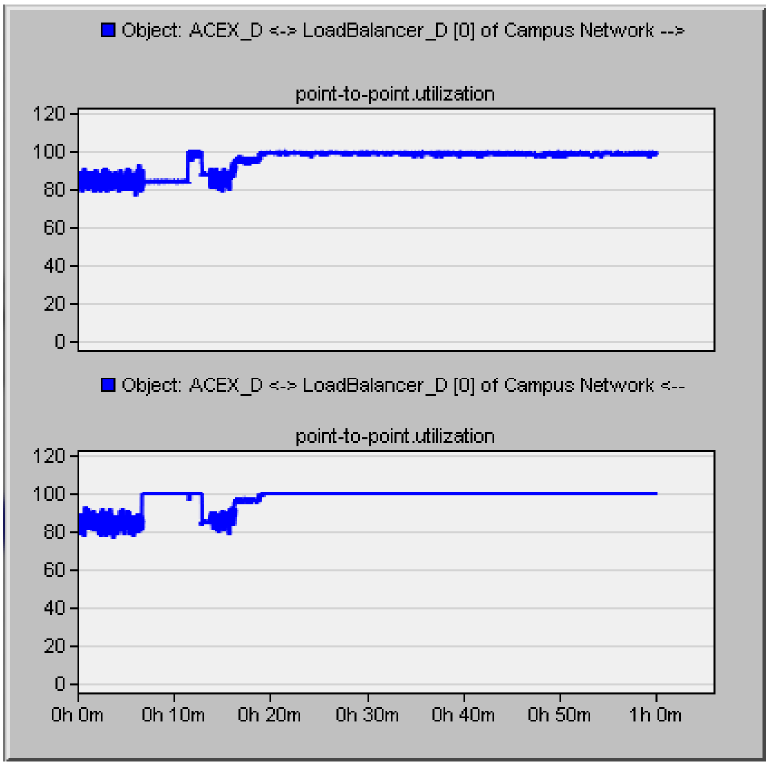

Figure 19: Link utilization provider link D

All above described scenarios are a selection from a number of possible use cases. Due to limited resources the simulation of huge numbers of transfers and large Image sizes was not possible in an arguable time. A number of significant scenarios had been simulated, examined and compared concerning their effectiveness and applicability in this study part and based on the recommendations in chapter 4.1 derived from the Intercloud demonstrator.

5 Conclusion and Outlook

In any case attention should be paid to the Live Cloning process for the discharge of overloaded servers. The server overload problem is of interest for a number of reasons: first, the cost of rejecting a request could not be ignored as compared to the cost of serving a request; second, the various server time-outs lead to many retransmissions in overload and amplify the situation; third, in case of SIP it has a server-to-server application level routing architecture. The SIP server-to-server architecture supports the deployment of a pushback SIP overload control solution and also the concept of elastic clouds. Elastic cloud systems can help to reduce the resources expenditure of the individual net providers and to lower the costs for redundant systems. Our investigation showed that, at present, network attacks are partly defendable however not avoidable. The statement regarding the hardening of the virtualization host systems resulting in a minor non-compliance resulting from an analysis against international standards, internal project goals, and best practices and how to handle overload situations was one of the bases for the architecture definition and for the definition of the key parameters for the validation of the process in a demonstrator and in a simulation environment.

As key parameters for the evaluation of the results of the ASMONIA demonstrator were defined the size of the live-cloned OpenSIPS VM and the bandwidth provided between the two Collaborative Clouds. The investigation on the live cloning process in a demonstrator with focus on the main security objective of the ASMONIA Intercloud demonstrator included how to maintain or increase the availability of cloud-based telecommunication infrastructures in case of a DoS attack.

The necessary steps and the time for each step like detecting the overload status, creating a server-image, the preparation for a transfer and the necessary time for the transfer of the image had been estimated on this demonstrator. These measured results were the basis for designing the simulation model. The findings from the simulation hardly depend on given requirements. In the context of this study we examined some chosen operating models based on a selection of possible use cases described in chapter 4.1. All simulations carried out on this model showed the advantages as well as the disadvantages of different distribution methods (sequential or parallel transfer of a VM-image to a backup server in the cloud).

Finally these investigations can only give some advices how you could prepare your environment against cyber-attacks but cannot give you a complete tutorial how you can do this. Nethertheless the simulated use cases helps you to get an understanding of the mechanisms and processes necessary for Intercloud networking. An applicable solution has to be defined according to the necessities of a provider or a group of cooperating providers. Further investigations are possible in more detail with additional information for defined use cases. The investigations showed that, by the employment of future technologies, cloud technologies will get increasing importance.

References

[ACM] http://dl.acm.org/citation.cfm?id=2046664

[CSA_TopThreats] Cloud Security Alliance, “Top Threats to Cloud Computing V1.0”, March 2010

[ASMONIA_D3.1] Asmonia D3.1, “Analysis of Requirements for the Deployment of Cloud Systems”

[ASMONIA_D3.2] Asmonia D3.2, “Architecture Concept for the Use of Cloud Systems”

[ASMONIA_D3.3] Asmonia D3.3, “Design and Implementation of an Intercloud demonstrator”

[ENISA_CloudRisks] ENISA, “Benefits, risks and recommendations for information security”, November 2009

[Fallenbeck2011] Niels Fallenbeck, “Virtual Machine Image Management for Elastic Resource Usage in Grid Computing”, PhD Thesis, Philipps-University Marburg, Germany, 2011

[Insinuator] http://www.insinuator.net/2012/05/vmdk-has-left-the-building/

[Nagios] Nagios, http://www.nagios.org, March 2013

[NIST_GLOSSARY] NIST IR 7298, "Glossary of Key Information Security Terms Richard Kissel", April 2006 <http://csrc.nist.gov/publications/nistir/NISTIR-7298_Glossary_Key_Infor_Security_Terms.pdf>

[NIST_SP800-144] NIST SP800-144, “Guidelines on Security and Privacy in Public Cloud Computing”, Wayne Jansen and Timothy Grance, December 2011

[NIST_SP800-145] NIST SP800-145, “The NIST Definition of Cloud Computing”, Peter Mell and Timothy Grance, September 2011

[OPNET Modeler] http://www.riverbed.com/products-solutions/products/network-planning-simulation/Network-Simulation.html, 2013

[VMWareSecAdv] http://www.vmware.com/security/advisories/VMSA-2012-0009.html

Abbreviations

Abbreviation	Explanation
ACEX	ASMONIA Cloud Exchanges
ACGW	ASMONIA Collaboration Gateway
ACM	Association for Computing Machinery
CPU	Central Processing Unit
CSA	Cloud Security Alliance
CSCF	Call Session Control Function
DoS	Denial of Service
ENISA	European Network and Information Security Agency
HSS	Home Subscriber Server
IP	Internet Protocol
NIST	National Institute of Standards and Technology
OVF	Open Virtualization Format
RAW	Raw image format
SIP	Session Initiation Protocol
TCP	Transmission Control Protocol
VDI	Virtual Disc Image
VM	Virtual Machine

Revision History

Version	Date	Changes
0.1	2013-04-04	Initial version.
0.2	2013-05-28	Review version
1.0	2013-07-01	Final version