Network Design Requirements for Disaster Resilience in IaaS Clouds

Rodrigo de Souza Couto[1], Stefano Secci[2],
Miguel Elias Mitre Campista[1],
Luís Henrique Maciel Kosmalski Costa[1]

[1] Universidade Federal do Rio de Janeiro - PEE/COPPE/GTA - DEL/POLI
P.O. Box 68504 CEP 21941-972, Rio de Janeiro, RJ, Brazil
[2] Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6
F-75005, Paris, France

Abstract

Many corporations rely on disaster recovery schemes to keep their computing and network services running after unexpected situations, such as natural disasters and attacks. As corporations migrate their infrastructure to the Cloud using the Infrastructure as a Service (IaaS) model, Cloud providers need to offer disaster-resilient services. This article provides guidelines to design a data center network infrastructure to support a disaster-resilient IaaS Cloud. These guidelines describe design requirements, such as the time to recover from disasters, and allow the identification of important domains that deserve further research efforts, such as the choice of data center site locations and disaster-resilient virtual machine placement. [1]

Index Terms

Cloud Networking, Disaster Resilience, Network Design, Infrastructure as a Service.

I. Introduction

Cloud Computing is revolutionizing the way IT services are deployed and consumed. Under the Infrastructure as a Service (IaaS) model, clients can outsource their entire IT infrastructure, running services inside virtual machines (VMs) hosted at the provider's substrate. To encourage IaaS subscriptions, Cloud providers usually employ resilient servers and network infrastructure [1].

Resilience of network services can be expressed as a Quality of Service (QoS) metric or, more specifically, as a Quality of Resilience (QoR) metric [2]. Typical QoR metrics are the service availability and time to recover from failures. QoS evaluation, on the other hand, addresses other metrics such as network latency and packet loss ratio. Generally, IaaS Cloud providers express their QoR in terms of VM availability over a given time interval, defining it as a Service Level Agreement (SLA). For example, Amazon Elastic Computer Cloud (Amazon EC2)[2] and Rackspace Cloud Servers[3] guarantee an IaaS availability of 99.95% and 100%, respectively. In such cases, the service is considered unavailable if all running VMs of a client have no external connectivity. The IaaS provider commitment is to refund the client proportionally to the experienced downtime. Some IaaS providers also define resilience in terms of the redundancy of their infrastructure. For example, Rackspace specifies that its physical servers are equipped with RAID 10 technology and have redundant power supply. Moreover, Rackspace's physical

[1]This paper was accepted for publication in the IEEE Communications Magazine October 2014 issue.

[2]Amazon EC2 SLA is online at http://aws.amazon.com/ec2-sla.

[3]Rackspace Cloud Servers SLA is online at http://www.rackspace.com/information/legal/cloud/sla.

network is dimensioned so that one failure in an upstream switch halves the bandwidth, instead of leaving the whole service down.

A common characteristic of all mentioned IaaS SLAs is that they do not cover failures out of the IaaS provider's control (e.g., a denial of service attack) and other force majeure events, such as hurricanes. In other words, a typical IaaS SLA does not consider disaster resilience. Nevertheless, an IaaS provider could be disaster-resilient, guaranteeing a given QoR after a disaster occurrence [3].

A disaster-resilient IaaS provider employs backup VMs in standby mode, which are only activated upon a disaster. Moreover, a working VM must be geographically isolated from its backup so that a disaster does not affect both. Hence, the DC needs to be geo-distributed and requires a Cloud network that is itself resilient to disasters and cost-effective. The design requirements for disaster-resilient IaaS scenarios are still an open issue, despite of their importance to allow Business Continuity Planning (BCP) for IaaS providers. BCP consists of several requirements, technical or non-technical, to guarantee that some services are available even when disasters occur. To make the IT infrastructure compliant with the organization's BCP, the IT staff must adopt a process called IT Service Continuity Management (ITSCM), which can be performed according to different frameworks, such as the set of procedures defined in the Service Design stage of the Information Technology Infrastructure Library (ITIL) [4] and the standard ISO/IEC 24762:2008 [5]. Implementation and testing of recovery schemes are examples of such procedures.

This work provides guidelines to design a DC network infrastructure supporting a disaster-resilient IaaS Cloud, organized as interrelated phases. The first one starts with the initial design considerations, such as assessing disaster risks and defining client requirements. In subsequent phases, disaster recovery mechanisms are chosen, as well as the network infrastructure and the VM placement scheme. It is important to note that our proposed guidelines do not intend to replace existing ITSCM frameworks, which have a broader scope, but act in conjunction with them to support a disaster-resilient IaaS Cloud. Moreover, we draw attention to incipient research topics, such as the physical design of a geo-distributed DC and the placement of VM backups.

II. Geo-distributed Data Center Network Design

The sites of a geo-distributed DC are spread over a geographic region and connected through a Wide Area Network (WAN). The geo-distribution increases the resilience of the Cloud network and makes DC sites closer to end users to reduce the access latency. Figure 1 illustrates a geo-distributed DC. Each dashed circle is a DC site, where servers are connected using an intra-DC fabric, composed of racks and Top of Rack (ToR) switches connected to core switches. DC sites can host different numbers of servers and are interconnected using long-haul links. Some sites may employ redundant links between each other to improve resilience. Finally, cloud users access DC services through gateways spread over the network.

Based on the literature about disaster resilience in optical WANs [6], [7], [8] and resilience in Clouds [1], [3], we draw guidelines to design a DC network infrastructure supporting a disaster-resilient IaaS Cloud, consisting of five interrelated phases summarized in Table I. "Planning" is the first phase, where all design requirements are defined. The "Modeling" phase is employed to describe the relationship between the requirements and the DC components designed on the next three phases. The "Modeling" phase should be executed after the end of each of the three upcoming phases, to improve the model according to the employed mechanisms. Based on the requirements, at the "Selection of Disaster Recovery Mechanisms" phase the DC designer makes, among other design decisions, the choice of the frequency in which backups are performed for each type of client. The "Site Placement and Topology Design" phase is employed to design the WAN infrastructure, based on the backup frequency and other requirements, by dimensioning the

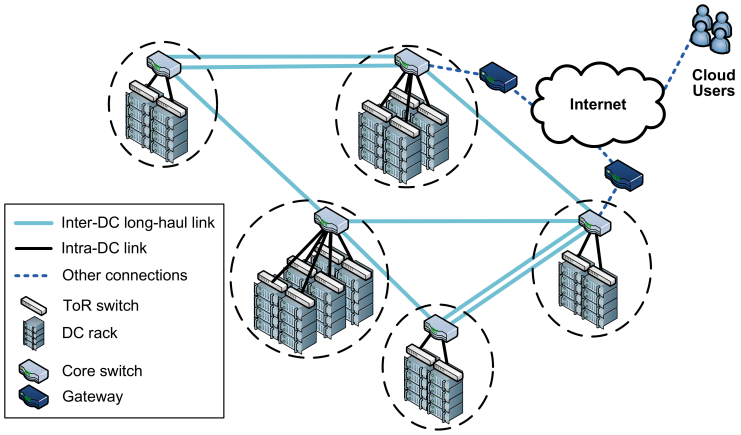

Fig. 1. Example of a geo-distributed DC composed of interconnected DC sites.

network capacity and the number of servers installed in each DC site. Finally, at the "Selection of VM Placement Mechanisms" phase the designer selects and configures mechanisms to place VMs based on the WAN and on the design requirements. Normally, the three last phases should be executed in the order given in Table I. Nevertheless, a designer might go back to a previous phase if its decisions precludes the accomplishment of the current phase. For example, when the chosen backup frequency demands an unfeasible WAN network capacity. Next, we detail each of the proposed phases. Some are organized in tasks, which consist of generic procedures important to accomplish each phase. Nevertheless, the list of tasks is non-exhaustive and thus more specific tasks can be added to each phase depending on the considered scenario and technologies.

TABLE I
GEO-DISTRIBUTED DC DESIGN PHASES.

Name	Design Goals
Planning	Assess possible disaster risks, define QoR and QoS requirements and budget constraints.
Modeling	Define network and failure models to be used on all design phases.
Selection of Disaster Recovery Mechanisms	Select the mechanisms of disaster detection, VM recovery, and network reconfiguration after disasters.
Site Placement and Topology Design	Define which locations in a geographical area are used by DC sites and design the interconnection WAN.
Selection of VM Placement Mechanisms	Select mechanisms to place VMs on the DC, specifying their policies regarding the isolation of backup and working VMs and the fulfillment of QoS and QoR requirements.

A. Planning

In this phase, the initial DC design planning is performed through the following tasks.

1) Definition of Disaster Risks: Disaster risks are disaster situations to be considered in the Cloud network design, such as large-scale power outages and hurricanes. To this end, all the possible disasters need to be listed and their effect on the DC infrastructure assessed. From this list, a subset of disaster types is selected for the next phases by analyzing the importance of each one. For example, the provider may want to ignore an unlikely disaster type.

As disasters are difficult to predict, more generic disaster situations can also be considered. For instance, one strategy could be to design a DC resilient to any entire-site or link failure, or

to failures on all elements inside a region [6].

2) Definition of Disaster Resilience Requirements: In this task, the Cloud provider defines the QoR and corresponding SLAs. The average values of QoR, such as the availability used by Amazon EC2, are generally not suitable for qualifying the disaster resilience of an infrastructure, since disasters can be very rare [1]. Instead, the most common disaster-related QoR metrics are the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO). RTO accounts for the time needed to restore a service after it has been affected by a disaster. For a given IaaS client, the RTO depends on the time to detect a failure, to restore the affected VMs from a backup site, to restart all the services running on these VMs, and to redirect the network traffic from the original site to the backup site. The other metric of interest, RPO, is the time lapse between the last backup of the service components (e.g., copy of virtual disks) and the disaster. The RPO gives an idea of the data loss after a disaster. Indeed, some services require a low RPO (e.g., banking transactions), and therefore continuous data replication. A low RPO implies high network bandwidth between DC sites to exchange large amounts of data. Both RTO and RPO levels can span from a few minutes to several hours [1].

3) Definition of Design Constraints: The design constraints are aspects to be considered regardless of the disaster resilience. Among the most important constraints are the QoS requirements, which influence the Quality of Experience (QoE). According to ITU-T Rec. P.10, QoE is "the overall acceptability of an application or service, as perceived subjectively by the end-user" [9], meaning that the QoE depends on all infrastructure elements, such as the network and physical servers, and on external factors such as the service price. Indeed, we have to ensure both QoR and QoS requirements to provide a good QoE. However, we need to guarantee the QoR without compromising the QoS, since disaster events are rare while QoS metrics are permanently perceived by end-users, directly or indirectly. For example, the DC geo-distribution, aiming to improve disaster resilience, increases the distance between DC sites and may increase the service latency, a QoS metric, when there are multiple VMs running across different sites. Therefore, in this task, the Cloud provider should list all of the QoS metrics to ensure that the next design steps consider these requirements. Moreover, QoR requirements that are unrelated to disasters, such as availability, must be considered if they appear in the SLAs.

The constraints also include other factors such as the maximum budget to build the DC, the possible geographical sites to install the DC, and other constraints related to the site physical capacity. As an example of the last one, the IaaS provider may need to install a minimum number of servers in a site according to the expected demand of a region.

B. Modeling

This phase defines models that capture the characteristics of the scenario defined in the "Planning phase". Also, the model describes how the DC components, defined on the three next design phases, affect the infrastructure characteristics (e.g., cost, QoS metrics, RPO, and RTO). Note that this phase stitches all design phases together, defining their relationship, which can vary depending of the considered scenario. The models defined in this phase basically take into account the disaster information and network parameters.

The disaster information, gathered on the Planning phase, is used to build disaster models. Disaster models for communication networks can be deterministic, probabilistic, or based on multi-layer networks [6]. A classical deterministic model is the utilization of Shared Risk Group (SRG) sets. An SRG is a set of infrastructure components susceptible to a common disaster situation. For example, considering power outages, an SRG is the set of DC sites served by one power plant. In opposition, probabilistic models consider that each component or set of components fail with a given probability independent of zones. As disasters and their impact on

the network are difficult to predict and are not frequent, deterministic models are preferable. The approach that considers multi-layer networks is an incipient research topic, which separately models failures in each network layer. For example, a single failure on the physical layer, such as cable cuts, can affect multiple IP routes, which can thus break several TCP connections. On the other hand, a recovery from a cable cut can be rapidly addressed by lower layers, being thus unnoticeable by upper layers. Multi-layer models are more complex and require more information of the environment than deterministic and probabilistic ones.

Network parameters, such as traffic distribution, length of network links, and available bandwidth are modeled by conventional network modeling approaches. For example, the network can be modeled as a graph, where the nodes are DC sites and the edges are the links between them. The edges can be weighted according to QoS parameters such as latency between DC sites. The latency could be found after running shortest path algorithms in the graph. Graph models can be combined with SRG information to capture disaster risks. In this case, the SRGs defined on the Disaster Model are composed of nodes and edges. Using graph theory, we can measure, for instance, which DC sites are affected by each SRG. A model that captures the resilience metrics based on network parameters is still an open research issue. This model captures, for example, how the increase of bandwidth between DC sites affects the RPO levels offered by the IaaS provider. In addition, the model could describe how the network reconfiguration and activation of backup VMs, described in the next section, affect RTO levels.

C. Selection of Disaster Recovery Mechanisms

In this phase, the Cloud provider chooses the recovery mechanisms, which impact directly the RTO and RPO, performing the following tasks.

1) Selection of Disaster Detection Mechanisms: Despite all failure detection mechanisms employed in the network layers (e.g., reaction of routing protocols), the DC must employ a mechanism to define when to migrate services from their working sites to backup sites. As the RTO depends on the reaction time to a disaster, failure detection plays an important role on disaster recovery. It can be done by periodically probing DC sites, using network alarms, etc. Obviously, the more frequent the probes and the network alarms, the shorter the RTO but at the cost of more control traffic.

2) Selection of VM Recovery Mechanisms: A suitable strategy to after-disaster VM recovery is to use VM snapshots [1]. A snapshot is a copy of the VM state at a given moment, which can include its disk, memory state, and settings. Most of the virtualization platforms support snapshots. Hence, the DC can maintain a snapshot of its VMs in backup sites and activate them after a disaster. Note that this scheme forces the services running on the VM to return to a previous state, affecting the RPO. More frequent snapshots translate to shorter RPO but spend more network resources for snapshot transfers. Indeed, the choice of the frequency to perform snapshots depends on the QoR classes and on the defined constraints.

3) Selection of Network Reconfiguration Mechanisms: When a VM starts running in another DC site after a disaster, the network infrastructure must re-route the traffic destined to this VM. This design task selects adequate mechanisms to perform network reconfiguration upon a disaster. IaaS providers generally employ Domain Name System (DNS) services to redirect the VM's traffic when it changes physical location. The Cloud DNS server is thus responsible for replying DNS queries with the current VM IP address. For example, Amazon Web Services (AWS) provides a DNS service called Amazon Route 53. The Cloud DNS services generally rely on anycast routing, where any node running the desired service can respond to requests,

enabling simple network reconfiguration after disasters[4]. Alternatively, providers can rely on cloud network overlay protocols supporting various forms of Ethernet or IP packet encapsulation to enable flexible VM location and adequate isolation of IaaS slices [10]. Note that network reconfiguration mechanisms have a high impact on the RTO, since they affect the period that VM network endpoints remain unreachable after a disaster.

D. Site Placement and Topology Design

In this phase, the DC inter-site WAN topology is chosen, as well as the location of each DC site. Although the intra-site local network is also a DC design decision, it is generally not concerned with disaster resilience, since a disaster typically causes the failure of an entire DC site. Hence, intra-site redundancy is generally employed to achieve high DC availability but not disaster resilience [11].

1) DC Site Placement: This task defines where to install DC sites in a geographic region to minimize the impact of disasters. Moreover, the site placement defines how many servers are installed in each DC site, and how many servers are left to host backup VMs.

Distributed DCs tend to span different SRGs being thus more disaster-resilient. Figure 2 illustrates different DC distribution levels, using the WAN topology of the French Research and Education Network (REN), RENATER. Each circle represents a Point of Presence (PoP). A DC site is represented by a server figure, and consists of a PoP with at least one installed server. In the three distribution levels (i.e., number of DC sites employed) shown in this figure, we spread 1,024 servers across the network, choosing a given number of DC sites to use. Considering a single-failure disaster model, where each DC site or link belongs to a different SRG, the figure indicates the fraction of servers available after the worst-case failure (i.e., failure that disconnects the largest number of servers). A server is considered available if it has at least one path to a gateway, represented by a triangle in the figure. As this example network is highly redundant in terms of paths to gateways, the worst-case SRG is always the failure of an entire DC site. Hence, a high distribution level makes the DC more robust to disasters, as each DC site tends to have fewer servers. Note, however, that the difference between "Medium Distribution" and "High Distribution" in terms of resilience is small since, after a given level of distribution, we cannot significantly improve the resilience to disasters.

Another advantage of geo-distribution is that fewer backup servers are needed as the DC becomes more resilient to disasters, as exemplified in Figure 3. Consider that a disaster may bring down an entire DC site and that any pair of DC sites is distant enough to avoid simultaneous failures. The figure illustrates different DC placement schemes to provide backup capacity to 12 servers. The backup is done by periodically performing VM snapshots on the working sites and sending them to backup sites. Consider the 1-site case, where the DC has only one working site and all VMs are copied to a backup site. If the working site suffers a disaster, its VMs start to run on the backup site after recovery. Since we have a single working DC site and we need to support the failure of the entire site, the backup site must have the same capacity as the working one in terms of supported VMs and disk storage. Consequently, a single DC site represents an expensive option in terms of VM capacity used for backup. This situation is not different from the case where no virtualization is used and an entire DC site is allocated for disaster recovery. However, as shown in the 2-site case, we can distribute the working DC in two different sites, each with half the number of servers of the 1-site case. As W1 and W2 do not fail together, the backup site does not need to run both W1 and W2. Hence, the backup site needs only to support half of VMs as compared to the previous case. Nevertheless, the storage

[4]Rackspace Cloud DNS Overview is online at http://www.rackspace.com/knowledge_center/article/rackspace-cloud-dns-overview.

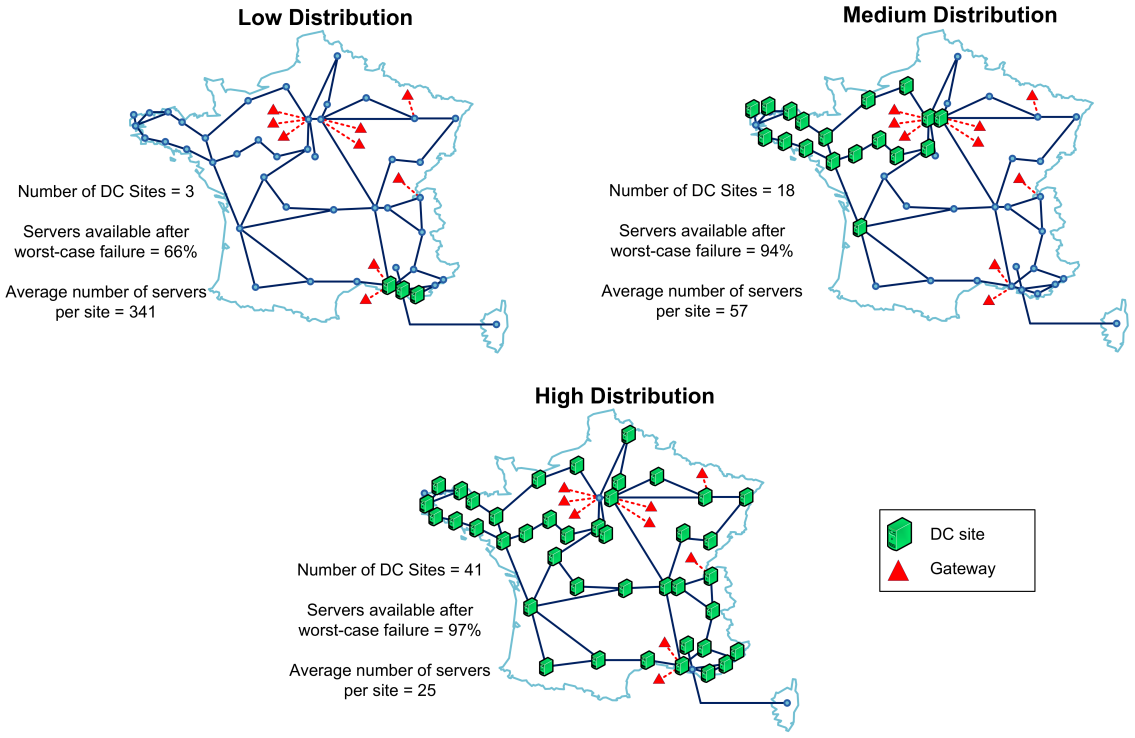

Fig. 2. Impact of DC distribution on the resilience.

capacity necessary is still the same since B1 must have the snapshot of all VMs. Using the same reasoning, the 4-site case reduces by four times the required server capacity as compared to the 1-site case.

Despite its advantages, DC distribution may be expensive. Referring again to Figure 3, each additional DC site requires one extra WAN link. Although the required capacity of WAN links decrease as we increase the distribution (i.e., fewer VM snapshots are transmitted over each WAN link), the site interconnection cost increases with the distance. The cost to install new DC sites should also be considered and depends on different factors, such as the security and availability concerns and the network capacity. Moreover, the cost of a single DC site may depend on its geographical location, being affected by factors such as the cost of the square meter on that location, cooling requirements given the weather conditions, local taxes, etc. For more information about DC costs, the reader may refer to the website The Cloud Calculator[5], which is an effort to estimate the cost of DC components according to different parameters.

Given the reasons above, the DC distribution should consider the constraints defined in the "Planning" phase, such as budget and QoS metrics. For example, in addition to performing the distribution considering the disaster resilience, the provider may prefer to install its sites closer to clients that have tighter latency requirements.

2) Design of the Inter-site WAN: This task is mostly characterized by the classical design of WAN networks for telecom operators. The literature on this problem is vast and generally addresses the design of optical networks [6]. One requirement of this network design is to improve the resilience of the network topology, by employing techniques such as path restoration and protection (e.g., provision of backup paths), multipath routing, and p-cycles [8]. Different from traditional telecom networks, in which the main goal is to connect several PoPs, in DC

[5]The Cloud Calculator is online at http://www.thecloudcalculator.com.

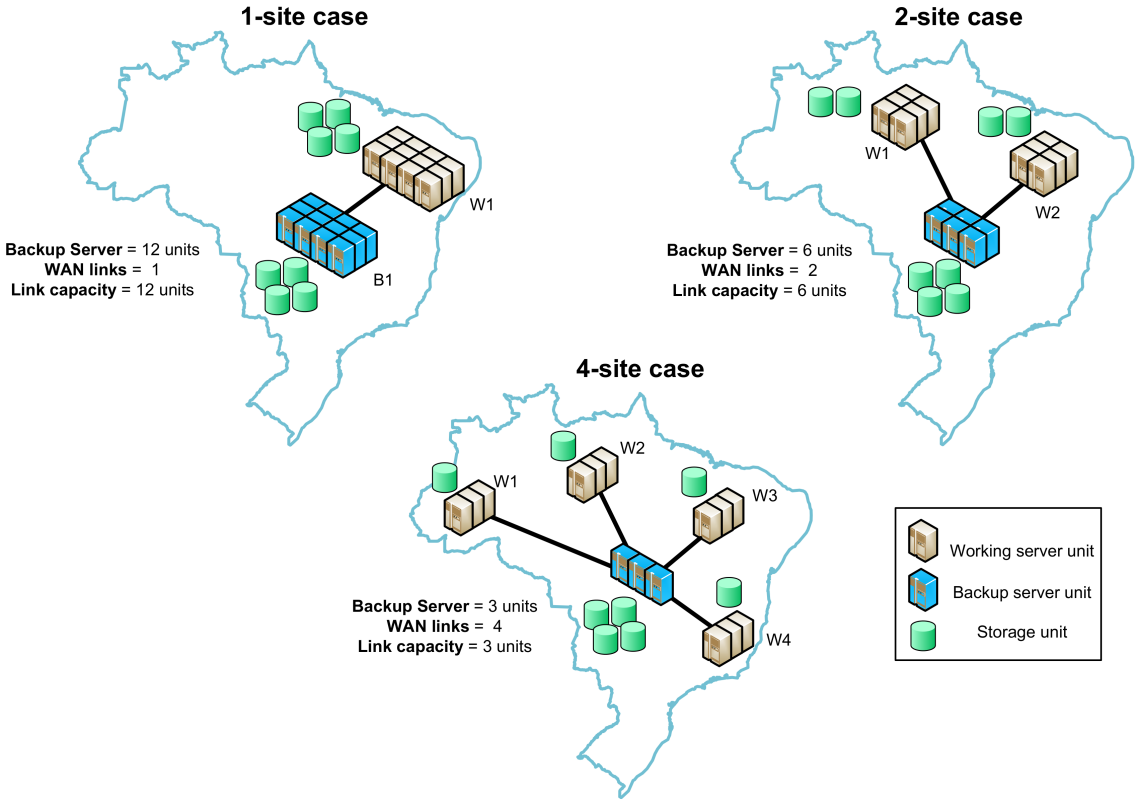

Fig. 3. Number of backup servers and DC distribution.

networks the goal is to provide VM hosting to clients spread over a region. Hence, the design of the inter-DC network must be correlated to the "DC site placement" task, since the network capacity and protection level provided to a DC site depends on the number of servers installed on it, as well as on its location. Also, this task defines the location of gateways, the capacity allocated to each link, the location of routers and switches used to interconnect sites, etc. [12].

E. Selection of VM Placement Mechanisms

Although the DC site placement and topology design play important roles with regards to disaster resilience, alone they do not guarantee the QoR requirements. QoR is also affected by the VM Placement Mechanism, which allocates VMs to physical nodes upon clients' requests. Generally, working VMs are allocated to maximize the provider's revenue and to meet the user requirements. However, when providing disaster recovery, IaaS providers must also allocate backup VMs for each working VM. Given that disaster resilience should not affect QoE under normal operation, the placement of backup VMs should be aware of user requirements such as QoS. A simple alternative to ensure this is by performing the VM placement in two phases. The first phase allocates the VMs requested by a client according to his/her requirements under normal operation; The second phase decides where to store the snapshots of each VM when they are covered by disaster recovery services.

Figure 4 shows a VM placement example. The DC is distributed across different sites in a region and the SRGs, circled by dashed lines, indicate which sites fail together. For each DC link, the available bandwidth between two sites is indicated. In the first phase, the VMs of a given client are allocated in two sites. This placement is performed according to the client QoS requirements or other QoR metrics that are not related to disasters, as the availability, not specified here. In the second phase, the placement mechanism decides where to store the

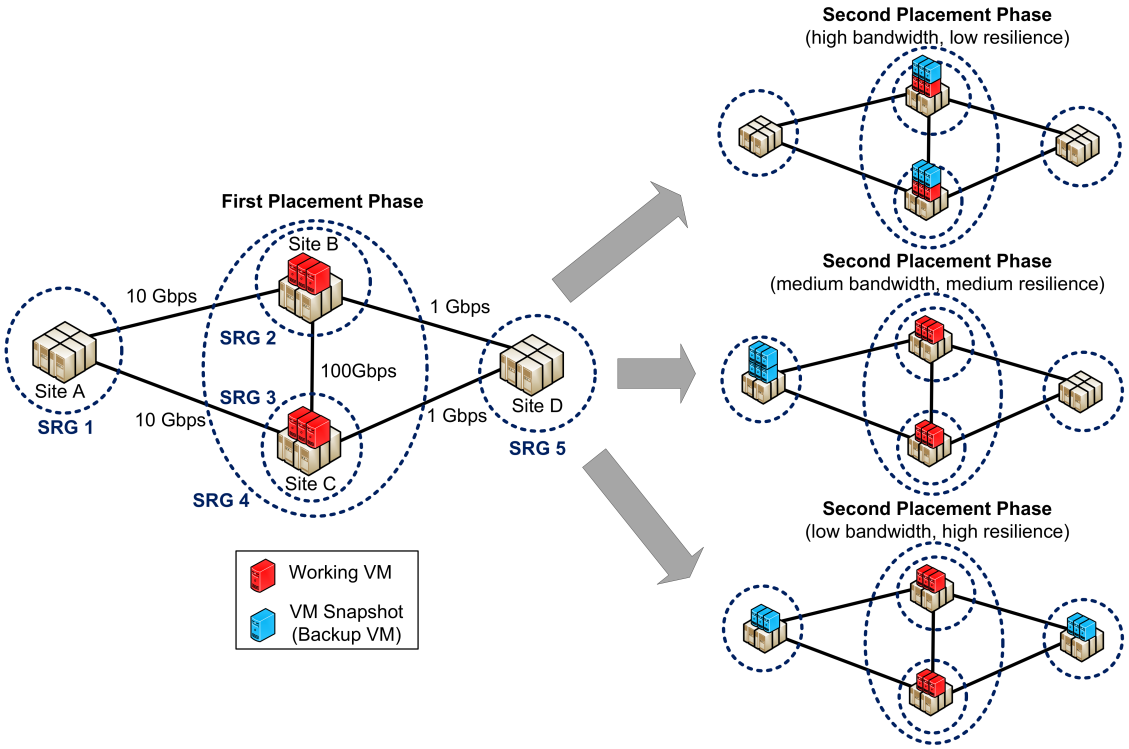

Fig. 4. Placement of working VMs and their snapshot locations.

backup of each VM. This placement must reduce the probability that the backup and working VMs are affected at the same time by disaster events. One approach is to isolate backup VMs, allocating them to sites belonging to SRGs different from the SRGs covering the sites hosting working VMs. This approach must consider the client QoR metrics, guaranteeing that the available resources are enough to meet the RTO and RPO requirements. Figure 4 shows three possible placements. Note that they have different bandwidth capacities between working sites and backup ones, which affects the RPO. In addition, they have different resilience levels depending on the number of SRGs spanned by backup sites, and the isolation in terms of SRG between working and backup sites.

III. Challenges and Research Directions

The design phases drawn in this work allow the identification of research directions in disaster resilience to IaaS Clouds. Although the "Selection of Disaster Recovery Mechanisms" phase is important to disaster resilience, it has a lot of intersections with other research areas such as high-availability (i.e., resilience to failures under provider's control), virtual machine mobility and network virtualization. On the other hand, the phases "Site Placement and Topology Design" and "Selection of the VM Placement Mechanisms" are the most challenging ones since they open a new research domain, which is the design of disaster-resilient cloud networks. Finally, the "Modeling" phase brings important challenges of stitching all DC components together, modeling their relationships with QoR and QoS metrics.

The "Site Placement and Topology Design" phase has the major challenge of jointly optimizing the inter-DC network design and the DC site placement. Works in this area investigate optimization algorithms to choose where to install each DC site, and in which sites each service is deployed as well as their backups [13], [7]. In addition, the current works design the underlying network between DC sites by configuring paths and setting protection schemes.

The current literature considers traditional DC distribution where services are replicated across a geo-distributed infrastructure, such as Content Delivery Networks (CDNs) [14], and assumes that each service is known at the time of DC construction. This assumption is not true in the IaaS case, since VM hosting demands are generally unknown a priori. Hence, the VM placement should be performed in a different phase, while the DC construction is based on the prediction of service demands. The state of the art addresses the service placement through the anycast principle. Hence, as the backups of each service are also operational, they can respond to requests. One drawback of the service replication performed in these works is the lack of backup synchronization among working copies, thus not considering RTO and RPO requirements. Regarding the "Selection of the VM Placement Mechanisms" phase, Bodík *et al.* [15] perform a resilient VM placement considering a single DC site and high availability requirements. Nevertheless, they do not consider geo-distributed DCs or backup and QoR to disasters.

IV. Conclusion

In this article, we have provided guidelines to design a DC network infrastructure supporting a disaster-resilient IaaS Cloud, based on the geographic redundancy of its components. We have described design phases, allowing us to draw potential research directions. In a nutshell, these directions concern the placement of nodes in a geo-distributed infrastructure, physically (e.g., DC sites) or virtually (e.g., VM snapshots), as well as how these nodes are interconnected through a WAN. We believe that the development of this new research area will allow IaaS providers to offer more sophisticated services, improving business continuity even when catastrophic events occur. Furthermore, a disaster-resilient Cloud motivates more corporations to migrate their IT infrastructure to an IaaS Cloud.

Acknowledgement

This work was partially supported by FAPERJ, CNPq, CAPES research agencies, and the Systematic FUI 15 RAVIR (http://www.ravir.io) project.

References

[1] E. Bauer and R. Adams, *Reliability and availability of cloud computing*. John Wiley & Sons, 2012.
[2] P. Cholda, J. Tapolcai, T. Cinkler, K. Wajda, and A. Jajszczyk, "Quality of resilience as a network reliability characterization tool," *IEEE Network*, vol. 23, no. 2, pp. 11–19, Mar. 2009.
[3] T. Wood, E. Cecchet, K. Ramakrishnan, P. Shenoy, J. Van der Merwe, and A. Venkataramani, "Disaster recovery as a cloud service: Economic benefits & deployment challenges," in *2nd USENIX Workshop on Hot Topics in Cloud Computing*, Jun. 2010.
[4] Cabinet Office, *ITIL®Service Design 2011 Edition*. TSO, 2008.
[5] "ISO/IEC 24762:2008, Information technology - Security techniques - Guidelines for information and communications technology disaster recovery services," ISO/IEC Standard, 2008.
[6] M. F. Habib, M. Tornatore, F. Dikbiyik, and B. Mukherjee, "Disaster survivability in optical communication networks," *Computer Communications*, vol. 36, no. 6, pp. 630–644, Mar. 2013.
[7] J. Xiao, H. Wen, B. Wu, X. Jiang, P.-H. Ho, and L. Zhang, "Joint design on DCN placement and survivable cloud service provision over all-optical mesh networks," *IEEE Transactions on Communications*, vol. 62, no. 1, pp. 235–245, Jan. 2014.
[8] W. D. Grover, *Mesh-based survivable transport networks: options and strategies for optical, MPLS, SONET and ATM networking*. Prentice Hall - PTR, 2004.
[9] R. Stankiewicz, P. Cholda, and A. Jajszczyk, "QoX: what is it really?" *IEEE Communications Magazine*, vol. 49, no. 4, pp. 148–158, Apr. 2011.
[10] M. Bari, R. Boutaba, R. Esteves, L. Granville, M. Podlesny, M. Rabbani, Q. Zhang, and M. Zhani, "Data center network virtualization: A survey," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 2, pp. 909–928, May 2013.
[11] R. S. Couto, M. E. M. Campista, and L. H. M. K. Costa, "A reliability analysis of datacenter topologies," in *IEEE GLOBECOM*, Dec. 2012, pp. 1890—1895.
[12] M. Pióro and D. Medhi, *Routing, flow, and capacity design in communication and computer networks*. Elsevier, 2004.

[13] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *Journal of Lightwave Technology*, vol. 30, no. 16, pp. 2563–2573, Aug. 2012.
[14] G. Pierre and M. van Steen, "Globule: a collaborative content delivery network," *IEEE Communications Magazine*, vol. 44, no. 8, pp. 127–133, Aug. 2006.
[15] P. Bodík, I. Menache, M. Chowdhury, P. Mani, D. A. Maltz, and I. Stoica, "Surviving failures in bandwidth-constrained datacenters," in *ACM SIGCOMM*, Aug. 2012, pp. 431–442.

Rodrigo S. Couto [S'11] (souza@gta.ufrj.br) received his *cum laude* Electronics and Computing Engineer degree from Universidade Federal do Rio de Janeiro (UFRJ) in 2011. Since October 2011 he is working toward a D.Sc. degree in Electrical Engineering from COPPE/UFRJ. Rodrigo did a one-year research internship (Oct/2012-Sep/2013) at LIP6, at the Université Pierre et Marie Currie (UPMC Paris VI). His research interests include cloud networks, network reliability and network virtualization. Rodrigo has been a member of IEEE Communications Society since 2011.

Stefano Secci [S'05-M'10] (stefano.secci@upmc.fr) is an associate professor at the Université Pierre et Marie Curie (UPMC Paris VI). He received a dual Ph.D. degree from the Politecnico di Milano and Telecom ParisTech. He covered positions also at NTNU, George Mason University, Fastweb Italia, and Ecole Polytechnique de Montral. His current research interests are about Internet resiliency and Cloud networking. He is vice-chair of the Internet Technical Committee, joint between the IEEE Communication Society and the Internet Society.

Miguel Elias M. Campista [S'05-M'10] (miguel@gta.ufrj.br) is an associate professor with Universidade Federal do Rio de Janeiro (UFRJ) since 2010. He received his Telecommunications Engineer degree from the Universidade Federal Fluminense in 2003 and his M.Sc. and D.Sc. degrees in Electrical Engineering from UFRJ, in 2005 and 2008, respectively. In 2012, Miguel has spent one year with at Université Pierre et Marie Curie, in France, as invited professor. His major research interests are in wireless networks, cloud computing, and social networks.

Luís Henrique M. K. Costa [S'99-M'01] (luish@gta.ufrj.br) received his electronics engineer and M.Sc. degrees in electrical engineering from Universidade Federal do Rio de Janeiro (UFRJ), respectively, and the Dr. degree from the Université Pierre et Marie Currie in 2001. Since August 2004 he has been an associate professor with COPPE/UFRJ. His major research interests are in the ares of routing, wireless networks, vehicular networks, and future Internet. Luís has been a member of the ACM and IEEE Communications Society since 2001.