Intercloud Security Considerations

David Bernstein
Huawei Technologies, USA
dbernstein@huawei.com

Deepak Vij
Huawei Technologies, USA
dvij@huawei.com

Abstract – **Cloud computing is a new design pattern for large, distributed datacenters. Service providers offering applications including search, email, and social networks have pioneered this specific to their application. Recently they have expanded offerings to include compute-related capabilities such as virtual machines, storage, and complete operating system services. The cloud computing design yields breakthroughs in geographical distribution, resource utilization efficiency, and infrastructure automation. These "public clouds" have been replicated by IT vendors for corporations to build "private clouds" of their own. Public and private clouds offer their end consumers a "pay as you go" model - a powerful shift for computing, towards a utility model like the electricity system, the telephone system, or more recently the Internet. However, unlike those utilities, clouds cannot yet federate and interoperate. Such federation is called the "Intercloud". Building the Intercloud is more than technical protocols. A blueprint for an Intercloud economy must be architected with a technically sound foundation and topology. As part of the overall Intercloud Topology, this paper builds on the technology foundation emerging for the Intercloud and specifically delves into details of Intercloud security considerations such as Trust Model, Identity and Access Management, governance considerations and so on.**

Keywords-component; Intercloud, Cloud Computing, Cloud Computing Security, Grid Security

I. INTRODUCTION

Cloud Computing has emerged recently as a new design pattern for a particular type of datacenter, or most commonly, a group of datacenters. Service providers offering applications including search, email, and social networks have pioneered this specific to their application. Recently they have expanded offerings to include compute-related capabilities such as virtual machines, storage, and complete operating system services.

Cloud Computing services as defined above are best exemplified by the Amazon Web Services (AWS) [1][2] or Google AppEngine [3][4]. Both of these systems exhibit all eight characteristics as detailed below. Various companies are beginning to offer similar services, such as the Microsoft Azure Service [5], and software companies such as VMware [6] and open source projects such as UCSB Eucalyptus [7][8] are creating software for building a cloud service.

For the purposes of this paper, we define Cloud Computing as a single logical datacenter which:

- May be hosted by anyone; an enterprise, a service provider, or a government.
- Implement a pool of computing resources and services which are shared amongst subscribers.
- Charge for resources and services using an "as used" metered and/or capacity based model.
- Are usually geographically distributed, in a manner which is transparent to the subscriber (unless they explicitly ask for visibility of that).
- Are automated in that the provisioning, upgrade, and configuration (and de-configuration and roll-back and un-provisioning) of resources and services occur on the "self service", usually programmatic request of the subscriber, occur in an automated way with no human operator assistance, and are delivered in one or two orders of seconds.
- Resources and services are delivered virtually, that is, although they may appear to be physical (servers, disks, network segments, etc) they are actually virtual implementations of those on an underlying physical infrastructure which the subscriber never sees.
- The physical infrastructure changes rarely. The virtually delivered resources and services are changing constantly.
- Resources and services may be of a physical metaphor (servers, disks, network segments, etc.; often called "Infrastructure as a Service" or IaaS) or they may be of an abstract metaphor (blob storage functions, message queue functions, email functions, multicast functions, all of which are accessed by running of code or script to a set of API's for these abstract services; often called "Platform as a Service" or PaaS). These may be intermixed.

The terms are well accepted now [9]. Use Cases and Scenarios for Cloud IaaS and PaaS interoperability

[10][11] have been detailed in the literature along with the challenges around actually implementing standards-based federation and hybrid clouds. The high level architecture for interoperability including a protocol suite and security approach was proposed where the term "Intercloud" was first coined [12].

Additional focus on security architecture was provided [13], and additional focus on how the overall architecture might be used to enable an exchange involving a marketplace was detailed and prototyped [14]. Overall Intercloud technical topology and protocol blueprints have been architected [15], and implementation approaches including presence and dialog, security approach, and semantic ontology model and directory, [16][17][18] have been defined. Finally, governance and market-making considerations have even been examined [19].

This paper briefly reviews this work and builds on that technology foundation, creating a business and services role definition for each element in the Intercloud topology. The paper goes on to describe the Intercloud security considerations and proposes a new Intercloud Trust Model, Identity and Access Management, Encryption and Key Management aspects and last but not the least the paper discusses the governance considerations of the overall Intercloud security environment.

II. INTERCLOUD TOPOLOGY

Cloud instances must be able to dialog with each other. One cloud must be able to find one or more other clouds, which for a particular interoperability scenario is ready, willing, and able to accept an interoperability transaction with and furthermore, exchanging whatever subscription or usage related information which might have been needed as a pre-cursor to the transaction. Thus, an Intercloud Protocol for presence and messaging needs to exist which can support the 1-to-1, 1-to-many, and many-to-many use cases. The discussion between clouds needs to encompass a variety of content, storage and computing resources.

The vision and topology for the Intercloud we will refer to is an analogy with the Internet itself: in a world of TCP/IP and the WWW, data is ubiquitous and interoperable in a network of networks known as the "Internet"; in a world of Cloud Computing, content, storage and computing is ubiquitous and interoperable in a network of Clouds known as the "Intercloud"; this is illustrated in Figure 1. The reference topology for realizing this vision is modeled after the public Internet infrastructure. Various providers will emerge in the enablement of the Intercloud. We first envision a community governed set of Intercloud Root providers who will act as brokers and host the Cloud Computing Resource Catalogs for the Intercloud computing resources, similar to DNS [20] would be utilized. One important difference for the cloud capabilities is that the root systems would be replicating and hierarchical, but would not replicate in a hierarchical fashion.

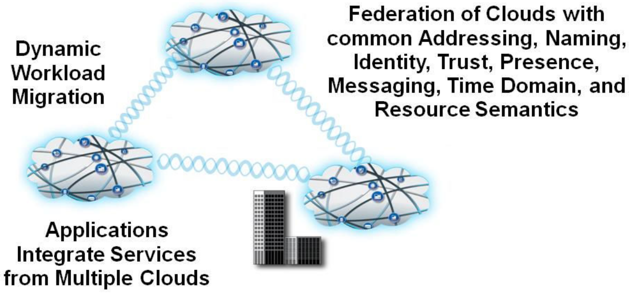
Dynamic Workload Migration
Federation of Clouds with common Addressing, Naming, Identity, Trust, Presence, Messaging, Time Domain, and Resource Semantics
Applications Integrate Services from Multiple Clouds

Figure 1. The Intercloud Vision

We propose that the roots replicate "sideways" and "upwards" using Peer to Peer technology [21] in order to scale. The sideways replication would be "master node" replication, as is common in P2P topologies, whereas the upwards replication would be to multiply interconnected peer replication, also as is common in P2P topologies.

The Intercloud Root instances will work with Intercloud Exchanges to solve the n^2 problem by facilitating as mediators for enabling connectivity among disparate cloud environments. This is a much preferred alternative to each cloud vendor establishing connectivity and collaboration among themselves (point-to-point), which would not scale physically or in a business sense.

Intercloud Exchange providers will facilitate the negotiation dialog and collaboration among disparate heterogeneous cloud environments, working in concert with Intercloud Root instances as described previously. Intercloud Root instances will host the root servers containing all presence information for Intercloud Root instances, Intercloud Exchange Instances, and Internet visible Intercloud capable Cloud instances. Intercloud Exchanges will host second-tier servers.

Individual Intercloud capable Clouds will communicate with each other, as clients, via the server environment hosted by Intercloud Roots and Intercloud Exchanges.

In order for the Intercloud capable Cloud instances to federate or otherwise interoperate resources, a Cloud Computing Resources Catalog system is necessary infrastructure. This catalog is the holistic and abstracted view of the computing resources across disparate cloud environments. Individual clouds will, in turn, will utilize this catalog in order to identify matching cloud resources by applying certain Preferences and

Constraints to the resources in the computing resources catalog. The technologies to use for this are based on the Semantic Web [22] which provides for a way to add "meaning and relatedness" to objects on the Web. To accomplish this, one defines a system for normalizing meaning across terminology, or Properties. This normalization is called an Ontology.

Due to the sheer size of global resources ontology information, a centralized approach for hosting the repository is not a viable solution due to the fact that one single entity can not be solely responsible and burdened with this humongous and globally dispersed task, single-point-of-failure, scalability and security ramifications, lack of autonomy as well as arguments related to trust and the authority on data. Instead, Intercloud Roots will host the globally dispersed computing resources catalog in a federated manner.

Intercloud Exchanges, in turn, will leverage the globally dispersed resources catalog information in order to match cloud resources by applying certain Preferences and Constraints to the resources. From overall topology perspectives, Intercloud Exchanges will provide processing nodes in a peer-to-peer manner on the lines of Distributed Hash Table (DHT) overlay based approach in order to facilitate optimized resources match-making queries. Ontology information in the DHT overlay nodes would be replicated from federated Intercloud Roots.

There has already been lot of work done on Semantic Peer-to-Peer based systems – GridVine[23], RDFPeers[24], Piazza[25], PIER[26], and "Distributed Overlay for Federation of Enterprise Clouds" [27].

All elements in the Intercloud topology contain some gateway capability analogous to an Internet Router, implementing Intercloud protocols in order to participate in Intercloud interoperability. We call these Intercloud Gateways. The entire topology is detailed in Figure 2.

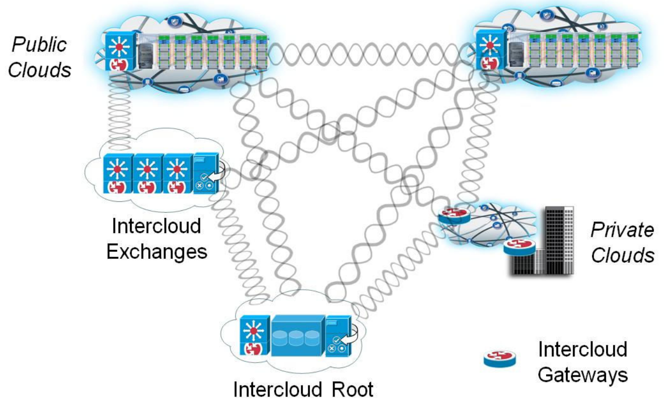

Figure 2. Reference Intercloud Topology and elements

The Intercloud Gateways would provide mechanism for supporting the entire profile of Intercloud protocols and standards. The Intercloud Root and Intercloud Exchanges would facilitate and mediate the initial Intercloud negotiating process among Clouds.

Once the initial negotiating process is completed, each of these Cloud instance would collaborate directly with each other via a protocol and transport appropriate for the interoperability action at hand; for example, a reliable protocol might be needed for transaction integrity, or a high speed streaming protocol might be needed optimized for data movement over a particular link.

III. INTERCLOUD TRUST MODEL

The diversity and flexibility of the capabilities envisioned by Intercloud enabled federated Cloud computing model, combined with the magnitudes and uncertainties of its components, pose difficult problems and challenges in effective provisioning and delivery of application services in an efficient and secured manner. Security is one of the most important and paramount elements of such a computing environment.

In an Intercloud cross-clouds federated environment, security concerns are even more important and complex. Intercloud paradigm or cloud computing paradigm, in general, will only be adopted by the users, if they are confident that their data and privacy are secured. Trust is one of the most fundamental means for improving security across heterogeneous independent cloud environments.

Currently, Public Key Infrastructure (PKI) based trust model is the most prevalent one. PKI trust model depends on a few leader nodes to secure the whole system. The leaders' validity certifications are signed by well established Certificate Authorities ("CA"s).

At a basic level, proposed Intercloud topology subscribes to the PKI based trust model. In accordance to the PKI trust model, the Intercloud Root systems will serve as a Trust Authority. In the currently proposed trust architecture, a Certificate issued by a Certificate Authority (CA) [28], must be utilized in the process to establish a trust chain. The CAs which provides certificates must provide them in specific formats, undergo annual security audits by certain types of accountancy corporations, and conform to a host of best practices known as Public Key Infrastructure [29]. These requirements can vary by country. The PKI best practices, the CA process, and the accountancy rules, need to be re-examined for cloud computing.

Certificates not only need to identify the clouds, but the resources the clouds offer, and the workloads that the cloud wishes federation with other clouds, to work upon. Where web sites are somewhat static, and a certificate can be generated to trust the identity of that web site, cloud objects such as resources and workloads are dynamic, and the certificates will have to be generated by a CA. As per the architecture of the CA, the Intercloud Exchange will need to be the intermediate CA, acting in a just-in-time fashion to provide limited lifetime trust to the transaction at hand.

The current PKI certificates based trust model is primarily all or nothing trust model and is unsuitable for Intercloud environment. According to the current PKI based trust model, once the CA authorizes the certificate for an entity, the entity is either trusted or non-trusted. This is more like a Boolean relationship. However, in the cloud computing environment, especially in the Intercloud environment, this model needs to be extended to have "Trust Index" to go along with the existing PKI based trust model. "Trust Index" is essentially a level of trust demonstrated by cloud providers. Depending on the level of trust (40%, 50%, 60%, or 100%), for example, one Intercloud provider might trust another provider to use its storage resources but not to execute programs using these resources. The *trust level* is specified within a given time because the *trust level* today between two entities is not necessarily the same *trust level* a year ago. Trust Level is something dynamic in nature as opposed to static PKI certificates.

From Intercloud topology perspectives, Intercloud Roots will provide static PKI CA root like functionality. On the other hand, Intercloud exchanges will be responsible for the dynamic "Trust Level" model layered on top of the PKI certificate based trust model. The overall trust model is more of a "Domain based Trust" model. It divides the cloud provider computing environment into several trust domains. Nodes in the same domain usually are much more familiar with each other, they have a higher degree of trust for each other.

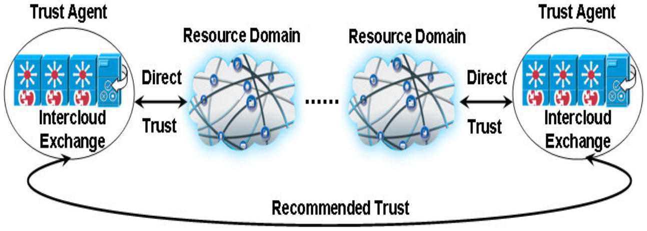

Figure 3: Intercloud Trust Management Model

Exchanges are the custodians/brokers of "Domain based Trust" systems environment for their affiliated cloud providers. Cloud providers rely on the Intercloud exchanges to manage trust. As Domain trust agents, Intercloud exchanges store other domains' trust information for inter-domain cooperation. Essentially, the trust information stored reflects trust value for a particular resource type (compute, storage etc.) for each domain. Exchanges also recommend other domains trust levels for the first time inter-domain interaction.

At a high level, we are working towards a trust algorithm framework in order to derive the "Trust Index" for a cloud provider. Essentially, the Intercloud Trust algorithm will evaluate the underlying security attributes of a cloud provider such as "Firewall Capabilities", "Intrusion Detection and Anti-Virus Capabilities" and so on. Additionally, cloud provider reputation parameters such as "Prior Success Rate", "Turnaround Time" and so on would be considered as part of the overall determination of "Trust Index". Accordingly, the fuzzy logic based aggregation algorithm will establish the "Trust Index" of a cloud provider.

Each Intercloud Exchange, as a Trust Agent, will discover the "Trust Index" of another cloud provider (via the corresponding Trust Agent) in a peer-to-peer manner on the lines of Distributed Hash Table (DHT) overlay based approach. The basic idea of DHT overlay system is to map a key space to a set of peers such that each peer is responsible for a given region of this space and storing data whose hash keys pertain to the peer's region. The advantage of such systems is their deterministic behavior and the fair balancing of load among the peers (assuming an appropriate hash function). Furthermore, they provide location transparency: queries can be issued at any peer without knowing the actual placement of the data.

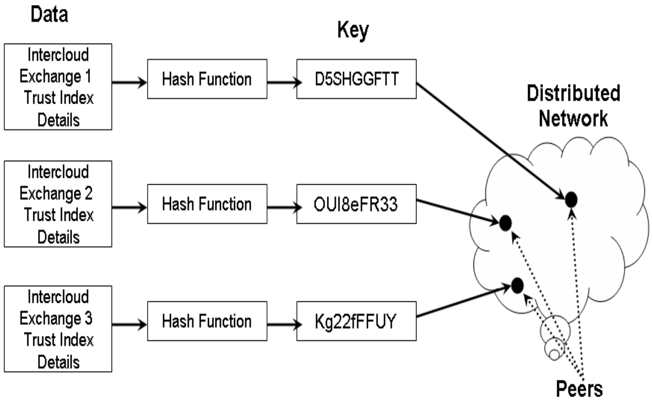
Data
Key
Intercloud Exchange 1 Trust Index Details
Hash Function
D5SHGGFTT
Intercloud Exchange 2 Trust Index Details
Hash Function
OUI8eFR33
Intercloud Exchange 3 Trust Index Details
Hash Function
Kg22fFFUY
Distributed Network
Peers

Figure 4: Distributed Hash Table.

The DHT peer-to-peer overlay is a self-organizing, distributed access structure, which associates logical peers representing the machines in the network with keys from a key space representing the underlying data structure. Each peer is responsible for some part of the overall key space and maintains additional routing information to forward queries to neighboring peers. As the number of machines taking part in the network and the amount of shared information evolve, peers opportunistically organize their routing tables according to a dynamic and distributed binary search tree.

The overall marketplace implications of this are quite interesting, in that the Intercloud Root looks a lot like a current Internet Root CA type of business, whereas the notion that the exchanges are also in the trust business as an adjunct to the actual exchange business.

IV. INTERCLOUD IDENTITY AND ACCESS MANAGEMENT

One of the key requirements to have success in effectively managing identities in the Intercloud environment is the presence and support for a robust standards based federated identity management capability using prevailing federation standards such as SAML [30], WS-Federation [31], and Liberty ID-FF [32]. Comprehensive Identity Management systems typically provide services such as: User Provisioning and User Management, Authentication and Authorization, Role Engineering, and Identity Data Integration/Virtualization.

In a typical federated identity model, in order for a cloud provider to establish secure communication with another cloud provider, it asks the trust provider service for a trust token. The trust provider service sends two copies of secret keys, the encrypted proof token of the trust service along with the encrypted requested token.

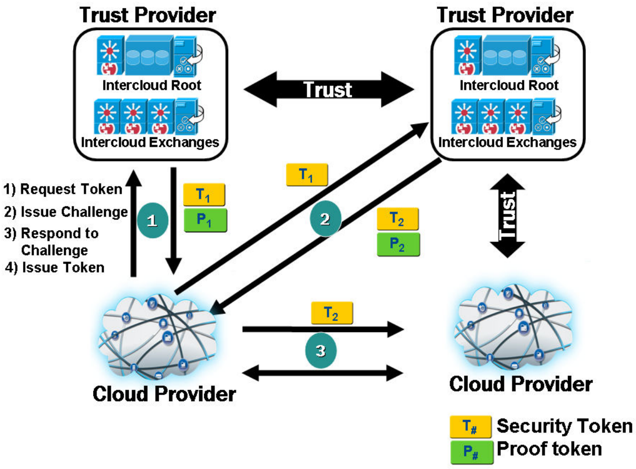

Figure 5: Intercloud Identity Federation Model.

As regards to granular level authorization in the Intercloud environment, support of XACML-compliant entitlement management is highly desirable. XACML [33] provides a standardized language and method of access control and policy enforcement. Currently, prevailing mechanism for granular level authorization is usually implemented in a proprietary non-standard fashion.

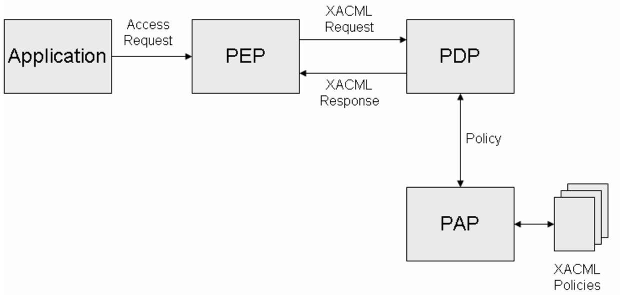

Figure 6: OASIS XACML Processing Environment.

XACML (eXtensible Access Control Markup Language) is an XML-based language for access control that has been standardized in OASIS. XACML describes both an access control policy language and a request/response language. The policy language is used to express access control policies (who can do what when). The request/response language expresses queries about whether a particular access should be allowed (requests) and describes answers to those queries (responses).

In a typical XACML usage scenario, a subject (e.g. human user, workstation) wants to take some action on a particular resource. The subject submits its query to the entity protecting the resource (e.g. filesystem, web server). This entity is called a Policy Enforcement Point (PEP).

The PEP forms a request (using the XACML request language) based on the attributes of the subject, action, resource, and other relevant information. The PEP then sends this request to a Policy Decision Point (PDP), which examines the request, retrieves policies (written in the XACML policy language) that are applicable to this request, and determines whether access should be granted according to the XACML rules for evaluating policies. That answer (expressed in the XACML response language) is returned to the PEP, which can then allow or deny access to the requester.

A Policy Administration Point (PAP) is used to get to the policies; the PDP uses the PAP where policies are authored and stored in an appropriate repository.

V. ENCRYPTION AND KEY MANAGEMENT

Encryption technology is a very key component of the overall Intercloud security framework. Security is designed in a manner so that data is encrypted “at rest” and “in transit”. In the Intercloud and cloud computing world in general, there is a radical paradigm shift specifically the way we think about computing by removing the specifics of location from its resources; cloud computing can be thought of as radical “deperimeterization”. However, in divorcing resources from location creates security issues that result from this lack of any perimeter. In such a world, there is an utmost need for securing the computing resources using strong encryption by leveraging underlying scalable and robust key management mechanism. However, encryption algorithms are as good as the underlying key management process.

Key management is not just about technology, it also includes People and Process elements as well. Failure or compromise of any of these components results in the failure or compromise of the whole system. Unlike traditional static internet environment, in an Intercloud environment there is a great need for separating the computing resources and the encryption keys, a chain of separation as well as a chain of custody with multiple parties involved at each step.

As computing resources in an Intercloud environment can potentially be anywhere. In order to be able to encrypt/decrypt these resources, the corresponding keys need to be retrieved. To help streamline the overall communication process between key management environment and cryptographic clients, we are evaluating interoperability standards such as recently announced OASIS Key Management Interoperability Protocol (KMIP [34]).

VI. Governance Considerations

When a business entrusts its data to a third party such as a cloud provider, it is vulnerable. Its data is sitting in cloud provider’s computing environment. In an Intercloud enabled federated cloud computing environment, it gets even more complex due to the involvement of more than one cloud provider. Many things can go wrong. The cloud service provider may go out of business or may decide to hold the data hostage if there is a dispute. It is important to understand in which country data will be hosted, because the location of the data directly affects the choice of the law that will govern the data.

If the data reside in China, it is likely that Chinese law will govern access to the servers where the data are hosted. If the client demands access to its data would Chinese law apply since the data are stored in China? Further, Chinese law may permit the Chinese Government to have unlimited access to the data stored in its territory whereas there might be stricter restrictions to access by the United States Government to data stored in the United States.

Controlling and governing who has access to the metadata associated with its data, or with the uses of its data, may be important. A company that holds sensitive personal data, company trade secrets, or other valuable information may wish to limit access to, or use of the traffic information associated with this data by the cloud service provider. For example, who looked at what information, and when or what queries or searches were run may have great value. The cloud service provider may want the ability to mine the company’s data or metadata for secondary uses, such as for marketing or market research purposes.

Numerous cloud service providers offer free access to their services or their applications with the view to mine the data in their custody in order to offer advertising services. In other cases, an organization or an individual may not mind the potential intrusion in their affairs if they determine that the financial benefit and ease of access to their information through the cloud outweighs the potential that third parties may access their files, pictures, or correspondence.

In an Intercloud federated environment, there are considerations and ramifications as far as prohibition again cross-border transfers of data assets. A global company that wishes to take advantage of cloud services will want to ensure that this use does not jeopardize its subsidiaries, clients, business partners and others which may be subject to foreign laws with different restrictions than those in effect in the United States. The US based company will want to know where the personal data of its employees, clients and others will be located, so that it can address the specific restrictions that foreign data protection laws may impose.

For example, a German subsidiary may not oppose the use of a cloud service provider in Argentina, but it will object to the transfer of its data to Turkey, Mexico, or the United States. Knowing where the cloud service provider will host the data is a prerequisite to implementing the required measures to ensure compliance with local laws that restrict the cross border flow of data.

Service providers will need a clear understanding of the complex restrictions and requirements created under the data protection laws of the European Union member states and of several other non-EU countries with similar laws. Cumbersome restrictions hamper the transfer of data outside of these countries. Their laws require data controllers (who originally collected the data) to inform individuals that their data will be processed abroad, and to obtain their consent to the transfer. In addition, the data controller and the recipient of the data may have to enter into special contracts that must be approved by the local Data Protection Authority.

VII. Related Work in the Grid Field

Similar work has been done in the Grid field which we are in the process of rationalizing the re-use of.

One particular area of interest is the Virtual Organization Management Service (VOMS) [35]. VOMS manages membership lists and roles for a virtual organization. It is more oriented towards users, rather than the clouds themselves, keeping track of each user and their roles, where users are described by their X.509 certificate's Distinguished Name property. We are investigating propagating this property into the cloud to cloud interoperability scheme we have proposed.

Another particular are of interest is the work of the International Grid Trust Federation (IGTF) [36]. The IGTF is a body to establish common policies and guidelines between its Policy Management Authorities (PMAs) members and to ensure compliance to this Federation Document amongst the participating PMAs. While the IGTF does not provide identity assertions, it maintains a list of trust anchors, root certificates and related meta-information for accredited authorities.

The Distribution contains Certificate Revocation List (CRL) locations, contact information, and signing policies. The IGTF consists of three member PMAs: the APGridPMA covering Asia and the Pacific, the EUGridPMA covering Europe, the Middle East and Africa, and TAGPMA covering Latin America, the Caribbean and North America. All registered members in each regional PMA are also members of the IGTF. These include identity providers, CAs, and their major relying parties, such as the various international grid projects.

We have approached IGTF to discuss leveraging them for Intercloud scheme we have proposed. It appears this organization may be a perfect answer to the governance and CA/key management challenges mentioned.

VIII. Conclusions

This paper reviewed the current state of the art in cloud computing federation, reviewing what has been come to be called the Intercloud at an overview technical level. Various aspects of security considerations were covered in the context of Intercloud federated environment.

The paper specifically proposes the new Intercloud "Trust Model" in conjunction with the prevalent PKI based "Trust Model". The paper also discussed security related governance considerations within an Intercloud computing environment.

We conclude with further refinement work needed in the area of Authorities and Governance, which we hope to leverage work from the Grid community.

Acknowledgements

We would like to thank Cloud Security Alliance (CSA) for their ongoing work on security guidance for critical areas of focus in cloud computing.

References

[1] Amazon Web Services, http://aws.amazon.com/
[2] Murty, James, Programming Amazon Web Services; S3, EC2, SQS, FPS, and SimpleDB, O'Reilly Press (2008).
[3] Google AppEngine, http://code.google.com/appengine/
[4] Ciurana, Eugene: Developing with Google App Engine, Firstpress (2009).
[5] Microsoft Azure, http://www.microsoft.com/azure/default.mspx
[6] VMware VCloud Initiative, http://www.vmware.com/technology/cloud-computing.html
[7] Nurmi D., Wolski R., Grzegorczyk C., Obertelli G., Soman S., Youseff L., Zagorodnov D.: The Eucalyptus Open-source Cloud-computing System, Proceedings of Cloud Computing and Its Applications, Chicago, Illinois (2008)
[8] Nurmi D., Wolski R., Grzegorczyk C., Obertelli G., Soman S., Youseff L., Zagorodnov D.: Eucalyptus: A Technical Report on an Elastic Utility Computing Architecture Linking Your Programs to Useful Systems, UCSB Computer Science Technical Report Number 2008-10 (2008)
[9] Youseff, L., Butrico, M. and Da Silva, D.: Toward a unified ontology of cloud computing, Proceedings of the GCE'08 Grid Computing Environments Workshop (2008).
[10] Mei, L., W.K. Chan, and Tse, T.H.: A Tale of Clouds: Paradigm Comparisons and Some Thoughts on Research Issues., APSCC pp.464-469, 2008 IEEE Asia-Pacific Services Computing Conference (2008).
[11] Cloud Computing Use Cases Google Group, http://groups.google.com/group/cloud-computing-use-cases , http://www.scribd.com/doc/18172802/Cloud-Computing-Use-Cases-Whitepaper
[12] Bernstein, D., Ludvigson, E., Sankar, K., Diamond, S., and Morrow, M., Blueprint for the Intercloud - Protocols and Formats for Cloud Computing Interoperability.. In Proceedings of ICIW '09, the Fourth International Conference on Internet and Web Applications and Services, pp. 328-336 (2009).
[13] Yildiz, M., Abawajy, J., Ercan, T., and Bernoth, A.: A Layered Security Approach for Cloud Computing Infrastructure, ISPAN, pp.763-767. In Proceedings of the 10th International Symposium on Pervasive Systems, Algorithms, and Networks (2009).
[14] Buyya, R., Pandey, S., and Vecchiola, C.: Cloudbus toolkit for market-oriented cloud computing. In Proceedings of the 1st International Conference on Cloud Computing (CloudCom) (2009).
[15] Bernstein, D.: The Intercloud: Cloud Interoperability at Internet Scale, In Proceedings of the 2009 NPC, pp. xiii (Keynote 2), Sixth IFIP International Conference on Network and Parallel Computing (2009).
[16] Bernstein, D., Vij, D.: Using XMPP as a transport in Intercloud Protocols In Proceedings of CloudComp 2010, the 2nd International Conference on Cloud Computing (2010).
[17] Bernstein, D., Vij, D.: Using Semantic Web Ontology for Intercloud Directories and Exchanges. In Proceedings of ICOMP'10, the 11th International Conference on Internet Computing (2010).
[18] Bernstein, D., Vij, D., Intercloud Directory and Exchange Protocol Detail using XMPP and RDF, Proceedings of IEEE 2010 International Workshop on Net-Centric Service Enterprises: Theory and Application (NCSE2010) (2010).

[19] Bernstein, D., Vij, D., An Intercloud Cloud Computing Economy - Technology, Governance, and Market Blueprints, Proceedings of ICSOC 2010, the 8th International Conference on Service Oriented Computing International Workshop on Net-Centric Service Enterprises (2010).

[20] Domain Names – Concepts and Facilities, and related other RFCs, http://www.ietf.org/rfc/rfc1034.txt

[21] Peer to Peer, http://en.wikipedia.org/wiki/Peer-to-peer

[22] W3C Semantic Web Activity, http://www.w3.org/2001/sw/

[23] Aberer, K., Cudr´e-Mauroux, P., Hauswirth, M., and Van Pelt, T.: GridVine: Building Internet-Scale Semantic Overlay Networks, International Semantic Web Conference, volume 3298 of Lecture Notes in Computer Science, pages 107–121. Springer (2004).

[24] Cai, M. and F. Martin.: RDFPeers: a scalable distributed RDF repository based on a structured peer-to-peer network., WWW '04: Proceedings of the 13th international conference on World Wide Web, pages 650–657, New York, NY, USA. ACM Press (2004).

[25] Tatarinov, I., Ives, Z., Madhavan, J., Halevy, A., Suciu, D., Dalvi, N., Dong, X., Kadiyska, Y., Miklau, G., and Mork, P., The Piazza Peer Data Management Project. SIGMOD Record, 32 (2003).

[26] Huebsch, R., Hellerstein, J., Lanham, N., Loo, B. T., Shenker , S., and Stoica, I., Querying the Internet with PIER., VLDB 2003, Proceedings of 29th International Conference on Very Large Data Bases, September 9-12, 2003, Berlin, Germany, pages 321–332. Morgan Kaufmann (2003).

[27] Ranjan, R., and Buyya, R., Distributed Overlay for Federation of Enterprise Clouds (2008).

[28] Certificate Authority, http://en.wikipedia.org/wiki/Certificate_authority

[29] Internet X.509 Public Key Infrastructure, Certificate Policy and Certification Practices Framework, http://tools.ietf.org/html/rfc3647

[30] Security Assertion Markup Language (SAML), http://saml.xml.org/saml-specifications

[31] Web Services Federation (WS-Federation), http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsfed

[32] Liberty ID-FF, http://projectliberty.org/liberty

[33] OASIS xEtensible Access Control Markup Language (XACML), http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml

[34] OASIS Key Management Interoperability Protocol (KMIP), http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=kmip

[35] VOMS Virtual Organization Management Service, http://vdt.cs.wisc.edu/components/voms.html

[36] International Grid Trust Federation (IGTF), http://www.igtf.net/