Cloud to Cloud Interoperability

Scott Dowell
Computer Sciences Corporation
San Diego, CA, USA.
sdowell@csc.com

Albert Barreto III, James Bret Michael, and Man-Tak Shing
Naval Postgraduate School
Monterey, CA, U.S.A.
{abarreto, bmichael, shing}@nps.edu

Abstract – *Cloud computing describes a new distributed computing paradigm that allows system of systems to access a shared pool of configurable computing resources (e.g., networks, servers, storage, data, applications, and services) that can be rapidly provisioned and released over the Internet with minimal user-management effort or cloud-provider interaction. Interoperability is central to enabling sharing of resources from a pool of cloud-service providers in a seamless fashion. In this paper we describe some of the challenges in achieving interoperability for cloud computing and recommend an adaptation of the U.S. Department of Defense's LISI Maturity Model to address cloud-to-cloud interoperability.*

Keywords: Cloud computing, interoperability, LISI

1 Introduction

Cloud computing is a catch-all phrase for referring to a distributed computing paradigm in which systems of systems access shared pools of configurable computing resources, treated as services, that can be rapidly provisioned and released over the Internet with minimal user-management effort or cloud-provider interaction [1]. Cloud computing holds the promise to revolutionize the way users collaborate over the Internet, similar to the way online shopping revolutionized the way books and music are distributed today, such as via the Apple iCloud. In [2], Foster *et al*. present a vision for moving toward the ultimate cloud-computing environment – a so-called Cloud Nirvana – where all unnecessary boundaries are removed in an evolutionary three-stage approach. Through the cloud, users may identify, select, and use the service that best delivers the given need at the time it is required. Cloud users must be able to move their personal property (e.g., data, applications, domain names, IP addresses) as well as runtime virtualized sessions across cloud provider systems and organizational boundaries in a seamless fashion. An enterprise will use services from different cloud platforms (both internal and external to the organization) for different applications to get the "best-of-breed." Cloud interoperability refers to the ease of migration and integration of applications and data between different providers' clouds.

As cloud computing matures, the ability to support interoperability becomes increasingly important. System-of-systems engineering will play a key role in determining how to provide for interoperability in cloud computing. A widely recognized model for system-of-systems interoperability is the Levels of Information System Interoperability (LISI) Maturity Model published by the Department of Defense (DoD) C4ISR Architecture Working Group [3]. LISI classifies the degree of sophistication with respect to exchanging and sharing information and services among systems in terms of PAID, an acronym for four closely interrelated attributes: Procedures, Applications, Infrastructure, and Data:

- The procedures (P) attribute reflects the degree of interoperability resulting from operational policies and processes, functional program development guidance, as well as compliance of technical and system architecture standards (e.g., hardware, system software, communications, data, and application standards).
- The application (A) attribute reflects the ability of the software applications to work on different systems and platforms as they progress through the interoperability maturity levels, ranging from stand-alone applications at the low end to applications that are designed for cross-discipline or cross-organizational boundaries at the high end.
- The infrastructure (I) attribute reflects the degree and form of connectivity between the systems and applications (e.g., point-to-point phone connection versus wide-area network across great variety of systems and communication protocols), and the way in which the systems interact with each other (e.g., application specific interface versus platform independent Web services).
- The data (D) attribute reflects the flexibility of the data format and the richness of the information being exchanged across systems and domains (ranging from files containing a single data type to integrated information space that supports all forms of data representation, presentation, and exploitation).

There are five levels of maturity in the LISI model:

- Level 0 – Isolated interoperability in a manual environment characterized by manual extraction and integration of data from multiple stand-alone systems.
- Level 1 – Connected interoperability in a peer-to-peer environment characterized by electronic connection,

separate data, separate applications, and homogeneous product exchange.

- Level 2 – Functional interoperability in a distributed environment characterized by local area networks, separate data, separate applications, heterogeneous product exchange, and basic collaboration.
- Level 3 – Domain-based interoperability in an integrated environment characterized by wide-area networks, shared data, separate applications, shared databases, and sophisticated collaboration.
- Level 4 – Enterprise-based interoperability in a universal environment characterized by wide-area networks, shared data, shared applications, cross-domain information sharing, and advanced collaborations.

The LISI model focuses on system-to-system information exchanges but falls short in providing a basis for assessing the maturity of cloud-to-cloud interoperability (C2CI). In particular, security and mobility across organizational boundaries and domains are important attributes that need to be considered when assessing the maturity level of C2CI, especially from a usability and acceptability-for-use perspective. This paper builds on the LISI model and presents a five-level maturity model for C2CI. The rest of the paper is organized as follows. Sections 2 and 3 present cloud interoperability challenges and a five-level model for C2CI. Section 4 discusses the use of the proposed model to assess C2CI maturity. Section 5 highlights current efforts to improve C2CI in the IT industry. Section 6 describes a Cloud Orchestration service for meeting the interoperability challenge and Section 7 presents a conclusion.

2 Cloud Interoperability Challenges

In this section, we discuss some of the challenges in making clouds interoperable.

2.1 Portability and Mobility

A question of interest to adopters of cloud computing is: "Can I deploy existing cloud artifacts (e.g., virtual server/desktop images, software applications, databases) on another service provider's services without modification to those artifacts?" In [4], Urquhart divides the image/application/data interoperability into two subcategories: portability and mobility. Urquhart defines portability as "the ability to move an image in a *down* state from one host to another, and then boot it at its destination." Mobility is defined as "the ability to move a live computer workload from one host to another without losing client connections or in-flight state." Portability/Mobility is a prime indicator of the degree of interoperability between clouds; mobility across cloud provider boundaries will be one of the targets of a mature interoperable cloud. It will require advancement in such things as open standards for virtual machine (VM) images and cloud-to-cloud application interfaces (APIs), as well as advancement in virtualization technologies to support the migration of live VM sessions, global IP addresses, and data services to static resources across cloud boundaries.

2.2 Cloud-Service Integration

In order to get the "best-of-breed" and to maintain control over the mission-critical operations and data, an enterprise may need to integrate both on-premise and the software-as-a-service (SaaS) applications to meet the business needs. The current practice of integrating software applications via an API requires a significant amount of coding as well as ongoing maintenance due to frequent modification and updates. Having both SaaS and on-premise applications interact via Web services and applying service-oriented architecture (SOA) principles to implement business logic via service composition can help solve the cloud-integration problem.

2.3 Security, Privacy and Trust

Cloud adopters also expect assurances from service providers that the provisioned services can be trusted to supply particular levels of security and privacy, such as controlling access by users via cloud services to personally identifiable information (PII) [5]. This will require effective solutions to the classical security problems that arise in multi-level security (MLS) and cross-domain systems, like federated identity management, active role-based access control (RBAC), as well as proper monitoring, logging and auditing as required by laws that govern data and application storage and usage and their movement across national, state and municipal boundaries.

With the potential to significantly reduce costs through consolidation and optimization of computing resources, cloud computing has introduced a unique set of security and privacy issues that must be addressed for cloud computing to be successful [6]. In addition, mature cloud computing environment providers will have to supply their customers with an appropriate level of security transparency to alleviate customers' reservations about the security and privacy afforded by the cloud [7]. At a minimum, cloud service providers should give their customers the same level of security assurance and transparency afforded by the non-cloud IT system of systems.

Security issues are strongly associated with administration of the cloud, including managing the users, resources and data, that are typically addressed via security policies for handling authentication, access control, session management and network communications. Migration from a legacy client-server model to a cloud-based model will mitigate some existing security issues while introducing new security issues. Failure to understand the new security issues or blindly attempting to apply legacy security policies and procedures for an enterprise's cloud migration will lead to problems. The Distributed Management Task Force (DMTF) Open Cloud Standards

Incubator Process and Deliverables model describes three components as work-in-progress [8]. Cloud security comprising management and control is among the three. As industry pushes more and more for cloud solutions, the gap between legacy policies (management) and procedures (control) and cloud security cannot be ignored. Having a well-defined cloud security policy is another indicator of the degree of interoperability between clouds. Hence, it will be necessary to have some agreed upon (and automated) means for objectively comparing the quality-of-security provided by one cloud service with that of another.

In [9], Nelson *et al*. pointed out that it is necessary to establish formal trust relationships between clouds for users to access and control remote resources across cloud boundaries. Sound methods for user authentication and authorization across cloud boundaries are needed for clouds to be interoperable. Currently, each enterprise and cloud provider has its own methods of proving identities and capabilities, and its own security policy labels that must be translated across organizational and cloud boundaries. Efficient and effective solutions to the domain composition problem will be another indication of mature cloud interoperability.

2.4 Management, Monitoring, and Audit

In addition to applying a single security and user-identity-management tool set to applications running on different cloud platforms, companies need a uniform tool set to automatically provision services, manage VM instances, and work with both cloud-based and enterprise-based applications.

Cloud users also need assurance that security and privacy policies are consistently applied and the service level agreements (SLA) are met as the cloud services migrate across cloud boundaries. Uniform processes and tool sets are needed to monitor and report the level of services and the compliance/violation of security/privacy policies in remote clouds.

3 The C2CI Maturity Model

In this section, we recommend some extensions to the PAID attributes and describe a new five-level model for assessing C2CI maturity.

3.1 Extensions to PAID

While retention of the PAID attributes is deemed appropriate as they apply to the cloud model and to legacy systems, we need to extend the meaning of the four attributes to address the interoperability concerns discussed in Section 2.

- The procedures (P) attribute will also reflect the availability of and adherence to uniform security and privacy policies and procedures that can be applied consistently across cloud boundaries, industrial standards for SLAs, standard procedures for cloud-services auditing, and technical and system architecture standards for cloud infrastructure and applications.
- The application (A) attribute will also reflect the ease of cloud-service integration, as well as the ability of the software applications to work and migrate seamlessly across cloud boundaries while maintaining the same quality-of-service (QoS) levels.
- The infrastructure (I) attribute will also reflect the degree of cloud mobility, availability of uniform tool sets for security (e.g., identity management), and cloud-services provisioning, management, monitoring, reporting and auditing.
- The data (D) attribute will also reflect the degree of the evolution from an application-centric to a data-centric view of information processing. Instead of today's artificial separation between data and applications, information in the cloud will be treated as artifacts, which are embodiments of data and their associated manipulators, mini programs that allow the user to process (e.g., view, edit, and print) the data [2]. Manipulators are dynamically configured and associated with an artifact, according to the artifact's state, and can provide access and security control.

3.2 A Five-Level C2CI Model

Since cloud computing builds on the premise that computing resources can be rapidly provisioned and released over the Internet, we can safely assume that, for any enterprise that is ready to migrate services to a cloud provider, the enterprise has surpassed levels 0 and 1 of the original LISI model and achieved the necessary networking and security maturity (e.g., protection of local area networks with firewalls and access control through local user authentication and file-access privileges) required to reach level 2. By removing Level 0, 1 and 2 from LISI, and adding three additional levels based on the degree of portability/mobility, security/privacy interoperability, ease of integration, and the availability of standard management, monitoring and audit procedures and tools, we maintain the components of LISI that are applicable to the cloud model while adding the appropriate levels necessary for evaluating C2CI. The proposed C2CI model consists of the following levels:

- Level 0 – Domain-based interoperability in an integrated environment characterized by wide-area networks, shared data, separate applications, shared databases, and sophisticated collaboration. Cloud services are confined to single provider clouds.
- Level 1 – Enterprise-based interoperability in a universal environment characterized by wide-area networks, shared data, shared applications, cross-domain information sharing, and advanced collaborations via the inter-cloud Web services.
- Level 2 – Portability interoperability in a public, private, or hybrid cloud environment where cloud

artifacts may traverse multiple providers in down states. Inter-cloud enforcement of security and privacy policies and SLA are based on pairwise agreements.

- Level 3 – Security interoperability in a public, private, or hybrid cloud environment where policies and procedures from one cloud provider will interact with other policies and procedures with other cloud provider(s) transparently and automatically using standardized protocols and cloud-wide formal trust relationships.
- Level 4 – Mobile interoperability in a public, private, or hybrid cloud environment where cloud artifacts may traverse multiple providers in in-flight states. There is no artificial separation between data and applications. Data in the cloud can be shared and manipulated by multiple applications on multiple platforms.

4 Applying the C2CI Model

With the recent announcement that the Pentagon has chosen Defense Information Systems Agency (DISA) as its top cloud computing supplier, and that DISA will not compete with private industry in the Army's bid for a cloud provider, it is becoming clear that a single DoD Cloud may not be possible [10]. Clearly, the C2CI model's three new levels address concerns such as in-flight migration from a DISA cloud to a non-DISA cloud, where a VM will have to somehow deal with a new IP space, perhaps with some kind of Network Address Translation (NAT) wrapper which allows it to communicate with the new IP space transparently as well as not violate security policies as it leaves one cloud, and is compliant as it settles into the new one. This level of interoperability is only attainable when Level 4 maturity of the C2CI model is achieved by all clouds participating in the exchange. The hybrid nature of a DoD and Army cloud would be able to successfully allow for this type of in-flight migration if compliant. Examining the PAID Infrastructure attribute in this scenario brings to light the connectivity and interactive nature of the two systems. Moving between IP spaces in an in-flight migration exposes protocol concerns such as IPv4 versus IPv6 addressing, QoS requirements as well as application and/or web services compatibility.

In a case where a VM might traverse clouds in a down state, and enter into a running state once on the destination cloud, the machine must be able to undergo a forensic analysis before it is allowed to communicate on the network and be available to the end-user. If DISA does become the main DoD cloud provider, and the Army does choose its own cloud provider, essentially creating a hybrid DoD cloud, a shared set of security policies and procedures (management and control) must be agreed upon and enforced across these two providers. If all clouds within the hybrid relationship have achieved C2CI Level 2 maturity, the forensic analysis required for down-state migration would be possible. Here, the PAID Procedures attribute highlights the importance of policies and procedures for maturity of systems and adherence to technical and system architecture standards for software, hardware, and applications. Additionally, the PAID Applications attribute reveals the need for interoperability of applications across different systems and platforms.

5 Current Industrial Efforts

In an enterprise of the size and scope of the DoD that has approximately 3.5 million users of unclassified systems, the ability to leverage infrastructure, platform, and software services is integral to the success of efficiency initiatives such as data-center consolidation. At the core of interoperability is workload portability and automated provisioning to dynamically adjust the cloud environment to the needs of the user of the service. With that said, interoperability extends to SLAs, service-level objectives (SLOs), QoS, accounting, and billing aspects of the IT architecture. Some of the working groups working in this arena are the Cloud Security Alliance (CSA), Cloud Computing Interoperability Forum (CCIF), the Open Grid Forum/ Open Cloud Computing Interface Working Group, and the DMTF incubator. Industry is addressing these requirements through cloud service architecture, cloud management tools, and image management.

1. Cloud Service Architecture – applications require resiliency to changes in connectivity and location. Interoperability includes the ability to support live-motion of an application, that is, retain a consistent state by maintaining the connections and contracts required for the application to remain viable. The National Institute of Standards and Technology (NIST) Cloud Computing Reference Architecture and Taxonomy Working Group is leading interested US Government agencies and industry partners to define a neutral cloud computing reference architecture and taxonomy to better understand various cloud services in the context of an overall Cloud Computing model, with the aim of using the reference architecture and taxonomy as a tool to communicate and analyze proposed standards for cloud security, interoperability, and portability and the reference implementations of the standards [11]. Today, cloud service architecture remains relatively immature. Most projects are focused on virtualization in private cloud scenarios with few enterprise-wide, multi-cloud environments working in an interoperable manner.
2. Cloud Infrastructure Management – defines the APIs to support the management and control of multiple cloud environments in public, private, hybrid, and community models. The management determines how services and data are shared between providers to include the starting and stopping of VMs and storage manipulation. Today, Cloud Infrastructure Management supports most of the characteristics of C2CI Level 2.

3. Image Management – defines how to provision cloud services and data on multiple hosts without changes. Common capabilities include the use of a service catalog and business rules reflecting the SLA and security levels. Images can be managed in a federated manner across the enterprise.

It is not uncommon to find organizations focused on standards are addressing infrastructure as a service (IaaS) and establishing a common API set and definitions. Examples of this approach include the Open Virtualization Format (OVF), which supports VM interoperability, and Cloud Data Management Interface (CDMI), which aims to standardize interfaces for cloud-based storage. Through OVF, VMs can be packaged and distributed across cloud environments and multiple hypervisors (e.g., Hyper-V, VMware, Xen). While useful, it represents C2CI Level 2.

Public providers attempt to attain Federal Information Security Management Act (FISMA) certification to an appropriate security level; however they provide limited visibility and transparency to support governance and security requirements of C2CI Level 3 and above.

The Federal Government and DoD are attempting to address C2CI maturity as well. The US Federal CIO Council launched the government-wide Federal Risk and Authorization Management Program (FedRAMP) in 2009 to provide a standard approach to accessing, authorizing and continuous monitoring of cloud computing services and products for all federal agencies. The National Institute of Standards and Technology (NIST) also launched the Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC) initiative in 2010 to develop and maintain a set of cloud system use cases through an open and ongoing process engaging industry, other Government agencies, and academia.

The bottom line: many proposed standards have yet to be fully explored or widely implemented and remain unproven. This void allows leading providers such as Amazon Web Services (AWS) APIs to be considered de facto standards.

6 Cloud Orchestration Service

Achieving semantic and syntactic interoperability is a challenge within the world of cloud services. Integral to C2CI is the ability to apply the right cloud service and data to solve a given business need at the time, cost, quality, and security level required. This ability is commonly referred to as cloud orchestration.

NIST explains that cloud orchestration is "the arrangement, coordination and management of cloud infrastructure to provide services to meet IT and business requirements." To accomplish cloud orchestration a "cloud broker" serves to intermediate, aggregate, and arbitrage services on behalf of the cloud consumer. In [12], NIST describes the broker function as the following:

- Intermediate: A cloud broker enhances a given service by improving some specific capability and provides the value-added service to cloud consumers.
- Aggregate: A cloud broker combines and integrates multiple services into one or more new services. The broker will provide data integration and ensure the secure movement of data among cloud consumer and multiple cloud providers.
- Arbitrage: Service arbitrage is similar to service aggregation, with the difference being that the services being aggregated are not fixed. Service arbitrage allows flexible and opportunistic choices for the broker. For example, the cloud broker can use a credit scoring service and select the best score from multiple scoring agencies.

Figure 1 shows the orchestration service that governs the provision and execution of cloud services across multiple domains (public, private, hybrid, internal, external).

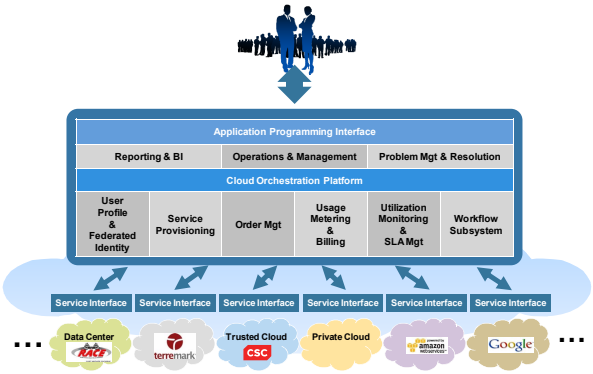

Figure 1. Cloud Orchestration Service

To achieve cloud orchestration, the cloud consumers require visibility and control across multiple domains to securely use cloud services that meet their policies and SLAs. To accomplish this, orchestration services require an integrated service catalog, provisioning, monitoring, and billing processes that are standardized across multiple IT services. Key to cloud orchestration is the ability to create trust in the services delivered by providers using governance, compliance, and regulation of the cloud services. This is accomplished through a common protocol that describes the service with such information as SLA and security profiles. Such a protocol can leverage the NIST Secure Content Automation Protocol (SCAP) to prescribe the key elements for cloud-trust. Current work on a cloud-trust protocol indicates that a common protocol to facilitate interoperability may be accomplished with as few as twenty-three elements. The open-source software community is aiding the advancement of cloud interoperability via adoption of the OpenStack and Cloud Foundry projects.

The cloud orchestration approach is essential to achieving C2CI. Today, organizations must use a combination of tools (e.g., enStratus, Tivoli, BMC, Puppet, Citrix, VMware) to form an initial C2CI capability. The

service orchestration lacks the policy-based automation of the cloud vision. One option to consider is establishing a common reference architecture and then leveraging the open-source movement with the cloud trust protocol as a way to achieve C2CI in the future.

7 Conclusion

This paper examines the characteristics of a mature interoperable cloud and presents a five-level model for assessing C2CI maturity. With the ongoing standards efforts in industry and government, we believe that it will not take too long to achieve Level 2: portability interoperability in a public, private, or hybrid cloud environment. However, major technological breakthroughs are needed to support the attainment of Levels 3 and 4.

As part of embracing cloud computing, enterprises themselves need to go through a cultural transformation to be able to implement the cloud-services paradigm. Most business units have been used to "infrastructure or application hugging." Sharing "their" infrastructure or application with others requires the building of trust and confidence in their IT organizations' ability to manage shared resources effectively and efficiently. This is also the case where services are procured through public cloud service providers. This trust and confidence has to deliver cost savings, agility, adequate security and other factors that the business cares about in this new cloud-computing environment. Furthermore, many enterprises have a track record of implementing "silo" computing environments. This silo mentality can result in multiple cloud environments (e.g., by business units, customers, applications) where there is little sharing of resources. Also, the "just in case" capacity planning mentality that has led to over-provisioning of resources should be supplanted by a systematic demand-management and capacity-planning process to avoid resource pools remaining underutilized and, thus, obviating the benefits that a cloud-computing environment provides.

Acknowledgement

The research was supported in part by a grant from the Office of the DoD Chief Information Officer. The views expressed in this report are those of the authors and do not reflect the official policy or position of the Computer Science Corporation, Department of Defense or the U.S. Government.

References

[1] P. Mell and T. Grance, *NIST Working Definition of Cloud Computing*, Version 15, 7 October 2009. Accessed on Jan 3, 2011: http://www.nist.gov/itl/cloud/upload/cloud-def-v15.pdf

[2] K. Foster, J. Shea, D. Drusinsky, J.B. Michael, T. Otani and M. Shing, "Removing the Boundaries: Steps toward a Cloud Nirvana," *Proc. 2010 IEEE International Conference on Granular Computing*, Silicon Valley, CA, 14-16 August, 2010, pp. 167-171.

[3] C4ISR Architecture Working Group Interoperability Panel, *Levels of Information Systems Interoperability (LISI)*, Department of Defense, Washington, DC, 30 March 1998.

[4] J. Urquhart, "Exploring cloud interoperability, part 2," news.cnet.com, May 7, 2009. Accessed on March 8, 2010: http://news.cnet.com/8301-19413_3-10235492-240.html

[5] M. Zhou, R. Zhang, W. Xie, W. Qian, and A. Zhou, "Security and Privacy in Cloud Computing: A Survey," *Proc. 6th International Conf. on Semantics Knowledge and Grid (SKG)*, Ningbo, China, 1-3 Nov. 2010, pp. 105-112.

[6] H. Takabi, J. Joshi, and G. Ahn, "Security and Privacy Challenges in Cloud Computing Environments," *Security & Privacy*, IEEE, Vol. 8, No. 6, Nov/Dec 2010, pp. 24-31.

[7] B. Michael and G. Dinolt, "Establishing Trust in Cloud Computing, *IAnewsletter*, Vol. 13, No. 2, pp. 4-8, 2010.

[8] DMTF, "Interoperable Clouds, A White Paper from the Open Cloud Standards Incubator". Accessed on April 7 2011:http://www.dmtf.org/sites/default/files/standards/documents/DSP-IS0101_1.0.0.pdf

[9] A.J. Nelson, G.W. Dinolt, J.B. Michael, and M.T. Shing, "A Security and Usability Perspective of Cloud File Systems," *Proc. 6th IEEE International System of Systems Engineering Conference*, Albuquerque, NM, 27-30 June 2011.

[10] B. Brewin, "DISA poised to become Pentagon's top cloud computing supplier," NextGov.com, January 3, 2011. Accessed on June 4, 2011: http://www.nextgov.com/nextgov/ng_20110103_7911.php?oref=topstory

[11] National Institute of Standards and Technology, NIST Reference Architecture & Taxonomy Working Group Draft Charter, January 7, 2011. Accessed on April 10, 2011: http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/ReferenceArchitectureTaxonomy/NIST_CCRATWG_001_Charter_010711.pdf

[12] National Institute of Standards and Technology, NIST Cloud Computing Reference Architecture, Version 1, March 30, 2011. Accessed on April 10, 2011: http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/ReferenceArchitectureTaxonomy/NIST_CC_Reference_Architecture_v1_March_30_2011.pdf