International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 2 Issue 4, April - 2013

Cloud Computing & Intrusion Tolerance With SOA

Harshal S. Deshpande Girish V. Patil Mahesh V. Shastri
Computer Engineering Dept.Computer Engineering Dept.Computer Engineering Dept.
Pdm. Dr. V. B. Kolte Pdm. Dr. V. B. Kolte Pdm. Dr. V. B. Kolte
College of Engineering College of Engineering College of Engineering
Malkapur, India. Malkapur, India. Malkapur, India.

Abstract

cloud computing and semantic web are challenging some of the assumptions made in the existing designs of intrusion tolerant systems. This paper provides an analysis of the changing landscape, describes the newly introduced risks and vulnerabilities, and briefly outlines research efforts that may point the way forward.

I. INTRODUCTION

Experience shows that attacks may never be completely prevented, and some attacks may not be detected accurately and on time. Consequently, intrusion tolerance, combining aspects of protection, detection and reaction, is currently considered the optimal way to address information security challenges. However, the architecture of intrusion-tolerant systems, integrating multiple layers of defenses, redundancy and diversity can be daunting, and is often viewed as heavyweight, costly to provision and difficult to dynamically re-provision. At the same time, the information technology landscape has been evolving with the introduction of new software technologies such as cloud computing [1], SOA [2] and Semantic Web [3].The new technologies present an opportunity. For example, cloud computing can reduce a lot of provisioning issues, and enable "on-click" dynamic provisioning of computing power and storage. The SOA concept implies that software building blocks, including security mechanisms, can now be thought of as services, potentially developed independently, to be connected to a service bus. Semantic Web envisions many of the tasks that require human comprehension of disparate data available in the network to be done by automated processing agents. Combining SOA and cloud has the potential to make intrusion tolerant architectures affordable in the same way safe deposit boxes in banks (instead of vaults in individual homes) made safe storage of valuables affordable. Similarly, semantic linking of disparate data can unlock inferences leading to new heights of cyber-defense situation awareness. However, indiscriminate migration to SOA and cloud computing (the "Someone Else's DataCenter" phenomenon) can also be potentially dangerous. In addition to compute power, storage orconnectivity, the cloud must offer a level of trust and protection. In SOA, the services must include security aspects in their service-level agreements in addition to "functionality" or "logic". But developing cloud or SOA services with customizable levels of security and trust is no different from developing trustworthy and secure computer programs—a problem that has not been solved completely yet.

II. EMERGING TECHNOLOGIES

Intrusion tolerant versions of distributed systems of various flavors (e.g., thin client, 3 tier, distributed objects, peer to peer, publish-subscribe) that are based on a vertical ownership structure, where a single organization has control over the software application, the CPU and memory resources it requires to run, as well as the access points for remote interactions, have been developed and experimented with [4, 5,6]. The tolerance of such systems is derived from the

protection, detection and redundancy mechanisms integrated into the vertical silos, controlled air-gapped communication among them, and adaptive management of the resulting defense-enabled silos. A typical example is shown in Figure 1, where Widgets' service is made available in the Internet via content delivery mechanisms such as Akamai. There is only one "cloud" in this scenario—the network. From the perspective of Widgets' customers, Widgets' services are available from the network cloud, whereas from Widgets' own perspective, the network cloud is a combination of its intranet (leased lines or tunnels through the public Internet connecting Widgets' corporate and partner sites) and the Internet (where Widgets' customers are). Widgets and its partner organizations can be expected to have multiple layers of defense to protect their own enclaves.

Cloud computing and SOA introduce a different kind of structure (see Figure 2). The "cloud" is not confined to the "network" anymore. Some of the software and storage that were on Widgets' corporate and partner sites will now be hosted in the cloud (e.g., Amazon's data centers). Instead of tunneling through the public Internet, Widgets and its partners can obtain high bandwidth connectivity from network service providers (e.g, Verizon) to link their premises to the cloud data centers. Providers like Amazon and Verizon can cater to many organizations like Widgets and its partners at the same time and possibly sharing the same resources creating horizontal layers that collect or co-locate communication,storage and computation from multiple sources. Widgets' customers on the other hand, will continue to view the network cloud as the source of Widgets' services.

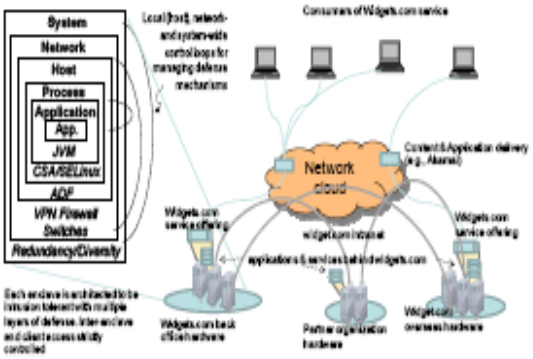

Figure 1: A networked distributed system

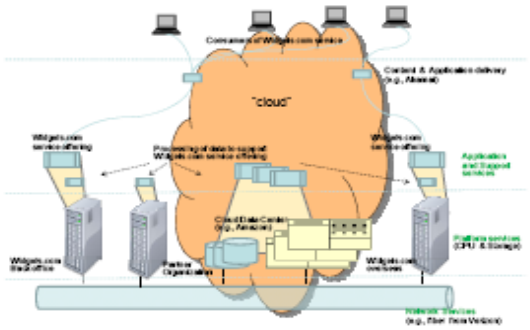

Figure 2: A system in a Cloud-SOA setting

The cloud data center or platform services providers offer services to start, advertise and connect hosted services to end consumers, migrating or load-balancing hosted services as necessary, and once again with certain properties (e.g., maintaining a standby, migrating or adding new instances if load increases etc.). Organizations like Widgets obviously need to worry about applications: buy vs. build, how to organize available building block services etc. In addition, they also need to worry about who accesses their data and computation hosted in the cloud, whether information exchanged within the cloud (data center or the network) are exposed to unauthorized entities or tampered during transit, how to trust the services building blocks found in the cloud, what level of QoS to negotiate with service providers (e.g., platform or network services providers)etc. Figure 3 illustrates the utility of semantic web technologies. Deriving answers to questions like the one posed there requires human interpretation of the data and services that are available in the network cloud. With the semantic web technology, automated agents can scour the network chasing semantic links to find the answer.

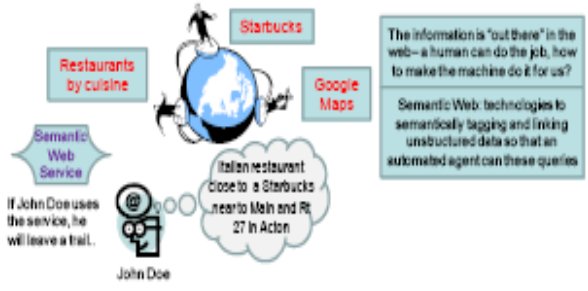

Figure 3: Example use of Semantic Web Technology

The confluence of cloud computing and SOA actually facilitates semantic linking and advanced data mining. In SOA, some services and information must be externalized (e.g., service description and discovery), some service transactions may leave a visible trace as they cross organizational boundaries, and furthermore, the information externalized this way may already be structured and tagged.

III. ANALYSIS OF SECURITY ISSUES

While the horizontal "services" stove-piping the "cloud" can be constructed to offer certain levels of security, we still need to worry about end-to-end security. For example, a cloud data center may offer storage or computing service with 99.9% availability, or the global information grid (GIG) [7] may offer core communication services with strong authentication and access control. But this security covers the interface between the "cloud" and its consumers (e.g., organizations like Widgets); end-users' interactions such as Widgets' customers logging in and using widgets.com are not covered, even though parts of the end-users' requests get processed in the cloud. Even from the perspective of an organization like Widgets, not everything is rosy and peachy- while it is easier to encrypt data to be stored in the cloud, no such technology exists to "encrypt" the computation that is delegated to the cloud. Semantic linking, and subsequent crawling and mining of such linked information and services may lead to information tied to the identity of individuals that the individuals and organizations may not want to share (i.e., violation of privacy). For instance, in the example shown in Figure 3, it is possible to track John Doe's eating habits by following the trail left by his use of the semantic web service (more damaging scenarios follow the same pattern of this benign example). It is not clear what an adversary, empowered with semantically linked data about the system, can do to an intrusion tolerant system that uses SOA and delegates some of its storage and computation to the cloud. It turns out that the introduction of SOA, cloud and semantic web technology can make some aspects of intrusion tolerance easier to realize.

A. Things likely to get better

1. Defense in depth**:** In a SOA-cloud setting, availability, confidentiality, integrity and access control can be embedded in each service layer imposing separation of concern and facilitating defense in depth. In this structure, network experts will worry about the network and platform experts will worry about storage and CPU availability (separation of concern).Systems configured by orchestrating services and resources with built-in security value-add will inherently include multiple independent layers of defense and containment boundaries.

2. Access control for resources:The SOA-cloud setting will enforce a level of access control to system resources and services that are not available today. Similar features at platform services providers will extend the scope of control to CPU usage and storage as well making certain kinds of denial of service attacks that plague the Internet today more difficult. Authentication and access control for individual services and resources also help building up system-wide defense in depth.

3. Reasoning about incident reports:Adoption of semantic web technology will enable semantic linkage and development of intelligent query processing capabilities.

B. Things likely to remain the same

1. Validation and trust:We argue that validating security claims, especially quantitative evaluation of security, will be at least as difficult as it is today in a SOA-cloud-semantic web setting. Separation of concern may help in constructing assurance cases, but this will be counterbalanced by the difficulty in evaluating the security claims made by the cloud services*2.*

*2.*Accountability*:*Accountability obviously is very useful as deterrence for insider threat as well as post-incident forensics. Execution of tasks that are internal to one organization today can span multiple organizations in SOA-cloud setting.

C. Things that need innovative solutions

1. Data protection:Today it is the data owner who accepts the terms and conditions of the cloud storage (e.g., when one uploads an album to Snapfish or Facebook). The data owner has no control over what a friend, who is authorized to access the photographs, does after he copies them. What can a platform service provider do to offer a confidentiality value-add ?In addition to loss of confidentiality, which is essentially

about data, semantic linking and data-mining that take advantage of such linkage will give rise to privacy issues , which is essentially about individuals.

2. Services management:In the SOA-cloud setting, a system is a collection of cooperating services including the cloud services (e.g., the network or platform services offering connectivity, CPU or storage), application services (implementing the business logic) and support services (providing among others, security functions). We argue that a specialized support service—the "services management" or SM service—will be needed to ensure end-to-end security and service delivery requirements.

3. Regulatory Issues:Suppose a terrorist organization buys a guaranteed service and uses encrypted communication between ingress A and egress B- the network operator will only have access to encrypted data, which is not helpful for prosecution. Similarly, a terrorist organization can store their secret information in the cloud in encrypted form. Law enforcement has already encountered similar issues with VOIP and peer-to-peer networks, despite the existence of laws like the Communications Assistance for Law Enforcement Act (CALEA).

D. An Emerging Opportunity

With two-way smart metering and intelligent devices in every home and distributed generation involving a larger percentage of green sources that are inherently unpredictable, electric grids of the future will become very large distributed interdependent cyber-physical systems requiring sophisticated algorithms processing huge amounts of data collected throughout the system that range from billing information and consumers' usage patterns to the internal state of generating stations and transmission lines and pricing data from energy market and carbon markets. And as recent news reports [12] indicate, it will also become an attractive target for cyber attacks. Various utilities and system operators have already embarked upon grid modernizing efforts. Many have adopted SOA for their advanced control center applications that obtain data and interact with each other by connecting to an enterprise service bus (ESB). In many cases telecom providers and new band-width-on demand (BoD) services connect control centers and other key elements—much like a cloud.

IV. SOLUTION APPROACHES

In this section we will briefly describe some work currently being done by us and other researchers that are relevant and may point the way forward.

1. Service-oriented security: Emerging standards and COTS products seem to exhibit an "everything is a service" theme. Some defenses that are typically part of an application will become externalized and shared in a SOA setting.

2. Trust and assurance: We have begun working on a framework of indicators from which it is possible to assess the assurance level of a system from various stakeholder perspectives. The indicators cover a range of static and organization-level aspects both internal and external to the system, as well as a number of dynamic properties of the system. The assessment is not in terms of absolute quantification; rather it provides a way to order various configurations of values and observations from the indicators in terms of the assurance concerns of a given stakeholder.

3. Data and information protection: Work in digital object identifier (DOI) system has developed a formalism to represent data stored in digital media as digital objects with unique identifier and associated metadata. We are exploring the possibility of encoding authentication and access control policies in a mark-up language, storing the policies with the digital objects, and enforcing them at the point of use.

V. CONCLUSIONS

SOA, cloud services and semantic web are three examples of emerging technologies that have the potential to alter the way survivable systems will be built in future. We showed where the existing intrusion tolerance technologies can help (e.g., supporting defense in depth), where they fall short (e.g., data protection and dynamic management of security), and also described promising lines of research that can help fill the gap.We argued that the emerging technologies will provide an opportunity to apply the existing intrusion tolerant technologies to a wider set of applications because they make provisioning and re-provisioning network, CPU and memory resources easier and more dynamic.

VI. REFERENCES

[1] Vaquero, L. M., Rodero-Merino, L., Caceres, J., and Lindner, M. A break in the clouds: towards a cloud definition. SIGCOMM Comput.Commun. Rev. 39, 1 (Dec. 2008), 50-55.
[2] Erl, T. Service-oriented Architecture: Concepts, Technology, and Design. Upper Saddle River: Prentice Hall PTR. 2005.
[3] Berners-Lee, T., Hendler, J., and Lassila, O. The Semantic Web. Scientific American Magazine.May 17, 2001.
[4] Wang, H., Liu, P., and Li, L. Evaluating the survivability of Intrusion Tolerant Database systems and the impact of intrusion detection deficiencies. Int. J. Inf. Comput. Secur. 1, 3 (Jun. 2007), 315-340.
[5] Valdes, A., Almgren, M., Cheung, S., Deswarte, Y., Dutertre, B., Levy, J., Saïdi, H., Stavridou, V., and Uribe, T. E. Dependable Intrusion Tolerance: Technology Demo. DARPA Information Survivability Conference and Exposition - Volume II, 2003
[6] Chong, J., Pal, P., Atighetchi, M., Rubel, P., Webber, F. Survivability Architecture of a Mission Critical System: The DPASA Example.ACSAC 2005: 495-504
[7] http://en.wikipedia.org/wiki/Global_Information_Grid
[8] Rushby, J. Design and Verification of Secure Systems. Proc. 8th ACM Symposium on Operating System Principles: 12–21, 1981
[9] http://altornetworks.com/products/vnf/
[10] https://www.trustedcomputinggroup.org/groups/
[11] Kissner, L., and Song, D. Privacy-preserving Set Operations. Advances in Cryptology, 2005.
[12] Wall Street Journal, Electriciy Grid in U.S. Penetrated by Spies (April2009): http://online.wsj.com/article/SB123914805204099085.html
[13] Trustworthy Cyber Infrastructure for the Power Grid (TCIP) home page:http://www.iti.illinois.edu/content/tcip-trustworthy-cyber-nfrastructurepower-grid
[14] CRitical UTility InfrastructurAL resilience (CRUTIAL) project home page: http://crutial.cesiricerca.it/