What Next in Intrusion Tolerance

Partha Pal, Rick Schantz, Michael Atighetchi, Joseph Loyall, Franklin Webber
BBN Technologies, Cambridge, MA
{ppal, rschantz, matighet, jloyall, fwebber}@bbn.com

Abstract

Emerging software technologies such as SOA, 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.

1. 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 Data Center" phenomenon) can also be potentially dangerous. In addition to compute power, storage or connectivity, 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. In addition, combination of SOA and cloud computing may unleash new security threats, and the power of semantic linking will make controlling access to information more difficult, threatening privacy of individuals and information owners. Unless these issues are well understood, and intrusion tolerance technologies are adapted to the new environment, new features and capabilities may have shorter time to market, but information systems of the future will become more vulnerable, and may actually fare worse against attackers than today's intrusion tolerant systems.

In this paper, we present our analysis of the potential impact SOA, cloud and semantic web technologies on intrusion tolerance. Our conclusions can be summarized as follows. A number of defenses and security techniques, especially those providing availability, integrity and confidentiality, can possibly be encapsulated in the cloud or within the services, and offered as value-add; but new capabilities (e.g., the ability to dynamically manage the security aspects of SOA services and cloud resources or support for privacy preserving interaction) will also be needed. Furthermore, some of the current security and intrusion tolerance challenges are likely to remain problematic, and may even be exacerbated, creating additional difficulties for law-enforcement in some cases.

2. 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.

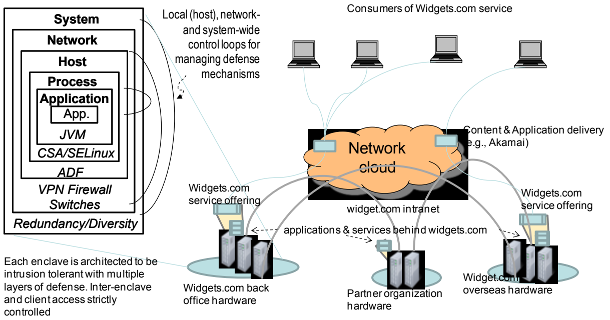
System
Network
Host
Process
Application
App.
JVM
CSA/SELinux
ADF
VPN Firewall Switches
Redundancy/Diversity
Local (host), network- and system-wide control loops for managing defense mechanisms
Each enclave is architected to be intrusion tolerant with multiple layers of defense. Inter-enclave and client access strictly controlled
Consumers of Widgets.com service
Network cloud
Content & Application delivery (e.g., Akamai)
Widgets.com service offering
widget.com intranet
Widgets.com service offering
applications & services behind widgets.com
Widgets.com back office hardware
Partner organization hardware
Widget.com overseas hardware

Figure 1: A networked distributed system

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.

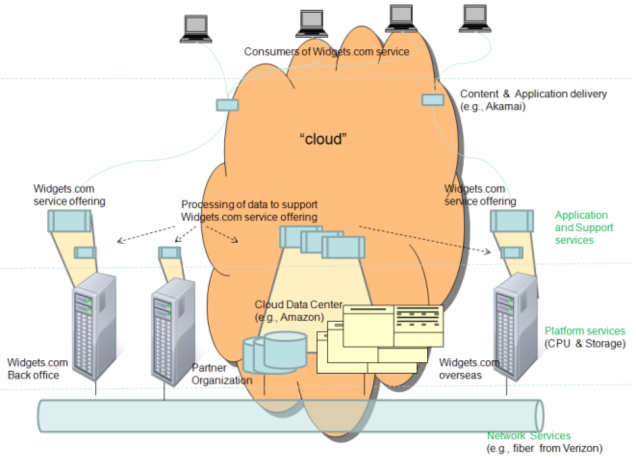
Consumers of Widgets.com service
Content & Application delivery (e.g., Akamai)
"cloud"
Widgets.com service offering
Processing of data to support Widgets.com service offering
Widgets.com service offering
Application and Support services
Cloud Data Center (e.g., Amazon)
Platform services (CPU & Storage)
Widgets.com Back office
Partner Organization
Widgets.com overseas
Network Services (e.g., fiber from Verizon)

Figure 2: A system in a Cloud-SOA setting

In this hybrid vision of cloud and SOA, the *network services providers* offer the service to establish communication paths that deliver bits from ingress to egress with certain properties (i.e., at a certain data rate -- guaranteed or best effort, unmodified, despite network failures etc). 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 semantic web technology, automated agents can scour the network chasing semantic links to find the answer. 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.

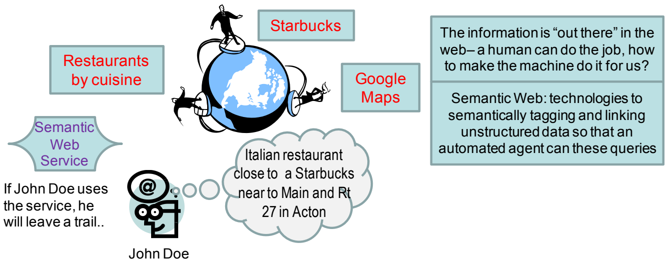
Starbucks
Restaurants by cuisine
Google Maps
The information is "out there" in the web– a human can do the job, how to make the machine do it for us?
Semantic Web: technologies to semantically tagging and linking unstructured data so that an automated agent can these queries
Semantic Web Service
If John Doe uses the service, he will leave a trail..
Italian restaurant close to a Starbucks near to Main and Rt 27 in Acton
John Doe

Figure 3: Example use of Semantic Web Technology

3. 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. At the same time, some current issues are likely to remain as problematic as it is today. But more importantly, we observe that the emerging software technologies will introduce additional complications and new challenges in a number of areas. We present some examples of each kind in the next three subsections.

3.1. Things likely to get better

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.

Availability mechanisms (e.g., redundancy and diversity) and redundancy-based protocols for checking or tolerating process integrity (e.g., Byzantine fault tolerant protocols) can be incorporated into the network services layer as well as the platform services layer. Mechanisms contained in a layer can also help contain the impact of failures within the layer (in order avoid violation of service level agreements). In this set up, different levels of diversity or redundancy independently for network and platform resources can be allocated for a given system configuration, and the allocation can be dynamically adjusted (i.e., adding or releasing bandwidth/CPU/memory) easily and nearly instantly based on urgency and cost considerations.

Existence of a network services layer can facilitate multiple layers of integrity protection for messages in transit. The network services layer can guarantee that what it received at ingress is not tampered before it exits at the egress, and at the same time, the messages can still be signed at the application layer. This will be similar to what we did in our prior survivable system work [6], but with the following differences. In prior work, the survivability architects needed to put in place network-level mechanisms at each of the communicating enclaves (recall, they owned the entire vertical), but in the SOA-cloud setting the capability can be bought for a price from the network provider. This commoditization also makes it easy to turn the network level mechanisms off or on dynamically.

Like availability and integrity, some aspects of confidentiality, in particular confidentiality of data can also be commoditized: the network or the cloud storage can offer encryption/obfuscation/isolation value-add that one can buy or dynamically negotiate.

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. Because unauthorized access can damage their revenue stream, the cloud vendors will aggressively control access and usage of their services and resources. For instance, a network services provider selling network bandwidth between two locations to various organizations can be expected to require the devices attaching to its network at these locations to authenticate themselves. In addition, it will likely enforce strict flow control to prevent bandwidth over-consumption, and block traffic types that were not contracted for. 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.

Reasoning about incident reports: Adoption of semantic web technology will enable semantic linkage and development of intelligent query processing capabilities. Such capabilities are believed to be helpful for managing and using vast amount of unstructured and distributed data. Applying semantic web technologies to the large volume of alerts and incident reports collected in intrusion tolerant systems today can therefore lead to improved cyber-defense situation awareness.

3.2. Things likely to remain the same

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.

How much trust should be placed on the services procured from the cloud? Do you want to actively assess the level of assurance provided by the platform or the network services layer? What are the mechanisms? How trustworthy is the assessment mechanism itself? Analogous questions arise today in peer-to-peer systems in the Internet, and it is likely that reputation or behavior based approaches used in the peer-to-peer context will be applicable in the SOA-cloud-semantic web setting as well.

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. Different organizations may monitor and track what they are interested in leaving holes in the end-to-end data and control flow. The audit trails may be incompatible with each other. But we argue that a similar situation exist today in systems and applications that involve multiple security domains.

3.3. Things that need innovative solutions

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. Clearly, this model will not work when the data owner has its own authentication and access control policy.

Even though breach of confidentiality is hard to detect, data in the cloud can be stored in encrypted form. But for computation delegated to the cloud, there is no such analog. This makes the computing processes in the cloud a weaker point in the data processing chain. 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. For instance, someone's social security number can be in the public—it is just a 9 digit number, but as soon as it is linked with an individual, it becomes private information. To preserve privacy, it is not the data itself, but the association of the data with an identity that needs to remain confidential—but it is not clear today how to control semantic linking or who has what rights to that link.

A related case can be demonstrated using "friend of a friend" or FOAF. On one hand, FOAF provides a level of access control: Alice, being a FOAF of Bob, obtains certain rights to access Bob's information. On the other hand, FOAF is often used as an indicator of trust, e.g., when Bob links his data with data that is linked with Alice (or a fixed number of indirections away from Alice) but not with data linked with Charlie, who is not a FOAF. In a semantically linked universe, FOAF is but one kind of semantic link that will be automatically formed, followed and mined by semantic web applications—however, mechanisms for 3-way authentication (between the data owners and the application exploiting the semantic link) and access control of the semantic links are not readily available yet. Meaningful auditing of semantic link chasing may even be very difficult (consider the difficulty of tracing back an attacker who hides behind a sequence of relays). It is likely that there will a need for 3rd party services for tracking and verification (like credit reporting agencies).

Finally, the discussion about data protection in SOA-cloud setting will remain incomplete without talking about metadata. Metadata forms the basis on which service consumers find services it needs. Unless there is a tamperproof way to associate metadata with services, and verify that association, new kinds of "phishing" attacks will appear.

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. We envision the SM service as a middleware service that takes requirements from the application owners (e.g., Widget) and provisions network, CPU and storage resources from network and platform services providers with appropriate level of security, and controls the security posture of all participating services.

The ease of dynamic re-provisioning in the SOA-cloud setting will open up the possibility of using more resources on demand—based on load (need to serve more requests) or threat (my services are being attacked, need additional redundant servers). The envisioned SM service should be able to react to security breaches and monitored load, and dynamically re-provision affected resources. However adding new resources is not free of side effects: it may temporarily suspend the ability to deliver service or impact the quality of service; the security posture of the services needs to be re-aligned. Therefore the SM service will need to dynamically coordinate with all the participating services as well.

It is easy to detect unavailability and integrity breaches. The service providers may be able to contain and mitigate the failures within their layers, but in cases when that is not possible or the mitigation is not sufficient meet the end-to-end security and service delivery requirement of the system, a system-wide response will be necessary. For example, Widget may decide to provision connectivity from a second network or platform services provider to maintain availability of its system to end customers while its primary providers work on restoring their services. The SM service is the natural choice for handling the coordination required for such system-wide actions. Obviously, the SM application itself needs to be intrusion tolerant.

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). How to distinguish between a terrorist taking advantage of the services and legitimate privacy and confidentiality requirements (e.g., Alice may rent cloud storage to safely store her medical or financial records, or establish a secure link between her and her doctor)? Issues like this will be at the forefront when SOA systems in the cloud will manage semantically linked information. To satisfactorily address the forensics, auditing and provenance requirements existing rules and regulations need to be revisited along with development of new technology hooks.

3.4. 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. New requirements such as the owner of a plug-in hybrid vehicle (PHV) being billed for charging the car in a public car park as well as a friend's house (charging at his home is no different from another household device) present the need for novel semantic linking of data. Overall, it is very likely that the electricity grid will undergo a level of disruptive transformation in the next 5 to 10 year period. Yet, at the same time, the grid must maintain the highest level of reliability. Under the current environment, maintaining reliability means it must operate through cyber attacks and provide continued service which may degrade initially, but need to recover quickly.

Keeping that requirement in mind, there is a recent surge of research and development activity both in the US [13] and EU [14] in the area of resiliency and protection of critical infrastructure such as the power grid. We argue that the electric grid will be an interesting *proving ground* where some of the issues we describe in this paper will be encountered, and hopefully, addressed.

In many ways, the transformation we anticipate in the power grid is reminiscent of the early days of the Internet. The ability to connect a large number of devices distributed over a large distributed area enabling bi-directional communication will lead to new applications and use cases, which in turn will lead to new requirements. However, networking and distributed system construction technology has progressed, and newer techniques such as SOA, cloud computing and semantic web are well poised to make these efforts comparatively easier and the evolution faster this time. It would be nice to be able to say the same thing with similar level of confidence about intrusion tolerance and security of the energy grids of the future.

4. 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.

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. For instance, instead of having their own internal authentication mechanisms, an Oracle database server and a JBOSS application server can share an authentication service in the cloud. But there are cases where such externalization will be risky and inefficient. For example, for session level encryption (i.e., after a consumer's session is established), it does not make sense for the end points to go to a third party service to encrypt their messages. Apart from inefficiency, this will raise the issue of trusting the encryption service. As part of our continued work on survivable systems, we have begun exploring the limits of externalizing defense mechanisms and developing the specialized SM middleware service we described earlier. In the SM service work we are leveraging our prior work in adaptive redundancy and multi-level resource management implemented as a middleware service as well.

A key point of concern in service-oriented security is deciding what security and survivability function needs to be a service. To illustrate, consider the issue of service corruption. In a traditional setting, a *voting protocol* among replicas can be used for this purpose. Should the voting algorithm be implemented as a *voting service* in a SOA setting? Which implementation choice offers better security: a distributed protocol where protocol libraries are embedded in each participant; or a voting service implementation where the participants interact with it? The latter can be a single point of failure, and will introduce another new service request-response interaction that needs to be protected.

Multiple Independent Levels of Security (MILS): Isolation and containment is a basic design principle of intrusion tolerant architectures. MILS [8] aim to provide a non-bypassable, evaluable, always-invoked, and tamperproof architecture, where components of various levels of trustworthiness can coexist. It is foreseeable that future cloud services will be constructed based on MILS. The quality of the MILS architecture can be used to gain confidence of the service consumers (for example, a vendor may claim that their separation kernel has a mean time t compromise of 7 days). Similarly, if the "service management" task can monitor the tamperproof mechanisms in the architecture, it can direct defensive responses to prevent further damage. New "separation mechanisms" will be needed as cloud providers seek to maximize their return on investment by using virtualization, new types of coding and multiplexing schemes to increase their resource utilization both in the network and platform services. Virtual machine firewalling techniques such as the one described in [9] can be useful in this context if organizations choose to delegate their applications as VMs, or the cloud provider sells computing resources as VMs. Virtualization at network devices such as routers, as well as novel networking technologies such as dynamic wave-division multiplexing (WDM) optical circuits over fiber backbone networks can provide multiple independent security in the network.

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. Trusted computing initiative [10] is another promising line of work that is extremely relevant in this context. Services built around trusted platform monitor (TPM) can be leveraged to assess whether a computation task can be handed off to platform resources in the cloud. We have also started looking at developing such services and facilitating user mode programs to safely use such services.

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. This technique will take advantage of semantic linking, availability of fast network and computing resources: request for a digital object will fetch the XML metadata to be process first; metadata processing may involve fetching DTD schema from remote sites, credentials checking, and producing a cryptographic code to unlock the actual data.

There has been quite a bit of mathematical work in privacy preserving computation [11], but developing a real world application as a privacy-preserving computation is impractically complex. We have started exploring simpler and more practical techniques such as substituting real data by a virtual "use once data" and hiding in the crowd to preserve a level of privacy. Such techniques will also leverage existing fast networks and semantic linking to establish and check association between the virtual "use once data" and its real counterpart, and also to formulate and accommodate the load of artificial "crowd" transactions.

5. 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. On the other hand, there are cases where the current SOA, cloud and semantic web offering has inherent vulnerability that can be exploited by a malicious adversary. Now is also the time to develop technologies to address them and infuse the emerging technologies with the appropriate security and survivability value-adds. We also argued that the electricity grid will provide a fertile ground to study some of the issues and validate some of the solution approaches we discussed here.

Finally, more awareness of the new challenges that arise at the confluence of SOA, cloud and Semantic web is clearly needed. New law enforcement requirements such as CALEA, HIPPA etc. further complicate the space of technical solutions. Un-informed adaptation of new technologies can be fatal and expensive (penalties, liabilities). We anticipate that innovative approaches to privacy, access control, accountability and trust management will be necessary to address these challenges.

6. 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 (April 8, 2009): 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-infrastructure-power-grid

[14] CRitical UTility InfrastructurAL resilience (CRUTIAL) project home page: http://crutial.cesiricerca.it/