Research Report

The Side Effects of Secure Web and Cloud Service Usage

NOWADAYS, ORGANIZATIONS ARE INCREASINGLY RELYING ON THIRD PARTIES TO SUPPLY VITAL IT SERVICES, WITH THE COMMUNICATION BEING PREDOMINATELY CONDUCTED VIA THE INTERNET. THIS, HOWEVER, POTENTIALLY EXHIBITS SENSITIVE BUSINESS INFORMATION TO OUTSIDERS. IN THIS ARTICLE, WE SHOW WHICH UNDESIRED SIDE EFFECTS STATE-OF-THE ART COUNTER MEASURES MAY HAVE.

Ulrich Lampe
Ralf Steinmetz
André Miede

Introduction

Over the last decade, service-orientation has been one of the major trends in the IT industry. Furthermore, it has inspired a variety of novel architectural paradigms, most notably, Service-oriented Architectures (SOA) and cloud computing. The general idea is to provide certain, often vital business functionality as a service from both internal and external providers, e.g., using Web services as a prominent and widely adopted implementation technology. Today, this development has cumulated in the popular and successful idea of providing *"Everything as a Service"* (XaaS), which in turn is one of the essential characteristics of cloud computing.

For the financial industry, this service-orientation is an important IT design principle and provides multiple benefits. First, it eases the integration of internal legacy IT systems through standardized interfaces. Second, it enables the seamless cooperation with external parties, e.g., which may be able to provide certain services substantially cheaper by exploiting specialization and economies of scale. In the latter case, the communication between a service user, such as a financial institution, and an external service provider is often conducted via public networks, most notably the Internet. To ensure typical security requirements such as confidentiality, the contents of the transferred messages will be encrypted, e.g., as recommended by the Federal Office for Information Security (2009).

However, an attacker – for instance, a competitor or a foreign government – can still relatively easily observe the fact that a communication process is conducted between the service user and service provider. Based on this information, an attacker may, for instance, identify optimal times for targeted Denial of Service attacks (e.g., busiest hours of the day), thus increasing potential damages. The attacker may also deduce business-relevant information about the service user (i.e., the financial institution), such as the success rate of certain transactions. On the other hand, attackers may gather detailed information on the service provider's business, e.g., its customer base, usage patterns, peak hours, and so on.

All this is possible without looking at the contents of the (encrypted) messages but only by monitoring communication relationships, e.g., between a financial institution and its service providers. Thus, a security goal that is commonly referred to as *relationship anonymity* is seriously threatened (Miede et al., 2011).

Fortunately, the research community has been concerned with this problem in different scenarios for many years. The combined efforts have resulted in so-called *anonymity systems*, which permit to obfuscate communication relations in public networks. The principle idea is to employ randomly selected relay nodes for the transfer of a message, which exacerbates eavesdropping and identifying communication relationships substantially. The principle is illustrated in Figure 1, where "U" denotes the service user and "P" denotes the service provider; "B" and "C" are randomly selected relay nodes in the anonymity system. "A" and "D" denote other possible nodes of the anonymity system that could have been chosen.

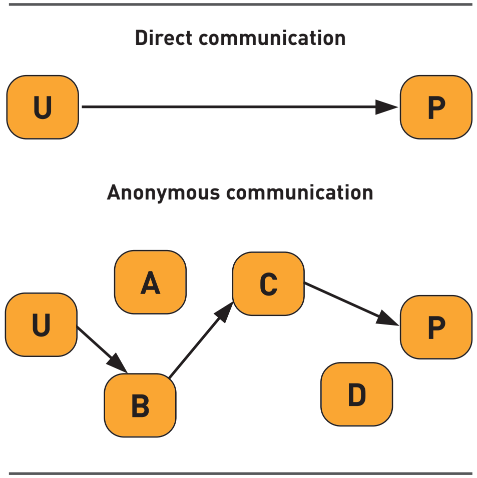

Figure 1: Direct vs. anonymous communication

Unfortunately, using such anonymity systems has side effects as well: through the introduction of relay nodes, the path of a message through the network may be significantly prolonged. This may inflict service consumption with respect to different Quality of Service (QoS) parameters, such as response time or throughput. Aside from the functional properties of a service, these non-functional QoS characteristics are of uttermost importance in the realization of business processes.

Based on these observations, we aim to answer the following research question: *What is the impact of the use of anonymity systems on the QoS of IT service consumption?*

Measurement Approach

In order to empirically examine the aforementioned research question, we have conducted a set of large-scale measurements. The principle aim was to quantify the impact of anonymity systems on the QoS of service consumption. In our experiments, we focused on the common and business-crucial parameters of response time, throughput, and availability.

To conduct our experiments, we implemented a reference service based on the popular Web services standards. This service was deployed on 12 globally distributed nodes, which represent a realistic set of worldwide service providers, based on studies on Web service provider distribution.

Over the course of four weeks, we conducted approximately two million invocations of the reference service. For the invocations, we used well-proven, state-of-the-art anonymity systems, i.e., JonDo and Tor. The first is a commercial system, where certified providers charge for data transfer, whereas the latter is operated by voluntary participants around the globe. Furthermore, we used the common direct, non-anonymized mode of access for reference purposes.

Our approach resembles the actual service consumption through a service user. Because the measurements were conducted using real networks, the results can be immediately transferred to practice.

Empirical Findings

The detailed results of our measurements are depicted in Figures 2, 3, and 4. In accordance with expectations, the application of anonymity systems for the secure usage of external services does have an effect on all regarded QoS parameters, i.e., response time, availability, and throughput.

To begin with, the effect is most notable with respect to the *response time* of external services (Figure 2). On average across all globally distributed providers, the response time using a non-anonymized, direct access corresponds to about 200 ms. Depending on the geographical location of the servers, the observed individual values range between about 25 ms (provider in Germany) and 370 ms (New Zealand). Please note that "US-1" and "US-2" denote two individual services in different parts of the United States, which explains the difference in response time.

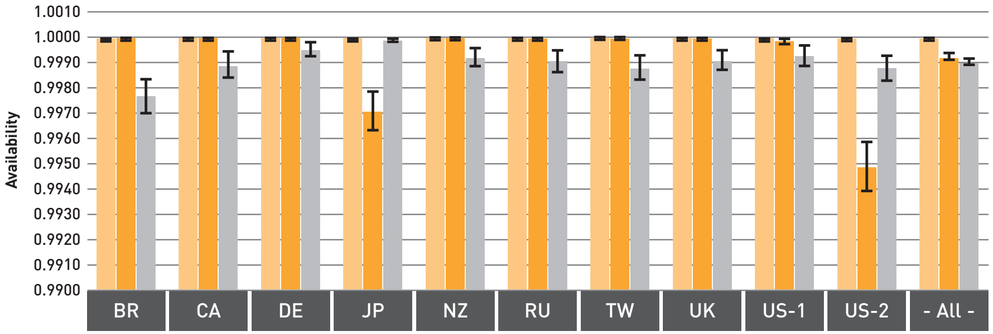

	BR	CA	DE	JP	NZ	RU	TW	UK	US-1	US-2	- All -
Direct	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
JonDo	1.0000	1.0000	1.0000	0.9971	1.0000	1.0000	1.0000	1.0000	0.9999	0.9949	0.9992
Tor	0.9977	0.9989	0.9995	0.9999	0.9992	0.9991	0.9988	0.9991	0.9993	0.9988	0.9990

Figure 3: Measurement results for the QoS parameter availability

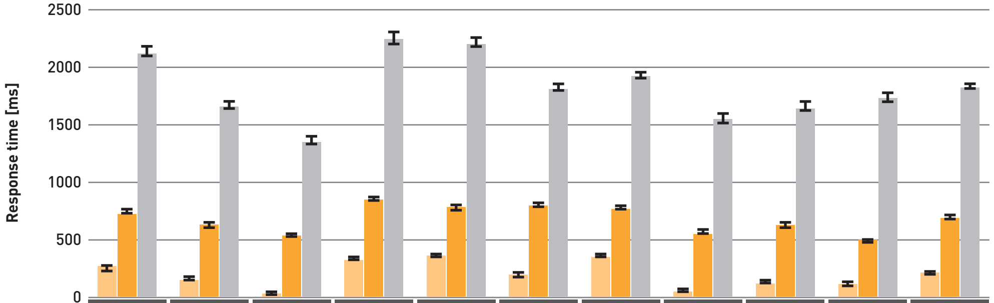

	BR	CA	DE	JP	NZ	RU	TW	UK	US-1	US-2	- All -
Direct	283.70	156.20	26.40	328.70	370.80	202.40	350.30	50.70	115.70	112.50	200.10
JonDo	742.30	640.20	535.60	864.70	789.30	798.70	772.40	549.20	626.20	498.60	681.90
Tor	2127.40	1670.20	1356.20	2247.30	2207.60	1809.70	1923.90	1560.60	1650.10	1736.40	1829.80

Figure 2: Measurement results for the QoS parameter response time

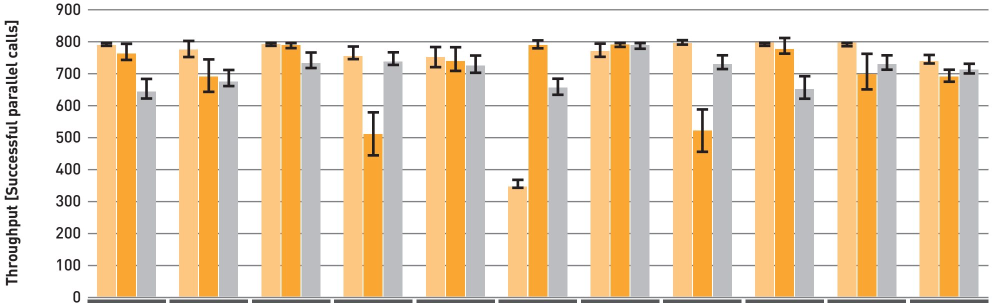

	BR	CA	DE	JP	NZ	RU	TW	UK	US-1	US-2	- All -
Direct	800.00	778.90	800.00	761.90	755.20	351.20	773.70	799.40	800.00	800.00	745.00
JonDo	766.00	693.40	794.70	513.90	743.60	791.80	798.80	521.90	785.30	704.80	696.10
Tor	649.90	684.10	741.60	746.30	728.70	661.90	789.70	735.00	655.60	732.70	715.00

Figure 4: Measurement results for the QoS parameter throughput

Using the commercial JonDo anonymity system, the average value increases to about 680 ms across all providers. Individual observations range between approximately 500 ms and 865 ms. This corresponds to a relative increase in response time of up to 2,000%. Thus, the effect, i.e., increase in response time compared to a direct access, is statistically significant.

The same is true for the Tor anonymity system, where the average response time reaches an average value of about 1,830 ms. In accordance, even for the geographically close providers in Western Europe (Germany, UK), the service response times exceed 1,000 ms. Thus, the relative increases are as high as 5,000%.

With respect to the QoS parameter of *availability*, the effects of anonymity systems appear negligible in absolute terms upon first sight (Figure 3). However, it should be recalled that most financial institutions have undertaken substantial and costly efforts in the past decades to achieve availabilities in their IT systems in the class of "five nines" (i.e., 99.999%) or more.

When using a direct mode of access, we did – in fact – observe a perfect availability of 100% across all servers, even in geographically remote locations. Once JonDo is applied however, we find a statistically significant reduction in service availability to about 99.92% on average, resulting in a mean availability in the class of merely "three nines". For Tor, the effect is even more pronounced and also consistently observable across all servers, resulting in a mean availability of about 99.9%, which would translate into about nine hours of service unavailability per year in practice.

Concerning the last QoS parameter in our experiments, *throughput*, we observed two oppositional effects: On the one hand, the throughput decreases compared to a direct access, because the anonymity nodes act as additional bottleneck between the service provider and the consumer. This effect is especially pronounced for the Tor network, where the nodes are operated by voluntary participants and may be subject to rather low bandwidth supply.

On the other hand, we also found an *increase* in throughput for a set of servers. The most likely explanation for this observation is that the anonymity networks artificially queue Web service requests. Thus, the series of parallel requests that we employed in our experiments become less "bursty" and can be more efficiently processed by those servers that possess less computing power.

Nevertheless, the overall effect of the two anonymity systems on the QoS parameter of throughput is significantly negative.

Thus, concerning our previously stated research question, we conclude that the use of anonymity systems does have a significant and negative impact on the QoS parameters of response time, availability, and throughput in the context of service consumption.

Practical Implications

As it has been suggested in the previous section, our findings have a number of practical implications.

First, the use of external services via public networks has to be generally reassessed. Cost savings and higher flexibility should be weighed against the disclosure of potentially sensitive business information to third parties. The latter aspect is of special importance in the financial institutions, which underlie rigid regulation with respect to data privacy.

Second, the application of anonymity systems as a potential countermeasure has to be thoroughly considered. While these mechanisms facilitate the security objective of relationship anonymity, they may also bring about substantial degradation in the QoS of business process execution. Depending on the requirements of the specific business case (e.g., time criticality of transactions), these effects may be problematic.

Third, the actual choice of an anonymity system can play an important role; our findings indicate that commercial systems have advantages with respect to QoS (specifically in respect to the parameter of response time), but they also impose additional cost for the service consumption.

Summary

Despite the use of standard security mechanisms such as encryption, the security objective of relationship anonymity may be threatened when external IT services are used over public networks, which allows attackers to infer potentially sensitive business information. With the spread of cloud computing, this issue gains further practical relevance, because an increasing number of services is offered via the Internet. State-of-the-art anonymity systems such as Tor or JonDo provide a suitable countermeasure, but do have a significant and negative impact on the QoS of service executions, resulting in, e.g., prolonged service response times and thus, higher latency in the execution of business processes.

Under these circumstances, financial institutions should thoroughly consider both the use of external services in general as well as the application of anonymity mechanisms, depending on the requirements of the specific business case and process. For this purpose, our research and empirical findings can provide a valuable decision support.

References

Miede, A.; Simsek, G.; Schulte, S.; Abawi, D. F.; Eckert, J.; Steinmetz, R.:
Revealing Business Relationships – Eavesdropping Crossorganizational Collaboration in the Internet of Services.
In: Proceedings of the Tenth International Conference Wirtschaftsinformatik (WI 2011).

Federal Office for Information Security:
SOA-Security-Kompendium – Sicherheit in Service-orientierten Architekturen – Version 2.0 (2009) [in German]
Available online:
https://www.bsi.bund.de/ContentBSI/Themen/SOA/StudienPublikationen/Studien_Publikationen.html