An Analysis of Anonymity Side Effects in the Internet of Services

Ulrich Lampe*, André Miede†, Tim Lusa*, Stefan Schulte‡, Ralf Steinmetz* and Schahram Dustdar‡
*Multimedia Communications Lab (KOM), Technische Universität Darmstadt, Darmstadt, Germany
Email: {ulrich.lampe, tim.lusa, ralf.steinmetz}@KOM.tu-darmstadt.de
†Fakultät für Ingenieurwissenschaften, HTW des Saarlandes, Saarbrücken, Germany
Email: andre.miede@htw-saarland.de
‡Distributed Systems Group, Vienna University of Technology, Vienna, Austria
Email: {s.schulte, dustdar}@infosys.tuwien.ac.at

***Abstract*—The Internet of Services will facilitate cross-organizational collaboration by allowing companies to utilize services from external providers. Even though standard security mechanisms such as message encryption may be in place, attackers could create detailed profiles of service consumers, providers, and marketplaces by monitoring communication endpoints. This threatens the security objective of relationship anonymity and potentially permits sensitive information about the underlying business processes or relationships between service consumers and providers to be revealed. While state-of-the-art countermeasures in the form of anonymity systems allow this problem to be addressed, they may have undesired side effects on the Quality of Service. This work provides a detailed empirical analysis of these side effects, based on an extensive measurement of the response time, availability, and throughput of representative, globally distributed services. Our experimental results are available to the interested public within the comprehensive dataset *WS-Anon*.**

***Keywords*-Internet of Services; Web Services; Security; Anonymity; Quality of Service**

I. Introduction and Motivation

In recent years, Service-oriented Architectures (SOAs) [1] have emerged as one important paradigm to realize tighter business and information technology (IT) alignment. SOA not only facilitates the integration of existing IT systems within a single organization, but also enables service-based cross-organizational workflows. While the SOA paradigm is technology-independent as such, Web service technologies are currently the most common way to implement service-oriented concepts [1], and have led to the vision of an *Internet of Services*, where services for all areas of life and business will be offered on the Web. The Internet of Services provides the foundation for complex business networks by supporting the composition and aggregation of existing services to value-added services, i.e., using marketplaces as intermediaries between service consumers and providers [2].

In order to enable such service-based cross-organizational collaboration, the security of the participating systems, exchanged messages, and utilized communication channels is a necessity. Regarding the security of Web service technology, substantial advancements have been achieved in recent years [3]. Nevertheless, we have identified several technology-independent and service-specific attacks in our past research, especially with respect to the Internet of Services [4].

In the work at hand, we address a specific security threat that relates to the security goal of "relationship anonymity" [5]: By monitoring communication at message exchange endpoints, attackers may be able to compile detailed profiles of service consumers, providers, and also of marketplaces. It is important to note that this attack does not target the actual *content* of the exchanged messages – which can be secured through standard security mechanisms, most notably encryption – but the principle fact of an active communication relationship between two parties. Depending on the monitoring means used, this attack may not even be illegal and can reveal important information about the underlying business processes and relationships between multiple parties [6].

Thus, due to the lack of protection by the currently used Web service security technologies, potentially very sensitive information is available to an attacker. However, standard anonymity systems are available for communication systems in order to achieve the goal of relationship anonymity. These mechanisms are principally applicable to Web services and could be used until dedicated solutions are available. Unfortunately, as an undesired side effect, these anonymity systems are very likely to have a negative impact on the Quality of Service (QoS) of Web service executions. Therefore, the overarching research question we seek to answer in this paper is: *"What is the impact of the use of anonymity systems on the Quality of Service of Web service-based communication?"*

The remainder of the paper is structured as follows: In the next section, we present the design and procedure of the experiments we conducted. Section III contains a discussion of the results and potential limitations of our experiments. Section IV gives an overview of related work, and Section V concludes the paper with a summary.

II. Experiment Design and Procedure

Due to the fact that Web services constitute a de facto standard implementation of SOA, we assume that cross-organizational collaboration in the Internet of Services can be realized by adopting this specific technology [1]. By applying standard Web service technologies like WSDL and SOAP,

communication is conducted using HTTP or HTTPS. Thus, it is possible to apply well-proven anonymity networks such as the ones presented within Section II-B.

In our experiments, we invoke Web services on the Internet and measure their QoS levels. We choose a testbed-based approach because it will be a good indicator for the potential QoS behavior in the Internet of Services. Specifically, we utilize real-world networks (i.e., the Internet), because this eliminates the need for a complex modeling of Internet traffic. We further employ real-world anonymity systems with real-world users, with which modeling errors are avoided, because complex systems (including user behavior) do not have to be rebuilt.

In contrast to our previous work [7], we do not employ existing, publicly available Web services, but implement our own customized Web service. This permits a more accurate measurement of QoS parameters. For instance, it can be verified that the partial unavailability of a Web service results from the network (or anonymity system), and not from a technical problem on the server. In addition, our own implementation permits a global distribution of servers and Web services, which is a key element of the envisioned Internet of Services [2].

Technically, the distribution of services has been achieved through deployment in the *PlanetLab* network [8]. PlanetLab is a globally distributed network of computers hosted by research institutions. As of mid-2012, it consisted of more than 1,000 nodes at approximately 500 sites. While no testbed can in the end represent the actual Internet [9], PlanetLab nevertheless offers an acceptable degree of abstraction for the scenario applied in this paper.

The next sections describe which QoS parameters and anonymity systems we chose for our experiments. We further present the measurement infrastructure and procedure.

A. Quality of Service Parameters

According to Schmitt, "QoS is the well-defined and controllable behavior of a system with respect to quantifiable parameters" [10]. Within service-oriented computing, QoS plays a very vital role: Service requesters require the service provider to fulfill a specific QoS level [11]. Hence, QoS guarantees are written down in *service level agreements* and are a major reason for selecting a particular service. We employ a QoS model by Repp [12], which has been specifically devised in the context of service-based systems.

Based on a seminal work on QoS-aware Web service composition by Menascé et al. [13], we select three quantitatively measurable QoS parameters from this model:

- *Response time* relates to the performance aspect of efficiency. It is measured as the time difference between the initiation of a Web service request and the reception of the corresponding reply.
- *Availability* is an aspect of dependability and trustability. It is measured as the ratio between the successful number of calls and the total number of calls of a Web service, i.e., it corresponds to the probability that a Web service is available.
- *Throughput* – in correspondence with response time – relates to the aspect of efficiency. It is measured by the maximum number of successful parallel invocations of a Web service, i.e., the number of service instances a service host is able to provide at the same time.

B. Anonymity Systems

Since the seminal work on anonymity by Chaum [14], a variety of both theoretical and live systems for providing anonymous communication has been developed and deployed [15]. In the work at hand, we focus on two low-latency anonymity systems that have also been evaluated in our previous work [7], namely *Tor*[1] ("The Onion Router") [16] and *JonDo*[2]. We deliberately do *not* set up our own anonymity network, because the very function principal of such network implies that the nodes are operated by third parties that cannot be linked to the initial sender of a message.

In contrast to our previous research, we did not examine *I2P*[3] ("Invisible Internet Project"). Although I2P improves continually, it is not mature enough for the scenarios described; in fact, the developers themselves point out the "relatively small size of the network and the lack of extensive academic review" on the project's Web site.

Tor is chosen because of its low-latency characteristics, world-wide distribution, and easy deployment. The Tor network is based on the participation of (anonymous) volunteers who operate nodes. At a specified time frequency, three (new) nodes are chosen by the sender at random – forming a so-called circuit of entrance, middle, and exit Tor router – to relay messages between the sender and the destination. Due to an "onion-style" encryption, each node only knows the preceding sender and subsequent receiver in the communication chain, but not the initial sender and receiver in conjunction [17].

JonDo is selected due to its strong security measures, such as a mandatory certification for node operators, its easy deployment, and its high reputation based on the research background of its developers. In JonDo, the client can choose between different "Mixes" or cascades of them, which are operated by certified providers. A Mix basically obfuscates the relationship between its input and output of messages so that an observer cannot link both message sets. JonDo provides the option of using the system for free or paying for additional features. Because the free version does not support the use of secure transport protocols like HTTPS, we only employ the commercial alternative in our experiments.

For more information regarding Tor and JonDo, we refer the interested reader to our previous work [7] or the respective project Web sites, which also provide conceptual comparisons of different systems.

C. Measurement Infrastructure and Procedure

For the realization of empirical measurements, we have implemented a prototypical infrastructure. As depicted in

[1] http://www.torproject.org/
[2] http://anonymous-proxy-servers.net/
[3] http://www.i2p2.de/

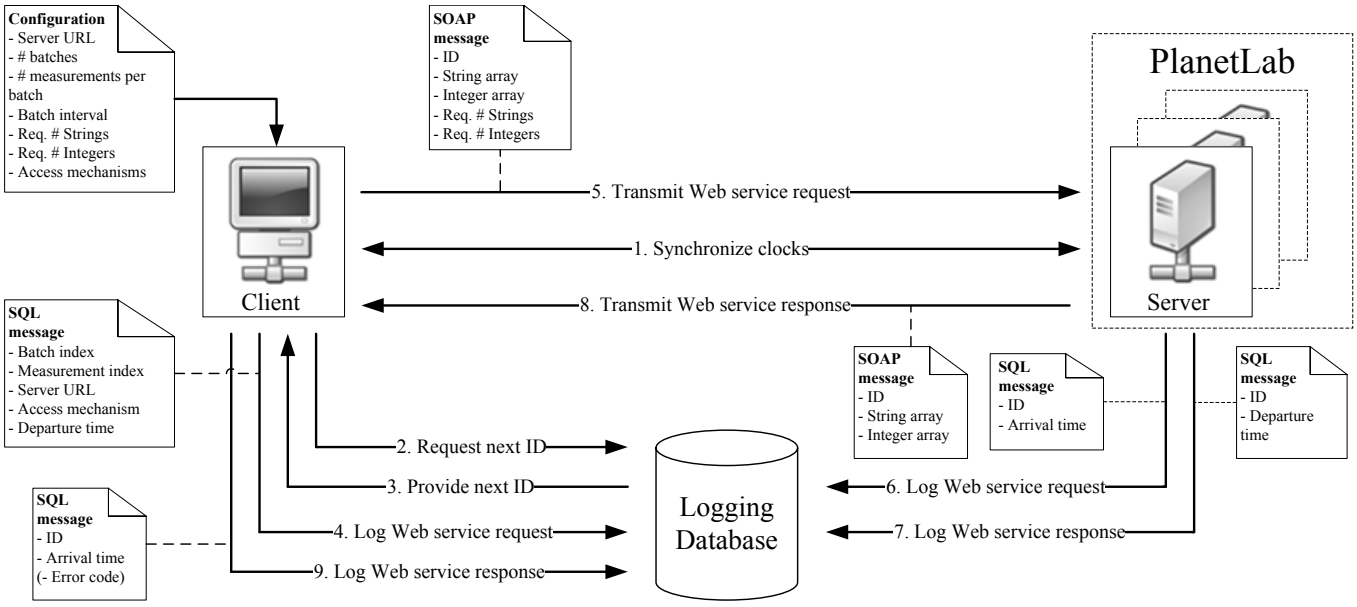

Figure 1. Schematic overview of the measurement infrastructure for the experiment.

Figure 1, the infrastructure comprises three components:

- A Java-based client tool, which runs on a dedicated machine located in Darmstadt, Germany, with a high bandwidth Internet connection, which conducts the Web service calls.
- A Java-based Web service, which is deployed on multiple nodes within the global PlanetLab network that serve the client. The worldwide distribution of servers – which is reflected by the country codes in Table I in the appendix – is inspired by empirical studies of public Web services [18], [19].
- A database, which is provided through a data center in Germany and is utilized for logging all results.

To initiate a new experiment, the client tool is provided with configuration parameters. These parameters include the number of test batches, each of which comprising of a number of individual measurements to be executed, and the time interval between the start of these batches. In addition, the target server and utilized *access mechanisms* – i.e., direct, JonDo, and/or Tor – are specified.

For each batch, client and target server first synchronize their clocks using the *Network Time Protocol*. Subsequently, for each individual measurement, a new ID is retrieved and a new log record is written into the database. Then, a SOAP message with random payload is generated. The payload is approximately 2 KB in size and comprises *String* and *Integer* arrays. These payload characteristics reflect the findings of an empirical study by Kim and Rosu [19]. The request is sent to our Web service on the target server, which responds with a SOAP message with new random data of the same structure. In accordance with our scenario, the content of all message transfers is encrypted using HTTPS. Upon transmission and receipt of a SOAP message on both client and server side, the log record is updated with a timestamp. If an error occurs in the initial transmission of the request, the client only logs the corresponding error code.

In general, the measurements are conducted in sequence, i.e., the next Web service request is initiated after the previous request has been completed or timed out. However, in order to measure throughput, a predefined number of Web service requests are transmitted in parallel. For this purpose, we set a maximum number of 800 parallel Web service requests, which we determined as an approximate throughput limit for the utilized server nodes.

From the log data, numeric values for each of the previously mentioned QoS parameters can be computed:

- *Response time* corresponds to the time difference between the initiation of a Web service request and the receipt of the respective response by the client. In accordance with Repp [12], the transfer time includes the processing time on both the client and server side (e.g., message serialization and encryption), as well as the network latency.
- *Availability* is determined based on the logged error code, i.e., it denotes the ratio of successful (error-free) Web service calls and overall Web service calls.
- *Throughput* is given by the number of successfully finished parallel Web service requests, i.e., the absolute count of requests within a batch of Web service invocations that are completed within a certain time interval.

Using this measurement infrastructure and procedure, it is now possible to carry out dedicated experiments, i.e., empirical measurements of QoS in direct and anonymous Web service invocations.

III. Output Analysis and Discussion

This section discusses the measurements and their implications. In order to minimize the influence of random events

based on, e.g., machine or network loads both on the client and server side, as well as in the anonymity systems, 2.25 million Web service calls were executed during different times of the day over a course of approximately three weeks.

Roughly three fourths of these calls belong to parallelized throughput measurements, in which each set of 800 calls constitutes one observation, resulting in 70 observations per server and access mechanism (i.e., 2,100 observations in total, based on 1,680,000 individual calls). The remainder belongs to sequential measurements, where each individual call constitutes one observation, which corresponds to 18,200 observations per server and access mechanism (i.e., 546,000 observations and calls in total).

We test for statistically significant differences in the observed QoS between the different access mechanisms using the following procedure: First, we determine the subsample of corresponding observations for each combination of QoS parameter and access mechanism. Subsequently, we compute the mean value and confidence interval (using $\alpha = 0.05$) across all observations in the subsample. We then apply an adapted version of the so-called "visual test", which tests for overlap in the confidence intervals; i.e., an impact or difference is deemed significant if there is no overlap between the confidence intervals. In case this leads to ambiguous results, we additionally employ a t-test [20] at the identical confidence level of $\alpha = 0.05$. We additionally apply the previously described procedure to the respective subsamples of observations for each individual server. That is, we further test for QoS differences between the access mechanisms for each individual server, rather than on average across all servers.

An overview of all measurements and the results of the statistical tests can be found in Table I. In addition, Figure 2 illustrates the cumulative distribution function for the QoS parameter response time, and Figure 3 depicts bar charts for the QoS parameters of response time and throughput.

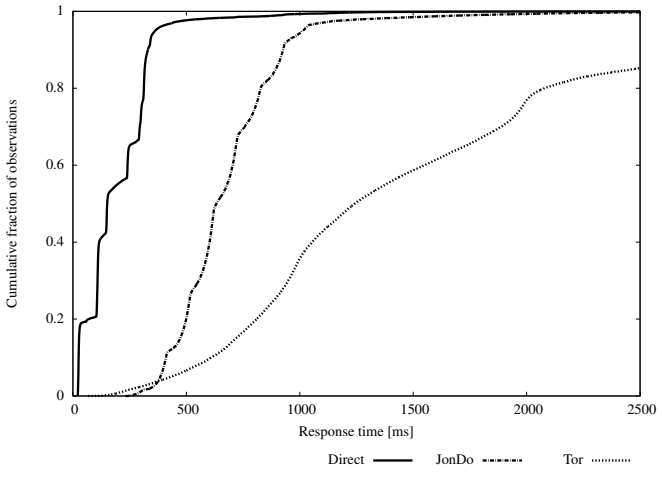

Figure 2. Cumulative distribution function for the QoS parameter response time (based on the aggregated sample across all servers).

A. Impact on Response Time

For the directly transmitted Web service requests, we observe a mean response time of about 200 ms. However, there are substantial variations, depending on the geographical location of the utilized server (cf. Figure 3a and Table I).

For the Web service requests using JonDo, we find significant differences in response time compared to a direct access for all ten regarded servers. On average, the use of JonDo leads to an increase in response time of about 500 ms. As can be seen from the cumulative distribution function in Figure 2, the increase is relatively constant compared to a direct access, resulting in a shift of the corresponding curve to the right. The largest increases in relative terms are found for servers that are geographically close to the client. Specifically, for the servers in Germany (DE) and the United Kingdom (UK), the increase amounts to up to 2,000%. In contrast, the smallest relative increases are observed for geographically remote servers, e.g., in New Zealand (NZ).

As a side note, it should be pointed out that we measured a substantially higher increase in response time of about 2,000 ms for the commercial version of JonDo in our previous work [7]. The difference can most likely be explained by the use of HTTPS in the experiments at hand, which appears to be transferred via a dedicated and apparently less populated proxy in the JonDo network. However, the observed change could also be attributed to (non-publicized) changes in the system architecture of JonDo in the time interval between our two experiments.

For the Web service requests using Tor, we also find significant differences in response time for all ten regarded servers compared to a direct access. On average, the use of Tor results in an increase in response time of approximately 1,600 ms. This increase is significantly higher than the change that can be observed for JonDo, both for each server individually and on average. In addition, Figure 2 indicates that the increase in response time is less constant (and, thus, less predictable) compared to a direct access. As in the case of JonDo, the highest relative increase in response time for Tor, amounting to more than 5,000%, can be observed for the geographically close server in Germany (DE). Again, the lowest relative increase occurs with a geographically remote server in Taiwan (TW).

In summary, both anonymity systems have a significant impact on the QoS parameter of response time even though they are deemed to be low-latency systems [21]. For both systems, this impact is negative, i.e., it corresponds to a substantial increase in response time, with Tor exhibiting a significantly stronger effect than JonDo.

B. Impact on Availability

For all Web service requests that were directly transmitted, we found an availability of 100% across all servers.

For the Web service requests that were transmitted via JonDo, we found significant differences in the availability of two regarded servers compared to a direct access (JP, US-2). For one server, we observed a slight, yet statistically non-significant decline in availability (US-1). For the remaining seven servers,

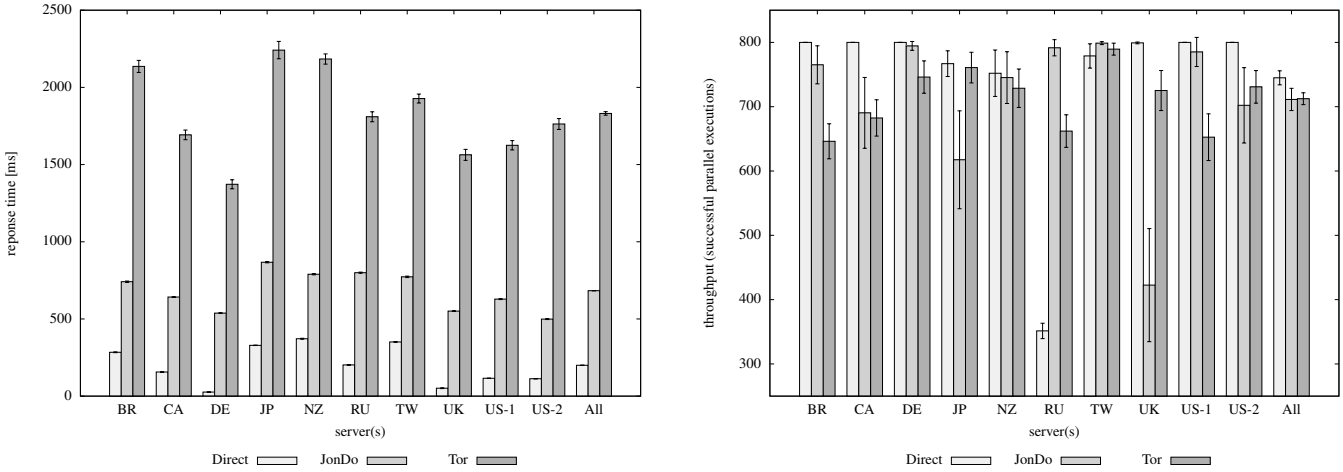

(a) Response time.

(b) Throughput.

Figure 3. Measurement results for selected QoS parameters per server (mean values with 95% confidence intervals).

an availability of 100% was found, thus providing results identical to those of a direct access.

In contrast, for the Web service requests using Tor, we find significant differences in availability for nine out of the ten regarded servers, compared to a direct access. For the remaining server (JP), no reduction in availability could be observed. For all but two servers (JP, US-2), Tor exhibits significantly worse results, with respect to availability, than JonDo. However, due to these outliers, the impact is not significantly different on average across all servers.

In summary, both anonymity systems have a significant impact on the QoS parameter of availability. This notion is supported by the average value across all servers. On a more detailed level, while the use of Tor affects all but one single server in our experiments, the impact of JonDo is limited to three servers. In any case, the average availability for all servers remains well in the class of *two nines*, i.e., above 99%, with both anonymity systems.

C. Impact on Throughput

For the Web service requests that were directly submitted, a mean throughput of approximately 750 parallel Web service requests was measured (cf. Figure 3b). The Russian-based server (RU) constitutes a notable outlier with a mean throughput of about 350 parallel requests. This outlier can likely be explained by insufficient computing power, i.e., the server is not capable of successfully handling more than the observed number of requests in parallel.

Using JonDo, we observe a significant change in throughput for seven out of ten servers. Interestingly, the change is positive for two of the servers (RU and TW). A potential explanation is that JonDo queues some of the Web service requests, resulting in a delayed transmission by the client and thus receipt by the server. Accordingly, the workload on the server may become less “bursty”, resulting in a smaller quantity of requests being discarded. For the remaining servers, the decline is usually in the range of 10% to 15%, with the notable exception of one server where we observe a more pronounced decline of more than 45% (UK).

For the Web service requests that were transmitted with Tor, the results are comparable to those of JonDo, with a significant change in throughput for seven servers. Again, the impact is positive for two servers (significantly for RU and non-significantly for TW), with the alleged explanation being the same as for JonDo.

Comparing JonDo with Tor, we obtain mixed results: While Tor delivers a significantly lower throughput for four servers (BR, DE, RU, US-1), the system also performs significantly better for two servers (JP, UK). Across all servers, we find no significant difference between the two anonymity systems. However, the average across all servers indicates marginal, yet non-significant advantages for Tor with respect to the QoS parameter of throughput.

In summary, the average throughput values across all servers support the notion that both anonymity systems have a significant impact on the QoS parameter of throughput. In general, the effect is negative, resulting in a reduction in throughput. However, for selected servers in our experiments, the change is positive.

D. Summary of Findings

The results of our experiments indicate that the use of anonymity systems does have a significant impact on the QoS parameters response time and availability. In our experiments, the effects were negative, resulting either in a prolongation of response times or a reduction in availability. In the case of throughput, we obtained mixed results: In some cases, the presumed queuing of Web service requests may lead to a higher throughput value. However, the average impact of anonymity systems on throughput is also negative, resulting in a reduction of the maximum number of parallel Web service calls.

Thus, in reference our research question from Section I, we

conclude: *"The impact of the use of anonymity systems on the Quality of Service of Web service-based communication is significant and negative regarding the parameters response time, availability, throughput."*

With respect to the comparison of the two regarded anonymity systems, we have obtained mixed results: There is a significant difference with respect to response time, where Tor introduced an additional latency per request of about 1,150 ms compared to JonDo. However, in the case of availability and throughput, the results differ on a server to server basis, thus not permitting any statistically significant statements about the two systems on average.

In summary, with respect to the two anonymity systems, we conclude: *"The impact of the use of Tor on the Quality of Service of Web service-based communication is significantly more negative compared to (the commercial version of) JonDo for the parameter of response time, but not significantly different for availability and throughput."* For reasons of fairness, it should be pointed out again that we employed the *commercial* version of JonDo in our experiments, which, in contrast to Tor, imposes a usage fee.

Overall, our analysis indicates that both the side effects of Web service anonymity and the choice of an anonymity system should be thoroughly considered whenever cross-organizational, service-based collaboration is carried out within the Internet of Services. In practice, the decision whether a gain in anonymity (and thus, security) should be traded for a degradation in QoS lies in the hand of the service consumer. This decision may be subject to various considerations, such as regulatory requirements, internal policies that govern security, or cost considerations. In this respect, our findings provide valuable decision support, because they permit reliable quantification of Web service anonymity side effects. Accordingly, we provide our experimental results to the interested public within the comprehensive dataset *WS-Anon*[4].

E. Limitations

To complete the output analysis and discussion, we want to acknowledge that our research approach has some potential and factual shortcomings and limitations.

First, we only employed *one* central client that acted as service consumer. Depending on the geographical location of a service consumer, the specific QoS effects of anonymity systems may differ. In our experiments, for instance, the JonDo system may have profited from the proximity of the Mix cascade; more specifically, both the client and the Mix were located in Germany, which potentially results in lower (additional) network latency. Similarly, in the case of Tor, the location of the client within the anonymity network topology may have a substantial impact on the overall performance. For example, if many Tor nodes are geographically close to the client, the network's overall performance behavior may be more advantageous due to reduced network latency. In fact, according to empirical research by McCoy et al. [17], a substantial share of nodes and bandwidth within Tor could be attributed to German users in 2008. Thus, from a global standpoint, the performance of both JonDo and Tor may have been overestimated in our experiments.

Second, all measurements in our experiments were taken between a central client and *individual services*. Thus, we assume that a workflow is centrally controlled and that each service is individually invoked by the middleware, thus following an *orchestration* approach [22]. The *end-to-end QoS* of the complete workflow can be computed based on the QoS properties of the individual services, using appropriate aggregation functions [23], [24]. However, it may also be assumed that the services within a workflow autonomously interact through direct message exchange, thus following an *choreography* approach [22]. In the context of service selection, such a model has been assumed by, e.g., Yu et al. [25]. In this case, the QoS properties of the communication links between pairs of services – rather than just between the client and each service – would be of relevance as well.

Third, the geographical distribution of servers as well as the size and structure of the exchanged messages in our experiments is based on empirical findings by Kim and Rosu [19] and Zheng et al. [18]. However, these authors have examined contemporary *public* Web services, which may not necessarily be representative of the future *commercial* services that will be offered in the Internet of Services.

Fourth, the PlanetLab nodes that acted as servers in our experiments are operated by third parties. Thus, these nodes were not under our full control in terms of performance behavior; specifically, the computational load on an individual node may be subject to the requirements of other PlanetLab users. This effect may be comparable to the load on an actual public Web server. However, because PlanetLab is primarily intended for scientific purposes, it does not necessarily fully resemble the behavior of a commercial Web service host in the Internet of Services.

IV. Related Work

The work at hand constitutes a substantial extension of our previous research [7]. In this past work, our aim was to quantify the impact of anonymity systems on the efficiency of real-world, public Web service executions, represented by the QoS attribute of response time. Now, we additionally regard the QoS attributes of availability and throughput. We also validate the significance of the observed QoS difference using common statistical methods. This work should also be seen as a complement to another previous publication [6], in which we aimed to quantify the risks of attacks on anonymity systems, but did not examine the side effects of countermeasures.

Issues related to QoS, especially performance aspects such as network latency, have been addressed previously both for Web services, e.g., [18], [26], and anonymity systems, e.g., [16], [21], [27]. However, the combination of Web services and anonymity systems has – to the best of our knowledge – not been covered by research so far, except by our previous work and the work at hand.

[4]http://www.kom.tu-darmstadt.de/ws-anon/

A general analysis of the Tor system has been conducted by McCoy et al. [17]. Their work focuses on information about the applications and users of Tor and highlighting potential security problems in the system. It lacks, e.g., a comparison to other systems such as JonDo. Dhungel et al. present an analysis of reasons for delays within Tor, but also do not provide comparable numbers for JonDo [27].

Empirical measurements in the Tor and AN.ON (the predecessor of JonDo) networks have been conducted by Wendolsky et al. [28]. In contrast to our work, the authors focus on general Internet applications and do not operate a dedicated infrastructure for measurements, but use public Web sites. Wendolsky et al. further restrict their analysis to the QoS parameters of latency and bandwidth in terms of the data transfer rate, thus omitting availability and throughput as observed within the study at hand. Experiments regarding latency and bandwidth in Tor alone have been conducted by, e.g., Dingledine and Murdoch [29]. However, the numbers observed are for general Web servers and not discussed in the context of the Internet of Services.

Besides, to the best of our knowledge, the issue of anonymous communication between the different organizational participants of an SOA has not received attention from the research community. Further aspects of anonymity such as the issue of anonymous Web service provision, as well as consumption are addressed, e.g., by Papastergiou et al. [30]. However, such functionality is rather undesirable in the context of cross-organizational collaboration in the Internet of Services. In fact, the situation that provider and consumer know and trust each other constitutes a prerequisite for functions such as billing, but also for compliance reasons.

In summary, to the best of our knowledge, none of the related work has previously examined the side effects of anonymity systems on different QoS parameters in the context of the Internet of Services. As major contribution, our work provides an empirical analysis based on a realistic setting with real-world anonymity systems and multiple distributed Web services.

V. Summary

In the Internet of Services, organizations will be able to realize (parts of) their business processes using services from external providers. This results in a variety of novel security threats. Specifically, the security goal of relationship anonymity may be threatened: By surveilling the message exchange between an organization and its service providers, attackers might unveil sensitive information about the underlying business processes, even if standard security mechanisms, such as encryption, are in place.

However, the goal of relationship anonymity may be realized through the use of standard anonymity systems, which obfuscate the communication relation between two parties. In the work at hand, we have examined and empirically analyzed to which extent such anonymity systems – specifically, the JonDo and Tor systems – may affect the Quality of Service of service executions.

For this purpose, we have implemented a testbed-oriented distributed measurement infrastructure. We have conducted experiments involving ten globally distributed server nodes using PlanetLab, thus collecting approximately 2.25 million individual measurements of Web service calls. In our analysis, we found that the use of anonymity systems significantly and negatively affects the Quality of Service parameters of response time, availability, and throughput. We also observed that Tor provides significantly more negative results with respect to response time, but could not find any statistically significant differences for the other two parameters.

In the future, we aim to exploit these findings through the extension of our existing QoS-aware service selection approach [23], [24]. This will permit the explicit consideration of anonymity requirements, but also the observed side-effects of anonymity, in the composition and execution of service-based workflows.

Acknowledgements

This work is partially supported by E-Finance Lab e. V., Frankfurt am Main, Germany (www.efinancelab.de) as well as the Austrian Science Fund (FWF): P23313-N23. We would like to thank Christian Gottron for his extensive support regarding PlanetLab and Sebastian Kaune for his valuable hints concerning the statistical methods.

References

[1] M. P. Papazoglou, P. Traverso, S. Dustdar, and F. Leymann, "Service-oriented Computing: State of the Art and Research Challenges," *Computer*, vol. 40, no. 11, pp. 38–45, 2007.

[2] C. Schroth, "The Internet of Services: Global Industrialization of Information Intensive Services," in *Proceedings of the Second IEEE International Conference on Digital Information Management (ICDIM 2007)*, 2007, pp. 635–642.

[3] R. Kanneganti and P. Chodavarapu, *SOA Security*. Manning Publications, 2008.

[4] A. Miede, T. Ackermann, N. Repp, D. F. Abawi, R. Steinmetz, and P. Buxmann, "Attacks on the Internet of Services – The Security Impact of Cross-organizational Service-based Collaboration," in *Proceedings of the Multikonferenz Wirtschaftsinformatik (MKWI 2010)*, 2010, pp. 425–426 (short) and 2151–2162 (full).

[5] A. Pfitzmann and M. Hansen, "A Terminology for Talking about Privacy by Data Minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management," http://dud.inf.tu-dresden.de/Anon_Terminology.shtml, Aug. 2010, v0.34. Last access on January 24, 2013.

[6] A. Miede, G. Şimşek, S. Schulte, D. F. Abawi, J. Eckert, and R. Steinmetz, "Revealing Business Relationships – Eavesdropping Cross-organizational Collaboration in the Internet of Services," in *Proceedings of the Tenth International Conference Wirtschaftsinformatik (WI 2011)*, vol. 2, 2011, pp. 1083–1092.

[7] A. Miede, U. Lampe, D. Schuller, J. Eckert, and R. Steinmetz, "Evaluating the QoS Impact of Web Service Anonymity," in *Proceedings of the Eighth IEEE European Conference on Web Services (ECOWS 2010)*, 2010, pp. 75–82.

[8] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "PlanetLab: An Overlay Testbed for Broad-Coverage Services," *ACM SIGCOMM Computer Communication Review*, vol. 33, pp. 3–12, 2003.

[9] N. Spring, L. Peterson, A. Bavier, and V. Pai, "Using PlanetLab for Network Research: Myths, Realities, and Best Practices," *ACM SIGOPS Operation Systems Review*, vol. 40, pp. 17–24, 2006.

[10] J. Schmitt, *Heterogeneous Network Quality of Service Systems*. Kluwer Academic Publishers, 2001.

[11] R. Berbner, M. Spahn, N. Repp, O. Heckmann, and R. Steinmetz, "Heuristics for QoS-aware Web Service Composition," in *Proceedings of the 3rd IEEE International Conference on Web Services (ICWS 2006)*, 2006, pp. 72–82.

[12] N. Repp, *Überwachung und Steuerung dienstbasierter Architekturen – Verteilungsstrategien und deren Umsetzung*. Books on Demand, 2009, in German.

[13] D. A. Menascé, "Composing Web Services: A QoS View," *IEEE Internet Computing*, vol. 8, no. 6, pp. 88–90, 2004.

[14] D. L. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, 1981.

[15] J. Ren and J. Wu, "Survey on Anonymous Communications in Computer Networks," *Computer Communications*, vol. 33, no. 4, pp. 420–431, 2010.

[16] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-generation Onion Router," in *Proceedings of the Thirteenth Conference on USENIX Security Symposium*, 2004, pp. 303–320.

[17] D. McCoy, K. Bauer, D. Grunwald, T. Kohno, and D. Sicker, "Shining Light in Dark Places: Understanding the Tor Network," in *Proceedings of the Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, 2008, pp. 63–76.

[18] Z. Zheng, Y. Zhang, and M. Lyu, "Distributed QoS Evaluation for Real-World Web Services," in *Proceedings of the 7th IEEE International Conference on Web Services (ICWS 2010)*, 2010, pp. 83–90.

[19] S. M. Kim and M.-C. Rosu, "A Survey of Public Web Services," in *Proceedings of the Fifth International Conference on E-Commerce and Web Technologies (EC-Web 2004)*, 2004, pp. 96–105.

[20] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley, 1991.

[21] S. J. Murdoch and R. N. Watson, "Metrics for Security and Performance in Low-Latency Anonymity Systems," in *Proceedings of the Eighth International Symposium on Privacy Enhancing Technologies (PETS 2008)*, 2008, pp. 115–132.

[22] N. M. Josuttis, *SOA in Practice: The Art of Distributed System Design*. O'Reilly Media, 2007.

[23] D. Schuller, A. Miede, J. Eckert, U. Lampe, A. Papageorgiou, and R. Steinmetz, "QoS-based Optimization of Service Compositions for Complex Workflows," in *Proceedings of the Eighth International Conference on Service-oriented Computing (ICSOC 2010)*, 2010, pp. 641–648.

[24] D. Schuller, U. Lampe, J. Eckert, R. Steinmetz, and S. Schulte, "Cost-driven Optimization of Complex Service-based Workflows for Stochastic QoS Parameters," in *Proceedings of the 19th International Conference on Web Services (ICWS 2012)*, 2012, pp. 66–74.

[25] T. Yu, Y. Zhang, and K. Lin, "Efficient Algorithms for Web Services Selection with End-to-End QoS Constraints1," *ACM Transactions on the Web*, vol. 1, no. 1, pp. 1–26, 2007.

[26] S. Rosario, A. Benveniste, S. Haar, and C. Jard, "Probabilistic QoS and Soft Contracts for Transaction-Based Web Services Orchestrations," *IEEE Transactions on Services Computing*, vol. 1, no. 4, pp. 187–200, 2008.

[27] P. Dhungel, M. Steiner, I. Rimac, V. Hilt, and K. W. Ross, "Waiting for Anonymity: Understanding Delays in the Tor Overlay," in *Proceedings of the IEEE Tenth International Conference on Peer-to-Peer Computing (P2P 2010)*, 2010, pp. 1–4.

[28] R. Wendolsky, D. Herrmann, and H. Federrath, "Performance Comparison of Low-Latency Anonymisation Services from a User Perspective," in *Proceedings of the Seventh International Symposium on Privacy Enhancing Technologies (PET 2007)*, 2007, pp. 233–253.

[29] R. Dingledine and S. J. Murdoch, "Performance Improvements on Tor or, Why Tor is slow and what we're going to do about it," https://svn.torproject.org/svn/projects/roadmaps/2009-03-11-performance.pdf, last access on January 24, 2013, 2009.

[30] S. Papastergiou, G. Valvis, and D. Polemi, "A Holistic Anonymity Framework for Web Services," in *Proceedings of the First International Conference on Pervasive Technologies Related to Assistive Environments (PETRA 2008)*, 2008, p. 38.

Table I

EXPERIMENTAL RESULTS FOR THE QOS PARAMETERS BY SERVER ("CI95" DENOTES THE HALF WIDTH OF THE 95% CONFIDENCE INTERVAL. SIGNIFICANT DIFFERENCES BETWEEN DIRECT ACCESS AND JONDO/TOR ARE MARKED WITH Δ, SIGNIFICANT DIFFERENCES BETWEEN JONDO AND TOR ARE INDICATED BY δ. SAMPLE SIZE FOR THE PARAMETERS RESPONSE TIME AND AVAILABILITY PER SERVER AND ACCESS MECHANISM IS $n = 18,200$, SAMPLE SIZE FOR THE PARAMTER THROUGHPUT PER SERVER AND ACCESS MECHANISM IS $n = 70$).

Server	RESPONSE TIME [MS]						AVAILABILITY				THROUGHPUT			
	Direct		JonDo		Tor		Direct		JonDo		Tor		Direct		JonDo		Tor	
Code	Mean	CI95	Mean	CI95	Mean	CI95	Mean	CI95	Mean	CI95	Mean	CI95	Mean	CI95	Mean	CI95	Mean	CI95
BR	284.0	2.7	741.1$^{\Delta}$	4.5	2,136.0$^{\Delta\delta}$	39.0	1.0000	0.0000	1.0000	0.0000	0.9977$^{\Delta\delta}$	0.0007	800.0	0.0	765.1$^{\Delta}$	29.6	646.2$^{\Delta\delta}$	27.2
CA	156.2	0.6	641.8$^{\Delta}$	2.9	1,692.4$^{\Delta\delta}$	31.2	1.0000	0.0000	1.0000	0.0000	0.9989$^{\Delta\delta}$	0.0005	800.0	0.0	690.4$^{\Delta}$	55.1	682.5$^{\Delta}$	28.2
DE	26.4	0.1	538.0$^{\Delta}$	2.4	1,372.3$^{\Delta\delta}$	29.3	1.0000	0.0000	1.0000	0.0000	0.9995$^{\Delta\delta}$	0.0003	800.0	0.0	794.5	6.9	746.1$^{\Delta\delta}$	25.2
JP	329.1	1.4	867.5$^{\Delta}$	4.1	2,241.6$^{\Delta\delta}$	56.3	1.0000	0.0000	0.9973$^{\Delta}$	0.0008	1.0000$^{\delta}$	0.0000	766.9	20.0	617.5$^{\Delta}$	76.1	760.8$^{\delta}$	23.8
NZ	371.7	3.2	789.8$^{\Delta}$	4.0	2,183.8$^{\Delta\delta}$	32.9	1.0000	0.0000	1.0000	0.0000	0.9992$^{\Delta\delta}$	0.0004	752.0	36.1	745.2	40.3	728.7	29.9
RU	202.0	1.8	799.6$^{\Delta}$	3.7	1,809.8$^{\Delta\delta}$	32.4	1.0000	0.0000	1.0000	0.0000	0.9991$^{\Delta\delta}$	0.0004	351.4	12.0	791.6$^{\Delta}$	12.6	662.1$^{\Delta\delta}$	25.3
TW	350.8	2.6	772.8$^{\Delta}$	4.7	1,927.7$^{\Delta\delta}$	29.2	1.0000	0.0000	1.0000	0.0000	0.9988$^{\Delta\delta}$	0.0005	778.9	18.9	798.8$^{\Delta}$	2.3	789.5	9.3
UK	51.3	2.6	551.0$^{\Delta}$	2.6	1,563.0$^{\Delta\delta}$	35.6	1.0000	0.0000	1.0000	0.0000	0.9991$^{\Delta\delta}$	0.0004	799.2	1.4	422.6$^{\Delta}$	87.9	725.1$^{\Delta\delta}$	31.1
US-1	115.7	0.5	628.5$^{\Delta}$	2.9	1,625.0$^{\Delta\delta}$	30.3	1.0000	0.0000	0.9999	0.0001	0.9993$^{\Delta\delta}$	0.0004	800.0	0.0	785.1	22.6	652.6$^{\Delta\delta}$	36.3
US-2	112.5	0.7	499.0$^{\Delta}$	3.0	1,762.9$^{\Delta\delta}$	35.1	1.0000	0.0000	0.9948$^{\Delta}$	0.0010	0.9987$^{\Delta\delta}$	0.0005	800.0	0.0	702.1$^{\Delta}$	58.6	730.8$^{\Delta}$	25.3
All	200.0	0.8	682.9$^{\Delta}$	1.3	1,831.4$^{\Delta\delta}$	11.4	1.0000	0.0000	0.9992$^{\Delta}$	0.0001	0.9990$^{\Delta}$	0.0001	744.9	10.9	711.3$^{\Delta}$	17.3	712.4$^{\Delta}$	9.2

BR=Brazil, CA=Canada, DE=Germany, JP=Japan, NZ=New Zealand, TW=Taiwan, UK=United Kingdom, US=United States