

- 1
2
3
4
5
6
7
8
9
- 10 • *Level of anonymity*: One needs to make some logical assumptions re-
11 garding the parties that would be able to access the ID of a given user.
12 For example, should the ID of some user be available merely to himself
13 and his home domain or only the owner of the ID will have access to
14 it? This is crucial in order to decide which user's sensitive information
15 contained in SIP headers needs to be protected and in which particular
16 case.
17
 - 18 • *Anonymity vs. pseudonymity*: This issue has to do with what is stated
19 in section 2.1. A person receiving a call from a UA using a pseudonym
20 can always return the call using the same pseudonym. This is not
21 feasible with totally anonymous schemes. Also, pseudonymity means
22 that all protected IDs are recoverable by the corresponding SIP en-
23 tities, which are the home servers of both parties. This means that
24 data retention policies do not need to change; service providers can log
25 connection information and recover a user ID upon request.
26
 - 27 • *Accountability*: Further to the previous point, if the anonymity scheme
28 does not support the standard SIP registration process then account-
29 ability (and billing) cannot be enforced.
30
 - 31 • *Cryptography*: Some schemes rely on cryptography to keep personal in-
32 formation private while others employ other means. Those that do not
33 use cryptography will probably be faster and have less administrative
34 requirements in terms of key management.
35
 - 36 • *Deployment cost*: This criterion has to do with the easiness of de-
37 ployment of a scheme. For instance, a scheme that requires the full
38 deployment of PKI, as that of S/MIME, presents a high cost.
39
 - 40 • *Depth of protection*: An answer on whether a scheme is capable of pro-
41 tecting user's ID private information leaking from other layers (apart
42 from the application one) is required.
43
 - 44 • *User-centric vs. Centralised IdM*: Two of the methods can be charac-
45 terised as user-centric; Anonymous URI, and partly the one given in
46 RFC 3323. SIPS and IPsec are based on the construction of secure
47 tunnels and thus IdM with reference to SIP is not applicable to them.
48 All the others offer centralised IdM.
49
- 50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9
- *Overhead*: This refers to the overhead imposed by the solution at hand. Leaving aside the increased resource consumption caused to servers, in SIP, a critical parameter is that of the user’s service time (latency). This parameter is only discussed in solutions (Karopoulos et al., 2011, 2010) and, as expected, was found to be closely related to the selected cryptographic scheme. The use of a symmetric algorithm like AES, resulted in insignificant delays. In contrary, a SIP request preparation delay may be increased to over 45 milliseconds (ms) per message operation in case of asymmetric algorithms (Karopoulos et al., 2010). On the other hand, the mean server response delay for a SIP server having a queue size of 1000 calls may be increased up to a maximum of 800 ms.
- 10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

Taking the above into account, we can provide a short but comprehensive comparison between the solutions described in the three previous subsections. Sometimes the caller wishes not to reveal their ID to the callee. This ID hiding option is offered by the PrivaSIP methods (Karopoulos et al., 2011, 2010), “Anonymous URI”, RFC 3323, and RFC 5767. Nevertheless, only the PrivaSIP ones are able to afford this feature while protecting the Digest username during the authentication process at the same time (see figure 2). Also, the same methods are in position to keep their privacy protecting features active while operating through untrusted domains. S/MIME can also protect the user ID, still it is unable to protect their username during authentication. Moreover, it cannot offer caller’s ID hiding from the callee. The protection of the home domain name of the caller can be only achieved by the use of “Anonymous URI”. However, as explained in section 3.1.1, this method has little practicality since it cannot support authentication. Regarding the IP addresses of the communicating parties it is evident that no method except SIPS, RFC 5767 and IPsec ones can effectively protect them from eavesdroppers. RFC 5767 and IPsec (under certain mode and algorithm of operation) can also provide privacy protection down to IP layer. Still, one has to carefully consider the special network architecture required by the RFC 5767 solution and the deployment cost imposed by the IPsec one.

26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

Another observation is that no scheme is able to protect the callee ID only. This may be useful for example in cases a user calls a certain hotline. As mentioned in section 1 this scenario could be of value for service providers as the available information to an observer would be “user U is calling a

52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 service from domain D2". Persistent traffic analysis is also not considered by
10 any of the above schemes. This, however, consists an important threat as in
11 any VoIP ecosystem the probability of a call between a couple of users does
12 not occur at a uniform rate, given that each user usually has a different set
13 of contacts. So, de-anonymisation attacks via the exploitation of long-term
14 statistic information are made possible (Zhang and Fischer-Hubner, 2013;
15 Danezis, 2003).
16
17

18 A last point of discussion, which is also brought up later on in section
19 3.2.4, is the utilisation of some anonymous communication system, like the
20 well-known Tor, to maximize the level of obfuscation achieved regarding SIP
21 messages. Note that such anonymisation systems are self-reliant, i.e., usu-
22 ally its operation does not depend on the protocols of upper layers, hence
23 they can be seamlessly combined with them. Taking Tor as an example,
24 the problem with SIP is that currently Tor only supports TCP for its trans-
25 port layer. So, although (Rosenberg et al., 2002) requires all SIP entities to
26 mandatory implement both UDP and TCP, many real-world VoIP applica-
27 tions rely solely on UDP for latency reasons. So, at least for the time being,
28 this is a serious impediment for VoIP users to enjoy strong anonymity to
29 real-time voice communication. Tunneling of the UDP traffic through Tor
30 does not really solve this issue because the traffic would be encapsulated
31 in TCP. The latency induced by Tor is also increased as the system relays
32 and mixes its traffic via multiple nodes. Despite that, recently, a first effort
33 to realise VoIP over Tor was materialised in an opensource product called
34 Torfone (TorFone, 2013). However, Torfone is not based on SIP but on an
35 (obsolete) version of zfone (ZRTP) (<http://zfoneproject.com/>). Its imple-
36 mentors do recognise this latency problem by stating that "*The payment for*
37 *anonymity is voice latency up to 2-4 seconds*". This observation is roughly
38 verified by some early and still ongoing experimental results of ours showing
39 an additional mean latency of about 700 ms when routing SIP traffic over
40 Tor. This time penalty is associated only to SIP signaling and it is perceived
41 starting from the moment the caller's UA sends out an invite until an OK
42 message is received by her. In any case, this is a quite interesting research
43 issue and it is sure to gain momentum as Tor network performance increases
44 over time, and some day it will eventually support UDP as well. For the
45 interested reader, a detailed analysis of similar to Tor solutions can be found
46 in (Edman and Yener, 2009; Ren and Wu, 2010; Ruiz-Martinez, 2012).
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

3.2. The case of Kerberos

As already pointed out, the Kerberos protocol (Neuman et al., 2005) is quite different in nature as compared to SIP. In fact, Kerberos can be seen as a service used for other services to authenticate users. Nowadays, Kerberos is one of the most well-established three-party authentication and key management protocols over open and insecure networks (MIT-Kerberos-Consortium, 2014). The protocol offers a SSO platform through the use of tickets, i.e., a piece of enciphered and integrity protected information that enables a user to be authenticated without re-entering their password. By capitalising on the extensive adoption of Kerberos by modern application services, Kerberos is also starting to gather considerable attention as a solution to provide federated access to any kind of application service through AAA infrastructure (Perez-Mendez et al., 2013).

The standard Kerberos protocol lacks of a mechanism to preserve user privacy. More precisely, in a similar way to SIP, Kerberos identifies the different participant entities via identifiers, which are in the form of “principal@realm”. For example, gkamb@AEGEAN.GR and printer/server.aegean.gr@AEGEAN.GR are valid IDs of a user and a service respectively within the AEGEAN.GR realm. Unfortunately, these principal IDs associated to both clients and services are communicated in cleartext. More specifically, the service identifier for which the ticket has been issued is conveyed in cleartext. Even more, the two messages of the AS exchange (Neuman et al., 2005) contain the client’s ID which is being authenticated by the AS module of the Key Distribution Center (KDC). The identity of the service the client is willing to access is also visible to any eavesdropper when monitoring the Ticket Granting Server (TGS) exchange (Neuman et al., 2005). Undeniably, this situation clearly violates the principle of user anonymity as an observer can straightforwardly learn the client’s real ID and discover which services are being accessed by them. Figure 3 depicts a typical Kerberos message flow concerning both single and cross-realm operation.

The basic service access model defined in Kerberos also contributes to privacy violations. This is due to the Kerberos atomic operation where the client first performs a message exchange with the KDC to acquire a ticket that is used in a subsequent exchange to access a service. In fact, this situation is present in the AS exchange (to provide the client a Ticket Granting Ticket (TGT) to be used in the TGS exchange) as well as in the TGS exchange (needed for the client to obtain a Service Ticket (ST) that is delivered

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

to the service via the AP exchange). So, an eavesdropper is able to easily correlate the different messages sent or received by a client to access a service. Even worse, a listener is in position to collect information about behavioral patterns of service access of given users in the network. This is true because typically the acquisition of an ST by a client to access a service is performed via the use of the same TGT used previously for obtaining other services. This simply means that service access unlinkability is not preserved, as by tracing the use of a given TGT, a malicious actor can figure out that the same - even anonymous - client is accessing these services (Tene, 2011; King and Jessen, 2010).

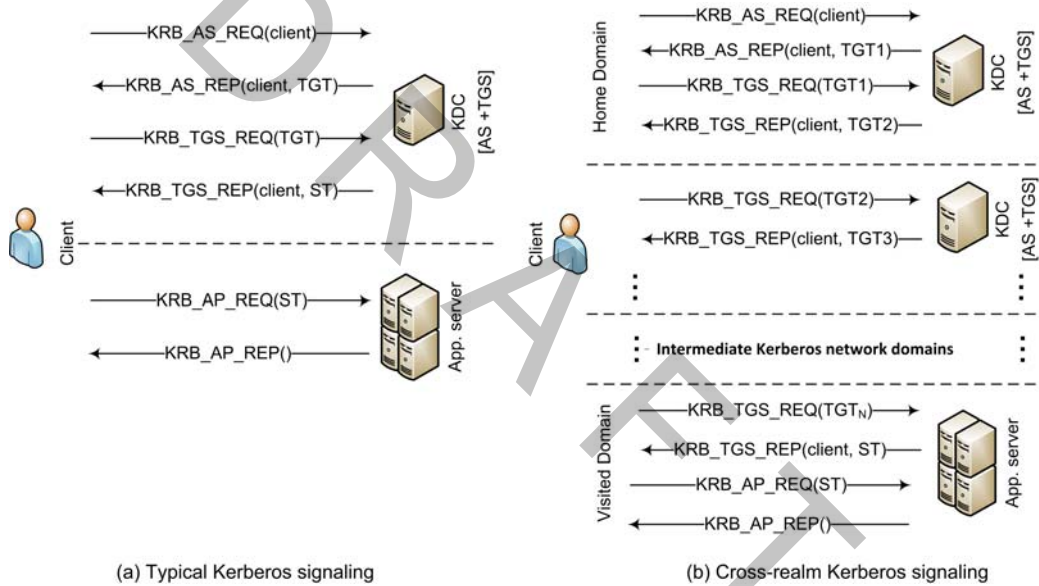


Figure 3: Kerberos protocol message flow

Taking all the above into account, to achieve user anonymity in Kerberos one must (a) guarantee that the real identifier (e.g., the username) of a given user involved in Kerberos transactions remain hidden, and (b) prevent malicious entities from being able to cross-relate the different messages sent and/or received by a specific user, thus providing message exchange unlinkability. Naturally, this anonymity facility is better to include multi-domain (e.g., federated) Kerberos environments as well. In the next subsections, a concise survey of the current privacy-enabling solutions in Kerberos is offered,

1
2
3
4
5
6
7
8
9 followed by a discussion of the findings.

10 11 12 *3.2.1. Standardisation efforts*

13 In an attempt to offer user anonymity, the work in (Medvinsky et al.,
14 1998; Zhu et al., 2011) enhances Kerberos protocol by introducing the anonym-
15 ous ticket concept. Instead of being assigned to a specific user regis-
16 tered in a realm, an anonymous ticket is associated to the anonymous user
17 (anon@anon). Therefore, the true client identity is not revealed neither
18 to the service nor to eavesdroppers. This, of course, pertains to a purely
19 anonymous scheme rather a pseudonymous one. However, for the anonym-
20 ous TGT acquisition, this solution requires the utilisation of anonymous
21 Public Key Cryptography for Initial Authentication in Kerberos (PKINIT)
22 (Zhu and Tung, 2006). Specifically, for the KDC to securely deliver to the
23 client the session key associated to the anonymous ticket, a secure channel
24 needs to be established by making use of certificate-based KDC authentica-
25 tion. This leads to a PKI-dependent solution which is not usually available,
26 especially in multi-domain environments. Moreover, the user remains com-
27 pletely anonymous even to KDC and services which are considered trusted.
28 This situation bears accountability problems as the organization controlling
29 the foreign domain typically needs for example to charge the visited user for
30 the services they obtained.

31
32 The Generalized Framework for Kerberos Pre-authentication
33 (Hartman and Zhu, 2011) comprises another solution towards address-
34 ing the client privacy issue in Kerberos. This framework is concerned with
35 the protection of client's identity in the messages transmitted from the
36 client to KDC, in such a way that the use of anonymous PKINIT is not
37 required. Simply put, by combining the anonymous ticket concept with the
38 security extensions defined in (Hartman and Zhu, 2011), clients are able
39 to obtain anonymous tickets. On the downside, this procedure mandates
40 the acquisition of a special ticket called armor TGT (contains a symmetric
41 key known to both the client and KDC to protect Kerberos exchanges),
42 which in turn, presents certain deficiencies. Note that an armor TGT must
43 be obtained before a client starts to utilise the pre-authentication extensions
44 with a given KDC. Specifically, three solutions are proposed. First, using its
45 real identity, a client is able to exercise a standard AS exchange to request
46 an armor TGT. This however allows listeners to easily correlate the client's
47 ID with the acquired armor TGT, meaning that when the client employs
48 the armor TGT towards requesting an anonymous ticket, an eavesdropper
49
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 can derive their real identity associated to the anonymous ticket. Second,
10 the user can obtain an armor TGT via the use of anonymous PKINIT.
11 This of course requires the KDC to own a valid certificate, which in turn
12 requires PKI. In case a PKI infrastructure is not present, a final method
13 is to acquire the armor TGT using anonymous PKINIT without KDC
14 authentication. Nevertheless, as stressed out by the authors, this option is
15 prone to man-in-the-middle attacks.
16
17

18 3.2.2. Custom Solutions

19 The work in (Pereñíguez-García et al., 2011) proposes a privacy frame-
20 work for Kerberos, coined as “PrivaKERB”. This framework does not re-
21 quire the existence of PKI or other infrastructure external to Kerberos. A
22 prominent feature of PrivaKERB is that along with user anonymity it of-
23 fers service access unlinkability. The latter refers to the granting of tickets
24 problem pointed out in section 3.2, that enables eavesdroppers to trace the
25 different services accessed by a specific user. It is to be noted that the
26 identity-enhancing functionalities by this work remain in total harmony with
27 user identification required by processes such as accounting and charging.
28 Particularly, to deliver user anonymity, the authors make use of temporary
29 client pseudonyms (transaction pseudonyms) only valid for a specific pe-
30 riod of time. The KDC is in control of pseudonym generation and a new
31 pseudonym is mandatorily delivered to the client each time a fresh home
32 TGT is issued. Moreover, service access unlinkability is imposed via the use
33 of extended anonymous tickets which include the client’s pseudonym in such
34 a way that is only accessible by trusted parties (i.e., KDCs, services). To
35 obstruct the TGT-based linkability that takes place when a client reuses the
36 same TGT several times to solicit access to multiple services, the authors
37 propose a new kind of single-use TGT called “self-renewed TGT”. Summa-
38 rizing, this solution accomplishes a fair level of unlinkability that prevents
39 eavesdroppers from linking the different service accesses performed by a spe-
40 cific anonymous user. This way, attackers cannot deduce whether separate
41 service accesses belong to the same or different anonymous clients.
42
43

44 On the negative side, PrivaKERB contributes little in protecting users
45 from observers attempting to cross-link the sequence of different messages
46 communicated by a specific - even anonymous - user to acquire a service.
47 Indeed, the series of messages starting with *TGT acquisition from the KDC*
48 and followed by *ST obtainment from the KDC* and *ST handing over to the*
49 *service* can be associated to the same anonymous user, and thus, exploited
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 by attackers to reveal the different Kerberos transactions in which a client
10 participated. This however leads to a privacy breach in multi-domain scenar-
11 ios as potential eavesdroppers have the chance to easily find out the service
12 visited by a client in a remote Kerberos domain. If these pieces of data are
13 systematically collected, it can be used toward disclosing important informa-
14 tion such as the most preferred services for roaming users in a realm and/or
15 the origin realm of clients visiting a specific Kerberos domain.
16
17

18 Motivated by the aforementioned insufficiency, the same au-
19 thors contributed a full-fledged anonymity framework, called KAMU
20 (Pereñíguez-García et al., 2013), able to achieve a full obfuscation of the
21 protocol’s messages from an eavesdropper point of view. To fix the foregoing
22 linkability problem, the authors came up with the specification of a mecha-
23 nism able to obscure the KDC distributed tickets, so as to hinder observers
24 from tracking the different tickets acquired by a client. This mechanism
25 camouflages by means of encryption the ticket (TGT or ST) sent by the
26 KDC to the client. By doing so, eavesdroppers cannot observe the ticket
27 and only the client can recover it. The implementation of this solution is
28 based on both normal Kerberos tickets as well as a new type of ticket,
29 called “fake ticket”. The solution requires normal tickets to be
30 transmitted in a way that remain confidentiality and integrity protected,
31 rather than having certain parts of them being visible (recall that in
32 standard Kerberos the service’s ID for which, say, an ST ticket has been
33 granted is transmitted in cleartext). In this way, attackers are blocked
34 from accessing or modifying a ticket. Also, according to this solution, a
35 normal ticket is placed in the padata field (Neuman et al., 2005) of the
36 message. This is an extensible field defined by Kerberos with the aim to
37 develop new functionalities or convey additional data.
38
39
40
41
42
43

44 On the other hand, a fake ticket is an entirely new type of ticket hav-
45 ing all of its fields belonging to the protected part (named EncTicketPart
46 (Neuman et al., 2005)) contain meaningless (null or randomly initialized) in-
47 formation. This however does not apply to the flags field; this is done to
48 enable all entities to recognise a fake ticket from a standard one by the pres-
49 ence of the fake flag. A fake ticket is intended to replace the standard one,
50 and thus, it is placed in the normal ticket field (Neuman et al., 2005) of every
51 reply message issued by the AS or TGS. As a consequence, no one except
52 the authorised entities are capable of accessing the real TGT or ST being
53 communicated to the client.
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32

3.2.3. Lower level Solutions

As in the case of SIP, one apparent solution to deliver anonymity in Kerberos could be the use of TLS tunnels (Josefsson, 2011). This way, eavesdroppers would be blocked from snooping on sensitive information such as the client and service identifier. It is obvious though that this solution presents the limitation of the hop-by-hop requirement; that is, the creation of a TLS tunnel between every pair of communicating entities is needed. This is however particularly defective in multi-domain scenarios as the clients would need to establish a TLS session with every intermediary KDC in the path from the home to visited domain. Still, no one can guarantee (or even acknowledge) that every network hop does afford (or establish) a TLS tunnel. A multi-domain PKI (Shimaoka et al., 2008) has also to be in place, since a pre-established trust relationship between the client and intermediary realms is not usually the case. These requirements are sure to significantly increase the deployment cost of the solution. Naturally, as already pointed out in section 3.1.3, the same shortcomings are to be taken for granted for lower layer tunnels, as that of IPsec.

3.2.4. Discussion

From the above discussion it becomes apparent that the tackling of the anonymity and unlinkability problem in Kerberos presents many similarities to that of SIP. It can be said that several critical privacy insufficiencies have been identified and a considerable mass of works are devoted in solving the problem. As in the case of SIP, we can perceive some standardisation efforts along with custom and generic solutions. Nevertheless, once more, it can be argued that generic solutions are just passing the privacy problem to a lower layer's protocol (e.g., TLS, IPsec). As already highlighted, while this solution works for virtually every superjacent protocol in the stack, it presents certain shortcomings mainly due to the need of external infrastructures and the hop-by-hop impediment. So, while none of the aforementioned solutions requires changes to the core Kerberos protocol, some of them are fully or partially based on PKI, which unfortunately - at least until now - is not the case for the majority of network realms.

Using the key points already identified in section 3.1.4 and the discussion given in sections 1 and 2 we can notice the following qualities:

- *Level of anonymity*: As pointed out in section 3.2, all principal IDs associated to both clients and services are visible to an observer when

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

in transit. So, any solution should guarantee that no one except the user herself and the service has access to their real identifiers. This is however a basic (or first) level of anonymity because due to the atomic operation of Kerberos explained in section 3.2 and illustrated in figure 3, a malicious actor is able to cross-relate information contained in the protocol's message flow, and thus identify a client even in case they remain anonymous. Therefore, as discussed in section 2.1, to obtain a stronger level of anonymity in Kerberos one requires to also impose message unlinkability. Furthermore, any given solution needs to consider single- as well as multi-realm network deployments in a way that no sensitive information is leaked out even in cases where some or all of the intermediate realms in the path collude. Once again, as the provision of anonymity and the protection of digital ID is primarily related to context, another important property for any solution here is to be able to support all levels of anonymity in an opt-in basis. Indeed, this property seems to be satisfied by both PrivaKerb and KAMU. Lastly, lower level solutions are offering both anonymity and unlinkability, but unfortunately they provide a lesser degree of flexibility and present certain shortcomings that need to be laboriously evaluated prior to deployment.

- *Anonymity vs. pseudonymity*: This issue has mainly to do with what has been discussed in section 2.1. So, while a totally anonymous solution have been proposed for Kerberos (Medvinsky et al., 1998; Zhu et al., 2011) it seems that it is conflicting with accountability. This is because the client does not reveal their ID even to KDC and the service which are considered a priori trusted. Also, as discussed in section 3.2.1, this proposal requires PKINIT. The Generalized Framework for Kerberos Pre-authentication tries to solve this latter issue but creates another deficiency pertaining to unlinkability. On the other hand, the two custom solutions namely PrivaKERB and KAMU are based on pseudonymity. However, in relation to section 2.1, these solutions make use of transaction pseudonyms at the client side which is not really “one-time identity” but valid for a specific period of time. In addition, KAMU exercises an interesting camouflaging ticket scheme along with encryption to work-around the problem and achieve full obfuscation of the protocol's message flow.

- 1
2
3
4
5
6
7
8
9
- 10 • *Accountability*: Elaborating on the previous point, and under the um-
11 brella of section 2.2, totally anonymous solutions cannot support ac-
12 countability which is normally required by the service provider. In spe-
13 cial cases, where for example a service is given to clients for free, full
14 anonymity may be desirable (and thus the existence of such a scheme
15 would become very handy). Nevertheless, a mechanism to elevate back
16 and forth to an accountable state is usually needed in this case, which
17 in turn adds complexity to the system, and thus such a decision is
18 normally densely interwoven with the particular case at hand.
 - 19 • *Cryptography*: As discussed in 3.2.3 the lower level solutions which are
20 in charge of constructing a secure tunnel for channelling all protocol's
21 sensitive information through it are ordinarily impose heavier cryp-
22 tography compared to those working entirely at the application layer.
23 Namely, the higher the layer of protection the greater the level of cus-
24 tomisation. In this respect, custom solutions such as KAMU employ
25 tailor-made strategies to protect only the information that matters.
26 This is verified by the results reported in the context of these works
27 and briefly outlined further down. Lastly, bear in mind that a main so-
28 licitude of the solutions discussed in section 3.2.2 was critical tasks, like
29 that of pseudonym generation, to be consigned to KDC care in an effort
30 to discharge the client from frequent operations that add overhead.
 - 31 • *Deployment cost*: As in the case of SIP, several solutions proposed for
32 Kerberos impose the use of some sort of PKI. This however comes at
33 a high cost and makes deployment far from being simple. Flexibility
34 and compatibility with current implementations is also key issues here
35 towards building a truly workable solution.
 - 36 • *Depth of protection*: As with SIP, Kerberos works at the application
37 layer, so the protection of private information about a user disclosed by
38 other layers is needed. From the foregoing discussion it becomes glaring
39 that apart from the SSL and IPsec solutions, all the others cope with
40 privacy at Kerberos level only. This situation results in the same worri-
41 ment spotted for SIP in section 3.1.4 (also briefly sketched in section 1).
42 So, the question here is what happens with privacy-sensitive informa-
43 tion belonging to the TCP/IP layer over which Kerberos is conveyed?
44 IP address, ports, domain name, packet contexts, sizes and timing,
45 and round-trip times are only certain pieces of information that can be
- 46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9 used towards identifying the communicating parties. Naturally, this is
10 a cross-layer privacy problem that calls for solutions like that of Tor.
11 This is recognised by the authors in (Pereñiguez-Garcia et al., 2013)
12 which point out several possible solutions to integrate with KAMU.
13 They particularly focus on Tor and explain that the formation of an
14 alliance between KAMU and Tor results in a robust cross-layer privacy-
15 preserving communication system able to effectively deal with a con-
16 siderable number of privacy attacks.
17
18

- 19
20 • *User-centric vs. Centralised IdM*: Under the scope of section 2.3 all the
21 anonymity solutions discussed in the present section offer centralised
22 IdM. Of course, as in SIP, TLS and IPsec driven solutions do not pro-
23 vide any kind of IdM Kerberos intrinsic solution.
24
- 25
26 • *Overhead*: To our knowledge, experimental results about the penalty in
27 terms of service times are available only for PrivaKERB and KAMU.
28 Specifically, according to the authors of these works, the first one is
29 found to augment the AS and TGS exchange time by 0.35 ms and 1.4
30 ms respectively. As expected, KAMU produces higher overheads due to
31 the extra time required to distribute the reinforced in terms of privacy
32 ticket. This however is translated to a negligible latency of about 0.89
33 and 2 ms for AS and TGS exchange correspondingly. Also, compared
34 to the message processing time for standard Kerberos, KAMU pro-
35 duces insignificant overheads. For instance, in TGS exchange which
36 represents the worst-case, the authors recorded an increment of 1.1
37 and 0.3 ms for the client and the KDC respectively. For further de-
38 tails on these metrics and experimental results the reader can refer to
39 (Pereñiguez-Garcia et al., 2011, 2013).
40
41
42
43
44

45 In summary, the KAMU solution seems to be the most complete in re-
46 gards to its privacy features. It not only allows the client to remain anony-
47 mous and untraceable from eavesdroppers, but also does not hinder the iden-
48 tification of clients when needed, e.g., for accounting and charging processes.
49 Moreover, its privacy features are preserved in both single- and multi-domain
50 scenarios without the need of PKI, as it simply relies on existing Kerberos
51 extensibility mechanisms. On the other hand, works like (Zhu et al., 2011;
52 Medvinsky et al., 1998) attempt to render the client fully anonymous and
53 thus fail to support important accounting operations performed by trusted
54 entities.
55
56
57
58

4. Conclusions

The need for anonymity is inevitably present in almost any protocol, application or service used in wired or wireless networks. Undoubtedly, the need of being innominate is an issue of great importance as it comprises the basis to protect fundamental human rights, such as the free expression of ideas and opinions, and allow people to perform their online activities in comfort and privacy. In this context, the goal of this paper is twofold. First off, it sheds light on the various issues revolving around anonymity and argues that it is a versatile concept that includes and affects several others, such as that of accountability, linkability, identity management, and so forth. Secondly, it conducts a short but comprehensive survey on the anonymity-preserving solutions proposed so far in the literature regarding SIP and Kerberos protocols. This serves as a dual case study for investigating the ways user's anonymity, and more general privacy, is confronted and dealt in the context of major, well-established protocols used at large in the cyberspace. In this respect, the survey part of the work at hand differs from the great mass of earlier ones which particularly focus on Web anonymity or anonymisation tools.

Also, bear in mind that the choice to include two application layer protocols (and not another from a lower layer) is not taken without due consideration. This is because providing anonymity and privacy in general at the application layer is usually harder to achieve and therefore more interesting. That is, any proposed solution needs to be tailored to the application, support accountability, and retain compatibility with current implementations. Moreover, we have in mind the case studies to be somehow comparable to each other. This would be problematic if we choose protocols lying in different layers of the Internet stack. It is therefore really interesting to observe that although the two aforementioned protocols have totally different usage, they were found to utilise quite similar methods to address user's privacy.

Moreover, this study has confirmed that every anonymity-preserving solution considers either directly or indirectly, and at least to some degree, aspects like accountability as those have been identified in the first part of the current work. It has been also exhibited that the research on this topic is active and constantly growing as anonymity and privacy in general for most protocols and services have not been treated in a "by-design" fashion. In this context, it seems that the most noteworthy issues for future designs to deal with is that of offering cross-layer privacy-preserving systems, the compatibility with

1
2
3
4
5
6
7
8
9 the base protocols, the support of anonymity across different, but somehow
10 federated network realms, and the smooth integration of anonymity with
11 vital underlying network operations.
12

13 14 **References**

- 15
16 ABC4Trust, Jan. 2014. Eu project - attribute-based credentials for trust.
17 URL <https://abc4trust.eu>
18
- 19 Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., Levkowitz, H., June 2004.
20 Extensible authentication protocol (eap). IETF RFC 3748.
21
- 22 Adjei, J., Olesen, H., 2011. Keeping identity private. IEEE Vehicular Tech-
23 nology Magazine 6 (3), 70–79.
24
- 25 Arias Cabarcos, P., Almenarez, F., Gomez Marmol, F., Andres, M., 2013. To
26 federate or not to federate: A reputation-based mechanism to dynamize
27 cooperation in identity management. Wireless Personal Communications,
28 1–18.
29
- 30 Belenkiy, M., Camenisch, J., Chase, M., Kohlweiss, M., Lysyanskaya, A.,
31 Shacham, H., 2009. Randomizable proofs and delegatable anonymous cred-
32 entials. In: CRYPTO. pp. 108–125.
33
- 34 Burkell, J., 2006. Anonymity in behavioural research: Not being unnamed,
35 but being unknown. University of Ottawa Law & Technology Journal 3 (1),
36 189–203.
37
- 38 Camenisch, J., Dubovitskaya, M., Lehmann, A., Neven, G., Paquin,
39 C., Preiss, F.-S., 2013. Concepts and languages for privacy-preserving
40 attribute-based authentication. In: IDMAN. pp. 34–52.
41
- 42 Camenisch, J., Lysyanskaya, A., 2001. An efficient system for non-
43 transferable anonymous credentials with optional anonymity revocation.
44 In: EUROCRYPT. pp. 93–118.
45
- 46 Cameron, K., Jan. 2006. The laws of identity.
47 URL <http://www.identityblog.com/?p=354>
48
- 49 Cao, Y., Yang, L., 2010. A survey of identity management technology. In:
50 IEEE International Conference on Information Theory and Information
51 Security (ICITIS). pp. 287–293.
52
53
54
55
56
57
58

- 1
2
3
4
5
6
7
8
9 Cavoukian, A., Oct. 2006. 7 laws of identity - the case for privacy-embedded
10 laws of identity in the digital age. White paper, Information & Privacy
11 Commissioner, Ontario, Canada.
12 URL http://www.ipc.on.ca/images/Resources/up-7laws_whitepaper.pdf
13
14
15 Chadwick, D. W., 2009. Federated identity management. In: Aldini, A.,
16 Barthe, G., Gorrieri, R. (Eds.), Foundations of Security Analysis and De-
17 sign V. Vol. 5705 of LNCS. Springer Berlin Heidelberg, pp. 96–120.
18
19
20 Chaum, D., 2003. Untraceable electronic mail, return addresses and digital
21 pseudonyms. In: Secure Electronic Voting. pp. 211–219.
22
23 Danezis, G., 2003. Statistical disclosure attacks. In: of the IFIP TC11 18th
24 International Conference on Information Security (SEC '03). Kluwer, pp.
25 421–426.
26
27
28 Davenport, D., 2002. Anonymity on the internet: why the price may be too
29 high. Commun. ACM 45 (4), 33–35.
30
31
32 Dolera Tormo, G., Gomez Marmol, F., Martinez Perez, G., 2013. Towards
33 the integration of reputation management in openid. Computer Standards
34 & Interfaces In Press, Accepted Manuscript, –.
35
36 E.C., 2012. Proposal for a regulation of the european parliament and of
37 the council on the protection of individuals with regard to the processing
38 of personal data and on the free movement of such data (general data
39 protection regulation). COM(2012) 11 final.
40 URL [http://ec.europa.eu/justice/data-protection/document/review2012/
41 com_2012_11_en.pdf](http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf)
42
43
44 Edman, M., Yener, B., 2009. On anonymity in an electronic society: A survey
45 of anonymous communication systems. ACM Comput. Surv. 42 (1).
46
47
48 E.U., Oct. 1995. European parliament - protection of individuals with regard
49 to the processing of personal data and on the free movement of such data.
50 URL [http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:
51 31995L0046:EN:HTML](http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML)
52
53
54 Fouque, P.-A., Poupard, G., Stern, J., 2000. Sharing decryption in the con-
55 text of voting or lotteries. In: Financial Cryptography. pp. 90–104.
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 Gartner, Feb. 2013. Half of new retail customer identities will be based on
10 social network identities by 2015.
11 URL <http://www.gartner.com/newsroom/id/2326015>
12
13 GEANT-project, 2014. edugain service.
14 URL <http://www.edugain.org>
15
16 Handley, M., Jacobson, V., Perkins, C., July 2006. Sdp: Session description
17 protocol. IETF RFC 4566.
18
19 Hansen, M., Tschofenig, H., Smith, R., Oct. 2011. Privacy terminology and
20 concepts.
21 URL <http://tools.ietf.org/html/draft-hansen-privacy-terminology-03>
22
23 Hartman, S., Zhu, L., April 2011. A generalized framework for kerberos pre-
24 authentication. IETF RFC 6113.
25
26 Hoepman, J.-H., May 2013. Revocable privacy.
27 URL <http://www.cs.ru.nl/~jhh/revocable-privacy/index.html>
28
29 ITU-T, Jan. 2009. Ngn identity management framework. Recommendation
30 Y.2720.
31
32 Jennings, C., Peterson, J., Watson, M., Nov. 2002. Private extensions to
33 the session initiation protocol (sip) for asserted identity within trusted
34 networks. IETF RFC 3325.
35
36 Johnson, C. Y., Sept. 2009. Project gaydar.
37 URL [http://www.boston.com/bostonglobe/ideas/articles/2009/09/20/
38 project_gaydar_an_mit_experiment_raises_new_questions_about_online_
39 privacy/](http://www.boston.com/bostonglobe/ideas/articles/2009/09/20/project_gaydar_an_mit_experiment_raises_new_questions_about_online_privacy/)
40
41 Joseffson, S., May 2011. Using kerberos version 5 over the transport layer
42 security (tls) protocol. IETF RFC 6251.
43
44 Karopoulos, G., Kambourakis, G., Gritzalis, S., 2011. Privasip: Ad-hoc iden-
45 tity privacy in sip. Computer Standards & Interfaces 33 (3), 301–314.
46
47 Karopoulos, G., Kambourakis, G., Gritzalis, S., Konstantinou, E., 2010. A
48 framework for identity privacy in sip. J. Network and Computer Applica-
49 tions 33 (1), 16–28.
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 Kim, M., 2010. The right to anonymous association in cyberspace: Us legal
10 protection for anonymity in name, in face, and in action. *SCRIPTed* 51
11 7 (1), 51–70.
- 12
13 King, N. J., Jessen, P. W., 2010. Profiling the mobile customer—privacy
14 concerns when behavioural advertisers target mobile phones—part i.
15 *Computer Law & Security Review* 26 (5), 455–478.
16 URL [http://www.sciencedirect.com/science/article/pii/
17 S0267364910001044](http://www.sciencedirect.com/science/article/pii/S0267364910001044)
- 18
19
20
21 Kügler, D., Vogt, H., 2002. Offline payments with auditable tracing. In: *Fi-
22 nancial Cryptography*. pp. 269–281.
- 23
24 Lessig, L., 2006. *Codev2*. Basic Books.
25 URL <http://www.codev2.cc/download+remix/Lessig-Codev2.pdf>
- 26
27 Mahy, R., Matthews, P., Rosenberg, J., April 2010. Traversal using relays
28 around nat (turn): Relay extensions to session traversal utilities for nat
29 (stun). IETF RFC 5766.
- 30
31
32 Maliki, T. E., Seigneur, J.-M., 2013. Online identity and user management
33 services - chapter 25. In: *Computer and Information Security Handbook*
34 (Second Edition), second edition Edition. Morgan Kaufmann, Boston, pp.
35 459–484.
- 36
37
38 Medvinsky, A., Cargille, J., Hur, M., March 1998. Anonymous credentials in
39 kerberos. IETF Internet Draft.
40 URL <http://tools.ietf.org/html/draft-ietf-cat-kerberos-anoncred-00>
- 41
42
43 MIT-Kerberos-Consortium, Jan. 2014. Mit kerberos & internet trust (mit-
44 kit) consortium. MIT Kerberos & Internet trust (MIT-KIT) Consortium.
45 URL <http://www.kerberos.org/>
- 46
47 Mukhamedov, A., Ryan, M. D., 2005. On anonymity with identity escrow.
48 In: *Formal Aspects in Security and Trust*. pp. 235–243.
- 49
50 Munakata, M., Schubert, S., Ohba, T., April 2010. User-agent-driven privacy
51 mechanism for sip. IETF RFC 5767.
- 52
53
54 Narayanan, A., Shmatikov, V., 2008. Robust de-anonymization of large
55 sparse datasets. In: *IEEE Symposium on Security and Privacy*. pp. 111–
56 125.
- 57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 Neuman, C., Yu, T., Hartman, S., Raeburn, K., July 2005. The kerberos
10 network authentication service (v5). IETF RFC 4120.
11
12 OASIS, March 2005. Assertions and protocols for the oasis security assertion
13 markup language (saml) v2.0.
14 URL <http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf>
15
16
17 OAuth, Aug. 2013. Oauth web site.
18 URL <http://oauth.net>
19
20
21 Ohm, P., 2009. Broken promises of privacy: Responding to the surprising
22 failure of anonymization. *UCLA Law Review* 57, 1701–.
23 URL <http://ssrn.com/abstract=1450006>
24
25
26 OpenID, Aug. 2013. Openid website.
27 URL <http://openid.net>
28
29 OpenID-Foundation, Jan. 2014. Openid connect v.1.0.
30 URL openid.net/connect/
31
32
33 Park, S., Park, H., Won, Y., Lee, J., Kent, S., Aug. 2009. Traceable anony-
34 mous certificate. IETF RFC 5636.
35
36
37 Pbd, Aug. 2013. Privacy by design web site.
38 URL <http://www.privacybydesign.ca/>
39
40
41 Pereñiguez-Garcia, F., Kambourakis, G., López, R. M., Gritzalis, S., Gómez-
42 Skarmeta, A. F., 2010. Privacy-enhanced fast re-authentication for eap-
43 based next generation network. *Computer Communications* 33 (14), 1682–
44 1694.
45
46
47 Pereñiguez-Garcia, F., Marin-Lopez, R., Kambourakis, G., Ruiz-Martinez,
48 A., Gritzalis, S., Skarmeta-Gomez, A., 2011. Privakerb: A user privacy
49 framework for kerberos. *Computers & Security*, 446–463.
50
51
52 Pereñiguez-Garcia, F., Marin-Lopez, R., Kambourakis, G., Ruiz-Martinez,
53 A., Gritzalis, S., Skarmeta-Gomez, A., 2013. Kamu: providing advanced
54 user privacy in kerberos multi-domain scenarios. *International Journal of*
55 *Information Security*, 1–21.
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9
10 Perez-Mendez, A., Pereñiguez-Garcia, F., Marin-Lopez, R., Lopez-Millan,
11 G., 2013. Out-of-band federated authentication for kerberos based on pana.
12 Computer Communications 36 (14), 1527–1538.
- 13
14 Peterson, J., Nov. 2002. A privacy mechanism for the session initiation pro-
15 tocol (sip). IETF RFC 3323.
- 16
17 Pfitzmann, A., Hansen, M., Aug. 2010. A terminology for talking about
18 privacy by data minimization: Anonymity, unlinkability, undetectability,
19 unobservability, pseudonymity, and identity management.
20 URL http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf
21
22
- 23 PRISM, 2013. Prism (surveillance program).
24 URL [http://en.wikipedia.org/wiki/PRISM_\(surveillance_program\)](http://en.wikipedia.org/wiki/PRISM_(surveillance_program))
25
- 26
27 Ramsdell, B., July 2004. Secure/multipurpose internet mail extensions
28 (s/mime) version 3.1 message specification. IETF RFC 3851.
- 29
30 Ren, J., Wu, J., 2010. Survey on anonymous communications in computer
31 networks. Computer Communications 33 (4), 420–431.
- 32
33 Rescorla, E., Modadugu, N., Jan. 2012. Datagram transport layer security
34 version 1.2. IETF RFC 6347.
- 35
36
37 Rosenberg, J., Oct. 2009. Obtaining and using globally routable user agent
38 uris (gruus) in the session initiation protocol (sip). IETF RFC 5627.
- 39
40 Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J.,
41 Sparks, R., Handley, M., Schooler, E., June 2002. Sip: Session initiation
42 protocol. IETF RFC 3261.
- 43
44
- 45 Ruiz-Martinez, A., 2012. A survey on solutions and main free tools for pri-
46 vacy enhancing web communications. Journal of Network and Computer
47 Applications 35 (5), 1473–1492.
- 48
49
50 Schneier, B., Jan. 2012. Anonymity won't kill the internet.
51 URL [http://www.wired.com/politics/security/commentary/
52 securitymatters/2006/01/70000?currentPage=all](http://www.wired.com/politics/security/commentary/securitymatters/2006/01/70000?currentPage=all)
53
- 54
55 Schulzrinne, H., Casner, S., Frederick, R., Jacobson, V., July 2003. Rtp: A
56 transport protocol for real-time applications. IETF RFC 3550.
- 57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 Shen, C., Schulzrinne, H. G., 2006. A voip privacy mechanism and its ap-
10 plication in voip peering for voice service provider topology and identity
11 hiding. Tech. rep., Department of Computer Science, Columbia University.
12 URL <http://hdl.handle.net/10022/AC:P:29476>
13
14
15 Shimaoka, M., Hastings, N., Nielsen, R., July 2008. Memorandum for multi-
16 domain public key infrastructure interoperability. IETF RFC 5217.
17
18 Simon, D., Aboba, B., Hurst, R., March 2008. The eap-tls authentication
19 protocol. IETF RFC 5216.
20
21
22 Such, J. M., Espinosa, A., Garcia-Fornes, A., Botti, V., 2011. Partial identi-
23 ties as a foundation for trust and reputation. *Engineering Applications of*
24 *Artificial Intelligence* 24 (7), 1128–1136.
25
26
27 Tene, O., 2011. Privacy: The new generations. *International Data Privacy*
28 *Law* 1 (1), 15–27.
29 URL <http://idpl.oxfordjournals.org/content/1/1/15.full>
30
31
32 TorFone, 2013. Tor fone: p2p secure and anonymous voip tool. V1.1b
33 (01.06.13).
34 URL <http://torfone.org/>
35
36
37 Tschofenig, H., July 2010. Federated authentication beyond the web: Prob-
38 lem statement and requirements. IETF ABFAB working group.
39 URL <http://tools.ietf.org/html/draft-tschofenig-moonshot-ps-01>
40
41
42 Wasserman, M., Hartman, S., Feb. 2014. Application bridging for federation
43 beyond the web (abfab) trust router protocol. IETF Internet-Draft.
44 URL <http://tools.ietf.org/search/draft-mrw-abfab-trust-router-02>
45
46
47 Winter, S., Salowey, J., 2013. Update to the eap applicability statement for
48 abfab. IETF Internet Draft.
49 URL <http://tools.ietf.org/html/draft-ietf-abfab-eapapplicability-06>
50
51
52 Wolff, J., 2013. Application-layer design patterns for accountable-
53 anonymous online identities. *Telecommunications Policy* In Press, Ac-
54 cepted Manuscript, –.
55
56
57 Zhang, G., Fischer-Hubner, S., 2013. A survey on anonymous voice over ip
58 communication: Attacks and defenses. *Electronic Commerce Research*, –.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Zhu, L., Leach, P., Hartman, S., April 2011. Anonymity support for kerberos.
IETF RFC 6112.

Zhu, L., Tung, B., June 2006. Public key cryptography for initial authentication in kerberos (pkinit). IETF RFC 4556.

D
R
A
F
T