

Elsevier Editorial System(tm) for Journal of Systems Architecture

Manuscript Draft

Manuscript Number:

Title: Authentication in Stealth Distributed Hash Tables

Article Type: Special Issue: Euromicro SEAA 2006

Keywords: Distributed Hash Table; Peer-to-Peer; Security; Public Key Infrastructure; Stealth DHT

Corresponding Author: Mr Andrew MacQuire, B.Sc (Hons.)

Corresponding Author's Institution: Lancaster University

First Author: Andrew MacQuire, B.Sc (Hons.)

Order of Authors: Andrew MacQuire, B.Sc (Hons.); Andrew Brampton, B.Sc (Hons.); Idris A Rai, Ph. D; Nicholas J Race, Ph. D; Laurent Mathy, Ph. D

Abstract: Most existing DHT algorithms assume that all nodes have equal capabilities. This assumption has previously been shown to be untrue in real deployments, where the heterogeneity of nodes can actually have a detrimental effect upon performance. We now acknowledge that nodes on the same overlay may also differ in terms of their trustworthiness. However, implementing and enforcing security policies in a network where all nodes are treated equally is a non-trivial task. We therefore extend our previous work on Stealth DHTs to consider the differentiation of nodes based on their trustworthiness rather than their capabilities alone.

# Authentication in Stealth Distributed Hash Tables

Andrew MacQuire      Andrew Brampton      Idris A. Rai  
Nicholas J. P. Race      Laurent Mathy

*Computing Department  
Lancaster University*

*{macquire,brampton,rai,race,laurent}@comp.lancs.ac.uk*

---

## Abstract

Most existing DHT algorithms assume that all nodes have equal capabilities. This assumption has previously been shown to be untrue in real deployments, where the heterogeneity of nodes can actually have a detrimental effect upon performance. We now acknowledge that nodes on the same overlay may also differ in terms of their trustworthiness. However, implementing and enforcing security policies in a network where all nodes are treated equally is a non-trivial task. We therefore extend our previous work on Stealth DHTs to consider the differentiation of nodes based on their trustworthiness rather than their capabilities alone.

*Key words:* Distributed Hash Table, Peer-to-Peer, Security, Public Key Infrastructure, Stealth DHT

---

## 1. Introduction

Distributed Hash Tables (DHTs) have been shown to be a useful form of decentralised, structured peer-to-peer overlay [24,29,31,20]. They allow for the provision of simple hash table functionality – that is, the ability to *put* and *get* pieces of data indexed via hash codes – across multiple nodes in a scalable and resilient fashion. Primarily, DHTs have been used as location substrates for many varied applications, such as large-scale file storage [9,7,12] and multicast [25,34], amongst others.

Theoretically, most DHTs consist of numerous nodes which organise themselves and behave in a well defined manner. Each node is associated with a unique identifier (ID), randomly selected from a large, sparsely-populated address space. When an object is *put* into a DHT, its contents are hashed in some way as to produce an identifier which also maps into this address space. The originating node then routes the object to the local node that it knows to

have the closest-matching identifier. If the recipient knows of an even closer node, it then forwards the data on. Eventually, the object reaches a node which does not know of any closer match. This node is considered the final destination, and must then take responsibility for the storage and/or any other handling of the object. Any individual who then wishes to *get* the same object at a later date can then do so based on the knowledge of the object's hash alone, as any message routed to the same hash should arrive at the same node that the object originally reached. Of course, this can only occur if all nodes follow the DHT protocol correctly.

Unfortunately, in a real-world deployment it would be naïve to assume that all nodes can be relied upon to conform to any prescribed behaviour. Without appropriate security policies in place, numerous problems may exist for public DHTs. For instance, a malicious node may examine, alter or deliberately drop messages passed through it (*i.e.* a sniffing, man-in-the-middle or denial of service

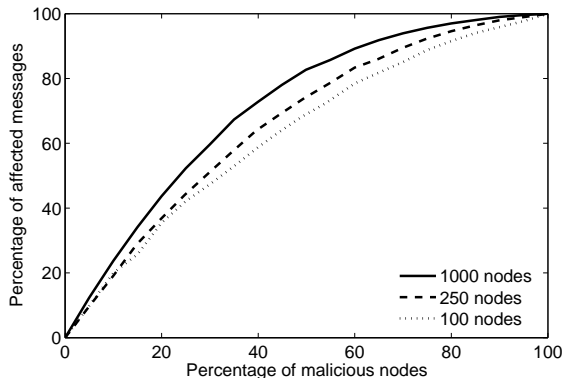


Fig. 1. The effect of malicious nodes forwarding messages

attack). Using the simulation platform described in Section 5 we can demonstrate how a relatively small number of untrustworthy nodes can have a large impact on a generic DHT. Fig. 1 shows the percentage of messages which would be affected as the corresponding percentage of malicious nodes increases. Observe that when a quarter of a 1,000 node DHT’s population is malicious, over half of the messages sent are subject to attack.

The ability to inject unsolicited messages into the DHT, or to alter those in transit, also allows untrustworthy nodes to corrupt the routing tables of others. Possible consequences of such actions could be legitimate nodes being denied service, or malicious nodes improving their standing on the network at the expense of others. In terms of content in the DHT, if all nodes are allowed to perform *put* operations, then the pollution of the system with unwanted or illegal data may become an issue. Furthermore, existing DHT algorithms used in peer-to-peer networks provide ideal environments for users to infringe intellectual property law because they allow users to easily duplicate copyrighted works or distribute illegal content. This is a major factor that hinders commercial use of existing DHT architectures, especially now as peer-to-peer application/service providers are increasingly liable for illegal activity carried out on their networks [11].

To avoid these security problems, it is clear that some method for handling untrustworthy nodes in DHTs is required, especially if DHTs are to be used commercially. Sadly, most traditional DHTs make the assumption of homogeneity amongst peers, treating them all as equal. This means that untrustworthy nodes are regrettably granted exactly the same privileges and responsibilities as the trustworthiness. A common approach to solving this problem

is to make use of an authentication mechanism to ensure that only trustworthy nodes are allowed to join the DHT. However, determining the veracity of a node can be a difficult process, and simply alienating all those who cannot prove themselves trustworthy may be unwise. Instead, such nodes should still be allowed to make use of the DHT, but in a limited capacity.

To this end, a means of separating sets of nodes on the same DHT proves necessary. We therefore demonstrate how our recently proposed Stealth DHT concept [3] can be used to provide this precise functionality with almost any existing DHT algorithm. Our original evaluation covered the benefits of this approach from a performance standpoint, whereas we now consider the advantages gained in terms of security.

The separation between nodes that a Stealth DHT provides can be exploited to enable secure content distribution, as when coupled with a suitable authentication mechanism, it aids in returning network control to the service provider. To clarify, the ability to create fine-grained permissions for nodes on the DHT means that security policies are much easier to enforce. For example, by having only authorised nodes storing and retrieving content, Stealth DHTs can ensure that only legitimate and useful content is served, thus providing a platform for Digital Rights Management (DRM) in peer-to-peer networks, as well as aiding in the prevention of pollution attacks.

The remainder of this paper is structured as follows: Section 2 gives a brief overview of the differences between traditional and Stealth DHTs. Section 3 discusses the structure of a typical Public Key Infrastructure (PKI). Section 4 then explains how a Stealth DHT and a PKI may interoperate. Section 5 then highlights and evaluates a number of implementation concerns for such a system, with further discussion of optional functionality in Section 6. Section 7 goes on to examine related work in the field of DHT authentication, and finally Section 8 concludes the paper.

## 2. Overview of a Stealth DHT

The join process for most DHT implementations involves a node first gathering state. Usually, this is achieved by routing a join message addressed to its own ID into the DHT via a bootstrap node (an already-connected node discovered through some alternate mechanism). Nodes along the message’s

path then reply directly with relevant routing information, as to allow the joining node to construct its routing tables. Once the joining node receives notification that its message has reached its destination, it announces its presence on the network so that other nodes may route messages through it.

Stealth DHTs [3] modify this procedure slightly to create two types of nodes on the network: *Service* and *Stealth*. Service nodes provide the routing infrastructure for the overlay, whereas stealth nodes communicate with and through service nodes only. This separation is achieved by halting the join procedure for stealth nodes after they have gathered state, but before they announce their presence on the DHT. The resultant effect is that stealth nodes do not appear in any routing tables, and thus are not used to forward any messages or store any keys. Therefore, they are incapable of interfering with message delivery or object storage in any direct manner. The routing data gathered by stealth nodes is used only for selecting the locally-optimal location to forward their own requests to, thus improving routing performance while removing the possibility of a single point of failure that many similar DHT super-peer schemes suffer from [16,33].

It is important to note that the assignment of the stealth and service node roles is application-dependent, and is not prescribed or constrained by the Stealth DHT itself. However, a Stealth DHT provides an individual or a service provider with the control to command such assignment. Therefore, from a security perspective, and since service nodes are responsible for handling all messages, they should consist of verifiably trustworthy machines. Conversely, any nodes which are potentially untrustworthy should be forced to join as stealth nodes, prohibiting them from interfering with DHT operations to any extent.

In contrast with traditional DHTs, the distinction between trustworthy and untrustworthy nodes provided by a Stealth DHT results in an architecture where the implementation of security policies is more straightforward. In addition, as service nodes never retain any knowledge of stealth nodes, Stealth DHTs are poised to offer significant advantages from a security perspective.

For example, when a stealth node joins or leaves the network, no service nodes need to update their routing tables, which prevents them from being affected by stealth nodes churning. This is especially important from a security standpoint, as the effects of heavy churn have been identified as being

particularly harmful to many DHT implementations [21,13]. Malicious individuals have accordingly used this knowledge as a means of facilitating Distributed Denial of Service (DDoS) attacks. By making numerous nodes under their control rapidly rejoin DHTs, an attacker can cause floods of maintenance messages as nodes struggle to keep their routing tables up to date, resulting in inconsistent routing and overloaded nodes. If, however, a Stealth DHT is used where the potentially untrustworthy are forced to join as stealth nodes, no routing table updates and so no maintenance messages are required. The inevitable cost of such an advantage comes in the form of increased stress upon the service nodes, although these are assumed to be relatively powerful machines capable of handling such load.

Of course, in a system such as this there still must be a means of ensuring that stealth nodes cannot masquerade as service nodes (*e.g.* by simply announcing their presence to other service nodes). This can be achieved in a Stealth DHT through the use of an appropriate authentication scheme to effectively enforce the separation between node types. Accordingly, this is the focus of this work, wherein we discuss how an authentication scheme based on a Public Key Infrastructure can be implemented in a Stealth DHT.

### 3. Overview of a Public Key Infrastructure

A Public Key Infrastructure (PKI) is a security platform which allows multiple users who have not previously exchanged any information to validate each other's identities, be sure of message integrity and even set up confidential communication. This is usually achieved via digital certification signed by mutually trusted third parties, where the certificates themselves are used to verify the identity of their owner through public/private key cryptography.

A typical PKI is composed of several logically separate entities, although the functionality offered by each may be contained within a single physical machine:

A **Registration Authority (RA)** is a trusted entity which acts as the first point of contact for an individual requesting certification. The RA is used to check the requestor's supplied credentials and, if deemed valid, pass them on to the Certification Authority.

A **Certification Authority (CA)** is a trusted

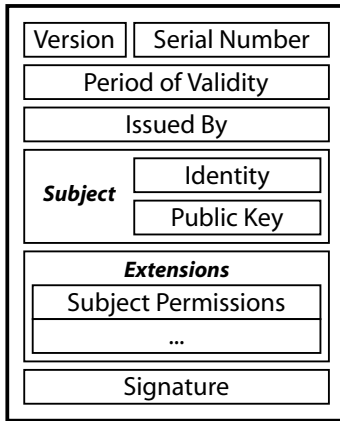


Fig. 2. Certificate Format

entity responsible for the creation and, if supported, revocation of certificates. As it is a mutually trusted third party, individuals may authenticate each other with confidence if they sign their messages using a certificate verifiably issued by a CA.

A **Certificate Repository** (CR) simply acts as a database of existing certificates. A CR need not be a trusted entity as the certificates it holds are immutable; if any attempt is made to alter an existing certificate, the digital signature will no longer match the contents. Note that if nodes are made responsible for the storage and dissemination of their own certificates, a dedicated CR may be redundant.

If supported, the CR also contains the **Certificate Revocation List** (CRL), indicating which certificates in the database have been forcibly revoked. Certificate revocation, however, is often an unsupported feature in actual PKIs due to implementation difficulties; an issue discussed further in Section 5.5.

Many PKI implementations use certificates which conform to the ITU-T X.509 standard [1], a simplified version of which can be seen in Fig. 2. Each certificate contains a version record for compatibility reasons, as well as a serial number to aid in certificate management. The certificate should also have two discrete dates associated with it, indicating its period of validity.

The most important element of any digital certificate, however, is the subject. That is, the individual or organisation whose identity the certificate may be used to authenticate. Typically this information is comprised of the owner’s name and any other pertinent information (address, organisation name *etc.*). The owner’s public key is also included in this section of the certificate. This key is crypto-

graphically paired with a corresponding private key that each individual must keep secret, as it serves as their means of creating digital signatures, as well as decrypting any messages encrypted with their public key.

Finally, the certificate must indicate the authority which issued it, and must also contain the signature of that authority. The key point here is that once the certificate has been signed in this manner, the data contained within cannot change without invalidating the signature. Note that certificates may also contain optional extension fields, allowing for application-specific additions.

It is also notable that as any individual who owns a signed certificate from a higher authority may also sign certificates themselves, lengthy certification hierarchies often exist. As requesting and verifying each level of certification separately may prove to be time-consuming, many PKIs allow for the *chaining* of certificates. That is, all certificates up to the initial, self-signed certificate (created by some intrinsically trustworthy entity) are included in a single collection. While such a certificate-chain will obviously be larger than any single certificate, it intuitively reduces overhead if verification up to the highest level is known to be required.

#### 4. Authentication in a Stealth DHT

A straightforward approach to implementing a PKI on a Stealth DHT is to require each service node to be issued with a certificate, as to prove its authority. In the simplest case, stealth nodes would not require certification. If, however, restrictions were placed on any of the operations that stealth nodes could perform, certification would be required to prove that a given stealth node is authorised to perform a particular operation.

The certificate extension fields may therefore contain a list of authorised operations that its owner may carry out (see Fig. 2). Simple examples could be the right to join as a service node, or the right to put content into the network. It would then be mandatory for the relevant DHT messages to contain a field which identifies the node’s certificate in some way, as well as a digital signature to prove that the same node was indeed the message’s creator. The node’s certificate or certificate-chain could also be attached to the message to aid in the authentication process.

#### 4.1. PKI Composition

There are two broad approaches to supplying the PKI's constituent elements for the Stealth DHT: *external* and *internal*. In the former, the RA, CA and CR all exist entirely separately to any part of the Stealth DHT itself, whereas in the latter the functionality is provided by a set of operations on the DHT, supported by some subset of the Stealth DHT's service nodes. This subset may simply be a single, well-known service node (*centralised PKI*) providing the entire PKI functionality for the Stealth DHT, or in contrast it may consist of several or all the existing service nodes (*distributed PKI*).

Exactly how the PKI elements are organised is entirely application-specific, and internal/external elements may be mixed. For example, as the RA and CA must be trusted entities, they may consist of a small number of highly-trusted service nodes with well-known identifiers. The CR, on the other hand, does not require a high level of trust due to the immutability of digital certificates. As a result, it could consist of several or all service nodes on the DHT, with certificates being hashed and stored as if they were normal DHT keys.

#### 4.2. Joining the DHT

The first step for a new user to take will typically be certificate acquisition. The user must therefore generate a public/private key pair, and then pass his or her public key along with any requested proof of identity to the Registration Authority. Following successful verification of these credentials by the RA, they are then passed to the Certification Authority. The CA then creates the certificate, signs it, and passes it back to the user via the RA. It may also be passed to a Certificate Repository, if necessary.

In a Stealth DHT with an external PKI that requires authenticated join operations, the user can simply send a join message containing his or her certificate to a suitable DHT bootstrap node, with the separate PKI server(s) being contacted as required. The same process is required for a Stealth DHT with an internal PKI, although the components of the PKI system are contained within the DHT instead. Note that in the latter case a new, uncertified user could also simply pass all his or her relevant details to a bootstrap node as their join message, with no need for any further action on their part. As all elements of the PKI are contained within the DHT,

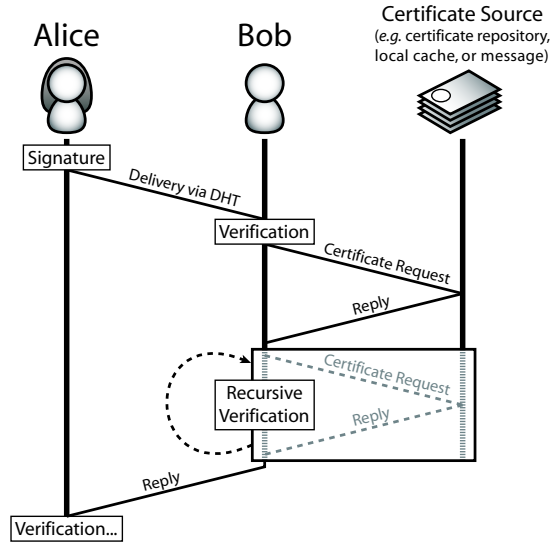


Fig. 3. Sequence of events for message authentication

a certificate can automatically be returned to them whilst an appropriate DHT join message with the newly created certificate attached is simultaneously forwarded, thus minimising join delay. In all cases, the DHT ID of the joining node is irrelevant (unless associated with their certificate); the real identity of a user is always determined by the certificate they own.

#### 4.3. Sending Messages

Following a successful join, nodes may then communicate as normal over the Stealth DHT, authenticating each other as necessary. As an example of this, assume we have two users, Alice and Bob, a stealth node and a service node respectively who have joined a Stealth DHT with an internal PKI. Alice wishes to send a message to Bob, who requires that messages be authenticated. The correct procedure (as shown in Fig. 3) would therefore be as follows:

Alice first creates a message, signs it, and delivers it to Bob via the DHT. Bob can then verify the signature, and thus the message integrity, using Alice's certificate. He may have acquired Alice's certificate from his certificate cache, a Certificate Repository, or from within the message itself. Bob can then recursively verify the issuers within the certificate chain, starting with Alice's certificate. This process continues until Bob reaches a certificate that he intrinsically trusts. At this point, the authentication

is complete, and Bob can continue to handle Alice’s message appropriately. Naturally, if the certificate chain does not eventually lead to a certificate that Bob trusts, his attempt to authenticate Alice fails. Following any reply from Bob, Alice may perform the same process on his message to authenticate his identity, but only if mutual authentication is required.

Further to this, if Alice and Bob wish to ensure that their messages are kept confidential from even the service nodes, they can simply use the public keys contained within each other’s certificates to encrypt the messages’ contents. To clarify, if Alice wants to send sensitive data to Bob, then she first uses the CR to retrieve his certificate beforehand. Following this, she uses his public key to encrypt the message contents and her own private key to sign the message. Only Bob’s private key can decrypt data encrypted in this manner, so Alice can be sure that only Bob is able to understand the message contents. Further to this, her signature ensures that Bob can be sure the message came from Alice, and that it was not tampered with.

Note that in an internal PKI, if the Certificate Repository functionality is spread across many service nodes, and certificate chains are not included in messages, users performing *get* operations may experience longer retrieval delays due to the need for the relevant certificates to also be requested from the DHT. This increased overhead should, however, be weighed against the cost of using a centralised PKI, which potentially represents a single point of failure for all nodes on the network.

## 5. Implementation Considerations

As Stealth DHTs may be tailored to a wide range of applications based on the assignment of roles to nodes, it follows that there are numerous application-specific decisions to be made regarding the implementation of any security system, as discussed in the following sections.

For evaluation purposes, we used our own discrete-event packet-level simulator based on a Stealth DHT implementation of Pastry [24] (validated in our previous work [3]). Each simulation was performed several times on a GT-ITM [4] generated transit-stub topology of 1,000 routers, with 4% transit nodes. Service, stealth and normal Pastry nodes were connected to this topology in a random fashion.

Results from our real-world C++ implementation of both Pastry and a Pastry Stealth DHT are also included in the following sections. Data was obtained from a randomly selected set of active PlanetLab nodes [18], which were connected to the DHT in a random fashion. Multiple iterations were run, with the network being given suitable time to reach a steady-state in all cases. The algorithm used in all authentication operations was the OpenSSL 0.9.8d implementation of RSA, with 1024 bit keys [10,22].

### 5.1. Certification Hierarchy

Some thought should be placed into how the certification structure is organised within the PKI. The simplest approach for a Stealth DHT could be to have a single globally trusted key, used to sign certificates for service nodes only. However, users may require a more complex hierarchy for economic, political or security reasons; again, this is an entirely application-specific issue. A possible example could be that each department within a typical commercial organisation is granted a certificate signed by a single, highly-trusted master key. It may be that the master key needs to be kept physically secure and is therefore inconvenient to access. By introducing this extra level of hierarchy, however, the need for it to be used to sign certificates on a regular basis is negated. This approach also allows for a finer level of control, as if the privileges of an entire department needed to be revoked, for instance, it is simply a matter of invalidating the associated departmental certificate.

### 5.2. Asymmetry vs. Symmetry

There are two obvious manners in which service and stealth nodes may authenticate each other: *asymmetric* (service or stealth node authentication only) or *symmetric* (mutual authentication). To clarify, service node authentication may be required if a stealth node is placing sensitive information into the network and wants to be sure of the recipient. Conversely, stealth node authentication may be required if a service node needs to verify if a stealth node is authorised to retrieve sensitive information from it. Finally, mutual authentication may be required if a combination of these factors arises. By taking this approach of authenticating the nodes only as required, processing and messaging overhead can therefore be minimised.

### 5.3. Authentication Granularity

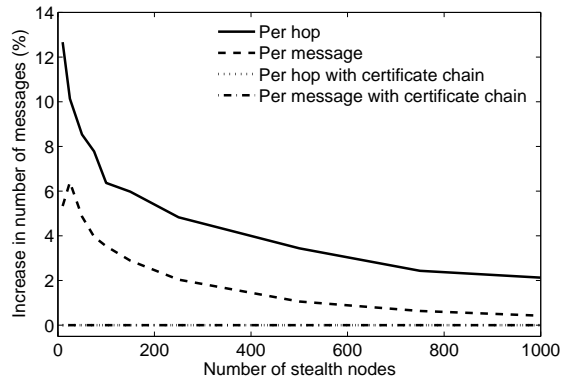
Exactly how authentication is performed on the DHT presents an issue of balance between overhead and security. A fine-grained approach would be to verify messages on a per-hop basis. For most DHTs, this results in all required authentication operations (such as retrieving appropriate certificates or checking revocation lists) being performed on average  $\log N$  times for each message, where  $N$  is the number of service nodes in the network. However, it also results in any invalid messages being dropped almost immediately, making it difficult for unauthenticated malicious nodes to get service nodes to pass invalid messages around; a tactic often used in DHT denial of service attacks, typically resulting in the network becoming overloaded with useless traffic.

Considering a slightly coarser approach, if service nodes were to simply validate messages upon receipt at their final destination, then all associated authentication operations need only be performed once per message. Obviously this means that any messages from unauthorised nodes may be unnecessarily routed by multiple service nodes on their journey, consuming bandwidth and processing time. However, it also means that the number of authentication operations performed may be significantly reduced.

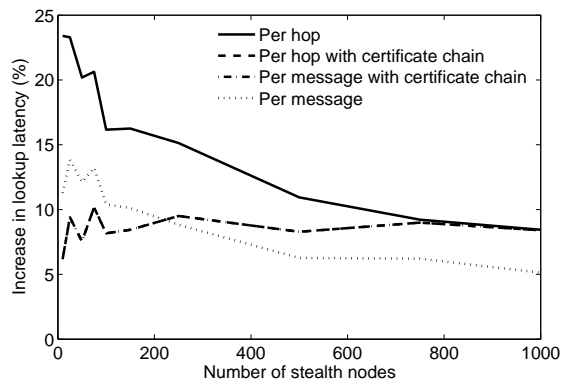
Beyond per-message authentication, service nodes may make use of the coarse concept of “sessions”. In other words, a typical stealth node may be permitted to consume a certain amount of resources, make use of a given service or be active on the DHT for a pre-determined length of time. Examples here could include a node being permitted to send or receive a set number of messages on the network, being given the ability to download a particular number of pieces of content or being provided with a “day-pass” to use the DHT, respectively.

Session-based authentication therefore requires that some state be stored and validated for stealth nodes. If the session is based on a length of time, then this can be as simple as issuing a certificate with a corresponding period of validity. Otherwise, it is likely that a more complex system is required, such as an accounting mechanism, as discussed in Section 6.1 or a session authentication protocol such as Kerberos [17].

Figs. 4 show results for simulations conducted with a fixed number of 100 service nodes, where



(a) Increase in the number of messages



(b) Increase in lookup latency

Fig. 4. Stealth DHT with 100 services nodes using an internal, fully distributed PKI relative to a Stealth DHT with no authentication. Note that the certificate-chain simulations are overlaid in these figures.

the number of stealth nodes was varied between 10 and 1,000. Every service node in the DHT initially held its own certificate as well as several content keys. During the simulation itself, stealth nodes performed random *get* requests for these keys at regular intervals. Authentication operations were performed at the source and destination in the per-message results and also performed at each intermediate node in the per-hop results. For both cases, simulations were run with and without certificate chains contained within messages.

Fig. 4(a) shows that both the per-hop and per-message authentication schemes result in no increase in the overall number of messages when certificate chaining is used. This is because all the certificates required to fully authenticate a message’s sender are included within the message itself, thus increasing the average message size, but resulting in no further requests to the CR.



On the other hand, an increase in the number of messages does exist when certificate chains are not used. This increase is due to the distributed nature of the Certificate Repository in an internal PKI; to acquire a certificate in order to perform authentication, a node must retrieve it from the service nodes through further DHT queries. Expectedly, per-hop authentication results in markedly more messages than per-message authentication. Also, as the number of stealth nodes increases, the percentage increase in the number of messages relative to a system without authentication falls. The reason for this is that the increased number of stealth nodes results in an accordingly increased number of requests for certificates. As all nodes cache certificates upon receipt, there is no need for certificates to be re-acquired.

Fig. 4(b) shows how lookup latency (the time elapsed between a node requesting, receiving and fully verifying a key from the DHT.) is affected by these factors. As expected, the cases which involved querying the network (*i.e.* those without certificate chains) result in increased lookup latencies. Note, however, that the cases with certificate chains also incur increased lookup latencies despite the lack of extra authentication messages. This is attributed to the larger average message size that occurs as a result of including certificate chains within messages.

The same conclusions can be drawn from Fig. 5, which shows the CDF of lookup latencies for our implementation running on PlanetLab. Simple certification hierarchies of two and three levels were used to ensure that certificate requests to an external server either were or were not generated (“requests” and “no requests” on the figure, respectively). We again witness that when requests for certificates are required, lookup latency is increased relative to a system without authentication. The larger message size required for authenticating messages also results in the “no request” scenario exhibiting slower lookups, although the difference is minor for the majority of cases.

As seen in Fig. 4(a), the relative increase in the number of required authentication messages decreases with larger numbers of stealth nodes. Fig. 4(b) therefore displays a correlated reduction in lookup latency. As the cases with certificate chains do not generate extra authentication message, they remain unaffected by the number of stealth nodes in the DHT.

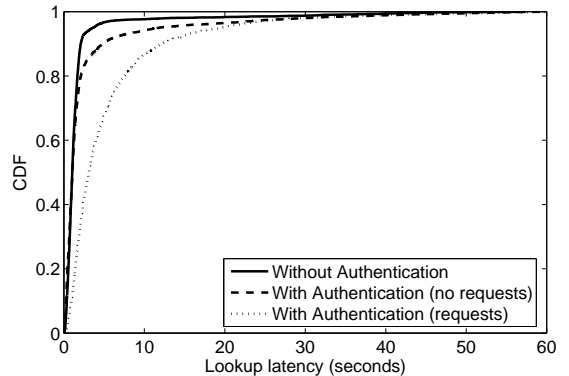


Fig. 5. Impact of per-hop authentication upon lookup latency with/without certificate requests

#### 5.4. Processing Overhead

An inevitable overhead involved with providing a PKI-based authentication system is the increase in processing time required for signing and verifying certificates. The effect this might have on a DHT is difficult to realistically quantify through simulation alone, and so we investigated the impact of using authentication on the performance of our real-world implementation.

As noted in Section 5.3, Fig. 5 shows the CDF of lookup latencies for our implementation running on PlanetLab, with and without authentication. From the figure, it is clear that approximately a quarter of the lookups made were noticeably slower when per-hop authentication (without requests) was used, relative to a DHT with no authentication mechanism. However, upon varying the authentication granularity to per-message (not shown), no improvement was witnessed. This infers that processing overhead may not be a major factor in the performance decrease, as the amount of processing in per-message authentication would be greatly decreased relative to per-hop.

Further investigation involved repeatedly benchmarking the authentication code on the PlanetLab nodes used, and it was found that over 90% were able to sign a message in less than 100 milliseconds (and 70% less than 50 milliseconds). In comparison with network latency, it therefore seems unlikely that processing overhead played a significant role in the performance decrease observed. Also, note that many of the values in this section are somewhat indicative of the high load on PlanetLab: the same test performed on a typical unloaded PC completed in less than 3 milliseconds on average.

It appears the actual reason for the longer lookup latencies is simply the delay involved in delivering the extra authentication data. Without authentication, messages were around 256 bytes in length; with authentication, messages were twice as large at 512 bytes. It is clear therefore, that a lack of processing power is not a major problem; instead, minimising the need to transmit certificates (*e.g.* by using suitable certification hierarchies) and ensuring network links are suitably provisioned are more important issues.

### 5.5. Certificate Revocation

Invalidating a given certificate at an arbitrary point in time within a PKI structure (*i.e.* revoking it) is traditionally a difficult problem, especially in distributed environments. However, the need for the ability to revoke certificate often outweighs the cost of any associated overhead. In other words, being able to remove a node from a network before it can do any lasting damage may be worth the extra associated costs. There are a number of possible methods for certificate revocation, each with advantages and disadvantages (see [32]).

An obvious solution is for stealth nodes to be simply issued with certificates with short expiration times, resulting in a quasi-revocation scheme in which a CA can simply refuse to re-issue a given certificate if it wishes to remove a given entity's privileges within the system. Unfortunately, this intuitively has a high maintenance overhead, as every individual that continues to exist on the network will require a new certificate to be generated and issued at a regular interval.

Another common approach is to use a CRL, or Certificate Revocation List. Any node verifying a certificate must then check the list as part of its normal authentication process. The list itself may be stored in one of a number of ways. For instance, and as with many of the other issues considered so far, a centralised server could be used. Again, the problems common to this sort of approach are that it creates a central point of failure, and could potentially result in the overloading of the server if many requests are regularly made. A distributed approach amongst service nodes could potentially be used to alleviate the load placed upon any one node, but this results in increased messaging overhead due to the added complexity of maintaining and retrieving the list. It is important to note that the list could

also be kept in its entirety on multiple nodes instead of split between them (*i.e.* replication *vs.* division, respectively).

There also exist several approaches for the manner in which nodes access such lists. In general terms, these fall into the *pull* and *push* distribution models. As an example of the pull model, nodes may poll the holder(s) of the CRL at regular intervals, or even for every transaction, depending on the level of granularity required. An example of the push model would be the Certificate Revocation List being broadcast to all nodes upon any update, or at a regular interval. Obviously, such an approach may result in a great deal of initialisation overhead if a large number of nodes which require the CRL exist; multicast overlays may have to be built, and ideally with some method of guaranteeing delivery. For small numbers of nodes, however, the cost of broadcasting updates and/or occasionally broadcasting the full list may be lower than having nodes repeatedly request it unnecessarily.

## 6. Optional Considerations

Beyond the mandatory decisions regarding the structure of the Stealth DHT PKI, there are also a number of additional considerations that may be required for certain applications, as discussed in the following sections.

### 6.1. Permissions Management

Managing the permissions of nodes on the network is yet another issue with multiple solutions. One possibility is to simply place them within each node's corresponding certificate, but this has two notable associated issues. Firstly, certificates are immutable after they are signed, so altering the permissions for a node would require a certificate to be re-issued. Secondly, the average DHT message size would have to increase to accommodate the larger certificates; exactly how large they are depends on the volume of data required to represent all the relevant details. However, such an approach does mean that the relevant permissions are immediately available to any node receiving a message containing the originating node's certificate.

An alternative approach would be to store permissions data within the network somehow, thereby avoiding the lack of flexibility and larger message sizes associated with storing them within certifi-

cates. The inevitable cost, however, arises in the form of extra messaging overhead; every time a node needs to check an individual's permissions, they would have to look up and also validate the relevant data.

The permissions for a given node in a system such as those discussed may require that some state is maintained for each individual. For instance, this may be based on a record of how much network resources the node has consumed. To provide such functionality, an accounting mechanism of some sort is therefore required.

As with most of the issues discussed so far, there are numerous possible approaches to the problem of providing an accounting mechanism for nodes in a Stealth DHT, each with advantages and disadvantages. Nevertheless, the common goal of each is identical: to provide a system which offers ACID properties. These are *Atomicity* (an action is either performed completely or not at all), *Consistency* (an action cannot place the data in an invalid state), *Isolation* (actions cannot interfere with each other's intermediate states) and *Durability* (the result of an action will persist).

The most straightforward approach is perhaps to just have a simple centralised server handle all transactions. This means messaging overhead is minimal, although it also means that a single point of failure exists for the system. Furthermore, dependent on the number of nodes considered, the server may be placed under load beyond its capabilities.

Another possibility would be to improve load-balancing by distributing the transaction processing across multiple nodes. For instance, this could be achieved by mapping a given node's state to a key in the DHT. However, it is important to note that without replication, such a system would be prone to losing data from time to time. With replication, however, the management overhead required to maintain the aforementioned ACID properties will intuitively increase due to the need for agreement amongst the replicated nodes.

The classical approaches for distributed accounting are divided into a few main categories, *local accounting*, *quorum-based* [6] and *token-based* [15,30,23]. Local accounting refers to the case when nodes do not collaboratively maintain a count, but only keep values locally instead. In contrast, quorum-based systems require a group of nodes to maintain the count in a cooperative fashion. Before the count is modified the quorum must reach a consensus on its old and new value. The more

nodes that exist in the quorum the more robust the system is to malicious nodes and network failures.

Token based systems may also use similar approaches, but the key difference is that each count is represented by many immutable tokens. Each token represents a currency that nodes can exchange for service. These tokens are cryptographically signed to restrict who can use them. Nodes acting as "banks" may be used to store and certify the tokens to ensure that double spending does not take place. Tokens have the benefit that they can either be self-issued [30], or can be issued by the system [23].

Variations on mechanisms such as these may also be used to provide logging functionality in the DHT, thus making users accountable for their actions.

## 6.2. Node Promotion

A key function of a Stealth DHT with authentication support is the strong separation between stealth and service nodes. Sometimes, however, it may be useful to allow nodes to alter their roles. For instance, if a Stealth DHT with a small number of service nodes is heavily loaded, the ability to 'promote' stealth nodes to a service role for offloading purposes could be advantageous.

In most cases, it would perhaps be unwise to offer a stealth node full privileges outright; without a careful selection process and suitable permissions management, the potential for abuse would be high. Determining suitable stealth nodes would most likely be a two stage process: first, discovering which nodes have the required resources available, and secondly, ascertaining which nodes within that set can be classed as trustworthy.

Finding stealth nodes with spare resources would not necessarily be difficult. On a heavily loaded DHT, service and stealth nodes would be communicating often, and so service nodes could simply attach a request for support to any replies they send. A suitable stealth node could then reply to this request, offering its spare capacity in whatever form it may take. Stealth nodes may also offer their services, either due to incentives placed by the service provider, or for altruistic reasons.

Exactly what resources are being offered has an important bearing on how rigorous the service node's subsequent examination of the stealth node should be. As a simple example, storing replicated certificates requires little to no trust at all; the certificates cannot be maliciously altered without it

being obvious, and should the stealth node disappear, the original certificate is not lost. In contrast, if a stealth node were to be enlisted to store the only copy of a piece of unprotected content, it would have to be trusted not to alter it, or to disconnect without replicating it. Note that neither of these examples actually require the stealth node to be part of the DHT's routing, as the stealth node could simply register with a service node as auxiliary storage space.

In the case of a former stealth node becoming involved in forwarding messages for the DHT, even greater care is required when validating the node's trustworthiness. For instance, if a malicious node is only used for storing protected data, the worst it can do is deny service for that data alone. If, however, the same malicious node is involved in DHT routing, it could supply legitimate service nodes with false routing data, potentially damaging the entire DHT.

Compromises are therefore required, and thus the concept of promotion is inextricably linked with the previous discussion on managing node permissions (see Section 6.1). Furthermore, it is important that any new privileges given to a stealth node can be revoked easily, an issue discussed in Section 5.5.

### 6.3. Digital Rights Management

As has been previously noted, a Stealth DHT with authentication support may prove to be a useful platform for content distribution. However, while a Stealth DHT as discussed thus far can ensure the content is delivered safely, commercial services may also require that the data be placed under certain restrictions. A typical example would be copy protection, where users are prevented from accessing content they are not licensed to use.

A common, generalised approach to this problem involves the use of licensing servers. All content is encrypted during the delivery process, and the licenses granted to users via the servers act as decryption keys. In some cases, every time the content is played back, a license has to be acquired. Naturally, this means that the license servers have to be reliable to ensure user satisfaction.

A natural extension for a Stealth DHT serving such content could therefore be to distribute licensing servers across the DHT. Many of the issues discussed thus far regarding Public Key Infrastructures also apply to DRM systems, as PKIs are often used within them. A Stealth DHT could therefore of-

fer both user authentication and content protection within a single system.

## 7. Related Work

Many previous works have discussed the varied problems associated with untrustworthy nodes in DHTs, noting that security is an issue commonly overlooked in algorithm proposals. For example, [27] and [5] both discuss possible DHT attacks and defenses. In the former, Sit and Morris broadly define three types of malicious behaviour: routing, storage and other miscellaneous attacks.

A commonly encountered technique often used in all three categories is the "Sybil" attack, as originally described by Douceur [8]. This refers to the situation when a single malicious node is able to masquerade as multiple distinct entities within the network in order to gain control of a substantial fraction of it. The conclusion is drawn that a suitable defense against such an attack is to have a logically centralised authority which is capable of certifying the identity of nodes in the network. Therefore, this paper can be said to support our approach of implementing a PKI in conjunction with a Stealth DHT.

Routing attacks in a DHT may refer to nodes deliberately providing incorrect lookups or producing incorrect routing updates. More specifically, an example could be the "Eclipse" strategy, as discussed in detail by Singh *et al.* [26]. This involves multiple malicious nodes deliberately attempting to partition peer-to-peer overlays in a form of Distributed Denial of Service (DDoS) attack. Again, the suggestion is made of a certification scheme as a straightforward solution, as with our approach to such a problem. Beyond this, the authors propose defenses such as constraining the entries placed in routing tables [5] or periodically auditing the connectivity of other nodes to detect anomalies which are symptomatic of those conducting such an attack.

Storage attacks may involve behaviour such as nodes refusing to store objects, corrupting them or simply denying their existence. More resourceful attackers may also use multiple malicious nodes to attempt to take control of specific pieces of content. Some may even try to make it impossible to access useful content by flooding the network with useless data [14]. Srivatsa and Liu suggested the approach to obfuscate the location on the DHT of specific keys from those not authorised to access them [28], although the most commonly suggested solution is to

make use of some sort of digital certification scheme, such as the one we have proposed. Furthermore, a Stealth DHT can ensure that potentially untrustworthy nodes are never even part of the DHT which is responsible for storing keys, although they may still access them if authorised to do so.

Of course, DHT-based storage systems have often considered security in their own right. PAST, for instance, uses the concept of “smartcards”, with which users hold associated public/private keys [9]. These smartcards are managed by brokers (trusted third parties). In other words, PAST is yet another system that takes the approach of using a Public Key Infrastructure for security purposes; again, the concept we have suggested and expanded upon in this work.

In terms of miscellaneous attacks, Sit and Morris also noted that an attacker may attempt to conduct a DDoS attack by causing multiple nodes under their control to rapidly join and leave the network, resulting in degradation of DHT performance [21,13]. Possible solutions to this problem are suggested in [5], such as forcing nodes to solve crypto-puzzles before they may join as a means of slowing down attackers attempting to run multiple logical nodes on a single physical machine. However, such an approach merely makes carrying out an attack slightly more difficult. Our Stealth DHT approach, however, means that stealth node churn has a significantly reduced effect on DHT performance relative to nodes churning in traditional DHTs.

Several works have also considered how such authentication systems may be implemented in a physically distributed fashion over peer-to-peer networks. For example, Aberer *et al.* discussed how a completely decentralised PKI based on a statistical approach could be deployed on many traditional DHTs (although they specifically use P-Grid [19]) [2]. The key difference in comparison with our work is that in this case, the authors consider a method that can function with a network consisting entirely of potentially untrustworthy nodes. However, they note that their system breaks down if more than 25% of nodes are actually malicious, and that it may not function with several DHTs, such as CAN or Chord [20,29]. We, however, believe that our system is implementable on any existing DHT, and should function regardless of the percentage of malicious stealth nodes.

## 8. Conclusion

The original goal of our Stealth DHT proposal was to provide a distinction between nodes of greater and lesser capabilities as a means of improving routing performance. Powerful nodes were responsible for handling message forwarding within the DHT, whereas the remaining, weaker nodes simply requested services from them. We have demonstrated that this separation can be extended to incorporate both verifiably trustworthy and potentially untrustworthy nodes. By selectively limiting the privileges of untrustworthy nodes on the network, on an individual basis if required, we can accordingly limit the numerous security problems associated with supplying service to them. By further augmenting our approach with a suitable Public Key Infrastructure to enforce the separation between node types, we have shown how a Stealth DHT can be used to supply a secure, resilient overlay that caters to both trustworthy and untrustworthy nodes simultaneously. Stealth DHTs do not necessarily need to deny access to potentially untrustworthy nodes as opposed to previous approaches that addressed security issues in DHTs. Instead, the operations that such nodes can perform need only be selectively limited. A Stealth DHT coupled with the authentication mechanisms described could therefore form an ideal secure location substrate for the numerous varied applications that DHTs can support.

## References

- [1] ITU-T Recommendation X.509, Information Technology Open Systems Interconnection - The Directory: Authentication Framework (August 1997).
- [2] K. Aberer, A. Datta, M. Hauswirth, A decentralized public key infrastructure for customer-to-customer e-commerce, *International Journal of Business Process Integration and Management* 1 (1) (2005) 26–33.
- [3] A. Brampton, A. MacQuire, I. A. Rai, N. J. P. Race, L. Mathy, Stealth Distributed Hash Table: A robust and flexible super-peered DHT, in: *Proc. of the 2nd Conference on Future Networking Technologies (CoNEXT)*, Lisbon, Portugal, 2006.
- [4] K. L. Calvert, M. B. Doar, E. W. Zegura, Modeling Internet topology, *IEEE Communications Magazine* 35 (6) (1997) 160–163.
- [5] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, D. S. Wallach, Secure routing for structured peer-to-peer overlay networks, in: *Proc. of the 5th Symposium on Operating Systems Design and Implementation (OSDI)*, 2002.

- [6] M. Castro, B. Liskov, Practical byzantine fault-tolerance, in: Proc. of the 3rd Symposium on Operating Systems Design and Implementation (OSDI), 1999.
- [7] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, I. Stoica, Wide-area cooperative storage with CFS, in: Proc. of the 18th ACM Symposium on Operating Systems Principles (SOSP), 2001.
- [8] J. R. Douceur, The Sybil attack, in: Proc. of the 1st International Workshop on Peer-to-Peer Systems (IPTPS), 2002.
- [9] P. Druschel, A. Rowstron, PAST: A large-scale, persistent peer-to-peer storage utility, in: Proc. of the 8th Workshop on Hot Topics in Operating Systems (HotOS), 2001.
- [10] <http://www.openssl.org>, OpenSSL: The open source toolkit for SSL/TLS.
- [11] <http://www.supremecourtus.gov/opinions/04pdf/04-480.pdf>, Supreme Court of the United States, no. 04.480. Argued March 29, 2005. Decided June 27, 2005.
- [12] J. Kubiawicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, B. Zhao, OceanStore: An architecture for global-scale persistent storage, in: Proc. of ACM ASPLOS, 2000.
- [13] J. Li, J. Stribling, T. M. Gil, R. Morris, M. F. Kaashoek, Comparing the performance of distributed hash tables under churn, in: Proc. of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS), 2004.
- [14] J. Liang, N. Naoumov, K. W. Ross, The index poisoning attack in P2P file sharing systems, in: Proc. of IEEE INFOCOM, 2006.
- [15] N. Liebau, V. Darlagiannis, A. Mauthe, R. Steinmetz, A token-based accounting scheme for p2p-systems, Tech. Rep. TR-2004-05, Technische Universitt Darmstadt (January 2004).
- [16] A. T. Mizrak, Y. Cheng, V. Kumar, S. Savage, Structured superpeers: Leveraging heterogeneity to provide constant-time lookup, in: Proc. of the 3rd IEEE Workshop on Internet Applications (WIAPP), 2003.
- [17] B. C. Neuman, T. Ts'o, Kerberos: An authentication service for computer networks, IEEE Communications 32 (9) (1994) 33–38.
- [18] L. Peterson, D. Culler, T. Anderson, T. Roscoe, A blueprint for introducing disruptive technology into the Internet, in: Proc. of the 1st Workshop on Hot Topics in Networks (HotNets-I), 2002.
- [19] C. G. Plaxton, R. Rajaraman, A. W. Richa, Accessing nearby copies of replicated objects in a distributed environment, in: Proc. of the 9th Annual ACM Symposium on Parallel Algorithms and Architectures, 1997.
- [20] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Shenker, A scalable content-addressable network, in: Proc. of ACM SIGCOMM, 2001.
- [21] S. Rhea, D. Geels, T. Roscoe, J. Kubiawicz, Handling churn in a DHT, in: Proc. of the USENIX Annual Technical Conference, 2004.
- [22] R. Rivest, A. Shamir, L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, Communications of the ACM 21 (2) (1978) 120–126.
- [23] R. L. Rivest, A. Shamir, PayWord and MicroMint: Two simple micropayment schemes, in: Security Protocols Workshop, 1996.
- [24] A. Rowstron, P. Druschel, Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems, in: Proc. of the 18th IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), 2001.
- [25] A. Rowstron, A.-M. Kermarrec, M. Castro, P. Druschel, SCRIBE: The design of a large-scale event notification infrastructure, in: Proc. of the 3rd International Workshop on Networked Group Communication (NGC), 2001.
- [26] A. Singh, T. Ngan, P. Druschel, D. Wallach, Eclipse attacks on overlay networks: Threats and defenses, in: Proc. of IEEE INFOCOM, 2006.
- [27] E. Sit, R. Morris, Security considerations for peer-to-peer distributed hash tables, in: Proc. of the 1st International Workshop on Peer-to-Peer Systems (IPTPS), 2002.
- [28] M. Srivatsa, L. Liu, Countering targeted file attacks using location keys, in: Proc. of the 14th USENIX Security Symposium, 2005.
- [29] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, H. Balakrishnan, Chord: A scalable peer-to-peer lookup service for Internet applications, in: Proc. of ACM SIGCOMM, 2001.
- [30] W. Thigpen, T. J. Hacker, L. F. McGinnis, B. D. Athey, Distributed accounting on the grid, in: Proc. of the 6th Joint Conference on Information Sciences, 2002.
- [31] B. Y. Zhao, J. D. Kubiawicz, A. D. Joseph, Tapestry: An infrastructure for fault-tolerant wide-area location and routing, Tech. Rep. UCB/CSD-01-1141, University of California, Berkeley, USA (April 2001).
- [32] P. Zheng, Tradeoffs in certificate revocation schemes, ACM SIGCOMM Computer Communication Review 33 (2) (2003) 103–112.
- [33] Y. Zhu, H. Wang, Y. Hu, A super-peer based lookup in structured peer-to-peer systems, in: Proc. of the 16th International Conference on Parallel and Distributed Computing Systems (PDCS), 2003.
- [34] S. Q. Zhuang, B. Y. Zhao, A. D. Joseph, R. H. Katz, J. Kubiawicz, Bayeux: An architecture for scalable and fault-tolerant wide-area data dissemination, in: Proc. of the 11th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV), 2001.