



UNIVERSITY OF AMSTERDAM

Master System and Network Engineering

Automatic SSH public key fingerprint retrieval and publication in DNSSEC

Research Project (1) report

Marc Buijsman, Pascal Cuylaerts
{marc.buijsman, pascal.cuylaerts}@os3.nl

6 February 2011

Contents

1	Introduction	1
1.1	Research question	2
2	Research	3
2.1	The desired mechanism	3
2.2	Shared secrets	4
2.3	Authentication without shared information	4
3	Mechanism design	6
3.1	The key retrieval mechanism	6
3.2	The key retrieval mechanism under attack	7
3.2.1	Attacker forwards messages	7
3.2.2	Attacker modifies messages	8
4	Implementation	9
4.1	Secrets file	9
4.2	Secret look up at the remote host	9
4.3	SSH connection	10
4.4	Local DNSSEC validation	11
4.5	Encryption	12
4.6	Pushing updates to the DNS	12
4.7	Existing list of SSHFP records	12
4.8	OpenSSH patch	12
4.9	System requirements	13
4.9.1	Overview	13
4.9.2	Description	14
5	Conclusion	15
	References	17
A	Program code and configuration files	i

A.1	For the administration machine	i
A.1.1	Application	i
A.1.2	Configuration file	viii
A.1.3	Usage	viii
A.1.4	Python interface to SSH client functionality	viii
A.2	For the remote host	xviii
A.2.1	Application	xviii
A.2.2	Configuration file	xxii

1 Introduction

The concept of trust is fundamental in computer network security. Although not everyone is aware of this, encrypted network connections are not much safer than unencrypted ones if the person that initiates such a connection does not explicitly trust that he or she is talking to the intended endpoint.

Take for example an HTTPS connection to a bank's website. If the browser shows that the SSL certificate used for the authentication of the bank's website is valid and that the channel is encrypted (e.g. by showing a lock icon or a green address bar), then one may trust that it really is that bank where he or she is sending sensitive data to (and that no one else can read it).

However, many people do not think about what their trust is actually based on. They implicitly trust the browser's maintainers who ship the browser with a list of what they think are trustworthy certificate authorities (CAs). These are the organisations (third parties) which ultimately need to be trusted since they are the ones signing the certificates, thereby claiming that the website is truly in hands of the associated organisation (the bank in this example).

Unfortunately, someone might not notice that a certificate has been signed by a CA which he or she does not actually trust, but which has been included by the browser's maintainers. If one cannot trust a "secure" connection to another machine, then it cannot be ruled out that an eavesdropper sits in between. Moreover, such a man-in-the-middle could maliciously impair the data flow.

The Secure Shell (SSH) protocol is, like SSL (or TLS), a way to have a secure (encrypted) connection between two computers. It is a protocol used for remotely accessing a machine's command line (shell) with end-to-end encryption. Unlike with SSL as in the example, with SSH one is likely to be confronted with the trust aspect more often. A machine's shell is usually supposed to be accessible by only a few people, whereas websites are aimed at serving many people. Because of this, it makes little sense to purchase a certificate from a CA to secure SSH connections. That is why an SSH client normally involves the user in the authentication process instead.

When one is connecting to a host using SSH for the first time, a so-called fingerprint derived from the remote host's public key is usually presented [1]. In the OpenSSH client this is a hexadecimal presentation of the MD5 hash of the public key. The user can either accept the fingerprint and continue connecting, or refuse it to abort the connection. This step is important. If the user believes that the public key belongs to the private key that is held by the intended host, then it is safe to continue. In the initialisation process the remote host must cryptographically prove that it possesses the private key to authenticate itself.

If the user does not trust the fingerprint, then it would be unwise to accept it. It could be the fingerprint of an eavesdropper for instance. To be able to verify the fingerprint, the user must have had contact with the host's administrator to retrieve it safely. If the user trusts the way the fingerprint has been retrieved, then this person can also trust that he or she is trying to connect to the intended host machine if the presented fingerprint matches the one retrieved out-of-band. The chance that there still is an eavesdropper in the middle is very small since it is hard to generate a public and private key pair with exactly the same fingerprint.

It would however be convenient for a person who is initiating an SSH connection to have a way of verifying the authenticity of a received SSH public key without his or her intervention. Possible human error when comparing fingerprints would also be eliminated. Such a means would require the person (if he or she cares about security) to trust an automated verification process, such that when the key is positively verified he or she can implicitly trust the key to be valid.

If this trust is based on a locally stored list of public keys or fingerprints that was composed by the person him- or herself then a simple automated lookup in this list would suffice. However, this solution is not very scalable. Every person has to compose his or her own list, and keys of previously unknown hosts still have to be verified manually.

The use of the Domain Name System (DNS) offers a better solution, as this is a single database that can be accessed by everyone. An administrator can publish a public key fingerprint in the DNS so that it is instantly publicly available, making it an easy way of distributing fingerprints.

The response to a DNS lookup request can be trusted if DNSSEC (DNS Security [2] [3] [4]) is used. If the retrieved resource record has been signed by an instance that is part of a DNSSEC chain of trust which is ultimately signed by a trusted instance (most commonly the DNS root), then the authenticity of the record can be verified. This would mean that a DNSSEC-validated SSH fingerprint resource record (SSHFP RR [5]) that is tied to a domain name can be trusted to be authenticated by the instance that has the authority over that domain name.

We earlier mentioned that a person would need to contact a host's administrator to retrieve the machine's fingerprint. This could however pose a problem if this person is an organisation's administrator him- or herself. If he or she administers only one machine, then it is not a big deal to walk to the machine, access it directly to retrieve its fingerprint and carrying it back to a workstation. This is the safest way to transport the fingerprint. But if there are many machines of which the fingerprints are yet unknown, then this becomes a cumbersome task.

For someone in such a situation it would be convenient to automate this task. A workstation can be used to collect the fingerprints, which could also push them to the DNS so that other workstations can easily retrieve them as well. When automating this whole process it is inevitable that the a potentially untrusted computer network will be used for the fingerprint retrieval. During our project we investigated a way of retrieving the fingerprints of remote machines securely over an insecure network in the situation where public keys are yet untrusted as a means of host authentication. Such a mechanism of validating a host's fingerprint opens the way for automated fingerprint retrieval and publication in DNSSEC.

1.1 Research question

Most of our research was focused on the problem of the insecure connection between an administrator's workstation and a remote machine whose SSH public key is unknown. We have investigated if this channel can be secured, and if so how this can be implemented in a software tool. We have also tried to enable this tool to automatically publish fingerprints in the DNS. This is the practical side of our project; to enable the tool to automatically collect fingerprints in a secure way, the research is a prerequisite for its implementation.

Our research question is:

How can SSH public key fingerprints be automatically collected from remote machines and published in DNSSEC in a secure way?

This can be further divided into the following subquestions:

- What are the possible solutions for secure data transfer over an untrusted network?
- Can we make use of existing methods or protocols to realise the possible solutions?
- How can these solutions be implemented in a tool that automates the collection of SSH public keys?
- How can we insert the SSH public key fingerprints into the DNS and sign them using DNSSEC in an automated way?

2 Research

In the introduction we explained how DNSSEC can be used to verify the validity of SSH fingerprints and therefore the validity of public keys. If a trust anchor was reached during the DNSSEC-validation of a resource record, then it can be trusted that this record has been authenticated by the instance that has the authority over the concerning domain name. Ultimately, this instance itself needs to be trusted as well. If one does not trust that the instance took great care of publishing the correct SSH fingerprint in the DNS, then doing DNSSEC validation makes little to no sense.

A DNS SSHFP record contains a SHA-1 hash (called “fingerprint”) of either an RSA or DSA public key [5]; both types can be used in the SSH authentication protocol. The hash is preceded by a number denoting the type of key used (1 for RSA and 2 for DSA) and a number denoting the used hashing algorithm (1 for SHA-1). An example is as follows:

```
domain.com IN SSHFP 2 1 d066788e581f8d91faf1e715954fca596324e851
```

2.1 The desired mechanism

We will be describing a mechanism for automatic public key retrieval from remote machines and fingerprint publication in the DNS. We focus for a large part on the situation where the public keys of the remote machines are not certain to belong to those machines. If one uses such a mechanism and he or she wants to be sure that the correct public key fingerprints are published, then there must be a way to verify that a received public key really belongs to the intended machine. After all, there could be an attacker in the middle with whom the actual SSH connection has been set up.

Since in such a situation one cannot be sure whether or not a public key belongs to a certain machine, it cannot be used for the authentication of the machine’s identity, even if the machine can prove that it possesses the corresponding private key. It is our goal to collect the public key in such a trustworthy way that it eventually can be used for this purpose. This is necessary for SSH connections where public key cryptography plays a central role in server authentication.

Therefore, some secure mechanism is needed to establish the *authenticity* and the *integrity* of a collected public key. That is, we want to make sure that a public key belongs to the machine with a certain identity, and we want to ensure that its integrity has been preserved during transfer to prevent a possible publication of a wrong or malicious fingerprint into the DNS.

The most secure way to collect public keys would be transporting them out-of-band from each machine separately. This would require a person to physically access these machines one by one to extract the public key. If there are many machines with many administrators, then this task can be simplified by asking each administrator to send their machine’s public key in a GPG [6] [7] signed email, for example. However, the senders’ GPG public keys first need to be trusted as well. If many machines are under control of a single administrator, this solution may not be workable because he or she still needs to physically access a relatively large number of machines.

In the last case, it would be very convenient to be able to automate the key retrieval process by a computer program without further human intervention needed. This will however need to be done over a potentially insecure network, because there is no other way a computer program can contact a remote machine. What we have here is a classic chicken-and-egg problem. We need to authenticate a machine for which we need its public key and we want the machine to prove that this really is its public key, but then we already need to have authenticated the machine. The machine therefore needs something else than a public and private key pair to be able to identify itself.

2.2 Shared secrets

In general, a person that needs to authenticate him- or herself, will need to know something (e.g. a passphrase), have something (e.g. a smartcard), be something (e.g. his fingerprint), do something (e.g. a signature) or a combination of these. The authority that is authenticating this person needs to be able to verify the provided information. In computer security, if two machines need to authenticate one another, they will often know each others public key and use challenge-response authentication combined with public key encryption. An alternative is to have both machines to know some shared secret such that each computer can prove somehow that it knows what the secret is, without revealing it to the outside world.

A shared secret can be seen as a passphrase. Just like passphrases, such a secret needs to stay secret between two parties to prevent a third party from misusing it. Unlike with public and private key pairs, both parties need to protect the secret since they both need to know it to be able to authenticate each other. If only one of the parties needs to authenticate itself to the other using a public and private key pair, then this party needs to protect the private key whereas the other party does not need to protect anything. It does need to know the public key, but since this key is publicly available it does not have to be protected from outsiders.

It could be easy to use a shared secret as a means of authentication in some cases though. A machine specific system identifier can be looked up by the machine itself or by someone having elevated privileges on the machine. A system's Universally Unique Identifier (UUID) for example is a good candidate for a machine identifier (as will be explained later), which is usually only readable by users who have root privileges. We decided to make use of a shared secret since some system identifiers might be listed on hardware inventory lists that are available within an organisation.

Having these numbers on paper already makes a walk to every machine within the organisation, to retrieve the identifiers manually, unnecessary. They could be entered in a computer file straight away. Once this has been done, a program can use this file to perform the automatic public key retrieval process. The only assumption we made is that the identifiers have not been copied by an untrusted party during the identifiers' retrieval process and that the inventory lists are stored safely, something that is important when using them as secrets but which we have not further investigated.

2.3 Authentication without shared information

At first, we tried to come up with a protocol that does not need any pre-shared data for the machines to be able to authenticate one another. In this case, there is no shared information that can be used for host authentication. Most of the possible solutions for this problem we have read about consisted of identity-based key agreement schemes that require a trusted third party to act as a key generation center (KGC [8]) that creates key pairs. Apart from the need for a trusted third party, these schemes were too complex for our application.

Methods to detect a man in the middle can also be used, such as the leap-of-faith method [9]. If there is someone in between during the first connection, then he must be in between during all the subsequent connections to prevent the administrator from being warned that the public key has changed. This could be hard to do for the attacker and therefore a second connection can be set up after a certain timespan to see if there will indeed be a warning. If so, then the administrator will know that there was someone in between either during the first connection or the second, making the received data during either connection untrusted.

An administrator could also make assumptions about the network between him or her and the remote host to determine if it will be safe enough to proceed without having the ability to authenticate the received data. Such an assumption can for example be that only the local area network (LAN) will be used which may be considered clear from intruders. Also, since our mechanism needs to be used only once to retrieve public keys, the risk of an attacker being present during the retrieval process is reduced

to only one connection for each host. This could be considered an acceptable risk.

However, since there is no information available to authenticate a remote host in these situations, data exchange can never be completely secure. We need information that can be used to authenticate a host to be able to set up a secure session with the host, so ensure that no malicious fingerprints will be published in the DNS. For our mechanism this information will be a pre-shared secret.

3 Mechanism design

3.1 The key retrieval mechanism

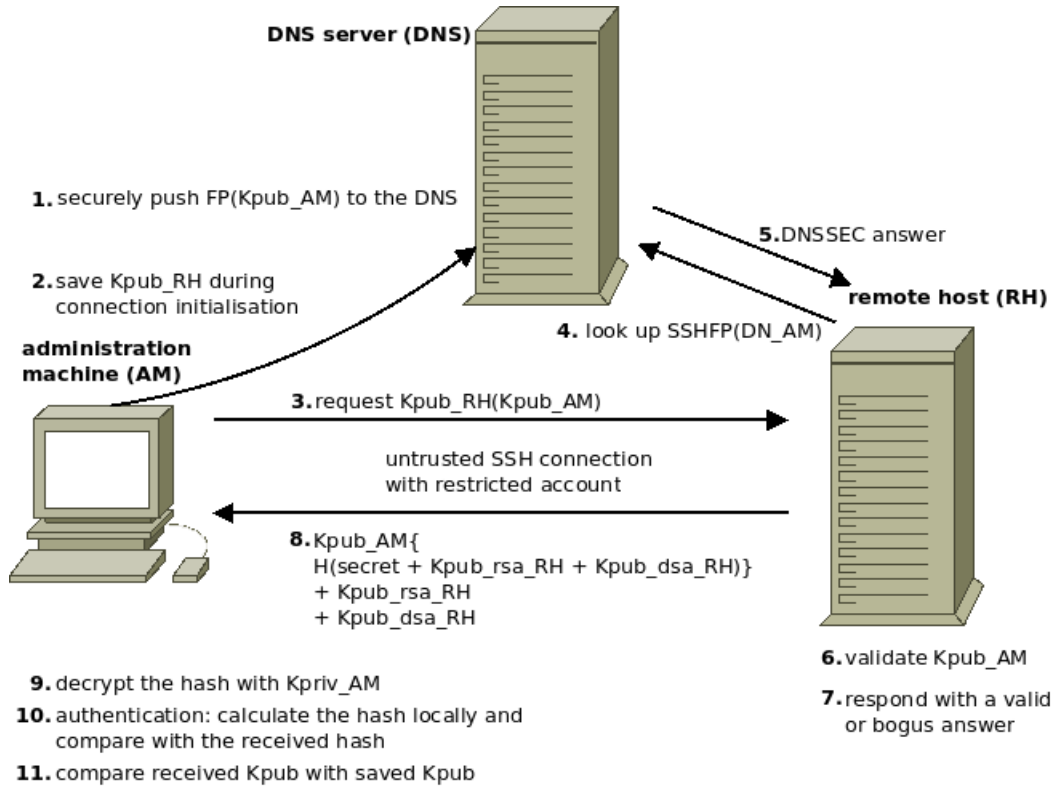


Figure 1: Key retrieval mechanism.

The mechanism we devised to securely retrieve remote hosts' fingerprints and publish them in the DNS (signed using DNSSEC) is illustrated in figure 1. This mechanism assumes that an administrator wants to collect the SSH public keys from a number of remote hosts (RHs) using one administration machine (AM).

To authenticate the responses that the AM will receive from the RHs, a list of shared secrets needs to be available on the AM with an entry for each RH. Because this shared secret is the only means for a remote host to authenticate its identity, this data needs proper protection and must at least be encrypted when stored on disk. Another requirement is that the fingerprint from the AM's SSH public key ($FP(K_{pub_AM})$) is stored in the secure domain name system (DNSSEC) (1).

The AM will contact a RH to retrieve its SSH public key (K_{pub_RH}) using SSH. This connection is untrusted and the account used to log in on the RH must have restricted permissions (since the credentials can be read by an eavesdropper). When the connection is being established, the AM will receive K_{pub_AM} and store it temporarily to use at the end of the process (2).

Once the connection has been established, the AM will send a request to the RH to ask for its public key and in this request the AM will include K_{pub_AM} (3). When the RH receives this request, it will look up the SSHFP records in the DNSSEC using the domain name of the AM (4) which needs to be pre-configured on the RH. The SSHFP records (with the associated RRSIGs) in the answer (5) will be validated locally and compared to the fingerprint derived from K_{pub_AM} (6). If the two fingerprints match, the RH will send a response to the AM which includes its secret and SSH public key. If the fingerprints did not match, the RH will respond with a bogus answer (7).

A valid response (8) is built up as follows:

$$K_{pub_AM}\{H(\text{secret} + K_{pub_rsa_RH} + K_{pub_dsa_RH})\} + K_{pub_rsa_RH} + K_{pub_dsa_RH}$$

The secret is concatenated with the RH's RSA and (if present) DSA public keys and this string is hashed. The resulting hash will be encrypted with K_{pub_AM} and then concatenated with the cleartext RSA and DSA public keys of the RH.

Upon retrieval of this response, the AM will decrypt the hash with its private key (K_{priv_AM}) (9) and calculate its own hash (10) with the received public keys and the secret it has stored locally. If the hashes match the AM can be sure that the response came from the RH he intended to contact and that the response has not been modified on the way back. The hash is therefore used to check the integrity of the public keys that were sent along. Since the secret is incorporated, the keys' authenticity can also be verified.

As an extra security check, the AM can now compare the public key he stored at the beginning of the process with the one he just received. If they do not match, the machine he was communicating with must have been an attacker that was performing a man-in-the-middle attack and who forwarded the request to the actual RH to let it respond with a valid answer. However, the keys' fingerprints can still be published in the DNS if the hashes match since that proves that the answer was not tampered with by the man in the middle.

3.2 The key retrieval mechanism under attack

3.2.1 Attacker forwards messages

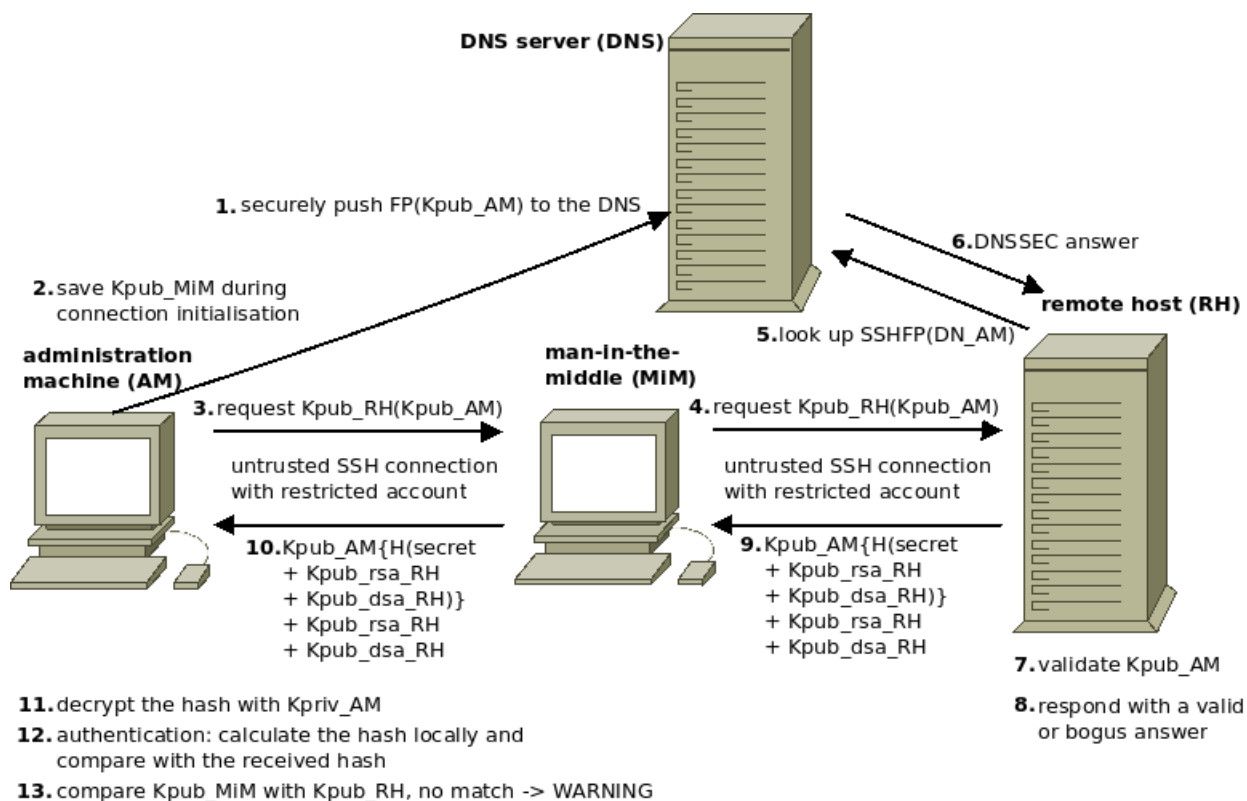


Figure 2: Key retrieval mechanism under MITM attack.

If this mechanism is under a man-in-the-middle attack, as illustrated in figure 2, the public key stored at the start (2) will be the one from the man in the middle (MiM). The MiM will just forward the request from the AM (3) to the RH (4) which will validate K_{pub_AM} using DNSSEC (5, 6, 7) and think it is really talking to the AM. As a result it will respond with a valid answer (8) but the hash of the secret concatenated with the public keys will be encrypted with the public key of the AM (9). This makes the intercepted response unreadable for the MiM because he does not know the AM's private key.

After the MiM forwarded the response to the AM (10), the AM will decrypt the hash (11) and calculate the hash itself with the received public keys and the secret it has stored locally (12). If the MiM has not tampered with the public keys and the hashes still match, the AM still does not have a clue that a third party was in the middle, which accepted the SSH connection and saw the login credentials passing by. But because the public key from the host he connected to was stored when the connection was set up (2), he can now check whether it matches the public key from the response. If not, he knows something suspicious happened.

3.2.2 Attacker modifies messages

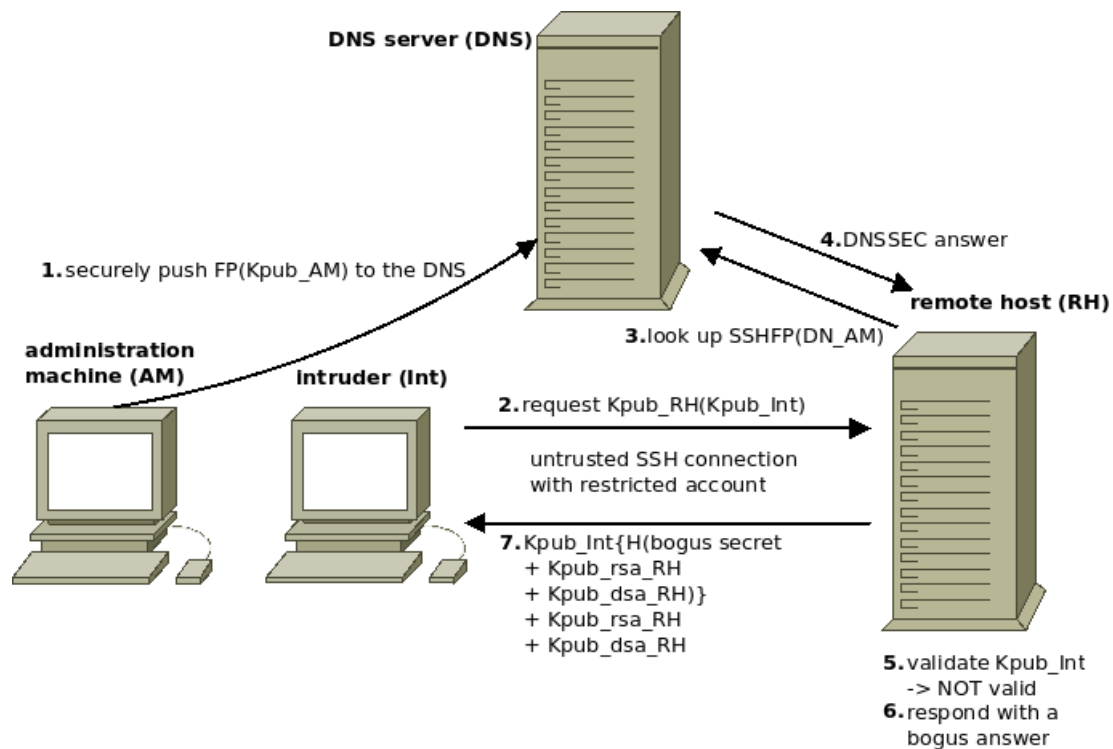


Figure 3: Key retrieval mechanism under attack directly.

When an intruder (Int) manages to log in directly into the RH and requests the secret (2) without forwarding the AM's request as illustrated in figure 3, the RH will notice that the Int is not the AM because the fingerprint of its public key does not match the one he looked up in the DNS (3, 4, 5). As a result, the RH will return a bogus answer (6, 7) encrypted with the Int's public key. This answer will contain a hash from a random string concatenated with the RH's public keys. The Int will be able to decrypt the hash and may assume he received a valid answer. He can now perform offline attacks in an attempt to recover the secret but he will only end up with a random string.

Not sending a response when the fingerprint of the intruder's public key (K_{pub_Int}) does not match the fingerprint found in the DNS would simplify the attacker's job since he would have less hashes to perform attacks on. It also prevents the AM to notice that something is going on, which would be a good thing to know such that the situation can be further investigated.

4 Implementation

We implemented this mechanism as a proof of concept with two programs written in Python (listing 1 in A.1.1 and listing 5 in A.2.1) for the Linux OS that will handle the communication between both parties. The programs need to be configured using a configuration file (listing 2 in A.1.2 and listing 6 in A.2.2).

The program that will be executed on the administration machine has two modes of operation. In the normal mode, the program will retrieve the public keys from the remote hosts and push their fingerprints to the DNS in the form of SSHFP records. The second mode takes a list of SSHFP records as input and pushes them directly to the DNS.

4.1 Secrets file

In normal mode, the program needs to have access to a file with one line of information for every host that needs to be contacted. This line will have the following format:

```
host.domain.org:4445434C-5700-1050-8034-B7C04F56344A:..CN7084106E00YU.Product Name
```

The first part is the domain name, then a strong secret followed by a weaker secret to which the program can fall back if the strong secret is not available. These are all separated by colons which we believe are acceptable separators for the types of secrets we had in mind.

We chose the system's Universally Unique Identifier (UUID) as the strong secret because of its selection from a large key space, making it hard to guess, and because it is the best information to uniquely identify virtual machines (VMs) and thus to authenticate them. The UUID is usually also listed in the configuration file of a VM which can easily be processed in an automated way to collect the UUIDs from all guest VMs if one has access to the host machine.

The weak secret is a concatenation of the serial number of the system's motherboard and its product name to enlarge the key space. Assuming that a detailed inventory is kept of all hardware used in an organisation's network with this kind of information, it should be easy to generate a list of physical hosts with their secrets. If the UUID of the machine is also listed in the inventory, then that is an advantage, because of the larger key space.

Note that the weak secret is more vulnerable to dictionary attacks. Building up a dictionary of known product names would be easy and a part of the motherboard's serial number also refers to the manufacturer, reducing the possible combinations. Any information that is already available to the administrator can be used to authenticate the remote hosts, and for our proof of concept we considered the serial number and product name identifiers secure enough.

4.2 Secret look up at the remote host

The remote host can find out its own secret from the output of the `dmidecode` command. This program will parse the contents of the system management BIOS (SMBIOS) table and present them in a human-readable format. The SMBIOS contains a description of the system's hardware components and other useful information such as serial numbers and details about the BIOS. Dmidecode will access the file `/dev/mem` to access this data. A user that wants access to this file will need elevated permissions.

The values read from the SMBIOS table are not always reliable, because manufacturers can leave values empty or can choose to fill in different kinds of information. The SMBIOS standard [10] is specified by the Distributed Management Task Force (DMTF) and not all the fields of the SMBIOS table are required to be filled in to comply to the standard. The UUID and the Product Name are required fields, but the

motherboard's serial number is not required. Although it may be empty according to the standard, we still chose to use the motherboard's serial as part of the secret, because it is a good identifier and most manufacturers seem to fill it in correctly.

Because the information in the secrets file is so critical for the authentication of a host, it should only be stored on disk at the administrator's side with proper encryption. Therefore our program will accept an AES encrypted file and prompt the administrator for the passphrase it needs to decrypt the file.

Please provide your credentials for the remote hosts.

Username:

Password:

Please provide the passphrase to decrypt the secrets file.

Passphrase:

On the remote hosts a restricted user account must be configured. This account will be used to set up the untrusted SSH connection over which the authenticated key retrieval will take place. Our program will prompt the administrator for these credentials at start up so he will not have to enter them in a configuration file in cleartext.

In order for the program that will be executed on the remote host to be able to read `/dev/mem`, the restricted user account needs to be able to run our program with elevated permissions. Therefore we added a line in the sudoers file `/etc/sudoers` and use `sudo` when executing the program.

```
untrusted ALL = (root) NOPASSWD: /path/to/program
```

This line means that the user `untrusted` can execute from `ALL` terminals, acting as `root` the program `/path/to/program` without being prompted for a password.

4.3 SSH connection

When the credentials have been entered and the secrets file could be decrypted (which is done using `gpg`), the program at the administrator's side will go through the secrets file line by line, creating one SSH session after the other with the host at each of the domain names. To be able to set up an SSH connection it uses a Python module that interfaces with the `libssh2` C library. We created this module ourselves (which we called `sshexec`, see listing 4 in A.1.4), with the basic functionality needed for an SSH session. It has been implemented using the Python C API [11] so that it could be included in the program.

In its current implementation only one SSH connection can exist at a time. It was largely based on an example source file that came with `libssh2`. When a connection has been initiated using the module, it returns the remote host's public key. This key will later be used to check if there was someone eavesdropping on the connection.

Once connected the program will ask the remote host to authenticate its RSA and (if present) DSA public keys using the type of shared secret. If both a strong and a weak secret are listed, then the strong secret will be used. It will do this by executing a command on the remote shell which will initiate the program at the remote machine's side. Once an answer has been received or when the execution timed out, the connection will be closed and the program will continue with the next line in the secrets file after it validated the public keys with the hash if that was sent back.

An answer consists of the encrypted hash concatenated with the RSA and DSA public keys in Base64 encoding, separated by colons (which do not occur in the Base64 encoding scheme). This string is

preceded by the response type, which can be "ANSWER", "ERROR" or "WARNING". Only with "ANSWER" a hash will be sent, the others will accompany a human readable message for alert and debug purposes.

An example of an "ANSWER" response is as follows:

```
ANSWER:saDp4JhJNNDttXgu9UidZEZDdq6VInS2Pyt1innR2SZLfBaFuZazzNns0vW2S9DkV/yng0Aee
t2dLuj1vJH3dV1bAPE4qQWj4uBdCJQE4oSUS3A5PjYnedZZYXpCjYQxzFDrKD166yqRUQdtFmpRbgI/bf
i+rEcn1YSU15pdVjuzQK/B3moYPuScCtj/7o9rn/Yn3auUCC3NzrlmPPibFi94ryLBcAQc3d0YW2N9S2
+0Fy1CZfdyRZiemr8g8P+W+gFeTKZEeSiG3GwZxeNuWxmLgkBsU+P4dViHR419dPayfeBTxcVD1T7PLX
e4/t3Q5GnzM4lzT6p47814TTmBg+w==:AAAAB3NzaC1yc2EAAAABIwAAAQEA+EVtkCxc1j1gI2J3HrH2
gkQFgg4dZXBwq6aV49330VGP6RRcn78RwkF+3zr1jnYhBCe1UmePQhmlsZH4ivXWY33XX27JX5ZZjsQ0
wPXXcS81wCb2p0Y4R2+pNKtp0uOM3YWSVXyLcANiABWBay+QPFnwyswcJ4o3AUhKuWz1hKUKpHGv10Is
2nIkYjY2Z1IcLbKlFEswur1Wf41ZRqqRkanS7T3UraxtrSC+Hz4aEuB9/WGJ4t/NReXpYBD1m78CgrfX
bje5LAMWGYr+Rri97KUB2vH/XN/aI7VV0u9ik7gH3PrLaeTsNOUMgSC45TQwiygaGIOuNUZPyx3ISX5K
gQ==:AAAAB3NzaC1kc3MAAACBAImVL4qXUVV0z1qYg/OaGvfXqEW3CIuJ3Dc0+ENO9ueKNu9p/RJ8+eZ
bn5vD8bEOwVWvg7/dirheKmMNVmUpDox99b7VaJaUufY8gZT80omN7NvBSQ64hWXuHA/xMbGdg6r6YDN
AmOPsnnLR90kWhLOWKIhN9INU68VtmcC8siGAAAAAFQDE6PYVTjb5XKtn1Uvs/jzYx+TenQAAAIBMBwb
2/AOE6/q/EZzWtp94oGNDDJ1VEWd6X7kdgsYjAXM0fk/eH2ri82+7X3JpeGS6LELaBqIhs3hG2HZp9wj
6bp5gLqjc1dWH8IKQpc0xJA/SDGDaH+xKklso1pxqIad/wivMAFo3I/+ch1777K/EKXN4uIzEETMUPL0
mq++nrAAAAIAjOUs3QZGpcpdWMFX8eVDnsrctvEcRjfgdUJx7pnr0sSX+NNNhTEB8JOXggHg5htfItEp
g2sBfp+Kpr9PpL+e1G14VTqNs47jJsadnvQZSRUJ5aZaKeX7VpEpyZxd98Cqcn4BOMLKLs5nEHTHyNoq
KqGVIOGB33+b2WLVa8dTpCg==
```

4.4 Local DNSSEC validation

The program at the remote host's side makes use of the `LibUnbound` [12] Python module to do local DNSSEC validation of the RSA public key it receives from the administration machine. The program looks up the SSHFP records of the domain name that was locally configured as the domain name of the administration machine. When the fingerprint of the received public key matches the fingerprint in an SSHFP record, and if that record has been validated using DNSSEC, the program will respond with the shared secret. If the fingerprint could not be validated, a bogus answer will be generated.

This step is important in the sense that it prevents an eavesdropper from discovering the host's secret. The generated hash will be encrypted with the received public key and only if this key belongs to the administration machine, that machine can decrypt the hash using its private key. If the key could not be verified as belonging to the administration machine, then it is possible that an eavesdropper is in between.

When the hash is encrypted by the eavesdropper's public key, he will also be able to decrypt it. By sending a bogus answer when the public key's fingerprint does not match, he will receive an invalid hash, making an offline attack on the hash to discover the secret pointless since it has not been involved in creating the hash. If the eavesdropper forwards the answer to the administration machine it can detect that something is wrong since the hash will not match the one it generates itself.

The DNSSEC validation must be done locally so that the whole validation process does not rely on the "last mile" between the DNS server and the host in which the DNSSEC answer could be forged to look valid when it is actually not. It is therefore necessary to have the certificate of a trust anchor installed at the host which in our case was the DNS root's certificate. One might consider to run Unbound as the local DNS resolver so that the root certificate is automatically updated when its key has been rolled over.

4.5 Encryption

As mentioned before a remote host uses the RSA public key of the administration machine to encrypt the hash. We included the `M2Crypto` [13] Python module for encryption functionality. A public key object is created from the RSA exponent and modulus that are extracted from the administrator's public key which is passed on to `M2Crypto` along with the hash to perform the encryption.

RSA “Optimal Asymmetric Encryption Padding” (OAEP) is applied just before the encryption to minimise the chance of a successful cryptographic attack [14]. This also causes the ciphertext to be different each time the same hash is being encrypted, making it impossible for an attacker to find out if an answer from the remote host is actually valid by trying to see if the answer stays the same after multiple identical requests (e.g. with a replay attack). Without the padding a valid answer would not change indeed, whereas a bogus answer is randomly generated at each rejected request.

At the administrator's side, `M2Crypto` is used again to decrypt the hash. The machine's private key is passed to the module, which is the reason why the program must run with root privileges since the private key is not world readable.

4.6 Pushing updates to the DNS

In case a list of SSHFP records is provided, the application will immediately try to push the new records to the DNS server, skipping the key retrieval process. Otherwise, the public keys are first retrieved from all the remote hosts whereafter SSHFP records are generated for the trusted keys. To perform dynamic DNS updates, we use `nsupdate` which is part of the package `bind9utils`.

Transaction signatures (TSIG) [15] are used to authenticate the updates. These signatures rely on a shared secret between the administration host and the DNS server. The secret key needs to be configured on the DNS server and the path to the local keyfile also needs to be configured in the configuration file of our application. Hash-based Message Authentication Codes (HMAC), HMAC-SHA512 in our implementation, are then used to ensure authenticity and integrity. We also force `nsupdate` to use TCP instead of UDP to ensure a successful update.

4.7 Existing list of SSHFP records

As mentioned before (2 Research), public keys can also be retrieved out-of-band or via encrypted email (GPG). We added the functionality to push an existing list of SSHFP records to the DNS, just by feeding the file to our administration application. The administrator just needs to offer the program a file with valid SSHFP records each on a new line. The help section of the application (listing 3 in A.1.3) shows how to use the arguments.

4.8 OpenSSH patch

The result of this whole process is of course more useful if one has a client application that actually looks up the SSHFP records in DNS and does local DNSSEC validation of the answers.

On the website <http://www.dnssec-tools.org/> one can find a whole suite of tools that make use of DNSSEC. First the `DNSSEC-Tools` package will need to be installed, which will install the DNSSEC-Tools resolver and validator libraries and headers on the system. Then OpenSSH [16] [17] can be patched with the patch included in the package. More detailed installation instructions can be found in the `README` file of the package, or on the website.

Once OpenSSH has been patched successful, a new option can be used, `StrictDnssecChecking`, in

`ssh_config`. This option can have the values `yes`, `no` and `ask`. One will also need to enable `VerifyHostKeyDNS`. This option is already available in the normal version of OpenSSH, but the patch is needed to add validation of the DNS answer using the RRSIG resource records.

When one tries to connect to a host whose fingerprint cannot be validated using DNSSEC, the following warning will be shown:

```

@ WARNING: UNTRUSTED DNS RESOLUTION FOR HOST KEY! @

```

If the key has also changed since the previous connection (according to the `known_hosts` file), an even stronger warning will be displayed:

```

@ WARNING: UNTRUSTED DNS RESOLUTION FOR HOST KEY! @
@ WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED! @
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA host key has just been changed.
The fingerprint for the RSA key sent by the remote host is
ba:7e:98:3c:42:96:54:b6:67:30:7a:3c:df:fd:33:7d.
Please contact your system administrator.
Add correct host key in /home/<user>/.ssh/known_hosts to get rid of this message.
Offending key in /home/<user>/.ssh/known_hosts:<line number>
RSA host key for host.domain.org has changed and you have requested strict checking.
Host key verification failed.

```

When the public key of the remote host can be trusted, a user will immediately be prompted for his or her credentials and will not be bothered with any message, not even the public key’s fingerprint.

4.9 System requirements

4.9.1 Overview

Administration machine

- Python application (listing 1 in A.1.1)
- dependencies (`argparse`, `M2Crypto`, `libssh2`, `bind9utils`)
- Python interface for `libssh2` C library (listing 4 in A.1.4)
- configuration file (listing 2 in A.1.2)
- encrypted secrets file
- shared (with DNS) key file

Remote host

- Python application (listing 5 in A.2.1)
- dependencies (argparse, M2Crypto, libunbound)
- configuration file (listing 6 in A.2.2)
- restricted user account
- edited sudoers file (see 4.2 Secret look up at the remote host)

DNS server

- SSHFP records for administration machine
- edited `named.conf`
- allow for dynamic updates (`nsupdate`)
- shared (with AM) key in `named.conf`

4.9.2 Description

The tools we created were meant as a proof of concept only intended to be used under a Linux OS. The two programs have their own dependencies and these can also have dependencies themselves. We have not tested any configurations other than our own, so it is always possible that one will need to have some library that is not listed in the overview above.

Dependencies

One will need to have at least the packages `python`, `python-argparse` and `python-M2Crypto` installed on the administration machine (AM) and the remote hosts (RH). The application at the AM needs `libssh2` in order to set up the SSH connections and `bind9utils` to perform the dynamic updates with `nsupdate`. On the RH, an installation of `libunbound` is required to do the DNSSEC local validation. For our application to be able to use the `libssh2` C library, the included Python interface we have developed needs to be present too.

Configuration

For both applications a configuration file is used to adjust the program to a specific implementation. On the RH a restricted user account needs to be configured and the sudoers file needs to be modified to allow the user to run our application with root permissions. For secure dynamic updates, a shared key needs to be present on the AM and the DNS server (in `named.conf`). The AM needs to be allowed to perform updates and the fingerprint of its public key needs to be published in the DNS beforehand.

5 Conclusion

The SSH protocol provides an encrypted channel with a remote host in order to securely use its shell. To authenticate the remote host it makes use of public key encryption. During the first connection setup with a remote host, the user of an SSH client program is usually asked to verify the host's public key fingerprint. However, this fingerprint may be unknown to the user. Normally, he or she should retrieve the fingerprint from the remote host's administrator out-of-band and check if it matches the one received over the network. If this is not the case, then a man in the middle could be listening on the line and modify the sent data if the user still accepts the fingerprint and proceeds with the connection.

It would be convenient to have a mechanism that can be used to retrieve and verify a yet untrusted public key without human intervention. In our project we have worked towards a solution in order to make that possible. In the introduction of this report we gave the research question of our project, divided into subquestions. The research question was:

How can SSH public key fingerprints be automatically collected from remote machines and published in DNSSEC in a secure way?

By answering the subquestions, the research question can be answered.

What are the possible solutions for secure data transfer over an untrusted network?

We wanted to have a way to authenticate data sent by certain remote hosts without the use of their public and private key pairs, since these are yet untrusted in the described situation. We also wanted to automate this process such that it would not be necessary to do this manually. If there are a lot of machines for which this needs to be done, then the solution for this problem offers the possibility of authenticating the hosts' public keys easily.

We have investigated what the possible solutions for this problem are without de need to rely on a trusted third party. We can distinguish two types of solutions: one type where the remote host's identity cannot be verified due to the lack of information about that host, and another type where such information is known such that a host's identity can be established.

The first type of solutions can never be completely secure. The administrator (who is initiating the automatic public key retrieval process) has to make some assumptions about the part of the network he or she uses and determine if it is safe enough to proceed without having the ability to authenticate the received data. Such an assumption can for example be that only the local area network (LAN) will be used which may be considered clear from intruders.

There are also methods for detecting man-in-the-middle attacks, such as the leap-of-faith method. If there is someone in between during the first connection, then he must be in between during all the subsequent connections to prevent the administrator from being warned that the public key has changed. This could be hard to do for the attacker and therefore a second connection can be set up after a certain timespan to see if there will indeed be a warning.

For the second type of solutions it is necessary to have certain information such that a host can be authenticated. As such, data sent by the host can be authenticated by the administrator to come from this host unaltered. It must be trusted that the part of the information that needs to be secret has not fallen into the wrong hands, though. This is the case with a public and private key pair, in which the private key has to be kept secret from everyone else. Since these cannot be used for authentication, we decided that a hard to guess pre-shared secret (e.g. the system's UUID) would be the best alternative.

We made use of shared secrets in our mechanism so that public keys could be authenticated, which subsequently could be used for secure data transfer. By creating hashes of the sent data concatenated with the secret, both the integrity of the data and its authenticity can be verified. By letting the remote

host verify the administrator's public key using DNSSEC and using this key to encrypt the hashes, it can be prevented that an eavesdropper does not get to see a hash in which the secret has been involved. If the public key could not be verified, a bogus answer can be sent back. Offline attacks to discover the secret will be pointless for the eavesdropper in that case.

Can we make use of existing methods or protocols to realise the possible solutions?

We have seen that most possible solutions to the key retrieval process involve trusted third parties. This is not desirable for this simple application. Soon it became clear that a pre-shared secret was the most feasible solution. The SSH protocol itself can be used in the retrieval mechanism. Using this protocol, an eavesdropper can be detected by comparing the public key received when the SSH connection was initiated and the public key received from the remote host later in the process. If the eavesdropper lets this last key unaltered, the two keys that the administrator received will not match. If he replaces the key with his own key, then there will be a match but then the hash cannot be validated. In both scenarios the administrator will be noticed that something is going on.

The DNS can be used to let the administrator's public key be verified by the remote hosts, by validating the key's fingerprint from the DNS with DNSSEC. If this is done locally and the public key is found to be valid, then it can be certain that a hash encrypted with this public key can only be decrypted by the administrator.

How can these solutions be implemented in a tool that automates the collection of SSH public keys?

We combined existing programs and libraries to implement the mechanism we came up with in a solution that requires a program on the administration machine to contact each host and execute of second program on this host in order to retrieve the public keys in a secure way. The mechanism makes use of the methods and protocols mentioned above. Our implementation also made it possible to automate this process for a list of hosts, given their domain name and a shared secret.

How can we insert the SSH public key fingerprints into the DNS and sign them using DNSSEC in an automated way?

For the BIND installation we used in our proof of concept, the easiest way of pushing dynamic updates to the DNS server was by using the program `nsupdate`. Authentication of the administration machine was enforced by using a pre-shared key and the updates themselves used transaction signatures to ensure authentication and integrity of the SSHFP resource records that needed to be inserted. The `nsupdate` program also makes sure that the new records are signed using DNSSEC, provided that it can find the private key needed for this process.

References

- [1] Ylonen & Lonvick, *The Secure Shell (SSH) Protocol Architecture*, RFC 4251, January 2006, <http://tools.ietf.org/html/rfc4251>.
- [2] Arends, et al., *DNS Security Introduction and Requirements*, RFC 4033, March 2005, <http://tools.ietf.org/html/rfc4033>.
- [3] Arends, et al., *Resource Records for the DNS Security Extensions*, RFC 4034, March 2005, <http://tools.ietf.org/html/rfc4034>.
- [4] Arends, et al., *Protocol Modifications for the DNS Security Extensions*, RFC 4035, March 2005, <http://tools.ietf.org/html/rfc4035>.
- [5] Schlyter & Griffin, *Using DNS to Securely Publish Secure Shell (SSH) Key Fingerprints*, RFC 4255, January 2006, <http://tools.ietf.org/html/rfc4255>.
- [6] Callas, et al., *OpenPGP Message Format*, RFC 4880, November 2007, <http://tools.ietf.org/html/rfc4880>.
- [7] GnuPG project homepage, <http://www.gnupg.org/>.
- [8] Al-Riyami & Paterson, (2003). *Certificateless Public Key Cryptography*, University of London.
- [9] Arkko, Jari, Nikander & Pekka, (2003). *Weak Authentication: How to Authenticate Unknown Principals without Trusted Parties*, Ericsson Research NomadicLab.
- [10] System Management BIOS, DMTF standard DSP0134, July 2010, <http://dmtf.org/standards/smbios>.
- [11] Python/C API Reference Manual, <http://docs.python.org/c-api/>.
- [12] Unbound documentation, <http://www.unbound.net/documentation/libunbound.html>.
- [13] Chandler Wiki: Me Too Crypto, <http://chandlerproject.org/bin/view/Projects/MeTooCrypto>.
- [14] *What is OAEP?*, <http://www.rsa.com/rsalabs/node.asp?id=2346>.
- [15] Vixie, et al., *Secret Key Transaction Authentication for DNS (TSIG)*, RFC 2845, May 2000, <http://tools.ietf.org/html/rfc2845>.
- [16] OpenSSH project homepage, <http://www.openssh.com/>.
- [17] OpenBSD. *OpenBSD Reference Manual (SSH)*. <http://www.openbsd.org/cgi-bin/man.cgi?query=ssh>.

A Program code and configuration files

A.1 For the administration machine

A.1.1 Application

Listing 1: tool_AM.py

```
1  #!/usr/bin/python
2
3  ### imports ###
4  import ConfigParser      # reading config files
5  import argparse         # parsing parameters
6  import subprocess       # spawning new processes
7  import shlex            # determining the correct tokenization for args
8  import hashlib          # computing hashes
9  import sys
10 import os
11 import base64           # base64 encoding/decoding
12 import logging          # will handle the logging of messages
13 import getpass          # password prompt, input is not printed
14 from M2Crypto import RSA
15 sys.path.append("lib")
16 from sshexec import *  # python code to access libssh's C library
17
18 ### default parameters ###
19 logger = None
20 logfile = "tool_AM.log"
21 username = ""
22 password = ""
23 RH_path_program = "tool_RH.py"
24 clear_secrets = ""
25 secrets_path = "secrets/secrets_aes.txt"
26 SSHFP_list = []
27 SSHFP_ttl = 1800
28 DN_DNS = "localhost"
29 DNS_zone = ""
30 DNS_update_file = "DNS_update.tmp"
31 Kpub_RH = ""
32 private_key_DNS_admin = ""
33 ###
34
35 ### functions ###
36 def decryptAES_File(secrets_file, passphrase):
37     global clear_secrets
38     logger.info("decrypting secrets file \"" + secrets_file + "\"...")
39     if os.access(secrets_file, os.F_OK): # if the file exists
40         command = subprocess.Popen(shlex.split("gpg --quiet --yes --logger-file /dev/null
41             --passphrase " + passphrase + " -d " + secrets_file), stdout = subprocess.PIPE)
42         clear_secrets = command.communicate()[0] # put the decrypted file in a global
43             variable
44         if clear_secrets == "":
45             logger.info("wrong passphrase...")
46             error_quit("the secrets file could not be decrypted...")
47         else:
48             logger.info("secrets decrypted...")
49     else:
50         error_quit("the secrets file \"" + secrets_file + "\" can not be accessed..")
51
52 def processList_Of_Hosts():
53     logger.info("start processing hosts...")
54     global clear_secrets
55     records = clear_secrets.splitlines()
56
57     for line in records:
58         processHost(line)
59     logger.info("all hosts processed...")
60
```

```

59 def processHost(record):
60     global username
61     global password
62     global RH_path_program
63
64     host = record.split(":")[0]
65     strong_secret = record.split(":")[1]
66     weak_secret = record.split(":")[2]
67     logger.info("processing host " + host + "...")
68
69     # which secret can be used?
70     secret_type = getSecret_Type(strong_secret, weak_secret)
71
72     # get the public key of the AM
73     public_key = getPublic_Key() # if public key not found -> program exits
74
75     # check parameters
76     allOK = True
77     if username == "":
78         allOK = False
79     if password == "":
80         allOK = False
81     if RH_path_program == "":
82         allOK = False
83     if host == "":
84         allOK = False
85     if secret_type == "":
86         allOK = False
87
88     # contact host
89     if allOK:
90         response = getAnswer_From_RH(RH_path_program, host, username, password, secret_type,
91                                     public_key) # [answers list , exit code]
92         if response is None:
93             logger.error("no valid answer received from remote host...")
94         else:
95             resp_list = response[0]
96             for resp in resp_list:
97                 # process answer, rep: <type>:<hash>:<rsa public key>:<dsa public key>
98                 msg = resp.split(":", 1)[1]
99                 msg_type = resp.split(":")[0]
100                if msg_type == "ERROR":
101                    logger.error(msg)
102                    break
103                elif msg_type == "WARNING":
104                    logger.info("WARNING: " + msg)
105                elif msg_type == "ANSWER":
106                    if secret_type == "strong":
107                        processAnswer(msg, strong_secret, host)
108                    elif secret_type == "weak":
109                        processAnswer(msg, weak_secret, host)
110                    break
111            else:
112                logger.error("one of the parameters was not set...")
113
114 def getSecret_Type(strong, weak):
115     secret_type = ""
116     if not strong == "":
117         secret_type = "strong"
118     elif not weak == "":
119         secret_type = "weak"
120     return secret_type
121
122 def getPublic_Key():
123     logger.info("locating public key...")
124     path_rsa = "/etc/ssh/ssh_host_rsa_key.pub"
125     path_dsa = "/etc/ssh/ssh_host_dsa_key.pub"
126     if os.access(path_rsa, os.F_OK):
127         return readFirst_Line(path_rsa).split()[1]
128     elif os.access(path_dsa, os.F_OK):

```

```

128     return readFirst_Line(path_dsa).split()[1]
129 else:
130     error_quit("the SSH public key file could not be accessed...")
131
132 def readFirst_Line(path):
133     f = open(path, 'r')
134     line = f.readline()
135     f.close()
136     return line
137
138 def getAnswer_From_RH(path, host, uname, passwd, secret_type, public_key):
139     global Kpub_RH
140     answer = None
141     IP_list = domainToIPs(host)
142     if (len(IP_list) == 0):
143         logger.error("domain name could not be resolved to an IP address...")
144         return None
145     IP = IP_list[0]
146     # connect through SSH
147     # need to add a timeout here
148     logger.info("connecting to " + host + " at " + IP)
149     SSH_connection = initConnection(IP)
150     Kpub_RH = SSH_connection[0] # put the key in the global variable
151     if SSH_connection:
152         logger.info("connection established...")
153         # log in
154         if loginPassword(uname, passwd):
155             logger.info("login succeeded...")
156             # execute command
157             answer = execCommand("sudo " + path + " -s " + secret_type + " -k " + public_key)
158
159             if answer is not None and len(answer[0]) == 0:
160                 answer = None
161             if answer is not None:
162                 logger.info("response received...")
163         else:
164             logger.error("login failed; the credentials were not accepted...")
165             # disconnect
166             closeConnection()
167             logger.info("connection closed...")
168         else:
169             logger.error("failed to set up a connection with the remote host...")
170     return answer
171
172 def processAnswer(answer, secret, host):
173     global Kpub_RH
174     key_type = base64.b64decode(Kpub_RH)[4:11]
175     logger.info("processing answer...")
176     logger.debug("answer:\n" + answer)
177
178     # parse answer # <hash>:<rsa public key>:<dsa public key>
179     untrusted_hash = answer.split(":")[0]
180     logger.info("decrypting hash...")
181     untrusted_hash = decryptRSA(untrusted_hash, key_type)
182     untrusted_rsa_key = answer.split(":")[1]
183     untrusted_dsa_key = answer.split(":")[2]
184
185     #compare the public keys
186     key_ok = False
187     if key_type == "ssh-rsa":
188         if Kpub_RH == untrusted_rsa_key:
189             logger.info("rsa public key matched...")
190         else:
191             logger.warning("the public key returned by the remote host doesn't match the key
192                 used to set up the SSH connection. You may be a victim of a man-in-the-middle
193                 attack... ")
194     elif key_type == "ssh-dss":
195         if Kpub_RH == untrusted_dsa_key:
196             logger.info("dsa public key matched...")
197         else:

```



```

196     logger.warning("the public key returned by the remote host doesn't match the key
197         used to set up the SSH connection. You may be a victim of a man-in-the-middle
198         attack... ")
199
200 # calculate the hash with local data
201 trusted_hash = makeHash(secret, untrusted_rsa_key, untrusted_dsa_key)
202 if trusted_hash == untrusted_hash:
203     logger.debug("hash " + trusted_hash + " is trusted...")
204     logger.info("hash is TRUSTED...")
205     # generate SSHFP records
206     makeSSHFP_Records(host, untrusted_rsa_key, untrusted_dsa_key)
207 else:
208     # warn admin
209     logger.warning("the hash received from host \"" + host + "\" is UNTRUSTED! The
210         remote host did NOT proof its knowledge of the secret. You may be a victim of a
211         man-in-the-middle attack, or your public key was not accepted. The retrieved
212         public key(s) won't be pushed to the DNS server..")
213
214 def decryptRSA(msg, key_type):
215     msg = base64.b64decode(msg)
216     Kpriv_AM_path = getPrivate_Key_Path(key_type)
217     try:
218         key = RSA.load_key(Kpriv_AM_path)
219     except:
220         error_quit("unable to load private key (wrong permissions?)")
221     decrypted_hash = key.private_decrypt(msg, RSA.pkcs1_oaep_padding)
222     return decrypted_hash
223
224 def makeHash(secret, rsa, dsa):
225     data = secret + rsa + dsa
226     return hashlib.sha512(data).hexdigest()
227
228 def makeSSHFP_Records(hostname, rsa_key, dsa_key):
229     global SSHFP_list
230     global SSHFP_ttl
231     logger.info("generating SSHFP records...")
232
233     # generate SSHFP records
234     SSHFP_rsa = hostname + " " + SSHFP_ttl + " IN SSHFP 1 1 " +
235         hashlib.sha1(base64.b64decode(rsa_key)).hexdigest()
236     SSHFP_dsa = hostname + " " + SSHFP_ttl + " IN SSHFP 2 1 " +
237         hashlib.sha1(base64.b64decode(dsa_key)).hexdigest()
238
239     logger.info("SSHFP records generated...")
240     logger.debug("SSHFP_rsa: " + SSHFP_rsa)
241     logger.debug("SSHFP_dsa: " + SSHFP_dsa)
242
243     # collect them in a list
244     SSHFP_list.append(SSHFP_rsa)
245     SSHFP_list.append(SSHFP_dsa)
246
247 def processList_Of_SSHFP_records(path):
248     global SSHFP_list
249
250     if os.access(path, os.F_OK):
251         f = open(path, 'r')
252         contents = f.read()
253         logger.debug("list of SSHFP records to push to DNS:\n" + contents)
254         for line in contents.splitlines():
255             SSHFP_list.append(line)
256         f.close()
257         logger.info("list read by program...")
258     else:
259         error_quit("the list of SSHFP records \"" + path + "\" could not be accessed...")
260
261 def testSSHFP_list(SSHFP_list):
262     notEmpty = False
263     if len(SSHFP_list) > 0:
264         notEmpty = True
265     else:

```

```

259     logger.info("no SSHFP records to be pushed to DNS...")
260     return notEmpty
261
262 def makeDNS_Update(path, server, zone, SSHFP_list):
263     logger.info("generating DNS update in temporary file \"" + path + "\"...")
264     f = open(path,"w")
265     f.write("server " + server + "\n")
266     f.write("zone " + zone + "\n")
267     for record in SSHFP_list:
268         f.write("update add " + record + "\n")
269     f.write("show \n")
270     f.write("send \n")
271     f.close()
272
273     # just for debugging
274     f = open(path,"r")
275     logger.debug("update:\n" + f.read())
276
277 def pushSSHFP_records(key, DNS_update):
278     if os.access(key, os.F_OK):
279         logger.info("trying to push SSHFP RR's to the DNS...")
280         command = subprocess.Popen(shlex.split("nsupdate -k " + key + " -v " + DNS_update),
281             stdout = subprocess.PIPE)
282         #response = command.communicate()[0]
283         output = command.communicate()
284         response = output[0]
285
286         # test response for errors
287         status = "ERROR"
288         for line in response.splitlines():
289             if "status: " in line:
290                 status = line.split(",")[1].split(":")[1].strip() # extract the status
291         if status == "NOERROR":
292             logger.info("DNS update was successful...")
293         elif response == "":
294             logger.error("DNS update was NOT successful..")
295             logger.error("no response from DNS server received or the DNS could not be
296                 contacted.")
297         else:
298             logger.error("DNS update was NOT successful..")
299             logger.debug("response:\n" + response)
300
301         # clean up
302         os.remove(DNS_update)
303         logger.info("temporary file \"" + DNS_update + "\" with DNS update removed...")
304     else:
305         logger.info("the private key file \"" + key + "\" could not be accessed,the DNS
306             update will not be executed...")
307
308 def error_quit(msg):
309     logger.error(msg)
310     logger.info("program has terminated...")
311     sys.exit(1)
312
313 def getPrivate_Key_Path(key_type):
314     logger.info("locating private key...")
315     path_rsa = "/etc/ssh/ssh_host_rsa_key"
316     path_dsa = "/etc/ssh/ssh_host_dsa_key"
317     if key_type == "ssh-rsa":
318         if os.access(path_rsa, os.F_OK):
319             return path_rsa
320         else:
321             error_quit("the SSH private key file could not be accessed...")
322     elif key_type == "ssh-dss":
323         if os.access(path_dsa, os.F_OK):
324             return path_dsa
325         else:
326             error_quit("the SSH private key file could not be accessed...")
327
328 ### main program ###

```

```

326 def main():
327     global logger
328     global logfile
329     global username
330     global password
331     global RH_path_program
332     global clear_secrets
333     global secrets_path
334     global SSHFP_list
335     global SSHFP_ttl
336     global DN_DNS
337     global DNS_zone
338     global DNS_update_file
339     global Kpub_RH
340     global private_key_DNS_admin
341
342     # parse arguments #
343     prog_description = "This tool can be used to retrieve the SSH public host keys from
        remote machines and push their fingerprints to a DNS server. If you already have a
        list of SSHFP records, you can feed them to this program and push them to DNS.
        This way you can skip the key retrieval process."
344     arg_parser = argparse.ArgumentParser(description = prog_description)
345
346     arg_parser.add_argument('-l',
347         required = False,
348         default = "",
349         dest = 'SSHFP_RR_list',
350         action = 'store',
351         help = 'The path to a list of SSHFP resource records, ready to push to the DNS
            server.')
```

```

352
353     arg_parser.add_argument('-q',
354         required = False,
355         default = False,
356         dest = 'quiet',
357         action = 'store_const',
358         const = True,
359         help = 'Quiet mode. No output will be printed to stdout.')
```

```

360
361     arg_parser.add_argument('-v',
362         required = False,
363         default = False,
364         dest = 'verbose',
365         action = 'store_const',
366         const = True,
367         help = 'Verbose mode. Debug info will also be printed to stdout.')
```

```

368
369     arg_parser._optionals.title = "flag arguments" # fixes the "optional arguments" in the
        help
370     arguments = arg_parser.parse_args()
371
372     conf = True
373     ## configuration ##
374     # parse config file #
375     config_file = "config/tool_AM.conf"
376     config_parser = ConfigParser.RawConfigParser()
377     if len(config_parser.read(config_file)) > 0:
378         # from config #
379         if config_parser.has_option('secrets file', 'path'):
380             secrets_path = config_parser.get('secrets file', 'path')
381
382         if config_parser.has_option('remote host', 'path to program'):
383             RH_path_program = config_parser.get('remote host', 'path to program')
384
385         if config_parser.has_option('DNS server', 'domain name DNS server'):
386             DN_DNS = config_parser.get('DNS server', 'domain name DNS server')
387
388         if config_parser.has_option('DNS server', 'private key admin'):
389             private_key_DNS_admin = config_parser.get('DNS server', 'private key admin')
390
```

```

391     if config_parser.has_option('DNS server', 'zone file'):
392         DNS_zone = config_parser.get('DNS server', 'zone file')
393
394     if config_parser.has_option('DNS server', 'ttl'):
395         SSHFP_ttl = config_parser.get('DNS server', 'ttl')
396
397     if config_parser.has_option('logging', 'path logfile'):
398         logfile = config_parser.get('logging', 'path logfile')
399 else:
400     conf = False
401
402 # from arguments #
403 SSHFP_list_path = arguments.SSHFP_RR_list
404 quiet = arguments.quiet
405 verbose = arguments.verbose
406
407 # configure logging #
408 # info levels: DEBUG (10) < INFO (20) < WARNING (30) < ERROR (40) < CRITICAL (50)
409 logger = logging.getLogger("standaard_log")
410 logger.setLevel(logging.DEBUG) # lowest level it will log
411 ch_stdout = logging.StreamHandler(sys.stdout)
412 if verbose:
413     ch_stdout.setLevel(logging.DEBUG)
414 else:
415     ch_stdout.setLevel(logging.INFO)
416 fm_stdout = logging.Formatter("%(levelname)s - %(message)s")
417 ch_stdout.setFormatter(fm_stdout)
418
419 ch_file = logging.FileHandler(logfile)
420 ch_file.setLevel(logging.INFO) # lowest level it will log -> omit DEBUG messages
421 fm_file = logging.Formatter("%(asctime)s - %(levelname)s - %(message)s")
422 ch_file.setFormatter(fm_file)
423
424 if not quiet:
425     logger.addHandler(ch_stdout) # log to stdout
426     logger.addHandler(ch_file) # log to file
427
428 if SSHFP_list_path == "":
429     # prompt user for credentials
430     print "\nPlease provide your credentials for the remote hosts."
431     username = raw_input("Username: ")
432     password = getpass.getpass("Password: ")
433     print ""
434     print "Please provide the passphrase to decrypt the secrets file."
435     Kdecrypt = getpass.getpass("Passphrase: ")
436     print ""
437
438 ## program flow ##
439 logger.info("program started...")
440 if not conf:
441     logger.warning("nothing read from configuration file")
442 if SSHFP_list_path == "":
443     logger.info("no SSHFP list provided, the public keys will be retrieved
444                 dynamically...")
445     # decrypt the secrets file
446     decryptAES_File(secrets_path, Kdecrypt)
447     # process each host in the secrets file
448     processList_Of_Hosts()
449     if testSSHFP_list(SSHFP_list):
450         # generate the DNS update command
451         makeDNS_Update(DNS_update_file, DN_DNS, DNS_zone, SSHFP_list)
452         # push the RR's to DNS
453         pushSSHFP_records(private_key_DNS_admin, DNS_update_file)
454 else:
455     logger.info("a list of SSHFP records is provided...")
456     # put the list in the global variable
457     processList_Of_SSHFP_records(SSHFP_list_path)
458     if testSSHFP_list(SSHFP_list):
459         # generate the DNS update command
460         makeDNS_Update(DNS_update_file, DN_DNS, DNS_zone, SSHFP_list)

```

```

460         # push the RR's to DNS
461         pushSSHFP_records(private_key_DNS_admin, DNS_update_file)
462         logger.info("program has terminated...")
463
464 if __name__ == "__main__":
465     main()

```

A.1.2 Configuration file

Listing 2: conf/tool_AM.conf

```

1 [secrets file]
2 path=path/to/secrets/file.txt
3
4 [remote host]
5 path to program=path/to/program.py
6
7 [DNS server]
8 domain name DNS server=dns.domain.org
9 private key admin=path/to/keyfile.private
10 zone file=zone.domain.org
11 ttl=1800 ;ttl for the SSHFP records in ms
12
13 [logging]
14 path logfile=path/to/logfile.log

```

A.1.3 Usage

Listing 3: ./tool_AM.py -h

```

1 usage: tool_AM.py [-h] [-l SSHFP_RR_LIST] [-q] [-v]
2
3 This tool can be used to retrieve the SSH public host keys from remote
4 machines and push their fingerprints to a DNS server. If you already have a
5 list of SSHFP records, you can feed them to this program and push them to DNS.
6 This way you can skip the key retrieval process.
7
8 flag arguments:
9   -h, --help            show this help message and exit
10  -l SSHFP_RR_LIST       The path to a list of SSHFP resource records, ready to
11                        push to the DNS server.
12  -q                    Quiet mode. No output will be printed to stdout.
13  -v                    Verbose mode. Debug info will also be printed to stdout.

```

A.1.4 Python interface to SSH client functionality

Listing 4: lib/source/sshexec.c

```

1 /*
2  ****
3  * sshexec.c (in) sshexec.so (out) *
4  * * *
5  * THIS IS A MODIFIED VERSION OF ssh2_exec.c FROM libssh2's EXAMPLE FILES. *
6  * IT WAS MEANT TO BE COMPILED TO A PYTHON MODULE WITH THE FOLLOWING COMMAND: *
7  * *
8  * gcc -shared -I/usr/include/python2.6/ -lpython2.6 -lssh2 -o sshexec.so sshexec.c *
9  * *
10 * SSH module for Python to execute a command on a remote host. *
11 * At the moment only one connection can exist at a time. *
12 ****
13 */
14
15 #include "libssh2_config.h"
16 #include <libssh2.h>
17 #include <Python.h>
18
19 #ifdef HAVE_WINSOCK2_H

```

```

20 # include <winsock2.h>
21 #endif
22 #ifdef HAVE_SYS_SOCKET_H
23 # include <sys/socket.h>
24 #endif
25 #ifdef HAVE_NETINET_IN_H
26 # include <netinet/in.h>
27 #endif
28 #ifdef HAVE_SYS_SELECT_H
29 # include <sys/select.h>
30 #endif
31 # ifdef HAVE_UNISTD_H
32 #include <unistd.h>
33 #endif
34 #ifdef HAVE_ARPA_INET_H
35 # include <arpa/inet.h>
36 #endif
37
38 #include <sys/time.h>
39 #include <sys/types.h>
40 #include <stdlib.h>
41 #include <fcntl.h>
42 #include <errno.h>
43 #include <stdio.h>
44 #include <ctype.h>
45 #include <netdb.h>
46 #include <unistd.h>
47 #include <pwd.h>
48 #include <string.h>
49 #include <time.h>
50
51 // #define LIBSSH2_ALLOC(session, count) session->alloc((count), &(session)->abstract)
52 #define TIMEOUT 20
53
54 const char *homedir = "";
55 int sock;
56 LIBSSH2_SESSION *session = NULL;
57 int auth = 0;
58
59 /*
60  * (c) Daniel Stenberg
61  *
62  * Found this function at
63  * http://www.mail-archive.com/libssh2-devel@lists.sourceforge.net/msg01630.html
64  */
65 size_t _libssh2_base64_encode(const char *inp, size_t insize, char **outptr) {
66     //extern LIBSSH2_SESSION *session;
67     const char table64[] =
68         "ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789+/";
69     unsigned char ibuf[3];
70     unsigned char obuf[4];
71     int i;
72     int inputparts;
73     char *output;
74     char *base64data;
75     const char *indata = inp;
76
77     *outptr = NULL; /* set to NULL in case of failure before we reach the end */
78
79     if(0 == insize)
80         insize = strlen(indata);
81
82     base64data = output = malloc(insize*4/3+4); //LIBSSH2_ALLOC(session, insize*4/3+4);
83     if(NULL == output)
84         return 0;
85
86     while(insize > 0) {
87         for (i = inputparts = 0; i < 3; i++) {
88             if(insize > 0) {
89                 inputparts++;

```

```

90         ibuf[i] = *indata;
91         indata++;
92         insize--;
93     } else {
94         ibuf[i] = 0;
95     }
96
97 }
98
99 obuf[0] = (unsigned char) ((ibuf[0] & 0xFC) >> 2);
100 obuf[1] = (unsigned char) (((ibuf[0] & 0x03) << 4) | \
101                          ((ibuf[1] & 0xF0) >> 4));
102 obuf[2] = (unsigned char) (((ibuf[1] & 0x0F) << 2) | \
103                          ((ibuf[2] & 0xC0) >> 6));
104 obuf[3] = (unsigned char) (ibuf[2] & 0x3F);
105
106 switch(inputparts) {
107     case 1: /* only one byte read */
108         snprintf(output, 5, "%c%c=",
109                 table64[obuf[0]],
110                 table64[obuf[1]]);
111         break;
112     case 2: /* two bytes read */
113         snprintf(output, 5, "%c%c%c=",
114                 table64[obuf[0]],
115                 table64[obuf[1]],
116                 table64[obuf[2]]);
117         break;
118     default:
119         snprintf(output, 5, "%c%c%c%c",
120                 table64[obuf[0]],
121                 table64[obuf[1]],
122                 table64[obuf[2]],
123                 table64[obuf[3]] );
124         break;
125 }
126 output += 4;
127 }
128
129 *output = 0;
130 *outptr = base64data; /* make it return the actual data memory */
131
132 return strlen(base64data); /* return the length of the new data */
133 }
134
135 static int waitsocket(int socket_fd, LIBSSH2_SESSION *session) {
136     struct timeval timeout;
137     int rc;
138     fd_set fd;
139     fd_set *writefd = NULL;
140     fd_set *readfd = NULL;
141     int dir;
142
143     timeout.tv_sec = 10;
144     timeout.tv_usec = 0;
145
146     FD_ZERO(&fd);
147
148     FD_SET(socket_fd, &fd);
149
150     /* now make sure we wait in the correct direction */
151     dir = libssh2_session_block_directions(session);
152
153     if(dir & LIBSSH2_SESSION_BLOCK_INBOUND)
154         readfd = &fd;
155
156     if(dir & LIBSSH2_SESSION_BLOCK_OUTBOUND)
157         writefd = &fd;
158
159     rc = select(socket_fd + 1, readfd, writefd, NULL, &timeout);

```

```

160
161     return rc;
162 }
163
164 static void closesession(void) {
165     extern int sock;
166     extern LIBSSH2_SESSION *session;
167     libssh2_session_disconnect(session, "Normal disconnect");
168     libssh2_session_free(session);
169     session = NULL;
170     close(sock);
171 }
172
173 static int closechannel(LIBSSH2_CHANNEL *channel, unsigned int to) {
174     extern int sock;
175     int exitcode = 127;
176     int rc;
177     time_t start;
178
179     // Close channel
180     start = time(NULL);
181     while ((rc = libssh2_channel_close(channel)) == LIBSSH2_ERROR_EAGAIN) {
182         // Time-out?
183         if (time(NULL) - start >= to) {
184             break;
185         }
186         waitsocket(sock, session);
187     }
188
189     // Get exit status
190     if (rc == 0) {
191         exitcode = libssh2_channel_get_exit_status(channel);
192     }
193
194     libssh2_channel_free(channel);
195
196     return exitcode;
197 }
198
199 static PyObject* py_domainToIPs(PyObject* self, PyObject* args) {
200     const char *domain;
201     struct hostent *he;
202     int i;
203     PyObject *ip;
204     PyObject *lst;
205
206     // Parse arguments
207     if (!PyArg_ParseTuple(args, "s", &domain)) {
208         return Py_None;
209     }
210
211     // Get addresses for host at domain
212     he = gethostbyname(domain);
213     if (!he) {
214         return PyList_New(0);
215     }
216
217     // Count number of addresses
218     for (i = 0; he->h_addr_list[i]; i++);
219     if (i == 0) {
220         return PyList_New(0);
221     }
222
223     // Create Python list
224     lst = PyList_New(i);
225
226     // Add addresses to list
227     i = 0;
228     while (he->h_addr_list[i]) {
229         ip = PyString_FromString(inet_ntoa(*(struct in_addr*)(he->h_addr_list[i])));

```



```

230         if (!ip) {
231             return Py_None;
232         }
233         PyList_SetItem(lst, i, ip);
234         i++;
235     }
236
237     return lst;
238 }
239
240 static PyObject* py_initConnection(PyObject* self, PyObject* args) {
241     extern int sock;
242     extern LIBSSH2_SESSION *session;
243     const char *ip;
244     char *khp = "/.ssh/known_hosts";
245     unsigned int to = TIMEOUT;
246     time_t start;
247     char *kh;
248     char check = 0;
249     unsigned long hostaddr;
250     struct sockaddr_in sin;
251     LIBSSH2_KNOWNHOSTS *nh;
252     int rc;
253     size_t len;
254     int type;
255     const char *key;
256     char *key_base64;
257     struct libssh2_knownhost *host;
258     PyObject *mismatch = Py_False;
259     PyObject *ret;
260
261     // Parse arguments
262     if (!PyArg_ParseTuple(args, "s|sI", &ip, &khp, &to)) {
263         return Py_None;
264     }
265
266     // Check if a session has already been initiated
267     if (session) {
268         return Py_None;
269     }
270
271     if (strcmp(khp, "/.ssh/known_hosts")) {
272         kh = khp;
273     } else {
274         kh = malloc(strlen(homedir)+strlen("/.ssh/known_hosts")+1);
275         strcpy(kh, homedir);
276         strcat(kh, khp);
277         check |= 1;
278     }
279
280     // Create socket and connect
281     hostaddr = inet_addr(ip);
282     sock = socket(AF_INET, SOCK_STREAM, 0);
283     sin.sin_family = AF_INET;
284     sin.sin_port = htons(22);
285     sin.sin_addr.s_addr = hostaddr;
286     if (connect(sock, (struct sockaddr*)&sin,
287               sizeof(struct sockaddr_in)) != 0) {
288         return Py_None;
289     }
290
291     // Create a session instance
292     session = libssh2_session_init();
293     if (!session) {
294         close(sock);
295         return Py_None;
296     }
297
298     // Tell libssh2 we want it all done non-blocking
299     libssh2_session_set_blocking(session, 0);

```

```

300
301 // Start it up. This will trade welcome banners, exchange keys,
302 // and setup crypto, compression, and MAC layers
303 start = time(NULL);
304 while ((rc = libssh2_session_startup(session, sock)) ==
305         LIBSSH2_ERROR_EAGAIN) {
306     // Time-out?
307     if (time(NULL) - start >= to) {
308         closesession();
309         return Py_None;
310     }
311 }
312 if (rc) {
313     closesession();
314     return Py_None;
315 }
316
317 // Check if the host's key is in the known-hosts file
318 nh = libssh2_knownhost_init(session);
319 if (!nh) {
320     closesession();
321     return Py_None;
322 }
323 key = libssh2_session_hostkey(session, &len, &type);
324 libssh2_knownhost_readfile(nh, kh, LIBSSH2_KNOWNHOST_FILE_OPENSSH);
325 if (check & 1) {
326     free(kh);
327 }
328 if (key) {
329     if (libssh2_knownhost_check(nh, (char *)ip, (char *)key, len,
330                                LIBSSH2_KNOWNHOST_TYPE_PLAIN |
331                                LIBSSH2_KNOWNHOST_KEYENC_RAW,
332                                &host) ==
333                                LIBSSH2_KNOWNHOST_CHECK_MISMATCH) {
334         mismatch = Py_True;
335     }
336 } else {
337     closesession();
338     libssh2_knownhost_free(nh);
339     return Py_None;
340 }
341 libssh2_knownhost_free(nh);
342
343 // Convert binary key into base64 format and return it
344 _libssh2_base64_encode(key, len, &key_base64);
345 ret = Py_BuildValue("(s,0)", key_base64, mismatch);
346 free(key_base64);
347 return ret;
348 }
349
350 static PyObject* py_loginPassword(PyObject* self, PyObject* args) {
351     extern LIBSSH2_SESSION *session;
352     extern int auth;
353     const char *username;
354     const char *password;
355     unsigned int to = TIMEOUT;
356     time_t start;
357     int rc;
358
359     // Parse arguments
360     if (!PyArg_ParseTuple(args, "ss|I", &username, &password, &to)) {
361         return Py_None;
362     }
363
364     // Check if there is an active session
365     if (!session) {
366         return Py_None;
367     }
368
369     // Try password login

```

```

370     start = time(NULL);
371     while ((rc = libssh2_userauth_password(session, username, password)) ==
372            LIBSSH2_ERROR_EAGAIN);
373         // Time-out?
374         if (time(NULL) - start >= to) {
375             return Py_False;
376         }
377     if (rc) {
378         return Py_False;
379     }
380
381     auth = 1;
382     return Py_True;
383 }
384
385 static PyObject* py_loginPublicKey(PyObject* self, PyObject* args) {
386     extern LIBSSH2_SESSION *session;
387     extern int auth;
388     const char *username;
389     char *puk = "/.ssh/id_rsa.pub";
390     char *pvk = "/.ssh/id_rsa";
391     char *pub;
392     char *prv;
393     char check = 0;
394     const char *passphrase = "";
395     unsigned int to = TIMEOUT;
396     time_t start;
397     int rc;
398
399     // Parse arguments
400     if (!PyArg_ParseTuple(args, "s|sssI", &username, &puk, &pvk, &passphrase, &to)) {
401         return Py_None;
402     }
403
404     // Check if there is an active session
405     if (!session) {
406         return Py_None;
407     }
408
409     // Construct path to public key
410     if (strcmp(puk, "/.ssh/id_rsa.pub")) {
411         pub = puk;
412     } else {
413         pub = malloc(strlen(homedir)+strlen("/.ssh/id_rsa.pub")+1);
414         strcpy(pub, homedir);
415         strcat(pub, puk);
416         check |= 1;
417     }
418
419     // Construct path to private key
420     if (strcmp(pvk, "/.ssh/id_rsa")) {
421         prv = pvk;
422     } else {
423         prv = malloc(strlen(homedir)+strlen("/.ssh/id_rsa")+1);
424         strcpy(prv, homedir);
425         strcat(prv, pvk);
426         check |= 2;
427     }
428
429     // Try public key login
430     start = time(NULL);
431     while ((rc = libssh2_userauth_publickey_fromfile(session, username, pub,
432            prv, passphrase)) ==
433            LIBSSH2_ERROR_EAGAIN) {
434
435         // Time-out?
436         if (time(NULL) - start >= to) {
437             // Free memory
438             if (check & 1) {
439                 free(pub);

```

```

440         if (check & 2) {
441             free(prv);
442         }
443         return Py_False;
444     }
445 }
446
447 // Free memory
448 if (check & 1) {
449     free(pub);
450 }
451 if (check & 2) {
452     free(prv);
453 }
454
455 // Check if succeeded
456 if (rc) {
457     return Py_False;
458 }
459
460 auth = 1;
461 return Py_True;
462 }
463
464 static PyObject* py_execCommand(PyObject* self, PyObject* args) {
465     extern int sock;
466     extern LIBSSH2_SESSION *session;
467     extern int auth;
468     const char *command;
469     unsigned int to = TIMEOUT;
470     time_t start;
471     int rc;
472     char buffer[0x4000];
473     int pos;
474     int exitcode;
475     int i;
476     int j;
477     LIBSSH2_CHANNEL *channel;
478     int lenanswers = 10;
479     char **answers;
480     char **temp;
481     int numanswers = 0;
482     PyObject *lst;
483
484     // Parse arguments
485     if (!PyArg_ParseTuple(args, "s|I", &command, &to)) {
486         return Py_None;
487     }
488
489     // Check if there is an active session and if the user has been logged in
490     if (!session || !auth) {
491         return Py_None;
492     }
493
494     answers = malloc(lenanswers*sizeof(char*));
495     if (answers == NULL) {
496         return Py_None;
497     }
498
499     // Exec non-blocking on the remote host
500     start = time(NULL);
501     while ((channel = libssh2_channel_open_session(session)) == NULL &&
502            libssh2_session_last_error(session, NULL, NULL, 0) ==
503            LIBSSH2_ERROR_EAGAIN) {
504         // Time-out?
505         if (time(NULL) - start >= to) {
506             if (channel != NULL) {
507                 closechannel(channel, to);
508             }
509             free(answers);

```

```

510         return Py_None;
511     }
512     waitsocket(sock, session);
513 }
514 if (channel == NULL) {
515     free(answers);
516     return Py_None;
517 }
518
519 // Execute command
520 start = time(NULL);
521 while ((rc = libssh2_channel_exec(channel, command)) ==
522        LIBSSH2_ERROR_EAGAIN) {
523     // Time-out?
524     if (time(NULL) - start >= to) {
525         closechannel(channel, to);
526         free(answers);
527         return Py_None;
528     }
529     waitsocket(sock, session);
530 }
531 if (rc != 0) {
532     closechannel(channel, to);
533     free(answers);
534     return Py_None;
535 }
536
537 // Loop until all answers have been received
538 start = time(NULL);
539 for (;;) {
540     // Loop until we block
541     do {
542         rc = libssh2_channel_read(channel, buffer, sizeof(buffer));
543         if (rc > 0) {
544             i = j = 0;
545
546             // Split answer on newlines and put every substring in the
547             // answers array
548             while (j < rc) {
549                 for (; buffer[j] != '\n' && j < rc; j++);
550                 pos = numanswers;
551                 numanswers++;
552
553                 // Check if there still is enough memory
554                 if (numanswers > lenanswers) {
555                     lenanswers *= 2;
556                     temp = realloc(answers, lenanswers*sizeof(char*));
557
558                     // If realloc failed, free memory and return
559                     if (temp == NULL) {
560                         numanswers--;
561                         for (i = 0; i < numanswers; i++) {
562                             free(answers[i]);
563                         }
564                         free(answers);
565                         closechannel(channel, to);
566                         return Py_None;
567                     }
568                     answers = temp;
569                 }
570                 answers[pos] = malloc((j-i+1)*sizeof(char));
571                 strncpy(answers[pos], &buffer[i], (j-i));
572                 answers[pos][j-i] = '\0';
573                 j++;
574                 i = j;
575             }
576         }
577     } while (rc > 0);
578 }
579

```

```

580         // This is due to blocking that would occur otherwise so we loop on
581         // this condition
582         if (rc == LIBSSH2_ERROR_EAGAIN) {
583             // Time-out?
584             if (time(NULL) - start >= to) {
585                 closechannel(channel, to);
586                 for (i = 0; i < numanswers; i++) {
587                     free(answers[i]);
588                 }
589                 free(answers);
590                 return Py_None;
591             }
592             waitsocket(sock, session);
593         } else {
594             break;
595         }
596     }
597
598     // Close channel
599     exitcode = closechannel(channel, to);
600
601     // Create Python list
602     lst = PyList_New(numanswers);
603
604     // Convert answers
605     for (i = 0; i < numanswers; i++) {
606         PyList_SetItem(lst, i, PyString_FromString(answers[i]));
607         free(answers[i]);
608     }
609     free(answers);
610
611     return Py_BuildValue("(0,i)", lst, exitcode);
612 }
613
614 static PyObject* py_closeConnection(PyObject* self, PyObject* args) {
615     extern LIBSSH2_SESSION *session;
616     extern int auth;
617
618     // Check if there is an active session
619     if (!session) {
620         return Py_False;
621     }
622
623     closesession();
624     auth = 0;
625
626     return Py_True;
627 }
628
629 static PyMethodDef sshexec_methods[] = {
630     {"domainToIPs", py_domainToIPs, METH_VARARGS},
631     {"initConnection", py_initConnection, METH_VARARGS},
632     {"loginPassword", py_loginPassword, METH_VARARGS},
633     {"loginPublicKey", py_loginPublicKey, METH_VARARGS},
634     {"execCommand", py_execCommand, METH_VARARGS},
635     {"closeConnection", py_closeConnection, METH_VARARGS},
636     {NULL, NULL}
637 };
638
639 void initsshexec() {
640     extern const char *homedir;
641     struct passwd *pw;
642
643     (void) Py_InitModule("sshexec", sshexec_methods);
644
645     // Get user's home directory
646     pw = getpwuid(getuid());
647     homedir = pw->pw_dir;
648 }

```

A.2 For the remote host

A.2.1 Application

Listing 5: tool_RH.py

```
1  #!/usr/bin/python
2
3  ### imports ###
4  import ConfigParser      # reading config files
5  import argparse         # parsing parameters
6  import subprocess       # spawning new processes
7  import shlex            # determining the correct tokenization for args
8  import hashlib          # computing hashes
9  import sys
10 import os
11 import string
12 import base64           # base64 encoding/decoding
13 import random
14 import math
15 import struct
16 from M2Crypto import RSA, DSA
17 from unbound import ub_ctx, RR_TYPE_SSHFP, RR_CLASS_IN
18
19 ### default parameters ###
20 TOOL_CONF = "conf/tool_RH.conf"
21 RESOLV_CONF = "/etc/resolv.conf"
22 TRUSTED_KEY = "/etc/unbound/root.key"
23 HOST_KEYS = "/etc/ssh"
24 AM_DOMAIN = "localhost"
25
26 ### functions ###
27 def warning(msg):
28     print "WARNING:" + msg
29
30 def error(msg):
31     print "ERROR:" + msg
32     sys.exit(1)
33
34 def answer(digest, rsa_key, dsa_key, am_key):
35     print "ANSWER:" + encrypt(digest, am_key) + ":" + rsa_key + ":" + dsa_key
36     sys.exit(0)
37
38 def encrypt(msg, key):
39     key = base64.b64decode(key)
40     fields = []
41
42     sb = key[0:4]
43     if len(sb) != 4:
44         error("bad key")
45     sd = struct.unpack(">I", sb)[0]
46     type = key[4:4+sd]
47     if len(type) != sd:
48         error("bad key")
49
50     if type == "ssh-dss":
51         error("RSA key required") # DSA cannot be used for encryption/decryption
52     elif type != "ssh-rsa":
53         error("bad key")
54
55     # Extract exponent and modulus
56     s = 4 + sd
57     for i in range(2):
58         sb = key[s:s+4]
59         if len(sb) != 4:
60             error("bad key")
61         sd = struct.unpack(">I", sb)[0]
62         val = key[s+4:s+4+sd]
63         if len(val) != sd:
```

```

64     error("bad key")
65     fields.append(sb + val)
66     s += 4 + sd
67
68     e = fields[0]
69     n = fields[1]
70
71     key = RSA.new_pub_key((e, n))
72
73     return base64.b64encode(key.public_encrypt(msg, RSA.pkcs1_oaep_padding))
74
75 def getRandomString(length):
76     return ''.join(random.choice(string.printable) for x in range(length))
77
78 def getSystemUUID():
79     command = subprocess.Popen(shlex.split('dmidecode -s system-uuid'),
80                                stdout=subprocess.PIPE)
81     return command.communicate()[0].rstrip()
82
83 def getSystemProductName():
84     command = subprocess.Popen(shlex.split('dmidecode -s system-product-name'),
85                                stdout=subprocess.PIPE)
86     return command.communicate()[0].rstrip()
87
88 # not required according to SMBIOS specification
89 def getMotherboardSerial():
90     command = subprocess.Popen(shlex.split('dmidecode -s baseboard-serial-number'),
91                                stdout=subprocess.PIPE)
92     return command.communicate()[0].rstrip()
93
94 def makeHash(secret, rsa_key, dsa_key):
95     secret += rsa_key + dsa_key
96     return hashlib.sha512(secret).hexdigest()
97
98 def getPublicKey_rsa():
99     try:
100        f = open(HOST_KEYS + '/ssh_host_rsa_key.pub', 'r')
101        key = f.readline().split()[1]
102    except IOError:
103        return ""
104    except:
105        key = ""
106
107    f.close()
108    return key
109
110 def getPublicKey_dsa():
111     try:
112        f = open(HOST_KEYS + '/ssh_host_dsa_key.pub', 'r')
113        key = f.readline().split()[1]
114    except IOError:
115        return ""
116    except:
117        key = ""
118
119    f.close()
120    return key
121
122 def getStrong_Secret():
123     return getSystemUUID()
124
125 def getWeak_Secret():
126     # motherboard_serial+system_product_name
127     return getMotherboardSerial()+getSystemProductName()
128
129 def getBogus_Secret():
130     # random string, with padding to minimize collisions
131     return "~@$$^*" + getRandomString(128) + " '!#%&("
132
133 def getSecretHash(secret_type, rsa_key, dsa_key):

```



```

131 secret=""
132 if secret_type == "strong":
133     secret=getStrong_Secret()
134 elif secret_type == "weak":
135     secret=getWeak_Secret()
136 elif secret_type == "bogus":
137     secret=getBogus_Secret()
138 else:
139     error("wrong type of secret")
140
141 if not secret:
142     error("wrong permissions")
143
144 return makeHash(secret, rsa_key, dsa_key)
145
146 def checkPublic_Key_AM(key, domain):
147     # validate the public key with the SSHFP record
148
149     types = {"ssh-rsa": 1, "ssh-dss": 2}
150
151     try:
152         key = base64.b64decode(key)
153     except:
154         error("bad key")
155
156     # Get key type
157     keytype = key[4:11]
158
159     if keytype not in types:
160         return False
161
162     keytype = types[keytype]
163
164     # Get key hash
165     digest = hashlib.sha1(key).hexdigest()
166
167     # Init Unbound
168     ctx = ub_ctx()
169     ctx.resolvconf(RESOLV_CONF)
170
171     # Read trusted (root) public key for DNSSEC validation
172     if (os.path.isfile(TRUSTED_KEY)):
173         ctx.add_ta_file(TRUSTED_KEY)
174
175     # Resolve SSHFP records for the domain name
176     status, result = ctx.resolve(domain, RR_TYPE_SSHFP, RR_CLASS_IN)
177
178     # Check if resolving succeeded and if the DNSSEC validation was positive
179     if status == 0 and result.havedata and result.secure:
180         sshfp = dict()
181
182         # Loop through the resolved SSHFP records
183         for record in result.data.address_list:
184             fp = record.split(".")
185
186             # Get public key type and digest type
187             pub = int(fp.pop(0))
188             dig = int(fp.pop(0))
189
190             # Digest algorithm must be SHA1; also no need to compute unused key types
191             if dig != 1 or pub != keytype:
192                 continue
193
194             conv = ""
195
196             # Convert FP from decimal to hexadecimal string
197             for num in fp:
198                 h = hex(int(num))[2:]
199                 if len(h) == 1:
200                     h = "0"+h

```

```

201         conv += h
202
203     # Store FP
204     if pub not in sshfp:
205         sshfp[pub] = []
206         sshfp[pub].append(conv)
207
208     # See if the fingerprints match
209     if digest in sshfp[keytype]:
210         return True
211
212     return False
213
214 ### main program ###
215 def main():
216     global TOOL_CONF
217     global RESOLV_CONF
218     global TRUSTED_KEY
219     global HOST_KEYS
220     global AM_DOMAIN
221
222     # parse arguments #
223     prog_description = "This tool will return the secret of this machine."
224     arg_parser = argparse.ArgumentParser(description=prog_description)
225     arg_parser.add_argument('-s',
226                             choices=['strong', 'weak'],
227                             required=True,
228                             dest='type_secret',
229                             action='store',
230                             help='The type of secret that must be returned "strong" or "weak".')
231     arg_parser.add_argument('-k',
232                             required=True,
233                             dest='rsa_public_key',
234                             action='store',
235                             help='The client\'s public key.')
236     arg_parser.add_argument('-c',
237                             required=False,
238                             default=sys.path[0]+"/"+TOOL_CONF,
239                             dest='path_to_conf',
240                             action='store',
241                             help='The path of the configuration file.')
242
243     arg_parser._optionals.title = "flag arguments" # fixes the "optional arguments" in the
244     help
245     arguments=arg_parser.parse_args()
246
247     ## configuration ##
248     # parse config file #
249     TOOL_CONF = arguments.path_to_conf
250     config_parser = ConfigParser.RawConfigParser()
251     if len(config_parser.read(TOOL_CONF)) > 0:
252         # from config #
253         if config_parser.has_option('administration machine', 'domain_name'):
254             AM_DOMAIN = config_parser.get('administration machine', 'domain_name')
255
256         if config_parser.has_option('key files', 'host_keys'):
257             HOST_KEYS = config_parser.get('key files', 'host_keys')
258
259         if config_parser.has_option('config files', 'resolv_conf'):
260             RESOLV_CONF = config_parser.get('config files', 'resolv_conf')
261
262         if config_parser.has_option('key files', 'trusted_key'):
263             TRUSTED_KEY = config_parser.get('key files', 'trusted_key')
264     else:
265         warning("nothing read from configuration file")
266
267     # from arguments #
268     Kpub_AM = arguments.rsa_public_key
269     TypeSecret = arguments.type_secret

```

```

270 # program flow #
271 rsa_key = getPublicKey_rsa()
272 dsa_key = getPublicKey_dsa()
273
274 if not rsa_key and not dsa_key:
275     error("no host key(s) found")
276
277 if checkPublic_Key_AM(Kpub_AM, AM_DOMAIN):
278     # return secret
279     answer(getSecretHash(TypeSecret, rsa_key, dsa_key), rsa_key, dsa_key, Kpub_AM)
280 else:
281     # return bogus answer
282     answer(getSecretHash("bogus", rsa_key, dsa_key), rsa_key, dsa_key, Kpub_AM)
283
284 if __name__ == "__main__":
285     main()

```

A.2.2 Configuration file

Listing 6: conf/tool_RH.conf

```

1 [administration machine]
2 domain_name=admin.domain.org
3
4 [key files]
5 host_keys=/etc/ssh
6 trusted_key=/etc/unbound/root.key
7
8 [config files]
9 resolv_conf=/etc/resolv.conf

```