

Public-Key Infrastructure (PKI) Lab

Copyright © 2018 by Wenliang Du.

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. If you remix, transform, or build upon the material, this copyright notice must be left intact, or reproduced in a way that is reasonable to the medium in which the work is being re-published.

1 Overview

Public key cryptography is the foundation of today's secure communication, but it is subject to man-in-the-middle attacks when one side of communication sends its public key to the other side. The fundamental problem is that there is no easy way to verify the ownership of a public key, i.e., given a public key and its claimed owner information, how do we ensure that the public key is indeed owned by the claimed owner? The Public Key Infrastructure (PKI) is a practical solution to this problem.

The learning objective of this lab is for students to gain the first-hand experience on PKI. SEED labs have a series of labs focusing on the public-key cryptography, and this one focuses on PKI. By doing the tasks in this lab, students should be able to gain a better understanding of how PKI works, how PKI is used to protect the Web, and how Man-in-the-middle attacks can be defeated by PKI. Moreover, students will be able to understand the root of the trust in the public-key infrastructure, and what problems will arise if the root trust is broken. This lab covers the following topics:

- Public-key encryption, Public-Key Infrastructure (PKI)
- Certificate Authority (CA), X.509 certificate, and root CA
- Apache, HTTP, and HTTPS
- Man-in-the-middle attacks

You may work with one partner on this lab.

Readings. Detailed coverage of PKI can be found in the following:

- Chapter 27, *Public Key Infrastructure*, of the SEED Book, *Computer & Internet Security: A Hands-on Approach*, 3rd Edition, by Wenliang Du. See details at [Hands on Security](#).

Related labs. A topic related to this lab is the Transport Layer Security (TLS), which is based on PKI. We have a separate lab for TLS. In addition, we have a lab called *RSA Public-Key Encryption and Signature Lab*, which focuses on the algorithm part of the public-key cryptography.

2 Submission

Submit a lab report in Canvas as a Word document. Your lab report should contain two sections:

- **Results Summary** Summarize the results/observations from each of the lab's tasks. Several questions in this lab are boxed and boldfaced. This section of your lab report should answer these questions. This section of your lab report may be written jointly with your partner(s). Your lab report must contain this section, even if you wrote it jointly with your partner(s).

- **Lessons Learned** Describe your major takeaways from completing this lab, using your results to illustrate your takeaways. I expect you to put significant thought into this section, as I am looking for you to reflect on the experience of working on the lab. Any sources that you use other than your lab results and the textbook must be carefully cited. This section of your lab report must be written entirely individually by you.

3 Lab Environment

Run these commands before starting this lab:

```
cd seed-labs
git pull upstream main
git push
```

In this lab, we will generate public-key certificates, and then use them to secure web servers. The certificate generation tasks will be conducted on the VM, but we will use a container to host the web server. The files for this lab are in the directory rooted at `category-crypto/Crypto_PKI` from `seed-labs`.

DNS setup. In this document, we use `www.bank32.com` as an example to show how to set up an HTTPS web server with this name. Students need to use a different name for their lab. Include your last name and the current year in the server name. For example, if Tom Kelliher does this lab in 2024, the server name should be `www.kelliher2024.com`. You do not need to own this domain; you just need to map this name to the container's IP address by adding the following entries to `/etc/hosts` (the first entry is required, otherwise, the example in this lab description will not work):

```
10.9.0.80 www.bank32.com
10.9.0.80 www.kelliher2024.com
```

`/etc/hosts` can only be written by the root user. Therefore, the command you use to edit this file must be run via `sudo`.

4 Lab Tasks

4.1 Task 1: Becoming a Certificate Authority (CA)

A Certificate Authority (CA) is a trusted entity that issues digital certificates. The digital certificate certifies the ownership of a public key by the named subject of the certificate. A number of commercial CAs are treated as root CAs; VeriSign is the largest CA at the time of writing. Users who want to get digital certificates issued by the commercial CAs need to pay those CAs.

In this lab, we need to create digital certificates, but we are not going to pay any commercial CA. We will become a root CA ourselves, and then use this CA to issue certificate for others (e.g. servers). In this task, we will make ourselves a root CA, and generate a certificate for this CA. Unlike other certificates, which are usually signed by another CA, the root CA's certificates are self-signed. Root CA's certificates are usually pre-loaded into most operating systems, web browsers, and other software that rely on PKI. Root CA's certificates are unconditionally trusted.

The Configuration File `openssl.conf`. Use `cd` to enter the `CA_Files` directory for the work of creating certificates and keys.

In order to use OpenSSL to create certificates, you have to have a configuration file. The configuration file usually has an extension `.cnf`. It is used by three OpenSSL commands: `ca`, `req` and `x509`. The manual page of `openssl.cnf` can be found from online resources. By default, OpenSSL uses the configuration file from `/usr/lib/ssl/openssl.cnf`. Since we need to make changes to this file, we will copy it into our current directory as `myCA-openssl.cnf`, and instruct OpenSSL to use this copy instead.

The `[CA_default]` section of the configuration file shows the default setting that we need to prepare. We need to create several sub-directories. Please uncomment the `unique_subject` line to allow creation of certifications with the same subject, because it is very likely that we will do that in the lab.

Listing 1: Default CA setting

```
[ CA_default ]
dir           = ./demoCA           # Where everything is kept
certs        = $dir/certs         # Where the issued certs are kept
crl_dir      = $dir/crl           # Where the issued crl are kept
database     = $dir/index.txt     # database index file.
#unique_subject = no              # Set to 'no' to allow creation of
                                # several certs with same subject.
new_certs_dir = $dir/newcerts     # default place for new certs.
serial       = $dir/serial        # The current serial number
```

For the `index.txt` file, simply create an empty file in the `demoCA` directory, which you will also have to create. You will also have to create a `newcerts` directory as a sub-directory of the `demoCA` directory. For the `serial` file, put a single number in string format (e.g. 1000) in the file. Once you have set up the configuration file `openssl.cnf`, you can create and issue certificates.

Certificate Authority (CA). As we described before, we need to generate a self-signed certificate for our CA. This means that this CA is totally trusted, and its certificate will serve as the root certificate. You can run the following command to generate the self-signed certificate for the CA:

```
openssl req -x509 -newkey rsa:4096 -sha256 -days 3650 \
            -keyout ca.key -out ca.crt
```

You will be prompted for a password. Do not lose this password, because you will have to type the passphrase each time you want to use this CA to sign certificates for others. You will also be asked to fill in the subject information, such as the Country Name, Common Name, etc. The output of the command are stored in two files: `ca.key` and `ca.crt`. The file `ca.key` contains the CA's private key, while `ca.crt` contains the public-key certificate.

We can use the following commands to look at the decoded content of the X509 certificate and the RSA key (`-text` means decoding the content into plain text; `-noout` means not printing out the encoded version):

```
openssl x509 -in ca.crt -text -noout
openssl rsa -in ca.key -text -noout
```

Run the above commands. Insert their output into your lab report. From the output, please identify the following:

- What part of the certificate indicates this is a CA's certificate?

- What part of the certificate indicates this is a self-signed certificate?

4.2 Task 2: Generating a Certificate Request for Your Web Server

A company called `bank32.com` (replace this with the name of your own web server) wants to get a public-key certificate from our CA. First it needs to generate a Certificate Signing Request (CSR), which basically includes the company's public key and identity information. The CSR will be sent to the CA, who will verify the identity information in the request, and then generate a certificate.

The command to generate a CSR is quite similar to the one we used in creating the self-signed certificate for the CA. The only difference is the `-x509` option. Without it, the command generates a request; with it, the command generates a self-signed certificate. The following command generate a CSR for `www.bank32.com` (you should use your own server name):

```
openssl req -newkey rsa:2048 -sha256 \
            -keyout server.key -out server.csr \
            -nodes
```

The command will generate a pair of public/private key, and then create a certificate signing request from the public key. We can use the following command to look at the decoded content of the CSR and private key files:

```
openssl req -in server.csr -text -noout
openssl rsa -in server.key -text -noout
```

Adding Alternative names. Many websites have different URLs. For example, `www.example.com`, `example.com`, `example.net`, and `example.org` are all pointing to the same web server. Due to the hostname matching policy enforced by browsers, the common name in a certificate must match with the server's hostname, or browsers will refuse to communicate with the server.

To allow a certificate to have multiple names, the X.509 specification defines extensions to be attached to a certificate. This extension is called Subject Alternative Name (SAN). Using the SAN extension, it's possible to specify several hostnames in the `subjectAltName` field of a certificate.

To generate a certificate signing request with such a field, we can put all the necessary information in a configuration file or at the command line. We will use the command-line approach in this task (the configuration file approach is used in another SEED lab, the TLS lab). We can add the following option to the "openssl req" command. It should be noted that the `subjectAltName` extension field must also include the one from the common name field; otherwise, the common name will not be accepted as a valid name.

```
-addext "subjectAltName = DNS:www.bank32.com, \
        DNS:www.bank32A.com, \
        DNS:www.bank32B.com"
```

Please add two alternative names to your certificate signing request. They will be needed in the tasks later.

Run the commands below and insert their output into your lab report.

```
openssl req -in server.csr -text -noout
openssl rsa -in server.key -text -noout
```

4.3 Task 3: Generating a Certificate for your server

The CSR file needs to have the CA's signature to form a certificate. In the real world, the CSR files are usually sent to a trusted CA for their signature. In this lab, we will use our own trusted CA to generate certificates. The following command turns the certificate signing request (`server.csr`) into an X509 certificate (`server.crt`), using the CA's `ca.crt` and `ca.key`:

```
openssl ca -config myCA_openssl.cnf -policy policy_anything \  
  -md sha256 -days 3650 \  
  -in server.csr -out server.crt -batch \  
  -cert ca.crt -keyfile ca.key
```

In the above command, `myCA_openssl.cnf` is the configuration file we copied from `/usr/lib/ssl/openssl.cnf` (we also made changes to this file in Task 1). We use the `policy_anything` policy defined in the configuration file. This is not the default policy; the default policy has more restriction, requiring some of the subject information in the request to match those in the CA's certificate. The policy used in the command, as indicated by its name, does not enforce any matching rule.

Copy the extension field. For security reasons, the default setting in `openssl.cnf` does not allow the "openssl ca" command to copy the extension field from the request to the final certificate. To enable that, we can go to our copy of the configuration file, uncomment the following line:

```
# Extension copying option: use with caution.  
copy_extensions = copy
```

After signing the certificate, please use the following command to print out the decoded content of the certificate, and check whether the alternative names are included.

```
openssl x509 -in server.crt -text -noout
```

Insert the output of the command above into your lab report.

4.4 Task 4: Deploying Certificate in an Apache-Based HTTPS Website

You will use a Docker container in this task and the following tasks. The Apache web server runs in this container. See the *Container Setup and Commands* Subsection at the end of this lab for a summary of the instructions. You will need to rebuild your container each time you change the container's `Dockerfile` or the files that are copied into the container.

In this task, we will see how public-key certificates are used by websites to secure web browsing. We will set up an HTTPS website based Apache. The Apache server, which is already installed in our container, supports the HTTPS protocol. To create an HTTPS website, we just need to configure the Apache server, so it knows where to get the private key and certificates. Inside our container, we have already set up an HTTPS site for `bank32.com`. Students can follow this example to set up their own HTTPS site.

An Apache server can simultaneously host multiple websites. It needs to know the directory where a website's files are stored. This is done via its `VirtualHost` file, located in the `/etc/apache2/sites-available` directory in the container. In our container, we have a file called `bank32_apache_ssl.conf`, which contains the following entry:

```
<VirtualHost *:443>  
  DocumentRoot /var/www/bank32
```

```
ServerName www.bank32.com
ServerAlias www.bank32A.com
ServerAlias www.bank32B.com
DirectoryIndex index.html
SSLEngine On
SSLCertificateFile /certs/bank32.crt ①
SSLCertificateKeyFile /certs/bank32.key ②
</VirtualHost>
```

The above example sets up the HTTPS site `https://www.bank32.com` (port 443 is the default HTTPS port). `bank32_apache_ssl.conf` is in the `image_www` directory. If you look at the `Dockerfile`, you will see the command that copies this file into its proper place in the container:

```
COPY ./bank32_apache_ssl.conf /etc/apache2/sites-available
```

Within `bank32_apache_ssl.conf`, the `ServerName` entry specifies the name of the website, while the `DocumentRoot` entry specifies where the files for the website are stored. Using the `ServerAlias` entries, we allow the website to have different names. You should also provide two alias entries.

We also need to tell Apache where the server certificate (Line ①) and private key (Line ②) are stored. In the `Dockerfile`, we have already included the commands to copy the certificate and key to the `/certs` folder of the container.

This file is in the `image_www` directory. Copy the file, naming the copy to correspond to your web site's name. For example, I named my copy `kelliher2024_apache_sss.conf`. There are comments in the file to help guide you. You will also have to add several lines to `Dockerfile` in the same directory. Again, there are comments in this file to help you.

In order to make the website work, we need to enable Apache's `ssl` module and then enable this site. They can be done using the following commands, which are already executed when the container is built.

```
# a2enmod ssl // Enable the SSL module
# a2ensite bank32_apache_ssl // Enable the sites described in this file
```

Building and Starting the Container The container contains the Apache Web Server. Build and run the container from the `Labsetup` folder, which is where the `docker-compose.yml` file is located. To build the container, follow the instructions in the **Container Setup and Commands** section. The container must be successfully built before it can be run. Once the container is built, it can be run using the `dcrun` command. The container can be stopped by giving the `dcrun` command the `ctrl-c` key combination

Starting the Apache server. The Apache server is automatically started in the container. If you should need to stop or restart Apache, you will need to get a shell on your running container. Open a new terminal window and following the instructions in **Container Setup and Commands** to run a command on the container to start Apache.

```
// Start the server
# service apache2 start

// Stop the server
# service apache2 stop

// Restart a server
# service apache2 restart
```

Once Apache is set up properly and running in the container, we can browse the web site, and all the traffic between the browser and the server will be encrypted.

Use the above example as guidance to set up an HTTPS server for your website.

Browsing the website.

Now, point the browser to your web server (note: you should put `https` at the beginning of your URL, instead of using `http`). Describe and explain your observations. You will not be able to succeed. Why? Explain. Following the instructions in the next paragraph, fix the problem and demonstrate that you can successfully visit the HTTPS website.

In the following, we provide instructions on how to load a certificate into Firefox. You have two certificates to choose from. Your site certificate and your CA certificate. Which is the correct certificate to load into Firefox? Explain. To manually add a certificate to the Firefox browser, type the following URL in the address bar, and click the **View Certificates** button on the page (scroll to the bottom).

```
about:preferences#privacy
```

In the **Authorities** tab, you will see a list of certificates that are already accepted by Firefox. From here, we can import our own certificates. After choosing the certificate file, please select the following option: “Trust this CA to identify web sites”. You will see that our certificate is now in Firefox’s list of accepted certificates.

4.5 Task 5: Launching a Man-In-The-Middle Attack

In this task, we will show how PKI can defeat Man-In-The-Middle (MITM) attacks. Figure 1 depicts how MITM attacks work. Assume Alice wants to visit `example.com` via the HTTPS protocol. She needs to get the public key from the `example.com` server; Alice will generate a secret, and encrypt the secret using the server’s public key, and send it to the server. If an attacker can intercept the communication between Alice and the server, the attacker can replace the server’s public key with its own public key. Therefore, Alice’s secret is actually encrypted with the attacker’s public key, so the attacker will be able to read the secret. The attacker can forward the secret to the server using the server’s public key. The secret is used to encrypt the communication between Alice and server, so the attacker can decrypt the encrypted communication.

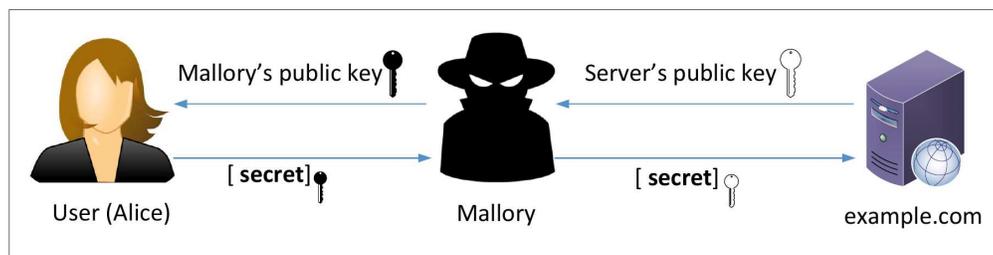


Figure 1: A Man-In-The-Middle (MITM) attack

The goal of this task is to help students understand how PKI can defeat such MITM attacks. In the task, we will emulate an MITM attack, and see how exactly PKI can defeat it. We will select a target website first. In this document, we use `www.example.com` as the target website, but in the task, to make it more meaningful, students should pick a popular website, such as a banking site or social network site.

Step 1: Setting up the malicious website. In Task 4, we have already set up an HTTPS website. We will use the same Apache server to impersonate `www.example.com` (or the site chosen by students). To achieve that, we will follow the instruction in Task 4 to add a `VirtualHost` entry to Apache's SSL configuration file: the `ServerName` should be `www.example.com`, but the rest of the configuration can be the same as that used in Task 4. Obviously, in the real world, you won't be able to get a valid certificate for `www.example.com`, so we will use the same certificate that we used for our own server.

Our goal is the following: when a user tries to visit `www.example.com`, we are going to get the user to land in our server, which hosts a fake website for `www.example.com`. If this were a social network website, The fake site can display a login page similar to the one in the target website. If users cannot tell the difference, they may type their account credentials in the fake webpage, essentially disclosing the credentials to the attacker.

Step 2: Becoming the man in the middle There are several ways to get the user's HTTPS request to land in our web server. One way is to attack the routing, so the user's HTTPS request is routed to our web server. Another way is to attack DNS, so when the victim's machine tries to find out the IP address of the target web server, it gets the IP address of our web server. In this task, we simulate the attack-DNS approach. Instead of launching an actual DNS cache poisoning attack, we simply modify the victim's machine's `/etc/hosts` file to emulate the result of a DNS cache poisoning attack by mapping the hostname `www.example.com` to our malicious web server.

```
10.9.0.80 www.example.com
```

Step 3: Browse the target website.

With everything set up, now visit the target real website, and observe how your browser reacts. Explain what you have observed.

4.6 Task 6: Launching a Man-In-The-Middle Attack with a Compromised CA

In this task, we assume that the root CA created in Task 1 is compromised by an attacker, and its private key is stolen. Therefore, the attacker can generate any arbitrary certificate using this CA's private key. In this task, we will see the consequence of such a compromise.

Design an experiment to show that the attacker can successfully launch MITM attacks on any HTTPS website. You can use the same setting created in Task 5, but this time, you need to demonstrate that the MITM attack is successful, i.e., the browser will not raise any suspicion when the victim tries to visit a website but land in the MITM attacker's fake website. Explain how you accomplished this.

5 Container Setup and Commands

Enter the `Labsetup` folder, and use the `docker-compose.yml` file to set up the lab environment. Detailed explanation of the content in this file and all the involved `Dockerfile` can be found from the user manual, [SEED Manual for Containers](#). If this is the first time you set up a SEED lab environment using containers, it is very important that you read the user manual.

In the following, we list some of the commonly used commands related to Docker and Compose. Since we are going to use these commands very frequently, we have created aliases for them in the `.bashrc` file (in our provided SEEDUbuntu 20.04 VM).

```
$ docker-compose build # Build the container images
$ docker-compose up    # Start the containers
$ docker-compose down  # Shut down the containers

// Aliases for the Compose commands above
$ dcbuild              # Alias for: docker-compose build
$ dcup                 # Alias for: docker-compose up
$ dcdown               # Alias for: docker-compose down
```

All the containers will be running in the background. To run commands on a container, we often need to get a shell on that container. We first need to use the `"docker ps"` command to find out the ID of the container, and then use `"docker exec"` to start a shell on that container. We have created aliases for them in the `.bashrc` file.

```
$ dockps              // Alias for: docker ps --format "{{.ID}} {{.Names}}"
$ docksh <id>        // Alias for: docker exec -it <id> /bin/bash

// The following example shows how to get a shell inside hostC
$ dockps
b1004832e275  hostA-10.9.0.5
0af4ea7a3e2e  hostB-10.9.0.6
9652715c8e0a  hostC-10.9.0.7

$ docksh 96
root@9652715c8e0a:/#

// Note: If a docker command requires a container ID, you do not need to
//       type the entire ID string. Typing the first few characters will
//       be sufficient, as long as they are unique among all the containers.
```

If you encounter problems when setting up the lab environment, please read the “Common Problems” section of the manual for potential solutions.