BLUE COAT

Network + Security + Cloud

CERTIFICATE PINNING

TECHNOLOGY PRIMER

INTRODUCTION

This document explains what certificate pinning is and how it is used and then considers the impact that certificate pinning has on enterprise Encrypted Traffic Management (ETM). While certificate pinning can prevent the use of some applications by an enterprise that is using ETM to allow security tools to see inside encrypted traffic flows it does not cause significant issues for most enterprises.

This document only addresses the use of certificate pinning by SSL/TLS traffic other protocols may use this technique but they are not considered here.

What is Certificate Pinning

During the handshake that takes place when an SSL/TLS connection is established the client can authenticate the server it is talking to by validating that the server certificate was issued by a Certificate Authority that the client trusts. To understand how this works we first need to explain the different types of certificates that a server can send to the client and how the server's certificate is created.

Certificates

All server certificates contain the server's public key which can be used by the client to encrypt handshake messages being sent back to the server. The server has the corresponding private key and this is never disclosed to anyone, the private key can be used to decrypt a handshake message encrypted using the public key.

The following types of certificate are relevant to understand the SSL handshake:

- **Self-signed server certificate.** This is a server certificate that has been generated by the server and that is not signed by a Certificate Authority (CA). Self-signed server certificates do not allow the client to authenticate the server as the certificate has not been issued by a trusted Certificate Authority.
- **Server certificates issued by a Certificate Authority (CA).** A server certificate issued by a trusted CA includes the server public key and details but also includes the CA details and a signature that is encrypted with the CA's private key.
- **Root CA Certificate.** A root CA certificate is issued by the CA itself and includes it's public key. The CA private key is never revealed to anyone. A root CA can sign a server certificate to indicate that the server certificate was issued by the CA to a specific server. The CA will only sign server certificates that are for a valid server. So, a trusted CA will not sign a server certificate for a server called server1.example.com if the certificate is being requested by an organization other than "example".
- **Intermediate CA Certificate.** This is a CA certificate that is issued by another CA, so the authenticity of the intermediate CA can be verified because it is signed by either another intermediate CA or by a root CA. Intermediate CAs can sign server certificates in exactly the same way a root CA can.

So, the server will always send a server certificate to the client as part of the SSL handshake and that server certificate will normally be signed by one or more CAs. It is not unusual to have the server certificate signed by an intermediate CA which itself is signed another intermediate CA that in turn is signed by the root CA. The set of CAs that are part of the server certificate represent the chain of trust that the client is relying on to authenticate the server.

In order for the client to authenticate a server certificate signed by one or more CAs it must have the CA certificates for the CAs so that it has access to the relevant CA public key in order to decrypt the signature and so verify it was created by the CA.

Server Authentication

The client will authenticate the server certificate it receives by doing the following:

- Identify the CA that signed the server certificate.
- Find that CA certificate in its trusted CA certificate store and obtain the public key.
- Use the public key to validate that the signature is genuine.
- If the CA is itself signed by a higher level CA then the client will find that CA in its trusted store and then use its public key to validate the signature on the first CA. This process my repeat until the top level CA in the chain of trust is reached – this will be the root CA.
- If the client does not have a local copy of any of the intermediate CA certificates then it will try to download them from the root CA.

If the chain of trust is authentic then the client will treat the server certificate as trusted. If the chain of trust is not authentic, perhaps because the client does not have a copy of the root CA in its trusted CA store, then the client will treat the server certificate as untrusted. Depending on how policy is set at the client an untrusted server certificate may result in warnings to the user or even in the client refusing to establish a connection to the server.

Most web browsers and operating systems include a default set of trusted public CA certificate that are automatically present in the local trusted CA store. In total there are a few hundred public Certificate Authorities whose CA certificates may be in the locally trusted CA store. In the case of Chrome and Internet Explorer these browsers use the OS trusted CA store but Firefox maintains its own trusted CA store separate from the OS store. New CA certificates may be added to the locally trusted store for a number of reasons:

- A new public CA may be created that adding to the local store.
- An existing CA certificate may be updated and so the new version needs to replace the existing version in the local CA store.
- An enterprise CA may need adding to the locally trusted store so that certificates signed by this CA can be authenticated.

Encrypted Traffic Management

In order for an ETM system to decrypt and re-encrypt traffic so that a attached security tools can detect any threats in the traffic it needs to intercept the server certificate sent by the server to the client. Once it has intercepted the server certificate it will replace the server public key with its own public key (which it has the corresponding private key to) and then it will sign the server certificate using an enterprise CA that is installed for this purpose.

As long as the client that receives the modified server certificate has the enterprise CA used to sign it installed in its trusted CA store then it will treat the server certificate as trusted. If the CA being used by the ETM to sign the modified server certificate is not in the client trusted CA store then it will be seen as untrusted and the client may issue warnings to the user or even prevent the session from happening. The fact that the client can apply different policies to trusted and untrusted server certificates is a key factor in certificate pinning.

Certificate Pinning

There are a number of different types of certificate pinning which will be discussed in a moment. However, all certificate pinning essentially relies on the following behavior:

- The client is pre-configured to know what server certificate it should expect.
- If the server certificate does not match the pre-configured server certificate then the client will prevent the session from taking place.

The different types of certificate pinning vary in what is pre-configured on the client and how it is matched with the server certificate received during the SSL/TLS handshake.

Hard Certificate Pinning

In this type of certificate pinning the client, normally an application, has the exact server certificate details pre-configured into the application and will check to see that the received server certificate matches the pre-configured certificate. If there is no match then the application will normally not function and will report an error to the user.

If this type of certificate pinning is in use then there is no way that an ETM system can provide visibility into the encrypted traffic as trying to do so will result in the server certificate being modified and so not matching the pre-configured certificate in the application. Depending on the security policy within an enterprise the ETM system may be configured to block this traffic because it cannot check it for threats or to cut it through untouched because it is a trusted application.

Hard certificate pinning is rarely used as it may result in an enterprise blocking the application which is not commercially good for the application vendor as it decreases the addressable market. There can also be complications if the server certificate ever needs to be changed as this will require that all the client software be updated to expect the new server certificate.

CA Pinning

In this type of pinning the client does not have the actual server certificate pre-configured, instead it has limited set of CA certificates that it is prepared to accept as authentication for a server certificate from the server. So, instead of trusting a server certificate if it is signed by anyone of the hundred plus publicly trusted CAs the client will only trust the server certificate if it is signed by a specific CA or by one of a small group of CAs. In fact the client can pin either on the server certificate or on the public key of the server but this detail does not alter what is described below.

This type of certificate pinning protects against the possibility that a public CA may be compromised. If this type of pinning is in use then you would expect that an ETM system could not provide visibility into the traffic but this is NOT the case. Clients that use this type of pinning explicitly allow for a server certificate that is signed by an enterprise CA as long as the enterprise CA is in the locally trusted CA store. An ETM would add the CA it is using to the client trusted CA store to prevent the client seeing warnings when accessing non pinned sites so there will be no problem with sites using CA pinning as the enterprise CA will be in the client's trusted store.

There are two ways that the set of CAs that will be trusted for a destination can be configured at the client. They may be part of the application itself, so for example in the case of Chrome it is hard coded to only trust a small set of public CAs to issue server certificates for the google.com domain. Another mechanism is the HTTP Public Key Pinning (HPKP) standard which provides a means for the server to indicate to the client what set of CAs it is going to use so that the client can test against this set. Browsers will disables pinning for certificate chains with private root certificates to enable ETM devices to operate. The RFC 7469 standard also recommends disabling pinning violation reports for such certificate chains.

BLUE COAT®

Network
+ Security
+ Cloud

Impact on ETM of certificate pinning

As noted earlier certificate pinning does not prevent an ETM device from decrypting and re-encrypting traffic if traffic is sent over the connection. Certificate pinning allows the client to decide whether or not to send traffic based on whether it trusts the server certificate. The client will typically not allow the application to function if it does not trust the server certificate. Note that clients may also apply policy on what to do with untrusted server certificates for flows where certificate pinning is not in use. The following list shows which server certificates will be trusted and which not by a client implementing certificate pinning.

- Original server certificate from the server unmodified by any ETM system. This will be trusted by clients using either type of certificate pinning.
- Server certificate modified by an ETM system and signed with an enterprise CA
  - not trusted by a client using hard certificate pinning
  - trusted by a client using CA pinning as long as the enterprise CA is present in the local trusted CA store
  - not trusted by a client using CA pinning if it does not have the enterprise CA in its local trusted CA store

As noted earlier, when an ETM system is in use by the enterprise, the CA that is being used to sign modified server certificates will be installed on the clients as a trusted CA meaning that the use of CA pinning does not prevent the ETM system from decrypting and re-encrypting the traffic, if the client sends any, in order for security tools to detect any threats that may be present.

Blue Coat Systems Inc.
www.bluecoat.com

Corporate Headquarters
Sunnyvale, CA
+1.408.220.2200

EMEA Headquarters
Hampshire, UK
+44.1252.554600

APAC Headquarters
Singapore
+65.6826.7000