

Securing Restful API Services

Lance Tyler, Matt Morris

CSC 458: Cryptography

Dr. Zachary Peterson

December 1, 2014

Contents

| | | |
|----------|---|----------|
| 1 | Introduction | 3 |
| 1.1 | Motivation | 3 |
| 1.2 | Securing Web Services | 3 |
| 2 | Design | 4 |
| 2.1 | Confidentiality | 4 |
| 2.2 | Authenticity and Integrity | 5 |
| 2.3 | Generating, and Transporting the Shared Key | 5 |
| 3 | Conclusions | 7 |

1 Introduction

1.1 Motivation

Many student developers find it difficult to turn their simple application idea into usable code. They may either not have the time it takes to develop a new code base from the ground up, or they may not have the expertise in the numerous systems required to create a new product. Our web service helps solve this problem by providing generalized functions used by social applications, such as managing user associations, managing users, and managing user data. Instead of having to setup their own database and server, application developers now simply use our public web service as a backend when creating their application.

As an extension to our senior project, we chose to implement some common practices in securing a web service for our final research project. Developers using our service will benefit by not having to design or implement any of the security measures required for a web service. Clients using the applications developed using our product will benefit by knowing that their data is secured by people with experience in cryptography.

1.2 Securing Web Services

When securing our web service there were five main attacks we chose to defend against: unauthorized access, parameter manipulation, network eavesdropping, disclosure of configuration data, and mes-

sage replay attacks. [2] The sections below summarize what these attacks are and how they apply to the web service we are protecting.

Unauthorized Access A web service should restrict access to any sensitive data that it stores. In our service these resources are the associations between two users, and data that a user owns. Only a properly authenticated user may create an association with another user, or set the visibility of their data. [2]

Parameter Manipulation In our service application developers program what parameters are passed to the web service. A malicious adversary may modify the uniform resource identifier (URI) in an attempt to gain access to restricted resources. Our URI is modeled after Restful resource naming practices:

```
/serverIP/ApplicationName/  
userID/resource
```

The resource is either associations, or data. This means that only the client who authenticates as the userID may access these resources.

Network Eavesdropping An adversary with access to the connection between the client and the server may be able to intercept and collect sensitive data. In todays world we can never be quite sure who can see our public information when using unencrypted means

of communication, so preventing passive eavesdroppers is a must.

Disclosure of Configuration Data

Many RESTful frameworks are set by default to display error information, most notably uncaught exceptions, to clients who causes the error in question. This can disclose potentially sensitive information to users of the service who shouldn't see that information. At a minimum, disclosing this information will allow adversaries to learn what specific framework your service is using, which could allow them to hone in on framework-specific vulnerabilities. In a worst case scenario, disclosing certain info could lead to something like a padding oracle attack.

Message Replay Even if requests to a RESTful service are encrypted, there is still nothing in place that stops some eavesdropper from collecting packets and simply re-sending copies of them to their destination. This would obviously be paramount in things like banking systems, but it is also important for general security for less critical applications as well.

2 Design

2.1 Confidentiality

Encryption In securing our web-service, we chose to provide confidentiality by running our service over TLS

1.2. The largest implications of this are that our service should theoretically have complete confidentiality of its communications with users of the service. The cipher suite used by the SSLContext is *tlsecdhe_rsa_with_aes128cbcsha*. Although this is not the best cipher suite due to the CBC component of it, this was unfortunately the best out of the box option available for our deployment server.

Replay Attacks Another benefit of using TLS is that our service is protected from message replay attacks. While our messages are still encrypted, nothing is stopping from an adversary from storing our request packets and re-sending them to contrive some sort of attack, or simply to cause harm to the system. Luckily, TLS inherently protects against message replay attacks by incorporating its own HMAC primitive. By including a sequence number in each packet that monotonically increases with each packet sent, the TLS protocol can effectively discard any replayed packets, since the replayed packet will have a sequence number which would be lower than the current sequence number (meaning it has been seen before). The TLS HMAC is computed from the TLS MAC secret, the sequence number, the message length, the message contents, and two fixed character strings. [6] Factoring this implication in, our service should now provide total encryption of our communications, as well as assurance that our encrypted packets cant simply be replayed for malicious purposes.

```

00 45 00 00 e3 60 21 40 00 80 06 00 00 7f 00 00 01 E...`!@. ....
10 7f 00 00 01 c9 f5 82 35 3c 47 2f bd 5a df 3d 7d .....5 <G/.Z,=}
20 50 18 00 1e b6 4d 00 00 50 4f 53 54 20 2f 63 68 P....M... POST /ch
30 65 63 6b 20 48 54 54 50 2f 31 2e 31 0d 0a 43 6f eck HTTP /1.1..Co
40 6e 74 65 6e 74 2d 4c 65 6e 67 74 68 3a 20 32 36 ntent-Le ngth: 26
50 0d 0a 43 6f 6e 74 65 6e 74 2d 54 79 70 65 3a 20 ..Conten t-Type:
60 74 65 78 74 2f 70 6c 61 69 6e 3b 20 63 68 61 72 text/pla in; char
70 73 65 74 3d 55 54 46 2d 38 0d 0a 48 6f 73 74 3a set=UTF- 8..Host:
80 20 6c 6f 63 61 6c 68 6f 73 74 3a 33 33 33 33 33 localho st:333333
90 0d 0a 43 6f 6e 6e 65 63 74 69 6f 6e 3a 20 4b 65 ..Connec tion: Ke
a0 65 70 2d 41 6c 69 76 65 0d 0a 55 73 65 72 2d 41 ep-Alive ..User-A
b0 67 65 6e 74 3a 20 52 65 73 74 43 6c 69 65 6e 74 gent: Re stClient
c0 2d 54 6f 6f 6c 0d 0a 0d 0a 43 61 6e 20 49 20 73 -Tool... .Can I s
d0 65 65 20 74 68 69 73 20 73 68 69 74 3f 3f 3f 3f ee this shit????
e0 3f 3f 3f ???

```

Figure 1: Wireshark screen shot of no encryption used

```

0000 45 00 00 fd 53 7d 40 00 80 06 00 00 7f 00 00 01 E...S}@. ....
0010 7f 00 00 01 c9 c7 82 35 36 8c 65 b2 df 04 65 0d .....5 6.e..e.
0020 50 18 00 1d 8f ec 00 00 17 03 01 00 d0 7b 7d fa P..... {}.
0030 01 04 40 81 40 8a bb 83 95 d7 bb d9 54 76 14 b6 ..@.@... .Tv..
0040 83 48 2b 0a 3c 83 69 4e 94 80 79 de 91 05 11 5d .H+.<.iN :y....]
0050 78 a5 91 f9 9d be df 5e 3a 34 f8 eb 7f 17 b3 72 x.....^ :4.....r
0060 4d d6 0a 01 f8 4b 0e 00 36 b5 32 35 a5 a7 07 c7 M....K.. 6.25....
0070 6d b0 ca f8 a7 66 b6 3d b8 12 25 b6 52 ef 2b b7 m....f.= .%.R.+
0080 42 3c 9e 2d 0e b1 c4 4c 87 7c 63 62 97 64 a1 2e B<-...L |cb.d..
0090 d5 51 d5 75 be ed c8 e7 fc b8 12 73 b9 59 d2 65 .Q.u.... .s.Y.e
00a0 5d d4 bb 01 86 e8 e8 eb 23 9f 54 bd 7b 6e c6 4d ]...... #.T.{n.M
00b0 08 17 97 92 b6 6f 38 61 ed 6e 65 54 6b 09 bf 5d .....o8a .neTk..]
00c0 4e 51 a5 71 22 bd 7f 0e 5e 05 6c cd c4 a3 7b 78 NQ.q"... ^.l...{x
00d0 5a 58 8b be e3 89 88 9c da fe 5f 1c cd e6 34 ed ZX.....4.
00e0 8a 01 46 1d fa 61 88 78 f1 42 6a 71 c8 0d a5 10 ..F..a.x .Bjq...
00f0 ec af 67 57 2f 65 f8 2e d4 ef 47 60 66 ..gw/e... .Gf

```

Figure 2: Wireshark screen shot of TLS encryption

Verification of TLS To ensure that TLS was turned on for our server we used Wireshark to sniff packets that were being sent to it. See figure 1 and 2. Al-

though "scrambled" data is not a good indication of if something is encrypted, we trust that the creators of the SSLContext class properly implemented TLS.

2.2 Authenticity and Integrity

The method we chose to prevent unauthorized access and to prevent parameter manipulation was to use the hashed based method authentication code (HMAC). An HMAC is a combination of a MAC and a hashing algorithm. According to authors of RFC2104 the cryptographic strength of the hash, the size and quality of the key, and the size of the output, all contribute to the strength of the

HMAC. [1] The HMAC can also provide authenticity of the origin of the resource request. The sections below illustrate the problems we encountered when designing a system that uses an HMAC, and the solutions to those problems.

2.3 Generating, and Transporting the Shared Key

Registration Before the client uses the service, they must first register over the TLS channel with a user ID, and a pass-

word. It is best practices to hash and salt a password before storing it in a database. For our implementation we used the SHA-256 hashing algorithm with a 128 bit salt generated from the CSPRNG found in `java.security.SecureRandom`. In the event that our database were to be compromised, user passwords would still be secured. This is due to the fact that a hash function is one way, meaning the original input of the function can never be derived from the output. Salting the user passwords protects against techniques such as rainbow table, or lookup tables. This is when an adversary pre-computes the hashes to passwords ahead of time, allowing them to simply do a reverse lookup of a hashed password to find the original inputted password. [5]

Authentication To get the shared key a user must first authenticate with their user ID password pair over TLS. To prevent against a brute force attack against a users password we have enforced a rate limit of 5 authentication attempts per user ID every 15 minutes. The goal of the rate limited authentication is to prevent any brute force guesses against the users password. This allows an adversary only $4(60min/15minslockedout) * 5(attempts) = 20attemptsperhour$. A password minimum length of 8 characters is also enforced. The result is the adversary does not have an efficient way to guess the entire $3.026 * 10^{15}$ passwords. [4]

Generating Shared Key As part of the requirement for generating an HMAC, a shared key must be established between the server and the client. During the authentication process the server generates a shared key, which it returns to the user, and stores in its database. These keys are generated using the `java.crypto.SecureRandom`, a FIPS compliant random number generator.[3] This key is ephemeral meaning it is only recognized as valid by the server for only 1 hour. The key the client receives may be incorporated into a session cookie by the application developer.

Generating the HMAC

$$HMAC(K, m) = H((K \oplus opad) | H((K \oplus ipad) | m))$$

$H = HashFunction$

$K = Key$

$m = message$

$opad = outerpadding(x5cblock)$

$ipad = innerpadding(x36block)$

[1]

An HMAC requires that the underlying hash function used to generate it does not have any targeted collisions. For generating the HMAC we chose the SHA-256 hash function which doesn't have any documented collisions. [1] To ensure a properly formatted HMAC we chose to use Java code from an AWS HMAC-SHA256 function. For the client we chose the `crypto-js` implementation of the HMAC-SHA256 found at: <https://code.google.com/p/crypto-js/>. Using the shared key acquired from

authenticating with the server, the client signs each requests (HTTP verb + URI of the resource + payload).

Timing Attack Prevention Upon receiving a resource request, the server will recreate the HMAC using the same information the client used, and compare the two HMACs. To prevent any timing attacks against the server we implemented our own .equals method for comparing the two HMACs that runs at constant time.

Restricting Resource Access In addition to ensuring that the client has not tampered with the parameters of the resource request, the HMAC also allows the server to verify the identity of a client requesting access to a resource. The server uses the following logic to determine if the request is allowed to access the resource:

```
userId -> from URI
HMAC -> found in HTTP header
SharedSecret ->
    dbLookup(userID)

if SharedSecret is null:
```

```
    unauthorized()
ServerHMAC ->
    computeHMAC(URI, SharedSecret)
if HMAC = ServerHMAC:
    allowAccess()
else:
    unauthorized()
```

If the server does not find a valid shared secret for that userID, it means the client requesting access to the resource is not authenticated. If the HMAC computed by the server does not match the HMAC sent by the client, then the client does not have the correct shared secret, and is not authorized access to the resource.

3 Conclusions

For this project we learned the importance of securing a web service. The big take away was that SLL/TLS is not that difficult to implement in a server, and should be used for protecting the users of a web service. We also learned how to authenticate users, as well as other security methods for managing user credentials.

References

- [1] R. C. H. Krawczyk, M. Bellare, “Hmac: Keyed-hashing for message authentication,” February 1997. [Online]. Available: <http://www.ietf.org/rfc/rfc2104.txt>

RFC 2104, HMAC, description

- [2] Microsoft, “Building secure web services,” June 2003. [Online]. Available: <http://msdn.microsoft.com/en-us/library/ff648643.aspx>

guide to building secure web services

- [3] Oracle, “Class securerandom.” [Online]. Available: <https://docs.oracle.com/javase/7/docs/api/java/security/SecureRandom.html>

API document on Java’s CPRNG

- [4] T. Pornin, “How many possible combinations in 8 character password.” [Online]. Available: <http://math.stackexchange.com/questions/739874/how-many-possible-combinations-in-8-character-password>

Stackexchange answer on number of passwords for 8 chars

- [5] D. Security, “Salted password hashing - doing it right.” [Online]. Available: <https://crackstation.net/hashing-security.htm>

How to salt and hash passwords correctly

- [6] E. R. T. Dierks, “The transport layer security (tls) protocol,” April 2006. [Online]. Available: <http://tools.ietf.org/html/rfc4346#appendix-F.2>

RFC 4346, TLS, description