

# Advanced Synapse setup with Let's Encrypt

[Originalartikel](#)

[Backup](#)

```
<html> <div alt=„article“>
```

    <p>So, you've installed and configured Synapse and started chatting from your very own Matrix home server? What's the next step? Well, right now you're probably accessing your new home server over plaintext HTTP, which is bad, particularly because you'll be sending your password over this connection when you log in. You could connect to Synapse's secure HTTP port, but your browser won't trust it by default because you'd normally need to pay for a certificate that your browser would recognise. That is, until recently!</p>

<p><a href="https://letsencrypt.org/" target="„\_blank"›Let's Encrypt</a› is a new initiative that issues SSL certificates free of charge, in an effort to make SSL universal on the Internet. In this blog post, we'll be walking through an example of how we can use this service to get ourselves a securely hosted Synapse.</p> <p>We're going to assume you have a Synapse installed and listening on the standard ports, 8008 and 8448. If not, follow the <a href="https://github.com/matrix-org/synapse/blob/master/README.rst" target="„\_blank"›Synapse README</a> and come back here when you're done. Everybody ready? Okay, good.</p> <p>So, in order to get a certificate from Let's Encrypt, we need to prove that we own our domain. The simplest way to do this is to host some special files on our web server. Now, Synapse won't do this. We could use a separate web server, but then we'd have to stop Synapse and start the other web server every time we renewed our certificate, and that means downtime. Instead, let's put our Synapse behind a proper web server and let that serve the files. This has added advantages, including that we can host our Matrix home server on the standard port 443 without having to run Synapse as root.</p> <p>For this example, we're going to use NGINX, so start by installing NGINX in whatever way your Linux distribution of choice recommends.</p> <p>Now, you should have a webroot for your new web server somewhere. Hopefully your helpful Linux distribution has started you off with a config file; let's see:</p> <pre># nano /etc/nginx/nginx.conf<br> </pre> <p>We're looking for the &#216;server&#217; section of that file. We need to make it look something like this:</p> <pre>server {<br>    # Make sure this is 0.0.0.0: no use listening on 127.0.0.1 or we'll only be<br>    # serving to ourselves! There's no port here, which means we'll listen on<br>    # port 80<br>    listen 0.0.0.0;<br>}<br>    server\_name example.com www.example.com;<br>}</pre>

```
    # Make sure this is 0.0.0.0: no use listening on 127.0.0.1 or we'll only be<br>    # serving to ourselves! There's no port here, which means we'll listen on<br>    # port 80<br>    listen 0.0.0.0;<br>}<br>    server_name example.com www.example.com;<br>}
```

&#13;

```
    server_name example.com www.example.com;<br>}
```

&#13;

```
access_log /var/log/nginx/example.com.access_log main;#13;
error_log /var/log/nginx/example.com info;#13;
```

&#13;

```
# This is where we put the files we want on our site#13;
root /var/www/examplecom/htdocs;#13;
```

&#13;

```
# Here's where it gets interesting: This will send any path that
starts#13;
# with /_matrix to our Synapse!#13;
```

&#13;

```
location /_matrix {#13;
    proxy_pass http://localhost:8008;#13;
}#13;
}#13;
```

</pre> <p>When you're happy with the look of that file, let's restart the server:</p> <pre># nginx -s reload#13; </pre> <p class=„p1“>Before we go any further, let's test our new configuration:</p> <pre>\$ curl [#13; {„old\\_verify\\_keys“:{},„server\\_name“:„example.com“,„signatures“:{„example.com“:{„ed25519:auto“:„RWb+w6vHUUokoDgElwG6Cg50ezZvBrzXtJmJIH8jEwl5x0JQ7prn3FwjhbKTH5jE7J8lly3HEc4COn4JCCvCA“}},„tls\\_fingerprints“:\[{„sha256“:„DMbzSZ5Uj7/6p/RT/UtQYJLHm5o0TwBSVYXsqpDdVDs“}\],„valid\\_until\\_ts“:1455203001035,„verify\\_keys“:{„ed25519:auto“:{„key“:„1YiTDjmE86AlmrblYE2lyqauV9wPo8jw2kxZAZFfl/Q“}}}}#13; </pre> <p class=„p1“>Those are your server's public keys! Now we have a web server running, we can get our SSL certificate. Let's Encrypt have their own client which will automate everything including rewriting your NGINX config file, however that means it has a large number of dependencies and needs to be run as root. For this example, we're going to use the much simpler <a href=„<https://github.com/diafygi/acme-tiny>“ target=„\\_blank“>acme\\_tiny.py</a>.&#160;I'm going to assume you have a user called, &#2016;letsencrypt&#212;, so, as root, let's set up the place for it to write its challenge files:</p> <pre># mkdir /var/www/examplecom/htdocs/.well-known/acme-challenge#13; # chown letsencrypt:users /var/www/examplecom/htdocs/.well-known/acme-challenge#13; </pre> <p class=„p1“>Now let's switch to our letsencrypt user:</p> <pre>\\$ ssh letsencrypt@example.com#13; </pre> <p class=„p1“>We'll start by getting ourselves a copy of acme\\_tiny.py:</p> <pre>\\$ git clone \[#13; </pre> <p class=„p1“>Now let's set up a directory structure \\(let's say we might want to manage more than one domain someday\\):</p> <pre>\\\$ mkdir examplecom#13; \\\$ cd examplecom#13; \\\$ ln -s /var/www/examplecom/htdocs/.well-known/acme-challenge challenges#13; </pre> <p class=„p1“>Now, we'll need to generate two keys for Let's Encrypt, and account key and a domain key. The former is what we use to identify ourselves to Let's Encrypt and the latter is the key we use to do the actual SSL.</p> <pre>\\\$\]\(https://github.com/diafygi/acme-tiny.git\)](http://example.com/_matrix/key/v2/server/auto)

```
openssl genrsa 4096 > letsencrypt_examplecom_account.key; Generating RSA private key,
4096 bit long modulus; ..++;
```

---

```
.....++; e is 65537 (0x10001);  
$ chmod 600 letsencrypt_examplecom_account.key; $ openssl genrsa 4096 > letsencrypt_examplecom_domain.key; Generating RSA private key, 4096 bit long modulus; ..++;
```

---

```
.....++; e is 65537 (0x10001); $ chmod 600 letsencrypt_examplecom_domain.key; </pre> <p class=„p1“>Now, store those keys somewhere safe! After you've done that, let's generate a certificate request for our domain. Note that we're requesting one for both example.com and www.example.com: this isn't strictly necessary for Matrix but could be useful if we want to host a website too.</p> <pre>$ openssl req -new -sha256 -key letsencrypt_examplecom_domain.key -subj ,/ -reqexts SAN -config &lt;(cat /etc/ssl/openssl.cnf &lt;(printf „[SAN]„nsubjectAltName=DNS:example.com,DNS:www.example.com)) &gt; examplecom.csr; </pre> <p class=„p1“>Okay, we have our keys, our certificate request, and somewhere to host our challenge files, so we're ready to request a certificate! Be careful about this part and make sure you've got everything right, because Let's Encrypt enforces strict rate limits on the number of certificates you can request for one domain. Here we go:</p> <pre>$ python ~/acme-tiny/acme_tiny.py -account letsencrypt_examplecom_account.key -csr examplecom.csr -acme-dir challenges; </pre> examplecom.crt; Parsing account key...; Parsing CSR...; Registering account...; Registered!; Verifying example.com...; example.com verified!; Verifying www.example.com...; www.example.com verified!; Signing certificate...; Certificate signed!; </pre> <p class=„p1“>Is that it, did it work? Well, let's see:</p> <pre>$ openssl x509 -in examplecom.crt -noout -text; Certificate:;
```

Data:  
Version: 3 (0x2);  
Serial Number:  
01:02:22:77:02:1b:eb:d5:3d:c3:14:6d:87:43:22:3d:fc:0f;  
Signature Algorithm: sha256WithRSAEncryption;  
Issuer: C=US, O=Let's Encrypt, CN=Let's Encrypt Authority X3;  
Validity;  
Not Before: Feb 6 21:37:00 2016 GMT;  
Not After : May 6 21:37:00 2016 GMT;  
Subject: CN=example.com;  
Subject Public Key Info:

```
[etc]</pre> <p class=„p1“>Congratulations, you have an official, signed certificate for your domain! Now, before we can use it, we need to add the Let's Encrypt certificate to it, because our web server needs to send both:</p> <pre>$ wget  
https://letsencrypt.org/certs/lets-encrypt-x3-cross-signed.pem -2016-02-06 23:38:55-  
https://letsencrypt.org/certs/lets-encrypt-x3-cross-signed.pem; Resolving letsencrypt.org...  
23.66.17.98, 2a02:26f0:60:489::2a1f, 2a02:26f0:60:481::2a1f&#13; Connecting to  
letsencrypt.org|23.66.17.98|:443... connected.&#13; HTTP request sent, awaiting response... 200  
OK&#13; Length: 1675 (1.6K) [application/x-x509-ca-cert]&#13; Saving to: &#8216;lets-encrypt-x3-  
cross-signed.pem&#8217;&#13; lets-encrypt-x3-cross-signed.pe  
100%[======>] 1.64K --KB/s in 0s&#13; 2016-02-06 23:38:55 (61.5 MB/s) - &#8216;lets-encrypt-x3-cross-
```

signed.pem"; saved [1675/1675]; \$ cat examplecom/examplecom.crt lets-encrypt-x3-cross-signed.pem >examplecom/examplecom\_cert\_chain.crt; </pre> <p class="p1">Now we let's symlink it in place, along with the domain key, so we can renew it easily later. We'll need to be root again for this:</p> <pre>\$ ssh root@example.com; # ln -s /home/letsencrypt/examplecom/examplecom\_cert\_chain.crt /etc/ssl/nginx/examplecom\_cert.pem; # ln -s /home/letsencrypt/examplecom/letsencrypt\_examplecom\_domain.key /etc/ssl/nginx/examplecom\_key.pem; </pre> <p class="p1">Now, one more crucial thing we have to do before using our SSL is to give NGINX some Diffie Hellman parameters. This is a good thing to do for any SSL configuration (it will increase your score on <a href="https://www.ssllabs.com/ssltest/" target="\_blank">SSL Labs</a>) but it's absolutely crucial for us because Synapse will only negotiate forward secret connections, so otherwise other Matrix home servers will refuse to talk to us! (Technically, Synapse also supports elliptic curve Diffie Hellman, which doesn't need DH parameters, but not all Synapses will support this.) You'll already have some Diffie Hellman parameters from your existing Synapse, so you could use them:</p> <pre># cp /home/synapse/synapse/matrix.example.com.tls.dh /etc/ssl/nginx/examplecom\_dhparams.pem; </pre> <p class="p1"># or you can generate your own. You'll probably want to do this on your desktop or laptop if you have OpenSSL installed, it will be much faster:</p> <pre>\$ openssl dhparam -out examplecom\_dhparams.pem 2048; Generating DH parameters, 2048 bit long safe prime, generator 2; This is going to take a long time; .....+.....[etc, etc]; \$ scp examplecom\_dhparams.pem root@example.com:/etc/ssl/nginx/examplecom\_dhparams.pem; </pre> <p class="p1">Now, let's get our new certificate in action! Open up your NGINX config file again, and add another server block that looks like this:</p> <pre> server {</pre>

```
listen 0.0.0.0:443;
server_name example.com www.example.com;
ssl on;
ssl_certificate /etc/ssl/nginx/examplecom_crt.pem;
ssl_certificate_key /etc/ssl/nginx/examplecom_key.pem;
ssl_dhparam /etc/ssl/nginx/examplecom_dhparams.pem;
ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
ssl_prefer_server_ciphers on;
# mozilla intermediate list, jan 2016;
ssl_ciphers "ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES256-GCM-SHA384:DHE-RSA-AES128-GCM-SHA256:DHE-DSS-AES128-GCM-SHA256:kEDH+AESGCM:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA256:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-AES256-SHA384:ECDHE-ECDSA-AES256-SHA384:ECDHE-RSA-AES256-SHA:ECDHE-ECDSA-AES256-SHA:DHE-RSA-AES128-SHA256:DHE-RSA-AES128-SHA:DHE-DSS-AES128-SHA:DHE-RSA-AES128-SHA256:DHE-RSA-AES256-SHA256:DHE-DSS-AES256-SHA:DHE-RSA-AES256-SHA:ECDHE-RSA-DES-CBC3-SHA:ECDHE-ECDSA-DES-CBC3-SHA:EDH-RSA-DES-CBC3-SHA:AES128-GCM-SHA256:AES256-GCM-SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:AES:CAMELLIA:DES-CBC3-SHA:!aNULL:!eNULL:!EXPORT:!DES:!RC4:!MD5:!PSK:!aECDH:!EDH-DSS-DES-CBC3-SHA:!KRB5-DES-CBC3-SHA";#;
ssl_session_cache shared:SSL:50m;
access_log /var/log/nginx/examplecom.ssl_access_log main;#;
```

```

    error_log /var/log/nginx/examplecom.ssl_error_log info;
    root /var/www/examplecom/htdocs;
    location /_matrix {
        proxy_pass http://localhost:8008;
    }
}

```

</pre> <p class=„p2“>It looks pretty similar to our previous server block, except for all that stuff about SSL in the middle. We're pointing NGINX at our certificate, key and Diffie Hellman parameter files and specifying what protocols and ciphers we want our server to talk. The long list here is taken from <a href=„[https://wiki.mozilla.org/Security/Server\\_Side\\_TLS](https://wiki.mozilla.org/Security/Server_Side_TLS)“ target=„\_blank“>Mozilla's Server Side TLS guidelines</a> and is their Intermediate list. See that page for more information on what that means, and choose a different list of ciphers if you prefer: just remember we must support at least the ephemeral Diffie Hellman ciphers, or other home servers won't talk to us!</p> <p class=„p2“>Now let's restart our NGINX and see if it works:</p> <pre># nginx -s reload</pre> <p class=„p2“>&#8230;and that command again, this time with https:</p> <pre>\$ curl [https://example.com/\\_matrix/key/v2/server/auto](https://example.com/_matrix/key/v2/server/auto)</pre> {„old\_verify\_keys“:{},„server\_name“:„example.com“„signatures“:{„example.com“:{„ed25519:auto“:„RWb+w6vHUUokoDgElwG6Cg50ezZvBrzXtJmJIH8jEwl5x0JQ7prn3FwjhgKTH5jE7J8Ily3HEc4COn4JCCvCA“}},„tls\_fingerprints“:[{„sha256“:„DMbzSZ5Uj7/6p/RT/UtQYJLHm5o0TwBSVYXsqpDdVDs“}],„valid\_until\_ts“:1455203001035,„verify\_keys“:{„ed25519:auto“:{„key“:„1YiTJjmE86AlmrblYE2lyqauV9wPo8jw2kxZAZFfl/Q“}}}</pre> <p class=„p1“>Hooray! You should now be able to open a browser to <https://example.com/matrix> and log in securely over SSL!</p> <h2 class=„p1“>Renewing Your Certificate</h2> <p class=„p1“>Now, there's one important step left, and that's to set up renewal for the certificate, otherwise we'll find our shiny new SSL will stop working in three months time. We can use the same acme\_tiny command to do this:</p> <pre>\$ python ~/acme-tiny/acme\_tiny.py -account letsencrypt\_examplecom\_account.key -csr examplecom.csr -acme-dir challenges/ &gt; examplecom.crt</pre> Parsing account key...</pre> Parsing CSR...</pre> Registering account...</pre> Already registered!</pre> Verifying example.com...</pre> example.com verified!</pre> Verifying [www.example.com](http://www.example.com)...</pre> [www.example.com](http://www.example.com) verified!</pre> Signing certificate...</pre> Certificate signed!</pre> \$ wget <https://letsencrypt.org/certs/lets-encrypt-x3-cross-signed.pem></pre> -2016-02-06 23:38:55- <https://letsencrypt.org/certs/lets-encrypt-x3-cross-signed.pem></pre> Resolving letsencrypt.org... 23.66.17.98, 2a02:26f0:60:489::2a1f, 2a02:26f0:60:481::2a1f</pre> Connecting to letsencrypt.org[23.66.17.98]:443... connected.</pre> HTTP request sent, awaiting response... 200 OK</pre> Length: 1675 (1.6K) [application/x-x509-ca-cert]</pre> Saving to: &#8216;lets-encrypt-x3-cross-signed.pem</pre> lets-encrypt-x3-cross-signed.pe 100%[=====]</pre> 1.64K -.KB/s in 0s</pre> 2016-02-06 23:38:55 (61.5 MB/s) - &#8216;lets-encrypt-x3-cross-signed.pem</pre> saved [1675/1675]</pre> \$ cat examplecom/examplecom.crt lets-encrypt-x3-cross-signed.pem &gt; examplecom/examplecom\_cert\_chain.crt</pre> <p class=„p1“>Synapse will automatically pick up the new certificate, but we'll need to tell NGINX to reload:</p> <pre># nginx -s reload</pre> <p class=„p1“>Setting up a cronjob to automate this is left as an exercise to the reader!</p> <h2 class=„p1“>Federation behind the HTTP Proxy</h2> <p class=„p1“>If you like, you can stop reading now: our clients can access our home server securely but other home server are still talking to our Synapse directly on port 8448. This is fine, and if you're happy with this, you can stop reading now. But remember how we made sure other Synapses could talk to our NGINX? Well, why not put federation behind our new web server too?</p> <p class=„p1“>Now, we need to do a couple of things to make this work: were you looking carefully at the JSON those curl commands returned? If you were, you might have noticed a key

called, `&#8216;tls_fingerprints&#8217;`. Our home server serves up a fingerprint of the TLS certificate its using from this API, and we`&#8217;`ve just given our web server a different certificate, so we need to give Synapse our new certificate.`</p> <p class=„p1“>`How are we going to tell other home servers to talk to our NGINX instead? Well, ultimately we`&#8217;`re going to change our DNS SRV record to point at port 443 instead of port 8448, but that change could take a while to propagate through caches, so let`&#8217;`s test it by having our NGINX listen on port 8448 temporarily. We can do this by copying that same block from above, but with a different port:`</p> <pre>`server {`&#13;`

```
listen 0.0.0.0:8448;&#13;
server_name example.com www.example.com;&#13;
[etc]&#13;
```

`</pre> <p class=„p2“>`Don`&#8217;t` restart NGINX just yet: we need to tell our Synapse to stop listening on that port first, so lets do that and give it our new certificate:`<pre>$ nano` /home/synapse/synapse/homeserver.yaml`&#13;` `<p class=„p2“>`Now we`&#8217;ll` want to find and edit the following lines:`<pre>`tls\_certificate\_path:  
`„/etc/ssl/nginx/examplecom_crt.pem“&#13; # We can comment this out, as long as we set no_tls to true below&#13; # tls_private_key_path: „/whatever/path/synapse/generated“&#13; # PEM dh parameters for ephemeral keys&#13; tls_dh_params_path:  
„/whatever/path/synapse/generated“&#13; # Turn off TLS everywhere (this overrides the listeners section below)&#13; no_tls: True&#13;`

1. port: 8008`&#13;`

tls: false`&#13;`

```
# We can bind to only localhost since only the local nginx needs to hit
this&#13;
bind_address: '127.0.0.1'&#13;
type: http&#13;
# Set this so that Synapse obeys nginx's X-Forwarded-For headers, then IP
addresses will be correct in Synapse's logs&#13;
x_forwarded: true&#13;
resources:&#13;
- names: [client, webclient]&#13;
compress: true&#13;
- names: [federation]&#13;
compress: false&#13;
```

`</pre> <p class=„p2“>`Note: if you have an old enough config file that you have `&#8216;bind_host&#8217;` and `&#8216;bind_port&#8217;` directives, now is the time to remove them.`</p> <p class=„p2“>`Now let`&#8217;s` restart Synapse and our web server to swap over what`&#8217;s` listening on our port 8448:`</p> <pre>$ synctl restart&#13; # nginx -s reload&#13;

</pre> <p class=„p2“>Now let&#8217;s try that test again on port 8448:</p> <pre>$ curl https://example.com:8448/\_matrix/key/v2/server/auto&#13;  
{„old_verify_keys“:{},„server_name“:„example.com“,„signatures“:{„example.com“:{„ed25519:auto“&#13;:„bdca31805e4209f6ff4d644251a29d0cb1dc828a4d6131c57cf8305288f337c0“}},„tls_fingerprints“:[{„sha256“:„1d9ec66599e229654a79f28e26675fdeb585027553af6d581926e821a6b6527c“}],„valid_u  
ntil_ts“:1455203001035,„verify_keys“:{„ed25519:auto“:{„key“:„1YiTDjmE86AlmrblYE2lyqauV9wPo8j`

w2kxZAZFfl/Q"}>}"}; </pre> <p class="p1">Notice anything different? The `tls_fingerprints` part has changed because we now have a different certificate. The `signatures/example.com/ed25519:auto` value has changed too: that's because that part is a signature of the rest of JSON object, so changing the `tls_fingerprints` has caused this to change too.</p> <p class="p1">And that's it! If you're happy everything is working, you can then change your DNS SRV record to point at port 443 instead of 8448, then after leaving a few days for the change to propagate through caches, remove the extra server block from your `nginx.conf` and restart to stop your `nginx` listening on port 8448.</p>

```
</div><div alt="article">
<p>Historically we've had two projects for iOS:</p>
```

<ul><li><strong>MatrixSDK</strong>: a low level library to interact with a Matrix homeserver</li>
<li><strong>Console</strong>: an example Matrix client based on MatrixSDK</li> </ul><p>The primary intention of Console was to demonstrate how to use MatrixSDK to write a Matrix client app.<br/>However, this split isn't helpful for developers who want higher level modules that provides `UIViewControllerAnimated`s ready to use in an existing app, with no need to manage low level communications with the Matrix homeserver.</p> <p>It is where the <strong>MatrixKit</strong> project started. MatrixKit sits between MatrixSDK and your existing iOS app.</p> <p>It provides customisable `UIViewControllerAnimated`s a developer can integrate in their app.&#160;&#160;If you want to add to your app a screen to chat in a room, you just need to use the `MXKRoomViewController`.</p>
<p>We made MatrixKit so that the components it provides are <strong>easy to integrate</strong> but also <strong>easy to customise</strong>. We do not have yet full samples of customisation as we've been focused on the library core, but here are a few examples:<br/></p> <div> &#160;&#160;</div> <p> You probably recognise the theme of the first one, as it's what we use in the Console app today.<br/>The second one is the iOS7-style look and feel from <a href="https://github.com/jessesquires/JSQMessagesViewController">JSQMessagesViewController</a> . With few lines of code we connected it to MatrixKit data models. Yes, data models provided by MatrixKit are reusable too.</p> <p>MatrixKit is also <strong>highly extensible</strong>. If you want to create new table cells to render messages, new views, new view controllers, etc, you will find a place to hook them into the MatrixKit code.</p> <p>MatrixKit repository is here:&#160;<a href="https://github.com/matrix-org/matrix-ios-kit">https://github.com/matrix-org/matrix-ios-kit</a>&#160;and it is available via CocoaPods (the MatrixKit pod).</p> <p>In parallel of MatrixKit, we did some spring-cleaning &#8211; the official Matrix.org iOS offerings are now split into three github repos. One for each deliverable:</p> <ul><li><a href="https://github.com/matrix-org/matrix-ios-sdk">https://github.com/matrix-org/matrix-ios-

sdk

</a></li> <li><a href=„<https://github.com/matrix-org/matrix-ios-kit>“><https://github.com/matrix-org/matrix-ios-kit>

</a></li> <li><a href=„<https://github.com/matrix-org/matrix-ios-console>“><https://github.com/matrix-org/matrix-ios-console>

</a></li> </ul><p><strong>Other releases:</strong></p> <p>Today, we released&#160;MatrixSDK 0.4.0 (<a href=„<https://github.com/matrix-org/matrix-ios-sdk/blob/master/CHANGES.rst#changes-in-matrix-ios-sdk-in-040-2015-04-23>“>changes</a>). Update your pods :)</p> <p>Console 0.4.0 (<a href=„<https://github.com/matrix-org/matrix-ios-console/blob/master/CHANGES.rst#changes-in-console-in-040-2015-04-23>“>changes</a>) is&#160;in the Apple submission process. This will be the first version of the app using MatrixKit. Aesthetically, there is no change since the previous version. The app is more stable due to all the data abstractions and management improvements provided by MatrixKit.</p> <p>If you&#8217;re an iOS developer, please have a go with MatrixKit and let us know on <a href=„<https://matrix.org/beta/#/room/#ios:matrix.org>“>#ios:matrix.org</a> how you get on!</p>

```
</div><div alt="article">
<p>Hi all,</p>
```

<p>What with the chaos of Mobile World Congress last week we seem to have been hoarding releases &#8211; so here&#8217;s what&#8217;s been happening!</p> <p><a href=„<https://github.com/matrix-org/synapse>“>Synapse 0.8.0</a> was released this afternoon. This is a major performance/stability release, with lots of good stuff going on &#8211; especially adding more customisable push notification support APIs for iOS/Android; support for registering accounts from mobile devices; extensions to the new Application Service API and lots of federation improvements and bug fixes. <b>Please upgrade at your earliest convenience.</b></p>

<p>Meanwhile, we quietly released the <a href=„<https://play.google.com/store/apps/details?id=org.matrix.androidsdk.alpha>“>Matrix Console</a> Android example app to the Google Play last week, alongside <a href=„<https://github.com/matrix-org/matrix-android-sdk>“>v0.2.2 of the Android SDK</a> &#8211; release notes below. There&#8217;ll be a new version of the Android Console app out tomorrow, but mentioning here for completeness and to share the release notes. Also, the <a href=„<https://github.com/matrix-org/matrix-ios-sdk>“>iOS SDK is now on v0.3.1</a>, with lots of performance and usability improvements.</p> <p>Finally, we have a whole new official <a href=„<https://github.com/matrix-org/matrix-js-sdk>“>Matrix client SDK for JavaScript</a>, all packaged up ready for use by Node developers and JS purists alike as an NPM with minimal dependencies. Meanwhile, the <a href=„<https://github.com/matrix-org/matrix-angular-sdk>“>matrix-angular-sdk</a> has been switched to use matrix-js-sdk from now on. You can use the plain JS API with a

```
npm install matrix-js-sdk
```

and then:</p> <pre>&#13; var sdk = require(„matrix-js-sdk“);&#13; var client = sdk.createClient(„<https://matrix.org>“);&#13; client.publicRooms(function(err, data) {&#13;

```
    console.log("Public Rooms: %s", JSON.stringify(data));&#13;
});&#13; &#13; </pre> <p> Meanwhile, release notes for all &amp; sundry lie below. </p>
<pre>&#13; &#13; Changes in synapse v0.8.0 (2015-03-06)&#13;
```

===== \* Add support for registration fallback. This is a page hosted on the server \* which allows a user to register for an account, regardless of what client; they are using (e.g. mobile devices);

\* Added new default push rules and made them configurable by clients: \* Suppress all notice messages; Notify when invited to a new room; Notify for messages that don't match any rule; Notify on incoming call;

\* Federation: \* Added per host server side rate-limiting of incoming federation requests. \* Added a ``/get\_missing\_events`` API to federation to reduce number of requests;

``/events`` requests;

\* Configuration: \* Added configuration option to disable registration: \* Added configuration option to change soft limit of number of open file descriptors: ``soft\_file\_limit``;

\* Make ``tls\_private\_key\_path`` optional when running with ``no\_tls``. \* Application services: \* Application services can now poll on the CS API ``/events`` for their events;

by providing their application service ``access\_token``;

\* Added exclusive namespace support to application services API. \* Changes in Matrix Android SDK in 0.2.2 (2015-02-27);

===== \* SDK; --;

;

----- \* Matrix Console; ----- \* Improvements: \* Exif management : the uploaded image is rotated according to the exif metadata;

(if the device has enough free memory);

\* Add a piechart while downloading an image; \* Add JSON representation of a message (tap on its row, &#8220;Message details&#8221;); \* The public rooms list is now sorted according to the number of members; \* Features: \* Add configuration and skeleton classes for receiving GCM messages; \* Add REST client for pushers API with add method for HTTP pushers; \* Add the account creation; \* Bug fixes: \* Reset the image

thumbnail when a row is reused. \* SYAND-30 Notification should be away when entering a room. \* Some images thumbnails were downloaded several times. \* Restore the foreground service. \* The medias cache was not cleared after logging out. \* The client crashed when joining #anime:matrix.org. \* SYAND-29 Messages in delivery status are not seen. \* Some user display names were their matrix IDs. \* The room name/ topic were invalid when inviting to a room. </pre> <pre> Changes in Matrix iOS SDK in 0.3.1 (2015-03-03); ======&#13; &#13; --&#13; SDK&#13; --&#13; Improvements:&#13; \* Improved push notifications documentation.&#13; \* MXSession: Slightly randomise reconnection times by up to 3s to prevent all&#13;

Matrix clients from retrying requests to the homeserver at the same time.&#13;

\* Improved logs. \* Bug fixes: \* SYIOS-90 - iOS can receive & display messages multiple times when on bad connections. -----&#13; Matrix Console. -----&#13; Improvements: \* Fixed warnings with 64bits builds. \* Room history: Improve scrolling handling when keyboard appears. \* Contacts: Prompt user when local contacts tab is selected if constact sync is disabled. \* Bug fixes: \* Fix crash when switching rooms while the event stream is resuming. \* SYIOS-69 - On Screen Keyboard can end up hiding the most recent messages in a room. \* SYIOS-98 - Crash when attempting to attach image on iPad. \* Changes in Matrix iOS SDK in 0.3.0 (2015-02-23); ======&#13; &#13; --&#13; SDK&#13; --&#13; Breaks:&#13; \* [MXSession initWithMatrixRestClient: andStore: ] and the onStoreDataReady argument in&#13;

[MXSession start:] has been removed. The SDK client can now use the asynchronous&#13; [MXSession setStore:] method to define a store and getting notified when the SDK can&#13; read cached data from it. (SYIOS-62)&#13;

\* MXStore implementations must now implement [MXStore openWithCredentials]. \* All MXRestClient methods now return MXHTTPOperation objects. \* Improvements: \* Created the MXSession.notificationCenter component: it indicates when an event must be&#13;

notified to the user according to user's push rules settings.&#13;

\* MXFileStore: Improved loading performance by 8x. \* Added an option (MXSession.loadPresenceBeforeCompletingSessionStart) to refresh presence&#13;

data in background when starting a session.&#13;

\* Created MXLogger to redirect NSLog to file and to log crashes or uncaught exception. \* MXRestClient: Added [MXRestClient registerFallback]. \* Logs: Make all NSLog calls follows the same format. \* Features: \* SYIOS-40 - Any HTTP request can fail due to rate-limiting on the server, and need to be retried. \* SYIOS-81 - Ability to send messages in the background. \* Bug fixes: \* SYIOS-67 - We should synthesise identicons for users with no avatar. \* MXSession: Fixed crash when closing the MXSession before the end of initial

Sync.&#13; </pre>

```
</div><div alt="article">
<p><strong>This post has now been edited into a guide
&#8211; you can find <a
href="https://github.com/matrix-org/matrix-doc/blob/master/supporting-docs/g
uides/2015-08-21-application_services.md">the source in github</a>, and the
<a href="http://matrix.org/docs/guides/application_services.html">formatted
guide on the matrix.org website</a>!</strong></p>
```

<hr/><p>Hi everyone. I'm Kegan, one of the core developers on Matrix. I'd like to explain a bit more about one of the upcoming features in Matrix: Application Services. This is an entirely new component in the Matrix architecture which gives great power (along with great responsibility!) to the wielder.</p><p>Application services are distinct modules which sit alongside a home server providing arbitrary extensible functionality decoupled from the home server implementation. Just like the rest of Matrix, they communicate via HTTP using JSON. Application services function in a very similar way to traditional clients, but they are given much more power than a normal client. They can reserve entire namespaces of room aliases and user IDs for their own purposes. They can silently monitor events in rooms, or any events directed at any user ID. This power allows application services to have extremely useful abilities which overall enhance the end user experience.</p><p>One of the main use cases for application services is for protocol bridges. As you may know, we have an IRC bridge bot on matrix.org which resides as a user on #matrix, #matrix-dev, #openwebrtc and #vuc which bridges into freenode. There is nothing special about this bot; it is just treated as a client. However, this limits the things the bot can do. For example, the bot cannot reserve the virtual user IDs it creates, and cannot lazily bridge arbitrary IRC rooms on-the-fly. By using application services, the IRC bot can do both of these things. This would allow any Matrix user to join a room alias which doesn't exist: say

```
#irc.freenode.python:matrix.org
```

, which would then tell the application service to create a new Matrix room, make the alias for it, join #python on freenode and bridge into it. Any IRC user on #python would then be provisioned as a virtual user e.g.

```
@irc.freenode.alice:matrix.org
```

. This also allows PMs to be sent directly to

```
@irc.freenode.alice:matrix.org
```

, no matter what channel Alice is on.</p><p>Application services have enormous potential for creating new and exciting ways to transform and enhance the core Matrix protocol. For example, you could aggregate information from multiple rooms into a summary room, or create throwaway virtual user accounts to proxy messages for a fixed user ID on-the-fly. As you may expect, all of this power assumes a high degree of trust between application services and home servers. Only home server admins can allow an application service to link up with their home server, and the application service is in no way federated to other home servers. You can think of application services as additional logic on the home server itself, without messing around with the book-keeping that home servers have to do. This makes adding useful functionality very easy.</p><p><strong>Example</strong></p><p>The application service (AS) API itself uses webhooks to communicate from the home server to

the AS:</p> <p>#8211; Room Alias Query API : The home server hits a URL on your application server to see if a room alias exists.<br/>&#8211; User Query API : The home server hits a URL on your application server to see if a user ID exists.<br/>&#8211; Push API : The home server hits a URL on your application server to notify you of new events for your users and rooms.</p> <p>A very basic application service may want to log all messages in rooms which have an alias starting with &#8220;#logged\_&#8221; (side note: logging won&#8217;t work if these rooms are using end-to-end encryption).</p> <p>Here&#8217;s an example of a very basic application service using Python (with Flask and Requests) which logs room activity:</p> <pre># app\_service.py:&#13; &#13; import json, requests&#160; # we will use this later&#13; from flask import Flask, jsonify, request&#13; app = Flask(name)&#13; &#13; @app.route(.,/transactions/&lt;transaction&gt;, methods=[„PUT“])&#13; def on\_receive\_events(transaction):&#13;

```
events = request.get_json()["events"]&#13;
for event in events:&#13;
    print "User: %s Room: %s" % (event["user_id"], event["room_id"])&#13;
    print "Event Type: %s" % event["type"]&#13;
    print "Content: %s" % event["content"]&#13;
return jsonify({})&#13;

&#13; if __name__ == „main“:&#13;

    app.run()&#13;
```

</pre> <p>Set your new application service running on port 5000&#160;with:</p> <pre>python app\_service.py&#13; </pre> <p>The home server needs to know that the application service exists before it will send requests to it. This is done via a registration YAML file which is specified in Synapse&#8217;s main config file e.g.

homeserver.yaml

. The server admin needs to add the application service registration configuration file as an entry to this file.</p> <pre># homeserver.yaml&#13; app\_service\_config\_files:&#13;

1. „/path/to/appservice/registration.yaml“&#13;

</pre> <p>NB: Note the &#8220;-&#8221; at the start; this indicates a list element. The registration file

registration.yaml

should look like:</p> <pre># registration.yaml&#13; &#13; # this is the base URL of the application service&#13; url: „<http://localhost:5000>“&#13; &#13; # This is the token that the AS should use as its access\_token when using the Client-Server API&#13; # This can be anything you want.&#13; as\_token: wfghWEGh3wgWHEf3478sHFWE&#13; &#13; # This is the token that the HS will use when sending requests to the AS.&#13; # This can be anything you want.&#13; hs\_token: uwg8243igya57aaABGFfgeyu&#13; &#13; # this is the local part of the desired user ID for this AS (in this case @logging:localhost)&#13; sender\_localpart: logging&#13; namespaces:&#13;

users: []&#13;

```
rooms: []#13;
aliases:#13;
- exclusive: false#13;
  regex: "#logged_.*"#13;
```

</pre> <p><strong>You will need to restart the home server after editing the config file before it will take effect.</strong></p> <p>To test everything is working correctly, go ahead and explicitly create a room with the alias &#8220;#logged\_test:localhost&#8221; and send a message into the room: the HS will relay the message to the AS by PUTing to /transactions/<tid> and you should see your AS print the event on the terminal. This will monitor any room which has an alias prefix of &#8220;#logged\_&#8221;, but it won&#8217;t lazily create room aliases if they don&#8217;t already exist. This means it will only&#160;log messages in the room you created before: #logged\_test:localhost. Try joining the room &#8220;#logged\_test2:localhost&#8221; without creating it, and it will fail. Let&#8217;s fix that and add in lazy room creation:</p>

```
<pre>@app.route(„/rooms/<alias>“)#13; def query_alias(alias):#13;
```

```
  alias_localpart = alias.split(":")[0][1:]:#13;
  requests.post(&#13;
    # NB: "TOKEN" is the as_token referred to in registration.yaml#13;
    "<a
      href="http://localhost:8008/_matrix/client/api/v1/createRoom?access_token=TOKEN"
      target="_blank">http://localhost:8008/_matrix/client/api/v1/createRoom?access_token=TOKEN</a>",&#13;
    json.dumps({#13;
      "room_alias_name": alias_localpart:#13;
    }),#13;
    headers={"Content-Type": "application/json"}#13;
  )#13;
  return jsonify({})#13;
```

</pre> <p>This makes the application service lazily create a room with the requested alias whenever the HS queries the AS for the existence of that alias (when users try to join that room), allowing any room with the alias prefix #logged\_ to be sent to the AS. Now try joining the room &#8220;#logged\_test2:localhost&#8221; and it will work as you&#8217;d expect. &#160;You can see that if this were a real bridge, the AS would have checked for the existence of #logged\_test2 in the remote network, and then lazily-created it in Matrix as required.</p> <p>Application services are powerful components which extend the functionality of home servers, but they are limited. They can only ever function in a &#8220;passive&#8221; way.&#160;For example, you cannot implement an application service which censors swear words in rooms, because there is no way to prevent the event from being sent. Aside from the&#160;fact that censoring will not work when using end-to-end encryption, all federated home servers would also need to reject the event in order to stop developing an&#160;inconsistent event graph. To &#8220;actively&#8221; monitor events, another component called a &#8220;Policy Server&#8221; is required, which is beyond the scope of this post. &#160;Also, Application Services can result in a performance bottleneck, as all events on the homeserver must be ordered and sent to the registered application services. &#160;If you are bridging huge amounts of traffic, you may be better off having your bridge directly talk the Server-Server federation API&#160;rather than the simpler Application Service API.</p> <p>I hope this post demonstrates how easy it is to create an application service, along with a few ideas of the kinds of things you can do with them. Obvious uses include build protocol bridges, search engines, invisible bots, etc. For more information on the AS HTTP API, check out the new <a

href="„<http://matrix.org/docs/spec/#application-service-api>“>Application Service API</a> section in the spec, or the raw drafts and spec in <a href="„<https://github.com/matrix-org/matrix-doc/>“><https://github.com/matrix-org/matrix-doc></a>. &#160;The AS API is not yet frozen, so feedback is very welcome!</p>

```
        </div><div alt="article">
        <p>One of the final remaining missing bits of Matrix today
is specifying and implementing the Application Service (AS) APIs which allow
you to easily extend Matrix with custom server-side functionality. The AS
APIs should let you perform any arbitrary manipulation on chatroom contents,
modulo end-to-end encryption constraints &#8211; e.g. machine translation;
archiving/searching contents; interactive automated services; conferencing;
firing push notifications and other hooks; etc. If you really want to look
behind the curtain, the bug tracking the development (somewhat out-of-date)
is at <a href="https://matrix.org/jira/browse/SPEC-34">SPEC-34</a>.</p>
```

<p>However, the most obvious use case for this is gatewaying Matrix rooms through to existing communication platforms (XMPP, SIP, non-standardised systems) &#8211; which is obviously key to Matrix&#8217;s overall goal of defragmenting communication. &#160;And the good news is that even though the AS APIs don&#8217;t yet exist, we can still make good progress here through the existing client-server API. &#160;Anyone who&#8217;s hung around chat systems like IRC should be familiar with the idea of bots &#8211; non-human users who participate in chatrooms and but perform arbitrary automated functionality &#8211; and you can go quite a long way in using the &#8216;bot&#8217; idiom to add automatic functionality to Matrix.</p> <p>[In fact the first AS API we&#8217;ll be adding will probably be simply extending the client-server API with some additional privileges to allow homeserver admins to hook up services to their server which are essentially privileged bots (e.g. ability to autojoin rooms created on that server with high power-level; ability to flag themselves as an invisible &#8216;service bot&#8217;;&#160;ability to monitor events from rooms without joining them: useful for read-only services such as sending push notifications, adding search/archive functionality; etc). &#160;This should also be familiar to IRC users, as it&#8217;s similar to&#160;the model that IRC Services uses.]</p> <p>So, we already have a few bots hanging around prototyping out bridging to other systems, which hopefully should evolve into full Application Services (where it makes sense; sometimes a bot is good enough). &#160;For instance, we have the <a href="„<https://github.com/tm604/Matrix-IRCBridge>“>Matrix/IRC Bridge</a>&#160;thanks to tm604 and LeoNerd. &#160;The way this works is simply a bot which joins IRC channels and their Matrix room equivalents; watching the messages on both sides of the bridge and relaying them to the other side, creating virtual users as required. &#160;In future we can be smarter about having the bridge talk on behalf of actual users, or letting actual users control their virtual users, but it&#8217;s good enough as a first cut.</p> <p>So for Friday&#8217;s <a href="„<http://www.voipusersconference.org/2014/vuc517-matrix-org/>“>VUC 517</a>, we decided at the last minute (on Tuesday) to make as much of VUC accessible via Matrix as possible. &#160;One side of this was hooking up the <a href="„<https://jitsi.org/Projects/JitsiMeet>“>Jitsi Video Bridge</a> to be accessible by talking to the underlying XMPP server &#8211; the other side was hooking up via SIP to the <a href="„<https://www.zipdx.info>“>ZipDX audio bridge</a>&#160;that is used for&#160;audio-only participants in the conference. &#160;Both of these would be done as Matrix bots &#8211; a virtual user that you could voice/video call 1:1 over Matrix which would then route your call into VUC appropriately.</p> <p>By Thursday, the Jitsi bot got to the point of being able to place calls and see a single video stream (picked at random), but the&#160;video uplink wasn&#8217;t getting through and actually selecting the right stream to watch (or viewing multiple

streams) wasn't in place either. I'm sure there'll be a similar blogpost once it's finished, so I'm not going to talk about it further here. Meanwhile, on Thursday night we hacked together a separate bot to jack into the ZipDX bridge via SIP. Tim Panton's suggestion was to use <a href="http://www.freeswitch.org">FreeSWITCH</a>';s <a href="https://freeswitch.org/confluence/display/FREESWITCH/mod\_verto">mod\_verto</a> for this, and after Anthony Minesalle and Mike Jerris from FreeSWITCH had popped up on <a href="http://matrix.org/beta">#matrix:matrix.org</a>';on Tuesday to find out what we're up to, this seemed like serendipity.</p> <p>We hadn't played with mod\_verto before, although had been pointed at it by someone on IRC shortly after we released Matrix it's a cool project from the FreeSWITCH dev team that exposes a simple JSON-RPC interface for call signalling over websockets, providing a much more suitable way for WebRTC developers to get calls in and out of FreeSWITCH than shoehorning a SIP stack into their browser. In this respect it's quite similar to Matrix, but there are some huge differences:</p>

<ul><li>Verto doesn't (yet) do federation; either for message-passing (like XMPP) or history-replication (like Matrix or XMPP FMUCs). In fact, Matrix fundamentally competes more with JSON-RPC at OSI layer 5 by defining a standardised RESTful API for federated state synchronisation; which so happens to define some datatypes for VoIP signalling; whereas Verto currently seems to be focused solely on the application-layer problem of VoIP signalling.</li> <li>Verto is a generic RPC API with method names like verto.invite, verto.answer, verto.media, verto.bye etc. for handling call signalling. It's obviously readily extensible to any other functionality expressed as an RPC invocation. The Matrix client-server API however passes around event objects within the context of a room; it's message passing rather than RPC.</li> <li>Matrix's VoIP events support trickle-ICE; announcing ICE candidates from WebRTC as and when they become available. This good is for speedy call-setup (you don't have to wait for all ICE to complete before setting up the call) and to support call continuity when roaming between different IP networks (in theory). However, Verto currently requires ICE candidates to be presented monolithically; it hasn't made sense to implement trickle-ICE as FreeSWITCH's core engine doesn't support it yet.</li> <li>Verto looks to be wired to speak JSON-RPC over Websockets, whereas Matrix deliberately uses plain old HTTP as its baseline for maximum simplicity and compatibility in PUTting and GETting JSON events around</li> <li>Verto could be an interoperable standard but the API hasn't been documented yet (unless I've totally missed it); to build the bot we looked at the websockets tab in Chrome's network inspector and drew some inferences from the <a href="https://github.com/matrix-org/synapse/blob/e43139ac5e1f337d9a82ee16d9a4f15195120ec3/contrib/vertobot/vertor-example.json">JSON seen in a call</a> placed using the <a href="https://webrtc.freeswitch.org/vertor">FreeSWITCH Verto demo site</a>, which was very intuitive and almost self-documenting. Meanwhile, the (minimal) doc for Matrix's events is up at:<a href="http://matrix.org/docs/spec/#voice-over-ip">http://matrix.org/docs/spec/#voice-over-ip</a>.</li> </ul> <p>Verto has a huge advantage however, in that FreeSWITCH has a mod\_verto today, and doesn't have a mod\_matrix; so one can use mod\_verto right now as an easy way to get VoIP calls in and out of FreeSWITCH from the web and bridge them to SIP! So, when writing a Matrix-&gt;SIP bridging bot for VUC, Verto turned out to be a really nice way to quickly get something up and running. The end result is at:<a href="https://github.com/matrix-org/synapse/blob/e43139ac5e1f337d9a82ee16d9a4f15195120ec3/contrib/vertobot/bot.pl">https://github.com/matrix-org/synapse/blob/develop/contrib/vertobot/bot.pl</a>'; a (precisely!) 300 line perl script built on LeoNerd's <a href="https://metacpan.org/release/Net-Async-Matrix">Net-Async-Matrix</a>'; and <a href="https://metacpan.org/release/Net::Async::Protocol::WebSocket">Net::Async::Protocol::WebSock

et</a> which logs into Matrix and routes any 1:1 audio calls it receives through to the defined mod\_verto service. &#160;Currently it gleefully hardcodes in the destination extension it calls and the caller-ID, but this could trivially be extended. &#160;It also chokes on SSL WebSockets, so your mod\_verto needs to be listening unencrypted on port 8081.</p> <p>The task of mapping between Matrix m.call.\* VoIP events and Verto verto.\* RPC methods was almost entirely trivial (although I hasten to add that Matrix&#8217;s and Verto&#8217;s were developed completely independently &#8211; it&#8217;s just that there are only so many ways to express VoIP signalling in JSON!)</p> <ul><li>Matrix&#8217;s m.call.invite is equivalent to the combination of verto.invite + verto.media (but may lack&#160;ICE candidates)</li> <li>Matrix&#8217;s m.call.candidates has no direct equivalent, so has to be coalesced and merged into verto.media</li> <li>Matrix&#8217;s m.call.answer is equivalent to verto.answer (but may lack ICE candidates)</li> <li>Matrix&#8217;s m.room.displayname is equivalent to verto.display (assuming I understand verto.display)</li> <li>Matrix&#8217;s m.call.hangup is equivalent to verto.bye</li> <li>&#8230;and these are the only verto RPCs we mapped.</li> </ul><p>For the demo itself, we obviously needed a FreeSWITCH with mod\_verto all up and running to hook into the ZipDX bridge: our friends at <a href=„<http://truphone.com>“>Truphone</a>&#160;were good enough to provide one at zero notice (Thanks James, Andy, Giacomo!), and we were up and running.</p> <p>Unfortunately we did hit some problems: Net::Async::Matrix has a few quirks which LeoNerd is working out currently; the bot doesn&#8217;t coalesce the trickle-ICE properly currently causing a race-condition that means ICE setup may fail; Matthew&#8217;s use of IO::Async was a bit buggy; and moreover we didn&#8217;t have time to implement DTMF&#8230; which was a bit of a shame as you can only unmute yourself on the ZipDX bridge via DTMF \*5!</p> <p>But in general, the mini-hackathon was a success and it was great fun to be able to listen into VUC via the bridge and demonstrate the first ever Matrix&lt;-&gt;SIP call! &#160;The bot ran as @vucbot:matrix.org, although is turned off now as there&#8217;s no VUC to listen to, and the FreeSWITCH &amp; bot were only deployed temporarily. &#160;Once the kinks mentioned above are sorted out we&#8217;ll hopefully set it running again permanently! &#160;And hopefully this little bot is an exciting precursor to more robust Matrix bridges and application services in the months to come&#8230;</p> <p>If you&#8217;re crazy enough to want to try to run the bot yourself, then it should actually be quite simple to get up and running:</p> <p>

```
# grab synapse if you don't have it already<br/>git clone  
https://github.com/matrix-org/synapse.git synapse-develop<br/>cd synapse-  
develop
```

</p> <p># you'll need the develop branch, as we haven't released a build with vertobot in it yet<br/>git checkout develop<br/>cd contrib/vertobot</p> <p># you'll need cpanm to install the perl dependencies<br/>cpan -i App::cpanminus<br/>cpanm -installdeps .</p> <p># manually install a develop version of Net::Async::Matrix as cpanm can't figure it out, seemingly<br/>cpanm -force PEVANS/Net-Async-Matrix-0.11\_002<br/># (you may need to also replace the 'croak' for the „Already have a room with ID“ error with 'warn' in Net::Async::Matrix if the bot crashes with this error)</p> <p># create a username account for your bot on a Matrix homeserver somewhere at this point</p> <p># set up a config file<br/>cp config.yaml mybot.yaml<br/># edit mybot.yaml to taste - at the least you must specify the login &amp; password &amp; homeserver for your bot!</p> <p># run it!<br/>./bot.pl -c mybot.yaml<br/></p> <p>Finally, huge thanks to everyone to helped make the VUC bridging escapade work out &#8211; Emil Ivov at Jitsi, James Body, Andy Smith and Giacomo Vacca at Truphone, Anthony Minesalle &amp; Mike Jerris &amp; Brian West at FreeSWITCH for writing freeswitch and mod\_verto, Tim Panton for the VUC intro and suggestion of mod\_verto, Randy Resnick &amp; Michael Graves at VUC itself, and of course the Matrix team for glueing our side of it together!</p> <p>Looking forward to lots more ambitious cross-protocol gatewaying and federation

in future!</p>

```
</div>
```

```
</html>
```

From:  
<https://schnipsl.qgelm.de/> - **Qgelm**

Permanent link:  
<https://schnipsl.qgelm.de/doku.php?id=wallabag:advanced-synapse-setup-with-lets-encrypt>

Last update: **2021/12/06 15:24**

