Julien Dubois **@juliendubois** *Tue Jun 02 21:19:55 +0000 2020*

The French #StopCovid application is out! Thanks to @mhausherr I had the URL to their backend server, so let's have a look at how they are hosting this ■■

The French press is currently reporting that the maintenance and hosting of this application will be billed "between 200,000 and 300,000 Euros" per month, see **https://t.co/PZiYEObuVT** so we can expect some really great stuff here

The application has been quite instable today, which made things a bit difficult at first. We're going to see why very soon ■

First thing, let's have a look at the SSL report -> **https://t.co/mVTJQcpRIx** and it's indeed not so bad! ■
The certificate is issued by Gandi (which I also use myself for all my websites), which is a great French company.

Second thing, hosting -> **https://t.co/8bLw56S7rx**
That's more interesting, their server is "Apache/2.4.6 (CentOS) OpenSSL/1.0.2k-fips PHP/7.4.6".
Now this is the kind of thing that shouldn't be public. I usually call "nginx" my Tomcat servers just to annoy hackers, for example.

Also why PHP? And why an old version of Apache, an old version of OpenSSL? This looks like a base image that was installed manually.
Searching for those versions: this is a CentOS 7. Not very recent, but still maintained, so that's good.

Please note that for something like that, I wouldn't use the OSS version. I would pay RHEL from RedHat, to have the support and the patches.
Remember the hosting costs? That's why you pay this kind of money.

Now let's see the host name: **https://t.co/blckz7Hlki**
This is good: the eu-west-2 region of Outscale. So this is indeed hosted by a French company (subsidiary of Dassault), and in France. See **https://t.co/bhIzzB4L85** where you can see they have (only!) 2 availability zones

OWS is for "Outscale Web Services", and that's a clone of AWS, compatible at the API level (they configure them with the AWS CLI). So we're just directly on a VM.
Having a look at their base images: **https://t.co/nsZFdlEiaq** we find our good old CentOS 7 from earlier.

So there's no API gateway or fancy stuff... We just arrive directly on a VM, using an old (but maintained) CentOS. And they have the default Apache running on it, which explains PHP.
This also explains today's downtime, there's just nothing specific for scaling or availability.

We have of course no details from their MongoDB instance or their PostgreSQL instance, but I'm guessing it's the same manual installation of everything.
BTW there are 3 collections in MongoDB and 1 table in PostgreSQL, so nothing huge.

So that was pretty boring, given the price tag I'm sure we were all expecting something a bit more fun!!
The good news: this is indeed hosted in France, hosted by a French company, as promised.

Interesting news from today! As commented in the thread they added a Kong gateway during the night.
https://t.co/zFtvZeUdCu
Let's have look!!! ■■ **https://t.co/dCn2OcHSyy**

This is the latest version of Kong, so definitely they "learnt" that using an API gateway was a good idea following yesterday's downtimes.
Looks pretty amateur-ish to me to do that in a hurry, that means they probably didn't read the full docs...

Indeed I hope they read **https://t.co/ewm0IYsv62** and its GDPR implications.
They had an agreement and an audit from **@CNIL** a few days ago and then they add this in a hurry?

And now back to my original thread on this application: they have a log in Kong with your IP addresses, and very likely the HTTP request payloads, which contains your keys.
This was installed in a hurry AFTER the audit, and not in the OSS code that was provided.

And today we have even more great news! As noted by **@bluxte** there are in fact 2 URLs. So you shouldn't do like me, and click on links without checking! So there is **https://t.co/AEMEUj7dy3** and **https://t.co/roEAiiEwI9** (yes, you need to have good eyes).

So my previous tweets might have been wrong, and we had those 2 URLs from the start - which is a bit weird but nothing more. It's hard to know, because maybe they have changed their setup in-between. The good news is that I had a closer look at their Kong URL ■

First good thing: they did secure the Spring Actuator endpoints.
I did warn them a few days ago, of course: **https://t.co/dFPuGk6UIO**
That's the issue when you don't Open Source all your application. There was no way to know they had this Kong server.

And now let's test their Kong setup, by doing a GET on **https://t.co/dPUvkhtFlB** - can you see those "rate limiting" HTTP headers? It is this plugin: **https://t.co/ZIkULPO3KG**

It's good to use rate limiting, and do it from France (unlike the stopcovid form that uses a US-based server for this!).
But this plugin also uses your IP for this, and if they did the full setup, those IPs are stored in a database.

So we have here a clear proof that IPs are logged somewhere on the server-side, and that isn't documented anywhere, especially not in the "Open Source" code that was provided.
The **@cnil** is pretty clear that this is personal data, as you can identify some people with IP.

Besides, we're talking about the French state here. They have build the Hadopi mechanism, so if you already got one of their e-mails (like me ■), you know they can identify you very easily with your IP address.

So really, how can the **@CNIL** validate such a system, that stores your IP address, maybe even in different systems (logs and rate limiting)? Did they only have access to the Open Source code, like us?

Oh that's incredible, I had an official answer (I guess, as anybody can create an account and answer on this system):
https://t.co/LnTb3guX4A
Let's have a closer look! ■■

Indeed, you can configure Kong to only do a "global" rate limiting, and that would be good.
But should we trust that person, given the fact that they added this gateway without describing it anywhere? Also, maybe they recently changed their config. But let's test.

If you do the following command you'll get those headers:
curl -sSkv **https://t.co/dPUvkhtFlB** 2>&1 | grep x-

And indeed, it seems to go down globally (I even tested using a VPN, to change my IP).
So that would be correct, then the lack of transparency is worrying. It would be so much easier if the app was OSS and fully documented!
But wait, now we know how many requests they have????

If you do those requests, you'll notice that it very rarely goes under 59805 available requests, out of 60000 per minute.
So they have 200 API calls every minute.
But yesterday they claimed to have 600,000 users?
https://t.co/FU2itI22m1

Also, if we do the math, and I took the medium price of 250,000 Euros per month for maintenance that was in the press - that's about 3 cents per API call. That's quite a good price ■

Some more (sad) **#StopCovid** news from today!
As noted by **@StopCovidData** the Kong data isn't available anymore... I'm guessing they are monitoring Twitter, too ■
Here are my final thoughts on this ■■

Despite everything that was announced ("transparent", "Open Source"), the government has hidden ON PURPOSE several important software components, that could be used to identify people, and are likely not GDPR compliant.

We saw them changing their configuration every day (tonight they added 2 new Kong servers), so even if that was audited by the **@CNIL** earlier, this audit can't be trusted anymore.

More importantly, the next day we report they have this unknown piece of software that can be used to track people, they remove the traces from the HTTP headers in order to cover it. Very transparent! That's totally what we do in Open Source...

Of course, some monitoring data was already gathered and we could see that close to nobody uses the application, and more importantly that the trend is stable. The government said they had 600,000 installations, but installation doesn't mean usage.

So let's just stop making advertisement for this application and go back to a normal life.
Just think that all of this money could have been used to pay French nurses better, and that they are the one saving lives for sure!