

HOW TO ANSWER TECHNICAL CONSTRAINTS ON CHINESE MARKET?

Introduction

With more than one billion inhabitants, China is an expanding market demanded but not as easy to access. The establishment of the Great Firewall restrains access to the market to Western Companies. You should come prepared when you want to gain a foothold in this market.

Different constraints will stand in front of you:

- On one hand, technological constraints due to hosting and testing solutions, accessibility of your website or mobile application.
- On the other hand, cultural constraints linked to Chinese population's consumption habits with the use of local online chat or payment system.

Hence this is primordial to take into consideration all of those factors to set up your business in China.

Following the opening of our subsidiary in Hong Kong, TheCodingMachine understands the constraints of this market.

In this white paper, we will explain the challenges provoked by the legal and technological constraints and how you can deal with them.

1 Cloud service and Great Firewall, how to make it match?

Great Firewall: definition and issues

You have probably already heard about what we call the Great Firewall.

It is a combination of legislative actions and technologies enforced by the Republic Of China to regulate the Internet domestically. The objective is to block the access to selected foreign websites and slow down cross-border internet traffic. It leads to:

- limiting access to foreign information sources
- blocking foreign Internet tools and mobile application (Google, Facebook, Wikipedia...)
- Requiring foreign companies to adapt to domestic regulations.

The Great Firewall has seen several transformations over the decades and has become complex and sophisticated. It is perpetually changing and tends to be more and more restrictive, especially on Internet traffic slowdown.

Servers and SaaS restrictions

You have a limited choice in China, but if you are looking to do business there, you better host your website locally for the different reasons expressed earlier. However, hosting in China will cost a bit more than another region in the world.

All of this process is known to be long and tedious, but we have found ways to facilitate access to the market.

2 Find the right hosting solution to answer technical constraints

Different solutions are offered to answer the technical constraints pushed by the Great Firewall.

Internet Content Provider (ICP)

An ICP license is a mandatory legal requirement by the Chinese Government. It is a state-issued registration that allows you to host your website on a server inside mainland China. Your ICP licence is given once the registration is made on a local Cloud Service Provider.

But you have to know some indications to evaluate if it is worth it or not:

- **Only sites within the Chinese mainland are required by law to have an ICP, if your site is hosted outside mainland, you do not need one. And it is not even possible to apply for one.**
- **ICP has nothing to do with whether or not your site is visible or blocked in China.**

With this information you can estimate if you really need an ICP or not. It actually depends on how serious you are about entering the Chinese Market and on what scale. If you are a small or medium business with a basic website that is only reaching Chinese users, then you can probably host in Hong Kong with no problem.

THE APPLICATION PROCESS

1. Register a domain in China and buy hosting solution which will be blocked until your ICP is validated.
2. Submit your ICP application form and documents to the web host. Your host will be informed and check with the Ministry of Industry and Information Technology (MIIT).
3. If the application is approved, MIIT notifies the host and it unlocks your account and you are good to go.

For more information, you can find all steps to follow [here](#)¹.

All of this process is known to be long and tedious. But you have other ways to access the Chinese market.

- **Host your website or mobile application in Hong Kong or Singapore**

Hong Kong or Singapore are excellent entry points to the Asian market. Actually, they are known to be the gateway to China and it is also the same case for web hosting (it can provide a fast and reliable hosting solution).

This is the best way to avoid the ICP licence process and so to host in China mainland.

- **Website speed:**

websites hosted outside the wall tend to be slower than the ones hosted in mainland China (due to server distances). Websites hosted in HK, Singapore, Japan or South Korea will be faster than ones hosted in the US or Western Europe.

- **Website access:**

websites are blocked by the Great Firewall only if they have problematic content, no matter where they are hosted. The Great Firewall does not work on where your website is hosted but on the content of it.

Several Chinese cloud service solutions are available in Hong Kong as Aliyun (Alibaba Cloud).

- **Cloud Enterprise Network**

Another way is to go through a Cloud Enterprise Network, a service offered by Aliyun (Alibaba Cloud).


Cloud Enterprise Network (CEN) through Alibaba is a service that allows you to create a global network for rapidly building a distributed business system with a hybrid cloud computing solution. You can have access to a set of virtual networking capabilities and it interconnects the Alibaba Cloud Network resources around the world.

This solution needs an ICP licence and is quite expensive to settle.

¹ Kendra Schaefer. *Chinese ICP Licensing: What, Why and How to Get Hosted in China*.
<https://bit.ly/2cSdDzM>

HOW IT WORKS?

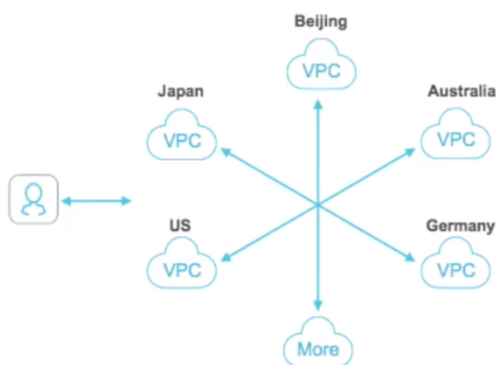
It is quite simple to do this. You have to create your CEN and attach network instances via service provider's private tunnel with desired bandwidth. CEN only allows communication between Aliyun data centers so unless you host your application in a european Aliyun data-center (like in Frankfurt), you will have to proxy the calls from the european Aliyun data-center to your servers.




HOW TO USE?

1. Create a CEN
2. Attach network Instances to CEN
3. Buy cross geography bandwidth package
4. Configure cross region bandwidth to intercommunicate between regions (only if they are not in the same region, otherwise it is not necessary).

Express Connect (point to point interconnection)



CEN (Access to interconnected global network)



Warning: with an average price of \$100/month for a 1Mbit/s bandwidth, the price of CEN is prohibitively high (you would need to pay \$2000/month to get the bandwidth of a poor ADSL connection)

- **AWS in China**

Amazon Web Services (AWS) has data centers in Beijing and Ningxia Regions to provide the best experience for customers in China and to comply with China legal and regulatory requirements. It has collaborated with China local partners for delivering cloud services, Sinnet for Beijing and NWCD for Ningxia.

If you are used to AWS, be aware that AWS China needs to be viewed as a completely separate account with its own billing and support. While some things will work such as Cloud Formation templates and Code Deploy scripts, assuming they only use the services available in China, most things will need to be setup again from scratch.

As for other Chinese Cloud Services, an ICP licence is mandatory. Nevertheless, if you already have an AWS account you have to register on a separate account from AWS China and vice versa.

You can find all the information needed [here](#).

3 Keeping your data synchronized

So, you decided to host your website in China? You are now facing a tough decision to make: how can you keep your data synchronized between your server in mainland China and your main server?

Based on your business model and on your technological stack, the answer might vary.

Split brain scenario

Do you need to synchronize data at all?

Maybe your business in China is sufficiently separated from your business outside of China that you don't need to synchronize data at all (or very seldomly through a bunch of batch files).

This is the happiest scenario as you will probably not suffer from the lag incurred by the great firewall at all.

One-way synchronization

If “split brain” is not an option, then you will need to rely on some form of synchronization. Calls through the great firewall can be delayed up to 30 seconds, so any real-time calls between China and the rest of the world is out of the question.

If your data flows in a predictable way (if you know that some tables of your data model will only be updated in your local server or only in China), you can set up one-way replication.

Depending on your architecture, this can be done in a variety of ways. Here are some examples:

- If you are using MySQL, you can set up some tables of your database in “master”/“slave” configuration
- If you are using a CQRS architecture, and if your data is not updated too often in China, you can transfer events asynchronously between China and the rest of the world. Tools like RabbitMQ federation can help you do this reliably
- Synchronizing files can be done relatively easily using a simple script with a “rsync” command
- ...

Two-ways synchronization

If you need two-ways synchronization, things are going to be way harder. Any solution allowing for active-active database replication will be unusable. Indeed, active-active clusters usually require a latency below the millisecond.

So you are in for a bit of work.

- Any technology that supports the notion of offline synchronization can be useful. For instance, some NoSQL databases like CouchDB have native support for offline synchronization. You will of course need to decide what happens if a record is updated at the same time in China and on your local server. Who wins?
- If your database does not support offline synchronization, you will have to code this at the application level: listen to all events modifying the database and send them through a bus to the other server.
- You will need to assess the likelihood of a conflict. Assessing that you have a latency of about 30 seconds between the servers, what is the likelihood that a particular object is updated at the same time in and out of China? If you assess that this likelihood is low, you can sometimes trick your database by setting up 2 “one-way” replications (see our case study below)

Finally, if you are operating on real time data with a high probability of updates on both sides, you are in the worst-case scenario. For instance, if you are writing a chat application linking a Chinese user and a non-chinese user, you will have to pay for guaranteed bandwidth. Check the “CEN” option in previous chapter.

4 How to test your product?

Your website is now hosted on a performant server, while you and your team are outside of China, the next step is to make sure that you can test your website or your application efficiently on Chinese internet.

You have two options:

1. Test in mainland or ask a local person to test.
2. Download VPN tool (ex: <https://www.speedin.com>) and connect to Chinese Internet:
 - For the Web on PC or Android on Mobile, work and debug with Google Chrome Browser by using developer tool
 - For iOS, work and debug by using Network Instrument tool on Xcode to monitor the performance

5 Integrate with Chinese payment technologies

Hong Kong as China Mainland are mobile first. The majority of the payment transactions are made with mobile phones. It is important to take this aspect into consideration to be sure not to miss your target.

The easiest way is to start with Chinese credit card payment "UnionPay", which is a Chinese financial services corporation, providing bank card services. With billions of credit cards in circulation worldwide, UnionPay is the most well-known payment system in the world, and used and privileged by Chinese consumers.

Alipay and WeChat pay are both Chinese means of E-payment, one was born from Alibaba and the other from the chat application WeChat. They are both highly used and inescapable in China but they are complicated to use for a foreign business, especially for cross-border business, as you have to go through a third-party service provider for payment settlement.

In Hong Kong the consumption is quite different. They mainly use Cash in daily life or their bank application such as PayMe or Faster Payment System (FPS) for e-payment or e-transactions. They also use PayPal which supports UnionPay. Other international payment API solutions are working as well in Hong Kong such as Stripe.

The CUP credit card owners can now have a PayPal account to purchase online. It enables merchants in European countries to accept payments from UnionPay credit and debit cards in China or Hong Kong.

6 Optimize your product

A common way to speed up an application is to use a Content Delivery Network (CDN). A CDN is a cluster of machines available world-wide that will serve your static content (images, CSS and JS files...)

A CDN can be incredibly useful in mainland China... if your CDN provider has servers in China! You should probably check Aliyun CDN solution for a fully "native" CDN.

Also, it is quite common for web developers to rely on resources hosted on third party CDNs. A typical example is the use of "Google Fonts" that are loaded directly from Google servers. Google servers are not accessible from China. If you use Google Fonts, the font loading will lock your website until the timeout expires (all your pages will take dozens of seconds to load making your website unusable).

A good rule of thumb is:

- Put all your resources (JS, CSS...) on your local server / local cluster
- If needed, use a chinese CDN to improve performance

CASE STUDY

Our client, a Parisian company, which links buyers and retailers (Walmart, Carrefour, Tesco...) with thousands of suppliers around the world.

It organized a huge summit in China in June 2019 and wanted to have high performances on their website and Web Application.

PROBLEM

Their server is currently hosted in Google Cloud in Europe. Due to the Great Firewall, there are high latencies when we try to access from China.

OUR SOLUTIONS

- **Option 1**

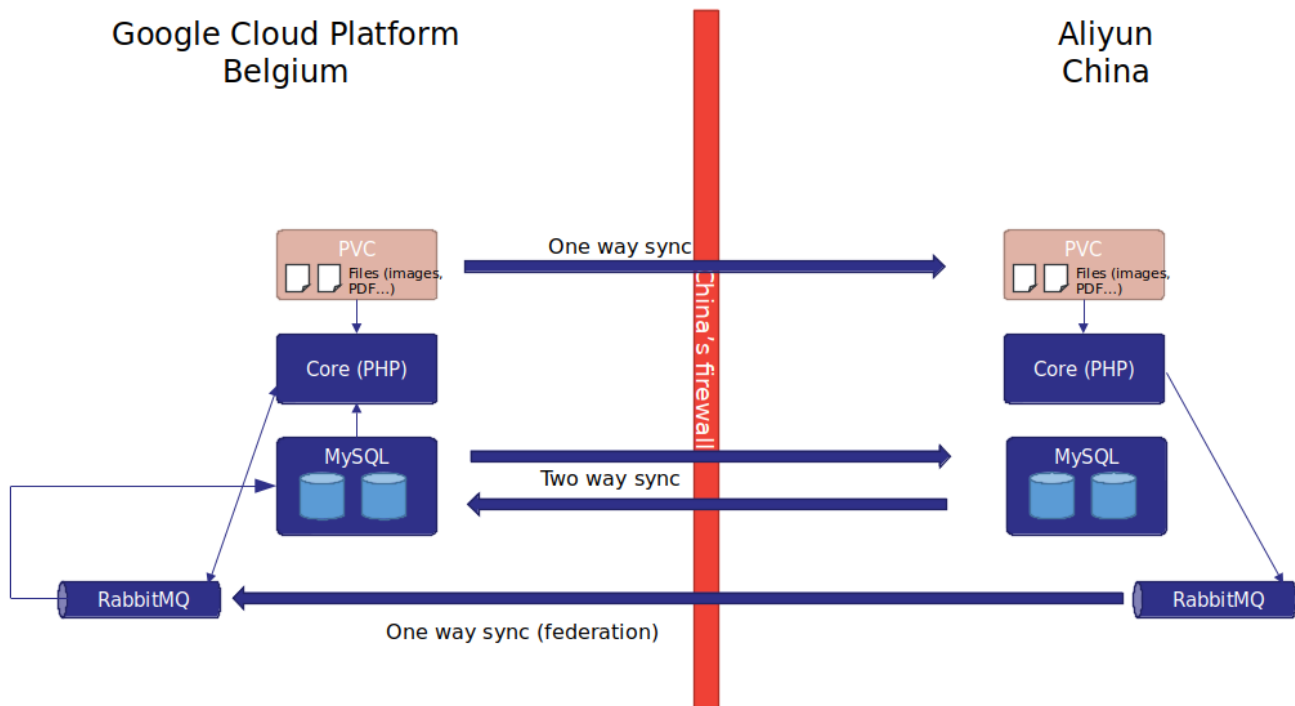
Use the Alibaba Cloud CEN (Cloud Enterprise Network), with proxy machines on both sides of the Great Firewall to accelerate the link between China and Europe, but at an extremely high cost (€1000/month for a speed of only 10MB/second... slower than an ADSL connection!).

- **Option 2**

The other solution was to replicate practically the entirety of our client platform in China and to synchronize MySQL databases through a bi-directional GTID replication and also static files (mono directional from Europe to China). This solution slashes costs but takes more time to configure. And an ICP licence is needed.

WE FINALLY CHOSE THE SECOND OPTION.

Simplified architecture of or client project replication



First steps

Our client had the chance to already be deployed on a Kubernetes cluster on GCP. Because the application was deployed on a Kubernetes cluster, deployment was already well automated. It was therefore not that difficult to recreate a complete new environment on a different Kubernetes cluster, in China.

We automated continuous integration and continuous delivery pipelines to automate images build for the containers and the deployment in China. Also, we created storage disks (aka persistent volumes in Kubernetes) in China.

The benefit is that Alibaba Cloud offers the same services as AWS and GDC, avoiding adaptation and security management issues.

The main challenge was then to manage to synchronize the state between Europe and China.

Databases synchronization (bi-directional)

Database synchronization is hard. It usually requires very low latency (less than 100µs) between servers for the servers to be able to provide ACID compliance. Obviously, with an average latency of 15-30 seconds due to the great firewall, true synchronisation is impossible.

Furthermore, the Great Firewall is always evolving and is completely unpredictable. You cannot be sure that your connection between the 2 datacenters you are using will be up. You could frequently suffer outages ranging from a few minutes to a few hours / days.

Therefore, you need to think of an architecture that can run separately and independently, and that can synchronize when the connection is up. It must be resilient to connection losses that are frequent.

Our client is using MySQL so we thought about using a Master-Slave topology. The chinese DB would have been the "slave". Read could have been possible on the chinese DB and write would have been done in Europe. This means that reads would be fast (local), but writes slow. This was not ideal. Also, the Chinese application would be down in case of connection loss. We needed another way.

We realized that we could configure a dual master-master replication using MySQL GTIDs. With an autoincrement trick, we can ensure that newly created content in Europe and China does not conflict. The solution is not perfect though. If a record is updated at the same time in Europe and China, the records will not be synchronized.

However, we estimated this was unlikely. Can you do it with your application? It really depends on your business. If your application requires frequent writes (if you are developing a chat or a video game), you will want to avoid this solution. The strategy you choose here really depends on your business.

[Do not hesitate to contact us if you want to discuss your particular needs!](#)

Files & assets synchronization (unidirectional from Europe to China)

A simple solution, rsync on the European server which connects via SSH on the Chinese server. It compares differentially the files and sends or updates only files missing or modified in a very efficient way.

RabbitMQ synchronization (unidirectional from China to Europe)

The last step was RabbitMQ synchronization, the most widely deployed open source message broker. Since it uses an asynchronous service bus, one good strategy is to centralize it on a single datacenter. So we simply duplicated the service but we made the European server listen to the Chinese one, thanks to Federated Exchanges.

Final result

After long hours of preparation, configuration, tests and reconfiguration, we finally achieved to update the production environment in Europe and China.

Our client has tested their application performance in China and they noticed a speed similar to the one they have in France.

Conclusion

The Chinese market is still in a high expansion and tend to be more and more profitable. However, accessing it is not a small job.

The Great Firewall obliges the rest of the world to adopt restrictive hosting solutions to answer those technical constraints. The good thing is that you will not fail of alternatives. The only issue is that they are more or less performant and easy to put in place. Make sure to adapt the solution to your project.

In terms of performance, you are not without testing solutions to be sure that your mobile or web app is working in China.

Moreover, in this mobile first country, where chat aps are inescapable, the adoption of local payment solutions is primordial for your business.

TheCodingMachine can help you to set up technically your business and to find the most efficient hosting solution to ensure the efficiency of your website and/or web application.

contact@thecodingmachine.com

TheCodingMachine
TCM://

