China Internet: why is it unreliable and how to fix it?

Internet in China is unreliable: it is probably less reliable than in some of the 20 poorest countries in the world. This fact is little known: most people believe that all sources of all their problems with Internet in China comes from the existence of the infamous “great firewall of China” or from government intervention. Nothing could be more wrong. The source of most problems actually come from the way Internet is being managed in China by telecommunication companies and possibly from the lack of government regulation enforcement based on the principle of network neutrality.

The current situation of Internet in China is problematic for Chinese economy: it hurts the digital transformation of industry, it slows down the development of startup companies. Cost of Internet transit within China is at least 10 times higher than in the rest of the world. Cost of hosting is between two and twenty times higher. Lead time to setup a server in a datacenter counts in weeks or months rather than days.

We will explain in this article the status of Internet in China and discuss possible solutions for companies to solve immediately the current situations until proper government regulation is implemented to enforce network neutrality within China.

Baidu is reliable

The first thing that we usually hear when we explain that “Internet in China is intrinsically unreliable” is doubt. Most people believe that the search engine www.baidu.com is a proof that Internet works well in China, since it always provides fast response about anywhere in the country with about any telecommunication company. Most people also believe that those web sites which can not be accessed reliably – such as www.google.com – are blocked by the so-called “great firewall of China”. Therefore, if a web site is not accessible in a country such as China with huge budgets for infrastructure and with world class telecommunication infrastructure companies such as Huawei, it can only be because of Chinese government intervention to block the site. Most people will then rightly demonstrate that by using a VPN software that circumvents – illegally – the “great firewall of China”, access to any web site suddenly becomes reliable.

This analysis – that great firewall makes Internet in China unreliable – is actually wrong in most cases. Many other countries in the world – including France, Turkey, Russia, etc. – have also implemented network inspection and filtering solutions similar to the “great firewall of China” but their Internet is much more stable and reliable.

The reason why Baidu and other very big web sites is reliable is also little known: those companies have created their own “private Internet” by leasing dedicated MPLS infrastructure to telecommunication companies and by setting up proxy servers in about any province or major city of the country and in every continent outside China. It is a bit like saying that Baidu is acting itself as a global telecommunication company, rather than relying on Chinese telecommunication companies for routing its data from users to its datacenters. Google actually does about the same in the rest of the world.

However not everyone has sufficient budget to implement a network strategy like Baidu. Google, Alibaba, Tencent, etc.

Chinese clouds “randomly” blocked inside China

We conducted in august 2015 an experiment to compare connectivity from different regions in China to 3 major cloud providers in China (Ucoud, QingCloud and Aliyun) . We also added servers in Canada at OVH, the largest European cloud provider and one of the cheapest in the world. Then we tested connectivity to those web sites using the https protocol from different parts of China using different Internet service providers (ISP). The tests we did were conducted both on web sites with ICP registration and web sites without ICP registration. For the record, ICP registration is requirement for any Web site hosted in China. ICP is to Web site what company registration is to companies.

The results are presented bellow.

ISP	**Ucloud (Guangzhou)**	**QingCloud (Beijing)**	**Aliyun (Hangzhou)**	**OVH (Canada)**
Yunnan FTTH	blocked	OK	OK	OK but slow
Yunnan 4G	OK	OK	OK	sometimes blocked
Guangzhou FTTH	OK	OK	50% packet loss	sometimes blocked
Guangzhou 4G	OK	OK	OK	sometimes blocked
Shanghai FTTH	OK	OK	OK	sometimes blocked
Shanghai 4G	OK	OK	OK	sometimes blocked
Paris FTTH	slow	slow	slow	OK

There was little or no difference between web sites with or without ICP registration: this probably implies that the web sites we tested were not explicitly blocked by the “great firewall”, unlike for example www.google.com which is obviously blocked in China.

Access from China to a web site hosted abroad is slow and sometimes blocked. There is apparently no clear rule. In regions like Yunnan, Internet access through FTTH to some foreign web sites was actually very stable albeit a bit slow. The same web site was accessible quite nicely from a 4G dongle then after one hour became unstable with “connection reset” errors similar to those we get when trying to access from China www.google.com. For further understanding of this unpredictable dynamic behavior, we recommend reading the article “VPN Gate: A Volunteer-Organized Public VPN Relay System with Blocking Resistance for Bypassing Government Censorship Firewalls” published by Daiyuu Nobori and Yasushi Shinjo of the University of Tsukuba[1]. This article explains

[1] https://www.usenix.org/system/files/conference/nsdi14/nsdi14-paper-nobori.pdf

quite well how access policy from China to web sites hosted abroad changes constantly and adapts within few days to any new technology intended to circumvent – illegally – Chinese Internet Law.

Access from abroad to China is slow but it is usually stable. We found some rare cases where a third party Chinese web site was blocked from abroad, but never with our own web sites. The slow speed comes probably from the lack of bandwidth that was allocated to our virtual machines combined with the high latency of Internet transit between Paris and China. We believe that by allocating more bandwidth and paying more, we could get good speed from abroad. Most cloud providers – including Ucloud and QingCloud – even provide multiple IP addresses with different geographic origin attached to a single virtual machine.

What is more surprising is that access from China to web sites also hosted within China is not stable. We experienced 50% packet loss from a Guangzhou ISP to Aliyun in some cases, and 100% packet loss from a Yunnan ISP to Ucloud. Both Aliyun and Ucloud were perfectly accessible that day from other provinces in China or through other Internet service providers (ISP) in the same province.

Please keep in mind that these observations were made during a short period of time. Blocked routes actually evolve over time: some routes which were blocked during our observations could be stable now; some routes which were stable during our observations could become blocked. A cloud provider which looks reliable one day may actually look unreliable the next day. We did find later some cases where QingCloud was blocked from certain ISP in Beijing. There is actually no way to predict or control which route is being blocked within China. We believe that the reason why this situation happens is simple: there is a lack of incentive in China to enforce stable routing across Internet service providers, combined with a lack of experience to maintain automatically correct routes between Internet service providers as a result of very high turnover of skilled staff in Chinese IT industry.

This situation is of course problematic: hosting a web site in any of the large cloud providers in China provides no guarantee at all that it will be reachable by everyone in China. Companies willing to deploy global Web applications all over China – in 300 retail stores for example – are likely to face difficulties if every day, 5% to 10% of their sites have no access to their Web application. Game companies can not ensure that all users will get reliable access to their game. This is why most companies in China rely on leased lines to interconnect all their offices at a cost which is more than 10 times higher than what is observed in the rest of the world. Overall, the poor quality of Internet routing in China leads of lack of competitiveness of the whole Chinese economy.

The only good news in our observations is that Internet access through 4G dongles is usually stable. I would therefore recommend anyone willing to access Internet in China and work on the Web to use 4G, no matter the cost. At least, it works.

Web site “blackmailing”

Another story was reported to us: it seems that from time to time, some telecommunication companies block access to certain web sites as part of a tough bargaining process. The web site that I was reported about had initially very stable access from all over China. Then the number of visits increased dramatically: suddenly, some parts of China could no longer access to it. After careful research, the owner of the web site found that one of the transit providers in China was filtering all accesses to his web site. He then had a discussion with the transit provider and – after some financial agreement was reached – the web site became accessible and stable again to everyone.

This kind of rogue behavior of transit providers in China probably means that there is a lack of enforcement of network neutrality in China. Telecommunication companies do not feel any risk of being prosecuted by the government and are thus creating all kinds of “smart” ways to generate more revenue, including blocking routes or filtering access.

These kinds of behavior can also be observed elsewhere in the world: but rather than blocking access, some telecommunication companies simply create an artificial shortage of bandwidth which makes web sites slow, until some financial agreement is reached. What happens in China, is simply a wilder version of this.

Short term solutions: technology

The short term solution to the current situation of Internet in China is technical. It is based on a multiple frontend architecture and on two key principles: frontend selection and route optimization. A frontend is a service that collects http request from the users and transfers them to the application. Multiple frontends are required to solve the problem of chine clouds being randomly blocked

Principle 1: frontend selection. A user of FTTH in Yunnan should for example access a frontend hosted at Aliyun (Hangzhou) to get best performance whereas a user in Guangzhou should access a frontend hosted in Ucloud (Guangzhou) and a user in Beijing should access for example QingCloud (Beijing). This selection evolves from day to day, from province to province and from ISP to ISP. For example, 4G users in Yunnan could get better performance with Ucloud (Guangzhou) even through FTTH users in the same province get blocked.

Principle 2: route optimization. Since the application and its database are usually hosted only in one single cloud provider, it is essential to ensure that communication from each frontend to the application hosted in other cloud is stable (which is not the case in China by default).

The architecture and the principles we just described are actually the same as those that Baidu is relying on to provide a reliable service. However, they require in the case of Baidu massive investment in dedicated infrastructure which is out of reach for most companies.

Companies called CDN – which stands for “Content Delivery Network” – provide part of these principles, by implementing an architecture similar to the one of Baidu which is then shared by multiple companies. However, the frontend selection algorithm of CDN providers is usually not efficient enough to overcome complex cases as the one in Yunnan with Ucloud or in Guangzhou with Aliyun which we described previously. And, unlike Baidu or Google which control their routing infrastructure, CDN providers can not always guarantee that routes from the frontends to the application backend will not be blocked under certain circumstances[2]. This is one of the reasons why no CDN can guarantee 100% availability in China[3]. Last, the use of a shared https frontend can

[2] Demystifying China Online – Reaching China with Your Website & Cloud Applications - http://www.cdnetworks.com/video/demystifying-china-online-reaching-china-with-your-website-cloud-applications/

[3] CDNetworks Reaching China with Your Website and Brand - The Hard Truth, p. 28 http://fr.slideshare.net/CDNetworks/cdnetworks-reaching-china-with-your-website-and-brand-the-hard-truth

be considered by some companies as a security leak since it reveals private corporate data to the CDN operator.

Grandenet

In order to solve problems that CDN providers can not solve, we have created a solution called “grandenet” which is now distributed in China and has been granted a license by Ministry of Industry and Information Technology of the Chinese government. It was designed in 2012 with the idea to provide corporate users with a solution to current problems of Internet in China while complying with Chinese Internet Law. We call it an “Application Delivery Network” (ADN) in order to highlight the fact that primary purpose is to ensure stable access to enterprise applications. We designed Grandenet IPv6 backbone network so that it does not encrypt any data it transfers: data encryption is therefore under responsibility of applications and of grandenet users.

Grandenet was implemented completely in September 2015 to provide stable access to the ERP that is used to manage all factories of a Chinese state owned company headquartered in Guangdong with factories in Yunnan, Hainan, Thailand, Malaysia and Indonesia. “grandenet” could solve a problem for which no solution had been found for more than 5 years. “grandenet” now enables this company to start its digital transformation in an complex international context.

The problem of frontend selection was solved by embedding in our web sites a small javascript that is capable to monitor the performance of each frontend and then redirects the user to the most reliable one. This Javascript is provided as an HTML5 offline application: once it is loaded, the user does not need to reload it. The first time the Javascript is loaded, we rely on the timeout of web browsers to select the best frontend. Web browsers support a feature called “multiple DNS entries” so that if one front end does not respond, another frontend is then selected using a round robin algorithm. This can take up to one minute, but it only has to be done once for all. The frontend selection process itself only takes a few seconds and can be run in background. In our opinion, our approach is reasonable for an application such as an ERP, a CRM or even for a web mail. Gmail for example takes up to one minute to load the first time. “grandenet” can achieve similar results.

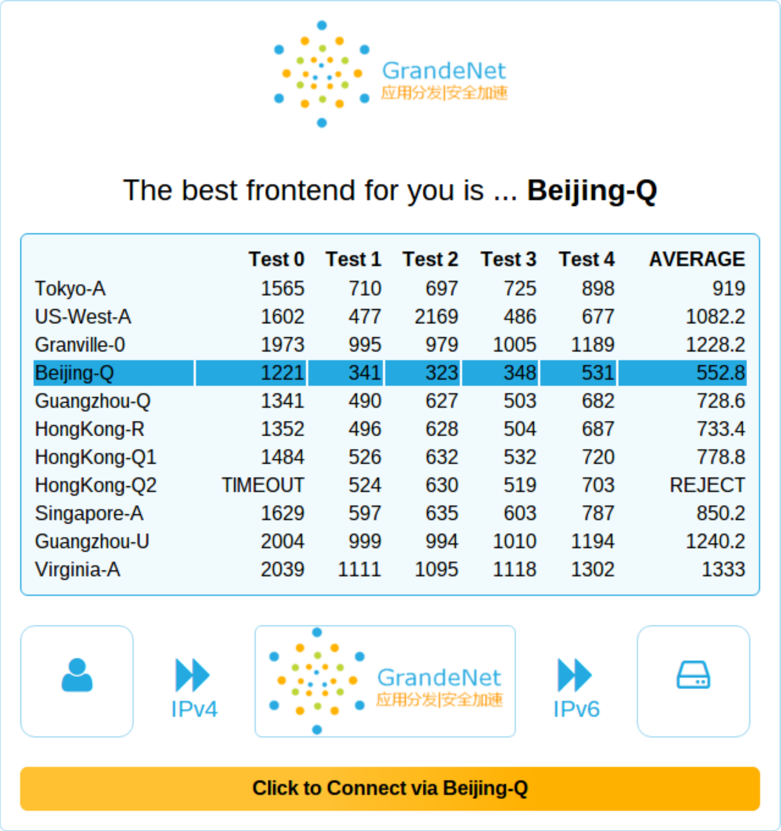

The best frontend for you is ... **Beijing-Q**

	Test 0	Test 1	Test 2	Test 3	Test 4	AVERAGE
Tokyo-A	1565	710	697	725	898	919
US-West-A	1602	477	2169	486	677	1082.2
Granville-0	1973	995	979	1005	1189	1228.2
Beijing-Q	1221	341	323	348	531	552.8
Guangzhou-Q	1341	490	627	503	682	728.6
HongKong-R	1352	496	628	504	687	733.4
HongKong-Q1	1484	526	632	532	720	778.8
HongKong-Q2	TIMEOUT	524	630	519	703	REJECT
Singapore-A	1629	597	635	603	787	850.2
Guangzhou-U	2004	999	994	1010	1194	1240.2
Virginia-A	2039	1111	1095	1118	1302	1333

Dynamic frontend selection in Grandenet – http://demoapp.node.grandenet.cn

The problem of connectivity from frontend to backend was solved by implementing a random mesh of multi-protocol Layer 2 tunnels with a latency optimizing routing protocol. Thanks to this mesh, each frontend can access the backend through billions of possible routes. If any route is blocked, the routing protocol we use (babel) finds a better route to access the backend. Because the number of possible routes is so huge, it is virtually impossible to lose connectivity. Further information can be found in this report: “Building a resilient overlay network : Re6stnet” by Ulysse Beaugnon[4].

Web sites which provide to “grandenet” a valid ICP registration will be assigned to frontends inside China and outside China. This way, we can ensure good performance no matter where the user is connection from: from China or from abroad. Thanks to routing optimization, this performance does not depend on where the application is hosted. Some users of “grandenet” host their application on Amazon or OVH outside China and use it in China. Some users of “grandenet” host their application inside China on Ucloud and use it outside China. And some users host their application on QingCloud but access to it through a frontend in Ucloud or Aliyun.

We also rely a on HTML5 offline technologies so that any static content is cached by the web browser for indefinite time. This reduces a lot the bandwidth usage and ensures much better performance for dynamic content of the Web application.

Overall, the “grandenet” approach is perfect for any application: ERP, CRM, web mail, online game, online shopping site, web POS, etc. It even solves the problem of CDN privacy by giving the option to deploy dedicated frontends on private infrastructure and by providing full source code: this is the best guarantee that “grandenet” will not act as a “man in the middle” unlike most CDNs.

[4] http://community.slapos.org/P-ViFiB-Resilient.Overlay.Network/Base_download

“grandenet” solution can also be used to provide reliable access to public web sites. However, in this case, the frontend selection algorithm previously described should be different. Our default recommendation is a geographic selection algorithm – similar to the one that Wikipedia has implemented – which can leverage the browser timeout ability. This is provided for free by “grandenet”. But in some cases, it is not sufficient. We then recommend to use advanced DNS providers such as Cedexis, which keep in real time a database of which frontend has best connectivity to which ISP anywhere in the world. Cedexis has just started to enter the Chinese market and is now adapting its solution which until now was relying on Amazon infrastructure.

Long term solution: stricter regulation

The current situation of poor Internet routing in China is in our opinion the consequence of a lack of Internet performance monitoring by Chinese government. Most countries in the world have faced or still face similar situation. Everywhere in the world except in Japan, Telecommunication oligopolies try to keep profits high by lowering Internet service quality rather than by innovating. Whereas cutting a highway immediately triggers police and government action, cutting a telecommunication route seems to trigger nothing, even though economic consequences are comparable. Unless government monitors efficiently the quality of Internet service providers and enforces clear rules such as “network neutrality”, the situation is unlikely to evolve favorably.

Monitoring the quality of service itself is a technical challenge considering the exponential number of routes and the difficulty to distinguish “human error” from “blackmail” whenever a route is being blocked. Enforcing rules on telecommunication companies requires to issue huge fines that are comparable to their profits. It also requires to break traditionally close social relationship which have existed everywhere in the world between government regulation agencies and telecommunication companies. It may be sometimes useful to grant a license to a new player that will pursue a radically different strategy: this is how Free and OVH in France modified radically the market situation, leading to a wide range of low cost, high quality services now available.

Until this happens, the only solution to overcome the lack of competitiveness of the Chinese Web hosting industry and at the same time monitor Internet performance in China is to use “grandenet” in combinations with cloud or hosting companies both inside China and abroad.