



[« Back to blog](#)

Why Google Went Offline Today and a Bit about How the Internet Works

November 6, 2012

Today, Google's services experienced a limited outage for about 27 minutes over some portions of the Internet. The reason this happened dives into the deep, dark corners of networking. I'm a network engineer at CloudFlare and I played a small part in helping ensure Google came back online. Here's a bit about what happened.

At around 6:24pm PST / 02:24 UTC (5 Nov. 2012 PST / 6 Nov. 2012 UTC), CloudFlare employees noticed that Google's services were offline. We use Google Apps for things like email so when we can't reach their servers the office notices quickly. I'm on the Network Engineering team so I jumped online to figure out if the problem was local to us or global.

Troubleshooting

I quickly realised that we were unable to resolve all of Google's services — or even reach 8.8.8.8, Google's public DNS server — so I started troubleshooting DNS.

```
$ dig +trace google.com
```

Here's the response I got when I tried to reach any of Google.com's name servers:

```
google.com.      172800      IN         NS         ns2.google.com.
google.com.      172800      IN         NS         ns1.google.com.
google.com.      172800      IN         NS         ns3.google.com.
google.com.      172800      IN         NS         ns4.google.com.
;; Received 164 bytes from 192.12.94.30#53(e.gtld-servers.net) in 152 ms

;; connection timed out; no servers could be reached
```

The fact that no servers could be reached means something was wrong. Specifically, it meant that from our office network we were unable to reach any of Google's DNS servers.

I started to look at the network layer, see if that's where the problems lay.

```
PING 216.239.32.10 (216.239.32.10): 56 data bytes
Request timeout for icmp_seq 0
92 bytes from 1-1-15.edge2-eqx-sin.moratelindo.co.id (202.43.176.217): Time to live exceeded
```

That was curious. Normally, we shouldn't be seeing an Indonesian ISP (Moratel) in the path to Google. I jumped on one of CloudFlare's routers to check what was going on. Meanwhile, others reports from around the globe on Twitter suggested we weren't the only ones seeing the problem.

Internet Routing

To understand what went wrong you need to understand a bit about how networking on the Internet works. The Internet is a collection of networks, known as "Autonomous Systems" (AS). Each network has a unique number to identify it known as AS number. CloudFlare's AS number is 13335, Google's is 15169. The networks are connected together by what is known as Border Gateway Protocol (BGP). BGP is the glue of the Internet — announcing what IP addresses belong to each network and establishing the routes from one AS to another. An Internet "route" is exactly what it sounds like: a path from the IP address on one AS to an IP address on another AS.



Welcome to the CloudFlare blog. CloudFlare provides performance and security for any website. Over 350,000 websites use CloudFlare. To learn more, please visit our [website](#).

Plans range from:

- Free
- Pro (\$20/month)
- Business (\$200/month)
- Enterprise (starting at \$3,000/month)

There is no hardware or software. CloudFlare works at the DNS level. It takes only 5 minutes to sign up.

[Sign up here.](#)

CloudFlare features

[Overview](#)
[CDN](#)
[Optimizer](#)
[Security](#)
[Analytics](#)
[Apps](#)
[Network map](#)
[System status](#)

Tags

[apps \(47\)](#)
[cloudflare \(34\)](#)
[HostingCon Interviews \(23\)](#)
[cdn \(22\)](#)
[data center \(19\)](#)
[ssl \(15\)](#)
[security \(13\)](#)
[web performance \(13\)](#)
[save the web \(12\)](#)
[DNS \(11\)](#)
[View all 641 tags](#)

BGP is largely a trust-based system. Networks trust each other to say which IP addresses and other networks are behind them. When you send a packet or make a request across the network, your ISP connects to its upstream providers or peers and finds the shortest path from your ISP to the destination network.

Unfortunately, if a network starts to send out an announcement of a particular IP address or network behind it, when in fact it is not, if that network is trusted by its upstreams and peers then packets can end up misrouted. That is what was happening here.

I looked at the BGP Routes for a Google IP Address. The route traversed Moratel (23947), an Indonesian ISP. Given that I'm looking at the routing from California and Google is operating Data Centre's not far from our office, packets should never be routed via Indonesia. The most likely cause was that Moratel was announcing a network that wasn't actually behind them.

The BGP Route I saw at the time was:

```
tom@edge01.sfo01> show route 216.239.34.10
```

```
inet.0: 422168 destinations, 422168 routes (422154 active, 0 holddown, 14 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
216.239.34.0/24      *[BGP/170] 00:15:47, MED 18, localpref 100
                    AS path: 4436 3491 23947 15169 I
                    > to 69.22.153.1 via ge-1/0/9.0
```

Looking at other routes, for example to Google's Public DNS, it was also stuck routing down the same (incorrect) path:

```
tom@edge01.sfo01> show route 8.8.8.8
```

```
inet.0: 422196 destinations, 422196 routes (422182 active, 0 holddown, 14 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
8.8.8.0/24          *[BGP/170] 00:27:02, MED 18, localpref 100
                    AS path: 4436 3491 23947 15169 I
                    > to 69.22.153.1 via ge-1/0/9.0
```

Route Leakage



[Download full size \(484 KB\)](#)

(Image Credit: The Simpsons)

Situations like this are referred to in the industry as "route leakage", as the route has "leaked" past normal paths. This isn't an unprecedented event. Google previously suffered a [similar outage](#) when Pakistan was allegedly trying to censor a video on YouTube and the National ISP of Pakistan null routed the service's IP addresses. Unfortunately, they leaked the null route externally. Pakistan Telecom's upstream provider, PCCW, trusted what Pakistan Telecom's was sending them and the routes spread across the Internet. The effect was YouTube was knocked offline for around 2 hours.

Get Updates

[Follow by email »](#)

Get the latest updates in your email box automatically.

[Subscribe via RSS »](#)

The case today was similar. Someone at Moratel likely "fat fingered" an Internet route. PCCW, who was Moratel's upstream provider, trusted the routes Moratel was sending to them. And, quickly, the bad routes spread. It is unlikely this was malicious, but rather a misconfiguration or an error evidencing some of the failings in the BGP Trust model.

The Fix

The solution was to get Moratel to stop announcing the routes they shouldn't be. A large part of being a network engineer, especially working at a large network like CloudFlare's, is having relationships with other network engineers around the world. When I figured out the problem, I contacted a colleague at Moratel to let him know what was going on. He was able to fix the problem at around 2:50 UTC / 6:50pm PST. Around 3 minutes later, routing returned to normal and Google's services came back online.

Looking at peering maps, I'd estimate the outage impacted around 3–5% of the Internet's population. The heaviest impact will have been felt in Hong Kong, where PCCW is the incumbent provider. If you were in the area and unable to reach Google's services around that time, now you know why.

Building a Better Internet

This all is a reminder about how the Internet is a system built on trust. Today's incident shows that, even if you're as big as Google, factors outside of your direct control can impact the ability of your customers to get to your site so it's important to have a network engineering team that is watching routes and managing your connectivity around the clock. CloudFlare works every day to ensure our customers get the optimal possible routes. We look out for all the websites on our network to ensure that their traffic is always delivered as fast as possible. Just another day in our ongoing efforts to [#savetheweb](#).

Update: Tuesday, November 6 11:00am PST

Moratel says the issue was caused by an unexpected hardware failure, causing this abnormal condition. This was not a malicious attempt. Moratel immediately shutdown the BGP peering with Google after contact was made while the hardware failure was being looked into.

Thanks for reading all the way to the end. If you enjoyed this post, take a second to [learn more about CloudFlare](#) or [nominate us for the 2012 Crunchie Award for Best Technical Innovation](#).

Posted by [Tom Paseka](#)

Like 3.2k

Tweet 1,945

- [125 responses](#)
- [Like](#)
- [Comment](#)

2 months ago [Lawrence Taur](#) responded:



Great job Tom.
Hey, Google! Buy CloudFlare! NOW!!!

2 months ago [Eduard](#) responded:



Great read! Tom, can you clarify in which shell environment you executed the "show route 216.239.34.10" command? Is there a similar tool for *nix?

2 months ago [moka](#) responded:



This is so bad ass.

2 months ago mary responded:



Do you mean that Google that you fixed faster than Google itself a problem related to Google?

What were google engineers were doing the same time?

2 months ago mary responded:



on the other hand this means that internet must not be concentrated into big players because then big parts of web fail.

Imagine if cloudflare is down, then how much of the internet will be down :

(

2 months ago [Danny](#) responded:



Interesting article!

I'm largely ignorant about networking principles so please tell me if this is a stupid question:

How is this different from a DNS poisoning attack? I believe they're different, and if so - could this theoretically be used to engineer similar attacks?

2 months ago [Christian Romney](#) responded:



Hey, Google—don't. It's nice to see competent small businesses thriving on their own. Heterogeneity is a good thing. Great post, Tom.

2 months ago [business model](#) responded:



nice article, you are on hacker news now

<http://news.ycombinator.com/item?id=4747910>

2 months ago Michael K responded:



Hey,

thank you for the great article.

Greetings from Germany

2 months ago [M](#) responded:



Thanks for the explanation. You described it clearly & directly. Much appreciated. Keep up the good work!

2 months ago Carlos Lopes responded:



Nice post!

i don't work in the network area so maybe what i will ask is a really dumb question :P

why the BGP doesn't check if the IP is already associated with other AS before accept it?

2 months ago [Tim Allard](#) responded:



Tom,

Thanks for the post. This was very cool to understand the big picture.

Nice job on getting things back in order. I learned a few thing here.

Tim

2 months ago [Joe](#) responded:



THAT. is awesome. You literally fixed the internet. Nice work.

2 months ago Coca responded:



Very interesting post. Can you recommend any books to learn more about internet routing? Thanks

2 months ago john conroy responded:



I hope Google at least buy you a fruit basket

2 months ago James responded:



It's not route leaking - this is route hijacking.

2 months ago Name responded:



This was actually pretty interesting to read. +1 (Also, you're on the front page of digg)

2 months ago Matt responded:



Nice work, Tom.

@lawrence The day that happens I'll promptly cancel my service. :-p

2 months ago Greg Hluska responded:



I've read many articles about how the web works, but this is not only one of the easiest to understand, it also manages to be compelling.

Wonderful work!

2 months ago Virendra Rajput responded:



does Google really need to buy Cloudfare now?

2 months ago Richard responded:



hmmmm...Nice stuff.

2 months ago fotodeka responded:



nice information...

2 months ago Tom Paseka responded:



@Eduard: This was output from a Juniper Router.

2 months ago Tom Paseka responded:



@Mary: I'm sure Google monitored and were investigating the issue, but usually back channels work faster than official paths.

2 months ago Paul Pichugin responded:



THIS is how all outages should be dealt with and explained.. thanks heaps for saving the internet :)

2 months ago Tom Paseka responded:



@Danny: this wasn't an attack, but this method can and has been used as an attack vector.

@Carlos Lopes: there was no mechanism in BGP when it was designed to check this. The folks at HE have some tools for checking multi-origin routes: <http://bgp.he.net/report/multi-origin-routes>.

@James: this wasn't a hijack, as much as a leak. there wasn't malicious attempt here and google's addresses were still announced from their own AS Number.

2 months ago Tony responded:



Hey Tom,

First off I want you to know that I really enjoyed reading this article. I am currently an undergraduate in computer science at the University of Colorado Boulder and my comp sci track is Networked Devices and Systems. I've always been interested in networks and how they work, but I've been noticing more and more how almost all companies that come to our school and offer career opportunities are looking for software engineers and it seems like almost no one is looking for network engineers. Is there anything you did in college to get your name out there or anything you can recommend for me in this case?

2 months ago [william](#) responded:



Totally bad ass! Good job!

2 months ago [Rana](#) responded:



Good effort!!

2 months ago [Forbes Burton](#) responded:



Wow, its quite silly how easy a simple miss-key or tiny fault can literally shake foundations, and this was some foundation! While it may not have been malicious in its effect, it does show the fragility of the internet. Great article!

2 months ago Simon responded:



Wow! This is one informed article

2 months ago [Tom Paseka](#) responded:



@Tony: To get yourself known, you'll need to get some exposure at network events like NANOG / www.nanog.com. This is where a large amount of the community meets.

2 months ago Kevin Phillips Bonnnng responded:



Great article, but no, BGP is not "literally" the glue of the Internet. That would have to be some sort of, you know, glue.

It's quite remarkable how often people misuse that idiom.

2 months ago Les responded:



Superb explanation! Thank you Tom.

You also expose the possibility of some "nasty" people infiltrating a network and doing what you said happened on purpose versus a possible "fat fingered" mistake.

2 months ago Carlos Lopes responded:



@Tom hm ok, i will take a look on that :)

2 months ago [Slav](#) responded:



Great article! We often forget how the internet really works - thanks for reminding us!

2 months ago [aespe](#) responded:



This is cool, innocent mistake solved quickly, great connection *with people :D* you have there tom :D

2 months ago [Stacy](#) responded:



Tom, very nice work. I am quite surprised that there isn't a small program at Google checking its routes to even its own server systems through various paths as a safeguard. This moreso considering that a few much smaller companies do have these kinds of alert systems set up (I only know of a couple of trading and PIS companies).

+10 on use of BitGlue to piece the internet back together.

2 months ago Ian responded:



What is needed is prefix lists, applied by the upstream provider, so that the peer in question cannot advertise routes not associated with their AS, or that they shouldn't be. Either that or route-maps. Due diligence on the part of the provider is required here, so that the weakness inherent in BGP's openness is compensated for.

2 months ago Giri responded:



Super !!! Great finding Tom
Greetings from India.

2 months ago [terbaru](#) responded:



internet.. trust no one, and you'll doom :D

2 months ago [Scott](#) responded:



The way I see it is the internet is an ever-evolving technology and this mishap is only another reminder of that.

2 months ago Patrick responded:



Excellent post, thanks for the insight and thanks for all your work. This really explains a lot of interesting things about the internet that I hadn't thought of before.

2 months ago Eddie responded:



Tom, I was wondering. If Moratel was injecting routes that are not behind them, wouldn't the AS path look more like "AS path: 4436 3491 23947 I"?

2 months ago [Tom Paseka](#) responded:



@Eddie: It looks like Moratel was forwarding Google's routes as originated by Google, but their actual routing didn't reflect what was in BGP, so connectivity was broken.

2 months ago [Florian MAURY](#) responded:



Good post. A word about how RPKI+ROA, or correct announce filtering would have helped to detect or even prevent the attack/misconfiguration impact would have been great, though. Fat fingers are a thing; trusting everything a peer say is another: it is sometimes called incompetence.

2 months ago [Rami GB](#) responded:



You do realize this is on the front page of reddit :P, anyways i felt happy for this reason since i my self use cloudflare free service and it's amazing i am planning to upgrading, but this is not why i wrote this comment

I just wanted to ask you Tom if you can provide me with a list of "trusted" titles on networks and DNS that would help me as a web developer, since most of my work with DNS and networking is a "guess" or "google" work.

I would love to understand networking at least at a level that allows me to create between web applications, thank you and i am waiting your reply :D.

Rami

2 months ago Keschy responded:



awesome article!

nice work =D

2 months ago varun responded:



Hi really this is very interesting article to read . or you made this interesting . explain so well. good job

2 months ago Sam Kear responded:



Thanks for the excellent post mortem analysis of the outage!

2 months ago bond responded:



Great article.. loved reading.. i learned something new today!

2 months ago James OK responded:



I work for google and I approve this message.

2 months ago Blog Tips responded:



I use CloudFlare on my website, so I was very proud to see that they fixed a problem of this magnitude. Awesome job Tom!

Google should literally have a reward for people who solve their problems before they do. Since Google is kind of an awesome company, I'm sure they'll make that happen. You might even get some Google glasses or a nice Christmas bonus, if you're lucky.

2 months ago Ribice responded:



This was one really interesting article for me and I learned something new today. Thank you!

2 months ago Peekker responded:



Great explanation. Reminds me of the needs of rice farmers. Don't divert the water in to the patches because farmers downriver need it too.

2 months ago alex responded:



yeah pretty awesome. i worked for google as well. not the classiest.

2 months ago Gordon responded:



We had the same thing happen in Australia in Feb, except it was one ISP taking down another (larger) ISP (unintentionally).

<http://www.itnews.com.au/News/291364,dodo-cops-blame-for-national-internet-ou...>

2 months ago Jian Gu responded:



I don't get it ... first, google's DNS servers network should be anycasted and AS15169 is very well connected, why your SFO router shows that to reach google's DNS server, you had to go through 3 ASes? second, I suppose AS15169 owns PI prefixes, why would Moratel/PCCW/nLayer filter routes coming from AS15169?

2 months ago David M responded:



Great Job Tom! I'll be passing this onto my class!

2 months ago AndyD responded:



My company peers with three IP carriers - AT&T, Level-3 and Cox Business. Each requires the IP addresses to be setup in the RIR. These carriers do not accept any advertisement I send without proper preregistration and administration. They filter the received advertisements based on that registration. I saw the filters come into play after the Paki-Youtube fiasco. I find it odd that the U.S. peer to Moratel does not filter Moratel's advertisements. The IPv4's are country/region specific. The U.S. carrier should know what IP's cannot come from Moratel and which ones can.

2 months ago [Jason McCreary](#) responded:



Great work. Good read.

2 months ago Jian Gu responded:



Great reading, but it makes sense to me that this was caused by a hardware failure where control plane and forwarding plane were out of sync in one of service providers routers. Why your sfo router has to go through 3 ASes to reach AS15169?

2 months ago [Tom](#) responded:



This is the most interesting thing I have read all day, thanks for the analysis. Not the normal kind of link for the front page of reddit!

2 months ago freeside responded:



what luck that the phone lines aren't going through the internets -else- how would you call them at tell them about the problem ... destination unreachable, please send by snail-mail

2 months ago Hellscreamgold responded:



Eduard responded:

Great read! Tom, can you clarify in which shell environment you executed the "show route 216.239.34.10" command? Is there a similar tool for *nix?

=====

Learn Cisco...it's your friend.

2 months ago Richard responded:



Good article.

I noted that your troubleshooting went from ns1 to ns2.google.com. Might be helpful for others to do a traceroute to see the hops where going to undersireable locations. The latency on those hops would be telling as well as you could clearly see a transcontinental path was taken.

2 months ago Hasan responded:



Great Article !

Best regard from Lebanon Tom

2 months ago [scott](#) responded:



So...now we know your username is tom, whats the password? 123?

2 months ago [Brian Olson](#) responded:



This is very interesting. My undergrad senior project was on BGP route hijacking. I agree that something like this was probably accidental. The real scary situations are more concerning. My project dealt with being able to hijack a route and then redirect the traffic intended for that route back to the destination network.

2 months ago Andy responded:



It was the mercury retrograde, duh!

2 months ago Jordan314 responded:



Good guy Tom saved google!

2 months ago Appreciative Reader responded:



This was awesome and extremely helpful. I'd like even a more detailed version if it's possible.

2 months ago cc responded:



Hello from Hong Kong. I'm using PCCW service and now I knew why Google failed yesterday morning. I think the guys at PCCW will never explain what the cause behind the issues.

Thanks you for the great explanation!

2 months ago [andrecht](#) responded:



Good job Tom !!!

greeting from Indonesia.

2 months ago [disclosure](#) responded:



You have just explained AS and BGP simply and clearly to non-techie like me using a real world problem. Certainly not the typical network engineer!

2 months ago [Budi Rahardjo](#) responded:



I very good and clear technical explanation! This is awesome. Thank you very much for taking the time to create the post. I know you don't have to, and it takes quite a bit of time and effort to do it, but posting like this what makes the internet is so useful. Again, thank you!

2 months ago oldcreek responded:



This whole thing did not make sense to me, Tom, care to explain in more detail what really happened? if your SF office could not reach google's DNS servers, then other people in SF/California/US should've likely affected also (but that was not the case); that Indonesia ISP Moratel is peering with Google, what exact route filtering Moratel fat-fingered? Even Google's DNS servers were wrongly routed to Mortatel, looks like you still had IP connectivity to it, so why your application broke?

I appreciate your response, so we all can learn from mistakes.

2 months ago Keyston responded:



yeah..daa daa dee dee.. o look the internet is broken... take this bitch..

FIXED

Google.. I think you should get with CloudFlare right now

2 months ago ra responded:



difficult to believe

2 months ago ga responded:



I cant believe that you got to know this before google engineers.

2 months ago [nico](#) responded:



I was experiencing this problem too (DNS timeout error). At first I think it was my internet, but after reading this post I know the real reason behind it. Good read.

2 months ago [Matthew Prince](#) responded:



@ga/@ra: we'll make sure to post pictures of the nice present the good folks at Google sent Tom as a thank you as soon as it arrives.

2 months ago Sam responded:



I wonder what google was doing ,were they even aware of the problem ? puts them in a really bad image. specially I have heard that they hire very intelligent people with very high IQ :D It's a joke.

2 months ago [Gigih Forda Nama](#) responded:



Thankyou for your explanation, this article was answer the question why our international connection through Moratel didn't established at past 5 Nov.

2 months ago John responded:



Those bashing google should understand that this issue has to be fixed the exact same way by google: by calling the indonesian ISP to remove the route. Perhaps google had to spend more time on a policy based trouble shooting, aka looking at application layers before moving down to network layer, we can't be sure of this. What is amazing is how a hardware failure can cause routes to be advertised to a bgp peer router, which requires more configuration than other protocols just to get a route advertised. In addition, assuming default BGP configuration, shouldn't BGP choose the shortest AS path? Cali to Cali should have a shorter path than Cali to Indo. Anyone know why the longer path was entered into the routing table?

2 months ago Fat Finger Jack responded:



So is it Fat Finger or hardware failure ? :D

2 months ago [Blog Tips](#) responded:



@Matthew Prince - I just got really exited, can't wait for the photos! Do I get a high-five for calling a gift in my previous post? Almost psychic...

/ Jim

2 months ago Rodrigo Arenas responded:



Great job.

But like others, I don't get why your router choose the wrong path when you say that your company isn't to far from the google's data center..

Can you repeat the "show.." commands and put the output...

Regards

2 months ago [Agung Sagita](#) responded:



btw, i wonder how long and how much money does it take to become trusted? and bring down the internet for 5 hours

is there any faster fixing other than contacting the originating isp to fix the error?

because, i dont know, if it only took 6 months and \$10000 to create an isp and got the trust, i think its cheap and short to create a massive downtime to the internet and the world losing billions in transactions right?

2 months ago CurRed responded:



Thanks for the explanation, insight and fix. Seriously.

2 months ago iwan responded:



why google web browser is laggy? is there impact in this case?

2 months ago Sara responded:



Tom, thank you for that interesting post. It's the first time I read something about networks that was interesting and easy to understand. Please write a book about networking. I will check on Amazon frequently.

2 months ago [Maria Burton](#) responded:



Your on [reddit](#) too

2 months ago Phil responded:



Tom, great blog and all but explain to people this outage was to your end and not all users accessing google.

2 months ago [Black](#) responded:



What kind of "an unexpected hardware failure" caused this abnormal condition? I don't know see how can a hardware make configuration changes.

2 months ago yuda responded:



I really enjoyed this Tom. As I expected from CloudFlare network engineer.

Okay, im voting CloudIFlare to Crunchies 2012. <3

2 months ago repairmanJack responded:



On behalf of those of us who are interested in these things but lack the technical foundation to interpret them, thank you for the exceptionally clear explanation.

It seems to me that the folks who can explain complicated things simply are the ones who really understand them. Tip 'O the hat for saving the intarwebs :-)

2 months ago [Lee Wei Yeong](#) liked this post.

2 months ago Joe responded:



Hardware failure? Really?

2 months ago joe responded:



hardware failure my foot

2 months ago [Amir Nasir](#) responded:



hmm....quite descriptive article...knew many things

2 months ago [Zulkifli](#) responded:



wow.. great information and good job Tom !

2 months ago Jan responded:



Yes, based on APNIC 's database import/export information Moratel AS 23947 peers with Google AS 15169. So they (Moratel) somehow broke their own connectivity towards Google peer. But why were you seeing that AS path: 4436 3491 23947 15169 as a valid and best one for those prefixies ? Don't you have shorter one ? In BGP selection criteria AS path is very powerfull one and you should see a closer one at least from US. In internet table mask for those Googles prefixies is in normal situation the same /24 as you listed here during the fault situation. So it's not that that they (Moratel) managed to advertise Googles prefixies with shorter masks. So why was that AS path chosen as a best one in your BGP table and futher to your routing table ?

2 months ago Rohith responded:



Your explanation is clear and well,ThankU

2 months ago [Nikhil Patel \(Web/Mobile Development\)](#) responded:



Great and timely fix. We recently started using cloudflare, and now I know that we made a good choice as we're in safe hands! :-)

2 months ago oldcreek responded:



@Jan possible explanation is that Cloudflare's SF office ISP nLayer 4436 has private peering with AS3491, there has higher BGP local preference for routes advertised from AS3491.

2 months ago [Tidy Design](#) responded:



Very very interesting, nice post and some awesome comments...

2 months ago [fajarm](#) responded:



Great explanation and thanks for your info. I also use Cloudflare for several my websites.

2 months ago heavycomputing-AS19117 responded:



We are AS19117. I've been involved in ISPs and running two of my own sequentially for the last 15 years. Just a heads up on my background for this statement:

Resist the calls to change this trust model. This ensures that techs are acting as humans and talking to eachother across countries, instead of relying on technical and automated security solutions. This encourages them to meet at conferences and form the most basic and humane functioning of networks: via personal communication. (This is part of the reason CERT was formed as well, to avoid purely technical solutions when all the tech is failing around you.)

While this method is subject to small glitches like this, they're pretty rare all things considered (or there'd be far more and bigger examples for Tom to draw on).

If an automated/centralized security model is adopted, many state agents will jump in and grab control of the buttons and start breaking things for political reasons. This will ensure things like the Arab Spring twitter revolutions will never happen again. And so is the same business as human rights and free speech case: businesses do not want a BIG RED SWITCH installed in the whitehouse, nor in Beijing.

Resist the fear mongering tales of a 'cyber war attack' that the 'pentagon will be best equipped to handle'. Disbelieve it. The exact reason the

internet works so well is exactly BECAUSE it is so decentralized (though that is decreasing day by day :(Artificially and purposely increasing centralization into any particular single hand will disadvantage smaller minority players. This includes small businesses, human rights groups, whistleblowers and dissidents in foreign nations, and smaller countries. Resist this temptation and stay aware of efforts to have the UN take control of ICANN (the one area of UN operations I do NOT support - Im Canadian and Canada has been a big UN support as I am personally).

2 months ago orson responded:



Tom you meant <http://nango.org> NOT <http://nanog.com> back there?

2 months ago johnd0e responded:



This is a very interesting read. Subscribed to the blog b/c of this. Thanks CF

2 months ago damanhuri_26 responded:



this was from my country
and my friends work at that ISP.
what an unexpected malicious!

2 months ago Jon Hartman responded:



A hardware failure that modifies route advertisements? As a fellow network engineer, I very much doubt you believe that for a second. The fact that they decided to lie about it, rather than just admit "we screwed up" shows a distinct lack of integrity.

2 months ago persona responded:



@orson: neither one nor another. what about <http://nanog.org> ? :-)

2 months ago Christo responded:



Once again highlights the risks around the BGP trust model. This is a good readable explanation of the problem. Well done!.

2 months ago Lars responded:



@eduard to expand a little more, this is something you will normally have to run in your border routers. The BGP table is generally constructed/refreshed from received information into the router's memory.

You might achieve a "check out" copy of the BGP in a program in some machine (you can get bgp daemons and feed them), but it'd rarely help serve anything, there is little BGP debugging afaik you could do on a pc you couldn't do on the routers themselves.

about 1 month ago Felipe responded:



Good stuff.. There has been a lot of talk and work on protecting the internet against this type of route leakage. At a recent conference there was a presentation on RPKI by Carlos Cagnazzo from LACNIC

https://www.dropbox.com/sh/4358z9df8dykn8l/Kz5NOrmf-n/carlos001_28102012-rpki...

about 1 month ago Andrea responded:



Great read. Thanks for sharing your experience.

about 1 month ago Fat Finger Jack responded:



@John Hartman In Indonesia .. self-esteem stand above all .. not only Politician have this but also judges, government .. and sadly .. engineers

..

about 1 month ago [ifbyairambulance](#) liked this post.

about 1 month ago [André Lima e Silva](#) responded:



Google, Gmail, Youtube are out right now in Brazil

<http://pic.twitter.com/r4Xf9gKA>

23 days ago Eric Kong responded:



Thank you very much for this article!

I have met the same problem. Here in China, through China-net network provider, I cannot connect to my aws us-west servers.

server domain is www.qianxun.net

here is the log:

#ping www.qianxun.net

PING misc-1244689530.us-west-1.elb.amazonaws.com

(54.241.14.167): 56 data bytes

36 bytes from 219.158.23.113: Time to live exceeded

Vr HL TOS Len ID Flg off TTL Pro cks Src Dst

4 5 00 5400 898c 0 0000 01 01 5466 192.168.213.118 54.241.14.167

Request timeout for icmp_seq 0

36 bytes from 219.158.23.113: Time to live exceeded

Vr HL TOS Len ID Flg off TTL Pro cks Src Dst

4 5 00 5400 a58e 0 0000 01 01 3864 192.168.213.118 54.241.14.167

Request timeout for icmp_seq 1

36 bytes from 219.158.23.113: Time to live exceeded

Vr HL TOS Len ID Flg off TTL Pro cks Src Dst

4 5 00 5400 a097 0 0000 01 01 3d5b 192.168.213.118 54.241.14.167

Request timeout for icmp_seq 2

36 bytes from 219.158.23.113: Time to live exceeded

Vr HL TOS Len ID Flg off TTL Pro cks Src Dst

4 5 00 5400 71bd 0 0000 01 01 6c35 192.168.213.118 54.241.14.167

Request timeout for icmp_seq 3

36 bytes from 219.158.23.113: Time to live exceeded

Vr HL TOS Len ID Flg off TTL Pro cks Src Dst

4 5 00 5400 feb0 0 0000 01 01 df41 192.168.213.118 54.241.14.167

Could you help to resolve it? I cannot find anyone that could help on this. I

read a translated article of your post and find you.

Thank you very much!

7 days ago [Aman](#) responded:



Wow, cloudflare rules the world :)

5 days ago [Ivan](#) responded:



Wow! Great article!

about 5 hours ago Bateman, Gerrie responded:



I am presently having trouble to access a web site where I take courses.

It says it is offline and for me to try a live site and retype in the address in the browser. The pop up window shows the symbol for Cloud Flare.

What should I do to get back this website?

A prompt reply would be nice so I can continue my studies.

Gerrie

Thank you

Leave a Comment

Name: Leave this field blank to
comment. Email:
 Homepage:

Want to skip this stuff?

Login with any of the following:

[Register or login to Posterous](#)



Comment: