Proto6: The Schema Was Not Supposed to Run

Proto6: The Schema Was Not Supposed to Run - Vladimir Tokarev, Cyera Research.

- Published: 2026-09-11
- Original: <https://www.cyera.com/research/proto6-the-schema-was-not-supposed-to-run>
- Preserved from: https://www.cyera.com/research/proto6-the-schema-was-not-supposed-to-run (live) on 2026-09-13
- 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.

Proto6: The Schema Was Not Supposed to Run | Cyera Research

TL;DR

Protobuf is structured data serialization format, protobuf.js is exactly that implemented in javascript - it has over 48 million weekly npm downloads. Cyera Research has found six vulnerabilities in it that let attackers achieve remote code execution or denial of service by poisoning schema-derived data.

Take the RCE chain. A node.js application accepts attacker-controlled input. That input reaches a prototype pollution gadget. Later, the same process uses protobuf.js to encode or decode a message. Because protobuf.js resolves type names through plain property lookups, a polluted Object.prototype can make an attacker-controlled string look like a valid protobuf primitive. protobuf.js then inserts that string into a generated encoder or decoder function and compiles it with Function(). The attacker gets arbitrary JavaScript execution inside the Node.js process.

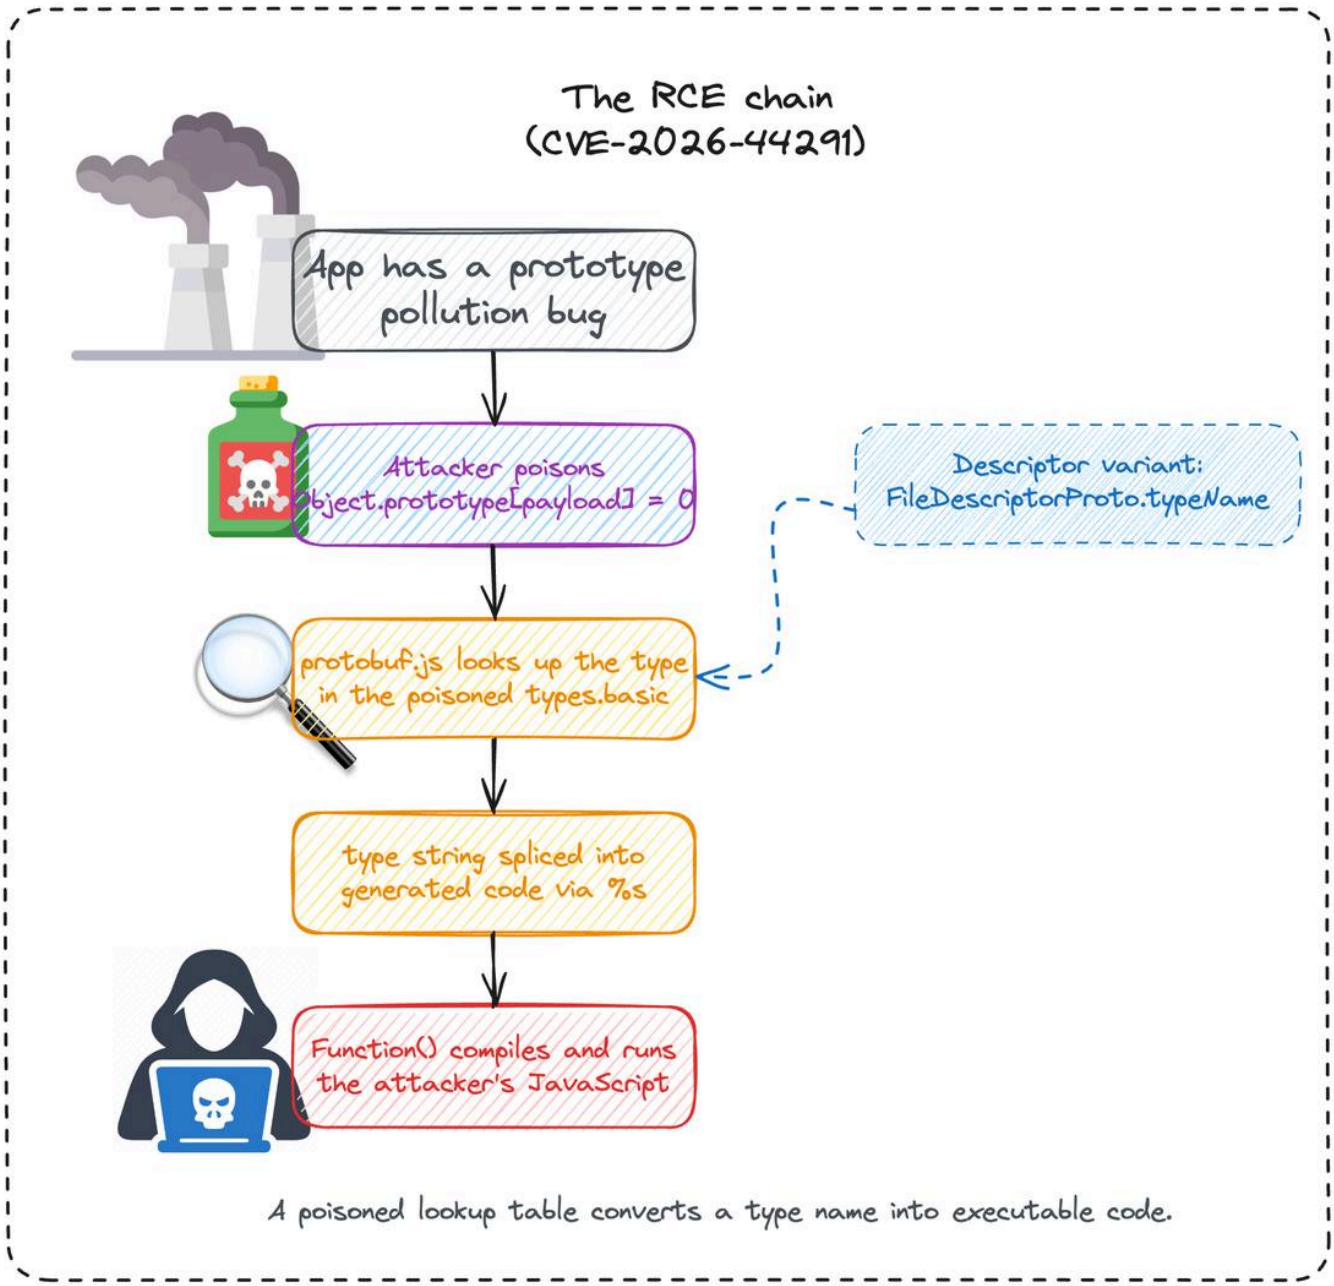

Figure 1: A poisoned prototype lets a type name pass the lookup and run as code inside Function().

A separate code injection path exists in pbjs (protobuf.js's code-generation CLI tool) static generation, with no prototype pollution prerequisite. Crafted schema names pass through to emitted JavaScript files and execute when those files are imported.

All issues were fixed in protobuf.js 7.5.6 and 8.0.2, and in protobufjs-cli 1.2.1 and 2.0.2.

CVEs assigned:

- CVE-2026-44289 (7.5 High): denial of service through unbounded protobuf recursion

- CVE-2026-44290 (7.5 High): process-wide denial of service when loading schemas with unsafe option paths
- CVE-2026-44291 (8.1 High): code generation gadget after prototype pollution
- CVE-2026-44292 (5.3 Moderate): prototype injection in generated message constructors
- CVE-2026-44294 (5.3 Moderate): denial of service from crafted field names in generated code
- CVE-2026-44295 (8.7 High): code injection in pbjs static output from crafted schema names

To summarize, the real gap is in the way protobuf.js operates. It works once, when the schema loads, it converts the result into a flat, specialized function with the type decisions already made. The JavaScript engine then compiles that function down to fast machine code, so each encode and decode does close to the minimum work. The question is why generate code at all? The alternative is to interpret the schema on every call: for each field, check its type, check whether it repeats, check whether it is a map, then handle it. That branching runs again on every message.

Crashing WhatsApp Bots and Google Cloud Functions

First, an attack you can build from these vulnerabilities - Baileys is the most popular open-source WhatsApp bot library, with over 2 million weekly npm downloads. It uses protobuf.js under the hood to decode binary messages from WhatsApp servers. The WhatsApp protocol schema (WAProto) has naturally recursive types: a Message contains an ExtendedTextMessage, which contains a ContextInfo, which contains a quoted Message which contains... well you get the picture, now recursion is part of the spec here.

We crafted a roughly 34KB binary payload containing around 10,000 nesting levels of that recursive structure. Sending it required nothing more than a WhatsApp message to a bot. When Baileys called Type.decode() on the incoming binary, protobuf.js recursed through every nesting layer until V8 ran out of stack space, which crashed the process.

The persistent part here is worse, because the malicious message sits in chat history. When the bot operator restarts the process and it catches up on missed messages, Baileys decodes the same payload again and gets the same crash again. The bot enters a restart loop that it cannot escape without manually clearing the message from the conversation.

We also confirmed CVE-2026-44292 against Baileys. By injecting **proto** keys carrying properties like isAdmin: true and canDelete: true into a message object, we could manipulate the prototype chain of individual message instances. The generated constructor copied those keys without filtering, giving a single message object properties its schema never defined.

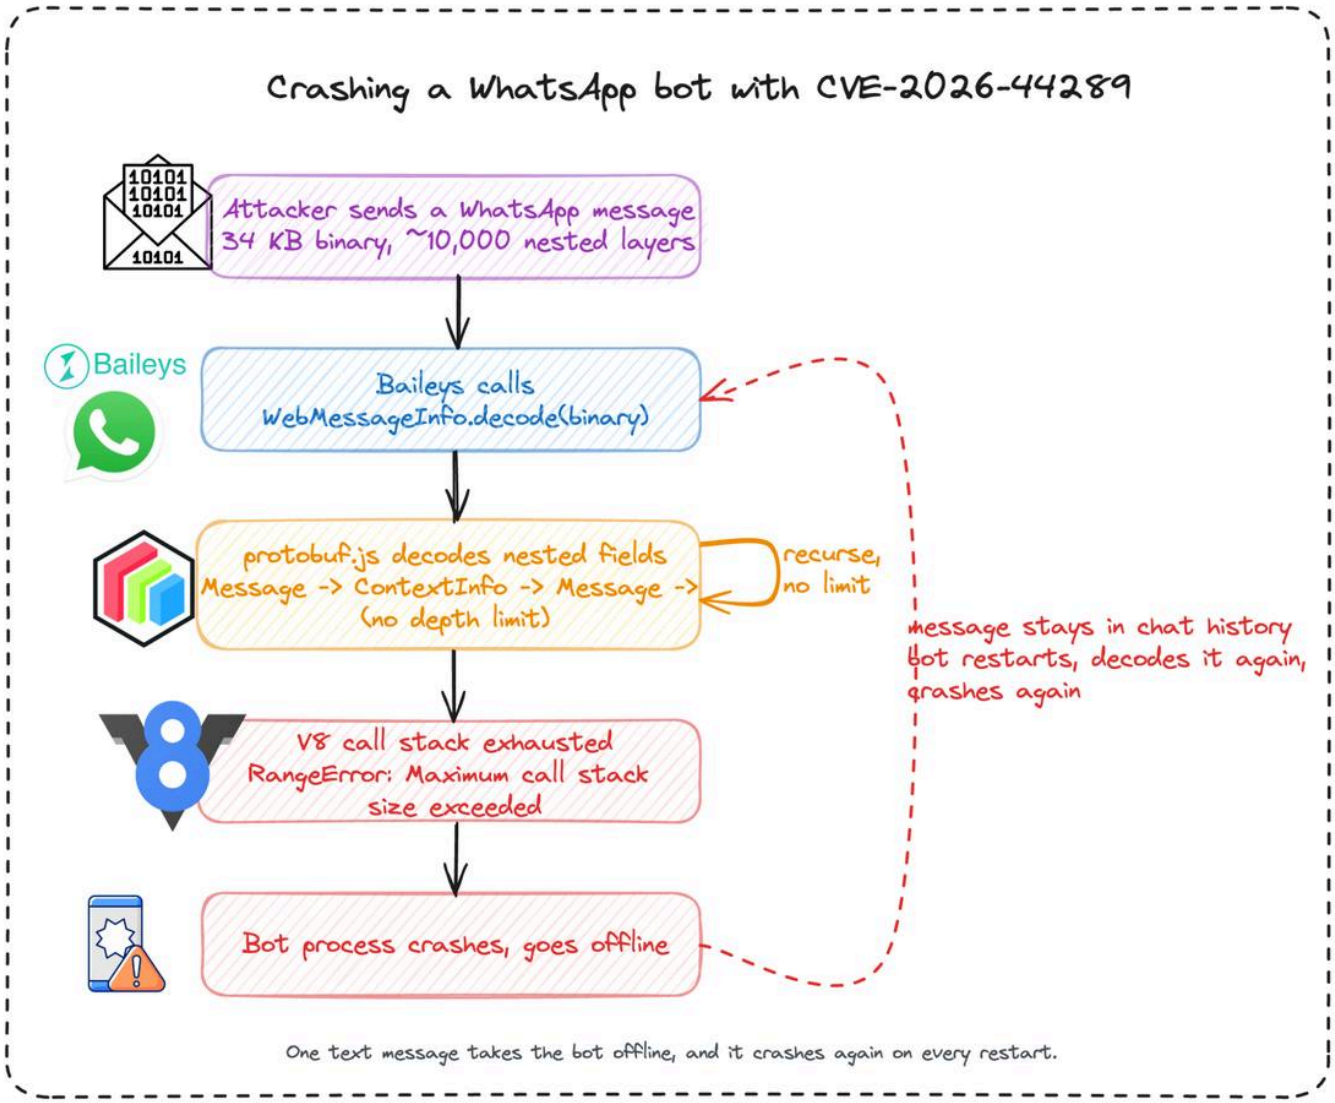

Figure 2: One nested message recurses until V8's stack overflows, and it crashes the bot again on every restart.

The same bug is not limited to WhatsApp bots. We reproduced the recursion crash against the Google Cloud Functions pattern that Google's own Firestore trigger documentation recommends: the same 34KB payload that crashed Baileys also crashes a Cloud Function decoding Pub/Sub events. Because the event is redelivered until it is acknowledged, the function crashes again on every retry until the message reaches the dead-letter queue. Want to see this in action? Watch the full PoC recording Proto6: The schema was not supposed to run.

What Are Protocol Buffers, and What Is protobuf.js?

Protocol Buffers (protobuf) let you define a schema for structured data and serialize it into a compact binary format, replacing verbose encodings like JSON with small, typed wire bytes. The format is everywhere: gRPC, internal service calls, mobile apps, message queues.

protobuf.js is the JavaScript implementation. It reads .proto files and JSON descriptors, turns them into in-memory JavaScript objects that describe each message type and field, and then generates specialized encoder and decoder functions on the fly. It also ships a CLI tool, pbjs, that can emit static JavaScript files from schemas ahead of time.

Most protobuf libraries compile schemas once at build time and ship the result. protobuf.js can do that too, but it also supports loading and compiling schemas dynamically. This changes how we should look at it. Think of it this way: protobuf.js does not just *read* schemas, it *generates custom JavaScript code from them*. If someone can tamper with the schema, the generated code changes, and so does what runs in your process.

The bugs located in that middle step, where schema metadata became code.

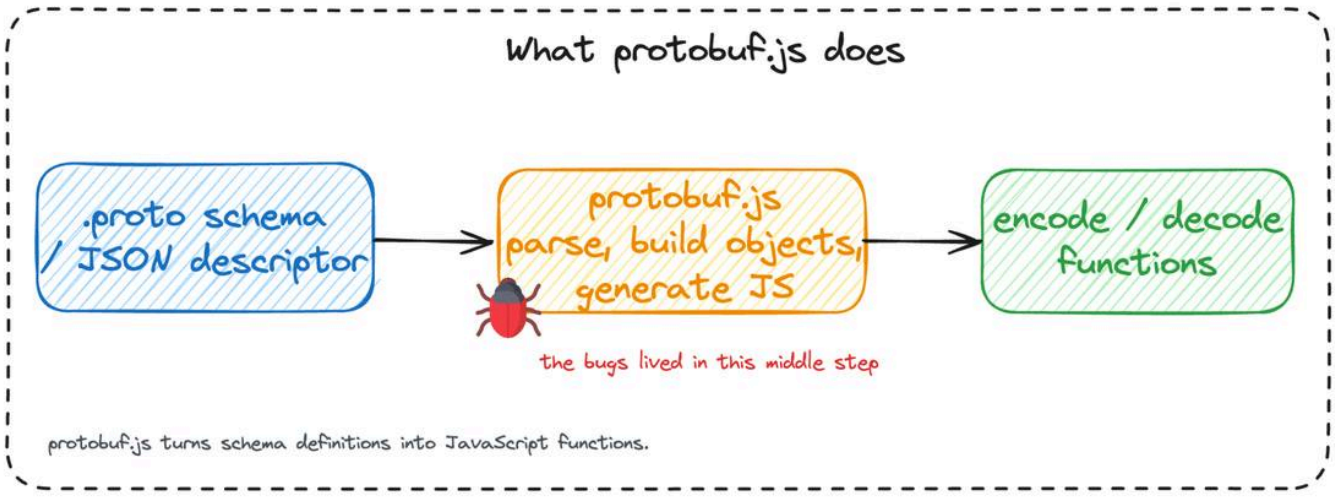

Figure 4: protobuf.js turns schema definitions into JavaScript functions. The bugs are located in that middle step, where metadata becomes code.

The Pattern We Kept Seeing

After the first finding, we started looking at every place schema data flowed into code. The same pattern kept showing up: protobuf.js treated schema values as trusted metadata, not as input that needed sanitization.

If your schemas are trusted, these bugs won't affect you. That's usually the case when your team writes the schemas itself and they aren't built from outside parameters. But protobuf.js is designed to load schemas from outside sources, and plenty of systems do exactly that. If you're in that second group, where a user can influence how a schema gets created, through params or some other control, you should keep reading and think about how to close that gap.

Once an attacker can influence a schema, "metadata" becomes input.

The six bugs came from watching where that input went:

- Type names ended up inside generated JavaScript functions, where they could run as code
- Namespace names ended up in static JavaScript files produced by the CLI tool
- Option paths walked through internal object structures and overwrote built-in functions
- Field names broke generated code by injecting raw control characters
- JSON object keys bypassed prototype safety checks in message constructors
- Nested binary messages triggered unlimited recursion that crashed the process

They all are different bugs but they have all the same smell and the schema was not supposed to run.

The RCE Chain

Start with the chain that anchored the research.

Think of a shared office phone book that every department uses. Prototype pollution is someone slipping a fake entry into that shared directory. Now when Johana from Sales looks up "the lawyer for the really sensitive stuff," instead of getting "no such contact," she finds the attacker's entry, and she calls it without a second thought, handing the details straight to the attacker. protobuf.js does the same thing: it looks up a type name, finds the attacker's planted entry instead of "not found," and trusts whatever it finds.

protobuf.js generates JavaScript functions at runtime for speed. When it loads a schema, it creates JavaScript objects that describe each message type and field. When you call encode() or decode(), protobuf.js writes the JavaScript source code for that function as a plain string, then passes it to Function() to turn it into a real, callable function. It is eval with extra steps. This makes encoding and decoding fast, but it also makes it a dangerous place for attacker-controlled strings to arrive.

Every field in a protobuf schema has a type, like uint32 or string. That type name controls which method gets written into the generated function. So we started there: where does field.type end up in the generated code? The relevant path starts in src/types.js, where protobuf.js implements lookup tables for primitive protobuf types:

The object o is created with {}. In JavaScript, that means it inherits from Object.prototype, a shared object that every plain object in the process can see. If an attacker has already added a property to Object.prototype (through a separate prototype pollution bug), then looking up a key in o can return the attacker's value instead of undefined.

protobuf.js stores its list of known primitive types (like uint32, string, bool) in objects built by this bake() function. Later, when it generates an encoder or decoder, it checks whether a field's type is in that list:

Normally, if field.type is not a valid primitive like uint32, string, or bool, wireType is undefined. protobuf.js then treats it as a message type and resolves it another way. But with prototype pollution, an attacker can make this true:

Now types.basic[payload] returns 0. Here is the important JavaScript detail: when code reads obj[key], JavaScript first checks whether key exists directly on obj. If it does not, JavaScript walks up to the object's prototype, then the prototype's prototype, and so on. Since types.basic was created as a normal object ({}), it inherits from Object.prototype. What does all that give the attacker? They injected a fake "type" into the shared prototype, and protobuf.js found it there. After pollution, the lookup behaves like this:

protobuf.js only checked whether the result was undefined. Since the polluted lookup returned 0, protobuf.js treated the attacker-controlled string as a known primitive type. This means the attacker successfully bypassed the guard. From here, the original field.type string gets pasted directly into the generated function body through a %s placeholder without any escaping and with no validation:

The %s placeholder in protobuf.js just calls String(value) and drops the result straight into the code. It does not escape anything, does not check whether the value is a safe identifier, and does not strip dangerous characters. Whatever string the attacker put in field.type lands in the generated function exactly as it is.

The generated code becomes something like:

That string is compiled with Function(), and at that point, the type name stops being a type name. It runs as JavaScript.

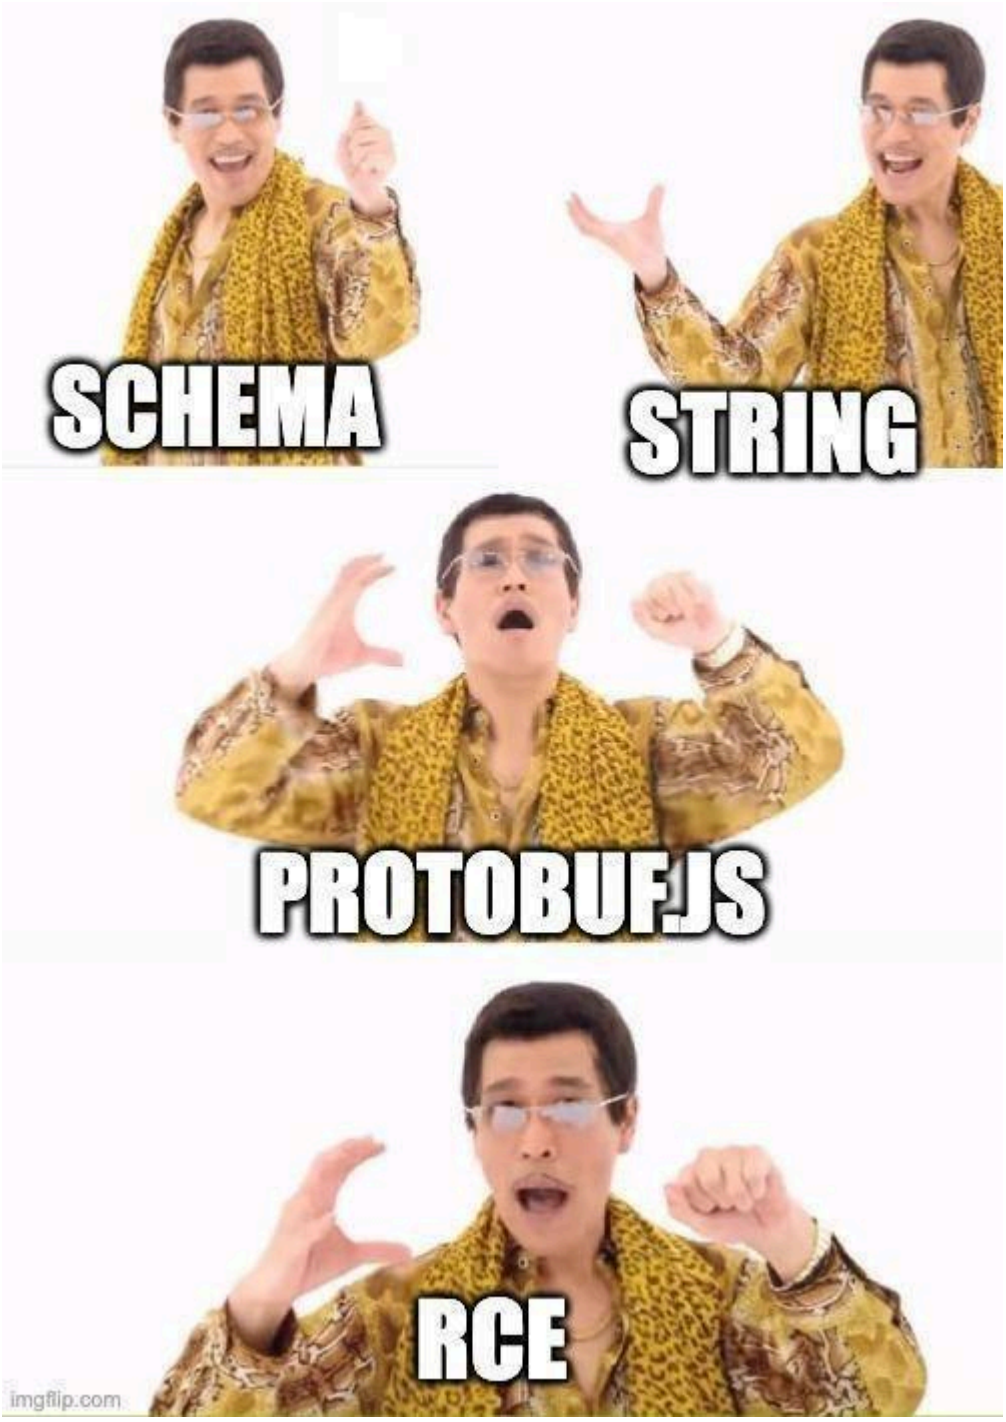

The Prerequisite Matters

This was not a standalone prototype pollution bug in protobuf.js. The RCE chain requires a separate pollution primitive in the application or one of its dependencies. protobuf.js provided the code generation gadget. Another bug has to provide the polluted prototype.

But in Node.js, prototype pollution has a long history. Libraries that merge JSON, parse query strings, flatten objects, or apply nested configuration paths have all had this class of issue. So the question becomes:

What happens if protobuf.js runs after pollution has already happened?

In affected versions, the answer was: a polluted type lookup could lead to code execution.

The same root cause also showed up through descriptors.

ext/descriptor converts protobuf descriptors into reflection objects. Field.fromDescriptor() copied descriptor.typeName directly into field.type:

So an attacker-controlled FileDescriptorProto could carry the payload in typeName. With the same prototype pollution prerequisite, that descriptor string reached the same types.basic lookup and the same %s codegen sink. Meaning that the path changed, but the sink did not.

A Second RCE Path: pbjs Static Code Generation

Imagine a photocopy machine that copies blueprints. You feed in a building plan, it prints a copy. But if someone hides extra text in the blueprint's title block, the copier faithfully reproduces it. Anyone who reads the copy follows those instructions without realizing they were not part of the original design.

The runtime codegen chain needed prototype pollution, but the pbjs static generation bug did not - what is pbjs ? It is the CLI tool that turns protobuf schemas into JavaScript files. A common flow looks like this:

Then the generated output.js is imported by the application or used in tests. Then the generated output.js is imported by the application or used in tests.

The vulnerable code is inside the cli/targets/static.js file. This static generator has a problematic escapeName() helper:

Despite the name, this did not really escape schema names. It only handled JavaScript reserved words. What does this lead to? Namespace, enum, service, and method names could contain newline characters, semicolons, quotes, and other syntax-significant characters. Some of those names were emitted through direct string concatenation:

A malicious JSON schema could define a namespace name like:

The generated JavaScript now contains the injected statement as real code. When the developer imports the generated file, the payload runs. In this case there is no second vulnerability required. The attacker-controlled input is the schema. The output is a JavaScript file. The trigger is importing that file. In any CI/CD pipeline where schema files arrive from external contributors (pull requests, shared registries, vendored dependencies), this turns a .proto or .json commit into arbitrary code execution on the build server, with access to secrets, deploy keys, and downstream artifacts.

The Other Four Bugs

The two bugs that lead to RCE are maybe the most dangerous of the six, but there are 4 more that are worth covering. All four are the same kind of problem we already saw: something that was supposed to be passive data (a name, an option, a payload) got treated as if it were safe/trusted.

util.safeProp() builds JavaScript property access strings for field names, choosing between dot notation and bracket notation depending on the name. It escaped backslashes and double quotes, which blocked the obvious injection routes. But it missed control characters like newlines (\n), carriage returns (\r), and null bytes (\0).

For example, a field named user\nname would produce generated code like m["user on one line and name"] on the next. That is not valid JavaScript. Function() rejected it with a SyntaxError. Because the quote escaping still held, this never became code execution. It became something more mundane and still painful: every call to encode, decode, or verify on that message type would fail, permanently, for any application that loaded the schema. Denial of service from a field name (CVE-2026-44294).

What happens when Object.keys stops working? Everything that calls it throws. util.setProperty() applied dot-separated option paths to objects and blocked **proto** and prototype as path segments, but it did not block constructor. On a plain object, constructor points to the global Object function. So a schema option with the path constructor.keys and value "corrupted" would walk from the options object {} to its .constructor (which is the global Object function), then set Object.keys = "corrupted". From that point on, every call to Object.keys() anywhere in the process throws a TypeError because it is no longer a function. This was reachable through setOption() and setParsedOption(), meaning any schema that set a custom option with a constructor-prefixed path could corrupt process-wide built-in functions. Not a local failure. A single malicious option path could break unrelated application code for the lifetime of the process (CVE-2026-44290).

The base Message constructor in protobuf.js already knew to skip **proto** when copying properties from input objects. The generated constructor for schema-defined types did not. When input arrives through JSON.parse(), **proto** becomes an own enumerable key on the resulting object, and the generated copy loop assigns it directly to this["**proto**"], changing the prototype of that specific message instance. This was not global pollution. It was per-instance prototype injection: enough to break code relying on prototype methods, inherited properties, or instanceof checks for that one object, but not enough to poison every object in the process (CVE-2026-44292).

Generated decoders for nested message fields called types[i].decode(r, r.uint32()) with no depth parameter, no counter, no maximum. A self-referencing message type is a perfectly valid protobuf. We built a roughly 15KB binary payload with thousands of nesting layers, and the decoder recursed until V8 threw RangeError: Maximum call stack size exceeded. Think of it like Russian nesting dolls, except there is no smallest doll. The decoder keeps opening dolls until it runs out of

desk space and the whole table collapses. This one becomes especially dangerous in the real-world chains below (CVE-2026-44289).

Real-World Impact

Now finding bugs in a library is one thing but it's a lot more interesting to see what one can actually achieve with those bugs. We present three exploit chains against systems people actually run in production.

Poisoning CI/CD Pipelines Through .proto Files

Many projects run pbjs --target static-module as a build step, converting .proto schemas into JavaScript files that get committed or consumed by tests. In open-source projects, those schemas often arrive through pull requests.

Here we will show an attack chain against this pattern. An attacker forks a project, submits a pull request containing a .proto file with a crafted namespace name, and waits. When CI runs the build step, pbjs generates an output.js file with injected JavaScript embedded in the namespace declaration. The injected code survives the beautify step. It executes the moment anything calls require('./generated.js'), whether that is the CI build itself, the test suite, or a downstream consumer importing the generated module.

The injection works with namespace names, enum names, and service names. Any schema element whose name reaches string concatenation in cli/targets/static.js is a viable carrier. The attack surface is every open-source project that accepts .proto contributions and runs pbjs in CI (CVE-2026-44295).

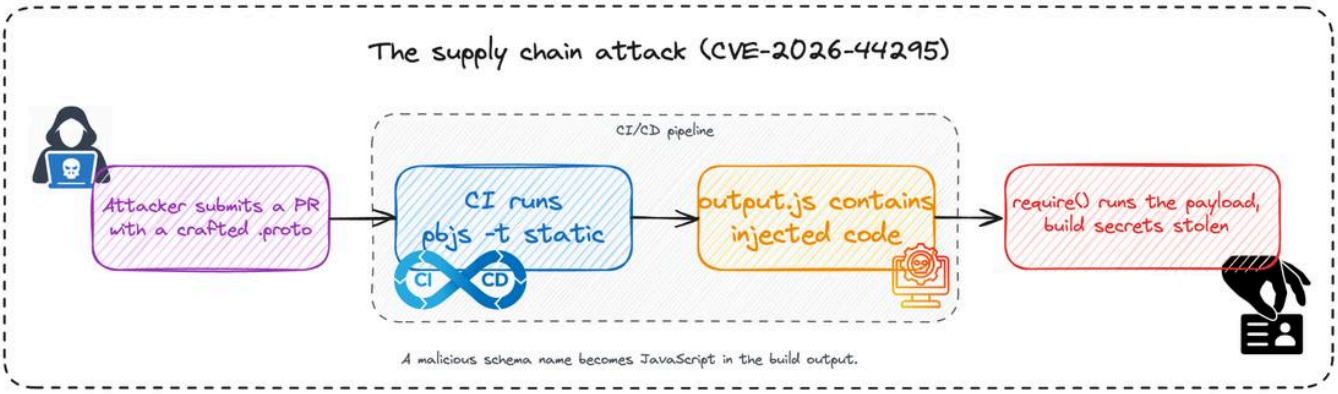

Figure 5: The supply chain attack. A crafted schema name in a pull request becomes JavaScript in the generated build output, and runs the moment CI imports it, exposing build secrets.

Because the injected payload executes during trusted build and testing workflows, an attacker may gain access to CI/CD runners, build secrets, signing credentials, cloud tokens, package registry credentials, and internal repositories. From there, the compromise can extend beyond the initial project, enabling attackers to tamper with release artifacts, poison downstream dependencies, and establish persistence within the software delivery pipeline itself.

When Are You Exposed?

Not every protobuf.js deployment was vulnerable to all six bugs. Exposure depended on which input surfaces the application used:

- **Schema from external sources.** Loading .proto files or JSON descriptors from users, plugins, tenants, or shared repositories exposed the field name DoS, option path corruption, and static codegen injection paths.
- **Binary messages from untrusted senders.** Decoding protobuf binary data with recursive message types exposed the stack overflow DoS. This includes gRPC services, message queues, and any endpoint accepting raw protobuf.
- **Prototype pollution elsewhere in the process.** If any dependency or application code had a separate prototype pollution bug, protobuf.js provided the gadget that turned it into code execution.
- **pbjs in CI on contributed schemas.** Running static code generation on pull-request content exposed the build pipeline injection path.

Applications that only used trusted schemas compiled at build time and only decoded flat, non-recursive messages from authenticated sources had minimal exposure. But protobuf.js is popular precisely because it is flexible, and the risky cases are the flexible ones.

The bigger lesson is not "all schemas are malicious." A library that turns schema metadata into code needs to treat that metadata as code-adjacent input, with the same validation, escaping, and distrust that any code generation boundary demands.

Fixes and Takeaways

Every vulnerability traced back to the same pattern: schema-controlled strings flowing into runtime structures without validation. The fixes are straightforward.

- **All six issues are fixed**. Upgrade protobuf.js to 7.5.6 or 8.0.2, and protobufjs-cli to 1.2.1 or 2.0.2. The mitigation work is tracked in PR #2163.
- **Type lookup tables built with Object.create(null)** eliminates inherited property access during runtime resolution, and validates typeName before interpolation.
- **All reflection names sanitized in static codegen**, not just Type.name but fields, namespaces, and services. String concatenation replaced with safe templating.
- **constructor blocked alongside *__proto__* and prototype** in option-path traversal.
- **Generated constructors match the base Message constructor**, **proto** keys skipped during property copy.
- **Recursive decoding enforces a configurable max nesting depth**, prevents stack exhaustion from crafted payloads.

None of these are exotic fixes. Most are small guardrails around places where data becomes structure, and structure becomes code.

Full details for each finding: CVE-2026-44289, CVE-2026-44290, CVE-2026-44291, CVE-2026-44292, CVE-2026-44294, CVE-2026-44295. Mitigation tracked in PR #2163.

Disclosure Timeline

We reported the advisories through GitHub Security Advisories, and the protobuf.js maintainers handled them promptly.

- Reports were opened in April 2026.
- The maintainers accepted the reports and linked the mitigation work to PR #2163.
- Fixed versions were released as protobuf.js 7.5.6 and 8.0.2.
- protobufjs-cli static generation fixes were released as 1.2.1 and 2.0.2.
- GitHub assigned the CVEs in May 2026.

Closing

The cool part about this research is that the same trust boundary kept showing up in different clothes. A type name became a method call. A namespace became a JavaScript statement. An option path walked into Object.keys. A field name broke a generated function. A binary payload became a recursive stack. What do all of those have in common? They started as metadata.

And metadata is fine, right up until the library starts running it.

September 1, 2026 ### PostGREShell: The database powering much of the internet had an open door for 12 years

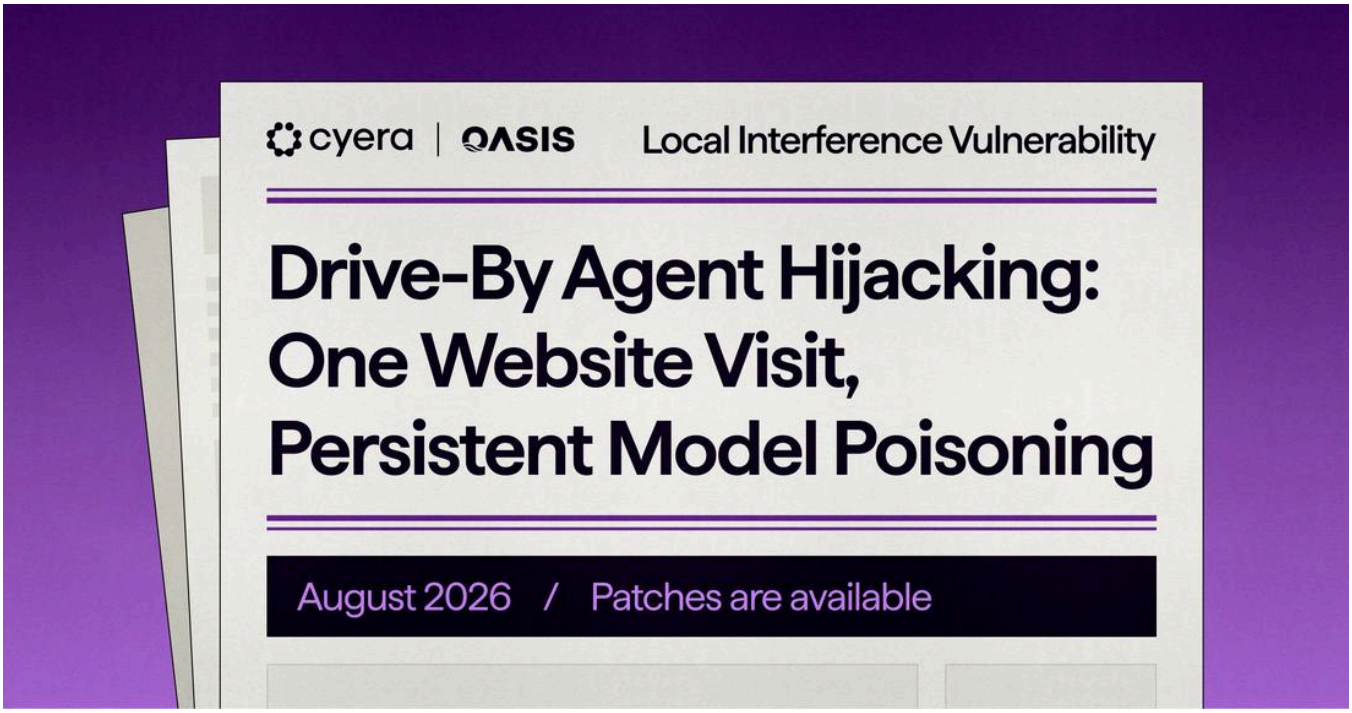

August 25, 2026 ### Drive-By Agent Hijacking: One Website Visit, Persistent Model Poisoning

[

[image: Cyera research banner with stylized purple icons representing AI concepts and a bridge over an abyss, alongside the text 'The Hidden Attack Surface of Agentic AI: Securing AI Agent Integration Platforms'.].png)

August 13, 2026

The Hidden Attack Surface of Agentic AI: Securing AI Agent Integration Platforms