



**black hat**<sup>®</sup>  
EUROPE 2023

DECEMBER 4-7  

---

EXCEL LONDON / UK



# HODOR: Reducing Attack Surface on Node.js via System Call Limitation

Speakers: Wenya Wang, Xingwei Lin

Contributors: Wenya Wang, Xingwei Lin, Jingyi Wang, Wang Gao, Dawu GuM



上海交通大學  
SHANGHAI JIAO TONG UNIVERSITY



蚂蚁集团  
ANT GROUP



浙江大學  
ZHEJIANG UNIVERSITY

#BHEU @BlackHatEvents




# Agenda

- [Introduction](#)
- Previous work & Remaining challenges
- **HODOR: system call level protection system for Node.js applications**
- Evaluation
- Conclusion & Takeaways

# Node.js

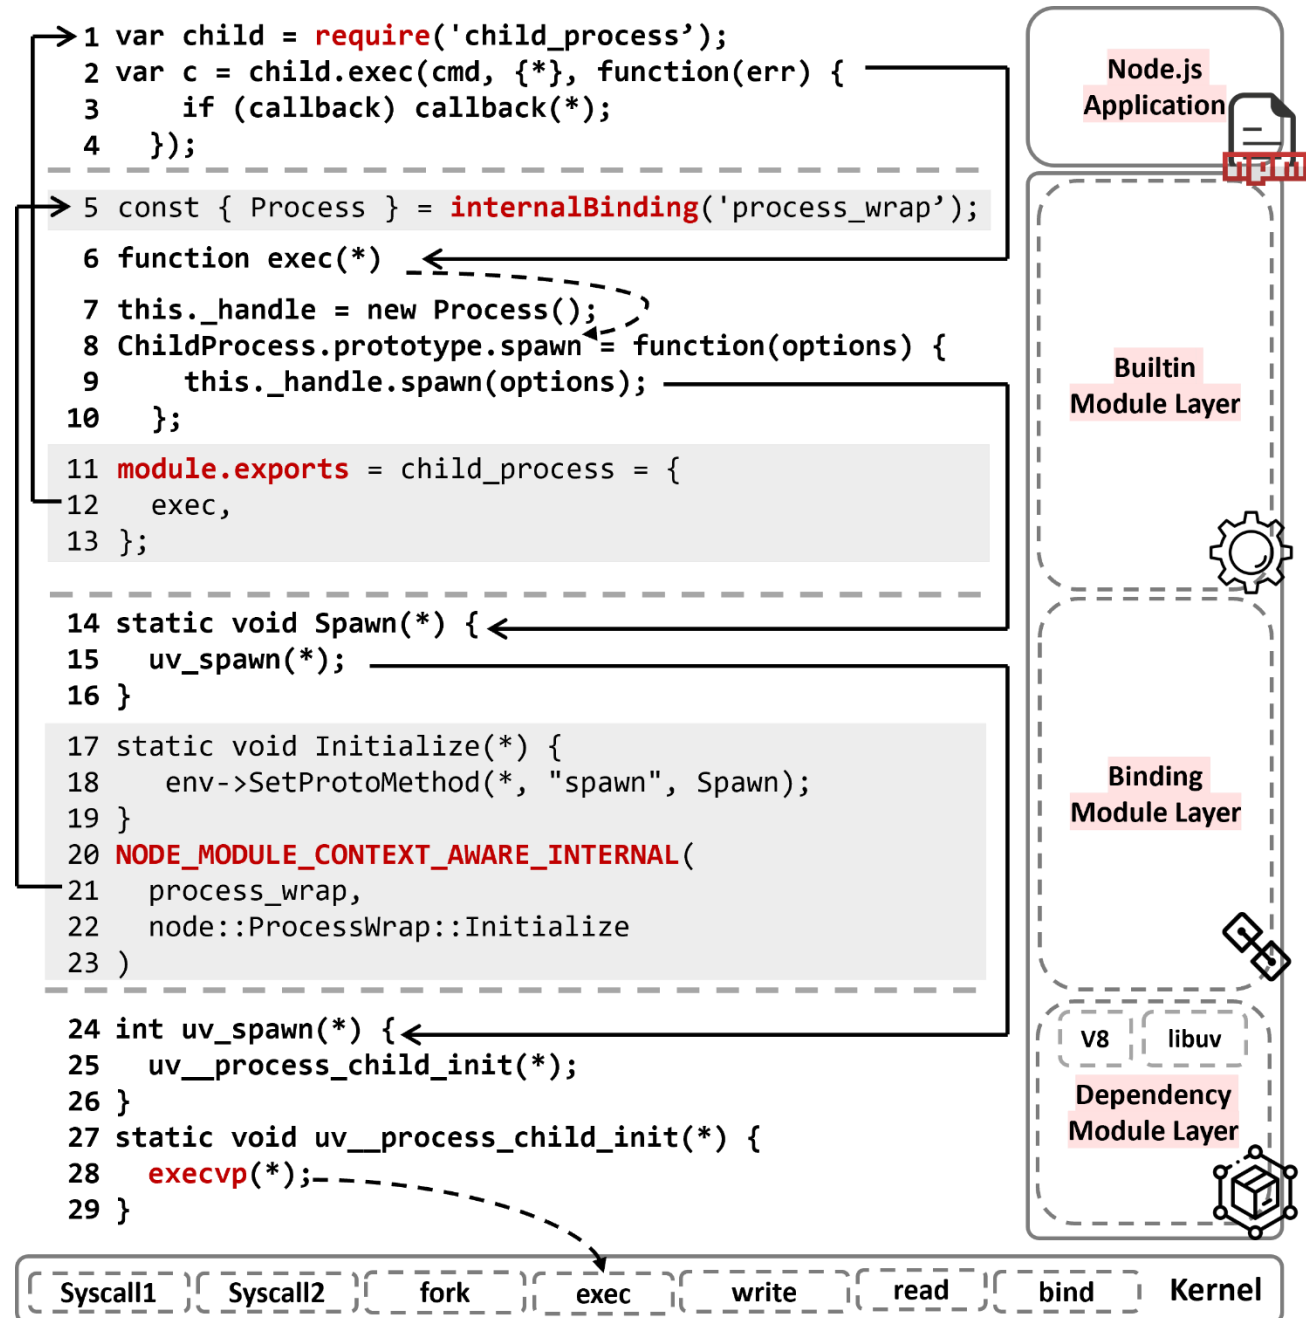
Node.js is an **open-source, cross-platform JavaScript runtime** environment.

- ✓ Asynchronous and Event-Driven
- ✓ Single-Threaded
- ✓ Cross-Platform
- ✓ NPM (Node Package Manager)
- ✓ JavaScript Everywhere



# Node.js architecture

- ✓ Node.js Applications (JS)
- ✓ Built-in Module Layer (JS)
- ✓ Binding Module Layer (C++)
- ✓ Dependency Module Layer (C)



# Motivation

- NPM is a package manager with over 1 million packages → The key to the success of Node.js
- 19.63% of packages in the NPM ecosystem depend on vulnerable packages, such as gadget chain attacks, inject-related attacks, and supply chain attacks. → Most of them may lead to ACE attacks.

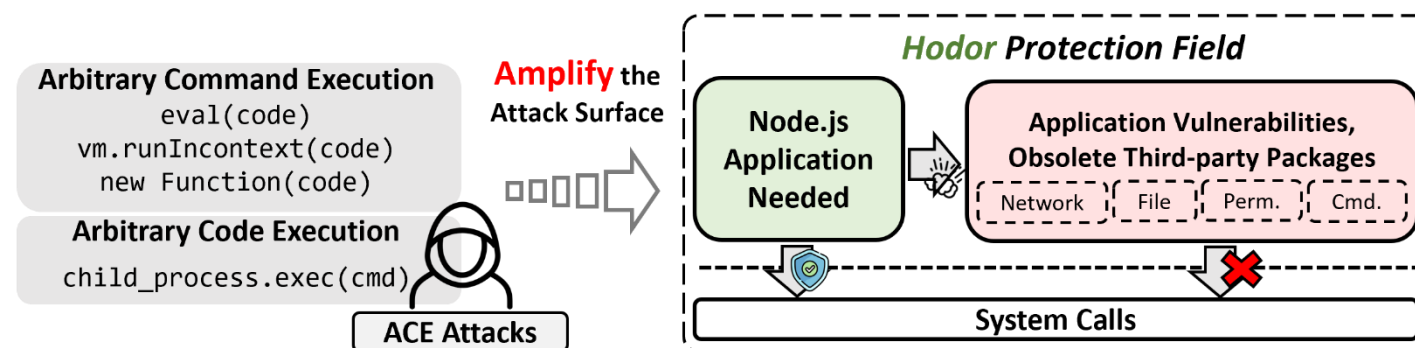
- Arbitrary Command/Code Execution: the attackers can perform arbitrary dangerous critical operations

- mail `cat / etc /passwd`
- mail `nc -l -e /bin/bash 8001`
- mail `su root`
- ...

```

1 // main.js
2 var growl = require("growl");
3 var message = 'You have mail!';
4 growl(message);
5
6 // ./lib/growl.js
7 exports = module.exports = growl; ←
8
9 var exec = require('child_process').exec
10 cmd = { pkg: "notify-send" };
11
12 function growl(msg, *) {
13   args = [cmd.pkg];
14   args.push(quote(msg)); ←
15   exec(args.join(' '),...);
16 };
  
```

**Growl Application (v1.8.0)**





# Agenda

- Introduction
- *Previous work & Remaining challenges*
- HODOR: system call level protection system for Node.js applications
- Evaluation
- Conclusion & Takeaways



# How to reduce the attack surface of ACE attacks for Node.js applications?





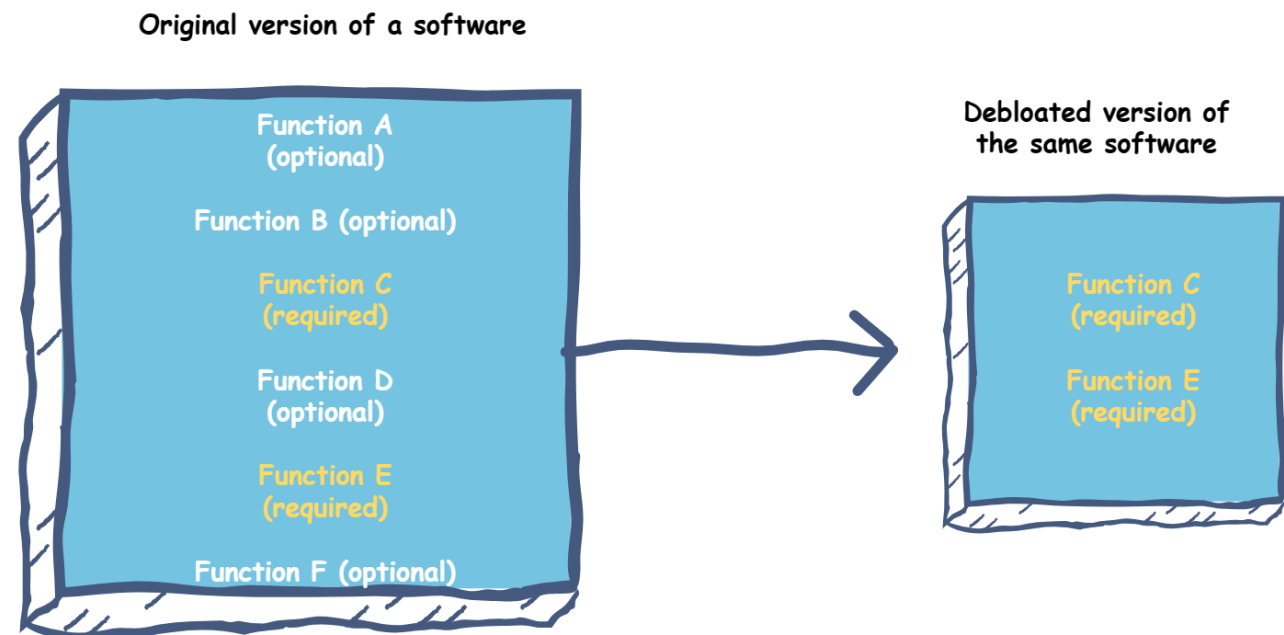
# How to reduce the attack surface of ACE attacks for Node.js applications?

## Threat Model

- ✓ Consider an attacker with *ACE ability*
- ✓ Not considered: preventing ACE, code vulnerabilities in binding layer/dependency layer, race condition, DOS attack, etc

# Existing Works: Software Debloating

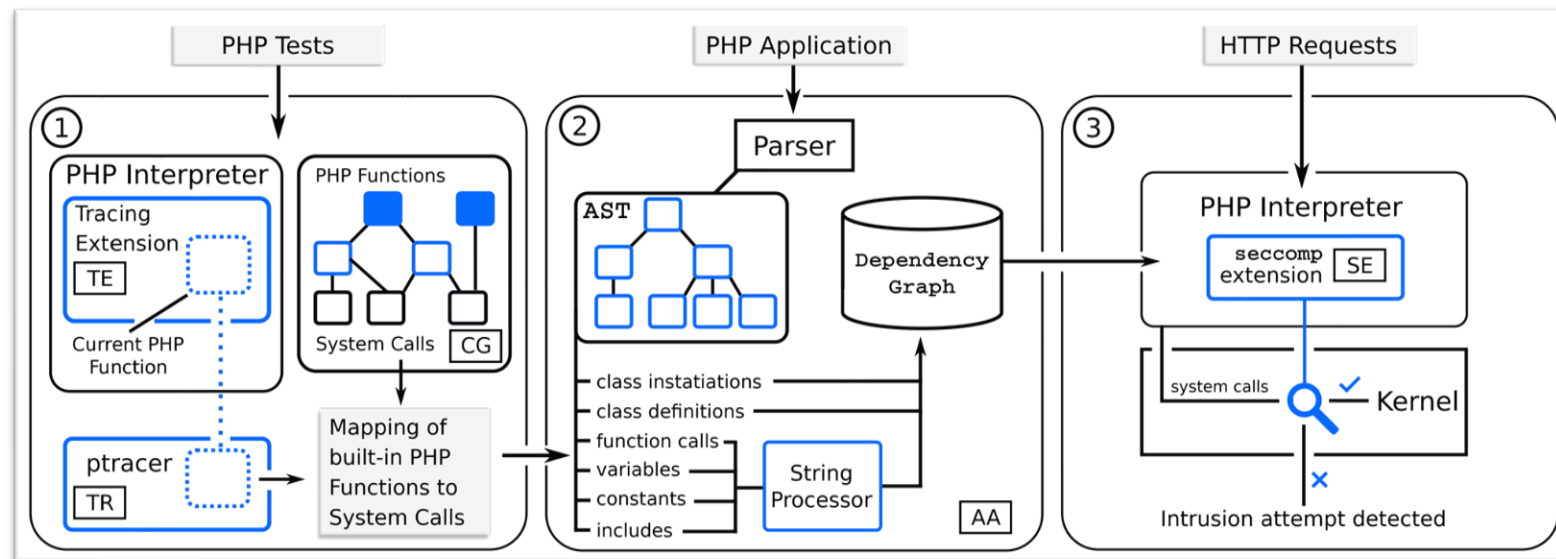
- Use program analysis to cut the useless code
  - ✓ (USENIX Sec'19) RAZOR: A Framework for Post-deployment Software Debloating
  - ✓ (USENIX Sec'19) Less is More: Quantifying the Security Benefits of Debloating Web Applications
  - ✓ (Useenix Sec'20) Slimium: Debloating the Chromium Browser with Feature Subsetting
  - ✓ (RAID'20) Mininode: Reducing the Attack Surface of Node.js Application



[What is software debloating? \(educative.io\)](https://www.educative.io/)

# Existing Works: System Call Limitation

- Restrict the system calls that can be used by the application
  - ✓ (USENIX Sec'20) Temporal System Call Specialization for Attack Surface Reduction
  - ✓ (RAID'20) Confine: Automated System Call Policy Generation for Container Attack Surface Reduction
  - ✓ (RAID'20) sysfilter: Automated System Call Filtering for Commodity Software
  - ✓ (PLDI'20) BlankIt Library Debloating Getting What You Want Instead of Cutting What You Don't
  - ✓ (USENIX Sec'21) Sapphire: Sandboxing PHP Applications with Tailored System Call Allowlists



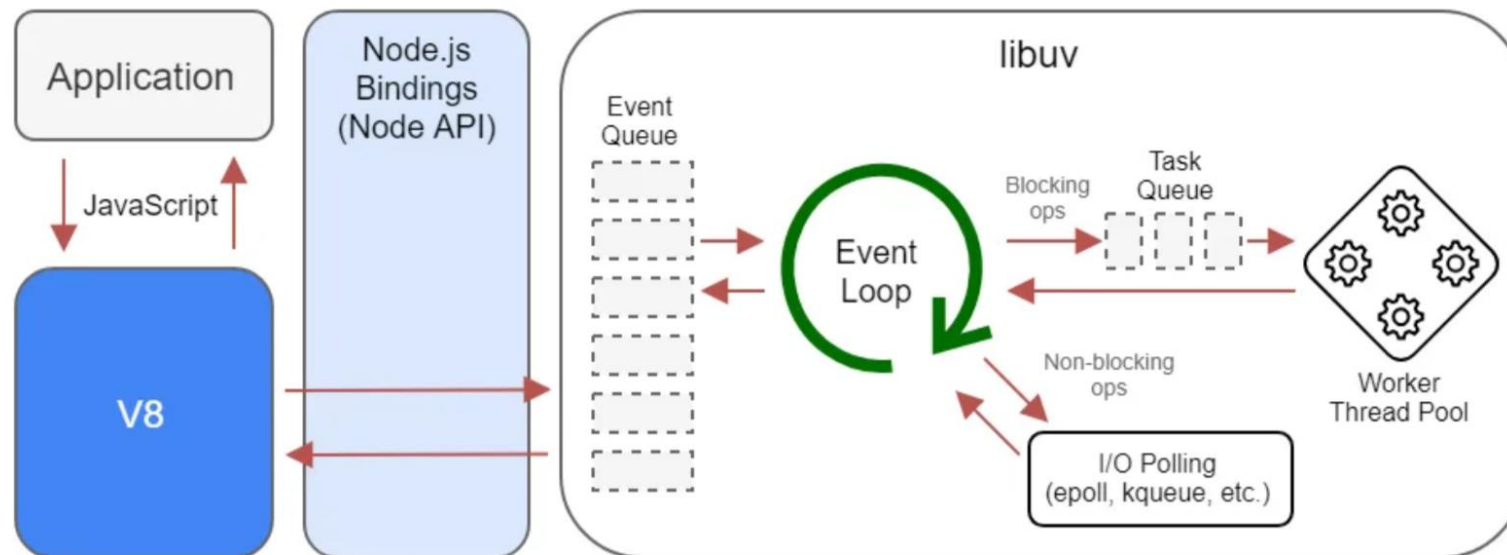
# Remaining Challenges

## 1. Cross-language mapping requirement

✓ JS code layer & C/C++ code layer

## 2. Integration with Node.js framework

✓ Node.js runs in a single process that creates two kinds of threads.



<https://medium.com/preezma/node-js-event-loop-architecture-go-deeper-node-core-c96b4cec7aa4>

# Problem Formulation

- The number of all system calls provided by the system:

$$S_{base} = |SYSCALL_{system}|$$

- The number of system calls in the whitelist:

$$S_{app} = |SYSCALL_{main-thread}| \cup |SYSCALL_{thread-pool}|$$

- The degree of attack surface reduction in the system call level:

$$SR = \frac{S_{app}}{S_{base}}$$

**Goal:** *minimize the attack surface in the system call level to prevent malicious critical operations, while not affecting the application's normal execution*

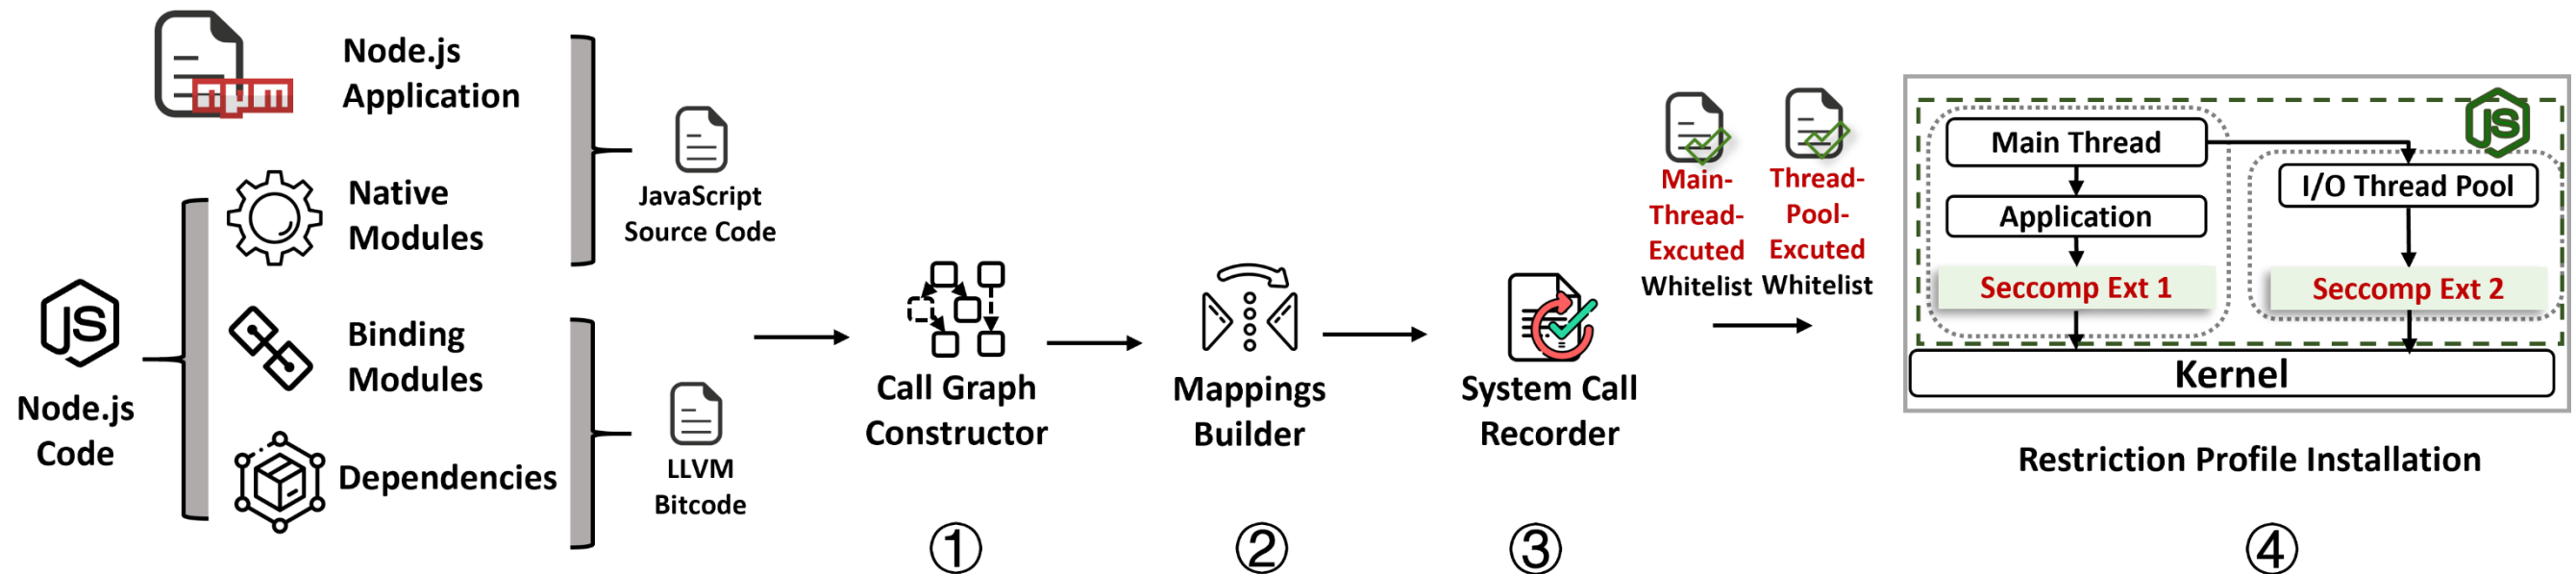


# Agenda

- Introduction
- Previous work & Remaining challenges
- *[HODOR: system call level protection system for Node.js applications](#)*
- Evaluation
- Conclusion & Takeaways

# Our approach: Hodor

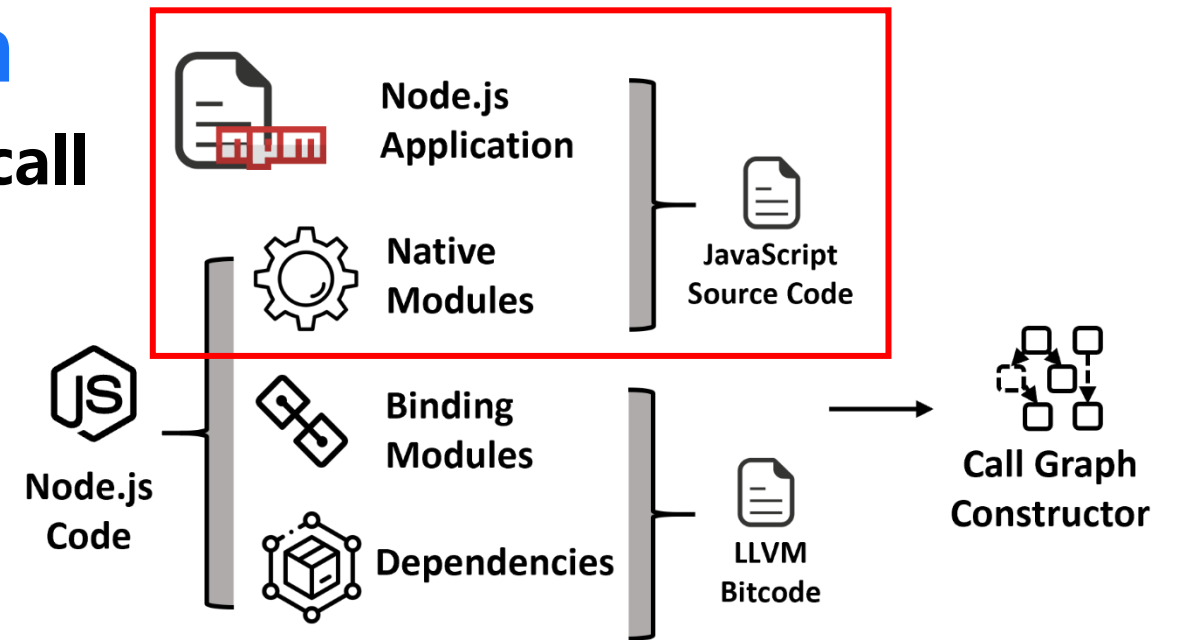
A lightweight runtime protection system.



# Step 1: Call Graph Construction

- **JS** > complement missing nodes/edges/syscall
  - ✓ Code features of built-in methods


- `[1,2,3].map(x => x * 2);`
- `fs.readFile(filename, CallbackFunc);`
- ✓ **Dynamic Analysis Refiner**
  - `let sum = new Function('a', 'b', 'return a+b');`
  - `eval("sum()");`
- ✓ **Dynamic Command Execution**
  - `child_process.exec("touch new file");`





# Step 1: Call Graph Construction

- **JS** > complement missing nodes/edges/syscall
  - ✓ **Code features of built-in methods**
    - `[1,2,3].map(x => x * 2);`
    - `fs.readFile(filename, CallbackFunc);`
  - ✓ **Dynamic Analysis Refiner**
    - `let sum = new Function('a', 'b', 'return a+b');`
    - `eval("sum()");`
  - ✓ **Dynamic Command Execution**
    - `child_process.exec("touch new file");`



# Step 1: Call Graph Construction

➤ **JS** > complement missing nodes/edges/syscall

✓ **Code features of built-in methods**

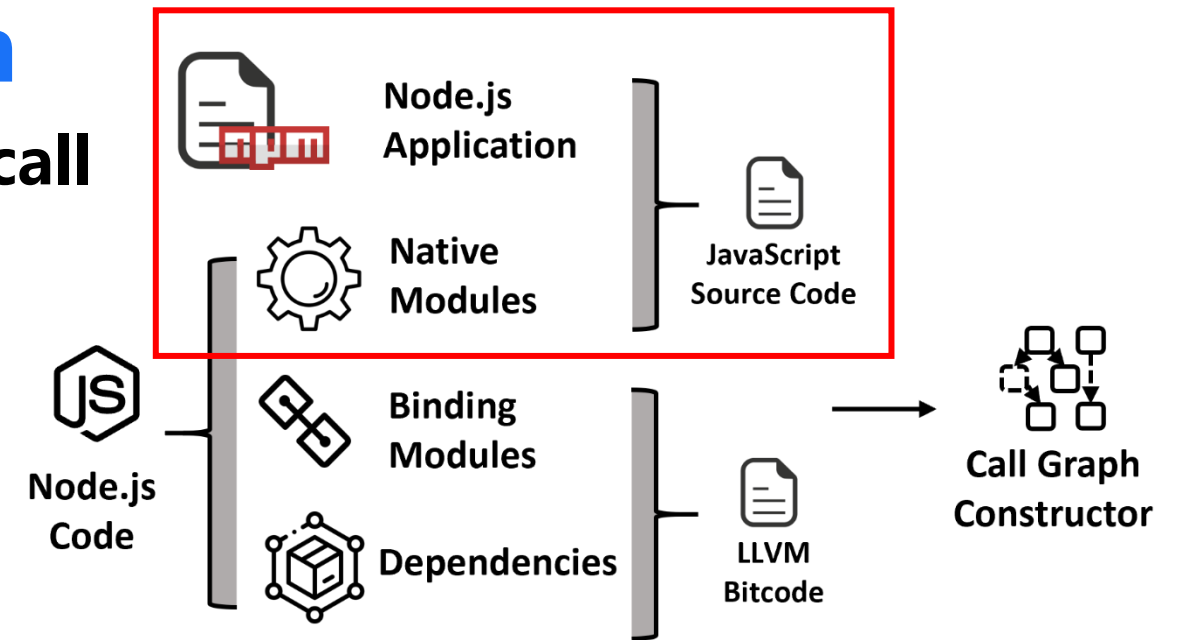
- `[1,2,3].map(x => x * 2);`
- `fs.readFile(filename, callbackFunc);`

✓ **Dynamic Analysis Refiner**

- `let sum = new Function('a', 'b', 'return a+b');`
- `eval("sum()");`

✓ **Dynamic Command Execution**

- `child_process.exec("touch new file");` ➔ **System Call Required**



# Step 1: Call Graph Construction

## ➤ JS > complement missing nodes/edges/syscall

### ✓ Code features of built-in methods

- `[1,2,3].map(x => x * 2);`
- `fs.readFile(filename, callbackFunc);`

### ✓ Dynamic Analysis Refiner

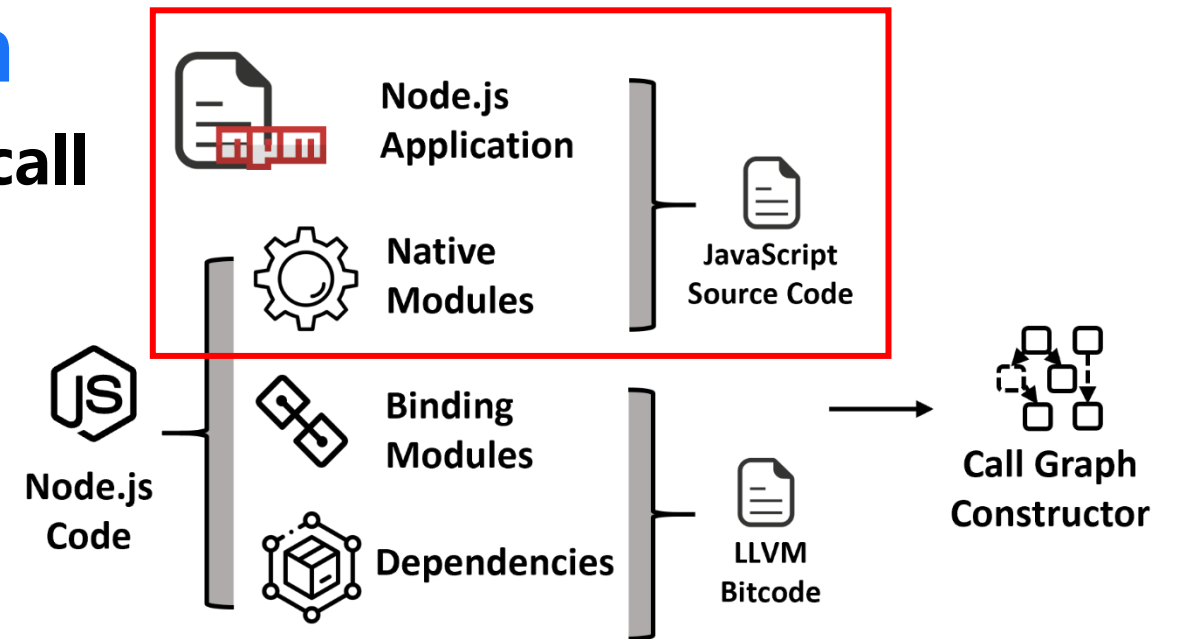
- `let sum = new Function('a', 'b', 'return a+b');`
- `eval("sum()");`

### ✓ Dynamic Command Execution

- `child_process.exec("touch new file");`

## ➤ Implementation

- ✓ Reimplement **JAM** and add in proposed optimizations
  - ISSTA'21 Modular call graph construction for security scanning of node.js applications
- ✓ Combine dynamic call graph tool **Nodeprof** and Linux **strace** utility



# Step 1: Call Graph Construction

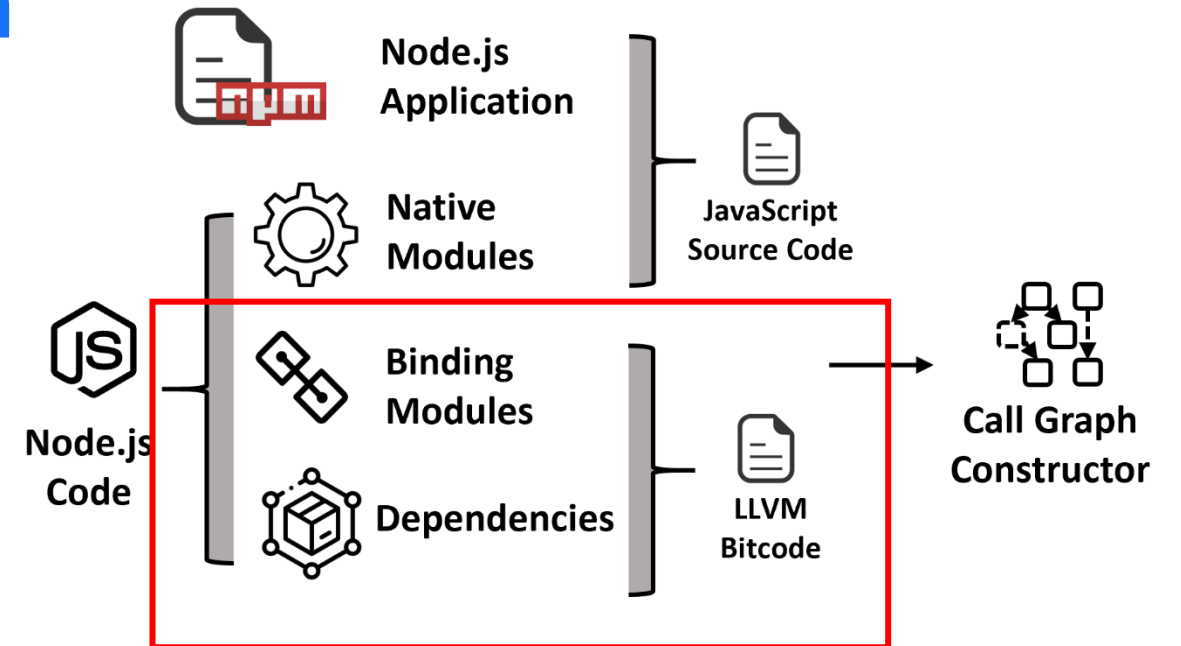
- C/C++ call graph > eliminate non-existing nodes/edges

- ✓ *Partial context-aware* analysis for **switch-case statement** & function pointer parameter

```

1 #define INIT(subtype) req->fs_type = UV_FS_ ## subtype;
2
3 int uv_fs_access(*) {
4     INIT(ACCESS);
5     POST;
6 }
7 int uv_fs_write(*) {
8     INIT(WRITE);
9     POST;
10 }
11
12 #define POST uv__fs_work(&req->work_req);
13
14 static void uv_fs_work(struct uv_work* w) { uv_fs_work.uv_fs_access
15     req = container_of(w, uv_fs_t, work_req);
16     switch (req->fs_type) {
17         X(ACCESS, access(req->path, req->flags));
18         X(WRITE, uv_fs_write_all(req));
19         ...
20     default: abort();
21     }
22 }

```



# Step 1: Call Graph Construction

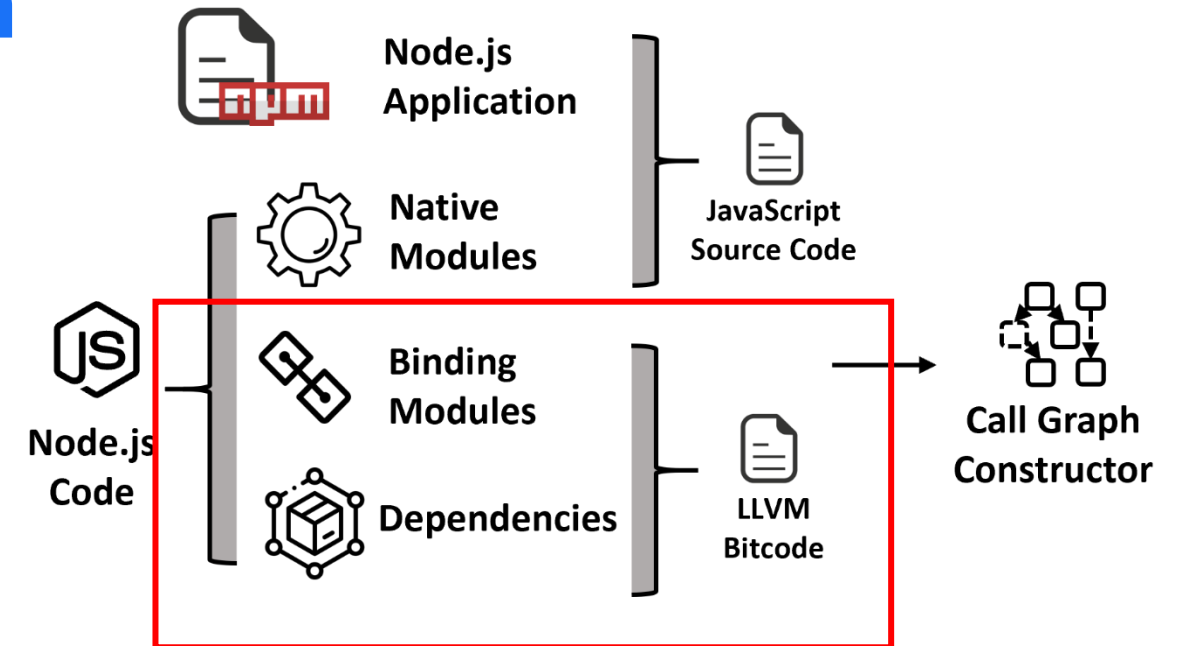
- C/C++ call graph > eliminate non-existing nodes/edges

✓ *Partial context-aware* analysis for switch-case statement & **function pointer parameter**

```

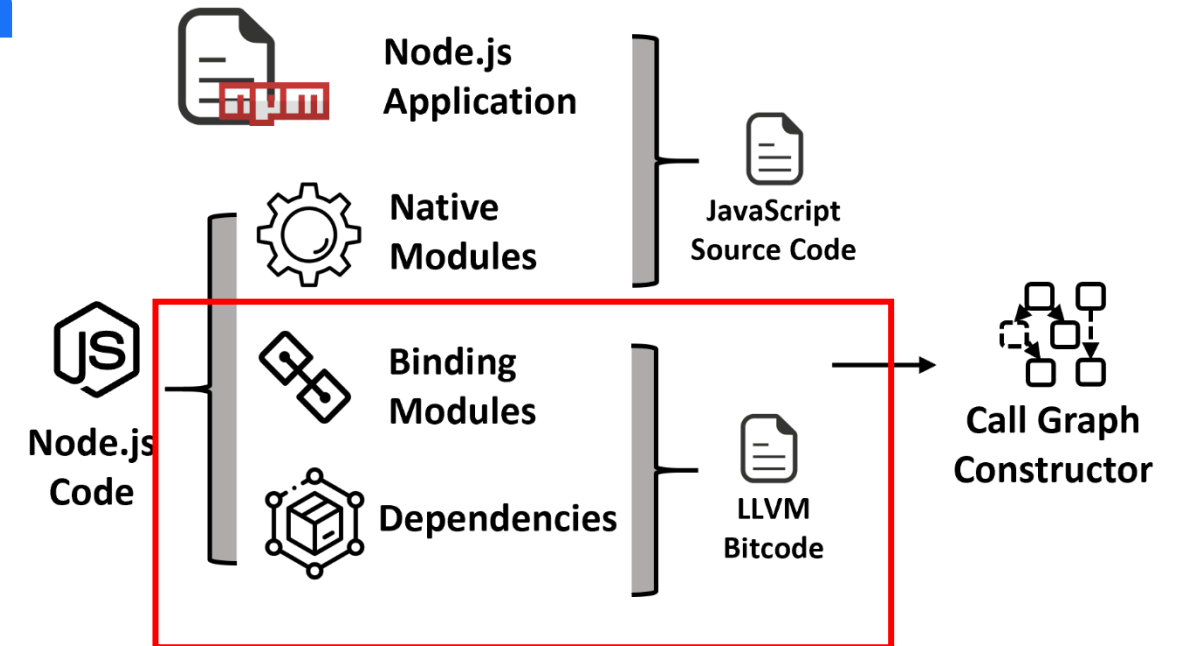
1 static void Read(const FunctionCallbackInfo<Value>& args) {
2   AsyncCall(..., uv_fs_read, ...);
3 }
4 static void Unlink(const FunctionCallbackInfo<Value>& args) {
5   AsyncCall(..., uv_fs_unlink, ...);
6 }
7 static void Rmdir(const FunctionCallbackInfo<Value>& args) {
8   AsyncCall(..., uv_fs_rmdir, ...);
9 }
10 FSReqBase* AsyncCall(..., Func fn, Args... fn_args) { AsyncCall.uv_fs_read ←
    [...]
11   return AsyncDestCall(..., fn, fn_args...);
12 }
13 FSReqBase* AsyncDestCall(..., Func fn, Args... fn_args) { AsyncDestCall.uv_fs_read ←
    [...]
14   req_wrap->Dispatch(fn, fn_args..., after);
15 }
16 int ReqWrap<T>::Dispatch(LibuvFunction fn, Args... args) { Dispatch.uv_fs_read ←
    [...]
17   CallLibuvFunction<T, LibuvFunction>::Call( fn, ...);
18 }
19 struct CallLibuvFunction<ReqT, int(*)>(uv_loop_t*, ReqT*, Args...> {
    [...]
20   static int Call(T fn, uv_loop_t* loop, ReqT* req, PassedArgs... args) {
21     return fn(loop, req, args...);
22   }
23 };

```



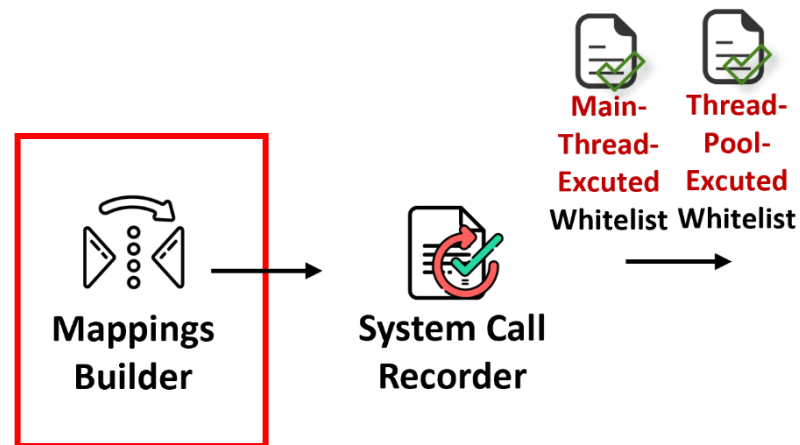
# Step 1: Call Graph Construction

- C/C++ call graph > eliminate non-existing nodes/edges
  - ✓ *Partial context-aware* analysis for **switch-case statement** & **function pointer parameter**
- **Implementation**
  - ✓ clang with wllvm > llvm link > **SVF ++**



## Step 2: Mapping Builder

- We build call graph traversal for call graphs of the Node.js application layer, Binding Module layer, and Dependency layer.
- We build LLVM Pass for the Built-in Module layer.
- We get mappings of different layers.



### Algorithm 1: Mapping generation

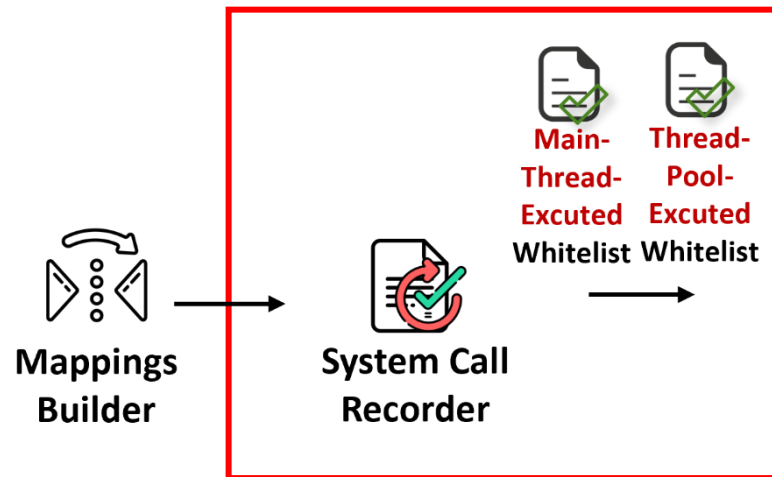
```

Data: Call graphs of builtin modules cg.builtin, call graphs of binding modules and
dependencies cg.bottom, call graphs of libc cg.libc, LLVM IR of binding modules
ir_bind
Result: Output mapping dict M
1 M.builtin ← {}; /* Mappings of builtin modules */
2 forall cg.Module ∈ cg.builtin do
3   forall method ∈ module.exports do
4     M.builtin.module.method ← {};
5     Callers C invoked by the method by traversing cg.Module;
6     forall c ∈ C do
7       if c == internalBinding then
8         M.builtin.module.method ← (module, method);
9 M.binding ← {}; /* Mappings of binding modules */
10 forall ir.Module ∈ ir_bind do
11   forall method_bind ∈ ir.Module do
12     M.binding.module.method ← func;
13 M.depend ← {}; /* Mappings of dependencies */
14 forall func ∈ M.bindings.module.method do
15   M_depend.module.method ← {};
16   Callers C invoked by the function by traversing cg.bottom;
17   forall c ∈ C do
18     if c == libc then
19       M.depend.module.method ← libc;
20 forall libc ∈ M.depend.module.method do
21   M.depend.module.method ← {};
22   Callers C invoked by the function by traversing cg.libc;
23   forall c ∈ C do
24     if c ∈ syscall then
25       M.depend.module.method.libc ← syscall;
26 return M.builtin, M.binding, M.depend;

```

## Step 3: System Call Recorder

- Based on mappings, we calculate the **system call whitelists** for the Node.js application.
- We **divide** the system call list into the system call list of **main thread** and the system call list of **the thread pool**.




---

### Algorithm 2: Whitelist generation

---

```

Data: Call graph of Node.js Application  $cg\_app$ , mapping sets  $M$ 
Result: Output whitelist  $W$ 
1  $wl.main \leftarrow \{\}$ ;
2  $wl.pool \leftarrow \{\}$ ;
3 Callers  $C$  invoked application by traversing  $cg\_app$ ;
4 forall  $c \in C$  do
5     | if  $c \in M.native.c$  then
6     |     | forall  $b \in M.native.c$  do
7     |     |     | if  $f \in M.bindings.b$  then
8     |     |     |     | forall  $sys \in M.depend.f$  do
9     |     |     |     |     | if  $b \in builtin\_threadpool$  then
10    |     |     |     |     |     |  $wl.pool \leftarrow sys$ ;
11    |     |     |     |     |     | else
12    |     |     |     |     |     |  $wl.main \leftarrow sys$ ;
13 return  $W$ 
  
```

---



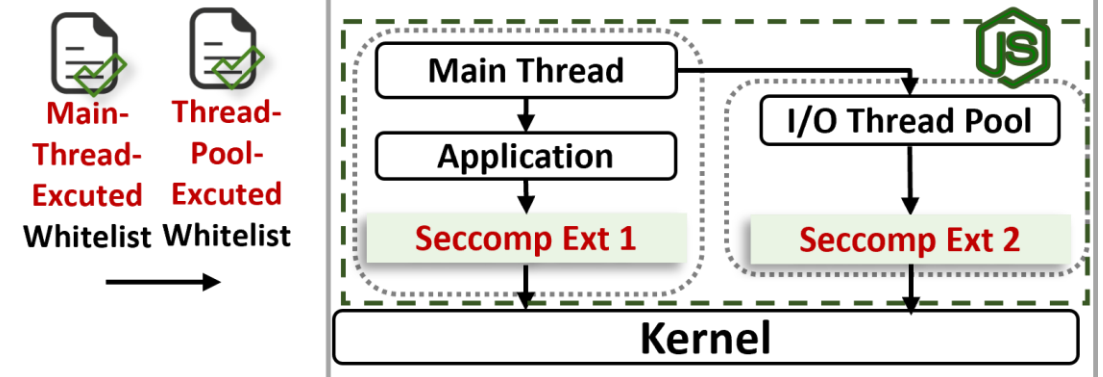
## Step 4: Hodor Installation

### ➤ Seccomp Implementation

- ✓ For ***thread pool required applications***, we ***first*** install the filter for the thread pool thread and ***then*** install the filter for the main thread to prevent the thread pool thread from inheriting the main thread filter.
- ✓ For thread pool dis-required applications, we ***only*** load the main thread filter.

### ➤ Read/write Permission Restrictions.

- ✓ ***Read*** and ***write system calls*** are widely used by Node.js engine.
- ✓ ***Chroot*** mechanism and Switch the ownership.



Restriction Profile Installation



# Agenda

- Introduction
- Previous work & Remaining challenges
- HODOR: system call level protection system for Node.js applications
- [Evaluation](#)
- Conclusion & Takeaways


# Evaluation

## ➤ Dataset

- ✓ 169 packages suffered from ACE attacks
- ✓ Three large-scale real-world applications (koa, express and json-server).
- ✓ Node.js core tests and 4 well-known web frameworks (koa, fastify, express, and connect).

## ➤ Total Result

- ✓ HODOR can **reduce the attack surface** of Node.js applications to **19.42%**.







# Evaluation - Comparison with Other Techniques

- ✓ HODOR can defend against a wider spectrum of attacks (additionally covering arbitrary command execution) with less runtime overhead.

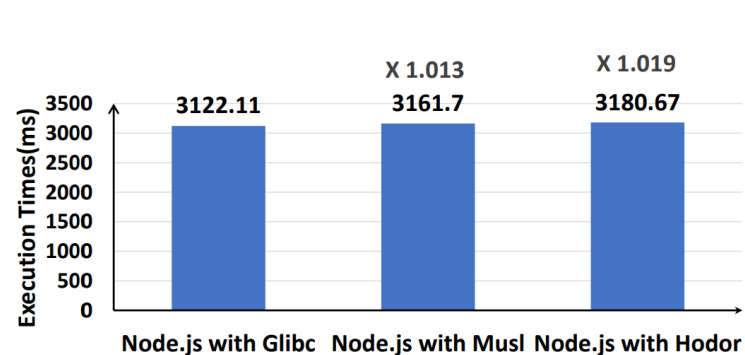
Table 7: Exploit execution for packages with arbitrary code execution attacks.

Package Name	Initial Attack	HODOR								MIR							
		Cmd		Permission		Net				Cmd		Permission		Net			
		Initial Attack	Exec	Fork	Setgid	Setuid	Connect	Listen	Bind	Initial Attack	Exec	Fork	Setgid	Setuid	Connect	Listen	Bind
accesslog	Print sSS	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
access-policy	Print 123	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
adm-zip	Path traversal	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
angular-expressions	Write file angular-expressions-success	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
cd-messenger	Print JHU	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
crypt	Print defconrussia	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
djv	touch HACKED	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ejs	Write file ejs-success	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
eslint-utils	Write file eslint-utils-success	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
front-matter	Print 1	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
grunt	Returns Date.now	✓	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
hot-formula-parser	Write file test	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
is-my-json-valid	Execute cat /etc/passwd	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
jben	Write file malicious	✓	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓
j-yaml	Returns Date.now	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
kmc	Write file kmc-success	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
marsdb	Write file marsdb-success	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
metacalc	Print process	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mixin-pro	Print hacked	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
m-log	Print injected	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mock2easy	Write mock2easy-success	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
modjs	Write modjs-success.txt	x	x	x	✓	✓	✓	x	x	✓	✓	✓	✓	✓	✓	✓	✓
modulify	Print hacked	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mol-prot	Write file mol-prot-success	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mongodb-query-parser	touch test-file	✓	✓	✓	✓	✓	✓	x	x	✓	✓	✓	✓	✓	✓	✓	✓
mongo-express	exec calculator	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mongoose-mask	Print "my evil code was run"	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mongo-parse	Write file hacked	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
morgan	Write file morgan-success	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
morgan-json	Print GLOBAL_CTF_HIT	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
music.js	Write file Song	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
node-extend	Print 123	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
node-extend	Print 123	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
node-extend	Print 123	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
node-serialize	Execute ls	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
notevil	Print pwned	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
pg	Print process.env	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
pid-class	Print 123	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
post-loader	Print rce	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
protojs	Write file protojs-success	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
realms-shim	Messed with Object.toString	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
reduce-css-calc	Read /etc/passwd	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
safe-eval	Return proces	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
safer-eval	Print id	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
sandbox	Print process.pid	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
serialize-javascript	Print 1	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
serialize-to-js	Execute ls	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
static-eval	Print hacked	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
tar	Overwrite file	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
themify	Write file Song	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
total.js	Touch HACKED	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
total.js	Touch HACKED	✓	x	x	✓	✓	✓	x	x	✓	✓	✓	✓	✓	✓	✓	✓
total.js	Touch HACKED	✓	x	x	✓	✓	✓	x	x	✓	✓	✓	✓	✓	✓	✓	✓
typed-function	Execute whoami	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
underscore	touch HELLO	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
value-censorship	Access the Function constructor	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
vm2	return process.env	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
zhaoyan91-eval-in-vm	return process.env	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
mobile-icon-resizer	Print hacked	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

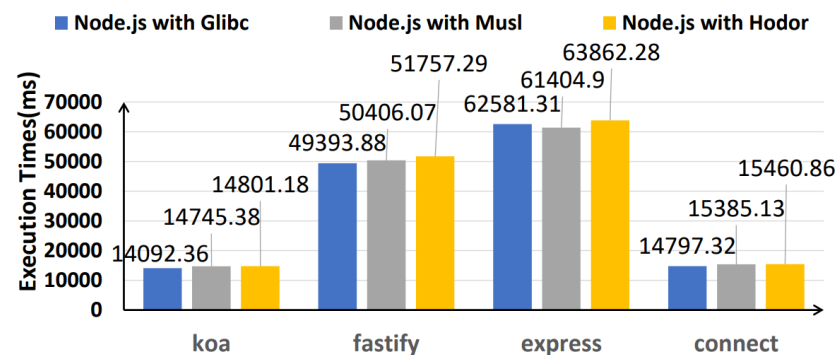
x: Exploits are executed; ✓: Exploits are blocked.

## Evaluation - Runtime Overhead

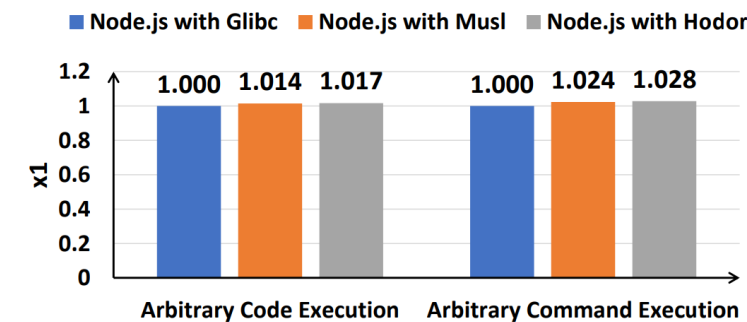
- ✓ The runtime overhead of HODOR is **0.61%** for **Node.js core tests**, **2.80%** for the **web framework**, and **0.39%** for all **168 packages**.



(a) Node.js Core Tests



(b) Web Framework



(c) Node.js Packages

Figure 8: Runtime overhead of Node.js core tests, web framework and applications under the protection of HODOR (RQ4).



# Agenda

- Introduction
- Previous work & Remaining challenges
- HODOR: system call level protection system for Node.js applications
- Evaluation
- [Conclusion & Takeaways](#)





# Conclusion & Takeaways

1. Attendees will learn a new call graph building methods for JavaScript code and C/C++ code.
2. Attendees will gain knowledge of a novel protection mechanism for Node.js applications, focusing on thread-level and system call-level security.
3. Attendees will develop an understanding of the hazards associated with vulnerabilities in the Node.js application ecosystem, with a particular emphasis on system call-level vulnerabilities.



# Thanks & Questions?

Wenya Wang, Xingwei Lin  @xwlin\_roy