

# A Dozen Years of Shellphish

## *From DEFCON to the Cyber Grand Challenge*

Antonio Bianchi

antonio@cs.ucsb.edu

Jacopo Corbetta

jacopo@cs.ucsb.edu

Andrew Dutcher

dutcher@cs.ucsb.edu



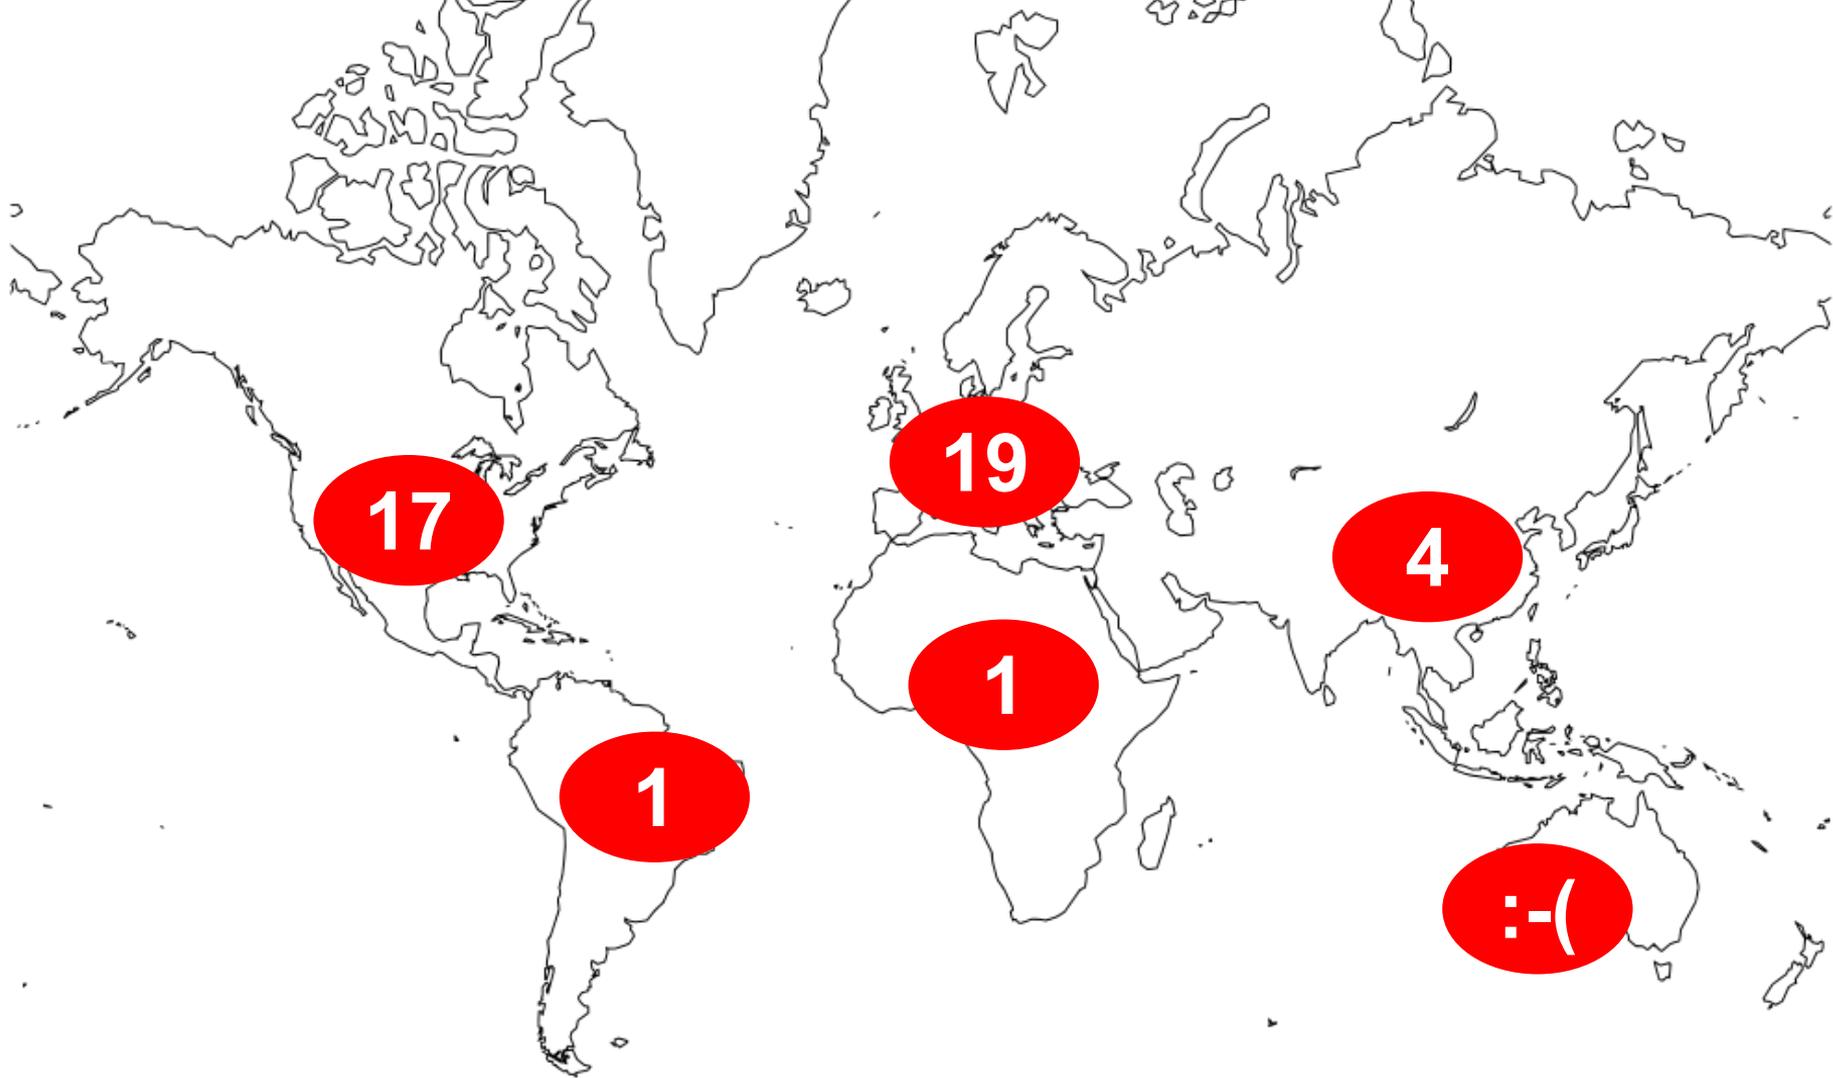
- Who are we?
  - A team of security enthusiasts
    - do research in System Security
    - play Capture the Flag competitions
    - released a couple of tools

- Started (in 2004) at:
  - SecLab: University of California, Santa Barbara



- expanded to:
  - Northeastern University: Boston
  - Eurecom: France
  - ...





- Security competition:
  - exploit a vulnerable service / website / device
  - reverse a binary
  - ...
- Different formats
  - Jeopardy – Attack-Defense
  - Online – Live
  - ...
- Basic idea: find the secret, submit to organizers, ... *profit*



- We do not only play CTFs
- We also organize them!
  - UCSB iCTF
    - Attack-Defense format
    - every year, since 2002!
  - Try to innovate with a different style every year
    - Site: [ictf.cs.ucsb.edu](http://ictf.cs.ucsb.edu)
    - Base: [github.com/ucsb-seclab/ictf-framework](https://github.com/ucsb-seclab/ictf-framework)
    - Vigna, et al., *"Ten years of ictf: The good, the bad, and the ugly."* 3GSE, 2014.

# Why we're here

---



SHERIFF



SHILLPHISH

- DARPA Cyber Grand Challenge (CGC)  
The (almost-)Million Dollar Baby
- Our Cyber Reasoning System (CRS)  
Fancy term for auto-playing a CTF
- Automated Vulnerability Discovery
- Live example using **angr**  
Open-source binary analysis framework
- Towards the Cyber Grand Challenge Finals (CFE)

# Cyber Grand Challenge (CGC)



- 2004: DARPA Grand Challenge
  - Autonomous vehicles



# Cyber Grand Challenge (CGC)



- 2014: DARPA **Cyber** Grand Challenge
  - Autonomous **hacking!**



# Cyber Grand Challenge (CGC)

---



- Started in 2014
- Qualification event: June 3rd, 2015, online
  - ~70 teams → 7 qualified teams
- Final event: August 4th, 2016 @ DEFCON (Las Vegas)
  - Winning CRS will also play against humans!



[cybergrandchallenge.com](http://cybergrandchallenge.com) / [cgc.darpa.mil](http://cgc.darpa.mil)

- Attack-Defense CTF
- ~~Solving security challenges~~ → Developing a system that solves security challenges
- Develop a system that **automatically**
  - Exploit vulnerabilities in binaries
  - Patch binaries, removing the vulnerabilities
- No human intervention
- Think it's trivial? How would you play?
  - And how would you organize it?

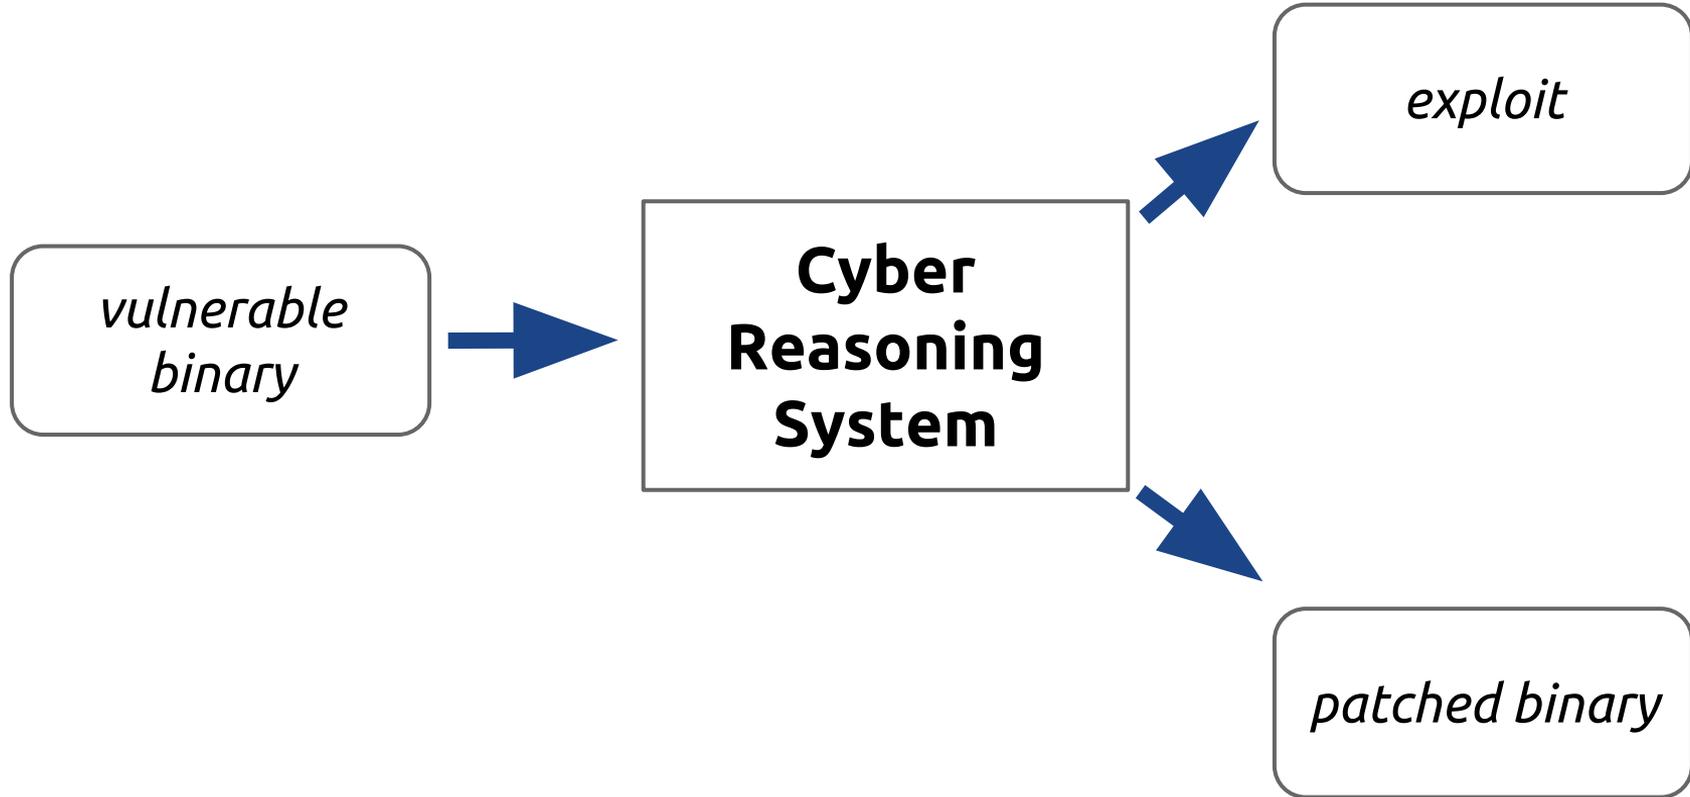
- Exploits
  - “getting a flag” (how? where?)
  - For the quals: exploit = **crash**
- Defend
  - `int main() { return 0; }`
    - Functionality checks
  - `SIGSEGV => exit(0)`
    - No easy “out-of-band” error handling
  - `QEMU-style interpreter: interrupts => exit(0)`
    - Performance cost (CPU, memory, file size)

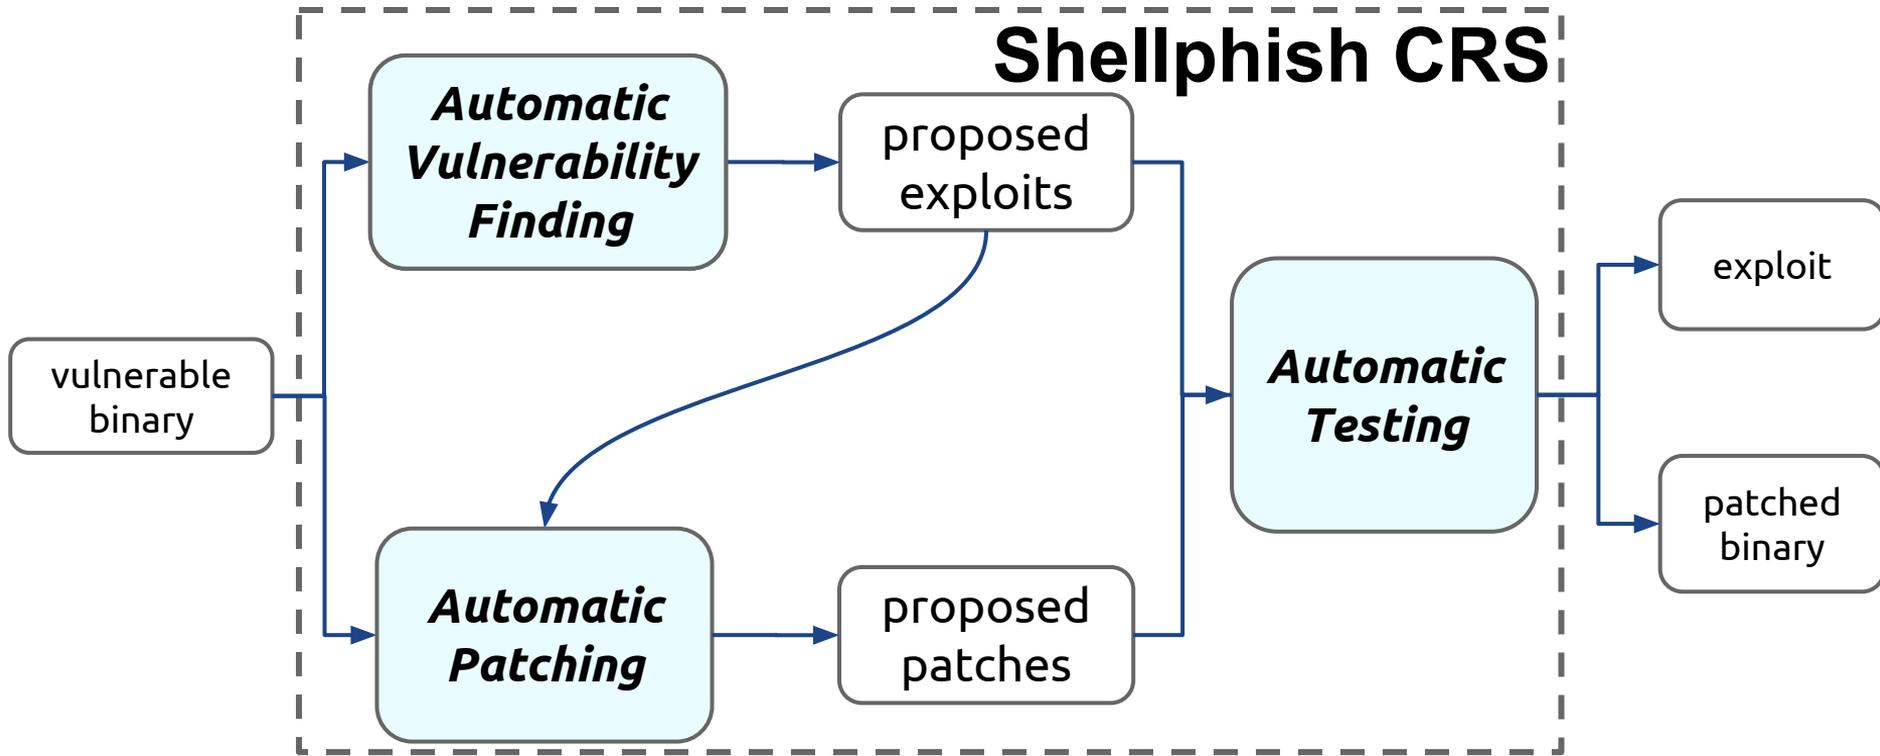


- Basic idea:
  - Real(istic) programs
  - No “extra” complications
    - Is modeling the entire POSIX API a good use of team resources? What about the file systems? Or horrible things like interruptible syscalls?

- Architecture: Intel x86, 32-bit
- OS: DECREE
  - Linux-like, but with 7 syscalls only
    - `transmit` / `receive` / `fdwait` ( $\approx$  `select`)
    - `allocate` / `deallocate` (even executable!)
    - `random`
    - `_terminate`
  - no signals, threads, shared memory
- “Bring Your Own Defense” approach (and pay for it)
  - Not even “the usual”: stack is executable, no ASLR, ...

- DARPA Cyber Grand Challenge (CGC)  
The (almost-)Million Dollar Baby
- **Our Cyber Reasoning System (CRS)**  
Fancy term for auto-playing a CTF
- Automated Vulnerability Discovery
- Live example using **angr**  
Open-source binary analysis framework
- Towards the Cyber Grand Challenge Finals (CFE)



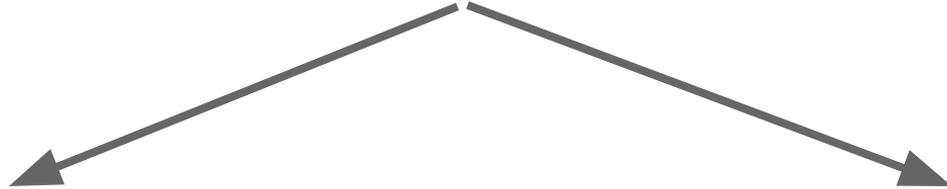


- DARPA Cyber Grand Challenge (CGC)  
The (almost-)Million Dollar Baby
- Our Cyber Reasoning System (CRS)  
Fancy term for auto-playing a CTF
- **Automated Vulnerability Discovery**
- Live example using **angr**  
Open-source binary analysis framework
- Towards the Cyber Grand Challenge Finals (CFE)

**“How do I crash a binary?”**



**“How do I reach state X in a binary?”**



***Dynamic Analysis/Fuzzing***

***Symbolic Execution***

- How do I reach the state: "You win!" is printed?

```
x = int(input())
if x >= 10:
    if x < 100:
        → print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

- How do I reach the state: “You win!” is printed?

```
x = int(input())
if x >= 10:
    if x < 100:
        → print "You win!"
        else:
            print "You lose!"
    else:
        print "You lose!"
```

- Try “1” → “You lose!”
- Try “2” → “You lose!”
- ...
- Try “10” → “You win!”

- How did we use Fuzzing for CGC?
- Coverage-guided fuzzing
  - Looking for “crashing inputs”
  - Based on AFL [lcamtuf.coredump.cx/afl/](http://lcamtuf.coredump.cx/afl/)
- In general, it cannot work in some cases
  - e.g., “magic numbers”, computations, ...

- How do I reach the state: "You win!" is printed?

```
x = int(input())
if x >= 10:
    if x^2 == 152399025:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

# Symbolic Execution



- Interpret the binary code and replace user-input with symbolic variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

|                             |
|-----------------------------|
| <b>State A</b>              |
| <b>Variables</b><br>x = ??? |
| <b>Constraints</b><br>{     |

# Symbolic Execution



- Interpret the binary code and replace user-input with symbolic variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

| State A                     |
|-----------------------------|
| <b>Variables</b><br>x = ??? |
| <b>Constraints</b><br>{     |

A red arrow originates from the `input()` function in the code block on the left and points to the `x = ???` entry in the 'Variables' section of the 'State A' table on the right.

# Symbolic Execution



- Follow all feasible paths, tracking "constraints" on variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

|                             |
|-----------------------------|
| <b>State A</b>              |
| <b>Variables</b><br>x = ??? |
| <b>Constraints</b><br>{     |

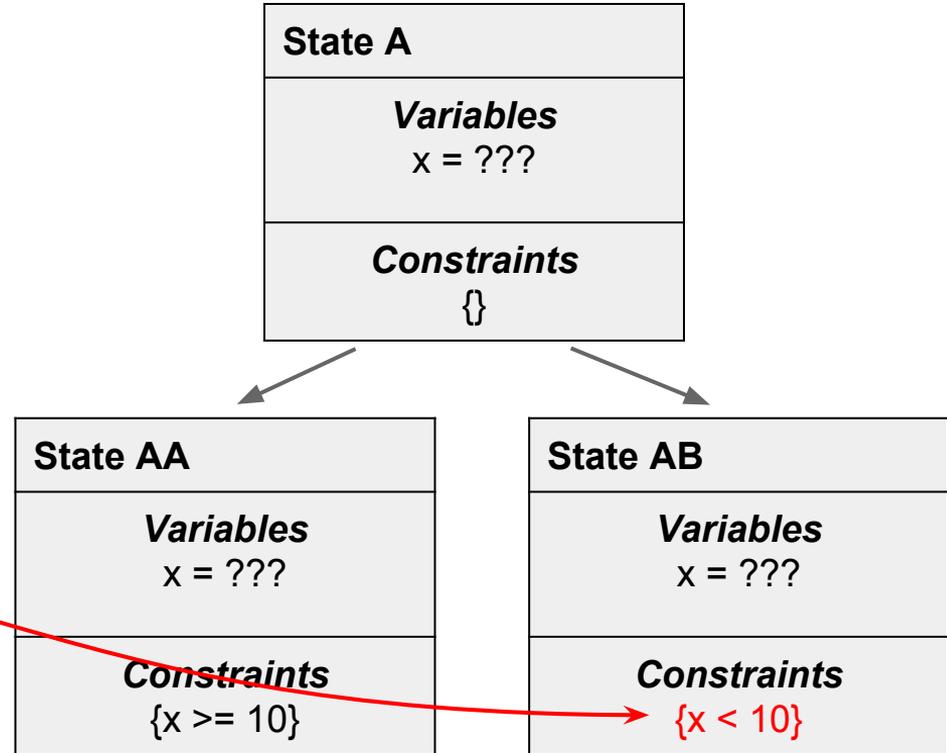
|                                 |
|---------------------------------|
| <b>State AA</b>                 |
| <b>Variables</b><br>x = ???     |
| <b>Constraints</b><br>{x >= 10} |

# Symbolic Execution



- Follow all feasible paths, tracking "constraints" on variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```



# Symbolic Execution



- Follow all feasible paths, tracking "constraints" on variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

| State AA                        |
|---------------------------------|
| <b>Variables</b><br>x = ???     |
| <b>Constraints</b><br>{x >= 10} |

| State AB                       |
|--------------------------------|
| <b>Variables</b><br>x = ???    |
| <b>Constraints</b><br>{x < 10} |

# Symbolic Execution



- Follow all feasible paths, tracking "constraints" on variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

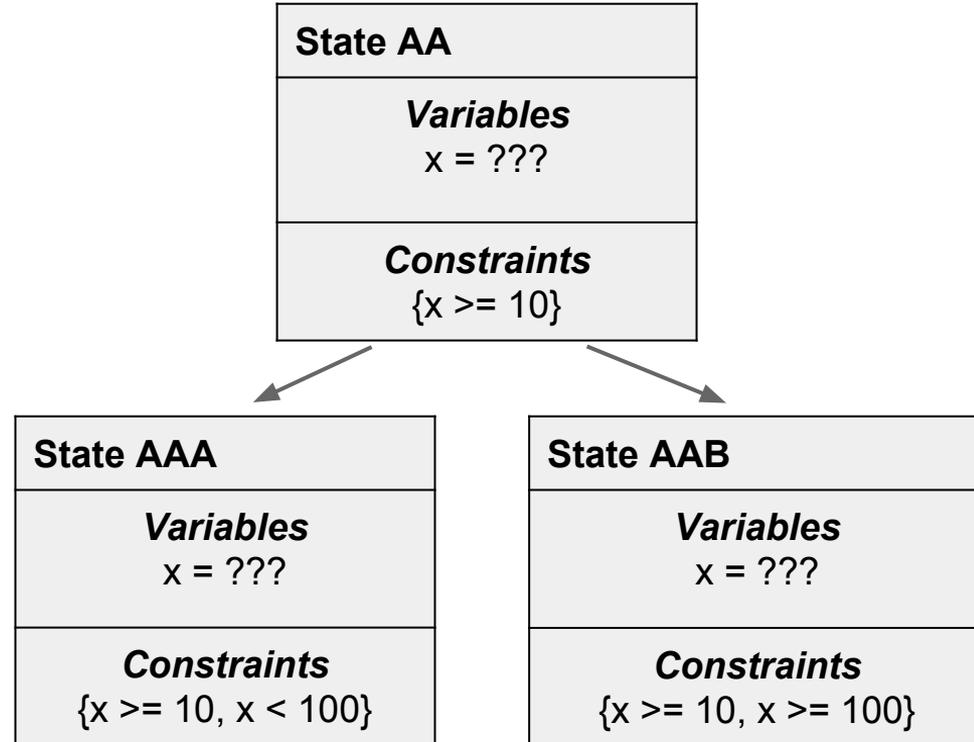
|                                 |
|---------------------------------|
| <b>State AA</b>                 |
| <b>Variables</b><br>x = ???     |
| <b>Constraints</b><br>{x >= 10} |

# Symbolic Execution



- Follow all feasible paths, tracking "constraints" on variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```



# Symbolic Execution



- Concretize the constraints on the symbolic variables

```
x = int(input())
if x >= 10:
    if x < 100:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

|  |
|--|
| State AAA                                |
| <b>Variables</b><br>x = ???              |
| <b>Constraints</b><br>{x >= 10, x < 100} |



**Concretization**

|                            |
|----------------------------|
| State AAA                  |
| <b>Variables</b><br>x = 99 |

# Symbolic Execution



- Concretize the constraints on the symbolic variables

```
x = int(input())
if x >= 10:
    if x^2 == 152399025:
        print "You win!"
    else:
        print "You lose!"
else:
    print "You lose!"
```

|  |
|--|
| <b>State AAA</b>                                     |
| <b>Variables</b><br>x = ???                          |
| <b>Constraints</b><br>{x >= 10,<br>x^2 == 152399025} |



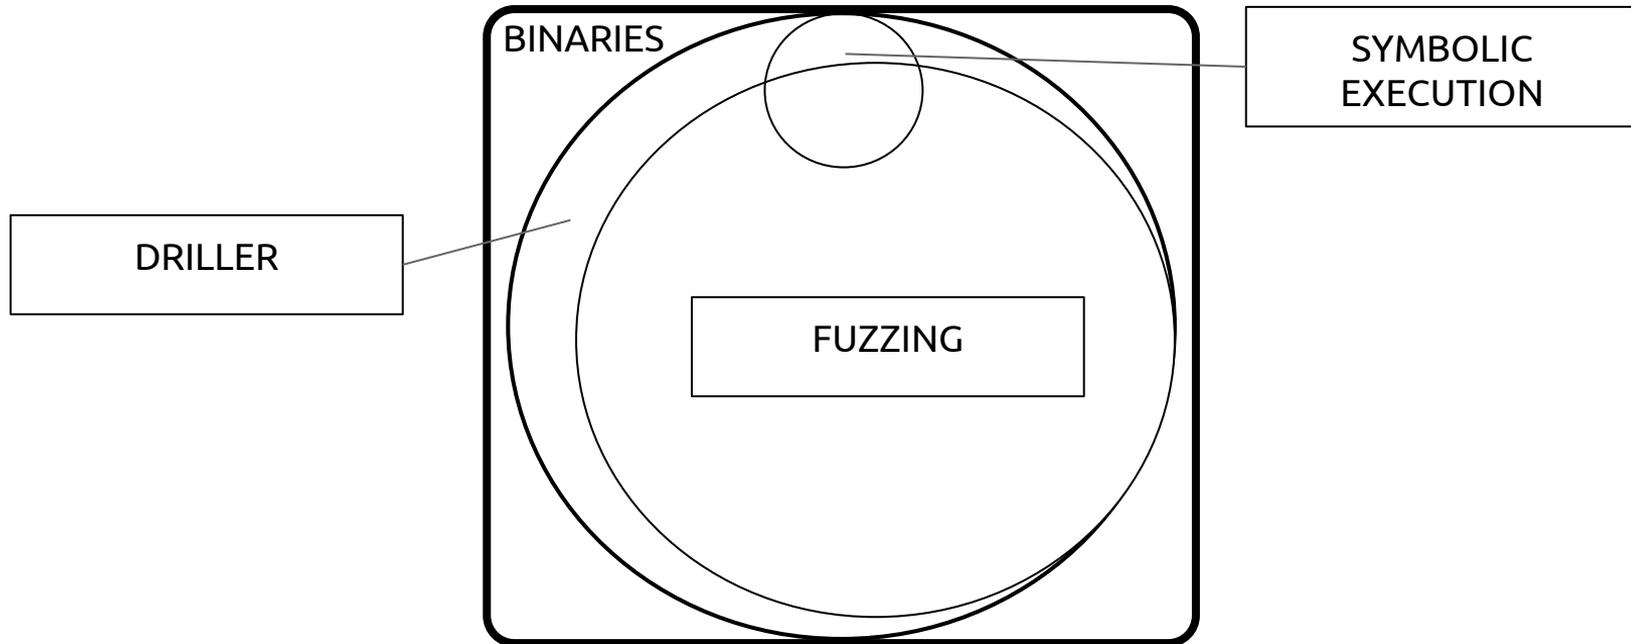
**Concretization**

|                               |
|-------------------------------|
| <b>State AAA</b>              |
| <b>Variables</b><br>x = 12345 |

- How did we use Symbolic Execution for CGC?
  - We used the symbolic execution engine of **angr**, the binary analysis platform developed at UCSB
- Symbolically execute the binaries looking for
  1. Memory accesses outside allocated regions
  2. “Unconstrained” instruction pointer (e.g., controlled by user input)
    - `eax = <user input>, jmp eax`
- If either 1. or 2. is true
  - we found an input that will make the program crash

- Combining the two approaches
- *“Driller: Augmenting Fuzzing Through Selective Symbolic Execution”*
  - Network and Distributed System Security Symposium (NDSS), February 2016, San Diego

- *“Driller: Augmenting Fuzzing Through Selective Symbolic Execution”*



# Coming up:

---



- DARPA Cyber Grand Challenge (CGC)  
The (almost-)Million Dollar Baby
- Our Cyber Reasoning System (CRS)  
Fancy term for auto-playing a CTF
- Automated Vulnerability Discovery
- **Live example using angr**  
**Open-source binary analysis framework**
- Towards the Cyber Grand Challenge Finals (CFE)

- Binary analysis platform, developed at UCSB
- Open source: [github.com/angr](https://github.com/angr) (star it!)
- Architecture independent
  - x86 (ELF, **CGC**, PE), amd64, mips, mips64, arm, aarch64, ppc, ppc64



- Written in Python!
  - installable with one (two?) command!
    - `mkvirtualenv angr` ← (optional, but don't complain if it's broken)
    - `pip install angr`
  - interactive shell (using IPython)
  - it has an interactive GUI

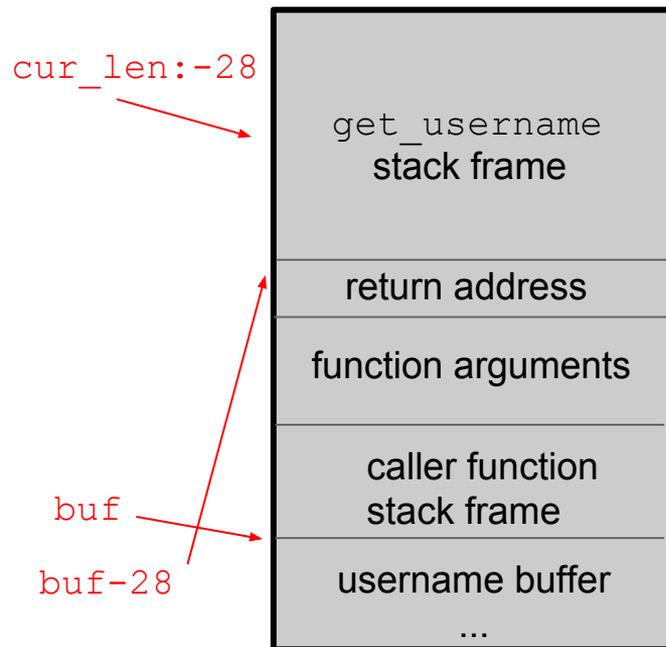
## Grub: “Back to 28” vulnerability

- Pressing backspace 28 times on the grub username prompt can get you a rescue shell
- <http://hmarco.org/bugs/CVE-2015-8370-Grub2-authentication-bypass.html>

```
if (key == '\b') // Does not checks underflows !!
{
    cur_len--; // Integer underflow !!
    grub_printf ("\b");
    continue;
}
```

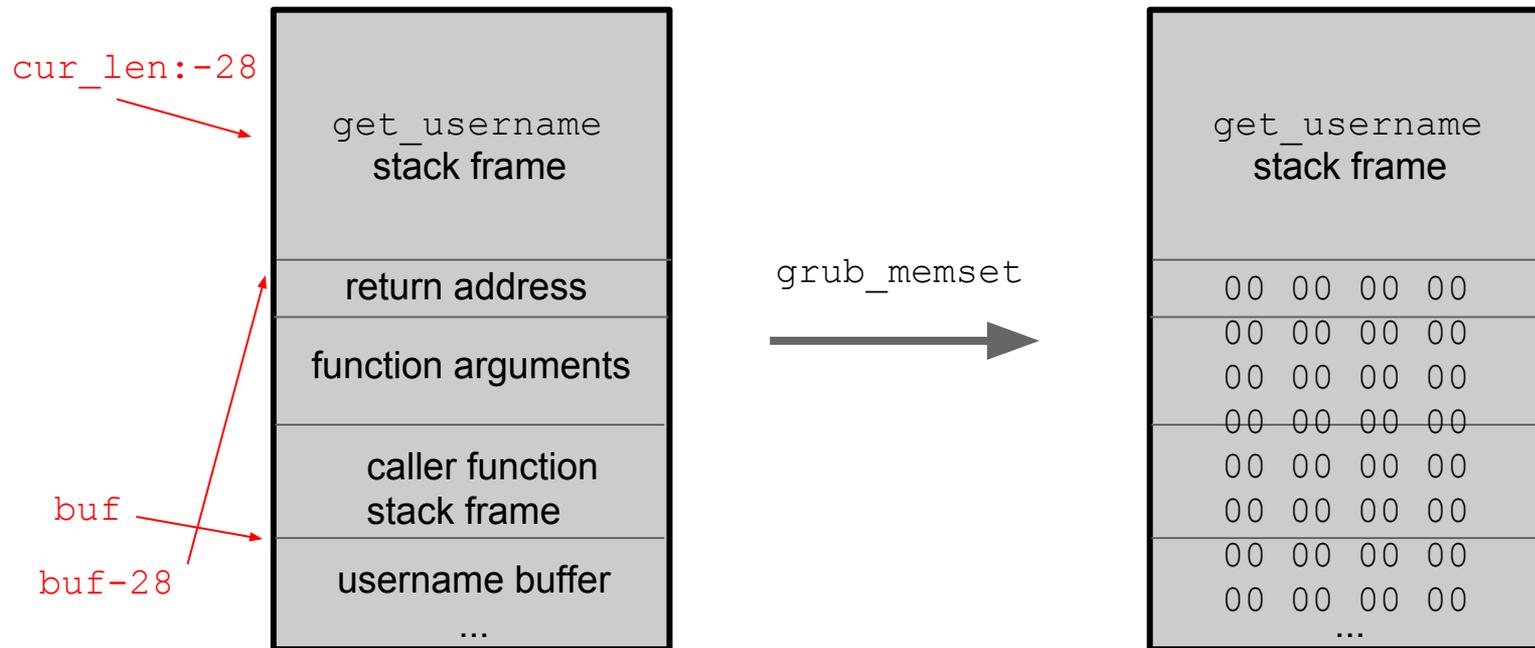
## Grub: “Back to 28” vulnerability

```
grub_memset( buf + cur_len, 0, buf_size - cur_len); // Out of bounds overwrite
```



## Grub: “Back to 28” vulnerability

```
grub_memset( buf + cur_len, 0, buf_size - cur_len); // Out of bounds overwrite
```



## Grub: “Back to 28” vulnerability

- Somehow, jumping to 00:0000 is completely exploitable
- Way beyond the scope of this demo

Although it seems quite difficult to build a successful attack just jumping to 0x0, we will show how to do it.

**Is there life after jumping to 0x0 ?**

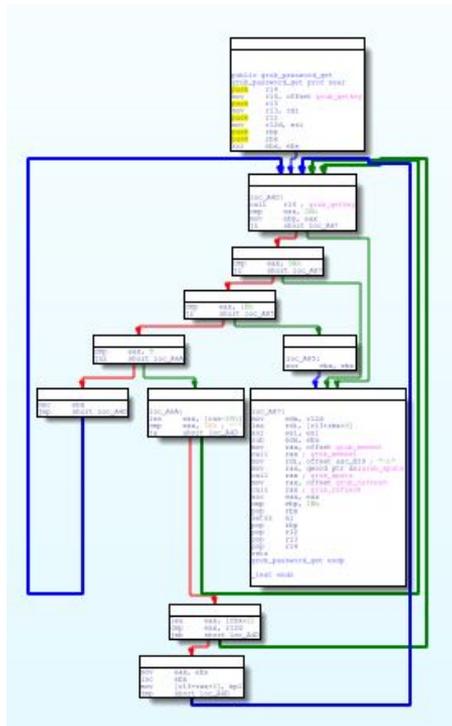
At address 0x0 resides the IVT (Interrupt Vector Table) of the processor. It contains a variety of pointers in the form of segment:offset.

```
0x0000: push %ebx
0x0001: incl (%eax)
0x0003: push %ebx
0x0005: incl (%eax)
0x0007: ret
0x0009: lodsb 0xb
0x000a: push %ebx
0x000b: incl (%eax)
0x000c: push %ebx
0x000d: incl (%eax)
0x000e: push %ebx
0x000f: incl (%eax)
0x0010: push %ebx
0x0011: incl (%eax)
0x0012: push %ebx
0x0013: incl (%eax)
0x0014: push %ebx
0x0015: incl (%eax)
0x0016: push %ebx
0x0017: incl (%eax)
0x0018: push %ebx
0x0019: incl (%eax)
0x001a: push %ebx
0x001b: incl (%eax)
0x001c: movsl %ds:(%esi), %es:(%edi)
0x001d: incb (%eax)
0x001e: xchg %ebp, %ecx
0x001f: add %dh, %al
0x0020: jge 0x0
```

The lowest part of the IVT interpreted as code.

At this early stage of the boot sequence, the processor and the execution framework is not fully functional. Next are the

## Grub: “Back to 28” vulnerability



- The correct path to the exploit goes around this loop 28 times, each of which has to follow a specific path
- The universe will grow old and die before naive symbolic execution finds this bug
- **Demonstration: this doesn't work, really!**
- A technique (implemented by angr) called veritesting<sup>1</sup> solves this problem in *some cases* by merging states when their instruction pointers converge, but in *this case* the complexity generated is too much for the constraint solver

<sup>1</sup><http://users.ece.cmu.edu/~aavgerin/papers/veritesting-icse-2014.pdf>

- Symbolic execution is powerful
- Symbolic execution is stupid
- You are incredibly weak
- You are very clever



Use angr to unlock your true potential

## Grub: “Back to 28” vulnerability

Manual examination of the state explosion tells you:

- Where the wasted computational power is going
- How to be more efficient

The naive approach is doing lots of weird things like entering letters and then deleting them again and again, or pressing the “home” key several times in a row, which don’t produce any interesting new states to analyze.

You can fix this!

Grub: “Back to 28” vulnerability

# Final demonstration

## Finding the bug

# Coming up:

---



- DARPA Cyber Grand Challenge (CGC)  
The (almost-)Million Dollar Baby
- Our Cyber Reasoning System (CRS)  
Fancy term for auto-playing a CTF
- Automated Vulnerability Discovery
- Live example using **angr**  
Open-source binary analysis framework
- **Towards the Cyber Grand Challenge Finals (CFE)**

- 7 teams passed the qualification phase
- Shellphish is one of them! :-)
- We exploited 44 binaries out of 131
- Every qualified team received \$750,000 !

- Different setup
  - Round-based attack-defense CTF
  - Probably, zero human intervention allowed
    - Not even bug fixing?
  - Data about previous rounds is available:
    - submitted exploits/patched binaries performance
    - (anonymized) network traffic
    - patches from other teams
  - Stealing patched binaries/exploits?

- Exploits are more realistic:
  - Two types:
    - Crash at a specific location and set a specific register to a specific value
    - Leak data from a specific memory page
  - We'll need a more realistic exploit generator:
    - **angr** automatic ROP-chain builder!
- Every team can also deploy network-level filtering rules

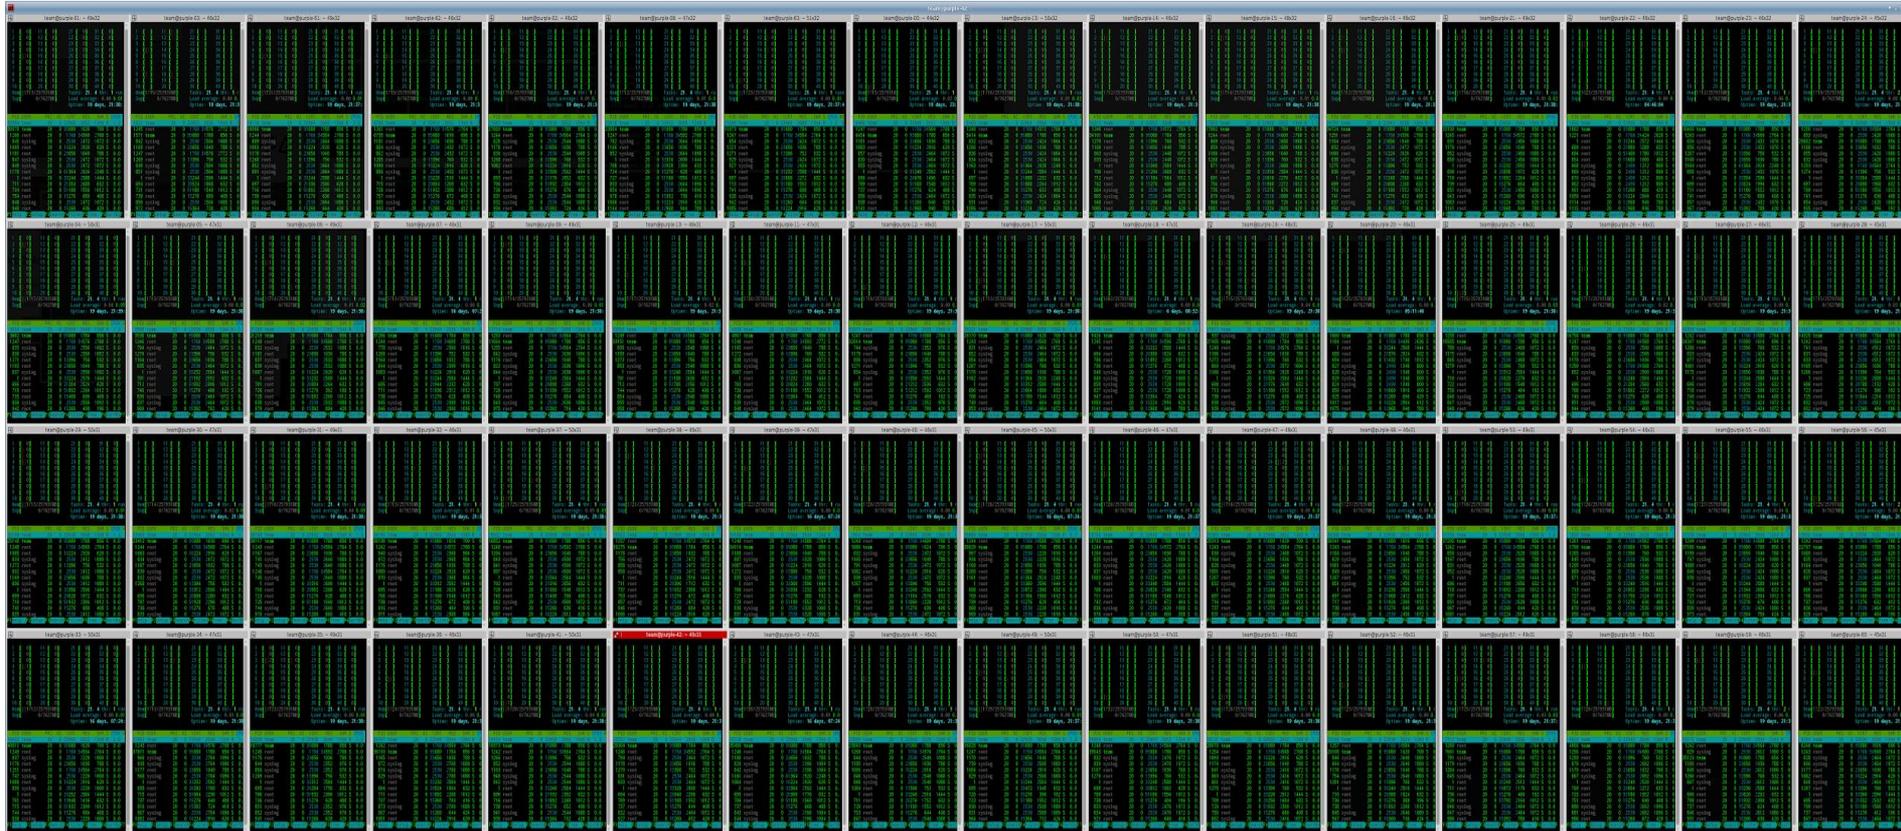
- Every team has access to a cluster of:
  - 1280 cores
  - 16 TB of RAM
  - 128 TB of storage



```
team@purple-42: ~ 46x31

1 [ ] 11 [ ] 21 [ ] 31 [ ]
2 [ ] 12 [ ] 22 [ ] 32 [ ]
3 [ ] 13 [ ] 23 [ ] 33 [ ]
4 [ ] 14 [ ] 24 [ ] 34 [ ]
5 [ ] 15 [ ] 25 [ ] 35 [ ]
6 [ ] 16 [ ] 26 [ ] 36 [ ]
7 [ ] 17 [ ] 27 [ ] 37 [ ]
8 [ ] 18 [ ] 28 [ ] 38 [ ]
9 [ ] 19 [ ] 29 [ ] 39 [ ]
10 [ ] 20 [ ] 30 [ ] 40 [ ]
Mem[1725/257931MB] Tasks: 21, 4 thr; 1 r
Swp[ 0/7627MB] Load average: 0.00 0.
Uptime: 19 days, 21:3

  PID USER      PRI  NI  VIRT   RES   SHR  S  CP
20357 team      20   0 22944  2048  1340  R   1
20304 team      20   0 95080  1784   856  S   0
 1246 root       20   0 176M  34592  2764  S   0
   794 syslog    20   0 253M  2464  1072  S   0
  1178 root       20   0 23656  1036   788  S   0
  1270 root       20   0 13396   752   532  S   0
   831 syslog    20   0 253M  2464  1072  S   0
   833 syslog    20   0 253M  2464  1072  S   0
    1 root       20   0 33328  2600  1444  S   0
   694 root       20   0 20940  2256   632  S   0
   709 root       20   0 51180  1552  1012  S   0
   737 root       20   0 15276   644   408  S   0
   832 syslog    20   0 253M  2464  1072  S   0
   927 root       20   0 15260   652   420  S   0
F1 Help F2 Setup F3 Search F4 Filter F5 Tree F6 Sort
```



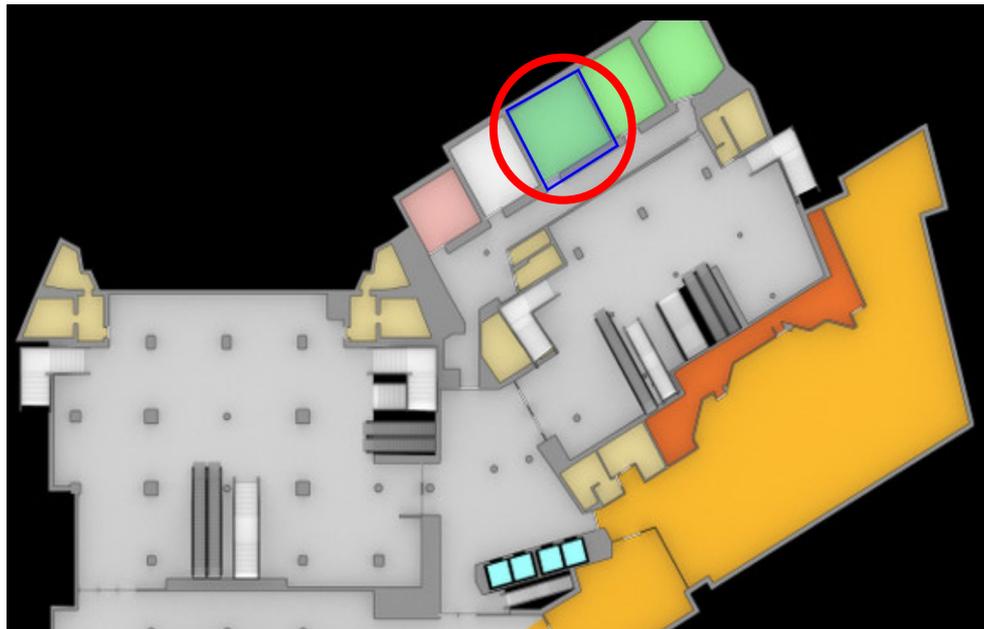
- Finals will take place on August 2016
  - DEFCON, Las Vegas
- Money prizes!
  - First place: \$2'000'000
  - Second place: \$1'000'000
  - Third place: \$750'000
- The winning team will compete against human teams at DEFCON CTF Finals :-)

# Shellphish CGC Team



# I want more...

- **angr** hands-on workshop
  - Just after this talk
  - Hall 13 (first floor)
  - Bring your laptop!



“That’s all folks!”



# *Questions?*

## **References:**

CGC: [cybergrandchallenge.org](http://cybergrandchallenge.org) – [cgc.darpa.mil](http://cgc.darpa.mil)

DARPA CGC presentation (DEFCON 2015): [youtu.be/gnyCbU7jGYA](https://youtu.be/gnyCbU7jGYA)

angr: [angr.io](http://angr.io) – [github.com/angr](https://github.com/angr)

emails: [antoniob@cs.ucsb.edu](mailto:antoniob@cs.ucsb.edu) – [jacopo@cs.ucsb.edu](mailto:jacopo@cs.ucsb.edu) – [dutcher@cs.ucsb.edu](mailto:dutcher@cs.ucsb.edu)