DARPA's Cyber Grand Challenge: Autonomous detection and patching of online vulnerabilities

Jeanne-Marie Musca

ABSTRACT

Early this August, the finals of DARPA's "Cyber Grand Challenge" (CGC) were held, in which autonomous computer programs competed against each other in a "Capture the Flag" challenge. These programs had to find vulnerabilities in software they were given, fix these flaws, and try to exploit flaws in the software given to other programs. We explore the inaugural CGC from the perspective of both those competing and those running the CGC.

MEET THE **WINNER**[1]

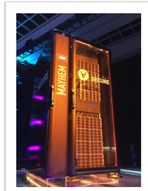

Mayhem is a cyber reasoning system built by ForAllSecure.

Without *any* human intervention, it protected a networked server for several hours during the finals for the CGC.

It detects vulnerabilities in binaries, and can then either patch them or exploit them.

TECHNIQUES

Mayhem uses several techniques**:**

OFFENSE	DEFENSE
Symbolic Execution	Hot Patching
Directed Fuzzing	Recompilation

Symbolic Execution

- Wrap program in code that takes symbolic input
- Symbolic input represents whole classes of inputs

Directed Fuzzing

- Generate input that is likely to trigger vulnerabilities

Hot Patching

- Insert a jump point right before suspicious code
- Redirect to code that prevents program from crashing

Recompilation

- Custom recompilation framework
- Preserves control flow of the program

DEF CON 24

Cyber Grand Challenge[2]

DARPA's Challenge:

How to make people interested in a competition that they can't actually see happening.

Solution:

Autonomous systems inhabit colorful servers on a stage. Several visualizations allow audience to watch the action.

Capture the Flag: Mayhem vs Humans[3]

Organizers had to adopt the CGC format to allow Mayhem to participate.

Mayhem finished last, but did have some moments in the competition when it was ahead of some human teams.

DARPA'S TESTING **TOOLS**[5]

DECREE

DARPA designed this environment just to test the autonomous systems in a secure environment.

- Allows only 7 system calls.
- Processes don't share memory.

Challenge Binaries

Programs that contain vulnerabilities from the CWE.

Competitors patch these binaries, and issue Proofs of Vulnerability by submitting input that compromise them.

THE **CYBER GRAND CHALLENGE**

The road to the finals was a (nearly) three year journey:

October 2013

DARPA announces its intent to hold the CGC

June 2015

Qualification event tests systems ability to analyze and patch programs.

Seven finalists are chosen.

August 2016

Seven autonomous systems play capture the flag against each other at DEF CON 24.

AREAS OF **EXCELLENCE**[4]

There are five Areas of Excellence in which DARPA tested the autonomous systems:

Analysis

- Discover the function of a program

Patching

- Fix vulnerabilities in programs

Vulnerability Scanning

- Generate input that proves a program is vulnerable

Service Resiliency

- Keep a program running and available

Network Defense

- Defend a network against real-time attacks

SOURCES

Image Credits:

Mayhem's portrait:
https://blog.forallsecure.com/2016/08/06/mayhem-wins-darpa-cgc/

Cyber Grand Challenge Stage:
http://www.darpa.mil/program/cyber-grand-challenge

References:

[1] https://blog.forallsecure.com/
[2] https://s3.amazonaws.com/cgcdist/cfe/cgc-final_event-cfe-brochure.pdf
[3] https://blog.legitbs.net/
[4] https://cgc.darpa.mil/CGC_Rules_18_Nov_14_Version_3.pdf
[5] https://github.com/CyberGrandChallenge