CALDERA

Automating Adversary Emulation

Andy Applebaum, Doug Miller

The MITRE Corporation

MITRE

Bios

Andy Applebaum

Lead Cyber Security Engineer

ATT&CK, AI Research

Doug Miller

Sr. Cyber Security Engineer

ATT&CK, CARET, Red teaming

Adversary Emulation?

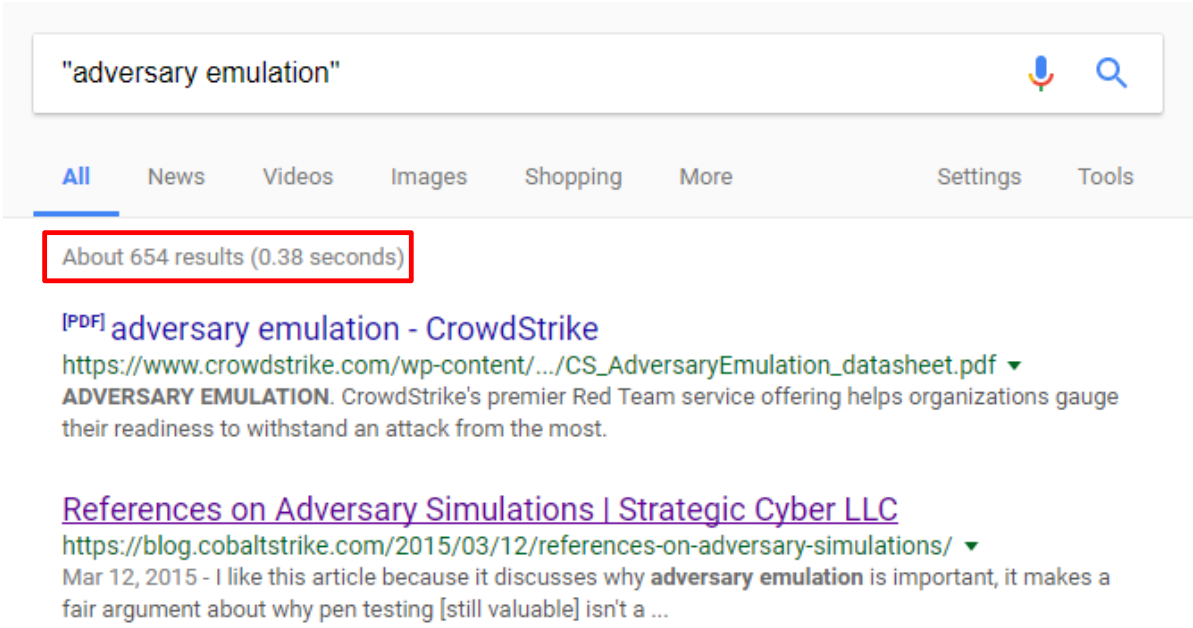

The False Negative Problem

As a defender, you have no idea what you miss

Cue: Adversary

- **Introduce a realistic* adversary on your network**

  *an emulated adversary

- ***Now you can determine what happens if an attacker gets on your network***
  - Did I detect them?
  - How far did they get?
  - How can I improve my detection and prevention?

Iterative Defensive Cycle

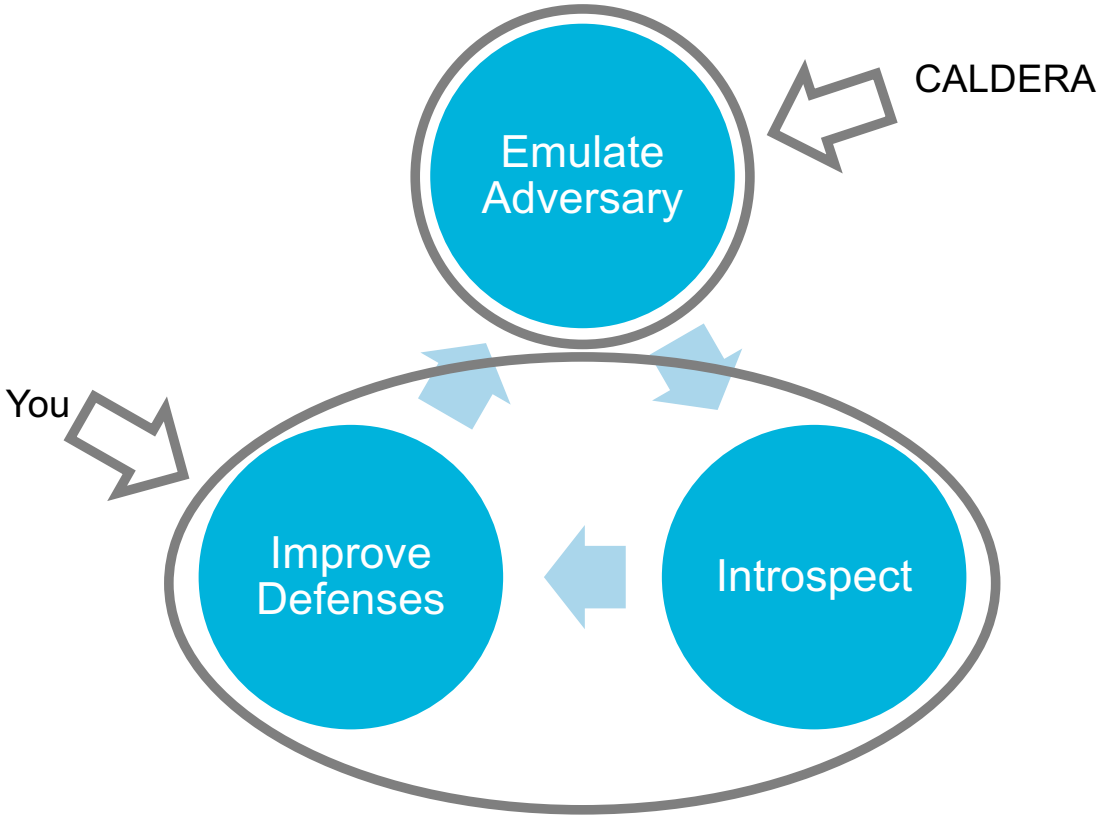

Successful Adversary Emulation

Make it real: Use the same techniques, tools, methods and goals of an attacker

End-to-End: Don't just look for holes or perform small attacks. Start from the initial compromise and go until objectives are accomplished

Repeatable: Be repeatable, so that your detection and prevention improvement (or degradation) can be measured over time

CALDERA:

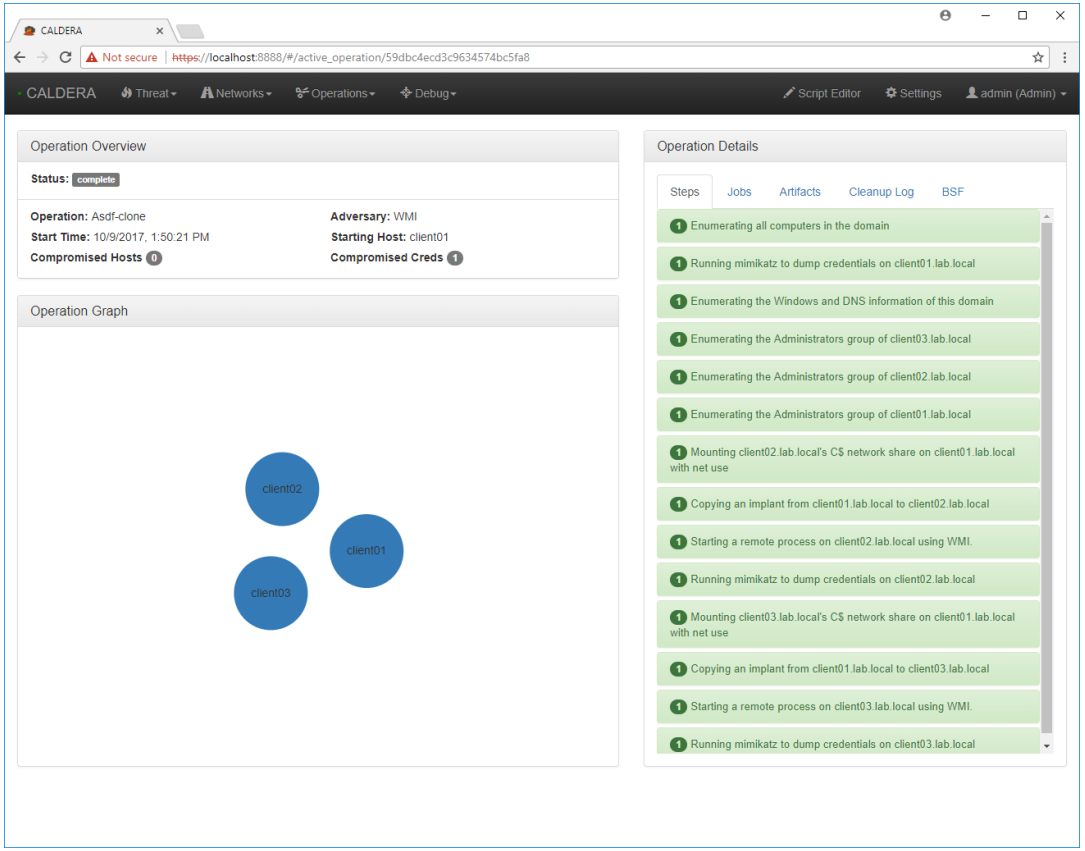

CALDERA – Conducting an Operation

1. **Load the CALDERA shim onto network hosts**
2. **Create an adversary by giving it capabilities**
3. **Launch the operation**

- **During the operation:**
  - CALDERA will report its activities, including artifacts created
  - CALDERA will automatically stop if it has exhausted its toolkit

- **After the operation:**
  - CALDERA will provide a report of what it did
  - CALDERA will automatically "reset" infected hosts, removing artifacts and stopping processes

Ingredients for Automated Adversary Emulation

What the adversary can do

- The adversary model
- MITRE ATT&CK™

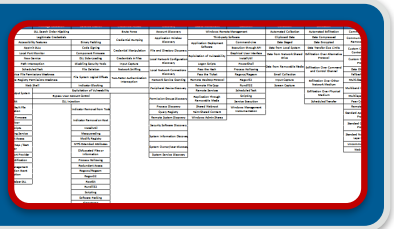

How the adversary chooses what to do

- CALDERA logic and decision engine

What the adversary needs to do it

- Infrastructure to support real adversary emulation
- Management server; client agents; web interface

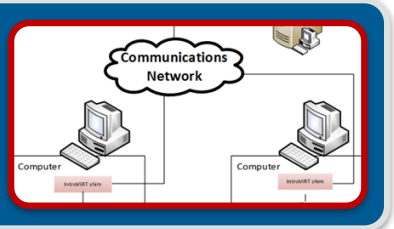

The Adversary Model

Choosing an Adversary Model

CALDERA emulates a real adversary after they get into a network

ATT&CK Matrix™ Tactics and Techniques

Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Execution	Collection	Exfiltration	Command and Control
DLL Search Order Hijacking			Brute Force	Account Discovery	Windows Remote Management		Audio Capture	Automated Exfiltration	Commonly Used Port
Legitimate Credentials			Credential Dumping	Application Window Discovery	Third-party Software		Automated Collection	Data Compressed	Communication Through Removable Media
Accessibility Features		Binary Padding			Application Deployment Software	Command-Line	Clipboard Data	Data Encrypted	
AppInit DLLs		Code Signing	Credential Manipulation	File and Directory Discovery		Execution through API	Data Staged	Data Transfer Size Limits	Connection Proxy
Local Port Monitor		Component Firmware			Exploitation of Vulnerability	Execution through Module Load	Data from Local System	Exfiltration Over Alternative Protocol	Custom Command and Control Protocol
New Service		DLL Side-Loading	Credentials in Files	Local Network Configuration Discovery			Data from Network Shared Drive		
Path Interception		Disabling Security Tools	Input Capture		Logon Scripts	Graphical User Interface		Exfiltration Over Command and Control Channel	Custom Cryptographic Protocol
Scheduled Task		File Deletion	Network Sniffing	Local Network Connections Discovery	Pass the Hash	InstallUtil	Data from Removable Media		
File System Permissions Weakness		File System Logical Offsets	Two-Factor Authentication Interception		Pass the Ticket	MSBuild			Data Encoding
Service Registry Permissions Weakness				Network Service Scanning	Remote Desktop Protocol	PowerShell	Email Collection	Exfiltration Over Other Network Medium	Data Obfuscation
Web Shell		Indicator Blocking		Peripheral Device Discovery	Remote File Copy	Process Hollowing	Input Capture		Fallback Channels
Authentication Package	Exploitation of Vulnerability				Remote Services	Regsvcs/Regasm	Screen Capture	Exfiltration Over Physical Medium	Multi-Stage Channels
	Bypass User Account Control			Permission Groups Discovery	Replication Through Removable Media	Regsvr32	Video Capture		Multiband Communication
Bootkit	DLL Injection					Rundll32		Scheduled Transfer	
Component Object Model Hijacking		Component Object Model Hijacking		Process Discovery	Shared Webroot	Scheduled Task			Multilayer Encryption
Basic Input/Output System		Indicator Removal from Tools		Query Registry	Taint Shared Content	Scripting			Remote File Copy
Change Default File Association		Indicator Removal on Host		Remote System Discovery	Windows Admin Shares	Service Execution			Standard Application Layer Protocol
Component Firmware		Install Root Certificate		Security Software Discovery		Windows Management Instrumentation			Standard Cryptographic Protocol
External Remote Services		InstallUtil		System Information Discovery					Standard Non-Application Layer Protocol
Hypervisor		Masquerading							
Logon Scripts		Modify Registry		System Owner/User Discovery					
Modify Existing Service		MSBuild							Uncommonly Used Port
Netsh Helper DLL		Network Share Removal		System Service Discovery					Web Service
Redundant Access		NTFS Extended Attributes		System Time Discovery					
Registry Run Keys / Start Folder		Obfuscated Files or Information							
Security Support Provider		Process Hollowing							
Shortcut Modification		Redundant Access							
Windows Management Instrumentation Event Subscription		Regsvcs/Regasm							
		Regsvr32							
		Rootkit							
Winlogon Helper DLL		Rundll32							
		Scripting							
		Software Packing							
		Timestomp							

https://attack.mitre.org

Supported Adversary Actions

Persistence
- Registry autorun keys
- Scheduled Task
- Services

Privilege Escalation
- Weak service perms
- Weak service file perms
- Unquoted paths (Path interception)

Defense Evasion
- Scripting
- Timestomping

Credential Access
- Credential Dumping

Lateral Movement
- Remote File Copy
- Windows Admin shares
- Pass the Hash
- PsExec

Discovery
- Remote System Discovery
- Local Network config
- Registry
- Account
- System information
- Processes/services
- System Owner
- Permission Group
- Files

Execution
- PowerShell
- Scheduled Task
- WMI
- SC (service controller)

Exfiltration
- HTTP/S

Decision Making for Automated Adversary Emulation

Early CALDERA

- **First version**
  - Finite-state machine (FSM) approach
  - Successfully tested within MITRE

- ***Hard to extend to new techniques***
  - Action needs to be coded into FSM
  - FSM logic needs to be recomputed
  - Inflexible in operation; hard to configure

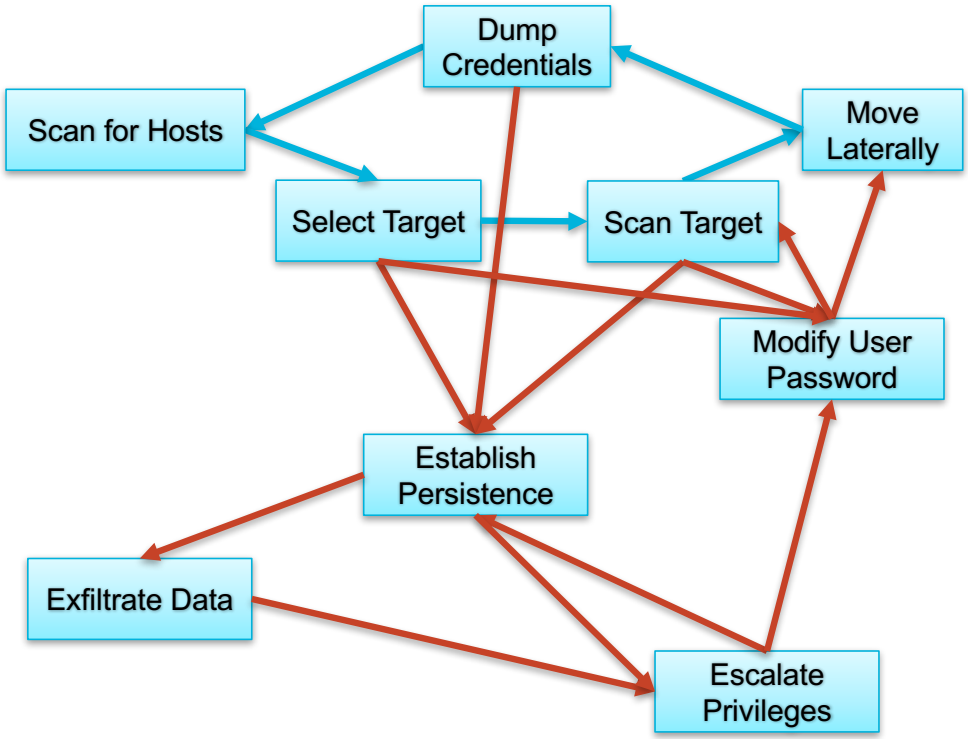

Early CALDERA

- **First version**
  - Finite-state machine (FSM) approach
  - Successfully tested within MITRE

- ***Hard to extend to new techniques***
  - Action needs to be coded into FSM
  - FSM logic needs to be recomputed
  - Inflexible in operation; hard to configure

- ***Predictability during execution***
  - Easy to spot and identify

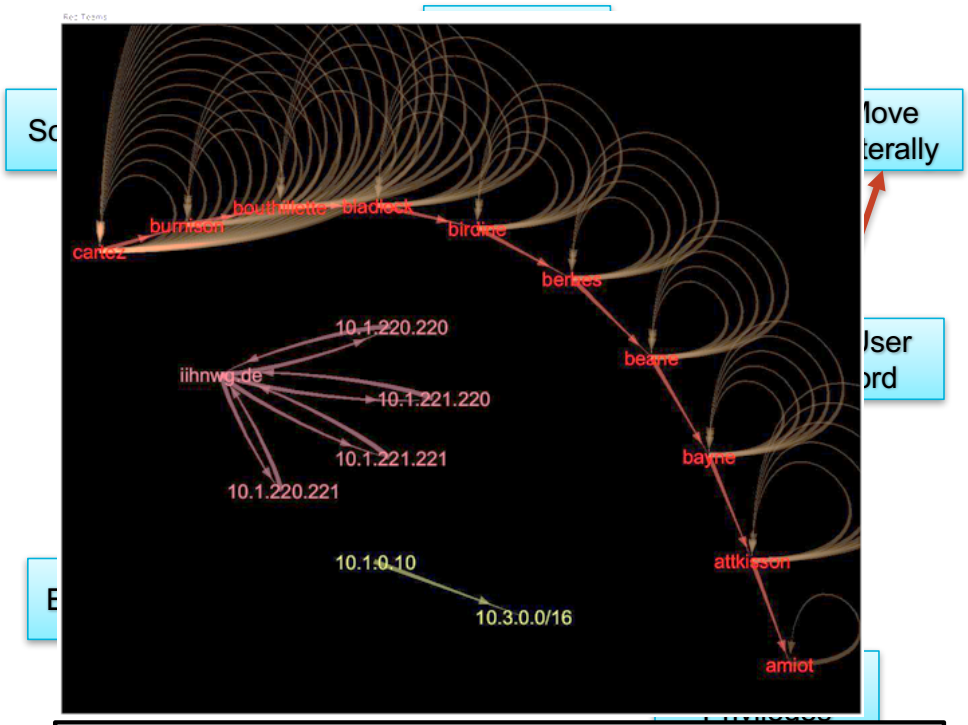

Visualized with the MITRE CyGraph tool

Noel, S., et al. "CyGraph: Graph-Based Analytics and Visualization for Cybersecurity." *Handbook of Statistics* 35 (2016): 117-167.

Designing an Adversary Decision Engine

- **Typical engagements have *human* operators dictating and controlling the assessment**
  - Targeting, TTP selection, domain inference…
  - … all needs to be fully automated!

- **Ideally, our automated adversary will:**
  - Make intelligent choices
  - Achieve tactical objectives
  - Easily incorporate new techniques
  - Work in new and unknown environments
  - Vary operations to test the defense
  - Chain weaknesses for maximum effect

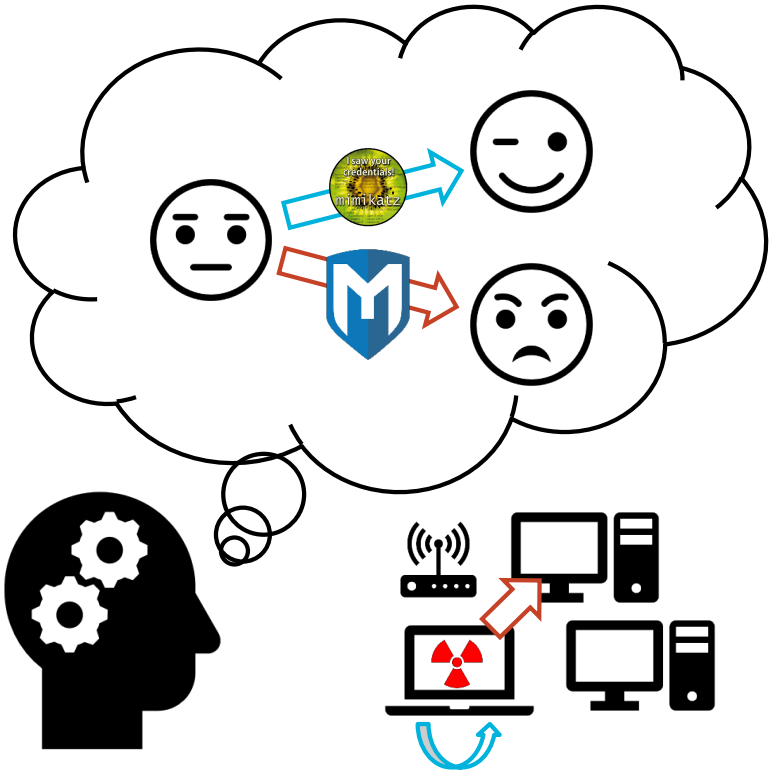

Example Scenario

- Start with code execution and a RAT on Host 1
- Identified sensitive files on Host 2
- *Goal*: exfiltrate sensitive data from Host 2

Example Scenario

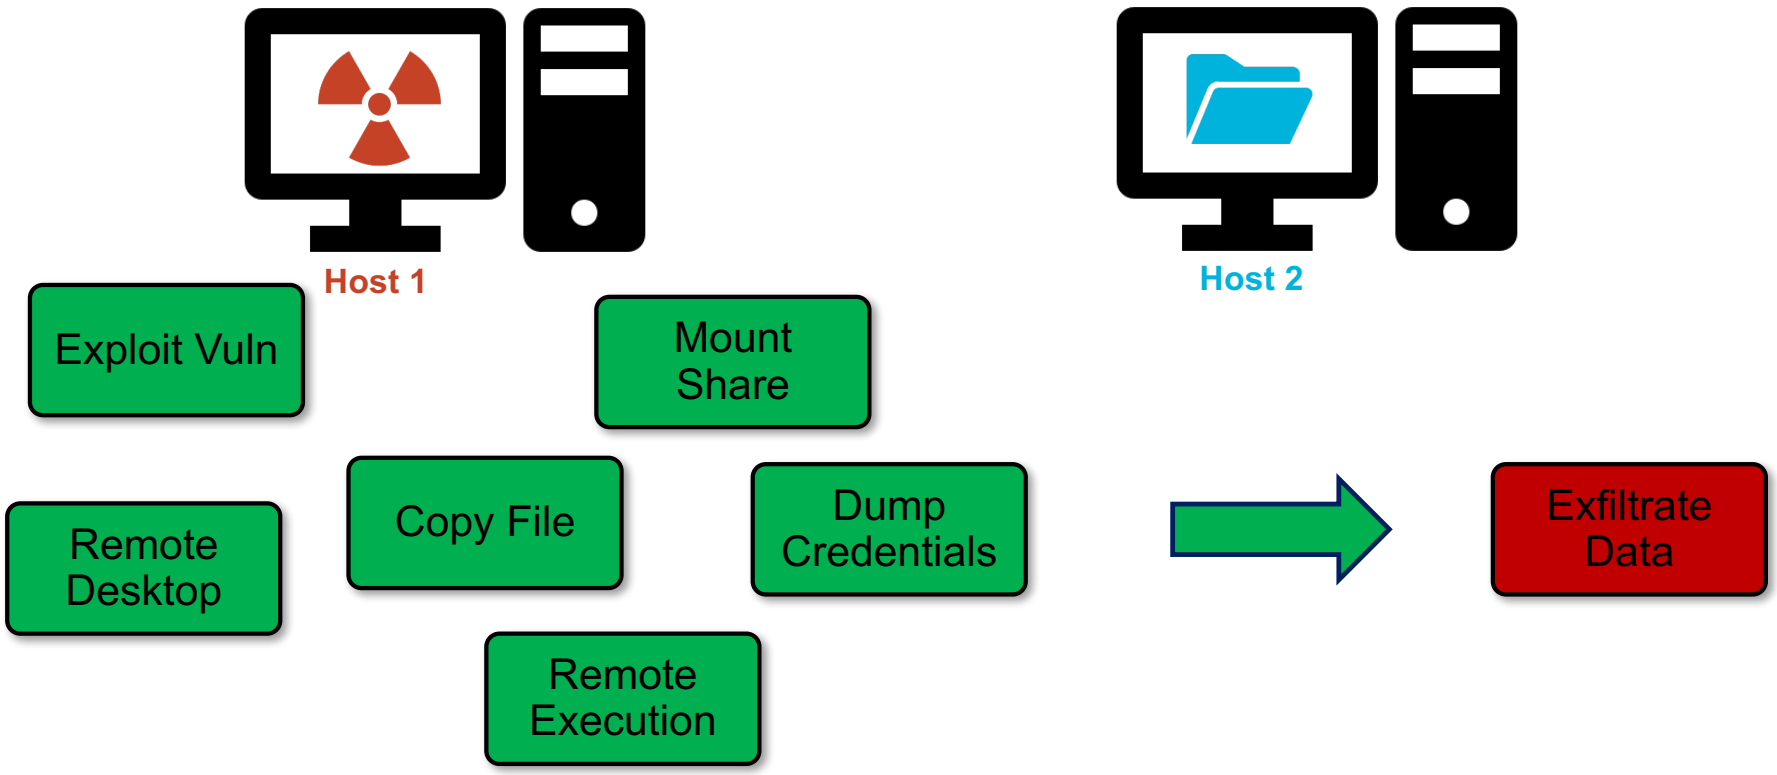

Analyzing Copying Over a File

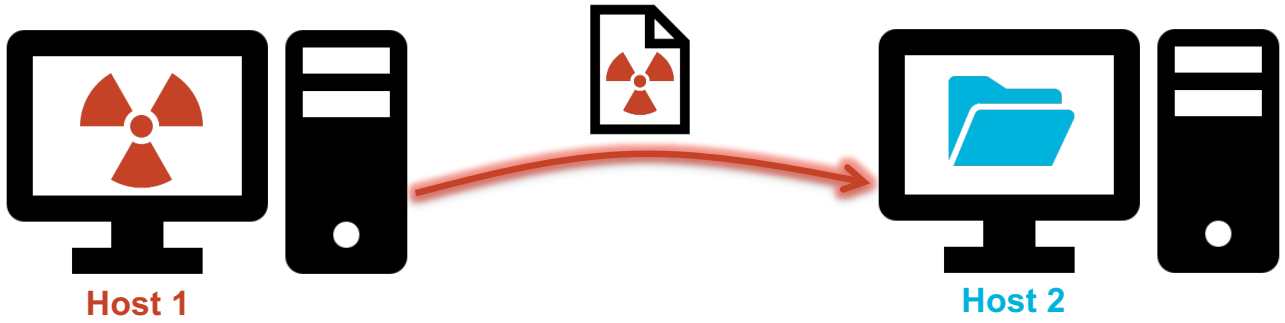

- **What do we need to do to copy a RAT over?**
  - Working RAT on source host
  - Mounted file share from target onto source host
  - Write access to file share

  } **Requirements, or *preconditions***

- **What happens after copying a RAT over?**
  - There will be a new file on the target host
  - That file will contain the RAT

  } **Consequences, or *postconditions***

Making a Plan to Copy a File

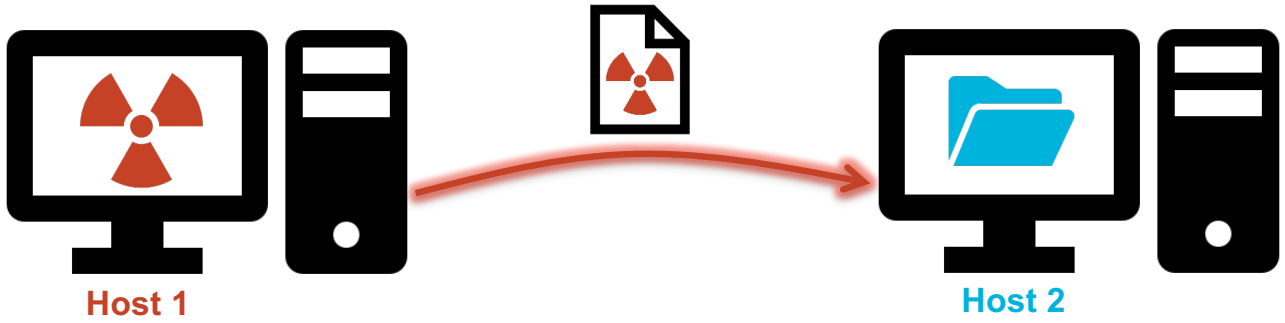

Dump Credentials
- **Need**: Elevated RAT
- **Get**: Credentials

Mount Share
- **Need**: Credentials
- **Get**: Mounted Share

Copy File
- **Need**: Mounted Share
- **Get**: File on Target

Sequence of actions, or *plan*

The Core CALDERA Idea

- **Move from an *explicit, prescribed* decision model towards a *dynamic, model-based one***

- **Tag actions with *semantic execution information*:**
  - **Preconditions** specify the requirements that must be true to execute a technique
  - **Postconditions** specify the consequences that will be true after executing a technique

- **No longer need to be explicitly told what to do!**
  - Instead, compare the current state to the available actions to determine which are valid

- **Added bonus: planning for the future**
  - If I dump credentials now, that can help me execute lateral movement in the future!

Fun With Preconditions

- **Preconditions tell you what you can do *now***
  - In chess: can tell you which moves are valid
  - Taken further: can tell you which moves are legal

- **In the emulation sense: given an escalated foothold on a host, we can:**
  - Dump credentials
  - Add/modify registry keys
  - Setup scheduled tasks
  - …

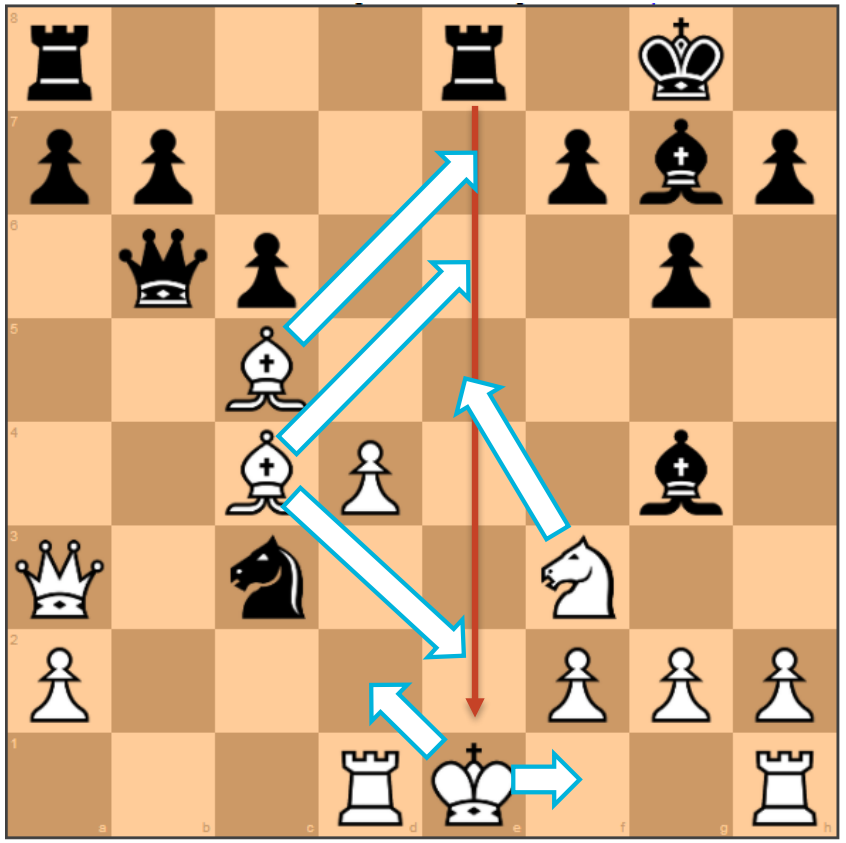

Fun With Postconditions

- **Postconditions tell you what will be true *after***
  - With preconditions, can *chain actions together* to plan for the future
  - Can evaluate *potential futures* to select the best immediate action

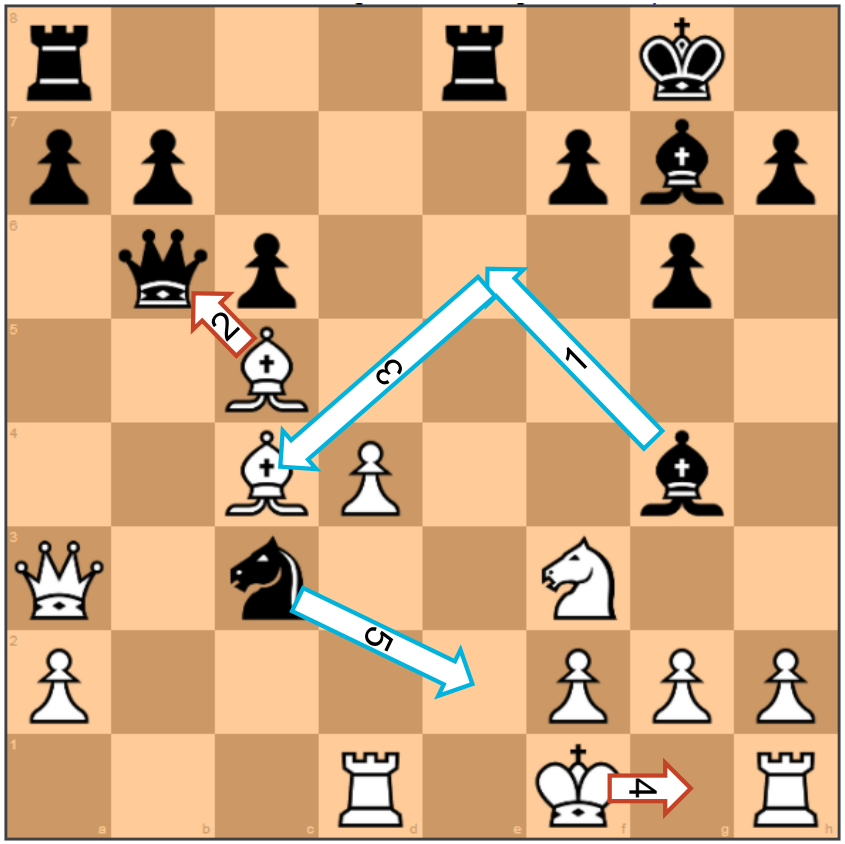

Fun With Postconditions

- **Postconditions tell you what will be true *after***
  - With preconditions, can *chain actions together* to plan for the future
  - Can evaluate *potential futures* to select the best immediate action

- **In the emulation sense: given an escalated foothold on a host, we can:**
  - Dump credentials and then laterally move
  - Add/modify registry keys and then dump credentials
  - Setup scheduled tasks and then add/modify registry keys
  - …

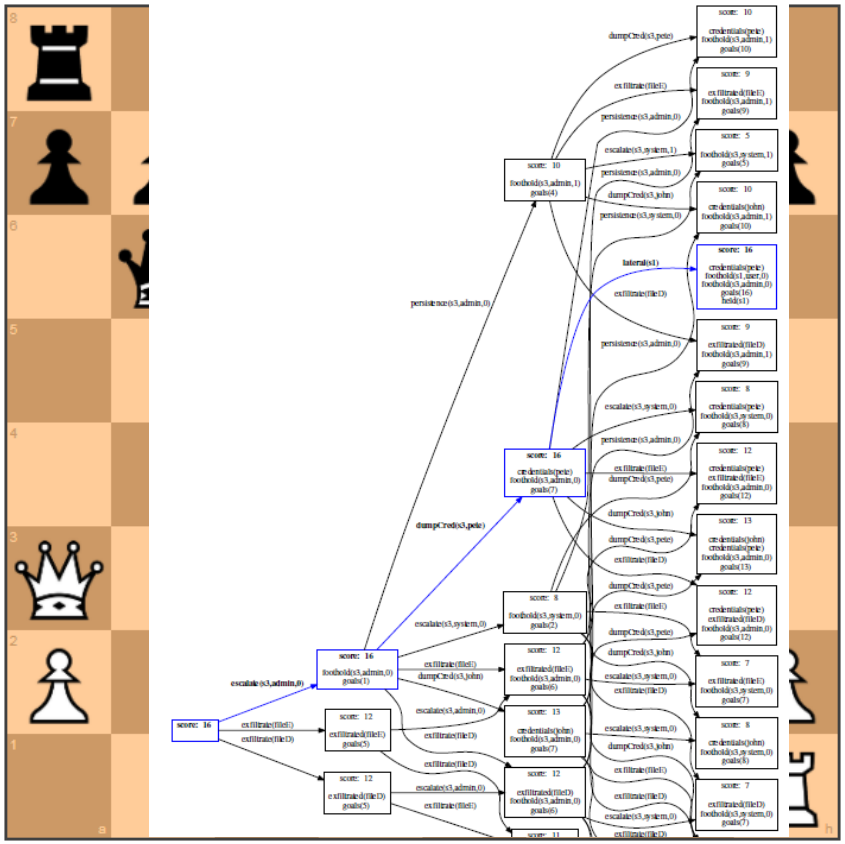

Making Progress

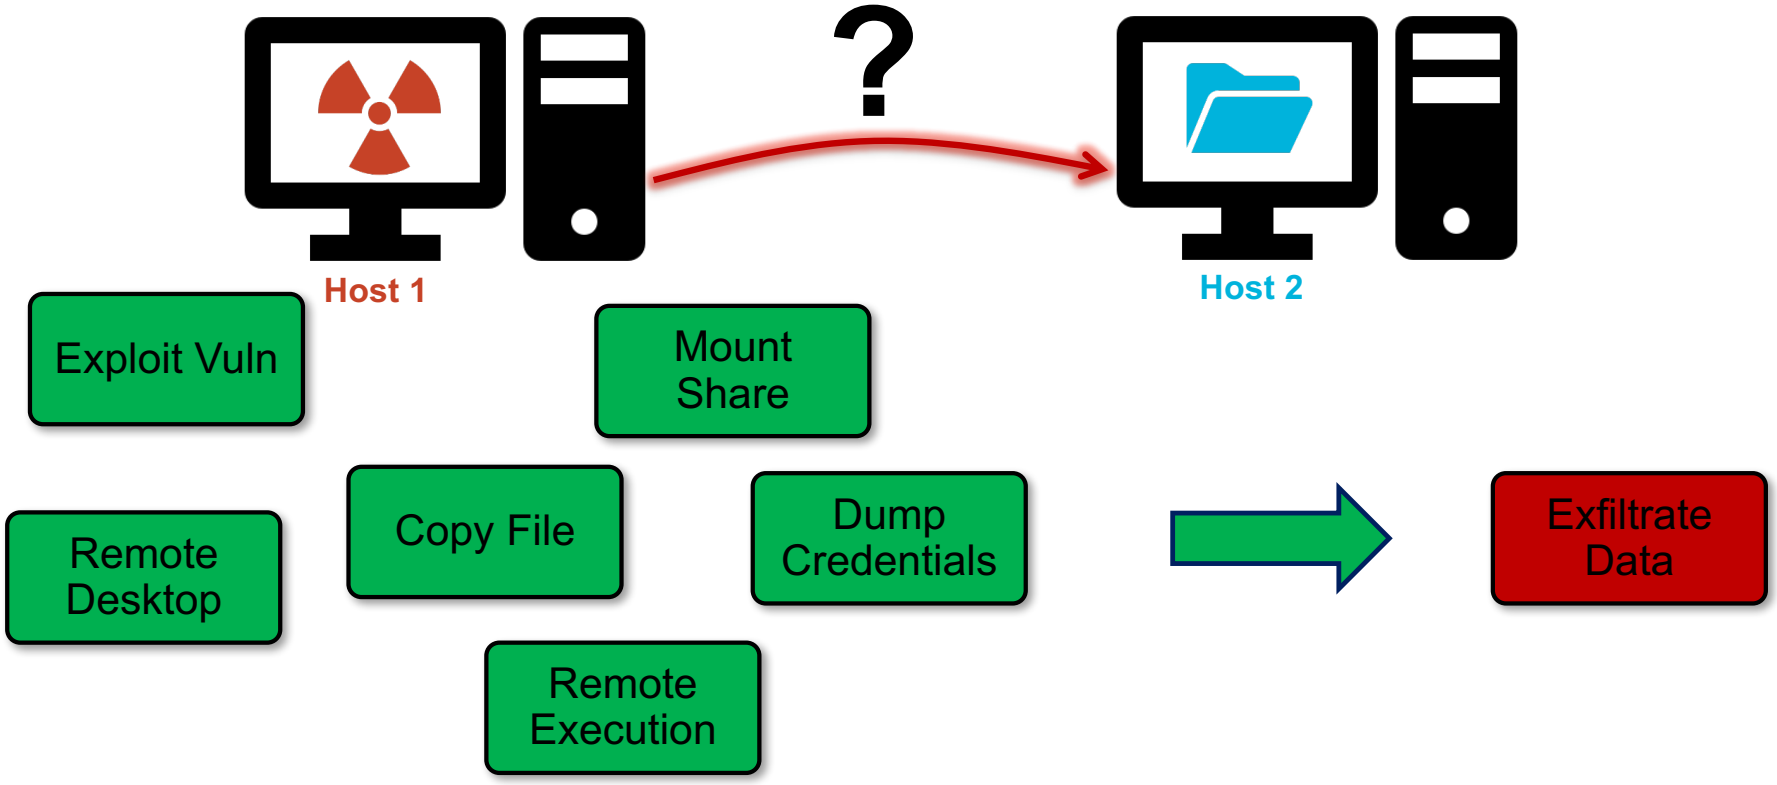

Building Plans: Copying a File

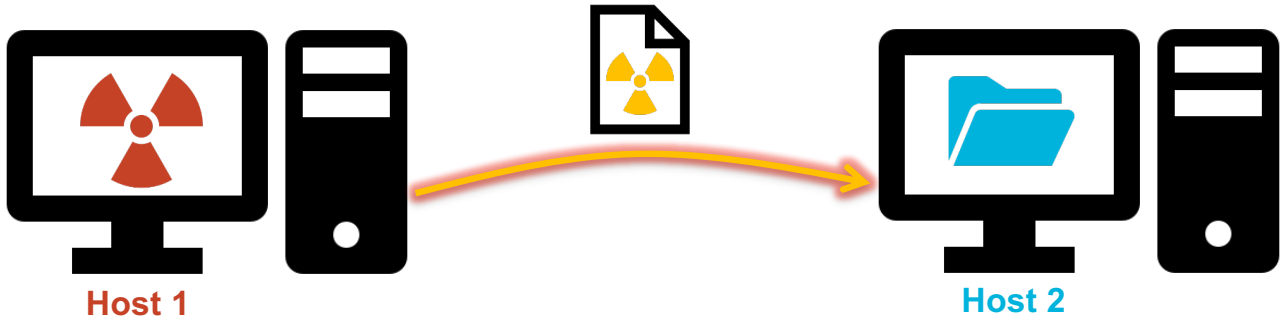

Dump Credentials → Mount Share → Copy File → Remote Execution → **Exfiltrate Data**

Building Plans: Exploiting a Vulnerability

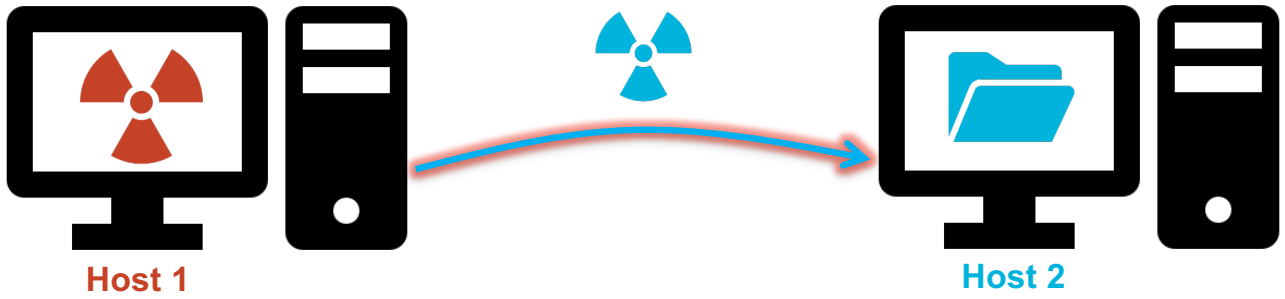

Building Plans: Remote Desktop Protocol

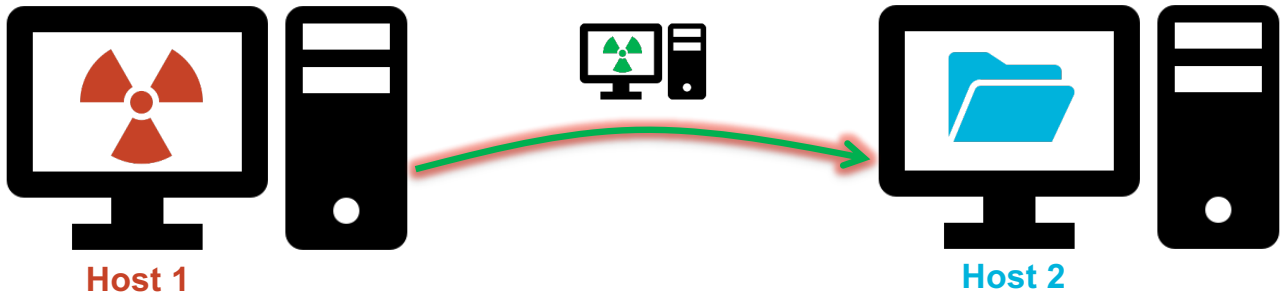

Dump Credentials → Mount Share → Copy File → Remote Execution → **Exfiltrate Data**

Exploit Vulnerability → **Exfiltrate Data**

Dump Credentials → Remote Desktop → **Exfiltrate Data**

Selecting the Right Plan

Host 1

Host 2

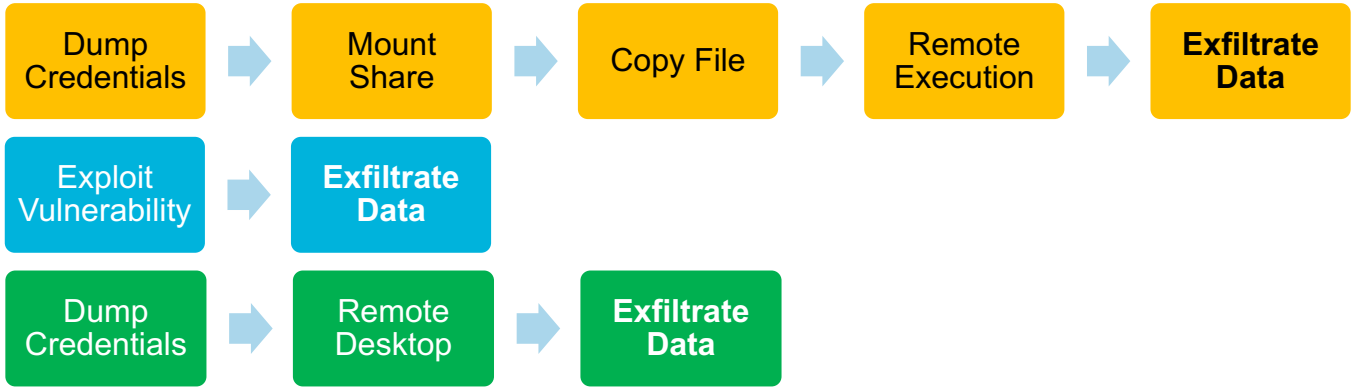

Selecting the Right Plan – The CALDERA Heuristic

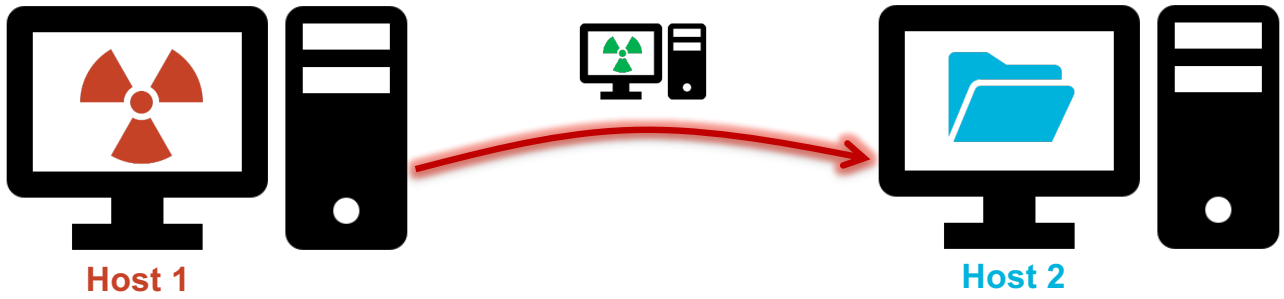

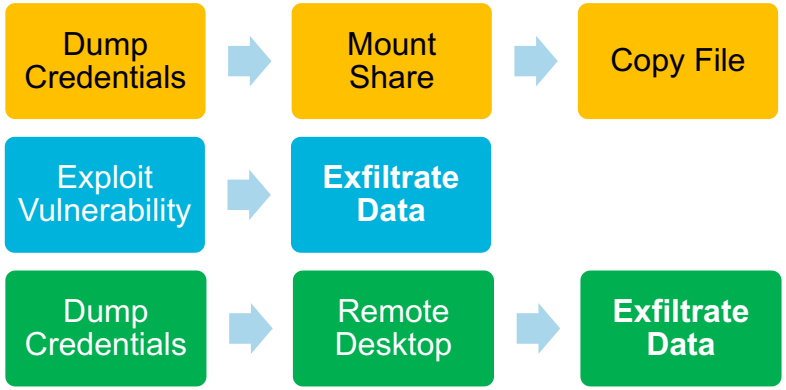

$$S(p) = \sum_{i=1}^{n} \frac{R(a_i)}{i}$$

- **Assign each action a *reward***
- **Score plans on summed *weighted rewards***

Action	Reward	Action	Reward
Exfiltrate Data	**100**	Copy File	5
Dump Credentials	*50*	Others	1

The Language of Pre/Postconditions: The Data Model

- **Need a way to logically encode what the pre and postconditions of techniques are**
  - Can specify requirements/consequences by specifying facts over a data model
- **CALDERA's language: *objects* and *typed fields***
  - Objects reference commonly used Windows components
  - Fields specify properties of objects, restricted by type
    - Constructed by default, some fields may not be defined
      - (this is important later!)

An Example Host Object

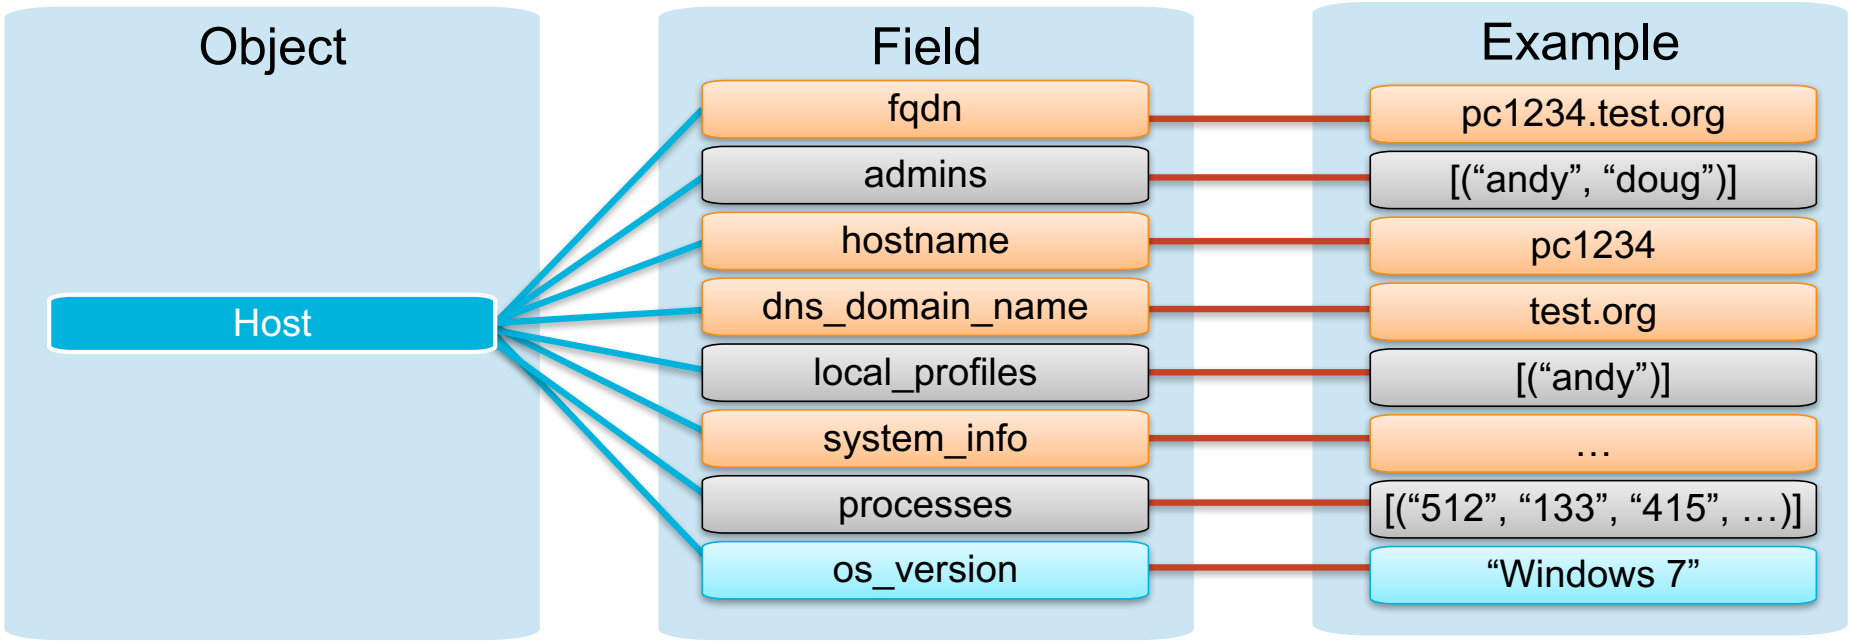

Object | String | Integer | DateTime | Boolean | Reference | List

Diving into the Data Model

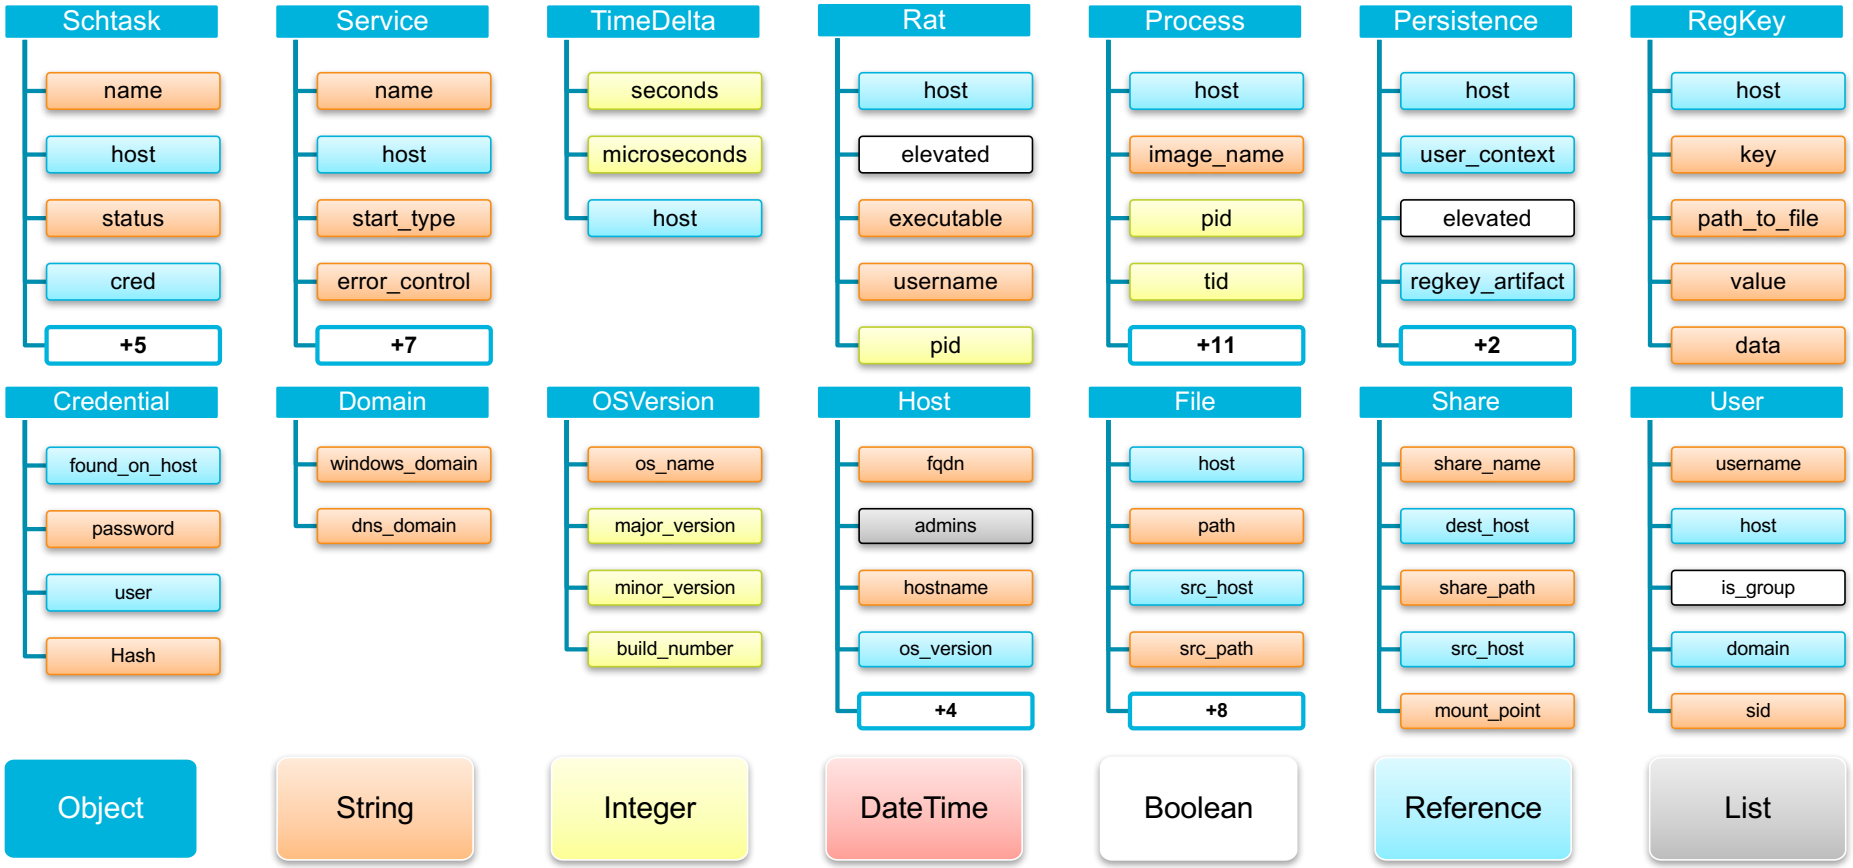

Declaring Actions

- **CALDERA provides a syntax to declare actions**
  - ***Preconditions*** specify the requirements
  - ***Postconditions*** specify the effects
  - ***Not_equal*** specifies inequality invariants
  - ***Preproperties*** specify that certain fields must be *defined* but not necessarily a specific value
  - ***Postproperties*** specify that certain fields will be defined after execution

Declaring Actions

- **CALDERA provides a syntax to declare actions**
  - ***P***
  - ***P***
  - ***N***
  - ***P***
  - ***P***

```
class NetUse(Step):
    value = 0
    preconditions = [("rat", OPRat),
                     ('host', OPHost),
                     ("cred", OPCredential({'$in': {'user': OPVar("host.admins")}})),
                     ('user', OPUser(OPVar("cred.user"))),
                     ('domain', OPDomain(OPVar("user.domain")))]

    postconditions = [('share_g', OPShare({"src_host": OPVar("rat.host"),
                                           "dest_host": OPVar("host"),
                                           'share_name': 'C$'}))]

    not_equal = [('host', 'rat.host')]

    preproperties = ['domain.windows_domain', 'cred.password', 'host.fqdn', 'user.username']
    postproperties = ["share_g.share_path", "share_g.mount_point"]

    deterministic = True
```

c value

- **Easy, right?**

CALDERA's Planning Algorithm

1. Update the world state
2. Figure out all valid actions to execute
3. Construct plans that lead off with those actions
   - Chain actions together by leveraging model
4. Run heuristic to determine *best* plan
5. Execute the first action in the best plan
6. Repeat

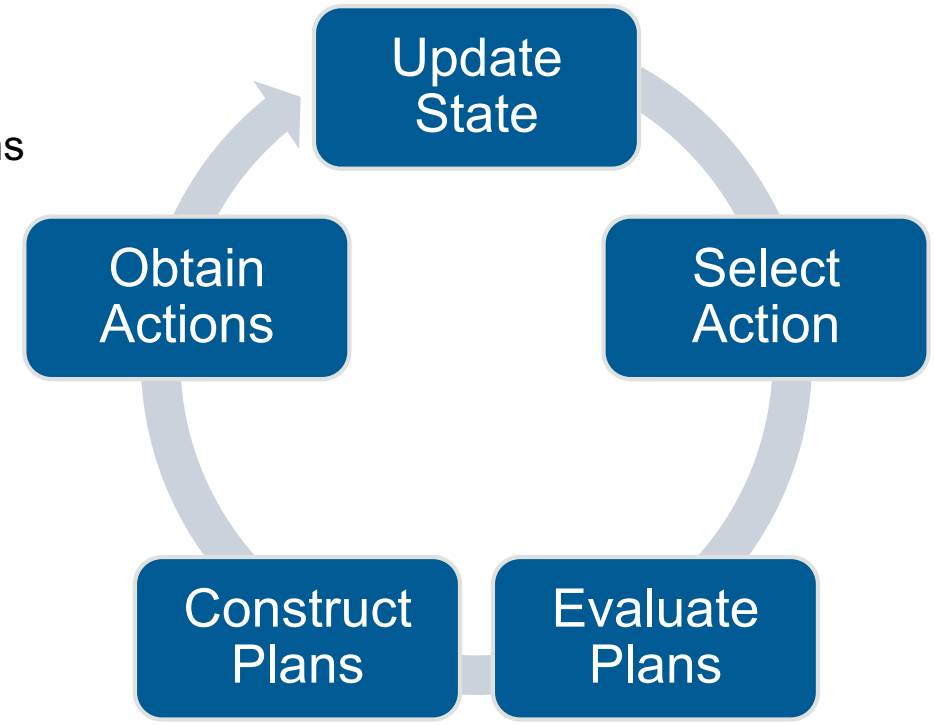

And we're done! …Right?

- **Using pre and postconditions we can dynamically construct plans and choose actions**
- **Meets our goals:**
  - Can construct plans to make intelligent choices
  - Can easily incorporate new techniques by defining pre/postconditions
  - Can modify rewards/included actions to vary operations
  - Chains weaknesses to achieve goals
  - Functions in new environments*
- **Wait – functions in new environments?**
  - Actually, maybe not: when constructing plans, there is a *significant* amount of uncertainty!
  - Consider dumping credentials:
    - Sometimes they're *great creds*
    - Sometimes there's none
- **In reality – handling uncertainty is a *very* hard problem!**

And we're done! …Right?

- **Using pre and postconditions we ca… actions**
- **Meets our goals:**
  - Can construct plans to make intelligent ch…
  - Can easily incorporate new techniques by…
  - Can modify rewards/included actions to va…
  - Chains weaknesses to achieve goals
  - Functions in new environments*
- **Wait – functions in new environment…**
  - Actually, maybe not: when constructing pla…
  - Consider dumping credentials:
    - Sometimes they're *great creds*
    - Sometimes there's none
- **In reality – handling uncertainty is a *very…***

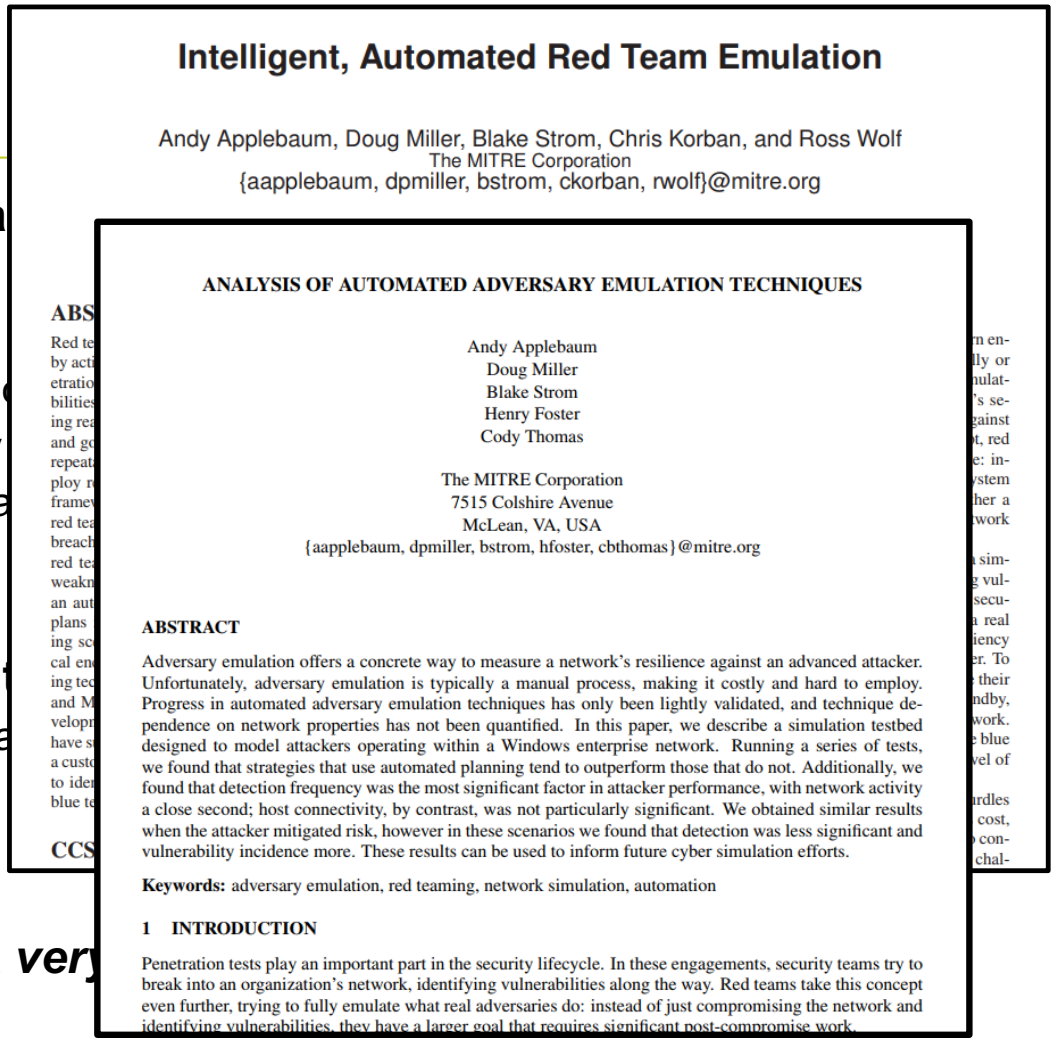

Intelligent, Automated Red Team Emulation

Andy Applebaum, Doug Miller, Blake Strom, Chris Korban, and Ross Wolf
The MITRE Corporation
{aapplebaum, dpmiller, bstrom, ckorban, rwolf}@mitre.org

ANALYSIS OF AUTOMATED ADVERSARY EMULATION TECHNIQUES

Andy Applebaum
Doug Miller
Blake Strom
Henry Foster
Cody Thomas

The MITRE Corporation
7515 Colshire Avenue
McLean, VA, USA
{aapplebaum, dpmiller, bstrom, hfoster, cbthomas}@mitre.org

ABSTRACT

Adversary emulation offers a concrete way to measure a network's resilience against an advanced attacker. Unfortunately, adversary emulation is typically a manual process, making it costly and hard to employ. Progress in automated adversary emulation techniques has only been lightly validated, and technique dependence on network properties has not been quantified. In this paper, we describe a simulation testbed designed to model attackers operating within a Windows enterprise network. Running a series of tests, we found that strategies that use automated planning tend to outperform those that do not. Additionally, we found that detection frequency was the most significant factor in attacker performance, with network activity a close second; host connectivity, by contrast, was not particularly significant. We obtained similar results when the attacker mitigated risk, however in these scenarios we found that detection was less significant and vulnerability incidence more. These results can be used to inform future cyber simulation efforts.

Keywords: adversary emulation, red teaming, network simulation, automation

1 INTRODUCTION

Penetration tests play an important part in the security lifecycle. In these engagements, security teams try to break into an organization's network, identifying vulnerabilities along the way. Red teams take this concept even further, trying to fully emulate what real adversaries do: instead of just compromising the network and identifying vulnerabilities, they have a larger goal that requires significant post-compromise work.

And we're done! …Right?

- **Using pr** … **ns and choose actions**
- **Meets ou**
  - Can cor
  - Can eas
  - Can mo
  - Chains
  - Functio
- **Wait – fu**
  - Actually … uncertainty!
  - Conside
    - Somet
    - Somet

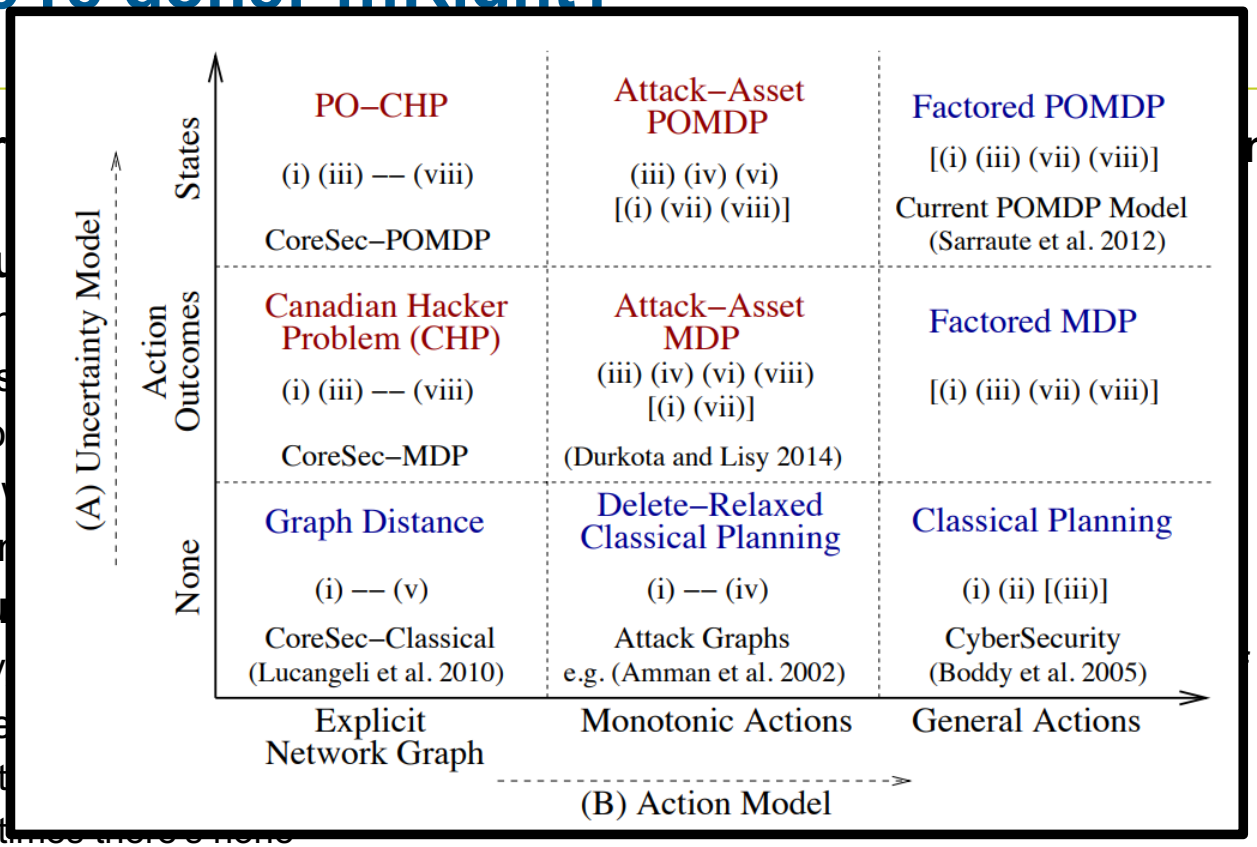

- **Hoffmann, Jörg. "Simulated Penetration Testing: From" Dijkstra" to" Turing Test++"." *ICAPS*. 2015.**

A Quick Fix with Hints

- **If we can't predict the outcome of an action, use hints**
- **Hints are crafted to be the "best" outcome of the action**
  - E.g. performing credential dumping gives me a "useful" credential

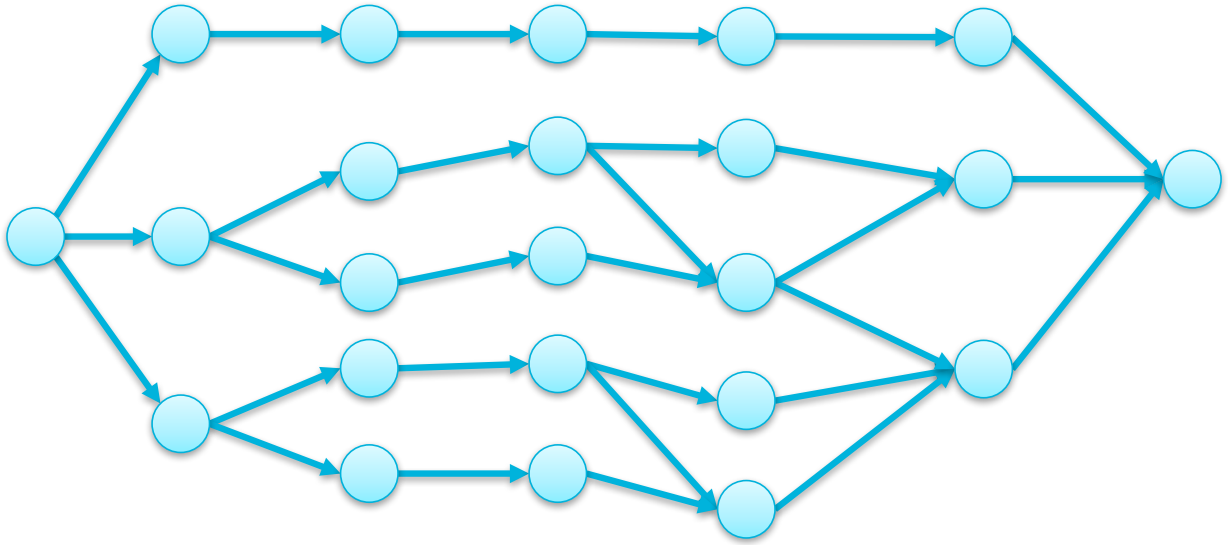

Architecture

System Architecture

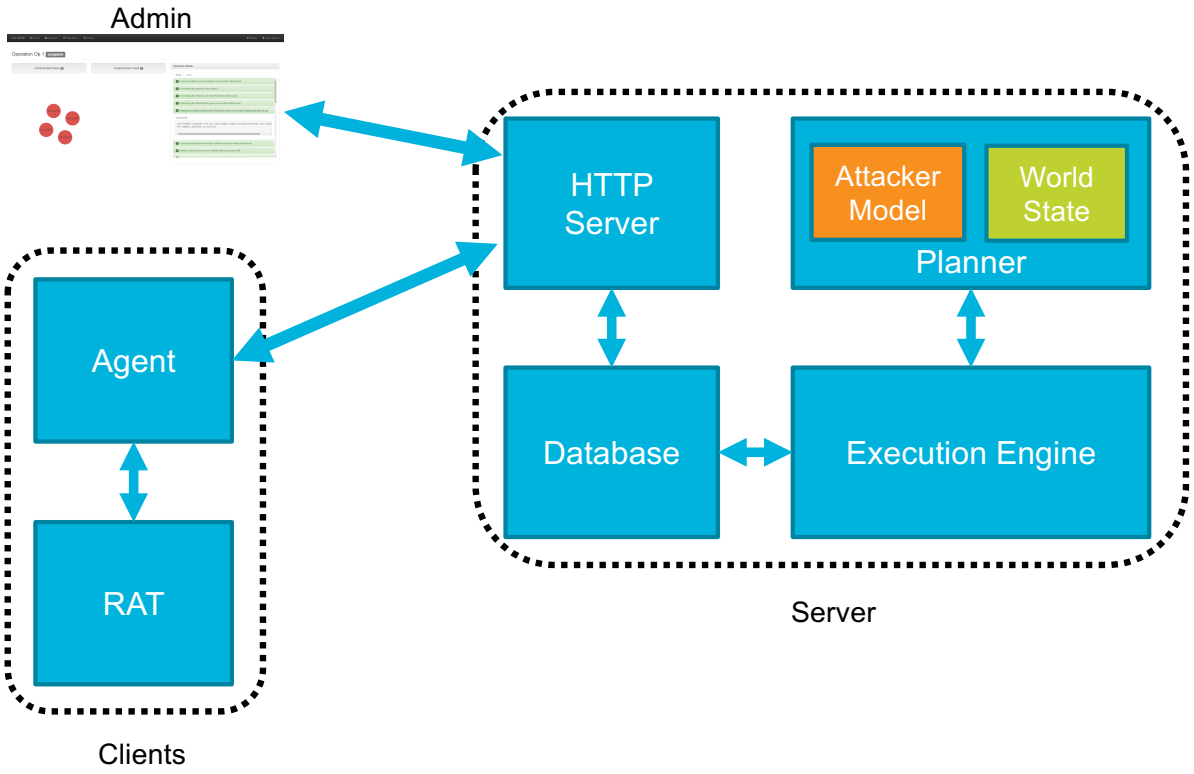

System Architecture

- Server and Agent written in Python 3
- Rat written in C#
- MongoDB
- Web interface is a JavaScript based web app
- pyDatalog logic backend

Demo

Walking Through an Operation

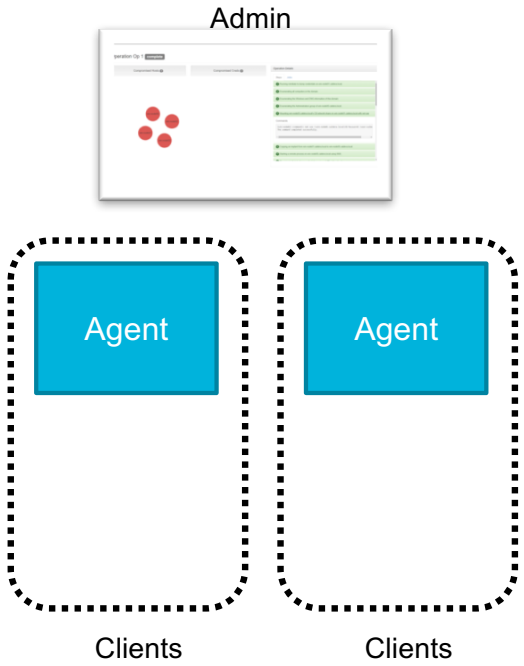

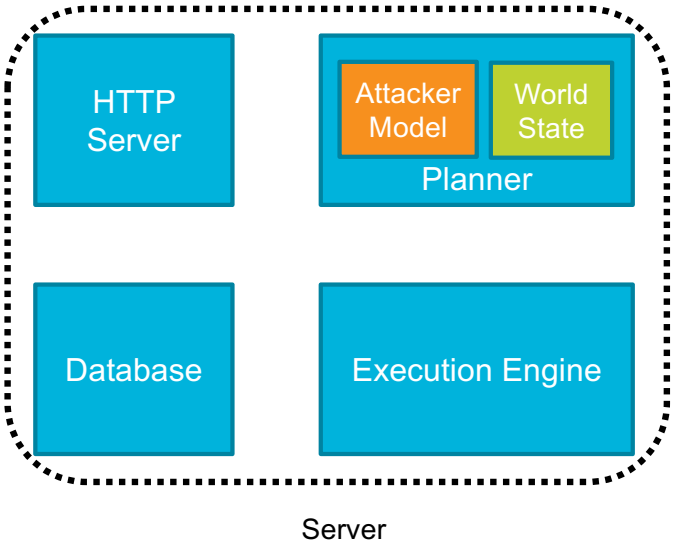

Walking Through an Operation

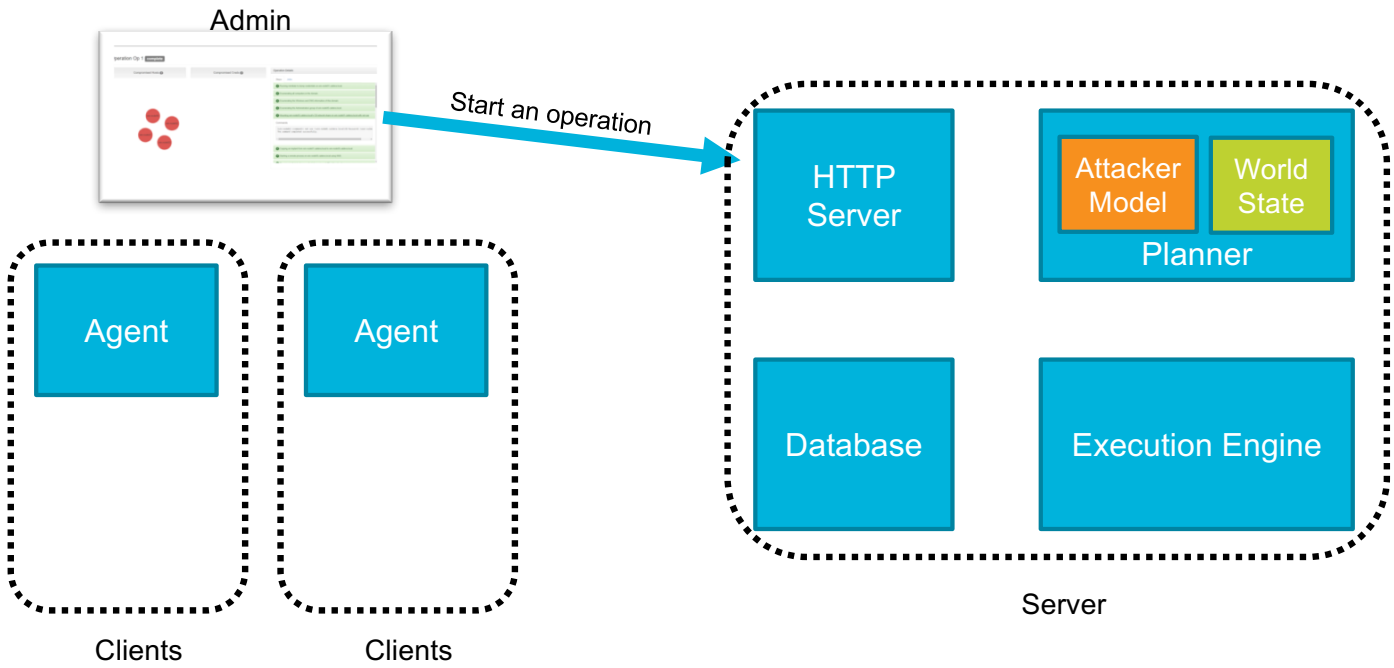

Walking Through an Operation

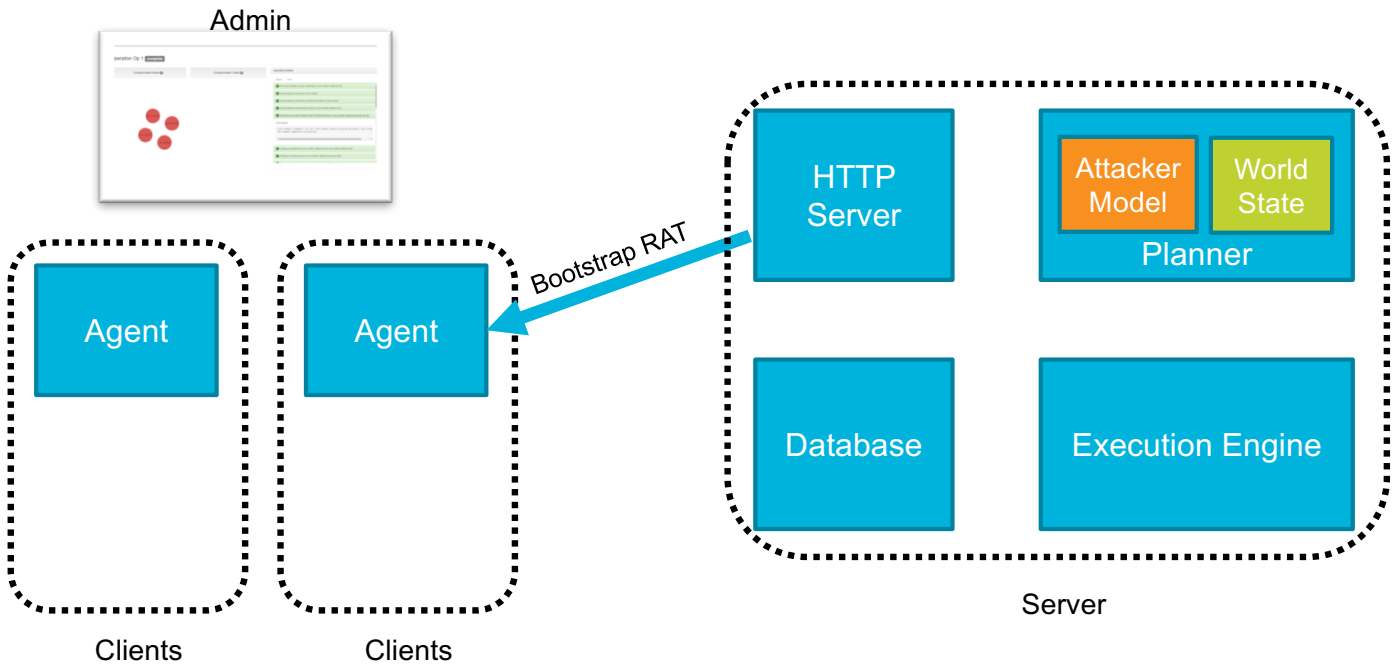

Walking Through an Operation

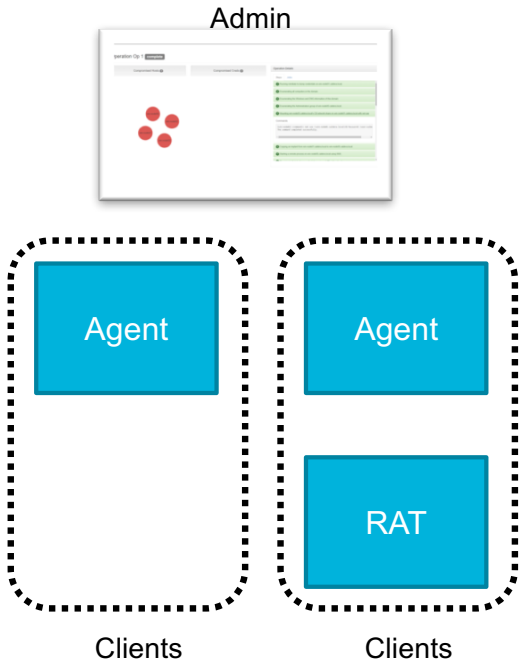

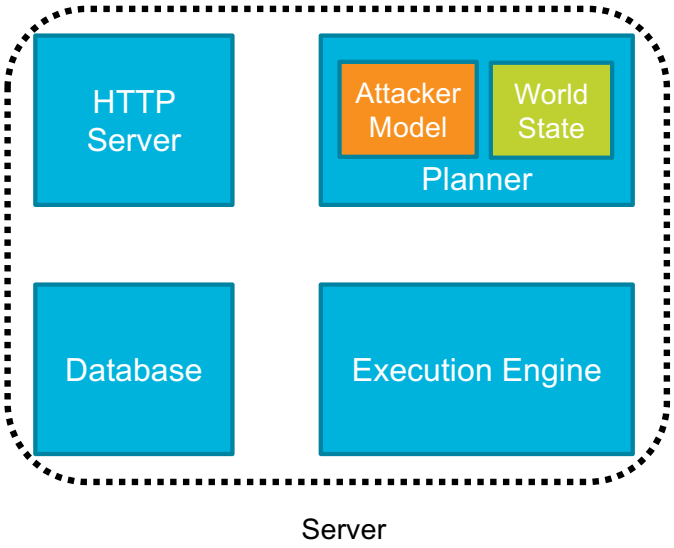

Walking Through an Operation

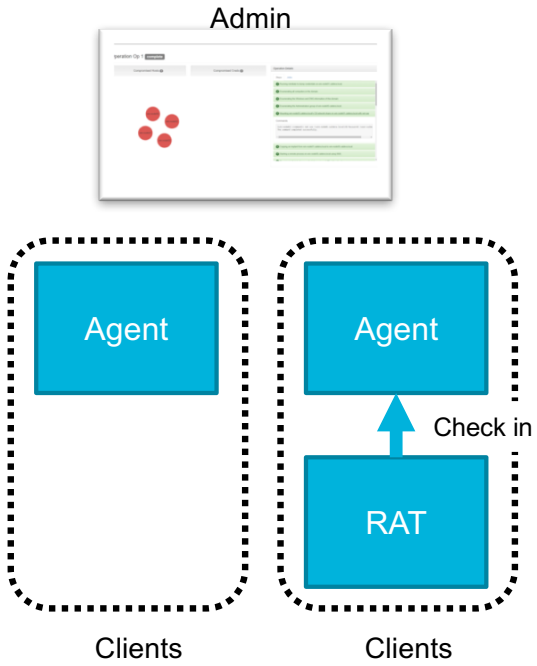

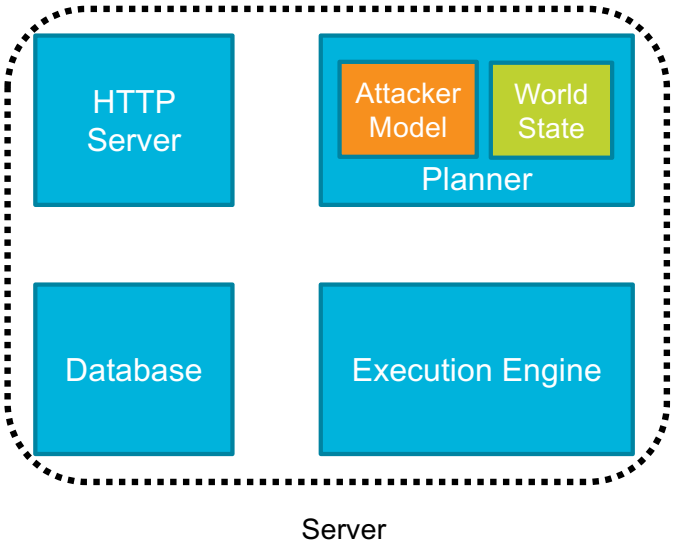

Walking Through an Operation

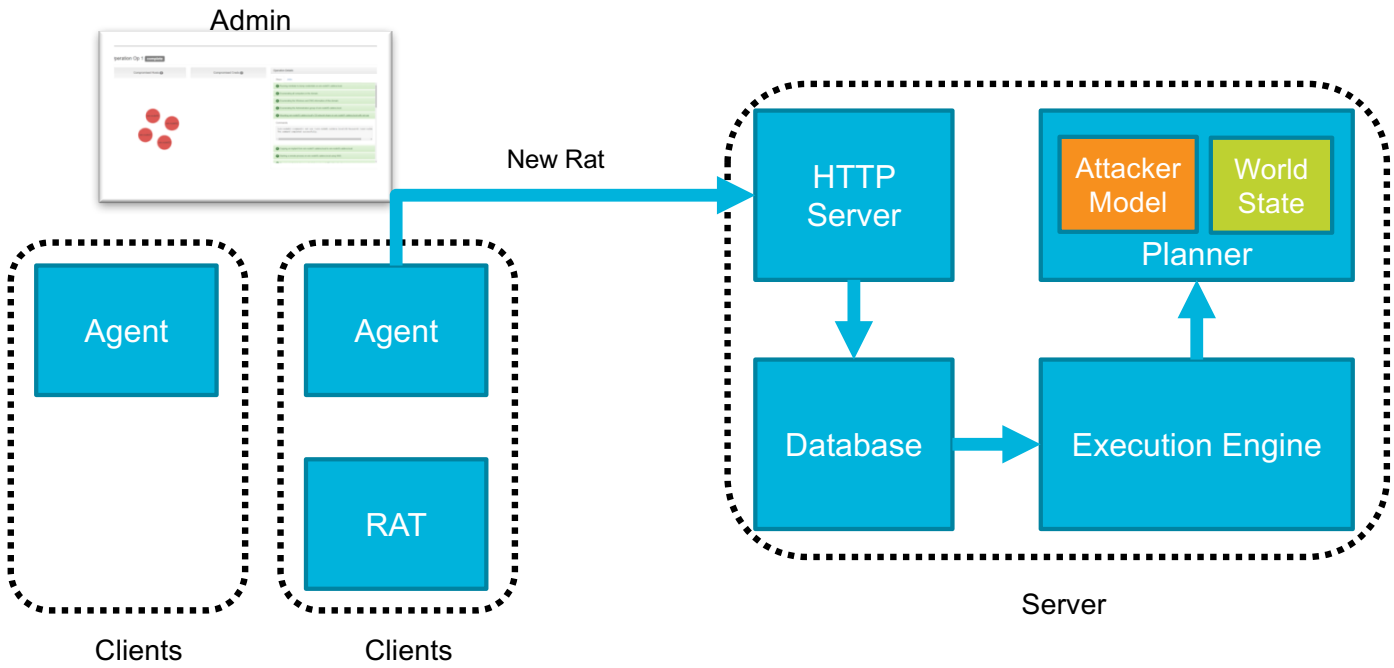

Walking Through an Operation

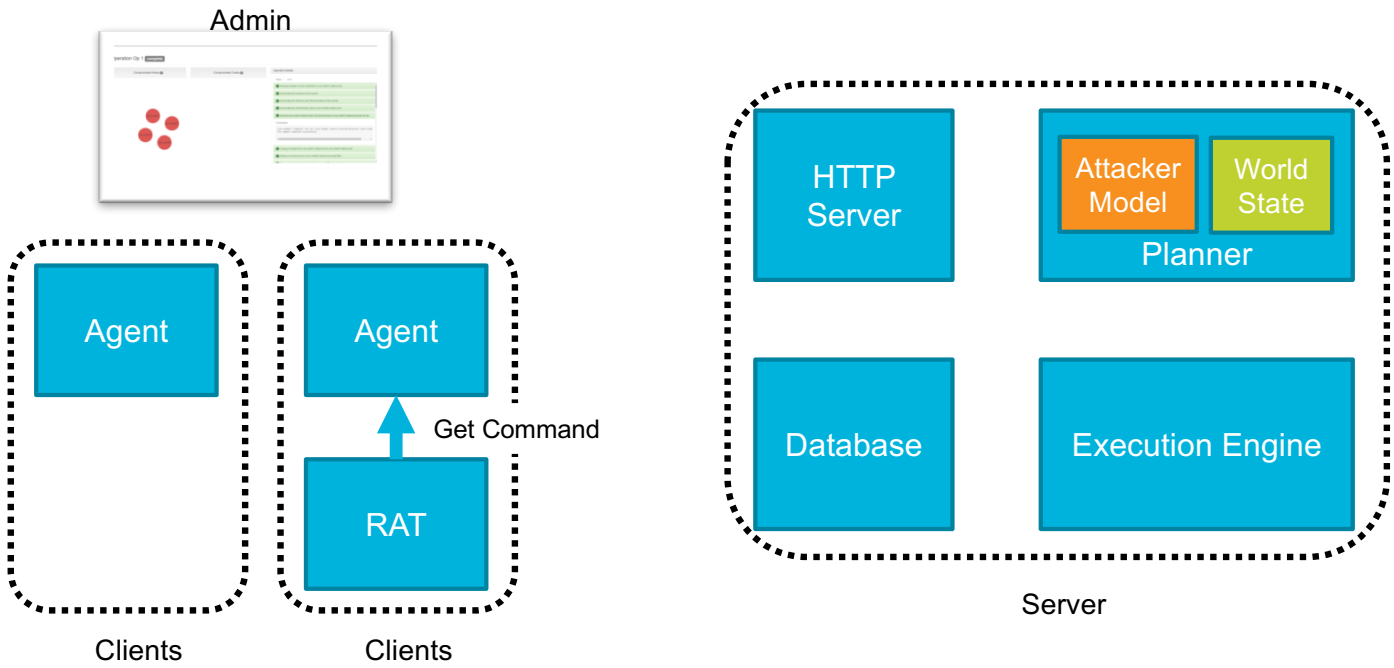

Walking Through an Operation

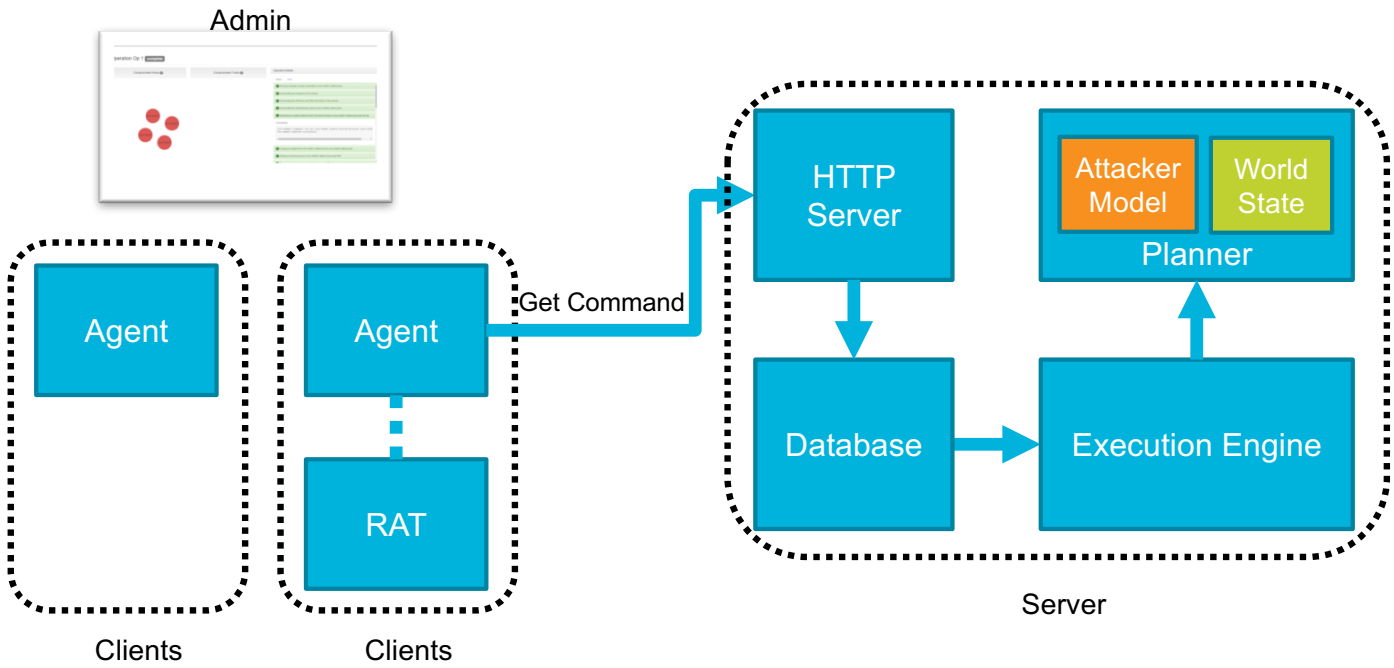

Walking Through an Operation

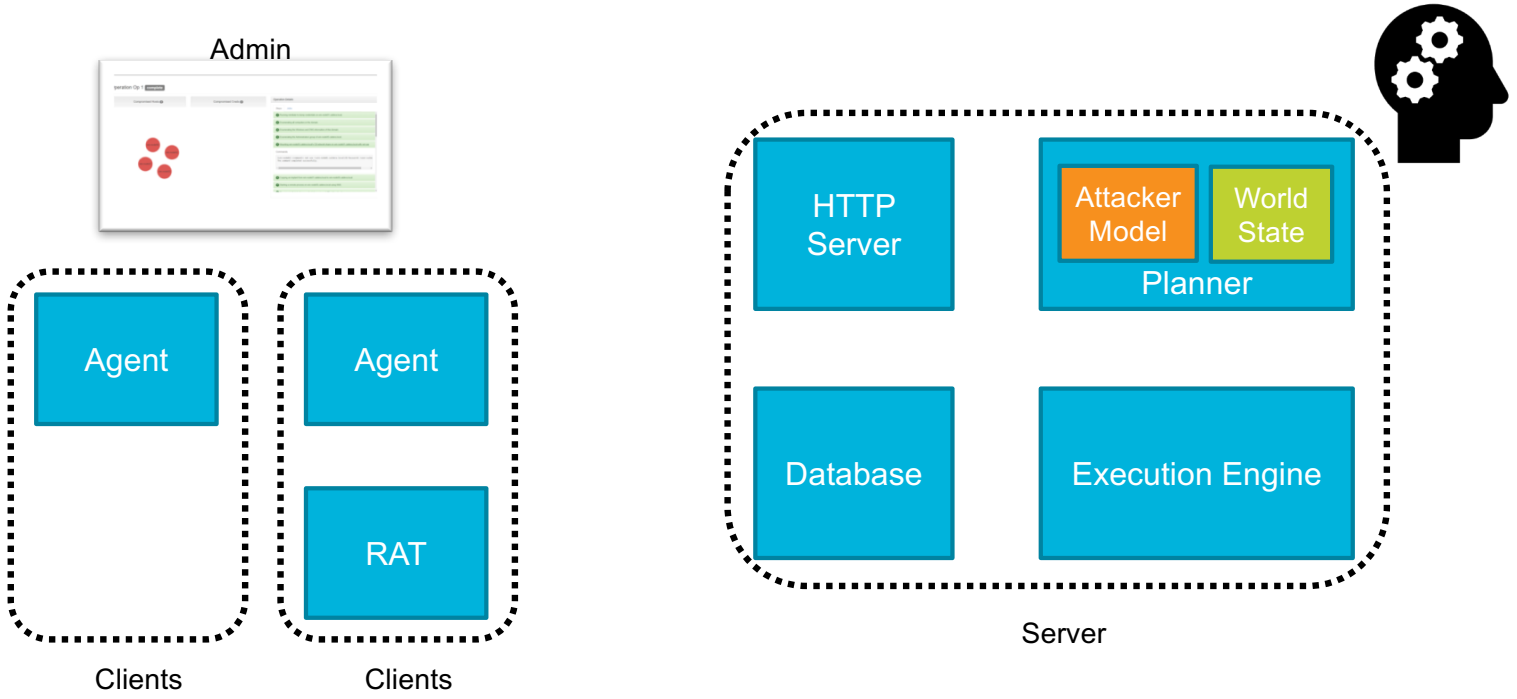

Walking Through an Operation

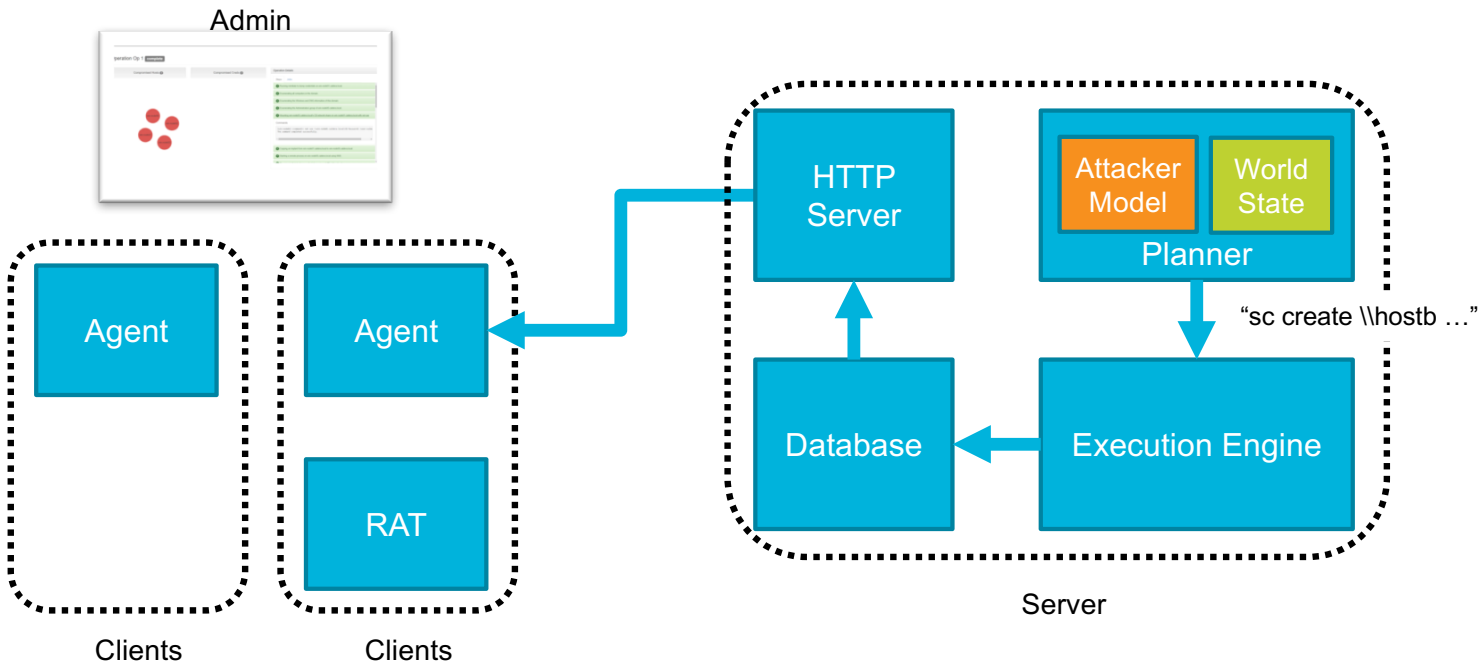

Walking Through an Operation

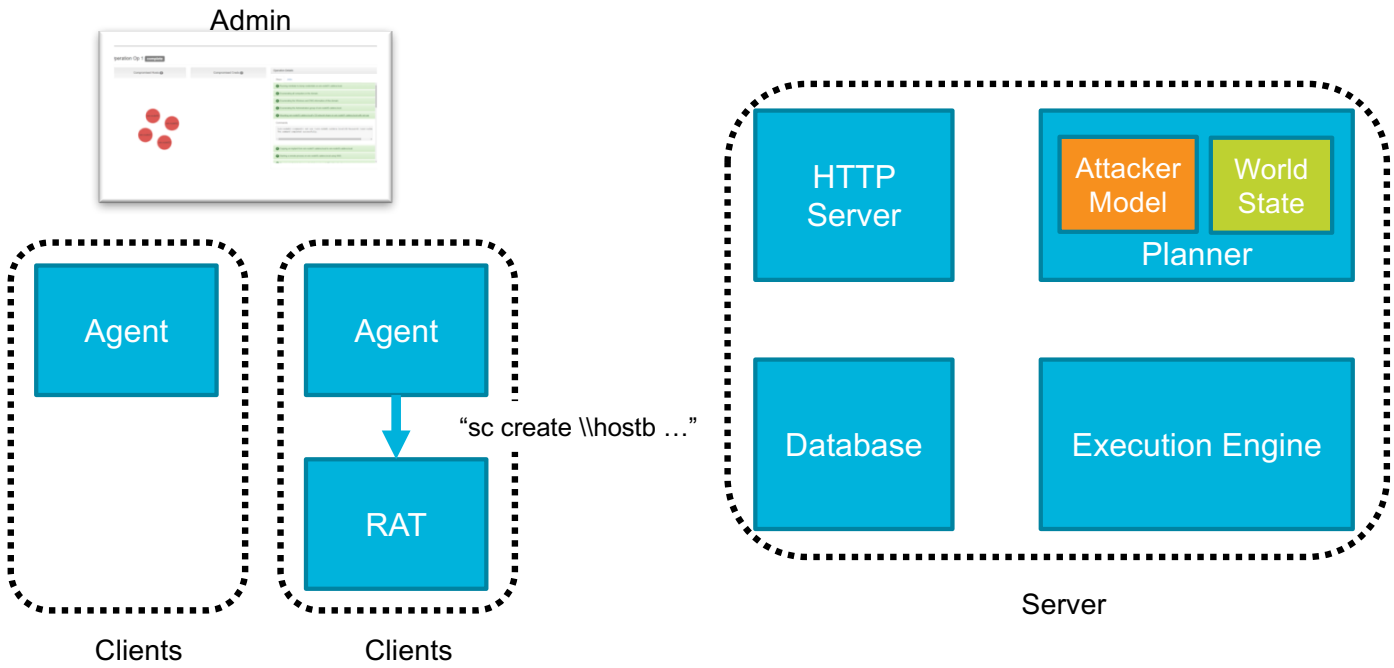

Walking Through an Operation

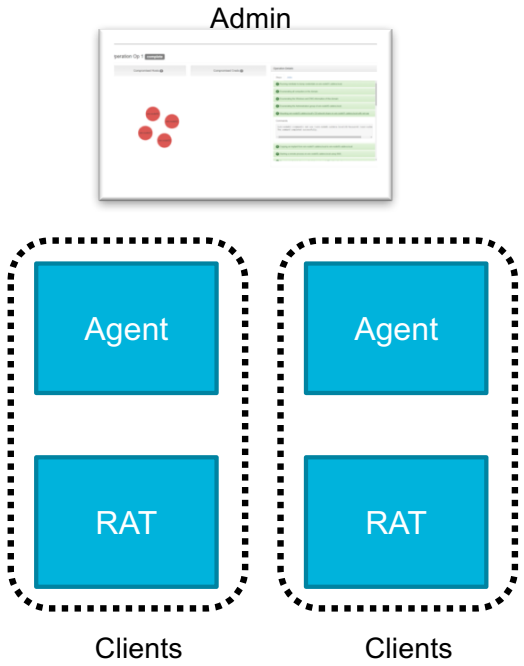

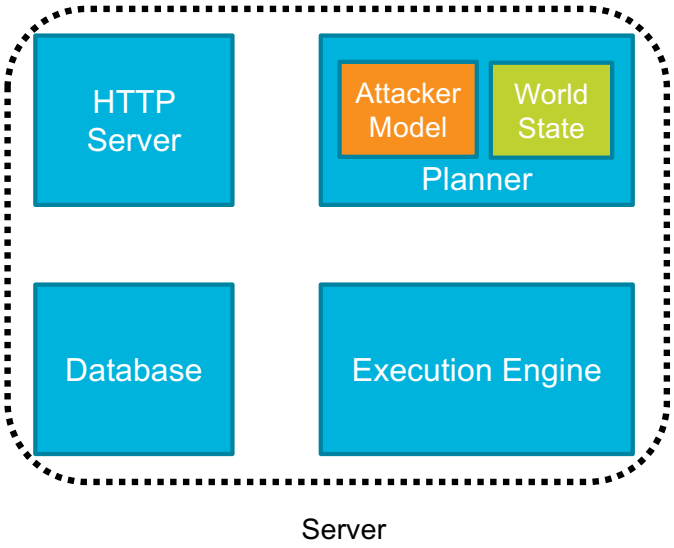

Walking Through an Operation

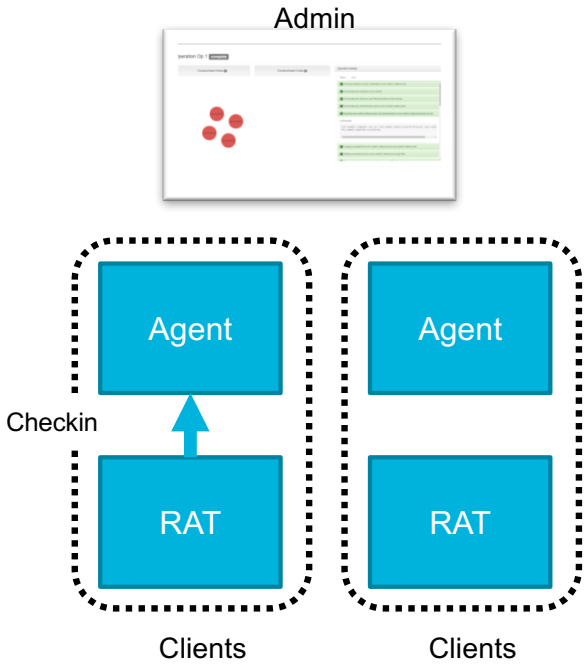

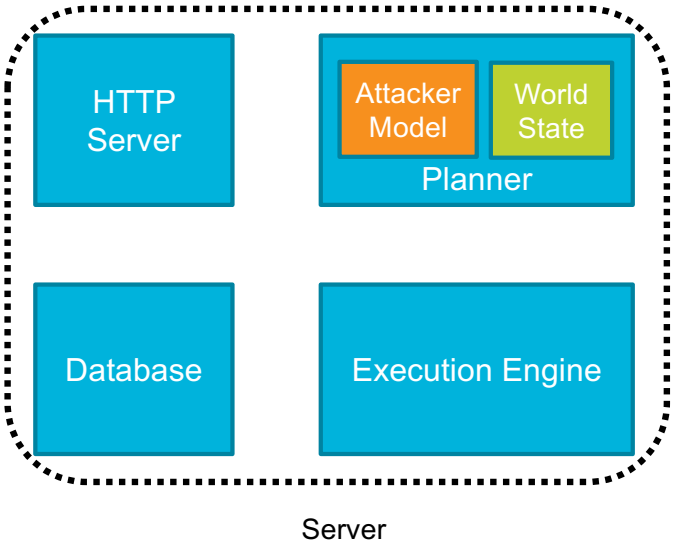

Walking Through an Operation

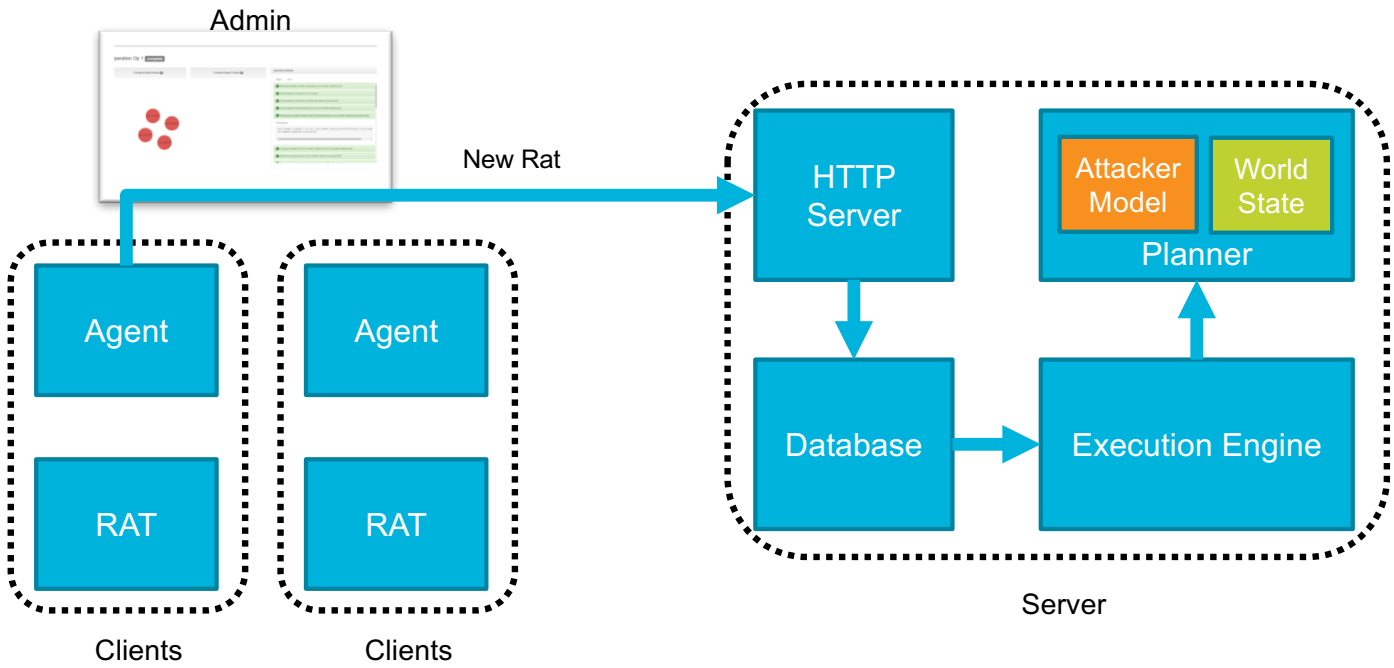

CALDERA Applications

- Testing analytics and sensors – Does my stuff work?
- Data generation – What does bad look like?
- Red/blue team training – I need practice!

Community Participation

- **Want to get involved with CALDERA? We accept Pull Requests**
  - Bugfixes
  - Implement a new adversary (ATT&CK) technique
  - Usability features
  - Integration with other tools and frameworks
  - Enhancements to our data model

Other (Free) Tools

BloodHound – Attack Path Generation

- https://github.com/BloodHoundAD/BloodHound

GoFetch – Automatic Execution of BloodHound paths with PowerShell tools

- https://github.com/GoFetchAD/GoFetch

ANGRYPUPPY – Automatic execution of BloodHound paths with Cobalt Strike

- https://github.com/vysec/ANGRYPUPPY

Death Star – Automatic Execution of attack paths with PowerShell Empire

- https://github.com/byt3bl33d3r/DeathStar

Atomic Red Team

- https://github.com/redcanaryco/atomic-red-team

Metta

- **<no url yet>**

(Probably more, sorry if we missed you)

Related (MITRE) Efforts

- **BRAWL: Automated Bot-vs-Bot Games**
  - Free data!
  - https://github.com/mitre/brawl-public-game-001
- **BRAWL Shared Format (BSF)**
  - Standardized format to correlate red bot vs blue bot cyber games
- **CASCADE: Automated Host-based Investigations**
  - https://github.com/mitre/cascade-server

Why this Matters

- The False Negative problem is real
- Offensive testing with Adversary Emulation can help
- Automation (like CALDERA) and human adversary emulation are complementary
- Pre and postconditions + planning are powerful
- Help Us!

Show me the code!

MITRE

MITRE is a not-for-profit organization whose sole focus is to operate federally funded research and development centers, or FFRDCs. Independent and objective, we take on some of our nation's—and the world's—most critical challenges and provide innovative, practical solutions.

Learn and share more about MITRE, FFRDCs, and our unique value at www.mitre.org