Carnegie Mellon University
Software Engineering Institute

Foundation of Cyber Ranges

Thomas G. Podnar
Geoffrey B. Dobson
Dustin D. Updyke
William E. Reed

May 2021

TECHNICAL REPORT
CMU/SEI-2021-TR-001
DOI: 10.1184/R1/13557566

CERT Division

http://www.sei.cmu.edu

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

This report was prepared for the SEI Administrative Agent AFLCMC/AZS 5 Eglin Street Hanscom AFB, MA 01731-2100

DM21-0209

Table of Contents

List of Tables

Abstract

This report details the design considerations and execution plan for building high-fidelity, realistic virtual cyber ranges that deliver maximum training and exercise value for participants. A cyber range is a fully interactive virtual instance of the IT infrastructure of a mid-to enterprise-level organization that is dedicated to cyberwarfare training and exercise. The Software Engineering Institute CERT Cyber Workforce Development (SEI/CERT/CWD) team at Carnegie Mellon University has provided a wide array of cyber ranges spanning various levels of complexity to the Department of Defense (DoD) since 2012. This report draws on tens of thousands of hours of experience and expertise in designing, developing, deploying, troubleshooting, maintaining, and improving cyber ranges. The most frequent use case for a cyber range is team exercise. Team exercises can focus on incident response, malware removal, network takeover, process improvement, and many other specialized use cases. This report summarizes the SEI/CERT/CWD current state of the art when it comes to delivering high-fidelity/realistic cyber ranges to DoD and non-DoD customers that support the complex operating environments required for their training and exercise purposes.

1 Introduction

Over the past decade, the Software Engineering Institute CERT Division at Carnegie Mellon University has conducted hundreds of successful high-fidelity team-based cyber range exercise events. Based on our research and first-hand experience, we strongly advocate the approach outlined in this report as an effective way to supplement and grow an organization's team-based cyber warfighting capability.

The SEI's CERT Division has a long history of literature that clearly delineates the value of team-based cyber exercises [Hammerstein 2010], along with a framework for providing optimal fidelity and realism [Dobson 2017]—as well as a technical report regarding supplemental cyber range technical tools [Updyke 2018].

Carnegie Mellon research teams are actively focused on improving all aspects of high-fidelity team-based cyber range exercises. Participation in realistic cyber range team-based exercises is critical to producing cyber operators capable of winning within the cyber warfare arena. Practicing cyber operations in a realistic range is particularly important given that adversaries are actively exercising their craft in real-world, production environments every day.

At Carnegie Mellon, we have successfully architected and implemented cyber ranges with thousands of virtual machines and a corresponding number of participants. The remainder of this report details a framework approach for each of the above considerations. While we largely focus on large-scale cyber ranges and how the above considerations influence exercise architecture and execution, we suggest that the framework effectively supports smaller engagements as well. The remainder of this report is organized into four parts: Section 1 summarizes the five key prerequisite organizational understandings necessary to successfully develop and deploy a complex cyber range. Section 2 details the core infrastructure requirements and present key planning considerations. Section 3 focuses on cyber range system architecture. Section 4 navigates the cyber range implementation process. In Section 5, implementation skill sets are summarized.

1.1 Cyber Range Prerequisite Understanding

The challenges of design, implementation, execution, and overall experience within an exercise will vary widely, based on the intended scenario, and is in many ways dependent upon the size and scope of the supporting cyber range. Despite this wide array of factors, it is advisable to review the following considerations prior to the design, implementation, and exercise execution phases of a cyber range:

- Use Cases - Define use-case scenarios from all participant perspectives (training audience, range maintenance, leadership, exercise engineers, role players).
- Goals - Set goals by clearly identifying the criteria for evaluating impact, effectiveness, and success—from a training and exercise perspective. Set goals for how the individuals and/or team will be assessed throughout various engagements.
- Technical Requirements - Document the architecture and specific technical requirements and the expected end state for the range and its artifacts post-event.

- Cost Analysis - Distinguish funding sources for initial implementation, ongoing maintenance, and exercise execution.
- Human Resources - Identify the availability of required technical and design skills and how to supplement them for shortfalls. Allocate staff and define roles throughout the design, implementation, and execution stages of the event. Be mindful of adequate support during exercise execution.

2 Core Infrastructure Foundations

Properly sizing the core technical infrastructure foundational layer is the most critical component of the entire cyber range architecture. A well architected solution will provide a realistic, highly consistent, and performant experience for exercise participants.

The majority of cyber ranges will be implemented using virtualization technologies. Therefore, the technical infrastructure foundation for a cyber range is very similar to the infrastructure foundation that would exist in mid-size to larger enterprise production data centers, for server virtualization capabilities and/or alignment with widely available cloud technologies.

2.1 Core Infrastructure Components

This description of the core infrastructure components includes explanation of requirements and planning considerations.

- Compute Power - Foundational physical server hardware resources must be able to adequately support the deployment and simultaneous load of hundreds to thousands of virtual machines from a CPU and memory perspective.
- Storage – Storage must be capable of providing consistent, low-latency I/O for the above compute power requirements, while at the same time providing tens of terabytes of long-term raw storage.
- Network - Network devices must be capable of providing high-bandwidth, low-latency connectivity to service the load produced by the above two components as well as meeting the added network complexity requirements of a cyber range implementation.
- Network Security and Isolation - Similar to a Test/QA environment, the operations of the compute power, storage, and network should be isolated and never impact an organization's production services. Network isolation is specifically important since the cyber range may contain malware or conduct malicious activities, as part of the cyber exercises, that would not be compatible with exposure to the production network.
- Configuration Management and Automation - Large-scale architecture component deployments must be automated and controlled via sets of pre-built configuration tools (i.e., "infrastructure as code" [IaC]).
- Backups - The cyber range architecture and configurations will likely be very unique. This environment should be protected with an effective backup strategy.
- "Easy Button" Style Resets - The cyber range architecture should be such that it can be quickly reverted or reset back to a baseline configuration. A single cyber warfare exercise can modify a cyber range in an undesirable manner or leave unintended artifacts behind. In order to maximize re-use of the range and limit its downtime, technical decisions that allow for large-scale resets are paramount.

3 Cyber Range Systems Architecture

The most effective cyber ranges are founded upon a realistic, but properly scaled instance of an organization's production infrastructure architecture. These architectural alignments support the proven training and exercise model of "train as you fight" [Dobson 2017]. An organization's technical teams will be able to leverage their existing experience in implementing the cyber range using this scaled replica concept. Similar "tribal knowledge" skill sets will also be applicable to the implementation of cyber range end-user workstations for Non Player Character roles (NPCs) and associated workstation security-based policies.

Beyond the initial requirements of the technical infrastructure foundational layer, the design requirements of a cyber range will have divergence compared to any other typical organizational infrastructure project. Likewise, technical team architecture and implementation skill sets will have more gaps and will be more challenged because of the divergence.

One example of the divergence is that an organization's production architecture will likely be at least partially reliant on public Internet connectivity. Therefore, replicating some portion of a "public Internet" in the cyber range will be beneficial for advanced cyber-warfare scenarios. The architecture and design of replicating some portions of the Internet, to include in the cyber range, will likely be a unique challenge for the technical implementation teams.

3.1 Zones and Teams Defined

Initially, as part of the United States Computer Security Defense Initiative, *blue teams* were defined as defensive counterparts to red "tiger teams" assessing systems within an organization. Even today, in United States warfighting simulation, U.S. forces are still considered blue, and the opposing force ("OPFOR") is red.

Red teams are concerned with the effective testing of an organization's overall security posture. They use the most up-to-date tools and emulate current adversary tactics, techniques, and procedures (TTP) models with their associated range of cyber threats, from simple script-based automation tools, all the way to nation state sophistication. They focus on the most likely attacks and look to realistically test the organization's response to what they might expect to encounter in their daily operations or under exceptional conditions, depending on the scenarios chosen by an organization's exercise stakeholders. Penetration testing is similar, but the attack patterns are typically known to the defending team, and can be of more limited scope. Most often, anything an attacker might have at its disposal or might do is within the rules of engagement for a red team. Given that time is on an attacking team's side in the real world—in similar spirit, red teams can execute long-term campaign-based reconnaissance compromise and lateral movement, encompassing several weeks or months of effort.

Today, blue team members think much like their potential enemies, both in terms of technology considerations and with regard to operations, process, and information access. They typically have at least some familiarity with current red teaming trends and exploits, as part of their training and knowledge-base. For the purposes of exercise, blue teams typically defend some portion of the overall range as their organizational enterprise network, which may include any combination of

user workstations, directory services/domains, email servers, and the like. Blue teams typically have full knowledge of their edge assets accessible to other zones and are aware of potential insider threat risks within their organizational boundaries. The sophistication of the blue team has moved from reactive to proactive protective methods. They too have complex TTPs for hunting adversaries and proactively finding vulnerabilities and indications of compromise in the infrastructure they are defending.

The optimal range must support the sophisticated needs of both teams. Once the range is complete, blue teams must be given time to set up and configure a myriad of potential tools, data collection, and sensors. Red teams must be able to coordinate and execute all steps in their planned phases of attacks.

3.1.1 Cyber Range Architecture Zones and Enclaves

The following section provides a more detailed overview of each of the possible deployed architectural zones of the cyber range. As a terminology reference, "In Game" zones would be components of the cyber range that the exercise participant would be defending from the exercise scenario adversary. Additionally, "In Game" zones would also be the areas utilized by the red team/adversaries and the interconnecting game networks. "Out of Game" zones would be all of the other supporting components that are required for exercise facilitation, but which are not all directly accessible to the participating members of the exercise.

3.1.1.1 "Out of Game" Zones

- Core infrastructure zone - provides the technical foundations of compute, storage, and network that all of the other zones are dependent upon
- Exercise administration and automation zone (aka "white team") - enables exercise administrators to manage and deploy exercise timeline events and coordinate technical resources used during the execution phase
  - Component details
    - web, script, and command-line-based interfaces that are "behind the scenes" controlling the exercise event list
- Participant access zone - provides a mechanism for cyber exercise participants to efficiently and effectively access and utilize the cyber range in the cyber warfare scenarios with which they are challenged
  - Component details
    - web-based interface with authentication and role-based mechanisms to enable team-based virtual machine views
    - mechanisms to upload files into the game space and copy/paste text/configurations
    - inter-team communications mechanisms for collaboration
- Metrics and evaluation zone - provides a mechanism for cyber exercise participants and administrators to evaluate the return on investment/effectiveness of the cyber exercise
  - Component details
    - white team evaluation criteria automation
    - participant evaluation/review/feedback

3.1.1.2 “In Game” Zones

- Adversary zone (aka “red team”) - provides an operational area for adversary teams to configure their tool sets and launch cyber-attacks against the areas being defended by participants of the exercise
  - Component details
    - network
      - layer 3 connectivity, via the “in-game” public Internet / grey zone, to the blue zones being defended
      - source IP diversity mechanisms for advanced attack scenarios
    - core operations
      - Linux and Windows servers configured with desired offensive operations tool sets
- Grey zone - provides in-game “public Internet” services for the networks being defended in advanced scenarios as well as providing a layer 3 conduit between the adversary (red) and defender (blue) zones
  - Component details
    - network
      - routing and IP address management (IPAM) services
    - core operations
      - Internet DNS – root-level and authoritative servers for top-level domains (TLDs)
      - Internet web servers (HTTP and HTTPS) with substantive content
      - email servers and email relaying
    - Internet service provider (ISP), and virtual private network (VPN) and partner services
- Defender zone (aka “blue team”) - closely models a scaled instance of an organization’s production infrastructure architecture. These are all components that will be targeted by the adversary.
  - Component details
    - network
      - routing, proxy, IPAM, dynamic host configuration protocol (DHCP), and firewall services
      - load balancing, content gateways
      - network segmentation, with isolated demilitarized zones (DMZs)
      - intrusion detection / prevention systems (IDS/IPS), switched port analyzer (SPAN/sensor ports, full packet capture
      - packet flow and other related network traffic flow metadata collection
    - core operations
      - system and network directory services
      - single sign-on (SSO) / identity management
      - software updates, licensing, and host management
      - domain (DNS) web, file, and application services
      - real-time event logging and forwarding
    - tools

    - incident response and case management workflow
    - endpoint management and process event tracing
    - change auditing and alerting
    - security information and event management (SIEM)
    - passive and active network vulnerability scanners
- organizational core applications
    - enterprise resource planning (ERP), customer relationship management (CRM), human resource management (HRM)
    - human-machine interface (HMI) / building automation systems (BAS)
    - software configuration management (SCM)
    - source code control systems (SCCS)

4 Architecture and Implementation Considerations

4.1 Systems Architecture Considerations

The defined zones each have many diverse components. At this point, creating a detailed architectural diagram of the proposed cyber range is recommended. This will guide the actual implementation of the components. Given the complexity of this phase, a detailed implementation project plan should also be created and followed.

4.2 Pre-Implementation Considerations

Pre-implementation factors should be considered. As noted earlier, a larger scale cyber range will ideally be implemented using automation and infrastructure-as-code (IaC). There are multiple technologies for enabling these significant efficiency-increasing mechanisms for either implementation in an organization's own data center or in the cloud. Both cloud and data center deployments have pros, cons, and considerations relating to available personnel, timelines, implementation and maintenance budgets, and available existing data center resources.

4.3 Machine Images

A significant initial pre-implementation challenge is building the required diverse upfront machine images and configurations to be compatible for automation and IaC-based implementations. Cloud-based implementations will have an edge in this area because the cloud vendors provide a large library of current up-to-date machine image templates for operating systems and virtual appliances that can be leveraged. In addition, pre-defined network topologies are also available in the cloud. Implementations within an organization's own data center can hopefully leverage existing automation and IaC tools already in place and being used at the data center. Otherwise, do-it-yourself (DIY) researching, learning, and testing several tiers of data-center-based IaC and automation technologies will be required.

Even with these tools, there will be requirements to understand the deeper technical aspects of each and every required operating system and virtual appliance that is used in the cyber range architecture. There are multiple considerations at this stage as detailed below.

- Low machine images count - Keep the total number of required images low while still satisfying the range architecture requirements. OS versions / patches and upgrades are continual. There should be a process to keep the patches and updates synced in the machine images. Fewer machine images will result in less time doing these auxiliary tasks.
- Generic machine images - Create machine images that can readily accept implementation configurations from the automation and IaC processes, while still aligned closely with existing production machine configurations. Virtualization enables the ability to make quick cloned copies of existing, pre-configured servers and workstations. Depending on the specific use case, this may be a very useful technique, especially for highly configured production servers. Ensure that sensitive production or customer data does not find its way into the cyber range on these clones.

- Accepting that complexity will still exist - Diverse machine images with complex configurations will still be required (networking appliances vs. user workstations vs. Linux or Windows servers). Even before the complex configurations are addressed, simply building this basic core, yet diverse set, can weigh heavily on available staffing resources.
- The risks of "too much" automation - "Too much automation" can sometimes be a bad thing relating to cyber range implementation. Virtual appliance images, which are often based on vendor-specific operating systems, can weigh heavily on cyber range implementation timelines. An example of these are commercial vendor virtual device machine images for firewalls, routers, proxies, IDS, and so on. Each is often a non-standard unique implementation case, where integration with deployment automation tools is challenging and, in some cases, not possible. The end result is that some components of the cyber range will need to have manual intervention that will be faster than the effort to fully automate the deployment.
- Adherence to commercial software licensing agreements - Similar to automation integration issues, software licensing management implementations will vary by vendor. Vendor-based virtual appliance images are often the most challenging cases in this area. Automation processes may not be able to license these devices directly. In some cases, the vendor will require "real" Internet access for the licensing, and this may be an issue for cyber ranges that are isolated from a network perspective.

On either path, cloud or local data center / DIY, even with a staff skilled in the suites of tools for automation, and IaC, there still exists a significant hurdle of creating, implementing, and validating original baseline machine image configurations across the diverse components of the planned cyber range architecture.

From the "in-game" blue zone perspective, the machine images can likely be based on production configurations. Despite this, in almost all cases, there will still need to be significant modifications to accommodate the changes required in the scaling differences between the cyber range architecture and an organization's production infrastructure architecture.

4.4 Network Architecture Considerations

Virtual networks are created within a cyber range by attaching virtual machines to specific port groups residing on a virtual switch. The virtual switch applies corresponding virtual local area network (VLAN) identifier (VID) tags to network traffic that traverses these port groups. An important consideration for a cyber range is to maintain a database of port group names and VID numbers. This mechanism enables IaC to automatically deploy virtual machines that are pre-connected to the correct virtual network segment. VLANs are typically allocated on a per-exercise or per-team basis within a range. However, VLANs may also need to be globally reserved for use in all exercises on a range, in the case of a force-on-force simulation that involves multiple team enclaves connected together across a common grey zone.

Traffic isolation between network segments is accomplished by VLAN containment. VLAN ingress and egress management is controlled by the configuration of virtual machine (VM) hypervisor port groups corresponding to the data in the cyber range database. Typically, trunk ports are configured to pass all VLANs on the physical network switches that interconnect the virtual switches across multiple hypervisors.

5 Implementation Order

A properly sequenced deployment of the cyber range zones is critical for meeting implementation timeline goals as outlined in the project plan. Based on experience, the following implementation order of the zones is advised.

1. core infrastructure zone
2. exercise administration and automation zone
3. defender / blue team zone - first phase
    a. L3 routing first, firewalls with no rules
    b. core services - DNS and directory services (Active Directory / LDAP)
4. grey zone / “in game” Internet
5. defender / blue team zone - second phase
    a. link to Internet connectivity - connect routing to the grey zone / “in-game” Internet
    b. core applications - web and app servers
    c. “in-game” management and cyber tools
    d. end users for NPC roles
    e. lock down firewalls based on realistic real-world network policies
6. adversary / red team zone
7. participant access zone
8. metrics and evaluation zone

5.1 Baseline / Minimal Configuration

A cyber range deployed using automation and IaC should be able to be deployed with at least the following baseline characteristics:

- Layer 2 network services are online (i.e., routing and DHCP).
- Layer 3 network connectivity is in place.
- Network devices and Windows and Linux servers and workstations are deployed from machine images.
- All network-connected devices have unique IP addresses assigned and are connected.
- Microsoft Active Directory Domain Services have been deployed with domain connected workstations.
- Basic LAN services like web servers and file servers are deployed with their services enabled.

Getting to this stage is a significant accomplishment for the cyber range being deployed. For some team-based exercise scenarios, this level of configuration may be sufficient. The downside of stopping here is that the fidelity level (i.e., how realistic the cyber range is) is considered low.

5.2 Increasing Realism

As mentioned previously, cyber ranges with high fidelity are one of the most effective ways to supplement and grow an organization's team-based cyber-war-fighting capabilities. Therefore, more work is required to facilitate higher fidelity alignment.

The best paths to increased realism lie in the level of detail of the configurations of the zones. This is most important for the blue zone because this is where the majority of the exercise participants will spend their time. At this stage of range implementation, reliance on subject matter expert (SME) knowledge becomes a necessity. The SMEs of each and every component will need to be consulted on how best to configure each aspect to closely align with what is used in an organization's production environment.

The cyber range implementation team must consider two important aspects of the concept of configurations:

- Tool configurations - These would be passive configurations and relate directly to how a tool is collecting, analyzing, and reporting on the metadata for which it is built (i.e., how event collectors are aggregating, filtering, and forwarding event metadata to a SIEM tool).
- Traffic generation - NPC tools and their configurations are used to generate continual network traffic and resultant log file events. Traffic generation has two main benefits: Tools and sensors have realistic meta data to collect and/or be logged. Secondly, adversary / red team traffic is intermixed on the network and in logs, and is not the only artifact present (making it otherwise relatively easy to identify adversarial behavior).

5.3 Zone Implementation Considerations

For effective implementation, it's important to consider factors particular to each zone.

5.3.1 Blue Zone

- IP addressing schemes closely match the production environment.
- Server and workstation naming conventions align with an organization's policies.
- Routers and firewalls are configured with production-level ACLs and firewall rules enabled.
- Forward and reverse web proxies have content filtering and ACLs enabled.
- Network sensors are properly placed to collect network traffic and associated metadata.
- Directory Services, such as Active Directory or LDAP, are populated with realistic usernames, groups, and group policies.
- Event logging metadata is collected, forwarded, and aggregated to the proper SIEM tools for reporting and analysis.
- Endpoint protection agents are deployed and online and active with current signatures.
- End-user workstations for NPCs are logged in and configured with realistic usernames from the Directory Services.
- NPC's are executing actions that contribute to generating realistic network traffic on the LAN and WAN (i.e., Internet web browsing, file services, and application execution).

5.3.2 Grey Zone

- Internet DNS is populated with real domain names with actual real-world IP addresses.
- Internet routing hops / interconnects are designed to mimic some scale of the real Internet.
- Some quantity of diverse websites has some iteration of content that is browsable by the NPCs.

5.3.3 Red Zone

- Source IP addresses, when geospatial IP is located by SIEM software, will properly place the source of the traffic from an adversarial country.

5.4 Striving for Realism - Challenges and Considerations on "not reinventing the wheel"

Core infrastructure and blue zone configurations should be fairly straightforward and are most closely aligned with what a deployment team would experience in its own organization's production environment. Grey, red, white, and access zones are going to be more challenging to implement. In this case, the SEI/CERT/CWD team has multiple open source tools available that can be used to effectively build out these other zones without having to reinvent the wheel with an organization's internal developers. These tools can be implemented in cloud or DIY datacenter implementations. To date, the SEI/CERT/CWD open source tools focus on exercise management, Internet space infrastructure, and NPC and other traffic generation.

5.5 Integrating Physical Devices into the Range

Some special-purpose systems, (such as internet of things [IoT] and industrial control systems / supervisory control and data acquisition [ICS/SCADA]), have physical components, which cannot be fully simulated in a virtualized range environment. Physical devices can be integrated into the virtual network topology by managing VLANs within the range and on the back-end networking hardware. Connecting a device that has the capacity to be managed via a TCP/IP network management interface is relatively straightforward, as the only requirement is to establish a network path for connectivity to the device. To integrate external special-purpose systems that lack built-in network management capabilities, technologies such as keyboard, video, mouse (KVM) to IP and serial to IP can be leveraged to bring control of these systems into the virtual space.

5.6 Implementation Team Skill Set Guidance

Deploying a larger scale realistic cyber range requires a very diverse set of skills. The following table aligns these skills with each of the defined zones. Having the diversity of skills and dedication of these resources, on the required timeline, will be a significant challenge of the implementation of the cyber range.

Table 1: Required Skills Aligned with Defined Zones

Skill Sets / Skill Set Holders	Core Zone	Blue Zone	Red Zone	Grey Zone	White Zone	Access Zone
Enterprise systems architect	x	x	x	x	x	x
Virtualization engineer	x	x	x	x	x	x
Storage engineer	x					
Network engineer	x	x	x	x		
Linux systems administrators	x	x	x	x		
Microsoft Windows systems administrators	x	x	x	x		
Tool-specific SMEs		x	x	x		
Configuration management/automation engineers	x	x		x		
Project managers	x	x	x	x	x	x
Software developers					x	x
Adversaries/attackers			x			
NPC SMEs		x	x		x	
Those monitoring performance and capacity	x					

6 Conclusion

In this report, we have introduced a comprehensive framework for designing, building, and executing high-fidelity team-based cyber ranges. We believe this framework enables effective realism and improves overall training value within a wide array of potential cyber warfare exercise scenarios. Using this framework, exercise coordinators have the fundamental guidance they need to create a challenging, highly realistic cyber exercise field that teams can utilize on their path to becoming elite cyber-operators.

Further, we anticipate an ongoing series of technical reports in the spirit of this report, highlighting our commitment to realism in training and exercise; covering more technical details of range build and delivery; replicating a reasonable facsimile of the Internet at large; prescribing how to properly create and manage large-scale machine deployments and how to mix cloud and on-premise environments. We also expect to deliver further information on bringing realistic user activity to life on a range, matching effective red team execution with maximal training value, and providing considerations for conducting and measuring elite-level team-based exercises.

References

URLs are valid as of the publication date of this document.

[Dobson 2017]
Dobson, Geoffrey B.; Podnar, Thomas G.; Cerini, Adam D.; & Osterritter, Luke J. *R-EACTR: A Framework for Designing Realistic Cyber Warfare Exercises.* CMU/SEI-2017-TR-005. Software Engineering Institute, Carnegie Mellon University. 2017. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=505224

[Hammerstein 2010]
Hammerstein, Josh & May, Christopher. *The CERT Approach to Cybersecurity Workforce Development*. Software Engineering Institute, Carnegie Mellon University. 2010. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=9697

[Updyke 2018]
Updyke, Dustin D.; Dobson, Geoffrey B.; Podnar, Thomas G.; Osterritter, Luke J.; Earl, Benjamin L.; & Cerini, Adam D. *Ghosts in the Machine: A Framework for Cyber-Warfare Exercise NPC Simulation.* CMU/SEI-2018-TR-005. Software Engineering Institute, Carnegie Mellon University. 2018. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=534316

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
(Leave Blank)	May 2021	Final

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Foundation of Cyber Ranges	FA8702-15-D-0002

6. AUTHOR(S)

Thomas G. Podnar, Geoffrey B. Dobson, Dustin D. Updyke, William E. Reed

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213	CMU/SEI-2021-TR-001

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
SEI Administrative Agent AFLCMC/AZS 5 Eglin Street Hanscom AFB, MA 01731-2100	n/a

11. SUPPLEMENTARY NOTES

12A DISTRIBUTION/AVAILABILITY STATEMENT	12B DISTRIBUTION CODE
Unclassified/Unlimited, DTIC, NTIS	

13. ABSTRACT (maximum 200 words)

This report details the design considerations and execution plan for building high-fidelity, realistic virtual cyber ranges that deliver maximum training and exercise value for participants. A cyber range is a fully interactive virtual instance of the IT infrastructure of a mid-to enterprise-level organization that is dedicated to cyberwarfare training and exercise. The Software Engineering Institute CERT Cyber Workforce Development (SEI/CERT/CWD) team at Carnegie Mellon University has provided a wide array of cyber ranges spanning various levels of complexity to the Department of Defense (DoD) since 2012. This report draws on tens of thousands of hours of experience and expertise in designing, developing, deploying, troubleshooting, maintaining, and improving cyber ranges. The most frequent use case for a cyber range is team exercise. Team exercises can focus on incident response, malware removal, network takeover, process improvement, and many other specialized use cases. This report summarizes the SEI/CERT/CWD current state of the art when it comes to delivering high-fidelity/realistic cyber ranges to DoD and non-DoD customers that support the complex operating environments required for their training and exercise purposes.

14. SUBJECT TERMS	15. NUMBER OF PAGES
high fidelity, realism, R-EACTR, cyber range, cyber warfare exercise, cyber operations	21

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89) Prescribed by ANSI Std. Z39-18
298-102