



Making the Most of an Offensive Security Exercise

Table of Contents

Using Offense for Strong Defense	3
Effective Preparation for Offensive Security Testing	4
Key Considerations Before Engaging in Offensive Security Testing	4
Other Aids to Maximize Offensive Security Testing	5
Key Terms and Concepts to Know for Offensive Security Testing	6
Choosing the Right Offensive Security Exercise	7
CrowdStrike Service Options and Descriptions	8
Web Application Penetration Test	9
External Penetration Test	9
Large Language Model (LLM) Application Penetration Test	9
Social Engineering Exercise	10
Red Team/Blue Team Exercise	10
Internal Red Team Exercise	11
Adversary Emulation Exercise	11
Focused Testing Exercises	11
Pre-Exercise Preparation	12
General Preparation	12
Penetration Test Preparation	12
Red Team Exercise Preparation	13
During the Exercise: Ensuring Seamless Execution	15
Effective Communication	15
Minimize Disruptions During Testing	15
Balance Remediation and Test Integrity	15
Plan for the Debrief	16
Post-Exercise: Turning Insights into Action	16
Finalizing the Report and Driving Action	16
Evaluating Results with a Strategic Lens	17
Debriefing Internally for Greater Impact	18
Prioritizing Findings and Planning Remediation	18
Validating Remediation Efforts and Retest	18
Maximizing Value from Offensive Exercises	19
Stopping Breaches Through Real-World Simulation	20

Using Offense for Strong Defense

Cybersecurity is a dynamic and high-stakes field. Organizations must continuously defend systems while investing heavily in metrics to quantify their security posture. However, measuring effectiveness remains a challenge. Ensuring these investments deliver the intended protection requires proactive testing against evolving threats. A single misstep can expose critical assets to adversaries operating with increasing sophistication.

Offensive security focuses on identifying and exploiting weaknesses before adversaries can take advantage of them. For a deeper exploration, refer to a related article:

[Red Teaming: How Red Team Testing Prepares You for Cyberattacks.](#)

It is important to distinguish offensive security from vulnerability scanning. While automated vulnerability scanners can help identify issues, they stop at identification. Offensive security takes a more active and broader approach. Issues found by offensive security teams are not simply limited to vulnerabilities with a known CVE identifier, they can include processes and people as well. Further, offensive security takes the next step after identification — exploitation. Simulating realistic attacks, offensive security teams attempt to exploit the flaws they find in systems to demonstrate the real risk of the issue.

This paper outlines how to maximize the value of an offensive security exercise, from preparation to post-exercise analysis and improvement.

Effective Preparation for Offensive Security Testing

When preparing for an offensive security exercise, a thoughtful and structured approach ensures meaningful results. While each of these considerations is discussed in more detail, some key areas to consider are:

- **Clearly define objectives internally:** Understanding what the organization hopes to achieve will help align testing efforts with security priorities, from selecting the right type of exercise to setting the exercise's objectives.
- **Communicate with the test team:** Its members are a resource designed to strengthen defenses.
- **Plan:** Establish scope lists and access requirements ahead of time to streamline the process and prevent unnecessary delays.
- **Give advance notice:** Engage the appropriate teams and stakeholders in advance to ensure findings are properly reviewed and integrated into security strategies.
- **Approach the results with the right mindset:** Findings should not be viewed as failures but as opportunities for improvement.
- **Plan for how findings will be tracked:** Ideally, this maintains consistency with other vulnerabilities and risks — introducing new, ad-hoc methods outside of existing processes can create confusion and inefficiencies.

Organizations should not wait for a “perfect” security state before leveraging offensive security, as real-world adversaries do not wait for ideal conditions to strike. Offensive security exercises are not one-time efforts — regular testing helps your organization stay informed and prepared for ever-evolving threats.

Key Considerations Before Engaging in Offensive Security Testing

Organizations of any size and geographic distribution can benefit from an offensive security exercise.

At its core, offensive security seeks to answer critical questions, such as: What could happen if an attacker gains access to a specific system? Could one system be leveraged to move laterally into another? Would an attack be detected in real time? These questions apply universally, whether an organization has 100 employees or 100,000.

That said, there are two key enablers organizations should consider before engaging in an offensive security exercise. At a minimum, organizations should have 1) a proactive security approach and 2) established risk and vulnerability tracking programs.

Offensive security, by definition, is proactive — it focuses on identifying weaknesses before adversaries can exploit them. However, security programs that are still maturing may not have the necessary foundation to fully benefit from this approach. If an organization is still addressing fundamental security gaps and best practices, an offensive exercise may divert critical resources from closing known vulnerabilities.

In such cases, alternative services like a Response Readiness Exercise or Strategic Advisory Services discussion may provide greater value. These services, detailed in the [CrowdStrike Professional Services Catalog](#), help organizations establish a strong security foundation before pursuing offensive testing.

Risk and/or vulnerability tracking is another essential component. Depending on the type of offensive exercise selected, you may need both. To maximize the value of the assessment, a structured process for tracking, prioritizing, and remediating these findings should already be in place. Findings from an offensive security exercise may not fit neatly into existing categories (e.g., a Red Team often generates findings that are not tied to a specific CVE), so some flexibility in existing processes may be needed. However, organizations should avoid creating new tracking processes, as that can often lead to less efficiency, and often, lost visibility. The post-exercise section explores key elements of this process in more detail.

Other Aids to Maximize Offensive Security Testing

Established threat intelligence: Having an established threat intelligence team and/or source, such as CrowdStrike's [threat intelligence offerings](#), will help match findings to threat groups targeting your industry. While real attackers will use the opportunities they have available to them, they are also humans and have preferences or playbooks. Knowing those preferences, tracked by threat intelligence, can help prioritize fixes with limited resources.

Leadership buy-in: Findings from offensive engagements can span a broad range of teams, from the security operations center (SOC) to the identity and access management (IAM) team. Having leadership engaged and supportive throughout the exercise can help drive findings to closure post-exercise when those findings expand beyond the purview of the security team.

Looking beyond compliance-driven objectives: Compliance can be a strong and necessary part of security. While offensive testing may help meet compliance requirements, approaching offensive testing with a broader perspective will help deal with findings that don't fall neatly into a compliance framework.

Viewing the Red Team as allies: While it may be tempting to view offensive exercises in terms of which side "won," both red and blue engineers are ultimately on the same team. Whether or not the Red Team succeeds in achieving the exercise objectives, if the activity improves the overall security posture, both teams win.

Collaboration: Offensive security exercises are almost always time-limited, but many adversaries can take all the time they want. To maximize value from an exercise, working with the Red Team to share information that makes the testing process more efficient can help shrink this gap. Sometimes, a critical finding is only a small bit of knowledge away.

Understanding known strengths and gaps: Organizations can use their known or suspected strengths and weaknesses as starting places for exercise objectives. This is also a good time for the organization to assess what it knows it doesn't have enough data to answer, using the exercise as a means to identify new or suspected risk areas.

Key Terms and Concepts to Know for Offensive Security Testing

The next sections on selecting and preparing for an exercise use the following terms and concepts. Each concept has both pros and cons that may lead an organization to one exercise type versus another.

- **Black-box:** A no-knowledge approach to the exercise, using only information that can be organically gathered without input from the target company
- **White-box:** An open-book approach to the exercise, providing access to things like source code, network diagrams, or other information relevant to the attack scenario
- **Gray-box:** A hybrid approach to the exercise, blending elements of secrecy with useful information
- **White-card:** A shortcut provided to testers to allow them to cover more ground — for example, granting access to a network region that the testers could not reach "organically" to allow further assessment
- **Assumed breach:** A method of starting or continuing a red teaming exercise, in which initial access to the target organization's network is granted to the Red Team, typically by either providing them with remote access credentials or by intentionally running "malware" provided by the Red Team

Choosing the Right Offensive Security Exercise

Offensive security can be broadly categorized into penetration testing and red teaming. Penetration testing focuses on identifying as many vulnerabilities as possible within a defined scope, while red teaming is designed to achieve specific objectives using a variety of tactics, often simulating real-world adversary behavior. These approaches are not mutually exclusive and can be used together to provide a more comprehensive assessment of an organization's security posture.

Selecting the right type of engagement requires careful consideration of key security concerns. Organizations should start by identifying their primary objectives and areas of uncertainty. This may include assessing the risk of adversary access to sensitive customer data, evaluating the effectiveness of IAM tools, or determining whether a specific application is vulnerable to data disclosure.

Organizations can consider the following questions to help guide the decision-making process:

- **Critical assets and business impact:** What are the organization's "crown jewels," and what types of incidents would cause the most significant harm to operations?
- **Visibility gaps:** Where do potential blind spots exist in monitoring and detection capabilities?
- **Risk demonstration:** Are there known security weaknesses that lack budgetary support for remediation, and would demonstrating risk help justify investment?
- **Detection versus response:** Is the primary goal to test the organization's ability to detect threats, respond to incidents, or both? Detection focuses on whether technical controls log and alert on suspicious activity, while response evaluates how well security teams remediate threats, such as quarantining compromised accounts. This can be stated differently: Should the exercise target technical or people-based controls?
- **Desired outcomes:** What is the ultimate goal of the engagement? Common objectives include identifying and remediating vulnerabilities, validating the effectiveness of detective controls, improving security engineering practices, enhancing response team capabilities, augmenting internal offensive testing, or answering a specific security question.
- **Desired outcomes:** What is the cost threshold for the exercise? The levels of stealth, complexity, and duration required for an engagement directly impact cost. Certain objectives, such as testing response capabilities in a stealthy manner, may require additional time and resources to execute effectively.

By addressing these questions, organizations can determine the most suitable type of offensive security exercise, whether it involves a penetration test, a red teaming engagement, or a combination of both. CrowdStrike offers a range of professional services designed to meet these needs. For a detailed overview, refer to the CrowdStrike [Professional Services Catalog](#) (see the Red Team Services section). Next is a summary of when each of CrowdStrike's most popular offerings may (or may not) align with the considerations already outlined.

CrowdStrike Service Options and Descriptions

CrowdStrike has a range of options that align with the considerations outlined in the previous section. The following table lists and summarizes each service option.

Exercise	Purpose	Scope	Primary Controls Addressed	Typical Duration
Web Application Pen Test	A test to identify and exploit flaws in a limited number of applications for data disclosure, privilege escalation, and/or remote code execution	1-3 web applications - No internal pivots	Technical	2 weeks
External Pen Test	A test to assess the external perimeter for exploitable flaws, focusing on remote code execution or data exposure	- External surface - No internal pivots	Technical	2 weeks
Social Engineering Exercise	An exercise that targets people and processes to achieve effects like password resets, remote code execution, or data disclosure	- Phishing/smishing/vishing - No internal pivots	- People - Processes	2 weeks
Red Team/Blue Team Exercise	A guided and collaborative exercise with both Red Team and Blue Team consultants focused on identifying detection opportunities and upskilling defensive team members with real-world attacks	Internal Active Directory domain	- Technical - Detection coverage - Process - People (training)	1 week
Internal Red Team Exercise	An intermediate-length internal network exercise focused on identifying as many attack paths to objectives as possible without stealth	Internal Active Directory domain and cloud environments	- Technical - Detection coverage	2 weeks
Adversary Emulation Exercise	A full scope, long-duration, stealth-focused exercise simulating a motivated attacker trying to breach the perimeter and achieve internal objectives	- External surface - Social engineering - Internal Active Directory domain and cloud environments	- Technical - People - Detection and response	4 weeks (minimum)
LLM Application/ Agent Pen Test	A test of LLM-based applications or agent deployments for objectives that could cause financial or security impacts	- One or more LLM-based applications or agents - Internally or externally facing	Technical	2 weeks (minimum)
Other Offerings	Cloud Breach Emulation and Response Exercise, Wireless Penetration Test, System Hardening Penetration Test, Network Segmentation Test, Mobile Application Penetration Test, Insider Risk Penetration Test, Security Controls Validation			

Web Application Penetration Test

The Web application testing focuses on one or more specific applications, conducting an in-depth evaluation of code-based security and potential business-logic flaws. This assessment can follow a black-box approach with zero knowledge, a white-box approach with access to source code, or a gray-box approach that combines both methods. Testing can also be conducted with or without credentialed access. A white-box approach is recommended to enhance efficiency and effectiveness.

The Web Application Penetration Test is well suited for scenarios where a small number of exposed applications are considered critical assets, and there are specific concerns about whether they could be leveraged to access sensitive data or other environments. This type of test helps strengthen engineering practices while addressing key questions about the confidentiality, integrity, and availability of the service the application provides.

The Web Application Penetration Test is not ideal when the objective is a comprehensive assessment of the external attack surface, as the scope remains limited to designated applications. It also does not address internal network security, as testing concludes upon achieving an effect. The team does not pivot to internal network objectives after achieving remote code execution.

External Penetration Test

The External Penetration Test focuses on reviewing and attacking the external attack surface. While certain exploits, such as memory-corruption attacks like buffer overflows, are excluded to ensure the safety of production servers, the team will attempt to identify and exploit as many vulnerabilities as possible within the internet-facing footprint. Any identified protocol, from HTTP to SSH to VoIP, is in scope for this activity.

The External Penetration Test is well suited for assessing patch management practices, uncovering unknown risks such as shadow IT, and gaining a broader understanding of an adversary's technical capability to breach the perimeter. This assessment pairs effectively with the [CrowdStrike Falcon® Exposure Management](#) module to test servers attributable to the organization.

Because the scope of the External Penetration Test is significantly broader than a single application, it does not provide the same depth of testing as the Web Application Penetration Test. This assessment also does not include social engineering attacks, as the focus remains solely on testing technical controls.

Large Language Model (LLM) Application Penetration Test

The LLM Application Penetration Test assesses how LLM-based applications can be compromised to create financial or security impacts. Testing can be conducted on internally or externally facing applications, depending on requirements. Similar to the Web Application Penetration Test, assessments can be performed from multiple role perspectives, such as unauthenticated versus authenticated access.

This testing provides the most value for organizations that have developed a custom application or agentic framework on top of a foundational model such as Claude or ChatGPT, or a proprietary LLM.

Evaluating third-party LLM services deployed within an organization can also be beneficial, provided they are integrated with internal resources rather than acting solely as pass-throughs.

This assessment does not focus on testing alignment mechanisms designed to prevent harmful or offensive content generation for individual users. While such findings may emerge as a byproduct, the primary objective is to identify vulnerabilities with tangible financial or security implications. Examples include manipulating a public-facing chatbot to approve an unreasonably low transaction price or accessing restricted information.

Social Engineering Exercise

The Social Engineering Exercise targets people-based controls and processes. Campaigns can include traditional email-based phishing, SMS-based phishing (smishing), and voicebased phishing (vishing) targeting various roles, such as the help desk.

This type of engagement is valuable for organizations looking to supplement automated phishing training or assess controls beyond purely technical measures. For example, it can help determine the effectiveness of safeguards preventing unauthorized password resets.

As with other penetration testing services, the assessment does not extend to internal objectives if access is gained. For a full kill-chain exercise, refer to the Adversary Emulation Exercise.

Red Team/Blue Team Exercise

The Red Team/Blue Team (RTBT) Exercise is an intensive, week-long engagement that brings together offensive and defensive consultants to test and enhance security operations

The Red Team conducts a broad-scope attack on the internal network, combining known, repeatable techniques with real-time opportunistic attacks. The goal is to achieve predefined objectives within the internal domain while emulating real-world adversaries.

Meanwhile, the Blue Team analyzes Red Team logs and works alongside the organization's defenders to hunt for indicators of compromise within logging and alerting platforms or to devise appropriate mitigations. This process helps identify detection gaps and provides hands-on training using live attack data.

Each day concludes with a debrief, bringing all teams together to discuss actions taken, key findings, and lessons learned. These sessions provide defenders with direct exposure to attacker methodologies, helping them better anticipate and counter real-world threats.

The RTBT Exercise is ideal for organizations looking to strengthen internal response teams in an openbook, guided format. Direct access to both defensive experts and Red Team operators provides a unique opportunity to sharpen detection and response capabilities. The exercise is limited to a single week, and white-card "shortcuts" may be introduced to maximize learning and ensure key objectives are met.

Internal Red Team Exercise

The Internal Red Team Exercise (IRTE) is a two-week exercise designed to achieve targeted objectives within an internal network. This exercise follows an “assumed breach” approach, starting with established internal access.

Objectives can be tailored to specific needs but often include goals such as full Active Directory domain compromise, on-premises-to-cloud pivots, or common eCrime scenarios like PII/PHI data theft or ransomware preparation. Custom objectives can also be defined to align with specific security concerns.

This exercise is ideal for validating known or suspected security gaps, assessing their potential impact if exploited, uncovering unknown vulnerabilities and misconfigurations, and evaluating the security of critical assets. It also provides an opportunity to enhance internal offensive capabilities by integrating dedicated Red Team resources with existing penetration testing teams.

The IRTE focuses on detective controls and overall security posture. Given the limited timeframe, the Red Team does not prioritize stealth while working toward exercise objectives. As a result, this exercise is not designed to test incident response capabilities.

Adversary Emulation Exercise

The Adversary Emulation Exercise is a full-scope engagement designed to simulate a targeted attack against an organization. Typically conducted as a black-box assessment, the team begins with no prior knowledge or scoping information (other than rules of engagement). It identifies the attack surface organically and attempts to gain access through any means necessary, including technical exploits and social engineering.

If access is obtained, the attack continues into the internal network, working toward a predefined objective while avoiding detection. If perimeter access cannot be achieved organically, shifting to an “assumed breach” model is recommended to assess internal security controls. This exercise is best suited for organizations that have already undergone other offensive testing and want to assess and enhance detection and response capabilities. By focusing on a specific objective and maintaining stealth, the engagement helps uncover blind spots in defensive coverage that might not be apparent in noisier attack scenarios.

Unlike the IRTE, the Adversary Emulation Exercise does not aim to provide broad coverage of internal attack paths. The full-scope nature and stealth requirements necessitate a longer engagement, often spanning four or more weeks.

Focused Testing Exercises

If you have a more focused request, CrowdStrike Services also provides support for cloud RTBT exercises, wireless penetration testing, system hardening testing, and network segmentation testing. Descriptions can be found in the [CrowdStrike Professional Services Catalog](#).

Pre-Exercise Preparation

Committing to an offensive exercise is a big step. To maximize its impact, several key elements should be considered beforehand. This section outlines essential preparations to ensure both the testing team and the organization are set up for success.

General Preparation

While approaching an exercise from the perspective of an unwitting attacker has value, real-world adversaries often have significantly more time to execute an attack than a consulting Red Team. Extending the duration of an engagement could help bridge this gap but would also increase costs unnecessarily.

A more efficient approach is to consider what information a determined attacker could realistically obtain with enough time and effort, then provide that data to the Red Team beforehand. This can reduce time spent on activities that are relevant to achieving the objective, but are not necessarily the focus of the exercise itself.

Another key consideration is maintaining an open mindset throughout the exercise. Hearing about security gaps in a hardened environment isn't always easy, but identifying and addressing those gaps is the primary goal. Ideally, no critical issues emerge, but when they do, the priority is delivering the insights needed to close them before a real attacker can exploit them.

Penetration Test Preparation

Prepare the Scope

Properly defining the scope is one of the most critical steps in preparing for a penetration test, especially for external assessments. This includes gathering key details such as IP ranges, root domains, and specific users targeted for social engineering.

Equally important is validating this information. Testing unauthorized systems is illegal, and while the Red Team conducts due diligence to confirm ownership, the assessment often relies on provided data until conflicting information arises. If scope lists include assets outside the organization's control, it can create significant complications.

Key areas to review during this phase include IP addresses, which can change ownership, and SaaS applications or web hosting services, where a domain might be in scope but the underlying infrastructure belongs to a third party that has not authorized testing.

Scoping should also account for any exclusions or special considerations. This could include users who should not be phished, restricted social engineering pretexts, or applications that support safety-critical functions requiring additional safeguards.

While the team will reach out with any questions, having a clearly defined scope from the outset ensures a smoother and more efficient engagement.

Prepare Access

For external, LLM, or web application penetration tests, setting up a test account and any necessary access accommodations — such as IP allowlisting for a web application firewall (WAF) — before the engagement begins helps ensure the test team can begin immediately on Day One.

Red Team Exercise Preparation

Red Team/Blue Team (RTBT) Exercise Participation

A key objective of the RTBT Exercise is to educate detection and response teams, threat hunters, and detection engineers. To maximize value, identify participants in advance and ensure they have time dedicated to the exercise.

Consistency is critical. Even if team members are rotated throughout the week, designating one or two individuals to stay consistently engaged for the full duration helps maintain continuity and knowledge transfer.

With a short execution window, having IT operations personnel involved or on standby is also beneficial. Their support ensures smooth facilitation of white-card events and helps provide full coverage across the attack lifecycle.

Active participation drives the most value, and open engagement with both Red Team and Blue Team consultants is highly encouraged. Creating a judgment-free environment for asking questions fosters learning and strengthens defensive capabilities.

Prepare Access Methods

For RTBT Exercises and Internal Red Team Exercises (IRTEs), the first day begins with establishing access through an “assumed breach.” The method is flexible, but remote access via a virtual desktop or virtual private network (VPN) is recommended. This setup allows the test team to troubleshoot issues and re-establish connectivity if Red Team command-and-control channels are disrupted.

Another approach is a “witting click,” where a designated individual executes Red Team-developed malware to simulate an initial compromise. This method provides a controlled way to establish command and control while aligning with the desired attack scenario.

Whether creating a test account or selecting an insider for a witting click, defining the attack narrative is key. A standard user profile typically provides a realistic starting point, but scenarios can be tailored — such as assessing the impact of a compromised software developer — by adjusting access levels accordingly.

For Adversary Emulation Exercises, if external access is not achieved within the agreed-upon timeframe — typically the first 1-2 weeks — shifting to an assumed breach approach ensures continued progress toward the exercise objectives.

Define Clear Objectives

Selecting an IRTE or Adversary Emulation Exercise means there’s already a general idea of the objectives, but coming into the kickoff meeting with well-defined goals helps establish clear ground rules from the start.

Consider what the exercise is meant to simulate — whether it's data theft, ransomware execution, or assessing an attacker's ability to breach a segregated manufacturing environment.

Focusing on no more than three primary objectives ensures a strong starting point. As the exercise progresses and goals are met, objectives can be adjusted or expanded to maximize value and ensure the most relevant data is captured within the available timeframe.

Internal Red Team Exercises — Stakeholder Communication

For IRTEs, where stealth is not a factor, proactive communication with key stakeholders helps streamline the exercise and maximize its impact. While every organization approaches this differently, some critical teams to engage include:

- **SOC and Incident Response teams:** Clear communication prevents unnecessary disruptions, ensuring testing time isn't lost to an overly aggressive response or unintended production impact.
- **Security Engineering teams:** Since findings will need to be addressed, engaging these teams early helps smooth the remediation process. They may also provide valuable input on objectives.
- **Risk and Vulnerability Management teams:** Offensive security findings don't always fit neatly into existing risk processes. Bringing in these teams early ensures a structured approach to integrating findings into remediation workflows.
- **Leadership:** If the exercise isn't leadership-driven, looping in key decision-makers helps align objectives with broader security goals and prepares the ground for any necessary executive debriefs. Leadership support is also critical for ensuring fixes are prioritized across multiple teams.

Advanced coordination with these groups reduces friction, improves efficiency, and ensures findings translate into meaningful improvements.

Adversary Emulation Exercises — Managing Deconfliction

The Adversary Emulation team operates with the goal of remaining undetected, but if discovery and eviction occur, deciding how to proceed ensures the exercise stays productive. One approach is to re-establish access through a new witting click, simulating an attacker phishing another user and continuing the engagement with stealth intact. Another option is transitioning to a more open approach, similar to an IRTE, where the focus shifts to identifying additional attack paths without strict stealth requirements.

There's no single correct answer, but having a plan in place reduces downtime and keeps the exercise aligned with its intended objectives.

During the Exercise: Ensuring Seamless Execution

While the preparation steps may seem extensive, they play a critical role in ensuring success, especially within a constrained timeline. The good news is that once the exercise begins, the majority of the workload shifts to the test team.

Even so, certain actions and behaviors can help maintain momentum and ensure the exercise runs smoothly. Staying engaged, facilitating quick decision-making, and addressing any roadblocks as they arise keep the operation effective and focused on achieving its objectives.

Effective Communication

Timely communication helps prevent roadblocks and minimizes downtime throughout the exercise. In the External Penetration Test, for example, the test team may uncover information that raises scope-related questions, potentially pausing operations until clarification is provided. Addressing these issues quickly keeps the engagement moving forward.

For RTBT Exercises, white-card scenarios may occasionally be needed to ensure defenders can observe the full attack lifecycle when organic progression isn't feasible. Coordinating with administrators to efficiently implement these "shortcuts" is essential given the limited timeframe.

A collaborative approach ensures the exercise delivers maximum value. Keeping communication open and responsive allows both teams to adapt as needed and maintain focus on achieving the core objectives.

Minimize Disruptions During Testing

For most exercises, unless response capabilities are being actively tested in an Adversary Emulation engagement, avoid blocklisting IP addresses, domains, or users, and ensure activity is deconflicted before quarantining hosts. While quarantines may still occur, repeated disruptions can significantly impact testing efficiency and limit progress.

Clear communication is essential to prevent unnecessary interruptions. Keeping both the test team and internal response personnel informed ensures that defenders are aware of the exercise and can distinguish simulated threats from real incidents, allowing the engagement to proceed as planned.

Balance Remediation and Test Integrity

It's natural to want to immediately fix security gaps as they're discovered, especially when a quick solution is available. However, remediating issues mid-exercise can significantly impact the validity of test data, particularly in internal domain assessments. For external exercises, remediation is generally more flexible, but it's a topic worth discussing with the test team during the kickoff meeting.

Consider a scenario where the test team gains access to a service account with administrative privileges due to a password stored in file shares. Resetting the password immediately removes the test team's ability to proceed to the next phase of operations without finding another path. In some cases, this might even eliminate the opportunity to assess downstream impacts that would otherwise be exploitable in a real-world attack if another path isn't available.

For serious concerns where vulnerability closure is paramount, engage the Red Team to identify alternative pathways to preserve the integrity of the test while addressing the issue.

Plan for the Debrief

The decision to hold a debrief with technical teams or executives is typically discussed during kickoff, but in many cases, the need (or lack of need) becomes clearer as the exercise progresses. By the second week, there is usually enough insight to determine whether a formal debrief would be valuable.

Scheduling these sessions early helps prevent calendar conflicts that could delay discussions or impact access to critical data, such as logs that may reach their end of life. Ensuring key stakeholders are available at the right time keeps momentum strong and allows for a more effective review of findings and next steps.

Post-Exercise: Turning Insights into Action

With the exercise complete, was everything 100% secure? That's always the goal, but in reality, there are usually areas that can be refined. Once the final report is delivered, much of the post-exercise work shifts to internal teams, but taking the right steps ensures the findings drive meaningful improvements. While the engagement itself may be over, the insights gained serve as a roadmap for strengthening defenses against real-world threats.

Finalizing the Report and Driving Action

Shortly after the active portion of the exercise concludes, a draft report will be provided outlining findings, observations, and recommendations for remediation. In the case of the IRTE or Adversary Emulation Exercise, it will also include narrative elements about other things that were tried but did not succeed. This is the opportunity to review the details, ask clarifying questions, and request adjustments to ensure the report accurately reflects detections, resolved issues, or any additional context relevant to the environment.

If a debrief with technical teams or leadership is needed, scheduling it within one to two weeks of receiving the draft helps maintain momentum while the data is still fresh, both in memory and in logging systems. Providing a few available time slots ensures a smooth discussion and timely action.

Once the draft addresses comments and concerns, the final report is delivered, marking the official close of the exercise. But the critical work begins after that: implementing the findings and strengthening defenses for what comes next.

Evaluating Results with a Strategic Lens

Offensive exercises are complex engagements with numerous variables, and understanding these factors is critical when reviewing findings.

Exercises Capture a Single Point in Time

A vulnerable server might have been offline for maintenance during testing but could be exposed the following week. Similarly, new exploits and attack techniques can emerge shortly after an engagement concludes, meaning a clean report today doesn't guarantee immunity tomorrow.

Exercises Are Time-constrained

While real-world adversaries may spend weeks or months identifying weaknesses, testing is conducted within a defined window. The team works thoroughly, but limitations on time mean findings represent only what could be identified within the engagement scope.

Not All Findings Indicate Persistent Risk

Just because an issue was exploited doesn't mean it poses an immediate or ongoing threat. For example, consider a scenario where the Red Team obtains administrative access on a server where they observe and hijack a session for a highly privileged domain user. The root cause — administrators leaving open RDP sessions — may be a systemic issue, but that exact attack path could disappear within hours. Assessing findings in context ensures remediation efforts focus on true long-term risks.

Every Organization Has Unique Risk Factors

The report provides risk ratings and estimated remediation efforts, but real-world impact varies based on internal controls, operational requirements, and business priorities. Integrating findings into standard risk and vulnerability management processes ensures they are prioritized appropriately.

Detection Capabilities Can Be Misleading

In the Internal Red Team Exercise, stealth is not the primary objective, allowing broader coverage within the available time. This can create a more optimistic view of detection capabilities than would be realistic against a stealthy adversary. Evaluating attacker behaviors rather than specific tools or techniques ensures a stronger defensive strategy.

Detection Does Not Guarantee Response

Alert fatigue is a real concern, and even a properly logged event could be overlooked. While the IRTE focuses on attack visibility rather than response, it still provides valuable insights. One way to bridge this gap is by having analysts log hypothetical response actions to simulate how incidents would be handled in a real scenario.

Understanding these factors helps ensure findings are interpreted correctly and can drive meaningful improvements in security posture.

Debriefing Internally for Greater Impact

Beyond the optional debrief with the CrowdStrike team, holding one or more internal debriefs with key stakeholders ensures alignment on findings and next steps. This is especially valuable if the exercise was conducted with stealth, as it provides an opportunity to share with those who were not directly involved.

These discussions help identify what worked well, where gaps were exposed, and what improvements are needed. They also serve as a starting point for prioritizing remediation efforts and establishing key points of contact for the next phase of action.

Prioritizing Findings and Planning Remediation

With the key considerations for interpreting offensive exercise results in mind, the next step is turning insights into action. Conducting an exercise is only valuable if the findings lead to real improvements — leaving them unaddressed defeats the purpose.

Each report includes suggested remediations and estimated effort levels to help initiate the process, but these are only recommendations. Every organization operates in a unique environment, so findings should be evaluated within the context of existing priorities, constraints, and risk tolerance.

Leveraging existing systems for tracking and remediation ensures findings receive proper attention. If a vulnerability or risk management program is already in place, integrating exercise results into those workflows prevents issues from being lost in a separate process. Offensive security findings may not always fit neatly into existing categories, but adapting them to current tracking mechanisms minimizes overhead and accelerates resolution.

Applying threat intelligence to findings helps prioritize remediation based on real-world risks. Not every vulnerability carries the same level of urgency, and aligning fixes with threats most relevant to the organization ensures efforts are focused where they matter most.

Avoid making snap decisions without considering broader business implications. While certain best practices — such as eliminating unencrypted FTP — seem straightforward, legacy systems may require it, making immediate remediation impractical or cost-prohibitive. Addressing security gaps should be a priority, but not all issues require immediate resolution. Balancing risk, feasibility, and operational impact leads to more effective long-term security improvements.

Validating Remediation Efforts and Retest

The purpose of offensive exercises — testing the effectiveness of defensive controls — extends to the fixes themselves. While some remediation efforts are straightforward to verify, such as closing a firewall port, others involve deeper complexities. In some cases, the chosen remediation may introduce new issues or fail to fully address the original problem, leaving systems exposed.

CrowdStrike offers an optional retest of findings from the exercise, typically focused on high- and critical-severity items to keep the engagement efficient and impactful. Retests are not entirely new exercises — instead, the team will conduct a thorough assessment against the original issue(s) to ensure remediation to a tolerable level. Retesting must take place within six months of the initial exercise, so consider that timeline when planning remediation efforts.

Maximizing Value from Offensive Exercises

Adopt an Opportunity Mindset

One of the biggest obstacles to improving security after an offensive exercise is the perception that poor performance reflects individual failure. This is especially true for teams outside security, such as network engineers, software developers, or social engineering targets. Rather than taking adverse action against individuals tied to specific findings, focus on using the results to strengthen the overall security posture.

Provide Broader Context

When possible, share details about the broader exercise rather than only individual findings with those responsible for remediation. Additional context can lead to better fixes, as subject matter experts can engineer more effective solutions. It also fosters greater engagement with security initiatives across teams. Always ensure compliance with internal restrictions, such as legal privilege, when sharing information.

Leverage Offensive Data for Future Exercises

Findings from an offensive exercise should inform future security engagements. Whether planning another hands-on-keyboard test or a strategic exercise like a tabletop, incorporating previous results provides a stronger starting point and enables exploration of areas that were not assessed in prior engagements due to time constraints or network placement. This approach improves efficiency and expands coverage beyond a fresh start each time.

Apply Offensive Insights to Defense

Offensive tools are not just for attackers — some can be repurposed for defensive use. While untrained personnel should not conduct attacks, many red teaming tools gather and analyze information that can be used to strengthen defenses. CrowdStrike provides [threat intelligence](#) to identify exposed credentials and common tactics, techniques, and procedures (TTPs), along with [CrowdStrike Falcon Exposure Management](#) for monitoring external assets.

Freely available tools, such as [BloodHound Community](#) and [Snaffler](#), can also help collect Active Directory data to map attack paths or uncover credentials stored in accessible file shares. These insights help teams proactively identify and close security gaps before adversaries can exploit them.

Stopping Breaches Through Real-World Simulation

CrowdStrike helps stop breaches by identifying how an organization can be compromised through simulated attacks. Offensive exercises provide critical insights into security posture, revealing both strengths and areas for improvement.

This white paper has covered key strategies to maximize the value of an engagement, from selecting the right type of exercise to planning, execution, and post-assessment actions.

Key takeaways:

- 1 Define clear objectives before testing begins to establish what the exercise should achieve.
- 2 Communicate with the test team, as they are a partner in improving security.
- 3 Plan ahead by addressing scope, access requirements, and logistics early.
- 4 Engage the right stakeholders to ensure findings reach those responsible for remediation.
- 5 Engage the right stakeholders to ensure findings reach those responsible for remediation.
- 6 Avoid creating unnecessary tracking methods — use existing processes for managing findings.
- 7 Offensive security should not be delayed until everything seems perfect. Testing is valuable even when gaps exist.
- 8 One exercise is not enough. Regularly recurring testing ensures continued visibility into the effectiveness of defenses.

**Browse the CrowdStrike →
Professional Services
Catalog**

Taking a proactive approach to offensive security ensures that testing efforts translate into meaningful improvements, strengthening defenses against real-world threats.

Learn more: <http://www.crowdstrike.com/services>

Contact us: services@crowdstrike.com

About CrowdStrike

CrowdStrike (Nasdaq: CRWD), a global cybersecurity leader, has redefined modern security with the world's most advanced cloud-native platform for protecting critical areas of enterprise risk — endpoints and cloud workloads, identity and data.

Powered by the CrowdStrike Security Cloud and world-class AI, the CrowdStrike Falcon® platform leverages real-time indicators of attack, threat intelligence, evolving adversary tradecraft and enriched telemetry from across the enterprise to deliver hyper-accurate detections, automated protection and remediation, elite threat hunting and prioritized observability of vulnerabilities.

Purpose-built in the cloud with a single lightweight-agent architecture, the Falcon platform delivers rapid and scalable deployment, superior protection and performance, reduced complexity and immediate time-to-value.

CrowdStrike: **We stop breaches**

Learn more: <https://www.crowdstrike.com/>

Start a free trial today: <https://www.crowdstrike.com/free-trial-guide/>

