

ROCHESTON®

THE RCCE FIELD MANUAL



**Tactical Tradecraft for Extreme Hacking and Continuous
Verification**

THE RCCE FIELD MANUAL

© 2023 Rocheston. All Rights Reserved.

RCCE® is a registered trademark of Rocheston in the United States and other countries.

No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission of Rocheston. This book is intended for informational and educational purposes only. The views expressed herein are the opinion of the author and should not be taken as professional advice. The author of this book and publisher are not responsible for any loss or damage resulting from the use of this book.

TABLE OF CONTENTS

| | |
|---|-----------|
| PURPOSE OF THIS MANUAL | 4 |
| CHAPTER 1: THE RCCE ROLE IN THE RCF ECOSYSTEM | 9 |
| CHAPTER 2: THINKING LIKE THE ADVERSARY | 12 |
| CHAPTER 3: ENTRY TECHNIQUES AND INITIAL ACCESS VALIDATION | 16 |
| CHAPTER 5: PERSISTENCE, STEALTH, AND SURVIVABILITY | 23 |
| CHAPTER 6: DETECTION ENGINEERING VALIDATION | 26 |
| CHAPTER 7: AUTONOMOUS DEFENSE AND SELF-HEALING TESTING | 29 |
| CHAPTER 8: EVIDENCE CAPTURE UNDER ADVERSARIAL CONDITIONS | 32 |
| CHAPTER 9: ROSECOIN VAULT INTEGRATION FOR FIELD OPERATIONS | 36 |
| CHAPTER 10: OPERATING DURING LIVE INCIDENTS | 39 |
| CHAPTER 11: POST-OPERATION REPORTING AND IMPROVEMENT LOOPS | 43 |

PURPOSE OF THIS MANUAL

This book exists for one reason: execution under pressure.

The Rochester Cybersecurity Framework defines what must be protected and why. This manual defines how that protection is tested, enforced, broken, and rebuilt in real environments by Rochester Certified Cybersecurity Engineers.

RCCE engineers are not auditors. They are not checkbox assessors. They are not compliance administrators. They are operators. Their mission is to subject security controls to adversarial pressure that mirrors what real threat actors will do, and to produce evidence that either validates those controls or exposes their failure.

This manual documents the offensive and defensive tradecraft used by RCCE engineers to continuously validate security posture against real adversary behavior. It turns abstract controls into living systems by attacking them, stressing them, and proving their resilience through evidence rather than assumption.

The operating principle is simple: if a control cannot survive RCCE testing, it is not considered real. There is no room for assumed security. There is no credit for intention. Either a control works under adversarial pressure, or it does not. This manual teaches engineers how to determine the truth and document it in a way that survives scrutiny.

Every technique, every methodology, every operational protocol in this book has been designed with a single priority: producing defensible, evidence-backed truth about the actual security posture of an environment. Not the hoped-for posture. Not the documented posture. The actual posture.

How This Manual Is Different from Traditional Security Books

Most security books teach tools. Some teach tactics. Very few teach operational doctrine.

The cybersecurity bookshelf is crowded with vendor-specific walkthroughs, certification study guides, and theoretical frameworks that describe what security should look like in ideal conditions. These books serve their purpose. This is not one of those books.

This manual is a field doctrine. It does not walk through tool configurations or explain how to run a scanner. It teaches how RCCE engineers think, how they model adversaries, how they break environments safely, how they validate controls continuously, and how they generate proof-grade evidence that survives audits, investigations, and courtrooms.

The difference is operational context. A tool manual tells you what buttons to press. A tactics guide tells you what sequence to follow. A doctrine manual tells you how to think when the plan breaks down, when the environment is not what the documentation says it is, and when you have thirty minutes to validate a critical control while an incident is active.

Everything in this book is written from five core assumptions that define the RCCE operational reality:

- The environment is hostile. Nothing can be trusted until validated.
- Controls will fail. The question is not whether but when and how.
- Attackers will adapt. Static testing produces static results that expire immediately.
- Time will be limited. There will never be enough time to test everything perfectly.
- Evidence must survive scrutiny. If the finding cannot be proven, it does not exist.

These assumptions are not pessimistic. They are realistic. They represent the conditions under which RCCE engineers actually operate, and they drive every decision about methodology, tooling, evidence capture, and reporting that this manual describes.

A traditional penetration testing engagement produces a report that says what was found during a snapshot in time. RCCE operations produce continuous streams of evidence that validate or invalidate security posture across the full lifecycle of an environment. The difference is not incremental. It is architectural.

Who This Book Is For

This book is written exclusively for practitioners who operate in live security environments and who have already achieved fluency in security fundamentals.

The primary audience includes Rochester Certified Cybersecurity Engineers who have completed the RCCE certification program and are executing operational security validation in the field. It also serves red team and blue team operators who conduct adversarial testing or defensive operations against real threats, security engineers responsible for the design, implementation, and validation of controls in production environments, incident responders and resilience engineers who must operate under crisis conditions while maintaining evidence integrity, and architects who design security controls that must survive determined adversarial pressure.

This book is not written for beginners. It assumes the reader has working knowledge of operating systems internals across Linux, Windows, and macOS. It assumes fluency in network protocols, routing, DNS, and traffic analysis. It assumes hands-on experience with at least one major cloud platform and its identity, networking, and logging subsystems. It assumes understanding of authentication protocols, identity federation, and access control architectures. It assumes familiarity with common attack techniques, vulnerability classes, and exploitation concepts.

If any of these assumptions do not hold, the reader should build that foundation first. The Rochester Cybersecurity Framework documentation and the RCCE certification program provide the structured learning path needed to reach this level. This manual is where that learning is applied under operational conditions.

Core RCCE Philosophy: Extreme Hacking

Extreme Hacking is the operational methodology that defines RCCE work. It is not a marketing phrase. It is not a branding exercise. It is a precise description of how RCCE engineers approach the validation of security controls.

The word extreme does not mean reckless. It means thorough to the point of discomfort. It means testing controls not just against known attack patterns but against creative, adaptive, multi-stage attack chains that mirror the behavior of sophisticated adversaries. It means refusing to accept surface-level findings when deeper investigation is possible. It means pursuing the truth of a system's security posture regardless of how uncomfortable that truth may be for stakeholders.

The word hacking does not mean unauthorized access. It means systematic, authorized adversarial pressure applied to systems, networks, identities, and processes for the purpose of

validation. Every action taken under the Extreme Hacking methodology is authorized, scoped, documented, and reversible.

RCCE engineers operate under a simple rule: if it can be exploited safely, it will be exploited intentionally, before a real adversary does. This is not recklessness. This is diligence. The alternative is to leave gaps that real adversaries will find and exploit without the courtesy of documentation.

What Extreme Hacking Focuses On

Extreme Hacking prioritizes the attack behaviors that cause actual organizational damage. Surface-level vulnerability scanning is not Extreme Hacking. Running a tool and reading its output is not Extreme Hacking. Checking a box on a compliance form is not Extreme Hacking.

Extreme Hacking focuses on real attacker paths rather than theoretical ones. Real adversaries do not attack systems the way vulnerability scanners do. They chain together misconfigurations, identity weaknesses, trust boundaries, and timing windows into multi-stage attack paths that reach critical assets. RCCE engineers model and execute these paths.

It focuses on privilege escalation rather than surface findings. Finding that a web application returns verbose error messages is a finding. Demonstrating that those error messages reveal internal architecture that enables lateral movement to a domain controller is Extreme Hacking. The emphasis is always on what an attacker can do with a finding, not the finding in isolation.

It focuses on lateral movement rather than isolated vulnerabilities. The most dangerous phase of any real attack is the expansion phase, where an attacker moves from initial access to critical assets. RCCE engineers validate whether the controls designed to prevent lateral movement actually function under pressure.

It focuses on persistence rather than one-time access. A real adversary who gains access to an environment does not leave after a single session. They establish persistence mechanisms that allow continued access. RCCE engineers validate whether environments can detect and remove persistent access.

It focuses on evidence capture at every stage. Every action taken during Extreme Hacking operations is documented with timestamps, artifacts, and chain-of-custody discipline. This is not bureaucracy. This is what separates professional adversarial validation from unauthorized access.

The Ethical Boundary

RCCE work is always authorized, scoped, and reversible. The goal is not disruption. The goal is truth.

Authorization means that every RCCE operation is conducted under explicit written authorization that defines the scope, timing, methods, and boundaries of the engagement. No exceptions. No implied authorization. No verbal agreements. Written authorization that names the systems, the methods, and the people.

Scoped means that RCCE engineers operate within defined boundaries. If a system is out of scope, it is not tested. If a technique is excluded from the engagement rules, it is not used. Scope boundaries exist for operational and legal reasons, and they are absolute.

Reversible means that RCCE engineers plan for rollback before they execute. Every action that modifies an environment has a documented reversal procedure. If something breaks unexpectedly, the engineer can restore the previous state. This is not optional. It is a fundamental discipline of the methodology.

CHAPTER 1: THE RCCE ROLE IN THE RCF ECOSYSTEM

Where RCCE Engineers Fit

The Rochester Cybersecurity Framework is a comprehensive security model built on twenty-four domains that cover every major surface of organizational security. These domains define what must be protected, how controls should be designed, what automation is expected, and what evidence of compliance looks like.

But frameworks do not protect organizations. Controls protect organizations. And controls only protect organizations when they actually work.

This is where RCCE engineers operate. They sit at the intersection of governance and reality, tasked with a single mission: determine whether the controls that exist on paper actually function under adversarial pressure. Their role is not to audit documentation. It is not to review configurations. It is to attack the environment the way a real adversary would and produce evidence of what happens.

In the RCF ecosystem, there are three distinct layers of security operation. The first is governance, which defines requirements, policies, and control specifications. The second is automation, which implements controls through technology including monitoring, enforcement, and response. The third is human-led adversarial validation, which is the RCCE layer.

Each layer depends on the others, but they are fundamentally different activities. Governance produces requirements. Automation implements those requirements. RCCE testing validates whether the implementation actually works. Without RCCE validation, governance is theory and automation is hope.

The Difference Between Assessment and Validation

The security industry has spent decades conflating assessment with validation. They are not the same thing.

Assessment involves reviewing configurations, policies, and architectures against a standard or benchmark. It answers the question: does this environment have the right things in place? Assessment is necessary but insufficient. An environment can have every recommended control deployed and configured exactly according to documentation, and still be completely vulnerable to a sophisticated attacker who chains together subtle weaknesses that no assessment checklist covers.

Validation involves testing whether controls actually prevent, detect, and respond to real attack behaviors. It answers the question: do these controls work when someone is actively trying to defeat them? Validation requires adversarial pressure. It requires creativity. It requires the willingness to break things in controlled ways to learn the truth.

RCCE engineers are validators, not assessors. They may use assessment data as input to their operations, but their output is fundamentally different. An assessor produces a compliance report. An RCCE engineer produces evidence of whether the environment survives attack.

RCCE Authority and Boundaries

RCCE engineers operate with specific authority that must be formally established before any operation begins. This authority is not assumed. It is not inherited from a general security role. It is explicitly granted for each engagement or operational period.

The authorization document for RCCE operations must define the specific systems, networks, and identities that are in scope. It must define the techniques that are authorized and any techniques that are explicitly excluded. It must define the time window during which operations may be conducted. It must define the escalation path for critical findings.

It must define the communication protocol between the RCCE engineer and the organization's operational staff. And it must define the evidence handling requirements, including who may access RCCE findings and under what conditions.

RCCE engineers have a responsibility to operate within their authorization and to halt operations immediately if they discover that continuing would violate scope, law, or safety boundaries. This is not a limitation on capability. It is a professional discipline that enables RCCE engineers to operate with maximum effectiveness within maximum trust.

Escalation Responsibility

RCCE engineers will discover things that require immediate action. A critical vulnerability in a production system. Evidence of active compromise. A misconfiguration that exposes sensitive data. These discoveries create an escalation responsibility that overrides the normal operation timeline.

The escalation protocol must be defined before operations begin, not improvised during a crisis. RCCE engineers must know exactly who to contact, through what channel, and within what timeframe when they discover critical findings. This protocol must account for the possibility that the normal escalation path may be compromised, which means that alternative escalation paths must also be defined.

Escalation responsibility also includes the obligation to preserve evidence before escalating. If an RCCE engineer discovers active compromise, the first action is to capture evidence. The second action is to escalate. This order matters because evidence may be destroyed or modified once the compromise is disclosed.

CHAPTER 2: THINKING LIKE THE ADVERSARY

Beyond Attack Lists

The cybersecurity industry has produced extensive catalogs of attack techniques. Frameworks like MITRE ATT&CK document hundreds of techniques organized by tactical category. These catalogs are valuable references, but they are not adversary models. An attack list tells you what techniques exist. An adversary model tells you which techniques a specific threat actor will use against a specific target, in what order, with what level of stealth, and with what objectives. The difference between the two is the difference between a dictionary and a conversation.

RCCE engineers do not memorize attack lists. They model intent. They ask: who is likely to attack this environment, what are they trying to achieve, what capabilities do they have, how much time are they willing to invest, and how much noise are they willing to make? The answers to these questions determine which techniques are relevant and how they should be chained together.

Motivation Analysis

Every adversary has a motivation. Understanding that motivation is the first step in building an accurate adversary model.

Financially motivated attackers optimize for monetization speed and scale. They target systems that provide direct financial access such as payment systems, banking applications, and cryptocurrency wallets. They also target systems that contain data with resale value such as personal information databases, healthcare records, and credentials. Their techniques prioritize efficiency over stealth because their operational model depends on volume. Espionage-motivated attackers optimize for access persistence and data exfiltration. They target intellectual property, strategic communications, research data, and decision-making systems. Their techniques prioritize stealth over speed because their operational model

depends on sustained access. They are willing to invest months or years in maintaining undetected presence within a target environment.

Sabotage-motivated attackers optimize for impact and disruption. They target operational technology, critical infrastructure, data integrity, and public-facing systems. Their techniques may prioritize destructive capability over stealth because their operational model depends on visible damage. However, sophisticated sabotage operations may combine long-term stealth access with timed destructive payloads.

Hacktivist-motivated attackers optimize for visibility and message amplification. They target systems where compromise can be publicized, such as public websites, social media accounts, and data stores that contain embarrassing information. Their techniques prioritize high-visibility impact over sustained access.

RCCE engineers must understand which motivations are most relevant to the environment they are validating. A financial services organization faces different adversary motivations than a defense contractor, and the validation approach must reflect those differences.

2.3 Capability Assessment

Motivation tells you what an adversary wants. Capability tells you what they can do. Adversary capabilities range from commodity tools available to anyone with an internet connection to custom zero-day exploits developed by nation-state programs. RCCE engineers must calibrate their testing to the capability level that represents a realistic threat to the environment.

Low-capability adversaries use publicly available tools, known exploits, and automated scanning. They rely on misconfigurations and unpatched vulnerabilities. Validating against this capability level is necessary but not sufficient for any environment that faces sophisticated threats.

Medium-capability adversaries combine public tools with custom scripts, social engineering, and moderate technical skill. They can chain together multiple weaknesses and adapt to defensive responses. They represent the threat level faced by most commercial organizations. High-capability adversaries develop custom tools, discover zero-day vulnerabilities, and operate with professional discipline including operational security, staged infrastructure, and

patient execution. They represent the threat level faced by critical infrastructure, defense, intelligence, and high-value commercial targets.

RCCE testing should be calibrated to at least the medium-capability level for any organization, with high-capability simulation for environments that face nation-state or advanced persistent threats. Testing only against commodity threats creates a false sense of security that evaporates against the first competent attacker.

2.4 Time and Persistence Modeling

Real adversaries operate on timescales that range from minutes to years. The amount of time an attacker is willing to invest fundamentally changes the techniques they use and the defenses that are relevant.

A smash-and-grab attacker with hours of operational time will use fast, noisy techniques and prioritize immediate objectives. The defenses that matter against this threat are detection speed and automated response.

An advanced persistent threat with months of operational time will use slow, quiet techniques and prioritize maintaining access. The defenses that matter against this threat are behavioral analytics, anomaly detection, and periodic security validation that catches slow-moving compromise.

RCCE engineers must model the time investment of their adversaries and adjust their validation approach accordingly. A two-week engagement that only tests fast attack paths validates only half the threat model. Continuous RCCE operations that test slow, persistent attack paths over months provide the other half.

2.5 Noise Tolerance and Stealth Tradeoffs

Every attack technique generates some level of observable activity. Some techniques are loud, generating log entries, network traffic patterns, and system events that are easy to detect. Others are quiet, blending into normal activity patterns and generating minimal observable evidence.

Adversaries make conscious tradeoffs between noise and speed. A loud technique that achieves the objective quickly may be preferable if the attacker believes detection response

will be slow. A quiet technique that takes longer may be preferable if the attacker believes detection capability is high.

RCCE engineers must test both ends of this spectrum. Loud techniques validate whether detection and response controls function at all. Quiet techniques validate whether detection and response controls can identify sophisticated activity that blends into normal patterns. Testing only loud techniques validates the floor of defensive capability. Testing quiet techniques validates the ceiling.

2.6 Decision-Making Under Uncertainty

Real adversaries operate with incomplete information. They do not have a network diagram, an asset inventory, or a list of deployed security controls. They must discover the environment as they move through it, making decisions based on partial knowledge and evolving understanding.

RCCE engineers should replicate this uncertainty in at least some of their operations. Black-box testing, where the engineer begins with no internal knowledge of the environment, forces the same discovery process that a real attacker would follow. This approach validates controls that the engineer might unconsciously avoid if they had full knowledge of the environment.

Decision-making under uncertainty also means accepting that not every attack path will succeed. Real adversaries encounter dead ends, trigger detections, and must adapt their approach in real time. RCCE engineers who practice decision-making under uncertainty develop the same adaptability, which makes their validation more realistic and their findings more credible.

CHAPTER 3: ENTRY TECHNIQUES AND INITIAL ACCESS VALIDATION

3.1 The Initial Access Problem

Initial access is the moment an adversary transitions from external observation to internal presence. It is the most studied phase of the attack lifecycle, the most heavily defended, and paradoxically, the most frequently bypassed.

Organizations invest heavily in perimeter security, email filtering, endpoint protection, and identity controls. Yet initial access continues to succeed against even well-defended environments. The reason is not that the controls are absent. It is that the controls contain gaps, misconfigurations, and assumptions that adversaries systematically discover and exploit.

RCCE engineers validate initial access controls by attempting to bypass them using techniques that mirror real adversary behavior. The goal is not to harvest credentials or steal data. The goal is to determine whether the controls that are supposed to prevent unauthorized access actually do so.

3.2 Identity Abuse Simulations

The most common initial access vector in modern environments is identity abuse. Stolen credentials, session tokens, and authentication bypass represent the path of least resistance for most adversaries.

RCCE engineers validate identity controls by testing credential theft resistance through simulated phishing campaigns that measure whether phishing-resistant authentication mechanisms actually prevent credential harvest. They test session management by validating

whether session tokens expire appropriately, whether session fixation is possible, and whether stolen sessions can be replayed from unauthorized locations. They test multi-factor authentication by verifying that MFA implementation cannot be bypassed through protocol downgrade, social engineering of helpdesk staff, or exploitation of MFA fatigue. They test identity federation by validating that federated identity systems properly validate assertions and that trust relationships cannot be abused.

Each identity abuse simulation must be conducted with specific preconditions documented, including the assumed capability of the attacker, the specific control being tested, and the expected result. The simulation must include safety controls that prevent actual credential compromise and rollback procedures that restore any modified identity configurations.

3.3 Misconfiguration Exploitation

Misconfigurations are the silent epidemic of enterprise security. A system can pass every compliance check and still contain misconfigurations that enable unauthorized access. RCCE engineers systematically probe for misconfigurations across exposed services, cloud resources, network devices, and application platforms. This includes testing default credentials on network devices, management interfaces, and service accounts. It includes validating that cloud storage resources are not publicly accessible due to permissive IAM policies. It includes testing for exposed management ports, debug interfaces, and administrative endpoints that should not be reachable from external networks. It includes validating that DNS configurations do not leak internal network information.

Misconfiguration exploitation is particularly important because misconfigurations are dynamic. A system that is correctly configured today may be misconfigured tomorrow due to a change management failure, a deployment script error, or a manual modification by an administrator under pressure. Continuous RCCE validation catches these drift-related misconfigurations before adversaries do.

3.4 Phishing-Resistant Control Bypass Testing

Phishing-resistant authentication has become a standard requirement, but implementation quality varies enormously. RCCE engineers test whether phishing-resistant controls actually resist phishing.

This includes testing FIDO2 and WebAuthn implementations for downgrade vulnerabilities where an attacker can force fallback to less secure authentication methods. It includes validating that phishing-resistant controls apply to all authentication flows, not just the primary login page. It includes testing whether administrative override procedures can bypass phishing-resistant requirements. It includes validating that phishing-resistant controls extend to service accounts and non-interactive authentication.

The emphasis is on real-world bypass, not theoretical vulnerability. If a phishing-resistant control can be bypassed through social engineering of a helpdesk operator, that is a valid finding regardless of the technical strength of the underlying authentication protocol.

3.5 Supply-Chain Access Modeling

Modern organizations depend on extensive supply chains of software vendors, service providers, and third-party integrations. Each supply chain relationship creates a potential initial access vector.

RCCE engineers validate supply chain controls by modeling how a compromised vendor could access the target environment. This includes testing API integrations for excessive permissions, validating that vendor remote access is properly scoped and monitored, testing whether compromised software updates could deliver payloads, and validating that supply chain security monitoring detects anomalous vendor behavior.

Supply chain validation is inherently complex because it involves systems and relationships that cross organizational boundaries. RCCE engineers must work within their authorization scope, which means they cannot directly test vendor systems. Instead, they model supply chain attacks by testing the target organization's controls against simulated vendor compromise scenarios.

3.6 Shadow Asset Discovery

Shadow assets are systems, services, and identities that exist in an environment but are not reflected in the organization's official inventory. They represent one of the most reliable initial access vectors because they are, by definition, not covered by the organization's security controls.

RCCE engineers perform shadow asset discovery as a standard part of initial access validation. This includes scanning for previously unknown external-facing services, identifying unmanaged cloud resources created outside the normal provisioning process, discovering test and development systems that have been connected to production networks, and identifying service accounts and identities that are not tracked in the organization's identity management system.

Shadow asset discovery often produces findings that are more impactful than any specific vulnerability. A fully patched, well-configured server that is not in the asset inventory is effectively invisible to the security operations team, which means it is also invisible to monitoring, alerting, and incident response.

3.7 Evidence Capture for Initial Access Testing

Every initial access technique used during RCCE operations must produce evidence that documents what was attempted, what succeeded, what failed, and what the implications are. Evidence for initial access testing includes timestamped records of every technique attempted, including techniques that did not succeed. It includes network captures showing the authentication flows tested. It includes screenshots and artifacts from successful access. It includes documentation of the specific control that failed and the conditions under which it failed. It includes rollback confirmation showing that any changes made during testing were reversed.

This evidence serves two purposes. First, it enables the organization to understand exactly what happened during testing and take appropriate remedial action. Second, it creates a defensible record that protects the RCCE engineer and the organization in the event of any dispute about what occurred during the engagement.

Chapter 4: Privilege Escalation and Lateral Movement

4.1 Why Expansion Is the Critical Phase

Initial access gets the headlines, but expansion is where the damage happens. An adversary who gains initial access to a workstation with standard user privileges has limited capability. An adversary who escalates to administrative privileges and moves laterally to domain controllers, database servers, and cloud management planes has the capability to destroy the organization.

The expansion phase, which encompasses privilege escalation and lateral movement, is the most critical phase for defensive validation because it is the phase where defensive controls have the most opportunity to detect and contain an attack. An initial access event may be a single authentication anomaly that blends into normal noise. An expansion campaign generates a pattern of activity across multiple systems that should be visible to monitoring, anomaly detection, and behavioral analytics.

RCCE engineers focus significant effort on validating expansion controls because these controls represent the organization's last line of defense before catastrophic compromise.

4.2 Identity Privilege Drift Validation

Privilege drift is the gradual accumulation of unnecessary permissions by users, service accounts, and automated identities over time. It is one of the most reliable paths from initial access to administrative control.

RCCE engineers validate privilege drift by mapping the actual permissions of identities across the environment and comparing them to the permissions those identities should have based on their role. This mapping often reveals service accounts with domain administrator privileges that were granted temporarily and never revoked, user accounts with permissions inherited from group memberships that are no longer appropriate, and automated identities with broad cloud IAM roles that were created during initial deployment and never scoped down.

The validation is not just finding the drift but demonstrating its impact. An RCCE engineer who identifies a service account with excessive privileges should demonstrate the attack path that those privileges enable, from initial compromise of the service account to the critical assets that become reachable through that compromise.

4.3 Service Account Abuse

Service accounts are the skeleton keys of modern environments. They often have broad permissions, they rarely have multi-factor authentication, their credentials are frequently stored in insecure locations, and their activity is often excluded from monitoring because it generates high volumes of expected traffic.

RCCE engineers validate service account security by testing whether service account credentials can be extracted from configuration files, environment variables, deployment scripts, or memory. They test whether service account permissions can be abused to access resources beyond the intended scope. They test whether service account activity is monitored and whether anomalous service account behavior triggers alerts. They test whether service accounts can be used to create additional identities that persist beyond the operational window.

Service account abuse is a particularly valuable RCCE testing focus because it validates multiple control domains simultaneously: identity management, secrets management, monitoring, and access control.

4.4 Cloud Role Chaining

Cloud environments introduce a new dimension of privilege escalation through role chaining. Cloud IAM systems allow identities to assume roles, and those roles can often assume additional roles, creating chains of privilege that may not be visible in any single configuration review.

RCCE engineers validate cloud role chaining by mapping the assume-role relationships across the cloud environment and identifying chains that lead from low-privilege identities to high-privilege roles. This includes testing whether a compromised developer workstation can chain through CI/CD service roles to reach production infrastructure. It includes validating whether cross-account role assumptions are properly constrained. It includes testing whether temporary credential lifetimes are appropriate for the permissions they grant.

Cloud role chaining validation is essential because the attack surface is not visible in traditional network-based security analysis. The privilege escalation path exists entirely within the cloud provider's identity system and may not generate any network traffic that traditional monitoring would detect.

4.5 Network Trust Boundary Abuse

Organizations define trust boundaries to segment their networks into zones with different security levels. The assumption is that traffic between zones is controlled and monitored. RCCE engineers test whether this assumption holds.

Network trust boundary validation includes testing whether micro-segmentation rules prevent lateral movement between workloads. It includes validating that network firewalls block unauthorized cross-zone traffic under all conditions, not just under normal conditions. It includes testing whether shared services such as DNS, NTP, and logging infrastructure can be abused to bridge trust boundaries. It includes validating that network monitoring detects anomalous cross-boundary traffic patterns.

Trust boundary abuse is often the most revealing test in an RCCE engagement because it exposes the gap between the network architecture as designed and the network architecture as implemented. Firewall rules accumulate exceptions over time. Micro-segmentation policies drift. Shared services create implicit trust paths that bypass explicit controls.

4.6 Cross-Domain Movement

Modern environments span multiple domains: on-premises Active Directory, cloud identity providers, SaaS applications, and operational technology networks. Each domain has its own identity system, its own access controls, and its own monitoring. The boundaries between domains are often the weakest points in the security architecture.

RCCE engineers test cross-domain movement by validating whether compromise in one domain enables access to another. This includes testing whether on-premises Active Directory compromise enables access to cloud resources through directory synchronization or federation. It includes validating whether SaaS application compromise enables access to internal systems through API integrations. It includes testing whether IT network compromise enables access to operational technology networks through converged infrastructure.

Cross-domain movement testing requires careful coordination and clear authorization because it may involve systems managed by different teams with different security policies. The RCCE engineer must ensure that authorization covers all domains that may be affected by the testing.

CHAPTER 5:

PERSISTENCE, STEALTH, AND SURVIVABILITY

5.1 Testing Survivability, Not Playing Attacker

Persistence testing validates a critical defensive question: if an attacker gains access, can they keep it? The answer determines whether the organization faces a recoverable incident or an ongoing compromise.

RCCE engineers test persistence not to demonstrate hacking skill but to validate whether the environment can detect and remove persistent access. The objective is not to stay hidden indefinitely. The objective is to validate whether monitoring, detection, and recovery controls function as designed when an adversary establishes persistent mechanisms. This distinction matters because it drives the methodology. An RCCE engineer testing persistence is not trying to win a hide-and-seek game against the blue team. They are systematically deploying known and novel persistence techniques and measuring whether each one is detected, how quickly it is detected, and whether the detection provides enough information for effective remediation.

5.2 Persistence Mechanism Validation

Persistence mechanisms range from simple scheduled tasks to sophisticated firmware implants. RCCE engineers validate the environment's ability to detect and remove persistence across the full spectrum.

Operating system persistence validation includes testing scheduled tasks and cron jobs, startup scripts and registry run keys, service installations and systemd unit files, WMI event subscriptions, and DLL search order manipulation. Each mechanism is deployed with proper documentation, and the time to detection is measured.

Identity persistence validation includes testing whether additional credentials can be added to existing accounts, whether new accounts can be created that blend into the normal identity population, whether API keys and tokens persist beyond expected lifetimes, and whether federated identity trust relationships can be manipulated to maintain access. Infrastructure persistence validation includes testing whether network devices can be modified to maintain access, whether cloud resources can be created that provide persistent access paths, and whether backup systems can be compromised to enable re-entry after remediation.

5.3 Detection Evasion Testing

Detection evasion testing validates whether security monitoring can identify adversary behavior that is specifically designed to avoid detection.

This includes testing whether endpoint detection and response systems can identify techniques that use native operating system tools rather than known malware. It includes validating whether network monitoring can detect data exfiltration that uses encrypted channels, steganography, or protocol tunneling. It includes testing whether behavioral analytics can identify anomalous activity when the adversary deliberately mimics normal user behavior patterns. It includes validating whether log integrity controls prevent an adversary from modifying or deleting evidence of their activity.

Detection evasion testing produces some of the most valuable findings in RCCE operations because it identifies the gaps between what the security monitoring is designed to detect and what it actually detects. These gaps represent the attack surface that sophisticated adversaries will exploit.

5.4 Log Integrity Stress Testing

Logs are the evidentiary foundation of security operations. If logs can be modified, deleted, or suppressed, then detection, investigation, and response all fail.

RCCE engineers validate log integrity by testing whether log collection can be interrupted without triggering alerts. They test whether log forwarding can be redirected to prevent security-relevant events from reaching the monitoring system. They test whether log storage can be accessed and modified by an attacker with elevated privileges. They test whether log

timestamps can be manipulated to create false timelines. They test whether critical events can be generated that do not appear in any log source.

Log integrity testing requires coordination with the monitoring team because some tests may temporarily affect log collection. The RCCE engineer must ensure that the testing itself does not create a monitoring gap that a real adversary could exploit.

5.5 Recovery Resistance Analysis

Recovery resistance measures how difficult it is for an organization to fully remove an adversary from their environment. Real-world incidents have shown that adversaries who are partially remediated frequently return because not all persistence mechanisms were identified and removed.

RCCE engineers test recovery resistance by deploying multiple persistence mechanisms of different types across different systems. They then inform the defensive team that the environment has been compromised and measure the completeness and speed of the remediation effort. The questions being validated include whether the defensive team can identify all persistence mechanisms, whether the remediation process is thorough enough to prevent re-entry, and whether the organization can verify that remediation is complete. Recovery resistance testing is often the most operationally sensitive RCCE test because it requires close coordination with the defensive team and clear communication about the purpose and scope of the exercise.

CHAPTER 6: DETECTION ENGINEERING VALIDATION

6.1 Aligning Offense with Defense

Detection engineering validation closes the loop between offensive operations and defensive expectations. RCCE engineers do not just attack environments. They measure whether the defensive systems designed to detect those attacks actually function.

This requires a methodology that maps offensive actions to expected defensive responses. For every technique executed during RCCE operations, there should be a corresponding detection that should trigger. The validation measures whether it does trigger, how quickly it triggers, whether the alert contains sufficient context for investigation, and whether the alert leads to an appropriate response.

Detection engineering validation transforms RCCE operations from pure red team exercises into integrated security validation. The output is not just a list of successful attacks but a comprehensive assessment of whether the defensive architecture performs as designed.

6.2 Detection Trigger Validation

The most fundamental question in detection engineering is whether detections trigger at all. RCCE engineers validate this by executing known detection scenarios and measuring the result.

Detection trigger validation follows a systematic process. The engineer identifies a specific detection rule or analytic in the security monitoring stack. The engineer executes the behavior that the detection is designed to identify. The engineer measures whether the detection triggers. If the detection does not trigger, the engineer investigates why, which may reveal gaps in log collection, rule logic errors, threshold misconfigurations, or environmental conditions that prevent the detection from functioning.

This validation must cover both high-fidelity detections that target specific techniques and behavioral analytics that target patterns of activity. High-fidelity detections should trigger reliably for their target techniques. Behavioral analytics may have higher false negative rates, but they should trigger for sustained adversary activity that generates a pattern over time.

6.3 Alert Context Validation

A detection that triggers but provides insufficient context for investigation is almost as useless as a detection that does not trigger. RCCE engineers validate whether alerts contain the information that analysts need to begin effective investigation.

Effective alert context includes the identity or system that triggered the alert, the specific behavior that was detected, the timestamp and duration of the behavior, related events before and after the triggering event, severity classification, and recommended investigation steps.

Alert context validation often reveals a common failure mode where detections trigger on the right behavior but present the alert in a way that does not enable efficient triage. The analyst sees that something happened but must spend significant time gathering additional context before they can determine whether the alert represents a true positive or a false positive. This delay is exploitable by adversaries who operate during the gap between detection and effective response.

6.4 Timeline Reconstruction Validation

Investigations require the ability to reconstruct what happened, in what order, across which systems. RCCE engineers validate whether the logging and monitoring infrastructure supports timeline reconstruction for the attack chains they execute.

Timeline reconstruction validation tests whether logs from different sources can be correlated by time. It tests whether the event sequence of an attack chain is visible when logs from all relevant sources are combined. It tests whether gaps in logging create blind spots that prevent complete timeline reconstruction. It tests whether the tools available to analysts enable efficient timeline assembly or whether the process is so manual and slow that it cannot keep pace with a real incident.

This validation is particularly important for multi-stage attacks that span multiple systems and identity boundaries. Each system may log its portion of the attack, but if those logs cannot be correlated into a coherent timeline, the defensive team cannot understand the full scope of the compromise.

6.5 Response Appropriateness Validation

The final link in the detection chain is response. A detection that triggers and provides good context still fails if the response action is inappropriate, disproportionate, or too slow. RCCE engineers validate response appropriateness by measuring what happens after a detection triggers. Do analysts investigate the alert within the expected timeframe? Do they take the correct containment actions? Do those containment actions actually prevent the attacker from continuing? Do the containment actions cause unnecessary disruption to legitimate operations?

Response appropriateness validation often requires coordination with the security operations team because it involves measuring human performance, not just technical capability. The RCCE engineer must document findings about response timeliness and effectiveness in a way that supports improvement rather than blame.

6.6 Documentation of Detection Failures

When detections fail, RCCE engineers document the failure as evidence, not blame. The documentation includes exactly what technique was used, what detection should have triggered, why it did not trigger, what the impact of the gap is, and what remediation is recommended.

Detection failure documentation serves as the input to detection engineering improvement. It tells the detection engineering team exactly what to build or fix, with specific evidence of the gap and the adversary behavior that exploits it. This creates a continuous improvement loop where RCCE operations drive detection engineering maturity.

CHAPTER 7:

AUTONOMOUS DEFENSE AND SELF-HEALING TESTING

7.1 Validating Automation Under Pressure

Modern security architectures increasingly rely on automated defense mechanisms that can contain threats without human intervention. These mechanisms include automated endpoint isolation, identity revocation, network segmentation enforcement, and service rollback. RCF Domains 17 and 21 specifically address autonomous defense and self-healing infrastructure.

Automated defenses are powerful but dangerous. An automated system that makes incorrect containment decisions can cause more disruption than the attack it is trying to prevent. RCCE engineers validate that automated defenses are effective, bounded, and safe.

7.2 Automated Containment Testing

Automated containment is the capability to isolate a compromised system or identity without waiting for human authorization. RCCE engineers validate automated containment by deliberately triggering high-confidence attack scenarios and measuring the response. The validation tests whether the automated containment action executes within the expected timeframe, whether it is effective at preventing the attacker from continuing, whether it is bounded so it does not cascade beyond the intended scope, and whether it generates sufficient notification for the security operations team to understand what happened and why.

RCCE engineers also test for false positive handling. They execute benign activities that are similar to attack behaviors and measure whether the automated system correctly distinguishes between real attacks and legitimate operations. Automated containment systems with high false positive rates will either be disabled by frustrated operations teams or will cause chronic disruption that undermines organizational productivity.

7.3 Identity Revocation Testing

Automated identity revocation is the capability to disable a compromised identity across all systems and services. RCCE engineers validate that identity revocation is complete, fast, and reversible.

Complete means that revoking an identity actually removes all access, not just access to the primary authentication system. An identity that is revoked in Active Directory but retains active sessions in cloud services and SaaS applications is not fully revoked. RCCE engineers test for these gaps by attempting to use revoked identities across all connected systems.

Fast means that revocation takes effect within seconds, not minutes or hours. Token lifetimes, session caching, and synchronization delays can create windows where a revoked identity retains access. RCCE engineers measure these windows and validate that they are within acceptable bounds.

Reversible means that if an identity is incorrectly revoked, the revocation can be undone without requiring the user to recreate their account, reset all their credentials, and rejoin all their groups. RCCE engineers validate that revocation recovery procedures work and are documented.

7.4 Endpoint Isolation Testing

Automated endpoint isolation removes a compromised endpoint from the network to prevent lateral movement. RCCE engineers validate that isolation is effective, that it does not break critical services, and that isolated endpoints can be safely investigated and recovered.

Effective isolation means that the isolated endpoint truly cannot communicate with other network resources. RCCE engineers test for isolation bypass through alternative network interfaces, DNS tunneling, and out-of-band communication channels. Ineffective isolation creates a false sense of containment while the adversary continues operating.

Service impact testing validates that isolating an endpoint does not break dependent services. If isolating a compromised server causes a cascade of failures in downstream services, the automated system must account for this dependency and either handle it gracefully or defer to human decision-making.

7.5 Rollback and Recovery Logic Testing

Self-healing infrastructure is designed to automatically restore systems to a known-good state after compromise. This includes container redeployment, configuration rollback, and infrastructure-as-code reconciliation.

RCCE engineers validate self-healing by compromising systems in controlled ways and measuring whether the self-healing mechanisms detect the modification and restore the correct state. This includes testing whether modified configurations are detected and reverted, whether replaced binaries are identified and restored, whether compromised containers are terminated and replaced, and whether the self-healing process itself can be targeted and disabled by an attacker.

The most important finding in self-healing validation is often whether the adversary can compromise the self-healing mechanism itself. If the infrastructure-as-code repository can be modified, or if the deployment pipeline can be poisoned, then the self-healing mechanism becomes a persistence mechanism for the attacker.

CHAPTER 8: EVIDENCE CAPTURE UNDER ADVERSARIAL CONDITIONS

8.1 Why Evidence Discipline Matters

RCCE operations produce findings that may be used in executive decision-making, regulatory reporting, legal proceedings, and insurance claims. Every finding must be supported by evidence that can survive scrutiny from hostile parties, including adversaries who want to dispute the findings, lawyers who want to challenge the methodology, and regulators who want to verify compliance.

Evidence discipline is not bureaucracy. It is the difference between a finding that changes organizational behavior and a finding that is dismissed as unsubstantiated opinion. RCCE engineers who produce evidence that cannot be challenged produce findings that cannot be ignored.

8.2 Chain-of-Custody Discipline

Chain of custody is the documented record of who collected evidence, when it was collected, how it was stored, and who has accessed it since collection. Any break in the chain of custody creates an opportunity for opposing parties to challenge the integrity of the evidence.

RCCE engineers maintain chain-of-custody discipline from the moment evidence is created. Every artifact, including screenshots, packet captures, log extracts, system images, and memory dumps, is immediately hashed using a cryptographic hash function. The hash,

timestamp, and collector identity are recorded in the evidence log. The artifact is stored in a location with access controls that prevent unauthorized modification.

Chain-of-custody documentation must be continuous and unbroken. If evidence is transferred from one storage location to another, the transfer is documented. If evidence is accessed for analysis, the access is documented. If evidence is provided to a third party, the provision is documented. The chain of custody must be complete enough that any party can trace the history of any piece of evidence from collection to presentation.

8.3 Evidence Minimization

Evidence minimization is the practice of collecting only the evidence needed to support a finding, rather than collecting everything possible. This practice exists for legal, ethical, and operational reasons.

Legal reasons include data protection regulations that restrict the collection and processing of personal data. Evidence that includes personal data beyond what is necessary to support the finding creates regulatory exposure.

Ethical reasons include respecting the privacy of individuals whose data may be incidentally captured during RCCE operations. A memory dump that captures a user's personal browsing history is not relevant to a finding about privilege escalation and should not be retained.

Operational reasons include the cost of storing, protecting, and managing large volumes of evidence. Every piece of evidence that is retained must be protected with the same chain-of-custody discipline, which means that unnecessary evidence creates unnecessary operational burden.

RCCE engineers practice evidence minimization by defining what evidence is needed before collecting it, collecting the minimum necessary to support the finding, redacting personal data that is not relevant to the finding, and destroying unnecessary evidence according to a documented retention schedule.

8.4 Context Preservation

Evidence without context is meaningless. A screenshot of a command prompt showing administrative access is powerful evidence if it is accompanied by documentation of how that access was achieved, what controls were bypassed, and what the implications are. Without that context, it is just a screenshot.

RCCE engineers preserve context by documenting the complete attack chain that led to each finding, including every step, every tool, every technique, and every decision point. They document the environmental conditions at the time of the finding, including what controls were active, what monitoring was in place, and what the network and identity state was. They document the expected behavior of the controls that were tested and the actual behavior that was observed. They document the implications of the finding in terms of the risk it represents to the organization.

8.5 Integrity Validation

Evidence integrity means that evidence has not been modified since collection. RCCE engineers validate integrity through cryptographic hashing, which produces a fixed-size fingerprint of the evidence that changes if any bit of the evidence is modified.

Integrity validation is performed at collection, when the hash is first computed and recorded. It is performed at storage, to verify that the evidence arrived intact. It is performed at analysis, to verify that the evidence has not been modified since storage. It is performed at presentation, to verify that the evidence being presented is identical to the evidence that was collected.

RCCE engineers use hash algorithms that are currently considered cryptographically secure. The specific algorithm used is documented as part of the evidence record. If a hash algorithm is later found to be insecure, the evidence must be re-hashed with a secure algorithm while the original hash is still verifiable.

8.6 Timing and Provenance

Timing is critical to evidence because it establishes when events occurred relative to each other and relative to the operational timeline. RCCE engineers must ensure that all evidence timestamps are synchronized to a reliable time source and that the time source itself is documented.

Provenance is the record of where evidence came from and how it was produced. For RCCE operations, provenance includes the specific tool and version used to generate the evidence, the command or action that produced the artifact, the system from which the artifact was collected, and the identity of the engineer who collected it.

Timing and provenance documentation enables independent verification. A third party should be able to review the provenance record and understand exactly how each piece of evidence was created, without needing to contact the RCCE engineer who created it.

CHAPTER 9: ROSECOIN VAULT INTEGRATION FOR FIELD OPERATIONS

9.1 The Immutability Requirement

RCCE findings must be tamper-proof. The credibility of RCCE operations depends on the ability to prove that findings have not been modified, backdated, or fabricated after the fact. Traditional evidence storage systems rely on access controls and organizational trust to prevent evidence tampering. Rosecoin Vault provides something stronger: mathematical proof of evidence integrity anchored to an immutable distributed ledger.

Rosecoin Vault integration adds an independent verification layer to RCCE operations. When an RCCE engineer anchors a finding to the Rosecoin blockchain, they create a record that cannot be modified by anyone, including the engineer, the organization, or Rochester itself. This creates trust through mathematics rather than authority.

9.2 Hashing Operational Artifacts

Every operational artifact produced during RCCE operations can be hashed and the hash anchored to the Rosecoin Vault. This includes evidence artifacts such as screenshots, packet captures, and log extracts. It includes operational logs that record every action taken during the engagement. It includes reports and findings documents. It includes authorization and scope documents.

The hashing process creates a fixed-size fingerprint of each artifact. The artifact itself is not stored on the blockchain, only the hash. This preserves the confidentiality of the evidence while proving its existence and integrity. Anyone who has access to the original artifact can verify that it matches the anchored hash, which proves that the artifact has not been modified since it was anchored.

9.3 Anchoring Timelines

Time is one of the most disputed aspects of security incidents and assessments. When did the vulnerability exist? When was it discovered? When was it reported? Rosecoin Vault anchoring provides independent timestamp verification that does not depend on the RCCE engineer's local clock or the organization's systems.

RCCE engineers anchor timeline events by hashing status records at key operational milestones. The beginning of operations, the discovery of each finding, the completion of testing, and the delivery of reports can all be anchored with blockchain timestamps that provide independent proof of when each event occurred.

This timeline anchoring is particularly valuable in dispute scenarios. If an organization claims that a vulnerability was reported later than it actually was, the blockchain anchor provides independent proof of the actual reporting time. If a regulatory body requires evidence of when a security assessment was conducted, the blockchain anchors provide that evidence with mathematical certainty.

9.4 Linking Actions to Controls

RCCE operations validate specific RCF controls. Rosecoin Vault integration enables each finding to be cryptographically linked to the specific control it validates, creating a verifiable chain from adversarial action to control assessment to evidence.

This linking serves multiple purposes. It enables auditors to verify that specific controls have been tested and to trace the evidence chain for each control assessment. It enables organizations to track control validation history over time, with each validation anchored to an immutable record. It enables regulatory bodies to verify compliance without relying on the organization's self-reporting.

9.5 Creating Independent Verifiability

The ultimate purpose of Rosecoin Vault integration is to enable any authorized party to independently verify RCCE findings without trusting the RCCE engineer, the organization, or any intermediary.

Independent verification works as follows. The verifier obtains the original evidence artifact from the organization or the RCCE engineer. The verifier computes the hash of the artifact. The verifier checks the Rosecoin Vault for a matching hash. If the hash matches, the verifier has mathematical proof that the artifact existed in its current form at the time the hash was anchored.

This verification process does not require any trust in the parties involved. It relies entirely on the mathematical properties of cryptographic hashing and the immutability of the blockchain. The result is evidence that is as close to indisputable as current technology allows.

CHAPTER 10:

OPERATING DURING LIVE INCIDENTS

10.1 The Hardest Operational Scenario

Operating during a live incident is the most demanding scenario an RCCE engineer will face. The environment is in crisis. Incident responders are working to contain and remediate an active threat. Systems may be compromised. Communications may be unreliable.

Emotions are high. Decisions must be made quickly with incomplete information. RCCE engineers operating during live incidents are not there to compete with incident responders. They are there to reinforce them. Their unique contribution is the ability to validate whether containment actions are effective, whether the adversary has been fully removed, and whether recovery can proceed safely.

This chapter defines the operational doctrine for RCCE work during live incidents, including coordination protocols, interference avoidance, and evidence-first operations.

10.2 Coordinating with Incident Response

Coordination with the incident response team must be established immediately and maintained continuously. The RCCE engineer must understand the current incident status, the actions being taken, and the constraints under which the response team is operating. The coordination protocol begins with establishing communication channels. The RCCE engineer must have a direct communication channel with the incident commander that is separate from normal operational communications. This channel must be reliable even if the primary communication systems are compromised.

The RCCE engineer must clearly communicate their role, their authority, and their operational plan. The incident response team must understand what the RCCE engineer is doing and why, so they do not confuse RCCE activity with adversary activity. This is critical. An RCCE engineer who is operating during an active incident without the incident response team's knowledge risks being identified as an additional threat, which wastes response resources and creates confusion.

The coordination protocol must include regular status updates where the RCCE engineer shares findings with the incident commander and adjusts their approach based on the evolving incident situation.

10.3 Avoiding Interference

The first rule of RCCE operations during a live incident is to avoid making things worse. Every action taken by the RCCE engineer must be evaluated against the risk of interfering with containment, destroying evidence, or creating additional confusion.

Actions that modify the state of compromised systems must be avoided unless specifically requested by the incident commander. Scanning or probing activities that could trigger additional alerts must be coordinated with the monitoring team. Any RCCE activity that could be mistaken for adversary behavior must be announced to the response team before execution.

Interference avoidance also means recognizing when to step back. If the incident response team is in the middle of a critical containment action, the RCCE engineer should pause their operations and resume when the containment action is complete. The RCCE engineer's validation can wait. Containment cannot.

10.4 Supporting Containment Decisions

RCCE engineers can provide unique support during containment by validating whether proposed containment actions will be effective. The incident response team may believe that isolating a specific subnet will contain the adversary, but the RCCE engineer's knowledge of attack paths and lateral movement techniques may reveal that the adversary has already moved beyond that subnet.

This support requires the RCCE engineer to share their analysis of adversary movement patterns, which may be based on their own testing of the environment during prior RCCE operations or on real-time analysis of the incident data.

The RCCE engineer does not make containment decisions. That authority belongs to the incident commander. The RCCE engineer provides the information and analysis that enables better containment decisions.

10.5 Evidence-First Operations

During a live incident, evidence can be destroyed in minutes. Systems are being isolated, reimaged, patched, and restored. Logs are rolling over. Memory is being flushed. Every minute that passes reduces the amount of recoverable evidence.

RCCE engineers operating during live incidents prioritize evidence capture above all other activities. Before validating containment effectiveness, they capture evidence. Before analyzing adversary behavior, they preserve the artifacts that document it. Before recommending additional actions, they ensure that the evidence supporting their recommendations is secure.

Evidence-first operations during incidents follow a strict protocol. Identify evidence sources before they are modified. Capture evidence using methods that preserve integrity. Hash and record evidence immediately after capture. Store evidence in a location that is isolated from the compromised environment. Continue with validation and analysis only after critical evidence is secure.

10.6 Post-Incident Validation

After the active phase of an incident concludes, RCCE engineers perform post-incident validation to determine whether remediation was complete and whether the adversary has been fully removed from the environment.

Post-incident validation includes scanning for persistence mechanisms that may have survived remediation, testing whether the access paths used by the adversary have been closed, validating that compromised credentials have been revoked across all systems, and verifying that the adversary cannot re-enter through backup or recovery systems that may contain compromised state.

Post-incident RCCE validation often discovers that remediation was incomplete. This is not a criticism of the incident response team. It reflects the reality that incident response under pressure is difficult and that some persistence mechanisms are designed to survive exactly the kind of remediation that incident responders perform. RCCE post-incident validation provides the additional layer of scrutiny that catches what the initial remediation missed.

CHAPTER 11: POST-OPERATION REPORTING AND IMPROVEMENT LOOPS

11.1 Purpose of RCCE Reporting

RCCE reports exist to improve systems, not to assign blame. This is not a platitude. It is a design principle that drives every aspect of how RCCE findings are documented, presented, and consumed.

A blame-oriented report creates defensiveness. Stakeholders who feel blamed for security failures will resist the findings, undermine the remediation process, and make future RCCE operations more difficult. An improvement-oriented report creates action. Stakeholders who understand that findings represent opportunities to strengthen their security posture will support remediation and welcome future testing.

RCCE engineers must internalize this distinction because it affects not just the tone of the report but its structure, its content, and its recommendations. Every element of an RCCE report should answer the question: what needs to change and how?

11.2 Report Structure

RCCE reports follow a standardized structure that ensures completeness, consistency, and actionability.

The executive summary provides a high-level overview of the operation, including scope, objectives, key findings, and overall security posture assessment. It is written for decision-makers who need to understand the implications without reading the technical details. The

executive summary is typically one to two pages and should enable the reader to make informed decisions about remediation priority and resource allocation.

The methodology section documents the approach used, including adversary models, techniques employed, and scope boundaries. This section provides the context necessary for evaluating the findings. It enables the reader to understand not just what was found but how the testing was conducted and what limitations apply.

The findings section documents each finding with a standardized format that includes a description of the finding, the technique used to discover it, the evidence supporting the finding, the RCF control or controls that are affected, the potential impact if exploited by a real adversary, and specific remediation recommendations.

The evidence appendix contains the detailed evidence artifacts supporting each finding, organized by finding reference number with chain-of-custody documentation.

The improvement recommendations section provides prioritized recommendations for strengthening the security posture, organized by urgency and effort. This section goes beyond individual finding remediation to address systemic improvements that would prevent similar findings in the future.

11.3 Actionable Findings

Every finding in an RCCE report must be actionable. A finding that describes a problem without providing a clear path to remediation is incomplete.

Actionable findings include a specific description of the vulnerability or control failure, stated in terms that the responsible team can understand without translation. They include reproduction steps that enable the responsible team to verify the finding independently. They include evidence that proves the finding is real and not a false positive. They include impact assessment that explains what an adversary could achieve by exploiting the finding. They include remediation recommendations that are specific, implementable, and testable. Remediation recommendations must be realistic. Recommending that an organization rebuild its entire identity infrastructure in response to a single finding is not actionable.

Recommending specific configuration changes, policy updates, or control deployments that address the finding within normal operational constraints is actionable.

11.4 Evidence-Backed Narratives

The most effective RCCE reports combine individual findings into attack narratives that show how multiple weaknesses chain together to create significant risk.

An attack narrative begins with initial access and traces the complete path to critical asset compromise, documenting each step, each control that was tested, each control that failed, and each control that succeeded. The narrative shows stakeholders not just individual vulnerabilities but the complete picture of how an adversary would operate in their environment.

Attack narratives are supported by evidence at every step. Each step in the narrative references the specific evidence artifact that documents it. This creates a verifiable story that can be traced from beginning to end with proof at every point.

These narratives are particularly powerful for executive communication because they translate technical findings into business risk. An executive may not understand the significance of a misconfigured IAM policy, but they will understand a narrative that shows how that misconfiguration enables an attacker to move from a compromised laptop to the financial reporting database in four steps.

11.5 Control-Mapped Results

Every RCCE finding maps to one or more RCF domains and controls. This mapping serves multiple purposes.

For the organization, control mapping enables prioritized remediation based on which controls are failing and which domains are most affected. If multiple findings map to the same control or domain, that indicates a systemic issue that may require architectural remediation rather than point fixes.

For the RCF maturity model, control mapping provides evidence of control effectiveness that feeds into the organization's maturity assessment. Controls that consistently pass RCCE validation demonstrate maturity. Controls that consistently fail demonstrate areas where investment is needed.

For regulatory compliance, control mapping provides the evidence chain from adversarial testing to specific control assessment that regulators require.

11.6 Time-Bound Recommendations

Remediation recommendations in RCCE reports must include timeframes. A recommendation without a timeframe is a wish, not a plan.

Timeframes are categorized based on urgency and risk. Critical recommendations address findings that represent immediate exploitable risk and should be remediated within days. High recommendations address findings that represent significant risk and should be remediated within weeks. Medium recommendations address findings that represent moderate risk and should be remediated within a defined planning cycle. Low recommendations address findings that represent minor risk or are opportunities for improvement that should be incorporated into normal operational improvement processes.

Timeframes must be realistic and account for the organization's operational constraints. A recommendation to remediate a critical finding within twenty-four hours is unrealistic if the remediation requires a change management process that takes a week. RCCE engineers must work with the organization to establish timeframes that balance urgency with operational reality.

11.7 Feeding Back into RCF Maturity

RCCE reports are not endpoints. They are inputs to a continuous improvement loop that drives organizational security maturity.

Each RCCE operation produces data about which controls work, which controls fail, and which controls have not been tested. This data feeds into the organization's RCF maturity model, providing evidence-based assessment of security posture that is updated with every RCCE engagement.

Over time, the accumulation of RCCE data creates a longitudinal view of security maturity. Trends become visible. Controls that consistently fail despite remediation indicate systemic issues that require different approaches. Controls that consistently pass indicate areas of strength that can be maintained with less investment. New attack paths discovered in

successive RCCE operations indicate an evolving threat landscape that requires adaptive security architecture.

This feedback loop is the mechanism through which RCCE operations create lasting security improvement rather than point-in-time assessments that are outdated before the ink dries.

What This Manual Produces

An RCCE engineer who follows the doctrine in this manual produces outcomes that are fundamentally different from traditional penetration testing or compliance assessment. They produce verified control effectiveness, which is evidence-backed proof that specific controls function or fail under adversarial pressure. Not assumed effectiveness. Not documented effectiveness. Verified effectiveness.

They produce evidence-backed attack narratives that show how real adversaries would operate in the target environment, supported by proof at every step of the attack chain. They produce validated detection and response paths that demonstrate whether the defensive architecture performs as designed when subjected to realistic adversary behavior. They produce continuous improvement inputs that feed into the organization's security maturity model, driving measurable progress over time rather than point-in-time snapshots. They produce defensible security truth that can survive scrutiny from auditors, regulators, lawyers, and executives, because it is backed by immutable evidence rather than opinion. This is the opposite of penetration testing as theater. It is the operational application of the Rochester Cybersecurity Framework through adversarial validation that produces measurable, verifiable, and actionable results.

Final Statement to the Reader

This manual is not safe. It is not comfortable. It is not designed to make organizations feel secure. It is designed to make them resilient.

If you follow this doctrine, you will break systems. You will expose assumptions that stakeholders believe are certainties. You will create uncomfortable truths that some people would prefer to remain hidden. You will demonstrate that controls trusted for years do not function as advertised. You will show that environments believed to be secure contain paths that lead directly to catastrophic compromise.

That is the job.

The alternative is to wait for a real adversary to discover these truths, at a time and in a manner of their choosing, without the courtesy of documentation, without the discipline of evidence capture, and without the opportunity for controlled remediation.

RCCE engineers choose the former. They choose to find the truth before the adversary does, to document it with evidence that cannot be disputed, and to present it in a way that drives improvement rather than despair.

Security that cannot survive RCCE scrutiny cannot survive a real adversary. This is not a theory. It is an operational fact demonstrated by every breach that exploited a weakness that existing security processes failed to identify.

This manual provides the doctrine. The Rochester Cybersecurity Framework provides the structure. The Rosecoin Vault provides the proof.

The rest is execution.