Cyber Power Through Propagation Control

Gregory P. Carpenter
CW PENSEC, LLC
Independent Researcher

gregory.carpenter@cwpensec.com

Abstract: Enterprise cyber defense operates in telemetry-saturated environments, yet major incidents continue to escalate through rapid, large-scale propagation. Failure rarely stems from lack of visibility; it reflects inability to suppress spread across interconnected systems before compromise becomes self-sustaining. Contemporary cyber operations exploit trust relationships and identity infrastructure to convert initial access into distributed control, overwhelming response capacity despite extensive sensing. This exposes a measurement gap: defenders lack a defensible method for assessing whether an environment can contain propagation once intrusion begins. We propose a cyber-epidemiological model of intrusion propagation that treats the enterprise as a population distributed across a directed graph of trust, privilege, and service relationships. Compromise is modeled as a dynamic process rather than a binary outcome, and defensive cyber operations are treated as endogenous counterforces that directly compete with attacker-driven spread. We formalize a compartmental S-E-I-Q-R model in which susceptible, exposed, infectious, quarantined, and recovered states correspond to operational phases of cyber campaigns, including foothold establishment, activation, lateral movement, isolation, and remediation. This formulation extends prior cyber-epidemic models by explicitly incorporating enterprise identity architecture and active defensive response dynamics central to modern cyber conflict. Building on this model, we derive a cyber reproduction number, R_o, using a next-generation matrix approach. R_o represents the expected number of secondary infectious compromises generated by a single infectious node in an otherwise susceptible environment. Interpreted as an environmental threshold, R_o distinguishes escalation from containment: when R_o exceeds unity, propagation outpaces defensive action; when R_o falls below unity, containment dominates and campaigns tend to collapse. R_o characterizes environmental susceptibility to sustained cyber operations, independent of attacker identity or tooling. To operationalize the framework without experiments, the paper proposes a telemetry-driven method for estimating model parameters from standard security operations center and incident response artifacts. The result is a defensible analytic lens linking security architecture, defensive operations, and campaign outcomes, supporting prioritization decisions that constrain adversary freedom of action, reduce coercive cyber power, and strengthen deterrence by denial at both enterprise and national scales.

Keywords: Cyber epidemiology, Reproduction number (R_o), Attack graph modelling, Propagation thresholds, Security telemetry, Cyber resilience

1. Introduction

Cyber defenses operating in telemetry-rich environments still fail at scale. Modern enterprises deploy endpoint detection and response, network telemetry, identity and access logging, cloud audit trails, case management systems, and threat intelligence feeds. This is a campaign-level failure rather than a single-host problem. Adversaries require time and exploitable trust pathways to convert initial access into durable operational reach.

The unresolved question is structural: can the environment suppress propagation once compromise begins? In other spread-dynamic domains, the central analytic question is threshold: does the system dampen escalation or amplify it until externally constrained (European Systemic Risk Board 2020)?

This paper proposes a cyber-epidemiological model of intrusion propagation that treats the enterprise as a population, compromise as a dynamic state, and defensive operations as endogenous counterforces. Prior compartmental models of malware propagation (Mishra and Saini 2007; Mishra and Jha 2010; Hernández-Guillén et al. 2017; Hernández-Guillén et al. 2019; Chernikova et al. 2023) often assume homogeneous infection and limited integration of identity architecture or live defensive response. Contemporary intrusions instead exploit identity systems, delegated privilege, and shared infrastructure as primary propagation substrates, while defenders respond under active contest.

The core contribution is adapting the epidemiological reproduction number to cyber defense. We adapt R_0 as an environment-level threshold for escalation versus containment. As in epidemiology, R_0defines a threshold: when $R_0 > 1$, spread exceeds response capacity and campaign-scale compromise becomes structurally likely; when $R_0 < 1$, defensive action suppresses spread.

This framing links defensive posture to systemic resilience. Architecture shapes transmission; defensive operations shape infectious duration and recovery. Environments that suppress propagation constrain adversary freedom of action, whereas persistently elevated R_o indicates structural vulnerability.

We develop a formal theoretical model without empirical validation. No experiments are conducted. The objective is to define a mathematically grounded framework, derive a cyber reproduction number using a next-generation matrix approach, and specify a telemetry-driven parameter estimation method using standard SOC and incident-response data.

2. Background and Related Work

2.1 Cyber-Epidemiological Modeling of Malware and Compromise

The analogy between biological epidemics and malware propagation is established. Compartmental models, including SIR, SEIR, and variants, have been applied to worms and malicious objects to capture latency, reinfection, and countermeasure effects (Mishra and Saini 2007; Mishra and Jha 2010; Hernández-Guillén et al. 2017). Subsequent work incorporated stability analysis and explicit defensive action, showing that containment and recovery reshape propagation dynamics (Hernández-Guillén et al. 2019). Recent studies further formalize behavioral assumptions and forecasting (Chernikova et al. 2023; Kovtun et al. 2024; Zhou et al. 2023).

A persistent limitation is operational mapping. Many models assume homogeneous substrates and collapse compromises into a single infected state, underrepresenting staged activation, identity-mediated transmission, live containment, and asymmetric recovery. This work retains the epidemiological framing but defines compartments and parameters in cyber-operational terms, enabling direct mapping from identity-mediated transmission signals and endogenous response actions observed in logs and incident timelines to model dynamics.

2.2 Attack Graphs, Lateral Movement, and Structural Pathways

In parallel, cybersecurity research uses graph-based abstractions to enumerate adversary pathways through attack graphs and lateral movement analysis. These models formalize reachability, privilege relations, and exploit chaining to surface structural exposure and prioritize mitigation. Recent work emphasizes identity relationships, particularly within enterprise directory services, as dominant enablers of lateral movement (Herranz Oliveros et al. 2024), alongside broader efforts to analyze cyber related graphs at scale (Kuikka et al. 2024).

The limitation is threshold reasoning. Attack graphs describe how campaigns may progress but provide no environment level criterion for when progression becomes self-sustaining under active defense. This work treats the graph as the transmission substrate and adds population dynamics to capture nonlinear escalation and containment.

2.3 Kill Chains, TTP-Centric Models, and Campaign-Scale Behavior

Operational cyber defense is commonly structured around kill chains, adversary TTPs, and stage-based sequencing. These frameworks organize detection and response but emphasize procedural progression rather than spread. Campaign outcomes hinge on whether initial access expands into distributed control, driven by propagation opportunity and defensive tempo rather than technique presence alone.

Campaign oriented research embeds cyber operations within structured representations to analyze adversary behavior at scale (Lee and Choi 2023). Automated attack emulation shows how multi-step operations can be planned and scaled, accentuating the need for models that represent propagation under adversarial optimization (Wang et al. 2024).

2.4 Cyber Resilience, Systemic Risk, and Architectural Defense

Cyber resilience frames security as the capacity to sustain operations, recover, and adapt under attack. Resilience engineering emphasizes architecture, response capability, and systemic dependencies as outcome determinants (Bodeau et al. 2011; Ross et al. 2021). Incidents escalate when shared infrastructure, including identity services, backups, and management planes, is compromised and recovery becomes protracted.

Systemic risk modeling offers an analogy. Instability is typically threshold driven rather than a linear accumulation of local failures. The European Systemic Risk Board cyber crash model emphasizes propagation and amplification across interdependent systems (European Systemic Risk Board 2020), motivating cyber defense equivalents to reproduction thresholds.

Modern security architecture guidance reflects this logic. Zero Trust treats implicit trust as the primary risk vector and advocates continuous verification, least privilege, segmentation, and assumed breach (NIST 2020).

In propagation terms, these controls reduce transmission probability, constrain reachability, and shorten infectious periods through faster containment.

2.5 Incident Handling, SOC Measurement, and Telemetry as Instrumentation

Cyber defense is executed through detection, response, and recovery. NIST incident handling guidance formalizes lifecycle management and emphasizes disciplined timelines, containment, and recovery validation (NIST 2025). Empirical SOC studies show effectiveness depends on decision and execution latency as well as detection capability (Forsberg and Frantti 2023).

Here, telemetry is instrumentation for system level inference. The objective is not attribution but estimating whether containment and recovery suppress propagation, preventing localized compromise from becoming campaign sustaining.

2.6 Summary of the Gap

A common limitation persists. Epidemic models capture nonlinear spread but lack enterprise-operational grounding. Attack graphs expose structure but provide no escalation threshold. Resilience frameworks emphasize architecture yet lack a scalar containment criterion. The reproduction-number framework integrates these perspectives by coupling graph structure, compartmental dynamics, and a threshold metric grounded in operational telemetry.

3. A Compartmental Epidemiological Model of Cyber Propagation

Reasoning about cyber intrusion at campaign scale requires moving beyond binary notions of compromise toward explicit modeling of intrusion progression. Cyber operations do not transition directly from access to impact; they evolve through foothold establishment, activation, lateral movement, and continuous interaction with defensive response. These dynamics map naturally to compartmental epidemiological models when defining compartments in operational cyber terms rather than biological analogy. Prior work has applied compartmental abstractions to malware and worm propagation, but typically assumes homogeneous infection and offers limited treatment of enterprise identity architecture and live defensive response dynamics central to modern intrusions (Mishra and Jha 2010; Hernández-Guillén et al. 2019; Chernikova et al. 2023).

3.1 System Representation

We model the enterprise as a directed graph $G = (V, E)$, where nodes represent hosts, identities, services, or control-plane elements, and edges represent feasible pathways for adversary action. Edges encode operational reachability and may arise from network connectivity, authentication trust, shared credentials, delegated privilege, or application-level access. This abstraction treats network topology and identity architecture as a single propagation substrate. While graph-based representations are well established in attack-graph and lateral-movement modeling, they are rarely integrated with population-dynamics frameworks capable of expressing system-level threshold behavior (Herranz-Oliveros et al. 2024; Kuikka et al. 2024).

At time t, each node occupies one of five compartments:

- **Susceptible (S):** Reachable nodes with exploitable conditions, such as exposed services, valid credentials, or trust relationships, that remain uncompromised.
- **Exposed (E): Nodes where initial access has been established via credential capture, payload delivery, or execution, but not yet used for propagation.**
- **Infectious (I):** Nodes actively leveraged for lateral movement, privilege escalation, command and control, or coordinated action.
- **Quarantined (Q):** Nodes isolated by defensive action, including segmentation enforcement, account restriction, or endpoint containment.
- **Recovered (R):** Nodes remediated to remove attacker access and reduce near-term susceptibility through reimaging, patching, or credential rotation.

Prior cyber-epidemic models frequently collapse compromise into a single infected state, obscuring dwell time, staged activation, and defensive interventions observed in real incidents (Mishra and Saini 2007; Hernández-Guillén et al. 2017). The compartmental distinctions introduced here reflect operational phases encountered in defensive cyber operations and incident response.

Notation Summary

$G = (V, E)$: Directed enterprise graph
S, E, I, Q, R: Compartment states
β: Effective transmission rate
σ: Activation rate
γ: Isolation rate
ρ: Recovery rate
R_0: Cyber reproduction number
$\rho(\cdot)$: Spectral radius

3.2 Transition Dynamics

Cyber propagation reflects a continuous contest between attacker-driven expansion and defender-driven isolation and recovery. This interaction is captured through four parameters governing transitions among compartments, each grounded in observable cyber-operational processes:

- **Effective transmission rate** (β)**:** The probability that an infectious node compromises a susceptible neighbor per unit time, conditioned on connectivity, vulnerability, and privilege.
- **Activation rate** (σ)**:** The rate at which exposed nodes transition to active use for propagation, reflecting attacker staging and dwell time.
- **Defensive isolation rate** (γ)**:** The rate at which infectious nodes are detected and isolated through defensive action.
- **Recovery or hardening rate** (ρ)**:** The rate at which compromised nodes are remediated such that attacker access is removed and near-term susceptibility reduced.

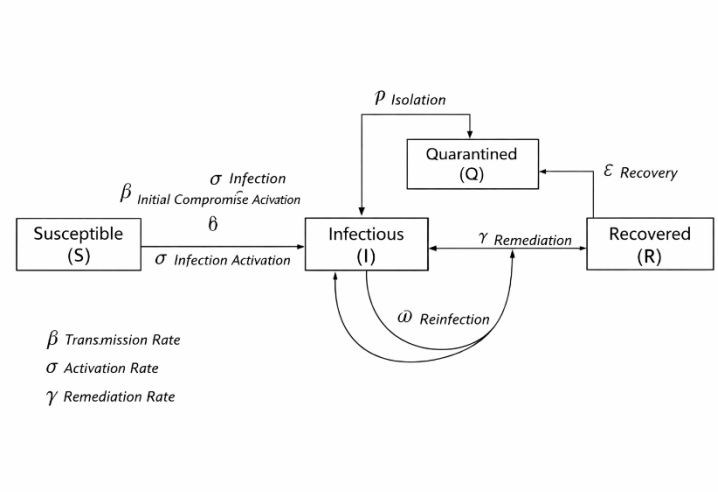

Figure 1: Compartmental model of cyber-attack propagation

In discrete time, the system is expressed as:

$$\begin{aligned} S_{t+1} &= S_t - \beta S_t I_t \\ E_{t+1} &= E_t + \beta S_t I_t - \sigma E_t \\ I_{t+1} &= I_t + \sigma E_t - (\gamma + \rho) I_t \\ Q_{t+1} &= Q_t + \gamma I_t \\ R_{t+1} &= R_t + \rho I_t \end{aligned}$$

Mixing and graph-structure note. The term $\beta S_t I_t$adopts a mean-field approximation, treating interaction between infectious and susceptible nodes as homogeneous across feasible edges. In practice, enterprise environments exhibit degree heterogeneity, privilege stratification, and control-plane dependencies represented by $G = (V, E)$. The mean-field formulation is used here to isolate threshold behavior and preserve tractable parameter estimation from operational telemetry. Section 4 interprets the resulting R_0 as an environment-level threshold under this approximation. Network-heterogeneous extensions, including spectral-radius–weighted thresholds, are left for future development.

3.3 Defensive Cyber Operations as Endogenous Forces

A defining feature of this framework is treating defensive cyber operations as endogenous to propagation rather than external mitigations applied after compromise. Detection, isolation, and remediation are modeled as forces that directly compete with attacker-driven spread at each timestep. Campaign escalation therefore depends not on detection in isolation, but on whether defensive action shortens the effective infectious period sufficiently to suppress further propagation.

From a security-architecture perspective, key model parameters are shaped by structural design choices. Flat networks, excessive privilege, and permissive trust relationships increase effective transmission β; segmentation, least-privilege, and constrained delegation reduce β by limiting reachable attack paths. Investment in SOC staffing, automation, and response orchestration primarily affects isolation $\gamma$by reducing time to containment. Together, operations and architecture act on complementary components of the propagation dynamic, consistent with cyber resiliency engineering that treats security as an emergent property of system design and adaptive response (Bodeau et al. 2011; Ross et al. 2021).

3.4 Campaign-Scale Interpretation

At campaign scale, the model captures behavior characteristic of contemporary offensive cyber operations. Advanced actors depend on conditions that sustain movement over time. Once multiple infectious nodes are established, propagation accelerates nonlinearly as lateral actions compound and response capacity saturates. This dynamic is poorly represented by linear kill-chain or TTP-centric models. By contrast, it emerges naturally in a population-dynamics formulation that explicitly models competing rates of propagation and containment.

4. Derivation and Interpretation of the Cyber Reproduction Number (R_o)

These transition dynamics provide the basis for deriving an environment-level propagation threshold, formalized in the following section as a cyber reproduction number. While compartmental dynamics describe how cyber intrusions evolve, defenders require a scalar measure that indicates whether a campaign is likely to escalate or collapse. In epidemiology, this role is served by the basic reproduction number, R_o. We adapt this construct to cyber defense by defining a cyber reproduction number that captures environmental susceptibility to propagation.

4.1 Definition and Derivation of Cyber R_o

Cyber $R_0$is defined as the expected number of secondary infectious nodes generated by a single infectious node introduced into an otherwise susceptible environment. Crucially, R_0 is a property of the defensive environment rather than a specific attacker, malware family, or exploit chain. As in epidemiology and systemic risk analysis, R_0serves as a threshold quantity: when it exceeds unity, propagation becomes self-sustaining; when it falls below unity, containment dominates and escalation collapses (European Systemic Risk Board 2020).

Following standard epidemiological practice, R_0 is derived using a next-generation matrix approach. Let the vector of epidemiologically active compartments be defined as

$$\boldsymbol{x} = \begin{bmatrix} E \\ I \end{bmatrix},$$

where Erepresents exposed nodes that have established an initial foothold but are not yet propagating, and Irepresents infectious nodes actively used for lateral movement and campaign expansion. Only infectious nodes generate new exposures.

The system dynamics are decomposed into two components: a matrix F, capturing the rate at which new exposures are generated, and a matrix V, capturing transitions among infected states and removal through isolation and recovery. Linearizing the system around the disease-free equilibrium yields:

$$\boldsymbol{F} = \begin{bmatrix} 0 & \beta \\ 0 & 0 \end{bmatrix}, \boldsymbol{V} = \begin{bmatrix} \sigma & 0 \\ -\sigma & \gamma + \rho \end{bmatrix}.$$

The next-generation matrix is defined as:

$$\boldsymbol{K} = \boldsymbol{F}\boldsymbol{V}^{-1},$$

and the cyber reproduction number is given by:

$$R_0 = \rho(\boldsymbol{K}),$$

where $\rho(\cdot)$denotes the spectral radius.

For this compartmental structure under the mean-field mixing assumption introduced in Section 3.2, the dominant eigenvalue admits a closed-form expression:

$$R_0 = \frac{\beta}{\gamma + \rho}.$$

In structured networks, the threshold generalizes to a spectral-radius form:

$$R_0 = \rho(\boldsymbol{FV}^{-1})$$

where the transmission component is weighted by the adjacency structure of G. The scalar expression presented here isolates the balance between transmission and removal while preserving analytic tractability and practical identifiability from telemetry data.

This result has a direct operational interpretation. The escalation threshold is governed by the balance between effective transmission and removal through containment and recovery. While the activation rate σinfluences the tempo of campaign progression, it does not alter the threshold condition itself. As a result, cyber R_0separates architectural exposure, captured by β, from defensive response capacity, captured by $\gamma$and ρ.

Interpreted operationally, $R_0 > 1$indicates that each compromised node generates more than one additional infectious node before isolation or remediation, making campaign-scale escalation structurally likely.

Conversely, $R_0 < 1$indicates that defensive action suppresses propagation faster than attackers can expand it, causing campaigns to collapse regardless of individual techniques or tooling. This threshold behavior mirrors empirical observations from large-scale ransomware and advanced persistent threat campaigns, where containment failure is systemic rather than exploit-specific (Lee and Choi 2023).

4.2 Relationship to Offensive and Defensive Cyber Operations

From an offensive perspective, cyber operations can be interpreted as increasing effective R_o. Techniques such as credential dumping, trust abuse, and living-off-the-land tooling increase $\beta$and σ while minimizing γ. Campaign success depends not on stealth alone but on sustaining R_o above unity long enough to generate irreversible effects (Wang et al. 2024).

From a defensive perspective, operations aim to suppress R_o. Detection, response automation, and remediation primarily act on γ and ρ, while architectural controls act on β. This clarifies the strategic value of architecture: reducing propagation potential before an incident begins rather than racing spread after compromise.

4.3 Cyber Power and Structural Resilience

At a strategic level, R_o provides a quantitative lens on cyber power and resilience. Environments that consistently maintain R_o below unity are inherently resistant to coercion through cyber means. Environments with high R_o invite repeated exploitation regardless of attacker identity. This framing aligns with national cyber strategy emphasis on sustained operational advantage rather than episodic compromise (U.S. Department of Defense, 2023; The White House 2023).

4.4 Practical Implications

The value of R_0 lies in decision support. Interval estimates are sufficient to determine threshold regime classification. By shifting focus from individual alerts to systemic thresholds, R_0 supports a more strategically grounded approach to defensive cyber operations.

5. Parameter Estimation from Security Telemetry

This paper proposes a telemetry-driven approach for parameterizing the cyber propagation model. The objective is to define a defensible and implementable method by which a security operations function can estimate model parameters in practice. In epidemiology and systemic risk analysis, model utility depends less on precise parameter values than on whether estimates are sufficiently grounded to distinguish regimes, specifically whether $R_0 > 1$or $R_0 < 1$, and to support comparative assessment of controls (Kovtun et al., 2024). The same principle applies to cyber defense.

The parameters of interest, effective transmission probability (β), exposure-to-compromise activation rate (σ), defensive isolation rate (γ), and recovery or hardening rate (ρ), are not abstract constants. They are operational properties of an environment, observable through standard defensive telemetry including identity logs, endpoint telemetry, network flow records, and incident timelines. The estimation workflow treats the SOC as the measurement instrument and the enterprise as the population under observation.

5.1 Telemetry Sources and Observability

Most enterprises already collect the data required to estimate these parameters, though rarely for system-level inference.

Identity and access telemetry: Authentication and authorization logs from directory services and cloud identity providers provide evidence of credential misuse, lateral authentication, privilege escalation, and anomalous access patterns. Directory artifacts such as group membership changes, delegated rights, and trust relationships expose structural conditions that influence transmission likelihood.

Endpoint and EDR telemetry: Process creation, command-line execution, module loading, and network connections indicate post-compromise activation (E→I), lateral tooling, and propagation attempts. Containment actions including host isolation and process termination provide timestamps relevant to defensive isolation (γ).

Network telemetry: NetFlow, firewall, proxy, DNS, and segmentation enforcement logs provide visibility into reachable paths, attempted connections, beaconing behavior, and blocked propagation, informing both graph construction and effective connectivity.

Incident response artifacts: Case management systems, playbook execution logs, and analyst notes provide high-fidelity timing data for detection-to-isolation and isolation-to-recovery intervals. Post-incident reviews support validation of assumptions and interpretation of anomalous dynamics.

These sources align with established guidance for incident handling, SOC performance measurement, and cyber resilience, though they are typically applied tactically rather than analytically (NIST, 2025; Forsberg and Frantti 2023; Ross et al. 2021).

Observability remains imperfect. Attackers may operate below logging thresholds, and telemetry coverage varies across environments. Accordingly, the estimation approach prioritizes conservative inference and explicitly accounts for uncertainty to avoid overstating defensive capability.

5.2 Estimating Effective Transmission (β)

In this framework, β denotes the probability that an infectious node generates a new exposed node per unit time along feasible edges of the enterprise graph $G = (V, E)$.

Let:

- $\mathcal{E}_t \subseteq E$ denote the set of feasible directed edges active during observation interval t,
- C_t denote the number of observed feasible propagation attempts along edges in $\mathcal{E}_t$,
- X_t denote the number of newly established exposed nodes attributable to infectious sources during t.

Under a Bernoulli transmission assumption along feasible contacts, conditional on observability, the maximum-likelihood estimator for β is:

$$\hat{\beta} = \frac{\sum_t X_t}{\sum_t C_t}$$

This estimator corresponds to a binomial likelihood model over observed feasible contacts:

$$X_t \sim \text{Binomial}(C_t, \beta)$$

Bias and Observability Adjustment

In practice, C_t is subject to undercount bias due to incomplete telemetry and unobserved propagation attempts. Let $\pi \in (0,1]$ denote estimated contact observability. A conservative lower-bound estimator is:

$$\hat{\beta}_{\text{lower}} = \frac{\sum_t X_t}{\sum_t C_t}$$

An upper-bound estimator adjusts for incomplete contact observability:

$$\hat{\beta}_{\text{upper}} = \frac{\sum_t X_t}{\pi \sum_t C_t}$$

where π may be estimated from logging coverage metrics, red-team ground truth exercises, or retrospective detection rates.

Thus β is treated as an interval parameter:

$$\beta \in [\hat{\beta}_{\text{lower}}, \hat{\beta}_{\text{upper}}]$$

Stratified Transmission

Because transmission probability varies by identity tier, network segment, and control-plane function, stratified estimators are recommended:

$$\hat{\beta}_k = \frac{\sum_t X_{t,k}}{\sum_t C_{t,k}}$$

where $k$indexes structural strata of G.

5.3 Estimating Activation (σ)

The activation rate σ captures the transition from exposure to infectious activity (E→I), reflecting the delay between initial foothold establishment and active propagation. σ can be estimated from the interval between first evidence of foothold establishment and first observed lateral or coordinated activity. Where telemetry resolution is limited, bounding σ rather than point estimation is often sufficient for threshold analysis.

σ is strongly shaped by offensive tradecraft. Hands-on-keyboard intrusions exhibit longer activation periods, while automated or monetization-driven campaigns shorten activation windows. Comparative analyses show ransomware operations consistently activate faster than espionage-focused intrusions, informing σ as an expectation of operational tempo rather than attribution (Lee and Choi 2023).

5.4 Estimating Defensive Isolation *(γ)*

The isolation rate γ represents the transition from infectious to quarantined state (I→Q), capturing the combined effects of detection, analyst decision-making, and containment execution. Operationally, γ reflects the inverse of the effective infectious duration prior to isolation.

Let $t_I^{(k)}$denote the timestamp of first observable infectious activity for incident instance k, and let $t_Q^{(k)}$denote the timestamp at which isolation or containment action is completed. Define the infectious duration prior to isolation as:

$$\Delta t_\gamma^{(k)} = t_Q^{(k)} - t_I^{(k)}.$$

Under an exponential removal assumption, a consistent estimator for γ is:

$$\hat{\gamma} = \frac{1}{\Delta \bar{t}_\gamma},$$

where $\Delta \bar{t}_\gamma$is the mean infectious-to-isolation interval across observed cases.

In practice, γ may be decomposed into detection latency, decision latency, and execution latency. This decomposition supports targeted operational improvement, as investments in automation, playbook refinement, or analyst staffing affect distinct components of isolation time. Empirical SOC studies indicate that reducing decision and execution latency frequently produces greater threshold impact than marginal improvements in detection accuracy (Forsberg and Frantti 2023; NIST 2025).

5.5 Estimating Recovery and Hardening (*ρ*)

The recovery rate ρ captures the transition from compromised to remediated state (I→R or Q→R), representing removal of adversary access and reduction of near-term susceptibility through reimaging, credential rotation, patching, and closure of exploited trust pathways.

Let $t_Q^{(k)}$denote the timestamp at which node $k$is isolated and let $t_R^{(k)}$denote the timestamp at which verified remediation and hardening are completed. Define the containment-to-recovery interval as:

$$\Delta t_\rho^{(k)} = t_R^{(k)} - t_Q^{(k)}.$$

Assuming exponential recovery dynamics, a consistent estimator for ρ is:

$$\hat{\rho} = \frac{1}{\Delta \bar{t}_\rho},$$

where $\Delta \bar{t}_\rho$is the mean isolation-to-recovery interval across observed cases.

Recovery duration varies substantially by asset class. Identity providers, backup infrastructure, and management planes frequently exhibit longer recovery intervals and disproportionately influence systemic resilience.

Stratified estimation by system criticality or functional tier is therefore recommended to prevent high-impact assets from being diluted in aggregate averages (Bodeau et al. 2011; Ross et al. 2021).

As with γ, recovery intervals may be right-censored when remediation remains ongoing or validation is delayed. Interval-based reporting and conservative classification prevent overstating effective hardening capacity.

5.6 Handling Uncertainty and Bias

Telemetry-based estimation is inherently biased by incomplete observability. Attackers may evade detection, defenders may identify compromise retrospectively, and logging coverage is uneven across systems. These effects introduce undercount bias in transmission events and right-censoring in isolation and recovery intervals.

To address these limitations, parameters are represented as bounded intervals rather than single-point estimates:

$$\beta \in [\beta_{\min}, \beta_{\max}], \gamma \in [\gamma_{\min}, \gamma_{\max}], \rho \in [\rho_{\min}, \rho_{\max}]$$

Bounds are derived from conservative event classification, earliest plausible infection timestamps, and latest plausible containment or remediation timestamps. Sensitivity analysis is conducted by propagating parameter intervals through the reproduction-number expression:

$$R_0 = \frac{\beta}{\gamma + \rho}$$

yielding an uncertainty interval:

$$R_0 \in \left[R_{0,\min}, R_{0,\max}\right]$$

Local Sensitivity

Local sensitivity of the threshold to parameter perturbation is given by:

$$\frac{\partial R_0}{\partial \beta} = \frac{1}{\gamma + \rho}$$
$$\frac{\partial R_0}{\partial \gamma} = -\frac{\beta}{(\gamma+\rho)^2}$$
$$\frac{\partial R_0}{\partial \rho} = -\frac{\beta}{(\gamma+\rho)^2}$$

These expressions permit ranking of architectural versus operational control leverage under perturbation.

The operational objective is regime classification rather than precise prediction. If $R_{0,\min} > 1$, escalation is structurally robust to uncertainty. If $R_{0,\max} < 1$, containment is structurally robust. When the interval spans unity, architectural or operational adjustments should be prioritized until threshold classification becomes unambiguous. This approach aligns with established robustness practice in epidemiology and systemic risk modeling (Kovtun et al. 2024).

5.7 Cyber Threat Intelligence as Parameter Refinement

Within this framework, cyber threat intelligence refines parameters rather than assigns attribution. Intelligence on tooling, preferences, and operational tempo informs expectations for β and σ, while lessons from prior campaigns inform *γ* and *ρ*. This positions intelligence as an input to defensive posture management rather than retrospective explanation, aligning with campaign-centric views of cyber operations (Lee and Choi, 2023).

5.8 Practical Deployment Considerations

The approach is deployable within existing SOC environments and requires no new sensors or experimental instrumentation. It relies on retrospective incident reconstruction and reassessment following architectural change. The model does not depend on attribution; it evaluates structural propagation conditions independent of actor identity or tooling.

The objective is a defensible determination of whether an environment is predisposed to outbreak-scale compromise. When R_o exceeds unity, the environment should be treated as pre-failed and architectural intervention prioritized. When $R_o < 1$, defensive operations can shift from crisis response toward sustained resilience.

6. Conclusion

Enterprise cyber defense failures are rarely attributable to insufficient visibility. They emerge when propagation outpaces containment once compromise is established. This paper reframes the problem as a threshold condition rather than isolated detections, advancing a cyber-epidemiological model that treats intrusion propagation as a population-level dynamic governed by architecture, trust relationships, and defensive response capacity.

By formalizing a compartmental S–E–I–Q–R model grounded in operational cyber states, this work extends prior cyber epidemic models to reflect contemporary enterprise conditions, particularly identity-driven propagation and active defensive intervention. The derivation of a cyber reproduction number, R_o, introduces a scalar, environment-level metric that discriminates between escalation and containment. When R_o exceeds unity, campaign-scale compromise becomes structurally probable independent of attacker tooling or intent. When R_o is driven below unity, defensive operations dominate and campaigns tend to fail.

The value of R_o lies not in precise prediction but in decision support. The telemetry-driven estimation approach presented here demonstrates that defenders can approximate propagation thresholds using existing operational data, without experiments, attribution, or additional sensing infrastructure. This enables comparative evaluation of architectural controls, response tempo, and recovery effectiveness, each directly linked to campaign-level outcomes.

Operationally, the framework shifts emphasis from alert pursuit to managing propagation potential. Strategically, it frames cyber resilience as a property of system design rather than reactive capability. Environments that sustain elevated R_o remain vulnerable to coercive cyber operations regardless of attribution or intent.

By introducing a mathematically grounded yet operationally implementable threshold lens, this work provides a missing analytical bridge between security architecture, defensive cyber operations, and campaign-scale risk. The result is a defensible basis to assess containment versus escalation before adversaries convert access into sustained control.

Ethics Declaration: This research is theoretical and conceptual in nature. It involved no human participants, operational systems, or experimental interventions, and therefore did not require ethical approval.

AI Declaration: AI-based tools were used to prepare this manuscript, supporting language refinement and structural clarity. All technical content, modeling decisions, analysis, and conclusions were developed, reviewed, and validated by the author, who retains full responsibility for the work.

References

Bodeau, D., Graubart, R., Picciotto, J. and McQuaid, R. (2011) *Cyber Resiliency Engineering Framework*. MITRE Technical Report.

Chernikova, A., Ivanov, A., Korolev, V. and Klyuev, S. (2023) 'Modeling self-propagating malware with epidemiological models', *Applied Network Science*, 8(1), art. 98.

European Systemic Risk Board (2020) *The making of a cyber crash: a conceptual model for systemic risk in the financial sector*. Occasional Paper Series No. 16.

Forsberg, J. and Frantti, T. (2023) 'Technical performance metrics of a security operations center', *Computers & Security*, 135, 103529. https://doi.org/10.1016/j.cose.2023.103529

Hernández-Guillén, J.D., Martín del Rey, A. and Hernández Encinas, L. (2017) 'Study of the stability of a SEIRS model for computer worm propagation', *Physica A*, 479, pp. 411–421. DOI: 10.1016/j.physa.2017.03.023

Hernández-Guillén, J.D., Martín del Rey, A. and Casado Vera, R. (2019) 'Security countermeasures of a SCIRAS model for advanced malware propagation', *IEEE Access*, 7, pp. 135472–135478. DOI:10.1109/ACCESS.2019.2942809

Herranz-Oliveros, D., García-Teodoro, P. and López-Rodríguez, C. (2024) 'Unsupervised learning for lateral-movement-based threat mitigation in Active Directory attack graphs', *Electronics*, 13(19), 3944. https://doi.org/10.3390/electronics13193944

Kovtun, V., Grochla, K., Al-Maitah, M., Aldosary, S. and Gryshchuk, T. (2024) 'Cyber epidemic spread forecasting based on the entropy-extremal interpretation of the SIR model', *Egyptian Informatics Journal*, 28, 100572. https://doi.org/10.1016/j.eij.2024.100572

Kuikka, V., Pykälä, L., Takko, T. and Kaski, K.K. (2024) 'Network modelling in analysing cyber-related graphs', *arXiv preprint*, arXiv:2412.14375. https://doi.org/10.3389/fcpxs.2025.1620260

Lee, I., and Choi, C. (2023) 'Camp2Vec: embedding cyber campaign with ATT&CK framework for attack group analysis', *ICT Express*, 9(6), pp. 1065–1070. https://doi.org/10.1016/j.icte.2023.05.008

Mishra, B.K. and Jha, N. (2010) 'SEIQRS model for the transmission of malicious objects in computer networks', *Applied Mathematical Modelling*, 34(3), pp. 710–715. https://doi.org/10.1016/j.apm.2009.06.011

Mishra, B.K. and Saini, D.K. (2007) 'SEIRS epidemic model with delay for transmission of malicious objects in computer networks', *Applied Mathematics and Computation*, 188(2), pp. 1476–1482. https://doi.org/10.1016/j.amc.2006.11.012

National Institute of Standards and Technology (NIST) (2020) *SP 800-207: Zero Trust Architecture*. Gaithersburg, MD: NIST.

National Institute of Standards and Technology (NIST) (2025) *SP 800-61r3: Computer Security Incident Handling Guide*. Gaithersburg, MD: NIST.

Ross, R., Pillitteri, V., Graubart, R., Bodeau, D. and McQuaid, R. (2021) *SP 800-160 Vol. 2 Rev. 1: Developing Cyber-Resilient Systems*. Gaithersburg, MD: NIST.

U.S. Department of Defense (2023) *Department of Defense Cyber Strategy*. Washington, DC: Department of Defense.

Wang, L., Zhang, Y., Li, Z., and Chen, X. (2024) 'From sands to mansions: towards automated cyberattack emulation with classical planning and large language models', *arXiv preprint*, arXiv:2407.16928.

The White House (2023) *National Cybersecurity Strategy*. Washington, DC: U.S. Government.