

EPic Launch



Critical Product Capabilities Brief

Company: R6 Security | Founder & CEO: Zsolt Nemeth



R6 Security, led by founder and CEO Zsolt Nemeth, is redefining cyber defense through its Phoenix platform powered by Automated Moving Target Defense (AMTD) technology. In an era of increasingly intelligent adversaries and ephemeral cloud-native environments, Phoenix offers a dynamic, proactive approach to cybersecurity. By combining chaos engineering with intelligent automation, R6 disrupts the attack surface in real-time—rendering threats ineffective before they can execute. This isn't just another layer in the defense stack; it's a constantly shifting foundation, purpose-built for the speed and complexity of AI-driven, containerized, multi-cloud workloads.

1. What makes R6 Security's Phoenix solution or AMTD (Automated Moving Target Defense) technology stand out from traditional security platforms?

Phoenix's Automated Moving Target Defense constantly randomizes system configurations and attack surfaces at runtime—something traditional static defense models cannot achieve. This shifting topology eliminates attacker dwell time and invalidates reconnaissance. Unlike signature-based or reactive systems, Phoenix proactively confuses and disables adversaries before an exploit path is even viable.

2. Can you elaborate on how Phoenix leverages chaos engineering and moving target defense for proactive protection?

Phoenix adopts principles of chaos engineering to inject controlled unpredictability into workloads, containers, and environments. By continually modifying elements like memory addresses, service ports, and network topology, it creates an environment where threats can't rely on fixed entry points or behavior patterns. The platform simulates failure conditions and adversarial tactics to harden systems before attackers can exploit them.



3. Why is moving target defense critical for protecting AI and containerized workloads today?

Modern AI workloads and containerized systems are inherently dynamic and scalable—but this also creates new vulnerabilities. Static security measures struggle to adapt in time. AMTD ensures that these assets never present a predictable surface, making it exponentially harder for attackers to gain persistence or lateral movement. This is especially critical for environments where traditional perimeter controls no longer apply.

4. How does R6 address the challenges of securing cloud-native environments, especially for enterprises with hybrid or multi-cloud setups?

R6 Phoenix is designed to be cloud-agnostic, seamlessly integrating into Kubernetes, serverless, and container orchestration environments across AWS, Azure, GCP, and hybrid architectures. It leverages policy-as-code and DevSecOps automation to ensure continuous protection and runtime flexibility regardless of where workloads reside.

5. How do you use AI for R6 and what aspects of threat detection and remediation are truly intelligent, versus engineered with smart rules?

Phoenix uses machine learning to identify anomalous behavior patterns in real time, allowing it to adapt to previously unseen attack vectors. Unlike simple heuristics or rule-based engines, its AI models evolve based on behavioral baselines and environmental signals, enabling autonomous threat disruption and intelligent remediation without human intervention.



6. What challenges did your team encounter when building Phoenix, and how did you overcome them?

One of the major challenges was engineering a platform that introduces high variability without compromising system performance or stability. The R6 team tackled this through extensive testing in simulated adversarial environments and through close collaboration with early enterprise partners to fine-tune performance thresholds while maintaining chaos-driven resilience.

7. How does Phoenix integrate with existing DevOps workflows, and what's the learning curve for security and platform teams?

Phoenix integrates directly into CI/CD pipelines and DevOps toolchains using familiar interfaces like Helm, Terraform, and Kubernetes CRDs. Its policy engine allows security teams to define runtime behaviors without impeding developers. Onboarding typically requires minimal training, and teams can see operational impact within days.

8. Who are your ideal customers, and what industries and use cases are showing the strongest demand for Phoenix and AMTD technology?

Ideal customers are enterprises with high-value intellectual property, national infrastructure, or regulated data environments—particularly in finance, defense, AI/ML ops, and healthcare. Use cases include protecting training datasets, containerized microservices, and operational workloads that require both agility and hardening.

9. Can you share examples or case studies demonstrating how your technology stopped sophisticated or novel attacks?

In one instance, Phoenix prevented an advanced persistent threat (APT) from maintaining command-and-control over a Kubernetes cluster by randomizing service ports and namespace bindings mid-session. In another case, a simulated ransomware payload was neutralized by breaking memory consistency during execution.



10. What feedback have you received from early adopters, and how has it shaped current product features?

Early adopters have praised Phoenix for being non-disruptive to development velocity while adding a measurable reduction in attack success rates. Their feedback directly led to the creation of a unified policy dashboard and improved support for multi-cloud service meshes.

11. How do you see adversaries evolving in response to dynamic defense solutions like yours?

As defenses become more dynamic, adversaries will likely shift toward real-time AI-driven attacks that attempt to re-map environments as quickly as they change. However, Phoenix is built for this arms race—its adaptive learning and runtime variability create a moving baseline that even AI-driven attacks struggle to anticipate.

12. What is your take on “AI washing” in cybersecurity, and how should buyers evaluate true AI-driven security products?

“AI washing” is rampant—many vendors simply wrap deterministic systems in machine learning buzzwords. Buyers should ask: Does the platform learn and adapt in production? Can it act autonomously in previously unknown situations? R6 encourages transparency through real-world validation, adversarial simulation, and demonstrable ML outcomes—not just claims.

