Rule/Query Generation from a Report with AI

Uncoder AI

Rule/Query Generation from a Report with AI

Released on April 16

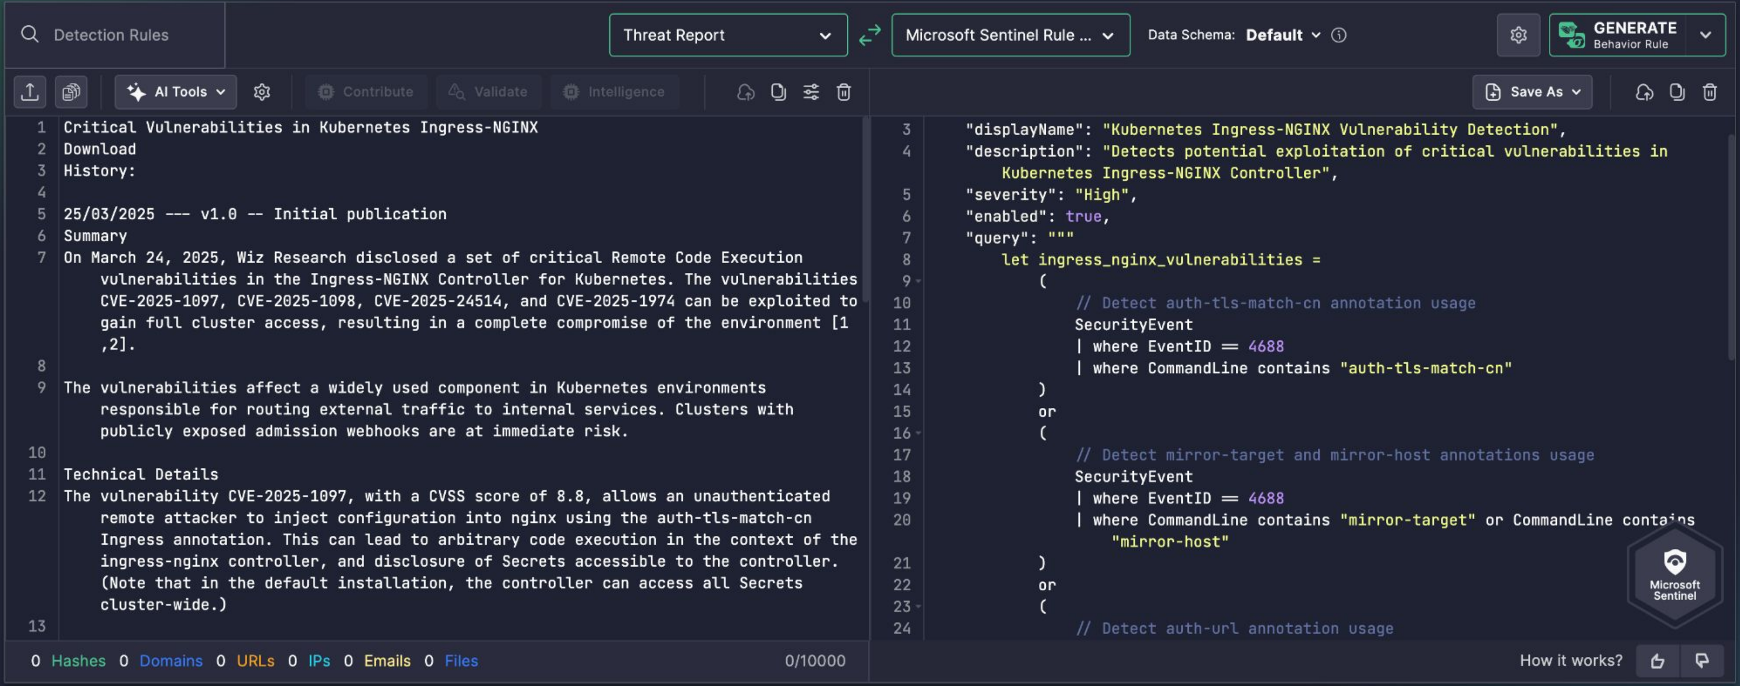

Critical Vulnerabilities in Kubernetes Ingress-NGINX
Download
History:

25/03/2025 --- v1.0 -- Initial publication
Summary
On March 24, 2025, Wiz Research disclosed a set of critical Remote Code Execution vulnerabilities in the Ingress-NGINX Controller for Kubernetes. The vulnerabilities CVE-2025-1097, CVE-2025-1098, CVE-2025-24514, and CVE-2025-1974 can be exploited to gain full cluster access, resulting in a complete compromise of the environment [1,2].

The vulnerabilities affect a widely used component in Kubernetes environments responsible for routing external traffic to internal services. Clusters with publicly exposed admission webhooks are at immediate risk.

Technical Details
The vulnerability CVE-2025-1097, with a CVSS score of 8.8, allows an unauthenticated remote attacker to inject configuration into nginx using the auth-tls-match-cn Ingress annotation. This can lead to arbitrary code execution in the context of the ingress-nginx controller, and disclosure of Secrets accessible to the controller. (Note that in the default installation, the controller can access all Secrets cluster-wide.)

```
    "displayName": "Kubernetes Ingress-NGINX Vulnerability Detection",
    "description": "Detects potential exploitation of critical vulnerabilities in
        Kubernetes Ingress-NGINX Controller",
    "severity": "High",
    "enabled": true,
    "query": """
        let ingress_nginx_vulnerabilities =
            (
                // Detect auth-tls-match-cn annotation usage
                SecurityEvent
                | where EventID == 4688
                | where CommandLine contains "auth-tls-match-cn"
            )
            or
            (
                // Detect mirror-target and mirror-host annotations usage
                SecurityEvent
                | where EventID == 4688
                | where CommandLine contains "mirror-target" or CommandLine contains
                    "mirror-host"
            )
            or
            (
                // Detect auth-url annotation usage
```

Rule/Query Generation from a Report with AI

Use case development life cycle may be long, sometimes too long. Uncoder AI analyzes the provided threat report and generates a rule/query to detect the described behavior. For this purpose, Uncoder AI uses Llama 3.3 customized for detection engineering and threat intelligence processing, hosted at SOC Prime SOC 2 Type II private cloud for maximum security, privacy, and IP protection.

- Multiple languages supported
- Data doesn't leave SOC Prime's infrastructure
- Reduces time spent translating threat intelligence into actionable detections
- Enables rapid prototyping of rules
- Allows small teams to cover broader threat landscapes

Rule/Query Generation from a Report with AI

56 languages supported

ANOMALI	APACHE kafka	ArcSight	CORTEX XSIAM BY PALO ALTO NETWORKS	CORTEX XDR BY PALO ALTO NETWORKS	AWS Athena	OpenSearch	DEVO
CROWDSTRIKE	FORTINET	REGEX	graylog	DATADOG	DNIF	ElastAlert	Elastic Stack
HUNTERS	LOGPOINT	LogRhythm	Falco	CrowdStrike Falcon LogScale	FIREEYE	Google Security Operations	hawksearch
Radar	LACEWORK	LIMA CHARLIE	LOGIQ	Microsoft Defender for Endpoint	Microsoft Sentinel	nviso	PowerShell
Qualys	GREP	ROOTA	RSA NETWITNESS	securonix	SentinelOne	snowflake	splunk>
SQL	SQLite	STIX	StreamAlert	sumo logic	Sysmon	uberAgent	Carbon Black.
onum	Logsign	SURICATA	TANIUM	SOPHOS EDR	logz.io	TREND MICRO XDR	exabeam