

Peeling Back the Layers of Security

SECURITY ONION



DEFCON908

Rish Pednekar

MARCH 19, 2026

AGENDA

[STATUS: PREPARING...]
[PRIORITY: MEDIUM]

1. What is Security Onion?
2. The Problem that Security Onion solves
3. Native AI Capabilities
4. Lab Architecture + Deployment
5. Challenges
6. Attack Traffic
7. Demo Results
8. Outcome for SOC Teams
9. Q&A



WHAT IS SECURITY ONION?

A *free and open-source* platform built by defenders for defenders.

A cohesive ecosystem designed for enterprise security monitoring, threat hunting, and log management.

2,000,000+

GLOBAL DOWNLOADS & DEPLOYMENTS

AVAILABLE ON MAJOR CLOUD MARKETPLACES



Suricata



High-performance NIDS for real-time threat detection and prevention.

Zeek



Deep protocol-level visibility and metadata extraction analysis.

Elasticsearch



Scalable search and analytics engine for indexing security data.



SOC Console



Custom web interface for alerts, dashboards, and threat hunting.

THE PROBLEM: ALERT FATIGUE

[STATUS: ANALYZING_SOC_METRICS]
[PRIORITY: CRITICAL]



83%

TIME SPENT ON REPETITIVE TRIAGE
- HELPNET SECURITY

Modern SOC's are **drowning in alerts**. Most are noise, causing burnout and missed threats.

- ⚙️ Thousands of alerts per day
- 🕒 High latency in response times
- 👤 High turnover & analyst fatigue

🤖 THE "AI INTERN"

🌙 **NEVER SLEEPS**
Triage 24/7/365 without degradation in accuracy or focus.

⚡ **INSTANT CONTEXT**
Automatically interprets cryptic rules into plain-English guidance.

🧠 **FORCE MULTIPLIER**
Handles the grunt work so humans can focus on strategic defense.

[NATIVE_AI_ENABLED]

100% UPTIME

NATIVE AI CAPABILITIES

[MODULE: SO_AI_CORE]
[VERSION: 2.4.210]

COMMUNITY TIER



AI Summaries

Plain-English explanations for 58,000+ detection rules. No more decoding cryptic Rule IDs or searching SIDs.



Investigation Playbooks

Tailored step-by-step guides for every alert. Built on the HCIP standard for consistent expert analysis.



Guided Analysis

Automated investigation flows. Clickable questions that execute live queries and return triage results instantly.

ENTERPRISE PRO



Onion AI Assistant

Conversational LLM interface for local data. Ask natural language questions like: *"Find suspicious process creation on the web server."*



MCP Server

Secure bridge using Model Context Protocol. Connect your own LLMs (e.g., Qwen, OpenAI) while maintaining data sovereignty.

PRO_LICENSE_REQUIRED

ENCRYPTION: AES-256

Enables advanced data exfiltration filtering and local token optimization.

58,000+

RULES SUMMARIZED

HCIP

STANDARDIZED WORKFLOWS

LOCAL-FIRST

DATA SOVEREIGNTY

ONION AI CAPABILITIES

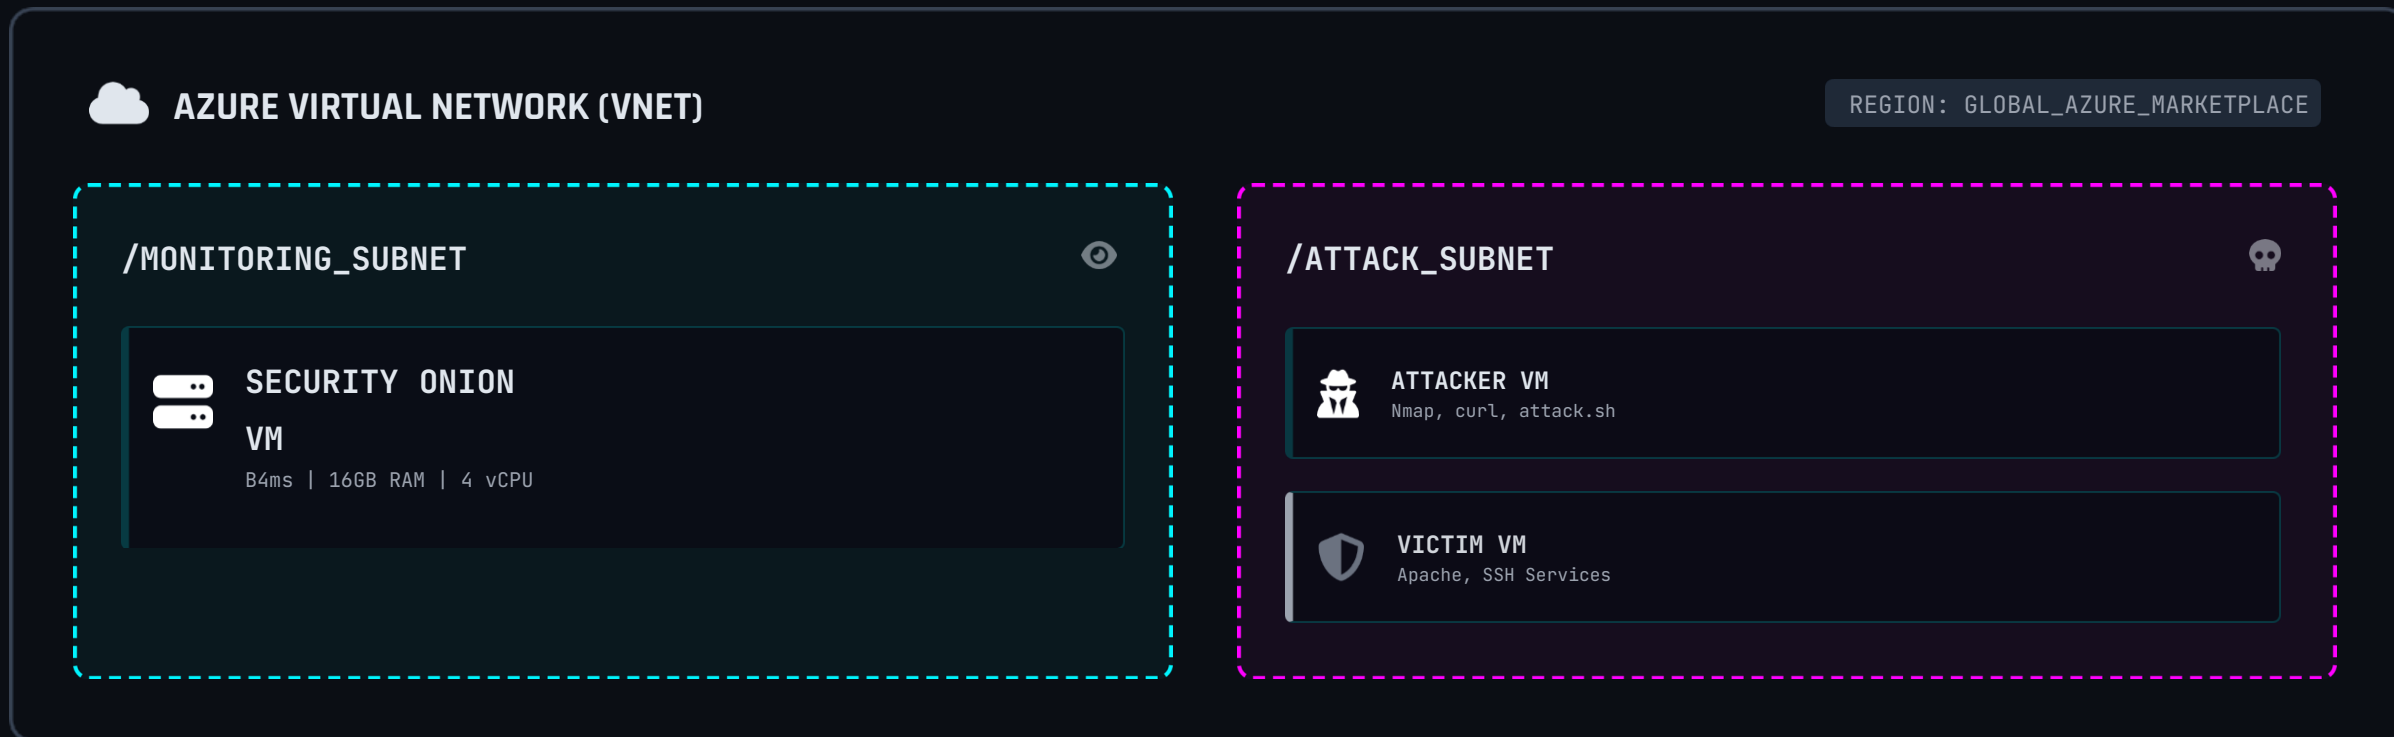
[MODULE: SO_AI_CORE]

[VERSION: 2.4.210]

Onion AI

EXTRA SPICY
Now with detection management!

LAB ARCHITECTURE: AZURE CLOUD



</> TERRAFORM AUTOMATED DEPLOYMENT

🚧 ZERO SUBNET MIRRORING (CHALLENGE)

🔪 REPRODUCIBLE ATTACK SIMULATIONS

ONE-COMMAND DEPLOYMENT

BASH — TERRAFORM APPLY

```
# Start deployment pipeline
$ terraform init
$ terraform apply -auto-approve
```

Plan: 12 to add, 0 to change, 0 to destroy.

```
azurerm_linux_virtual_machine.so: Creating...
azurerm_network_security_group.so: Creating...
azurerm_virtual_network.lab: Creating...
```

Apply complete! Resources: 12 added.

Outputs:

```
so_console_url = "https://example.ip."
attacker_ssh   = "ssh admin@example.ip"
```

01

PLAN
Terraform Init

02

PROVISION
Azure Resources

03

CONFIGURE
Cloud-Init

04

READY
SOC Live

AZURE MARKETPLACE

Utilizes official Security Onion 2.4.210 images for rapid, verified node provisioning.

CLOUD-INIT AUTOMATION

Pre-configures Attacker (Nmap, scripts) and Victim (Apache, SSH) VMs automatically upon boot. More scripts found in the GitHub!

NSG ACCESS CONTROL

Strict, predefined firewall rules ensure only necessary traffic reaches the monitoring NIC.


THE AZURE NETWORKING CHALLENGE

[VNET_SEGMENTATION: ENABLED]
[TRAFFIC_MIRRORING: NOT_NATIVE]

THE VISIBILITY GAP

Azure Virtual Networks do **not** natively mirror traffic between VMs on the same subnet.

- ✘ No built-in SPAN or TAP capability at the VNET level.
- ✘ Intra-subnet attack traffic remains invisible to sensors.
- ✘ Traditional NIC promiscuous mode is ineffective in cloud.



TACTICAL WORKAROUND



Forcing traffic to the sensor NIC to validate detection logic.

DEFCON CHALLENGE

Achieve full enterprise-grade visibility without direct targeting. Your mission:

Azure Virtual Network TAP

Configure traffic mirroring to capture ALL traffic between Attacker and Victim VMs. Solve the cloud visibility puzzle.

[GOAL: 100% VISIBILITY]

[REWARD: SKILL_UP]

NODE:
AZURE_US_WEST

NIC_STATUS:
LISTENING

ALERT:
SAME_SUBNET_BLINDSPOT_DETECTED

GENERATING MULTI-STAGE ATTACK TRAFFIC

```
./attack.sh --target 10.0.2.10
```

SCRIPT: ATTACK.SH // TARGET: MONITORING NIC

STAGE 1: RECON

Comprehensive port scanning to identify open services and OS footprint.

- ✔ SYN / FIN / XMAS Scans
- ✔ NULL / ACK Scans
- ✔ UDP Port Discovery
- ✔ OS Fingerprinting

STAGE 2: PROBING

Active exploitation attempts and protocol-specific brute force attacks.

- ✔ Telnet & FTP Auth
- ✔ SMB Vulnerability Checks
- ✔ SNMP Enumeration
- ✔ SSH Brute Force

STAGE 3: EXFIL

Advanced C2 and data exfiltration techniques using DNS protocols.

- ✔ DGA Domain Generation
- ✔ DNS Tunneling (Base64)
- ✔ DNS Zone Transfers
- ✔ Malicious C2 Callbacks

T+0M

T+15M

T+30M

SIMULATION TIMELINE: FULL ATTACK CYCLE

ENGINE: SURICATA RULES: ETOPEN 58,000+

TRAFFIC_GEN: ACTIVE

DEMO: SOC CONSOLE OVERVIEW

SITUATIONAL AWARENESS // REAL-TIME EVENT STREAM

ACTIVE DETECTIONS

12

SENSOR STATUS

ONLINE

	Count	rule.name	event.module	event.severity_label	rule.uuid
>	151	ET INFO Server Responded with Vulnerable OpenSSH Version (CVE-2024-6409)	suricata	medium	2054407
>	48	ET SCAN NMAP OS Detection Probe	suricata	medium	2018489
>	44	ET SCAN Potential SSH Scan OUTBOUND	suricata	medium	2003068
>	30	GPL ICMP PING *NIX	suricata	low	2100366
>	30	GPL ICMP PING speedera	suricata	low	2100480
>	30	Security Onion - Grid Node Login Failure (SSH)	sigma	high	923421c7-9b1e-45d4-80cc-e21d060c87
>	18	ET SCAN Suspicious inbound to PostgreSQL port 5432	suricata	medium	2010939
>	18	ET SCAN Suspicious inbound to MySQL port 3306	suricata	medium	2010937
>	16	ET SCAN Suspicious inbound to MSSQL port 1433	suricata	medium	2010935
>	16	ET SCAN Suspicious inbound to Oracle SQL port 1521	suricata	medium	2010936
>	3	ET SCAN Potential SSH Scan	suricata	medium	2001219
>	3	ET SCAN Potential VNC Scan 5800-5820	suricata	medium	2002910
>	2	GPL SNMP private access udp	suricata	medium	2101413
>	2	GPL SNMP public access udp	suricata	medium	2101411

Items per page: 500 1-14 of 14

	48	ET SCAN NMAP OS Detection Probe
ALERT DETAILS		
@timestamp	2026-03-01T04:14:31.255Z	
@version	1	
cloud.account.id	fe52b075-ca3d-4c59-8916-bb1164f92deb	
cloud.instance.id	ecc2a58f-2bb3-4fb1-a69e-e96fb47b03c3	
cloud.instance.name	so-ai-lab-securityonion	
cloud.machine.type	Standard_B4ms	
cloud.provider	azure	
cloud.region	eastus	
cloud.service.name	Virtual Machines	
data_stream.dataset	suricata	
data_stream.namespace	so	
data_stream.type	logs	
destination.ip	10.0.2.10	
destination.port	44523	
dns.query.name	CCCCCCCCCCCCCCCCCCCC	
ecs.version	8.0.0	
elastic_agent.id	6c0efd9d-d786-44fc-84fa-59f5e9c1a8f1	
elastic_agent.snapshot	false	
elastic_agent.version	8.18.8	
event.category	network	
event.dataset	suricata.alert	
event.ingested	2026-03-01T04:14:31.327Z	
event.module	suricata	
event.severity	2	
event.severity_label	medium	
input.type	log	
log.file.path	/msm/suricata/eve-2026-03-01-04.13.json	
log.id.uid	1824212977749030	
log.offset	73833	

DEMO: AI SUMMARIES IN ACTION

Overview

ET SCAN Potential SSH Scan OUTBOUND

Summary

This rule detects potential outbound SSH scanning activity by monitoring for multiple SYN packets sent from an internal network to external addresses over TCP port 22. If a source initiates five or more connection attempts within 120 seconds, it triggers an alert, indicating possible reconnaissance activity, such as a brute-force attack attempt on SSH services.

Status: Enabled

Overview

ET SCAN NMAP OS Detection Probe

Summary

This rule detects an Nmap operating system detection probe being sent to a network. The probe is characterized by specific UDP traffic with a destination port of 10000 or higher and a data size of 300. The sequence includes multiple repeated "C" characters and specific patterns within the data payload, which are indicative of an attempted reconnaissance activity using Nmap.

Status: Enabled

Overview

ET INFO Server Responded with Vulnerable OpenSSH Version (CVE-2024-6409)

Summary

This rule detects network traffic in which a server responds with a version of OpenSSH that is vulnerable to CVE-2024-6409. It identifies the presence of an OpenSSH version 8.7 or 8.8, which are known to be affected, in established connections directed towards client hosts. This activity suggests a potential exploit attempt targeting systems running these specific versions of OpenSSH.

Status: Enabled

PLAIN ENGLISH
No more decoding cryptic Suricata/Zeek rule syntax or obscure SIDs.

ZERO RESEARCH
Contextual information is provided inline, eliminating external Google searches.

MASSIVE SCALE
Native support for the entire 58,000+ ETOPEN NIDS ruleset out-of-the-box.

DEMO: AUTOMATED INVESTIGATION PLAYBOOKS

[MODULE: PLAYBOOK_ENGINE_V2]

[FRAMEWORK: HCIP_STANDARD]

ET SCAN Potential SSH Scan

OVERVIEW

OPERATIONAL NOTES

DETECT

Some playbooks were generated by AI and it's possible that they may not be 100% accurate. Please let us know if you see any issues.

```
name: ET SCAN Potential SSH Scan
id: "1261968"
type: detection
description: |
  Detects multiple SYN packets to SSH port 22 from external sources within a short timeframe.
  May indicate SSH brute force scanning or legitimate connection retries from remote systems.
created: 2024-01-15T00:00:00Z
detection_id: "2001219"
detection_category: ""
detection_type: nids
contributors:
  - SecurityOnionSolutions
questions:
  - question: What specific SSH connection attempts triggered this scanning alert?
    context: Shows the exact connection patterns that matched the scanning threshold.
    range: +/-15m
    answer_sources: []
    query: |
      aggregation: false
      logsource:
        category: network
        service: connection
      detection:
        selection:
          community_id: '{network.community_id}'
        condition: selection
      fields:
        - src_ip
        - dst_ip
        - dst_port
  - question: Does this source IP normally attempt SSH connections to our network?
    context: Determines if this scanning activity represents a deviation from normal patterns.
    range: -7d
    answer_sources: []
    query: |
      aggregation: true
      logsource:
        category: network
        service: connection
      detection:
        selection:
          dst_ip: '{destination.ip}'
        condition: selection
      fields:
        - dst_ip
  - question: What is the timing pattern of SSH connection attempts from this source?
    context: Reveals whether connections follow automated scanning patterns or manual attempts.
    range: +/-30m
    answer_sources: []
    query: |
      aggregation: false
      logsource:
        category: network
        service: connection
      detection:
        selection:
          src_ip: '{related.ip}'
          dst_ip: '{related.ip}'
        condition: selection
      fields:
        - src_ip
        - dst_ip
        - dst_port
        - network.protocol
        - event.duration
        - client_ip_bytes
        - server_ip_bytes
        - connection.state.description
  - question: Are multiple SSH servers being targeted by this source IP?
    context: Identifies the scope of scanning activity across SSH infrastructure.
```

ET MALWARE Observed Win32/Lumma Stealer Related Domain (undimik .you) in TLS SNI

OVERVIEW

OPERATIONAL NOTES

DETECTION SOURCE

TUNING (0)

PLAYBOOKS (1)

HISTORY

Some playbooks were generated by AI and it's possible that they may not be 100% accurate. Please let us know if you see any issues.

```
name: ET MALWARE Category Playbook
id: "160002"
type: Detection
description: |
  This playbook is designed to help investigate alerts from the ET MALWARE category, which detects malicious software activity on the network.
created: 2025-04-30T00:00:00Z
modified: 2025-04-30T00:00:00Z
detection_id: ""
detection_category: ET MALWARE
detection_type: nids
contributors:
  - SecurityOnionSolutions
questions:
  - question: What is the specific type of malware activity being detected?
    context: Different malware alerts may indicate different stages of infection (delivery, execution, C2, etc.). Understanding the specific activity helps prioritize response. A good starting point for this context is the detection name and description.
    answer_sources:
      - alert
    query: |
      aggregation: false
      logsource:
        category: alert
        product: suricata
      detection:
        selection:
          document_id: '{doc_id}'
        condition: selection
      fields:
        - rule.name
        - src_ip
        - src.port
        - dst_ip
        - dst.port
  - question: Is this the first time this specific malware alert has been seen for the internal hosts?
    context: Understanding if this is a new infection or ongoing activity helps determine the infection timeline and scope.
    range: -30d
    answer_sources:
      - alert
    query: |
      aggregation: true
      logsource:
        category: alert
        product: suricata
      detection:
        selection:
          rule.name: '{rule.name}'
        filters:
          - src_ip: '{network.private_ip}'
          - dst_ip: '{network.private_ip}'
        condition: selection and filter
      fields:
        - rule.name
        - src_ip
        - dst_ip
  - question: What process is associated with the network connection?
    context: Identifying the process responsible for triggering the alert can help determine the potential infection vector and extent. Pivoting off of the process GUID can help find other related events.
    range: +/-15m
    answer_sources:
      - network
    query: |
      aggregation: false
      logsource:
        category: network
      detection:
        selection:
          community_id: '{network.community_id}'
        filters:
          image.exists: true
        condition: selection and filter
      fields:
        - host.name
        - user
        - process.guid
        - event.action
        - image
```

SO_CONSOLE:~/playbooks/active

ALERT_ID:
4029184

INTEGRITY:
VERIFIED

LIVE_FEED

DEMO: GUIDED ANALYSIS TRIAGE

Alert Details AI Summary Playbook **Guided Analysis**

48 ET SCAN NMAP OS Detection Probe suricata medium 2018489

ALERT DETAILS GUIDED ANALYSIS

Guided Analysis is a new experimental feature. Some of these playbooks were generated by AI and it's possible that they may not be 100% accurate. Please let us know if you see any issues.

- > What was the complete UDP packet that triggered this NMAP detection?
- > Does this host normally receive UDP traffic on port 10000?
- > What is the pattern of UDP traffic from this source IP?
- > What other UDP ports are being probed by this source?
- > Is the scanning system a known-good system?
- > What type of scan is being detected?
- > Is this part of a broader scanning campaign?
- > What is the scanning pattern and frequency?
- > Are there any successful connections following the scan?
- > Are there any other alerts associated with the scanning system?
- > What other network connections occurred from the scanning source?
- > Are other internal hosts receiving similar NMAP probes?
- > What other external connections occurred to the target host?
- > Are there related reconnaissance alerts across the organization?
- > Did any processes respond to the UDP probe on port 10000?
- > What process is associated with this specific network connection?
- > Are there any associated DNS queries from the scanning source?

Is this part of a broader scanning campaign?

Context
Multiple scan attempts from the same source across different targets or services may indicate a systematic reconnaissance effort.

Date Range: -6h

tags:alert AND (source.ip:10.0.2.20 AND rule.name:"SCAN") | groupby rule.name* destination.ip destination.port

Instant Insight (max 5 results)
Original event occurred at 2026-02-28 23:14:31.255 -05:00

Count	rule.name	destination.ip	destination.port
4	ET SCAN NMAP OS Detection Probe	10.0.2.10	30677
4	ET SCAN NMAP OS Detection Probe	10.0.2.10	31056
4	ET SCAN NMAP OS Detection Probe	10.0.2.10	31261
4	ET SCAN NMAP OS Detection Probe	10.0.2.10	33894
4	ET SCAN NMAP OS Detection Probe	10.0.2.10	34933
35	ET SCAN Potential SSH Scan OUTBOUND	10.0.2.10	22
18	ET SCAN Suspicious inbound to PostgreSQL port 5432	10.0.2.10	5432
18	ET SCAN Suspicious inbound to MySQL port 3306	10.0.2.10	3306
16	ET SCAN Suspicious inbound to MSSQL port 1433	10.0.2.10	1433

WHAT THIS MEANS FOR SOC TEAMS

[TRANSFORMATION_ANALYSIS_REPORT]

FASTER TRIAGE

Automated interpretation of alerts removes the "What does this SID mean?" phase. Drastically reduces Mean Time to Triage (MTTT).

STATUS: OPTIMIZED

JUNIOR EMPOWERMENT

Junior analysts gain immediate access to investigation frameworks -- thus, giving them accelerated onboarding and skill leveling.

STATUS: UPSKILLED

CONSISTENT QUALITY

Every alert gets the same rigorous, structured investigation via HCIP standards. No missed critical steps during shifts.

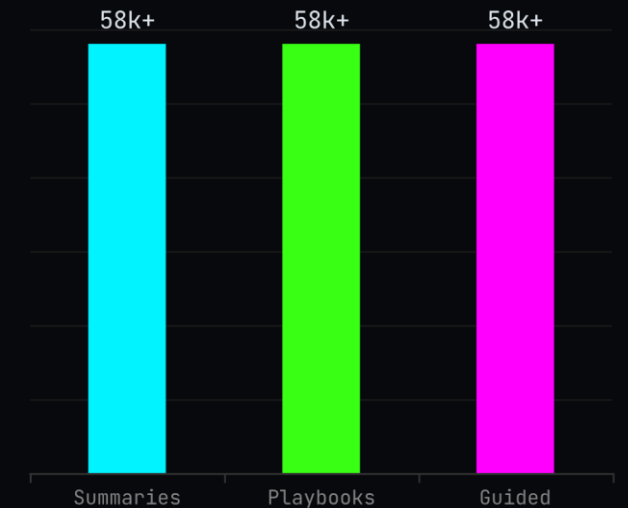
STATUS: STANDARDIZED

FREE TIER VALUE

Summaries, Playbooks, and Guided Analysis are 100% free. Enterprise-grade AI support without the budget barrier.

STATUS: ACCESSIBLE

AI-ENHANCED RULE COVERAGE



TOTAL RULE-SPECIFIC PLAYBOOKS

58,000+

[ANALYST_CAPACITY: +400%]

[ERROR_RATE: MINIMIZED]

[DATA_SOVEREIGNTY: SECURED]

SECURITY ONION AI // FORCE MULTIPLIER

WHAT THIS ALSO MEANS FOR SOC TEAMS

[TRANSFORMATION_ANALYSIS_REPORT]



Possible Over-reliance

Analysts may rely solely on the tool for their investigations.



Possible False Confidence in Investigation Skills

Analysts may overinflate their investigation skill levels.



Resource Hungry

Tool requires 16GB RAM before adding AI processing

TRY IT YOURSELF: OPEN SOURCE ACCESS



[ACCESS_GRANTED: LAB_RESOURCES]

GITHUB REPOSITORY

Access all Terraform configurations, attack scripts, and lab documentation to recreate the demo environment.

```
https://github.com/rishp66/security_onion_ai
```

QUICK DEPLOY

Run the following in your terminal to spin up the Azure lab:

```
$ terraform apply
```

ESTIMATED SETUP TIME: ~12 MINUTES

TAKE-HOME CHALLENGE

Azure VNET TAP

Configure traffic mirroring to achieve 100% visibility between your Attacker and Victim VMs.

Solve the Visibility Gap

Azure doesn't natively mirror traffic. Can you route it correctly to the Security Onion sensor?

Have fun!!

I hope you enjoy building!

QUESTIONS & DISCUSSION

LAB REPOSITORY

Access the full Azure Terraform lab and attack scripts:

```
https://github.com/rishp66/security_onion_ai
```

GET IN TOUCH

Speaker: Rish Pednekar

Email: `rishped.31@gmail.com`

LinkedIn:

`linkedin.com/in/rish-pednekar`



THE FLOOR IS YOURS

Ask about AI

implementation, Azure
visibility, or Security
Onion features.



TAKE-HOME CHALLENGE

Configure **Azure Virtual Network TAP** to solve the visibility gap!