PSTA
Public Safety Threat Alliance
Public Safety ISAO

Vulnerability Advisory

TLP:CLEAR

ESXi Vulnerabilities Actively Exploited, Disrupting Emergency Services

Disclosure Protocol: CLEAR — Disclosure is not limited
Date of Writing: 01 February 2024
PSTA Severity Level: HIGH
Risk Assessment: Exploited in the Wild, Disruption of Mission Critical Systems, Patches Available

Summary

Threat actors are actively exploiting unpatched ESXi vulnerabilities in public safety networks leading to compromises of a radio network, computer-aided dispatch system, and a large portion of a municipal network. In three instances over two days, extortion threat actors exploited unpatched flaws in out-of-date VMware ESXi servers allowing them to encrypt and disrupt critical emergency functions.

In the first instance, the *Akira* extortion syndicate attacked a United States dispatch center on 23 January 2024, forcing operators to employ backup systems and disrupting the victim's access to national crime databases. In the second attack two days later, unknown threat actors attacked a large U.S. city, exploiting an out-of-date VMware ESXi server, allowing adversaries to access and encrypt the broadband radio network.

In the third instance, in a neighboring county to the large U.S. city, threat actors were able to gain access through unpatched ESXi servers allowing them to encrypt a large portion of the municipal functions. At this time, it is believed the two organizations which are in proximity were compromised independently through similar methods.

The three attacks in close succession mark a concerning trend of at least two highly persistent cybercriminal groups gaining access to public safety networks via vulnerable ESXi servers, which has resulted in degradation of emergency functions.

Assessed Exploited Vulnerabilities

At the time of writing, it is not confirmed which VMware/ESXi flaws *Akira* and the unknown threat actor exploited when attacking public safety networks. However, some flaws are likely candidates due to their in-the-wild exploitation status, severity, or facilitation of remote code execution.

- CVE-2023-34048
  Status: *Exploited*
  VMware vCenter Server contains an out-of-bounds write vulnerability. A malicious actor with network access to vCenter Server may trigger an out-of-bounds write potentially leading to remote code execution.
  - VMware vCenter Server
    - Version 8.0
    - Version 7.0 and earlier
  - VMware Cloud Foundation (VMware vCenter Server)
    - Version 5.x
    - Version 4.x

- CVE-2023-34056
  Status: *No Exploitation Observed*
  VMware vCenter Server contains a partial information disclosure vulnerability. A malicious actor with nonadministrative privileges to vCenter Server may leverage this issue to access unauthorized data.
  - VMware vCenter Server
    - Version 8.0
    - Version 7.0 and earlier
  - VMware Cloud Foundation (VMware vCenter Server)
    - Version 5.x
    - Version 4.x

- CVE-2023-20887
  Status: *Exploited*
  VMware Aria Operations for Networks (formerly vRealize Network Insight) contains a command injection vulnerability that allows a malicious actor with network access to perform an attack resulting in remote code execution.
  - VMware Aria Operations Networks
    - Version 6.x

Akira Indicators-of-Compromise (IOCs)

The following indicators-of-compromise are associated with the *Akira* extortion syndicate's ransomware strain. These IOCs do not represent the whole of *Akira* IOCs, only the ones recently observed. These IOCs are not directly associated with the above public safety compromises, but organizations should ensure they add them to their detections.

Value	Tags	Creation Date
7b1a706bfee14a0072bcccaba5dfb336b66320eafad85904f45ec7576c9e7727	Ransomware	12-January-2024
3c92bfc71004340ebc00146ced294bc94f49f6a5e212016ac05e7d10fcb3312c	Ransomware	12-January-2024
7b295a10d54c870d59fab3a83a8b983282f6250a0be9df581334eb93d53f3488	Ransomware	12-January-2024
8631ac37f605daacf47095955837ec5abbd5e98c540ffd58bb9bf873b1685a50	Ransomware	12-January-2024
c4d103fbf2699c4bb2a8cb0f879b1993d340d9ce0af105c5e55565bed0d3aa99	Ransomware	11-December-2023
5009343ce7e6e22a777b22440480fe2eb26098d4a2ecc62e6df4498819e26b5c	Ransomware	11-December-2023
6abc0e6ef8d728a6269f8bd16881b7617dae01e032d38d583fbbbf5fcb6cac73	Ransomware	04-December-2023

Sources / Further Reading

1. https://www.cisa.gov/known-exploited-vulnerabilities-catalog
2. https://news.sophos.com/en-us/2023/12/21/akira-again-the-ransomware-that-keeps-on-taking/
3. https://www.vmware.com/security/advisories/VMSA-2023-0023.html

Recommended Mitigations

The following recommendations are suggested as part of a robust defensive strategy. These recommendations focus on mitigations for vulnerability exploitation and ransomware activity.

- ☐ **Protect Internet-facing Services**
  Ensure assets on the public internet expose no exploitable services, such as ESXi-based servers or applications. Where these services must be exposed, appropriate compensating controls are implemented to prevent common forms of abuse and exploitation. All unnecessary OS applications and network protocols are disabled on internet-facing assets.

- ☐ **Patch Known Exploited Flaws, Prioritizing ESXi**
  Ensure all known exploited vulnerabilities (listed in CISA's Known Exploited Vulnerabilities Catalog, such as the aforementioned CVE-2023-34048) in internet-facing systems are patched or otherwise mitigated within a risk-informed span of time, prioritizing ESXi flaws and more critical assets first.

- ☐ **Establish Regular Data Backups**
  Ensure all systems that are necessary for operations are regularly backed up on a regular cadence (no less than once per year). Backups are stored separately from the source systems and tested on a recurring basis, no less than once per year.

Appendix A: Assessment and Response Standard Operating Procedures

Levels of Analytic Confidence

High Confidence	Moderate Confidence	Low Confidence
Generally indicates judgments based on high-quality information, and/or the nature of the issue makes it possible to render a solid judgment. A "high confidence" judgment is not a fact or a certainty, however, and still carries a risk of being wrong.	Generally means credibly sourced and plausible information, but not of sufficient quality or corroboration to warrant a higher level of confidence.	Generally means questionable or implausible information was used, the information is too fragmented or poorly corroborated to make solid analytic inferences, or significant concerns or problems with sources existed.

Appendix B: Traffic Light Protocol for Disclosure

As part of the PSTA, agencies and other members are encouraged to share their own cybersecurity threat experiences to improve the awareness and readiness of the overall group. Submitting agencies should stipulate the level of disclosure required for their submissions according to the PSTA Traffic Light Protocol (TLP), based upon the CISA Traffic Light Protocol guidance, which helps all members submit and leverage insights while being respectful of the submitting agency's preferences.

RED: Restricted to the immediate PSTA participants only

- When should it be used? Sources may use TLP:RED when information cannot be effectively acted upon by additional parties, and could lead to impacts on a party's privacy, reputation, or operations if misused.
- How may it be shared? Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed. In the context of a meeting, for example, TLP:RED information is limited to those present at the meeting. In most circumstances, TLP:RED should be exchanged verbally or in person.

GREEN: Restricted to the community

- When should it be used? Sources may use TLP:GREEN when information is useful for the awareness of all participating organizations as well as with peers within the broader community or sector.
- How may it be shared? Recipients may share TLP:GREEN information with peers and partner organizations within their sector or community, but not via publicly accessible channels. Information in this category can be circulated widely within a particular community. TLP:GREEN information may not be released outside of the community.

AMBER: Restricted to participants' organizations

- When should it be used? Sources may use TLP:AMBER when information requires support to be effectively acted upon, yet carries risks to privacy, reputation, or operations if shared outside of the organizations involved.
- How may it be shared? Recipients may only share TLP:AMBER information with members of their own organization, and with clients or customers who need to know the information to protect themselves or prevent further harm. TLP:AMBER+STRICT Restricts sharing to the organization only.

CLEAR: Disclosure is not limited

- When should it be used? Sources may use TLP:CLEAR when information carries minimal or no foreseeable risk of misuse, in accordance with applicable rules and procedures for public release.
- How may it be shared? Subject to standard copyright rules, TLP:CLEAR information may be distributed without restriction.