Case Study

Linux Attacks in Real World Where Should We Begin?

Recently, damage caused by attacks targeting Linux servers has been on the rise. Threat actors target Linux servers because they are connected to numerous customer desktops and store vast business-critical data. As Linux-targeting malware and ransomware continue to increase, they now cause severe consequences such as business disruption. Organization must prioritize protecting their Linux servers that contain their critical assets and customer data.

AhnLab has been providing strong security across all endpoint systems, including Linux, through a security architecture that integrates our anti-malware (V3 Net for Linux), sandbox (AhnLab MDS), and EDR (AhnLab EDR) solutions. This architecture has been adopted by many customers and has contributed to their enhanced security. Our platform-centric solutions have also demonstrated outstanding technical capabilities by achieving excellent results in renowned security evaluations.

1. Statistics on Linux server attacks

Our threat intelligence unit, AhnLab SEcurity intelligence Center (ASEC), conducts statistical analysis by detecting and categorizing brute force and dictionary attacks targeting poorly managed Linux SSH (Secure Shell) servers. Poorly managed refers to systems where account credentials are vulnerable to cyber-attacks.

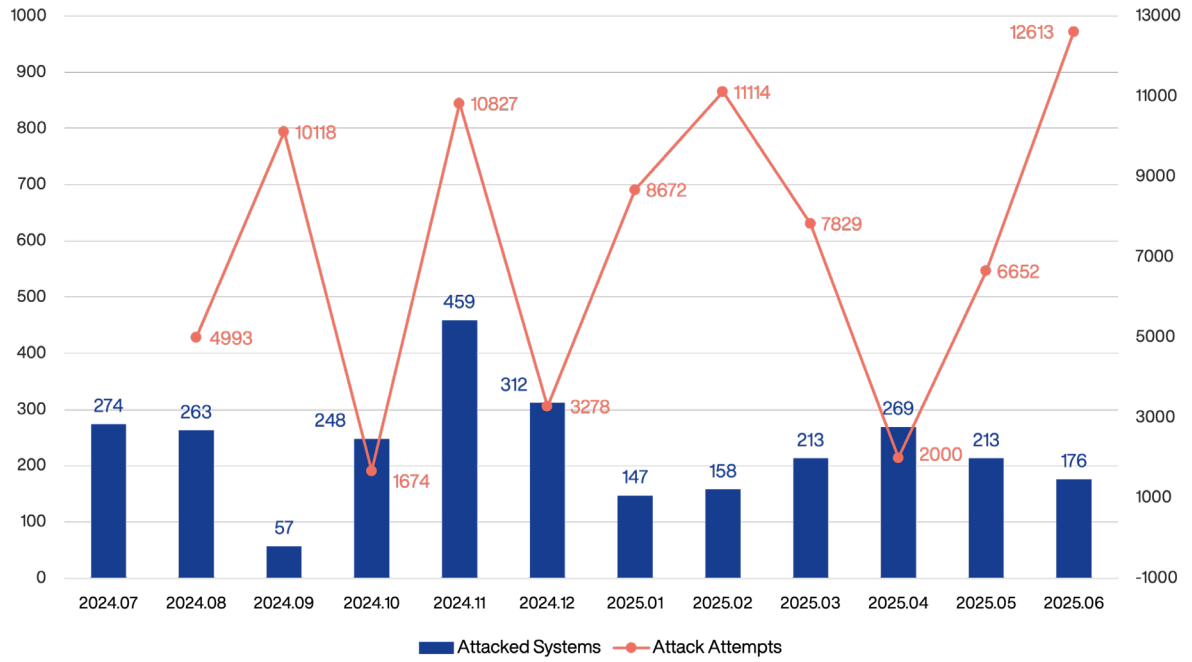

Figure 1. One-Year Linux Server Attack Statistics As of June 2025

Figure 1 shows statistics of attacks targeting Linux servers over the past year, as of June 2025.

"Attacked Systems" refers to the number of systems that threat actors exploited, with explicit evidence of malware deployment. If a threat actor successfully logs in with an admin account on a poorly managed system, they gain complete control over that system.

"Attack Attempts" refers to the frequency of threat actors attempting cyber-attacks on “Attacked Systems”. Linux SSH server attacks typically begin with scanning, followed by brute force or dictionary attacks to obtain account credentials or gather background information. "Attack Attempts" represents cases that malware deployment logs were discovered after conducting procedures described above.

The types of malware used in these attacks encompass DDoS bots, coin miners, backdoors, and ransomware. Below is a brief explanation of each malware.

DDoS Bot: Malware that allows threat actors to control infected systems and carry out DDoS attacks. It often installs additional payloads or executes additional malicious commands.

CoinMiner: Malware using a resource of the infected system to mine cryptocurrency.

Backdoor: Malware enabling threat actors to access the system and perform malicious activities.

Ransomware: Malware encrypting files on the infected system to demand ransom.

Figure 1 shows that over 12,000 attacks were performed on 176 systems in June. Over the past year, approximately 100 systems were compromised monthly, with the number of attack attempts reaching several thousand to over ten thousand in some cases. Organizations recently began paying attention to Linux server attacks because of major hacking incidents, but these attacks have been active for quite some time.

Figure 2 makes the trend even clearer. The number of Linux vulnerabilities jumped from 290 in 2023 to 3,529 in 2024, more than tenfold. By May 2025, with the year not even halfway over, more than 60,000 new Linux-targeting malware had been identified.

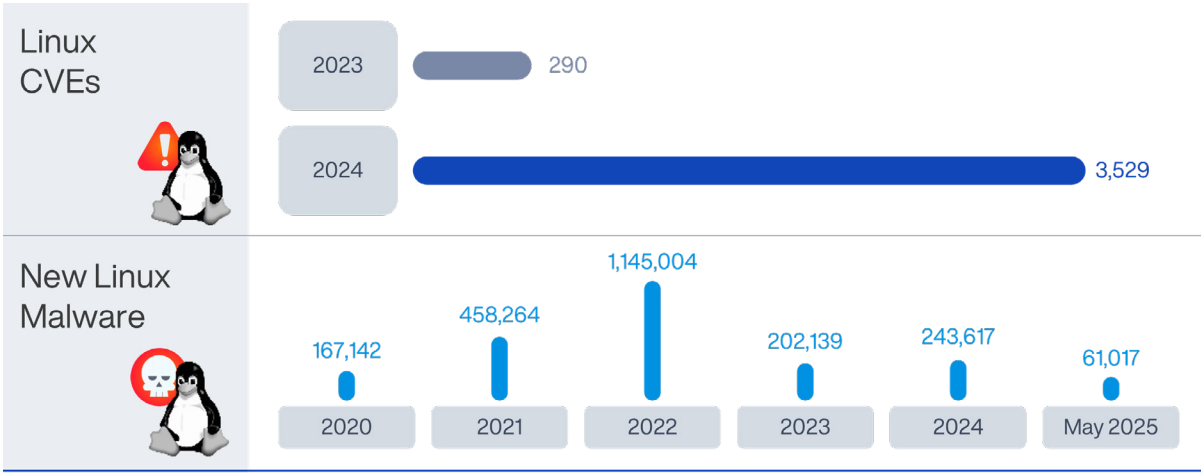

Figure 2. Number of Linux Vulnerabilities and New Malware

Hackers have been eyeing on Linux servers as they are connected to numerous customer desktops and store a vast amount of business-critical data. As Linux-targeting malware and ransomware continue to increase, they now cause severe consequences such as business disruption.

2. Linux server attack cases

Let's explore two real-world examples to understand Linux server attacks better.

Case 1: Data theft by BPFDoor (Backdoor)

Berkeley Packet Filter (BPF) is a mechanism developed for network packet filtering. It is installed in the kernel area and determines whether to allow packets from outside. BPFDoor is a Linux backdoor that exploits BPF's packet filtering feature. By adding packet filtering rules, it checks if a manipulated packet called a "magic packet" has been received and then performs malicious behaviors.

Last October, we detected and analyzed BPFDoor using AhnLab EDR, our EDR solution. Details are on the ASEC Blog.

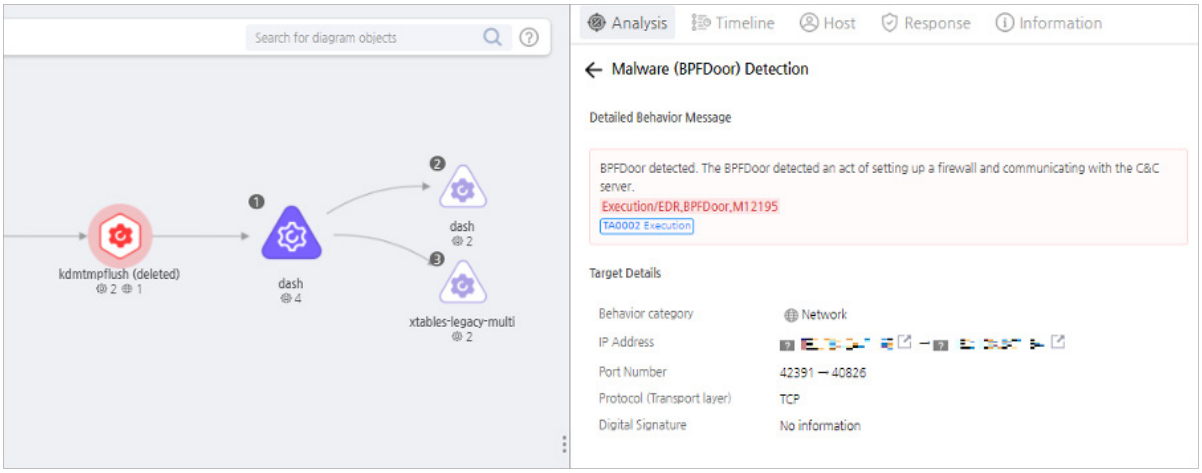

Figure 3. BPFDoor Detection of AhnLab EDR in October 2024

According to our analysis, BPFDoor copies itself to the path "/dev/shm" with the name "kdmtmpflush" and then deletes itself. The path is where a memory-based file system is located, storing and processing temporary data. Threat actors often utilize this normal system as it only runs in memory without writing to disk.

Then, it renames itself by selecting one of the strings shown in Figure 4 and disguises itself as a legitimate process. At this stage, the prctl() function is used.

```
char *self[] = {
        "/sbin/udevd -d",
        "/sbin/mingetty /dev/tty7",
        "/usr/sbin/console-kit-daemon --no-daemon",
        "hald-addon-acpi: listening on acpi kernel interface /proc/acpi/event",
        "dbus-daemon --system",
        "hald-runner",
        "pickup -l -t fifo -u",
        "avahi-daemon: chroot helper",
        "/sbin/auditd -n",
        "/usr/lib/systemd/systemd-journald"
};
```

Figure 4. Strings Used to Disguise As Legitimate Processes

BPFDoor registers the BPF filter and waits. Once the operator sends a command with the magic packet, the BPF filter passes it along, and BPFDoor executes each command required for the attack. This is how the magic packet source code works. If the password is "justforfun", it provides a reverse shell by connecting to the IP and port in the magic packet. If the password is "socket", it provides a bind shell by opening a new port and configuring the firewall to establish an internal connection for the threat actor. Finally, if the password does not match the expected values, it replies "1" to the operator. In this way, the threat actor can confirm whether malware successfully infected the target system. Once it establishes the internal connection, the hacker can perform malicious actions such as data exfiltration.

BPFDoor recently compromised a South Korean telecommunications company and leaked USIM data of its customers. The one used in telco attacks was a variant of what we analyzed last year, but they operate essentially the same way. Read our BPFDoor case study for more details.

Case 2: Unauthorized system access via proxy installation

In the second quarter of 2025, we identified cases where Linux servers were compromised with proxies being installed. In these cases, threat actors installed normal tools like TinyProxy or Sing-box. Since there were no additional attack traces, we assume the attackers' objectives were to use the infected systems as proxy nodes.

After successfully logging into the Linux server, threat actors downloaded a malicious Bash malware and installed TinyProxy. Then, they edited the configuration file (/etc/tinyproxy/tinyproxy.conf or /etc/tinyproxy.conf) by deleting access control rules that began with "Allow" and "Deny". They added the rule "Allow 0.0.0.0/0" for unrestricted access from external sources. By accessing port 8888, which TinyProxy serves on, the threat actor could abuse the infected system as a proxy.

```
200 #Allow 127.0.0.1
201 #Allow ::1
202 #Allow 192.168.0.0/16
203 #Allow 172.16.0.0/12
204 #Allow 10.0.0.0/8
205

325 #
326 #ReverseBaseURL "http://localhost:8888/"
327
328
329
330
331 # Added by script - WARNING: Allows all connections!
332 Allow 0.0.0.0/0
```

Figure 5. The Commented-Out and Inserted TinyProxy Configuration

The following case involves the installation of a Sing-box on a compromised system. Sing-box is a multi-purpose proxy tool that supports protocols such as vmess-argo, vless-reality, Hysteria2, and TUICv5. According to its GitHub description, the tool can be used to bypass protections on services like ChatGPT and Netflix. In this case, the threat actor gained unauthorized access to a Linux system and installed a Sing-box, likely with the intent of conducting further activities or generating financial profit.

Recent Linux server attacks show a trend of abusing normal tools, such as TinyProxy and Sing-box, instead of using an obvious malware. By turning an infected system into a proxy, threat actors can hide their identity while they perform subsequent attacks. Also, they may profit illegally by selling access to these proxy nodes.

3. What do we need for Linux server security?

Organizations cannot counter these latest cyber threats with a fragmented and single-product approach. Modern cyber-attacks traverse various security domains across endpoints and email, using new malware and its variants. Also, these attacks are not one-off events as they can recur anytime. Therefore, organizations must have the following capabilities for successful security.

- A security framework that enables threat detection, analysis, and response
- A seamless integration of security products to protect multiple security domains
- A threat hunting strategy that goes beyond detection and prevention

It would be ideal to operate many solutions properly, but many organizations often lack the resources to do so. For those looking for optimal solutions, we recommend embracing antivirus, sandbox, EDR, and MDR services.

#1. Antivirus: Antivirus (AV) detects and blocks malware in advance with its signatures and behavior analysis. It is the most fundamental yet essential solution for endpoint security. AV products are available for desktops and servers, and are tailored to environments across Windows, macOS, and Linux. Organizations should look for an AV optimized for the environment for Linux server security.

#2 Sandbox: A sandbox solution collects files from various security domains and performs behavioral analysis in a virtual machine (VM). For example, it can execute and manipulate document files in VMs to detect hidden malicious behaviors or unknown threats. It can also detect malware and ransomware that bypasses antivirus solutions. Sandbox is best viewed as complementary to antivirus, quickly detecting known malware based on signatures.

#3. EDR: EDR (Endpoint Detection & Response) monitors all behaviors and events at endpoints and collects data for incident investigation. EDR proactively tracks and analyzes cyber threats based on the behavioral data, helping organizations design a long-term threat response strategy. Simply put, it functions like a CCTV that monitors everything.

#4. MDR service: MDR (Managed Detection & Response) is a service in which experts deliver threat detection, analysis, and response using EDR. The service helps them reduce the burden of security operations while strengthening their detection and response capabilities. MDR services have seen an increasing demand as cyber threats grow more sophisticated, and EDR-based detection and analysis require a high level of expertise.

By unifying the operation of antivirus, sandbox, and EDR with MDR services, organizations can build a robust threat response framework against ever-evolving cyber threats – significantly enhancing the security posture of Linux servers.

4. Our optimal Linux security architecture

For resource-stretched customers, we offer optimal security solutions across endpoint and email so customers can build a robust cross-domain threat response system. From isolating malicious emails and analyzing files via network traffic mirroring to delivering a full-package endpoint security features powered by a managed service, we offer near-complete security against Linux-targeting attacks.

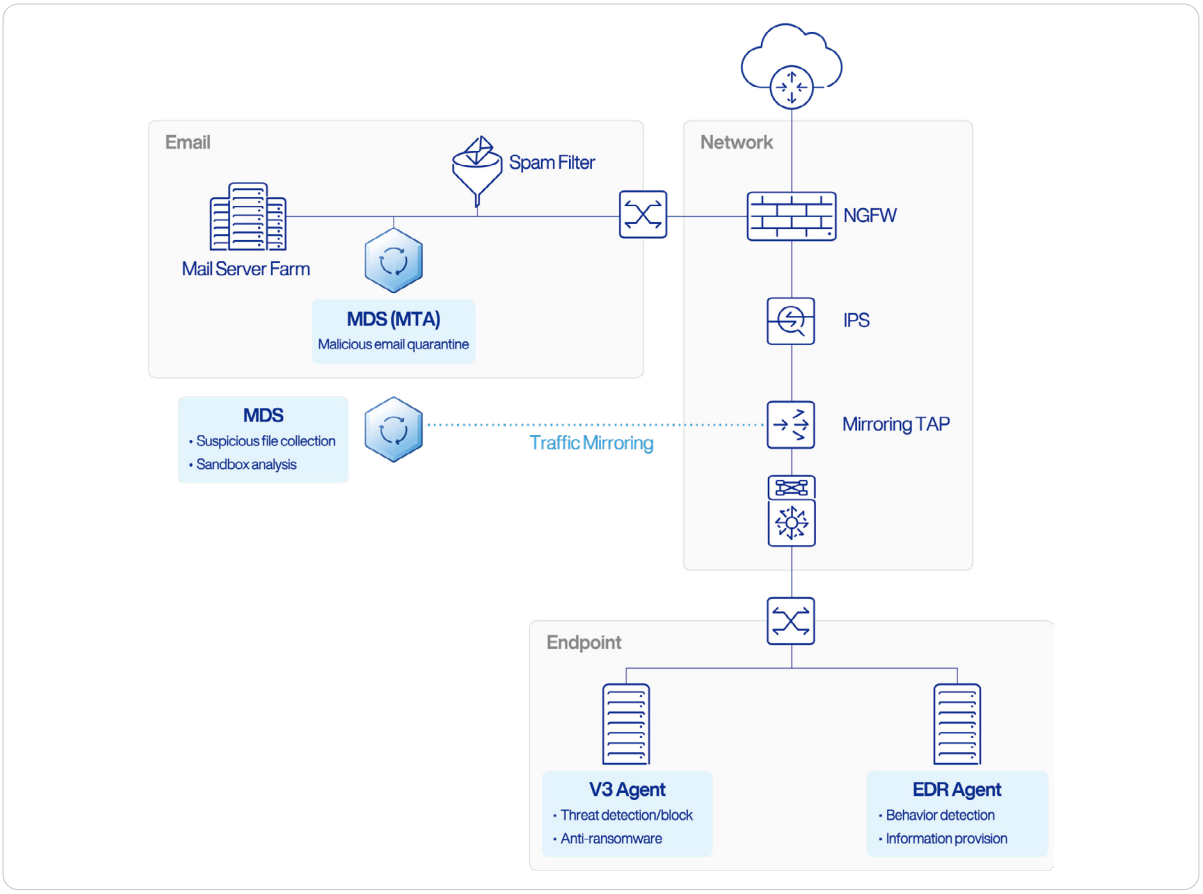

Figure 6. Our Linux Server Security Architecture

Above is our Linux server security architecture that delivers powerful protection against major cyberattacks, including ransomware. Solutions within the architecture seamlessly interact with each other to deliver maximum security capabilities; they are powered by comprehensive detection technologies, including static, reputation, dynamic, and behavior detection and analysis.

The role of each solution is summarized below.

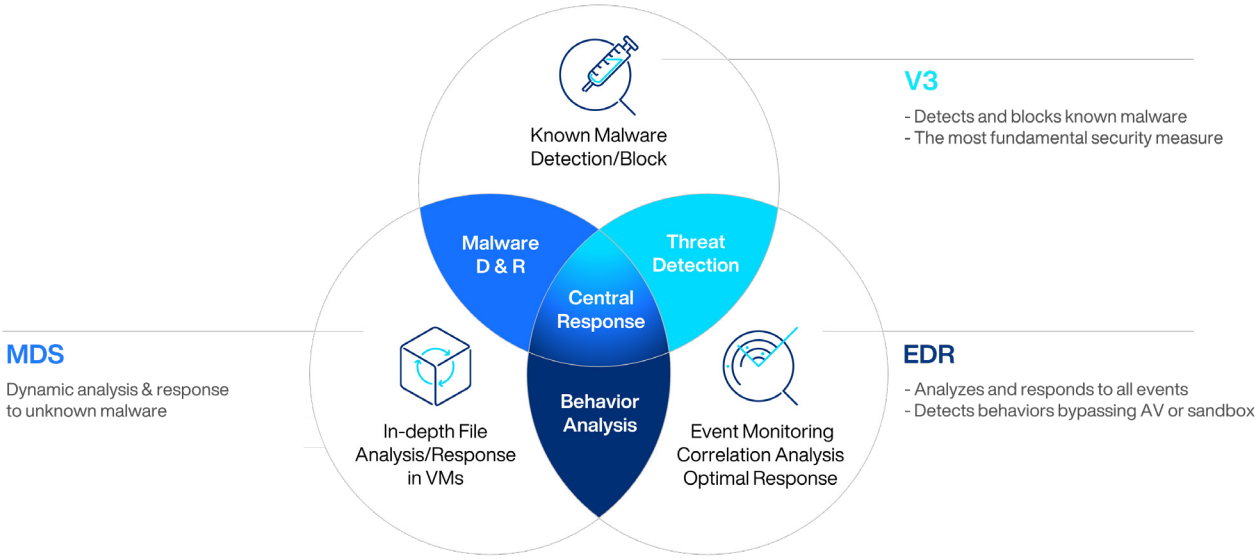

Figure 7. A Role of Each Solution

#1. AhnLab V3: Detection & blocking of known malware

AhnLab V3, our anti-malware, is backed by thirty years of history and technical expertise. It accurately detects and blocks known malware using its proprietary database containing billions of malware signatures. By leveraging thousands of malicious behavior patterns, it can properly prevent known and unknown malware.

It also delivers ransomware-specific features such as file decoy, application isolation & scan, and a ransomware security folder that quarantines potential ransomware. It also detects and blocks fileless malware by employing process memory detection AMSI (Anti-Malware Scan Interface) technologies.

AhnLab V3 provides solutions for different endpoint devices and operating systems. For Linux servers, V3 Net for Linux provides optimal detection and prevention of Linux-based malware, such as BPFDoor. It also supports cloud security features such as Docker container scan, helping customers achieve true hybrid cloud security.

#2. AhnLab MDS: Blocking malicious emails and analyzing unknown malware

From a Linux security perspective, AhnLab MDS performs sandbox-based file inspection across various domains, including email gateway and inter-network environments.

In the email domain, AhnLab MDS MTA (Mail Transfer Agent) performs a comprehensive scan of the email header, body, URLs, and attachments. It directly accesses URLs in the body to detect cyber threats and analyzes attachments in a sandbox environment. Malicious emails are quarantined to prevent them from entering the system.

The sandbox collects suspicious files via network traffic mirroring and executes them in VMs for in-depth analysis. Linux files are often executed with specific parameters, and AhnLab MDS allows users to directly input files and parameters to check the file's behavior. It can also simultaneously analyze multiple files in a VM to accelerate malware detection and analysis.

#3. AhnLab EDR: Detection and response to all endpoint events

AhnLab EDR monitors all events across various devices, such as servers and desktops, to detect and respond to endpoint threats. Its core features encompass endpoint behavioral data collection, event correlation analysis, and unified response via AhnLab V3 and MDS integration. The solution is designed to manage threats across all endpoints, minimize the dwell time of unknown threats, and prevent potential damage and recurrence.

AhnLab EDR performs detailed analysis based on the MITRE ATT&CK framework and uncovers various factors, including inflow paths, major behaviors, correlations, severities, and others. Then, it represents the analyzed data in diagrams, timelines, and process trees, allowing customers to easily understand the underlying context.

Its dedicated console, called "EDR Analyzer," optimized for advanced security management, enables users to accurately identify and respond to cyber threats while configuring policies optimized to the organization.

#4. MDR Service: Maximizing the effectiveness of EDR

We deliver the MDR service with our EDR by default to support more effortless operation and an enhanced security experience. Customers of AhnLab EDR can take advantage of services like real-time monitoring, groundbreaking analysis, risk prioritization, optimal response guide, and customized analysis reports delivered by our industry-leading security experts. Customers seeking a higher-level service can use the "EDR Premium" option with an advanced MDR Service. It contains extra services, such as more extensive log analysis and the creation of custom detection rules tailored to the customer.

5. Why Our Offering?

There are numerous solutions available on the market for anti-malware, sandbox, and EDR. Here are three reasons why customers find our solutions appealing.

#1. Seamless integration and platform-centric approach

Since launching AhnLab V3 over thirty years ago, we have developed every solution that constitutes the security architecture. Our native development experience has become a source of seamless integration and more robust cybersecurity capabilities. Customers can leverage these solutions in a unified manner to achieve even stronger security outcomes.

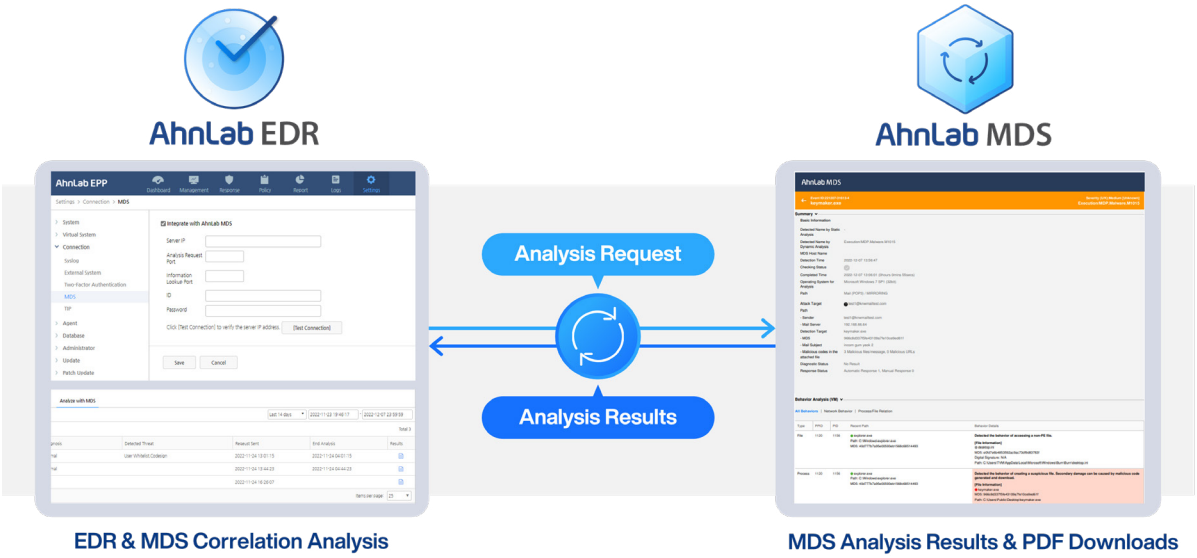

Figure 8. Correlated Analysis Between AhnLab EDR and MDS

Our approach results in significantly greater capabilities in both analysis and response. For example, if AhnLab V3 identifies a file as malware, EDR can provide detailed analysis to enable threat hunting and proactive prevention of similar future threats. Additionally, when analyzing various processes and files collected by EDR, customers can request deeper analysis to AhnLab MDS and retrieve results to achieve a more precise and comprehensive understanding of cyber threats.

#2. Proven excellence backed by customer references

Our security architecture has demonstrated its effectiveness over many years through widespread adoption by customers. Many customers have embraced two or all three solutions to achieve integrated security benefits. Each solution plays a distinct but complementary role in raising the level of security together as a platform.

In addition, these solutions extend beyond Linux to provide integrated, cross-platform security for Windows, helping customers build secure business environments.

#3. Outstanding results in various security evaluations

Our solutions have consistently demonstrated technical excellence by achieving exceptional results in globally accredited security evaluations.

In particular, AhnLab EDR has been evaluated in MITRE ATT&CK Evaluation, one of the most authoritative global security product tests, for four consecutive years. The evaluation rigorously assesses threat detection and response capabilities of security products by emulating tactics and techniques of major threat groups.

In the recent round 6, AhnLab EDR achieved 95% visibility by detecting 56 out of 59 substeps in Ransomware scenarios emulating CL0P and LockBit campaigns across Windows and Linux platforms. AhnLab EDR provided high-quality evidence with comprehensive and contextual analysis into sophisticated threat behaviors.

In addition, AhnLab EDR delivered 49 "Techniques" among 56 substeps detected. This demonstrates that our customers can thoroughly understand the underlying context (how and why) of malicious behaviors and make informed decisions by referring to the evidence of our solutions. The result is even more meaningful as context-aware detection is key to triggering an optimal response, especially when dealing with sophisticated and ever-evolving modern cyber threats.

For more details of our round 6 results, please read the result analysis report.

6. Conclusion

We utilize our time-renowned technology and knowledge built up over the past thirty years to provide strong and integrated security across different security domains, including Linux servers. The greatness of our solutions has been proven through numerous customer case studies, while their technological excellence has been validated through various security evaluations.

If you are searching for solutions to conquer ever-evolving Linux-targeting attacks, we assure you that our architecture and products are among the best options that you can possibly explore.

Visit our website to learn more about our products for Linux server security.

- V3 Net for Linux Server
- AhnLab MDS
- AhnLab EDR
- MDR service

AhnLab

AhnLab, Inc.
220, Pangyoyeok-ro, Bundang-gu, Seongnam-si, Gyeonggi-do, 13493, South Korea
www.ahnlab.com/en | global.sales@ahnlab.com