DATA SHEET

Zscaler Cloud Sandbox

Stop unknown attacks in seconds with the industry's first AI-powered inline sandbox

Zscaler Cloud Sandbox delivers unlimited, latency-free inspection to block zero day threats before they reach users and endpoints.

Legacy security solutions struggle to keep pace with today's evolving threat landscape. Traditional defenses such as firewalls and endpoint solutions often let malicious files slip through, fail to detect threats hidden within encrypted traffic, and can't scale effectively. With 87.2% of internet traffic now encrypted[1], threats delivered over secure channels have become more prevalent. Alarmingly 86.5% of these attacks involve ransomware or other malware, leaving organizations increasingly exposed.

Adversaries are leveraging artificial intelligence (AI) to launch faster, more sophisticated attacks, making it much easier to execute complex cyberattacks with a lower skill ceiling. In this environment, where AI-powered attacks and encrypted threats are the norm, protecting your organization demands modern solutions that can efficiently detect and neutralize threats before they slip through the cracks.

Business Benefits:

- **Prevent Zero Day Threats in Seconds:** Stop unknown file-based threats with inline malware and advanced threat detection, including AI-driven instant verdicts.
- **Bolster Security and Preserve Productivity:** Maximize security while keeping users productive by automatically detecting and quarantining threats—integrating Zero Trust Browser Isolation with sandbox capabilities.
- **Optimize SOC Workflows:** Seamlessly integrate malware protection into your SOC workflows with out-of-band file analysis, third-party threat detection tools, and malware analysis using both unpatched and fully patched VMs for efficient threat investigation.
- **Deploy Easily and Scale Globally:** Reduce costs and eliminate the hassle of outdated hardware and software. Simple policy configurations deliver immediate value, driving strong ROI and enabling strategic growth.
- **Maintain Data Residency and Compliance with Global Sandbox Nodes:** Capitalize on our global edge presence to ensure data residency and compliance. Zscaler Cloud Sandbox is fully integrated with Zscaler Internet Access™ and a part of the Zero Trust Exchange™, allowing fully integrated protection and an unmatched user experience.

[1] ThreatLabz 2024 Encrypted Attacks Report

Zscaler Cloud Sandbox

Built on a cloud native, proxy-based architecture, Zscaler Cloud Sandbox is the world's first AI-driven inline malware prevention engine that blocks unknown, file-based threats before they reach endpoints. With real-time analysis, instant AI-driven verdicts, and continuous threat intelligence sharing, our sandbox helps organizations stay ahead of sophisticated attacks without compromising productivity.

To cater to varying customer needs, the Zscaler Cloud Sandbox offers two options:

- **Cloud Sandbox** included with Zscaler Internet Access (ZIA), provides essential threat prevention capabilities for baseline protection.
- **Advanced Cloud Sandbox** delivers enhanced features such as AI Instant Verdict, granular policy control, API-driven analysis, and in-depth threat investigation for customers requiring more robust defenses.

Unlike traditional sandbox solutions that depend on network firewalls or endpoint-heavy agents, Zscaler Cloud Sandbox operates directly from the cloud. It eliminates the burden on endpoints, requires no reliance on network traffic flows, and seamlessly integrates into SOC workflows for faster, more effective threat detection and response.

The sandbox operates in two key modes to ensure flexible and comprehensive security:

- **Inline Inspection:** Files are analyzed as traffic flows through the cloud, and AI delivers real-time verdicts to block threats without latency.
- **Out-of-Band Inspection:** Files can be submitted via APIs for a quick scan or in-depth analysis as needed.

By offering foundational and advanced options, Zscaler ensures that customers of all sizes can benefit from layered security tailored to their needs.

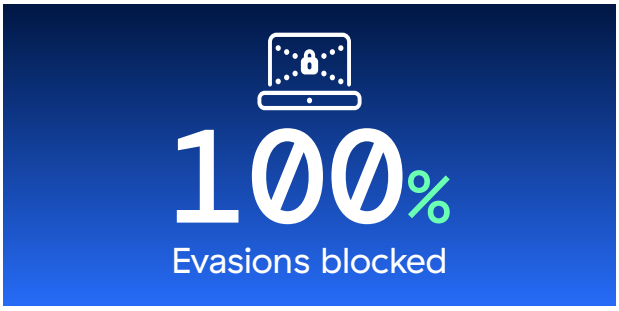

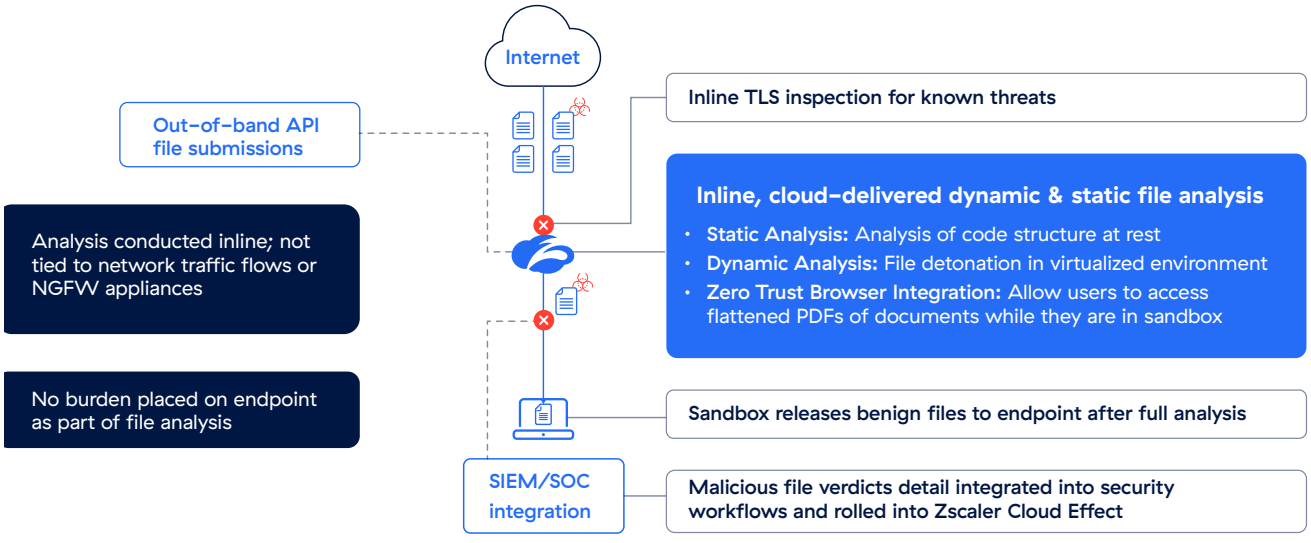

How It Works: Architecture and Key Processes

The unknown or suspicious file is first sent through a prefiltering analysis engine that checks the file contents against 40+ threat feeds, antivirus signatures, YARA rules, and AI/ ML models to render a quick verdict, blocking similarly known threats. After the initial triage, the file then undergoes robust static, dynamic, and secondary analysis that includes file execution in a controlled, isolated environment to reach an actionable verdict. The last step is post- processing, which updates the Zscaler threat database and customer policy enforcement.

With AI-based verdicts, benign files are delivered instantly while malicious files are blocked for all Zscaler global users because of the shared protection from the cloud effect. This stops patient-zero infections and emerging threats for all users, regardless of device or location.

Key Capabilities

AI-Powered Security

Harness the power of AI/ML models trained on over 600 million samples to deliver instant, high-confidence verdicts on unknown files. The AI Instant Verdict engine predicts whether a file is malicious or benign in real time, allowing immediate blocking or rapid delivery to users without delay. For cases where files are flagged as "suspicious," our customizable Minimum Threat Score enables users to fine-tune policy behavior, balancing security, and productivity. All confirmed threats are further validated through the Advanced Cloud Sandbox, which generates detailed reports and comprehensively analyzes the threat.

Additionally, PDF files are fully supported, expanding on existing file types to broaden coverage against evolving threats. Complementing this is AI-enabled PhishCatch, part of the broader Zscaler Internet Access, which strengthens protection against AI-powered phishing attacks at the network level, providing holistic, multilayered defense.

Layered Malware Detection

Zscaler Cloud Sandbox delivers multilayered threat detection powered by real-time updates from over 500 trillion daily threat signals and the Zscaler Cloud Effect, which instantly blocks known threats for all customers once identified. The process begins with static analysis to inspect the file's code structure while at rest, followed by

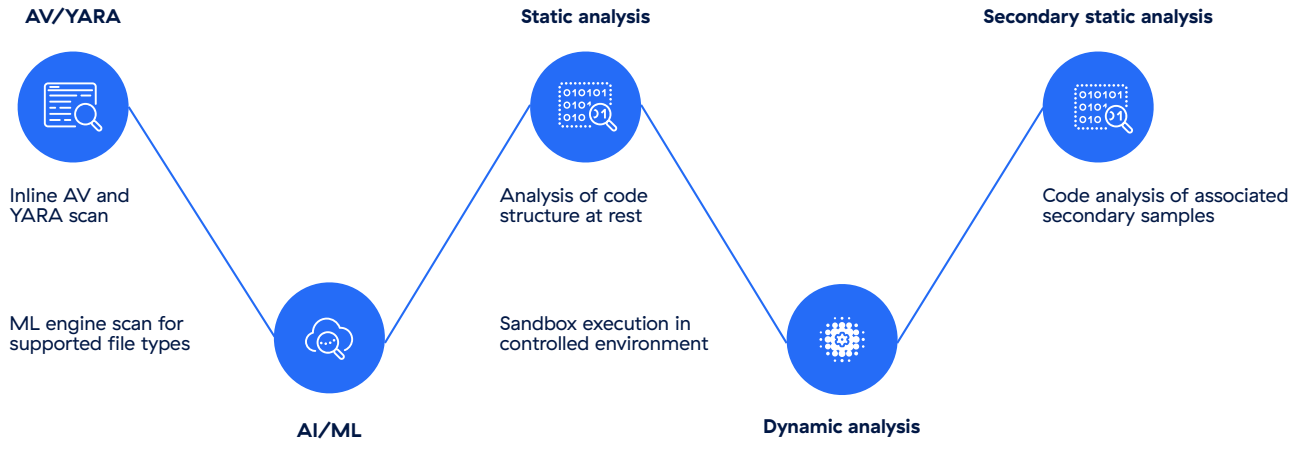

dynamic analysis where the file is detonated in a virtualized environment to observe its behavior. A secondary static analysis examines any hidden or secondary payloads.

This layered approach, combined with inline protection and unlimited TLS/SSL inspection using our Single Scan, Multi-Action engine, ensures low-latency, high-performance threat prevention for even the most sophisticated, evolving threats.

Zero Trust Browser Integration

The Zero Trust Browser integrates seamlessly with the Advanced Cloud Sandbox to ensure users remain productive without compromising security. This feature allows users to securely interact with original files in a cloud-based, isolated environment while the files undergo sandbox analysis. Risky or unknown files are automatically quarantined, but users can still access them securely during the analysis. If the sandbox analysis determines the file is safe, users can download the original version without delays.

If the file is flagged as malicious, users have options: they can download a flattened PDF version with all active content removed (CDR-1) or, when appropriate, download an editable version with only the harmful content removed (CDR-3, may require additional licensing). This flexible integration ensures that users remain productive while zero day threats are thoroughly examined and blocked when necessary. It is important to note that Zero Trust Browser is an optional module and works in tandem with the Advanced Cloud Sandbox to enhance the overall user experience.

API-Driven Analysis

Zscaler Cloud Sandbox's API integration streamlines investigations by allowing files to be sent directly to the sandbox for analysis. Security teams can easily access sandbox results via APIs, ingest them into SIEMs, SOAR platforms, or share with EDR solutions, enabling seamless SOC workflows. After detonation, the sandbox generates comprehensive reports, highlighting polymorphism, callback behaviors, and attack lifecycles, mapped to the MITRE ATT&CK framework. These insights can be operationalized to enhance threat detection, response, and proactive threat hunting.

For flexibility, organizations can use the **Advanced Sandbox Submission API** for thorough analysis of files out-of-band supporting on demand submission and reporting of unknown threats, or the **Deep Inspection Scan API** to receive an instant verdict on a file sample. Both APIs help reduce investigation times and enhance collaboration across the security stack.

Granular Policy Control and Reporting

With Zscaler Cloud Sandbox, administrators can configure and enforce granular security policies tailored to specific user roles, locations, or application categories—all within just a few clicks. The intuitive interface allows for fast, simple policy creation and deployment, minimizing time and resource investment. Policies, including rule orders for precise execution, are easily managed and automatically follow users or groups, even when they change locations.

For deeper control, admins can configure custom hash blocklists, YARA rules, and automated JA3 fingerprinting to enhance threat detection. Contextual and pre-configured reporting, including MITRE ATT&CK mapping, helps meet compliance and audit requirements. As a cloud-delivered solution, there is no hardware or software to manage, ensuring organizations can reduce operational complexity while gaining powerful, flexible security enforcement.

In just 20 minutes, Zscaler Advanced Cloud Sandbox with AI Instant Verdict enhanced file security and ensured productivity for a global enterprise:

- **91%** of files were instantly delivered as safe, keeping workflows uninterrupted.
- **9%** of files required advanced dynamic analysis to determine their nature.
- **5%** of files were flagged as malicious and blocked, preventing potential breaches, and ensuring enterprise-wide protection.

Use Cases

Defend Against Ransomware and Other Malware

Stop file-based ransomware, spyware, and other advanced malware by leveraging multilayered threat detection techniques, including dynamic file analysis, hash blocklists, and heuristic scanning. By identifying and blocking malicious files before they can spread, Zscaler Cloud Sandbox helps organizations maintain a stronger proactive security posture.

Empower Your Security Operations Center (SOC)

Equip your SOC with actionable threat intelligence, malware insights, and in-depth analysis reports mapped to the MITRE ATT&CK framework. This enables faster detection of threats, accelerates investigation, and enhances response times, all supported by the massive scale and intelligence of the world's largest security cloud.

Enable Out-of-Band API Analysis for Malware Inspection

Detect hidden threats within files through out-of-band API-driven submissions, enabling faster malware inspection without disrupting operations. Seamlessly inspect files from third-party sources, cloud repositories, or during business-critical events like mergers and acquisitions, ensuring comprehensive threat coverage and minimal risk to core systems.

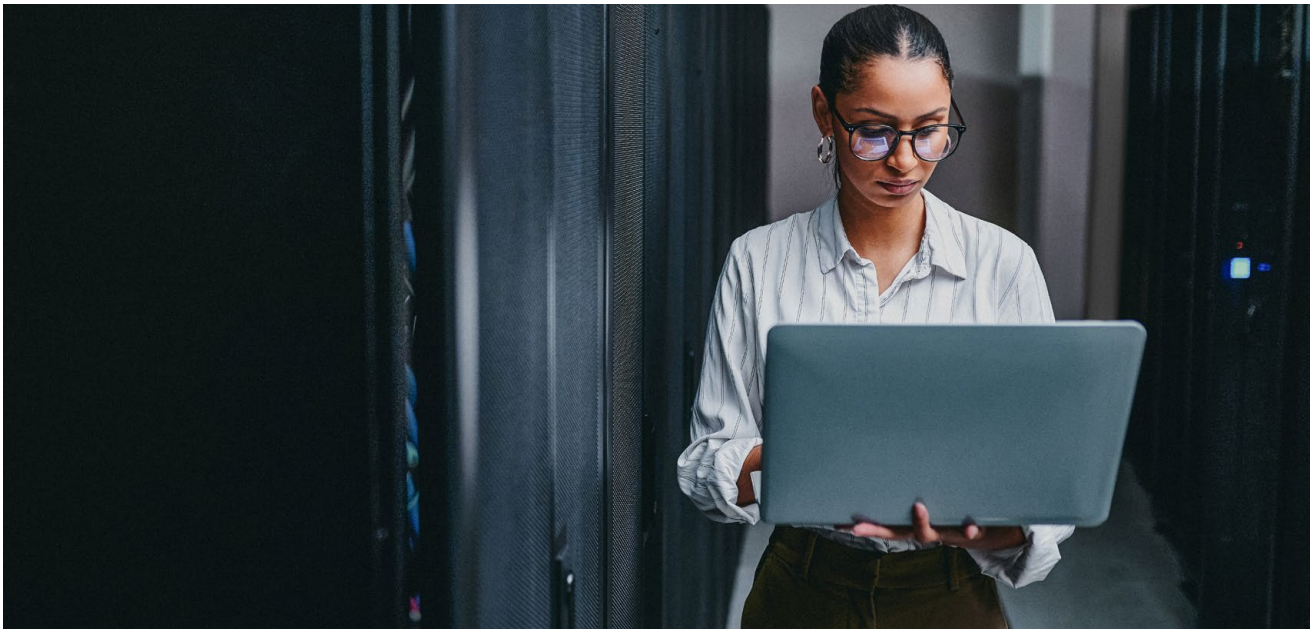

Cloud Sandbox (Standard) vs Advanced Cloud Sandbox

Feature	Value	Standard Cloud Sandbox	Advanced Cloud Sandbox
File type support	Broad file type coverage—expand protection against evolving threats	Windows Executables Windows Libraries Zip archives	Adobe Flash Android Package Kit Archives HTML Application Optical Disc Image Java Applet Microsoft Installer Microsoft Office Portable Document Format Python 2 & 3 Files Visual Basic Script Windows Batch Files Windows Executables Windows Libraries Windows PowerShell Zip archives with images Zip archives with scripts
AI Instant Verdict (AI/ML Quarantine and Blocking)	Allows benign or blocks malicious files in seconds with high confidence	Not available	Available for select file types
Granular policies	Tailor security policies using multiple criteria–– users, locations, URL categories, and more	Not available	Granular policies fitting many use cases can be created with ease
Reporting	Detailed reports with MITRE ATT&CK matrix mappings, TTPs, malware behavior and intent, IOCs, PCAPs, and more for actionable insights	Not available	Multiple reports and dashboards are available, provides rich insights
API	Out-of-band API for on demand Sandbox file analysis	Not available	Advanced Sandbox Submission API* (includes dynamic analysis) includes 3000 files per month by default Deep Inspection Scan API* (not including dynamic analysis) includes 5 files per day *Additional licensing may be required based on API usage requirements
Zero Trust Browser Integration	Automatically quarantine risky files allowing users to safely access the file during sandbox analysis	Not available	Available for PDF and Microsoft Office files

Capabilities	Description
Prefiltering analysis engine	AV, file hash blocklists, YARA rules, automated JA3 fingerprinting detections, and ML/AI model
Static, dynamic, and secondary static analysis	All files go through 3 stages of sandbox analysis: static analysis, dynamic analysis, and secondary static analysis of all files dropped during dynamic analysis
File support	7z, .apk, .applet, .bat, .bz, .bz2, .class, .cmd, .dll, .dll64, .doc, .docm, .docx, .dotm, .dotx, .exe, .exe64, .gtar, .hta, .iso, .jar, .msi, .ocx, .p, .pdf, .pickle, .pkl, .potm, .potx, .ppam, .ppsm, .ppsx, .ppt, .pptm, .pptx, .ps1, .py, .pyd, .pyw, .rar, .rtf, .scr, .slk, .swf, .sys, .tar, .tgz, .vbs, .wsc, .wsf, .wsh, .xla, .xlam, .xls, .xlsb, .xlsm, .xlsx, .xltm, .zip, .zip with images (.jpeg & .png), &.zip with scripts (.cmd, .hta .js, .lnk, .ps1, .svg, .vbs)
SSL inspection	Unlimited capacity
File retention	Zscaler Cloud Sandbox operates solely in memory. Files are stripped of identifiable information during analysis. After analysis is completed, benign files are purged from memory while malicious files are encrypted and stored indefinitely, sharing insights across all Zscaler's users for continuous protection (Cloud Effect).
Protocol support	HTTP, HTTPS, FTP, FTP over HTTP
Files per day	Unlimited capacity (inline)
Maximum file size supported; no reconfiguration needed	APK: 50 MB Adobe Flash: 2 MB Archives, .zip with images, & .zip with scripts: 50 MB HTML Application: 5 MB Java Applet: 5 MB Microsoft Office: 20 MB Portable Document Format: 20 MB Scripts: 5 MB Windows dlls, exes, libraries, & MSIs: 50 MB Windows Scripts: 5 MB
Deployment method	Cloud native, no on-premises hardware or NGFW hooks needed
Threat intel integration	40+ security partner threat intel feeds
Management and reporting	Full reporting including malware behavior and intent, indicators of compromise (IOCs), dropped files, PCAPs

Forensics	Initial sample, secondary payloads, PCAPs
API support	Robust API support, two sample submission APIs and report retrieval via API in JSON format in summary or detailed formats
Granular policies	Easy to use and configure policies for users, location, location groups, file types, user groups, departments, URL categories, and protocols
Privacy and compliance certifications	Compliant with rigorous global Commercial and Government risk, privacy, and compliance
Industry and data privacy regulations	Compliance adherence to industry-specific and in-country data privacy regulations. Privacy **data sheet**
3rd Party Independent Validation	AAA rating by CyberRatings

zscaler™ | Experience your world, secured.™

About Zscaler
Zscaler (NASDAQ: ZS) accelerates digital transformation so that customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SASE-based Zero Trust Exchange is the world's largest inline cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.