

Sandbox Service



Contents

HOW THE SANDBOX SERVICE WORKS: AN OVERVIEW	3
CLOUD PREFILTER: SAVING TIME AND MONEY	3
SANDBOX SERVICE: AI-POWERED DYNAMIC MALWARE ANALYSIS.....	4
THE REPORT: AS COMPREHENSIVE AS IT GETS	5
DYNAMIC ANALYSIS OVERVIEW	5
EVENT LOGGING AND MITRE SUPPORT	6
NETWORK ACTIVITY	6
BEHAVIOR VISUALIZATION	7
INDICATORS OF COMPROMISE	7
EASY INTEGRATION	7
THE CASE FOR A SANDBOX SERVICE.....	8
WHY BITDEFENDER	8
CONTACT US	8

How The Sandbox Service Works: An Overview

Bitdefender's Sandbox Service is a powerful environment to run in-depth analysis of unknown files and URLs. It can protect against evasive zero-day threats, APTs and other sophisticated attacks thanks to advanced AI-powered detection engines.

To access the highly scalable service, partners can submit a sample via a straightforward API or intuitive GUI. The analysis can take a few seconds to a few minutes, depending on the service used (prefilter and/or detonation) and the file itself.

If the partner wants, the sample can be sent to a prefilter before detonation. It uses machine learning algorithms, reputation analysis, multiple unpacking tools, and our award-winning anti-malware engine to return a verdict quickly.

By default, files that receive a definite verdict from our prefilter don't get sent to the sandbox, to save costs and valuable time for partners. However, API calls can be customized to force detonation in the sandbox, or bypass prefiltering entirely.

When a file makes it to the sandbox, it's detonated in a fine-tuned virtual machine that simulates every aspect of a regular endpoint. The detonation has a dynamic execution time, running for as long as it's needed to extract all the relevant data from the sample.

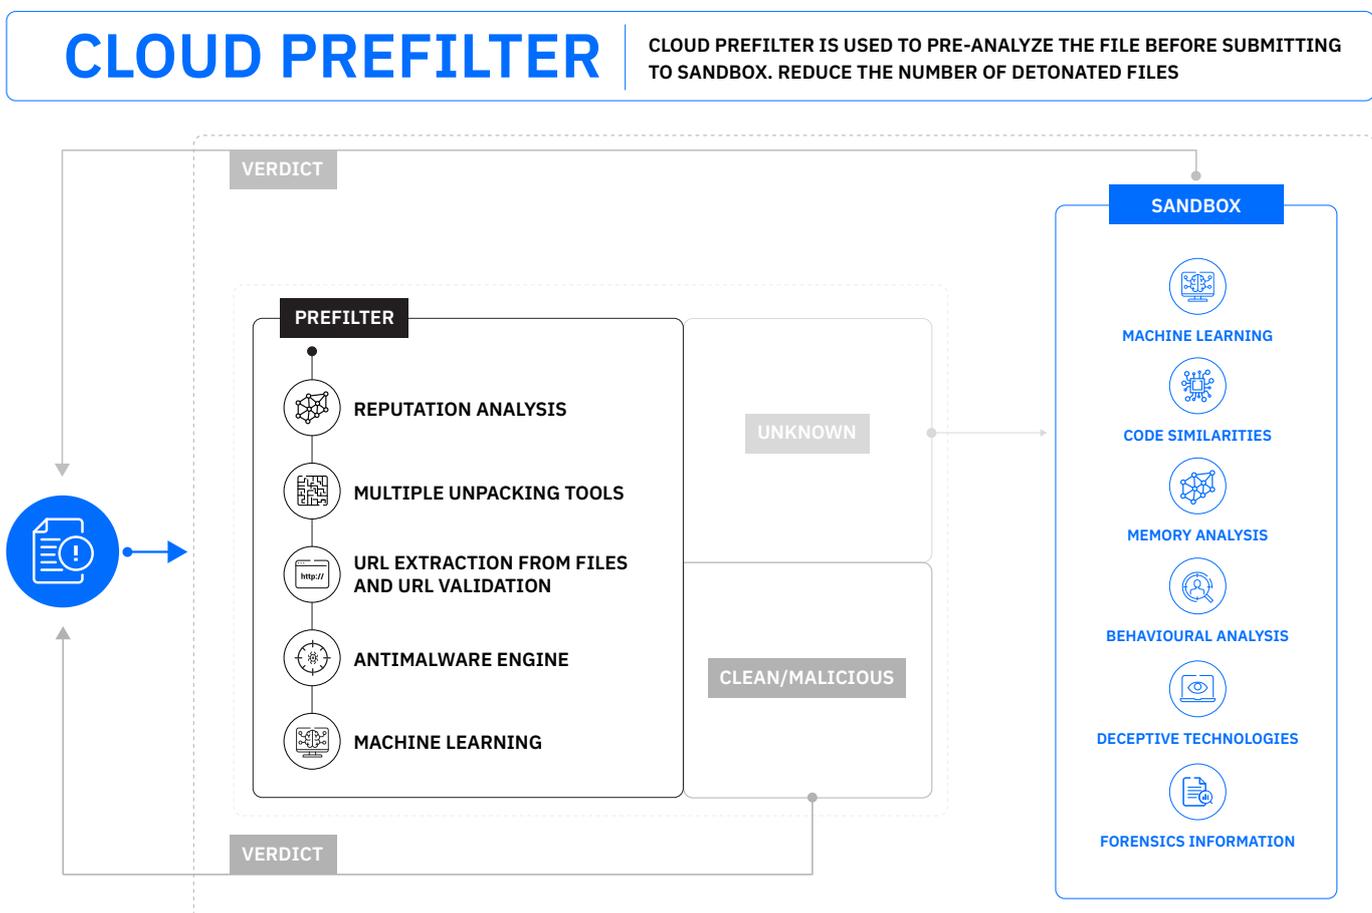
Then, a comprehensive report is returned to the client. It includes in-depth behavioral analysis, event logs, network activity mapping, MITRE references, a list of IoCs detected, file similarities and more.

Alternatively, samples can be submitted via the IntelliZone portal, our one-stop shop for security researchers to consult advanced threat intelligence and use the sandbox service.

Cloud Prefilter: Saving Time and Money

Bitdefender's cloud prefilter uses highly aggressive malware detection engines and heuristics to filter files before detonation.

The prefilter employs deep learning algorithms trained on vast datasets of malicious files. The fine-tuned models empower it to detect potential threats or clean samples with pinpoint accuracy. If the prefilter can return a reliable verdict, the sample is not sent to the sandbox, saving time and money for the partner.

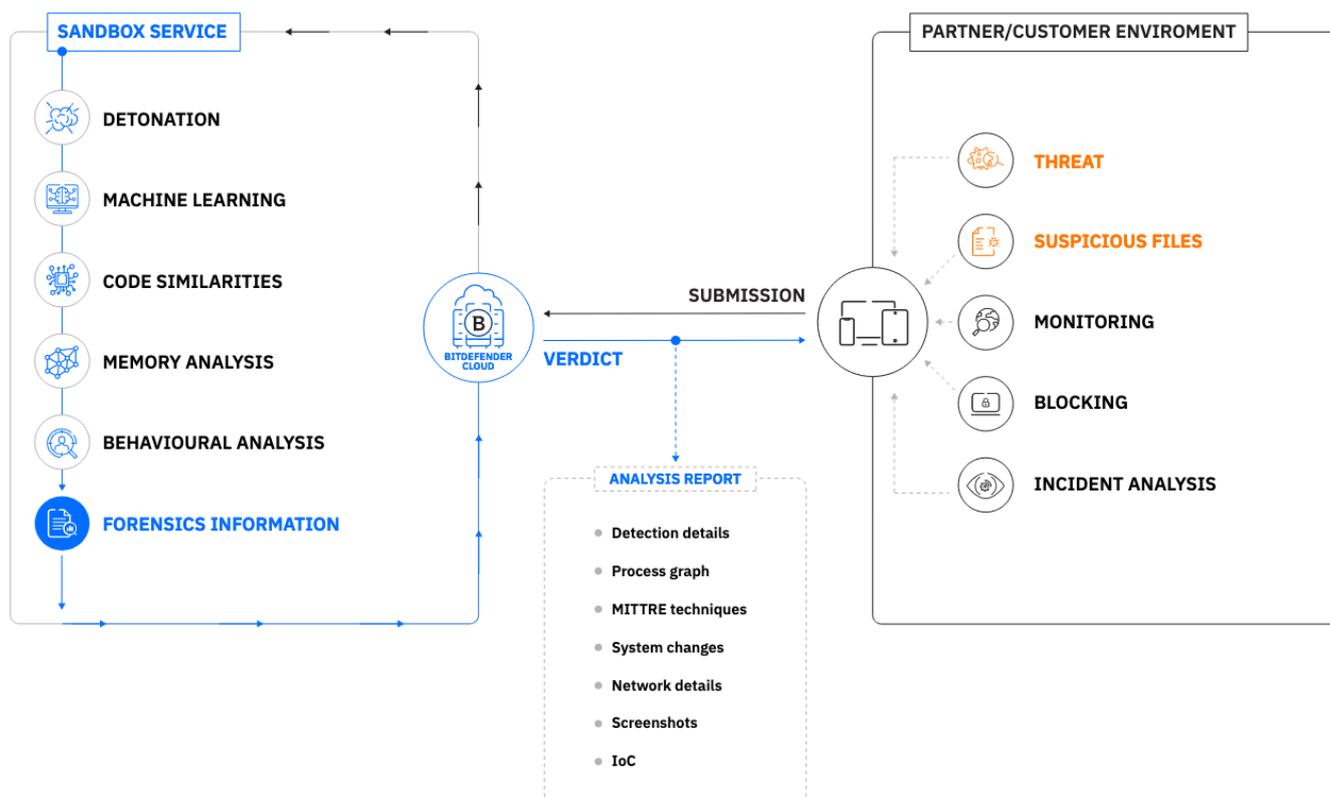


For SOC analysts or researchers that want to extract digital forensics from a sample, the cloud prefilter can be bypassed, or detonation can be forced regardless of the prefilter's verdict, depending on the partner's license.

Sandbox Service: AI-Powered Dynamic Malware Analysis

Bitdefender's Sandbox Service is one of the most advanced dynamic malware analysis tools. It can effectively log every detail about a file's behavior, neutralize sandbox evasion techniques, and even adapt execution time to the sample's behavior.

By default, samples are executed for 3 minutes. Detonation time might be extended, however, if we detect any activity after 3 minutes, including attempts to delay execution.



When executing a sample, the environment closely resembles a real endpoint, complete with common applications and files. Throughout the process, our monitoring technology remains hidden while files are executed in native applications. PDF files are opened by Adobe Reader, documents by MS Office, and various anti-evasion techniques are applied to make sure samples behave as they would in the wild.

Here's an overview of the kind of events logged by the sandbox:

- ↳ Writing, deleting, moving, duplicating, or replacing files on the system and on removable drives.
- ↳ Execution of newly created files.
- ↳ Changes to the file system.
- ↳ Changes to the applications running inside the virtual machine.
- ↳ Changes to the Windows taskbar and Start menu.
- ↳ Creating, terminating, or injecting processes.
- ↳ Writing or deleting registry keys.
- ↳ Creating mutex objects.
- ↳ Creating, starting, stopping, modifying, querying, or deleting services.
- ↳ Changing browser security settings.
- ↳ Changing Windows Explorer display settings.
- ↳ Adding files to the firewall exception list.
- ↳ Changing network settings.
- ↳ Enabling execution at system startup.
- ↳ Connecting to a remote host.
- ↳ Accessing certain domains.

- ↳ Transferring data to and from certain domains.
- ↳ Accessing URLs, IPs and ports through various communication protocols.
- ↳ Checking the indicators of the virtual environment.
- ↳ Checking the indicators of monitoring tools.
- ↳ Creating snapshots.
- ↳ SSDT, IDT, and IRP hooks.
- ↳ Memory dumps of suspicious processes.
- ↳ Windows API function calls.
- ↳ Becoming inactive for a certain period to delay execution.
- ↳ Creating files with actions to be executed at certain time intervals.

On top of logging, the sandbox service uses machine learning to detect malicious actions, extract IoCs, and provide actionable insight to researchers.

When URLs are submitted to the sandbox, one of two things can happen. If the URL points to a file, the sandbox will download it and process it as usual. If the URL is a simple webpage, the sandbox will open it in a browser, and try to detect malicious behavior.

The Bitdefender Sandbox Service is also getting better every day. It's already stopping APTs, zero-day exploits, or sophisticated ransomware from infecting the infrastructure of Bitdefender OEM partners and B2B clients. Since it's constantly queried from all over the world, the service has access to fresh threat data. This enhances its deep learning models, so it can detect complex modern threats with ease.

The Report: As Comprehensive as It Gets

The sandbox service automatically generates a report, in HTML format, about every detonation. It contains crucial details about the sample, like IoCs, event logs, network activity, and more.

Dynamic Analysis Overview

At the top of the report, you'll see a breakdown of the sample, complete with some context on the threat actor, and links to copy the hash of the file, or view it in VirusTotal.

Threat Analysis:



Threat: **Trojan**
Family: **wannacry**
Severity: **99**

 [48ef93a1de0269871907684440722421a77c6221dfeafcb37add89d80aa88f67.exe](#)

Threat Actor:	Lazarus Group
Also known as:	Hastati Group, APTC26, NewRomanic Cyber Army Team, Labyrinth Chollima, Whois Hacking Team, Group 77, Zinc, Hidden Cobra, Nickel Academy, Guardians of Peace, APT-C-26, ITG03, TA404, DEV-0139, Gods Apostles, Gods Disciples, Diamond Sleet
Target countries:	
Target sectors:	Cryptocurrency, Engineering, Financial, Government, Technology
Confidence:	medium

Description: Lazarus Group is a state-sponsored group believed to be run by the North Korean government. It was first recognized in 2009 and is infamous for often destructive attacks. While initial attacks were straightforward (mostly DDoS attacks against South Korean organizations), the group has become more sophisticated over time. The most notable attacks include the widely reported Sony Pictures breach (2014), Bangladesh Bank heist (2016), and WannaCry ransomware attack (2017). During the COVID-19 pandemic, pharmaceutical companies became a primary target (2020). North Korean groups are known for code-sharing, and there is a significant overlap between various groups. Another specific aspect of

This breakdown includes:

- ↳ Submission details, such as the size of the sample, analysis time and document summary.
- ↳ An overview of detections by Bitdefender engines, which can be expanded for more details. For example, the sample from the report highlighted above triggered our margin linear regression classifier, and Kernel-based machine learning algorithms.
- ↳ Alerts, such as attempts to evade anti-malware products, or access restricted directories

Event Logging and MITRE Support

Every action taken by the sample is logged and highlighted in the report, outlining things like modified keys, file operations or network accesses.

This in-depth breakdown can be used by researchers to analyze a file manually, but it's also the basis for a lot of other elements in the report. For example, we use this breakdown to:

- ↳ Map the sample's action to different MITRE techniques and highlight the specific flow that matches known TTPs.
- ↳ Highlight major system changes, such as file operations, network activity, or DNS requests.
- ↳ Outline the sample's file operations, like modified, deleted, temporary, or new files either affected or created by the malware.

Network Activity

Based on network-related events, the sample's network activity is broken down geographically:

General
Incoming Traffic: 22.27 KB (529 packet(s))
Outgoing Traffic: 519.75 KB (9716 packet(s))
Accessed IPs: 5656
Used remote ports: 1
Downloaded Files: 0
Malicious connections found: YES

TCP
Packets: 9716 sent, 529 received

ICMP
Packets: 104 requests

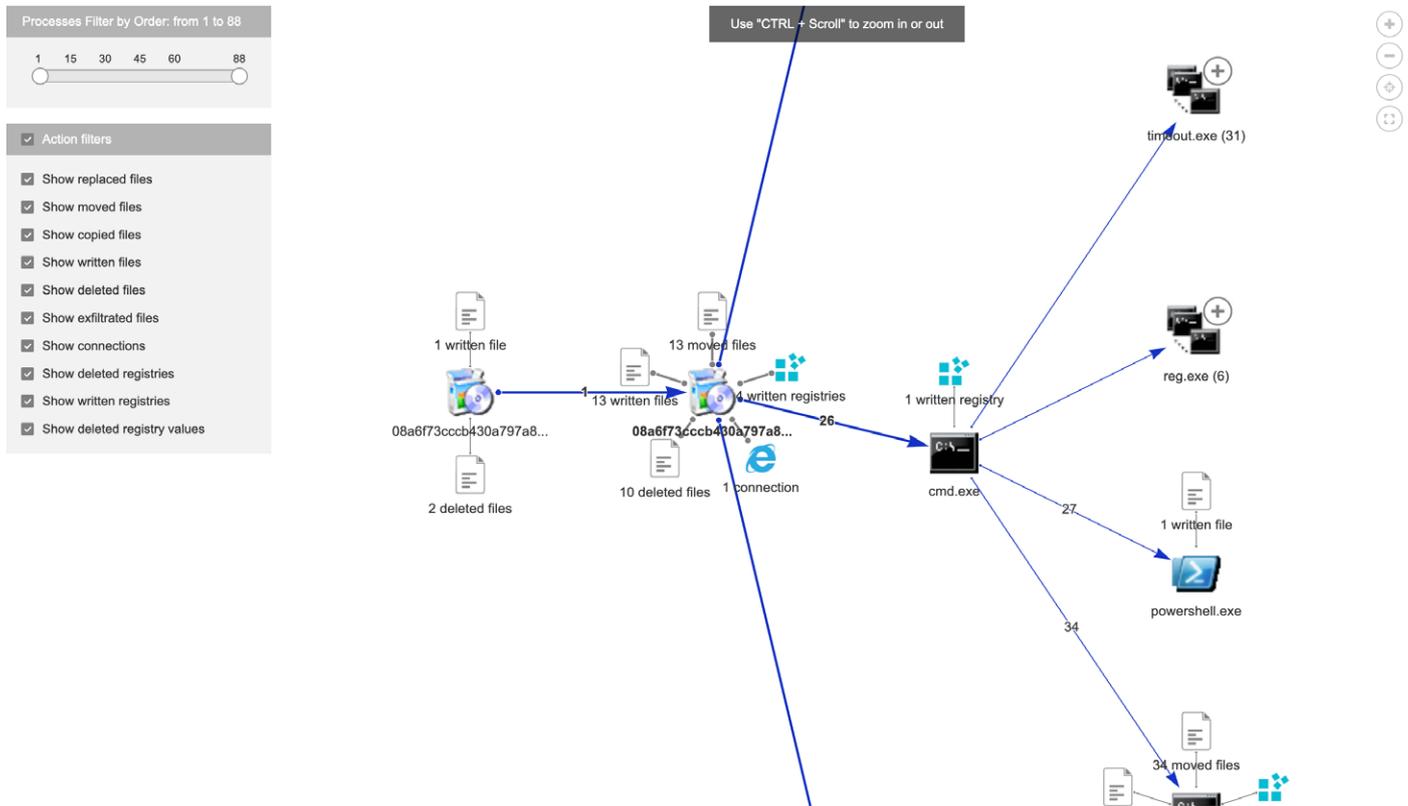


All network activity, like DNS requests, HTTP requests, and network traffic is further detailed beneath the map. Here are some of the details highlighted in the report:

- ↳ A verdict of malicious/clean for IPs, URLs and domains extracted from this network activity. They're all checked against threat intelligence from our telemetry, to detect and provide context on malicious IoCs.
- ↳ A verdict of malicious/clean for network traffic, which is scanned with deep packet inspection.
- ↳ Relevant data like the type of requests made by the sample, HTTP methods used, remote ports accessed, response codes, local IP addresses used, and more.

Behavior Visualization

To help analysts understand malware, the report features several visualization tools, including a timeline view, a process tree view, and a customizable graph.



Researchers can zoom in on the graph and filter actions, for example by hiding file operations if they just want to analyze registry changes. Clicking on any node will reveal additional details, such as process ID, command line script, path, or the MD5 hash.

Indicators of Compromise

Each IoC extracted from this analysis is available to download in plain text, OpenIOC or STIX format. Lastly, the report creates a gallery with relevant screenshots highlighting malicious activity.

Easy Integration

There are two ways for partners to access the Sandbox Service. They can either drag-and-drop files in the IntelliZone portal or submit them using an async API. In either case, it's very easy to access the service.

Files or URLs are submitted via HTTP to our highly available server. This initial request returns a unique Job ID and some context in JSON format. Based on the ID, partners can query the analysis results or report with another HTTP request.

Partners can add their own branding to any HTML report, and request a JSON summary of the results, which includes:

- ↳ The file and submission
- ↳ Threat details
- ↳ Detections
- ↳ System changes
- ↳ MITRE techniques
- ↳ IoCs

Moreover, partners can download all the IoCs extracted from a sample in either XML for OpenIOC formats, or JSON for raw data and STIX.

The Case for a Sandbox Service

As cybercriminals develop more sophisticated techniques to remain elusive, the cost and complexity of managing threats is growing exponentially. Zero-day malware has become more prevalent than ever, often bypassing existing security layers.

Businesses of all sizes are facing zero-day exploits, targeted attacks, and advanced persistent threats that have never been seen before and are specifically designed to evade traditional malware defenses.

A sandbox service is crucial to stop these kinds of threats before it's too late. Sophisticated malware will only be noticed once it's already done damage to an organization's infrastructure. By detonating unknown files in a sandbox, this damage is inconsequential, limited to a secure virtual environment.

But increasingly clever threat actors aren't the only reason to invest in a sandbox. Analysts and security researchers are in turn overwhelmed by alert fatigue, and an ever-expanding threat surface, especially with the move to cloud infrastructures.

A sandbox solution can help them detect, extract, triage and analyze threat intelligence with ease. Quick verdicts help limit the number of incidents or alerts, while in-depth analysis provides actionable insights.

Why Bitdefender

Bitdefender has been at the forefront of cybersecurity innovation for decades. Our products constantly top the charts of independent tests, with high detection rates and minimal false positives.

Most importantly, we know what a good sandbox needs because we use it ourselves. Dynamic analysis in the sandbox is a crucial part of our award-winning EPP products. Plus, Bitdefender operates 2 SOCs, employing hundreds of analysts and security researchers.

A reliable sandbox is indispensable to our day-to-day operations. The Sandbox Service is efficient and insightful because it was **built by and for security experts**.

Contact Us

For more information regarding Bitdefender Sandbox Service please visit our website at bitdefender.com/oem/sandbox-service.html.

Evaluating the Bitdefender Sandbox Service is free of charge and includes technical support. Scan the QR code to contact us and request a free trial:

