SentinelOne

SentinelOne Technical Brief

SentinelOne unifies prevention, detection and response in a fundamentally new approach to endpoint protection, driven by behavior-based threat detection and intelligent automation.

For today's enterprise user endpoints and data centers alike, SentinelOne's approach addresses the entire threat execution lifecycle (pre-execution, on execution, and post-execution) to detect advanced malware, exploits and sophisticated attacks, and respond to any threat at machine speed. This Technical Brief will provide a detailed overview of the key features and underlying technologies employed by SentinelOne's Endpoint Protection and Data Center Protection platforms.

PLATFORM DESCRIPTION

The SentinelOne Endpoint Protection Platform (EPP) and Data Center Protection Platform (DCPP) are both highly-scalable, lightweight agent-based solutions. All detection and preventative measures are performed on the endpoint device, and agents are all connected to a central management console.

The SentinelOne EPP protects Windows, Mac OS X and Linux-based endpoint devices, and SentinelOne DCPP deploys across physical, virtual, and cloud-based servers running Windows and Linux. System requirements are detailed in a separate section at the end of this document.

KEY CAPABILITIES AND PLATFORM TECHNOLOGY

SentinelOne Endpoint Agent

The SentinelOne agent is a lightweight, small-footprint module that is installed on the endpoint device or server. On a user endpoint device, it monitors all activity at both the kernel level and in user space. For servers, the agent does not sit in-line, as a means of preserving server performance and flexibility. In general, the minimal overhead incurred with monitored operations is 4 micro seconds. For an endpoint device used in a typical way over a 24-hour period, this amounts to a total delay of just one second. SentinelOne's monitoring process runs at low priority on the system, and consumes between 0%-4% CPU cycles. It's memory footprint is about 20MB and the agent occupies approximately 200MB on disk. Agents can be deployed using a standard MSI/PKG package.

Activity Monitoring

On a user endpoint device, the SentinelOne agent taps every process and thread on the system. It extracts all relevant operations data: system calls, network, IO, registry (on Windows), and more. This is so that it can track the behavior of every process executing on the system. Traditional antivirus (along with other prevention solutions that use inline processes) uses static signatures or other reputation methods to evaluate executing binaries for the purpose of determining whether or not a file is malicious. By contrast, SentinelOne's approach doesn't require inline placement. The agent automatically "taps" to obtain operation data, allowing the running process to continue while monitoring everything the process does during and after execution.

The monitoring module asynchronously sends endpoint operation data to a preprocessing module, which analyzes the data to build a full context around every process. This stage translates the raw monitored operations data log into a much more structured, abstract operation language.

©2016 SentinelOne. All Rights Reserved.

Dynamic Behavior Analysis

The analyzing module works constantly in the background and runs sophisticated pattern matching algorithms to detect malicious behaviors against a full context of normal process operations. It looks system-wide at operations, as well as at historical information.

Patterns are derived from analyzing attack behaviors and techniques. SentinelOne's cyber threat researchers focuses on reverse-engineering thousands of malware samples and other vectors of attack daily. In the lab, samples are clustered, and then threat behaviors are analyzed and scored. The analyzing module scores every suspicious pattern detected during process execution, and once the aggregate score exceeds a threshold, the process is considered malicious. Suspicious patterns of execution are typically different techniques or interactions with the operating system that a threat employs throughout its execution lifecycle. This lifecycle typically includes the following stages: exploitation, obfuscation, persistence, collection, and exfiltration.

Mitigation

When a process is considered malicious, the mitigation module executes the actions that stop the threat from moving laterally. Via the SentinelOne management console, a user can either manually kill or quarantine a malicious process or file, or create individual mitigation policies to be automatically executed upon threat detection at the endpoint. Automated mitigation policies can be applied to customized groups of endpoints, as well.

Remediation

Following an attack, SentinelOne can easily restore deleted files, and roll back modified files to their previous trusted states. This capability leverages Windows Shadow Copy.

Immunization

Each time a new, never-before-seen threat has been detected on a user or server endpoint through SentinelOne's dynamic behavior tracking engine, it is instantly signed and all other SentinelOne agents on the network are notified with the update. This makes the entire network immune to the unknown attack, which is prevented from spreading and from running on other machines.

Prevention

Though antivirus is an antiquated solution, static prevention is still effective in blocking known threats. SentinelOne provides a layer of preemptive protection by leveraging multiple leading cloud reputation services. With the EPP/DCPP's 'Cloud intelligence' setting, SentinelOne sends hashes from executed binaries that exhibit suspicious behavior. This feature also uses several leading scan engines to check the file's reputation. Binaries identified as 'malicious' are proactively blocked, while benign ones are added to the whitelist to minimize false positives.

Detailed, Real-Time Forensics

Constant monitoring of all processes on the endpoint enables SentinelOne to provide real-time forensics and a 360° view of attacks via a single management console, accessible from any device, anywhere. Security or Incident Response analysts can quickly access forensic data, and investigate to determine the root cause of an attack and accelerate incident response activities. All of the activity monitoring data collected by the agent is sent back to the management console over an encrypted SSL link, and stored on the management server in encrypted file systems (for details on types of data collected, refer to the Appendix).

The SentinelOne EPP and DCPP use this data to compile real-time forensic information to identify where attacks originated, and how they unfolded in terms of processes, affected files, etc. In addition, SentinelOne's forensic data can be easily offloaded to popular SIEM systems, including Splunk, LogRhythm, for further investigation or sent to network security devices for proactively blocking threats at the gateway.

©2016 SentinelOne. All Rights Reserved.

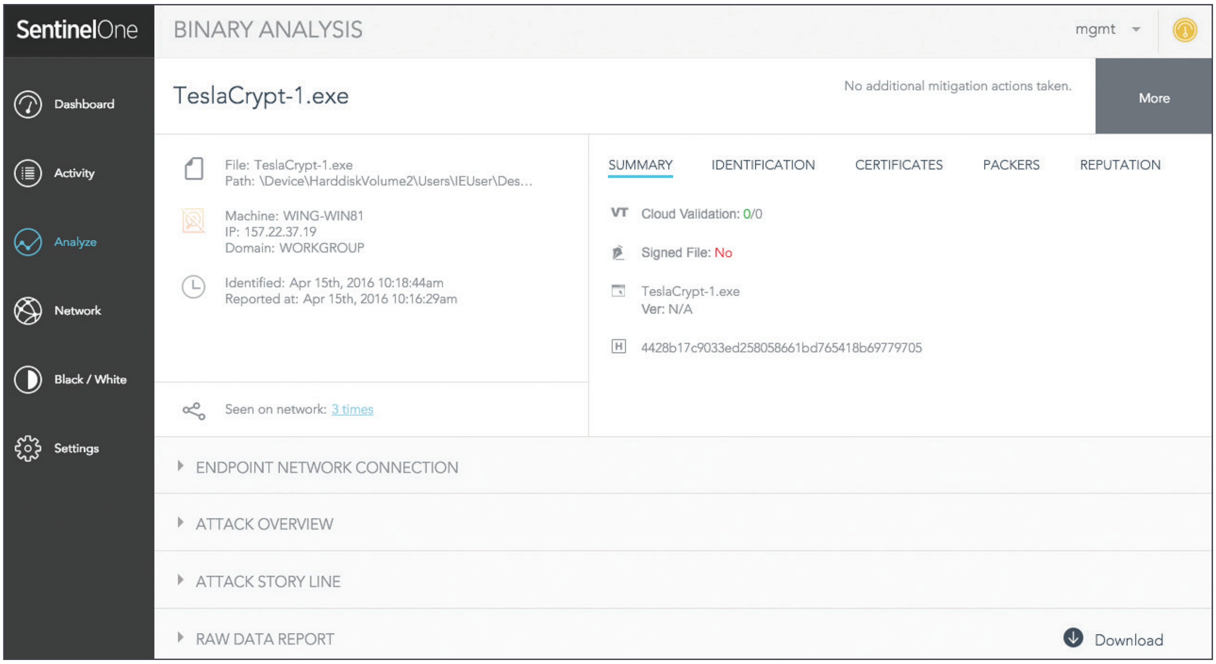

TeslaCrypt-1.exe

No additional mitigation actions taken.

File: TeslaCrypt-1.exe
Path: \Device\HarddiskVolume2\Users\IEUser\Des…

Machine: WING-WIN81
IP: 157.22.37.19
Domain: WORKGROUP

Identified: Apr 15th, 2016 10:18:44am
Reported at: Apr 15th, 2016 10:16:29am

Seen on network: 3 times

Cloud Validation: 0/0

Signed File: No

TeslaCrypt-1.exe
Ver: N/A

4428b17c9033ed258058661bd765418b69779705

ENDPOINT NETWORK CONNECTION

ATTACK OVERVIEW

ATTACK STORY LINE

RAW DATA REPORT

360° view of attacks

SentinelOne EPP and DCPP provide a 360° view of attacks including:

SUMMARY INFORMATION

This section of the SentinelOne EPP or DCPP management console outlines the basic attack details, including attack statistics, dwell time, file information, path, machine name, IP, domain, along with information about where else on the network the attack has been seen. In addition, this section shows cloud reputation validation, certificate information (if the file is signed or not), and advanced attack details (such as a list of known packers that may have been used).

ATTACK OVERVIEW

Detailed information about indicators SentinelOne used to determine if a process was malicious, including capturing attack statistics and dwell time. See the table below for a complete explanation of the different event categories.

ATTACK STORY LINE

The Attack Story Line is an intuitive visualization of how an attack propagated during its execution. This graph depicts the processes the attack created, terminated, or tainted, the low-level kernel and API calls it made, the files it dropped, altered, deleted or created, the registry keys it changed, created, or deleted (along with their values), and which inbound or outbound network connections were made.

RAW DATA

A comprehensive line-by-line detailed view of changes made to the system, files, processes, and registry settings.

The forensic reports are accessed through the management console and provide rich, visual details in real time that simplify collection and analysis of security incident data to accelerate investigative efforts. This information enables analysts to easily determine if other machines on the network were also compromised.

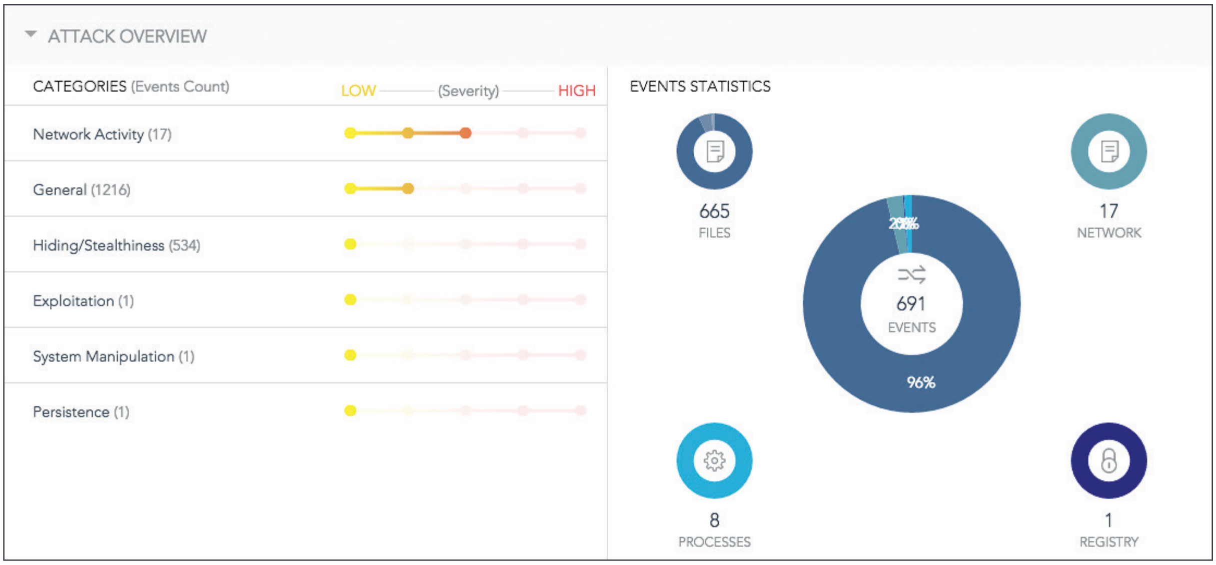

ATTACK OVERVIEW

CATEGORIES (Events Count)

Network Activity (17)

General (1216)

Hiding/Stealthiness (534)

Exploitation (1)

System Manipulation (1)

Persistence (1)

Attack Overview

The Attack Overview section provides a breakdown of the different malicious behaviors that were detected, along with their associated risk levels. In addition, it reports key activities performed by the attack, the attack's dwell time, and the number of network calls made.

CATEGORIES	AGENT MONITORS MALWARE ATTEMPTS TO:
HIDING/STEALTHINESS	Hide operation from traditional antivirus solutions, as well as from the user. Common methods include: modifying registry keys or file attributes, using obscure file names and code obfuscation. Other techniques the agent monitors are: sophisticated code injections, in memory encryption/ decryption, and the use of commercial or custom/modified packers.
PROCESS OPERATIONS	Manipulate process operations by performing remote code injections to other processes, hiding processes and services, as well as elevating or manipulating processes.
SPYING	Track user behavior (e.g., log keystrokes, take screenshots) through API, sys, or IO calls.
ANTI-DETECTION	Evade detection from standard anti-virus solutions through obfuscation techniques such as deleting its own files or leveraging packers.
GENERAL	Perform behaviors that may not be strictly malicious in isolation, but provides additional context to help determine whether the process is part of an attack flow or not.
EXPLOITATION	Take advantage of vulnerabilities through memory manipulations, privileged function calls, or buffer overflows.
SYSTEM MANIPULATION	Manipulate operating system files that typically do not change often (e.g., registry settings, task scheduler, etc). This enables malware to take advantage of the system to avoid detection, persist, collect data, and mitigation.
NETWORK ACTIVITY	Connect to command and control servers. The purpose is to allow malware to download additional components or exfiltrate data.
PRIVILEGE ESCALATION	Elevate user privilege levels to gain access to system resources. This would allow malware to perform unauthorized actions including modifying files and settings or access to system resources.
PERSISTENCE	Persist on the system using a number of approaches such as, loading itself after a system reset through operating system manipulation (e.g., task scheduler, registry settings, launch agents, etc), injecting into existing system libraries, and modifying the master boot record.

©2016 SentinelOne. All Rights Reserved.

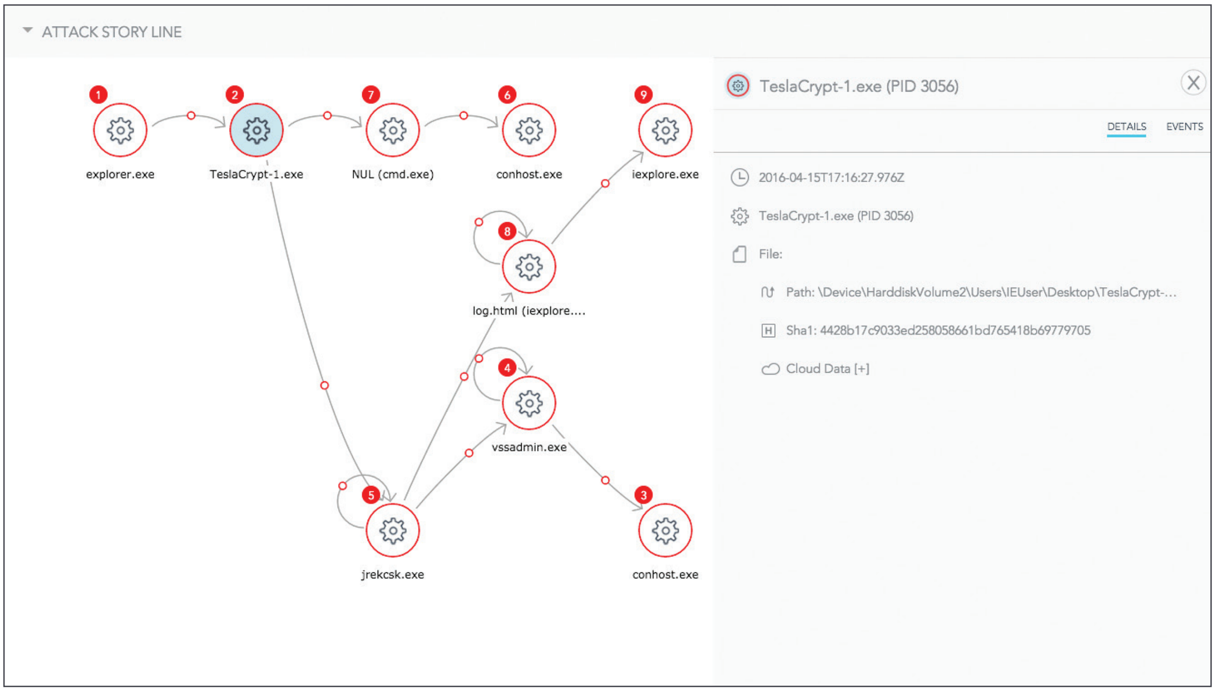

TeslaCrypt-1.exe (PID 3056)

2016-04-15T17:16:27.976Z

TeslaCrypt-1.exe (PID 3056)

File:

Path: \Device\HarddiskVolume2\Users\IEUser\Desktop\TeslaCrypt-...

Sha1: 4428b17c9033ed258058661bd765418b69779705

Cloud Data [+]

Attack Story Line

The Attack Story Line report provides a detailed graphical view of the threat execution flow including the sequence of events, malicious behaviors, and affected system components. The unique visual format of the report graphically correlates chain-related events of affected systems. This helps analysts minimize efforts needed to investigate security incidents and take action in response. Details provided with this view include the names of the malicious processes, actions taken (e.g., creating, modifying, or deleting other system files, including registry settings or processes), and the sequence of the execution flow. In addition, users can select a specific process on the attack story line and view the individual network, file, process, or data actions that were taken.

©2016 SentinelOne. All Rights Reserved.

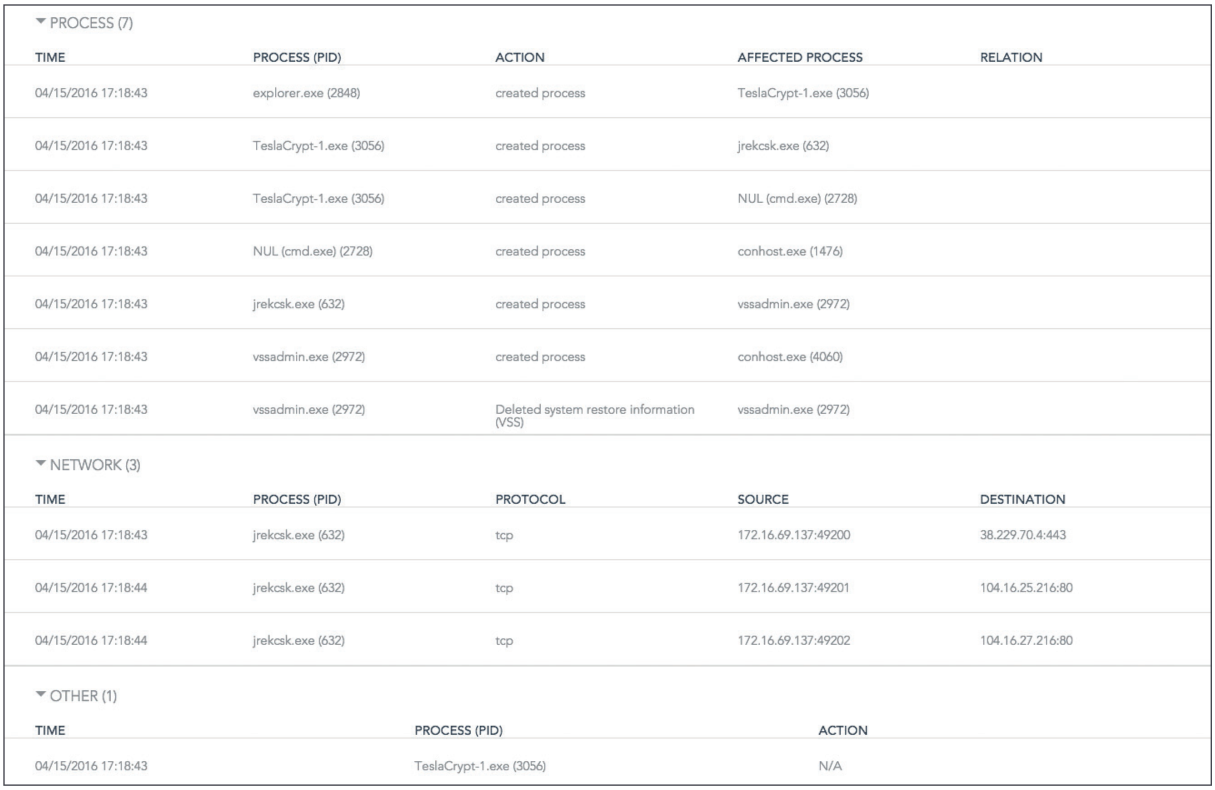

▾ PROCESS (7)

TIME	PROCESS (PID)	ACTION	AFFECTED PROCESS	RELATION
04/15/2016 17:18:43	explorer.exe (2848)	created process	TeslaCrypt-1.exe (3056)	
04/15/2016 17:18:43	TeslaCrypt-1.exe (3056)	created process	jrekcsk.exe (632)	
04/15/2016 17:18:43	TeslaCrypt-1.exe (3056)	created process	NUL (cmd.exe) (2728)	
04/15/2016 17:18:43	NUL (cmd.exe) (2728)	created process	conhost.exe (1476)	
04/15/2016 17:18:43	jrekcsk.exe (632)	created process	vssadmin.exe (2972)	
04/15/2016 17:18:43	vssadmin.exe (2972)	created process	conhost.exe (4060)	
04/15/2016 17:18:43	vssadmin.exe (2972)	Deleted system restore information (VSS)	vssadmin.exe (2972)	

▾ NETWORK (3)

TIME	PROCESS (PID)	PROTOCOL	SOURCE	DESTINATION
04/15/2016 17:18:43	jrekcsk.exe (632)	tcp	172.16.69.137:49200	38.229.70.4:443
04/15/2016 17:18:44	jrekcsk.exe (632)	tcp	172.16.69.137:49201	104.16.25.216:80
04/15/2016 17:18:44	jrekcsk.exe (632)	tcp	172.16.69.137:49202	104.16.27.216:80

▾ OTHER (1)

TIME	PROCESS (PID)	ACTION
04/15/2016 17:18:43	TeslaCrypt-1.exe (3056)	N/A

Raw Data Report

For a more in-depth look at each of the events associated with a security incident, the Raw Data report provides comprehensive attack-related details including activity for files, network, processes, and registry (Windows only).

FILE

The File section provides further detail about files involved in an attack including the timestamp, file names, actions executed, and the file location.

PROCESS

The Process section contains details on processes involved in an attack including the timestamp, process name/ID, process actions executed, the names of impacted processes, and the relationship of those processes.

NETWORK

The Network section includes details about connections a malicious process attempted to make. It lists the protocol used, the source and destination addresses, and the time of when these connections were attempted.

REGISTR

The Registry section provides specific information about the registry key associated with the attack as well as the action performed, the time the action took place, and the location of the registry key.

©2016 SentinelOne. All Rights Reserved.

SYSTEM REQUIREMENTS

CLIENTS	
OPERATING SYSTEMS	SentinelOne Endpoint Protection Platform • Windows 7, 8, 8.1 • .Net 4.5 • OS X 10.9x, 10.10x • Red Hat Linux, CentOS 6.5 or higher SentinelOne Data Center Protection Platform • Windows Server 2008 R2, 2012 R2 • Red Hat Linux, CentOS 6.5 or higher
VIRTUAL ENVIRONMENTS	• vSphere, Microsoft Hyper-V, Citrix Xen Server, Xen Desktop, Xen App

MANAGEMENT SERVER ***(ON PREMISE)***	
OPERATING SYSTEM	• Linux Ubuntu 14.04 LTS Server
HARDWARE	• 4-core Intel Xeon E5-2680v2 2.8 GHz or better • 8 GB RAM • 500 GB free disk space

APPENDIX - DATA COLLECTION

The following sections list the types of data collected by the SentinelOne agent.

HARDWARE DATA	• CPU data (ID, architecture, # of cores, clock speed) • RAM size • Disk size • Hardware device info • Device type (Desktop/Server/Mobile)
USER DATA	• User name • Machine name • Workgroup/domain
VERSION DATA	• Installed OS version • Installed SentinelOne EDR agent version
PROCESS ACTIVITY	• Time of machine activity • Running processes (name, ID, CPU usage, memory) • Low level System calls • User space API calls • For each process the SentinelOne EDR agent collects: - File access, metadata only (full path, file type, type of access, time of access etc.) - Network access, metadata only (IP, protocol used, time of access etc.) - Memory access, metadata only (memory addresses, permissions, sources, targets) - Registry access [Windows only] (keys created, altered, deleted, values) - Registry modified content [Windows only] (values of new or modified keys)
NETWORK	• Internal network IP address, domain name, DNS server • Public IP address (if running cloud-based management) • URLs accessed • Inbound/Outbound connections, metadata only (source, target, and application)

©2016 SentinelOne. All Rights Reserved.