CLOUDPHYSICS

DATA-DRIVEN INSIGHTS FOR SMARTER IT

CloudPhysics Observer Security FAQ

CloudPhysics has applied our extensive security expertise to ensure customer data is protected at all points of the data transactions. This document answers some of the most common security questions. For a more extensive discussion of our security procedures, please contact us.

Q. What is the CloudPhysics Observer?

CloudPhysics collects data from your environment using a virtual appliance called the CloudPhysics Observer. This virtual appliance is a minimum resource appliance designed to collect data from within your VMware vCenter and cloud environment through read-only APIs, process the data, and share the data to CloudPhysics through secure means.

Additional levels of data collection are available with elevated privileges for guest process discovery using VMware Tools APIs and limited guest credentials at the discretion of the vSphere Admin.

Q. What are the system requirements of the CloudPhysics Observer?

The virtual appliance requires the following resources:

8GB of RAM, 2 Virtual CPU's, and 20GB of disk space when deployed. Total network traffic resources will be approximately 5MB per hour per 100 VMs in the datacenter.

The CloudPhysics observer will also require internet access to send collected data to the public cloud through an encrypted connection to the internet domain:

entanglement.cloudphysics.com

This domain is used for API calls only and has no public accessible web pages.

These communications will occur on Port 443. The Virtual Appliance must be on a network LAN segment that has access to VMware vCenter for VMware vSphere data collection. Data collection for cloud providers will be collected directly by CloudPhysics from the public cloud through published APIs.

Q. Does CloudPhysics deploy any agents to hosts or VM guest operating systems?

CloudPhysics does not deploy any probes or agents to VMware ESXi Hosts or any guest OS. All communications are achieved through existing management interfaces and results in no additional load to the host environment. CloudPhysics can take advantage of VMware Tools to collect process details within a guest if already deployed but is not required for infrastructure data collection.

Q. How is Data Collected?

CloudPhysics Observer collects data from VMware vCenter and cloud providers by public APIs. For VMware vCenter, CloudPhysics collects performance, configuration, and other metadata from the VMware vCenter on a defined schedule. Natively, vCenter collects performance and configuration data from its managed resources on a 20-second granularity. This data is typically rolled-up and destroyed once data is an hour old by vCenter. Before data is rolled-up and destroyed, CloudPhysics collects this performance and configuration data frequently enough directly from vCenter to maintain the 20-second granularity. This data collection process is agentless and has no impact on the VMs or hosts being analyzed since it already exists in vCenter. For cloud providers, CloudPhysics collects configuration and performance history data from the public APIs once per day. For VMware vCenter, CloudPhysics requires a Read-Only account with access to list and read configurations of the virtual environment.

For VMware vCenter, these credentials are detailed in the CloudPhysics install guide located at https://www.cloudphysics.com/installing-cloudphysics/ and for Amazon Web Services, the policy details can be found at https://www.cloudphysics.com/connectaws/

Q. What type of data is Collected?

Infrastructure Configuration Data
This data describes either the virtual datacenter or the cloud environment under observation by CloudPhysics. This data defines the environment to be monitored including the vCenter and its configurations as well as the resources consumed by the systems and resources under management by vCenter. This data does not include network topology or data to recreate the network architecture. For VMware vCenter v4.0 and above, this data will consist of vCenter details, datacenter details, VM details, host details, virtual domain details, datastore details, network port details, virtual network details, and resource group details.

Performance Data
This data will consist of CPU, Storage, Network, and RAM usage details. Utilization, peak performance, bandwidth, and characteristics of these will all make up the performance data. CloudPhysics will also generate derivatives of this data for averages, means, 99th Percentile, and 95th Percentiles.

CLOUDPHYSICS

Task Data
Task data provides a view of major events and scheduled services in the environment such as resources starting and stopping, vMotions, and environmental changes. These events and tasks often include the event, a brief description.

Metadata and Tags
Many resources contain metadata to describe a service, its role, and provide context to its relationship to other objects in the environment. The most common metadata collected are tags used for managing objects in the environment to offer classification and organization of resources, data, and services.

Running Processes within VMware VMs
With the addition of the Fall 2018 Observer Refresh, CloudPhysics provides administrators the option to collect inventories of running processes within a VM. This data collection is achieved as a guest request through VMware tools and allows the VMware tools to return a list of processes currently running on the host to help classify applications and services associated with VMs.

Q. Where is my data stored?

Data collected by the observer is quickly processed and parsed to remove unnecessary data before being compressed, encrypted, and sent to the CloudPhysics servers for data processing. The most recent data collections will be held in the Observer until they can be delivered to the CloudPhysics cloud. CloudPhysics stores each customer's data in dedicated logical containers until the data can be queued, verified, and loaded into the CloudPhysics data lake for analysis.

Q. How is data protected in transit to the Cloud?

CloudPhysics communicates over TLS 1.2 for current observers on Port 443 from the CloudPhysics Observer to CloudPhysics. Communications to Amazon Web Service will occur over secure REST API communications over HTTPS (TLS) on Port 443. All communications are encrypted using the latest supported secure standards for data communications.

Q. Can the CloudPhysics Observer operate through a network proxy?

Yes, the CloudPhysics Observer supports proxies implementing the HTTP Proxy protocol. Unauthenticated and authenticated proxies using Basic, Digest or NTLM authentication are supported. SOCKS and Transparent (intercepting) proxies are not supported. When connecting through a proxy the Observer will use the HTTP CONNECT method to connect directly (and exclusively) to entanglement.cloudphysics.com on port 443. Loading a custom TLS Certificate Authorities is not supported and the appliance will fail to function if the proxy attempts to intercept TLS traffic.

Q. Is any personal identifiable information collected?

CloudPhysics collects data center configuration and performance data. As a result, there is minimal exposure of personal identifiable information collected. CloudPhysics only collects user information for portal account access and invitations of new users by existing users. This data will consist of company, name, and email address only. This data is used by the organization for user account management and credentialing.

Q. How large of an environment will CloudPhysics support?

CloudPhysics is not limited to the number of VMs, Hosts, Servers or Clouds. The current cloud model allows for one CloudPhysics Observer per vCenter to allow scalability. Data sent to the cloud will queue for processing and the environment will scale dynamically to accommodate capacity.

Q. How is the Observer secured and how often is it updated?

The CloudPhysics Observer is a hardened guest. All unnecessary services, packages and users have been removed. Collection code runs in separate process and network namespaces from the base appliance and these namespaces are deleted and recreated from an immutable base image on each reboot of the appliance.

Q. How does CloudPhysics monitor availability and integrity of hosts within our environment?

CloudPhysics does not monitor the availability of systems within the customer organizations beyond the most recent communication between the CloudPhysics Observer and CloudPhysics cloud services. We utilize internal and third party services to monitor availability and functioning of hosts within our infrastructure.

Q. What credentials are required to be granted to CloudPhysics to access the vCenter?

CloudPhysics needs a limited access account that has read and list capabilities against VMware vCenter. Details for security and policy requirements for vCenter are detailed at https://www.cloudphysics.com/installing-cloudphysics/

Q. What connectivity and protocols are used by the Observer?

CloudPhysics communicates over TLS 1.2 for current observers on Port 443 from the CloudPhysics Observer to CloudPhysics. Communications to Amazon Web Service will occur over secure REST API communications on TLS and HTTP on Port 443. All communications are encrypted using the latest supported secure standards for data communications.

CLOUDPHYSICS

Guest Process and Application Discovery Questions

Q. Am I required to configure the guest process collection or dependency mapping collection?

No. Data collection within guest operating systems is entirely optional and definable during the CloudPhysics Observer installation and configuration.

Q. Can I disable guest process collection and network dependency mapping?

Yes. Dependency mapping and guest process collection are options that require dedicated credentials during the setup of the CloudPhysics Observer. If no credentials are provided, the collection process will not be executed.

Q. How are the guest processes collected?

Guest processes are collected with a VMware Tools feature to collect guest processes. This request originated from the CloudPhysics Observer to VMware vCenter. Upon request, vCenter will attempt to issue the command to the VMware tools within the guest OS. The VMware vCenter will initiate a process collect command under the identity of the guest account specified in the CloudPhysics observer during the Observer setup process. The VMware Tools will issue the command as the specified guest user every six hours. If the guest OS allows the guest user, the process list from the host is collected and stored in a guest user home directory. Upon completion of the collection, Output of command execution is collected by CloudPhysics Observer using vSphere API that in turn uses VMware tool to collect the command output temporarily stored in the output file.

Assuming user have access to their own home directory, the application data will be written to the user home directory and removed upon data collection. If the user does not have sufficient rights to delete their temp files, the file will be overwritten with each collection to ensure the volume storage is minimal.

Q. How is dependency mapping data collected?

Dependency Mapping is derived from a network analysis tool called NetStat. CloudPhysics issues a request to VMware vCenter for details form the guest OS. VMware vCenter can direct queries to the guest OS if VMware Tools is deployed and enabled. The request will be a simple command to issue a NetStat command and direct the output to a temporary file located in the guest user's home directory. The NetStat command will collect all open network communications and report the source IP Address, Destination IP Address, TCP/UDP, as well as port. This data is directed into a local temp storage file where it is processed and sent to the VMware vCenter by VMware Tools.

Q. How frequently is my data collected?

CloudPhysics will collect both guest process and network dependency data independently on a defined schedule. Initial releases will collect data every six hours.

Q. How do you create a dependency map?

Dependency maps are generated based on source and destination IP Address and ports identified by NetStat during the dependency mapping data analysis online. These data will identify all major network communications by the guest OS and map IP addresses to other VMs. Any VM that talks outside of the private network ranges would be considered communications outside of your data center.

Q. What credentials are required to collect guest processes and Dependency Data?

A domain guest ID is best for collection of data. This user credential does not need to be a domain admin or have root access within a guest OS. For mixed environments, ensure the same user id and password exists in both Linux and Windows environments.

Q. What data is collected for guest process?

A simple table of process ID and Process Name is generated when the vSphere API command is issued. This command returns back a simple text list of all processes currently running in the guest OS.

Q. What data is collected for Dependency mapping?

NetStat returns a text output of the source, destination, port, and potentially protocol information from the guest. This data varies slightly from the operating system to operating systems but typically. Additional data may include packet count, state, or world ID.

Q. What is the data flow during collection?

CloudPhysics issues a request for data to VMware VCenter for a specific guest OS. VMware vCenter will issue the credential and command to the guest OS. If the command is allowed to execute, VMware tools will direct all output from the command to a temp file in a guest user home directory. Upon completion of the command, VMware tools retrieve the temp file and direct the output back to VMware vCenter as a temporary variable for the guest OS. CloudPhysics will then collect the temp variable from the VMware vCenter on the next data collection cycle. If the data collection fails or an error is generated, this data is also reported back to the VMware vCenter for collection by the CloudPhysics observer.

Rev. March 18, 2020