HPE CloudPhysics

HPE CloudPhysics Security FAQ for VMware vSphere Virtual Appliance

HPE 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 at cloudphysics@hpe.com.

Q. How is Data Collected?

The HPE CloudPhysics virtual appliance collects data from VMware vCenter and cloud providers by public APIs. For VMware vCenter, HPE CloudPhysics collects performance, configuration, and other metadata from the VMware vCenter on a defined schedule. Natively, VMware 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 the VMware vCenter. Before data is rolled-up and destroyed, HPE 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 VMware vCenter. For cloud providers, HPE CloudPhysics collects configuration and performance history data from the public APIs once per day. For VMware vCenter, HPE 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 HPE CloudPhysics install guide located at https://www.cloudphysics.com/installing-cloudphysics/#Credentials

Q. What is the HPE CloudPhysics virtual appliance?

HPE CloudPhysics collects data from your environment using a virtual appliance called the HPE CloudPhysics virtual appliance. 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 HPE 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. Does HPE CloudPhysics deploy any agents to hosts or VM guest operating systems?

HPE 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. HPE 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 the virtual appliance secured and how often is it updated?

The HPE CloudPhysics virtual appliance is a hardened linux based guest os. 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. What are the system requirements of the HPE CloudPhysics virtual appliance?

The virtual appliance requires the following resources:

8GB of RAM, 2 Virtual CPU's, and 120GB of disk space when deployed. Total network traffic resources will be approximately 5MB per hour per 100 VMs in the datacenter. This data may be revised periodically and will be posted to the Installing CloudPhysics support page.

Q. Can the HPE CloudPhysics virtual appliance operate through a network proxy?

Yes, the HPE CloudPhysics virtual appliance 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 virtual appliance 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. Can I install the HPE CloudPhysics virtual applaince in a network without internet Access?

No. The HPE CloudPhysics virtaul applaince requires internet access to send colelcted data and download security updates. The virtaul applaince will need to talk both to the VMware vCenter and the public internet.

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

HPE 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/#Credentials

Q. What connectivity and protocols are used by the virtual appliance?

HPE CloudPhysics communicates over TLS 1.2 for current observers on Port 443 from the HPE CloudPhysics virtual appliance to HPE CloudPhysics. Communications to our serrvers and 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.

HPE HPE CloudPhysics Virtual Applaince Security FAQ for VMware vCenter (continued)

Q. What type of data is Collected?

Infrastructure Configuration Data
This data describes either the virtual datacenter or the cloud environment under observation by HPE 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. HPE CloudPhysics will also generate derivatives of this data for averages, means, 99th Percentile, and 95th Percentiles.

Task Data
Task data provides a view of major events and scheduled services in the environment such as resources starting and stopping, VMware 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 virtual appliance Refresh, HPE 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. How long is data kept?

HPE CloudPhysics will retain all metadata indefinitely for aggregated and anonymous statistical and historical trending analysis. This metadata is used to help compare users to the global data set to identify inefficiencies. Machine metadata can be deleted at the discretion of HPE CloudPhysics.

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 HPE CloudPhysics servers for data processing inside of Google Cloud. The most recent data collections will be held in the virtual appliance until they can be delivered to the HPE CloudPhysics cloud. HPE CloudPhysics stores each customer's data in dedicated logical containers until the data can be queued, verified, and loaded into the HPE CloudPhysics data lake for analysis.

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

All data sent to HPE CloudPhysics is compressed and encrypted before sending to the cloud. All communications to the cloud are secured with an TLS connection on Port 443. Authentication data is one-way hashed at rest. Access is controlled via Google Cloud IAM, SSH public key auth, and firewalls. The current release of the HPE CloudPhysics appliance utilizes TLS v1.2.

Q. Is any personal identifiable information collected?

HPE CloudPhysics collects data center configuration and performance data. As a result, there is minimal exposure of personal identifiable information collected. HPE 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. Any data placed in VMware vCenter TAGS, Host names, or VM name will not be removed.

Q. How is my data secured in the cloud?

Unique pseudo-random hashed identifiers are created to represent each organization. The metadata is processed/stored in Google Cloud in a multi-tenant structure. Google Cloud data security services are detailed at https://cloud.google.com/security. Strict firewall rules, role based access, and two-factor authentication are used to limit access to customer data by HPE CloudPhysics.

Q. Who has access to my data?

Authorized HPE CloudPhysics employees, the customer and anyone external authorized by the customer through invitations or assessments.

Q. How does HPE CloudPhysics mitigate vulnerabilities and data risk?

HPE CloudPhysics monitors security vulnerability disclosures and maintains our environment accordingly. Full disk encryption is mandated for all employee computers and mobile devices that have access to HPE CloudPhysics documents/data. Administrative access is via two-factor authentication (SSH public key with password-protected keys).

Guest Process and Application Discovery Questions (Optional)

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 HPE CloudPhysics virtual appliance installation and configuratio By not provising a Guest OS login credential in the virtual appliance, you can skip the the guest process and dependency map collection process.

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 HPE CloudPhysics virtual appliance. 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 HPE CloudPhysics virtual appliance 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 HPE CloudPhysics observer during the virtual appliance 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 HPE CloudPhysics virtual appliance 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 guest OS network analysis tool called NetStat or similar command. HPE 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.

A detailed Dependnecy Map FAQ can be found here:

https://www.cloudphysics.com/dependency-map-faq

Q. How frequently is my Guest OS data collected?

HPE CloudPhysics will collect both guest process and network dependency data independently on a defined schedule. Initial releases will collect guest OS data every six hours to reduce VMware vCenter utilization.

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?

HPE 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. HPE 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 HPE CloudPhysics observer.

Q. How long is my data kept?

All data will be retained indefinitely for a user account to allow for historical analysis by the users. This data will remain part of your account as long as your account is active and will not be deleted until a deletion request is received.

Q. Can I remove or delete my data?

HPE CloudPhysics keeps all anonymized metadata for global comparison of performance, configurations, and is used to compare users against the global dataset. Organization data and account identifiable data cen be deleted upon request. Any identifiable data is maintained in your account only.

Data and account deletion can be initiated here:
https://www.cloudphysics.com/deletion-request/