Hewlett Packard Enterprise

Check if the document is available in the language of your choice.

HPE CLOUDPHYSICS SECURITY

Frequently asked questions

CONTENTS

INTRODUCTION

The extensive security expertise behind the HPE CloudPhysics platform from Hewlett Packard Enterprise is focused on protecting customer data at all points of the data transactions. This document answers some of the most commonly asked security questions. For a more extensive discussion of HPE CloudPhysics security procedures, contact Cloudphysicssupport@hpe.com.

BASIC HPE CLOUDPHYSICS QUESTIONS

What is the HPE CloudPhysics Observer?

The HPE CloudPhysics platform collects data from your environment by using a virtual appliance called the HPE CloudPhysics Observer. The Observer is a minimum resource appliance designed to collect data from your VMware vCenter® instance through read-only APIs. It processes the data and shares it with HPE CloudPhysics through secure means.

At the discretion of the VMware vSphere® administrator, additional levels of data collection can be made available with elevated privileges for guest process discovery by using VMware Tools™ APIs and limited guest credentials.

What are the system requirements of the HPE CloudPhysics Observer?

The virtual appliance requires 8 GB of RAM, 2 virtual CPUs, and 20 GB of disk space when deployed. Total network traffic resources total approximately 5 MB per hour per 100 VMs in the data center.

The HPE CloudPhysics Observer also requires internet access to send collected data to the public cloud through an encrypted connection to the internet domain at entanglement.cloudphysics.com. This domain is used for API calls only and has no publicly accessible web pages.

These communications 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.

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

HPE CloudPhysics does not deploy probes or agents to VMware ESXi™ hosts or to any guest OS. All communication is achieved through existing management interfaces, so minimal load is placed on the host environment. If VMware Tools is already deployed, HPE CloudPhysics can take advantage of it to collect process details within a guest OS, but VMware Tools is not required for infrastructure data collection.

How is data collected?

The HPE CloudPhysics Observer uses public APIs to collect data from VMware vCenter. HPE CloudPhysics requires a read-only account, with access to list and read configurations of the virtual environment. It collects performance and configuration data and other metadata from the vCenter instance on a defined schedule. vCenter collects performance and configuration data natively from its managed resources at a 20-second granularity. This data is typically rolled up and destroyed by vCenter after it is an hour old. Before that happens, HPE CloudPhysics collects the performance and configuration data directly from vCenter frequently enough to maintain the 20-second granularity. This data collection process is agentless and has no impact on the VMs or hosts being analyzed because it already exists in vCenter.

For information about the credentials required for VMware vCenter, see Installing the HPE CloudPhysics Virtual Appliance.

What type of data is collected?

HPE CloudPhysics collects several types of data:

- **Infrastructure configuration data** describes the virtual data center that is 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 the resources managed by vCenter. This data does not include network topology or data to recreate the network architecture. For VMware vCenter v4.1 and above, this data consists of details about the vCenter instance, the data center, the VM, the host, the virtual domain, the datastore, the network port, the virtual network, and the resource group.
- **Performance data** consists of CPU, storage, network, and RAM usage details. Utilization, peak performance, bandwidth, and characteristics of these components all make up the performance data. HPE CloudPhysics also generates derivatives of this data for averages, for means, and for 99th and 95th percentiles.
- **Task data** provides a view of major events and scheduled services in the environment, such as resources starting and stopping, vMotion actions, and environmental changes. These events and tasks often include the event and a brief description.
- **Metadata and tags** are part of the metadata that many resources include to describe a service and its role, and to provide context for its relationship with other objects in the environment. The most common type of metadata collected is the tags that are used for managing objects in the environment and for classifying and organizing resources, data, and services.

- **Inventories of running processes** within VMware® VMs can optionally be collected by administrations. This data collection is achieved by a guest request through VMware Tools. The request allows VMware Tools to return a list of processes that are currently running on the host to help classify applications and services associated with VMs.

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. The most recent data collections are held in the Observer 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.

How is data protected in transit to the cloud?

HPE CloudPhysics communicates over TLS 1.2 for current Observers on port 443 from the HPE CloudPhysics Observer to HPE CloudPhysics. Communications to Amazon Web Services occur through secure REST API communications over HTTPS (TLS) on port 443. All communications are encrypted by using the latest supported secure standards for data communications.

Can the HPE CloudPhysics Observer operate through a network proxy?

Yes, the HPE CloudPhysics Observer supports proxies that implement the HTTP Proxy protocol. It supports both unauthenticated and authenticated proxies that use basic, digest, or NTLM authentication. SOCKS and transparent (intercepting) proxies are not supported. When connecting through a proxy, the Observer uses the HTTP CONNECT method to connect directly (and exclusively) to entanglement.cloudphysics.com on port 443. HPE CloudPhysics does not support loading a custom TLS Certificate Authority, and the appliance will fail to function if the proxy attempts to intercept TLS traffic.

Is any identifiable personal information collected?

Because HPE CloudPhysics collects data center configuration and performance data, there is minimal exposure of collected personally identifiable information. HPE CloudPhysics collects user information only for portal account access and for invitations to new users from existing users. This data consists of company, name, and email address only. This data is used by the organization for user account management and credentialing.

How large an environment does HPE CloudPhysics support?

HPE CloudPhysics is not limited by the number of VMs, hosts, or servers. The current cloud model allows one HPE CloudPhysics Observer per vCenter to enable scalability. Data sent to the cloud is queued for processing, and the environment scales dynamically to accommodate capacity.

How is the Observer secured, and how often is it updated?

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

How does HPE CloudPhysics monitor availability and integrity of hosts within the customer's environment?

HPE CloudPhysics does not monitor the availability of systems in customer organizations beyond the most recent communication between the HPE CloudPhysics Observer and HPE CloudPhysics cloud services. It uses internal and third-party services to monitor availability and functioning of hosts within its infrastructure.

What credentials are required for HPE CloudPhysics to access the vCenter instance?

HPE CloudPhysics requires a limited-access account that has read and list capabilities against VMware vCenter. For more information about security and policy requirements for vCenter, see Installing the HPE CloudPhysics Virtual Appliance.

What connectivity and protocols does the Observer use?

HPE CloudPhysics communicates over TLS 1.2 for current Observers on port 443 from the HPE CloudPhysics Observer to HPE CloudPhysics. All communications are encrypted by using the latest supported secure standards for data communications.

GUEST PROCESS AND APPLICATION DISCOVERY QUESTIONS

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 installation and configuration of the HPE CloudPhysics Observer.

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 Observer. If no credentials are provided, the collection process will not be executed.

How are the guest processes collected?

A feature of VMware Tools is used to collect guest processes. The request to VMware vCenter originates from the HPE CloudPhysics Observer. When it receives the request, vCenter tries to issue the command to the VMware Tools feature within the guest OS. VMware vCenter initiates a process collect command under the identity of the guest account that was specified in the HPE CloudPhysics Observer when it was set up.

VMware Tools issues 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. After the collection is complete, the HPE CloudPhysics Observer collects the output of the command execution. The HPE CloudPhysics Observer uses a vSphere API, which in turn uses VMware Tools to collect the command output that is temporarily stored in the output file.

Assuming that the user has access to their own home directory, the application data is written to the user's home directory and removed during data collection. If the user does not have sufficient rights to delete their temp files, the file is overwritten with each collection to keep the volume storage minimal.

How is dependency mapping data collected?

Dependency mapping is derived from a network analysis tool called **NetStat**. HPE CloudPhysics issues a request to VMware vCenter for details from the guest OS. vCenter can direct queries to the guest OS if VMware Tools is deployed and enabled. The request is 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 collects all open network communications and reports source and destination IP addresses, TCP/UDP, and port. This data is directed into a local temp storage file, where it is processed and sent to VMware vCenter by VMware Tools.

How frequently is my data collected?

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

How do you create a dependency map?

Dependency maps are generated based on source and destination IP addresses and on the ports identified by NetStat during the online dependency mapping data analysis. This data identifies all major network communications by the guest OS and map IP addresses to other VMs. If any VM communicates outside of the private network ranges, those communications are considered to be outside of your data center.

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 administrator or to have root access within a guest OS. For mixed environments, confirm that the same user ID and password exist in both Linux® and Windows environments.

What data is collected for the guest process?

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

What data is collected for dependency mapping?

NetStat returns a text output of source, destination, and port, and potentially of protocol information from the guest. This data varies slightly from OS to OS; typically, however, additional data might include packet count, state, or world ID.

What is the data flow during collection?

HPE CloudPhysics issues a request for data to VMware VCenter for a specific guest OS. VMware vCenter issues the credential and command to the guest OS. If the command is allowed to execute, VMware Tools directs all output from the command to a temp file in a guest user home directory. Upon completion of the command, VMware Tools retrieves the temp file and directs the output back to VMware vCenter as a temporary variable for the guest OS. HPE CloudPhysics then collects the temp variable from 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.

Who has access to the guest process and dependency mapping data?

All users with access to your HPE CloudPhysics account can use analytics that derive data from the data collection process. The processes are used as tags in the HPE CloudPhysics environment to enable quick classification of applications and guest OS instances. NetStat dependency mapping data is available only through dependency mapping cards that are enabled to account users.

Can I remove or delete my data?

All account data can be removed by sending a request to cloudphysicssupport@hpe.com. HPE CloudPhysics keeps all anonymized metadata for global comparison of performance and configurations. The metadata is used to compare users against the global dataset.

How do I access my guest process data in the HPE CloudPhysics portal?

HPE CloudPhysics makes all guest process data available today in Card Builder for the VM object called **guest processes**. In addition, some guest processes are used to generate tags or events for some analytics. For example, the Microsoft SQL Server process is used to identify guest operating systems that have SQL databases installed, and it can be used to automatically generate tags associated with these VMs.

Frequently asked questions

Resources, contacts, or additional links

Cloudphysics@hpe.com

Cloudphysicssupport@hpe.com

entanglement.cloudphysics.com

Installing the HPE CloudPhysics Virtual Appliance

LEARN MORE AT

hpe.com/storage

Make the right purchase decision.
Contact our presales specialists.

Get updates

© Copyright 2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. SQL Server and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware, VMware ESXi, VMware Tools, VMware vCenter, and VMware vSphere, are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.

a00117306ENW, September 2021

Hewlett Packard Enterprise