Cloudamize

Cloudamize Agents FAQ

Cloudamize is a cloud infrastructure analytics platform that provides data analysis and recommendations to speed and simplify cloud migration and management. Our platform helps you choose your best-fit cloud vendor; automates application discovery and dependency mapping to design a precise migration plan; analyzes your performance metrics and usage patterns on an ongoing basis to ensure your cloud is always right-sized; and provides clear visibility into cloud costs for better control. Armed with these insights, you can more quickly make accurate cloud decisions, achieve your optimal cloud environment, and maximize the value of your cloud investments.

Q: What are Cloudamize Agents?

A: Cloudamize Agents automatically collect on-premises and in-cloud system and application level data that is essential to accurately projecting cloud costs, designing migration plans based on detailed application dependency mapping, and identifying the best-fit cloud configuration for right-sizing your cloud. Because Cloudamize Agents are purpose-built for moving to and operating in the cloud, they collect and provide only data relevant to the cloud journey. They capture application inventory, dependencies, and performance metrics, and send this data in an encrypted format to the Cloudamize server, where it is stored in an encrypted data repository.

Q: What data do Cloudamize Agents capture?

A: The agents capture infrastructure and application level information from hosts, including inventory of installed software applications, system and process performance, resource utilization, and network dependencies between workloads. Cloudamize Agents only gather system and application level data pertaining to discovery, dependencies, and performance, and never collect sensitive or personally identifiable information PII from the servers they are installed on. See here for a detailed list of metrics that the agents collect: http://www.cloudamize.com/hubfs/Data_Collected_by_Cloudamize_Agents.pdf.

Q: Do Cloudamize Agents require a reboot?

A: No, Cloudamize Agents do not require you to reboot your machines.

Q: How do I pause/stop data collection?

A: Cloudamize portal allows you to stop all agents with one click or stop individual agents as needed. Additionally, Cloudamize will automatically stop data collection at the completion of the assessment. You will be further required to uninstall Cloudamize Agents upon conclusion of assessment.

Q: How do Cloudamize Agents compare to agentless solutions from other vendors?

A: Here are the following ways Cloudamize Agents compare to agentless solutions from other vendors:

Security Protection

Agentless solutions run a service on the endpoints that listens for incoming connections. The service must listen on ports and open inbound connections through a firewall (10+ inbound ports per endpoint), which poses a security risk. The remote data collector must be allowed to communicate with the endpoints on different inbound ports, and it may also need to be installed with domain administration privileges to be able to access the endpoints.

Cloudamize Agents are much more secure than agentless monitoring. The agent to application/OS communications are handled via outbound SSL or via internal proxy. Hence, no additional firewall rules need to be configured. Also, Cloudamize Agents run their processes on the endpoint that initiates a connection to the Cloudamize server. Therefore, the security burden of listening on the port is shifted to the Cloudamize server side rather than the endpoint. All of the data Cloudamize collects is encrypted and protected by Secure Sockets Layer (SSL) and Transport Layer Security (TLS) and sent to the Cloudamize server that is hosted in public cloud (AWS, Azure, GCP).

Cloudamize Agents Architecture

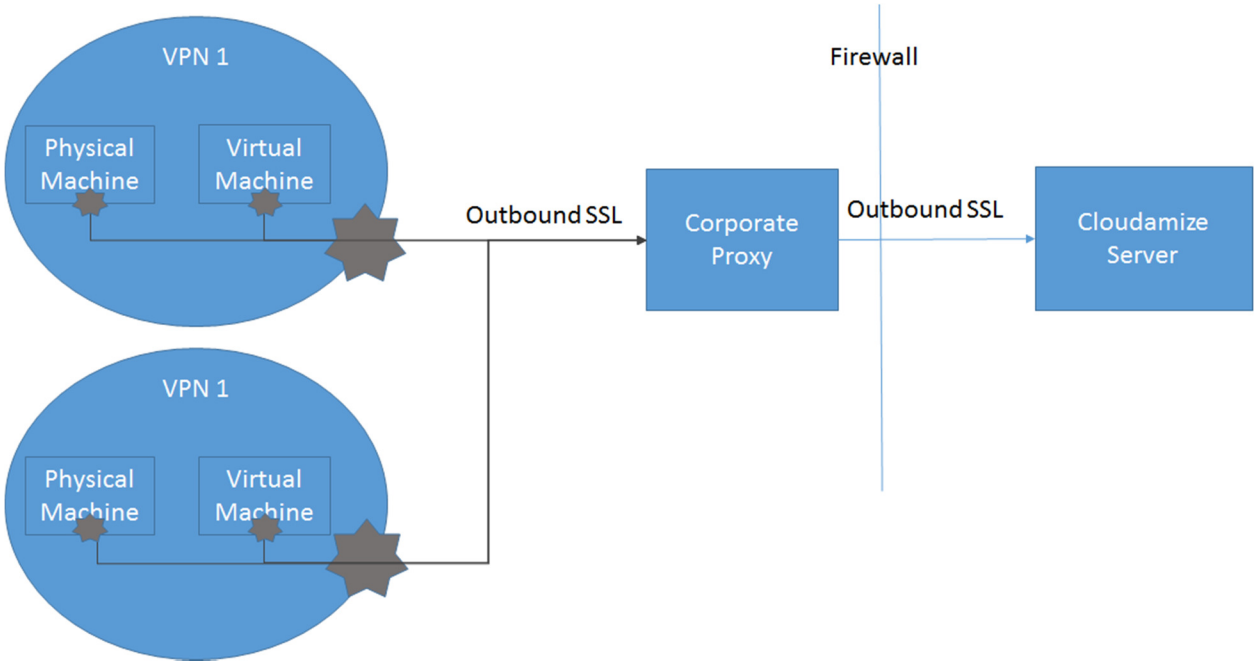

Agent-less Architecture (from other vendors)

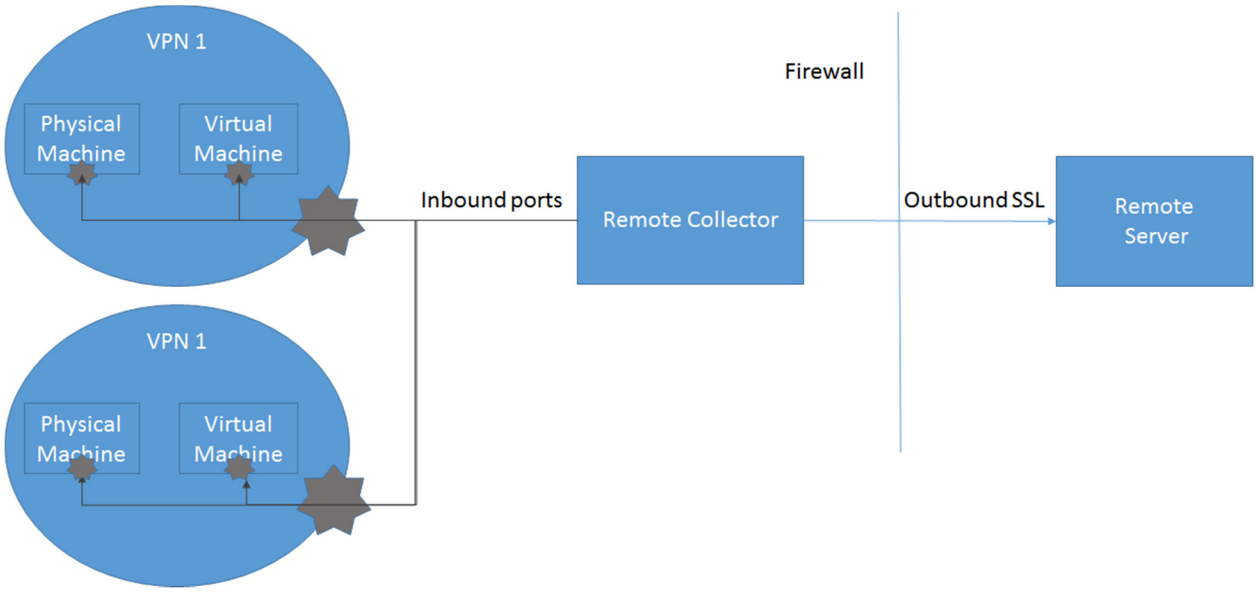

Ease of Deployment

Many believe agentless solutions are easier to deploy than agents, but in reality, they are not. Agentless solutions require SNMP/WMI to collect data, and because many endpoints do not support SNMP/WMI, it takes a considerable amount of time to install and configure the agentless service on each endpoint. Also, because many inbound ports must be opened for each endpoint for every VLAN/Subnet (10+ per end point), it takes significant time to gain security approval and implement firewall changes.

Cloudamize Agents are very easy and quick to deploy. Our customers typically deploy the agents to servers globally. Depending on the size of the infrastructure, it can take a few minutes to few hours, and customers receive all data within just two weeks. Only a single outbound port (TCP 443) to a specific server is required. Generally this port is open and no additional firewall changes are required. Also, the Agent traffic can be routed via authenticated corporate proxy and no firewall rule change is required in that case - only an outbound SSL from proxy to a server is required and that is generally open.

The following tools can be used for fast deployment of Cloudamize Agents:

- Microsoft system center
- Active Directory
- vCenter
- Puppet
- Chef

Resource Consumption

Cloudamize Agents are lightweight, typically using less than 0.5% CPU utilization. Also, Cloudamize Agents run two processes: data collection and watchdog. The watchdog process monitors every system metric and halts data collection if the CPU utilization goes beyond 2%. Data collection will not resume again until the watchdog process tells it to. Cloudamize agents will therefore never bog down your

network or disrupt your day to day operations. Also, Cloudamize Agents encrypt and compress all data on the host before sending it, and as a result use minimal resources.

Agentless solutions are very network-intensive. They rely on looking at the network and consume significant CPU and memory. Because agentless solutions use standard network protocols, they cannot compress the data before sending it, thereby bogging down the network.

Accuracy

Right-Sizing Your Cloud

Cloudamize Agents collect in-depth performance metrics (such as peak CPU utilization, allocated and peak RAM usage, storage capacity and current occupancy, disc IOPS and bandwidth, and throughput) and usage patterns to find best cloud configuration for right-sizing each of your workloads. These metrics are collected at high frequency - in 30 second intervals and include peaks and valleys, NOT averages.

Why is this important?

Right-sizing your cloud, which includes choosing your optimal cloud vendor, instance types, storage options, and required capacity, cannot be done without these highly precise analytics. Agentless solutions simply cannot collect this level of detail. First, they cannot measure memory, so choosing the right instance is almost impossible. Second, they only look at hypervisor level data rather than specific instances. Their recommendations around improving cost and performance will be off because they're based on what hypervisor is allocating to the instance rather than what the instance is actually using. Third, agentless solutions do not measure CPU metrics at high frequency, and they average the metrics they do collect. If you size your cloud environment based on averages, your infrastructure will suffer serious performance degradation when you hit peaks, and you will incur unnecessary costs during slow periods.

Application Discovery and Dependency Mapping

Cloudamize Agents can see all commands and web server requests, allowing them to automatically identify every application in your infrastructure. On the other hand, agentless solutions cannot read UDP protocol - only TCP. Because they identify applications by looking at protocols, they miss every application using UDP, which can be up to half of all your applications.

Agents collect detailed application dependencies, including the other applications and servers each application is communicating to, how often they communicate, whether the communication is bi-directional or unidirectional, and what the allowable latency is in communications between each. This guides your migration process, including which applications you will decide to migrate and the best order in which to move and pilot them.

Agentless solutions, however, do not collect application dependencies. There are two reasons for this. First, 80% of application and server connections are bursty/short-lived, and because agentless solutions cannot collect data at high frequencies, they miss these dependencies. Second, because agentless solutions cannot read UDP protocol, they not only miss these applications, but also their connections/dependencies.

Why is this important?

Without identifying all of your applications and their dependencies, you will inevitably group applications incorrectly, which will triple migration cycles, increase the number of applications that break during migration, and compound de-bugging issues once in the cloud. The data that agents provide significantly reduce migration time by avoiding massive trial and error.

Summary

Overall, agent-based solutions get you to your goal of realizing the benefits of the cloud faster: they get you your data more quickly, and they provide in-depth metrics collected at high frequency to ensure you have all the data you need for accurately migrating applications and right-sizing your cloud.

Benefits	**Cloudamize Agent-Based**	**Agentless (Other Vendors)**
Security Protection	High: no new firewalls are opened.	Low: firewalls are opened.
Ease of Deployment	Easy: only a single outbound port is required and it's usually open. Can also use corporate proxy. No firewall changes.	Hard: many inbound ports need to be opened and requires firewall changes. Also requires configuration on each endpoint.
Resource Consumption	Low: data is collected locally and compressed data is sent over the internet. Data collection is halted if exceeds allocated CPU.	High: introduces additional network traffic as the raw performance data is transported over the network to a remote data collector.
Accuracy	High: high frequency collection of both CPU and memory metrics. Discovers all applications and their dependencies in detail.	Medium to Low: Only CPU averages are collected. Not all applications and dependencies are identified due to protocol limitations and low frequency data collection.
Scalability	High: server can handle greater amount of concurrent processes.	Low: server has a limit in terms of how many connections it can concurrently handle.
Robustness	High: process initiation and stream handling is done by agents. Caching mechanism is available in the event of network disruption and hence robust solution.	Low: process initiation and stream handling is done on the server. Affected by networking issues, e.g. if the network connection is unstable, it affects the data collection.