ARIS Process Performance Manager
CLOUD INFRASTRUCTURE

Version 10.2

April 2018

Contents

1   Text conventions

Menu items, file names, etc. are indicated in texts as follows:

- Menu items, key combinations, dialogs, file names, entries, etc. are displayed in **bold**.

- User-defined entries are shown in **<bold and in angle brackets>**.

- Single-line example texts (for example, a long directory path that covers several lines) are separated by ↵ at the end of the line.

- File extracts are shown in this font format:

  ```
  This paragraph contains a file extract.
  ```

2   General

This document provides a quick overview of the installation and operation of ARIS Process Performance Manager (PPM) in a cloud environment. A cloud environment can be hosted by Amazon EC2 (Elastic Cloud Version 2), Microsoft Azure or any other cloud provider that supports the operating systems needed for PPM. The document guides you in configuring and operating PPM in a cloud environment or cloud instance.

3 PPM Infrastructure

This chapter describes the basic infrastructure components of PPM. A detailed description is provided in the document **PPM System architecture**.

3.1 Network infrastructure

The following figure shows all relevant components in the PPM infrastructure. It is divided into client and server components. The server components can be accessed by the load balancer using HTTP or HTTPS protocol only.
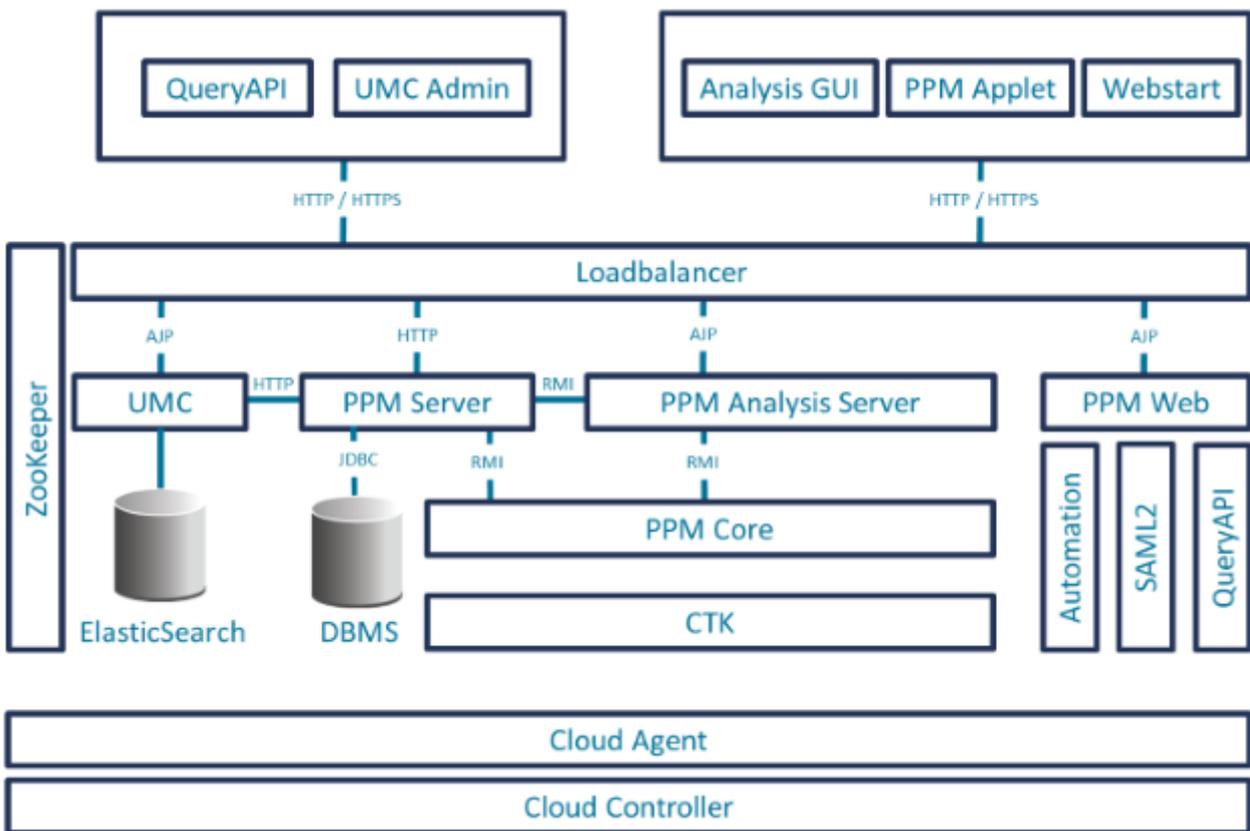
3.2 Connection parameters

You have three options to access an PPM server installation.

▪ Web browser (UMC Admin & PPM Applet)

▪ Java Web start (PPM Web start)

▪ Java application (PPM Analysis GUI)

Those client components connect to PPM only using HTTP or HTTPS via load balancer. Different server components, for example, UMC, PPM Web or PPM server are resolved using different load balancer contexts.

You can access the PPM system using the load balancer base URL:

▪ http://<cloud instance IP/FQDN>:4080/ or

- https://<cloud instance IP/FQDN>:4443/

The following table shows the connection parameters for the different server components.

Component	Server URL
UMC	http://<cloud instance IP/FQDN>:4080/umc https://<cloud instance IP/FQDN>:4443/umc
PPM Server (via Applet)	http://<cloud instance IP/FQDN>:4080/ppm/html https://<cloud instance IP/FQDN>:4443/ppm/html
PPM Query API (old version)	http://<cloud instance IP/FQDN>:4080/API_<client name>/urlquery/… https://<cloud instance IP/FQDN>:4443/API_<client name>/urlquery/…
PPM Query API (rest api)	http://<cloud instance IP/FQDN>:4080/ppmserver/API_<client name>/… https://<cloud instance IP/FQDN>:4443/ppmserver/API_<client name>/…
PPM Server (via Webstart)	http://<cloud instance IP/FQDN>:4080/ppm/html/ppm.jnlp https://<cloud instance IP/FQDN>:4443/ppm/html/ppm.jnlp
PPM Documentation	http://<cloud instance IP/FQDN>:4080/ppm/html/help/ppm/en/overview/index.htm https://<cloud instance IP/FQDN>:4443/ppm/html/help/ppm/en/overview/index.htm
PPM Online Help	http://<cloud instance IP/FQDN>:4080/ppm/html/help/ppm/en/handling/index.htm https://<cloud instance IP/FQDN>:4443/ppm/html/help/ppm/en/handling/index.htm

The IP address or fully qualified domain name (FQDN) is the unique entry point to gain access to a PPM system running in a cloud environment. This IP address is required to install PPM on a cloud instance (for details see chapter Installation (Page 11)).

All internal PPM components (for example, PPM registries, PPM servers, PPM Web, etc.) register with their internal host name at the Apache ZooKeeper. Other components that need access to those internal PPM components will request the proper address from the Apache ZooKeeper. Apache ZooKeeper centrally manages all relevant host name/IP configurations in the PPM system. Host name properties in all relevant settings files (for example, RMIServerURL or CorbaServerURL) were removed and placed in the runnable **ppm_core** and can be changed using standard Cloud Controller commands. It is not necessary to change multiple settings files to change the IP or host name of the PPM system.

After a change of host name or IP address issued in Cloud Controller, all PPM components must be restarted.

4 Cloud infrastructure

4.1 Overview

Cloud computing is computing in the form of a service rather than a product. Shared resources, software, and data are provided to computers and other devices via a network (typically the Internet) that basically acts like a utility (like the electricity grid). The cloud provides computation, software, data access, and storage services that do not require end-user knowledge of the physical location and configuration of the system that delivers the services. This concept meets a wide-spread need of IT: a way to increase capacity or add capabilities on the fly without investing in new infrastructure, training new personnel, or licensing new software. Cloud computing encompasses any subscription-based or pay-per-use service that extends IT's existing capabilities in real time using the Internet. The services provided by cloud computing can be categorized in three different models:

- **IaaS** – Infrastructure as a Service: The user can access and use virtual hardware resources such as desktop, server, network or memory. Choosing IaaS, users create their own cloud-based computer cluster and thereby account for the selection, installation, and operation of the software themselves.

- **PaaS** – Platform as a Service: Cloud instances provide net access for programming or runtime environments using flexible and customizable computing and data capacities. PaaS can be used to develop your own software applications or execute them in a predefined environment provided and maintained by your service provider.

- **SaaS** – Software as a Service – cloud instances provide net access to software applications offered by your service provider. These applications are operated in predefined and ready-to-use environments or infrastructures. SaaS is sometimes also described as "software on demand".

This document only relates to SaaS, that is, providing ARIS services in a cloud instance with access to these services via the Internet from anywhere. For this solution, we distinguish between three provisioning models:

- **Public cloud** - Offers access to abstract IT infrastructures for the public Internet. Public-cloud service providers allow their customers to use or rent these infrastructures on a monetary basis with users only having to pay for actual usage (pay-as-you-go), without investing money in computing or data center infrastructures.

- **Private cloud** - Offers access to abstract IT infrastructures within your own organization. The connection between the cloud instance and your home office network can be established via VPN (VPC at Amazon).

- **Local Cloud** - Offers access to abstract IT infrastructures on your local computer. Most likely you only receive a predefined image or instance of the software for download and you need to run it locally on a virtualization platform (for example, VMWare or Virtual Box) hosted on your computer. The connection to that image or instance can be established using local network interfaces on your computer or the hosting application you are running the image or instance on.

The following sections provide more details on the two different concepts and all necessary configurations for PPM.

4.2   Public cloud

This chapter deals with the SaaS service model using a public-cloud infrastructure. It describes the network infrastructure and configurations required to operate the relevant applications in that environment.

4.2.1   Network infrastructure

The following figure illustrates a rough overview of the network structure of a public-cloud instance. Usually, the client is hidden behind a firewall and proxy, which can be established by a company or at home within the private network. All cloud instances have two IP addresses: a private one, which is only visible in the cloud network itself and a public IP, which is visible on the Internet. The latter is normally not published in the runtime environment of the cloud instance itself, which is commonly called Network Address Translation (NAT). Moreover, the public IP address is not fixed, it changes every time the instance is started or rebooted.
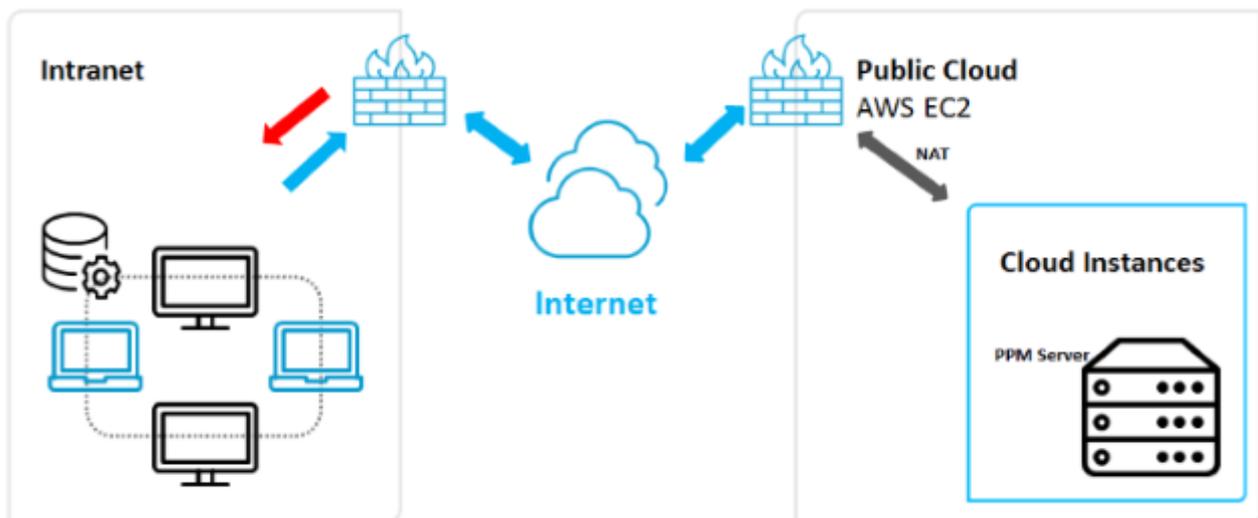
The changing public IP address of the cloud instance may cause trouble during the configuration of software applications. There are two way of avoiding this:

▪ Provide an easy-to-use script or software that automatically configures the application correctly using the new IP address. Depending on the application, this approach may not be necessary. However, it is always associated with the need to reconfigure the client's user access to that system because the URL for the applications changes with the IP address.

▪ Use another service provided by Amazon called "elastic IP". This is a kind of static IP address provided by Amazon. You can choose from a bundle of IP addresses associated with your Amazon account and then assign the IP address you selected to your cloud instance every time you have to start or reboot it. This service is not free as the IP addresses have to be reserved for your account. It is payable in addition to any other service costs your instance incurs. For the installation of PPM, an "elastic IP" is required.

4.2.2   Connection Parameters

If you connect to a PPM system in a public cloud, you only need the base URL for the load balancer containing the public IP address or fully qualified domain name (FQDN) of the cloud instance. See chapter Public cloud (Page 6) for details on URL. In most cases, default ports of PPM (HTTP 4080) are blocked by the cloud provider's firewall. In this case, you need to reconfigure the ports after installation. For more details, see chapter Reconfigure host and port parameters (Page 13).

4.3   Private cloud

This chapter deals with the SaaS service model using a private-cloud infrastructure. It describes the network infrastructure and configurations required to operate the relevant applications in that environment.

4.3.1   Network infrastructure

The following figure provides a rough overview of the network structure of a private-cloud instance. The private-cloud instance (or more than one) is embedded in the local area network of your company. It may also be considered a DMZ for your local network.
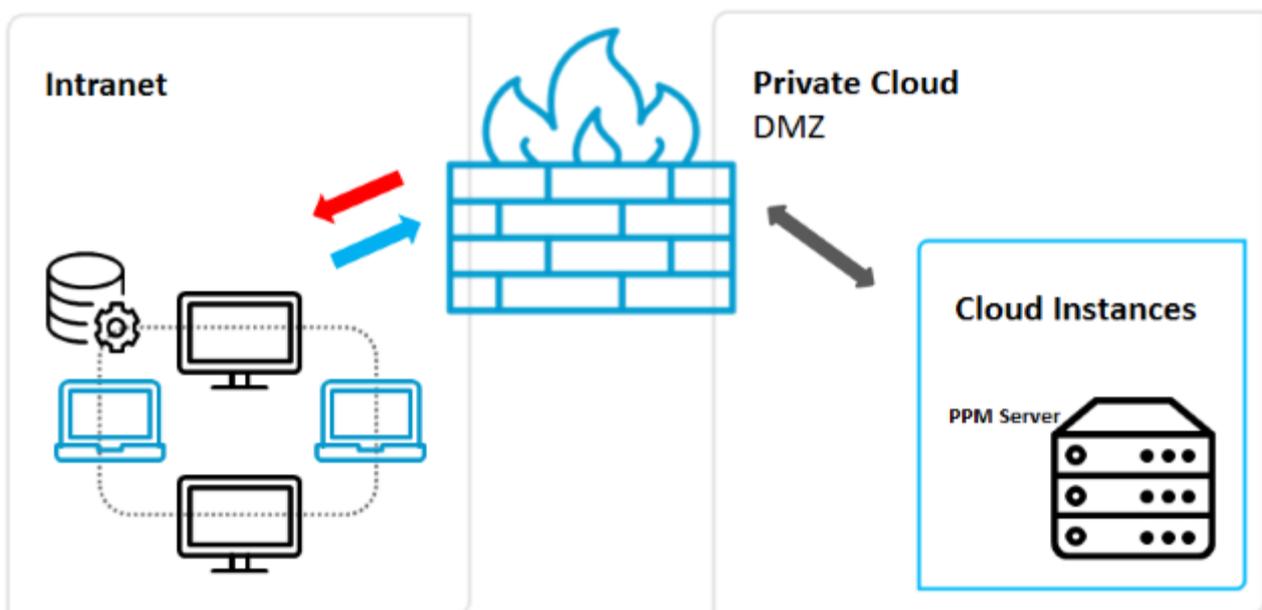
Depending on the solution you want to apply, various characteristics of each application need to be considered. They are described in the following sections. Please note that the private-cloud scenario requires much more effort in configuring your network or all routers involved, as both traffic directions (outbound and inbound) need to be covered. This scenario needs to be embedded by network administrators in order to create and operate a secure network tunnel between your cloud instance and your LAN, in which some resources of your LAN are also accessible through the cloud instance.

4.3.2   Connection Parameters

Connecting to a PPM system in a private cloud is no different from connecting to a public one. You still need the base URL for the load balancer containing the IP address or fully qualified domain name (FQDN) of the cloud instance. See chapter Public cloud (Page 6) for details on URL. As private cloud access to the instance might be restricted to local networks only, access from the internet is usually not possible. Furthermore, default ports used during installation (HTTP 4080) can be used, since these ports are not usually blocked in private cloud scenarios.

4.4   Local cloud

This chapter deals with cloud infrastructures running on your local machine using preconfigured images or instances prepared for virtual environment applications like VMWare or Virtual Box. It describes the network infrastructure and configurations needed to operate PPM in such an environment.

4.4.1   Network Infrastructure

The following figure provides a rough overview of the network structure of a virtual machine running on a local computer or server. The hosting application (for example, VMWare Player or Virtual Box Manager) provides a virtualization platform to run cloud images or instances on your local machine. These images or instances can be provided as ready-to-use appliances or you can create your own images. These images are containers for a complete independent operating system also incorporating the application you want to run on that operating system. The hosting application provides access to physical resources of your host like network interfaces, memory and hard drive capacity. Access to the containers is possible using specially configured network interfaces and the hosting application itself, which is illustrated in the following figure.

Intranet
Host or any other local
client

VMWare Player
Locally installed guest system

VM Instances

PPM Server
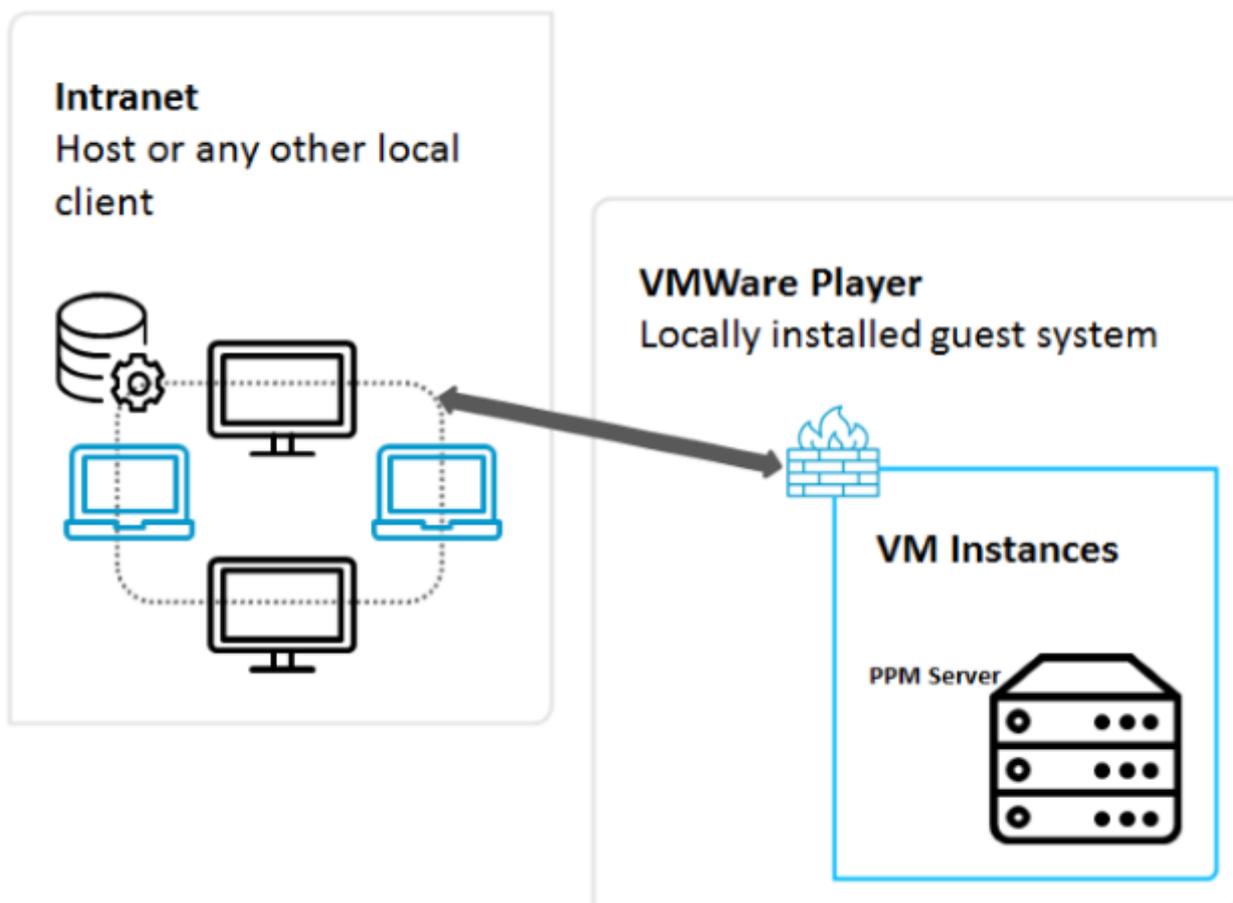
4.4.2   Connection parameters

Connection parameters for this scenario depend on the network configuration of the used virtual machine/cloud instance (VM) on the hosting application. Basically, you have up to four ways of configuring access to the physical network resources on the host:

▪ Internal networking (intranet)

▪ This configuration provides a totally isolated network, where only the client VMs can communicate with each other over a network interface. Even the host is not part of this network. Client VM's IP addresses need to be configured statically.

- Host only networking (Host only)

  Similar to the internal networking, but the host provides some convenience services like DHCP, which means that you do not have to configure the client's IP addresses statically. The host is also part of the network.

- Indirect access to the intranet (NAT)

  This network type is only used for client applications used in the VM, that is, no services are provided within the VM that need to be called externally.

- Direct access to the intranet (Bridged)

  With a bridged network configuration, your VM is integrated in the local network infrastructure and can provide services with its own IP address and network name

Installing PPM on such a VM image, you need to know the used network configuration of the image. In most cases, "Host Only" is sufficient and you can use the IP address or given name (FQDN) of the image during installation (see chapter Local installation (Page 11) for details). This information is provided by the host to the image.

The base URL of the load balancer contains this IP address or name. See chapter Public cloud (Page 6) for details on URL.

5   Installation

5.1   Installation requirements

You have two options to install PPM on a cloud instance:

- Local installation on the cloud instance
- Remote installation on the cloud instance (currently only supported for Linux)

For both options, you must have access to the cloud instance in order to copy the setup and application repository to the hard drive of the cloud instance.

REQUIREMENTS

- Using an Amazon EC2 cloud instance, configure external access with the ports **4443** and **4080** to access the default ports of the load balancer (see Connection Parameters (Page 7)). If it is not possible to reconfigure Amazon EC2 firewalls, you must use the default ports **80** (HTTP) and **443** (HTTPS) during further installation steps.
- Determine the correct FQDN or IP address of the cloud instance. The external default FQDN on Amazon EC2 is "ec2-<elastic-ip>.compute-1.amazonaws.com" with "-" instead of the dots in the IP address, for example, "ec2-111-222-33-44.compute-1.amazonaws.com" if the elastic IP address is 111.222.33.44. Is this IP address or the FQDN is not available, it can be retrieved in the EC2 cloud instance by using the following URL in a browser: http://instance-data/latest/meta-data/public-hostname.
- For public cloud instances, the public IP address will change every time the system is rebooted. If you do not want to purchase an elastic IP, you must reconfigure your PPM system every time the IP and hence the FQDN changes. See chapter Reconfigure host and port parameters (Page 13) for details.

5.2   Installation steps

5.2.1   Local installation

Copy the PPM setup folder to the cloud instance and any available PPM patch setup that is currently available to install PPM locally on a Windows operating system-based cloud instance.

- Start the PPM setup program and install PPM.
- Enter the given FQDN in the PPM setup on dialog **External IP Address**.
- By default, only HTTP is initially available. Choose a port of your choice in the corresponding dialog. Default port for HTTP connections is 4080.
- Install the currently available fix provided in the PPM patch setup.
- Use the FQDN during any further PPM usage (client creation, etc.). Do not use the internal or external IP address.

5.2.2　　　Remote installation

To install PPM on a Linux os-based cloud instance, you must copy the Cloud Agent RPM package to the cloud instance. RPM is a software packaging system initially provided by Red Hat (Red Hat Package Manager). This RPM system must be supported on the chosen Linux OS. The Cloud Agent can be installed using the following command:

▪　sudo rpm -i <cloud agent file>.rpm

▪　for example: sudo -i ppm10-cloud-agent-10.2.0.0.1210538-1.x86_64.rpm

After the Cloud Agent has been installed and executed, you must start the Windows setup for PPM on another Windows-based host. This can be another Windows-based cloud instance or your local Windows machine. Ensure that you have a valid network connection to the cloud instance. In the Windows setup, choose remote installation in the corresponding dialog and enter the URL of the Cloud Agent installed on the Linux system. The URL looks like this:

▪　http://<cloud instance IP/FQDN>

By default, only HTTP is initially available. Choose a port of your choice in the corresponding dialog. Default port for HTTP connections is 4080.

5.2.3　　　Distributed installation

To establish a Master-Subserver-System with PPM in a cloud scenario, you must install PPM separately on the selected cloud instances. For example, you must install a master server on one cloud instance and two subservers separately on two other cloud instances.

In such a complex scenario, it is not sufficient to provide access on the separate instances using the base URL of the load balancer. On the subserver instances, you also must open the RMI ports on which the master communicates with the subservers. Those are:

▪　RMI registry port of the subserver

▪　RMI object port of the subserver

The ports must be configured in the firewall of the master (outbound connection) and the subservers (inbound connection). Additionally, the RMI registry must be bound to the internal host name of the cloud instance. This internal host name can be retrieved using the following URL in a local browser:

http://instance-data/latest/meta-data/local-hostname

To configure this using Cloud Controller, the following commands must be issued on the subserver cloud instances:

▪　reconfigure ppm_core +ppmrmi.zookeeper.application.instance.host=<dedicated host>

▪　reconfigure ppm_core +ppmrmi.zookeeper.application.instance.port=<dedicated port>

In a public cloud scenario, this host name is usually different from the public host name, reachable from the outside world (internet).

6   Operations

The following sections describe important operation steps after PPM has been installed successfully in a cloud instance.

6.1   Security settings

By default, only HTTP protocol is configured during the PPM installation. HTTPS is disabled as you need a valid certificate for the load balancer. Perform the following steps to enable HTTPS.

- Create your own certificate for your server and have it signed by an official certification authority (CA). See PPM Operation Guide for details.

- Enhance the load balancer with this certificate. The certificate consists of two files (a key and a **crt** file) packed into a zip archive. The Cloud Controller command looks like this:
  enhance loadbalancer_x with sslCertificate local file "<path to ZIP file>"

- Finally enable HTTPS protocol and make it the default protocol that should be used:
  reconfigure loadbalancer_x +HTTPD.ssl.port=4443
  reconfigure loadbalancer_x +zookeeper.application.instance.port=4443
  reconfigure loadbalancer_x +zookeeper.application.instance.scheme=https

The given port is the default HTTPS port used in PPM installation and **loadbalancer_x** represents the load balancer ID.

6.2   Reconfigure host and port parameters

6.2.1   Change host parameters

Especially in public cloud scenarios, the public host name of the cloud instance may change after a reboot of the system or the public host name was not correctly provided during the installation process. In those cases, it is necessary to reconfigure that public host name in the load balancer. Otherwise, the PPM system will not work properly and at worst case not reachable from outside the cloud instance.

The public host name is stored in the load balancer parameters **HTTPD.servername** and **zookeeper.application.instance.host**.

The parameters can be displayed using the Cloud Controller command:

- show instance loadbalancer_x config

If required, you can change the public host name using the following command:

- reconfigure loadbalancer_x +HTTPD.servername=<new hostname>

- reconfigure loadbalancer_x +zookeeper.application.instance.host=<new hostname>

The values of these parameters must be identical.

6.2.2   Change port parameters

The public access ports for the PPM system are stored in the following load balancer parameters:

- HTTPD.port (HTTP Port)

- HTTPS.port (HTTPS Port)

These ports can be changed using the following Cloud Controller commands:

- reconfigure loadbalancer_x +HTTPD.port=<new HTTP-Port>

- reconfigure loadbalancer_x +HTTPD.ssl.port=<new HTTPS-Port>

6.3   Usage restrictions

Assuming a system set-up and operation as described above, all client features of ARIS Process Performance Manager are available. However, there are several server features that may cause problems with the underlying source systems:

- **LDAP**: If you are considering using an LDAP system for user authentication purposes in PPM, the LDAP system must be available via the Internet and directly accessible by your cloud instance. As this configuration is basically impossible in a public-cloud scenario due to security issues, we recommend to use the default PPM user administration. LDAP is not available in this scenario.

- **SAP-2-PPM extraction**: For this data extraction, an SAP system needs to be available either in the cloud or via a direct IP connection outside the cloud in your local network. Both scenarios are basically unable to handle cloud instances.

- **CSV-2-PPM extraction**: All CSV files must be available locally in the cloud instance, which only leaves you with the option to copy the source data into the cloud using FTP, sFTP, SCP, or any other transfer protocol.

- **JDBC-2-PPM extraction**: For this data extraction, any of the supported database systems must be available and accessible through the JDBC standard protocol. If the database is located in the cloud itself or on the same instance, this does not represent a problem. But accessing a database located in your secured local area network is basically impossible because you have to access the database directly via the Internet.

- **PPM data import**: All data either in plain XML or compressed ZIP format must be available locally on the PPM server for the import process. If you do not extract data from source systems also hosted in the cloud or another instance, you are left with the option of transferring the extracted data files to the cloud instance using FTP, sFTP, or SCP (assuming you run an SSH server in the cloud instance).