BigPrivacy®

Installation Guide

User Guide

Rev 20183008

Information contained herein is subject to change and is not guaranteed to be error-free.

Anonos has been actively engaged in research and development to advance the state of the art in global data protection, privacy and security technology since 2012. Anonos BigPrivacy systems and technology are protected by an intellectual property portfolio that includes, but is not limited to: (2018) SYSTEMS AND METHODS FOR ENHANCING DATA PROTECTION BY ANONOSIZING STRUCTURED AND UNSTRUCTURED DATA AND INCORPORATING MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE IN CLASSICAL AND QUANTUM COMPUTING ENVIRONMENTS – NO. 10,043,035; (2017) SYSTEMS AND METHODS FOR ANONOSIZING DATA – NO. 9,619,669; (2016) SYSTEMS AND METHODS FOR CONTEXTUALIZED DATA PROTECTION – NO. 9,361,481; (2015) DYNAMIC DE-IDENTIFICATION AND ANONYMITY – NO. 9,129,133; (2015) DYNAMIC DE-IDENTIFICATION AND ANONYMITY – NO. 9,087,216; (2015) DYNAMIC DE-IDENTIFICATION AND ANONYMITY – NO. 9,087,215; plus 60+ additional U.S. and international patent applications.

See anonos.com/patents for more information.

Contents

About the Installation Guide

This document is a comprehensive guide to installing and configuring the BigPrivacy® SaveYourData® computer software application (“SYD”). Notes on proper preparation and planning are included to help ensure that your installation runs smoothly.

SYD provides for connecting to source (ingest) databases, extracting data, and taking the initial step in transitioning “to GDPR compliance by basing data processing on a different lawful basis while ensuring that continued processing is fair and accounted for" by supporting "legitimate interests" as a lawful basis for ongoing processing.

Document Conventions

The following conventions are used throughout the Installation Guide to highlight important information, labels, cautions, and notes.

- The BigPrivacy SaveYourData software application is routinely referred to herein as “SaveYourData” or “SYD.”
- Labels of clickable elements in the user interface, like button names, are in **Bold.**
- Screen names are capitalized, like Audit screen.
- Text entered at the command line or in configuration files uses the `Courier New font.`
- Important notes are tagged with the notes icon 📝

Additional Resources

For stakeholders in the installation process, or admin personnel responsible for maintaining the system, please also see,

- SaveYourData User Guide

Topics Covered

The Installation Guide provides installation steps for installing the SYD server application, planning and preparing beforehand, and a test and validation plan after base installation is complete.

The major topics covered in this guide are,

- About the Installation Guide
- Introduction
- Preparation and Planning
- Installation Step by Step
- Test and Validation
- Maintenance and Troubleshooting

Who Should Read This Guide?

The primary audience for this guide is made up of the technicians and project managers from the systems integrator tasked with installing SYD at the customer site. Portions of this document may also be useful to share with the customer technicians responsible for providing an installation-ready environment.

And any stakeholder with an interest in the structure, functionality, and design of SaveYourData may find portions of this document informative.

Version and Release Notes

This guide supports SaveYourData version 1.0.0, the first GA release.

Release 1.0.0 Highlights

- **Regulatory Compliance**

  Pseudonymisation of ingest datasets up to 50 terabytes.

- **Extensive Database Support**

  Full JDBC read-only support for ingest (source) databases.

- **Saved Connection Profiles**

  Ability to save database connection profiles for repeated use.

- **Audit Record**

  Logging of events and user activity.

- **Table Selection**

  Option to select all tables, or a subset of tables for ingestion.

- **Active Directory Integration**

  Authentication against Active Directory for user login.

Getting Technical Support

Support is available from Hitachi by phone or email.

Phone	+44-203-608-9365
Email	hitachisupport@bigprivacy.com

To accelerate resolution of your issue, please confirm basic connectivity and availability of customer-owned components prior to contacting support.

Important Concepts

This section defines some of the key concepts that are referenced throughout this document. Familiarity with the topics listed will not only lend clarity to much of the content in this guide but will help lay the groundwork for a deeper understanding of the installation process.

Connection Profile

The data record SYD uses to connect to ingest databases in the end-user IT environment. Connection profiles are displayed in the user interface of the application, on the Main screen. Each profile is represented visually as a user-named 'card', each card representing a single database and table selection.

Customer

The end-user organization that purchased SYD. The customer provides the host environment for SYD. They are typically managed and supported by the systems integrator.

Destination database

The customer-provided PostgreSQL(10) database that SYD writes pseudonymised data to. It contains the output data from an ingest job. That data is referred to variously as "pseudonymised" data.

ETL

Extract, Transform, and Load. The general type of processing performed by SYD.

GDPR

The EU General Data Protection Regulation that defines "pseudonymisation" and provides guidance that defines what is and what is not compliant data.

Ingest database

The database used as the original source of data process by an ingest job.

Ingest Job

The process of connecting to a database, ingesting its dataset, processing that dataset for pseudonymisation, and writing the compliant version of the dataset to the destination database.

Installation Team

The person or persons form the systems integrator who perform the on-site installation for the customer.

JDBC

Java database connector, an API that abstracts logical database operations from the platform.

Mosaic Effect

The “Mosaic Effect” occurs when a person is indirectly identifiable due to a phenomenon referred to by the Article 29 Working Party as “unique combinations” where, notwithstanding the lack of identifiers that directly single out of a particular person, the person is still “identifiable” because that information may be combined with other pieces of information known to relate to the same individual (whether the latter is retained by the data controller or not) to create a “mosaic” of the person, enabling the individual to be distinguished from others. To help address this issue, SaveYourData replaces each occurrence of the same data element with a different dynamically generated token to separate the information value of data from the risk of re-identification via the Mosaic Effect.

MFA

Multi-factor authentication.

Pseudonymisation

GDPR Article 4(5) defines “Pseudonymisation” as requiring separation of the information value of data from the risk of re-identification. To benefit from GDPR statutory/regulatory incentives and rewards for pseudonymisation, this separation is necessary. Replacing multiple occurrences of the same personal data elements with “static” (or persistent) tokens fails to separate the information value of data from the risk of re-identification because re-identifying correlations and linkage attacks (aka the “Mosaic Effect”) are possible because “static” (or persistent) identifiers are used instead of dynamic de-identifiers. The “Mosaic Effect” occurs when a person is indirectly identifiable due to a phenomenon referred to by the Article 29 Working Party as “unique combinations” where, notwithstanding the lack of identifiers that directly single out of a particular person, the person is still “identifiable” because that information may be combined with other pieces of information known to relate to the same individual (whether the latter is retained by the data controller or not) to create a “mosaic” of the person, enabling the individual to be distinguished from others. To help address this issue, SaveYourData replaces each occurrence of the same data element with a different dynamically generated token to separate the information value of data from the risk of re-identification via the Mosaic Effect.

SYD

An acronym for SaveYourData

SYD Access Group

The group defined in the corporate directory (AD or other LDAP 3+). Its membership consists of all the users permitted to access SYD in the customer's organization.

SYD Installation Image

VMware ESXi-compatible SaveYourData (SYD) virtual machine (VM) image provided by Anonos to the systems integrator.

Systems Integrator

The umbrella term for the type of company that manages and performs the SYD installation, manages the customer, and provides front-line support for the customer and their users. (also, system integrator).

Introduction

TOPICS

- What is SaveYourData?
- Architectural and Deployment Notes
- Linux System Requirements

1

What is SaveYourData?

SYD keeps an organization’s options open for potential ongoing use of personal data collected using (now under the GDPR) legally non-compliant broad-based consent (“Legacy Personal Data”) without requiring (i) deletion or (ii) anonymisation of the data.

Regulatory guidelines issued in connection with the GDPR require organizations that previously relied upon consent to review their pre-GDPR consents to ensure they are compliant with new heightened GDPR requirements for consent. If the review reveals that the prior consent is not valid under the GDPR, data controllers are provided the following “one off” opportunity to get GDPR compliant consent or to change the lawful basis on which the relevant processing takes place:

> *If a controller finds that the consent previously obtained under the old legislation will not meet the standard of GDPR consent, then controllers must undertake action to comply with these standards, for example by refreshing consent in a GDPR-compliant way. Under the GDPR, it is not possible to swap between one lawful basis and another.* ***If a controller is unable to renew consent in a compliant way and is also unable –*** ***as a one off situation*** *–* ***to make the transition to GDPR compliance by basing data processing on a different lawful basis while ensuring that continued processing is fair and accounted for,*** *the processing activities must be stopped. In any event the controller needs to observe the principles of lawful, fair and transparent processing.*

Many organizations are being advised to delete Legacy Personal Data because they are unaware that technical solutions such as SaveYourData exist to enable data controllers and processors to “pseudonymise” data in accordance with the requirements of Article 4(5) of the EU General Data Protection Regulation (GDPR) to help support lawful secondary uses of data, like iterative analytics and artificial intelligence (AI), when consent does not provide a valid legal basis under the GDPR because processing cannot be described with specificity and unambiguity at the time of data collection.

If an organisation is subject to regulatory retention or reporting obligations, this may mean locking up Legacy Personal Data so that it is accessible only in response to regulator inquiry. Alternatively, organizations are being advised to “delete” or "anonymise" their Legacy Personal Data so that relinking (directly or indirectly) to identifying data is no longer possible. In either situation, this means that access to Legacy Personal Data for analytics, artificial intelligence, machine learning, or digital transformation may no longer be possible.

SaveYourData provides a third option instead of (i) deleting or (ii) anonymising Legacy Personal Data. SYD represents a “one off” opportunity to transform Legacy Personal Data for potential future secondary use such as analytics, artificial intelligence, machine learning or digital transformation. SaveYourData accomplishes the initial step in transitioning “to GDPR compliance by basing data processing on a different lawful basis while ensuring that continued processing is fair and accounted for" by pseudonymising personal data to help support "legitimate interests" as a lawful basis for processing.

Architectural and Deployment Notes

The SYD application is designed to support ingest and pseudonymisation of large data sets from RDBMS sources that contain sensitive private data. It works with any data platform that supports the JDBC API for platform-agnostic, read-only access.

SYD runs in the customer's VM environment in containerized modules that are lightweight and straightforward to maintain and update.

The three major containers deployed during installation are,

1. **Web UI Application**
   This is the presentation layer of the application, and the single interface for end-users.
2. **Pseudonymisation Application**
   This is the core engine responsible for pseudonymising ingest data and writing it to the destination database.
3. **Configuration and Ingest Test Database**
   A container with two database schemas – a lightweight configuration database with preconfigured values, and a data source to use for initial post-install validation tests.

The configuration tables also store the connection profiles displayed on the Ingest Data page, and the audit events that are shown on the Audit page.

Note that some configuration parameters are also stored separately, in flat files that are edited at the time of installation.

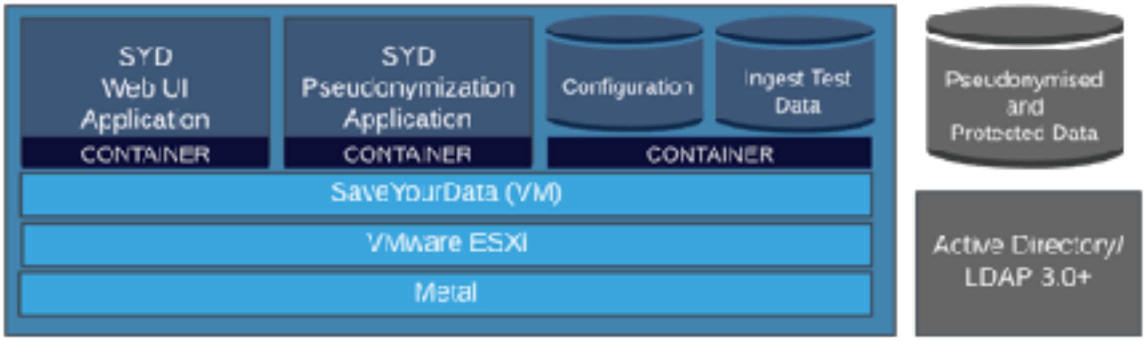

Figure 1 **Major SYD Architectural Components**

The diagram below summarizes the data flow in a deployed system, from the customer's data center, through pseudonymisation, and ultimately back into the data center, specifically the destination database, in pseudonymised form.

SYD End-to-End Data Path

Level1 DFD

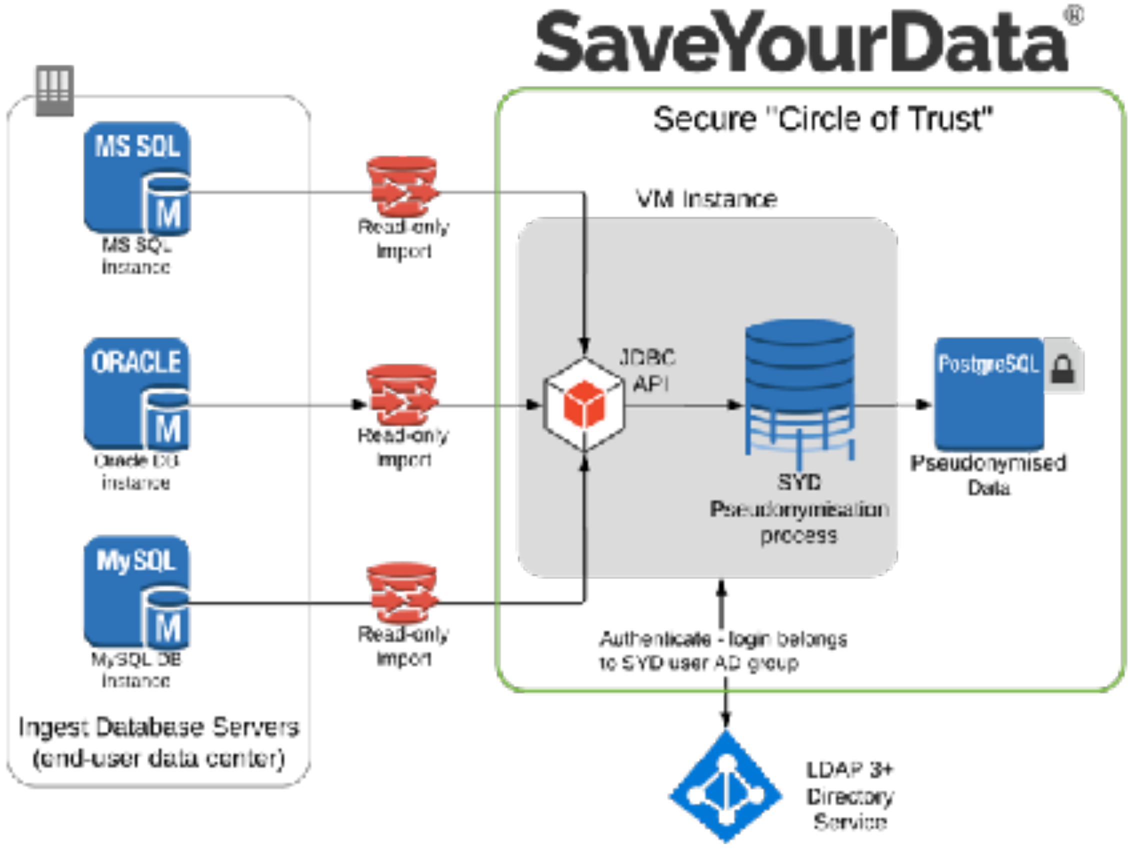

Figure 2 **SYD Data Flow**

NOTES

1. The system connects with read-only access to any JDBC-compliant RDBMS in the customer's data-center. Only a single connection is active at a given time, however the system can store connection profiles for N total ingest data platforms without restriction.

2. Active Directory or LDAP 3.0+ can be used to authenticate users against the group authorized for SYD access

3. SYD Pseudonymisation process (or "ingest job") is run by a user, and pseudonymised data is written to the designated output RDBMS in the end-user data center.

Support for Specific Data Types

The SYD ingest system supports those data types specified in the SQL-92 standard:

Supported Numeric Types

```
SMALLINT, INTEGER, INT
DECIMAL, NUMERIC, FLOAT
REAL, DOUBLE PRECISION
BIT, BIT VARYING
```

Supported String Types

```
CHARACTER, CHAR
NATIONAL CHARACTER, NATIONAL CHAR, NCHAR
CHARACTER VARYING, CHAR VARYING, VARCHAR
NATIONAL CHARACTER VARYING, NATIONAL CHAR VARYING, NCHAR VARYING
```

Supported Date and Time Types

```
DATE, TIME, INTERVAL
TIME WITH TIME ZONE, TIMESTAMP, TIMESTAMP WITH TIMEZONE
```

When the system encounters unsupported data types, like XML or LOB variations, the system will either convert to string or disregard from the ingest job.

Linux System Requirements

SYD uses a Java-based stack running on Linux with PosgreSQL as the back-end (destination) database for storing pseudonymised data. JDBC is used to connect to ingest data sources in the customer's data center.

The SYD application resides in the customer's VM environment, which must meet or exceed the following specifications.

- VMware ESXi 6.7 host with available resources:
  - 4 x 2.0+ GHz CPUs
  - 16 GB RAM
  - 100 GB storage space

- PostgreSQL 10 database with a maximum 150 TB storage capacity.
  - This is the output database for the pseudonymised data.
  - Data volumes may be stored on encrypted file systems per customer’s internal requirements.

- Microsoft Active Directory Server (2003 or later) or an LDAP v3-compatible directory server. This is required for user authentication.

- SSL certificate (optional)
  If no certificate is provided, SYD uses a “dummy” certificate that still protects network traffic with strong encryption but lacks a valid certificate authority.

Storing Output Data in Encrypted Format
Although SYD doesn't manage encryption natively, if the customer requires that output data be encrypted they can enable that in PostgreSQL without affecting the functionality of SYD.

See online PostgreSQL documentation for more information.

Preparation and Planning

TOPICS

- Preparation and Planning Overview
- Roles and Responsibilities
- Pre-Installation Requirements

2

Preparation and Planning Overview

To help facilitate efficient implementation, this section provides an orientation for the installation technicians on the systems integrator (SI) side, and outlines important pre-install tasks, not only for the SI team, but for the customer IT technicians tasked with supporting implementation.

Installation Process in Brief

Prior to implementation, Anonos provides the systems integrator with a VMware ESXi-compatible SYD VM image, the "SYD Installation Image".

Because SYD is hosted in the customer's IT environment, the systems integrator guides the customer in the preparation of the installation host as well as supporting systems, like the PostgreSQL (10) database, and the configuration of an LDAP group in their existing corporate directory for users to authenticate against.

During installation, the SI technician deploys the SYD installation image in the customer’s VMware ESXi environment. The resulting VM contains pre-installed Docker containers that provide all of SYD's UI and ingest functionality. The installation technician also performs several configuration tasks, including installing an SSL certificate, applying configuration values to several flat files, and defining access information for the corporate LDAP or Active Directory server.

After completing configuration, the containers are started, and a test plan is executed to verify that the system is functioning properly.

In production, the customer’s end-users browse to a TLS-secured URL to log in and access SYD's web UI. They use their corporate IT username and password, typically the same ones they use for their email if they use Active Directory. Provided the user belongs to the designated SYD access group, they are logged in with full access to the SYD UI and all its functions.

Roles and Responsibilities

The installation team typically belongs to an Anonos-authorized systems integrator, VAR, or other technology partner. As the OEM, Anonos isn't directly involved in installation.

Essentially, the SI takes the lead in the implementation, and their installation team is responsible for managing the customer and the installation process. Anonos provides an installation image, documentation, and tier 2/3 escalation support after installation.

Allocation of Responsibilities for OEM (Anonos) and Installation Team

Anonos	Systems Integrator (SI) Installation Team
Application Software Provide ESXi VM image, including containers, for SaveYourData to the Anonos-authorized installation technician.	**VM Install** Install the SaveYourData VM image on customer ESXi cluster
Tier 2+ Support Ongoing tier 2/3+ support	**Front-line Support and Maintenance** Ongoing maintenance and tier 0/1/2 support
	Project Coordination and Management Coordinate deployment of infrastructure pseudonymised data database (PostgreSQL) to store JITI data: • Create a DB user with full access to the database • Configure DB username, password, hostname, port, and database name in SYD
	SSL Cert Install customer SSL certificate on VM (customer may elect not to provide certs)
	Coordinate LDAP Coordinate configuration of LDAP directory and connector: • Coordinate networking (static/dynamic IP configuration on the SYD; DNS hostname for end users for their login URL)

Pre-Installation Requirements

There are several vital prerequisites to on-site installation, and it's primarily the responsibility of your installation team to communicate them accurately to the customer.

All pre-installation tasks rely primarily on customer IT resources. Whether or not your organization provides staff on-site for these tasks is largely a business and logistical decision, but as the SI, you're responsible for ensuring the customer is ready for installation.

This section discusses exactly that–the tasks that must be completed prior to installation. They are all "blocking" in regard to the tasks required of the implementation team, which can't physically begin their tasks until these prerequisites are met.

Customer Pre-Installation Tasks

1. **Prepare VM Environment**
   Verify that VM Environment meets specifications and is installation-Ready

2. **Deploy PostgreSQL 10 150TB (maximum) Destination Database**
   Destination Database (PostgreSQL 10) is installed, empty, and accessible from the VM environment

3. **SYD Access Group Configured in LDAP/Active Directory**
   Active Directory or LDAP v3+ Directory Server configured

4. **Update JDBC Drivers**
   Update JDBC Drivers on VM if required (by default, a PostgreSQL driver is installed)

5. **Provide SSL Certificate (optional)**
   SYD uses a dummy certificate if none is provided.

Figure **3** on the following page depicts the components required external to the VM.

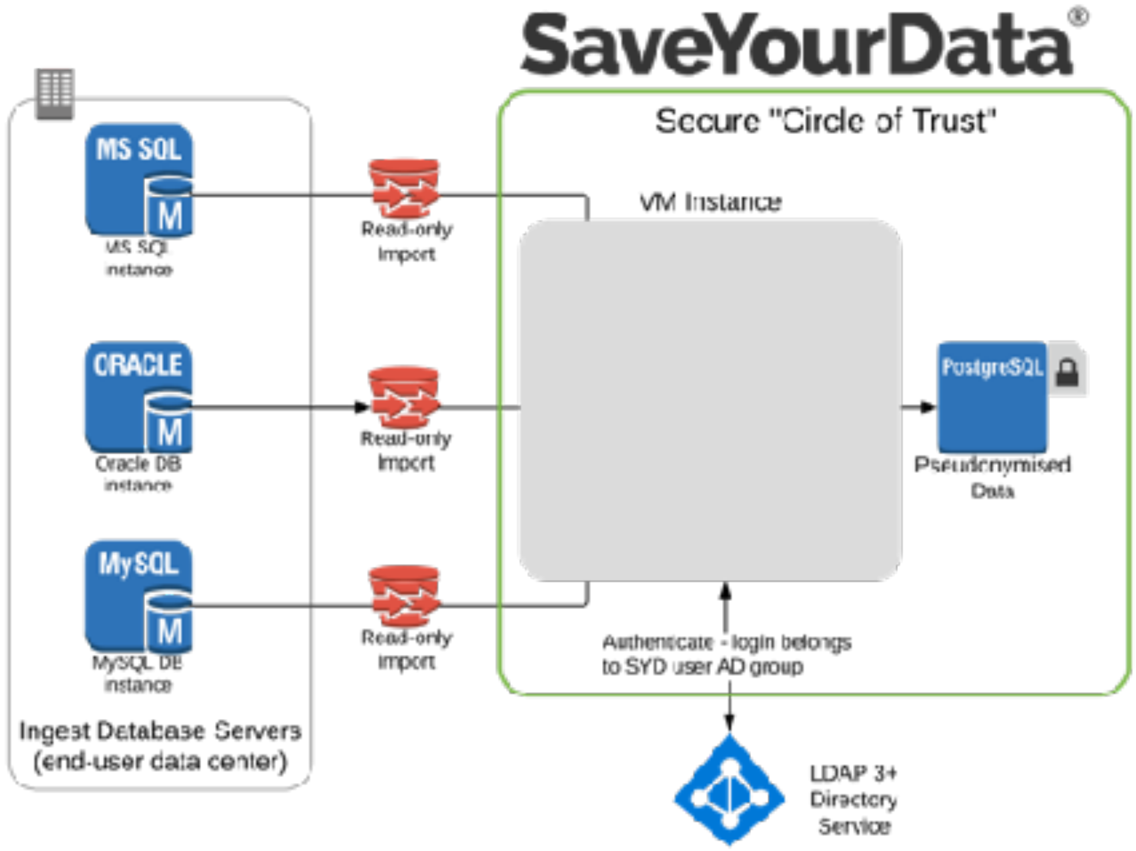

Figure 3 **Components configured and managed by the customer**

Prepare VM Environment

The VM environment, where the SYD image is deployed, must meet minimum requirements prior to installation.

Who Prepares the VM and What are the Specs?

The customer's IT team is typically responsible for installing this component, unless alternate arrangements are made.

VM Environment Specifications

- VMware ESXi 6.7 host with available resources:
  - 4 x 2.0+ GHz CPUs
  - 16 GB RAM
  - 100 GB storage space

- Microsoft Active Directory Server (2003 or later) or an LDAP v3-compatible directory server (required for user authentication)

Customer checklist for VM Environment

✓ **4 CPUs** Ensure availability of 4 2.0+ GHz CPUs

✓ **16 GB RAM** Ensure availability of 16GB of RAM

✓ **100GB Storage** Ensure availability of 100GB of storage space

If you or the customer has questions about the compatibility of the VM environment with ESXi 6.7, VMware maintains definitive specifications that can be found at

`https://www.vmware.com/resources/compatibility/search.php`

Optimizing the Host Environment for Performance

The role of network performance is another point of education worth bringing to the customer's attention early in the project. Prior to installation, the customer should be made aware of the crucial role their infrastructure plays in the performance of SYD ingest jobs, especially when ingesting data volumes on a terabyte scale.

Recall that an "ingest job" is the end-to-end process of pseudonymising a customer-owned dataset – all the steps, from connecting to the ingest database, ingesting the data, and processing and writing the resulting pseudonymised data to the destination database. It's the customer's infrastructure that ultimately dictates the duration of ingest jobs. More specifically, it's the maximum data transfer rate across the LAN that is the primary limiting factor in overall system performance. Knowing the implications of data transfer rate on job durations may drive your customer's decisions about kind of resources they allocate to the SYD VM host.

Table 1 Ingest Job Durations below shows some sample ingest volumes vs. transfer rates, and a grid of durations for each combination. Note the bottom row in particular, and the multi-day job durations required for a 50TB dataset over 5 days at 1 Gbps.

Ingest Data Volume - Rate	1 Gbps	10 Gbps	40 Gbps
1 GB	0:00:00:09	0:00:00:01	>0:00:00:01
2 GB	0:00:00:17	0:00:00:02	>0:00:00:01
50 GB	0:00:07:09	0:00:00:43	0:00:00:11
100 GB	0:00:14:19	0:00:01:26	0:00:00:21
1 TB	0:02:26:36	0:00:14:40	0:00:03:40
2 TB	0:04:53:12	0:00:29:19	0:00:07:20
50 TB	5:02:10:05	0:12:13:00	0:03:03:15

Format: days:hh:mm:ss

Table 1 **Ingest Job Durations**

Additionally, some factors that affect network performance are transient, so it can't be assumed that your data transfer rate will always perform at its benchmarked maximum.

If the customer is made aware of the implications of data transfer rates, they'll be better equipped to make decisions about the VM deployment environment, and to make the trade-offs appropriate to optimize transfer rate based on their priorities.

A convenient resource for calculating transfer times in various units can be found at https://www.expedient.com/file-transfer-time-calculator.

Talking with the Customer About Job Duration
Everyone on the installation team should be prepared to answer common questions about the system, its components, performance, and fundamental functionality. Especially when processing terabytes of data, the customer may express concerns about the running times for such jobs.

The simple fact is that the data transfer rate across the network is *the* gating factor governing job duration, and that rate is entirely dependent on the IT infrastructure at the customer's site.

It's important not to leave the customer with the impression that the SYD application – its architecture or algorithms – is a limiting factor in the performance of the system.

Deploy the PostgreSQL Destination Database

The destination database is where the pseudonymised output of ingest jobs is stored, and it must be deployed, empty, and accessible across the network before the first ingest job will run.

Who Configures It and What are the Specs?

The customer's IT team is typically responsible for installing this component, unless alternate arrangements are made. The destination database is a PostgreSQL 10 database with a maximum 150 TB capacity (empty at installation). On the SYD side, the destination database location and credentials are described as key-value pairs in the `/opt/saveyourdata/config/ingest.properties` file, which is configured based on customer-provided information by the installation team, as part of the install process. Note also that if the customer intends to use a port other than the PostgreSQL default of 5432, it's important to ensure that the matching configuration parameter is set in the `ingest.properties` file during installation.

Encryption Options

Customers who require that stored pseudonymised data be encrypted can do so using the file system partition encryption options described in PostgreSQL documentation. Although SYD provides no native encryption faculties, it is completely agnostic to storage encryption on the server.

Customer Checklist for PostgreSQL Installation

- ✓ **PostgreSQL Destination Database is Production-Ready** Ensure the database is reachable from the installation VM environment and that it's blank.
- ✓ **150TB** Verify capacity is no more than 150TB.
- ✓ **Set Permissions** Database user must have `ALL` and `CREATE` permissions enabled.
- ✓ **PostgreSQL 10** Verify version is PostgreSQL 10 destination database
    - ○ **Database URL** Provide database URL to the installation team
    - ○ **Username** Provide username to the installation team
    - o **Password** Provide password to the installation team
    - ○ **Port 5432** If the production IP port will not be 5432, the installation team must make the necessary changes to the `ingest.properties` file during installation.

Setting Permissions on the Destination Database

The user account assigned to the SYD application for access to the destination database must have `ALL` and `CREATE` permissions enabled.

This is accomplished by executing the following commands in PostgreSQL:

```
CREATE ROLE [USER] LOGIN ENCRYPTED PASSWORD '<password>';

GRANT ALL ON DATABASE [DATABASE_NAME] TO [USER];

GRANT CREATE ON DATABASE [DATABASE_NAME] TO [USER];
```

Deploying JDBC Drivers to the SYD VM

Customers should deploy JDBC drivers on the SYD VM to enable SYD to connect to the input database systems using JDBC.

Drivers can be found on the website for each RDBMS OEM, and the license for each driver (specific to a particular RDBMS platform) must be accepted individually by the end-user at installation. Below are some example resources, but check the web to ensure you have the most current information.

- For Oracle:
  - (Version 12c) Oracle 12.1.0.2.0 JDBC 4.1 or higher within 12.x.x.x.x. See http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html for information.
  - (Version 11g) See http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html for information.
  - Sterling B2B Integrator supports JDBC Type-4 drivers on a single node of a database except with Oracle Real Application Clusters (RAC). The JDBC Type-4 drivers to can be used to connect with multiple nodes of an Oracle RAC.

- For Microsoft SQL Server:
  - Microsoft SQL Server 2014 - Use SQL Server JDBC Driver 4.x
  - Microsoft SQL Server 2012 - Use SQL Server JDBC Driver 4.x
  - Microsoft SQL Server 2008 - Use SQL Server JDBC Driver 3.0
  - Regardless of the Microsoft SQL Server version, when using the Lightweight JDBC Adapter, use SQL Server JDBC Driver 4.x

To obtain the driver, go to the Microsoft Download Center at http://www.microsoft.com/en-us/download/default.aspx and search for the required SQL Server JDBC driver version.

- For DB2®, see http://www.ibm.com/support/docview.wss?uid=swg21363866 for information.

The JDBC Drivers for your ingest databases are deployed to the SYD VM by copying them to the `/opt/drivers` directory. Multiple JDBC driver files may be copied to this directory to support the various ingest database engines you may have in your environment.

Logistical Notes for Systems Integrators

Your organization already has processes and procedures in place for delivering technology to your customers. They're built around best practices and the installation requirements specific to each application or solution you deliver. The following notes are provided to inform the procedures you develop around SYD for communicating with your customer and managing the implementation.

First, we'll look at a sample deployment timeline to illustrate the sequence of key events and the responsible party, if not the actual schedule. For convenience, the next page also contains a centralized list of the information required from the customer in order to complete installation.

Sample Deployment Timeline

Period	**Customer** Action Items	**SI / Installation Team** Action Items	**Anonos** Action Items
Week 1		• Issues PO to Anonos • Kickoff with customer • Issues specs, prerequisites, schedule to customer	
Week 2	• Builds VM deployment environment to Anonos specs. • Configures SYD authorized user group in LDAP. • Builds PostgreSQL destination database. • Procures SSL cert (optional)		Delivers SYD install image to installation team
Week 3	• Reviews site readiness with installation team • Reviews availability of customer IT resources for on-site installation with installation team • Issues access info to Installation Team for destination database and LDAP	• Reviews site readiness with customer • Reviews customer support staff availability for on-site installation with customer	
Week n	• IT staff supports installation team on-site.	• Installation team on-site. • Notifies Anonos of install and start of support period.	• Acknowledges install and start of support period.

SYD Installation Data Provided by the Customer

Below is a comprehensive list of data that the customer provides to the installation team based on their pre-installation tasks. This data must be available to the installation team before SYD can be installed and brought on line.

Component	Item	Value
VM Environment access Administrative access to the pre-installed ESXi server or to the underlying vSphere cluster	URL	
	username	
	password	
	IP port	
Destination database access (PostgreSQL)	URL	
	username	
	password	
	IP port	
LDAP Directory access SYD users are authenticated against members of the SYD access group in the directory.	URL	
	username	
	password	
	SYD access group name*	
Test User (must be in directory and belong to SYD auth group)	Test User username	
	Test User password	

Installation Step-By-Step

TOPICS

- Installation Procedure
- Installation Part 1
- Installation Part 2
- Installation Part 3

3

SYD Installation Procedure

With the prerequisites from the previous section properly addressed, the system is ready for on-site installation.

Everyone on the installation team should be familiar with the procedures in this section prior to the scheduled site visit.

SYD solution installation is divided in to three major parts:

1. VM Image Deployment
2. Configuration Prior to Container Startup
3. Startup Containers

Installation Part 1 – VM Image Deployment

The VM image provided by Anonos is installed through the VMware installation tool. To install the VM image, you must have administrative access to the pre-installed ESXi server or to the underlying vSphere cluster. OS-level "root" access not required for any ESXi operations.

Use the following procedure to deploy the installation image.

TO CONFIGURE SYD PARAMETERS IN VMWARE

1. Open VMware ESXi in the installation environment.
2. Select **Create / Register VM** in the VMware ESXi interface.

   The screen shown in *Figure 4* is displayed.

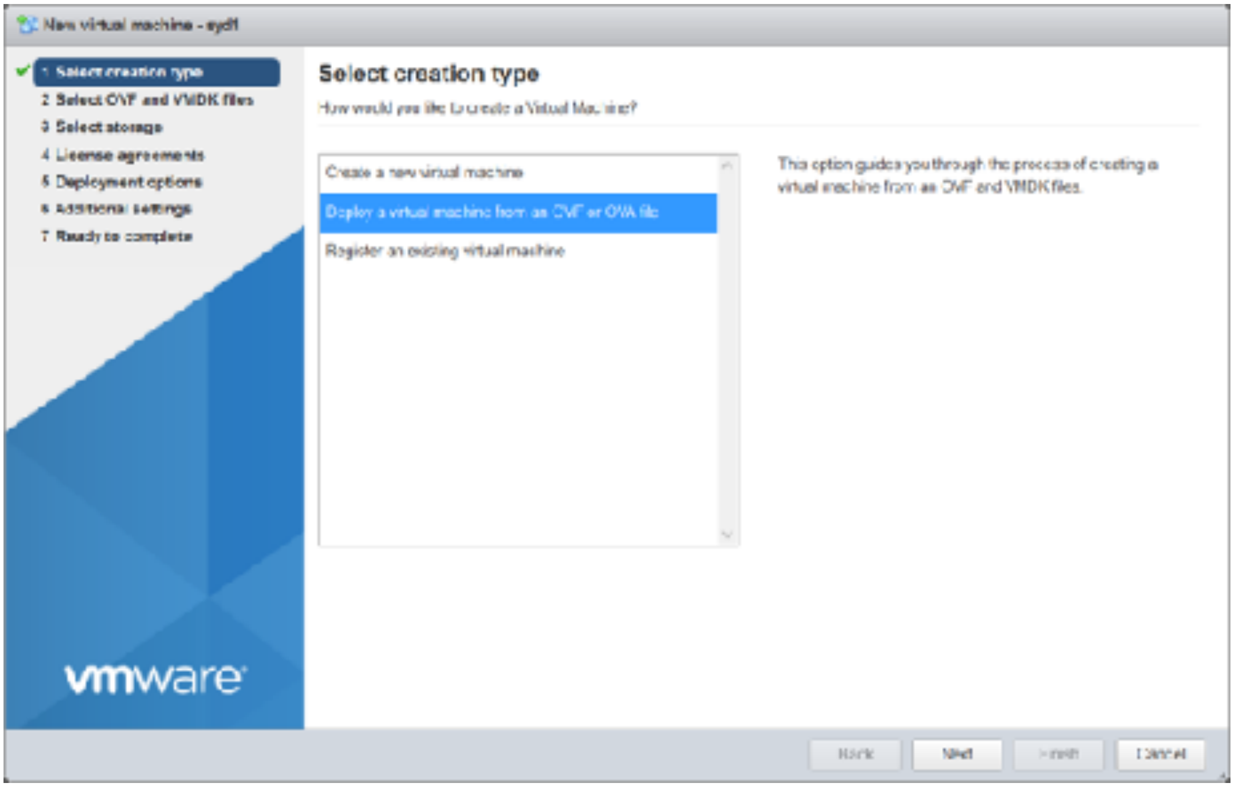

New virtual machine - syd1

1 Select creation type
2 Select OVF and VMDK files
3 Select storage
4 License agreements
5 Deployment options
6 Additional settings
7 Ready to complete

Select creation type

How would you like to create a Virtual Machine?

Create a new virtual machine
Deploy a virtual machine from an OVF or OVA file
Register an existing virtual machine

This option guides you through the process of creating a virtual machine from an OVF and VMDK files.

Figure 4 **VMware configuration for new virtual machine**

3. Select “Deploy a virtual machine from an OVF or OVA file”, then click **Next**.

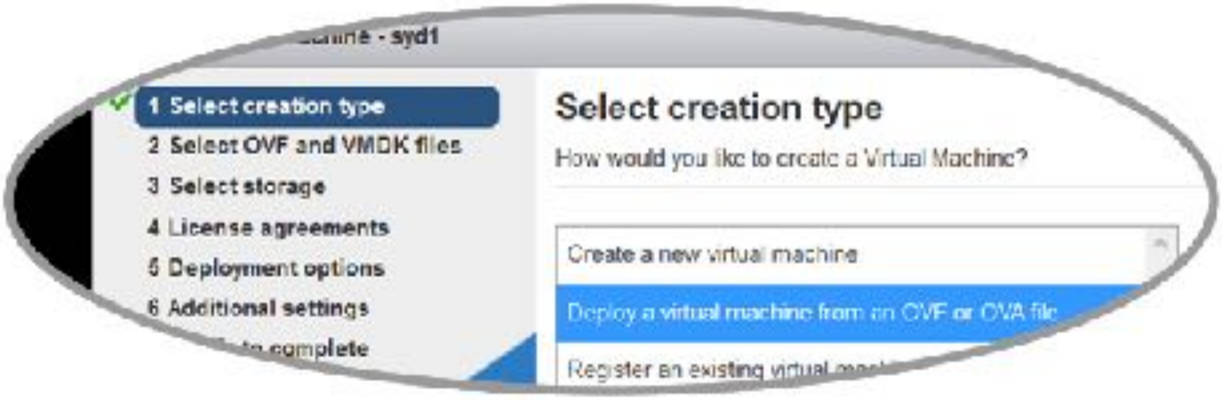

achine - syd1

1 Select creation type
2 Select OVF and VMDK files
3 Select storage
4 License agreements
5 Deployment options
6 Additional settings
 to complete

Select creation type

How would you like to create a Virtual Machine?

Create a new virtual machine
Deploy a virtual machine from an OVF or OVA file
Register an existing virtual mach

4. Enter a name for the virtual machine and click the light blue box to open a file selection dialog.

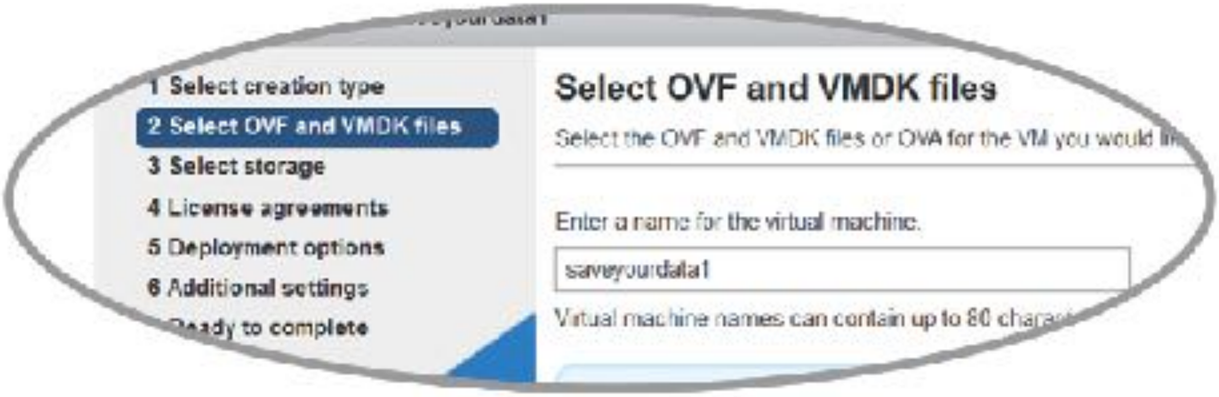

yourdata1

1 Select creation type
2 Select OVF and VMDK files
3 Select storage
4 License agreements
5 Deployment options
6 Additional settings
 Ready to complete

Select OVF and VMDK files

Select the OVF and VMDK files or OVA for the VM you would li

Enter a name for the virtual machine.

saveyourdata1

Virtual machine names can contain up to 80 charac

5. Select the .ovf and .vmdk files provided by Anonos, then select **Next**.

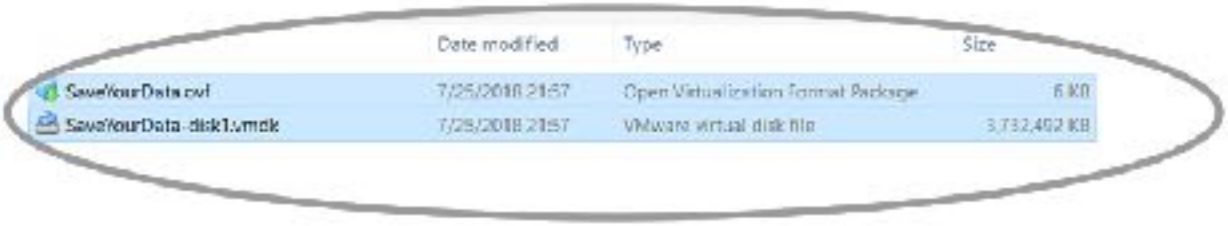

	Date modified	Type	Size
SaveYourData.ovf	7/25/2018 21:57	Open Virtualization Format Package	6 KB
SaveYourData-disk1.vmdk	7/25/2018 21:57	VMware virtual disk file	3,732,492 KB

6. Select the data store where the virtual machine is to be located. The data store must have at least 100 GB free. Select **Next**.

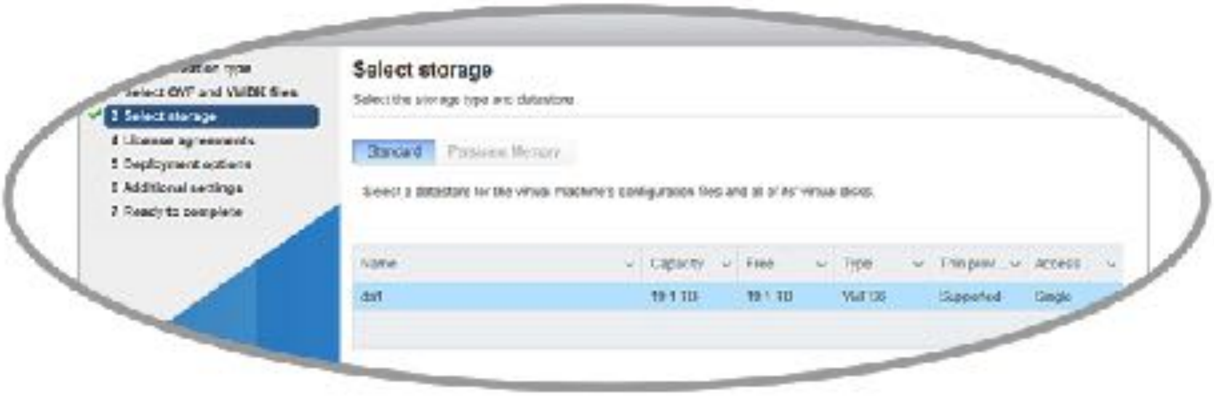

3 Select storage
4 License agreements
5 Deployment options
6 Additional settings
7 Ready to complete

Select storage

Select the storage type and datastore

Select a datastore for the virtual machine's configuration files and all of its virtual disks.

Name	Capacity	Free	Type	Thin prov...	Access
ds1	19.1 TB	19.1 TB	VMFS6	Supported	Single

7. Select your preferred network and disk provisioning types according to the customer's local IT configurations & policies. Select **Next**.

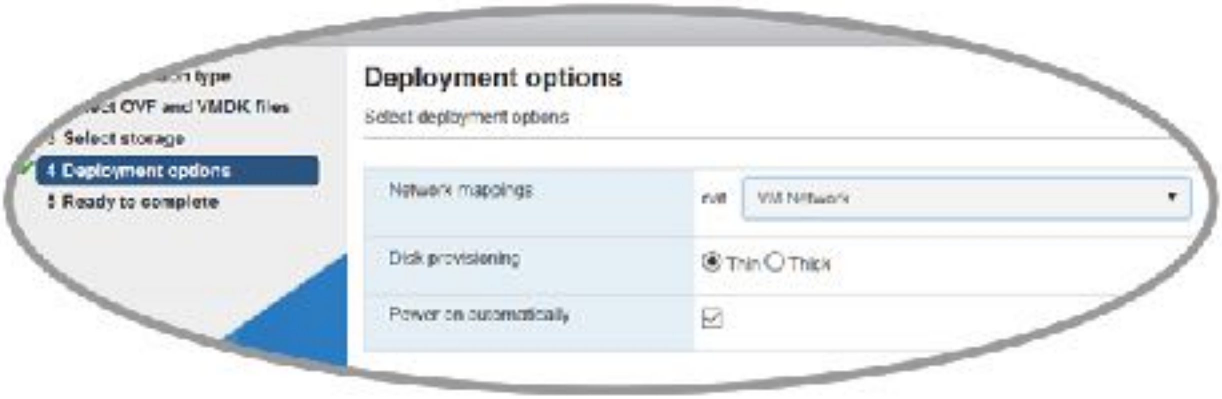

Select OVF and VMDK files
3 Select storage
4 Deployment options
5 Ready to complete

Deployment options

Select deployment options

Network mappings	nat: VM Network ▾
Disk provisioning	● Thin ○ Thick
Power on automatically	☑

8. Review the VM configuration and make any necessary corrections, then select **Finish**.

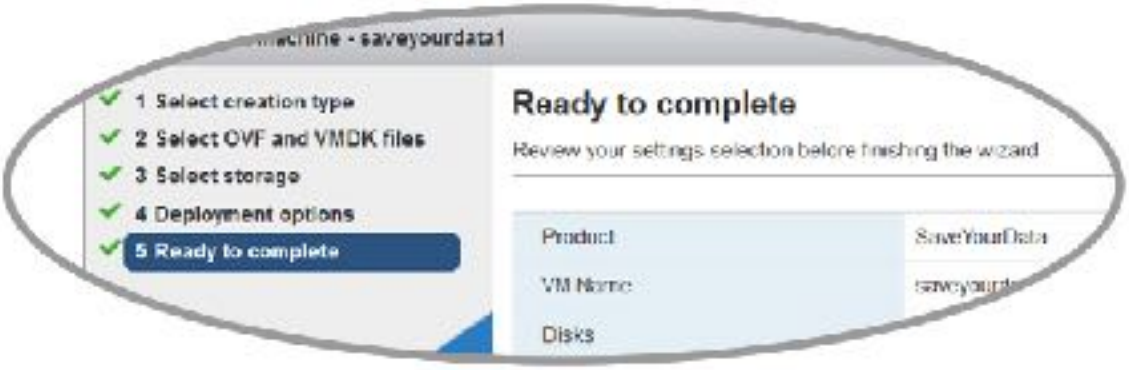

machine - saveyourdata1

✓ 1 Select creation type
✓ 2 Select OVF and VMDK files
✓ 3 Select storage
✓ 4 Deployment options
✓ 5 Ready to complete

Ready to complete

Review your settings selection before finishing the wizard

Product	SaveYourData
VM Name	saveyourd
Disks	

It may take several minutes for the VM to deploy, depending on the existing load on shared resources.

CAUTION! Avoid using the browser until deployment is complete
It's important to let the deployment process run to completion before navigating to another page, or otherwise interacting with the browser.

Doing so may cause unpredictable system behavior or installation failure.

Upon completion of image installation, the VM and application containers are ready for configuration.

Before starting the containers, two SYD-specific configuration files must be configured, and Ubuntu networking must be properly defined, as detailed in the *Installation Part 2 – Configuration* procedure on the next page.

If you experience errors installing the SYD image, see section *5 Maintenance and Troubleshooting*.

Shouldn't Docker be Installed to Support Containers?
If you've worked with container images on other installations, you might have experience installing Docker or a similar application to support those containers. Anonos pre-installs both Docker and the required containers in VMware ESXi, providing a single load-and-go image that removes the need for installation technicians to handle those components separately.

Installation Part 2 – Configuration

The two files that describe installation-time configuration in key-value pairs must be edited prior to starting the containers. Both are flat files that are included in the image you just deployed in Part 1. They each contain parameters that are vital to the proper interaction of SYD with the customer's IT environment:

- **`.env`**
  `PATH: /opt/saveyourdata/config/.env file`
  HTTPS and access management configuration parameters are defined as key-value pairs.

- **`ingest.properties`**
  `PATH: /opt/saveyourdata/config/ingest.properties file`
  Pseudonymised data database location configuration is specified as key-value pairs in the `/opt/saveyourdata/config /ingest.properties` file.

TO CONFIGURE SYD

1. Log into the SYD VM.
    a. If you have network access to the SYD VM and prefer to log in via SSH, please refer to the instructions in Appendix A in the back of this document.

    b. If you have console access to the SYD VM via the ESXi user interface, you may log into the system with the username “ubuntu” and the password “ubuntu”. See the image below for reference. Please note these console credentials are unavailable for use over SSH connections.

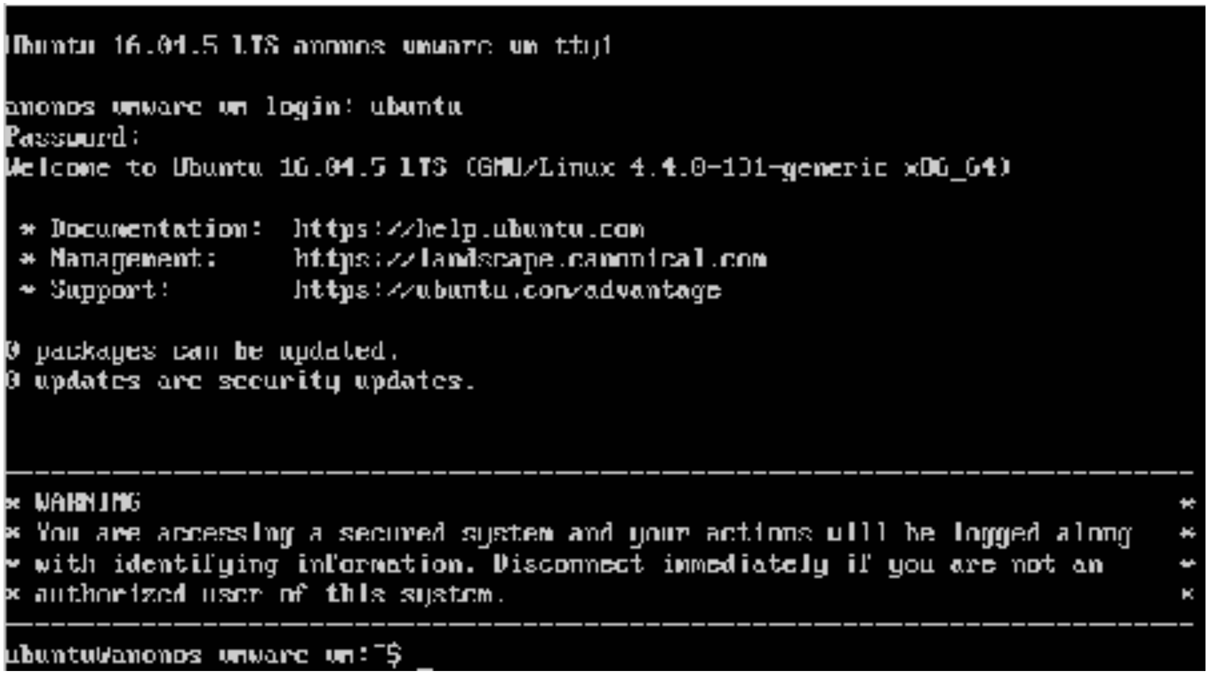

```
Ubuntu 16.04.5 LTS anonos vmware vm tty1

anonos vmware vm login: ubuntu
Password:
Welcome to Ubuntu 16.04.5 LTS (GNU/Linux 4.4.0-131-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/advantage

0 packages can be updated.
0 updates are security updates.


------------------------------------------------------------------------
* WARNING                                                              *
* You are accessing a secured system and your actions will be logged along  *
* with identifying information. Disconnect immediately if you are not an    *
* authorized user of this system.                                      *
------------------------------------------------------------------------
ubuntu@anonos vmware vm:~$ _
```

2. Enter super-user mode by running `sudo su -`. When prompted, enter the password “ubuntu”.

3. Ensure that networking is properly defined in the interfaces file. Edit the file by running `nano /etc/network/interfaces` (shown below). Also note,

    - The format of this file is defined in standard Debian & Ubuntu Linux file `/etc/network/interfaces`. Refer to Linux manual page interfaces(5) for syntax.
    - The SYD VM is configured to receive networking settings via DHCP by default.
    - Consult the end-user IT contact for appropriate configuration settings.

4. Open SSH and log in to the SYD VM using "ubuntu " for both the username and password.

5. Ensure that networking is properly defined in the interfaces file. Edit the file by running `sudo nano /etc/network/interfaces` (shown below). Also note,

    - The format of this file is defined in standard Debian & Ubuntu Linux file `/`

`etc/network/interfaces`. Refer to Linux manual page interfaces(5) for syntax.

- Consult the end-user IT contact for appropriate configuration settings.

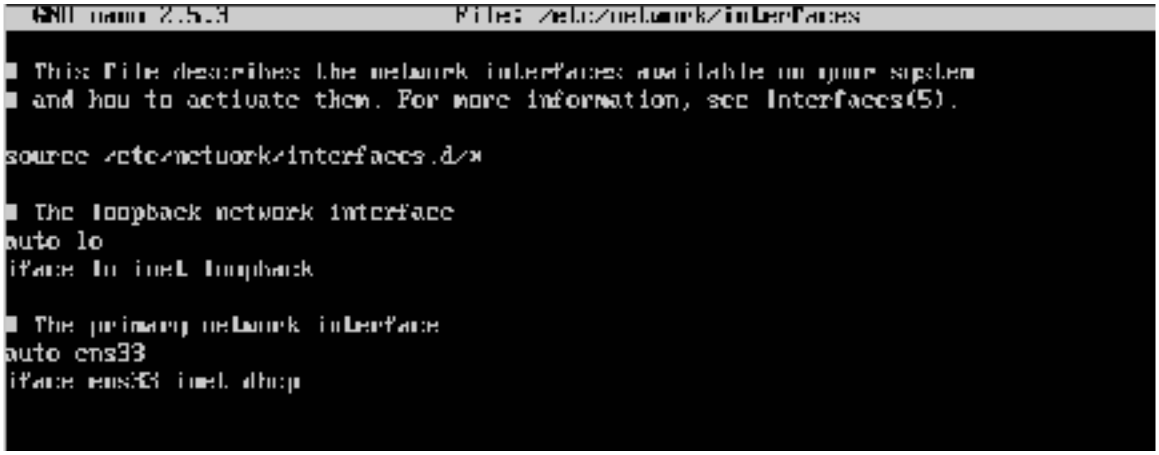

```
  GNU nano 2.5.3                File: /etc/network/interfaces

# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).

source /etc/network/interfaces.d/*

# The loopback network interface
auto lo
iface lo inet loopback

# The primary network interface
auto ens33
iface ens33 inet dhcp
```

6. If you modified /etc/network/interfaces in the previous step, restart SYD VM networking services by running the command `service networking restart`. If you originally logged into the SYD VM via SSH, your connection may be interrupted. In the case of DHCP-based IP address assignment, the IP address of the SYD VM may change at this time, and it may be retrieved at the ESXi interface.

7. Set HTTPS and access management configuration values by editing the `.env` file. Edit the file by running `nano /opt/saveyourdata/config/.env`. Only modify parameters specified in Table 2, the *Configuration Parameters in the .env File* on page 49. Save the file. (Consider backing up the file for reference.)

   In addition to the consideration of other configuration parameters, if the user has provisioned an SSL certificate and private key to use with this SYD VM, specify the locations for those files in this step and follow the next two steps to upload the files to the SYD VM.

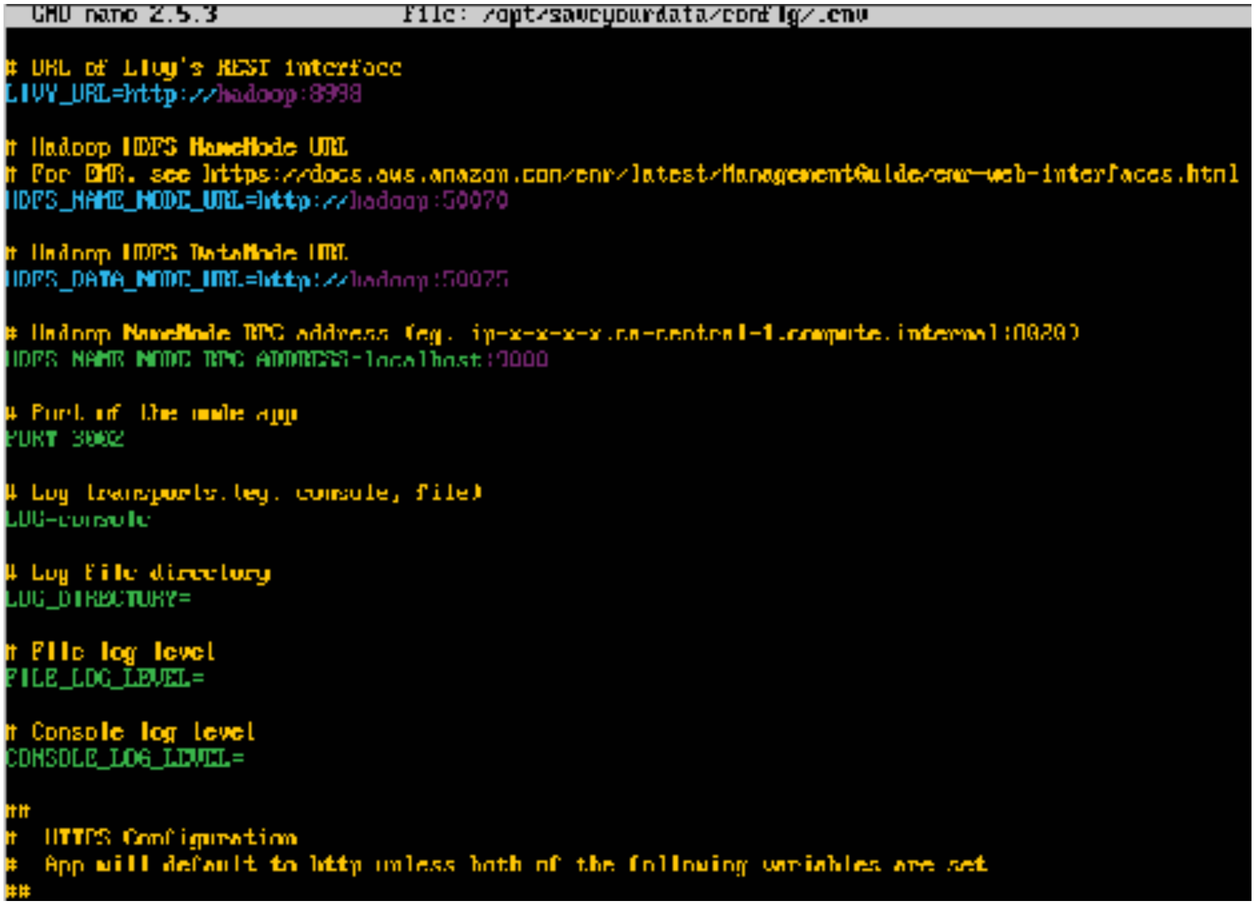

```
  GNU nano 2.5.3               File: /opt/saveyourdata/config/.env

# URL of Livy's REST interface
LIVY_URL=http://hadoop:8998

# Hadoop HDFS NameNode URL
# For EMR, see https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-web-interfaces.html
HDFS_NAME_NODE_URL=http://hadoop:50070

# Hadoop HDFS DataNode URL
HDFS_DATA_NODE_URL=http://hadoop:50075

# Hadoop NameNode RPC address (eg. ip-x-x-x-x.ca-central-1.compute.internal:8020)
HDFS_NAME_NODE_RPC_ADDRESS=localhost:9000

# Port of the node app
PORT=3002

# Log transports (eg. console, file)
LOG=console

# Log file directory
LOG_DIRECTORY=

# File log level
FILE_LOG_LEVEL=

# Console log level
CONSOLE_LOG_LEVEL=

##
#  HTTPS Configuration
#  App will default to http unless both of the following variables are set
##
```

8. (If the user has provisioned an SSL certificate and private key) Using SCP from an external computer, upload the user's SSL certificate and private key from your computer (or the computer where those files are stored) to the SYD VM. For example, execute the following command from a Linux or Mac OS terminal to upload a certificate and key for `int.example.com` to the `/home/ubuntu` directory of the SYD VM at the IP address 10.0.42.200: `scp int.example.com.crt int.example.com.key ubuntu@10.0.42.200:`

9. (If the user has provisioned an SSL certificate and private key) Back on the SYD VM, copy the SSL certificate and private key to the paths and filenames you specified as values for the `HTTPS_CERT` and `HTTPS_KEY` parameters in the `.env` file. Using the example files from the previous step, the commands are:

   a. `cp /home/ubuntu/int.example.com.crt /etc/ssl/certs/`

   b. `cp /home/ubuntu/int.example.com.key /etc/ssl/private/`

10. Set configuration values for data ingest in the `ingest.properties` file. Edit the file by running,
    `nano /opt/saveyourdata/config/ingest.properties`.

    Only modify parameters specified in Table 3 *Configuration Parameters in the ingest.properties File* on page 50. Save the file. (Consider backing up the file for reference.)

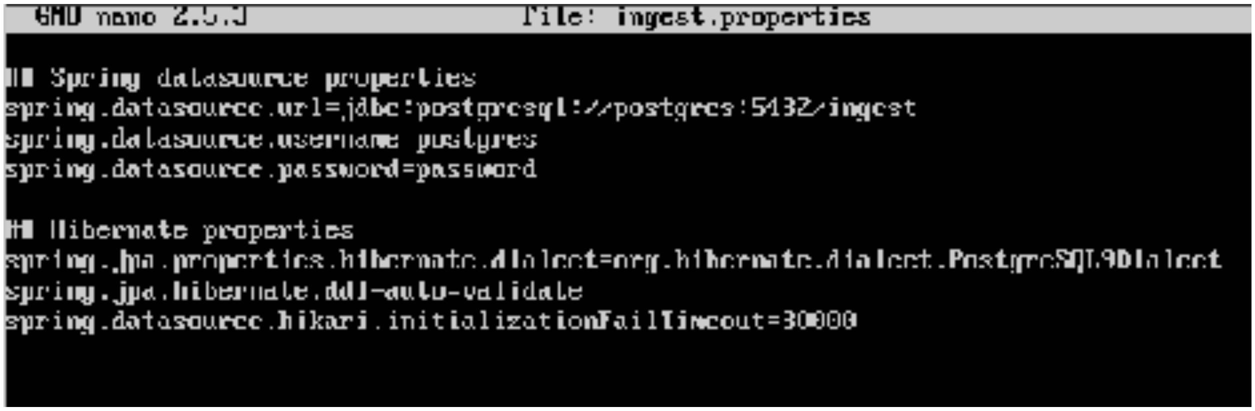

```
  GNU nano 2.5.3                     File: ingest.properties

## Spring datasource properties
spring.datasource.url=jdbc:postgresql://postgres:5432/ingest
spring.datasource.username=postgres
spring.datasource.password=password

## Hibernate properties
spring.jpa.properties.hibernate.dialect=org.hibernate.dialect.PostgreSQL9Dialect
spring.jpa.hibernate.ddl-auto=validate
spring.datasource.hikari.initializationFailTimeout=30000
```

Upon completion of configuration, the containers should be ready to start.
We'll look at that process next, in *Installation Part 3 – Container Startup*.

Configuration Settings in the .env File

(Use this as a reference to complete step 2 of *Installation Part 2*.)

Parameter	Notes
`HTTPS_CERT=<cert_path>`	Absolute path to the SSL certificate. Default: `/etc/ssl/certs/example.crt` SSL certificates must be stored in the `/etc/ssl/certs` directory.
`HTTPS_KEY=<key_path>`	Absolute path to the private key for the SSL certificate. Default:`/etc/ssl/private/example.key` SSL certificate private keys must be stored in the `/etc/ssl/private` directory.
`DIRECTORY_URL=<url>`	URL to the Active Directory or LDAP 3+ Server Ex: `ldap://corp.company.net`
`DIRECTORY_BASE_DN=<dn>`	Base DN Ex: `OU=corp,DC=corp,DC=company,DC=net`
`DIRECTORY_USER=<user>`	Simple Authentication Distinguished Name to bind. The user name SYD will pass to the AD server to log in. Ex: `CORP\\Admin`
`DIRECTORY_PASSWORD=<password>`	Simple Authentication password. The password SYD will pass to the AD server to log in. Ex: `Temp123$`
`DIRECTORY_GROUP=<group_name>`	Group name for SaveYourData users. Members of this group may access SYD. Ex: `syd_group1`

Table 2 **.env file configuration settings**

Configuration Settings in the ingest.properties File

(Please use this as a reference to complete step 4 of *Installation Part 2*.)

Parameter	Notes
`spring.datasource.url=<jdbc_url>`	JDBC URL to an empty PostgreSQL 10 database with 150 TB capacity Ex: jdbc:postgresql://syd1.internal.example.com:5432/syd-protected
`spring.datasource.username=<username>`	Username for the database at spring.datasource.url
`spring.datasource.password=<password>`	Password for the database at spring.datasource.url

Table 3 **ingest.properties file configuration settings**

Installation Part 3 – Container Start Up

After confirming that all the parameters are properly entered in the configuration files in *Installation Part 2*, you are ready to start the SYD containers.

TO START THE SYD APPLICATION CONTAINERS

1. While logged in to the VM, enter the following three commands to start the containers.

    1. (If you have not already authenticated as a super-user) `sudo su -`

    2. `cd /opt/saveyourdata`

    3. `docker-compose up -d`

    You will see the following services start.

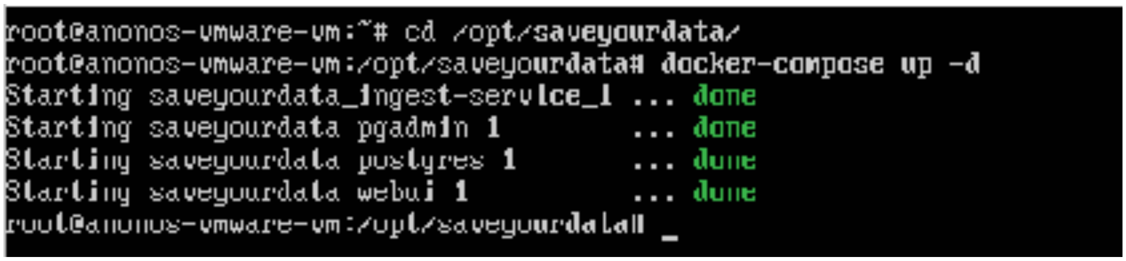

```
root@anonos-vmware-vm:~# cd /opt/saveyourdata/
root@anonos-vmware-vm:/opt/saveyourdata# docker-compose up -d
Starting saveyourdata_ingest-service_1 ... done
Starting saveyourdata_pgadmin_1        ... done
Starting saveyourdata_postgres_1       ... done
Starting saveyourdata_webui_1          ... done
root@anonos-vmware-vm:/opt/saveyourdata# _
```

At this point, barring any errors or obvious issues, please proceed to the test and validation procedure in the next section, to confirm that SYD is now in production and that the installation is complete.

If you experience errors while bringing the containers on line, see section *5 Maintenance and Troubleshooting*.

Test and Validation

4

Functional Test Plan: Verifying Core Functionality

The recommended functional test should take about 10-20 minutes if everything is working correctly. If you're not familiar with the application, the SaveYourData User Guide will help you run through the test plan.

Goals

The functional test is designed to confirm all the core features available to authorized users from the UI. That includes login, defining connectivity profiles, connecting to a database, and running an "ingest job".

It should confirm application readiness and connectivity.

Test Case Sequence

1. Initial Validation
2. Connect and Run an Ingest Job on Test Data
3. Connect and Run an Ingest Job on Customer Data
4. User Login Scenarios
5. Connect and Run an Ingest Job on Customer Data with Selected Tables

TEST CASE 1: Initial Validation

GOALS

- Verify user access to the UI

WHAT YOU'LL NEED

- Access to the *SaveYourData User Guide*.
- Credentials for the test user.
- Customer IT technician standing by

1. Log in to the SYD web interface using credentials for the test user. The SYD web interface is available via the HTTPS protocol at the IP address or hostname of the SYD VM. Replacing “XYZ” with the IP address or hostname of the SYD VM, connect to the SYD VM user interface in a web browser at: `https://XYZ/`

   success criteria

   1. You should see the main SYD screen as shown in the User Guide.
   2. No errors should appear.

   PASS___
   FAIL___

   Notes:

2. Navigate to the Audit screen by clicking Audit on the LHS panel.

   success criteria

   1. You should see the Audit screen as shown in the User Guide.
   2. An event showing your login should appear with the date-time stamp accurately reflected.

   PASS___
   FAIL___

   Notes:

TEST CASE 2: Create Profile, Connect and Run Ingest Job on Test Data

GOALS

- Verify connection and Ingest Functionality using included Test Data

WHAT YOU'LL NEED

- Access to the *SaveYourData User Guide*.
- Credentials for the test user.
- Customer IT technician standing by

1. Log in to the application using credentials for the test user, OR navigate back to the main screen.

   Follow the instructions in the *SaveYourData User Guide* to set up a profile for the pre-installed test database, a single-table schema with the following connection parameters:

   Database Type: `PostgreSQL`
   Hostname: `postgres`
   Port: `5432`
   Database: `financial`
   Username: `postgres`

   Click **CONNECT**.
   Enter "password" at the prompt for a password.
   success criteria
   1. The **CONNECT** option on the profile card now says **EDIT CONNECTION**.

PASS___ FAIL___

Notes:

2. Click **EDIT TABLES** on the test database profile.
3. On the Edit Tables screen, click **Save** without changing the checked tables.
4. Click **Ingest Databases** in the left-hand navigation pane to get back to the connection profiles.
5. Now click **START DATA INGEST**.

   success criteria
   1. A green popup notification saying "Job started successfully" should appear.
   2. You may see an "Ingestion in Progress" notification, although the test dataset should process very quickly.

PASS___ FAIL___

Notes:

6. Click **INGEST STATUS** on the test database profile.

   success criteria
   1. The status screen should show the job you just finished, *N* "minute(s) ago". Along with the volume of data processed.

PASS___ FAIL___

Notes:

TEST CASE 3: Connect and Run Ingest Job on Customer Data

GOALS

- Verify connection and ingest functionality using a database in the customer's data-center.

WHAT YOU'LL NEED

- Access to the *SaveYourData User Guide*.
- Credentials for an ingest database, preferably <50GB of data, with three or more tables.
- Customer IT technician to run the test with valid end-user credentials. NOTE: A customer IT tech is suggested because actual user and database credentials are being entered.

1. Ask the customer IT tech to log in to the application using credentials for a valid end-user.

   success criteria

   1. You should see the main SYD screen as shown in the User Guide.
   2. No errors should appear.

PASS___ FAIL___

Notes:

2. Click **ADD NEW DATABASE**, and set up a new connection profile as shown in the user guide.
3. Click **EDIT TABLES**. \
4. On the Edit Tables screen click **Save**, without changing the checked tables.
5. Click **Ingest Databases** in the left-hand navigation pane to get back to the connection profiles.
6. Now click **START DATA INGEST**.

   success criteria

   1. A green popup notification saying "Job started successfully" should appear.
   2. You may see an "Ingestion in Progress" notification, although the test dataset should process very quickly.

PASS___ FAIL___

Notes:

7. Click **INGEST STATUS** on the test database profile.

   success criteria

   2. The status screen should show the job you just finished, *N* "minute(s) ago". Along with the volume of data processed.

PASS___ FAIL___

Notes:

TEST CASE 4: User Login Scenarios

GOALS

- Verify authorized user access to the UI
- Verify the appropriate message is displayed based on LDAP configuration.

WHAT YOU'LL NEED

- Access to the *SaveYourData User Guide*.
- Credentials for the test user.
- Customer IT technician standing by

1. Log in (or ask the customer IT tech to) to the application using credentials for a user with LDAP credentials but not belonging to the SYD access group.
   success criteria
   1. A popup message appears on the login attempt saying "not part of <GROUP NAME>".
   2. You should NOT see the main SYD screen, or have any access to the application.
   3. No other errors should appear.

PASS___
FAIL___

Notes:

2. Log in to the application using credentials that are not in LDAP.
   success criteria
   1. A popup message appears on the login attempt saying "credentials are not valid."
   2. You should NOT see the main SYD screen, or have any access to the application.
   3. No other errors should appear.

PASS___
FAIL___

Notes:

TEST CASE 5: Connect and Run Ingest Job on Customer Data w/ Selected Tables

GOALS

- Verify connection and ingest functionality using a database in the customer's data-center with only a subset of its tables.

WHAT YOU'LL NEED

- Access to the *SaveYourData User Guide*.
- Credentials for an ingest database, preferably <50GB of data, with three or more tables.
- Customer IT technician to run the test with valid end-user credentials. NOTE: A customer IT tech is suggested because actual user and database credentials are being entered.

After login,

1. Click **ADD NEW DATABASE**, and set up a connection profile to the chosen database as shown in the user guide.
2. Click **EDIT TABLES**.
3. On the Edit Tables screen, select or deselect tables as desired, leaving at least one table selected.
4. Click **Save**.
5. Confirm that the profile shows the tables you selected.
6. Click **Ingest Databases** in the left-hand navigation pane to get back to the connection profiles.
7. Now click **START DATA INGEST**.

success criteria

1. A green popup notification saying "Job started successfully" should appear.
2. You may see an "Ingestion in Progress" notification.

PASS___ FAIL___

Notes:

8. Click **INGEST STATUS** on the test database profile.

success criteria

1. The status screen should show the job you just finished, *N* "minute(s) ago". Along with the volume of data processed.

PASS___ FAIL___

Notes:

Validation Test Plan: Running SaveYourData on VMware Infrastructure

This section of the test plan is adapted for use with SYD validation from a VMware template.

Goals

The following test plan outlines the specific steps and configurations to be tested in the validation process for VMware Infrastructure. The goals are to quantify the likely performance of SYD, and to confirm that support of such a configuration will not introduce new risk to our mutual customers.

Outline additional goals/deliverables for this exercise below:

Application Configuration and Requirements:

This section provides details on the specific components of the application. Where possible, a graphical representation of the application should also be provided, with mappings of various functional requirements to specific hardware hosts or virtual machines.

Application Component	**Software Version**	**Hardware Requirements**
Application Server	*Product Version, OS*	*CPU, RAM, DISK*
Database Server		
Web Server		

VMware Infrastructure Testing (Optionally)

In order to understand how [*ISV Product*] works with higher level functionality of VMware Infrastructure, we will perform one or more of the following tests.

VMware vMotion Testing
While running the workload, VMware vMotion is used to execute manual migration of the database virtual machine from one VMware ESX host to another. During this test, response time and transaction rates are monitored, and any observed slowdown in performance is measured. Five such operations will be executed, and averages are then determined across the five.

VMware Distributed Resource Scheduler (DRS) Testing
During this test, virtual machines are assigned to VMware ESX hosts such that the majority of the load will be on one host. In the first test, set the aggressiveness level of VMware DRS to “Conservative.” Start up the test and monitor how VMware DRS moves virtual machines across a cluster to balance the load. Monitor transaction throughput, as well as CPU utilization of the various hosts in the cluster. You should see CPU utilization balance across the hosts with little decrease in throughput. Now run the same test after setting the aggressiveness of VMware DRS to a substantially higher value. Typical deliverables are CPU charts from VMware vCenter Server that reflect the balancing of the load during these tests.

VMware High Availability (HA) Testing
Run the workload on a clustered resource pool, and then do a hard shutdown of one host of the configuration. Note how the virtual machines come up on another host in the cluster. Restart the application components if necessary, and functionally verify that the application is working again. Measure the time it takes for the virtual machine to start accepting work again. This work is often most interesting in the case of services that start automatically at reboot, such as Web servers, such that the application will automatically be ready to work as soon as the virtual machine is restarted on another host.

Wrap-up and Documentation

Where possible, results of testing are documented in a standard format provided by VMware. This will then be entered into a template for a deployment guide to be used as a starting place for documenting the results of the testing. Where possible, the virtual machines used for testing should be archived for possible use in the future. Any outstanding issues should be documented for further follow-up with either VMware or other ISV development teams.

Maintenance & Troubleshooting

5

Maintenance Notes

SYD requires regular maintenance to ensure appropriate security in the underlying technology stack.

Regular Maintenance

Three major architectural components SYD must be patched at regular intervals to maintain appropriate levels of security. Those are:

1. **VM Security:** The SYD VM requires regular Ubuntu Linux 16.04 LTS security patches, which may be applied by the installation vendor or the customer

2. **Container Security:** The Docker containers in the SYD VM require regular Ubuntu Linux 16.04 LTS security patches, which are provided by the installation vendor installing replacement Docker containers on the SYD VM.

3. **Database Security:** Any infrastructure databases used by the SYD must have regular security patches applied by the installation vendor or the customer.

Recommended Maintenance Schedule

Component	Maintenance Trigger	Maintenance Task	Notes
VMware	New patch release	Security Patch	
Docker Containers	New patch release	Security Patch	
RDBMS ingest data sources	New patch release	Security Patch	
JDBC Drivers	New driver release	JDBC driver update	

JDBC Driver Maintenance

SYD uses the JDBC standard API to connect to both ingest databases and the output database. JDBC drivers must be maintained according to vendor recommendations to prevent connectivity issues in production.

Update JDBC drivers in the VM.

- For Oracle:
  - (Version 12c) Oracle 12.1.0.2.0 JDBC 4.1 or higher within 12.x.x.x.x. See http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html for information.
  - (Version 11g) See http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html for information.
  - Sterling B2B Integrator supports JDBC Type-4 drivers on a single node of a database except with Oracle Real Application Clusters (RAC). You can use the JDBC Type-4 drivers to connect with multiple nodes of an Oracle RAC.

- For Microsoft SQL Server:
  - Microsoft SQL Server 2014 - Use SQL Server JDBC Driver 4.x
  - Microsoft SQL Server 2012 - Use SQL Server JDBC Driver 4.x
  - Microsoft SQL Server 2008 - Use SQL Server JDBC Driver 3.0
  - Regardless of the Microsoft SQL Server version, when using the Lightweight JDBC Adapter, use SQL Server JDBC Driver 4.x

  To obtain the driver, go to the Microsoft Download Center at http://www.microsoft.com/en-us/download/default.aspx and search for the required SQL Server JDBC driver version.
- For DB2®, see http://www.ibm.com/support/docview.wss?uid=swg21363866 for information.

TO UPDATE A JDBC DRIVER ON THE SYD VM

1. Download the desired *.jar driver file to your workstation.

2. Using SCP, upload the *.jar driver to the SYD VM.

   For example, to upload a driver file named `example.jar` from a Linux or Mac OS terminal to the `/home/ubuntu` directory of the SYD VM at the IP address 10.0.42.200, enter: `scp example.jar ubuntu@10.0.42.200:`

3. In super-user mode, copy the *.jar file to the `/opt/saveyourdata/drivers` directory.

Monitoring Log Files

SYD logs important system events formatted as JSON objects in flat files. End-users have the option to tail any or all of these files to identify critical system errors right when they happen.

There are two methods of viewing log files – using tail, or accessing Docker logs by specifying a container name.

Method 1 Using Tail

The first way to access logs is to view or tail any of the available log files explicitly. The raw log files are stored under the `/var/lib/docker/containers` directory, and each log line is a JSON object representing the log event. Run `find . -name '*-json.log'` from the /var/lib/docker/containers directory to enumerate available log files.

Method 2 Using Docker Logs

The second way to access and view logs is to run the following:

```
docker logs -f <container_name>
```

...where container_name is one of the four containers we started with `docker-compose up -d`, namely:

- `saveyourdata_webui_1`
- `saveyourdata_postgres_1`
- `saveyourdata_ingest-service_1`
- `saveyourdata_pgadmin_1`

This method outputs a log tail as one would normally expect, with one event per line, for example:

```
2018-07-26 01:14:13.762 [main] INFO
com.anonos.bigprivacy.Application - Started Application in 8.324
seconds (JVM running for 9.24)
```

(edited)

Log Format

When tailing raw log files, each line of the log file is a JSON object. This section describes how to decode and understand those JSON objects.

Object Keys

log	The string object with event details
stream	Describes the output channel used to post the entry
time	Time-stamp

Sample Log Entries

```
{
"log":"2018-07-26 01:49:19.361  INFO 6 --- [           main]
o.s.b.w.embedded.tomcat.TomcatWebServer  : Tomcat initialized
with port(s): 8080 (http)\n",
"stream":"stdout",
"time":"2018-07-26T01:49:19.361395461Z"
}

{
"log":"2018-07-26 01:49:19.409  INFO 6 --- [           main]
o.apache.catalina.core.StandardService   : Starting service
[Tomcat]\n",
"stream":"stdout",
"time":"2018-07-26T01:49:19.410110291Z"
}
{
"log":"2018-07-26 01:49:19.410  INFO 6 --- [           main]
org.apache.catalina.core.StandardEngine  : Starting Servlet
Engine: Apache Tomcat/8.5.31\n",

"stream":"stdout",

"time":"2018-07-26T01:49:19.41082241Z"
}
```

Troubleshooting SYD Application Issues

Because the installation process leverages off-the-shelf products like VMware and Docker, the approach to troubleshooting will differ greatly between two major components –production issues in the SYD application, and VMware issues that arise during installation.

This section addresses issues with SYD and connected systems only. The next section discusses resources for resolving issues directly related to VMware.

SYD relies heavily on network infrastructure for performance and proper functioning, so any fault in the chain of connectivity has the potential not only to interrupt proper functioning, but to manifest itself in a different component that it originated in.

Because of the diverse components involved in data flow, the general strategy for troubleshooting issues is to first isolate the component that is the root cause, prior to focusing on any one element. The customer should be encouraged to confirm basic network connectivity and functionality before opening a support ticket to accelerate the resolution process.

The following troubleshooting issues are addressed.

- Troubleshooting Login Issue 1: User Does not Belong to Group
- Troubleshooting Login Issue 2: User Credentials not Valid
- Troubleshooting Database Connection Issues

Troubleshooting Login Issue 1: User does not Belong to Group

SYD authenticates user login attempts against the customer's LDAP 3-compliant server. For access to SYD, users must have both a valid username and password, and belong to the group designated for authorized SYD users.

Symptoms

When the user attempts to log in, an error appears that user does not belong to the group required for SYD access.

Typical Root Cause(s)

The error indicates that,

- the system is connecting to the LDAP properly and reading user records.
- the system found the user credentials in LDAP, but that user is not a member of the SYD access group.

Resolution

If ALL users get this message,

- Verify they are members of the SYD access group in the customer's LDAP system.
    - If not, add them to that group and retry.
    - If they are... next step.

- Verify that the name of the SYD access group in the customer's LDAP system matches the name assigned to `DIRECTORY_GROUP` in the .env file.
    - If not, add them to that group and retry.
    - Otherwise...next step.

If only some users get this message,

- Verify that the affected users are members of the SYD access group in the customer's LDAP system.
    - If not, add them to that group and retry.
    - If they are... next step.

- Consider other restrictions in LDAP that may prevent the user form logging in as an artifact of other settings.

Troubleshooting Login Issue 2: User Credentials not Valid

SYD authenticates user login attempts against the customer's LDAP 3-compliant server. For access to SYD, users must have both a valid username and password, and belong to the group designated for authorized SYD users.

Symptoms

When the user attempts to log in, an error appears saying user credentials are not valid.

Typical Root Cause(s)

- The user is not configured in the LDAP directory used to authenticate SYD logins.
- The user appears to be configured in the LDAP directory used to authenticate SYD logins, but is entering the wrong username and/or password.

This error indicates that the user record is not found at all in the directory. It is unrelated to group membership.

Resolution

If ALL users get this message,

- Verify that both the username and password SYD uses to connect to LDAP is correct in the .env file. The strings assigned to `DIRECTORY_USER` and `DIRECTORY_PASSWORD` should match the correct username and password for directory access.
  - ○ If not, correct the entry or advise the user of their proper credentials and retry.
  - ○ If they are... next step.

If only some users get this message,

- Verify that both the username and password they enter at login matches what is in LDAP.
  - ○ If not, correct the entry or advise the user of their proper credentials and retry.
  - ○ If they are... next step.

- Consider other restrictions in LDAP that may prevent the user form logging in as an artifact of other settings.

Troubleshooting Database Connection Issues

Symptoms

An error appears when you click **CONNECT** on a connection profile, and the profile **CONNECT** link doesn't switch to **EDIT THIS CONNECTION**.

Typical Root Cause(s)

- LAN connectivity is not available
- The connection is available, but the connection string is not valid.
- The RDBMS platform is down.

Resolving connectivity issues

1. Check the database name and credentials and try test connection again.
2. Test access to the RDBMS target outside of the SYD interface.
   a) ping the endpoint to ensure that it is reachable.
   b) try a manual connection to the database with your DBMS application.

A failure to connect can be anywhere in the integrated system.

- Retry once to verify that it wasn't just a transient error that caused the issue.
- Verify that the database URL and username are correct.
- Verify that outside of SYD, you can see the URL.
- Verify that outside of SYD, you can connect to the database on login using your database management application

- SaveYourData uses JDBC to maintain standards-based compatibility with any data platform that supports JDBC. Nevertheless, platform-specific drivers must be installed prior to use.

The application fails to connect to an ingest data source after the user clicks the connect link.

The majority of failure scenarios involve driver incompatibility or LAN/WAN infrastructure external to SaveYourData.

Troubleshooting Installation and VMware Issues

If you experience issues during installation, VMware offers extensive troubleshooting literature online. Two VM troubleshooting articles are duplicated below for convenient reference during installation.

Additional resources are available on-line from the VMware knowledge base at https://kb.vmware.com/s/

Troubleshooting an ESX/ESXi host installation failure (1003649)

This article is adapted from VMware support material. It guides you through the process of troubleshooting the installation of an ESX/ESXi Server host.

Symptoms

- ESX/ESXi Server host installation stops with an error before completion.
- ESX/ESXi Server host installation fails to complete.
- Unable to complete ESX/ESXi Server installation.
- You may receive one or more of these errors during an installation or upgrade:
    - Your hardware is not compatible with ESX
    - Failed to load vmkernel: 0xbad0013
    - Failed to import upgrade. Error was: failed upload
    - Unable to connect to the remote server
- Encountered error FileIOError: The error data is: Filename - None Message - I/O Error (28) on file : [Errno 28] No space left on device Errno - 10 Description - Unable to create, write or read a file as expected.

Purpose

If you are experiencing these issues when using Update Manager to upgrade your ESX/ESXi host, see Troubleshooting issues when Update Manager cannot upgrade an ESX/ESXi host (1032626).

Resolution

Validate that each troubleshooting step below is true for your environment. Each step provides instructions or a link to a document, in order to eliminate possible causes and take corrective action as necessary. The steps are ordered in the most appropriate sequence to isolate the issue and identify the proper resolution. Do not skip a step.

- Verify that the minimum system requirements for installing an ESX/ESXi Server host have been met. For more information, see Minimum system requirements for installing ESX/ESXi (1003661).

- Verify that the hardware being installed on is certified. For more information, see Confirming ESX/ESXi host hardware (System, Storage, and I/O) compatibility (1003916).

- Verify that your hardware is working as expected. For more information, see Verifying your hardware is functioning correctly (1003690)
- Verify the problem is not being caused by the installation media. For more information, see Checking media integrity for ESX/ESXi host installations or upgrades (1003674).
- Verify that your firmware and BIOS are at the latest revision. For more information, see Checking your firmware and BIOS levels to ensure compatibility with ESX/ESXi (1037257).
- Verify a problem with your mouse and/or keyboard is not preventing you from completing the installation. For more information, see Troubleshooting issues when a mouse or keyboard stop functioning during installation or upgrade of an ESX/ESXi host system (1003868).
- When installing on top of an existing installation, you may need to wipe all data from the disks, and complete a fresh install using the ISO image if your installation continually fails. For more information, see Completing a clean installation using the ISO image if your installation or upgrade fails of ESX/ESXi (1004038).
- Installations can also fail if your ESX/ESXi host is full on disk space. For more information, see Investigating disk space on an ESX or ESXi host (1003564).
- If you are experiencing these issues when using Update Manager to upgrade your ESX/ESXi, see Troubleshooting issues where Update Manager cannot upgrade an ESX/ESXi host (1032626).

Note: If your problem persists after performing the steps in this article:

Gather the VMware Support Script Data. For more information, see Collecting Diagnostic Information in a VMware Virtual Infrastructure Environment (1003689) .

File a support request with VMware Support and note this KB Article ID in the problem description. For more information, see How to Submit a Support Request .

Related Information

For more information on installing or upgrading ESX/ESXi for your version, see the Installation Guides from the VMware Product Documentation.

- Investigating disk space on an ESX or ESXi host
- Minimum system requirements for installing ESXi/ESX
- Checking media integrity for ESX/ESXi host installations or upgrades
- Verifying your hardware is functioning correctly
- Troubleshooting issues when a mouse or keyboard stop functioning during installation or upgrade of an ESX/ESXi host system
- Confirming ESX/ESXi host hardware (System, Storage, and I/O) compatibility
- Completing a clean installation using the ISO image if your installation or upgrade fails of ESX/ESXi
- Overview of upgrading or migrating from vCenter Server 4.x to vCenter Server 5.0
- Troubleshooting when VMware Update Manager cannot upgrade an ESXi/ESX host
- Checking your firmware and BIOS levels to ensure compatibility with ESX/ESXi

Confirming ESX/ESXi host hardware (System, Storage, and I/O) compatibility (1003916)

Details

This article provides links to ESX/ESXi host Hardware Compatibility Documents (HCLs) so that you can verify your System, Storage, and I/O devices are on the VMware Certified and Supported Hardware Compatibility Lists.

Additionally, you can also verify if your systems and hardware require specific BIOS and firmware versions. If your System, Storage, or I/O devices are not listed or no specific BIOS or firmware versions are listed, contact your OEM or third party vendor for further verification and support.

Solution

VMware Hardware Compatibility Guides

Compare your hardware information with the VMware ESX/ESXi Server / Systems, I/O, and SAN Compatibility guides located at VMware Hardware Compatibility Guides.
Review these lists to verify correct system BIOS and firmware levels.
If you have any additional questions, contact your OEM hardware vendor directly to verify that your hardware has the recommended BIOS and firmware versions for all hardware installed in your system and storage devices.
Confirming Hardware Compatibility
To confirm hardware compatibility:

Connect to the ESX/ESXi host with an SSH session using root credentials. For more information, see Using ESXi Shell in ESXi 5.x and 6.0 (2004746) and Tech Support Mode for Emergency Support (1003677).

Run this command to display the system information:

```
# esxcfg-info | less -I
```

You see an output similar to:

```
    |----Product
Name..............................................ProLiant DL380
G6
    |----Vendor
Name...............................................Hewlett-
Packard
```

Identify the SCSI shared storage devices with the following command:

For ESX/ESXi 4.x, ESXi 5.x and 6.0, run the command:

```
# esxcfg-scsidevs -l | egrep -i 'display name|vendor'
```

You see output similar to:

```
Display Name: Local ServeRA Disk (mpx.vmhba0:C0:T0:L0)
  Vendor: ServeRA Model: 8k-l Mirror Revis: V1.0
```

Run this command to find additional peripherals and devices:

```
# lspci -vvv
```

You see an output similar to:

```
    02:0e.0 RAID bus controller: Dell Computer Corporation
PowerEdge Expandable RAID Controller 4E/SI/DI (rev 06)
    Subsystem: Dell Computer Corporation: Unknown device 016d
    Flags: bus master, stepping, 66Mhz, medium devsel, latency
64, IRQ 24
    Memory at d80f0000 (32-bit, prefetchable) [size=64K]
    Memory at dfdc0000 (32-bit, non-prefetchable) [size=256K]
    Expansion ROM at dfe00000 [disabled] [size=128K]
    Capabilities: [c0] Power Management version 2
    Capabilities: [d0] Message Signalled Interrupts: 64bit+
Queue=0/1 Enable-
    Capabilities: [e0] PCI-X non-bridge device.

    06:07.0 Ethernet controller: Intel Corporation 8254NXX
Gigabit Ethernet Controller (rev 05)
    Subsystem: Dell Computer Corporation: Unknown device 016d
    Flags: bus master, 66Mhz, medium devsel, latency 32, IRQ 25
    Memory at dfae0000 (32-bit, non-prefetchable) [size=128K]
    I/O ports at ecc0 [size=64]
    Capabilities: [dc] Power Management version 2
    Capabilities: [e4] PCI-X non-bridge device.

    07:08.0 Ethernet controller: Intel Corporation 8254NXX
Gigabit Ethernet Controller (rev 05)
    Subsystem: Dell Computer Corporation: Unknown device 016d
    Flags: bus master, 66Mhz, medium devsel, latency 32, IRQ 26
    Memory at df8e0000 (32-bit, non-prefetchable) [size=128K]
    I/O ports at dcc0 [size=64]
    Capabilities: [dc] Power Management version 2
    Capabilities: [e4] PCI-X non-bridge device.
```

Compare your hardware information to the VMware ESX/ESXi Server / Systems, I/O, and SAN Compatibility guides located at VMware Hardware Compatibility Guides.

APPENDIX A

Instructions for Initial Access to SaveYourData VM

A

Purpose

This appendix is intended to be used as a guide to gain access to the SaveYourData VM appliance on VMware. Typical console access is available with provided credentials as well as methods for importing OpenSSH public keys directly into the appliance from Windows or Linux on VMware ESXi hypervisors or VMware Workstation.

Quick Start

If you are already familiar with Microsoft Powershell and have a working setup, please jump ahead to the relevant section:

- Use the VMware PowerCLI on ESXi environments to add an SSH Public Key
- Use the VMware Workstation environments to add an SSH Public Key

Installation

The following sections address Powershell installation on various versions of both Windows and Linux.

Install Powershell on Windows

Powershell comes pre-installed on Windows. You may need to update your Powershell if you're running an out of date version.

On Windows 10, Windows Server 2016, or later

Powershell 5.1 comes pre-installed.

On Earlier Versions of Windows

Please follow Microsoft documentation on upgrading Powershell, found at `https://docs.microsoft.com/en-us/powershell/scripting/setup/ installing-windows-powershell?view=powershell-6`

Powershell 5.1 is available for download via WFM5.1 on:

- Windows 7 SP1
- Windows 8.1
- Windows Server 2008 R2 SP1
- Windows Server 2012
- Windows Server 2012 R2

Install Powershell on Linux

Ubuntu

1. Import the public repository GPG keys from Microsoft

```
curl https://packages.microsoft.com/keys/microsoft.asc | sudo
apt-key add -
```

2. Download the Apt sources for Microsoft Ubuntu Repository

```
curl https://packages.microsoft.com/config/ubuntu/16.04/prod.list
| sudo tee /etc/apt/sources.list.d/microsoft.list
```

3. Update your apt cache

```
sudo apt-get update
```

4. Install Powershell package

```
sudo apt-get install -y powershell
```

5. Start Powershell

```
powershell
```

RHEL / CentOS

1. Register the Microsoft RedHat repository

   ```
   curl https://packages.microsoft.com/config/rhel/7/prod.repo |
   sudo tee /etc/yum.repos.d/microsoft.repo
   ```

2. Install PowerShell

   ```
   sudo yum install -y powershell
   ```

3. Start PowerShell

   ```
   powershell
   ```

General Linux Install

To download and install the binaries manually, please follow the Microsoft instructions found at

```
https://docs.microsoft.com/en-us/powershell/scripting/setup/
installing-powershell-core-on-linux?view=powershell-6#binary-
archives
```

Install VMware PowerCLI on Powershell

Once you have installed Powershell on your Operating System, follow the instructions here.

1. Open Powershell.

   Always use `Run as Administrator` privileges on Windows or `sudo` on Linux. On Linux, the executable name is `pwsh`, so you would execute `sudo pwsh`

2. Check which version of Powershell is being used, and if applicable update it to the latest release (version 5.1 at time of writing). Powershell 5.1 is required for these installation steps to work.

   Check powershell version: `$psversiontable`

3. Download the VMware PowerCLI module:

   ```
   Save-Module -Name VMware.PowerCLI -Path $env:TEMP
   ```

   You can verify the files have downloaded by executing `dir $env:temp` in Powershell and locating the downloaded module.

4. Install the module into Powershell:

   ```
   Install-Module -Name VMware.PowerCLI
   ```

   You will be asked to trust the installation, select `[Y] Yes:`

5. Once completed, test the PowerCLI module is installed and available to PowerShell:

   ```
   Get-Module VMware.PowerCLI –ListAvailable
   ```

   You should see VMware.PowerCLI listed.

Use the VMware PowerCLI to add an SSH Public Key

1. Open Powershell

   On Windows: Open Powershell from the start menu by right-clicking and selecting `Run As Administrator`

   On Linux: `sudo pwsh`

2. Connect to your vCenter server via IP or DNS hostname:

   ```
   connect-viserver <VCENTER_IP_ADDRESS>
   ```

   Note: For the above command, depending on if you have a valid certificate or not, you may need to set powercli to ignore an invalid certificate with the following command:

   2a. `Set-PowerCLIConfiguration -InvalidCertificateAction Ignore`

   2b. You may be asked to log in, if you are enter your credentials for the vCenter server

   You are connected to vCenter when you see output from the shell. It displays the vCenter IP, the User connected and the port used.

3. Use the VM Name you assigned to get the VM object in Powershell:

   ```
   $vm = Get-VM -Name <VM_NAME>
   ```

4. Copy the OpenSSH public key file into the VM:

   ```
   Copy-VMGuestFile -Source <OPENSSH_PUBLIC_KEY_FILE> -Destination /home/
   ubuntu/.ssh/authorized_keys -VM $vm -GuestToLocal -GuestUser ubuntu -
   GuestPassword ubuntu
   ```

   *Note: The authorized_keys file must contain one public key per line in the OpenSSH format and not the PuTTY format. An example file with multiple keys might look like:

   ```
   ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDRCiaXYCUPstuhKT/
   i1RezbPDQooD5cTw2KDDhMAEe1vejJFIEkTO+dGg6YZ8z7Wu1YoW6wB4y1AREqnEXRrcoWG
   2uhLeOs+k7AWys5ruEBwQdasoD0UgIzArkiKMu1vwDGt8G5a5vcTLhZJjyI08bhmSM+8+H/
   rf7PL++YBd/mKf+YAzR6S92Sc2kxZ4n2w3+xlQxrViyNopUfsdYgCVuNXs3rO6ACexE/
   z4A7NWmB3FKyhm0NxlmcoN1yMo2eTShCLdUcP6bkemhG7QHx+XbZUejC66eDSOmrLWEHAJR
   PScXUF8stedFiOKB5A0EMHHovB06u9x9RFZRhBbidWRJ user1@server1
   ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQCVTGj8HBh9y6xuQE9jvySU/igu/0/
   icJV/Bf6ovwFy7Rjl4RJUu1UahNMODkFsCGOMCk1KhoHmmE52Zp0rmmev/pI4f4uFk/
   ziDkZj0rrkaZRQLG9YAn19dS1cvMNavuWBNzXY0zMd8o64shkUe2df5wnnYCS9wFloGUBcr
   Fmd7oWphiR1wW2/
   X+GgCY9BwIWncfYDW7de589YbSzyitg2dUrUaCr+1tShNrXGovXWvSLQhwQ8C5D0A53Lpsm
   cjDvwJkpzNvoN7LYqJFBebr253moUth67CkN9iaRF6rZHxY9q1aiMJ1VXyEnp8d8uWch0mP
   MZ0M68eW4tMD1V4EMJ user2@server1
   ```

Use the VMware Workstation to add an SSH Public Key

1. Using the VMware Workstation Console, login to the ubuntu desktop login prompt with the following credentials:

   Username: `ubuntu` Password: `ubuntu`

2. Create the .ssh directory and assign access permissions:

   ```
   mkdir ~/.ssh
   chmod 0700 ~/.ssh
   ```

3. Create authorized_keys file and paste in your openSSH public key using your preferred editor

   ```
   vim ~/.ssh/authorized_keys
   ```

   *Note: The authorized_keys file must contain one public key per line in the OpenSSH format and not the PuTTY format. An example file with multiple keys might look like:

   ```
   ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDRCiaXYCUPstuhKT/
   i1RezbPDQooD5cTw2KDDhMAEe1vejJFIEkTO+dGg6YZ8z7Wu1YoW6wB4y1AREqnEXRrcoWG
   2uhLeOs+k7AWys5ruEBwQdasoD0UgIzArkiKMu1vwDGt8G5a5vcTLhZJjyI08bhmSM+8+H/
   rf7PL++YBd/mKf+YAzR6S92Sc2kxZ4n2w3+xlQxrViyNopUfsdYgCVuNXs3rO6ACexE/
   z4A7NWmB3FKyhm0NxlmcoN1yMo2eTShCLdUcP6bkemhG7QHx+XbZUejC66eDSOmrLWEHAJR
   PScXUF8stedFiOKB5A0EMHHovB06u9x9RFZRhBbidWRJ user1@server1
   ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQCVTGj8HBh9y6xuQE9jvySU/igu/0/
   icJV/Bf6ovwFy7Rjl4RJUu1UahNMODkFsCGOMCk1KhoHmmE52
   ```

Information contained herein is subject to change and is not guaranteed to be error-free.