

IBM PowerAI Vision

Version 1.1.2

*PowerAI Vision Guide*

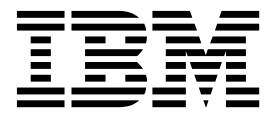




IBM PowerAI Vision

Version 1.1.2

*PowerAI Vision Guide*



**Note**

Before using this information and the product it supports, read the information in “Notices” on page 109.

This edition applies to IBM PowerAI Vision Version 1.1.2 and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2018.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>About this document</b> . . . . .	<b>v</b>
Highlighting . . . . .	v
ISO 9000. . . . .	v
<b>IBM PowerAI Vision overview</b> . . . . .	<b>1</b>
Use cases . . . . .	2
What's new. . . . .	6
IBM PowerAI Vision Trial . . . . .	6
<b>PowerAI Vision concepts</b> . . . . .	<b>11</b>
<b>Planning for PowerAI Vision</b> . . . . .	<b>13</b>
<b>License Management in IBM License Metric Tool</b> . . . . .	<b>15</b>
<b>Installing, upgrading, and uninstalling PowerAI Vision</b> . . . . .	<b>17</b>
Prerequisites for installing PowerAI Vision . . . . .	17
Installing PowerAI Vision stand-alone . . . . .	21
Installing PowerAI Vision with IBM Cloud Private . . . . .	24
Upgrading PowerAI Vision . . . . .	25
Uninstalling PowerAI Vision stand-alone . . . . .	27
<b>Checking the application and environment</b> . . . . .	<b>29</b>
Checking Kubernetes services status . . . . .	29
Checking Kubernetes node status . . . . .	30
Checking Kubernetes storage status . . . . .	35
Checking application deployment . . . . .	37
Checking system GPU status . . . . .	41
<b>Logging in to PowerAI Vision</b> . . . . .	<b>43</b>
<b>Working with the user interface</b> . . . . .	<b>45</b>
<b>Training and working with models</b> . . . . .	<b>47</b>
Creating and working with data sets . . . . .	47
Labeling objects . . . . .	48
Training a model . . . . .	50
Working with custom models . . . . .	52
Deploying a trained model . . . . .	60
PowerAI Vision REST APIs . . . . .	60
Testing a model . . . . .	61
Refining a model . . . . .	61
Automatically labeling objects . . . . .	61
Augmenting the data set . . . . .	63
Importing and exporting PowerAI Vision information . . . . .	64
Understanding metrics . . . . .	66
<b>Scenario: Detecting objects in a video</b> . . . . .	<b>69</b>
<b>Scenario: Classifying images</b> . . . . .	<b>75</b>
<b>Scenario: Detecting objects in images</b> . . . . .	<b>77</b>
<b>Administering PowerAI Vision.</b> . . . . .	<b>81</b>
Managing users . . . . .	81
Installing a new SSL certificate . . . . .	83
<b>PowerAI Vision Inference Server</b> . . . . .	<b>85</b>
Inference on embedded edge devices . . . . .	88
<b>Troubleshooting and contacting support</b> . . . . .	<b>91</b>
Troubleshooting known issues - PowerAI Vision standard install . . . . .	91
Troubleshooting known issues - PowerAI Vision Inference Server . . . . .	99
Gather PowerAI Vision logs and contact support . . . . .	100
Getting fixes from Fix Central . . . . .	102
Contacting IBM Support. . . . .	102
<b>Notices</b> . . . . .	<b>105</b>
Trademarks . . . . .	106
Terms and conditions for product documentation . . . . .	107
<b>Notices</b> . . . . .	<b>109</b>
Trademarks . . . . .	111



---

## About this document

This document provides you with information about installing and using IBM® PowerAI Vision to create a dataset that contains images or videos.

---

## Highlighting

The following highlighting conventions are used in this document:

<b>Bold</b>	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Bold highlighting also identifies graphical objects, such as buttons, labels, and icons that the you select.
<i>Italics</i>	Identifies parameters for actual names or values that you supply.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or text that you must type.

---

## ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.



---

## IBM PowerAI Vision overview

The IBM PowerAI Vision platform, built on cognitive infrastructure, is a new generation of video/image analysis platforms. The platform offers built-in deep learning models that learn to analyze images and video streams for classification and object detection.

PowerAI Vision includes tools and interfaces for anyone with limited skills in deep learning technologies. You can use PowerAI Vision to easily label images and videos that can be used to train and validate a model. The model can then be validated and deployed in customized solutions that demand image classification and object detection.

The following are the main features of PowerAI Vision:

### **Streamlined model training**

You can use existing models that are already trained as starting point to reduce the time required to train models and improve trained results.

### **Single-click model deployment**

After you create a training model, you can deploy an API with one click. You can then develop applications based on the model that you deployed.

### **Data set management and labeling**

You can manage both raw and labeled data.

### **Video object detection and labeling assistance**

Videos that you import can be scanned for objects and the objects can be automatically labeled.

## Architecture overview

The architecture of PowerAI Vision consists of hardware, resource management, deep learning computation, service management, and application service layers. Each layer is built around industry-standard technologies.



Table 1. Overview of the architecture layers

Architectural Layer	Description
Infrastructure Layer	Consists of hardware systems that support PowerAI Vision, including virtual machines (containers), accelerators (GPUs/FPGAs), storage systems, networks, and so on.
Resource Management Layer	Coordinates and schedules all computing resources.
Deep Learning Calculation Layer	Consists of deep learning algorithms, including data processing modules, model training modules, and inference modules.
Service Management Layer	Manages user projects in a graphical interface, including image preprocessing, data annotation management, training task management, model management, and API management.
Application Service Layer	Located on the top of the PowerAI Vision platform, it is responsible for managing all application-related services, including image labeling and preprocessing services, video annotation services, customized image classification services, and customized object detection services.

## Use cases

IBM PowerAI Vision includes these main use cases to demonstrate its capabilities:

### Static image classification

Determine whether an image belongs to one or more classes of images based on overall image contents. For example, determining the species of dog in an image.

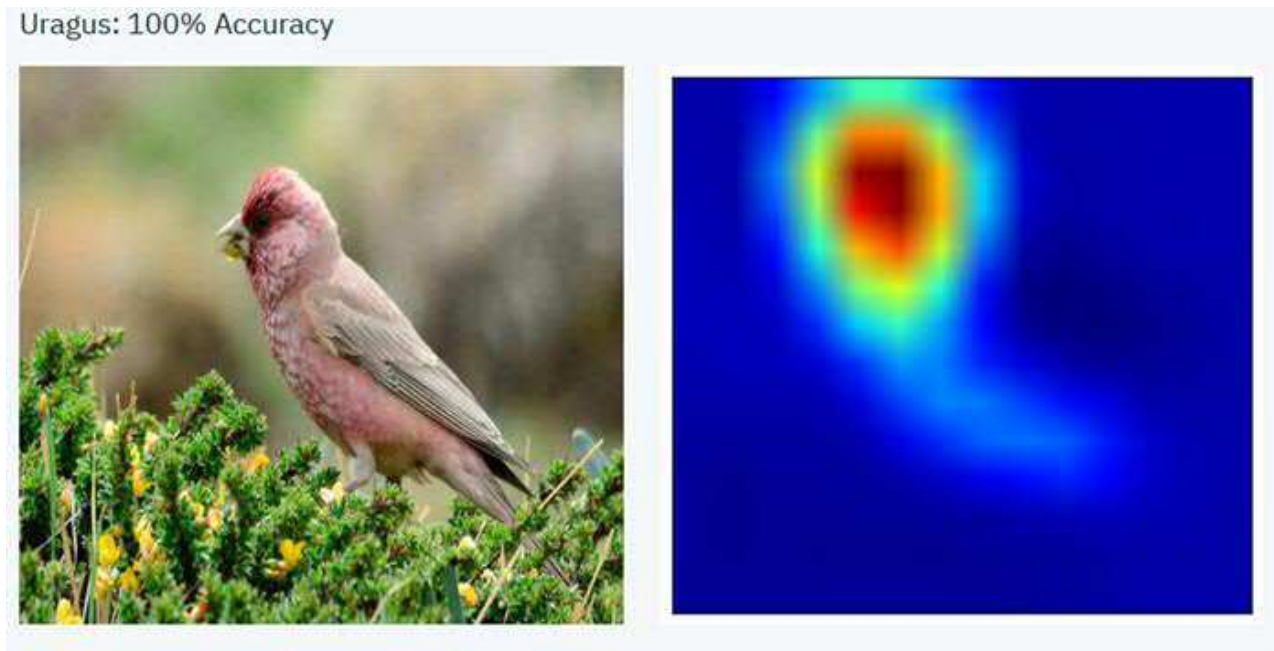


Figure 1. Detecting the overall contents of an image, based on custom training data

### Static image detection

Determine and label the contents of an image based on user-defined data labels. For example, finding and labeling all dogs in an image.

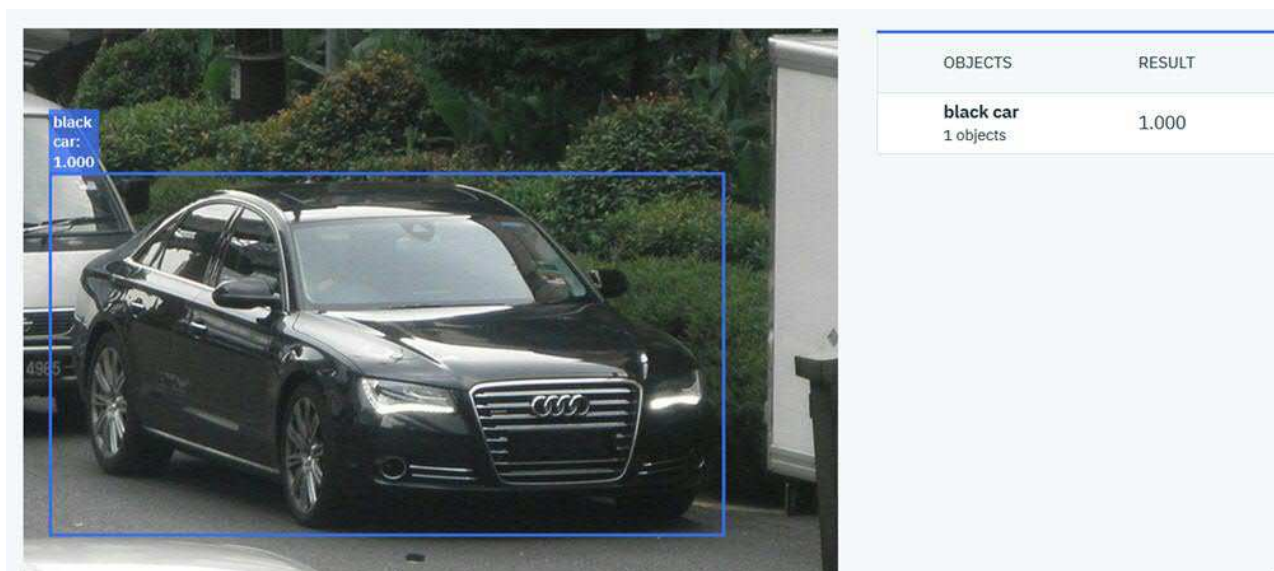


Figure 2. Detecting and labeling instances of objects within an image based on custom training data

### Video object detection

Determine and label the contents of an uploaded video or live video stream based on user-defined data labels. For example, finding and labeling all dogs in a video.



*Figure 3. Detecting and labeling instances of objects within an image based on custom training data*

#### **Auto label an image or video**

After deploying a model for object detection, you can improve its accuracy by using the Auto label function. This function improves the model's accuracy by quickly adding more data to the data set.

System-added tags are green, while manually added tags are blue.

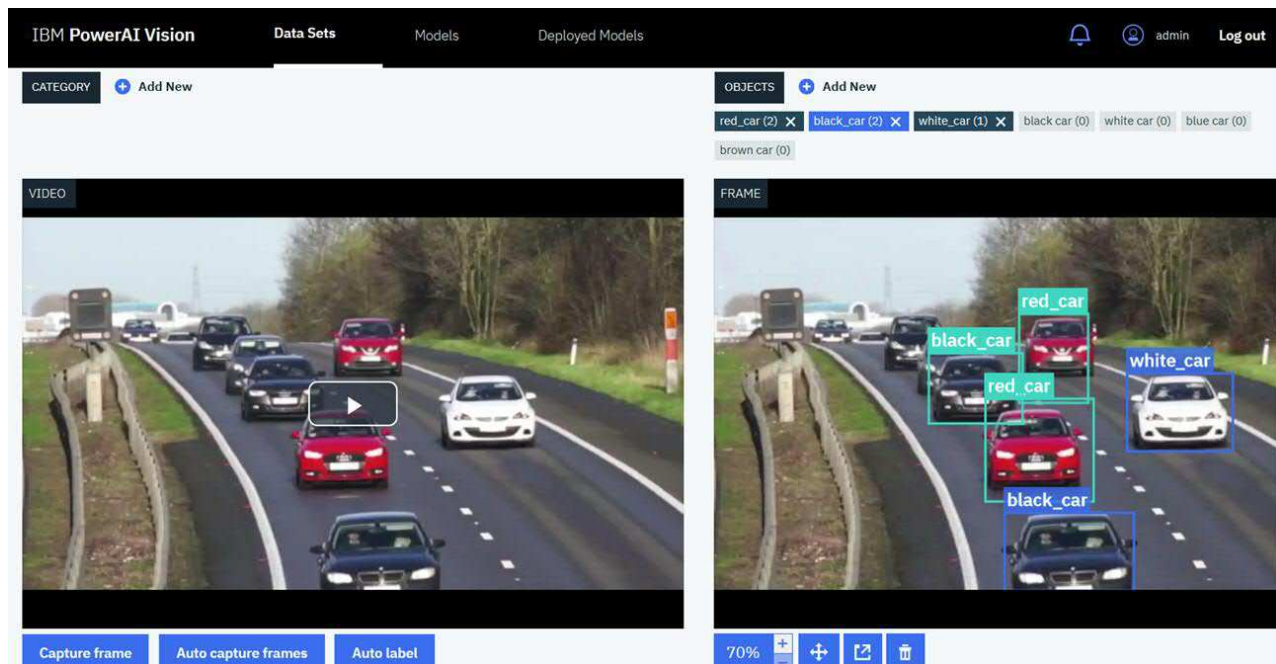


Figure 4. Auto labeled video

### Data augmentation

After deploying a model, you can improve the model by using data augmentation to add modified images to the data set, then retraining the model. *Data augmentation* is the use of filters, such as blur and rotate, to create new versions of existing images. When you use data augmentation, a new data set is created that contains all of the existing images, plus the newly generated images.

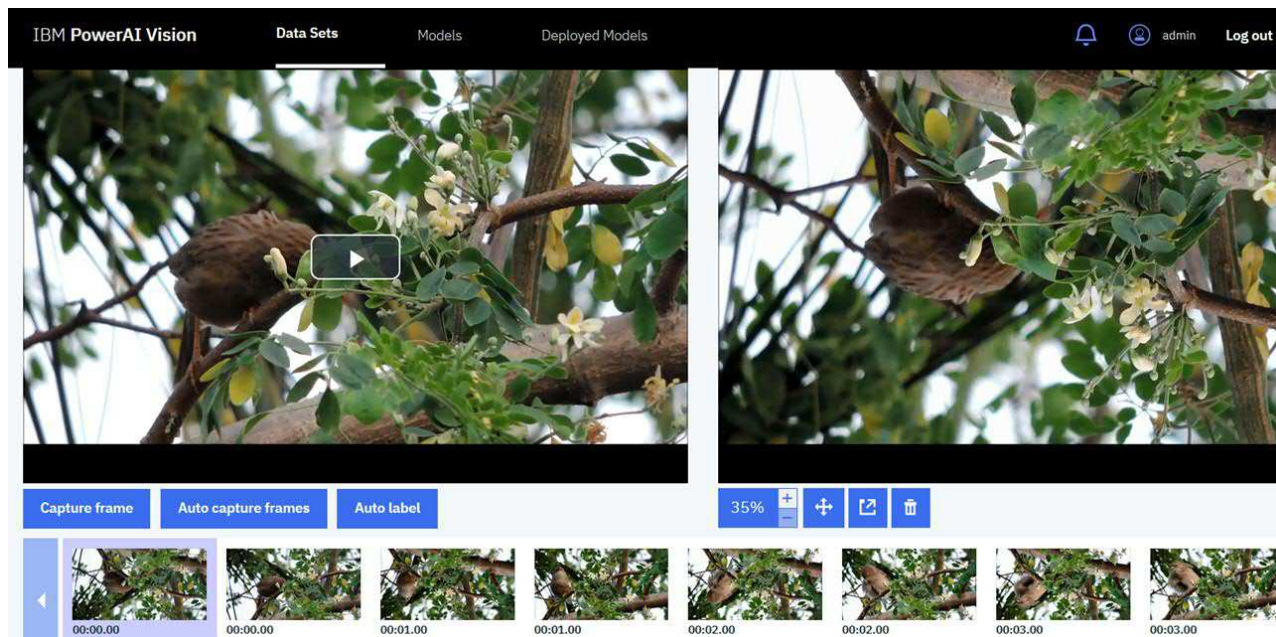


Figure 5. Augmented video

---

## What's new

The following functions, features, and support have been added for PowerAI Vision Version 1.1.2:

### Custom models

You can now import custom TensorFlow models into PowerAI Vision and use them for training. For details, see “Working with custom models” on page 52.

### Advanced metrics

New metrics are shown after a model is trained. For details, see “Understanding metrics” on page 66.

### Deployment to Enterprise FPGA accelerators (technology preview)

Models trained using the “Tiny Yolo v2” model for object detection can be deployed to enterprise-class FPGA accelerators for inferencing. See “Deploying a trained model” on page 60 for more details.

### Multiple category classification results

The results of image classification inferences now include multiple categories with the confidence levels for each category. See “Scenario: Classifying images” on page 75 for more information.

### Support for Ubuntu 18.04

The PowerAI Vision Training and Inference product is now supported on the Ubuntu 18.04 platform.

### Inference for Servers expanded coverage

The “Inference for Servers” platform coverage is now expanded with support for:

- Ubuntu 18.04
- CPU only deployments
- x86 deployment

See “PowerAI Vision Inference Server” on page 85 for more details.

### Support for new versions of IBM Cloud™ Private

PowerAI Vision can now be installed with IBM Cloud Private 3.1 and later.

### Non-administrator users

You can now create users that are not administrators. These users can view and manage only resources that they created. See “Managing users” on page 81 for details.

---

## IBM PowerAI Vision Trial

PowerAI Vision offers a trial version of the product. It has full functionality, but is not licensed for production use.

- “Installing the trial version”
- “What happens when the trial expires?” on page 7
- “Upgrading to the full version of PowerAI Vision” on page 8

## Installing the trial version

**Attention:** You cannot install PowerAI Vision stand-alone on the same system that has the following software installed:

- IBM Data Science Experience (DSX)
- IBM Cloud Private
- Any other Kubernetes based applications

1. You must complete the following installation prerequisites steps before you install PowerAI Vision.
  - a. Complete all steps in the “Prerequisites for installing PowerAI Vision” on page 17 topic.

- b. Your system must have a proper subscription and repository that provides you with updated packages. For information, see the Red Hat Subscription Manager documentation.
  - c. Turn on Extra Packages for Enterprise Linux (EPEL). For information, see the EPEL website.
2. Go to PowerAI Vision Trial download site. Download the .tar file and the .rpm files as instructed.
3. Untar the product tar file, and run the installation command for the platform you are installing on.

#### RHEL

```
sudo yum install ./<file_name>.rpm
```

#### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

4. From the directory that contains the downloaded tar file, run the appropriate script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

5. (RHEL only) Open ports for the firewall to access PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/sbin/firewall.sh
```

6. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

A user named admin is created with a password of passw0rd. For instructions to change these values, see “Managing users” on page 81.

You must read and accept the license agreement that is displayed before you can use PowerAI Vision. It can take several minutes to start PowerAI Vision. To check the status of the startup process, run this script:

```
sudo /opt/powerai-vision/bin/helm.sh status vision
```

In the output from the **helm.sh status vision** script, you can verify which PowerAI Vision components are available by locating the Deployment section and identifying that the AVAILABLE column has a value of 1 for each component. The following is an example of the output from the **helm.sh status vision** script that shows all components are available:

```
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
powerai-vision-mongodb              1         1         1             1         3h
powerai-vision-portal                1         1         1             1         3h
powerai-vision-postgres              1         1         1             1         3h
powerai-vision-taskanalyzer          1         1         1             1         3h
powerai-vision-ui                    1         1         1             1         3h
powerai-vision-video-nginx           1         1         1             1         3h
powerai-vision-video-portal          1         1         1             1         3h
powerai-vision-video-rabmq           1         1         1             1         3h
powerai-vision-video-redis           1         1         1             1         3h
powerai-vision-video-test-nginx       1         1         1             1         3h
powerai-vision-video-test-portal      1         1         1             1         3h
powerai-vision-video-test-rabmq       1         1         1             1         3h
powerai-vision-video-test-redis       1         1         1             1         3h
```

## What happens when the trial expires?

You can see how much time is left in the trial by reviewing the count down in the upper-right corner of the user interface. When the timed trial expires, the product will cease to work, including any running training, inference, import, or export operations. However, if you purchase a license, you will automatically regain access to all of your data sets, models, and so on.

If the trial expires and you want to purchase PowerAI Vision, follow the instructions in “Upgrading to the full version of PowerAI Vision.”

If the trial expires and you do not decide to purchase PowerAI Vision, follow these steps:

1. Remove previously installed images and data by running the following script:

```
sudo /opt/powerai-vision/bin/purge_image.sh
```

2. Remove PowerAI Vision by running the following command:

```
sudo yum remove powerai-vision
```

3. Delete the data directory by running the following command:

```
sudo rm -rf /opt/powerai-vision/
```

## Upgrading to the full version of PowerAI Vision

When you are ready to purchase PowerAI Vision, you can buy a license from PowerAI Vision Marketplace. Use one of these methods to upgrade to the full version. Your data is not deleted when the product is uninstalled. You will automatically regain access to all of your data sets, models, and so on.

1. Stop the current instance of PowerAI Vision by running the following script:

```
sudo /opt/powerai-vision/bin/powerai_vision_stop.sh
```

2. Obtain and install PowerAI Vision:

### Install PowerAI Vision from IBM Passport Advantage

- a. Download the product tar file from the IBM Passport Advantage website.
- b. Untar the product tar file, and run the installation command for the platform you are installing on.

#### RHEL

```
sudo yum install ./<file_name>.rpm
```

#### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

- c. From the directory that contains the downloaded tar file, run the appropriate script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

- d. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

### Install PowerAI Vision from AAS

- a. Download the product tar.gz file from Advanced Administration System (AAS). This system is also called Entitled Software Support (ESS).
- b. Unzip and untar the tar.gz file by running this command. The install files are extracted to powerai-vision-aas-1.1.2.1/.  

```
gunzip -c file_name.tar.gz | tar -xvf -
```
- c. Untar the product tar file, and run the installation command for the platform you are installing on.

#### RHEL

```
sudo yum install ./<file_name>.rpm
```

#### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

- d. From the directory that contains the extracted tar file, run this script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

- e. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

**Related concepts:**

“Uninstalling PowerAI Vision stand-alone” on page 27

You must uninstall PowerAI Vision stand-alone on your system, before you can install IBM Cloud Private, IBM Data Science Experience Local, or other Kubernetes-based applications.



---

## PowerAI Vision concepts

PowerAI Vision provides an easy to use graphical user interface (GUI) that you can use to quickly create computer vision-related artificial intelligence (AI) solutions.

You must be familiar with the following concepts before you can start using PowerAI Vision:

### **Data set**

A data set is a collection of images and videos that you uploaded to PowerAI Vision. An example of a data set would be images of cars.

### **Category**

A category is used to classify an image. The image can belong to only a single category. An example of a category for a data set that contains cars would be car manufacturers (Toyota, Honda, Chevy, and Ford).

### **Object**

An object is used to identify specific items in an image or specific frames in a video. You can label multiple objects in an image or a frame in a video. An example of objects in an image of cars might be wheel, headlights, and windshield.

**Model** A model is a set of tuned algorithms and that produces a predicted output. Models are trained based on the input that is provided by a data set to classify images or video frames, or find objects in images or video frames.



---

## Planning for PowerAI Vision

You must meet the software and hardware requirements and understand the supported file types before you can install PowerAI Vision.

- “Hardware requirements”
- “Software requirements”
- “Disk space requirements” on page 14
- “Supported web browsers” on page 14
- “Supported image types” on page 14
- “Supported video types” on page 14

### Hardware requirements

PowerAI Vision requires the following hardware:

- POWER8 S822LC (8335-GTB) or POWER9 AC922 with at least one NVIDIA NVLink capable GPU
- 128 GB of memory
- 40 GB of storage
- Ethernet network interface
- 40 GB of storage. See “Disk space requirements” on page 14 for details.
- If **Optimized for speed (tiny YOLO v2)** is selected when training the model, there are multiple options for deploying the model for testing. Deploying a model to a Xilinx FPGA requires the Xilinx Alveo U200 Accelerator card.

### Software requirements

You must install the following software before you install PowerAI Vision:

#### Linux

- Red Hat Enterprise Linux (RHEL) 7.5 (little endian).
- Ubuntu 16.04 or later.

#### NVIDIA CUDA

10.0 or later drivers. For information, see the NVIDIA CUDA Toolkit website.

#### Docker

- RHEL - Docker Version 1.13, or later, which is the version of Docker that is installed with RHEL 7.5.
- Ubuntu - Docker CE 18.06.01

### Networking requirements

Your environment must meet the following networking requirements:

- A default route must be specified on the host system.
- Docker 0 must be in a trusted firewall zone.
- IPv4 port forwarding must be enabled.
- IPv6 must be enabled.

## Disk space requirements

PowerAI Vision has the following storage requirements for the initial product installation and for the data sets that will be managed by the product.

### Standalone installation

- /var - The product installation requires at least 25 Gb of space in the /var file system for the product Docker images. PowerAI Vision also generates log information in this file system.

**Recommendation:** If you want to minimize the root (/) file system, make sure that /var has its own volume. The /var file system should have at least 50 Gb of space, more if additional applications are being run on the system that use this file system for log data and so on.

- /opt - PowerAI Vision data sets are stored in this file system. The storage needs will vary depending on the data sets and the contents - i.e., video data can require large amounts of storage.

**Recommendation:** If you want to minimize the root (/) file system, make sure that /opt has its own volume. The /opt file system should have at least 40 Gb of space, although this value might be more depending on your data sets.

### IBM Cloud Private installation

The PowerAI Vision product will use the configured persistent storage for the deployment, the requirements are documented in Installing PowerAI Vision with IBM Cloud Private.

## Supported web browsers

The following web browsers are supported:

- Google Chrome Version 60, or later
- Firefox Quantum 59.0, or later

## Supported image types

The following image formats are supported:

- JPEG
- PNG

## Supported video types

The following video formats are supported:

### Can be played in the PowerAI Vision GUI:

- Ogg Vorbis (.ogg)
- VP8 or VP9 (.webm)
- H.264 encoded videos with MP4 format (.mp4)

### Supported by API only:

- Matroska (.mkv)
- Audio Video Interleave (.avi)
- Moving Picture Experts Group (.mpg or .mpeg2)

### Not supported:

Videos that are encoded with the H.265 codec.

---

## License Management in IBM License Metric Tool

The IBM PowerAI Vision product is licensed per *Virtual Server* ("Learn about software licensing - Virtual Server"). When the product is installed, a software license metric (SLM) tag file is created to track usage with the IBM License Metric Tool.

The license metric tag is an XML file, with extension `.slmtag`. The IBM License Metric Tool discovers the license metric tag file and provides license consumption reports that, compared with license entitlements, allow IBM to verify license compliance. The tag file is human-readable and can therefore be interpreted by individuals for audit purposes.

The license metric tag file has a standard format and consists of two parts:

### Header information

Contains:

#### SchemaVersion

Identifies the schema version of the license metric tag file.

#### SoftwareIdentity

Identifies the software identity instance that provides license metric data. Contains:

- **Name**  
Name of the software identity - *IBM PowerAI Vision Training and Inference* or *IBM PowerAI Vision Inference for Servers*
- **PersistentId**  
Unique identifier of the software identity. For IBM PowerAI Vision 1.1.2, the assigned **PersistentId** is:
  - **IBM PowerAI Vision Training and Inference** - c5bc710c34df4e6c9e683fdb621eefb
  - **IBM PowerAI Vision Inference for Servers** - 572a77adfa984ccb89c52a88a4af85b5
- **InstanceId**  
Identifies the instance of the software identity that provides metrics by the path of the software for which *SLMTag* is generated - `/opt/powerai-vision`.

### Metrics information

IBM PowerAI Vision 1.1.2 is licensed per Virtual Server, so the values are:

- **Type** - `VIRTUAL_SERVER`
- **Period** - **StartTime** is the time of install/deploy, **EndTime** is set to date '9999-12-31' so that the IBM License Metric Tool will understand that it as a perpetual license.



---

## Installing, upgrading, and uninstalling PowerAI Vision

Use the information in these topics to work with the product installation. You can install PowerAI Vision by using the command line (stand-alone) or by using IBM Cloud Private.

Only the most current level of each release of IBM PowerAI Vision should be installed, where version numbers are in the format *version.release.modification*.

After installing PowerAI Vision, you can optionally change the SSL certificate by following the steps in this topic: “Installing a new SSL certificate” on page 83.

---

### Prerequisites for installing PowerAI Vision

Before you can install either PowerAI Vision stand-alone or PowerAI Vision with IBM Cloud Private, you must configure Red Hat Enterprise Linux (RHEL), enable the Fedora Extra Packages for Enterprise Linux (EPEL) repository, and install NVIDIA CUDA drivers.

**Note:** Neither IBM PowerAI nor PowerAI Enterprise are required for running PowerAI Vision.

- “Red Hat Enterprise Linux operating system and repository setup”
- “Ubuntu operating system and repository setup” on page 18
- “NVIDIA Components: IBM POWER9 specific udev rules (Red Hat only)” on page 18
- “GPU driver (Red Hat only)” on page 19
- “GPU driver, docker, nvidia-docker2 (Ubuntu only)” on page 20

### Red Hat Enterprise Linux operating system and repository setup

1. Enable common, optional, and extra repo channels.

IBM POWER8®:

```
sudo subscription-manager repos --enable=rhel-7-for-power-le-optional-rpms
```

```
sudo subscription-manager repos --enable=rhel-7-for-power-le-extras-rpms
```

```
sudo subscription-manager repos --enable=rhel-7-for-power-le-rpms
```

IBM POWER9™:

```
sudo subscription-manager repos --enable=rhel-7-for-power-9-optional-rpms
```

```
sudo subscription-manager repos --enable=rhel-7-for-power-9-extras-rpms
```

```
sudo subscription-manager repos --enable=rhel-7-for-power-9-rpms
```

2. Install packages needed for the installation.

```
sudo yum -y install wget nano bzip2
```

3. Enable Fedora Project EPEL (Extra Packages for Enterprise Linux) repo:

```
wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

```
sudo rpm -ihv epel-release-latest-7.noarch.rpm
```

4. Load the latest kernel or do a full update:

- Load the latest kernel:

```
sudo yum update kernel kernel-tools kernel-tools-libs kernel-bootwrapper
```

```
reboot
```

- Do a full update:

```
sudo yum update
```

```
sudo reboot
```

**Important:** RHEL 7.6 was released at the end of October, but is not yet supported by PowerAI. Running just **yum update** might upgrade a 7.5 system to 7.6. In order to avoid this, customers with a standard RHEL subscription might use:

```
sudo subscription-manager release --set=7.5
```

Customers should consult Red Hat if they're unsure how to avoid unintended upgrade.

## Ubuntu operating system and repository setup

1. Install packages needed for the installation

```
sudo apt-get install -y wget nano apt-transport-https ca-certificates curl software-properties-common
```

2. Load the latest kernel

```
sudo apt-get install linux-headers-$(uname -r)
```

```
sudo reboot
```

3. Or do a full update

```
sudo apt-get update
```

```
sudo apt-get dist-upgrade
```

```
sudo reboot
```

## NVIDIA Components: IBM POWER9 specific udev rules (Red Hat only)

Before you install the NVIDIA components, the udev Memory Auto-Onlining Rule must be disabled for the CUDA driver to function properly. To disable it, follow these steps:

1. Copy the `/lib/udev/rules.d/40-redhat.rules` file to the directory for user overridden rules.

```
sudo cp /lib/udev/rules.d/40-redhat.rules /etc/udev/rules.d/
```

2. Edit the `/etc/udev/rules.d/40-redhat.rules` file.

```
sudo nano /etc/udev/rules.d/40-redhat.rules
```

3. Comment out the following line and save the change:

```
SUBSYSTEM=="memory", ACTION=="add", PROGRAM="/bin/uname -p",
```

```
RESULT!="s390*", ATTR{state}=="offline", ATTR{state}="online"
```

4. Optionally, delete the first line of the file, since the file was copied to a directory where it cannot be overwritten.

```
# do not edit this file, it will be overwritten on update
```

5. Restart the system for the changes to take effect.

```
sudo reboot
```

## Remove previously installed CUDA and NVIDIA drivers (Red Hat only)

Before installing CUDA 10, uninstall any previous installations of CUDA and NVIDIA drivers. Follow these steps:

1. Run the following command:

```
sudo yum remove libglvnd*
```

Note that removing `libglvnd*` also uninstalls the Nvidia drivers.

2. Verify that the drivers were uninstalled:

```
sudo yum list installed | grep cuda
```

If a previous local repo is found, uninstall it:

```
sudo rpm -e cuda-repo-rhel7-9-2-local-9.2.148-1.ppc64le
```

3. Finally run

```
sudo yum clean all
```

If you get a message to remove the yum cache, run

```
sudo rm -rf /var/cache/yum
```

## GPU driver (Red Hat only)

Install the driver by following these steps:

1. Download NVIDIA CUDA 10

- Select **Operating System**: Linux.
- Select **Architecture**: ppc64le.
- Select **Distribution**: RHEL.
- Select **Version**: 7.
- Select **Installer Type**: rpm (network).
- Follow the *Linux on POWER* installation instructions in the CUDA Quick Start Guide, including the steps that describe how to set up the CUDA development environment by updating PATH and LD\_LIBRARY\_PATH.

2. Install CUDA and the GPU driver.

**Note:** For AC922 systems: OS and system firmware updates are required before you install the latest GPU driver.

a. Install the CUDA Base repository rpm.

1) Install the RPM by running the **sudo yum install cuda-repo-rhel7\*.ppc64le.rpm** command.

b. Run this command to install CUDA and the GPU driver:

```
sudo yum install cuda
```

c. Ensure that nvidia-persistenced is enabled. For instructions, see NVIDIA Persistence Daemon in the PowerAI system setup topic.

d. Restart to activate the driver.

e. Verify that the CUDA drivers are installed by running the `/usr/bin/nvidia-smi` application.

### Example output

```
# nvidia-smi
Thu Sep  6 17:57:38 2018

+-----+
| NVIDIA-SMI 396.15                  Driver Version: 396.15          |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|  0   Tesla P100-SXM2...    Off   | 00000002:01:00:0 Off |             0         |
| N/A   30C   P0      38W / 300W | 747MiB / 16280MiB |      0%    Default   |
+-----+-----+
|  1   Tesla P100-SXM2...    Off   | 00000003:01:00:0 Off |             0         |
| N/A   31C   P0      31W / 300W | 10MiB / 16280MiB |      0%    Default   |
+-----+-----+
|  2   Tesla P100-SXM2...    Off   | 0000000A:01:00:0 Off |             0         |
| N/A   31C   P0      32W / 300W | 10MiB / 16280MiB |      0%    Default   |
+-----+-----+
|  3   Tesla P100-SXM2...    Off   | 0000000B:01:00:0 Off |             0         |
| N/A   30C   P0      30W / 300W | 10MiB / 16280MiB |      0%    Default   |
+-----+-----+

+-----+
| Processes:                         GPU Memory |
|   GPU       PID    Type    Process name                     Usage |
+-----+-----+
|    0      124097     C   /usr/bin/python                     729MiB |
+-----+-----+
```

For help understanding the output, see “Checking system GPU status” on page 41.

For more information, see the *Linux POWER®* installation instructions in the *CUDA Quick Start Guide*. It includes steps for setting up the CUDA development environment by updating PATH and LD\_LIBRARY\_PATH.

## NVIDIA Persistence Daemon (Red Hat only)

The NVIDIA Persistence Daemon may be automatically started for POWER9 installations. Check that it is running with the following command:

```
systemctl status nvidia-persistenced
```

If it is not active, run the following command:

```
sudo systemctl enable nvidia-persistenced
```

## GPU driver, docker, nvidia-docker2 (Ubuntu only)

To run PowerAI within docker containers, only the GPU driver needs to be installed on the host.

1. Download NVIDIA driver 410.72 from <http://www.nvidia.com/Download/index.aspx>.

- Select **Product Type:** Tesla
- Select **Product Series:** P-Series
- Select **Product:** Tesla P100
- Select **Operating System:** Linux POWER LE Ubuntu 18.04 (If Linux POWER LE Ubuntu 18.04 is not available, click **Show all Operating Systems**)
- Select **CUDA Toolkit:** 10.0
- Click **Search** to go to the download link

2. Install the GPU driver repository deb package and cuda-drivers.

```
sudo dpkg -i nvidia-driver-local-repo-ubuntu1804-410.72_1.0-1_ppc64el.deb
```

```
sudo apt-get update
```

```
sudo apt-get install cuda-drivers
```

3. Edit the nvidia-persistenced file.

```
sudo systemctl edit --full nvidia-persistenced
```

Replace the contents with the following lines:

[Unit]

Description=NVIDIA Persistence Daemon

Wants=syslog.target

[Service]

Type=forking

PIDFile=/var/run/nvidia-persistenced/nvidia-persistenced.pid

Restart=always

ExecStart=/usr/bin/nvidia-persistenced --verbose

ExecStopPost=/bin/rm -rf /var/run/nvidia-persistenced

TimeoutSec=300

[Install]

WantedBy=multi-user.target

4. Set nvidia-persistenced to start at boot  
sudo systemctl enable nvidia-persistenced
5. Restart your system.
6. For Ubuntu platforms, a Docker runtime must be installed. If there is no Docker runtime installed yet, install Docker-CE on Ubuntu.  
curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo apt-key add -  
  
sudo add-apt-repository "deb [arch=ppc64el] https://download.docker.com/linux/ubuntu bionic stable"  
  
sudo apt-get update  
  
sudo apt-get install docker-ce

---

## Installing PowerAI Vision stand-alone

You use the command line to install PowerAI Vision stand-alone.

### PowerAI Vision stand-alone installation prerequisites

You must complete the following installation prerequisites steps before you install PowerAI Vision.

1. Complete all steps in the “Prerequisites for installing PowerAI Vision” on page 17 topic.
2. Your system must have a proper subscription and repository that provides you with updated packages. For information, see the Red Hat Subscription Manager documentation.
3. Turn on Extra Packages for Enterprise Linux (EPEL). For information, see the EPEL website.

**Attention:** You cannot install PowerAI Vision stand-alone on the same system that has the following software installed:

- IBM Data Science Experience (DSX)
- IBM Cloud Private
- Any other Kubernetes based applications
- “Install PowerAI Vision from IBM Passport Advantage”
- “Install PowerAI Vision from AAS ” on page 22
- “Install PowerAI Vision trial mode” on page 23

### Install PowerAI Vision from IBM Passport Advantage

To install PowerAI Vision stand-alone, complete the following steps:

1. Download the product tar file from the IBM Passport Advantage website.
2. Untar the product tar file, and run the installation command for the platform you are installing on.

#### RHEL

```
sudo yum install ./<file_name>.rpm
```

#### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

3. From the directory that contains the downloaded tar file, run the appropriate script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

4. (RHEL only) Open ports for the firewall to access PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/sbin/firewall.sh
```

5. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

A user named admin is created with a password of passw0rd. For instructions to change these values, see “Managing users” on page 81.

You must read and accept the license agreement that is displayed before you can use PowerAI Vision. It can take several minutes to start PowerAI Vision. To check the status of the startup process, run this script:

```
sudo /opt/powerai-vision/bin/helm.sh status vision
```

In the output from the **helm.sh status vision** script, you can verify which PowerAI Vision components are available by locating the Deployment section and identifying that the AVAILABLE column has a value of 1 for each component. The following is an example of the output from the **helm.sh status vision** script that shows all components are available:

```
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
powerai-vision-mongodb             1        1        1            1          3h
powerai-vision-portal              1        1        1            1          3h
powerai-vision-postgres            1        1        1            1          3h
powerai-vision-taskanalyzer        1        1        1            1          3h
powerai-vision-ui                  1        1        1            1          3h
powerai-vision-video-nginx         1        1        1            1          3h
powerai-vision-video-portal        1        1        1            1          3h
powerai-vision-video-rabmq         1        1        1            1          3h
powerai-vision-video-redis         1        1        1            1          3h
powerai-vision-video-test-nginx    1        1        1            1          3h
powerai-vision-video-test-portal   1        1        1            1          3h
powerai-vision-video-test-rabmq    1        1        1            1          3h
powerai-vision-video-test-redis    1        1        1            1          3h
```

6. Install any available fix packs. For instructions see “Getting fixes from Fix Central” on page 102.

## Install PowerAI Vision from AAS

1. Download the product tar.gz file from Advanced Administration System (AAS). This system is also called Entitled Software Support (ESS).

2. Unzip and untar the tar.gz file by running this command. The install files are extracted to powerai-vision-aas-1.1.2.1/.

```
gunzip -c file_name.tar.gz | tar -xvf
```

3. Untar the product tar file, and run the installation command for the platform you are installing on.

### RHEL

```
sudo yum install ./<file_name>.rpm
```

### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

4. From the directory that contains the extracted tar file, run this script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

5. (RHEL only) Open ports for the firewall to access PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/sbin/firewall.sh
```

6. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

A user named admin is created with a password of passw0rd. For instructions to change these values, see “Managing users” on page 81.

You must read and accept the license agreement that is displayed before you can use PowerAI Vision. It can take several minutes to start PowerAI Vision. To check the status of the startup process, run this script:

```
sudo /opt/powerai-vision/bin/helm.sh status vision
```

In the output from the **helm.sh status vision** script, you can verify which PowerAI Vision components are available by locating the Deployment section and identifying that the AVAILABLE column has a value of 1 for each component. The following is an example of the output from the **helm.sh status vision** script that shows all components are available:

```
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
powerai-vision-mongodb              1        1        1            1          3h
powerai-vision-portal                1        1        1            1          3h
powerai-vision-postgres              1        1        1            1          3h
powerai-vision-taskanalyzer          1        1        1            1          3h
powerai-vision-ui                    1        1        1            1          3h
powerai-vision-video-nginx           1        1        1            1          3h
powerai-vision-video-portal          1        1        1            1          3h
powerai-vision-video-rabmq           1        1        1            1          3h
powerai-vision-video-redis           1        1        1            1          3h
powerai-vision-video-test-nginx       1        1        1            1          3h
powerai-vision-video-test-portal      1        1        1            1          3h
powerai-vision-video-test-rabmq       1        1        1            1          3h
powerai-vision-video-test-redis       1        1        1            1          3h
```

7. Install any available fix packs. For instructions see “Getting fixes from Fix Central” on page 102.

## Install PowerAI Vision trial mode

1. Go to PowerAI Vision Trial download site. Download the .tar file and the .rpm files as instructed.
2. Untar the product tar file, and run the installation command for the platform you are installing on.

### RHEL

```
sudo yum install ./<file_name>.rpm
```

### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

3. From the directory that contains the downloaded tar file, run the appropriate script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

4. (RHEL only) Open ports for the firewall to access PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/sbin/firewall.sh
```

5. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

A user named admin is created with a password of passw0rd. For instructions to change these values, see “Managing users” on page 81.

You must read and accept the license agreement that is displayed before you can use PowerAI Vision. It can take several minutes to start PowerAI Vision. To check the status of the startup process, run this script:

```
sudo /opt/powerai-vision/bin/helm.sh status vision
```

In the output from the **helm.sh status vision** script, you can verify which PowerAI Vision components are available by locating the Deployment section and identifying that the AVAILABLE

column has a value of 1 for each component. The following is an example of the output from the **helm.sh status vision** script that shows all components are available:

```
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
powerai-vision-mongodb              1        1        1            1          3h
powerai-vision-portal                1        1        1            1          3h
powerai-vision-postgres              1        1        1            1          3h
powerai-vision-taskanalyzer          1        1        1            1          3h
powerai-vision-ui                    1        1        1            1          3h
powerai-vision-video-nginx           1        1        1            1          3h
powerai-vision-video-portal          1        1        1            1          3h
powerai-vision-video-rabmq           1        1        1            1          3h
powerai-vision-video-redis           1        1        1            1          3h
powerai-vision-video-test-nginx      1        1        1            1          3h
powerai-vision-video-test-portal     1        1        1            1          3h
powerai-vision-video-test-rabmq      1        1        1            1          3h
powerai-vision-video-test-redis      1        1        1            1          3h
```

### Related concepts:

“Logging in to PowerAI Vision” on page 43

Follow these steps to log in to PowerAI Vision.

## Installing PowerAI Vision with IBM Cloud Private

If you have more than one IBM Power Systems server available, you can use IBM Cloud Private 2.1.0.3 or 3.1.0 to install a single instance of PowerAI Vision that has access to all the Power Systems GPUs across the entire cluster.

If you have only a single IBM Power Systems server and do not have an existing IBM Cloud Private environment, you should use the PowerAI Vision stand-alone process. For more information, see the “Installing PowerAI Vision stand-alone” on page 21 topic.

To install PowerAI Vision with IBM Cloud Private, complete the following steps:

### Notes:

- If IBM Cloud Private is already installed and configured in your environment, you can go to step 4.
- The links to IBM Cloud Private go to the 3.1.0 Knowledge Center. To go to a different version, click the link, then click **Change version**.

1. Install IBM Cloud Private. For more information, see the Installing IBM Cloud Private topic.
2. Install the IBM Cloud CLI. For more information, see the Install IBM Cloud CLI topic.
3. Authenticate to your master node in your IBM Cloud Private environment. For more information, see the Configuring authentication for the Docker CLI topic.

To log in to the IBM Cloud Private cluster, run the appropriate command:

- In an IBM Cloud Private 2.1.0 environment, run:  

```
bx pr login -a https://<cluster-domain-name>:8443/ --skip-ssl-validationCopy
```
- In an IBM Cloud Private 3.1.0 environment, run:  

```
cloudctl login -a https://<cluster-domain-name>:8443/ --skip-ssl-validationCopy
```

4. Download the appropriate tar file from IBM Passport Advantage.
5. Untar the `powerai-vision-ppa-1.1.2.0.tar` tar file. It contains install packages for the standalone product, as well as the tar file with the containers that must be loaded for the IBM Cloud Private installation.
6. To make PowerAI Vision available in IBM Cloud Private catalog, run the appropriate command:
  - IBM Cloud Private 2.1.0.3:  

```
bx pr load-ppa-archive --archive file_name.tar [--clustername <cluster_CA_domain>]
```
  - IBM Cloud Private 3.1.0 or later:

```
cloudctl catalog load-archive --archive file_name.tar
```

Where `cluster_CA_domain` is the certificate authority (CA) domain. If you did not specify a CA domain, the default value is `mycluster.icp`.

7. Review the Chart README for PowerAI Vision carefully. It documents prerequisites, requirements, and limitations of PowerAI Vision in IBM Cloud Private.
8. Verify that you have a minimum of 40 GB of persistent storage. If your IBM Cloud Private installation has dynamic provisioned storage, you can use it for your 40 GB of persistent storage. To manually create persistent volumes in IBM Cloud Private, see the [Creating a Persistent Volume](#) topic. After you create the persistent volume, you must make the volume sharable across all nodes in the cluster.

**Note:** Do not use `HostPath` for the persistent storage unless you have only one node in your cluster. See [Creating a Persistent Volume](#) in the IBM Cloud Private documentation for details.

9. To install PowerAI Vision from the IBM Cloud Private catalog, from the navigation menu select **Catalog > Helm Charts**.
10. In the search box, enter `vision` and click **powerai-vision**. Review the information.
11. Click **Configure** and enter information for the **Release name** and the **Namespace** fields. The default user name is `admin` and the default password is `passw0rd`. For instructions to change these values, see “Managing users” on page 81. For information about namespaces, see [Namespaces](#) in the IBM Cloud Private Knowledge Center.
12. Click **Install**.
13. For information about accessing PowerAI Vision, see [Logging into PowerAI Vision](#).

**Important:** NFS volumes should have the “`no_root_squash`” flag set in `/etc/exports`:

```
/var/nfs *(rw,no_root_squash,no_subtree_check)
```

---

## Upgrading PowerAI Vision

When upgrading to the latest version of PowerAI Vision, your data from the previous release will not be lost, as long as you are upgrading to the same type of install. For example; from the stand-alone version to the stand-alone version. However, you will need to delete and redeploy any deployed models after upgrading.

- Upgrade the stand-alone version
- Upgrade PowerAI Vision with IBM Cloud Private

### Upgrade the stand-alone version

1. Stop the current instance of PowerAI Vision by running the following script:

```
sudo /opt/powerai-vision/bin/powerai_vision_stop.sh
```

2. Obtain and install PowerAI Vision:

#### Install PowerAI Vision from IBM Passport Advantage

- a. Download the product tar file from the IBM Passport Advantage website.
- b. Untar the product tar file, and run the installation command for the platform you are installing on.

##### RHEL

```
sudo yum install ./<file_name>.rpm
```

##### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

You will be prompted to accept the upgrade of the product if you are running an interactive install.

- c. From the directory that contains the downloaded tar file, run the appropriate script as root or with `sudo` privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

- d. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

A user named admin is created with a password of passw0rd. This user owns all data that was created in the previous version. If other users existed in the previous version, they will no longer exist. For instructions to change the created user name or password, and to learn how to create new users, see “Managing users” on page 81.

- e. Install any available fix packs. For instructions see “Getting fixes from Fix Central” on page 102.

### Install PowerAI Vision from AAS

- a. Download the product tar.gz file from Advanced Administration System (AAS). This system is also called Entitled Software Support (ESS).
- b. Unzip and untar the tar.gz file by running this command. The install files are extracted to powerai-vision-aas-1.1.2.1/.  

```
gunzip -c file_name.tar.gz | tar -xvf  
_
```
- c. Untar the product tar file, and run the installation command for the platform you are installing on.

#### RHEL

```
sudo yum install ./<file_name>.rpm
```

#### Ubuntu

```
sudo dpkg -i ./<file_name>.deb
```

- d. From the directory that contains the extracted tar file, run this script as root or with sudo privileges:

```
sudo /opt/powerai-vision/bin/load_images.sh -f ./file_name.tar
```

**Note:** The installation process can take some time to complete.

- e. After the installation is complete, you can start PowerAI Vision by running this script:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

A user named admin is created with a password of passw0rd. . This user owns all data that was created in the previous version. If other users existed in the previous version, they will no longer exist. For instructions to change the created user name or password, and to learn how to create new users, see “Managing users” on page 81.

- f. Install any available fix packs. For instructions see “Getting fixes from Fix Central” on page 102.
3. Delete any deployed models from the previous version of the product. To delete a deployed model, click **Deployed Models**. Next, select the model that you want to delete and click **Delete**. The trained model is not deleted from PowerAI Vision.
4. Redeploy trained models as necessary.
  - a. Click **Models** from the menu.
  - b. Select the model you want to deploy and click **Deploy**.
  - c. Specify a name for the model, and for models that were trained with the **Optimized for speed (tiny YOLO v2)** model, choose the accelerator to deploy to. You can choose GPU, CPU, or Xilinx FPGA - 16 bit (technology preview).
  - d. Click **Deploy**. The Deployed Models page is displayed. When the model has been deployed, the status column displays **Ready**.
  - e. Click the deployed model to get the API endpoint, to view details about the model, such as the owner and the accuracy, and to test other videos or images against the model.

### Upgrade PowerAI Vision with IBM Cloud Private

1. Download the product tar file from the IBM Passport Advantage website.
  2. To make PowerAI Vision available in IBM Cloud Private catalog, run the appropriate command:
    - IBM Cloud Private 2.1.0.3:

```
bx pr load-ppa-archive --archive file_name.tar [--clustername <cluster_CA_domain>]
```
    - IBM Cloud Private 3.1.0 or later:

```
cloudctl catalog load-archive --archive file_name.tar
```

Where `cluster_CA_domain` is the certificate authority (CA) domain. If you did not specify a CA domain, the default value is `mycluster.icp`.
  3. Navigate to your Helm Release. Click **Upgrade** and the upgrade to the new PowerAI Vision images starts.
- Note:** The upgrade process can take some time to complete.
4. As part of the upgrade process, PowerAI Vision is restarted and a user named `admin` is created with a password of `passw0rd`. This user owns all data that was created in the previous version. If other users existed in the previous version, they will no longer exist. For instructions to change the created user name or password, and to learn how to create new users, see “Managing users” on page 81.

---

## Uninstalling PowerAI Vision stand-alone

You must uninstall PowerAI Vision stand-alone on your system, before you can install IBM Cloud Private, IBM Data Science Experience Local, or other Kubernetes-based applications.

To uninstall PowerAI Vision, complete the following steps:

**Note:** If you run the following commands, all the data that you gathered is deleted. Export your data sets and models before you run the following commands.

1. Stop the current instance of PowerAI Vision by running the following script:

```
sudo /opt/powerai-vision/bin/powerai_vision_stop.sh
```
2. Remove previously installed images and data by running the following script:

```
sudo /opt/powerai-vision/bin/purge_image.sh
```
3. Remove PowerAI Vision by running the following command:

```
sudo yum remove powerai-vision
```
4. Delete the data directory by running the following command:

```
sudo rm -rf /opt/powerai-vision/
```
5. Verify that PowerAI Vision was uninstalled by running the following command:

```
rpm -q powerai-vision
```



---

## Checking the application and environment

After installation of PowerAI Vision, you can check the status of the application and environment by using commands documented in these topics. The Kubernetes commands `helm.sh` and `kubect1.sh` are installed in the `bin` directory of the product install path. (default: `/opt/powerai-vision`).

---

### Checking Kubernetes services status

The Kubernetes infrastructure is used to run the PowerAI Vision application. The `kubect1` command can be used to check the status of these underlying services, using the `--namespace kube-system` option.

- “Using `kubect1 get pods` to check kube-system”
- “Using `kubect1 describe pods` to check kube-system”

### Using `kubect1 get pods` to check kube-system

The `kubect1` command is used to show the detailed status of the Kubernetes pods deployed to run the PowerAI Vision application.

#### Example output

```
# kubect1 get pods --namespace kube-system
NAME                                READY    STATUS    RESTARTS   AGE
default-http-backend-77c86f88b4-bm9nq  1/1      Running   0           4d
kube-dns-c5b9d46b-fffn7              3/3      Running   0           4d
nginx-ingress-lb-ppc64le-djkt6        1/1      Running   0           4d
tiller-deploy-5f954f4845-9sr64        1/1      Running   0           4d
```

#### Interpreting the output

- When the Kubernetes system is running correctly, each of the pods should have:
  - In the `READY` column all pods should be counted - for example, “1/1” or “3/3”.
  - A value of “Running” in the `STATUS` column.
- A `STATUS` value other than “Running” indicates an issue with the Kubernetes infrastructure.
- A non-0, and growing, value in the `RESTARTS` column indicates an issue with that Kubernetes pod.

### Using `kubect1 describe pods` to check kube-system

The `kubect1 describe pods` command provides detailed information about each of the pods that provide Kubernetes infrastructure. If the output from a specific pod is desired, run the command `kubect1 describe pod pod_name --namespace kube-system`.

#### Example output

The output from the command is verbose, so sample output from only one pod is shown:

```
# kubectl describe pods --namespace kube-system
...
Name:          tiller-deploy-5f954f4845-9sr64
Namespace:     kube-system
Node:          127.0.0.1/127.0.0.1
Start Time:    Mon, 17 Sep 2018 12:26:33 -0500
Labels:        app=helm
               name=tiller
               pod-template-hash=1951090401
Annotations:   kubernetes.io/created-by={"kind":"SerializedReference","apiVersion":"v1","reference":{"kind":"ReplicaSet","namespace":"ku
Status:        Running
IP:            172.17.0.4
Created By:    ReplicaSet/tiller-deploy-5f954f4845
Controlled By: ReplicaSet/tiller-deploy-5f954f4845
Containers:
  tiller:
    Container ID:  docker://f049f19b4180d0406c04fa7d5ca8993ac1ef596a29d8d0096a54eb504182dd0b
    Image:         ibmcom/tiller-ppc64le:v2.6.0
    Image ID:      docker-pullable://docker.io/ibmcom/tiller-ppc64le@sha256:6dc8e12643d0c78b268f221205e00751ae20d37de31d45af2b21065652f0
    Port:          44134/TCP
    State:         Running
      Started:     Mon, 17 Sep 2018 12:26:38 -0500
    Ready:         True
    Restart Count: 0
    Liveness:      http-get http://:44135/liveness delay=1s timeout=1s period=10s #success=1 #failure=3
    Readiness:     http-get http://:44135/readiness delay=1s timeout=1s period=10s #success=1 #failure=3
    Environment:
      TILLER_NAMESPACE: kube-system
    Mounts:
      /var/run/secrets/kubernetes.io/serviceaccount from default-token-xs2b4 (ro)
Conditions:
  Type          Status
  Initialized   True
  Ready         True
  PodScheduled  True
Volumes:
  default-token-xs2b4:
    Type:          Secret (a volume populated by a Secret)
    SecretName:    default-token-xs2b4
    Optional:      false
QoS Class:       BestEffort
Node-Selectors:  beta.kubernetes.io/os=linux
Tolerations:     node.alpha.kubernetes.io/notReady:NoExecute for 300s
                  node.alpha.kubernetes.io/unreachable:NoExecute for 300s
Events:          <none>
```

## Interpreting the output

Significant fields providing status of the Kubernetes pods include:

- The **Status** field should be “Running” - any other status will indicate issues with the environment.
- In the **Conditions** section, the **Ready** field should indicate “True”. Any other value indicates that there are issues with the environment.
- If there are issues with any pods, the **Events** section of the pod should have information about issues the pod encountered.

## Checking Kubernetes node status

Use these commands to check the status of the nodes in the environment.

- “kubectl.sh get pods” on page 31
- “kubectl describe nodes command” on page 31
- “kubectl describe pods command” on page 33

## kubect1.sh get pods

The `kubect1` command is used to show the detailed status of the Kubernetes pods deployed to run the PowerAI Vision application.

### Example output

```
[root@aprilmin8 bin]# /opt/powerai-vision/bin/kubect1 get pods
NAME                                READY   STATUS    RESTARTS   AGE
powerai-vision-infer-ic-f873e5ff-821b-472a-bec0-32b1cdc4b59ns7j  1/1     Running   0           14d
powerai-vision-infer-od-fae2aa93-e394-4a04-bec7-f072b81f2678b2m  1/1     Running   0           14d
powerai-vision-keycloak-987f9698d-kdvp2                          1/1     Running   1           14d
powerai-vision-mongodb-686ffbd4b9-cq2mx                          1/1     Running   0           14d
powerai-vision-portal-768c4ffdc7-9ppnt                           1/1     Running   0           14d
powerai-vision-postgres-7c88467b4c-szw8m                         1/1     Running   0           14d
powerai-vision-taskanalyzer-788bdcc5cc-dc5cp                     1/1     Running   0           14d
powerai-vision-ui-844bb5f7f9-t55bg                               1/1     Running   0           14d
powerai-vision-video-nginx-845544c55b-w8ks6                      1/1     Running   0           14d
powerai-vision-video-portal-7499b577c7-csvs8                     1/1     Running   0           14d
powerai-vision-video-rabmq-85b486bc9d-nqdmn                      1/1     Running   0           14d
powerai-vision-video-redis-67d9766fb6-gw7hd                      1/1     Running   0           14d
powerai-vision-video-test-nginx-b5cb6b759-fkwp5                 1/1     Running   0           14d
powerai-vision-video-test-portal-756f7f9b99-mtwwj                1/1     Running   0           14d
powerai-vision-video-test-rabmq-578c775d4-chbtd                  1/1     Running   0           14d
powerai-vision-video-test-redis-779f7545b4-dffws                 1/1     Running   0           14d
```

### Interpreting the output

- When the application is running correctly, each of the pods should have:
  - A value of 1/1 in the READY column
  - A value of Running in the STATUS column
- In the above example output, pods with `infer` in the name are created when a model is deployed. These will only appear if there are models deployed in the instance of the application running on the system.
- A STATUS value other than Running indicates an issue with the pod.
- A non-0 and increasing value in the RESTARTS column indicates an issue with that pod.

If there are indications of issues with pods, see “Troubleshooting known issues - PowerAI Vision standard install” on page 91.

## kubect1 describe nodes command

The `kubect1 describe nodes` command provides status information regarding the Kubernetes environment used to run the PowerAI Vision application.

### Example output

```

# /opt/powerai-vision/bin/kubectl.sh describe nodes
Name: 127.0.0.1
Roles: <none>
Labels: beta.kubernetes.io/arch=ppc64le
        beta.kubernetes.io/os=linux
        gpu/nvidia=TeslaP100-SXM2-16GB
        kubernetes.io/hostname=127.0.0.1
Annotations: node.alpha.kubernetes.io/ttl=0
              volumes.kubernetes.io/controller-managed-attach-detach=true
Taints: <none>
CreationTimestamp: Tue, 14 Aug 2018 00:54:26 -0500
Conditions:
  Type                Status  LastHeartbeatTime             LastTransitionTime            Reason                       Message
  ----                -
  OutOfDisk            False   Wed, 12 Sep 2018 18:26:03 -0500   Tue, 14 Aug 2018 00:54:26 -0500   KubeletHasSufficientDisk    kubelet has s
  MemoryPressure       False   Wed, 12 Sep 2018 18:26:03 -0500   Tue, 14 Aug 2018 00:54:26 -0500   KubeletHasSufficientMemory  kubelet has s
  DiskPressure         False   Wed, 12 Sep 2018 18:26:03 -0500   Tue, 14 Aug 2018 00:54:26 -0500   KubeletHasNoDiskPressure    kubelet has n
  Ready               True    Wed, 12 Sep 2018 18:26:03 -0500   Tue, 14 Aug 2018 00:54:36 -0500   KubeletReady                 kubelet is po

Addresses:
  InternalIP: 127.0.0.1
  Hostname: 127.0.0.1
Capacity:
  alpha.kubernetes.io/nvidia-gpu: 4
  cpu: 128
  memory: 1067492928Ki
  pods: 110
Allocatable:
  alpha.kubernetes.io/nvidia-gpu: 4
  cpu: 128
  memory: 1067390528Ki
  pods: 110
System Info:
  Machine ID: 833e0926ee21aed71ec075d726cbcfe0
  System UUID: 101537A
  Boot ID: 6edbbcle-475d-4ac7-b8dc-18227ff6a6f4
  Kernel Version: 3.10.0-862.6.3.el7.ppc64le
  OS Image: Debian GNU/Linux 8 (jessie)
  Operating System: linux
  Architecture: ppc64le
  Container Runtime Version: docker://1.13.1
  Kubelet Version: v1.8.3+icp
  Kube-Proxy Version: v1.8.3+icp
  ExternalID: 127.0.0.1
Non-terminated Pods: (20 in total)
  Namespace      Name
  -----
  default         powerai-vision-infer-ic-2c9713b3-a357-466e-a10c-596e71e735c6wvm
  default         powerai-vision-infer-od-2b2c5f0d-f6ae-4f46-8f0b-b946c2bc18kznb6
  default         powerai-vision-keycloak-987f9698d-kdvp2
  default         powerai-vision-mongodb-686ffbd4b9-cq2mx
  default         powerai-vision-portal-768c4ffdc7-9686z
  default         powerai-vision-postgres-7c88467b4c-szw8m
  default         powerai-vision-taskanaly-788bdcc5cc-dc5cp
  default         powerai-vision-ui-844bb5f7f9-t55bg
  default         powerai-vision-video-nginx-845544c55b-w8ks6
  default         powerai-vision-video-portal-7499b577c7-csvs8
  default         powerai-vision-video-rabmq-85b486bc9d-nqdmn
  default         powerai-vision-video-redis-67d9766fb6-gw7hd
  default         powerai-vision-video-test-nginx-b5cb6b759-fkwp5
  default         powerai-vision-video-test-portal-756f7f9b99-mtwwj
  default         powerai-vision-video-test-rabmq-578c775d4-chbtd
  default         powerai-vision-video-test-redis-779f7545b4-dffws
  kube-system     default-http-backend-77c86f88b4-g8sp9
  kube-system     kube-dns-c5b9d46b-67dvg
  kube-system     nginx-ingress-lb-ppc64le-pdzs6
  kube-system     tiller-deploy-5f954f4845-kkv8z

CPU Requests  CPU Limits  Memory Requests
-----
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)
0 (0%)        0 (0%)      0 (0%)

Allocated resources:
(Total limits may be over 100 percent, i.e., overcommitted.)
CPU Requests  CPU Limits  Memory Requests  Memory Limits  NvidiaGPU Limits
-----
0 (0%)        0 (0%)      0 (0%)          0 (0%)         2 (50%)
Events: <none>
#

```

## Interpreting the output

- Most of the information is informational regarding the system resources (CPUs, GPUs, memory) and version information (OS, Docker, Kubernetes).
- The **Conditions** section can indicate whether there are system resource issues that will affect the running of the application. For example, if any of the **OutOfDisk**, **MemoryPressure**, or **DiskPressure** conditions are **True**, there are insufficient system resources to run PowerAI Vision. For example, the following **Conditions** section shows a system that does not have sufficient disk space available, indicated by **DiskPressure** status of **True**:

Conditions:					
Type	Status	LastHeartbeatTime	LastTransitionTime	Reason	Message
----	-----	-----	-----	-----	-----
OutOfDisk	False	[...]	[...]	KubeletHasSufficientDisk	kubelet has su
MemoryPressure	False	[...]	[...]	KubeletHasSufficientMemory	kubelet has su
DiskPressure	<b>True</b>	[...]	[...]	KubeletHasDiskPressure	kubelet has di
Ready	True	[...]	[...]	KubeletReady	kubelet is pos

- The **Events** section will also have messages that can indicate if there are issues with the environment. For example, the following events indicate issues with disk space that have led to Kubernetes attempting to reclaim resources (“eviction”) which can affect the availability of Kubernetes applications:

Events:				
Type	Reason	Age	From	Message
----	-----	----	----	-----
Normal	NodeHasDiskPressure	5m	kubelet, 127.0.0.1	Node 127.0.0.1 status is now: NodeHasDiskPressure
Warning	<b>EvictionThresholdMet</b>	3s (x23 over 5m)	kubelet, 127.0.0.1	Attempting to reclaim nodefs

## kubect! describe pods command

The `kubect! .sh describe pods` command provides detailed information about each of the pods used by the PowerAI Vision application. If the output from a specific pod is desired, the command `kubect! .sh describe pod podname`. To determine the values for *podname* look at the output from `kubect! .sh get pods`.

### Example output

The output from the command is verbose, so sample output from only one pod is shown:

```

# kubectl describe pods
...
Name:          powerai-vision-ui-844bb5f7f9-khx1m
Namespace:     default
Node:          127.0.0.1/127.0.0.1
Start Time:    Mon, 17 Sep 2018 12:27:18 -0500
Labels:        app=powerai-vision
               chart=ibm-powerai-vision-prod-1.1.0
               component=powerai-vision-ui
               heritage=Tiller
               pod-template-hash=4006619395
               release=vision
Annotations:   checksum/config=b131ad79a4838feecc85cde3b422bc82ef3214a437462e02b775df6d3582a4f6
               kubernetes.io/created-by={"kind":"SerializedReference","apiVersion":"v1","reference":{"kind":"ReplicaSet","namespace":"de
               productID=5737-H10
               productName=IBM PowerAI Vision
               productVersion=1.1.1.0
Status:        Running
IP:            172.17.0.14
Created By:    ReplicaSet/powerai-vision-ui-844bb5f7f9
Controlled By: ReplicaSet/powerai-vision-ui-844bb5f7f9
Containers:
  powerai-vision-ui:
    Container ID:  docker://16867f9922458d4a517018f52edc6c319e1e9d6408cc333cf242be618e179425
    Image:         powerai-vision-ui:1.1.1.0
    Image ID:      docker://sha256:c327077b7d605518f6b4651141ab864459212d34022fabad844073e5e66e9b9c
    Port:          80/TCP
    State:         Running
      Started:     Mon, 17 Sep 2018 12:27:33 -0500
    Ready:         True
    Restart Count: 0
    Liveness:      http-get http://:http/powerai-vision/index.html delay=240s timeout=5s period=10s #success=1 #failure=3
    Readiness:     http-get http://:http/powerai-vision/index.html delay=5s timeout=1s period=10s #success=1 #failure=3
    Environment:
      CONTEXT_ROOT:      <set to the key 'CONTEXT_ROOT' of config map 'powerai-vision-config'>      Optional: false
      DLAAS_API_SERVER:  <set to the key 'DLAAS_API_SERVER' of config map 'powerai-vision-config'>      Optional: false
      SERVER_HOST_VIDEO_TEST: <set to the key 'SERVER_HOST_VIDEO_TEST' of config map 'powerai-vision-config'>      Optional: false
      SERVICE_PORT_VIDEO_TEST: <set to the key 'SERVICE_PORT_VIDEO_TEST' of config map 'powerai-vision-config'>      Optional: false
      WEBROOT_VIDEO_TEST:  <set to the key 'WEBROOT_VIDEO_TEST' of config map 'powerai-vision-config'>      Optional: false
    Mounts:
      /opt/powerai-vision/data from data-mount (rw)
      /var/run/secrets/kubernetes.io/serviceaccount from default-token-hr9zz (ro)
Conditions:
  Type             Status
  Initialized       True
  Ready            True
  PodScheduled     True
Volumes:
  data-mount:
    Type:          PersistentVolumeClaim (a reference to a PersistentVolumeClaim in the same namespace)
    ClaimName:     powerai-vision-data-pvc
    ReadOnly:      false
  default-token-hr9zz:
    Type:          Secret (a volume populated by a Secret)
    SecretName:    default-token-hr9zz
    Optional:      false
QoS Class:        BestEffort
Node-Selectors:   beta.kubernetes.io/arch=ppc64le
Tolerations:      node.alpha.kubernetes.io/notReady:NoExecute for 300s
                  node.alpha.kubernetes.io/unreachable:NoExecute for 300s
Events:           <none>
...

```

## Interpreting the output

Significant fields providing status of the application pods include:

- Information about the product name and version are given in **productName** and **productVersion**.
- The **Status** field should be Running. Any other status indicates problems with the application pod.

- If there are issues with a pod, the **Events** section of the pod should have information about problems encountered.

## Checking Kubernetes storage status

The PowerAI Vision application requires disk storage for activities including data set storage. The disk space requirements are described using Kubernetes Persistent Volume configuration. The `kubectl` command can be used to examine the *pv* (PersistentVolume) and *pvc* (PersistentVolumeClaims) resources.

**Note:** The storage requirements described in the **PersistentVolume** and **PersistentVolumeClaims** are not enforced in the standalone deployment. Therefore, the requested space might not be available in the underlying storage of the system. See “Disk space requirements” on page 14 for information about product storage requirements.

- “Using `kubectl get pv` and `pvc` commands”
- “Using the `kubectl describe pv` command”
- “Using the `kubectl describe pvc` command” on page 36

## Using `kubectl get pv` and `pvc` commands

The `kubectl get pv` and `kubectl get pvc` commands can be used to see what PersistentVolume and PersistentVolumeClaim have been defined for the application.

### Example output

```
# /opt/powerai-vision/bin/kubectl.sh get pv
NAME                CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM                                STORAGECLASS  REASON
powerai-vision-data  40Gi      RWX           Retain          Bound   default/powerai-vision-data-pvc
# /opt/powerai-vision/bin/kubectl.sh get pvc
NAME                STATUS    VOLUME          CAPACITY  ACCESS MODES  STORAGECLASS  AGE
powerai-vision-data-pvc  Bound    powerai-vision-data  40Gi      RWX           default/powerai-vision-data-pvc  48d
```

### Interpreting the output

The above output shows information about the Persistent Volume and Persistent Volume Claim for PowerAI Vision. The application currently has a capacity claim of 40G and it is successfully “Bound”. If the **STATUS** is not “Bound”, the application does not have access to the necessary storage.

## Using the `kubectl describe pv` command

The `kubectl describe pv` command is used to see detailed information about the Persistent Volume used by the application.

### Example output

```

# /opt/powerai-vision/bin/kubectl.sh describe pv
NAME                CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM                STORAGECLASS  REASON  AGE
powerai-vision-data  40Gi      RWX           Retain          Bound   default/powerai-vision-data-pvc
[root@dlf01 ~]# /opt/powerai-vision/bin/kubectl.sh describe pv
Name:                powerai-vision-data
Labels:              assign-to=powerai-vision-data
                    type=local
Annotations:         pv.kubernetes.io/bound-by-controller=yes
StorageClass:
Status:              Bound
Claim:               default/powerai-vision-data-pvc
Reclaim Policy:      Retain
Access Modes:        RWX
Capacity:            40Gi
Message:
Source:
  Type:               HostPath (bare host directory volume)
  Path:               /opt/powerai-vision/volume/
Events:               <none>

```

## Interpreting the output

The above output shows more details about the Persistent Volume used by the application. The **Source** section has the critical configuration values for **Type** and **Path**. The **Events** section will have information about Error events if there were issues with the Persistent Volume.

## Using the kubectl describe pvc command

The `kubectl describe pvc` command is used to see detailed information about the Persistent Volume Claim for the application.

### Example output

```

[root@dlf01 ~]# /opt/powerai-vision/bin/kubectl.sh describe pvc
Name:                powerai-vision-data-pvc
Namespace:           default
StorageClass:
Status:              Bound
Volume:              powerai-vision-data
Labels:              app=powerai-vision
                    chart=ibm-powerai-vision-prod-1.1.0
                    heritage=Tiller
                    release=vision
Annotations:         pv.kubernetes.io/bind-completed=yes
                    pv.kubernetes.io/bound-by-controller=yes
Capacity:            40Gi
Access Modes:        RWX
Events:              <none>

```

## Interpreting the output

The above output shows more details about the Persistent Volume Claim used by the application. The **Volume** section references the underlying Persistent Volume, and the **Status** should be “Bound” if it has been successfully allocated to the application. The **Events** section will show if there were issues with the Persistent Volume Claim.

---

## Checking application deployment

PowerAI Vision processes require a Kubernetes environment. Use these commands to verify that the Kubernetes environment was deployed correctly and that all nodes are configured appropriately.

- “helm.sh”
- “kubectl describe deployment” on page 39

### helm.sh

The `helm.sh` command shows the status of the full Kubernetes environment of the PowerAI Vision application.

### Example output

```
# helm status vision
LAST DEPLOYED: Tue Aug 14 02:26:56 2018
NAMESPACE: default
STATUS: DEPLOYED

RESOURCES:
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
powerai-vision-keycloak            1         1         1             1         14d
powerai-vision-mongodb              1         1         1             1         14d
powerai-vision-portal               1         1         1             1         14d
powerai-vision-postgres             1         1         1             1         14d
powerai-vision-taskanalyst          1         1         1             1         14d
powerai-vision-ui                   1         1         1             1         14d
powerai-vision-video-nginx          1         1         1             1         14d
powerai-vision-video-portal         1         1         1             1         14d
powerai-vision-video-rabmq          1         1         1             1         14d
powerai-vision-video-redis          1         1         1             1         14d
powerai-vision-video-test-nginx     1         1         1             1         14d
powerai-vision-video-test-portal    1         1         1             1         14d
powerai-vision-video-test-rabmq     1         1         1             1         14d
powerai-vision-video-test-redis     1         1         1             1         14d

==> v1beta1/Ingress
NAME                                HOSTS  ADDRESS  PORTS  AGE
powerai-vision-ing *                127.0.0.1  80      14d

==> v1/Secret
NAME                                TYPE    DATA  AGE
powerai-vision-secrets Opaque  6      14d

==> v1/ConfigMap
NAME                                DATA  AGE
powerai-vision-config 52     14d

==> v1/PersistentVolumeClaim
NAME                                STATUS  VOLUME          CAPACITY  ACCESSMODES  STORAGECLASS  AGE
powerai-vision-data-pvc Bound   powerai-vision-data  40Gi      RWX          14d

==> v1/Service
NAME                                CLUSTER-IP  EXTERNAL-IP  PORT(S)          AGE
powerai-vision-keycloak            10.10.0.221 <none>       8080/TCP,8443/TCP 14d
powerai-vision-mongodb              10.10.0.120 <none>       27017/TCP         14d
powerai-vision-portal               10.10.0.57  <none>       9080/TCP          14d
powerai-vision-postgres             10.10.0.123 <none>       5432/TCP          14d
powerai-vision-taskanalyst          10.10.0.157 <none>       5000/TCP          14d
powerai-vision-ui                   10.10.0.72  <none>       80/TCP            14d
powerai-vision-video-nginx          10.10.0.208 <none>       8081/TCP          14d
powerai-vision-video-portal         10.10.0.226 <none>       8080/TCP,8081/TCP 14d
powerai-vision-video-rabmq          10.10.0.117 <none>       5672/TCP          14d
powerai-vision-video-redis          10.10.0.189 <none>       6379/TCP          14d
powerai-vision-video-test-nginx     10.10.0.148 <none>       8083/TCP          14d
powerai-vision-video-test-portal    10.10.0.12  <none>       8080/TCP,8081/TCP 14d
powerai-vision-video-test-rabmq     10.10.0.26  <none>       5672/TCP          14d
powerai-vision-video-test-redis     10.10.0.27  <none>       6379/TCP          14d

NOTES:

Find the PowerAI Vision UI URL by running the following commands:

export NODE_IP=$(kubectl get ing powerai-vision-ing --namespace default -o jsonpath="{.status.loadBalancer.ingress[0].ip}")
echo https://${NODE_IP}/powerai-vision/
```

## Important fields in the output

### STATUS

The value for STATUS should be DEPLOYED after a successful installation.

## RESOURCES

The status of individual Kubernetes pods is displayed in this section. The CURRENT and AVAILABLE values for each pod should be equal to or greater than the DESIRED value.

```
RESOURCES:
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
...
powerai-vision-portal              1        1        1           1          14d
...
```

## kubectl describe deployment

The `kubectl describe deployment` command provides verbose status information about each of the deployed nodes in the Kubernetes environment that is being used to run PowerAI Vision.

### Example output

The following shows the output from one of the nodes. The full output for all nodes is much longer and has similar entries for each node.

```
# /opt/powerai-vision/bin/kubectl.sh describe deployment
...
Name:                powerai-vision-ui
Namespace:           default
CreationTimestamp:    Mon, 17 Sep 2018 12:27:00 -0500
Labels:              app=powerai-vision
                    chart=ibm-powerai-vision-prod-1.1.0
                    heritage=Tiller
                    release=vision
                    run=powerai-vision-ui-deployment
Annotations:         deployment.kubernetes.io/revision=1
Selector:            run=powerai-vision-ui-deployment-pod
Replicas:            1 desired | 1 updated | 1 total | 1 available | 0 unavailable
StrategyType:        RollingUpdate
MinReadySeconds:     0
RollingUpdateStrategy: 1 max unavailable, 1 max surge
Pod Template:
  Labels:            app=powerai-vision
                    chart=ibm-powerai-vision-prod-1.1.0
                    component=powerai-vision-ui
                    heritage=Tiller
                    release=vision
                    run=powerai-vision-ui-deployment-pod
  Annotations:       checksum/config=b131ad79a4838feecc85cde3b422bc82ef3214a437462e02b775df6d3582a4f6
                    productID=5737-H10
                    productName=IBM PowerAI Vision
                    productVersion=1.1.1.0
Containers:
  powerai-vision-ui:
    Image:            powerai-vision-ui:1.1.1.0
    Port:             80/TCP
    Liveness:         http-get http://:http/powerai-vision/index.html delay=240s timeout=5s period=10s #success=1 #failure=3
    Readiness:        http-get http://:http/powerai-vision/index.html delay=5s timeout=1s period=10s #success=1 #failure=3
    Environment:
      CONTEXT_ROOT:    <set to the key 'CONTEXT_ROOT' of config map 'powerai-vision-config'>      Optional: false
      DLAAS_API_SERVER: <set to the key 'DLAAS_API_SERVER' of config map 'powerai-vision-config'>      Optional: false
      SERVER_HOST_VIDEO_TEST: <set to the key 'SERVER_HOST_VIDEO_TEST' of config map 'powerai-vision-config'> Optional: false
      SERVICE_PORT_VIDEO_TEST: <set to the key 'SERVICE_PORT_VIDEO_TEST' of config map 'powerai-vision-config'> Optional: false
      WEBROOT_VIDEO_TEST: <set to the key 'WEBROOT_VIDEO_TEST' of config map 'powerai-vision-config'> Optional: false
    Mounts:
      /opt/powerai-vision/data from data-mount (rw)
  Volumes:
    data-mount:
      Type:            PersistentVolumeClaim (a reference to a PersistentVolumeClaim in the same namespace)
      ClaimName:       powerai-vision-data-pvc
      ReadOnly:        false
Conditions:
  Type          Status  Reason
  ----          -
  Available      True   MinimumReplicasAvailable
OldReplicaSets: <none>
NewReplicaSet:  powerai-vision-ui-844bb5f7f9 (1/1 replicas created)
Events:         <none>
...
```

## Interpreting the output

- The **Replicas** line shows information regarding how many images are desired and available (similar to the output from `kubectl get pods`):  
Replicas: 1 desired | 1 updated | 1 total | 1 available | 0 unavailable  
The “available” value should be equal to the “desired” value.
- The **productVersion** value indicates the level of PowerAI Vision installed:  
productVersion=1.1.1.0
- The **Image** value provides information about the Docker container:  
Image: powerai-vision-ui:1.1.1.0
- The **Conditions** section has important information about the current status of the image, and any reasons if the status is “failure”.

## Checking system GPU status

In PowerAI Vision, GPUs are used to train and deploy models. Use these commands to verify that GPUs are set up and available.

`nvidia-smi`

The `nvidia-smi` command is a NVIDIA utility, installed with the CUDA toolkit. For details, see “Prerequisites for installing PowerAI Vision” on page 17. With `nvidia-smi`, you can view the status of the GPUs on the system.

### Example output

```
# nvidia-smi
Thu Sep  6 17:57:38 2018

+-----+
| NVIDIA-SMI 396.15                  Driver Version: 396.15          |
+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|  0   Tesla P100-SXM2...    Off   | 00000002:01:00:0 Off |             0        |
| N/A   30C    P0     38W / 300W | 747MiB / 16280MiB |      0%    Default   |
+-----+-----+
|  1   Tesla P100-SXM2...    Off   | 00000003:01:00:0 Off |             0        |
| N/A   31C    P0     31W / 300W | 10MiB / 16280MiB |      0%    Default   |
+-----+-----+
|  2   Tesla P100-SXM2...    Off   | 0000000A:01:00:0 Off |             0        |
| N/A   31C    P0     32W / 300W | 10MiB / 16280MiB |      0%    Default   |
+-----+-----+
|  3   Tesla P100-SXM2...    Off   | 0000000B:01:00:0 Off |             0        |
| N/A   30C    P0     30W / 300W | 10MiB / 16280MiB |      0%    Default   |
+-----+-----+

+-----+
| Processes:                         GPU Memory |
|   GPU       PID    Type    Process name      Usage    |
+-----+-----+
|      0      124097    C   /usr/bin/python      729MiB   |
+-----+
```

### Interpreting the output

The above output shows the following:

- The system has 4 (0-3) **Tesla P100** GPUs.
- In the last portion of the output, it shows that GPU 0 has a process deployed and running. This can indicate a PowerAI Vision training task or a deployed model. Any GPUs with running jobs are not available for training jobs or deployment of trained models from the user interface.
- The output should correctly display the memory configuration of the GPUs. For example, “Unknown error” indicates an issue with the driver setup or configuration. See “GPUs are not available for training or inference” on page 93 for more information.



---

## Logging in to PowerAI Vision

Follow these steps to log in to PowerAI Vision.

**Note:** PowerAI Vision is supported on these browsers:

- Google Chrome Version 60, or later
- Firefox Quantum 59.0, or later

1. Enter the appropriate PowerAI Vision URL in a supported browser:

**PowerAI Vision stand-alone URL**

`https://hostname/powerai-vision/`, where *hostname* is the system on which you installed PowerAI Vision.

**PowerAI Vision with IBM Cloud Private URL**

`https://proxyhost/powerai-vision-RELEASE/`, where *proxyhost* is the host name of your IBM Cloud Private proxy server, and *RELEASE* is the name you specified in the Release name field when you deployed the Helm chart.

2. Enter your user name and password. A default user name (admin) and password (passw0rd) was created at install time. For instructions to change these values, see “Managing users” on page 81.

**Related concepts:**

“Managing users” on page 81

There are two kinds of users in PowerAI Vision: administrators, and everyone else. The way you work with users and passwords differs, depending on how PowerAI Vision is installed.



---

## Working with the user interface

The PowerAI Vision user interface is made up of these basic parts: the title bar, the side bar, the action bar, the data area, and the notification center.

### Interface areas

The user interface is made up of several different areas:

#### The title bar

The title bar lets you choose a work space, access the notification area, or work with your profile.

#### The action bar

This is where you find the actions that you can take on images, videos, data sets, and models in the current data area. The available actions differ depending on what type of object you are working with.

#### The side bar

Data sets and models have a side bar with filtering options. Filtering helps you specify which objects to include in the data area.

**Navigating:** If the side bar is long, for example, if you have a data set with a lot of different types of objects, you can scroll through the side bar content. To scroll, hover over the appropriate content and use your mouse roller or keyboard arrow keys. If the mouse pointer is right over the categories, for example, scrolling moves you through that list. If the mouse pointer is further to the right, on the edge of the side bar, scrolling moves you through all of the content on the side bar.

#### The data area

This is where you find the objects that you can act on. It lists the objects of the selected type, or displays the data included in the data set.

#### Filtering:

With large data sets, you might need to filter the files that are shown in the data area. By default, your whole data set is shown.

#### Filter by Images / Videos

When you deselect a file type, those files are no longer shown in the data area. Therefore, if you only have Images selected, only images are shown in the data area.

#### Categories / objects

When you select categories and / or objects, *all* files of the specified type that belong to any of the selected categories, or contain the selected objects, are shown.

For example, assume you have a data set with two categories: Cats and Dogs. Also assume that you tagged these types of objects: Face, Collar, and Tail. Then if you select Images, the category Dogs, and the object Collar, you will see all images that are dogs *or* contain a collar. This will include images of cats if they have a collar as well as images of dogs with no collar.

#### Using filtering and “Select all” with video data

When you capture frames from a video, these frames always maintain a child / parent relationship with the original video. That has some selection and filtering implications.

- When using the filter on the side bar, if *any* video frame matches the filter criteria, both the frame and its parent video are selected and are shown in the data area.

- If you click the “Select” box in the action bar, everything in the data area is selected. Therefore, if there is a video shown in the data area, it, and all of its child frames, are selected. Any action performed in this situation applies to all selected images, the video, and all of its child frames.

### Example

A user has captured 50 frames from a video file Cars Video. Fourteen frames of the 50 have no labels.

1. The user selects **Unlabeled** in the Objects filter in the sidebar. The 14 frames with no labels and their parent video, Cars Video, are shown in the data area.
2. The user clicks **Select** in the action bar. The frames and the video are all selected.
3. The user clicks **Delete**, intending to delete the unlabeled frames. However, because the video was selected, it, and the 36 labeled frames, are also deleted.

To delete only the unlabeled frames, the user should click **Select** in the action bar to quickly select all 14 frames, then deselect the video file before clicking **Delete**.

### The notification area

Click the bell icon in the title bar to access the notification area. This allows you to view and work with messages. Click the arrow to return to your previous view.

### Related concepts:

“Training and working with models” on page 47

Use these processes to create, deploy, and refine models.

“Scenario: Detecting objects in a video” on page 69

In this fictional scenario, you want to create a deep learning model to monitor traffic on a busy road. You have a video that displays the traffic during the day. From this video, you want to know how many cars are on the busy road every day, and what are the peak times that have the most cars on the road.

### Related tasks:

“Training a model” on page 50

After the data set has all of the object labels added, you can train your deep learning model. Trained models can then be deployed for use.

“Creating and working with data sets” on page 47

Before you can work with videos or images, you need to create a data set. A data set is a group of images, videos, or both that you will use to train a deployable model.

“Deploying a trained model” on page 60

Deploy a trained model to get it ready to use within PowerAI Vision or a different program, such as IBM PowerAI. Deploying a model creates a unique API endpoint based on that model for inference operations.

“Scenario: Classifying images” on page 75

The goal of this example is to train a model to classify images of birds into groups based on their physiological similarities. Once the model is trained with a known dataset, users can upload new data sets to auto classify the birds into their respective categories. We will prepare the data, create a data set, train the model, and test the model.

---

## Training and working with models

Use these processes to create, deploy, and refine models.

You can only see and work with objects (data sets, files, trained models, and deployed models) that you own. An object is owned by the user who created it.

---

### Creating and working with data sets

Before you can work with videos or images, you need to create a data set. A data set is a group of images, videos, or both that you will use to train a deployable model.

To create a data set and add content to it, follow these steps:

1. Log in to PowerAI Vision.
2. Click **Data Sets** in the title bar to open the Data Sets page. There are several ways to create your new data set:
  - To create an empty data set, click **Create new data set**.
  - If you have a previously exported data set, click **Import .zip file**.
  - If you want to copy an existing data set, select the data set and click **Copy**.

**Note:** PowerAI Vision limited support for Pascal VOC annotations. Annotations for multiple files residing in a common XML file are not supported. In other words, each annotation XML file can only contain annotations for a single image, identified by the filename attribute.

If you have a single XML annotation file containing annotations for multiple images in the data set to be imported, the annotations need to be split out into separate XML files before PowerAI Vision can import the annotations successfully.

3. Click the data set you just created to open it. Add images and videos by using **Import file** or by dragging them to the + area. If you do not follow these considerations, your upload will fail and a message will be shown on the screen. For details about why the upload failed, click the bell icon at the top of the page to open the Notifications center.

#### Upload considerations:

- You can select multiple image or video files, or a single .zip file that contains images and videos, but you cannot upload a folder that contains images or videos.
- You can play only the following video types in PowerAI Vision:
  - Ogg Vorbis (.ogg)
  - VP8 or VP9 (.webm)
  - H.264 encoded videos with MP4 format (.mp4)
- You cannot navigate away from the PowerAI Vision page or refresh until the upload completes. You can navigate to different pages within PowerAI Vision during the upload.
- There is a 2GB size limit per upload session. This limit applies to a single .zip file or a set of files. You can, however upload 2 GB of files, then upload more after the original upload completes.
- The models used by PowerAI Vision have limitations on the size and resolution of images that can be handled. If the original data is high resolution, then the user must consider:
  - If the images do not need fine detail for classification or object detection, they should be downsampled to 1-2 megapixels.
  - If the images do require fine detail, they should to be divided into smaller images of 1-2 megapixels each.

- High resolution images will be scaled to a maximum of 1000 x 600 pixels.

## Working with data sets

After your data set has been created, select it in the Data Sets page to duplicate, rename, delete it, and so on. To work with the images and videos contained in the data set, click the name of the data set to open it.

---

## Labeling objects

One of the most important steps is to ensure that you properly label objects by adding tags to your data.

You must label at least five separate instances of each object for training. For example, if you want to train the data set to recognize cars and you have three images and one video, you must add the “car” tag to each image and at least two frames of the video. Tagging five cars in one image would not be adequate. If this requirement is not met and you train the model, it will not be trained to recognize that type of object.

**Note:** A data set with a variety of representative objects tagged will train a more accurate model. The exact number of images and objects cannot be specified, but some guidelines recommend as many as 1,000 representative images for each class. However, you might not need a data this large to train a model with satisfactory accuracy.

If your data set does not have many images or sufficient variety for training, consider using the Augmentation feature to increase the data set.

- “Labeling videos”
- “Labeling images” on page 49

## Labeling videos

1. Select the video from your data set and select **Label Objects**.
2. Capture frames by using one of these options:
  - **Auto capture frames** - PowerAI Vision captures a video frame every  $n$  seconds, where  $n$  is specified in the **Capture Interval (seconds)** field.

**Note:**

- Depending on the length and size of the video and the interval you specified to capture frames, the process to capture frames can take several minutes.
- When performing multiple auto label operations on the same video, it is possible to get multiple frames with the same time offset. This situation can occur when the intervals overlap and labels have been edited on the frames at the overlap points.

For example, labeling at a 10 second interval, editing some of the labels on those frames, and then labeling again at 5 a second interval has an overlap every 10 seconds. There might be duplicate images at each of the 10 second intervals with edited labels.

- **Manually capture frames** - use **Capture frame** to capture relevant frames.
3. If required, manually add new frames to an existing data set. This might happen if **Auto capture frames** does not produce enough frames with a specific object type. To manually add new frames, follow these steps:
    - a. Play the video and when the frame you want is displayed, click the pause icon.

**Tip:** You can use the video player's status bar to find a frame you want.

- b. Click **Capture Frame**.

4. Create new object labels for the data set by clicking **Add object** in the left pane. To add multiple object labels, enter one label, click **+**, then enter the next until you are done. Label names cannot contain any special characters other than the underscore ( `_` ). For example, characters such as these are not allowed: `"/ \ | { } ( ) ; ,`
5. Tag the objects in the images by following these steps.
  - a. Select the first frame in the carousel.
  - b. Select the correct object label.
  - c. Click and hold down the left mouse button to draw a box around the object.

Follow these guidelines when identifying and drawing objects in video frames:

  - Do not label part of an object. For example, do not label a car that is only partially in the frame.
  - If a frame has more than one object, you must label all objects. For example, if you have cars and motorcycles in a video frame, you must label the cars and the motorcycles. Label objects with a consistent approach.
  - Draw a box around each individual object. Do not draw a box around groups of objects. For example, if two cars are right next to each other, you must draw a box around each car.
  - Draw the box as close to the objects as possible. Do not leave blank space around the objects.
  - You can draw boxes around objects that overlap. For example, if one object is behind another object, you can label them both. However, it is recommended that you only label objects if the majority of the object is visible. You can also draw boxes that touch each other.
  - You can zoom in on the frame to make it easier to draw the box around objects.
  - Boxes cannot extend off the edge of the frame.
6. After all objects are labeled in all of the images in the carousel, click **Done editing**.

## Labeling images

Follow these steps to label images in your data set:

1. Create new object labels for the data set by clicking **Add object** in the left pane. To add multiple object labels, enter one label, click **+**, then enter the next until you are done. Label names cannot contain any special characters other than the underscore ( `_` ). For example, characters such as these are not allowed: `"/ \ | { } ( ) ; ,`
2. Tag the objects in the images, following these guidelines:
  - Do not label part of an object. For example, do not label a car that is only partially in the image.
  - If an image has more than one object, you must label all objects. For example, if you have cars and motorcycles in an image, you must label the cars and the motorcycles. Label objects consistently.
  - Draw a box around each individual object. Do not draw a box around groups of objects. For example, if two cars are right next to each other, you must draw a box around each car.
  - Draw the box as close to the objects as possible. Do not leave blank space around the objects.
  - You can draw boxes around objects that overlap. For example, if one object is behind another object, you can label them both. However, it is recommended that you only label objects if the majority of the object is visible. You can also draw boxes that touch each other.
  - You can zoom in on the image to make it easier to draw the box around objects.
  - Boxes cannot extend off the edge of the image.

### Related tasks:

“Automatically labeling objects” on page 61

After deploying a model for object detection, you can improve its accuracy by using the Auto label function. This function improves the model's accuracy by quickly adding more data to the data set.

---

## Training a model

After the data set has all of the object labels added, you can train your deep learning model. Trained models can then be deployed for use.

1. From the Data set page, click **Train**.
2. In the Train data set window, fill out the values as appropriate, then click **Train**:

### Type of training

#### Image classification

Choose this if you want to use the model to categorize images as belonging to one of the types that you defined in the data set.

**Note:** To train a model for classification, the data set must meet these requirements:

- There must be at least two categories.
- Each category must have at least five images.

#### Object detection

Choose this if you want to use the model to label objects within images.

**Note:** To train a model for object detection, the data set must contain at least five images with an object tagged for each defined label. For example, if the data set has "Car" and "Motorcycle" defined as objects, at least five images need to have "Car" tagged, and at least five images need to have "Motorcycle" tagged.

### Model selection

#### System default - Image classification

Models trained with the system default model can only be run on a GPU

#### Optimized for accuracy - Object detection

Models optimized for accuracy can only be run on a GPU.

#### Optimized for speed - Object detection

Models optimized for speed can be run anywhere, but might not be as accurate as those optimized for accuracy. These models use "you only look once" (YOLO) V2 and will take several hours to train.

You will choose the accelerator to deploy to when deploying the model. You can choose GPU, CPU, or Xilinx FPGA - 16 bit (technology preview).

#### Custom model

Select an imported model to use for training.

### Advanced options

#### Base model

PowerAI Vision comes with several common models such as flowers, faces, and so on, that you can use to help classify your data. This option is available only when training for image classification.

#### Model hyperparameters

For advanced users, these settings are available to help fine-tune the training.

##### Max iteration

The maximum number of times the data is passed through the training algorithm.

##### Momentum (Object detection only)

This value increases the step size used when trying to find the minimum value of the error curve. A larger step size can keep the algorithm from stopping at a local minimum instead of finding the global minimum.

**Ratio (Object detection only)**

PowerAI Vision automatically “splits” the data set for internal validation of the model’s performance during training. The default value of 80/20 will result in 80% of the test data (at random) being used for training, and 20% being used for measurement / validation.

**Test iteration (Image classification only)**

The number of times data is passed through the training algorithm before possible completion. For example, if this value is 100, and Test interval is 50, the model is run through the algorithm at least 100 times; being tested ever 50 times.

**Test interval (Image classification only)**

The number of times the model is passed through the algorithm before testing. For example, if this value is 50, the model is tested every 50 iterations. Each of these tests becomes a data point on the metrics graphs.

**Learning rate**

This option determines how much the weights in the network are adjusted with respect to the loss gradient. A correctly tuned value can result in a shorter training time. However, it is recommended that only advanced users change this value.

**Weight decay**

This value specifies regularization in the network. It protects against over-fitting and is used to multiply the weights when training.

**Note:** PowerAI Vision assigns one GPU to each active training job or deployed deep learning API. For example, if a system has four GPUs and you have two deployed web APIs, there are two GPUs available for active training jobs. If a training job appears to be hanging, it might be waiting for another training job to complete, or there might not be a GPU available to run it. For information to fix this, see “PowerAI Vision cannot train a model” on page 94.

3. (Optional) Stop the training process by clicking **Stop training** > **Keep Model** > **Continue**.

You can wait for the entire training model process complete, or you can stop the training process when the lines in the training graph start to flatten out. This is because improvements in quality of training might plateau over time. Therefore, the fastest way to deploy a model and refine the data set is to stop the process before quality stops improving.

**Understanding the model training graph**

As PowerAI Vision trains the model, the graph shows the relative performance of the model over time. The model should converge at the end of the training with low error and high accuracy.

In the figure, you can see the Loss CLS line and the Loss Bbox lines start to plateau. In the training graph, the lower the loss value, the better. Therefore, you can stop the training process when the loss value stops decreasing. The training model has completed enough iterations and you can continue to the next step.

## Loss VS Iteration

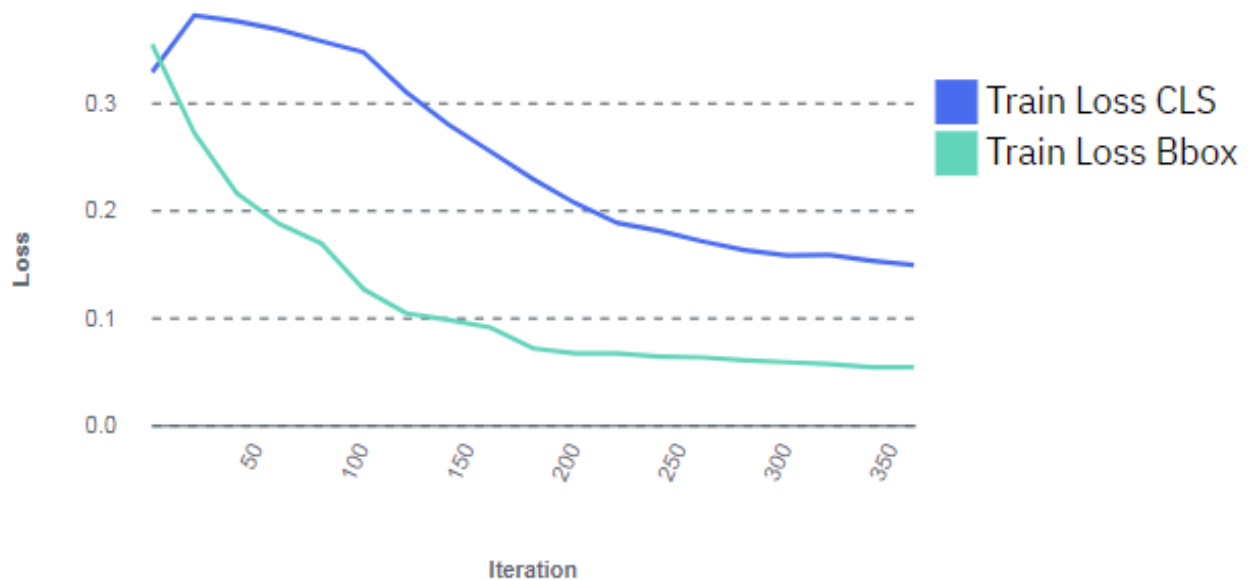


Figure 6. Model training graph

**Important:** If the training graph converges quickly and has 100% accuracy, the data set does not have enough information. The same is true if the accuracy of the training graph fails to rise or the errors in the graph do not decrease at the end of the training process. For example, a model with high accuracy might be able to discover all instances of different race cars, but might have trouble differentiating between specific race cars or those that have different colors. In this situation, add more images or video frames to the data set, tag them, then try the training again.

### Related concepts:

“Understanding metrics” on page 66

PowerAI Vision provides several metrics to help you measure how effectively your model has been trained.

## Working with custom models

You can save time and resources by using your own TensorFlow based custom models (also referred to as *custom networks*) with PowerAI Vision. In general, custom models work the same as any other model in PowerAI Vision. However, there are some differences you should understand.

When you upload a custom model to the Custom Models page, you can use the model to train a data set in PowerAI Vision and generate a PowerAI Vision trained model. Use the information in this topic to prepare a PowerAI Vision trained model by using a custom TensorFlow model: “Preparing a model that will be used to train data sets in PowerAI Vision.”

### Preparing a model that will be used to train data sets in PowerAI Vision

If your custom model will be used to train data sets in the PowerAI Vision framework, your custom model must meet the following requirements.

After the model is properly prepared, upload it to PowerAI Vision by opening the Custom Models page and clicking **Browse files**. You can then use it to train a data set. Follow these instructions to train a data set; selecting **Custom model**: “Training a model” on page 50.

### Custom model requirements:

- It must be TensorFlow based.
- It must implement the MyTrain Python class.
  - The MyTrain implementation must reside in a file named `train.py` in the top level directory of the zip file contents.
  - The following import must be added to the `train.py` file in order to define the training callbacks:  
`from train_interface import TrainCallback`
  - The class name must be `MyTrain`.

### MyTrain Template:

```
class MyTrain(TrainCallback):
    def __init__():
        pass
    def onPreprocessing(self, labels, images, workspace_path, params):
        pass
    def onTraining(self, monitor_handler):
        pass
    def onCompleted(self, model_path):
        pass
    def onFailed(self, train_status, e, tb_message):
        pass
```

### class MyTrain(TrainCallback):

Use the MyTrain API to prepare a TensorFlow model that will be used to train data sets with PowerAI Vision.

- “Template”
- “def onPreprocessing(self, labels, images, workspace\_path, params)”
- “def onTraining(self, monitor\_handler)” on page 54
- “def onCompleted(self, model\_path)” on page 55
- “def onFailed(self, train\_status, e, tb\_message):” on page 55
- “Monitoring and reporting statistics” on page 55

### Template

This is a template you can use for the MyTrain API:

```
class MyTrain(TrainCallback):
    def __init__():
        pass
    def onPreprocessing(self, labels, images, workspace_path, params):
        pass
    def onTraining(self, monitor_handler):
        pass
    def onCompleted(self, model_path):
        pass
    def onFailed(self, train_status, e, tb_message):
        pass
```

### def onPreprocessing(self, labels, images, workspace\_path, params)

Callback for data set preprocessing.

### Input

#### *labels* (dict)

Image categories and index.

Example: {'safety\_vest': 1, 'helmet': 0, 'no\_safety\_vest': 2, 'no\_helmet': 3}

### *images*

- *image classification* (dict): Image path and its category.

Example: {'/dataset/Acridtheres/001.jpg': 'Acridtheres', '/dataset/Butorides/002.jpg': 'Gallinula', '/dataset/Butorides/003.jpg': 'Butorides'}

- *object detection* (list): List of annotation objects; including the image name and annotation.

Example:

```
[annotation[0] annotation[1] ...]
image filename
annotations[0].filename: /dataset/safety-detection/ee1fba93-a5f0-4c8b-8496-ce7605914651.jpg
image size [width, height, depth]
annotations[0].size: [450, 330, 3]
bounding box #0 label
annotations[0].object[0].label: helmet
# bounding box #0 position [xmin, ymin, xmax, ymax]
annotations[0].object[0].bbox: [111, 16, 205, 106]
annotations[0].object[1].label: helmet
annotations[0].object[1].bbox: [257, 42, 340, 140]
annotations[0].object[2].label: safety_vest
annotations[0].object[2].bbox: [40, 105, 215, 291]
annotations[0].object[3].label: safety_vest
annotations[0].object[3].bbox: [207, 124, 382, 309]
```

### *workspace\_path* (string)

Temporary workspace path recommended to be used in all training life cycles.

Example: "/tmp/workspace"

### *params* (dict)

Hyper parameters for training. These parameters are available to the custom model, but they are not required.

- Object detection example:

```
{ 'max_iter' : 4000, 'learning_rate' : 0.001, 'weight_decay' : 0.0005, 'momentum' :
0.9 , 'train_test_ratio' : 0.8 }
```

- Classification example:

```
{ 'max_iter' : 4000, 'learning_rate' : 0.001, 'weight_decay' : 0.0005,
'test_iteration' : 100, 'test_interval' : 20}
```

### **Output:**

None

### **def onTraining(self, monitor\_handler)**

Callback for training.

### **Input**

*monitor\_handler* (MonitorHandler): Handler for train/test status monitoring.

### **Output**

None

**def onCompleted(self, model\_path)**

Callback for training completed. A training task is terminated either with `onCompleted()` or with `onFailed()`. You need to save the trained model in this callback.

**Input**

*model\_path* (String): The absolute model path and file.

**Output**

None

**def onFailed(self, train\_status, e, tb\_message):**

Callback for training failed. A train task is terminated either with `onCompleted()` or with `onFailed()`

**Input**

*train\_status* (string)

Training status when the failure occurred.

*e* (Exception object)

Programming exception object.

*tb\_message* (string)

Formatted traceback message.

**Output**

None

**Monitoring and reporting statistics**

The `onTraining` API passes a `monitor_handler` object. This object provides callbacks to report both training and test messages back to PowerAI Vision. Depending on the type of training being performed, classification or object detection, the appropriate callback must be used.

*Object detection callbacks:*

Use this callback when the custom model is trained for object detection.

- “`def updateTrainMetrics(current_iter, max_iter, loss_cls, loss_bbox, epoch)`”
- “`def updateTestMetrics(mAP)`” on page 56

**def updateTrainMetrics(current\_iter, max\_iter, loss\_cls, loss\_bbox, epoch)**

Handler for status updates from the training process. This should be called **actively** by your custom code to post training status to the PowerAI Vision user interface.

**Input**

*current\_iter* (int)

Current iteration in the epoch

*max\_iter* (int)

Maximum iterations in one epoch

*loss\_cls* (float)

Training loss of classification

***loss\_bbox (float)***

Training loss of bounding box prediction

***epoch (int)***

Current training epoch

### Example

```
monitor_handler.updateTrainMetrics(current_iter, max_iter, loss_cls, loss_bbox, epoch)
```

### Output

None

### **def updateTestMetrics(mAP)**

Handler for status updates from the testing process. This should be called **actively** by your custom code to post testing status to the PowerAI Vision user interface.

### Input

*mAP* (float): Testing mean average precision

### Example

```
monitor_handler.updateTestMetrics(mAP)
```

### Output

None

*Classification callbacks:*

Use this callback when the custom model is trained for image classification.

- “def updateTrainMetrics(current\_iter, max\_iter, loss, epoch)”
- “def updateTestMetrics(current\_iter, accuracy, loss, epoch)” on page 57

### **def updateTrainMetrics(current\_iter, max\_iter, loss, epoch)**

Handler for status updates from the training process. This should be called **actively** by your custom code to post training status to the PowerAI Vision user interface.

### Input

***current\_iter (int)***

Current iteration in the epoch

***max\_iter (int)***

Maximum iterations in one epoch

***loss (float)***

Training loss

***epoch (int)***

Current training epoch

### Example

```
monitor_handler.updateTrainMetrics(current_iter, max_iter, loss, epoch)
```

### Output

None

```
def updateTestMetrics(current_iter, accuracy, loss, epoch)
```

Handler for status updates from the testing process. This should be called **actively** by your custom code to post testing status to the PowerAI Vision user interface.

### Input

*current\_iter* (int)

Current iteration in the epoch

*accuracy* (float)

Testing accuracy

*loss* (float)

Training loss

*epoch* (int)

Current training epoch

### Example

```
monitor_handler.updateTrainMetrics(iter_num, accuracy, loss, epoch_num)
```

### Output

None

## Preparing a model that will be deployed in PowerAI Vision

If your custom model will be deployed in the PowerAI Vision framework, your custom model must meet the following requirements.

After the model is properly prepared, import it to PowerAI Vision by navigating to the Models page and clicking **Import .zip file**. To deploy the model, on the Models page, select the model and click **Deploy model**.

### Custom model requirements:

- It must be TensorFlow based.
- It must implement the MyDeploy Python class.
  - The MyDeploy implementation must reside in a file named `deploy.py` in the top level directory of the zip file contents.
  - The following import must be added to the `deploy.py` file in order to define the training callbacks:  
`from deploy_interface import DeployCallback`
  - The class name must be `MyDeploy`.

### MyDeploy Template:

```
class MyDeploy(DeployCallback):  
    def __init__(self):  
        pass  
    def onModelLoading(self, model_path, labels, workspace_path):  
        pass  
    def onTest(self):
```

```

    pass
def onInference(self, image_url, params):
    pass
def onFailed(self, deploy_status, e, tb_message):
    pass

```

### **class MyDeploy(DeployCallback):**

Use the MyDeploy API to prepare a TensorFlow model that will be deployed in PowerAI Vision.

### **Template**

This is a template you can use for the MyDeploy API:

```

class MyDeploy(DeployCallback):
def __init__(self):
    pass
def onModelLoading(self, model_path, labels, workspace_path):
    pass
def onTest(self):
    pass
def onInference(self, image_url, params):
    pass
def onFailed(self, deploy_status, e, tb_message):
    pass

```

### **def onModelLoading(self, model\_path, labels, workspace\_path)**

Callback for load model.

### **Input**

#### ***model\_path* (string)**

Model path. The model must be decompressed before this callback.

#### ***workspace\_path* (string)**

Temporary workspace path recommended to be used in all deploy activities.

#### ***labels* (dict)**

The label index to name mapping.

Example: {1: 'safety\_vest', 0: 'helmet', 2: 'no\_safety\_vest', 3: 'no\_helmet'}

### **Output:**

None

### **def onTest(self)**

Test API interface with a custom message.

### **Input**

None

### **Output**

*message* (string): Output message.

**def onInference(self, image\_url, params)**

Inference with a single image.

### Input

**image\_url (string)**

Path of the image for inference.

**params (dict)**

Additional inference options.

**heatmap (string)**

Request a heat map. This is only supported for classification. Possible values:

- "true" : A heat map is requested.
- "false" : A heat map is not requested.

**conf\_threshold (float)**

Confidence threshold. Value in the range 0.0 - 1.0, to be treated as a percentage. Only results with a confidence greater than the specified threshold are returned. The smaller confidence threshold you specify, the more results are returned. If you specify 0, many, many results will be returned because there is no filter based on the confidence level of the model.

### Output (classification)

result({"label": "apple", "confidence": 0.9, "heatmap": "\_value\_}): predicted label and its score  
label (string) : predicted label name  
confidence (float) : number for certainty. between 0 and 1  
heatmap (string) : heatmap return

### Output (object detection)

result([{"confidence": 0.95, "label": "badge", "ymax": 145, "xmax": 172, "xmin": 157, "ymin": 123}]): predicted  
confidence(float): number for certainty. between 0 and 1  
label(string): predicted label name  
ymax(int): the max Y axis of bounding box  
xmax(int): the max X axis of bounding box  
ymin(int): the min Y axis of bounding box  
xmin(int): the min X axis of bounding box

**def onFailed(self, deploy\_status, e, tb\_message)**

Callback for deploy failed. A deploy task is terminated with onFailed().

### Input

**deploy\_status (string)**

Deploy status when the failure occurred.

**e (Exception object)**

Programming exception object.

**tb\_message (string)**

Formatted traceback message.

### Output

None

---

## Deploying a trained model

Deploy a trained model to get it ready to use within PowerAI Vision or a different program, such as IBM PowerAI. Deploying a model creates a unique API endpoint based on that model for inference operations.

**Note:** PowerAI Vision assigns one GPU to each active training job or deployed deep learning API. For example, if a system has four GPUs and you have two deployed web APIs, there are two GPUs available for active training jobs.

To deploy the trained model, follow these steps:

1. Click **Models** from the menu.
2. Select the model you want to deploy and click **Deploy**.
3. Specify a name for the model, and for models that were trained with the **Optimized for speed (tiny YOLO v2)** model, choose the accelerator to deploy to. You can choose GPU, CPU, or Xilinx FPGA - 16 bit (technology preview).

**Note:** Deploying a model to a Xilinx FPGA requires the Xilinx Alveo U200 Accelerator card.

4. Click **Deploy**. The Deployed Models page is displayed. When the model has been deployed, the status column displays **Ready**.
5. Click the deployed model to get the API endpoint, to view details about the model, such as the owner and the accuracy, and to test other videos or images against the model. For information about using the API see Vision Service API documentation.

**Note:** When using the API, the smaller confidence threshold you specify, the more results are returned. If you specify 0, many, many results will be returned because there is no filter based on the confidence level of the model.

6. If necessary, you can delete a deployed model. To delete a deployed model, click **Deployed Models**. Next, select the model that you want to delete and click **Delete**. The trained model is not deleted from PowerAI Vision.

### Related concepts:

“Working with the user interface” on page 45

The PowerAI Vision user interface is made up of these basic parts: the title bar, the side bar, the action bar, the data area, and the notification center.

“Understanding metrics” on page 66

PowerAI Vision provides several metrics to help you measure how effectively your model has been trained.

### Related information:

 [Vision Service API documentation](#)

---

## PowerAI Vision REST APIs

You can use REST APIs to work with PowerAI Vision data sets and models, such as performing training and deployment. You can also use them to perform administrative tasks, such as monitoring events. These APIs allow you to bypass the user interface and automate PowerAI Vision processes or solutions.

For information about using the API see Vision Service API documentation. There are also examples of using the APIs for different actions, published at <https://github.com/IBM/powerai/tree/master/vision/tools/vapi/cli>.

---

## Testing a model

After deploying your model, you should test it against other images and videos to make sure that it works as expected.

1. Click **Deployed Models** from the menu.
2. Click the model you want to test. The model opens in the Deployed model page.
3. Use the **Test Videos** (object detection models only) or **Test Images** areas to upload images and videos, one at a time.
4. The results are shown on the bottom of the window.
  - If you used an image to test an image classification model, test result displays the uploaded picture and test result heat map, and gives the classification and the confidence of the classification. Multiple classes are returned with the decreasing levels of confidence for the different classes. The returned heat map is for the highest confidence classification and can help you determine whether the model has correctly learned the features of this classification. The red area of the heat map corresponds to the highest relevance areas of the picture.
  - If you used an image to test an object detection model, the identified objects are labeled in the image, with the calculated precision.
  - If you used a video to test an object detection model, the video is processed, then as you watch the processed video, the identified objects are labeled as they appear in the video. Processing the video might take a while, depending on its size.
5. If you are satisfied with the results, the model is ready to be used in production. Otherwise, you can refine the model by following the instructions in this topic: “Refining a model.”

---

## Refining a model

After deploying a model, you can improve its accuracy by supplying more data. There are several methods you can use to add more data to the model.

You can add more data by using any combination of the following options:

1. Upload new images or videos to the data set and classify or label them as appropriate.
2. For an existing video, capture more frames and classify or label them as appropriate.
3. Use data augmentation. *Data augmentation* is the use of filters, such as blur and rotate, to create new versions of existing images. When you use data augmentation, a new data set is created that contains all of the existing images, plus the newly generated images. For instructions, see “Augmenting the data set” on page 63.
4. For models trained for object detection, you can use the Auto label function to identify more objects in the existing data. See “Automatically labeling objects” for instructions.

After adding more data, train the model again.

## Automatically labeling objects

After deploying a model for object detection, you can improve its accuracy by using the Auto label function. This function improves the model's accuracy by quickly adding more data to the data set.

### Notes:

- You can automatically label images or videos that have not had labels manually added. If any labels have been manually added, that image or frame is skipped.
- If labels have been added through auto label, those images and frames are reprocessed. The previous labels are removed and new labels are added.
- When performing multiple auto label operations on the same video, it is possible to get multiple frames with the same time offset. This situation can occur when the intervals overlap and labels have been edited on the frames at the overlap points.

For example, labeling at a 10 second interval, editing some of the labels on those frames, and then labeling again at 5 a second interval has an overlap every 10 seconds. There might be duplicate images at each of the 10 second intervals with edited labels.

- You can manually add labels to images and frames that have been auto labeled. If you do so, and then run **Auto label** again, those images and frames are skipped because they now contain manually added labels. The automatically added labels are not removed.

When you automatically label objects, an existing trained model is used to generate new labels in the data set.

1. Open the data set that you want to add more data to and select **Auto label**.
2. Select the trained model to use, then click **Auto label**.
3. Labels are added to existing images or videos have not been manually labeled. The automatically added labels are green. Manually added labels are blue. For videos, if frames have already been captured, those frames are used for auto labeling. If frames have not been captured, the video is divided into frames and then labeled.

### Automatically labeling objects in a data set

When you auto label a data set, an existing trained model is used to generate labels for images and video frames that have not been manually labeled.

#### Notes:

- All images and frames that were not manually labeled are processed. Therefore, objects that contain only labels that were added by using the auto label function are reprocessed. The previous labels are removed and new labels are added.
- When auto labeling a data set, only images and frames are auto labeled. Therefore, any videos that do not have captured frames are skipped.

Follow these steps to generate new labels in the data set.

1. Open the data set that you want to add more data to and select **Auto label**.
2. Select the trained model to use, then click **Auto label**.
3. Labels are added to existing images or videos have not been manually labeled. The automatically added labels are green. Manually added labels are blue.

You can manually add labels to images and frames that have been auto labeled, and you can manipulate (move, resize) the labels that were automatically generated. If you do so, and then run **Auto label** again, those images and frames are skipped because the frame was manually edited. The automatically added labels are not removed.

### Automatically labeling videos

When using the auto label function on a data set, only frames and images processed. Videos are ignored. However, you can run the auto label function on an individual video.

**Note:** Any frames that were previously captured by using auto capture and were not manually labeled are deleted before auto labeling. This helps avoid labeling duplicate frames. Manually captured frames are not deleted.

Follow these steps to run the auto label function on a video.

1. Open the data set that contains the video.
2. Select the video and click **Label objects**.
3. Click **Auto label** then choose the time interval to capture frames and the trained model to use for labeling, then click **Auto label**.
4. Frames are captured at the specified interval and labels are added by using the specified trained mode. The automatically added labels are green.

After processing, you can manually add labels to the frames that have been auto labeled and you can manipulate (move, resize) the labels that were automatically generated. However, when performing multiple auto label operations on the same video, it is possible to get multiple frames with the same time offset. This situation can occur when the intervals overlap and labels have been edited on the frames at the overlap points.

For example, labeling at a 10 second interval, editing some of the labels on those frames, and then labeling again at 5 a second interval has an overlap every 10 seconds. There might be duplicate images at each of the 10 second intervals with edited labels.

## Augmenting the data set

After deploying a model, you can improve the model by using data augmentation to add modified images to the data set, then retraining the model. *Data augmentation* is the use of filters, such as blur and rotate, to create new versions of existing images. When you use data augmentation, a new data set is created that contains all of the existing images, plus the newly generated images.

To augment a data set, follow these steps:

1. Open the data set for a deployed model.
2. Select the images to use for augmentation, then click **Augment data**. If you select a video, every captured frame is used for augmentation. If you select some, but not all, frames in a video, only the selected frames are used for augmentation.
3. Choose any combination of filters to apply to your data set, then click **Continue**.  
Each filter generates one or more new versions of each selected image; the filters are not cumulative. For example, if you select **Sharpen** and **Flip horizontal**, six new images are generated; one flipped and five sharpened.  
When you select a filter, you can see an example of what that filter would do to an image. This sample image is **not** a live preview of the filter. It is an example of what an image might look like with that filter applied. Some filters, such as Blur and Sharpen, have additional settings you can choose.
4. Specify a name for the new data set and click **Create data set**.
5. The new data set, containing the original images, is created immediately. The augmented images are added after all processing completes. After the new data set is created, you can train a model based on the new data set. See this topic for instructions: “Training a model” on page 50.

## Augmentation settings

These settings are available when augmenting data.

Each filter generates one or more new versions of each selected image; the filters are not cumulative. For example, if you select **Sharpen** and **Flip horizontal**, six new images are generated; one flipped and five sharpened.

**Note:** When you select a filter, you can see an example of what that filter would do to an image. This sample image is **not** a live preview of the filter. It is an example of what an image might look like with that filter applied.

**Blur** Select the maximum amount of Gaussian and motion blur. Gaussian blur makes the entire image appear out of focus by reducing detail and noise. Motion blur makes the image appear as if it (or the camera) is in motion.

Five new images are generated in the range of each nonzero selection. For example, if Motion = 25 and Gaussian = 10, then five images are generated by applying a motion blur filter in random strengths in the range 0-25, and five additional images are generated by applying a Gaussian blur filter in the range 0-10.

### **Sharpen**

Select the maximum amount of sharpening to apply. Some noise will be introduced. Five new images are generated in the specified range. For example, if Sharpness = 25, five new images are generated by applying the sharpen filter in random strengths in the range of 0-25.

**Color** Select the maximum amount of change in the image's brightness, contrast, hue, and saturation. Five new images are generated by using randomly selected values in the selected ranges. The resultant values can be either positive or negative.

For example, if Brightness = 30, Contrast = 15, Hue = 5, and Saturation = 10, five images are generated that have brightness changed by (-30, 30)% , contrast is changed by (-15, 15)%, and so on.

**Crop** Select the maximum percentage of the image that should remain. For example, selecting 25 means that at most 25% of the original image remains and 75% is removed. Five new images will be generated that are cropped in the selected range. The crop is centered at a random point.

For example, if Crop = 25, five images are generated cropped to retain 100% - 25% of the original image.

### **Vertical flip**

Create a new image by flipping the existing image across the top edge. That is, the top of the image becomes the bottom.

### **Horizontal flip**

Create a new image by flipping the existing image across the side edge. That is, the left side of the image becomes the right side.

**Rotate** Select the maximum value of rotation for the new images. Rotation can be either clockwise or counter-clockwise. Five new images are generated that are rotated by this amount. For example, if this value is 45, five new images are generated that are rotated either clockwise or counter-clockwise by a random number in the range 0-45.

**Noise** Select the maximum amount of noise to add to the new images, specified as a percentage of what PowerAI Vision determines to be a reasonable amount of noise for the images to remain usable. Therefore, if you select 100, none of the generated images will have 100% noise added. Instead, the output images will possibly have the maximum amount of noise added while still remaining usable.

Five new images are generated with noise added in the specified range. For example, if this value is 25, five new images are created with a random amount of noise added in the range 0 - 25% of a reasonable amount of noise.

---

## **Importing and exporting PowerAI Vision information**

You can import and export PowerAI Vision models and data sets. This allows you to save them for archiving then use them later, use them on a different PowerAI Vision install, and so on.

- "Exporting"
- "Importing" on page 65

### **Exporting**

#### **Export a data set**

To export a data set, open the Data sets page, open the data set you want to export, then click **Export** in the action bar. The data set is saved in your default download directory as `data_set_name.zip`. This zip file contains the images as well as any tags or categories you have assigned.

**Notes:**

- When exporting a data set, any objects that are not used in the data set are not contained in the exported data set. Therefore, they are not included when the data set is imported.  
For example, if the object or label “car” is defined but is not used in any of the images in the data set, the exported data set does not include the “car” object or label. When the data set is imported, the “car” object or label is not created.
- In PowerAI Vision 1.1.1, any information about augmented images is lost on export. Therefore, if the data set is later imported (regardless of the product version), the augmented images will be in the data set, but they will no longer be marked as augmented.

### Export a model

When you export custom trained model (a model that was trained by using a custom model), the generated zip file is not encrypted or password protected, unlike with other models exported from PowerAI Vision.

To export a model, open the Models page, select the model you want to export, then click **Export** in the left pane. The model is saved in your default download directory as *character\_string.zip*; where *character\_string* is randomly generated by the system.

#### Note:

If the model is not a Custom model that was imported, the exported model can only be used with PowerAI Vision. It can be imported into the Inference Server product and deployed with the Inference Server product.

It is not recommended that you use an exported model with an earlier version of the product than it was exported from. Additionally, a model from a prior version will not have support for features that were added to later versions of the product. That is, if you export a model from version *x.1* and import it into *x.2*, features that were added in *x.2* will not be supported on the imported model.

## Importing

### Import a data set

1. Navigate to the Data sets page.
2. Drag and drop an exported data set .zip file onto the **Create** box.

**Important:** After the upload starts, do not close the PowerAI Vision tab or refresh the page. Doing so stops the upload.

3. After the upload completes, the data set has its original name.

#### Notes:

- In PowerAI Vision 1.1.1, any information about augmented images is lost on export. Therefore, if the data set is later imported (regardless of the product version), the augmented images will be in the data set, but they will no longer be marked as augmented.
- The data set associated with a model is not preserved when it is exported. Therefore, for imported models, the Data set field is set to “Not found”.

### Import a model

Instead of using PowerAI Vision to train a new model, you can *import* a model that was previously trained with PowerAI Vision, and was then exported. This lets you streamline data processing by offloading training tasks and allowing you to reuse models on multiple systems. After the model is imported to the Models page, you can deploy it in PowerAI Vision. Use the information in this topic to prepare a custom model that will be deployed in PowerAI Vision: “Preparing a model that will be deployed in PowerAI Vision” on page 57.

1. Navigate to the Models page.

2. Drag and drop a previously exported model .zip file onto the **Import** box.

**Important:** After the upload starts, do not close the PowerAI Vision tab or refresh the page. Doing so stops the upload.

3. After the upload completes, the model has its original name.

---

## Understanding metrics

PowerAI Vision provides several metrics to help you measure how effectively your model has been trained.

To understand these metrics, you must understand these terms:

### True positive

A *true positive* result is when PowerAI Vision correctly labels or categorizes an image. For example, categorizing an image of a cat as a cat.

### False positive

A *false positive* result is when PowerAI Vision labels or categorizes an image when it should not have. For example, categorizing an image of a cat as a dog.

### True negative

A *true negative* result is when PowerAI Vision correctly does not label or categorize an image. For example, not categorizing an image of a cat as a dog.

### False negative

A *false negative* result is when PowerAI Vision does not label or categorize an image, but should have. For example, not categorizing an image of a cat as a cat.

Of course, for a model in production, the values for true negative / positive and false negative / positive can't accurately be known. These values are the expected values for these measurements.

- “Metrics for image classification (Trained for accuracy)”
- “Metrics for object detection (Trained for accuracy)” on page 67
- “Metrics for object detection using the Tiny Yolo model (Trained for speed)” on page 68
- “Metrics for custom models” on page 68

## Metrics for image classification (Trained for accuracy)

### Accuracy

Measures the percentage of correctly classified images. It is calculated by (true positives + true negatives) / (true positives + true negatives + false positives + false negatives).

### PR curve (Advanced)

The precision-recall (PR) curve plots precision vs. recall (sensitivity). Because precision and recall are typically inversely related, it can help you decide whether the model is appropriate for your needs. That is, do you need a system with high precision (fewer results, but the results are more likely to be accurate), or high recall (more results, but the results are more likely to contain false positives)?

### Precision

Precision tells describes how "clean" our population of hits is. It measures the percentage of images that are correctly classified. That is, when the model classifies an image into a category, how often is it correct? It is calculated by true positives / (true positives + false positives).

### Recall

The percentage of the images that were classified into a category, compared to all images

that should have been classified into that category. That is, when an image belongs in a category, how often is it identified? It is calculated as  $\text{true positives} / (\text{true positives} + \text{false negatives})$ .

### Confusion matrix (Advanced)

The confusion matrix is used to calculate the other metrics, such as precision and recall. Each column of the matrix represents the instances in a predicted class (those that PowerAI Vision marked as belonging to a category). Each row represents the instances in an actual category. Therefore, each cell measures how many times an image was correctly and incorrectly classified.

You can view the confusion matrix as a table of values or a heat map. A heat map is a way of visualizing the data, so that the higher values appear more “hot” (closer to red) and lower values appear more “cool” (closer to blue). Higher values show more confidence in the model.

This matrix makes it easy to see if the model is confusing categories, or not identifying certain categories.

## Metrics for object detection (Trained for accuracy)

### Accuracy

Measures the percentage of correct image classifications. It is calculated by  $(\text{true positives} + \text{true negatives}) / \text{all cases}$ .

### Mean Average precision (mAP)

The average over all classes of the maximum *precision* for each object at each *recall* value. Precision measures how accurate the model is. That is, the percent of the classified objects that are correct. Recall measures how well the model returns the correct objects. For example, out of 100 images of dogs, how many of them were classified as dogs?

To calculate this, first, the PR curve is found. Then, the maximum precision for each recall value is determined. This is the maximum precision for any recall value greater than or equal to the current recall value. For example, if the precision values range from .35 to .55 (and then never reach .55 again) for recall values in the interval .3 - .6, then the maximum precision for every recall value in the interval .3 - .6 is set to .55.

The mAP is then calculated as the average of the maximum precision values.

### IoU (Intersection over union)

The accuracy of the location and size of the image label boxes.

It is calculated by the intersection between a ground truth bounding box and a predicted bounding box, divided by the union of both bounding boxes; where the intersection is the area of overlap, a *ground truth* bounding box is the hand drawn box, and the *predicted bounding box* is the one drawn by PowerAI Vision.

### Confusion matrix (Advanced)

The confusion matrix is used to calculate the other metrics, such as precision and recall. Each column of the matrix represents the instances in a predicted class (those that PowerAI Vision marked as belonging to a category). Each row represents the instances in an actual category. Therefore, each cell measures how many times an image was correctly and incorrectly classified.

You can view the confusion matrix as a table of values or a heat map. A heat map is a way of visualizing the data, so that the higher values appear more “hot” (closer to red) and lower values appear more “cool” (closer to blue). Higher values show more confidence in the model.

This matrix makes it easy to see if the model is confusing categories, or not identifying certain categories.

### PR curve (Advanced)

The precision-recall (PR) curve plots precision vs. recall (sensitivity). Because precision and recall are typically inversely related, it can help you decide whether the model is appropriate for your

needs. That is, do you need a system with high precision (fewer results, but the results are more likely to be accurate), or high recall (more results, but the results are more likely to contain false positives)?

**Precision**

Precision tells describes how "clean" our population of hits is. It measures the percentage of objects that are correctly identified. That is, when the model identifies an object, how often is it correct? It is calculated by  $\text{true positives} / (\text{true positives} + \text{false positives})$ .

**Recall** The percentage of the images that were labeled as an object, compared to all images that contain that object. That is, how often is an object correctly identified? It is calculated as  $\text{true positives} / (\text{true positives} + \text{false negatives})$ .

## **Metrics for object detection using the Tiny Yolo model (Trained for speed)**

**Accuracy**

Measures the percentage of correctly classified objects. It is calculated by  $(\text{true positives} + \text{true negatives}) / (\text{true positives} + \text{true negatives} + \text{false positives} + \text{false negatives})$ .

## **Metrics for custom models**

When a custom model is imported and deployed, the following metric is shown:

**Accuracy**

Measures the percentage of correct categorizations. It is calculated by  $(\text{true positives} + \text{true negatives}) / (\text{true positives} + \text{true negatives} + \text{false positives} + \text{false negatives})$ .

---

## Scenario: Detecting objects in a video

In this fictional scenario, you want to create a deep learning model to monitor traffic on a busy road. You have a video that displays the traffic during the day. From this video, you want to know how many cars are on the busy road every day, and what are the peak times that have the most cars on the road.

The video file used in this scenario is available for download here: [Download car video](#).

To create a deep learning model, you will perform the following steps:

1. Importing a video
2. Labeling objects in a video
3. Training a model
4. Deploying a model
5. Automatically label frames in a video

### Import a video and create a data set

First, create a data set and add videos to it.

1. Log in to PowerAI Vision.
2. Click **Data Sets** in the title bar to open the Data Sets page. There are several ways to create your new data set:
3. From the **Data set** page, click the icon and name the data set Traffic Video.
4. To add a video to the data set, click the Traffic Video data set and click **Import file** or drag the video to the + area.

**Important:** You cannot navigate away from the PowerAI Vision page or refresh until the upload completes. You can navigate to different pages within PowerAI Vision during the upload.

### Labeling objects in a video

The next step is to label objects in the video. For object detection, you must have at minimum five labels for each object. We will create Car and Motorcycle objects and will label at least five frames in the video with cars and at least five frames with motorcycles.

1. Select the video from your data set and select **Label Objects**.
2. Capture frames by using one of these methods:
  - Click **Auto capture frames** and specify a value for **Capture Interval (Seconds)** that will result in at least five frames. We will select this option and specify 10 seconds.

**Note:** Depending on the length and size of the video and the interval you specified to capture frames, the process to capture frames can take several minutes.

- Click **Capture frame** to manually capture frames. If you use this option, you must capture a minimum of five frames from the video.
3. If you used **Auto capture frames**, verify that there are enough of each object type in the video frames. If not, follow these steps to add new frames to the existing data set.

In this scenario, the motorcycle is only in a single automatically captured frame at 40 seconds. Therefore, we must capture at least four more frames with the motorcycle. The motorcycle comes into view at 36.72 seconds. To correctly capture the motorcycle in motion we will create extra frames at 37.79 seconds, 41.53 seconds, and 42.61 seconds.

- a. Play the video. When the frame you want is displayed, click pause.

- b. Click **Capture Frame**.
  4. Create new object labels for the data set by clicking **Add object** in the left pane. Enter Car, click **+**, then enter Motorcycle, then click **Add**.
  5. Tag the objects in the frames:
    - Select the first frame in the carousel.
    - Select the correct object label, for example, "Car".
    - Click and hold down the left mouse button to draw a box around the object.
- Review the following tips about identifying and drawing objects in video frames and images:
- Do not label part of an object. For example, do not label a car that is only partially in the frame.
  - If a frame has more than one object, you must label all objects. For example, if you have cars and motorcycles in a video frame, you must label the cars and the motorcycles. Label objects with a consistent approach.
  - Draw a box around each individual object. Do not draw a box around groups of objects. For example, if two cars are right next to each other, you must draw a box around each car.
  - Draw the box as close to the objects as possible. Do not leave blank space around the objects.
  - You can draw boxes around objects that overlap. For example, if one object is behind another object, you can label them both. However, it is recommended that you only label objects if the majority of the object is visible. You can also draw boxes that touch each other.
  - You can zoom in on the frame to make it easier to draw the box around objects.
  - Boxes cannot extend off the edge of the frame.
6. After all objects are labeled in all of the images in the carousel, click **Done editing**.

The following figure displays the captured video frame at 41.53 seconds with object labels of **Car** and **Motorcycle**. Figure 1 also displays a box around the five frames (four of the frames were added manually) in the carousel that required object labels for the motorcycle that is in each frame.



Figure 7. Labeling objects in PowerAI Vision

## Training a model

With all the object labels that are identified in your data set, you can now train your deep learning model. To train a model, complete the following steps:

1. From the Data set page, click **Train**.
2. Fill out the fields on the Train Data set page, ensuring that you select **Object Detection**. We will choose **System Default** for **Performance Type**
3. Click **Train**.
4. (Optional) Stop the training process by clicking **Stop training > Keep Model > Continue**.

You can wait for the entire training model process complete, or you can stop the training process when the lines in the training graph start to flatten out. This is because improvements in quality of training might plateau over time. Therefore, the fastest way to deploy a model and refine the data set is to stop the process before quality stops improving.

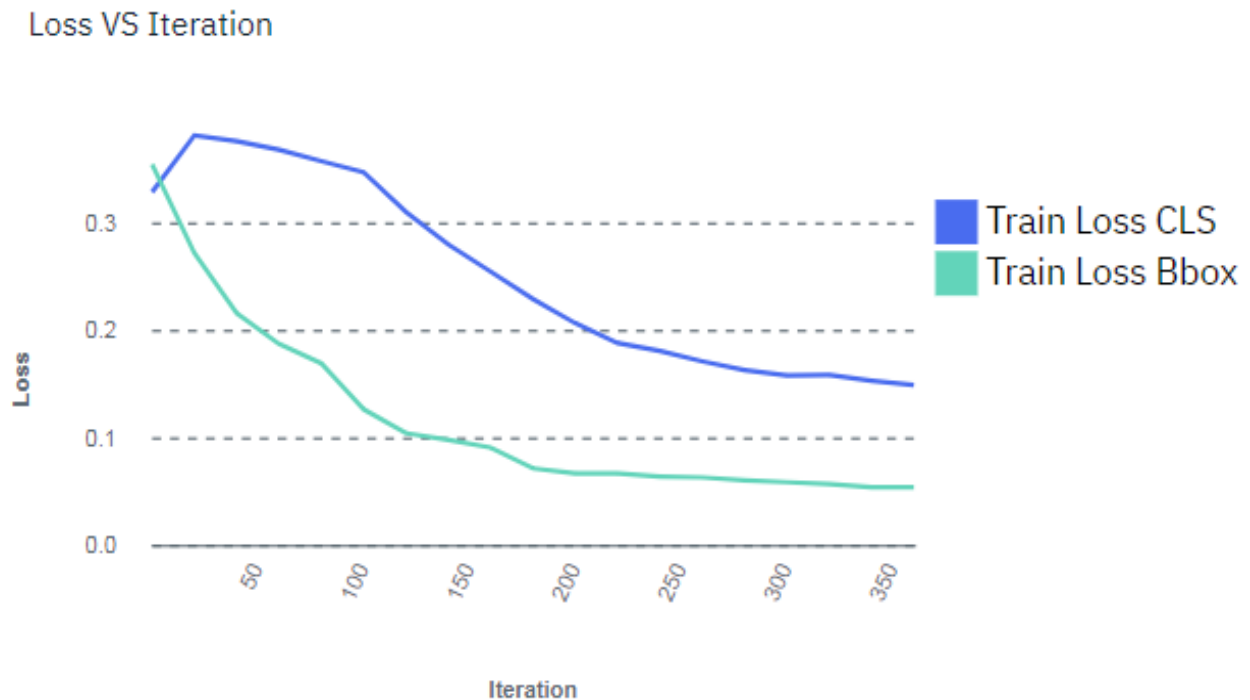


Figure 8. Model training graph

**Important:** If the training graph converges quickly and has 100% accuracy, the data set does not have enough information. The same is true if the accuracy of the training graph fails to rise or the errors in the graph do not decrease at the end of the training process. For example, a model with high accuracy might be able to discover all instances of different race cars, but might have trouble differentiating between specific race cars or those that have different colors. In this situation, add more images or video frames to the data set, tag them, then try the training again.

## Deploying a trained model

A deployed model uses one GPU on the system. To deploy the trained model, complete the following steps:

1. Click **Models** from the menu.
2. Select the model you created in the previous step and click **Deploy**.

3. Specify a name for the model, and click **Deploy**. The Deployed Models page is displayed, and the model is deployed when the status column displays **Ready**.
4. Double-click the deployed model to get the API endpoint and test other videos or images against the model. For information about using the API see Vision Service API documentation.

## Automatically label frames in a video

You can use the auto label function to automatically identify objects in the frames of a video after a model has been deployed.

In this scenario, you have only nine frames. To improve the accuracy for your deep learning model, you can add more frames to the data set. Remember, you can rapidly iterate by stopping the training on a model and checking the results of the model against a test data set. You can also use the model to auto label more objects in your data set. This process improves the overall accuracy of your final model.

To use the auto label function, complete the following steps:

**Note:** Any frames that were previously captured by using auto capture and were not manually labeled are deleted before auto labeling. This helps avoid labeling duplicate frames. Manually captured frames are not deleted.

1. Click **Data sets** from the menu, and select the data set that you used to create the previously trained model.
2. Select the video in the data set that had nine frames, and click **Label Objects**.
3. Click **Auto label**.
4. Specify how often you want to capture frames and automatically label the frames. Select the name of the trained model that you deployed in step 3, and click **Auto label**. In this scenario, you previously captured frames every 10 seconds. To improve the accuracy of the deep learning model by capturing and labeling more frames, you can specify 6 seconds.
5. After the auto label process completes, the new frames are added to the carousel. Click the new frames and verify that the objects have the correct labels. The object labels that were automatically added are green and the object labels you manually added are in blue. In this scenario, the carousel now has 17 frames.

## Next steps

After processing, you can manually add labels to the frames that have been auto labeled and you can manipulate (move, resize) the labels that were automatically generated. However, when performing multiple auto label operations on the same video, it is possible to get multiple frames with the same time offset. This situation can occur when the intervals overlap and labels have been edited on the frames at the overlap points.

For example, labeling at a 10 second interval, editing some of the labels on those frames, and then labeling again at 5 a second interval has an overlap every 10 seconds. There might be duplicate images at each of the 10 second intervals with edited labels.

You can continue to refine the data set as much as you want. When you are satisfied with the data set, you can retrain the model by completing steps 1 - 3. This time when you retrain the model, you might want to train the model for a longer time to improve the overall accuracy of the model. The loss lines in the training model graph should converge to a stable flat line. The lower the loss lines are in the training graph the better. After the training completes, you can redeploy the model by completing steps 1 - 3. You can double-click the deployed model to get the API endpoint and test other videos or images against the model.

**Related concepts:**

“Working with the user interface” on page 45

The PowerAI Vision user interface is made up of these basic parts: the title bar, the side bar, the action bar, the data area, and the notification center.

“Understanding metrics” on page 66

PowerAI Vision provides several metrics to help you measure how effectively your model has been trained.

**Related information:**

 [Vision Service API documentation](#)



---

## Scenario: Classifying images

The goal of this example is to train a model to classify images of birds into groups based on their physiological similarities. Once the model is trained with a known dataset, users can upload new data sets to auto classify the birds into their respective categories. We will prepare the data, create a data set, train the model, and test the model.

### 1. Prepare the data.

Data preparation consists of gathering two types of data, *training data* and *test data*. Training data is used to teach the neural network features of the object so that it can build the classification model. Test data is used to validate the accuracy of the trained model. Our data will include pictures of different types of birds.

#### Notes:

- Different images should be used for training data and test data.
  - Images must be in one of these formats:
    - JPEG
    - PNG
- ### 2. Create a data set. Log in to the PowerAI Vision user interface, click Data Sets in the title bar, click **Create new data set** and name the data set Birds.
- ### 3. Populate the data set.
- a. In the left pane, expand Categories, click **Add category**. Add the “Acridotheres” category and click **Add**, then click **OK**.
  - b. Upload images of Acridotheres by dragging the images onto the **Drag files here** area.
  - c. In the left pane, click “Uncategorized”. The newly uploaded files are shown.
  - d. Click the **Select** box to select the images you just uploaded, then click **Assign category** and choose “Acridotheres”.
  - e. Repeat the above steps for the other categories.

**Note:** To train a model for classification, the data set must meet these requirements:

- There must be at least two categories.
  - Each category must have at least five images.
- ### 4. From the Data set page, click **Train**. In the Train data set window, choose **Image classification** and keep the default values for all other settings, then click **Train**.
- ### 5. After training is complete, click **Deploy model**.

**Important:** Each deployed model uses one GPU.

### 6. Test the trained model. On the Deployed models page, open the model you just deployed. Scroll down to the Test Images area and input a test image.

The test result displays the uploaded picture and test result heat map, and gives the classification and the confidence of the classification. Multiple classes are returned with the decreasing levels of confidence for the different classes. The returned heat map is for the highest confidence classification and can help you determine whether the model has correctly learned the features of this classification. The red area of the heat map corresponds to the highest relevance areas of the picture.

If you are not satisfied with the result, use the information in this topic to refine the model: “Refining a model” on page 61. Otherwise, the model is ready to be used in production.

#### Related concepts:

“Working with the user interface” on page 45

The PowerAI Vision user interface is made up of these basic parts: the title bar, the side bar, the action bar, the data area, and the notification center.

“Understanding metrics” on page 66

PowerAI Vision provides several metrics to help you measure how effectively your model has been trained.

**Related information:**



Vision Service API documentation

---

## Scenario: Detecting objects in images

In this fictional scenario, you want to create a deep learning model to determine the make and model of a car caught by a traffic camera.

The image file used in this scenario is available for download here: [Download car image](#).

To create a deep learning model, you will perform the following steps:

1. "Import images and create a data set"
2. "Labeling objects in an image"
3. "Training a model" on page 78
4. "Deploying a trained model" on page 79

### Import images and create a data set

First, create a data set and add images to it.

1. Log in to PowerAI Vision.
2. Click **Data Sets** in the title bar to open the Data Sets page. There are several ways to create your new data set:
3. From the **Data set** page, click the icon and name the data set Traffic camera.
4. To add an image to the data set, click the Traffic image data set and click **Import file** or drag the image to the + area.

**Important:** You cannot navigate away from the PowerAI Vision page or refresh until the upload completes. You can navigate to different pages within PowerAI Vision during the upload.

### Labeling objects in an image

The next step is to label objects in the images. For object detection, you must have at minimum five labels for each object. We will create "Black car" and "White car" objects and will label at least five images as black cars, and at least five as white cars.

1. Select the images from your data set and click **Tag Objects**.
2. Create new object labels for the data set by clicking **Add object** in the left pane. Enter Black car, click +, then enter Black car, then click **Add**.
3. Tag the objects in the images:
  - a. The first image is open in the data area, with thumbnails of all the selected image on the left side. Select the correct object label, for example, "Black car".
  - b. Click and hold down the left mouse button to draw a box around the object. Continue labeling objects in the image.
  - c. Select the thumbnail of the next image to open it. Add the appropriate tags, and continue through the rest of the images.

Review the following tips about identifying and drawing objects in images:

- Start tagging from the top of the image and work your way to the bottom. This will help keep labels from covering objects that have not been labeled yet.
- Do not label part of an object. For example, do not label a car that is only partially in the image.
- If an image has more than one object, you must label all objects. For example, if you have black cars and white cars in an image, label them all. Label objects with a consistent approach.

- Draw a box around each individual object. Do not draw a box around groups of objects. For example, if two cars are right next to each other, you must draw a box around each car.
  - Draw the box as close to the objects as possible. Do not leave blank space around the objects.
  - You can draw boxes around objects that overlap. For example, if one object is behind another object, you can label them both. However, it is recommended that you only label objects if the majority of the object is visible. You can also draw boxes that touch each other.
  - You can zoom in on the image to make it easier to draw the box around objects.
  - Boxes cannot extend off the edge of the frame.
4. After all objects are labeled in all of the image, click **Done editing**.

## Training a model

With all the object labels that are identified in your data set, you can now train your deep learning model. To train a model, complete the following steps:

1. From the Data set page, click **Train**.
2. Fill out the fields on the Train Data set page, ensuring that you select **Object Detection**. We will choose **System Default** for **Performance Type**
3. Click **Train**.
4. (Optional) Stop the training process by clicking **Stop training > Keep Model > Continue**.

You can wait for the entire training model process complete, or you can stop the training process when the lines in the training graph start to flatten out. This is because improvements in quality of training might plateau over time. Therefore, the fastest way to deploy a model and refine the data set is to stop the process before quality stops improving.

### Loss VS Iteration

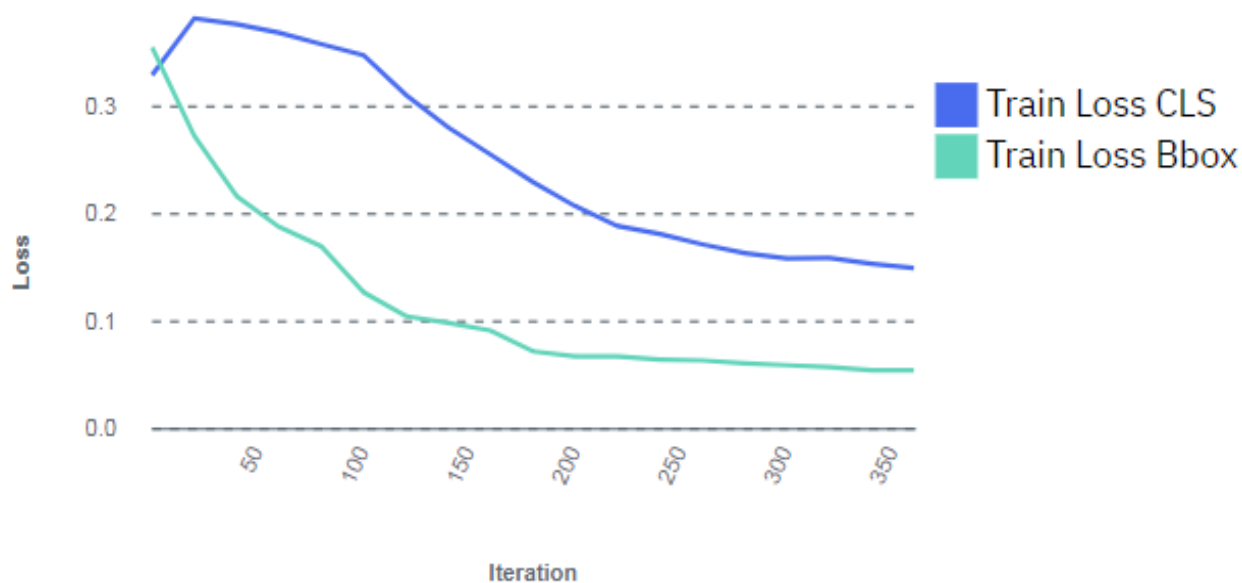


Figure 9. Model training graph

**Important:** If the training graph converges quickly and has 100% accuracy, the data set does not have enough information. The same is true if the accuracy of the training graph fails to rise or the errors in the graph do not decrease at the end of the training process. For example, a model with high accuracy

might be able to discover all instances of different race cars, but might have trouble differentiating between specific race cars or those that have different colors. In this situation, add more images or video frames to the data set, tag them, then try the training again.

## Deploying a trained model

A deployed model uses one GPU on the system. To deploy the trained model, complete the following steps:

1. Click **Models** from the menu.
2. Select the model you created in the previous step and click **Deploy**.
3. Specify a name for the model, and click **Deploy**. The Deployed Models page is displayed, and the model is deployed when the status column displays **Ready**.
4. Double-click the deployed model to get the API endpoint and test other videos or images against the model. For information about using the API see Vision Service API documentation.

## Next steps

You can continue to refine the data set as much as you want. When you are satisfied with the data set, you can train the model again. This time when you train the model, you might want to train the model for a longer time to improve the overall accuracy of the model. The loss lines in the training model graph should converge to a stable flat line. The lower the loss lines are in the training graph the better. After the training completes, you can deploy the model again. You can double-click the deployed model to get the API endpoint and test other images or images against the model.



---

## Administering PowerAI Vision

Use this information to administer PowerAI Vision, such as stopping, starting, and determining the status of the pods.

### Start or stop PowerAI Vision

There are several situations when you might need to stop and start PowerAI Vision. For example, when upgrading or performing maintenance on the product or on the system, when troubleshooting a problem, and so on. Use these commands to start or stop PowerAI Vision, as appropriate:

```
powerai-vision-stop.sh  
powerai-vision-start.sh
```

### Determine the status of PowerAI Vision pods

When troubleshooting a problem with PowerAI Vision, you might need to check the status of the Docker pods that are part of PowerAI Vision. For example, if the product does not start, if it is returning errors, or if actions are not completing. Run `kubectl get pods` to see the status. For example:

```
$ /opt/powerai-vision/bin/kubectl get pods  
NAME                                READY   STATUS    RESTARTS   AGE  
powerai-vision-mongodb-764f99fcf6-12nzd    1/1     Running   0          12h  
powerai-vision-portal-76fbc7db68-7rr47     1/1     Running   0          12h  
powerai-vision-postgres-55c6f7fcf6-42fht   1/1     Running   0          12h  
powerai-vision-taskanalyzer-55bfb587d4-cvzln 1/1     Running   0          12h  
powerai-vision-ui-845d8c8d8-bmf7          1/1     Running   0          12h  
powerai-vision-video-nginx-8474f7c44c-qmxm4 1/1     Running   0          12h  
powerai-vision-video-portal-5b76558784-8mb8d 1/1     Running   0          12h  
powerai-vision-video-rabmq-5d5d786f9f-nz7pn 1/1     Running   0          12h  
powerai-vision-video-redis-59c557b69-hf8pg  1/1     Running   0          12h  
powerai-vision-video-test-nginx-5dc6887666-19tb8 1/1     Running   0          12h  
powerai-vision-video-test-portal-54d85ff65b-945gp 1/1     Running   0          12h  
powerai-vision-video-test-rabmq-6858cc749-grhgm 1/1     Running   0          12h  
powerai-vision-video-test-redis-75977cdd8f-lbljb 1/1     Running   0          12h
```

If one or more pods is not running, try stopping and restarting PowerAI Vision.

---

## Managing users

There are two kinds of users in PowerAI Vision: administrators, and everyone else. The way you work with users and passwords differs, depending on how PowerAI Vision is installed.

- “Types of users”
- “PowerAI Vision installed as stand-alone” on page 82
- “PowerAI Vision installed with IBM Cloud Private” on page 82

### Types of users

#### Non-administrator users

Users other than the administrator can only see and edit resources that they created.

#### Administrator

The administrator user (admin) can see and manage all resources in PowerAI Vision regardless of who owns it. A default user name of admin with a password of `passw0rd` are created at install time. You can add, remove, or modify users by using the `kubectl` command. You should be aware of the following considerations when working with admin users:

#### Data sets

- The administrator can see and edit all data sets. That is, this user can add and delete files, create labels, assign categories, duplicate, rename, and delete the data set.
- If the administrator uploads a file to a different user's data set, it is listed as being owned by the data set owner.
- If the administrator duplicates a data set, the duplicate data set is owned by the administrator.

### Models

- The administrator can see, rename, and delete all models, including after they are deployed.
- If the administrator trains a model, the training task and the generated model is owned by the administrator.
- If the administrator deploys a model, the deployed model is owned by the administrator.

## PowerAI Vision installed as stand-alone

If you installed PowerAI Vision stand-alone, you can use the `powerai_vision_users.sh` script in the `/opt/powerai-vision/bin/` directory to create, delete, modify, and list users.

### Usage

```
powerai_vision_users.sh [command] [ --user name ] [ --password password ]
```

### Command

Specifies the action to take.

#### create

Create a user in the PowerAI Vision instance. The user argument is required for this operation. You can set the password by one of these methods:

- Specify it with the command by using the password argument.
- Store it in the environment variable, `VISION_USER_PASSWORD`.

#### delete

Delete a user from the PowerAI Vision instance. The user argument is required for this operation.

#### list

#### modify

Modifies the user's password. The user argument is required for this operation. You can set the new password by one of these methods:

- Specify it with the command by using the password argument.
- Store it in the environment variable, `VISION_USER_PASSWORD`.

### Name

The user name on which the command is to operate on.

### Password

Optionally set a user's password when creating or modifying a user.

## PowerAI Vision installed with IBM Cloud Private

PowerAI Vision uses Keycloak for user management and authentication. All users and passwords are maintained by Keycloak and stored in a Postgres database. A default user name of `admin` with a password of `passw0rd` are created at install time. You can add, remove, or modify users by using the `kubecttl` command.

1. Authenticate to the cluster, so that you can run `kubecttl` commands. For example:

- In an IBM Cloud Private 2.1.0 environment, run:
 

```
bx pr login -a https://<cluster-domain-name>:8443/ --skip-ssl-validation
```

- In an IBM Cloud Private 3.1.0 environment, run:  
`cloudctl login -a https://<cluster-domain-name>:8443/ --skip-ssl-validation`
- 2. Note your release name. In the example below, this is `aivision`.
- 3. To manage users, run the following command:  
`kubect1 run --rm -i --restart=Never usermgt --image=powerai-vision-usermgt:version -- action --user newusername --password password --release release`

The above command has the following variables:

- `action` can be one of these values: `create`, `delete`, `modify`, or `list`.
- `version` is the release number of the PowerAI Vision product. For example, `1.1.2.0`. To find the correct value, view the configmap. For example:

```
$ kubect1 get cm
NAME                               DATA   AGE
powerai-vision-v1.1.2-config      52      56d
```

The password argument is optional. You can set the password in one of these ways:

- The `--password` argument in `powerai-vision-usermgt`.
- The `--env` option for `kubect1` with the `VISION_USER_PASSWORD` environment variable. For example, add `--env="VISION_USER_PASSWORD=${MY_PASS}"` to the `kubect1 run` command.

**Example:** To create `customusername` with password `custompassw0rd1234` on release `aivision`, run:

```
$ kubect1 run --rm -i --restart=Never usermgt --image=powerai-vision-usermgt:1.1.2.0
-- create --user customusername --password custompassw0rd1234 --release aivision
Created user: customusername
```

**Example:** To list users in the PowerAI Vision 1.1.2 deployment, run:

```
$ kubect1 run --rm -i --restart=Never usermgt --image=powerai-vision-usermgt:1.1.2.0 -- list --release v111
If you don't see a command prompt, try pressing enter.
admin
testuser1
testuser2
```

#### Notes:

- If running in the non-default namespace, make sure to specify the `--namespace` option.
- The version tag on the container should match `image.releaseTag` in the `values.yaml` file.
- The argument `release` should match the release name you assigned when deploying the chart.
- There is not a typo with the spacing of the `--` before `create`. It should be `--<SPACE>create<SPACE> --user username...`. This is intentional and an artifact of how the commands are passed into the user management tool.

## Installing a new SSL certificate

PowerAI Vision ships with a self-signed certificate that is used by default, but this can be replaced with a certificate generated for PowerAI Vision for secure communications. If you wish to use your own certificate, you can create a Kubernetes secret with the new certificate, then point to the secret. An example is shown in the Kubernetes documentation.

To replace the default self-signed certificate in PowerAI Vision, follow these steps:

1. Create a Kubernetes secret with the contents of the certificate by running this command, where `/path/tls.key` and `/path/tls.crt` are the new key / certificate pair and `secret_name` is the name of the new secret:  
`kubect1 create secret tls secret_name --key /path/tls.key --cert /path/tls.crt`

For example, to create a secret named “foo-secret” in the directory “tmp”, run this command:

```
kubectl create secret tls foo-secret --key /tmp/tls.key --cert /tmp/tls.crt
```

2. Shut down PowerAI Vision:

```
sudo /opt/powerai-vision/bin/powerai-vision-stop.sh
```

3. Edit `powerai_vision_start.sh` and add the secret to the Helm upgrade command, where *secret\_name* is the name of the secret and *host\_name* is a host that is defined in the certificate for PowerAI Vision access.

**Before**

```
$BIN_DIR/helm.sh upgrade --install --set nameOverride=powerai vision ${CHART} &> $LOG_DIR/helm_install.log
```

**After**

```
$BIN_DIR/helm.sh upgrade --install --set nameOverride=powerai vision --set tls.secretName=secret_name --set ingress.hosts
```

For example, if you have a secret named “foo-secret” and the certificate defines hosts “host-a” and “host-a.domain.com”, make the following change:

**After**

```
$BIN_DIR/helm.sh upgrade --install --set nameOverride=powerai vision --set tls.secretName=foo-secret --set ingress.hosts
```

4. Start PowerAI Vision:

```
$ sudo /opt/powerai-vision/bin/powerai-vision-start.sh
```

---

## PowerAI Vision Inference Server

With a PowerAI Vision Inference server, you can quickly and easily deploy multiple trained models to a single server. These models are portable and can be used by many users and on different systems. This allows you to make trained models available to others, such as customers or collaborators.

- “Hardware requirements”
- “Software requirements”
- “Installing” on page 86
- “Deploying a trained model” on page 86
- “Inference” on page 87
- “Output” on page 87
- “Stopping a deployed model” on page 88

### Hardware requirements

#### Hardware requirements

- For deployment, the amount of memory required depends on the type of model you want to deploy.

**Note:** When initially deployed, the model uses less memory than will be required when inferences occur using the model. The memory requirements below are based on the amount of memory required when the model is used to perform inferences. To avoid errors, you should always follow these memory requirements.

- A classification model (GoogLeNet) requires about 750 MB GPU memory. For example, for a system with 16 GB memory GPUs, 19-20 image classification models can be deployed.
- An object detection model (Faster R-CNN) requires about 2 GB GPU memory. For example, for a system with 16 GB memory GPUs, 8 object detection models can be deployed.
- A object detection model (tiny YOLO V2) requires about 750 MB GPU memory. For example, for a system with 16 GB memory GPUs, 19-20 image classification models can be deployed.
- A custom classification model based on TensorFlow will take all memory available on a GPU. However, you can deploy it to a GPU that has at least 2GB memory.
- A custom object detection model based on TensorFlow will take all memory available on a GPU. However, you can deploy it to a GPU that has at least 2GB memory.

### Software requirements

#### Linux

- Red Hat Enterprise Linux (RHEL) 7.5 (little endian).
- Ubuntu 16.04 or later.

#### NVIDIA CUDA

- x86 - 9.2 or later drivers. For information, see the NVIDIA CUDA Toolkit website.
- ppc64le - 10.0 or later drivers. For information, see the NVIDIA CUDA Toolkit website.

#### Docker

- RHEL - Docker Version 1.13, or later, which is the version of Docker that is installed with RHEL 7.5.
- Ubuntu - Docker CE 18.06.01
- When running Docker, nvidia-docker2 is supported. For support of nvidia-docker2 on Docker CE, see “Using nvidia-docker2 with PowerAI Vision Inference Server” on page 100.

**Unzip** The unzip package is required on the system to deploy the zipped models.

## Installing

1. Download the install files by using one of these methods:
  - Download the product tar file from the IBM Passport Advantage website.
  - Download the product tar.gz file from Advanced Administration System (AAS). This system is also called Entitled Software Support (ESS).
2. Run the appropriate commands to install the product, depending on the platform you are installing on. There are RPM files for installation on RHEL (x86 and ppc64le) and DEB files for installation on Ubuntu (amd64 and ppc64le).

**RHEL** `rpm -i file_name.rpm`

**Ubuntu**  
`dpkg -i file_name.deb`

Load the product Docker images with the appropriate container's tar file. The file name has this format: powerai-vision-inference-*<arch>*-containers-*<release>*.tar, where *<arch>* is x86 or ppc64le, and *<release>* is the product version being installed.

`/opt/powerai-vision/dnn-deploy-service/bin/load_images.sh -f <tar_file>`

PowerAI Vision Inference Server will be installed at `/opt/powerai-vision/dnn-deploy-service`.

## Deploying a trained model

The following types of models can be deployed: object detection using Faster R-CNN (default), tiny-YOLO V2, and custom TensorFlow models. Image classification using GoogLeNet (default) and custom TensorFlow models. To deploy a model, run this command:

`/opt/powerai-vision/dnn-deploy-service/bin/deploy_zip_model.sh`

### Notes:

- The first time you run this command, you are prompted to accept the license agreement.
- On a RHEL system with SELinux enabled (default), the loaded model files must have an appropriate SELinux context to be loaded into a container. To ensure a model has the proper context, run:

`sudo chcon -t svirt_sandbox_file_t <model_path>`

### Usage:

`./deploy_zip_model.sh -m <model-name> -p <port> -g <gpu> zipped_model_file`

#### **model-name**

The docker container name for the deployed model.

**port** The port to deploy the model to.

**gpu** Optional: The GPU to deploy the model to. If specified as -1, the model will be deployed to CPU.

#### **zipped\_model\_file**

The full path and file name of the trained model that was exported from PowerAI Vision. It can be an image classification model or an object detection model, but must be in zip format.

### Examples:

`/opt/powerai-vision/dnn-deploy-service/bin/deploy_zip_model.sh --model dog --port 6001 --gpu 1 ./dog_classification.zip`  
`/opt/powerai-vision/dnn-deploy-service/bin/deploy_zip_model.sh --m face -p 6002 /home/user/mydata/face.zip`

## Inference

Inference can be done by using the deployed model with a local file or an image URL.

### Example 1 - Classification:

```
curl -F "imagefile=@/home/testdata/cocker-spaniel-dogs-puppies-1.jpg" http://localhost:6001/inference
```

### Example 2 - Object detection:

```
curl -G -d "imageurl=https://assets.imgix.net/examples/couple.jpg" http://localhost:6002/inference
```

### Example 3 – Object detection of tiny YOLO model with confidence threshold:

```
curl -F "imagefile=@/home/testdata/Chihuahua.jpeg" -F "confthre=0.8" http://localhost:6001/inference
```

**Note:** Confidence threshold works for FRCNN and tiny YOLO object detection models and google-net image classification models. The confidence threshold is a value in the range 0.0 - 1.0, treated as a percentage. Only results with a confidence greater than the specified threshold are returned. The smaller confidence threshold you specify, the more results are returned. If you specify 0, many, many results will be returned because there is no filter based on the confidence level of the model. The default confidence threshold is 0.5.

## Output

The PowerAI Vision Inference Server can deploy both image classification models and object detection models.

### Image classification model

A successful classification will report something similar to the following:

#### Example 1 output - success

```
{"classified": {"Cocker Spaniel": 0.93}, "result": "success"}
```

The image has been classified as a Cocker Spaniel with a confidence of .93.

#### Example 1 output - fail

```
{"result": "fail"}
```

The image could not be classified. This might happen if the image could not be loaded, for example.

### Object detection model

A successful detection will report something similar to the following:

#### Example 2 output - success

```
{"classified": [{"confidence": 0.94, "ymax": 335, "label": "face", "xmax": 576, "xmin": 424, "ymin": 160, "attr": []}], "result": "success"}
```

The faces in the image are located at the specified coordinates. The confidence of each label is given.

#### Example 2 output - success

```
{"classified": [], "result": "success"}
```

Object detection was carried out successfully, but there was nothing to be labeled that has confidence above the threshold.

#### Example 2 output - fail

```
{"result": "fail"}
```

Objects could not be detected. This might happen if the image could not be loaded, for example.

## Stopping a deployed model

To stop the deployed model, run the following command. When you stop the deployed model, the GPU memory is made available.

```
docker stop <model-name>; docker rm <model-name>
```

### Example 1:

```
docker stop dog; docker rm dog
```

### Example 2:

```
docker stop face; docker rm face
```

---

## Inference on embedded edge devices

Using edge computing for your inference models helps save processing time by removing latency issues, ensures security, and also decreases bandwidth usage. This topic describes how to use PowerAI Vision with embedded edge devices.

- A full end to end use case is available if you use the DeepRed FPGA by V3 Technology that takes camera input, analyzes the video, and outputs the video with bounding-boxes on an HDMI attached device. See the section “Inference on DeepRED” for details.
- If you want to create your own solution or have a different FPGA board, then use the information in this section: “Inference with a custom solution” on page 89.

## Inference on DeepRED

DeepRED is an embedded artificial intelligence development system that supports PowerAI Vision. It lets you quickly deploy the trained model for testing and production.

Generate an IP core for use with DeepRED:

1. Perform customization. For DEEPRED this is not optional. If they want to use DEEPRED, they need to change the ZC706 to DEEPRED in the configmap (both instances of it). They should not add custom DSP\_NUM, etc.
2. Perform optional customization.
  - a. On the PowerAI Vision host operating system, run the appropriate command.
    - For a standard install:  

```
$ /opt/powerai-vision/bin/kubect1.sh edit configmap powerai-vision-config
```
    - For PowerAI Vision installed on IBM Cloud Private:  

```
$ kubect1 edit configmap powerai-vision-release_name-config
```
  - b. Find the row beginning EMB\_COD\_IMAGE in the configuration file and replace ZC706 with DEEPRED:  
"EMB\_COD\_IMAGE": ["**DEEPRED,DEEPRED**,powerai-vision-dnn-edge:1.1.2.0"],
  - c. Save and exit.
3. Restart PowerAI Vision by running the appropriate command. The deleted pods will automatically restart.
  - For a standard install:  

```
$ /opt/powerai-vision/bin/kubect1.sh delete pod -l app=powerai-vision
```
  - For PowerAI Vision installed on IBM Cloud Private:  

```
$ kubect1 delete pod -l app=powerai-vision-release_name
```
4. Train your model.
  - On the Train data set page, for **Type of training**, select **Object detection**.

- Under **Advanced options**, choose **Optimized for speed**
5. Copy the IP core file for compilation. The generated FPGA IP core is named *UUID-ipcore.zip*, where *UUID* is the UUID of the trained model. It is stored in the following location:
    - For a standard install: `/opt/powerai-vision/volume/data/trained-models`.
    - For PowerAI Vision installed on IBM Cloud Private, it is stored in your Persistent Volume under `<PATH_TO_VOLUME>/data/trained-models`.

## Inference with a custom solution

Using a custom solution requires appropriate hardware and software, as well as FPGA development skills. You must be able to:

- Take an existing IP core and use Vivado to merge it into a custom solution. Refer to the [PIE DNN Accelerator IP Integration Guide.pdf](#) for instructions to integrate the generated DNN IP core into your project.
- Set up and use Vivado, Petalinux, and other software.

### Environment requirements

- A chip set that can provide enough BRAM, such as Xilinx 7035 or later
- A board with PL side (not just PS side) DRAM so that it can provide sufficient bandwidth between the FPGA and DRAM.

Follow these steps to generate an IP core for use with a custom solution. The examples included are for a ZC706 card:

1. Perform optional customization.

By default, PowerAI Vision is configured to use the following resources on a ZC706 card. However, you can customize these values.

```
DSP_NUM=700
RAM18E_NUM=800
DDR_BANDWIDTH=80000.0
DDR_DATA_WIDTH=512
FPGA_TYPE=xc7z045ffg900-2
```

- a. On the PowerAI Vision host operating system, run the appropriate command.

- For a standard install:
 

```
$ /opt/powerai-vision/bin/kubect1.sh edit configmap powerai-vision-config
```
- For PowerAI Vision installed on IBM Cloud Private:
 

```
$ kubect1 edit configmap powerai-vision-release_name-config
```

- b. Find the row beginning `EMB_COD_IMAGE` in the configuration file and input your custom values.

For example, for a ZC706 card, replace **ZC706** with the appropriate values for your card: `"EMB_COD_IMAGE": ["ZC706,ZC706,powerai-vision-dnn-edge:1.1.2.0"]`, as shown here:

```
"EMB_COD_IMAGE": ["ZC706,DSP_NUM=700:RAM18E_NUM=800:DDR_BANDWIDTH=80000.0:
DDR_DATA_WIDTH=512:FPGA_TYPE=xcvu9p12fsgd2104e,powerai-vision-dnn-edge:1.1.2.0"],
```

- c. Save and exit.

2. Restart PowerAI Vision by running the appropriate command. The deleted pods will automatically restart.

- For a standard install:
 

```
$ /opt/powerai-vision/bin/kubect1.sh delete pod -l app=powerai-vision
```
- For PowerAI Vision installed on IBM Cloud Private:
 

```
$ kubect1 delete pod -l app=powerai-vision-release_name
```

3. Train your model.

- On the Train data set page, for **Type of training**, select **Object detection**.
  - Under **Advanced options**, choose **Optimized for speed**
4. Copy the IP core file for compilation. The generated FPGA IP core is named *UUID*-ipcore.zip, where *UUID* is the UUID of the trained model. It is stored in the following location:
- For a standard install: /opt/powerai-vision/volume/data/trained-models.
  - For PowerAI Vision installed on IBM Cloud Private, it is stored in your Persistent Volume under `<PATH_TO_VOLUME>/data/trained-models`.

---

## Troubleshooting and contacting support

To isolate and resolve problems with your IBM products, you can use the following troubleshooting and support information. This information contains instructions for using the problem-determination resources that are provided with your IBM products, including PowerAI Vision.

---

### Troubleshooting known issues - PowerAI Vision standard install

Following are some problems you might encounter when using PowerAI Vision, along with steps to fix them.

- “The PowerAI Vision GUI does not work”
- “Resource pages are not being populated in the user interface”
- “Unexpected / old pages displayed when accessing the user interface” on page 92
- “PowerAI Vision does not play video” on page 92
- “Auto detection video does not play in Firefox browser” on page 92
- “Out of space error from load\_images.sh” on page 92
- “GPUs are not available for training or inference” on page 93
- “I forgot my user name or password” on page 93
- “PowerAI Vision cannot train a model” on page 94
- “Training or deployment hangs - Kubernetes pod cleanup” on page 95
- “Model training and inference fails” on page 95
- “Auto labeling of a data set returns "Auto Label Error"” on page 95
- “PowerAI Vision does not start” on page 96
- “PowerAI Vision fails to start - Kubernetes connection issue” on page 96
- “PowerAI Vision startup hangs - helm issue” on page 97
- “Helm status errors when starting PowerAI Vision” on page 98
- “Uploading a large file fails” on page 99
- “Some PowerAI Vision functions don't work ” on page 99

#### The PowerAI Vision GUI does not work

##### Problem

You cannot label objects, view training charts, or create categories.

##### Solution

Verify that you are using a supported web browser. The following web browsers are supported:

- Google Chrome Version 60, or later
- Firefox Quantum 59.0, or later

#### Resource pages are not being populated in the user interface

##### Problem

Resource pages, such as data sets and models, are not being populated. Notifications indicate that there is an error obtaining the resource. For example, "Error obtaining data sets."

##### Solution

Check the status of the powerai-vision-portal pod. This pod provides the data to the user interface, and until it is ready ( 1/1) with a status of Running, these errors will occur. See “Checking Kubernetes node status” on page 30 for instructions.

If the application is restarting, there is an expected delay before all services are available and fully functioning. Otherwise, this may indicate an unexpected termination (error) of the `powerai-vision-portal` pod. If that happens, follow these instructions: “Gather PowerAI Vision logs and contact support” on page 100.

## Unexpected / old pages displayed when accessing the user interface

### Problem

After updating, reinstalling, or restarting PowerAI Vision, the browser presents pages that are from the previous version or are stale.

### Solution

This problem is typically caused by the browser using a cached version of the page. To solve the problem, try one of these methods:

- Use a Firefox Private Window to access the user interface.
- Use a Chrome Incognito Window to access the user interface.
- Bypass the browser cache:
  - In most Windows and Linux browsers: Hold down `Ctrl` and press `F5`.
  - In Chrome and Firefox for Mac: Hold down `⌘` `Cmd` and `⇧` `Shift` and press `R`.

## PowerAI Vision does not play video

### Problem

You cannot upload a video, or after the video is uploaded the video does not play.

### Solution

Verify that your video is a supported type:

- Ogg Vorbis (.ogg)
- VP8 or VP9 (.webm)
- H.264 encoded videos with MP4 format (.mp4)

If your video is not in a supported format, transcode your video by using a conversion utility. Such utilities are available under various free and paid licenses.

## Auto detection video does not play in Firefox browser

### Problem

The Firefox browser reports “The media playback was aborted due to a corruption problem or because the media used features your browser did not support”. This happens in versions of the Firefox browser that do not support YUV444 chroma subsampling, which prevents the video from being played successfully.

### Solution

Use a version of Firefox that supports YUV444 chroma subsampling or use a different browser (such as Chrome) that does support it.

## Out of space error from `load_images.sh`

### Problem

When installing the product, the `load_images.sh` script is used to load the PowerAI Vision Docker images. The script might terminate with errors, the most frequent issue being insufficient disk space for loading the Docker images.

For example, the `/var/lib/docker` file system can run out of space, resulting in a message indicating that an image was not fully loaded. The following output shows that the Docker image `powerai-vision-dnn` was not able to be fully loaded because of insufficient file system space:

```

root@kottos-vm1:~# df --output -BG "/var/lib/docker/"
Filesystem      Type Inodes IUsed IFree IUse% 1G-blocks    Used Avail Use% File      Mounted on
/dev/vda2      ext4 8208384 595697 7612687    8%      124G     81G   37G   70% /var/lib/docker/ /
root@kottos-vm1:~#

*****
892d6f64ce41: Loading layer [====>] 21.26MB/21.26MB
785af1d0c551: Loading layer [====>] 1.692MB/1.692MB
dc102f4a3565: Loading layer [====>] 747.9MB/747.9MB
aac4b03de02a: Loading layer [====>] 344.1MB/344.1MB
d0ea7f5f6aab: Loading layer [====>] 2.689MB/2.689MB
62d3d10c6cc2: Loading layer [====>] 9.291MB/9.291MB
240c4d86e5c7: Loading layer [====>] 778MB/778MB
889cd0648a86: Loading layer [====>] 2.775MB/2.775MB
56bbb2f20054: Loading layer [====>] 3.584kB/3.584kB
3d3c7acb72e2: Loading layer [====>] 2.117GB/3.242GB
Error processing tar file(exit status 1): write /usr/bin/grops: no space left on device

[ FAIL ] Some images failed to load
[ FAIL ] Failure info:
          Loading the PowerAI Vision docker images...
root@kottos-vm1:~#

```

This situation can also be noted in the output from `/opt/powerai-vision/bin/kubect1 get pods`. This command is described in “Checking the application and environment” on page 29, which shows images that could not be loaded with a status of `ErrImagePull` or `ImagePullBackOff`.

## Solution

The file system space for `/var/lib/docker` needs to be increased, even if the file system is not completely full. There might still be space in the file system where `/var/lib/docker` is located, but insufficient space for the PowerAI Vision Docker images. There are operating system mechanisms to do this, including moving or mounting `/var/lib/docker` to a file system partition with more space.

After the error situation has been addressed by increasing or cleaning up disk space on the `/var/lib/docker/` file system, re-run the `load_images.sh` script to continue loading the images. No clean up of the previous run of `load_images.sh` is required.

## I forgot my user name or password

### Problem

You forgot your user name or password and cannot log in to the PowerAI Vision GUI.

### Solution

PowerAI Vision uses an internally managed users account database. To change your user name or password, see “Logging in to PowerAI Vision” on page 43.

## GPUs are not available for training or inference

### Problem

If PowerAI Vision cannot perform training or inference operations, check the following:

- Verify that the `nvidia smi` output shows all relevant information about the GPU devices. For example, the following output shows Unknown error messages indicating that the GPUs are not in the proper state:

Mon Dec 3 15:43:07 2018

```

+-----+
| NVIDIA-SMI 410.72      Driver Version: 410.72      CUDA Version: 10.0      |
+-----+-----+-----+
| GPU  Name      Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+-----+

```

0	Tesla V100-SXM2...	Off	00000004:04:00.0	Off	0
N/A	31C	P0	49W / 300W	Unknown Error	0%
					Default

...

- Verify that the **nvidia-persistenced** service is enabled and running (active) by using the command `sudo systemctl status nvidia-persistenced`:  

```
# systemctl status nvidia-persistenced
* nvidia-persistenced.service - NVIDIA Persistence Daemon
   Loaded: loaded (/etc/systemd/system/nvidia-persistenced.service; enabled; vendor preset: disabled)
   Active: active (running) since Tue 2018-11-13 08:41:22 CST; 2 weeks 6 days ago
```

## Solution

If the GPU status indicates errors and the **nvidia-persistenced** service is not enabled and active, enable and start the service:

1. Enable the service:  

```
sudo systemctl enable nvidia-persistenced
```
2. Start the service:  

```
sudo systemctl start nvidia-persistenced
```

## PowerAI Vision cannot train a model

### Problem

The model training process might fail if your system does not have enough GPU resources.

### Solution

- If you are training a data set for image classification, verify that at least two image categories are defined, and that each category has a minimum of five images.
- If you are training a data set for object detection, verify that at least one object label is used. You must also verify that each object is labeled in a minimum of five images.
- Ensure that enough GPUs are available. PowerAI Vision assigns one GPU to each active training job or deployed deep learning API. For example, if a system has four GPUs and you have two deployed web APIs, there are two GPUs available for active training jobs. If a training job appears to be hanging, it might be waiting for another training job to complete, or there might not be a GPU available to run it.

To determine how many GPUs are available on the system, run the **sudo /opt/powerai-vision/bin/kubectrl.sh describe nodes** script and review the **nVidiaGPU Limits** column.

The following is an example of the output from **sudo /opt/powerai-vision/bin/kubectrl.sh describe nodes** that shows two GPUs currently in use:

Name:	127.0.0.1
Roles:	<none>
Labels:	beta.kubernetes.io/arch=ppc64le beta.kubernetes.io/os=linux gpu/nvidia=TeslaV100-SXM2-16GB kubernetes.io/hostname=127.0.0.1
Annotations:	node.alpha.kubernetes.io/ttl=0 volumes.kubernetes.io/controller-managed-attach-detach=true...
Allocated resources:	(Total limits may be over 100 percent, i.e., overcommitted.) CPU Requests CPU Limits Memory Requests Memory Limits <b>nVidiaGPU Limits</b>
	0 (0%) 0 (0%) 0 (0%) 0 (0%) <b>2 (50%)</b>
Events:	<none>

If all the systems GPUs are in use, you can either delete a deployed web API (making the API unavailable for inference) or you can stop a training model that is running.

- To delete a deployed model, click **Deployed Models**. Next, select the model that you want to delete and click **Delete**. The trained model is not deleted from PowerAI Vision. You can redeploy the model later when more GPUs are available.
- To stop a training model that is running, click **Models**. Next, select the model that has a status of **Training in Progress** and click **Stop Training**.

## Training or deployment hangs - Kubernetes pod cleanup

### Problem

You submit a job for training or deployment, but it never completes. When doing training or deployments, sometimes some pods that are running previous jobs get out of sync with the vision-service MongoDB and they hang forever instead of getting terminated within a minute or so. In turn, they hold GPUs so no new training or deployment jobs can complete. They will be in the Scheduled state forever.

To verify that this is the problem, run `kubectl get pods` and review the output. The last column shows the age of the pod. If it is older than a few minutes, use the information in “Solution” to solve the problem.

### Example:

```
kubectl get pods
powerai-vision-infer-ic-06767722-47df-4ec1-bd58-91299255f6hxxzk 1/1 Running 0 22m
powerai-vision-infer-ic-35884119-87b6-4d1e-a263-8fb645f0addqd2z 1/1 Running 0 22m
powerai-vision-infer-ic-7e03c8f3-908a-4b52-b5d1-6d2befec69ggqw5 1/1 Running 0 5h
powerai-vision-infer-od-clc16515-5955-4ec2-8f23-bd21d394128b6k4 1/1 Running 0 3h
```

### Solution

Follow these steps to manually delete the deployments that are hanging.

1. Determine the running deployments and look for those that have been running longer than a few minutes:  
`kubectl get deployments`
2. Delete the deployments that were identified as hanging in the previous step.  
`kubectl delete deployment deployment_id`
3. You can now try the training or deploy again, assuming there are available GPUs.

**Note:** When a deployment is manually deleted, vision-service might try to recreate it when it is restarted. The only way to force Kubernetes to permanently delete it is to remove the failing model from PowerAI Vision.

## Model training and inference fails

### Problem

The NVIDIA GPU device is not accessible by the PowerAI Vision Docker containers. To confirm this, run `kubectl logs -f _powerai-vision-portal-ID_` and then check `pod_powerai-vision-portal-ID_powerai-vision-portal.log` for an error indicating error == cudaSuccess (30 vs. 0):  
 F0731 20:34:05.334903 35 common.cpp:159] Check failed: error == cudaSuccess (30 vs. 0) unknown error  
 \*\*\* Check failure stack trace: \*\*\*  
 /opt/py-faster-rcnn/FRCNN/bin/train\_frcnn.sh: line 24: 35 Aborted (core dumped) \_train\_frcnn.

### Solution

Use sudo to alter SELINUX permissions for all of the NVIDIA devices so they are accessible via the PowerAI Vision Docker containers.

```
sudo chcon -t container_file_t /dev/nvidia*
```

## Auto labeling of a data set returns "Auto Label Error"

### Problem

Auto labeling cannot be performed on a data set that does not have unlabeled images, unless some of the images were previously labeled by the auto label function.

## Solution

Ensure that the **Objects** section of the data set side bar shows there are objects that are “Unlabeled”. If there are none, that is, if “Unlabeled (0)” is displayed in the side bar, add new images that are unlabeled or remove labels from some images, then run auto label again.

## PowerAI Vision does not start

### Problem

When you enter the URL for PowerAI Vision from a supported web browser, nothing is displayed. You see a 404 error or Connection Refused message.

### Solution

Complete the following steps to solve this problem:

1. Verify that IP version 4 (IPv4) port forwarding is enabled by running the **/sbin/sysctl net.ipv4.conf.all.forwarding** command and verifying that the value for **net.ipv4.conf.all.forwarding** is set to 1.

If IPv4 port forwarding is not enabled, run the **/sbin/sysctl -w net.ipv4.conf.all.forwarding=1** command. For more information about port forwarding with Docker, see UCP requires IPv4 IP Forwarding in the Docker success center.

2. If IPv4 port forwarding is enabled and the docker0 interface is a member of the trusted zone, check the Helm chart status by running this script:

```
sudo /opt/powerai-vision/bin/helm.sh status vision
```

In the script output, verify that the PowerAI Vision components are available by locating the Deployment section and identifying that the AVAILABLE column has a value of 1 for each component. The following is an example of the output from the **helm.sh status vision** script that shows all components are available:

```
RESOURCES:
==> v1beta1/Deployment
NAME                                DESIRED  CURRENT  UP-TO-DATE  AVAILABLE  AGE
powerai-vision-mongodb              1         1         1             1         4d
powerai-vision-portal                1         1         1             1         4d
powerai-vision-postgres              1         1         1             1         4d
powerai-vision-taskanalyzer          1         1         1             1         4d
powerai-vision-ui                    1         1         1             1         4d
powerai-vision-video-nginx           1         1         1             1         4d
powerai-vision-video-portal          1         1         1             1         4d
powerai-vision-video-rabmq           1         1         1             1         4d
powerai-vision-video-redis           1         1         1             1         4d
powerai-vision-video-test-nginx      1         1         1             1         4d
powerai-vision-video-test-portal     1         1         1             1         4d
powerai-vision-video-test-rabmq      1         1         1             1         4d
powerai-vision-video-test-redis      1         1         1             1         4d
```

If you recently started PowerAI Vision and some components are not available, wait a few minutes for these components to become available. If any components remain unavailable, gather the logs and contact IBM Support, as described in this topic: “Gather PowerAI Vision logs and contact support” on page 100.

3. If the docker0 interface is a member of a trusted zone and all PowerAI Vision components are available, verify that the firewall is configured to allow communication through port 443 (used to connect to PowerAI Vision) by running this command:

```
sudo firewall-cmd --permanent --zone=public --add-port=443/tcp
```

## PowerAI Vision fails to start - Kubernetes connection issue

### Problem

If the host system does not have a default route defined in the networking configuration, the Kubernetes cluster will fail to start with connection issues. For example:

```
$ sudo /opt/powerai-vision/bin/powerai_vision_start.sh
INFO: Setting up GPU...
[...]
Checking kubernetes cluster status...
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #1:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #2:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #3:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #4:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #5:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #6:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #7:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #8:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #9:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #10:
The connection to the server 127.0.0.1:8080 was refused - did you specify the right host or port?
INFO: Probing cluster status #11:
ERROR: Retry timeout. Error in starting kubernetes cluster, please check /opt/powerai-vision/log/kubernetes for logs.
```

## Solution

Define a default route in the networking configuration. For instructions to do this on Red Hat Enterprise Linux (RHEL), refer to 2.2.4 Static Routes and the Default Gateway in the Red Hat Customer Portal.

## PowerAI Vision startup hangs - helm issue

### Problem

PowerAI Vision startup hangs with the message "Unable to start helm within 30 seconds - trying again." For example:

```
root> sudo /opt/powerai-vision/bin/powerai_vision_start.sh
Checking ports usage...
Checking ports completed, no conflict port usage detected.
[ INFO ] Setting up the GPU...
Init cuda devices...
Devices init completed!
Persistence mode is already Enabled for GPU 00000004:04:00.0.
Persistence mode is already Enabled for GPU 00000004:05:00.0.
Persistence mode is already Enabled for GPU 00000035:03:00.0.
Persistence mode is already Enabled for GPU 00000035:04:00.0.
All done.
[ INFO ] Starting kubernetes...
Checking kubernetes cluster status...
Probing cluster status #1: NotReady
Probing cluster status #2: NotReady
Probing cluster status #3: NotReady
Probing cluster status #4: Ready
Bootting up ingress controller...
Initializing helm...
[ WARN ] Unable to start helm within 30 seconds - trying again. If this continues, contact support.
[ WARN ] Unable to start helm within 30 seconds - trying again. If this continues, contact support.
[ WARN ] Unable to start helm within 30 seconds - trying again. If this continues, contact support.
[ WARN ] Unable to start helm within 30 seconds - trying again. If this continues, contact support.
```

## Solution

To solve this problem, you must follow these steps exactly as written:

1. Cancel PowerAI Vision startup by pressing ctrl+c.
2. Stop PowerAI Vision by running this command:

```
sudo /opt/powerai-vision/bin/powerai_vision_stop.sh
```

3. Modify the RHEL settings as follows:

```
sudo nmcli connection modify docker0 connection.zone trusted
sudo systemctl stop NetworkManager.service
sudo firewall-cmd --permanent --zone=trusted --change-interface=docker0
sudo systemctl start NetworkManager.service
sudo nmcli connection modify docker0 connection.zone trusted
sudo systemctl restart docker.service
```

4. Start PowerAI Vision again:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

If the above commands do not fix the startup issue, check for a cgroup leak that can impact Docker. A Kubernetes/Docker issue can cause this situation, and after fixing the firewall issue the start up can still fail if there was cgroup leakage.

One symptom of this situation is that the `df` command is slow to respond. To check for excessive cgroup mounts, run the `mount` command:

```
$ mount | grep cgroup | wc -l
```

If the cgroup count is in thousands, reboot the system to clear up the cgroups.

## Helm status errors when starting PowerAI Vision

### Problem

There is an issue in some RHEL releases that causes the startup of PowerAI Vision to fail after restarting the host system. When this is the problem, the system tries to initialize Helm at 30 second intervals but never succeeds. Therefore, the startup never succeeds. You can verify this status by running the Helm status vision command:

```
# /opt/powerai-vision/bin/helm status vision
```

Result:

```
Error: getting deployed release "vision": Get https://10.10.0.1:443/api/v1/namespaces/kube-system/configmaps[...]: dial tcp 10.1
```

### Solution

To solve this problem, you must follow these steps exactly as written:

1. Cancel PowerAI Vision startup by pressing `ctrl+c`.

2. Stop PowerAI Vision by running this command:

```
sudo /opt/powerai-vision/bin/powerai_vision_stop.sh
```

3. Modify the RHEL settings as follows:

```
sudo nmcli connection modify docker0 connection.zone trusted
sudo systemctl stop NetworkManager.service
sudo firewall-cmd --permanent --zone=trusted --change-interface=docker0
sudo systemctl start NetworkManager.service
sudo nmcli connection modify docker0 connection.zone trusted
sudo systemctl restart docker.service
```

4. Start PowerAI Vision again:

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

If the above commands do not fix the startup issue, check for a cgroup leak that can impact Docker. A Kubernetes/Docker issue can cause this situation, and after fixing the firewall issue the start up can still fail if there was cgroup leakage.

One symptom of this situation is that the `df` command is slow to respond. To check for excessive cgroup mounts, run the `mount` command:

```
$ mount | grep cgroup | wc -l
```

If the cgroup count is in thousands, reboot the system to clear up the cgroups.

## Uploading a large file fails

When uploading files into a data set, there is a 2GB size limit per upload session. This limit applies to a single .zip file or a set of files. When you upload a large file that is under 2 GB, you might see the upload start (showing a progress bar) but then you get an error message in the user interface. This error happens due to a Nginx timeout, where the file upload is taking longer than the defined 5 minute Nginx timeout.

Despite the notification error, the large file has been uploaded. Refreshing the page will show the uploaded files in the data set.

## Some PowerAI Vision functions don't work

### Problem

PowerAI Vision seems to start correctly, but some functions, like automatic labeling or automatic frame capture, do not function.

To verify that this is the problem, run `/opt/powerai-vision/bin/kubect1.sh get pods` and verify that one or more pods are in state `CrashLoopBackOff`. For example:

```
kubect1 get pods
```

NAME	READY	STATUS	RESTARTS	AGE
...				
powerai-vision-video-rabmq-5d5d786f9f-7jfk9	0/1	CrashLoopBackOff	2	54s

### Solution

PowerAI Vision requires IPv6. Enable IPv6 on the system.

---

## Troubleshooting known issues - PowerAI Vision Inference Server

Following are some problems you might encounter when using PowerAI Vision, along with steps to fix them.

- “On a RHEL system, the container terminates soon after deployment”
- “Problems installing an rpm on a RHEL system with Docker CE” on page 100
- “Using nvidia-docker2 with PowerAI Vision Inference Server” on page 100

### On a RHEL system, the container terminates soon after deployment

#### Problem

Shortly after model deployment, the container terminates.

For example, deploy a model:

```
sudo /opt/powerai-vision/dnn-deploy-service/bin/deploy_zip_model.sh --model dog --port 6001 --gpu 1 dog_classifi
```

Then after about one minute, run the following command. It returns empty, which means that the container terminated:

```
sudo docker container ps | grep dog
```

#### Solution

To solve this problem, follow these steps:

1. Run these commands, where *path/file\_name* is the zip file of the model and the directory of the collected Nvidia driver:

```
sudo chcon -t container_file_t path/file_name.zip
```

```
sudo chcon -t container_file_t /opt/powerai-vision/dnn-deploy-service/bin/nvidia-driver/410.48
```

2. Try the deployment again. For example:

```
sudo docker rm dog
sudo /opt/powerai-vision/dnn-deploy-service/bin/deploy_zip_model.sh --model dog --port 6001 --gpu 1 dog_classifi
```

## Problems installing an rpm on a RHEL system with Docker CE

### Problem

When installing an rpm on a RHEL system with Docker CE, you see this error: Error: Failed dependencies: docker is needed by <file\_name.rpm>. For example:

### Solution

To install an rpm on a system with Docker CE instead of Docker, force install the rpm by the following command

```
rpm --nodeps -i file_name.rpm
```

## Using nvidia-docker2 with PowerAI Vision Inference Server

### Problem

Deploying a model with PowerAI Vision Inference Server fails when using the nvidia-docker2 runtime on a Docker CE based system.

By default, installing nvidia-docker2 does not set it as the default runtime.

### Solution

These steps apply to a Docker CE based system only.

1. Ensure that you are running Docker CE by running the appropriate command:

- RHEL system: `rpm -qa | grep 'docker-ce'`
- Ubuntu system: `dpkg --get-selections | grep 'docker-ce'`

If you got non empty returns, then the system is using Docker CE. Example:

```
# dpkg --get-selections | grep 'docker-ce'
ii  docker-ce                5:18.09.0~3-0~ubuntu-bionic      amd64      Docker: the open-source container management system
ii  docker-ce-cli             5:18.09.0~3-0~ubuntu-bionic      amd64      Docker CLI: the open-source command line interface
```

2. Modify the `deploy_zip_model.sh` script to set nvidia-docker2 to be the default runtime. By default, this file is in `/opt/powerai-vision/dnn-deploy-service/bin`. Look for this section in the script, typically near line 159, and add the correct runtime:

#### Before:

```
else
  CMD="docker run -d \
    -e BASE \
    -v ${ZIPPED_MODEL_PATH}:/config/dropins/model.zip \
    --name ${CONTAINER_NAME} \
    -p ${PORT_TO}:${PORT} "
fi
```

#### After:

```
else
  CMD="docker run -d \
    -e BASE \
    -v ${ZIPPED_MODEL_PATH}:/config/dropins/model.zip \
    --name ${CONTAINER_NAME} \
    --runtime=nvidia \
    -p ${PORT_TO}:${PORT} "
fi
```

## Gather PowerAI Vision logs and contact support

Sometimes you cannot solve a problem by troubleshooting the symptoms. In such cases, you must collect diagnostic data and contact support.

Collecting and inspecting data before you open a problem management record (PMR) can help you to answer the following questions:

- Do the symptoms match any known problems? If so, has a fix or workaround been published?
- Can the problem be identified and resolved without a code fix?
- When does the problem occur?

To gather logs for support, follow these steps:

1. Collect logs from the PowerAI Vision application.

- **Standalone installation:** Run the **sudo /opt/powerai-vision/bin/collect\_logs.sh** script. The directory where the log file is saved is listed in the **INFO: FFDC Collected** section, as shown in the following example:

```
INFO: Collecting PowerAI Vision Application Logs...
INFO: Collecting PowerAI Infrastructure Logs...
INFO: Collecting configuration information...
INFO: Collecting System Details...
INFO: Collecting Platform Logs...
INFO: FFDC Collected below:
-rw-r--r--. 1 root root 95477342 May 22 18:15 /var/log/powerai-vision/powerai-vision.logs.18_15_11_May_22_2018.tgz
```

- **IBM Cloud Private installation:**

- Enable the **kubectl** command. For instructions, see this topic in the IBM Cloud Private Knowledge Center: [Accessing your IBM® Cloud private cluster by using the kubectl CLI](#).
- Collect logs from the PowerAI Vision pods using the `<release_name>` specified when installing/deploying:

```
kubectl get pods > /tmp/kubectl_pod_status.txt
for i in $(kubectl get pods -o name | grep <release_name>); do
  kubectl logs $i > /tmp/kubectl_logs_$i.txt
  kubectl describe pods $i > /tmp/kubectl_describe_pods_$i.txt
done
kubectl describe deploy > /tmp/kubectl_deploy.txt
kubectl describe configmap > /tmp/kubectl_cm.txt
```

The log files to provide are generated here: **/tmp/kubectl\_logs\***.

2. Optionally, you can obtain the logs for a single pod of the application.

- Use the **kubectl get pods** command to view the running pods for the application. See “**kubectl.sh get pods**” on page 31. For example:

```
[root@aprilmin8 bin]# /opt/powerai-vision/bin/kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
powerai-vision-infer-ic-f873e5ff-821b-472a-bec0-32b1cdc4b59ns7j  1/1     Running   0           14d
powerai-vision-infer-od-fae2aa93-e394-4a04-bec7-f072b81f2678b2m  1/1     Running   0           14d
powerai-vision-keycloak-987f9698d-kdvp2                          1/1     Running   1           14d
powerai-vision-mongodb-686ffbd4b9-cq2mx                         1/1     Running   0           14d
powerai-vision-portal-768c4ffdc7-9ppnt                          1/1     Running   0           14d
powerai-vision-postgres-7c88467b4c-szw8m                       1/1     Running   0           14d
powerai-vision-taskanalzy-788bdcc5cc-dc5cp                     1/1     Running   0           14d
powerai-vision-ui-844bb5f7f9-t55bg                             1/1     Running   0           14d
powerai-vision-video-nginx-845544c55b-w8ks6                    1/1     Running   0           14d
powerai-vision-video-portal-7499b577c7-csvs8                   1/1     Running   0           14d
powerai-vision-video-rabmq-85b486bc9d-nqdmn                    1/1     Running   0           14d
powerai-vision-video-redis-67d9766fb6-gw7hd                    1/1     Running   0           14d
powerai-vision-video-test-nginx-b5cb6b759-fkwp5                1/1     Running   0           14d
powerai-vision-video-test-portal-756f7f9b99-mtwwj               1/1     Running   0           14d
powerai-vision-video-test-rabmq-578c775d4-chbtd                1/1     Running   0           14d
powerai-vision-video-test-redis-779f7545b4-dffws               1/1     Running   0           14d
```

- Run **kubectl.sh logs <pod-name> > <outputfile>**, where `<pod-name>` is obtained from the **kubectl.sh get pods** command. For example, using the above output, to collect logs in the file **vision-portal.log** run the command: **\$ kubectl.sh logs powerai-vision-portal-9455fc5cc-k9qpq > vision-portal.log**

3. Submit the problem to IBM Support in one of the following ways:

- Online through the IBM Support Portal: <http://www.ibm.com/software/support/>: You can open, update, and view all of your service requests from the Service Request portlet on the Service Request web page.
- By phone: For the phone number to call in your region, see the Directory of worldwide contacts web page: <http://www.ibm.com/planetwide/>.

---

## Getting fixes from Fix Central

You can use Fix Central to find the fixes that are recommended by IBM Support for various products, including PowerAI Vision. With Fix Central, you can search, select, order, and download fixes for your system with a choice of delivery options. A PowerAI Vision product fix might be available to resolve your problem.

To find and install fixes:

1. Obtain the tools that are required to get the fix. If it is not installed, obtain your product update installer. You can download the installer from Fix Central: <http://www.ibm.com/support/fixcentral>. This site provides download, installation, and configuration instructions for the update installer.

**Note:** For more information about how to obtain software fixes, from the Fix Central page, click **Getting started with Fix Central**, then click the Software tab.

2. Under **Find product**, type “PowerAI Vision” in the **Product selector** field.
3. Select PowerAI Vision. For **Installed version**, select **All**. For **Platform**, select the appropriate platform or select **All**, then click **Continue**.
4. Identify and select the fix that is required, then click **Continue**.
5. Download the fix. When you download the file, ensure that the name of the maintenance file is not changed, either intentionally or by the web browser or download utility.
6. Stop PowerAI Vision by using this script:
 

```
sudo /opt/powerai-vision/bin/powerai_vision_stop.sh
```
7. Install the RPM that was downloaded by running this command:
 

```
sudo yum install ./<fixpack-rpmfile>.rpm
```
8. Log in as root or with sudo privileges, then load the images provided in the TAR file that was downloaded by running this script:
 

```
sudo /opt/powerai-vision/bin/load_images.sh ./<fixpack-tarfile>.tar
```
9. Start PowerAI Vision by running the following script. You must read and accept the license agreement that is displayed before you can use PowerAI Vision.
 

```
sudo /opt/powerai-vision/bin/powerai_vision_start.sh
```

---

## Contacting IBM Support

IBM Support provides assistance with product defects, answers FAQs, and helps users resolve problems with the product.

After trying to find your answer or solution by using other self-help options such as technotes, you can contact IBM Support. Before contacting IBM Support, your company or organization must have an active IBM software maintenance agreement (SWMA), and you must be authorized to submit problems to IBM. For information about the types of available software support, see the Support portfolio topic in the *“Software Support Handbook”*.

To determine what versions of the product are supported, refer to the Software lifecycle page.

To contact IBM Support about a problem:

1. Define the problem, gather background information, and determine the severity of the problem. For software support information, see the Getting IBM support topic in the *Software Support Handbook*.
2. Gather diagnostic information.
3. Submit the problem to IBM Support in one of the following ways:
  - Using IBM Support Assistant (ISA):
  - Online through the IBM Support Portal: You can open, update, and view all of your service requests on the Service Request page.
  - By phone: For the phone number to call in your region, see the Directory of worldwide contacts web page.

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM Support website daily, so that other users who experience the same problem can benefit from the same resolution.



---

## Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:  
© (your company name) (year).  
Portions of this code are derived from IBM Corp. Sample Programs.  
© Copyright IBM Corp. \_enter the year or years\_.

---

## Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Java<sup>™</sup> and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.



---

## Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

### Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

### Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

### Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

### Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



---

## Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing*

*IBM Corporation*

*North Castle Drive, MD-NC119*

*Armonk, NY 10504-1785*

*US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing*

*Legal and Intellectual Property Law*

*IBM Japan Ltd.*

*19-21, Nihonbashi-Hakozakicho, Chuo-ku*

*Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing*

*IBM Corporation*

*North Castle Drive, MD-NC119*

*Armonk, NY 10504-1785*

*US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

#### **COPYRIGHT LICENSE:**

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these

programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. \_enter the year or years\_.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries..







Printed in USA