Ruckus SPoT and vSPoT Release Notes

Release 4.7.0 and Release 3.7.0

Contents

Release Information

Overview

This section contains release information for SPoT 4.7.0 and vSPoT 3.7.0.

This document provides release information on the Ruckus SmartPositioning Technology (SPoT) and virtual SPoT (vSPoT). SPoT/vSPoT requires ZoneDirector or SmartZone wireless LAN controller to communicate between the SPoT location server and the access points.

Please refer to the ZoneDirector, SmartZone and Access Point Release Notes for information on Ruckus controllers and access points.

Most documents are available in PDF format from the Ruckus Support portal: https://support.ruckuswireless.com

Supported Controllers

This SPoT release supports ZoneDirector and SmartZone controllers.

The list is as follows:

ZoneDirector 9.8 and above (ZD 1100, 1200, 3000, 5000) SmartZone 3.0 and above (SZ100, SCG200, vSZ and SZ300)

> **NOTE**
> To manage multiple vSPoT venues on a single ZoneDirector controller, use firmware 9.12.2 and above. This will allow the controller to manage multiple vSPoT Venue IDs using a single IP.

> **NOTE**
> To manage multiple vSPoT venues on a single SmartZone controller, use firmware 3.1.1 patch 1 and above. This will allow the controller to manage multiple vSPoT Venue IDs using a single IP.

Supported Access Point Models

This release supports ZoneFlex and SmartZone Access Points except Unleashed and Cloud-based Access Points.

> **NOTE**
> SPoT support on the 5 GHz radio of 802.11ac APs is available as of ZoneDirector firmware version 9.10 and SmartZone version 3.x

vSPoT Release Information

The following table maps the vSPoT release numbers to the corresponding cloud SPoT releases.

vSPoT Release	SPoT Release
1.0	2.0
1.0.1	2.0
1.1.0	2.1
1.2	2.2
1.3.1	2.3.1
1.4	2.4
1.5	2.5
1.6	2.6
1.7	2.7
1.7.1	2.7.1
1.8	2.8
2.0	3.0
2.3.4	3.3.4
2.3.7	3.3.7
2.4.0	3.4.0
2.4.2	3.4.2
2.4.3	3.4.3
2.4.4	3.4.4
3.0.0	4.0.0
3.1.3	4.1.3
3.1.7	4.1.7
3.2.0	4.2.0
3.3.0	4.3.0
3.4.0	4.4.0
3.4.2	4.4.2
3.4.3	4.4.3
3.5.0	4.5.0
3.6.0	4.6.0
3.7.0	4.7.0

New Features and Enhancements

New Features and Enhancements

This section lists the new features and enhancements in this release.

Enhancements

System

- Location APIs are now available when hashing is turned on. Client MAC retrieved will be hashed.
- Patches for Meltdown vulnerability.

Venue

- None

vSPoT

- None

Caveats, Limitations, and Known Issues

Caveats, Limitations, and Known Issues

This section lists the caveats, limitations and known issues in this release, along with workarounds if applicable.

Limitation

- **None**

Resolved Issues

Resolved Issues

This section contains the list of resolved issues.

- Filtering of random MAC is enabled when hashing is turned on, such that unique visitor is the same with and without hashing. **[ER-5971]**

Upgrading to a New Version

This chapter lists important information that you must be aware of when upgrading vSPoT.

NOTE
A full HOST OS upgrade is required for any versions of vSPoT 3.2.0 or prior versions, since it includes changes related to CentOS7 based docker container with systemd support. System commands to start and stop application data services is different for systemd in comparison to previous vSPoT versions.

Full Upgrade

The following is the upgrade procedure based on the installation type.

CAUTION
In order to upgrade to vSPoT version 3.x, you need to first ensure that the existing vSPoT is on version 2.4.0 and above. If your vSPoT version is below 2.4.0, you cannot directly upgrade to vSPoT 3.x. Do refer to *vSPoT 2.4.4 Installation Guide* on the support website (https://support.ruckuswireless.com) to first upgrade to vSPoT 2.4.4 before upgrading to 3.x.

1. vSPoT on VMware Upgrade from Version 3.x to Latest on page 14

2. vSPoT on VMware Upgrade Version 2.4.x to 3.x (Latest) on page 16

3. vSPoT on AWS Upgrade on page 24

Key Features

The following is the distinction between the two upgrade procedures.

Full Upgrade

NOTE
Full upgrade refers to vSPoT on VMware Upgrade from Version 3.x to Latest on page 14, vSPoT on VMware Upgrade Version 2.4.x to 3.x (Latest) on page 16 or vSPoT on AWS Upgrade on page 24

* Provides better security since it contains host operating system improvements and security updates

* Provides a more atomic upgrade since the host operating system and application containers are built and tested together as part of the release process

* Easier for VMware administrators

Rolling Upgrade

NOTE
This is a mandatory procedure if you are upgrading to vSPoT 3.3.0 or later versions from vSPoT 3.2.0 or earlier versions. Perform a full upgrade to take advantage of the security and performance improvements to the Host OS.

> **NOTE**
> vSPoT upgrades from version 3.3.0 onwards can be performed via rolling upgrades.

- Decreased downtime during the upgrade procedure since there is no switch to virtual disks or having to restart the VM. This saves an enormous amount of downtime
- Easier upgrade procedure since access and storage configuration of VMware is not required
- Decreased manual configuration since networking and NTP server configuration is persisted within the VM, therefore this does not require reconfiguration
- Easier for Linux administrators

vSPoT on VMware Upgrade from Version 3.x to Latest

The following is the procedure for a full upgrade of vSPoT from version 3.x to the latest version.

1. Log in to the vSphere Client:

   a. Launch an instance of VMWare vSphere Client (e.g. Windows VMWare vSphere Client).

   b. In the VMWare vSphere Client, enter the IP Address (or Hostname) and administrative credentials to login to your instance of ESXi server running the source and target vSPoT instances where you want to migrate the data storage volume.

2. Copy the AP MAC addresses seen on the Access Points page (navigate to **SPoT Admin Portal** > **Access Points**), in order to retain a backup of all the AP MAC addresses.

3. Power Off Source and Target vSPoT instances:

   > **NOTE**
   > As an example, the source instance is vSPoT 3.0.0 and the target instance is vSPoT 3.2.

   a. On the left panel, expand the list of installed VM instances, and locate the source and target vSPoT instances that will be swapping data storage volumes.

   b. Right click on each and **Power-Off** each instance successively (shortcut: **Ctrl-E**).

4. Detach source vSPoT instance data storage volume:

   a. Once both source and target instances have been Powered Off, navigate to the source VM instance (note the name of the source vSPoT instance - this will be used to relocate and mount the storage from the source to the target vSPoT instance in later steps). Right-click and choose the '**Edit Settings...**' menu item.

   b. On the child window that appears, on the left panel, locate and identify the item under the 'Hardware' column, named '**Hard disk 3**', and click to select it.

   c. From the same child window, with the 'Hard disk 3' item selected, locate the '**Remove**' button at the top of the left panel of the child window.

   > ⚠️ **CAUTION**
   > **On Removal Options displayed on the right panel, select 'Remove from virtual machine' ONLY. (Do Not use the other option, as this would wipe out the data and therefore make it impossible to migrate the data volume to the new target instance.)**

   d. Notice that the 'Hard disk 3' item is struck out and the 'Summary' column indicates it as 'Removed'; Click '**OK**' at the bottom right hand corner of the child window to proceed.

5. Remove target vSPoT instance data storage volume:

    a. From the list of VM instances on the parent window left panel, locate the target vSPoT instance item, and select and right click on the '**Edit Settings...**' menu item.

    b. Repeat Steps (3b) to (3d) above but for the target vSPoT instance, and click '**OK**' to end the disk removal process, so that the data storage volume to be attached may take its place.

6. Attach source vSPoT instance data storage volume to target vSPoT instance:

    a. From the list of VM instances on the parent window left panel, again locate the target vSPoT instance item, right-click and select the '**Edit Settings...**' menu item.

    b. On the 'Virtual Machine Properties' child window, now locate and click the '**Add...**' button at the top of the left panel of the child window.

    c. On the new 'Add Hardware' child window, select the '**Hard Disk**' item in the middle panel of this child window and click on the '**Next**' button located at the bottom of this child window.

    d. In the 'Select a Disk' option under the 'Add Hardware' child window, select the option '**Use an existing virtual disk - Reuse a previously configured virtual disk**' in the 'Disk' option, then click on the '**Next**' button located at the bottom of this child window.

    e. On the 'Select Existing Disk' option under the 'Add Hardware' child window, click on the '**Browse...**' button next to the 'Disk File Path' text entry field.

    f. From the new Windows file explorer dialog, locate and select the '**Datastores**' file type that was created for your ESXi server, (e.g. esxi-local-storage-1), select the item, then click '**Open**'.

    g. From the expanded list in the 'Browse Datastores' dialog, scroll to and locate the 'Name' identical to your source vSPoT instance that you recorded in Step (3a) above (e.g. *vSPoT2-3.0.0-build-543_vmx*), and select the item, then click '**Open**'.

    h. From the refreshed 'Browse Datastores' dialog, locate and select the item ending with '**_vmx_2.vmdk**', and click '**OK**'.

    i. Back in the parent '**Add Hardware**' -> '**Select Existing Disk**' child window, notice the filled out 'Disk File Path' file based on your selections from the above steps, then click the '**Next**' button.

    j. On the '**Add Hardware**' -> '**Advanced Options**' child window, right side panel, locate the '**Virtual Device Node**' section, and select the item directly beneath the 'SCSI (0:1) Hard disk 2' item (e.g. '**SCSI (0:2)**'), then click the '**Next**' button.

    k. Review the '**Add Hardware**' -> '**Ready to Complete**' -> '**Options**' summary information, and click the '**Finish**' button to complete the data storage volume transfer action, OR, click '**Back**' to alter/edit any previously entered values/choices, OR, click '**Cancel**' to ABORT the operation.

    l. Finally, click the '**OK**' button at the bottom right-hand corner of the '**Virtual Machine Properties**' window to finish the VM settings changes applied/changed/aborted.

    Power On your target vSPoT instance VM and verify the data volume migration action has successfully completed.

7. Verify successful upgrade.

    a. Verify that the web service is running.

    b. View and validate application state from *http://<vSPoT VM IP>:8443*.

    c. Check the application version (example, latest version), and login to the vSPoT administration dashboard using a valid user account and admin account.

8. Verify admin and analytics dashboards.

    a. From the Admin pages, validate the data presence of radio maps, access points, venues, etc.

    b. From the analytics page, verify real-time and historical data (heat maps and daily visit counts)

vSPoT on VMware Upgrade Version 2.4.x to 3.x (Latest)

The following is the procedure to upgrade a vSPoT deployment from version 2.4.x to vSPoT 3.x (latest).

> **NOTE**
> This upgrade process will result in application downtime due to the need for migrating data across different and incompatible storage engines. During the migration process, vSPoT should not be running, otherwise data could be lost during migration.

⚠️ **CAUTION**
In order to upgrade to vSPoT version 3.x, you need to first ensure that the existing vSPoT is on version 2.4.0 and above. If your vSPoT version is below 2.4.x, you cannot directly upgrade to vSPoT 3.x. Do refer to *vSPoT 2.4.4 Installation Guide* on the support website (https://support.ruckuswireless.com) to first upgrade to vSPoT 2.4.4 before upgrading to 3.x.

Prerequisites

- You are required to open port 22 of your vSPoT host operating system (Host OS) on a public facing IP address.

Upgrade Procedure

The upgrade procedure has the following list of instructions.

> **NOTE**
> Ruckus recommends that you follow the procedure for a successful upgrade.
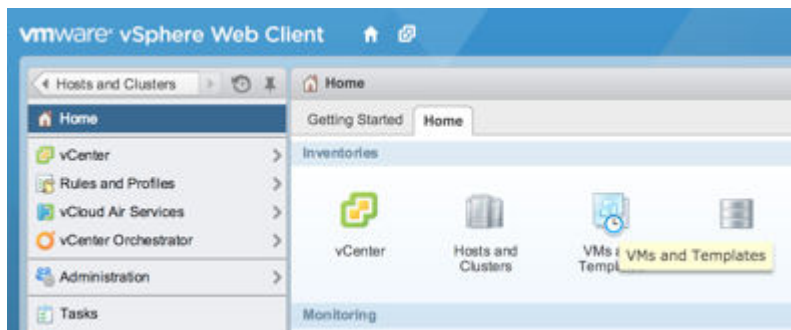
Backingup the Current vSPoT (Cloning)

This section describes the procedure to back up an existing vSPoT instance using the cloning method.

1. Log in to the vSPoT vCenter using the vSphere Web Client and click **Home**.
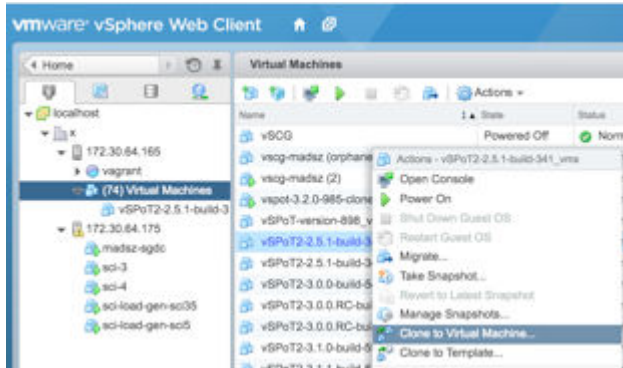
   The **Home** page is displayed.

   FIGURE 1 Viewing the Home Page

2. Click **VMs and Templates**.

   The **Virtual Machines** page is displayed.

   FIGURE 2 Viewing the Virtual Machines

3. Right click the 2.x vSPoT instance and click **Clone to Virtual Machines**.

   The **Clone Existing Virtual Machine** page is displayed.

   FIGURE 3 Cloning an Existing Virtual Machine Instance: Edit Settings
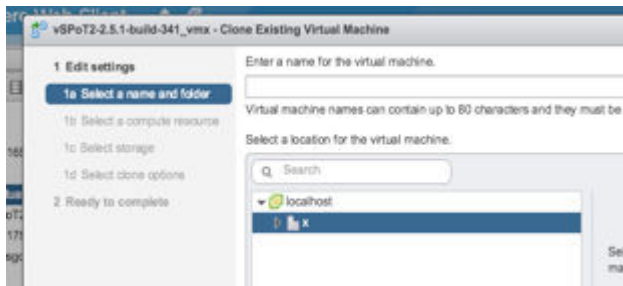
4. Under **Edit Settings**, complete the following steps:

   a) Select the name and folder and click **Next**.

   b) Select a compute resource and click **Next**.

   c) Select the storage and click **Next**.

   d) Select the clone options and click **Next**.

   > **NOTE**
   > Do not select any of the check boxes resulting in changes to the OS and VM configurations.

5.   Under **Ready to Complete**, click **Finish**.
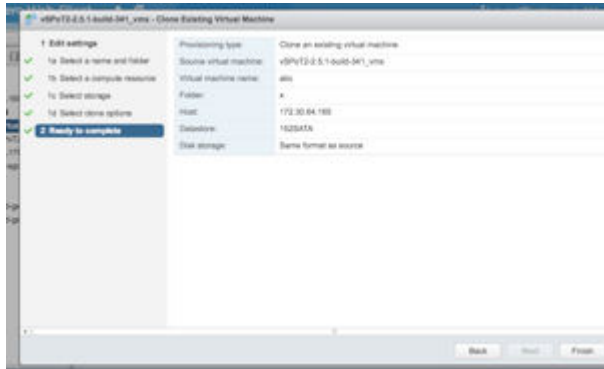
FIGURE 4 Cloning an Existing Virtual Machine Instance: Ready to Complete

The cloned vSPoT VM can be started as a normal vSPoT instance in case the full upgrade fails.

Venue Identifiers

After a successful backup of the vSPoT instance, ensure you have an updated list of Venue identifiers. This is required during migration.

Download vSPoT 3.x

Download the latest vSPoT build version 3.x (example -vSPoT-3.2.0-build-985.tar) from the support site (https://support.ruckuswireless.com)

Stop VM Server Running vSPoT

Stop the existing VM server running vSPoT version 2.4.4 by running the following CLI commands.

1.   SSH to the vSPoT host operating system and to the vSPoT docker container.

> **NOTE**
> Refer to Enabling SSH Access on page 29 if you have not yet configured SSH access.

```
ssh admin@<vspot_host_os_public_ip>
admin@vspotappliance:~$ spot enter
```

2.   Shut down all the running services other than Mongod server.

> **NOTE**
> Keep the Mongod server process up and running.

```
root@vspot:~# stop rk-venue
root@vspot:~# stop resque
root@vspot:~# service nginx stop
```

3.   Exit from the docker container and return to the vSPoT host operating system (Host OS).

Backup the MongoDB

Backup the MongoDB data directory and clean the existing MongoDB data directory contents with the following steps.

1. From vSPoT Host OS, enter the container and download the scripts *vspot_venue_dump.sh* and *vspot_venue_restore.sh* from Ruckus vSPoT AWS S3 migration_scripts, and give the scripts execution rights.

```
admin@vspotappliance:~$ spot enter
root@vspot:~# wget -O /storage/vspot_venue_dump.sh https://s3-us-west-2.amazonaws.com/ruckuslbs/
public/vspot/vspot_venue_dump.sh
root@vspot:~# wget -O /storage/vspot_venue_restore.sh https://s3-us-west-2.amazonaws.com/ruckuslbs/
public/vspot/vspot_venue_restore.sh
root@vspot:~# wget -O /usr/bin/mongodump32 https://s3-us-west-2.amazonaws.com/ruckuslbs/public/vspot/
mongodump32
root@vspot:~# chmod +x /storage/{vspot_venue_dump.sh,vspot_venue_restore.sh} /usr/bin/mongodump32
```

2. Backup the MongoDB data for a specified venue (example *rksg-dev*) by using the below procedure.

> **NOTE**
> If you have more than one venue to migrate, repeat this step for each venue as per your list of Venue Identifiers on page 18.

```
root@vspot:~# bash /storage/vspot_venue_dump.sh rksg-dev --full-dump

Response:
THIS SCRIPT MUST BE RUN IN vSPoT docker container!
IT IS MEANT TO DO A vSPoT MongoDB and radio maps dump for a vSPoT data migration purpose!
+ '[' '!' -f /etc/envvars ']'
+ source /etc/envvars

Creating /storage/venue_dump/dump_mongo and /storage/venue_dump/dump_radio_maps directories ...

DONE.

Proceed to dump groups ...
+ /usr/bin/mongodump32 --gzip --port 27017 --db rk-system_production --query '{'\''_id'\'':
ObjectId("562677639180054fd7000002")}' --collection groups --out /storage/venue_dump/dump_mongo
2016-07-26T14:52:43.841+0800    writing rk-system_production.groups to
2016-07-26T14:52:43.843+0800    done dumping rk-system_production.groups (1 document)

DONE.

Proceed to dump venue rksg-dev ...
+ /usr/bin/mongodump32 --gzip --port 27017 --db rk-system_production --query '{'\''_id'\'': '\''rksg-
dev'\''}' --collection venues --out /storage/venue_dump/dump_mongo
2016-07-26T14:52:43.862+0800    writing rk-system_production.venues to
2016-07-26T14:52:43.863+0800    done dumping rk-system_production.venues (1 document)

DONE.

Proceed to dump venue rksg-dev access_points ...
2016-07-26T15:09:20.821+0800    rk-system_production.impressions  28173377
2016-07-26T15:09:21.962+0800    rk-system_production.impressions  28285371
2016-07-26T15:09:21.963+0800    done dumping rk-system_production.impressions (28285371 documents)

DONE.

Analytics data dump completed DONE.

Copying radio maps for rksg-dev into /storage/venue_dump/dump_radio_maps ...
+ cp -rpfv /opt/spot/system/private/venues/rksg-dev /storage/venue_dump/dump_radio_maps
`/opt/spot/system/private/venues/rksg-dev' -> `/storage/venue_dump/dump_radio_maps/rksg-dev'
`/opt/spot/system/private/venues/rksg-dev/foreground_maps' - `/storage/venue_dump/dump_radio_maps/
rksg-dev/foreground_maps'

DONE.
+ cd /storage/venue_dump
```

```
Creating gzip tarball at /storage/rksg-dev_dump.tar.gz of above mongodumps and radio maps ...
+ tar -zcvf /storage/rksg-dev_dump.tar.gz dump_mongo dump_radio_maps
dump_mongo/
dump_mongo/rk-system_production/
dump_mongo/rk-system_production/groups.metadata.json.gz

dump_radio_maps/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/floor_6.png
dump_radio_maps/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/floor_3.png
dump_radio_maps/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/floor_6_annotated.png

DONE.
Deleting mongodumps and radio maps dump temp dir /storage/venue_dump ...
+ rm -rf /storage/venue_dump
+ ls -alrth /storage/rksg-dev_dump.tar.gz
-rw-r--r-- 1 root root 919M Jul 26 15:10 /storage/rksg-dev_dump.tar.gz

ALL DONE.
```

3. Verify the existence of non-empty archive file.

```
root@vspot:~#ls -lah /storage/rksg-dev_dump.tar.gz
```

Remove Redis Append Only File (AOF)

Upgraded version of vSPoT does not use Redis AOF.

1. This must be removed to prevent data corruption during Redis server startup by executing the following command:

```
admin@vspotappliance:~$ sudo rm -f /storage/redis/data/rk_venue.aof
```

2. Exit and return to the vSPoT Host OS.

Prepare the Existing MongoDB Data Directories

Prepare existing MongoDB data directories by executing the following commands.

1. From vSPoT Host OS, enter vSPoT docker container and stop the MongoDB server process.

```
admin@vspotappliance:~$ spot enter
root@vspot:~# service mongod stop
```

2. Prepare a new MongoDB data directory from the vSPoT docker container and a new working directory.

```
root@vspot:~# mv /storage/mongo /storage/mongo24
root@vspot:~# mkdir -p /storage/mongo/{data,log}
root@vspot:~# chown mongod:mongod -R /storage/mongo
```

3. Exit back to vSPoT Host OS and shut down vSPoT docker container.

```
admin@vspotappliance:~$ spot stop 244
```

Detach Storage from vSPoT

From vSphere user interface, detach the storage volume for the current VM running vSPoT version 2.4.4.

⚠️ **CAUTION**
Do not delete the detached storage volume.

For more details, refer step 4 in vSPoT on VMware Upgrade from Version 3.x to Latest.

Install and Import the vSPoT 3.x VM Image

Install and import the VMDK image into your existing vSphere.

The procedure is similar to that described in vSPoT on VMware Upgrade from Version 3.x to Latest on page 14 but keep in mind that this installation is for vSPoT 3.x.

Attach Storage to vSPoT

From vSphere user interface, attach the storage volume to the current VM running the installed vSPoT 3.x.

Start vSPoT 3.x

Start the vSPoT 3.x VM after the storage volume has been successfully attached.

Migrate Data to New Storage Engine Format

Migrate MongoDB data to be compatible with the new storage engine used in vSPoT 3.x by following the below procedure.

1. Check vSPoT 3.x docker container is running by running the command:

   ```
   admin@vspotappliance:~$ spot list
   ```

   ```
   admin@vspotappliance:~$ spot list
   CONTAINER ID        IMAGE                                                            COMMAND        CREATED        STATUS
   847391fb7cfb        registry.internal.ruckuslbs.com/ruckus/vspot:3.2.0-999   "/sbin/init"     2 weeks ago      Up 18 hours      0.0

   admin@vspotappliance:~$ spot enter
   ```

2. Shut down all the running services other than Mongod server.

   ```
   root@vspot:~# systemctl sop rk-venue.target
   root@vspot:~# systemctl stop resque.target
   root@vspot:~# service nginx stop
   root@vspot:~# service redis-rk_venue stop
   root@vspot:~# service redis-rk_system stop
   root@vspot:~# service mosquitto stop
   ```

3. Check that all processes that was stopped in the above step have been terminated

   ```
   #The following should not return any lines

   root@vspot:~# ps -ef | grep  -E 'ruby|redis|mosquitto|nginx' | grep -v grep

   # If any process is returned, repeat the commands to shutdown the respective process.
   ```

Restore the Data

Restore the MongoDB data with the following steps.

1. Enter the Docker container

   ```
   admin@vspotappliance:~$ spot enter
   ```

2. Run the data restore script to first restore the MongoDB data, then subsequently create database indexes.

NOTE
This may take a long time to complete, based on the migration data set.

```
root@vspot:~# bash /storage/vspot_venue_restore.sh /storage/rksg-dev_dump.tar.gz
```

NOTE
If you had previously exported more than one venue, you will need to repeat this data restore step for each venues as per your list of

Response
```
> THIS SCRIPT MUST BE RUN IN vSPoT docker container!
> .
> ONLY A MongoDB and radio maps restore from a valid vSPoT venue dump will be performed from /
storage/rksg-dev_dump.tar.gz file.
> .
> ADDITIONAL MONGODB DATA CHANGES MUST BE PERFORMED FOR AN AWS SPoT to vSPoT data restoration!
> .
> .
> Deleting any previous /storage/venue_dump dir and contents...
> + rm -rf /storage/venue_dump
> .
> DONE.
> + cd /storage
> + '[' -f /storage/rksg-dev_dump.tar.gz ']'
> + mkdir -p /storage/venue_dump
> .
> Exploding /storage/rksg-dev_dump.tar.gz into /storage/venue_dump ...
> + tar -zxvf /storage/rksg-dev_dump.tar.gz -C /storage/venue_dump
> dump_mongo/
> dump_mongo/rk-system_production/
> dump_mongo/rk-system_production/groups.metadata.json.gz
> dump_mongo/rk-system_production/visits.metadata.json.gz
> dump_mongo/rk-system_production/locations.bson.gz
> .
> .
> dump_mongo/rk-system_production/groups.bson.gz
> dump_mongo/rk-system_production/access_points.metadata.json.gz
> dump_radio_maps/
> dump_radio_maps/rksg-dev/
> dump_radio_maps/rksg-dev/foreground_maps/
> dump_radio_maps/rksg-dev/foreground_maps/pPEVNY51Jtutq2vW8Yxq2w.png
> dump_radio_maps/rksg-dev/foreground_maps/hjiaM6CtqdqMBLb96YkxqA.jpg
> .
> .
> dump_radio_maps/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/floor_6_annotated.png
> .
> DONE.
> .
> Restoring DB data dump ...
> + mongorestore --gzip --port 27017 --noIndexRestore /storage/venue_dump/dump_mongo
> 2016-07-26T16:11:54.938+0800    building a list of dbs and collections to restore from /storage/
venue_dump/dump_mongo dir
> 2016-07-26T16:11:54.940+0800    reading metadata for rk-system_production.visits from /storage/
venue_dump/dump_mongo/rk-system_production/visits.metadata.json.gz
> .
> .
> 2016-07-26T16:26:54.940+0800    [########################]    rk-system_production.visits  4.2
GB/393.5 MB  (1100.9%)
> 2016-07-26T16:26:54.940+0800    [########################]  rk-system_production.impressions  7.4
GB/267.1 MB  (2848.0%)
> 2016-07-26T16:26:54.940+0800
> 2016-07-26T16:26:57.356+0800    [########################]  rk-system_production.impressions  7.4
GB/267.1 MB  (2853.7%)
> 2016-07-26T16:26:57.356+0800    no indexes to restore
> 2016-07-26T16:26:57.356+0800    finished restoring rk-system_production.impressions (28285371
documents)
> 2016-07-26T16:26:57.939+0800    [########################]  rk-system_production.visits  4.2 GB/
393.5 MB  (1104.0%)
> .
```

```
>  .
> 2016-07-26T16:48:15.702+0800     [########################]  rk-system_production.visits  7.9 GB/
393.5 MB  (2062.0%)
> 2016-07-26T16:48:15.702+0800     no indexes to restore
> 2016-07-26T16:48:15.702+0800     finished restoring rk-system_production.visits (28138897 documents)
> 2016-07-26T16:48:15.702+0800     done
>  .
> DONE.
> + cd /opt/spot/system
>  .
> Create indexes for Analytics DB ...
> + bundle exec rake mi:create_indexes
> Run options: --seed 10490
> # Running:
>  .
> Finished in 0.002529s, 0.0000 runs/s, 0.0000 assertions/s.
> 0 runs, 0 assertions, 0 failures, 0 errors, 0 skips
>  .
> DONE.
>  .
> Restoring radio maps dump ...
> + cp -rpfv /storage/venue_dump/dump_radio_maps/rksg-dev /opt/spot/system/private/venues/
> `/storage/venue_dump/dump_radio_maps/rksg-dev/foreground_maps/xIhaFqAniklXPImRY1dewg.png' -> `/opt/
spot/system/private/venues/rksg-dev/foreground_maps/xIhaFqAniklXPImRY1dewg.png'
>  .
>  .
> `/storage/venue_dump/dump_radio_maps/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/venue.svg' ->
`/opt/spot/system/private/venues/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/venue.svg'
> `/storage/venue_dump/dump_radio_maps/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/
floor_3_inside.png' -> `/opt/spot/system/private/venues/rksg-dev/radio_maps/56383dbb3d6ae907d7000012/
floor_3_insid
e.png'
>  .
> DONE.
>  .
> ALL DONE.
root@vspot:~#
```

Verify Successful Migration

Start the required services for verification.

```
root@vspot:~# service mosquitto start
root@vspot:~# service redis-rk_venue start
root@vspot:~# service redis-rk_system start
root@vspot:~# service nginx start
```

View and validate application state from http://<vSPoT VM IP>:8443. Check the application version, and its ability to login to vSPoT version 3.x administration dashboard using a valid user account and admin account.

Start Data Services

Start the data services and background services.

```
root@vspot:~# systemctl start rk-venue.target
root@vspot:~# systemctl start resque.target
```

> **NOTE**
> This will begin to receive live data from the access points.

From the Admin pages, validate the data presence of radio maps, access points, venues, etc. From the analytics page, verify real-time and historical data (heat maps and daily visit counts).

Clean Directories and Backup Snapshots

After successfully migrating to vSPoT version 3.x follow the below procedure

1. Clean the earlier database directory by running the below command.

    ```
    admin@vspotappliance:~$ sudo rm -fr /storage/mongo24
    ```

2. Exit to vSPoT Host OS and from VMWare vCenter web interface.

3. You may want to delete or cleanup your VMWare snapshots after you have ascertained that the upgrade is successful. You may no longer need to revert to the previous version of vSPoT VM.

Re-Import vSPoT License File

Re-import the vSPoT license to get the full working license purchased for your vSPoT instance.

vSPoT on AWS Upgrade

Virtual SPoT in AWS can be upgraded with the following steps for version 2.4.4 to 3.x latest.

> ⚠️ **CAUTION**
> **This upgrade process will result in application downtime due to the need for data migration across different and incompatible storage engines. During the migration process, vSPoT should not be running, otherwise data could be lost during migration.**

Prerequisites

You are required to get full administration access to AWS billing account and management console.

Upgrade Procedure

The following are the steps for upgrading vSPoT in AWS for version 2.4.4 to 3.x.

Backup the Current vSPoT

This section describes the procedure to backup the current vSPoT.

1. Login to the AWS EC2 management console and obtain the identifier **INSTANCE ID** of your vSPoT instance, which is in the format *i-ffffffff*.

2. Navigate to **Elastic Block Store** > **Volumes** to search for volumes using the identifier **INSTANCE ID** obtained in Step 1. In this example, the identifier is *i-baf7cda2*.

FIGURE 5 Create Snapshot
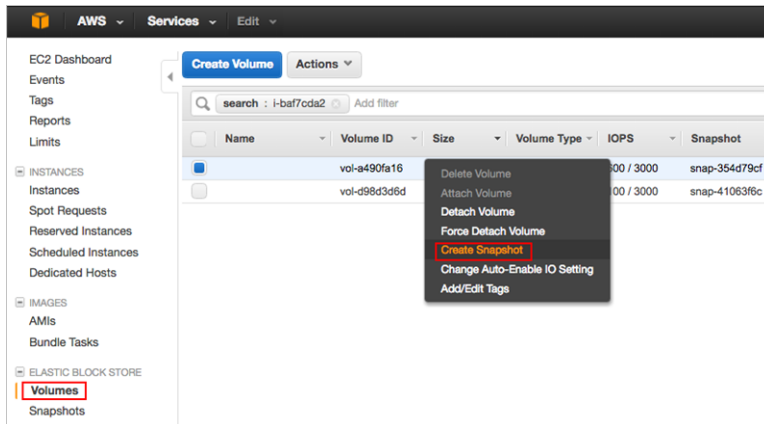
3. Select the vSPoT instance and right-click to select **Create Snapshot**.

4. After snapshot creation is triggered, the process will run in the background. You can proceed to the next step.

Venue Identifiers

After a successful backup of the vSPoT instance, ensure you have an updated list of Venue identifiers. This is required during migration.

Download vSPoT 3.x Template

Download the latest **Ruckus vSPoT Software Release** template for AWS from (https://support.ruckuswireless.com)

FIGURE 6 Download Details

Stop vSPoT Server

Stop the existing server running vSPoT version 2.4.4 by using the following CLI commands.

> **NOTE**
> Ensure you have the IP address and SSH private key file for accessing your AWS EC2 instance.

1. Refer to Accessing vSPoT using AWS CLI on page 30 to login using SSH to the vSPoT AWS EC2 instance host operating system (Host OS). Enter the vSPoT docker container.

```
ssh centos@IP_ADDR -i $SSH_PRIVATE_KEY_FILE
centos@ip-addr:~$ spot enter
```

2. Shut down all the running services other than Mongod server.

```
root@vspot:~# stop rk-venue
root@vspot:~# stop resque
root@vspot:~# service nginx stop
```

Backup the MongoDB

Backup the MongoDB data directory and clean the existing MongoDB data directory contents with the following steps.

1. From the vSPoT docker container, download the Ruckus vSPoT AWS S3 migration scripts, and give it execution rights..

```
admin@vspotappliance:~$ spot enter
root@vspot:~# wget -O /storage/vspot_venue_dump.sh https://s3-us-west-2.amazonaws.com/ruckuslbs/
public/vspot/vspot_venue_dump.sh
root@vspot:~# wget -O /storage/vspot_venue_restore.sh https://s3-us-west-2.amazonaws.com/ruckuslbs/
public/vspot/vspot_venue_restore.sh
root@vspot:~# wget -O /usr/bin/mongodump32 https://s3-us-west-2.amazonaws.com/ruckuslbs/public/vspot/
mongodump32
root@vspot:~# chmod +x /storage/{vspot_venue_dump.sh,vspot_venue_restore.sh} /usr/bin/mongodump32
```

2. Backup the MongoDB data for a specified venue (example *rksg-dev*) by using the below procedure.

> **NOTE**
> If you have more than one venue to migrate, repeat this step for each venue as per your list of Venue Identifiers on page 25.

```
root@vspot:~# bash /storage/vspot_venue_dump.sh rksg-dev --full-dump
```

3. Verify the existence of non-empty archive file.

```
root@vspot:~#ls -lah /storage/rksg-dev_dump.tar.gz
```

Remove Redis Append Only File (AOF)

Upgraded version of vSPoT does not use Redis AOF.

1. This must be removed to prevent data corruption during Redis server startup by executing the following command:

```
root@vspot:~# rm -f /storage/redis/data/rk_venue.aof
```

Prepare the Existing MongoDB Data Directories

Prepare existing MongoDB data directories by executing the following commands.

1. From vSPoT Host OS, enter vSPoT docker container and stop the MongoDB server process.

```
admin@vspotappliance:~$ spot enter
root@vspot:~# service mongod stop
```

2. Prepare a new MongoDB data directory from the vSPoT docker container and a new working directory.

```
root@vspot:~# mv /storage/mongo /storage/mongo24
root@vspot:~# mkdir -p /storage/mongo/{data,log}
root@vspot:~# chown mongod:mongod -R /storage/mongo
```

3. Exit back to vSPoT Host OS and shut down vSPoT docker container.

```
admin@vspotappliance:~$ spot stop 244
```

Remove vSPoT 2.4.4. Stack Template

⚠️ **CAUTION**
Do not delete the detached storage volume from AWS EC2 management console.

1. Refer to Deleting vSPoT on AWS on page 31 to remove the vSPoT 2.4.4 stack. This will also detach the storage volume. Do make a note of the storage volume identifier for later use.

Create a new EC2 Instance with vSPoT 3.x Stack Template

To create a new EC2 instance:

1. Refer to *Automated Setup using CloudFormation* of the installation guide to install vSPoT 3.x using the vSPoT 3.x stack template downloaded . You will also need to specify the StorageVolumeId field as seen in the figure below. The identifier is in the format of *vol-ffffffff* or *vol-ffffffffffffffff*.

FIGURE 7 Specifying Options in the Template

Migrate Data to New Storage Engine Format

Migrate MongoDB data to be compatible with the new storage engine used in vSPoT 3.x by following the below procedure.

1. Check vSPoT 3.x docker container is running by running the command:

```
centos@ip-addr:~$ spot list
```

2. Shut down all the running services other than Mongod server.

```
centos@ip-addr:~# spot enter
root@vspot:~# systemctl stop rk-venue.target
root@vspot:~# systemctl stop resque.target
root@vspot:~# service nginx stop
```

```
root@vspot:~# service redis-rk_venue stop
root@vspot:~# service redis-rk_system stop
root@vspot:~# service mosquitto stop
```

3. Check that all processes that was stopped in the above step have been terminated.

```
root@vspot:~# ps -ef | grep -E 'ruby|redis|mosquitto|nginx' | grep -v grep
```

> **NOTE**
> If any process is returned, repeat the commands to shutdown the processes.

Restore the Data

Restore the MongoDB data with the following steps.

1. Run the data restore script to first restore the MongoDB data, then subsequently create database indexes.

    > **NOTE**
    > This may take a long time to complete, based on the migration data set.

```
root@vspot:~# bash /storage/vspot_venue_restore.sh /storage/rksg-dev_dump.tar.gz
```

> **NOTE**
> If you had previously exported more than one venue, you will need to repeat this data restore step for each
> venues as per your list of Venue Identifiers on page 25.

Verify Successful Migration

Start the required services for verification.

```
root@vspot:~# service mosquitto start
root@vspot:~# service redis-rk_venue start
root@vspot:~# service redis-rk_system start
root@vspot:~# service nginx start
```

View and validate application state from http://<vSPoT VM IP>:8443. Check the application version, and its ability to login to
vSPoT version 3.x administration dashboard using a valid user account and admin account.

Start Data Services

Start the data services and background services.

```
root@vspot:~# systemctl start rk-venue.target
root@vspot:~# systemctl start resque.target
```

> **NOTE**
> This will begin to receive live data from the access points.

From the Admin pages, validate the data presence of radio maps, access points, venues, etc. From the analytics page, verify real-
time and historical data (heat maps and daily visit counts).

Clean Directories and Backup Snapshots

After successfully migrating to vSPoT version 3.x follow the below procedure.

1. Clean the earlier database directory by running the below command.

   ```
   admin@vspotappliance:~$ sudo rm -fr /storage/mongo24
   ```

2. Exit from the SSH terminal session.

3. You may want to delete or cleanup your EC2 EBS snapshots after you have ascertained that the upgrade is successful and you may no longer need to revert to the previous version of vSPoT VM.

Re-Import vSPoT License File

Re-import the vSPoT license to get the full working license purchased for your vSPoT instance.

Sub Sections to Upgrade Procedure

Enabling SSH Access

The options in this section help you enable remote SSH access on logging to the vSPoT CLI console.

- SSH key based authentication
- Enabling password based authentication

SSH Key Based Authentication

This is a preferred and secure option. Execute the following CLI commands for enabling SSH access.

1. After logging in to vSPoT CLI, create the SSH directory in the administrator's home directory.

   ```
   admin@vspotappliance:~$ mkdir -p ~/.ssh
   ```

2. Obtain the SSH public key for your local machine. If you do not have a SSH public key, generate it by following Step 1 of the below reference.

   ```
   https://www.centos.org/docs/5/html/5.2/Deployment_Guide/s3-openssh-rsa-keys-v2.html
   ```

3. Add your local machine's SSH public key to your vSPoT instance's SSH *authorized_keys* file. This can be done in either of the following two ways:

   Example 1: Copy your local machine's SSH public key into the vSPoT instance

   ```
   admin@vspotappliance:~$ echo 'ssh-rsa
   AAAAB3NzaC1yc2EAAAADAQABAAABAQCjfXGGdeNsqTPPWdI6Iext/
   DhZ5GrewC6yh6HZAYjlgDamdZebbAvzIwWvWJjxsOGHkhdF5eY9qGlKdZih
   WVOTaM1oKrUSshAEEPQnYUBa/nF3J3q4nIX4rOiKsBrT44/
   sYKbt+TKgeQ9x5RfwqjqlxCQeq4UbV8K2xfojLCjF6if8jRXYoyUNIf4t0S
   mRzGkEtQ8UqnYDmOGlgu4N+kewcT1laz3ty0YfMK1HToN1+
   RXL256ZDoLT8w0TAo5h/lpiLthgr8t8+UYHSeejEuoR+
   zC2E+37Dr8JVbeTzXAT5zRZhxE7uLkjPCj8HKLekHqyG2225lrdKr
   dWeYmrK/Xp joe' >> ~/.ssh/authorized_keys
   ```

Example 2: From your vSPoT instance, download your SSH public key from a publicly accessible URL

```
admin@vspotappliance:~$
curl ${SSH_KEY_PUBLIC_URL} >> ~/.ssh/authorized_keys

Example: SSH_KEY_PUBLIC_URL
https://raw.githubusercontent.com/mitchellh/vagrant/master/keys/vagrant.pub
```

Enabling Password Based Authentication

This is a less secure option. Execute the following CLI commands for enabling SSH access.

```
admin@vspotappliance:~$ sudo sed -i 's/^PasswordAuthentication.*/PasswordAuthentication yes/'
/etc/ssh/sshd_config

admin@vspotappliance:~$ sudo service sshd restart

Redirecting to /bin/systemctl restart sshd.service

admin@vspotappliance:~$
```

Accessing vSPoT using AWS CLI

This section describes certain technical operations that will require you to log into vSPoT through the shell console using a SSH client.

Follow these steps to login using CLI.

1. Use the SSH command with your private key for the vSPoT instance.

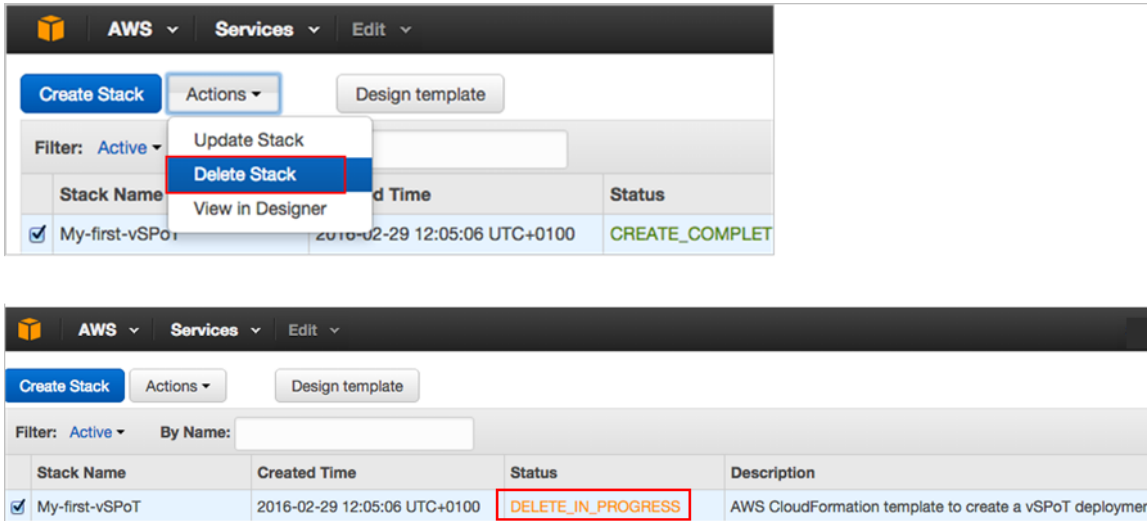2. The generated or imported SSH key is used as the login to the system as seen in the figure below.

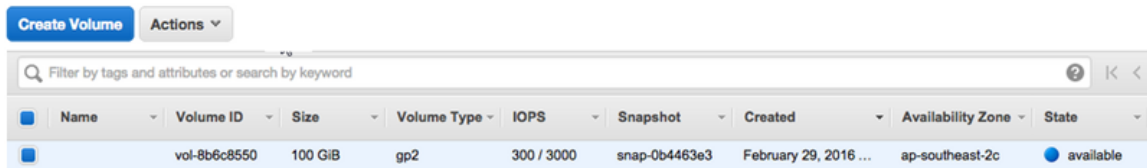**FIGURE 8** SSH Identifier
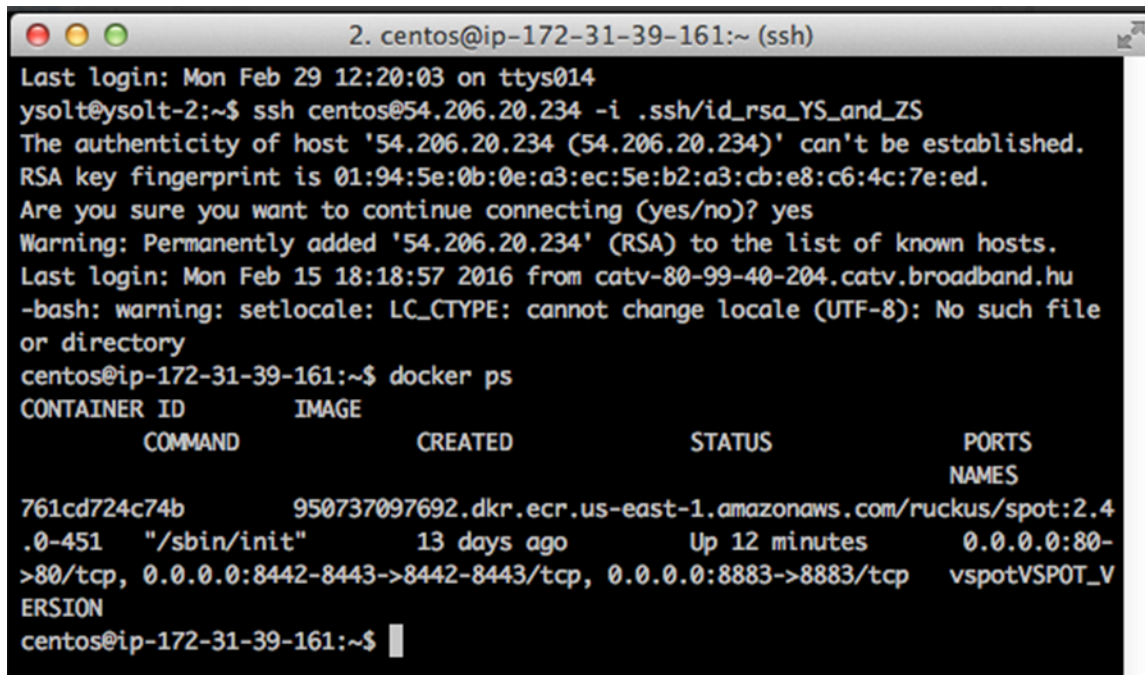
Deleting vSPoT on AWS

To delete a vSPoT instance navigate to **Actions** drop down and click the **Delete Stack** button.

FIGURE 9 Delete AWS CloudFormation Stack

Delete stack retains data storage volume that has been created. To retain the previous historical data navigate to **Amazon Web Services > Compute > EC2 > Elastic Block Store > Volumes** to store or delete the vSPoT data storage volume.

FIGURE 10 AWS Data Storage Volume