Sushy Tools Documentation

Release 2.2.1.dev9

OpenStack Foundation

Mar 02, 2026

CONTENTS

This is a set of simple simulation tools aimed at supporting the development and testing of the Redfish protocol implementations and, in particular, the Sushy library (https://docs.openstack.org/sushy/). It is not designed for use outside of development and testing environments. Please do not run sushy-tools in a production environment of any kind.

The package ships two simulators - the static Redfish responder and the virtual Redfish BMC (which is backed by libvirt or OpenStack cloud).

The static Redfish responder is a simple REST API server which responds with the same things to client queries. It is effectively read-only.

The virtual Redfish BMC resembles the real Redfish-controlled bare metal machine to some extent. Some client queries are translated to commands that actually control VM instances simulating bare metal hardware. However, some of the Redfish commands just return static content, never touching the virtualization backend and in this regard, the virtual Redfish BMC is similar to the static Redfish responder.

- Free software: Apache license

- Documentation: https://docs.openstack.org/sushy-tools

- Source: http://opendev.org/openstack/sushy-tools

- Bugs: https://storyboard.openstack.org/#!/project/openstack/sushy-tools

DOCUMENTATION

1.1 Installation

The sushy-tools Python package can be downloaded and installed with *pip*:

```
$ pip install sushy-tools
```

Or, if you have virtualenvwrapper installed:

```
$ mkvirtualenv sushy-tools
$ pip install sushy-tools
```

The *Virtual Redfish BMC* tool relies upon one or more hypervisors to mimic bare metal nodes. Depending on the virtualization backend you are planning to use, certain third-party dependencies should also be installed.

The dependencies for the virtualization backends that should be installed for the corresponding drivers to become operational are:

- *libvirt-python* for the libvirt driver

- *openstacksdk* for the nova driver

> **Note**
>
> The dependencies for at least one virtualization backend should be satisfied to have the *Virtual Redfish BMC* emulator operational.

1.2 Configuring emulators

1.2.1 Running emulators in the background

The emulators run as interactive processes attached to the terminal by default. You can create systemd services to run the emulators in the background. For each emulator, create a systemd unit file, and update `<full-path>` to the `sushy-static` or `sushy-emulator` binary, and adjust the arguments as necessary, for example:

```
[Unit]
Description=Sushy Libvirt emulator
After=syslog.target
```

```
[Service]
Type=simple
ExecStart=/<full-path>/sushy-emulator --port 8000 --libvirt-uri "qemu:///
↪system"
StandardOutput=syslog
StandardError=syslog
```

If you want to run the emulators with different configurations, for example, the `sushy-static` emulator with different mockup files, then create a new systemd unit file.

You can also use `gunicorn` to run `sushy-emulator`, for example:

```
ExecStart=/usr/bin/gunicorn sushy_tools.emulator.main:app
```

1.2.2 Using configuration file

Besides command-line options, *sushy-emulator* can be configured via a configuration file. This tool uses the Flask application configuration infrastructure. Emulator-specific configuration options are prefixed with **SUSHY_EMULATOR_** to make sure that they don't collide with Flask's own configuration options.

The configuration file itself can be specified through the *SUSHY_EMULATOR_CONFIG* environment variable.

The full list of supported options and their meanings can be found in the sample configuration file:

```
# sushy emulator configuration file built on top of Flask application
# configuration infrastructure: http://flask.pocoo.org/docs/config/

# Listen on all local IP interfaces
SUSHY_EMULATOR_LISTEN_IP = u''

# Bind to TCP port 8000
SUSHY_EMULATOR_LISTEN_PORT = 8000

# Serve this SSL certificate to the clients
SUSHY_EMULATOR_SSL_CERT = None

# If SSL certificate is being served, this is its RSA private key
SUSHY_EMULATOR_SSL_KEY = None

# If authentication is desired, set this to an htpasswd file.
SUSHY_EMULATOR_AUTH_FILE = None

# The OpenStack cloud ID to use. This option enables OpenStack driver.
SUSHY_EMULATOR_OS_CLOUD = None

# If image should created via file upload instead of web-download based␣
↪image
# import OpenStack cloud virtual media
SUSHY_EMULATOR_OS_VMEDIA_IMAGE_FILE_UPLOAD = False
```

```python
# Blank non-bootable image used by the Openstack driver virtual media.
# In "ejected" state the cdrom device is rebuilt with this image.
SUSHY_EMULATOR_OS_VMEDIA_BLANK_IMAGE = 'sushy-tools-blank-image'

# When set to true, the instance rebuild on virtual media eject
# is delayed until the next RedFish power action.
SUSHY_EMULATOR_OS_VMEDIA_DELAY_EJECT = False

# The OpenStack cloud ID to use for Ironic. This option enables Ironic
→driver.
SUSHY_EMULATOR_IRONIC_CLOUD = None

# The libvirt URI to use. This option enables libvirt driver.
SUSHY_EMULATOR_LIBVIRT_URI = u'qemu:///system'

# Instruct the libvirt driver to ignore any instructions to set the boot
→device,
# allowing the UEFI firmware to instead rely on the EFI Boot Manager.
# Note: This sets the legacy boot element to dev="fd" and relies on the
→floppy
# not existing. It likely won't work if your VM has a floppy drive.
SUSHY_EMULATOR_IGNORE_BOOT_DEVICE = False


# The map of firmware loaders dependent on the boot mode and system
# architecture. Ideally the x86_64 loader will be capable of secure boot or
→not
# based on the chosen nvram.
SUSHY_EMULATOR_BOOT_LOADER_MAP = {
    'UEFI': {
        'x86_64': u'/usr/share/OVMF/OVMF_CODE.secboot.fd',
        'aarch64': u'/usr/share/AAVMF/AAVMF_CODE.fd'
    },
    'Legacy': {
        'x86_64': None,
        'aarch64': None
    }
}

# nvram templates to use on x86_64 to enable or disable secure boot
SUSHY_EMULATOR_SECURE_BOOT_ENABLED_NVRAM = '/usr/share/OVMF/OVMF_VARS.
→secboot.fd'
SUSHY_EMULATOR_SECURE_BOOT_DISABLED_NVRAM = '/usr/share/OVMF/OVMF_VARS.fd'

# This map contains statically configured Redfish Chassis linked up with the
# Systems and Managers enclosed into this Chassis.
#
# The first chassis in the list will contain all other resources.
#
```

```
# If this map is not present in the configuration, a single default Chassis
↪is
# configured automatically to enclose all available Systems and Managers.
SUSHY_EMULATOR_CHASSIS = [
    {
        u'Id': u'Chassis',
        u'Name': u'Chassis',
        u'UUID': u'48295861-2522-3561-6729-621118518810'
    }
]

# This map contains statically configured Redfish IndicatorLED resource
↪state
# ('Lit', 'Off', 'Blinking'), keyed by UUIDs of System and Chassis
↪resources.
#
# If this map is not present in the configuration, each System and Chassis
↪will
# have their IndicatorLED `Lit` by default.
#
# The Redfish client can change IndicatorLED state. The new state is
↪volatile,
# i.e. it's maintained in process memory.
SUSHY_EMULATOR_INDICATOR_LEDS = {
#    u'48295861-2522-3561-6729-621118518810': u'Blinking'
}

# This map contains statically configured virtual media resources.
# These devices ('Cd', 'Floppy', 'USBStick') will be exposed by the
↪Manager(s)
# and possibly used by the System(s) if system emulation backend supports
↪boot
# image configuration.
#
# This value is ignored by the OpenStack driver, which only supports the 'Cd
↪'
# device. If this map is not present in the configuration, the following
# configuration is used for other drivers:
SUSHY_EMULATOR_VMEDIA_DEVICES = {
    u'Cd': {
        u'Name': 'Virtual CD',
        u'MediaTypes': [
            u'CD',
            u'DVD'
        ]
    },
    u'Floppy': {
        u'Name': u'Virtual Removable Media',
        u'MediaTypes': [
```

```
            u'Floppy',
            u'USBStick'
        ]
    }
}

# Instruct the virtual media insertion to not verify the SSL certificate␣
↪when
# retrieving the image.
SUSHY_EMULATOR_VMEDIA_VERIFY_SSL = False

# The libvirt storage pool to use.
SUSHY_EMULATOR_STORAGE_POOL = 'default'

# This map contains statically configured Redfish Storage resources linked␣
↪up
# with the Systems resources, keyed by the UUIDs of the Systems.
SUSHY_EMULATOR_STORAGE = {
    "da69abcc-dae0-4913-9a7b-d344043097c0": [
        {
            "Id": "1",
            "Name": "Local Storage Controller",
            "StorageControllers": [
                {
                    "MemberId": "0",
                    "Name": "Contoso Integrated RAID",
                    "SpeedGbps": 12
                }
            ],
            "Drives": [
                "32ADF365C6C1B7BD"
            ]
        }
    ]
}

# This map contains statically configured Redfish Drives resources. The␣
↪Drive
# objects are keyed in a composite fashion using a tuple of the form
# (System_UUID, Storage_ID) referring to the UUID of the System and Id of␣
↪the
# Storage resource, respectively, to which the Drive belongs.
SUSHY_EMULATOR_DRIVES = {
    ("da69abcc-dae0-4913-9a7b-d344043097c0", "1"): [
        {
            "Id": "32ADF365C6C1B7BD",
            "Name": "Drive Sample",
            "CapacityBytes": 89527000000,
            "Protocol": "SAS"
```

```
            }
        ]
}

# This map contains dynamically configured Redfish Volume resources backed␣
↪by
# the libvirt virtualization backend of the dynamic Redfish emulator.
# The Volume objects are keyed in a composite fashion using a tuple of the␣
↪form
# (System_UUID, Storage_ID) referring to the UUID of the System and ID of␣
↪the
# Storage resource, respectively, to which the Volume belongs.
#
# Only the Volumes specified in the map or created via a POST request are
# allowed to be emulated upon by the emulator. Volumes other than these can
# neither be listed nor deleted.
#
# The Volumes in the map missing from the libvirt backend will be created
# dynamically in the pool name specified (provided the pool exists in the
# backend). If the pool name is not specified, the Volume will be created
# automatically in the pool specified by SUSHY_EMULATOR_STORAGE_POOL, by
# default 'default'.
SUSHY_EMULATOR_VOLUMES = {
    ('da69abcc-dae0-4913-9a7b-d344043097c0', '1'): [
        {
            "libvirtPoolName": "sushyPool",
            "libvirtVolName": "testVol",
            "Id": "1",
            "Name": "Sample Volume 1",
            "VolumeType": "Mirrored",
            "CapacityBytes": 23748
        },
        {
            "libvirtPoolName": "sushyPool",
            "libvirtVolName": "testVol1",
            "Id": "2",
            "Name": "Sample Volume 2",
            "VolumeType": "StripedWithParity",
            "CapacityBytes": 48395
        }
    ]
}

# This list contains the identities of instances that the driver will␣
↪filter by.
# It is useful in a tenant environment where only some instances represent
# virtual bare metal.
SUSHY_EMULATOR_ALLOWED_INSTANCES = [
    "437XR1138R2",
```

```
    "1",
    "529QB9450R6",
    "529QB9451R6",
    "529QB9452R6",
    "529QB9453R6"
]


# Disable the ability to power off the node, in line with NCSI enablement in
# Ironic
SUSHY_EMULATOR_DISABLE_POWER_OFF = False
```

1.3 Using Redfish emulators

The sushy-tools package includes two emulators - static and dynamic.

The static emulator can be used to serve Redfish mocks in the form of static JSON documents. The dynamic emulator relies upon the *libvirt*, *OpenStack* or *Ironic* virtualization backends to mimic nodes behind a Redfish BMC.

1.3.1 Static Redfish BMC

The static Redfish responder is a simple REST API server which serves static contents down to the Redfish client. The tool emulates the simple read-only BMC.

The user is expected to supply the Redfish-compliant contents, perhaps downloaded from the DMTF web site. For example, this .zip archive includes Redfish content mocks for Redfish 1.0.0.

```
curl -o DSP2043_1.0.0.zip \
    https://www.dmtf.org/sites/default/files/standards/documents/DSP2043_1.
↪0.0.zip
unzip -d mockups DSP2043_1.0.0.zip
sushy-static -m mockups/public-rackmount
```

Once you have the static emulator running, you can use it as if it was a read-only bare metal controller listening at *localhost:8000* (by default):

```
curl http://localhost:8000/redfish/v1/Systems/
{
    "@odata.type": "#ComputerSystemCollection.ComputerSystemCollection",
    "Name": "Computer System Collection",
    "Members@odata.count": 1,
    "Members": [
        {
            "@odata.id": "/redfish/v1/Systems/437XR1138R2"
        }
    ],
    "@odata.context": "/redfish/v1/$metadata#Systems",
    "@odata.id": "/redfish/v1/Systems",
    "@Redfish.Copyright": "Copyright 2014-2016 Distributed Management Task␣
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
```

```
↪dmtf.org/about/policies/copyright."
}
```

You can mock different Redfish versions as well as different bare metal machines by providing the appropriate Redfish contents.

1.3.2 Virtual Redfish BMC

The Virtual Redfish BMC emulator is functionally similar to the Virtual BMC tool, except that the frontend protocol is Redfish rather than IPMI. The Redfish commands coming from the client are handled by one or more resource-specific drivers.

Feature sets

The emulator can be configured with different feature sets to emulate different hardware. The feature set is supplied either via the `SUSHY_EMULATOR_FEATURE_SET` configuration variable or through the `--feature-set` command line flag.

Supported feature sets are: * `minimum` - only Systems with Boot settings and no other optional fields. * `vmedia` - `minimum` plus Managers, VirtualMedia and EthernetInterfaces. * `full` - all features implemented in the emulator.

Systems resource

For the *Systems* resource, the emulator maintains two drivers relying on a virtualization backend to emulate bare metal machines by means of virtual machines. In addition, there is a fake driver used to mock bare metal machines.

The following sections will explain how to configure and use each of these drivers.

Systems resource driver: libvirt

The first thing you need is to set up some libvirt-managed virtual machines (AKA domains) to manipulate. The following command will create a new virtual machine i.e. libvirt domain *vbmc-node*:

```
tmpfile=$(mktemp /tmp/sushy-domain.XXXXXX)
virt-install \
   --name vbmc-node \
   --ram 1024 \
   --disk size=1 \
   --vcpus 2 \
   --os-type linux \
   --os-variant fedora28 \
   --graphics vnc \
   --print-xml > $tmpfile
virsh define --file $tmpfile
rm $tmpfile
```

Next you can fire up the Redfish virtual BMC which will listen at *localhost:8000* (by default):

```
sushy-emulator
 * Running on http://localhost:8000/ (Press CTRL+C to quit)
```

Now you should be able to see your libvirt domain among the Redfish *Systems*:

```
curl http://localhost:8000/redfish/v1/Systems/
{
    "@odata.type": "#ComputerSystemCollection.ComputerSystemCollection",
    "Name": "Computer System Collection",
    "Members@odata.count": 1,
    "Members": [

        {
            "@odata.id": "/redfish/v1/Systems/vbmc-node"
        }

    ],
    "@odata.context": "/redfish/v1/$metadata#ComputerSystemCollection.
↪ComputerSystemCollection",
    "@odata.id": "/redfish/v1/Systems",
    "@Redfish.Copyright": "Copyright 2014-2016 Distributed Management Task␣
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
↪dmtf.org/about/policies/copyright."
}
```

You should be able to flip its power state via the Redfish call:

```
curl -d '{"ResetType":"On"}' \
    -H "Content-Type: application/json" -X POST \
     http://localhost:8000/redfish/v1/Systems/vbmc-node/Actions/
↪ComputerSystem.Reset

curl -d '{"ResetType":"ForceOff"}' \
    -H "Content-Type: application/json" -X POST \
     http://localhost:8000/redfish/v1/Systems/vbmc-node/Actions/
↪ComputerSystem.Reset
```

You can have as many domains as you need. The domains can be concurrently managed over Redfish and some other tool like *Virtual BMC*.

Simple Storage resource

For emulating the *Simple Storage* resource, some additional preparation is required on the host side.

First, you need to create, build and start a libvirt storage pool using virsh:

```
virsh pool-define-as testPool dir - - - - "/testPool"
virsh pool-build testPool
virsh pool-start testPool
virsh pool-autostart testPool
```

Next, create a storage volume in the above created storage pool:

```
virsh vol-create-as testPool testVol 1G
```

Next, attach the created volume to the virtual machine/domain:

```
virsh attach-disk vbmc-node /testPool/testVol sda
```

Now, query the *Simple Storage* resource collection for the *vbmc-node* domain in a closely similar format (with 'ide' and 'scsi', here, referring to the two Redfish Simple Storage Controllers available for this domain):

```
curl http://localhost:8000/redfish/v1/vbmc-node/SimpleStorage
{
    "@odata.type": "#SimpleStorageCollection.SimpleStorageCollection",
    "Name": "Simple Storage Collection",
    "Members@odata.count": 2,
    "Members": [


            {
                "@odata.id": "/redfish/v1/Systems/vbmc-node/
↪SimpleStorage/ide"
            },

            {
                "@odata.id": "/redfish/v1/Systems/vbmc-node/
↪SimpleStorage/scsi"
            }

    ],
    "Oem": {},
    "@odata.context": "/redfish/v1/$metadata#SimpleStorageCollection.
↪SimpleStorageCollection",
    "@odata.id": "/redfish/v1/Systems/vbmc-node/SimpleStorage"
}
```

UEFI boot

By default, *legacy* or *BIOS* mode is used to boot the instance. However, the libvirt domain can be configured to boot via UEFI firmware. This process requires additional preparation on the host side.

On the host you need to have OVMF firmware binaries installed. Fedora users could pull them as *edk2-ovmf* RPM. On Ubuntu, *apt-get install ovmf* should do the job.

Then you need to create a VM by running *virt-install* with the UEFI-specific *–boot* options:

Example:

```
tmpfile=$(mktemp /tmp/sushy-domain.XXXXXX)
virt-install \
   --name vbmc-node \
   --ram 1024 \
   --boot loader.readonly=yes \
   --boot loader.type=pflash \
   --boot loader.secure=no \
   --boot loader=/usr/share/OVMF/OVMF_CODE.secboot.fd \
   --boot nvram.template=/usr/share/OVMF/OVMF_VARS.fd \
   --disk size=1 \
```

```
  --vcpus 2 \
  --os-type linux \
  --os-variant fedora28 \
  --graphics vnc \
  --print-xml > $tmpfile
virsh define --file $tmpfile
rm $tmpfile
```

This will create a new *libvirt* domain with the path to OVMF images properly configured. Let's take a note on the path to the blob by running *virsh dumpxml vbmc-node*:

Example:

```
<domain type="kvm">
  ...
  <os>
    <type arch="x86_64" machine="q35">hvm</type>
    <loader readonly="yes" type="pflash" secure="no">/usr/share/edk2/ovmf/
↪OVMF_CODE.secboot.fd</loader>
    <nvram template="/usr/share/edk2/ovmf/OVMF_VARS.fd"/>
    <boot dev="hd"/>
  </os>
  ...
</domain>
```

Because now we need to add this path to the emulator's configuration matching the VM architecture we are running. It is also possible to make Redfish calls to enable or disable Secure Boot by specifying which nvram template to load in each case. Make a copy of the stock configuration file and edit it accordingly:

```
$ cat sushy-tools/doc/source/admin/emulator.conf
...
SUSHY_EMULATOR_BOOT_LOADER_MAP = {
    'Uefi': {
        'x86_64': '/usr/share/OVMF/OVMF_CODE.secboot.fd',
        ...
}
SUSHY_EMULATOR_SECURE_BOOT_ENABLED_NVRAM = '/usr/share/OVMF/OVMF_VARS.
↪secboot.fd'
SUSHY_EMULATOR_SECURE_BOOT_DISABLED_NVRAM = '/usr/share/OVMF/OVMF_VARS.fd'
...
```

Now you can run *sushy-emulator* with the updated configuration file:

```
sushy-emulator --config emulator.conf
```

> **Note**
>
> The images you will serve to your VMs need to be UEFI-bootable.

Settable boot image

The *libvirt* system emulation backend supports setting custom boot images, so that libvirt domains (representing bare metal nodes) can boot from user images.

This feature enables system boot from virtual media device.

The limitations:

- Only ISO images are supported

See *VirtualMedia* resource section for more information on how to perform virtual media boot.

Systems resource driver: OpenStack

You can use OpenStack cloud instances to simulate Redfish-managed bare metal machines. This setup is known under the name of OpenStack Virtual Baremetal. We will largely reuse its OpenStack infrastructure and configuration instructions. After all, what we are trying to do here is to set up the Redfish emulator alongside the openstackbmc tool which is used for exactly the same purpose at OVB with the only difference being that it works over the *IPMI* protocol as opposed to *Redfish*.

The easiest way is probably to set up your OpenStack Virtual Baremetal cloud by following its instructions.

> **Note**
>
> The OVB documentation includes Nova patches to support the `libvirt:pxe-first` metadata flag for PXE network booting. If you cannot patch your Nova deployment, see the *PXE boot via Nova rescue mode* section for an alternative approach that works with vanilla Nova.

Once your OVB cloud is operational, you log into the *BMC* instance and *set up sushy-tools* there.

Next you can invoke the Redfish virtual BMC pointing it to your OVB cloud:

```
sushy-emulator --os-cloud rdo-cloud
 * Running on http://localhost:8000/ (Press CTRL+C to quit)
```

By this point you should be able to see your OpenStack instances among the Redfish *Systems*:

```
curl http://localhost:8000/redfish/v1/Systems/
{
    "@odata.type": "#ComputerSystemCollection.ComputerSystemCollection",
    "Name": "Computer System Collection",
    "Members@odata.count": 1,
    "Members": [

        {
            "@odata.id": "/redfish/v1/Systems/8dbe91da-4002-4d61-a56d-
↪1a00fc61c35d"
        }

    ],
    "@odata.context": "/redfish/v1/$metadata#ComputerSystemCollection.
↪ComputerSystemCollection",
```

(continues on next page)

```
    "@odata.id": "/redfish/v1/Systems",
    "@Redfish.Copyright": "Copyright 2014-2016 Distributed Management Task.
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
↪dmtf.org/about/policies/copyright."
}
```

And flip an instance's power state via the Redfish call:

```
curl -d '{"ResetType":"On"}' \
    -H "Content-Type: application/json" -X POST \
     http://localhost:8000/redfish/v1/Systems/vbmc-node/Actions/
↪ComputerSystem.Reset

curl -d '{"ResetType":"ForceOff"}' \
    -H "Content-Type: application/json" -X POST \
     http://localhost:8000/redfish/v1/Systems/vbmc-node/Actions/
↪ComputerSystem.Reset
```

You can have as many OpenStack instances as you need. The instances can be concurrently managed over Redfish and functionally similar tools.

Creating Openstack instances for virtual media boot

When creating Openstack instances for virtual media boot the instances must be configured to boot from volumes. One volume configured with `device_type: disk` and `boot_index: 1`. A second volume configured with `device_type: cdrom`, `disk_bus: scsi` and `boot_index: 0`.

The `cdrom` volume should initially be created with small (1 Megabyte) "blank" non-bootable image so that the server boots from the `disk`. On insert/eject, this volume will be rebuilt. Following is an example showing how to create a "blank" image, and upload to glance:

```
qemu-img create -f qcow2 blank-image.qcow2 1M
openstack image create --disk-format qcow2 --file blank-image.qcow2 \
  --property hw_firmware_type=uefi --property hw_machine_type=q35 \
  --property os_shutdown_timeout=5 \
  sushy-tools-blank-image
```

The following is an example show `block_device_mapping` that can be used to when creating an instance using create_server from the Openstack SDK.

```
block_device_mapping=[
    {
        'uuid': IMAGE_ID,
        'boot_index': 1,
        'source_type': 'image',
        'destination_type': 'volume',
        'device_type': 'disk',
        'volume_size': 20,
        'delete_on_termination': True,
    },
    {
```

```python
        'uuid': BLANK_IMG_ID,
        'boot_index': 0,
        'source_type': 'image',
        'destination_type': 'volume',
        'device_type': 'cdrom',
        'disk_bus': 'scsi',
        'volume_size': 5,
        'delete_on_termination': True,
    }
]
```

The following is an example Openstack heat template for creating an instance:

```yaml
ironic0:
  type: OS::Nova::Server
  properties:
    flavor: m1.medium
    block_device_mapping_v2:
      - device_type: disk
        boot_index: 1
        image_id: glance-image-name
        volume_size: 40
        delete_on_termination: true
      - device_type: cdrom
        disk_bus: scsi
        boot_index: 0
        image_id: sushy-tools-blank-image
        volume_size: 5
        delete_on_termination: true
```

PXE boot via Nova rescue mode

The OpenStack driver supports PXE network booting of instances by leveraging Nova's rescue mode. This feature allows instances to boot from a network-boot image such as iPXE.

This approach works with vanilla Nova deployments without requiring custom patches[1] to support the `libvirt:pxe-first` metadata flag, making it suitable for standard OpenStack environments.

Configuration

To enable this feature, configure the following in your emulator configuration:

```python
# Enable PXE boot via rescue mode
SUSHY_EMULATOR_OS_RESCUE_PXE_BOOT = True

# Specify rescue images for BIOS and UEFI boot modes
SUSHY_EMULATOR_OS_RESCUE_PXE_IMAGE_BIOS = 'ipxe-rescue-bios'
SUSHY_EMULATOR_OS_RESCUE_PXE_IMAGE_UEFI = 'ipxe-rescue-uefi'
```

---

[1] https://opendev.org/openstack/openstack-virtual-baremetal/src/branch/stable/2.0/patches/nova

```
# Optional: Directory for persistent state storage
SUSHY_EMULATOR_STATE_DIR = '/var/lib/sushy-emulator'
```

The rescue images should be pre-created in Glance and contain the necessary PXE boot environment (e.g., iPXE).

> **Note**
>
> When `SUSHY_EMULATOR_STATE_DIR` is configured, boot mode settings persist across emulator restarts.

Requirements

- Nova compute service must support the rescue operation
- Rescue images must be available in Glance
- Instances must be created with boot mode metadata (BIOS or UEFI)

Usage

Setting boot mode: When a Redfish client sets `BootSourceOverrideTarget` to `Pxe`, the instance is configured to boot from the network on the next power cycle.

Power on: When powering on an instance configured for PXE boot:

- The instance boots into rescue mode using the appropriate rescue image (BIOS or UEFI) based on the instance's boot mode
- If the instance is already in rescue mode, it remains in rescue mode (no new rescue operation is triggered)

Power off: Powering off an instance in rescue mode automatically exits rescue mode before stopping the instance.

Reboot: Rebooting an instance in rescue mode keeps it in rescue mode.

Changing boot device: Setting `BootSourceOverrideTarget` back to `Hdd` configures the instance to boot normally (without rescue mode) on the next power cycle.

> **Note**
>
> During rescue and unrescue operations, the emulator may return `503 Service Unavailable` (SYS518) if the operation is still in progress. Clients should retry the request after a short delay.

Example workflow

```
# Set boot device to PXE
curl -X PATCH -H 'Content-Type: application/json' \
    -d '{"Boot": {"BootSourceOverrideTarget": "Pxe"}}' \
    http://localhost:8000/redfish/v1/Systems/<uuid>
```

```
# Power on (will boot via rescue with PXE image)
curl -d '{"ResetType":"On"}' \
    -H "Content-Type: application/json" -X POST \
    http://localhost:8000/redfish/v1/Systems/<uuid>/Actions/ComputerSystem.
↪Reset

# Later, set boot back to disk
curl -X PATCH -H 'Content-Type: application/json' \
    -d '{"Boot": {"BootSourceOverrideTarget": "Hdd"}}' \
    http://localhost:8000/redfish/v1/Systems/<uuid>

# Reboot (will exit rescue and boot normally)
curl -d '{"ResetType":"ForceRestart"}' \
    -H "Content-Type: application/json" -X POST \
    http://localhost:8000/redfish/v1/Systems/<uuid>/Actions/ComputerSystem.
↪Reset
```

> **Note**
>
> Boot device changes take effect on the next power cycle, similar to physical BMC behavior.

Limitations

- Volume-backed instances are not supported (Nova does not support rescue for volume-backed instances)

Systems resource driver: Ironic

You can use the Ironic driver to manage Ironic baremetal instance to simulated Redfish API. You may want to do this if you require a redfish-compatible endpoint but don't have direct access to the BMC (you only have access via Ironic) or the BMC doesn't support redfish.

Assuming your bare metal cloud is set up you can invoke the Redfish emulator by running:

```
sushy-emulator --ironic-cloud baremetal-cloud
 * Running on http://localhost:8000/ (Press CTRL+C to quit)
```

By this point you should be able to see your Bare metal instances among the Redfish *Systems*:

```
curl http://localhost:8000/redfish/v1/Systems/
{
    "@odata.type": "#ComputerSystemCollection.ComputerSystemCollection",
    "Name": "Computer System Collection",
    "Members@odata.count": 1,
    "Members": [

        {
            "@odata.id": "/redfish/v1/Systems/<uuid>"
```

```
        }
    ],
    "@odata.context": "/redfish/v1/$metadata#ComputerSystemCollection.
↪ComputerSystemCollection",
    "@odata.id": "/redfish/v1/Systems",
    "@Redfish.Copyright": "Copyright 2014-2016 Distributed Management Task␣
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
↪dmtf.org/about/policies/copyright."
}
```

And flip an instance's power state via the Redfish call:

```
curl -d '{"ResetType":"On"}' \
    -H "Content-Type: application/json" -X POST \
    http://localhost:8000/redfish/v1/Systems/<uuid>/Actions/ComputerSystem.
↪Reset

curl -d '{"ResetType":"ForceOff"}' \
    -H "Content-Type: application/json" -X POST \
    http://localhost:8000/redfish/v1/Systems/<uuid>/Actions/ComputerSystem.
↪Reset
```

Or update their boot device:

```
curl -d '{"Boot":{"BootSourceOverrideTarget":"Pxe"}}' \
    -H "Content-Type: application/json" -X PATCH \
    http://localhost:8000/redfish/v1/Systems/<uuid>

curl -d '{"Boot":{"BootSourceOverrideTarget":"Hdd"}}' \
    -H "Content-Type: application/json" -X PATCH \
    http://localhost:8000/redfish/v1/Systems/<uuid>
```

Systems resource driver: fake

The `fake` system driver is designed to conduct large-scale testing of Ironic without having a lot of bare metal machines or being able to create a large number of virtual machines. When the Redfish emulator is configured with the `fake` system backend, all operations just return success. Any modifications are done purely in the local cache. This way, many Ironic operations can be tested at scale without access to a large computing pool.

System status notifications

The `fake` driver may need to simulate components that run on the VMs to test an end-to-end deployment. This requires a hook interface to integrate external components. For instance, when testing Ironic scalability, Ironic needs to communicate with the Ironic Python Agent (IPA). A fake IPA can be implemented and synchronized with the VM status using this hook, which notifies the fake IPA whenever the VM status changes.

To enable notifications, set `external_notifier` to `True` in the fake System object:

```
{
    "uuid": "7946b59-9e44-4fa7-8e91-f3527a1ef094",
    "name": "fake",
    "power_state": "Off",
    "external_notifier": True,
    "nics": [
        {
            "mac": "00:5c:52:31:3a:9c",
            "ip": "172.22.0.100"
        }
    ]
}
```

After this, whenever the fake driver updates this System object, it will send an HTTP `PUT` request with the new system object as `JSON` data. The endpoint URL can be configured with the parameter `EXTERNAL_NOTIFICATION_URL`.

Filtering by allowed instances

It is not always desirable to manage every accessible virtual machine as a Redfish System, such as when an OpenStack tenant has many instances which do not represent virtual bare metal. In this case it is possible to specify a list of UUIDs which are allowed.

```
$ cat sushy-tools/doc/source/admin/emulator.conf
...
SUSHY_EMULATOR_ALLOWED_INSTANCES = [
    "437XR1138R2",
    "1",
    "529QB9450R6",
    "529QB9451R6",
    "529QB9452R6",
    "529QB9453R6"
]
...
```

Managers resource

Managers are emulated based on Systems: each *System* has a *Manager* with the same UUID. The first manager (alphabetically) will pretend to manage all *Chassis* and potentially other resources.

Managers will be revealed when querying the *Managers* resource directly, as well as other resources they manage or have some other relations.

```
curl http://localhost:8000/redfish/v1/Managers
{
    "@odata.type": "#ManagerCollection.ManagerCollection",
    "Name": "Manager Collection",
    "Members@odata.count": 1,
    "Members": [

        {
```

```
            "@odata.id": "/redfish/v1/Managers/58893887-8974-2487-2389-
↪841168418919"
        }

    ],
    "@odata.context": "/redfish/v1/$metadata#ManagerCollection.
↪ManagerCollection",
    "@odata.id": "/redfish/v1/Managers",
    "@Redfish.Copyright": "Copyright 2014-2017 Distributed Management Task␣
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
↪dmtf.org/about/policies/copyright."
```

Chassis resource

For emulating the *Chassis* resource, the user can statically configure one or more imaginary chassis. All existing resources (e.g. *Systems*, *Managers*, *Drives*) will pretend to reside in the first chassis.

```
SUSHY_EMULATOR_CHASSIS = [
    {
        "Id": "Chassis",
        "Name": "Chassis",
        "UUID": "48295861-2522-3561-6729-621118518810"
    }
]
```

By default a single chassis with be configured automatically.

Chassis will be revealed when querying the *Chassis* resource directly, as well as other resources they manage or have some other relations.

```
curl http://localhost:8000/redfish/v1/Chassis
{
    "@odata.type": "#ChassisCollection.ChassisCollection",
    "Name": "Chassis Collection",
    "Members@odata.count": 1,
    "Members": [
        {
            "@odata.id": "/redfish/v1/Chassis/48295861-2522-3561-6729-
↪621118518810"
        }
    ],
    "@odata.context": "/redfish/v1/$metadata#ChassisCollection.
↪ChassisCollection",
    "@odata.id": "/redfish/v1/Chassis",
    "@Redfish.Copyright": "Copyright 2014-2017 Distributed Management Task␣
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
↪dmtf.org/about/policies/copyright."
```

Indicator resource

The *IndicatorLED* resource is emulated as a persistent emulator database record, observable and manageable by a Redfish client.

By default, the *Chassis* and *Systems* resources have emulated *IndicatorLED* sub-resources attached and *Lit*.

Non-default initial indicator state can optionally be configured on a per-resource basis:

```
SUSHY_EMULATOR_INDICATOR_LEDS = {
    "48295861-2522-3561-6729-621118518810": "Blinking"
}
```

Indicator LEDs will be revealed when querying any resource having *IndicatorLED*:

```
$ curl http://localhost:8000/redfish/v1/Chassis/48295861-2522-3561-6729-
↪621118518810
{
    "@odata.type": "#Chassis.v1_5_0.Chassis",
    "Id": "48295861-2522-3561-6729-621118518810",
    "Name": "Chassis",
    "UUID": "48295861-2522-3561-6729-621118518810",
    ...
    "IndicatorLED": "Lit",
    ...
}
```

Redfish client can turn *IndicatorLED* into a different state:

```
curl -d '{"IndicatorLED": "Blinking"}' \
    -H "Content-Type: application/json" -X PATCH \
    http://localhost:8000/redfish/v1/Chassis/48295861-2522-3561-6729-
↪621118518810
```

Virtual media resource

The Virtual Media resource is emulated as a persistent emulator database record, observable and manageable by a Redfish client.

By default, a *VirtualMedia* resource includes two emulated removable devices: *Cd* and *Floppy*. Each *Manager* resource gets its own collection of virtual media devices as a *VirtualMedia* sub-resource.

If the currently used *Systems* resource emulation driver supports setting the boot image, the *VirtualMedia* resource will apply the inserted image onto all the systems being managed by this manager. Setting the system boot source to *Cd* and boot mode to *Uefi* will cause the system to boot from the virtual media image.

The user can change virtual media devices and their properties through emulator configuration (except for the OpenStack driver which only supports *Cd*):

```
SUSHY_EMULATOR_VMEDIA_DEVICES = {
    "Cd": {
        "Name": "Virtual CD",
        "MediaTypes": [
            "CD",
```

(continues on next page)

```
            "DVD"
        ]
    },
    "Floppy": {
        "Name": "Virtual Removable Media",
        "MediaTypes": [
            "Floppy",
            "USBStick"
        ]
    }
}
```

Virtual Media resource will be revealed when querying System resource:

```
curl -L http://localhost:8000/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia
{
    "@odata.type": "#VirtualMediaCollection.VirtualMediaCollection",
    "Name": "Virtual Media Services",
    "Description": "Redfish-BMC Virtual Media Service Settings",
    "Members@odata.count": 2,
    "Members": [

        {
            "@odata.id": "/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia/Cd"
        },

        {
            "@odata.id": "/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia/Floppy"
        }

    ],
    "@odata.context": "/redfish/v1/$metadata#VirtualMediaCollection.
↪VirtualMediaCollection",
    "@odata.id": "/redfish/v1/Systems/58893887-8974-2487-2389-841168418919/
↪VirtualMedia",
    "@Redfish.Copyright": "Copyright 2014-2017 Distributed Management Task.
↪Force, Inc. (DMTF). For the full DMTF copyright policy, see http://www.
↪dmtf.org/about/policies/copyright."
}
```

Redfish client can insert a HTTP-based image into the virtual device:

```
curl -d '{"Image": "http://localhost.localdomain/mini.iso", "Inserted":.
↪true}' \
    -H "Content-Type: application/json" \
    -X POST \
    http://localhost:8000/redfish/v1/Systems/58893887-8974-2487-2389-
```

```
↪841168418919/VirtualMedia/Cd/Actions/VirtualMedia.InsertMedia
```

On insert the OpenStack driver will:

- Upload the image directly to glance from the URL (long running)

- Store the URL, image ID and volume ID in server metadata properties *sushy-tools-image-url*, *sushy-tools-import-image*.

- Rebuild the volume with *boot_index: 0* using the image from Glance.

Redfish client can eject image from virtual media device:

```
curl -d '{}' \
     -H "Content-Type: application/json" \
     -X POST \
     http://localhost:8000/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia/Cd/Actions/VirtualMedia.EjectMedia
```

On eject the OpenStack driver will:

- Look up the imported image from instance metadata *sushy-tools-import-image*.

- Delete the imported image.

- Reset the instance metadata.

- Rebuild the server volume with *boot_index: 0* with a "blank" (non-bootable) image. The "blank" image used is defined in the configuration using *SUSHY_EMULATOR_OS_VMEDIA_BLANK_IMAGE* (defaults to: *sushy-tools-blank-image*)

Virtual media boot

> **Note**
>
> With the OpenStack driver the cloud backing the server instances must have support for rebuilding a volume-backed instance with a different image. This was introduced in 26.0.0 (Zed), Nova API microversion 2.93.

To boot a system from a virtual media device, the client first needs to figure out which Manager is responsible for the system of interest:

```
$ curl http://localhost:8000/redfish/v1/Systems/281c2fc3-dd34-439a-9f0f-
↪63df45e2c998
{
...
"Links": {
    "Chassis": [
    ],
    "ManagedBy": [
        {
            "@odata.id": "/redfish/v1/Managers/58893887-8974-2487-2389-
↪841168418919"
```

```
        }
    ]
},
...
```

Exploring the Redfish API links, the client can learn the virtual media devices being offered:

```
$ curl http://localhost:8000/redfish/v1/Systems/58893887-894-2487-2389-
↪841168418919/VirtualMedia
...
"Members": [
{

    "@odata.id": "/redfish/v1/Systems/58893887-8974-2487-2389-841168418919/
↪VirtualMedia/Cd"
},
...
```

Knowing the virtual media device name, the client can check out its present state:

```
$ curl http://localhost:8000/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia/Cd
{

    ...
    "Name": "Virtual CD",
    "MediaTypes": [
        "CD",
        "DVD"
    ],
    "Image": "",
    "ImageName": "",
    "ConnectedVia": "URI",
    "Inserted": false,
    "WriteProtected": false,
    ...
```

Assuming that the *http://localhost/var/tmp/mini.iso* URL points to a bootable UEFI or hybrid ISO, the following Redfish REST API call will insert the image into the virtual CD drive:

```
$ curl -d \
    '{"Image":"http:://localhost/var/tmp/mini.iso", "Inserted": true}' \
    -H "Content-Type: application/json" \
    -X POST \
    http://localhost:8000/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia/Cd/Actions/VirtualMedia.InsertMedia
```

Querying again, the emulator should have it in the drive:

```
$ curl http://localhost:8000/redfish/v1/Systems/58893887-8974-2487-2389-
↪841168418919/VirtualMedia/Cd
{
```

```
    ...
    "Name": "Virtual CD",
    "MediaTypes": [
        "CD",
        "DVD"
    ],
    "Image": "http://localhost/var/tmp/mini.iso",
    "ImageName": "mini.iso",
    "ConnectedVia": "URI",
    "Inserted": true,
    "WriteProtected": true,
    ...
```

Next, the node needs to be configured to boot from its local CD drive over UEFI:

```
$ curl -X PATCH -H 'Content-Type: application/json' \
    -d '{
      "Boot": {
          "BootSourceOverrideTarget": "Cd",
          "BootSourceOverrideMode": "Uefi",
          "BootSourceOverrideEnabled": "Continuous"
      }
    }' \
    http://localhost:8000/redfish/v1/Systems/281c2fc3-dd34-439a-9f0f-
↪63df45e2c998
```

Note

With the OpenStack driver the boot source is changed during insert and eject, so setting *BootSourceOverrideTarget* to *Cd* or *Hdd* has no effect.

By this point the system will boot off the virtual CD drive when powering it on:

```
curl -d '{"ResetType":"On"}' \
    -H "Content-Type: application/json" -X POST \
     http://localhost:8000/redfish/v1/Systems/281c2fc3-dd34-439a-9f0f-
↪63df45e2c998/Actions/ComputerSystem.Reset
```

Note

The ISO files to boot from must be UEFI-bootable. libvirtd should be running on the same machine with sushy-emulator.

Storage resource

For emulating *Storage* resource for a System of choice, the user can statically configure one or more imaginary storage instances along with the corresponding storage controllers which are also imaginary.

The IDs of the imaginary drives associated with a *Storage* resource can be provided as a list under *Drives*.

The *Storage* instances are keyed by the UUIDs of the System they belong to.

```
SUSHY_EMULATOR_STORAGE = {
    "da69abcc-dae0-4913-9a7b-d344043097c0": [
        {
            "Id": "1",
            "Name": "Local Storage Controller",
            "StorageControllers": [
                {
                    "MemberId": "0",
                    "Name": "Contoso Integrated RAID",
                    "SpeedGbps": 12
                }
            ],
            "Drives": [
                "32ADF365C6C1B7BD"
            ]
        }
    ]
}
```

The Storage resources can be revealed by querying the Storage resource for the corresponding System directly.

```
curl http://localhost:8000/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
↪d344043097c0/Storage
{
    "@odata.type": "#StorageCollection.StorageCollection",
    "Name": "Storage Collection",
    "Members@odata.count": 1,
    "Members": [
        {
            "@odata.id": "/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
↪d344043097c0/Storage/1"
        }
    ],
    "Oem": {},
    "@odata.context": "/redfish/v1/$metadata#StorageCollection.
↪StorageCollection",
    "@odata.id": "/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-d344043097c0/
↪Storage"
}
```

Drive resource

For emulating the *Drive* resource, the user can statically configure one or more Drives.

The *Drive* instances are keyed in a composite manner using (System_UUID, Storage_ID), where System_UUID is the UUID of the System and Storage_ID is the ID of the Storage resource to which that particular Drive belongs.

```
SUSHY_EMULATOR_DRIVES = {
    ("da69abcc-dae0-4913-9a7b-d344043097c0", "1"): [
        {
            "Id": "32ADF365C6C1B7BD",
            "Name": "Drive Sample",
            "CapacityBytes": 899527000000,
            "Protocol": "SAS"
        }
    ]
}
```

The *Drive* resource can be revealed by querying it via the System and the Storage resource it belongs to.

```
curl http://localhost:8000/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
→d344043097c0/Storage/1/Drives/32ADF365C6C1B7BD
{
    ...
    "Id": "32ADF365C6C1B7BD",
    "Name": "Drive Sample",
    "Model": "C123",
    "Revision": "100A",
    "CapacityBytes": 899527000000,
    "FailurePredicted": false,
    "Protocol": "SAS",
    "MediaType": "HDD",
    "Manufacturer": "Contoso",
    "SerialNumber": "1234570",
    ...
}
```

Storage Volume resource

The *Volume* resource is emulated as a persistent emulator database record, backed by the libvirt virtualization backend of the dynamic Redfish emulator.

Only the volumes specified in the config file or created via a POST request are allowed to be emulated upon by the emulator and appear as libvirt volumes in the libvirt virtualization backend. Volumes other than these can neither be listed nor deleted.

To allow libvirt volumes to be emulated upon, they need to be specified in the configuration file in the following format (keyed compositely by the System UUID and the Storage ID):

```
SUSHY_EMULATOR_VOLUMES = {
    ('da69abcc-dae0-4913-9a7b-d344043097c0', '1'): [
```

```
        {
            "libvirtPoolName": "sushyPool",
            "libvirtVolName": "testVol",
            "Id": "1",
            "Name": "Sample Volume 1",
            "VolumeType": "Mirrored",
            "CapacityBytes": 23748
        },
        {

            "libvirtPoolName": "sushyPool",
            "libvirtVolName": "testVol1",
            "Id": "2",
            "Name": "Sample Volume 2",
            "VolumeType": "StripedWithParity",
            "CapacityBytes": 48395

        }

    ]

}
```

The Volume resources can be revealed by querying the Volumes resource for the corresponding System and Storage.

```
curl http://localhost:8000/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
↪d344043097c0/Storage/1/Volumes
{
    "@odata.type": "#VolumeCollection.VolumeCollection",
    "Name": "Storage Volume Collection",
    "Members@odata.count": 2,
    "Members": [
        {
            "@odata.id": "/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
↪d344043097c0/Storage/1/Volumes/1"
        },
        {
            "@odata.id": "/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
↪d344043097c0/Storage/1/Volumes/2"
        }
    ],
    "@odata.context": "/redfish/v1/$metadata#VolumeCollection.
↪VolumeCollection",
    "@odata.id": "/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-d344043097c0/
↪Storage/1/Volumes",
}
```

A new volume can also be created in the libvirt backend via a POST request on a Volume Collection:

```
curl -d '{"Name": "SampleVol",\
         "VolumeType": "Mirrored",\
         "CapacityBytes": 74859}' \
    -H "Content-Type: application/json" \
```

```
    -X POST \
    http://localhost:8000/redfish/v1/Systems/da69abcc-dae0-4913-9a7b-
↪d344043097c0/Storage/1/Volumes
```

1.4 Contributing

If you would like to contribute to the development of OpenStack, you must follow the steps in this page:

> http://docs.openstack.org/infra/manual/developers.html

If you already have a good understanding of how the system works and your OpenStack accounts are set up, you can skip to the development workflow section of this documentation to learn how changes to OpenStack should be submitted for review via the Gerrit tool:

> http://docs.openstack.org/infra/manual/developers.html#development-workflow

Pull requests submitted through GitHub will be ignored.

Bugs should be filed on Launchpad, not GitHub:

> https://bugs.launchpad.net/sushy

1.4.1 Cloning the sushy-tools repository

If you haven't already, the sushy-tools source code should be pulled directly from git.

```
# from the directory where you want the source code to reside
git clone https://opendev.org/openstack/sushy-tools
```

1.4.2 Running the emulators locally

Activate the virtual environment and run the emulator of your choice. For instance, to run the dynamic emulator:

```
tox -e venv -- sushy-emulator
```

For more information on running the emulators, refer to the user docs for the *dynamic-emulator* and the *static- emulator*.

INDICES AND TABLES

- genindex

- modindex

- search