



## **Deploying BlueField Software Using ISO with PXE**

# Table of contents

PXE Server Setup

---

PXE Sequence with Redfish

---

BlueField software, including Ubuntu OS, NIC firmware and BMC software, can be deployed using an ISO image similar to the standard Ubuntu ISO deployment method. The BlueField ISO image is based on the standard Ubuntu ISO image for Arm64, with an updated kernel and added DOCA packages.

PXE booting the BlueField device with the ISO image results in Arm OS installation (including DOCA packages), NIC firmware and BMC firmware update (if `BMC_PASSWORD` is provided). This is equivalent to using the corresponding version of bf-bundle BFB.

An Ubuntu ISO image can be used if the RShim interface is not available or if there is an existing deployment system in place that can handle a standard Ubuntu ISO image.

## PXE Server Setup

Mount the ISO:

```
$ mount bf-bundle-2.7.0085-1-2024-06-14-22-36-50.iso /mnt
$ cp /mnt/casper/vmlinuz /var/lib/tftpboot/boot/
$ cp /mnt/casper/initrd /var/lib/tftpboot/boot/
```

Example of `grub.cfg`:

```
menuentry "Install BF OS" {
    linux /boot/vmlinuz autoinstall fsck.mode=skip no-snapd
    console=hvc0 console=ttyAMA0 earlycon=pl011,0x13010000
    net.ifnames=0 biosdevname=0 iommu.passthrough=1 ip=dhcp
    url=http://<HTTP server IP>/jammy/ISO/bf-bundle-2.7.0085-1-2024-06-14-22-36-50.iso
    bfnet=eth0:dhcp bfks=http://<HTTP server IP>/jammy/ISO/bfks
    initrd /boot/initrd
}
```

The `bf.cfg` file can be used to customize the installation procedure. To create `bf.cfg` on the BlueField to be used for the installation use the `bfks` parameter to point to the script located on HTTP server that will create `bf.cfg` file:

bfks example:

```
cat > /etc/bf.cfg << 'EOF'
BMC_PASSWORD="..."
EOF
```

Standard automatic Ubuntu installation using `autoinstall.yaml` is also supported. See [Introduction to autoinstall - Ubuntu installation documentation](#).

Example of `autoinstall.yaml` that can be used to customize the installation and modify `bf.cfg`:

Example of a `grub.cfg` with `autoinstall.yaml`:

```
menuentry "Install BF OS" {
    linux /boot/vmlinuz autoinstall fsck.mode=skip no-snapd
    console=hvc0 console=ttyAMA0 earlycon=pl011,0x13010000
    net.ifnames=0 biosdevname=0 iommu.passthrough=1 ip=dhcp
    url=http://<HTTP server IP>/jammy/ISO/bf-bundle-2.7.0085-1-2024-06-14-22-36-50.iso force-
    ai=http://<HTTP server IP>/jammy/ISO/autoinstall.yaml cloud-config-url=/dev/null
    initrd /boot/initrd
}
```

Example of `autoinstall.yaml`:

```
version: 1

apt:
  preserve_sources_list: false
  conf: |
    Dpkg::Options {
```

```
    "--force-confdef";
    "--force-confold";
};

storage:
  swap:
    size: 0
  grub:
    reorder_uefi: true
  config:
- id: nvme0n1
  type: disk
  ptable: gpt
  path: /dev/nvme0n1
  name: osdisk
  wipe: superblock-recursive

- id: nvme0n1-part1
  type: partition
  device: nvme0n1
  number: 1
  size: 50MB
  flag: boot
  grub_device: true

- id: nvme0n1-part1-fs1
  type: format
  fstype: fat32
  label: efi
  volume: nvme0n1-part1

- id: nvme0n1-part2
  type: partition
  device: nvme0n1
  number: 2
  size: -1
```

```
- id: nvme0n1-part2-fs1
  type: format
  fstype: ext4
  label: root
  volume: nvme0n1-part2

- id: nvme0n1-mount
  type: mount
  path: /
  device: nvme0n1-part2-fs1
  options: defaults
  passno: 0
  fstype: auto

- id: nvme0n1-boot-mount
  type: mount
  path: /boot/efi
  device: nvme0n1-part1-fs1
  options: umask=0077
  passno: 1

reporting:
  builtin:
    type: print

# Add user-data so that subiquity doesn't complain about us not
# having a identity section
user-data:
  debug:
    verbose: true
  write_files:
    - path: /etc/iptables/rules.v4
      permissions: '0644'
      owner: 'root:root'
      content: |
```

```

*mangle
:PREROUTING ACCEPT [45:3582]
:INPUT ACCEPT [45:3582]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [36:4600]
:POSTROUTING ACCEPT [36:4600]
:KUBE-IPTABLES-HINT - [0:0]
:KUBE-KUBELET-CANARY - [0:0]
COMMIT
*filter
:INPUT ACCEPT [41:3374]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [32:3672]
:DOCKER-USER - [0:0]
:KUBE-FIREWALL - [0:0]
:KUBE-KUBELET-CANARY - [0:0]
:LOGGING - [0:0]
:POSTROUTING - [0:0]
:PREROUTING - [0:0]
-A INPUT -j KUBE-FIREWALL
-A INPUT -p tcp -m tcp --dport 111 -j REJECT --reject-
with icmp-port-unreachable
-A INPUT -p udp -m udp --dport 111 -j REJECT --reject-
with icmp-port-unreachable
-A INPUT -i lo -m comment --comment MD_IPTABLES -j ACCEPT
-A INPUT -d 127.0.0.0/8 -m mark --mark 0xb -m comment --
comment MD_IPTABLES -j DROP
-A INPUT -m mark --mark 0xb -m state --state
RELATED,ESTABLISHED -m comment --comment MD_IPTABLES -j ACCEPT
-A INPUT -p tcp -m tcp ! --dport 22 ! --tcp-flags
FIN,SYN,RST,ACK SYN -m mark --mark 0xb -m state --state NEW -m
comment --comment MD_IPTABLES -j DROP
-A INPUT -f -m mark --mark 0xb -m comment --comment
MD_IPTABLES -j DROP
-A INPUT -p tcp -m tcp --tcp-flags
FIN,SYN,RST,PSH,ACK,URG FIN,SYN,RST,PSH,ACK,URG -m mark --mark

```

```

0xb -m comment --comment MD_IPTABLES -j DROP
    -A INPUT -p tcp -m tcp --tcp-flags
FIN,SYN,RST,PSH,ACK,URG NONE -m mark --mark 0xb -m comment --
comment MD_IPTABLES -j DROP
    -A INPUT -m mark --mark 0xb -m state --state INVALID -m
comment --comment MD_IPTABLES -j DROP
    -A INPUT -p tcp -m tcp --tcp-flags RST RST -m mark --mark
0xb -m hashlimit --hashlimit-above 2/sec --hashlimit-burst 2 --
hashlimit-mode srcip --hashlimit-name hashlimit_0 --hashlimit-
htable-expire 30000 -m comment --comment MD_IPTABLES -j DROP
    -A INPUT -p tcp -m mark --mark 0xb -m state --state NEW -
m hashlimit --hashlimit-above 50/sec --hashlimit-burst 50 --
hashlimit-mode srcip --hashlimit-name hashlimit_1 --hashlimit-
htable-expire 30000 -m comment --comment MD_IPTABLES -j DROP
    -A INPUT -p tcp -m mark --mark 0xb -m conntrack --ctstate
NEW -m hashlimit --hashlimit-above 60/sec --hashlimit-burst 20 --
hashlimit-mode srcip --hashlimit-name hashlimit_2 --hashlimit-
htable-expire 30000 -m comment --comment MD_IPTABLES -j DROP
    -A INPUT -m mark --mark 0xb -m recent --rcheck --seconds
86400 --name portscan --mask 255.255.255.255 --rsource -m comment --
comment MD_IPTABLES -j DROP
    -A INPUT -m mark --mark 0xb -m recent --remove --name
portscan --mask 255.255.255.255 --rsource -m comment --comment
MD_IPTABLES
    -A INPUT -p tcp -m tcp --dport 22 -m mark --mark 0xb -m
conntrack --ctstate NEW -m recent --set --name DEFAULT --mask
255.255.255.255 --rsource -m comment --comment MD_IPTABLES
    -A INPUT -p tcp -m tcp --dport 22 -m mark --mark 0xb -m
conntrack --ctstate NEW -m recent --update --seconds 60 --
hitcount 50 --name DEFAULT --mask 255.255.255.255 --rsource -m
comment --comment MD_IPTABLES -j DROP
    -A INPUT -p tcp -m tcp --dport 443 -m mark --mark 0xb -m
conntrack --ctstate NEW -m recent --set --name DEFAULT --mask
255.255.255.255 --rsource -m comment --comment MD_IPTABLES
    -A INPUT -p tcp -m tcp --dport 443 -m mark --mark 0xb -m
conntrack --ctstate NEW -m recent --update --seconds 60 --

```

```

hitcount 10 --name DEFAULT --mask 255.255.255.255 --resource -m
comment --comment MD_IPTABLES -j DROP
    -A INPUT -p udp -m udp --dport 161 -m mark --mark 0xb -m
conntrack --ctstate NEW -m recent --set --name DEFAULT --mask
255.255.255.255 --resource -m comment --comment MD_IPTABLES
    -A INPUT -p udp -m udp --dport 161 -m mark --mark 0xb -m
conntrack --ctstate NEW -m recent --update --seconds 60 --
hitcount 100 --name DEFAULT --mask 255.255.255.255 --resource -m
comment --comment MD_IPTABLES -j DROP
    -A INPUT -p tcp -m tcp --dport 22 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p tcp -m tcp --dport 443 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p tcp -m tcp --dport 179 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p udp -m udp --dport 68 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p udp -m udp --dport 122 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p udp -m udp --dport 161 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p udp -m udp --dport 6306 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p udp -m udp --dport 69 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
    -A INPUT -p udp -m udp --dport 389 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT

```

```
-A INPUT -p tcp -m tcp --dport 389 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --dport 1812:1813 -m mark --mark 0xb
-m conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --dport 49 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p tcp -m tcp --dport 49 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --sport 53 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p tcp -m tcp --sport 53 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --dport 500 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --dport 4500 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --dport 1293 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p tcp -m tcp --dport 1293 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p udp -m udp --dport 1707 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```
-A INPUT -p tcp -m tcp --dport 1707 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
```

```

-A INPUT -i lo -p udp -m udp --dport 3786 -m conntrack --
ctstate NEW,ESTABLISHED -m comment --comment MD_IPTABLES -j
ACCEPT
-A INPUT -i lo -p udp -m udp --dport 33000 -m conntrack --
ctstate NEW,ESTABLISHED -m comment --comment MD_IPTABLES -j
ACCEPT
-A INPUT -p icmp -m mark --mark 0xb -m comment --comment
MD_IPTABLES -j ACCEPT
-A INPUT -p udp -m udp --sport 5353 --dport 5353 -m mark --
mark 0xb -m conntrack --ctstate NEW,ESTABLISHED -m comment --
comment MD_IPTABLES -j ACCEPT
-A INPUT -p udp -m udp --dport 33434:33523 -m mark --mark
0xb -m comment --comment MD_IPTABLES -j REJECT --reject-with
icmp-port-unreachable
-A INPUT -p udp -m udp --dport 123 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
-A INPUT -p udp -m udp --dport 514 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
-A INPUT -p udp -m udp --dport 67 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
MD_IPTABLES -j ACCEPT
-A INPUT -p tcp -m tcp --dport 60102 -m mark --mark 0xb -m
conntrack --ctstate NEW,ESTABLISHED -m comment --comment
"MD_IPTABLES: Feature HA port" -j ACCEPT
-A INPUT -m mark --mark 0xb -m comment --comment
MD_IPTABLES -j LOGGING
-A FORWARD -j DOCKER-USER
-A OUTPUT -o oob_net0 -m comment --comment MD_IPTABLES -j
ACCEPT
-A DOCKER-USER -j RETURN
-A LOGGING -m mark --mark 0xb -m comment --comment
MD_IPTABLES -j NFLOG --nflog-prefix "IPTables-Dropped:" --nflog-group
3

```

```

        -A LOGGING -m mark --mark 0xb -m comment --comment
MD_IPTABLES -j DROP
        -A PREROUTING -i oob_net0 -m comment --comment
MD_IPTABLES -j MARK --set-xmark 0xb/0xffffffff
        -A PREROUTING -p tcp -m tcpmss ! --mss 536:65535 -m tcp ! -
-dport 22 -m mark --mark 0xb -m conntrack --ctstate NEW -m comment
--comment MD_IPTABLES -j DROP
    COMMIT
    *nat
    :PREROUTING ACCEPT [1:320]
    :INPUT ACCEPT [1:320]
    :OUTPUT ACCEPT [8:556]
    :POSTROUTING ACCEPT [8:556]
    :KUBE-KUBELET-CANARY - [0:0]
    :KUBE-MARK-DROP - [0:0]
    :KUBE-MARK-MASQ - [0:0]
    :KUBE-POSTROUTING - [0:0]
        -A POSTROUTING -m comment --comment "kubernetes postrouting rules" -
j KUBE-POSTROUTING
        -A KUBE-MARK-DROP -j MARK --set-xmark 0x8000/0x8000
        -A KUBE-MARK-MASQ -j MARK --set-xmark 0x4000/0x4000
        -A KUBE-POSTROUTING -m mark ! --mark 0x4000/0x4000 -j RETURN
        -A KUBE-POSTROUTING -j MARK --set-xmark 0x4000/0x0
        -A KUBE-POSTROUTING -m comment --comment "kubernetes service traffic
requiring SNAT" -j MASQUERADE --random-fully
    COMMIT
    users:
        - name: ubuntu
          lock_passwd: False
          groups: adm, audio, cdrom, dialout, dip, floppy, lxd,
netdev, plugdev, sudo, video
          sudo: ALL=(ALL) NOPASSWD:ALL
          shell: /bin/bash
          plain_text_passwd: 'ubuntu'
    chpasswd:
        list: |

```

```

    ubuntu:ubuntu
    expire: True
no_ssh_fingerprints: true
runcmd:
  - [ /usr/sbin/netfilter-persistent, start ]
  - [
/opt/mellanox/doca/services/telemetry/import_doca_telemetry.sh ]
  - [ /usr/bin/bfrshlog, "INFO: DPU is ready" ]

late-commands:
# write release file
- |
cat << EOF > /target/etc/bf-release
BF_NAME="Mellanox Bluefield"
BF_PRETTY_NAME="Mellanox Bluefield"
BF_SWBUILD_TIMESTAMP="2024-06-12-12-47-25"
BF_SWBUILD_VERSION="2.7.0085-1"
BF_COMMIT_ID="7fce146"
BF_PLATFORM="BlueField SoC"
BF_SERIAL_NUMBER="1332723060006"
EOF

# mount cdrom
- mkdir -p /target/tmp/cdrom
- mount --bind /cdrom /target/tmp/cdrom || true
- |
cat << EOF > /target/etc/apt/sources.list
deb [check-date=no] file:///tmp/cdrom/jammy main restricted
EOF

# avoid running flash kernel after install kernel
- mkdir -p /target/run/systemd
- echo docker > /target/run/systemd/container

# Install packages
- curtin in-target -- apt update -y

```

```
- curtin in-target -- apt remove -y --purge `dpkg --get-selections | grep openipmi | awk '{print $2}'`
```

```
- curtin in-target -- /bin/bash -c "DEBIAN_FRONTEND=noninteractive RUN_FW_UPDATER=no apt-get install --no-install-recommends -y acpid bc binutils bridge-utils build-essential cracklib-runtime dc docker.io flash-kernel i2c-tools ifenslave iperf3 iptables-persistent iputils-arping iputils-ping iputils-tracepath kexec-tools libpam-pwquality libssl-dev lldpad lm-sensors net-tools network-manager nfs-common nvme-cli openssh-server python3.10 python3-pyinotify python3-pip rasdaemon rsync sbsigntool shim-signed tcpdump watchdog doca-runtime doca-devel containerd kubelet runc nv-common-apis nvidia-repo-keys linux-bluefield-modules-bluefield linux-image-5.15.0-1042-bluefield"
```

```
# rewrite sources
```

```
- |  
cat << EOF > /target/etc/apt/sources.list  
deb http://ports.ubuntu.com/ubuntu-ports/ jammy main restricted universe multiverse  
deb http://ports.ubuntu.com/ubuntu-ports/ jammy-updates main restricted universe  
multiverse  
deb http://ports.ubuntu.com/ubuntu-ports/ jammy-security main restricted universe multiverse  
EOF
```

```
# Allow cloud-init to configure networking
```

```
- find /target/etc/cloud/cloud.cfg.d/ -type f ! -name README !  
-name 05_logging.cfg ! -name 90_dpkg.cfg -delete || true;  
- curtin in-target -- cloud-init clean
```

```
# Post-installation steps
```

```
# Create bf.cfg
```

```
- |  
cat << EOF > /target/etc/bf.cfg  
# UPDATE_ATF_UEFI - Updated ATF/UEFI (Default: yes)  
# Relevant for PXE installation only as while using RSHIM interface  
ATF/UEFI  
# will always be updated using capsule method  
UPDATE_ATF_UEFI="yes"
```

```
#####  
# BMC Component Update
```

```
#####  
# BMC_USER - User name to be used to access BMC (Default:  
root)  
BMC_USER="root"  
  
# BMC_PASSWORD - Password used by the BMC user to access BMC  
(Default: None)  
BMC_PASSWORD=""  
  
# BMC_IP_TIMEOUT - Maximum time in seconds to wait for the  
connection to the  
# BMC to be established (Default: 600)  
BMC_IP_TIMEOUT=600  
  
# BMC_TASK_TIMEOUT - Maximum time in seconds to wait for BMC  
task (BMC/CEC  
# Firmware update) to complete (Default: 1800)  
BMC_TASK_TIMEOUT=1800  
  
# UPDATE_BMC_FW - Update BMC firmware (Default: yes)  
UPDATE_BMC_FW="yes"  
  
# BMC_REBOOT - Reboot BMC after BMC firmware update to apply  
the new version  
# (Default: no). Note that the BMC reboot will reset the BMC  
console.  
BMC_REBOOT="no"  
  
# UPDATE_CEC_FW - Update CEC firmware (Default: yes)  
UPDATE_CEC_FW="yes"  
  
# UPDATE_DPU_GOLDEN_IMAGE - Update BlueField Golden Image  
(Default: yes)  
UPDATE_DPU_GOLDEN_IMAGE="yes"
```

```

# UPDATE_NIC_FW_GOLDEN_IMAGE- Update NIC firmware Golden
Image (Default: yes)
UPDATE_NIC_FW_GOLDEN_IMAGE="yes"

# pre_bmc_components_update - Shell function called by BFB's
install.sh before
# updating BMC components (no communication to the BMC is
established at this
# point)

# post_bmc_components_update - Shell function called by BFB's
install.sh after
# updating BMC components

#####
# NIC Firmware update

#####
# WITH_NIC_FW_UPDATE - Update NIC Firmware (Default: no)
WITH_NIC_FW_UPDATE="yes"
EOF

# Run post-installation script to update ATF/UEFI, NIC
firmware and BMC components
- curtin in-target -- /bin/bash -c "device=/dev/nvme0n1 /usr/local/sbin/bfiso-
post-install.sh || true"

- curtin in-target -- systemctl disable snapd

```

## PXE Sequence with Redfish

HTTP boot configuration can be done using the BlueField BMC's Redfish interface.

ISO upgrade via Redfish to set UEFI HTTPs/PXE boot by setting UEFI first boot source.

To set the UEFI first boot source using Redfish:

1. Follow the instructions under section "[Deploying BlueField Software Using BFB with PXE](#)".
2. Check the current boot override settings by performing a GET on the `ComputerSystem` schema over 1GbE to the BlueField BMC. Look for the `"Boot"` property.

```
curl -k -X GET -u root:<password> https://<BF-BMC-IP>/redfish/v1/Systems/<SystemID>/ | python3 -m json.tool
{
  ...
  "Boot": {
    "BootNext": "",
    "BootOrderPropertySelection": "BootOrder",
    "BootSourceOverrideEnabled": "Disabled",
    "BootSourceOverrideMode": "UEFI",
    "BootSourceOverrideTarget": "None",
    "UefiTargetBootSourceOverride": "None",
    .....
  },
  ....
  "BootSourceOverrideEnabled@Redfish.AllowableValues": [
    "Once",
    "Continuous",
    "Disabled"
  ],
  "BootSourceOverrideTarget@Redfish.AllowableValues": [
    "None",
    "Pxe",
    "UefiHttp",
    "UefiShell",
    "UefiTarget",
    "UefiBootNext"
  ],

```

```
.....  
}
```

### Info

Boot override enables overriding the first boot source, either once or continuously.

3. The example output above shows the `BootSourceOverrideEnabled` property is `Disabled` and `BootSourceOverrideTarget` is `None`. The `BootSourceOverrideMode` property should always be set to `UEFI`. Allowable values of `BootSourceOverrideEnabled` and `BootSourceOverrideTarget` are defined in the metadata (`BootSourceOverrideEnabled@Redfish.AllowableValues` and `BootSourceOverrideTarget@Redfish.AllowableValues` respectively).
4. If `BootSourceOverrideEnabled` is set to `Once`, then boot override is disabled after the first boot, and any related properties are reset to their former values to avoid repetition. If it is set to `Continuous`, then on every reboot the BlueField keeps performing boot override (HTTPBoot).
5. To perform boot override, perform a PATCH to pending settings URI over the 1GbE to the BlueField BMC.

```
curl -k -X PATCH -d '{"Boot":  
{"BootSourceOverrideEnabled": "Once",  
"BootSourceOverrideMode": "UEFI", "BootSourceOverrideTarget":  
"UefiHttp", "HttpBootUri": "http://<HTTP-Server-  
Ip>/Image.iso"}}' -u root:<password> https://<BF-BMC-  
IP>/redfish/v1/Systems/<SystemID>/Settings | python3 -m  
json.tool
```

For example:

```
curl -k -X GET -u root:<password> https://<BF-BMC-IP>/redfish/v1/Systems/<SystemID>/ | python3 -m json.tool
{
  ...
  "Boot": {
    "BootNext": "",
    "BootOrderPropertySelection": "BootOrder",
    "BootSourceOverrideEnabled": "Once",
    "BootSourceOverrideMode": "UEFI",
    "BootSourceOverrideTarget": "UefiHttp",
    "UefiTargetBootSourceOverride": "None",
    .....
  },
  .....
}
```

6. After performing the above PATCH successfully, reboot the BlueField using the Redfish Manager schema over the 1GbE to the BlueField BMC:

```
curl -k -u root:<password> -H "Content-Type: application/json" -X POST https://<BF-BMC-IP>/redfish/v1/Systems/Bluefield/Actions/ComputerSystem.Reset -d '{"ResetType" : "GracefulRestart"}'
```

7. Once UEFI has completed, check whether the settings are applied by performing a GET on `ComputerSystem` schema over the 1GbE OOB to the BlueField BMC.

 **Note**

The `HttpBootUri` property is parsed by the Redfish server and the URI is presented to the BlueField as part of DHCP lease when the BlueField performs the HTTP boot.

## Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF

ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

## **Trademarks**

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2025, NVIDIA. PDF Generated on 11/20/2025