



## Network

# Table of contents

BlueField Host Network Interface	3
OOB Network 3-Port Switch Control	7

This section contains the following pages:

- [BlueField Host Network Interface](#)
- [OOB Network 3-Port Switch Control](#)

---

# BlueField Host Network Interface

Under URI `redfish/v1/Systems/Bluefield`, there is a collection called `EthernetInterfaces` representing the data ports and the OOB port of the BlueField. It is read-only and contains network information (e.g., IP addresses, MAC addresses).

## Note

These abilities are supported when operating in DPU mode only.

## Displaying Network Interfaces Collection

```
curl -k -u root:'<password>' -H 'Content-Type: application/json'
-X GET
https://<bmc_ip>/redfish/v1/Systems/Bluefield/EthernetInterfaces
{
  "@odata.id" :
  "/redfish/v1/Systems/Bluefield/EthernetInterfaces",
  "@odata.type" :
  "#EthernetInterfaceCollection.EthernetInterfaceCollection",
  "Description": "Collection of EthernetInterfaces of the host",
  "Members": [
    {
      "@odata.id" :
      "/redfish/v1/Systems/Bluefield/EthernetInterfaces/eth0"
    },
    {
```

```

    "@odata.id" :
"/redfish/v1/Systems/Bluefield/EthernetInterfaces/oob0"
  }
],
"Members@odata.count" : 2,
"Name" : "Ethernet Network Interface Collection"
}

```

## Displaying Network Interface Object

### Info

The interface object has a field called `LinkStatus` which is determined by the following rules:

- If the interface is the OOB port (i.e., `oob_net0`), `LinkStatus` would display `LinkUp` if the port is configured up using `ifconfig/ip` command.
- If the interface is a data port (i.e., `eth0/eth1` or `ib0/ib1`), `LinkStatus` would display `NoLink` if no QSFP cable is connected. If a QSFP transceiver is connected, the link would appear as `LinkUp` if the port is configured as up using the `ifconfig/ip` commands. If not, it displays `LinkDown`.

```

curl -k -u root:'<password>' -H 'Content-Type: application/json'
-X GET
https://<bmc_ip>/redfish/v1/Systems/Bluefield/EthernetInterfaces/oob0
{
  "@odata.id" :
"/redfish/v1/Systems/Bluefield/EthernetInterfaces/oob0",
  "@odata.type" : "#EthernetInterface.v1_6_0.EthernetInterface",

```

```
"DHCPv4": {
  "DHCPEnabled": true,
  "UseDNSServers": false,
  "UseDomainName": false,
  "UseNTPServers": false
},
"DHCPv6": {
  "OperatingMode": "Stateful",
  "UseDNSServers": false,
  "UseDomainName": false,
  "UseNTPServers": false
},
"Description": "Host Network Interface for port oob0",
"IPv4Addresses": [
  {
    "Address": "10.345.41.97",
    "AddressOrigin": "Static",
    "Gateway": "0.0.0.0",
    "SubnetMask": "255.255.240.0"
  }
],
"IPv4StaticAddresses": [
  {
    "Address": "10.345.41.97",
    "AddressOrigin": "Static",
    "Gateway": "0.0.0.0",
    "SubnetMask": "255.255.240.0"
  }
],
"IPv6AddressPolicyTable": [],
"IPv6Addresses": [
  {
    "Address": "fe80::a278:c2ff:fe0e:87a4",
    "AddressOrigin": "Static",
    "AddressState": null,
    "PrefixLength": 64
  }
]
```

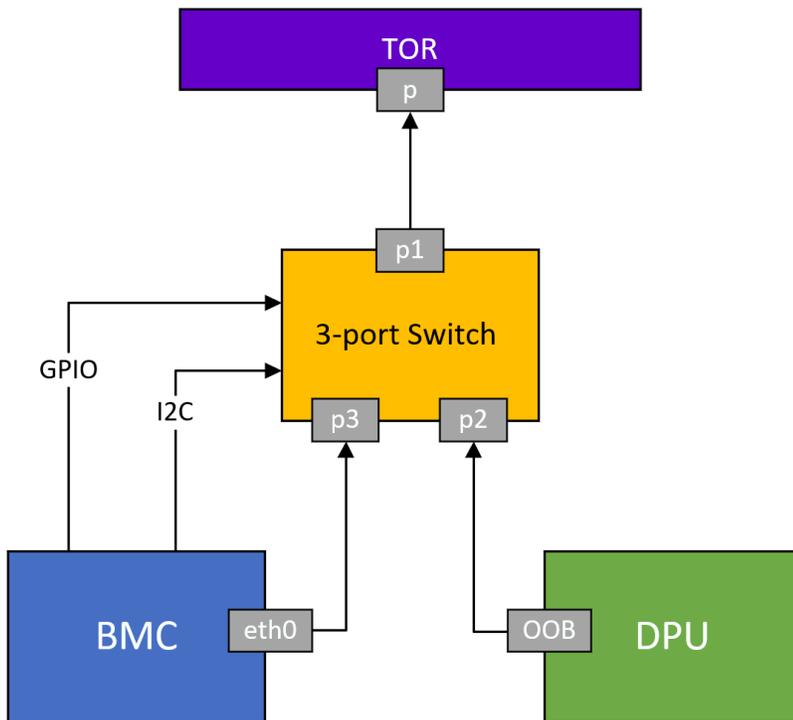
```
    }
  ],
  "IPv6DefaultGateway": "0:0:0:0:0:0:0:0",
  "IPv6StaticAddresses": [
    {
      "Address": "fe80::a278:c2ff:fe0e:87a4",
      "PrefixLength": 64
    }
  ],
  "Id": "oob0",
  "InterfaceEnabled": true,
  "LinkStatus": "LinkUp",
  "MACAddress": "a0:88:a2:0e:87:a4",
  "MTUSize": 1500,
  "Name": "Host Ethernet Interface",
  "NameServers": [],
  "SpeedMbps": 1000,
  "StaticNameServers": [],
  "Status": {
    "State": "Disabled"
  }
}
```

**i Note**

If the user changed the BlueField's IP information dynamically, rebooting the BMC should show the updated IP info.

# OOB Network 3-Port Switch Control

To enable both the BMC and the Arm on NVIDIA® BlueField® to access the out-of-band (OOB) network management interface, an L2, 3-port switch has been incorporated into the system. This switch acts as a bridge, connecting the RJ45 port (OOB), the BMC, and the Arm in the DPU. It is important to note that the switch is exclusively managed by the DPU's BMC through a dedicated I2C line and a GPIO signal that controls the switch's reset function.



## 3-Port Switch IPMI Commands

netfunc	cmd	data	Description
0x32	0x97	N/A	Get 3-port switch ports mode. On success, it returns: <ul style="list-style-type: none"><li>0x00 – all ports are allowed access to RJ45</li></ul>

netfunc	cmd	data	Description
			<ul style="list-style-type: none"> <li>0x01 – only BMC is allowed access to RJ45</li> </ul>
0x32	0x98	<ul style="list-style-type: none"> <li>0x00 – all ports are allowed access to RJ45</li> <li>0x01 – only BMC is allowed access to RJ45</li> </ul>	Set 3-port switch ports mode. <ul style="list-style-type: none"> <li>Setting this command is only possible while the user is logged on to the BMC, this command is not supported over the network interfaces (IPMI nor Redfish)</li> <li>Setting is persistent across power cycle and switch reset command</li> </ul>
0x32	0xA1	0x3	Reset on-board 3-port switch

### Info

In all these use cases, the internal pathway connecting the DPU and the BMC remains operational. This enables communication between the BMC and the DPU over the internal network.

Example for disabling the OOB network of the DPU Arm:

```
#bmc> ipmitool raw 0x32 0x98 0x1
```

## 3-Port Switch Redfish Commands

### Getting 3-port Switch Ports Mode

```
curl -k -u root:<password> -H 'Content-Type: application/json' -X GET
https://<bmc_ip>/redfish/v1/Systems/Bluefield/Oem/Nvidia/Switch
```

Example output:

```
{
  "LinkStatus": {
    "BMC": "LinkUp",
    "DPU": "LinkUp",
    "RJ45": "LinkUp"
  },
  "TorSwitchMode": {
    "BmcOobEnabled": true,
    "DpuOobEnabled": true
  }
}
```

Where:

- `LinkStatus` – displays the link status of each port on the 3-port switch
  - For the `RJ45` and `DPU` ports, the link status is taken from their `PHY Basic Status Register` and from `PHY Basic Control Register` for the port's power down status on the 3-port switch
  - The BMC link is considered always `LinkUp`
- `TorSwitchMode`:
  - `BmcOobEnabled` – if `true`; enables the BMC to access the out-of-band network
  - `DpuOobEnabled` – if `true`; enables the BlueField to access the out-of-band network

## Setting 3-port Switch Port Mode

```
curl -k -u root:'<password>' -H 'Content-Type: application/json' -X PATCH -d
'{"TorSwitchMode": {"BmcOobEnabled": <Port State>, "DpuOobEnabled": <Port State>}}'
https://<bmc_ip>/redfish/v1/Systems/Bluefield/Oem/Nvidia/Switch
```

Where `Port State`:

- True – Enable the port to access the out-of-band network
- False – Disable the port to access the out-of-band network

### **Note**

The internal pathway connecting the BMC and RJ45 is not allowed to be disabled using a Redfish command. Therefore, the parameter `BmcOobEnabled` should be set as `true` when setting 3-port switch ports mode, otherwise the Redfish command would return an error.

### **Note**

For the patch command request, both `BmcOobEnabled` and `DpuOobEnabled` must be set.

The following is an example of how to set only BMC is allowed access to RJ45:

```
curl -k -u root:'<password>' -H 'Content-Type: application/json' -X PATCH -d
'{"TorSwitchMode": {"BmcOobEnabled": true, "DpuOobEnabled": false}}'
https://<bmc_ip>/redfish/v1/Systems/Bluefield/Oem/Nvidia/Switch
```

Example output:

```
{
  "@Message.ExtendedInfo": [
    {
      "@odata.type": "#Message.v1_1_1.Message",
      "Message": "The request completed successfully.",
      "MessageArgs": [ ],
      "MessageId": "Base.1.15.0.Success",
      "MessageSeverity": "OK",
      "Resolution": "None"
    }
  ]
}
```

## Resetting On-board 3-port Switch

```
curl -k -u root:'<password>' -H 'Content-Type: application/json' -X POST
https://<bmc_ip>/redfish/v1/Systems/Bluefield/0em/Nvidia/Switch.Res
```

Example output:

```
{
  "@Message.ExtendedInfo": [
    {
      "@odata.type": "#Message.v1_1_1.Message",
      "Message": "The request completed successfully.",
      "MessageArgs": [ ],
      "MessageId": "Base.1.15.0.Success",
      "MessageSeverity": "OK",
      "Resolution": "None"
    }
  ]
}
```



```
}  
  ]  
}
```

## 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 NON-INFRINGEMENT, 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. PDF Generated on 12/15/2025