

R3 Installation Document of IEC Type 5: SmartNIC for Integrated Edge Cloud (IEC) Blueprint Family

- [Introduction](#)
- [License](#)
- [How to use this document](#)
- [Deployment Architecture](#)
- [blocked URL](#)
- [Pre-Installation Requirements](#)
 - [Hardware Requirements](#)
 - [Minimum Hardware Requirements](#)
 - [Recommended Hardware Requirements](#)
 - [Software Prerequisites](#)
 - [Database Prerequisites](#)
 - [Schema scripts](#)
 - [Other Installation Requirements](#)
 - [Jump Host Requirements](#)
 - [Network Requirements](#)
 - [Bare Metal Node Requirements](#)
 - [Execution Requirements \(Bare Metal Only\)](#)
- [Installation High-Level Overview](#)
 - [SmartNIC Server Preparations](#)
 - [Basic SmartNIC Management Operations](#)
 - [SmartNIC Modes](#)
 - [Bare Metal Deployment Guide](#)
 - [Install Bare Metal Jump Host](#)
 - [Creating a Node Inventory File](#)
 - [Creating the Settings Files](#)
 - [Running](#)
 - [Virtual Deployment Guide](#)
 - [Standard Deployment Overview](#)
 - [Snapshot Deployment Overview](#)
 - [Special Requirements for Virtual Deployments](#)
 - [Install Jump Host](#)
 - [Verifying the Setup - VMs](#)
 - [Upstream Deployment Guide](#)
 - [Upstream Deployment Key Features](#)
 - [Special Requirements for Upstream Deployments](#)
 - [Scenarios and Deploy Settings for Upstream Deployments](#)
 - [Including Upstream Patches with Deployment](#)
 - [Running](#)
 - [Interacting with Containerized Overcloud](#)
- [Verifying the Setup as defined the Akraio validation feature project plus any additional testing specific to the blue print](#)
- [Developer Guide and Troubleshooting](#)
 - [Utilization of Images](#)
 - [Post-deployment Configuration](#)
 - [Debugging Failures](#)
 - [Reporting a Bug](#)
- [Uninstall Guide](#)
- [Troubleshooting](#)
 - [Error Message Guide](#)
 - [Server Unable to Find the SmartNIC](#)
 - [The SmartNIC no Longer Works](#)
 - [SmartNIC Stopped Working After Installing Another](#)
 - [Link Indicator Light is Off](#)
 - [Link Light is On but No Communication is Established](#)
- [Maintenance](#)
 - [Blue Print Package Maintenance](#)
 - [Software maintenance](#)
 - [Hardware maintenance](#)
 - [Blueprint Deployment Maintenance](#)
- [Frequently Asked Questions](#)
- [License](#)
- [References](#)
- [Definitions, acronyms and abbreviations](#)

Introduction

BlueField's Ethernet SmartNIC solutions combine Arm® processing power with advanced network offloads to accelerate a multitude of security, networking and storage applications.

Based on the BlueField system-on-chip (SoC) programmable network engine, BlueField SmartNICs integrate Mellanox ConnectX®-5 network controller with Arm multicore processors and cryptology options to deliver an innovative and high-performance programmable platform. Through this article, you will be able to quickly set up a simple test lab which will allow you to experiment and explore BlueField SmartNIC capabilities. The setup includes a bare metal server hosting BlueField 25GbE SmartNIC connected back-to-back to another bare metal server with ConnectX-5 NIC.

License

N/A

How to use this document

this document is used to guide how to install the OVS-DPDK into SmartNIC and integrate with the server

Deployment Architecture

The following diagram shows the physical configuration between the BlueField SmartNIC and the ConnectX-5 NIC.

[blocked URL](#)

Pre-Installation Requirements

Hardware Requirements

Minimum Hardware Requirements

Qty.	Item	Comments
2	x86 servers running CentOS 7.5	-
1	ConnectX-5 NIC (FW: 16.24.0328 and up)	-
1	BlueField 25GbE SmartNIC with Arm processor (FW: 18.24.1000 and up)	-
1	25GB/s SFP28 networking cable	-
1	QSA28 Adapter (100GB-to-25GB Adapter)	-
1	UART-to-USB cable	For SmartNIC USB console access
1	Mini USB Type B cable	For SmartNIC USB Console Access
1	bfb image to be loaded on the BlueField SmartNIC Arm.	Image contains: <ul style="list-style-type: none">• CentOS 7.5 over Arm architecture• MLNX_OFED_LINUX-4.4-2.5.11.0 (compiled with DPDK and BlueField flags)• DPDK libraries version 18.11• Pre-configured network settings for easy SmartNIC management

1. Insert ConnectX-5 NIC into **Server A**
2. Insert BlueField SmartNIC into **Server B**
3. Connect the SmartNIC UART cable from SmartNIC to its hosting **Server B** USB port #1
4. Connect Mini USB cable from SmartNIC to its hosting **Server B** USB port #2
5. Insert QSA28 Adapter into **Server B** 100GB port
6. Connect one of the SmartNIC ports to the ConnectX-5 using 25GB cable
7. Connect both servers to a management LAN

Recommended Hardware Requirements

N/A

Software Prerequisites

MLNX_OFED

Note

1. Use automatic FM update during MLNX_OFED installation ONLY if the current BlueField FW version is 18.24.1000 (GA) or above. In case its older, manually burn the FW after MLNX_OFED installation.

2. In case you plan to run DPDK on the host itself, install MLNX_OFED with "--dpdk --upstream-libs" flags
3. The "BlueField" flag is not required in some of the MLNX_OFED versions and OS releases

OVS-DPDK

1. We've open-sourced a fork of DPDK based Open vSwitch, which supports the VXLAN decap/encap full offload: <https://github.com/bytedance/ovs-dpdk>
2. The code is fully compatible with upstream Open vSwitch, you can follow the same procedures to build them from the source. See the link: <http://docs.openvswitch.org/en/latest/intro/install/>
3. Or you can use the ansible scripts in the Akaino Gerrit to install both BlueField Card and OVS-DPDK. Currently, the scripts are not available, as the scripts work by relying on many company's internal services. We are refactoring the code that you can use it in your own SmartNIC setups.

Database Prerequisites

N/A

Schema scripts

N/A

Other Installation Requirements

N/A

Jump Host Requirements

N/A

Network Requirements

N/A

Bare Metal Node Requirements

N/A

Execution Requirements (Bare Metal Only)

N/A

Installation High-Level Overview

SmartNIC Server Preparations

- Install useful packages:

```
yum install -y screen minicom net-tools tcpdump usbutils net-utils sshpass iptables-services fio epel-release iperf wget git unzip perl pciutils lsof tcl libxml2-python tk unbound gtk2 atk cairo gcc-gfortran tcsh nvme-cli
```

- Install MLNX_OFED with the mentioned flags below.

```
tar -xzvf MLNX_OFED_LINUX-4.4-2.5.11.0-rhel7.5-x86_64.tgz
cd MLNX_OFED_LINUX-4.4-2.5.11.0-rhel7.5-x86_64
./mlnxofedinstall --with-nvme --bluefield --without-fw-update
```

Note

1. Use an automatic FW update during MLNX_OFED installation ONLY if the current BlueField FW version is 18.24.1000 (GA) or above. In case its older, manually burn the FW after MLNX_OFED installation.
2. In case you plan to run DPDK on the host itself, install MLNX_OFED with "--dpdk --upstream-libs" flags
3. The "BlueField" flag is not required in some of the MLNX_OFED versions and OS releases

- Enable SR-IOV on FW (Optional):

```
mlxconfig -y -d /dev/mst/mt41682_pciconf0 s SRIOV_EN=1
mlxconfig -y -d /dev/mst/mt41682_pciconf0 s NUM_OF_VFS=8
```

- Uncomment the following line in `/etc/modprobe.d/rshim.conf` to set the default RShim access method to the SmartNIC:

```
options rshim backend_driver=rshim_usb
```

Note

1. When this line is commented, the first available driver will be used.
2. When the mini USB cable is not used, the driver should be `rshim_pcie`

- Create a file named `ifcfg-tmfifo_net0` under `/etc/sysconfig/network-scripts/` directory, with the following content:

```
NAME="tmfifo_net0"
HWADDR="00:1a:ca:ff:ff:02"
ONBOOT=yes
NETBOOT=yes
IPV6INIT=yes
BOOTPROTO=static
TYPE=Ethernet
IPADDR=192.168.100.1
NETMASK=255.255.255.0
```

- `ifcfg-tmfifo_net0` is the network interface used for IP communication between the x86 host and the SmartNIC with `rshim_net` driver.
- Enable ip forwarding and set a NAT rule to allow the SmartNIC OS to communicate with the outer world via the host management interface:

```
sysctl -w net.ipv4.ip_forward=1
systemctl enable iptables
systemctl start iptables
iptables -F
iptables -t nat -A POSTROUTING -o <host_mgmt_interfac> -j MASQUERADE
service iptables save
```

- Disable Firewall and Network Manager services

```
systemctl disable firewalld.service
systemctl disable NetworkManager
```

- Reboot the server. Once the server is up again, run some configuration sanity tests:
- To verify BlueField FW version:

```
mlxfwmanager --query
```

- To verify RShim modules are loaded:

```
lsmod | grep -i rshim
```

- To identify the backend active RShim driver:

```
cat /dev/rshim0/misc
```

- To verify the NAT rule for SmartNIC communication with the outer world:

```
iptables -L -n -t nat
```

- Download the BlueField Quick Start bfb image: <https://mellanox.box.com/s/q44f501xq1e29yb1bjpsb9vzqqs89a76>
- Place the bfb image on the host and load it on the BlueField SmartNIC using the following command:

```
cat /home/CentOS7.5ALT-aarch64-QuickStart.bfb > /dev/rshim0/boot
```

Use one of the console access methods described in the "SmartNIC Management Methods" section below to follow the image installation process.

When bfb image installation process is completed, connect to the SmartNIC with the console or SSH, as described below, and test connectivity from the SmartNIC OS to the outer world.

- SmartNIC Management Methods

To manage the SmartNIC from the hosting server, use one of the following methods

[blocked URL](#)

Access Credentials: root|centos

Console over USB-Serial / Serial Port properties: Rate 115200, HW Flow Control OFF.

```
minicom -D /dev/ttyUSB0
```

Console over USB (rshim_usb) / Serial Port properties: Rate 115200, HW Flow Control OFF

```
minicom -D /dev/rshim0/console
```

SSH (rshim_net)

```
ssh root@192.168.100.2
```

Note

SmartNIC OS is configured to obtain internet access via its hosting server.

Basic SmartNIC Management Operations

How-to soft reset the SmartNIC via host server:

Issue the following command on the server:

```
echo "SW_RESET 1" > /dev/rshim0/misc
```

Follow the boot process using console access:

```
minicom -D /dev/ttyUSB0
```

Note

This is not a power cycle of the SmartNIC.

How-to load BlueField SmartNIC OS image from the hosting server:

Place bfb image on the server

Issue the following command on the server:

```
cat <image>.bfb > /dev/rshim0/boot
```

Follow the installation and boot process using the console access:

```
minicom -D /dev/ttyUSB0
```

SmartNIC Modes

BlueField SmartNIC can operate in two modes:

SEPERATED HOST (default mode)

In this mode, both the x86 hosting server and the SmartNIC Arm OS act as separated entities; Both can use and communicate with each other or the network via the ConnectX-5 module of the SmartNIC.

EMBEDDED

In this mode, the x86 hosting server communicates with the outer world only through the SmartNIC Arm.

[blocked URL](#)

Note

In order to identify the SmartNIC current mode, issue the following commands on the x86 hosting server:

```
x86_host#mst start
x86_host#mlxconfig -d /dev/mst/mt41682_pciconf0 q | grep -i model
```

Switching to EMBEDDED Mode from SEPERATED_HOST Mode

On the x86 host:

Enable EMBEDDED mode:

```
x86_host #mst start
x86_host #mlxconfig -d /dev/mst/mt41682_pciconf0 s INTERNAL_CPU_MODEL=1
```

Reboot.

Verify:

```
x86_host#mst start
x86_host#mlxconfig -d /dev/mst/mt41682_pciconf0 q | grep -i model
```

On the Arm:

Issue the following commands to enable ECPF parameters (applied per port):

Note

The ECPF parameters will be available for setting only after completing the previous steps.

```
Snic#mst start
Snic#mst status -v
Snic#mlxconfig -d /dev/mst/mt41682_pciconf0 s ECPF_ESWITCH_MANAGER=1 ECPF_PAGE_SUPPLIER=1
Snic#mlxconfig -d /dev/mst/mt41682_pciconf0.1 s ECPF_ESWITCH_MANAGER=1 ECPF_PAGE_SUPPLIER=1
```

Reboot the x86 host.

Verify configuration:

Once the procedure is completed, issue the following command on both the x86 host and the Arm:

```
/opt/mellanox/iproute2/sbin/rdma link
```

On x86 host: the output should list two network devices.

```
1/1: mlx5_0/1: state ACTIVE physical_state LINK_UP netdev ens1f0
2/1: mlx5_1/1: state DOWN physical_state DISABLED netdev ens1f1
```

On Arm: the output should list four representors:

```
1/1: mlx5_0/1: state ACTIVE physical_state LINK_UP netdev rep0-ffff
2/1: mlx5_1/1: state ACTIVE physical_state LINK_UP netdev rep0-0
3/1: mlx5_2/1: state DOWN physical_state DISABLED netdev repl-ffff
4/1: mlx5_3/1: state DOWN physical_state DISABLED netdev repl-0
```

Issue the following commands on the Arm to see the correlation between the PCI device and its representors:

```
mst start
mst status -v

PCI devices:
-----
DEVICE_TYPE MST PCI RDMA NET NUMA
BlueField(rev:0) NA 03:00.1 mlx5_3,mlx5_2 net-eth2,net-rep1-ffff -1
BlueField(rev:0) NA 03:00.0 mlx5_1,mlx5_0 net-rep0-ffff,net-eth1 -1
```

Note

The rep0-0 / rep1-0 are the representors facing the x86 host and the rep0-ffff / rep1-ffff are the representors facing the network.

Switching to SEPERATED_HOST Mode from EMBEDDED Mode

On the x86 host:

Enable SEPERATED_HOST mode:

```
x86_host #mst start
x86_host #mlxconfig -d /dev/mst/mt41682_pciconf0 s INTERNAL_CPU_MODEL=0
```

Reboot.

Verify:

```
x86_host#mst start
x86_host#mlxconfig -d /dev/mst/mt41682_pciconf0 q | grep -i model
```

Testing Traffic in EMBEDDED Mode using DPDK

In Embedded mode, traffic from the x86 server hosting the SmartNIC to the remote x86 server hosting the ConnectX-5 is going via the SmartNIC Arm.

DPDK application is used to forward the traffic and bypass Arm OS kernel. Follow the configuration steps below to test it.

On the Arm

Stop OVS:

```
systemctl stop openvswitch
```

Verify hugetlbfs is mounted:

```
mount | grep -i hugetlbfs
```

if not, mount it:

```
mount -t hugetlbfs nodev /dev/hugepages
```

Set Hugepages:

```
echo 4 > /sys/kernel/mm/hugepages/hugepages-524288kB/nr_hugepages
```

Verify hugepages were allocated (look for HugePages_Total and Free):

```
cat /proc/meminfo | grep -i huge
```

Identify the SmartNIC PCI device ID:

```
lspci | grep -i mellanox
```

Start testPMD application **while probing the representors** (pci device is 03:00.0):

```
cd /usr/bin/  
./testpmd -w 03:00.0,representor=[0,65535] - -i  
testpmd> start
```

Set IPs on both x86 host and ConnectX-5 host and test connectivity with ping (it will be forwarded by the Arm DPDK)

On the **ConnectX-5 external host** start iperf server:

```
iperf -s
```

On the **x86 host**, start iperf client:

```
iperf -c 172.16.0.2 -P12
```

While the test is running, you can check traffic statistics on **Arm testPMD**:

```
testpmd> show port stats all
```

Bare Metal Deployment Guide

Install Bare Metal Jump Host

N/A

Creating a Node Inventory File

N/A

Creating the Settings Files

N/A

Running

N/A

Virtual Deployment Guide

N/A

Standard Deployment Overview

N/A

Snapshot Deployment Overview

N/A

Special Requirements for Virtual Deployments

N/A

Install Jump Host

N/A

Verifying the Setup - VMs

N/A

Upstream Deployment Guide

N/A

Upstream Deployment Key Features

N/A

Special Requirements for Upstream Deployments

N/A

Scenarios and Deploy Settings for Upstream Deployments

N/A

Including Upstream Patches with Deployment

N/A

Running

N/A

Interacting with Containerized Overcloud

N/A

Verifying the Setup as defined the Akraino validation feature project plus any additional testing specific to the blue print

N/A

Developer Guide and Troubleshooting

Utilization of Images

N/A

Post-deployment Configuration

N/A

Debugging Failures

N/A

Reporting a Bug

N/A

Uninstall Guide

not applicable

Troubleshooting

Error Message Guide

Server Unable to Find the SmartNIC

- Ensure that the SmartNIC is placed correctly
- Make sure the SmartNIC slot and the SmartNIC are compatible

- Install the SmartNIC in a different PCI Express slot
- Use the drivers that came with the SmartNIC or download the latest
- Make sure your motherboard has the latest BIOS
- Try to reboot the server

The SmartNIC no Longer Works

- Reseat the SmartNIC in its slot or a different slot, if necessary
- Try using another cable
- Reinstall the drivers for the network driver files may be damaged or deleted
- Reboot the server

SmartNIC Stopped Working After Installing Another

- Try removing and re-installing all SmartNICs
- Check that cables are connected properly
- Make sure your motherboard has the latest BIOS

Link Indicator Light is Off

- Try another port on the switch
- Make sure the cable is securely attached
- Check you are using the proper cables that do not exceed the recommended lengths
- Verify that your switch and SmartNIC port are compatible

Link Light is On but No Communication is Established

- Check that the latest driver is loaded
- Check that both the SmartNIC and its link are set to the same speed and duplex settings

Maintenance

Blue Print Package Maintenance

Software maintenance

N/A

Hardware maintenance

N/A

Blueprint Deployment Maintenance

N/A

Frequently Asked Questions

N/A

License

N/A

References

N/A

Definitions, acronyms and abbreviations

OVS : Open Virtual-Switch

