

GV I/O Virtualized Video Servers

Csaba Kormendy, December 2020

Introduction

With the GV I/O 4.0 release, Grass Valley offers customers the option of running the GV I/O video server product in a virtualized environment. The GVIO-SW-VM product supports SD, 720p and 1080i with currently supported codecs as a 2in/1out configuration via SMPTE ST 2110 I/O only. SDI, UHD and 1080p are not supported at this time. GV I/O virtualized video servers are *ONLY* supported with VMware vSphere (ESXi) environments. The VM product is delivered in the convenient Open Virtualization Format (OVF), which can easily be deployed onto a vSphere host or cluster. This document provides a comprehensive overview of everything that is required for a successful virtualized deployment of the GV I/O VM product.

Table of Contents

GV I/O VM Hardware Requirements	2	Deploying GV I/O VM	5
GV I/O VM Provisioning	2	VMware ESXi Host BIOS Configuration	5
Servers	2	VMware ESXi Host Software	5
Storage	2	Deploying Mellanox ConnectX-5 Software/Firmware to VMware ESXi	6
Storage Best Practice Required to Achieve Acceptable Performance	3	Configure Mellanox Network Ports for SR-IOV in vSphere	9
Storage for vSphere Datastores	3	Importing GV I/O VM To ESXi Host	9
Proxy File Storage	3	Mellanox ConnectX-5 SR-IOV Related Configuration of GV I/O VM that must be Performed with Customer Hardware	12
GV I/O VM Networking	3	Deploying Drivers and Other Software to the GV I/O VM	14
Networking Best Practices	3	Smoke Testing GV I/O VM Operation	15
Network Naming Conventions	3		
Manufacture and Distribution of VMs	4		
Distribution Method	4		
On-site Installation Prerequisites and Preparation	4		
Host Licenses	4		
vCenter	4		
Network	4		
Storage	4		

GV I/O VM Hardware Requirements

The GV I/O product processes real-time video signals. This requires a high-performance, low-latency system design in order to function effectively. It is essential to pay careful attention to the resource requirements specified throughout this document when designing a vSphere environment into which a GV I/O virtualized solution can be deployed.

GV I/O VM Provisioning

One instance of the GV I/O VM product requires the resources summarized below in order to support a 2in/1out channel configuration for SD, 720p and 1080i video formats:

CPU

- 26vCPU — two virtual sockets with 13 cores per socket
- Reservation: 57044 MHz
- Shares: Normal

Memory

- 80 GB
- Reservation: 81920 (Reserve all guest memory)
- Shares: Normal

Provisioned Storage

- 200 GB

VM Advanced Options

- Latency sensitivity set to high

Servers

Grass Valley has currently validated only a single server solution for virtualized GV I/O installations.

- Cisco UCS C240 M5 Rack Server

Hardware components critical to the performance of the GV I/O VM product as validated by Grass Valley are summarized in the table below along with their corresponding Cisco part numbers. Additional Cisco components, VMware products and licenses will be needed as defined by the customer to complete the bill of materials to build a complete server unit.

The DDR4 RAM modules should be deployed such that each memory channel of the dual-CPU system is populated. This configuration is required to maximize memory bandwidth.

The Mellanox network card must be installed to one of the PCIe x16 slots. These are identified as PCIe slots 2, 4 or 5 in the Cisco documentation.

Storage

Grass Valley has validated a number of options for storage in GV I/O virtualized solutions. These storage configurations have been tested and found to provide the high performance and low latency that GV I/O video servers require.

Grass Valley has validated the following storage for VM images:

- 4x 1 TB 12G SAS 7.2K RPM HDD in a RAID-10 configuration internal to the Cisco UCS C240 M5 server unit

Grass Valley has validated two third-party options for storing VM images and Proxy media:

- NetApp hybrid arrays, with SSD cache enabled
- Tegile hybrid arrays, with SSD cache enabled

Grass Valley has validated the following storage for high-resolution and proxy media use :

- GV AMS Pro (Isilon)

Qty	Cisco Part Number	Description
1	UCSC-C240-M5SX	UCS C240 M5 24 SFF + 2 rear drives w/o CPU,mem,HD,PCIe, PS
2	UCS-CPU-I6238R	Intel 6238R 2.2 GHz/165W 28C/38.5 MB DDR4 2933 MHz
12	UCS-MR-X16G1RT-H	16 GB DDR4-2933 MHz RDIMM/1Rx4/1.2v
1	UCSC-PCI-1-C240M5	Riser 1 including 3 PCIe slots (x8, x16, x8); slot 3 required CPU2
1	UCSC-PCI-2A-240M5	Riser 2A including 3 PCIe slots (x8, x16, x16) supports GPU
1	UCSC-P-M5D100GF*	MELLANOX CX-5 MCX516A-CDAT 2x 100 GbE QSFP PCIe NIC

* Grass Valley has validated the GV I/O VM product with Cisco part number UCSC-P-M5D100GF, Mellanox part number MCX516A-CDAT for SMPTE ST 2110 I/O. This is a network card from the Mellanox ConnectX-5 product family. Grass Valley anticipates recommending a network card from the newer Mellanox ConnectX-6 family in the near future.

GV I/O Virtualized Video Servers

Storage Best Practice Required to Achieve Acceptable Performance

- Provide a separate, dedicated network path for the vSphere datastores
- Provide a separate, dedicated network path for SMB storage
- Minimize the number of network hops between the VMware host(s) and the storage system
- GV I/O virtualized video servers require multiple types of data storage:
 - Storage for vSphere datastores
 - Storage for high-resolution media files
 - Storage for proxy-quality media files

Storage for vSphere Datastores

vSphere datastores are used to hold the files that comprise the actual virtual machines. The following options for vSphere Datastores have been tested and are supported.

- 4x 1 TB 12G SAS 7.2K RPM HDD in a RAID-10 configuration internal to the Cisco UCS C240 M5 server unit
- NFS Datastore on NetApp hybrid array, SSD cache enabled
- NFS Datastore on Tegile hybrid array, SSD cache enabled

Proxy File Storage

SMB storage is used to contain the lower resolution “proxy-quality” versions of the video assets managed by a GV STRATUS system, also known as proxy storage. These files are generated by a K2 Summit client, a GV I/O client or a GV STRATUS Render Engine and are served by the GV STRATUS Proxy Server. Excellent read and write performance to/from this share is important to maintain good performance of a GV STRATUS system. Additional requirements:

- 10 GB network interface
- Proxy storage share name must be “proxy”
- The service accounts used on the Proxy Server and all GV I/O clients require read/write access
- The following options for proxy storage have been tested and are supported:
 - SMB share on Network Appliance hybrid array, SSD cache enabled
 - SMB share on Tegile hybrid array, SSD cache enabled
 - SMB share on GV AMS Pro (Isilon) array, SSD cache enabled

GV I/O VM Networking

A virtualized GV I/O video server requires the same networking considerations as a physical GV I/O video server. The network requirements are summarized below.

Server Role	VM Storage*	Control	FTP	Media 1	Media 2	S2110-1	S2110-2
GV I/O	Yes, 10 GB	Yes, 1 GB	No	Yes, 10 GB	No	Yes, 100 GB	Yes, 100 GB

* VM Storage on vSphere datastores

Networking Best Practices

For best performance and highest reliability, the following best practices are recommended:

- Be sure that all physical switches are on the vSphere supported equipment list
- Use redundant physical networking whenever possible
- Use multiple vNICs when possible
- Isolate traffic by using separate virtual switches and vNICs for each network
- Use jumbo frames for accessing VM datastores if equipment permits
- Minimize the number of network hops between the VMware hosts, storage and other devices they interact with

Network Naming Conventions

GV I/O VM video servers come from the factory with network interface names as shown below. During deployment you are presented with an opportunity to map each factory network name to a network name found in the vSphere environment at the installation site. The network names at the user installation site may differ from those shown below.

Factory Network Name	Purpose
Control Connection	Control network, database access, access to corporate network
Media Connection	Access to high-resolution media files on GV AMS Pro storage
HBR Media Port 1	SMPTE ST 2110 media I/O
HBR Media Port 2	SMPTE ST 2110 media I/O

Manufacture and Distribution of VMs

Each GV I/O VM is manufactured and configured by Grass Valley to the specifications of a customer's order. This ensures that the VM is properly configured and that the Microsoft products on the VM have been correctly licensed with the unique license code sold as part of that VM.

During the manufacturing process, the appropriate GV I/O software is installed and configured on each VM, and the requested IP addresses have been put in place.

On-site Installation Prerequisites and Preparation

The vSphere and network infrastructure at the deployment site must be in place before the GV I/O VM components can be deployed. The following should all be in place, be healthy and have sufficient resources to support the GV I/O VMs.

In most cases, the user is responsible for purchasing, installing and configuring the VMware ESXi hosts. Please verify that the ESXi hosts match the supported hardware configuration.

The GV I/O VM product requires VMware 6.7 Update 3 or later. While VMware 7.0 is available, GV I/O VM has only been validated with VMware 6.7 Update 3.

Host Licenses

Normally the user is responsible for purchasing and installing the vSphere licenses used on the ESXi host, or on the vCenter server in a multi-host environment. Take care to validate that the vSphere licenses for each host matches a supported version and a supported edition as specified. The GV I/O VM product uses SR-IOV technology which requires VMware Enterprise Plus licensing.

vCenter

In a vSphere environment with more than one ESXi host, a vCenter server is required. vCenter is a crucial part of any vSphere environment that has more than one host — it is the portal used to manage the hosts, virtual networks, storage and virtual machines.

Distribution Method

After manufacturing and testing, each VM is exported to an OVF file and placed on a USB drive that is delivered to the customer. Each OVF file represents a single VM. If an order includes multiple GV I/O VM line items, multiple USB drives will be shipped. It is important to handle the USB drives with care, if the drive becomes damaged from mishandling and a replacement drive is needed you could lose valuable time in getting your system online. It is also important to be aware that all necessary Microsoft license keys come pre-installed on GV I/O VMs and the proof of evidence for those licenses is included with each USB drive. Embedded license keys in GV I/O VMs are not transferable to other VMs or devices, and must be retained for audit purposes. Please keep your USB drives in a safe place!

Network

In any VMware environment there are two parts to the network:

- The virtual network defined within the ESXi hosts
- The physical network that connects the ESXi hosts to the other Grass Valley equipment and the rest of the user's network

In addition, VMware environments that use the GV I/O VM product will have a third portion of the network used for SMPTE ST 2110 media over IP I/O.

The system report, which accompanies the order, specifies which servers have interfaces on which networks and what IP each of those interfaces have. It should be verified that:

- All networks that will be used by the VMs have been defined on the virtual switches on each host
- There is a record of the VLAN name for each network

VMs are shipped from the factory with network names that match those shown on the system report. When deploying the VMs, there will be a prompt to map those names to a VLAN name defined on the vSphere infrastructure. Have that information ready before attempting to deploy the VMs.

It is essential to verify that the physical network is correctly built.

All vNICs associated with each VLAN must have the recommended bandwidth and redundancy as specified in the GV I/O VM documentation.

Storage

Proper storage for the VMs, for proxy files and for high-resolution media files is critical for good system performance. All storage should comply with supported storage options specified by Grass Valley.

Identify early in the deployment process the storage that will hold the VMs since this information is required when deploying each VM.

Deploying GV I/O VM

The GV I/O VM product processes real-time video signals. This requires a high-performance, low-latency system design in order to function effectively. This portion of the document details configuration settings and procedures that must be executed for the GV I/O VM product to function effectively.

VMware ESXi Host BIOS Configuration

For the servers hosting GV I/O VMs, please ensure the following BIOS settings are configured as noted below:

SR-IOV Support = Enabled

This setting should be enabled by default. Single Root I/O Virtualization (SR-IOV) allows multiple VMs running a variety of guest operating systems to share a single PCIe network adapter within a host server.

Hyper-Threading [ALL] = Disable

Hyperthreading allows parallel processing of multiple threads and will be enabled by default. For the GV I/O VM product this setting must be configured to a disabled state.

Intel (R) VT = Enable

This setting will be enabled by default. Intel Virtualization Technology (VT) enables CPU features that allow hardware-accelerated virtualization functionality required for VMware ESXi 6.7 and later.

Numa = Enable

This setting enables or disables nonuniform memory access (NUMA) and will be enabled by default.

ADDC Sparing = Disable

Patrol Scrub = Disable

These settings are enabled by default and have to do with scanning, reporting, and correcting bit errors in RAM for applications that require high reliability and fault-tolerance. For the GV I/O VM product we want maximum memory bandwidth and minimum latency.

Intel VT for Directed I/O (VT-d) = Enable

This setting should already be enabled by default. VT-d provides enhanced performance for SR-IOV. If the "Intel VT for Directed I/O (VT-d)" setting is being changed, enabling for GV I/O VM, the C240 server unit must be powered off and then powered on again in order for the setting to take effect.

SpeedStep (Pstates) = Disable

Turbo Mode = Disable

These settings are enabled by default and provide a capability for the CPU to adjust its clock frequency. For the GV I/O VM product this setting must be configured to a disabled state.

Hardware P-States = Disable

This setting is set to "Native Mode" by default. For the GV I/O VM product we want maximum performance and minimum latency.

Autonomous Core C-State = Disable

CPU C6 report = Disable

Enhanced Halt State (C1E) = Disable

OS ACPI Cx = ACPI C2

For the GV I/O VM product we are configuring these power management settings for maximum performance and minimum latency.

Package C State = C0/C1 state

This setting is set to "C0/C1 State" by default. For the GV I/O VM product we want maximum performance and minimum latency.

ENERGY_PERF_BIAS_CFG mode = Performance

For the GV I/O VM product we want maximum performance and minimum latency.

VMware ESXi Host Software

For the servers hosting GV I/O VMs, make sure the following items are addressed:

- VMware 6.7 Update 3 is installed and configured
- NTP is disabled
- Installer Privileges. The installation process requires administrator privileges on the target machine
- Disable Hyperthreading on ESXi
- ESXi Host Power Configuration = High performance

GV I/O Virtualized Video Servers

Deploying Mellanox ConnectX-5 Software/Firmware to VMware ESXi

Using the NVIDIA/Mellanox document “[RDG: VMware based Mellanox SMPTE 2110 Media Cloud](#)” and documents that it in turn references, install drivers and firmware for the Mellanox ConnectX-5 network card installed in the server hosting the GV I/O VMs.

Please make sure to use software/firmware build numbers referenced in this document rather than the versions in the third-party documents referenced in this section. The build numbers referenced in this document are the build numbers that the GV I/O VM product has been validated with.

VMware ESXi 6.7 Driver for nmlx5_core 4.17.15.16 NIC for Mellanox ConnectX4/5 Ethernet Adapters

The file name is MLNX-NATIVE-ESX-ConnectX-4-5_4.17.15.16-10EM-670.0.0-13000354.zip

The Mellanox NIC driver installation procedure is based on the Mellanox document “[How-to: Mellanox driver upgrade on VMware ESXi 6.5 and above](#)” with the modifications as documented below:

At Step 9 “Install the driver”

The form of the command is given as:

```
esxcli software vib install -d <path>/<bundle_file>
```

but for the GV I/O VM product needs to be:

```
esxcli software vib install -d <path>/<bundle_file> --no-sig-check
```

and so the full command ends up being something like:

```
ESXi Console
esxcli software vib install -d /tmp/MLNX-NATIVE-ESX-ConnectX-4-5_4.17.15.16-10EM-670.0.0-13000354.zip --no-sig-check
```

Continue with the next steps in the Mellanox document, rebooting the server and verifying the driver installation.

Mellanox Firmware Tools (MFT) version 4.15.0 for "VMware ESX Server version 6.7 Native x64"

The file names are:

```
mft-4.15.0.104-10EM-650.0.0.4598673.x86_64.vib
```

```
nmst-4.15.0.104-1OEM.650.0.0.4598673.x86_64.vib
```

The Mellanox Firmware Tools installation procedure is based on the Mellanox document “[How-to: Mellanox driver upgrade on VMware ESXi 6.5 and above](#)” with the modifications as documented below:

At Step 6 “Install the vib package”

Again a “--no-sig-check” needs to be appended to the command and so the full commands end up being something like.

```
ESXi Console
esxcli software vib install -v /tmp/mft-4.15.0.104-10EM-650.0.0.4598673.x86_64.vib --no-sig-check
```

```
ESXi Console
esxcli software vib install -v /tmp/nmst-4.15.0.104-1OEM.650.0.0.4598673.x86_64.vib --no-sig-check
```

Continue with the next steps in the Mellanox document, rebooting the server and starting the mst driver.

ConnectX-5 Ethernet Card Firmware version 16.28.1002 for MCX516A-CDAT (PSID MT_000000013)

The file name is fw-ConnectX5-rel-16_28_1002-MCX516A-CDA_Ax-UEFI-14.21.16-FlexBoot-3.6.101

The Mellanox NIC Firmware update procedure is based on the Mellanox document "[How-to: Firmware update for Mellanox ConnectX-4/5 adapter on VMware ESXi 6.5 and above](#)" with the modifications as documented below:

At Step 12 "Burn the firmware image to the Adapter Card"

The form of the command is given as:

```
~ flint -d <device_name> -i <binary image> b
```

but for the GV I/O VM product needs to be:

```
~ flint -d <device_name> -i <binary image> -allow_psid_change burn
```

and so the full command ends up being something like:

```
ESXi Console
/opt/mellanox/bin/flint -d mt4121_pciconf0 -i /tmp/Fw-ConnectX5-rel-16_28_1002-MCX516A-CDA_Ax-UEFI-14.21.16-FlexBoot-3.6.101 -allow_psid_change burn
```

Continue with the next steps in the Mellanox document, rebooting the server and verifying the firmware version.

Enable Hardware Timestamp and Packet Pacing

1. Log in to ESXi command line interface with root permissions.
2. Get the NIC device ID.

```
ESXi Console
/opt/mellanox/bin/mst status
MST devices:
-----
mt4121_pciconf0
```

3. Using the NIC device ID obtained in the previous command, execute commands as noted below. For simplicity, the device ID from the example step above will be used in the example steps below.

```
ESXi Console
/opt/mellanox/bin/mlxconfig -d mt4121_pciconf0 set VF_VPD_ENABLE=1
Device #1:
-----
Device type:   ConnectX5
Name:          MCX516A-CDA_Ax
Description:   ConnectX-5 Ex EN network interface card; 100GbE dual-port QSFP28; PCIe4.0 x16; tall bracket; ROHS R6
Device:        mt4121_pciconf0.1
Configurations:
                Next Boot      New
VF_VPD_ENABLE  True(1)      True(1)
Apply new Configuration? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.
```

4. After rebooting, the setting can be verified via the command:

```
ESXi Console
/opt/mellanox/bin/mlxconfig -d mt4121_pciconf0 q | grep VPD
VF_VPD_ENABLE          True(1)
```

5. Execute the following command to write the value 3 to the register at 0xd8068.

```
ESXi Console
/opt/mellanox/bin/mcra mt4121_pciconf0 0xd8068 3
```

6. To check the parameter, execute:

```
ESXi Console
/opt/mellanox/bin/mcra mt4121_pciconf0 0xd8068
0x00000003
```

7. Create a file with the text:

```
MLNX_RAW_TLV_FILE
0x00000004 0x0000010c 0x00000000 0x00000001
```

For example, using the commands:

```
ESXi Console
echo "MLNX_RAW_TLV_FILE" > /tmp/mlxconfig_raw.txt
echo "0x00000004 0x0000010c 0x00000000 0x00000001" >> /tmp/mlxconfig_raw.txt
```

8. The file contents can be reviewed via:

```
ESXi Console
cat /tmp/mlxconfig_raw.txt
MLNX_RAW_TLV_FILE
0x00000004 0x0000010c 0x00000000 0x00000001
```

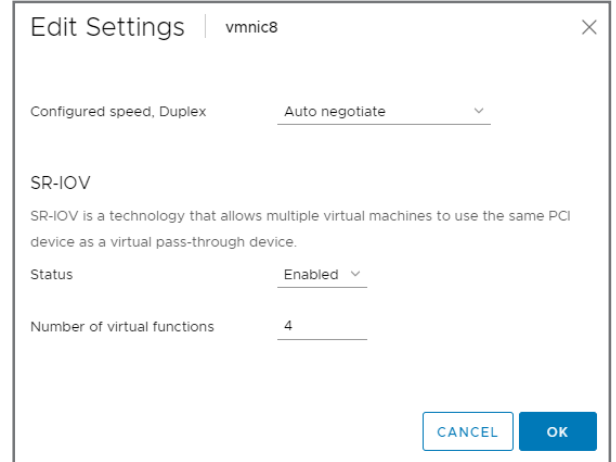
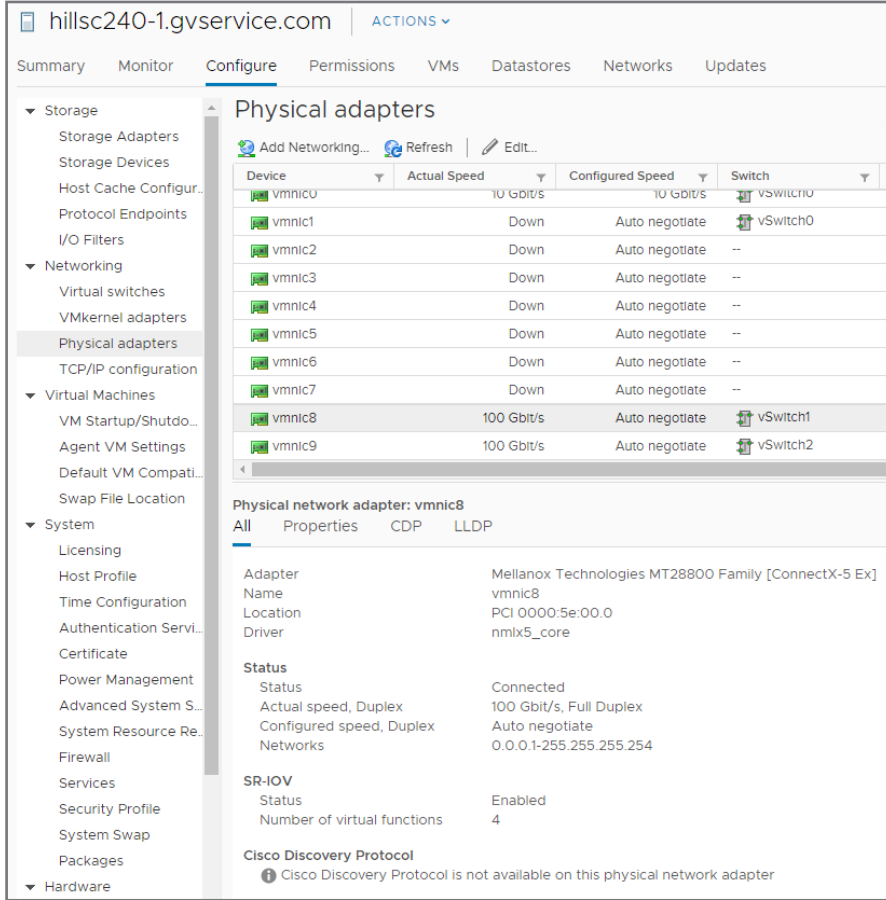
9. Write this raw configuration file to the NIC:

```
ESXi Console
/opt/mellanox/bin/mlxconfig -d mt4121_pciconf0 -f /tmp/mlxconfig_raw.txt set_raw
Raw TLV #1 Info:
Length: 0x4
Version: 0
OverrideEn: 0
Type: 0x0000010c
Data: 0x00000001
Operation intended for advanced users.
Are you sure you want to apply raw TLV file? (y/n) [n] : y
Applying... Done!
-I- Please reboot machine to load new configurations.
```

10. Reboot the server.

Configure Mellanox Network Ports for SR-IOV in vSphere

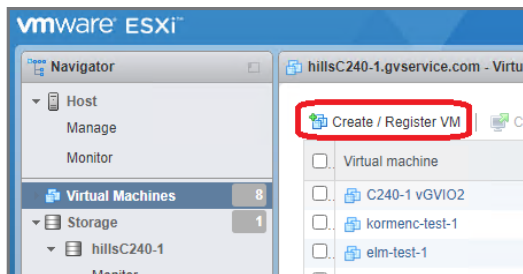
Select one of the Mellanox network ports, click on **Edit**, and enable SR-IOV for the network port. Make sure to leave speed, duplex set to "Auto negotiate". We've observed problems when setting the port speed to other values. Set "Number of virtual functions" to 4. Repeat the process for the other Mellanox network port.



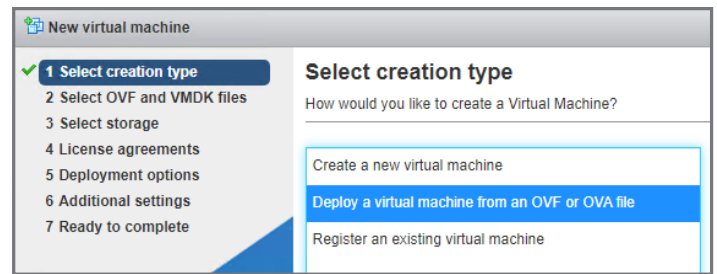
Importing GV I/O VM To ESXi Host

The GV I/O VM product is delivered to customers via a USB drive in Open Virtualization Format (OVF), which can easily be deployed onto a vSphere host or cluster.

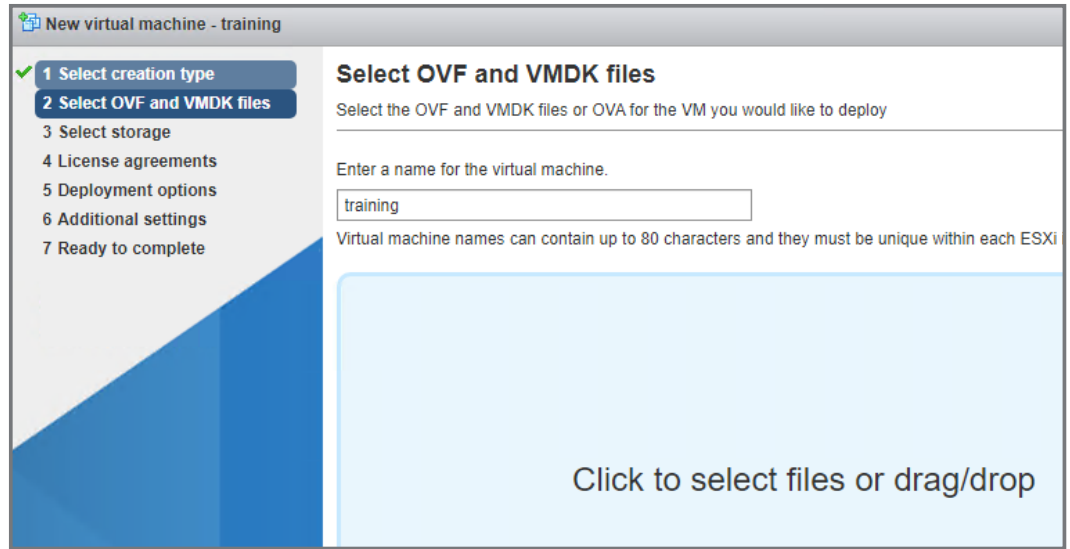
1. In vSphere, click "Create / Register VM".



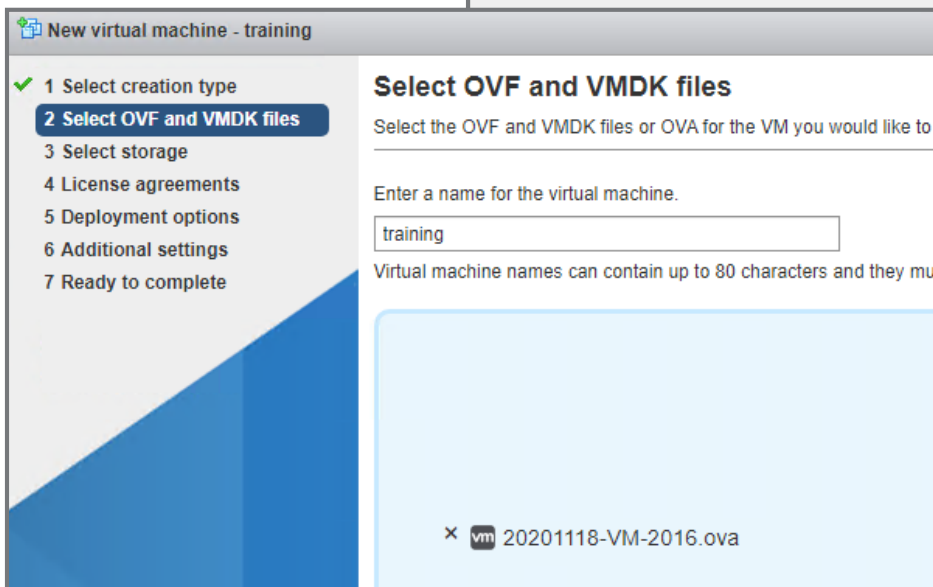

2. Choose "Deploy a virtual machine from an OVF or OVA file".



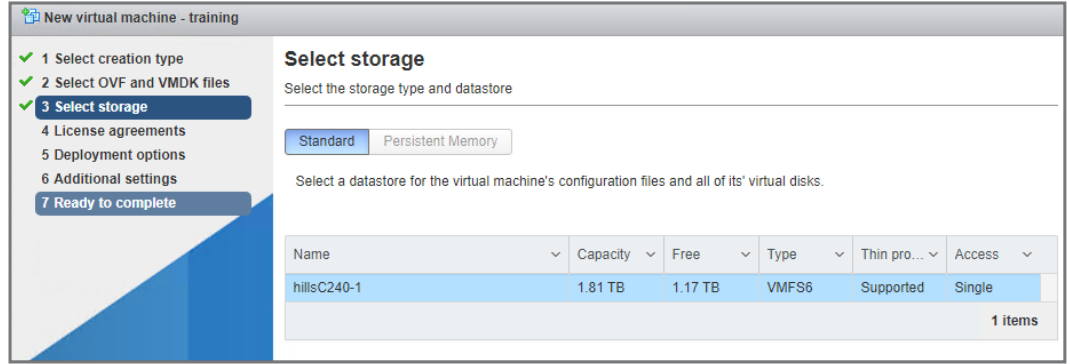
3. Enter a name for the virtual machine.



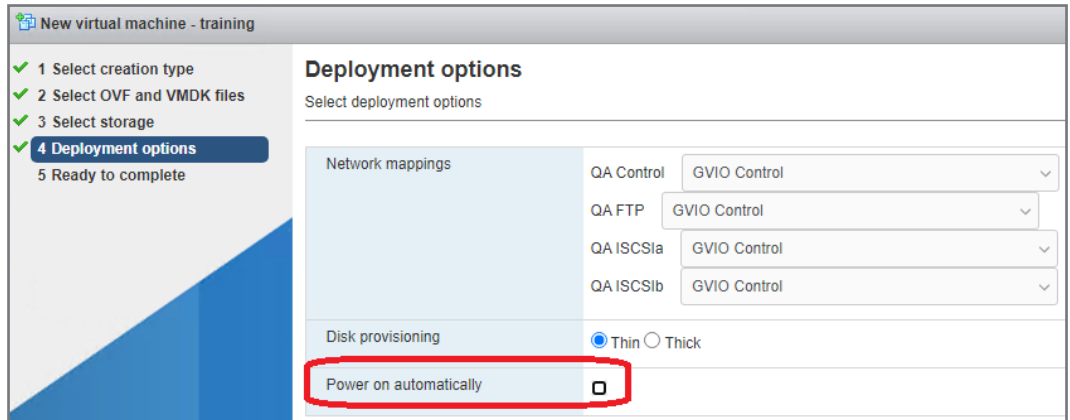
4. Choose the OVA file.



5. Select storage. In this example, we are choosing the C240 unit's internal storage.

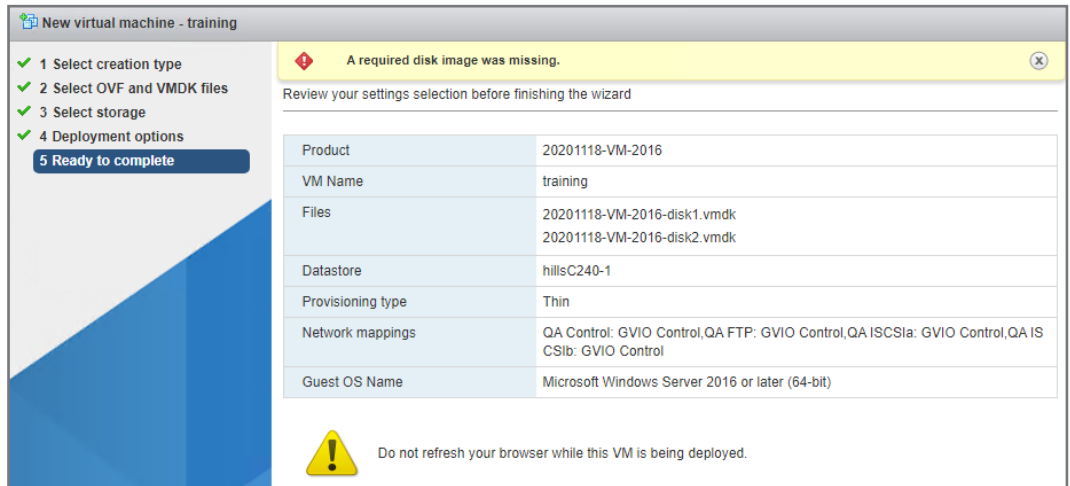


6. Leave network mappings as is for now. Leave disk provisioning at default "Thin" setting. Uncheck power on automatically. Additional VM configuration is required before powering on the VM for the first time.



7. Review summary. Ignore the warning in this screen shot.

8. Finish the process and wait for the OVA to finish importing. It can take on the order of 15 minutes to complete.



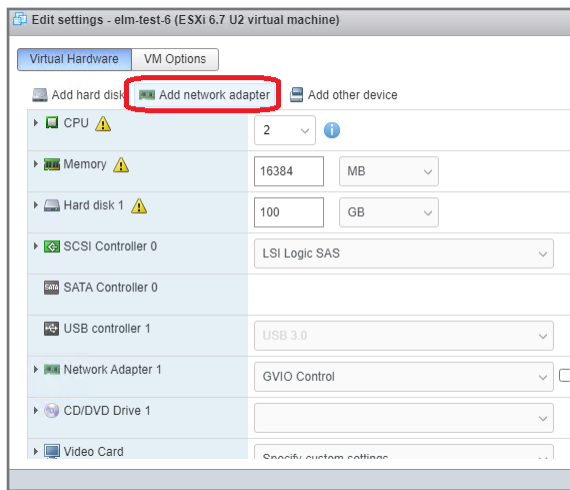
Mellanox ConnectX-5 SR-IOV Related Configuration of GV I/O VM that must be Performed with Customer Hardware

After the VM has been copied into the VMware environment, VM configuration steps need to be executed to match the host’s hardware. The GV I/O VM product uses a technology referred to as Single Root I/O Virtualization (SR-IOV). When SR-IOV is used with VMware/vSphere, some features of VMware/vSphere do not function when SR-IOV is enabled for the VM (<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.networking.doc/GUID-E8E8D7B2-FE67-4B4F-921F-C3D6D7223869.html>). The following features are not available for virtual machines configured with SR-IOV:

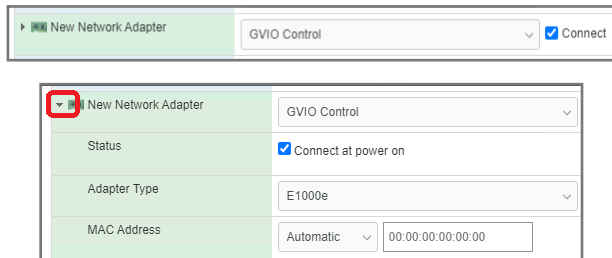
- vSphere vMotion
- Storage vMotion
- vShield
- NetFlow
- VXLAN Virtual Wire
- vSphere High Availability
- vSphere Fault Tolerance
- vSphere DRS
- vSphere DPM
- Virtual machine suspend and resume
- Virtual machine snapshots
- MAC-based VLAN for passthrough virtual functions
- Hot addition and removal of virtual devices, memory and vCPU
- Participation in a cluster environment
- Network statistics for a virtual machine NIC using SR-IOV passthrough

Add SR-IOV NIC ports to VM instance


1. Select the VM and edit its settings. Click on “Add network adapter.”



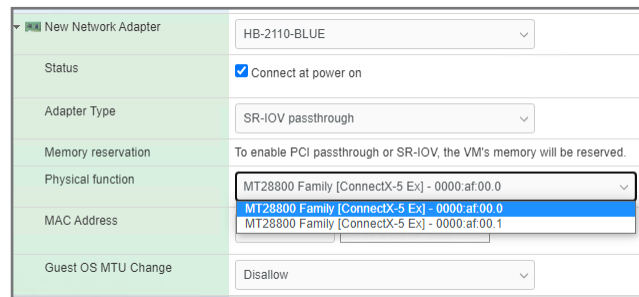
2. Expand the network adapter settings.



3. Choose the network name and for “Adapter Type,” choose SR-IOV Passthrough.



4. For “Physical function,” choose the entry that matches the desired network name. For additional information, see “How To Choose Physical Function” below.



5. Leave “MAC Address” set to Automatic and “Guest OS MTU Change” set to Disallow.


6. Save the settings.

7. Repeat the process to add a second SR-IOV NIC port to the VM, for the other physical port of the Mellanox network adapter.

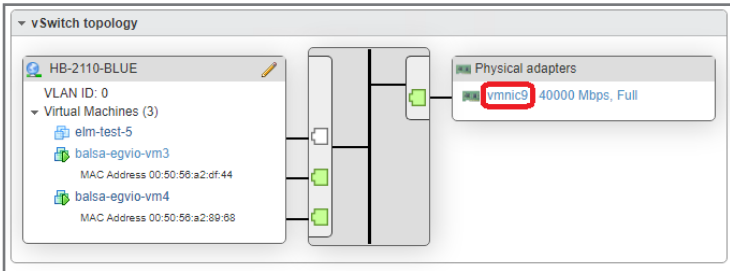
GV I/O Virtualized Video Servers

How To Choose Physical Function


1. Click on Networking and select the “Virtual switches” tab. Click on the vSwitch that goes with one of the Mellanox NIC ports.



2. Click on the vmnic that represents the physical port.



3. Make note of the PCI address entry.



4. This tells us that “0000:af:00:1” goes with “vmnic9” and the “HB-2110-BLUE” network name.

GV I/O Virtualized Video Servers

Deploying Drivers and Other Software to the GV I/O VM

A number of software components and configuration steps must be completed using the hardware where an instance of the GV I/O VM product is deployed.

1. Change Embedded Security Manager to Update Mode

2. Mellanox ConnectX-5 driver installation

Driver installation must be performed after deploying VM because the driver installer is capable of updating firmware on Mellanox ConnectX-5 card. The firmware update can take on the order of 30 minutes. Since the Mellanox ConnectX-5 card is shared by multiple VMs, a firmware update to the network card could potentially be disruptive to other VMs using the network card, so plan accordingly.

Obtain the Mellanox ConnectX-5 driver installer (MLNX_WinOF2...) from the same GV I/O prerequisites file used for R640 based GV I/O units. The GV I/O prerequisites file can be downloaded via <ftp://ftp.grassvalley.com/GVIO/>

You do not need to install Mellanox WinMFT64 component used with R640-based GV I/O. It is not relevant for GV I/O VM. Similarly, “mst status -v” and “rivermax_config” steps should be omitted for GV I/O VM. If executed with GV I/O VM, these programs will benignly fail as they will be unable to communicate with the Mellanox card.

After the driver installation has completed, Mellanox network adapter ports to “HBR Media Port 1” and “HBR Media Port 2” to follow the same convention as for R640 based GV I/O units.

Configure IP addresses for the Mellanox network adapter ports as appropriate.

3. Rivermax license file

Check for presence of c:\license\rivermax.lic file. The license file is required in order to enable SMPTE ST 2110 media I/O related hardware acceleration provided by the Mellanox ConnectX-5 network card. The license file should have been deployed to VM by GV Manufacturing. If the license file is not present, obtain the “GV global” license file from <ftp://ftp.grassvalley.com/Networking/Mellanox/>

The accompanying CSV file contains a list of Mellanox network cards that the Rivermax license is valid for. One can compare the NIC serial number vs. CSV file. If the serial number is not present in the CSV file, please contact your Grass Valley representative to obtain Grass Valley part number GVIO-SW-OPT-RMAX, which should have been part of the sales order for the GV I/O VM purchase.

Only one instance of GVIO-SW-OPT-RMAX is required per Mellanox ConnectX-5 NIC. If two VMs are deployed to the hardware, only one GVIO-SW-OPT-RMAX is needed. The same c:\license\rivermax.lic file is used for both VM instances. So, if a second VM is being added to a system where a working rivermax.lic file is already present, one can copy the rivermax.lic file from one GV I/O VM to the other.

4. Execute NetAdapterConfig_*.ps1 script

Obtain the script from the same GV I/O prerequisites file used for R640-based GV I/O units. The GV I/O prerequisites file can be downloaded via <ftp://ftp.grassvalley.com/GVIO/>

The script performs Mellanox network adapter-related tuning that must be performed after the Mellanox NIC driver has been installed or updated. Executing this script more than once is benign.

5. VyDriver related optimization

Open Network and Sharing Center in the Control Panel, then right-click and select Properties on the first Mellanox network port. Disable all protocol and services except Vy and Internet Protocol Version 4 (TCP/IPv4). Go to the second Mellanox network port, then disable all protocol and services except Vy and Internet Protocol Version 4 (TCP/IPv4). For all non-Mellanox network ports, right-click on Properties and uncheck Vy for each network port.

6. Execute WeightNetworkAdapters*.ps1 script

Obtain the script from C:\Grass Valley\ as present in the GV I/O VM's file system, just like for a R640-based GV I/O unit. Please refer to WeightNetworkAdapters_README.txt for details about executing this script.

The script performs general network adapter-related tuning that must be performed after any network driver has been installed or updated for the GV I/O VM instance, or if new network ports are added to the VM. Executing this script more than once is benign.

7. Change Embedded Security Manager to Enabled Mode

8. Reboot the GV I/O VM instance

Smoke Testing GV I/O VM Operation

At this point, the GV I/O VM product is ready for basic review and commissioning-related configuration and testing. IP addresses, etc. should be configured for each of the VM's network interfaces.

1. Review the control network for connectivity/functionality
2. Review the media storage network for connectivity/functionality

Verify that a V:\ drive has been mapped and the user account to be used with the GV I/O software has read/write access to the media storage location.

3. Review SMPTE ST 2110 I/O network for connectivity/functionality.

After configuring the S2059 Client application to for the SMPTE ST 2110 I/O network ports, review the PTP status information and verify that the GV I/O VM is able to synchronize with PTP.

4. Review that AppServer is able to start successfully

Review GV I/O logs for Rivermax related errors and address them by making sure that a valid Rivermax license is present.



WWW.GRASSVALLEY.COM

Join the Conversation at [GrassValleyLive](#) on Facebook, Twitter, YouTube and [Grass Valley](#) on LinkedIn.



www.grassvalley.com/blog

This product may be protected by one or more patents. For further information, please visit: www.grassvalley.com/patents.

Grass Valley®, GV® and the Grass Valley logo are trademarks or registered trademarks of Grass Valley USA, LLC, or its affiliated companies in the United States and other jurisdictions. Grass Valley products listed above are trademarks or registered trademarks of Grass Valley USA, LLC or its affiliated companies, and other parties may also have trademark rights in other terms used herein.

Copyright © 2021 Grass Valley Canada. All rights reserved. Specifications subject to change without notice.