Copyright(c) 2013 - 2020 Intel Corporation

This release includes the native ixgben VMware ESX Driver for Intel(R) Ethernet Controllers 82599, x520, x540, x550, x552 and x553 family

Driver version: 1.8.9.0

Supported ESXi release: 7.0

==========

Contents

- Important Notes
- Supported Features
- New Features
- New Hardware Supported
- Bug Fixes
- Known Issues and Workarounds
- Command Line Parameters
- Previously Released Versions

==========

Important Notes:

- VMware vSphere Hypervisor (ESXi) 7.0 Support:

  Added VMware vSphere Hypervisor (ESXi) 7.0 support. VMware vSphere Hypervisor (ESXi) 7.0 introduces changes related to:

  - How SR-IOV VFs are created.
  - How driver module parameters function.
  - How driver modules are upgraded.

  Consult VMware vSphere Hypervisor (ESXi) 7.0 release notes for a complete list of new features and changes. To avoid serious problems, carefully review VMware hardware requirements before installing or upgrading to VMware vSphere Hypervisor (ESXi) 7.0.

- Upgrading from VMware vSphere Hypervisor (ESXi) 6.5 to 7.0:

  Uninstall device drivers for Intel Ethernet Adapters from VMware vSphere Hypervisor (ESXi) 6.5 host prior to starting VMware vSphere Hypervisor (ESXi) 7.0 upgrade process (failing to do so causes device drivers to stop loading on VMware vSphere Hypervisor (ESXi) 7.0). Upgrade to VMware vSphere Hypervisor (ESXi) 7.0. Install device drivers compiled with VMware vSphere Hypervisor (ESXi) 7.0 DDK.

  To download device drivers for VMware vSphere Hypervisor (ESXi) 7.0, visit the VMware VCG download site at:

https://www.vmware.com/resources/compatibility/search.php?deviceCategory=io.

- SR-IOV Virtual Function (VF) Creation:

  VMware vSphere Hypervisor (ESXi) 7.0 WebGUI allows users to instantiate VFs for each Network Adapter port. Creating VFs using VMware vSphere Hypervisor (ESXi) 7.0 WebGUI triggers immediate device driver reload that removes other device driver settings, like LLDP, RSS, VMDQ, etc. that might be enabled. This might cause loss of network connectivity. The device driver might fail to reload if SR-IOV VFs are configured and at least one VF is assigned to an active VM. Reboot the server after creating VFs to avoid this scenario. VMware vSphere Hypervisor (ESXi) 7.0 ignores VF creation using "max_vfs" module parameter if the VFs are created using VMware vSphere Hypervisor (ESXi) 7.0 WebGUI.

- Receive Side Scaling (RSS)
  To enable parallel traffic processing, RSS allows the driver to spread ingress network traffic over multiple receive queues associated with individual CPUs. RSS is enabled by default. RSS can be managed using the "esxcli system module parameters" command with RSS, DRSS and DevRSS parameters. If the admin configure only DRSS â€" then RSS is disabled despite having higher priority. If the admin tries to configure RSS and DRSS at the same time â€" then NetQ RSS is given precedence over DefQ RSS since those flavors are mutually exclusive. Once DevRSS is set, the following RSS connected config is ignored and virtualization (VMDQ / SR-IOV) is disabled on the PF. If any of RSS connected parameters is not enabled directly (default load) - then the RSS is enabled.

  NOTE: Reboot is needed after setting RSS mode.

- Recovery Mode
  A device will enter recovery mode if a device's NVM becomes corrupted. If a device enters recovery modeÂ because of an interrupted NVM update, you should attempt to finish the update.
  If the device is in recovery mode because of a corrupted NVM, use the nvmupdate utility to reset the NVM back to factory defaults.

  NOTE: You must power cycle your system after using Recovery Mode to completelyÂ reset the firmware and hardware.

- Backplane devices
  Backplane devices are operating in auto mode only, and thus the user cannot manually change speed settings.

Supported Features:

- Rx, Tx, TSO checksum offload
- Netqueue (VMDQ)
- Default Queue, NetQueue and Device RSS
- Hardware VLAN filtering
- Rx Hardware VLAN stripping
- Tx Hardware VLAN inserting

- Interrupt moderation
- SR-IOV (support one queue per VF, VF MTU, and VF VLAN)
   Valid range for max_vfs:
   1-61 (VMDQ default)
   1-63 (VMDQ set to 0 or 1)
   When VMDQ is set to 2 or more, the maximum number of VFs supported is 63 â€“ VMDQ.
- VMDQ and SR-IOV co-existence
- Link Auto-negotiation
- Flow Control
- Management APIs for CIM Provider, OCSD/OCBB
- Firmware recovery mode.
- Wake on LAN (WoL) support.
- Support for handling PCIe link failure or surprise removal of LAN controllers.

New Features:

   - Default Queue, NetQueue and Device RSS

New Hardware Supported:

   - None

Bug Fixes:

   - None

Known Issues and Workarounds:

- When user adds guest OS tagged and untagged adapters to a portgroup set in VLAN trunk mode (4095), untagged traffic is dropped.
   Workaround: None

- User is not able to create 63 VFs for one port on the adapter, 62 are available.
   Workaround: None

- Link flapping after powering off the VM or during vMotion.
   Workaround: use module parameter to disable the driver workaround for a kernel issue (lost VLAN connectivity once the last VM in a port group has been shut down).
   The kernel bug has been fixed in ESXi 6.5 U1 (and later releases), in which case the workaround can be safely disabled. The workaround is already disabled in builds for ESXi 6.5.

- On 82599 adapters emulated interfaces in a VLAN ALL portgroup may not be able to communicate with VF interfaces of the same PF. 82599 adapters lack a HW feature that

allows the driver to configure VLAN ALL mode for loopback traffic in the emulated data path.
   Workaround: none

- Incoming VLAN traffic is dropped after enabling software emulation of VLAN tagging and untagging for given PF. This has been introduced intentionally to drop VLAN tagged packets which do not have an active destination endpoint.
   Workaround: none

- Very low throughput when sending IPv6 to a Linux VM that uses a VMXNET3 adapter
   Workaround: Please look at the VMware Knowledge Base 2057874

- The VF VM using ixgbevf driver may crash if the user is constantly changing the MTU size for an adapter that is being used for a heavy traffic. The issue was be observed when decreasing the MTU from higher values to the default 1500B.
   Workaround: Disable the "large buffers" feature on the VF driver by issuing the below command on the guest OS: "ethtool --set-priv-flags eth0 legacy-rx on"

- In certain conditions VF driver may receive incorrect state of the PF. It may also result with a lost frame before executing a command or a lost answer for the command. The issue is visible in stress test setups with many operations on the VMs. In such environment the VF reset procedure may fail causing the traffic to break for some amount of time or until the next VF reset is performed. The user may see the information about TX hanging in the system logs.
   Workaround: none

Command Line Parameters:
----

ethtool is not supported for native driver.
Please use esxcli, vsish, or esxcfg-* to set or get the driver information, for example:

- Get the driver supported module parameters
  esxcli system module parameters list -m ixgben

- Set a driver module parameter (clearing other parameter settings)
  esxcli system module parameters set -m ixgben -p RSS=0,1

- Set a driver module parameter (other parameter settings left unchanged)
  esxcli system module parameters set -m ixgben -a -p RSS=0,1

- Get the driver info
  esxcli network nic get -n vmnic1

- Get an uplink stats
  esxcli network nic stats -n vmnic1

- Get the Netqueue stats
  vsish -e get /net/pNics/vmnic1/stats

- Disable kernel VLAN issue workaround
  esxcli system module parameters set -m ixgben -a -p "VlanRemoveWorkaround=0,0"

==========

Previously Released Versions:

- Driver Version: 1.8.7
   Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
   Supported ESXi releases: 6.0 and 6.7
   Compatible ESXi version: 6.5
   New Features Supported:
      - None
   New Hardware Supported:
      - None
   Bug fixes:
      - Added priority field to VLAN header
       - Added notification of link partner when Link Flow Control settings are changed
       - Reduced initialization time after uplink is set down then up

- Driver Version: 1.7.20
   Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
   Supported ESXi releases: 6.0 and 6.7
   Compatible ESXi version: 6.5
   New Features Supported:
      - None
   New Hardware Supported:
      - None
   Bug fixes:
      - Fixed high CPU usage when SFP+ module is not inserted in the NIC.

- Driver Version: 1.7.17
   Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
   Supported ESXi releases: 6.0 and 6.7
   Compatible ESXi version: 6.5
   New Features Supported:
      - Added Wake on LAN (WoL) support.
   New Hardware Supported:
      - None
   Bug Fixes:
      - Fixed driver version reporting as unavailable in iDRAC.
      - Fixed incorrect branding strings for specific supported devices.
   Known Issues:
      - If the VF's guest VLAN interface and the VF's portgroup have the same VLAN ID, packets appear on the VF's guest VLAN interface for 82599 and x540 adapters due to the HW limitation.
         Workaround: none
      - On 82599 adapters emulated interfaces in a VLAN ALL portgroup may not be able to communicate with VF interfaces of the same PF. 82599 adapters lack a HW feature

that allows the driver to configure VLAN ALL mode for loopback traffic in the emulated data path.
Workaround: none
- Incoming VLAN traffic is dropped after enabling software emulation of VLAN tagging and untagging for given PF. This has been introduced intentionally to drop VLAN tagged packets which do not have an active destination endpoint.
Workaround: none
- Very low throughput when sending IPv6 to a Linux VM that uses a VMXNET3 adapter
Workaround: Please look at the VMware Knowledge Base 2057874

- Driver Version: 1.7.15
Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
Supported ESXi releases: 6.0 and 6.7
Compatible ESXi version: 6.5
New Features Supported:
- None
New Hardware Supported:
- None
Bug Fixes:
- Fixed VLAN tagged packets accepted after turning off last VM
- Fixed VF connection issues when MAC address is changed
- Fixed missing OROM version in 'esxcli network nic get -n <vmnic>' output
- Fixed untagged packets being received by a VF in a VLAN-tagged portgroup
- Fixed VF being able to send VLAN-tagged packets despite being in a untagged portgroup
- Fixed intermittent Tx hang due to a race condition
Known issues:
- If the VF's guest VLAN interface and the VF's portgroup have the same VLAN ID, packets appear on the VF's guest VLAN interface for 82599 and x540 adapters due to the HW limitation.
Workaround: none
- On 82599 adapters emulated interfaces in a VLAN ALL portgroup may not be able to communicate with VF interfaces of the same PF. 82599 adapters lack a HW feature that allows the driver to configure VLAN ALL mode for loopback traffic in the emulated data path.
Workaround: none
- Incoming VLAN traffic is dropped after enabling software emulation of VLAN tagging and untagging for given PF. This has been introduced intentionally to drop VLAN tagged packets which do not have an active destination endpoint.
Workaround: none
- Very low throughput when sending IPv6 to a Linux VM that uses a VMXNET3 adapter
Workaround: Please look at the VMware Knowledge Base 2057874

- Driver Version: 1.7.10
Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
Supported ESXi releases: 6.0 and 6.7
Compatible ESXi version: 6.5
New Features Supported:

- Added firmware recovery mode.
   New Hardware Supported:
      - None
   Bug Fixes:
      - Fixed traffic hang between VM with VF adapter and VM with vmxnet3 after disabling/enabling vmnic.
      - Fixed an issue with VM's being able to communicate over VLANs in VGT mode when Port Group had VLANs set to None(0)
      - Fix for silicon errata #26: TX hang observed on some queues during regular traffic with VFLR on the fly.
  Please see X550 specification update for more information.
      - Fix for dropped tagged loopback traffic originated from VF.
      - Fix for lost VLAN connectivity once the last VM in a port group has been shut down.
      - Reduced driver's memory footprint
      - Fix for duplicate multicast / broadcast packets during heavy traffic.
      - Fix for X552/X557-AT adapters linking at 1G speed after a NIC down/up cycle.
      - Fix for the NIC down procedure hanging when heavy traffic is running.
      - Fix for lost connectivity when VMDQ is disabled and SR-IOV is enabled
   Known Issues:
      - Unable to reload VF driver on SLES 12SP2 and ESXi 6.0 update 3
         Workaround: upgrade to ESXi 6.0 update 3a or ESXi 6.5
      - Very low throughput when sending IPv6 to a Linux VM that uses a VMXNET3 adapter
         Workaround: Please look at the VMware Knowledge Base 2057874

- Driver Version: 1.6.5
   Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
   Supported ESXi release: 6.0
   Compatible ESXi version: 6.5 and 6.7
   New Features Supported:
      - None
   New Hardware Supported:
      - Added new devices support for specific OEMs
   Bug Fixes:
      - Add a VF hang detection and take an appropriate action to prevent infinite reset loop
      - Show proper link speed using dual speed optics module connected to 1G link partner
      - Fix the issue where the driver may fail to load after an unexpected system reboot
      - Fix for Multiple Link Down messages after unplug a cable
      - Fix VF link status where it was not updated when PF link state changed
   Known Issues:
      - Unable to reload VF driver on SLES 12SP2 and ESXi 6.0 update 3
         Workaround: upgrade to ESXi 6.0 update 3a or ESXi 6.5

- Driver Version: 1.5.3
   Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
   Supported ESXi release: 6.0

Compatible ESXi version: 6.5 and 6.7
New Features Supported:
- Added support for OEM firmware version. If OEM specific adapter is detected, OEM product version in firmware
version string is reported instead of etrack id.
- Added ability to get/set Tx/Rx ring buffer size using esxcli command
New Hardware Supported:
- Added new devices support for specific OEMs
Bug Fixes:
- Fixed possible Tx/Rx queue race condition when there are heavy resets followed by link down/up cycles, which lead to an IOMMU fault
- Fixed Tx hang issue when enabling 64 VFs
- Fixed vmknic and VF communication issue when both are configured using the same PF port
- Fixed zero temperature value for hardware that supports temperature sensors
Known Issues:
- None

- Driver Version: 1.4.1
Hardware Supported: Intel(R) Ethernet Controllers 82599, x520, x540, x550, and x552 family
Supported ESXi release: 6.0
Compatible ESXi version: 6.5 and 6.7
Features Supported:
- Rx, Tx, TSO checksum offload
- Netqueue (VMDQ)
- SR-IOV (support one queue per VF, VF MTU, and VF VLAN)
- VMDQ and SR-IOV co-existence
- Netqueue UDP RSS for VxLAN
- Hardware VLAN filtering
- Rx Hardware VLAN stripping
- Tx Hardware VLAN inserting
- Link auto-negotiation and link flow control
- Interrupt moderation
- Management APIs for CIM Provider, OCSD/OCBB
New Hardware Supported:
- Added new devices support for specific OEMs
Bug Fixes:
- None
Known Issues:
- There is no traffic when vmknic and VF are configured using the same PF port
Workaround: none
- Enabling 64 VFs results in TX hang condition
Workaround: Reduce the number of VFs to 63 or lower