

# vNIC Configuration Guide

---

Release, August 2011



2051 Mission College Blvd.  
Santa Clara, CA 95054  
[www.bladenetwork.net](http://www.bladenetwork.net)

Copyright © 2011 BLADE Network Technologies, an IBM company, 2051 Mission College Blvd., Santa Clara, California, 95054, USA. All rights reserved. Part Number Release.

This document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of BLADE Network Technologies, Inc. Documentation is provided “as is” without warranty of any kind, either express or implied, including any kind of implied or express warranty of non-infringement or the implied warranties of merchantability or fitness for a particular purpose.

U.S. Government End Users: This document is provided with a “commercial item” as defined by FAR 2.101 (Oct. 1995) and contains “commercial technical data” and “commercial software documentation” as those terms are used in FAR 12.211-12.212 (Oct. 1995). Government End Users are authorized to use this documentation only in accordance with those rights and restrictions set forth herein, consistent with FAR 12.211- 12.212 (Oct. 1995), DFARS 227.7202 (JUN 1995) and DFARS 252.227-7015 (Nov. 1995).

BLADE Network Technologies reserves the right to change any products described herein at any time, and without notice. BLADE Network Technologies assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by BLADE Network Technologies. The use and purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of BLADE Network Technologies.

BLADE Network Technologies, the BLADE logo, BLADEHarmony, BNT, NMotion, RackSwitch, Rackonomics, RackSwitch Solution Partner, ServerMobility, SmartConnect and VMready are trademarks of BLADE Network Technologies in the United States and certain other countries. Cisco® and EtherChannel® are registered trademarks of Cisco Systems, Inc. in the United States and certain other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. Any other trademarks appearing in this manual are owned by their respective companies.

Originated in the USA.

# vNIC Configuration Guide

---

## vNIC Enhancements

---

### Global Configuration Control of vNIC Uplink Sharing

When uplink sharing is disabled (dedicated uplink mode), which is the current implementation, there is only one vNIC group using the uplink. The switch mandatorily inserts the vNIC group VLAN tag for ingress packets and removing the vNIC group VLAN tag for egress packets on the uplink port.

When uplink sharing is enabled (shared uplink mode), there can be multiple vNIC groups using the same uplink. The switch will let the uplink switch (another upstream switch connected to the uplink) handle the VLAN tags, similar to what vNIC ports are doing – let NIC handle the VLAN tags. To support multiple vNIC groups using the same uplink port, the “tag” configuration must be enabled on the uplink port.

---

**Note** – In shared mode, by leaving “tagpvid” and “pvid” user configurable for uplink ports, the user has the chance to carry traffic for one vNIC group in untagged packets. The user can do this by configuring the uplink port “pvid” to one of the vNIC group VLANs and disable “tagpvid”. This is different from vNIC ports which require that all packets coming in from vNICs must carry their vNIC group VLAN explicitly as the outer tag.

---

**Note** – The shared mode allows sharing the same uplink with multiple vNIC groups. But the user is still allowed to configure just one vNIC group on an uplink even though the shared mode is enabled. Even if there is only one vNIC group using an uplink port, whether shared mode or dedicated mode is used can have some different implications on the use of VLAN tags as described above.

---

The ISCLI commands will be under configure terminal mode:

```
[no] vnic uplink-share
```

## Per vNIC-Group Configuration Control of LACP Trunk Uplink

For each vNIC group, the administrator can now add or remove an LACP trunk as the uplink. The ISCLI commands will be under vNIC group configuration mode (vnic vnicgroup x):

- key <admin-key>     Add the uplink LACP trunk to the vNIC group
- no key                 Remove the uplink LACP trunk from the vNIC group

## Topology Used




Figure 1: Topology Used

## Areas Covered for Enhancements Validation

---

- Adding an LACP key as uplink in a vNIC group  
An LACP key was added as an uplink to a vNIC group. The LACP was defined and aggregated prior to adding it to the vNIC group.
- Adding the same LACP key as uplink to multiple vNIC groups  
The same LACP key was defined as the uplink to 3 vNIC groups. Prior to configuring this, the LACP was defined and aggregated. Also uplink share option was enabled on the switch side
- Removing an LACP key from a vNIC group  
An LACP key previously added to multiple vNIC groups was removed from the vNIC group. The LACP will still be defined as the uplink for the other vNIC groups.
- Sending ping/traffic between the vNICs in the same vNIC group  
Ping and traffic tests were performed to verify connectivity between vNICs in the same group
- Sending ping/traffic between the vNICs and uplink servers  
Using the topology describe above, ping and traffic tests were performed between the vNIC servers and the uplink server. Proper VLAN configuration needs to be performed on the uplink switch in order to allow traffic.
- Enable/disable uplink share  
Tried to toggle uplink share option while multiple vNIC groups were sharing the same uplink. Same test was performed when each vNIC group had its own uplink. When multiple vNIC groups are sharing an uplink, the user is not allowed to disable uplink share option.
- Copy/retrieve vNIC configuration from TFTP server  
This test is to ensure that configuration is not corrupted when transferring the configuration from/to a TFTP server
- Changing of the vNIC group VLAN  
Changing the vNIC group VLAN is properly propagated on the server side. The uplink switches/servers need to be reconfigured to reflect the new VLAN configuration.
- Failover and failback on the vNIC group when LACP is used as uplink  
Failover and failback tests were performed using the topology described above and an LACP key defined as the uplink. The LACP was created and aggregated prior to adding it as the uplink. Also, failover was enabled on the vNIC group. Upon, uplink failure a DCBX vlink down is sent to all the vNICs in the vNIC group. At failback, the switch will send vlink up on all vNICs from the vNIC group.

- Failover and failback of vNICs when single LACP is used as uplink for multiple vNIC groups

Same test as above was performed by adding the same LACP key to 3 vNIC groups. The LACP was created and aggregated prior to adding it as the uplink. Also, failover was enabled on all vNIC groups. Upon, uplink failure a DCBX vlink down is sent to all the vNICs in all vNIC groups sharing the uplink. At failback, the switch will send vlink up on all vNICs from the vNIC groups with the shared uplink.

- Failure of one of the LACP links has no effect on the traffic

After adding a pre-configured LACP key to a vNIC group, some of the LACP ports are toggled (not all). In this case, traffic should be load balanced across all the remaining members of the LACP portchannel without interrupting the traffic flow.

## Best Practices And Recommendations

---

We recommend the following as the best practices for configuring this solution:

- When adding a port to a LACP portchannel, the user must make sure the port will have the same settings as the ports already in the LACP portchannel. For example, the settings generally include “pvid”, “tag”, “tagpvid”, VLAN membership, STP state, etc. If a certain port setting is automatically managed by switch software under certain scenarios, then the user doesn’t need to worry about it. For example, in uplink shared mode, vNIC uplink “tag” is automatically enabled by switch software. So the user does not need to worry about “tag” configuration of the port. For another example, in uplink shared mode, the setting “pvid” is user-configurable for uplink ports. When adding a port to an existing LACP portchannel, the user must change the “pvid” of the port to the same as the other ports in the LACP portchannel if they are different.
- The STP states of vNIC uplink ports are user-configurable. Since most users will not need STP on the vNIC uplink ports at the edge of the network, switch software will automatically turn off the STP state of uplink ports when added to a vNIC group (see note below). After that, the user can choose to turn on the STP state manually for the uplink ports if necessary. Similarly, switch software will automatically turn on STP state for uplink ports removed from a vNIC group.

---

**Note – Note 1:** STP is only disabled on uplink LACP if the LAG is established before the key is added to the vNIC group. By default, STP is disabled on the ports being added as uplinks in a vNIC group. For LACP, STP will only be turned off if the LACP portchannel is created prior to adding the key to the vNIC group. Otherwise, STP is not turned off. In this case, if the user wants to have STP off on these ports, he/she can either manually turn off STP on these ports or create the LACP portchannel before adding the key to the vNIC group. (In GA release, STP will be automatically turned off for uplink ports even if they are added to the LACP portchannel after the key is added to the vNIC group.)

---

## Known Issues And Workarounds

---

### Members cannot be added to LACP portchannel after adding the portchannel to the vNIC group

#### Problem:

If the user attempts to add a port to an LACP portchannel after the portchannel has already been added to the vNIC group, an STP mismatch will be displayed.

#### Workaround:

In order to add the port to the portchannel the user needs to do the following:

- Disable STP on the corresponding STGs (STGs that are assigned to the vNIC VLANs) interface port <x>  

```
no spanning-tree stp <y> enable
```
- Change the PVID of the LACP ports to 1 (if PVID is not 1)
- Add the new port to the LACP portchannel (by changing the LACP key to match the portchannel's)
- Change the PVID of the LACP ports back to the desired value

### STP is only disabled on uplink LACP if the LAG is established before the key is added to the vNIC group

#### Problem:

By default, STP is disabled on the ports being added as uplinks in a vNIC group. For LACP, STP will only be turned off if the LACP portchannel is created prior to adding the key to the vNIC group. Otherwise, STP is not turned off.

#### Workaround:

Create the LACP portchannel prior to adding it to the vNIC groups or manually disable STP on the uplink ports.

## STP is not re-enabled on a port that is removed from an LACP assigned to a vNIC group

### Problem:

If the port is removed from an LACP portchannel that belongs to a vNIC group, the STP will not be re-enabled on the respective STG

### Workaround:

User can manually enable STP on that STG

```
interface port <x>
spanning-tree stp <y> enable
```

## STP is not disabled on the corresponding STG when changing the vNIC VLAN

### Problem:

When changing the vNIC VLAN, all uplink ports are re-assigned to the new VLAN. However, STP is not disabled on the new STG

### Workaround:

User can manually disable STP on the proper STG

## A user defined VLAN cannot be added as a vNIC VLAN

### Problem:

If a VLAN has been defined on the switch using “vlan <x>” or port “pvid <X>” command, it cannot be added as vNIC VLAN regardless of the ports that are assigned to this VLAN. This is the current vNIC implementation

### Workaround:

User needs to delete the VLAN using “no vlan <x>” command. After that, the VLAN can be added to the vNIC group