



# **NetXtreme/NetXtreme II RSS Indirection Table Configuration Guide on Microsoft® Windows Server 2012**

## Revision History

<i>Revision</i>	<i>Date</i>	<i>Change Description</i>
NX1_NX2_RSS-SWUM101-R	05/01/13	<b>Update:</b> <ul style="list-style-type: none"><li>• Removing references to Windows Server 2008 R2.</li></ul>
NX1_NX2_RSS-SWUM100-R	04/10/13	Initial release

Broadcom Corporation  
5300 California Avenue  
Irvine, CA 92617

© 2013 by Broadcom Corporation  
All rights reserved  
Printed in the U.S.A.

Broadcom®, the pulse logo, Connecting everything®, and the Connecting everything logo are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

## Table of Contents

<b>About This Document</b> .....	5
Purpose and Audience .....	5
Acronyms and Abbreviations.....	5
Document Conventions .....	5
<b>Technical Support</b> .....	5
<b>Requirements</b> .....	6
<b>Configuring the RSS Indirection Table</b> .....	7
<b>Summary</b> .....	14

## List of Figures

Figure 1: Get-NetAdapterRss Output .....	8
Figure 2: Get NetAdapterRss Command .....	9
Figure 3: Device Manager's Driver Key Values .....	12
Figure 4: RssBaseProcNumber and RssMaxProcNumber .....	13
Figure 5: CPU Core and Processor Numbers .....	14

---

## About This Document

### Purpose and Audience

This document provides instructions on how to configure Ethernet Networking Receive Side Scaling (RSS) Base and Maximum CPU Core assignment in the RSS Indirection table of the Broadcom® NetXtreme® (NX1) 5719/20 and NetXtreme II® (NX2) 1G 5709/5716 and NX2 10G 57712/578xxS device family.

This document is written for the network administrator who wishes to hand-tune the Broadcom network controller RSS Indirection table on Microsoft® Windows Server® 2012 multi-CPU core systems.

### Acronyms and Abbreviations

In most cases, acronyms and abbreviations are defined on first use.

For a comprehensive list of acronyms and other terms used in Broadcom documents, go to: <http://www.broadcom.com/press/glossary.php>.

### Document Conventions

The following conventions may be used in this document:

<i>Convention</i>	<i>Description</i>
<b>Bold</b>	User input and actions: for example, type <b>exit</b> , click <b>OK</b> , press <b>Alt+C</b>
Monospace	Code: <code>#include &lt;iostream&gt;</code> HTML: <code>&lt;td rowspan = 3&gt;</code> Command line commands and parameters: <code>w1 [-1] &lt;command&gt;</code>
< >	Placeholders for <i>required</i> elements: enter your <code>&lt;username&gt;</code> or <code>w1 &lt;command&gt;</code>
[ ]	Indicates <i>optional</i> command-line parameters: <code>w1 [-1]</code> Indicates bit and byte ranges (inclusive): <code>[0:3]</code> or <code>[7:0]</code>

---

## Technical Support

Broadcom provides customer access to a wide range of information, including technical documentation, schematic diagrams, product bill of materials, PCB layout information, and software updates through its customer support portal (<https://support.broadcom.com>). For a CSP account, contact your Sales or Engineering support representative.

In addition, Broadcom provides other product support through its Downloads and Support site (<http://www.broadcom.com/support/>).

## Requirements

To configure RSS-capable networking device's RSS Indirection tables on a Windows Server system, the following must be true:

- The server system must have multiple CPU cores.
- The server system must have Windows Server 2012 installed.
- The NX1 5719/20 and NX2 5709/16/57712/578xxS must have the latest release firmware and drivers installed from the OEM's web site.
- RSS must be enabled on the desired network adapter ports/partitions.
- Hyper Threading or Logical Processors should be disabled in the system BIOS.



**Note:** VMQ and RSS are incompatible with each other. Do not enable VMQs when RSS is enabled.

The system in these examples uses:

- Eight quad-core CPUs (total of 32 cores)
- Two dual 10GbE ported 57810 NICs
  - One in normal Single Function (SF) mode (networking connections one and two)
  - One in Switch Independent NIC Partitioning (NPAR) mode (networking connections three through ten).

The SF mode 57810 can use a maximum of 16 RSS queues per port. The NPAR mode 57810 can use a maximum of four RSS queues per partition (four partitions per port).

Windows PowerShell (version 3.0) is used to real-time configure (no rebooting) and monitor each individual network adapter ports/partitions RSS Indirection tables. The two commands primarily used are `Get-NetAdapterRss` and `Set-NetAdapterRss`.

## Configuring the RSS Indirection Table

The simplest way to configure a specific NX1/NX2 adapter port or 57712/578xx NPAR partition RSS Indirection table is to use the PowerShell `Set-NetAdapterRss` command's `BaseProcessorNumber` and `MaxProcessorNumber` parameters.

In the following example, there are 32 cores available. The first 57810 (in SF mode) will have each of its two ports (one and two) given their own four distinct RSS cores. The second 57810 (in NPAR mode) will have each of its eight NPAR partitions (four partitions per port, ports three through ten) given their own two distinct RSS cores. The first eight cores will not be assigned to an adapter so they can be used by the host OS for other computing processes.

In PowerShell (before making any changes) use the `Get-NetAdapterRss` command to see the range of each adapter's RSS Indirection table; they should initially be set from core 0 to core 31. The actual selected RSS cores assigned to each adapter will vary with each host, but they will overlap if the total number of receive queues (`NumberOfReceiveQueues`) values for all RSS adapters is greater than the number of CPU cores in the host system, which is the case in this example as it has  $2 \times 16$  (first two 57810 SF mode ports) plus  $8 \times 4$  (second two 57810 NPARed ports) for a total of 64 which is greater than the 32 available.

The `-ifdesc` parameter allows you to select which adapter to view by filtering on the string in the "InterfaceDescription" field. If no parameters are provided, then all RSS adapters are outputted. The registry `MaxProcessors` value (in this output) is not the `*MaxRssProcessors` value from the device driver INF and registry entry, instead it is the lower of the two `NumberOfReceiveQueues` values (from the driver INF or registry `*NumRssQueues`) and the difference between `BaseProcessor Number` (from the registry `*RssBaseProcNumber` if present, if not present it defaults to 0) and `MaxProcessor Number` (from the registry `*RssMaxProcNumber` if present; if not present, it defaults to the highest CPU core number). The following displays the sixteen CPU cores assigned to this specific port's RSS indirection table (0-1, 4, 6, 10-12, 15-17, and 22-27). These CPU cores are the ones that are used for all RSS/TSS of traffic streams on this port.

```
Get-NetAdapterRss -ifdesc *#1*
Name:Ethernet 1
InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #1
Enabled: True
NumberOfReceiveQueues: 16
Profile: NUMAStatic
BaseProcessor: [Group:Number]: 0:0
MaxProcessor: [Group:Number]: 0:31
MaxProcessors: 16
RssProcessorArray: [Group:Number/NUMA Distance] : 0:0/0 0:1/0 0:2/0 0:3/0 0:4/0 0:5/0 0:6/
0 0:7/0 0:8/0 0:9/0 0:10/0 0:11/0 0:12/0 0:13/0 0:14/0 0:15/0 0:16/0 0:17/0 0:18/0
0:19/0 0:20/0 0:21/0 0:22/0 0:23/0 0:24/0 0:25/0 0:26/0 0:27/0 0:28/0 0:29/0 0:30/0
0:31/0
```





Get-NetAdapterRss  
 Name: Ethernet 1  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #1  
 Enabled: True  
 NumberOfReceiveQueues: 16  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:8  
 MaxProcessor: [Group:Number]: 0:11  
 MaxProcessors: 4  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:8/0 0:9/0 0:10/0 0:11/0  
 IndirectionTable: [Group:Number]: 0:80:90:100:11

Name: Ethernet 2  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #2  
 Enabled: True  
 NumberOfReceiveQueues: 16  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:12  
 MaxProcessor: [Group:Number]: 0:15  
 MaxProcessors: 4  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:12/0 0:13/0 0:14/0 0:15/0  
 IndirectionTable: [Group:Number]: 0:120:130:140:15

Name: Ethernet 3  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #3  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:16  
 MaxProcessor: [Group:Number]: 0:17  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:16/0 0:17/0  
 IndirectionTable: [Group:Number]: 0:160:17

Name: Ethernet 4  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #4  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:18  
 MaxProcessor: [Group:Number]: 0:19  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:18/0 0:19/0  
 IndirectionTable: [Group:Number]: 0:180:19

Name: Ethernet 5  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #5  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:20  
 MaxProcessor: [Group:Number]: 0:21  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:20/0 0:21/0  
 IndirectionTable: [Group:Number]: 0:200:21

Name: Ethernet 6  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #6  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:22  
 MaxProcessor: [Group:Number]: 0:23  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:22/0 0:23/0  
 IndirectionTable: [Group:Number]: 0:220:23

Name: Ethernet 7  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #7  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:24  
 MaxProcessor: [Group:Number]: 0:25  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:24/0 0:25/0  
 IndirectionTable: [Group:Number]: 0:240:25

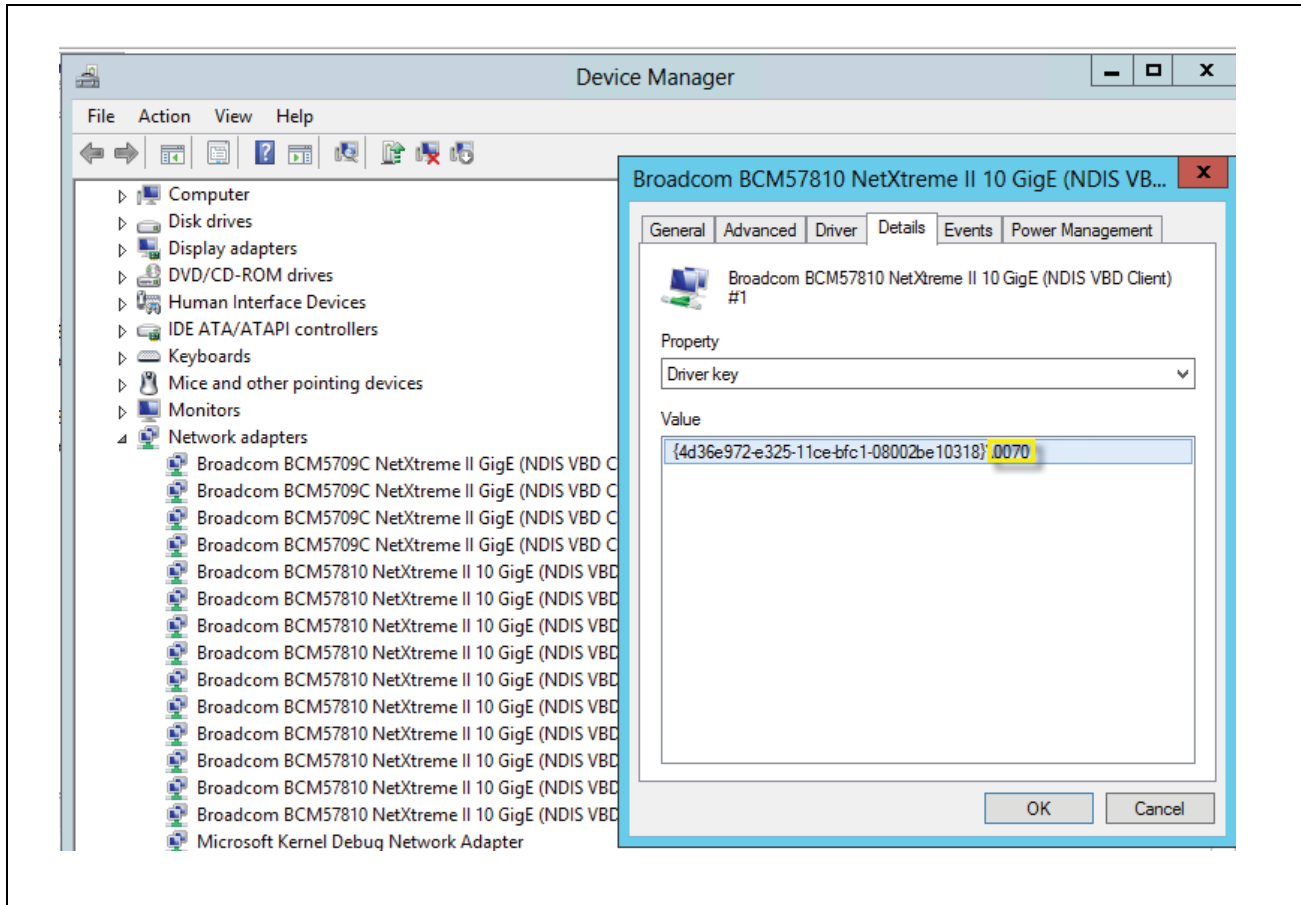
Name: Ethernet 8  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #8  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:26  
 MaxProcessor: [Group:Number]: 0:27  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:26/0 0:27/0  
 IndirectionTable: [Group:Number]: 0:260:27

Name: Ethernet 9  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #9  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:28  
 MaxProcessor: [Group:Number]: 0:29  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:18/0 0:19/0  
 IndirectionTable: [Group:Number]: 0:280:29

Name: Ethernet 10  
 InterfaceDescription: Broadcom BCM57810 NetXtreme II 10 GigE (NDIS VBD Client) #10  
 Enabled: True  
 NumberOfReceiveQueues: 4  
 Profile: NUMAStatic  
 BaseProcessor: [Group:Number]: 0:30  
 MaxProcessor: [Group:Number]: 0:31  
 MaxProcessors: 2  
 RssProcessorArray: [Group:Number/NUMA Distance]: 0:30/0 0:31/0  
 IndirectionTable: [Group:Number]: 0:300:31

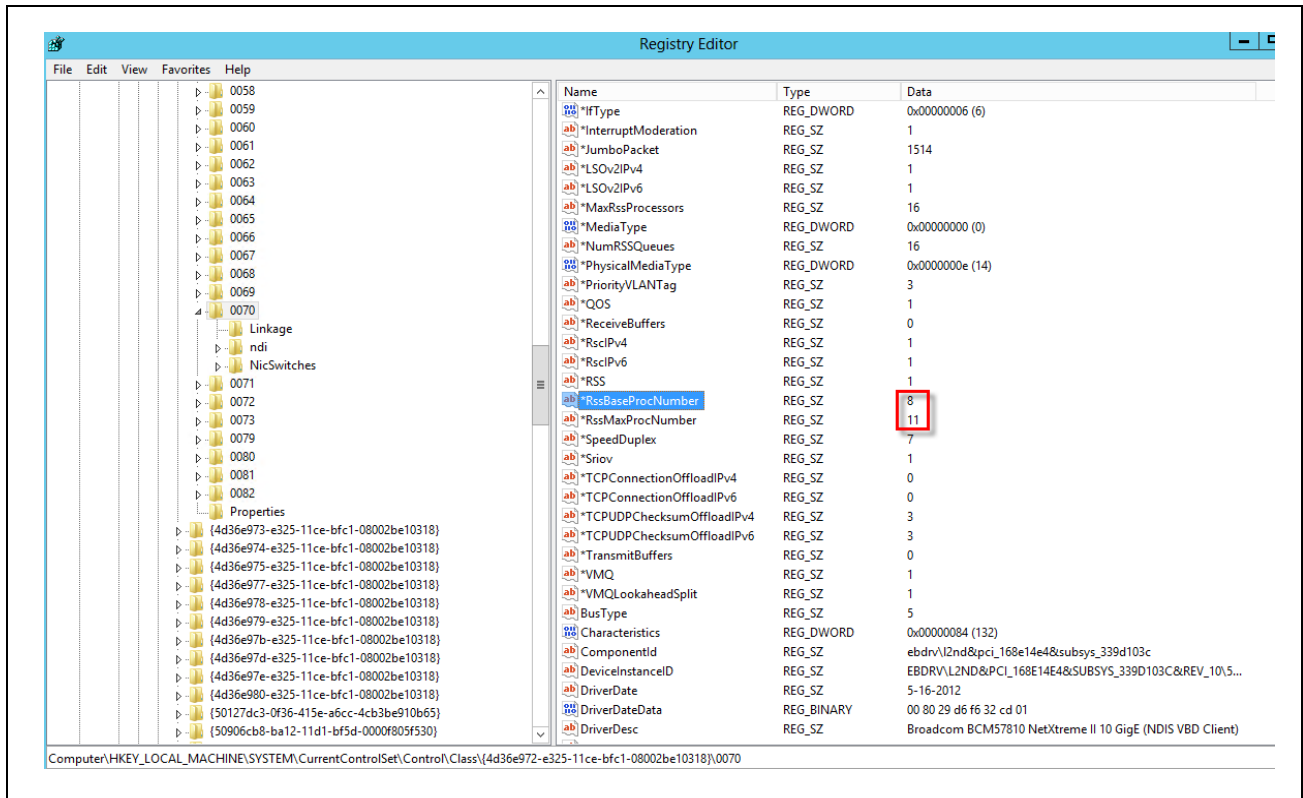
These changes are assigned instantly and no host system reboot is needed. The values can also be seen in each adapter's registry entry. To find a specific adapters registry index value, use the Windows Device Manager's Properties Details tab Driver Key values. The first port's index value is 0070.

**Figure 3: Device Manager's Driver Key Values**



In the registry ...\\HKEY\_LOCAL\_MACHINE\\SYSTEM\\CurrentControlSet\\Control\\Class\\{4d36e972-e325-11ce-bfc1-08002be10318}\\0070 entry, the newly created keyword RssBaseProcNumber has the value of eight and RssMaxProcNumber has the value of eleven.

Figure 4: RssBaseProcNumber and RssMaxProcNumber



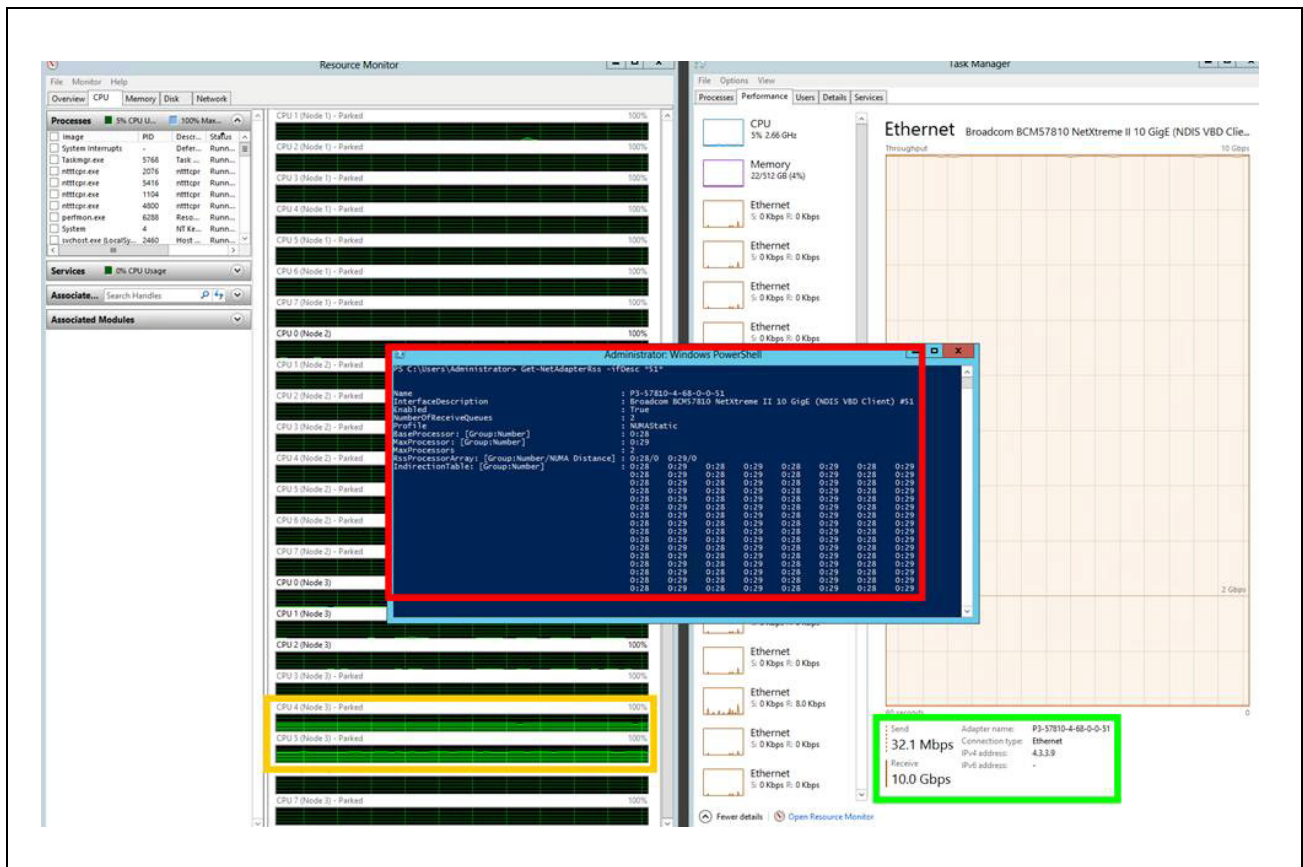
Further verify these cores get the received traffic assigned to them by receiving traffic on the designated port/partition (using a TCP test traffic-generating application such as NTTTC or iPerf) and opening a resource monitor window and observing which CPU cores are in use when the traffic starts and stops. In this system, there are four separate CPUs and each has eight cores. The numbering in the resource monitor has the CPU core number followed by the processor node number.

1. CPU 0 (Node 0)
2. CPU 1 (Node 0)
3. CPU 2 (Node 0)
4. CPU 3 (Node 0)
5. CPU 4 (Node 0)
6. CPU 5 (Node 0)
7. CPU 6 (Node 0)
8. CPU 7 (Node 0)
9. CPU 0 (Node 1)
10. CPU 1 (Node 1)

...

- 29. CPU 4 (Node 3)
- 30. CPU 5 (Node 3)
- 31. CPU 6 (Node 3)
- 32. CPU 7 (Node 3)

Figure 5: CPU Core and Processor Numbers



## Summary

The RSS indirection table of each RSS capable network adapter on a system can be reconfigured using Microsoft's PowerShell (version 3.0) to fine-tune a system's CPU core to networking adapter connection utilization and to allow host processes (such as SQL Server) to have undivided usage of the lower starting CPU cores.

Broadcom® Corporation reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom Corporation is believed to be accurate and reliable. However, Broadcom Corporation does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Connecting  
**everything®**



**BROADCOM CORPORATION**

5300 California Avenue

Irvine, CA 92617

© 2013 by BROADCOM CORPORATION. All rights reserved.

NX1\_NX2\_RSS-SWUM101-R May 1, 2013

Phone: 949-926-5000

Fax: 949-926-5203

E-mail: [info@broadcom.com](mailto:info@broadcom.com)

Web: [www.broadcom.com](http://www.broadcom.com)