



IBUtils2 Utility Release Notes v2.16

Table of Contents

1	Release Notes Update History	3
2	Overview	4
2.1	Software Download	4
2.2	Packages.....	4
3	Changes and New Features	5
3.1	Changes in Command Line Parameters	8
4	Bug Fixes	12
5	Legal Notices and 3rd Party Licenses	13

1 Release Notes Update History

Revision	Date	Description
2.16	February 08, 2024	Initial release of this Release Notes version. This version introduces Changes and New Features .

2 Overview

ibdiagnet scans the fabric using directed route packets and extracts all the available information regarding its connectivity and devices. It then produces the following files in the output directory.

2.1 Software Download

Please visit [InfiniBand Management Tools](#) page.

2.2 Packages

Package	Version
MLNX_OFED	24.01-0.3.3.1
UFM	6.16.x

3 Changes and New Features

v.2.16.0	
PHY Plugin	Added support for MRCS register.
PHY Plugin	Added support for PRTL register.
PHY Plugin	Updated PDDR register.
Cables	Reduced false-positive errors about an invalid port in the cable plugin for splitted switches.
Cables	Changed cable info fetch priority from cable plugin to PHY plugin.
v2.15.0	
Counters	Added support for PPCNT InfiniBand general counters.
DGX-H100	Added support for DGX-H100 systems in Rail Validation
db_csv File	Added the option to report to the db_csv file if no BER threshold is found for a port.
SLRG Register	Excluded the SLRG register from the default PHY list.
MSPS Register	Updated MSPS register (added the <code>power_consumption</code> field to the dump).
v2.14.0	
PCIe Connectivity Health Report	Introduced a new report on PCIe connectivity health based on a comparison of enabled/active speed and width.
Cables	Removed validation of transceiver firmware versions on the same cable.
v2.13.0	
Counters	Added support for "Fast Link Recovery" counters.
Counters	Added support for On-Demand-Paging (ODP) counters.
Reports	Added reports for mismatch in cable firmware versions
Reports	Added reports for cable length.
v2.12.0	
Programable Port Congestion Control	Added support for fetching Programable Port Congestion Control Counters (PPCC) at GA level.
FLID	Added support for FLID at GA level.
SHARP	Added support for "ad-hoc" trees for SHARP.
iblinkinfo	Added new dump file in format if iblinkinfo.
v2.11.0	
Programable Port Congestion Control	Added support for fetching Programable Port Congestion Control Counters (PPCC). This feature is at Technical Preview level.
FLID	Added support for FLID. This feature is at Technical Preview level.

GMP MADs	Added support for sending GMP MADs from user-space (Kernel bypass). This feature is at Alpha level.
In-band auto-discovery	In-band auto-discovery for switch management IP. This feature is available only in the UFM package.
PHY plugin	Added support for IB cable/transceiver diagnostics: SNR and EOM (PEMI).
ibdiagnet's runtime statistics	Added ibdiagnet's runtime statistics (CPU utilization, Stage duration, MAD statistics).
v2.10.0	
Cable Information	Fetching cable info from Common Management Interface Specification (CMIS) compatible cables connected to HCA.
smparquery Utility	The smparquery utility is moved from the deprecated package "ar_info" to the "ibutils2" package. Now, location of the binary file is "/usr/bin/smparquery" , instead of "/usr/sbin/smparquery".
topodiff Tool	The topodiff tool supports now more than 4 HCA systems.
Routers stage	"Routers stage" will be launched by default.
Auto-detection of Outgoing IB Port	Improved the auto-detection of outgoing IB port.
v2.9.0	
Proactive FRN (pFRN)	Added support for Proactive FRN (pFRN) configuration and counters.
HashBasedForwarding	Added support for HashBasedForwarding (HBF).
PortVLXmitWait	Added support for vendor spec PortVLXmitWait 64bit.
Hierarchy Info	Added support for Hierarchy Info (A15 InfiniBand spec): <ul style="list-style-type: none"> • Port Hierarchy Info • Physical Hierarchy Info
PHY plugin	Added support for up to 16 fans (START_FANS_SPEED section).
SHARP trees validation	Added new validations for SHARP trees in fat-tree topology with "parallel" links between switches.
VL Arbitration	Added support for VL Arbitration.
Cables	Modified the format for temperature thresholds. How it uses human readable format.
PMPortSamplesControl	Added PMPortSamplesControl to db_csv.
FDBS file	Enabled FDBS file creation by default.
DFP dump file	Improved DFP dump file.
Dump file	Performance improvements in dump file creation
v2.8.0	
Cables	Added support for CMIS cable.
PortVLXmit	Added support for PortVLXmitFlowCtlUpdateErrors.
PortVLXmit	Added support for PortVLXmitWait.

PHY plugin	Updated SLRG_16 and SLTP_16 registers.
PHY plugin	Added support for PCI Diagnostic Data Pages.
Adapter Cards: Socket-Direct	Added a list of Socket-Direct HCAs to DB_CSV.
Adapter Cards: Socket-Direct	Socket-Direct HCAs are excluded from Rail-Optimized Topology validation.
PM Stage Reports Overflow & Threshold	PM Stage reports overflow & threshold for 3 counters only: <ul style="list-style-type: none"> • SyncHeaderErrorCounter • UnknownBlockCounter • port_fec_uncorrectable_block_counter
Fat Tree Validation	Improved "Connectivity group" detection.
ibnetdiscover	ibnetdiscover file will not be generate in case of the scope file usage.
Fabric Summary	Fabric Summary now includes number of Socket-Direct HCAs
Error/Warning Reporting	Error/warning reporting to screen and log files is limited by 5 of each type, DB_CSV includes all errors
v2.7.0	
Link Speed	Added support for NDR InfiniBand link speed.
Fat Tree Topology	Enabled a new Fat Tree Topology validation tool.
Virtualization Stage	Redesigned the Virtualization stage to asynchronous mode.
AGUID	AGUID stage is disabled by default. To enable it, use the '--aguid' parameter.
db_csv	db_csv will now contain information about connected ports only.
v2.6.1	
ibdiagnet	<ul style="list-style-type: none"> • Enabled Adaptive Routing validation • Performance improvement of routing checking • Added new counters and diagnostic information: <ul style="list-style-type: none"> • SHARP: "SAT" ("Streaming Aggregation") counters • PHY: Maximum PLR per second field in PPCNT • PHY: SLLM register • PHY: New SymbolBER thresholds • Adaptive Routing: PortARTrails counter • Added support for SHARP security (AMKEY) • Enabled reporting port counter differences when using "--pm_pause_time" (PM_DELTA section in db_csv) • Added Dragonfly+ Topology Validation (--dfp, --dfp_opt [<max_cas>]) • Added the option to report the version on the screen and in log file
ibtopodif	Added support for "stable" names of RDMA devices based on PCI/slot/function location.
ibnetdiscover dump file (created by ibdiagnet)	Now the dump file includes Virtual Port info
v2.5.1	

Rail Optimized Topology Validation	Checks links between compute nodes and leaf switches to provide rail optimized topology (<code>--rail_validation, --rail_validation_opt [<regex>]</code>).
Service Level	Customization SL for GMP MADs in ibdiagnet (<code>--sl</code>).
General	<ul style="list-style-type: none"> • Added output similar to the results of the "ibnetdiscovery" • ibdiagnet: Disabled the output by default for AR/FDBS files in routing stage • Moved SHARP performance counters to db_csv • Dumped by default FEC_MODE section in db_csv • Improved AR connectivity check
v2.4.0	
General	<ul style="list-style-type: none"> • Flexible output control options (<code>--enable_output</code> , <code>--disable_output</code> , <code>--path</code>) • Discovery only mode (<code>-discovery_only</code>) • Support for MLNX Congestion Control counters • "Fabric Summary" is enabled by default. • Dates and versions in "nodes_info" file are printed in human-readable format. • Dump files include ibdiagnet version and command line parameters. • Added split mode to IBNL for "InfiniBand Smart Director Switches" (CS8500)
Performance Improvements	<p>The performance of following steps in ibdiagnet have been improved:</p> <ul style="list-style-type: none"> • Routing validation • Network discovery • Virtualization stage • Dump creation

3.1 Changes in Command Line Parameters

Version	Tool	Parameter Name	Status	Description
2.16.0	ibdiagpath	--lids_list	New	Using the lids utility for creating scope file.
2.15.0	ibdiagnet	--skip	Deprecated	vs_cap_smp & vs_cap_gmp values will be ignored. SMP & GMP capabilities will be retrieve every time.
	ibdiagpath	--adaptive_routing	New	Using adaptive routing tables to look up for possible paths.
2.12.0	ibdiagnet	--sharp_opt	Changed	Added a new option "- [ad_hoc]" which indicates the ad-hoc trees support in SHARP, and prevents warnings for tree_id duplication in the fabric Value: <[csc][dsc][dscp] [ad_hoc]>

2.11.0	ibdiagnet	--ppcc	New	<p>This parameters enables fetching PPCC (Port Programable Congestion Control) counters.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • File path: In this cases ibdiagnet loads from file PPCC Algorithms • Folder path: ibdiagnet loads all files from the directory • Wildcard: ibdiagnet loads files according to wildcard matching. <p>For more information on the supported wildcard syntax refer to the manual page by typing <code>`man 7 glob`</code></p>
		--verbs	Experimental	Send and receive GMPs via ibverbs instead of ibumad library.
2.9.0	ibdiagnet	--r_opt	Deprecated	The same functionality is supported by --disable_output option Value: skip_far
			Removed	Values: vs, far, rn, drnc, crnc
		--pm_get_all	New	Get all PM counters by activate the following flags: --per_svl_cntrs --sc --extended_speeds pm_per_lane
		--pm_clear_all	New	Clear all PM counters by activate the following flags: --scr --pc
		--ft_roots_regex_opt	New	Regular expression to select topology root nodes. To be applied to switch descriptions. Value: <regular expression>
2.8.0	--r_opt	dump_only_skip_routing_tables	Added	Skips routing tables (LFTs) retrieving
	rn	Deprecated	Dumps routing notification data to file (enabled by default)	
	drnc	Deprecated	Dumps routing notification port counters to file (enabled by default)	
2.7.0	ibdiagnet	--ft	New	Enables Fat Tree Topology Validation (default - disabled)
		--aguid	New	Collects AGUIDs
		--enable_spst	Removed	SPST mode is enabled by default. Option was deprecated in 2.6.1

2.6.1	ibdiagnet	--smp_window	Upper limit and default are changed	Max: 256; default: 16
		--gmp_window	Upper limit and default are changed	Max: 16384; default: 256
		--am_key	New	Specifies default AMKEY for the fabric
		--am_key_file	New	Specifies the path to file (AMKEY per GUID)
		--smdb	New	Specifies the path to OpenSM SMDB file (required for Adaptive Routing & Dragonfly+ Topology validation)
		--ber_threshold_table	New	Specifies the path to BER thresholds table file (BER per FEC)
		--create_ber_threshold_table	New	Creates template file of BER threshold table
		--enable_spst	Deprecated	SPST mode is enabled by default
		--dfp	New	Enabling DFP Topology Validation (default - disabled)
		--dfp_opt	New	Specifies comma separated DFP Topology Validation options
	dfp_opt	<max_cas>	-	Specifies max number of CAs for "Root" switch in Dragonfly+ island (default: 1). This parameter is mutual exclusive with --smdb
	ibdmchk	--FAR	New	Adds support for input FAR file
2.5.1	ibdiagnet	--sl	New	Specifies the SL to be used (default=0)
		--rail_validation	New	Enabling Rail Optimized Topology Validation (default - disabled)
		--rail_validation_opt	New	Specifies comma separated Rail Optimized Topology Validation options
		--clear_congestion_counters	New	Dumping Congestion Counters and clearing them
		--fec_mode	Deprecated	FEC_MODE section will be dump to "db_csv" by default
	--rail_validation_opt	<regex>	-	Specifies regular expression to filter HCA nodes from reports. To be applied to HCAs node descriptions
2.4.0	ibdiagnet	--enable_output	New	Enable creation of specific dump file
		--disable_output	New	Disable specific dump file

	--path	New	Set custom path for specific dump file
	--discovery_only	New	Discover IB fabric, save topology information into "db_csv" file and exit
	--smp_window	Upper limit is changed	New max is 128
	--gmp_window	Upper limit is changed	New max is 8192
ibtopodiff	--ibnl_dir	New	Set path for IBNL files

4 Bug Fixes

2.16.0	<ul style="list-style-type: none"> Fixed an issue related to the execute command from the switches Fixed an issue related to fetching multiple Congestion Control counters
2.15.0	<ul style="list-style-type: none"> Fixed an issue related to loading the 'HWInfo.technology' field from a file. Removed 'PERF' records duplicates in the 'db_csv' file.
2.14.0	<ul style="list-style-type: none"> Fixed an issue related to retrieving RX/TX Power fields in the cable report
2.13.0	<ul style="list-style-type: none"> Fixed a crash in SHARP stage that occurred when one or more switches was reset
2.11.0	<ul style="list-style-type: none"> Fixed the DB_CSV file which contained non-ASCII characters
2.9.0	<ul style="list-style-type: none"> Fixed PCIe plugin behavior to support scope file Fixed the wrong number of SHARP active jobs in fabric summary
2.8.0	<ul style="list-style-type: none"> Fixed an unexpected exit when detected HCAs with invalid GUID in discovery stage Fixed an issue that corrupted the ibdagnet reporting for other nodes when encountering an HCA with an invalid LID Fixed a possible crash in SHARP validation mode when SHARP version is not specified in command line parameters Fixed the wrong values in "switchguid" and in "remote port" fields in ibnetdiscover file
2.6.1	<ul style="list-style-type: none"> Fixed an issue that caused ibtopodiff not to recognize IB Routers Fixed the reset flow for PHY/PCI counters
2.5.1	<ul style="list-style-type: none"> Fixed an issue that resulted in GUIDs being printed in Mkey Mmanager in HEX format Fixed an issue that resulted in error messages being printed with the correct node GUID instead of the value "zero" Fixed an issue that loaded duplicates and half links from the neighbors file of MKey manager Fixed an issue that printed errors/warning status in stages Fixed an issue that resulted in a critical error in case of duplicate names with HCA/mlx patterns
2.4.0	<ul style="list-style-type: none"> Fixed an issue that caused ibdiagnet to skip PCIe counters for some HCAs, if more than one HCA was installed in the host. Fixed an issue that caused ibdiagnet to report a wrong value for SLRG 16nm register. Fixed an issue that caused ibdiagnet to report a wrong value in SMP access register. Fixed an issue that caused ibdiagnet to report a wrong number of "Aggregation Nodes" in "Fabric Summary" table

5 Legal Notices and 3rd Party Licenses

Product	Version	Legal Notices and 3rd Party Licenses
ibutils2	2.16	<ul style="list-style-type: none">• License• 3rd Party Notice• 3rd Party Unify Notice

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/



or Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2024 NVIDIA Corporation & affiliates. All Rights Reserved.

