

Intel® Omni-Path Architecture (Intel® OPA)

Packet Integrity Protection and Local Link Integrity Counter White Paper


October 2017

In 2020, Cornelis Networks acquired the Omni-Path business of Intel Corporation, including the Intel® Omni-Path and Intel® True Scale products.

Cornelis Networks is now the manufacturer of [these](#) products, renamed Omni-Path and True Scale.

As Cornelis Networks creates and transitions to its own written materials for the products, you can reference the attached legacy materials.

All questions and product support requests should be directed to support@cornelisnetworks.com.



Legal Disclaimer

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm>

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at <http://www.intel.com/> or from the OEM or retailer.

No computer system can be absolutely secure.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2017, Intel Corporation. All rights reserved.



Contents


1	Intel® Omni-Path Architecture Packet Integrity Protection	5
1.1	Background	5
1.2	Intel® Omni-Path Implementation.....	5
1.3	Relevant Counters	6
1.3.1	Link Quality Indicator (LQI)	6
1.3.2	LocalLinkIntegrityErrors (LLI) Counter	6
1.3.3	PortRcvErrors (RxE) Counter	7
1.3.4	ExcessiveBufferOverrunErrors (EBO) Counter	7
1.3.5	LinkErrorRecovery (LER) Counter.....	7
1.3.6	LinkDowned (LD) Counter.....	7
1.3.7	UncorrectableErrors (Unc) Counter	7
1.3.8	FMConfigErrors (FMC) Counter	7
1.4	Counter Calculations.....	8
1.4.1	Integrity Category Calculations.....	8
1.4.2	Default Weight Values	8
1.4.3	Resolution Threshold	8
1.5	Corrective Actions	9
1.5.1	Verify OPA Fabric Status and Topology	9
1.5.2	Verify Physical Links	10



Revision History

Date	Revision	Description
October 2017	1.0	First release of document.

§



1 *Intel® Omni-Path Architecture Packet Integrity Protection*

1.1 Background

In the Quad Data Rate (QDR) InfiniBand* (IB) timeframe, packet integrity was protected with a 16-bit Link Cyclic Redundancy Check (LCRC) and a 32-bit Invariant Cyclic Redundancy Check (ICRC). If a bit error occurred in a packet on any link en route to the final destination, the packet would be dropped or marked as bad, and the destination Host Channel Adapter (HCA), upon receiving a packet out of sequence, would initiate an end-to-end retry. Similarly, if the sending HCA never received an ack, after a timeout it would initiate an end-to-end retry. Since a high frequency of end-to-end retries generated anywhere in a cluster can be damaging to cluster performance, the QDR products were designed to generate significantly lower bit error rates than allowed by the IB spec of $1e-12$. While a Bit Error Rate (BER) $< 1e-15$ was guaranteed, most of the QDR IB links ran error-free, and so there was a tolerable, near zero number of end-to-end retries generated in a cluster.

As it moves to 100 Gb links signaling at 25 Gbps and beyond, the industry has recognized that deploying nearly error-free links is not economically feasible. Therefore, specifications have been developed requiring link-level correction, enabling tolerance of raw bit error rates as high as $5e-06$ or more. A raw BER of $5e-06$ would generate 500,000 bit errors per second on each link, but each error would be immediately corrected prior to packet forwarding to the next link, so the error rate can be tolerated without suffering end-to-end retries.

100 Gigabit Ethernet* (GbE) and Enhanced Data Rate (EDR) IB chose to use forward error correction (FEC) as the link-level correction mechanism. While this can correct a large number of raw errors, it has the unfortunate side-effect of adding significant latency to every hop. So, for example, a 90 ns switch latency can jump to over 140 ns when FEC is turned on. This is a serious performance issue for High Performance Computing (HPC) environments considering EDR IB as the interconnect.

1.2 Intel® Omni-Path Implementation

Intel® Omni-Path specifies a link-level correction mechanism called Packet Integrity Protection that corrects bit errors, but does not add latency to every packet, making it a much more suitable scheme for HPC environments.

Packet Integrity Protection allows for rapid and transparent recovery of transmission errors between a sender and a receiver on an Intel® Omni-Path Architecture link. Given the very high Intel® OPA signaling rate (25.78125G per lane) and the goal of supporting large scale systems of a hundred thousand or more links, transient bit errors must be tolerated while ensuring that the performance impact is insignificant. Packet Integrity Protection enables recovery of transient errors whether they occur between a host and switch or between switches. This eliminates the occurrence of



transport level timeouts and end-to-end retries in a transient error scenario. (These mechanisms are used to recover from other situations, such as links going down and dropping packets in the fabric.) This is done without the heavy latency penalty associated with alternate error recovery approaches.

Packet Integrity Protection is capable of handling bit error rates of up to 1e-07 without performance impacts. Intel specifies a worst case bit error rate of 1e-08, which allows up to 1000 corrected link integrity errors per second per link without impacting performance.

Most links will perform much better than this target, but the design handles that level of raw error without issue; some links will likely approach this raw bit error rate. For example, Active Optical Cable (AOC) manufacturers may screen for this BER, and while most will perform far better, some will be in this neighborhood and will of course still be acceptable.

1.3 Relevant Counters

The indicator and counters described here are most relevant for Packet Integrity Protection. For complete details, refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Port Counters Overview section.

1.3.1 Link Quality Indicator (LQI)

This is a status indicator, similar to the signal strength bar display on a mobile phone, that enumerates link quality as a range of 0-5, with 5 being very good. Values in the lower part of the range may indicate hardware problems with components such as ports and cables that surface as signal integrity issues, leading to performance and other problems.

Link Quality Values and Description

Link Quality Value	Description
5	Working at or above preferred link quality, no action needed.
4	Not used.
3	Working on low end of acceptable link quality, recommended corrective action on next maintenance window.
2	Working below acceptable link quality, recommend timely corrective action.
1	Working far below acceptable link quality, recommend immediate corrective action.
0	Link down

1.3.2 LocalLinkIntegrityErrors (LLI) Counter

This counter indicates the number of retries initiated by a link transfer layer receiver.



The retry rate is represented by the Link Quality Indicator. A link that is meeting performance requirements has a Link Quality of 5, which corresponds to 1000 or fewer replays per second.

1.3.3 PortRcvErrors (RxE) Counter

This counter indicates the total number of packets containing an error that were received by the port, including Link Layer protocol violations and malformed packets. It indicates possible misconfiguration of a port, either by the Subnet Manager (SM) or by user intervention. It can also indicate hardware issues or extremely poor link signal integrity.

1.3.4 ExcessiveBufferOverrunErrors (EBO) Counter

This counter, associated with credit management, indicates an input buffer overrun. It indicates possible misconfiguration of a port, either by the Subnet Manager (SM) or by user intervention. It can also indicate hardware issues or extremely poor link signal integrity.

1.3.5 LinkErrorRecovery (LER) Counter

This counter indicates the number of times the link has successfully completed the link error recovery process.

Link Quality Indicator is the primary indicator for link quality to use. This counter is factored into the value reported for Link Quality Indicator. This counter may be non-zero for a properly functioning link.

1.3.6 LinkDowned (LD) Counter


This counter indicates the total number of times the port has failed the link error recovery process and downed the link. These events can cause disruptions to fabric traffic. A typical cause of LinkDowned events is the rebooting of the host or switch on the other end of the link.

1.3.7 UncorrectableErrors (Unc) Counter

This counter indicates the number of unrecoverable device errors. This may indicate a defect in the reporting device.

1.3.8 FMConfigErrors (FMC) Counter

This counter reports inconsistent configurations of the low-level Subnet Management Agent (SMA) on either side of the link. It indicates possible misconfiguration of a port, either by the Subnet Manager (SM) or by user intervention.



1.4 Counter Calculations

The calculations described here occur within the Performance Manager (PM) subsystem of the Fabric Manager (FM). These computed values are most relevant for Packet Integrity Protection. For complete details, refer to the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*, PM Counters Calculations appendix.

1.4.1 Integrity Category Calculations

The integrity category calculation applies weights to each of the following counters before summing them for the final value:

- LocalLinkIntegrityErrors
- PortRcvErrors
- LinkErrorRecovery
- LinkDowned
- UncorrectableErrors
- FMConfigErrors
- Neighbor's ExcessiveBufferOverruns
- LinkWidthDowngrade
Use "LinkWidth.Active – LinkWidthDowngrade.RxActive" for delta
- LinkQualityIndicator
Use "2(5-LQI) – 1" for delta

Note: All counters are the delta value from the previous sweep unless otherwise noted.

1.4.2 Default Weight Values

Within the PM defaults, the LocalLinkIntegrityErrors and LinkErrorRecovery counters have weights of 0 in the calculation. These values are factored into the LQI values reported by each port, so it is redundant to include these values in the weighted calculation of Integrity.

This approach also allows each port to have algorithms for LQI that are properly based on the link speed, width, and physical connector/cable type as appropriate.

1.4.3 Resolution Threshold

In a typical fabric, very few ports have any errors. To optimize PM performance, the PM only gets the detailed error counters from ports where the error counters have changed. In order to perform this optimization, the PMA query that returns the basic port counters (that is, the typical data movement counters) also returns an ErrorCounterSummary value. This value is a sum of all the error counters for the port.

The PM compares this single value to the previous sweep's value and only fetches the error counters from ports where this value has changed. For ports where the value



has not changed, the PM does not fetch error counters and reports the same values for the error counters as in its previous sweep.

Since the LocalLinkIntegrity and LinkErrorRecovery counters might not be zero on good links, the PM specifies a “Resolution” to the PMA. The PMA divides the actual values for these counters by the specified resolution and includes the resulting reduced values for these two counters in its ErrorSummary total.

- Default Resolution value for LocalLinkIntegrity = 10000
- Default Resolution value for LinkErrorRecovery = 1024

These defaults are chosen based on the BER values discussed earlier and the PM default sweep time of 10 seconds.

With these default values, a port with an average of 1000 LocalLinkIntegrity events per second will only have its LocalLinkIntegrity value fetched when it crosses the next 10,000 units boundary. Therefore, the value shown for the counter in the PM/PA, `opareport`, Fabric Manager Graphical User Interface (FM GUI), `opatop`, and other tools that access the counters from the PM may be 10,000 less than its true value.

To see the true values, use one of the following methods:

- Use `opareport` with the `-M` option, which fetches the raw counters directly from the PMA.
- Use the `opapmaquery` tool that directly accesses the PMA.
- Use assorted other tools that use `opareport`, for example `opaextracterrors`, which also supports the `-M` option.

1.5 Corrective Actions

This section describes two corrective actions.

1.5.1 Verify OPA Fabric Status and Topology

For complete details, refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Verifying OPA Fabric Status and Topology section.

1. Start the Intel FastFabric OPA Tools using either the `opafastfabric` or `opaconfig` commands.
2. From the main menu, type 4 to open the FastFabric OPA Host Verification/Admin menu.
3. From the FastFabric OPA Host Verification/Admin menu, type 4.
The menu item changes from [Skip] to [Perform].

Note: More than one menu item may be selected. The operations are performed individually and in sequence with the menu.

4. (Optional) Type a to select the Start or Stop Bit Error Rate Cable Test option.
5. Type p to begin the operation.
6. At each prompt, provide the required information and press **Enter**. You can optionally check error rate, downgrades, and other items.



When the tests complete, a punchlist that summarizes the issues found is generated. The punchlist file is located in `$FF_RESULT_DIR/punchlist.csv`, typically `/root/punchlist.csv`. The issues are also shown and saved to a more detailed results output file.

If desired, you can directly use the `opalinkanalysis` and `opacabletest` commands. Refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide* for details.

1.5.2 Verify Physical Links

The corrective action described here is most relevant for resolving physical link issues. For complete details, refer to the *Intel® Omni-Path Fabric Staging Guide*, Debug Intel® Omni-Path Physical Link Issues section.

Intel recommends the following approach for physical link resolution:

1. Unplug and re-insert each end of a physical cable. Check that the cable actually clicks into place. It may be useful to do this step separately for each end of the cable. Re-run the verification procedure and verify whether the issue has been resolved or not.

Note: This step has resolved more link issues in fabric installs than all others.
2. Replace the questionable cable with a known good cable to isolate whether it is an Intel® Host Fabric Interface (HFI)/Switch issue or cable issue.
3. If step 2 did not resolve the issue, then connect the known good cable to an alternate switch port.
 - a. If the alternate switch port links up, then replace the switch.
 - b. If the alternate switch port fails to link, replace the HFI/Switch at the other end of the link.
4. If step 2 resolved the issue, then install the questionable cable into another location that previously had a known good link, and verify whether it works.
 - a. If the questionable cable works in its new location, then the issue may be a mechanical latching issue on the HFI/Switch connector. Repeat step 1 in the original location and see if the issue persists.
 - b. If the questionable cable does not work, contact Intel Customer Support about returning the cable.