

iSER: Frequently Asked Questions

The iSCSI Extensions for RDMA or iSER protocol is an iSCSI translation layer for operation over RDMA transports, such as InfiniBand or iWARP/Ethernet. The most unexpected fact about iSER is that, despite its name, it is not compatible with iSCSI and does not interoperate with the large iSCSI installed base. This FAQ clarifies this and other questions that are associated with this protocol.

Is iSER compatible with iSCSI?

NO – iSER is effectively an emulation layer that translates iSCSI to RDMA transactions that cannot be understood by non-RDMA peers: iSER transports iSCSI control messages in RDMA SEND/RECV commands and iSCSI payload in RDMA READ/WRITE operations, altering the wire formats in the process. The following figure shows the protocol layering for iSCSI vs. iSER, revealing the number of additional layers introduced by the latter, and the resulting incompatibility on the wire.

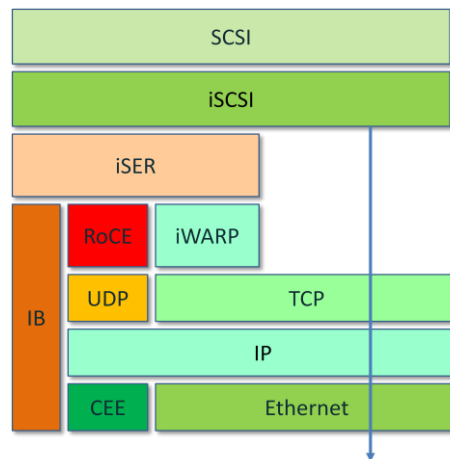


Figure 1 – iSCSI and iSER Protocol Layers

Does iSER interoperate with iSCSI peers?

NO – an iSER client or server can only communicate with other iSER end nodes. Interoperating in this case means disabling iSER extensions, and losing offload support and all performance benefits.

Does iSER work with software peers?

NO – iSER requires RDMA enabled hardware on both ends of the communication. In contrast, iSCSI hardware implementations are not only fully interoperable with software peers, but also preserve all the local performance benefits regardless of the peer type.

Does iSER provide good performance?

MAYBE – iSER can provide improved efficiencies at large I/O sizes, but benchmarks show degraded performance at small I/O sizes due to the additional overheads introduced. iSCSI offload suffers from no such overheads and thus provides predictably better performance.

Does iSER provide lower CPU utilization?

YES – Compared to software implementations, applications get lower CPU utilization thanks to hardware offload and direct data placement (DDP), which are also both provided by iSCSI offload HBAs.

Does iSER provide true zero copy?

YES – Again, true zero copy is also provided by iSCSI offload HBAs and is an integral part of the iSCSI protocol itself.

Does iSER work over Ethernet?

YES – iSER can utilize RDMA transports that operate over Ethernet, such as iWARP and RoCE. However, only iWARP works over standard Ethernet and provides the routability, scalability and reach of iSCSI.

Does iSER over RoCE provide out-of-order placement and re-ordering?

NO – RoCE requires a lossless network with expensive DCB equipment because it cannot deal with packet loss and reordering.

Is iSER deployed widely?

NO – Unlike iSCSI, which is natively supported in all major operating systems, iSER requires large software and system integration investments that limit its usability, similarly to the now deprecated SCSI RDMA Protocol (SRP) before it that failed to gain market acceptance.

Does iSER replace iSCSI?

NO – iSER only supports the data movement phase of iSCSI, the rest of the protocol is unchanged, and must be preserved and maintained within the implementation. It is thus not possible to simplify the storage stack by adopting iSER.

Is there any advantage to iSER over iSCSI?

NOT over ETHERNET – iSER is needed in fabrics with inefficient TCP/IP support, such as InfiniBand, but provides little benefit in Ethernet compared to hardware offloaded iSCSI. iSER requires an extra layer on top of iSCSI for the same CPU offload, true zero copy, and enhanced data protection as iSCSI HBAs.

On the other hand, iSER over RoCE suffers from the known issues associated with the latter, including expensive infrastructure with CEE (DCB) support, complex configuration, lack of congestion control and short reach. The situation is further confused by the many versions of RoCE that have different limitations and do not interoperate among themselves, even from the same vendor.

Related Links

[RoCE FAQ](#)

[Heritage and Future: iSCSI vs. iSER](#)

[Chelsio Technical Library](#)

References

[1] [RFC 5040 - A Remote Direct Memory Access Protocol Specification](#)

[2] [RFC 7145 - iSCSI Extensions for the RDMA Specification](#)