

Low-Latency UDP

for Myri-10G 10-Gigabit Ethernet NICs

DBL (Datagram Bypass Layer) is an optional, user-level, software interface for accelerating applications that depend on UDP/IP communication. DBL was developed principally for financial PubSub applications, but is also useful for other applications in which performance depends on UDP latency.

DBL software with Myri-10G 10-Gigabit Ethernet NICs leverages user-level, kernel-bypass, messaging techniques originally developed for HPC applications, and applies them to UDP communication over Ethernet in a multicast and/or unicast environment. DBL is interoperable and wire-compatible with other UDP/IP implementations.

Unlike UDP communication through the host operating system, which allows the full set of IP networking services to potentially hundreds of user application threads, but at the cost of latency in the operating-system protocol stack, DBL takes advantage of kernel-bypass firmware in Myri-10G NICs to allow high-priority user threads to send and receive UDP frames directly. The typical user-level latency of DBL on modern hosts is $\sim 2.5\mu\text{s}$.

DBL is not intended to replace UDP communication through the host operating system, nor to improve performance in situations where overall system performance depends on multiplexing many user application threads over fewer available computing resources (or CPU cores). Rather, DBL will improve the throughput, response time, and transaction rates of user application threads that have dedicated computing resources. As such, DBL best suits cases where target latencies, response times, and jitter would otherwise be hampered by the operating-system protocol stack and scheduling policies.

When used with Myri-10G 10-Gigabit Ethernet NICs, DBL can be used selectively to accelerate certain UDP sockets while allowing the remaining TCP/IP and UDP/IP traffic to be handled through the conventional Ethernet driver and operating-system protocol stack. The DBL software cooperates with and extends the standard BSD Sockets API for UDP packets. The DBL API allows users to create datagram sockets and bind them to DBL-capable interfaces. There are function calls to implement the datagram sockets `sendto()` and `recvfrom()` semantics, but with an interface that reduces software overhead in order to achieve low latency and high message rate.

DBL software distributions are currently available for Linux, Windows, Solaris, and FreeBSD. The full DBL API is accessible from C and Java on all supported platforms. In addition, DBL can be used via the Java NIO Channel interface on Linux, Solaris, and FreeBSD. The Channel interface allows existing applications to achieve dramatic latency and transaction-time improvements with only minimal changes to existing source code.

Access to this optional software for Myri-10G NICs requires a mutual NDA. Complete documentation is provided in the software distribution.

Myricom[®] is a registered trademark of Myricom, Inc.