

## Market Requirements Document

### Feature Name: Infiniband Support

Version: 1

Date Submitted: 11/10/09

Completed By: Leon Guzenda

### Description of the Problem

Infiniband™ is an increasingly popular switched fabric communications link that is used in a large number of High Performance Computing (HPC) environments. Objectivity/DB can run without change on an operating system and hardware that use Infiniband, but it takes no advantage of its advanced performance and architectural features.

### Background

InfiniBand features include quality of service and failover and it is designed to be scalable. The InfiniBand architecture specification defines a connection between processor nodes and high performance I/O nodes such as storage devices. InfiniBand forms a superset of the Virtual Interface Architecture.

Infiniband provides:

- User-level access to message passing.
- Remote direct memory access (both read and write modes)
- Up to 2 GB of data can be moved in a single request.

Objectivity/DB could exploit Infiniband to move data (pages) between clients and Advanced Multithreaded Servers (AMS).

### Description of the Requested Feature

Direct kernel support for Infiniband (or the Virtual Interface Architecture) that exploits its performance features. Besides using its fast messaging capability (instead of standard TCP/IP) the kernel and AMS could also be configured to:

- Allow a client to request a page, receive a remote memory pointer, then move one or more pages directly from a pre-cached AMS memory to the regular client side cache. This would considerably improve streaming speeds.
- Allow a client to move blocks of data (one or more pages) directly into a remote AMS memory cache for a particular file. This would improve data ingest and most commit operation speeds.
- Allow AMS to directly read one or more pages from its local disk and write them directly into a client side Infiniband remote (to the AMS) memory cache.

**Part of an existing feature or does it require another feature, if so, which one?**

- Enhances performance in configurations that use Infiniband hardware.

**How is this problem being solved now, and why isn't that acceptable?**

Objectivity/DB has been run in configurations that use Infiniband, but without exploiting the performance features.

**What languages must support this capability?**

- C++
- Java
- SQL++

**Which platforms must be supported?**

- Linux, Windows and Solaris.

**Do any competitors already have this feature?**

- No.

**Customers who require this feature**

- HPC applications.

**Revenue at risk, or which could be won**

- This will allow us to leverage relationships with hardware providers, such as DDN.
- It will give us a competitive performance edge in HPC environments.

**When is this required?**

- Release 11.



**Additional Notes****1. Related Material**

We will also need:

- Marketing collateral, including a Press Release.
- Technical Publications.
- Release Note.
- Extra Quality Assurance Material, particularly performance tests versus regular AMS and local connections.

**2. Hardware Requirements**

Engineering will require access to Infiniband hardware, either within the network or at a partner. A 20-port Infiniband switch costs around \$6K. Each computer connected to it will need a port card.