

MVAPICH2 1.0.3 User Guide

MVAPICH TEAM

NETWORK-BASED COMPUTING LABORATORY
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
THE OHIO STATE UNIVERSITY

<http://mvapich.cse.ohio-state.edu>

Copyright ©2003-2008
Network-Based Computing Laboratory,
headed by Dr. D. K. Panda.
All rights reserved.

Last revised: July 2, 2008

Contents

1	Overview of the Open-Source MVAPICH Project	1
2	How to use this User Guide?	1
3	MVAPICH2 1.0 Features	2
4	Installation Instructions	6
4.1	Download MVAPICH2 source code	6
4.2	Prepare MVAPICH2 source code	6
4.3	Downloading MVAPICH2 Source Code from Anonymous SVN	6
4.4	Build MVAPICH2	7
4.4.1	Build MVAPICH2 with OpenFabrics Gen2-IB and iWARP	7
4.4.2	Build MVAPICH2 with uDAPL	8
4.4.3	Build MVAPICH2 with VAPI	9
4.4.4	Build MVAPICH2 with TCP/IPoIB	10
5	Basic Usage Instructions	11
5.1	Compile MPI Applications	11
5.2	Setting MPD Environment	11
5.3	Run MPI Applications Using mpiexec with OpenFabrics Gen2-IB or VAPI Device	12
5.3.1	Run MPI Applications using Shared Library Support	13
5.3.2	Run MPI Application using TotalView Debugger Support	13
5.4	Run MPI Application using OpenFabrics Gen2-iWARP Device	14
5.5	Run MPI Application using mpiexec with uDAPL Device	14
5.6	Run MPI Application using mpiexec with TCP/IP	15
5.6.1	The MPI Job Uses IPoIB	15
5.6.2	Both MPD And the MPI Job Use IPoIB	15
5.7	Run MPI Applications using SLURM	16
6	Advanced Usage Instructions	17

6.1	Run MPI applications on Multi-Rail Configurations (for OpenFabrics Gen2-IB and Gen2-iWARP Devices)	17
6.2	Run MPI application with Customized Optimizations (for OpenFabrics Gen2-IB and Gen2-iWARP Devices)	17
6.3	Run MPI application with Checkpoint/Restart Support (for OpenFabrics Gen2-IB Device)	18
6.4	Run MPI application with RDMA CM support (for OpenFabrics Gen2-IB and Gen2-iWARP Devices)	20
6.5	Run MPI application with Shared Memory Collectives	21
6.6	Run MPI Application with Hot-Spot and Congestion Avoidance (for OpenFabrics Gen2-IB Device)	21
6.7	Run MPI Application with Network Fault Tolerance Support (for OpenFabrics Gen2-IB Device)	22
7	Using OSU Benchmarks	23
8	FAQ and Troubleshooting with MVAPICH2	24
8.1	General Questions and Troubleshooting	24
8.1.1	Invalid Communicators Error	24
8.1.2	Are <code>fork()</code> and <code>system()</code> supported?	24
8.1.3	Cannot Build with the PathScale Compiler	24
8.1.4	Cannot find <code>mpd.conf</code>	24
8.1.5	<code>sched_setaffinity: Bad address</code>	25
8.1.6	Multi-threaded programs seem to run sequentially	25
8.1.7	Running MPI programs built with <code>gfortran</code>	25
8.2	With Gen2 Interface	25
8.2.1	Cannot Open HCA	25
8.2.2	Checking state of IB Link	25
8.2.3	Undefined reference to <code>ibv_get_device_list</code>	26
8.2.4	Creation of CQ or QP failure	26
8.2.5	Hang with the HSAM Functionality	26
8.2.6	Failure with Automatic Path Migration	26
8.2.7	Error opening file	26

8.2.8	RDMA CM Address error	26
8.2.9	RDMA CM Route error	27
8.3	With Gen2-iWARP Interface	27
8.3.1	Error opening file	27
8.3.2	RDMA CM Address error	27
8.3.3	RDMA CM Route error	27
8.4	With VAPI Interface	27
8.4.1	Cannot pass MPI_Init	27
8.4.2	Cannot Open HCA	27
8.4.3	Cannot include vapi.h	28
8.4.4	VAPI_RETRY_EXEC_ERROR	28
8.4.5	ld:multiple definitions of symbol _calloc error on MacOS	28
8.4.6	No Fortran interface on the MacOS platform	28
8.5	With UDAPL Interface	29
8.5.1	Cannot Open IA	29
8.5.2	DAT Insufficient Resource	29
8.5.3	Cannot find libdat.so	29
8.5.4	Cannot find mpd.conf	29
8.6	The MPD mpiexec fails with “no msg recvd from mpd when expecting ack of request.”	29
8.7	Checkpoint/Restart	29
9	Scalable features for Large Scale Clusters and Performance Tuning	31
9.1	RDMA Based Point-to-Point tuning	31
9.2	Shared Receive Queue (SRQ) Tuning	31
9.3	Shared Memory Tuning	32
9.4	On-demand Connection Management Tuning	32
10	MVAPICH2 Parameters	33
10.1	MV2_CKPT_FILE	33
10.2	MV2_CKPT_INTERVAL	33
10.3	MV2_CKPT_MAX_SAVE_CKPTS	33

10.4	MV2_CKPT_MPD_BASE_PORT	34
10.5	MV2_CKPT_MPIEXEC_PORT	34
10.6	MV2_CKPT_NO_SYNC	34
10.7	MV2_CM_RECV_BUFFERS	34
10.8	MV2_CM_SPIN_COUNT	35
10.9	MV2_CM_TIMEOUT	35
10.10	MV2_DAPL_PROVIDER	35
10.11	MV2_DEFAULT_MTU	35
10.12	MV2_ENABLE_AFFINITY	36
10.13	MV2_GET_FALLBACK_THRESHOLD	36
10.14	MV2_IBA_EAGER_THRESHOLD	36
10.15	MV2_INITIAL_PREPOST_DEPTH	37
10.16	MV2_MPD_RECVTIMEOUT_MULTIPLIER	37
10.17	MV2_NDREG_ENTRIES	37
10.18	MV2_NUM_HCAS	37
10.19	MV2_NUM_PORTS	38
10.20	MV2_NUM_QP_PER_PORT	38
10.21	MV2_NUM_RDMA_BUFFER	38
10.22	MV2_ON_DEMAND_THRESHOLD	38
10.23	MV2_PREPOST_DEPTH	39
10.24	MV2_PUT_FALLBACK_THRESHOLD	39
10.25	MV2_RDMA_CM_ARP_TIMEOUT	39
10.26	MV2_RNDV_PROTOCOL	39
10.27	MV2_R3_THRESHOLD	40
10.28	MV2_R3_NOCACHE_THRESHOLD	40
10.29	MV2_SHMEM_COLL_MAX_MSG_SIZE	40
10.30	MV2_SHMEM_COLL_NUM_COMM	40
10.31	MV2_SRQ_LIMIT	40
10.32	MV2_SRQ_SIZE	41
10.33	MV2_USE_APM	41

10.34MV2_USE_APM_TEST	41
10.35MV2_USE_BLOCKING	41
10.36MV2_USE_COALESCE	42
10.37MV2_USE_HSAM	42
10.38MV2_USE_IWARP_MODE	42
10.39MV2_USE_LAZY_MEM_UNREGISTER	42
10.40LAZY_MEM_UNREGISTER	43
10.41MV2_USE_RDMA_CM	43
10.42MV2_USE_RDMA_FAST_PATH	43
10.43MV2_USE_RDMA_ONE_SIDED	43
10.44MV2_USE_RING_STARTUP	44
10.45MV2_USE_SHARED_MEM	44
10.46MV2_USE_SHMEM_ALLREDUCE	44
10.47MV2_USE_SHMEM_BARRIER	44
10.48MV2_USE_SHMEM_COLL	44
10.49MV2_USE_SHMEM_REDUCE	45
10.50MV2_USE_SRQ	45
10.51MV2_VBUF_POOL_SIZE	45
10.52MV2_VBUF_SECONDARY_POOL_SIZE	45
10.53MV2_VBUF_TOTAL_SIZE	46
10.54SMP_EAGERSIZE	46
10.55SMPI_LENGTH_QUEUE	46
10.56SMP_NUM_SEND_BUFFER	46
10.57SMP_SEND_BUF_SIZE	47

1 Overview of the Open-Source MVAPICH Project

InfiniBand and 10GbE/iWARP are emerging as high-performance interconnects delivering low latency and high bandwidth. They are also getting widespread acceptance due to their *open standards*.

MVAPICH (pronounced as “em-vah-pich”) is an *open-source* MPI software to exploit the novel features and mechanisms of InfiniBand, iWARP and other RDMA-enabled interconnects and deliver best performance and scalability to MPI applications. This software is developed in the [Network-Based Computing Laboratory \(NBCL\)](#), headed by [Prof. Dhabaleswar K. \(DK\) Panda](#).

Currently, there are two versions of this MPI: MVAPICH with MPI-1 semantics and MVAPICH2 with MPI-2 semantics. This *open-source* MPI software project started in 2001 and a first high-performance implementation was demonstrated at Supercomputing ’02 conference. After that, this software has been steadily gaining acceptance in the HPC and InfiniBand community. As of June 10th, 2008, more than 700 organizations (National Labs, Universities and Industry) world-wide have downloaded this software from OSU’s web site directly. In addition, many InfiniBand and iWARP vendors, server vendors, and systems integrators have been incorporating MVAPICH/MVAPICH2 into their software stacks and distributing it. Several InfiniBand systems using MVAPICH/MVAPICH2 have obtained positions in the TOP 500 ranking. MVAPICH and MVAPICH2 are also available with the Open Fabrics Enterprise Distribution (OFED) stack. Both MVAPICH and MVAPICH2 distributions are available under BSD licensing.

More details on MVAPICH/MVAPICH2 software, users list, mailing lists, sample performance numbers on a wide range of platforms and interconnect, a set of OSU benchmarks, related publications, and other InfiniBand- and iWARP-related projects (parallel file systems, storage, data centers) can be obtained from the following URL:

<http://mvapich.cse.ohio-state.edu>

This document contains necessary information for MVAPICH2 users to download, install, test, use, tune and troubleshoot MVAPICH2 1.0.3. As we get feedbacks from users and take care of bug-fixes, we introduce new tarballs and also continuously update this document. Thus, we strongly request you to refer to our web page for updates.

2 How to use this User Guide?

This guide is designed to take the user through all the steps involved in configuring, installing, running and tuning MPI applications over InfiniBand using MVAPICH2 1.0.3.

In Section 3 we describe all the features in MVAPICH2 1.0.3. As you read through this section, please note our new features (highlighted as NEW) in the 1.0 series. Some of these features are designed in order to optimize specific type of MPI applications and achieve greater scalability. Section 4 describes in detail the configuration and installation steps. This section enables the user to identify specific compilation flags which can be used to turn some of the features on or off. Basic usage of MVAPICH2 is explained in Section 5. Section 6 provides instructions for running MVAPICH2 with some of the advanced features. Section 7 describes the usage of the

OSU Benchmarks. If you have any problems using MVAPICH2, please check Section 8 where we list some of the common problems people face. In Section 9 we suggest some tuning techniques for multi-thousand node clusters using some of our new features. Finally in Section 10 we list all important run-time parameters, their default values and a small description of what that parameter stands for.

3 MVAPICH2 1.0 Features

MVAPICH2 (MPI-2 over InfiniBand) is an MPI-2 implementation based on [MPICH2](#) ADI3 layer. It also supports all MPI-1 functionalities. MVAPICH2 1.0.3 is available as a single integrated package (with MPICH2 1.0.5p4).

The current release supports the following five underlying transport interfaces:

- **OpenFabrics Gen2-IB:** This interface supports all InfiniBand compliant devices based on the [OpenFabrics](#) Gen2 layer. This interface has the most features and is most widely used. For example, this interface can be used over all Mellanox InfiniBand adapters, IBM eHCA adapters and Qlogic adapters.
- **OpenFabrics Gen2-iWARP:** This interface supports all iWARP compliant devices supported by OpenFabrics. For example, this layer supports Chelsio T3 adapters with the native iWARP mode.
- **uDAPL:** This interface supports all network-adapters and software stacks which implement the portable DAPL interface from the [DAT Collaborative](#). For example, this interface can be used over all Mellanox adapters, Chelsio adapters and NetEffect adapters. It can also be used with Solaris uDAPL-IBTL implementation over InfiniBand adapters.
- **VAPI:** This interface supports Mellanox VAPI interface available from the IB Gold software distribution. This interface can be used for all Mellanox adapters.
- **TCP/IP:** The standard TCP/IP interface (provided by MPICH2) to work with a range of network adapters supporting TCP/IP interface. This interface can be used with IPoIB (TCP/IP over InfiniBand network) support of InfiniBand also. However, it will not deliver good performance/scalability as compared to the other interfaces.

MVAPICH2-1.0.3 delivers the same level of performance as MVAPICH 1.0.1, the latest release package of MVAPICH supporting MPI-1 standard. In addition, MVAPICH2 1.0.3 provides support and optimizations for other MPI-2 features, multi-threading and fault-tolerance (Checkpoint-restart).

A complete set of features of MVAPICH2 1.0.3 are:

- Design for scaling to multi-thousand nodes with highest performance and reduced memory usage.

- On-demand Connection Management: This feature enables InfiniBand connections to be setup dynamically, enhancing the scalability of MVAPICH2 on clusters of thousands of nodes.
 - * Native InfiniBand Unreliable Datagram (UD) based asynchronous connection management for OpenFabrics Gen2-IB interface.
 - * (NEW) RDMA CM based on-demand connection management for OpenFabrics Gen2-iWARP and OpenFabrics Gen2-IB interfaces.
 - * (NEW) uDAPL on-demand connection management based on standard uDAPL interface.
 - (NEW) Message coalescing support to enable reduction of per Queue-pair send queues for reduction in memory requirement on large scale clusters. This design also increases the small message messaging rate significantly. Available for OpenFabrics Gen2-IB interface.
 - (NEW) Hot-Spot Avoidance Mechanism (HSAM) for alleviating network congestion in large scale clusters. Available for OpenFabrics Gen2-IB interface.
 - (NEW) RDMA Read utilized for increased overlap of computation and communication for OpenFabrics device. Available for OpenFabrics Gen2-IB and iWARP interfaces.
 - Shared Receive Queue (SRQ) with flow control. This design uses significantly less memory for MPI library. Available for OpenFabrics Gen2-IB interface.
 - Adaptive RDMA Fast Path with Polling Set for low-latency messaging. Available for OpenFabrics Gen2-IB and iWARP interfaces.
- Fault tolerance support
 - Checkpoint-restart support for application transparent systems-level fault tolerance. [BLCR-based](#) support using OpenFabrics Gen2-IB interface.
 - (NEW) Application-initiated system-level checkpointing is also supported. User application can request a whole program checkpoint synchronously by calling special MVAPICH2 functions.
 - * Flexible interface to work with different files systems. Tested with ext3 (local disk), NFS and PVFS2.
 - (NEW) Network-Level fault tolerance with Automatic Path Migration (APM) for tolerating intermittent network failures over InfiniBand.
 - Optimized intra-node communication support by taking advantage of shared-memory communication. Available for all interfaces.
 - Efficient Buffer Organization for Memory Scalability of Intra-node Communication
 - Multi-core optimized.
 - * [Performance benefits of multi-core optimized support can be found here.](#)
 - Optimized for Bus-based SMP and NUMA-Based SMP systems.
 - Utilizes Processor Affinity.

- Shared memory optimizations for collective communication operations. Available for all interfaces.
 - Efficient algorithms and optimizations for barrier, reduce and all-reduce operations
 - * [Performance benefits of optimized collective support can be found here.](#)
- Integrated multi-rail communication support. Available for OpenFabrics Gen2-IB and iWARP interfaces.
 - Multiple queue pairs per port
 - Multiple ports per adapter
 - Multiple adapters
 - Support for both one-sided and point-to-point operations
 - (NEW) Support for OpenFabrics Gen2-iWARP interface and RDMA CM (for Gen2-IB).
- Multi-threading support. Available for all interfaces, including TCP/IP.
- High-performance optimized and scalable support for one-sided communication: Put, Get and Accumulate. Supported synchronization calls: Fence, Active Target, Passive (lock and unlock).
 - Direct RDMA based One-sided communication support for OpenFabrics Gen2-iWARP and RDMA CM (with Gen2-IB)

Available for all interfaces.

- Two modes of communication progress
 - Polling
 - (NEW) Blocking (enables running multiple MPI processes/processor). Available for Open Fabrics Gen2-IB interface.
- Scalable job startup using in-band IB communication with MPD
- Advanced AVL tree-based Resource-aware registration cache
- Memory Hook Support provided by integration with `ptmalloc2` library. This provides safe release of memory to the Operating System and is expected to benefit the memory usage of applications that heavily use `malloc` and `free` operations.
- High Performance and Portable Support for multiple networks and operating systems through uDAPL interface.
 - InfiniBand (tested with)
 - * uDAPL over OpenFabrics Gen2-IB on Linux
 - * uDAPL over IBTL on Solaris

This uDAPL support is generic and can work with other networks that provide uDAPL interface. Please note that the stability and performance of MVAPICH2 with uDAPL depends on the stability and performance of the uDAPL library used. Starting from version 1.0.3, MVAPICH2 supports both uDAPL v2 and v1.2 on Linux.

- Support for TotalView debugger.
- Shared Library Support for existing binary MPI application programs to run.
- ROMIO Support for MPI-IO.
- Single code base for the following platforms (Architecture, OS, Compilers, Devices and InfiniBand adapters)
 - Architecture: (tested with) EM64T, Opteron and IA-32; IBM PPC and Mac G5
 - Operating Systems: (tested with) Linux and Solaris; and Mac OSX
 - Compilers: (tested with) gcc, intel, pathscale, pgi and sun studio
 - Devices: (tested with) OpenFabrics Gen2-IB, OpenFabrics Gen2-iWARP, and uDAPL; VAPI and TCP/IP
 - InfiniBand adapters:
 - * (tested with) Mellanox adapters with PCI-X and PCI-Express (SDR and DDR with mem-full and mem-free cards) and ConnectX
 - * PathScale adapter (through OpenIB/Gen2 support)
 - * IBM ehca adapter (through OpenIB/Gen2 support)
 - 10GigE adapters:
 - * (tested with) Chelsio T3 adapter with iWARP support
- Integrated and easy-to-use build script which automatically detects system architecture and InfiniBand adapter types and optimizes MVAPICH2 for any particular installation.

The MVAPICH2 1.0.3 package and the project also includes the following provisions:

- [Public SVN](#) access of the codebase
- A set of micro-benchmarks (including multi-threading latency test) for carrying out MPI-level performance evaluation after the installation
- Public [mvapich-discuss](#) mailing list for mvapich users to
 - Ask for help and support from each other and get prompt response
 - Enable users and developers to contribute patches and enhancements

4 Installation Instructions

4.1 Download MVAPICH2 source code

The MVAPICH2 1.0.3 source code package includes the MPICH2 1.0.5p4 version. All the required files are present as a single tarball.

You can go to the [MVAPICH2 website](#) to obtain the source code.

4.2 Prepare MVAPICH2 source code

Untar the archive you have downloaded from the web page. Given 1.0.3 version, the following command can run on most Unix machines:

```
$ tar xzf mvapich2-1.0.3.tar.gz
```

You will have a directory named `mvapich2-1.0.3`.

4.3 Downloading MVAPICH2 Source Code from Anonymous SVN

The MVAPICH2 source code is also available for download through a public SVN:

- To view the latest MVAPICH2 source code:
`https://mvapich.cse.ohio-state.edu/svn/mpi/mvapich2`
- There are three directories (trunk, tags and branches) in the MVAPICH2 SVN repository. Please keep in mind the following guidelines before deciding which version of MVAPICH2 to download:
 - “trunk” will contain the latest source code as we enhance and improve MVAPICH2. It may contain newer features and major bug fixes. However, it is likely to be lightly tested.
 - If you want a stable version with major bug fixes, please use the “1.0” branch from the branches directory. Please be aware that while the branches directory contains bug fixes, it will not have any new features, and as such is used for maintenance purposes.
 - If you need exactly the version released, You may also use the “1.0.3” tag from the tags directory. However, there will be no updates to this tags directory.
- In order to check out the current MVAPICH2 head from the SVN, you need to have SVN installed on your machine. Please execute the following commands (based on which version you want to use, as indicated above):
 - To check out the latest trunk (with latest features):
 - * `svn co https://mvapich.cse.ohio-state.edu/svn/mpi/mvapich2/trunk mvapich2`
 - To check out a release branch (containing bug fixes):

- * `svn co https://mvapich.cse.ohio-state.edu/svn/mpi/mvapich2/branches/1.0/ mvapich2`
- To check out a release tag:
 - * `svn co https://mvapich.cse.ohio-state.edu/svn/mpi/mvapich2/tags/1.0.3/ mvapich2`

- The directory `mvapich2` under your current directory contains MVAPICH2.

4.4 Build MVAPICH2

There are several options to build MVAPICH2 1.0.3 based on the underlying InfiniBand libraries you want to utilize. In this section we describe in detail the steps you need to perform to correctly build MVAPICH2 on your choice of libraries, namely OpenFabrics Gen2-IB, OpenFabrics Gen2-iWARP, Mellanox VAPI, uDAPL and TCP/IP.

4.4.1 Build MVAPICH2 with OpenFabrics Gen2-IB and iWARP

There are several methods to configure MVAPICH2 and optimize it for a given platform.

- Using Default configuration

Go to the `mvapich2-1.0.3` directory. We have included a single script for Gen2 (`make.mvapich2.ofa`) that takes care of different platforms, compilers and architectures. By default, the compilation script uses `gcc`. In order to use a specific compiler, set the `CC`, `CXX`, `F77`, and `F90` variables to the appropriate commands. MVAPICH2 currently supports the GCC, Intel, PathScale and PGI compiler suites for Gen2.

- Customizing MVAPICH2 to use additional features and optimization schemes

MVAPICH2 has many optimization schemes to enhance performance. In 1.0.3, many of those optimization schemes can be turned on/off by runtime variables. Please refer to Section 6.2 for examples to run MVAPICH2 with those runtime variables and Section 10 for a full list of the variables controlling other optimization schemes. However, for performance reasons, there are still some schemes controlled by compilation flags which are listed in the variable `CFLAGS` (in the default compilation script) and/or various configure options.

- *Memory Hook Support:* The variable `PTMALLOC` in the compilation script `make.mvapich2.ofa` is used for enabling/disabling the memory hooks. By default, this is enabled. In order to disable it, change the value of the variable to “no”. The memory hook support based on the `ptmalloc2` library allows MVAPICH2 to return memory back to the system upon calls to `free`. This is expected to improve memory consumption on long running applications which perform several `malloc/free` operations.
- *Header Caching:* Controlled by `-DUSE_HEADER_CACHING`. Provides better performance for small messages. In order to disable it, you can remove this from the compilation script.
- *Shared-memory communication:* Controlled by `-D_SMP_`. Intra-node communication uses shared-memory operations instead of loopback from the HCA.

- *Shared library support*: To enable shared library support, add the flag `--enable-sharedlibs=kind` to configure in the default `make.mvapich2.ofa` script, where “kind” should be “gcc” for the GCC, Intel, PathScale, and PGI compilers on Intel and Opteron Hardware and “solaris86-cc” for the Sun Studio compiler on Solaris Opteron hardware for the uDAPL build script specifically. Please refer to Section 5.3.1 for more details.
- *TotalView Debugger support*: This enables the use of TotalView debugger for your MPI applications. In order to enable this feature, you need to pass `--enable-sharedlibs=kind` as noted above for shared library support and `--enable-g=dbg` options to configure in the `make.mvapich2.ofa` script. Please refer to Section 5.3.2 for more details.
- *Checkpoint/Restart support*: This feature is available only for the OpenFabrics Gen2-IB interface. The variable `ENABLE_CKPT` in the compilation script `make.mvapich2.ofa` is used for enabling/disabling the checkpoint/restart support. By default, this is disabled. In order to enable it, change the value of the variable to “yes”. Checkpoint/Restart support in MVAPICH2 is based on Berkeley Lab Checkpoint/Restart (BLCR) toolkit, the system must be installed with BLCR first. Before compilation, the environment variable `BLCR_HOME` must be set to the installation path of BLCR. When `ENABLE_CKPT` is set to “yes”, the Shared-memory communication is disabled automatically, and the Memory Hook Support is required. Please refer to Section 6.3 for more details.
- *RDMA CM support*: The variable `RDMA_CM_SUPPORT` in the compilation script `make.mvapich2.ofa` is used for enabling and disabling the basic RDMA CM support. This flag is enabled by default. The value of this variable needs to be set to “no” for disabling this support. Please refer to Section 6.4 for more details on usage of RDMA CM. (Note: In order to use RDMA CM support on OFED 1.1 software stack, please add the flag `-DOFED_VERSION_1.1` to the cflags in the compilation script `make.mvapich2.ofa`).

After setting all the parameters, the script `make.mvapich2.ofa` configures, builds and installs the entire package in the directory specified by the variable `PREFIX`.

4.4.2 Build MVAPICH2 with uDAPL

There are several methods to configure MVAPICH2 and optimize it for a given platform.

- *Use of the Default configuration*

Go to the `mvapich2-1.0.3` directory. We have provided a single script `make.mvapich2.udapl` for building over uDAPL.

Because there are many varieties of uDAPL packages, the script `make.mvapich2.udapl` requires two parameters: `dat_library_path` and `dat_include_path`, before it can start building the package.

By default, the compilation script uses `gcc`. In order to use a specific compiler, set the `CC`, `CXX`, `F77`, and `F90` variables to the appropriate commands. MVAPICH2 currently supports the GCC, Intel, PathScale, PGI, and Sun Studio compiler suites for uDAPL. In addition, it may prompt for options, such as cluster size, IO bus and link speed. The script then proceeds to configure, build and install the package.

- Customize MVAPICH2 to achieve the best performance on your cluster.

MVAPICH2 has many optimization schemes to enhance performance. These schemes can be turned on/off by the compilation flags which are listed in the variable CFLAGS (in the default compilation script) and/or various configure options.

- *Header Caching*: Controlled by `-DUSE_HEADER_CACHING`. Provides better performance for small messages. Note that you must also define `-DRDMA_FAST_PATH` to enable this scheme.
- *Shared-memory communication*: Controlled by `-D_SMP_`. Intra-node communication uses shared-memory operations instead of loopback from the HCA.
- *Optimizations for one sided communication*: Controlled by `-DONE_SIDED`. One sided operations can be directly built on RDMA operations. Latency and throughput of one sided operations can be improved.
- *Shared library support*: To enable shared library support, add the flag `-enable-sharedlibs=kind` to configure in the default `make.mvapich2.udapl` script, where `kind` should be “gcc” for the GCC, Intel, PathScale, and PGI compilers on Intel and Opteron Hardware and “solaris86-cc” for the Sun Studio compiler on Solaris Opteron hardware for the uDAPL build script specifically. Please refer to Section 5.3.1 for more details.
- *TotalView Debugger support*: This enables the use of TotalView debugger for your MPI applications. In order to enable this feature, you need to pass `-enable-sharedlibs=kind` as noted above for shared library support and `-enable-g=dbg` options to configure in the `make.mvapich2.udapl` script. Please refer to Section 5.3.2 for more details.

After setting all the parameters, the script `make.mvapich2.udapl` configures, builds and installs the entire package in the directory specified by the variable PREFIX.

4.4.3 Build MVAPICH2 with VAPI

There are several methods to configure MVAPICH2 and optimize it for a given platform.

- Use of the Default configuration

Go to the `mvapich2-1.0.3` directory. We have included a single script for VAPI (`make.mvapich2.vapi`) that takes care of different platforms, compilers and architectures. By default, the compilation script uses `gcc`. In order to use a specific compiler, set the `CC`, `CXX`, `F77`, and `F90` variables to the appropriate commands. MVAPICH2 currently supports the GCC, Intel, PathScale and PGI compiler suites for VAPI. The platform/architecture is detected automatically. The script prompts the user for parameters like I/O Bus type (PCI-X, PCI-Ex) and InfiniBand link speed (SDR, DDR).

- Customize MVAPICH2 to achieve the best performance on your cluster. Most of the configuration options for MVAPICH2 over VAPI are similar to the ones over uDAPL (Section 4.4.2).

After setting all the parameters, the script `make.mvapich2.vapi` configures, builds and installs the entire package in the directory specified by the variable PREFIX.

4.4.4 Build MVAPICH2 with TCP/IPoIB

Go to the `mvapich2-1.0.3` directory. We have provided a script `make.mvapich2.tcp` for building MVAPICH2 over TCP/IP intended for use over IPoIB (IP over InfiniBand). By default, the compilation script uses `gcc`. In order to use a specific compiler, set the `CC`, `CXX`, `F77`, and `F90` variables to the appropriate commands. MVAPICH2 currently supports the GCC, Intel, PathScale, PGI, and Sun Studio compiler suites for TCP/IPoIB. Simply execute this script (e.g. `./make.mvapich2.tcp`) for completing your build.

5 Basic Usage Instructions

5.1 Compile MPI Applications

MVAPICH2 provides a variety of MPI compilers to support applications written in different programming languages. Please use `mpicc`, `mpif77`, `mpiCC`, or `mpif90` to compile applications. The correct compiler should be selected depending upon the programming language of your MPI application.

These compilers are available in the `MVAPICH2_HOME/bin` directory. MVAPICH2 installation directory can also be specified by modifying `$PREFIX`, then all the above compilers will also be present in the `$PREFIX/bin` directory.

5.2 Setting MPD Environment

Prerequisites: `ssh` should be enabled between the front nodes and the computing nodes.

Please follow these steps to setup MPD:

- Please ensure that you have `.mpd.conf` and `.mpdpasswd` in your home directory. They are typically a single line file containing the following:

```
secretword=56rtG9.
```

The content of `.mpd.conf` and `.mpdpasswd` should be exactly the same. (Ofcourse, the secretword should be different.)

- Please include MPD path into your path

```
$ export MPD_BIN=$MVAPICH2_HOME/bin
$ export PATH=$MVAPICH2_HOME/bin:$PATH
```

`$MVAPICH2_HOME` is the installation path of your MVAPICH2, as specified by `$PREFIX` when you configure MVAPICH2.

- Prepare hostfile

Specify the hostnames of the compute nodes in a file. If you have a hostfile like:

```
n0
n1
n2
n3
```

then one process per node will be started on each one of these compute nodes.

- Start MPD Daemons on the compute nodes

```
$ mpdboot -n 4 -f hostfile
```

Note: The command, `mpdboot`, also takes a default hostfile name `mpd.hosts`. If you have created the hostfile as `mpd.hosts`, you can omit the option “-f hostfile”.

- Check status of MPD Daemons on the compute nodes (optional)

```
$ mpdtrace
```

This should list all the nodes specified in the hostfile, not necessarily in the order specified in the hostfile.

Up to now, we have specified setting up the environment, which is independent of the underlying device supported by MVAPICH2. In the next sections, we present details specific to different devices.

5.3 Run MPI Applications Using mpiexec with OpenFabrics Gen2-IB or VAPI Device

To start multiple processes, mpiexec can be used in the following fashion:

```
$ mpiexec -n 4 ./cpi
```

Four processes will be started on the compute nodes n0, n1, n2 and n3. mpiexec can also be run with several options. `$ mpiexec --help` lists all the possible options. A useful option is to specify a machinefile which holds the process mapping to machines. It can also be used to specify the number of processes to be run on each host. The machinefile option can be used with mpiexec as follows:

```
$ mpiexec -machinefile mf -n 4 ./cpi
```

where the machine file "mf" contains the process to machine mapping. For example, if you want to run all the 4 processes on n0, then "mf" contains the following lines:

```
$ cat mf
n0
n0
n0
n0
```

Environmental variables can be set with mpiexec as follows:

```
$ mpiexec -n 4 -env ENV1 value1 -env ENV2 value2 ./cpi
```

Note that the environmental variables should be put immediately before the executable file. The mpiexec command also propagates exported variables in its runtime environment to all processes by default. Exporting a variable before running mpiexec has the same effect as explicitly passing its value with the -env command line option. The command above could be done in the following manner when using a Bourne shell derivative:

```
$ export ENV1=value1
$ export ENV2=value2
$ mpiexec -n 4 ./cpi
```

5.3.1 Run MPI Applications using Shared Library Support

MVAPICH2 provides shared library support. This feature allows you to build your application on top of MPI shared library. If you choose this option, you still will be able to compile applications with static libraries. But as default, when you have shared library support enabled your applications will be built on top of shared libraries automatically. the following commands provide some examples of how to build and run your application with shared library support.

- To compile your application with shared library support. Run the following command:
`$ mpicc -o cpi cpi.c`
- To execute an application compiled with shared library support, you need to specify the path to the shared library by putting `LD_LIBRARY_PATH=path-to-shared-libraries` in the command line. For example, `$ mpiexec -np 2 -env LD_LIBRARY_PATH $MVAPICH2_BUILD/lib/shared ./cpi` Again, note that “`LD_LIBRARY_PATH=path-to-shared-libraries`” should be put immediately before the executable file. As with previous cases, exporting the `LD_LIBRARY_PATH` environment variable also works.
- To disable MVAPICH2 shared library support even if you have installed MVAPICH2, run the following command:
`$ mpicc -noshlib -o cpi cpi.c`

5.3.2 Run MPI Application using TotalView Debugger Support

MVAPICH2 provides TotalView support. The following commands provide an example of how to build and run your application with TotalView support. Note: running TotalView requires correct setup in your environment, if you encounter any problem with your setup, please check with your system administrator for help.

- Define `ssh` as a `TVDSVRLAUNCHCMD` variable in your default shell. For example, with `bashrc`, you can do:
`$ echo "export TVDSVRLAUNCHCMD=ssh" >> ~/.bashrc`
- Configure MVAPICH2 with the configure options `--enable-g=dbg --enable-sharedlibs=kind --enable-debuginfo` in addition to the default options and then build MVAPICH2.
- Compile your program with a flag `-g`:
`$ mpicc -g -o prog prog.c`
- Define the correct path to TotalView as the `TOTALVIEW` variable. For example, under `bash` shell:
`$ export MPIEXEC_TOTALVIEW=<path_to_TotalView>`
- Run your program:
`$ mpiexec -tv -np 2 -env LD_LIBRARY_PATH $MVAPICH2_BUILD/lib/shared: $MVAPICH2_BUILD/lib prog`

- Troubleshooting:
 - X authentication errors: check if you have enabled X Forwarding
`$ cat "ForwardX11 yes" >> $HOME/.ssh/config`
 - ssh authentication error: ssh to the computer node with its long form hostname, for example, ssh i0.domain.osu.edu

5.4 Run MPI Application using OpenFabrics Gen2-iWARP Device

In MVAPICH2, Gen2-iWARP support is enabled with the use of the run time environment variable ‘MV2_USE_IWARP_MODE’.

In addition to this flag, all the systems to be used need the following one time setup for enabling RDMA CM usage.

- **Setup the RDMA CM device:** RDMA CM device needs to be setup, configured with an IP address and connected to the network.
- **Setup the Local Address File:** Create the file (`/etc/mv2.conf`) with the local IP address to be used by RDMA CM. (Multiple IP addresses can be listed (one per line) for multirail configurations).
`$ echo 10.1.1.1 >> /etc/mv2.conf`

Programs can be executed as follows:

```
$ mpiexec -n 4 -env MV2_USE_IWARP_MODE 1 -env ENV1 value1 prog
```

The iWARP device also provides totalview debugging and shared library support. Please refer to section 5.3.1 and 5.3.2 for shared library and totalview support, respectively.

5.5 Run MPI Application using mpiexec with uDAPL Device

MVAPICH2 can be configured with the uDAPL device, as described in the section 4.4.2 . To compile MPI applications, please refer to the section 5.1. In order to run MPI applications with uDAPL support, please specify the environmental variable MV2_DAPL_PROVIDER. As an example,

```
$ mpiexec -n 4 -env MV2_DAPL_PROVIDER ib0 ./cpi
```

or:

```
$ export MV2_DAPL_PROVIDER=ib0
```

```
$ mpiexec -n 4 ./cpi
```

Please check the `/etc/dat.conf` file to find all the available uDAPL service providers. The default value for the uDAPL provider will be chosen, if no environment variable is provided at runtime. For OFED 1.2, please use `OpenIB-cma` as the uDAPL provider.

The uDAPL device also provides totalview debugging and shared library support. Please refer to section 5.3.1 and 5.3.2 for shared library and totalview support, respectively.

5.6 Run MPI Application using mpiexec with TCP/IP

If you would like to run an MPI job using IPoIB but your IB card is not the default interface for ip traffic you have two options. For both of the options , assume that you have a cluster setup as the following:

#hostname	Eth Addr	IPoIB Addr
compute1	192.168.0.1	192.168.1.1
compute2	192.168.0.2	192.168.1.1
compute3	192.168.0.3	192.168.1.1
compute4	192.168.0.4	192.168.1.1

5.6.1 The MPI Job Uses IPoIB

In this scenario, you will start up mpd like normal. However, you will need to create a machine file for mpiexec that tells mpiexec to use a particular interface. Example:

```
$ cat - > $(MPD_HOSTFILE) compute1
compute2
compute3
compute4
```

```
$ mpdboot -n 4 -f $(MPD_HOSTFILE)
compute1 ifhn=192.168.1.1
compute2 ifhn=192.168.1.2
compute3 ifhn=192.168.1.3
compute4 ifhn=192.168.1.4
```

The ifhn portion tells mpiexec to use the interface associated with that ip address for each machine. You can now run your MPI application using IPoIB similar to the following.

```
$ mpiexec -n $(NUM_PROCESS) -f $(MACHINE_FILE) $(MPI_APPLICATION)
```

5.6.2 Both MPD And the MPI Job Use IPoIB

In this scenario you will start up mpd in a modified fashion. However, you will not need to create a machine file for mpiexec. Your hostsfile for mpdboot must contain the ip addresses, or hostnames mapped to these addresses, of each machine's IPoIB interface. The only exception is that you do not list the ip address or hostname of the local machine. This will be specified on the command line of the mpdboot command using the -ifhn option. Example:

```
$ cat - > $(MPD_HOSTFILE)
192.168.1.2
```

```
192.168.1.3
192.168.1.4
```

```
$ mpdboot -n 4 -f $HOSTSFILE --ifhn=192.168.1.1
```

The `-ifhn` option tells `mpdboot` to use the interface corresponding to that ip address to create the `mpd` ring and run MPI jobs. You can now run your MPI application using IPoIB similar to the following.

```
$ mpiexec -n $(NUM_PROCES) $(MPI_APPLICATION)
```

Note: For both options, you can replace the IPoIB addresses with aliases

5.7 Run MPI Applications using SLURM

SLURM is an open-source resource manager designed by Lawrence Livermore National Laboratory. SLURM software package and its related documents can be downloaded from:

<http://www.llnl.gov/linux/slurm/>

Once SLURM is installed and the daemons are started, applications compiled with MVAPICH2 can be launched by SLURM, e.g.

```
$ srun -n2 --mpi=none ./a.out
```

The use of SLURM enables many good features such as explicit CPU and memory binding. For example, if you have two processes and want to bind the first process to CPU 0 and Memory 0, and the second process to CPU 4 and Memory 1, then it can be achieved by:

```
$ srun --cpu_bind=v,map_cpu:0,4 --mem_bind=v,map_mem:0,1 -n2 --mpi=none ./a.out
```

For more information about SLURM and its features please visit SLURM website.

6 Advanced Usage Instructions

In this section, we present the usage instructions for advanced features provided by MVAPICH2.

6.1 Run MPI applications on Multi-Rail Configurations (for OpenFabrics Gen2-IB and Gen2-iWARP Devices)

MVAPICH2 has integrated multi-rail support. Run-time variables are used to specify the control parameters of the multi-rail support; number of adapters with `MV2_NUM_HCAS` (section 10.18), number of ports per adapter with `MV2_NUM_PORTS` (section 10.19), and number of queue pairs per port with `MV2_NUM_QP_PER_PORT` (section 10.20). Those variables are default to 1 if you do not specify them. Following is an example to run multi-rail support with two adapters, using one port per adapter and one queue pair per port:

```
$ mpiexec -n 2 -env MV2_NUM_HCAS 2 -env MV2_NUM_PORTS 1 -env MV2_NUM_QP_PER_PORT 1 prog
```

Note that you don't need to put `-env MV2_NUM_PORTS 1 -env MV2_NUM_QP_PER_PORT 1` since they default to 1, so you can type:

```
$ mpiexec -n 2 -env MV2_NUM_HCAS 2 prog
```

6.2 Run MPI application with Customized Optimizations (for OpenFabrics Gen2-IB and Gen2-iWARP Devices)

In MVAPICH2-1.0.3, run-time variables are used to switch various optimization schemes on and off. Following is a list of optimizations schemes and the control environmental variables, for a full list please refer to the section 10:

- **Adaptive RDMA fast path:** using RDMA write to enhance performance for short messages. Default: on; to disable:

```
$ mpiexec -n 2 -env MV2_USE_RDMA_FAST_PATH 0 prog
```

- **Shared-receive Queue:** This feature is available only with Gen2-IB devices. This is targeted for using Shared Receive Queue (SRQ). Default: on; to disable:

```
$ mpiexec -n 2 -env MV2_USE_SRQ 0 prog
```

- **Optimizations for one sided communication:** One sided operations can be directly built on RDMA operations. Currently this scheme will be disabled if on-demand connection management is used. Default: on; to disable:

```
$ mpiexec -n 2 -env MV2_USE_RDMA_ONE_SIDED 0 prog
```

- **Lazy memory unregistration:** user-level registration cache. Default: on; to disable:
`$ mpiexec -n 2 -env MV2_USE_LAZY_MEM_UNREGISTER 0 prog`

6.3 Run MPI application with Checkpoint/Restart Support (for OpenFabrics Gen2-IB Device)

MVAPICH2 provides system-level checkpoint/restart functionality for the OpenFabrics Gen2-IB interface. Three methods are provided to invoke checkpointing: Manual, Automated and Application Initiated Synchronous Checkpointing. In order to utilize the checkpoint/restart functionality there a couple of steps that need to be followed.

- Download and install the BLCR (Berkeley Lab's Checkpoint/Restart) package. The packages can be downloaded from [this webpage](#).
- Make sure the BLCR packages are installed on every node and the `LD_LIBRARY_PATH` must contain the path to the shared library of BLCR, usually `$BLCR_HOME/lib`.
- MVAPICH2 needs to be compiled with checkpoint/restart support, see section 4.4.1.
- BLCR's kernel modules must be loaded on all the compute nodes.
- Make sure the `PATH` contains the path to the executables of BLCR, usually `$BLCR_HOME/bin`.
- The extended version of Multi-Purpose Daemon (MPD), provided in MVAPICH2 package, needs to be used as the process manager.

And users are strongly encouraged to read the [Administrators guide](#) of BLCR, and test the BLCR on the target platform, before using the checkpointing feature of MVAPICH2.

Now, your system is set up to use the Checkpoint/Restart features of MVAPICH2. There are several parameters related to MVAPICH2 to be setup to control the configuration and useage of this feature.

- `MV2_CKPT_FILE`: This parameter specifies the path and the base filename for checkpoint files of MPI processes. Please note that File System performance is critical to the performance of checkpointing. This parameter controls which file system will be used to store the checkpoint files. For example, if your PVFS2 is mounted at `/mnt/pvfs2`, using `export MV2_CKPT_FILE=/mnt/pvfs2/ckptfile` will let the checkpoint files being stored in pvfs2 file system. See Section 10.1 for details.
- `MV2_CKPT_INTERVAL`: This parameter can be used to enable automatic checkpointing. See Section 10.2 for details.
- `MV2_CKPT_MAX_SAVE_CKPTS`: This parameter is used to limit the number of checkpoints saved on file system. See Section 10.3 for details.

- `MV2_CKPT_NO_SYNC`: This parameter is used to control whether the program forces the checkpoint files being synced to disk or not before it continues execution. See Section 10.6 for details.
- `MV2_CKPT_MPD_BASE_PORT`, See Section 10.4 for details.
- `MV2_CKPT_MPIEXEC_PORT`, See Section 10.4 for details.

In order to provide maximum flexibility to end users who wish to use the checkpoint/restart features of MVAPICH2, we've provided three different methods which can be used to take the checkpoints during the execution of the MPI application. These methods are described as follows:

- **Manual Checkpointing:** In this mode, the user simply launches an MPI application and chooses when to checkpoint the application. This mode can be primarily used for experimentation during deployment stages. In order to use this mode, the MPI application is launched normally using `mpiexec`. When the user decides to take a checkpoint, the users can issue a BLCR command called `'cr_checkpoint'` with the process id (PID) of the `mpiexec` process. In order to simplify the process, the script `mv2_checkpoint` can be used. This script is available in the same directory as `mpiexec`.
- **Automated Checkpointing:** In this mode, the user can launch the MPI application normally using `mpiexec`. However, instead of manually issuing checkpoints as described in the above bullet, a parameter (`MV2_CKPT_INTERVAL`) can be set to automatically take checkpoints and user-defined intervals. Please refer to Section 10.2 for a complete usage description of this variable. This mode can be used to take checkpoints of a long running application, for example every 1 hour, 2 hours etc. based on user's choice.
- **Application Initiated Synchronous Checkpointing:** In this mode, the MPI application which is running can itself request for a checkpoint. Application can request a whole program checkpoint synchronously by calling `MVAPICH2.Sync_Checkpoint`. Note that it is a collective operation, and this function must be called from all processes to take the checkpoint. Synchronous checkpointing is enabled by default when `ENABLE_CKPT=yes` is set in make script. To disable synchronous checkpointing, remove `-DSYNC_CKPT` from `CR_FLAG` in the make script. This mode is expected to be used by applications that can be modified and have well defined synchronization points. These points can be effectively used to take checkpoints. An example of how this mode can be activated is given below.

```
#include "mpi.h"
#include <unistd.h>
#include <stdio.h>

int main(int argc, char *argv[])
{
    MPI_Init(&argc, &argv);
    printf("Computation\n");
    sleep(5);
}
```

```

MPI_Barrier(MPI_COMM_WORLD);
MVAPICH2_Sync_Checkpoint();
MPI_Barrier(MPI_COMM_WORLD);
printf("Computation\n");
sleep(5);
MPI_Finalize();
return 0;
}

```

To restart a job from a checkpoint, users need to issue another command of BLCR, ‘‘cr_restart’’ with the checkpoint file name of the MPI job console as the parameter, usually `context.<pid>`. The checkpoint file name of the MPI job console can be specified when issuing the checkpoint, see the ‘‘cr_checkpoint --help’’ for more information. Please note that the names of checkpoint files of the MPI processes will be assigned according to the environment variable `MV2_CKPT_FILE`, (`$MV2_CKPT_FILE.<number of checkpoint>.<process rank>`).

Please refer to the Section 8.7 for troubleshooting with Checkpoint/Restart.

6.4 Run MPI application with RDMA CM support (for OpenFabrics Gen2-IB and Gen2-iWARP Devices)

In MVAPICH2, for using RDMA CM the runtime variable `MV2_USE_RDMA_CM` needs to be used as described in 10. (Note: In order to use RDMA CM support on OFED 1.1 software stack, please add the flag `-DOFED_VERSION_1_1` to the cflags in the compilation script `make.mvapich2.ofa`)

In addition to these flags, all the systems to be used need the following one time setup for enabling RDMA CM usage.

- **Setup the RDMA CM device:** RDMA CM device needs to be setup, configured with an IP address and connected to the network.
- **Setup the Local Address File:** Create the file (`/etc/mv2.conf`) with the local IP address to be used by RDMA CM. (Multiple IP addresses can be listed (one per line) for multirail configurations).

```
$ echo 10.1.1.1 >> /etc/mv2.conf
```

Programs can be executed as follows:

```
$ mpiexec -n 2 -env MV2_USE_RDMA_CM 1 prog
```

6.5 Run MPI application with Shared Memory Collectives

In MVAPICH2, support for shared memory based collectives has been enabled for MPI applications running over OpenFabrics Gen2-IB, Gen2-iWARP and uDAPL stack. Currently, this support is available for the following collective operations:

- `MPI_Allreduce`
- `MPI_Reduce`
- `MPI_Barrier`

Optionally, these feature can be turned off at runtime by using the following parameters:

- `MV2_USE_SHMEM_COLL` (section 10.48)
- `MV2_USE_SHMEM_ALLREDUCE` (section 10.46)
- `MV2_USE_SHMEM_REDUCE` (section 10.49)
- `MV2_USE_SHMEM_BARRIER` (section 10.47)

Please refer to Section 10 for further details.

6.6 Run MPI Application with Hot-Spot and Congestion Avoidance (for OpenFabrics Gen2-IB Device)

MVAPICH2 supports hot-spot and congestion avoidance using InfiniBand multi-pathing mechanism. This support is available for MPI applications using OpebFabrics stack and InfiniBand adapters.

To enable this functionality, a run-time variable, `MV2_USE_HSAM` (Section 10.37) can be enabled, as shown in the following example:

```
$ mpiexec -n 2 -env MV2_USE_HSAM 1 ./cpi
```

This functionality automatically defines the number of paths for hot-spot avoidance. Alternatively, the maximum number of paths to be used between a pair of processes can be defined by using a run-time variable `MV2_NUM_QP_PER_PORT` (Section 10.20).

We expect this functionality to show benefits in the presence of at least partially non-overlapping paths in the network. OpenSM, the subnet manager distributed with OpenFabrics supports LMC mechanism, which can be used to create multiple paths:

```
$ opensm -l4
```

will start the subnet manager with LMC value to four, creating sixteen paths between every pair of nodes.

6.7 Run MPI Application with Network Fault Tolerance Support (for OpenFabrics Gen2-IB Device)

MVAPICH2 supports network fault recovery by using InfiniBand Automatic Path Migration mechanism. This support is available for MPI applications using OpebFabrics stack and InfiniBand adapters.

To enable this functionality, a run-time variable, `MV2_USE_APM` (section 10.33) can be enabled, as shown in the following example:

```
$ mpiexec -n 2 -env MV2_USE_APM 1 ./cpi
```

MVAPICH2 also supports testing Automatic Path Migration in the subnet in the absence of network faults. This can be controlled by using a run-time variable `MV2_USE_APM_TEST` (section 10.34). This should be combined with `MV2_USE_APM` as follows:

```
$ mpiexec -n 2 -env MV2_USE_APM 1 -env MV2_USE_APM_TEST 1 ./cpi
```

7 Using OSU Benchmarks

If you have arrived at this point, you have successfully installed MVAPICH2. Congratulations!! In the `mvapich2-1.0.3/osu_benchmarks` directory, we provide these basic performance tests:

- One-way latency test (`osu_latency.c`)
- One-way bandwidth test (`osu_bw.c`)
- Bi-directional bandwidth (`osu_bibw.c`)
- Multiple Bandwidth / Message Rate test (`osu_mbw_mr.c`)
- One-sided put latency (`osu_put_latency.c`)
- One-sided put bandwidth (`osu_put_bw.c`)
- One-sided put bi-directional bandwidth (`osu_put_bibw.c`)
- One-sided get latency (`osu_get_latency.c`)
- One-sided get bandwidth (`osu_get_bw.c`)
- One-sided accumulate latency (`osu_acc_latency.c`)
- One-way multi-threaded latency test (`osu_latency_mt.c`) - Multi-threading support must be compiled in to run this test.

These benchmarks are also available on our project's web page. Sample performance numbers for these benchmarks on representative platforms with InfiniBand and iWARP adapters are also included on our projects' web page. You are welcome to compare your performance numbers with our numbers. If you see any big discrepancy, please let us know by sending an email to mvapich-discuss@cse.ohio-state.edu.

8 FAQ and Troubleshooting with MVAPICH2

Based on our experience and feedback we have received from our users, here we include some of the problems a user may experience and the steps to resolve them. If you are experiencing any other problem, please feel free to contact us by sending an email to mvapich-discuss@cse.ohio-state.edu.

MVAPICH2 can be used over five underlying transport interfaces, namely OpenFabrics (Gen2), OpenFabrics (Gen2-iWARP), VAPI, UDAPL and TCP/IP. Based on the underlying library being utilized, the troubleshooting steps may be different. However, some of the troubleshooting hints are common for all underlying libraries. Thus, in this section, we have divided the troubleshooting tips into four sections: General troubleshooting and Troubleshooting over any one of the five transport interfaces.

8.1 General Questions and Troubleshooting

8.1.1 Invalid Communicators Error

This is a problem which typically occurs due to the presence of multiple installations of MVAPICH2 on the same set of nodes. The problem is due to the presence of `mpi.h` other than the one, which is used for executing the program. This problem can be resolved by making sure that the `mpi.h` from other installation is not included.

8.1.2 Are `fork()` and `system()` supported?

`fork()` and `system()` is supported for the OpenFabrics device as long as the kernel is being used is Linux 2.6.16 or newer. Additionally, the version of OFED used should be 1.2 or higher. The environment variable `IBV_FORK_SAFE=1` must also be set to enable fork support.

8.1.3 Cannot Build with the PathScale Compiler

There is a known bug with the PathScale compiler (before version 2.5) when building MVAPICH2. This problem will be solved in the next major release of the PathScale compiler. To work around this bug, use the the “`-LNO:simd=0`” C compiler option. This can be set in the build script similarly to:

```
export CC='pathcc -LNO:simd=0'
```

Please note the use of double quotes. If you are building shared libraries and are using the PathScale compiler (version below 2.5), then you should add “`-g`” to your `CFLAGS`, in order to get around a compiler bug.

8.1.4 Cannot find `mpd.conf`

If you get this error, please set your `.mpd.conf` and `.mpdpasswd` files.

8.1.5 sched_setaffinity: Bad address

MVAPICH2 supports CPU affinity for multi-way SMP systems. This feature requires a kernel version of 2.6. The above error reports that CPU affinity cannot be enabled on the system. This feature can be disabled by setting `-env MV2_ENABLE_AFFINITY 0`.

8.1.6 Multi-threaded programs seem to run sequentially

MVAPICH2 uses CPU affinity to have better performance for single-threaded programs. For multi-threaded programs, however, it may schedule all the threads of a process to run on the same CPU. CPU affinity should be disabled in this case to solve the problem, i.e. `set -env MV2_ENABLE_AFFINITY 0`.

8.1.7 Running MPI programs built with gfortran

MPI programs built with gfortran might not appear to run correctly due to the default output buffering used by gfortran. If it seems there is an issue with program output, the `GFORTTRAN_UNBUFFERED_ALL` variable can be set to “y” and exported into the environment before using the `mpiexec` command to launch the program, as done in the bash shell example below:

```
$ export GFORTRAN_UNBUFFERED_ALL=y
```

8.2 With Gen2 Interface

8.2.1 Cannot Open HCA

The above error reports that the InfiniBand Adapter is not ready for communication. Make sure that the drivers are up. This can be done by executing the following command which gives the path at which drivers are setup.

```
% locate libibverbs
```

8.2.2 Checking state of IB Link

In order to check the status of the IB link, one of the following commands can be used:

```
% ibstatus
```

```
or
```

```
% ibv_devinfo.
```

8.2.3 Undefined reference to `ibv_get_device_list`

Add `-DGEN2_OLD_DEVICE_LIST_VERB` macro to `CFLAGS` and rebuild `MVAPICH2-gen2`. If this happens, this means that your Gen2 installation is old and needs to be updated.

8.2.4 Creation of CQ or QP failure

A possible reason could be inability to pin the memory required. Make sure the following parameters are set.

In `/etc/security/limits.conf` add the following line:

```
* soft memlock phys_mem_in_KB
```

After this, add the following line to `/etc/init.d/sshd` and restart `sshd`

```
ulimit -l phys_mem_in_KB
```

8.2.5 Hang with the HSAM Functionality

HSAM functionality uses multi-pathing mechanism with LMC functionality. However, some versions of OpenFabrics Drivers (including OpenFabrics Enterprise Distribution (OFED) 1.1) and using the Up*/Down* routing engine do not configure the routes correctly using the LMC mechanism. We strongly suggest to upgrade to OFED 1.2, which supports Up*/Down* routing engine and LMC mechanism correctly.

8.2.6 Failure with Automatic Path Migration

`MVAPICH2` provides network fault tolerance with Automatic Path Migration (APM). However, APM is supported only with OFED 1.2 onwards. With OFED 1.1 and prior versions of OpenFabrics drivers, APM functionality is not completely supported. Please refer to Section 10.33 and section 10.34

8.2.7 Error opening file

If you configure `MVAPICH2` with `RDMA_CM` and see this error, you need to verify if you have setup up the local IP address to be used by `RDMA_CM` in the file `/etc/mv2.conf`. Further, you need to make sure that this file has the appropriate file read permissions. Please follow Section 6.4 for more details on this.

8.2.8 RDMA CM Address error

If you get this error, please verify that the IP address specified `/etc/mv2.conf` is correctly specified with the IP address of the device you plan to use `RDMA_CM` with.

8.2.9 RDMA CM Route error

If see this error, you need to check whether the specified network is working or not.

8.3 With Gen2-iWARP Interface

8.3.1 Error opening file

If you configure MVAPICH2 with RDMA_CM and see this error, you need to verify if you have setup up the local IP address to be used by RDMA_CM in the file `/etc/mv2.conf`. Further, you need to make sure that this file has the appropriate file read permissions. Please follow Section 5.4 for more details on this.

8.3.2 RDMA CM Address error

If you get this error, please verify that the IP address specified `/etc/mv2.conf` is correctly specified with the IP address of the device you plan to use RDMA_CM with.

8.3.3 RDMA CM Route error

If see this error, you need to check whether the specified network is working or not.

8.4 With VAPI Interface

8.4.1 Cannot pass MPI_Init

If your MPI application cannot pass `MPI_Init`, please make sure of the following things:

- If you have enabled `ssh` based startup, make sure that you have set up `ssh` keys for login into all the nodes without any password prompt.
- Please run some network-level programs to make sure the network used is working. E.g. running `perf_main` (distributed with IBGD) is a good choice to test whether InfiniBand is working.

8.4.2 Cannot Open HCA

The above error reports that the InfiniBand Adapter is not ready for communication. Make sure that the drivers are up. This can be done by executing

```
$ locate libvapi
```

which gives the path at which drivers are setup.

8.4.3 Cannot include vapi.h

This error is generated during compilation, if the correct path to the InfiniBand library installation is not given.

Please setup the environment variable MTHOME as

```
/usr/local/ibgd/driver/infinihost
```

Please setup the environment variable MTHOME as

```
$ export MTHOME=/usr/local/ibgd/driver/infinihost
```

If the problem persists, please contact your system administrator or send an email to mvapich-discuss@cse.ohio-state.edu.

8.4.4 VAPI_RETRY_EXEC_ERROR

This error usually indicates that all InfiniBand links the MPI application is trying to use are not in the PORT_ACTIVE state. Please make sure that all ports show PORT_ACTIVE with the VAPI utility `vstat`.

8.4.5 ld: multiple definitions of symbol _calloc error on MacOS

Please make sure that the environmental variable "MAC_OSX" is set before your configuration. If you use manual configuration and not `make.mvapich2.vapi`, you must configure MVAICH2 in the following way:

```
% export MAC_OSX=yes % ./configure ..with option.. % make % make install
```

If you encounter this problem compiling your own applications,

```
"ld: multiple definitions of symbol _calloc
/usr/lib/libm.dylib(malloc.So) definition of _calloc
/tmp/mvapich2-0.9.3/lib/libmpich.a(dreg-g5.o)
definition of _calloc in section (__TEXT,__text)
ld: multiple definitions of symbol _free
/usr/lib/libm.dylib(malloc.So) definition of _free
/tmp/mvapich2-0.9.3/mvapich/lib/libmpich.a(dreg-g5.o)
definition of _free in section (__TEXT,__text) "
```

it is likely that you have explicitly included "-lm". You should remove that.

8.4.6 No Fortran interface on the MacOS platform

To enable Fortran support, you would need to install the IBM compiler located at (there is a 60-day free trial version) available from [IBM](#).

Once you unpack the tarball, you can customize and use `make.mvapich2.vapi` to compile and install the package or manually configure, compile and install the package.

8.5 With UDAPL Interface

8.5.1 Cannot Open IA

If you configure MVAPICH2 with uDAPL and see this error, you need to check whether you have specified the correct uDAPL service provider (at the end of Section 5.3). If you have specified the uDAPL provider but still see this error, you need to check whether the specified network is working or not. For OFED 1.2, please use `OpenIB-cma` as the uDAPL provider.

8.5.2 DAT Insufficient Resource

If you configure MVAPICH2 with uDAPL and see this error, you need to reduce the value of the environmental variable `RDMA_DEFAULT_MAX_WQE` depending on the underlying network.

8.5.3 Cannot find libdat.so

If you get the error: “error while loading shared libraries, libdat.so”, The location of the dat shared library is incorrect. You need to find the correct path of libdat.so and export `LD_LIBRARY_PATH` to this correct location. For example:

```
$ export LD_LIBRARY_PATH=/path/to/libdat.so:$LD_LIBRARY_PATH
$ mpiexec -n 2 ./a.out
```

8.5.4 Cannot find mpd.conf

If you get this error, please set your `.mpd.conf` and `.mpdpasswd` files as mentioned in Section 5.3.

8.6 The MPD mpiexec fails with “no msg recvd from mpd when expecting ack of request.”

This failure may be an indication that there is a problem with your cluster configuration. If you are confident in the correctness of your cluster configuration, then you can tune the timeout with `MV2_MPD_RECVTIMEOUT_MULTIPLIER`.

8.7 Checkpoint/Restart

Please make sure the following things for a successful restart:

- The MPD must be started on all the compute nodes and the console node before a restart.
- The BLCR modules must be loaded on all the compute nodes and the console node before a restart
- The checkpoint file of MPI job console must be accessible from the console node.
- The corresponding checkpoint files of the MPI processes must be accessible from the compute nodes using the same path as when checkpoint was taken.

The following things can cause a restart to fail:

- The job which was checkpointed is not terminated or the some processes in that job are not cleaned properly. Usually they will be cleaned automatically, otherwise, since the pid can't be used by BLCR's to restart, restart will fail.
- The processes in the job have opened temporary files and these temporary files are removed or not accessible from the nodes where the processes are restarted on.

FAQ regarding Berkeley Lab Checkpoint/Restart (BLCR) can be found at <http://upc-bugs.lbl.gov/blcr/doc/html/FAQ.html> And the userguide for BLCR can be found at http://upc-bugs.lbl.gov/blcr/doc/html/BLCR_Users_Guide.html

If you encounter any problem with the Checkpoint/Restart support, please feel free to contact us as mvapich-discuss@cse.ohio-state.edu.

9 Scalable features for Large Scale Clusters and Performance Tuning

MVAPICH2 provides many different parameters for tuning performance for a wide variety of platforms and applications. These parameters can be either compile time parameters or runtime parameters. please refer to Section 9 for a complete description of all these parameters. In this section we classify these parameters depending on what you are tuning for and provide guidelines on how to use them.

9.1 RDMA Based Point-to-Point tuning

The following parameters are important in tuning the memory requirements for adaptive rdma fast path feature.

- MV2_VBUF_TOTAL_SIZE
- MV2_NUM_RDMA_BUFFER
- MV2_RDMA_VBUF_POOL_SIZE

MV2_RDMA_VBUF_POOL_SIZE is a fixed number of pool of vbufs. These vbufs can be shared among all different connections depending on the communication needs of each connection.

On the other hand, the product of MV2_VBUF_TOTAL_SIZE and MV2_NUM_RDMA_BUFFER generally is a measure of the amount of memory registered for eager message passing. These buffers are not shared across connections.

In MVAPICH2-1.0, MV2_VBUF_TOTAL_SIZE is adjustable by environmental variables. Please refer to Section 10.53 for details.

9.2 Shared Receive Queue (SRQ) Tuning

The main environmental parameters controlling the behavior of the Shared Receive Queue design are:

- MV2_SRQ_SIZE (10.32)
- MV2_SRQ_LIMIT (10.31)

MV2_SRQ_SIZE is the maximum size of the Shared Receive Queue. You may increase this to value 1000 if the application requires very large number of processors (4K and beyond).

MV2_SRQ_LIMIT defines the low watermark for the flow control handler. This can be reduced if your aim is to reduce the number of interrupts.

9.3 Shared Memory Tuning

MVAPICH2 uses shared memory communication channel to achieve high-performance message passing among processes that are on the same physical node. The two main parameters which are used for tuning shared memory performance for small messages are `SMPI_LENGTH_QUEUE` (10.55) and `SMP_EAGER_SIZE` (10.54). The two main parameters which are used for tuning shared memory performance for large messages are `SMP_SEND_BUF_SIZE`(10.57) and `SMP_NUM_SEND_BUFFER` (10.56).

`SMPI_LENGTH_QUEUE` is the size of the shared memory buffer which is used to store outstanding small and control messages. `SMP_EAGER_SIZE` defines the switch point from Eager protocol to Rendezvous protocol.

Messages larger than `SMP_EAGER_SIZE` are packetized and sent out in a pipelined manner. `SMP_SEND_BUF_SIZE` is the packet size, i.e. the send buffer size. `SMP_NUM_SEND_BUFFER` is the number of send buffers.

9.4 On-demand Connection Management Tuning

MVAPICH2 uses on-demand connection management to reduce the memory usage of MPI library. There are 4 parameters to tune connection manager: `MV2_ON_DEMAND_THRESHOLD` (10.22), `MV2_CM_RECV_BUFFERS` (10.7), `MV2_CM_TIMEOUT` (10.9), and `MV2_CM_SPIN_COUNT` (10.8). The first one applies to Gen2-IB, Gen2-iWARP and uDAPL devices and the other three only apply to Gen2 device.

`MV2_ON_DEMAND_THRESHOLD` defines threshold for enabling on-demand connection management scheme. When the size of the job is larger than the threshold value, on-demand connection management will be used.

`MV2_CM_RECV_BUFFERS` defines the number of buffers used by connection manager to establish new connections. These buffers are quite small and are shared for all connections, so this value may be increased to 8192 for large clusters to avoid retries in case of packet drops.

`MV2_CM_TIMEOUT` is the timeout value associated with connection management messages via UD channel. Decreasing this value may lead to faster retries but at the cost of generating duplicate messages.

`MV2_CM_SPIN_COUNT` is the number of the connection manager polls for new control messages from UD channel for each interrupt. This may be increased to reduce the interrupt overhead when many incoming control messages from UD channel at the same time.

10 MVAPICH2 Parameters

10.1 MV2_CKPT_FILE

- Class: Run Time
- Default: /tmp/ckpt
- Applicable device(s): Gen2

This parameter specifies the path and the base filename for checkpoint files of MPI processes. The checkpoint files will be named as \$MV2_CKPT_FILE.<number of checkpoint>.<process rank>, for example, /tmp/ckpt.1.0 is the checkpoint file for process 0's first checkpoint. To checkpoint on network-based file systems, user just need to specify the path to it, such as /mnt/pvfs2/my_ckpt_file.

10.2 MV2_CKPT_INTERVAL

- Class: Run Time
- Default: 0
- Unit: minutes
- Applicable device(s): Gen2

This parameter can be used to enable automatic checkpointing. To let MPI job console automatically take checkpoints, this value needs to be set to the desired checkpointing interval. A zero will disable automatic checkpointing. Using automatic checkpointing, the checkpoint file for the MPI job console will be named as \$MV2_CKPT_FILE.<number of checkpoint>.auto. Users need to use this file for restart.

10.3 MV2_CKPT_MAX_SAVE_CKPTS

- Class: Run Time
- Default: 0
- Applicable device(s): Gen2

This parameter is used to limit the number of checkpoints saved on file system to save the file system space. When set to a positive value N, only the last N checkpoints will be saved.

10.4 MV2_CKPT_MPD_BASE_PORT

- Class: Run Time
- Default: 24678
- Applicable device(s): Gen2

This parameter specifies the ports of socket connections to pass checkpointing control messages between MPD manager and MPI process. Users need to have a set of unused ports starting with \$MV2_CKPT_MPD_BASE_PORT on the compute nodes. The used port will be the \$MV2_CKPT_MPD_BASE_PORT + <process rank> for each MPI processes.

10.5 MV2_CKPT_MPIEXEC_PORT

- Class: Run Time
- Default: 14678
- Applicable device(s): Gen2

This parameter specifies the port of the socket connection for passing checkpointing control messages on MPI job console node. Users need to have an unused port to be set to \$MV2_CKPT_MPIEXEC_PORT on the console node.

10.6 MV2_CKPT_NO_SYNC

- Class: Run Time
- Applicable device(s): Gen2

When this parameter is set to any value, the checkpoints will not be required to sync to disk. It can reduce the checkpointing delay in many cases. But if users are using local file system, or any parallel file system with local cache, to store the checkpoints, it is recommended not to set this parameter because otherwise the checkpoint files will be cached in local memory and will likely be lost upon failure.

10.7 MV2_CM_RECV_BUFFERS

- Class: Run Time
- Default: 1024
- Applicable device(s): Gen2

This defines the number of buffers used by connection manager to establish new connections. These buffers are quite small and are shared for all connections, so this value may be increased to 8192 for large clusters to avoid retries in case of packet drops.

10.8 MV2_CM_SPIN_COUNT

- Class: Run Time
- Default: 5000
- Applicable device(s): Gen2

This is the number of the connection manager polls for new control messages from UD channel for each interrupt. This may be increased to reduce the interrupt overhead when many incoming control messages from UD channel at the same time.

10.9 MV2_CM_TIMEOUT

- Class: Run Time
- Default: 500
- Unit: milliseconds
- Applicable device(s): Gen2

This is the timeout value associated with connection management messages via UD channel. Decreasing this value may lead to faster retries but at the cost of generating duplicate messages.

10.10 MV2_DAPL_PROVIDER

- Class: Run time
- Default: ofa-v2-ib0 (Linux DAPL v2.0), OpenIB-cma (Linux DAPL v1.2), ibd0 (Solaris)
- Applicable device(s): uDAPL

This is to specify the underlying uDAPL library that the user would like to use if MVAPICH2 is built with uDAPL.

10.11 MV2_DEFAULT_MTU

- Class: Run time
- Default: MTU1024

- Applicable device(s): Gen2, uDAPL, VAPI

The internal MTU used for IB. This parameter should be a string instead of an integer. Valid values are: MTU256, MTU512, MTU1024, MTU2048, MTU4096.

10.12 MV2_ENABLE_AFFINITY

- Class: Run time
- Default: 1
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

Enable CPU affinity by setting MV2_ENABLE_AFFINITY=1 or disable it by setting MV2_ENABLE_AFFINITY=0. MV2_ENABLE_AFFINITY does not take effect when _AFFINITY_ is not defined.

10.13 MV2_GET_FALLBACK_THRESHOLD

- Class: Run time
- This threshold value needs to be set in bytes.
- This option is effective if we define ONE_SIDED flag.
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This defines the threshold beyond which the MPI_Get implementation is based on direct one sided RDMA operations.

10.14 MV2_IBA_EAGER_THRESHOLD

- Class: Run time
- Default: Architecture dependent (12KB for IA-32)
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This specifies the switch point between eager and rendezvous protocol in MVAPICH.

10.15 MV2_INITIAL_PREPOST_DEPTH

- Class: Run time
- Default: 10
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This defines the initial number of pre-posted receive buffers for each connection. If communication happens for a particular connection, the number of buffers will be increased to RDMA_PREPOST_DEPTH.

10.16 MV2_MPD_RECVTIMEOUT_MULTIPLIER

- Class: Run time
- Default: 0.05

The multiplier to be added to the MPD mpiexec timeout for each process in a job.

10.17 MV2_NDREG_ENTRIES

- Class: Run time
- Default: 1000
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This defines the total number of buffers that can be stored in the registration cache. It has no effect if MV2_USE_LAZY_MEM_UNREGISTER is not set. A larger value will lead to less frequent lazy de-registration.

10.18 MV2_NUM_HCAS

- Class: Run time
- Default: 1
- Applicable device(s): Gen2, Gen2-iWARP

This parameter indicates number of InfiniBand adapters to be used for communication on an end node.

10.19 MV2_NUM_PORTS

- Class: Run time
- Default: 1
- Applicable device(s): Gen2, Gen2-iWARP

This parameter indicates number of ports per InfiniBand adapter to be used for communication per adapter on an end node.

10.20 MV2_NUM_QP_PER_PORT

- Class: Run time
- Default: 1
- Applicable device(s): Gen2, Gen2-iWARP

This parameter indicates number of queue pairs per port to be used for communication on an end node. This is useful in the presence of multiple send/recv engines available per port for data transfer.

10.21 MV2_NUM_RDMA_BUFFER

- Class: Run time
- Default: Architecture dependent (32 for EM64T)
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

The number of RDMA buffers used for the RDMA fast path. This *fast path* is used to reduce latency and overhead of small data and control messages. This value will be ineffective if MV2_USE_RDMA_FAST_PATH is not set.

10.22 MV2_ON_DEMAND_THRESHOLD

- Class: Run Time
- Default: 32
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL

This defines threshold for enabling on-demand connection management scheme. When the size of the job is larger than the threshold value, on-demand connection management will be used.

10.23 MV2_PREPOST_DEPTH

- Class: Run time
- Default: 64
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This defines the number of buffers pre-posted for each connection to handle send/receive operations.

10.24 MV2_PUT_FALLBACK_THRESHOLD

- Class: Run time
- This threshold value needs to be set in bytes.
- This option is effective if we define ONE_SIDED flag.
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This defines the threshold beyond which the MPI_Put implementation is based on direct one sided RDMA operations.

10.25 MV2_RDMA_CM_ARP_TIMEOUT

- Class: Run Time
- Default: 2000 ms
- Applicable device(s): Gen2, Gen2-iWARP

This parameter specifies the arp timeout to be used by RDMA CM module.

10.26 MV2_RNDV_PROTOCOL

- Class: Run time
- Default: RPUT
- Applicable device(s): Gen2, Gen2-iWARP

The value of this variable can be set to choose different Rendezvous protocols. RPUT (default RDMA-Write) RGET (RDMA Read based), R3 (send/recv based).

10.27 MV2_R3_THRESHOLD

- Class: Run time
- Default: 4096
- Applicable device(s): Gen2, Gen2-iWARP

The value of this variable controls what message sizes go over the R3 rendezvous protocol. Messages above this message size use MV2_RNDV_PROTOCOL.

10.28 MV2_R3_NOCACHE_THRESHOLD

- Class: Run time
- Default: 32768
- Applicable device(s): Gen2, Gen2-iWARP

The value of this variable controls what message sizes go over the R3 rendezvous protocol when the registration cache is disabled (MV2_USE_LAZY_MEM_UNREGISTER=0). Messages above this message size use MV2_RNDV_PROTOCOL.

10.29 MV2_SHMEM_COLL_MAX_MSG_SIZE

- Class: Run Time
- Applicable device(s): Gen2, Gen2-iWARP

This parameter can be used to select the max buffer size of message for shared memory collectives.

10.30 MV2_SHMEM_COLL_NUM_COMM

- Class: Run Time
- Applicable device(s): Gen2, Gen2-iWARP

This parameter can be used to select the number of communicators using shared memory collectives.

10.31 MV2_SRQ_LIMIT

- Class: Run Time
- Default: 30

- Applicable device(s): Gen2, Gen2-iWARP

This is the low watermark limit for the Shared Receive Queue. If the number of available work entries on the SRQ drops below this limit, the flow control will be activated.

10.32 MV2_SRQ_SIZE

- Class: Run Time
- Default: 512
- Applicable device(s): Gen2, Gen2-iWARP

This is the maximum number of work requests allowed on the Shared Receive Queue.

10.33 MV2_USE_APM

- Class: Run Time
- Applicable device(s): Gen2

This parameter is used for recovery from network faults using Automatic Path Migration. This functionality is beneficial in the presence of multiple paths in the network, which can be enabled by using lmc mechanism.

10.34 MV2_USE_APM_TEST

- Class: Run Time
- Applicable device(s): Gen2

This parameter is used for testing the Automatic Path Migration functionality. It periodically moves the alternate path as the primary path of communication and re-loads another alternate path.

10.35 MV2_USE_BLOCKING

- Class: Run time
- Default: 0
- Applicable device(s): Gen2

Setting this parameter enables mvapich2 to use blocking mode progress. MPI applications do not take up any CPU when they are waiting for incoming messages.

10.36 MV2_USE_COALESCE

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP

Setting this parameter enables message coalescing to increase small message throughput

10.37 MV2_USE_HSAM

- Class: Run Time
- Applicable device(s): Gen2

This parameter is used for utilizing hot-spot avoidance with InfiniBand clusters. To leverage this functionality, the subnet should be configured with lmc greater than zero. Please refer to section 6.6 for detailed information.

10.38 MV2_USE_IWARP_MODE

- Class: Run Time
- Default: unset
- Applicable device(s): Gen2, Gen2-iWARP

This parameter enables the library to run in iWARP mode. The library has to be built using the flag -DRDMA_CM for using this feature.

10.39 MV2_USE_LAZY_MEM_UNREGISTER

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL

Setting this parameter enables mvapich2 to use memory registration cache.

10.40 LAZY_MEM_UNREGISTER

- Class: Compile time
- Default: not set
- Applicable device(s): VAPI

Setting this parameter enables mvapich2 to use memory registration cache. We recommend not to set this variable for the vapi device.

10.41 MV2_USE_RDMA_CM

- Class: Run Time
- Default: Network Dependant (set for Gen2-iWARP and unset for Gen2)
- Applicable device(s): Gen2, Gen2-iWARP

This parameter enables the use of RDMA CM for establishing the connections. The library has to be built using the flag -DRDMA_CM for using this feature.

10.42 MV2_USE_RDMA_FAST_PATH

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL

Setting this parameter enables mvapich2 to use adaptive rdma fast path features for Gen2 devices and static rdma fast path features for uDAPL device.

10.43 MV2_USE_RDMA_ONE_SIDED

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL

Setting this parameter allows mvapich2 to use optimized one sided implementation based RDMA operations.

10.44 MV2_USE_RING_STARTUP

- Class: Run time
- Default: set
- Applicable device(s): Gen2

Setting this parameter enables mvapich2 to use ring based startup.

10.45 MV2_USE_SHARED_MEM

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

Use shared memory for intra-node communication.

10.46 MV2_USE_SHMEM_ALLREDUCE

- Class: Run Time
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This parameter can be used to turn off shared memory based MPI_Allreduce for Gen2 over IBA by setting this to 0.

10.47 MV2_USE_SHMEM_BARRIER

- Class: Run Time
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This parameter can be used to turn off shared memory based MPI_Barrier for Gen2 over IBA by setting this to 0.

10.48 MV2_USE_SHMEM_COLL

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

Use shared memory for collective communication. Set this to 0 for disabling shared memory collectives.

10.49 MV2_USE_SHMEM_REDUCE

- Class: Run Time
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This parameter can be used to turn off shared memory based MPI_Reduce for Gen2 over IBA by setting this to 0.

10.50 MV2_USE_SRQ

- Class: Run time
- Default: set
- Applicable device(s): Gen2, Gen2-iWARP

Setting this parameter enables mvapich2 to use shared receive queue.

10.51 MV2_VBUF_POOL_SIZE

- Class: Run time
- Default: 512
- Applicable device(s): Gen2, Gen2-iWARP

The number of vbufs in the initial pool. This pool is shared among all the connections.

10.52 MV2_VBUF_SECONDARY_POOL_SIZE

- Class: Run time
- Default: 128
- Applicable device(s): Gen2, Gen2-iWARP

The number of vbufs allocated each time when the global pool is running out in the initial pool. This is also shared among all the connections.

10.53 MV2_VBUF_TOTAL_SIZE

- Class: Run time
- Default: Architecture dependent (6 KB for EM64T)
- Applicable device(s): Gen2, Gen2-iWARP

The size of each vbuf, the basic communication buffer of MVAPICH2.

10.54 SMP_EAGERSIZE

- Class: Run time
- Default: Architecture dependent
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This has no effect if macro `_SMP_` is not defined. It defines the switch point from Eager protocol to Rendezvous protocol for intra-node communication. If macro `_SMP_RNDV_` is defined, then for messages larger than `SMP_EAGERSIZE`, SMP Rendezvous protocol is used. Note that this variable should be set in KBytes.

10.55 SMPI_LENGTH_QUEUE

- Class: Run time
- Default: Architecture dependent
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This has no effect if macro `_SMP_` is not defined. It defines the size of shared buffer between every two processes on the same node for transferring messages smaller than or equal to `SMP_EAGERSIZE`. Note that this variable should be set in KBytes.

10.56 SMP_NUM_SEND_BUFFER

- Class: Run time
- Default: Architecture dependent
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This has no effect if macro `_SMP_` is not defined. It defines the number of internal send buffers for sending intra-node messages larger than `SMP_EAGERSIZE`.

10.57 SMP_SEND_BUF_SIZE

- Class: Compile time
- Default: Architecture dependent
- Applicable device(s): Gen2, Gen2-iWARP, uDAPL, VAPI

This has no effect if macro `_SMP_` is not defined. It defines the packet size when sending intra-node messages larger than `SMP_EAGERSIZE`.