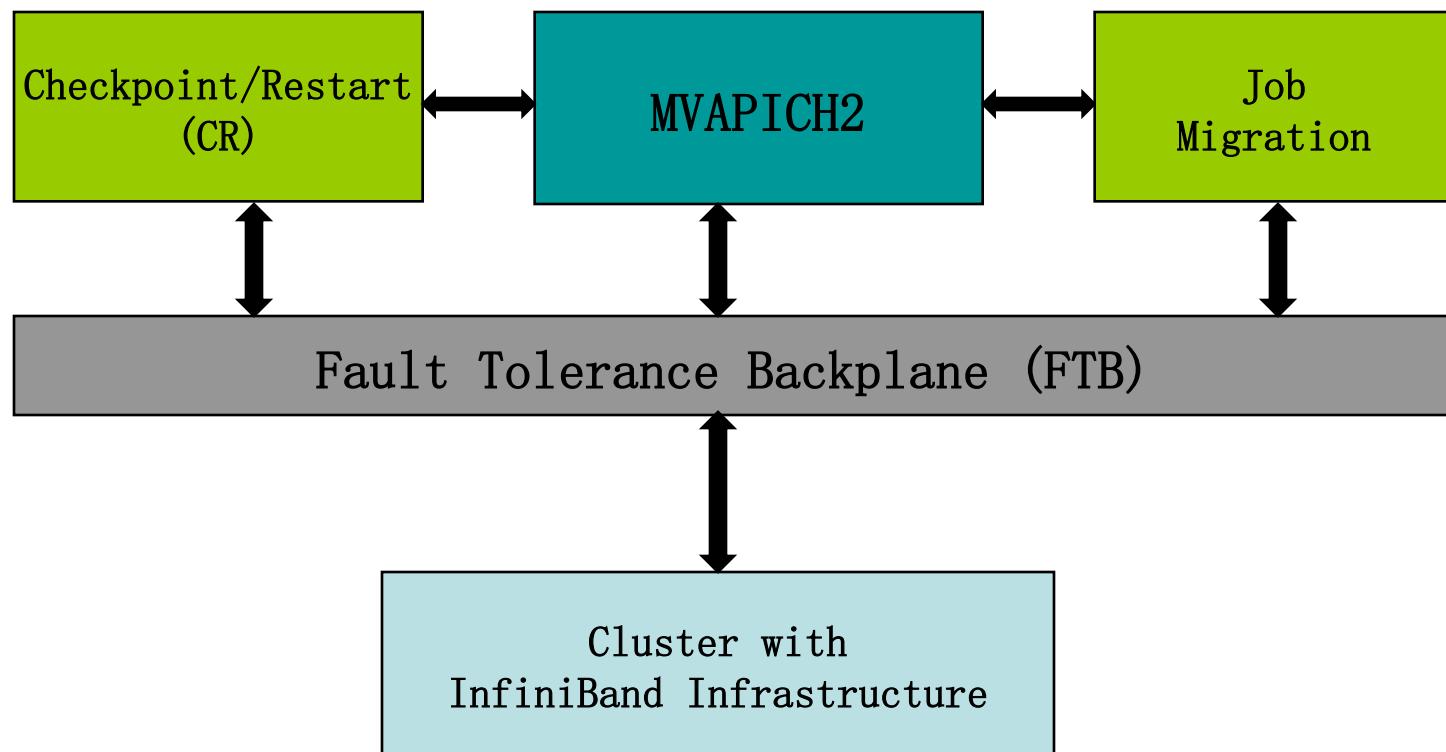


# FTB-Enabled Fault Tolerance in MVAPICH2

The Ohio State University



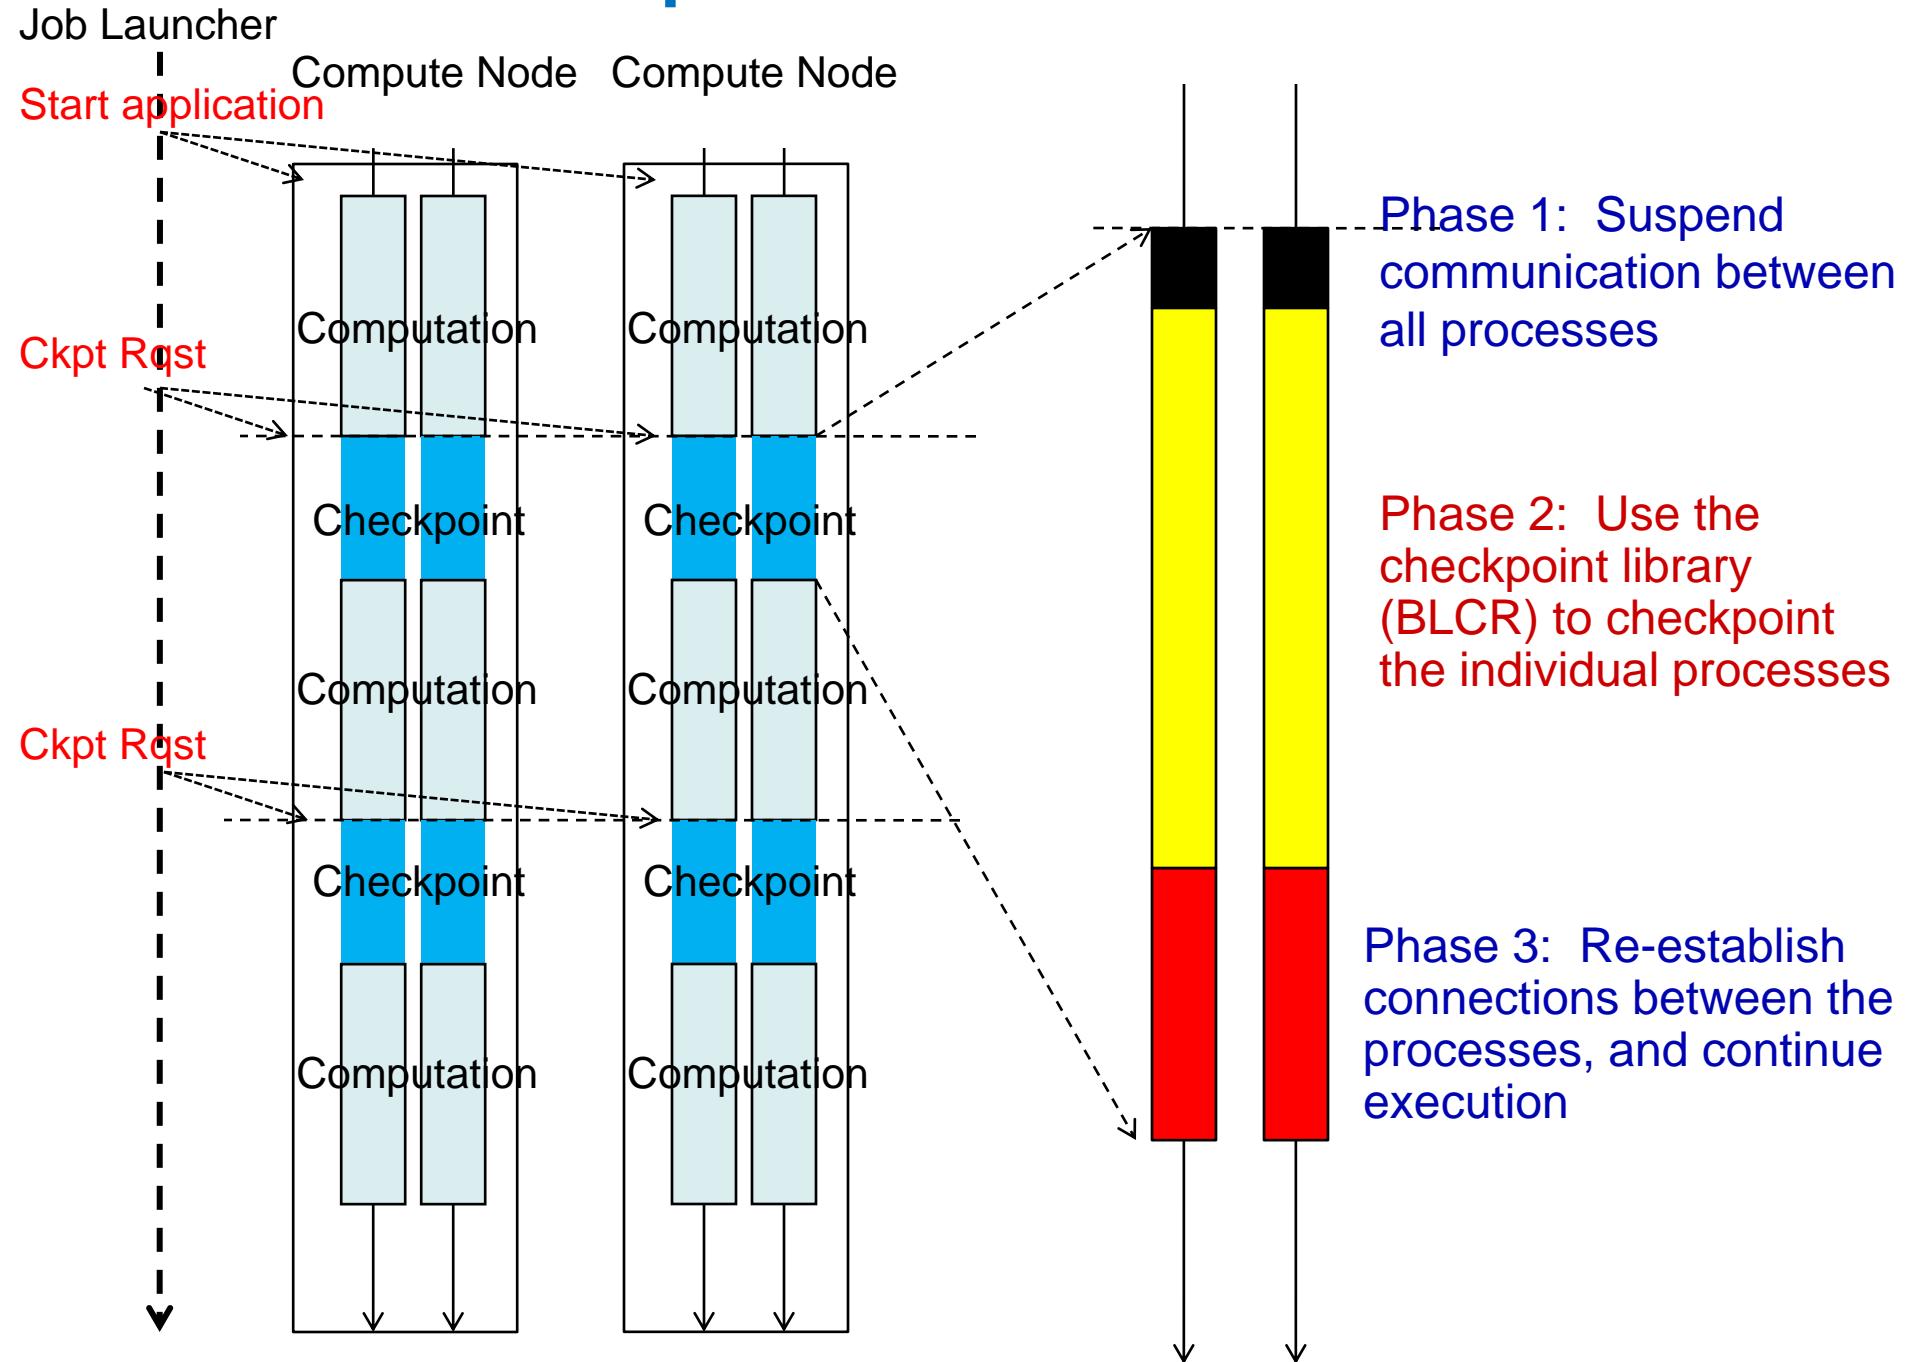
# Fault Tolerance with MVAPICH2



# FTB Workflows in MVAPICH2

- Checkpoint-Restart Mechanism
- Fast Check pointing by Write Aggregation with Dynamic Buffer Interleaving
- RDMA-Based Job Migration Framework for MPI over InfiniBand

# Checkpoint-Restart Workflow



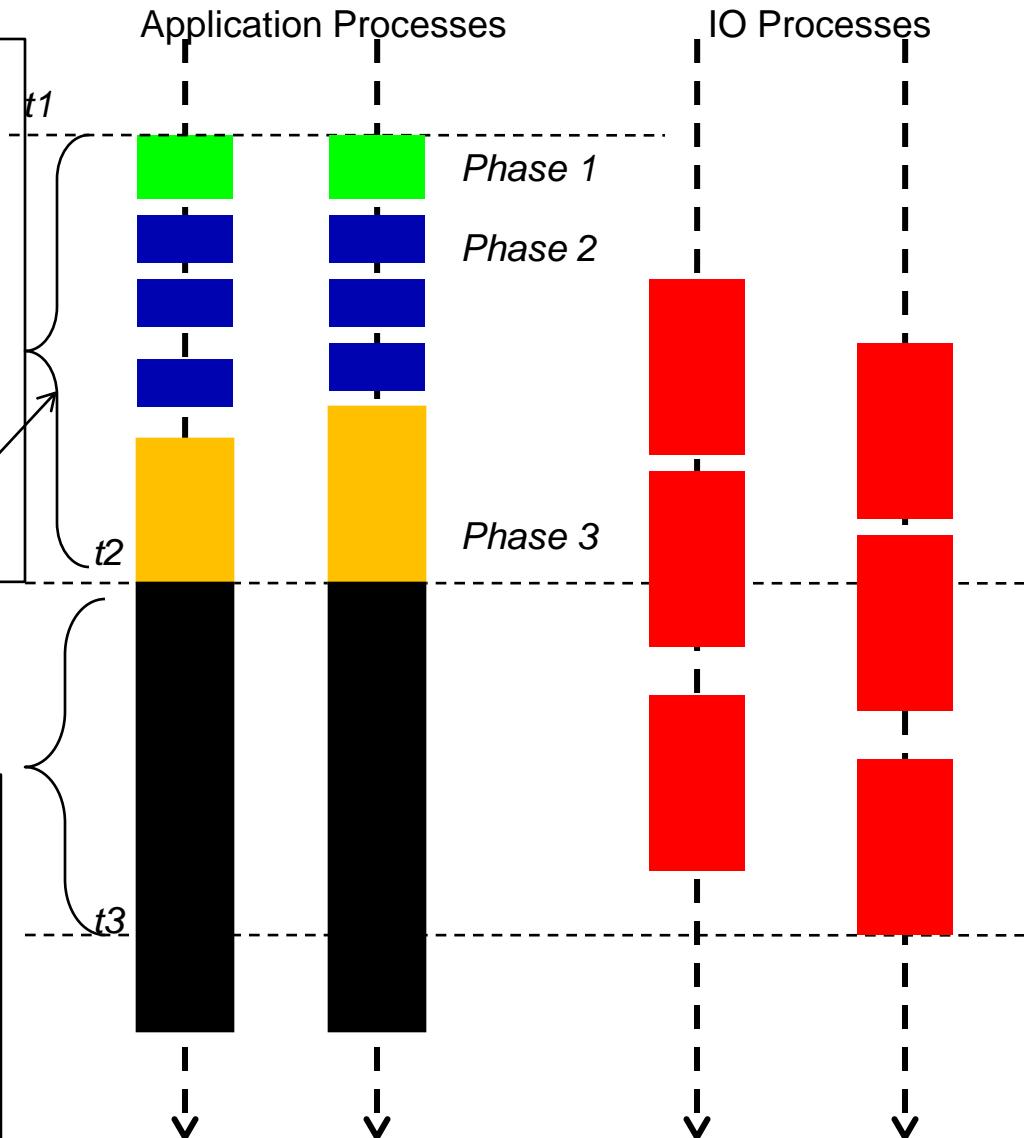
# Write Aggregation Workflow

- ✓ Aggregate many VFS writes into larger chunk writes
  - Reduce disk seeks, improve bandwidth utilization
- ✓ Overlaps application progress with slow checkpoint file writing
  - Improve application execution time

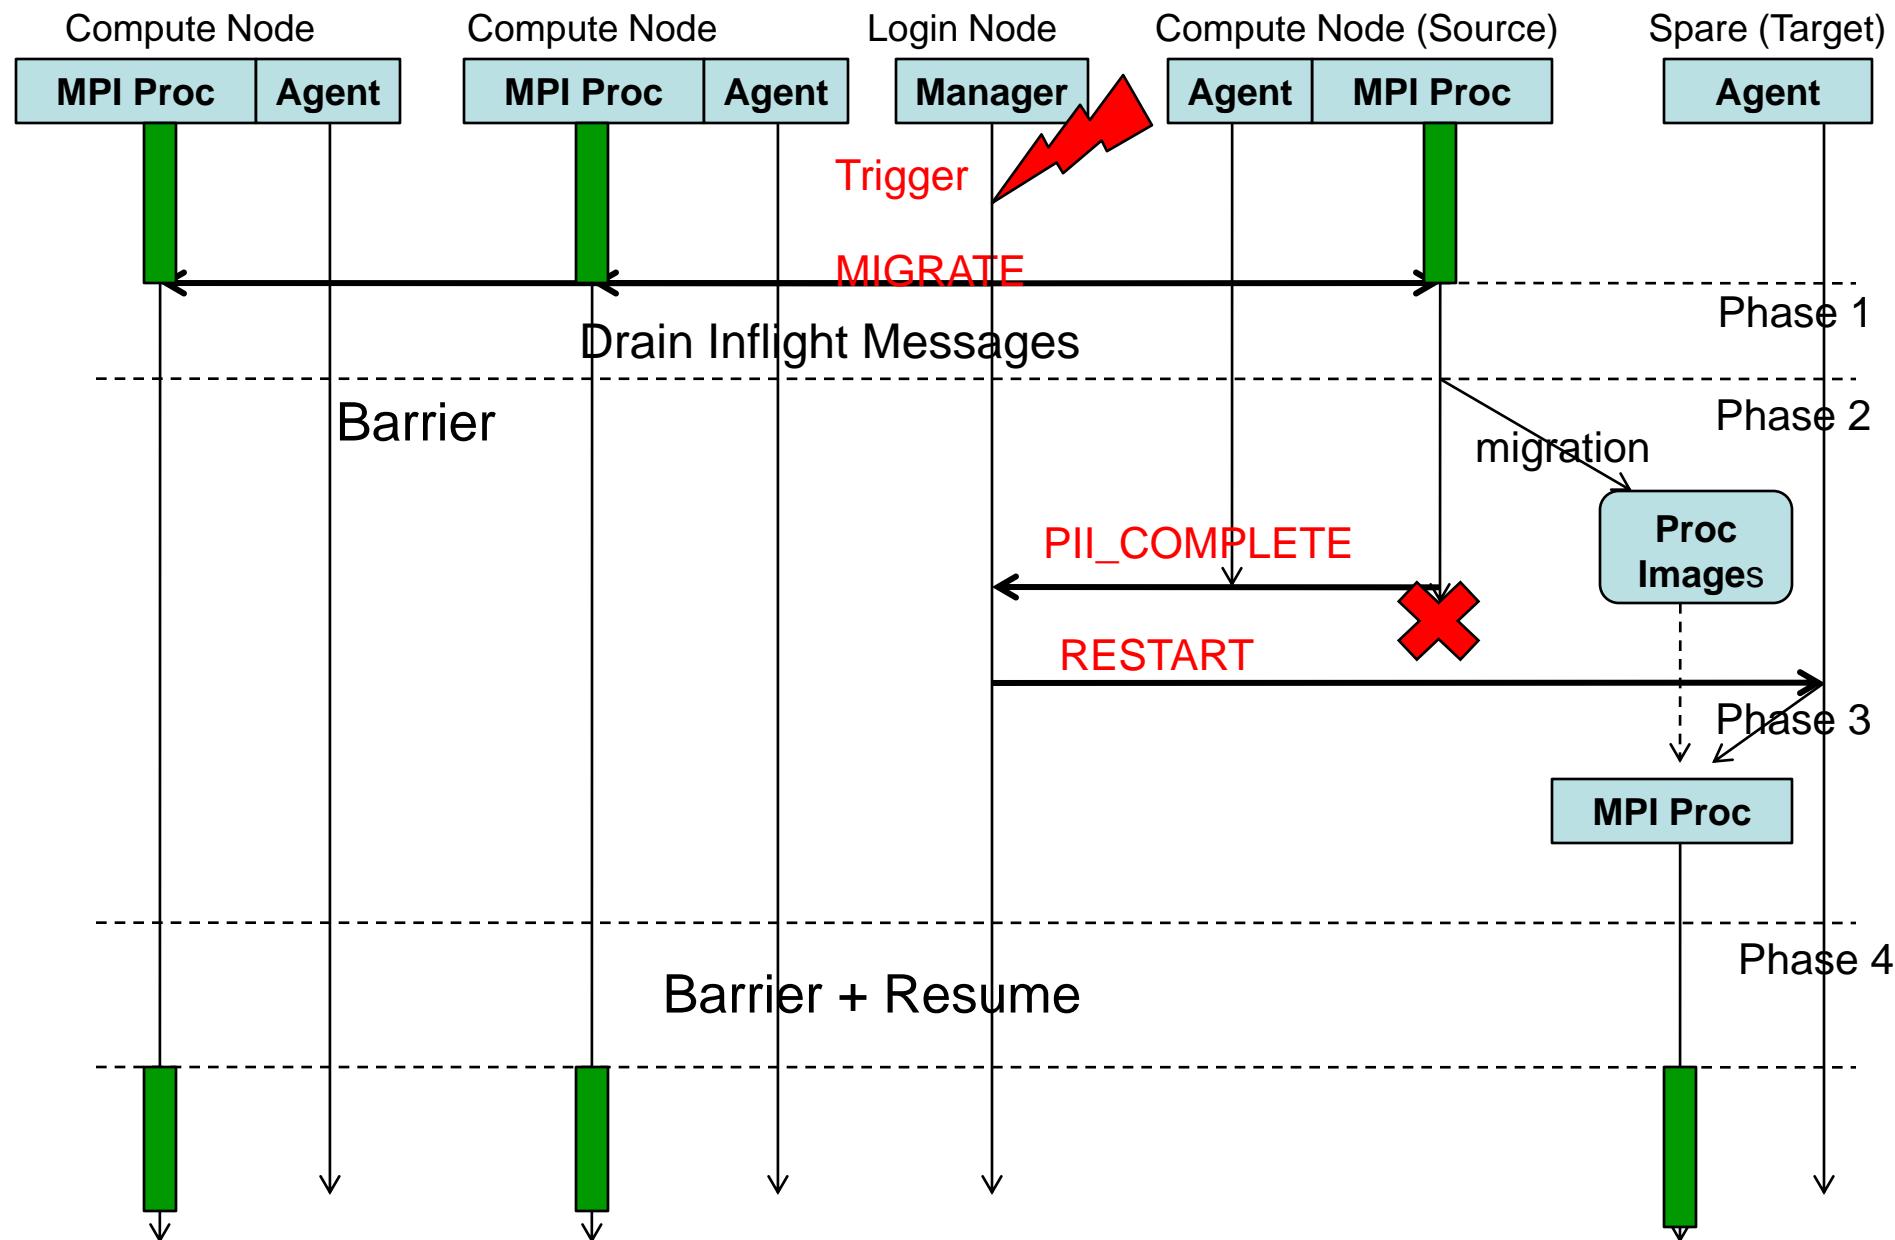
• Application-perceived checkpoint time =  $t_2 - t_1$

•  $(t_3 - t_2)$  is overlapping between Application-Progress and File-Write

• Vulnerable to failures in this window  
• Provide an interface to poll for File-Write completion



# Job Migration Workflow



# FTB Events in MVAPICH2 CR Framework

- CR\_FTB\_CHECKPOINT
  - Req. a Checkpoint of the MPI Job
  - Sent by mpirun\_rsh
- CR\_FTB\_CKPT\_DONE
  - Checkpoint completed successfully
  - Sent by MPI processes that were able to checkpoint
- CR\_FTB\_CKPT\_FAIL
  - Checkpoint failed
  - Sent by MPI Processes that failed to take the checkpoint
- CR\_FTB\_RSRT\_DONE
  - Restart completed successfully
  - Sent by MPI Processes that were able to restart
- CR\_FTB\_RSRT\_FAIL
  - Restart failed
  - Sent by MPI Processes that failed to restart
- CR\_FTB\_CKPT\_FINALIZE
  - CR module has shutdown
  - Sent by all MPI Processes doing MPI\_Finalize()
- CR\_FTB\_APP\_CKPT\_REQ
  - Sent by MPI Process from where the user requested a checkpoint through MVAPICH2\_Sync\_Checkpoint command

# FTB Events in MVAPICH2 Migration Framework

- PREDICTOR\_NODES\_FAILURE
  - An FTB-aware failure predictor informs the Migration Manager about a possible failure
  - Thrown by: Failure Predictor
- REQ\_MIGRATE
  - The Migration manager arbitrates a spare node from a given set and provides it to mpispawn
  - Thrown by: Migration Manager
- MIGRATE\_DONE
  - The mpispawn informs the mig. manager about successful migration
  - Thrown by: mpispawn on target node
- MIG\_NODE\_MIGRATED
  - Migration successfully completed
  - Thrown by: Migration Manager
- MIG\_NODE\_MIGRATE\_FAILED
  - Migration failed
  - Thrown by: Migration Manager

# Availability and Future Plans

- **Availability**
  - FTB-CR was incorporated in MVAPICH2 1.4
  - FTB-enabled CR with write aggregation and Job Migration has been incorporated in MVAPICH2 1.6
  - It is available at <http://mvapich.cse.ohio-state.edu>
  - Detailed guidelines for using FTB support in MVAPICH2 is available at [https://mvapich.cse.ohio-state.edu/support/user\\_guide\\_mvapich2-1.5.1.html#x1-400006.6](https://mvapich.cse.ohio-state.edu/support/user_guide_mvapich2-1.5.1.html#x1-400006.6)
- **Future Plans**
  - Optimize the design by analyzing the restart overhead at the target node