

# Towards Building a Highly-Available Cluster Based Model for High Performance Computing

Azzedine Boukerche† Raed Al-Shaikh† and Mirela Sechi Moretti Notare††

PARADISE Research Laboratory

†SITE, University of Ottawa

††Barddal University, Brasil

Emails: {boukerch, rshaikh}@site.uottawa.ca, mnotare@ieee.org

## Abstract

In recent years, we have witnessed a growing interest in high performance computing (HPC) using a cluster of workstations. However, many challenges remain to be resolved before these systems become dependable. One of the challenges in a clustered environment is to keep system failure to the minimum level and while achieving the highest possible level of system availability. High-Availability (HA) computing attempts to avoid the problems of unexpected failures through active redundancy and preemptive measures. In this paper, we propose to build HA-clusters based model for high performance computing. Our model is based on combination of both HPC and HA concepts, we also propose to investigate further the *hardware* and the *management* layers of the HA-HPC cluster design, and the *parallel-applications* layer (i.e. FT-MPI implementations). In this work, we focus upon the latter layer. We discuss our model, and present our simulation experiments we have carried out to evaluate our proposed model.

## 1. Introduction

Commodity-hardware clusters offer several advantages over the traditional supercomputers. First, High Performance Clusters are intended to be a cheaper replacement for the more complex/expensive supercomputers to run traditional technical applications such as simulations, biotechnology, financial market modeling, data mining and stream processing [1]. Second, cluster computing can scale to very large systems. Hundreds or even thousands of machines can be networked to suit the application needs. In fact, the entire Internet can be viewed as one truly huge cluster [3]. The third advantage is availability; replacing a "faulty node" within a cluster is trivial compared to fixing a faulty SMP component, resulting in a lower mean-time-to-repair (MTTR) for carefully designed cluster configuration [18].

On the other hand, HPCs are not problems-free. First, clusters have higher network latency with a lower network bandwidth compared to SMP and supercomputers. However, as we will see in section 3, this network difference is becoming insignificant, thanks to the advancements in the network interconnects. The second potential problem is the frequency of hardware failures (Mean-time-to-failure, or MTTF). Because of many heterogeneous commodity hardware involved to build an HPC

cluster, the probability of a hardware failure is higher than an SMP machine. Therefore, most clusters are unable to handle runtime system configuration changes caused by transient failures and require a complete restart of the entire machine [18].

## 2. Motivation

One of the challenges in an HPC clustered environment is to keep system failure to the minimum and to provide the highest possible level of system availability. Due to the fact that very large and complex applications are being run on increasingly larger scale distributed computing environments, High-Availability (HA) computing has become critically important to the fundamental concept of High Performance Computing (HPC). This is because commodity hardware is employed to construct these clusters, and to a certain extent, the application code's runtime exceeds the hardware's mean-time-between-failures (MTBF) rate for the entire cluster [15]. Thus, in order to efficiently run these very large and complex applications, HA computing techniques must be employed in the HPC environment.

In this paper, we combine both HPC and HA models and layout the design of a HA-HPC cluster, considering all possible measures. In particular, we explore the *hardware* and the *management* layers of

the HA-HPC cluster design, as well as a more focused study on the parallel-applications layer (i.e. FT-MPI implementations).

The rest of the paper is organized as follows: In the next section, we describe the general architecture of the Beowulf cluster, which is becoming the standard design of any HPC cluster. In section 4, we demonstrate the modifications needed in the hardware level, in order to make a large-scale Beowulf architecture fault tolerant. In section 5, we discuss the fault tolerant issues in the management level by demonstrating the High Availability Open Source Cluster Application Resource (HA-OSCAR) that aims toward non-stop services in the HPC environment. Next and in Section 6 we explore and compare, in more details, the fault tolerant techniques in the application level, mainly the MPI layer. In section 7, we report our performance and test results. The last section states our conclusions and summary.

### 3. The Beowulf Cluster Architecture

Originally developed at NASA, Beowulf clusters are developed worldwide to support scientific computing [22]. Fundamentally, it is a design for HPC clusters on inexpensive personal computer hardware. A typical “large-scale” Beowulf cluster consists of the following major system components (figure 1):

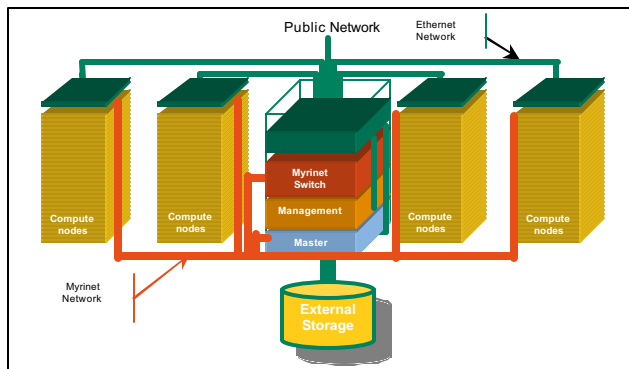


Figure-1: A general “large-scale” Beowulf cluster layout

- 1) A master (or control) server: a master server is responsible for serving user requests and distributing them to clients via scheduling/queuing software, such as Portable Batch System (PBS) or LSF [2]. To control access to the cluster, users are only permitted to access this node and are blocked from accessing the other nodes. In a simple Beowulf cluster, the master node is considered to be a single point of failure. In particular, a master node failure can render

the entire cluster unusable. Therefore, redundancy should exist in order to provide HA to the master node, as we will see in the next section.

- 2) Multiple identical client (compute) nodes: these clients or compute nodes are normally dedicated to computation. Normally, users are blocked from direct access to these compute nodes. An HPC cluster may be as simple as two networked identical compute nodes or as complex as thousands of nodes connected together via high speed network.
- 3) The management node: this node is used for administrative purposes, such as installing, configuring and administering all other client nodes.
- 4) Network Interconnect: currently, there are several network interconnects that provide low latency (less than 5 Micro seconds) and high bandwidth (multiple Gbps). The suitable HPC interconnect is determined by the application that is intended to run on the cluster. Two of the leading products are Myrinet [7] and Quadrics. More recently, InfiniBand [7] has entered the high performance computing market. Unlike Myrinet and Quadrics, InfiniBand was initially introduced as a generic interconnect for inter-process communication and I/O. Nevertheless, its high performance and scalability make it also attractive as a communication layer for high performance computing. Tests [12,19] show that for small MPI messages, InfiniBand outperform other interconnects due to the higher bandwidth, while Quadrics fits small MPI messages the best due to the very low latency. Table 1 shows each interconnect option with its peak latency and throughput measures.

|                   | <b>Gbit Ethernet</b> | <b>Myrinet D-Cards</b> | <b>Quadrics QsNet</b> | <b>InfiniBand</b> |
|-------------------|----------------------|------------------------|-----------------------|-------------------|
| <b>Throughput</b> | 120MB/s              | 473MB/s                | 308 MB/s              | 841MB/s**         |
| <b>Latency*</b>   | 12.1                 | 6.8                    | 4.6                   | 6.7               |
| <b>Price</b>      | Cheap                | Moderate               | Moderate-to-expensive | Expensive         |

\* Measures are in Microseconds.

\*\* limited by the PCI-X bus

Table-1: Network interconnects comparison

In the next section, we start our proposed design by introducing redundancy and failover components in the lowest level of the HA-HPC clusters.

#### 4. FT-HPC in the Hardware Layer

As we have seen from the standard Beowulf design, there are certain components where they are considered to be a single point of failure, such as the master node and the network switches. In earlier times, *cost* was a valid reason for systems designers to emphasize on the application needs and minimize the cost of building an HPC cluster by removing all the costly redundant hardware components. However, as the cost of commodity hardware is declining, combining HPC and HA architectures is becoming feasible to achieve HA cluster that is used for High Performance Computing. For instance, various techniques are currently available to have redundant master nodes [18]. These implementations include active/active, active/hot-standby, and active/cold-standby nodes. In the active/active, both master nodes simultaneously provide services to external requests and once one node is down, the other will take over total control. Whereas, a hot standby head node monitors system health and only takes over control when there is an outage at the primary node. The cold-standby architecture is very similar to the hot-standby, except that the backup head is activated from a cold start.

Figure 2 shows a completely hardware-modified FT-Beowulf architecture. Each master node has two network interfaces connected to two separate virtual networks (VLANs); one is connected to the Internet by a public VLAN, and the other NIC is connected to the private LAN. In this setup, the standby server monitors the primary server and waits for taking over when a failure in the primary server is detected. Note that the Local LAN Switches and the management nodes are completely redundant and the cluster may be operational even if one switch fails.

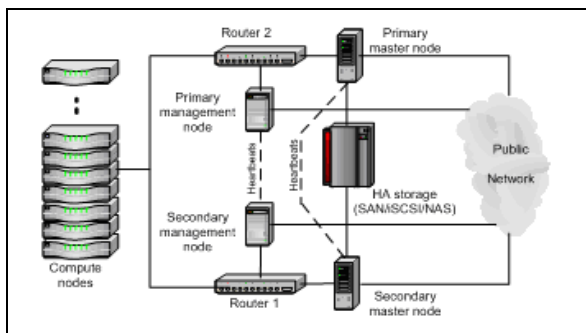


Figure2: A general fault-tolerant HPC cluster architecture

#### 5. FT-HPC in the Management Layer

Systems that have the ability to hot-swap hardware components need a management software layer that understands the concept of dynamic system configuration, in order to keep the system alive. One of these management applications is the High Availability Open Source Cluster Application Resource (HA-OSCAR). HA-OSCAR is an Open Source project that aims toward non-stop services in the HPC environment through a combined power of High Availability and Performance Computing solutions [2].

##### 5.1 HA-OSCAR

The OSCAR project is built by a mixture of academic and research members including: Bald Guy Software (BGS), DELL, IBM, Intel, MSC.Software, Indiana University, the National Center for Supercomputing Applications (NCSA), Oak Ridge National Laboratory (ORNL), and University of Sherbrooke [4]. The HA-OSCAR project's primary goal is to enhance a Beowulf cluster system for mission-critical applications and sensitive HPC infrastructures. Its basic package includes a set of "core" toolkits needed to build and maintain a cluster [4]. Other included tools cover most, if not all, commonly used HPC applications, such as: LAM/MPI, PVM and MPICH for running parallel applications, Maui Portable Batch System (PBS) for batch job scheduling, OpenSSH for secure remote login, and System Installation Suite (SIS).

HA-OSCAR addresses service level faults via the Adaptive Self Healing technique (ASH) [4]. ASH MON daemon monitors service(s) availability at every interval (default is 5 seconds) and triggers alerts upon failure detection. When a failure is triggered, a hot-standby node takes over. This hot-standby node is a clone of the active node that contains the entire HA-OSCAR software bundle and will process user requests when the active master node fails.

To see how the master node's hot-swapping process works with HA-OSCAR, consider the following example that is illustrated in figure 3: HA-OSCAR assigns a primary server public IP address to be used for external accesses and a private IP address to connect to the compute nodes (a private VLAN as users are not allowed to access the compute nodes). For the standby server, the public IP address is initially unassigned and its private IP address is configured as shown in the figure:




Figure -3: HA-OSCAR failover process to the master node

When a primary server failure occurs, all its network interfaces will go down. HA-OSCAR will trigger the network disconnection and the standby server takes over and clones the primary server network configuration. The standby server will automatically mimic its both public and private network interfaces to be the same as the original IP addresses for the primary master node. This IP cloning process only takes 3 to 5 seconds [4].

## 6. FT-HPC in the Application Layer

Up to this point, we have considered the hardware and management elements in designing a FT-HPC cluster. Now we present, in more details, the HA issues in the application level. Although other implementation exists such as PVM (Parallel Virtual Machine), we focus our study on the Message Passing Interface (MPI) implementations, mainly because it is becoming the de facto standard communication protocol for parallel processing and HPC, especially when performance is the main concern [6].

### 6.1 Introduction to MPI

MPI is a library of routines that can be called from any programming language (mainly C or FORTRAN), on a distributed memory system. It is designed to allow a network of heterogeneous machines to be used as a single distributed parallel processor. MPI's advantage over older message passing libraries is that it is both portable and fast (because each MPI implementation is optimized for the hardware it runs on) [6].

The need for a fault tolerant MPI standard arises from the growing concern with the reliability of processors, communication and systems' structure. As current HPC systems increase in size, fault tolerance in MPI becomes a very important concern for critical high performance applications using the MPI library. W. Gropp and E. Lusk [11] define the requirements of a fault-tolerance MPI as:

- Failure can be detected
- Information needed to continue the computation is available

Moreover, they see that the highest level of MPI "survival" is that when MPI implementation automatically recovers from faults while the MPI program continues without significant change to its behavior. Their definition to the second level of survival is when the applications are notified of the problem and are prepared to handle it. In each of two cases, the application may proceed without restarting because the program has enough information about the failing process and can assign another running process to pick up the work from there. A different level of survival is that an application aborts and restarts from a checkpoint. Here, the states of all processes are saved outside the processes themselves, in a stable storage for example.

In the next section, we analyze some of the techniques that are used to build such fault tolerant MPI packages, namely: Checkpointing, Message Logging and the Worker/Manager algorithm. We compare these concepts and propose minor modifications to make them suit our complete HA-HPC design. Later, we review some of the existing fault tolerant MPI packages that are based on the mentioned techniques. Finally, we conclude our work by benchmarking a small-scale cluster and stating our findings and comments.

### 6.2 MPI Error Handlers

Before moving to the various fault tolerant MPI techniques, we first present the pre-defined MPI error-handlers and describe their routines.

MPI error-handlers specify the action to be taken when the MPI program runs through a failure. The specified error handling routine is used for any MPI exception that occurs during an MPI call for a communication with the communicator. The set of errors calls that are handled by MPI is implementation-dependent. Therefore, a stable MPI implementation will attempt to handle as many errors as possible. Errors that are not handled by MPI will be either handled by the error handling mechanisms of the language run-time or the operating system.

The MPI error-handlers do not necessarily allow the user to continue to use MPI. Instead, they allow a user to issue user-defined error messages and to take actions unrelated to MPI (such as flushing I/O buffers, having a memory dump or calling other outside routines) before a program exits. However, it is not required from a non-fault tolerant MPI implementation to use error handlers.

MPI error-handlers can be either built in or user defined [17]. The built-in error-handlers are `MPI_ERRORS_ARE_FATAL` (The default error-handler) and `MPI_ERRORS_RETURN`. The first error-handler indicates that if an MPI function returns unsuccessfully then all the processes in the communicator will abort. The latter handler indicates that MPI functions will attempt to return an error code. User defined error-handlers are attached to the MPI communicators. The ability to support user-defined error-handlers is important for developers when building their own MPI interfaces.

### 6.3 Fault Tolerant MPI Techniques

In this section, we cover the commonly used techniques for implementing fault tolerant MPI. In particular, we study checkpointing (Global, independent and coordinated checkpointing), Message Logging, and Manager/Worker techniques.

#### 6.3.1 MPI Checkpointing

Checkpointing is a technique where the status of the computation is saved in different stages of running, allowing the application to be restarted from that point in the event of a failure. The primary advantage of checkpointing is that it is easy to implement. However, it is often considered expensive because the time taken to do a checkpoint can rapidly grow. Therefore, fast I/O media is usually used to connect the cluster's nodes to the shared storage (e.g. fiber links), and checkpointed data are saved in a reliable storage that should not be affected by the application failure. Normally, a parallel file system, such as GPFS or SNFS [13], is used for checkpointing, since more than one process will be writing to the same storage. Checkpointing can be classified as user-directed and system-directed. In user-directed checkpointing, the programmer manually writes out any data that will be needed to restart the application. The user has to ensure that all needed data is saved, which might not be an easy task. Moreover, the checkpoints must be taken at particular points in the program, typically when there is no interaction between processes, which again can be difficult, especially for programs that are not well-structured. Certain tools and APIs exist to assist programmers to determine where and when to checkpoint [8]. Although user-directed checkpointing seems complicated, system-directed checkpointing is much harder to implement because the processes states might be scattered throughout the cluster [2].

|                   | Locally-stored Global Checkpointing      | Independent Checkpointing           | Coordinated Checkpointing      |
|-------------------|------------------------------------------|-------------------------------------|--------------------------------|
| <b>Checkpoint</b> | All at once                              | Independent of each other           | All at once                    |
| <b>Restart</b>    | All at once                              | All at once                         | Subset of processes            |
| <b>Storage</b>    | Local to each process                    | Centralized                         | Centralized                    |
| <b>Complexity</b> | Easy to checkpoint, difficult to restart | Difficult to checkpoint and restart | Easy to checkpoint and restart |

Tabele-2 checkpointing techniques comparison

#### 6.3.2 Message Logging

As checkpointing can grow rapidly and become very expensive, message logging is developed in order to reduce the checkpointing cost, but still enable recovery. The basic idea underlying message logging is to log all the operations of message transmissions, and replay them in case of failures to reach a globally consistent state, without having to restore that state from stable storage. In initial checkpoint state is needed as a starting point, and then all messages that have been sent since are simply replayed from the log and retransmitted accordingly.

Message logging technique has the advantage of avoiding to checkpoint the whole application, and therefore save storage space. However, it is clear that message logging technique has to coexist with checkpointing. That is, a starting checkpoint should be available to implement message logging. Therefore, checkpointing performance affects the overall message logging performance as well.

#### 6.3.3 The Manager/Worker Technique

In the generic architecture, the MPI process does not connect directly to the other ones. Instead, it uses a communication daemon that handles sending, receiving, recording messages, and establishing connections with all components of the system. In the standard MPI, the failure of any one MPI process affects all processes in the communicator, even those that are not in direct communication with the failed process. In contrast, in non-MPI client-server programs, the failure of a client does not effect on the server, because each peer is aware about the status of the other peer. For example, the client can easily recognize that the server has failed and stop communicating with it. W. Gropp and E. Lusk [11] tried to mimic this scheme and apply it in the MPI context. In this algorithm, a manager process use `MPI_Comm_spawn` to create the workers and submits small tasks to them, while keeping track of them.

Then, workers return their results to the manager, simultaneously requesting a new task. This sort of communication makes it easy to recover after faults because the manager keeps a copy of the task specification and can simply re-assign it to another worker if one dies. In particular, the manager marks the failing communicator as invalid and does not use it again. When a worker dies, the `MPI_Comm_spawn` routine is used to replace the worker and continue processing with no fewer workers. According to [11], this program may not work on every implementation of MPI, because the MPI implementation must be able to return a non-success return code in the case of a communication failure such as an aborted process or failed network link.

A simple modification to the worker/manager technique in order to make it work is to add a timeout period in which the manager assumes the death of the worker process. This way, the manager process is notified about the dying process (or worker) and assigns the task again to another worker. Moreover, we may reduce the amount of messages exchanged by having the manager and all the workers write their status into a single shared storage, using a reliable parallel file system [2,3]. This way, when a worker fails and is assigned as a bad node, the manager allocates a new stand-by node, which reads the last status of the failing node and picks up the work from there.

## 6.4 Related Work on Fault Tolerant MPI

Many fault tolerant MPI implementations exist, such as LAM/MPI, Open MPI, WMPI (Windows implementation), and FT-MPI ...etc. The main difference between these implementations is the way they react to process or nodes failures in a way beyond that of the standard MPI interface. In particular, several implementations direct their fault tolerant techniques to the application level, while other techniques target their implementation to the transport and data-link levels [21]. Now we study one of these implementations in more detail:

### 6.4.1 StarFish MPI

The initial implementation of StarFish runs on Linux and supports both Myrinet and Ethernet communication links [19]. Each node in a Starfish cluster runs a daemon, and all Starfish daemons form a process group. Starfish daemons maintain some data for each application process running on the machine, as well as some shared state that defines the current cluster configuration and settings. These daemons are

responsible for interacting with clients, spawning the application processes, tracking and recovering from failures, and maintaining the system configuration [2]. Further, each application process is composed of 5 major components. These are: a group handler, which is responsible for communicating with the daemon, an application part, which includes the user supplied MPI code, a checkpoint/restart module, an MPI module, and a virtual network interface. These components communicate using an object bus based listener model [14]. To guarantee low latency and minimal impact on performance, the application part has a separate fast data path to and from the MPI module.

Starfish offers two forms of fault-tolerance for applications: The main fault-tolerant mechanism employed by Starfish is checkpoint/restart. The checkpoint/restart module of Starfish is capable of performing both coordinated and uncoordinated checkpoint, which is either system driven or application driven. Thus, when a node failure occurs, Starfish can automatically restart the application from the last checkpoint. The other form of fault tolerance offered by Starfish is more application dependent, and is suitable mostly to applications that can be trivially parallelized. For such applications, whenever a node that runs one of the application processes crashes, the event is delivered to all surviving application processes. Once the surviving members find out about the failure of a node, they repartition the data sets on which each process computes, and continue to run without interruption.

When an application is submitted to Starfish, the client determines the fault tolerant policy that should be applied to this application, i.e., should automatic restart or process notifications be used, and some rules regarding how to choose the node on which a process will be started after a partial failure.

### 6.4.3 Other Fault Tolerant Message Passing Implementations

As mentioned previously, MPI has a rich set of communication functions, which makes MPI favored over other implementations [2]. However, there are other popular parallel interfaces, such as PVM (Parallel Virtual Machine), and its various fault tolerant implementations, such as DynamicPVM and MPVM that provide the same MPI functionality. PVM is different than MPI in a way that it is built around the concept of a virtual machine, so it has the advantage when the application is going to run over a networked collection of hosts, especially if the hosts are heterogeneous. Moreover, PVM contains resource

management and process control functions that are important for creating portable applications that run on clusters of workstations. G. Geist and J. A. Kohl P. M. Papadopoulos in [22] explore more on the differences between PVM and MPI. For completeness, we view one of PVM implementations and study how it handles fault tolerant in parallel applications.

### 6.4.3.1 DynamicPVM

In general, PVM transmit communication messages using daemons, i.e. a message is first transferred to the sender's daemon, then forwarded to the daemon of the receiver and then delivered to the receiver [6].

While standard PVM offers only a static process assignment to the application, DynamicPVM provides dynamic process assignment and task scheduling, so that processes are migrated during runtime during failures. In particular, when a process failure is triggered in DynamicPVM, the local daemon on a new node prepares itself to receive the messages from the failing node, and sets its message buffer. The routing information of the local daemon on the old node gets updated so that messages which are still being sent to the old node are forwarded to the daemon on the new node. The sender daemon is informed about the new location of the process so that, in future, it sends directly to this process. One limitation in the current DynamicPVM implementation is that it is only possible to migrate one process at a time.

## 7. Experimental Results

In this section, we analyze our small-scale HA-HPC functionality and evaluate its performance. In our evaluation, we used eight-clustered nodes, each equipped with a Pentium-4 3.2GHz processor, 1.0 GHz memory and run Linux RedHat 9.0. Two of these nodes are set as the master nodes, a primary and a hot-standby, and are managed by HA-Oscar. The other six nodes are configured as the compute nodes and are interconnected by a 100MB/s network. We used DynamicPVM and FT-MPI interfaces to handle our parallel jobs, with POV-Ray package [1] on top. POV-Ray is a 3D ray-tracing engine that takes users' input and simulates the way light interacts with the defined objects to create 3D pictures and animation. POV-Ray has the ability to distribute a rendering across multiple systems, in the form of one-master and many-slave task. The master node has the responsibility of dividing the image up into small blocks, which are assigned to the slaves. When the slaves finish rendering the blocks, they are sent back to the master, which combines them to form the final image. In our


experiment, we used a 2600 bytes script file to render a 1024x768 picture with an output size of 2.3MB.

First, we ran our benchmarks with varying number of nodes, while keeping the number of processes in each node fixed, as shown in figure-4a. Clearly, the best performance/nodes was obtained while rendering on four nodes. This performance increase rate has dropped as we increased the nodes to eight. We expect that this performance gain would be flattened with 12 nodes.

Next, we ran other benchmarks while varying the number of nodes and the running processes at the same time, as in figure-4b. The best run cases were when the number of processes matched the number of nodes that are running (i.e. one process on each node). By increasing the number of processes beyond the number of nodes, the inter-process communication overhead became noticeable, causing the image to take more time to render.

Figure-4c shows how much TCP traffic is generated on the master node. Obviously, the traffic increases as we increase the number of nodes, due to the further processes' distribution among nodes. In Figure-4d, we show how image resolution affects the rendering time. While the dashed curve demonstrates the rendering time when we enforce POV-Ray to evenly distribute the processes among the nodes, the continuous curve shows the rendering time when we let the program freely distribute the processes among nodes.

Figure-4e shows the performance comparison between FT-MPI and DynamicPVM when rendering the same image on both message passing interfaces, using 200 processes in each node. FT-MPI showed a performance increase compared to DynamicPVM, especially with low number of nodes, but they performed comparatively with 8 nodes. Finally, we tested FT-MPI and DynamicPVM performance with a single node fault. Figure-4f shows the time taken to render the image while dropping one node during the process (for example, 2-1 in the figure means dropping from 2 nodes to one). Both interfaces showed approximately the same performance, even though they use different recovery schemes, as we discussed in section-2.



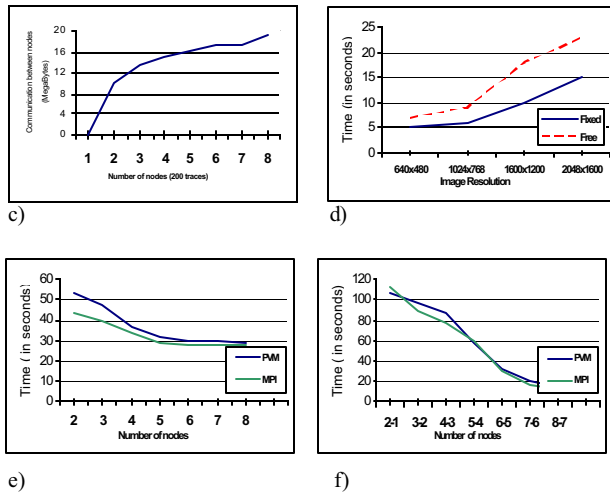


Figure-4: a) Number of nodes vs. rendering time b) number of processes vs. rendering time c) Communication traffic node on the master d) Image resolution vs. rendering time e) DynamicPVM vs. FT-MPI f) DynamicPVM vs. FT-MPI (with faults).

## 8. Conclusions

In this paper, we have presented an architecture model that can be used to build a FT-HPC cluster. These include maintaining high availability at the hardware layer, cluster manageability and fault tolerance in the application level. In particular, we explored the hardware and the management layers of the HA-HPC cluster design, as well as a more focused study on the parallel-applications layer (i.e. fault tolerant MPI implementations). We also showed that the Manager/worker algorithm that was proposed by [1] can be improved to suit most of the FT-MPI implementations. Finally, we have presented a small-scale fault tolerant HPC cluster using HA-OSCAR and different MPI implementations, to study the behavior of such a system. Our results show that combining HPC and HA architectures is feasible, in order achieve HA cluster that is used for High Performance Computing.

## REFERENCES

[1] F. Pister, L. Hess and V. Lindenstruth, Fault Tolerant Grid and Cluster Systems, Kirchhoff Institute of Physics (KIP), University Heidelberg, Germany.  
 [2] Fagg, G., Dongarra, J., Building and using a Fault Tolerant MPI implementation, In't Journal of High Performance Applications and Supercomputing, 2004.  
 [3] I. Haddad, C. Leangsuksun, R. Libby, T. Liu, Y. Liu, S. Scott, Highly Reliable Linux HPC Clusters: Self-awareness Approach, Proc. of the 2nd International

Symposium on Parallel and Distributed Processing and Applications, 2004.

[4] J. Mugler, T. Naughton, S. Scott, ... C. Leangsuksun, "OSCAR Clusters", Proceeding of The Linux Symposium 2003, July 23rd-26th, 200

[5] P. Lemarinier, A. Bouteiller, T. Herault, G. Krawezik, F. Cappello., Improved Message Logging versus Improved Coordinated Checkpointing for fault tolerant MPI, in: "Proceedings of the 6th international conference on Cluster Computing, 2004.

[6] Fagg, G. and Dongarra, J., Fault-tolerant MPI: Supporting dynamic applications in a dynamic world, 7th European PVM/MPI Users' Group Meeting No 1908 in Springer LNCS, pp. 346-353, 2000

[7] Fagg, G.E. and Dongarra, J.J. Building and using a fault tolerant MPI implementation, Int'l J. of High Perf. Computer Appl and Supercomputing, 2004

[8] Stellner, G. 1996. CoCheck: checkpointing and process migration for MPI, in Proceedings of IPPS '96. pp. 526-531.

[9] Snir, M., Otto, S.W., Huss-Lederman, S., Walker, D.W., and Dongarra, J. 1998. MPI—The Complete Reference: V. 1, The MPI Core, MIT Press.

[10] Li, K., Naughton, J.F., and Plank, J.S. 1994. Low-latency, Concurrent Checkpointing for Parallel Programs. IEEE TPDS, 5(8): 874-879

[11] W. Gropp, E. Lusk. Fault Tolerant in MPI Programs. Argonne National Laboratory.

[12] A. Gidenstam et. al., Dynamic and Fault Tolerant Cluster Management.

[13] A. Hasegawa, Hiroshi Matsouka, K. Nakanishi. Clustering Software for Linux-Based HPC.

[14] B. Polgar. Designing the Reconfiguration Strategies of Fault Tolerant Servers. Technical University of Budapest.

[15] P. Sobe. Fault Tolerant Web Services on a Computing Cluster. Institute of Computer Engineering. Medical University of Luebeck.

[16] Sultan, F. Bohra, A. Smaldone, S. Pan, Y. Gallard, P. Neamtiu, I. Iftode, L. Recovering Internet service sessions from operating system failures. Internet Computing, IEEE. 2005. ((2) . pages: 17- 27.

[17] Rob T. Aulwes, et. al., . Architecture of LA -MPI, a network-fault-tolerant MPI. Proc. 18th Int'l Parallel and Distributed Processing Symposium, pages 26-30.

[18] A. Azagury, D. Dolev, G. Gofit, John M. Marberg, J. Satran: Highly Available Cluster: A Case Study. FTCS 1994: 404-413.

[19] A. Agbaria, R. Friedman. Starfish: Fault-Tolerant Dynamic NIP1 Programs on Clusters of Workstations.

[20] S. Rao, L. Alvisi, and H.M. Vin. "Egida: An Extensible Toolkit for Low-overhead Fault-tolerance," Int'l Fault-Tolerant Computing Symp. 48-55, 1999.