How to Build a Cluster

Intel® Server Board S3000PT

Recipe ID: 24PLTF240000000100-01

Contents

Introduction

This Intel ESAA recipe describes the steps required to build a High Performance Computing (HPC) cluster based upon the Intel S3000PT server board; from hardware assembly through performance testing and cluster application development. It includes various options for hardware design, interconnects, management, BIOS and OS provisioning, job schedulers, cluster tools and cluster application development packages and runtime environments.

Note: It is recommended that the entire document is read prior to any installation activity

Overview

As mentioned above, this recipe illustrates the beginning to end setup and implementation of a HPC cluster. The basic outline of the recipe is as follows:

1) Cluster hardware setup – racks, power, cabling, switches, etc.

2) Server management – application options for pre boot and post boot system management

3) Node BIOS provisioning – a PXE environment to boot nodes to DOS for updates

4) Node OS provisioning and Cluster Management with Platform Open Cluster Stack (OCS) – An OS provisioning solution for easily deploying, running and managing cluster operating environments

5) Cluster Application Development and Application Execution Environments – Complete packages for running, developing and optimizing cluster applications, and testing cluster performance.

Each of these sections will include either references to locations of existing documentation or detailed instructions on how to setup specific installation requirements.

Hardware Components

Quantity per node	Item	Manufacturer	Model
2 (per chassis)	Intel® Server Board	Intel	S3000PT
1 Chassis Per 2 Nodes Required	Third Party Reference Chassis http://www.evercase.com	Evercase	1U
4GB Per Server Board	Memory	Any supported	Please refer to the Tested Memory list
1 Per Server Board	Intel® Xeon® Processor	Intel	Please refer to the Qualified and Supported Intel® Processors list at
1 Per Server Board	1U passive heatsink (light weight with heat pipe)	Intel	AXXUPHS
1 Per Server Board	SATA 3.5" hard drive	Any supported	Please refer to the Intel® Tested Hardware and Operating System List
As Needed	External USB 2.0 DVD drive	Any supported	Please refer to the Intel® Tested Hardware and Operating System List
1 Per Server Board	PCI-E InfiniBand* Host Channel Adapter	SilverStorm (Qlogic)	HCA 9000
1 Optional Per Server Board	InfiniBand* Host Channel Adapter	Qlogic	QLE 7140
1 Optional Per Server Board	InfiniBand* Host Channel Adapter	Intel	AXXIBDDRPT
As Needed	InfiniBand cables	Any supported	
1 (min per cluster)	InfiniBand switch	Any supported	

1 (min), As Needed	Ethernet cables	Any supported	
1 (min per cluster)	Ethernet switch	Any supported	
1 (min per cluster)	Keyboard, Video, Mouse (KVM) switch	Any supported	
1	KVM cables	Any supported	

Table 1 – Hardware Bill of Materials

Note: The AXXUPHS has been mechanically validated for use with the S3000PT board with direct chassis attach implementations

Software Used in the Installation

Dist. By	Description	File Name
Platform Computing	Open Cluster Stack (OCS) 4.1.1-1.1	
Intel	Intel® Tools 4.1.9.	

Table 2 - Software Bill of Materials

Hardware Installation

Server racks

Depending on the size of the cluster being built, the frontend, switches and any management nodes should be centrally located if multiple racks are being installed. Cluster nodes should follow a clear and organized numbering scheme to aid in node identification. Cable, power, and cooling plans are very important and must be carefully considered. These plans drive the ultimate layout of equipment in the racks. A typical configuration involves switches and servers in the same rack. An example of a large rack layout is illustrated by figure 1 below.

Note: Platform compute node numbering begins at 0. Rack identification numbers begin at 1.*

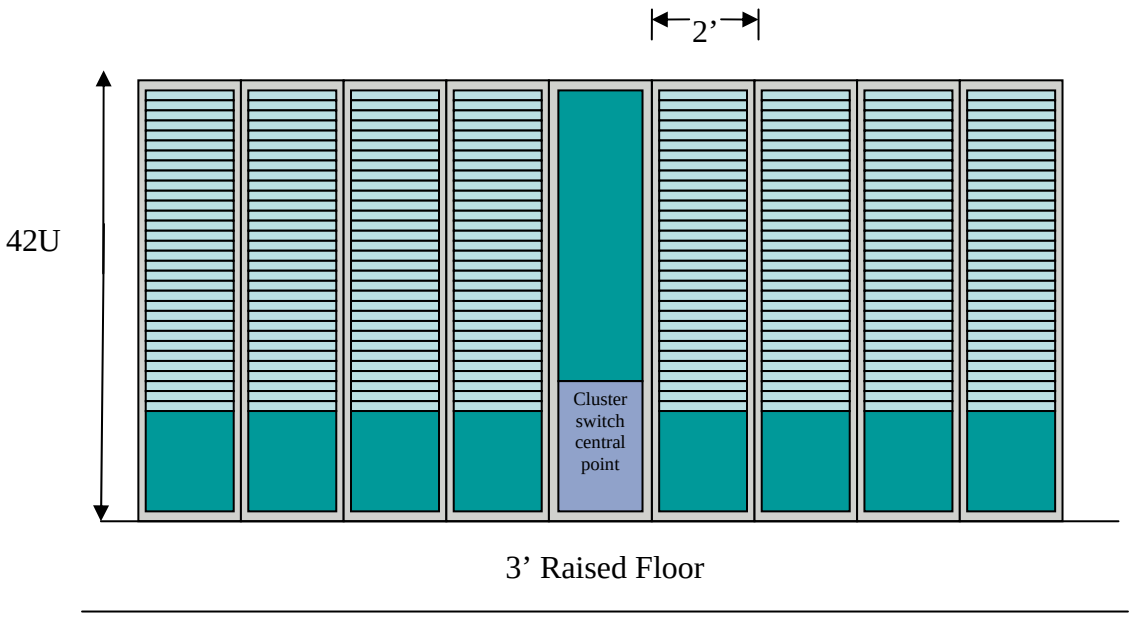

Figure 1. Server racks