

Customizable Diskless Solution for HPC Clusters

Emanuel Mariano Ravera, Facundo Ferrer

Argentina Software Development Center (ASDC)
Intel Argentina

Abstract. Diskless clusters provide a stable, easily replicable and auto-deployable solution for High Performance Computing (HPC). On the other hand, the Intel® Cluster Ready program simplifies the selection, deployment, and management of HPC clusters by establishing a standard specification for them ensuring component and application interoperability. This work proposes to use a framework for on-the-fly creation of diskless solutions that allows users with little or no knowledge to create and deploy customized HPC clusters based on the Intel® Cluster Ready Specification.

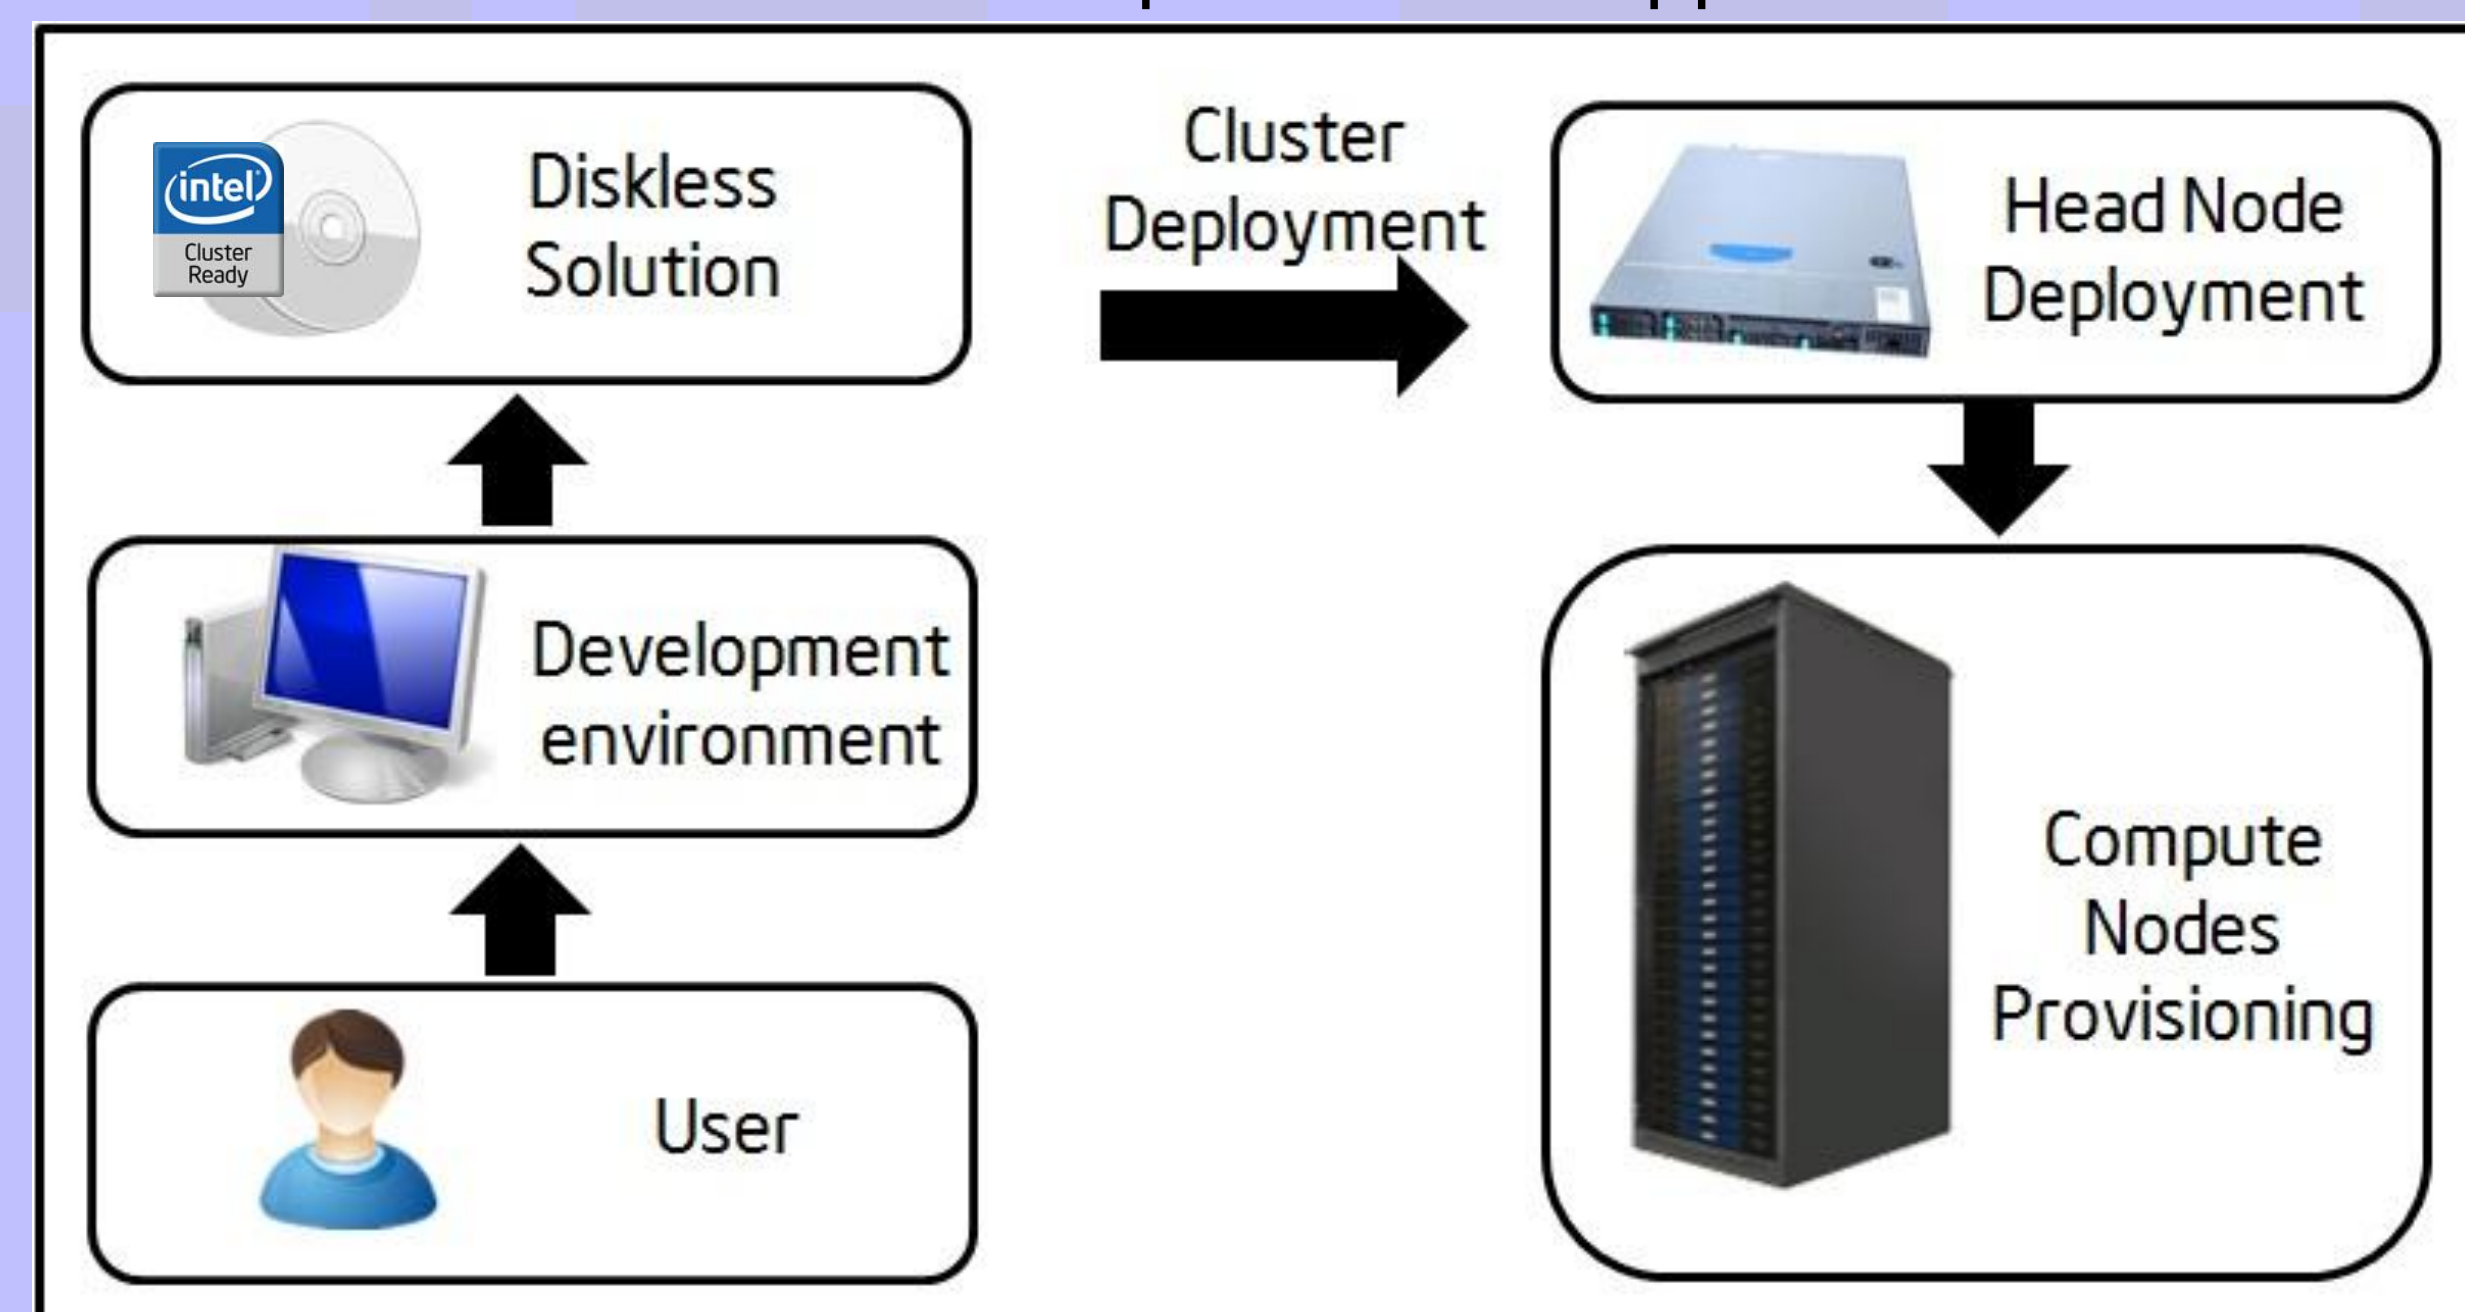
Introduction and Motivation

The installation of HPC clusters requires trained people and lots of exploration effort. Installation and maintenance of the cluster have high costs including the need of specific expertise.

Diskless clusters provides a stable, easily replicable and auto-deployable solution over a Linux architecture through a live CD system.

The Intel® Cluster Ready architecture establishes a standard specification for HPC clusters. The Intel Cluster Checker® tool, a key component of the Intel Cluster Ready program, is an automated tool that validates a cluster settings against the Intel Cluster Ready specification, and checks the general wellness of the cluster.

The proposed customization framework enables the user to engineer a personalized diskless solution from scratch based on specific requirements and taking as foundation the Intel® Cluster Ready architecture. The framework provides the capability to create different software stacks containing the Linux operating system plus customized configurations. This allows the user for instance, to deploy a cluster using a well-known and pre-loaded software stack to run specific HPC applications.

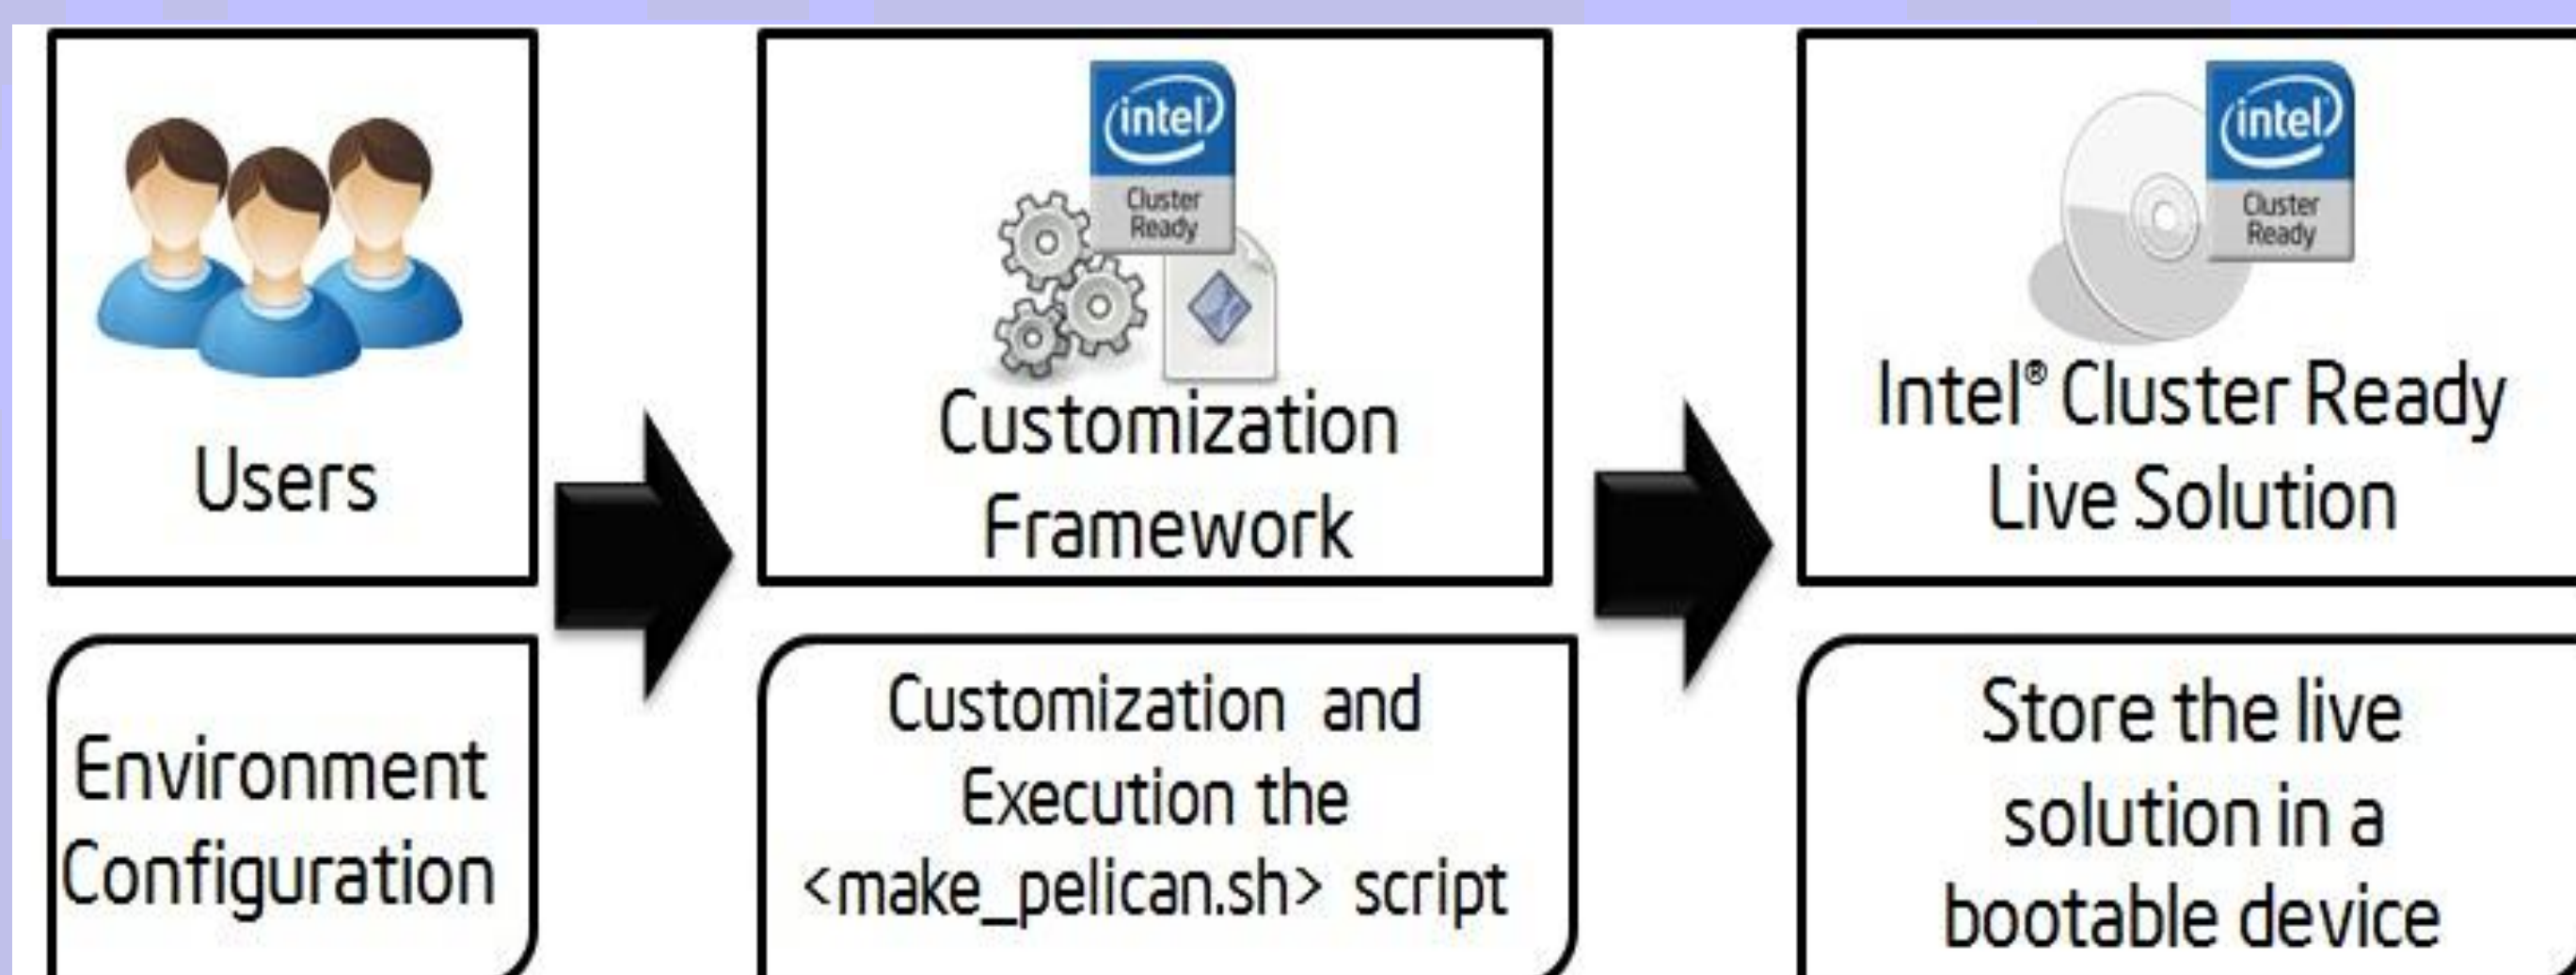


Generator Selection

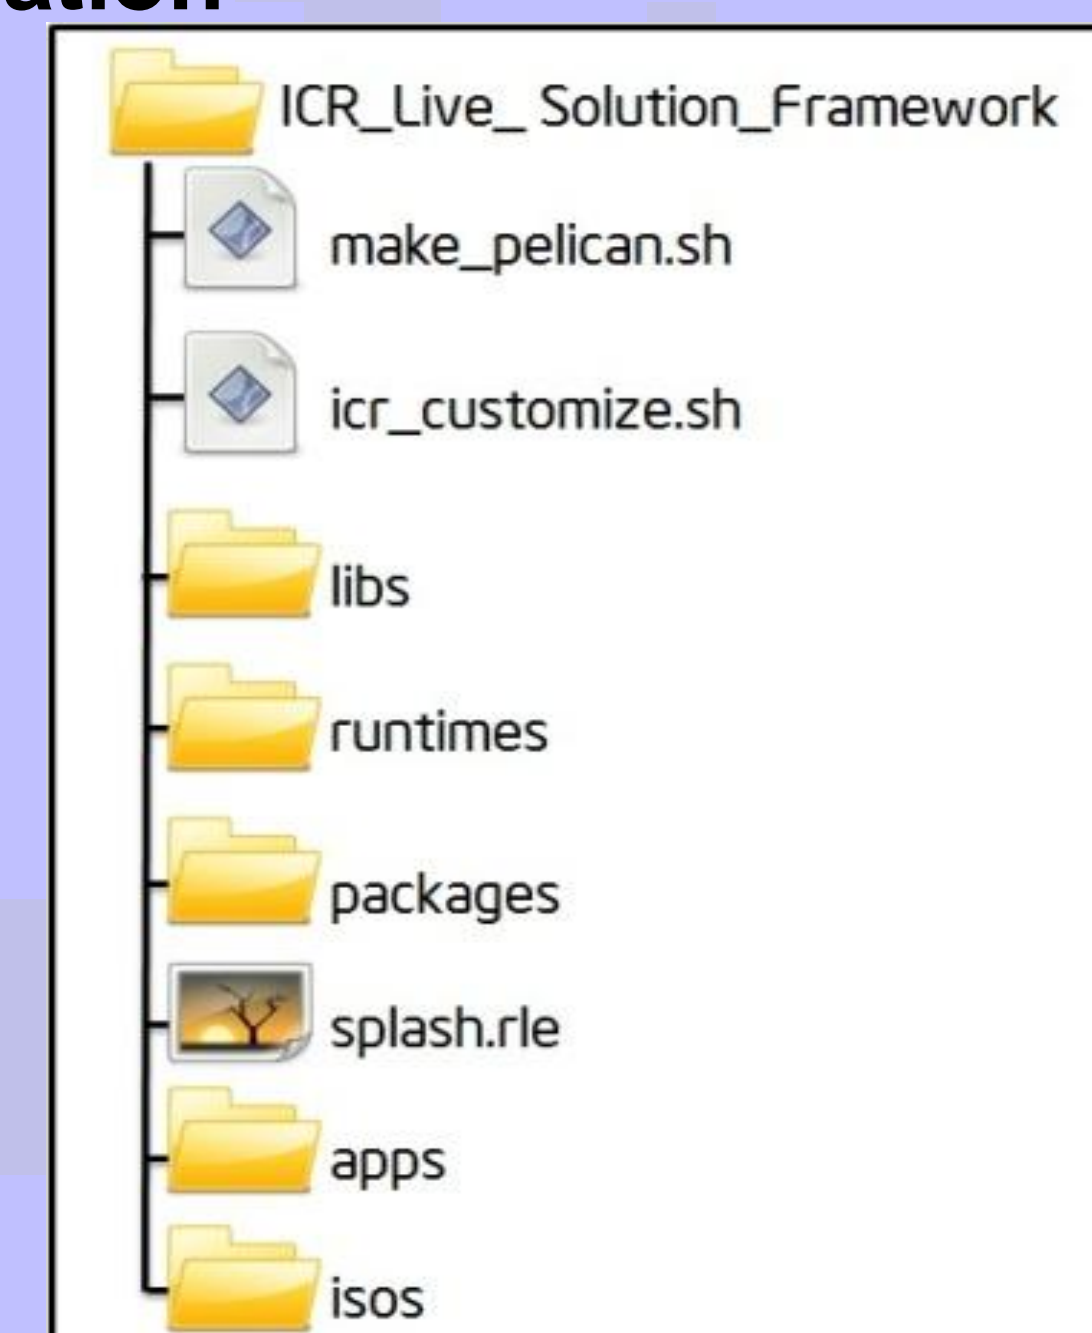
The simplicity showed in the generation and customization process, plus the flexibility and good results obtained using the Intel® Cluster Checker tool, made us select PelicanHPC as foundation of the customization framework. PelicanHPC is a Debian-based diskless solution that uses the live-helper* framework created by the Debian* Live Project.

Diskless Solution Generation Process

The generation process involves components and actions to create a diskless solution from scratch. The figure illustrates the relation among these components and how the process works with the inputs and outputs. The framework simplifies the customization of a diskless clustering solution to meet user needs by setting parameters and adding new capabilities.



Framework Implementation



Diskless Solution Validation

- **Deployment Time Reduction:** Average deployment time for a 32 node cluster using the diskless solution was 15 minutes against 3 or more hours required by an expert.
- **Easy to use:** Previous knowledge to use the framework was not required.
- **Always Intel® Cluster Ready:** Non-expert users can develop a customized diskless solution that matches Intel® Cluster Ready specification.
- **Infiniband Support:** the OFED software stack is included and configured as part of the generated diskless solution.
- **Easy to run applications:** the customization framework allows HPC applications to be included in the solution when the cluster is deployed.

Conclusions

The main result of this work is to provide users with a simple method to enter the HPC world. The developed framework gives the user a customized diskless solution that enables him to create and deploy a diskless clustering solution with little or no need of expertise in HPC.

Another advantage is that the deployed cluster also complies with the Intel® Cluster Ready architecture. The architecture guarantee that registered HPC applications in the program will be executed successfully.

These customized diskless solutions generates a great opportunity to approach the HPC world of computing and their application to industry, increasing the expansion of the usage of HPC in emerging markets.

This opens a good opportunity to use these customized diskless solutions to deploy on-demand HPC systems. This usage model allows the user to create their customized image and deploy a cluster that fit his requirements.

Next Steps

The research will continue moving to use the diskless solution generated by the customization framework as base of a compute on demand model that satisfy specific end user needs and uses rented hardware from third parties.

Thanks

- To Intel for giving us the opportunity of making this research as part of an Internship activity.
- To the ACE (Argentina Cluster Engineering) Team that helped in the research, specially to Ricardo Medel, Dan Hirsch, Cesar Martinez and Andres More that unselfishly gave their time in the review of this work.