

A Simple User Administration View of Computing Clusters

Valeria M. Bastos, Myrian A. Costa, Matheus Ambrozio, Nelson F. F. Ebecken

Abstract—In this paper a very simple and effective user administration view of computing clusters systems is implemented in order of friendly provide the configuration and monitoring of distributed application executions. The user view, the administrator view, and an internal control module create an illusionary management environment for better system usability. The architecture, properties, performance, and the comparison with others software for cluster management are briefly commented.

Keywords—Big data, computing clusters, administration view, user view.

I. INTRODUCTION

THIS document presents EPĀAK, a simple user administration view of computing clusters. A computer cluster simply gives you more processing power than you have with one computer alone. For normal computer use and typical user applications, one computer is all you need. But for people who desire to write their own algorithms to solve specific problems they are working on, they sometimes find that just having one computer for crunching numbers will tie up their single computer for hours, days, or weeks before it converges on a solution.

A computer cluster is just a bunch of computers connected to a local network or LAN. An application written to run on a cluster needs to be designed specifically to utilize the cluster. Usually there is the "main" application running on one main computer that breaks the problem up into manageable chunks and sends them out to the "worker" applications running on the other nodes of the cluster.

A cluster manager usually is a backend graphical user interface (GUI) or command-line software that runs on one or all cluster nodes (in some cases it runs on a different server or cluster of management servers.) The cluster manager works together with a cluster management agent. These agents run on each node of the cluster to manage and configure services, a set of services, or to manage and configure the complete cluster server itself (supercomputing.) In some cases, the cluster manager is mostly used to dispatch work for the cluster (or cloud) to perform. In this last case, a subset of the cluster manager can be a remote desktop application that is used not for configuration but just to send work and get back work results from a cluster. In other cases, the cluster is more

related to the availability and the load balancing than to computational or specific service clusters.

EPĀAK aims to provide users a global view of computing clusters systems, enabling the configuration and monitoring of users environments for execution of their applications and distributed jobs. It allows the management of several services, in independent environments among users. EPĀAK was developed for CentOS 6 [1] or Cloudera [2] based environment, in order to manage an integrated system of one or more clusters, with different configurations or purposes, as shown in Fig. 1, and different administrators and users can be assigned to each of them.

In typical high performance computing environments, several libraries and languages are necessary for the development and execution of parallel and distributed jobs, such as: MPI, OpenMP, R, scientific and machine learning libraries, and Hadoop/MapReduce [3] packages. The challenge of managing these services and libraries is in system administrator role, restricting, in certain situations, the end user's ability to configure the most suitable working environment.

II. USER INTERFACE

EPĀAK provides support for two different views of the computing environment: a User View and an Administrator View. User View lets users to handle all essential services for their applications and jobs (Fig. 2). In Administrator View, system administrators effectively control the entire system and can take corrective actions if necessary.

A friendly web interface shows the information collected from open-source tools, such as hardware health, processor and memory loads, and network and disk states. User visualizes the overall and historic state of clusters systems, such as CPUs usage, memory load, disks access and network throughput. Besides that, all jobs currently running and waiting in processing queues, active and available services for users are shown.

The system handles information from open-source tools:

- Ganglia [4]: Hardware resources graphs and machine states are create by gmon and collected by gmetad.
- Torque/Maui [5]: Job submission and management interface uses torque/maui as job scheduler.
- Software Collections [6]: Organize and manage different versions of each software and/or service installed on the clusters.
- Supervisor [7]: Make automated control of services, such as initialization, finalization, and restarting.

N. F. F. Ebecken, M. A. Costa, M. Ambrozio and V. M. Bastos are with the Federal University of Rio de Janeiro, RJ, CEP21945-970 Brazil (phone: +552139387389; e-mail: nelson@ntt.ufrj.br, myrian@ntt.ufrj.br, valeriab@ntt.ufrj.br, mambrozio@poli.ufrj.br).

Once users can see the workload in each computational node of each cluster, they are able to decide in which of them their jobs will be executed, optimizing their jobs time. Furthermore, users can reconfigure their views and choose the essential services for them. User configuration can be saved for future use in the next login.

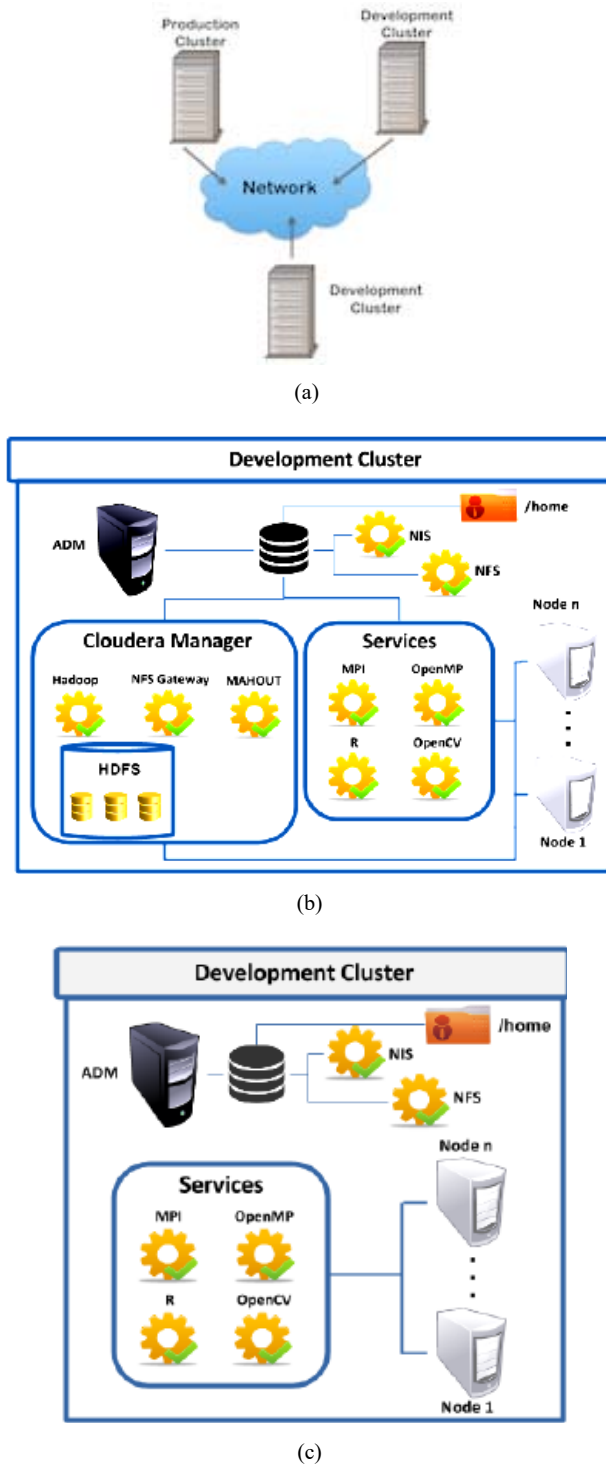


Fig. 1 (a) Computing Clusters System, (b) Development Cluster, (c) Development Cluster, (d) Production Cluster

When a user needs a software or library that is not available for him, he can place an order containing the version of software or library. The administrator packages the software and use the web interface to Epiak administrators to install that version in all compute nodes at once with log recording. After the installation, the user receives an email informing that the new package is ready to use and can be chosen for his environment.





Fig. 2 (a) Login Screen, (b) Hardware health and components loads, (c) Services management

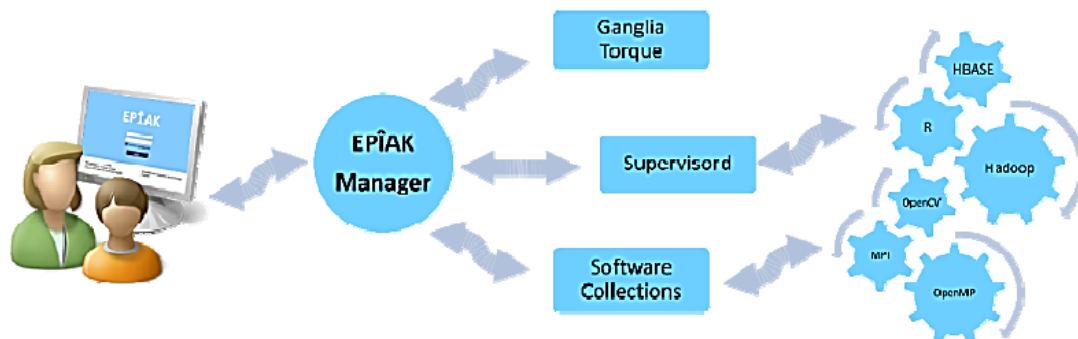


Fig. 3 EPIAK Architecture

III. EPIAK MANAGER

An internal control module, EPIAK Manager, manages the whole system and determines what users can do and view, creating an illusion of management for them. The Manager supplies users with a web interface where they manage their specific environment. It acquires the overall state of the system and services from open-source tools, and makes it available for users.

The manager handles each user request, verifying the state of specific service or library that are related to the desired operation or authorize the job or application execution. Any settings made by a user do not impact the system utilization from others. If a service stop is requested, the manager shows to user that this service is stopped, even if another user is viewing this service as active.

The environment module tool is utilized to manage multiple versions of software and libraries. When receiving a request for installation of a software version, the cluster administrator makes a number of choices, sending source code in tar format, macros, and environment variables through web interface setting. Several automatic scripts install the necessary software and library and their dependencies, creating proper directories and environment variables. After installation completion, Epiak sends mail to the administrator informing the success or failure (with failure point). In case of success, an email is sent to the user that the application software can now be used. This installation procedure encapsulates software outside of the standard system tree, thus making maintenance easier for installation of packages versions, without generating conflict problems.

The Epiak can be installed in more than one cluster and can set different administrators and users to each of them. The administrator view can be changed reflecting the specific configuration of each cluster.

Others management cluster software and packages has similar behavior in administrator view but they don't have the user configurable view that Epiak has. For example, OpenStack [8] is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control. It provides the resource provisioning for the user (IaaS - infrastructure-as-a-service).

Epiak project does not share that goal. It seeks an ease of use of a compute cluster by users, assisting in the selection and use of libraries and other services so that the user does not worry about the settings for a specific environment and often laborious to install, upgrade and maintain working.

Epiak not seek to control the physical resources of clusters, but assist the "Computing Environment" user, improving the view that the user has one or more clusters so that it can decide which machine(s) to run your code.

Epiak does not manage virtual machines, network routing and storage. It provides the machines usage information for the user to know the status of one or more clusters.

Epiak project through a web interface, it provides everything a user of a compute cluster needs to run its applications, including the selection of "computing

environments" pre-configured. For an administrator, the Epiak provides most of the cluster management also via a web interface.

IV. CONCLUSION

Based on the experience of this year, we can conclude that:

- The integration system of several clusters brings scalability and flexibility for production environments;
- The end user's ability to configure the most suitable working environment improves system usability;
- Creating independent and different user environments enhances system utilization enabling shared use of resources and more served users.

The future work that we are planning includes:

- Data flow management among clusters with synchronization of selected datasets;
- Allow users to create multiple environments for different usage profiles;
- Increase the number of managed services and applications.

ACKNOWLEDGMENT

This paper was supported by FAPERJ, the Rio de Janeiro Research Foundation, and CNPq the National Research Council.

REFERENCES

- [1] CentOS System Administration Essentials, Andrew Mallett, Packt Publishing, 2014.
- [2] Cloudera Administration Handbook, Rohit Menon Packt Publishing, 2014.
- [3] Hadoop: The Definitive Guide, Tom White, O'Reilly Media, 4th Edition – 2015.
- [4] Ganglia Monitoring System - <http://ganglia.info> 07/23/2015.
- [5] Torque Resource Manager <http://www.adaptivecomputing.com/products/opensource/torque> 07/21/2015.
- [6] Software Collections <https://www.softwarecollections.org> 07/21/2015.
- [7] Supervisor - <http://supervisord.org/index.html> 07/27/2015.
- [8] OpenStack: The Open Source Cloud Operating System <https://www.openstack.org/software/> 07/23/2015.