

# Crowbar: Deploying Big Data Benchmark Configurations

## Abstract

Big Data configurations typically use commodity hardware to provide fast and efficient access to data by moving jobs to where data resides. Given the size of the dataset, the cluster has to provide a large number of worker nodes to process data and to host partition replicas. In that regard, for a similarly sized dataset, Big Data configurations require more servers than those that would be used by traditional Data Warehouse (DW) clusters where data typically resides on disk farms external to the processing servers.

In order to obtain the maximum performance of a System under Test (SUT), benchmark' configurations have to be sized with more hardware than actual production systems. With that reality, it is expected that Big Data benchmark SUTs will be more complex and more difficult to install and configure. For ease of benchmarking, there will be a need to provide a tool that will ease the pain and the time of deploying such benchmark configurations.

Crowbar is an open source software framework, developed by the Dell Cloud Computing solutions team to deploy large scale Cloud clusters. It has been successfully used to deploy Dell OpenStack cloud and Dell Cloudera Hadoop solutions.

## Architecture

Crowbar provides a modular platform containing the building blocks to provision, monitor, and operate a large scale cloud deployment. Starting with bare metal installation, Crowbar automates all the required installation and configuration tasks. The core capabilities provided are:

- Hardware configuration – updating and configuring BIOS and BMC boards.
- Deployment of base operating system.
- Deployment of cloud/big data components.
- Providing core network infrastructure services (NTP, DNS, DHCP).
- Monitoring availability and performance of all deployed components.

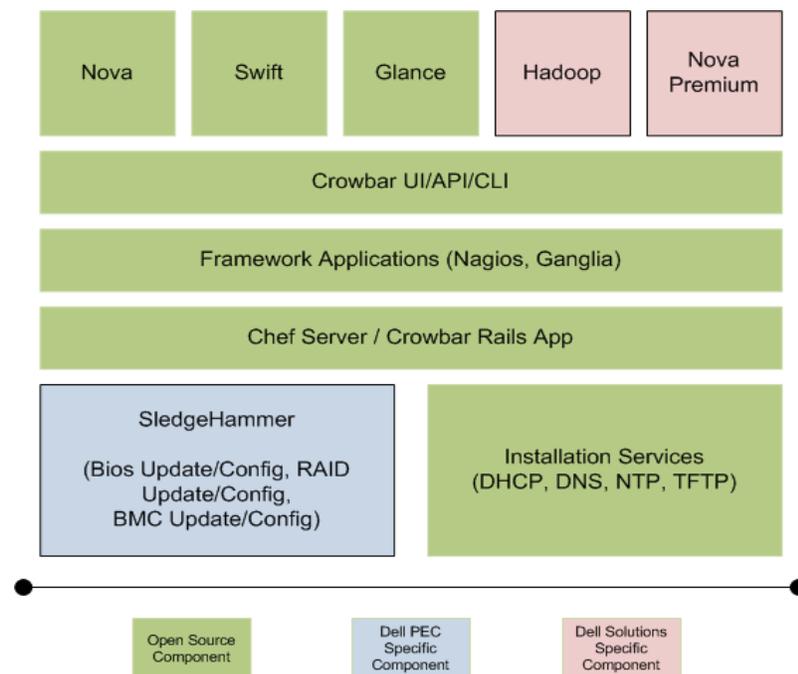


Figure 1: Crowbar high-level components

Each function in the cluster is controlled by a Crowbar component called a Barclamp. There are Barclamps for Nagios, Ganglia, NTP, and a variety of other basic services. Each Barclamp is responsible for all the aspects of the underlying technology required to make it usable.

Crowbar makes extensive use of Opscode Chef Server which provides an environment where abstract definitions are written as source code that describes how each part of the infrastructure is to be built, and applied to individual servers.

## Deployment

To allow easy integration into an existing environment, Crowbar allows customization of its components. For instance, the default monitoring tools (Nagios and Ganglia) can be disabled if other existing tools are preferred. Internal cloud services can be connected to extant services; for example, the cloud's NTP service can be configured to synchronize with existing servers.

To control the operation of a Barclamp, a proposal for the Barclamp is created or edited (if one is already in place). A proposal comprises several parts:

- Parameters to customize the operation of the function; for example: upstream DNS resolvers.
- List of machines in the deployment that fulfill the different roles in the function.
- Internal system information.

When provisioning a function, a proposal generated by Crowbar has to be created. Each core service running on the admin server has a default proposal included as part of the Crowbar installation. These proposals have to be edited before installing these services on the admin node.

When a proposal is committed, Crowbar configures the Chef server and other components in the systems (TFTP, DHCP, and so on) to build the setup described in the proposal. Machines in the deployment affected by the proposals have their configuration updated using Chef client-side components. At the end of the process, the function described by the proposal is ready for use.

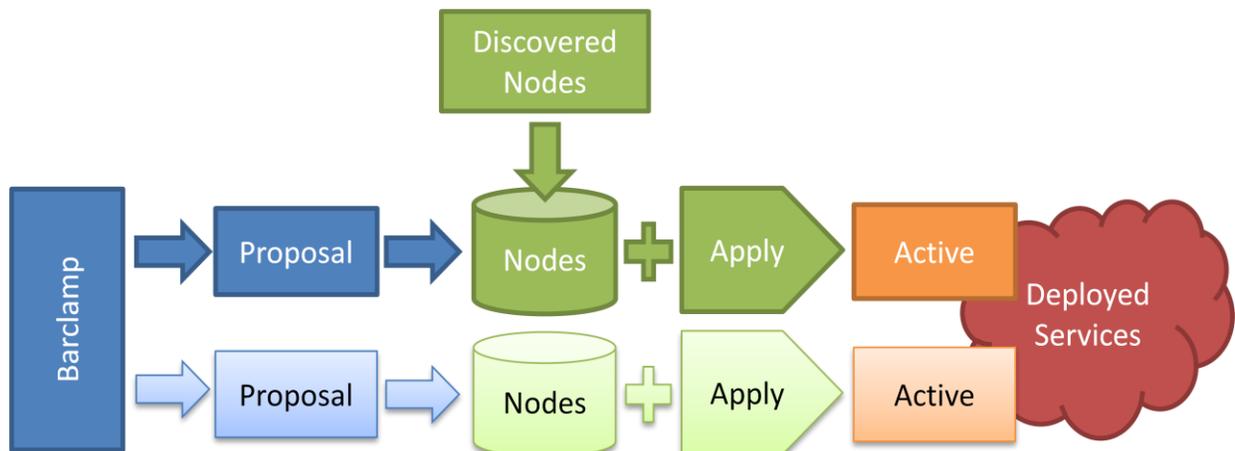


Figure 2: Crowbar process

Finally, a cloud deployment is dynamic. Machines come and go, break down, or get repurposed. Crowbar's operational model makes sure that machines are hooked into the key infrastructure services. Critical services (for example Nagios monitoring) are installed automatically on newly provisioned machines by default, and those machines may be easily allocated for use with any additional Crowbar services desired.

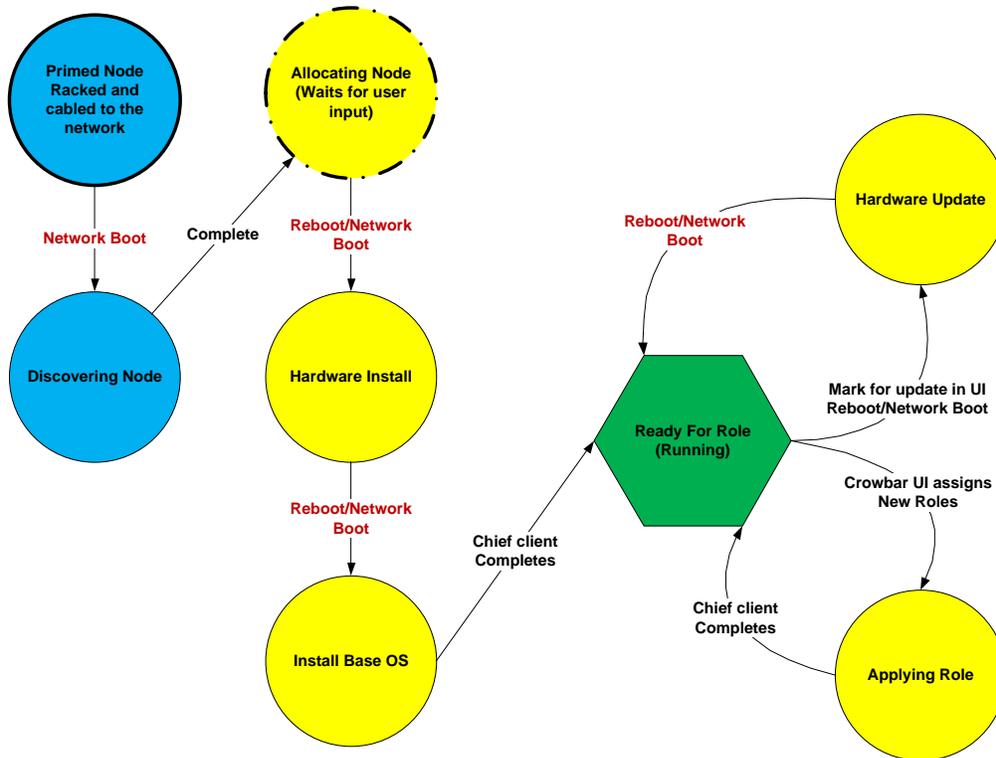


Figure 3: Deployment life cycle

## Nicholas Wakou

Nicholas is a Systems Engineer/Senior Consultant with the Cloud Computing team of the Dell Next Generation Compute Solutions Group. He has been actively involved in developing and running TPC benchmarks since the year 2000. He is Dell's primary representative at the Transaction Processing Performance Council (TPC) and SPECcloud development committee. He is a past Chair of the TPC Public Relations Committee and a past member of the TPC Technical Advisory Board (TAB).