


Scientific Cluster Deployment & Recovery Using puppet to simplify cluster management

V. Hendrix¹, D. Benjamin², Y. Yao¹
¹Lawrence Berkeley National Lab, ²Duke University

- A cluster deployment and recovery process based on the puppet configuration management engine which allows a part-time FTE to easily deploy and recover entire clusters with minimal effort

The Problem




- A university **scientist** needs to wear an additional **system administrator** hat.
- This extra duty of maintaining a **scientific data analysis cluster (DAC)** causes the scientist to be diverted from practicing science.

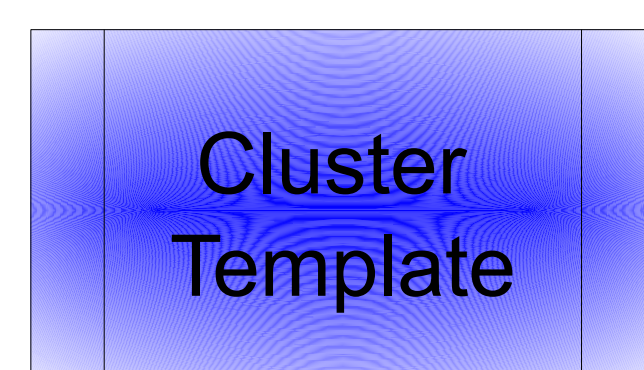
The Solution

A cluster administrator (the scientist) uses a **cluster template** and a **cluster tool** to generate a **cluster definition** that is used for turnkey DAC deployment

- Cluster Template** a set of cluster node roles defined by service definitions and their relationships
- Cluster Tool**: a python application that understands the cluster template structure and produces a cluster definition
- Cluster Definition**: a set of kickstart installation scripts combined with a complete puppet server configuration.



The Components



A **cluster template** is created by a *cluster designer*. It has

- configuration parameters
- templates
- role and service definitions

A **role** is the function of a cluster node such as HEAD and WORKER nodes in a batch cluster. A role is defined by

- The services it provides
- The services it is a client of
- How those services interact

A **service** is managed by a puppet module definition.

- Puppet defines the target state of a node with its declarative language.
- Puppet modules can be published in a version control systems such as git or SVN.

node
role

Service

Cluster
Definition

A **cluster definition** is used to deploy a scientific analysis cluster

- Kickstart install and other configuration scripts
- Puppet server configuration

A **kickstart** configuration file is used to perform kickstart installations for each node via USB key or PXE boot.

kickstart.cfg

Other configuration **scripts** are generated depending on the cluster needs such as shell scripts for installing virtual machines.

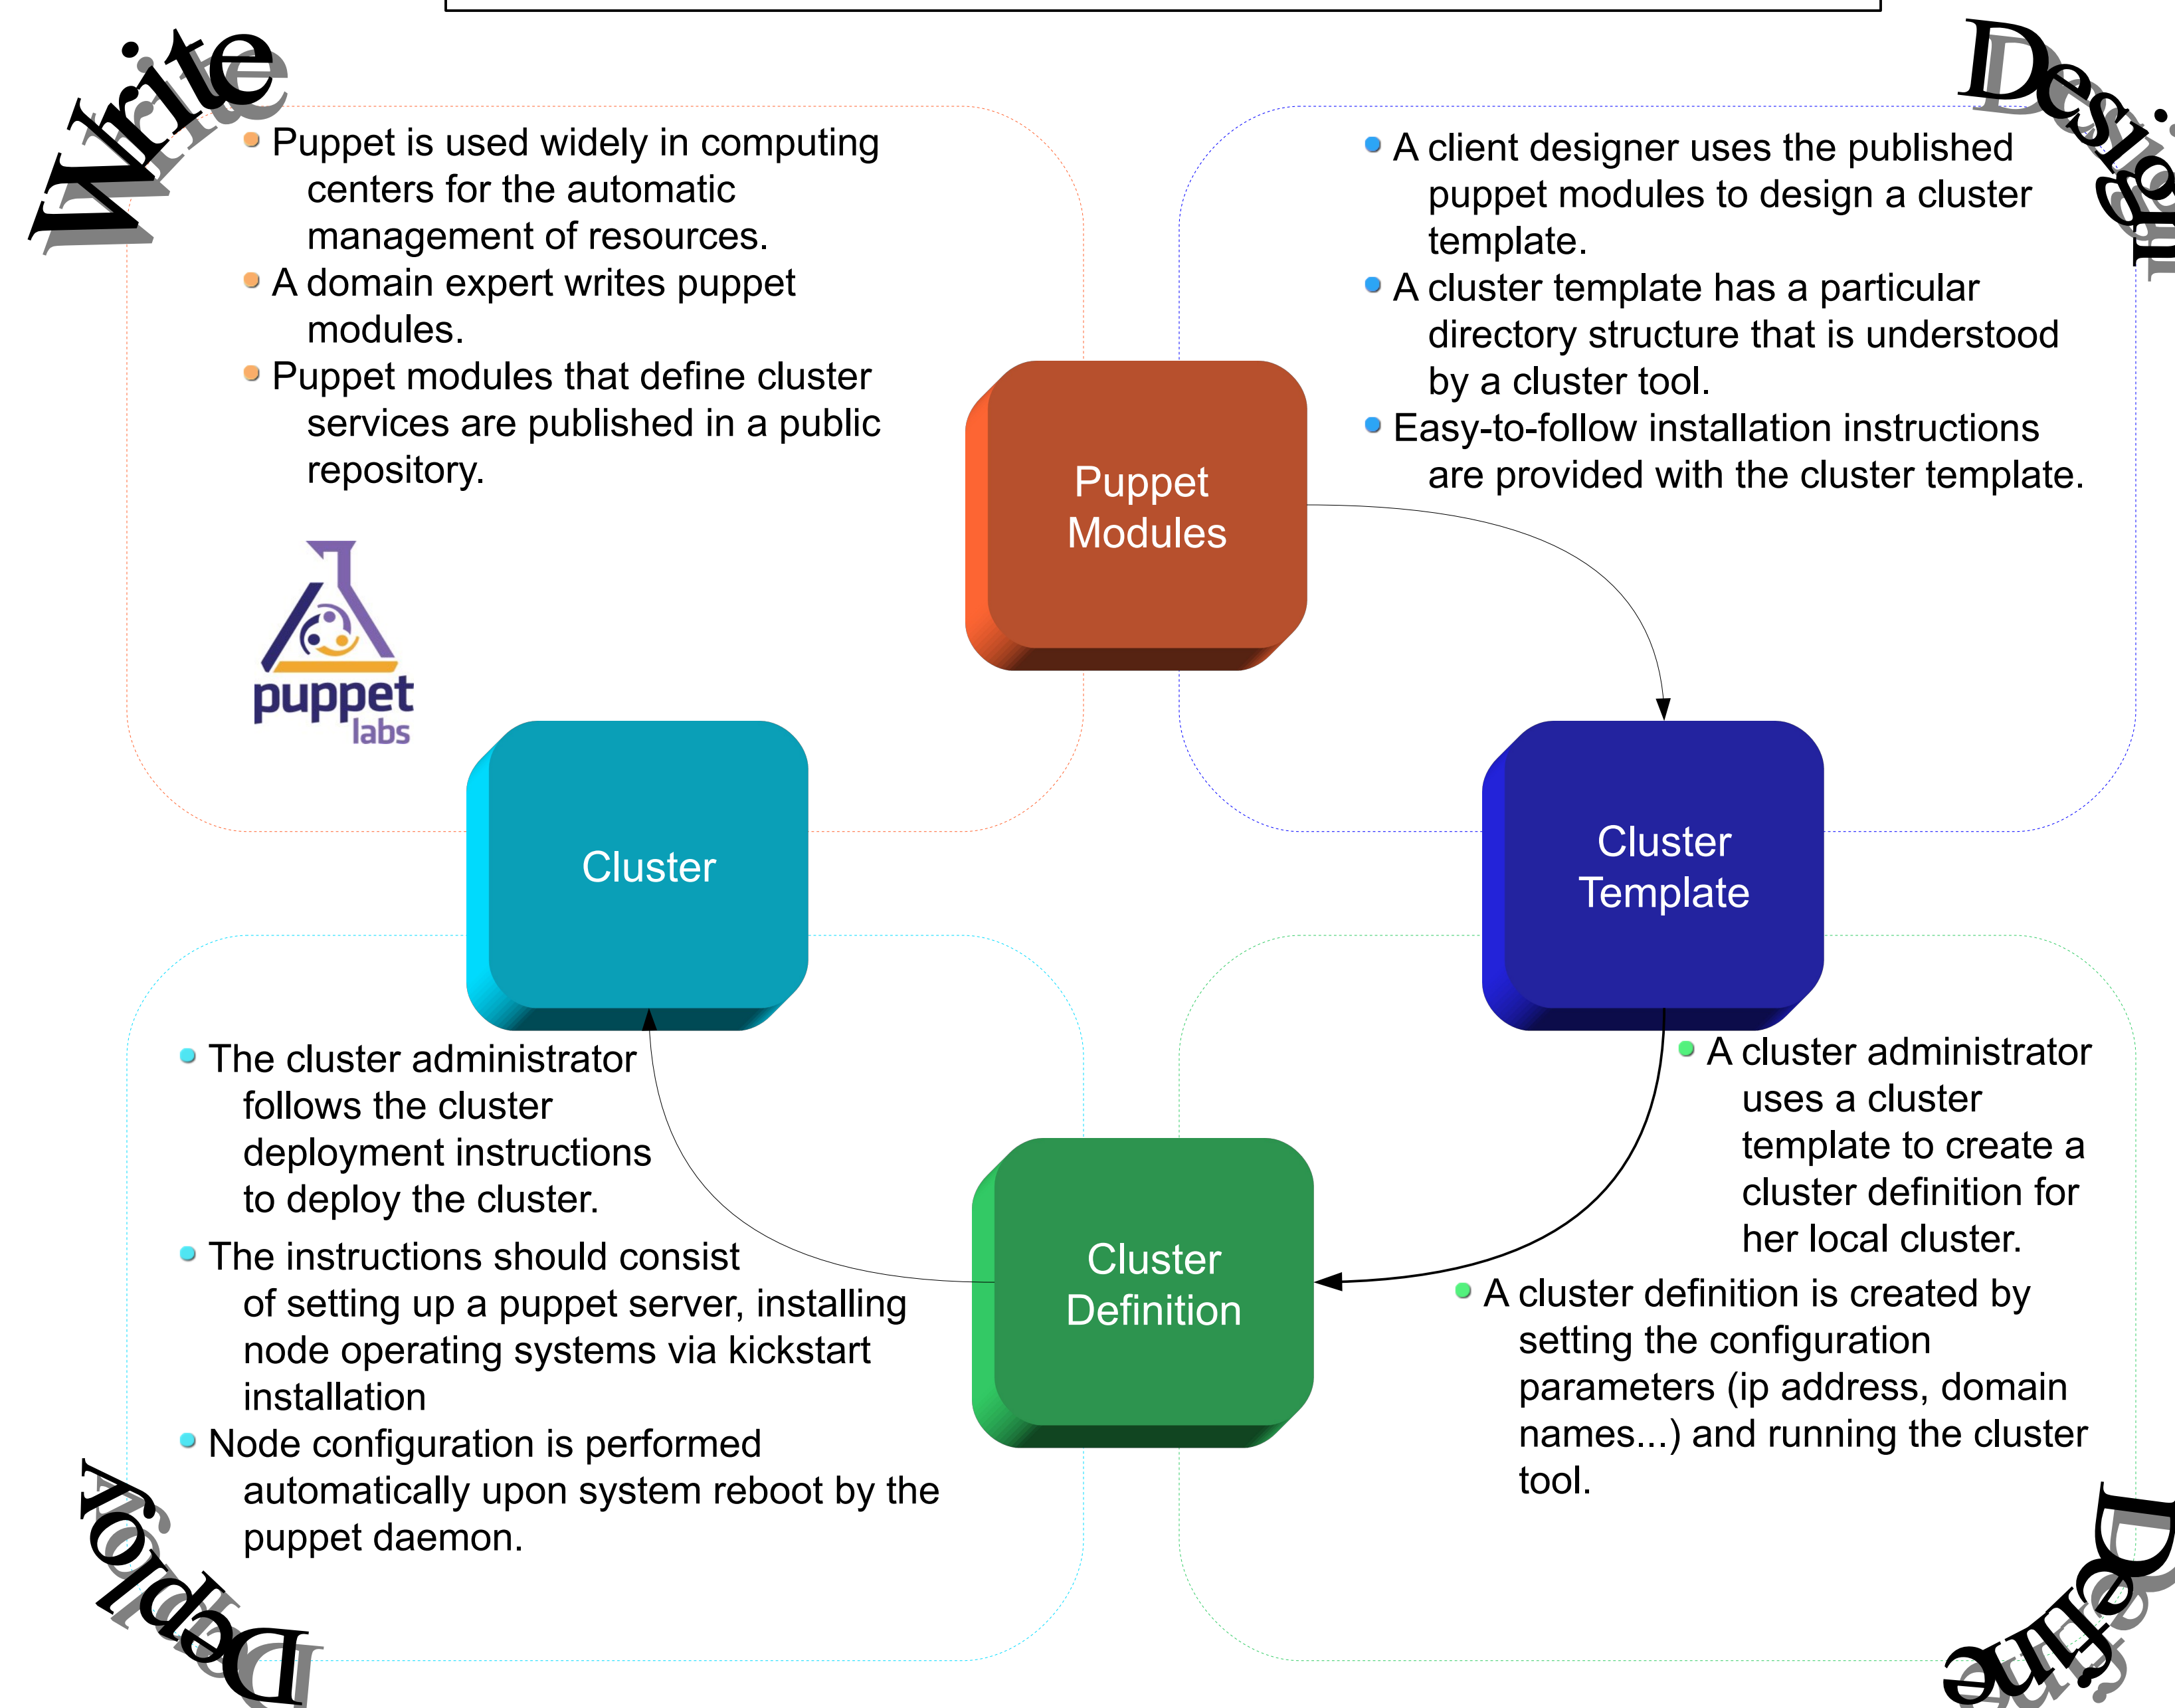
scripts

The Actors

There are several actors in this process from the domain expert who understand the intricacies of the clusters services to the end user who submits analysis jobs on the cluster.


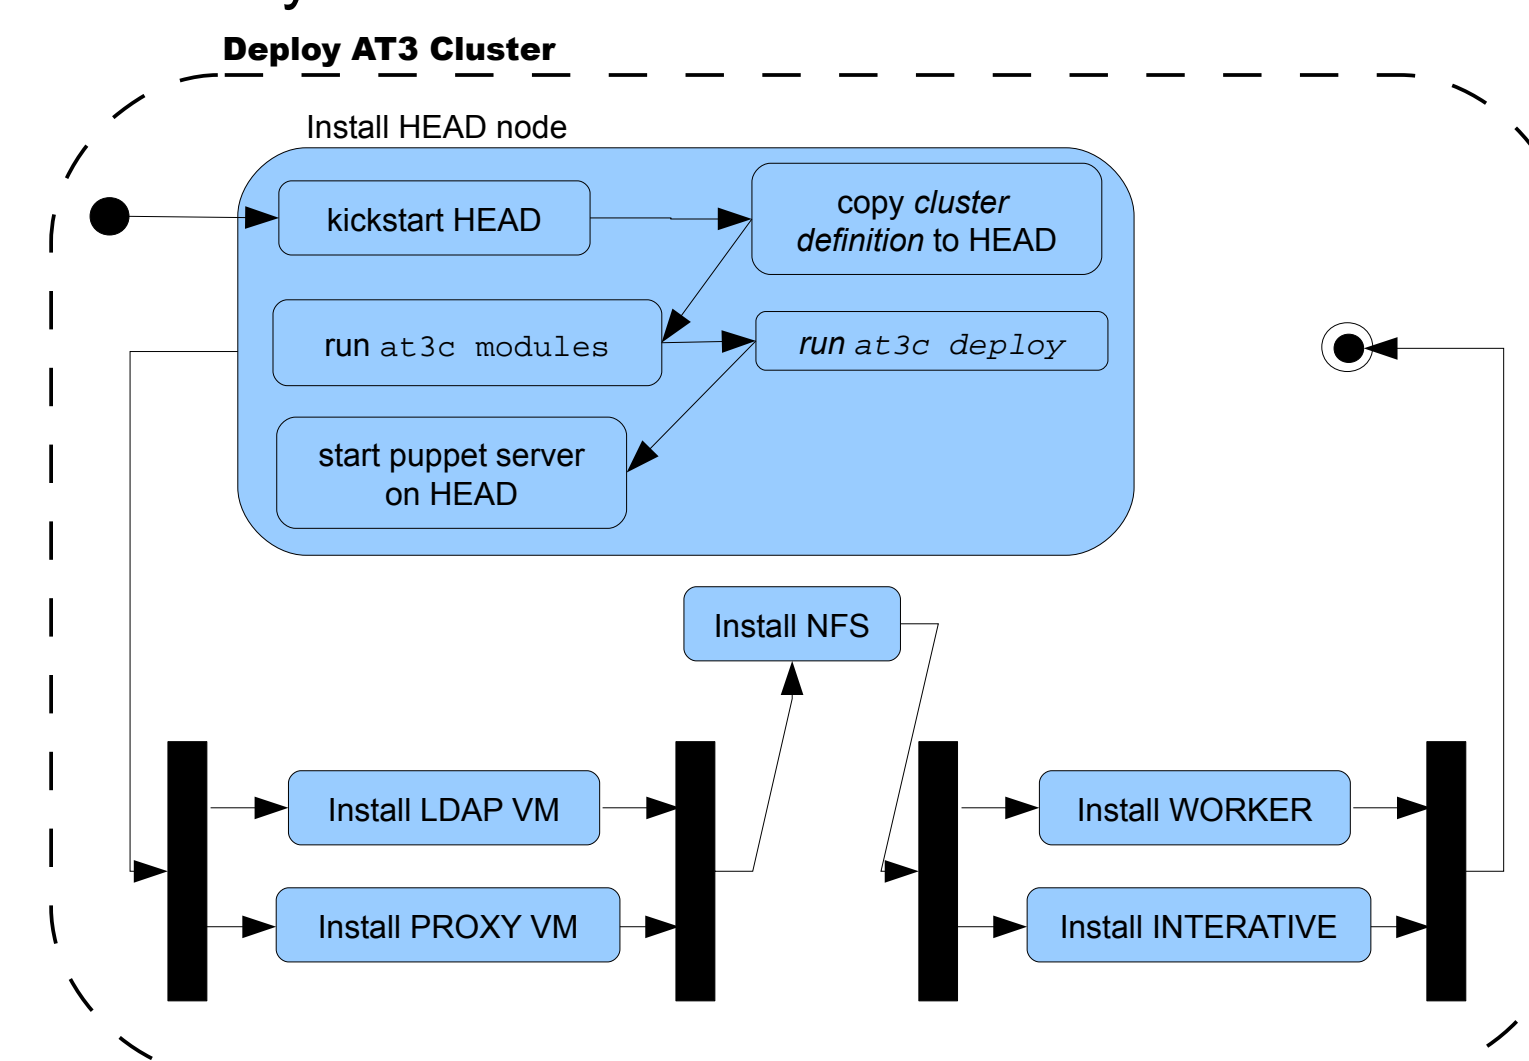
- Domain Expert**: a domain expert writes puppet modules for configuring and managing services used by data analysis clusters (DAC).
- Cluster Designer**: a cluster designer creates a cluster template with the puppet modules designed by the domain expert.
- Cluster Manager**: A cluster manager procures hardware, sets up networking according to the cluster topology and deploys/manages the DAC.
- Cluster User**: A cluster user logs into the cluster and performs a task. In the case of DAC, a user would run analysis jobs.

The Process



Deploy a Cluster

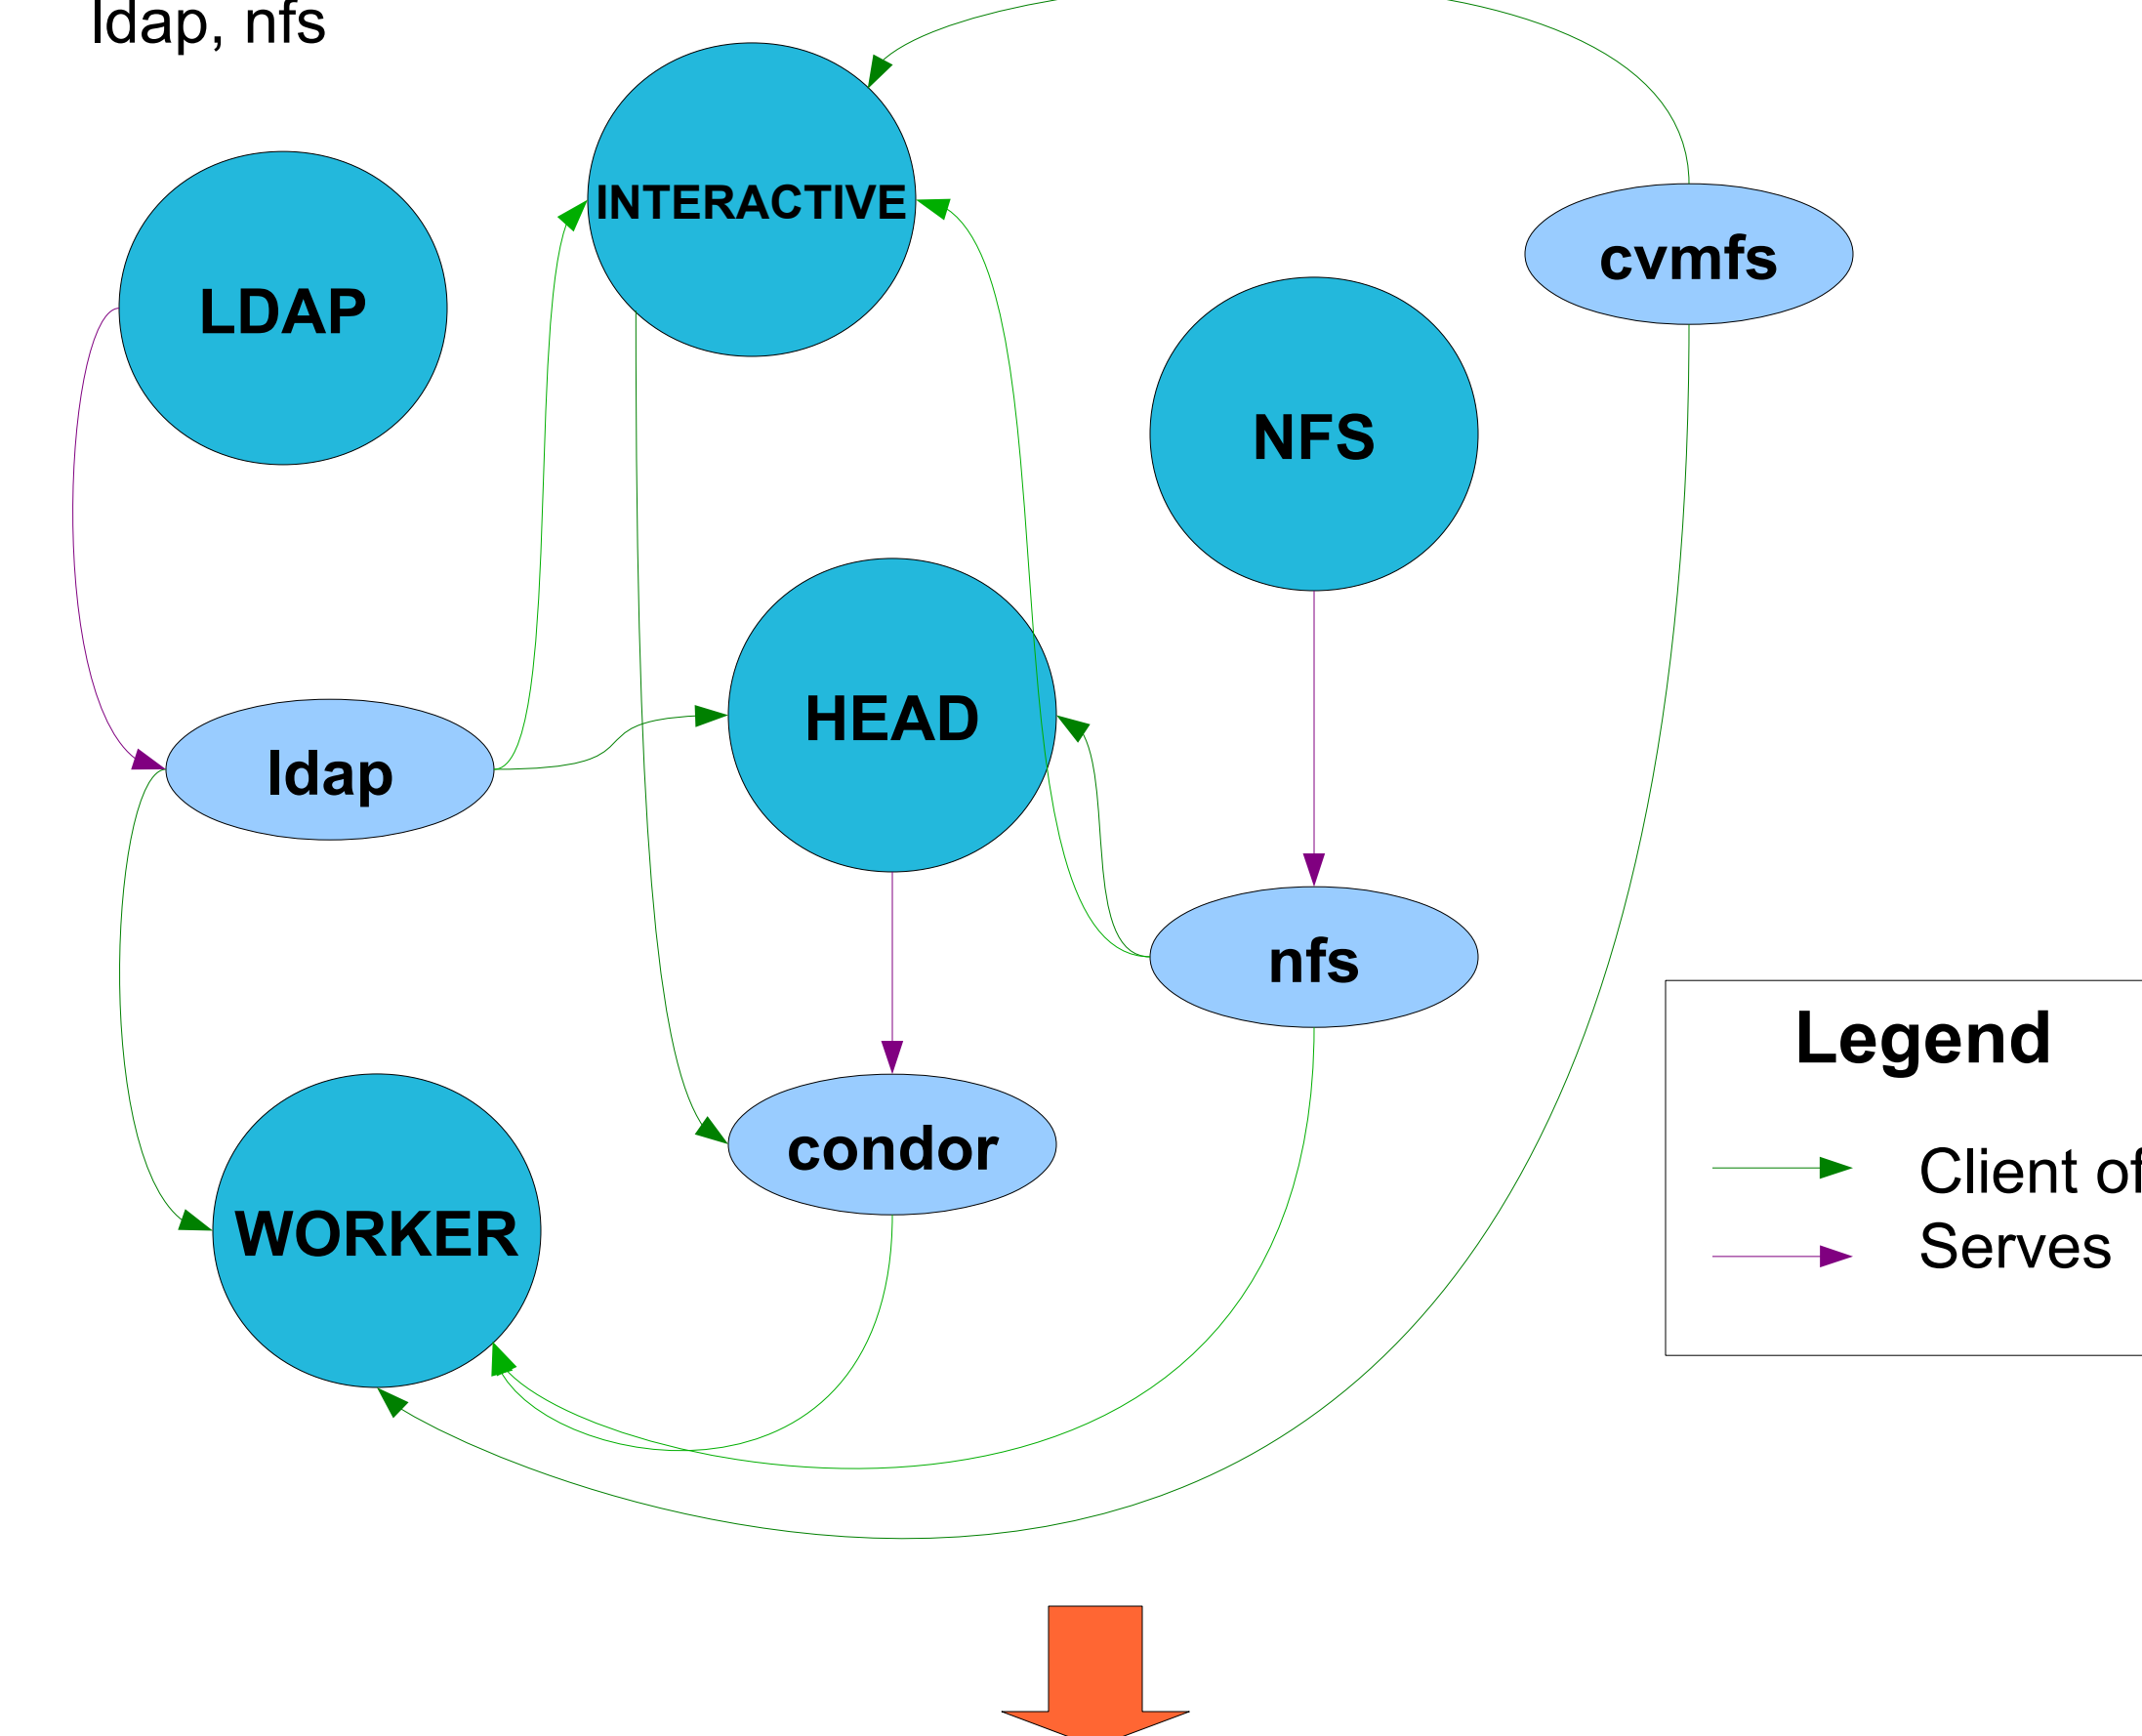
Deploying the cluster is a straight forward process for each node. Using the bootable USB stick, the cluster administrator installs the operating system via kickstart installation. After the operating system installation is complete, the system is rebooted and the puppet client configures the system.



ATLAS Tier 3, an Example

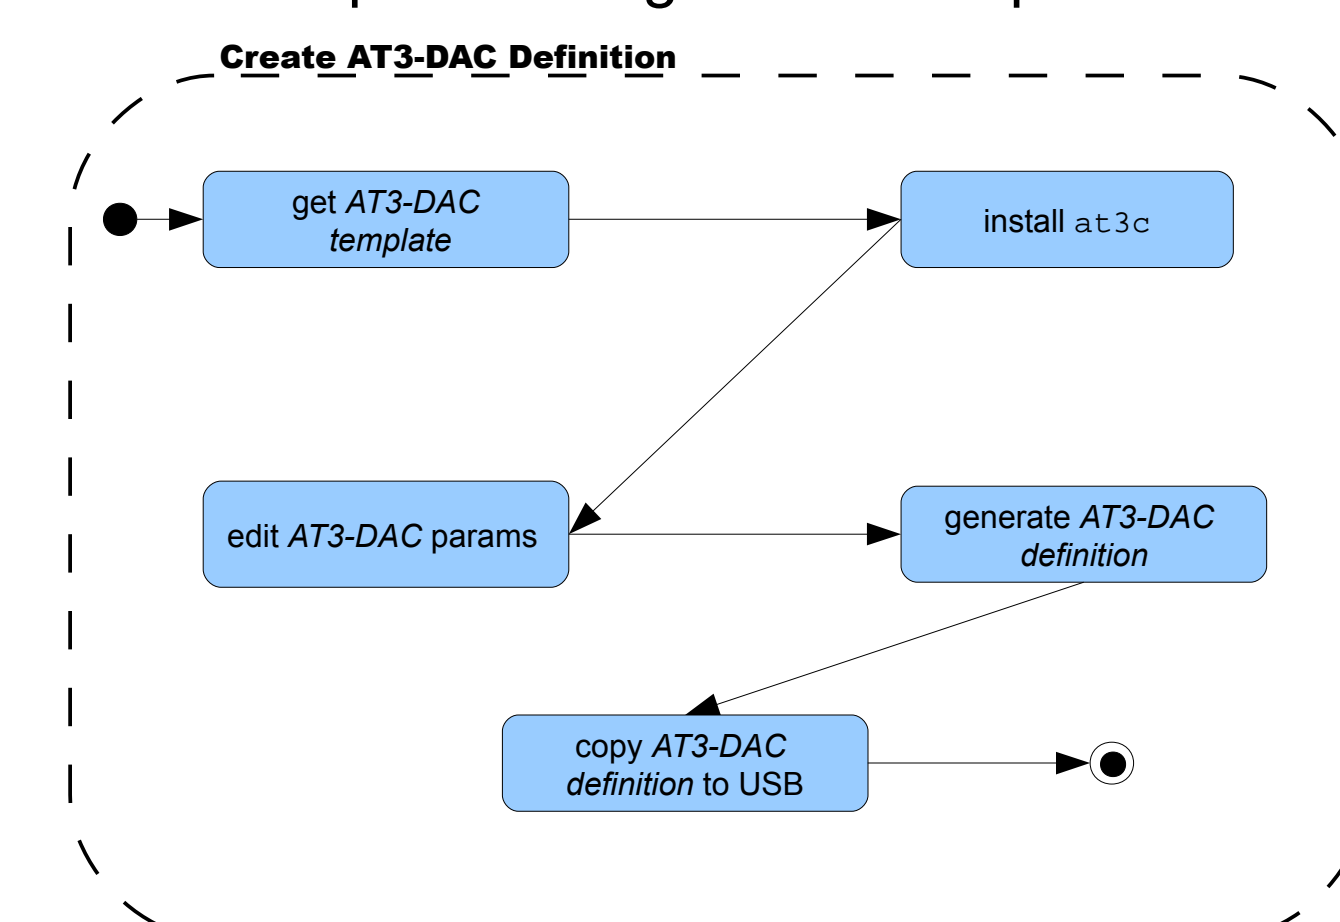
A simplified ATLAS Tier 3 Analysis cluster. The diagram shows the relationship between node roles and their services

- | Roles | Services |
|--|--|
| ■ HEAD : serves: condor head; client of: nfs ldap | ■ ldap : user authentication |
| ■ NFS : serves: nfs service | ■ condor : job submission/execution |
| ■ LDAP : serves: ldap service | ■ nfs : distributed filesystem |
| ■ INTERACTIVE : client of: condor, cvmfs, ldap, nfs | ■ cvmfs : ATLAS software distribution |
| ■ WORKER : client of: condor, cvmfs, ldap, nfs | |



Create a Cluster Definition

ATLAS Tier 3 Cluster tool (at3c) is a set of python commands that facilitates the deployment of ATLAS Tier 3 data analysis clusters. It is a command line tool for creating a cluster definition from a pre-existing cluster template.



Atlas Tier 3 Cluster

- Install the cluster tool from <http://svnweb.cern.ch/guest/atustier3/at3c/trunk>
 - Get the cluster template <http://svnweb.cern.ch/guest/atustier3/at3c-auto-cluster/trunk>
 - Any questions?
- Val Hendrix, vchendrix@lbl.gov