

Ansible Architecture:

There are two main node roles in Ansible. The controller node is the node that runs the Ansible software and from which the operator issues Ansible commands. The controller node can be a server running Linux, an operator laptop, or a system running Ansible Tower. The only requirement is that the controller node needs to be Linux.

From the controller node, the managed nodes are addressed. On the controller node, an inventory is maintained to know which managed clients are available. Ansible doesn't require the use of any agents. That means it can reach out to managed nodes without a need to install anything. To do so, Ansible uses native remote access solutions that are provided by the managed node. On Linux, remote access is realized by using SSH; on Windows, it is realized by using Windows Remote Management (WinRM); and on network devices, it can be provided by using SSH or API access.

To configure the managed clients, Ansible uses playbooks. A playbook is written in YAML and contains one or more plays. Each play consists of one or more tasks that are executed on the managed nodes.

To implement the tasks, Ansible uses modules. Modules are the pieces of code that do the actual work on the managed nodes, and many modules are available—more than 3,000 already, and the number is increasing. Ansible also provides plug-ins. Ansible plug-ins are used to extend Ansible functionality with additional features.

Ansible playbooks should be developed to be idempotent. That means a playbook will always produce the same results, even if it is started multiple times on the same node. As a part of the idempotency, playbooks should also be self-containing and not depend on any other playbooks to be successful