



## Comparative Analysis of Configuration Management Tools: Chef vs. Ansible, SaltStack, and Puppet

---

Harsh Katariya, Madhuri Gedam, Roopali Lolage, Shruti Patil  
and Ujjval Shrivastav

EasyChair preprints are intended for rapid  
dissemination of research results and are  
integrated with the rest of EasyChair.

January 24, 2025

# Comparative Analysis of Configuration Management Tools: Chef vs. Ansible, SaltStack, and Puppet

Mr. Harsh Katariya<sup>1</sup>, Mrs. Madhuri N. Gedam<sup>2</sup>, Dr. Roopali Lolage<sup>3</sup>, Mrs. Shruti Patil<sup>4</sup>, Mr. Ujjval Shrivastav<sup>5</sup>

B.E Student, Department of Information Technology, Maharashtra, India<sup>1</sup>

Asst. Professor, Dept. of Information Engineering, Shree LR Tiwari College of Engineering, Mumbai, India<sup>2</sup>

B.E Student, Department Computer Engineering, Maharashtra, India<sup>5</sup>

[harsh.katariya@slrtce.in](mailto:harsh.katariya@slrtce.in)<sup>1</sup>, [madhuri.gedam@slrtce.in](mailto:madhuri.gedam@slrtce.in)<sup>2</sup>, [roopali.lolage@slrtce.in](mailto:roopali.lolage@slrtce.in)<sup>3</sup>, [shruti.patil@slrtce.in](mailto:shruti.patil@slrtce.in)<sup>4</sup>

**Abstract:** Configuration Management (CM) tools play a critical role in automating infrastructure and maintaining consistency in modern IT systems. This paper presents a comparative analysis of Chef, Ansible, SaltStack, and Puppet, four prominent tools in the domain of Configuration Management. Each tool has unique features, strengths, and challenges that cater to diverse organizational needs. Chef, known for its flexibility and integration capabilities, offers a robust solution for managing large-scale environments but comes with a steep learning curve. Ansible is praised for its simplicity and agentless architecture, making it beginner-friendly. SaltStack excels in speed and scalability, while Puppet's mature ecosystem provides extensive support for complex configurations. This study evaluates these tools based on criteria such as ease of use, scalability, integration, security, and community support. Through case studies and analysis, we discuss how these tools are implemented in real-world scenarios and their effectiveness in addressing key challenges like configuration drift, dependency management, and security compliance. The findings aim to provide insights for organizations and developers in selecting the most suitable CM tool for their needs. Ultimately, while each tool has its trade-offs, the choice depends on specific project requirements and team expertise.

**Keywords:** Configuration Management, Chef Automation Tool, Infrastructure as Code (IaC), Continuous Integration (CI), Dependency Management, Configuration Drift.

## I. Introduction

Software Configuration Management (SCM) is a foundational practice in software engineering, crucial for ensuring the consistency, reliability, and reproducibility of software systems. As software projects grow in complexity, managing configurations manually becomes impractical and error-prone [1]. This has led to the adoption of automated tools that can streamline the management of configurations across diverse environments. Among these tools, Chef has emerged as a powerful solution for automating infrastructure management through code, also known as Infrastructure as Code (IaC) [1]. Chef, developed by Opscode (now Chef Software, Inc.), is an automation platform designed to manage infrastructure on physical, virtual, and cloud environments. By treating infrastructure as code, Chef allows developers and system

administrators to define the desired state of their systems using a simple and declarative language [2]. This approach not only enhances the reproducibility and reliability of configurations but also integrates seamlessly into modern development practices, such as Continuous Integration (CI) and Continuous Deployment (CD) [4].

Despite its advantages, implementing Chef for SCM is not without challenges. Organizations often encounter issues such as managing complex dependencies, ensuring security compliance, and addressing the steep learning curve associated with the tool [4]. This paper explores these challenges in detail and proposes solutions to overcome them. We will examine the principles of Chef, its role in automating infrastructure management, and its integration into CI/CD pipelines [5].

## Advantages of Chef Tool

**Flexibility and Customization:** Chef's use of a Ruby-based DSL allows for highly customizable and flexible configurations, enabling the management of complex scenarios.

**Integration Capabilities:** Chef integrates seamlessly with various CI/CD tools, cloud platforms, and container orchestration systems, facilitating a cohesive automation environment.

**Compliance and Security:** Chef Automate includes robust compliance features, enabling automated security policy enforcement and regulatory compliance checks.

**Scalability:** Chef's master-agent architecture efficiently manages large-scale infrastructures, making it suitable for enterprises with extensive and complex environments.

## II. Chef Architecture

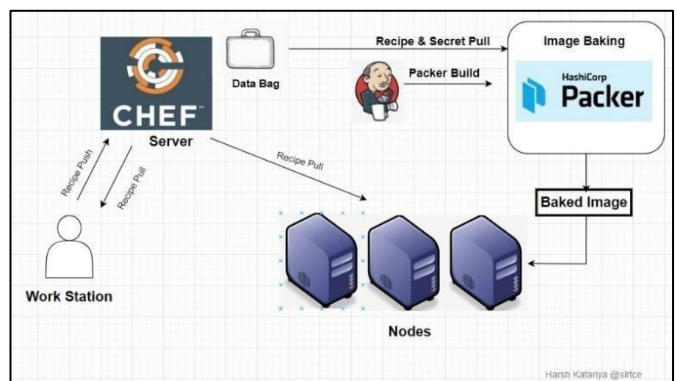


Fig.1. Chef Configuration [6]

### III. Components in Chef Architecture

#### 3.1 Chef Infra:

Chef Infra is the core component of the Chef ecosystem, designed to automate the management of infrastructure. It uses a master-agent model where the Chef server acts as the central hub, managing all the nodes in the system.

**Chef Server:** Acts as a central repository for storing configuration data, cookbooks, policies, and other artifacts required for configuration management.

**Chef Client:** Installed on each node, the Chef client periodically polls the Chef server for updates and applies the necessary configurations to ensure the node's state matches the desired state defined in the Chef cookbooks.

#### 3.2 Cookbooks and Recipes:

**Cookbooks:** Collections of recipes, attributes, files, templates, libraries, and resources. They define how Chef manages system configurations.

**Recipes:** Specific configuration details written in Ruby, detailing the desired state of a node, such as installing packages, creating files, and managing services.

### IV. Related Work

The most popular SCM tools are Puppet, Chef, Ansible, and Salt Stack. Their general characteristics are presented in Table 1. Among them the system development language, Configuration Description Language (CDL), system architecture, software distribution license, and supported platform [5]. Within the execution of distributed application packages, an important requirement is a support for managing Windows and Linux OSes families. The paper does not discuss systems whose development is suspended or discontinued, as well as proprietary systems [5].

**Table 1. General Configuration of SCM Tools [2]**

System	System Development Language	License	Architecture	Config on description language	Platform support	
					Linux	Window
Ansible	Python	GPLv3+	Agentless	YAML-like JSON	+	+
Chef	Ruby,Erlang	Apache 2.0	Client-server	Ruby	+	+
Puppet	C++,Clojure	Apache 2.0	Client-server	Ruby-like proprietary	+	+
SaltStack	Python	Apache 2.0	Agentless/Client-server	YAML-like	+	+

All four systems provide a web-interface for configuration management. It can be used to create reports and visualize the infrastructure configuration. In addition, all of the systems have the ability for connecting to an external monitoring system. SaltStack uses its own implementation for that [5]. The features of the considered SCM tools are reflected in the system characteristics, description specifics, delivery methods, and installation of system configurations. These System characteristics are shown in Table 2.

**Table 2. System characteristics of the SCM tools [6]**

System	Database	Configuration on delivery	Supported server OSes	Step-by-step insallation	Delivery Method
Ansible	-	SSH	IBM AIX,BSD,Linux	+	Push
Chef	PostgreSQL	RabbitMQ	Linux Only	+	Pull
Puppet	PuppetDB	Mcollective	Linux Only	-	Push, Pull
SaltStack	-	ZeroMQ	Linux, BSD	-	Pull

On deploying Puppet and Chef, the Puppet DB and PostgreSQL database systems are used, respectively. These databases provide centralized configuration storage. The performance and scalability of each system directly depend on the database used. In addition, the database server requires additional maintenance.

Enterprise version of Ansible uses PostgreSQL. In addition, it is allowing to install Mongo DB (MySQL) to build fault-tolerant architectures and for storing logs.

### V. Chef vs Puppet

Chef and Puppet are leading configuration management tools, each with distinct approaches and features [6]. Chef uses a Ruby-based imperative language, offering flexibility and detailed control over infrastructure configurations, making it suitable for developers familiar with Ruby but presenting a steeper learning curve for others. Puppet, on the other hand, employs a declarative DSL, focusing on specifying the desired state of the infrastructure, which can simplify the learning process and make it more accessible for new users. Both tools utilize a master-agent architecture, though Chef also offers a standalone mode with chef solo [6].

and Puppet provides Puppet Apply for local execution. In terms of performance, Puppet is recognized for its efficiency in managing large environments through its resource abstraction layer, while Chef's performance depends on the complexity of its recipes. Community support is robust for both, with Chef Supermarket and Puppet Forge offering extensive repositories of pre-written configurations.

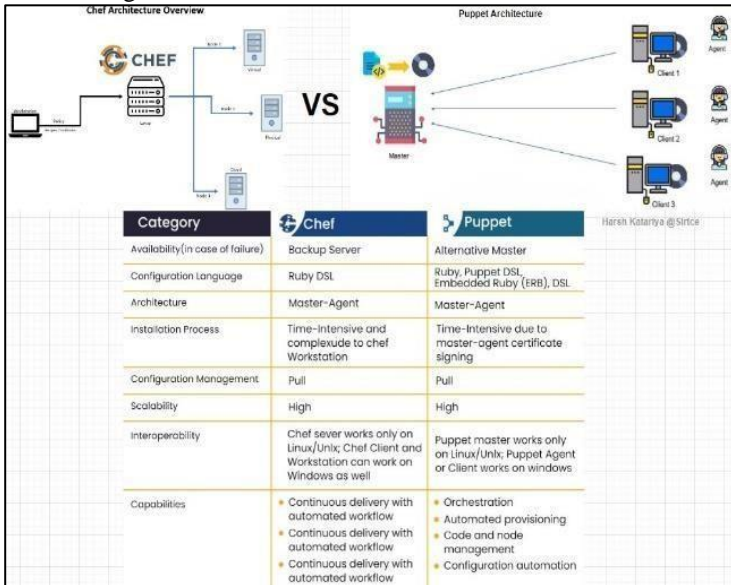


Fig.2. Chef vs Puppet [6]

### I. Chef vs Ansible

Chef and Ansible are both popular configuration management tools, but they differ in several key ways. Chef uses a Ruby-based domain-specific language (DSL) and follows an imperative approach, defining the specific steps to achieve the desired state. It operates on a master-agent architecture, which can introduce complexity in setup and maintenance. Ansible, on the other hand, uses YAML for its playbooks and adopts a declarative approach, specifying the desired end state without detailing the steps to achieve it. Ansible is agentless, leveraging SSH for communication, which simplifies setup and reduces overhead.

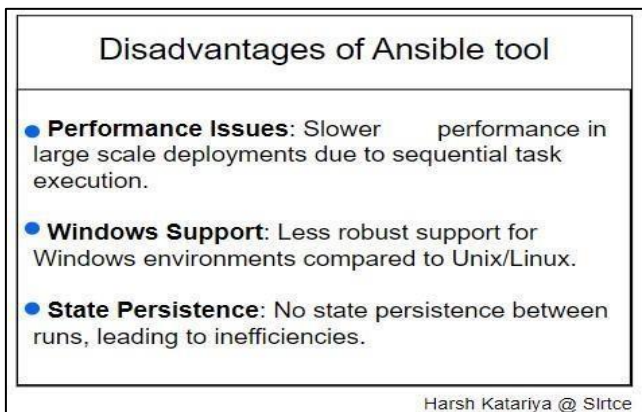


Fig.3. Weakness of Ansible Tool

Imperative approach, requiring users to script the specific steps to achieve a desired state. It operates on a master-agent architecture, which necessitates setting up a Chef server and agents on each managed node. This architecture provides powerful automation capabilities but can introduce complexity in management and maintenance.

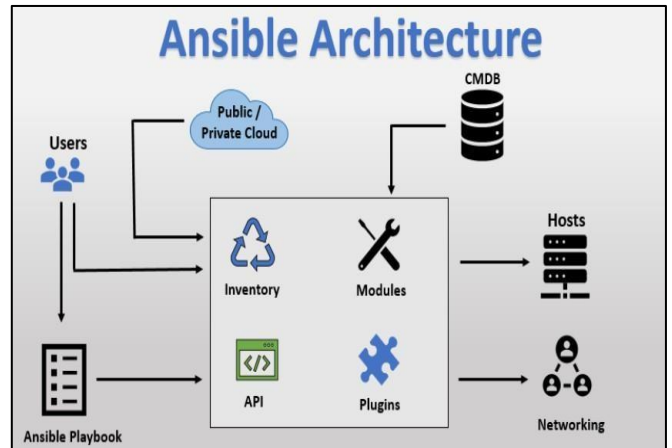


Fig.4. Ansible Architecture

### VII. Chef vs Saltstack

While both Chef and SaltStack are sophisticated tools for automation and configuration management, they serve distinct operational purposes. Chef specifies the precise actions to do in order to reach the desired system state using an imperative approach and a Ruby-based DSL. It uses a master-agent approach, in which agents deployed on nodes are managed by Chef servers.

This offers strong configuration options, but it also means extra setup and upkeep is needed. On the other hand, SaltStack uses an event-driven architecture and a Python-based language to allow agent-based (master-minion) as well as agentless (via SSH) operations.

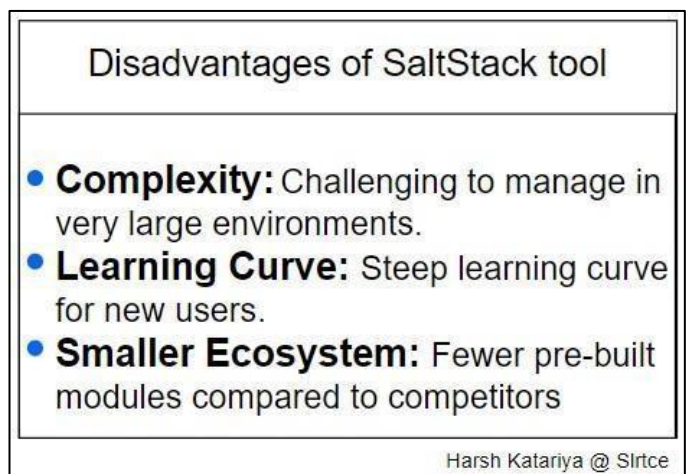


Fig.5. Weakness of Salt Stack Tool

SaltStack can be challenging to configure for beginners due to its steep learning curve and YAML-based syntax, which can become error-prone in large deployments. While SaltStack is designed for scalability, its performance can degrade in very large environments, particularly with high-frequency state runs.

### VIII. Literature Survey

**Table 3. Survey of Various tools [6]**

Tool	Language	Approach	Architecture	Key Features	Strengths	Challenges
Puppet	DSL	Declarative	Master-Agent	PuppetDB, PostgreSQL, Puppet Forge	Simplifies configuration	Setup complexity
Chef	Ruby	Imperative	Master-Agent	Chef Server, Chef Supermarket	Flexibility, detailed control	Steep learning curve
Ansible	YAML	Declarative	Agentless (SSH)	Simple playbooks, minimal setup	Ease of use, fast learning curve	Performance in large environments
SaltStack	Python	Mixed	Master-Minion/Agentless	Event-driven, real-time updates	Speed, real-time capabilities	Learning curve for complex setups

A succinct comparison of the top Software Configuration Management (SCM) tools—Puppet, Chef, Ansible, and SaltStack—can be found in the literature table 3. It draws attention to their salient traits, such as language, strategy, architecture, features, advantages, and disadvantages. This comparison makes it easier to comprehend the unique benefits and constraints of each tool, enabling well-informed decision-making for the selection of the best SCM tool in accordance with particular organizational demands and infrastructure requirements.

### IX. Methodology

**1. Purpose and Objectives:** The main objective is to increase infrastructure management's consistency, efficiency, and scalability by utilizing Chef's automated capabilities to improve Software Configuration Management (SCM). This entails dealing with typical SCM issues like dependency management and configuration drift.

**2. Innovative Approach:** Use a cutting-edge technique that blends Chef with contemporary DevOps procedures. In order to automate end-to-end deployment pipelines as well as system setup.

**3. Integration and Automation:** Boost Chef's integration with current CI/CD tools and monitoring systems by creating unique plugins or extensions. Put in place automated processes that dynamically modify setups in response to input and real-time measurements.

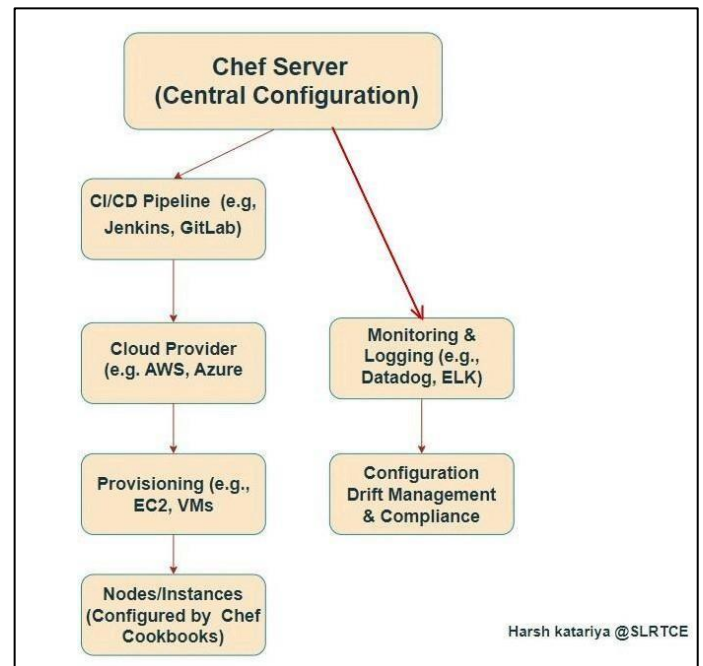
**4. Testing and Validation:** Provide sophisticated testing techniques, such as automated tests linked into the CI/CD pipeline for configuration consistency and compliance checks. Predictive analytics may be used to identify and resolve any configuration problems before they have an influence on operations.

**5. Impact analyses and case studies:** Use this innovative technique to conduct case studies with companies that have used Chef. Analyze the effects on configuration drift reduction, scalability, and operational efficiency. Collect both quantitative and qualitative data in order to evaluate advancements and pinpoint optimal methodologies.

### X. Result

Numerous node kinds, including as servers, network devices, cloud virtual machines, and containers, may be managed using Chef. It oversees several systems, including Windows, Linux, and mainframes. The solution is designed to facilitate application deployment on IT infrastructure by allowing developers and IT operations specialists to collaborate.

#### X.1 Application of Chef



**Fig. 6. Chef Tool Integration in Real-Time Infrastructure Management**

## XI. About Application

The central store for all configuration information, policies, and recipes is the Chef Server.

**CI/CD Pipeline:** Integrates with Chef to guarantee that settings are updated when code changes are pushed to the repository and to automate deployments.

**Cloud Provider:** Manages the Chef settings for the provisioning of servers or instances (such as EC2 in AWS).

**Monitoring and logging:** Gathers metrics and data in real-time from nodes to make sure settings are followed correctly and to spot any irregularities.

**Provisioning:** In a cloud environment, Chef automates the setup and configuration of servers or instances.

**Configuration Drift Management:** Takes care of any deviations from the established policies to guarantee that every node maintains the appropriate configuration state.

**Nodes/Instances:** These are the real servers or instances that Chef manages and configures. Chef applies and maintains their configuration through the use of recipes.

This Application in Figure 6 shows how Chef Tools are used in real-time inside a contemporary infrastructure. It demonstrates how Chef Server interacts with cloud providers for provisioning, centralizes configuration management, and interfaces with CI/CD pipelines for automated deployments. The graphic also shows how logging and monitoring technologies can control drift and ensure configuration compliance. Chef configures nodes or instances to provide automated and standardized infrastructure management.

## XI. Conclusion

This paper underscores the transformative impact of Chef in Software Configuration Management (SCM), highlighting its ability to enhance consistency, reliability, and efficiency in managing complex infrastructures. Chef's automation capabilities, combined with its integration into Continuous Integration (CI) and Continuous Deployment (CD) pipelines, offer significant benefits, including reduced configuration drift, improved scalability, and enhanced operational efficiency. Despite challenges such as managing dependencies, ensuring security, and navigating the learning curve, the proposed solutions and best practices provide actionable strategies to address these issues. Case studies demonstrate successful implementations of Chef, reaffirming its value in modern software development environments. Overall, Chef proves to be a powerful tool for effective SCM, offering substantial improvements in software robustness and agility when its challenges are strategically managed.

## XII. References

- [1] Waldemar Hummer, Florian Rosenberg, Fábio Oliveira, Tamar Eilam "Automated Testing of Chef Automation Scripts" 09 December 2013, <https://dl.acm.org/doi/10.1145/2541614.2541632>.
- [2] Christof Ebert, Gorka Gallardo, Josune Hernantes, and Nicolas Serrano, "DevOps", 25 April 2016, <https://ieeexplore.ieee.org/document/7458761>.
- [3] Stefano Dalla Palma, Dario DiNuccia, Fabio Palombab, Damian Andrew Tamburric, "The Journal of Systems & Software" 3 April 2020, <https://www.sciencedirect.com/journal/journal-of-systems-and-software>.
- [4] Stefan Meyer, Philip Healy, Theo Lynn, John Morrison, "Quality Assurance for Open Source Software Configuration Management" 2013, <https://ieeexplore.ieee.org/document/6821183>.
- [5] AKOND RAHMAN, MD RAYHANUR RAHMAN, CHRISPARNIN, LAURIEWILLIAMS, "Security Smells in Ansible and Chef Scripts" January 2021, <https://dl.acm.org/doi/10.1145/3408897>.
- [6] E. Luchian, C. Filip, A. B. Rus, I. -A. Ivanciu and V. Dobrota, "Automation of the infrastructure and services for an openstack deployment using chef tool," 2016 15th RoEduNet Conference: Networking in Education and Research, Bucharest, Romania.
- [7] Arturs Bartusevics, Leonids Novickis, "Models for Implementation of Software Configuration Management", 2015, <https://www.sciencedirect.com/science/article/pii/S1877050914015701>.
- [8] Jacky Estublier, David Leblang, André van der Hoek, Reidar Conradi, Geoffrey Clemm, Walter Tichy, and Darcy Wiborg-Weber. 2005. "Impact of software engineering research on the practice of software configuration management". ACM Trans. Softw. Eng. Methodol. 14, 4 (October 2005), 383–430. <https://doi.org/10.1145/1101815.1101817>.
- [9] Roman Kostromin, "Survey of software configuration management tools of nodes in heterogeneous distributed computing environment", 2020, [DOI:10.47350/ICCS-DE.2020.15](https://doi.org/10.47350/ICCS-DE.2020.15).
- [10] Hong Mei, Lu Zhang, Fuqing Yang, "A Software Configuration Management Model for Supporting Component-Based Software Development", ACM SIGSOFT, March 2001, <https://dl.acm.org/doi/abs/10.1145/505776.505790>.