

Hyperscale Web Company Moves Closer to Automated, Scalable, Repeatable Cluster Software Deployment



CASE STUDY

Challenge

- Locate alternative software deployment solution to semi-automated tools heavily reliant on administrator expertise and manual scripting for deployment and management of Hadoop clusters with thousands of nodes
- Accelerate cluster deployment (and upgrades) from weeks (or months) to under an hour
- Enable automated, scalable, repeatable processes for deployment, monitoring, and fixes

Solution and Benefits

Rocks+ deployment and management software with the Hadoop software Roll

- Centralized, fully-automated deployment of hundreds to thousands of application clusters on heterogeneous devices
- Central point of management enables customized configurations, fixes, patches, and changes to individual nodes through XML node files
- BitTorrent-like installer simplifies large-scale deployments
- Easy-to-use, simplified command line interface (CLI) commands
- Reusable configurations for fast deployment of new Hadoop clusters

Growing its user base tenfold would stress the IT environment of any company. But when that company is hosting over a billion transactions, requiring more than 50 billion database requests every day, the agility and stability of the IT environment underlie the company's very viability. In this case, the stress was compounded by having multiple data centers – each running Hadoop clusters with over a thousand nodes – that relied on a single administrator skilled in the use of DevOps tools. With his departure, the company recognized that their reliance on a semi-automated solution that is dependent on hand crafted software scripts for each cluster jeopardized the speed and scalability of future software deployment and management. StackIQ ran a proof of concept trial to demonstrate to the company how Rocks+ with the Hadoop Roll could configure 110 heterogeneous compute nodes within 30 minutes (and scale beyond this with a very small increase in deployment time thanks to Avalanche, the Rocks+ parallel installer).

Challenge: Replace Complex, Time-Intensive Cluster Software Deployment and Management

To support massive growth, the web company's data centers deployed the Apache Hadoop software framework that supports distributed applications with large data sets. Clusters of servers run Hadoop on standard racks with processors running commodity x86 CPUs. The clusters are in multiple data centers and each cluster has up to 3,000 nodes. A separate test cluster is set up in the company's development center, running the latest versions of Hadoop.

Using a mixture of commodity x86 servers from multiple Tier 1 and Tier 2 hardware vendors, the company required a software deployment and management solution for clusters running the high-demand analytics applications underlying their web platform.

Previously, provisioning and managing server clusters in "Big Infrastructure" environments, where applications rely on very large data sets and high volumes of processing, was a time-consuming endeavor. Administrators relied on manual processes based on writing software scripts to provision and manage the images on each server. Later, semi-automated tools such as Puppet were created to package scripts and deployment and configuration instructions for specific software. However these tools still often depend on additional manual scripting, which requires a level of technical proficiency among server administrators and slows down deployment, upgrades, and fixes.

The company had relied on one Puppet expert to guide the deployment and management of all Hadoop clusters. For each cluster, a team of administrators would deploy an operating system on bare metal. Custom Puppet scripts would subsequently be used to configure Hadoop on each cluster. For new deployments, including newer versions of Hadoop or for new



CASE STUDY

hardware, the Puppet scripts would have to be updated. The OS and scripts would have to be redeployed from scratch. This could take weeks to months, even when a hardware vendor deployed an OS image on the servers beforehand. Without the Puppet expert, there was an immediate knowledge void and deployment of new clusters was curtailed until another solution could be found.

The search for a solution was vital as the company was continuing to grow and was considering purchasing thousands of new servers from a Tier 1 hardware vendor, who did not have a software solution that met their needs.

Solution: A Centralized, Automated, Scalable, Easy-to-use Alternative

In search of an alternative to Puppet, the company launched an internal development effort to create a home-grown server provisioning and management system. Then, at the Red Hat and JBoss World Summit in Boston in 2011—which focused on open source cloud computing and platform, virtualization, middleware, and management technologies—representatives from the web company's IT department met the developers of Rocks+. They learned that Rocks+ is a comprehensive software suite designed to automate the deployment and management of Big Data, analytics, or high performance computing environments using the Linux operating system, cluster management middleware, libraries, compilers, and monitoring tools.

When they saw a Rocks+ demo and learned that it can deploy 1,000-node Hadoop clusters within an hour, they were intrigued, as they were already on the verge of expanding their Hadoop footprint and found out that their hardware vendor of choice was already an authorized Rocks+ reseller, lending additional credibility.

Rocks+ Hadoop Proof of Concept

In the summer of 2011, StackIQ designed a proof of concept for deployment of a Hadoop cluster with some of the company's custom software on 110 heterogeneous compute nodes. The software facilitates ad placement on their web site.

The proof of concept was conducted at one of the company's large collocation facilities. Rocks+ was installed on the hardware along with the StackIQ Hadoop Roll with the company's custom components integrated. It took approximately 40 minutes for installation of the frontend, which is responsible for understanding the network topology, gathering data on the number of disk drives on back end machines, and creating a repository to conduct the Hadoop installation. It then took 20 minutes for deployment of the Hadoop infrastructure on the 110 back end compute nodes (from bare metal), more than meeting the company's challenge for deployment of the nodes within half an hour.

By identifying the MAC addresses of the nodes, Rocks+ was able to quickly transfer that information to the Rocks+ database, identifying the nodes to be included in the cluster. The other nodes in neighboring racks continued to handle other workloads unaffected while Rocks+ deployed the Hadoop software to the designated infrastructure. Expanding the cluster to other nodes is as simple as adding their MAC addresses.

Deciding Between a Commercial Product and a Home-grown Development

The company's IT team particularly liked the central point of management in Rocks+ and how automated the entire deployment process was. They were impressed with how scaling Hadoop clusters across thousands of nodes could be done quickly with the Rocks+ Avalanche installer and its BitTorrent-like features.



The company was also impressed with how reusable Rocks+ and the Hadoop Roll is. Once it is configured according to the company's specifications, it can be used to easily and quickly deploy new nodes and new clusters. Everything is portable and reusable for multiple data centers. Additionally, the simplicity of Rocks+ was very attractive. CLI commands such as "rocks add roll" or "rocks start hadoop" makes it intuitive and easy to understand and use, not requiring coding skills and years of server development or operations experience.

For the web company, the Hadoop infrastructure running analytics applications is considered a profit center, contributing directly to transactions on the site. Therefore a server deployment and management solution is a priority. They must now choose between an unproven, home-grown solution, which has yet to be completely developed, and the fully-automated Rocks+, which is based on years of development and global use of the open source Rocks cluster provisioning and management solution within high performance computing environments that is now primed for Big Infrastructure enterprise data centers.

More Information

To find out more about StackIQ and Rocks+ go to www.stackIQ.com



4225 Executive Square #1000
La Jolla, CA 92037
858.380.2020
888.400.3966 fax
info@stackiq.com