

Optimizing Data Centers for Big Infrastructure Applications



Whether you need to analyze large data sets or deploy a cloud, building big infrastructure is a big job. This paper discusses how Rocks+ makes the job easier.

Contents

- The Era of Big Infrastructure Applications 2
- Challenges in Deploying Big Infrastructure Software 3
- Simpler and More Powerful Automation for Big Infrastructure 4
- Preconfigured Software Modules ("Rolls") 6
- Summary 7
- Key Benefits of Rocks+ 7
- About StackIQ 8

Abstract.

Many application workloads today have grown to encompass hundreds or thousands of servers and massive amounts of data. Pioneers of these applications among research organizations and businesses like Google and Facebook have developed proprietary tools and techniques to manage their compute environments. Server vendors have also created solutions for deploying and managing large homogeneous server farms. But increasingly, open source software and heterogeneous commodity hardware are having a major impact on enterprise applications and are being deployed in the biggest application environments. This white paper provides an overview of the growing wave of large-scale, 'Big Infrastructure' applications and explores the challenges inherent in provisioning, configuring, and deploying software for clusters running massive data workloads. The Rocks+ infrastructure management system for clustered data center and cloud environments from StackIQ is described. Rocks+ features (including centralized management, dynamic one-step provisioning of heterogeneous appliance types, and an open, modular framework) constitute a paradigm shift in Big Infrastructure management techniques. This newer approach to provisioning and management is contrasted with legacy and alternate methods that are more time-intensive, inefficient, and costly.



The Era of Big Infrastructure Applications

A February 2011 IDC report forecasts that data use is expected to grow by as much as 44 times in the next 10 years, amounting to more than 35 zettabytes globally. (This follows an IDC study in 2009 that estimated that at that time the global Internet contained about 500 exabytes or half a zettabyte.) The amount of information managed by enterprise data centers will grow by a factor of 50 in the coming decade. The number of both virtual and physical servers worldwide will grow by a factor of 10. By 2015, cloud computing will encompass nearly 20% of information from less than 2% in 2011.

“Storing this data is cheap, and it can be mined for valuable information.”

Computer science professor Joe Hellerstein of the University of California, Berkeley, has proclaimed this era, “...the ‘Industrial Revolution of Data,’ where the majority of data will be stamped out by machines: software logs, cameras, microphones, RFID readers, wireless sensor networks and so on. These machines generate data a lot faster than people can, and their production rates will grow exponentially with Moore’s Law. Storing this data is cheap,

and it can be mined for valuable information.” Aside from data gathered by machines, billions of PC and smartphone users accessing social media, ecommerce, and communicating with each other generate huge amounts of data in their wake.

Companies from Google and Facebook to General Motors and Walmart are mining this data in search of sales opportunities and trends, using large data sets, the size of which was once the exclusive domain of scientific research, national security, and oil and gas exploration organizations. Today these applications can be run on hundreds or thousands of clustered servers instead of single-purpose storage appliances from Teradata, EMC, and Oracle, or supercomputers from manufacturers such as Cray and IBM. These newer Big Infrastructure environments can be built on ‘bare metal servers’ in company data centers or leased from virtual cloud environments such as Amazon’s Elastic Compute Cloud (EC2). With the availability of commodity servers and open source software, the costs associated with ‘Big Data’ applications and Big Infrastructure have declined, spurring greater investment [Figure 1].

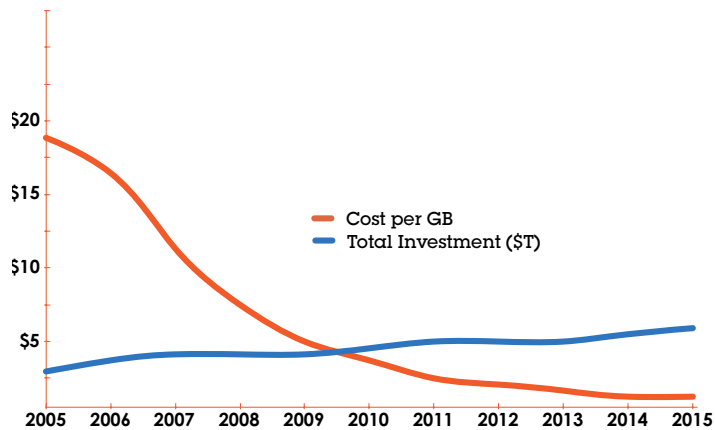


Figure 1. Falling Costs and Rising Investment in Computing Infrastructure
 Source: IDC's Digital Universe Study sponsored by EMC, June 2011

While cloud computing accounts for less than 2% of IT spending today, IDC estimates that by 2015 nearly 20% of data on the Internet will be touched by cloud computing service providers. But the same research report also cautions that effectively handling the rising tide of Big Data applications is already proving challenging to the development and operations personnel within data centers. The study forecasts that, "The skills, experience, and resources to manage all these bits of data will become scarcer and more specialized, requiring a new, flexible, and scalable IT infrastructure..."

Challenges in Deploying Big Infrastructure Software

Over the years, high-performance computing environments handling applications with large data sets have evolved specialized tools to manage their Big Infrastructure applications environments. These environments include powerful, inexpensive servers with x86-based processors, open source operating systems such as Linux, and home grown server provisioning and management software. By contrast, business enterprises have mostly continued to use proprietary server appliances and accompanying software for data warehouse, business intelligence, and other applications from specialized vendors such as Teradata, EMC, and Oracle.

As businesses transition to using commodity hardware and open source software to gain scalability, flexibility, and price-performance benefits, they have discovered that deploying and managing Big Infrastructure server environments is extremely challenging. To date, a variety of legacy management software and newer applications tools are being used. Often separate tools handle deployment, scalability, and management functions and approach them in different ways that add complexity to the environment. Most of the tools manage changes to server configurations and other fixes through software scripts. Making changes to one appliance or several appliances through scripting is a manual, time-intensive, error-prone process. It can leave the administrator unsure if the changes have been implemented throughout the cluster. Attempts at using proprietary software solutions from EMC, Teradata, Oracle, and others has proven very expensive, and these solutions often do not scale to meet customer needs.

A new group of tools that automate the development and operations functions within Big Infrastructure server environments in data centers (hence the collective term 'DevOps tools' for the category) include Chef and Puppet. These tools help businesses automate the configuration and management of infrastructure in cloud computing

environments or internally across the servers in a company's data center. Server configuration and management are accomplished through the use of programming languages like Ruby or proprietary scripting languages. Some of these tool vendors have sought to simplify the use of these management scripting languages for less seasoned systems administrators by packaging software scripts into specific server deployment and configuration instructions (Chef calls them 'recipes').

Rocks+ — Simpler and More Powerful Automation for Big Infrastructure

Originally developed in 2000 by researchers at the San Diego Supercomputer Center at the University of California, San Diego, Rocks is an open source Linux cluster provisioning and management solution. Available free, it was designed to enable end users to easily, quickly, and cost-effectively build, manage, and scale application clusters and grid endpoints. Thousands of environments around the world now use Rocks.

Rocks+ from StackIQ is the commercial edition of open-source Rocks, [Figure 2]. It was developed by the same creators of open source Rocks but designed for the world of Big Infrastructure. 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.

With Rocks+, physical or virtual infrastructure can be quickly provisioned, deployed, monitored, and managed. System administrators manage the entire system using a single pane of glass. New nodes are also configured automatically from bare metal—with a single command—without the need for complex administrator assistance. If a physical or virtual node needs an update, it will be completely re-provisioned by Rocks to ensure it boots into a known-good state. Since Rocks+ places every bit, on every physical and virtual node, administrators have complete control and consistency across the entire infrastructure. Rocks "Rolls" and Rocks "Appliances" give administrators the control they need to define site-specific solutions to meet their enterprise requirements.

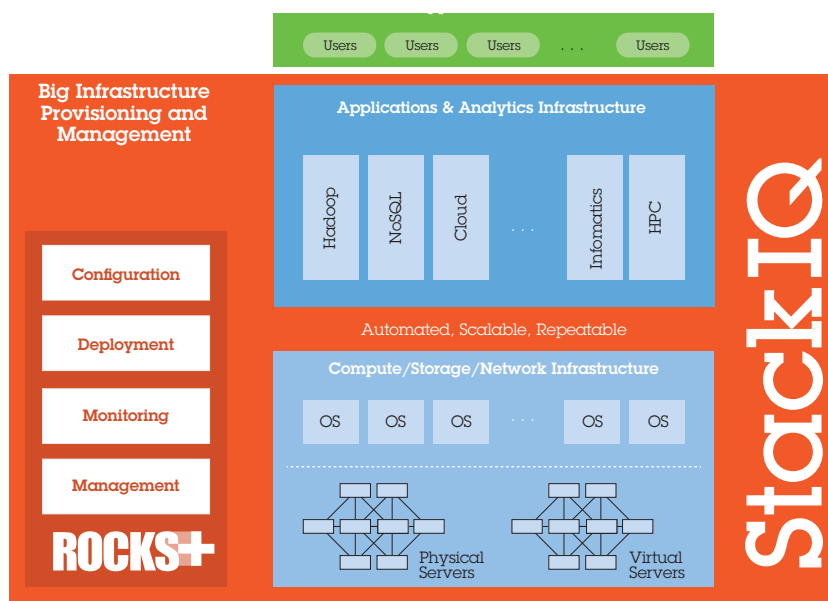
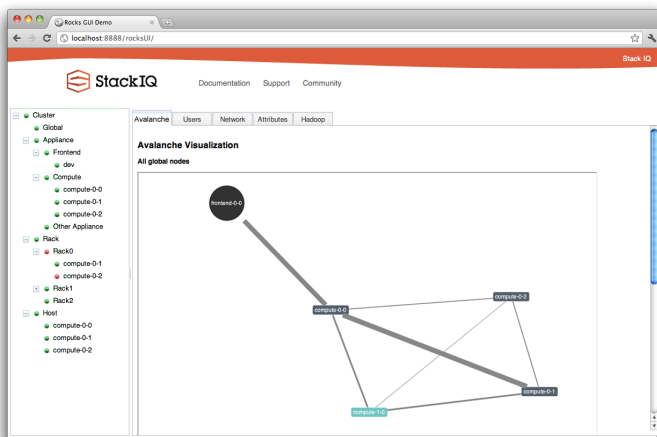


Figure 2. Rocks+ for Physical and Virtual Heterogeneous Server Cluster environments

Another component of Rocks+ designed for rapid provisioning of Big Infrastructure is the Avalanche Installer. This specialized installer avoids the bottleneck that results from delivering packages from a central distribution server to multiple simultaneously-installing nodes. When an installing node

downloads a package using Avalanche, it stores the package locally, then sends a message to a custom BitTorrent-inspired tracker that places the package in a peer-to-peer cache. When a subsequent installing node wants to download the same package, the node sends a message to the tracker, which responds with a list of peers. The installing node will download the package from a peer, rather than burdening the central distribution server. By implementing this peer-to-peer protocol, the Avalanche Installer dramatically reduces the network pressure on the central distribution server and enables large-scale simultaneous node installations.



Rocks+ Avalanche Installer Visualization

As compared to other Big Infrastructure server automation tools and approaches, with Rocks+:

- The focus is not on manual processes or scripts; instead each configured server cluster is deployed from the database
- Changes or fixes to software are made centrally and then redeployed to the entire cluster
- Pre-packaged, automatically configured data center or cloud software stacks called "Rolls" are available to simplify Big Infrastructure deployments on physical or virtual servers
- An automated, accelerated installation process based on an open architecture enables the dynamic provisioning of heterogeneous appliance types across massive data center environments
- Massively scalable deployments are possible with BitTorrent-style package sharing from the Avalanche Installer, which delivers packages from a central distribution server to multiple simultaneously-installing nodes

Preconfigured Software Rolls

Rocks+ users can optionally utilize pre-packaged software modules called Rolls that provide an end-to-end Big Infrastructure software stack while also integrating software components for site-specific requirements.

Rolls are software modules that may be selected and automatically configured in Rocks+ [Figure 3].

Rolls for Rocks+

Rocks+ Hadoop Roll This preconfigured Roll provisions Apache Hadoop, including MapReduce, the Hadoop Distribution File System (HDFS), ZooKeeper, Hbase, and Hive. It is easy to use, highly scalable, and allows for deployment of Hadoop clusters of all shapes and sizes (including heterogeneous hardware support, parallel disk formatting, and multi-distribution support). Typically, the installation and management of a Hadoop cluster requires a long, manual process. The end user or deployment team must install and configure each component of the software stack by hand. The setup time for such systems and the ongoing management can be problematic and time-intensive. Security and reliability may be impacted. As more enterprises process Big Data with Hadoop, the need for a supported and robust platform that is easy and fast to deploy, manage, and operate becomes more essential. The Rocks+ Hadoop Roll completely automates the process.

Additional Big Data Rolls include MongoDB and Cassandra.

Rocks+ Workload Management Rolls The Moab Roll adds the Moab Cluster Suite to Rocks+ and provides an automated configuration for powerful workload

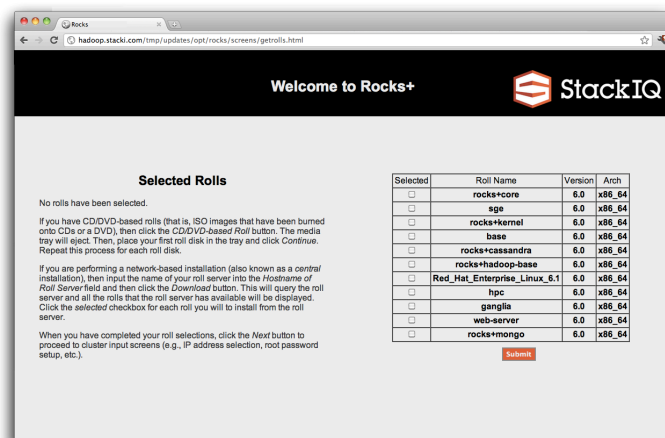


Figure 3. Rocks+ Rolls Selection Screen

and resource management that integrates scheduling monitoring, and reporting for cluster workloads. The Univa Grid Engine Roll for Rocks+ provides a number of great features, including priority policies and dynamic resource management. The PBS Pro Roll from Altair from Rocks+ combines intelligent scheduling algorithms with application-aware user portals and powerful visualization tools.

Rocks+ High Performance Computing and Analytics

Rolls Rocks+ originated in the world of high performance computing (HPC) and StackIQ has created Rolls for large scientific, research, and analytics applications. The HPC and Open Fabrics Enterprise Distribution (OFED) Rolls were designed as core components of Rocks+. They are continually updated to support the latest stable OpenMPI and OFED distributions and includes drivers and firmware for Mellanox InfiniBand host channel adapters (HCAs) and Intel Internet Wide Area RDMA (Remote Direct Memory Access) Protocol (iWarp), including fabric management tools, and multiple Message Passing Interface (MPI) environments.



Key Benefits of Rocks+

- Faster time to deployment
- Automated, consistent deployment
- Simplified operation that can be quickly learned without systems administration experience
- Reduced downtime due to configuration errors
- Choice of stack components to tailor solutions to specific requirements

The Solarflare Roll for Rocks+ includes a preconfigured Open-Load® network stack. It dramatically reduces application-to-application message latency, which is an essential feature for high-frequency trading and scientific analytics applications.

Rocks+ for Public Cloud Environments The same powerful tools that administrators have been using to manage hyperscale physical infrastructure for years are now available to build robust, enterprise-class cloud infrastructure. Rocks+ is available on Amazon's EC2, enabling customers that choose to lease virtual cloud resources to quickly design, configure, and manage Big Infrastructure without having to build out a physical data center. Visit <http://www.StackIQ.com/ec2> to learn more.

Summary

Big Data applications are proliferating. The move to open source software and commodity hardware within the Big Infrastructure environments running these applications have left a void in cost-effective, comprehensive deployment and management tools. Rocks+ is the enterprise version of Rocks, an open-source management tool originally developed to set up application clusters for Big Data applications in research and academia. Rocks+ provides an array of benefits that simplify Big Infrastructure deployment and management within clustered data center environments and public cloud infrastructure such as Amazon's Elastic Compute Cloud 2 (EC2). Its modular design and unique approach to server cluster automation lower operational costs and increase efficiency in the deployment and management of Big Infrastructure.

About StackIQ

StackIQ is a leading provider of Big Infrastructure management software for clusters and clouds. Based on open-source Rocks cluster software, StackIQ's Rocks+ product simplifies the deployment and management of highly scalable systems. StackIQ is based in La Jolla, California, adjacent to the University of California, San Diego, where the open-source Rocks Group was founded. Rocks+ includes software developed by the Rocks Cluster Group at the San Diego Supercomputer Center at the University of California, San Diego, and its contributors. Rocks® is a registered trademark of the Regents of the University of California.



4225 Executive Square
Suite 1000
La Jolla, CA 92037
858.380.2020
info@stackIQ.com

StackIQ and the StackIQ Logo are trademarks of StackIQ, Inc. in the U.S. and other countries. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between StackIQ and any other company.