EMA Impact Brief: Cisco Introduces Workload Automation for Big Data

Introduction

In October 2012, Cisco Systems announced version 6.1 of Cisco Tidal Enterprise Scheduler (Cisco TES). Cisco TES is a *Value Leader* in Enterprise Management Associates' (EMA) 2012 Workload Automation Radar Report and received the EMA award for "Best Implementation and Configuration." Version 6.1 will be available for general purchase in November 2012 and includes new capabilities in four key areas:

1. **Big Data:** Hadoop adapter for scheduling Big Data workloads
2. **Public Cloud:** Support for Amazon EC2 and S3 storage
3. **Self-service:** Web portal for business users
4. **Mobile access:** iPhone app for operations management

This EMA Impact Brief examines the significance of these new capabilities for existing and new Cisco TES customers.

Cisco TES 6.1: The New Features under the Microscope

Workload automation (WLA) has come a long way, since the days of static mainframe job scheduling. Today WLA constitutes the backbone of the datacenter, aligning numerous enterprise applications along the lines of business processes. Cloud, DevOps, and Big Data are today's most talked about IT topics, all of which are heavily reliant on WLA. To efficiently support Cloud, DevOps, and Big Data projects, WLA software must have the capability to automate and orchestrate existing enterprise systems and data silos. In line with this requirement, Cisco TES 6.1 adds four new capabilities: Hadoop Adapter, Amazon EC2/S3 integration, Self-service Portal, and iPhone App. These four capabilities facilitate and simplify the end-to-end management of business processes that involve cloud resource provisioning and management, Big Data processing and analysis, and narrowing the divide between developers and IT operations.

Hadoop Adapter

The new Hadoop Adapter, available for Cisco TES 6.1, facilitates the end-to-end management of complex Hadoop workflows from within the familiar Cisco TES user interface. Customers can now include complex Big Data job streams with their greater workload automation workflows, without leaving Cisco TES. Centrally controlling workflows in their entirety is essential, when managing business services as a whole, in an SLA centric manner (see Figure 1).

Customers can now include complex Big Data job streams with their greater workload automation workflows, without leaving Cisco TES.

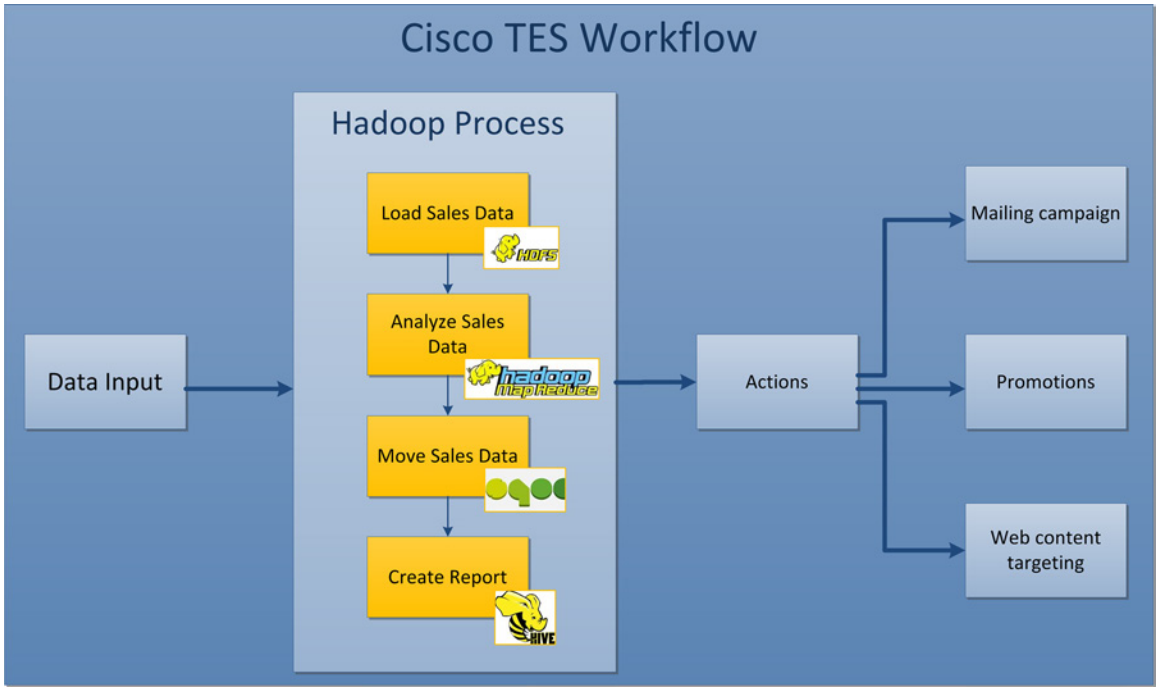

Figure 1 - Hadoop as Part of a Cisco TES managed Workflow

The Cisco TES adapter for Hadoop includes API integrations for:

- **HDFS:** The first step of almost any Hadoop-based analysis is the import of data into the Hadoop Distributed File System (HDFS). This can be any type of structured and unstructured data from almost any data source. Web clicks, search data, chat records, file system data, emails, surveillance cam videos, sensor data, loyalty card use, and ERP transactions are only a few examples of the data that can be imported into HDFS via the Cisco TES Data Mover.
- **Map Reduce:** Map Reduce is the Hadoop programming framework where unstructured data is processed and business-relevant patterns identified. One of the most famous examples of this type of analysis is a major department store's ability to determine when a specific customer is pregnant, based on her purchase history, compared to the purchase behavior of customers who signed up for the store's baby registry[1].
- **Sqoop:** Sqoop is the Hadoop ETL (Extract Transfer Load) tool for moving data from SQL database sources into HDFS and vice versa. Sqoop can also be used to move data from HDFS to Hive, for further analysis and reporting.
- **Hive:** Hive makes Hadoop and other data available for SQL-like queries, providing "structured" data tables within Hadoop. Making Big Data available to SQL savvy IT and business staff significantly broadens the use the organization will draw from strategy-relevant data that has gone mainly untapped for decades.

Taking full advantage of Big Data constitutes a tremendous competitive advantage for many organizations. Cisco TES 6.1 now facilitates the central end-to-end management of the entire process: data acquisition, analysis, manipulation, and reporting. Based on the outcome of the analysis, Cisco TES triggers further workflows that result in immediate actions, exploiting what was learned from the Big Data analysis.

1 For an excellent overview of the power of Big Data, please review the following article: http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

Amazon EC2 and S3 Integration

Cisco TES 6.1 is now able to create Amazon EC2/S3 compute and storage resources as parts of a greater workflow. Cisco TES 6.1 workflows can include conditions for when specific jobs or workflows should be executed on Amazon resources and define how the required EC2/S3 resources are provisioned and how the corresponding data is moved to the Amazon environment.

Centrally managing and moving workloads to and from Amazon EC2/S3, without manual intervention, tremendously increases scalability and elasticity of the corporate workload automation capabilities. Controlling Amazon EC2/S3 management via Cisco TES also allows organizations to benefit from Cisco TES security management, facilitating authorization, authentication, data and application security.

Centrally managing and moving workloads to and from Amazon EC2/S3, without manual intervention, tremendously increases scalability and elasticity of the corporate workload automation capabilities.

Self-Service Portal

The Cisco TES 6.1 self-service portal is based on newScale technology (Cisco acquired newScale in 2011) rebranded as the Cisco Workplace Portal. By using this self-service portal with Cisco TES, business users can monitor performance and health of relevant workloads via a web browser and without having to make phone calls to enterprise IT. Business users are also provided with a number of basic job control capabilities, for launching, restarting, cancelling, pausing, and resuming jobs. The Cisco TES 6.1 self-service portal includes granular user management and keeps a full audit track of each business user's action.

iPhone App

The ability to access and manage workloads from anywhere is essential in today's business service focused organization, where productivity has to be ensured around the clock. The new Cisco TES iPhone app enables workload engineers and administrators to monitor and control workloads managed by multiple Cisco TES instances. The iPhone app (available on the Apple App Store as Cisco Enterprise Scheduler for iOS) provides easy access to performance and health data, shows an overview of scheduled jobs, and facilitates the management of job queues, system logs and alerts. Access to the iPhone app is managed through the existing Cisco TES security model.

EMA Perspective

The list of additional capabilities that can be found in Cisco TES 6.1 reads like an enumeration of today's hottest topics in enterprise IT. In other words, the 6.1 release adds tremendous customer value to an already robust and EMA award winning workload automation solution. From a strategic point of view, Cisco TES 6.1 achieves the following two central goals:

a) **Improved end-to-end management:** Cisco TES now allows customers to manage Hadoop and Amazon EC2/S3 components as part of their greater job workflows. Centrally managing Big Data and public cloud capabilities further extends the reach of Cisco TES and is a precondition for customers to easily tie workloads to SLAs.

b) **Business accessibility of WLA:** The self-service portal and the iPhone app both make Cisco TES easier to manage and more transparent to the business. Rescuing WLA from the hands of the "data center mafia" and tying it closer to the business is one of today's most significant goals of enterprise IT. EMA commends Cisco for making an important step into this direction.

Cisco TES 6.1 is a significant release, allowing customers to further bridge today's gap between enterprise IT capabilities and business requirements. Customers can leverage their existing Cisco TES investment, to manage a bigger share of these requirements. Especially, as Hadoop moves away from experimental deployments and into the center of many organizations' business strategy, it is essential to incorporate Big Data into the corporate processes and workflows. Managing Hadoop from within Cisco TES applies corporate SLAs and policies to Big Data.

Customers can leverage their existing Cisco TES investment, to manage a bigger share of these requirements.

Allowing customers to more efficiently use Amazon EC2/S3, incorporate Hadoop into their existing workflows, and reduce support cost by offering self-service and a mobile app, are new Cisco TES 6.1 capabilities illustrating Cisco's ability to deliver essential cutting edge functionality.