Transforming Management for Modern Scale-Out Infrastructure

"Software is eating the world."
– Marc Andreessen, WSJ 2011

That famous quote from venture capitalist and Netscape founder Marc Andreessen was true then, and is even more true today. More powerful, complex software is driving the transformation of almost every industry – hospitality (AirBnB), retail (Amazon), and transportation (Uber) – it's even driving the cars themselves! And the most significant impact of software becoming an overwhelming force is the change it has brought to the very nature of computing.

The increasingly complex nature of such innovative software solutions required more IT resources which in turn accentuated the need for more efficient IT solutions. The simple one-app/one-server model had to go. It had survived far too long with absurdly low utilization rates.

Initially, server virtualization, with multiple instances of virtual machines, revolutionized the way companies implemented data centers, allowing them to consolidate computing resources and gain server utilization efficiencies - along with the added benefits of energy and footprint savings.

But that was only the first piece of a much bigger IT transformation puzzle. Virtualization birthed massive horizontal scale-out computing, cloud computing, and all the flavors of "as a Service" (SaaS, PaaS, etc.). Virtualized compute resources were eventually accompanied by corresponding virtual storage solutions (e.g. VSAN, Storage Spaces) and inevitably virtual networking (NFV).

Virtualizing all aspects of IT has resulted in truly software-defined data centers (SDDC), where all the resources needed for a particular workload can be (virtually) mustered from anywhere in the physical datacenter - across machines, operating systems, vendor brands – transparently to the application. The applications, users, and business owners now operate in a world that is simpler, faster and more efficient for them. Thanks to these innovations they can readily scale up to thousands of nodes.

Management at scale: Bigger is better…but also harder to manage

But now, transformation is needed in the management of large scale IT environments - not just hyperscale data centers, but growing scale customers as well. Imagine the challenges for companies like AirBnB and Uber. Their innovative software has opened a floodgate of demand from an exponentially growing customer base - worldwide. These companies must now scale deployment and resource management at a fast enough rate to stay up with their business growth. This exact same challenge faces many large scale data centers today that are implementing cloud computing – from traditional service providers and carriers to innovative companies building new cloud-native apps, often termed Platform 3 (P3) computing. To justify IT outlays these companies need to "get to live" or "get to money" more rapidly than has been available in the past. They need a simpler, faster way to deploy and manage massively-scaled IT environments.

This paper describes how Dell EMC and Intel® have collaborated using Intel® Rack Scale Design and OpenStack to allow for more simplified and automated management of scale-out resources on the latest rack scale solution from Dell EMC – the DSS 9000. Dell EMC has integrated the software components to provide an easier and seamless "kickstart" for workload deployment. This paper details the various software components involved, how they interact, and the new orchestration capabilities they bring to rack scale data centers. These components include:

- **OpenStack** – a leading open-source cloud computing platform
- **A resource orchestration tool for OpenStack** - For example MAAS, OpenStack Fuel, etc.

This paper describes how Dell EMC® and Intel® have collaborated using Rack Scale Design and OpenStack to allow for more simplified and automated management of scale-out resources on the latest rack scale solution from Dell EMC – the DSS 9000.

- **Redfish** – an open, standard API for infrastructure management published by DMTF
- **Intel® Rack Scale Design (RSD)** – logical architecture that allows efficient pooling and utilization of disaggregated compute, storage and network resources
- **Dell EMC kickstart integration components** – a set of tools to enable Intel® RSD for OpenStack clouds
- **Dell EMC DSS 9000** – a highly flexible cloud-ready rack infrastructure solution

These components and their interaction are described in the following sections. The last section of the paper walks through a simplified example of cloud management that shows how these components can easily discover and deploy resources at rack scale.

OpenStack: Open, flexible, innovative

OpenStack, the open-source software platform for cloud computing, is increasingly seen as a solution for massively scale out private cloud/hybrid cloud environments. There are a number of reasons for this:

- It is an open industry standard with more than 60 leading companies contributing and participating - and with new OpenStack clouds regularly being introduced worldwide, there is no lock-in to a single proprietary vendor.
- It allows data centers to abstract functions from the hardware layer in a modular fashion and assign IT resources (server, storage, and networking) to them independently
- Its modular design allows it to readily integrate both new and third party technologies as well as legacy infrastructure – allowing easy flexibility to address specific business needs.
- It has proven its reliability and efficiency on some of the biggest public and private clouds in the world.
- It derives accelerated innovation benefits from the efforts of a worldwide developer community.

OpenStack was the initial choice of cloud software platform for the DSS 9000 kickstart for two reasons. First, many customers that are implementing OpenStack have voiced the need for better or faster deployment. Secondly, the ability to build on open source innovation and integrate with open APIs enabled Dell EMC to extend the management capabilities and build kickstart integration software for extreme rack scale environments.

RedFish: An open API for hardware management

Redfish is a single, open API intended for the management of all hardware - compute, storage and networking. It is being designed and implemented over time. It replaces IPMI-over-LAN and is a modern, scalable, secure management API.

Redfish is the key to disaggregating hardware resources, is a foundational component of Intel® Rack Scale Design, and is central to the Dell EMC transformation of extreme scale computing.

It enables seamless, programmable insight, and control of resources. It provides access for client applications and browser based GUIs, allowing for the automation of hyperscale architectures.

Intel® Rack Scale Design: Dynamically assemble infrastructure

Intel® Rack Scale Design is another piece of the puzzle for transforming hyperscale environments. It is a software-defined architecture that allows for disaggregation of compute, storage, and network hardware resources, and introduces the ability to more efficiently pool and utilize these resources. It is built specifically to address the needs of the largest hyperscale customers –carriers and service providers – who want the ability to treat hardware resources as a service.

RSD is a set of APIs and software – with Intel® Pod Manager (PODM) at its heart. With Pod Manager, administrators don't need to know the physical location of infrastructure components. Instead, when an administrator requests resources for a workload, they are located (from anywhere in the rack or "pod") through an automated discovery process, and those resources are allocated to the workload dynamically.

Figure 1: Architecture for OpenStack/RSD/Kickstart Extreme Scale Management

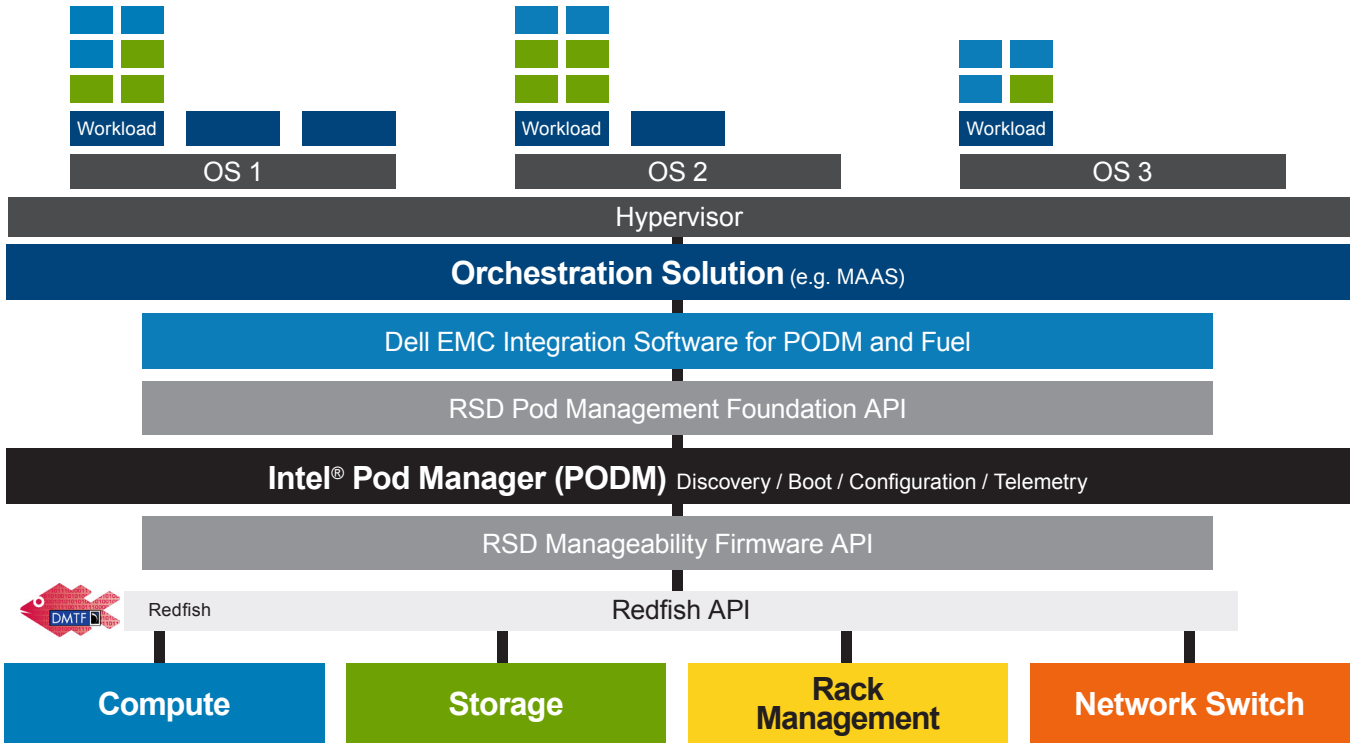

Intel® Pod Manager acts as a resource manager for disaggregated pools of hardware resources by communicating using standard Redfish APIs to hardware-aware software within the DSS 9000 called the Pooled System Management Engine (PSME). PODM provides configuration management of the pools of compute, storage and networking resources to administrators and applications through a vendor-agnostic orchestration (RESTful) API. (See Figure 2).

Pod Manager provides:

- Asset and location discovery (find the assets, create an inventory)
- Management of disaggregated resources (enable allocation, etc.)
- System support for dynamically assembled infrastructure

Figure 2: Pod Manager enables the pooling of disaggregated resources from across one or more racks and allows administrators to assemble them as needed and assign them to specific workloads

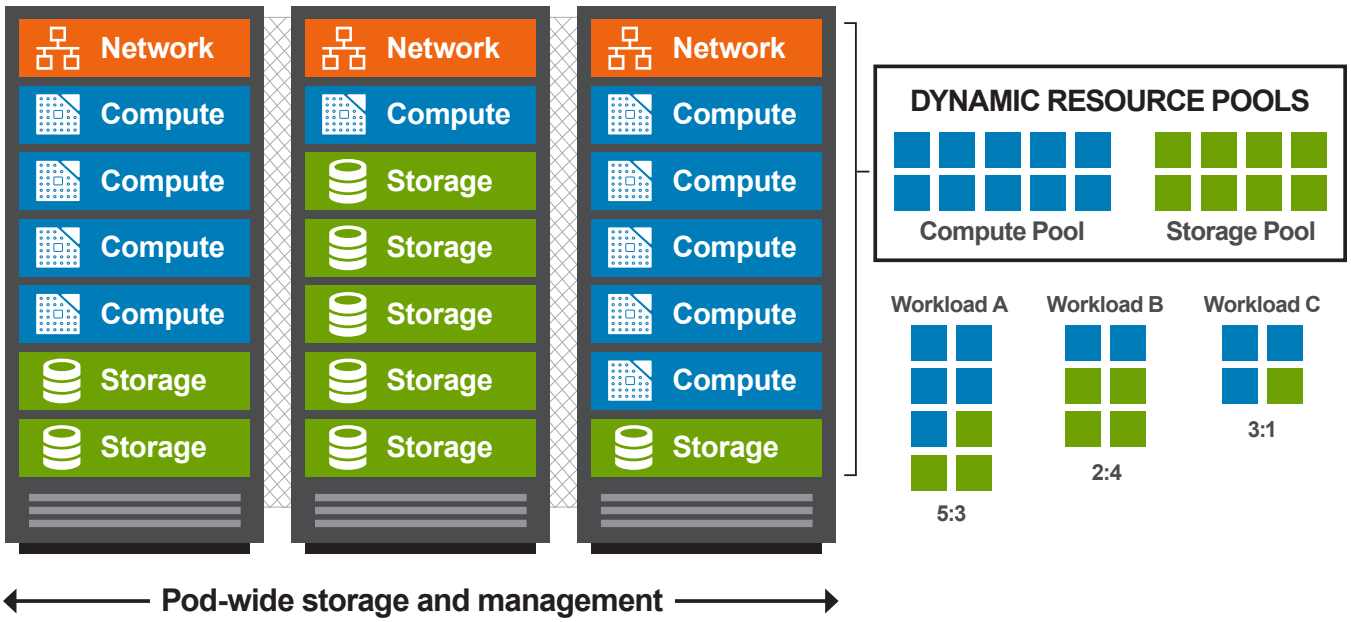

While this paper describes a Dell EMC-enabled OpenStack cloud implementation, Intel® Rack Scale Design can also support other cloud frameworks (e.g., Microsoft™ or VMware™) allowing administrators to manage solutions running several clouds from different vendors – all running on the same underlying hardware, and all with the ability to define, allocate and re-allocate resources quickly and independently, as necessary.

NOTE: A "pod" is defined as the logical extended scaled environment – it may be all the resources in a single rack, or it may include all the resources across many racks.

Orchestration Tools: Cloud Control

OpenStack orchestration tools provide intuitive, Web-driven interfaces for deployment and management of the OpenStack platform. (See Figure 3) Note that a CLI interface is also available.

Figure 3: Web UI-driven interface for deployment and management of the OpenStack platform.

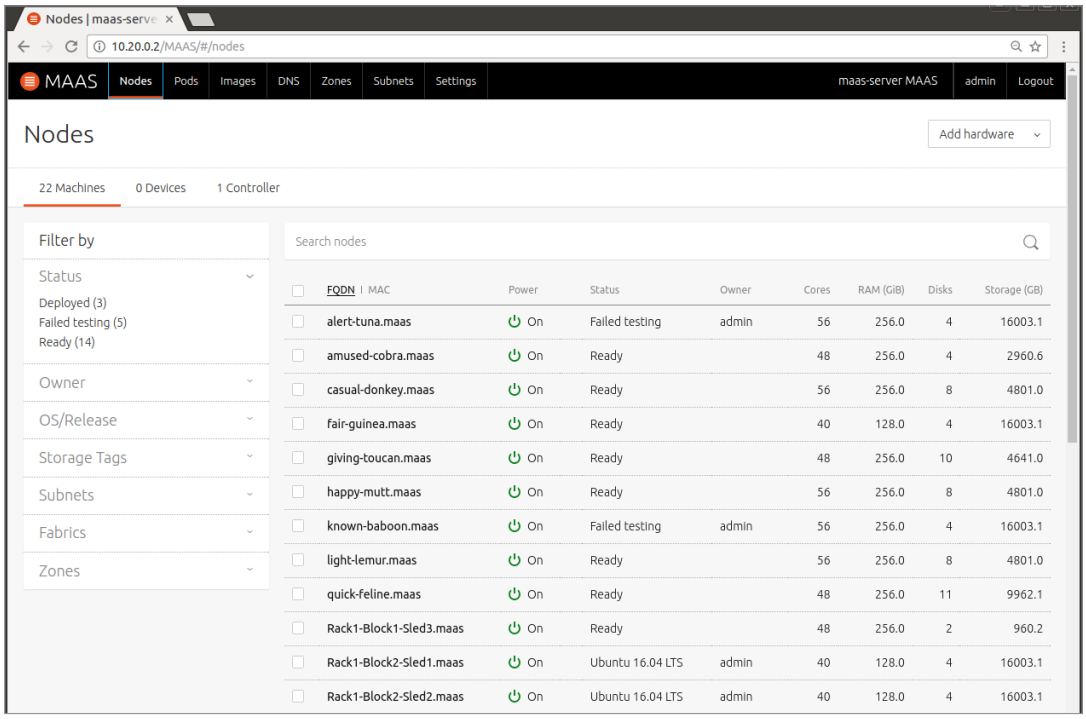

These tools are designed to address easier, more efficient ways to manage cloud resources. They simplify the time-consuming, complex, and error-prone process of deploying and managing a large number of different OpenStack platforms/clusters/nodes in a highly scalable environment.

Key features of OpenStack orchestration tools are:

- Automatic discovery of bare metal and virtual nodes
- Enablement of roles and resources assignments across the cloud
- Validation of network configuration

But additional software is needed for orchestration tools to work seamlessly with Pod Manager, so Dell EMC provides the kickstart integration software components to make this transparent to the user.

Dell EMC kickstart integration software: Helping deploy scale-out workloads quickly and easily

While orchestration tools have the ability to recognize disaggregated resources and provide resource management capabilities, there is still a need to associate the workloads/resources with specific hardware.

As stated previously, Dell EMC kickstart integration software components work with Intel® RSD Pod Manager (PODM) and orchestration tools to speed the deployment and management of OpenStack clouds using Dell EMC DSS 9000 infrastructure. (See Figure 4)

This software allows administrators to request specific hardware resources for workloads, have them automatically allocated (by PODM and the orchestration tool), and then to specifically configure them (set BIOS, RAID, firmware), before deploying them in an OpenStack environment.

Figure 4: Dell EMC supplies integration software to make the interaction of Intel® RSD and OpenStack software transparent.

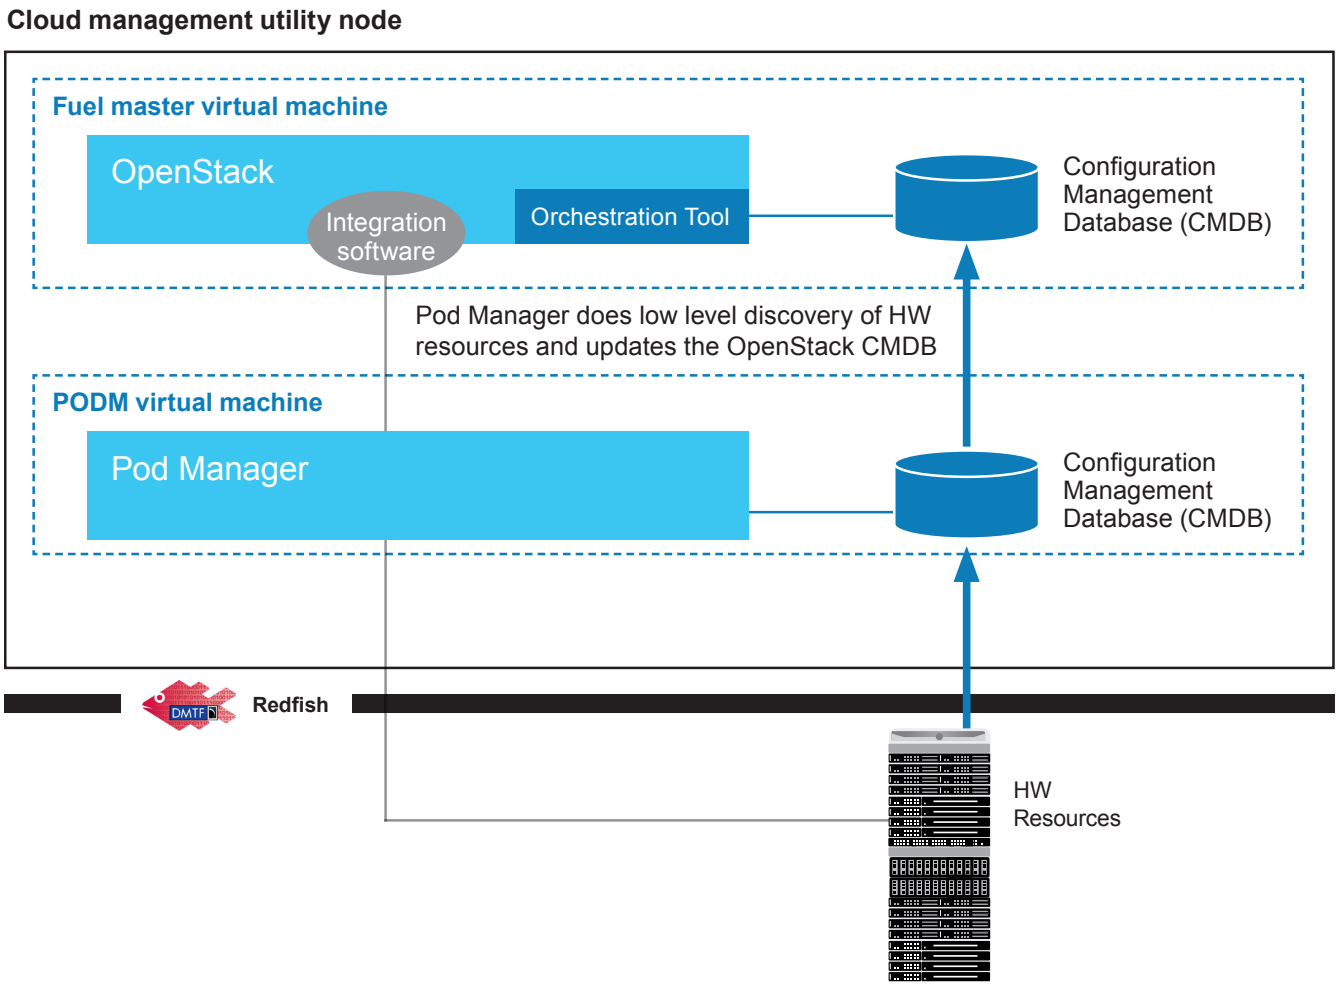

The software can also list a full, detailed physical inventory of a rack or pod – including CPUs, cores, memory, hard drives, SSDs, etc.

One (simplified) way to think about this is that this software takes generic orchestration requests (allocate CPU, storage, etc.) and generates specific DSS 9000 hardware responses. You can create multiple, resourced cluster workloads very quickly, each with different, specific characteristics – and all independent of locale.

There are essentially five Dell EMC kickstart components:

- a PODM virtual box VM
- PODM tools (to be installed on the utility node that drives deployment)
- kickstart integration software to run on the orchestration master node
- a customized bootstrap image for the orchestration tool to support the DSS 9000
- firmware/BIOS update tools (to apply firmware and BIOS configuration to multiple nodes)

The *Dell EMC OpenStack Kickstart Guide for the DSS 9000* has more details on the behavior and use of these components. Contact a Dell EMC sales representative to obtain a copy.

DSS 9000: Designed for rack scale deployment

The Dell EMC DSS 9000 is a rack-level solution that incorporates insights from proven hyperscale design – for example, full disaggregation of hardware resources - to specifically address the unique computing and IT management needs of scale companies, particular telecom carrier companies and Cloud Service Providers (CSPs).

Each DSS 9000 rack allows combinations of both half and full-width server sleds and full-width storage sleds, and incorporates shared rack-scale power, cooling and networking to provide maximum configuration flexibility.

The combination of resources are managed at the rack-level, through the Rack Manager, a hardware controller enabled with APIs that conform to the Redfish API and the Redfish extension for PODM. The Rack Manager provides a single point of control across the rack.

The DSS 9000 ships with all hardware pre-integrated, ready to be allocated, provisioned and deployed as a fully functional cloud platform. With a server to act as a utility node (e.g., a Dell EMC PowerEdge R430) and a Dell EMC S4048-ON switch as the Top of Rack switch, an administrator can use the software described above to rapidly deploy an OpenStack cloud across all the resources on the DSS 9000.

Figure 5: An example configuration DSS 9000

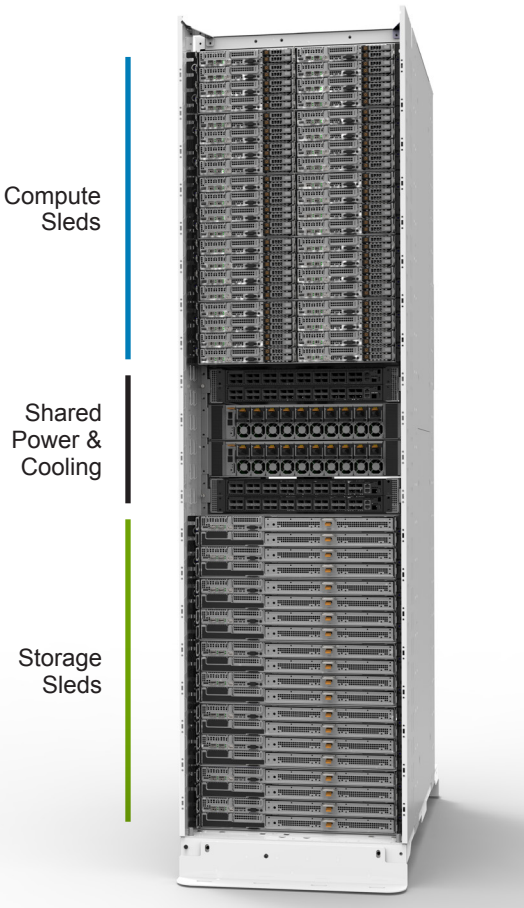

NOTE: The DSS 9000 is currently sold only to select large scale customers (typically more than 10 racks). To inquire about the DSS 9000 and Extreme Scale Infrastructure, contact your Dell EMC sales representative or email ESI@dell.com.

Much easier administration

With the software innovations described above, administering an OpenStack cloud environment becomes much easier. A single command allows you to allocate resources, with specific characteristics, from anywhere in the rack – without needing to know their physical locale. For example:

python *podm_tool.py* node allocate “Sample-Workload” 3 CPU=2 MEM=64000

Once resources are allocated, you can assemble them with another simple command, then use the orchestration tool to configure BIOS, firmware and RAID, and ultimately deploy the newly created nodes on the OpenStack cluster. The previously mentioned *Dell EMC OpenStack Kickstart Guide for the DSS 9000* makes deployment even easier by providing a step-by-step description of how to allocate, provision and deploy nodes on an OpenStack cloud.

Conclusion

As cloud computing environments become ever larger and more complex, Dell EMC is working with Intel® to simplify cloud management at scale, so that businesses can easily and rapidly respond to the demands of their customers and spend more of their effort and resources on their own pressing business issues rather than IT management challenges. ■