IBM

Move to the Cloud Easily with IBM SmartCloud Entry V3.1 and IBM Common Cloud Stack

- Lower your total cost of computing with this simple-to-use product
- Use the power and resources of the cloud with ease
- Create and manage a Microsoft Hyper-V cloud

Juscelino Candido de Lima Junior
Louis Fuka

ibm.com/redbooks

Redpaper

International Technical Support Organization

Move to the Cloud Easily with IBM SmartCloud Entry V3.1 and IBM Common Cloud Stack

December 2013

REDP-5030-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (December 2013)

This edition applies to Version 3.1 of IBM SmartCloud Entry (product numbers are dependent upon configuration) and to the IBM Common Cloud Stack, which is powered by the OpenStack grizzly release from the OpenStack Foundation.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Active Memory™
BladeCenter®
DB2®
IBM®
IBM Flex System™
IBM Flex System Manager™
IBM SmartCloud®
Informix®
Power Systems™
PowerVM®
PureFlex™
Redbooks®
Redpaper™
Redbooks (logo)  ®
System i®
System x®
Systems Director VMControl™
Tivoli®
XIV®

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redpaper™ publication describes our newest product for helping businesses enter the world of cloud technology: IBM SmartCloud® Entry, Version 3.1. This software includes the IBM Common Cloud Stack, which is powered by open source OpenStack from the OpenStack Foundation. These products work together to help you move your business to the cloud. This paper explains the ways that you can take advantage of this extensive and powerful technology.

With the power of cloud computing, you can enhance your existing services, extend your market reach, and create new markets for your services. IBM SmartCloud Entry V3.1 is easy to use, even for first-time cloud service users. It can enable you to benefit from cloud technology in less time and with less effort than you might think.

This Redpaper features two scenarios that demonstrate the ease of carrying out processes with IBM SmartCloud Entry software. The information is directed to two primary audiences. Chapter 1 is directed mainly to decision-makers, such as CEOs, CIOs, and CFOs, who need to know about cloud technology and the power that it offers. The remainder of the paper is directed to IT professionals, such as information architects, business intelligence administrators, and database administrators, who need to know about the functions and capabilities of SmartCloud Entry and Common Cloud Stack.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center, in New York state.

Juscelino Candido de Lima Junior is a certified Expert IT Specialist and a workload automation expert. He supports multiple clients and designs architectures worldwide. He has assisted clients as an international consultant for IBM Tivoli® Workload Automation products. Juscelino's areas of expertise include systems architecture, virtualization, networks, distributed systems, production control, and systems administration. He holds numerous certifications on many products and heterogeneous technologies.

Louis Fuka is a Client Technical Specialist with 15 years of experience at IBM. He holds a Master of Science degree in Electrical Engineering from the University of Texas at Austin and a Bachelor of Science degree in Electrical Engineering from the University of Notre Dame. His areas of expertise include cloud computing, virtualization, IBM Common Cloud Stack, Tivoli software, technical sales, and technical instruction.

Thanks to the following people for their contributions to this project:

- Jeff Boleman, Cloud Architect, Growth Initiatives
- Jay S. Bryant, Linux Developer, OpenStack Development
- Ella Buslovich, IBM Redbooks® Graphic Specialist
- Jim Carrigan, Manager, ATS Specialist, IBM System x®, Cloud and Clusters, Technical Sales
- Wissam Elriachy, SmartCloud Entry on Power Product Owner
- Michael J. Fork, Lead Architect for Hosted Private Cloud, CloudFirst Innovation
- Daniel L. Hiebert, Software Developer, OpenStack Development
- Prashanth S. Kaggali, Technical Consultant
- Karen Lawrence, IBM Redbooks Technical Writer
- Andre Eduardo Botura Lentini, Manager, Production Support
- Jim Moody, Advanced Technical Support Specialist, System x, Linux, Technical Sales, North America
- Eddie Munive, Certified IT Specialist, Cloud, 1350Clusters, System x, IBM BladeCenter® ATS Americas
- Matt Riedemann, Advisory Software Engineer, OpenStack Development
- Gregory Ross, Program Manager, Worldwide Internal Enablement
- Boden Russell, Advanced Cloud Solutions and OpenStack Development
- Yuri Henrique Sierakowski, Tivoli Web Services Tool Support
- Richard Theis, Cloud Computing and Emerging Solutions Development
- Randy Zhang, Delivery Project Manager, Service Operations Delivery, IBM Certified Senior Project Manager, Precision Machine Products Association

Thanks to the following authors of an earlier IBM Redpaper publication, *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908, published in February 2013:

- Jeff Boleman
- Martin Gudmundsen
- Yu Zhong Cao
- Randy Zhang

Now you can become a published author, too!

Here is an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form:

  ibm.com/redbooks

- Send your comments in email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Move to the cloud with IBM SmartCloud Entry V3.1

With the power of cloud technology, you can enhance your existing services, extend your market reach, and create new markets for your services. IBM SmartCloud Entry Version 3.1 is easy to use, even for first-time cloud users. By using SmartCloud Entry software, you can quickly benefit from cloud technology and add value to your business.

In this chapter, we briefly describe cloud technology and the resources that it makes available. Then we describe the IBM Common Cloud Stack, which is the software that connects your environment to the cloud. It is powered by OpenStack from the OpenStack Foundation. Lastly, we describe the SmartCloud Entry appliance, its features, and how typically cumbersome implementation and migration tasks can be carried out with ease.

1.1 Tap into the power of the cloud

IBM SmartCloud Entry 3.1 makes it easy to tap into the power of the cloud. Your business can benefit from the efficiencies of the cloud and from the features of SmartCloud Entry Version 3.1 software, the IBM Common Cloud Stack, and the SmartCloud Entry appliance.

In addition to describing cloud technology and the resources that it makes available, this paper includes two scenarios to demonstrate how easy SmartCloud Entry is to use for key tasks:

- Process steps for integrating SmartCloud Entry with the Microsoft Hyper-V appliance, the newest appliance that is supported by this software
- Process steps for image management, such as building, importing, and activating images

This Redpaper is for two primary audiences. Chapter 1 is directed to decision-makers, such as CEOs, CIOs, and CFOs, who need to know about cloud technology and the power that it offers. The remainder of the paper is directed to IT professionals, such as information architects, business intelligence administrators, and database administrators, who need to know about the SmartCloud Entry functions and capabilities.

1.1.1 The OpenStack Foundation

We extend our gratitude to the OpenStack Foundation for their efforts in developing open source OpenStack. It is because of their efforts that the resources are available to further develop the product into IBM Common Cloud Stack by adding capabilities that are based on lessons learned, techniques from our experience in implementing cloud solutions, and other differentiators, including tools for these tasks:

- Monitoring
- Billing
- Accounting
- Business process management
- Connecting to the Microsoft Hyper-V server

For the purposes of this paper, when a property or properties is shared by OpenStack from the OpenStack Foundation and the IBM Common Cloud Stack, we describe them collectively as the *OpenStack products*. Otherwise, we refer to them individually as *open source OpenStack* and *IBM Common Cloud Stack*.

1.1.2 Resources that cloud computing makes available to you

As defined by the National Institute of Standards and Technology (NIST)[1], cloud computing is “a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (for example, networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.” Cloud computing enables convenient, on-demand access to resources on your network, such as servers, storage, and software. By either definition, you can see that the cloud offers multiple services, all of which can result in a lower total cost of computing. According to NIST, cloud computing includes these characteristics:

- On-demand self-service: Users can request resources when they want them.
- Network access: Requested resources are available from various locations or devices.

[1] See the National Institute of Standards and Technology (NIST) website: http://www.nist.gov/itl/cloud/index.cfm

- Resource pooling: Resources are made available from a pool. When a user finishes using it, the resource is returned for reallocation. Virtualization means that the resources are selected that are best-suited to meet the immediate need.
- Rapid elasticity: You can rapidly react to changing user requirements by changing the allocation of resources.
- Measured service: The quality and amount of resources that are used by each user can be measured accurately and dynamically. This enables you to account for, and charge for, allocated resources.

1.1.3 Technologies that enable cloud computing

Cloud computing builds upon several existing technologies, including virtualization, automation, and service management. Used in both public (Internet-based) and private (intranet-based) clouds, these technologies are essential to the efficiency and lower total cost of computing that are possible with the cloud.

Virtualization

Virtualization is the process of simulating computing resources and then using those resources to perform computing functions. The simulated resources must be mapped to physical (non-virtualized) resources. This mapping can be manipulated to provide flexibility and to make the best use of the available resources.

Traditionally, computing resources (central processing units, or CPUs, for example) have been tied to physical objects. In traditional distributed computing (in which resources are shared among networked computers), physical computing machinery was allocated only to applications. With the capabilities and advances of virtualization, these services are expanded to include RAM, storage, and networking resources in addition to CPUs.

Applications do not typically use all of the allocated processing capacity all of the time, so mapping an application to a physical set of processors is wasteful. By virtualizing processing capabilities, multiple virtual processors can be mapped to the same physical processors. The total simulated capacity of these virtual processors can exceed the capacity of the physical processor to which they are mapped. However, because not all of the applications use the maximum allocated processing capacity at the same time, performance remains acceptable, and the physical processing capacity is used more efficiently. This results in lower total cost.

As an example, a simulated data store can initially be mapped to a hard disk drive array (storage that contains multiple hard disks). To improve data access time, the data store can be remapped to a solid-state drive that has faster access times. To improve reliability, the simulated data store can be mapped to several redundant physical devices that can be located around the world. The application that is using the simulated data store interacts with the virtualized data store and does not detect the actual location of the data.

Virtualization of computing resources is becoming commonplace because it increases flexibility, efficiency, and reliability. Overall, it is virtualization that provides the resource pooling and rapid flexibility that cloud computing requires. However, IT organizations must develop and implement new business processes to provide, track, control, account for, and charge for the services that are built on these virtualized assets.

This is where your organization benefits from SmartCloud Entry software.

Automation

Provisioning computing resources requires many steps in a heterogeneous environment where processors, memory, storage, networking, software, licensing, and many other components provide services. To provide interconnected, functioning services, many tasks must be completed to ensure that these components are compatible.

These tasks can be automated with a series of tools, some simple and some complex. Scripts, installers, schedulers, and other automation tools can be set up to handle these tasks with a minimum of intervention by administrators. These tools provide repeatability and reliability while reducing the cost and the time that is required to deliver IT services.

Cloud computing relies on automation to ensure that requested services are delivered in a timely manner. However, as with the virtualization process, the automation process must be managed. Processes must be encoded, documented, and maintained. Incidents of improper function must be corrected or compensated for before the service is restored. Ongoing or recurring problems must be analyzed and addressed.

Service management

Although cloud computing is now commonplace, it is rapidly evolving. The role of open source code is growing, which lowers total cost of computing and promotes shared community standards. IBM continually refines its offerings to provide more sophisticated services and embraces several strategies to maintain technical leadership:

- Use open standards and technologies
- Contribute to the OpenStack Foundation and other organizations
- Provide a range of cloud delivery options
- Offer solutions that work with open source and third-party components

As described in 1.1.1, “The OpenStack Foundation” on page 2, SmartCloud Entry uses IBM Common Cloud Stack, which is powered by open source OpenStack, to interface with the cloud environment. Common Cloud Stack adds value by incorporating monitoring, billing, accounting, and business process management capabilities in cloud environments, plus tools for connecting to the Microsoft Hyper-V server.

1.1.4 Standards for cloud architecture and implementation

Now we describe some of the standards that are used to maintain consistency and interoperability when implementing cloud environments. These standards can help you more fully understand the products and processes that we cover in this paper.

The IBM Cloud Computing Reference Architecture

The IBM Cloud Computing Reference Architecture (CCRA) is a guide for designing and implementing cloud computing. It defines the cloud building blocks, such as the architectural elements that compose the cloud solution and the relationships between these elements.

Roles in a cloud implementation

Several roles included in the CCRA apply specifically to a cloud environment:

- Cloud service *creator*: Builds cloud computing services by using service creation tools to assemble elements into a service offering. Sets up methods to manage the service lifecycle. An example of a cloud service creator is a team of services personnel that gathers requirements and implements an IT function as a cloud-delivered service.

- Cloud service *provider*: Responsible for day-to-day operation of the cloud service. An example of a cloud service provider is a Managed Service Provider that provides virtual desktop infrastructure services to a client.
- Cloud service *consumer*: Requests and uses the cloud-based service. An example of a cloud service consumer is a software test team that requests cloud-based virtual machines with specific software versions to test a software release.

Roles might overlap, and one person or entity can fulfill more than one role. A Managed Service Provider, for example, might create and provide a cloud-based service on a cloud infrastructure that the client owns.

Service models in a cloud implementation

These are the typical service models that cloud providers use:

Infrastructure as a service (IaaS)
The cloud service provider makes fundamental computing resources available to the cloud service consumer. Resources can include processing power or storage. The consumer uses these resources to run operating systems and applications. For example, a data store might be provided to the consumer, who uses the data store as a mounted volume on a virtual system.

Platform as a service (PaaS)
The provider makes an environment available where the consumer runs applications. The environment generally includes virtualized hardware assets, an operating system, and perhaps programming tools. An example of PaaS is provisioning a customer with a Linux virtual machine that the customer installs and runs as a web server.

Software as a service (SaaS)
The provider makes available the applications and databases that the consumer needs, including virtual hardware and an operating system. An example of this is a computer-aided design (CAD) system that incorporates a physical graphical processing unit that is attached to a virtual machine that runs specialized design software.

Business process as a service (BPaaS)
The provider makes an environment available to implement a competed business function. An example is a payroll processing or expense reporting system.

SmartCloud Entry functions as an IaaS, which is the building block for cloud implementation.

For more information about the CCRA, see Appendix A, “Cloud architecture standards” on page 61.

Common Cloud Management Platform

The Common Cloud Management Platform (CCMP) is a set of functions that deliver and manage cloud services. CCMP functions are used to define the appropriate service models for cloud services (IaaS, PaaS, SaaS, or BPaaS).

The IBM CCRA specifies three steps for creating a detailed design of a cloud computing solution:

1. Define the functional and nonfunctional requirements.

   Functional requirements are those that are necessary to build the cloud service. Nonfunctional requirements are constraints within which the cloud service creator and management must operate.

2. Derive the logical definition.

   The architecture team uses the CCRA and the requirements that are defined in Step 1 to determine which architectural components are needed and which products are used to implement these components.

3. Create the physical design.

   The layout of the infrastructure (nodes, networks, and software) is determined.

The CCRA incorporates several work products. These work products describe the planning tasks that are necessary for a successful implementation of cloud services.

For more information about CCMP, see Appendix A, “Cloud architecture standards” on page 61.

1.2 IBM SmartCloud Entry Version 3.1

SmartCloud Entry 3.1 makes it easy to implement your first cloud environment. For clients who have an established cloud environment and are already automated, virtualized, and service-oriented, this software can help you make cloud management activities easier, too. You can migrate your current systems and monitor your systems and functions.

In the following sections, we describe the SmartCloud Entry appliance and the advantages, features, and architecture of the SmartCloud Entry Version 3.1 migration tool.

1.2.1 What SmartCloud Entry can do for you

SmartCloud Entry is easy-to-deploy, simple-to-use software that features a self-service portal for workload provisioning, virtualized image management, and monitoring. It is innovative and cost-effective, and it includes security, automation, basic metering, and integrated platform management.

This software has several advantages:

- Fast time to value. It can be installed and configured in a matter of hours.
- Works with your existing infrastructure. You can integrate it with many existing cloud environments.
- Self-service interface. With this interface and with configurable email notification policies, you can generate, modify, extend, and withdraw cloud service requests.
- Heterogeneous cloud management. A hypervisor is the component of cloud management that manages virtual machines. SmartCloud Entry manages hypervisors that run on IBM System x, IBM Power Systems™, IBM PureFlex™ Systems, and IBM Flex Systems, and it does so all from a single graphical user interface (GUI). Managed environments include VMControl, kernel-based virtual machines (KVMs), and VMware. New in SmartCloud Entry V3.1 is the management of Microsoft Hyper-V.
- Project-level customizations. These customizations enable the cloud to conform to your particular operational structure.
- Configurable expiration and approval policies. These policies alleviate virtual image proliferation across a data center.
- Billing capabilities. These capabilities provide a way to determine the cost of the cloud services that are provided.
- Metering capabilities. These capabilities provide feedback about resource use.

1.2.2 Features of IBM SmartCloud Entry

Version 3.1 continues to provide features from SmartCloud Entry V2.4.

Note: SmartCloud Entry V3.1 refers to *instances* rather than *workloads* (from V2.4). Therefore, we use that terminology here.

Summary of SmartCloud Entry V2.4 features

Version 2.4 includes the following features[2]:

- A single SmartCloud Entry V2.4 appliance can manage multiple environments:
  - IBM Power Systems compute nodes and x86 compute nodes.
  - Multiple IBM Power Systems and KVM environments that use IBM Systems Director with VMControl, or IBM Flex System™ managers on the IBM PureFlex System and IBM Flex Systems.
  - VMware environments with multiple vCenter instances.
- A self-service GUI from which cloud service consumers can generate and manage requests for cloud services and administrators can manage SmartCloud Entry.
- Discovers existing virtual machine (VM) images and instances.
- Control of VM image configuration. Administrators can set default values for VM size (CPU, storage, and RAM), target host, and network configuration. The administrator has the option to display these settings to the cloud services consumer for modification at request time or keep them hidden to make these values mandatory.
- Integration with LDAP.
- Ability to manage approval policies and expiration policies by cloud or by project.
- Custom network configuration templates can be assigned to a single cloud environment or used across all managed cloud environments.
- Options to enable metering and view metrics by elapsed time, CPU hours, storage hours (in gigabytes), and memory hours (in gigabytes) for each VM instance.
- Billing, including the following tasks:
  - Enable billing
  - Set up accounts, add funds, and assign cloud service consumers to accounts
  - Choose policies for VM instances in delinquent accounts (stop, delete, take no action)
  - Establish pricing for CPU, RAM, and storage

Features of IBM SmartCloud Entry V3.1, summarized

Version 3.1 adds the following capabilities:

- Manage Microsoft Hyper-V cloud environments by using IBM Common Cloud Stack, which includes these features:
  - Support for Microsoft Windows and Linux images
  - Flavor management to define sizing for VM instances
  - Secure Shell (SSH) key management
  - IPv4 and IPv6 network management
- Use images that are generated with the IBM Image Construction and Composition Tool. These now enable secure access through public and private key pairs.

[2] See *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908, for a full list of V2.4 features and four detailed deployment use cases, including before and after architectural diagrams.

- Pass user data to a provisioned instance that, in turn, provides a mechanism for configuring the instance.
- Suspend and resume a workload on all cloud types.
- Deploy multiple instances of an image at one time.
- Define multiple configurations of single virtual images. These image configurations can exist in different projects, which allows for cross-project sharing of VM images.
- Set user SSH keys during deployment to VMware and Hyper-V environments.
- Enable secure access to LDAP servers, a requirement that is necessary to meet most security guidelines and standards.
- Use enterprise web browsers.
- Resize a disk at deployment time for KVM or IBM Power Systems compute nodes.
- Enable a remote restart for KVM or Power Systems compute nodes.
- Gain the following capabilities:
  - When integrated with VMware cloud environments, SmartCloud Entry V3.1, you get the following options:
    - Customize user data for a deployed instance.
    - Set a user password and SSH key.
    - Delay powering off a system to allow completion of the shutdown process.
    - Deploy an instance that is a linked clone of an image.
  - When integrated with IBM Flex System Manager™ Advanced, you can use SmartCloud Entry V3.1 for these tasks:
    - Set a workload priority, which is useful during host evacuation and for workload mobility.
    - Pin a workload to a host to prevent it from being moved during relocation for workload balancing.
  - When using Flex System Manager Advanced for Power Systems compute nodes, SmartCloud Entry V3.1 provides these options:
    - Attach storage by using N_Port ID Virtualization (NPIV) when you are deploying storage to a system pool.
    - Deploy IBM System i® images to system pools.
    - Enable IBM Active Memory™ Expansion and change settings during deployments.
- Integrate SmartCloud Entry V3.1 appliances with VMware, KVM, and Hyper-V. These appliances are functionally equivalent, which gives the implementer flexibility in where to place the appliance.

1.2.3 The SmartCloud Entry appliance

SmartCloud Entry can be delivered as a prepackaged virtual machine, with networking preconfigured. Like SmartCloud Entry V3.1, the SmartCloud Entry appliance runs on VMware, KVM, and Hyper-V hypervisors. These versions all have the same functionality for managing clouds.

Figure 1-1 shows the components of the SmartCloud Entry appliance.

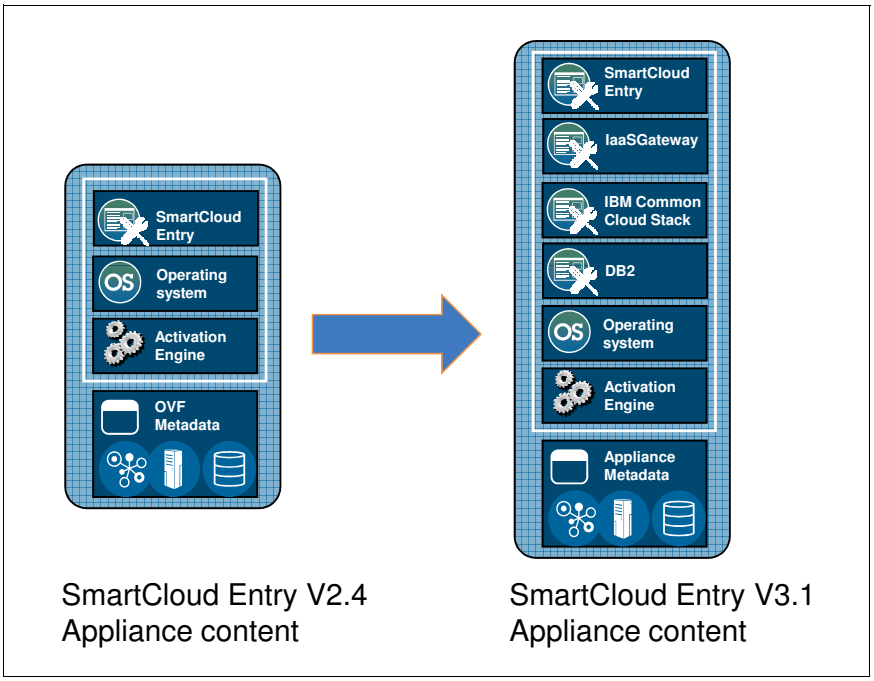

Figure 1-1 The IBM SmartCloud Entry appliance

The following components are included in the appliance:

- An activation engine to allow configuration of the appliance when it is first started after deployment
- Virtual Machine Administration

  Administrative functions are performed by using the appliance management utility named *sceappmgr* that is used for the following functions:

  - Enable, stop, start, restart, and view IBM Common Cloud Stack services
  - Stop, start, restart, and view the status of SmartCloud Entry software and its services
  - Stop and restart the SmartCloud Entry appliance
  - Migrate from SmartCloud Entry V2.4
  - Apply fix packs
  - Manage passwords and credentials
  - Collect logs
  - Manage storage volume groups to expand and manage the appliance storage space
- An instance of IBM DB2® database
- An instance of IBM Common Cloud Stack software
- The IaaS gateway that is used by SmartCloud Entry software to interface with IBM Common Cloud Stack

1.2.4 The SmartCloud Entry architecture

Figure 1-2 shows the IBM SmartCloud Entry architecture.

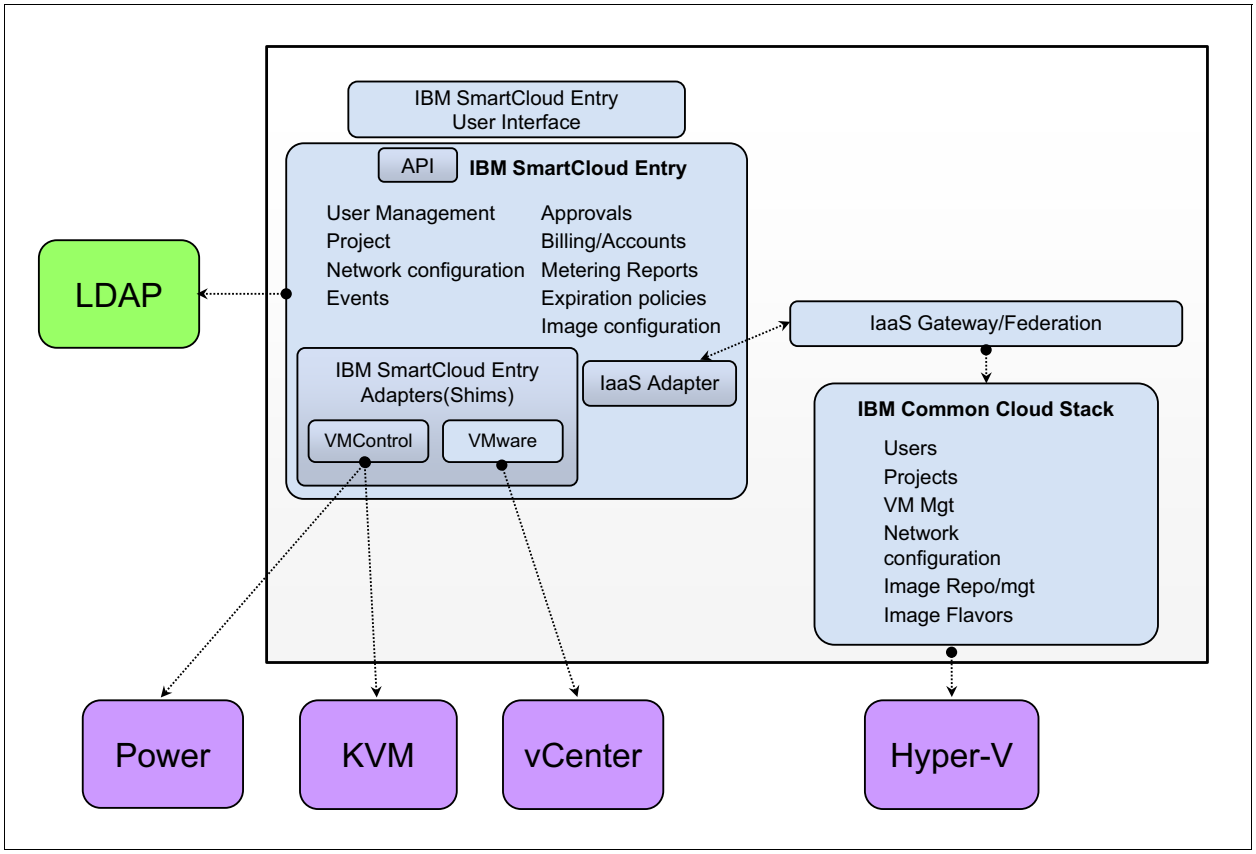

Figure 1-2 IBM SmartCloud Entry architecture

SmartCloud Entry includes these components:

- Self-service GUI: Used for administration activities and for cloud service consumers to request cloud services.
- Application programming interface (API): Integrates with custom GUIs and third-party applications.
- Proprietary adapters (also called *shims*): Code that enables SmartCloud Entry software to interface with cloud environments. These shims are included:
  - VMware shim: This shim enables SmartCloud Entry to interface with one or more instances of VMware vCenter to manage VMware cloud environments.
  - VMControl shim: This shim enables SmartCloud Entry to interface with IBM Systems Director VMControl™ to manage IBM Power Systems and KVM cloud environments.
- IaaS adapter and IaaS gateway: This adapter and gateway enable SmartCloud Entry to interface with IBM Common Cloud Stack to manage Microsoft Hyper-V Cloud environments.

Note: Because SmartCloud Entry uses three different adapters, the management functionality of the three different types of clouds (VMware, KVM/Power, and Hyper-V) is slightly different. For example, instances that are deployed in Hyper-V clouds are sized according to specifications called *flavors*, but CPU, RAM, and storage are set individually for instances deployed in VMware clouds.

1.2.5 SmartCloud Entry migration

You can use the SmartCloud Entry migration tools to migrate data, including user data, projects, accounts, and network templates from an instance of SmartCloud Entry V2.4 to V3.1. Properties files are also migrated to preserve the behavior of SmartCloud Entry after migration. For example, when an SmartCloud Entry V2.4 system has billing enabled, the migrated SmartCloud Entry system has billing enabled also.

The preferred migration method is to install the SmartCloud Entry appliance, and then migrate the data from the instance of SmartCloud Entry V2.4. (The instance of SmartCloud Entry V2.4 can be either an appliance or natively installed.) This uses the IBM Common Cloud Stack instance that is installed on the SmartCloud Entry appliance.

It is not necessary that the old and new servers be running at the same time. Therefore, you can store the necessary data from V2.4 and reuse the IP address from the V2.4 server for the SmartCloud Entry appliance.

1.2.6 How SmartCloud Entry uses IBM Common Cloud Stack

When managing KVM, VMware, and IBM Power hypervisors, SmartCloud Entry V3.1 uses shims to connect to the hypervisors. For more information about the use of SmartCloud Entry to manage KVM, VMware, and Power System clouds, see *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908.

SmartCloud Entry V3.1, however, delivers with IBM Common Cloud Stack, the software that manages connections to Microsoft Hyper-V hypervisors. The use of IBM Common Cloud Stack to integrate with Microsoft Hyper-V is a new feature of SmartCloud Entry V3.1. The process for carrying out the integration is explained in detail in Chapter 3, “Integrating IBM SmartCloud Entry with Microsoft Hyper-V” on page 25.

Note: SmartCloud Entry V3.1 is required for managing Hyper-V clouds. It is bundled with IBM Common Cloud Stack on the V3.1 appliance. Managing Hyper-V clouds requires that both SmartCloud Entry V3.1 and IBM Common Cloud Stack are installed.

2

Architecture of the OpenStack products

This chapter describes some of the commonalities between IBM Common Cloud Stack software and open source OpenStack and the areas where the features that IBM adds to the IBM Common Cloud Stack make it unique and more valuable. We also include troubleshooting tips to demonstrate how easy troubleshooting can be.

As described in 1.1.1, “The OpenStack Foundation” on page 2, IBM Common Cloud Stack is built upon open source OpenStack software. Throughout this paper, when these two products are identical (for example, in the overall architecture), we describe the software collectively as the *OpenStack products*. Otherwise, the software is described individually as *open source OpenStack* and *IBM Common Cloud Stack*.

2.1 The OpenStack products

Some facets of the OpenStack products are identical, although some differ. This section explains some of these similarities and differences. Following a brief description of these, we review the architecture of the OpenStack products.

2.1.1 Architecture

Overall, the architecture among the OpenStack products is shared. The advances that are built into the IBM Common Cloud Stack result in more value, including the enablement of the enhanced IBM version to interact with today's prominent hypervisors (VMware and KVM) and Microsoft Hyper-V.

2.1.2 Releases and fixes

As with other software organizations and communities, the OpenStack Foundation releases code fixes and enhancements periodically. The current version of IBM Common Cloud Stack software is based on the open source OpenStack Grizzly release (2013.1). Releases of Common Cloud Stack are tightly coupled to releases of open source OpenStack from the OpenStack Foundation.

2.2 Architecture of the OpenStack products

The OpenStack products use the same module names for their individual components. IBM Common Cloud Stack adds unique features and modifications. Therefore, some components differ slightly from those of open source OpenStack software.

2.2.1 Overall architecture of the OpenStack products

Figure 2-1 shows the architectural components of the OpenStack products.

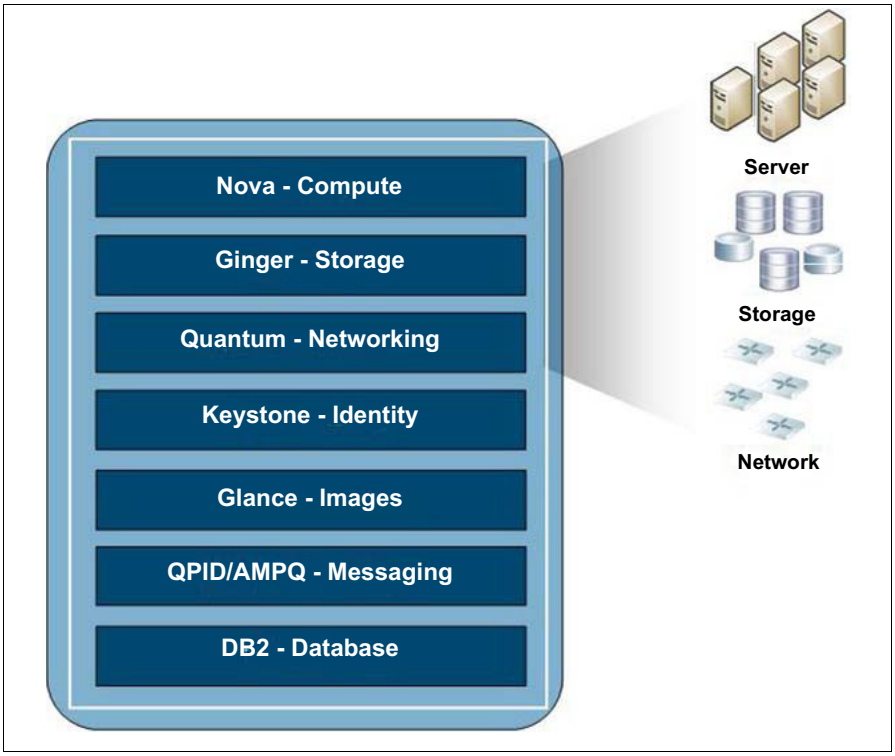

Figure 2-1 Architecture of the OpenStack products

In this model, open source OpenStack provides basic functionality for managing virtual systems. Base functionalities are extended by using IBM assets and code to produce IBM Common Cloud Stack, which is delivered to clients as part of IBM offerings in solution-specific configurations.

The sections that follow describe the components of the products. Because the OpenStack products are constantly evolving technologies, the components are referred to by their build names. For example, the build name for the open source OpenStack compute node is Nova. Both OpenStack products use the same build names and components. The IBM Common Cloud Stack provides additional value by supporting the IBM DB2 database.

2.2.2 Compute node: Nova

The underlying open source OpenStack project name is Nova. This is the central component of the OpenStack products, and it is where OpenStack communication with the hypervisors occurs. Nova uses every component of the open source OpenStack product to provision and manage virtual machines.

Primary features

Nova includes the following primary features:

- Manage virtualized server resources
- Manage local area networks (LANs)
- API with rate-limiting and authentication
- Distributed and asynchronous architecture
- VM image management
- Live guest migration
- Live VM management
- Floating IP addresses
- Security groups
- Role-based access control (RBAC)
- Projects and quotas
- Virtual Network Computing (VNC) proxy through a web browser
- Support for booting VMs from storage based on the Cinder infrastructure (Cinder is described in “Storage service: Cinder” on page 16)

In Nova, all communications are routed by using queue messages, and API requests are validated and placed in those queues. Users listen to the messages that are based on role and host names, and, following that, the responses are dispatched back through the queue.

The compute service manages the individual hypervisors and compute nodes. This is directly implemented by hypervisor APIs and results in the following benefits:

- Avoids abstraction layers that bring least-common-denomination support
- Enables easier exploitation of hypervisor differentiators:
  - The service instance runs on every physical compute node, which helps minimize failures if you have nodes spread throughout the network
  - Support for security groups that define firewall rules

2.2.3 Storage service: Cinder

Storage in the OpenStack products is not unified. The following types are typically available:

- A dedicated service code named Swift that provides object storage service on top of a distributed infrastructure. As an object storage service, it cannot be used directly as virtual disks for virtual machines. A typical use of Swift is to store the disk images.

- The common block storage service called Cinder. It controls the storage devices, such as SAN storage, and services, such as a Network File System (NFS) directly, and it manages storage volumes lifecycles. Storage volumes are typically used by the virtual servers as virtual disks. However, not all hypervisors and their associated drivers support attaching storage volumes managed by Cinder directly to virtual servers.
- Storage that is managed by the hypervisors directly. OpenStack products do not provide interfaces to directly manipulate this storage. Therefore, the virtual disks are created by the corresponding hypervisor driver code when virtual servers are created. Based on each configuration, the hypervisor management tools attach storage devices or services to the hypervisor. One example is the OpenStack VMWare driver, which is used to create the virtual disks on hypervisor-managed Virtual Machine File Systems (VMFSs) by default. The IBM PowerVM® driver has a similar option.

Figure 2-2 shows the various Cinder components and how they communicate.

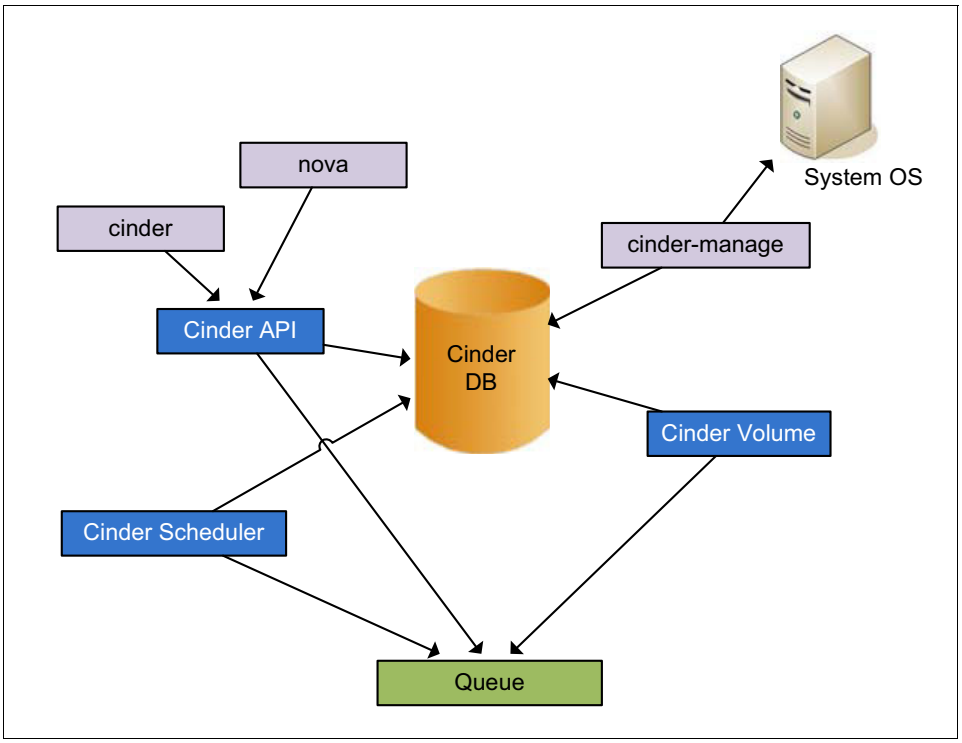

Figure 2-2 Storage components of the OpenStack products

As Figure 2-2 shows, Cinder includes the following components:

- Cinder API listens on Port 8776 by default and takes requests from the Cinder client and other operating system services.
- Cinder Scheduler interacts with Cinder volumes to load balance the volume creation.
- Cinder Volume runs on each system that serves and manages volumes.
- Cinder-manage provides control of Cinder users, projects, roles, shell selection, VPN connections, and floating IP address configuration.
- Cinder command-line interface (CLI) communicates with the Cinder API to create, delete, and list requests from the Cinder client and other operating system services.
- Nova is still the interface for attaching and detaching volumes.

Primary features

Cinder includes these primary features and functions:

- Provides persistent block-level storage devices for use with OpenStack compute instances
- Manages the creation, attachment, and detachment of the block devices to servers
- Supports booting virtual machines from Cinder-backed storage
- Provides snapshot and restore functionality
- Supports multiple volume drivers:
  - Logical volume management (LVM)-backed volumes, Internet Small Computing System Interface (iSCSI)
  - IBM XIV® Storage System (iSCSI)
  - SAN Volume Controller (SVC) (both iSCSI and Fibre Channel)
  - NetApp (iSCSI and NFS), EMC Corporation (iSCSI)
  - Hewlett-Packard Company Lefthand (iSCSI)
  - Reliable Autonomic Distributed Object Store (RADOS) block devices (for example, Ceph, a storage platform for a distributed file system).

A full list of supported volume drivers is available on the Volume Drivers page of the OpenStack.org website:

http://docs.openstack.org/grizzly/openstack-block-storage/admin/content/ch-volume-drivers.html

2.2.4 Networking: Quantum

Quantum provides networking as a service to OpenStack. It provisions and manages virtualized network resources, such as networks, ports, and attachments.

Quantum was previously part of the Nova component (nova-network). Quantum began as an incubator project to overcome the limitations that are caused by being a part of nova-network. In the Grizzly release, Quantum replaced nova-network.

Among the many advantages of Quantum are that you can create multiple private networks and control the IP addresses on those networks. API extensions can be added to enable greater control. Security features, quality of service, and network monitoring can also be configured in Quantum.

Quantum (OpenStack networking) includes the following main components:

- Quantum-server: This is the main process in OpenStack networking. It listens for API requests and passes them to the configured plug-in. By default, it listens on Port 9696. It interfaces with the database for persistent configuration, and, unlike other services, it sets up the database when the appliance is first started.
- Quantum-plugin-openvswitch-agent: This is the agent plug-in that runs on each hypervisor (compute system) to perform local Virtual Switch (vSwitch) configurations. Not all plug-ins require an agent.
- quantum-dhcp-agent: This provides Dynamic Host Configuration Protocol (DHCP) service to VMs. It is common across all plug-ins.
- quantum-l3-agent: This provides external network access to VMs through L3 network access translation (L3/NAT) forwarding. It is common across all plug-ins.

Figure 2-3 illustrates a realistic setup of Quantum services. All of these services can run on one node or on all nodes. The best configuration depends on your network load and needs.

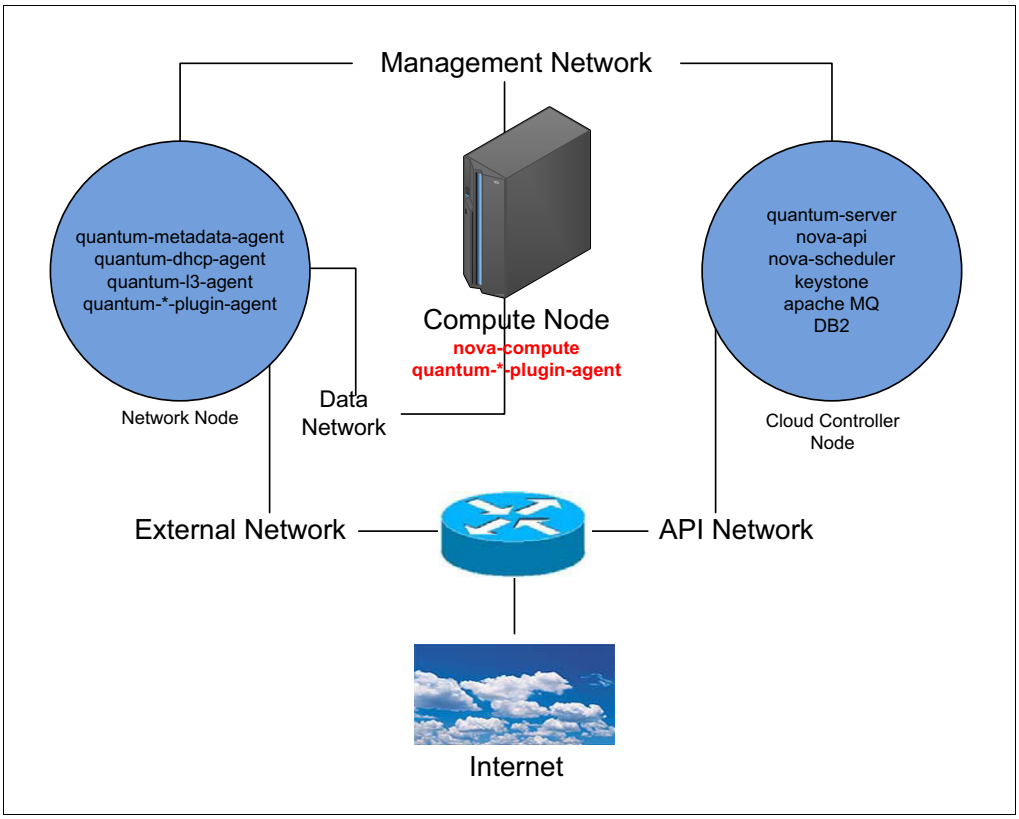

Figure 2-3 Networking in the OpenStack products

Primary features

Quantum enables the following primary features and functions:

- Flexible networking models to suit the needs of different applications or user groups
- Creation and deletion of project-specific L2 networks
- Attachment or detachment of a host from a network
- L3 support (Internet protocol address management, or IPAM: Dedicated static and DHCP, floating IPs, DHCP, routing)
- L4-7 support (load balancers)
- Extension framework to deploy and manage more network services, such as intrusion detection systems (IDS), load balancing, firewalls, and virtual private networks (VPNs)
- Support for the following technologies and products:
  - Open vSwitch
  - OpenFlow (Big Switch, Floodlight, NEC controllers)
  - Cisco Nexus and numerous SDN and network virtualization providers (for example, Nicira, Midokura, PLUMgrid, Brocade, and Mellanox)

2.2.5 Identity service: Keystone

Keystone is the identity and security service that is used by the OpenStack products for authentication. Keystone has two primary functions:

- User management: This module maintains user permissions and access levels.
- Service catalog: This module contains a catalog of which services are available and where the API endpoints are stored.

These are the main Keystone user authorization concepts:

- Users: This represents a user who has identifiable information, such as name, password, and email address.
- Tenants: A *tenant* is another term for a project. When an OpenStack service is requested, a tenant must be specified.
- Roles: Roles capture which operations a user has permission to perform for a tenant.

The main concepts of identity service management are *services* and *endpoints*:

- Services: This is an OpenStack service, such as Quantum or Nova. The services provide endpoints through which users can access resources and perform certain operations.
- Endpoints: This is where Keystone accesses the service, typically through a URL.

Primary features

Keystone includes the following primary features:

- Identity service: Authenticates user and password requests from multiple back ends (for example, SQL or LDAP)
- Token service: Validates and manages tokens that are used after initial user name and password verification
- Service catalog: An endpoint registry of available services
- Policy service: Authorizes API requests and includes a rule-based authorization engine and the associated rule management interface
- Access control: Domain, project, and user model with role-based access control (RBAC) for access to compute, storage, networking functions

2.2.6 Image service: Glance

Images in OpenStack are managed by using Glance services. The simplest way to describe Glance is as a file store, where each file can be described by using metadata that is stored in Glance. The compute service pulls the image from Glance on a start operation, and then it takes appropriate action to deploy the selected image.

The OpenStack Image service provides discovery, registration, and delivery services for disk and server images. The ability to copy or take a snapshot of a server image and immediately store it is a powerful capability of the OpenStack cloud software. If you are provisioning multiple servers, you can use stored images as a template for getting new servers up and running quickly and more consistently, rather than installing a server operating system and individually configuring more services.

Image service can also be used to store and catalog an unlimited number of backups.

The following image formats are supported by the Image service:[1]

- Unstructured disk image format (Raw)
- Amazon Machine Image (AMI) by Amazon
- Amazon Ramdisk Image (ARI) by Amazon
- Amazon Kernel Image (AKI) by Amazon
- Virtual Hard Disk (VHD) by Microsoft
- Virtual Disk Image (VDI) by VirtualBox
- qcow2 (QEMU) by Linux-KVM
- VMware (VMDK) by EMC Corporation
- International Organization for Standardization (ISO)

Note: The IBM PowerVM driver supports only the Raw format, and Hyper-v supports only the VHD format.

The following Image containers are supported by Glance:

- Open Virtualization Format (OVF), an open standard
- Amazon Machine Image (AMI) by Amazon
- Amazon Ramdisk Image (ARI) by Amazon
- Amazon Kernel Image (AKI) by Amazon
- BARE format, typically used in IBM SmartCloud Entry (indicates that there is no associated configuration file)

Primary features

Glance includes these primary features:

- Image registry is available, but storage is optional and delegated to a configurable store.
- Administrators can create base templates from which users can start new compute instances.
- Users can choose from available images or create their own from existing servers.
- Snapshots can also be stored in the Image service so that virtual machines can be backed up quickly.

2.2.7 Messaging service: Qpid

Qpid is a messaging system that implements the Advanced Message Queuing Protocol (AMQP), which lets programs communicate by exchanging messages. Those messages offer delivery, speed, security, and freedom from spam.

Routing can be designed in a flexible way, which enables easy support from common messaging paradigms, such as point-to-point, fan-out, publish-subscribe, and request-response.

The compute management stack control plane is built on the queue and database, which provide the following services:

- Communications route through the queue.
- API requests are validated and placed in the queue.
- Users listen to queues that are based on role or role plus host name.
- Responses are dispatched back through the queue.

[1] See the Disk and Container Formats page on the OpenStack.org website: http://docs.openstack.org/developer/glance/formats.html

2.2.8 OpenStack database

The OpenStack products require that a database store data for the main components. This is often used to track the current state of the system and to maximize the available resources. Each component (Nova, Cinder, Glance, Keystone, and Quantum) has a database for storing data and information.

A primary feature, specifically of the IBM Common Cloud Stack, is that DB2 is fully supported, beginning with the Grizzly release (IBM SmartCloud Entry V3.1). It is the default database option for IBM Common Cloud Stack, because DB2 is a well-known database that is supported by extensive documentation.

2.2.9 Process flow of the components

Figure 2-4 provides an overview of managers, services, and APIs. The program is written in Python, and the blue background indicates that the components are under the same Python runtime environment.

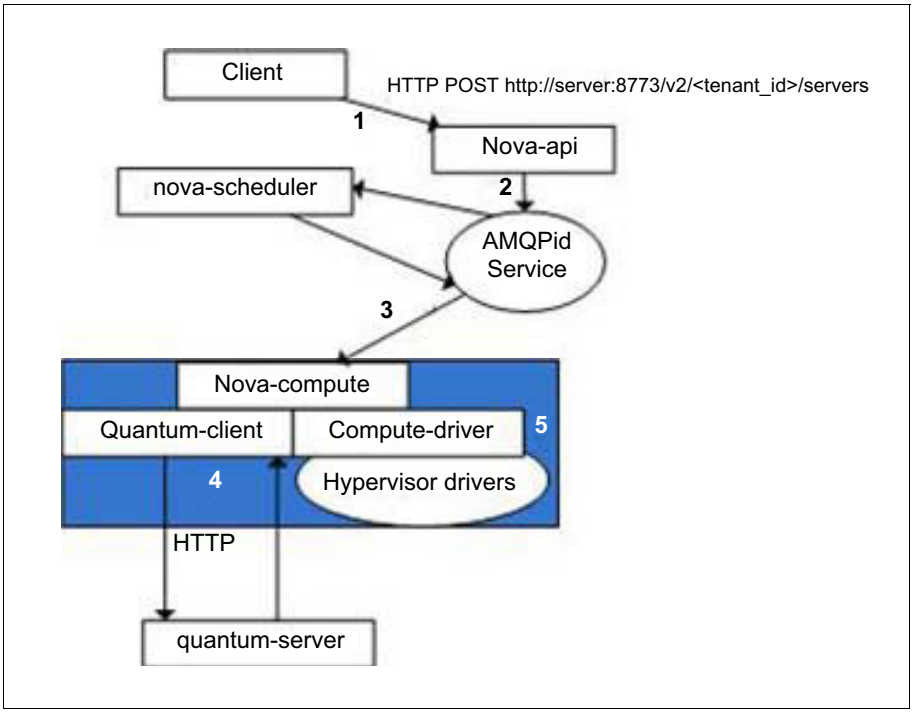

Figure 2-4 Nova components of the OpenStack products

The diagram shown in Figure 2-4 illustrates the following process flow:

1. A request is made to Nova-API by HTTP.
2. The Nova-API is multicast to the Nova scheduler (for placement).
3. The Nova-scheduler is cast to the chosen Nova-compute node.
4. The Nova-compute network information calls Quantum.
5. The Nova-compute node creates images in the hypervisor on the compute node.

2.3 Troubleshooting made easy

IBM SmartCloud Entry is designed for ease-of-use and ease-of-troubleshooting. Table 2-1 shows a list of log files to use for troubleshooting purposes. These particular log files are for use with issues that are related to the setup and maintenance of IBM SmartCloud Entry V3.1 with IBM Common Cloud Stack enabled.

Table 2-1 Summary of log files for services of the OpenStack products

IBM Common Cloud Stack and open source OpenStack service	Log files	Debug locale log files (defaults)
Keystone	/var/log/keystone/keystone.log	/var/log/keystone/keystone-debug.log
Cinder	/var/log/cinder/api.log /var/log/cinder/cinder-manage.log /var/log/cinder/scheduler.log /var/log/cinder/volume.log	/var/log/cinder/cinder-api-debug.log /var/log/cinder/cinder-manage-debug.log /var/log/cinder/cinder-scheduler-debug.log /var/log/cinder/cinder-volume-debug.log
Nova	/var/log/nova/api.log /var/log/nova/compute.log /var/log/nova/conductor.log /var/log/nova/nova-manage.log /var/log/nova/scheduler.log	/var/log/nova/nova-api-debug.log /var/log/nova/nova-compute-debug.log /var/log/nova/nova-conductor-debug.log /var/log/nova/nova-manage-debug.log /var/log/nova/nova-scheduler-debug.log
Quantum	/var/log/quantum/dhcp-agent.log /var/log/quantum/l3-agent.log /var/log/quantum/openvswitch.log /var/log/quantum/server.log	/var/log/quantum/quantum-server-debug.log
Glance	/var/log/glance/api.log /var/log/glance/registry.log	/var/log/glance/glance-api-debug.log /var/log/glance/glance-registry-debug.log
Image Management: Deploy / activation	/var/log/nova /var/log/glance /var/log/quantum /var/log/iaasgateway	NA
Image Management: Deploy / Delete	/var/log/glance	NA
Compute node	$INSTALL_DIR\IBM\SmartCloud Entry\Hyper-V Agent\log	NA

Several issues can be resolved with these log files:

- Image management issues: Review the IBM SmartCloud Entry logs for issues that are related to deploying or deleting images.
- Activation issues: You can start the debug process with Nova logs.
- Image deployment: After you deploy an image, it is installed on the compute node. For issues, review the Nova and Quantum logs at the compute node level.

For Hyper-V troubleshooting, follow these troubleshooting tips:

- To create an installation log file, start the installer by using the **msiexec** command with the logging arguments, as shown in Example 2-1 on page 24.

Example 2-1 Install log command

```
msiexec /i "IBM SmartCloud Entry Hyper-V Agent.msi" /L*V "C:\tmp\install.log
```

- To create an uninstallation log file, start the uninstaller by using **msiexec** with the logging arguments as shown in Example 2-2.

Example 2-2 Uninstall log command

```
msiexec /x "IBM SmartCloud Entry Hyper-V Agent.msi" /L*V "C:\tmp\install.log"
```

- For difficulties with the Nova or Quantum agent, inspect the following log files (see Table 2-1 on page 23):

  a. Verify that the Microsoft Windows services shown in Figure 2-5 are running.

Figure 2-5 Microsoft Windows services panel that displays IBM SmartCloud Entry V3.1 services

  b. Verify and analyze the Windows 2012 event logs for insightful information about Hyper-V, particularly in live migration scenarios. Look for any errors for the Hyper-V services. Figure 2-6 shows an example.

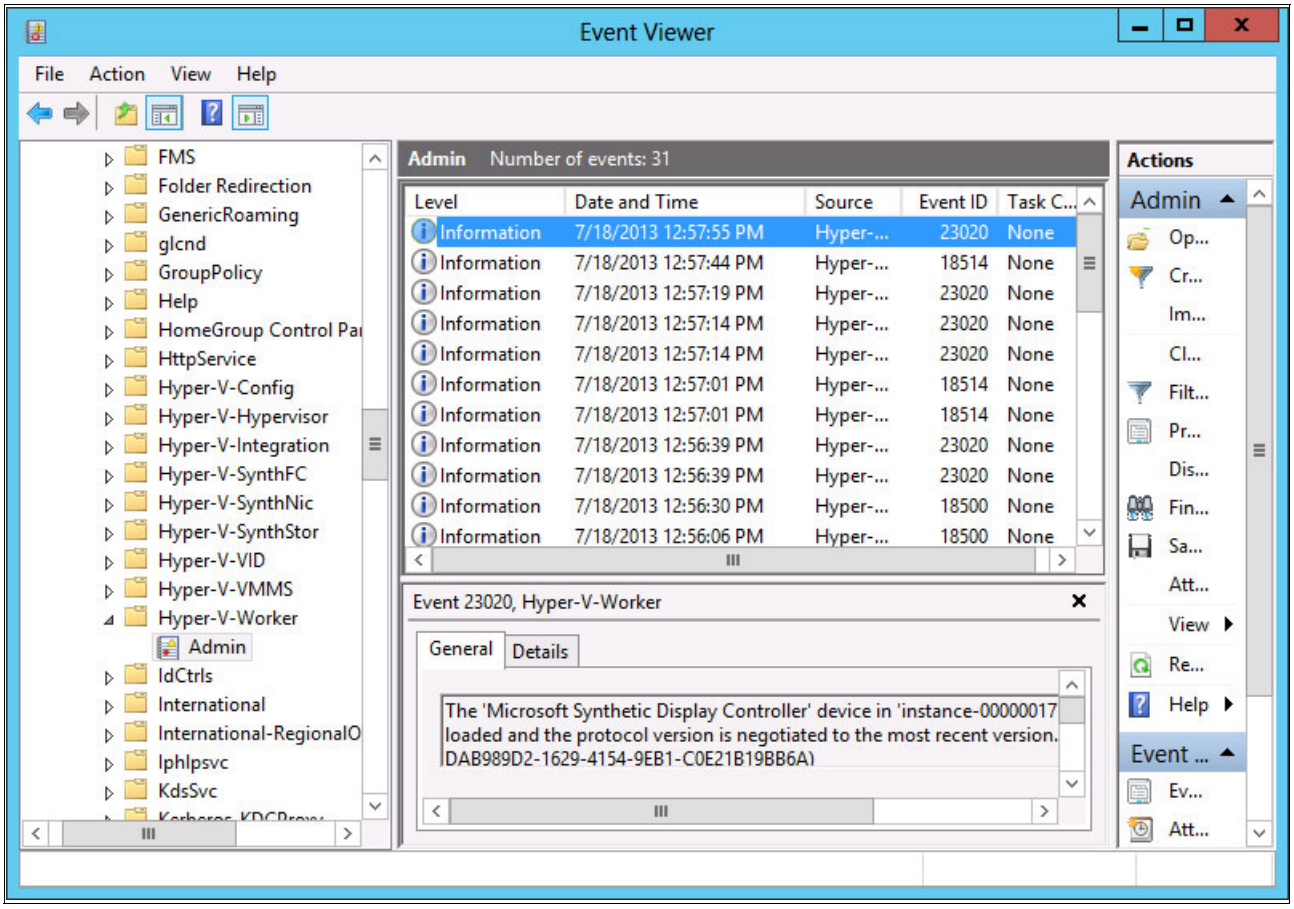

Figure 2-6 Event log

3

Integrating IBM SmartCloud Entry with Microsoft Hyper-V

This chapter describes a scenario in which we set up an IBM SmartCloud Entry appliance and demonstrate the ease of integration of IBM SmartCloud Entry software with IBM Common Cloud Stack and Microsoft Hyper-V software.

3.1 Scenario environment

This scenario is an extension of the scenario that is outlined in *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908[1]. In that Redpaper, we describe a scenario in which the client uses IBM SmartCloud Entry Version 2.4 to manage a VMware cloud that contains multiple ESXi servers. Now, with the ability of SmartCloud Entry Version 3.1 to manage Microsoft Hyper-V, the client wants to integrate the Hyper-V cloud management service. With both VMware and Hyper-V cloud management capabilities, this can attain a maximum potential for efficiencies.

Figure 3-1 shows the environment in which this scenario was implemented.

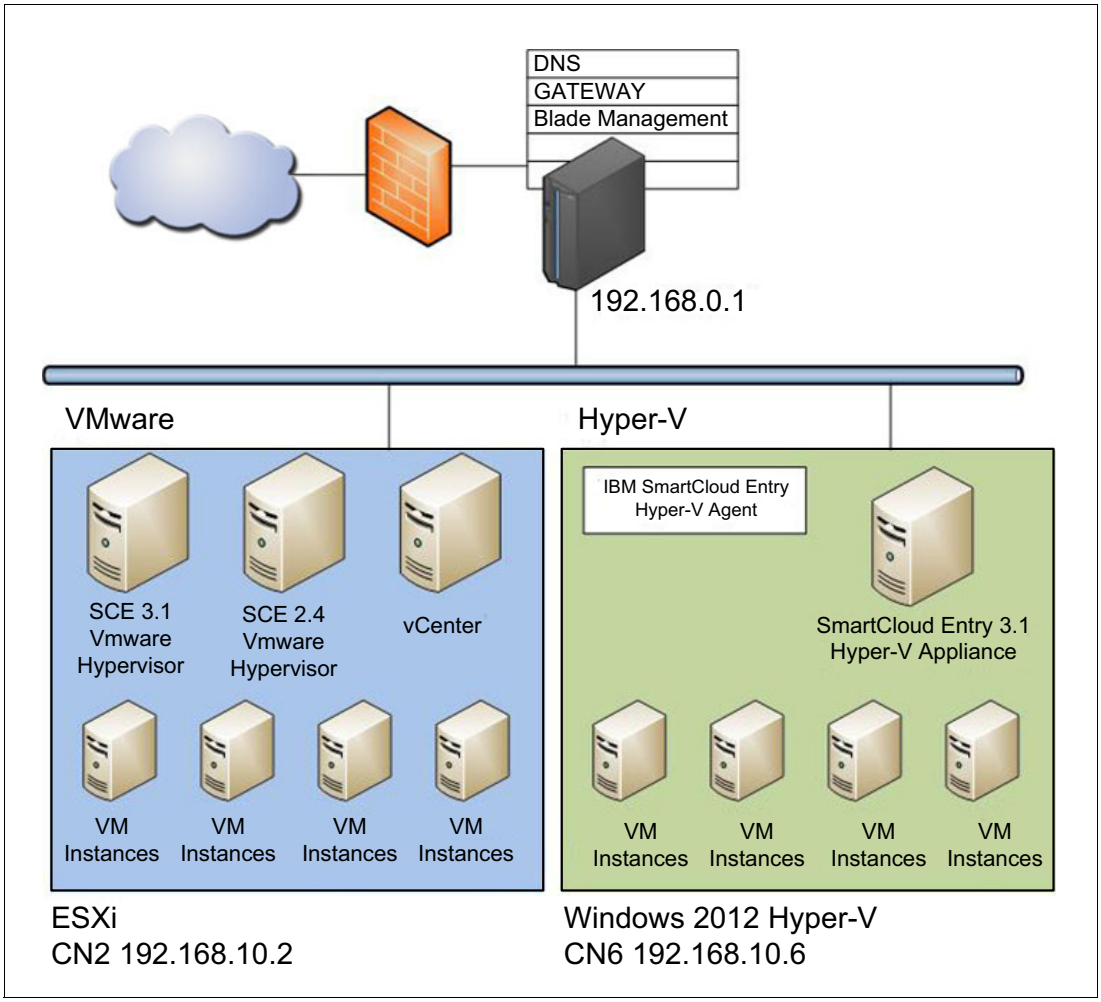

Figure 3-1 A development environment that uses IBM SmartCloud Entry V3.1 to manage Hyper-V

Note: Although the scenario in Figure 3-1 shows multiple VMware ESXi nodes and multiple Microsoft Hyper-V nodes, this scenario takes place in a development environment. A production scenario can have multiple VMware ESXi nodes and multiple Microsoft Hyper-V nodes.

[1] *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908: http://www.redbooks.ibm.com/abstracts/redp4908.html

3.2 Scenario overview

The client uses SmartCloud Entry V2.4 to manage a VMware cloud that contains multiple ESXi servers, as shown in Figure 3-2.

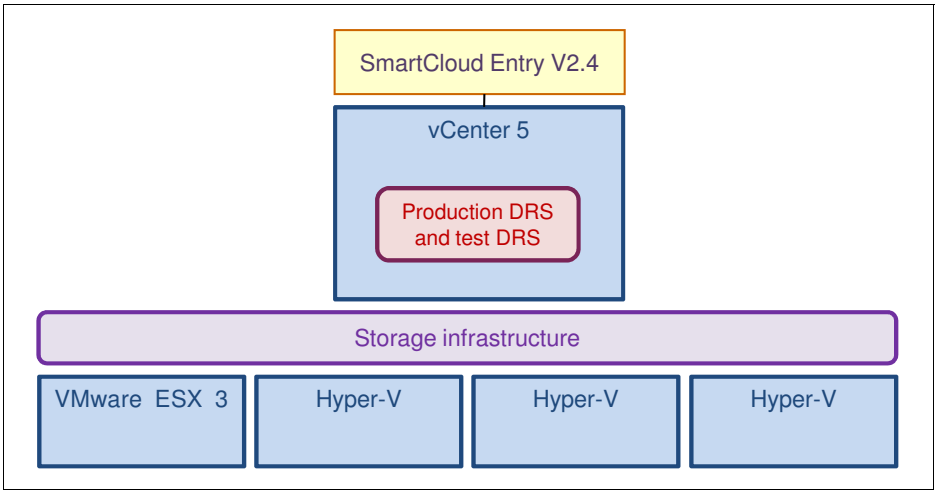

Figure 3-2 Before the migration, IBM SmartCloud Entry V2.4 manages a VMware cloud

Managing both the Microsoft Hyper-V and VMware cloud environments produces the architecture that is shown in Figure 3-3.

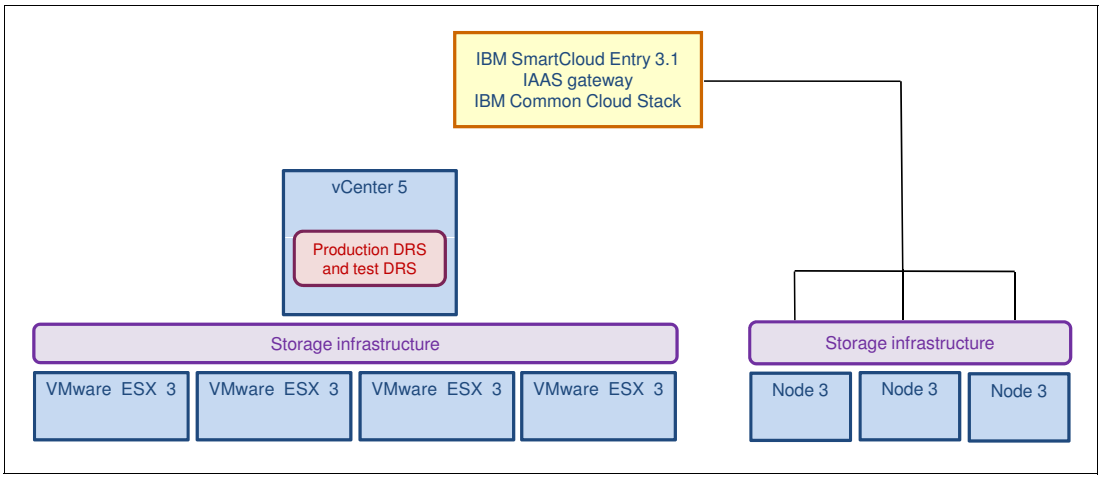

Figure 3-3 After the migration, IBM SmartCloud Entry V3.1 manages VMware and Microsoft Hyper-V

The process for upgrading to the newer version is described in 3.3, “Migration process” on page 27.

3.3 Migration process

We begin the migration with a SmartCloud Entry V2.4 environment with VMware ESXi servers, projects, users, and accounts. Rather than re-creating this environment in SmartCloud Entry V3.1, we migrate this data.

Because the SmartCloud Entry V3.1 appliance delivers with IBM Common Cloud Stack, we can install the V3.1 appliance, and then transfer the data from the V2.4 appliance. This is more efficient than updating the instance of SmartCloud Entry V2.4 to V3.1. In this scenario, we have a unique IP address for each appliance so that appliances can run at the same time.

Note: If you prefer to use the same IP address for each appliance, the process is simple:

1. Download the `/home/sysadmin/.SCE24` directory to a temporary storage location.
2. Shut down the SmartCloud Entry V2.4 appliance.
3. Install the SmartCloud Entry appliance.
4. Upload the `.SCE24` directory.

When you are using two IP addresses, the process requires the following steps:

1. Install the SmartCloud Entry appliance.

   Procedural steps for the installation are in the IBM SmartCloud Entry V3.1 Administrator Guide, which is available on this web page:

   `http://ibm.co/1c54U5v`
2. Apply fix packs to the appliance.
3. Transfer the `.SCE24` directory from the V2.4 appliance to the V3.1 appliance.
4. Migrate the data in the `.SCE24` directory from the V2.4 appliance to the V3.1 SmartCloud Entry appliance.
5. Enable an IBM Common Cloud Stack on the SmartCloud Entry appliance.
6. Install the IBM Hyper-V agent on a node that is running Microsoft Hyper-V.
7. Create an IBM Common Cloud Stack cloud in IBM SmartCloud Entry.

3.3.1 Install the IBM SmartCloud Entry appliance on a VMware ESXi node

Installing the SmartCloud Entry V3.1 appliance is nearly identical to installing the V2.4 appliance. Process steps for installing V2.4 are included in IBM SmartCloud Entry 2.4 Deployment Use Cases, REDP-4908[2]). The installation of V3.1 is the same except for two important points:

- The CPU, RAM, and storage requirements for the SmartCloud Entry appliance are significantly larger to support the instance of IBM Common Cloud Stack. Because IBM Common Cloud Stack is required to manage the Microsoft Hyper-V cloud environment, we need to allow for the following requirements:
  - 4 CPUs
  - 60 GB free disk space
  - 8 GB physical memory
- In addition to the management network and the optional customer network, there is a data network. This data network is used by the IBM Common Cloud Stack Quantum agent. Only IBM Common Cloud Stack can manage the Microsoft Hyper-V cloud environment, so we use this data network.

[2] See *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908 `http://www.redbooks.ibm.com/abstracts/redp4908.html`

3.3.2 Install the IBM SmartCloud Entry appliance on a Hyper-V node

This section covers deploying the SmartCloud Entry appliance to a Microsoft Hyper-V node.

Note: Although there are three appliances (one each for KVM, VMware, and Microsoft Hyper-V), it is necessary to deploy only one SmartCloud Entry V3.1 appliance. All three appliances have the same capabilities. A single appliance can manage multiple instances of all supported clouds. Therefore, a single appliance that is deployed either to a VMware cloud environment or a Microsoft Hyper-V environment, is sufficient to manage both environments. This paper describes the deployment process of both types of appliance.

Deploy the IBM SmartCloud Entry appliance

To deploy the SmartCloud Entry appliance, perform the following steps:

1. On the target node, create a `scevm` folder in the root directory.
2. Download the appliance installation file, and decompress it onto the target node. You can place it anywhere in the directory structure. In this example, we placed it in the root folder of the `C:\` drive. Figure 3-4 shows the display for the `scevm` folder and the decompressed folder with the appliance.

Figure 3-4 Folders for IBM SmartCloud Entry appliance

The decompressed folder displays as shown in Figure 3-5.

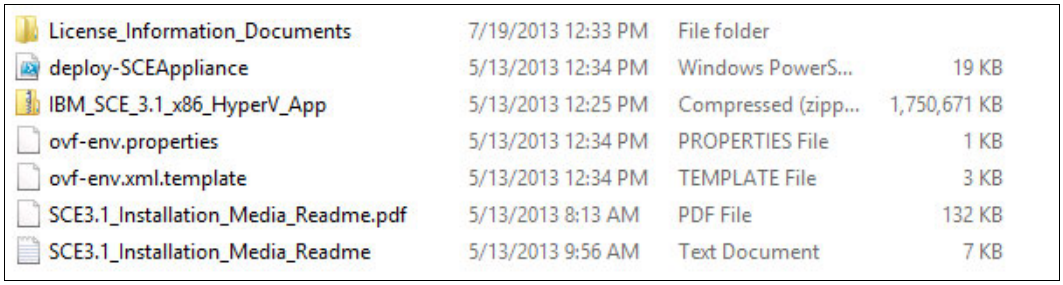

Figure 3-5 Contents of the appliance folder

Note: The `IBM_SCE_3.1_x86_HyperV_App.zip` file in Figure 3-5 is another compressed file that contains the appliance virtual machine files. Do not decompress this file. The installation program uses it to install the appliance.

3. Make a backup copy of the `ovf-env.properties` file, and then edit the file to set the properties that are used to configure the appliance. For more information about setting these properties, see the IBM SmartCloud Entry V3.1 Administrator Guide:

   `http://ibm.co/1c54U5v`

4. Run the **`deploy-SCEAppliance`** command to install the appliance, as this example shows:

   `.\deploy-SCEAppliance.ps1 -inputFile .\ovf-env.properties`

   When the appliance is installed, you can start it by using the Microsoft Hyper-V management interface.

3.3.3 Apply fix packs to the IBM SmartCloud Entry appliance

After you install the appliance, apply any relevant fix packs by completing these steps:

1. Download the fix pack from the IBM Fix Central website:

   http://www.ibm.com/support/fixcentral/

2. In the Fix Central window, select the following values in the fields indicated:

   a. Product Group: **Other Software**
   b. Other Software: **IBM SmartCloud Entry**
   c. IBM SmartCloud Entry: **3.1.0**
   a. Platform: **All**

The completed form displays as shown in Figure 3-6.

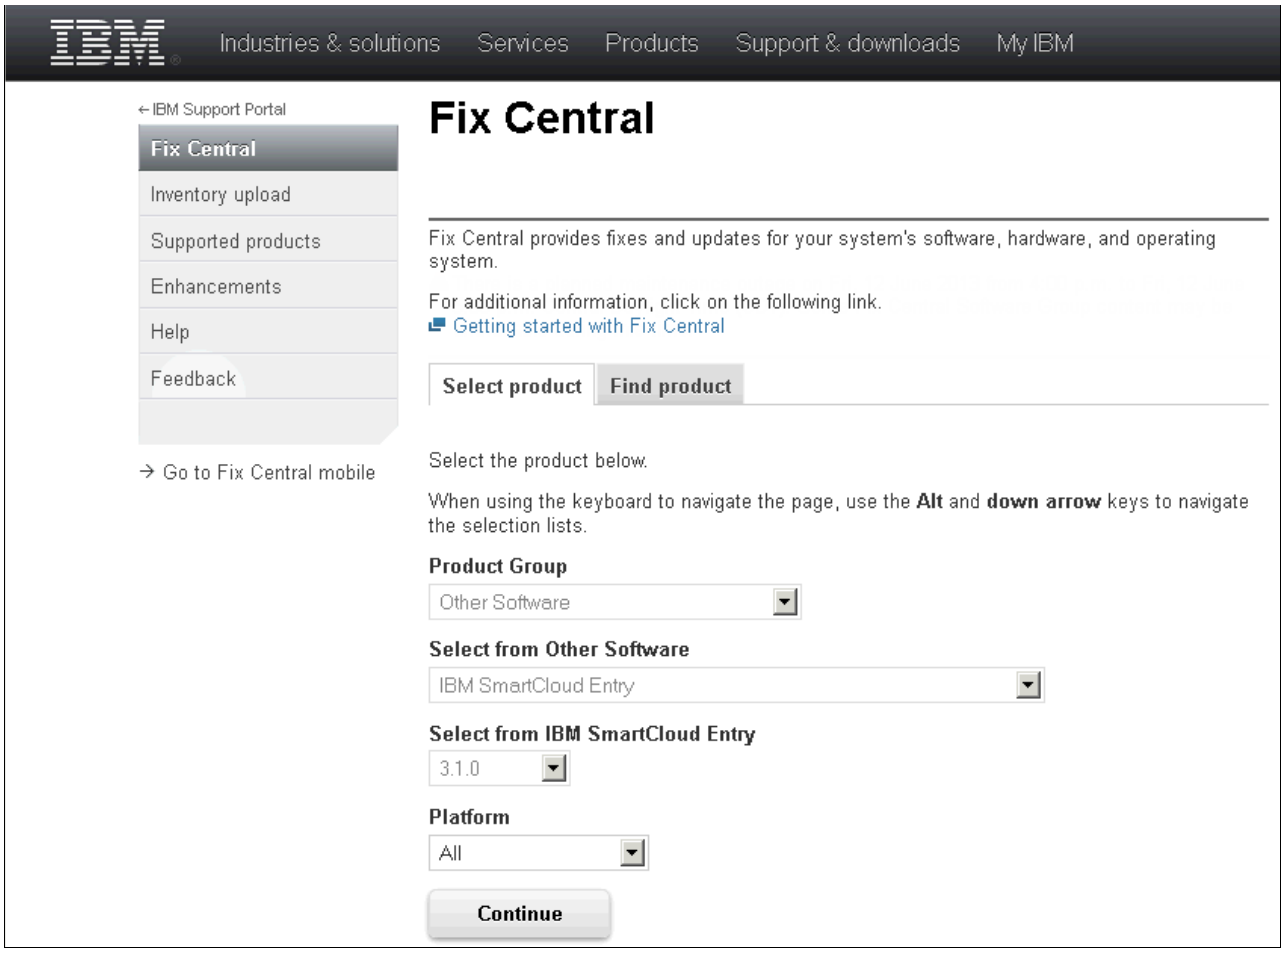

Figure 3-6 Completed form for IBM Fix Central

3. When you click **Continue** to submit the form, the Identify Fixes window displays.
4. Select **Browse for Fixes**, click **Continue**, and the Select Fixes window displays.

5. Select the needed fixes. We chose to download the 3.1.0.1 fix for appliances and the latest Hyper-V agent, as shown in Figure 3-7.

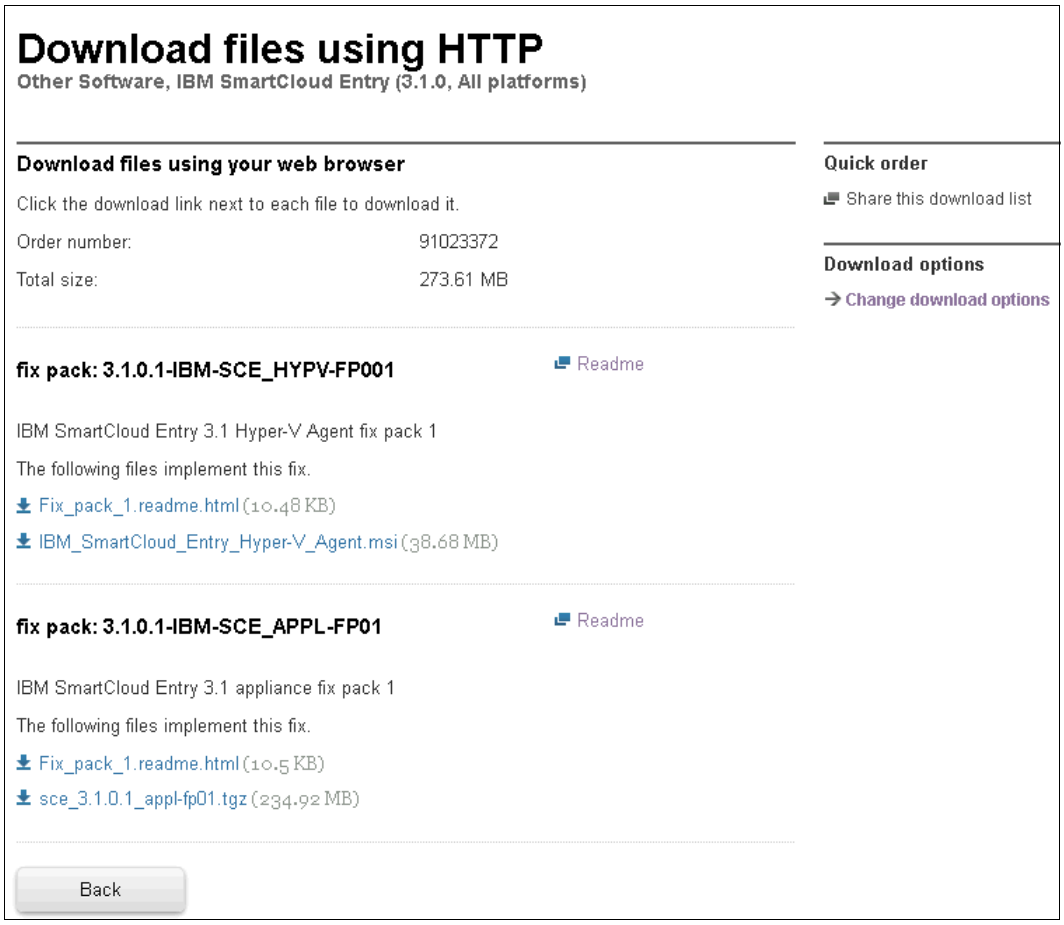

Figure 3-7 Selected fixes ready for download

6. Select **HTTP** or **Download Director** for the download method.

   Figure 3-7 indicates that we are downloading by using HTTP. If you prefer to use Download Director, click the **Change download options** link.

Note: You must have a browser that uses Java to use Download Director.

7. Transfer the files from your download directory to the SmartCloud Entry appliance `/tmp` directory. You can use Secure Copy Protocol (SCP) utilities to transfer the files.

   The `sce_3.1.0.1_appl-fp01.tgz` file is the fix pack to apply to the SmartCloud Entry 3.1 server. The `IBM_SmartCloud_Entry-Hyper-v_Agent.msi` file is the new version of the agent to install on the Microsoft Hyper-V node.

Note: Do not extract the `sce_3.1.0.1_appl-fp01.tgz` file. The fix pack installation utility needs that file in `.tgz` format.

8. On the SmartCloud Entry appliance, use the `sceappmgr` utility and follow these steps to install the fix pack:
   a. Log in to the SmartCloud Entry appliance as `sysadmin`.
   b. At the prompt, type `sceappmgr` to display a menu of options.
   c. Select option 5, **`Support and Maintenance`**, and the next menu of options displays.
   d. Select option 2, **`Install Fix Pack`**, and you are prompted for the path for the fix pack.
9. Enter the full path for the fix pack. For this example, this is the full path:

   `/tmp/sce_3.1.0.1_appl-fp01.tgz`
10. At the warning that displays, enter `1` (one) to continue the installation.
11. When the fix pack is installed, the SmartCloud Entry splash screen shown in Figure 3-8 displays.

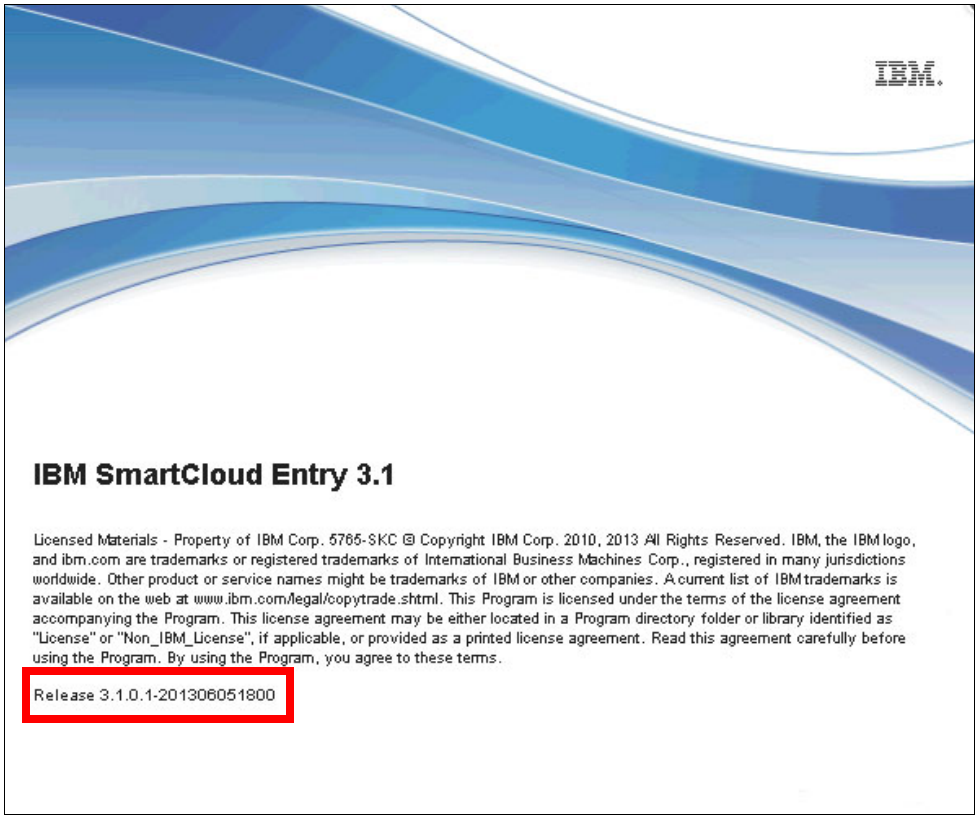

Figure 3-8 IBM SmartCloud Entry V3.1 splash screen

In this scenario, we updated SmartCloud Entry to release 3.1.0.1-201306051800.

3.3.4 Migrate data from SmartCloud Entry V2.4 to V3.1

After updating the SmartCloud Entry appliance to the newest version, we migrate our data. To do so, we follow these steps:

1. Copy the `.SCE24` directory from the `/home/sysadmin` directory on the SmartCloud Entry V2.4 appliance to the `/home/sysadmin` directory on the SmartCloud Entry V3.1 appliance. When this step is finished, the SmartCloud Entry appliance has both an `.SCE24` and a `.SCE31` directory in the `/home/sysadmin` folder, as shown in Figure 3-9.

```
IBM-SCE> ls -a
.               .bash_logout    .keystonerc                      .SCE24  .ssh
..              .bash_profile   lost+found                       .SCE31
.bash_history   .bashrc         migrateConfigAndDatabase.sh      sce.sh
```

Figure 3-9 Results of copying .SCE24 to the IBM SmartCloud Entry

2. On the SmartCloud Entry appliance, type `sceappmgr`, and a menu of options displays.
3. Select option **5**, **Support and Maintenance**, and the next menu of options displays.
4. Select option **3**, **Migrate from SmartCloud Entry 2.4**.
5. Select option **1** to confirm.
6. Select the language that you prefer, and the migration proceeds.
7. After the migration is finished, log in to SmartCloud Entry V3.1 to confirm that the data has migrated successfully.

Note: The migration sets the state for billing, metering, and so on to reflect the state of those functions in V2.4. For example, if billing was enabled in V2.4, the migration process enables billing for the SmartCloud Entry V3.1 appliance also, in addition to creating all of the V2.4 accounts.

3.3.5 Enable IBM Common Cloud Stack services

Next, we enable the IBM Common Cloud Stack services. This enables SmartCloud Entry to create IBM Common Cloud Stack clouds to manage Microsoft Hyper-V nodes.

IBM Common Cloud Stack is bundled with the SmartCloud Entry V3.1 appliance. IBM Common Cloud Stack cannot be migrated to another node, nor can SmartCloud Entry V3.1 interface with a version of OpenStack that is not IBM Common Cloud Stack. Enabling IBM Common Cloud Stack requires that you allocate the following resources to the appliance:

- 4 CPUs
- 60 GB free disk space
- 8 GB physical memory

Complete the following steps to enable IBM Common Cloud Stack:

1. Log in to the SmartCloud Entry V3.1 appliance as `sysadmin`.
2. On the SmartCloud Entry appliance, enter `sceappmgr`.
3. Select option **`1`** (**`Enable OpenStack Services`**).
4. When prompted, select option **`1`** to continue.

The IBM Common Cloud Stack services are now enabled and started.

3.3.6 Install the Microsoft Hyper-V agent on a managed node

IBM Common Cloud Stack interacts with Microsoft Hyper-V nodes by using an agent. The installer for the agent is included with the SmartCloud Entry software and is in Microsoft Windows installer (`.msi`) format. Check the IBM Support: Fix Central website to determine whether an updated agent installer is available:

`https://www.ibm.com/support/fixcentral/`

For more information, see 3.3.3, “Apply fix packs to the IBM SmartCloud Entry appliance” on page 30.

Note: Each Microsoft Hyper-V node must have the agent installed.

Although the agent is named the SmartCloud Entry Hyper-V agent, it actually interfaces with the instance of IBM Common Cloud Stack.

Figure 3-10 shows the Nova and Quantum agents installed on a Hyper-v node to enable manage-to support for Hyper-v. After the agent is applied, Hyper-v is referred to as a *compute node* from an IBM Common Cloud Stack perspective.

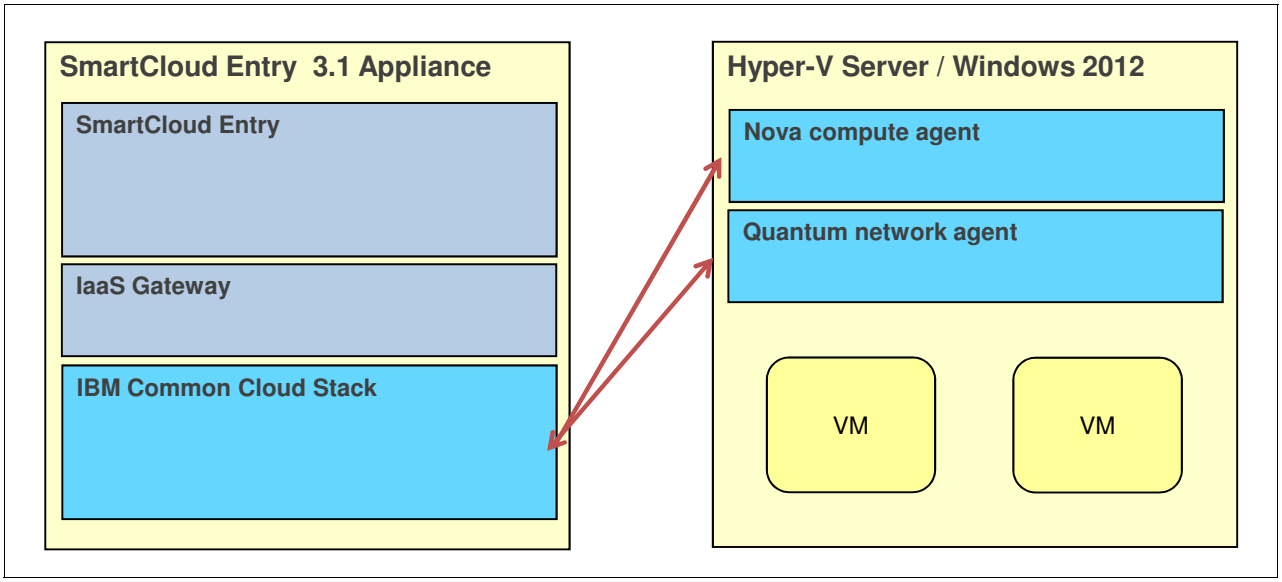

Figure 3-10 Nova and Quantum agents installed on a Hyper-V node

Verify the prerequisites

The following prerequisites must be met to install the Hyper-V agent on a managed node:

- Operating system requirements must be met, as described in Table 3-1 on page 35).

Table 3-1 Supported Windows and Hyper-V platforms

Operating system	Version	Notes
Windows Server 2012	Standard and Datacenter edition	► With the latest fix pack, both the GUI and the silent installation are supported ► Hyper-V Role must be enabled
Hyper-V Server 2012		► With latest fix pack ► Headless min-GUI ► Only silent installation is supported

- The SmartCloud Entry V3.1 appliance must be deployed, and the IP address of the SmartCloud Entry V3.1 appliance must be known.
- Confirm network connectivity between the Hyper-V system and the SmartCloud Entry appliance.
- Confirm that Network Time Service (NTP) is in sync with the SmartCloud Entry appliance (see OpenStack Foundation documentation at `http://docs.openstack.org/` and Windows help).
- Confirm that the person who is performing the installation is in the Admin group on the Hyper-V node.
- *Shared nothing* live migration requires additional setup, such as ensuring that the Hyper-V host is part of a domain.

Install and configure the agent

Complete the following steps to install and configure the Microsoft Hyper-V node:

1. Transfer the `.msi` file to the Microsoft Hyper-V node. Any standard method of file transfer is appropriate.
2. Navigate to the directory that contains the `.msi` file, and run the installer by using the following command:

   ```
   ./"IBM SmartCloud Entry Hyper-V Agent.msi"
   ```

3. After the installer runs, click **Next** on the Welcome window.
4. Accept the user license agreement.
5. Set the directory for the installation (the default directory is acceptable).

6. Select the **Complete** installation.

   The Nova Compute Configuration window that is shown in Figure 3-11 displays.

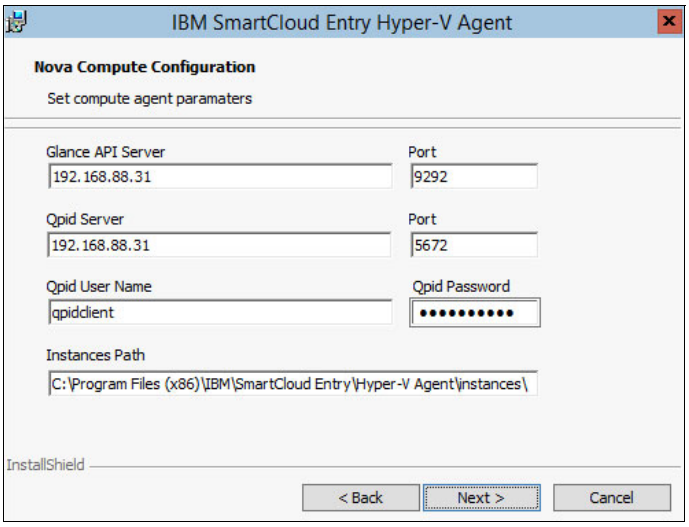

Figure 3-11 The Nova Compute Configuration window

7. In the Nova Compute Configuration window, enter these parameters:
   - Glance API Server: The IP address or DNS-resolvable name of the SmartCloud Entry server
   - Glance API server port: Leave as the default (`9292`)
   - Qpid server: The IP address or DNS-resolvable name of the SmartCloud Entry server
   - Qpid server port: Leave as default (`5672`)
   - Qpid user name: `qpidclient`
   - Qpid password: The assigned password (the default password is `openstack1`)
   - Instances path: Accept the default
8. Click **Next** to proceed to the Nova Compute Advanced Configuration window that is shown in Figure 3-12.

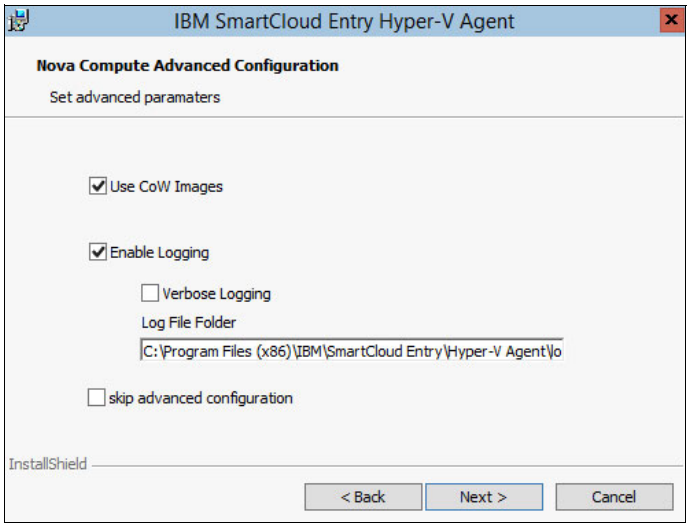

Figure 3-12 The Nova Compute Advanced Configuration window

9. Optional: Select **Use CoW Images** (that is, Copy on Write) and **Enable Logging**.

> **Note:** Selecting **Use CoW Images** can speed up image deployment.

10. Ensure that the **skip advanced configuration** box is clear.
11. Click **Next** to proceed to the Quantum Network Configuration window that is shown in Figure 3-13.

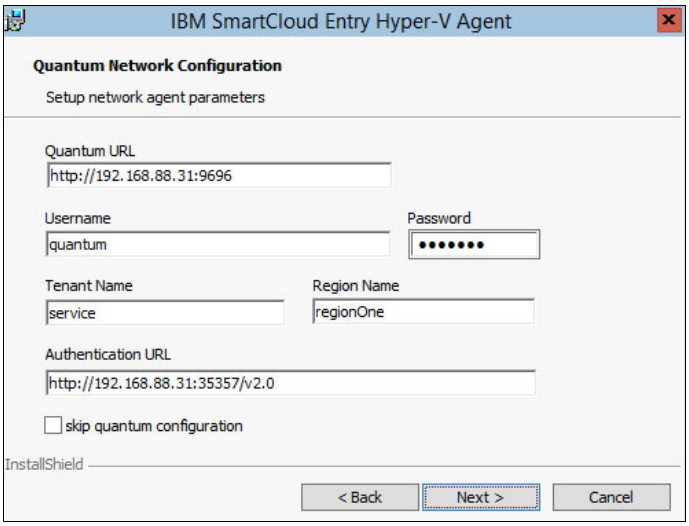

Figure 3-13 The Quantum Network Configuration window

12. Enter the following parameters in the Quantum Network Configuration window:
    - Quantum URL (follow this example to insert your server IP address or DNS name):

      `http://<IBM SmartCloud EntryServer IP Address or DNS Resolvable name>:9696`
    - Username: `quantum`
    - Password: quantum user password (the default password is `quantum`)
    - Tenant Name: `service`
    - Region Name: `regionOne`
    - Authentication URL (insert your server IP address or DNS name):

      `http://<IBM SmartCloud Entry server IP address or DNS-resolvable name>:35357/v2.0`
    - Ensure that the **skip quantum configuration** box is clear.

13. Click **Next** to proceed to the Hyper-V Live Migration settings window that is shown in Figure 3-14.

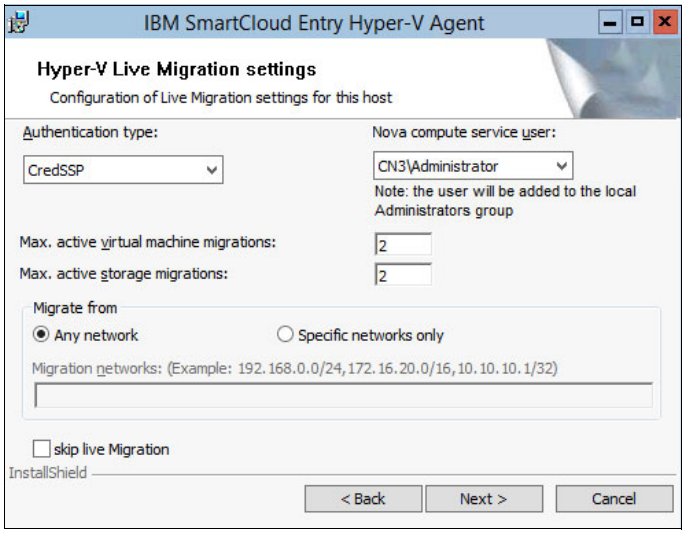

Figure 3-14 The Hyper-V Live Migration settings window

14. Your Microsoft Hyper-V node must be a member of a domain to configure live migration.

    If it is not, select the **skip live Migration** check box, and click **Next**.

    If your Microsoft Hyper-V node is already a member of a domain and you want to enable live migration, enter the following parameters:

    – Authentication type: **CredSSP**
    – Nova compute service user: The domain name of the user with the authorization to manage the SmartCloud Entry Hyper-V agent. This is typically an administrator or the user who installed the agent.
    – Maximum number of active virtual machine migrations: The default value is 2.
    – Maximum number of active storage migrations: The default value is 2.
    – Which networks to migrate from: If unsure, select the **Any Network** option.

15. Select **Next** to proceed to the first Virtual Switch Configuration window, which is shown in Figure 3-15.

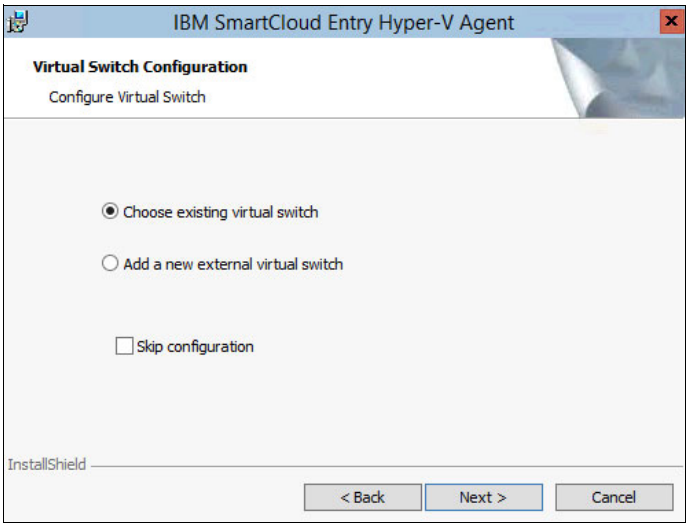

Figure 3-15 First Virtual Switch Configuration window

16. In that first Virtual Switch Configuration window, select either **Choose existing virtual switch** or **Add a new external virtual switch**. The typical selection is to **Choose an existing virtual switch**.
17. Ensure that the **Skip configuration** check box is clear.
18. Click **Next** to proceed to the second Virtual Switch Configuration window, which is shown in Figure 3-16.

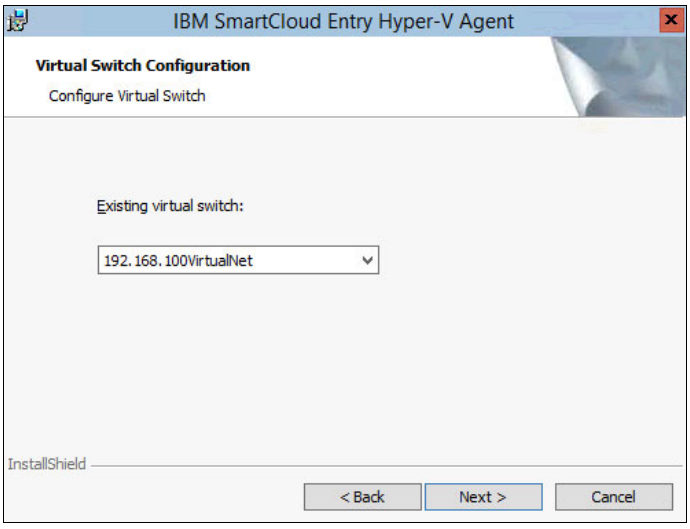

Figure 3-16 Second Virtual Switch Configuration window

19. In the second Virtual Switch Configuration window, select a virtual switch from the drop-down menu.
20. After you click **Next** to proceed to install the SmartCloud Entry Hyper-V agent, the installation process is complete.

Verify success of the installation

To verify that the SmartCloud Entry Hyper-V agent has installed successfully, log in to the SmartCloud Entry appliance as `sysadmin`, and run the **`nova hypervisor-list`** command. The output is similar to this example:

```
+----+---------------------+
| ID | Hypervisor hostname |
+----+---------------------+
| 1  | cn6                 |
+----+---------------------+
```

Note: You might receive an authentication error when you run the **`nova`** command. If you do, set the *OS_PASSWORD* environment variable to the Nova password (the default is `openstack1`) by using the following command:

```
export OS_PASSWORD=openstack1
```

Alternatively, pass the Nova password by using the **`--os_password`** argument, for example:

```
nova hypervisor-list --os-password openstack1
```

You have successfully integrated a Microsoft Hyper-V node with the IBM Common Cloud Stack environment that is running on the SmartCloud Entry server. Repeat this process for any additional Microsoft Hyper-V nodes that you want to add to your managed cloud.

3.3.7 Create and manage a Hyper-V cloud in IBM SmartCloud Entry

Now that you have enabled IBM Common Cloud Stack and integrated Microsoft Hyper-V nodes into the IBM Common Cloud Stack, you can manage this cloud by using SmartCloud Entry.

Follow these steps to manage the Microsoft Hyper-V cloud:

1. Log in to the SmartCloud Entry self-service GUI as `admin`.
2. Navigate to the Configuration tab, and select **Clouds**.

3. Click the **New Cloud** icon to create a cloud.

   The Add Cloud Configuration window displays, as shown in Figure 3-17.

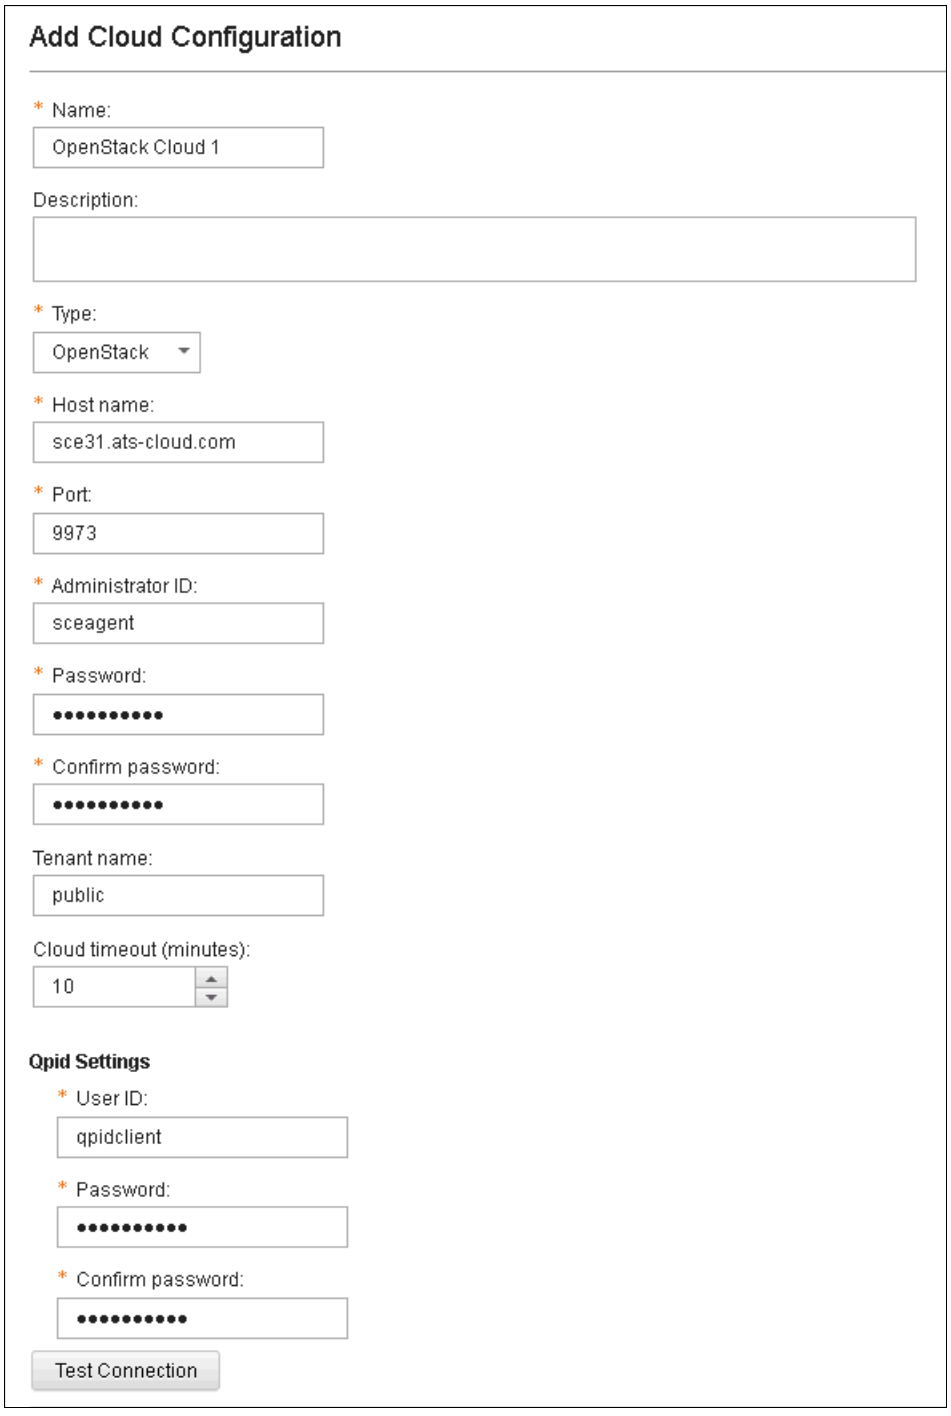

Figure 3-17 The Add Cloud Configuration window for a Microsoft Hyper-V cloud

4. Enter the following parameters:
   - Name: The name for the cloud as it displays in the self-service GUI
   - Description: An optional description of the cloud
   - Type: `OpenStack`
   - Host Name: Enter either the IP address or the DNS-resolvable name of the SmartCloud Entry appliance
   - Port: `9973`
   - Administrator ID: `sceagent`
   - Password and Confirm Password fields: The password for the `sceagent` user (the default value is `openstack1`)
   - Tenant: `public`
   - Cloud timeout: Accept the default of 10 minutes
   - Qpid settings:
     - User ID: `qpidclient`
     - Password and Confirm Password fields: The qpidclient password (the default value is `openstack1`)
5. Click **Test Connection**, and you receive a notice that the connection is successful. If you do not receive the notice, double-check these items:
   a. Ensure that you set the parameters correctly.
   b. Ensure that the passwords are correct and user names are spelled correctly.
   a. Ensure that the IBM Common Cloud Stack processes are running by using **`scappmgr`**.
6. When the connection test is successful, click **Add Cloud** to add the cloud.

You are now ready to import an image into the IBM Common Cloud Stack cloud. For more information about this process, see Chapter 4, “Image management” on page 43.

4

Image management

With IBM SmartCloud Entry software, you can create and store images on clouds that are created with IBM Common Cloud Stack software.

In this chapter, we define some of the key image management concepts for SmartCloud Entry software, such as activation processes, strategies, and deployment. Then we describe the steps for building an image.

4.1 Building images

Several products are available for building images, including the IBM Image Construction and Composition Tool, VMware Studio, SUSE Studio, and open source OpenStack images. All of these provide an extensive set of functions and features. The drawback is that users typically must have extensive knowledge of the particular software to fully use the functions and features.

It might be to your advantage to build images by using IBM SmartCloud Entry, primarily because it is easy to use. When you build images manually with this software, deployment can be faster for tasks that typically require extensive product knowledge.

Next, we describe how easy it is to use IBM SmartCloud Entry to build an image manually and then activate and deploy the image.

4.1.1 Build an image manually

In this scenario, we describe the process of manually building an image for deployment. We use Microsoft Windows 2008 workstation. The goal is to build a Microsoft Windows Server 2012 image with activation software for eventual deployment on an IBM SmartCloud Entry Version 3.1 Hyper-V cloud.

Ensure that you have the installation media for Microsoft Windows 2012 available, and then complete the following steps:

1. Log in to IBM SmartCloud Entry as administrator.
2. From the Disk Management panel, click **Create a Virtual Hard Disk (VHD)**, and then format and initialize the media.
3. Use the **imagex** command to apply the installation to the newly created virtual hard disk (VHD).
4. After the image is created in VHD format, boot the image and run the Microsoft Windows **sysprep** command. For more information about Sysprep, see “How to use the Sysprep tool...” on the Microsoft Support website:

   http://support.microsoft.com/kb/302577

For more information about building images, see *Chapter 5. Creating images manually*, on the OpenStack.org website:

http://docs.openstack.org/grizzly/openstack-image/content/ch_creating_images_manually.html

4.1.2 Import images

IBM SmartCloud Entry can import and delete images, although only IBM Common Cloud Stack images are supported.

To import images from IBM SmartCloud Entry, complete the following steps:

1. In IBM SmartCloud Entry, select **Images** → **more** → **Import Image** to display the Import Image window shown in Figure 4-1.

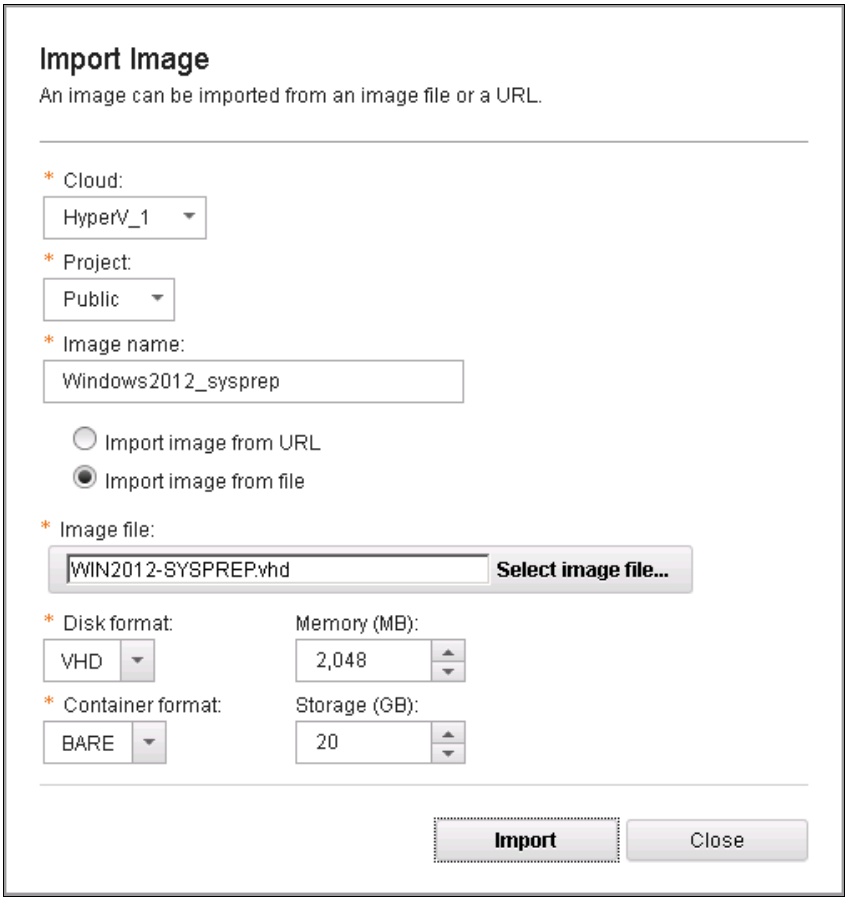

Figure 4-1 IBM SmartCloud Entry Import Image window

2. Specify this information:
   - Name of the cloud
   - Project to which the image will be imported
   - Image name
   - Image source (either URL or local)

Note: IBM SmartCloud Entry V3.1 supports only the VHD disk format for the Hyper-V hypervisor, and the following guest operating systems and platforms:

- Microsoft Windows 8
- Microsoft Windows 7
- Microsoft Windows 2012
- Microsoft Windows 2008 R2
- Red Hat Enterprise Linux 6.4
- SUSE Linux Enterprise Server 11.2

The container format is BARE. For more information, see *Disk and Container Formats for Images* on the OpenStack.org website:

http://docs.openstack.org/grizzly/openstack-compute/admin/content/image-formats.html

Hyper-V images cannot be decreased in size, so we suggest providing the minimum amount of memory and storage that the image consent requires for deployment.

At deployment, you can also assign a *flavor*, which is a new feature in SmartCloud Entry V3.1. With flavors, you can dynamically allocate the amount of RAM, CPU, and storage. For more information, see 4.2, “Image flavors” on page 56.

4.1.3 Configuration strategies

Configuration strategy is a feature of IBM Common Cloud Stack. This feature is used for activation of a deployed image. It is similar to open source OpenStack configuration drive support with cloud-init software, an early initialization tool.

For more information about configuration strategies, see the IBM SmartCloud Entry V3.1 Administrator Guide:

http://ibm.co/1c54U5v

SmartCloud Entry supports two main types of configuration plug-in strategies:

- Open Virtualization Format (OVF)
- Microsoft Windows System Preparation (Sysprep)

The image must be built with the correct software package for the configuration strategy to automatically access and apply the customizations to the deployed image. IBM Virtual System Activation Engine is used for OVF deployment, and Microsoft Sysprep is used for Sysprep deployment.

After the image is deployed with one of these configuration methods, this pluggable configuration strategy feature acknowledges the customizations and how they were made available. The customizations that are provided by IBM Common Cloud Stack come from two sources, server and user metadata:

- Server metadata that is provided by the OpenStack products, as shown in Example 4-1. For a full list, see the SmartCloud Entry V3.1 Administrator Guide:

  http://ibm.co/1c54U5v

  See the the example that follows.

Example 4-1 Server metadata sample

```
server.hostname
server.domainname
server.dns-client.pri_dns
server.network.[n].mac
server.network.[n].[v4|v6].address
server.network.[n].[v4|v6].netmask
server.network.[n].[v4|v6].cidr
server.network.[n].[v4|v6].gateway
```

- Server metadata that is provided by the user who deploys the image. The Sysprep or OVF configuration strategy includes the following parts:

  Type (required)
  Choose either OVF or Sysprep.

  Template (required)
  This refers to the OVF descriptor or `unattend.xml` file (to feed Sysprep). Although copying and pasting content to the template box is allowed, the preference is to import Sysprep from the XML file.

  Mapping (required)
  This defines how to map the server metadata that is provided by both the OpenStack product and the user who deploys the image to the appropriate elements or parts of the template. This is a JavaScript Object Notation (JSON) array.

  User metadata (optional)
  This determines how the deployment properties for the configuration strategy are defined and displayed in IBM SmartCloud Entry V3.1 and is a JSON array. If there is no metadata, the display reads, “Please provide a user password.”

To add, edit, or delete configuration strategies from IBM SmartCloud Entry, follow these steps:

1. Navigate to the Images tab, and select the image.
2. Select **More** → **Configuration Strategy**. Figure 4-2 on page 48 shows the Configuration Strategy window for this image.

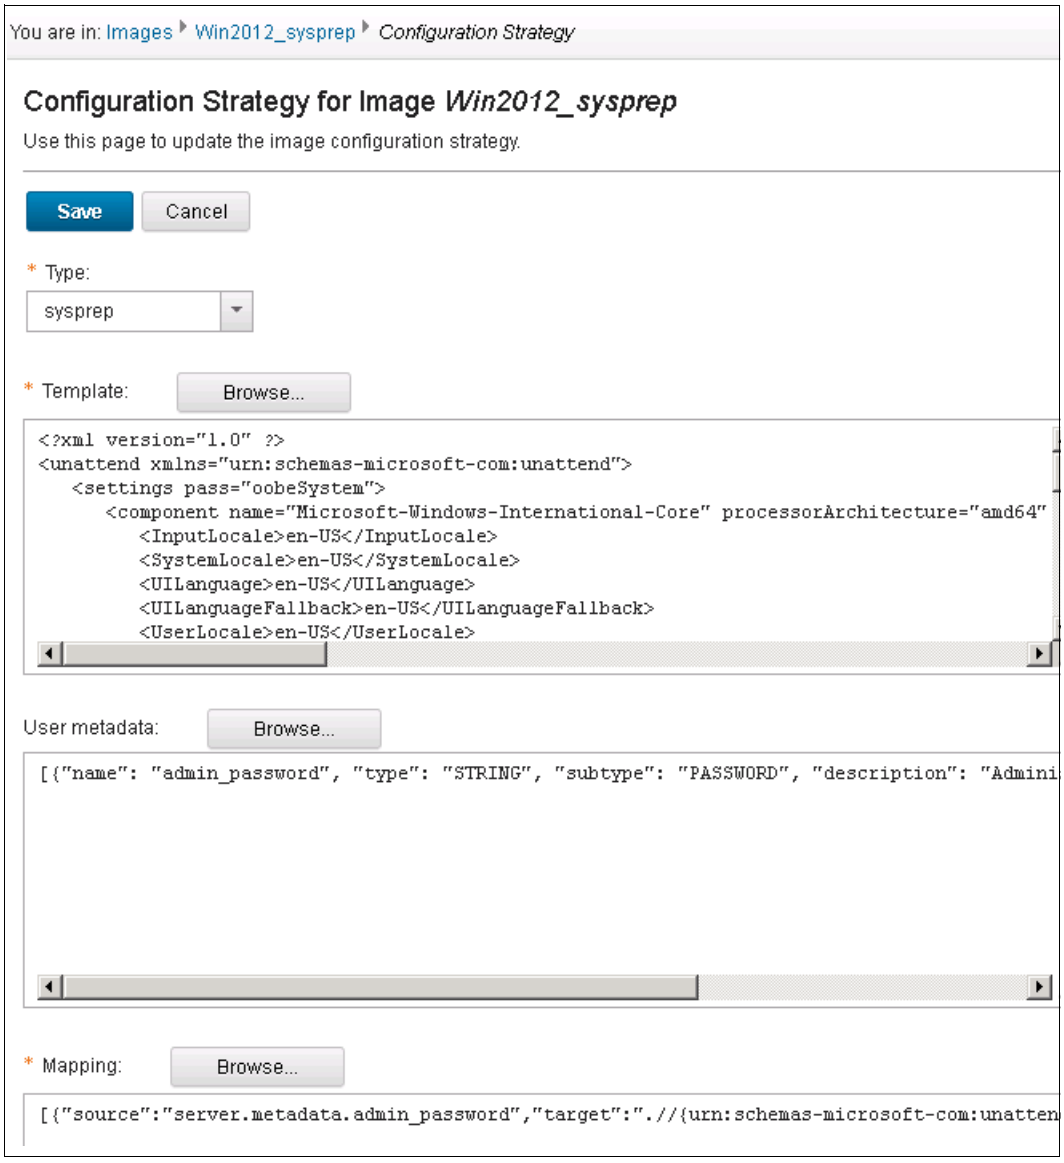

Figure 4-2 Defining the configuration strategy

Example 4-2 shows the user metadata to use when you deploy an image with Sysprep as the configuration strategy.

Example 4-2

```
[{"name": "admin_password", "type": "STRING", "subtype": "PASSWORD",
"description": "Administrator Password", "required" : "true"}]
```

Example 4-3 on page 48 shows the mapping to use when you deploy an image with Sysprep as the configuration strategy.

Example 4-3

```
[{"source":"server.metadata.admin_password","target":".//{urn:schemas-microsoft-co
m:unattend}AdministratorPassword/{urn:schemas-microsoft-com:unattend}Value"}]
```

To use the cloud-init early initialization tool on your images, follow these steps:

1. Install the cloud-init early initialization tool on the bare image.
2. Customize the configuration file.
3. Provide user data when you are deploying the image, as shown in Figure 4-3.

   Figure 4-3 shows an image with cloud-init user data deployment, a specific key pair, and a personality file to be deployed in the `/tmp/personality_1` path.

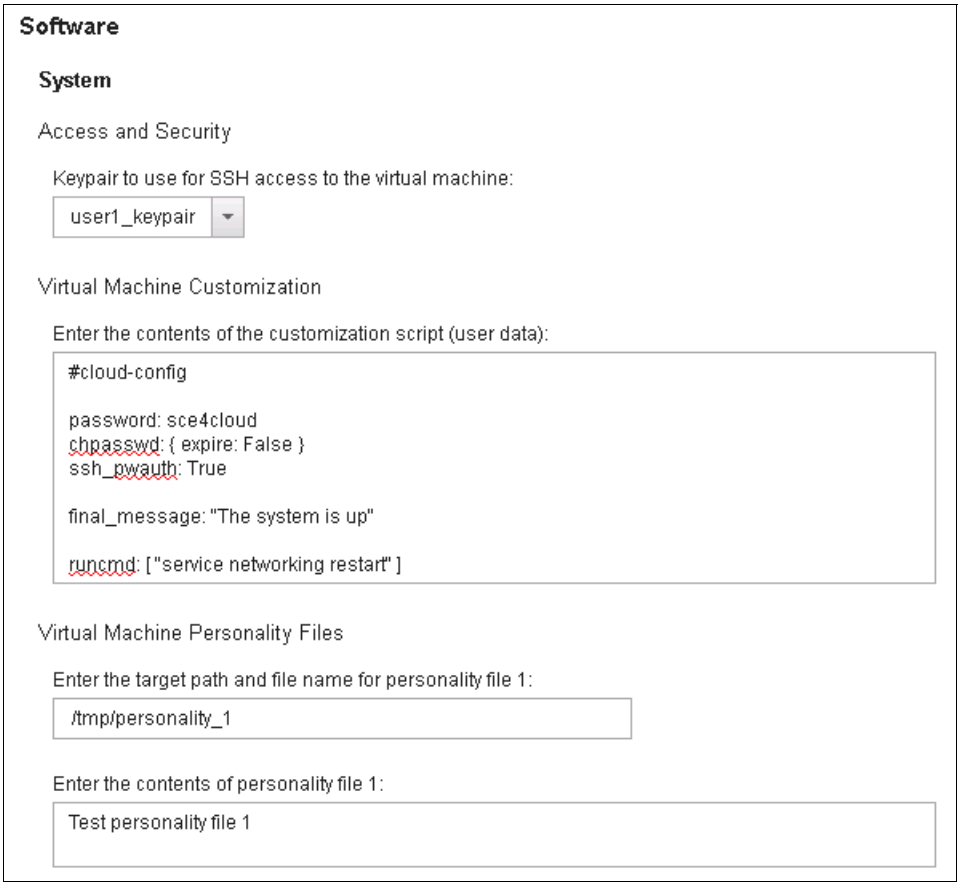

Figure 4-3 Image deployment by using the cloud-init early initialization tool

4.1.4 Image activation

IBM Common Cloud Stack relies on image activation software to start and customize the image, and provides data on a configuration drive (config drive) to allow for image activation. The most common way to activate IBM Common Cloud Stack images is to use the cloud-init early initialization tool.

You might need to rebuild your images before you use the cloud-init tool. IBM SmartCloud Entry V3.1 can activate images by using IBM Virtual System Activation Engine (which is included with VMControl Prepare for Capture, the IBM Image Construction and Composition tool, and IBM Informix® Virtual Appliance Factory Service). Alternatively, you can use Microsoft Sysprep, which is useful when you have already built images with these activation software packages, or when the administrator has limited experience with cloud-init.

SmartCloud Entry V3.1 supports the following types of image activation software:

- Cloud-init: This tool is common to both OpenStack products, and it is used for a variety of tasks, including image activation.
- Microsoft Windows cloud-init
- IBM SmartCloud init
- IBM Virtual System Activation Engine (VSAE): This tool is included with VMControl Prepare for Capture in VM Control, the IBM Image Construction and Composition Tool, and the IBM Informix Virtual Appliance Factory Service.
- Microsoft Sysprep: This is useful when you have images already built with these activation software packages or when the administrator has limited experience with cloud-init.

Note: Do not use two activation software packages on the same images. Doing so can cause discrepancies among the images.

Figure 4-4 on page 50 illustrates some of the differences between the activation software packages. Figure 4-4 on page 50 shows four images, and each has h particular activation software:

- Image A uses cloud-init software for activation.
- Images B and C have metadata (configuration strategy) assigned, which is stored on IBM SmartCloud Entry.
- Image D has no activation software.

As the bottom row in Figure 4-4 on page 50 shows, four instances were deployed.

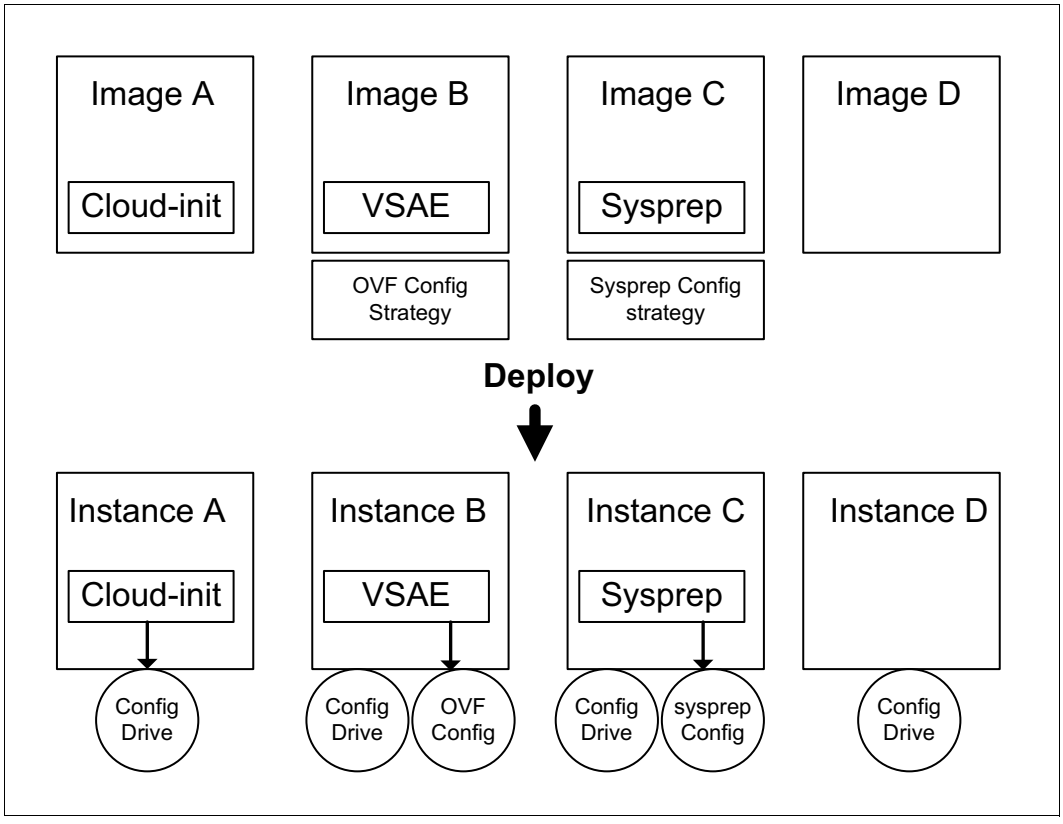

Figure 4-4 Diagram containing different activation software

Each of the four instances that were deployed has a configuration drive. SmartCloud Entry requires that there be a configuration drive for every image that is deployed, even though they do not use cloud-init nor any activation software. The details about the instances in Figure 4-4 make this information clearer:

- Instance A uses cloud-init, so the configuration drive is read by cloud-init, and the customizations are applied.
- Instance B has a configuration drive installed, but the drive does not read the instance. Instead, Instance B reads the OVF configuration drive. The same situation happens for instance C, except that Instance C reads the Sysprep configuration drive.
- Instance D has a configuration drive that is deployed by IBM Common Cloud Stack, but no activation software is installed to read the drive. No customizations are applied.

The configuration drive can be accessed by any guest operating system that can mount an ISO 9960 file system. The configuration drive contains these contents:

- User data
- Personality files
- Secure Shell (SSH) key pair
- Network information (for static networks)

For information about other server metadata, see the open source OpenStack configuration drive documentation on their Config drive page:

`http://docs.openstack.org/grizzly/openstack-compute/admin/content/config-drive.html`

When you are deploying an image by using OVF format or the Sysprep configuration strategy, IBM Common Cloud Stack prompts you to identify the type of configuration strategy. This can be accessed from the guest operating system, similarly to how the configuration drive that is used by cloud-init is accessed. The process for accessing a config drive that uses cloud-init is shown in Figure 4-2 on page 48 and in Figure 4-4 on page 50.

The OVF configuration strategy drive contents usually include an ovf-env.xml file with configuration strategy mappings applied. The Sysprep configuration strategy drive contents include an unattend.xml file, also with configuration strategy mappings applied.

Figure 4-5 shows two ISO 9960 mounted drives on the newly deployed instance. Drive `D:\` is the config drive (`config-2`). Drive `E:\` is the Sysprep configuration drive.

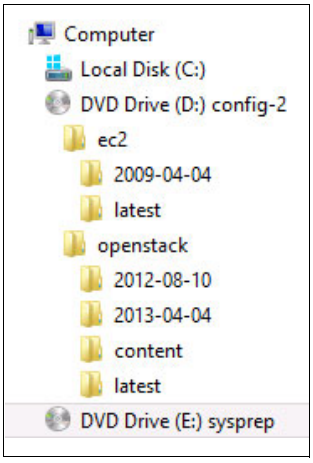

Figure 4-5 Image configuration drive contents (Sysprep)

4.1.5 Image copy and multi-instance deployment

IBM SmartCloud Entry V3.1 introduces the ability to copy image settings across multiple projects without having to import the entire image into the cloud.

The following details about image copying are helpful to know:

- The configuration settings that are applied to the image are copied to the deployed image.
- Image settings are copied and stored separately from the original base image, which also includes the project.
- Deleting the copied image does not affect the original image.
- Deleting the original image also deletes all copies.
- Images can be copied multiple times.

Copy a single image

Only a few steps are necessary to copy an image in IBM SmartCloud Entry V3.1:

1. In the Images window, select the image name.
2. Choose the image that you want to copy. Provide a name and a description for the copied image, and then click **Copy**, as shown in Figure 4-6.

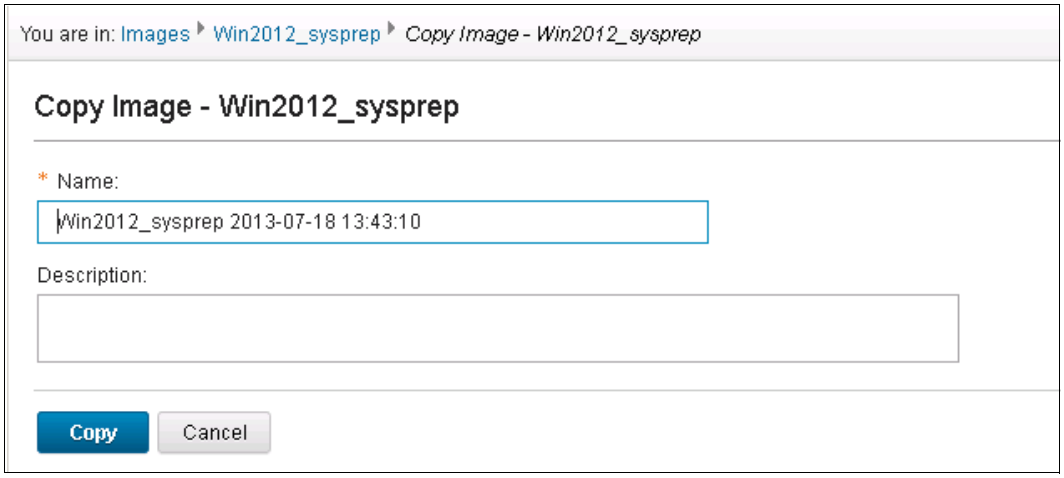

Figure 4-6 IBM SmartCloud Entry V3.1 Copy Image window

Figure 4-7 shows how information about a copied image displays. The Base images and Related images fields indicate that the image that was deployed is not the original image.

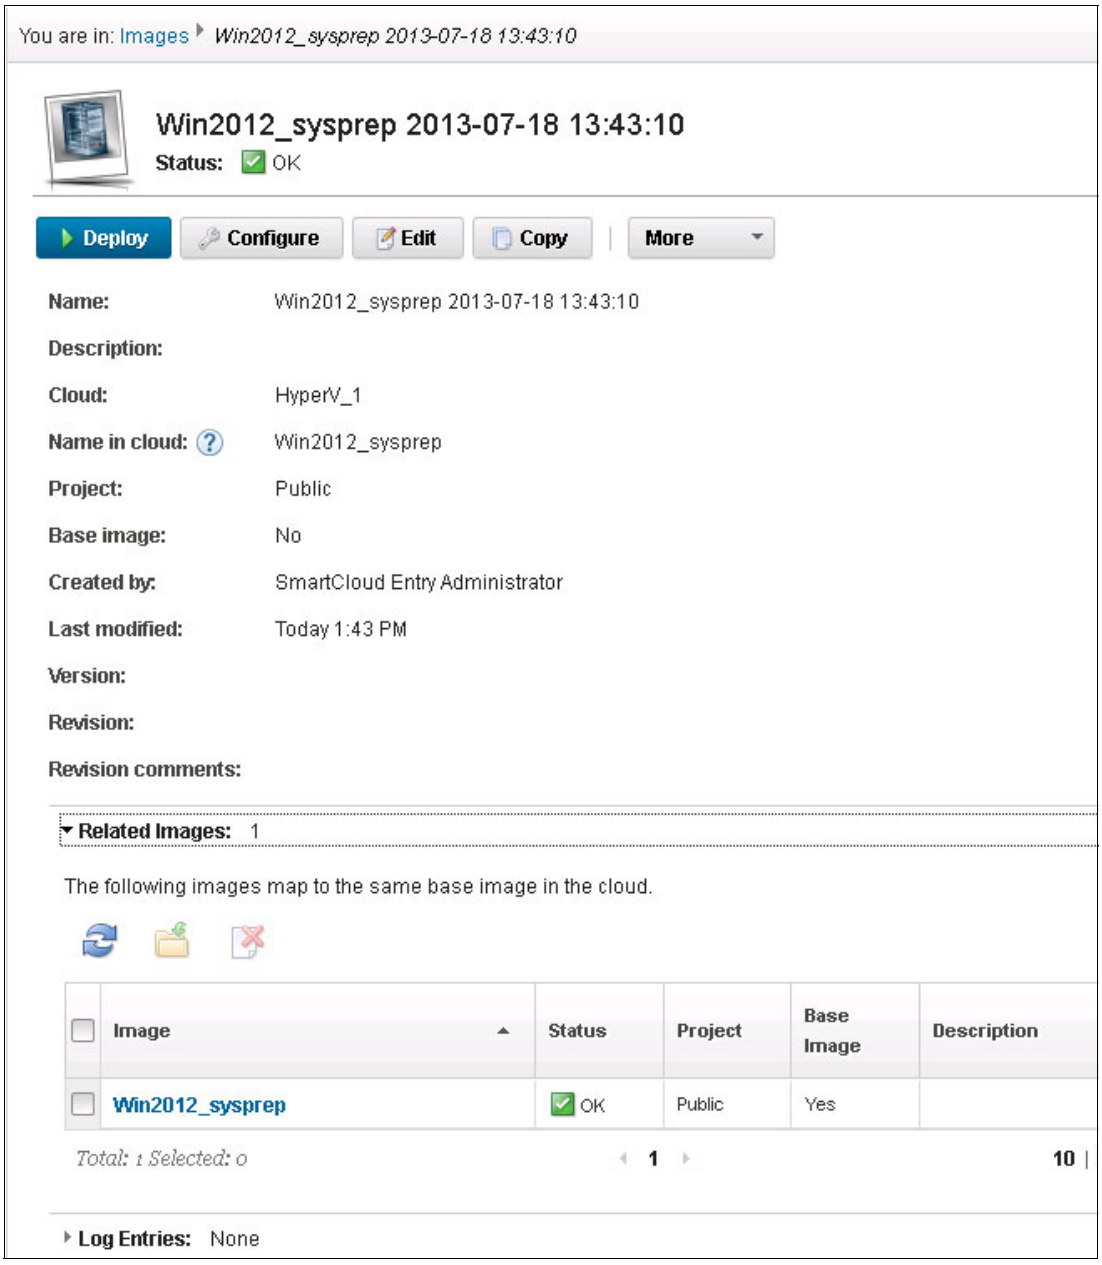

Figure 4-7 IBM SmartCloud Entry V3.1 copied image window

Deployment of multiple instances

IBM SmartCloud Entry can deploy multiple instances of an image from one image. Figure 4-8 on page 54 shows the image deployment window. The Instances field shows that up to five instances can be set for deployment simultaneously.

Consider the following details about multi-instance deployment:

- A *-X* is added to each base name (where *X* is the instance number). Aside from this, creating an instance is the same as carrying out separate deployments (see Figure 4-9 on page 55).
- A maximum of five instances can be simultaneously deployed by setting the value in the SmartCloud Entry GUI, as shown in Figure 4-8. You can increase the number by modifying the deployment.properties file in the SmartCloud Entry appliance.

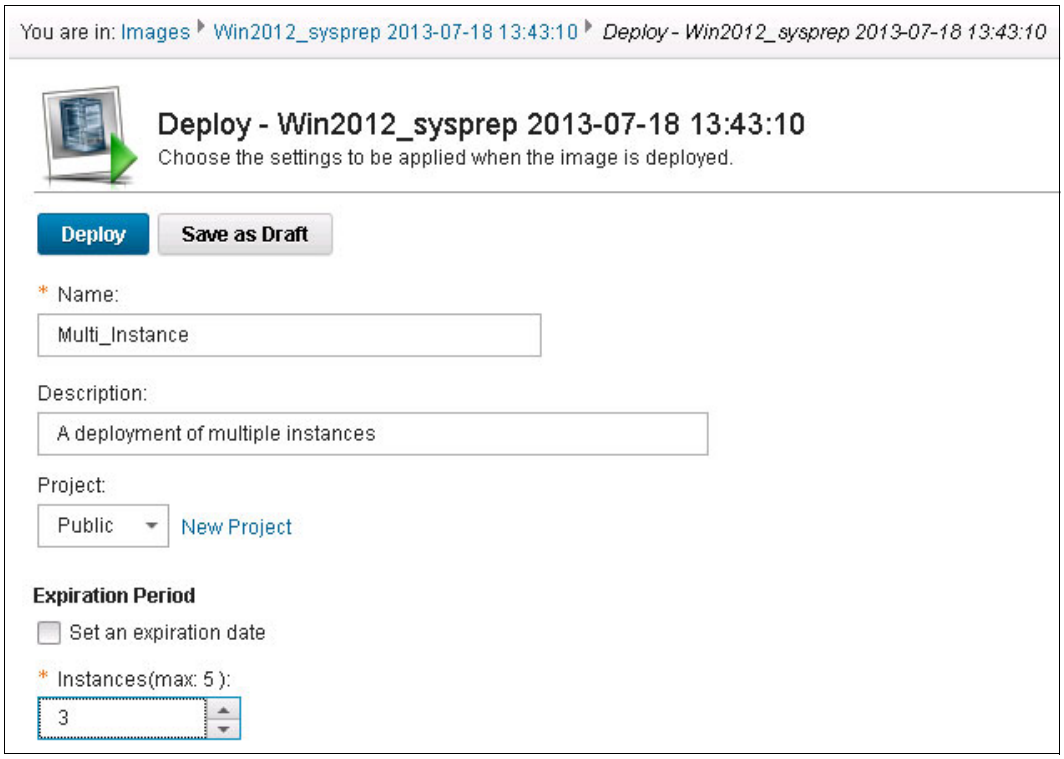

Figure 4-8 Deploying multiple images by using IBM SmartCloud Entry V3.1

- No relationship is maintained between multiple instances. They are separately managed (start, stop, delete, and so on).
- The multi-instance deployment feature is enabled by default. It can be disabled by using the **`com.ibm.cfs.deployments.multi.enabled`** property in the `deployment.properties` file.

Figure 4-9 shows the three instances that are deployed at the same time.

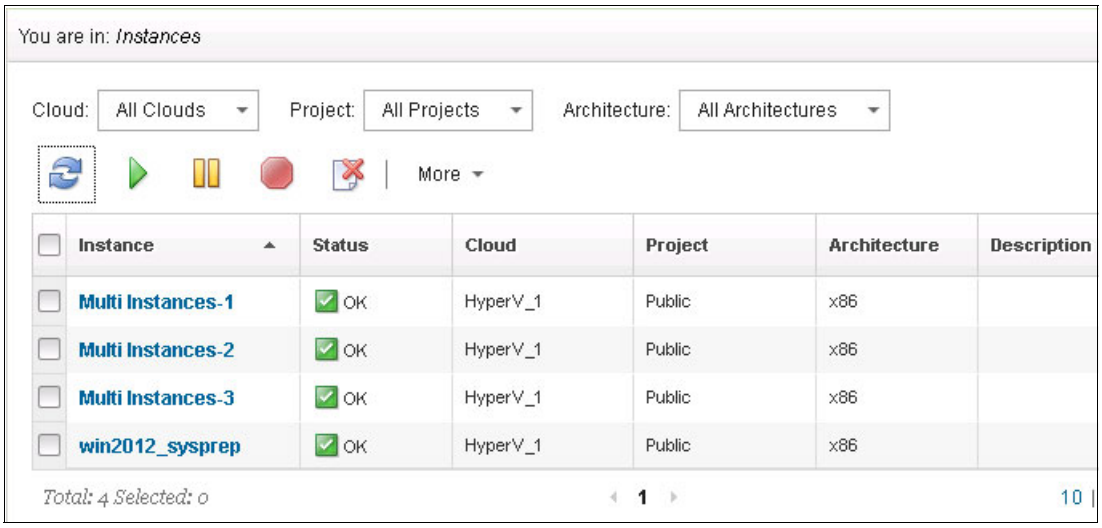

Figure 4-9 Instances window, which indicates that multiple instances can be deployed at one time

Figure 4-10 shows the images that are deployed on the Hyper-V Manager panel for the node.

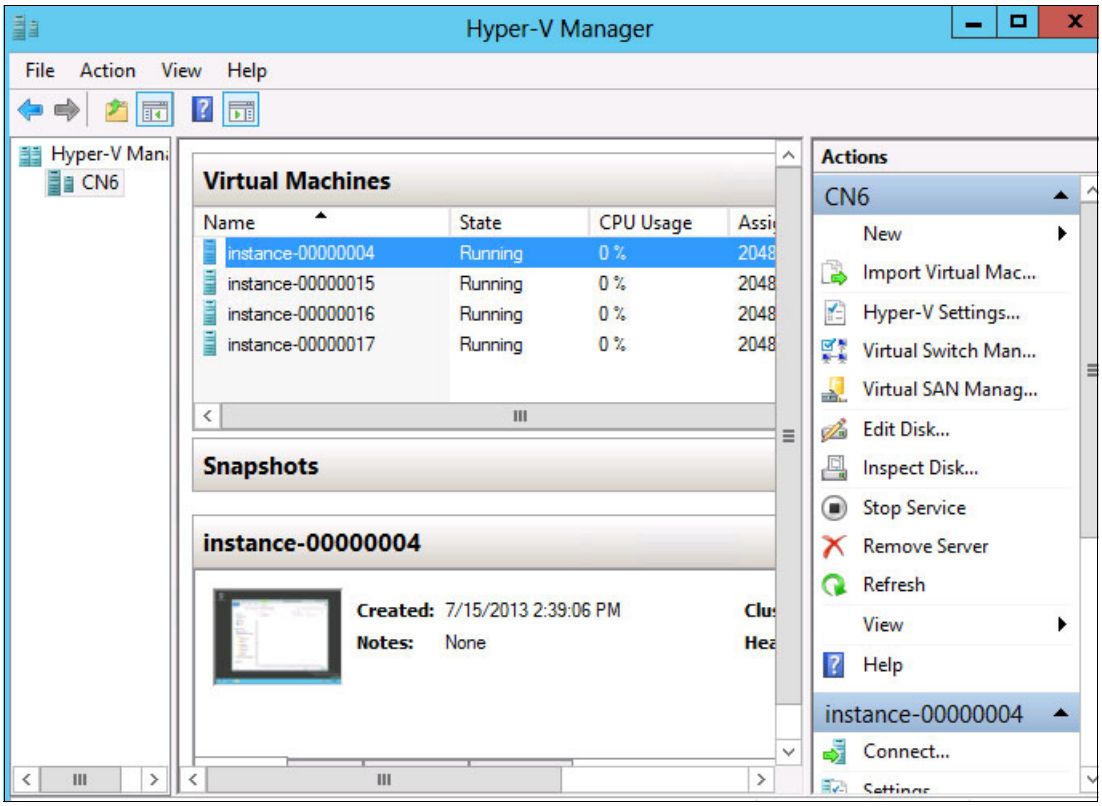

Figure 4-10 Instances that are deployed on the Microsoft 2012 Server (Hyper-V)

4.2 Image flavors

When you are setting the image to deploy, you can dynamically allocate the amount of RAM, CPU, and storage. This is known as setting a *flavor*. Figure 4-11 shows the flavors that are predefined in SmartCloud Entry V3.1, ranging from *tiny* to *extra large* (`xlarge`).

.

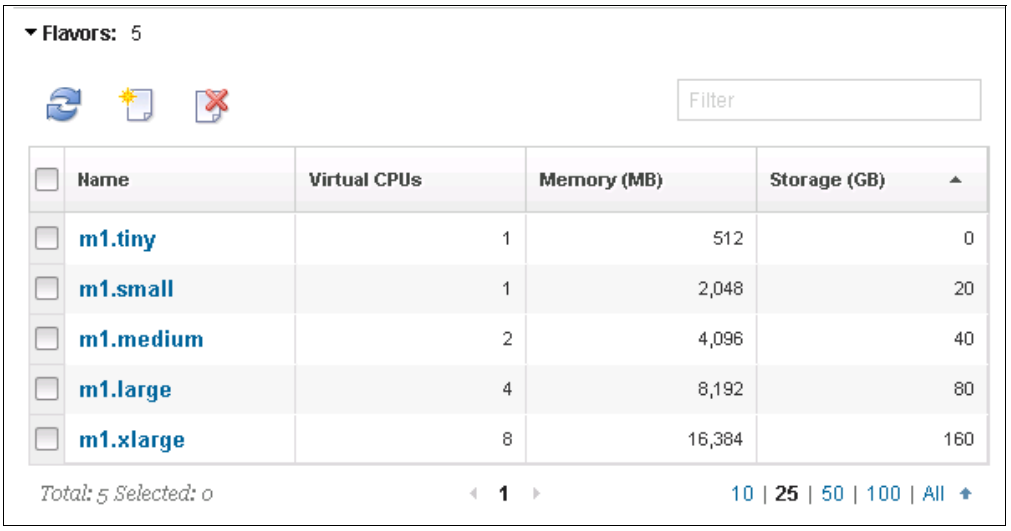

Figure 4-11 Image flavors available by default on IBM SmartCloud Entry V3.1

4.2.1 Create a flavor

You can create a flavor to use your physical and virtual resources more effectively by following these steps:

1. Click **Clouds** → **OpenStack cloud** → **Edit Cloud**.
2. Under Flavors, click the **New** icon.
3. Click **Save**.

Figure 4-12 shows the New Flavor window.

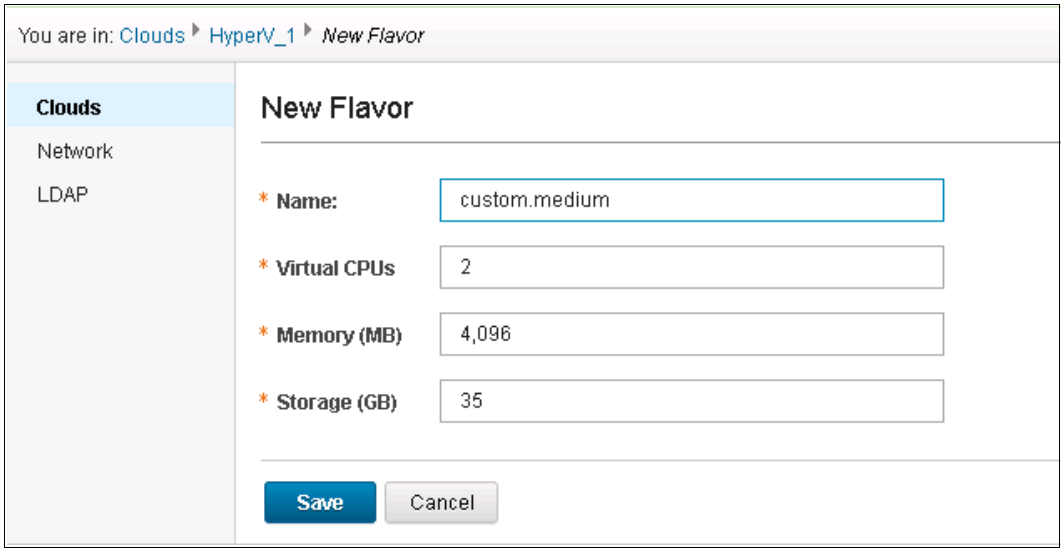

Figure 4-12 IBM SmartCloud Entry V3.1 custom flavor creation window

For more information about flavors, see the Instance Building Blocks section in the *OpenStack Compute Administration Guide - Grizzly, 2013.1*:

http://docs.openstack.org/grizzly/openstack-compute/admin/content/instance-building-blocks.html

4.3 SSH key management

With IBM SmartCloud Entry, you can generate or import an SSH key pair to access your newly deployed instance or instances. We begin with the steps for creating and importing an SSH key pair. Then, we describe how key pair management works (see 4.3.3, "Use IBM SmartCloud Entry key pairs" on page 59).

4.3.1 Create an SSH key pair

In IBM SmartCloud Entry, complete the following steps to create an SSH key pair:

1. Click **Menu** → **Access** → **Key Pairs** → **Create** icon.
2. Assign a name to the key pair, and click **Save**.

Figure 4-13 shows the New Key Pair window. IBM SmartCloud Entry stores the public key and provides the private key for download. The private key is required when you are connecting to the instance that you deployed.

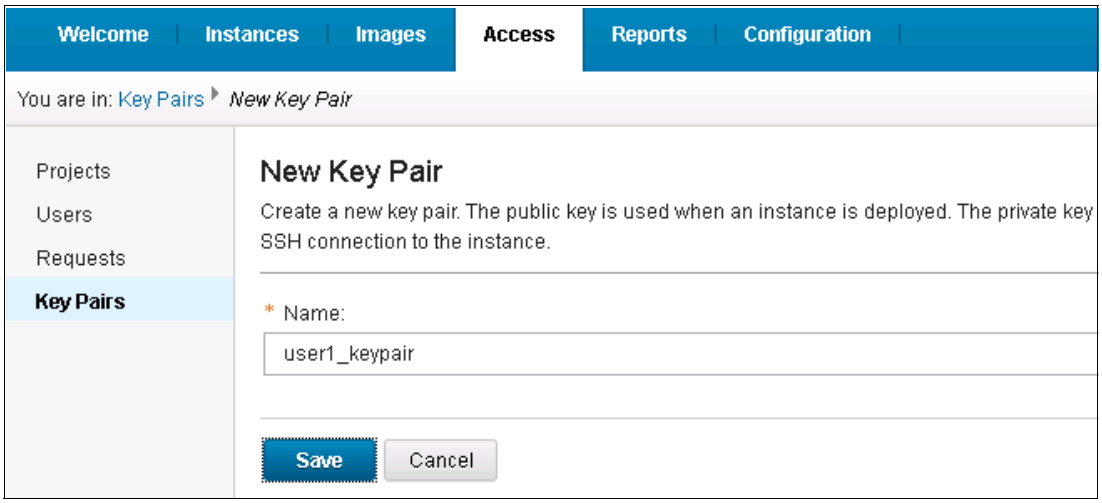

Figure 4-13 Creating an SSH key pair

4.3.2 Import an SSH key pair

If you have a private key, you can import it into IBM SmartCloud Entry by completing the following steps:

1. Click **Menu** → **Access** → **Key Pairs** → **Import** icon.
2. Select the key pair to import.

Figure 4-14 on page 58 shows the Import Public Key window.

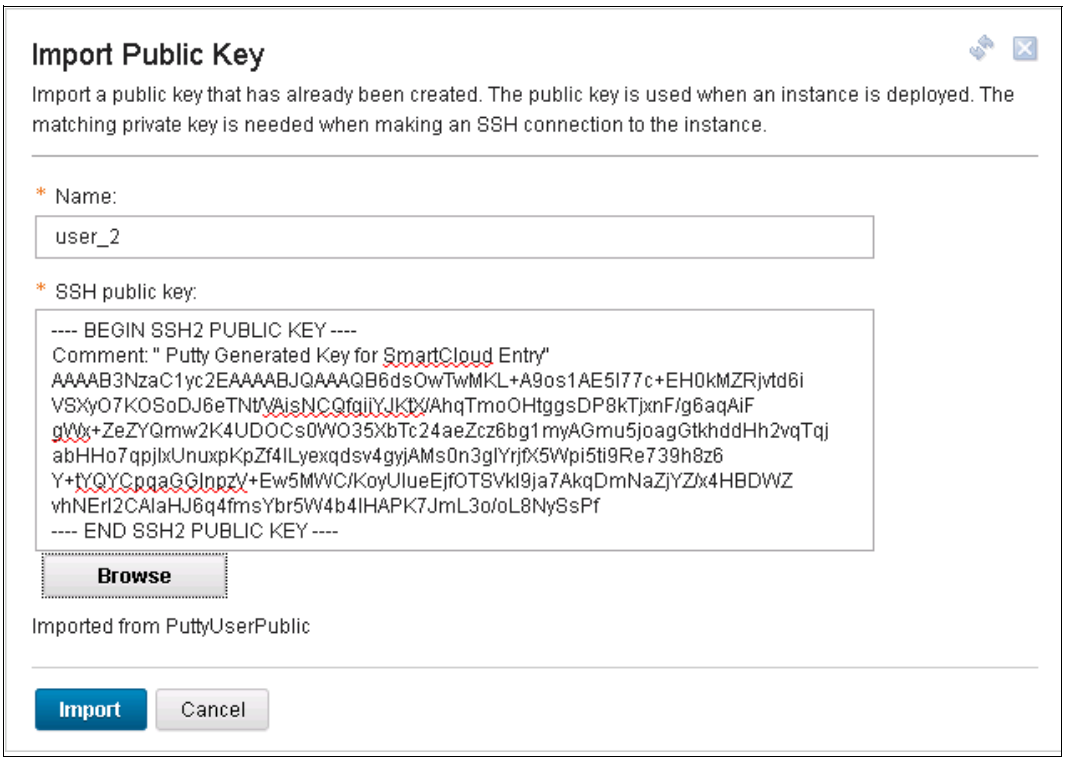

Figure 4-14 Import Public Key window

SSH key pairs must meet these requirements:

- IBM Common Cloud Stack stores only the public key in the SSH-authorized key file format.
- The SSH key pair is linked to the user who created it. No other users have access to the SSH key pair, including administrators.
- Standard SSH format is required and must be compatible with, for example, PuTTY and OpenSSH.
- The SSH key pair can be used outside of IBM Common Cloud Stack.

Figure 4-15 shows how the flow of a key pair connection might look. In this flow diagram, the instance file was created from an image file by using cloud-init activation software. The user deployed the image and chose a public key pair at deployment, so the image has the public key file stored on the configuration drive. Therefore, the public key file was copied to the instance by the cloud-init activation software.

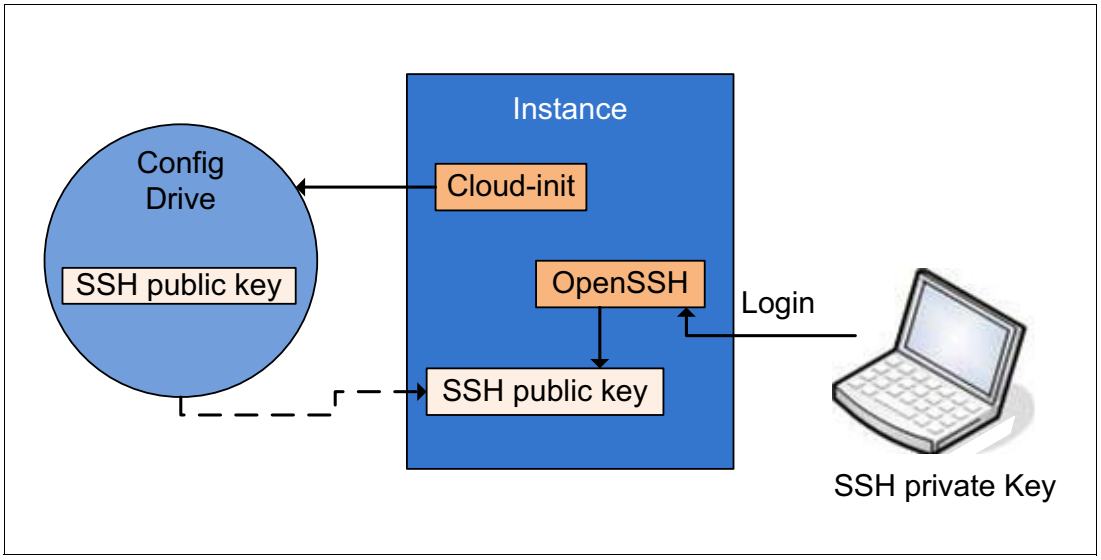

Figure 4-15 IBM Common Cloud Stack key pair management

After this process is carried out once, you can log in to the instance by using the private key. For more information about using the key pairs, see 4.3.3, “Use IBM SmartCloud Entry key pairs” on page 59.

4.3.3 Use IBM SmartCloud Entry key pairs

If you have a key pair that was generated by using IBM SmartCloud Entry and an instance that was deployed with the public key on it, you can access the deployed instance by using OpenSSH with the command that is shown in Example 4-4.

Example 4-4 SSH command to access an instance

```
ssh -i user1OpenStackKeyPair.pem user@192.168.0.10
```

You can also use PuTTY to access the instance. However, PuTTY allows only certain types of key pairs. Therefore, if you created a key pair with IBM SmartCloud Entry, you might need to convert this first, and then import the private key to PuTTY.

1. Complete these steps to open the PuTTY Key Generator:
   a. Click **Conversions** → **Import key**.
   b. Select the private key that you downloaded from IBM SmartCloud Entry.
   c. (Optional) Add a pass phrase for the private key.
   d. Click **Save private key**.
2. Complete these steps to open the PuTTY authentication agent:
   a. Click **Add Key**.
   b. Select the private key that was converted and saved.
   c. Open **PuTTY Configuration**, and log in to the system.

A

Cloud architecture standards

Standards that are related to cloud architecture are briefly explained in 1.1.4, “Standards for cloud architecture and implementation” on page 4. Further details are provided in this Appendix. Developers in your organization are most likely to be interested in this chapter.

IBM Cloud Computing Reference Architecture

The IBM Cloud Computing Reference Architecture (CCRA) is a guide for designing and implementing a cloud computing solution. The CCRA helps to determine the preferred implementation in line with the functional and non-functional requirements of a prospective implementation. It defines the basic building blocks, such as the architectural elements that compose the cloud solution and the relationships between these elements. The CCRA also defines basic principles of managing cloud services.

The IBM CCRA includes the following additional characteristics:

- Vendor and technology neutral: The IBM CCRA outlines the necessary functions for delivering cloud services, but it does not specify which tools to use to implement them. This increases flexibility for cloud service creators and providers. An example of this flexibility is that organizations can incorporate their existing infrastructure with solutions from multiple vendors in building and providing cloud services.
- Focused on open standards: The IBM CCRA promotes standards-based rather than proprietary tools, allowing organizations to select from a broad array of IBM and non-IBM hardware and software.
- Based on service-oriented architecture (SOA): SOA is defined by The Open Group (http://www3.opengroup.org/subjectareas/soa) to be “an architectural style that supports service orientation. (It is) a way of thinking in terms of services, service-based development, and the outcomes of services.”

In the IBM CCRA, cloud computing solutions as inherently SOA solutions, although not all SOA solutions are cloud computing solutions. Figure A-1 provides an overview of the IBM CCRA.

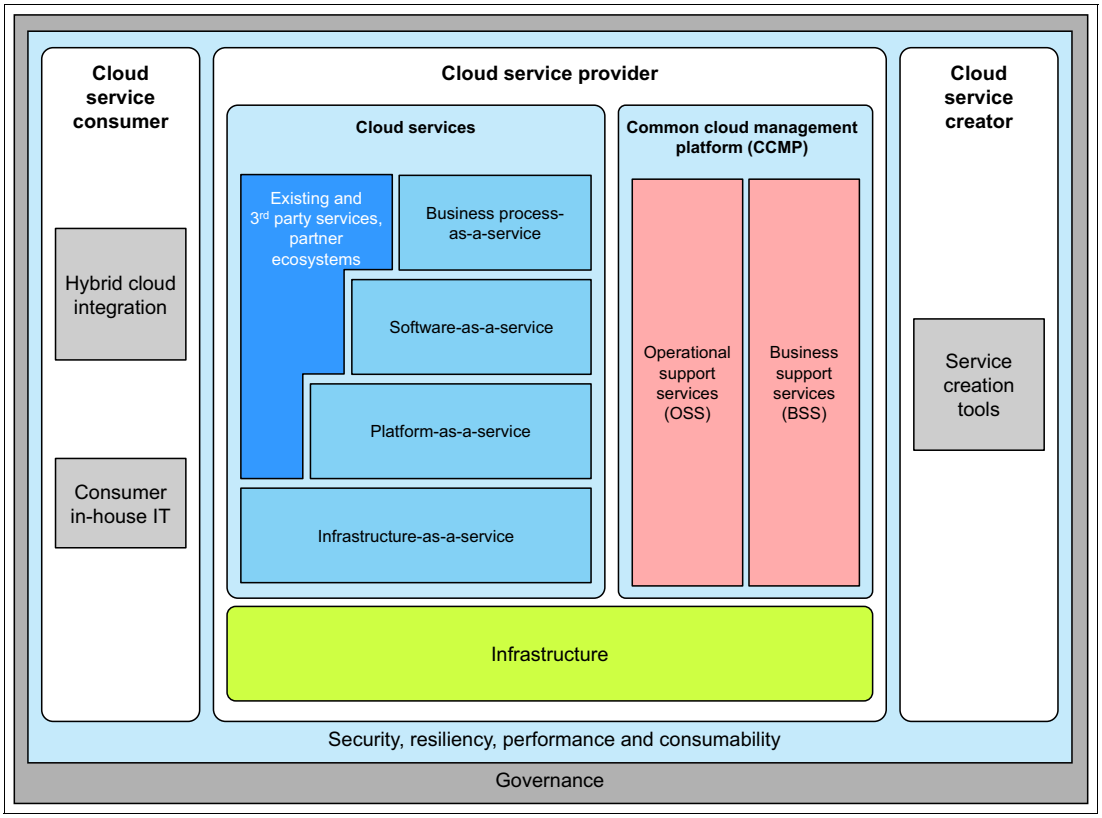

Figure A-1 The Cloud Computing Reference Architecture

In line with the CCRA are the roles and service models described in “The IBM Cloud Computing Reference Architecture” on page 4.

Common Cloud Management Platform

The CCMP is a set of functions that delivers and manages cloud services. A cloud services provider uses a CCMP to make cloud services available. The CCMP does not specify which software products or tools to use to meet these requirements. The selection of tools is made during the architecture process of a cloud service solution. Therefore, the CCMP functions are used to define the appropriate services models (IaaS, PaaS, SaaS, or BPaaS) of cloud services. Figure A-2 shows the CCMP platform.

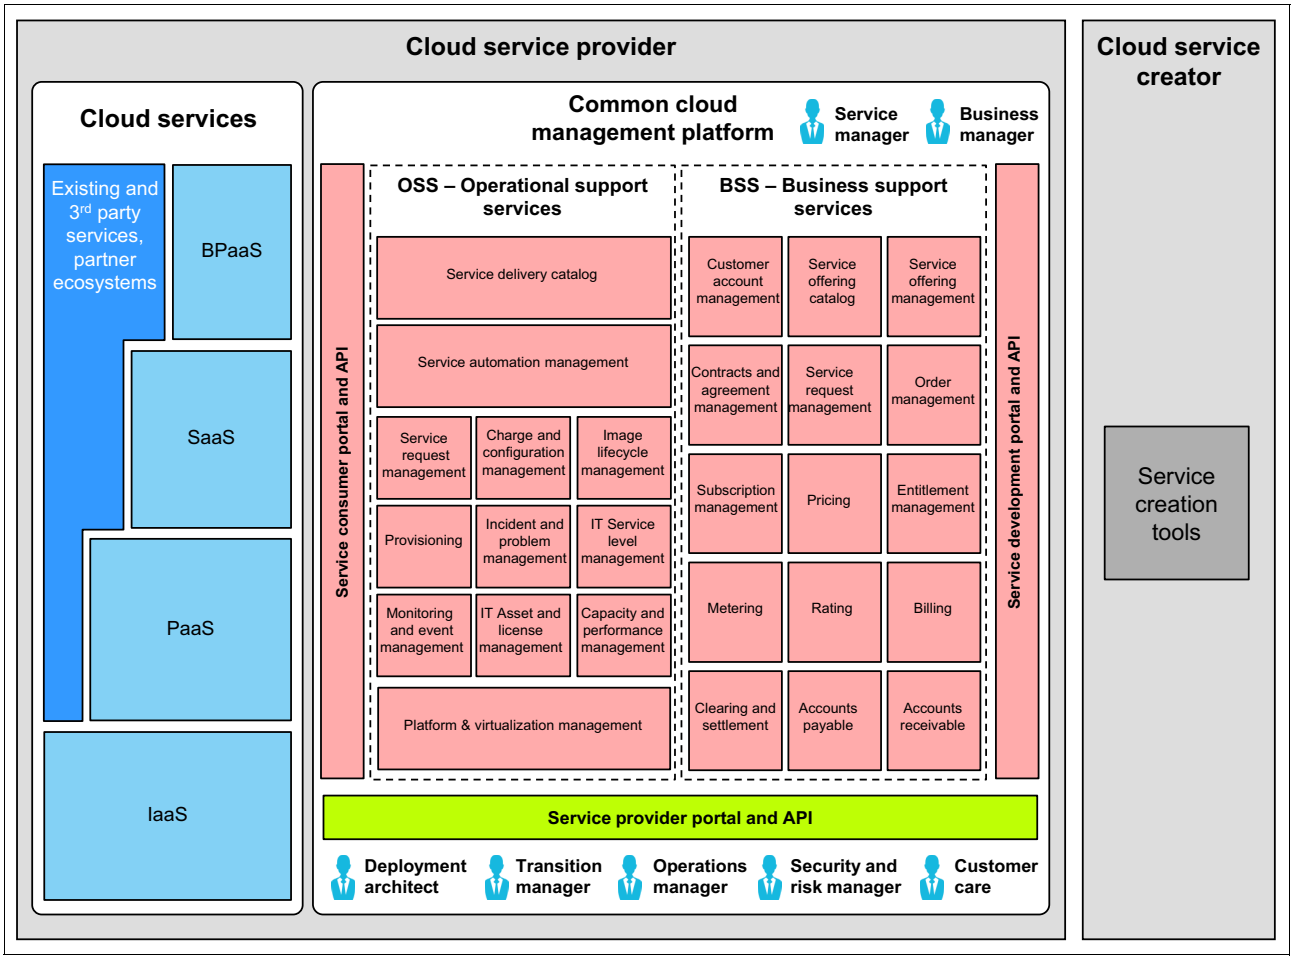

Figure A-2 The Cloud Computing Management Platform

The CCMP includes the following components:

- Business Support Services (BSS): BSS functions are for the business management of cloud services, including these services:
  - Determining pricing for cloud services
  - Billing for cloud services
  - Publishing cloud service offerings in a manner that people who use cloud service can easily access and request them
  - Determining whether cloud service requests are approved
  - Accounts payable and receivable

Not all business support services that are listed in the CCRA are part of every cloud services solution. As an example, in a private cloud, your company can choose not to process charge-backs to users, and so accounts payable and accounts receivable services are not necessary.

- Operational Support Services (OSS): OSS functions are for the operational management of cloud services, including the following services:
  - Provisioning assets to fulfill the cloud service request.
  - Managing assets and licenses that are associated with fulfilling a cloud services request.
  - Resolving incidents and fixing root causes of problems that interfere with delivery of cloud services.
  - Managing and virtualizing the underlying infrastructure.
  - Monitoring the cloud environment.
  - It is possible that some OSSs might not be a part of a particular cloud solution. For example, a simple private cloud might not have the capacity and performance management tools implemented.

For more information, see *Introduction and Architecture Overview, IBM Cloud Computing Reference Architecture 2.0* on The Open Group website:

https://www.opengroup.org/cloudcomputing/uploads/40/23840/CCRA.IBMSubmission.02282011.doc

Move to the Cloud Easily with IBM SmartCloud Entry V3.1 and IBM Common Cloud Stack

Redpaper™

Lower your total cost of computing with this simple-to-use product

Use the power and resources of the cloud with ease

Create and manage a Microsoft Hyper-V cloud

This IBM Redpaper publication describes our newest product for helping businesses enter the world of cloud technology: IBM SmartCloud Entry, Version 3.1. This software includes the IBM Common Cloud Stack, which is powered by open source OpenStack from the OpenStack Foundation. These products work together to help you move your business to the cloud. This paper explains the ways that you can take advantage of this extensive and powerful technology.

With the power of cloud computing, you can enhance your existing services, extend your market reach, and create new markets for your services. IBM SmartCloud Entry V3.1 is easy to use, even for first-time cloud service users. It can enable you to benefit from cloud technology in less time and with less effort than you might think.

This Redpaper features two scenarios that demonstrate the ease of carrying out processes with IBM SmartCloud Entry software. The information is directed to two primary audiences. Chapter 1 is directed mainly to decision-makers, such as CEOs, CIOs, and CFOs, who need to know about cloud technology and the power that it offers. The remainder of the paper is directed to IT professionals, such as information architects, business intelligence administrators, and database administrators, who need to know about the functions and capabilities of SmartCloud Entry and Common Cloud Stack.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

REDP-5030-00