

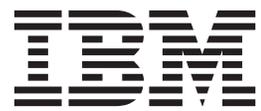
IBM Cloud Orchestrator
Content Pack for Infrastructure as a Service
Version 2.4

*Content Pack for Infrastructure as a
Service*



IBM Cloud Orchestrator
Content Pack for Infrastructure as a Service
Version 2.4

*Content Pack for Infrastructure as a
Service*



Note

Before using this information and the product it supports, read the information in "Notices" on page 19.

Contents

Preface	v
Audience	v

IBM Cloud Orchestrator Content Pack for Infrastructure as a Service 1

Installing and configuring	2
Upgrading	4
Customizing the installation packages of the management agents	4
Making the toolkit writable	6
Customizing the list of available flavors	6
Customizing the list of IBM Tivoli Storage Manager policies and backup frequencies.	7
Customizing the list of available IBM Endpoint Manager compliance groups	7
Customizing the YUM repository for the deployed systems	8
Toolkit scenarios	8
Creating an offering to deploy a cloud service	8
Decommissioning a cloud service	11
Deleting an offering to deploy a cloud service.	12
Deploying a cloud service	12
Registering a management server configuration	13
Unregistering a management server configuration	13
Modifying a management server configuration	14
Toolkit developer's reference.	14
Environmental variables	14
storehouse_iaas_templates	15
user_interaction_type	15

Business processes definitions	15
Create Service Deployment Template	15
Decommission a Cloud Service	15
Delete Service Deployment Template	15
Run Service Deployment Template	15
Register Management Server	15
Modify Management Server	15
Unregister Management Server	15
Human services	15
Create Service Template	15
Decommission a Cloud Service	15
Delete Service Deployment Template	16
Deploy Service Template	16
Register Management Server	16
Modify Management Server	16
Unregister Management Server	16
Coach views	16
Integration services.	16
General system services	16
Samples about how to use the available services and views	16
Troubleshooting	16
Configuring an SMTP server in IBM Cloud Orchestrator	17

Notices 19

Trademarks and Service Marks 21

Preface

This publication documents how to use the IBM Cloud Orchestrator Content Pack for Infrastructure as a Service.

Audience

This information is intended for content developers of IBM Cloud Orchestrator who implement or customize the content pack for Infrastructure as a Service.

IBM Cloud Orchestrator Content Pack for Infrastructure as a Service

The IBM Cloud Orchestrator Content Pack for Infrastructure as a Service provides a set of reusable artifacts in IBM Cloud Orchestrator to implement common customer scenarios in the context of cloud infrastructure lifecycle: deployment of virtual servers, enablement of specific services bound to the customer infrastructure, maintenance, and decommissioning of a service.

The main requirements addressed by this content pack are:

- Support the lifecycle of cloud services fully integrated in the customer infrastructure:
 - Setup or decommissioning of elements of the customer infrastructure, for example, agents, specific software, and so on.
 - Consume or provide data to or from external systems, for example, CMDB, LDAP, IPAM, and so on.
 - Maintenance of the cloud service enabled by means of base operations as stop or start of the service, for example, offline, and so on.
 - Expose services to external systems for enterprise level integration.
- Provide a library of reusable artifacts, building blocks, that can be used to assemble the cloud automation scenarios described in the previous point (BPM artifacts) in IBM Cloud Orchestrator.
- Provide a few IBM Cloud Orchestrator Business Process Manager working samples implementing the most common cloud automation scenarios and including a large set of connected building blocks. These samples are used to enable a development model based more on customization of existing artifacts than on the creation of new ones.
- Enable IT administrators to publish a cloud automation (service) controlling the configuration elements exposed to the end users by hiding, defaulting, and customizing the parameters of the cloud automation without needing the skills typical of service developers.

This content pack covers only a subset of supported lifecycle scenarios. It only addresses the deployment or configuration of a service, while all the interactions with external systems (for example, CMDB, LDAP, IPAM and son on) are not included and will be made available in a future refresh.

To understand the available scenarios, familiarize yourself with some of the concepts:

Cloud service

The cloud service is the final virtual system that must be deployed and that is configured according to the information contained in the service template. Using this content pack, you can create the following kinds of cloud services according to the base pattern used to configure the service:

- OpenStack Heat templates
- Classic Virtual System Patterns
- Advanced Virtual System Patterns

Service template

The service template is a document that describes how a given service

must be deployed and configured. It contains all the information needed to configure the virtual system in terms of default parameters, optional components, and additional management capabilities that must be configured on top of it. It also maintains information about which options, when a given service is instantiated, must be exposed to the final users and which ones must be hidden and filled with default values. Service templates are stored in the Storehouse library as JSON data structures. When a service template gets created, a companion Self-Service Offering is created as well. The link between a service template and the corresponding Offering is represented by the ID of the Offering that is stored as part of the template name.

Flavor

The flavor is a sort of logical template that defines the virtual hardware and environmental characteristics of the instance to be deployed. At a lower level, a flavor is an appropriate combination of an Environment Profile, a Cloud Group, an IP Group, and an OpenStack flavor that, once selected, are used to instantiate the virtual system. This element has been introduced to simplify the user interaction when creating the offering (and its associated template) and when triggering it.

After the installation of this content pack, a number of default flavors are available, represented by a basic combination of the first Environment Profile, Cloud Group, IP Group and the first three OpenStack flavors defined in the environment. Based on the organization of the environment where this content pack gets installed, such default list of flavors might need to be updated. Instructions about how to do this are provided later in this document.

Note: Flavors are available only when you create cloud services that leverage Virtual System Classic patterns.

Offering to deploy a cloud service

This is the Self-Service Offering created along with a service template and that, once triggered, drive the deployment of the service as described by the template. This Offering is always bound to the same pair of Human Service and Business Process Definitions. The Human Service is able to automatically select the template corresponding to the triggered Offering and to pass it to the service deployment and configuration logic.

Supported versions

- IBM Cloud Orchestrator version 2.4
- IBM Endpoint Manager versions 9.0 and 9.1
- IBM Tivoli Storage Manager version 6.2
- IBM Tivoli Monitoring version 6.3 Fix Pack 2

Note: Currently, only Linux x86 machines are supported, as nodes of the deployed systems.

Installing and configuring

A good knowledge of IBM Cloud Orchestrator and of the Business Process Manager programming model is required for using the toolkit as a software development kit (SDK) for building new content.

This content pack comes already installed in IBM Cloud Orchestrator V2.4. This means that the toolkit and the related offerings and actions are already available in the IBM Cloud Orchestrator environment. However, a set of configuration actions must be performed before you use the content pack.

From the Cloud marketplace, download the external version of this content pack because it contains a few additional artifacts required to fully configure it. You can download the content pack here: <http://www.ibm.com/software/ismlibrary?NavCode=1TW10SO27>.

After you download the compressed file from the marketplace, extract the following archive that contain prebuilt template packages to install and configure the management agents supported by this content pack:

- itm-client.zip
- iem-client.zip
- tsm-client.zip

You must customize these files by adding the necessary agent installation packages and then publishing to a reachable web server. See “Customizing the installation packages of the management agents” on page 4.

Before starting to use this content pack, you must:

- If it has never been done in your environment before, run the “Initialize the data Persistence Service” offering located in the “Infrastructure management” category. This is required to initialize the “IBM Cloud Orchestrator Content Pack for Persistency Services, which this toolkit depends on.
- Register at least one management server in the environment, from among the supported ones. Different offerings must be used for this registration activity depending on the kind of the management server configuration to create. In particular:

To register a management server configuration for IBM Tivoli Monitoring and IBM Tivoli Storage Manager:

Run the “Register Management Server” offering located in the “Infrastructure management” category.

To register a management server configuration for IBM Endpoint Manager:

Run the “Register IBM Endpoint Manager Server” offering located in the “Infrastructure management” category.

- Review and customize the list of available default IBM Endpoint Manager compliance groups, Tivoli Storage Manager policies, and backup frequencies.
- If you want to enable the “Backup & Restore” capability, ensure that a working YUM repository is correctly configured in the virtual images that you are going to deploy. It must reference the operating system package repository. This is required in case the deployed instances miss some of the operating system dependencies required for the installation of the management agents. Alternatively, you can specify an explicit YUM repository to be configured on the deployed systems instead of requiring an existing one. This requires some customization to be applied to the toolkit.
- Ensure that the *unzip* utility is installed in the virtual images that you are going to deploy. This is required for the installation of the management agents.
- Ensure that the virtual images included in the pattern being deployed exist in the OpenStack region where the pattern is deployed.

- If you are going to configure additional volumes in your cloud service, ensure that the OpenStack Cinder component is correctly configured and that at least one volume type is available.

Specifically for cloud services that leverage Virtual System Classic patterns:

- You can customize the user interaction type behavior that can be modified through the `user_interaction_type` toolkit variable. See `user_interaction_type` in “Environmental variables” on page 14.
- If the user interaction type has been set to “simplified,” the definition of the Flavors must be customized to fit the resources available in the environment . See “Customizing the list of available flavors” on page 6.
- If the user interaction type has been set to “simplified,” ensure that all the patterns that you are going to deploy through this solution have the parameters of each part be assigned a default value and are locked. Both at offering creation time and at offering deploy time it is not possible to customize them and if a required parameter does not have a value the deploy fails. Be aware that if the pattern contains parts with multiple NICs, all of them are assigned to the same IP group.

Upgrading

If you already have the toolkit installed and are upgrading from SmartCloud Orchestrator V2.3 to IBM Cloud Orchestrator V2.4, the toolkit is automatically upgraded. However, because of changes in the structure of the self-service catalog and new version 2.4 architecture, you must make some adjustments manually to clean up the environment so that the toolkits can run.

Make the following adjustments:

- Remove all the offerings under the "Manage Cloud Services" category. See **IBM Cloud Orchestrator > Managing orchestration workflows > Orchestration workflows > Self-service offerings** in the IBM Knowledge Center.
- In moving from SmartCloud Orchestrator V2.3 to IBM Cloud Orchestrator V2.4, several changes have been implemented in the internal architecture of the product; the main one being the removal of the dependency against the Chef infrastructure for installing software modules and its replacement with the IBM Endpoint Manager one. As a consequence, cloud services definitions (for example, offerings) created with the version 2.3 toolkit, do not work in IBM Cloud Orchestrator V2.4. You must manually remove the old definitions (offerings) and recreate new ones through the 2.4 version of the toolkit.

Customizing the installation packages of the management agents

Three archives are provided to handle the installation and configuration of the management agents on the deployed systems.

Before using these archives, you must customize them to fit specific deployment needs:

iem-client.zip

The IBM Endpoint Manager Client installable file, for example, `BESAgent-XXXX.rpm`, must be added to the package. You can find the list of IBM Endpoint Manager packages at <http://support.bigfix.com/bes/install/downloadbes.html>.

Here is the typical content of the `iem-client.zip` file after it has been fully customized:

```
Archive: iem-agent.zip
  BESAgent-9.0.649.0-rhe5.x86_64.rpm
  install-iem-client.sh
  readme.txt
```

itm-client.zip

Export the IBM Tivoli Monitoring agent bundle for your platform into a temporary directory and add its content to a compressed file. This compressed file must also contain the parent directory where the bundle files are located. Then, embed it in the `itm-client.zip` file. Finally, modify the reference to the name of the created compressed file, that is specified by the `lnx_itmAgentBundle` shell variable, inside the `install-itm-client.sh` script.

The IBM Tivoli Monitoring agent bundles can be extracted, from the from the IBM Tivoli Monitoring server machine, using the `tacmd exportBundles` command. For documentation about the command, see http://publib.boulder.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3/cmdref/exportbundles.htm.

For example, to extract the agent bundle for the Linux x86 platform, you can use the following CLI:

```
tacmd exportBundles -e c:\export -o LOCAL -t lz -p lx8266
```

If you need to handle the decommissioning flow of the IBM Tivoli Monitoring agent, then the IBM Tivoli Monitoring Tivoli Enterprise Services User Interface Extensions (namely, the `tacmd` CLI) component must be installed with the IBM Tivoli Monitoring agent. To do this, extract the IBM Tivoli Monitoring CLI component for the platform for which you need it into a temporary directory and add its content into a compressed file. This compressed file must also contain the parent directory where the bundle files are located. Then, embed it in the `itm-client.zip` file. Finally, modify the reference to the name of the created compressed file, that is specified by the `lnx_itmCliBundle` shell variable, inside the `install-itm-client.sh` script. For example, to extract the CLI bundle for the Linux x86 platform, use the following command:

```
tacmd exportBundles -e c:\export -o LOCAL -t ue -p lx8266
```

The KSH shell component must be available on the nodes where the IBM Tivoli Monitoring agent must be installed. The agent installation script tries to find the KSH installable RPM package among the files contained in the `itm-client.zip` file. If that component is not available by default on the target systems, the corresponding RPM file can be retrieved from the operating system package repository and added to the compressed file. Once there, it is automatically installed before the IBM Tivoli Monitoring agent.

Here is the typical content of the `itm-client.zip` file after it has been fully customized:

```
Archive: itm-agent.zip
  install-itm-client.sh
  ksh-20100621-19.e16.x86_64.rpm
  lz_063002000_lx8266.zip
  readme.txt
  silent_config.txt
  ue_063002000_lx8266.zip
```

tsm-client.zip

Download the IBM Tivoli Storage Manager agent bundles for your platform from the IBM site and add them into this archive. For a list of the IBM Tivoli Storage Manager agent packages, see <http://www-01.ibm.com/support/docview.wss?&uid=swg24034630>.

Select the package for your platform and accept the terms and conditions. Download the .tar file containing the agent bundle. For example, for the Linux x86 platform, download the file: 6.2.5.0-TIV-TSMBAC-LinuxX86.tar. The actual name may change according to the version available at the moment you download it.

After the package has been locally downloaded, it must be extracted into a temporary directory and its content added to a compressed file. Then, this compressed file must be added to the tsm-client.zip file. Finally, the reference to the created compressed file, that is assigned to the TSM_PACKAGE shell variable, inside the install-tsm-client.sh script must be updated to match the name of the created compressed file.

Here is the typical content of the tsm-client.zip file after it has been fully customized:

```
Archive: tsm-agent.zip
  dsm.opt.template
  dsm.sys.template
  install-tsm-client.sh
  readme.txt
  TSM-6.2.zip
```

where TSM-6.2.zip contains all the files that have been extracted from the tar file containing the agent bundle. In this example, 6.2.5.0-TIV-TSMBAC-LinuxX86.tar.

Afterwards, you must upload these archives to a web server where they can be downloaded from the deployed systems. Network connectivity between the deployed systems and the web server must be ensured to allow a successful installation and configuration of the agents.

Making the toolkit writable

By default, after it is installed, the SCOrchestrator_IAAS_Assets toolkit can be accessed in read-only mode through the Business Process Designer. To apply the customization described in the following sections, you must make the toolkit writable.

To make the toolkit writable:

1. Open the Business Process Designer tool.
2. In the **Toolkits** tab, select the SCOrchestrator_IAAS_Assets toolkit and click **Manage**.
3. Select the **Allow users to update the toolkit** check box.

Customizing the list of available flavors

To modify the list of available flavors, use the Business Process Manager Process Designer Tool.

Open the SCOrchestrator_IAAS_Assets toolkit and select the **BuildServiceFlavors** General System Service.

This service generates the definitions of the available flavors. To add a new flavor, you must add a new block calling the BuildServiceFlavor Integration Service by cloning an existing one and chaining it to the already existing ones. Then, in the **Data Mapping > Input Mapping** section, you can customize the definition of the new flavor by specifying the name of the Environment Profile, Cloud Group, IP Group, and OpenStack flavor that are associated with it.

You can also remove or modify an already existing flavor similarly.

When adding a new flavor, be sure to copy and paste an already-defined flavor instead of adding a new block from scratch so that it inherits all the required configurations.

Afterwards, save the changes to make them available in IBM Cloud Orchestrator.

Customizing the list of IBM Tivoli Storage Manager policies and backup frequencies

To modify the list of backup policies and frequencies, use the Business Process Manager Process Designer Tool.

The lists, which you can see in the UI when you configure the **Backup and Restore** capability, are defined in the **Extract Params** block of the **Config Backup and Restore** Human Service. Each name of a displayed backup policy is mapped into an IBM Tivoli Storage Manager Policy domain. You must update this list to reflect the values of the backup policy domains defined in the IBM Tivoli Storage Manager server available in the environment.

The same applies for the backup frequencies. Each name of a displayed backup frequency is mapped into a Tivoli Storage Manager backup frequency. This list must be updated to reflect the values of the backup frequencies defined in the Tivoli Storage Manager server available in the environment.

If policies and backup frequencies are not correctly defined in the Tivoli Storage Manager server, the installation and registration of the Tivoli Storage Manager agent fail.

Customizing the list of available IBM Endpoint Manager compliance groups

To modify the list of compliance groups, use the Business Process Manager Process Designer Tool.

This list, which you can see in the UI when you configure the **Security and Compliance** capability, is defined in the **Config Security & Compliance** Human Service. To access and modify the definition of the list, open the “Variables” section of the Human Service and select the “availableGroups” private variable.

Each name of a displayed compliance group is mapped into an IBM Endpoint Manager client setting that is set on the agents once installed.

You must update this list to reflect the values of the properties defined in the IBM Endpoint Manager server available in the environment to define the related automatic groups.

More specifically, when a compliance groups is selected, a client setting called `_BESCLIENT_GROUP_NAME` is created on each agent with a value equal to the one of the selected compliance groups. To have these agents automatically subscribed to an IBM Endpoint Manager automatic group when creating that group, set the property **Relevance Expression** to **is true** and specify the following expression:

```
exists setting "_BESClient_GROUP_NAME" of client AND value of setting  
    "_BESClient_GROUP_NAME" of client = "MY_VALUE"
```

This way, the group includes any computer that has the `_BESCLIENT_GROUP_NAME` setting set to `MY_VALUE` where `MY_VALUE` is the value associated to one of the defined compliance groups.

Customizing the YUM repository for the deployed systems

By default, when the IBM Tivoli Storage Manager agent must be installed on the deployed systems, a YUM repository referencing the operating system install packages is required to install any missing dependency.

Alternatively, an explicit YUM repository can be configured at runtime to the deployed system.

To enable this capability, you must use the Business Process Manager Process Designer Tool to customize the **Config Backup and Restore** Human Service.

The **Extract Params** block must be edited by adding a Javascript snippet similar to:

```
if(tw.local.yumRepoUrl == null) {  
    tw.local.yumRepoUrl = "http://162.168.2.3/repo/rhel";  
}
```

where the value assigned to the `tw.local.yumRepoUrl` variable, must refer to the YUM repository of the operating system packages available in the environment.

Afterwards, the installation and configuration of the IBM Tivoli Storage Manager agent leverages that repository to install any missing dependency.

Toolkit scenarios

There are a number of scenarios that are immediately available from the toolkit.

The following items are the main building blocks of the toolkit:

- “Creating an offering to deploy a cloud service”
- “Decommissioning a cloud service” on page 11
- “Deleting an offering to deploy a cloud service” on page 12
- “Deploying a cloud service” on page 12
- “Registering a management server configuration” on page 13
- “Unregistering a management server configuration” on page 13
- “Modifying a management server configuration” on page 14

Creating an offering to deploy a cloud service

You can configure and create a new entry in the self-service catalog that can later be used to trigger the deployment of a new service in the cloud.

Through this offering, the Catalog Editor can compose a new service by customizing a set of basic deployment options and specifying the additional management capabilities that must be enabled in the deployed instance. At the end of this process, a new self-service offering is created in the catalog that can be used to immediately trigger the deployment of the configured system, by exposing a UI with a limited set of customization options.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the **Design and create cloud service offerings** category.
2. Start the **Create an Offering to Deploy a Cloud Service** offering.
3. After the offering starts:
 - a. If at least one of the supported management servers (IBM Tivoli Monitoring, IBM Endpoint Manager, IBM Tivoli Storage Manager, or Opscode Chef) is not registered in the environment, you receive a message specifying that it is not possible to configure the corresponding management capability. Additionally, you cannot enable the corresponding capability for the cloud service being configured in the UI.
 - b. Specify the general options related to the offering that is created:
 - Provide a name and optional description.
 - Select an icon to assign to the offering.
 - Specify the category under which the offering is created.
 - Specify whether an approval is requested for deploying the offering.
 - c. If the approval has been enabled, select the role of the users that receive the approval request when the offering is triggered. Only the users having that role in the same project of the logged-on user triggering the offering receive the request and have the possibility of approving it.
 - d. Provide your own OpenStack Heat template or select the pattern to deploy from the available patterns.
 - e. If you provide the Heat template:
 - 1) Select the Region.
 - 2) Enter the Heat template.
 - 3) Configure the Heat template.
 - f. If you select a Classic Virtual System Pattern and the active user interaction type has been set to “simplified”:
 - 1) Select a Flavor to size the instance to deploy. The list of available Flavors depend on the customization done in the environment where this content pack is installed.
 - 2) Optionally, specify whether the selected Flavor can be shown to the final users and modified by them by checking **Hide this configuration option from users**.
 - g. If you select a Classic Virtual System Pattern and the active user interaction type has been set to “full”:
 - 1) Configure the general parameters of the selected pattern, such as, the Environment Profile.
 - 2) For each part of the selected pattern:
 - a) Specify the Cloud Group and the Flavor.
 - b) For each virtual NIC, specify the IP Group and, if required, the IP address and host name.
 - c) Specify any custom parameter associated to the part.

- d) Specify any custom parameter for the script packages associated to the part.

Note: For all the parameters, you can specify whether they can be shown to the final users and modified by them by checking **Hide this configuration option from users**.

- h. If you select an Advanced Virtual System Pattern:
- 1) Select the Deployment Configuration.
 - 2) Configure each node.
- i. After you set the options related to the chosen pattern, specify the credentials to be used to remotely accessing the virtual servers being deployed:
- If you select **Connect through SSH keys**:
 - The administrative user
 - An OpenStack key pair name or your own private key

Note: To create suitable key pairs to allow access to virtual servers you can use the offering “Register a key to enable access to virtual servers” available in the “Infrastructure Management” category.

- If you select **Connect through Credentials**:
 - The administrative user name
 - The password of the administrative user
- j. Select the management and additional capabilities to be configured in the deployed system:
- **Security and Compliance:** Enabled only if at least one IBM Endpoint Manager server has been registered
 - **Backup and Restore:** Enabled only if at least one IBM Tivoli Storage Manager server has been registered
 - **Monitoring:** Enabled only if at least IBM Tivoli Monitoring server has been registered
 - **Storage Volumes:** Through the OpenStack Cinder service.
 - **Additional Software Installation:** Enabled only if at least one Opscode Chef Server has been registered
- k. For each of the management services enabled, a page is displayed where you can configure them:

Security and Compliance

Specify the location of the installation package of the agent, the customized `iem-agent.zip` file. Select the IP address or host name of the IBM Endpoint Manager server. Additionally, you can select an optional compliance group, representing the name of an IBM Endpoint Manager automatic group where the installed agents are subscribed.

Backup and Restore

Specify the location of the installation package of the agent, the customized `tsm-agent.zip` file. Select the IP address or host name of the IBM Tivoli Storage Manager server. Optionally, you can associate the node to a backup policy and define the backup frequency.

Monitoring

Specify the location of the installation package of the agent, the customized itm-agent.zip file. Select the IP address or host name of the IBM Tivoli Monitoring server.

Storage Volumes

The size of the volume to be created and attached to the system must be provided along with the volume type, the file system type, and the mount point. You can specify if they must be exposed to users after the created offering is triggered.

Note: You can only create and attach volumes to the deployed service when you select a flavor that refers to KVM cloud groups.

Additional Software Installation

Select the Chef server. This step is optional and is required only if more than one Chef servers are registered in the environment. Then specify the software modules to be installed on specific parts of the selected pattern. For each entry, configure the corresponding parameters for the specified software.

4. Review the summary information and submit the operation. Click **Create** to start the creation of the offering. At the end of the operation, a new self-service offering is available under the category specified.

Table 1. Typical configuration of the offering

Name	Create an Offering to Deploy a Cloud Service
Category	Design and create cloud service offerings
Process	Create Service Deployment Template (SCOrchestrator_IAAS_Assets)
User interface	Create Service Template (SCOrchestrator_IAAS_Assets)

Decommissioning a cloud service

You can decommission a cloud service from the system.

In the context of this operation, any agent installed on top of the virtual machines of the service are stopped and unregistered from the corresponding management server. Each attached storage volume is detached and deleted from the system. Then, the virtual system instance of the service gets removed.

Note: Because of a defect in the IBM Endpoint Manager V9.0 application, the decommissioning of the IBM Endpoint manager agents does not remove the definition of the computer from the IBM Endpoint Manager server. A fix for this issue will be made available shortly in IBM Endpoint Manager.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the **Deploy cloud services** category.
2. Start the **Decommission a Cloud Service** offering.
3. After the offering starts:
 - a. Select the service instance to decommission.
 - b. Review the summary information and submit the operation. When you click **Delete**, the operation starts and its outcome is available in **My Requests**.

At the end of the operation, a notification report with the detailed outcome of the operation is delivered to the user that requested the operation, provided that a

valid email address and that a working Simple Mail Transfer Protocol (SMTP) server have been configured with IBM Cloud Orchestrator.

Table 2. Typical configuration of the offering

Name	Decommission a Cloud Service
Category	Deploy cloud services
Process	Decommission Cloud Service (SCOrchestrator_IAAS_Assets)
User interface	Decommission Cloud Service (SCOrchestrator_IAAS_Assets)

Deleting an offering to deploy a cloud service

You can remove an entry from the self-service catalog created through the **Create an Offering to Deploy a Cloud Service** offering along with its template stored in the Storehouse.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the **Design and create cloud service offerings** category.
2. Start the **Delete an Offering to Deploy a Cloud Service** offering.
3. After the offering starts:
 - a. Select the offering to be removed.
 - b. Submit the operation.

At the end of the operation, the offering and its associated template is removed from the system.

Table 3. Typical configuration of the offering

Name	Delete an Offering to Deploy a Cloud Service
Category	Design and create cloud service offerings
Process	Delete Service Deployment Template (SCOrchestrator_IAAS_Assets)
User interface	Delete Service Deployment Template (SCOrchestrator_IAAS_Assets)

Deploying a cloud service

You can launch a self-service catalog created through the **Create an Offering to Deploy a Cloud Service** offering.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the category where the offering is located.
2. Start the offering.
3. After the offering starts:
 - a. Specify a name for the instance that is being created.
 - b. Based on the configuration of the Offering, the user is asked to review and provide a value for all those parameters that have been marked to be shown to the final users. Review and if necessary, assign a different value to those parameters.
4. Review the summary info and submit the operation. Click **Deploy** to start the offering.

At the end of the operation, a new virtual system instance is available and configured according to the options you specified.

If an approval has been enabled for the selected offering, the deployment and configuration of the service occur only if the request gets approved. More specifically, all the users matching the approver selection criteria receive a dedicated request in their inbox that they can approve or reject. Based on this decision, the actual deployment of the service occurs or is rejected by the system. The deployment flow is suspended until one of those users takes ownership of the request and accepts or rejects it.

At the end of the operation, a notification report with the detailed outcome of the operation is delivered to the user that requested the operation, provided that a valid email address and that a working Simple Mail Transfer Protocol (SMTP) server have been configured with IBM Cloud Orchestrator.

Registering a management server configuration

You can register the basic configuration parameters in IBM Cloud Orchestrator to access an IBM Tivoli Monitoring or IBM Tivoli Storage Manager server.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the **Infrastructure management** category.
2. Start the **Register Management Server** offering.
3. After the offering starts:
 - a. Specify the type of server configuration to be created: **IBM Tivoli Monitoring** or **IBM Tivoli Storage Manager**.
 - b. Specify the server host name or IP address.
 - c. Specify the user name of an administrative user to connect to the server.
 - d. Specify the password for the user specified.
 - e. Submit the offering.

Table 4. Typical configuration of the offering

Name	Register Management Server
Category	Infrastructure management
Process	Register Management Server (SCOrchestrator_IAAS_Assets)
User interface	Register Management Server (SCOrchestrator_IAAS_Assets)

Unregistering a management server configuration

You can unregister the basic configuration parameters from IBM Cloud Orchestrator to access a configured IBM Tivoli Monitoring or IBM Tivoli Storage Manager server.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the **Infrastructure management** category.
2. Start the **Unregister Management Server** offering.
3. After the offering starts:
 - a. Select the server configuration to be removed.
 - b. Submit the offering.

Table 5. Typical configuration of the offering

Name	Unregister Management Server
-------------	------------------------------

Table 5. Typical configuration of the offering (continued)

Category	Infrastructure management
Process	Unregister Management Server (SCOrchestrator_IAAS_Assets)
User interface	Unregister Management Server (SCOrchestrator_IAAS_Assets)

Modifying a management server configuration

You can update the basic configuration parameters in IBM Cloud Orchestrator to access a configured IBM Tivoli Monitoring or IBM Tivoli Storage Manager server.

To run this scenario, perform the following steps:

1. Go to the **Self-Service Catalog** and open the **Infrastructure management** category.
2. Start the **ModifyManagement Server** offering.
3. After the offering starts:
 - a. Select the server configuration to modify.
 - b. Review and modify the server connection port.
 - c. Review and modify the username to connect to the server.
 - d. Review and modify the password for the user specified.
 - e. Submit the offering.

Table 6. Typical configuration of the offering

Name	Modify Management Server
Category	Infrastructure management
Process	Modify Management Server (SCOrchestrator_IAAS_Assets)
User interface	Modify Management Server (SCOrchestrator_IAAS_Assets)

Toolkit developer's reference

This section contains reference information about the Business Process Manager artifacts exposed in the toolkit contained in the content pack. It is intended to be used by IBM Cloud Orchestrator content developers to extend the already available scenarios or to write new scenarios leveraging the building blocks available from the toolkit.

The following items are the main building blocks of the toolkit:

- “Environmental variables”
- “Business processes definitions” on page 15
- “Human services” on page 15
- “Coach views” on page 16
- “Integration services” on page 16
- “General system services” on page 16

Environmental variables

Some environment variables are available so you can customize some behaviors when you run the use cases.

Environmental variables are defined in the toolkit in **Toolkit Settings > Environment**.

storehouse_iaas_templates

This variable contains the default path in the Storehouse library where all the configuration templates are stored. Never modify the value of this variable. Its default value is set to: /storehouse/sco_iaas_templates/.

user_interaction_type

This variable allows to switch the UI to configure an offering and to deploy it from the simplified form into the full one. Allowed values are: "simplified", "full". The default value is: "simplified". The value of this variable has effect only when you select Virtual System Classic patterns.

Business processes definitions

There are a number of public Business Process Manager business objects defined in the toolkit.

Create Service Deployment Template

This is the business process that handles the creation of a deployment template and the creation of a new self-service offering associated with the template.

Decommission a Cloud Service

This is the business process that handles the decommission of a service.

Delete Service Deployment Template

This is the business process that handles the deletion of a service deployment template and its associated self-service offering, if any.

Run Service Deployment Template

This is the business process that handles the entire end-to-end deployment flow according to the configuration specified in the template.

Register Management Server

This is the business process that handles the registration of the management server configuration of the supported type: IBM Tivoli Monitoring, IBM Tivoli Storage Manager.

Modify Management Server

This is the business process that handles the modification of the registered management server configuration.

Unregister Management Server

This is the business process that handles the unregistration of the registered management server configuration.

Human services

There are a number of Human Services artifacts available in the toolkit.

Create Service Template

This is the Human Service that is used to collect all the parameters required to create a new service deployment template and its associated self-service offering.

Decommission a Cloud Service

This is the Human Service used to collect the information for decommissioning a cloud service.

Delete Service Deployment Template

This is the Human Service used to collect the information for deleting an already existing deployment template and its associated self-service offering.

Deploy Service Template

This is the Human Service used to collect all the parameters needed to run the deployment of a service template. This service is able to automatically select the deployment template for which the offering associated to it has been created and to make a request to the user for only the subset of parameters that have not been marked as *hidden* by the operator that created it.

Register Management Server

This is the Human Service used to collect the information of the management server configuration: IBM Tivoli Monitoring, IBM Tivoli Storage Manager.

Modify Management Server

This is the Human Service used to modify the registered management server configurations.

Unregister Management Server

This is the Human Service used to unregister the management server configurations.

Coach views

No public coach views are currently available.

Integration services

No public integration services are currently available.

General system services

No public general system services are currently available.

Samples about how to use the available services and views

The services that are delivered as part of this content pack can be used as starting points and as samples for developing new content.

Using samples can give you an idea about how to invoke and combine the Integration Services delivered in the content pack. You can clone and adapt them to better fit your needs. After you are familiar with them, you can start creating your new offerings and actions.

For information about how to develop new content, see the IBM Cloud Orchestrator Content Development Guide available in the IBM Knowledge Center:

- IBM Cloud Orchestrator > Managing orchestration workflows > IBM Cloud Orchestrator content development

Troubleshooting

There might be some problems or limitations.

If there is a failure running a scenario, the information available on the UI or through the mail report might not be enough for you to completely understand the root cause of the failure. If this is the case, review the Business Process Manager

and IBM Cloud Orchestrator traces to see more specific error messages. For more information, see the following topic in the IBM Cloud Orchestrator documentation:

IBM Cloud Orchestrator > Troubleshooting.

Configuring an SMTP server in IBM Cloud Orchestrator

To use the email notification feature, you must configure Business Process Manager to use an appropriate SMTP relay server inside your organization. Contact your mail server administration to obtain the host name of the SMTP server.

Business Process Manager does not support any type of SMTP authentication or non-standard SMTP ports.

Create a file called `100Custom.xml` with the following content:

```
properties>
  <server merge="mergeChildren">
    <email merge="mergeChildren">
      <smtp-server merge="replace">${SMTP_HOST_NAME}</smtp-server>
    </email>
  </server>
</properties>
```

Replace `SMTP_HOST_NAME` with your SMTP server host name.

Put the `100Custom.xml` file in the directory: `/opt/ibm/BPM/v8.5/profiles/DmgrProfile/config/cells/PCCell11/nodes/Node1/servers/SingleClusterMember1/process-center/config`

If the file is already there, edit it adding or updating the “server/email” section with the provided information. Afterwards, restart the Business Process Manager server using the command: `service bpm restart`.

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.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks and Service Marks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. 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>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are trademarks or registered trademarks of Adobe Systems, Incorporated, in the United States and/or other countries.

Intel, the Intel logo, Intel Inside, the Intel Inside logo, Intel Centrino, the Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Chef is a registered trademark of Opscode in the United States and other countries.



Printed in USA