Hewlett Packard Enterprise

Technical white paper

Check if the document is available in the language of your choice.

HPE GREENLAKE FOR PRIVATE CLOUD

Instances and blueprints feature overview and best practices

CONTENTS

HPE GREENLAKE FOR PRIVATE CLOUD SERVICE OVERVIEW

HPE GreenLake for private cloud service is one of the services enabled by HPE GreenLake cloud services platform. HPE GreenLake for private cloud service provides the agility and velocity of the public cloud in an on-premises, pay-per-use solution that's ideal for your cloud-native applications and traditional applications. It is designed for enterprise customers who need to host applications on-premises or co-located data center for various reasons: compliance, governance, security, data gravity, latency, application dependency, cost, and so on. The service is accessed from HPE GreenLake Central, an intuitive, self-service, as-a-service portal and operations console that allows enterprises to easily access, trial, and consume HPE GreenLake services and hybrid cloud environments. HPE GreenLake Central also offers consumption analytics service for all HPE GreenLake cloud services. Consumption analytics meters usage of a customer's HPE GreenLake infrastructure and tracks available capacity, providing visibility and forecasting service for IT admins and line of business users. With HPE GreenLake Central, you can achieve a unified cloud experience across your hybrid estate through a single portal.

HPE GreenLake for private cloud service includes hardware, software, and services, fully managed by Hewlett Packard Enterprise. Benefits of the solution are:

- Hardware and software are preintegrated at HPE factory and delivered to your colocation facility or data center in as little as 14 days. HPE installs and configures the solution.
- The pricing model is consumption-based pay-per-use model. It offers full transparency of usage and costs and there is no big up-front investment. It's scaled for your current business demands, but there's built-in buffer capacity to address any unexpected spikes in resource demand.
- The solution is fully managed by HPE. HPE monitors and manages usage and performance in real time. You don't need to worry about infrastructure maintenance, upgrade, and so on. And since your usage is metered and managed by HPE, HPE helps you scale up your capacity ahead of business growth.
- It provides cloud experience for your on-premises data center. Role-based access and policy features allow company IT to flexibly manage access of the resources for different functional teams. The intuitive self-service portal and extensive APIs allow users to quickly spin up VMs. It also provides a rich set of capabilities and automation features to orchestrate and manage application lifecycle. This reduces dependency on company IT organization and improves speed and agility to better respond to business needs.

With the rich set of capabilities such as instance type, blueprints, integration with automation scripts, and workflows, the HPE GreenLake for private cloud service not only provides VM infrastructure as a service, it also enables you to easily, flexibly, and reliably create and deploy applications on top of the VM infrastructure. It is a powerful cloud application management platform. In this white paper, we describe these capabilities and best practices on how to utilize these features to realize the full potential of the application management with HPE GreenLake for private cloud service.

INSTANCES

One of the capabilities for cloud application management is the **Instances** feature. In many cloud management platforms, for example, Amazon Web Services (AWS), an Instance is a representative of a singular object such as a **virtual machine (VM)**. In HPE GreenLake for private cloud, an instance can be a single VM running base OS or can be an application service running in a single VM or multiple VMs. For example, we can have an instance running CentOS 7.9.2009 OS, or we can have a MariaDB database instance with one VM running the database service or a MariaDB database instance with three VMs running the database service.

Instances are provisioned from instance types. HPE GreenLake for private cloud provides you flexible ways and many options to build your own instance type, then make it available in the instance types catalog, for example a specific version of CentOS, Ubuntu instance type, or an application instance type, for example, MariaDB, Jenkins, and so on. Before illustrating when and how to create a custom instance type, let's understand the design of instance types first.

Instance types

Figure 1 shows the hierarchical design of **instance types**.

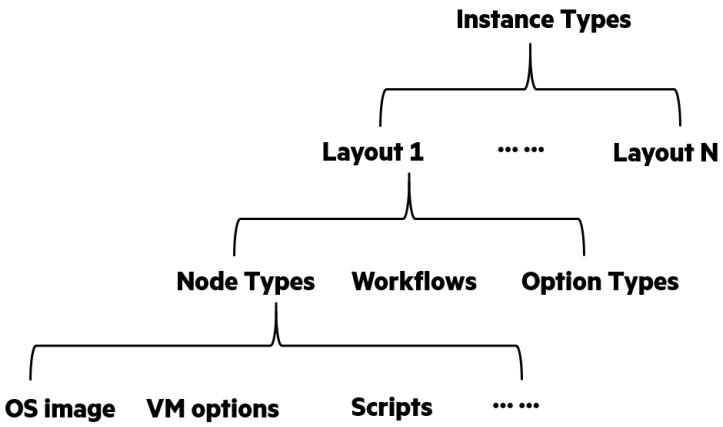

FIGURE 1. Instance types hierarchical design

- An instance type can have one or many layouts attached to it; a layout can only be attached to a single instance type.
- In each layout, you define layout name, description, version, and so on. There are some other key configuration options in layout:
  - Node Types: OS image, VM options, and automation scripts to run along with OS image installation, and such
  - Workflows: List of scripts or tasks grouped together, can be selected to run at provision time or executed on existing instances
  - Option Types: Custom input fields that can be added to instance type and Layout, for example, username, password that are specific to an instance, and so on.

As you can see, an instance type contains a base OS image, and a lot of other configuration options such as scripts, workflows, options types, which can be customized to build an application-specific new instance type.

Once you build your own instance type and make it available in the instance types catalog, you can create your instance using the create instance wizard.

From HPE GreenLake for private cloud dashboard, click the **Provisioning** tab from the main menu, select **Instances**, and click the green **+ADD** icon to add a new instance. In the **CREATE INSTANCE** wizard, the **TYPE** tab shows available instance types from instance types catalog. Figure 2 shows a screenshot example.

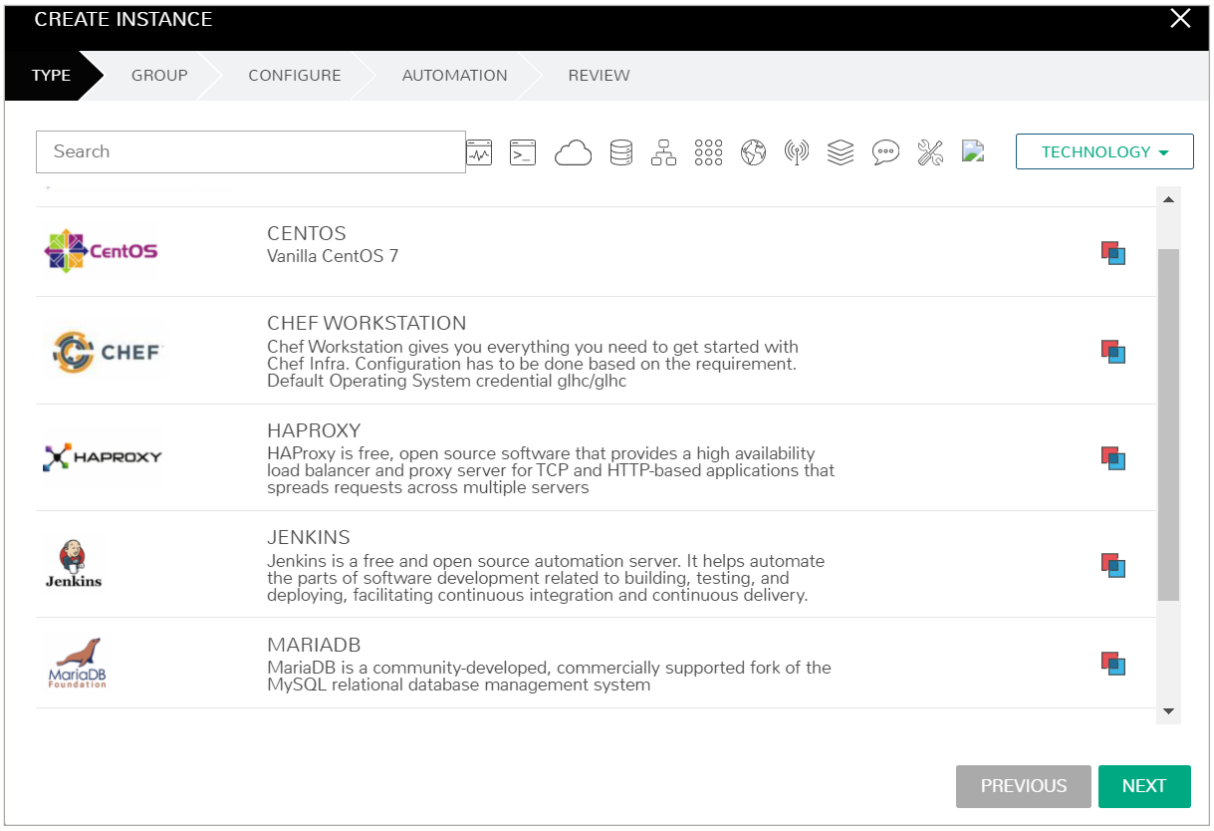

FIGURE 2. CREATE INSTANCE wizard

By default, HPE GreenLake for private cloud provides some prebuilt instance types, such as CentOS, Apache. Figure 2 shows some custom-built application instance types for illustration purpose, for example, Chef, HAProxy, Jenkins, MariaDB. You need to build these types of instance types based on your application requirement: base OS version, application software packages versions, and such. We recommend you follow the best practices on creating application instance types and instances.

Best practices of creating application instances

As we mentioned earlier, one of the benefits of HPE GreenLake for private cloud is the intuitive self-service portal for end users. To serve the HPE GreenLake for private cloud users more efficiently, your company IT can build a set of golden images based on company standard and business requirement, for example, a set of latest major OS version of the guest OS (Windows, CentOS, Ubuntu, SLES, Red Hat® Enterprise Linux® [RHEL], and such); with specific requirement of network, security, NIC, and so on; or some base application servers, for example, Apache Web Server. Then, IT uploads these golden images to HPE GreenLake for private cloud virtual images store. Using HPE GreenLake for private cloud self-service portal, the application developers can use the same golden OS images to create their own application instance types, for example, database developers create MariaDB instance type, DevOps developers create Chef Workstation instance type, or HAProxy load balancer and proxy server instance type, or Jenkins instance type, and so on. Using this best practice, the company IT team only needs to maintain the golden images; the application developers have great flexibility to build their own application instances with speed and agility, thereby reducing dependency on company IT.

Next, let's look at how to use the different configuration options to build an application instance type on top of the golden OS images.

Node Types with OS image and scripts example

In this section, we have an example of MariaDB instance type with Node Type screen capture and Layout screen capture. For complete step-by-step instructions on how to build a MariaDB instance type using HPE GreenLake for private cloud service, see "HPE GreenLake for private cloud deploy a 2-Tiered WordPress application using scripts and blueprints" white paper.[1]

[1] HPE GreenLake for private cloud: Deploy a Two-Tiered WordPress Application using scripts and blueprints h20195.www2.hpe.com/v2/Getdocument.aspx?docname=a50003251enw

In Figure 3, the **VM IMAGE** field is set with **CentOS_Training**. The **VM IMAGE** drop-down menu lists all the uploaded or synced golden images from the HPE GreenLake for private cloud virtual image store. In this example, the **CentOS_Training** image is a CentOS 7 64-bit image that is uploaded in the HPE GreenLake for private cloud virtual image store by IT team. You can check the virtual image store catalog under **Provisioning > Virtual Images**.

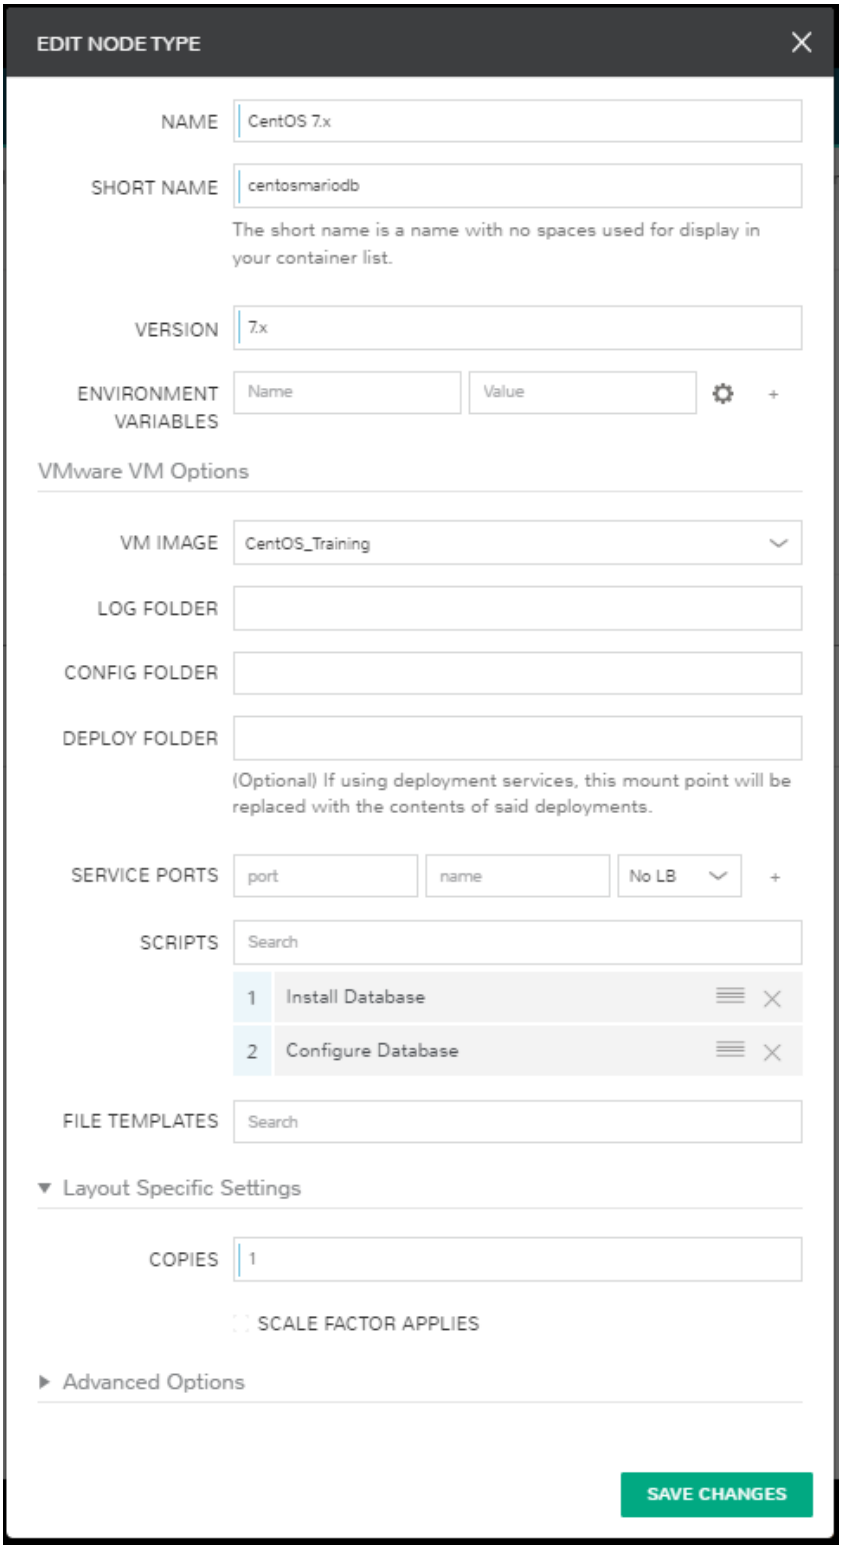

FIGURE 3. MariaDB node type example

To build a MariaDB node type, beside the base golden OS image **CentOS_Training**, it has **Install Database** and **Configure Database** scripts added in the **SCRIPTS** option. These two scripts are Bash scripts to install and configure the MariaDB database on top of the CentOS 7 64-bit OS image. Note that scripts are created under **Provisioning > Library > Scripts** from the HPE GreenLake for private cloud console. While creating scripts, the **PHASE** drop-down menu gives you a lot of options on when the scripts are executed: Pre Provision, Provision, or Post Provision, and so on, shown in Figure 4.

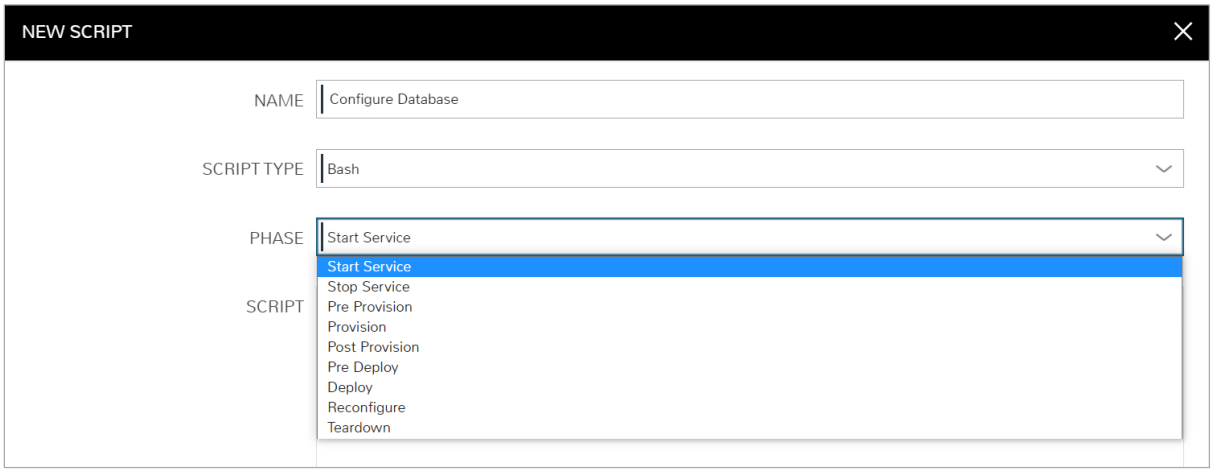

FIGURE 4. Configuration options for creating new script

Figure 5 shows the Bash script example of **Configure Database**. In this example, the **PHASE** is set as **Provision**, which means the script is going to be executed during instance provisioning phase.

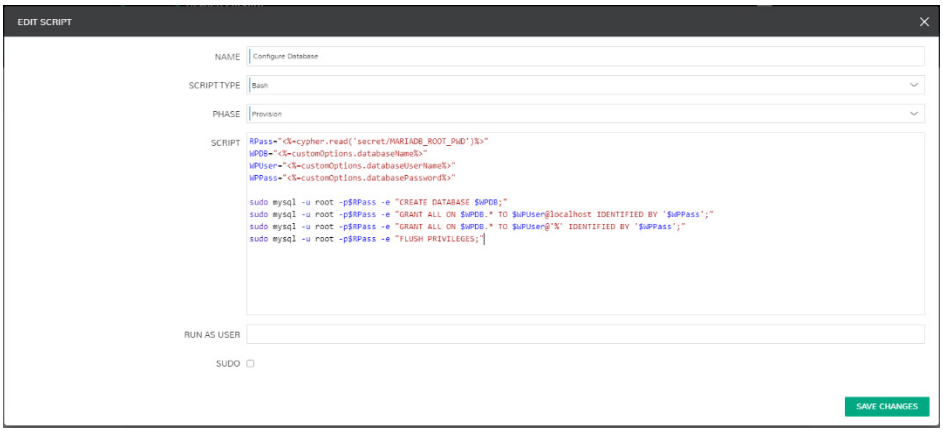

FIGURE 5. Configure database bash script example

Layout example

Figure 6 shows a Mario_DB_CentOS layout example. As you can see, this layout includes node types, option types, and workflow configurations, which are highly customizable. In this example, the option types are the Mario database custom input parameters namely username, password, and such.

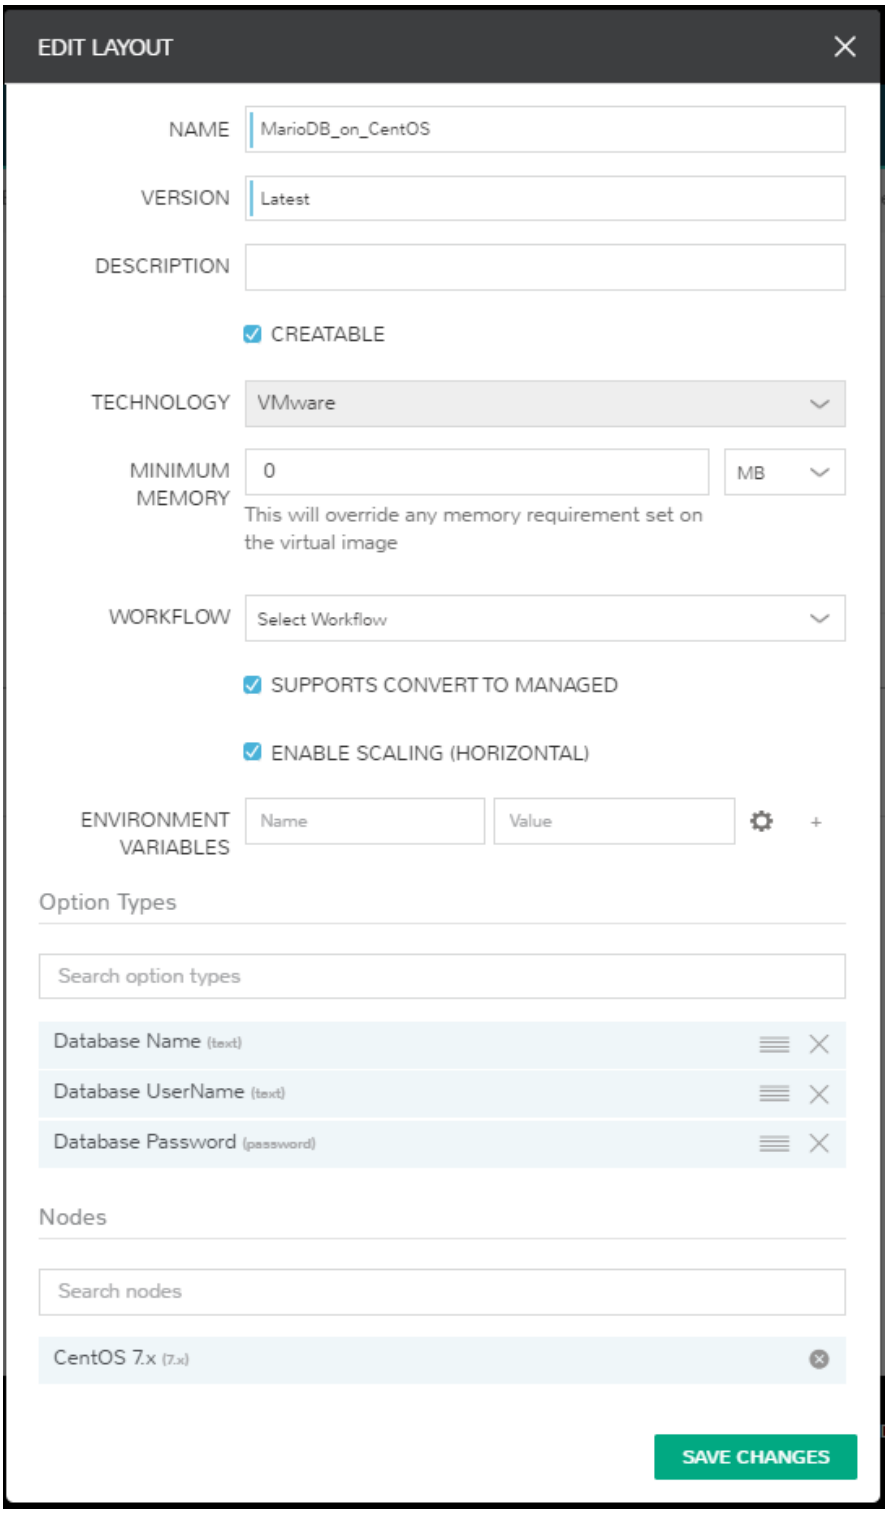

FIGURE 6. Mario_DB_CentOS layout example

Instance with auto-scaling example

There is another feature called auto-scaling for Instances. You can define the auto-scaling policy threshold under **Provisioning > Automation > Scale** Thresholds. Figure 7 shows a CPU threshold policy example. In this example, when instance VM CPU utilization is above 75%, a new VM will be automatically provisioned. Note that you need to update your load balancer to direct traffic to this new VM.

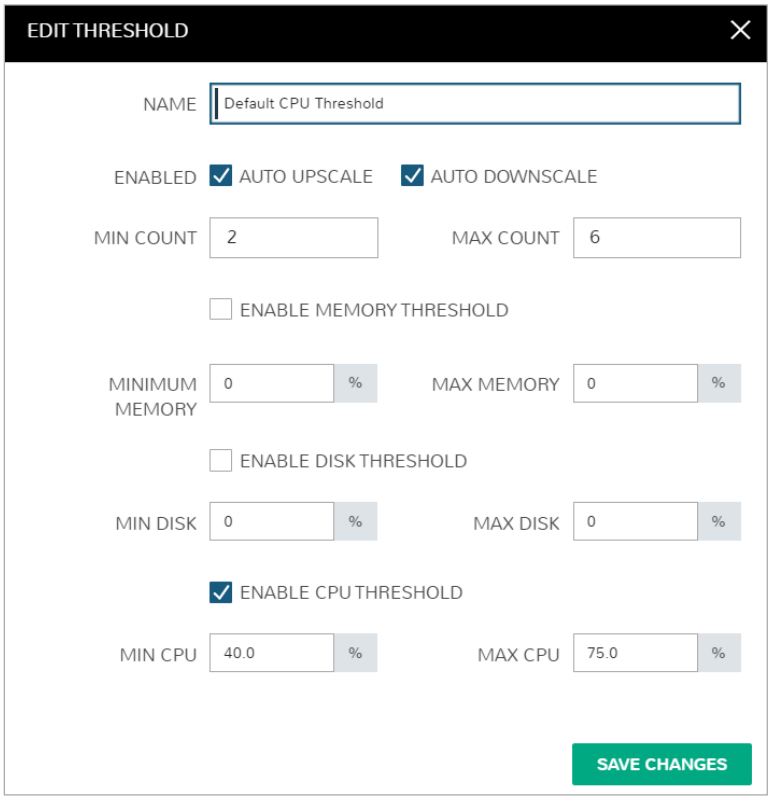

FIGURE 7. Define auto-scaling policy

When provisioning a new instance, we select the scale type and threshold policy to enable auto-scaling for this new instance, shown in Figure 8.

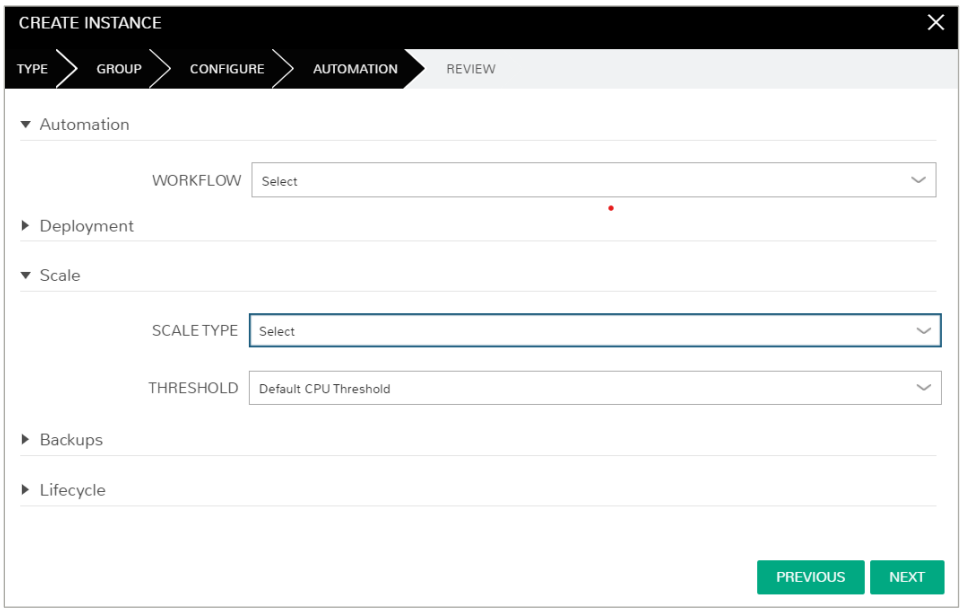

FIGURE 8. Enable auto-scaling when creating instance

Instance with multinodes example

In the "Instances" section, we mentioned that in HPE GreenLake for private cloud, an instance can be a single VM running base OS or can be an application service running in a single VM or multiple VMs. In this section, we show you an Apache Web Server instance example with three VMs running the web service. To create a multinode Apache instance, first, go to **Provisioning > Instances** to add a new Apache Web Server instance using the **CREATE INSTANCES** wizard, select the Apache instance type, and follow the wizard to fill in required parameters in each step. Wait until the Apache instance comes up. To add more VMs in the Apache instance, click the green **ACTIONS** button, and then select **Add Node** as shown in Figure 9.

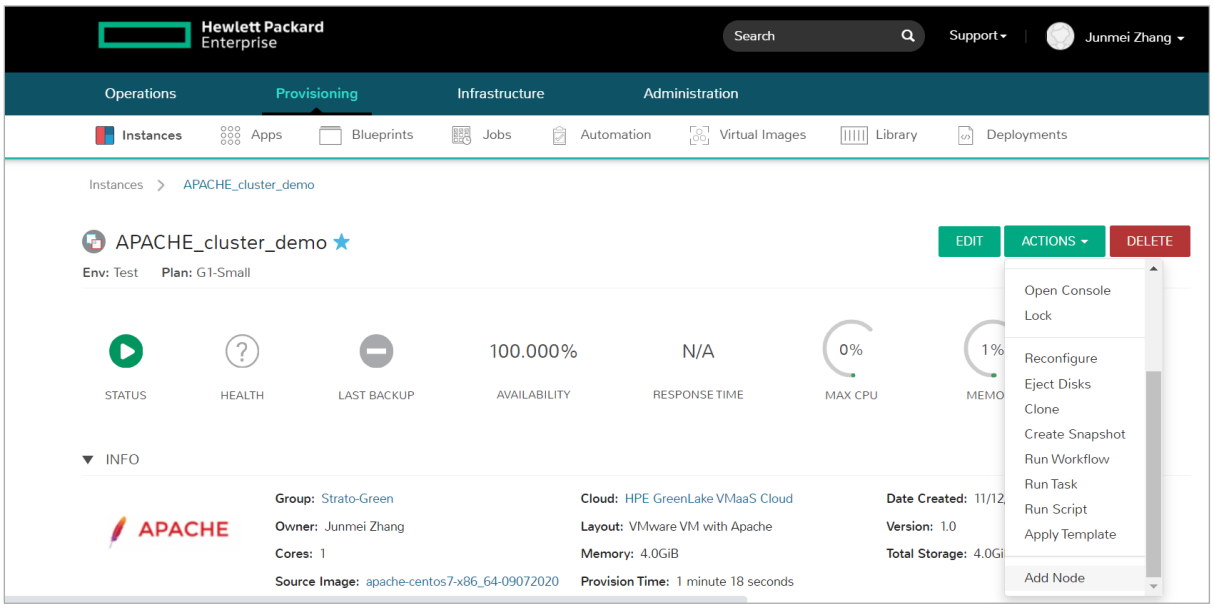

FIGURE 9. Add more nodes to Apache single VM instance

EXECUTE INSTANCE ACTION? window pops up as shown in Figure 10.

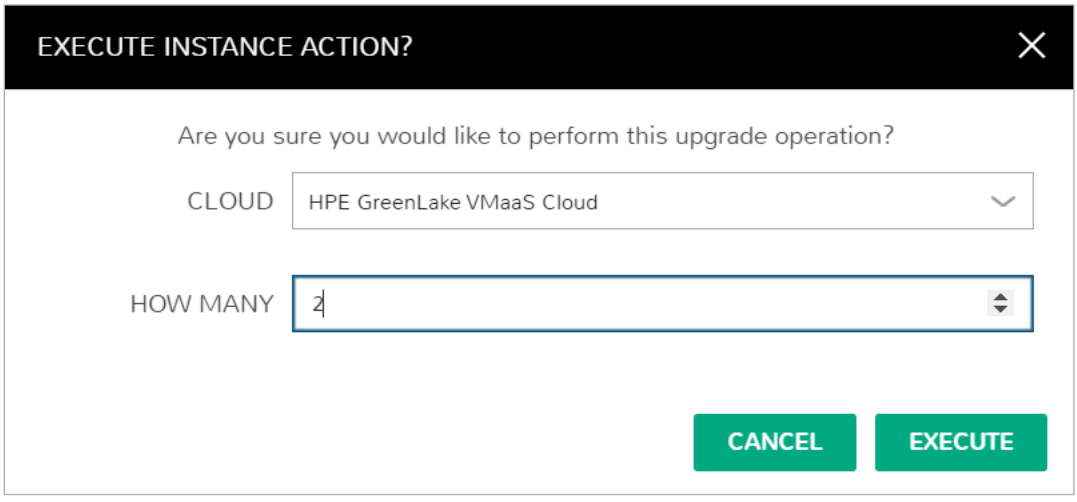

FIGURE 10. EXECUTE INSTANCE ACTION

Specify the number of nodes to be added to this instance and click EXECUTE to finish. Figure 11 shows the Apache Web Server instance with three VMs running.

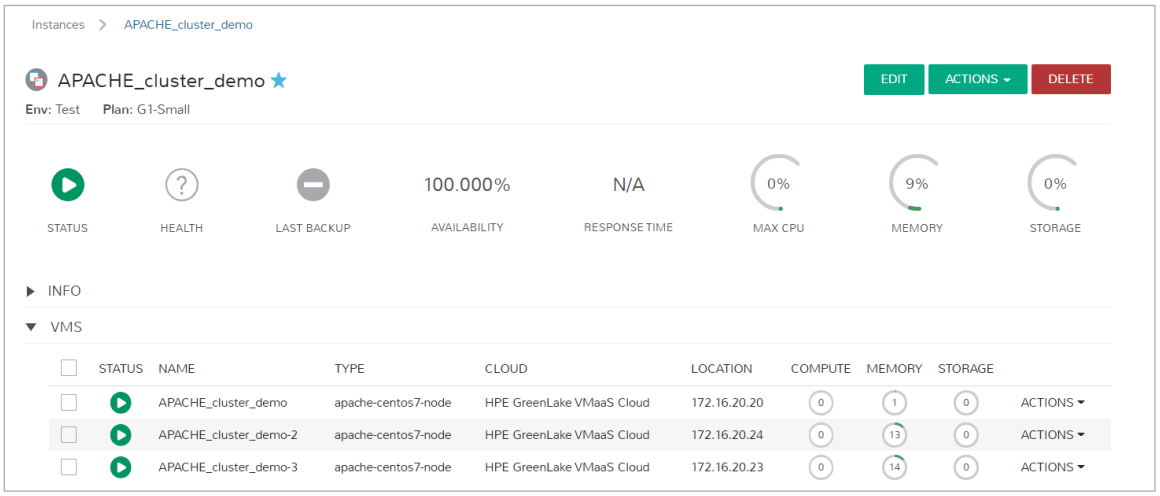

FIGURE 11. Apache Web Server instance with three VMs

Automation tasks, workflows, and Ansible playbook examples

HPE GreenLake for private cloud also provides a rich set of automation features, for example, tasks, workflows, integration with industry's various automation tools namely Puppet, Chef, Ansible, Ansible Tower, and such. Tasks are individual automation elements (for example, an individual playbook or a script). Workflows consist of one or more tasks. Tasks and workflows can be integrated with instances, blueprints at different provisioning phases, for example, run tasks or workflows at Pre Provision phase, Provision phase, or Post Provision phase. For detailed examples of how Ansible playbooks, tasks and workflows are integrated while provisioning an application, see "HPE GreenLake for private cloud: Utilizing application blueprints to simplify deployment of a multinode application using GreenLake for private cloud" white paper.[2]

BLUEPRINTS AND APPLICATION

Another powerful feature that enables cloud application management for HPE GreenLake for private cloud is the blueprints feature. You can instantiate applications quickly and reliably using blueprint. Figure 12 shows the blueprint hierarchical design.

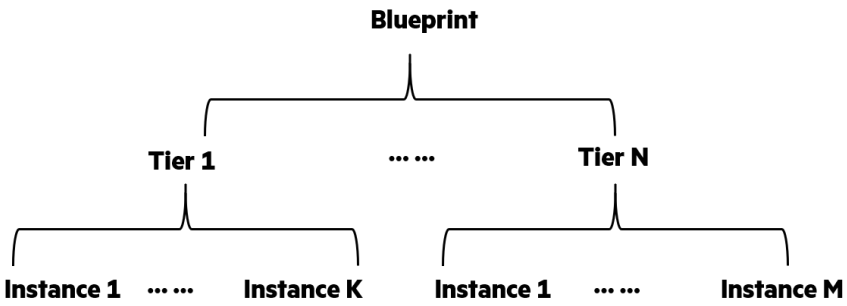

FIGURE 12. Blueprint hierarchical design

In simple term, blueprint is an application template; it defines the structure of an application. An application can be a single tier application or multitier application, for example, web tier, app tier, and database tier linked together to provide a service. Blueprint is structured with tiers and instance types. A blueprint can have many tiers, for example, web, app, database, cache, messaging, and such. A tier may have a single instance type or a group of instance types. The tiers can be marked as connected such that network communication rules can appropriately be defined. The tiers can also set booting sequence based on application requirement. An application is an instantiation of a blueprint. Figure 13 shows a 2-tier WordPress application blueprint example, using WordPress instance type at app tier and MariaDB instance type at database tier.

[2] HPE GreenLake for private cloud: Utilizing application blueprints to simplify deployment of a multinode application using GreenLake for private cloud developer.hpe.com/uploads/media/2020/6/gl4pc_eshop_bp_v1_35-1593186155592.pdf

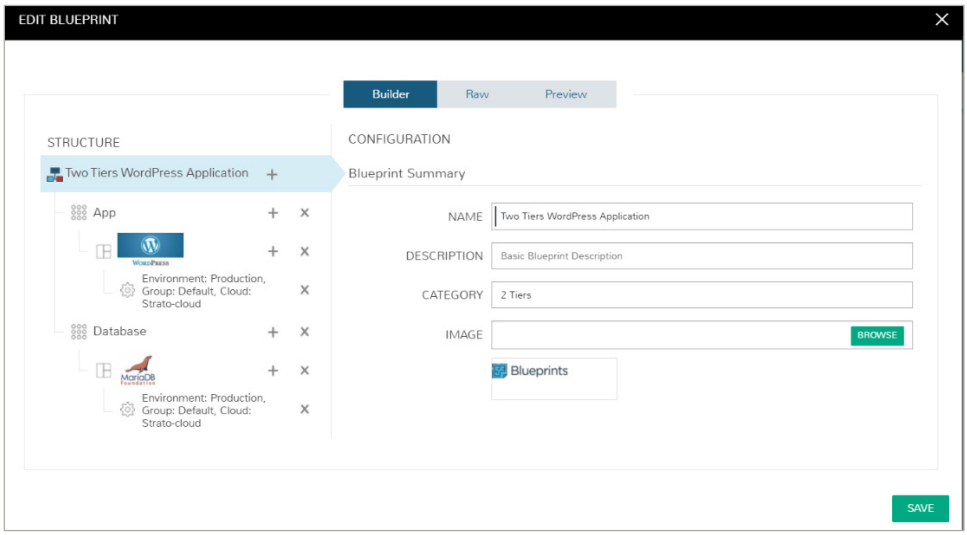

FIGURE 13. 2-tier WordPress application blueprint example

Blueprints can also be exported as YAML or JSON code and created with the API and CLI. Once you have a blueprint built for HPE GreenLake for private cloud, it is easy to export the blueprint and import to another HPE GreenLake for private cloud environment or public cloud.

Blueprint best practices

Now, we understand what blueprint is. Let's look at the best practices on how we utilize blueprint feature to quickly create and deploy applications. As we described in the "Instances" section, application developers have the flexible and rich configuration options to build instance types based on specific application requirements using golden base OS image. With blueprint feature, application developers can easily build a multitier blueprint using the different instance types. Once you build a blueprint for the application, you can save it to the blueprint catalog. Then, you can navigate to **Provisioning > Apps**, click the green **+ADD** button, and use the **NEW APP** wizard to quickly provision the application by selecting corresponding blueprint from the catalog. Figure 14 shows an example of the **NEW APP** wizard with the blueprints catalog.

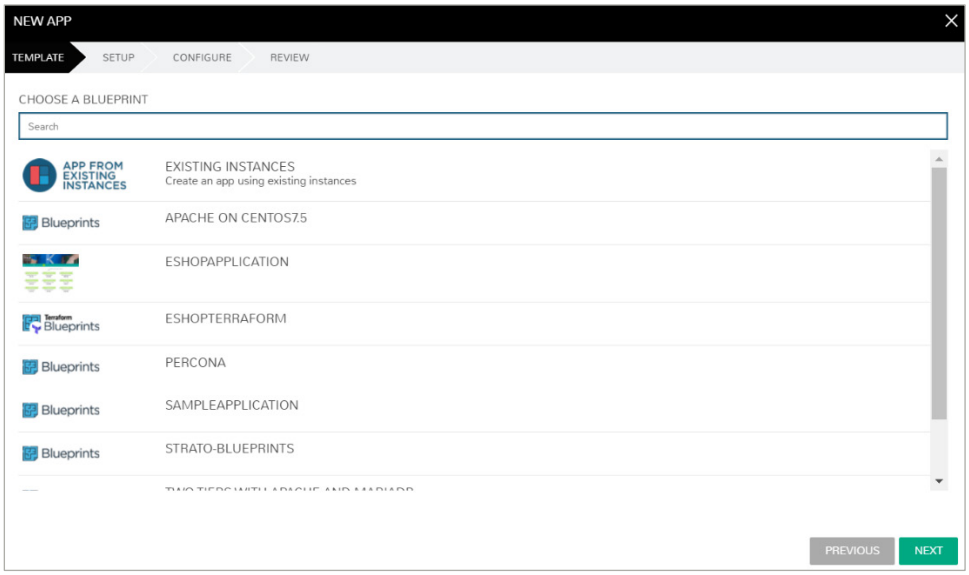

FIGURE 14. Create a new application wizard

With the blueprint feature, it is easy to create and deploy the application with reproducibility. White paper "HPE GreenLake for private cloud: Utilizing application blueprints to simplify deployment of a multinode application using GreenLake for private cloud" shows a blueprint use case example with step-by-step instruction.

SUMMARY

With the rich set of capabilities such as instance type, blueprint, integration with automation scripts, Ansible playbooks, tasks, and workflows, HPE GreenLake for private cloud solution not only provides VM infrastructure as a service, it also enables you to easily, flexibly, and reliably create and deploy applications on top of the VM infrastructure. It is a powerful cloud application management platform.

Learn more by exploring HPE GreenLake cloud services page or signing up for a free trial at HPE GreenLake for private cloud.

REFERENCES

1. **HPE GreenLake for private cloud:** Deploy a Two-Tiered WordPress Application using scripts and blueprints h20195.www2.hpe.com/v2/Getdocument.aspx?docname=a50003251enw
2. **HPE GreenLake for private cloud:** Utilizing application blueprints to simplify deployment of a multinode application using GreenLake for private cloud developer.hpe.com/uploads/media/2020/6/gl4pc_eshop_bp_v1_35-1593186155592.pdf

LEARN MORE AT

hpe.com/us/en/greenlake/private-cloud-vms.html

a50003040ENW, January 2021

Hewlett Packard Enterprise