


IBM PureApplication System

Cloud configuration: IP groups, cloud groups and storage volumes, environment profile, and deployment target



© 2012 IBM Corporation

This presentation covers the details of how to configure different clouds within IBM PureApplication™ Systems. Part of the cloud configuration involves understanding and setting IP groups, cloud groups, and storage volumes. Once, these resources are created, the deployment can then be targeted to the cloud. In PureApplication System, deployments of the workload are on Environment profiles. This presentation includes details on the Environment profiles and how to deploy the workloads.

Table of contents

- Overview of cloud configuration
- IP groups
 - Overview
 - Creating new IP groups
 - Add IP range, modify attributes
- Cloud groups
 - Overview
 - Creating new cloud groups
 - Modify cloud group attributes
 - Best practice
- Storage volumes
- Environment profile: Virtual system and virtual application pattern deployment target
- Summary

The agenda covers an overview of cloud configuration and then dives into IP groups, cloud groups, storage volumes, and then how to deploy the virtual application and virtual system patterns to the configured clouds through the environment profiles, discussed in the presentation.

Overview – cloud configuration

To deploy virtual applications and virtual systems in IBM PureApplication System, you need to configure the cloud as follows:

- Create **Cloud Groups** that have one or more compute nodes where virtual machines (VMs) are deployed
- Create a logical grouping of one or more IP Addresses, called an **IP Group**, that are then assigned to a cloud group
 - Deployed VMs are assigned IP addresses from the IP group
- Optionally, you can create **Storage Volumes** and attach to the deployed VMs for additional storage



IBM PureApplication System offers ability to create different isolated clouds, called cloud groups. Each cloud group can have its own access permission and service level agreements or SLAs. Cloud groups contain one or more Compute nodes where virtual Machines are deployed. Compute nodes can belong to one cloud group at any given time. If you want to move the Compute node to another cloud group, you must first remove it from the owning cloud group and then add it to the other cloud group.

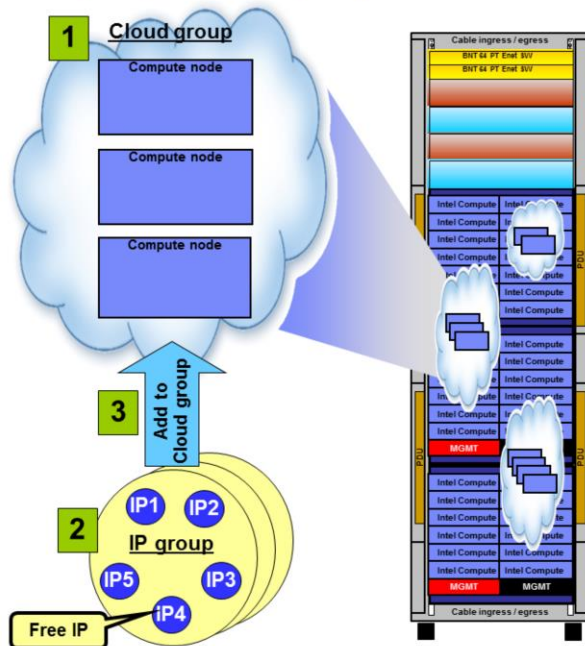
IP groups define one or more IP addresses. As Virtual machines are deployed in a cloud group, they need to be assigned an IP address. This IP address comes from the IP groups. Cloud groups need to have one or more IP groups. IP groups can be assigned to one Cloud group at any given time. A cloud group can have multiple IP groups assigned to it.

If the virtual machine needs additional storage outside its disk space, storage volumes can be defined and attached to the VM. The Client would then use the additional storage within the VM as any available disk.

The next few slides will pictorially walk through the process of creating Clouds within IBM PureApplication System.

Cloud configuration: Creating cloud groups and IP groups

1. Create cloud groups with one or more free compute nodes
2. Create one or more IP groups containing a set of IP addresses that will be assigned to VMs during deployment
3. Associate the IP groups to the Cloud group



4

Cloud configuration

© 2012 IBM Corporation

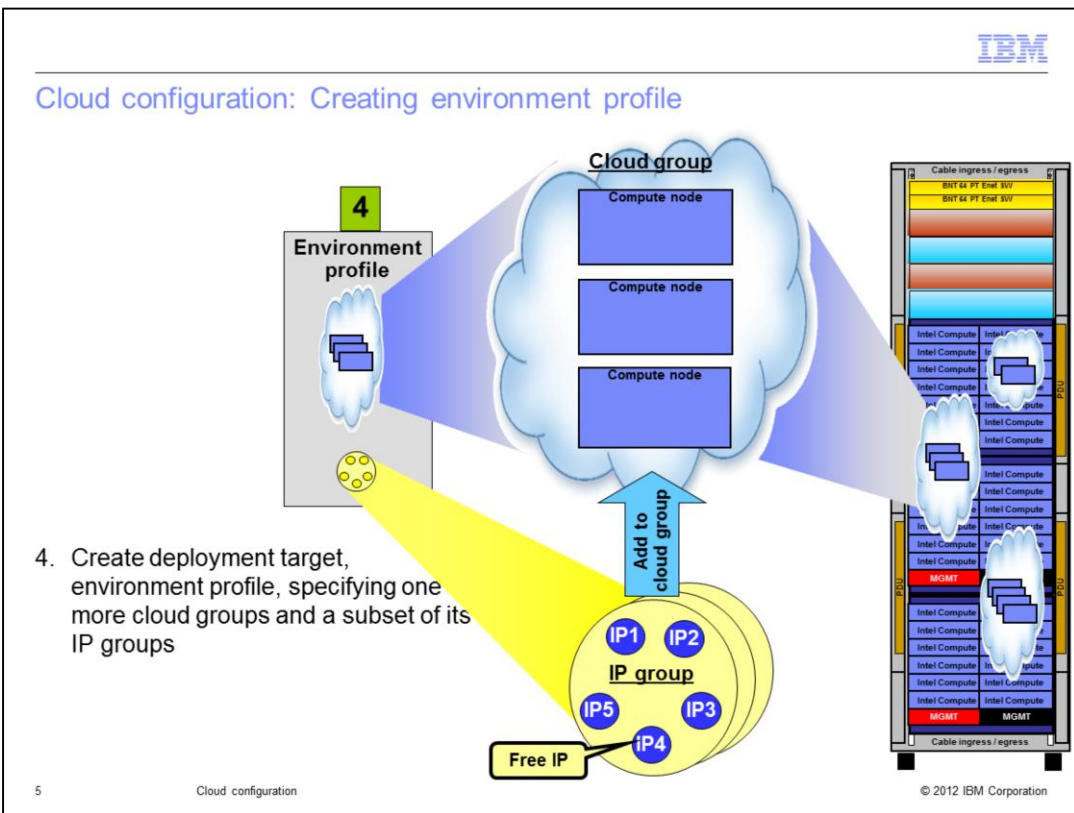
This slide shows the Cloud groups, IP groups and association of IP group to the cloud group.

To create different clouds, you must create a cloud group and add one or more free compute nodes to the cloud group. These Compute nodes will be the target for instances deployed to this cloud group. PureApplication System will manage and balance the load across the different Compute nodes based on SLAs provided by you. The second step is to create IP groups. IP groups contain one or more IP addresses.

The third step is to associate or add the IP group to the cloud group. Once the IP group is added to the cloud group, the IP addresses from the IP group are used to assign to the virtual machine as they get deployed to this cloud group. You can have more than one IP group associated with a cloud group. As VMs are removed, the IP address is freed up and returns to the pool to be used again for a future deployment.

Both, the Compute node as well as an IP group can belong to only one cloud group. If needed, they can be removed from one cloud group and added to another.

Cloud configuration: Creating environment profile

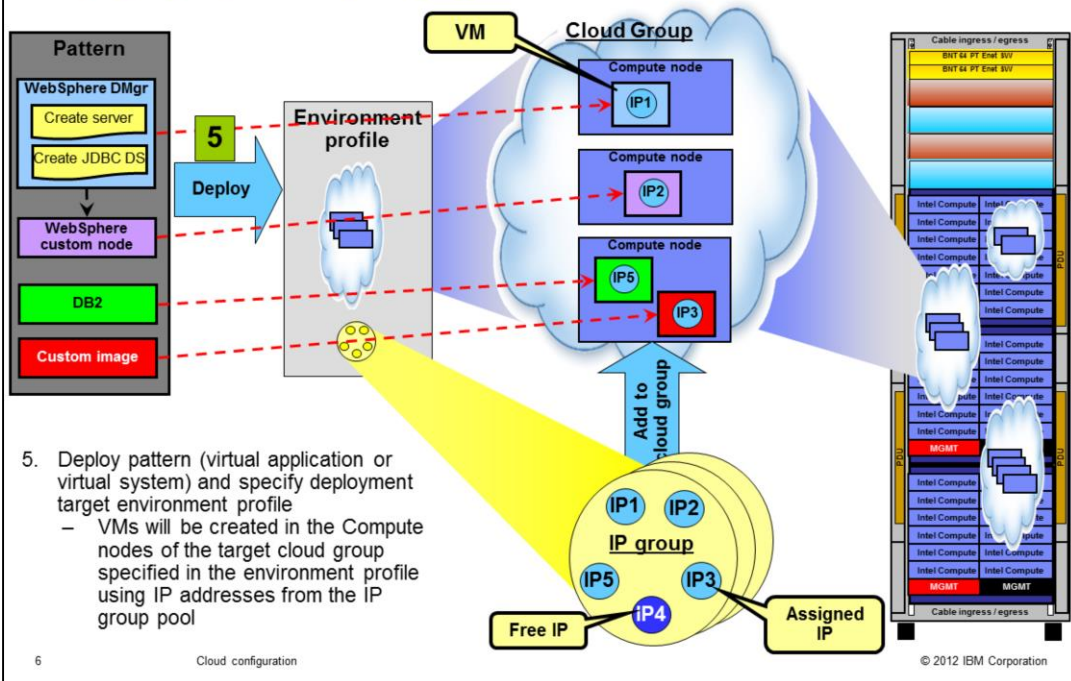


In PureApplication System, the patterns are deployed using environment profiles. Environment profiles allow client to define one or more cloud groups and a subset of its IP groups as a deployment target. In addition, one can specify the priorities of the VMs in that cloud group.

Step 4 is to create the environment profile from the Workload console, Cloud tab.

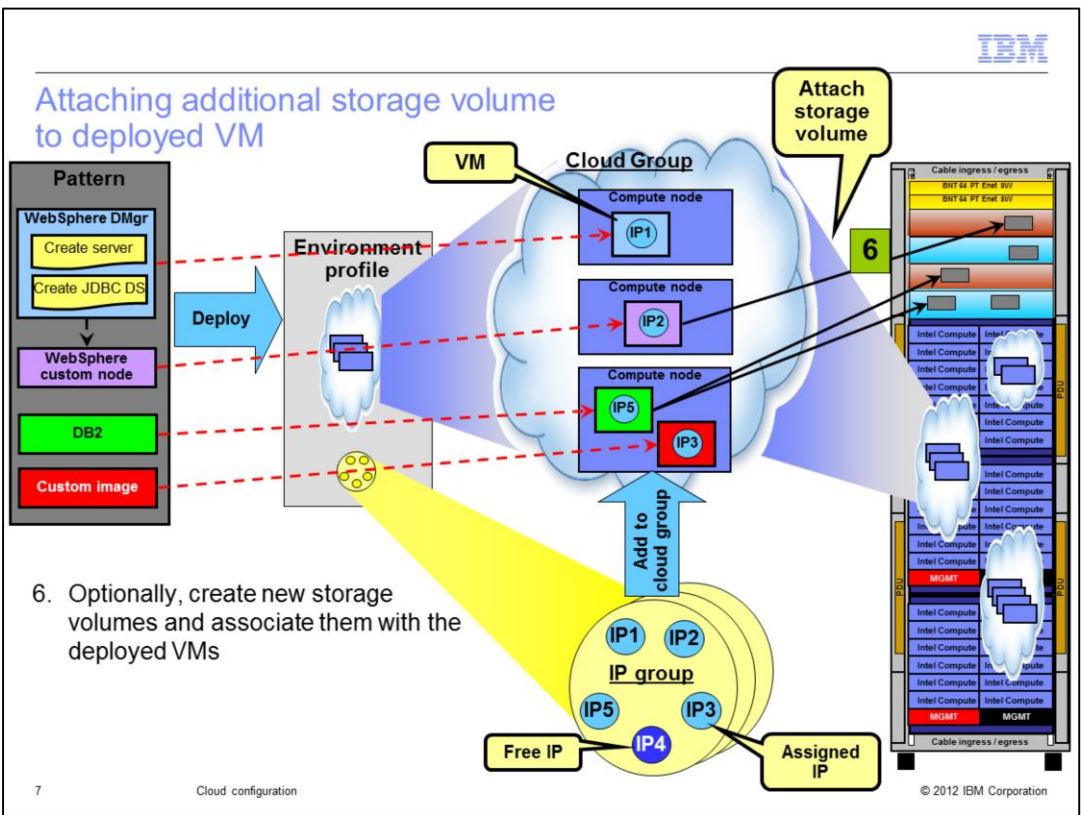
The deployer of a virtual application or virtual system pattern will need to specify the environment profile that is the target of the deployment.

Deploying a pattern specifying environment profile



The next step is to deploy the virtual system or virtual application pattern. At the start of the deployment, the deployer specifies the environment profile that is the target of the deployment.

Through the environment profile, the cloud group and the IP group are determined. Through the cloud group, the compute nodes are determined. Through the IP group, the IP addresses to assign to the virtual machines are determined. PureApplication System knows where to deploy the virtual machines that make up the deployed instance of the pattern and what IP address to use for those VMs.



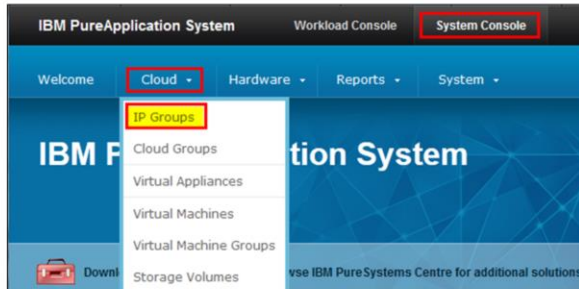
If you needed more storage for your virtual machines that make up the deployed instance, PureApplication System allows you to create a storage volume and associate that with the virtual machine. While a storage volume can only be assigned to one VM, a VM can have multiple storage volumes. Once assigned, the storage volume will appear as a disk available within the VM. This is shown in the slide by step 6.

Cloud configuration: IP groups

The next section discusses the IP groups and how to create and manage it.

IP groups - overview

- IP groups is a logical grouping of one or more IP addresses with the networking info (DNS, subnet, etc.)
- IP groups (and by default its IP addresses) are assigned to a cloud group
 - Cannot share IP groups across cloud groups
 - Cloud groups can be assigned multiple IP groups
- VMs of the deployments targeted to a cloud group will be assigned the IP address from the IP groups assigned to the cloud group
- IP groups is accessed using **System Console** → **Cloud** → **IP Groups**, as shown below



9


Cloud configuration

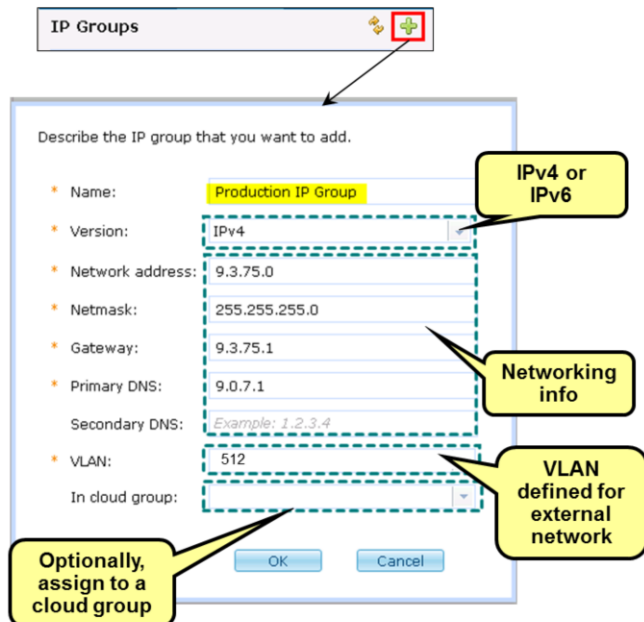
© 2012 IBM Corporation

An IP group is a pool of IP addresses. When IBM PureApplication System deploys a pattern into the cloud it will use IP addresses from the IP group that are defined as part of the IP groups. The IP groups are assigned to the cloud group.

IP groups can be accessed from the System Console's Cloud tab and IP groups.

Create new IP group

- Click the  sign to create new IP group
- Provide the must-have fields (*) as shown
- Specify the networking info
- Specify the VLAN that connects to external network
- Optionally can associate this IP group to the cloud group here or later in cloud group panel



IP Groups

Describe the IP group that you want to add.

* Name: Production IP Group

* Version: IPv4

* Network address: 9.3.75.0

* Netmask: 255.255.255.0

* Gateway: 9.3.75.1

* Primary DNS: 9.0.7.1

Secondary DNS: Example: 1.2.3.4

* VLAN: 512

In cloud group:

OK Cancel

IPv4 or IPv6

Networking info

VLAN defined for external network

Optionally, assign to a cloud group

10

Cloud configuration

© 2012 IBM Corporation

Once you open the IP groups, you will see the previously defined IP groups. The green plus icon is used to create a new IP group. Clicking on the plus icon will bring up a window in which you can enter initial values for your IP group.

The name of the IP group needs to be any unique value that you want IBM PureApplication System to manage this IP group by.

You can select IPv4 or IPv6 and based on the selection, additional parameters will need to be filled like Subnet address, netmask, gateway, primary DNS and secondary DNS. These are standard networking concepts when defining a subnet.

The next value to fill the VLAN that is defined for external networking configuration. Through the VLAN, the external connection is linked to the IP addresses specified in the IP group.

The last optional step is to add the IP group to the cloud group. You may add the IP group later or within the cloud group console. Note that the IP addresses can only be assigned to the virtual machines once its been added to the cloud group.

IP group Attributes

TestIPGroup1 Delete Report

Jobs: Pending jobs: 0 Started Jobs: 0 [View details...](#)

Description: (none)

State: Available

Created on: Jul 10, 2012 11:55:40 AM

Updated on: Jul 20, 2012 4:21:26 PM

Version: ipv4

Network address: [172.40.1.0](#)

Netmask: [255.255.255.0](#)

Gateway: [172.40.1.1](#)

Primary DNS: [172.40.1.1](#)

Secondary DNS: (none)

VLAN: 621

In cloud group: (none) [\[Edit \]](#)

IP addresses: 11 total - 0 available, 11 in use.

- [172.40.1.10](#) (172.40.1.10)
- [172.40.1.11](#) (172.40.1.11)
- [172.40.1.12](#) (172.40.1.12)
- [172.40.1.13](#) (172.40.1.13)

[\[show more \]](#)

Access granted to: [admin](#) [\[all \]](#) [\[remove \]](#)

Add by: IP Range

start ip:

end ip:

Modify network parameters

Modify VLAN

Specify IP addresses

Delete

Report

Add/remove cloud group

IP addresses already added

Add by IP range or host name

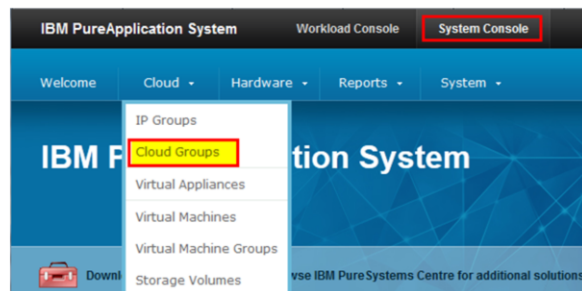
Once you have created the IP group, the IP group attribute panel comes up. Here you can modify some of the parameters like subnet information, VLAN and the cloud group, that you entered when creating the IP group. You can add/remove the IP address that is part of this IP group. IP address can be added as a single IP address or range of IP addresses.

Cloud configuration: creating cloud groups

The next section discusses the cloud groups and how to create and manage them.

Cloud groups - overview

- Logical grouping of computing resources (compute nodes) to target your deployments in PureApplication System
- Cloud groups contains one or more compute nodes
 - Cannot share compute nodes across cloud groups
- Cloud groups have the following customization attributes
 - Compute nodes to deploy VMs
 - IP groups for assigning IP to VMs for deployment
 - Definition of virtual machine and volume configurations you can choose for deployment
- Additional info on current resources in the cloud group is shown in the cloud group console panel
- Cloud groups are accessed through the **System Console** → **Cloud** → **Cloud Groups**, as shown below



13

Cloud configuration

© 2012 IBM Corporation

A Cloud group is a logical group of compute nodes on which pattern deployments can be targeted. By having multiple cloud groups, PureApplication System provides multi-tenancy and a level of isolation for different teams to deploy their patterns.

Compute nodes can only belong to one Cloud group. Cloud groups also contain one or more IP groups for assigning IP address to virtual machines of deployed patterns. One can also defined virtual machine and storage volume configuration that can be used during deployment. These definitions allows pre-canned size and capacity for the VMs and storage. More details on these configurations will be discussed later in this presentation.

Creating new cloud group

- Click “+” and in the panel, provide the needed attributes
- Name of the cloud group
- Type: set how much CPU/memory resources to allocate for the VMs in the compute nodes of this cloud group. Values are ‘Average’ or ‘Dedicated’
- Optional management VLAN: Internal VLAN used to isolate all management traffic
 - Keeps the internal communication within a cloud group isolated to that specific cloud group

The screenshot shows the 'Cloud Groups' panel in the IBM PureApplication System. A red box highlights the '+' icon in the top right corner of the panel. An arrow points from this icon to a dialog box titled 'Describe the cloud that you want to create.' The dialog box contains the following fields:

- Name: Production CG
- Description: Cloud for Production environment
- Type: Average (highlighted with a red box, with a callout showing 'Average' and 'Dedicated' options)
- Management VLAN ID: 224

At the bottom of the dialog box are 'OK' and 'Cancel' buttons. The page number '14' is visible in the bottom left corner, and the copyright notice '© 2012 IBM Corporation' is in the bottom right corner.

During the creation of a cloud group, you will need to specify the name, type, collocation policy and Management VAN id.

The Type field lets PureApplication System know how to pack the VMs in the compute nodes of the cloud group.

The management VLAN is chosen by you to be the internal VLAN PureApplication System used to isolate all management traffic. This allows the management traffic isolated from other communication.

Cloud group type - details

- Allows you to determine how to pack the workloads (VMs) on the compute nodes and allocate the processing power and memory to the workloads
- W1500

Cloud group types	vCPU versus pCPU	Limit of vCPUs per compute node	Memory	Usage
Dedicated	1 vCPU=1pCPU	16 vCPUs	No overcommit	This cloud group provides the best access to CPU and memory resources with a limited number of virtual machines.
Average	8 vCPU=1pCPU	128vCPUs	No overcommit	This cloud group type is the most recommended because it supports a reasonable amount of allocated CPUs.

- W1700

Cloud group types	vCPU versus pCPU	Limit of vCPUs per compute node	Memory	Usage
Dedicated	1 vCPU=1pCPU	28 vCPUs	No overcommit	This cloud group provides the best access to CPU and memory resources with a limited number of virtual machines.
Average	10 vCPU=1pCPU	280vCPUs	No overcommit	This cloud group type is the most recommended because it supports a reasonable amount of allocated CPUs.

On Intel there are 16 physical cores per compute node. On Power there are 32 cores per compute node, but 4 are reserved.

Through the cloud group type attribute, you can dictate how many VMs you want to pack in the compute nodes that are part of the cloud group. There are 2 possible values – dedicated and average.

In the dedicated mode, the virtual CPU specified for a VM is reserved and dedicated to the VM. On Intel if your VM had specified 16 virtual CPUs, that VM will take over an entire compute node. This can be typically used for production environment for applications that require some guaranteed CPU allocation.

Alternatively, the average type packs more VMs in the compute node. In the same Intel example, if your VM specifies 16 virtual CPU's, then it will only be allocated one eighth of the CPU on the compute node. With average then you can pack more VMs in the compute node. This is a recommended type for most development use cases.

Placement of workloads within a cloud (group)

- PureApplication System determines which compute node within the cloud group to place the VMs, based on several factors, some under your control and some based on the type of the workloads
 - Same types of nodes run on different compute nodes (where possible) within a cloud group
 - Minimize licensing usage –
 - Licensing is based on # of cores used
 - Pack more VMs of the same products on the same compute node, as possible
 - Balancing resources based on CPU, disk, network, and so on
- PureApplication System performs periodic rebalance of the VMs to spread the resources used across the compute nodes within a cloud group, while satisfying some of the factors listed here

Several factors play a role in deciding which compute node within a cloud group will host the VM. PureApplication System tries to spread the load of the VMs across the different nodes within the cloud group, while maintaining some restrictions. For example, nodes of the same type within a deployment is spread across compute nodes to better handle node failures. The algorithm tries to minimize the license usage since some software licensing is based on the number of cores used. It tries to pack more VMs of the same products on the same compute node as possible to optimize the storage I/O.

Over time, there can be imbalance of the compute node resources being used as deployments are started and stopped.

PureApplication System performs periodic rebalance of the VMs to spread the resources used across the compute nodes within a cloud group, while satisfying some of the factors listed here. The rebalance does not affect the access to the VM or its applications.

Cloud group attributes

The next section discusses the attributes of the Cloud groups.

Cloud group attributes

Cloud Groups

Search...

- Consolidation ⚠
- Dedicated ⚠
- Shared 🟢
- VirtualManagement ⚠

Shared Delete Report

Jobs: Pending jobs : 0 Active jobs : 0 View details...

Description: Default PureApplication Shared Cloud Group

State: ✔ Available

Created on: 6/4/12 9:12 PM

Updated on: 6/23/12 10:21 AM

Type: Dedicated

Cloud group policy: Availability

Fault tolerant: Active

Management VLAN ID: 3105

IP groups:

IP Groups	VLAN	Action
ipgroupRack15	635	[remove]

Add more...

Access granted to:

- ⊕ **Compute nodes:**
- ⊕ **Hardware information:** Physical CPUs allocation: 0%
Physical memory allocation: 0%
- ⊕ **Virtualization statistics:** Total virtual machines: 0
- ⊕ **Storage usage statistics:** Total volumes: 0
Storage utilization: 0%
- ⊕ **Virtual machines:**
- ⊕ **Virtual machine configurations:** 13 total
- ⊕ **Volume configurations:** (none)

- Can view/manage cloud group attributes with Cloud Administrator Read or Full permission

18
Cloud configuration
© 2012 IBM Corporation

You can view the cloud group attributes if you have the cloud read or full permission. Certain attributes are only available with full permission. The panel shows the list of all cloud groups. Clicking a cloud group will open the attribute panel. Cloud group administrator with full permission can delete the cloud group. Other administrators can view the reports of the cloud group.

Many of these attributes are described in details over the next few slides.

Cloud group attributes - General info, IP groups and access control

Shared

Jobs: Pending jobs : 0 Active jobs : 0 [View details...](#)

Description: (none)

State: Available

Created on: 6/16/12 4:43 PM

Updated on: 6/29/12 9:52 PM

Type: Average

Cloud group policy: Availability

Fault tolerant: Active

Management VLAN ID: 1236

IP groups:	IP Groups	VLAN	Action
	IPGroup24Subnet	624	[remove]

Add more...

Access granted to: admin [all]
dn [read]

Compute nodes:

19 Cloud configuration © 2012 IBM Corporation

Add/remove IP groups for this cloud group

Grant access

Cloud group attribute panel shows all the details of the cloud groups and the functions that can be performed. This and the next few slides will show the details of the Cloud group attributes. You can view the jobs associated with this Cloud group, set the type to Average, dedicated, view the VLAN ID specified when creating the Cloud group. You can add or remove the IP groups that are associated with this cloud group. The Cloud administrator can grant read-only or all access to other users and groups. The next few slides will get into more details of other attributes.

Cloud group attributes – Add or remove compute nodes

- Compute nodes can be added or removed from the cloud group, as shown below
 - Can only add compute nodes that are not assigned to any other cloud group
- Shows status of each compute node – whether its running, stopped, or in quiesce mode
- Links let you dive into the details of the compute node

Compute nodes:

Status	Compute nodes	Physical CPU	Physical memory	Action
Available	SN#23ZWR94	1%	1%	
Quiesce	SN#23AWVL6	0%	1%	
Available	SN#23AWTM9	1%	1%	
Available	SN#23AWTC9	6%	1%	
Available	SN#23AWVD6	2%	1%	
Available	SN#23AWTY7	2%	1%	

1 - 6 of 6 items

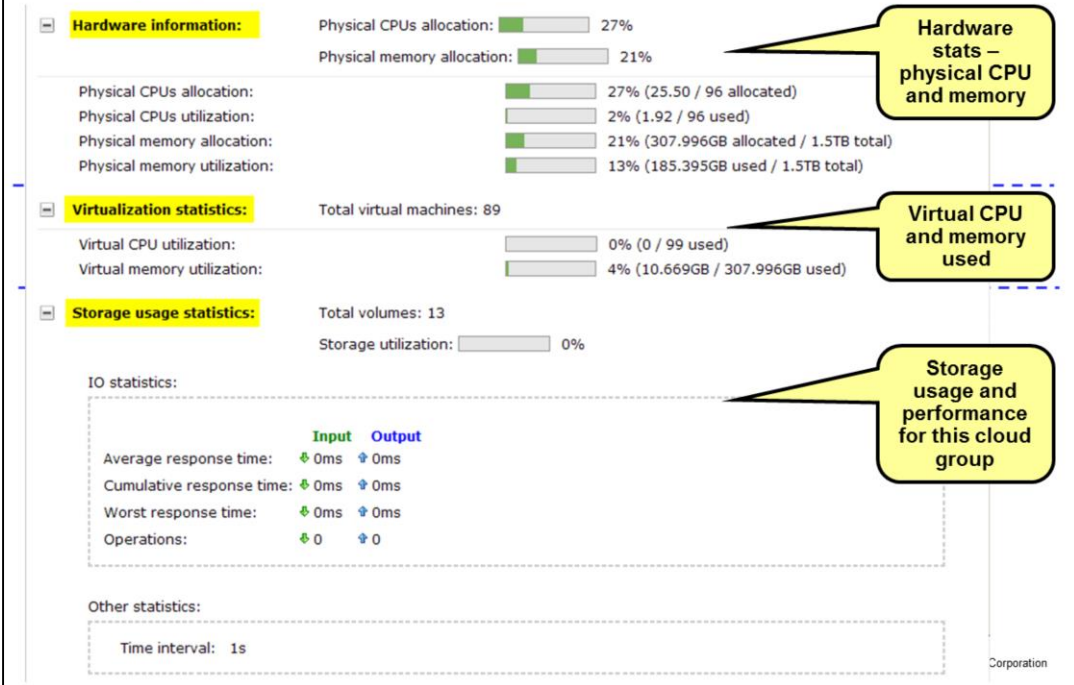
Callouts:

- Compute node status:** Points to the status column.
- Compute node names and links to its details:** Points to the link in the 'Compute nodes' column.
- Add compute node to cloud group:** Points to the 'Add more...' button.
- Remove compute node:** Points to the 'Remove' icon in the 'Action' column.

20 Cloud configuration © 2012 IBM Corporation

Cloud groups contain one or more compute nodes where the workloads run. The compute node section of the cloud group attribute panel shows how to add or remove compute nodes from the cloud group. It has the name of the compute node and link that takes you to the details of a selected compute node. It also shows the status of the compute node whether its running, stopped or in quiesce mode. It also shows the overall resources in terms of physical CPU and memory used for each of the compute nodes.

Cloud group attributes – Performance stats



Part of the attribute panel provides additional info on hardware utilization, virtualization statistics and storage usage statistics, as shown on the slide.

Hardware information shows the aggregate hardware usage for all the compute nodes in the cloud group for physical CPUs and Memory.

Virtualization statistics shows the virtual resource utilization for CPU and memory along with the number of VM instances.

The storage usage statistics shows the number of storage volumes allocated in the cloud group and the aggregate performance of the storage for read and write operations.

Cloud group attributes – Virtual machines

Virtual machines:

Name	Status	Virtual memory utilization	Virtual CPU utilization
ipas-lpar-024-029-1391-Enterprise_Application-was.21340999116520	Running	5%	0%
ipas-lpar-024-033-1392-ICCT-icct.11341071152305	Running	20%	0%
ipas-lpar-024-034-Standalone-GRJExtendwIfix-1393	Running	10%	0%
ipas-lpar-024-035-Standalone-8003-ESX32-RHEL64-WEB20-SS-1429	Stopped	0%	0%
ipas-lpar-024-036-DMGR-yiwen.70023_dev_cluster-1471	Running	11%	1%
ipas-lpar-024-037-1431-Caching-Master.11341195338703	Running	2%	1%
ipas-lpar-024-202-OS Node-ISSW_ICCT_VS-1319	Stopping	0%	0%
ipas-lpar-024-203-Standalone-Test_VSP-1303	Stopping	0%	0%

1 - 8 of 89 items

Virtual machine name and links to its details

The virtual machines section displays all the VMs in this Cloud group along with its status, and virtual resource utilization. By clicking in the VM name, you can go to the details of the VM and perform authorized functions.

Virtual machine (VM) configuration and storage volume configuration

VM configurations

- VM configurations specify the CPU (physical and virtual) and virtual memory to be used when creating a VM during virtual system pattern deployment
- You select pre-defined VM configuration to suit the needs of the workload on the VM
- Currently not being exploited by Virtual System patterns

Storage volume configurations

- Volume configuration specify the storage volume disk size to be used when creating new storage for the cloud group
- Every cloud group has a set of default VM configurations and volume configurations, or, you can create new configurations
- For virtual applications, PureApplication System will select the best VM configuration for the needs of the load

Virtual machine configuration allows you to specify the physical and virtual CPU and memory for the different VMs during deployment. PureApplication System provides few default configurations and allows defining custom configurations.

Storage volume configuration allows you to specify a pre-defined storage volume size that then can be used when creating new storage volumes for this cloud group.

These configurations provide an easy way to create the VMs and storage volumes based on some standard configurations that you can define for your organization.

Cloud group attributes – Virtual machine configurations

The screenshot displays the 'Virtual machine configurations' panel in the IBM Cloud console. At the top, a red box highlights the title 'Virtual machine configurations: total: 1'. Below this is a table with the following data:

Name	Virtual memory	CPU count	Description	Action
perftest	3 GB	2	perftest	

Below the table, a red box highlights the '+ Create a new configuration' button. A yellow callout bubble labeled 'Pre-defined VM configurations' points to the table. Another yellow callout bubble labeled 'Delete VM configuration' points to the delete icon in the 'Action' column. A third yellow callout bubble labeled 'Create new VM configuration' points to the '+ Create a new configuration' button. Below the table, a dialog box titled 'Describe the virtual machine configuration that you want to create.' is shown. It contains the following fields:

- Name:
- Description:
- CPU count:
- Virtual memory(MB):

At the bottom of the dialog are 'OK' and 'Cancel' buttons. The page number '24' is in the bottom left, 'Cloud configuration' is in the bottom center, and '© 2012 IBM Corporation' is in the bottom right.

The panels define the existing virtual machine configurations and can create new virtual machine configuration definitions. You specify the virtual CPU count, the virtual memory in MB and the physical CPU count as part of the definition.

Cloud group – Storage volume configurations

- Use volume configurations to pre-define the storage disk size of the VMs within this cloud group – used during storage volume creation

The screenshot displays the 'Volume Configurations' section of the IBM Cloud management console. It features a table with two existing configurations and a 'Create a new configuration' button. A callout box labeled 'Delete volume configuration' points to the delete icon in the 'Action' column of the 'volumeConfig2' row. Another callout box labeled 'Create new volume configuration' points to the 'Create a new configuration' button. To the right, a dialog box titled 'Describe the volume configuration that you want to create.' is shown, with fields for 'Name' (filled with 'volumeConfig3'), 'Description', and 'Size(MB)' (filled with '2000'). 'OK' and 'Cancel' buttons are at the bottom of the dialog.

Name	Size	Action
volumeConfig2	1.95 GB	
volumeConfig1	1000 MB	

[+ Create a new configuration](#)

Describe the volume configuration that you want to create.

Name:

Description:

Size(MB):

25

Cloud configuration

© 2012 IBM Corporation

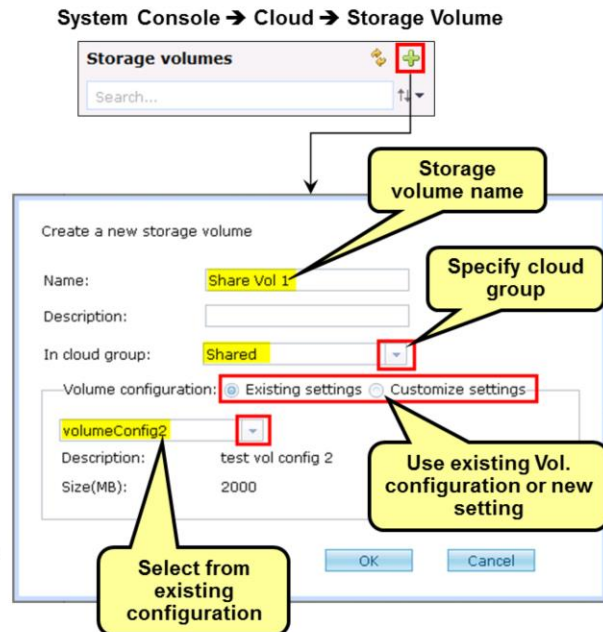
The panels show the existing storage volume configurations and can create new storage volume configuration definitions. You specify the name and the size in MB. When creating a new storage volume, you can then use one of the pre-defined configurations.

Cloud configuration: Storage volumes

The next section discusses how to create and manage the storage volumes.

Storage volumes – Additional storage for virtual machine

- You can create storage volumes and attach them to deployed VMs
 - Provides additional storage outside the VM
 - Attachment done in the VM panel
- Storage volumes are defined for a cloud group
- Creating new storage volume
 - Provide Name
 - Specify cloud group
 - Can use pre-defined volume configuration (defined as part of cloud group) or use new customize settings
- Storage volume can only be assigned to one VM
 - Can be re-attached to VM in the same cloud group



27

Cloud configuration

© 2012 IBM Corporation

There can be situations where you need additional storage for a running VM. This provides additional storage outside the VM. This is accomplished by creating a storage volume in the cloud group that has the running VM, and then attaching it to a deployed VM.

To create a new storage volume, go to System Console, Cloud tab and then storage volumes. To create a new storage volume, click the “+” sign.

In the create a new storage volume panel, specify the name for the storage volume, select the cloud group which will contain the storage volume, and specify a storage volume configuration.

The storage volume configuration can be one of the pre-defined configurations, or you could use customize setting to specify a different storage volume disk size.

Once the storage volume is created, you can go to the virtual machine panel and attach the storage volume to the VM. Storage volume can only be assigned to one VM at any given time. You can detach a storage volume from a VM, and re-attach to another VM in the same cloud group.

Storage volume attribute panel

The screenshot shows the 'Storage volume attribute panel' for a storage volume named 'Shared Vol1'. The panel includes the following fields and actions:

- Jobs:** Pending jobs : 0, Active jobs : 1. A 'View details...' link is present.
- Description:** (none)
- Status:** Available
- Created on:** 7/6/12 12:20 PM
- Updated on:** 7/6/12 12:20 PM
- Virtual machines:** (none)
- Size:** 2000
- In cloud group:** Shared
- Storage controller:** 78N0K3

Callouts provide additional context:

- 'Delete or view report' points to the 'Delete' and 'Report' buttons at the top right.
- 'View jobs details' points to the 'View details...' link.
- 'VM this storage is allocated' points to the '(none)' value in the 'Virtual machines' field.
- 'Storage size in MB' points to the '2000' value in the 'Size' field.
- 'Owning cloud group' points to the 'Shared' value in the 'In cloud group' field.
- 'Controller for this storage volume' points to the '78N0K3' value in the 'Storage controller' field.

28

Cloud configuration

© 2012 IBM Corporation

The panel on this slide shows the details of the storage volume. You can delete the storage volume or view its report. You can view the status of the storage volume, the owning cloud group, the size of the storage and the controller for this storage volume. This storage volume can be allocated to any VM in that cloud group. If this storage volume is allocated to a VM, the owning VM is specified in this panel.

Attaching storage to deployed VM

- Storage is attached from the VM attribute panel
- Storage is mounted at the specified mount point in the deployed VM
- Can reach a given VM attribute panel from **System Console → Cloud → Virtual Machine**
- Alternatively, can access VM panel from **Workload Console → Instances → <Your instance> → Virtual Machines → Manage**

The image shows a screenshot of a cloud management console. On the left, a table lists virtual machines. One VM is selected, and a 'More' button is clicked, leading to an 'Attach volume' dialog box. The dialog box has a 'Mount point' field containing '/share_vol' and an 'Attach' button. A callout bubble points to this field with the text 'Attach to a mount point within the VM'. On the right, the VM's detail panel is shown. It includes fields for Status (Running), Created Time, Updated on, Virtual CPU, Virtual memory, In cloud group, On compute node, Virtual Appliance, Virtual machine configurations, Virtual machine group, Storage volumes, and IP addresses. The 'Storage volumes' field is a dropdown menu currently set to 'no volume'. A callout bubble points to this dropdown with the text 'Select storage volume to attach to this VM'.

This slide shows how to attach a storage volume to a deployed VM. From the VM panel, you can select any un-attached storage volumes available in the cloud group where the VM is running. Once the free storage volume is selected, you will need to specify the mount point at which the storage is attached within the deployed VM.

You can get to the VM detail panel from System Console to the Cloud tab and then to the virtual machine.

You can also reach the VM detail panel from the Workload Console Virtual System or Virtual Instance panel. Expanding the virtual machines for a virtual system or virtual Instance, you can select Manage and then “More.” That link will take you to the VM detail panel from where you can attach storage volume to the VM.

***Environment profiles
and
deploying virtual application and system patterns***

This section provides details for the environment profile which are deployment targets for virtual application and virtual system patterns

Environment profiles

- Different users and groups in an organization using the same PureApplication System, need or are given different cloud resources on which they can deploy their workloads
- In PureApplication System, this is achieved using environment profiles – it allows the cloud administrator to customize the following:
 - Placement of virtual machines – through cloud group
 - Assignment of IP addresses to virtual machines
 - Names used for deployed VMs
 - Limits on the resources used by the VMs
 - Specify the priority of the VMs
- Different groups are then assigned and given permission to one or more environment profiles
- During deployment of virtual application and virtual system patterns, deployer will specify the environment profile to be used as the deployment target

Environment profile provides a way for a Cloud Administrator to assign different Cloud resources to different teams and groups within an organization that require different resources and priorities.

In PureApplication System, environment profile allows you to customize the placement of the VMs through the use of cloud groups, specify how to assign IP addresses, give flexibility to customize VM names, specify the limits on the resources consumed by the VM using this profile, and specify the deployment priorities for the VMs.

Different environment profiles can then be assigned to different groups thus allowing control over where they can deploy and how much resources they can consume.

When virtual application and virtual system patterns are deployed, the deployer is required to specify the environment profile for which they have been granted permission.

Creating environment profiles

- Create environment profiles from **Workload Console** → **Cloud** → **Environment Profile** menu
- Specify the name and environment attribute
 - Can use the environment attribute for filtering and classification purposes during pattern deployment
- Once created, complete the environment profile definition from the details panel
- Need Cloud Administrator full permission to create a new environment profile

Environment Profiles

Describe the environment profile you want to create

Name:

Description:

Hypervisor type:

Environment:

OK Cancel

32

© 2012 IBM Corporation

Cloud Administrator can create a new environment profile from the Workload Console, Cloud tab. Required field is name and environment. Environment attribute is used for filtering and classification purposes during deploying patterns. Some examples of environment attribute are Development, Test, Production, and so on.

Environment profile – details

- Modify the environment attribute for filtering
- Specify virtual machine name format
- Specify how IP addresses is provided to the VM
 - From the subset of IP groups assigned to the cloud
 - Or to be decided by pattern deployer
- Deployment priorities
- Deploy to cloud group and subset of its IP groups
- Environment limits
- Access permission - give read or write permission to other users/groups

Hypervisor type: PureScale_ESX

Environment: All (Attribute for filtering)

Created on: May 27, 2012 7:22:43 AM

Current status: ✔ Environment profile can now be used for deployments

Updated on: May 31, 2012 11:34:01 AM (Customize VM name)

Virtual machine name format: None provided (Specify how IP addresses are assigned to VMs)

IP addresses provided by: IP Groups (Specify how IP addresses are assigned to VMs)

Deployment priority: Platinum - High(17) Medium(12) Low(7) (Priority)

Deploy to cloud groups:		Name	Alias
<input type="checkbox"/>	Shared		Shared
<input checked="" type="checkbox"/>	In use	MyIPGroupRack4	MyIPGroupRack4
			Subnet
			172.17

Remove current cloud group first

Environment limits (VM limits)

Access granted to: Default Admin [owner] (Grant access)

Add more... (Grant access)

Deployment cloud group

33 Cloud configuration © 2012 IBM Corporation

Configure the environment profile after creating it. In the environment profile details panel, you can modify the environment attribute used for filtering. You can specify the format of the virtual machine name. You can specify how the IP address are assigned to the VMs deploying using this environment profile. It can be either from the selected IP groups or can be decided by the pattern deployer at deployment time. If a VM needs to be at a fixed IP address, then the deployer can choose to specify the IP address at deployment time.

You can specify a priority of the VMs using this environment profile. This priority is defined as Platinum, Golden, Silver and Bronze. This priority along with the deployment time high, medium and low priority are combined to give an overall priority. The overall priority is then used for optimization, placement and moving the VMs across the compute nodes when and if the need arises. More on how the priorities are used is discussed later.

You can control the limits for the VMs being deployed through this environment profile. VMs cannot be deployed using this environment profile if their resource demands are larger than the limits set in the environment profile. This allows a Cloud Administrator who creates the environment profile more control on the max resources being used by the VMs.

Lastly, as an owner of the environment profile, you can grant access read or write permission to other users and groups.

Deploying virtual applications/system patterns

- When deploying patterns, you select an environment profile, not a cloud group
- Environment profiles specifies the cloud group and IP groups for deployment

Refresh Deploy Edit Clone Lock Delete

Deploying virtual application or virtual system pattern

Describe the virtual system you want to deploy.

Virtual system name:
Day Trader VSys Instance

Choose environment
 IP version: IPv4
 IP version: IPv6

Choose profile
 Type: All
 Profile: BPG profile
 Priority: High

Schedule deployment
 Configure virtual parts

OK Cancel

Select environment profile as the deployment target

When clicking the “Deploy” button to deploy a virtual system or virtual Instance pattern, one of the required parameters is the environment profile to be used for VM deployment when the pattern is provisioned.

The panel shows the deploy options and the selection of the environment profile. The Type field allows you to filter the environment profiles. The deployment priority of High, Medium and Low is used along with the environment profile deployment priority to determine the overall priority of the VMs.

Priority of VMs

- The priorities of VMs play a role when moving the VM from one compute node to another within the same cloud group
- VM priorities are calculated from these factors:
 - Max priority defined in environment profile – Platinum, Gold, Silver and Bronze
 - VM priority defined during deployment – High, Medium, Low
- Effective priority is a combination of both priorities
 - Example: Platinum-High, Platinum-Low, Silver-Medium, Bronze-Low, and so on
- Platinum-High > Gold-High > Silver-High > Bronze-High > Platinum-Med > Gold-Med > ... > Platinum-Low > Gold-Low > Silver-Low > Bronze-Low
- If no space available in other compute nodes, some VMs with low priority are not moved
 - Client will have to make space and restart the VM on the other compute nodes

The environment profile priority, defined as Platinum, Golden, Silver and Bronze along with the deployment time priority defined as High, Medium and Low are combined to give an overall priority.

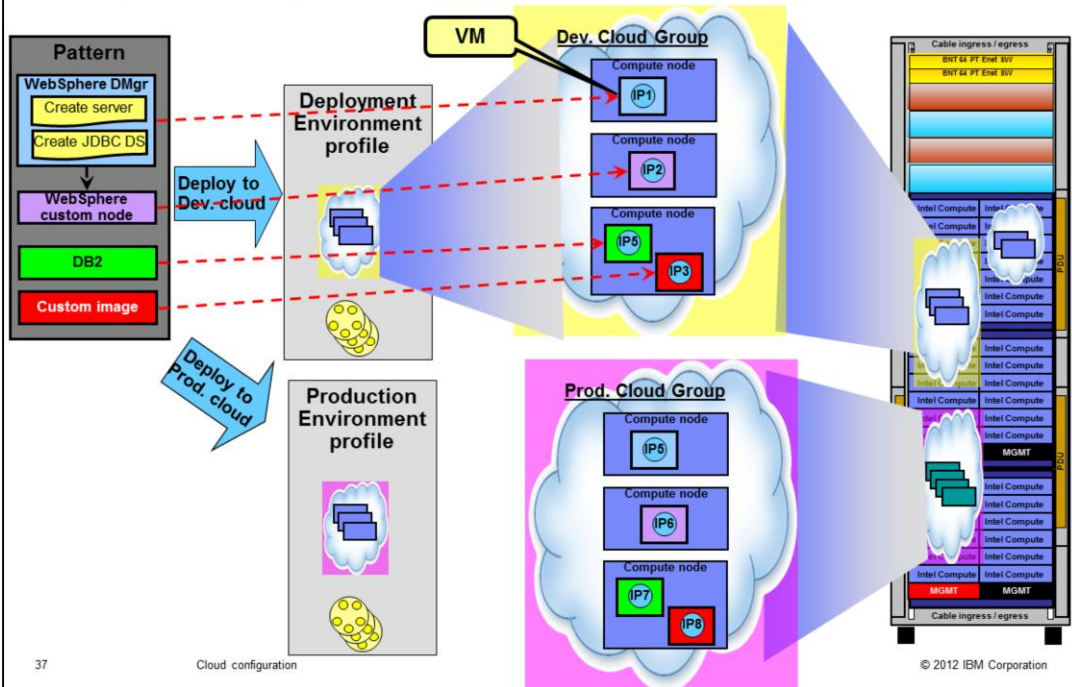
The priorities are from High to Medium to Low, and within each category, its from Platinum to Gold to Silver to Bronze. For example, the Platinum-High priority is higher then Gold-High which is higher than Silver-Low which is higher than Platinum-Medium which is higher than Gold-High which is higher than Silver-high and so on.

This overall priority is then used to decide which VM within the cloud group moves across compute nodes before other lower priority VMs, in case of failures.

Carving out PureApplication System into different clouds

This section discusses some of the best practices.

Example – carving out cloud groups for different purposes



PureApplication System gives you the flexibility to carve out the compute nodes into different clouds and how much you can allocated to each of the clouds based on the type of workloads being deployed to those clouds. The example shows how one can possibly carve out different cloud groups for different purposes, for example development and production purposes, as shown in the slide.

While the slide shows a virtual system pattern, the same applies to other deployment patterns like virtual applications or virtual appliance.

As shown on the slide, one possible option is to create different cloud groups for each purpose, and create a separate environment profile. The pattern can first be tested in the development cloud group and once verified, the same pattern can then be deployed to the production. You control how you want to carve out the different cloud groups in the rack and policies that drive your deployment SLAs.

While not shown on the slide, you can also create different environment profiles specifying the same cloud group with same or different IP addresses and same of different SLAs. This allows using the same cloud group to deploy workloads for different purposes, like development and system test, each having their own IP addresses and SLAs but using the same cloud.

Section

Best practice

This section discusses some of the best practices.

Best practices and suggestions

- For production and critical applications, have two or more compute nodes in a cloud group for failover of VMs within the cloud group
 - When a compute node fails, PureApplication System moves the VMs to other compute nodes (if available) within the same cloud group
 - Try to have the compute nodes in different chassis to further improve the resiliency
- In order to have the ELB proxies use a fixed set of IP addresses, create a dedicated IP pool for ELB proxies and use the environment profile with the dedicated IP pool for ELB Proxy deployment
 - This will help the front end routers to know the IP addresses of the ELB proxies, even if ELB proxies were being scaled up and down

Here are some best practice tips. In production and for more critical applications, use a cloud group with at least two or more compute nodes. In case of a compute node failure, PureApplication System will automatically move the VM to other compute nodes within the cloud group with affecting the application or the user. The VMs are moved based on the priorities discussed before. If there is no room for the VMs of lower priority, those VMs will be stopped. The user will then have to make some space and start the VMs.

It also further improves the resiliency and failover capability if the compute nodes were from different chassis, to the extent possible.

The front end sprayers that send requests to the ELB proxies with PureApplication System need to know the IP addresses of the ELB proxies so that as the ELB proxies are scaled up and down, or if there is a failure in one of the ELB proxy, the Frontend IP sprayer can try the proxies on the other IP addresses. One good practice would be to create a dedicated small IP group for use of ELB proxies. If you know you have a certain number of ELB proxy instances based on the scaling policy you supplied, you can have a IP group with slightly higher number of IP addresses than the max ELB proxy instances. The environment profile that is used to target the ELB proxies can then use the dedicated IP group. This way, changes to the ELB proxies from IP addresses perspective can be handled by the IP sprayer.

Section

Summary

This next section provides the summary of this presentation.

Summary

- Cloud groups are groupings of compute nodes where deployments are targeted
- PureApplication System IP groups are pools of IP addresses specified by the client
- IP groups are assigned to cloud groups from where deployed VMs are assigned IP address
- Environment profile is used to target virtual system and virtual application deployments
- Though the environment profile, you specify the cloud group and how to use the IP address for the deployed VMs
- Some best practice tips were shared

In summary, many concepts of how to create clouds in PureApplication System were covered, around cloud groups, IP groups, storage volumes and environment profiles. This presentation also covered the deployment of workloads using environment profiles and some best practices were shared.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and PureApplication are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.