



Updated: June 2020

ARM template version 6.2 is based on:

- [VMware Tanzu Greenplum Database version 6.8.1](#)
- Or
- [VMware Tanzu Greenplum Database version 5.27.1](#)

Overview

VMware Tanzu Greenplum on Azure uses Azure Resource Manager (ARM) template to automate the deployment of a cluster in the Azure cloud.

Cloud Native and Cloud Managed Tools and Features

Utility/Feature	Description
bouncer	Command line utility that simplifies using pgBouncer connection pooler.
gpcompute	Grow or shrink compute independently of storage.
gpprow	Grow storage independently of compute.
gpmaintain / gpcronmaintain	Automates routine maintenance of VMware Tanzu Greenplum such as vacuum and analyzing tables.
goptional	Installation tool for installing optional components such as MADlib and Command Center.
gppower	Automates the Pausing and Resuming a cluster to save on infrastructure costs.
gprelease / gpcronrelease	Automates the upgrade of VMware Tanzu Greenplum, any optional components installed, and the cloud utilities.
gpsnap / gpcronsnap	Automates the execution of snapshot backups and restores using the cloud vendor's native snapshot utilities.
Self-Healing	Automatic instance replacement and recovery in a failure event.

Licensing

Greenplum on Azure BYOL is licensed by the number of cores deployed and it is important to note that in Azure, **1 vCPUs equals 1 core**. Customers will total the number of vCPUs deployed in the cluster and then purchase that many subscription cores from VMware Tanzu Software.

Calculating Usable Storage

To determine the storage size of the Deployment, multiply the number of segment instances in the cluster times the number of disks per instance, and times the size of disk chosen to get the raw storage. Divide the raw storage by two because of mirroring and then multiply this by .7 to leave space for transaction and temp files.

Instance Type	Instance Count	Disk Size
Standard_H8	64	2TB

(64 Segment Instances * 2 Disks * 2TB Disks) = 256 TB Raw Storage

(256 / 2) * .7 = 89.6 TB Usable Storage

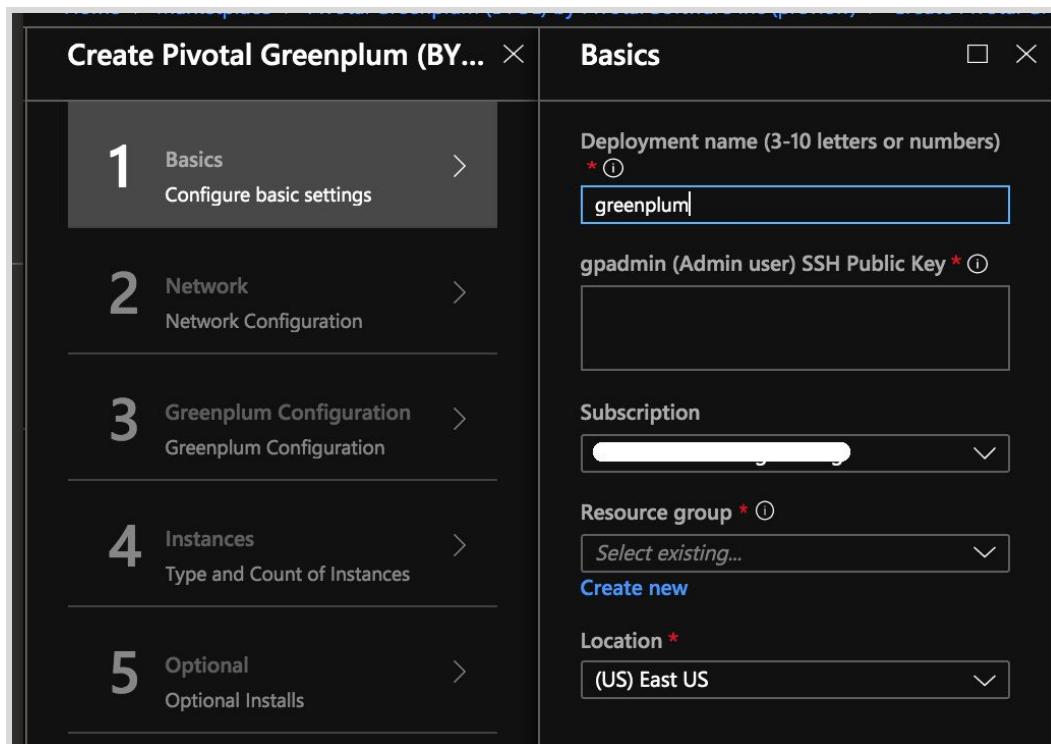
Overview	1
Cloud Native and Cloud Managed Tools and Features	1
Licensing	1
Calculating Usable Storage	2
Deploying Greenplum on Azure	6
Parameters - Basics	6
Deployment Name	6
gpadmin (Admin user) SSH Public Key	6
Subscription	6
Resource Group	6
Location	7
Parameters - Network	7
Internet Access	7
Subnet Mask	7
Parameters - Greenplum Configuration	7
Database Version	7
Database Name	8
Disk Type	8
Parameters - Instances	8
Master Instance Type	8
Master Instance Disk Size	9
Segment Instance Type	9
Segment Instance Disk Size	9
Segment Instance Count	9
ARM Template	9
Connecting	11
phpPgAdmin	12
SSH Access	12
Client Tool	13
Additional Resources	13
Deployment Logs	13
Validation	13
Greenplum on Azure Additional Features	13
bouncer	13
gpcompute	14
gpgrow	15
gpmaintain / gpcronmaintain	15
goptional	16
gppower	16

gprelease / gpcronrelease	16
gpsnap / gpcronsnap	17
Self Healing	18
Segment Healing	18
Standby-Master Healing	18
Master Healing	19
The PXF Extension Framework (PXF)	19
Patching Linux on Azure	19
Time Zone	20
Core Dumps	20
Greenplum on Azure Technical Details and History	21
Azure Resources	21
Image	21
Resource Group	21
Virtual Network	21
Subnet	21
Network Security Group	21
Storage Account	21
Public IP Address	21
Network Interface	22
Availability Set	22
Virtual Machines	22
Storage	22
Root and Swap	22
Data Storage	22
Data Storage	22
Master Storage	22
Diagram	23
Version History	23
Version 6.2	23
Fixes	23
Enhancements	23
Version 6.1	23
Fixes	23
Enhancements	24
Version 6.0	24
Enhancements	24
Version 5.1	24
Fixes	24
Enhancements	24

Version 5.0	24
Fixes	24
Enhancements	25
Version 4.2	25
Enhancements	25
Version 4.1	25
Fixes	25
Enhancements	26
Version 4.0.0	26
Fixes	26
Enhancements	26
Version 3.5	26
Enhancements	26
Version 3.4	27
Fixes	27
Enhancements	27
Version 3.3	27
Enhancements	27
Version 3.2	27
Enhancements	27
Version 3.1	27
Fixes	27
Enhancements	27
Version 3.0	28
Fixes	28
Enhancements	28

Deploying Greenplum on Azure

Parameters - Basics



Deployment Name

This identifies the VMware Tanzu Greenplum Deployment. The default is "greenplum".

gpadmin (Admin user) SSH Public Key

This is your public key used to ssh to the Master instance. Password authentication has been disabled on all instances so the only way to connect via ssh is by creating a key-pair. Use a utility such as ssh-keygen to create your public and private key pairs. Once created, provide the public key as the parameter and reference the private key when connecting via ssh.

Subscription

Pick from one of your existing Azure Subscriptions.

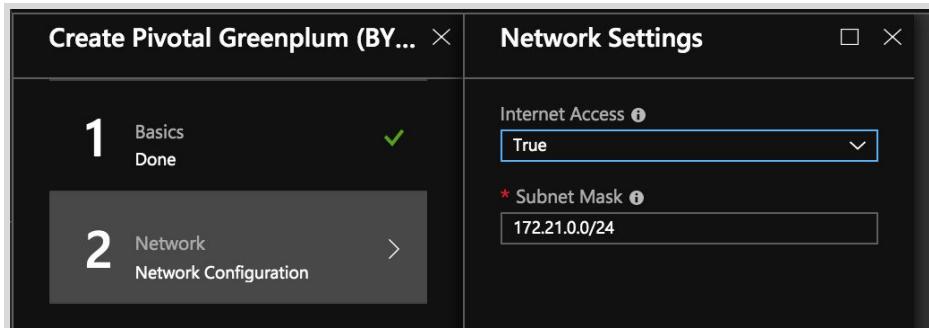
Resource Group

Microsoft only allows a Deployment to be created in an empty Resource Group so either create a new one or pick an existing, empty Resource Group for your deployment.

Location

Provide the location in Azure where you want your cluster to be deployed.

Parameters - Network



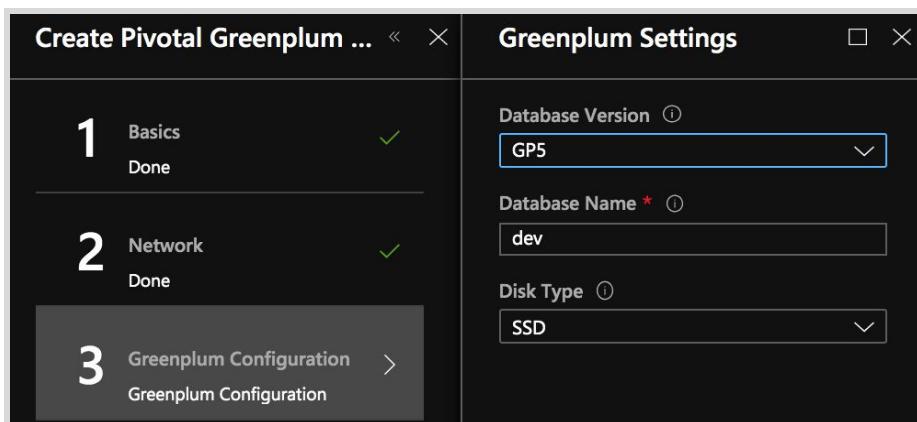
Internet Access

True means a Public IP address will be created for the Master instance with ports 22, 5432, 28080, and 28090 open to the Internet. False means the Master will not have a Public IP address created and a jump box will be needed to access the cluster.

Subnet Mask

The Subnet CIDR block that the subnet will cover. You can typically use the default since the deployment is using a dedicated vnet but this added feature gives you more control over how the IP addresses are assigned in the subnet.

Parameters - Greenplum Configuration



Database Version

Pick the major database version you wish to use. GP5 and GP6 are currently available with this release.

Database Name

This defines the default database name for your Greenplum cluster. Optional components will be installed in this database.

Disk Type

Standard HDD and Standard SSD options are available. Recent performance benchmarks indicate that SSD disks will provide better performance.

Helpful Tip

You can leverage the **gpsnap** utility to change the disk type after your deployment has been completed. This is done by creating a snapshot with gpsnap and then restoring that snapshot but using a different disk type.

Parameters - Instances

The screenshot shows the 'Create Pivotal Greenplum ...' wizard with five steps:

- 1 Basics**: Done
- 2 Network**: Done
- 3 Greenplum Configuration**: Done
- 4 Instances**: Type and Count of Instances (highlighted)
- 5 Summary**: Pivotal Greenplum (BYOL) by Piv...

In the 'Instance Settings' step, the following configuration is shown:

- Master Instance Type ***: 1x Standard H8
8 vcpus, 56 GiB memory
[Change size]
- Master Instance Disk Size**: 500
- Segment Instance Type ***: 1x Standard H8
8 vcpus, 56 GiB memory
[Change size]
- Segment Instance Disk Size**: 2000
- Segment Instance Count**: 0

Master Instance Type

It is recommended to use the HPC instance types if available in your region. The Master node typically does not need a large number of vCPUs or RAM so picking the smallest instance type available is usually sufficient.

Master Instance Disk Size

Specify the size of the Data Disk(s) attached to the Master instance. In a Single-Instance deployment, the Master will have two disks while in a Multi-Instance deployment, there will only be one.

Segment Instance Type

It is recommended to use the HPC instance types if available in your region. Furthermore, pick the same "family" of instance type. If you decide to use the HPC instance type for your Master, also use HPC for the Segment Instances.

Use the Standard_H8 or Standard_D13_v2 instance types for moderate concurrency needs and use the Standard_H16 instance type when the cluster has a high concurrency demand.

Instance Type	Memory	vCPUs	Data Disks	Segs Per Instance
Standard_H8	56	8	2	2
Standard_D13_v2	56	8	2	2
Standard_H16	112	16	2	2

Helpful Tip

You can leverage the **gpcompute** utility to change the instance type after your deployment has been completed. This is currently available when using the Standard_H8 and Standard_H16 instance types.

Segment Instance Disk Size

Each Segment instance will have two data disks and this specifies how large each disk is. The sizes listed are based on the cost structure in Azure.

Segment Instance Count

The number of Segment instances in the Deployment can range from 0 up to 64. When picking 0, the Master will act as a Single-Instance cluster with Master and Segment processes in one virtual machine.

ARM Template

Deployment is very simple in the Azure Marketplace. Simply provide the parameters in the user interface and then submit the ARM template to create the deployment.

Once the deployment has been submitted, go to the new Resource Group that you created for this deployment.

Next, click on Deployments in your Resource Group.

There are two Deployments that are created. One is used by Microsoft to track deployments while the second is the Greenplum cluster. Click on the second deployment (the one that doesn't start with "pid-").

Once complete, Click on the other Deployment and then "Outputs".

	Value	Action
database Admin UserName	gpadmin	
command Center UserName	gmon	
master Host	40.87.6.8	
database Port	5432	
database Password	ssh to the Master Host as gpadmin with your private key ...	
phpPgAdmin URL	http://40.87.6.8:28090	
command Center URL	http://40.87.6.8:28080	
optionalInstalls	Install Optional Components with the goptional CLI utilit...	

The above shows the information needed to get started using the new VMware Tanzu Greenplum on Azure cluster.

Connecting

Connecting can be done with the web based GP Browser database client, ssh or with an external database client tool like pgAdmin 4. The Deployment Output for Master Instance, Port, Admin UserName, and Password information used to connect to Greenplum.

The database password is no longer available in the Deployment output and must be retrieved from the Master host. Use your private key and ssh as gpadmin to the Master host and then "cat ~/.pgpass".

Example:

```
cat ~/.pgpass
localhost:5432:pgbouncer:gpadmin:vmcbigojqen2e
mdw:6432:*:gpadmin:vmcbigojqen2e
*:6432:gpperfmon:gmon:vmcbigojqen2e
```

The randomly generated password in this example is vmcbigojqen2e. If you wish to change the gpadmin database password, you need to do the following:

1. Change the database password

```
psql -c "alter user gpadmin password 'newpassword';"
```

2. Update the .pgpass file with the new password

```
localhost:5432:pgbouncer:gpadmin:newpassword
mdw:6432:*:gpadmin:newpassword
*:6432:gpperfmon:gmon:newpassword
```

phpPgAdmin

This is a VMWare Tanzu enhanced version of the popular phpPgAdmin web based SQL tool. It has been configured to start automatically on the Master Instance. The URL is provided in the Deployment Output. Connection is simple by using the admin "gpadmin" user and the provided database password.

The screenshot shows the phpPgAdmin interface. On the left, there's a tree view of databases: 'Pivotal Database' (selected), 'dev' (selected), 'Schemas' (expanded), containing 'gp_toolkit' and 'public' (expanded), with 'Tables', 'Views', 'Sequences', 'Functions', 'Full Text Search', and 'Domains'. Below the tree is 'gpperfmon'. On the right, the 'Query Results' panel shows a table with 10 rows of segment data. The table has columns: dbid, content, role, preferred_role, mode, status, port, hostname, address, and datadir. The first few rows show segments 1, 6, and 5. Row 10 is highlighted. Below the table is a 'SQL?' button and a SQL editor with the query 'select * from gp_segment_configuration'. At the bottom, there are buttons for 'Edit SQL' and 'D'. A message bar at the top says '– You are logged in as user "gpadmin"'. The URL in the browser bar is 'Not Secure | 40.117.197.36:28090 Greenplum Command Center - HTML5'.

Helpful Tip

phpPgAdmin uses Apache HTTP server and does not include an SSL certificate. If you plan on using this tool across the Internet, it is highly recommended that you configure Apache with an SSL certificate.

SSH Access

Password authentication has been disabled so you must use the Key Pair you created for the deployment. To connect with ssh, use your private key that matches the public key you provided when creating the deployment. Connect as "gpadmin" which is the administrator user for Greenplum. The message of the day provides detailed information about the deployment as shown below.

```
*****
Greenplum Version: 6.8.1
Master Data Directory: /data1/master/gpseg-1
Master: mdw
Standby: sdw1
Segment Hosts: sdw1, sdw10, sdw11, sdw12, sdw13, sdw14, sdw15, sdw16, sdw17,
sdw18, sdw19, sdw2, sdw20, sdw21, sdw22, sdw23, sdw24, sdw25, sdw26, sdw27,
sdw28, sdw29, sdw3, sdw30, sdw31, sdw32, sdw33, sdw34, sdw35, sdw36, sdw37,
sdw38, sdw39, sdw4, sdw40, sdw41, sdw42, sdw43, sdw44, sdw45, sdw46, sdw47,
sdw48, sdw49, sdw5, sdw50, sdw51, sdw52, sdw53, sdw54, sdw55, sdw56, sdw57,
sdw58, sdw59, sdw6, sdw60, sdw61, sdw62, sdw63, sdw64, sdw7, sdw8, sdw9
Total Segments: 128
Computed storage in GB: 46930
*****
Note: Be sure to switch to gpadmin:
sudo su - gpadmin

[gpadmin@mdw ~]$
```

Client Tool

Connecting with a remote client utility like pgAdmin 4 is also very easy to do using the Master public IP address and password provided in the Deployment Output.

Additional Resources

Installation of VMware Tanzu Greenplum on Azure includes detailed logs plus supplemental installs and validation scripts that can be executed after the initial installation is complete.

Deployment Logs

Logs of the deployment can be found in: `/opt/pivotal/greenplum/rollout.log`. These logs are on every instance but the Master instance will have more detailed logs of the database initialization.

Validation

Validation includes scripts to run industry standard benchmarks of TPC-H and TPC-DS. It also includes scripts to validate the disk and network performance of the deployment using the VMware Tanzu Greenplum utility "gpcheckperf".

Greenplum on Azure Additional Features

bouncer

pgBouncer is a load balancing utility that is included with Greenplum. This utility allows far greater connections to the database with less impact on resources. It is recommended to use pgBouncer instead of connecting directly to the database. More information on pgBouncer is available in the Greenplum documentation.

The bouncer utility automates the starting, stopping, pausing, and resuming of pgBouncer. It is recommended to use this utility to manage pgBouncer on your cluster.

pgBouncer is configured to listen on port 5432 which is the default port usually used by Greenplum. Greenplum has been configured to listen on port 6432.

Authentication has been configured to use "md5" which is an encrypted password. Create users and assign passwords in Greenplum as normal and pgBouncer will authenticate users with the database passwords you set.

Pooling has been configured for "transaction" with max client connections of 1000 and max database connections to 10. These settings can be changed but these defaults provide a good starting point for most installations.

Configuration and logs for pgBouncer are located in /data1/master/pgbouncer on the Master Instance.

Note that for JDBC connections, you may need to "search_path" to the ignore_startup_parameters configuration item in the ini file.

Connections can optionally be made with SSL to secure connections from your client to the database.

Actions available with bouncer include start, stop, pause, and resume.

```
bouncer <action>
```

Note: pgBouncer may not be compatible with some applications. You can disable pgBouncer by making these changes.

1. Stop pgBouncer with "bouncer stop"
2. Stop Greenplum Command Center with "gpcc stop"
3. Stop the database with "gpstop -a"
4. On the mdw, change \$MASTER_DATA_DIRECTORY/postgresql.conf port=6432 to port=5432.
5. On mdw, change /home/gpadmin/.bashrc PGPORT environment variable to 5432.
6. On mdw, change all 6432 values to 5432 in /home/gpadmin/.pgpass
7. Source /home/gpadmin/.bashrc
8. Start the database with "gpstart -a"
9. Start Greenplum Command Center with "gpcc start"

gpcompute

The "gpcompute" utility enables you to add or reduce the compute capacity of your cluster independently of your storage. The command completely automates and integrates with Azure to change the instance type of the segment instances.

The status command will show you the current instance type of your segment instances and also show the instance type of the current launch template.

```
gpcompute status
```

The command will accept these two different instance types in Azure: Standard_H8 and Standard_H16.

```
gpcompute <instance_type>
```

The Standard_H8 instance type is ideal for single-user and moderate concurrency workloads while the Standard_H16 instance type provides nearly the same single-user performance but an increase in concurrency performance.

The "gpcompute" utility enables you to add or reduce the compute capacity of your cluster independently of your storage but is not yet available for Azure.

Note: Clusters with the Standard_D13_v2 instance type do not support gpcompute. Performance with the Standard_D14_v2 with 2x more RAM and 2x more CPU did not increase the performance by a significant enough amount to warrant supporting it.

gpgrow

Increase the segment instances' disk size in your cluster completely online with this utility. Segment Instances have 2 data disks per each and the data size is specified in GB. Azure only allows increasing a disk size and this is not an online activity. The command will restart the segment instances automatically to complete the task.

Note: This utility does not support growing the Master Instance disk size.

```
gpgrow <new disk size>
```

gpmaintain / gpcronmaintain

The gpmaintain utility automates basic maintenance of Greenplum. It analyzes all tables using analyzedb, vacuums tables that are in need, and maintains the catalog.

The gpcronmaintain utility is executed via cron and looks at the schedule file in /usr/local/greenplum-cloud/conf/gpcronmaintain.conf. By default, this runs at 4:00 AM every Saturday and can be configured to run more or less often.

goptional

```
*****
Pivotal Greenplum on Azure Optional Installer
*****
[INFO] The Optional installer can restart Greenplum during the installation process.
[INFO] Database components will be installed in "dev".
[INFO] Use PGDATABASE environment variable to set database name.

Type the number that corresponds to the software you would like to install.
Picking software to install that is already installed will result in a re-installation.

1) phpPgAdmin [INSTALLED]          6) plr
2) CommandCenter [INSTALLED]        7) postgis
3) DataSciencePython               8) plcontainer
4) DataScienceR                   9) pivotal_greenplum_backup_restore
5) madlib                         10) Exit
#?
```

This utility simplifies installing optional components after the deployment has been completed. Simply run "goptional" to see the optional installation options. This tool is also used in conjunction with gprelease to upgrade or reinstall already installed optional packages.

phpPgAdmin and Command Center are installed with every deployment automatically while MADlib, Data Science Python, Data Science R, PL/R, PostGIS, PL/Container, and Backup/Restore can optionally be installed with goptional.

gpower

The gpower command makes it easy to pause (stop) a cluster to save on infrastructure costs. The command stops the database and all segment instances with gpower pause and resume a cluster with gpower resume.

Pausing and suspending with gpower:

```
gpower pause
```

```
gpower resume
```

This command will stop/start and deallocate/allocate the segment instances, stop/start the database, stop/start PXF (if it was running when paused), stop/start Command Center (if it was running when paused), and pause/resume the Segment Auto-Scaling-Group.

gprelease / gpcronrelease

This gprelease utility upgrades a Greenplum on Azure cluster to the latest minor release available. The tool automatically downloads the binaries, copies it to the instances in the cluster, stops the cluster, installs the new version, and then starts the cluster again. The tool automatically executes goptional so that optionally installed packages are re-installed or upgraded to a compatible version.

This gpcronrelease utility is the scheduling component of gprelease. It checks to see if a new release is available and updates the message of the day to indicate a new version is available. By default, this runs in cron weekly on Sunday at 1:00 AM in the local timezone.

gpsnap / gpcronsnap

The gpsnap utility leverages Azure disk snapshots to quickly create a database backup. In order to get a consistent view of the database, the command first stops the database and then executes Azure commands in parallel to create a snapshot for each disk volume. After the disk snapshots are created, the database is restarted. This usually only requires a database outage of 30 seconds to 2 minutes.

Each snapshot is tagged so that when a restore is desired, the disk volumes get attached and mounted to the correct instances. This is done automatically with gpsnap.

Due to restrictions in Azure, you can not use your snapshot backup with another cluster.

Due to restrictions in Azure, you can not copy snapshots from one Region to another.

This utility manages creating, listing, deleting, and restoring snapshots. Please note that creating or restoring a snapshot will restart the database.

The actions available with gpsnap are list, create, delete, and restore. Delete requires an additional parameter of the snapshot_id you wish to delete, and restore requires the snapshot_id and the disk type you wish to use.

```
gpsnap list
gpsnap create
gpsnap delete <snapshot_id>
gpsnap restore <snapshot_id> <disk_type>
```

Helpful Tip

You can leverage the **gpsnap** utility to change the disk type after your deployment has been completed. This is done by creating a snapshot with gpsnap and then restoring that snapshot but using a different disk type.

This gpcronsnap utility manages the automatic execution of gpsnap. By default, there will be a cron job that runs every 10 minutes and uses the configuration file:
`/usr/local/greenplum-cloud/conf/gpcronsnap.conf` to determine if a snapshot is needed or not.

`gpcronsnap.conf`

```
# This utility will only execute gpsnap once on the day(s) of the week you
specify below.
```

```
#maximum number of snapshots; delete the oldest when max reached  
max_snapshots=4  
  
#snapshot day of week (1..7); 1 is Monday  
#to specify daily, use (1 2 3 4 5 6 7)  
snapshot_dow=(7)  
  
#time of day to run the snapshot  
#do not schedule a time where the snapshot may not finish until the next day  
#do not schedule a time that conflicts with gpcronmaintain  
snapshot_time=03:00
```

As shown above, the default schedule is a weekly snapshot on Sunday at 3:00 AM in the local timezone. Four snapshots will be retained before the oldest snapshot will be automatically deleted.

Self Healing

Azure automatically has "service healing" which is a process where a bad instance gets replaced automatically. The bad instance will be shut down and a new instance will be brought online automatically. The new instance will retain all disks as well.

VMware Tanzu Greenplum on Azure will automatically run the commands needed to heal the database. It is done by the startup script /opt/pivotal/greenplum/azure_heal.sh.

Monitoring of the execution of the healing script on each instance can be done with this operating system command:

```
sudo journalctl -u greenplum.service
```

Segment Healing

When a segment fails, the instance on restart will automatically detect that the database is up but needing to recover the one instance. It will run gprecoverseg automatically and when ready to rebalance, it will then pause pgBouncer so current queries can complete, and then it will rebalance the cluster. Lastly, it will resume pgBouncer.

Standby-Master Healing

In the event that the first segment instance were to fail which also has the Standby-Master process, the self-healing process will recover automatically. The Standby-Master process is restarted automatically. Database activity may continue during the Standby-Master Healing process.

Master Healing

In the event that the Master were to fail, the new instance executes a few Greenplum utilities. The process first will fail over to the Standby-Master process running on the first segment instance followed immediately by a database shutdown. The Standby-Master is then returned back to a Standby-Master role and the new Master instance is set as the Master.

The database is up and operational at this point but the database statistics were lost in this process so users are still not allowed to connect. The pgBouncer load balancer is restarted and normal database activity may resume.

The PXF Extension Framework (PXF)

The PXF Extension Framework (PXF) provides parallel, high throughput data access and federated queries across heterogeneous data sources via built-in connectors that map a Greenplum Database external table definition to an external data source. This Greenplum Database extension is based on PXF from Apache HAWQ.

PXF has been pre-configured in your cluster with settings stored in /usr/local/greenplum-pxf/. To start PXF in your cluster, simply execute this command as gpadmin on the Master (mdw) instance.

```
pxf cluster start
```

And to stop PXF:

```
pxf cluster stop
```

Self-Healing and upgrades via gprelease have been enhanced to restart PXF if the process was found to be running.

More information about PXF can be found here:

https://gpdb.docs.pivotal.io/6-8/pxf/overview_pxf.html
https://gpdb.docs.pivotal.io/5270/pxf/overview_pxf.html

Patching Linux on Azure

After you have deployed your cluster in Azure, you may need to patch the operating system to address vulnerabilities and in order to do that, you have to take a few extra steps compared to an on-premise installation.

First stop the database, pgBouncer, PXF (if started), and Command Center (if started).

```
gpstop -a  
bouncer stop
```

```
pxf cluster stop  
gpcc stop
```

Next, use "gpssh -f all_hosts.txt" to connect to all instances via ssh. The all_hosts.txt file is in the gpadmin home directory on the Master Instance. Execute the command to run the yum update.

```
gpssh -f /usr/local/greenplum-cloud/conf/all_hosts.txt "sudo yum update all -y"
```

Note: Be patient as this executes on all instances in parallel. You will not see any output until the command completes on all instances.

Once complete, you can restart the virtual machines.

Now you can start the database, pgBouncer, PXF (if desired), and Command Center (if desired).

```
gpstart -a  
bouncer start  
pxf cluster start  
gpcc start
```

Time Zone

The timezone for a deployment defaults to UTC. If you wish to change the timezone, first stop the database, pgBouncer, PXF (if started), and Command Center (if started).

```
gpstop -a  
bouncer stop  
pxf cluster stop  
gpcc stop
```

The following example changes the timezone to EST. Be sure to pick a valid value from /usr/share/zoneinfo/.

```
gpssh -f /usr/local/greenplum-cloud/conf/all_hosts.txt "sudo rm -rf  
/etc/localtime; ln -s /usr/share/zoneinfo/EST /etc/localtime"
```

Now you can start the database, pgBouncer, PXF (if desired), and Command Center (if desired).

```
gpstart -a  
bouncer start  
pxf cluster start  
gpcc start
```

Core Dumps

Core dumps are enabled and it makes it easier for Support to debug issues with the database.

Each virtual machine instance has a /var/coredumps directory owned by the gpadmin user. These dumps are retained for up to 7 days and then will get automatically purged by the cronjob script /opt/pivotal/greenplum/dump_maintenance.sh.

Greenplum on Azure Technical Details and History

Azure Resources

Image

The Image is based off of CentOS 7.6 with Accelerated Networking enabled. The image has all the software packages and pre-reqs for installing VMware Tanzu Greenplum and necessary add-ons including Intel-enhanced networking drivers.

Resource Group

This is a collection of resources in Azure that you want to manage as a single entity. It also specifies the Region. Only one Greenplum Deployment is allowed in the same Resource Group.

Virtual Network

A dedicated Virtual Network is used when deploying Greenplum on Azure. This isolates traffic for Greenplum and helps to ensure it remains secure.

Subnet

Specifies the IP address range for instances deployed in the Virtual Network. This is the IP address range used by the Interconnect traffic for Greenplum. DHCP is used to assign IP addresses to instances deployed in the Subnet.

Network Security Group

This specifies which ports and protocols are allowed between instances in the cluster as well as from the Internet. All traffic is allowed between all instances on all protocols. TCP Ports 28080 (HTTP), 22 (ssh), and 5432 (Greenplum) are exposed on the 1 to 2 Master instances.

Storage Account

This is used to enable boot diagnostics in Azure. Boot diagnostics are needed to enable serial access to the virtual machines.

Public IP Address

Static public IP Addresses are assigned to the 1 to 2 Master instances.

Network Interface

All instances in the cluster have a Network Interface, or NIC attached to it. The Master instance has a NIC with Public IP Address while the Segment instances do not.

Availability Set

An Availability Set is defined for the deployment and it specifies the number of Fault and Update domains. This is done to minimize the impact to Greenplum in the event of Azure restarting an instance.

Virtual Machines

These are the instances deployed in the cluster. It has Disks attached to each instance and specified in the ARM template. It also includes the deployment execution of the VMware Tanzu scripts executed on each instance.

Console access has been enabled by setting the root password. The default root password is the same as the randomly generated password for gpadmin but with the additional "#" character at the end.

Storage

Root and Swap

Storage for the root partition is fixed at 32 GB each. The swap partition uses the automatic temp disk created by Azure and is sized the same as RAM unless RAM is greater than 32 GB and then swap is set to 32 GB.

Data Storage

Disk are mounted with "rw,noatime,nobarrier,nodev,inode64,allocsize=16m 0 2" and blockdev read ahead of 16385. The scheduler is set to deadline.

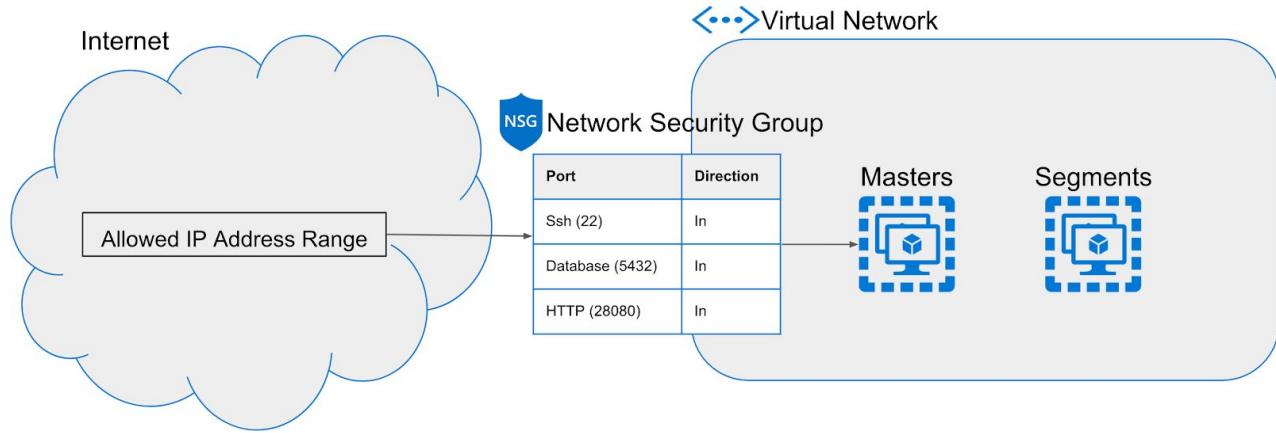
Data Storage

The storage option has been configured with Standard_LRS disks which are optimized for throughput instead of IOPs. Disk configuration has been optimized based on VM Type.

Master Storage

The data storage needed on the Master is much less than on the Segment instances so there will be only 1 "data" mount on the Master while the Segment instances will have 4.

Diagram



Version History

Version 6.2

Fixes

- Fixed condition where gpmaintain might fail due to a catalog lock.

Enhancements

- Updated documentation to reflect VMware Tanzu instead of Pivotal Software.
- Upgraded VMware Tanzu Greenplum GP6 to version 6.8.1 and GP5 to 5.27.1.
- Changed default database version to GP6.
- Upgraded Command Center versions to 6.2 and 4.10 for GP6 and GP5 respectively.
- Upgraded PostGIS version to 2.5.4 build 3 for GP6.
- Upgraded PL/Container to version 2.1.2 for GP6.
- Upgraded MADlib to version 1.17 for GP5 and GP6.
- Upgraded Cloud goptional to now install Backup/Restore utility.
- Added operating system tools for debugging any problems that might happen (gdb, strace, pstack, sysstat, lsof, dstat, iotop and atop).
- Updated operating system to the latest version.

Version 6.1

Fixes

- Installation of PL/Container during the initial deployment of a cluster could cause a timeout because of the time it takes to install all of the docker images. To prevent this, PL/Container can only be installed with the command line utility goptional rather than as a parameter in the template.

- Some web browsers prevent the use of self-signed SSL certificates so phpPgAdmin has switched to http rather than https. It is recommended that you install your own certificate and use SSL post-installation.
- The bouncer utility was displaying the wrong program name in the INFO messages. This issue has been resolved.
- The wrong database version would be displayed in the message of the day after a database upgrade with gprelease. This issue has been resolved.

Enhancements

- Upgraded Pivotal Greenplum GP6 to version 6.6.0 and GP5 to 5.26.
- Updated kernel to latest version available.
- Command Center is now installed by default.
- Azure now supports the gpcompute utility with the HPC (Standard_H8 and Standard_H16) instance types.

Version 6.0

Enhancements

- Upgraded Pivotal Greenplum GP6 to version 6.4.0 and GP5 to 5.25.
- Added Database Name parameter so customers can specify the default database name.
- Added core file generation to make it easier to debug issues with Support.

Version 5.1

Fixes

- Self-healing now restarts Command Center after a segment instance failure.

Enhancements

- Added PL/Container optional install.
- Added ability to specify the disk type when restoring a snapshot.
- Simplified the disk formatting process during deployment.
- Ran a yum update to update for new security patches.
- Upgraded Pivotal Greenplum GP6 to version 6.3.0 and GP5 to 5.24.
- Changed the count of segment instances from a multiple of 4 to 2.

Version 5.0

Fixes

- Updated logic for scheduled maintenance to ensure it doesn't execute while it is already running.
- Changed when the gpmon password is set during Command Center installation so that it prevents fatal login messages in the GUI.

- Updated pg_hba.conf file settings automatically during deployment and recovery so that it always has hostnames and not ip addresses. This makes snapshot restoration much more reliable.
- "set-hostname" was executed twice during a deployment for each instance.
- Added "create extension" when initializing PXF.
- Swap file would be deleted on rebooting machines.

Enhancements

- Added DatabaseVersion parameter to indicate either GP5 or GP6 major version to install.
- Added Pivotal Greenplum version 6.1.0 and upgraded GP5 to 5.23.
- Added ability to grow storage independently from compute with gpgrow.
- Added ability to pause / resume a cluster with gppower.
- Removed Standard_H16 and Standard_D14_v2 instance types and doubled the number of instances supported.
- Ran yum update to update to CentOS 7.7.
- Removed analyzedb when Master Instance is recovered as this is no longer recommended.
- Disabled gpperfmon for Command Center as starting with version 4.8, it is recommended to not use gpperfmon_install.
- Updated operating system limits settings.
- Added SSD disk option for data disks.

Version 4.2

Enhancements

- Upgraded to Pivotal Greenplum 5.20.1 and Command Center 4.7.0.
- Upgraded CentOS kernel to 7.6.
- Added deployment parameter to specify if the Master node will have a public ip address or not.
- The goptional tool now displays which database packages will be installed.
- Snapshot name enhancements so that a snapshots will be filtered by the database version (GP5 vs GP6).

Version 4.1

Fixes

- gpcrondelete maintains n snapshots determined in the gpcronsnap.conf file. The delete process would delete snapshots from other clusters to maintain the determined number of snapshots. Starting with 4.1, gpcronsnap will only delete snapshots that exceed the number and have been created by the currently running cluster.
- Changed disk entries in fstab file to use UUID which ensures proper mounting on reboots and snapshot restores.

Enhancements

- Created a new binary installer for the greenplum-cloud utilities. This makes it easier to upgrade these utilities.
- GP Browser renamed back to phpPgAdmin. It has enhancements for a Pivotal template and handling of partitioned tables. Installation is done via an RPM.
- Add caching of ssh keys between the segments back to the master.
- Added SSL to pgBouncer
- gprelease now updates cloud utilities
- Added new sysctl.conf entries for better memory management
- Added disk size parameter
- Changed drop down list of Subnet CIDRs to a text box to give more flexibility
- Reboots now reset the blockdev read ahead for data disks
- Enable Resource Groups by default instead of Resource Queues
- Automated maintenance with gpmaintain and gpcronmaintain
- Enabled console access by setting root password

Version 4.0.0

Fixes

- Disabled unnecessary nf_conntrack which could cause queries to fail under heavy concurrent load.

Enhancements

- Upgraded Greenplum to 5.17.0 and Greenplum Command Center to 4.6.0.
- Added support for PXF and included OpenJDK.
- Enhanced Greenplum Command Center installer to better handle future upgrades.
- Added Subnet parameter.
- Improved check_disk and check_network scripts in /opt/pivotal/greenplum/validation/performance

Version 3.5

Enhancements

- Upgraded Greenplum to 5.16.0 and Greenplum Command Center to 4.5.1.
- Enhanced gprelease/goptional to migrate existing packages that are installed and upgrade packages if needed.
- Improved database initializing performance so deployments are completed faster.

Version 3.4

Fixes

- Resolved issue where gprelease failed if a package directory exists but no packages are installed.

Enhancements

- Upgraded Greenplum to 5.13.0
- Removed standby-master host and process now runs on first segment host
- Added GP Browser SQL utility
- Added Standard_D14_v2 as an option for regions that don't support HPC instance types

Version 3.3

Enhancements

- Upgraded Greenplum to 5.12.0
- Upgraded Greenplum Command Center to 4.4.2
- Added gpsnap and gpcronsnap for executing database backups using Disk snapshots

Version 3.2

Enhancements

- Upgraded Greenplum to 5.10.2.
- Upgraded Greenplum Command Center to 4.3.1
- Added bouncer pause and resume functions
- Added Self-Healing
- Renamed gpupgrade and gpcronupgrade to gprelease and gpcronrelease
- gprelease and gpcronrelease enhanced for better integration to optionally installed components
- goptional tool created to make it easier to install optional components and also upgrade existing components.

Version 3.1

Fixes

- Identified and disabled UDP port conflict in Azure

Enhancements

- Upgraded Greenplum to 5.9.0.
- Upgraded Greenplum Command Center to 4.2.0.
- Added Data Science Python and R packages.
- Increased the number of segment nodes to 32.

- Optimized and standardized offering to Standard_H16 and Standard_H8 instance types.

Version 3.0

Fixes

- Patched the operating system for the Meltdown and Spectre vulnerabilities.

Enhancements

- Rewrite of Template and Scripts to align scripts, tools, and user experience of Greenplum running on other clouds.
- Manage Greenplum Database upgrades with two new tools; gpupgrade and gpcronupgrade
- Optional installs are now available as parameters to the ARM Template and visible in the Marketplace. Optional installs are still available post-installation.