IBM Db2 pureScale on Azure

Deployment Guide

Table of Content

Db2 pureScale on Azure Overview

IBM® Db2 pureScale is the high-end OLTP solution from the Db2 LUW family that provides the highest availability and scalability for both RHEL and SLES Linux environments.

IBM® Db2 pureScale reduces the risk and cost of business growth by providing extreme capacity, continuous availability, and application transparency.

These features include:

- Extreme Capacity - Start with a small cluster, grow as you need online.
- Continuous Availability – Design with integrated cluster manager to orchestrate automatic recovery on components failure, allowing transactions to continue other healthy members in the event of planned or unplanned member outage.
- Application Transparency - Avoid the risk and cost of application changes. As you grow the system, you do not need to make the application cluster-aware in order to get good scalability.
- Workload Consolidation - Map multiple workloads or databases to different subset members within the same Db2 pureScale instance with the capability to shift members from one workload to another in response to workload demands.

The self-managed IBM® Db2 pureScale on Azure supports Remote Direct Memory Access (RDMA) protocol over Converged Ethernet (RoCE) network along with TCP/IP. Db2 pureScale is a clustered system that facilitates collaborative and scalable processing of database requests. Db2 pureScale shares the same Db2 z/OS Sysplex architecture and design features which is the gold standard for availability. It delivers exceptional levels of database availability that satisfy the uniqueness of many applications. With Db2 pureScale, we can achieve enhanced performance, reliability and shared data management by adding members to the pureScale cluster.

Benefits of deploying Db2 pureScale on Azure

Some benefits of using Db2 pureScale on Azure –

1. Flexible and competitive pricing - You only pay for the resources you consume (i.e. CPU, storage, memory). You can resize the VM by adjusting its size.
2. Agility and scalability - Vertical and horizontal scaling allows you to scale the number of running instances up or down

3. Ease of deployment - Create a Db2 pureScale instance in minutes by launching the deployment from the Db2 pureScale Azure Marketplace page in Azure portal as simple as a click of a button
4. Premium storage and optimization - The majority of the Db2 pureScale Azure deployments come with NVME disk for the operating system and Elastic SAN for Db2 filesystems with High performance and low-latency disk support.

Architecture Diagrams

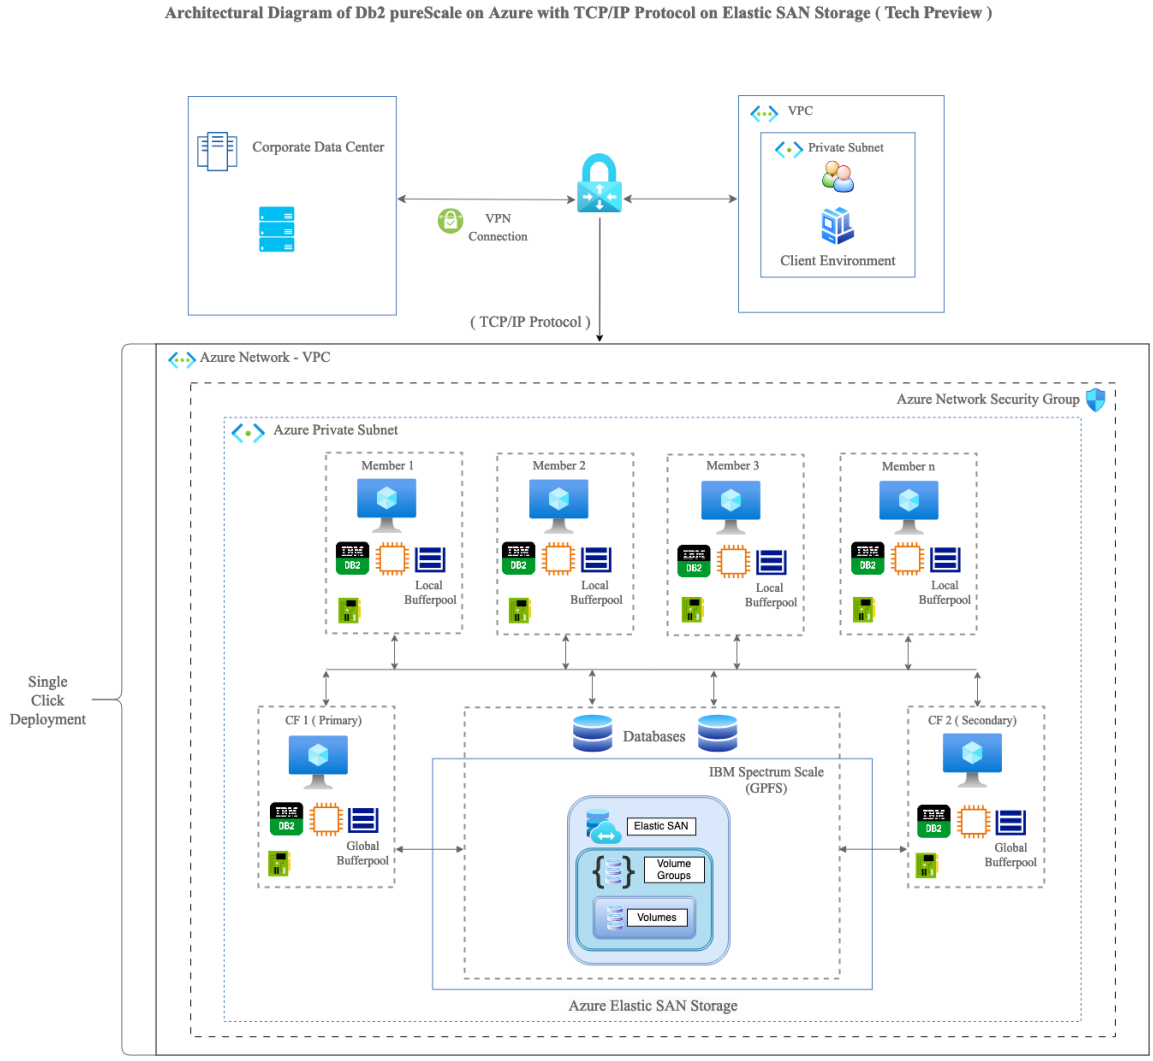

Figure 1 IBM Db2 pureScale architecture on Azure with TCP/IP network using Elastic SAN storage

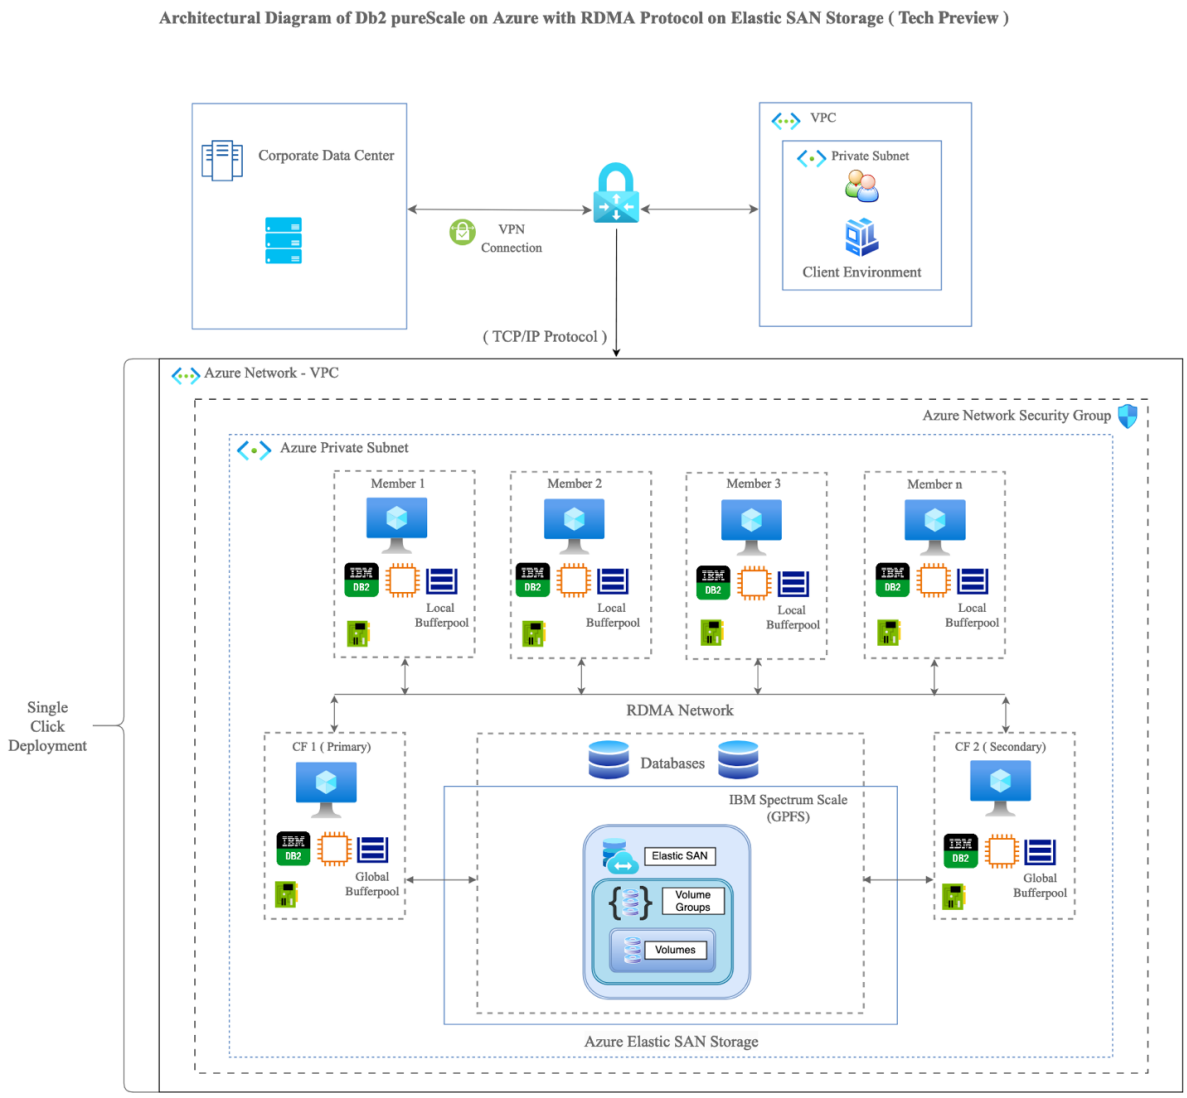

Figure 2 IBM Db2 pureScale architecture on Azure with RDMA network using Elastic SAN storage

Details in the above figure:

- A Virtual Network is configured with private subnets according to Azure best practices
- Virtual Machines as members or caching facilities, each with the following:
  - Elastic SAN storage
  - Network Interface private subnets
- Elastic SAN storage is attached to all Virtual Machines

Pre-Requisite

- An Azure account with **Contributor** role. If you don't already have an Azure account, create one at https://azure.microsoft.com/ by following the on-screen instructions. Part of the sign-up process involves receiving a phone call and entering a PIN using the phone keypad. Your Azure account is automatically signed up for all Azure services. You are charged only for the services you use. Ensure the Azure account is properly configured as required.
- SSH Key Pair – link to generate ssh key pair.

Deployment Procedure

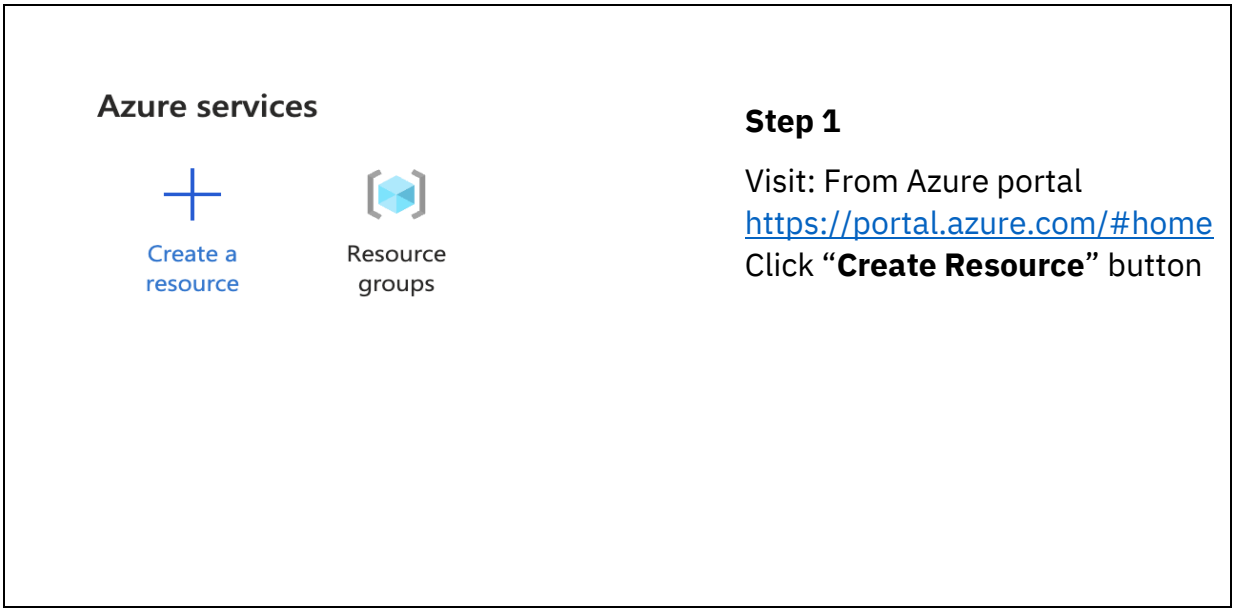

Step 1

Visit: From Azure portal https://portal.azure.com/#home
Click "**Create Resource**" button

Step 2

Now, click on "**See more in Marketplace**"

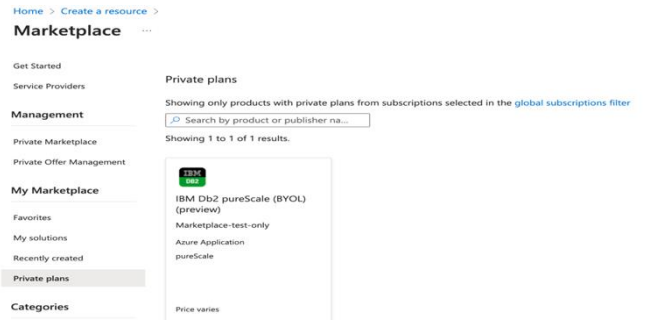

Step 3

Go to private plans and select "**IBM Db2 pureScale BYOL**".

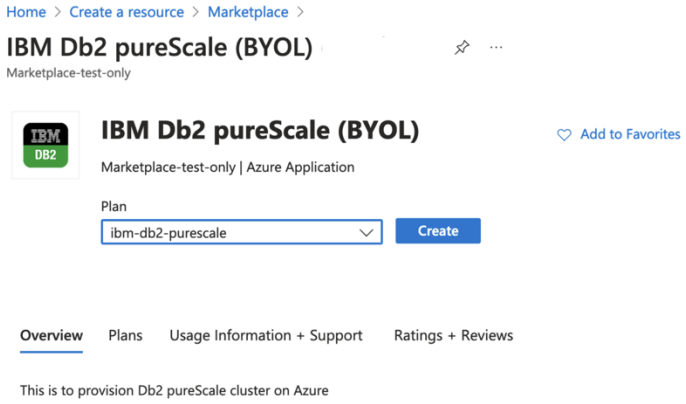

Step 4

Select the plan and then click on "**Create**".

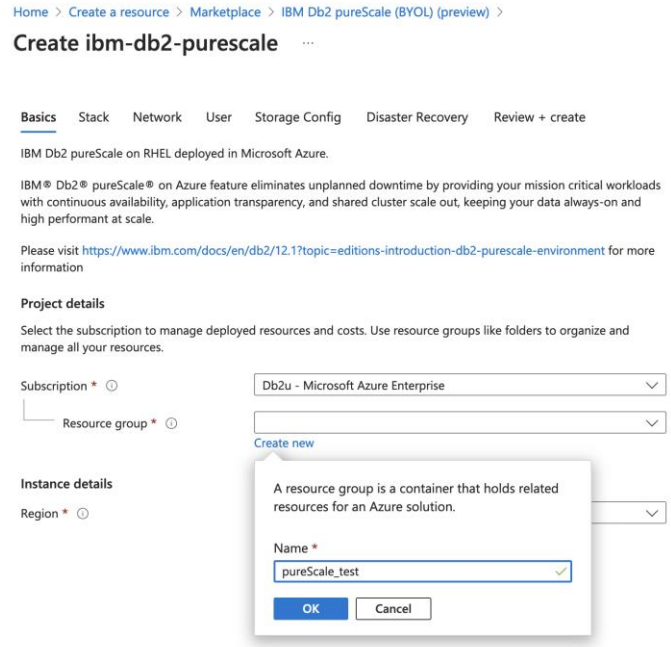

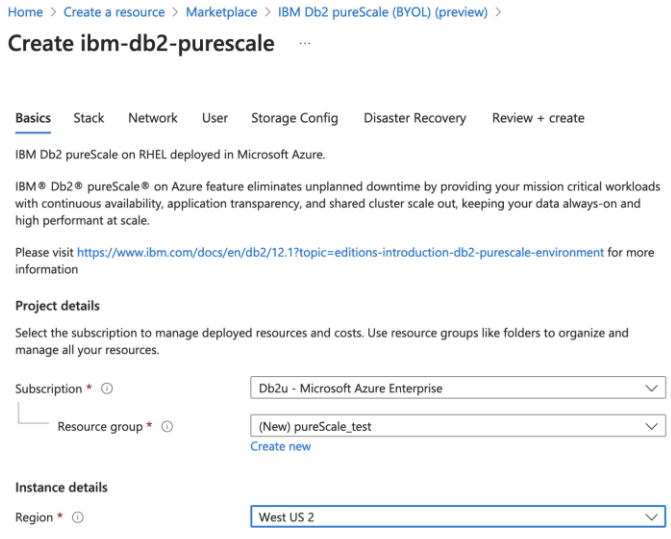

Step 5

Click on “**Create New**” for Resource group, enter name of choice and click **OK**.
After resource group, select **Region** to use for cluster creation. Then click **Next**.

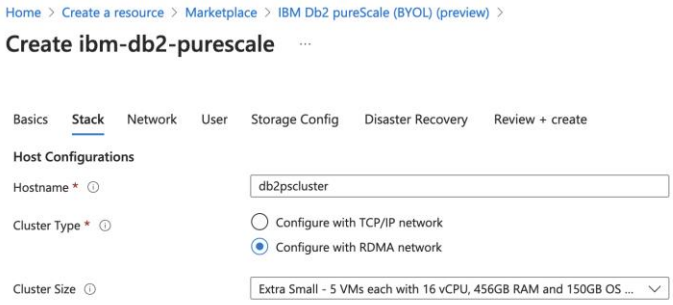

Step 6

Select cluster type as “**TCPIP**” or “**RDMA**” and Cluster Size per requirement among Extra Small/ Small/Medium/Large/Extra Large. Then click **Next**.

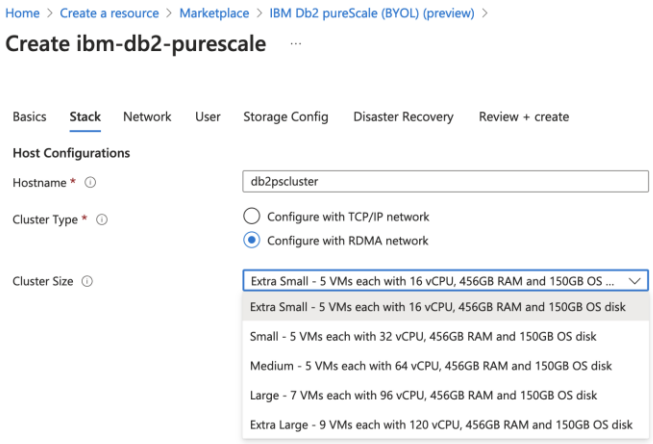

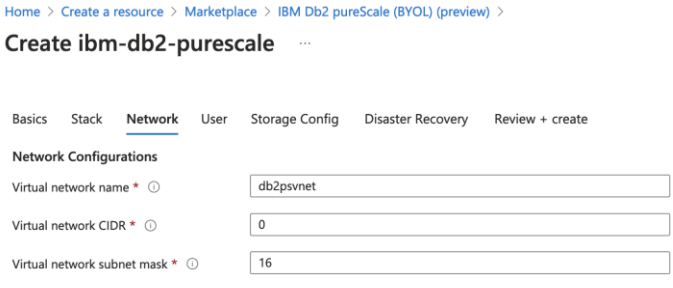

Step 7

Update the Virtual Network configurations. Then click **Next**.

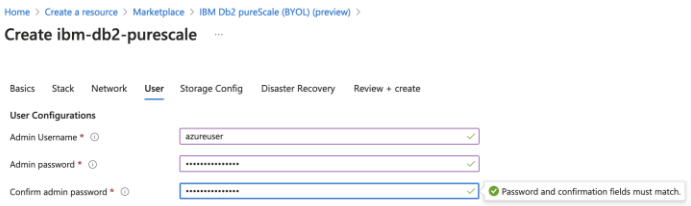

Step 8

Setup User-ID, Password to use for login. Then click **Next**.

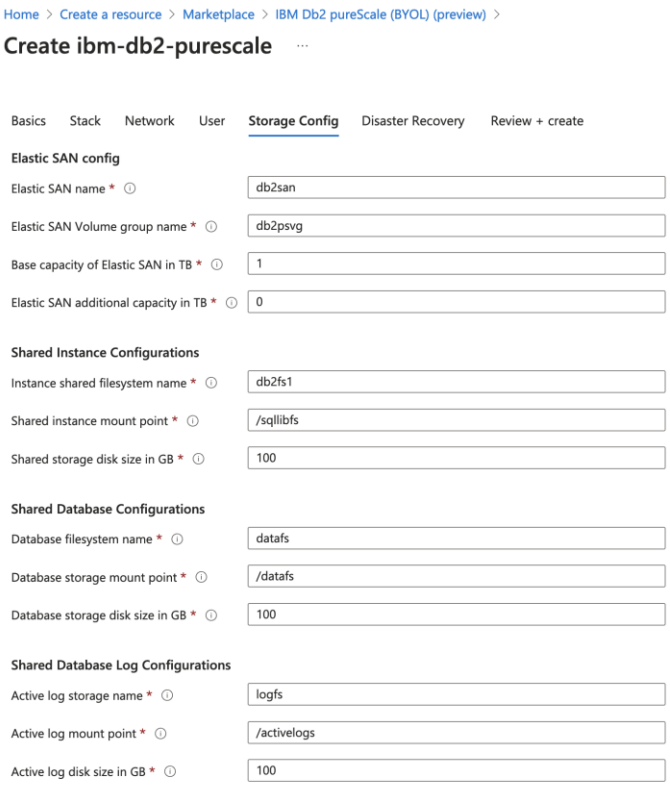

Step 9

On the Storage Configuration page, review the default settings and customize them as necessary. When you finish reviewing and customizing the parameters. Then click **Next**.

Parameter Label	Default	Description
HostName	-	Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Cluster Type		User can select any one from Configure with TCP/IP network or Configure with RDMA network
Cluster Size:	Extra Small	Choose the size of deployment - TCP: Extra Small (3 members, 2 cf, 16 vCPU, 32GB RAM, 12.5GbE), Small (3 Members, 2cf, 32 vCPU, 64GB RAM, 16GbE), Medium (3 Members, 2cf, 48 vCPU, 96GB RAM, 21GbE), Large (5 Members, 2cf, 64 vCPU, 128GB RAM, 28GbE) Extra Large (7 Members, 2cf, 72 vCPU, 144GB RAM, 30GbE) RDMA: Extra Small (3 members, 2 cf, 16 vCPU, 456GB RAM, 200GbE), Small (3 Members, 2cf, 32 vCPU, 456GB RAM, 200GbE), Medium (3 Members, 2cf, 64 vCPU, 456GB RAM, 200GbE), Large (5 Members, 2cf, 96 vCPU, 456GB RAM, 200GbE) Extra Large (7 Members, 2cf, 120 vCPU, 456GB RAM, 200GbE)

Virtual Network Name	db2psvnet	User can define a name of choice
Admin Username	-	*User can define any name of choice*
Admin Password	-	*Password to be set per password criteria*
Elastic SAN Name	db2san	*Any name of choice*
Elastic SAN Volume Group name	db2psvg	*Any name of choice*
Base capacity of Elastic SAN in TB	-	*Required capacity of the cluster in TB* *(example below for 6 TB)*
Elastic SAN additional capacity in TB	0	*Specify only if any additional capacity is required*
Instance shared filesystem name	db2fs1	Name defaults to db2fs1 only.
Instance shared mount point	/sqllibfs	Set the directory of the shared instance mount point.
Shared storage disk size in GB	200	Set the size of the shared storage disk for the cluster (in GB).
Database filesystem name	datafs	Set the name of the shared database file system.
Database storage mount point	/datafs	Set the name of the mount point for database shared storage.
Database storage disk size	1000	Set the size of the database shared storage for the cluster (in GB).
Active log storage name	logfs	Set the name of the filesystem for log shared storage.

Active log mount point	/logfs	Set the name of the mount point for log shared storage.
Active log storage size	500	Set the size of the shared storage disk for the cluster (in GB).
Archive log storage name	archlogfs	Set the name of the filesystem for archive log shared storage.
Archive log mount point	/archlogfs	Set the name of the mount point for archive log shared storage.
Archive log storage size	500	Set the size of the shared storage disk for the cluster (in GB).
Archive log storage name	backupfs	Set the name of the filesystem for archive log shared storage.
Archive log mount point	/backup	Set the name of the mount point for backup storage.
Archive log storage size	1000	Set the size of the shared storage disk for the cluster (in GB).

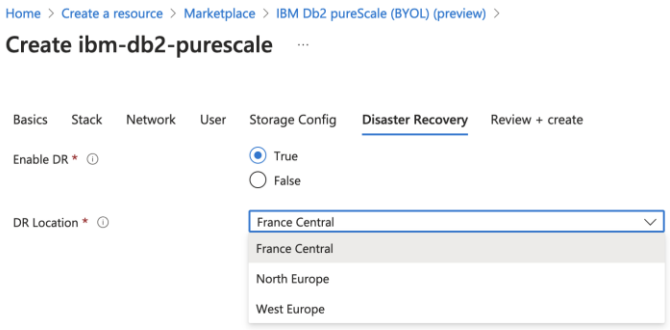

Step 10

On the Disaster Recovery page, please select "**Enable DR**" if you are looking for Disaster Environment. After that please select the "**Disaster Recovery Location**". Then click **Next**.

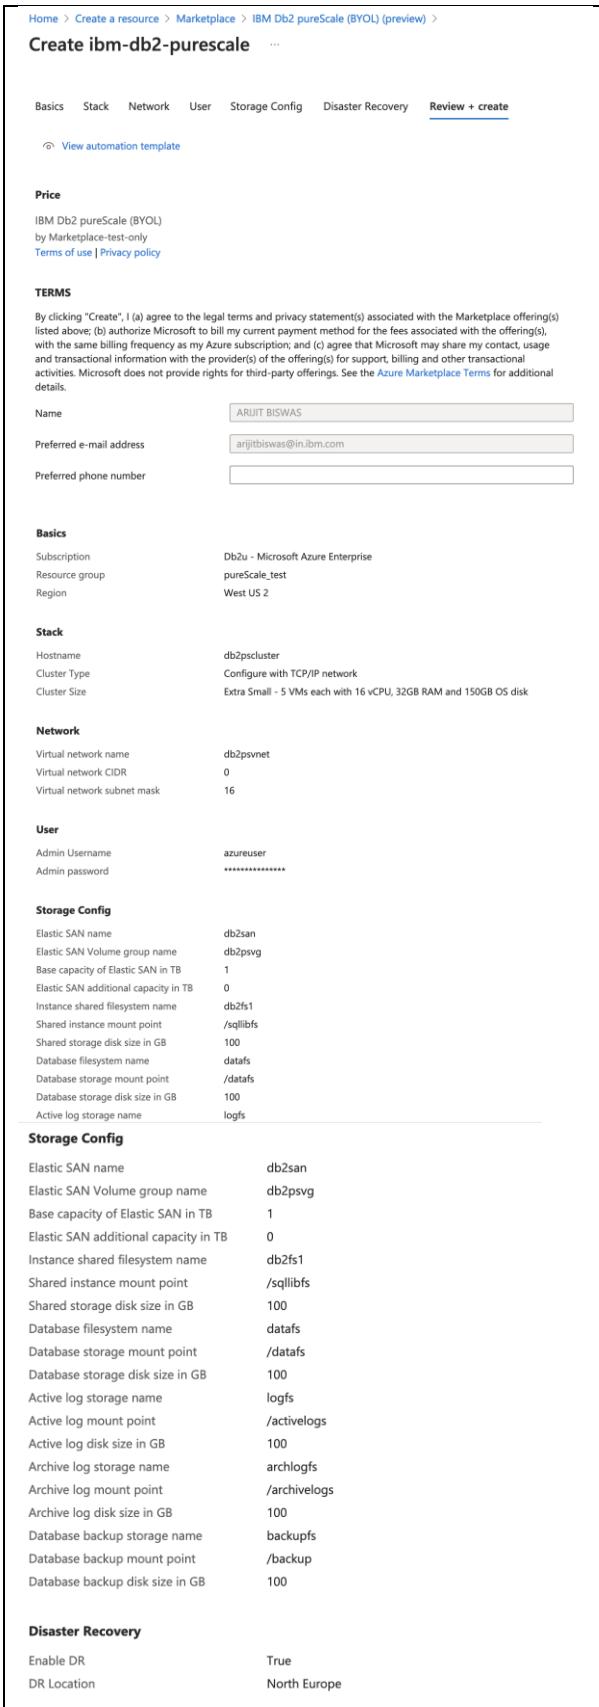

Step 11

On the "**Review + Create**" page, review and confirm the template then click on **Create.**

It will take approximately 30-35 min for cluster deployment to complete.

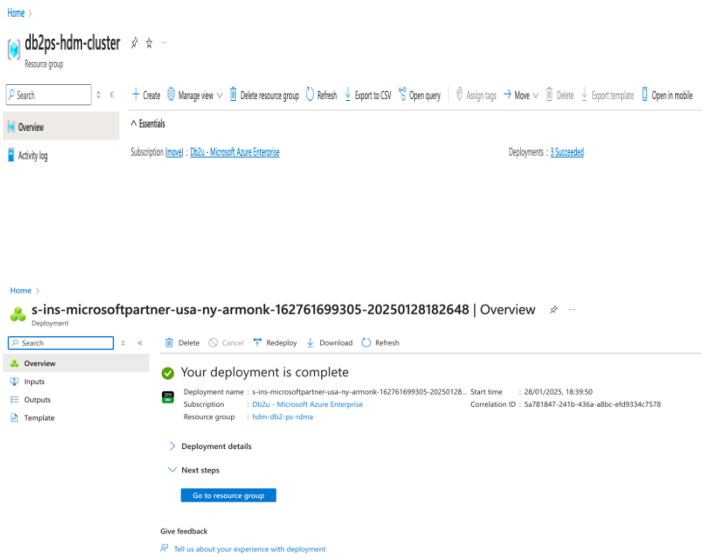

Step 12

From portal home page, go to created resource group and then you can monitor the in-progress or completed deployments.

Step 13

Approximately 30-35 minutes after the deployment has been completed, verify the cluster has successfully installed Db2 pureScale by connecting to the first (hostname-1) host in the cluster.

Step 14

Switch to the instance user and run "**db2instance -list**" and validate that all members have been added. It will show that the Db2 pureScale cluster is up.

Step 15

Update the Db2 instance user password for each host.

Set Up Client Connectivity – 3 Ways to connect

1. Site to Site connectivity

   https://learn.microsoft.com/en-us/azure/vpn-gateway/tutorial-create-gateway-portal

2. Azure Bastian service

   https://learn.microsoft.com/en-us/azure/bastion/bastion-overview

3. Creating Bastian VM

   https://learn.microsoft.com/en-us/azure/virtual-machines/linux/quick-create-portal?tabs=ubuntu

   1. From Azure portal, under created resource update bastion-nsg files with source and destination with Bastian public and private IP
   2. Make sure bastian .ssh/authorized_key is updated with your ssh pub key (provided during bastian VM cretion).

4. If you want to use Bastian VM as a Db2 client server

   Install Db2 Client in Bastian VM

Linux

Linux (64-bit)	Linux (32-bit)	Linux on Z	Linux on Power
IBM Db2 Universal Fixpack	--	IBM Db2 Universal Fixpack	IBM Db2 Universal Fixpack
IBM Db2 Server Fix Pack	--	IBM Db2 Server Fix Pack	IBM Db2 Server Fix Pack
IBM Data Server Client	IBM Data Server Client	IBM Data Server Client	IBM Data Server Client
IBM Data Server Runtime Client	IBM Data Server Runtime Client	IBM Data Server Runtime Client	IBM Data Server Runtime Client
IBM Data Server Driver	IBM Data Server Driver	IBM Data Server Driver	IBM Data Server Driver
IBM Data Server Driver for ODBC and CLI	IBM Data Server Driver for ODBC and CLI	IBM Data Server Driver for ODBC and CLI	IBM Data Server Driver for ODBC and CLI
IBM Data Server Driver for JDBC and SQLJ	--	--	--

   Once installation is completed, Update ssh pub key for Bastian server root and instance id in all clusters .ssh/authorized_key.

   For non-java clients, update db2dsdriver.cfg with cluster host details
   Java client can use JDBC/ODBC connection string.

Creating Additional Members for the Existing Db2 pureScale Cluster

Adding a new member to IBM® Db2® pureScale® on Azure can be done automatically from the Azure Marketplace. The new member will be added using the Azure Automation templates.

Adding members can be done online without having to bring down the entire Db2 pureScale cluster.

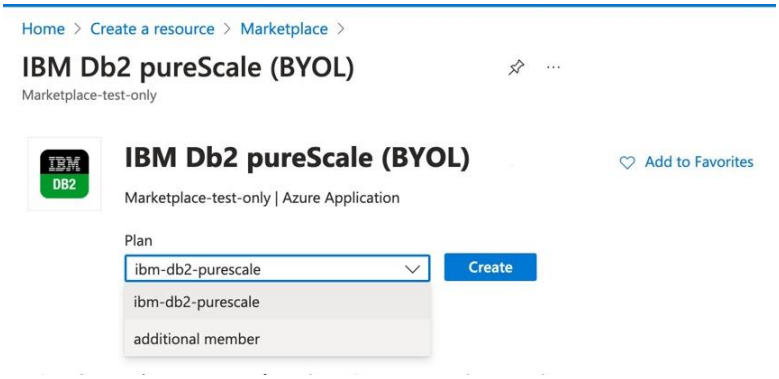

Step 1

From the Azure portal, Click “**Create Resource**” button. Then go to private plans and select “**IBM Db2 pureScale BYOL**”.

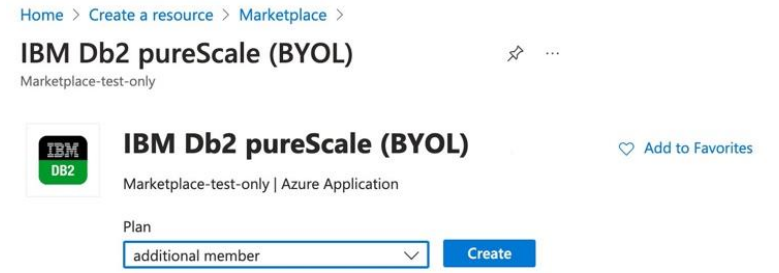

Step 2

Select the “**additional member**” plan and then click on “**Create**”.

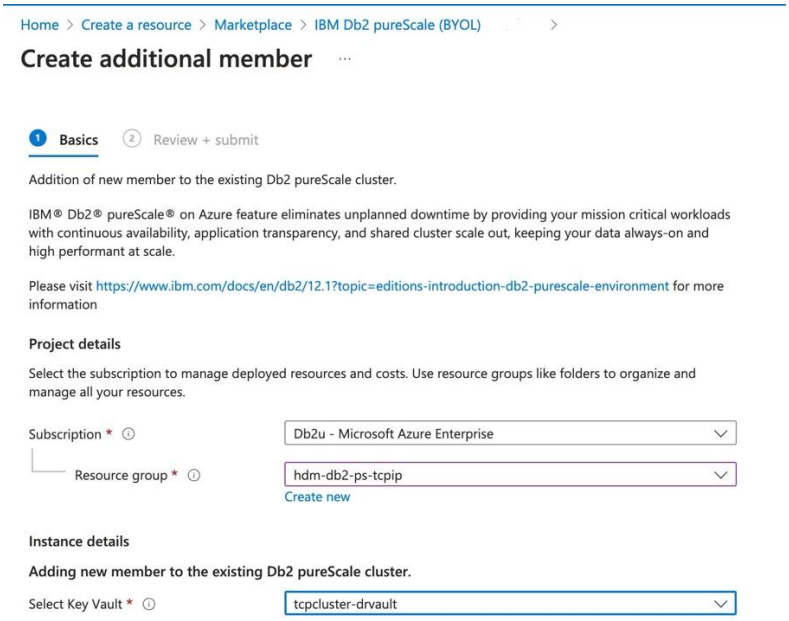

Step 3

Please mention the **same resource group**, under which the existing Db2 pureScale cluster resides. Also, select the instance details from the **Key Vault.**

Then click **Next**.

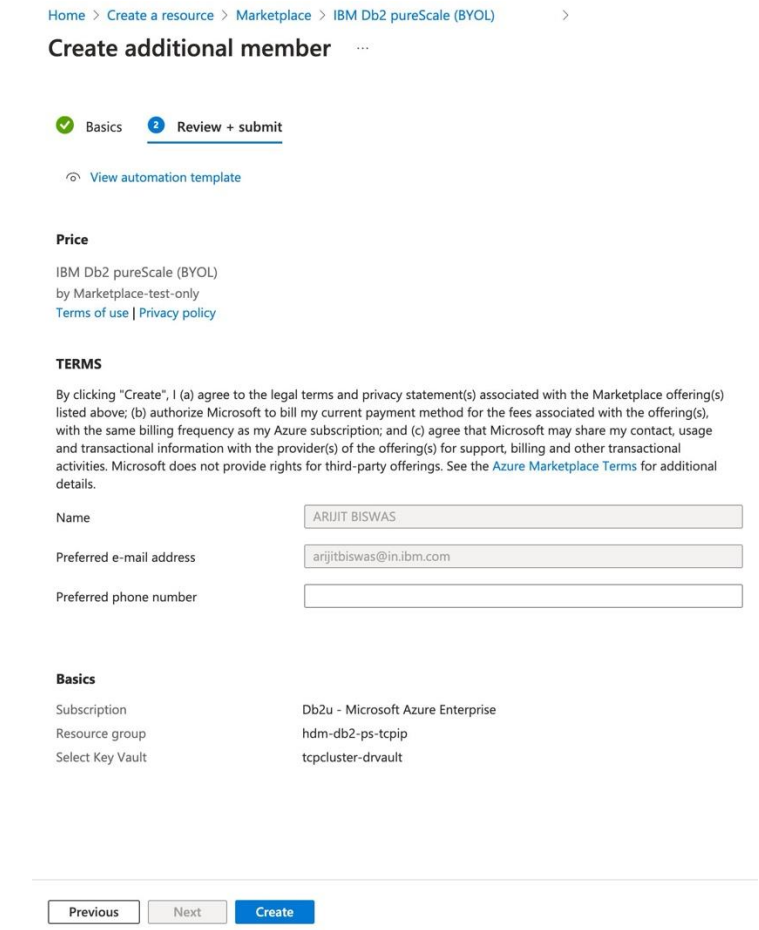

Step 4

On the "**Review + Submit**" page, review and confirm the template then click on **Create.**

It will take approximately 10-15 minutes to add the additional new member to the existing Db2 pureScale cluster.

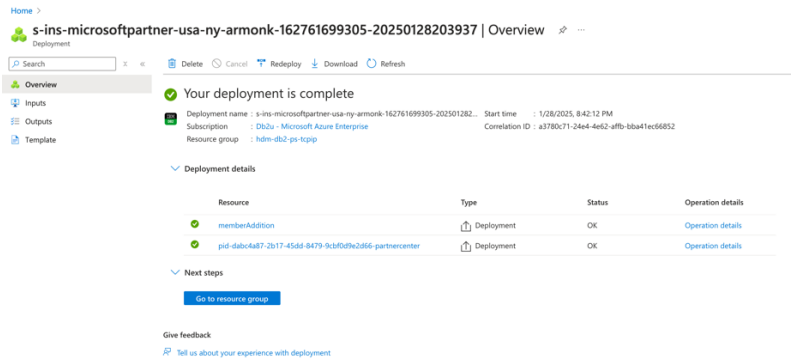

Step 5

Switch to the instance user and run "**db2instance -list**" and validate that the new member has been added to the cluster.

Technical Specifications

Region	France Central North Europe West Europe West US 2
OS	64-bit Linux
Host	FB v2 (TCPIP) and HB v2 (RDMA)
Storage	Elastic SAN
Database	Db2 v12.1.1.0 (Latest Db2 release)
Instances	Extra-Small, Small, Medium, Large, Extra-Large
License	The IBM® Db2 pureScale Feature on Azure employs Bring Your Own License (BYOL) model. New IBM Db2 license can be purchased from Passport Advantage

Support

For any queries or support requirement raise a case with details here