

The Lenovo logo is displayed in white text on a black rectangular background.

Introduction to Azure Arc for Data Services on Lenovo ThinkAgile MX Series

Last update: 1 December 2021
Version 1.1

Presents a SQL Server use
case for Lenovo ThinkAgile
MX Series

Describes Azure Arc for data
services for running SQL
server on ThinkAgile MX VMs

Describes the option to run
Azure SQL Managed Instance
SaaS on ThinkAgile MX with
Lenovo TruScale

Provides instructions for setting
up Azure SQL MI on Azure
Kubernetes Service on
ThinkAgile MX

Laurentiu Petre
Vinay Kulkarni



Table of Contents

1	Introduction	1
2	Business problem and business value	2
2.1	Business problem	2
2.2	Business value.....	2
3	Azure Arc for data services	3
3.1	Azure data services available currently.....	3
3.2	Azure SQL MI implementation on Lenovo ThinkAgile MX.....	3
3.3	Azure Arc Data Controller deployment.....	3
3.4	Creating the Azure SQL managed instance on Azure Arc.....	7
	Resources	9

1 Introduction

Customer IT environments these days have become very complicated with lots of applications running on diverse infrastructure spread over on-premises data centers, on the edge and probably in multiple clouds. There are probably different tools and frameworks in use and multiple technologies like DevOps and Kubernetes.

These IT resources are hard to view, manage and secure. A lot of innovation is happening in the cloud today. Microsoft Azure Arc is a set of technologies that allows customers to view and manage both their on-premises and cloud resources with a single pane of glass self-service portal. Azure Arc allows you to run innovative and advanced Azure services on your on-premises infrastructure. The Azure Arc for data services solution allows you to specifically run Azure SQL database and Postgres database services on-premises.

This white paper walks you through the steps required to set up and run Azure SQL database service on Lenovo ThinkAgile MX for Azure Stack HCI.

2 Business problem and business value

The following section provides a summary of the business problems that this white paper is intended to help address, and the value that the listed solutions can provide.

2.1 Business problem

As an organization you need to ensure you keep your data safe and secure and your workloads and applications keep running while incorporating the latest updates and innovations and the costs manageable. An organization is always looking for better management and policy enforcement on their infrastructure systems.

2.2 Business value

When you invest in a Lenovo ThinkAgile MX solution for Microsoft Azure Stack HCI to meet your hybrid cloud needs, there are several options available to keep your data safe and secure along with built-in capabilities in a multi-node high availability solution.

Azure Arc for data services from Microsoft provides you the ability to take advantage of Azure SQL innovation in the cloud by running the SQL Server Managed Instance (SQL MI) service using Azure Kubernetes Service (AKS) on Lenovo ThinkAgile MX for Azure Stack HCI.

Azure Arc for data services lets you run the popular Azure SQL MI service on premises in the data center or at the edge, and in multi-cloud environments on reliable ThinkAgile MX infrastructure solutions from Lenovo. You can benefit from the latest innovation in Azure, bursting feature with elastic scale, built-in high availability feature, and unified management for data workloads with or without a direct Azure connection.

3 Azure Arc for data services

3.1 Azure data services available currently

Two data services are available currently through Azure Arc. The general-purpose tier of Azure Arc-enabled SQL Managed Instance became generally available on July 30, 2021. Azure Arc-enabled PostgreSQL Hyperscale service is available in public preview. This paper focuses on Azure SQL management instance on Azure Kubernetes Service on Lenovo ThinkAgile MX for Azure Stack HCI.

3.2 Azure SQL MI implementation on Lenovo ThinkAgile MX

There are several approaches to having a working Azure SQL Managed Instance and we have focused on having it run on the local, on-premises infrastructure. This approach can comply with certain regulatory requirements and provide a cost competitive option.

The main prerequisites for having Azure SQL Managed Instance running are the following:

- A working Azure Stack HCI cluster
- Azure Kubernetes Services enabled on the Azure Stack HCI cluster
- A Kubernetes cluster with Azure Arc enabled deployed on Azure Kubernetes Services
- An Azure Data Controller deployed on the Kubernetes cluster

As the first three prerequisites were detailed in other previous documents, we`re going to focus on settings up the Azure Data Controller and enabling Azure SQL MI.

3.3 Azure Arc Data Controller deployment

From the Azure portal select *Kubernetes Services* and once this is selected the cluster should be visible:

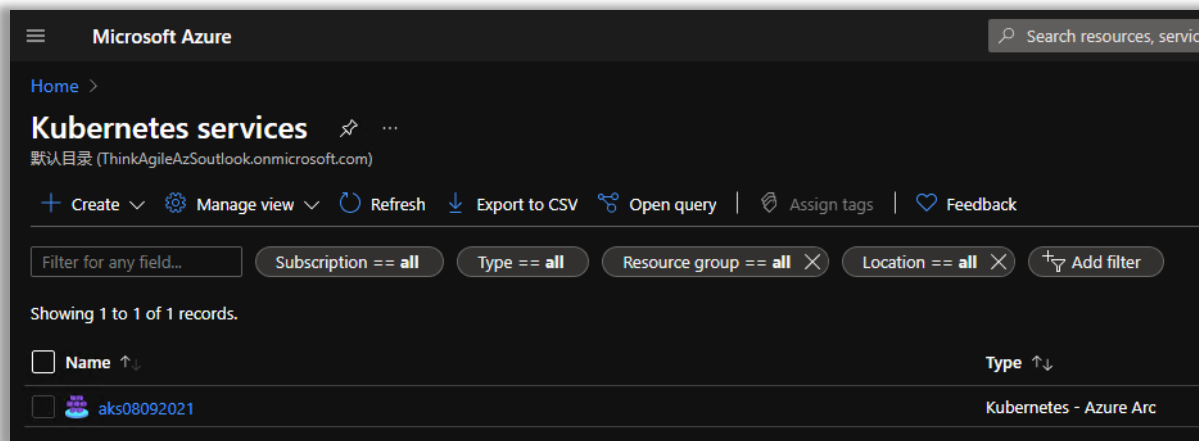


Figure 1 Kubernetes cluster

This should be the same cluster that is visible from Windows Admin Center, Azure Kubernetes Service section:



Figure 2 AKS in WAC

The next step would be to select the cluster from Azure portal and go to the *Extensions* section:

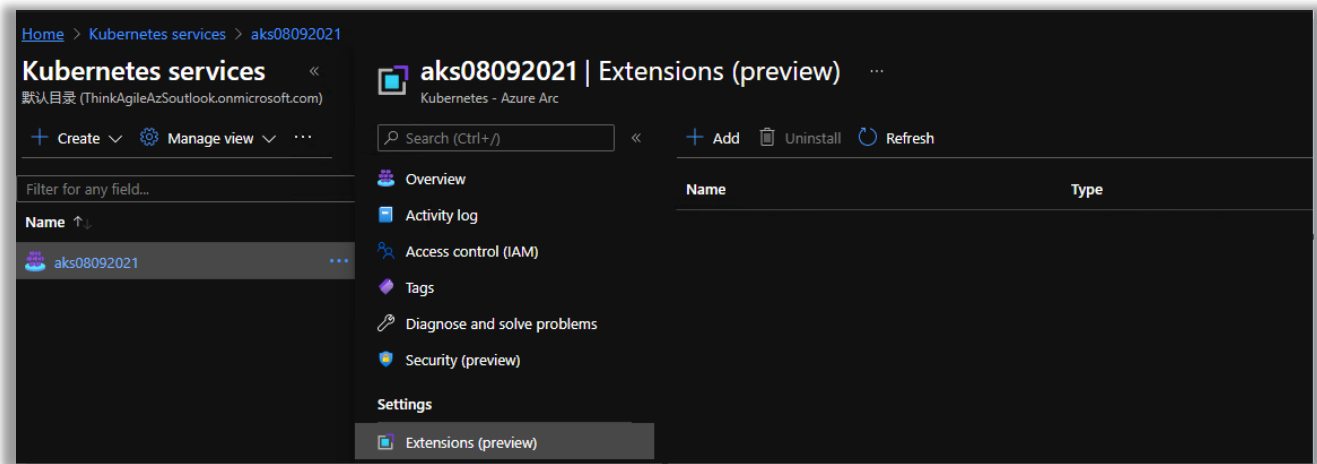


Figure 3 Extensions

Here, select *Add* and pick the *Azure Arc data controller* resource:

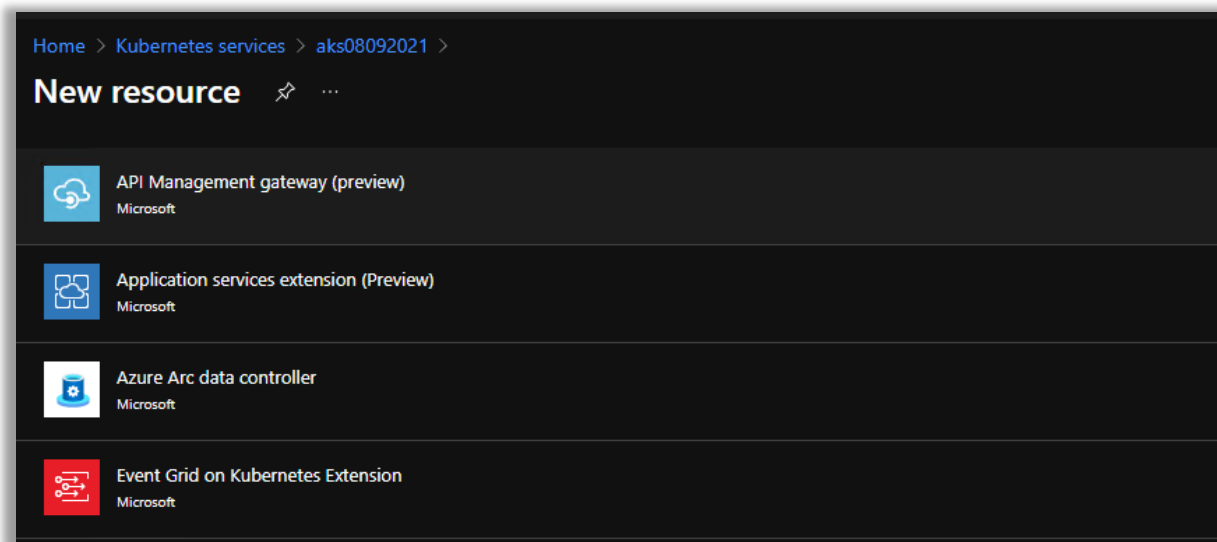


Figure 4 Azure Arc data controller resource

In the following information screen select *Create* and additional information regarding the Data Controller will be required:

Create an Azure Arc data controller to enable Azure data services in the Kubernetes environment of your choice.

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ Microsoft Azure Enterprise - Azure Stack Solutions Developmnet ✓

Resource group * ⓘ aks08092021 ✓

Data controller details

Provide a name to identify your data controller for remote management and monitoring.

i This data controller will be installed in direct connectivity mode to an existing Azure Arc-enabled Kubernetes cluster. This connectivity mode will allow you to create and manage Arc-enabled data services, such as Arc-enabled SQL Managed Instance and Arc-enabled PostgreSQL Hyperscale, directly from the Azure portal.

Data controller name * datacontroller2009 ✓

Custom location

A custom location is an Azure resource that represents the namespace on your Kubernetes cluster where the data controller will be hosted. [Learn more about custom locations](#)

Custom location * ⓘ

i Only custom locations for cluster 'aks08092021'

Kubernetes configuration

Select a template appropriate for your cluster

Kubernetes configuration template * ⓘ

Administrator account

Data controller login * ⓘ

Create new

Create new custom location

Name * customlocation2009 ✓

Namespace * default ✓

Create Cancel

Figure 5 Data controller settings

For the *Data controller name* field any name compliant with the requirements should be ok and on the *Custom location* field we chose the *default* namespace.

From the other options available, special care should be taken to select *azure-arc-aks-hci* on the *Kubernetes configuration template* and a service principal should be available to enter its data:

Kubernetes configuration

Select a template appropriate for your cluster configuration.

Kubernetes configuration template * ⓘ azure-arc-aks-hci ✓

Infrastructure onpremises ✓

Data storage class * ⓘ default ✓

Log storage class * ⓘ default ✓

Service type *

Node port: Exposes the service on each node's IP at a static port.

Load balancer: Exposes the service externally through a load balancer.

Administrator account

Data controller login * ⓘ administrator ✓

Password * ⓘ ✓

Confirm password * ✓

Upload Service Principal

Service Principal * ⓘ Manually enter service principal information ✓

Client Id * ⓘ 5c81952a-24fc-478d-97ff-7e6a42dbdf16 ✓

Client Secret * ⓘ

< Previous Next : Additional Settings >

Figure 6 Data Controller settings

Once all the data has been provided navigate thru the menus and start the deployment. Once it is done the status can be checked in the *Extensions* section:




Figure 7 Data Controller status

Also, going into *Azure Arc data controllers* we can check the status of the data controller:

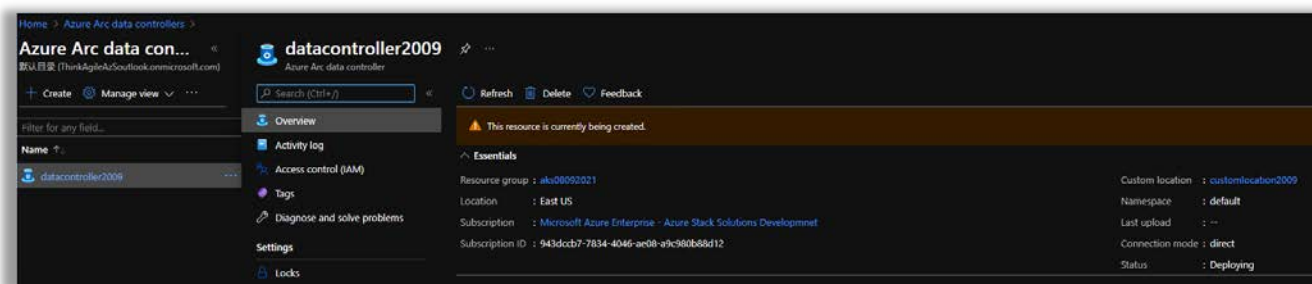


Figure 8 Data Controller resource

Once the deployment is complete the next step is to deploy the Azure SQL Managed instance.

3.4 Creating the Azure SQL managed instance on Azure Arc

As the SQL MI is going to be created over cli with Azure CLI it is required to have the following tools installed locally on the host that will be used for management:

- Azure CLI (the guide for the installation can be found on Microsoft's page: <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>)
- Arcdata Azure CLI extension (the instructions can be found here: <https://docs.microsoft.com/en-us/azure/azure-arc/data/install-arcdata-extension>)
- The following extensions for Azure CLI connectedk8s, k8s-extension, customlocation (the instructions for the installation can be found here: <https://docs.microsoft.com/en-us/azure/azure-arc/kubernetes/custom-locations#prerequisites>)

Once the prerequisites have been installed the installation can be initiated by running the following command in a powershell instance:

```
az sql mi-arc create -n sqldemo --k8s-namespace default --use-k8s
```

where sqldemo is the name that we chose and default is the namespace used.

3.6 Connecting to the Kubernetes cluster

One of the ways in which we can connect to the Kubernetes cluster is using the kubectl.exe utility available on Kubernetes page:

<https://kubernetes.io/docs/tasks/tools/install-kubectl-windows/>

Once the utility is downloaded the next step would be to download the configuration file for the Kubernetes cluster from WAC by selecting the "Download cluster kubeconfig":




Figure 9 cluster kubeconfig

After the configuration file and utility are saved, we need to configure kubectl to connect to the Kubernetes cluster. One of the ways to do this is to run the following powershell command making sure to replace the file path with your environment:

```
$ENV:KUBECONFIG=("C:\Users\user\Documents\kubectl\kubeconfig.xml")
```

Once this is done the connection can be tested by running

```
Kuberctl.exe get ns
```

And the output should look like this:

```
PS C:\Users\lpetre\Documents\kubectl> .\kubectl.exe get ns
NAME           STATUS    AGE
azure-arc      Active   3d23h
default        Active   3d23h
kube-node-lease Active   3d23h
kube-public    Active   3d23h
kube-system    Active   3d23h
```

Figure 10 kubernetes namespaces

Once this configuration is set the data controller deployment can be monitored by using several commands presented in the following article:

<https://docs.microsoft.com/en-us/azure/azure-arc/data/create-data-controller-direct-azure-portal#monitor-the-creation-from-your-kubernetes-cluster>

Change history

December 1, 2021 (version 1.1):

- Updated section 3.6 with additional steps

Resources

For more information about the topics that are described in this document, see the following resources:

- Lenovo ThinkAgile MX solution:
<https://www.lenovo.com/us/en/data-center/software-defined-infrastructure/ThinkAgile-MX-Certified-Node/p/WMD00000377>
- Microsoft Azure Stack HCI:
<https://azure.microsoft.com/en-us/products/azure-stack/hci/>
- Azure Arc for data services
<https://azure.microsoft.com/en-us/services/azure-arc/hybrid-data-services/#overview>

Trademarks and special notices

© Copyright Lenovo 2021.

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
ThinkAgile

The following terms are trademarks of other companies:

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

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.