```
{
 "cells": [
  {
   "cell_type": "markdown",
   "id": "d2ab2199",
   "metadata": {},
   "source": [
    "# Lab2: Cloud Computing Services and Setup\n",
    "\n",
    "Maryam R.Aliabadi, August 7th, 2023"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "16c7018f",
   "metadata": {},
   "source": [
    "## Learning objectives\n",
    "\n",
    "The goal of this lab is to launch and configure an Azure Virtual
Machine for use within your team. You will learn:\n",
    "\n",
    "- Setting up an Azure Virtual Machine\n",
    "- Logging into Azure Virtual Machine\n",
    "- Setting up a common space in Azure Virtual Machine\n",
    "- Setup your JupyterHub\n",
    "- Installing packages in TLJH"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "efc4550e",
   "metadata": {},
   "source": [
    "## Category of services in Azure\n",
    "\n",
    "Here are the main categories of services available in Azure:\n",
    "\n",
    "-
[Compute](https://learn.microsoft.com/en-us/azure/architecture/guide/techno
logy-choices/compute-decision-tree)\n",
    "-
[Storage](https://learn.microsoft.com/en-us/azure/storage/common/storage-in
troduction)\n",
    "-
[Database](https://azure.microsoft.com/en-us/products/category/databases/)\
```

```
n",
    "- [Networking and Content
Delivery](https://learn.microsoft.com/en-us/azure/networking/fundamentals/n
etworking-overview)  \n",
    "- [Security, Identity, and
Compliance](https://learn.microsoft.com/en-us/azure/security/fundamentals/i
dentity-management-best-practices) \n",
    "\n",
    "Having an understanding of different cloud services is crucial in
deciding the most suitable solution for your business. While it's
impossible to cover every category and service within Azure, we will be
focusing on the fundamental and commonly used services that serve as the
building blocks for most Azure services. Every year they add around 10 new
services, so again the most important skill is to adapt to new technologies
and tools.\n",
    "\n",
    "For a comprehensive list of services and their details, please refer
to [this](https://azure.microsoft.com/en-us/services/) link.\n",
    "\n",
    "## Compute (Azure Virtual Machines)\n",
    "A web service that provides secure, resizable compute capacity in the
cloud. It is designed to make ***web-scale computing*** easier for
developers.\n",
    "\n",
    "### The littlest JupyterHub (TLJH)\n",
    "\n",
    "We are all using JupyterLab on our personal computers, but when we
move to the industry, things change, and you will mostly use JupyterHub,
which is a hub of Jupyter where employees log in with their username and
password. We have a similar setup here at UBC known as
https://ubc.syzygy.ca/. In our example, we will set up a JupyterHub so that
your team members can use it to log in and create a collaborative
environment. This entire JupyterHub setup is made easy by
[TLJH](https://tljh.jupyter.org/en/latest/install/azure.html), which we
will be using. TLJH...\n",
    "\n",
    "- Greatly simplifies the deployment process.\n",
    "- Enables a much faster setup/teardown.\n",
    "- Provides the same ability to customize authentication, etc.\n",
    "- Defines the user environment by installing packages as admin."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "b5e8abd1",
   "metadata": {},
```

```
   "source": [
    "### Setting up an Azure Virtual Machine\n",
    "\n",
    "You can use the following resources as a reference for setting up an Azure Virtual Machine:\n",
    "\n",
    "- [Quickstart: Create a Windows virtual machine in the Azure portal](https://learn.microsoft.com/en-us/azure/virtual-machines/windows/quick-create-portal)\n",
    "- [Create a Windows virtual machine in Azure - Training](https://learn.microsoft.com/en-us/training/modules/create-windows-virtual-machine-in-azure/)\n",
    "- [9 Steps to Set Up Azure Virtual Machine - WaferWire](https://www.waferwire.com/blog/how-to-set-up-azure-cloud-environment-virtual-machine/)\n",
    "- [Create Free Windows Virtual Machine in Azure - GeeksforGeeks](https://www.geeksforgeeks.org/create-free-windows-virtual-machine-in-azure/)\n",
    "\n",
    "These resources provide step-by-step instructions on how to create and deploy a Windows Virtual Machine in Microsoft Azure. They cover everything from logging into your Microsoft Azure account, to selecting the appropriate services, creating a new Virtual Machine, and connecting to it."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "f82d19b6",
   "metadata": {},
   "source": [
    "1. **Project details**: Choose your subscription and resource group. If you don't have a resource group, create one. Then, choose a region near you or near other services you may need to access. For example, if you want to access a database in the same region, choose the same region for your virtual machine.\n",
    "2. **Instance details**: \n",
    "   - **Virtual machine name**: Give your virtual machine a name.\n",
    "   - **Security type**: Choose \"Standard\".\n",
    "   - **Application and OS Images**: Choose an Azure Machine Image (AMI). Ubuntu Server 20.04 LTS is a good choice.\n",
    "   - **Architecture**: Keep the default architecture (x64)\n",
    "   - **Instance type**: Choose an Instance Size. Click on see all sizes and choose a size that fits your needs. For example, Standard B1s is a good choice for a small project (7.59$/mo). Keep in mind about the expenses of each machine, you have limited resources and it's not
```

```
recommended to use a large machine for a small project.\n",
    "3. **Administrator account**: \n",
    "   - **Authentication type**: Use authentication type as SSH.\n",
    "   - **Username**: Enter a username that will be used to login in Jupyter Notebook\n",
    "   - **Key pair name**: Enter a name for your key pair.\n",
    "4. **Inbound port rules**: Select the ports you want to open. For example, open port 22 for SSH, 80 for HTTP, 443 for HTTPS. You can always add more ports later.\n",
    "\n",
    "5. **Advanced Details**: For the advanced details, get the configuration code in the [link](https://tljh.jupyter.org/en/latest/install/azure.html) and replace \"admin-user-name\" (remove < > as well) with your Jupyter username.\n",
    "\n",
    "After filling in these details, you can proceed to the other tabs (Disks, Networking, Management, Monitoring, Advanced, Tags) for further customization of your virtual machine. Once you've reviewed everything, click on \"Review + create\" to provision your virtual machine.\n",
    "\n",
    "Below is the configuration code I have used for my virtual machine"
   ]
  },
  {
   "cell_type": "raw",
   "id": "96b66e73",
   "metadata": {},
   "source": [
    "#!/bin/bash\n",
    "curl -L https://tljh.jupyter.org/bootstrap.py \\\n",
    "  | sudo python3 - \\\n",
    "    --admin maryam"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "4a1e5b1e",
   "metadata": {},
   "source": [
    "6. Search for your virtual machine under virtual machines to see if it’s running. Give it 15 - 20 minutes as it takes time to set up JupyterHub even if it shows running. You can also check system logs during this time."
   ]
  },
  {
   "cell_type": "markdown",
```

```
   "id": "f980b8bd",
   "metadata": {},
   "source": [
    "### Azure extension FTW\n",
    "\n",
    "You can also login into Azure virtual machine right from your vs code terminal all thanks to the azure extension.\n",
    "\n",
    "1. **Install the Azure extension**: If you haven't already, install the Azure extension for Visual Studio Code. This will allow you to interact with Azure directly from your IDE.\n",
    "\n",
    "2. **Open Visual Studio Code**: Launch Visual Studio Code on your local machine.\n",
    "\n",
    "3. **Connect to Azure**: In Visual Studio Code, click on the Azure icon in the Activity Bar on the side. This will open the Azure extension. If you're not already logged in, you'll be prompted to log in to your Azure account.\n",
    "\n",
    "4. **Find your Virtual Machine**: After logging in, click on the refresh button in the Azure extension. Under \"Virtual Machines\", you should see your newly created virtual machine.\n",
    "\n",
    "\n",
    "If you have not generated your own system's key-value pair, you can follow these steps to generate one:\n",
    "\n",
    "1. Open your terminal.\n",
    "2. Navigate to the directory where you want to save your key pair.\n",
    "3. Type in the following command to generate a key pair: `ssh-keygen`\n",
    "\n",
    "This will start the key generation process as follows"
   ]
  },
  {
   "cell_type": "raw",
   "id": "05e0c068",
   "metadata": {},
   "source": [
    "C:\\Users\\barya>ssh-keygen\n",
    "Generating public/private rsa key pair.\n",
    "Enter file in which to save the key (C:\\Users\\barya/.ssh/id_rsa): demo\n",
    "Enter passphrase (empty for no passphrase):\n",
```

```
    "Enter same passphrase again:\n",
    "Your identification has been saved in demo\n",
    "Your public key has been saved in demo.pub\n",
    "The key fingerprint is:\n",
    "SHA256:F7TnDrlxXcCD/b6KiCpDGdLemB91mwXXQy1L2PjtUMc xyz@Aryan\n",
    "The key's randomart image is:\n",
    "+---[RSA 3072]----+\n",
    "|          .*.. . |\n",
    "|        ..+.* o E|\n",
    "| .      .ooo.* . |\n",
    "|. o   ...+.*+ .  |\n",
    "| o * . .S+= oo   |\n",
    "|  * o  .o= + ..  |\n",
    "| . . .  o . . .  |\n",
    "|  o .  . . . .   |\n",
    "|   o... . . ..o. |\n",
    "+----[SHA256]-----+"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "b28aa1f8",
   "metadata": {},
   "source": [
    "After you have generated your key pair, you can add it to your Azure Virtual Machine for secure SSH access. Here's how:\n",
    "\n",
    "1. In Visual Studio Code, right-click on your virtual machine in the Azure extension.\n",
    "2. Click on \"Add SSH Key\".\n",
    "3. Navigate to the .ssh folder in your home directory and select your public key to upload.\n",
    "\n",
    "Once you've done this, you should be able to securely connect to your Azure Virtual Machine using SSH with the key pair you generated."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "212d9c34",
   "metadata": {},
   "source": [
    "### Interacting with the Azure Virtual Machine \n",
    "\n",
    "Now we have our Azure Virtual Machine ready let's see some things that we can do; You can Stop, Terminate, Start the instance and Monitor System
```

```
logs. You can also create an image of your instance and scale it up if you need more processing power.\n",
    "Assuming someone from your team has set up this Azure Virtual Machine, there are several questions we need to consider. In the following three sections, you will find answers to these questions:\n",
    "\n",
    "- How can you and your team members log in to the Azure Virtual Machine?\n",
    "- How can you make directories accessible to your team members\n",
    "- How can team members log in to the JupyterHub?\n",
    "\n",
    "In the following, you will find the answer to the above questions:"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "8be4b904",
   "metadata": {},
   "source": [
    "### Logging into the Azure VM instance \n",
    "\n",
    "There are 3 ways to connect to the Azure Virtual Machine:\n",
    "1. **Azure portal**: You can connect to the Azure Virtual Machine through the Azure portal. This is the easiest way to connect to the Azure Virtual Machine. You can connect to the Azure Virtual Machine through the Azure portal by clicking on the \"Connect\" button on the overview page of your virtual machine. Select the ***`SSH using Azure CLI`*** tab and follow the instructions there. \n",
    "\n",
    "2. **SSH**: You can connect to the Azure Virtual Machine through SSH. This is the most common way to connect to the Azure Virtual Machine. \n",
    "Remember the private .pem key you downloaded when you created your Azure Virtual Machine? You will need this to connect to the Azure Virtual Machine through SSH.\n",
    "\n",
    "SSH by using the following command in your terminal:\n",
    "\n",
    "   `ssh -i <path-to-private-key> <username>@<public-ip-address>`\n",
    "\n",
    "   For example, if your username is \"ubuntu\", your public IP address is \"13.82.123.456\", and your private key is located in the .ssh folder in your home directory, you would use the following command:\n",
    "\n",
    "   `ssh -i ~/Downloads/mytestkey.pem ubuntu@13.82.123.456`\n",
    "\n",
    "3. **Azure extension**: You can connect to the Azure Virtual Machine
```

through the Azure extension in Visual Studio Code. This is the most convenient way to connect to the Azure Virtual Machine. You can connect to the Azure Virtual Machine through the Azure extension by right-clicking on your virtual machine in the Azure extension and clicking on \"Connect with SSH\". The setup was done above. This method involves adding your computer's public key to the Azure Virtual Machine. This is a one-time process and you will not need to do it again."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "e33deaa4",
   "metadata": {},
   "source": [
    "The person who sets up the Azure Virtual Machine (maryam in this case) will only have access to the server, as that person will have the private key. Actually, Maryam has the private key of the root user `ubuntu` (the one I created when setting up Azure Virtual Machine and downloaded at last stage) and is logging into the Azure Virtual Machine as `ubuntu`. If you have not yet logged into the Azure Virtual Machine, please check out the last step in the Azure Virtual Machine creation process. The following is how you logged into the Azure Virtual Machine:"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "8957aa89",
   "metadata": {},
   "source": [
    "As you can see, I logged into the instance as ubuntu. If others need to access the instance, a new user account must be created, and a keypair should be generated for them to use.\n",
    "\n",
    "For example, imagine a team with 3 members: Maryam, Maria, and Nicole. Maryam creates an Azure Virtual Machine and therefore has exclusive access to it, using the private key. But how can Maria and Nicole access it? It is not advisable for Maryam to share her private key with others. Instead, Maryam needs to create individual user accounts for Maria and Nicole and register their public keys into the system.\n",
    "\n",
    "Here is the conversation that happens;\n",
    "\n",
    "Maryam: Hey Maria - I have created an Azure Virtual Machine for us to use. I want to give you access to it. Can you provide me your public key?\n",
    "Maria: Sure, I don't have anything at the moment but I will create a

```
keypair and provide you with the public key. \n",
    "Maryam: Great, Once I have it, I will add your public key to the
server so you can access it."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "7c7626f9",
   "metadata": {},
   "source": [
    "### Adding a new user to your Azure Virtual Machine\n",
    "\n",
    "1. Connect to your Azure Virtual Machine using SSH.\n",
    "\n",
    "2. Use the adduser command to add a new user account to an Azure
Virtual Machine (replace new_user with the new account name). The following
example creates an associated group, home directory, and an entry in the
/etc/passwd file of the instance."
   ]
  },
  {
   "cell_type": "raw",
   "id": "a176d389",
   "metadata": {},
   "source": [
    "\n",
    "$ sudo adduser new_user\n"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "044dd092",
   "metadata": {},
   "source": [
    "The home directory might not be created by default in some
configurations. Verify that the home directory was created before
continuing."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "ddf2cd6e",
   "metadata": {},
   "source": [
    "Note: If you add the new_user to an Ubuntu instance, then include the
--disabled-password option to avoid adding a password to the new account:"
```

```
   ]
  },
  {
   "cell_type": "raw",
   "id": "0af6bafe",
   "metadata": {},
   "source": [
    "$ sudo adduser new_user --disabled-password"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "37f59d24",
   "metadata": {},
   "source": [
    "3.    Change the security context to the new_user account so that
folders and files you create have the correct permissions:"
   ]
  },
  {
   "cell_type": "code",
   "execution_count": null,
   "id": "117c102c",
   "metadata": {
    "vscode": {
     "languageId": "plaintext"
    }
   },
   "outputs": [],
   "source": [
    "$ sudo su - new_user"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "e25e04eb",
   "metadata": {},
   "source": [
    "Note: When you run the sudo su - new_user command, the name at the top
of the command shell prompt changes to reflect the new user account context
of your shell session."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "b692cc58",
```

```
   "metadata": {},
   "source": [
    "\n",
    "4.    Create a .ssh directory in the new_user home directory:"
   ]
  },
  {
   "cell_type": "raw",
   "id": "bf1badcd",
   "metadata": {},
   "source": [
    "$ mkdir .ssh"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "b3c31f9a",
   "metadata": {},
   "source": [
    "5.    Use the chmod command to change the .ssh directory's permissions
to 700. Changing the permissions restricts access so that only the new_user
can read, write, or open the .ssh directory"
   ]
  },
  {
   "cell_type": "raw",
   "id": "e0eacb84",
   "metadata": {},
   "source": [
    "$ chmod 700 .ssh"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "58f09701",
   "metadata": {},
   "source": [
    "6.    Use the touch command to create the authorized_keys file in the
.ssh directory:"
   ]
  },
  {
   "cell_type": "raw",
   "id": "4df112ae",
   "metadata": {},
   "source": [
```

```
    "$ touch .ssh/authorized_keys"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "0241e6c5",
   "metadata": {},
   "source": [
    "7.    Use the chmod command to change the .ssh/authorized_keys file permissions to 600. Changing the file permissions restricts read or write access to the new_user."
   ]
  },
  {
   "cell_type": "raw",
   "id": "3b98041f",
   "metadata": {},
   "source": [
    "$ chmod 600 .ssh/authorized_keys"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "14006eea",
   "metadata": {},
   "source": [
    "\n",
    "After the successful completion of the process, nicole and maria will be able to login like below:\n"
   ]
  },
  {
   "cell_type": "raw",
   "id": "acb9e551",
   "metadata": {},
   "source": [
    "# example nicole\n",
    "ssh -i <path_to_nicole_private_key> nicole@13.82.123.456\n",
    "# example maria\n",
    "ssh -i <path_to_maria_private_key> maria@13.82.123.456"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "a63e43ad",
   "metadata": {},
```

```
   "source": [
    "### Setting up a common space in Azure Virtual Machine\n",
    "After creating accounts for your team members, you'll notice folders for each user in the \"/home\" directory. However, the person who set up the Azure Virtual Machine (in my case maryam) doesn't have a user account. Instead, Maryam is using the ubuntu account (and has access to the ubuntu private key). Also, here ubuntu is the `root user`."
   ]
  },
  {
   "cell_type": "raw",
   "id": "27efbe33",
   "metadata": {},
   "source": [
    "maria@ip-172-31-21-101:~$ cd /home\n",
    "maria@ip-172-31-21-101:/home$ ls -ltrh\n",
    "total 12K\n",
    "drwxr-xr-x 2 nicole nicole 4.0K Mar 21 07:32 nicole\n",
    "drwxr-xr-x 5 ubuntu ubuntu 4.0K Mar 21 07:39 ubuntu\n",
    "drwxr-xr-x 5 maria  maria  4.0K Mar 21 07:39 maria\n"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "c40419d0",
   "metadata": {},
   "source": [
    "Setup a common data folder to download data, and this folder should be accessible by all users in the JupyterHub. Following commands, make a folder and make it accessible to everyone. \n"
   ]
  },
  {
   "cell_type": "code",
   "execution_count": null,
   "id": "04eab0b3",
   "metadata": {
    "vscode": {
     "languageId": "plaintext"
    }
   },
   "outputs": [],
   "source": [
    "sudo mkdir -p /srv/data/my_shared_data_folder\n",
    "sudo chmod 777 /srv/data/my_shared_data_folder/"
   ]
```

```
  },
  {
   "cell_type": "markdown",
   "id": "a16d70f4",
   "metadata": {},
   "source": [
    "#### Setup your JupyterHub on Azure Virtual Machine\n",
    "\n",
    "We installed our TLJH while setting up the Azure Virtual Machine. Currently, we only have one admin user - the person who set up the Azure Virtual Machine - and the admin username is what you provided in your `Azure user data` (in my case maryam). We will use that username to log in to JupyterHub. \n",
    "\n",
    "1) Under description, check for \"Public IP address\" and paste the IP address in your browser for your JupyterHub. (make sure it is not automatically prepended with `https://`, we want to access it over `http://`)\n",
    "\n",
    "2) Enter your username (what you provided in your `Azure user data`) and use a strong password & remember it, as what you enter here will be the admin password.\n"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "e18c9daa",
   "metadata": {},
   "source": [
    "We want to add other team members as admin users to the jupyterhub. [Here](https://tljh.jupyter.org/en/latest/howto/admin/admin-users.html?highlight=adding%20user#adding-admin-users-from-the-jupyterhub-interface) are the instructions that you can use. I strongly recommend you check it out, but I will list down the main points."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "997d4a25",
   "metadata": {},
   "source": [
    "In your JupyterHub, go to \"Control Panel\" --> \"admin.\" Here add other members of your group, use their names and make them admins.\n",
    "\n",
    "3) Check if other members can log in to the JupyterHub from their machines by giving them the URL to connect. Step 2 is applicable here for
```

```
other members."
   ]
  },
  {
   "cell_type": "markdown",
   "id": "de3228ee",
   "metadata": {},
   "source": [
    "\n",
    "##### Install packages in TLJH \n",
    "\n",
    "You have to install the packages that are needed. You ***should*** 
refer this TLJH [document](
https://tljh.jupyter.org/en/latest/howto/env/user-environment.html).\n",
    "\n",
    "Don't forget to add option -E. This way, all packages that you install 
will be available to other users in your JupyterHub. For example, open up 
the terminal from the jupyter, and install packages that you need;"
   ]
  },
  {
   "cell_type": "raw",
   "id": "043abfa4",
   "metadata": {},
   "source": [
    "sudo -E pip install pandas\n",
    "sudo -E pip install s3fs\n",
    "sudo -E pip install pyarrow"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "4a34e81f",
   "metadata": {},
   "source": [
    "You surely want to install the packages mentioned above. Also, install 
any other packages that are required for your data wrangling. Please note 
that the setup and server we have are too small for a fully-fledged 
collaborative space. Therefore, please only install what is necessary, or 
it may lead to a system crash.\n"
   ]
  },
  {
   "cell_type": "markdown",
   "id": "1fe9f1b2",
   "metadata": {},
```

```
   "source": [
    "### How to create an image in Azure?\n",
    "\n",
    "Creating an image of your Azure Virtual Machine is similar to creating an image in EC2. An image in Azure is a copy of your virtual machine (VM) that you can use to create new VMs with the same configuration.\n",
    "\n",
    "Here are the steps to create an image of your VM in Azure:\n",
    "\n",
    "1. In the Azure portal, select your VM.\n",
    "2. In the left-hand menu, under 'Settings', select 'Capture'.\n",
    "3. Provide a name for the image and select a storage account where you want to store the image.\n",
    "4. Click on 'Create'.\n",
    "\n",
    "This will create an image of your VM with all its configurations and installed packages. You can use this image to create new VMs with the same setup.\n",
    "\n",
    "Please note that creating an image will generalize the VM, which deallocates and marks it as generalized in the Azure portal. After a VM has been generalized, it cannot be started or redeployed until it has been re-provisioned."
   ]
  }
 ],
 "metadata": {
  "kernelspec": {
   "display_name": "Python 3 (ipykernel)",
   "language": "python",
   "name": "python3"
  },
  "language_info": {
   "codemirror_mode": {
    "name": "ipython",
    "version": 3
   },
   "file_extension": ".py",
   "mimetype": "text/x-python",
   "name": "python",
   "nbconvert_exporter": "python",
   "pygments_lexer": "ipython3",
   "version": "3.9.13"
  },
  "vscode": {
   "interpreter": {
```

```
    "hash":
"6e9e0baa62560f8a3b402c12d339bdad33c58a25305700ec7e7682c0b6251f68"
   }
  }
 },
 "nbformat": 4,
 "nbformat_minor": 5
}
```