Installing HVR on AWS using HVR Image

HVR Image for AWS is a pre-configured Linux-based AMI (Amazon Machine Image) to run HVR's remote listener (as a daemon) including Redshift drivers. HVR Image for AWS supports the use as both a remote HVR agent (for capture and/or integration) as well use as a hub machine using either a PostgreSQL database on the machine, an RDS database, or a database running on a different machine.

HVR Image for AWS is available from the AWS Marketplace and includes all necessary components to connect to Amazon Redshift, Oracle and PostgreSQL on RDS, S3 or any HVR supported target database on EC2, enabling replication from/into all supported on-premises platforms.

HVR Image for AWS is currently available only on Linux. Connectivity to SQL Server requires an HVR installation on Windows (optionally running as an agent).

Using HVR Image for AWS to create a remote listener agent

To use the HVR Image for AWS:

1. Sign in to the AWS portal and select **EC2** under the **Services** menu.
2. Click the **Launch Instance** button.
3. On the right side menu bar, click **AWS Marketplace** and type 'HVR' in the search field.
4. Select the HVR for AWS offering suitable for you.

   For optimum efficiency make sure HVR in AMI runs in the same region as the database or service taking part in real-time data movement.

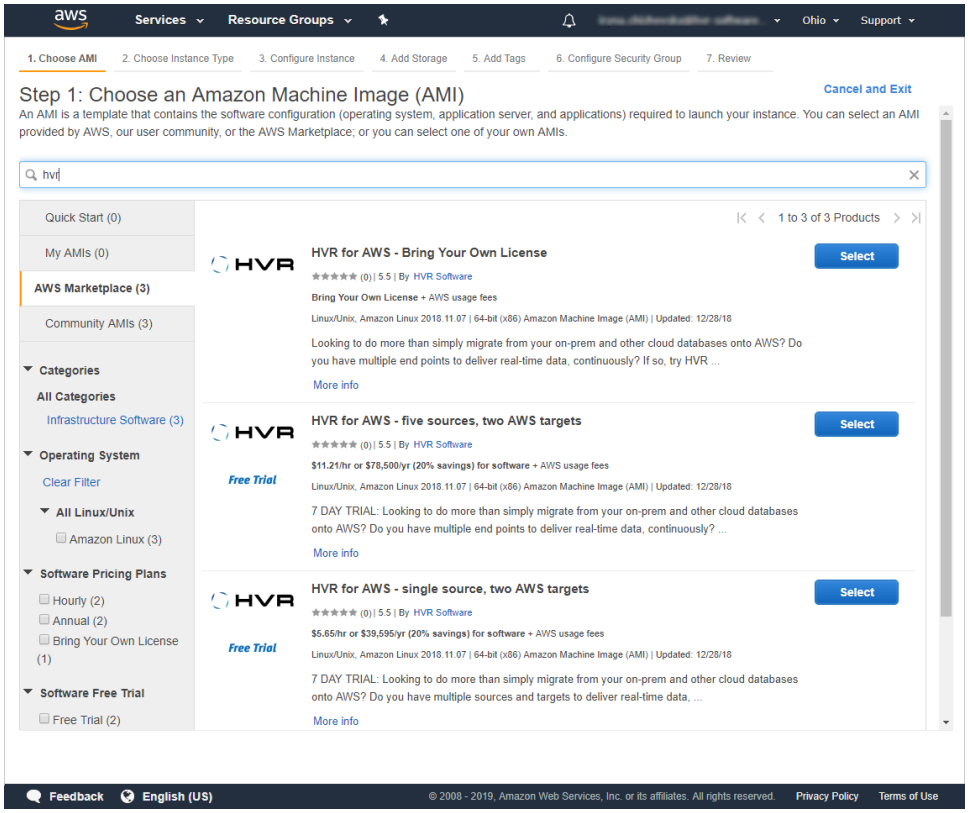

5. HVR will show you the product details, click **Continue**.

6. Under the **Choose an Instance Type** tab, select a required one and click **Next: Configure Instance Details**.

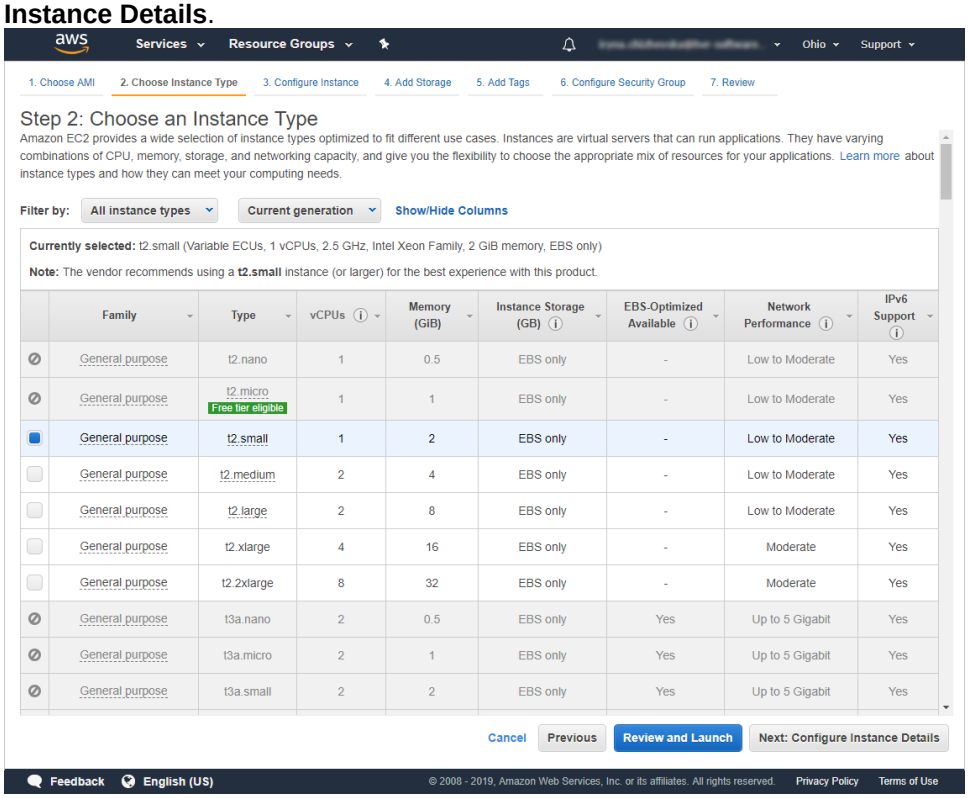

7. If appropriate, under the **Configure Instance** tab, select a preferred network (VPC) subnet or use the default one.

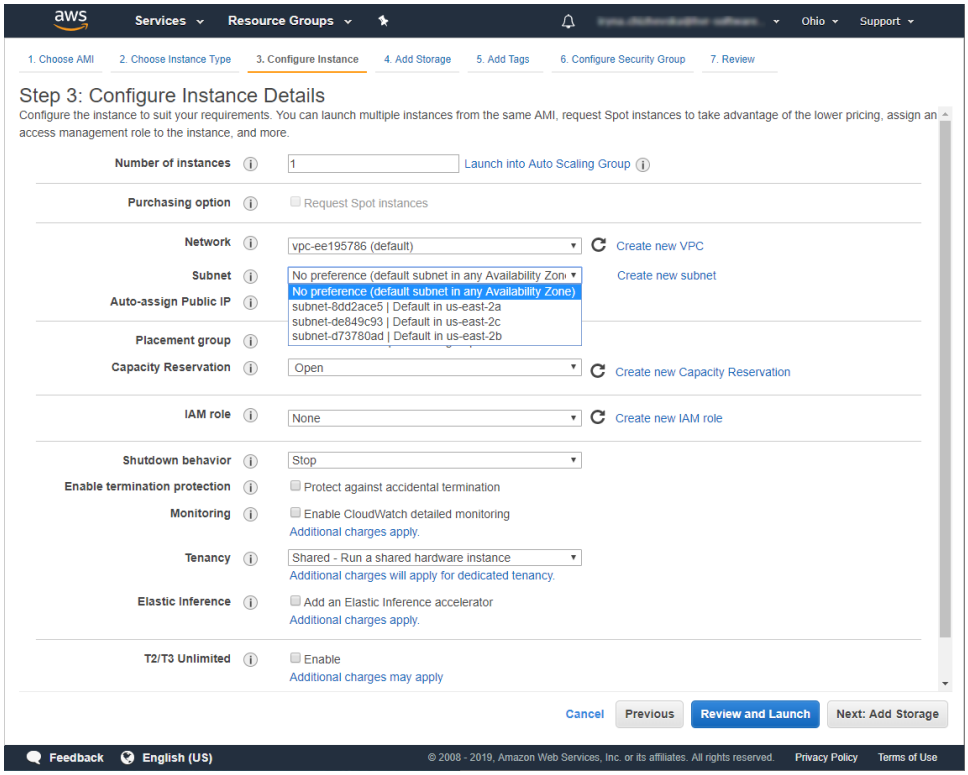

8. Under the **Add Storage** and **Add Tags** tabs, proceed with default settings.
9. Under the **Configure Security Group** tab, set up traffic rules for your instance. HVR uses TPC/IP connectivity on port 4343. Configure an incoming connection for the HVR listener daemon on this port and limit the IP range as narrow as possible. Click **Review and Launch**.

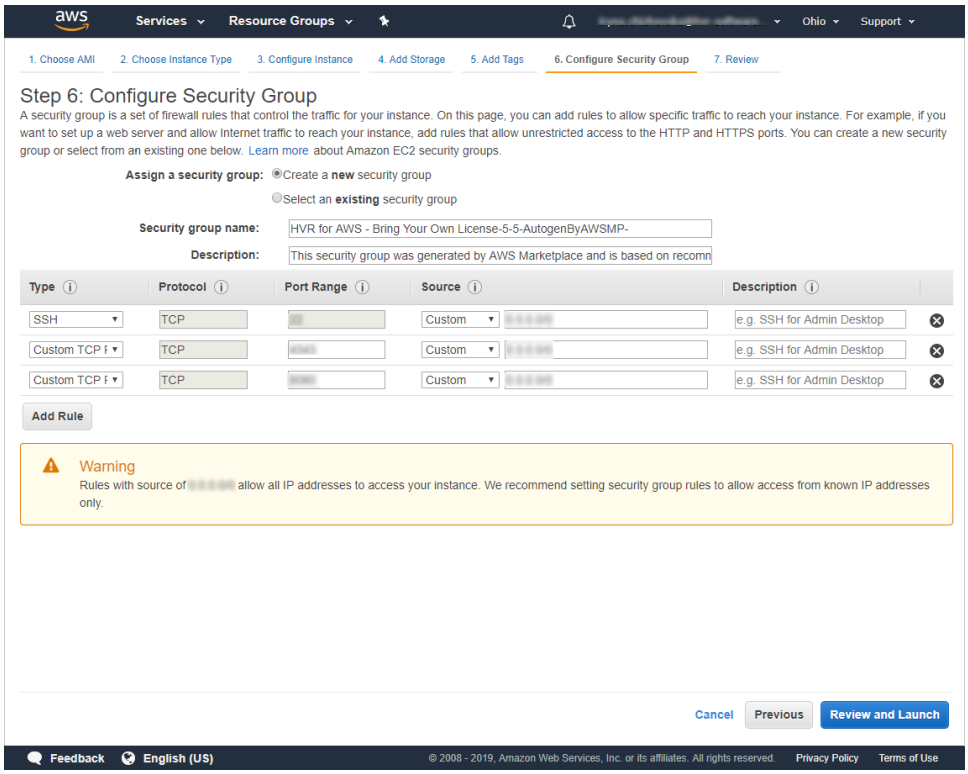

10. Under the **Review** tab, review your instance launch details and click **Launch** to assign a key pair to your instance and complete the launch process. An EC2 key pair will be used to securely access the AMI. If you have an existing EC2 key pair defined, you can use that. If you don't have a key pair defined, you can create one. Once created, the associated private key file will be downloaded through your web browser to your local computer.

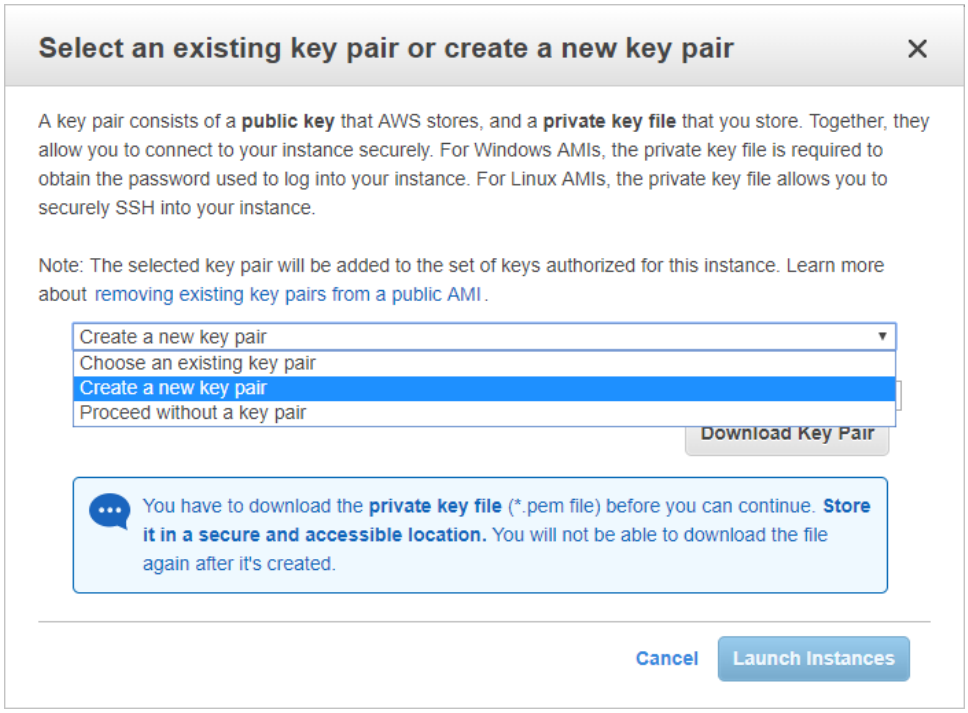

11. You will now see the AMI being deployed in the Marketplace. This may take a few minutes. The details of the created AMI can be accessed in the AWS EC2 console. Note down the Public DNS name or IP address of the AMI as this will be required when connecting to the AMI.

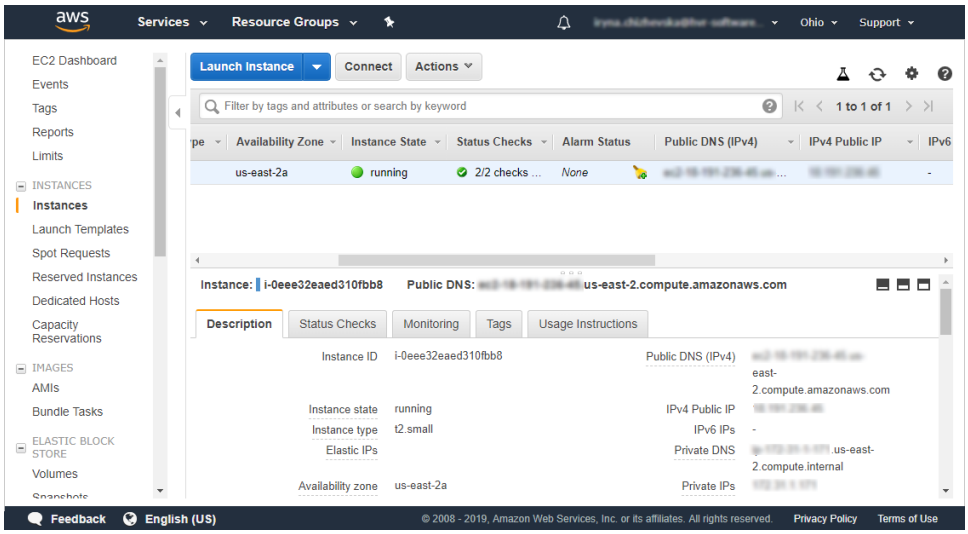

12. Once the AMI has been deployed and is running you can start an ssh session to the AMI instance to obtain information on how to proceed. To connect, use the **pem** key that was used when the image was created. Connect as a user named **ec2-user**:

```
ssh -i <key-pair-name>.pem ec2-user@<Public-DNS-name>
```

13. The Message of the Day that you see when connecting to the EC2 instance contains important information for using this image. In particular, it provides you with the file location of the SSL public certificate you need for an on-premises HVR installation to connect to the HVR AMI and the location on the AMI from which you can download the HVR software to install on-premises.

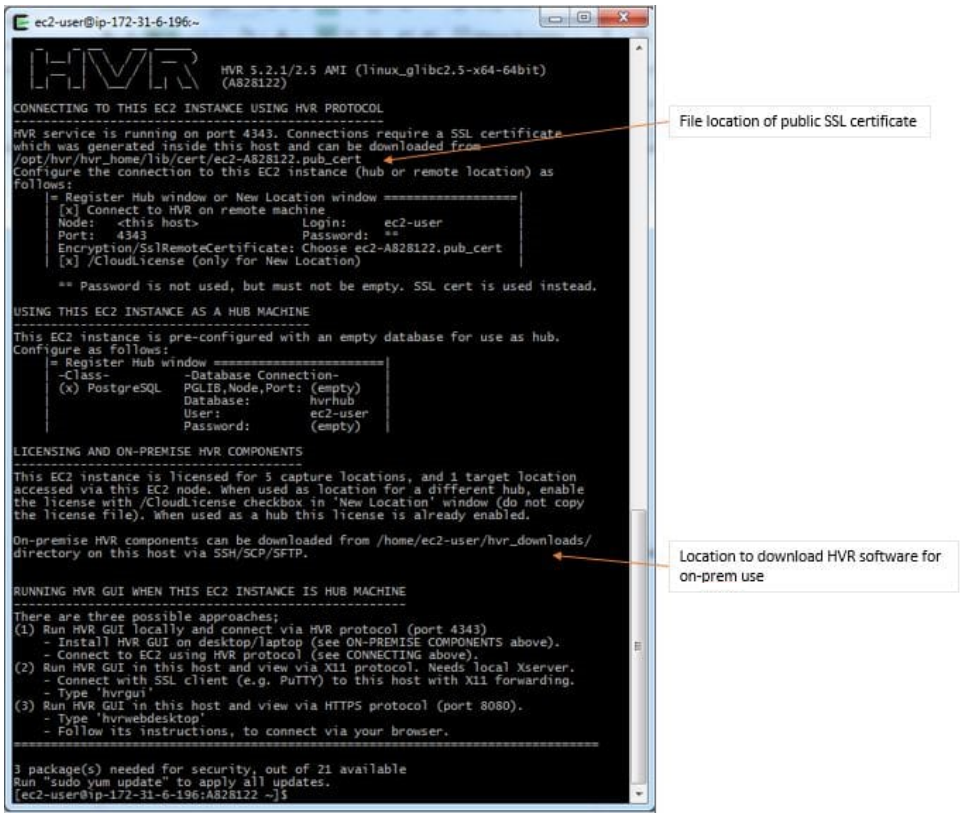

14. opy the public certificate **.pub_cert** file to your on-premises hub machine and download the Linux and Windows HVR packages. To do this, use your preferred file transfer tool, such as sftp, scp or a GUI-equivalent application like WinSCP.