



# Building Render Farms in the Cloud—Using Amazon EC2 for On-Demand Rendering

Alex Lorman – TITAN Salvage

## VI5718-R

Not everyone can afford either the cost or infrastructure of a permanently installed server farm used for completing projects. This course will teach users the most effective methodologies for quickly creating a server farm (based on the Backburner application) using Amazon.com, Inc.'s, Elastic Compute Cloud (EC2) servers, which are paid on an on-demand basis. We will cover basic use of Amazon's system, and we will focus on how to configure specific parameters within the Backburner app so that rendering can run successfully on a virtualized on-demand network.

## Learning Objectives

At the end of this class, you will be able to:

- Know the basic operation of the Amazon EC2 control panel
- Know the differences between deploying the Backburner application locally and remotely
- Be able to troubleshoot issues between computers on the Amazon EC2 farm
- Activate an entire farm on Amazon EC2 on demand

## About the Speaker

Alex Lorman earned a BS in architecture from the Catholic University of American in 2011, and he worked in architecture until he became bored and managed to talk TITAN Salvage's managing director, Rich Habib, into hiring him on a freelance basis. He has worked on many TITAN projects over the last 2 years, including the ongoing removal of the Costa Concordia from the rocks in Italy. Alex uses any and all software that can help him to do his job with increased ease, quickness, and accuracy, and he specializes in taking ideas from conception through their construction in steel—on time, on budget, and to specification, by whatever means necessary. Whether he's in front of a computer or in coveralls with a welding rig, large-scale projects and short deadlines are just parts of the challenge. Great software can always make it better.

Email: [alex.lorman@titansalvage.com](mailto:alex.lorman@titansalvage.com)

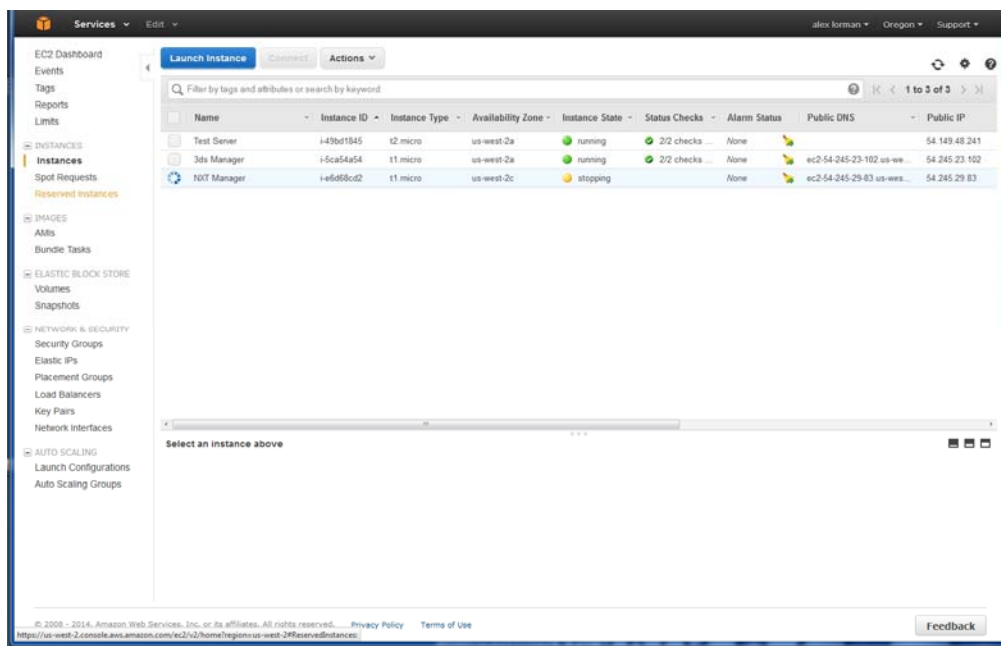
LinkedIn: Alex Lorman

*This outline assumes a basic knowledge of the following concepts and thus they are not covered in full detail:*

- *Virtual Machine terminology*
- *Remote Desktop Connection Usage in Windows*
- *Installation of the Backburner and Associated software on machines.*
- *Usage of Backburner on a local LAN*

### The EC2 Control Panel

The EC2 Panel is the web interface provided by Amazon to allow for control of groups of virtual machines. ***These virtual machines are known in Amazon terminology as Instances.***



Above: This shows the basic EC2 control panel when viewed in a web browser. This screen shows three instances in existence with two running and one powering off.




We can also see the instances above are listed as either t2.micro or t2.micro. This refers to the specification of the hardware associated with the given Instance.

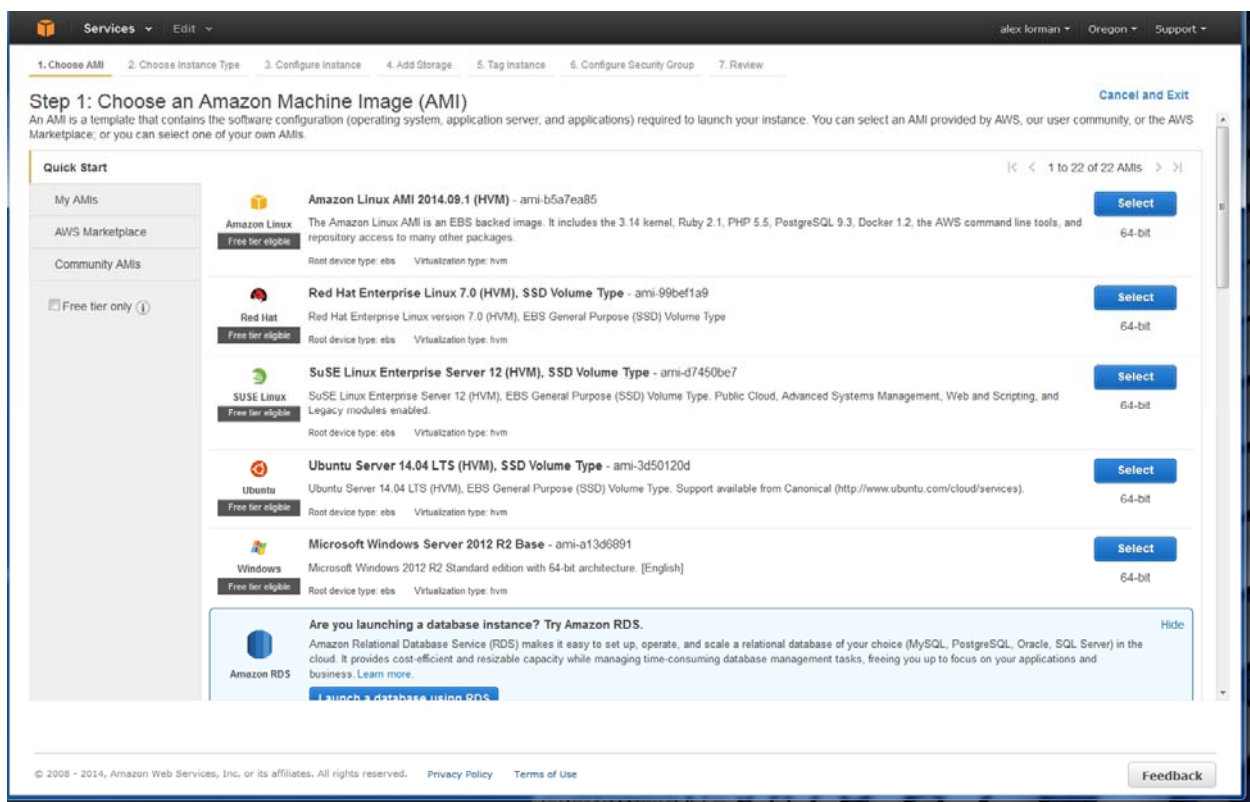
Instances can also be renamed using the name field. We can see that these are named “Test Server” et al.

### Launching an Instance


To create a new Instance press the Launch Instance button



You are then presented with the following selection to determine both the OS configuration and power specifications of the new machine.



Above we can see that we can create a new instance in almost any OS. We will be using Windows Server exclusively for Backburner. ***These machines are all “bare-metal” setups, they contain nothing but the operating system.*** We would use these when creating a Server or Manager machine from scratch. If we already had configured a machine and wanted to simply create more instances of it we would use the “My AMIs” tab.

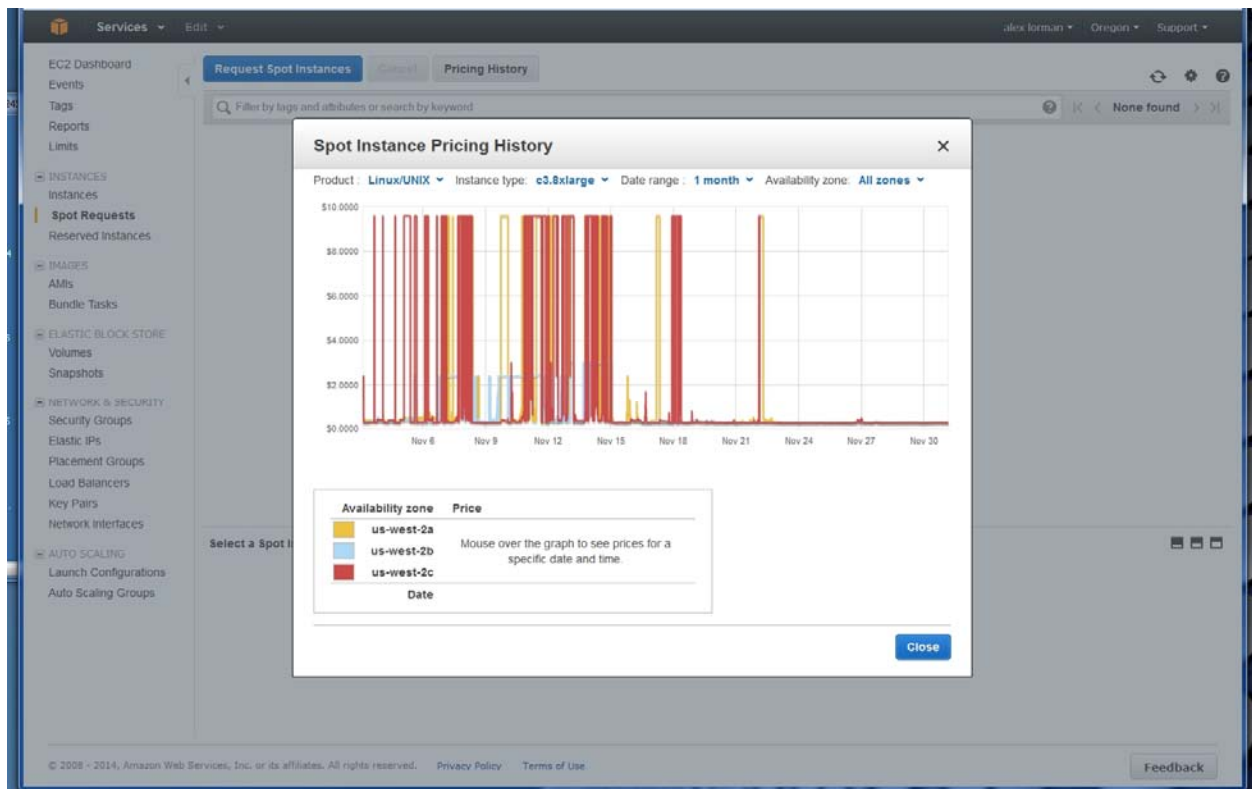


The next screen will guide us on to Specifications of the Instance.

NB: For Backburner, the Compute Optimized will work the best for CPU heavy rendering.

### Instance Pricing

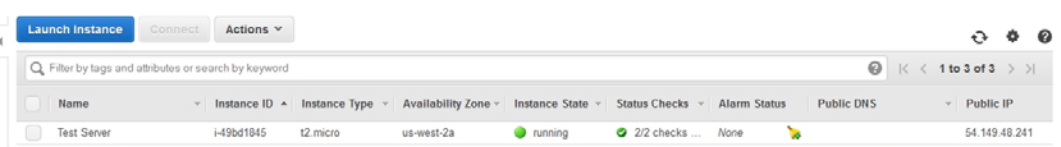
Instances are priced in two ways; On-Demand and Spot. On-Demand Instances are priced substantially higher per hour (all EC2 time is billed by the hour, rounded up the nearest full hour) however is guaranteed for uptime and availability. Spot Instances are a way of bidding on the unused computing capacity EC2 may have at a given time, ***however if your pricing instructions will not go above a certain point your Instances may be terminated.***



The above shows the price fluctuations of a Computer Optimized Instance over a given time frame. Clearly the pricing is market driven.

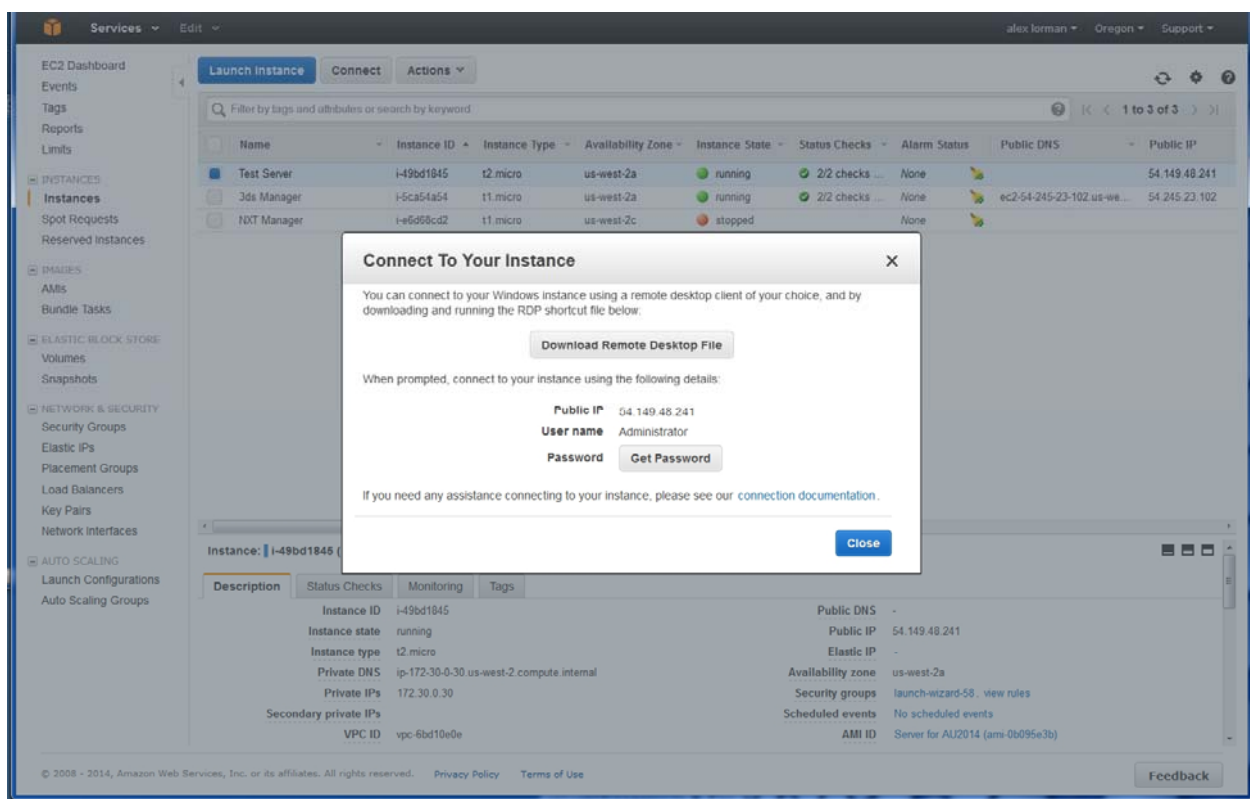
### Connecting to your Instances

Once you have launched an Instance and it has gone through the preflight checks it will display like the below;



Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS	Public IP
Test Server	i-49bd1845	t2.micro	us-west-2a	running	2/2 checks ...	None		54.149.48.241

Highlight the Instance and press “Connect” and the following screen will appear



**Connect To Your Instance**

You can connect to your Windows instance using a remote desktop client of your choice, and by downloading and running the RDP shortcut file below.

[Download Remote Desktop File](#)

When prompted, connect to your instance using the following details:

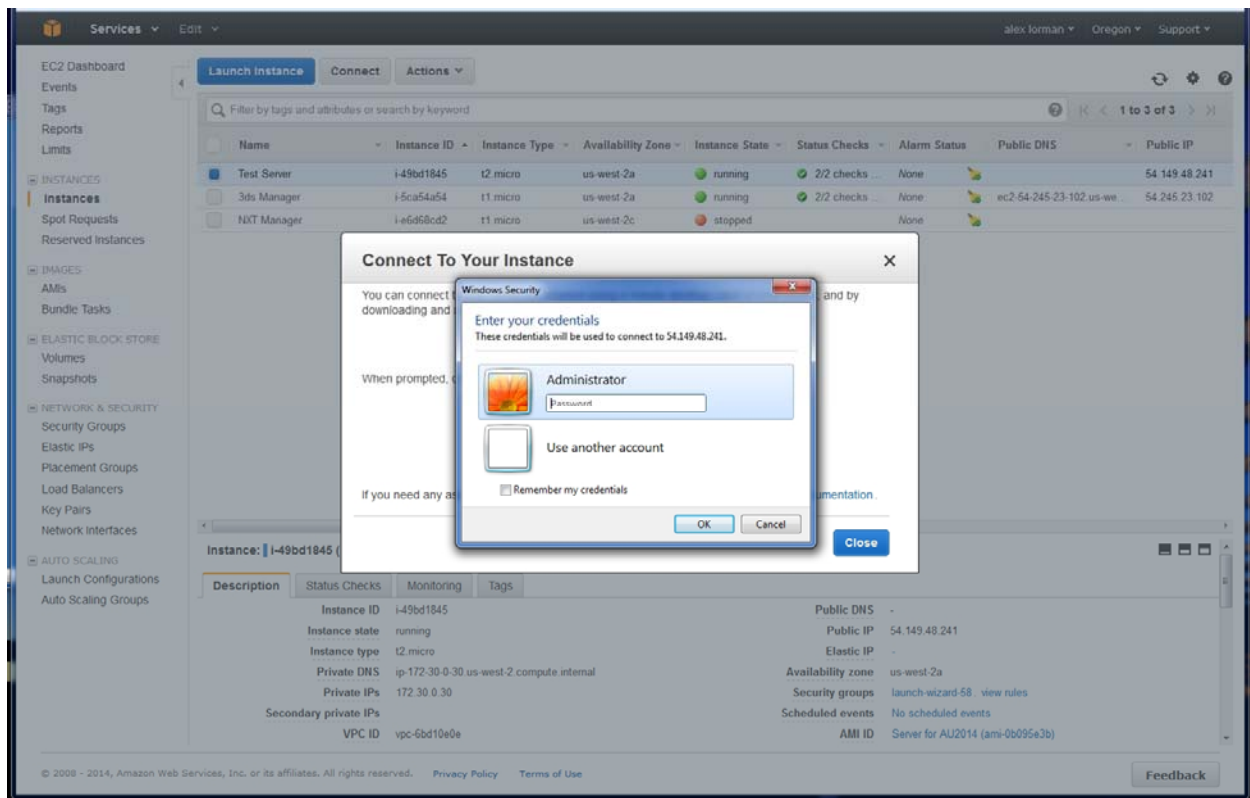
**Public IP** 54.149.48.241  
**User name** Administrator  
**Password** [Get Password](#)

If you need any assistance connecting to your instance, please see our [connection documentation](#).

[Close](#)

Downloading the remote desktop file is by far the easiest way to connect. It will then auto-launch Windows Terminal Services Client and behave exactly like a traditional remote server.

You will then be asked to log in via MSTSC.



There are several different ways to deal with passwords on Instances.

- Key pairs – Amazon has a facility within the EC2 environment that can allow the machines to decrypt a password, on the fly, from an encryption key provided by the user to the web interface. Amazon will walk you through this process. This way the password is different for each machine although very secure.
- Hard coded passwords – Although less secure since any instance launched from the same AMI (images, mentioned in just a moment) will have the same hard-coded password, for something such as a render farm where the machines are likely to be all the same and to be used fairly ephemeraly this is an easy option,
- No Password – This may or may not work with MSTSC (since sometimes it requires a password) and is not recommended in either case.

### What are AMIs?

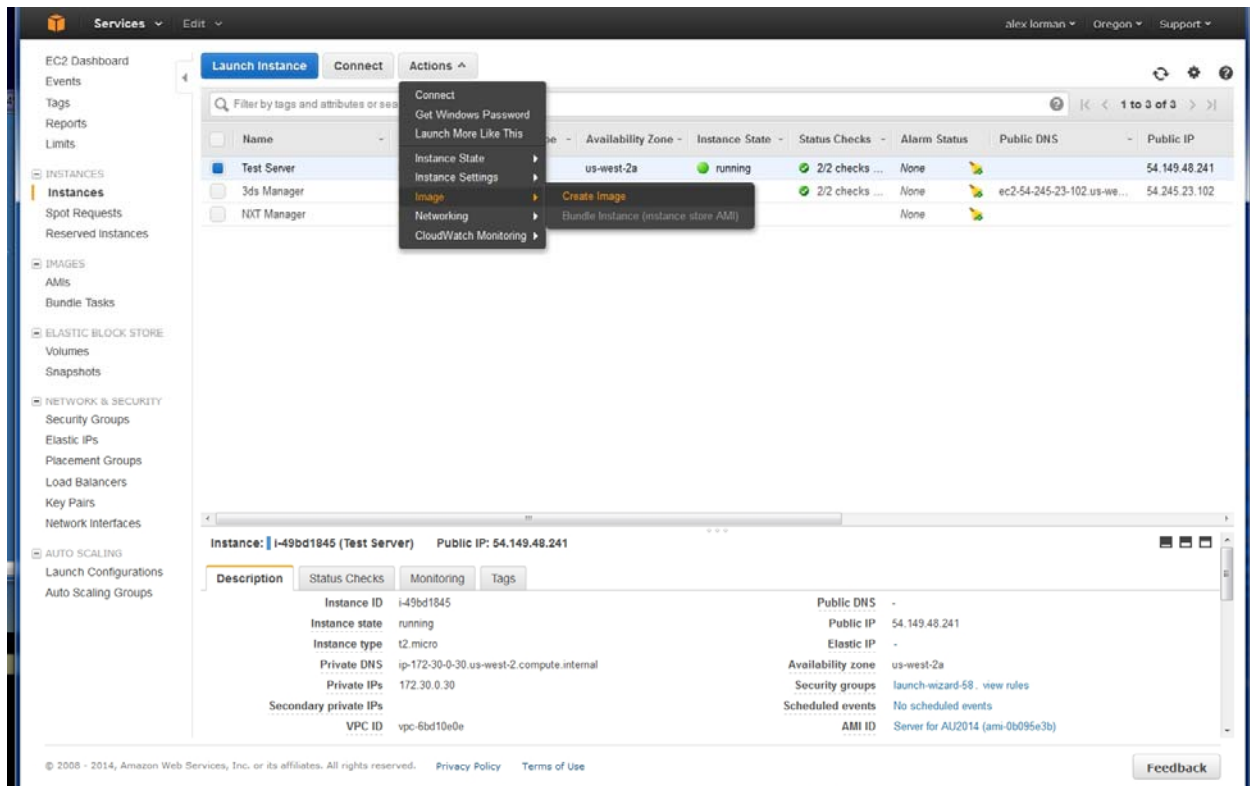
AMIs (in Amazon parlance) are system images that have been captured from built-up Instances and can be then poured into a new Instance as requested. **These function exactly like images or snapshots taken from traditional hard drives.** The easiest workflow to creating machines that can be re-used with no additional setup is as follows;

- Create a new Instance (easiest and cheapest to have it be a Micro Instance while building up the system)
- Install needed software
- Change setting as desired

- Capture machine as an AMI
- Re-Launch as many clones from that AMI as needed.

### Creating an AMI


Once you have built a machine to your satisfaction you can click on it and create an AMI through the following screen;



This takes just a few minutes but does involve a system reboot.

### Launching Instances based on AMIs

When configuring new Instances and you wish to use the instance you have created (i.e. one with Backburner configured properly) you need simply press the following button when launching a new Instance.





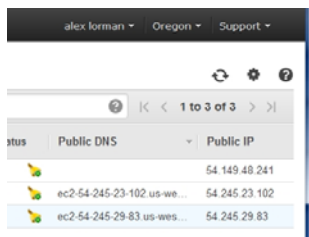
## Snapshots, AMIs and Volumes

A small point of clarification;

- AMIs represent an entire virtual machine, in a ready-to-run state. If set up, an AMI can contain multiple drives of varying sizes.
- Snapshots are created periodically and automatically and are simply a snapshot of a specific drive. These can be OS based drives or storage based drives. Snapshots cannot be used to launch Instances.
- Volumes are simply Amazon's way of referring to drives. Again these can be for storage or OS use and can be of arbitrary size. They are generally priced by the GB.

## Elastic IPs and Instances

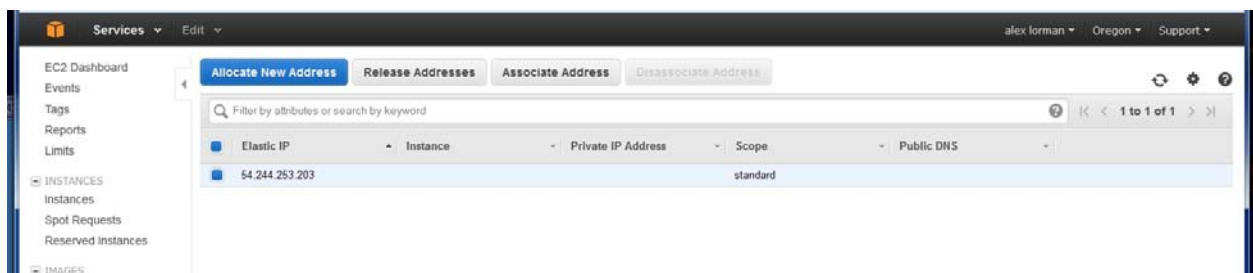
When creating instances it is important they be reachable on the correct ports so that RDC and similar services can connect. All Instances have a Public DNS address that can be used to connect by RDC. This address, however, can change at any time. By default firewalling and port forwarding is enabled for remote access to the machines.



Instance ID	Public DNS	Public IP
ec2-54-245-23-102 us-we...	ec2-54-245-23-102.us-west-2.amazonaws.com	54.245.23.102
ec2-54-245-29-83 us-we...	ec2-54-245-29-83.us-west-2.amazonaws.com	54.245.29.83

Sometimes it becomes either necessary or far easier to have a computer have a traditional IPV4 external facing IP address. This serves as a permanent way of calling the machine by that address and while it can be changed between machines it will not expire or be arbitrary.

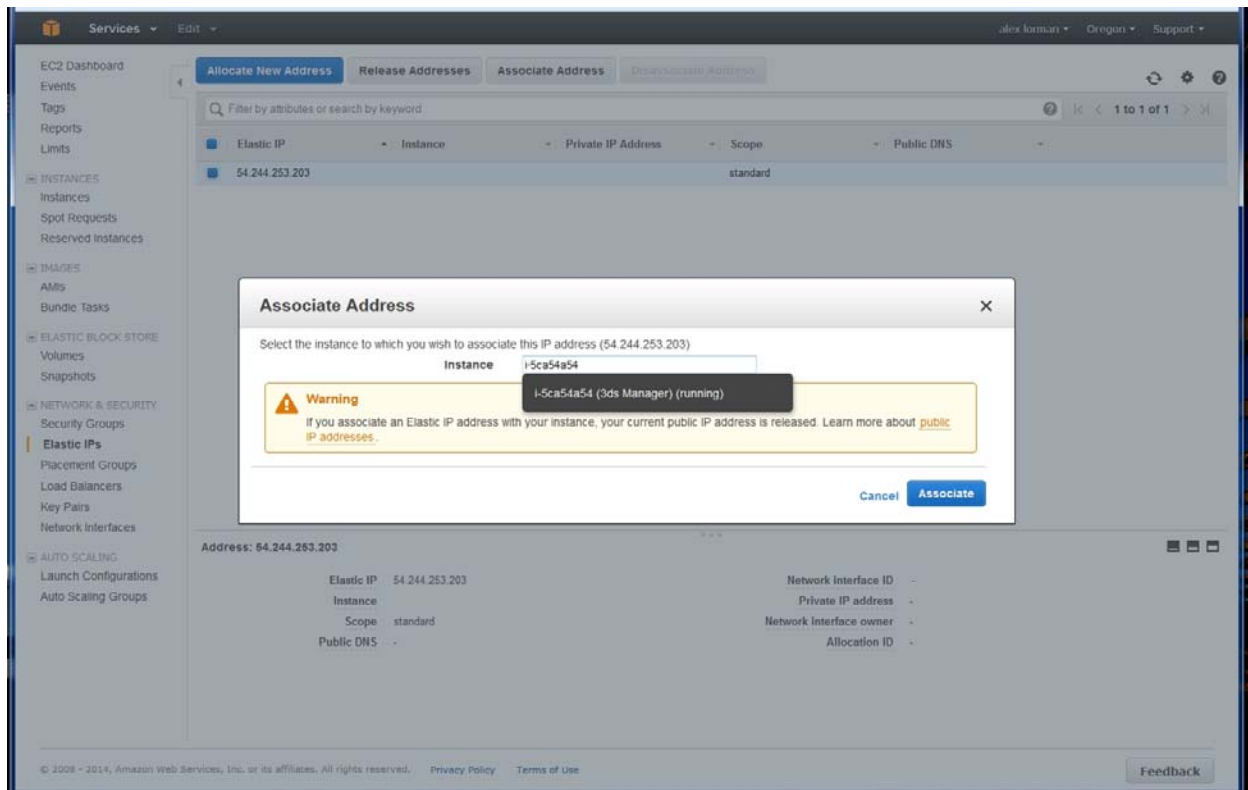
Here we can see the one Allocated Address that we have although it is not currently associated with an Instance.



Elastic IP	Instance	Private IP Address	Scope	Public DNS
54.244.253.203			standard	

Associating it with a running Instance is quite easy;



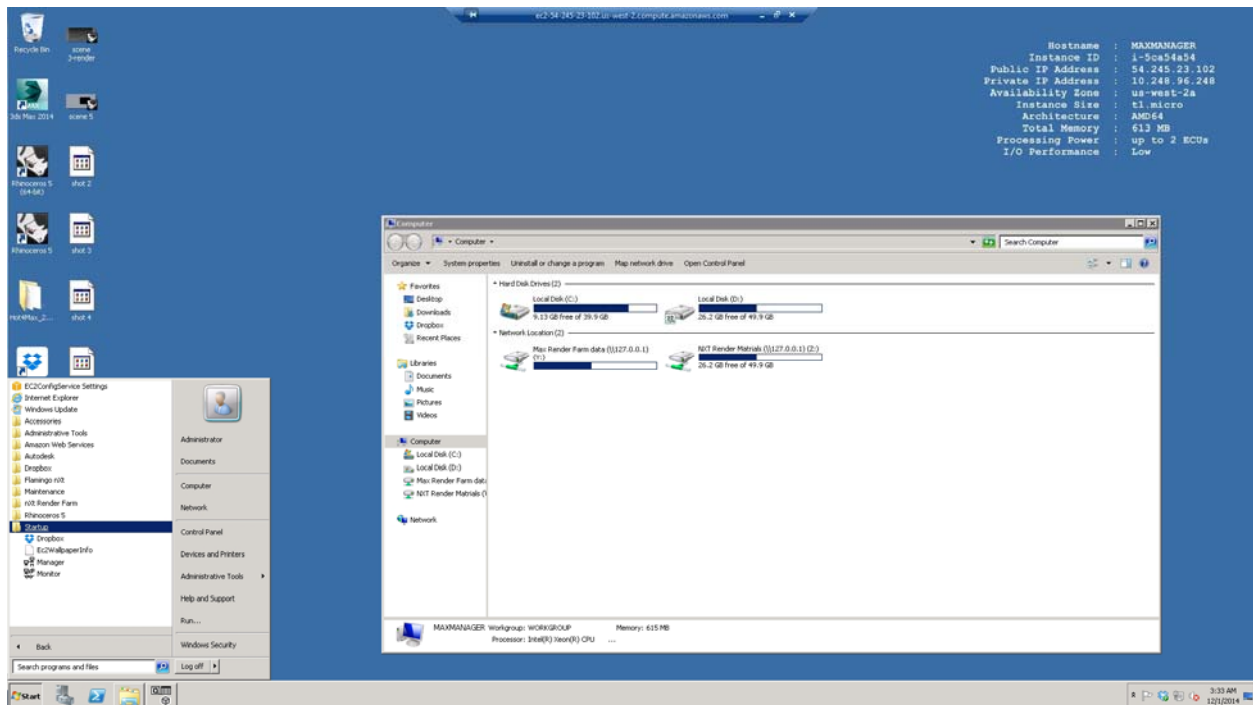


***NB. I have found it works best to associate an Elastic IP with the manager machine then call that machine using that address within all Backburner uses, rather than having the Servers look for a machine. The only caveat with this is that Amazon will only allow you 6 Elastic IP addresses without further approval, so running multiple farms on the same network might be a challenge.***

### Deploying Backburner on Instances

Once bare-metal Instances are set up installing Backburner on them is relatively straightforward. Once the installer is on the Instance it is simply a matter of running as one would a normal server.

Below is a desktop capture of a running Instance with Backburner installed.

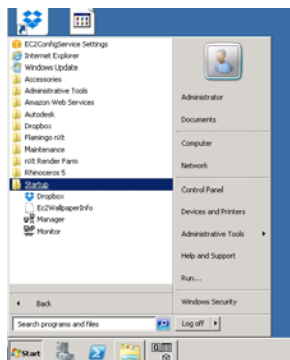


### Setting up the Manager

The Manager machine will be a slightly different configuration than the Server nodes and a few considerations need to be taken into account. The Manager computer, in this case, will also be our source of storage for the returning images and base render files.

***When building Manager from an Instance (or building fresh) be sure to equip it with an extra drive of suitable size. This will then be a logical drive in windows and be shared as normal.***

After you have installed Backburner (at Startup, rather than as a system service) make sure that Manager (and Monitor) are in the Startup folder so that they will run on windows startup.



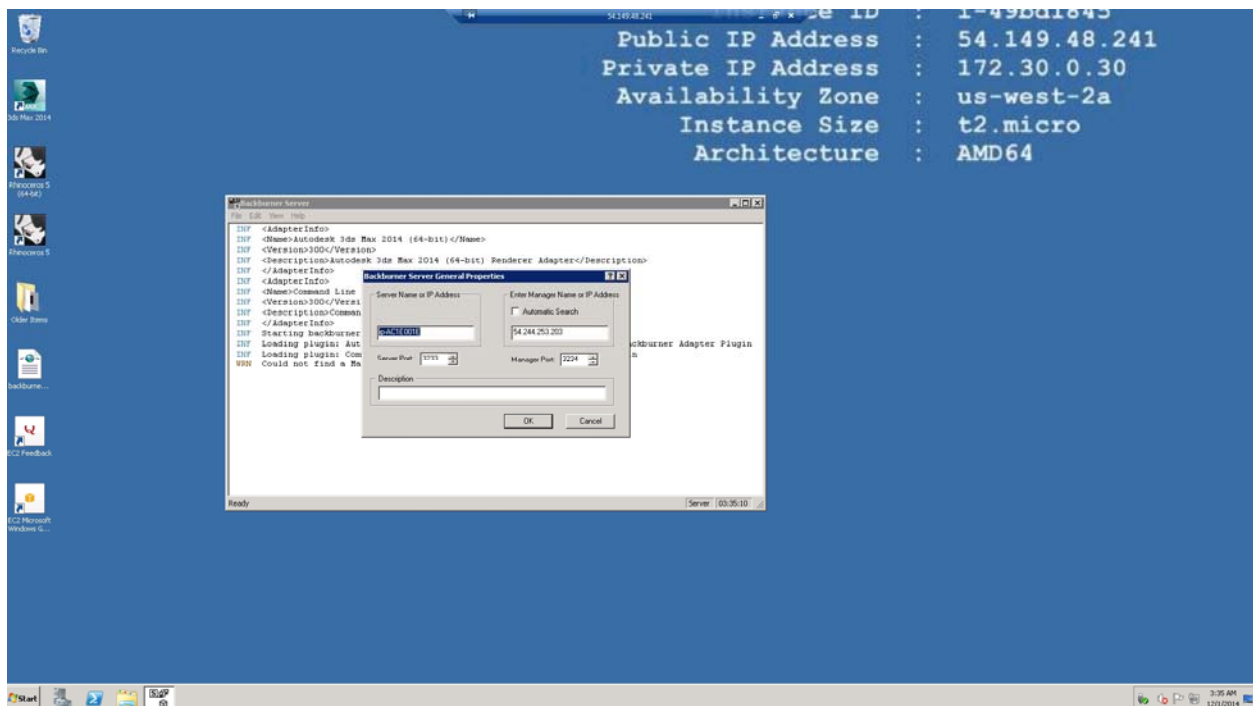
The stock settings for Manger should be sufficient here, although any specific requirements will be known to you when doing your own setup.

**NB; Be sure to set up a network share of the storage drive and allow Read and Write access so that the Server machines read and write to it. This is more of a general Backburner discussion rather than something specific to EC2**

## Setting up the Server Nodes

Once a Manager machine has been set up, the next step is to launch another bare-metal Instance and configure it to be the Server. This involves installing Backburner as normal except configuring Server to start at startup instead of Manager.

After installing Server the configuration panel should look like this:



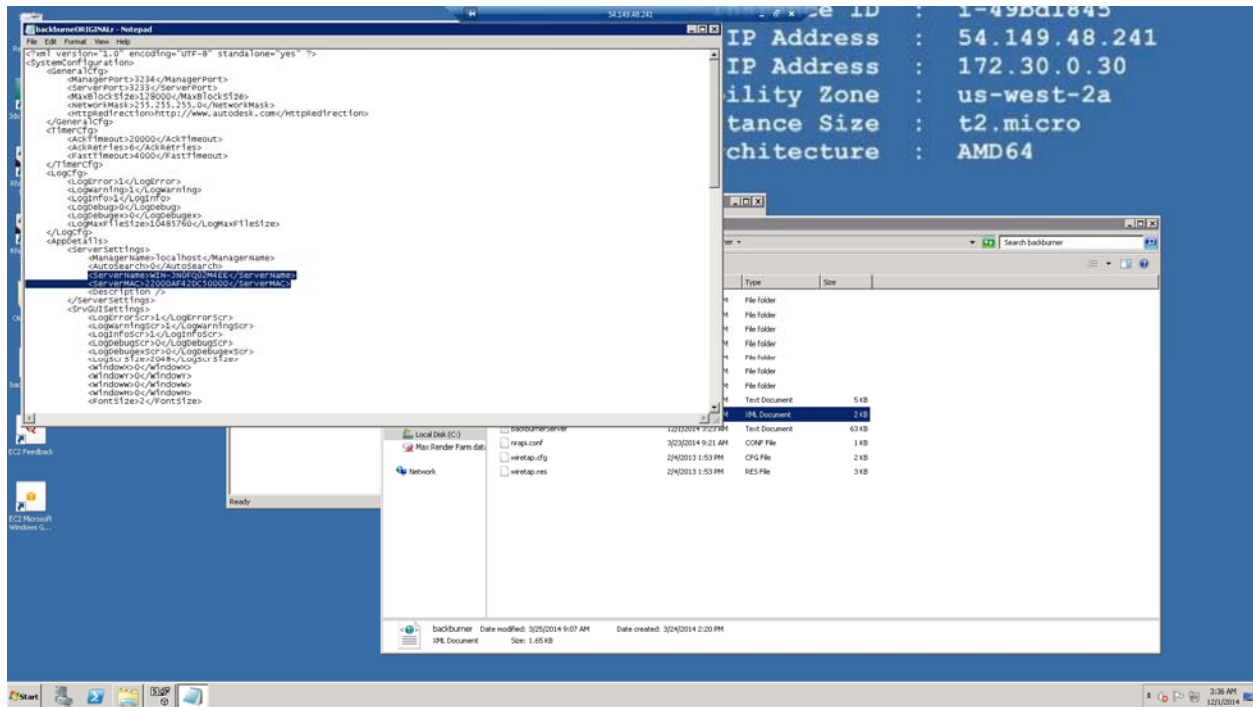
The Manager name, in this case, is hard coded into the configuration so that it will most easily find it.

**NB The one key difference between configuring Server on a single Instance and cloning that Instance onto other machines becomes Machine naming. Machines created from the same AMI will, by default, have the same Computer Name in Windows. This creates problems in Manager since it will not detect the different machine names.**

To change the above problem you need to open and edit the Backburner.xml file. Located in C:>Users>XX>AppData>Local>Backburner>Backburner.xml

(Where XX is the given user. Generally Administrator)

Seen here the following lines need to be deleted..



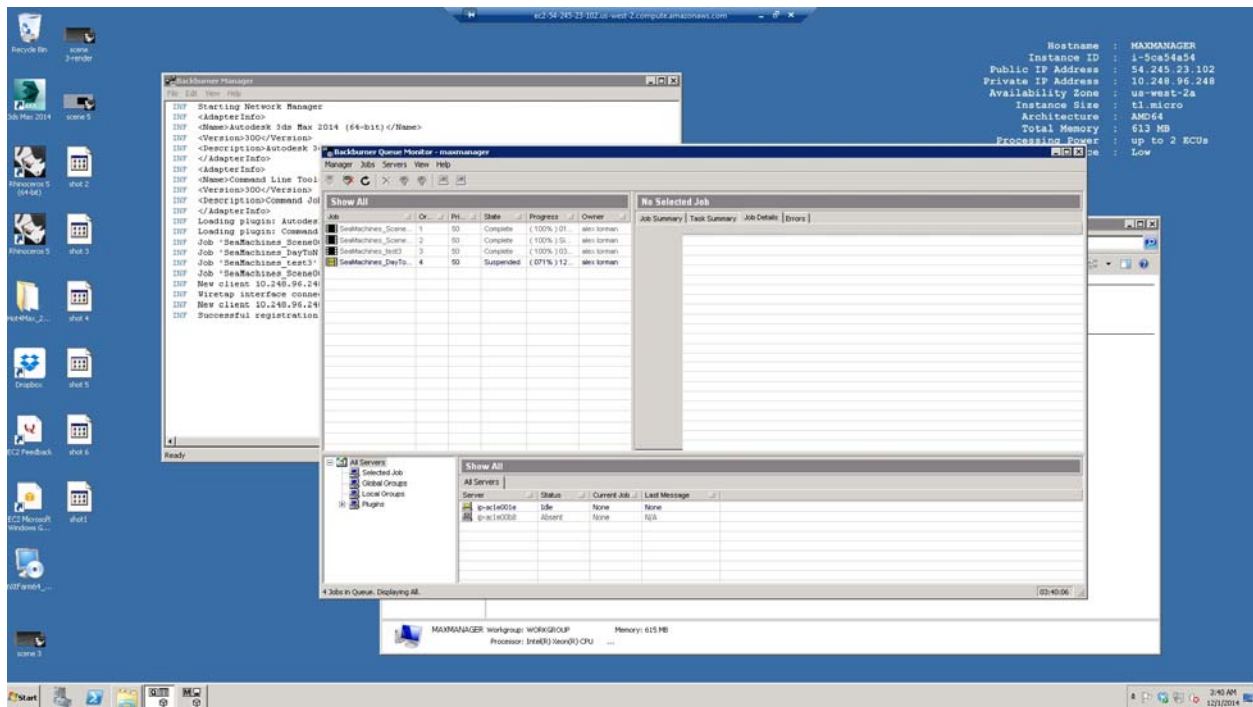
Then re-save the file and create an AMI before restarting Backburner. This will allow each machine, when started, to write in their own unique machine name and will allow them to communicate properly with Manager.

When configuring the Server nodes, before creating the final AMI, make sure to have connected the network drives to the manager so that they load when the machine launches. This will allow it to write the data back to Manager Machine.

## Activating a Farm

Once Manager and Server have both been imaged and launched, the server should register with the Manager as normal and you should see a screen similar to this on the Monitor.

## Building Render Farms in the Cloud—Using Amazon EC2 for On-Demand Rendering



From here you are in a position to launch as many of your Server AMIs as you deem needed for the rendering job. Either using Spot or Normal Instance pricing depending on your needs.

From here jobs are managed as normal within Backburner and jobs should render normally on as many machines as you have launched.