

MANJRASOFT PTY LTD



Installation and User Guide

Aneka 3.0

Manjrasoft

5/13/2013

This document describes in detail the steps involved in installing and configuring an Aneka Cloud. It covers the prerequisites for the installation, the installation process, setting up and configuring Aneka Daemons and Aneka Containers, and running some of the sample applications to test your new cloud deployment.

Table of Contents

1	Introduction.....	1
1.1	What is an Aneka Cloud composed of?.....	1
2	System Requirements.....	3
2.1	Hardware Requirements	3
2.2	Software Requirements	3
3	Installation	4
3.1	Installing Aneka Cloud Management Studio	4
3.1.1	Step 1 - Run the installer package to start the Setup Wizard	4
3.1.2	Step 2 - Specifying the installation folder	4
3.1.3	Step 3 - Confirm and start the installation	5
4	Aneka Cloud Management Studio	8
4.1	Starting up Management Studio	8
4.2	Shutting down Aneka Management Studio	9
4.2.1	The Configuration File	9
4.3	Cloud Explorer - Overview	10
5	Infrastructure Management.....	12
5.1	Adding a new Machine	12
5.2	Importing and Exporting Machine Lists	12
5.3	Repository Management	13
5.4	Machine Credentials.....	15
5.5	Access Denied Machines	17
5.6	Uninstalled Machines.....	18
5.7	Installed Machines	19
6	Installing the Aneka Daemon	20
6.1	Step 1 - Repository and Port Configuration	20
6.2	Step 2 - Review Summary and Install	21
7	Creating an Aneka Cloud.....	25
7.1	Installing the Master Container	25
7.1.1	Step 1 - Container Type and Port Configuration	25
7.1.2	Step 2 - Security Configuration.....	26
7.1.3	Step 3 - Persistence Configuration.....	27
7.1.4	Step 4 - Cost and Software Appliances	30
7.1.5	Step 5 - Failover Configuration.....	31
7.1.6	Step 6 - Service Configuration.....	32

7.1.7	Step 7 - Summary	33
7.2	Installing Worker Containers	34
7.2.1	Step 1 - Container Type and Port	34
7.2.2	Step 2 - Security Configuration.....	36
7.2.3	Step 3 - Persistence Configuration.....	36
7.2.4	Step 4 - Cost and Software Appliances	37
7.2.5	Step 5 - Failover	38
7.2.6	Step 6 - Service Configuration.....	39
7.2.7	Step 7 - Summary	40
8	Cloud Management.....	42
8.1	Master Containers	42
8.2	Quarantined Containers	43
8.3	Orphan Containers.....	44
8.4	Container Management	46
8.5	Real-Time Performance Monitoring	47
8.6	Reporting and Accounting	48
8.7	User Management.....	52
8.8	Host Name Conversion	52
9	Testing the Cloud with Sample Applications	56
9.1	Thread Programming Model	566
9.2	Task Programming Model	577
9.3	MapReduce Programming Model.....	599
	Appendix A - Service Configuration.....	62

1 Introduction

Aneka is a Cloud Application Development Platform (CAP) for developing and running compute and data intensive applications. As a platform it provides users with both a runtime environment for executing applications developed using any of the three supported programming models, and a set of APIs and tools that allow you to build new applications or run existing legacy code. The purpose of this document is to help you through the process of installing and setting up an Aneka Cloud environment. This document will cover everything from helping you to understand your existing infrastructure, different deployment options, installing the Management Studio, configuring Aneka Daemons and Containers, and finally running some of the samples to test your environment.

1.1 What is an Aneka Cloud composed of?

An Aneka Cloud is composed of a collection of services deployed on top of an infrastructure. This infrastructure can include both physical and virtual machines located in your local area network or Data Centre. Aneka services are hosted on *Aneka Containers* which are managed by *Aneka Daemons*. An Aneka Daemon is a background service that runs on a machine and helps you to install, start, stop, update and reconfigure Containers.

A key component of the Aneka platform is the Aneka Management Studio, a portal for managing your infrastructure and clouds. Administrators use the Aneka Management Studio to define their infrastructure, deploy Aneka Daemons, and install and configure Aneka Containers. The figure below shows a high-level representation of an Aneka Cloud, composed of a Master Container that is responsible for scheduling jobs to Workers, and a group of Worker Containers that execute the jobs. Each machine is typically configured with a single instance of the Aneka Daemon and a single instance of the Aneka Container.




Figure 1 - Management Scenario.

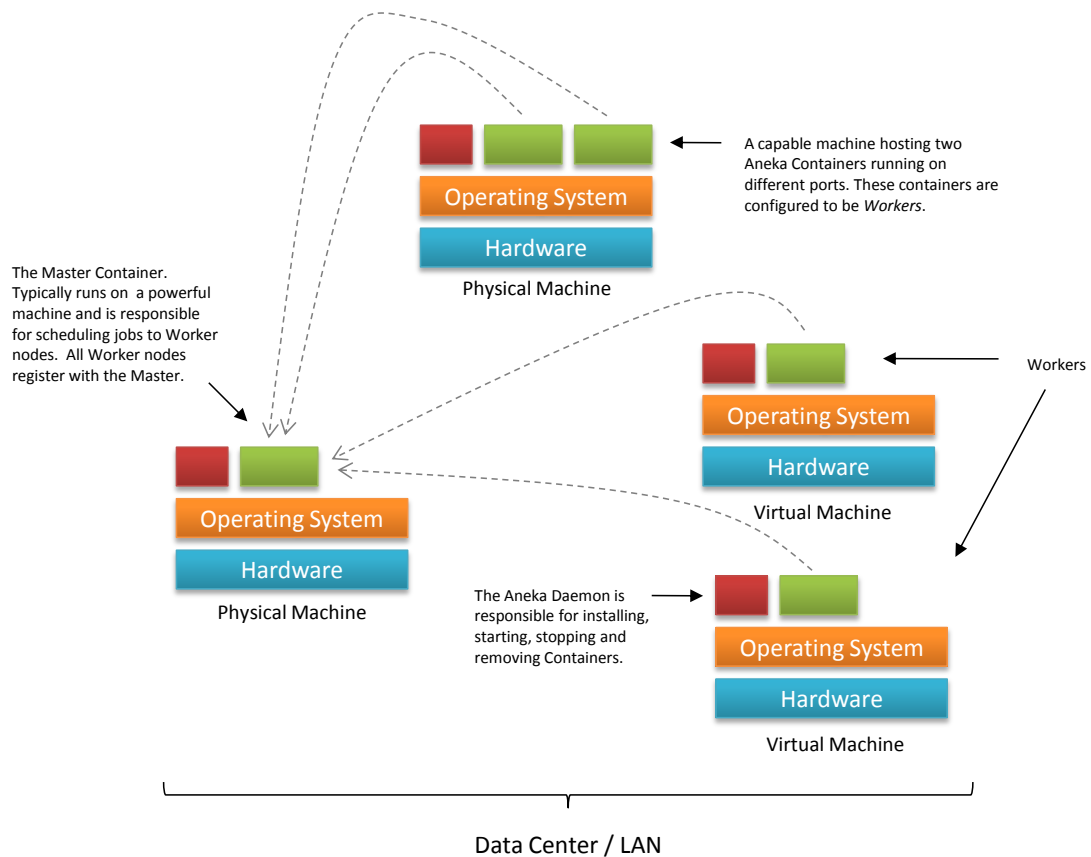


Figure 2 - A high-level view of an Aneka Cloud

The following sections provide detailed instructions on installing and configuring your Clouds.

2 System Requirements

This section describes the hardware and software requirements for creating your cloud environment. For a more detailed and accurate description of the requirements see release notes.

2.1 Hardware Requirements

Disk Space: Installation of Aneka using the distributed Microsoft Installer Package (MSI) requires approximately 60 megabytes of disk space. This installation includes the Aneka Cloud Management Studio, binaries for installing Aneka Daemons and Containers, the Design Explorer, sample applications and documentation. The Aneka Cloud Management Studio is your portal for setting up an Aneka Cloud environment. This involves installing the Aneka Daemon and Aneka Containers on remote machines. The Aneka Daemon requires approximately 10 megabytes of disk space, and a Container requires approximately 20 megabytes.

NOTE: While the above describes the minimum disk space requirements for installation, running Aneka requires additional disk space for creating temporary files such as logs, storing data files from users that are required for execution, and creating new data output files as a result of the execution. This additional disk space required therefore depends on how you plan to use Aneka. See sections below on Storage Service and Execution Services for ideal disk space requirements.

CPU and Memory: The CPU and memory requirements differ based on what you are installing. Running Aneka Cloud Management Studio on your machine requires little resources and a typical machine with 1 gigabyte of RAM would suffice. Your Master Container is responsible for scheduling jobs and keeping track of all the Workers in the network. It should ideally run on a powerful machine such as one with a dual-core processor and 3 gigabytes of RAM. Your Workers are responsible for executing jobs and a similarly capable machine would be ideal.

2.2 Software Requirements

Operating Systems: Aneka is built on Microsoft .Net Framework 2.0 and is capable of running on most Windows operating systems. Aneka is certified to run on Windows 2000, Windows XP, Windows 2003, Windows Vista, Windows 7 and Windows 2008.

ECMA Runtime Environment: Aneka has been designed to run on Microsoft .Net Framework 2.0 or the Mono Runtime 2.6 or higher. As such, Aneka is also capable of running on a number of Linux platforms.

Database: It is highly recommended that you use database persistence in a production environment. Memory persistence can be useful for evaluating and testing purposes when you need to get a Cloud up and running quickly. Aneka supports Microsoft SQL Server 2005, Microsoft SQL Server 2008 and MySQL 5.1+ for database persistence.

3 Installation

This section assumes that you have a copy of the Aneka distribution with you. If you do not have a copy already, you can download the latest version from Manjrasoft's [Website](#).

3.1 Installing Aneka Cloud Management Studio

Aneka installation begins with installing Aneka Cloud Management Studio. The Cloud Management Studio is your portal for creating, configuring and managing Aneka Clouds. Installing Aneka using the distributed Microsoft Installer Package (MSI) is a quick process involving three steps as described below.

3.1.1 Step 1 – Run the installer package to start the Setup Wizard

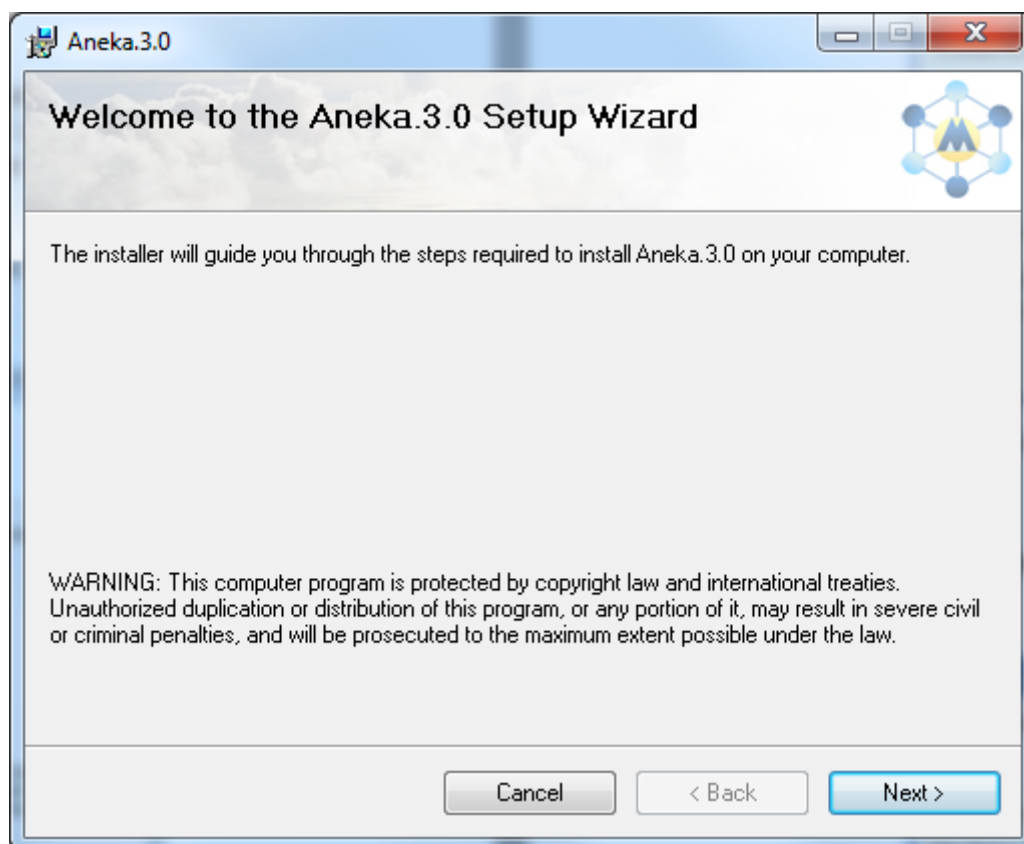


Figure 3 - Welcome Page

The Welcome Page is self-explanatory and you can proceed by clicking next.

3.1.2 Step 2 – Specifying the installation folder

In Step 2 you specify the installation folder. By default Aneka is installed in C:\Program Files\Manjrasoft\Aneka.3.0.

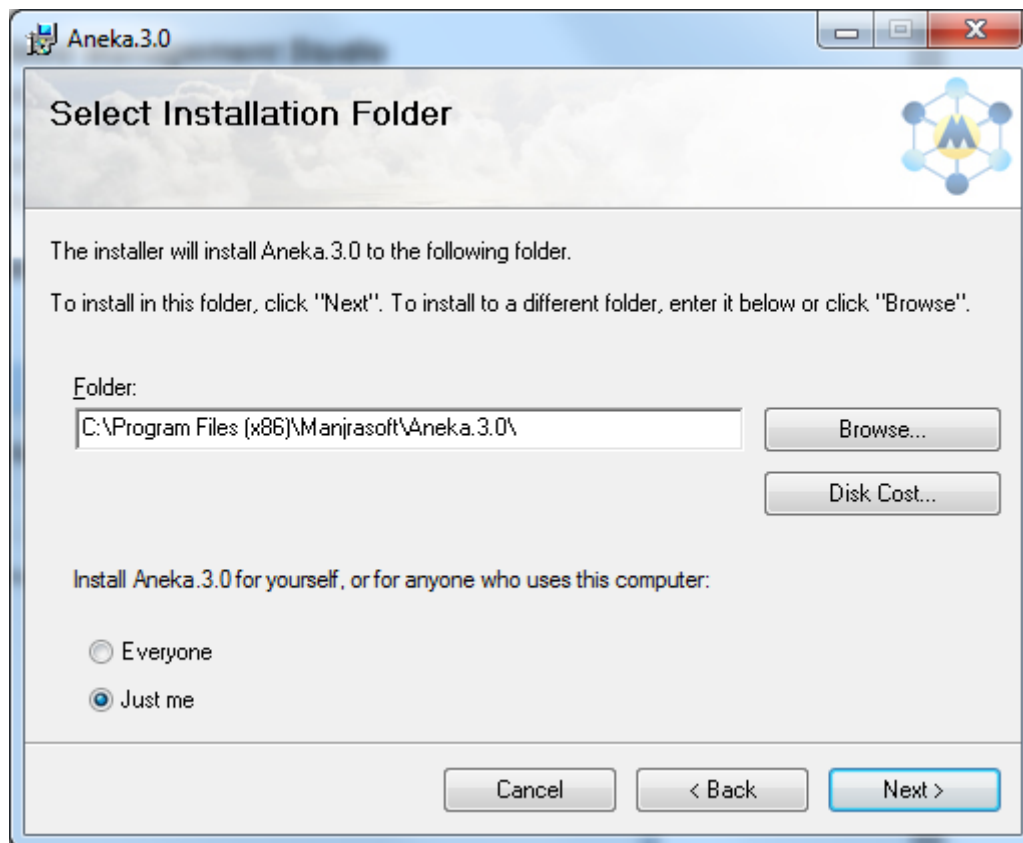


Figure 4 - Specifying the installation folder

3.1.3 Step 3 – Confirm and start the installation


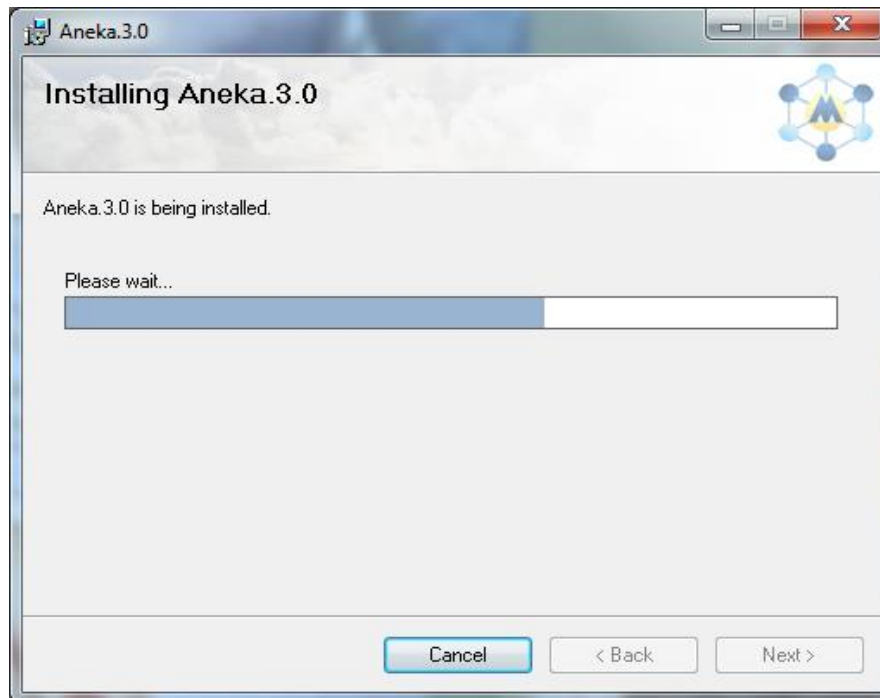
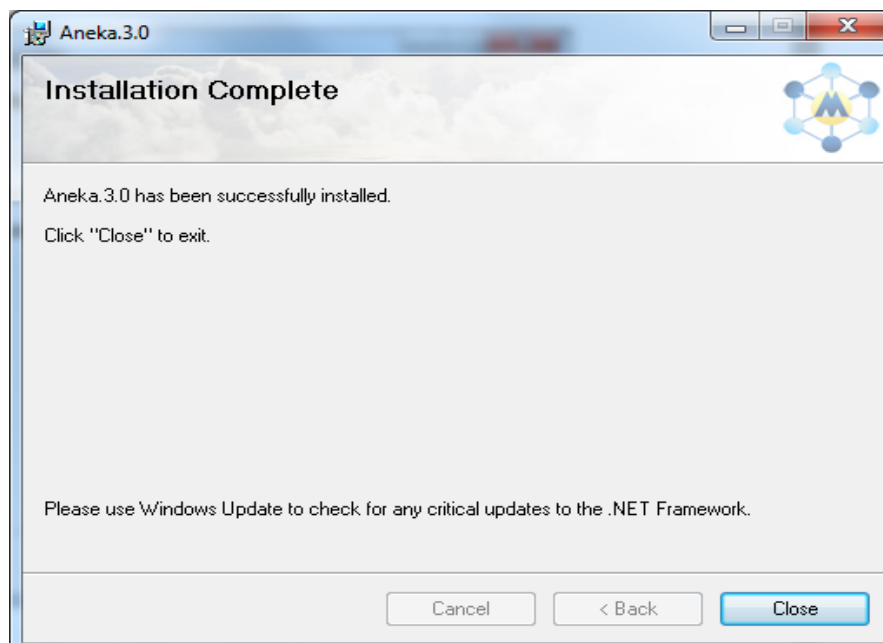


Figure 5 - Confirm Installation

At this point you are ready to begin the installation. Click “Next” to start the installation or “Back” to change your installation folder.

*Figure 6 - Installation Progress**Figure 7 - Installation Complete*

Once the installation is complete, close the wizard and launch Aneka Management Studio from the start menu.

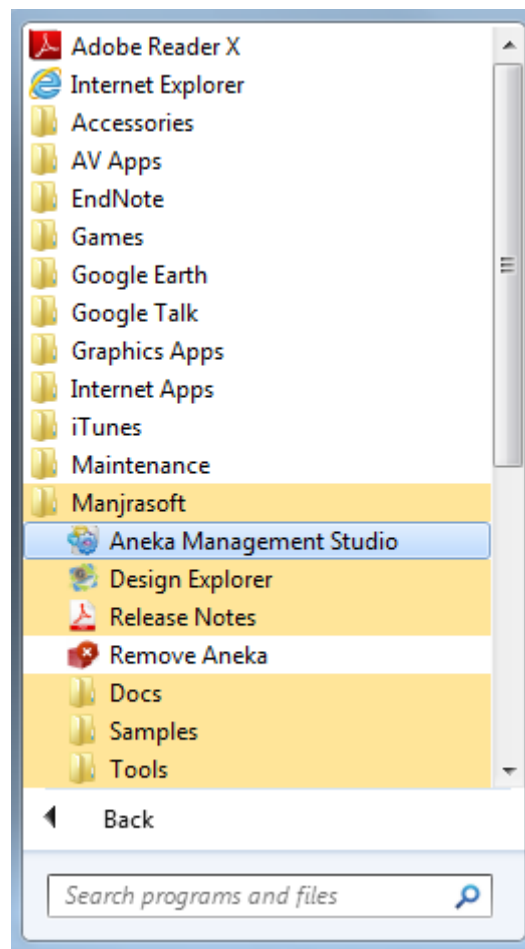


Figure 8 - Start Menu

4 Aneka Cloud Management Studio

The Aneka Cloud Management Studio is your portal for managing your infrastructure and clouds. It provides facilities for defining your underlying cloud infrastructure and creating one or more Aneka Clouds on top. It lets you create and manage Aneka user accounts, monitor the overall performance of your Cloud, obtain detailed reporting information on resource usage, data transfers, billing and application (job) execution. It also provides facilities for troubleshooting your deployments by allowing you to access and examine remote logs.

4.1 Starting up Management Studio

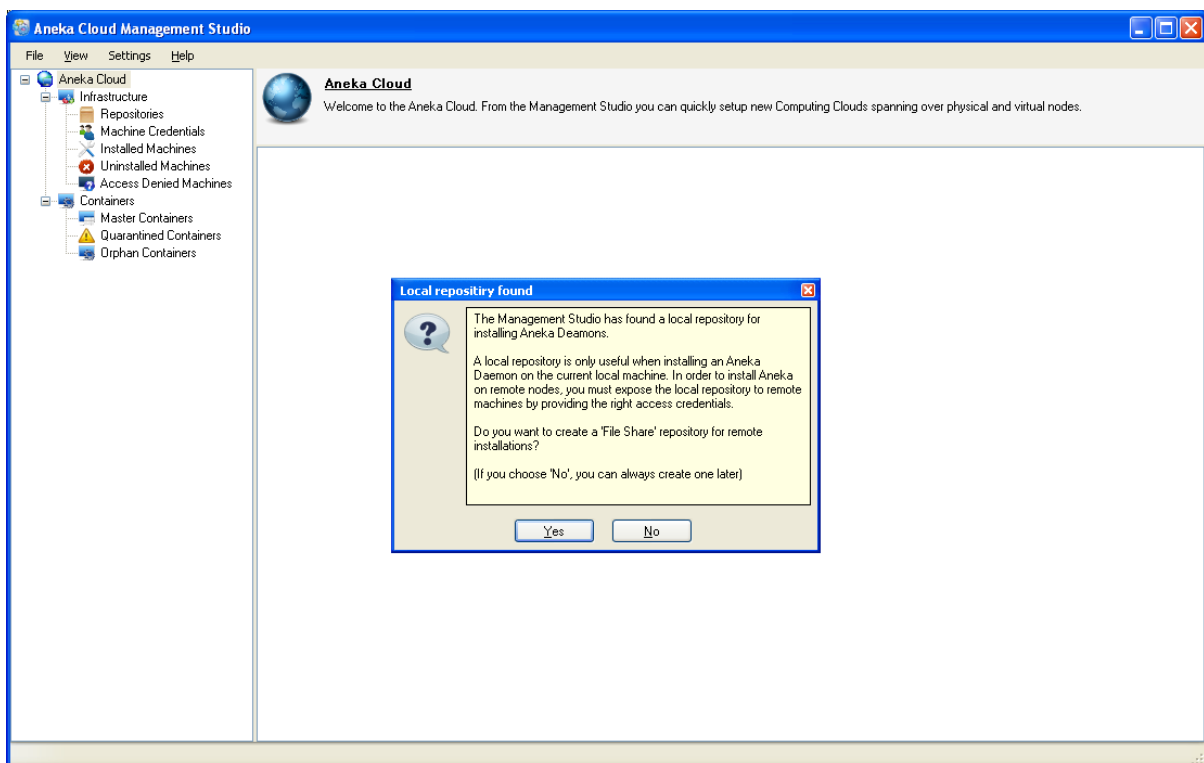


Figure 9 - Starting Aneka Cloud Management Studio for the first time.

When Aneka Cloud Management Studio is started up for the first time you'll be asked to create a *Remote Repository* for performing remote installations. Setting up a *Remote Repository* requires selecting a suitable repository type and supplying valid credentials which remote machines can use to connect and download required files. You may however choose to create this repository at a late time before making remote installations. If no repository is defined, you will be restricted to making local installations only.

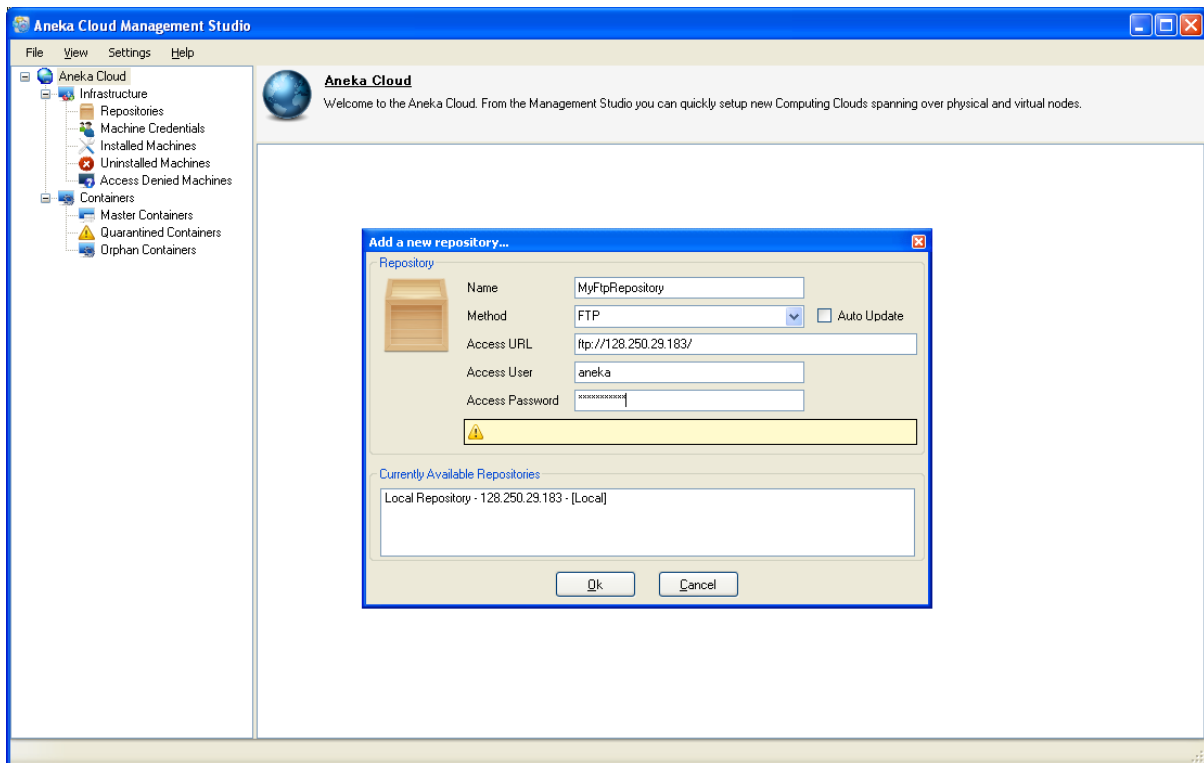


Figure 10 - Creating a repository for remote installations

If you choose to create a remote repository at this point, you will be restricted to using either FTP or Windows File Share. See section 5.3 on *Repository Management* for more details.

4.2 Shutting down Aneka Management Studio

When attempting to shut down Aneka Management Studio, you will be given the choice of saving all configuration data from the current session. It is *highly recommended* that you save this information and restore it the next time you start using the Management Studio.

4.2.1 The Configuration File

The configuration file, `ManagementStudio.config`, contains all information that describes your infrastructure, your Clouds, the machine credentials, repositories and authentication keys (see section on installing the Master Container) that you defined when using Aneka Management Studio. It is recommended that you save this information when you exit Management Studio so that you can restore it at a later session, and get up-to-speed with your Cloud management without having to redefine all settings again. Some configuration information, such as authentication keys, must be maintained safely if you are to add new Containers to your existing Cloud. Losing an authentication key however, is not detrimental as you will be able to reconfigure your clouds with a new key.

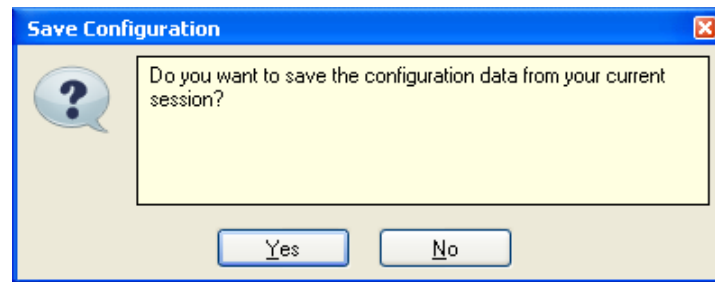


Figure 11 - Request to save configuration data when closing Management Studio

The configuration file is always encrypted before being written to disk for security reasons. When saving configuration data you will be required to specify a password as show in Figure 12.




Figure 12 - Password to encrypt configuration data

When starting up Aneka Management Studio at a later session, you will be given the choice of restoring your configuration data. If you choose to do so, you must re-enter the same password you used when saving.




Figure 13 - Request to restore configuration data

4.3 Cloud Explorer - Overview

The figure below shows the *Cloud Explorer* in Aneka Management Studio. The Cloud Explorer lets you manage your infrastructure (the physical or virtual machines in your network or Data Centre) and the Aneka Clouds that run on top of this infrastructure. It provides quick navigation for defining repositories, credentials, viewing the status of the

machines in your infrastructure, and the Containers that makeup Aneka Clouds. The sections below describe the infrastructure and cloud management features in Aneka in greater detail.

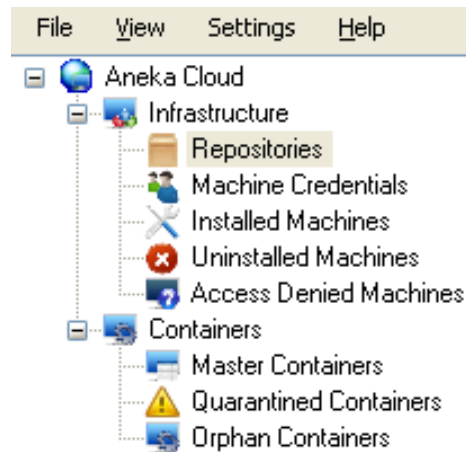


Figure 14 - Cloud Explorer

5 Infrastructure Management

Your infrastructure is composed of the physical and virtual machines in your Local Area Network or Data Centre. Before creating an Aneka Cloud you must define your infrastructure for hosting the Cloud. The Aneka Cloud Management Studio provides facilities for defining and managing your infrastructure.

5.1 Adding a new Machine




Figure 15 - Adding a Machine

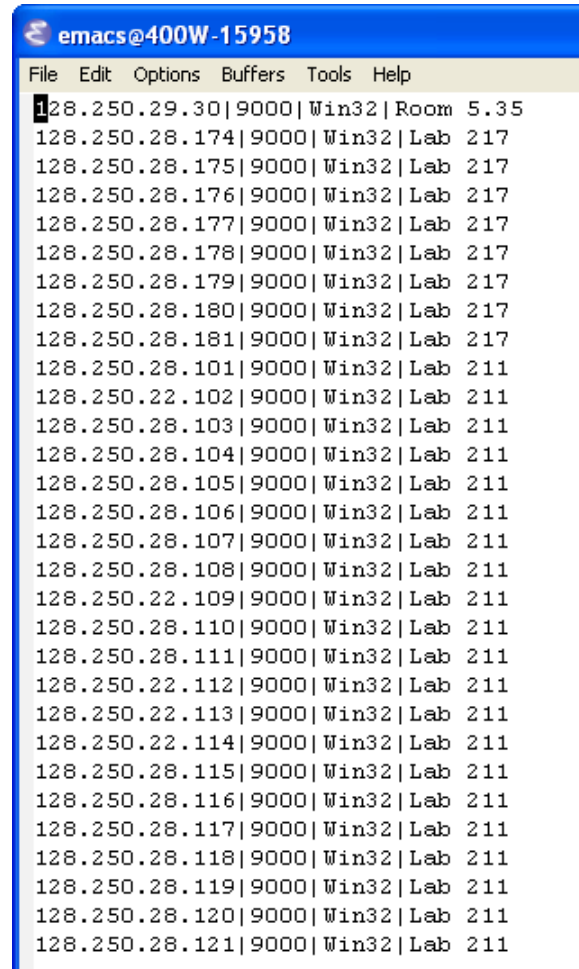
The simplest way to represent your infrastructure in Management Studio is to add each machine individually. You do this by selecting the *Add Machine* option in the *File* menu. This approach is simple enough if you have a few machines but can be cumbersome when dealing with a large network or Data Centre. You may alternatively create and import a *machine list* file as described below.

5.2 Importing and Exporting Machine Lists

A machine list file provides a convenient way to import a number of machines that define your infrastructure. Each machine is represented by a single line in the file in the following form:

```
Machine [IP Address or Name]|Aneka Daemon Port|Platform|Label
```

Figure 16 below shows a sample machine list file, all of which are Windows platforms (currently the only supported platform by the Aneka Management Studio) and labelled according to their location in the network. You may choose to use any label you wish to describe your machines in a meaningful way. The ports refer to the Aneka Daemon port (whether a daemon has been installed or not). Once you have created your machine list file you can import it by selecting *Import Machine List* from the *File* menu. You may also make changes to your infrastructure by adding more machines or removing existing ones in Management Studio, and export a new machine list file for later use.



```

128.250.29.30|9000|Win32|Room 5.35
128.250.28.174|9000|Win32|Lab 217
128.250.28.175|9000|Win32|Lab 217
128.250.28.176|9000|Win32|Lab 217
128.250.28.177|9000|Win32|Lab 217
128.250.28.178|9000|Win32|Lab 217
128.250.28.179|9000|Win32|Lab 217
128.250.28.180|9000|Win32|Lab 217
128.250.28.181|9000|Win32|Lab 217
128.250.28.101|9000|Win32|Lab 211
128.250.22.102|9000|Win32|Lab 211
128.250.28.103|9000|Win32|Lab 211
128.250.28.104|9000|Win32|Lab 211
128.250.28.105|9000|Win32|Lab 211
128.250.28.106|9000|Win32|Lab 211
128.250.28.107|9000|Win32|Lab 211
128.250.28.108|9000|Win32|Lab 211
128.250.22.109|9000|Win32|Lab 211
128.250.28.110|9000|Win32|Lab 211
128.250.28.111|9000|Win32|Lab 211
128.250.22.112|9000|Win32|Lab 211
128.250.22.113|9000|Win32|Lab 211
128.250.22.114|9000|Win32|Lab 211
128.250.28.115|9000|Win32|Lab 211
128.250.28.116|9000|Win32|Lab 211
128.250.28.117|9000|Win32|Lab 211
128.250.28.118|9000|Win32|Lab 211
128.250.28.119|9000|Win32|Lab 211
128.250.28.120|9000|Win32|Lab 211
128.250.28.121|9000|Win32|Lab 211

```

Figure 16 - A sample Machine List file

5.3 Repository Management

An Aneka Repository is a location containing all required binaries for installing Aneka Daemons and Containers. When the Aneka Daemon is installed on a machine, the binaries for both the Daemon and the Container are downloaded from the selected repository to a local repository folder in the target machine and installed from there.

When the Aneka Management Studio is installed on a machine for the first time, a *Local Repository* is created by default. The local repository is not accessible over the network and can only be used to perform local installations. That is, for installing the Aneka Daemon and Containers on the same machine in which the Management Studio is installed.

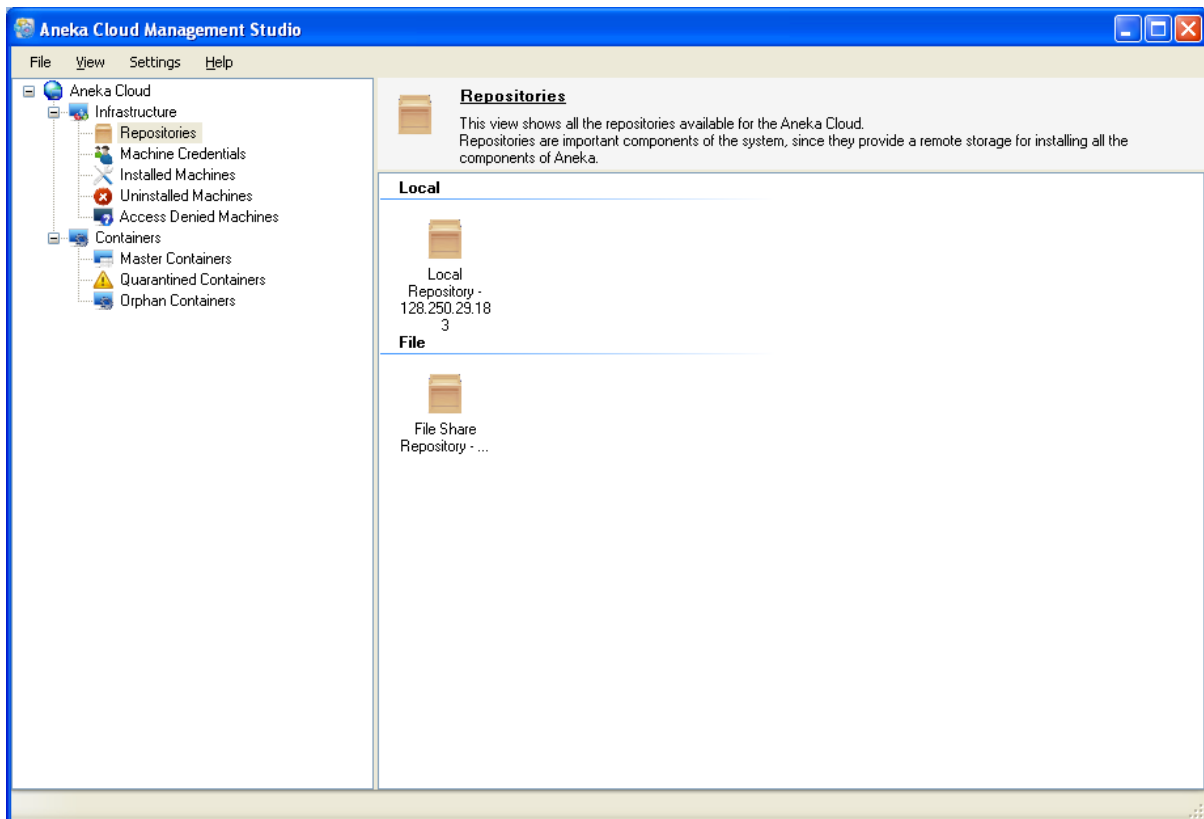


Figure 17 - List of Aneka Repositories for installing Aneka Daemons and Containers

As illustrated in Figure 9 in the previous section, when the Management Studio is started for the first time you are given the opportunity to create a *Remote Repository* for performing remote installations. If you chose *No* you may create one now if you intend to install Aneka on remote machines.

Creating a new repository is easy and can be done by selecting *Add* from the context menu for *Repositories*. The figure below shows the dialog for adding a new repository. The *Name* field refers to what you would like to call your new repository. *Method* refers to the mode of access to your repository. The three supported methods include *Local*, *Windows File Share*, and *FTP*. The *Access URL* refers the location of your repository. In the case of a *Local Repository*, this would typically refer to *Repository* folder in your Aneka installation directory. In the case of a *Windows File Share Repository* the URL must be a network accessible path, in the following form:

Format \\<host>\<drive letter>\$\<path>

Example \\marg\C\$\Aneka\Repository

In the case of an FTP repository, the URL is the path to a repository folder in the FTP server. Depending on the *Method* you may also be required to specify a valid access credential that remote machines can use to connect and download the required binaries.

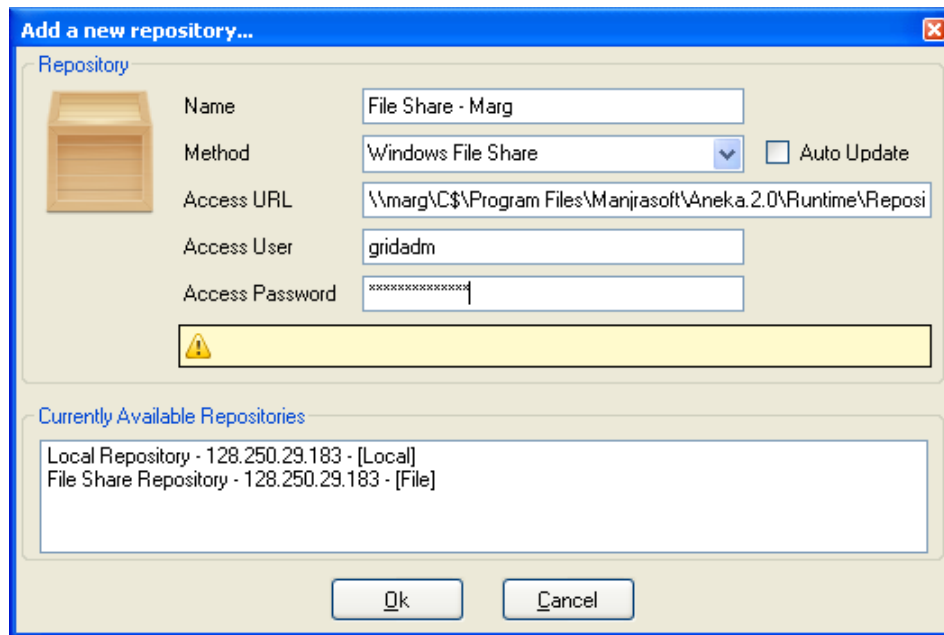


Figure 18 - Creating a new File Share Repository

NOTE: We discourage the use of a Windows File Share Repository when installing Aneka on a large number of nodes. Remote machines connect to the Windows File Share Repository using an IPC\$ session. As Windows limits the number of concurrent IPC\$ sessions, connections from some of the remote machines may fail. You may alternatively install Aneka on the remote machines in smaller batches depending on the limit set on your Windows machine hosting the repository. We recommend however, that you use an FTP repository which poses no such limitations.

When you accept your repository settings by pressing *Ok*, the Management Studio will begin preparing your *Remote Repository*. This involves copying the binaries from your local repository in the installation folder to the remote location. This may take a few minutes.

5.4 Machine Credentials

The Windows *Machine Credentials* are the credentials you use to access and manage your remote machines. These include activities such as installing the Aneka Daemon on remote machines, and starting, stopping and uninstalling the Daemon service. These credentials must thus have administrative privileges on the target machine.

When you start up Management Studio for the first time, there will be no credentials defined for you by default. Before you starting managing your infrastructure you must define one or more credentials. To create a new credential select *Add* from the context menu for *Machine Credentials*. This will bring up a dialog requesting the user name and password for the credential as show in Figure 20.

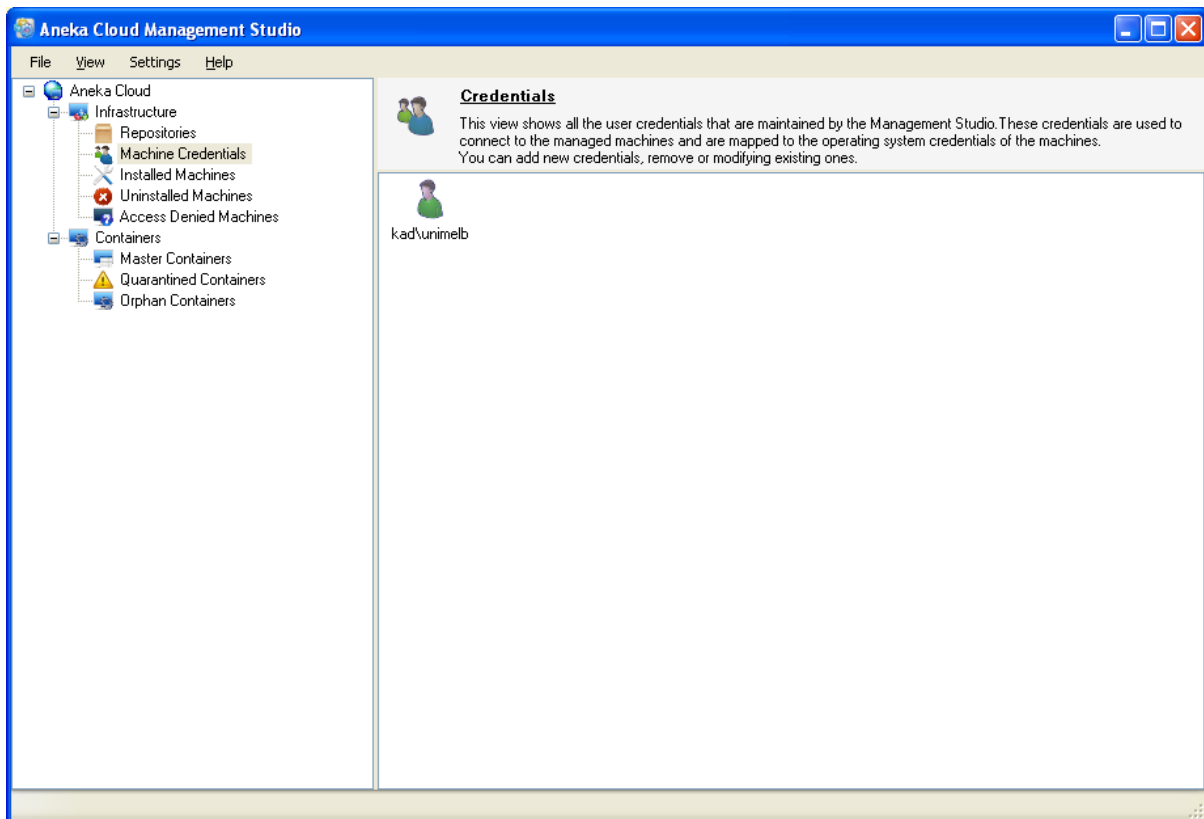


Figure 19 - A list of Machine Credentials

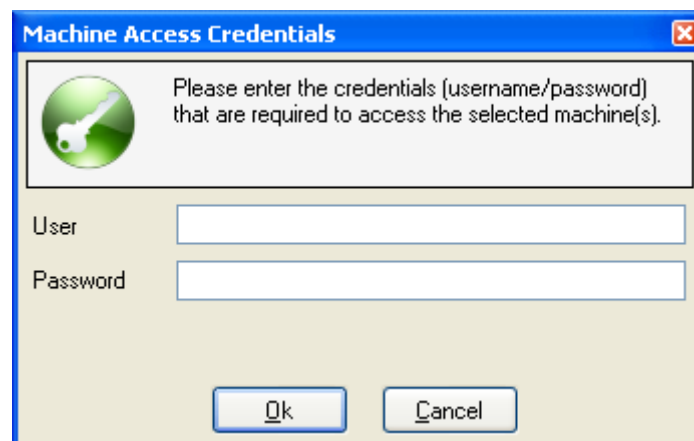


Figure 20 - Creating a new credential

As noted earlier, the user account must have administrative rights to the target machine, and can either be a local machine account or a domain account. When specifying a domain account you must also prefix the domain name to the user name, for example: *mydomain\anekaadmin*.

5.5 Access Denied Machines

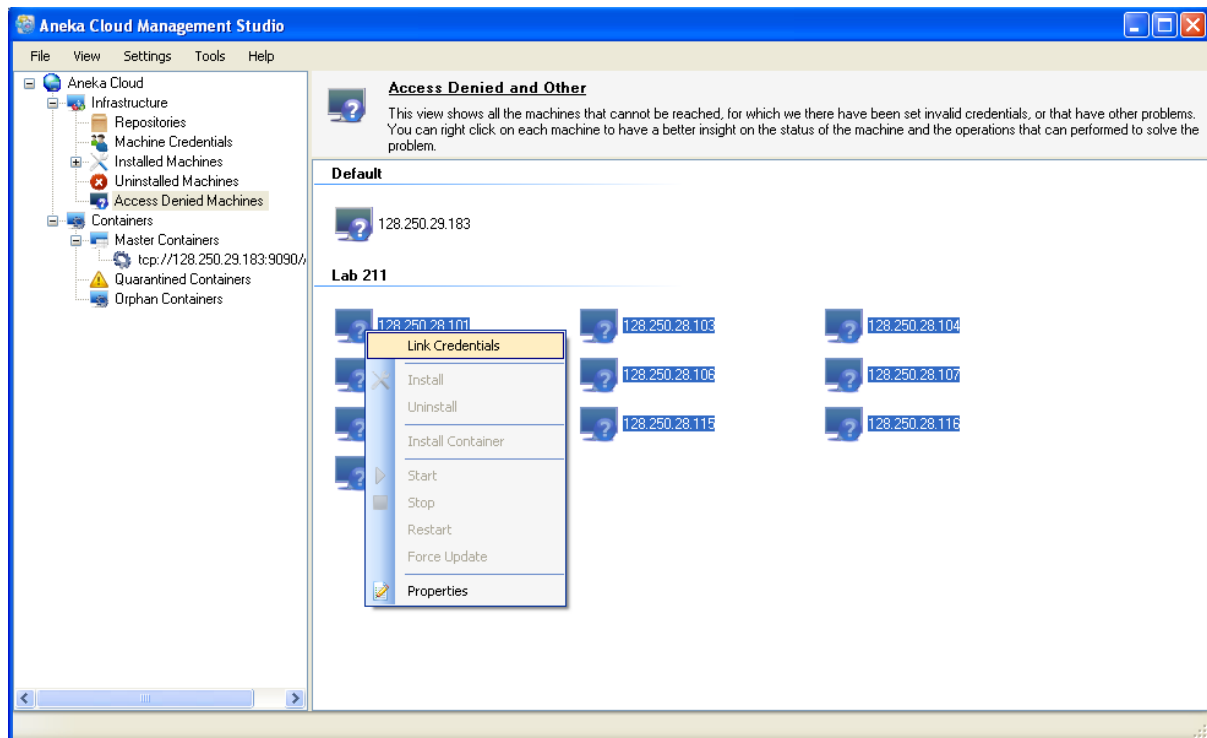


Figure 21 - Access Denied Machines

When you add or import a list of machines for the first time, they will be listed under *Access Denied Machines* as shown above. This is because the Management Studio does not know what credentials to use to connect to the remote machines. In order to associate a *Machine Credential*, select all relevant machines and then select *Link Credentials* from the context menu. This pops up a dialog allowing you to select one of the credentials that you defined earlier, or create a new one if required as shown below.

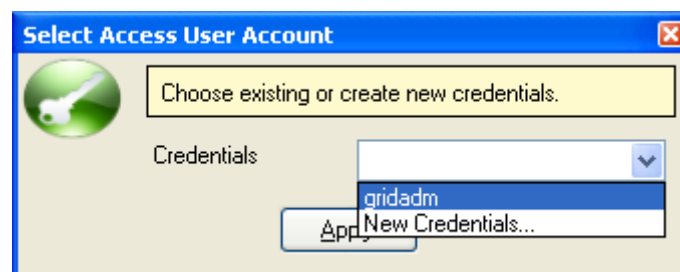


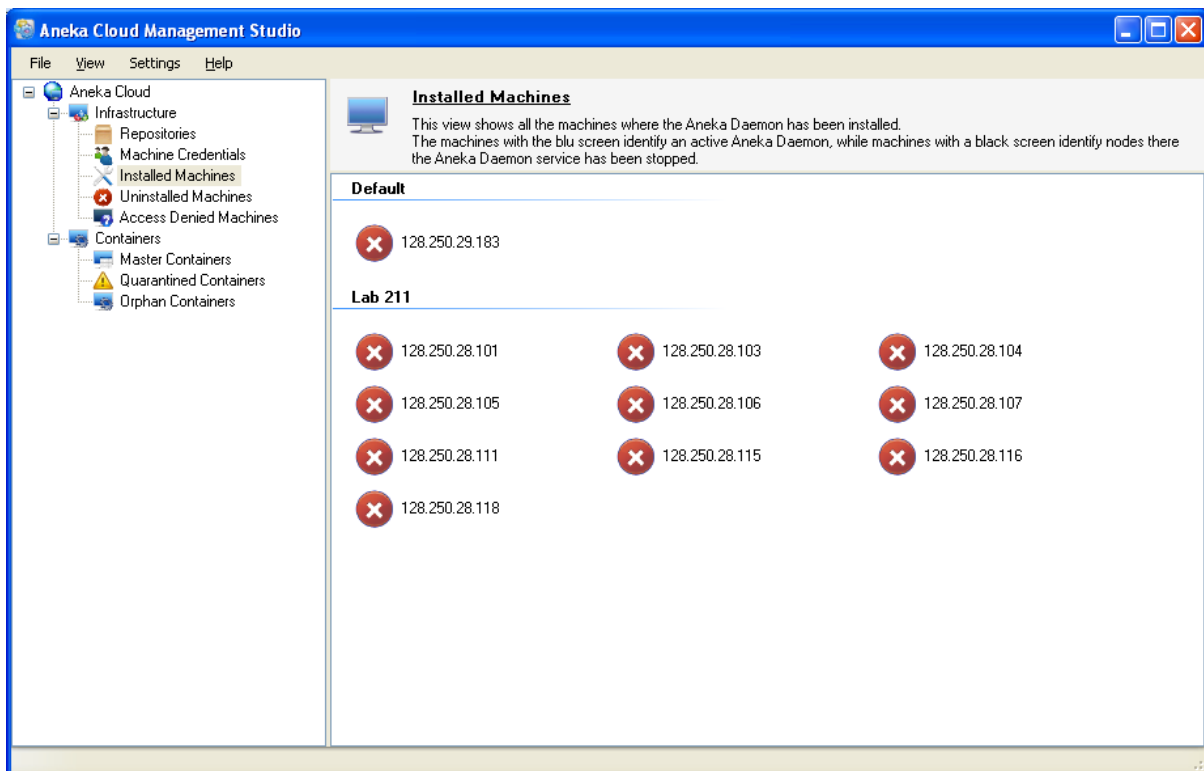
Figure 22 - Selecting a Credential

Lab 211*Figure 23 - Network Not Accessible*

Note that after linking credentials you may occasionally encounter situations where a machine is still not accessible via the network, either because there is a problem with the network or the machine is down. Figure 23 shows a list of machines that are not accessible. In these situations there is no choice but manually troubleshoot the problem outside Management Studio.

5.6 Uninstalled Machines

An *Uninstalled Machine* is one which has no Aneka Daemon running. This implies that either Aneka was never installed on the machine or a previous installation was removed. The figure below shows a list of *Uninstalled Machines* after having linked a valid *Machine Credential*. It is important to realize that the Management Studio can only determine the status of a machine after linking to a valid credential. To install Aneka Daemons on *Uninstalled Machines* see section below on *Installing the Aneka Daemon*.

*Figure 24 - Uninstalled Machines*

5.7 Installed Machines

An *Installed Machine* is one which has an instance of the Aneka Daemon running. An *Installed Machine* may also have one or more Container instances running. You may perform a number of different activities on an *Installed Machine* such as starting and stopping the Aneka Daemon service, installing a Container, updating the binaries, viewing the machine's properties, or uninstalling the Aneka Daemon (and Containers) from the machine.

Once you have a collection on *Installed Machines* at your disposal, you are ready to create your *Aneka Cloud*.

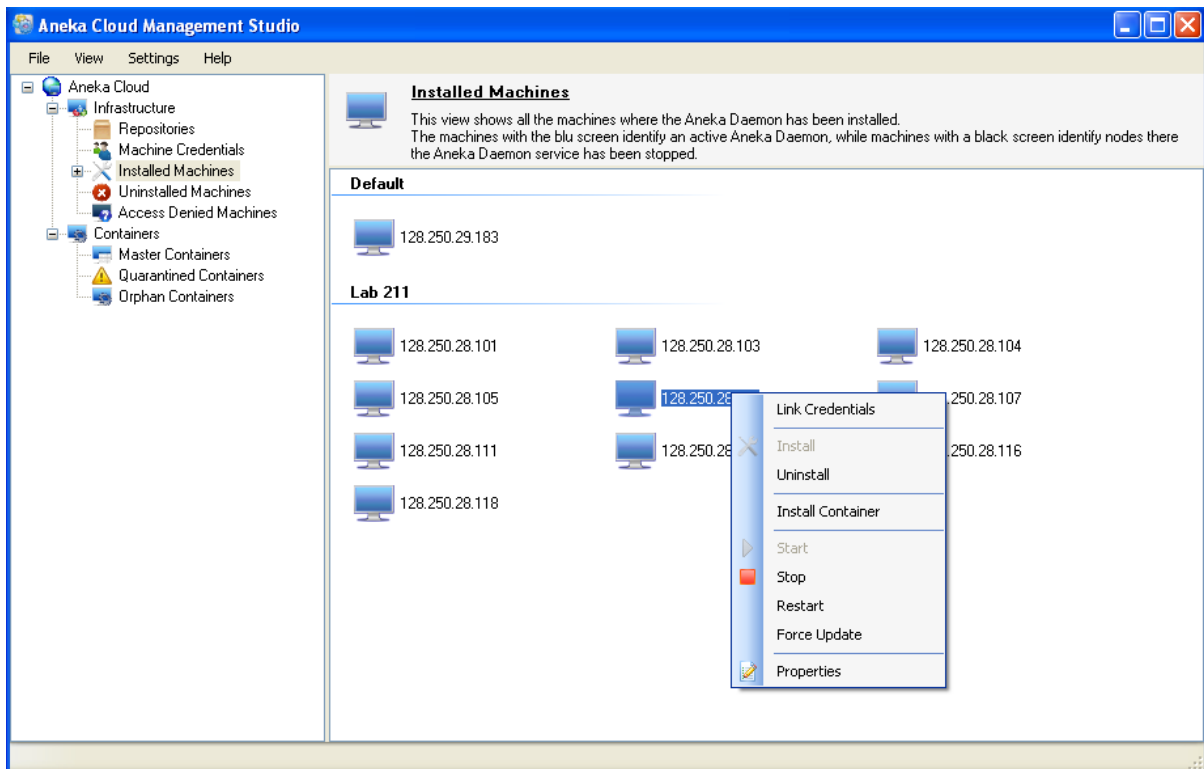


Figure 25 - Installed Machines

6 Installing the Aneka Daemon

The previous sections looked at the process of defining your infrastructure for creating Aneka Clouds. We looked at adding and importing machines, creating *Aneka Repositories* and *Machine Credentials* for performing remote installations, and the various states your machines can be in (*Access Denied*, *Uninstalled*, *Installed*).

This section will introduce you to the process of installing Aneka Daemons, the prerequisite for creating Aneka Clouds. The Aneka Daemon is responsible for managing Aneka Containers that make up your cloud. This includes installing new Containers, starting, stopping and uninstalling Containers. Aneka Daemons provide the underlying management infrastructure for Aneka Clouds. In order to install the Aneka Daemon, select the list of *Uninstalled Machines* that you wish to install the Daemon on, and then select *Install* from the context menu as shown in the figure below. This brings up the *Aneka Daemon Installation Wizard* as shown in the Figure 27.

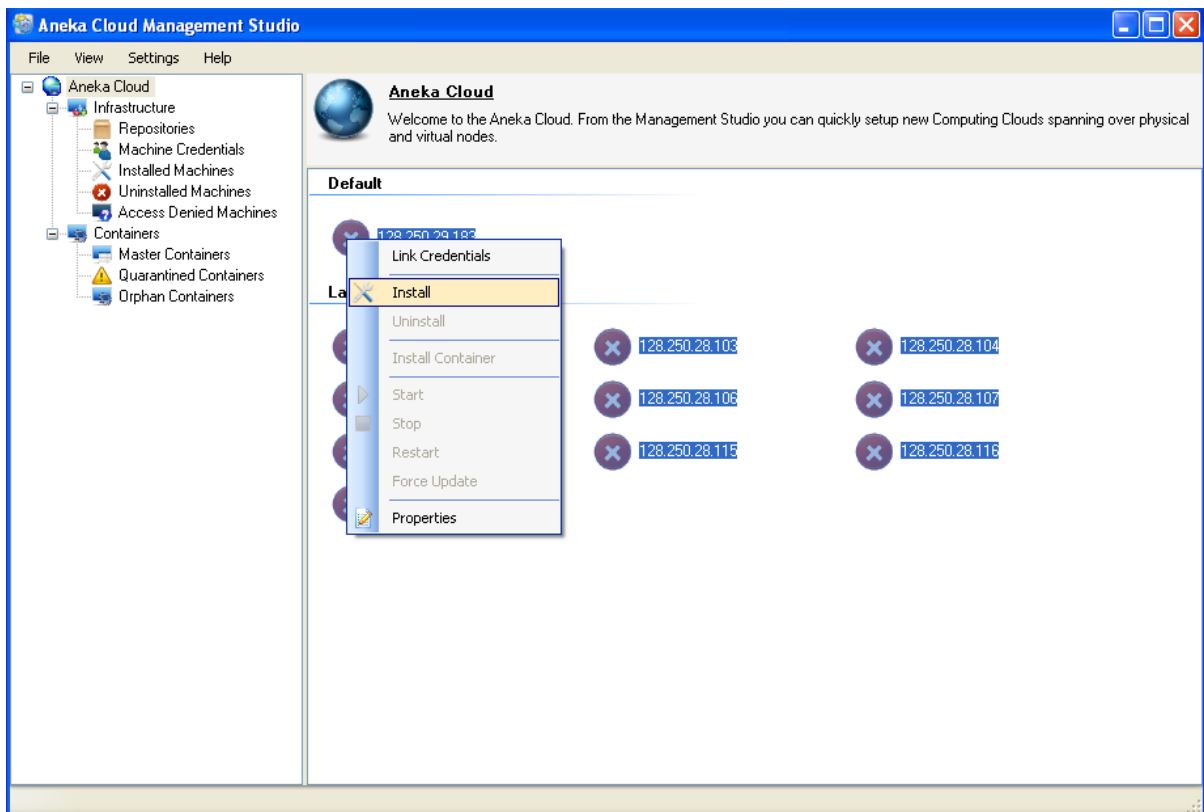


Figure 26 - Installing the Aneka Daemon on a group of machines

6.1 Step 1 – Repository and Port Configuration

The Wizard will automatically probe all selected machines on start up to find if the default port, 9000, is free. If the port in any one of the machines is not free, as a result of being used by another program, you must either ensure that it becomes free, or change

the port number for the Aneka Daemon. If you wish to verify whether your new port is free on all target machines, press the *Probe* button again.

After selecting a valid port, you must select a suitable repository to install the Daemon from. As mentioned previously, a *Local Repository* is only valid for performing local installations. When performing remote installations you must select one of the valid *File Share Repositories* you defined earlier. If all settings are correct you may now click *Next* to view a summary of the installation. If either your port or repository selection is incorrect the Wizard will inform you of this and prevent you from proceeding to the next step.

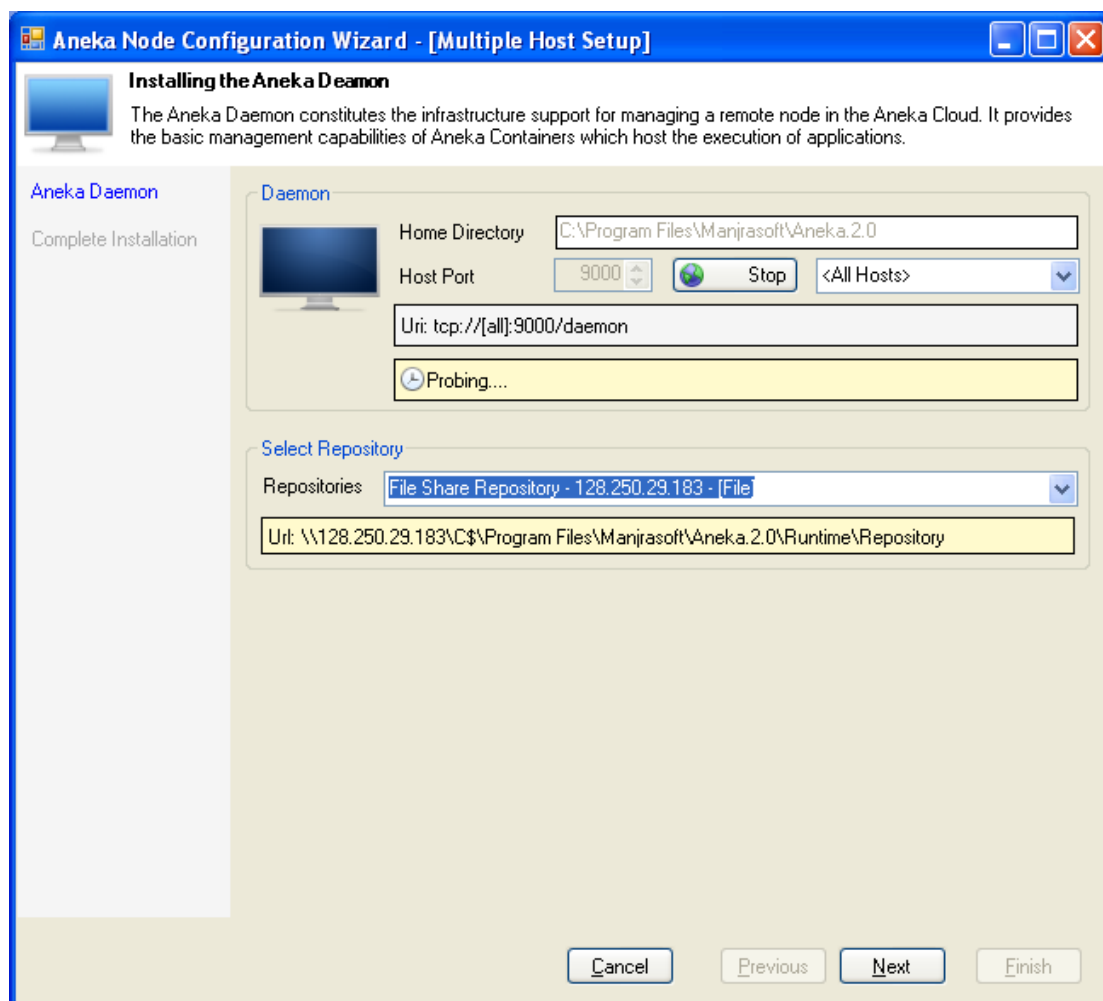


Figure 27 - Aneka Daemon Installation Wizard

6.2 Step 2 – Review Summary and Install

Figure 28 shows a summary of the settings prior to installing the Daemon. If you are satisfied with the installation settings you may start the installation by pressing *Finish*. To change the installation settings press *Previous*.

When you press *Finish* the Wizard will begin installing the Aneka Daemon on all of the selected machines. This process might take a few minutes depending on the number of

machines selected. When the Daemons are installed on the target machines they are automatically started up as shown in Figure 29. You might occasionally notice some of the machines being marked with a *warning* icon. This is because the Management Studio is unable to determine the status of the Aneka Daemon on those machines, the most common reason being that the Daemon service might have not finished starting up. Select *Refresh* in the *View* menu after completing the installation in order to see the proper status as shown in the figures below.

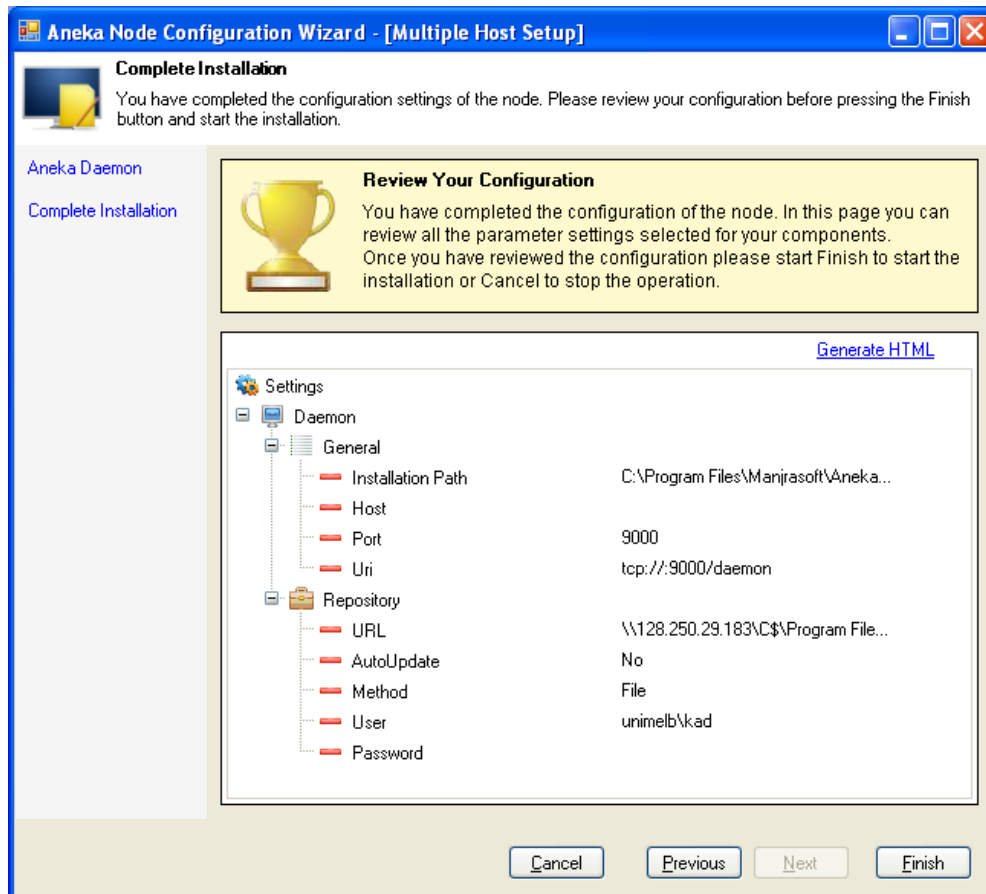


Figure 28 - A summary of the Aneka Daemon installation settings

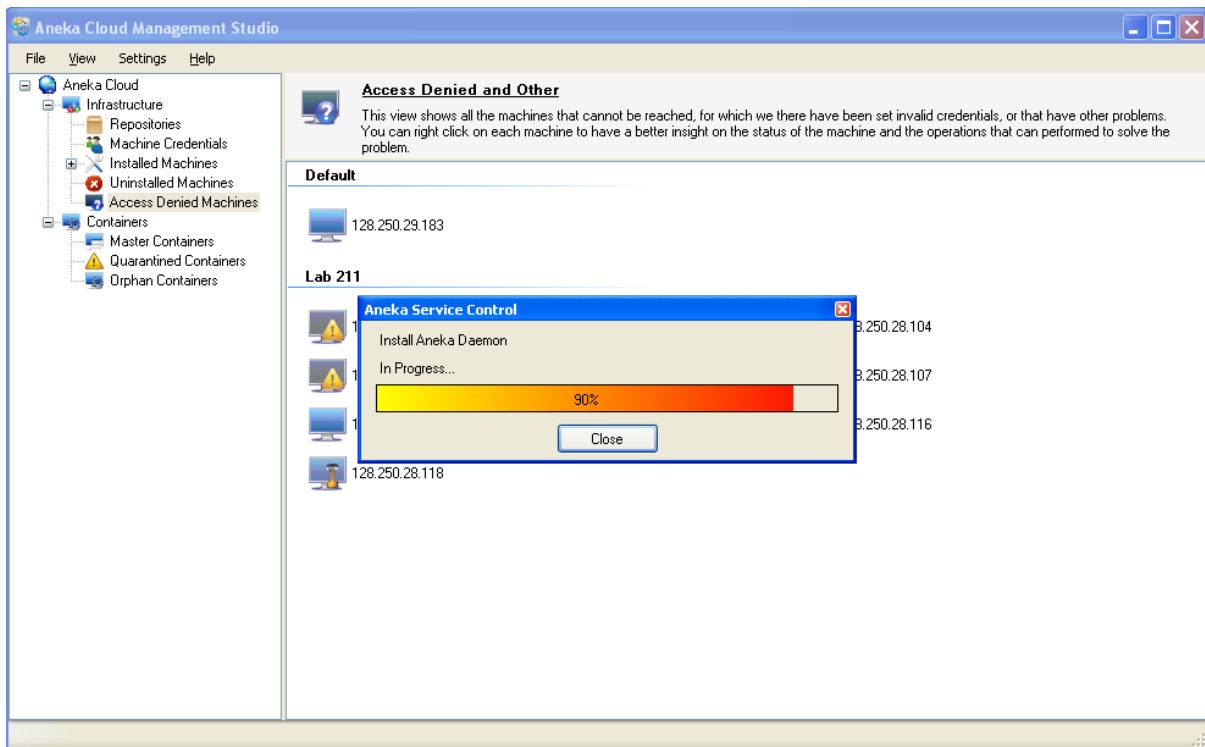


Figure 29 - Installation progress of the Aneka Daemon on the selected machines

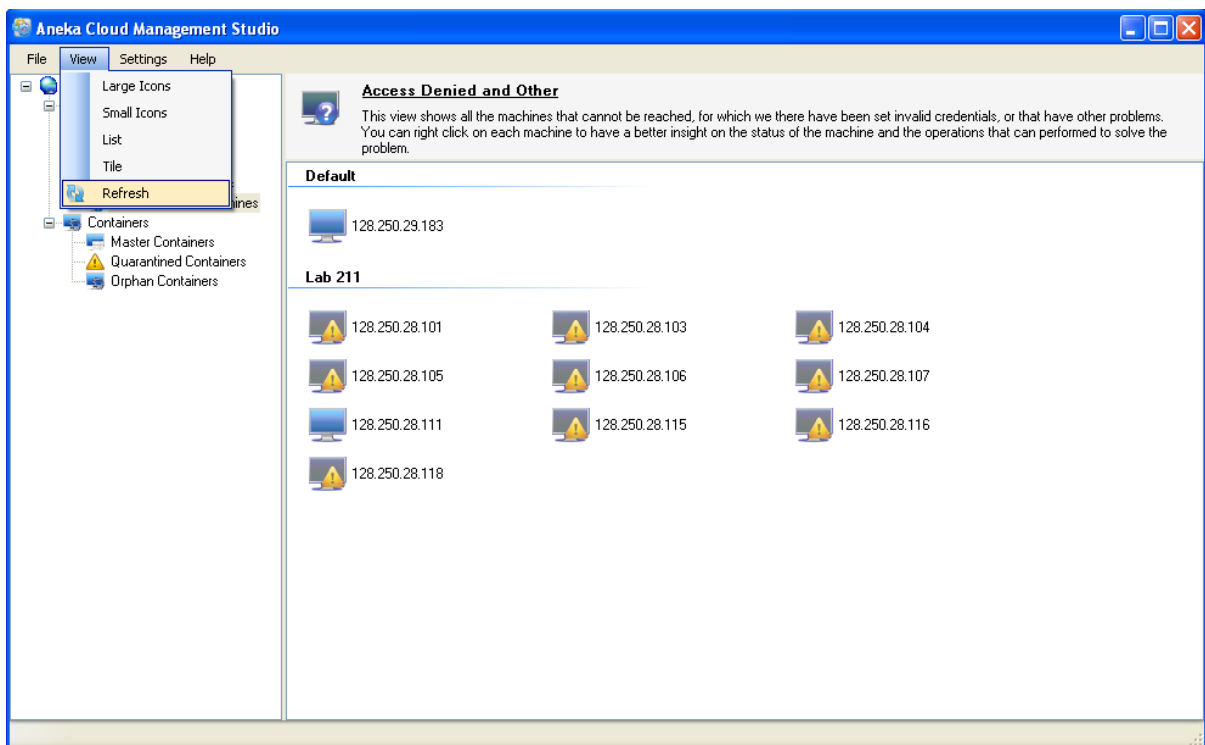


Figure 30 - Refreshing the status of the machines after installation

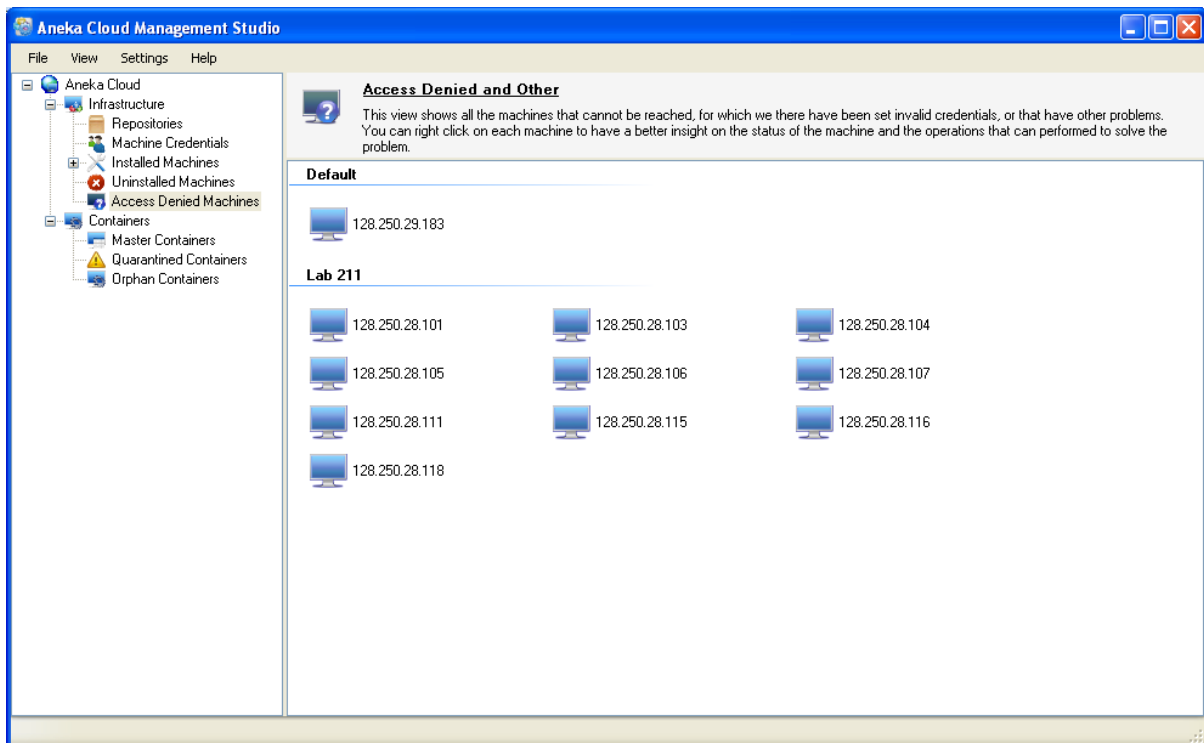


Figure 31 - Installed Machines with the Aneka Daemon running

7 Creating an Aneka Cloud

Once you have the Aneka Daemon running on all selected machines you are now ready to create an Aneka Cloud. As described earlier, an Aneka Cloud is composed of a *Master Container* and group of *Worker Containers*. All Workers are registered to the Master, which acts as a gateway to the Cloud. End-users submit their applications, composed of a number of jobs, to the Master which in turn schedules them to Workers.

To begin creating your cloud, you must first decide on the machine that will host the *Master Container*. Ideally, this must be a machine capable of processing requests from a number of clients while also managing a large number of Workers. The capability of this machine will of course depend on the size of your cloud and the expected usage. To install the Master Container, select *Install Container* from the context menu on the selected machine, as show in the figure below. This brings up the Container installation Wizard.

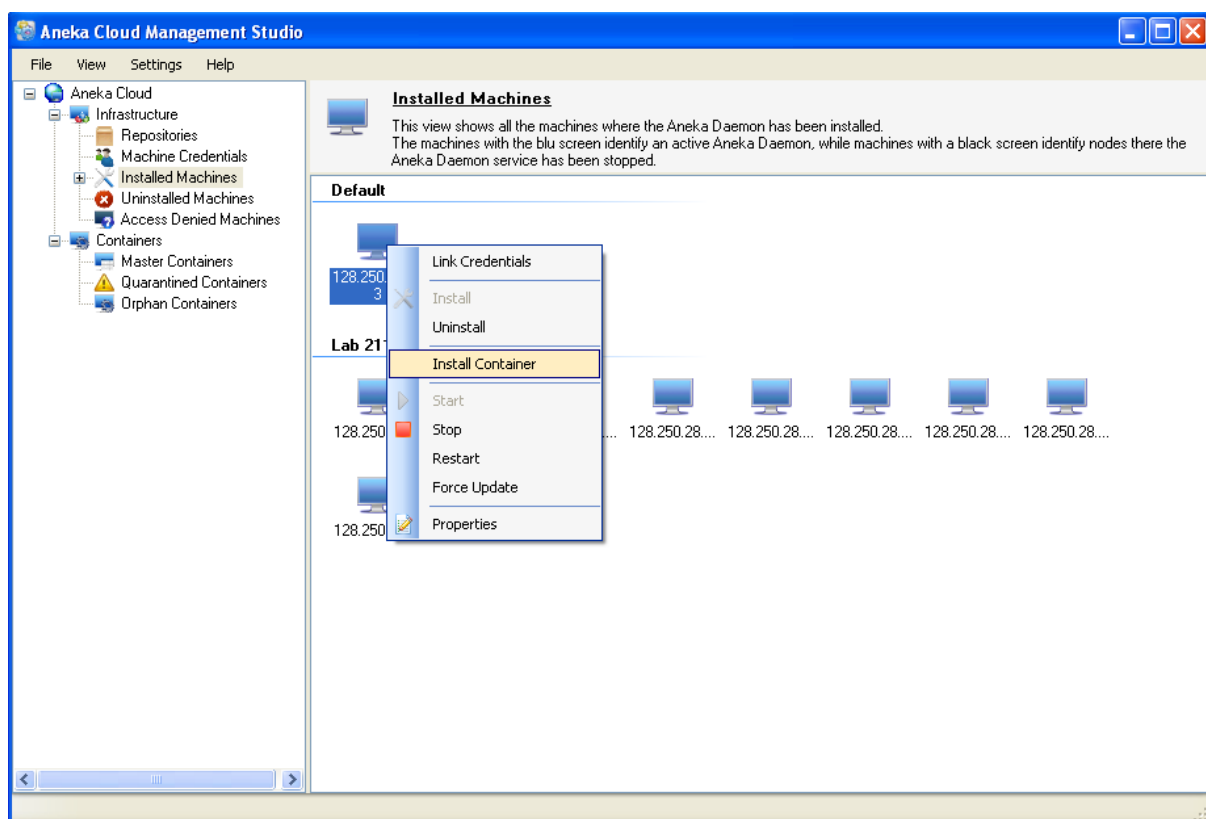


Figure 32 - Installing the Master Container on the Master Node

7.1 Installing the Master Container

7.1.1 Step 1 – Container Type and Port Configuration

The Wizard will automatically probe the default Container port, 9090, on start up in order to verify that it is free. If the port is not free, as a result of being used by another program, you must either ensure that it becomes free, or assign a new port number for the Container. If you wish to verify whether your new port is free, press the Probe button again. Make sure that the selected *Container Type* is *Master* and click *Next* to proceed.




Figure 33 - The Container Installation Wizard for installing the Master

NOTE: If the port is being occupied by an existing Container on the target machine, the Wizard will notify you of its existence and ask if you want to reconfigure it. If your intention is to install a new container, you must change the Container port. On the other hand, if you intend to reconfigure the existing Container you may proceed. To reiterate, the Container Installation Wizard can either be used to install a new Container or reconfigure an existing Container.

7.1.2 Step 2 – Security Configuration

This step helps you setup security for your cloud. It is highly recommended that you enable security in a production environment. For testing or evaluation purposes however, you may choose to disable security. Enabling secure communication ensures that all users

are required to supply a valid user name and password to use the Aneka Cloud. It also ensures that the Master and Workers authenticate each other and all traffic exchanged between them is encrypted. In order to create Aneka user accounts see section below on creating Aneka users.

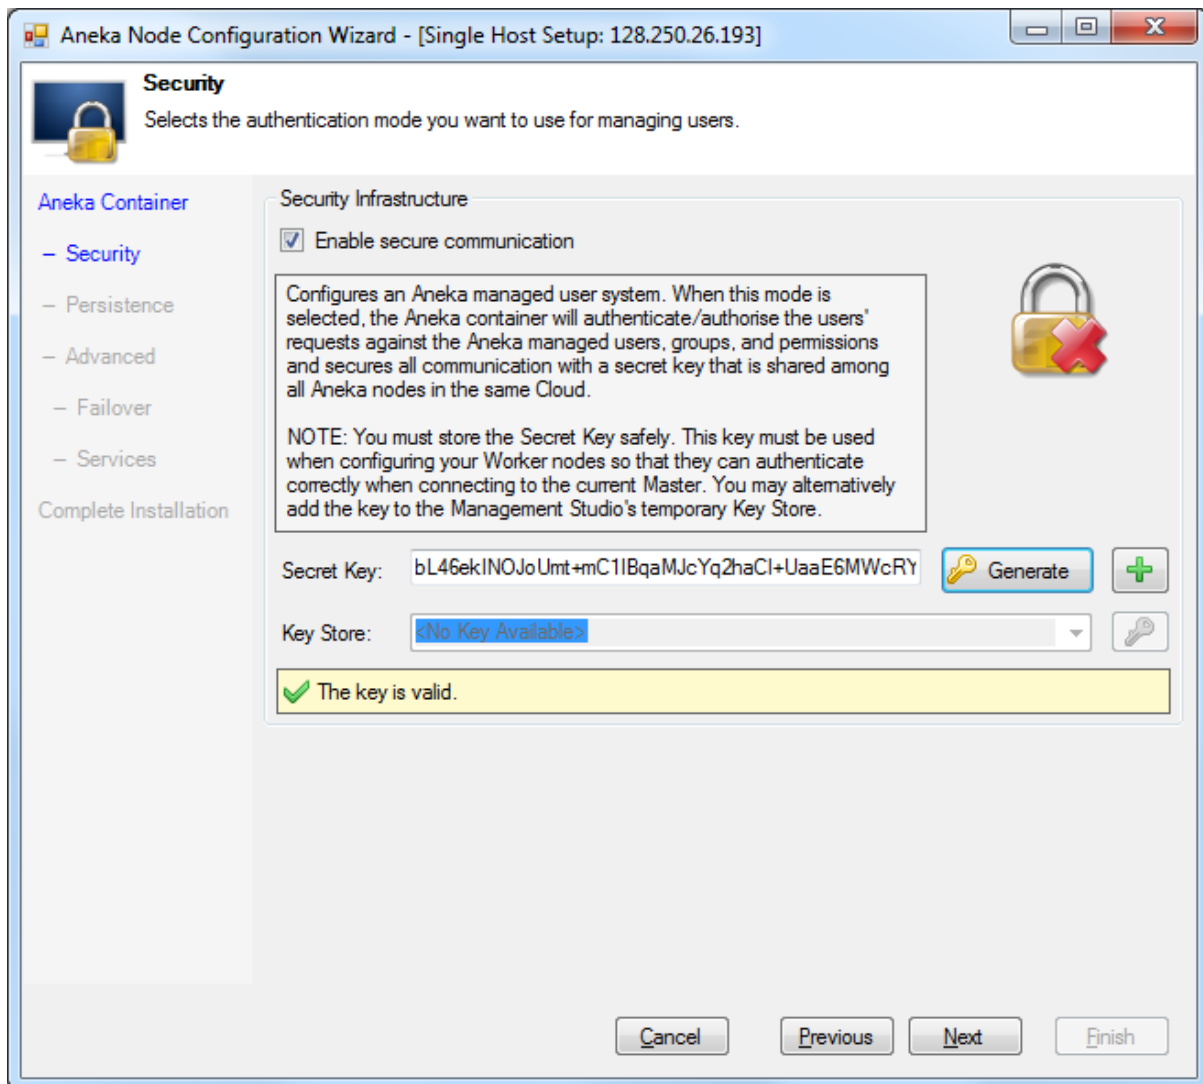


Figure 34 - Configuring security for your cloud

The *Secret Key* is a shared security key between the Master and Worker Containers for authentication and encryption. It is mandatory that you generate a key if you enable secure communication. It is also recommended that you add this key to the *Key Store* with a friendly name, as you will need it later when configuring a Worker Container. The *Key Store* is written to the Management Studio's configuration file and is restored during later sessions. (See The Configuration File discussed previously). Click *Next* to proceed.

7.1.3 Step 3 – Persistence Configuration

An Aneka Cloud maintains data for different purposes including user management, application and job management, resource management, accounting, billing and

monitoring to name a few. This step lets you setup the persistence requirements for your cloud.

There are two *Persistence Modes* available to choose from, *Relational Database* and *Memory*. It is highly recommended that you choose *Relational Database* for production environments. Memory persistence is typically useful for testing or evaluation purposes, and will store all data in volatile memory which will be lost when the Container restarts or crashes.

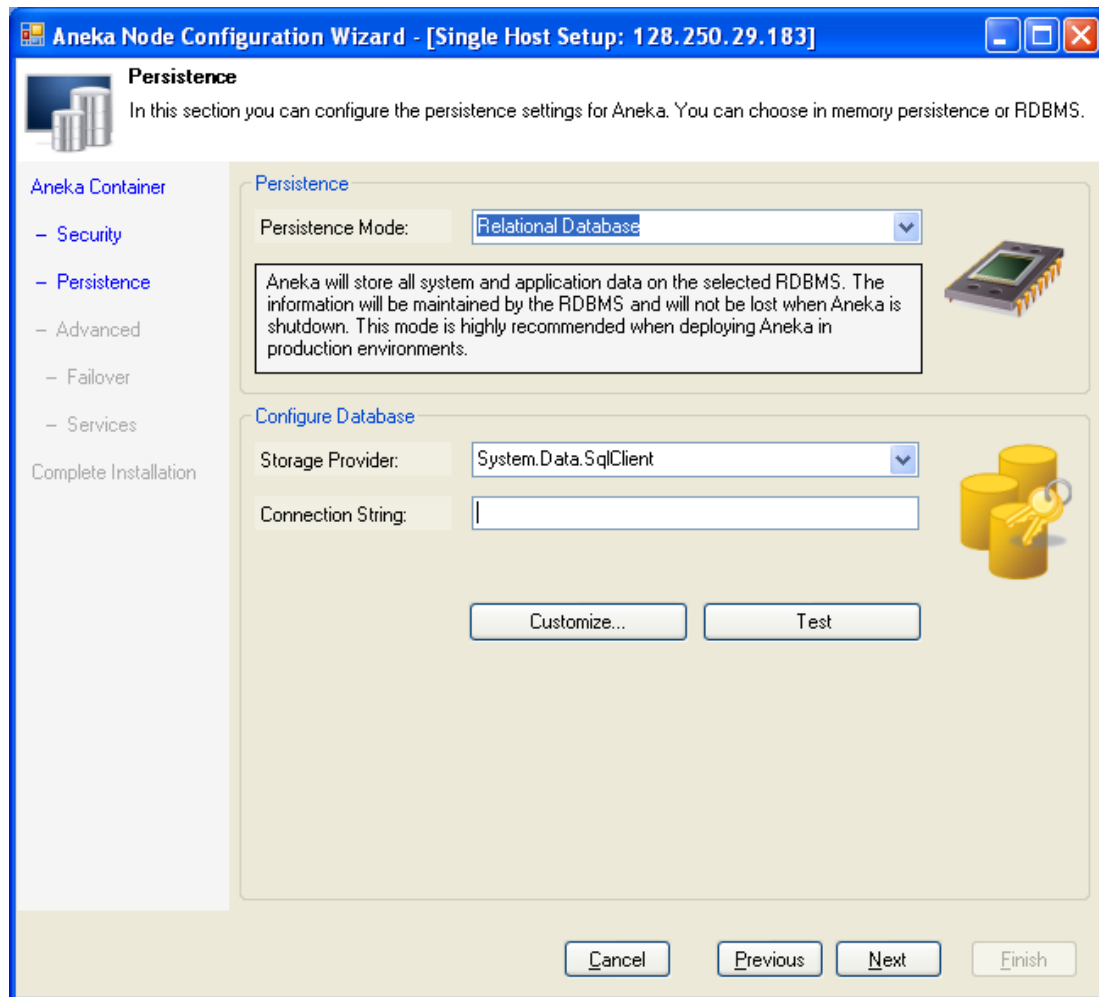


Figure 35 - Setting up persistence for the Aneka Cloud

Selecting *Memory* persistence requires no further configuration. Selecting *Relational Database* requires that you select one of the two supported *Storage Providers*. Aneka currently supports SQL Server 2005/2008 and MySQL 5.1 or greater. You must next configure the *Connection String*. Pressing *Customize* brings up the Connection String configuration dialog as shown in the figure below.

This configuration dialog differs depending on whether you select *System.Data.SqlClient* or *MySql.Data.MySqlClient* as your data provider. If you select *System.Data.SqlClient*

provider for SQL Server 2005/2008 database, you are required to provide, at the least, the *User ID*, *Password*, *Data Source*, and *Initial Catalog*. If you select *MySQL.Data.MySqlClient* provider for MySQL databases you are required to provide, at the least, the *User Id*, *Password*, *Database* and *Server*. Once you have setup the connection string, you can test it by clicking the *Test* button. If a successful connection can be made to the database server you can proceed by clicking *Ok*.

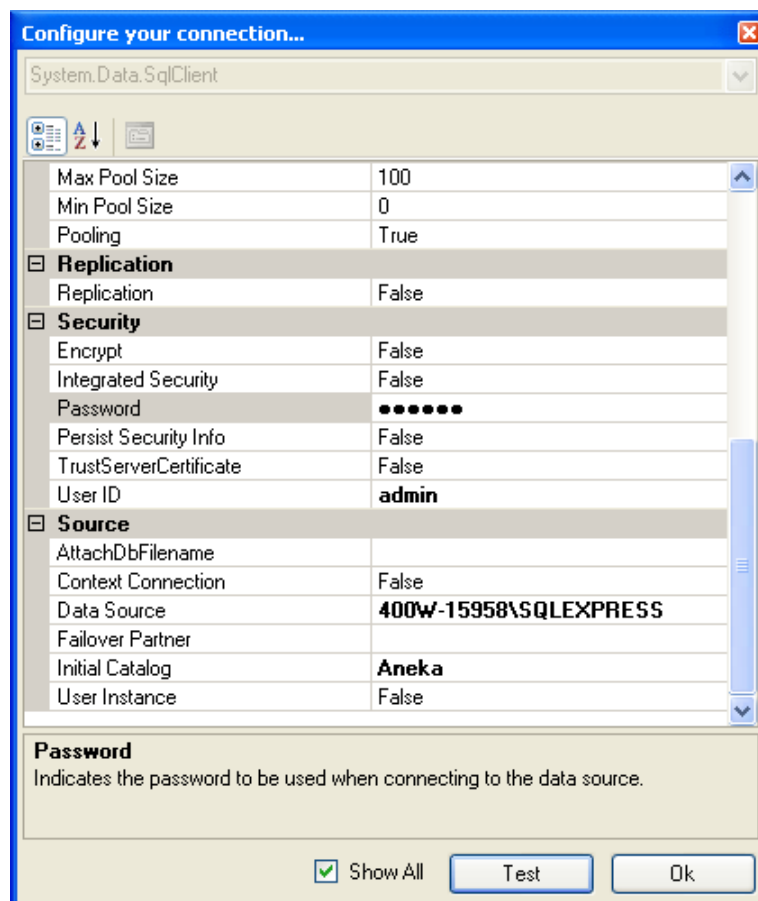


Figure 36 - Connection String configuration dialog

NOTE: If the database store for Aneka was setup during a previous installation, you are still required to specify the provider and connection string. However, the installation Wizard will not create a new store. If you wish to re-create the new store you must drop all tables and stored procedures in your database by running the *CleanStore.sql* script found in the *\Misc\Data* folder in your Aneka installation.

7.1.4 Step 4 – Cost and Software Appliances

You have now completed the basic settings for installing the Container. To review the installation settings and proceed with installing the Container, click on the link “[here](#)” as shown in Figure 37. For advance configuration, continue reading.

Aneka Node Configuration Wizard - [Single Host Setup: 128.250.29.183]

Advanced Configuration Settings
In this page you can set up the settings that further customize the container, such as the name of the instance, its cost per hour, and the installed software appliances.

Aneka Container

- Security
- Persistence
- Advanced
- Failover
- Services
- Complete Installation

The basic installation of the container has been completed!

Click [here](#) to review the installation settings and start the installation, or continue and configure the advanced properties of the Container.

Container Properties

Instance Name: 128.250.29.183-90 ☒ Use IP as Name

Cost (Unit/Hour): 0

Software Appliances

Software appliances currently configured

Name	Vendor	Version
POV-Ray	Persistence of Vision RayTracer Pty. Ltd.	3.1

Details

Name: POV-Ray

Vendor: Persistence of Vision RayTracer Pty. Ltd.

Version: 3

Figure 37 - Cost and Software Appliances

This step lets you configure additional container properties such as the *Instance Name* and *Cost*. The *Instance Name* refers to the name of the Container. You make choose to leave the *Use IP as Name* option checked by default if you would like to use the IP address and port number of the Container as its name, or uncheck to specify a different name. The *Cost* is an integral value stated in units per hour and is used to specify the *billing cost* for executing a job. This property only makes sense for Worker Containers running *Execution Services*. A job that runs for 30 minutes on a Worker Container for example, will be billed at half the specified cost. The aggregated cost of running all jobs across all the Worker Containers in the Cloud for a given application is thus the total cost of executing that application. End-users can be billed accordingly by converting this cost into a meaningful currency value.

Software Appliances are applications that must exist on the Worker node for a job to execute successfully. This can include software such as the Java Runtime Environment (JRE) or rendering engines such as POV-Ray and Maya. Users submitting Jobs to Aneka can state any application requirements for their jobs to execute successfully. The Scheduler will then dispatch these jobs to those Worker nodes that have the required *Software Appliance*.

Software Appliances are therefore typically setup when installing Worker Containers that run *Execution Services*. To add a Software Appliance to a Container you must specify the Name, Vendor and Version as shown in Figure 37, and click on the *Add* button. Note that the same values must be specified when submitting an application with special software requirements.

7.1.5 Step 5 – Failover Configuration

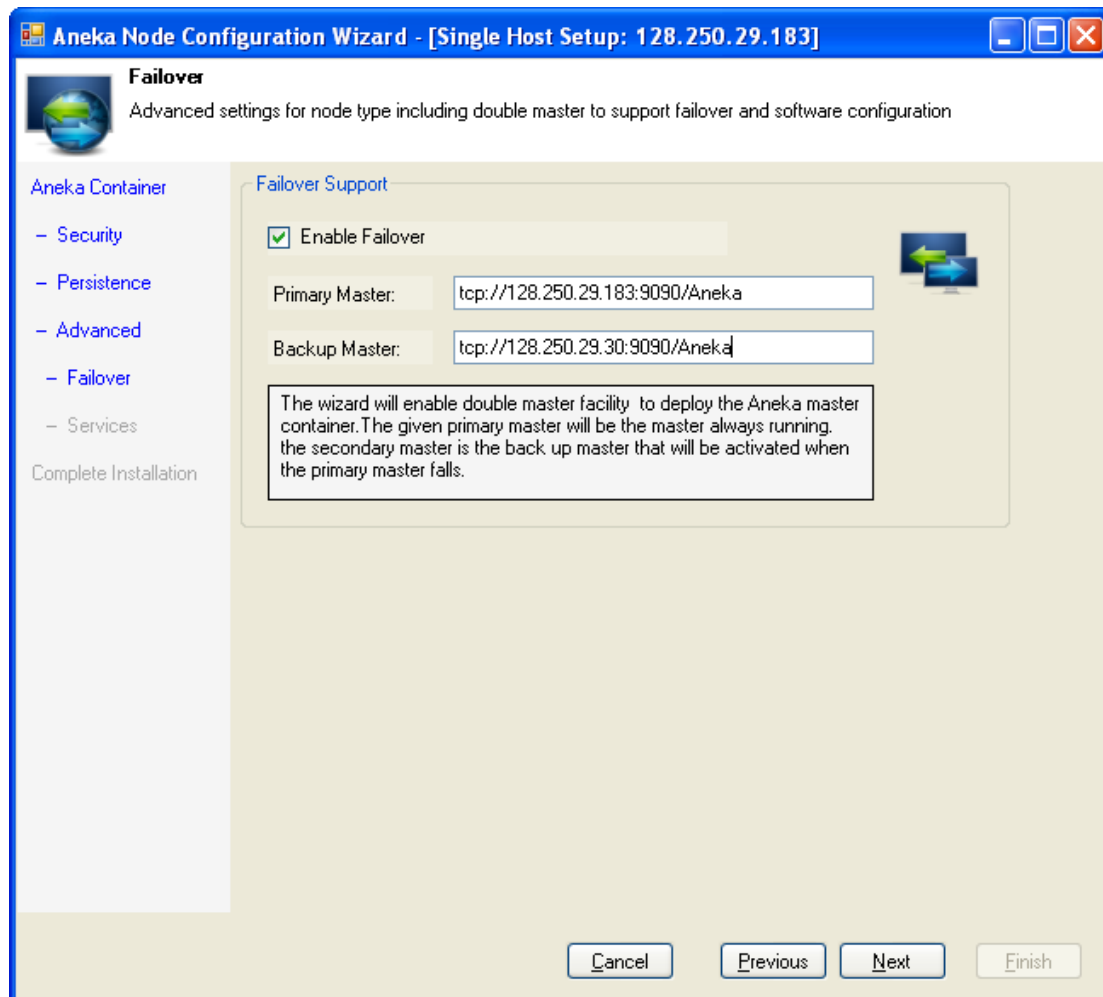


Figure 38 - Failover configuration

A failover system increases the availability of your Aneka Cloud. This step lets you configure your cloud for failover by specifying a *Primary* and *Backup Master*. In the event the *Primary Master* Container should fail, the *Backup Master* will take control of the cloud so that it could continue to function. In order to enable the failover system you must have multicasting working on your network. When installing the Backup Master Container, you must specify the same URLs for *Primary Master* and *Backup Master* as shown in the figure.

7.1.6 Step 6 – Service Configuration

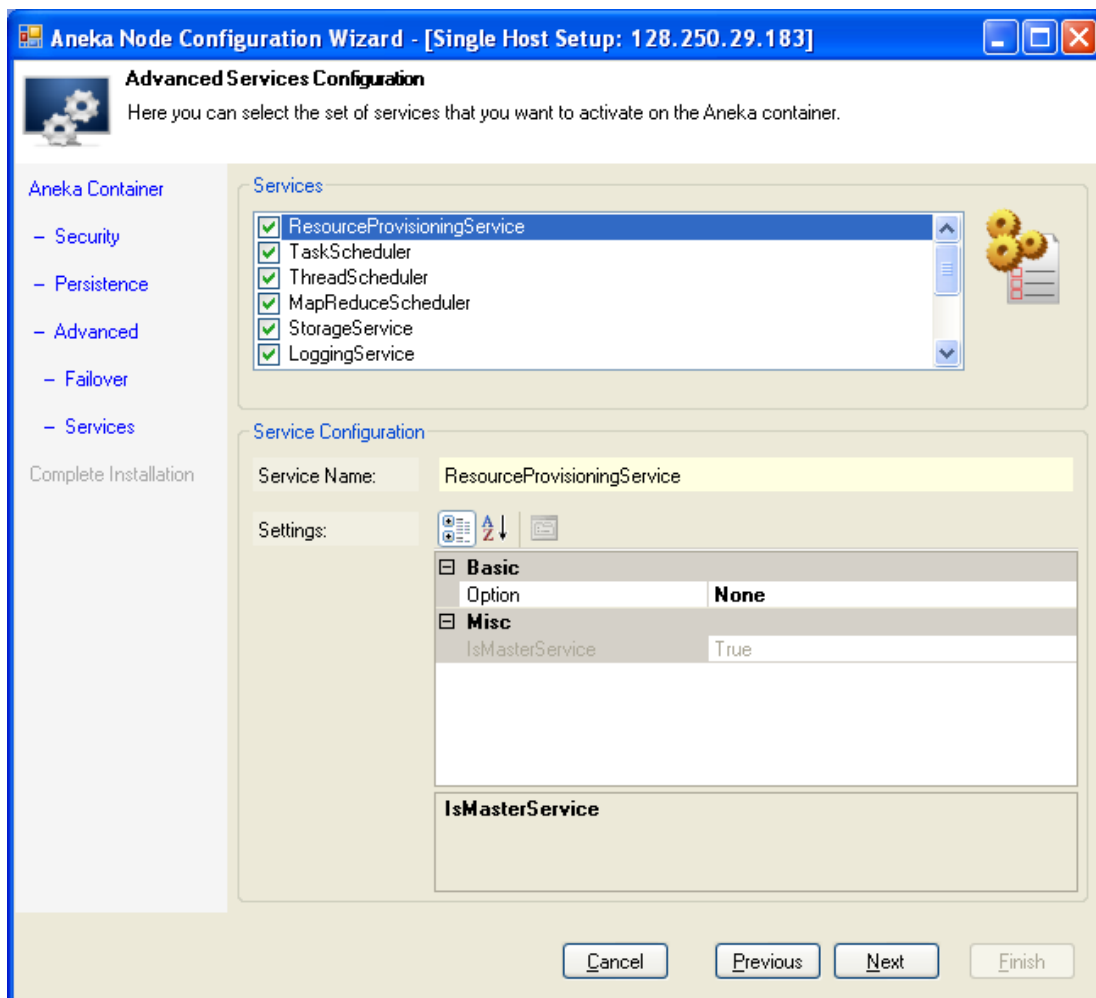


Figure 39 - Service Configuration

This step allows you to select the list of services you wish to host on the Master Container. You may leave the default selection intact for the most common service configuration for the Master Container. This includes the following list of services:

- ResourceProvisioningService
- TaskScheduler

- ThreadScheduler
- MapReduceScheduler
- StorageService
- LoggingService
- ReportingService
- MonitoringService

To configure the properties of each to the listed services, select the service and then edit the properties in the *property box* below. For a more detailed description of each of these services and properties, see Appendix A on Service Configuration. You may also choose to install the corresponding *Execution Services* for each of the scheduling services listed above. This is not a suggested practice, as this will increase the computation load of the Master Container leading to a performance bottleneck in the system. You may however install these services for testing and evaluation purposes.

7.1.7 Step 7 - Summary

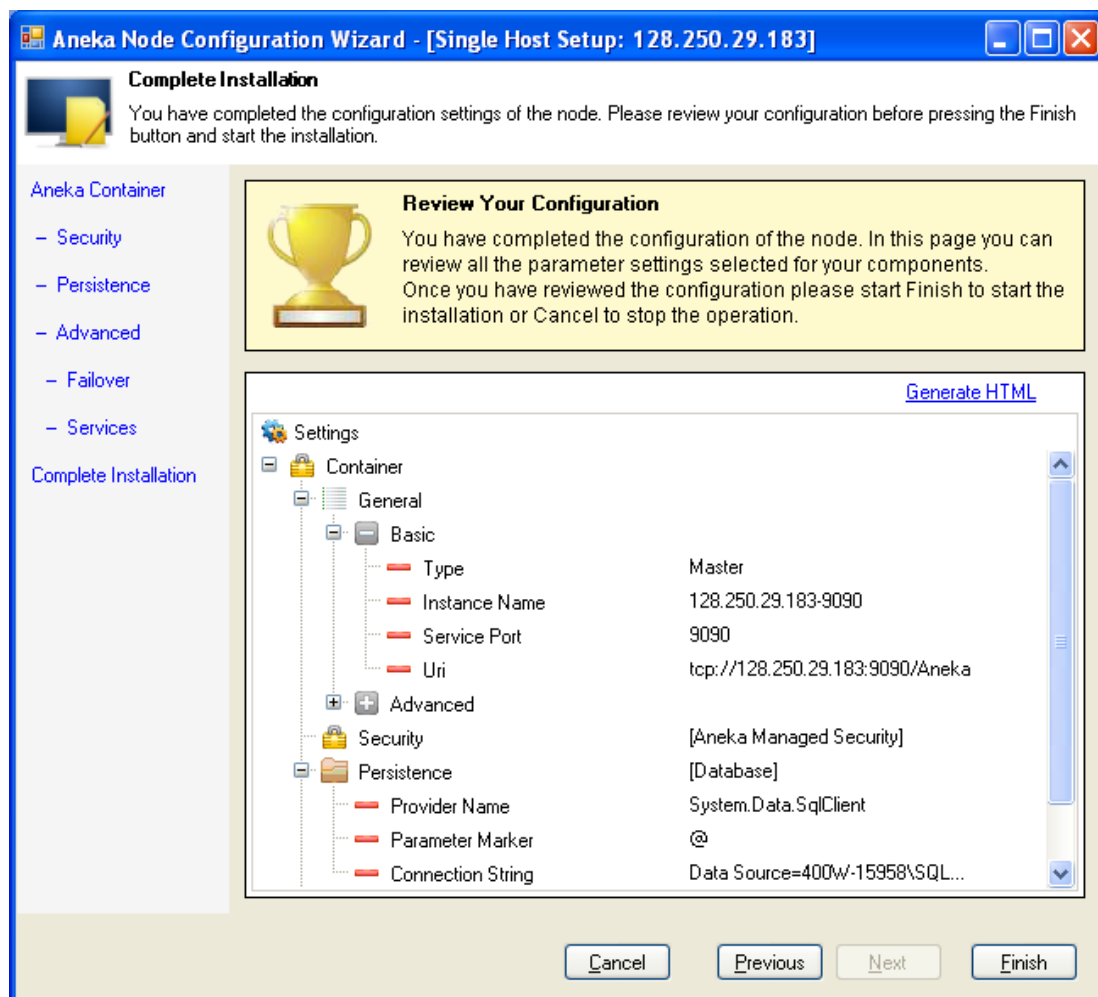


Figure 40 - Summary of the installation settings

The final step of the Wizard presents a summary of the installation settings. Browse through the settings to ensure that everything is correct. In order to change any of the settings click *Previous* until you reach the corresponding step of the Wizard. To start installing the Container click *Finish*.

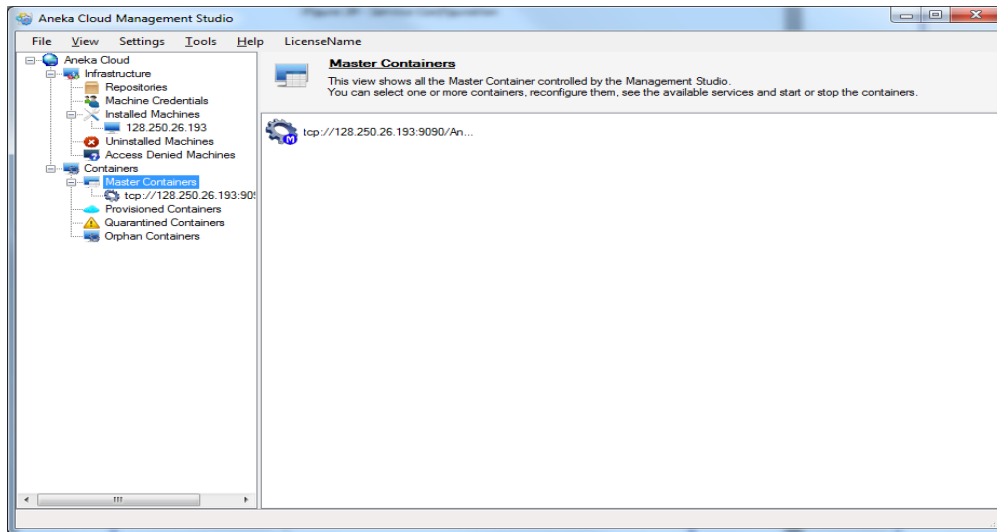


Figure 41 - The installed Master Container

To view the installed Master Container, click on *Master Containers* in the Cloud Explorer. The figure above shows the installed Master Container up and running. Now that you have the Master Container ready, you can proceed to setting up your Workers.

7.2 Installing Worker Containers

The process of installing Worker Containers is similar to that of installing the Master. Select the list of machines that you wish to install Workers on, and then select *Install Container* from the context menu as shown in the figure below.

7.2.1 Step 1 – Container Type and Port

Selecting *Install Container* on the chosen machines brings up the Container Installation Wizard. This Wizard will begin by automatically probing the default Container port, 9090, on all the selected hosts in order to verify that the port is free. If the port is not free, as a result of being used by another program, you must either ensure that it becomes free, or assign a new port number for the Containers. If you wish to verify whether your new port is free, press the *Probe* button again.

NOTE: If the port is being occupied by an existing Container on the target machine, the Wizard will notify you of its existence and ask if you want to reconfigure it. If your intention is to install a new container, you must change the Container port. On the other hand, if you intend to reconfigure the existing Container you may proceed. To reiterate, the Container Installation Wizard can either be used to install a new Container or reconfigure an existing Container.

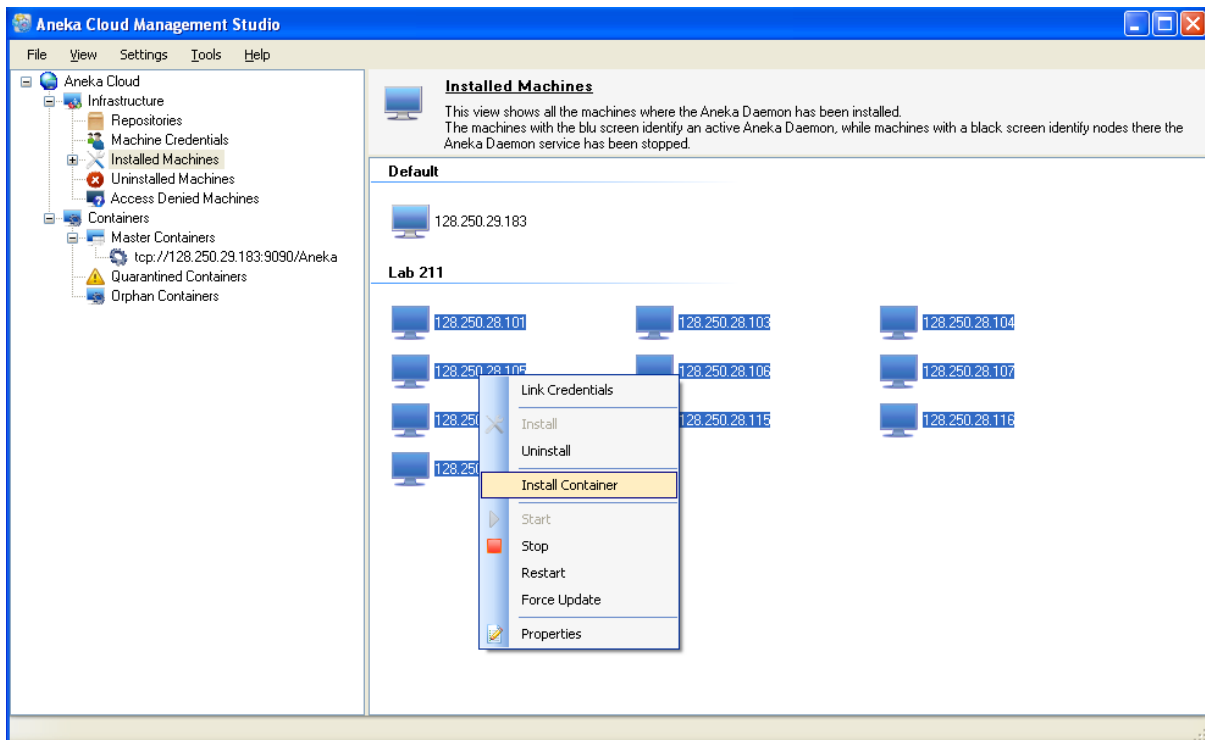


Figure 42 - Installing Worker Containers

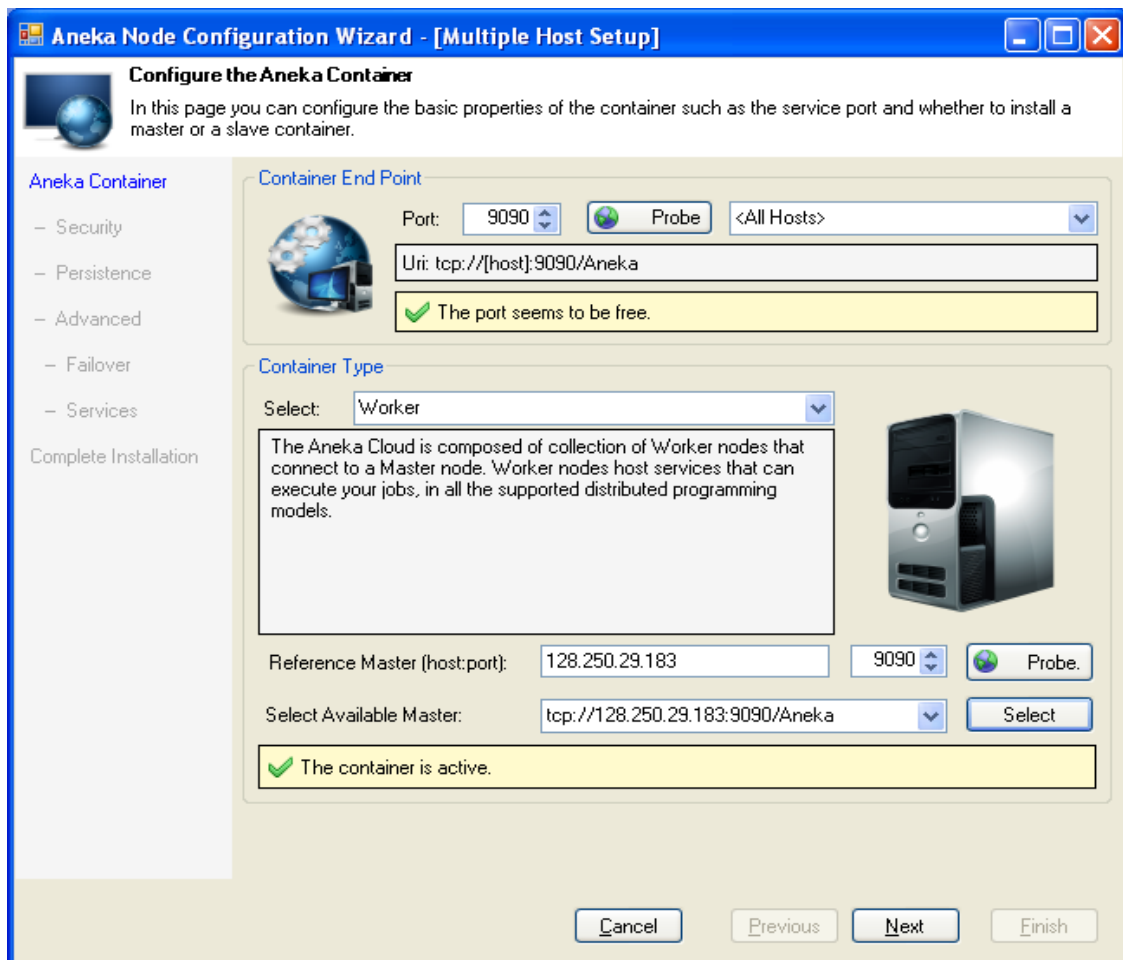


Figure 43 - The Container Installation Wizard for installing Workers

Ensure that the selected *Container Type* is *Worker*. Next select from one of the available Masters, or manually enter the host IP and port number of the Master. *Probe* the *Reference Master* you just entered in order to verify that it is up and running. Click *Next* to proceed.

7.2.2 Step 2 – Security Configuration

If you chose to disable security when configuring the Master Container, you must also disable security for your Workers. However, it is highly recommended that you enable security in a production environment. This ensures that the Master and Workers authenticate each other and all traffic exchanged between them is encrypted.

To setup security for your Workers, you must either type in the same *Secret Key* that you used for your Master or select the same key from the *Key Store*. Click *Next* to proceed.

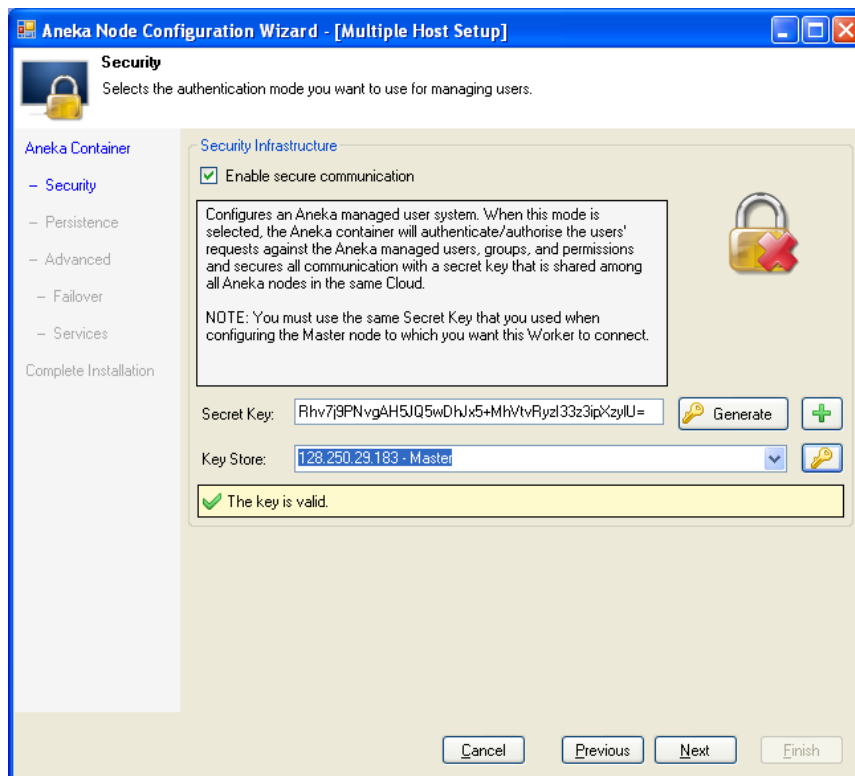


Figure 44 - Configuring security for your cloud

7.2.3 Step 3 – Persistence Configuration

No special persistence configuration is required for Workers. Click *Next* to proceed. (See figure in the next page).

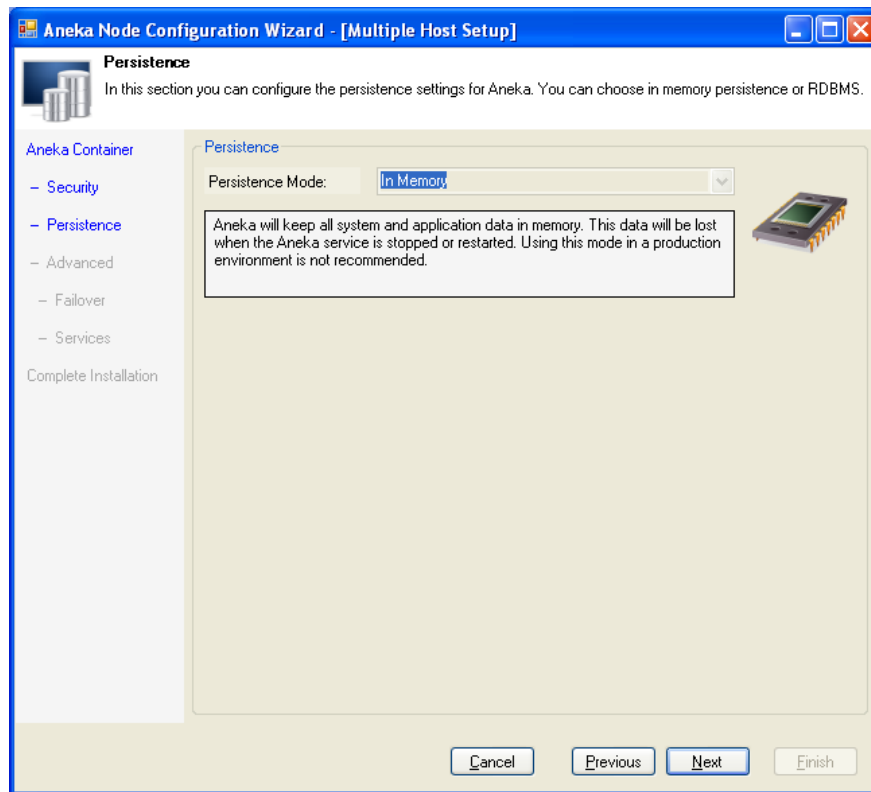


Figure 45 - No persistence configuration is required for Workers

7.2.4 Step 4 – Cost and Software Appliances

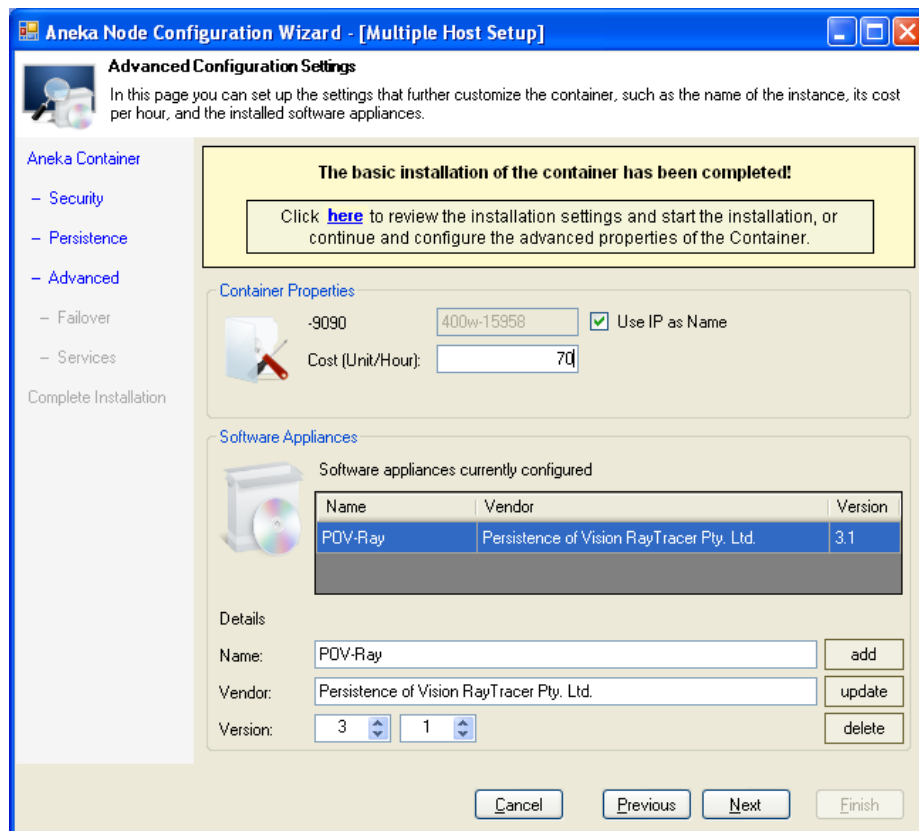


Figure 46 - Cost and Software Appliances

You have now completed the basic settings for installing the Container. To review the installation settings and proceed with installing the Containers, click on the link “[here](#)” as shown in the figure above. For advance configuration, continue reading.

This step lets you configure additional container properties such as the *Instance Name* and *Cost*. The *Instance Name* refers to the name of the Container. You make choose to leave the *Use IP as Name* option checked by default if you would like to use the IP address and port of the Container as its name, or uncheck to specify a different name. The *Cost* is an integral value stated in units per hour and is used to specify the *billing cost* for executing a job. This property only makes sense for Worker Containers running *Execution Services*. A job that runs for 30 minutes on a Worker Container for example, will be billed at half the specified cost. The aggregated cost of running all jobs across all the Worker Containers in the Cloud for a given application is thus the total cost of executing that application. End-users can be billed accordingly by converting this cost into a meaningful currency value.

Software Appliances are applications that must exist on the Worker node for a job to execute successfully. This can include software such as the Java Runtime Environment (JRE) or rendering engines such as POV-Ray and Maya. Users submitting Jobs to Aneka can state any application requirements for their jobs to execute successfully. The Scheduler will then dispatch these jobs to those Worker nodes that have the required *Software Appliance*.

Software Appliances are therefore typically setup when installing Worker Containers that run *Execution Services*. To add a Software Appliance to a Container you must specify the Name, Vendor and Version as shown in Figure 46, and click on the *Add* button.

7.2.5 Step 5 – Failover

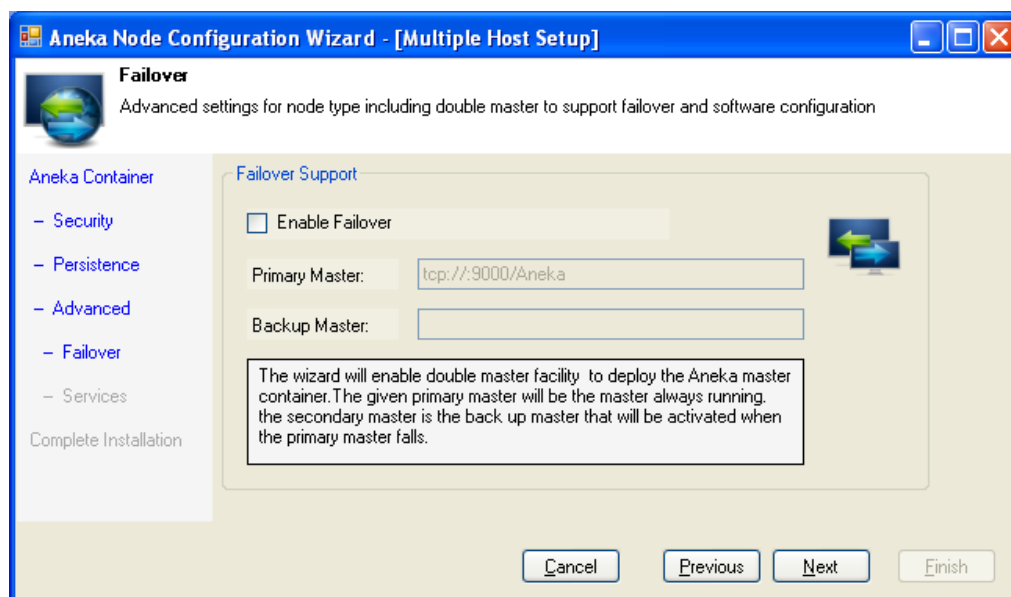


Figure 47 - No failover configuration required for Workers

No failover configuration is required for Workers. Click *Next* to proceed.

7.2.6 Step 6 – Service Configuration

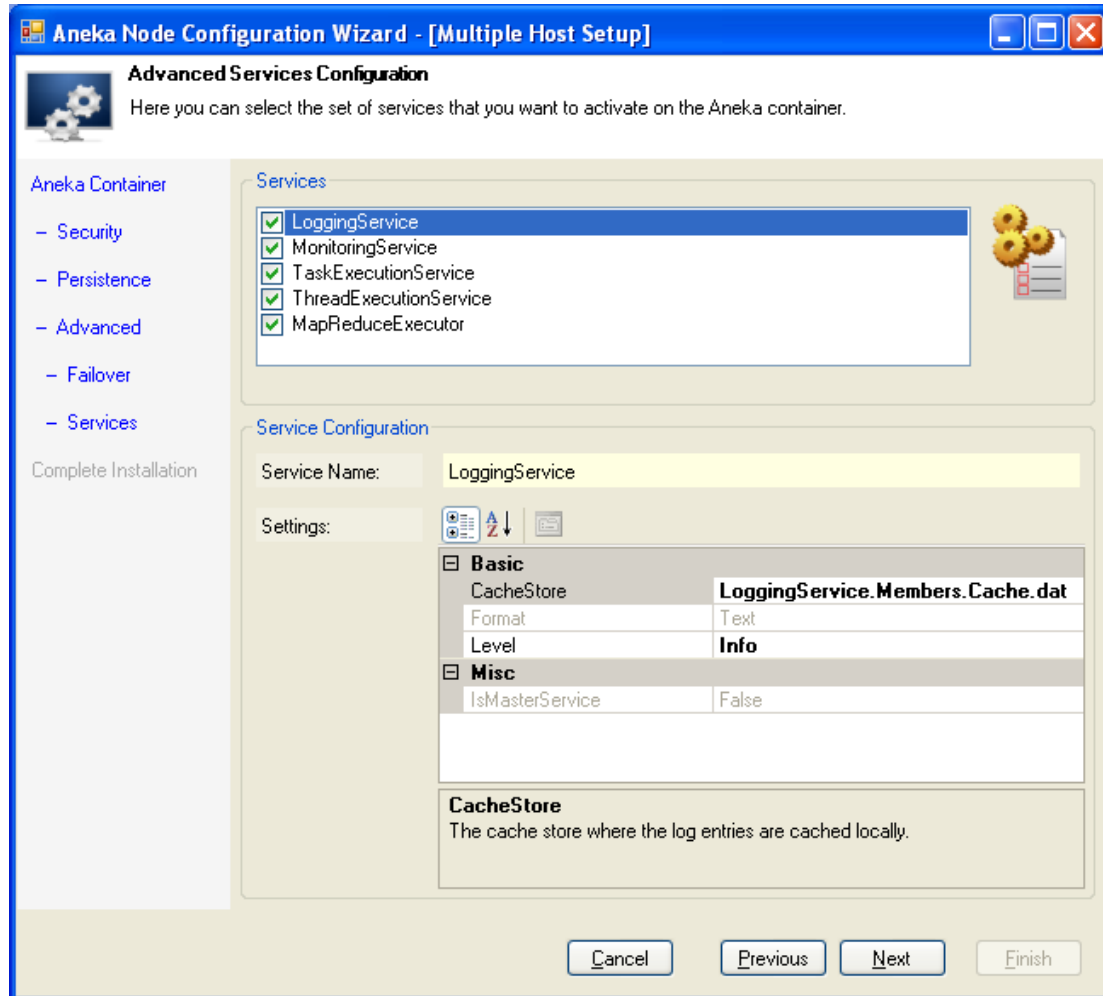


Figure 48 - Service Configuration

This step allows you to select the list of services you wish to host on the Worker Container. You may leave the default selection intact for the most common service configuration for Worker Containers. This includes the following list of services:

- LoggingService
- MonitoringService
- TaskExecutionService
- ThreadExecutionService
- MapReduceExecutor

To configure the properties of each to the listed services, select the service and then edit the properties in the *property box* below. For a more detailed description of each of these services and properties, see Appendix A on Service Configuration. Click *Next* to proceed.

7.2.7 Step 7 - Summary

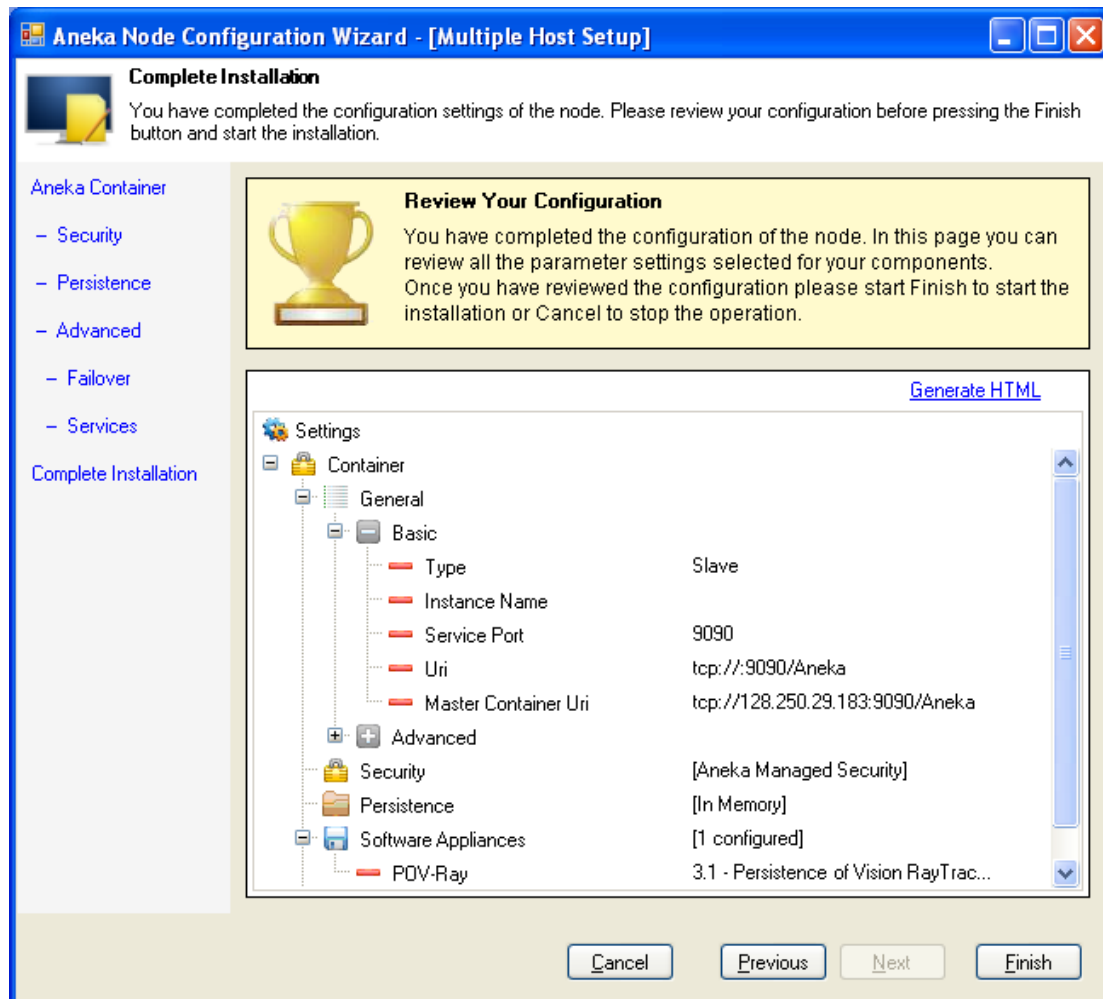


Figure 49 - Summary

The final step of the Wizard presents a summary of the installation settings. Browse through the settings to ensure that everything is correct. In order to change any of the settings click *Previous* until you reach the corresponding step of the Wizard. To start installing the Container click *Finish*.

Figure 50 below shows the installation progress. This process can take a few minutes depending on the number of machines you are installing the Workers on. Once the installation has completed, click on the Master Container in Cloud Explorer to see a list of all new Workers connected to the Master Container.

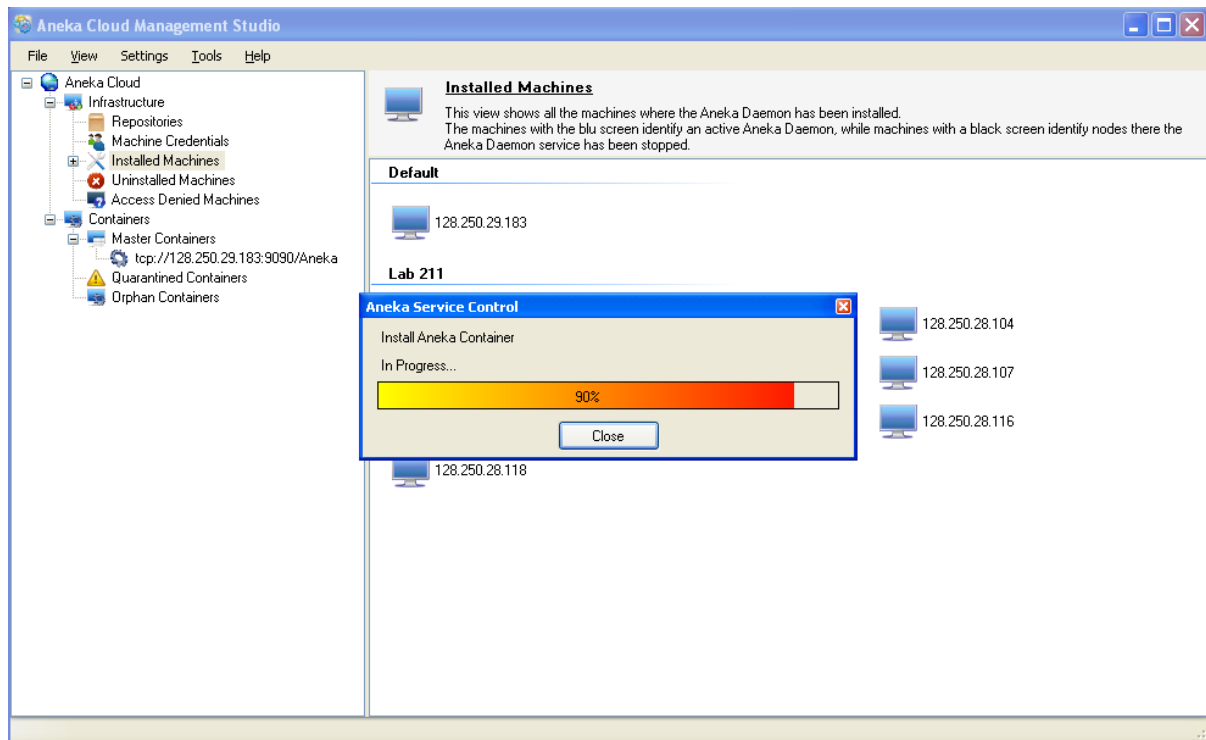


Figure 50 - Installation of Worker Containers

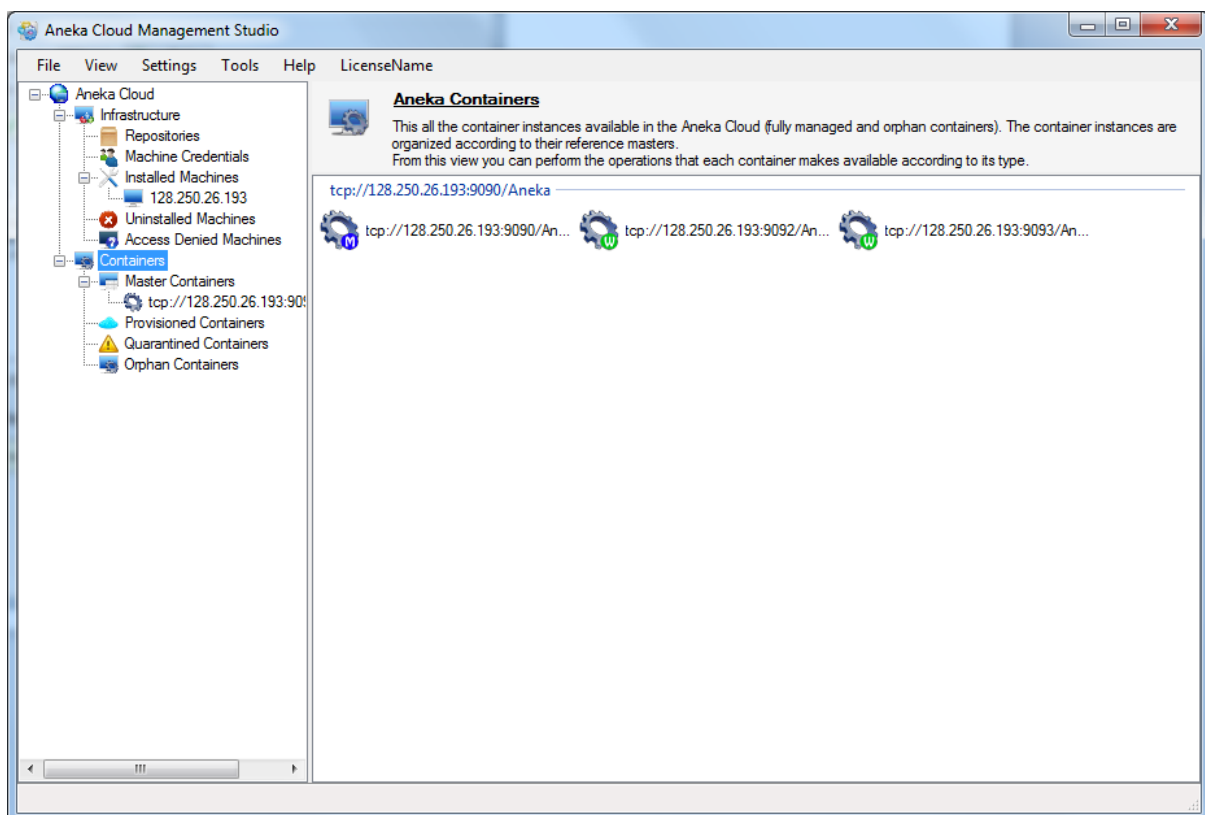


Figure 51 - Installed Worker Containers

8 Cloud Management

Aneka Management Studio provides a number of other features that help you manage your Clouds. This includes features for managing Aneka user accounts, monitoring the overall performance of your Cloud, and obtaining detailed reporting information on resource usage, data transfers, billing and application (job) execution. It also provides facilities for starting, stopping, quarantining Containers and troubleshooting your deployments by viewing remote logs.

8.1 Master Containers

To view list of all Aneka Clouds deployed on your infrastructure, click on *Master Containers* in the *Cloud Explorer*. Each *Primary Master Container* identifies a separate Aneka Cloud. To view the list of Workers in the Cloud click on the corresponding Master Container in Cloud Explorer as shown in the figure below.

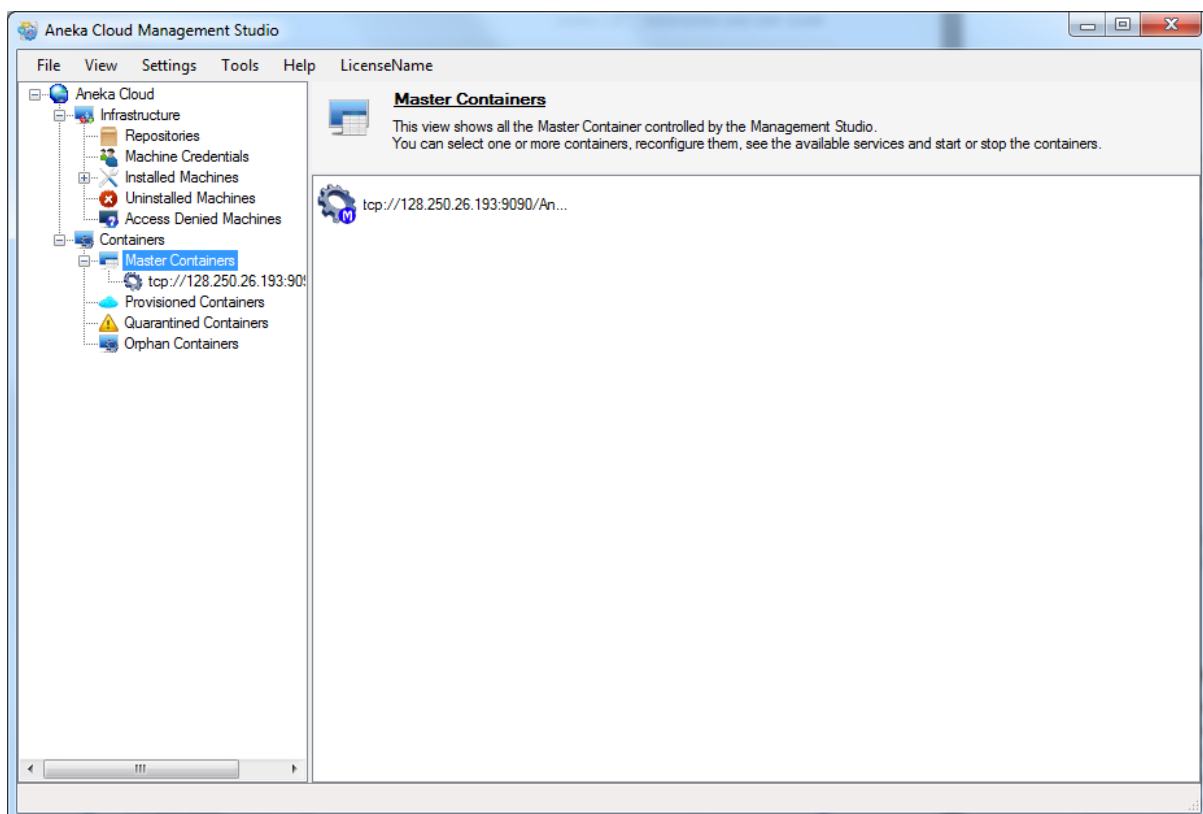


Figure 52 - A List of Master Containers

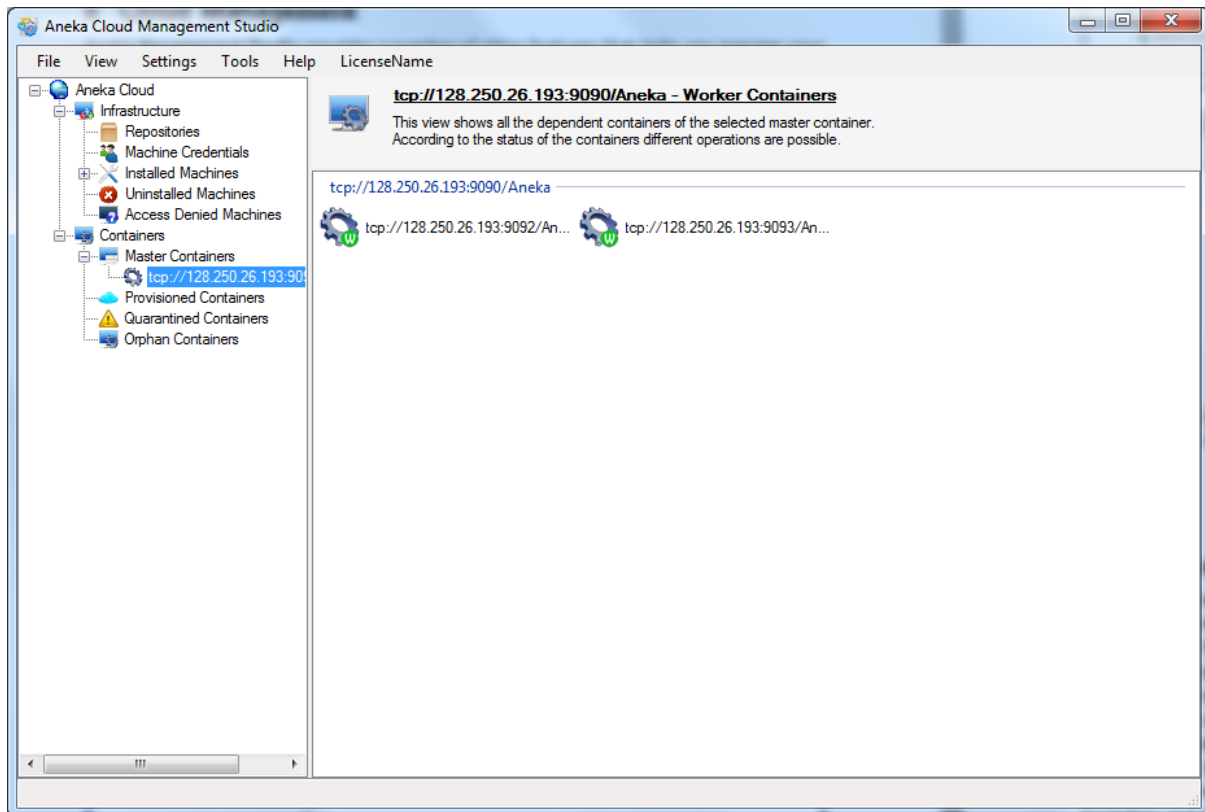


Figure 53 - The Workers belonging to a Master

8.2 Quarantined Containers

The *Quarantine* action allows you to put Containers under *Quarantine* or remove them from *Quarantine*. *Quarantine* is a status that prevents a Container from being selected by the scheduler for assigning jobs to be executed. When a container is under *Quarantine* it will never receive any job to execute. This feature is useful when the administrator detects that some Containers do not behave properly and wants to temporarily exclude them from the set of eligible Containers for job execution until the problem is fixed. Once a container has been put under *Quarantine* it will disappear from the list of Worker Containers in the cloud, and will be put under the *Quarantined Containers* in Cloud Explorer. To remove a container from the *Quarantine*, simply select the *Quarantined Container* and then select *Remove from Quarantine* from the context menu. Note that you are not allowed to quarantine the Master Container.

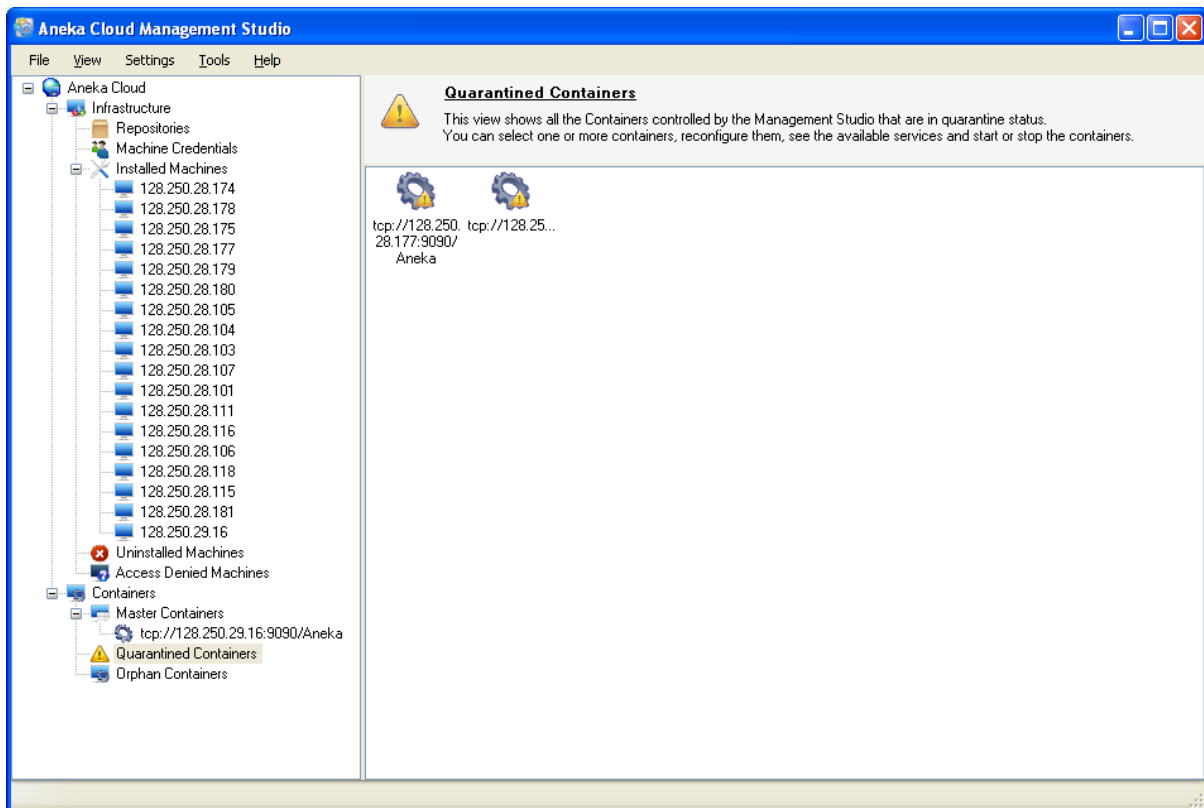


Figure 54 - Quarantined Containers

8.3 Orphan Containers

Containers become *Orphans* when the Aneka Daemon responsible for managing them is no longer available. This commonly happens when the Aneka Daemon crashes, but leaves the Containers intact and usable. *Orphan Containers* can still receive jobs but provide limited management capabilities because of the lack of control. It is recommended that you troubleshoot the problem and resume the Aneka Daemon service. Figure 55 shows a sample list of *Orphan Containers*.

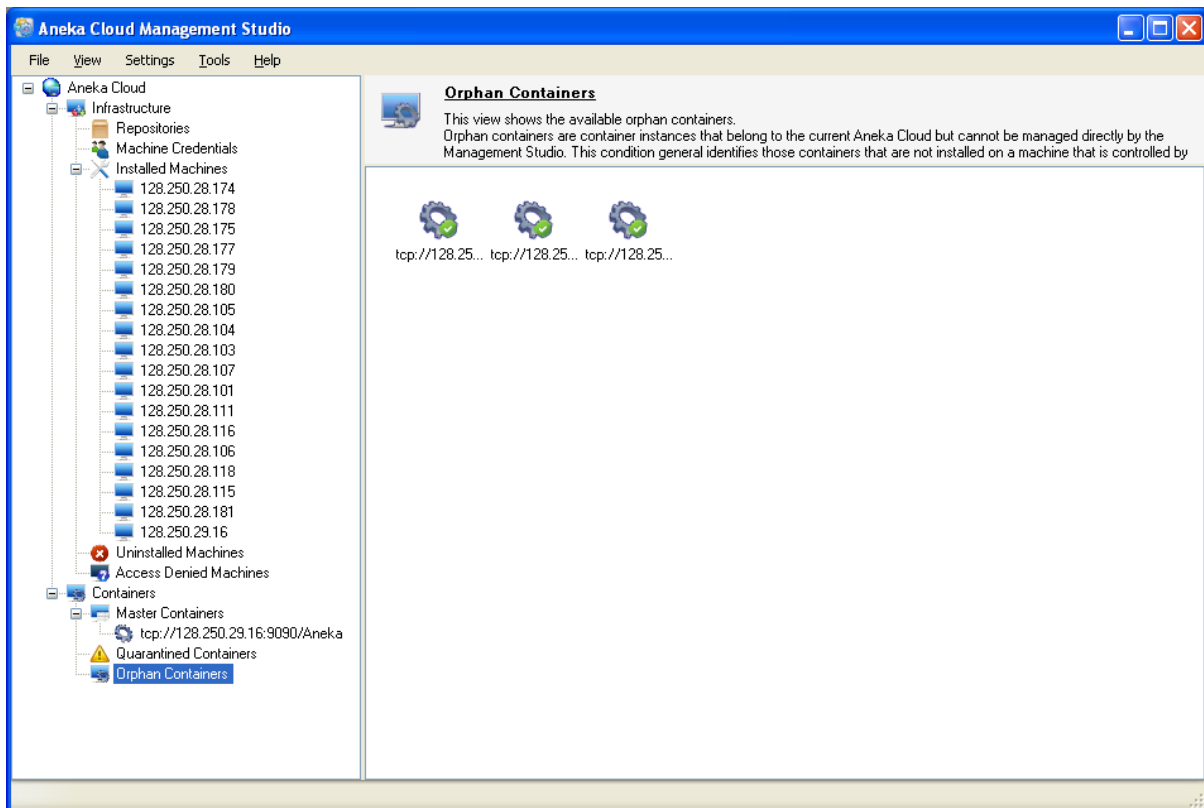


Figure 55 - Orphan Containers

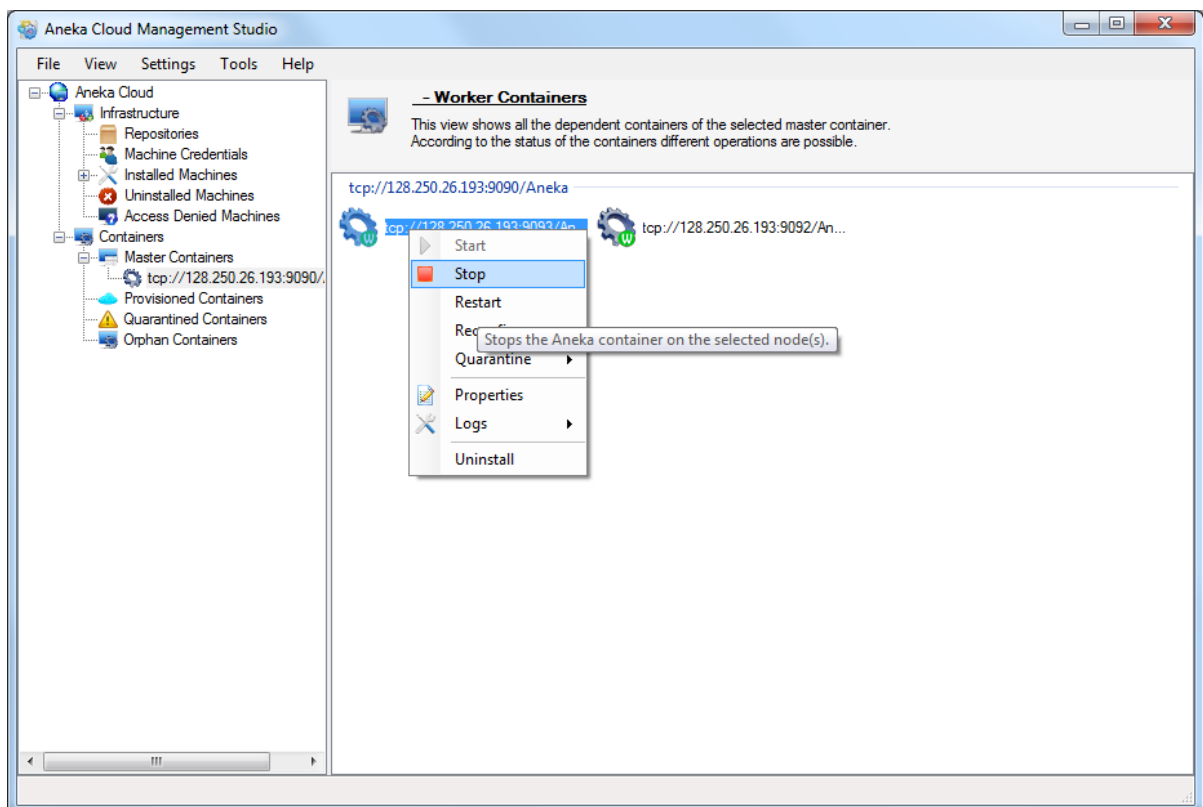


Figure 56 - Managing a Container

8.4 Container Management

The following is a list of actions that can be performed on Containers.

Start: Starts the selected Containers.

Stop: Stops the selected Containers.

Restart: Similar to performing *Stop* followed by *Start*. This action restarts the selected Containers.

Reconfigure: Reconfigure the selected Containers. This action launches the Container installation Wizard once again, allowing you to reconfigure the Container.

Quarantine: As discussed earlier, this action quarantines the Container.

Properties: This action only applies to a single Container and it pops up a dialog as shown in Figure 57 showing the properties of the Container. The dialog displays multiple tabs presenting information on general system properties of the Container, the list of Services running in the Container, and the performance of the Container.

Logs: This action allows you to see real time monitoring of logs for a particular Container, or to view the log archives. You can also specify interested logging levels (info, debug, warn and error) using the *Logging Listener Port* dialog in the *Settings* menu.




Figure 57: Aneka Container Properties

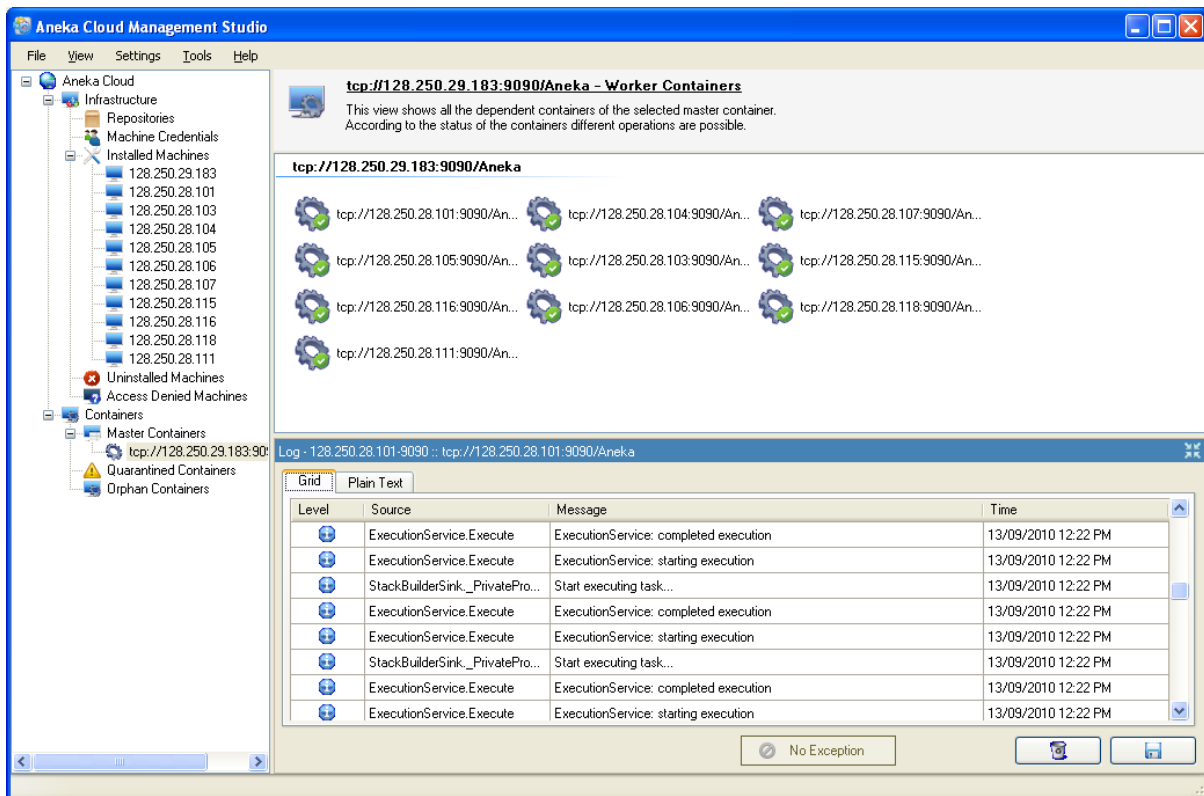


Figure 58: Container Logging (Real time monitoring and log archives)

8.5 Real-Time Performance Monitoring

Aneka Management Studio allows you to monitor the overall performance of your Cloud. Selecting *Statistics* from the *Tools* menu brings up the *System Performance Monitor* showing statistical data on CPU and memory availability and current usage in the Cloud as shown in the figure below. This data is constantly updated every few seconds, reflecting the most current status.

This graphical view of your Cloud can be useful when you need to observe the current load in the Cloud. If you are interested in analysing resource usage over a longer period of time see section on *Reporting and Accounting* below.




Figure 59 - The overall performance of an Aneka Cloud

8.6 Reporting and Accounting

The *Reporting and Accounting* dialog provides detailed information on resource usage, file transfers, billing and application execution. You can access this dialog by selecting *Reporting and Accounting* from the Tools menu. Accessing *Reporting and Accounting* information on an Aneka Cloud requires administrative privileges bringing up an access credentials dialog as shown in the figure below. The default Administrator password for all Aneka Clouds is an empty string. You can thus leave the password field blank, and click *Connect*.

The figure shows a 'Connect to the Aneka Cloud' dialog box. It contains a section titled 'Aneka User Credentials' with the following fields:

- Username: Administrator
- Password: (empty)
- Master Uri: tcp://128.250.29.183:9090/Anek
- Cloud Name: (empty)

At the bottom of the dialog are two buttons: 'Connect' and 'Cancel'.

Figure 60 - Dialog to enter Administrator password

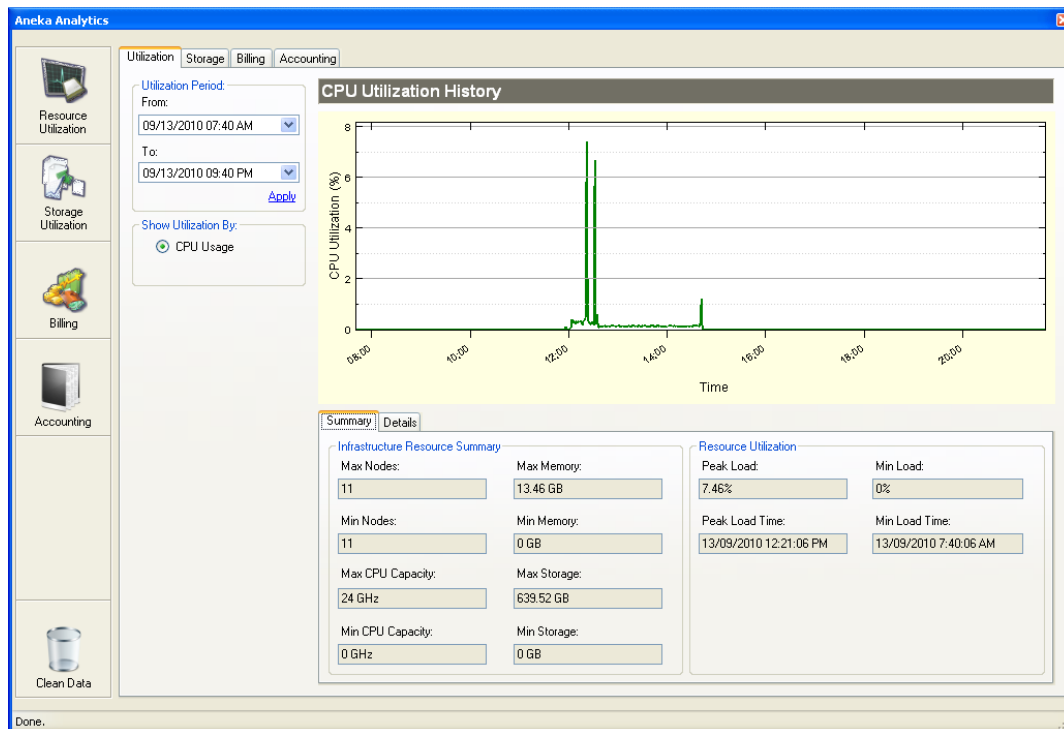


Figure 61 - Reporting and Accounting - CPU Utilization

Figure 61 shows the *Reporting and Accounting* dialog. To view the CPU utilization history for your Cloud, select the desired period in the *Utilization* tab and click *Apply*. This shows a graphical view of the CPU utilization along with a summary of the infrastructure and resources. For a detailed view of the performance information reported by the various Containers in the Cloud click on the *Details* tab.

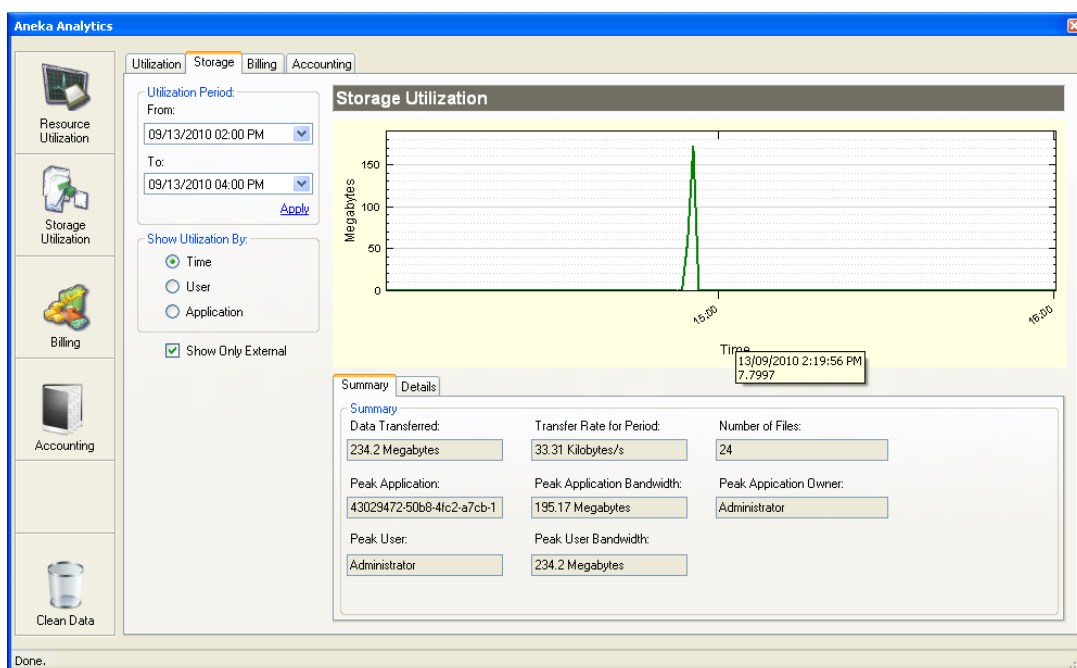


Figure 62 - Reporting and Accounting - Storage Utilization

To view the Storage utilization in your Cloud, switch to the *Storage* tab in the *Reporting and Accounting* dialog, select the desired period and click *Apply*. This shows a graphical view of the Storage Utilization based on *Time*, *User* or *Application*. By default, all internal file transfers between the Storage Service and Worker Containers are not shown. Instead, this dialog presents information on files exchanged between Clients and the Storage Service. To view internal transfers uncheck *Show Only External*. To view detailed information on individual transfers, click on the *Details* tab.

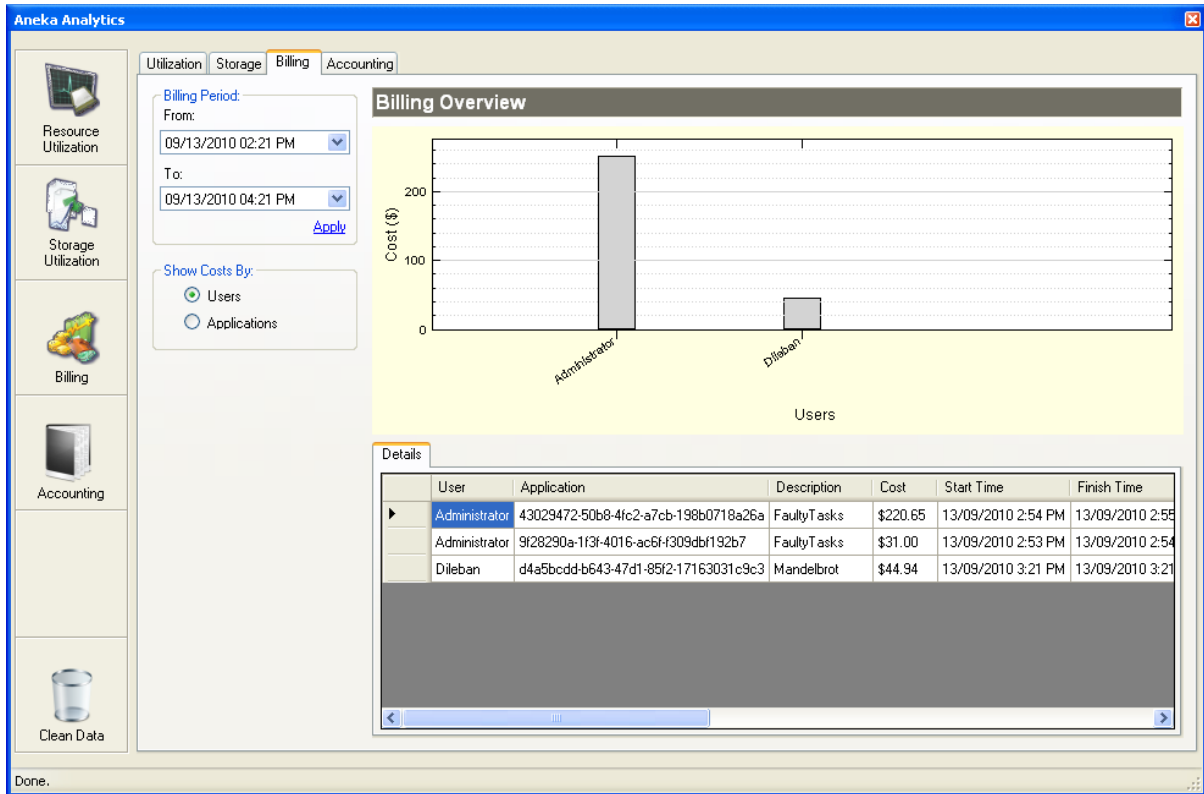


Figure 63 - Reporting and Accounting - Billing

Figure 63 shows Billing information for Aneka users. To view Billing information switch to the *Billing* tab in the *Reporting and Accounting* dialog, select the desired billing period and click *Apply*. This shows a graphical view of the costs incurred by different users. To view costs for each application, selection *Applications* in the *Show Costs By* group box. The *Details* tab at the bottom of the screen shows more information on the each of the applications, including its total cost, start and finish times, numbers of resources used, and the number of successful and failed tasks.

Figure 64 and Figure 65 shows *Accounting* information for applications running on Aneka. This allows users analyse and monitor running applications, view historical data on job execution, and statistics on the application as a whole. The *Accounting* tab consists of *Applications* and *Information* tabs as shown below. The *Applications* tab lists all applications currently running on the Cloud. An Administrator can terminate any applications if he or she wishes to. Click *Refresh* to update the list of running applications.

The *Information* tab shows historical and statistical information on application execution, including the jobs submitted and the resources utilized. The *Statistics* tab presents a graphical representation of the resources used by an application in the form of a job distribution chart, showing the number of work units completed on each of the resources. End-users will only be able to see his/her application details. The administrator on the other hand can view details of all the applications.

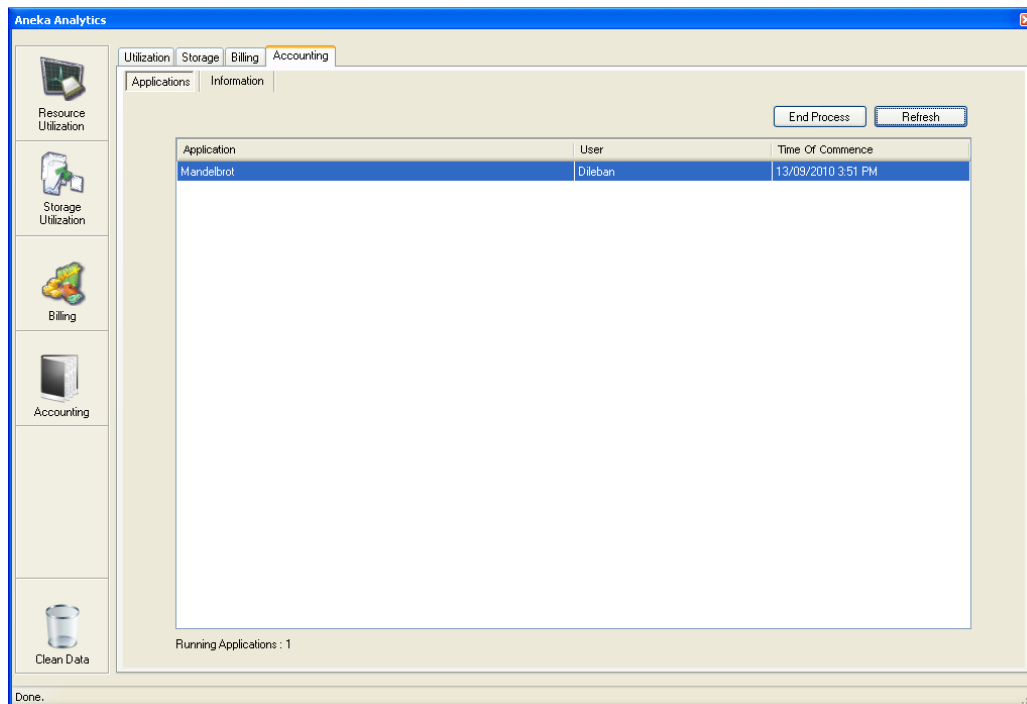


Figure 64 - Reporting and Accounting - Running Applications

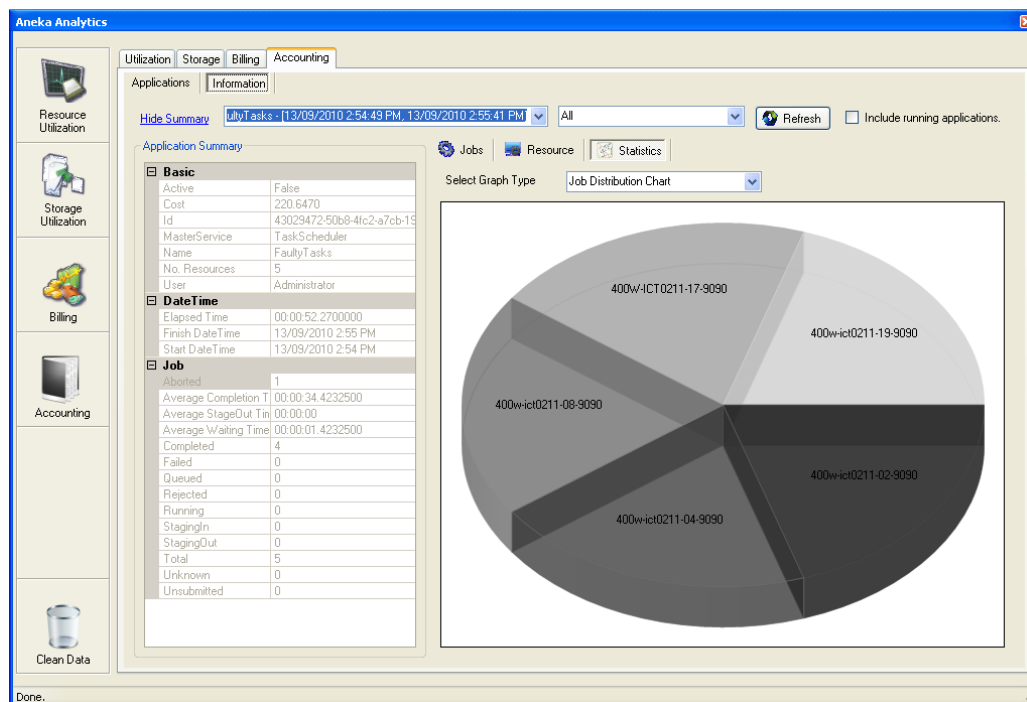


Figure 65 - Reporting and Accounting - Application Statistics

8.7 User Management

The *User Management* dialog can be accessed by selecting *User Management* from the Tools menu. This dialog is used for managing Aneka user accounts, user groups and user permissions. This dialog is only available if the Aneka Cloud has been configured to use security.

The first tab, *User Management* is used for creating, updating and deleting Aneka user accounts. End-users must have a valid account in order to submit applications if security has been enabled on the Cloud. Only an administrator is able to create and delete user accounts. End-users are allowed to modify their own accounts.

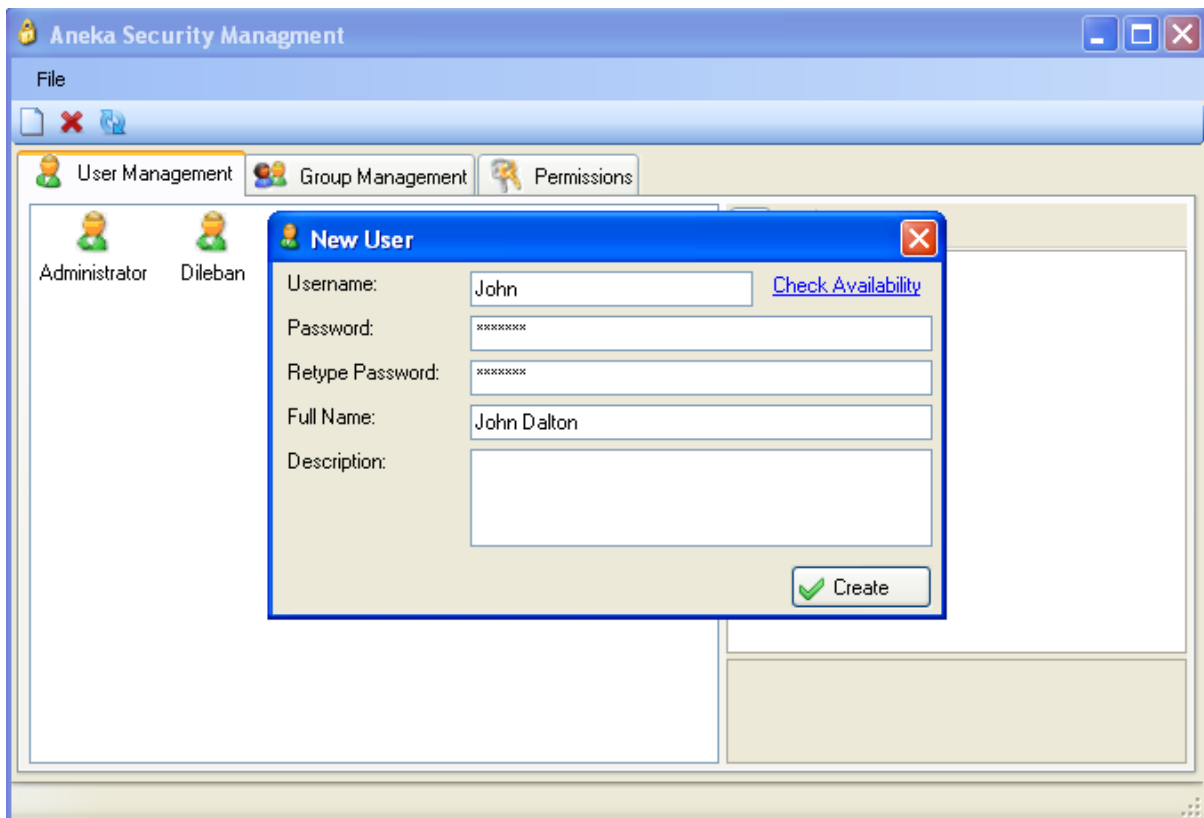


Figure 66: Creating an Aneka user

The second tab, *Group Management*, is used to create user groups that link different user accounts and user permissions, as shown in Figure 67. Only Administrators of Aneka can manage user groups.

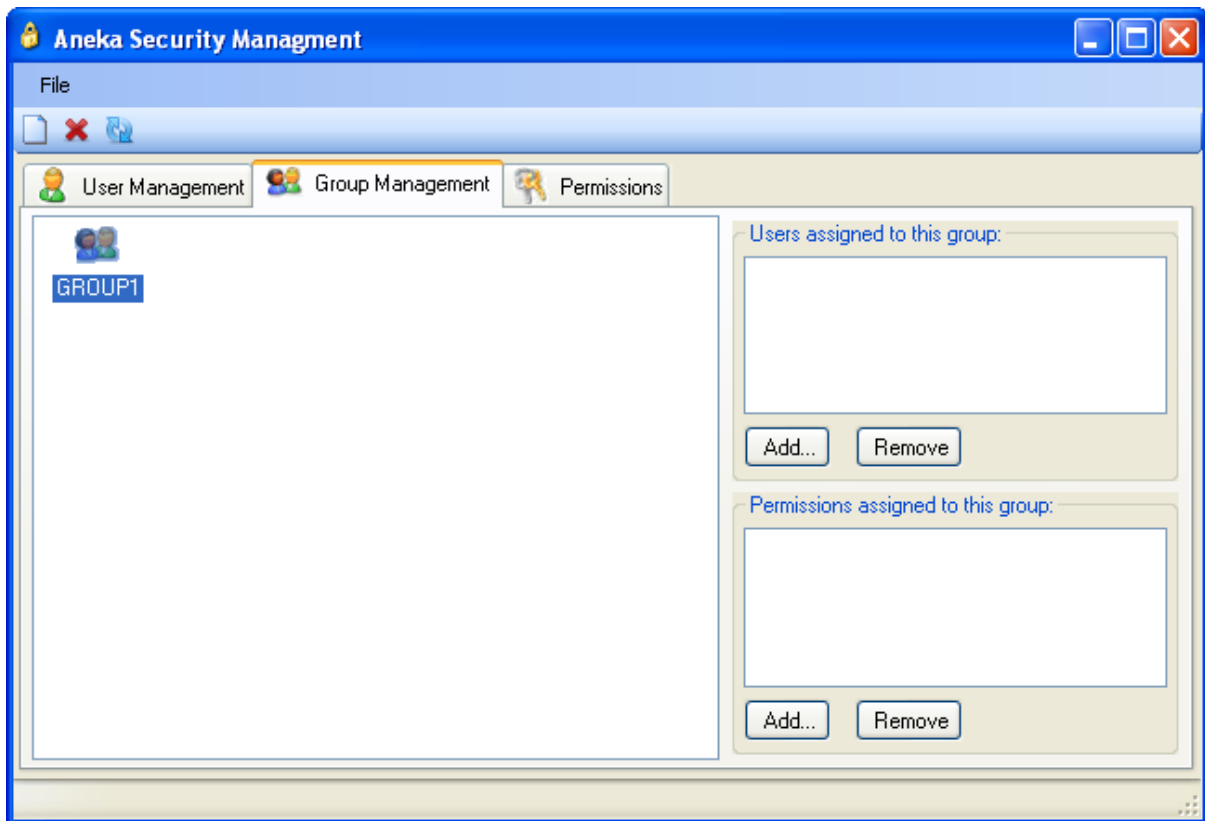


Figure 67: User group creation

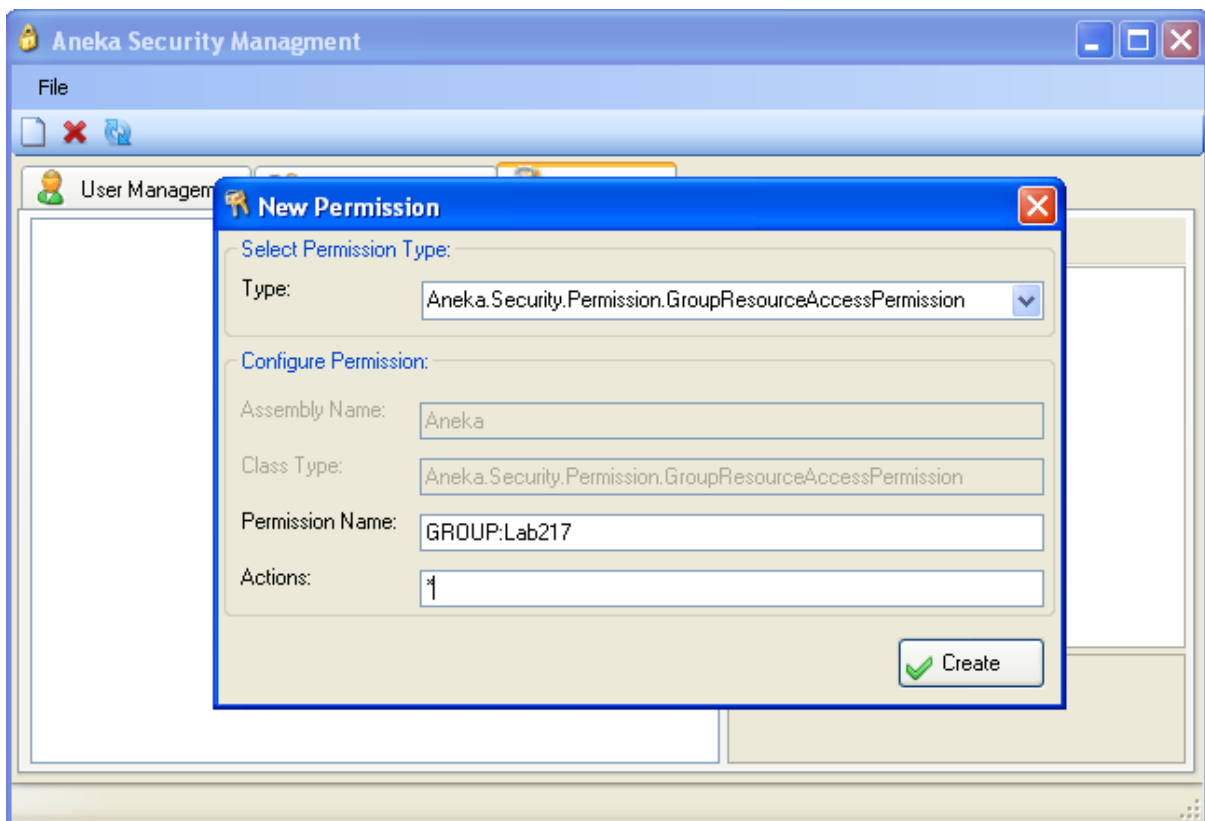


Figure 68 - Creating a new permission

The last tab, *Permission Management*, is used to create various user permissions that apply certain restrictions on Aneka related activities. One of the permissions implemented in current version of Aneka is the *GroupAccessPermission*. It basically treats Aneka Containers in different groups, and only allows access to certain groups of Containers for user accounts attached to the permission. The user permission management example is shown in the figure above. The *GroupAccessPermission* uses two parameters to set the permission for the user. The first parameter is the Permission Name, it should be in the form of `GROUP:[groupName]`; The name of the group is the name specified in the group in first page of configuration wizard when installing a container. The second parameter is the machine list of the group that are allowed to access, it can be in the three formats ('*' means all the machines in the group, '-' means no machine is allowed, and [machine-name1,machine-name2,...,machine-nameN] as a list of machines that are explicitly allowed, the name of the machine is configured as the instance name field in the first page of configuration wizard when installing a container).

8.8 Host Name Conversion

The user can easily recognize the machine by seeing its hostname by clicking the Host Name from the view menu.

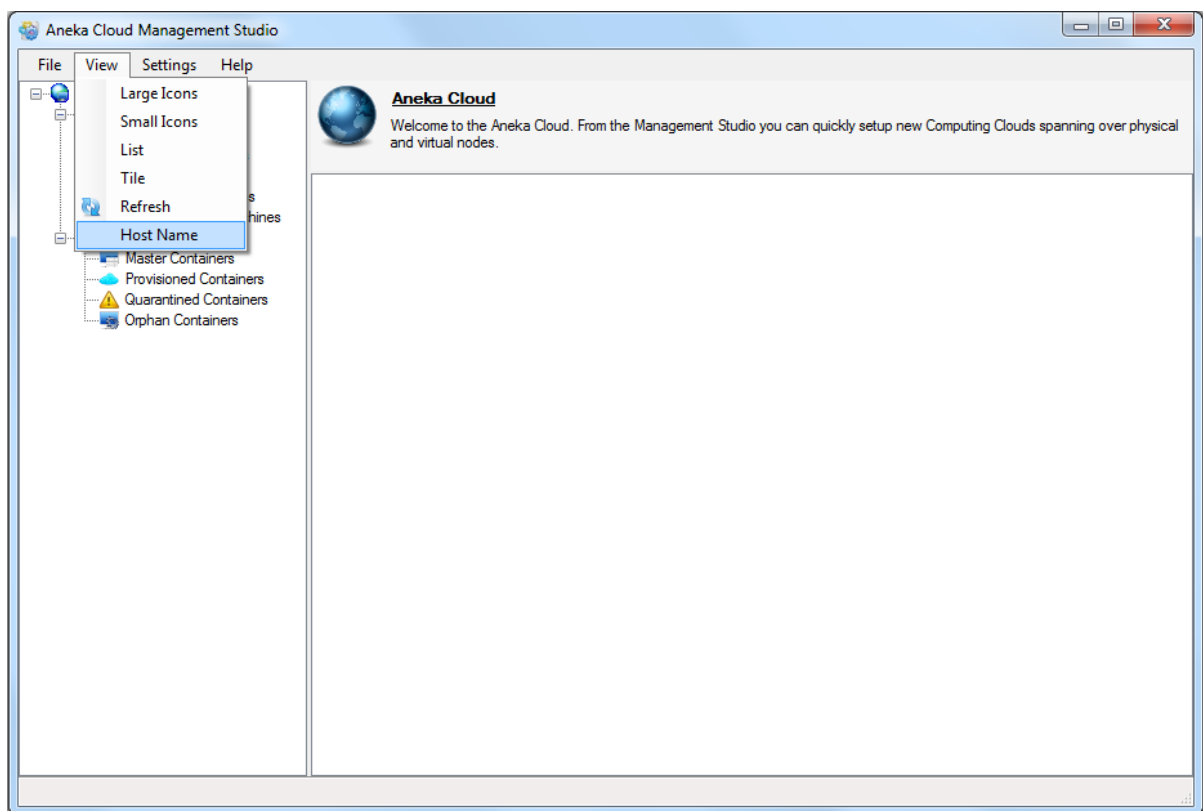


Figure 69 - Host Name option

Note: The Host Name option will only get displayed if and only if the host name is properly configured with the domain name in the system and user should be able to get the IP address from the hostname and the hostname from the IP address in the command prompt.

While clicking on the Host Name option under the view menu it will convert the daemon and all the containers's IP address into the hostname. Now the system will behave the same as it was working with IP address.

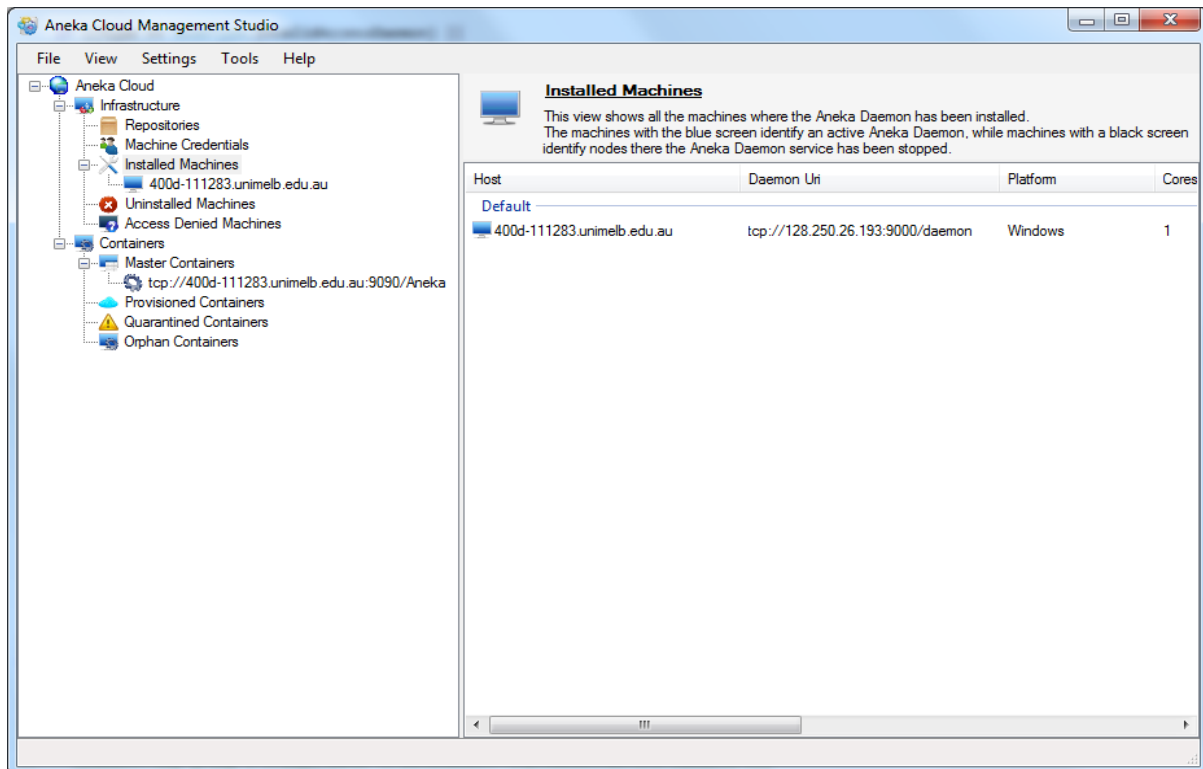


Figure 70 - Converted the IP address into Host Name

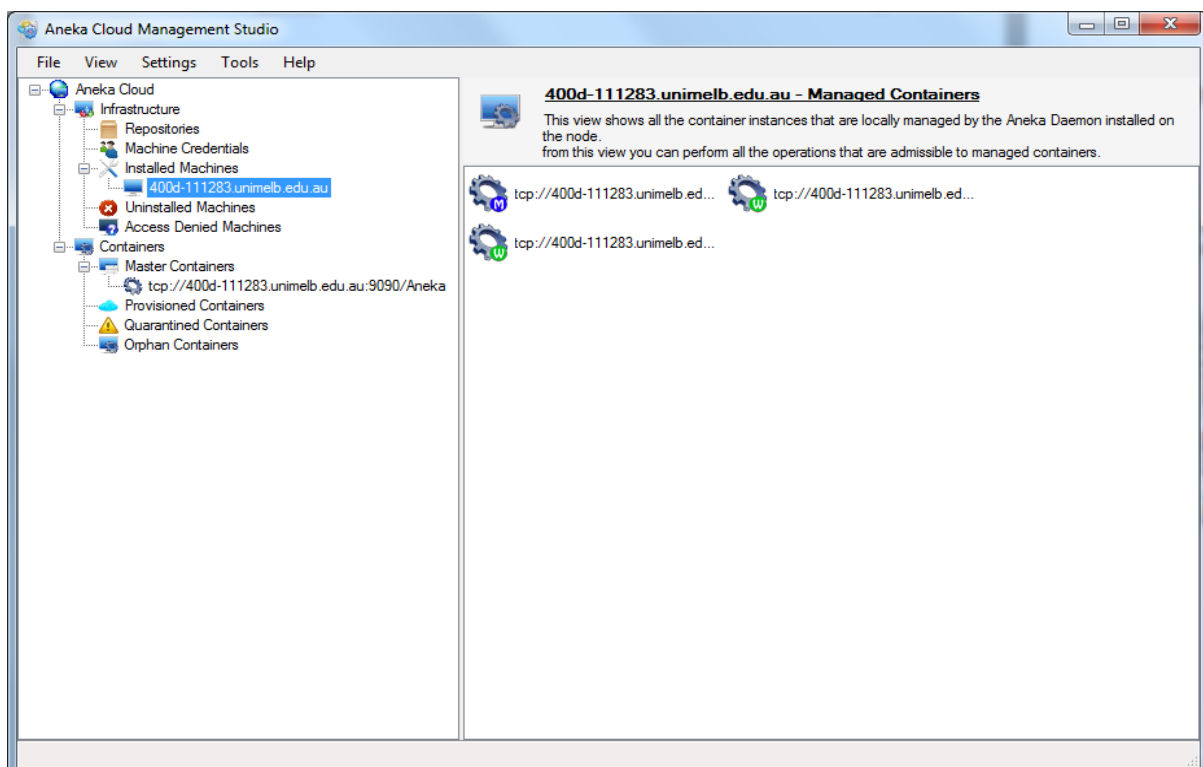


Figure 71 - Converted the containers and Daemon's IP address into Host Name

9 Testing the Cloud with Sample Applications

Once you have successfully deployed an Aneka Cloud, you can test it by running some of the sample applications shipped along with the distribution. These include samples for the Thread Programming Model, Task Programming Model, and MapReduce Programming Model. The following sections guide you through running some of the examples.

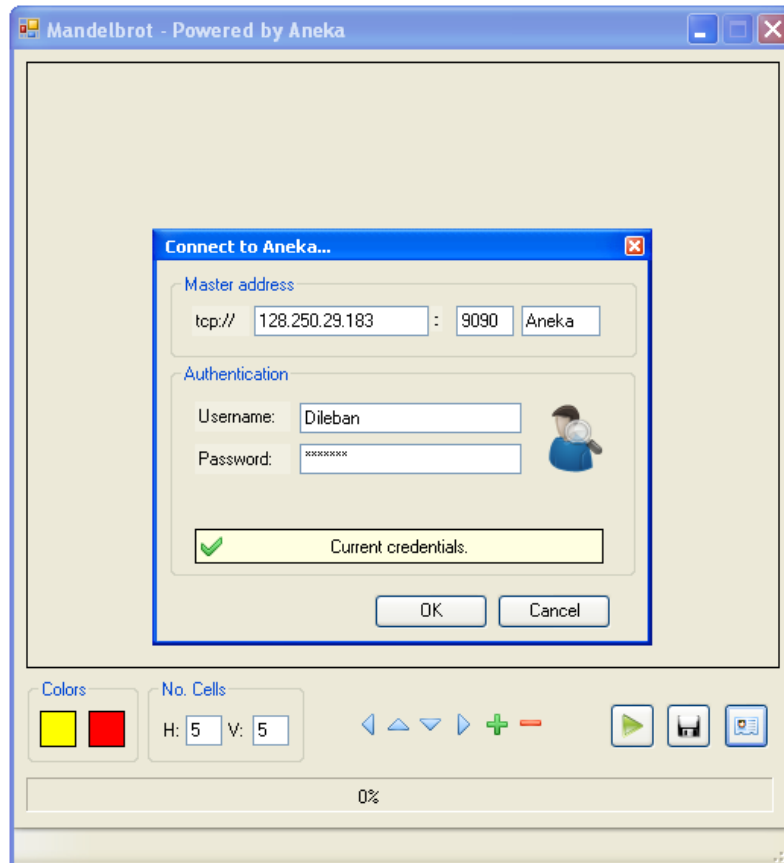


Figure 72 - Specifying the credentials for the Mandelbrot application

9.1 Thread Programming Model

The Mandelbrot sample application generates a rendering of the *Mandelbrot Set* in parallel. The Mandelbrot Set, named after the mathematician Benoît Mandelbrot, is a fractal formed by a set of points on a complex plane, determined using a mathematical function. This sample application was implemented using the Thread Programming Model.

Launch the Mandelbrot application from the start menu. To run the sample application, you must specify the address of the Master Container along with a valid Aneka user credential. If you have not setup user accounts yet, you may use the Administrator account with the default empty password. To run the application, click on the *Play* button. The figure below shows a sample run of the application.




Figure 73 - The Mandelbrot application running

9.2 Task Programming Model

Image convolution is the transformation of the pixels in an image using mathematical functions, to produce various results such as Gaussian blur, edge detection, and embossing. The sample convolution application was implemented using the Task Programming Model.

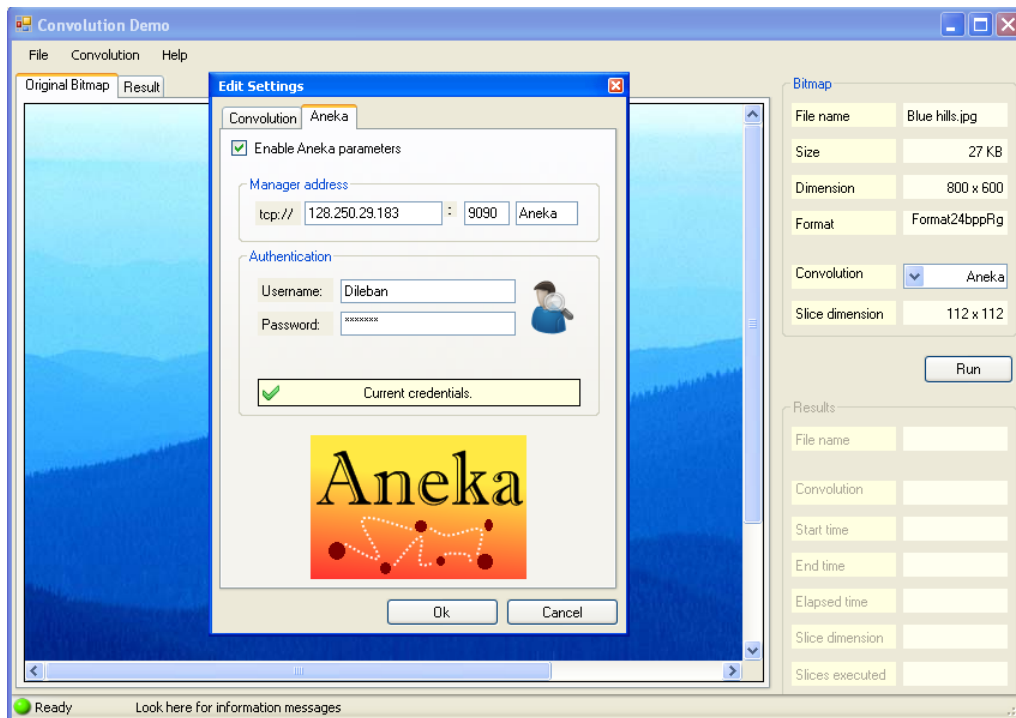


Figure 74 - Specifying the credentials for the convolution sample application

Launch the sample application from the start menu and specify the address of the Master Container along with a valid Aneka user credential. If you have not setup user accounts yet, you may use the Administrator account with the default empty password. To run the application, click on the *Run* button. The figure below shows a sample run of the application embossing the original image.

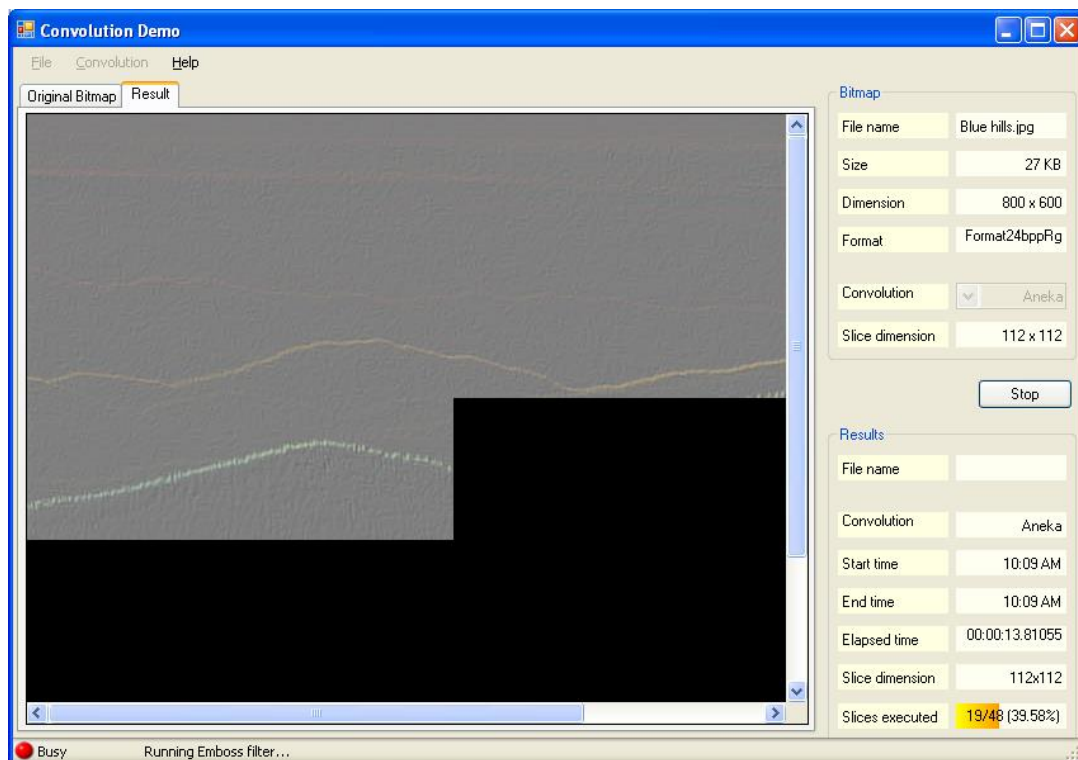


Figure 75 - The convolution example embossing the original image

9.3 MapReduce Programming Model

MapReduce is a programming model and an associated implementation for processing and generating large data sets. Users specify a map function that processes a key/value pair to generate a set of intermediate key/value pairs, and a reduce function that merges all intermediate values associated with the same intermediate key. For more details on MapReduce see the [original paper](#) from Google.

Aneka ships with two sample applications for the MapReduce Programming Model, the PiCalculator and WordCounter. Both these samples can be found in your Aneka installation directory in the following path:

<installation directory>\Aneka.3.0\Examples\Tutorials\Map Reduce Model\

This section will guide you through running the WordCounter example, but the process is similar for the PiCalculator.

The WordCounter example is a canonical application of MapReduce which counts the occurrences of each word in a given set of documents.

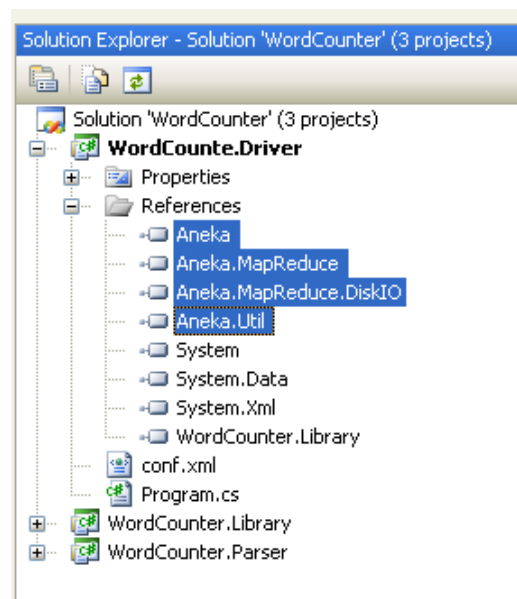


Figure 76 - Re-linking the libraries in the WordCounter example

To run the example, launch the Visual Studio Solution for WordCounter, and re-link the broken references to the Aneka libraries for all three projects as shown in the figure above. The Aneka libraries can be found in the SDK folder in your installation directory. When you are done, build the solution. Next edit the configuration file, *conf.xml*, in the debug folder for your WordCounter.Driver project, and change the *SchedulerUri* to the

address of your Master Container as shown Figure . You may also change the default user credentials to another Aneka user account if you wish to do so.

Next open the command prompt, and switch to the debug folder in WordCounter.Driver project folder and run the WordCounter example as shown in Figure . This program uses the sample Web pages, found in the *Workspace* directory in the same debug folder, to count the occurrences of each word. Once it completes running, it launches a results file as shown in Figure 7869.

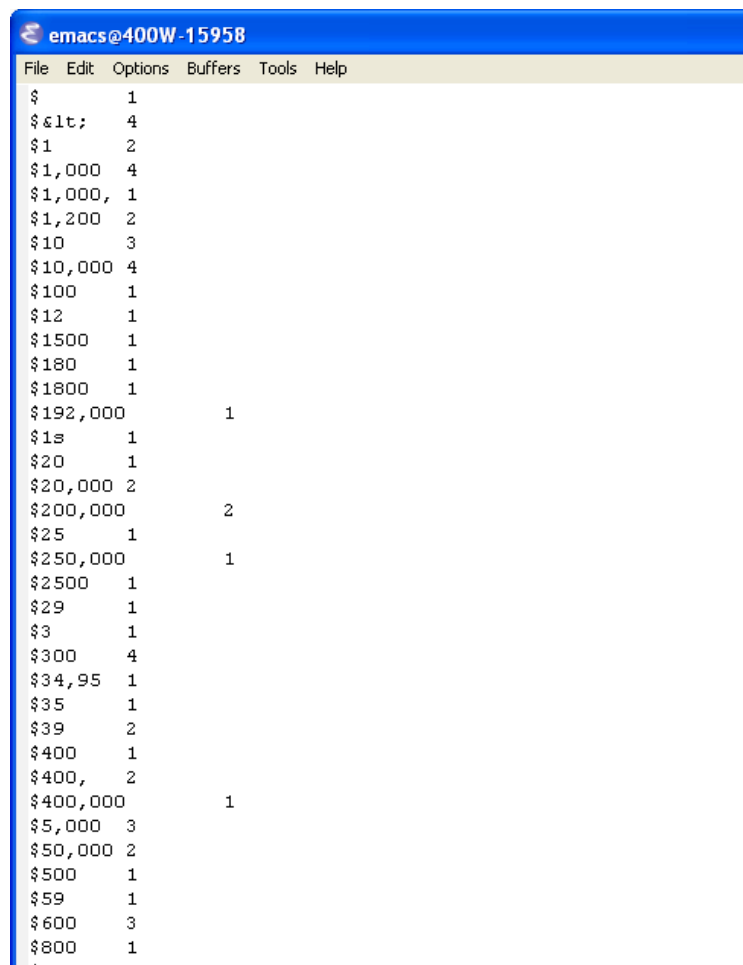
```
<?xml version="1.0" encoding="utf-8" ?>
<Aneka>
  <UseFileTransfer value="false" />
  <Workspace value="Workspace" />
  <SingleSubmission value="false" />
  <ResubmitMode value="AUTO" />
  <PollingTime value="1000" />
  <LogMessages value="true" />
  <SchedulerUri value="tcp://128.250.29.183:9090/Aneka" />
  <UserCredential type="Aneka.Security.UserCredentials" assembly="Aneka.dll">
    <UserCredentials username="Administrator" password="" />
  </UserCredential>

  <Groups>
    <Group name="MapReduce">
      <Property name="LogFile" value="WordCounter.log" />
      <Property name="FetchResult" value="true" />
      <Property name="UseCombiner" value="true" />
      <Property name="SynchReduce" value="true" />
      <Property name="IsInputReady" value="false" />
      <Property name="Partitions" value="1"/>
      <Property name="Attempts" value="3" />
    </Group>
  </Groups>
</Aneka>
```

Figure 77 - Editing the configuration file for the WordCounter example

```
WordCounter.Driver\bin\Debug>MapReduce.WordCounter.exe -c conf.xml
```

Figure 79 - Running the WordCounter example



File	Edit	Options	Buffers	Tools	Help
\$		1			
\$<		4			
\$1		2			
\$1,000		4			
\$1,000,		1			
\$1,200		2			
\$10		3			
\$10,000		4			
\$100		1			
\$12		1			
\$1500		1			
\$180		1			
\$1800		1			
\$192,000			1		
\$1s		1			
\$20		1			
\$20,000		2			
\$200,000			2		
\$25		1			
\$250,000			1		
\$2500		1			
\$29		1			
\$3		1			
\$300		4			
\$34,95		1			
\$35		1			
\$39		2			
\$400		1			
\$400,		2			
\$400,000			1		
\$5,000		3			
\$50,000		2			
\$500		1			
\$59		1			
\$600		3			
\$800		1			
----		.			

Figure 7869 - The results of running the WordCounter example

Appendix A - Service Configuration

TaskScheduler

The TaskScheduler service provides scheduling for the Task Programming Model. This service is typically installed on the Master Container.

Property	Description
AlgorithmSelection	Supports the following scheduling algorithms: <ul style="list-style-type: none">- FIFO- ClockRatePriority- FixedPriorityPreemption- TimeCostQoSOptimization- FixedQueueProvisioning- DeadlinePriorityProvisioning

ThreadScheduler

The ThreadScheduler service provides scheduling for the Thread Programming Model. This service is typically installed on the Master Container.

Property	Description
AlgorithmSelection	Supports the following scheduling algorithms: <ul style="list-style-type: none">- FIFO- ClockRatePriority- FixedPriorityPreemption- TimeCostQoSOptimization- FixedQueueProvisioning- DeadlinePriorityProvisioning

MapReduceScheduler

The MapReduceScheduler service provides scheduling for the MapReduce Programming Model. This service is typically installed on the Master Container.

Property	Description
Username	A Windows account name on the scheduling node to access the storage of MapReduce applications.
Password	The password for the account name given above.
StorageDirectory	The directory where all the required files (input and output files) for MapReduce applications are stored.

TaskExecutionService

The TaskExecutionService service executes tasks received from the TaskScheduler service. This service is typically installed on the Worker Container.

Property	Description
ExecutionDirectory	The name of the folder in the Worker's home directory for executing tasks. Each task will have its own execution directory below this directory for storing all files required to execute the task.
KeepFiles	If 'True', all files within the task's execution directory are <i>not</i> deleted when the task is completed. Setting this property to True is useful for debugging purposes.
MazimumTime	Sets the maximum time after which task executions will be terminated. This provides a mechanism to prevent infinite task executions that occupy the resources indefinitely.

ThreadExecutionService

The ThreadExecutionService service executes threads received from the ThreadScheduler service. This service is typically installed on the Worker Container.

Property	Description
----------	-------------

ExecutionDirectory	The name of the folder in the Worker's home directory for executing threads. Each thread will have its own execution directory below this directory for storing all files required to execute the thread.
KeepFiles	If 'True', all files within the thread's execution directory are <i>not</i> deleted when the thread is completed. Setting this property to 'True' is useful for debugging purposes.
MazimumTime	Sets the maximum time after which thread executions will be terminated. This provides a mechanism to prevent infinite thread executions that occupy the resources indefinitely.

MapReduceExecutor

The MapReduceExecutor service executes jobs received from the MapReduceScheduler service. This service is typically installed on the Worker Container.

Property	Description
Username	A Windows account name on the executing node to access the storage of MapReduce applications.
Password	The password for the account name given above.
StorageDirectory	The directory where all the required files (input and output files) for MapReduce applications are stored. NOTE: If the MapReduceExecutor is configured to run on the same node as the scheduler, these directories must be different.

StorageService

Supports file transfers between Clients and the Aneka Cloud. This service is typically installed on the Master Container. You may alternatively host the Storage Service on a separate Container to reduce the load on the Master. Make sure you have sufficient disk space, at least in the order of a few gigabytes, on the machine hosting the Storage Service.

Property	
----------	--

	Description
DbConnectionString	The connection string for connecting to a database (this value is automatically set choose to use a Relational Database for persistence)
DbProvider	The provider for the database connection (this value is automatically set choose to use a Relational Database for persistence).
FileTransferProtocol	The different transfer protocols. Currently FTP is supported.
UseRDBMS	Set to 'True' if the Storage Service is to use a Database to store information.

ResourceProvisioningService

Dynamically provisions resources from public and private cloud providers such as Amazon, GoGrid and XenServer. This service is typically installed on the Master Container.

Property	Description
Option	None or Managed. If the Managed option is selected, users can specify different pool as shown in the figures below

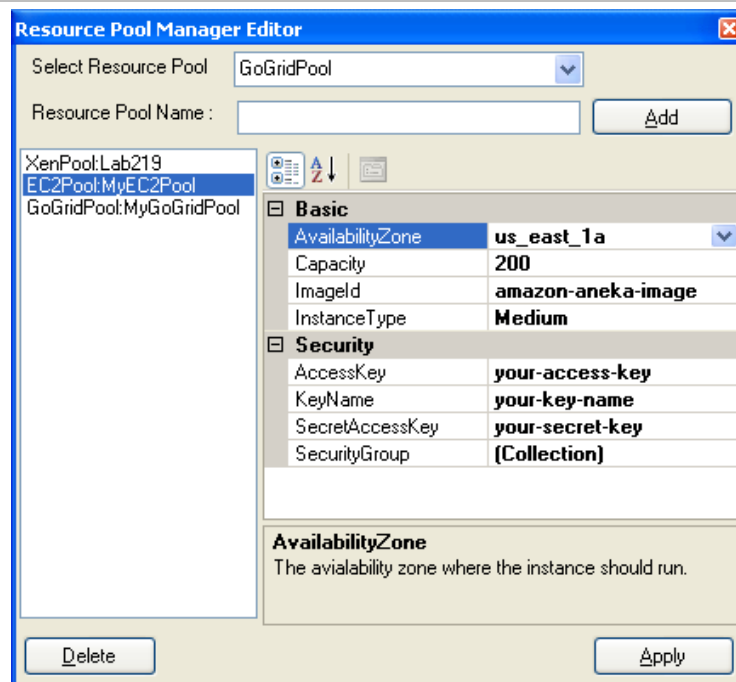
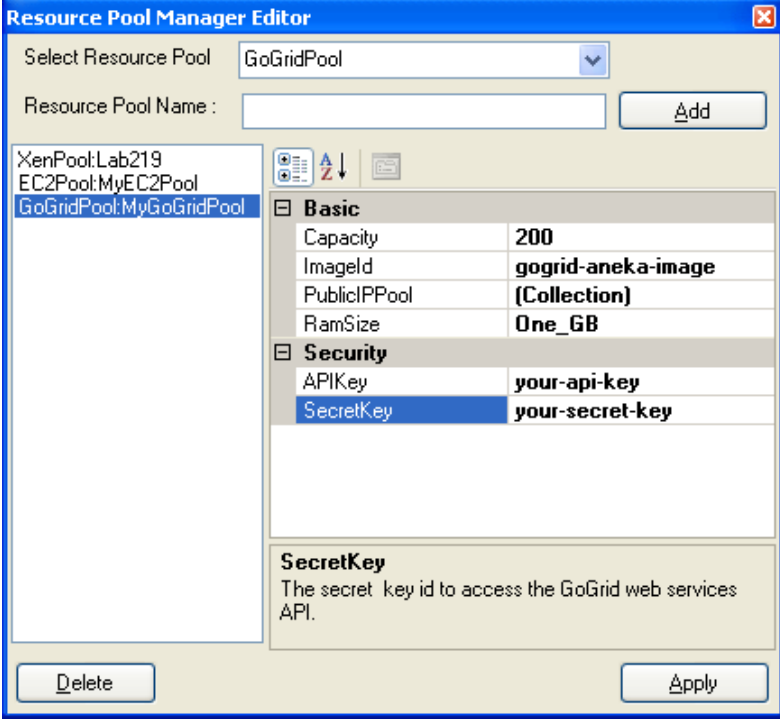


Figure 80: Configuration Wizard -Amazon Pool Configuration



Resource Pool Manager Editor

Select Resource Pool: GoGridPool

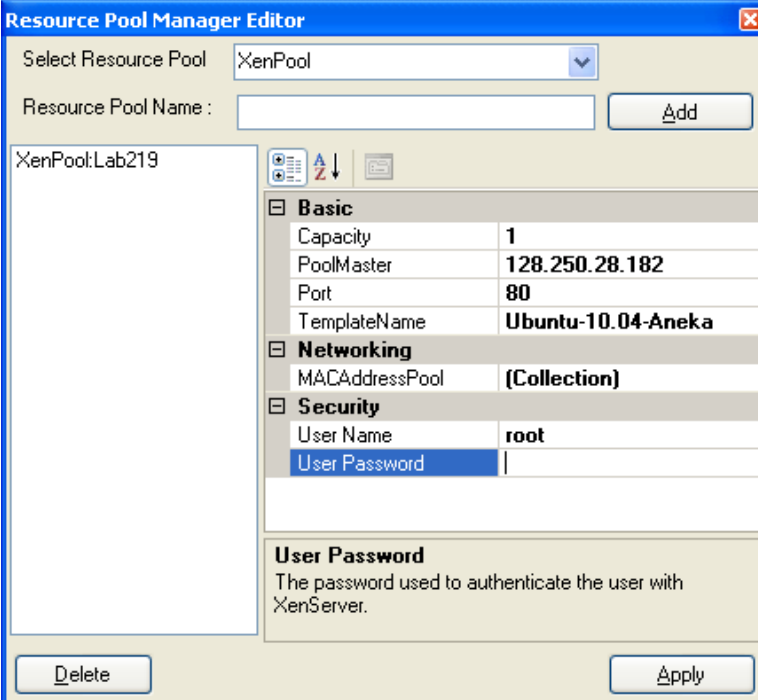
Resource Pool Name:

XenPool:Lab219
 EC2Pool:MyEC2Pool
GoGridPool:MyGoGridPool

Basic	
Capacity	200
ImageId	gogrid-aneka-image
PublicIPPool	(Collection)
RamSize	One_GB
Security	
APIKey	your-api-key
SecretKey	your-secret-key

SecretKey
The secret key id to access the GoGrid web services API.

Figure 81: Configuration Wizard - GoGrid Pool Configuration



Resource Pool Manager Editor

Select Resource Pool: XenPool

Resource Pool Name:

XenPool:Lab219

Basic	
Capacity	1
PoolMaster	128.250.28.182
Port	80
TemplateName	Ubuntu-10.04-Aneka
Networking	
MACAddressPool	(Collection)
Security	
User Name	root
User Password	<input type="password"/>

User Password
The password used to authenticate the user with XenServer.

Figure 82 - Configuration Wizard - XenServer Pool Configuration

ReportingService

Collects monitoring data from Worker Containers and provides access to various reporting information. You must enable this service in order view reporting information through the Aneka Management Studio. This service is typically installed on the Master Container.

Property	Description
DbConnectionString	The connection string for connecting to a database (this value is automatically set choose to use a Relational Database for persistence)
DbProvider	The provider for the database connection (this value is automatically set choose to use a Relational Database for persistence).
FlushPeriodInMinutes	The period after which all reporting information will be flushed. Use this value carefully in order to ensure that database (or disk) storing the reporting data does not grow indefinitely.
ReportingFormat	The format of the reporting data. The supported formats include XML and Binary. The default recommended format is XML. Use binary if you need to save storage space.
UseRDBMS	Set to 'True' if the Reporting Service is to use a Database to store information.

MonitoringService

Collects and reports various monitoring data to the Reporting Service. You must enable this service to have access to reporting information such as CPU performance and storage utilization (file transfers). This service is typically installed on both the Master and Worker Containers.

Property	Description
ReportingServiceUri	The Uri of the ReportingService. If the ReportingService is hosted on the Master Container, use the Uri of the Master Container.

LoggingService

Provides remote access to a Container's logging data. This service also allows you to access the Container's log archives, and can be useful for troubleshooting problems. This service is typically run on both, the Master and Worker Containers.

Property	Description
CacheStore	The location where the logging entries cache will be stored.
Format	The format of the logging entries that will be stored
Level	The logging level which the logging entries will be recorded including Debug, Info, Warn, Error and All.