



MS Windows Adaptive Agent For Windows 2008



Copyright © 2000 - 2009 KEMP Technologies, Inc. All rights reserved.

KEMP Technologies, Inc. reserves all ownership rights for the LoadMaster product line including software and documentation.

The use of the LoadMaster Load Balancer is subject to the license agreement.

Information in this guide may be modified at any time without prior notice.

Company names used in examples are fictitious unless otherwise noted.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective owners.

International copyright law protects this document. All rights are reserved. Copyright in the content, in the screens displaying the content, and their arrangement, is owned by KEMP Technologies, Inc. unless otherwise indicated. The content and data found on this document may not be redistributed, copied, sold, transferred, or modified, without the express written permission of KEMP Technologies, Inc. De-compilation, reverse engineering or disassembly of document is strictly prohibited.



Preface

This guide demonstrates how to install the KEMP Technologies, Inc. MS Windows Adaptive Agent with Internet Information Server 7.0, the adaptive agent is not coupled to IIS 7. You can install the agent with a web server of choice as long as it supports Common Gateway Interface (CGI).

Overview

A KEMP LoadMaster determines how to best distribute load among available real servers by a variety of methods. These can include the availability of the servers (by doing health checking), as well as both active connections and connection rates to the real servers.

It may be beneficial however to also have information on the performance characteristics of the real servers themselves. This is possible through the “adaptive” load balancing method in the load balancer. In conjunction with an agent, the LoadMaster can distribute traffic based on the load on the actual server, which can include CPU load, disk load, and other performance metrics.

The LoadMaster Server Agent works by residing on a server and reporting back to the LoadMaster on how loaded the server is. The metrics used to determine the system's load are entirely configurable by the user.

The LoadMaster Agent simply reports a number, between 1 and 99, with 1 being the most idle, and 100 being the most busy. If the reported value is at 100, the agent will automatic report back 101. The 101 value is treated as a failed health check and will remove the server from the pool. The number is determined by the parameters specified in the configuration file. It can include only one performance metric (such as CPU utilization), or be a combination of up to eight different metrics (such as CPU utilization, memory utilization, disk I/O).

Pre-requisite

1. Transfer the agent to the server
2. Administrative access to the server

Outline

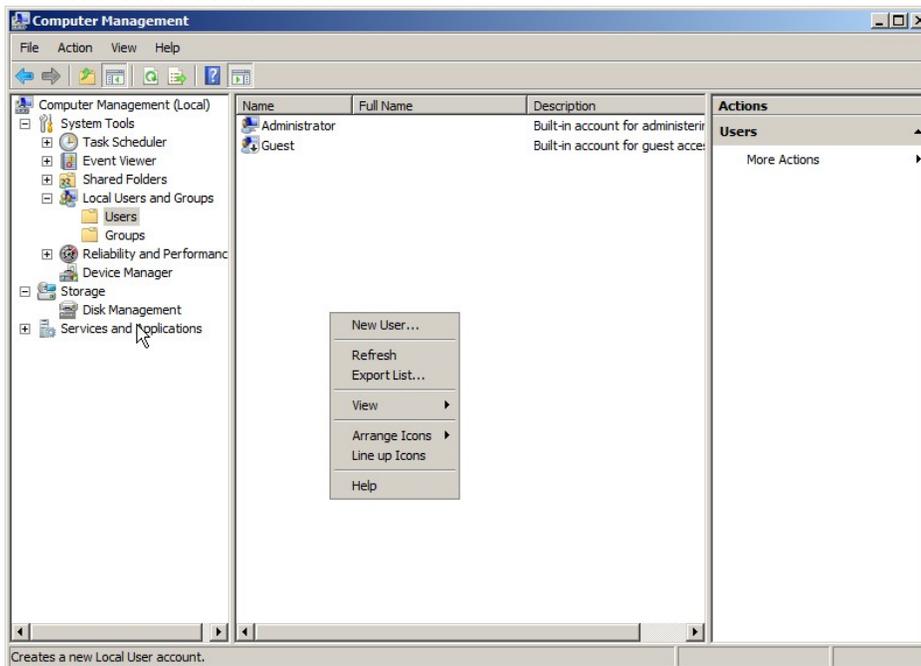
1. Create a “Imagent” user account
2. Install MS Windows Adaptive Agent
3. Install web server (IIS 7 default)
4. Configure web server
5. Configure MS Windows Adaptive Agent
6. Test/Troubleshoot
7. Repeat on each Real Server
8. Configure LoadMaster

Create MS Windows Adaptive Agent user account

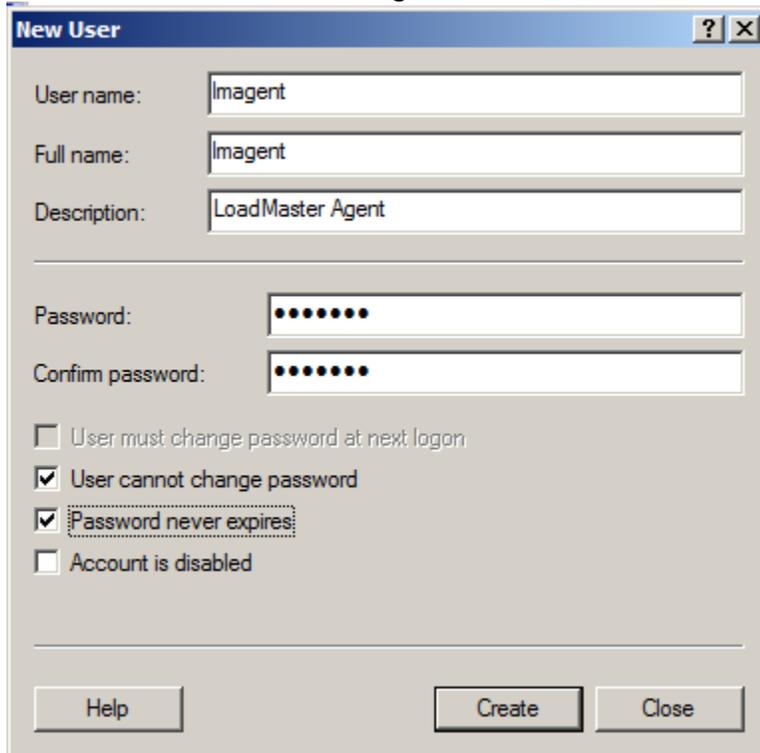


Create a dedicated user for executing the MS Windows Adaptive Agent. Create a user named “Imagent” and make sure the user is a member of the “Performance Monitor Users”.

1. Login as a Administrator
2. From the “Start” menu select the “Run” option. Enter the “compmgmt.msc” and click “Okay”
3. Expand “Local Users and Groups”. Select “Action” from application menu and then “New User”



4. Set the User Name as “Imagent”



The screenshot shows a 'New User' dialog box with the following fields and options:

- User name: Imagent
- Full name: Imagent
- Description: LoadMaster Agent
- Password: [masked]
- Confirm password: [masked]
- User must change password at next logon
- User cannot change password
- Password never expires
- Account is disabled

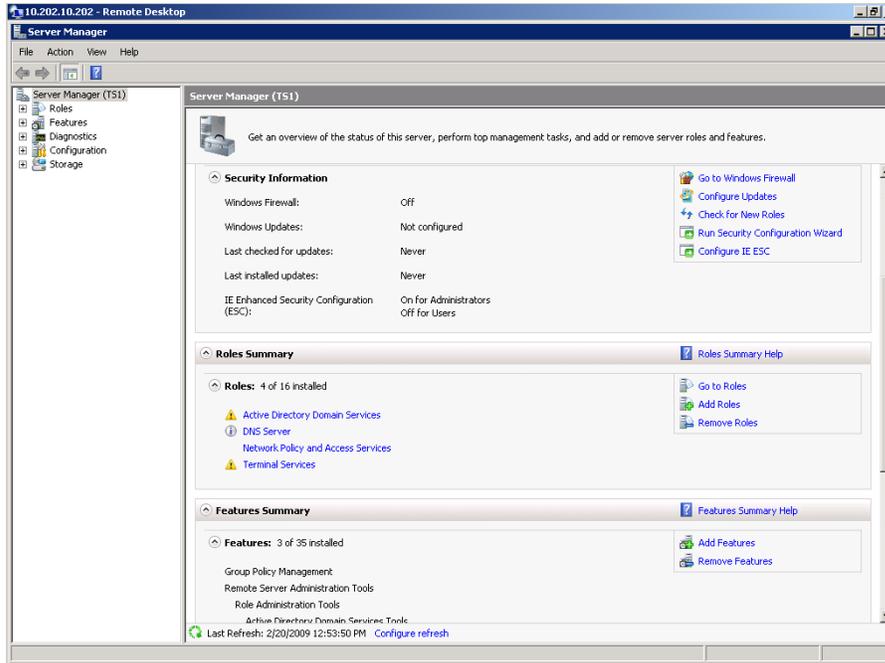
Buttons at the bottom: Help, Create (highlighted), Close.

5. Once the “Imagent” user has been created, select the properties. Add the user to the “Performance Monitor Users” group.

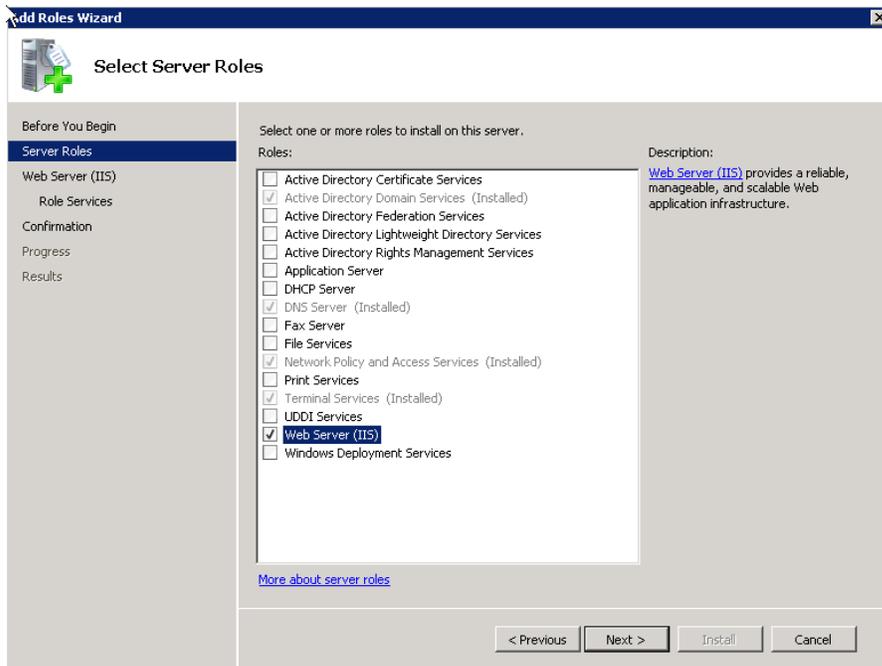


Install Internet Information Server 7

As Administrator install Internet Information Server 7 following Microsoft guidelines. Start by accessing the “Server Manager” console. From the “Start” menu select the “Run” option. Enter the “CompMgmtLauncher” and click “Okay”. You can start the process using the Windows 2008 Server Manager, the select the “Add Roles”.

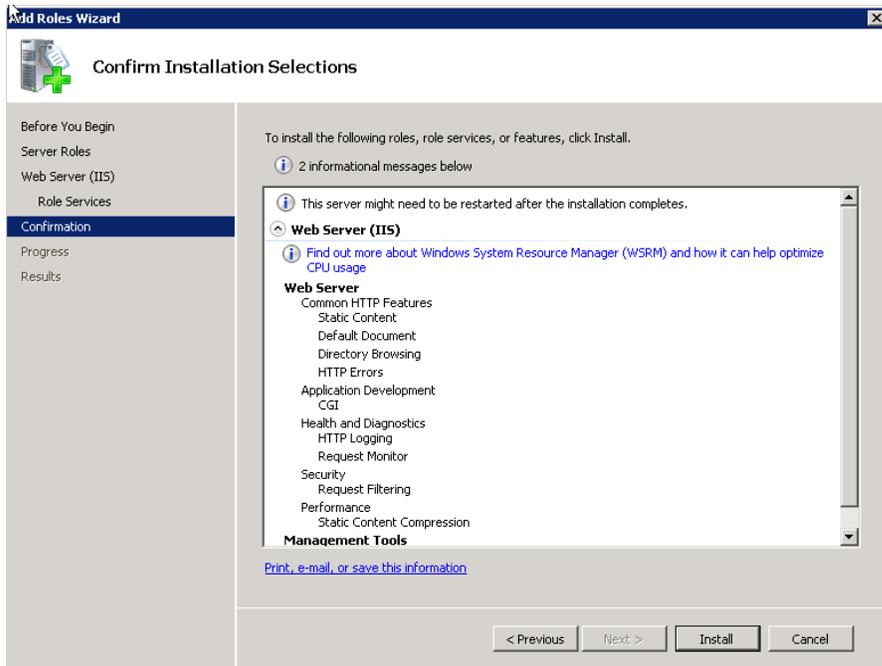
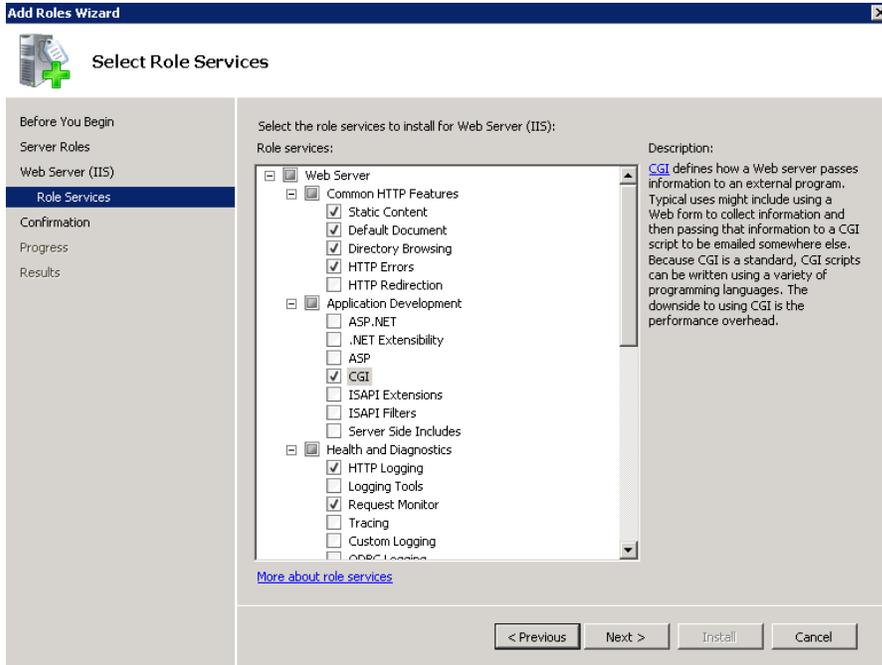


Check the “Web Server (IIS)” check box and select the “Next >” button.



Make sure to install the “CGI” support by checking the box, look under “Application Development” and select the “Next >” button.





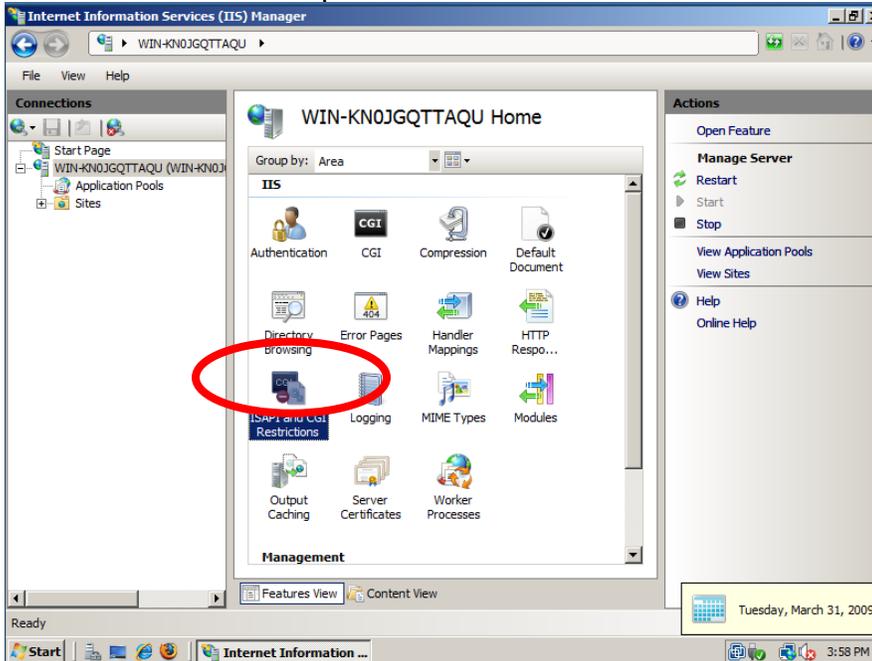
Verify your installation of IIS by browsing <http://localhost>

Install MS Windows Adaptive Agent

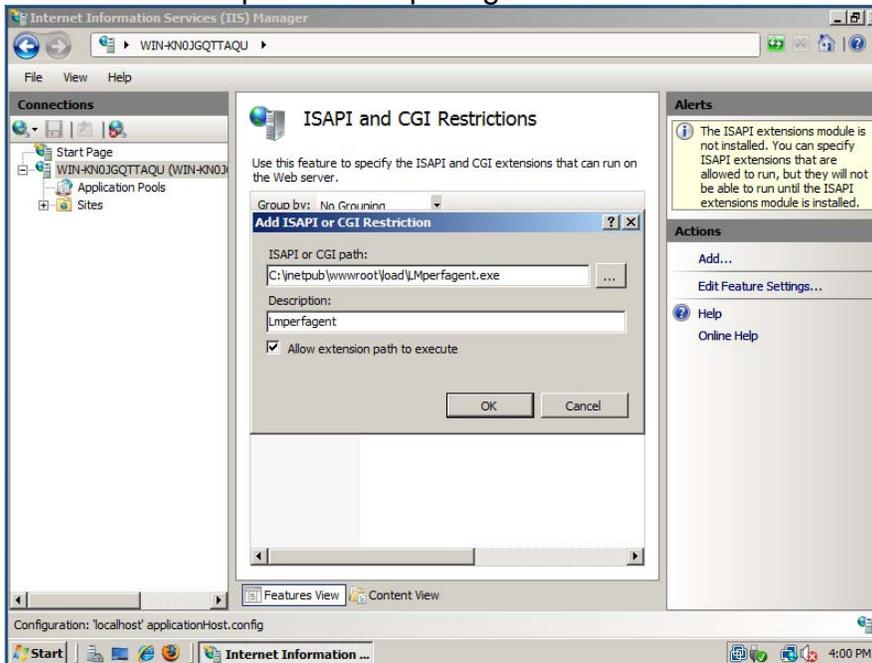
1. Login as a Administrator
2. Create a folder named "load" in the default IIS web root, commonly C:\inetpub\wwwroot



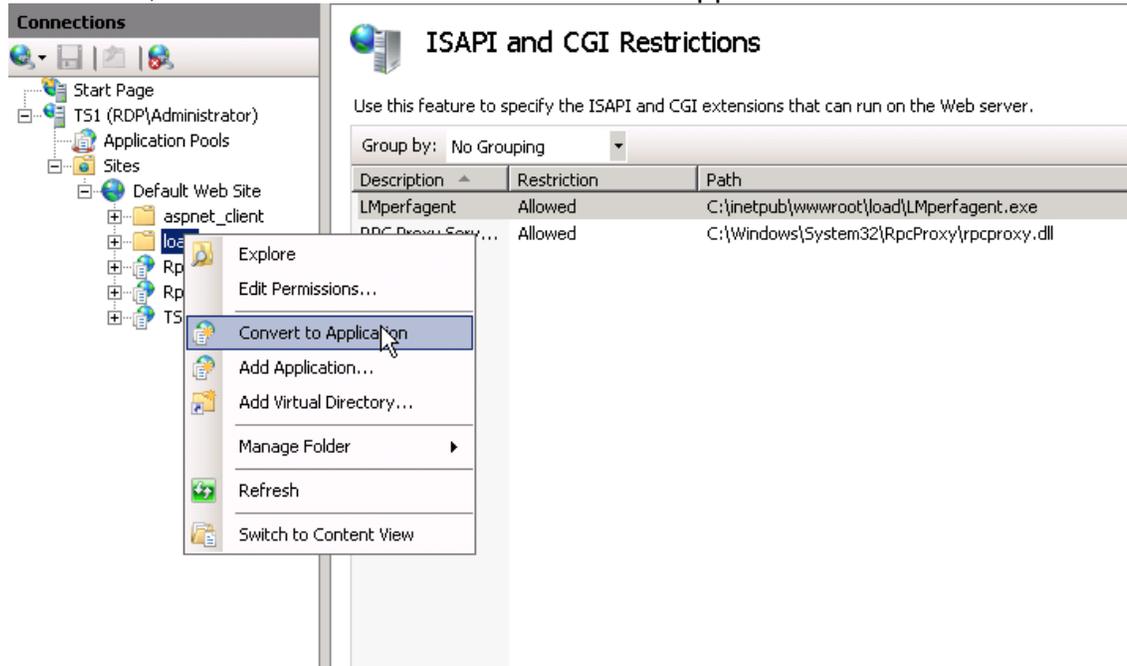
3. Uncompress the agent files; copy only Lmperfagent.exe and LMperfagent-config.txt into the “load” directory.
4. Start Internet Information Services (IIS) Manager
5. Select the IIS Group “ISAPI and CGI Restrictions”



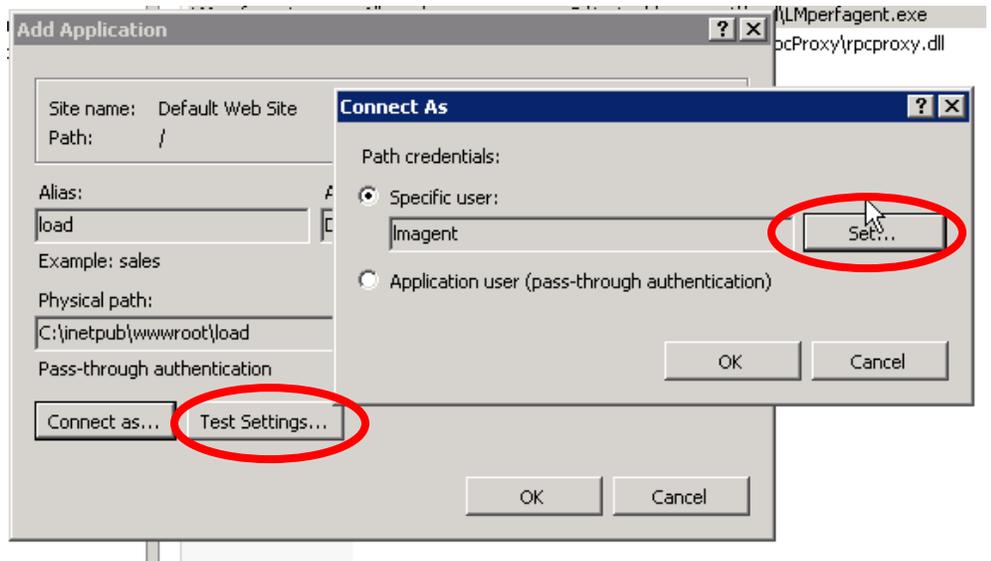
6. Under Actions add a new extension, select the Lmperfagent.exe file and set the description to “Lmperfagent”.



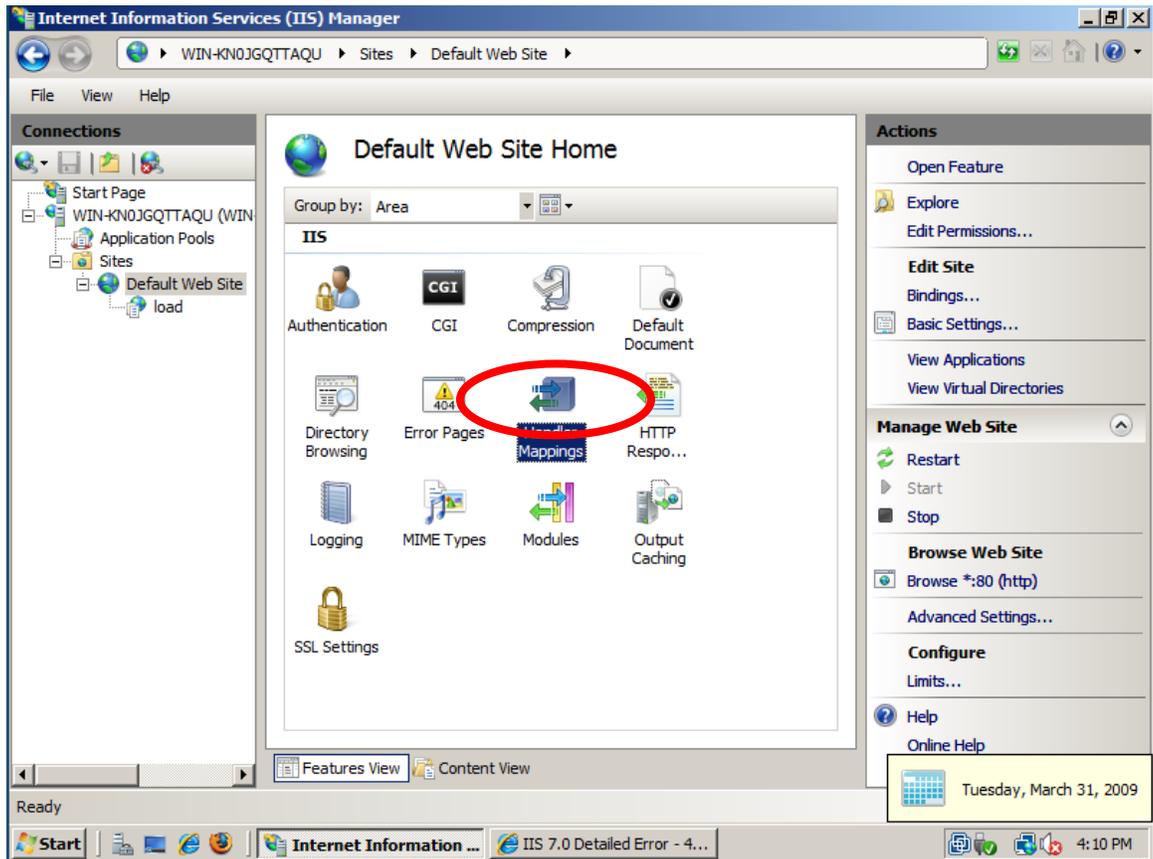
7. Locate the “load” directory under the tree navigation under the “Default Web Site”, select the folder and convert it to an application.



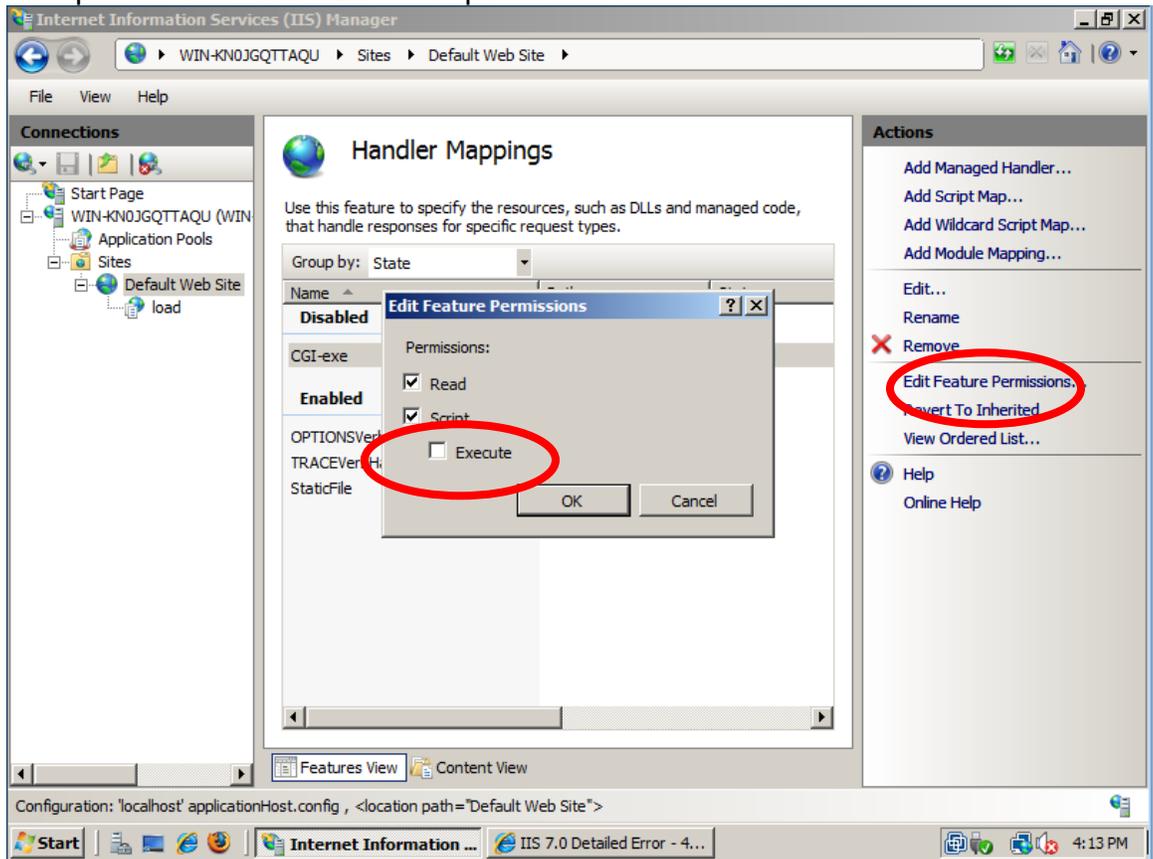
8. When adding the new application make sure to configure the “Connect as” to be use the “Imagent” user. (“Imagent” user should have been created as part of the pre-requisites) Verify the authentication is correct by using the “Test Settings” feature.



9. Locate the “Default Web Site” and then select the “Handler Mappings” feature.



10. Select the “CGI-exe” mapping and then select “Edit Feature Permission ...” option and enable “Execute” permission.



11. Test using <http://localhost/load/LMperfagent.exe> the web page will display a value between 0 and 99



Configure LoadMaster

In the LoadMaster Web User Interface, select the Rules and Checking link in the left navigation, and then the “Check Parameters” menu. In there, you'll find the configuration options for adaptive metrics.

The screenshot shows the LoadMaster Web User Interface. On the left is a navigation menu with the following items: Home, Virtual Services, Real Servers, Rules & Checking (expanded), Content Rules, Check Parameters (highlighted), Statistics, Certificates, and System Configuration. The main content area is titled 'Adaptive Parameters' and contains the following fields:

Adaptive Interval (sec)	7
Adaptive URL	/load <input type="button" value="Set URL"/>
Port	80 <input type="button" value="Set Port"/>
Min. Control Variable Value (%)	5
Min. Weight Adjustment Value (%)	5

Below these fields is a button labeled 'Reset values to Default'. The next section is titled 'Service Check Parameters' and contains the following fields:

Check Interval(sec)	14
Connect Timeout (sec)	7
Retry Count	2

Below these fields is another button labeled 'Reset values to Default'.

Adaptive Interval determines how often the adaptive value is pulled. The Adaptive URL tells the LoadMaster where to find the value, and port specifies which port the web server is answering on. If your webroot for “/” is C:\inetpub\wwwroot\site1\, and your agent located in C:\inetpub\wwwroot\load\LMperfagent.exe”, then you would change Adaptive URL to /load/LMperfagent.exe.

When you enter the Adaptive URL, hit enter, and the changes will be saved. The Min. Control Variable Value is a percentage that specifies a threshold below which the balancer will switch to static weight-based scheduling, i.e. normal Weighted Round Robin. The value is a percentage of the maximum load, but maximum value permitted is 50. The default is 5. The Min. Weight Adjustment Value is a percentage value is 10, then the weight will never be set to below 10. A value of 5 is recommended, and the maximum value is 50.

Virtual Service Configuration

Once the global settings are configured and the agent installed on the real servers, you can enable adaptive load balancing on a per-virtual service basis. In the properties for each virtual service, you'll see an option for scheduling method. **Select “resource based (adaptive)” to enable adaptive load balancing.**

The screenshot shows a dropdown menu for 'Scheduling Method' with the option 'resource based (adaptive)' selected.



Troubleshooting

Common pitfalls arise around permission. Login with the “Imagent” user and try running the agent via the command line.

Start the Command Windows and change to the “load” directory. Execute LMperfagent.exe

Expected value is between 0 and 101.

1. Normal range 0 to 99
2. Disable Real Server 101

Potential Error Codes:

- 1 = Unable to open local LMperfagent-config.txt file
- 2 = Unable to open performance query
- 3 = Error collecting performance query data
- 4 = Error closing performance query

Check the memberships of the “Imagent” user and verify the access level is high enough. Make sure the configuration file for the agent can be read by the “Imagent” user and is located in the “load” directory.

Configuration File

If you look at the configuration file (either the txt file in a text editor, or the xls file in a spreadsheet editor), you'll see that each metric has its own line, and each line consists of five different fields.

The first field is the weight that the specified metric has on the overall performance value reported. The second field is the top range of the value reported by that metric. For example, if you were to measure system memory in terms of megabytes, and the system had 2 GB of internal memory, the maximum value would be 2048. If the memory utilization was 1024, then the value reported would be “50”, since the memory is 50% utilized.

The third field is whether the value is an up or down value. A good example to illustrate the difference between an up or down metric would be a metric that reports memory. If the metric reports memory available, then it would be a “down” metric. As the metric value goes down, the value reported would go up. A system with 2048 MB of RAM that reports 0 memory available would be a fully utilized system, so it would report 101.

If instead the metric reported memory utilized, a value of 2048 MB of RAM would mean no RAM was being used, and the value reported would be 1. If the memory utilized shot up to 2048, then the value reported would be 101. The fourth field is the hard limit value. Setting a value will cause the agent to override all other metrics, and ignore their weights, and report a 100% loaded



condition if this threshold is reached. An example would be for CPU utilization. If you have CPU weighted to 50%, and you have disk I/O and pages per second each weighted at 25%, you can have a situation where disk I/O and pages per second are nearly at zero, but CPU hits 100%. In that situation, the agent would report 50% utilization to the LoadMaster. If you had a hard limit set for 100 for CPU, if CPU hit 100, then the agent would report 101.

And the fifth field is the metric itself. A list of available metrics can be found in the example configuration file. Additional information is available in the README file

