1 RAC Server Virtualization

1.1 Windows 2008 Hyper-V

1.2 Introduction

Remote Audio Compressor Server is used by Interaction Recorder to move and Compress Audio files. This takes stress off the xIC Server as the RemoteAudioCompressor threads will routinely use 100% of the CPU during compression. This is not real-time audio processing so using Hyper-V was possible and did not affect the audio quality of the compressed files.

1.3 Setup and Configuration

No Reservation needs to be made for a RAC on the Hyper-V Server Host. We gave our RAC 1 CPU and 1024MB of Ram but going beyond this would be fine. In IA we assigned 3 Threads via the configuration under Server Configuration\Audio Compression.

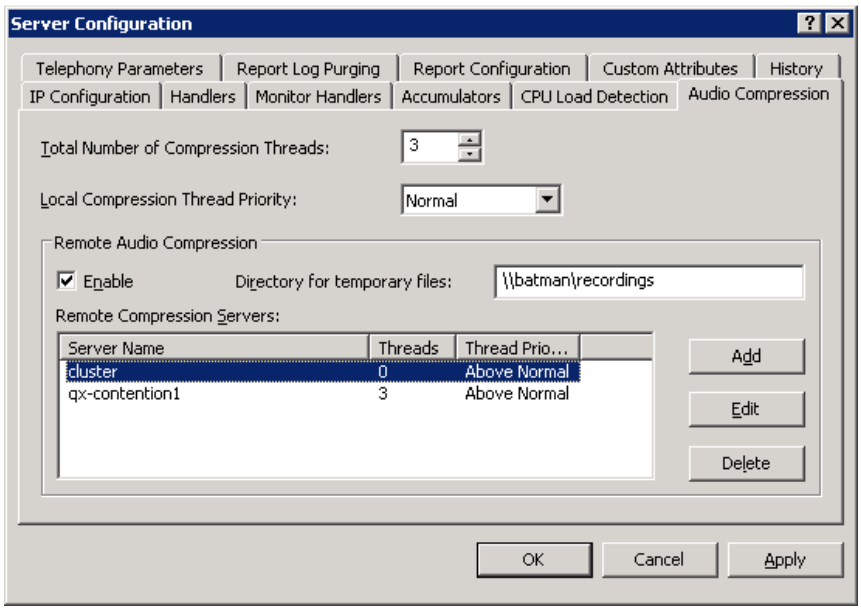

1.4 Test Results

We compared a Hyper-V Server guest with a dedicated RAC Server (Real Machine). We ran 1000 recordings though the RAC and then scored the Audio Files with an automated process. We had a number of competing Hyper-V Server guests that purposely took the CPU of the Hyper-V Server Host to 99% CPU as to starve the RAC of CPU time.

We were able to prove that the CPU contention from other machines running on the Virtual Machine Host would not cause problems with the audio processing.

Testing Results with 3.0 SU3 on Hyper-V 64bit Enterprise

RAC Server compresses to TrueSpeech

Windows 2003 R2 Standard Server Guest

Cps	Reservation?	Avg MOS	LOW	HIGH	Number of Recordings Rec/Sent
4	No, and Contention	3.51877	3.39501	3.72434	1000/1000

Edit

RAC Server compresses to TrueSpeech with Encryption

Windows 2003 R2 Standard Server Guest

Cps	Reservation?	Avg MOS	LOW	HIGH	Number of Recordings Rec/Sent
4	No, and Contention	3.51891	3.39023	3.72964	1000/1000

Edit

Testing Results with RAC Server 3.0 SU3 on a Real Machine

RAC Server compresses to TrueSpeech codec

Windows 2003 R2 Standard Server (3.0Ghz HT CPU with 1024MB RAM)

Cps	Reservation?	Avg MOS	LOW	HIGH	Number of Recordings Rec/Sent
4	N/A	3.52261	3.39249	3.72964	1000/1000

1.5 Additional Information

The results above are a concise summary of many test runs and other analysis of performance.

Audio Quality Measurement

Audio quality verification is performed using two different algorithms. The PESQ algorithm (ITU Recommendation ITU-T P.862) is based off of a known good recording and all recording are compared to this base recording. The 3SQM algorithm (ITU Recommendation ITU-T P.563) is based on models of human vocal tract and the human perception of abnormalities on a vocal signal. Note that the PESQ score is on a -0.5 to 4.5 scale. All PESQ scores have 0.5 added to them before Mean and Standard Deviation calculations to provide a more direct comparison to 3SQM scores.