



Tuning WebHound™ 4.0 and SAS® 8.2 for Enterprise Windows Systems

Martin Powdrill

ES7000 Performance Center

Unisys Corporation

Overview

- **Determine maximum “scalability” of WebHound™ application on Enterprise Windows System (ES7000).**
- **Provide Commodity / Enterprise System Comparisons**
- **Document results and provide sizing guidelines as a result of the benchmarking activities**

SAS is a registered trademark or trademark of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or Trademarks of their respective companies.

Scalability

“Achieving Improved Performance Benefits when Increasing System Capacity (CPUs, I/O, Memory)”

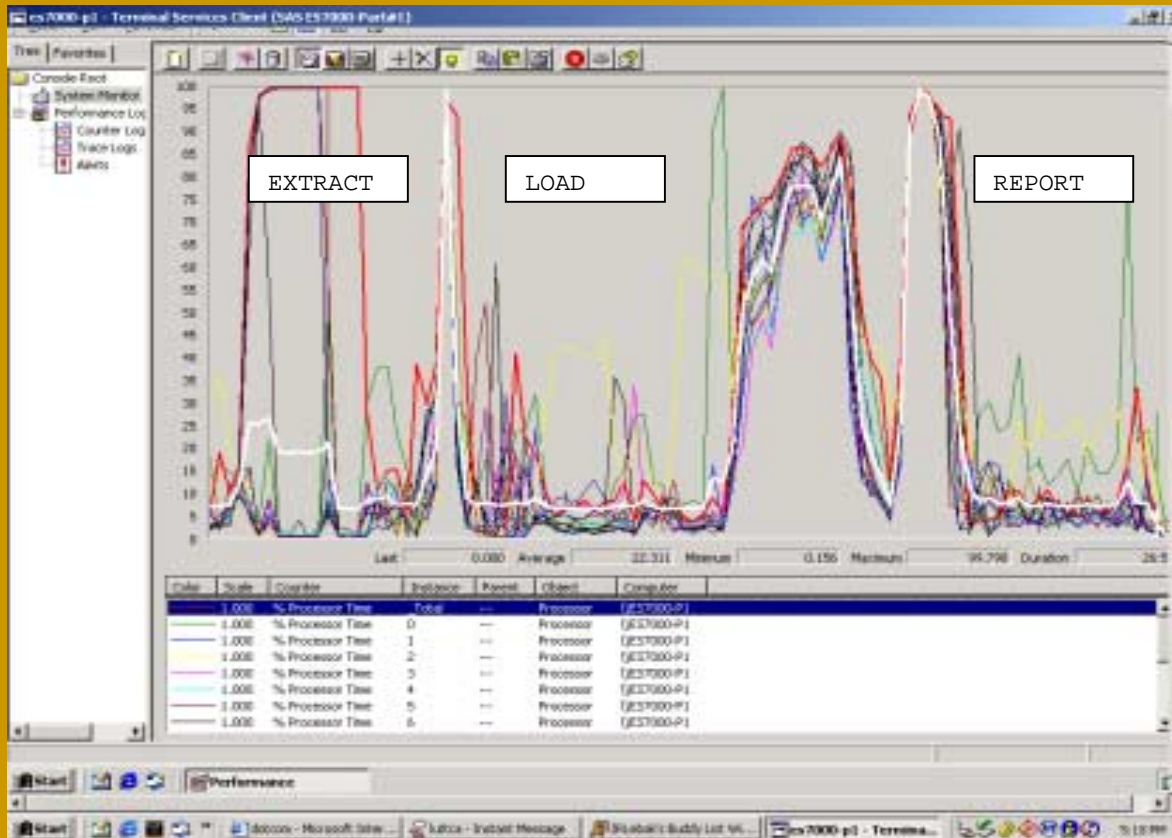
Topics to be Covered

- **WebHound™ Architecture and Applicability to Enterprise Systems**
- **Benchmarking Testing Description and Methodology**
- **Results achieved**
 - **Commodity .vs Enterprise Servers**
- **Sizing Guidelines for Windows Enterprise Systems**

WebHound™ Architecture

- **Batch-oriented Processing**
- **Command-driven**
- **3 distinct processing phases**
 - **Extract**
 - **Load**
 - **Report**
- **Parallel Processing Capability**
 - **SAS/CONNECT® w/ MP Connect Feature**
 - **Provides potential for “scalability” results**
 - **Parent/Child relationship provide concurrent SAS sessions**

WebHound™ MP Connect Processing



Testing Methodology

- Define Workload(s)
- Determine Key Performance Metric(s) to Measure
 - Total Processing Time
 - *Extract, Load and Report Phases*
- Establish Baseline for Comparisons
 - Unisys 8 CPU Commodity Windows Server
 - ES5085 Server
- System and Application -Level Tuning Only (No Code Changes!)

Workload Characterizations

9 different workloads varied by:

1. **Number of days to process**

- One Day
- Seven Days
- Thirty Days

2. **Web Log**

- Sizes and Number of files to process

Web Log Data Input Sizes: (Each one a workload)

	<u>1 DAY</u>	<u>7 DAY</u>	<u>30 DAY</u>
SMALL	41MB	798MB	3.13GB
MEDIUM	219 MB	2.61GB	11.0GB
LARGE	2.9 GB	17.3 GB	40.1 GB

Baseline System Configuration

8 CPU Commodity Server

- **Unisys ES5085 Server**
 - 8 700 MHZ XEON Processors
 - 8 GB RAM
 - 2 I/O Channels
 - Emulex Fibre channel controllers
 - Disk Subsystem
 - 1 ESM7800 EMC Clariion
 - 20 36 GB Drives (720 GB) configured as RAID 1/0
- **Windows 2000 Advanced Server SP2**

Target System Configuration

ES7000 Enterprise Server

- **Unisys ES7000 16 CPUs**
 - 16 700 MHZ Xeon Processors
 - 16 GB RAM
 - 2 Fibre Channel Paths to 720 GB Disk Subsystem
 - Hardware RAID 1/0 (Striped then Mirrored) Disks
 - 2 Brocade SAN Switches
- **Windows 2000 Datacenter Server SP2**

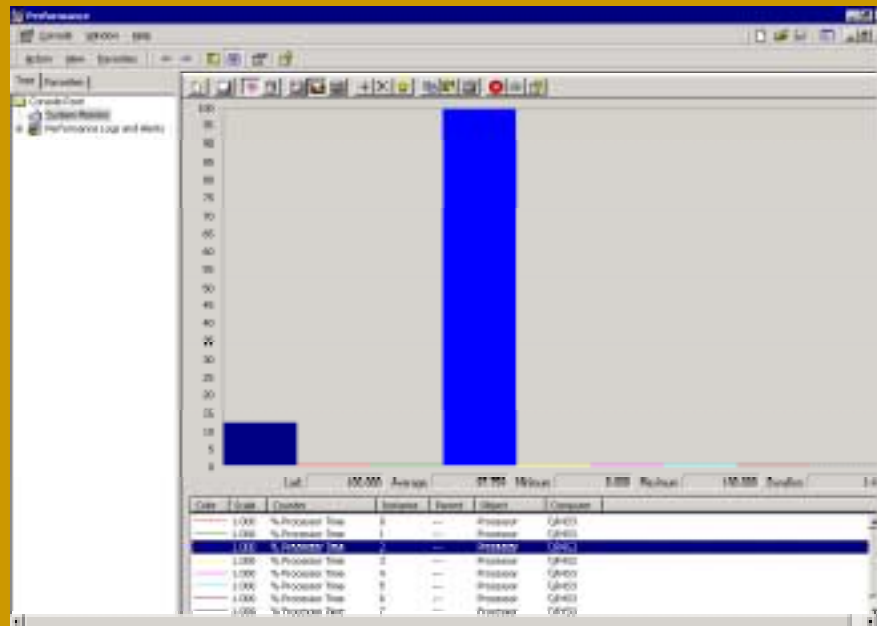
Tuning Methodology

- Establish a Baseline (ES5085 Server)
- Proper Monitoring
 - W2K Datacenter Performance Monitor
 - W2k Datacenter Task Manager
 - W2K/Unisys Process Manager (Affinity)
- Make Changes Individually (One at a Time)
- Repeat Tests.....Repeat Tests.....Repeat Tests.....


Where are the the Bottleneck(s)

- **CPU**
 - **Processor Utilization**
 - CPU Queuing
 - Thread Management
- **I/O**
 - **Disk Queuing**
 - I/O Per Second
- **Memory**
 - Pages/sec
 - Available Bytes

Results ...CPU Scalability....



1 Job CPU Utilization (No parallel processing)




14 Job (SAS Sessions) CPU Utilization

Results I/O


- **High Disk Queuing**
 - Multiple SAS sessions produces contention for SASWORK Area
 - Initially prevented improved performance when testing more than 8 CPUs

Majority of Tuning activity focused on resolving this issue to obtain higher scalability results


Final Processing Time Comparison... (One Day)



Seven Day Processing Results...



30 Day's Processing...



Obtaining Optimal Performance...

- **Robust Disk Configuration**
 - Multiple I/O Channels
 - External Disk Subsystem
 - Multiple Logical Units (LUNs)
 - Hardware RAID
 - Software RAID (RAID over RAID)
- **8-16 Processor (Symmetrical Multi-Processing) Server**
 - Most Effectively handles multiple SAS sessions

ES7000 Guidelines

Hardware

- **Processors**
 - Most effective when number of SAS child sessions is 2 less than total number of CPUs
- **Memory**
 - 16 GB
 - /3 GB Boot.ini Switch
 - Allows SAS sessions more than 2 GB
- **I/O**
 - 2 or More Fibre-Channel Controllers
 - Storage Area Networks (SAN)
 - Multiple (at least 2) Disk Subsystems
 - Split LUNs (Disks) across all channels

WebHound™ Application Guidelines

- **SASWORK (Temporary Area used by ALL SAS Sessions)**
 - Hardware RAID across multiple spindles
 - Create at least 4 Logical Units (LUNs)
 - Split LUNs (I/O activity) across multiple (>2) channels
 - Group LUNs together in Windows as single Volume
 - Provides Software Striping
- **SAS MART (STAY Area)**
 - Raid 1/0

WebHound™ Sizing...

- **File sizing guidelines are provided in the SUGI paper proceedings...**
- **Engage SAS®/Unisys CTC and COE for assistance**

Conclusion

- **WebHound™ 4.0 benefits from scalable hardware in Enterprise Windows Environments**
- **Reduced Processing Time is the end result**
- **Enterprise Systems provide maximum benefits for larger workloads**

Acknowledgements

- **Huong Le (SAS – WebHound™ Testing)**
- **Kevin DeBruhl (SAS - WebHound™ Development)**
- **Margaret Crevar (SAS – CTC)**
- **Scott Moore (Unisys – CTC/COE)**
- **James Lebak (Unisys – W2K Integration Architect)**
- **Rob Hamm (SAS – TAM)**

Questions?

About the Speaker

Speaker **Martin Powdrill**
Technical Account Manager

Location **Unisys Corporation**
104 Acorn Hill Lane
Apex NC 27502

Telephone **(919) 367-8770**

E-Mail **Martin.Powdrill@unisys.com**