




DesertWiz
Consulting



SHARE
Technology • Connections • Results

Increasing Efficiency and Reducing System Overhead

Norman Hollander
System z and zPerformance Specialist
DesertWiz Consulting

zNorman@DesertWiz.Biz
Norman.Hollander@DesertWiz.Biz

SHARE 115, Winter 2010
Boston, MA

August 3, 2010



SHARE in Boston



DesertWiz
Consulting



Trademarks & Product Names

- Trademarks and Registered Trademarks referenced in this Presentation are the Property of the Respective Owners, and should be regarded as appearing with the appropriate ® or ™ symbols at their first reference.
- Observations and Recommendations are solely the opinion of the author.

Abstract



DesertWiz
Consulting



- Today's System z and zEnterprise Server mainframes running z/OS are the most tunable business computing platform ever created. Building on four and a half decades of hardware and software innovation, not only is it pound-for-pound the least expensive environment to run your IT workload, but as it keeps getting more functional and affordable. You can play a key role in making it even more efficient in the context of your environment. This session reviews the mission-enabling hardware, software and z/OS performance characteristics of the platform and offers insight, approaches and solutions for getting even more bang for your IT buck.



DesertWiz
Consulting



Overview

- Mainframe environments continue to grow and become more complex at a rapid pace
- System z Hardware and associated peripherals are challenging for most customers
- Keeping Up with z/OS (z/VM, Linux for System z) technology is time consuming
- The number of IT professionals at most customer sites has NOT grown with the environment
- Many customers have discovered that their previously well-behaved systems and applications are just not performing well
- Many customers don't really know what is causing service levels to miss their goals

Reality Check



DesertWiz
Consulting



“Users constantly call to complain about slow or poor system performance. The phone never rings off the hook with calls of system performance that is too good or too fast.”

-ndh



DesertWiz
Consulting



Agenda

- The Elements of Performance
- System z and zEnterprise Technology
- Delays and Degradation
- Performance Enhancement Features
- Is Tuning Worth It?
- Success Stories
- Q&A




DesertWiz
Consulting



The Elements of Performance

- There are only 3 Elements of Performance
 - Processor
 - General Purpose Engine
 - Coupling Facility Engine
 - zAAP Engine
 - zIIP Engine
 - Crypto Engine
 - Future Offload Engines
 - Memory (Paging)
 - Real Storage
 - 64-bit Architecture
 - Input/Output
 - Peripheral Devices
 - Coupling Facilities
 - Communications Devices
 - Security Devices





DesertWiz
Consulting



Perceived Performance

- There are many performance metrics available to indicate how well the 3 elements of performance are doing
 - CPU busy
 - CPU time
 - Paging
 - I/O rate
- Service Level Agreements typically use something like these as a measure for success
 - Number of transaction per second
 - xx% of specific application business operation completed in yy seconds
 - All jobs updating application's xx database must complete by 3am



DesertWiz
Consulting



Performance Metrics

- All metrics are interesting
- Some are good; some are not so relevant
- Performance metrics should relate to in-place Service Level Agreements
- Accounting and Chargeback metrics are not necessarily complete performance indicators
- In a widget factory, the number of widgets manufactured per hour is critical



DesertWiz
Consulting



Today's Processors

- Today's z/Operating Systems run only on System z Hardware
- z196, z10 (EC, BC), z9 (EC, BC)
 - zEnterprise
 - Model 2817, 125 models
 - 5.2 GHz quad core processor chip
 - z10 EC is 4.4 GHz (~18% increase)
- Specialty Processors
 - zIIPs, zAAPs, ICFs, IFLs
 - Crypto, Compression
 - BladeCenter Extension (zBX)
 - Model 2458
 - Up to 112 Power7 blades

**The best CPU
would have 10%
more MIPS on a
single Processor
than would ever
be needed**

-ndh



DesertWiz
Consulting



Today's Processors

- Big Challenge in Faster, Fewer Engine Environments
 - May be an increase of MIPS
 - May cause queuing complications
 - May release resources for latent demand
 - Uni-Processor environments should be avoided
- CPU time is highly variable
- Specialty engines may cause overhead
- New cache structure may cause overhead



DesertWiz
Consulting



Today's I/O

- Response Time to a Device is key
- FICON channels are typical
 - High performance FICON improves throughput
 - ESCON is going away
 - Parallel and byte multiplexer channels no longer available
- Highly Cached Storage Devices (1-3 milliseconds). Delays?
- High-speed Communication Devices (OSA, Hipersockets)
- Bigger, Faster Tape Devices (and virtualized)
- Faster, and Virtual, CF Links (3-5 microseconds)
- Faster Security Devices
- Separate Compression Engines
- Larger than 32K Blocks of Data





DesertWiz
Consulting



Today's Memory/Paging

- Only Real Memory Available
 - 2 GB was the limit prior to z/OS
 - z196 has up to 3TB of memory
 - 1 TB max per LPAR
 - Large Memory Images now common for DB2, USS, Java, and Websphere Application Services
- 64-bit Architecture
 - Above the Line- 16MB
 - Above the Bar- 2GB
- Paging to DASD Devices
 - Must have sufficient Paging volumes for several 2GB simultaneous Dumps





DesertWiz
Consulting



What's our course of action?

- Performance comes in 2 flavors:
 - Look for the external, environmental opportunities
 - There are many knobs to turn at the operating system level
 - There are many opportunities with the infrastructure components
 - Long term solutions may have new features or options that can improve performance and throughput
 - It may be time to replace outdated solutions with newer more efficient ones



DesertWiz
Consulting



What's our course of action?

- Performance comes in 2 flavors:
 - Look for the internal application opportunities
 - There many be coding opportunities
 - There may be I/O avoidance opportunities
 - There are specific solutions to help identify application inefficiencies




DesertWiz
Consulting



Delays and Degradation

- The Real challenge is discovering:
 - “What am I waiting for?”
 - “I have plenty of work to do.”
 - “I should NOT be waiting.”
- Back to 3 elements
 - CPU, Storage, I/O
- What’s holding me up?
 - Waiting for Device
 - Waiting for Volume
 - Waiting for Resource
 - Waiting for System Services
 - Waiting for JES Services
 - Waiting for Coupling Facility Process
 - Waiting for Communication Process
 - Waiting for Security Process
 - Waiting for Another Application
 - Waiting on Myself
 - Waiting on the UNKNOWN





DesertWiz
Consulting



What's Holding Up the Application?

- Only 2 areas to look into
 - Internally
 - Within the application (faulty code?)
 - Within some system service
 - Externally (Environmental)
 - Within the same image
 - In other systems

Could optimization or modernization benefit?

What's Holding Up the Operating System?



DesertWiz
Consulting



SHARE
Technology - Connections - Results

- Only 2 Areas to look into
 - Internally
 - Within a Server Subsystem
 - Device Allocation, HSM, SMS
 - Tape Library
 - XCF/XES
 - JES
 - Dump
 - Workload Manager
 - Reserve, Enqueues
 - Catalogue
 - DBM, RRS
 - LLA, VLF, DLF
 - Capacity
 - Externally (Environmental)
 - Within the Same Sysplex
 - Reserve, ENQ
 - Workload Manager





DesertWiz
Consulting




What's holding up the I/O?

- Only 2 Areas to Look Into
 - Internally
 - Device Allocation, Dynamic Allocation
 - Reserve, ENQ
 - Catalogue
 - Cache Controller Issue
 - Cached Volume, DFW
 - IOSQ, PAV
 - Fragmented Volume
 - Fragmented Data Set
 - Within the Application (faulty code?)
 - Within a System Service
 - Externally
 - GDPS
 - Network
 - XRC (Global Mirroring)
 - Other Images doing the above





DesertWiz
Consulting



SHARE
Technology • Connections • Results

Performance Enhancing Features of z196/z10 and z/OS



DesertWiz
Consulting



HiperDispatch

- Allows z/OS to manage work across fewer logical processors (z10/z196 only)
 - Limits the number of active logical processors based of partition weight, workload demand, and available capacity
 - Works with enhanced PR/SM microcode to build strong affinity with logical and physical processors (Parked Processors)
 - Helps reduce Multi-Processor (MP) effect
 - Improves throughput performance by 2-10%
 - Dynamically activated with a SET command
- What is needed
 - Revisions to WLM Service Policy to insure proper Workload Prioritization
 - SYSSTC (typically contains many short-running local SRBs)
 - SRBs classified in SYSSTC can execute on ANY available LP even in HD mode
 - Revisions to LPAR Controls (weight, LCPs) to insure enough processors and access to them



DesertWiz
Consulting



Large Page Support

- Special performance enhancing feature (z10/z196 only)
 - Intended for specific applications that are long running and memory access intensive
 - Large Page size is 1 MB, backed by 256 contiguous 4K real storage frames, available only for 64-Bit Virtual Private storage
 - Memory Object Pages are backed at allocation time and are implicitly fixed; they are NOT pageable
 - Throughput performance is variable based on application
 - Statically activated at IPL time, enabled by IEASYSxx LFAREA specification
- What is needed
 - Revisions to LPAR definitions to provide sufficient memory (1 GB min.)
 - Changes to Application code; IARV64 with PAGEFRAMESIZE option
 - Authorized callers must run under a UserID with read access to the Facility Class Resource IARRSM.LRGPAGES

High Performance FICON for System z (zHPF)



DesertWiz
Consulting



- z/Architecture performance enhancing feature (z10/z196 only)
 - Intended for specific I/O intensive applications (OLTP)
 - Substantially reduces the number of Channel Command Words (CCWs) and Information Units (IU sequences)
 - CCWs replaced by TCWs
 - I/O optimized for number access methods, including DB2 (synch I/Os, some Log I/Os, deferred writes), VSAM, ISPF Browse/Edit, PDSE, HFS, zFS
 - Supports FICON Express2, Express4, and new Express8
 - I/O performance improvement in several areas (estimates)
 - 30-70% reduction in number of CCW executions, 10-30% reduction Channel utilization
 - Up to 100% improvement in I/Os per seconds, Up to 50% improvement in Mega Bytes per second
 - Dynamically activated with a SET command
- What is needed
 - No Charge LIC feature on z10 at Driver 76
 - zHPF Chargeable feature on DS8000 System Storage

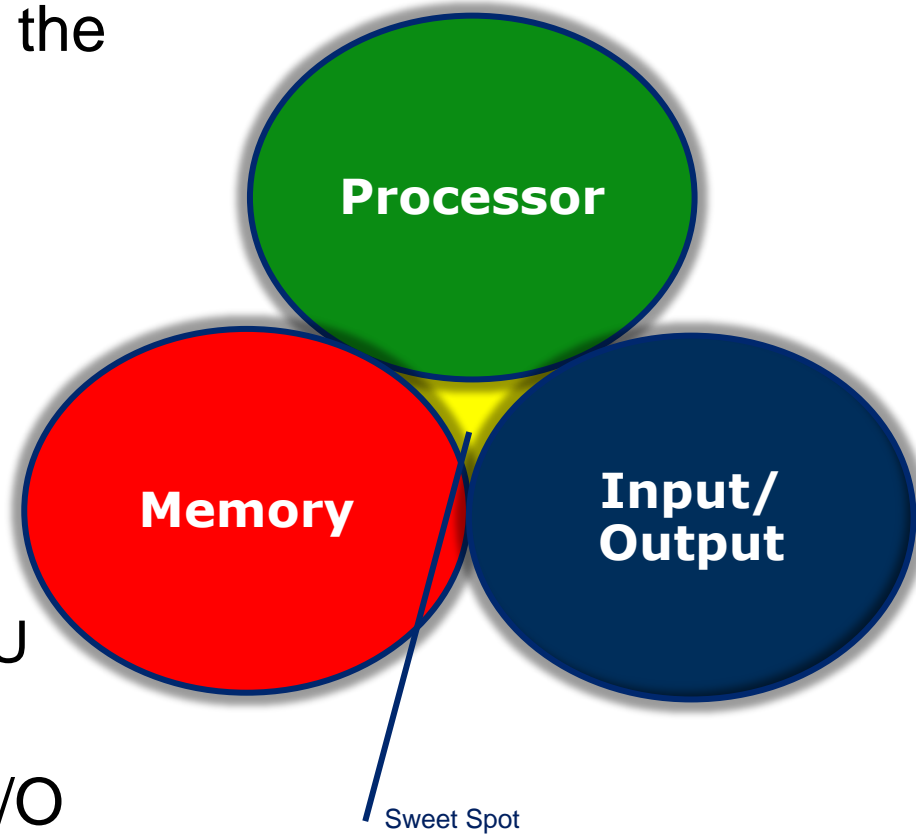


DesertWiz
Consulting



Why bother with all of this?

- Tuning efforts help optimize the balance of our 3 elements
- Use Memory to improve I/O
- Use Memory to improve Processor utilization
- Avoid I/O to optimize Processor utilization
 - There are no “roll-over” CPU minutes
- Use Processor to optimize I/O





DesertWiz
Consulting



Is tuning worth it?

- YES!
- Many low effort, low risk, high value options to improve performance
- Hardware is certainly inexpensive these days; software licenses are not
- Using the technologies of System z and z/OS can improve performance, help achieve goals, avoid/defer upgrades
- Shortening the nightly Batch window improves availability
- Increasing the number of mission critical transactions may allow an organization to improve the Total Cost of Ownership (TCO) of the platform



DesertWiz
Consulting



Success Stories

- **Many small short running transactions**

In a banking environment, 1 CICS transaction that was tuned only saved 15 seconds from the original. Except that 15-20 tellers in 500 branches executed this transaction every hour. $20 \times 500 = 10,000$ transactions per hour, times 15 seconds = 150,000 Seconds every hour. Over an 8 hour day, that's 5.5 hours

- **Weekend Batch database jobs**

In a S&L environment, a Database Batch job ran nearly 43 hours over the weekend. This prevented new loans from being originated during that period. By modifying the code to divide the online files into east and west coast versions, the Batch process was shortened to 11 hours. Further looking at the application, it was determined that the Job was constantly opening a file, reading 1 record, and Closing it. There was also no VSAM buffering specified. Correcting these coding problems took the 11 hour process down to 3 hours, 23 minutes. That returned 8 hours of time that loans could be originated on the weekend. The estimate for that lack of revenue was \$75K/hour; or \$600K/week, or \$31.2 million per year. This was in the mid-90s, so the dollar had a different value then. Even if the estimate was off by 50%, was effort was far worth it?

- **I/O avoidance effort**

Another bank in the west implemented several I/O avoidance solutions. This led to an avoidance of about 14.2 billion I/Os incurred by a "top" banking product. This saved many hours of elapsed time to process the work at an estimated savings of about \$12.1 million per year.

Summary




- There are many performance opportunities to be found in today's mainframe environments. Even with the low cost of hardware today, avoiding software increases, or hardware upgrades makes tuning and worthwhile exercise.
- Balancing the 3 elements of performance can yield big results. In particular, the shrinking of the Batch window, or the increase in transaction throughput.
- There are many performance management solutions to help report on the metrics of interest, help improve I/O avoidance, and find inefficient application code.




DesertWiz
Consulting

Q&A





DesertWiz
Consulting



SHARE
Technology • Connections • Results

Increasing Efficiency and Reducing System Overhead

Norman Hollander
System z and zPerformance Specialist
DesertWiz Consulting

zNorman@DesertWiz.Biz
Norman.Hollander@DesertWiz.Biz

SHARE 115, Winter 2010
Boston, MA

August 3, 2010



SHARE in Boston