Cheryl's Hot Flashes #15

Cheryl Watson

Session 2509; SHARE 106 in Seattle

March 10, 2006

Watson & Walker, Inc.

home of **Cheryl Watson's TUNING Letter, CPU Chart, BoxScore** and **GoalTender**

Agenda

- Survey Questions
- What's New?
  - Dedicated LPARs
  - Assembler Programmer Help
  - New Family of Mainframes!
- User Experiences
- Interesting APARs
- Downloads
- Publications
- 6-Month Update
- This SHARE

Survey Questions - Hardware

- Current Server Type
  - z800, z900, z890?
  - z990?
  - z9-109?
  - Older Hardware?
- Using zAAP Processors *(10 at last SHARE)*?
- Thinking about a zIIP Processor?
- Activated IRD CPU Management *(15)*?
- Have Used On/Off Capacity on Demand?
- Doing Heavy Cryptographic Work *(10)*?
- Using Variable WLC Pricing (*25*)?
- Using WebSphere on z/OS (*35*)?

Survey Questions - Software

- Operating System
  - z/OS.e *(5)*?
  - z/OS 1.4 or 1.5 (*180-190*)?
  - z/OS 1.6 (*20*)?
  - z/OS 1.7?
  - Earlier than z/OS 1.4 (*12*)?
  - Note: End of Service for z/OS 1.4 & 1.5 is March 2007
- Using VSAM RLS for CICS (*12*)?
- Using Transactional VSAM (*0*)?
- Using type 60 records for catalog recovery?

Survey Questions – z/OS Migration

- Installation of new z/OS releases:
  - Every release (each year)?
  - Every other release?
  - Less frequent?
- Length of roll-out of new z/OS releases:
  - 3 months or less?
  - 3 to 6 months?
  - 6 to 9 months?
  - more than 9 months?
- Thanks to Paul Feller of Aegon

Dedicated LPARs

- Dedicated LPAR - an LPAR where all of the CPs are dedicated to that LPAR only
- Once upon a time…
  - Dedicated LPARs performed very closely to the CEC with the same number of CPs.
  - For example, a 9672-RX5 (10-way rated at 445 MIPS) would provide less capacity than two 9672-R55s (two 5-ways rated at total of 535 MIPS). You could get a 20% increase in capacity by dedicating.
  - In the past, many sites created dedicated LPARs to increase capacity.

Dedicated LPARs

- Dedicated LPARs have a cost:
  - You lose the opportunity to share the CPs when needed.
  - On the other hand, some sites use them to reduce I/O delay.

- zPCR became available last fall
  - This free tool from WSC provides a way to estimate the capacity of new configurations.
  - Customer noted that zPCR estimated very little improvement when changing a shared LPAR to dedicated.

- Thanks to David Crowley of Progressive Insurance

Dedicated LPARs

- Another customer tried to gain capacity using dedicated
  - They took a 7-way with three LPARs, and dedicated 4 CPs to one LPAR.
  - All three LPARs got less capacity.
  - zPCR predicted *2% gain* in CEC capacity; but result was *4% loss* in CEC capacity.

- Thanks to Joe Montana and Jerry Urbaniak of Acxiom

- We didn't know why this happened, so we investigated…

Dedicated LPARs

- IBM's Kathy Walsh, Gary King, and Walt Caprice provided the answers:
  - The difference in performance between dedicated and shared CPs has been steadily declining (starting with the 9672 G5 and G6).
  - Performance of shared CPs has improved due to enhancements in hardware, software (z/OS), and millicode (PR/SM). Note the improved MP efficiency (on a 9672, each new CP added between 3 and 6% overhead; on a z9-109, each new CP adds less than 1% overhead).

Dedicated LPARs

- Answers (cont.):
  - Meanwhile, the performance of dedicated CPs has been dropping as the HW design has evolved to support ever larger processors with more CPs, LPARs and memory. Thus, there is more sharing of HW resources (such as the memory cache hierarchy) among all LPARs, dedicated or shared.
  - This resulted in zPCR's estimate of 2% improvement for the case just mentioned.
  - Result – dedicated LPARs no longer buy you enough capacity improvement to outweigh the cost.

Dedicated LPARs

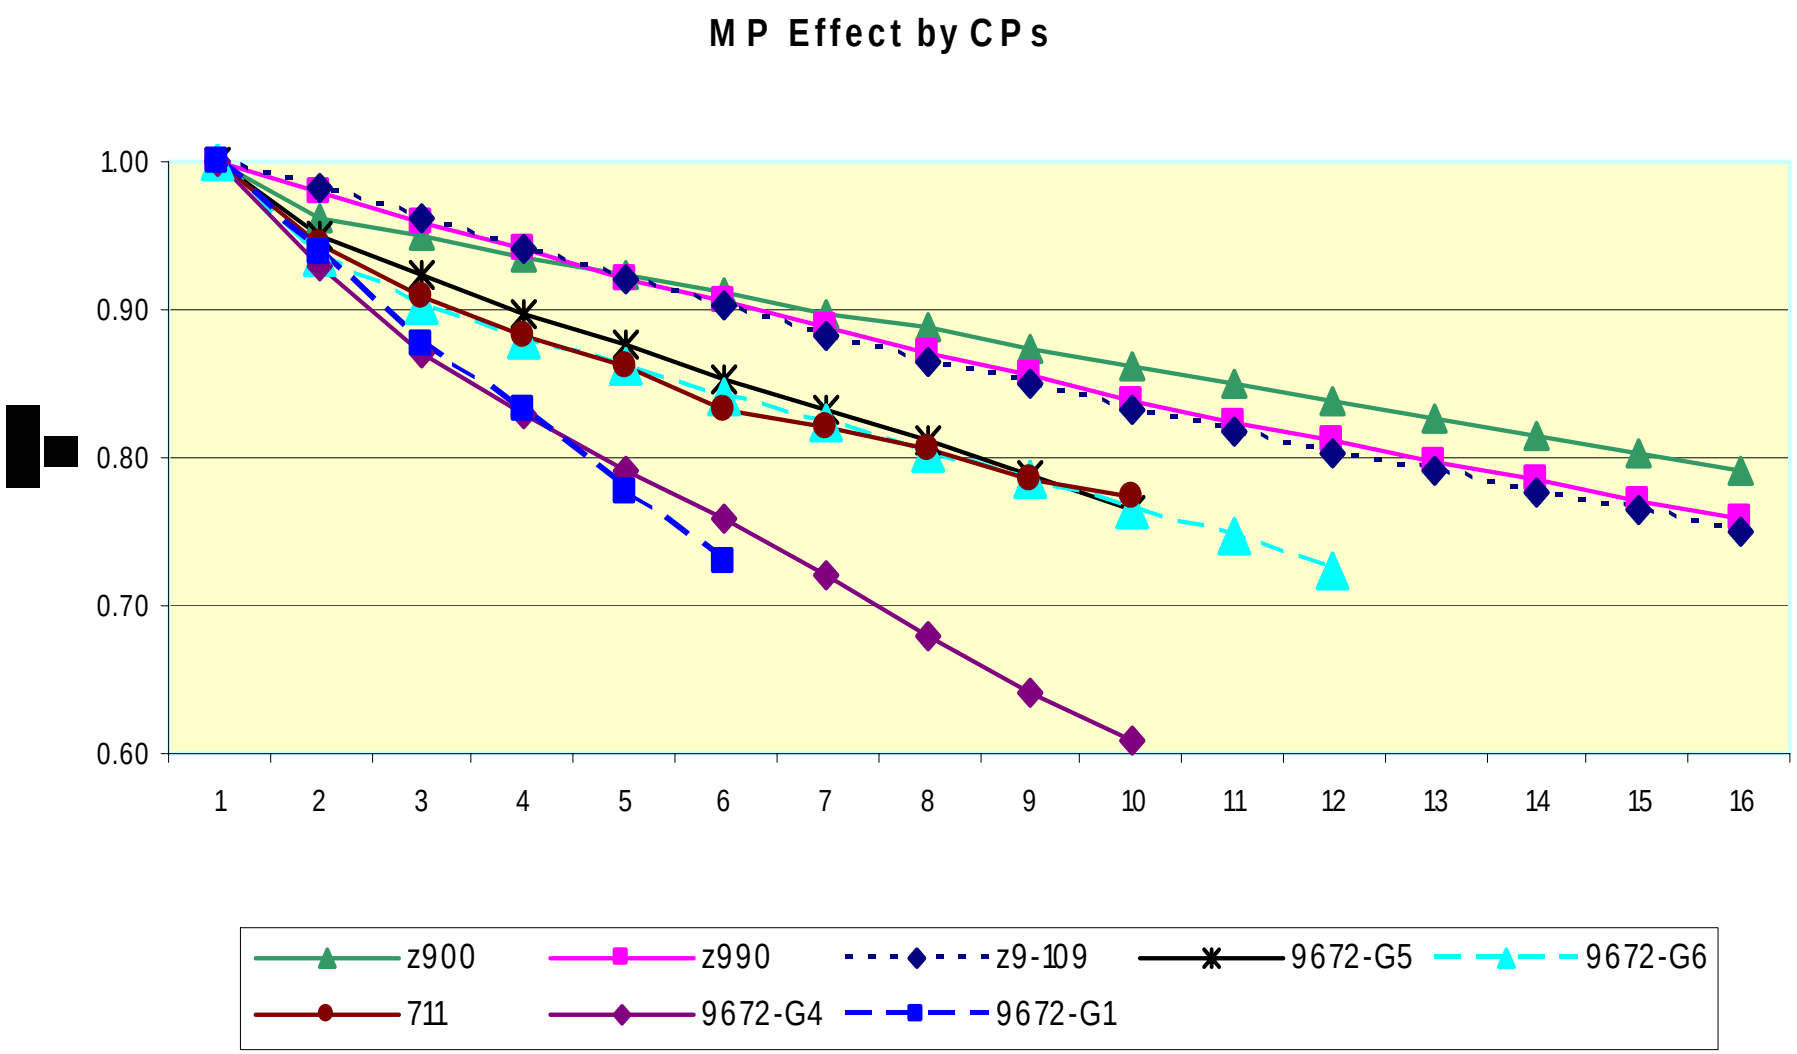

Dedicated LPARs

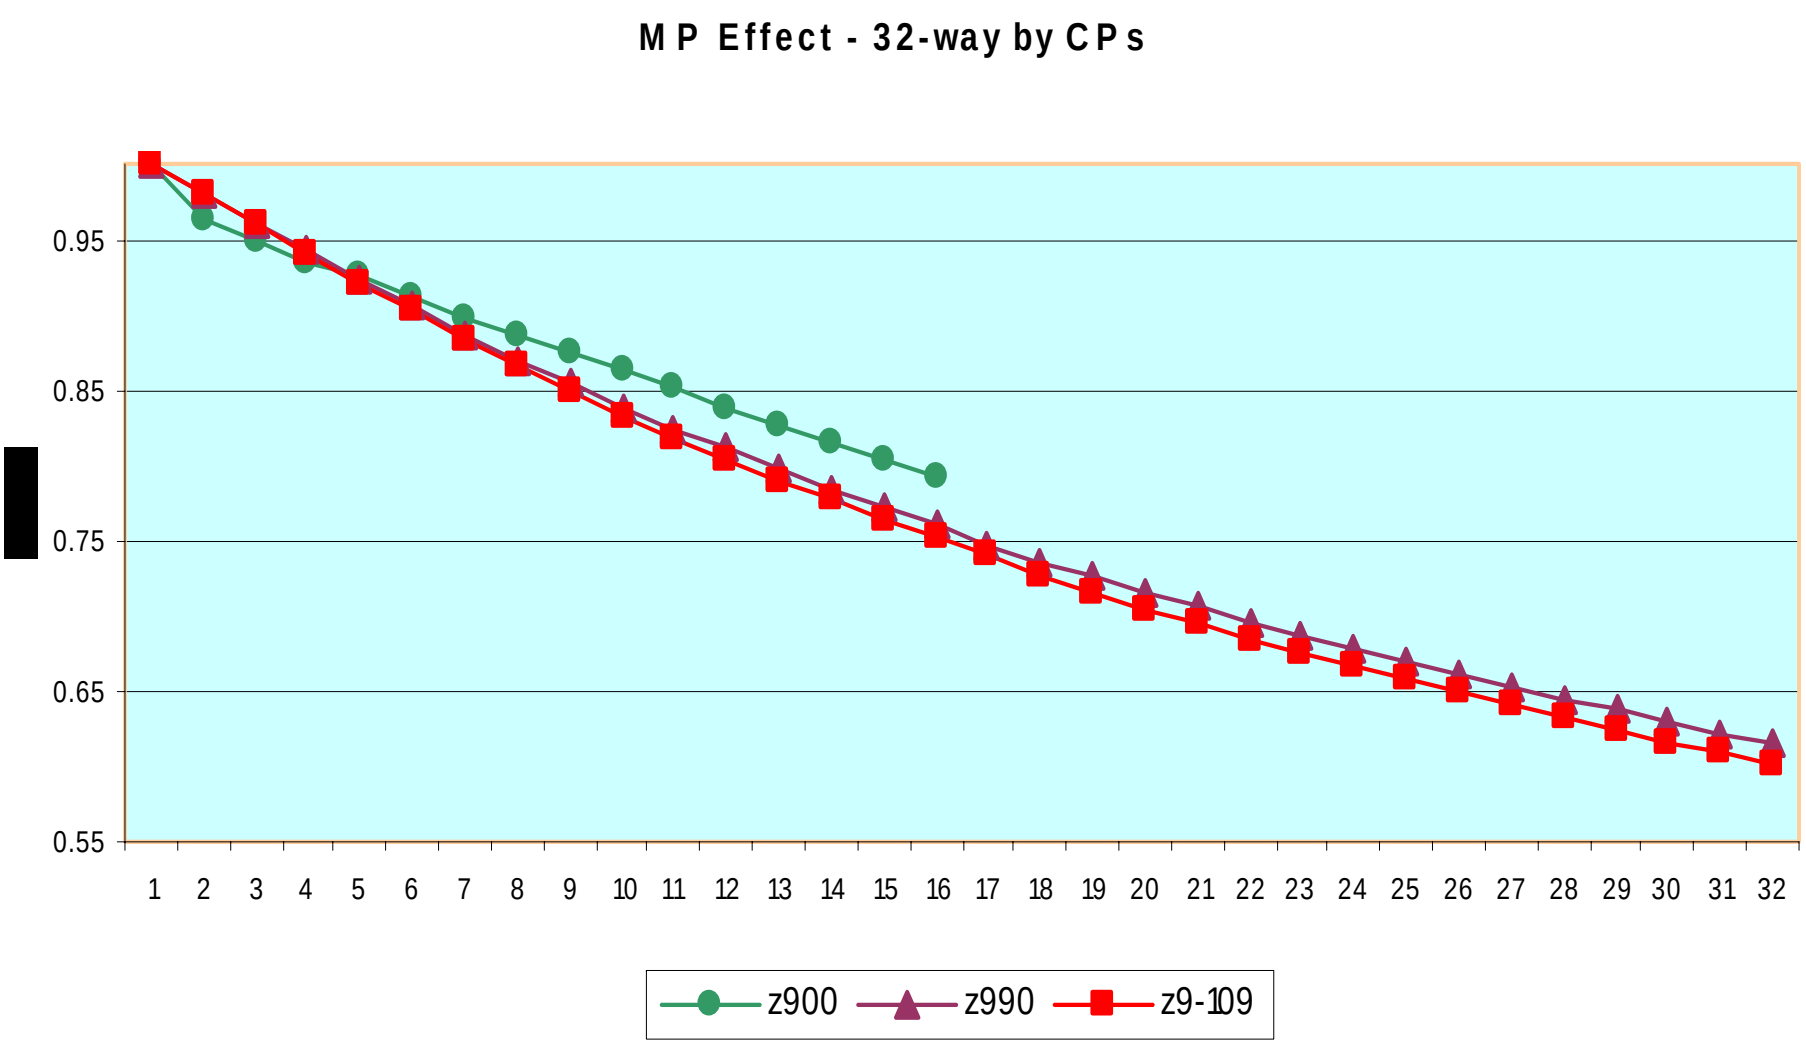

Dedicated LPARs

- A zPCR Example of z990 16-way (2084-316):
  - 1 LPAR, 16 dedicated CPs = 5400 MIPS
  - 1 LPAR, 16 shared CPs = 5379 -0.4%
  - 2 LPARs, 16 shared CPs on each = 5223 -3.3%
  - 2 LPARs, 8 shared CPs on each = 5504 +1.9%
  - 2 LPARs, 8 dedicated CPs on each = 5549 +2.7%
- But…
  - 2 2084-308s (1 LPAR, 8 ded CPs each) = 6025 +11.6%

Dedicated LPARs

- Now Consider Service Units (z990)
  - 16 LPs - 14,953.2 SU/Sec
  - 8 LPs - 17,777.8 SU/Sec (18.9% difference)
  - But depending on whether the 8-way LPAR is on a 16-way or 8-way, the difference is only from -.04% and +11.6% in MIPS (from previous slide).
- If you use service units for capacity planning or chargeback, the service units can be greatly over-stated for anything but a stand-alone machine.

Dedicated LPARs

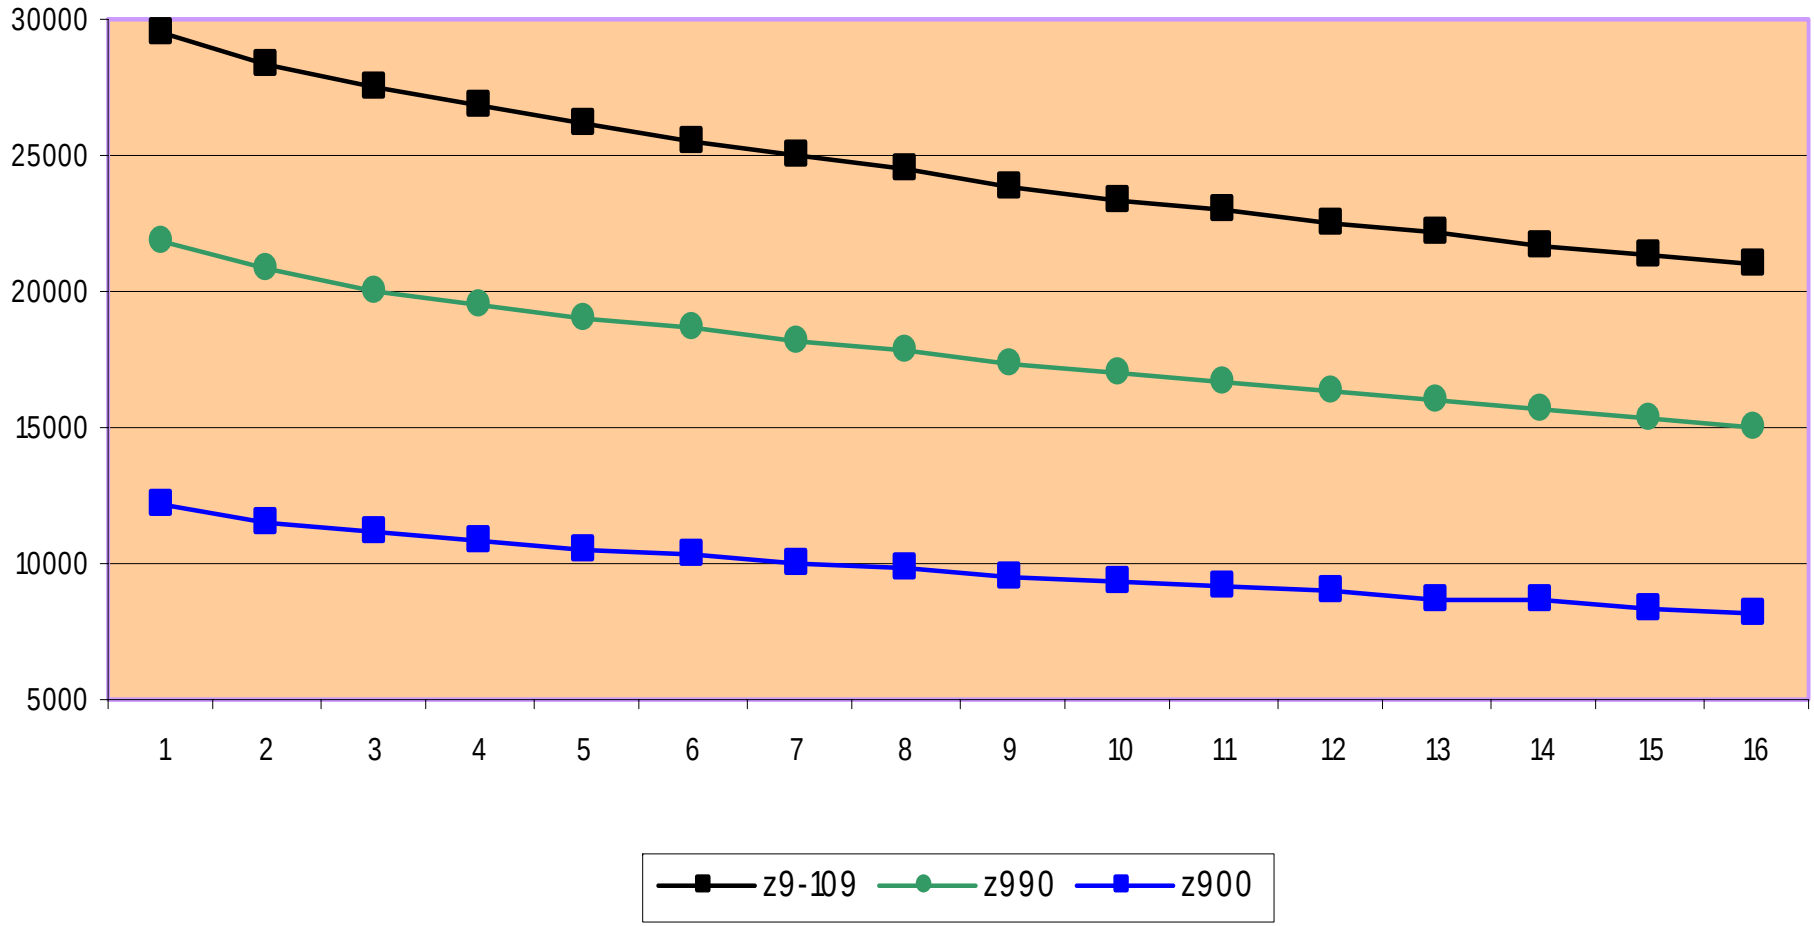

Dedicated LPARs

- Summary
  - Don’t expect to gain much performance by using dedicated LPARs
  - Use zPCR before implementing any configuration change
  - Don’t use service units for capacity planning or chargeback

Assembler Programmer Help

- A major complaint we hear is:
  - We have no assembler programmers any more
  - We have no one to update our exits

- While at CMG in December
  - I saw a product that might help these sites
  - This is not a recommendation (we haven't tried the product or talked to users), but we think you might want to know about it

Assembler Programmer Help

- *OS/EM* (Awesome) from Trident Services
  - It's an Operating System Manager
  - Many features, including HSM optimization, but the one we want to mention is the creation of user exit code by providing ISPF panels for parameter input
  - Almost 300 exits are supported: allocation, DFP, DFHSM, ISPF, JES2, JES3, RACF, SAF, SMF, and TSO/E
  - Allows dynamic changing of exits and LPA modules controlled by exits (i.e. without IPL)
  - See www.triserv.com for more details

Assembler Programmer Help

- *Easy/Exit* from DTS Software
  - Provides exits for SMF, DFHSM, and ABARS
  - Instead of assembler, use CLIST-like statements
  - Several samples provided (e.g. assigning priority to DFHSM RECALLs, rather than using first-come, first-served logic)
  - See www.DTSsoftware.com for more details

New Family of Mainframes!

- Platform Solutions, Inc.
  - Started by ex-Amdahl engineers in 1999
  - Very impressive credentials for funding and management (ex-Amdahl and ex-IBM)
  - PSI provides an Intel Itanium2 based processor to compete with IBM's processors
  - Can support z/OS, Linux, UNIX, & Windows simultaneously
  - www.platform-solutions.com

New Family of Mainframes!

- Platform Solutions, Inc.
  - First trial done by L. L. Bean (presented at last SHARE, available under ‘Literature’ on PSI site); Multiple ESP sites in Europe started in January ‘06
  - Provides alternative to IBM hardware, while still running z/OS
  - Allows interfaces between z/OS, Windows, and UNIX running on the same processors
  - Will have up to 64-way machines initially, followed by 128-way machines

User Experiences – DB2 V8

- Migration from DB2 V7 to DB2 V8
  - Customer experienced a 15% increase in DB2-CPU
  - Applied major maintenance upgrade to DB2 V8, which reduced the increase to only 5%
  - Later upgrade to DB2 V8.1, with its maintenance, brought the CPU usage back to previous levels, although selected rebinding needed to occur first (in order to make use of more efficient V8 optimizer)
  - Best thing to do is to page fix buffer pools, which can provide significant CPU savings in V8
  - We're interested in anyone else's experience – technical@watsonwalker.com

User Experiences – Paging Spike

- Spike in Paging
  - Customer noticed a spike in paging (average 8 pages per second over an hour, but normally close to zero)
  - Cause was tracked down to a DB2 dump
  - Other reasons for paging spikes:
    - Sorting
    - Large subsystem retrieving seldom used data
  - If you want to know the reason, use RMF/CMF Workload Activity Report to see paging by service class

- Thanks to Stephen McAuliffe of Bank of America

User Experiences – SRB Time

- Do sites charge for SRB time?
  - We posted query on IBM-Main & Cheryl's List
  - Most sites charge for TCB plus SRB time at the same rate
  - Five sites said that they didn't charge for SRB time because
    - That's the way we've always done it (don't know why)
    - SRB times aren't repeatable
  - SMF for new processors (zAAPs, maybe zIIPs) won't have separate field for nonquiescable SRBs
  - Recommendation: Treat SRB and TCB times as one for chargeback

User Experiences – PAVs

- On WLM Options Panel:
  - if Dynamic alias management=YES and I/O priority management=YES, then PAVs managed to goals
  - if Dynamic alias management=YES and I/O priority management=NO, then PAVs managed to minimize overall IOS queueing
  - If Dynamic alias management=NO and I/O priority management=YES, then WLM might reduce I/O queueing delays by changing I/O priority
  - If both are =NO (default), then I/O priority is the same as dispatch priority and PAVs aren't dynamically managed

User Experiences – PAVs

- Rumor from CMG
  - Setting WLM to manage PAVs based on goals instead of I/O queueing caused increase in WLM CPU time to increase by a factor of four
    - i.e. Both options were set to YES

- We'd like feedback from attendees who are using dynamic PAVs – technical@watsonwalker.com

Interesting APARs - GRS

- APAR **OA11382**
  - GRS will be enhanced to enable you to specify the contention notifying system (CNS) for GRS Star
  - Long-requested function
  - Still OPEN, will be available for z/OS 1.7 and incorporated in z/OS 1.8

Interesting APARs - BPXBATCH

- APAR OA11699
  - BPXBATCH enhancement based on old SHARE requirement (submitted by Clark Kidd several years ago)
  - Available in z/OS 1.8, or in z/OS 1.5-1.7 with APAR (20Dec2005)
  - Enhancements:
    - Increase the maximum characters allowed for parameter input.
    - Allow use of MVS Files for STDOUT and STDERR.
    - Allow the invocation of APF authorized z/OS UNIX programs in the same address space as the originating job.
  - The SHARE requirement process DOES work.

Interesting APARs - CICS

- APAR PK20133
  - CICS prior to 3.1 running under z/OS 1.7
  - Problems with EIBTIME being overlaid by the CEECEOWN routine.

- Thanks to Gary Maditz of Navy Federal

Interesting APARs - COBOL

- APAR PQ95214/PK15432
  - PQ95214 (z/OS 1.4-1.6, 24Jun2005; included in z/OS 1.7) - *COBOL Runtime Update*. Changes are made to Language Environment in support of programs compiled with the Enterprise COBOL for z/OS Version 3 Release 4 compiler. Now PE.
  - PK15432 (z/OS 1.4-1.7, 25Jan2006) - *Exec Binary SEARCH ALL ... WHEN .. Gives Different Result After PQ95214 if Search Argument is Longer Than (the key) 06/01/19 PTF PECHANGE*.
  - Users could see incorrect output

Interesting APARs - COBOL

- APAR PQ95214/PK15432
  - If you have applied PQ95214, you should back it off or apply PK15432.
  - If you are running the z/OS 1.7 version of Language Environment, you should apply PK15432.
  - As you convert programs to Enterprise COBOL V3R4, make sure to pay attention to warning messages issued during the compile or during application run time.

Interesting APARs - ASM

- APARS relating to ASM's Suspend/Resume with PAVs
  - OA09675 (z/OS 1.4+, 16Feb2006) - *Hang Occurred After Fdrpas Swap of PLPA/Commom Paging Device*; pre-REQs OA14248
  - Any processing that UNBINDs aliases (such as TDMF, HIPERSWAP, IOS Dynamic pathing failures, etc.) is subject to this failure.
  - The problem exists when any suspended channel program is on an alias when the alias becomes unbound. Functions such as swapping will unbind aliases prior to continuing with their swap, thus, they are more susceptible to the problem.

Interesting APARs - ASM

- APARs relating to ASM's Suspend/Resume
  - OA14248 (z/OS 1.4+, 7Feb2006) – *New Function*.
  - As processors, control units, DASD have increased in speed the benefits of suspend/resume have decreased to the point where it has been determined the difference between suspend/ resume and start subchannel are no longer relevant. Therefore it has been decided to no longer use suspend/resume.
  - Discussions on IBM-Main bring into question whether IBM is just removing suspend/resume as a method to solve a difficult problem.
  - I believe those IBMers who say that in today's configurations that suspend/resume simply doesn't buy you much.

Downloads

- Download z/OS 1.7 PDF Files
  - www.ibm.com/servers/eserver/zseries/zos/bkserv/r7pdf/mvs.html

- For the oldtimers:
  - Interview regarding birth of computing (ENIAC) in ComputerWorld, 20Feb2006
  - www.computerworld.com/databasetopics/data/datacenter/story/0,10801,108790,00.html?from=story%5Fkc
  - Thanks to Lawrence Jermyn

Downloads

- SAS Newsletter
  - From Systems Seminar Consultants, Inc.
    - www.sys-seminar.com
    - Company provides training and consulting
    - Provides free newsletter on SAS techniques
      - Called ***The Missing Semicolon***
      - Select archives from left side of home page
      - It's EXCELLENT!
    - Also find their conference presentations online

- Thanks to **Jerry Urbaniak** of Acxiom

Downloads

- OMEGAMON® z/OS® Management Console
  - No charge monitoring tool to look at Health Checker, system status, and configuration data
  - Designed for those new to z/OS
  - Announcement letter 205-329 on 13Dec2005
  - See the description in HOT Topics, found in your SHARE bag
  - It's *VERY COOL!*

Publications

- Redbooks
  - **www.redbooks.ibm.com**
  - **SG24-6472-02** (Updated 1Feb2006)
    - *System Programmer's Guide to: Workload Manager*
  - **REDP-4014** (2Mar2006)
    - *Coupling Facility Performance: A Real World Perspective*
  - **REDP-3940** (21Mar2005)
    - *JES2 Performance and Availability Considerations*

Publications

- Techdocs – The Technical Sales Library
  - **www.ibm.com/support/techdocs**
  - **TD102705** (7Nov2005)
    - Managing Logical Processors on the IBM eServer zSeries z990 (describes multibook considerations)
  - **TD102670** (25Oct2005)
    - Dynamic ICF Dispatching

Publications

- Watson & Walker Web Site
  - Updated SMF Reference Summary (up to z/OS 1.7)
    - www.watsonwalker.com/articles.html
    - Enhanced with corresponding exit names
    - This new summary is 8.5”x11”, but will be uploading a smaller version soon
  - Results of Survey (for those not participating)
    - Cheryl’s List
    - Free listserver with TUNING Letter status and occasional tips and hints – sign up at www.watsonwalker.com/feed.html
    - We’ll send out the survey results next week

6-Month Update – VSAM RLS

- Asked for experiences from large VSAM RLS shops
  - Few responses
  - Alan Gray, Lead Systems Architect at CareFirst BlueCross BlueShield responded:
    - Used RLS to support biggest CICS region (primarily on a 1C8, but some on a z104), worked well
    - The RLS usage was seen as an increase in supervisor CPU time for the CICS region
    - Applications area moved one of biggest file systems to DB2, at a cost much greater than RLS

6-Month Update – VSAM (Repeat)

- Obsolete VSAM Attributes
  - IMBED, REPLICATE and KEYRANGE
  - Not Supported for New VSAM Data Sets
  - Open will Fail for EXISTING Data Sets in a Future z/OS Release (sometime after z/OS 1.8)
- **II13894**
  - Describes Free Tool to Identify Existing Data Sets
- **OA10952**
  - Enhancement to IDCAMS EXPORT and IMPORT
  - Removes IMBED and REPLICATE
  - Use to Fix Existing Data Sets

6-Month Update - VSAM (Repeat)

- APAR **OA11334**
  - Adds New Messages to IDCAMS EXAMINE INDEXTEST
    - IDC11773I
      - Number of Keys on Each Index Level
      - Average Key Length After Compression
    - IDC11774I
      - Current CI Size
      - Recommended Minimum CI Size
    - IDC11775I
      - Number of Unreachable CI Blocks (if any)
  - Should be Useful for VSAM KSDS Tuning
    - Values Difficult to Obtain Elsewhere
  - Bad News: Currently Only Available for z/OS 1.6+

This SHARE

- Interesting Sessions
  - 1057, *"Why CICS? Why the Mainframe? Why > > Now?"* by Steve Ware, University of Florida – provides compelling reasons for maintaining CICS on the mainframe
  - 2500, *"zOS Performance 'Hot' Topics"* by Kathy Walsh – latest in WSC information and performance APARs (always my favorite session at SHARE)
  - 3801, *"What's New with GDPS in 2006?"* by Angelo Corridori – Great summary of GDPS and terminology (plus new enhancements)

This SHARE

- From Harv Emery, session 2867
  - New z9-109 HSA Estimation Tool
  - Available at no charge on Resource Link
  - You can now estimate HSA requirement *before* ordering storage
  - www.ibm.com/servers/resourcelink

This SHARE

- 'Make IT Easy' from IBM
  - www.ibm.com/servers/eserver/zseries/zos/eou
  - Provides website for 'z/OS novices'
  - Includes:
    - z/OS Basics – Excellent manual!
    - Pointer to z/OS Wizards website - www.ibm.com/servers/eserver/zseries/zos/wizards
    - Many other links

This SHARE

- From Kathy Walsh, session 2500
  - WSC has seen several real storage problems
  - Most can be prevented by specifying in IEAOPTxx
    - MCCAFCTH=(4000,4500)
  - See APARs OA12185 (SUG, 18Oct2005) and OA14318 (OPEN), both spin loops
  - See RSM APARs OA06417-OA06419, which manage new resource thresholds to prevent storage shortage below 2GB if 3GB+ is online

This SHARE

- IBM wants to keep us on our toes:
  - New term: System z = zSeries and/or z9-109
  - Also seen this week:
    - DB2 for z/OS V8
      - DB2 V8 or z/OS V8 – you guess!
    - EWLM for z/OS 2.1
      - Oops! Did I miss a z/OS announcement?
    - GDPS/XRC is referred to as z/OS Global Mirror
      - This is different from GDPS/Global Mirror
      - It's also different from GDPS/PPRC which is referred to as Metro Mirror

More From This SHARE!

- zAAPs and IFLs
  - Once you purchase a zAAP or IFL processor, that purchase is for life when you do an MES upgrade (unless IBM changes the rules)
  - So if you buy a zAAP on a z990 and swap out the machine for a z9, you get the new & faster zAAP on the z9 at no additional charge
  - Knowing this should help in planning and timing of purchases (e.g. you might not need to wait until upgrade time)
  - Status of zIIPs and ICFs are unknown (by me)

More From This SHARE!

- HATS (Host Access Transformation Services)
  - Provides quick way to change green screen applications to Web-based with up-to-date look and feel
  - High use in customer support or data entry sites where green-screen training is a problem
  - Customers say that with training (mostly free with the product) that you can have an application converted in a week. Once trained, you can convert smaller applications in a matter of hours.

- Thanks to Jim Marshall

Traditional Green Screen – From IBM Handout

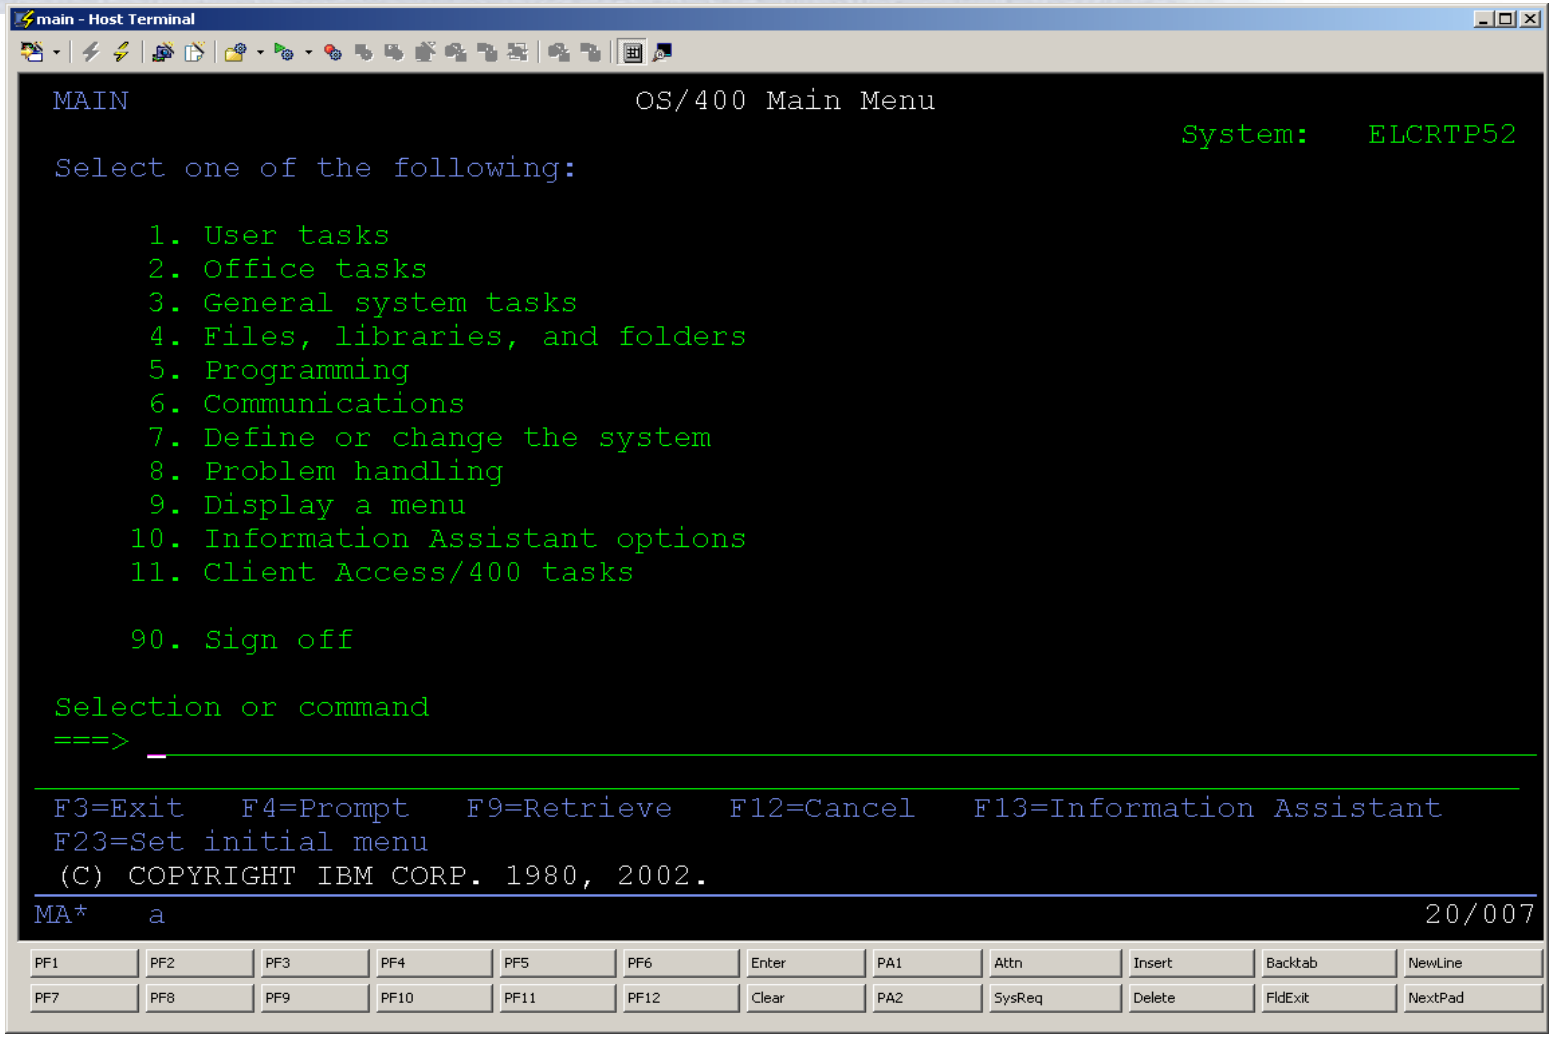

HATS Transformation Using Default Rules – from IBM Handout

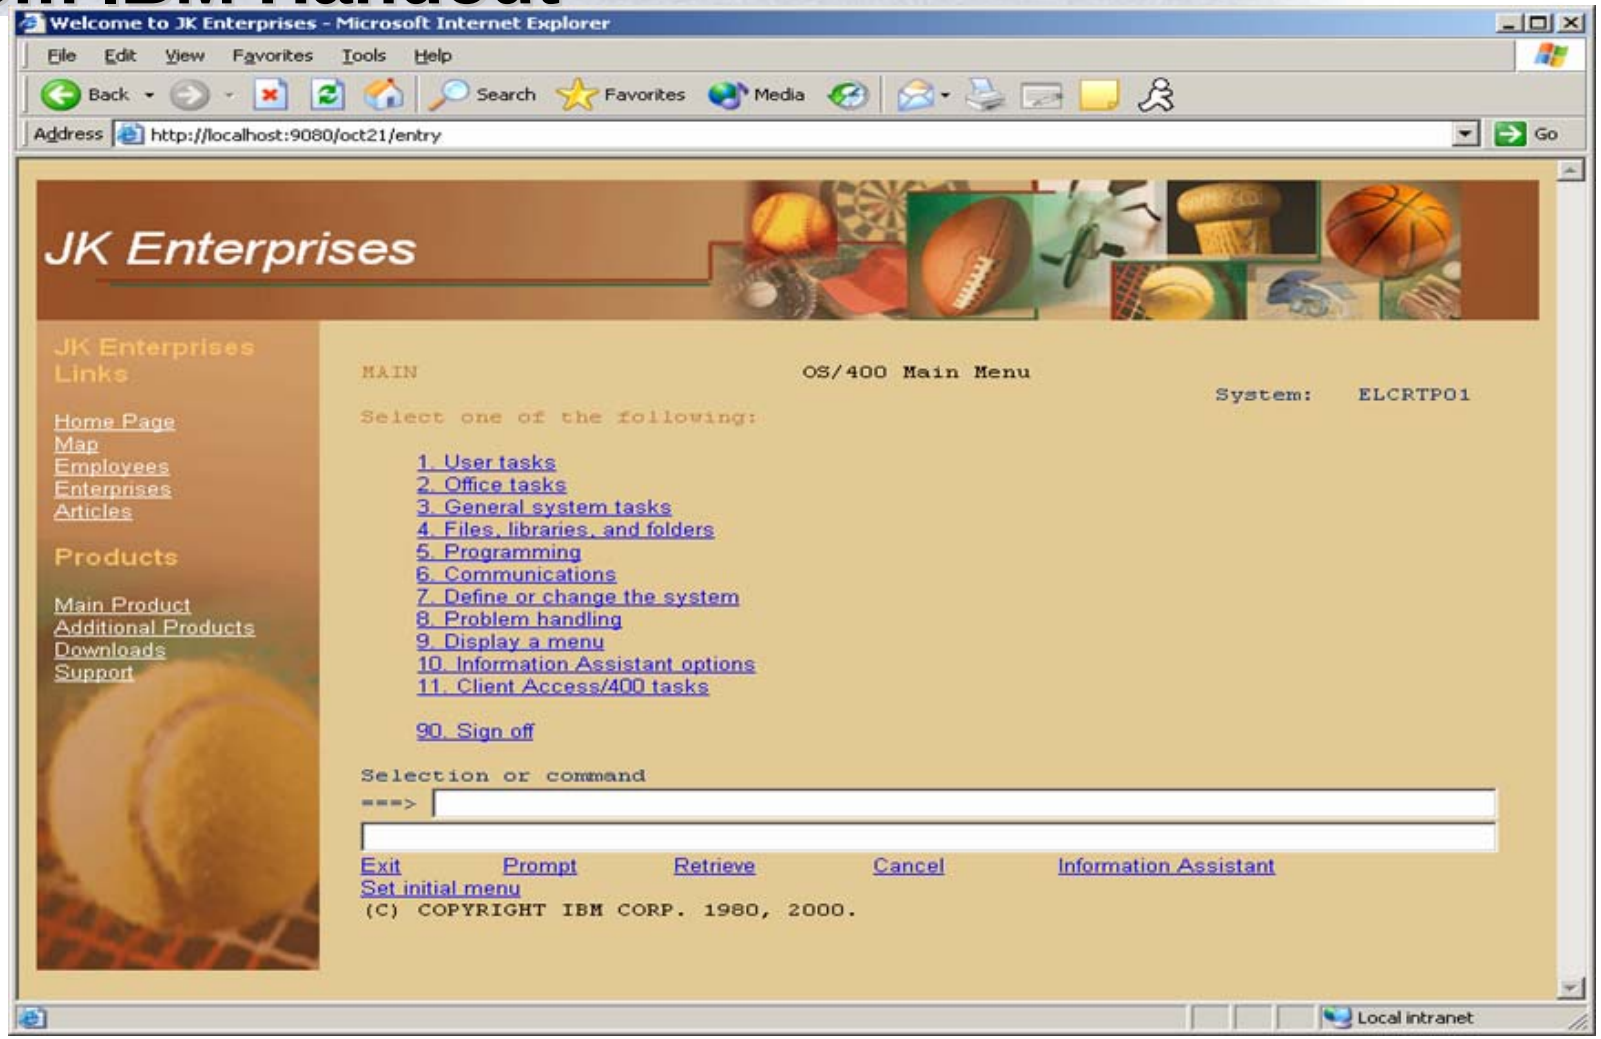

Combining Screens – From IBM Handout

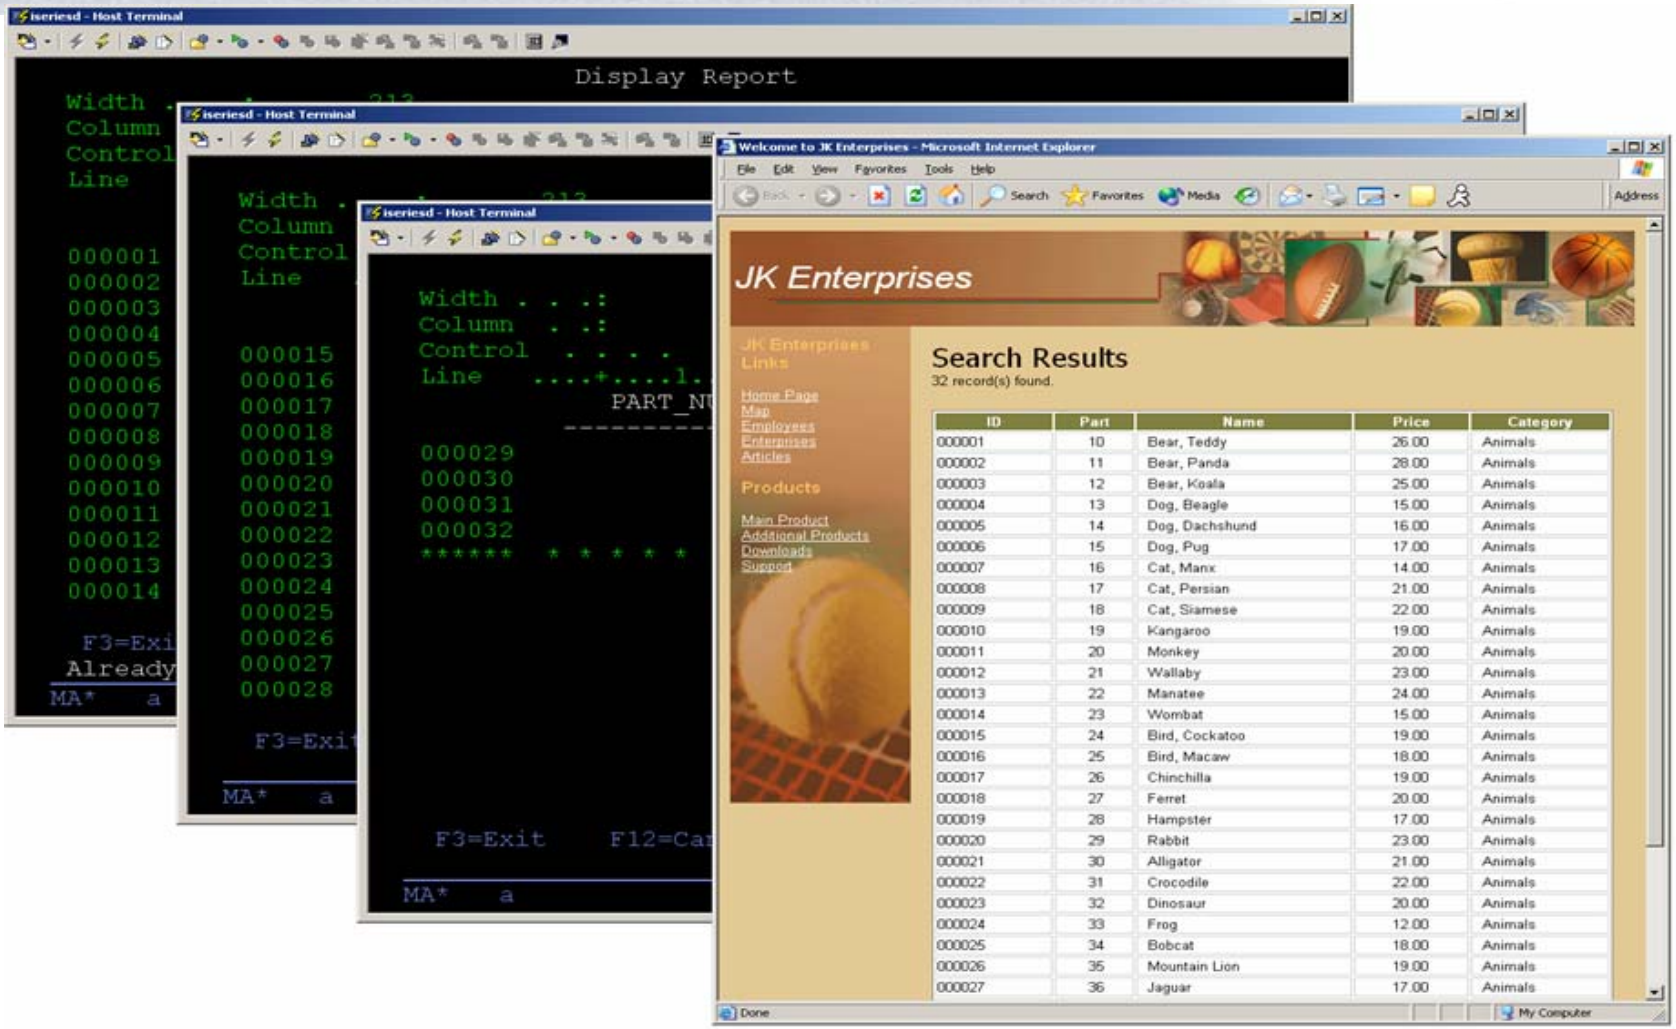

ID	Part	Name	Price	Category
000001	10	Bear, Teddy	26.00	Animals
000002	11	Bear, Panda	28.00	Animals
000003	12	Bear, Koala	25.00	Animals
000004	13	Dog, Beagle	15.00	Animals
000005	14	Dog, Dachshund	16.00	Animals
000006	15	Dog, Pug	17.00	Animals
000007	16	Cat, Manx	14.00	Animals
000008	17	Cat, Persian	21.00	Animals
000009	18	Cat, Siamese	22.00	Animals
000010	19	Kangaroo	19.00	Animals
000011	20	Monkey	20.00	Animals
000012	21	Wallaby	23.00	Animals
000013	22	Manatee	24.00	Animals
000014	23	Wombat	15.00	Animals
000015	24	Bird, Cockatoo	19.00	Animals
000016	25	Bird, Macaw	18.00	Animals
000017	26	Chinchilla	19.00	Animals
000018	27	Ferret	20.00	Animals
000019	28	Hampster	17.00	Animals
000020	29	Rabbit	23.00	Animals
000021	30	Alligator	21.00	Animals
000022	31	Crocodile	22.00	Animals
000023	32	Dinosaur	20.00	Animals
000024	33	Frog	12.00	Animals
000025	34	Bobcat	18.00	Animals
000026	35	Mountain Lion	19.00	Animals
000027	36	Jaguar	17.00	Animals

HATS

- HATS
  - Very inexpensive and powerful
  - Many customers are unaware of this option
  - Lots and lots of help and documentation
  - www.ibm.com/software/webservers/hats
  - Request CD-ROM with literature and presentations
  - Join the Web world today!

- Also see sessions by Pratik Nanavati in proceedings

One Last APAR

- OA14222
  - *ADEND15F RC190 After a Policy Activation*, z/OS 1.4-1.7, 14Feb2006
  - If using CPU critical and a policy activation occurs, it's possible for the DP to be set incorrectly and either an abend or spin loop occurs

- Thanks to Lawrence Jermyn

See You in Baltimore!

- Email:
  - technical@watsonwalker.com
- Web site: www.watsonwalker.com