POWERED BY HP...
HPe3000

MPE/iX 6.5

High-End Performance Release

Presented by:

Kevin Cooper

Hewlett-Packard
Commercial Systems Division
19447 Pruneridge Avenue #47UA
Cupertino, CA 95014 U.S.A.
Phone: (408)447-4004
FAX: (408)447-4441
Email: kevin_cooper@hp.com

MPE/iX 6.5 Performance - Overview of Presentation

- 6.5 Performance and Capacity Objectives
- 6.5 Capacity Increases
- 6.5 Performance Enhancements
- New Performance Options Since 5.5
- What Should You Expect with 6.5?

MPE/iX 6.5
Target Customers

- Customers reaching the performance or capacity limits of the 997/800 or 989/650 on MPE/iX releases 5.5 or 6.0.
- HP e3000 Server Consolidations
- High-end AMISYS application customers
- Customers who need other new features available only in the 6.5 release.

MPE/iX 6.5 Performance and Capacity Objectives

- Performance: A 30% increase for
  a Series 997/1200 running 6.5
  over a Series 997/800 running 6.0
- Capacity: Support enough users on the new Series 997/1200 system to take advantage of all its performance.
- Result: 44% performance increase, with lots more users!

What MPE/iX 6.5 Is - and Is Not

- MPE/iX 6.5 is a HIGH-END performance release, to enable larger HP e3000s.
- 6.5 does NOT provide a performance improvement for every HP e3000 system, but its “performance” changes were made in ways so they would not degrade performance on any HP e3000 systems.
- Some systems may notice an increase in overhead from other new 6.5 features.

MPE/iX 6.5
Capacity Increases

- More Concurrent Users and Processes
- Files > 4GB
- Memory > 4GB
- More Disks (up to 511 now supported)

More Concurrent Users and Processes

- Virtual Terminal (VT) connections increase from 2000 to 2600.
- Data Segment Table (DST) entries, used by the system for every logon, increase from 16K to 64K.
- The :NEWCI command allows more user processes on a system, by replacing each user's Command Interpreter (CI) process.

Breaking the 4GB Barrier - Larger Files

- On 6.5, two types of files can now contain more than 4GB of data:
  - MPE fixed-length record files
  - New KSAM64 files
- There is some increase in file system overhead on 6.5 for all fixed-length files.
- TurboIMAGE datasets can be up to 80GB in size, when using the maximum block size. They still use “chunks” of up to 4GB.

Breaking the 4GB Barrier - More Memory

- The Series 997 now supports up to 16GB. Best performance is with the new 8GB "Large Memory" carrier cards, which hold memory in 1GB increments.
- 8GB is now supported on the Series 979, 989, 929/030, and 939/030, which all have 64-bit PA-8000 or PA-8200 processors.
- Older systems are still limited to 3.75GB, due to their 32-bit processors.

How Much Memory Do 6.5 Systems Need?

- Rumors that 6.5 takes huge amounts of memory to run are simply NOT true.
- Each user (session, network connection, or job) should require ABOUT the same amount of memory on 6.5 as on 5.5 or 6.0.
- The 6.5 operating system itself may use a few megabytes more memory than 6.0 did, as a result of its new features.

How Much Memory Does -YOUR- 6.5 System Need?

- 6.5 should NOT require more memory if:
  - your system RUNS WELL on 5.5 or 6.0,
  - you have at least 128MB, and
  - you are NOT adding anything else.
- If you are ADDING users, processors or software, or if you see memory pressure on 5.5 or 6.0, you should add memory.
- If you have 96MB or less, you are more likely to notice an impact. You should also check for patches before updating.

6.5 Performance Enhancements

- Memory Manager
- Dispatcher
- Reorganized NL
- Sort

A Performance Analogy: An Airport Runway

- Only one plane at a time can be on the runway (the critical resource).
- Everything works fine if there are only a few planes needing the runway (light demand for the resource).
- Problems arise during peak periods or bad weather (demand exceeds capacity).

Analogy: Runway Congestion

- Some planes must wait on the ground (blocked, waiting for the resource).
- Some planes have to circle in the air (spinning, waiting for the resource).
- Other planes end up waiting at a distant airport, with some flights being cancelled (to decrease the demand).
- We need to find a way to get more planes through the airport!

Memory Manager

- Problem: Before 6.5, all processors on a system shared a single resource lock for all Memory Manager activity (to fetch or post a page, or change its status).
- Any processor trying to lock this resource went into a tight “spin wait” loop if it was already locked. The plane was circling.
- One of the biggest causes of spin waits on MPE/iX 5.5 is the checkpoint process of Transaction Manager (often PIN 11).

Memory Manager Solution

- 6.5 introduces a new locking mechanism in the Memory Manager, which increases the “granularity” of its locks.
- This could be compared to adding new runways, which cross each other. Now only the intersections must be controlled.
- This greatly reduces the spin waits.
- On 6.5, many processors can be doing memory management at the same time.

Dispatcher (Process Scheduler)

- Problem: Before 6.5, all processors shared a single “Run Queue” of processes ready to run. Activity in this queue was too high on busy on-line systems, where many processes run for a short time and then wait for a terminal read or disk I/O.
- Any processor trying to lock this queue (to add, remove, or change the priority of a process) went into a tight “spin wait” loop if the queue was already locked.

Dispatcher Solution

- 6.5 introduces “Multiple Run Queues” on Series 997/1000 and 997/1200 systems. Each processor gets its own run queue.
- Back to the airport: We added enough non-intersecting runways to handle all the planes!
- 6.5 also includes a new “Load Balancer” to keep these queues properly balanced.

Reorganized NL.PUB.SYS

- Problem: Series 997 and 979 systems were having too many TLB (Translation Lookaside Buffer) Misses - when virtual addresses were not found in the TLB.
- It turned out the NL was using too many of the available TLB entries on PA-8000s.
- Solution: The 6.5 NL was reorganized so it would take fewer TLB entries on these systems, by combining routines with like access rights into fewer, larger modules.

Sort

- The intrinsics that perform Sort routines had to be modified to handle files >4GB.
- While modifying this code, HP found a way to reduce the overhead of Sort.
- Sorting a file >4GB works fine, but expect it to take a while, even stand-alone on today's fastest HP e3000 systems.

New Performance Options Since MPE/iX Release 5.5

- TurboIMAGE
- Transaction Manager
- Networking
- 5.5 Performance Patches

TurboIMAGE (TI) Options

- Dependency Semaphore (DSem):
  Changes TI internally so it locks a subset of related datasets, instead of the whole database, while doing modifications to chains. (TI adds intersecting runways.)

- Prefetch:
  Changes TI so it first fetches all pages from disk that may need to be modified, before it issues its internal locks. (TI spends less time using the runways.)

Transaction Manager (XM)

- ALTERCHKPTPRI command (VOLUTIL) on 6.5 allows you to configure the priority of the XM checkpoint process (PIN 11). [A 5.5 patch can change it from 30 to 152.]
- ALTERCHKPTSTAT command (VOLUTIL) decreases XM checkpoint overhead by changing how file pages are selected for posting to disk (on 6.0 only).
- XM log file size is configurable as of 6.0.

Networking

- Increased TCP performance
- Increased socket connections (TCP, UDP)
- 100 Base T network links

5.5 Performance Patches (included in 6.0 and 6.5)

- Terminal Writes - Reduced contention in the routine “tm_lock_acb”, by no longer requiring every terminal write on the system to lock a single semaphore.
- Timer Semaphore - Changed this new semaphore, introduced in 5.0 for the :SETCLOCK command, so it would be “priority boosting.”

What Performance Should You Expect With MPE/iX 6.5?

- You will see the greatest performance benefits from MPE/iX 6.5 on Series 997 10-way and 12-way systems.

- You may see some benefit on other systems - most likely those with:
  - 6-way or 8-way multi-processing
  - 64-bit PA-8000 or PA-8200 processors (Series 997, 979, 989).