

VM Performance History

SEAS

7Oct1986

Lynn Wheeler
lynn@garlic.com

360/67

- 750 ns memory
- Simplex - 256k-1024k
- Duplex – 512k-2048k
- Adresss Translation
 - 24bit
 - 32bit

CP/67 Release 1

- 10 level dispatcher
 - Based on cpu use
 - High processor overhead
- No real storage throttle

CP/67 Release 2

- Two level dispatcher (Q1 & Q2)
 - Essentially round-robin
 - Every pass thru dispatcher look at all tasks multiple times
 - 30-40 users, 10% cpu overhead
- Thashing throttle based on fixed table
- Kernel storage management significant
 - Upwards 1/3rd CP overhead

Work as undergraduate

- Fastpath
 - Kernel SVC linkage reduced 185ms to 65ms
 - Virtual SVC reflection 25ms eliminating dispatch
 - Program interrupt handler
 - Dispatcher
- BALR linkages
- Dynamic adaptive controls
- Local LRU & Working Set thrashing throttle
- Pageable kernel
- Custom I/O for CMS file i/o
- Ordered seek & chained requests

68 Share Presentation

- OS/360 MFT14
 - Base 322secs elapsed time
- Base & Modified CP/67
 - Reduced CP CPU from 533secs to 113secs

	W/CP	ratio	CP CPU
original	855	2.2	533
changed	435	1.35	113

CP67 Release 3

- FastPath
 - Fast redispatch
 - Fast SVC reflect
 - Fast instruction simulation restart
- Restructured Dispatcher
 - Separate in-q chain
- Kernel storage subpool
- CMS Diagnose I/O

CP/67 Release 3.1

- Dynamic Adaptive Scheduler
- Working set page thrash
- Global LRU Replacement
- Limited avail. Support for V=R

CP67 Release 3.2

- Internal 3.1L
 - Development group focused on vm370
- Ordered seek queuing
- PCI interrupts
- Chained requests
 - Increased 2301 from 80/sec to 300/sec

VM370 Release 1

- Lots of Simplification
- No Fastpath, dynamic adaptive, etc
 - Q1 always ahead of Q2
 - Quanta limit only virtual CPU
 - One case Q1 ran for 20mins
- Shared segments
- Local (not global) LRU
- Approx. round-robin

VM370 Release 2

- Some fastpath (release1 PLC9)
- VMA Hardware support
- CPU use based on both virtual and CP
- VM/VS1 handshaking

CSC/VM (release 2)

- Relocatable Shared Segments
- Paging Access Method
- CP67 Dynamic Adaptive
- CP67 working set & global LRU
- CP67 Fastpath
- Restructure page supervisor
- Page & Swappable migration
- Q3
- Autolog

SHARE VM Scheduler White Paper

- VM370 is regression from the best of CP67
- Proposed
 - Additional I/O measurements
 - Resource targets include I/O
 - Runlist limit include more than working set
 - Group scheduling

VM370 Release 3

- VMA support for VM with shared segments
- DCSS subset of CSC/VM (but w/o PAM)
- AUTOLOG command

Resource Manager (VM370 3.4)

- CP67 Dynamic Adaptive
- CP67 working set and global lru
- CP67 fastpath
- Restructure page supervisor
- Page & swappable migration
- Q3
- Reliability and cleanup of over 60 modules

VM370 Release 3.8 ECPS

- Enhanced VMA function
- Virtual timer support
- E6 opcodes
 - Kernel code moved to microcode for 10:1 improvement
 - Top 6k bytes of code, approx.80% of CP cpu

VM370 Release 4 SMP

- Dependent on lots of code in resource manager
- Resource manager was 1st priced SCP
- Free SMP code couldn't require priced code
- Moved possibly 80% of code into free base w/o changing RM price

SJR/VM Release 5

- Block Pre-paging (track previous pages)
- SYSPAG
- Single chain for eligible list
- Simple group scheduling
- Restructure IOS for performance & availability
- CMS chained terminal I/O
- Restructure CMS sysgen
 - Multiple file directories in shared segments

SJR/VM Release 5.12

- Extensive timing measurements
 - Runnable & non-runnable measured
 - Restructured Q3 controls
 - IOBLOK queued and service times
 - PAGE I/O queued and service times

VM370 Release 6 (SEPP)

- CMS EDF
- Additional ECPS support for CMS
- S&Y directories in shared segment
- Uniprocessor V=R support

PAM CDF & EDF

- Typical application do some file operation
 - 4k/EDF I/Os for physical I/Os, PAM I/O are no. 4k page transfers

	CPU	I/O	elapsed
4k/EDF	3.72	1958	82
PAM/EDF	3.41	3836	56
PAM/CDF	3.02	3790	52

VM/SP-HPO Release 1

- CP does block 3270 for multiple simulated line
- SMP reworked for single, non-SMP guest
- Support for hardware “protect”
- Significant increase in CP overhead masked by other changes

CP67-3.2 v. HPO2.5

machine	360/67	3081	ratio
MIPS			40-50
pages	105	700	66
users	80	320	4
channels	6	24	4
drums	12m	72m	*6
Page i/o	150	600	4

VM/SP-HPO Release 3.x

- Cache sensitive kernel storage for SMP
- “true” runlist
- 300 mills. Queue drop delay
- Big pages
- SYSPAG

VM/HPO 3.4

- Global LRU

VM Migration Aid

- Originally internal use only for MVS/XA development
 - Originally VM370 killed and all people moved to support MVS/XA
 - Endicott managed to save VM370 product mission but had to recreate group from scratch

Cluster z/VM

- Minor mention from Hillgang 2009 presentation
- From the Annals of Release No Software Before Its Time
- Internal US HONE datacenters (world-wide sales & marketing support) consolidated in silicon valley in mid-70s.
- Implemented load-balancing and fall-over across large loosely-coupled cluster with large shared disk farm.