


LPAR VS VM Preferred Guests

2000 VM/VSE Technical Conference
Session 1D1

Romney White
S/390 Software
Romney@vnet.ibm.com
USIB64SN at IBMMAIL
607-755-8276



Abstract

In non-VM installations, hardware Logical Partitioning (LPAR) support is widely viewed as an alternative to VM/ESA and its Multiple Preferred Guest (MPG) facilities. However, VM aficionados know that there must be a catch, since VM offers so much more. In fact, as this session explains, running VM in an LPAR can be a viable configuration option.

Because LPAR and MPG are based on many of the same technologies and concepts, it isn't surprising that they are often compared and sometimes confused. In this session, we try to present a factual comparison of LPAR and MPG, with an eye to helping customers decide which solution is the right one for them.

Trademarks and Service Marks

- ESCON
- Processor Resource/ Systems Manager
- PR/SM
- S/390
- VM/ESA
- VM/XA

Agenda

- Objectives
- Background
- PR/SM
- Logical Partitioning
- Multiple Preferred Guests
- Comparison
- Conclusions

Objectives

- Outline functions
- Compare and contrast capabilities
- Recommend appropriate use
- Not addressing VM in an LPAR

Background

- VM or LPAR?
 - Longstanding question
- Customer perspective
 - VMer: LPAR is just VM in microcode
 - MVSer: LPAR is VM for production systems
- Both have value: when and how much?

Processor Resource/Systems Manager (PR/SM)

- Hardware function
- Resource partitioning feature
- Replaced VM/XA MHPGS (Multiple High-Performance Guests Support)
- Required for MPG or LPAR

Logical Partitioning

- A response to customer demand
 - Amdahl MDF
 - Larger systems
 - MP effects
 - Pricing models
- Up to 15 partitions/system
- Supported by all S/390 processors (except P/390)
- Exploits PR/SM hardware

Multiple Preferred Guests

- A better way than V=V to run guest production
- Up to six preferred guests
 - One V=R guest
 - SIE I/O Assist
 - Guest recovery
 - Bypass CCW translation
 - Up to six (five) V=F guests
 - SIE I/O Assist
- Exploits PR/SM hardware

Comparisons

- CPU
- Storage
- I/O
- Coupling
- Resource management
- Configuration management
- Performance
- Performance management

Acronyms

- CP - Central Processor
- LP - Logical Processor
- LP - Logical Partition
- HSA - Hardware System Area
- EMIF - ESCON Multiple Image Facility
- CF - Coupling Facility
- ICMF - Integrated Coupling Migration Facility

CPU

LPAR

- Dedicated or Shared
- Dynamic
- Dedicated CPs are reserved
- Some VM assists not available for VM guests

MPG

- Dedicated or Shared
- Dynamic
- Dedicated CPs are dynamic

LPAR Dispatching

- Relative weight per logical processor
- Rolling 32-interval average refreshed every 50 ms
- +/- 1.8% accuracy at full utilization
- Capping is hard limit on a logical CP basis
- Do not cap unnecessarily or if asymmetric logical processor configuration

LPAR Scheduling

- Event-driven
 - Wait detection
 - Spin loop notification (Diagnose X'44')
 - SIGP interpretation
 - Preemption for pending I/O and timer interruptions
- Time-driven

$$\text{time slice} = \frac{25 \text{ ms} * \text{number of shared CPs}}{\text{number of LPs started}}$$

LPAR Scheduling Examples

- Five 10-way LPARs on a 10-way; shared CPs

$$\text{time slice} = \frac{25\text{ms} * 10}{5*10} = 5 \text{ ms}$$

- One 10-way + four 3-way LPARs on a 10-way; shared CPs

$$\text{time slice} = \frac{25\text{ms} * 10}{1*10+4*3} = 11 \text{ ms}$$

MPG Dispatching and Scheduling

- More complicated due to
 - Potential for 1000s of users
 - Interactive service needs
 - short time slice
 - consistent response time
 - CP functions
- (wave hands a lot)

Storage

LPAR

- Initial and reserved allocations
- Dynamically configurable
- Some complexity
- No sharing

MPG

- Default and maximum allocations
- Only Expanded Storage dynamic
- Logoff/Logon considerations

LPAR Storage Planning

HSA		16M
Unused Addressability		
LP_B Reserved	48M	
LP_B Initial	192M	
LP_A Reserved	48M	
LP_A Initial	256M	

2 x storage size (1024M)

512M Real

LPAR Storage Allocation

HSA	16M
LP_A Reserved	48M
LP_B Initial	192M
LP_A Initial	256M

real storage (512M)

MPG Storage Use

HSA	
V=V Area	
CP Nucleus	
V=F Area	VM4
	VM3
	VM2
V=R Free	
V=R Area	VM1

I/O

LPAR

- Dedicated channel paths
 - may be configurable
- Shared ESCON paths (EMIF)
- I/O device partitioning

MPG

- Dedicated or shared devices
- Reconfigure at device level
- I/O throttling
- Virtual devices
- Minidisk cache

Coupling

LPAR

- CF LPARs
- Sender channels
 - dedicated
 - shared
- Receiver channels
 - dedicated
- ICMF

VM

- VM/ESA V2R3 includes guest coupling support
 - CF virtual machines
 - virtual channels
- Session 29H: User Experience Testing an OS/390 Parallel Sysplex Under VM/ESA (Thursday, 8:00 am)

Resource Management

LPAR

- Dedicated CPs
- Processor weights
 - relative
- Resource capping
 - hard limit
- Time-driven or event-driven

MPG

- Dedicated CPs (dynamic)
- SHARE settings
 - relative
 - absolute
- SHARE limits
 - hard
 - soft

Configuration Management

LPAR

- IOCP
- HCD
- LPxxx frames
- Some complex rules

MPG

- User directory
- Dynamic I/O
- ATTACH, DETACH, DEFINE, LINK

Performance

LPAR

- Close to native
- CPU sharing and capping can constrain
- No penalty for shared paths

MPG

- Close to native
- CPU sharing and capping can constrain
- I/O sharing considerations
- Virtual devices and MDC can improve over native

Performance Management

LPAR

- Hardware data
- SAD frame
- RMF or VM/PRF in a Logical Partition

MPG

- VM monitor data
- INDICATE
- RTM/ESA, VM/PRF

Measurement Basis: Timers

- **TOD Clock**
 - Moves with real time
 - LPAR and VM/ESA allow guest to set
- **Clock Comparator**
 - Compares with TOD Clock
 - Moves with real time
- **Processor Timer**
 - Runs when dispatched
 - Stopped when pre-empted

Measurement Problem Basis: Timers

- Processor timer stopped in involuntary wait state
- Results vary depending on LPAR configuration and processor contention
- Some facilities (e.g., RMF, CP INDICATE) calculate CPU utilization incorrectly
- Others (e.g., RTM/ESA, VM/PRF) correctly use elapsed time as denominator in computations

Observations (1)

- Many VM systems run in LPARs
 - Fewer variables
 - Configuration
 - Performance
 - Operations viewpoint
 - dynamic = uncontrolled
 - flexible = unpredictable

Observations (2)

- Some MVS systems run beside VM LPARs rather than as guests
 - MVS orientation
 - Historical VM instability (now folklore)
 - Guest survival works very well
 - Historical performance issue
 - SIE made a dramatic difference
 - Don't notice MVS if CPUs dedicated

Observations (3)

- LPAR primarily for MVS shops
 - Few images
 - Near native performance
 - Thin layer
 - dispatcher
 - resource manager
 - Hardware sharing
 - independent
 - secure

Conclusions (1)

- Use LPAR if
 - No VM skills
 - No need for VM facilities
 - large numbers of users
 - flexibility
 - virtualization
 - CMS
 - Production sysplex

Conclusions (2)

- Use VM MPG if
 - VM already installed
 - Performance benefits from
 - virtual devices
 - virtual disk in storage
 - minidisk cache
 - Variable configuration
 - Resources available to dedicate
 - Number of CPUs > number of CPs

References

- Processor Resource/ Systems Manager Planning Guide GA22-7236
- VM/ESA: Running Guest Operating Systems SC24-5755