

# Thin Provisioning for DB2: The New Storage Provisioning Paradigm

Paul Pendle  
EMC Corporation


August 7, 2012 3:00PM-4:00PM  
Session Number: 12065



# Agenda

- Traditional storage provisioning
  - Challenges
  - Trends
- Thin provisioning concepts
- DB2 considerations
- Recommendations
- Conclusions

# Traditional Storage Provisioning



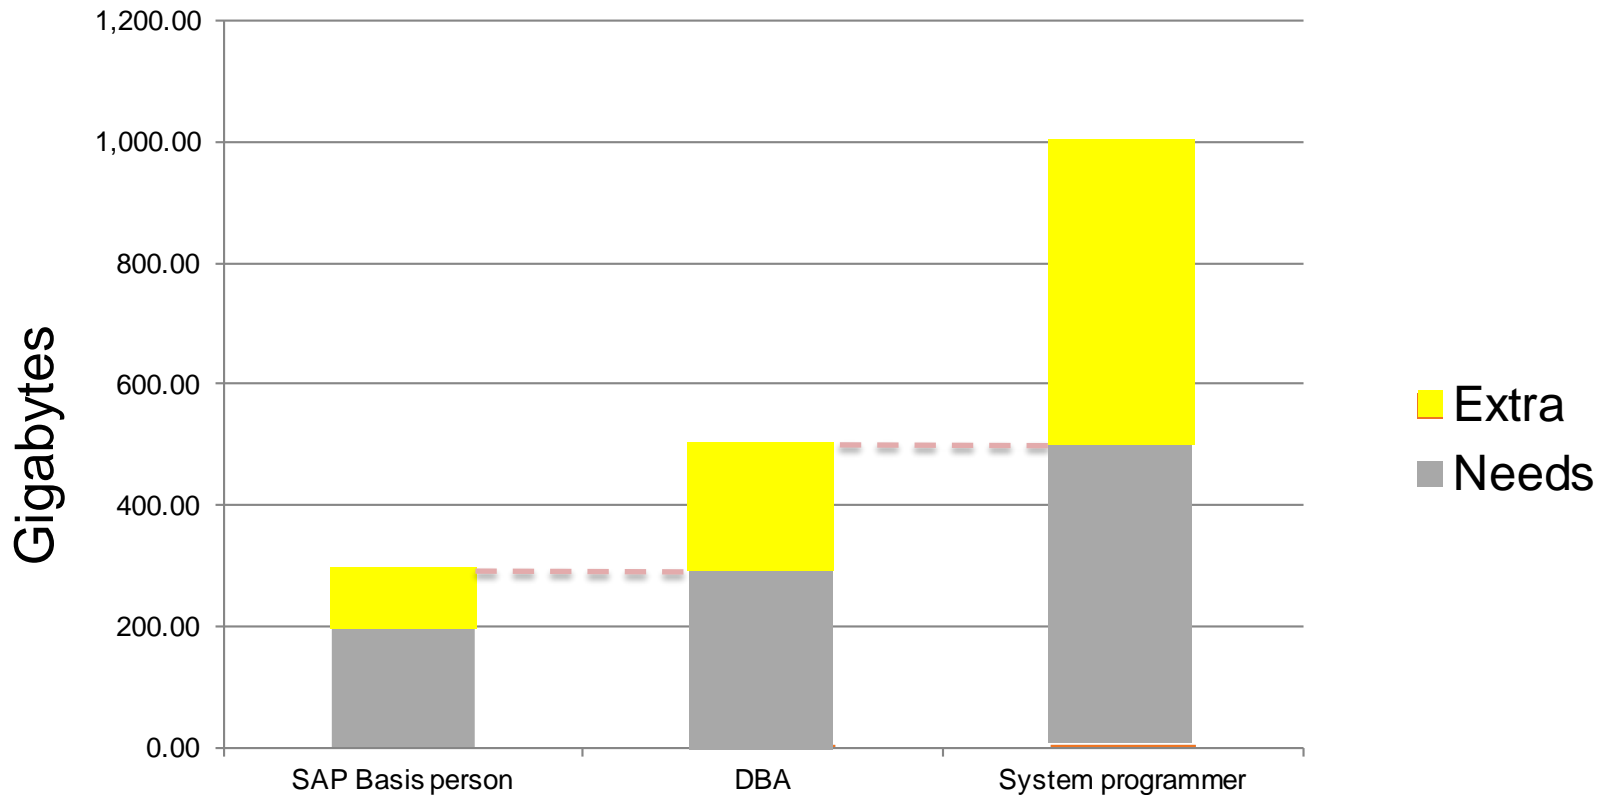
# Traditional Provisioning Issues

- It's a slow process
- Overallocation opportunities
- Complexity of storage planning
- Asymmetric configurations
- Very large physical disks
  - Sibling-pend issues
  - Reduced IOPS density
- Vertical striping DB2 partitions

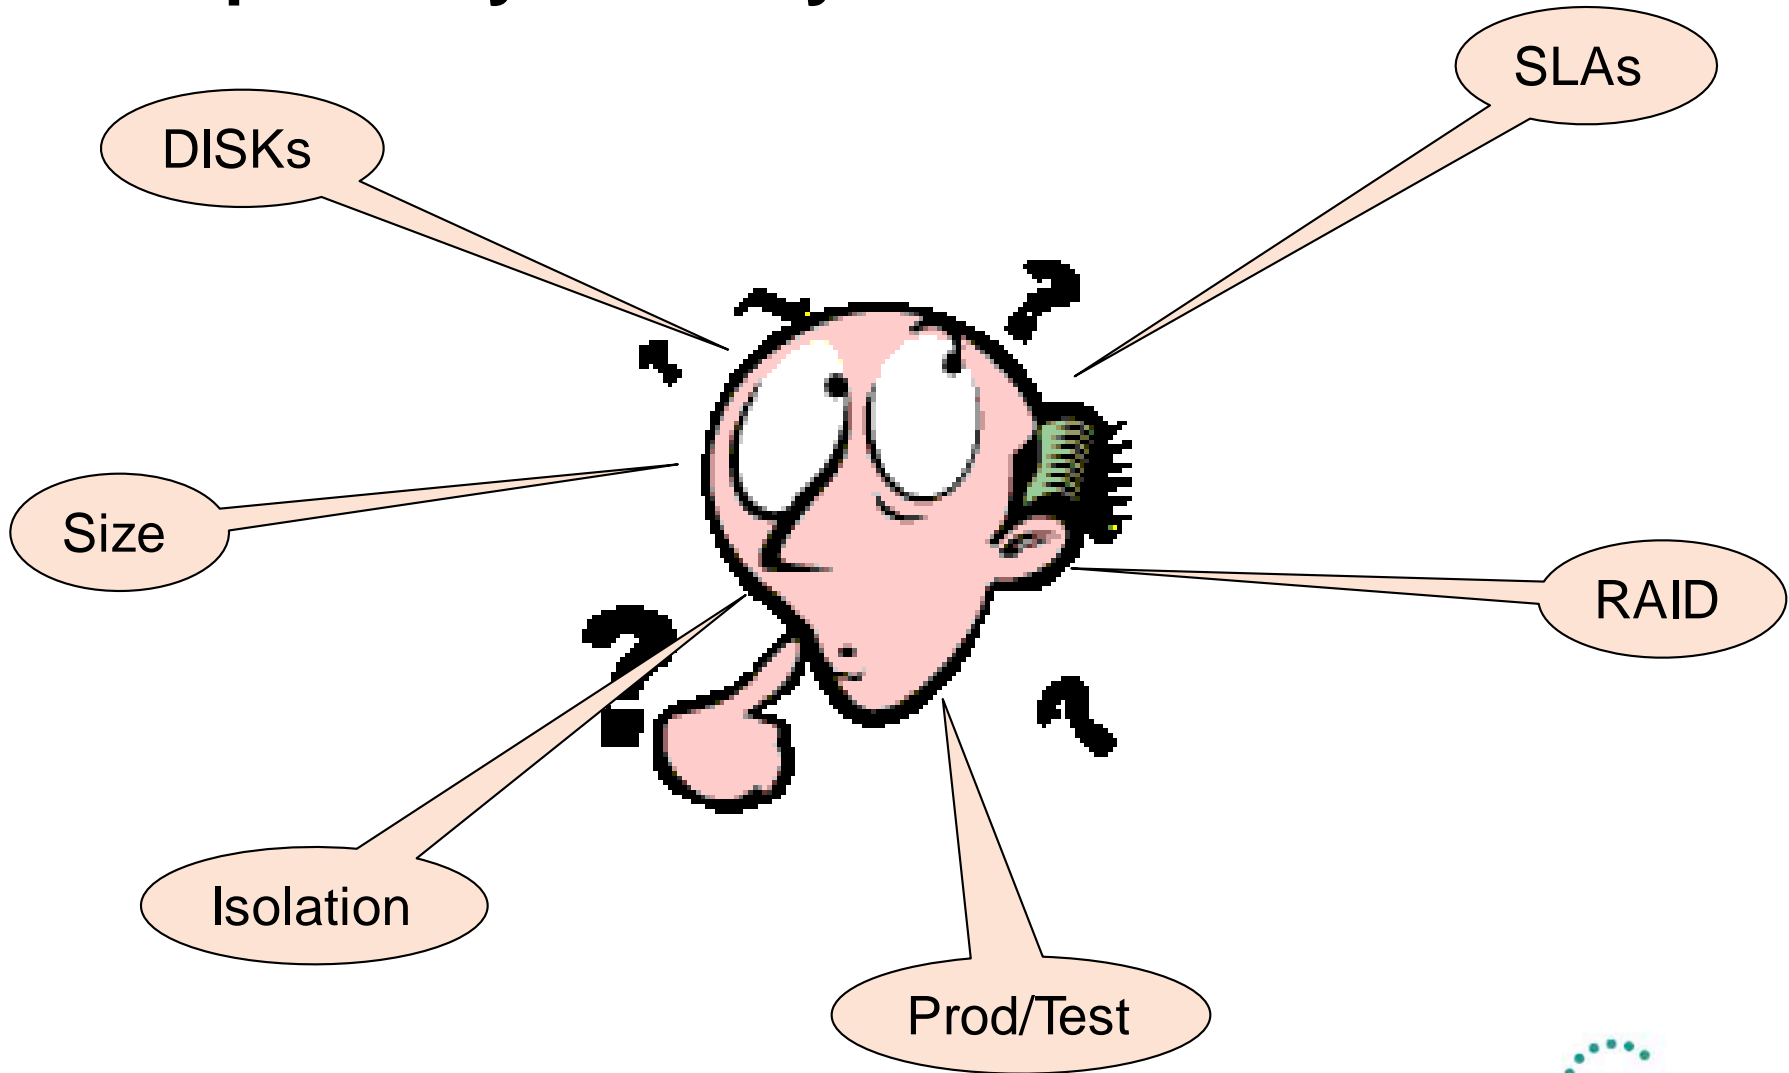
# Slow Process


- Purchase orders
- Possible updates to IODF
- SMS
  - Storage group
  - ACS routine (maybe)
- System backup
  - Copy pools
- SRDF/PPRC
  - Replication management
- Change control!!

# Overallocation Opportunities



# Complex Physical Layout






**SHARE**  
Technology · Connections · Results

# Asymmetric Layouts

Application A

Application B



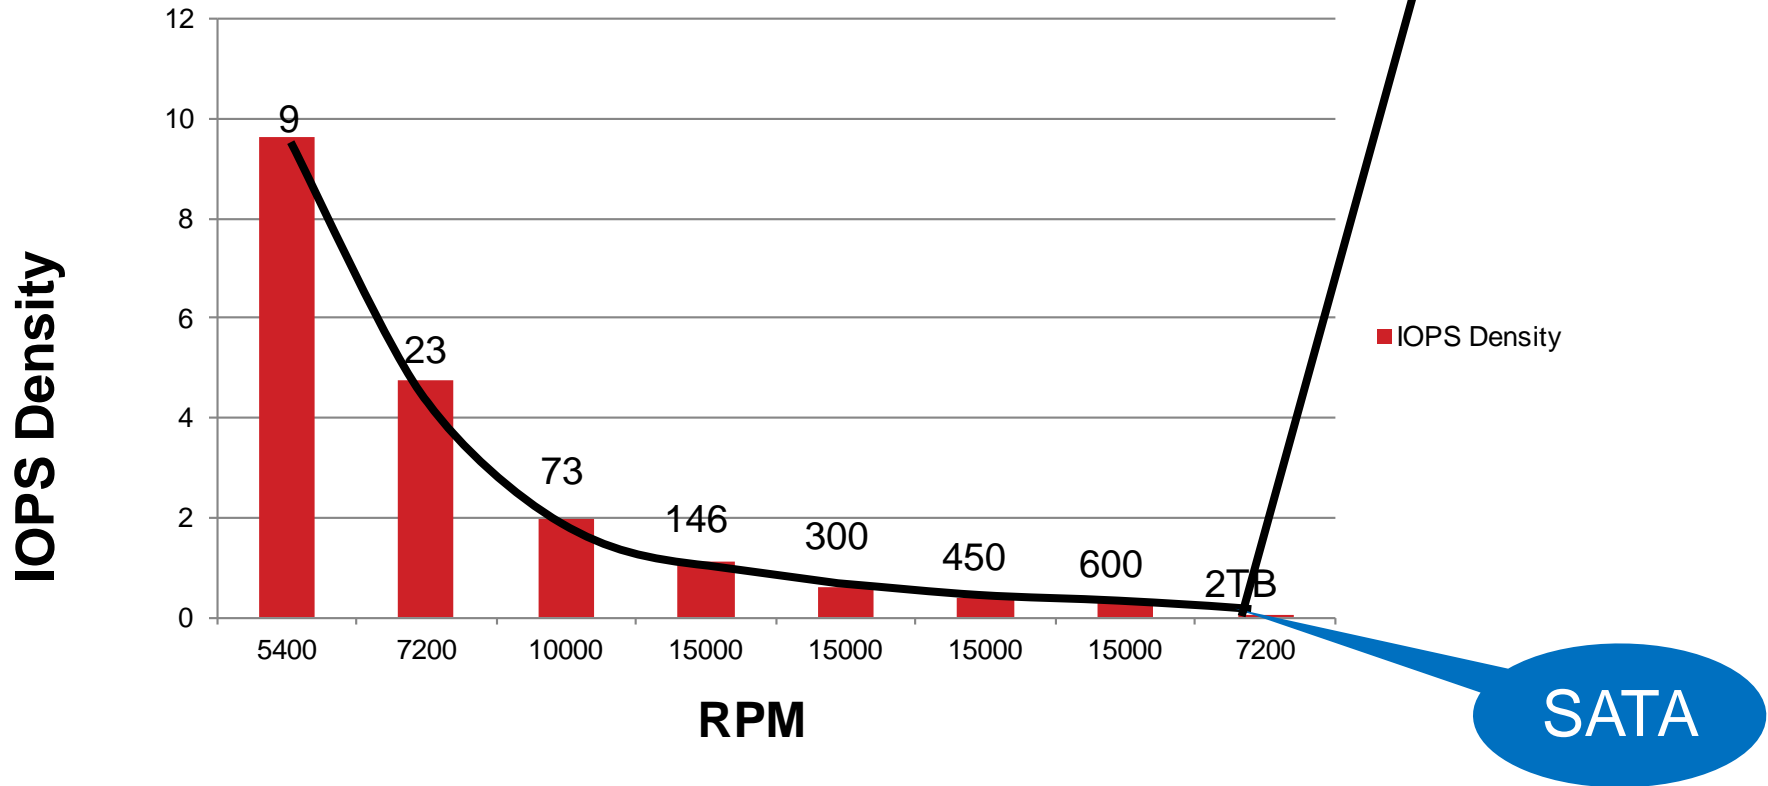


# Very Large Physical Disks

- FC drives are now 600GB
- How to avoid sibling pend / sibling contention?
- How to segregate work loads on the back-end?
  - Production and test sharing the same disk
- No faster in terms of RPM than 10 years ago

# IOPS Density Trend for Disk Drives

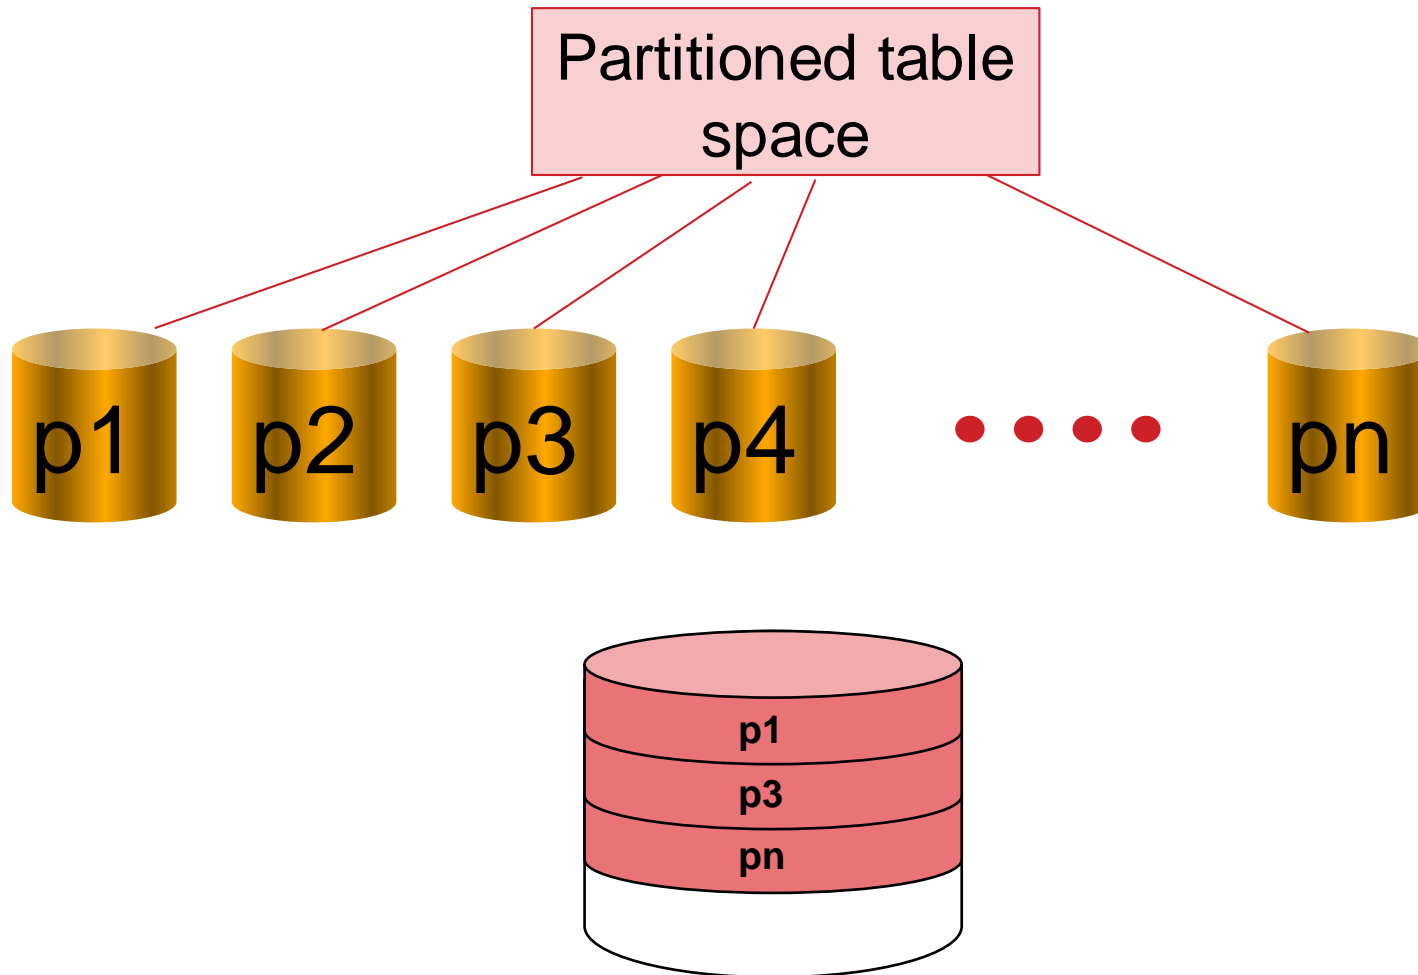
$$\text{IOPS Density} = \text{IOPS capability} / \text{Size GB}$$



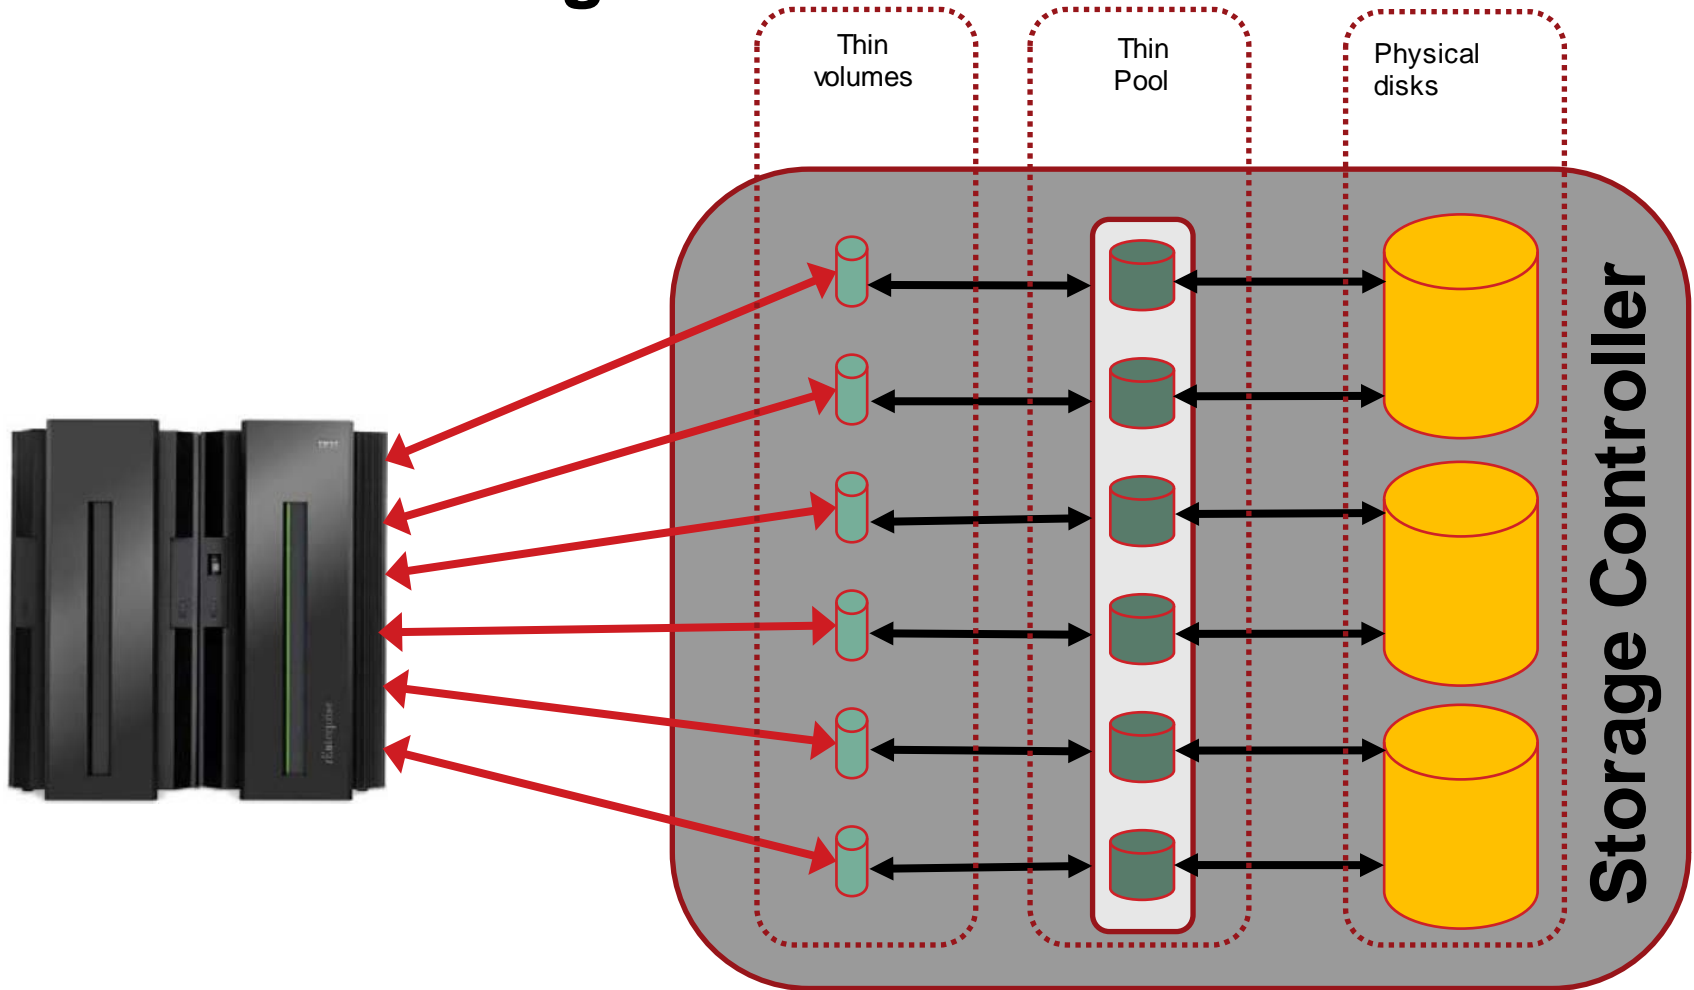
# Large Volume Performance

- UCB constraints forcing you to larger volumes > MOD54
- EAVs
- How to get the right service levels?
- How to get enough disks behind the volume to support the I/O activity


# Vertical Striping Table Spaces




# Thin Provisioning



# Thin Provisioning and Tiered Storage



# Thin Provisioning and Tiered Controllers





# Thin Provisioning Advantage

Traditional Provisioning	Thin Provisioning
Slow process	Lightning fast, no host activity
Overallocation	Exact allocation
Complexity	Simplicity
Asymmetric configurations	Balanced configurations
Large physical disks	Wide striping
Vertical striping for DB2 partitions	Simplified table space allocation




# Ease of Provisioning

10 Thin Mod 27s



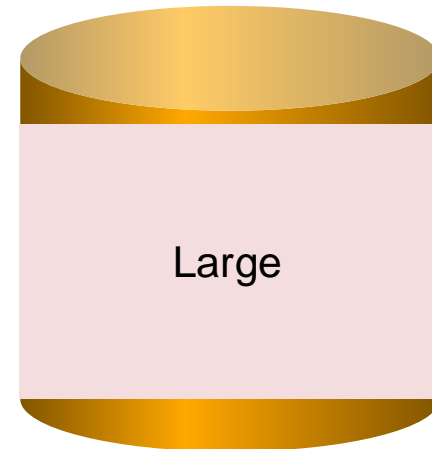
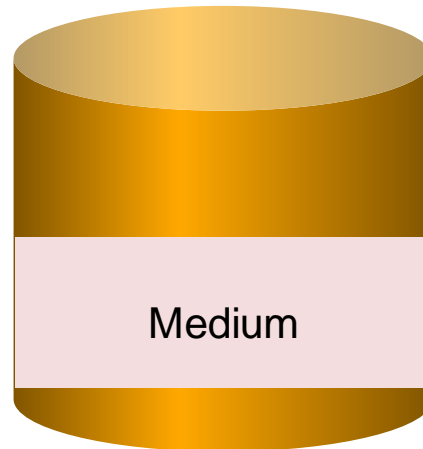
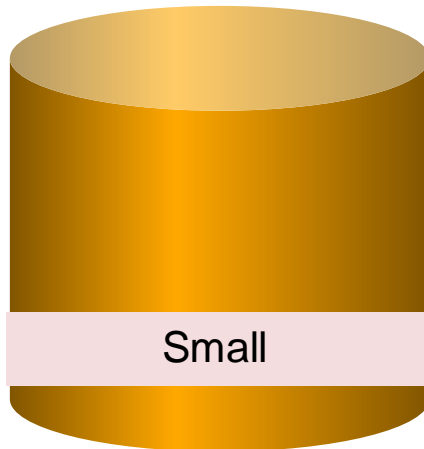
10 Thick Mod 9s



# DB2 Table Space Allocation

One size fits all:  
PRIQTY (average)  
SECQTY (average)

PRIQTY = -1  
DSNZPARM  
TSQTY(avg)



# Thin Provisioning Introduction


- “Thin Provisioning” is the industry term
  - Virtualization layer
  - Storage on demand
  - Wide striping
- Storage Vendors
  - EMC: Virtual Provisioning
  - IBM: Extent Space Efficient
  - Hitachi: Dynamic Provisioning
- **“Storage on demand”**

# Thin Provisioning Operations


# Reads from Thin Devices

- Thin devices are virtual devices and contain only pointers to assigned blocks on data devices in a thin pool
- When a read is performed on a thin device, the data being read is retrieved from the appropriate data devices
- Reads directed to an area of a thin device that has not been assigned do not trigger assignment operations
  - Reading an unwritten block returns “standard record 0”

# Writes to Thin Devices



# Wide Striping



**Aggregate IOPS capability**

$$\sum_{X=1}^n \text{IOPS (disk}_x\text{)}$$

# Chunk size

- EMC Virtual Provisioning
  - 684KB (called a track group)
- IBM ESE (Extent Space Efficient)
  - 1GB == 1113 Cyls == Mod 1 (called an extent)
- Hitachi Dynamic Provisioning
  - 38MB (called a “page”)
- Considerations
  - Performance
  - Tiering
  - Space consumption




# Thin Provisioning Considerations

- Pool management
  - Re-balancing
  - Draining devices
  - Space de-allocation
- Thin-friendly applications/processes
- Thin-unfriendly applications/processes
- Oversubscription
  - Monitoring
- Pre-allocating volumes

# Thin Pool Rebalancing


- Pool Rebalancing redistributes pool capacity
- Used mostly when expanding a pool
- Avoids hot spots on newly added devices



Thin Pool

# Draining Pool Devices

- What if a pool is over-allocated?
  - Changing business requirements could leave a pool overallocated
  - Application staff overstated their requirements
- Draining devices can help



Thin Pool

# Space De-allocation

- Space de-allocation/reclamation returns unused thin extents back to the pool
- Reclamation operations are run against individual thin devices
- Each thin device might have space that was once written to but is not longer used
  - Dataset moved
  - DB2 online REORG
  - Dataset deleted
  - &&TEMP files


# Thin-Friendly Applications/Programs

- Thin-friendly:
  - Doesn't write a lot of non-data
  - Doesn't pre-format space
  - Doesn't write a lot of meta-data throughout the dataset
- DB2 table space allocations
  - V8: formats 2 cylinders (2 > PRIQTY)
  - V9 and V10: formats 16 cylinders (16 > PRIQTY)
- ICKDSF quick INIT

# Thin-Unfriendly Applications/Processes


- DB2 active logs
- DB2 10 hash table spaces
- ADABAS Associator / Data storage

# Monitoring Oversubscription



# Space Pre-Allocation

- You like wide-striping and balance
- You like the ease of management
- You don't want over-subscription and its risk
- You don't want to manage the pool space
- You should “pre-allocate”







# Performance considerations

- Indirection layer overhead
- New chunk overhead
- Sequential vs. random
- Wide striping benefit

# Recommendations

- Evaluate your vendor's thin provisioning offering
- Implement thin provisioning with a tiering solution
  - Increase IOPS density
  - Target active DB2 data to SSDs devices
- Keep some SMS storage groups “thick”
  - DSNDB07 – for temp data
  - &&TEMP data
  - DB2 LOG data
- Consider impact on remote replication solutions
  - SRDF, PPRC, etc


# Thin Provisioning Summary

- Storage virtualization for mainframe
  - Helps performance through wide striping
  - Eases management and provisioning
  - Balances capacity and eliminates skew
  - Increases utilization
- The foundation of tiering strategy
  - In a controller
  - Between controllers
  - Into the cloud

# THANK YOU

Paul Pendle: [paul.pendle@emc.com](mailto:paul.pendle@emc.com)

Session: 12065



**SHARE**  
Technology • Connections • Results



Complete your sessions evaluation online at [SHARE.org/AnaheimEval](http://SHARE.org/AnaheimEval)