

# 75 Ways to SAVE IO & 25 Ways to Save CPU

**Lee Goddard**  
*DBI Software*

2013-11-19 | Platform: LUW



## Disclaimer

- All is my opinion... take it for what it's worth...
- Your mileage will vary
- Facts only valid just prior to writing them down...
  
- Study on Retaining Information:
  - See or hear = 20-30%
  - Discuss = 70%
  - Teach = 90-95%

## Agenda

- Session Objectives
- Introduction
- Not just a checklist but a systematic methodology for profiling your system and then ATTACKING the problem
- How do I save I/O – 75 Ways Discussion
  - System, Storage, Middleware (incl. DB)
- How do I save CPU – 25 Ways Discussion
  - System, Storage, Middleware (incl. DB)

## Session Objectives

- Demonstrate a Systematic Methodology for profiling your system
- Discuss issues
- Provide a roadmap / checklist to **SAVE resources**
  - CPU
  - Memory
  - I/O
  - Time
  - Energy
  - Upgrades
  - Software costs
  - ... and more

## Background, Role and Contact Information

- Lee Goddard Indiana, USA
- Served 26+ years IBM
  - Executive / Sr. Certified Consulting SWIT Specialist) and Team Lead Midwest BI Best Practice Team and Team Lead for IM Tech Team
- Last 2+ decades, Lee led > **200 clients projects**
  - w/80+ being Terabyte Class 4 + 12 Petabyte level designs
- Certified IBM Teraplex VLDB Integration Specialist – Certified in HW, SW, & Svc w/ Info. Mgmt. and AIX tech certs
- Currently: **Chief Performance Architect** at DBI Software
- [www.linkedin.com/in/leegoddardin/](http://www.linkedin.com/in/leegoddardin/)
- [Lee.Goddard@dbiSoftware.com](mailto:Lee.Goddard@dbiSoftware.com) petabytedw@yahoo.com

## Who is in the Audience???

- DBA's?
- Architects?
- Managers?
- CXX?
  
- Performance > 30% of job?

## Who is using a Virtualized Environments????

- LPARs?
- zLinux? IFLs?
- VMWare?
- Microsoft?
  
- Others?

## DB2Night Show™

- How many know about DBI's free DB2 educations on the DB2Night Show™?
- How many people know about the replays?
- How many people know about the LUW and DB2 on z?

## DB2Night Show™ Hundreds of hours of Free Education

- 15 Shows Scheduled for Season 5!
- LUW (119) and DB2 for zOS (41)
  - <http://www.dbisoftware.com/db2nightshow/>
- **Replays: almost 400,000 downloads!**
  - <http://www.dbisoftware.com/blog/db2nightshow.php>

### **SEASON #5: Schedule and Registration Links**

## What are the performance issues? Ref: DB2Night Show™

### QUICKPOLL

## The DB2 LUW database I care the most about tends to be:

Poll Results (multiple answers allowed):



# Taming the Data Tsunami: 17 Laws of Building Petabyte Level DB Systems & excerpts from: I like Big Data and I Cannot Lie... Getting Your Arms Around the Chaos

**Lee Goddard**      **DB2Night Show™**  
*dbi Software*      *Season 5 - #118*



*Also - IDUG Australia 2013:*

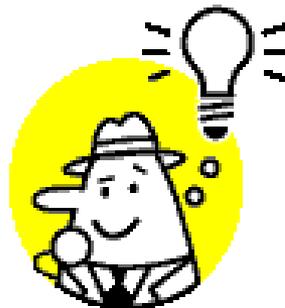
Session Code: C12 ½ Day Workshop

2013-9-12 8:30-12:30 | Platform: LUW – I (Rm 11.03)

*Shameless Plug* 😊

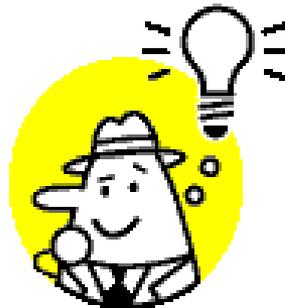
# 17 Laws of Building Terabyte Class DB Systems

1. 3 Absolutes
2. Requirements need to be well understood
3. Governance, sponsorship and funding matters
4. You cannot boil the ocean
5. You can only compress time so far
6. 3 Most Important things in MPP (Scalability)
7. I/O Subsystem matters
8. Crayne's Law
9. 3 things clients always ask



## 17 Laws of Building TB Class DB Sys. (cont.)

10. There will be speed bumps
11. Few uses FS compression or near line storage
12. The Moose on the table
13. Balanced Configuration Unit (BCU)
14. Be: Proactive, continual and methodic capacity planning matters
15. Organization matters
16. Relationships matters
17. Leadership matters



## Observations:

- As leaders in the industry....
  - WE need to do more faster...
  - WE need more test systems!
  - WE need **more research projects**
  - WE need more testing and **discipline** and less
  - WE need to collect **value of a transaction & costs**
  - WE need to **mentor** more
  - WE need to **write and give back** more
  - WE need to grow more **SKILLS!** ... AND NOT settle for 'black boxes' e.g. Storage ... "You don't understand, the Storage guy works for another manager and another agenda"

## First you need a plan... a process... a roadmap!



1. Triage which system first
2. Determine what needs to be fixed on that system
  - Sort and fix the 'BIG ROCKs FIRST!'
3. Fix the issue!

**Here**

**There**

## First you need a plan... a process... a roadmap

- **tri·age**

<http://dictionary.reference.com/browse/triage?s=t>

- [tree-ahzh] Show IPA noun, adjective, verb, tri·aged, tri·ag·ing.

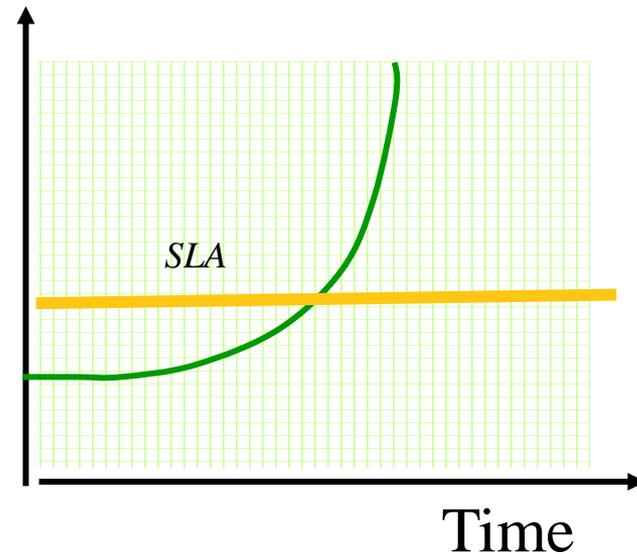
- noun

- the process of sorting victims, as of a battle or disaster, to determine medical priority in order to increase the number of survivors.
- **THE DETERMINATION OF PRIORITIES FOR ACTION IN AN EMERGENCY.**

- **Origin:** 1925–30; < French: sorting, equivalent to tri ( er ) to sort (see try) + -age -age

## Why Bother?

- **Crayne's Law: All computers wait at the same speed**
  - "The Serious Assembler" circa 1985 for 8086/8088
- Once there is a bottleneck... **....OR LOCKS!**
- **EVERYTHING WAITS**
  - Network
  - HW
  - Operating System
  - Middleware
    - DB
    - Messaging
    - Etc.
  - Application



## Audience Quiz???

- What is the best way of conserving CPU and IO resources?

Tell me about an I/O issue you've gotten over?

Tell me about an CPU issue you've gotten over?

## First you need a plan... a process... a roadmap

### 1. Triage which system first

PROCESS?

**2007** Score Blog - [http://www.dbisoftware.com/blog/Scott\\_Hayes.php?show=0,08,2007](http://www.dbisoftware.com/blog/Scott_Hayes.php?show=0,08,2007)

### 2. Determine what needs to be fixed on that system

- Sort and fix the ‘BIG ROCKS FIRST!’
- <http://www.youtube.com/watch?v=to-DyF3AL3o>

### 3. Fix the issue!

## Being Proactive Matters: The Value of a Good Triage Process

- What about the 2:00 A.M. Call or page that backup did not run or there is a production problem!?
  - Is your head groggy / fuzzy / not thinking straight at 2:00 AM?
  - Chaos and long weekends or go home on time and kiss the spouse and tickle your children



OR →



## Being Proactive Matters: The Value of a Good Triage Process

- What about the 2:00 A.M. Call or page that backup did not run or there is a production problem!?
  - Is your head groggy / fuzzy / not thinking straight at 2:00 AM?
  - Chaos and long weekends or go home and kiss the spouse and tickle your children
- Need a well defined or written process so you don't make mistakes
  - Don't you think it's important to have a process?
  - **What items are important?**
  - What system issues need identified and escalated?
- Save HW upgrade? Black Friday!?
  - 132 categories of savings!

## First you need a plan... a process... a roadmap

### 1. Triage which system first

PROCESS?

2007 Score Blog - [http://www.dbisoftware.com/blog/Scott\\_Hayes.php?show=0,08,2007](http://www.dbisoftware.com/blog/Scott_Hayes.php?show=0,08,2007)

### 2. Determine what needs to be fixed on that system

- Sort and find the 'BIG ROCKS FIRST!' - Steven Covey
- <http://www.youtube.com/watch?v=to-DyF3AL3o>

### 3. Fix the issue!

## Audience Poll Question???

- Who has done a research projects?
  - Last 6 months?
  - Why?
  - What have you learned?
  - Proactive vs Reactive?

## First you need a plan... a process... a roadmap

### 1. Triage which system first

PROCESS?

**2007** Score Blog - [http://www.dbisoftware.com/blog/Scott\\_Hayes.php?show=0,08,2007](http://www.dbisoftware.com/blog/Scott_Hayes.php?show=0,08,2007)

### 2. Determine what needs to be fixed on that system

- Sort and fix the ‘BIG ROCKS FIRST!’
- Big Rocks First - <http://www.youtube.com/watch?v=to-DyF3AL3o>

### 3. Fix the issue!

- Even if you know what needs to be fixed, wouldn't it be nice to have a checklist sorted by type of fix? So STAY TUNED!

# BI Infrastructure Optimization – Value of Balance

## Traditional Approach



- 30%+ overhead on I/O
- High I/O waits
- Lower process utilization
- BI performance problems are 60%+ I/O related

## ISAS – InfoSphere Smart Analytic Sys = Balanced (BCU) Approach



Query performance improvements of 50%+  
UK

\*base slide via Simon Field IBM

## Indexes, Indexes, Indexes

- Unique / non-unique
- Compound
- Clustering
- Allow Reverse scans
- Collect statistics for indexes w/ Extended detail statistics
- Ascending
- Descending
- % free
- % of minimum amount of used space to be left on index page
- Include column
- Others...

**Webinar Index Design Best Practices and Case Studies?**



<http://www.dbisoftware.com/media/20090812INDEXES.wmv>

## Access Path Matters: Doing Less Work w/Better Path

- Optimizer level
  - Influence which index?
  - Which path?
- Who is using Index Advisor?
- Who is using Workload Query Tuner?
- Other tools?

## Disk Layout Matters: Case Study - Better Performance

- 20% improvement by laying out arrays differently
- Ratio of Disk to CPU
  - Observation: 4-6 usually
  - Observation: 70% of 'emergency phone calls' for performance issues – I/O
  - Best Practice: 10 per Proc!
- 50% improvement in large tables
- Daily MQT's recalc > 24 hours = 12 min.!
- Don't just accept the virtual disk they give you! 

## Efficiency Matters and NOT Ratios

- In an OLTP system How many rows should you have to access before you get your row needed?
- Expert Advice
  - E.g. **Advice: Index Read Efficiency**
  - <http://www.dbisoftware.com/brother-eagle/advice/db2dbiref.php>

## Efficiency Matters and NOT Ratios

- In an OLTP system How many rows should you have to access before you get your row needed?
- Expert Advice
  - E.g. Advice: Index Read Efficiency
  - <http://www.dbisoftware.com/brother-eagle/advice/db2dbiref.php>



- If you had 80# of pack
- 50# of Sand (aggregated sub sec queries - granules represent the small queries!)

***“PHYSICAL DESIGN FLAWS OBFUSCATED BY LARGE MEMORY POOLS ARE THE NUMBER ONE CPU KILLER IN AMERICA AND AROUND THE GLOBE” Scott Hayes***

[http://www.dbisoftware.com/blog/db2\\_performance.php?id=102](http://www.dbisoftware.com/blog/db2_performance.php?id=102)

## How do you know where the GOAL LINE IS?

- Observation:
  - MOST teams do **NOT** have a good base line...



## How many I/O's can your CPU Generate?

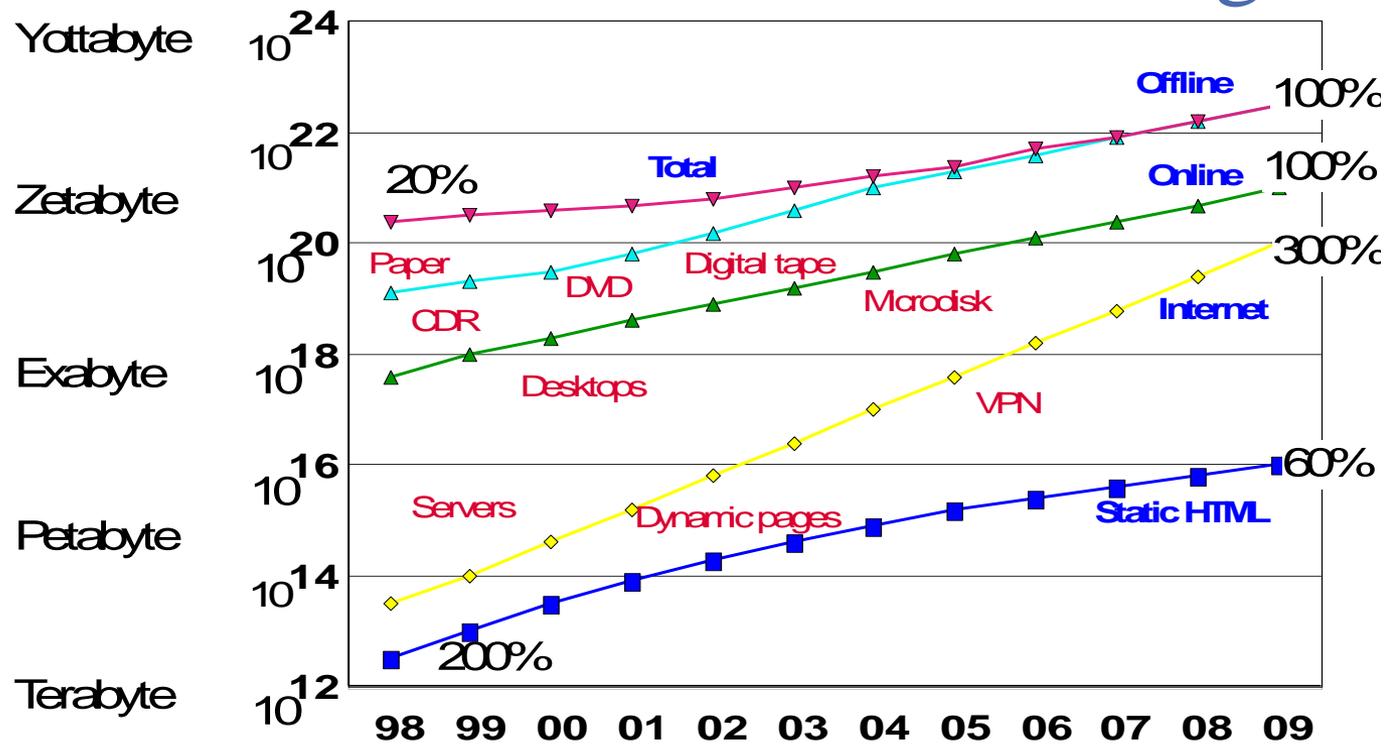
- Do you know how many I/O's your CPU can generate?
  - Have you tested?
  - Have you researched – TPC-C, TPC-H, Specint, other benchmarks?
    - Do you realize that it they can generate X,000,000 I/O's with X HW
      - Why not divide the I/O's by the # CPUs = MAX\_IO\_PER\_POWERX

## How many I/O's can your IO Subsystem Generate?

- \*\*\* Send me an e-mail  
→ Spreadsheet reference
- What EXERCISER programs do you use?
  - nstress
    - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power+Systems/page/nstress>

	A	B	C	D	E	F	G	H	I	J	K	L	M
1		IOPS per RAID Group				Total IOPS per disk types							
2		EFD	15K SAS	10K SAS	7200K SAS	RG #	EFD	RG #	15K SAS	RG #	10K SAS	RG #	7200K SAS
3	RAID 1 (4+4)	23,800	1,224	816	544	2	47,600	-	-	-	-	-	-
4	RAID 1 (4+4)	21,000	1,080	720	480	2	42,000	-	-	-	-	-	-
5	RAID 5 (4+1)	13,563	698	465	310	-	-	2	1,395	-	-	-	-
6	RAID 5 (4+1)	10,938	563	375	250	-	-	-	-	2	750	-	-
7	RAID 5 (8+1)	24,413	1,256	837	558	-	-	2	2,511	-	-	-	-
8	RAID 5 (8+1)	19,688	1,013	675	450	-	-	-	-	2	1,350	-	-
9	RAID 6 (6+2)	21,000	1,080	720	480	-	-	-	-	-	-	-	2
10	RAID 6 (6+2)	16,333	840	560	373	-	-	-	-	-	-	-	2
11	RAID 6 (14+2)	42,000	2,160	1,440	960	-	-	-	-	-	-	-	2
12	RAID 6 (14+2)	32,667	1,680	1,120	747	-	-	-	-	-	-	-	2
13													
14		Formulas											
15		Manual Entry											
16		Data Points											
17													
18													
19													

# Think big, plan ahead...Faster Gargantuan volume increases are coming...



Internet changes from static pages to databases, channels, VPN,  
 Online data grows with pervasive devices; accessible from afar  
 CDR, DVD, digital tape take over the analog world

## I/O Profile Differences Between OLTP, DW, Big Data, OLAP, BLU

- OLTP
  - CPU per tx low
  - Indexing – want to be very efficient IREF\* <10
  - Performance Objects
    - Indexes
      - Lower optimization
- DW / Big Data
  - CPU per query very high
  - Good design
  - Performance Objects
    - Indexing
      - High optimization
      - Layered Data Architecture (Hot Warm Cold)
    - MQTs
    - Partitioning

## I/O Profile Differences Between OLTP, DW, Big Data, OLAP, BLU

Characteristic	OLTP	DW	Big Data	OLAP	BLU
CPU					
Memory					
I/O					
# Users					
Size of Query					

# Profile Differences Between OLTP, DW, Big Data, OLAP, BLU

Characteristic	OLTP	DW	Big Data	OLAP	BLU
CPU – Sys Lev	High	High	High	Medium	High
per tx	Low	Very High	High	Very Low	Med-High
Memory	Low	High	Depends	Very High	High
per tx	Low	High	High-VH	Low–Med	Med-High
I/O	Very Low	Very High	Very High	Very Low	Low-Med
# Users	Very High	Low to Medium	Very Low	High	Med-High
Size of Query	Very Small	Very Large	Very Large	Small – Medium	Small – V Large
ETL	Short tx	VH - Load batch / N Real Time	H Parallel / Light Schema	INTENSE	Med-High

# You Need a SOLID Foundation



## Sizing Systems

*Why don't we spend more time planning?*

**Q:** Why does your systems ran out of 'gas' to early?

How do you do capacity planning?



## How do you monitor CPU / IO today?

- NMON?
- db2top?
- Scripts?
- Tools?
  - Db2batch
- Auto?
- Config advisor?
- Presets like BLU, SAP?
- How do you preserve history?
- Trend / graph?
- Triage?
- Methodology?
- Roadmap?

## What FEATURES MATTER for I/O?

...BY Version.Release?

Stay tuned... DB2Night Show™ = JAN

## V8.1 Performance Tuning

- "[DB2 Universal Database Enterprise Server Edition V8.1: Basic Performance Tuning Guidelines](#)"
- "[Best practices for tuning DB2 UDB v8.1 and its databases](#)" (developerWorks: April 2004): Get optimal performance out of your DB2 UDB database and its applications.
- "[A Quick Reference for Tuning DB2 Universal Database EEE](#)" (developerWorks, May 2002)
- "[DB2 Tuning Tips for OLTP Applications](#)" (developerWorks, January 2002)
- "[Initial Tuning and Design Considerations](#)" (developerWorks, May 2002)
- "[Setting the Volume to Eleven - Tuning DB2 Configuration Parameters](#)" (developerWorks: April 2002)
- "[Top 10 performance tips](#)" (developerWorks, March 2004)
- Redbooks:
  - <http://www.redbooks.ibm.com/abstracts/sg246012.html>
  - <http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246949.html> - DB2 for Content Manager

## Your Selection of SQL Matters: MAJOR I/O SAVINGS

- ROLLUP
- CUBE
- Grouping Sets
- Grouping Set - 15 level Benchmark 2X processors vs Oracle but 1,000x FASTER!!!!!!



**SINGLE PASS OF THE TABLE!**

## DB2 V10.1 Days

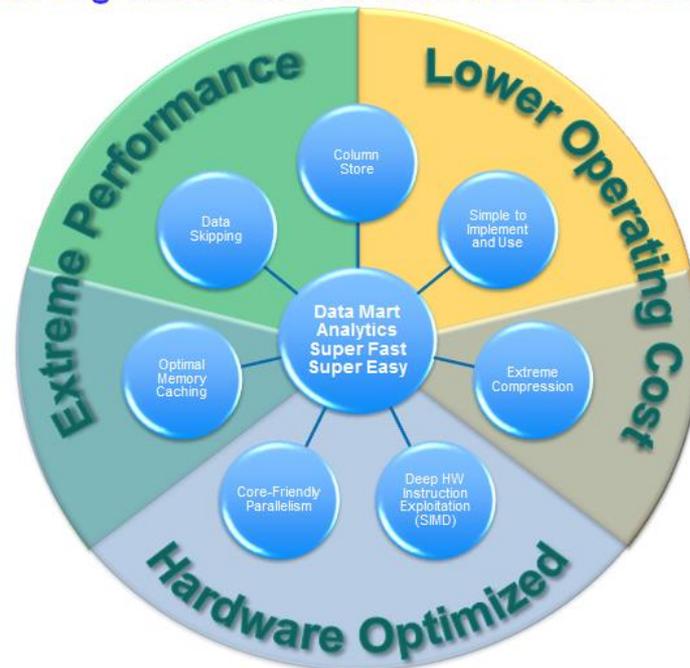
- Still more compression
- Multi-Temperature Storage
  - Included in DB2 ESE, AESE, Enterprise Developer Edition
- Time Travel queries
- Announcement:
  - <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS212-074>

# IBM's Kepler DB2 V10.5 - 7 Big Ideas - How many Save I/O and CPU?



## The Seven Big Ideas of DB2 with BLU Acceleration

- STILL MORE COMPRESSION
  - COLUMN &...
- Data Skipping... and more ... see MANY sessions.



## Performance and Tuning Matters: Indexes

- During index creation or Needs LARGE UTILITY HEAP (**util\_heal\_sz** – DB cfg)
- Reduce sort overflows – increase **sheapthres** DBM cfg
- Separate tablespaces for relational indexes
- If you use ORDER BY's, Group BY's, or Distinct the optimizer might NOT choose and indexed if:
  - Index clustering poor (Syscat.indexes – CLUSTERRATIO and CLUSTERFACTOR)
  - Small tables cheaper to scan
  - Competing index choices

Ref: DB2 Performance Tuning and Troubleshooting guide

## Question???

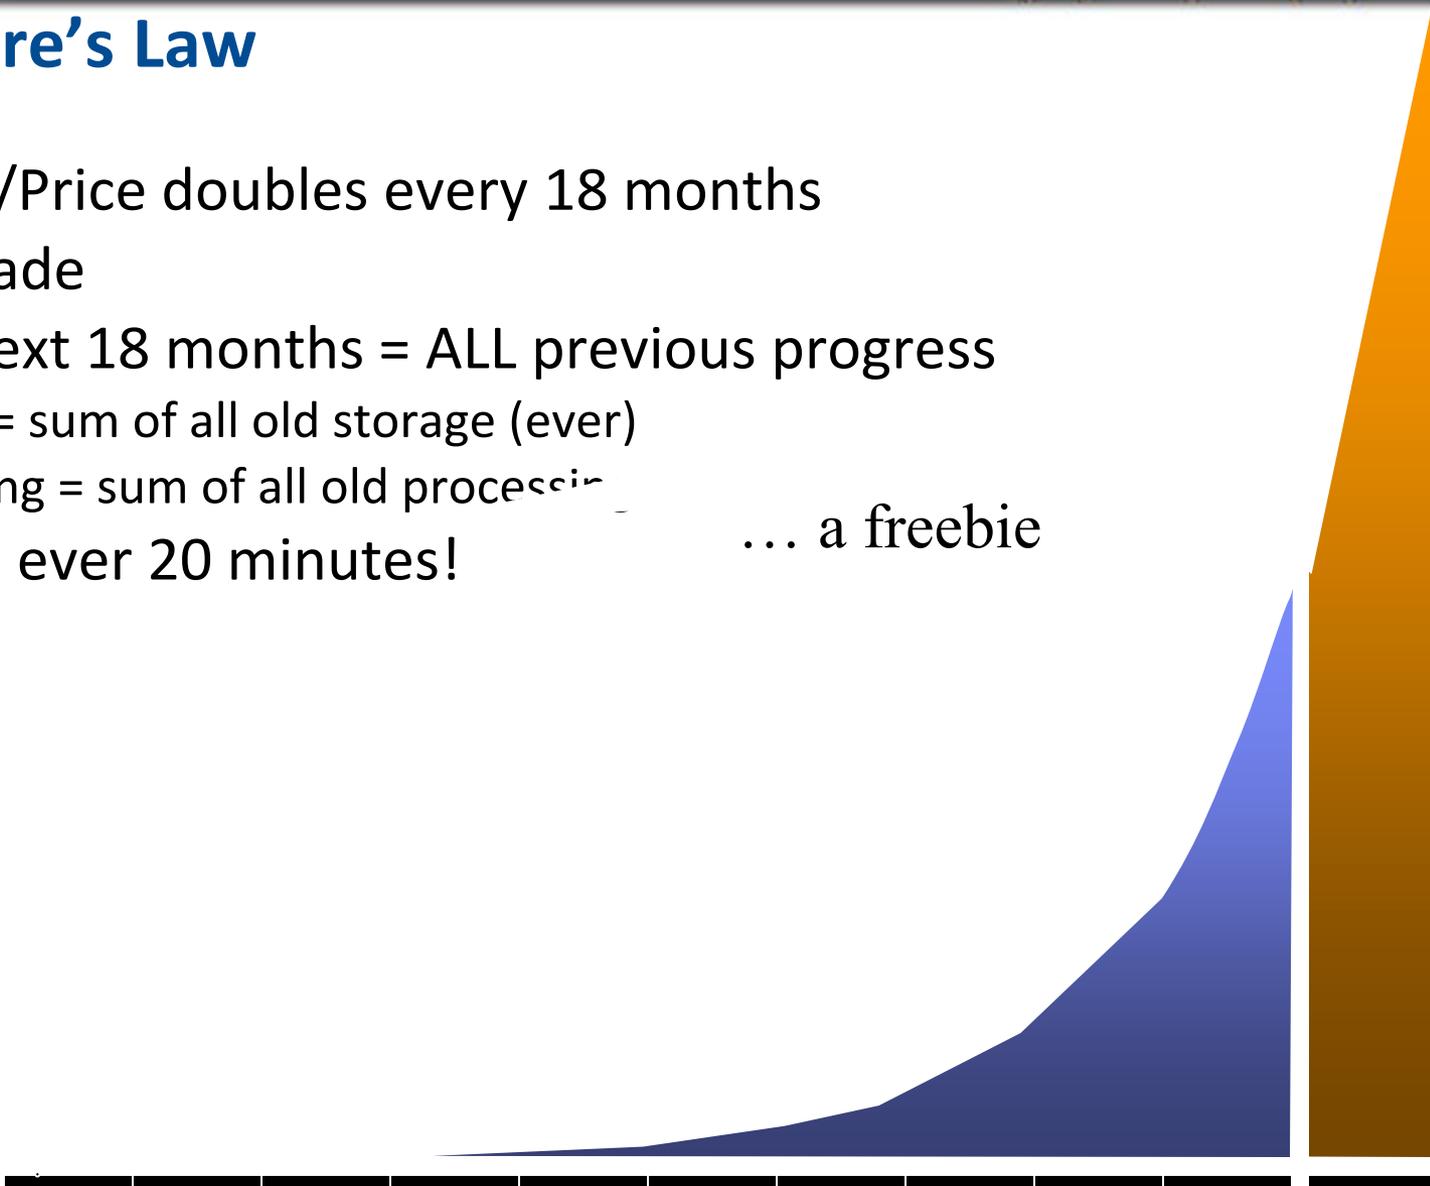
- What year will there be no more spinning disk

Gartner said “2017 70+%”

## Trends: Moore's Law

- Performance/Price doubles every 18 months
- 100x per decade
- Progress in next 18 months = ALL previous progress
  - New storage = sum of all old storage (ever)
  - New processing = sum of all old processing ... a freebie
- E. coli double ever 20 minutes!

15 years ago

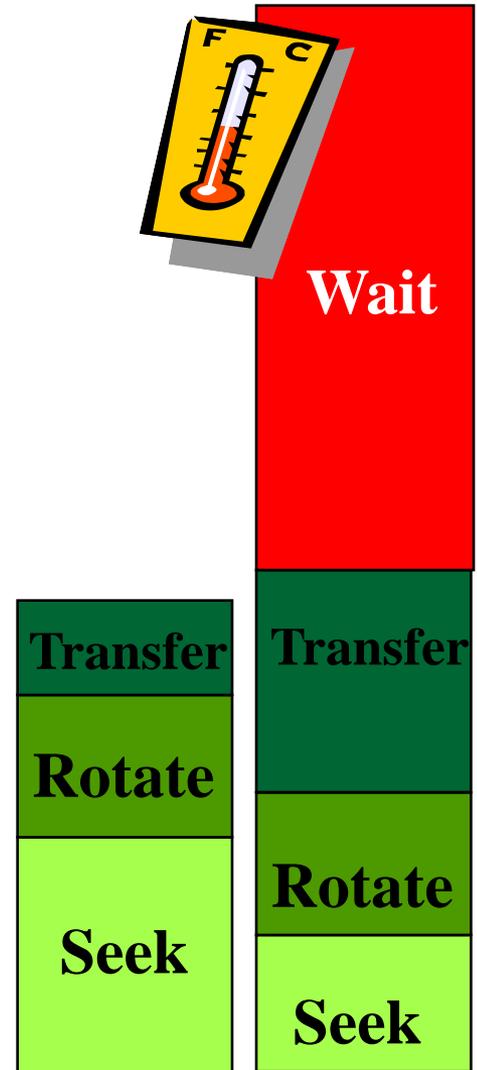


## The IO Access Time Myth

- The Myth: seek or pick time dominates
- The reality: (1) Queuing dominates  
(2) Transfer dominates large block IO  
(3) Disk seeks often short
- Conclusion: many cheap disks better than one/several fast expensive disk
  - shorter queues
  - parallel transfer
  - lower cost/access and cost/byte
- Spindles make a difference!
  - IO requests spread between more disks

**Quiz: What is the CPU Myth?**

Jim Gray & Gordon Bell: VLDB 95 Parallel Database Systems Survey



# Architectural backdrop – Amdahl’s Law

The Constraint Law: *“The slowest component constrains the faster components”*

▪ A simple expression of the maximum potential benefit from parallel processing was defined by Amdahl’s Law. The model separates the tasks that can be performed in parallel from the tasks that need to be serialized. The speedup time on N processing,  $S(N)$ , is defined as the ratio of execution time on a single processor  $T(1)$  to the execution time on N processors  $T(N)$ . Amdahl subdivides the time T into serial processing time  $T_s$  and parallel processing time  $T_p$ , and then expresses the optimal speedup as:

$$S(N) = \frac{T(1)}{T(N)}$$

▪ However, even tasks that are well suited to parallel processing often do not improve in performance with additional CPUs. The major limitation is generally called “overhead”, and includes factors such as context switching, set up time (stack operators, thread or process initialization), bus contention etc. When the overhead time is large compared to the processing time one can expect parallel processing to perform poorly and in some cases even cause performance to get worse. This overhead is often non-linear, for example the cost of bus contention, and context switching can increase in non-linear ways as the system resource becomes saturated. If we model the overhead as  $T_o(N)$ , then the speedup becomes approximately:

$$S(N) = \frac{T(1)}{T_s + T_o(N) + \frac{T_p}{N}}$$

Gene Myron Amdahl is best known for his work in mainframe computers at IBM in the 1950’s and 1960’s. He founded Amdahl Corporation in 1970. Amdahl Corp manufactured "plug-compatible" mainframes, shipping its first machine in 1975. The 1975 Amdahl 470 V6, was a less expensive, more reliable and faster computer than IBM’s System 370/165. By 1979 Amdahl had over 6000 employees and had sold over 1 billion USD in mainframe equipment. Gene left Amdahl in 1979 and went on to found and consult on numerous other technology firms.

## I/O Subsystem Matters

- 400 MB, 1GB, 2 GB, 4 GB, 9.1 GB, 18 GB, 36 GB, 72 GB, 146 GB, 300GB, 750GB, 1TB, 2TB, 3 TB, 4TB ... and **now New SSD / Hybrid drives!** 590,000 IOPS Arrays!
  - **Yikes** where will it stop?
- **RATE OF CHANGE** increasing
  - *“In the next 2.x years, we will create more information, than has been created since the inception of mankind”*
  - Increase in **VOLUME, VARIETY, VELOCITY, VULNERABILITY** of data being stored
  - Increasing # of users & complexity of queries
  - Implementation of new applications and retiring old ones
  - Individual disk sizes increasing
  - # of spindles per tablespace or per processor decreasing. Go look @ TPC-H!
  - CACHE marketing ‘promises’ but not delivering!
- **Storage Subsystems**
  - Debug harder
  - Multi-tier and moving hot blocks make it virtually impossible to tune!
  - SAN and Mixed Workloads (DSS and OLTP) in same system
  - Access pattern will be ‘all over the board
  - On Demand will blur planning
  - CLOUD, Grid, and Utility model could dramatically complicate pattern
  - RAID is the norm – More going to more sophisticated RAID Styles 6 & \*
    - Physical disk is invisible to DB2
    - Stripped Container and Parallel\_IO are helpful but need variable knob to drive subsystem harder
- **Tools and understanding lacking**
  - BEZNext has modeling capability
  - dbi Software – “5 Clicks to Success”
  - Storage, DB, System folks have some tools – Need TOTAL SYSTEM VIEW, NEED Business input, SPC, by LOB
  - Operating system is ‘corn-fused’... reporting at what it thinks is the hdisk (vpaths, dynamic vpath load balancing, etc.)
  - Debug is to say the least... extremely hard through all of the layers!

**“Disk planning on sizes and layout is a pervasive issue”  
IBM CAC - Petabyte Council  
IBM Technical Competency Team  
AND anyone else that has done this before !**

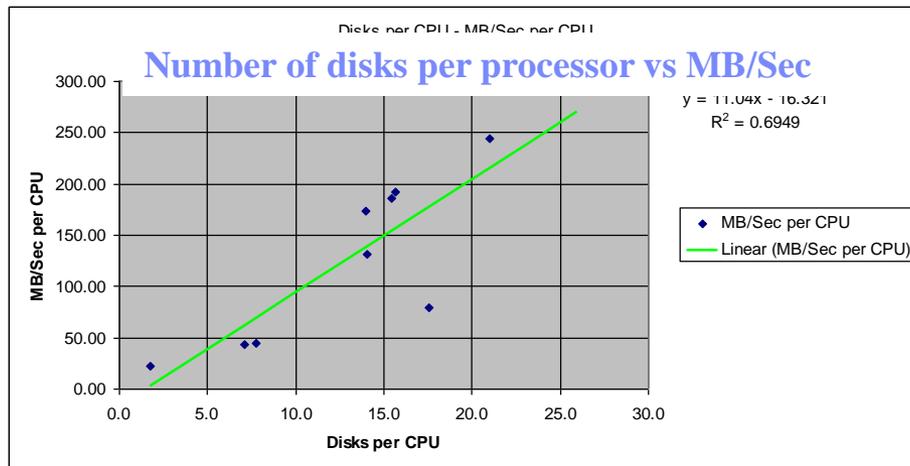
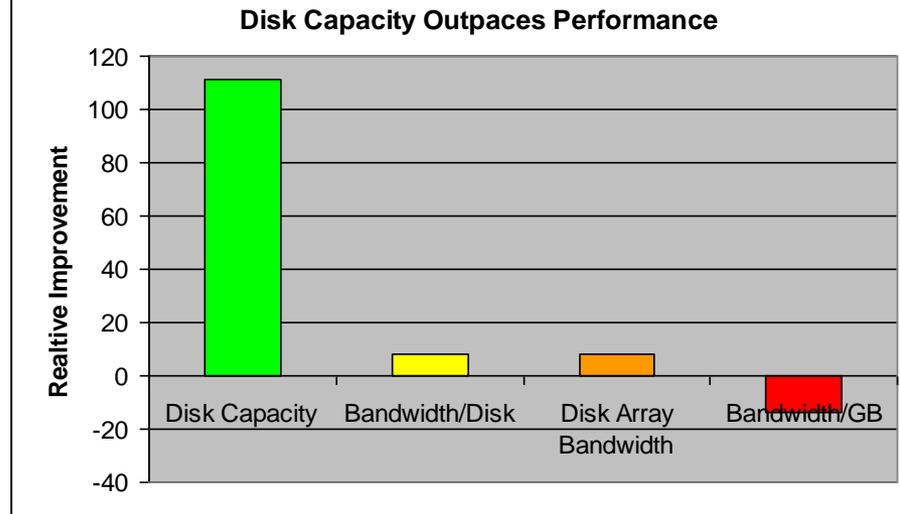
**#7**



# Hardware Technology Trends & Directions

- Ever more powerful CPUs
  - ▶ Moore's law still applies
    - 2x number transistors every 18 months
  - ▶ Similar improvements in Memory
  - ▶ Larger SMP servers (up to 128 CPUs)
    - Largest pSeries SMP capable of ~13TB raw data
  
- Larger drive storage capacities
  - ▶ 9 GB → 18 GB → 36 GB → 73 GB → 146 GB → 300 GB → 1TB
  
- Performance per GB of storage remains static
  
- Enterprise Storage architectures
  - ▶ Monolithic storage vs Modular storage
  - ▶ Enterprise Managed (SAN)
  - ▶ Storage consolidation
  
- Resource/Skill Silos
  - ▶ Storage Management
  - ▶ Systems Administrators
  - ▶ Database Administrators
  - ▶ Application Developers

## 6+Yr Relative Disk Performance Improvements



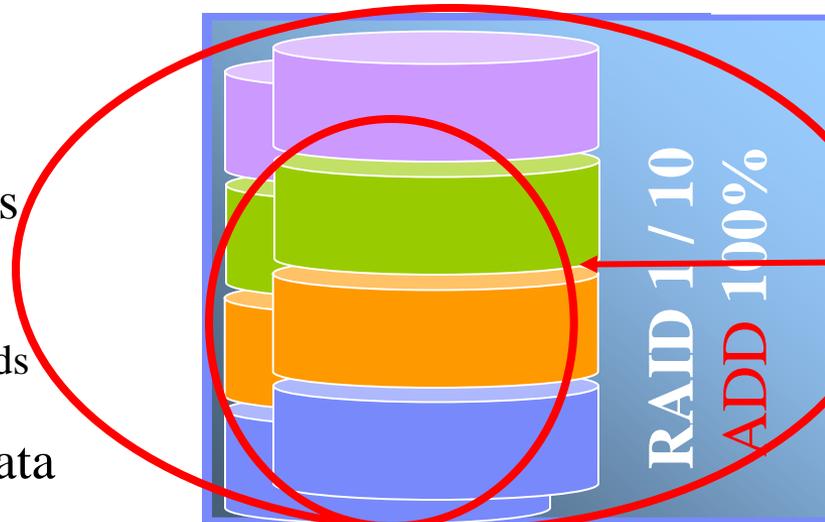
## Quiz: What does your disk cost you?

## Quiz: What does your CPU cost you?

# Balancing out Disk Requirement

## Step 1: Raw Data to Useable Disk

Temp space  
Performance Structures  
Aggregates, MQTs & Indexes  
Working Space  
Staging, Logs, storage overheads  
Raw data



Number of Rows \* Average Row Length

- **Total disk space before RAID = 4  
x Raw**
- **What is your RAW to Disk Req?**
- **What is your Active RAW Data?**

## Quiz: What's your Disk RATIO number?

- If you have one row or 1 MB of data to load, how many rows and MB does that mean for your data base growth?

## Saving I/O on ETL – Best Practices

-  • SKILL CHECK!?
  - Understand you full stack
  - Be able to understand the terms and acronyms
-  • Aggregation Metadata?
-  • Partitioning?
  - Partition as early in the process as possible
-  • To SORT or NOT to SORT, that is the Question?

# What's your ETL number (ratio)?

## Requirements need to be well understood

- No Redos - Nothing burns the time line like redoing something
  - Re-layout the disks due to bottlenecks
  - Change Data Integration processes from daily to part day to Near Real Time
  - Scope creep!
    - New application
      - Disk space is off by 30 - 50%
      - Adding disk and reorganization is almost never optimal



... UP FRONT! ... REF: 15 WAYS TO RUN OUT OF DISK

#2

## A Full Stack Perspective Matters”

- I/O
- Balance
- Suggested Building Block
- Double or triple striped?
- CPU
  - Reduce Processor Stall – Cache miss- BLU

## 15 Ways to Run out of disk! ... besides data envy!

- New applications
- Unexpected growth within an application
- MANY copies of tables
- Multiple copies
  - Entire DB, tablespace, table, HA, DR
  - Non-Prod Test, QA, Perf
- Requirements not well understood - REDO
- Delays / Debug copies
- ETL Ratio creep
- SAS extracts
- User space never cleaned up or regulated
- RAID1, 5, 6, 8, 10
- WAY to many indexes
- MDC on wrong column
- New Versions, Releases, Fixpacks
- Someone forgot to CLEANUP / prune – dumps, logs, old versions, tar files
- UNICODE

## Governance, sponsorship & funding

- Governance
- Sponsorship
  - “If the lead dog is not leading, the sled will not stay on track”
  - Business Intelligence has proven itself
- Funding
  - \$Green is not a 4 letter word and should not be a one time event!
    - Don't 'bean count' / 'pinch pennies' good business decisions!
  - Equipping: Nothing derails a project faster than not having the proper resources
  - Consulting
  - Equipment
  - Time Line
  - EXPECTATION!

## 75 Ways to Save I/O – System Level

- Precept – Do less work and things will run faster!
- Overallocation = paging = "performance goes into the toilet"
- Concurrent I/O and VMTUNE  
-P -p
- Network: MTU size, send / rec sizes
- Block sizes to large - DPF
- Solid State and Ramdisk

Do only **1 level of striping**-->  
DB2 ONLY

File Opens

Match sizes of I/O from:

- App->DB->OS->dsk

Benchmarking and caching

Virtualization **"OVER-HEAD"**

Other Benefits:

- Kernel Resources, paging, I/O & mem buss
- Less reboot time

## Saving CPU at the System Level \* Freebie

- Shift batch to batch window! Free cycles Vs dribbling CPU on floor
- Don't pin processes to processor
- Smaller LPARs means smaller memory footprint to sort
- Codepage Conversion / comparison

## 75 Ways to Save I/O – DB Level

- Blue Triangle Demo
  - Locality of data (more good data and less bad data)
- Card trick - Partition Elimination and N squared effect
- “Sort is a 4 letter word”
  - Sort or order on ETL?
  - RCT – Range Cluster Tables
  - ITC – Insert Time Cluster
  - Index (forward and reverse, partition key)
  - Sort overflows
    - Ratio of Rows/GB for total memory
    - Sort is non-linear
  - Heaps – careful and error message are confusing

## 75 Ways to Save I/O – DB Level (cont.)

- Indexes
  - Include Columns
  - RCT
  - MDC
  - ITC (Insert Time Cluster)
  - MQT
  - Forward and Backward
  - Don't OVER INDEX! -  
impact on ETL and  
deletes
  - Don't use RI → ETL?
- Clustered
  - % free (table and index)
  - Block (Did I mention  
MDCs?)
  - Index and Temp  
Compression (V9.7)

## 75 Ways to Save I/O – DB Level (cont.)

- Leg Elimination
  - Union all views
  - Table Partition / roll in roll out
  - Hash
- WLM (WorkLoad Mgmt)
- Buffered Inserts
- For Fetch Only
- Fetch first 100 rows
- More Memory or More
- ETL Related
  - LOAD vs Insert
  - HPU vs export
  - DataStage deep parallelism
  - Commit Frequency
  - Shift batch to batch window!
  - LDA vs Marts
  - ELT vs ETL
  - Append
  - Compatible processes

**DON'T DIM THE LIGHTS by extracting the whole table to app (DataStage, SAS, ETL, app, mining, etc.)!**

## 75 Ways to Save I/O – DB Level (cont.)

- SQL Tuning
  - Bad queries – cartesian joins
  - Generated
  - 'perceived penalty / cost'
  - “120 page explain”
  - SAVE OFF ACCESS PLANS
    - BEFORE THE UPGRADE!
  - DB Design
  - Optimizer Level / profiles / statistical views
  - Static vs Dynamic
- RUN RUNSTATS!
- Correlation
- Distribution Statistics
- Locking (UR)
- Access Path
  - Optimizer levels
  - Replicated tables
- CUBE
- ROLLUP
- Grouping Sets

## 75 Ways to Save I/O – DB Level (cont.)

- Cache
  - Disk
  - Disk Subsystem
  - CIO / DIO vs File System
  - Buffer Pool
  - Big block reads
  - BE CAREFUL with  
AUTOMATIC  
SETTINGS...
  - TEST, TEST TEST
  - BE Careful of defaults

## Saving CPU – DB Level \* Freebie

- MQT – CPU & I/O
- MDC
- ITC (Insert Time Cluster)
- WLM

## Homework Matters

- ISAS = IBM SMART ANALYTICS SYS (DPF) – BALANCE
- zLinux
- Pricing and packaging
- Compression value
- Be a student of tips and best practices
- Write about what you learned (Amber Crooks, Scott Hayes)
- Layered Data Architecture
- Temporal / Time Travel??? ITC = Insert Time Cluster

## Homework Matters – Cont.

How is memory consumption regulated?

WLM?

How do you tell what **application** is using what disk?

Prediction vs Trending?

Reactive vs proactive = SPM= Strategic Performance Management



“A prescriptive” approach



## Planning Matters

- SPM – Strategic Performance Management
  - Be Proactive, Continual, Methodic
  - EXPECTATIONS
- Meta Group article on SPM
- What is the value of accurate sizings?
  - A web based company in Florida was open for 2 days before it's system imploded, and so did the company!
  - After a hurricane a large insurance company got 10X the # of transactions and, the Data Architect said“...their system fell over”

**DO the Homework**

## What is this Application's Workload / Profile?

- Per Application
  - # Users?
  - Concurrency?
  - IO/Sec requirement?
  - CPU consumption?
  - Growth?
    - Existing apps?
    - New apps?
    - New users?
      - GEO?

## 75 Ways to Save I/O - DB

- Blue Triangle Demo
  - Locality of Data
  - Leg Elimination
  - Drive all the resources
    - $N^2$  effect

**# 1 ISSUE:** I/O is undersized #  
of spindles per processor (I/O  
per sec req.)

Sort overflow

RAW data ratio to:

Processor

Partition

GB of Mem

**# 2 AFTER** application design is  
access plan = Optimizer

Opt settings

**#3 SQL generated?**

## LAW: There will be speed bumps

- You have to plan for problems... rest assured, they will come
- Some known things:
  - Fixpacks!
  - Version Upgrades!
  - Data volumes increases!!
  - More data sources!!!!
  - Near Real Time pressures changing ETL processes!
  - Delays or part of the team not meeting schedule
- There are many others and unknown 'rocks'

*Inventive thought takes time!*



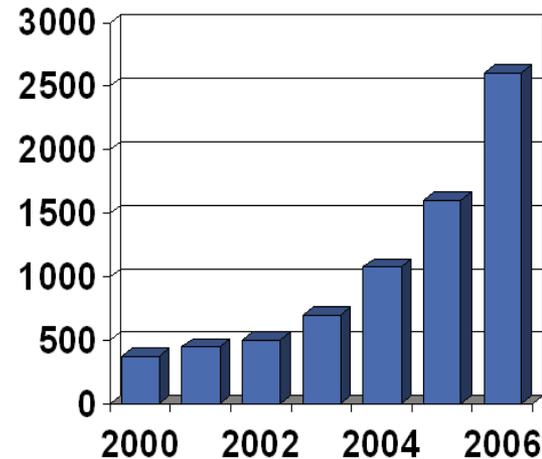
15 Ways to Run out of Disk

#10

# Why didn't more people use compression or near line storage? ... Now they are!

- Why don't people use compression at the file system level?
- Why not near line storage?
  - Conclusion: Disk prices are falling so no one cares... Do not get sucked into this.
- Survey???
  - If you could compress large amounts of data by a significant amount (trading off CPU and time for the convenience), would you use it?

Data Overload:  
Petabytes of Disk shipments



Source: Fortune Magazine:  
3/17/03 Storage shipments up 57% in past 3 yrs and is expected to rise 430% by 2006

 Don't get 'Data Envy' You need several years of history...fine... but no one deletes!

## Recipe for a Successful TB Class System

- **Find a business opportunity**
- Supply a solution by: finding the 'Pans' - (HW), and the following 'Ingredients' - (SW and people)
- Starting with 2 metric Tons of brain cycles
  - Apply business application liberally
- **Develop a plan that includes:**
  - A Sound vision from the leadership to:
    - Define the quantity of 'servings' you will need for clients
    - Define what it will take in infrastructure costs
    - Defined what it will take in integration of the idea to industrialization and
    - Creating a quickly scaleable, web enabled and a gargantuan warehouse!
- **Defined Key success factors:**
- Start with 6 pounds of elbow grease from the right skill, at the right time. Tempered with world class consultants that have 'Scars of several successes under their belts are key'
- **Baking Instructions:**
  - Establish a good name brand
  - Build the capital to acquire market share
  - Hire the best!
  - Get infrastructure and staff on board as quickly as possible.
  - Schedule the additional skilled folks you need.



### Baking Instructions (cont.)

Stir in the ideas from the 'brain cycles' slowly while mixing in talented people that have done each component before, and avoid 'lumps', by mixing well.

Add 'Sugar'(capital) as needed.

Bake at an reasonable (not aggressive) time line with scope changing as the business dictates

In parallel, Shake vigorously, the need for clients, advertising, Sales, inventory control, market basket analysis, cost of goods sold analysis, supply chain analysis, hiring, firing, growing, investing, divesting, and acquiring.

Set the production date at an aggressive time, while balancing the dollars available.

Enjoy!

## Summary:

- Need = Not just a checklist but a systematic methodology for profiling your system (even by application)
  - Prioritizing
  - Then ATTACKING the problem
- Research and Science projects
- How do I save I/O – 75 Ways Discussion
  - System, Storage, Middleware (incl. DB)
- How do I save CPU – 25 Ways Discussion
  - System, Storage, Middleware (incl. DB)
- More to come by version / release – Jan DB2Night Show™

## Questions????



Questions????

Don't forget the DB2Night Show™ downloads

# 118



**Taming the Data Tsunami:  
17 Laws of building Petabyte Level DB Systems**

**Evaluations**



## Great Performance URLs:

- [www.db2nightshow.com](http://www.db2nightshow.com) Free DB2 Education from DBI
- <http://www.dbisoftware.com/brother-eagle/advice/db2dbiref.php>
- <http://www.dbisoftware.com/brother-eagle/advice/db2dbssdwt.php>
- <http://db2commerce.com/> Ember Crooks
- Manual City:
  - <http://www-01.ibm.com/support/docview.wss?uid=swg27009474>
  - By version, .PDFs, best practices
- <http://www.ibm.com/developerworks/data/library/techarticle/dm-0404mcarthur/>
- <http://www.mydeveloperconnection.com/pdf/DB2-SQL-Tuning.pdf>
- <http://www.ibm.com/developerworks/data/library/techarticle/anshum/0107anshum.html>
- <http://performancewiki.com/db2-tuning.html>
- <http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=%2Fcom.ibm.db2.udb.admin.doc%2Fdoc%2Fc0007586.htm>
- Redbooks:
  - <http://www.redbooks.ibm.com/abstracts/sg246012.html>
  - <http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246949.html>
- <http://www.db2dean.com/Previous/PerformTools.html>

## 10.5 Performance Related Publications!

- [DB2PerfTuneTroubleshooting-db2d3e1050.pdf](#)
- [DB2AdminRoutinesViews-db2are1050.pdf](#)
- [DB2-SQL-Tuning.pdf](#)
- [DB2Monitoring-db2f0e1050.pdf](#)
- [DB2Partitioning-db2pce1050.pdf](#)

# Lee Goddard

dbi Software

*Lee.Goddard@dbiSoftware.com*

75 Ways to Save I/O and 25 Ways to Save CPU

