Self-Managing Technology in Database Management Systems

Surajit Chaudhuri, Microsoft Research
<surajitc@microsoft.com>
Benoit Dageville, Oracle
<benoit.dageville@oracle.com>
Guy Lohman, IBM Almaden Research Center
<lohman@almaden.ibm.com>

Agenda

- **Motivation**
- Aspects of Self-Managing DBMSs
- Architectural trade-offs
- Research Issues
- Self-Managing Features in Current Products
  - Microsoft SQL Server
  - Oracle
  - IBM DB2

Once Upon A Time…

- **Computers were**
  - **Large**
  - **Unreliable**
  - **Expensive (Millions of $$$)**
- **People were cheap (comparatively)**

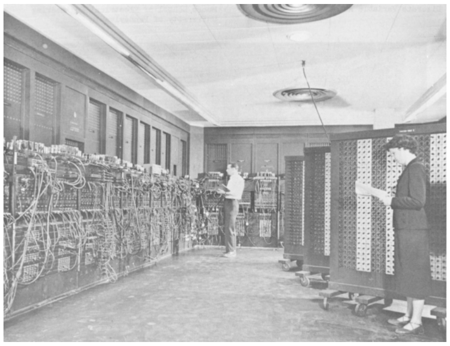

Technology Marches On!

- Chips (CPUs and memory) got
  - Smaller
  - Faster
  - More reliable
  - Cheaper (tens of $)!

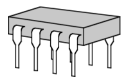

- Storage got
  - Bigger
  - Faster
  - More reliable
  - Cheaper (fractions of pennies)!

- Communication got
  - Faster
  - More reliable
  - Cheaper!
- Systems got a lot more complex!
- People got
  - Not appreciably faster or more reliable
  - (a lot) More expensive!

$1000 Buys…

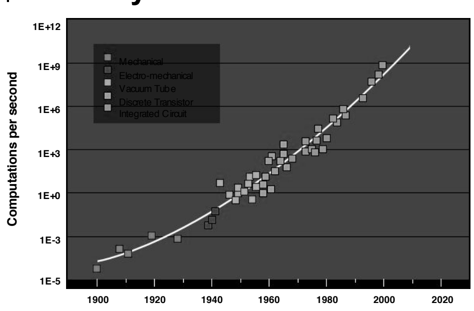

Reference:
"Mind Children: The Future of Robot and Human Intelligence," Hans Moravec, Harvard University Press, 1988,
"The Age of Spritual Machines," Ray Kirzweil, Viking, 1999.

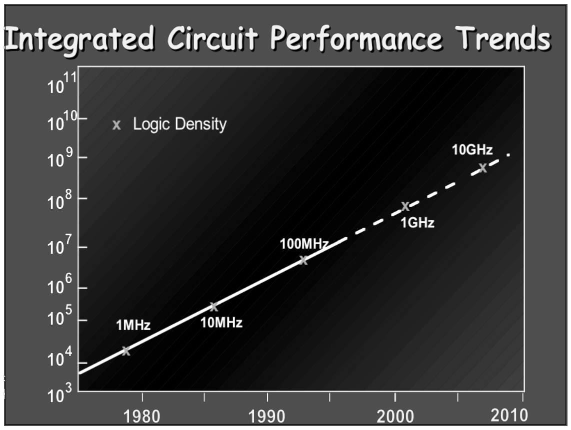

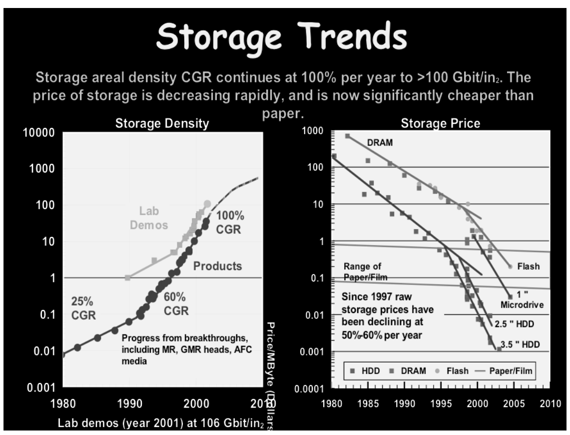

Today

- Disks on laptops have more capacity than most need
  - 1 Terabyte for $1199: http://www.lacie.com/products/product.htm?id=10118
- CPUs cost less than a good meal
  - Complete "bare bones" machines for $200 (retail)
  - Example: http://shop1.outpost.com/product/3847537
- Network capacity glut permits streaming voice and video
- But people, …

Cost of Labor Rarely Decreases (Despite Outsourcing)

Employment Cost Index (1989 = 100):

Total comp., Professional, Specialty, & Technical Occupations (all civilian)

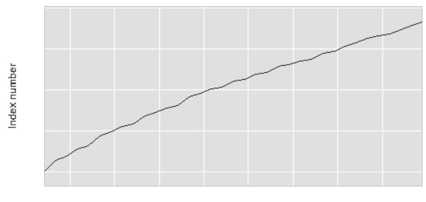

- **Source:** U.S. Bureau of Labor Statistics (http://data.bls.gov/servlet/SurveyOutputServlet)
  Series Id: ECU11121I

The High Cost of I/T Management

For example: the cost to manage storage is typically <u>twice</u> the cost of the actual storage system.

Storage: What $3 million bought in 1984 and 2000.

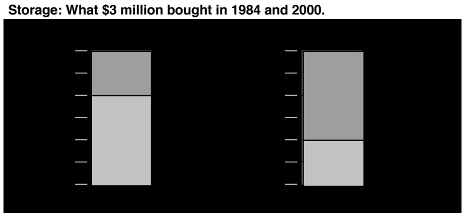

(1) J. P. Gelb, "System-managed storage," IBM Systems Journal, Vol 28, No. 1, 1989 pp. 77-103.
(2) "Storage on Tap: Understanding the Business Value of Storage Service Providers", ITCentrix report, March 2001.
(3) "Server Storage and RAID Worldwide" (SRRD-WW-MS-9901), Gartner Group/Dataquest report, May 1999.

Human Costs Dominate in Database, Too

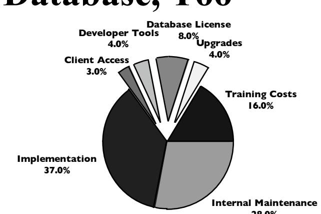

81% is "People Cost"

Source: The AberdeenGroup, 1998
http://relay.bvk.co.yu/progress/aberdeen/aberdeen.htm

Houston, we have a problem …
Complex heterogeneous infrastructures are the norm!

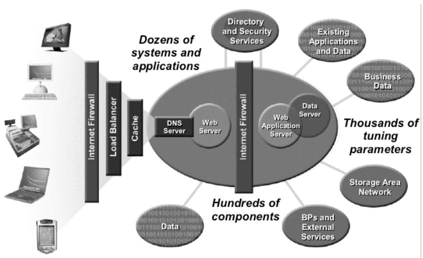

Making the Front Page

eBay
Outage: 22 hours 12 June 1999
Operating System Failure
Cost: $3 million to $5 million revenue hit and 26% decline in stock price

E*Trade
3 February 1999 through 3 March 1999: Four outages of at least five hours
System Upgrades
Cost: ????
22 percent stock price hit on 5 February 1999

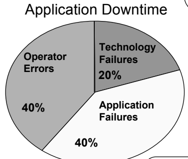

AT&T
13 April 1998 outage: Six to 26 hours
Software Upgrade
Cost: $40 million in rebates
Forced to file SLAs with the FCC (frame relay)

Dev. Bank of Singapore
1 July 1999 to August 1999:
Processing Errors
Incorrect debiting of POS due to a system overload
Cost: Embarrassment/loss of integrity; interest charges

America Online
6 August 1996 outage: 24 hours
Maintenance/Human Error
Cost: $3 million in rebates
Investment: ???

Charles Schwab & Co.
24 February 1999 through 21 April 1999:
Four outages of at least four hours
Upgrades/Operator Errors
Cost: ???; Announced that it had made $70 million in new infrastructure investment.

Source: Gartner Group

Does this look familiar?

"If it's so efficient, why doesn't it fix itself!"

Reducing the TCO

- Management costs a major part of total IT spending
  - Cost of HW decreasing while cost of managing systems is increasing
  - IT System form core of business today
    - Customers and suppliers deal directly with IT systems over the web
    - Reliable IT Infrastructure is critical to success
    - ☞ IT Performance = Business Performance
  - Increased reliance on IT and explosion in data volume require more administrative staff
  - Limited availability of skilled labor results in spiraling DBA salary
- Increased business competitiveness requires reduction in operating expenses
  - ☞ IT Managers being asked to do more with less $$

Managing Increasing Complexity

- Increase in Complexity & Size of Applications
  - Database workloads are more mixed (e.g. OLTP and complex reporting).
  - Database workloads are more dynamic.
  - Data size is growing rapidly
    - ☞ Multi-terabytes are no longer the exception!
- DBMS vendors have responded to these challenges by
  - Enlarging the scope of existing features
    - New access structures, complex optimizations
    - Complex hardware architectures like clusters or MPPs
  - Adding new features in the server
    - Objects, XML, OLAP, data mining, ETL
    - Replication, high-availability, …

☞ Managing/tuning a modern database system requires a very high degree of expertise!

Agenda

- Motivation
- **Aspects of Self-Managing DBMSs**
- Architectural trade-offs
- Research Issues
- Self-Managing Features in Current Products
  - Microsoft SQL Server
  - Oracle
  - IBM DB2

Autonomic Computing

The Idea

Wouldn't it be great if your Database (and entire system!) were as easy to maintain and as self-controlled as your refrigerator?

What Is The Self-Managing Vision?

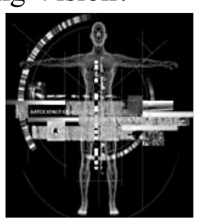

- **“Intelligent” open systems that…**
  - § **Manage complexity**
  - § **“Know” themselves**
  - § **Continuously tune themselves**
  - § **Adapt to unpredictable conditions**
  - § **Prevent and recover from failures**
  - § **Provide a safe environment**

Frees your business to focus on business, not infrastructure

Huge Scope of DBA Responsibilities

- Initial Design & Layout
  - Hardware configuration
  - Logical database design
  - Physical data layout (partitioning, allocation to nodegroups, clustering)
  - Auxiliary data structures (indexes, view materializations)
  - Configuration parameters (hundreds!)
  - Security policies, groups, userids

- Dynamic Monitoring & Adjustment
  - Database statistics to collect and when
  - Clustering and Reorganization
  - Memory allocation, esp. buffer pool sizes
  - System / query status
  - Problem determination (deadlocks, bad plans, ...)
  - Visualization of all the above

Where DBA's spend their time

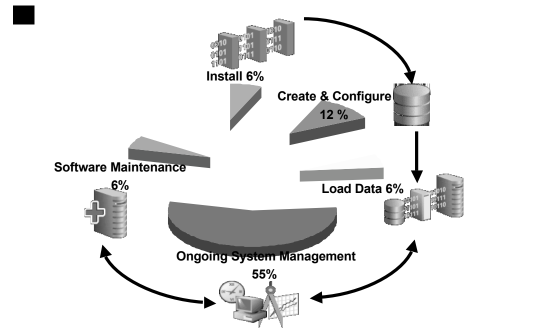

Source: International Oracle User Group (IOUG) 2001 DBA Survey

Ongoing System Management

55% of DBA's time is spent in ongoing management, monitoring and tuning

- Performance Diagnosis & Troubleshooting
- SQL & Application Tuning
- System Resource Tuning
- Space & Object Management
- Backup

Source: International Oracle User Group (IOUG) 2001 DBA Survey

Manageability Challenges - Today

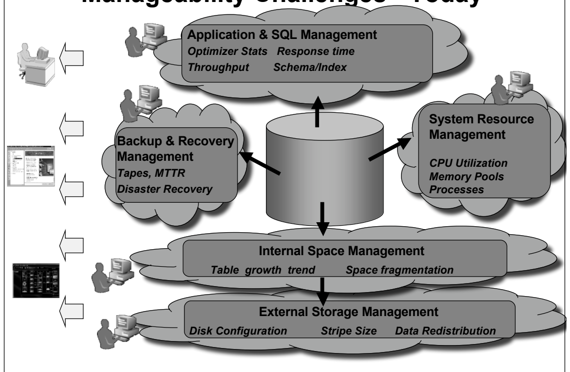

Core Capabilities for Enabling Self-Managing Systems

- Problem Determination
- Common System Administration
- Adaptive Monitoring
- Solution Install
- Policy-based Management

- Complex Analysis
- Heterogeneous Workload Management

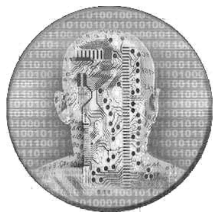

A Self-Managing Taxonomy (Self-CHOP)

Increase Responsiveness

Adapt to dynamically changing environments

Business Resiliency

Discover, diagnose, and act to prevent disruptions

Operational Efficiency

Tune resources and balance workloads to maximize use of IT resources

Secure Information and Resources

Anticipate, detect, identify, and protect against attacks

Self-Managing Deployment Model

Evolution, not revolution

	Basic Level 1	Managed Level 2	Predictive Level 3	Adaptive Level 4	Autonomic Level 5
Characteristics	Multiple sources of system generated data	Consolidation of data and actions through management tools	System monitors, correlates and recommends actions	System monitors, correlates and takes action	Integrated components dynamically managed by business rules/policies
Skills	Requires *extensive, highly skilled* IT staff	IT staff *analyzes and takes actions*	IT staff *approves and initiates actions*	IT staff *manages performance* against SLAs	IT staff *focuses* on enabling business needs
Benefits	Basic Requirements Met	Greater system awareness Improved productivity	Reduced dependency on deep skills Faster/better decision making	Balanced human/system interaction IT agility and resiliency	Business policy drives IT management Business agility and resiliency

Manual → Autonomic

Agenda

- Motivation
- Aspects of Self-Managing DBMSs
- **Architectural trade-offs**
- Research Issues
- Self-Managing Features in Current Products
  - Microsoft SQL Server
  - Oracle
  - IBM DB2

Core Building Blocks for an open architecture

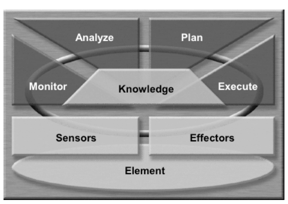

- **An autonomic element contains a continuous control loop that monitors activities and takes actions to adjust the system to meet business objectives**
- **Autonomic elements learn from past experience to build action plans**
- **Managed elements need to be instrumented consistently**

Architectural Trade-Offs (1 of 2)

- What granularity for such "autonomic elements"?
  - Per database?
  - Per CPU?
  - Per component (e.g., DBMS, App Server, …)?
  - Per complete system?
- Distributed?
  - \+ Local control
  - \+ Simpler
  - \+ Scalable
  - \- Don't have the "big picture"
  - \- Unstable "Tug of war" with other components possible
- Centralized?
  - \+ Have broader view of cause & impact
  - \- Won't scale well
  - \- Relies on communication speed, availability, & standards

Architectural Trade-Offs (2 of 2)

- Hybrid (hierarchical)?
  - Blend: both distributed & centralized control elements
  - Communicate only necessary info to
    - Other components
    - Central controller

  + Have broad view as well as local control
  + Scalable
  - Relies on communication speed, availability, & standards
  - Complex interactions between controllers
  - Can still have unstable conflicts

Multiple Contexts for Self-Managing Behavior

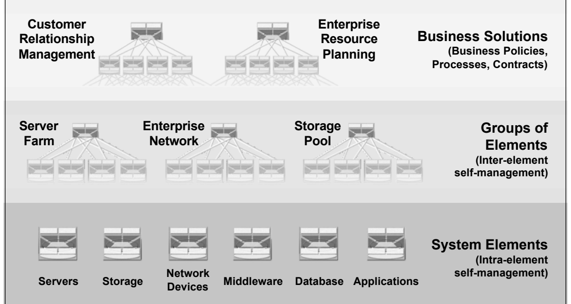

Agenda

- Motivation
- Aspects of Self-Managing DBMSs
- Architectural trade-offs
- **Research Issues**
- Self-Managing Features in Current Products
  – Microsoft SQL Server
  – Oracle
  – IBM DB2

Research Topics / Issues (1 of 2)

- Capacity planning (modeling & estimation)
  - How model systems with limited specification?
  - How maintain model with evolving HW & SW?
- Installation
  - Dependency graph of prerequisite versions, configurations
- Database Design
  - Logical Design (application design, normalization)
  - Physical Design – how to decide:
    - Selection of indexes, materialized views, etc.
    - Data placement (clustering, partitioning, etc.)
    - Dynamic storage provisioning
- Performance tuning
  - How automate determination of poor performers?
  - How dynamically re-configure system in response to load changes?
- Maintenance – when / how to perform
  - Backups?
  - Reorganizations?
  - Statistics collection?
  - Upgrades?

Research Topics / Issues (2 of 2)

- Self-Healing
  - How much monitoring data to collect?
  - How do you know if your system is "firing on all cylinders"?
  - How do you isolate problems from noise of diagnostics?
  - How do you correlate logs from components on different machines w/ diff. clocks?
  - How do you isolate root cause from cascading error messages?
  - Fuzzy searching of symptom databases
  - How do you automatically generate diagnostics to resolve ambiguous problems?
  - How do you model and determine the cause & repair for problems never before seen?
  - How do you determine the best fix for a problem, even if the cause is known?
  - How do you build repair rules automatically from past successes & failures?
- System Control
  - Scheduling & prioritization of tasks
  - How do you resolve conflicting rules & priorities?
  - How do you make progress on maintenance without impacting production workload?
  - How do you avoid instability and "thrashing" (control theory)?
  - How much monitoring is enough to resolve problems but not impact production?
  - How do you learn from past successes & failures?

Conclusions

- Systems management dominates Total Cost of Ownership (TCO)
  - HW & SW costs decreasing
  - DBA (and other people) costs increasing
  - Complexity and size of systems increasing
- Only solution is Self-Managing DBMSs!
- Some self-managing features in existing products (remainder of this tutorial)
- Many challenging research issues remain!

Agenda

- Motivation
- Aspects of Self-Managing DBMSs
- Architectural trade-offs
- Research Issues
- **Self-Managing Features in Current Products**
  - Microsoft SQL Server
  - Oracle
  - IBM DB2