



“THINK” About Data 2008

Data Lifecycle Management

Driving and Optimizing Storage

Stewart Bazneh
Tivoli Storage Sales Leader


IBM Tivoli Software

© 2008 IBM Corporation

Green issues & energy usage

- IT is in the spotlight


- Datacentres account for approx 4% of UK corporate electricity usage
- The UK’s Climate change bill will legally enforce:
 - 26-32% reduction in UK carbon emissions by 2012
 - 60% reduction by 2050
- EU emissions trading scheme will force polluters to buy additional carbon credits whilst “good citizens” can sell credits
- Power has changed from a cheap commodity to a business-limiting factor
- A Double-Whammy for polluters – expensive power + purchase of carbon credits



Green issues & energy usage


- IT is in the spotlight

- Excessive heat and insufficient power are the two most critical problems in data centers today.
 - Michael Bell, "Use Best Practices to Design Data Center Facilities," Gartner
- Energy consumption is currently the number one cost for datacentres and virtualisation could be the key to make the Irish IT industry more energy efficient
 - David McAuley, Sustainable Energy Ireland
- A crisis is looming ... around the most fundamental of requirements, the need for power... the average energy costs of running a UK corporate datacentre is currently £5.3 million per year and is set to increase to £11 million over the next 5 years
 - Broadgroup Consultants



This Information is Driving the Need *For Data Center Transformation*


Data Types



Today

80% unstructured data


Data Value



By 2010...

1000x storage per image

Data Growth



Through 2012...

54% annual storage growth

And by 2020...

Personal data footprints will be 16x current levels

Storage challenges driven by multiple sources




***“Buying more hardware” is no longer the answer;
it now requires a comprehensive information centric approach***




Storage Power Landscape

Components of Data Center Power Consumption




Source: IBM

Storage Power Consumption/GB




Source: IBM

Data Center Storage Usage External PB Shipped



Source: IDC

Data Center Storage Power Growth


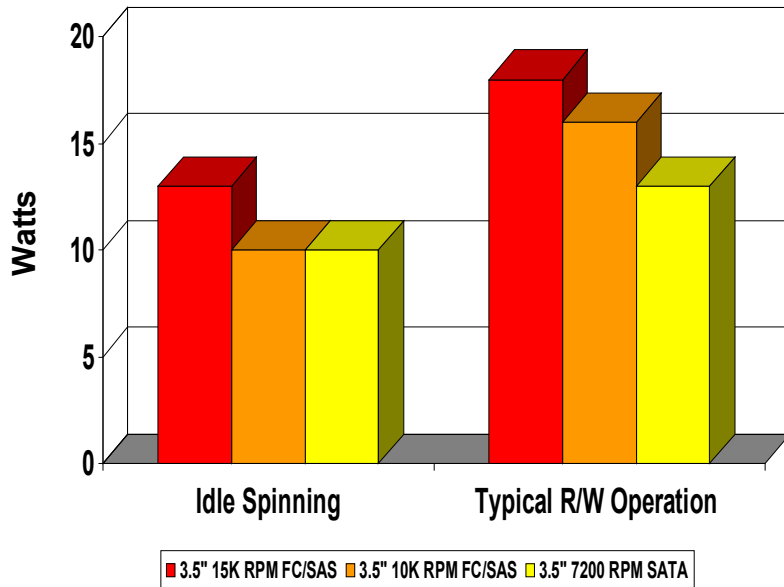


Source: IBM

Deploy More Power Efficient Storage - *Action*

Place the data on the most power efficient (slowest) storage that meets your needs


Drive Power Use



Speed Kills: Best server class drive in Watts/TB is 7,200 RPM 500 GB drive

Source: IBM

Storage – What’s the truth?



Deploy More Power Efficient Storage – *Facts*

Tape Power & Cooling is Dramatically Better Than Disk

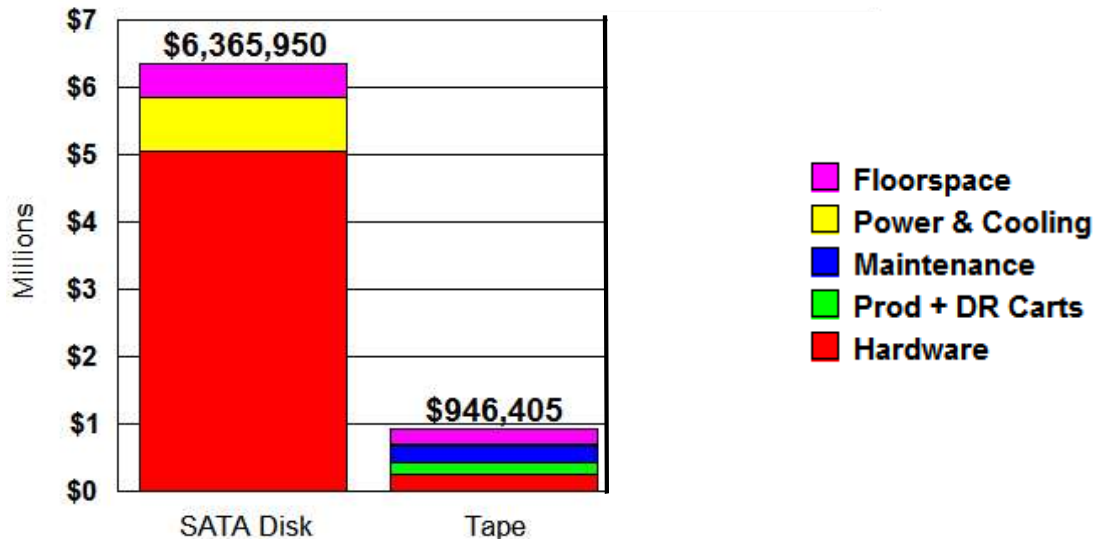
| | SATA Disk | Tape | D:T Ratio |
|-----------------------|--------------------|------------------|-----------|
| Hardware | \$5,067,652 | \$269,771 | 19 |
| Prod + DR Carts | \$0 | \$168,750 | 0 |
| Maintenance | \$0 | \$240,036 | 0 |
| Power & Cooling | \$805,098 | \$40,396 | 20 |
| Floorspace | \$493,200 | \$227,452 | 2 |
| Estimated Cost | \$6,365,950 | \$946,405 | 7 |

← Tape is very “green” @ 20x less energy expense

- Customer Storage Goals:**
- Performance
 - Compliance
 - Data Security
 - Disaster Protection
 - Reduce TCO and energy costs

All disk or all tape may not address all goals

10 Year TCO Analysis



- Scenario:**
- Store 250TB
 - 25% Growth Rate
 - Over 10 Years
 - DS4700 SATA Disk
 - LTO 4 Tape Library

Deploy More Power Efficient Storage – Action

Combine Tape & Disk to Address Goals and Best “Green” Practice

| | SATA Disk | Tape | Blended D+T | D:T Ratio | D:B Ratio |
|-----------------------|--------------------|------------------|--------------------|-----------|-----------|
| Hardware | \$5,067,652 | \$269,771 | \$1,376,693 | 19 | 4 |
| Prod + DR Carts | \$0 | \$168,750 | \$126,250 | 0 | 0 |
| Maintenance | \$0 | \$240,036 | \$208,699 | 0 | 0 |
| Power & Cooling | \$805,098 | \$40,396 | \$211,612 | 20 | 4 |
| Floorspace | \$493,200 | \$227,452 | \$332,092 | 2 | 1 |
| Estimated Cost | \$6,365,950 | \$946,405 | \$2,255,346 | 7 | 3 |

Customer Storage Goals:

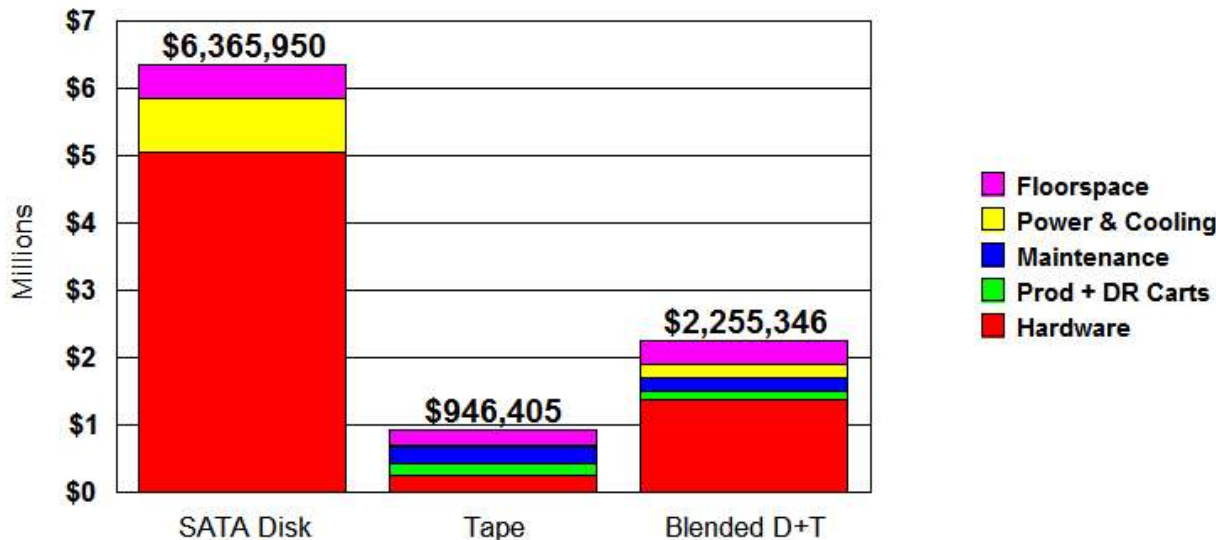
- Performance
- Compliance
- Data Security
- Disaster Protection
- Reduce TCO and energy costs

Blended disk and tape can address the goals

IBM Blended Disk and Tape Products

- TS7500 Virtualization Engine for Linux, Unix and Windows systems
- TS7700 Virtualization Engine for System z
- DR550 and WORM Tape
- Tivoli Storage Manager

10 Year TCO Analysis



Storage Management Overview

Optimizing storage infrastructure for business continuity and compliance


- Visibility:** View storage capacity, utilization, performance and resource configuration to optimize use and availability
- Control:** Control storage management costs by more effectively leveraging capacity and tiering of information storage resources
- Automation:** Automate data protection and retention based on policies, automate provisioning or configuration changes and automate data migration



Data Protection

Protect business critical application data ensuring rapid recovery, data reliability and integrity

Prevent legal exposure through data encryption and key management



Storage Resource and Infrastructure Management

Enable visibility and automate control through centralized management and automated provisioning

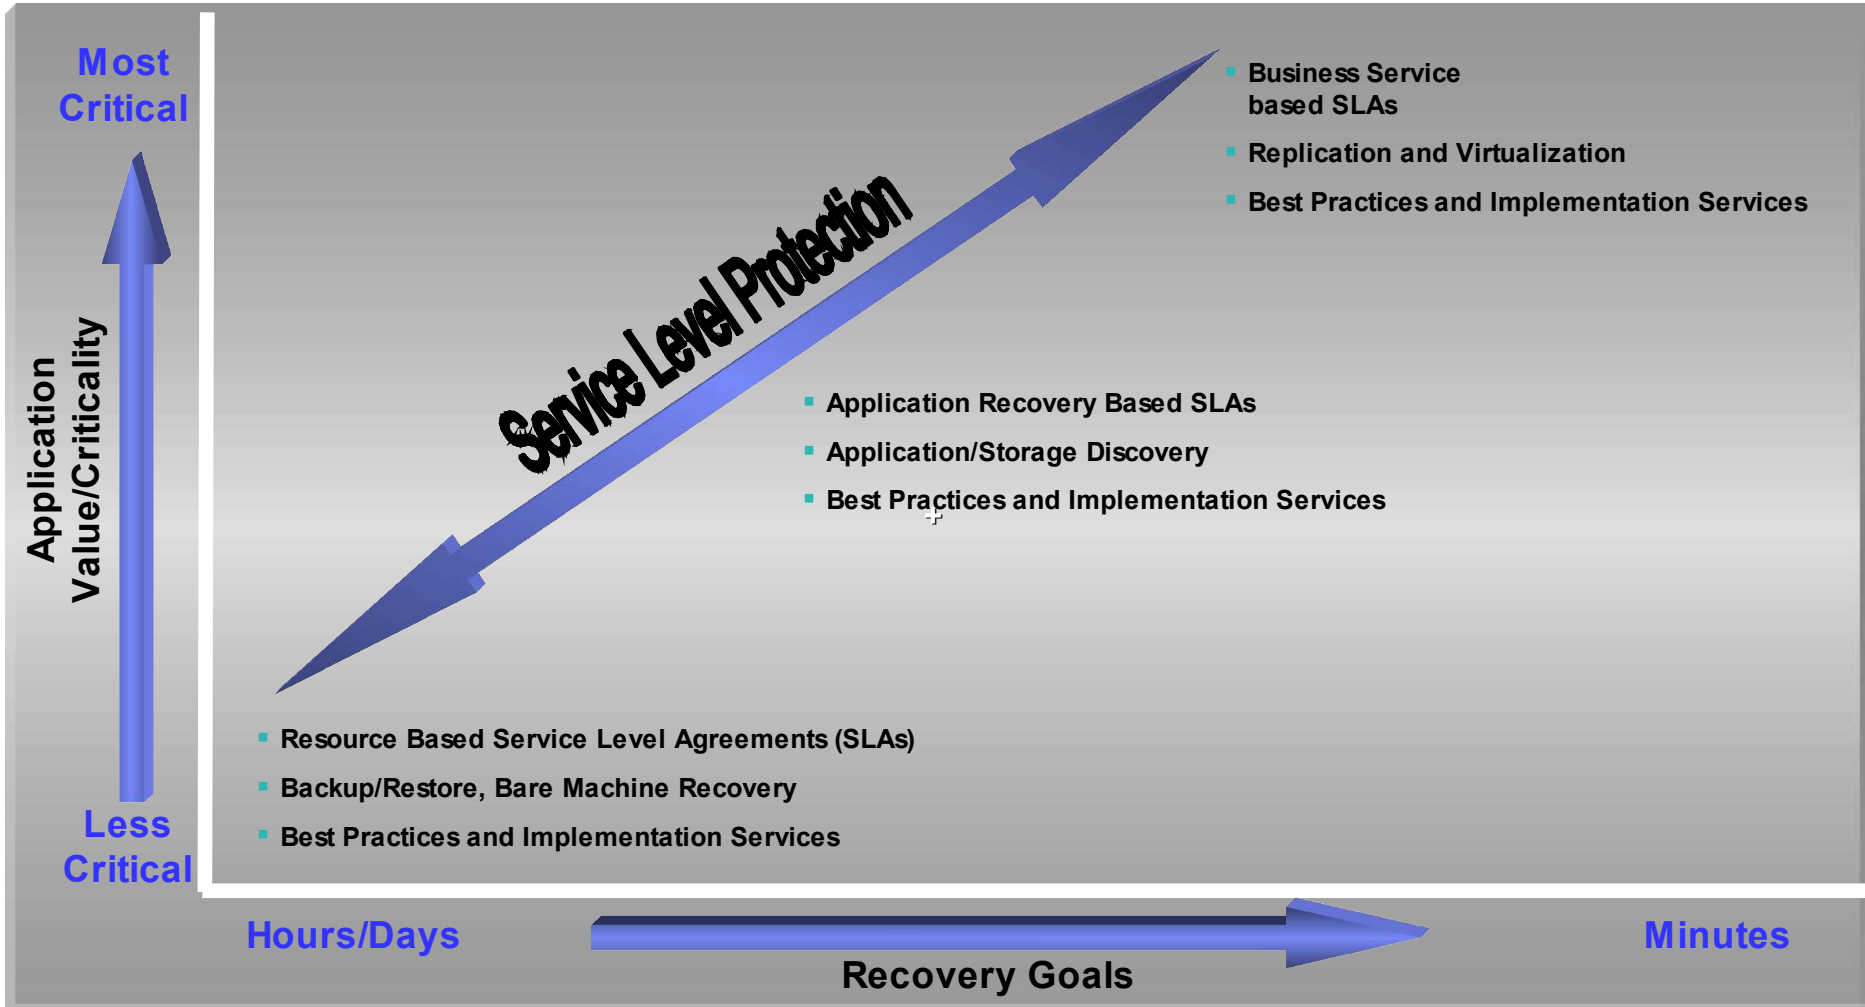
Optimize storage performance and capacity utilization, lowering costs and improving service levels

Information Retention and Lifecycle Management

Keep information at the most efficient and cost effective tier for the current stage of its life

Keeps copies of active or inactive data for long-term storage for regulatory or bookkeeping requirements

Business Continuity for Service Level Protection




IBM and its Partners have capabilities to meet the full continuum of Service Level Objectives

Virtualization takes power efficiency to the next level

- Virtualization creates largest impact to Power/cooling in Data Centers
 - Virtualization improves server and storage utilization which:
 - Reduces the number of servers and therefore, Space, Power, & Cooling required
 - Allows for targeted thermal solutions for efficiencies
 - Allows for more efficient power distribution, backup, and regulation
 - Enables dynamic resource management for
 - Deployment and Operational Optimizations

- Some examples
 - Base Consolidation
 - Policy Based Consolidation
 - Service Oriented Resource Allocation
 - Simplified Technology Transitions
 - Policy Based Provisioning



Clients that drive the best results from ILM initiatives focus on 6 best practices

Data Rationalization

Storage Virtualization

Tiered Storage Environment

Information Management

**Storage Process, Organization,
Technology, and Governance
Model**

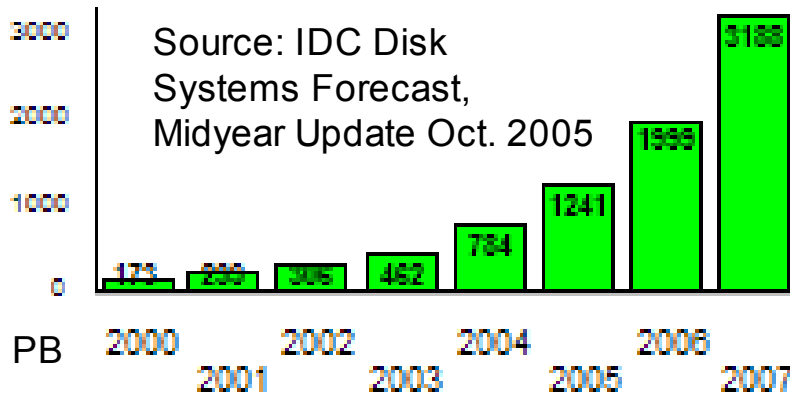
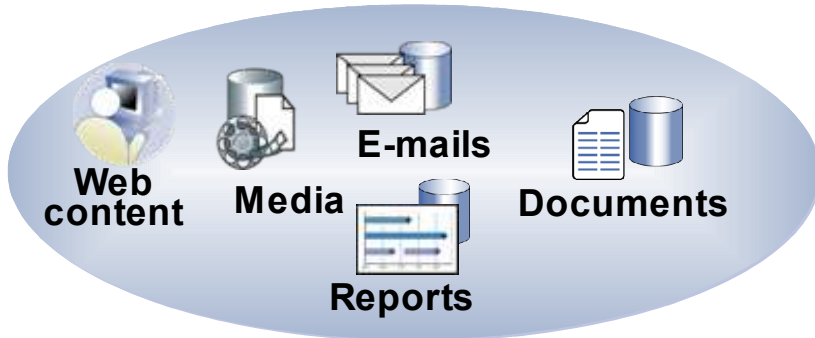
**Archiving and Information
Retention**



- ✓ Cost Reduction & Simplification
- ✓ Improve Efficiency
- ✓ Manage Risk and Streamline Compliance

These things are fundamental to a green storage agenda

Challenges – Variety, Volume, and Velocity



■ Variety of Information

Information Technology holds the promise of bringing a variety of new types of information to the people who need it

■ Volume of Data

Data is growing exponentially. IDC estimates continued 60% yearly growth of new disk PB shipped

■ Velocity of Change

IT Organizations are under tremendous pressure to deliver the right IT services. 85% of problems are caused by IT staff changing something. 80% of problems not detected by IT staff until reported by users.


80% of IT problems are reported by end users




Storage Management Challenges - Real and Growing

Like labour costs, storage management is taking an ever larger share of the budget

Disk storage is growing rapidly




Storage-related expenditures*, as a % of IT budgets, is also growing rapidly



Source: International Technology Group, Sept 2003

* hardware, software, storage networking, personnel, backup operations, recovery, security

Four Fundamental Truths About Data



1. All data, when created, does not have equal value
 2. Data changes in business value and in service level requirements over time
- IT resources should be allocated according to the value of data
 - Data must be managed and leveraged effectively throughout its entire lifespan ... **data outlives media**

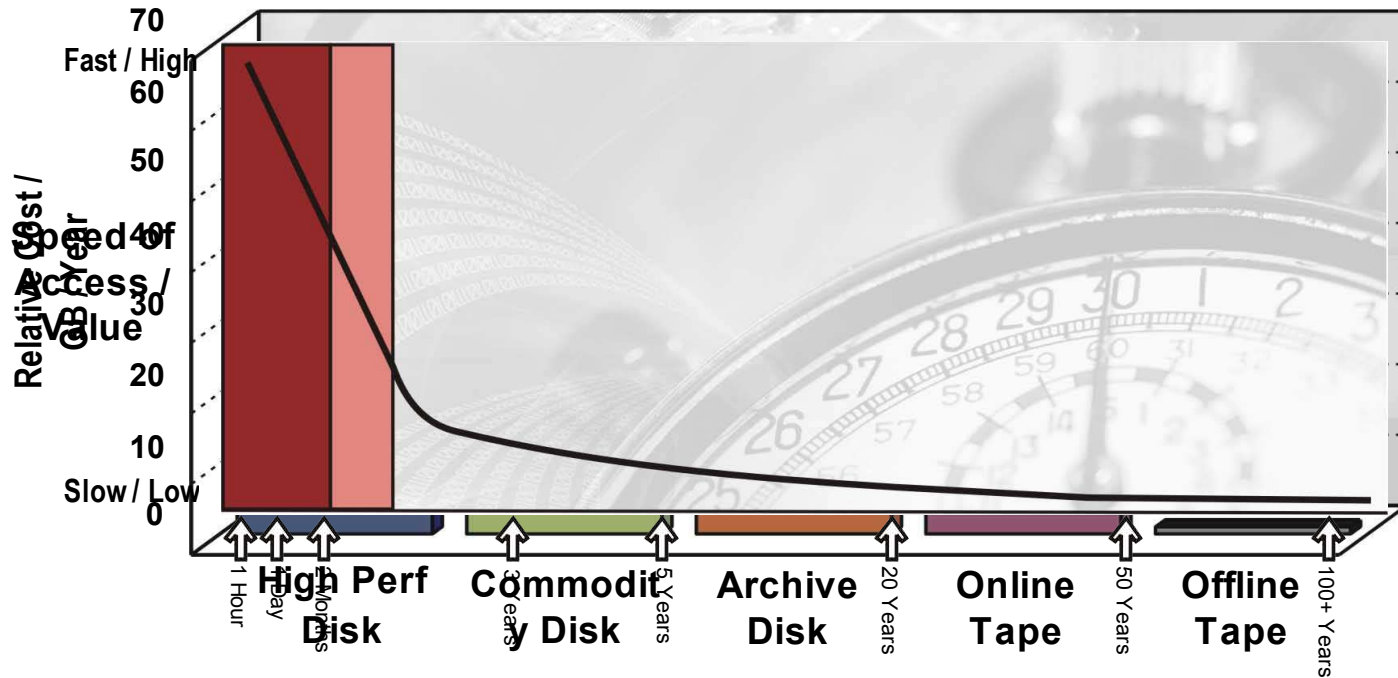
Six Best Practices to Consider

| | |
|---|---|
| Data Rationalization | <ul style="list-style-type: none"> Establish groups of valid and invalid data Determines opportunities to reclaim and consolidate storage |
| Storage Virtualization | <ul style="list-style-type: none"> Combines physical capacity from multiple storage systems into a single logical storage pool which can be centrally managed |
| Tiered Storage Environment | <ul style="list-style-type: none"> Align variable cost hardware types with information classes and classes of services |
| Information Management | <ul style="list-style-type: none"> Leverage information assets to improve business decisions – offering better access to data integrated across the organization |
| Storage Process, Organization, Technology, and Governance Model | <ul style="list-style-type: none"> Establish a governance model comprised of process, organization, technology and service management |
| Archiving and Information Retention | <ul style="list-style-type: none"> Enhance systems performance while enabling organizations to better manage risk and streamline regulatory compliance |



The Different Stages in Data’s Lifecycle

As data loses its value ... it should be moved onto lower cost storage



But this should be a well managed process:

- keeping within service management priorities placed on the data by the business, and ...
- complying with security, regulatory or other external requirements



“THINK” About Data 2008

The Steps and Components of IBM's Data Lifecycle Management Solution

IBM Tivoli Software

© 2008 IBM Corporation

The Steps in Data Lifecycle Management

- **Step1: Planning and Assessment**

 - Categorize Data according to its Value
 - Identify, Evaluate, Control & Predict
 - Establish Policies

- **Step 2: Active Data Management**

 - Pooling of storage by class of service
 - Virtualize physical storage
 - Policy based data placement & migration

- **Step 3: Inactive Data Management**

 - Business Continuity & Storage Management
 - Application Protection, HSM & Archiving
 - Long Term Retention for Compliance

Visibility

Know your data

TotalStorage Productivity Center

Control

Manage your data

SAN Volume Controller
Tivoli Storage Manager

Automate

Dispose of your data

Tivoli Storage Manager
Archive Manager & DR550

Data Lifecycle Solutions From IBM

- TSM
 - Tivoli Storage Software (world-wide) exceeded market growth for 11 consecutive quarters
 - TSM transactional revenue grew 50% YoY in NE IOT in 2007
 - TSM & HSM is No.1 in the HSM market
 - TSM has been a leading Enterprise class backup tool for distributed systems since its introduction in 1993

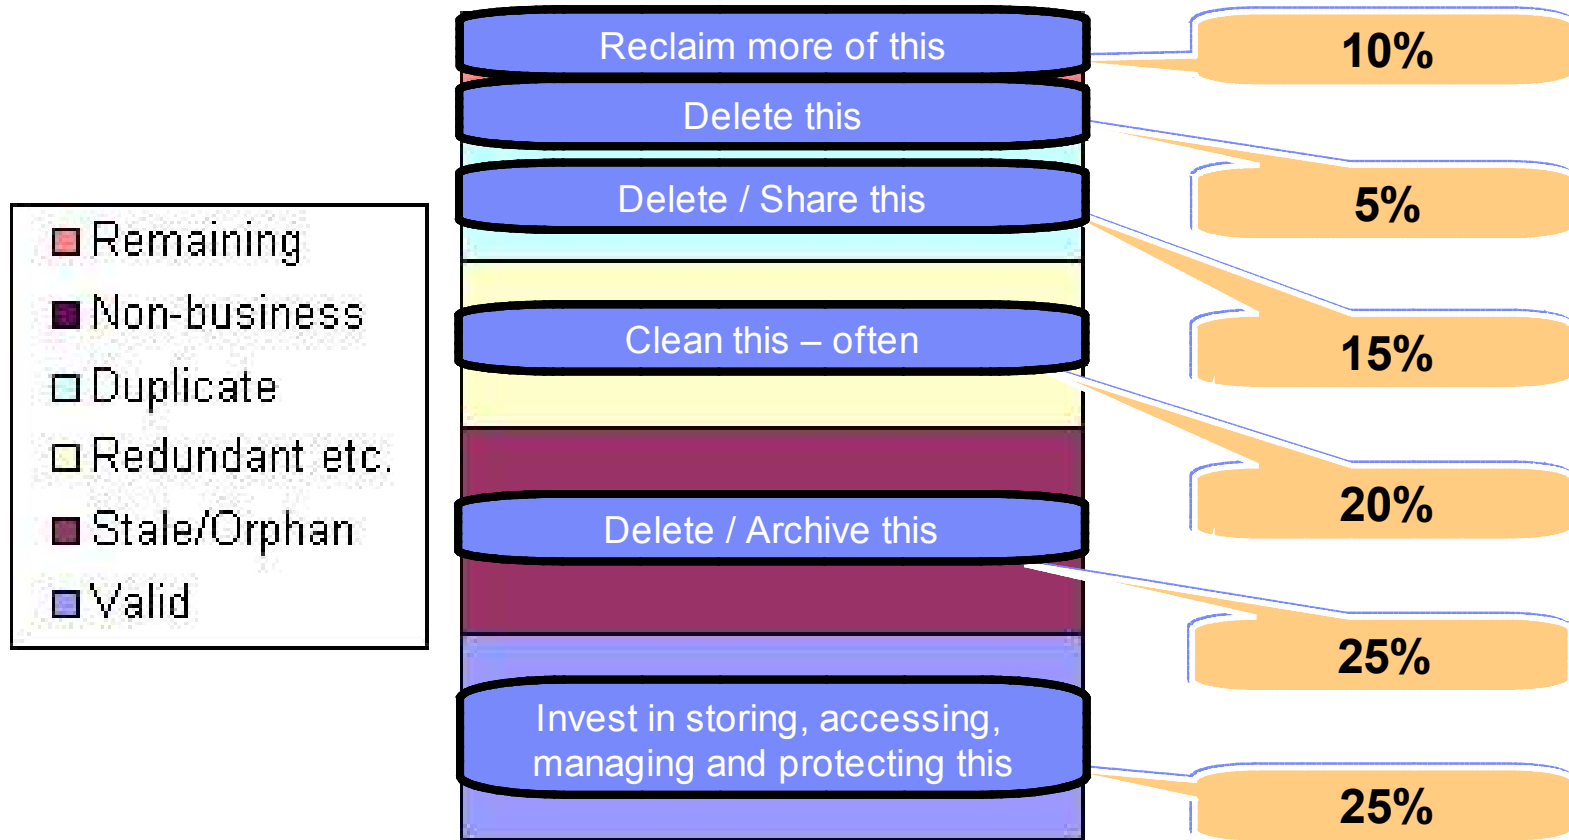
- TPC
 - TPC is leading the Gartner Magic Quadrant for SRM Tools (March 2007)
 - TPC transactional revenue grew 92% YoY in NE IOT in 2007
 - TPC is built to Open Standards and works with all SM-S compliant vendors
 - The market is only 20% penetrated
 - TPC CAGR of 35%

- SVC
 - Over 10,000 SVC nodes have been shipped since 2003
 - IBM is a leader in virtualised disc arrays
 - SVC transactional revenue grew 79% YoY in NE IOT in 2007
 - Enterprise Strategy Group reports that early virtualization adopters on average *every year* save:
 - 24% on hardware costs
 - 16% on software costs
 - 19% on SAN administration costs
 - SAN Volume Controller demonstrates scalability with the fastest Storage Performance Council benchmark results for any disc controller from any vendor
 - SAN Volume Controller can virtualize IBM and non-IBM storage (over 120 systems from IBM, EMC, HP, HDS, Sun, Dell, NetApp, Fujitsu, NEC, Bull)

DLM Step 1: Planning and Assessment

Visibility

Categorize Data according to its Value



By defining information as “categories” or “classes” of data, DLM enables the creation of effective data management standards and policies

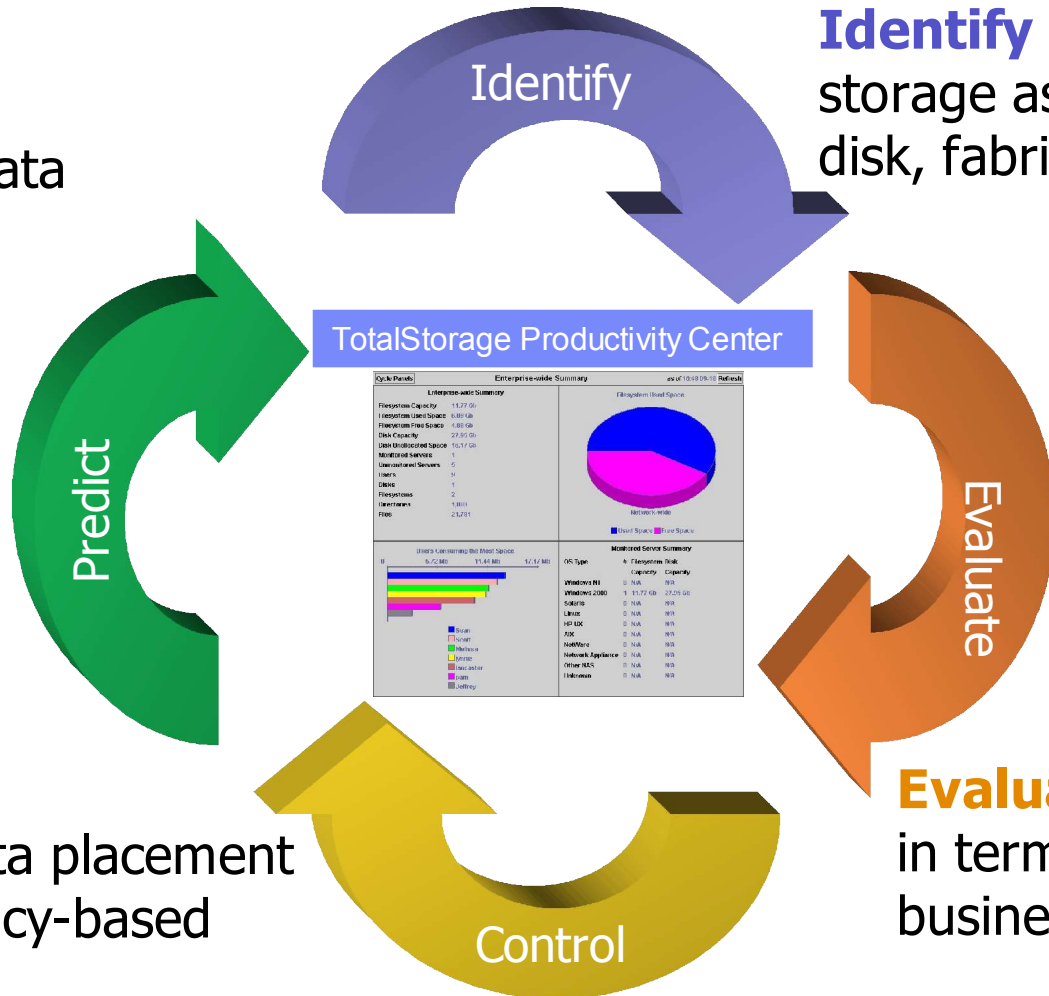
DLM Step 1: Planning and Assessment

Identify, Evaluate, Control & Predict

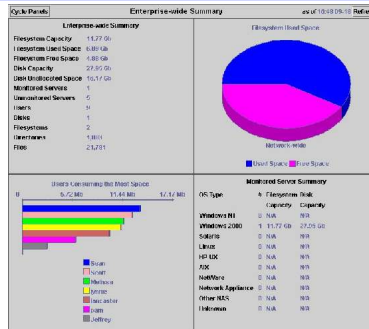
Visibility

Predict data usage and growth

Identify enterprise storage assets (data, disk, fabric, ...)



TotalStorage Productivity Center



Control data placement through policy-based automation


Evaluate data in terms of its business value

DLM Step 2: Active Data Management

Control

Virtualize physical storage

Make changes to the storage without disrupting host applications




SAN

Manage the storage pool from a central point



Apply copy services across the storage pool

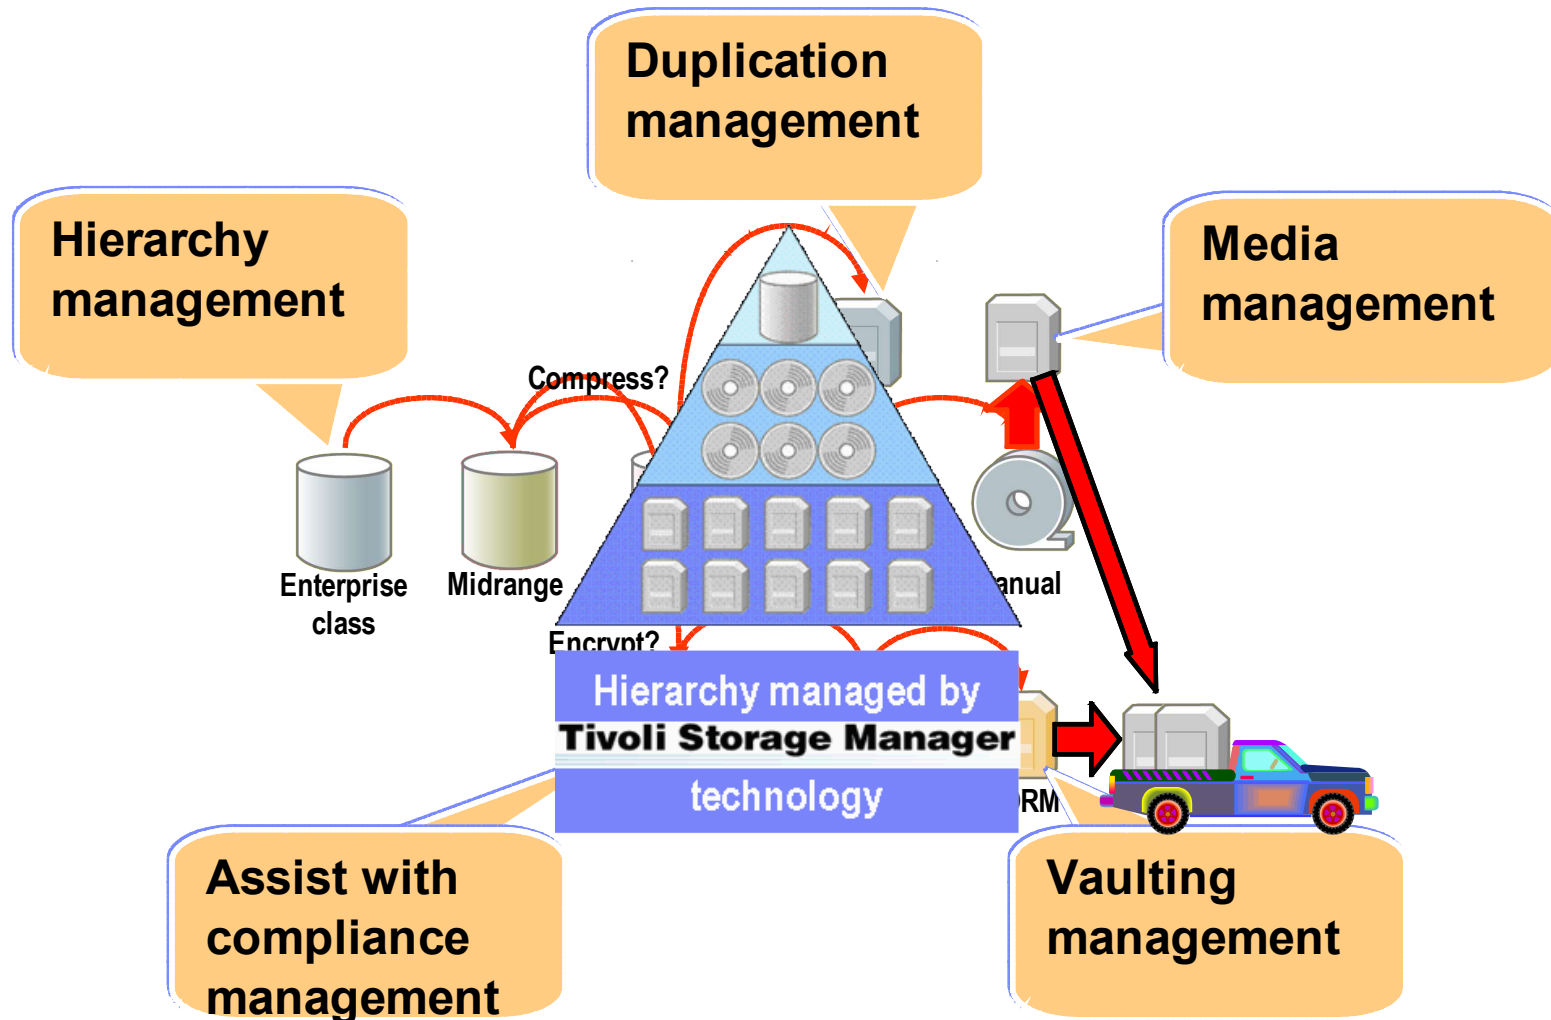


Storage Pool



Combine the capacity from multiple arrays into a single pool of storage

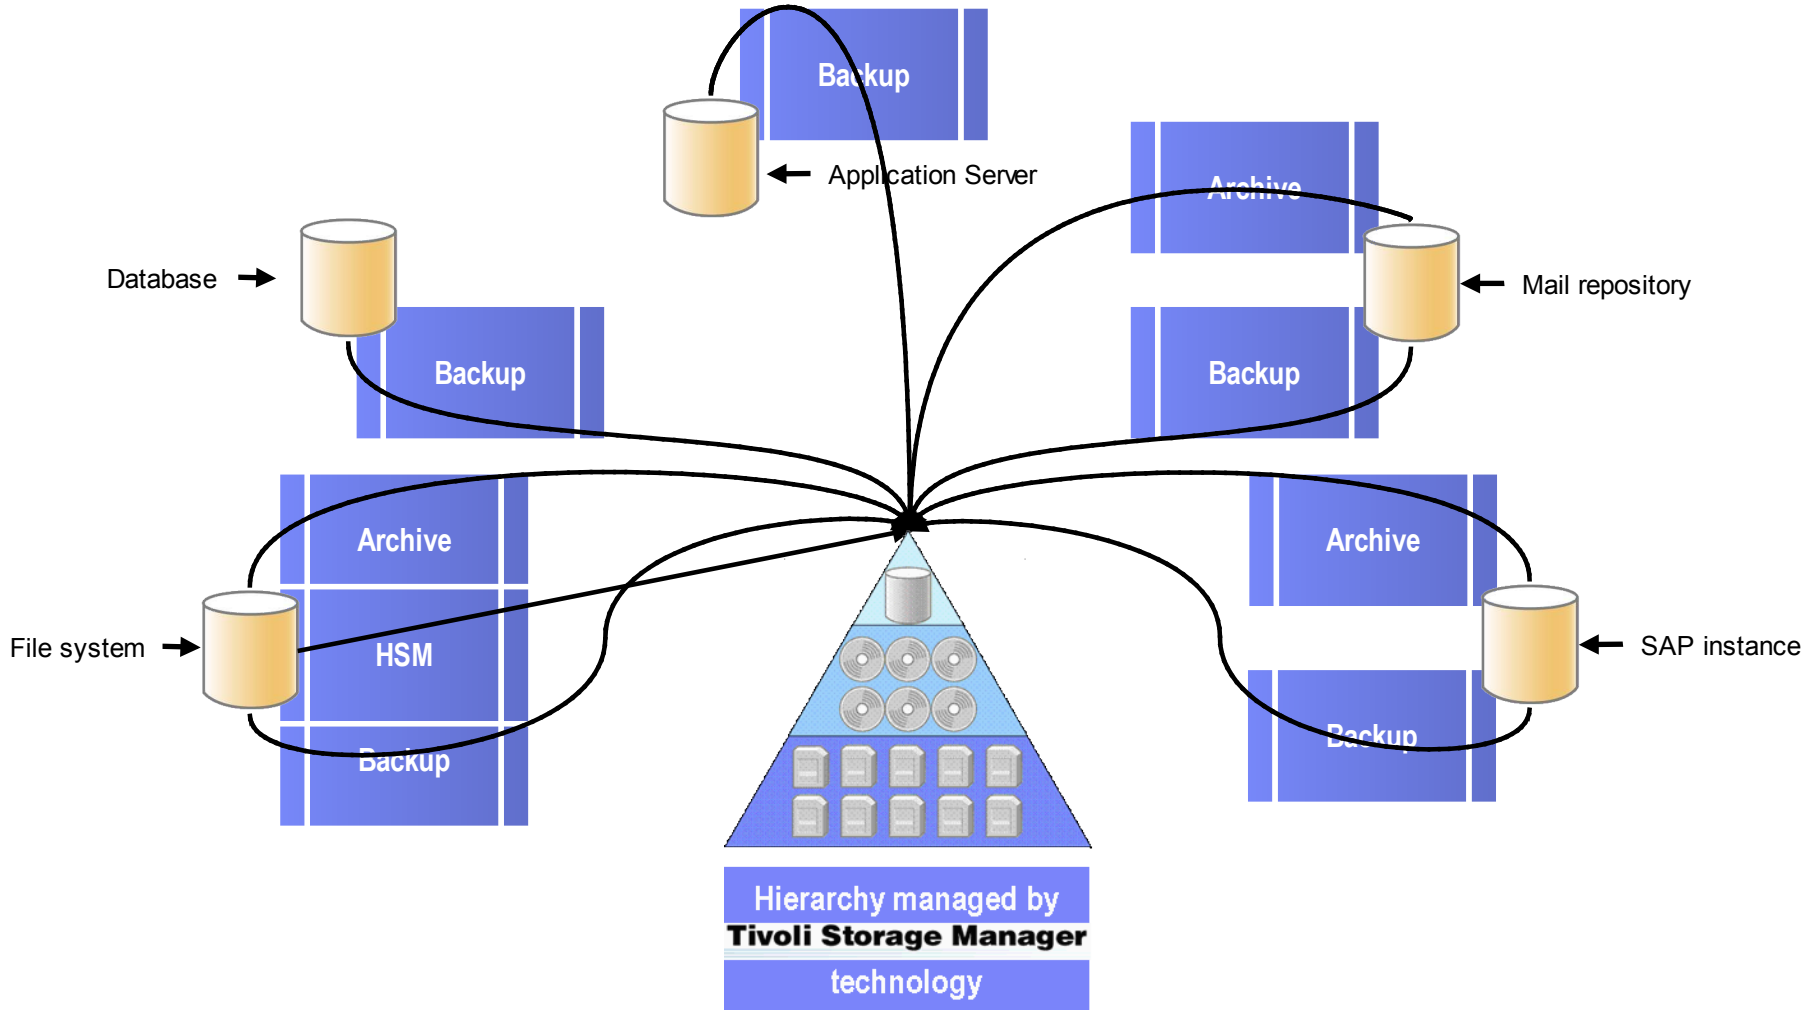
DLM Step 3: Inactive & Active Data Management **Automate** Business Continuity & Storage Management



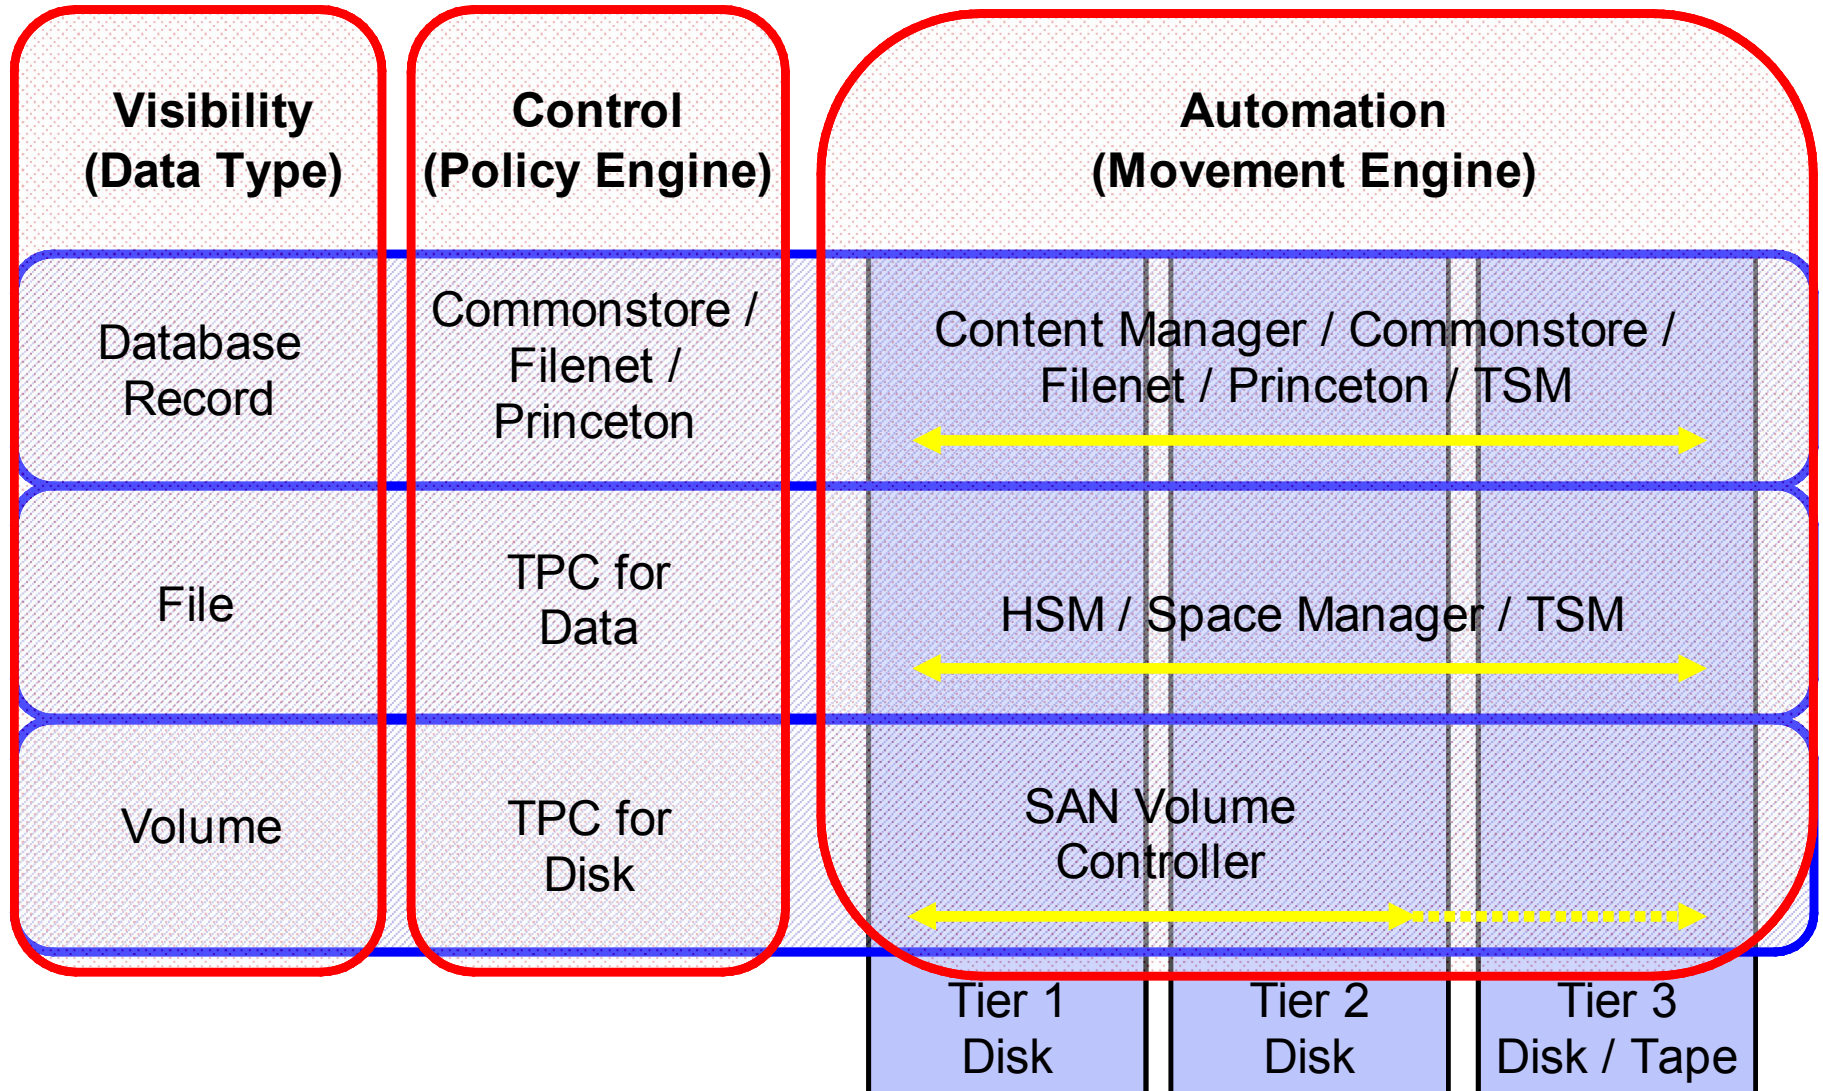
DLM Step 3: Inactive Data Management

Application Protection, HSM & Archiving

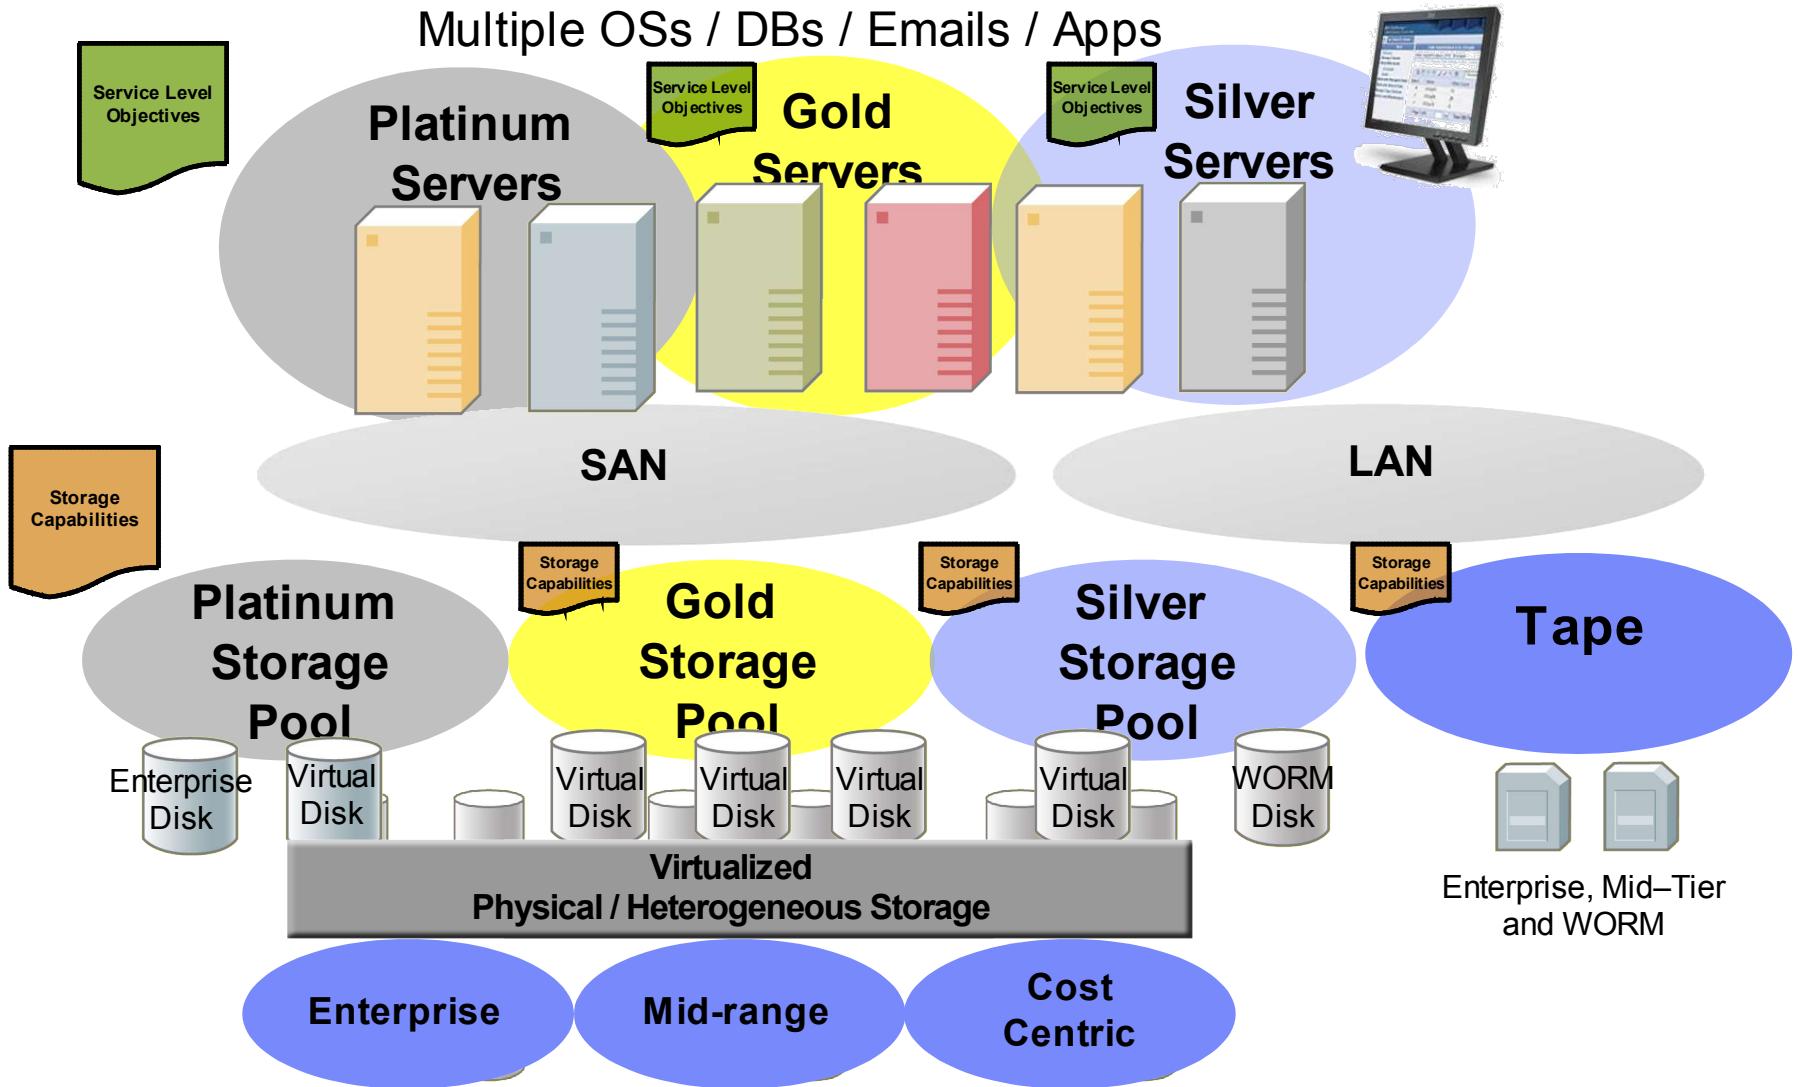
Automate



Making DLM Work: Visibility-Control-Automation



Tying it all together with Service Management





“THINK” About Data 2008

The Business Value of Data Lifecycle Management

IBM Tivoli Software


© 2008 IBM Corporation

DLM Helps Lower Storage TCO

Average TCO for Storage Systems

Source: Gartner Group

- Better manage explosive data growth
- Simplify storage management across tiers of heterogeneous storage
- Reduce time to provision storage



- Monitor and automate backup and archive of corporate files

- Lower hardware costs
- Storage resource optimization
- Make more efficient storage purchases based on capacity demands

- Predict out of space conditions
- Detect availability issues in SAN
- Monitor SAN and Disk performance hot spots

DLM Helps Realize Ongoing Savings

“Hard” Savings:

- **Storage Management Labour Savings:**
 - The reduction in workload for current IT staff managing IT and business systems storage and tasks, and the ability for the staff to scale more effectively

- **Storage Purchase Avoidance:**
 - The reduction in the need for on-line storage through the automatic and intelligent management of company information to near-line and off-line storage resources

- **Tape System Purchase Avoidance:**
 - The reduction in the need for tape system and media through more efficient management of backup and compression technology

- **Network Bandwidth Investment Avoidance:**
 - The elimination of backup window issues, leading to a reduction or elimination of planned network bandwidth improvements to support backup window requirements

“Soft” Savings:

- **Restore Time Benefits:**
 - A reduction in lost business and productivity through faster data restores

- **Backup Coverage Risk Avoidance:**
 - An increase in backup coverage resulting in reduced risk of data loss

- **Increased Availability:**
 - Reduced unplanned downtime for business systems, applications and infrastructure components

Additional Savings:

- **Data Migration:**
 - Migrate data from one device to another without taking storage offline – better reallocation, scalability and upgradability without disruption

- **Competitive Savings:**
 - Cost effective, flexible copy services can replace multiple costly existing tools

How Are DLM “Hard” Savings Achieved?

Storage Purchase Avoidance savings are achieved by:

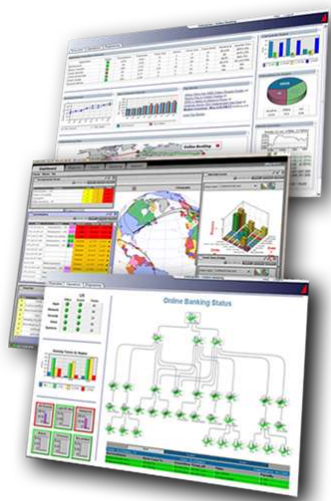
- Identifying unused data that can be archived or deleted, which helps reclaim wasted space and reduces the need to purchase additional storage
- Using hierarchical storage management (HSM) techniques to significantly reduce wasted space in expensive online storage devices by systematically migrating rarely used files according to pre-defined policies
- Treating all storage as a common pool, which greatly improves overall storage utilization
- Supporting multiple storage profiles and advanced copy services, which facilitates the use of less expensive storage for secondary copies

IT Operations Labour savings are achieved by:

- Improving capacity and performance management and planning, by proactively monitoring and enforcing storage usage policies, and by gathering key configuration, network statistics, and other status information
- Simplifying storage asset management through extensive reporting, and helping inventory control and recording by automatically discovering resources
- Easing administration through the use of a central user interface for configuring and managing all supported devices on the storage estate
- Providing monitoring and alerting, provisioning and error prediction capabilities

Summary

Data Lifecycle Management is a set of Storage Solutions that are perfectly aligned to the IBM Service Management Strategy



***Respond faster
and make better
decisions***



***Improve quality
and reduce risk***



***Lower costs
and build
agility***

Enabling customer innovation through **Visibility, **Control**, and **Automation****