

# Storage, data, and information: a brief overview

**John Wilkes, Beth Keer,  
Christopher Hoover, Alistair  
Veitch, Pankaj Mehra**

HP Labs, Palo Alto, California  
December 2006



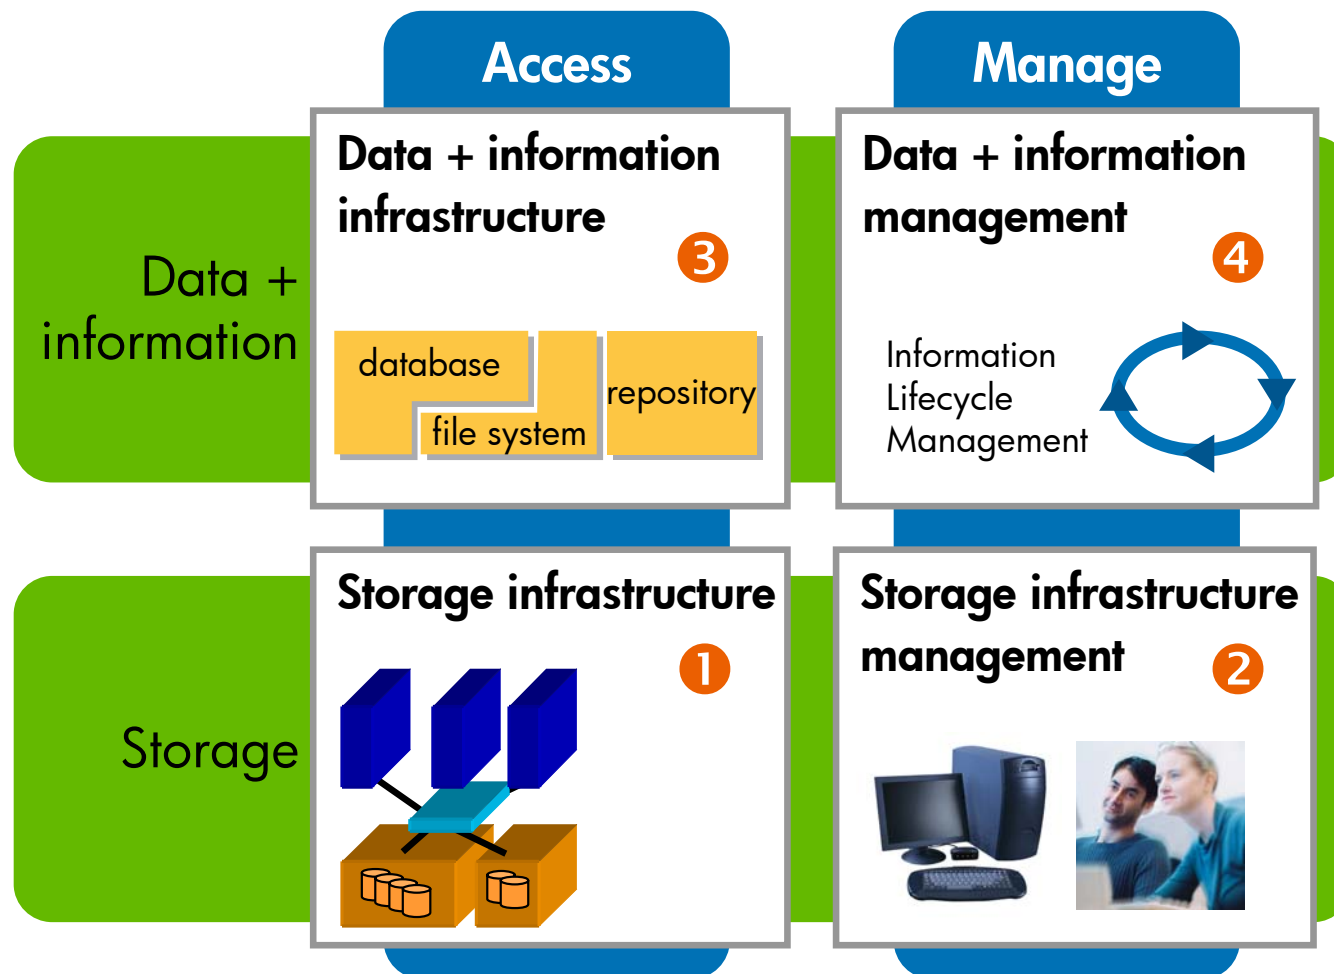
# Roadmap



Enterprises  
and people



Business logic, processes & applications



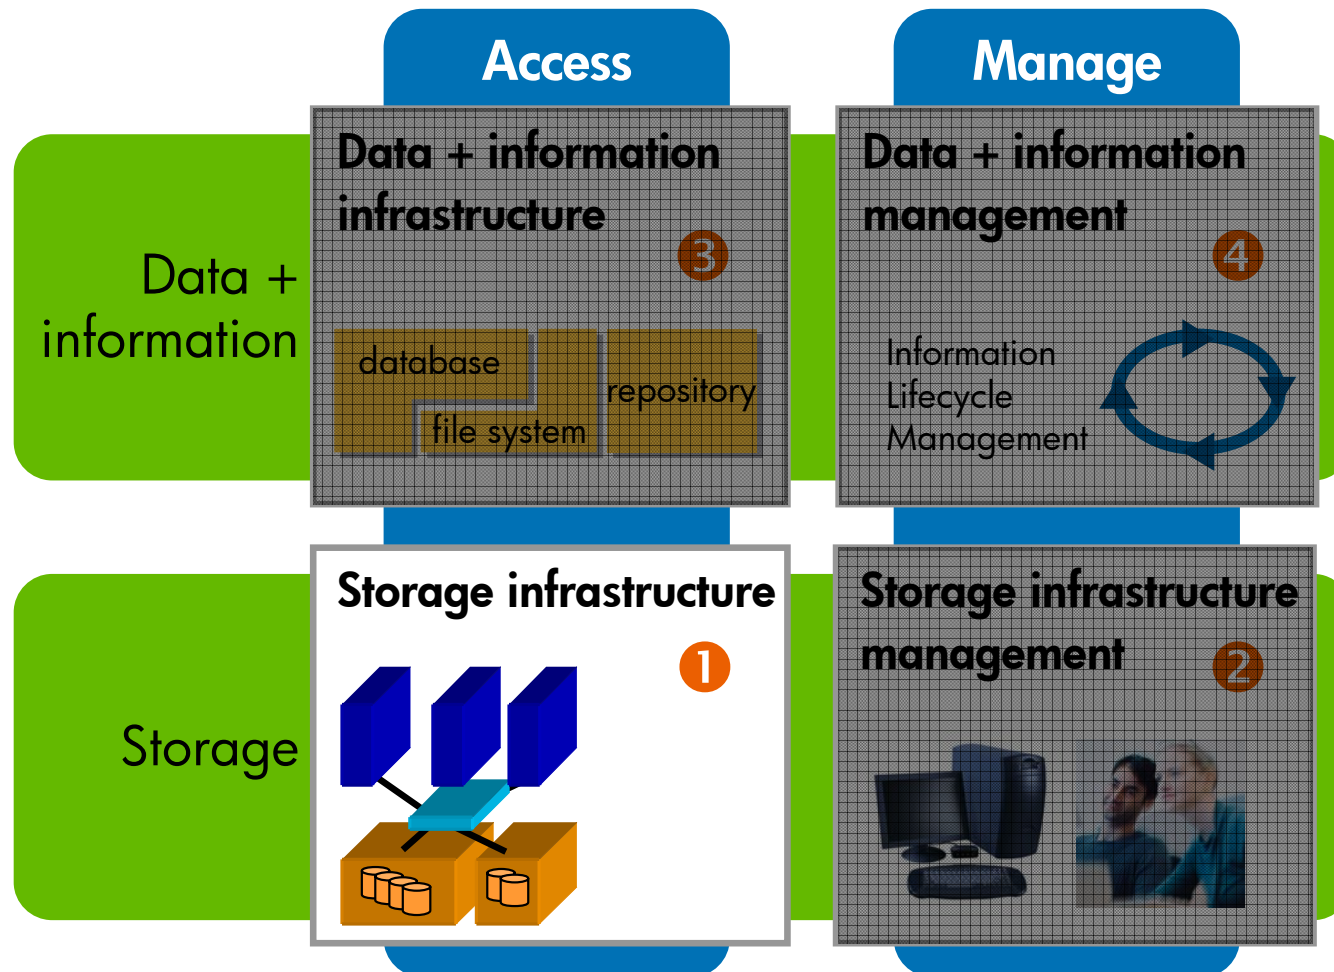
# Storage infrastructure



Enterprises  
and people




Business logic, processes & applications






# Storage infrastructure: it's all about blocks

- contiguous set of bytes (characters)
- fixed length (typically 512 bytes)
- stored on disks or arrays of disks (and tapes)

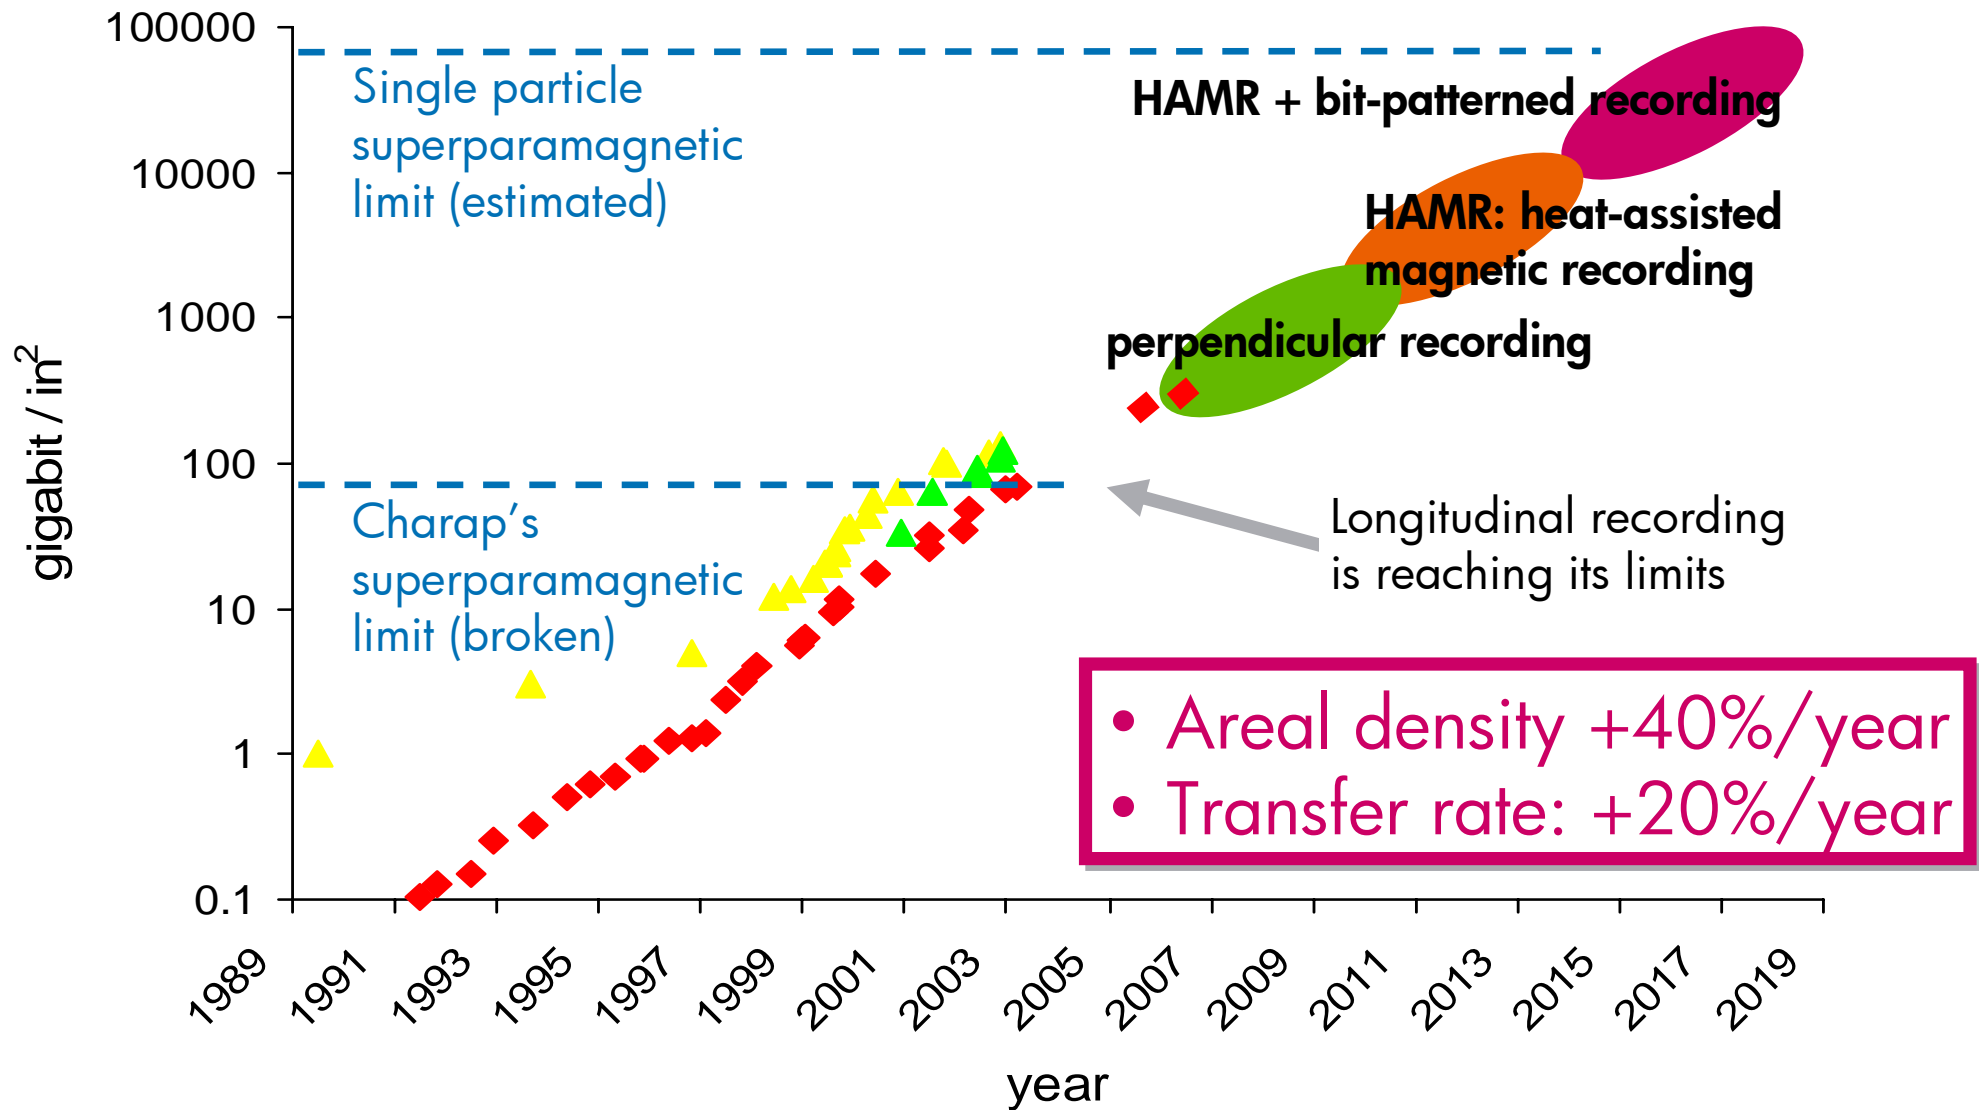




# Disk drives simplified (2)



# Disk drive trends: areal density



source: Seagate

December 8, 2006





# Disk drive trends: scaling up a disk access

	<i>real time</i>	<i>scaled-up time</i>
cache	1ns	look-up on paper: few seconds
memory	100ns	retrieve book from room: 1 minute
disk	5ms	get entire bookcase delivered: <b>35 days</b> new bookcases: 300/day



# Types of drives

- Segments
  - Consumer Electronics
  - Mobile
  - Desktop
  - Enterprise
- Differentiators
  - capacity
  - performance
  - reliability
  - form factor (size)
  - interface
  - price (\$/GB)
  - supplier support

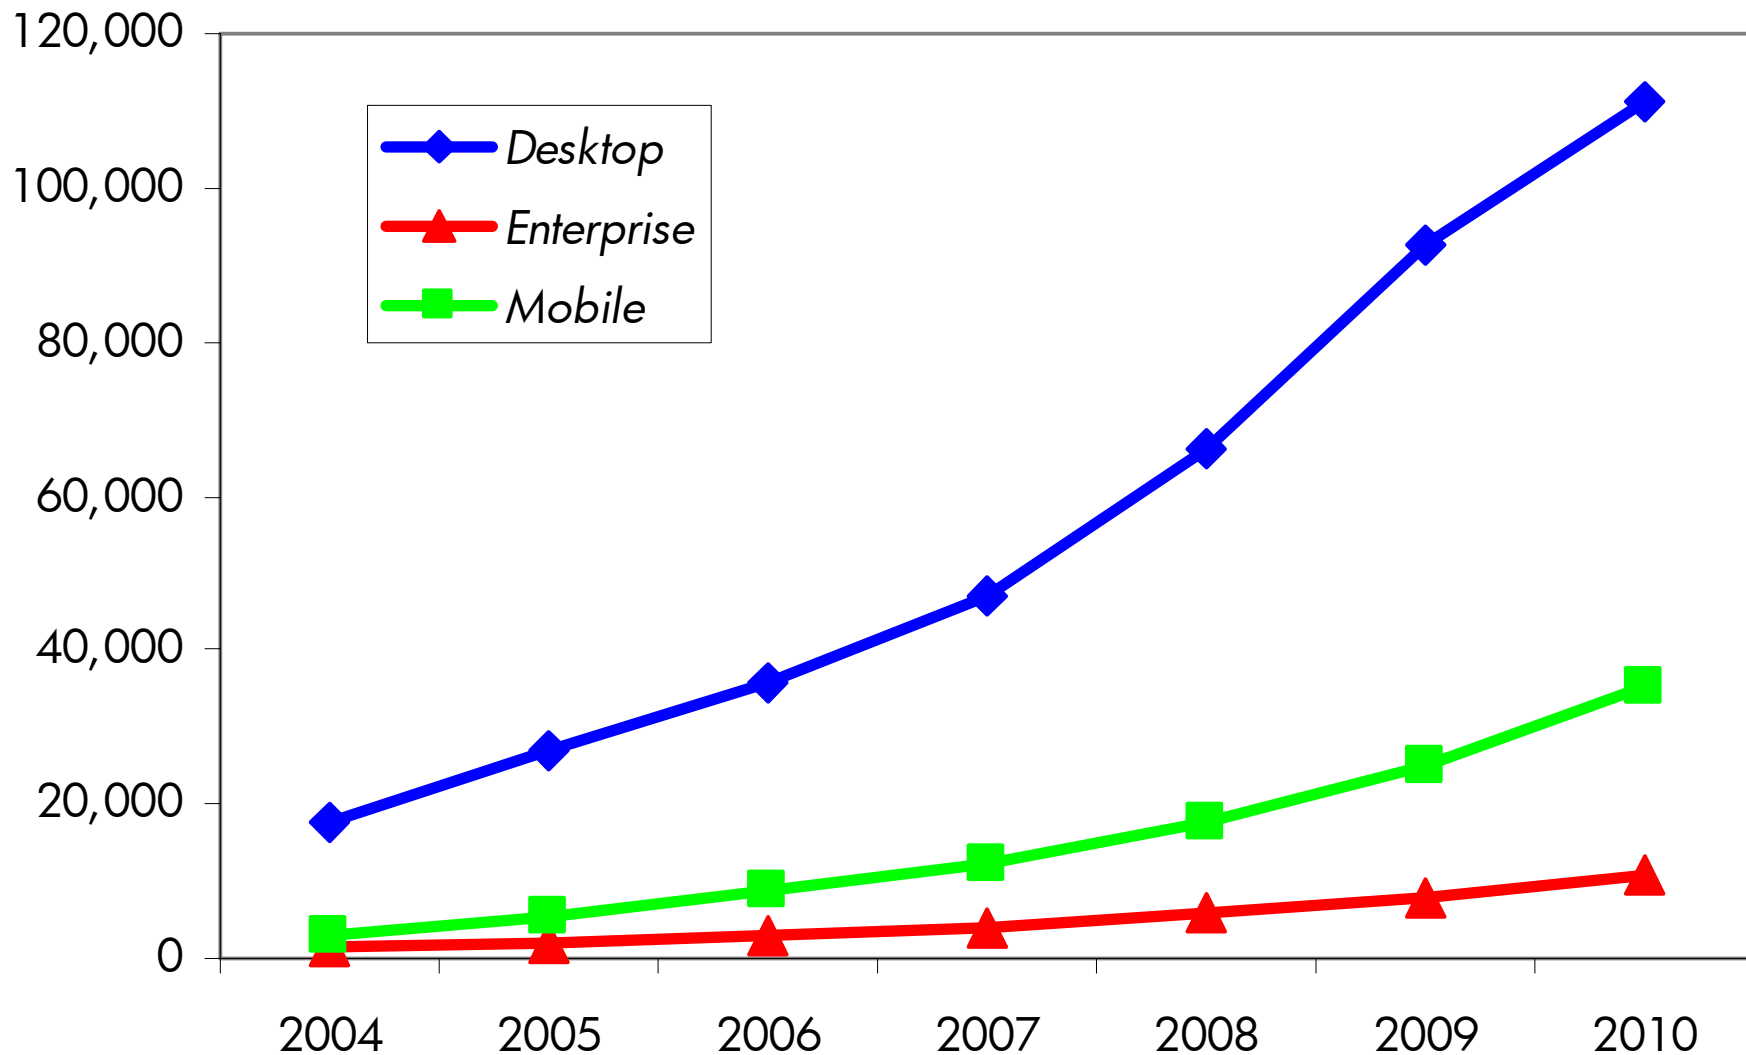




# Disk drive shipments (capacity) by drive type



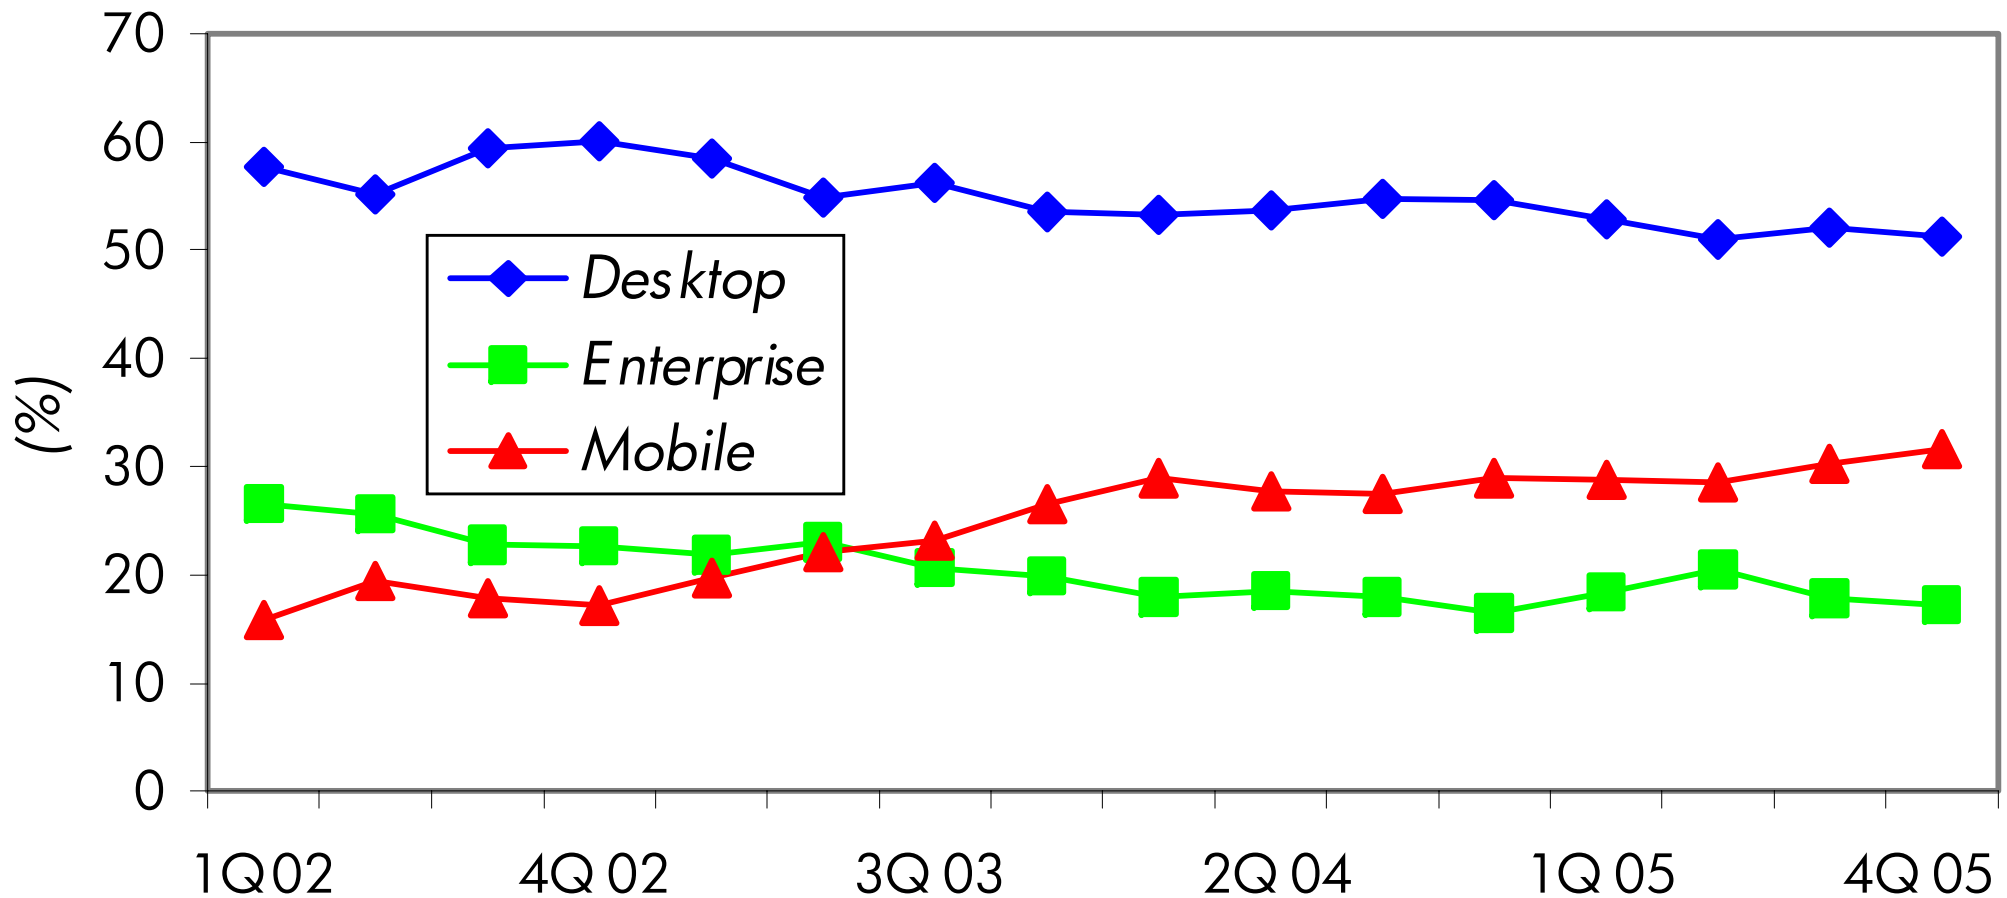
Petabytes shipped worldwide



source: IDC

December 8, 2006

# Disk drive shipments (revenue) by drive type

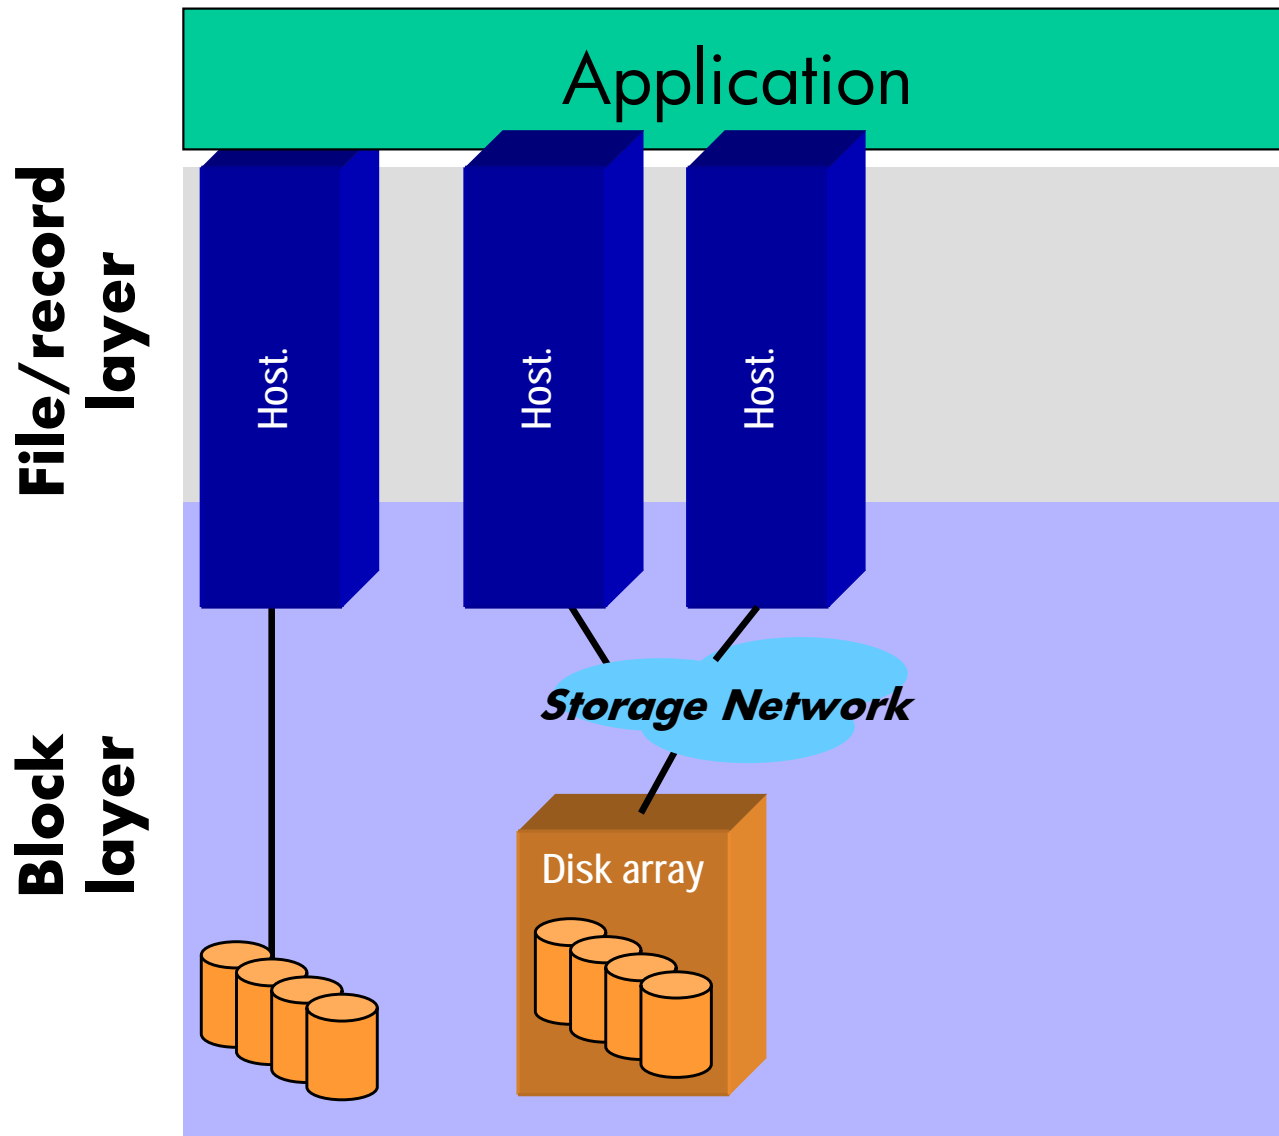


source: IDC

December 8, 2006



# Block device attachment types





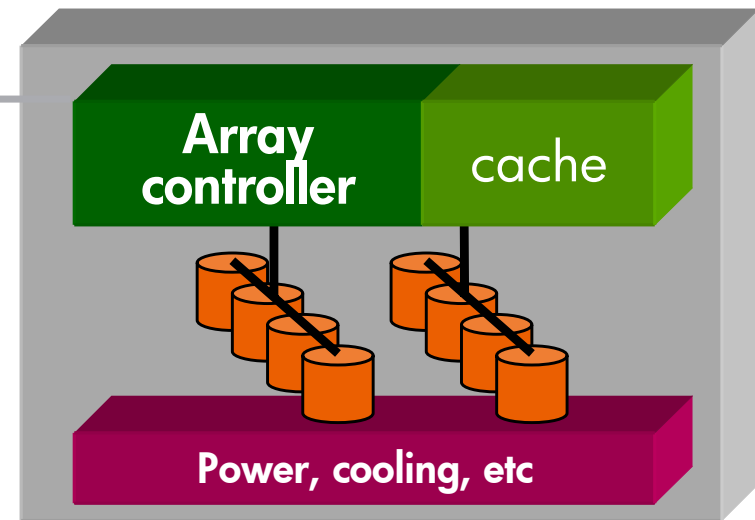
# Disk arrays

- A disk array is just:
  - mechanical enclosure
  - power and cooling
  - controller(s)
  - network connection
- That offers:
  - failure protection
  - aggregation (capacity, performance)
  - virtualization



**HP MSA disk array**


to host ←





# Disk arrays

**HP XP1024 disk array**

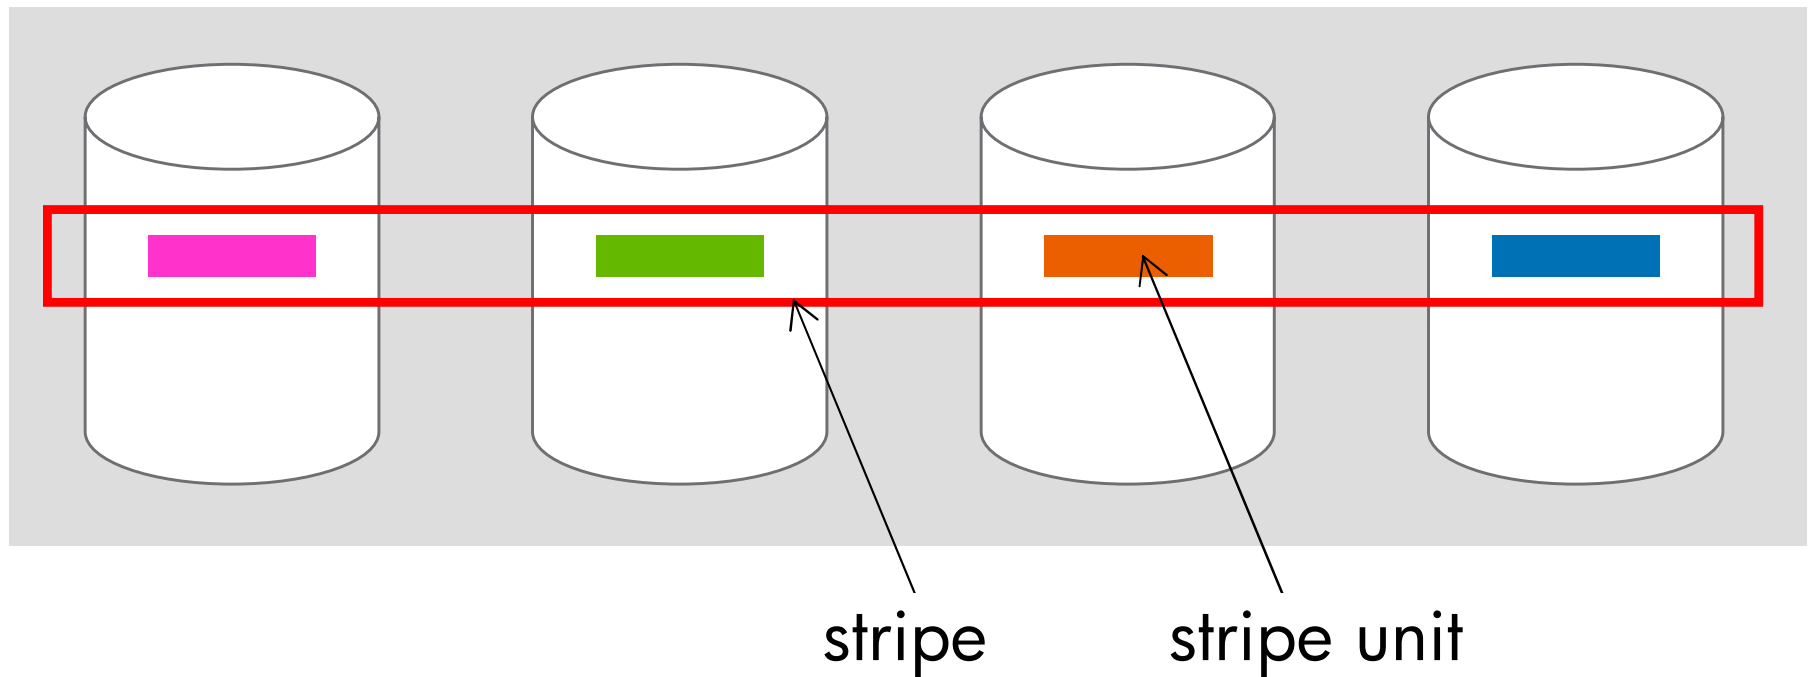


# Disk arrays

## Protection techniques



RAID0: **striping** — no redundancy






# Disk arrays

## Protection techniques



RAID1: **mirroring** — full redundancy






# Disk arrays

## Protection techniques

RAID5: **parity-protection** —  
partial redundancy





# Disk arrays aggregation + virtualization

- “Slice and dice” block storage into *volumes*
  - larger or smaller than a single disk
- Usually done by software on the disk array
  - but it can also happen in other places



# Don't forget tapes ...





# Storage infrastructure

## key messages

- It's all about *blocks*
- The physical characteristics of *disk drives* dictates a lot of the technology for block storage
  - performance
  - failure tolerance
- *Software* adds value to and differentiates disk arrays via aggregation and virtualization

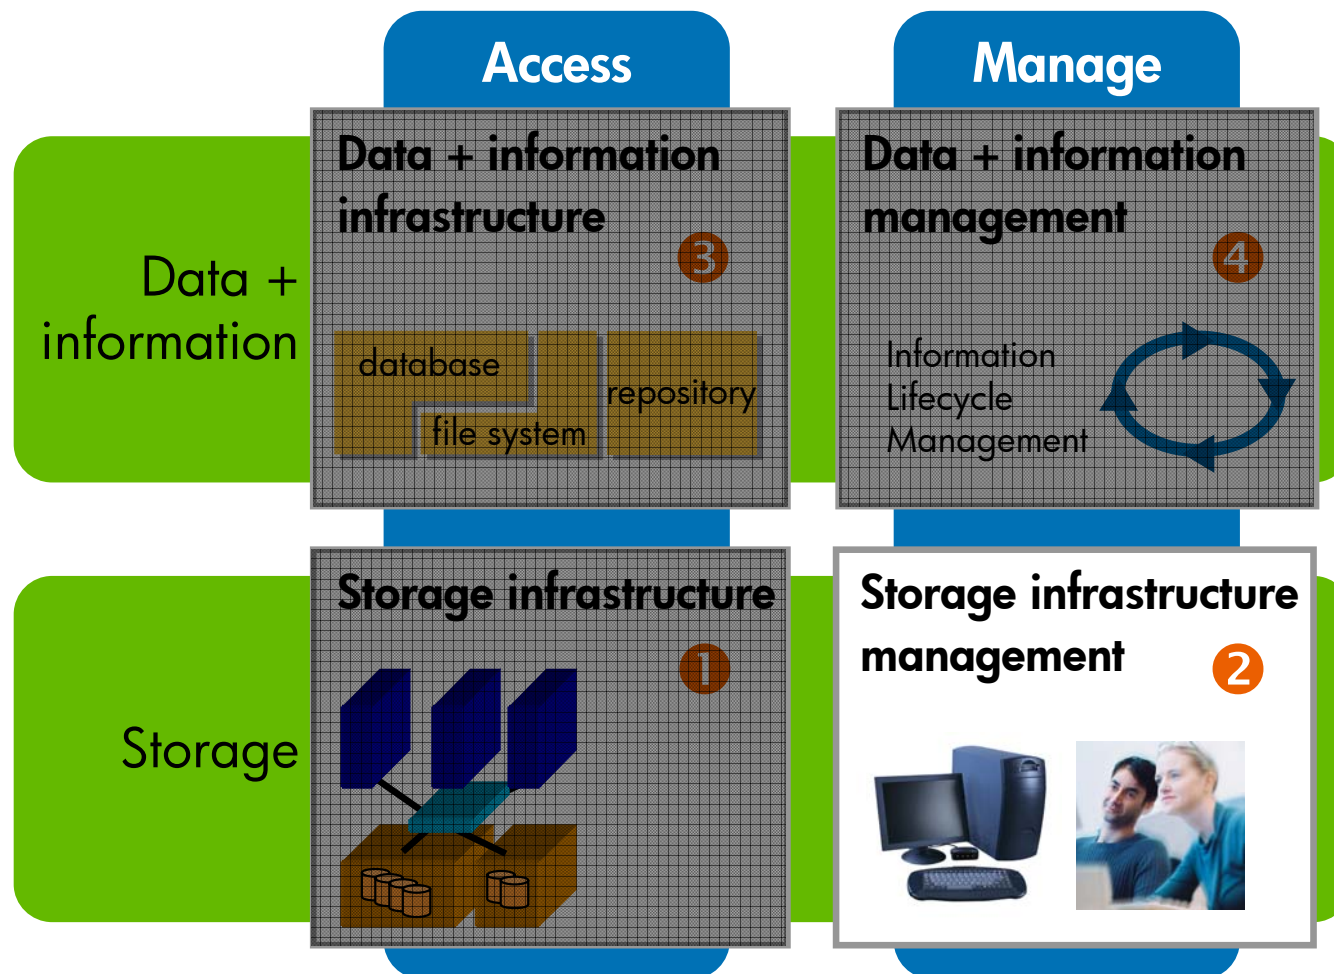
# Storage infrastructure management



Enterprises  
and people



Business logic, processes & applications









# Q: what is storage management?

## A: all of these ...

- 1 Activate licensed features in fabric elements
- 2 Add SAN resource domain (fabric + devices) to existing installation
- 3 Add host to existing FC fabric
- 4 Add hub to existing FC loop/fabric
- 5 Add peripheral disk device to bridge
- 6 Add peripheral disk device to storage array
- 7 Add port to storage array
- 8 Add switch to existing FC fabric
- 9 Add tape drive or library to bridge
- 10 Analyze SAN topology for single points of failure
- 11 Analyze SAN topology for traffic hot spots
- 12 Analyze device behavior to predict failures
- 13 Assign IP addresses to SAN components
- 14 Assign OS to run in partition/on platform
- 15 Assign action for event response
- 16 Assign free volume to OS/application
- 17 Audit actual configuration against planned/intended config
- 18 Audit firmware configuration
- 19 Audit software configuration
- 20 Boot OS in partition/on platform
- 21 Change OS or OS FC driver revision
- 22 Change cabling to service/management modem(s)
- 23 Change cabling to service/management network hub
- 24 Change cabling to service/management serial hub
- 25 Change cabling to service/management server(s)
- 26 Change fabric cabling to HBA
- 27 Change fabric cabling to use spare port
- 28 Change fabric internal topology (ISL's)
- 29 Configure and compile OS kernel
- 30 Convert existing fabric to cascaded fabric
- 31 Convert existing fabric to fully redundant fabric
- 32 Convert host bus adapter from FC-SW to FC-AL or vice versa
- 33 Convert single-initiator SCSI bus to multi-initiator
- 34 Convert two existing fabrics into a single fabric
- 35 Diagnose I/O errors
- 36 Diagnose directed path/device I/O (online, offline)
- 37 Diagnose system crash/hang
- 38 Download FC host bus adaptor firmware
- 39 Download FC switch firmware
- 40 Download storage array firmware
- 41 Download tape library firmware
- 42 Failover broken host bus adapter
- 43 Failover broken intra-switch port or trunk (ISL)
- 44 Failover broken storage array port or link
- 45 Failover broken switch port or link
- 46 Find physical location of specific device or fabric element
- 47 Install new FC-AL loop
- 48 Install new FC-SW fabric
- 49 Install new host
- 50 Install service/management software (servers, agents)
- 51 Install software, patches, service packs
- 52 Install storage array (Shark, EMC, HDS, Clariion)
- 53 Install tape system with shared drives
- 54 Install tape system with unshared drives and shared robotics
- 55 Mount OS file systems
- 56 Online/offline FC-SCSI bridge
- 57 Online/offline OS volume manager objects (mirrored, concatenated, etc)
- 58 Online/offline host bus adapter
- 59 Online/offline intra-switch trunk (ISL)
- 60 Online/offline path in multipath-capable OS
- 61 Online/offline peripheral device
- 62 Rebuild system for disaster recovery
- 63 Replace FC-AL hub
- 64 Replace FC-SCSI bridge (SAN Data Gateway, NUMA-Q FC Bridge)
- 65 Replace FC-SW switch (single switch fabric, multiple switch fabric)
- 66 Replace SAN management server
- 67 Replace failed director/controller in storage array
- 68 Replace host bus adaptor
- 69 Replace host
- 70 Replace peripheral device
- 71 Replace platform management server
- 72 Replace tape library robotics
- 73 Reserve tape media and storage slots within tape library
- 74 Reset/power-cycle FC-SCSI bridge
- 75 Reset/power-cycle entire installation (power-fail, first bringup)
- 76 Reset/power-cycle host platform
- 77 Reset/power-cycle peripheral devices (on bridge)
- 78 Reset/power-cycle storage array
- 79 Run offline diagnostics (using idle/disused system components)
- 80 Run online diagnostics (using "active" system components)
- 81 Sanitize used fabric elements to safely reuse in new fabric (clear NVRAM)
- 82 Set/view "POST" diagnostic level
- 83 Set/view "business continuation volumes" (BCV)
- 84 Set/view OS configuration files/registry
- 85 Set/view OS volume manager volumes
- 86 Set/view SNMP trap destination
- 87 Set/view backup schedule
- 88 Set/view event reporting threshold
- 89 Set/view event-/error-report destination
- 90 Set/view online diagnostics error threshold trigger
- 91 Set/view phone-home/email-home destination
- 92 Set/view service/management authentication (passwords)
- 93 Set/view storage array LUN masking and LUN mapping
- 94 Set/view storage array volume definition
- 95 Set/view switch ISL topology
- 96 Set/view switch zoning
- 97 Set/view system boot parameters (device, flags, etc)
- 98 Set/view vital product data (diary RAM)
- 99 Test (acceptance) post-install/-repair
- 100 View/search system logs (OS, platform, fabric element, etc ...)

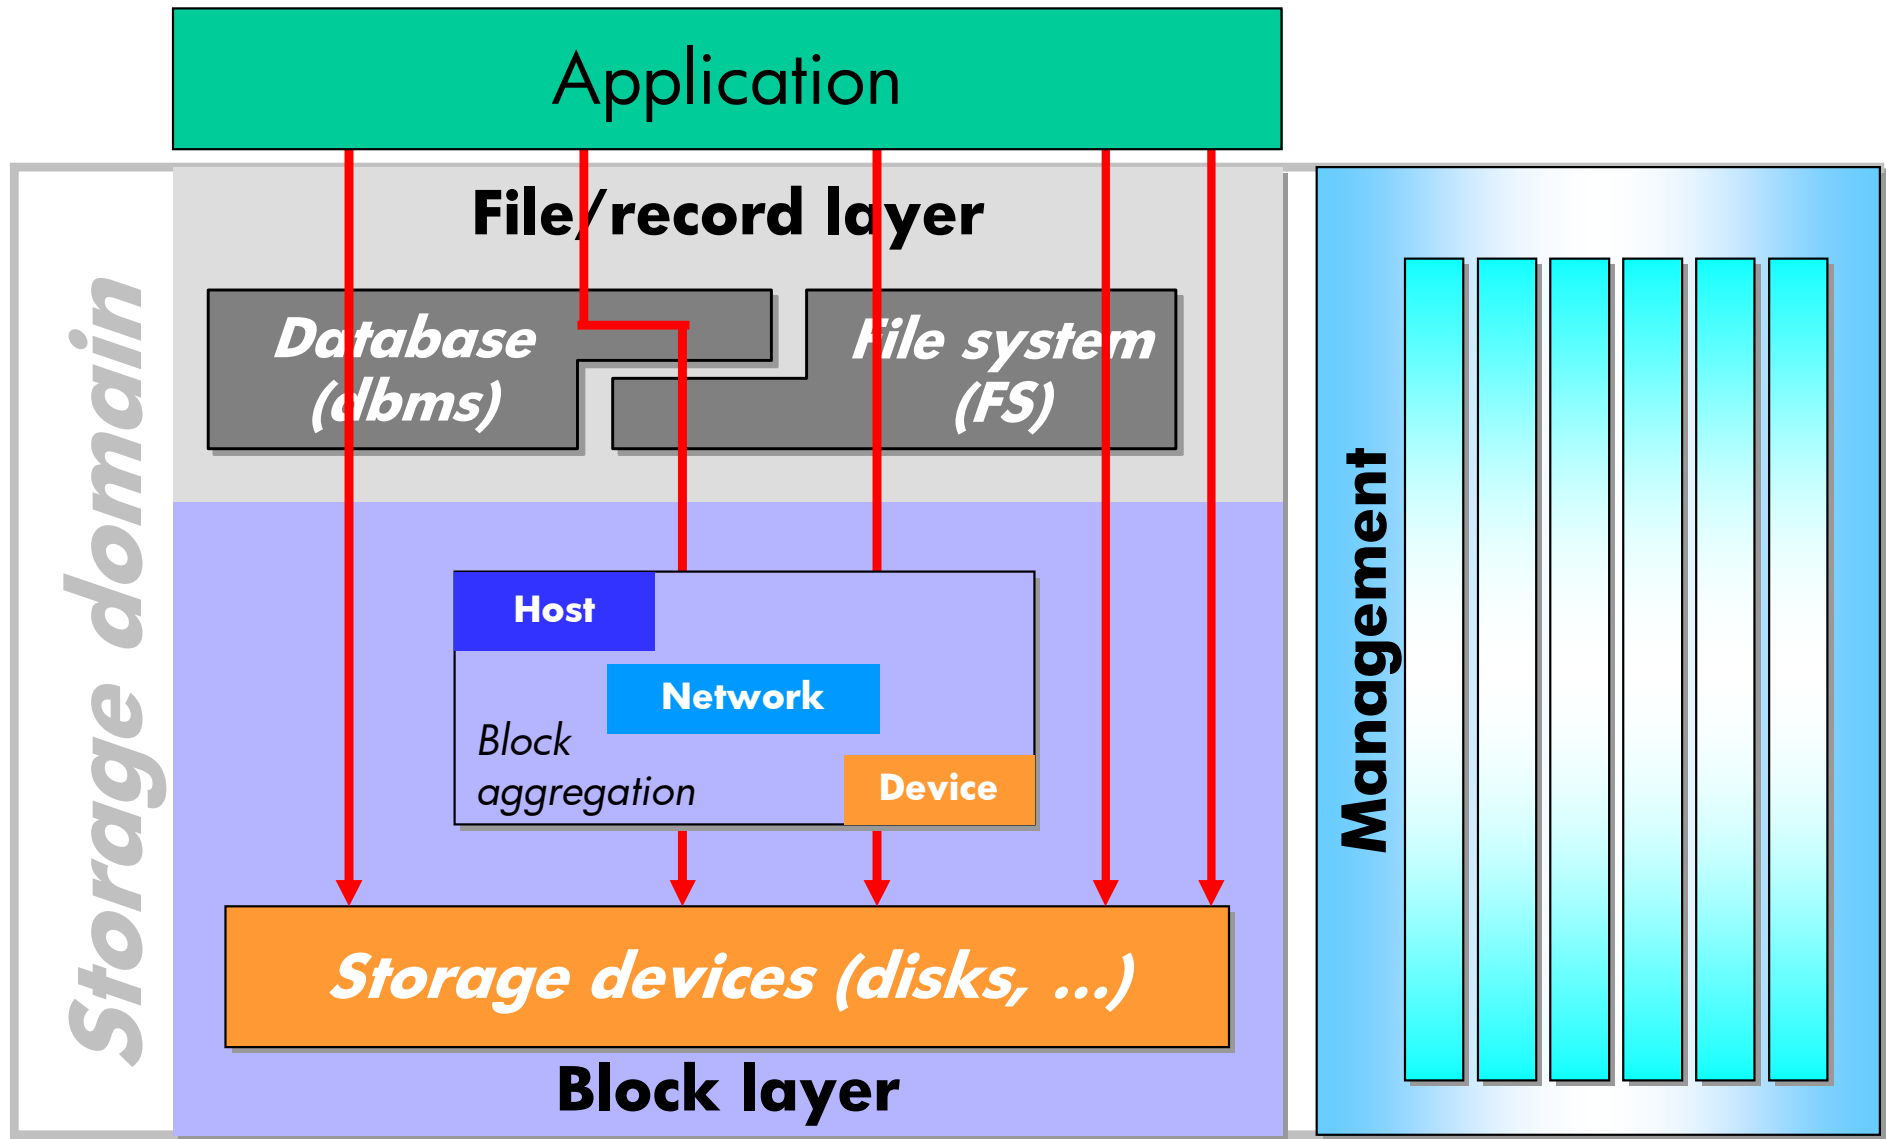
list from  
Stuart Friedberg,  
Veritas







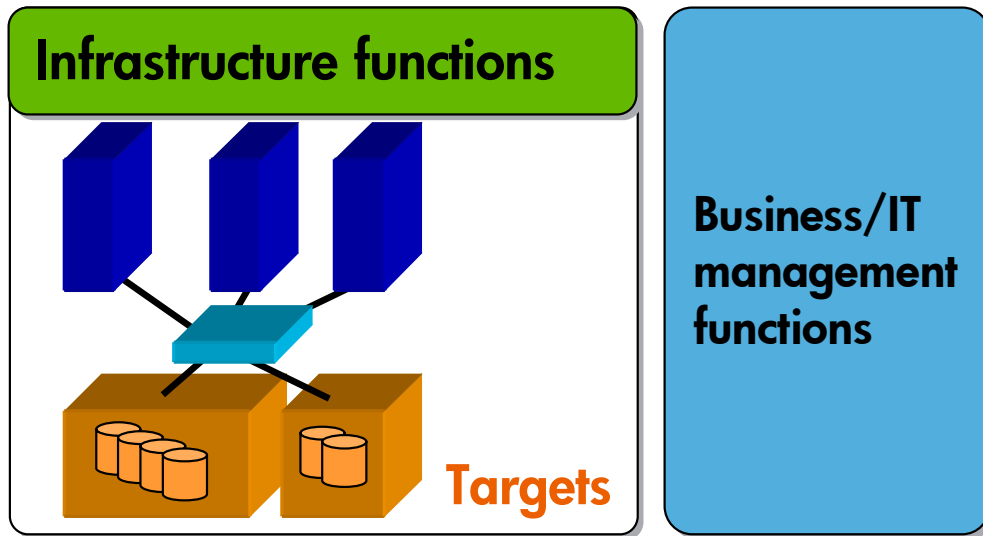
# SNIA Shared Storage Model





Q: what **is** storage management?

A: divide the space



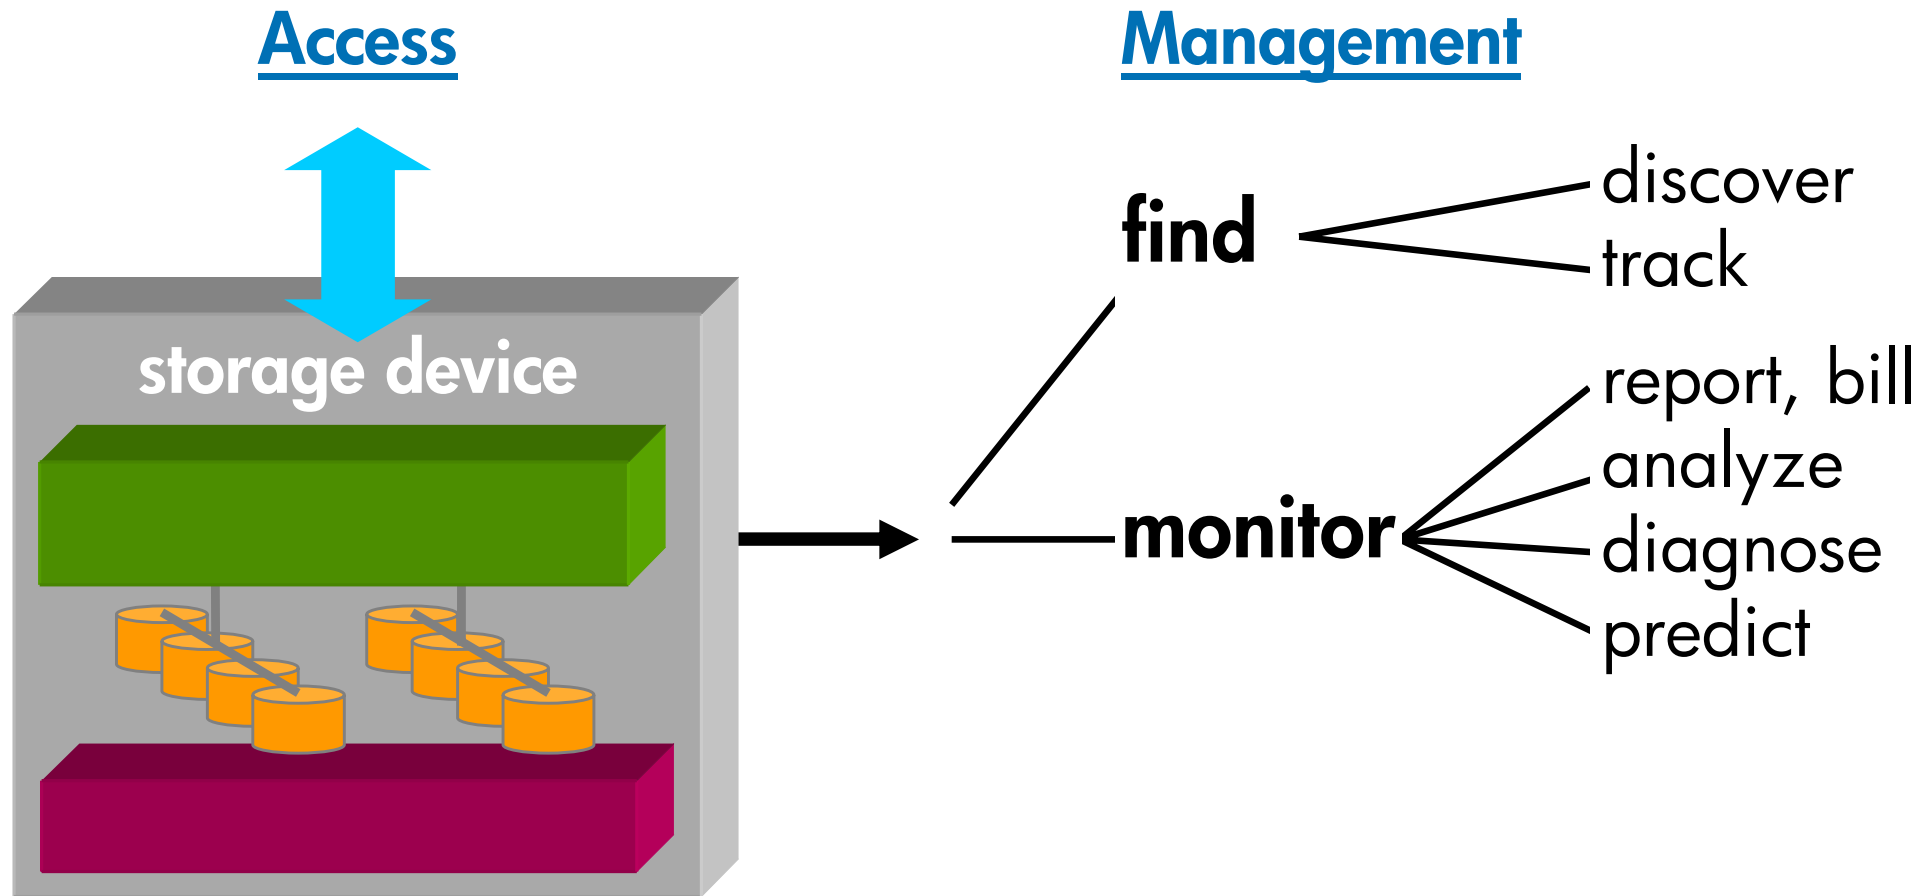
- **targets**
  - managing *things*
- **infrastructure functions**
  - managing *capabilities*
- **business/IT management functions**
  - achieving *goals*





# Target management

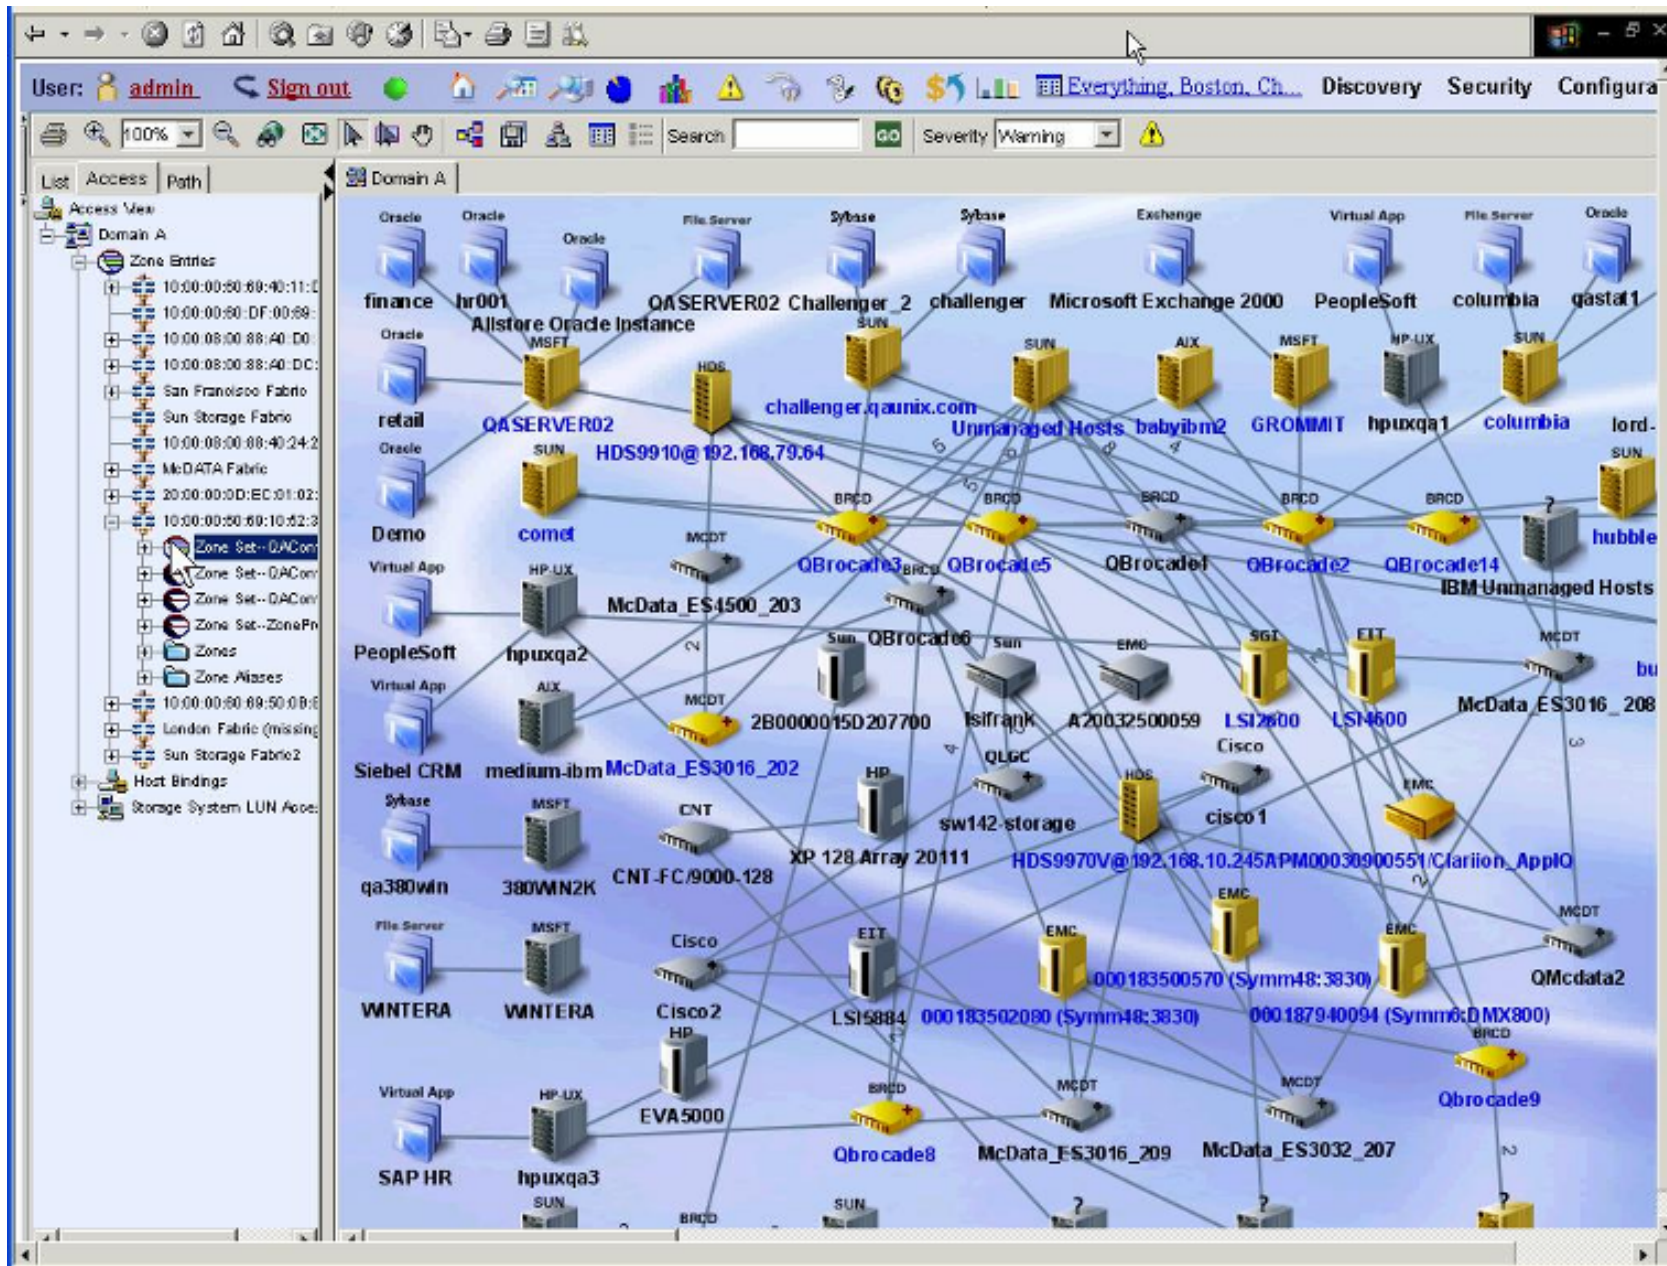
## Access and management





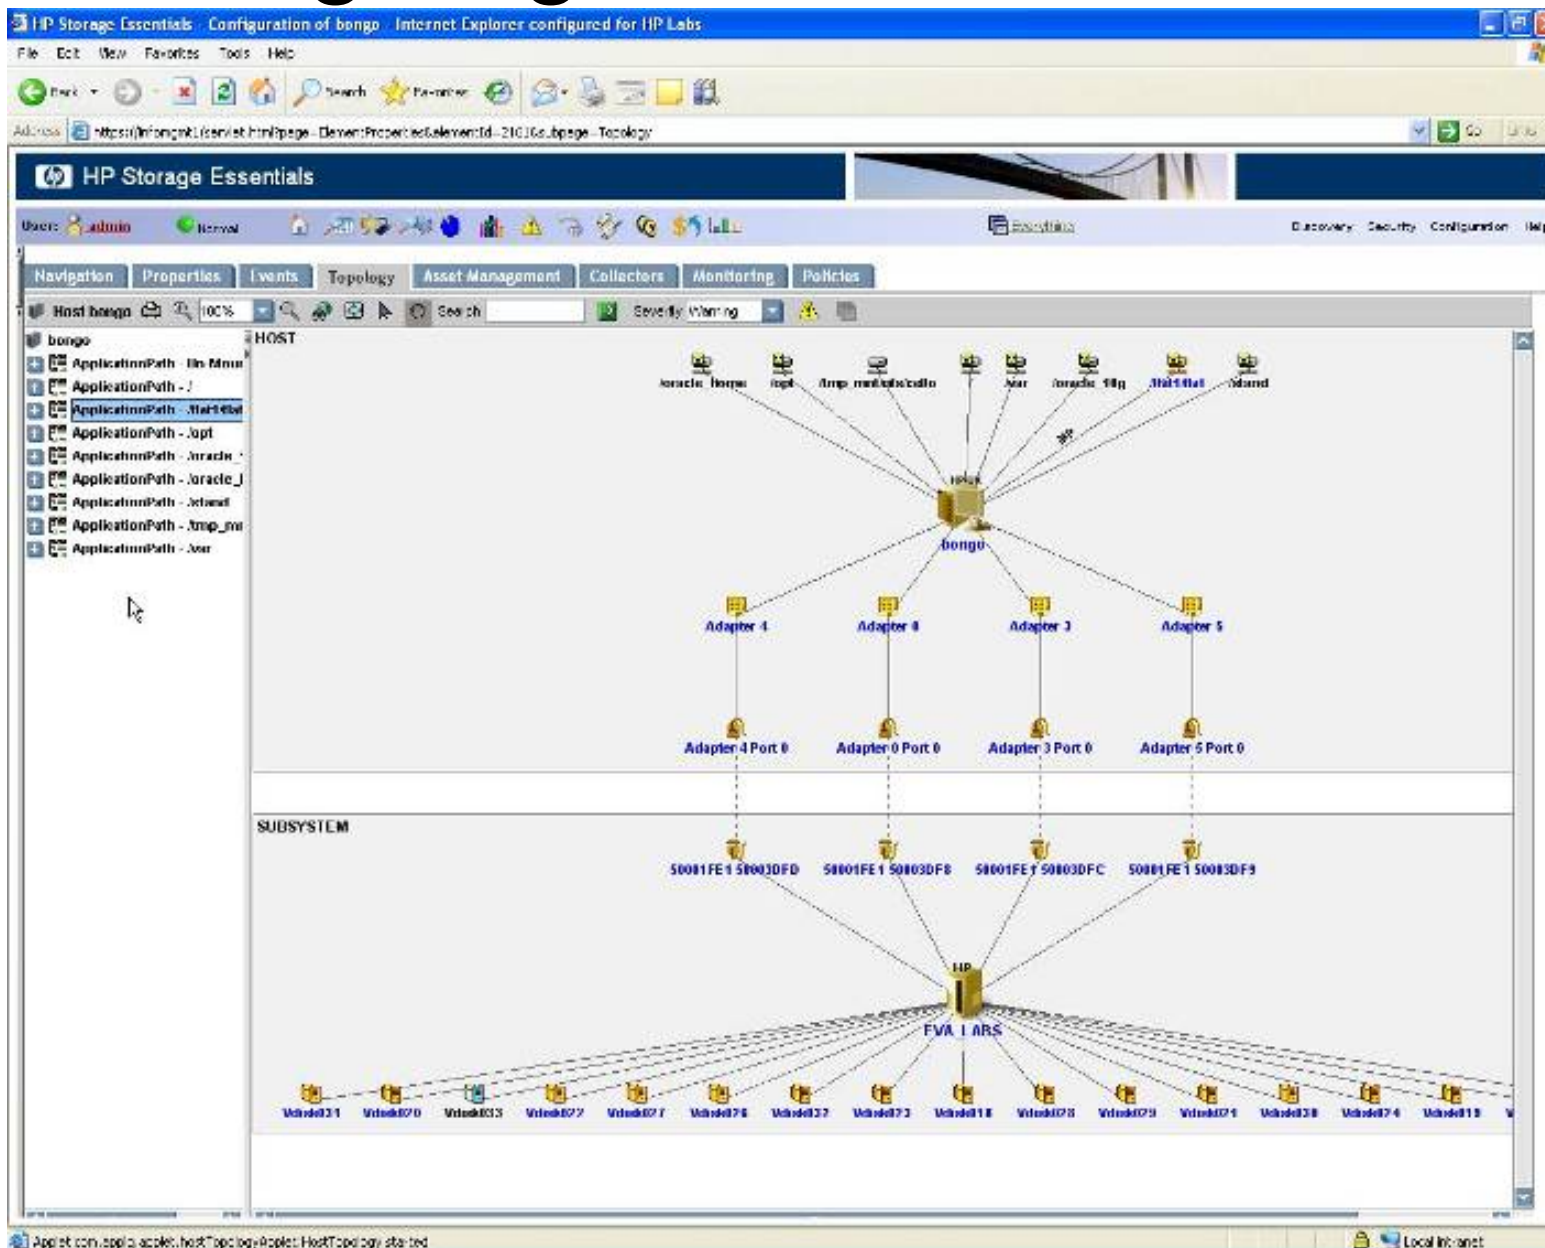
# Target management

## Finding targets



# Target management

## Finding targets

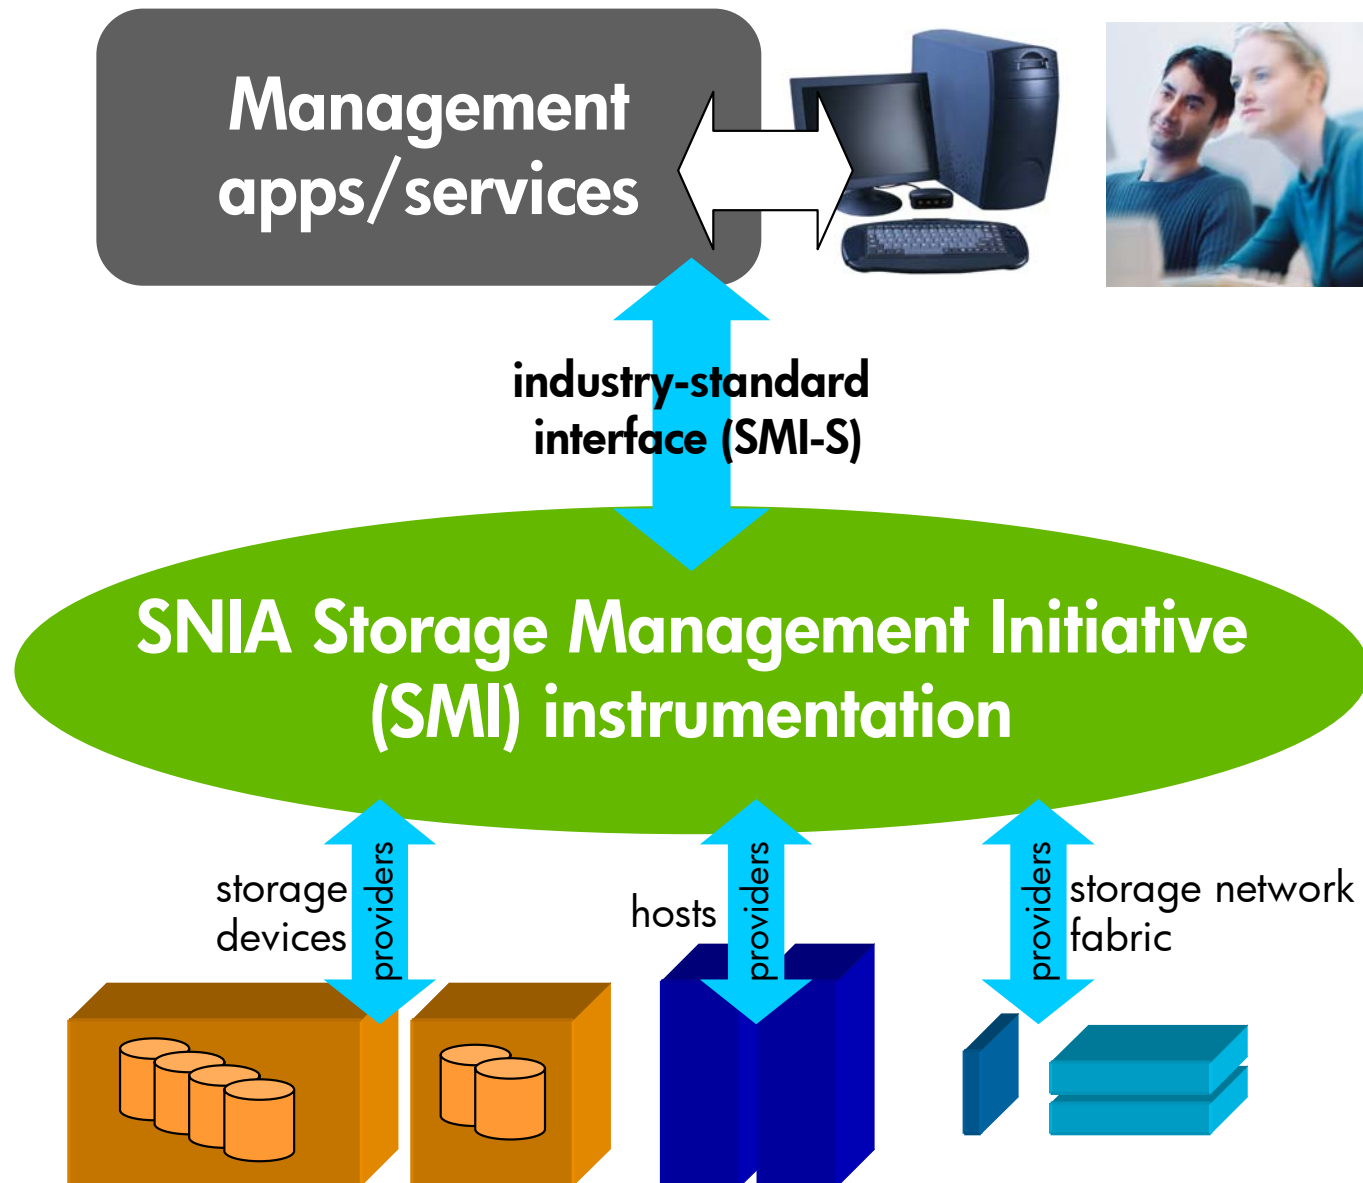






# Target management

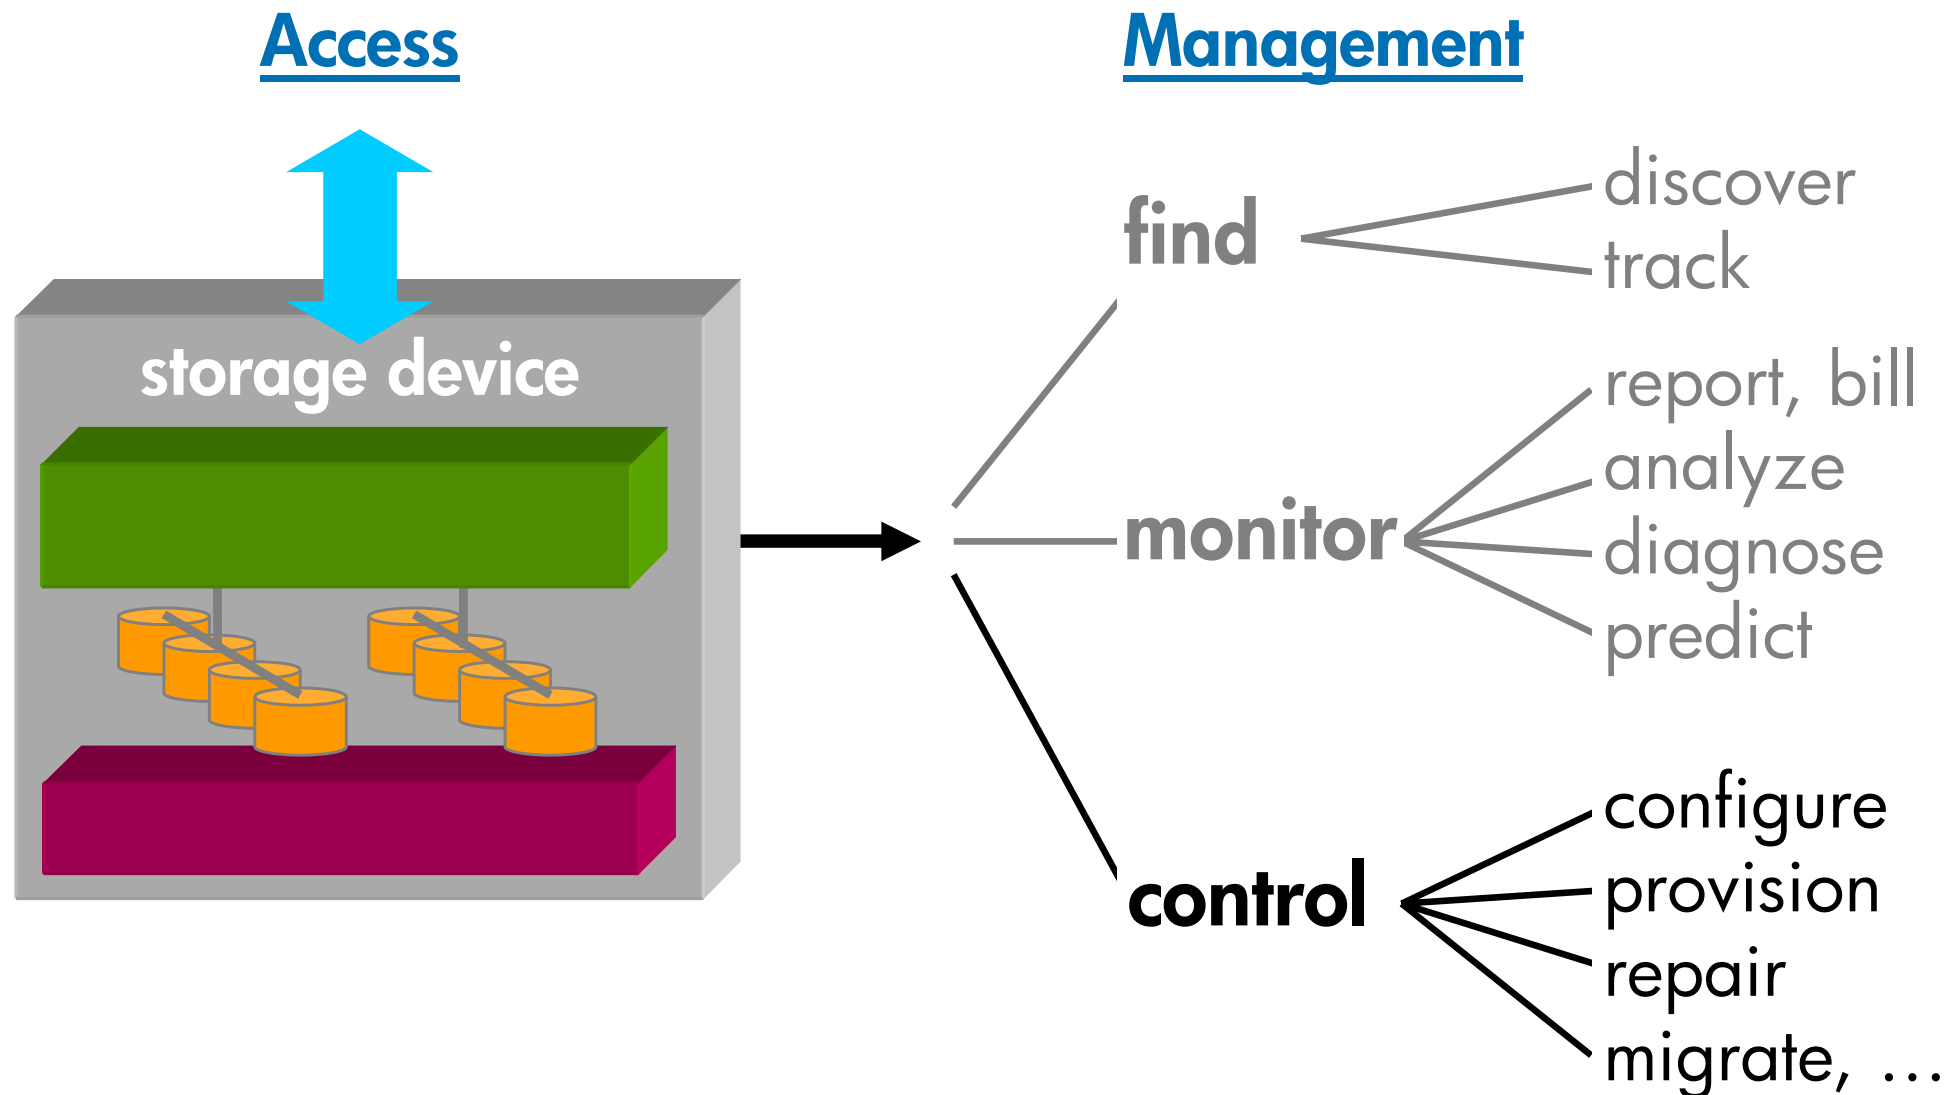
## SMI: a standard monitoring platform





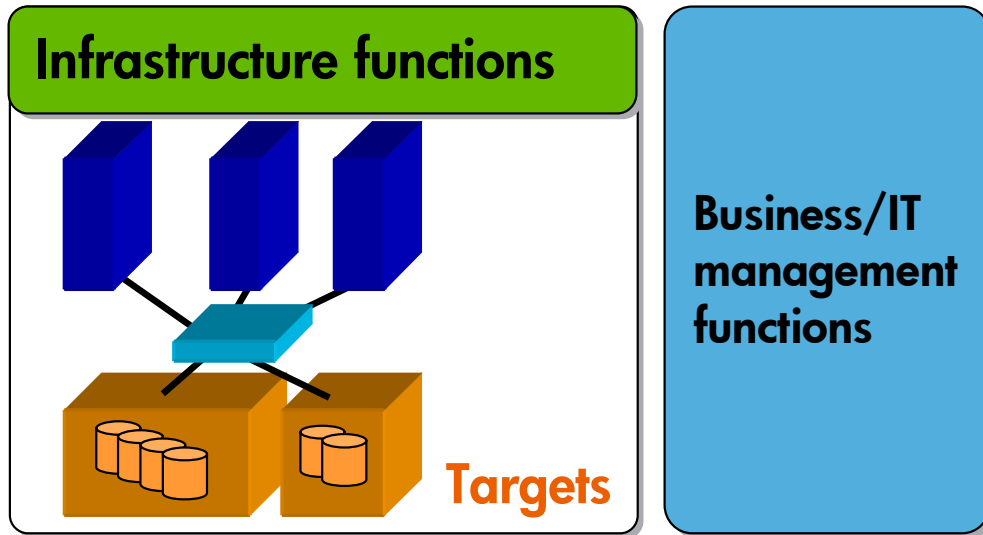
# Target management

## Access and management





Q: what is storage management?  
A: divide the space



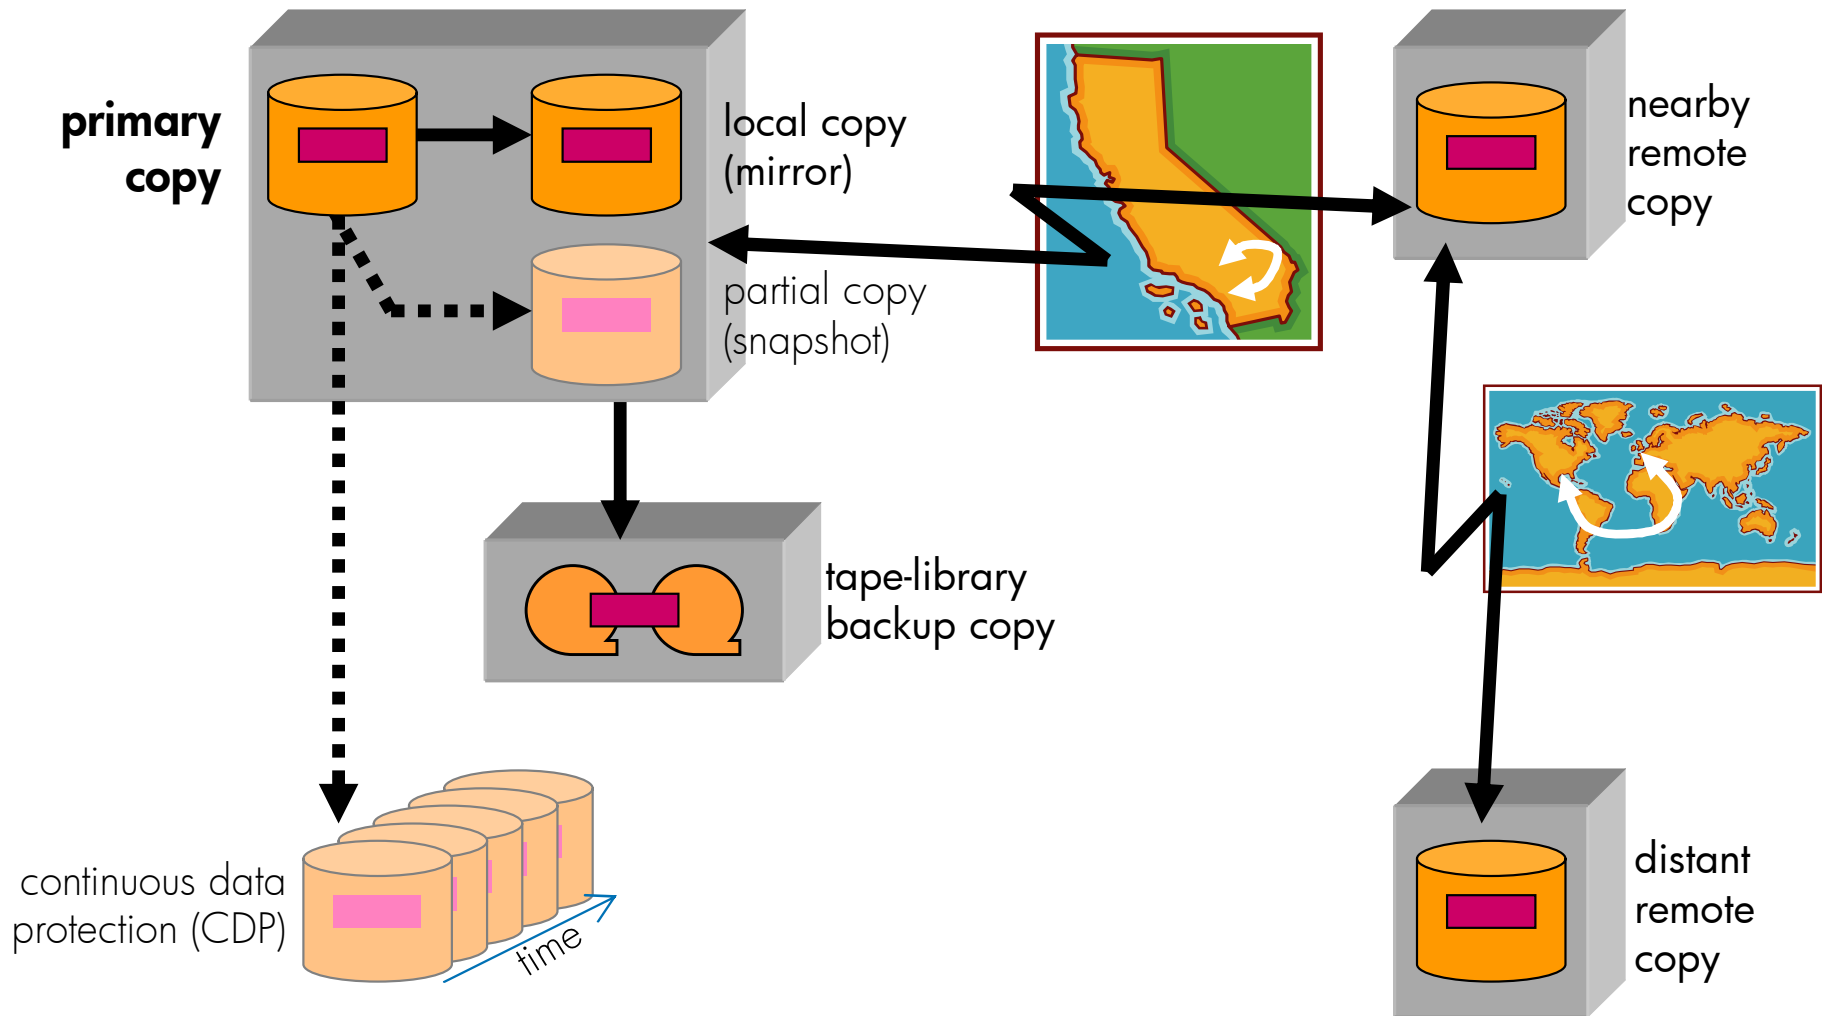
- **targets**
  - managing *things*
- **infrastructure functions**
  - managing *capabilities*
- **business/IT management functions**
  - achieving *goals*





# Infrastructure functions


## Data replication





Q: what is storage management?

A: divide the space



- **targets**
  - managing *things*
- **infrastructure functions**
  - managing *capabilities*

- **business/IT management functions**
  - achieving *goals*



# Business/IT management functions

## Why replicate data?




### Isolate separate copies

- for performance
- for protection

### Keep active data accessible

- **Business continuity** (business always up)
- **Failure/disaster tolerance** (hide faults)
- **Failure/disaster recovery** (pick up after faults)



### Get data close to its users

- e.g., caching

### Migrate data

- while it is being accessed

# Business/IT management functions

## Quality of Service (QoS) objectives

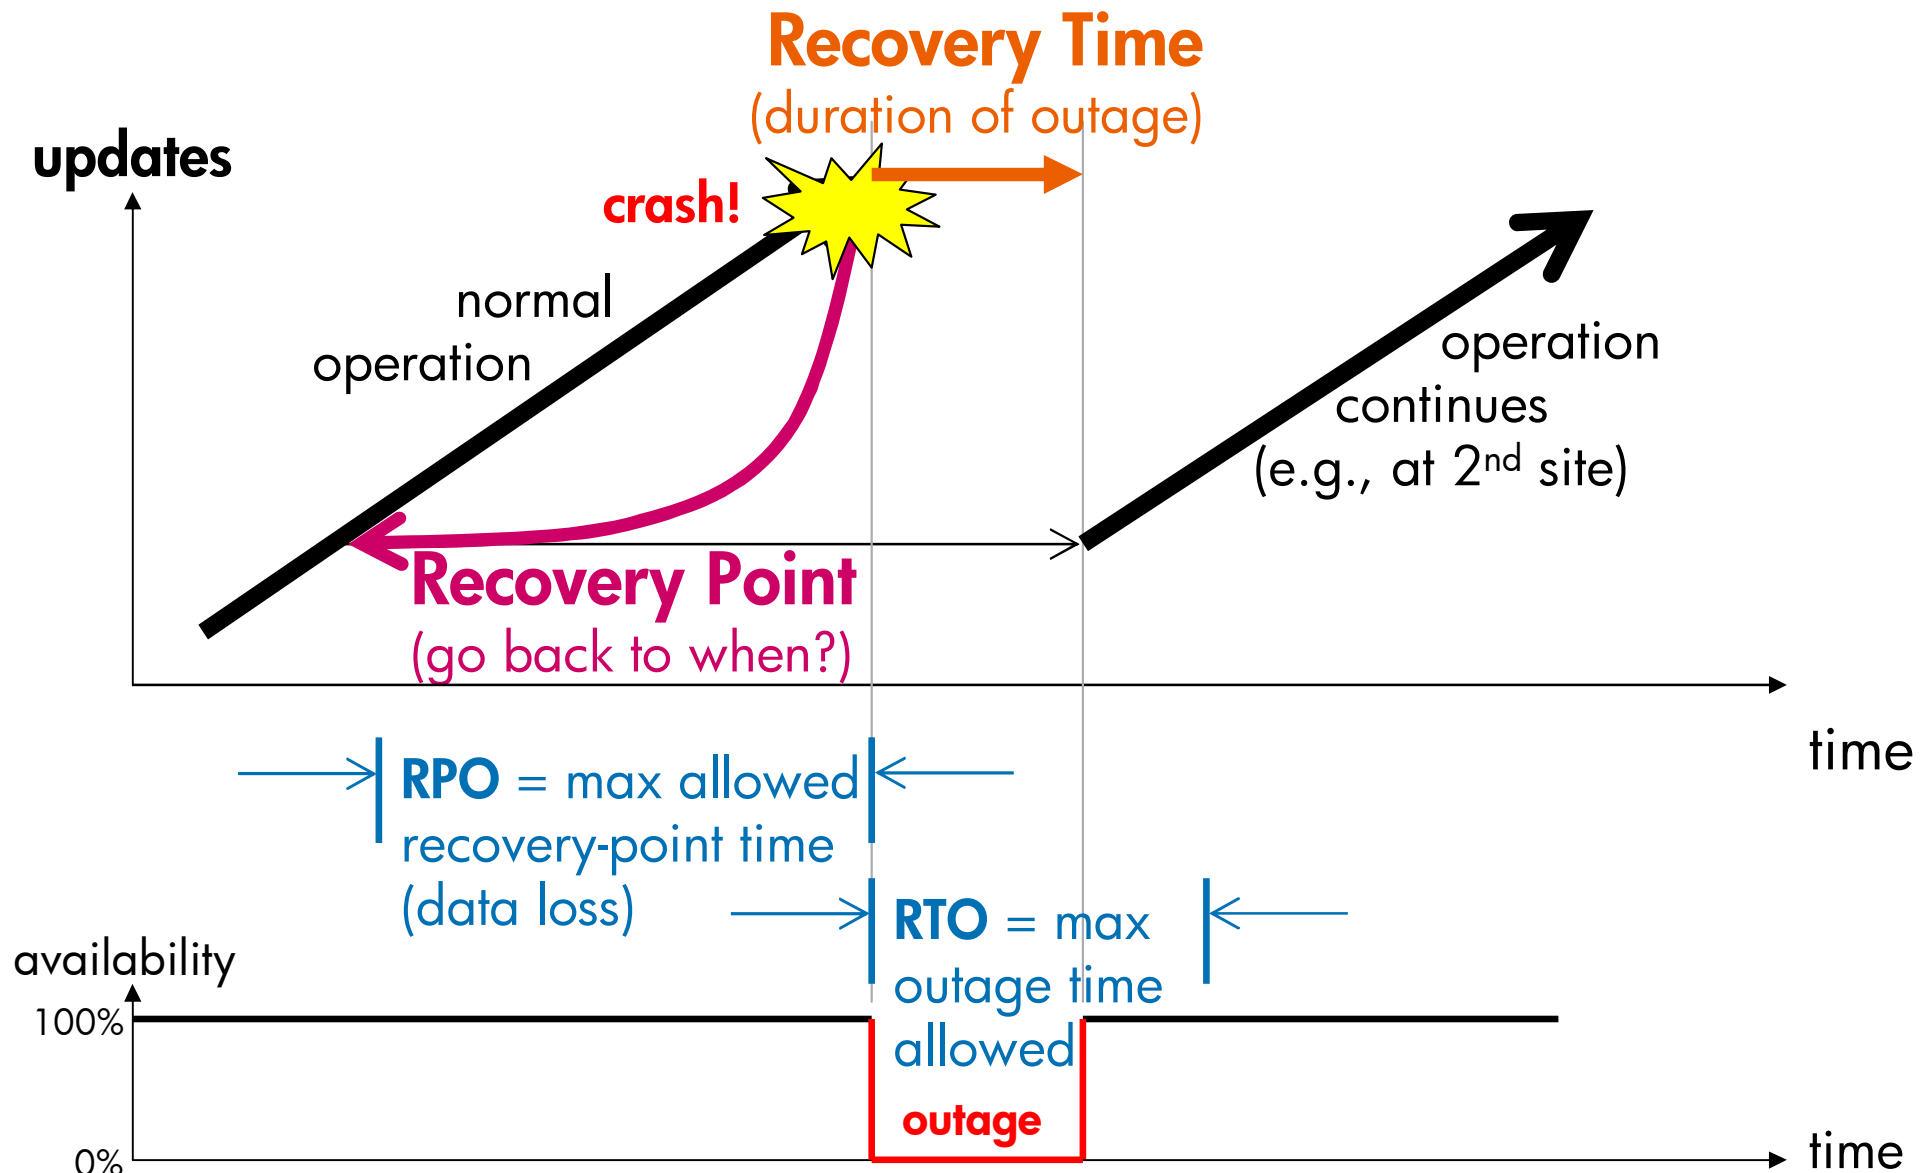


- **QoS goals** = “I want ...”
  - security, performance
  - reliability, availability
  - consistency, correctness
- **QoS measurement** = “what is happening?”
  - monitor, report, analyze, diagnose, predict
- **QoS enforcement** = “make sure that ...”
  - bad stuff doesn’t happen (stop it)
  - good stuff does (adapt)



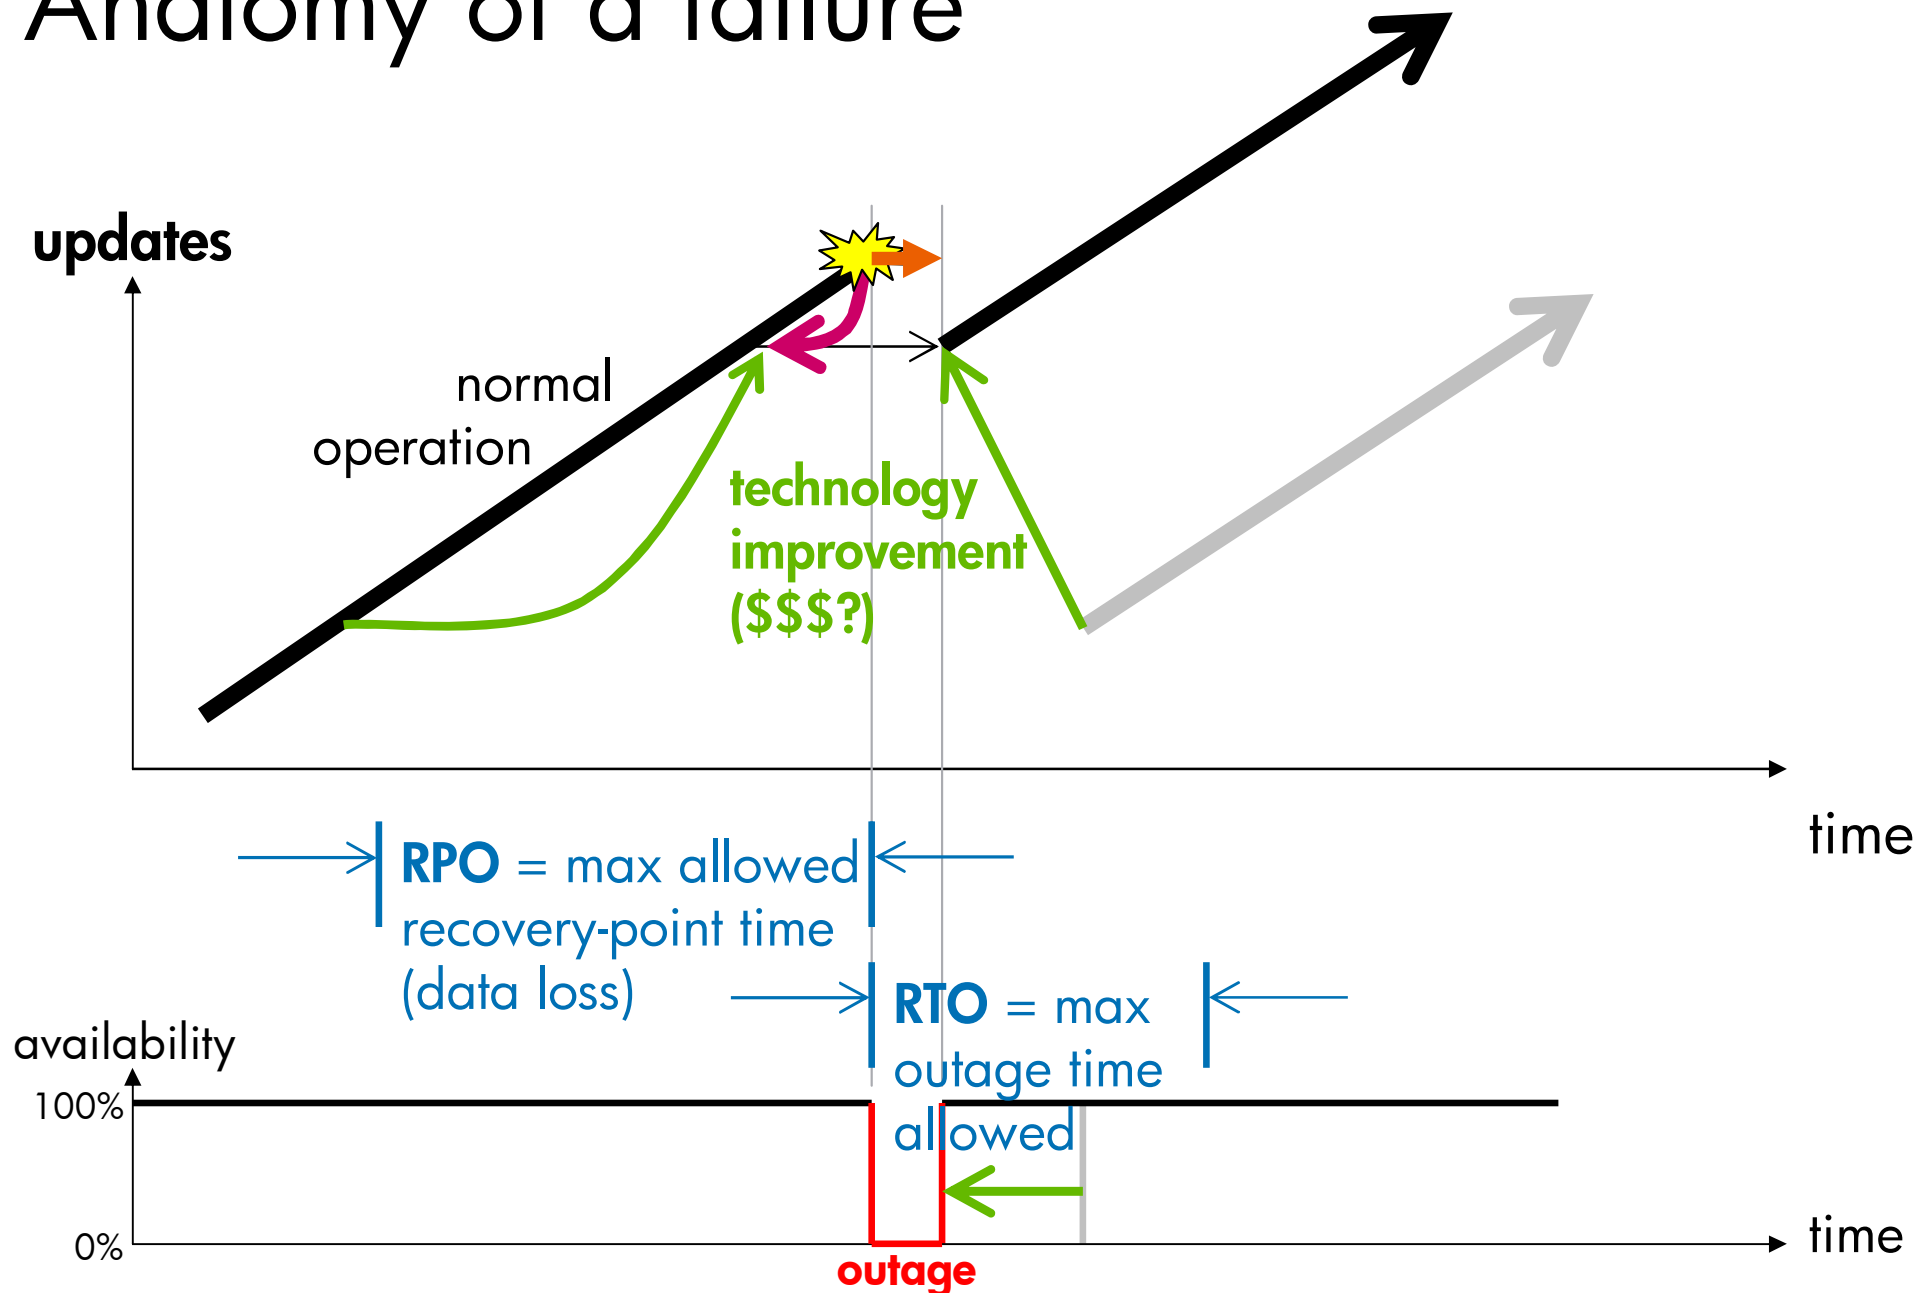
# Management function: protect data

## Anatomy of a failure



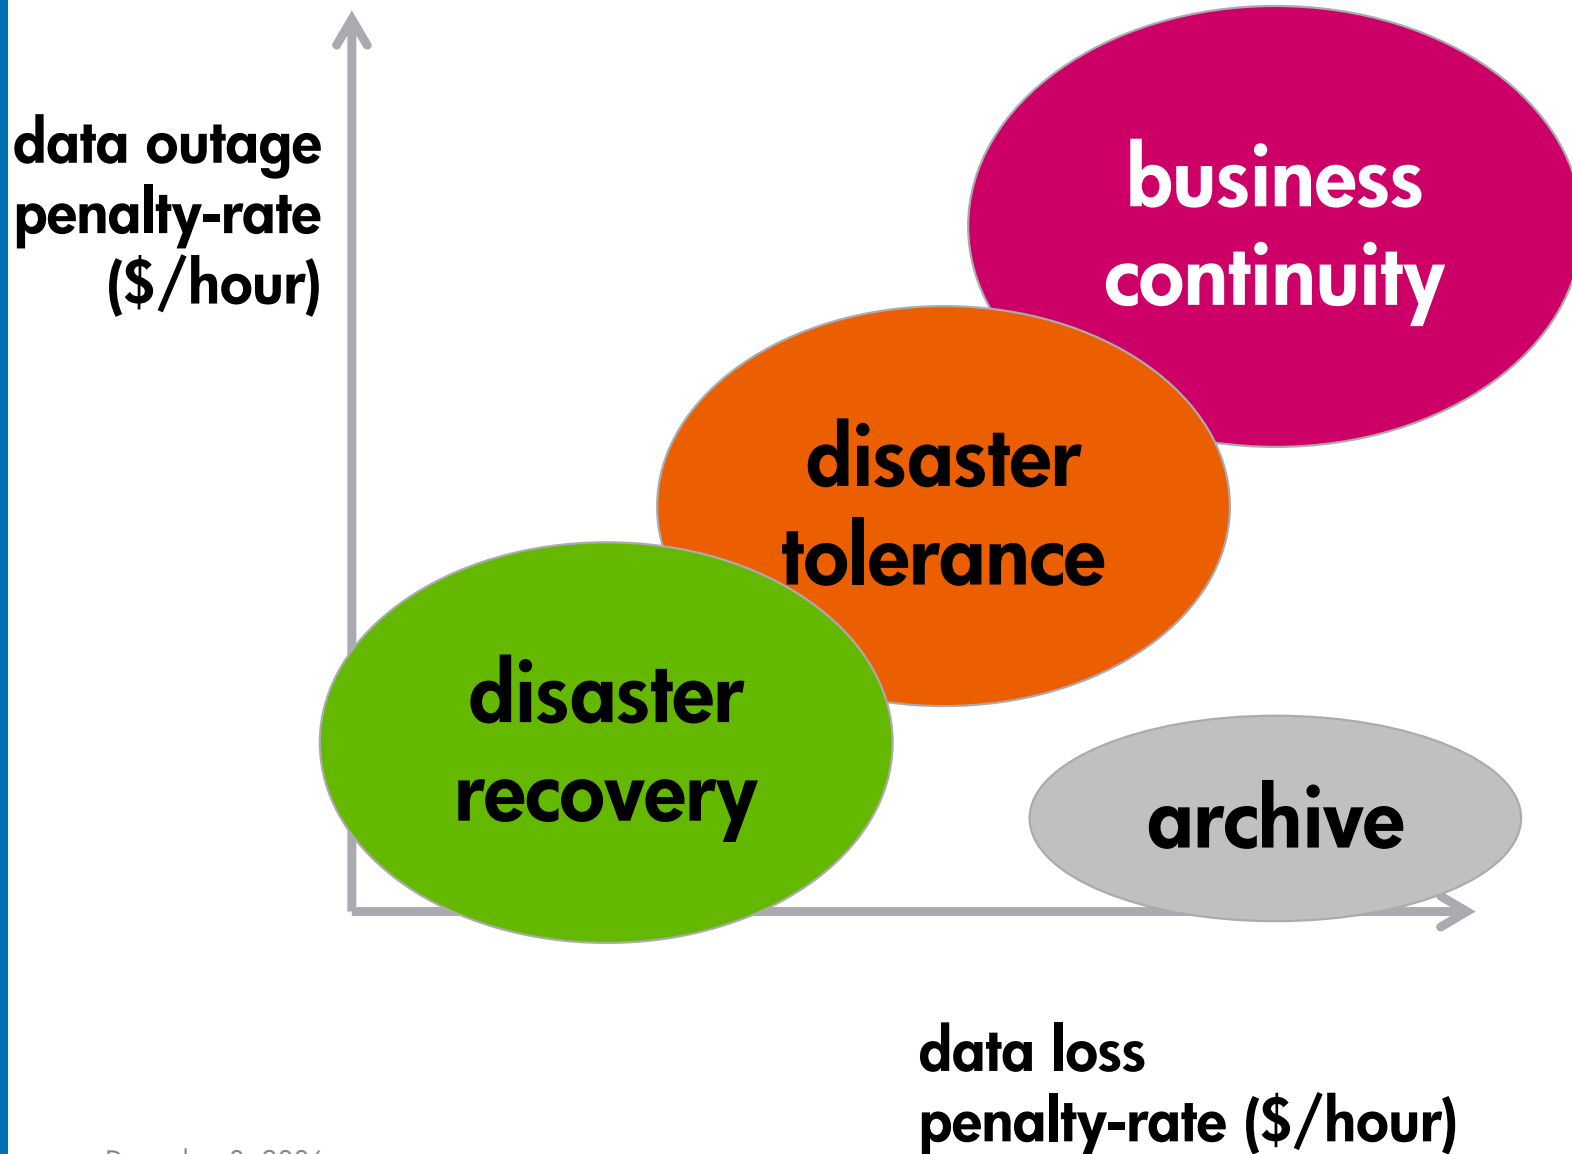
# Management function: protect data

## Anatomy of a failure



# Management function: protect data


## Goals defined by penalties





# Management function: protect data

## Automation simplifies choosing



adapted from *Designing for disasters*,  
K. Keeton et al, FAST'04, March-April 2004


December 8, 2006



# Storage-management challenges

## Administrator costs


- Storage costs are dropping
  - 1995: ~\$5000/GB raw
  - 2005: \$0.5/GB raw
- Administrator costs are not
  - 2004–5 salary: \$68k





# Storage-management challenges

## Add a block-volume to a host



### Touch points

- logical volume manager
- storage-network interface cards
- storage network switches (zones)
- disk array ports (LUNs)
- logical unit (LU)
- physical volume usage

# Storage-management challenges

## Modern tools help ...



HP Storage Essentials - Path Provisioning - Internet Explorer configured for HP Labs

File Edit View Favorites Tools Help

Address <https://infomgmt1/servlet.html?page=AutoProvisionApplet>

**HP Storage Essentials**

User: **admin** Normal

Volume Creation + LUN Security + Zone Operation Refresh

**Provision Configuration**

1. Storage System	2. Host	3. Volume	4. LUN	5. Zone												
<table border="1"><thead><tr><th>S</th><th>Name</th><th>Unmapped</th><th>Raw</th></tr></thead><tbody><tr><td><input type="checkbox"/></td><td>EVA_LABS</td><td>1,445,744 MB</td><td>0 MB</td></tr><tr><td><input checked="" type="checkbox"/></td><td>EVA_SSD1</td><td>152,768 MB</td><td>88,484 MB</td></tr></tbody></table>	S	Name	Unmapped	Raw	<input type="checkbox"/>	EVA_LABS	1,445,744 MB	0 MB	<input checked="" type="checkbox"/>	EVA_SSD1	152,768 MB	88,484 MB	<p>Hosts</p> <ul style="list-style-type: none"><li>bongo</li><li>Host_2677 (discovered)</li><li>Host_2680 (discovered)</li><li>Host_2651 (discovered)</li><li>INFOMGMT1</li></ul>	<p>Volumes</p> <ul style="list-style-type: none"><li>Mapped</li><li>Unmapped</li><li>Free Extents</li></ul>	<p>LUNs</p> <ul style="list-style-type: none"><li>50001FE1500221E8 Total LUN=8 WWN=50</li><li>50001FE1500221E9 Total LUN=8 WWN=50</li><li>50001FE1500221EC Total LUN=8 WWN=50</li><li>50001FE1500221ED Total LUN=8 WWN=50</li></ul>	<p>Zones</p> <ul style="list-style-type: none"><li>10:00:00:60:68:22:3C:76</li></ul>
S	Name	Unmapped	Raw													
<input type="checkbox"/>	EVA_LABS	1,445,744 MB	0 MB													
<input checked="" type="checkbox"/>	EVA_SSD1	152,768 MB	88,484 MB													
Selected Storage System: EVA_SSD1	Host data loaded.	Volume data loaded.	Host Security Group data loaded.	Zone data loaded.												
<b>Step 1</b>	<b>Step 2</b>	<b>Step 3</b>	<b>Step 4</b>	<b>Step 5</b>												

Provision Jobs

M	ID	Description	State	Schedule
<input type="checkbox"/>				
<input type="checkbox"/>				

JobConsole MessageConsole

Path 1 Create Job

An example Storage Essentials provisioning screen

# Storage-management challenges

## ... but finicky details still remain



HP Storage Essentials - Disk Drives for bongo - Internet Explorer configured for HP Labs

Address: https://hpstorage.com/.../NavigationItem=DiskDrives

HP Storage Essentials

Navigation: Properties Events Topology Asset Management Collectors Monitoring Policies

Host bongo

HOST		Dependencies	
Host Bus Adapters	Volume Management	Storage Volumes	Applications
HBA Paths	Multipathing	Disk Partitions	Switches
Bindings		Disk Drives	Storage Systems
			NAS Hosts

<a href="#">/dev/dsk/c47t13d0</a>	47	0 (0x0)	<a href="#">/dev/vq_swap:c47t13d0/vol1</a>
<a href="#">/dev/dsk/c47t14d0</a>	47	0 (0x0)	<a href="#">/dev/vq_opt:c47t14d0/vol1</a>
<a href="#">/dev/dsk/c47t15d0</a>	47	0 (0x0)	<a href="#">/dev/vq_oraclehome:c47t15d0/vol1</a>

Disk Drives

Name	SCSI Bus	OS LUN (hex)	Disk Partitions
.../c47t13d0	100	D(0x0)	c100t14d0
.../c47t14d0	1	D(0x0)	/dev/vq_opt:c47t14d0/vol1, /dev/vq_swap:c47t14d0/vol1, /dev/vq_oraclehome:c47t14d0/vol1
.../c47t15d0	2	D(0x0)	/dev/vq_opt:c47t15d0/vol1, /dev/vq_swap:c47t15d0/vol1, /dev/vq_oraclehome:c47t15d0/vol1
.../c47t16d0	44	1(0x1)	/dev/vq_opt:c47t16d0/vol1, /dev/vq_swap:c47t16d0/vol1, /dev/vq_oraclehome:c47t16d0/vol1
.../c47t17d0	44	2(0x2)	/dev/vq_opt:c47t17d0/vol1, /dev/vq_swap:c47t17d0/vol1, /dev/vq_oraclehome:c47t17d0/vol1
.../c47t18d0	44	3(0x3)	/dev/vq_opt:c47t18d0/vol1, /dev/vq_swap:c47t18d0/vol1, /dev/vq_oraclehome:c47t18d0/vol1
.../c47t19d0	44	4(0x4)	/dev/vq_opt:c47t19d0/vol1, /dev/vq_swap:c47t19d0/vol1, /dev/vq_oraclehome:c47t19d0/vol1
.../c47t20d0	47	D(0x0)	/dev/vq_opt:c47t20d0/vol1, /dev/vq_swap:c47t20d0/vol1, /dev/vq_oraclehome:c47t20d0/vol1
.../c47t21d0	47	D(0x0)	/dev/vq_opt:c47t21d0/vol1, /dev/vq_swap:c47t21d0/vol1, /dev/vq_oraclehome:c47t21d0/vol1
.../c47t22d0	47	D(0x0)	/dev/vq_opt:c47t22d0/vol1, /dev/vq_swap:c47t22d0/vol1, /dev/vq_oraclehome:c47t22d0/vol1
.../c47t23d0	49	1(0x1)	/dev/vq_opt:c47t23d0/vol1, /dev/vq_swap:c47t23d0/vol1, /dev/vq_oraclehome:c47t23d0/vol1
.../c47t24d0	49	2(0x2)	/dev/vq_opt:c47t24d0/vol1, /dev/vq_swap:c47t24d0/vol1, /dev/vq_oraclehome:c47t24d0/vol1
.../c47t25d0	49	3(0x3)	/dev/vq_opt:c47t25d0/vol1, /dev/vq_swap:c47t25d0/vol1, /dev/vq_oraclehome:c47t25d0/vol1
.../c47t26d0	49	4(0x4)	/dev/vq_opt:c47t26d0/vol1, /dev/vq_swap:c47t26d0/vol1, /dev/vq_oraclehome:c47t26d0/vol1
.../c47t27d0	49	5(0x5)	c29d0
.../c47t28d0	50	1(0x1)	/dev/vq_opt:c47t28d0/vol1, /dev/vq_swap:c47t28d0/vol1, /dev/vq_oraclehome:c47t28d0/vol1

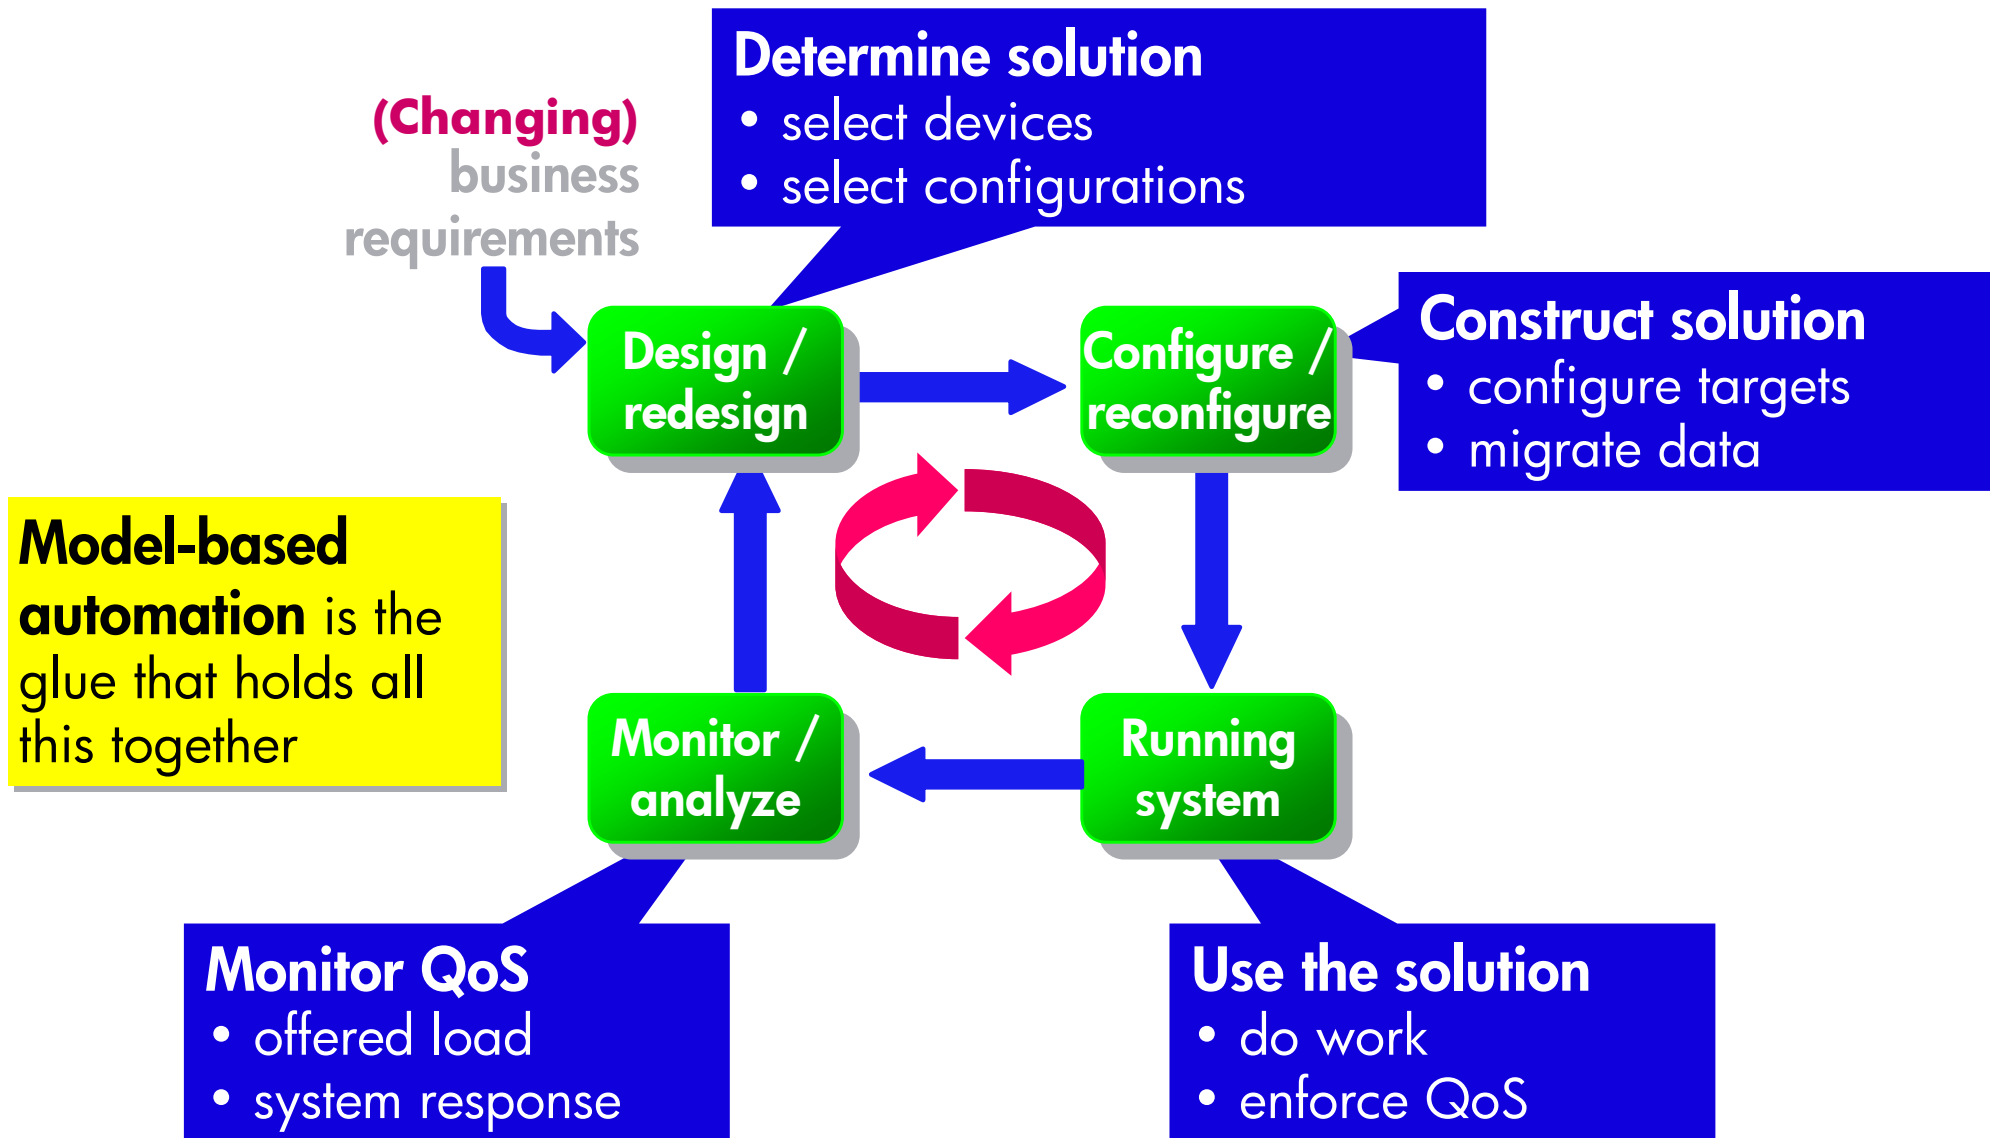
Transpose one digit, and you wipe out the Oracle database software!






# Storage-management challenges

## Solution: automation







# Storage infrastructure management: **summary**



- storage management appears complicated, because it **is complicated**
- monitoring is pretty much under control; **control** is next
  - example: data replication
- **automation** is the key to making IT business-effective
  - *virtualization* makes this simpler

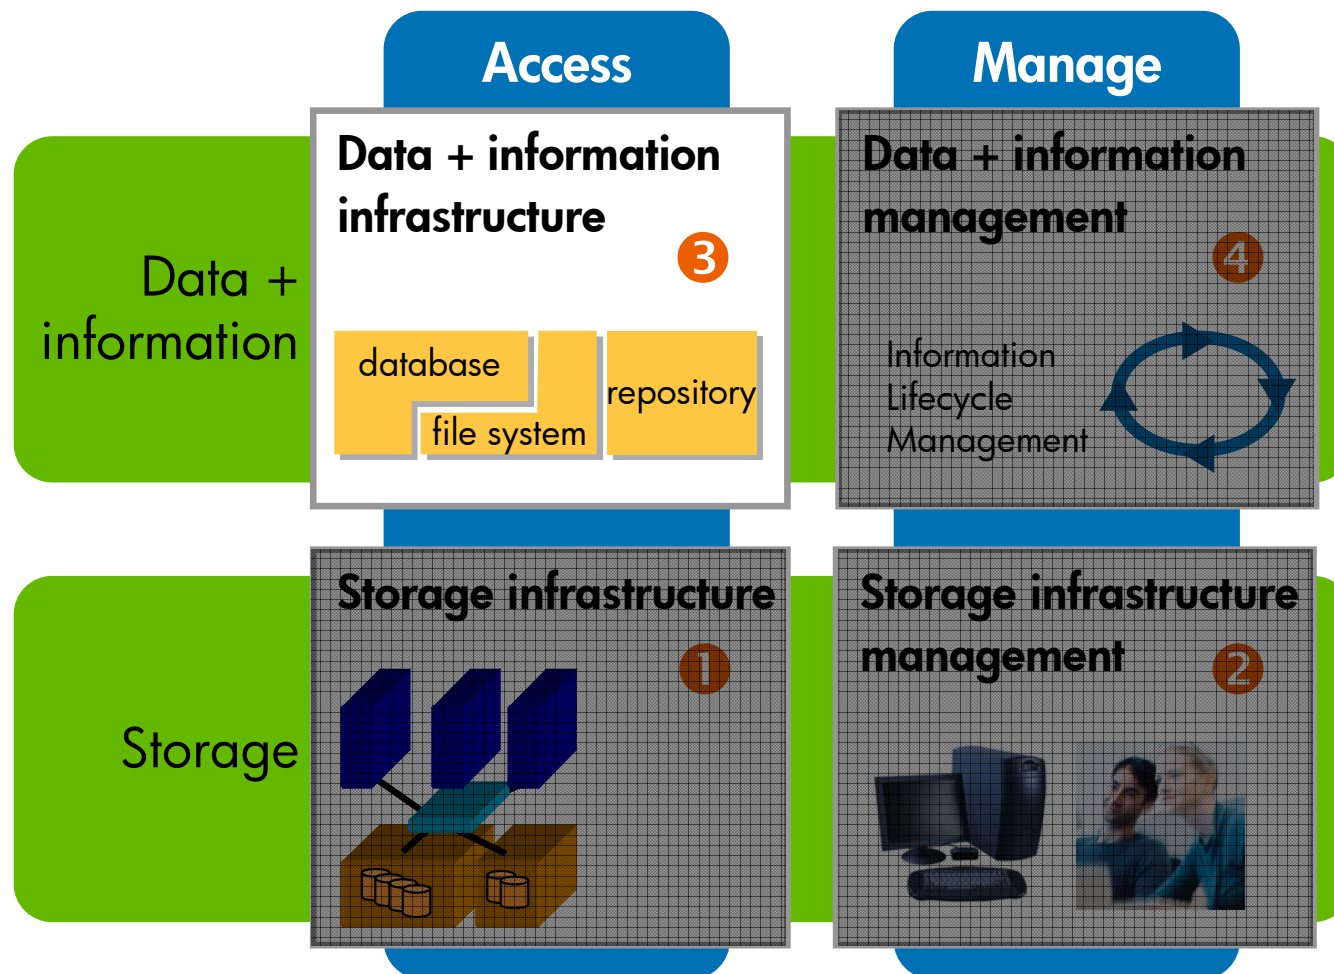
# Data/info infrastructure



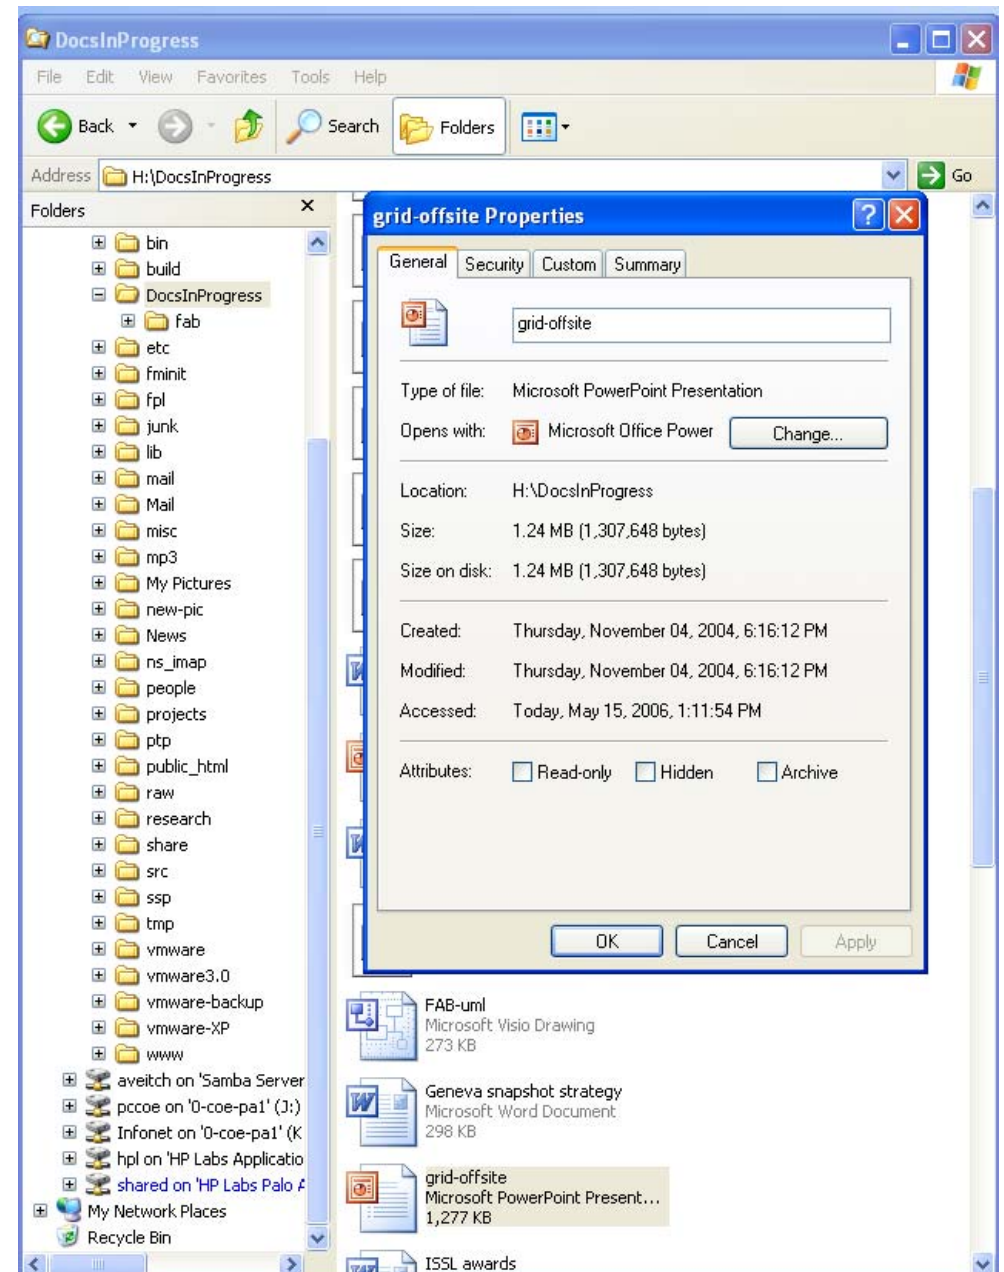
Enterprises  
and people



Business logic, processes & applications

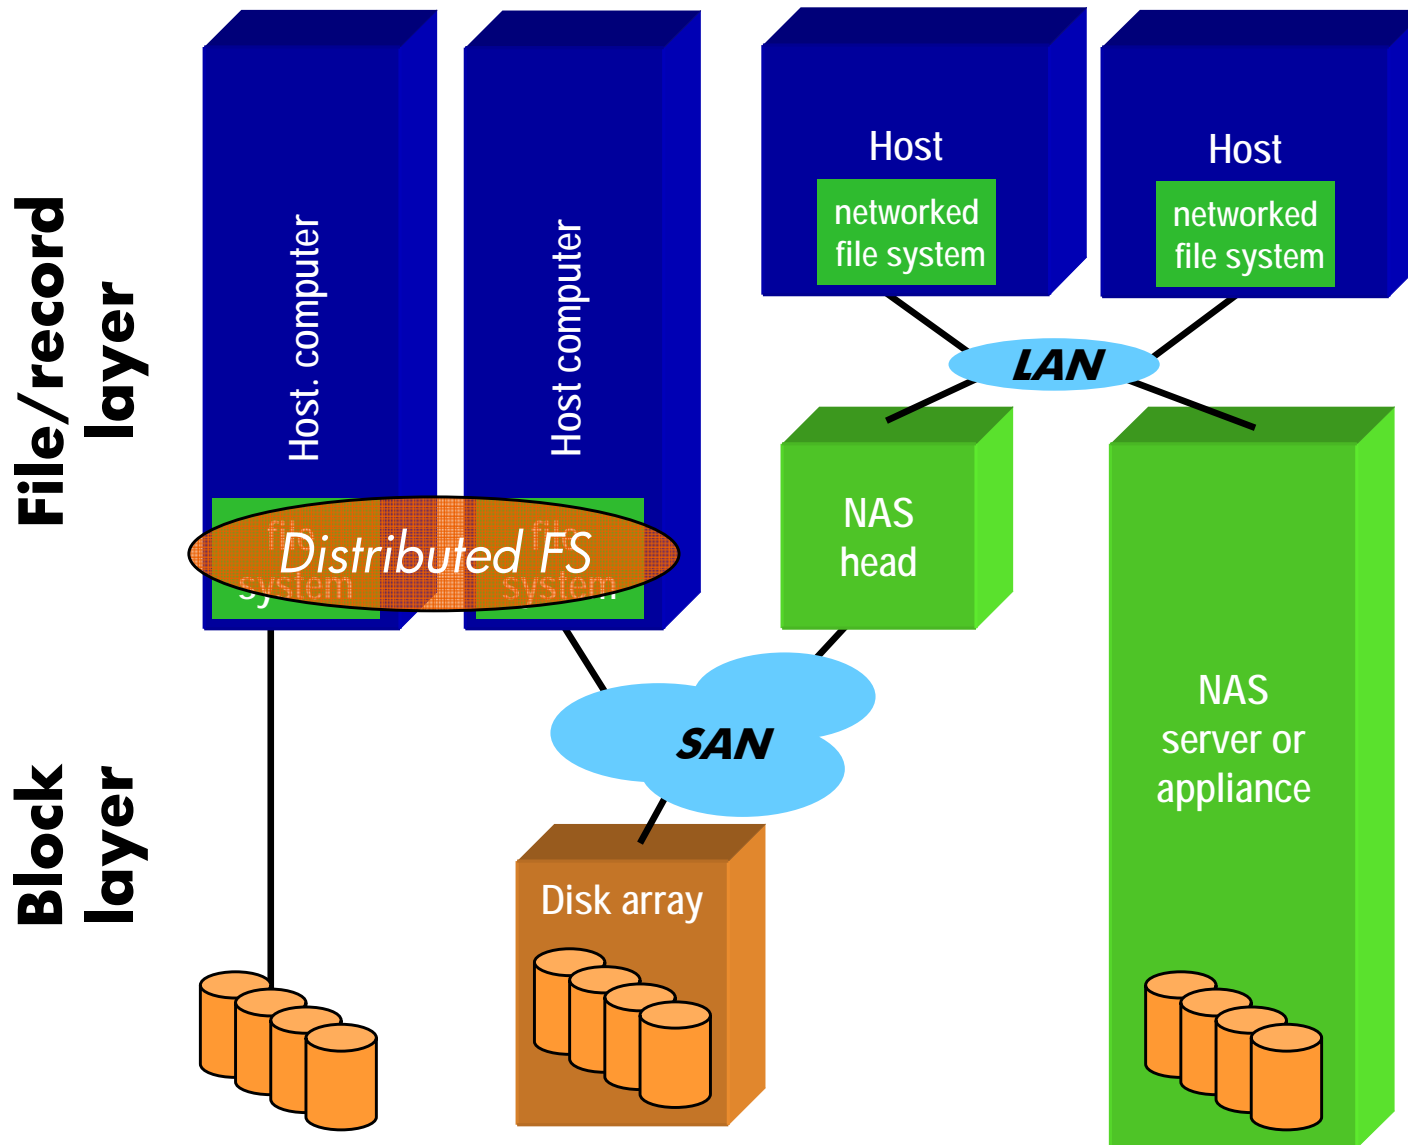


# File systems





# File systems and NAS



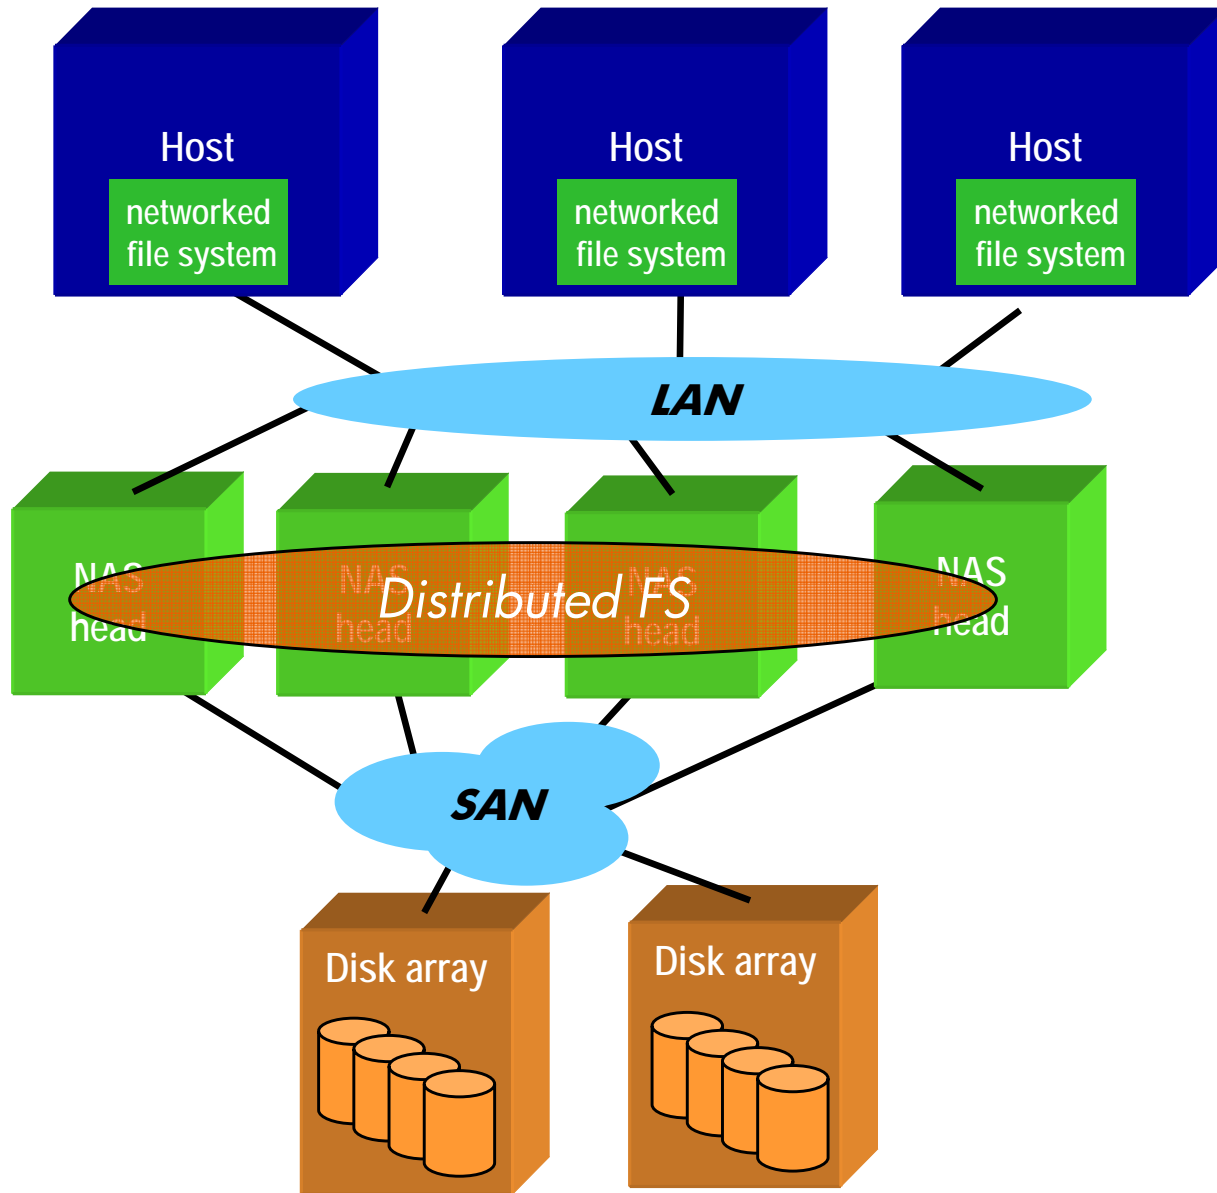
SAN  
(Storage Area Network)

LAN  
(Local Area Network)

NAS  
(Networked Attach Storage)

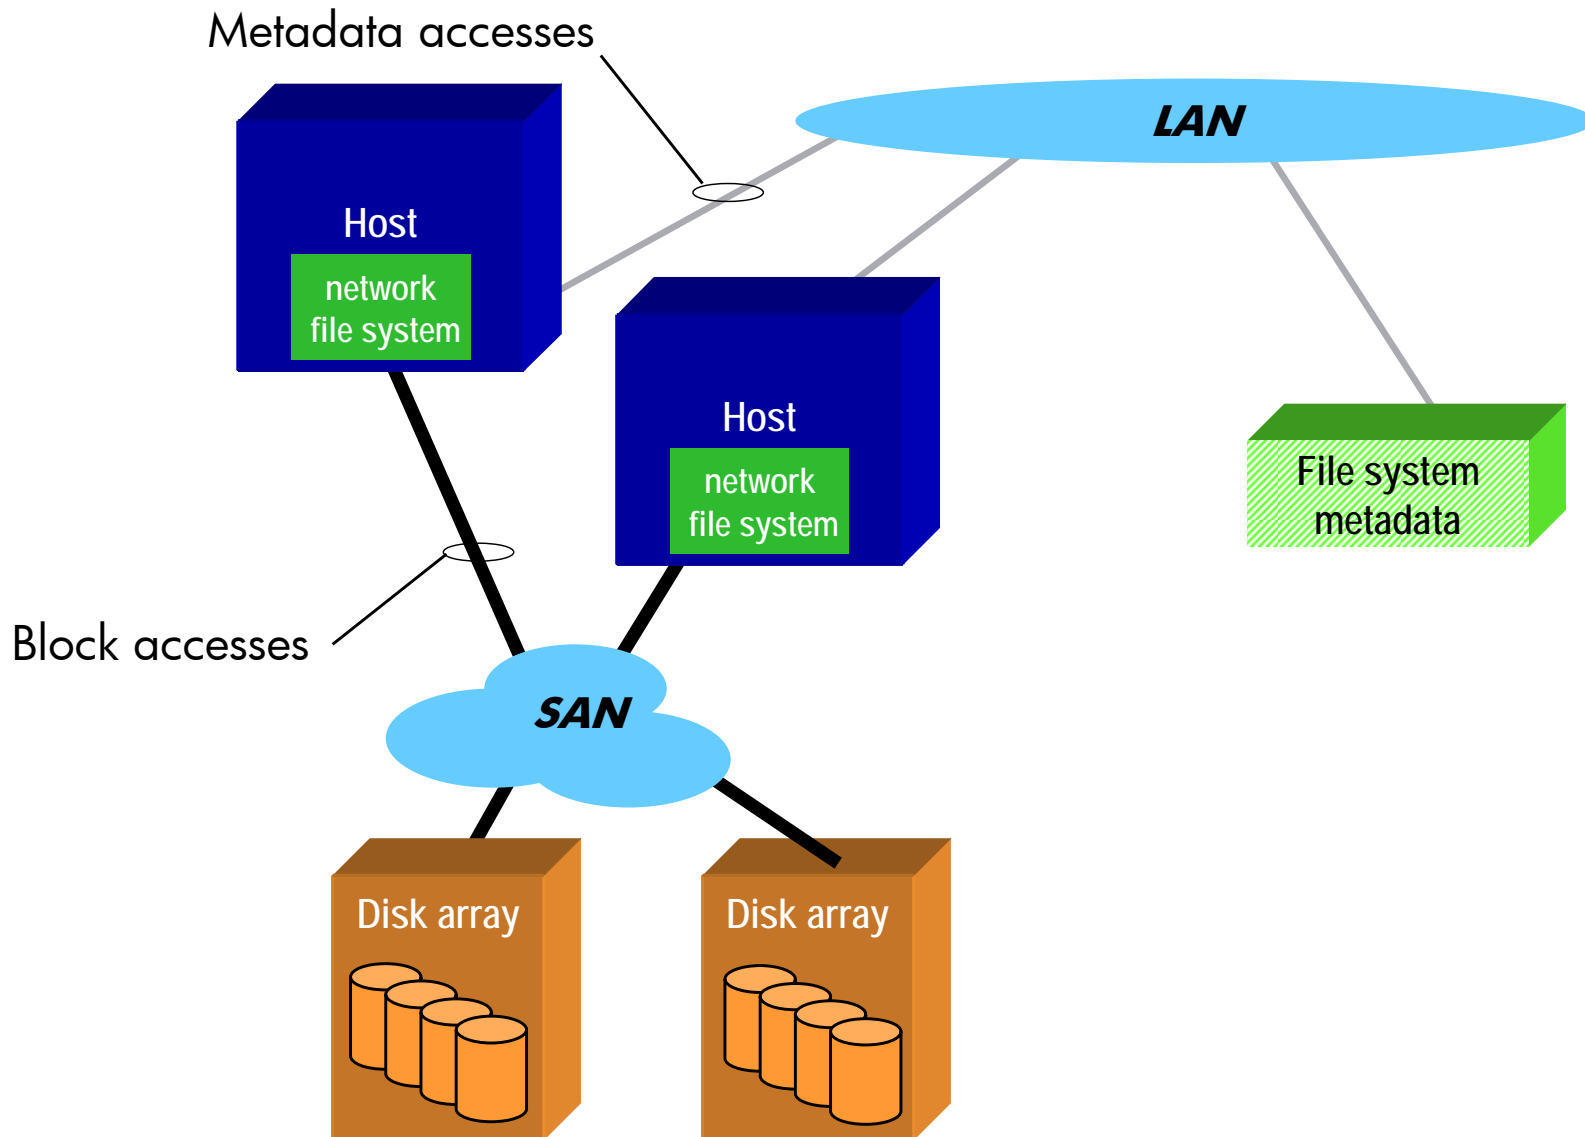


# Clustered NAS (1)





# Clustered NAS (2)





# Databases

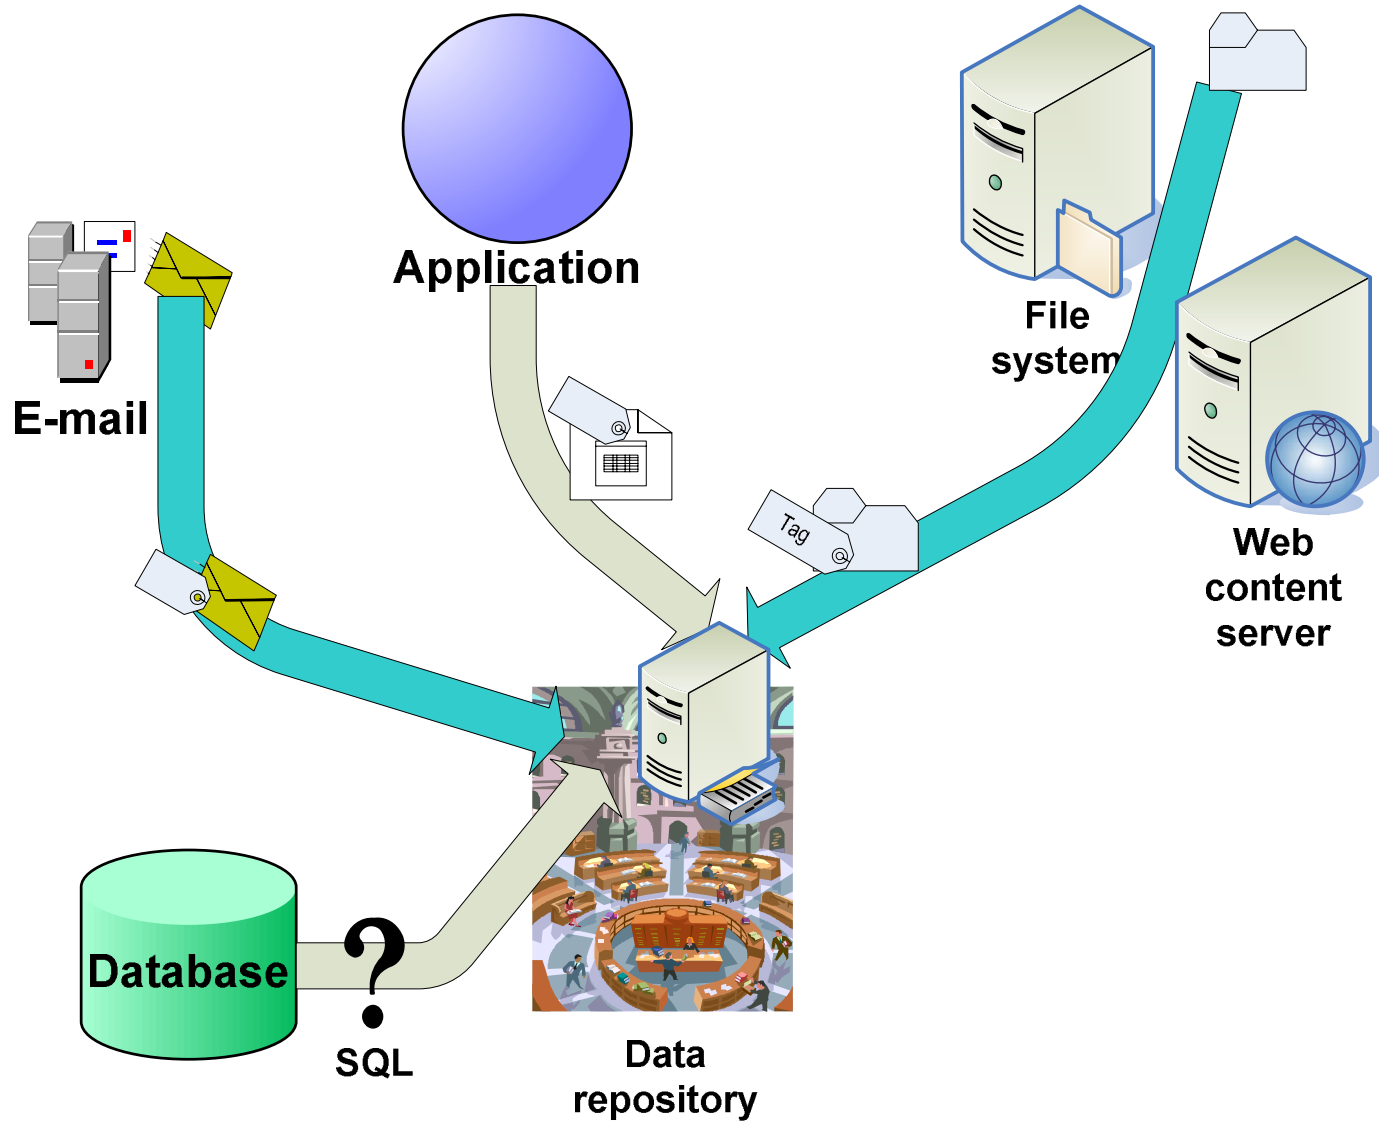
## employee table

Public	John	Q	12789	123-45-6789	\$5,000
Doe	Jane	P	34939	234-56-7890	\$7,500
Flintstone	Fred	R	00212	345-67-8901	\$2,789
Rubble	Barney	F	00314	456-78-9102	\$3,012
Coyote	Wile	A	15715	567-89-0123	\$12,000
Runner	Road	A	15714	678-90-1234	\$12,001



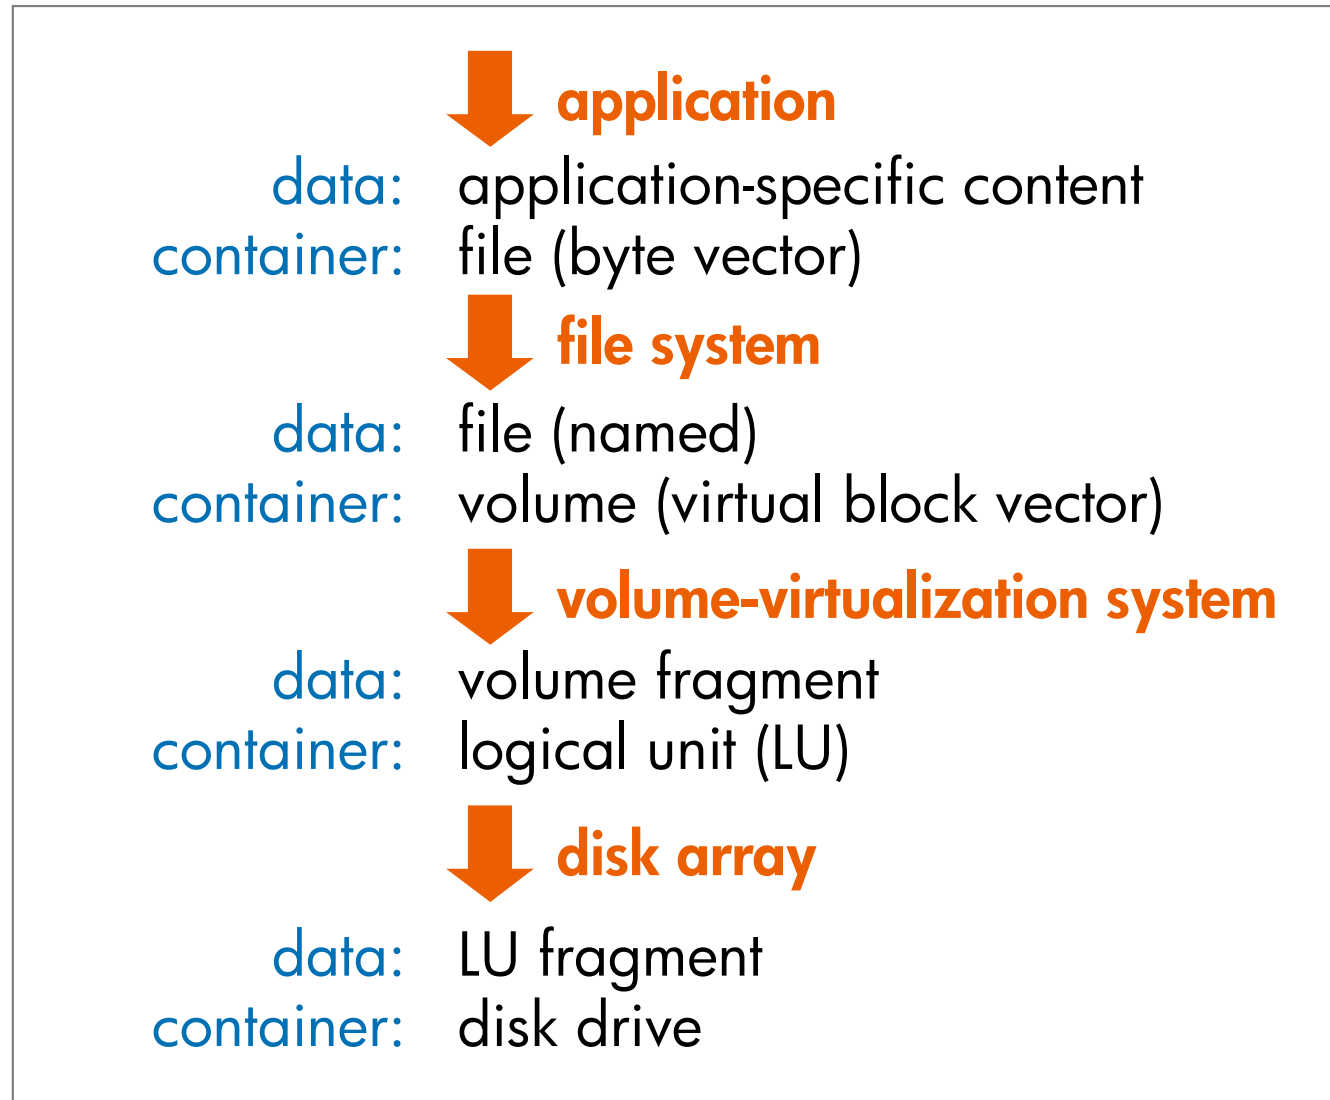


# Repositories



# Data and information infrastructure

## Contents and containers



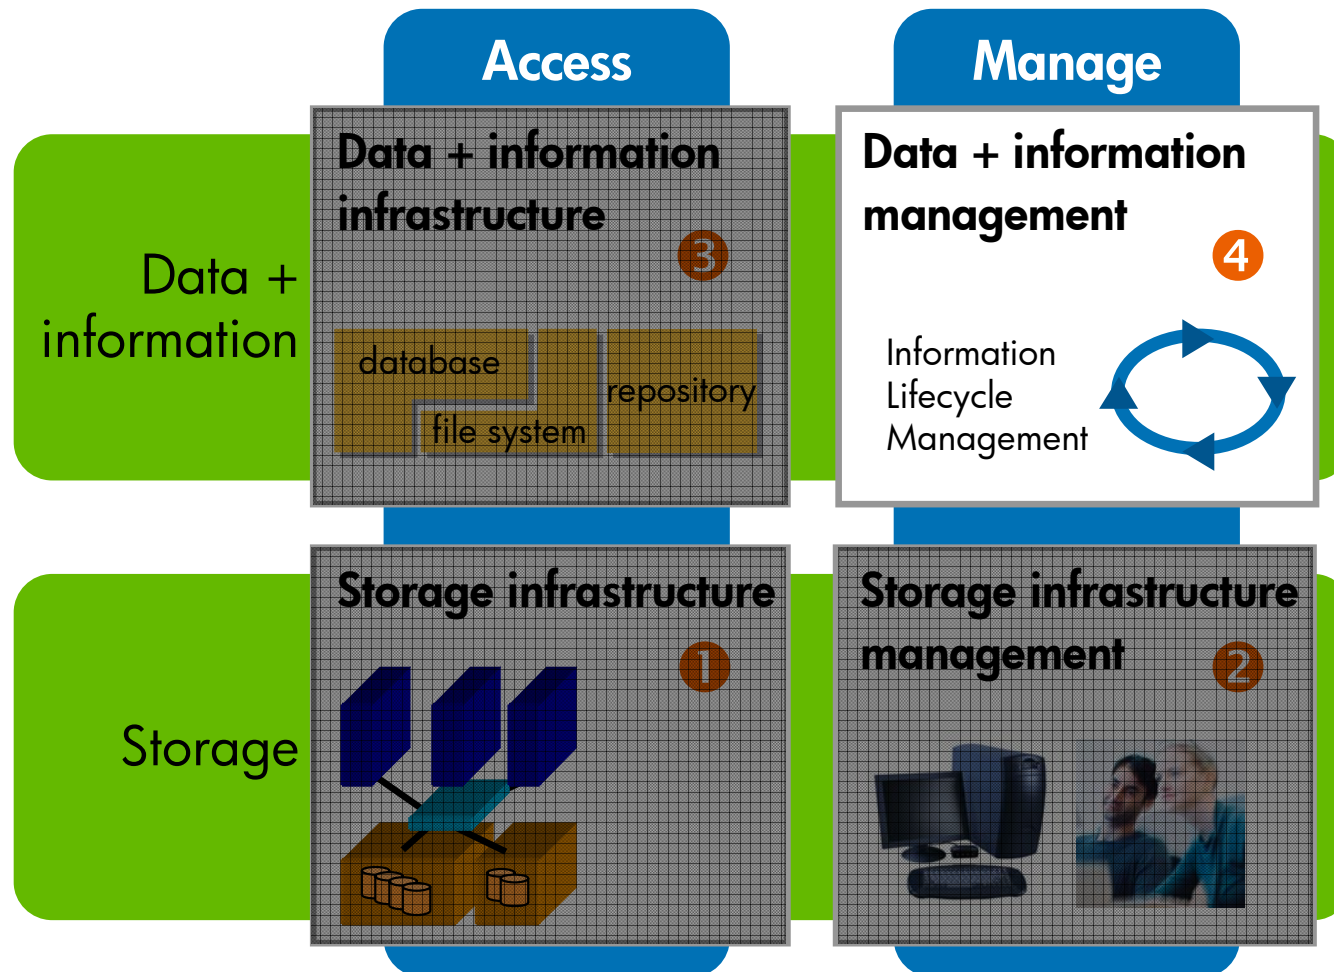
# Data/info infrastructure management



Enterprises  
and people



Business logic, processes & applications




# Information Lifecycle Management (ILM)



- What to keep
- How long to keep it
- Where to keep it
  
- Drivers:
  - value of information
  - law/regulations (compliance)
  - security & privacy needs

# Information-management market size

worldwide compliance-driven info-management revenue



source: IDC, 2006.

# A typical information lifecycle: unstructured data



## Operational

- frequently updated during 72 hours after creation

## Transitional

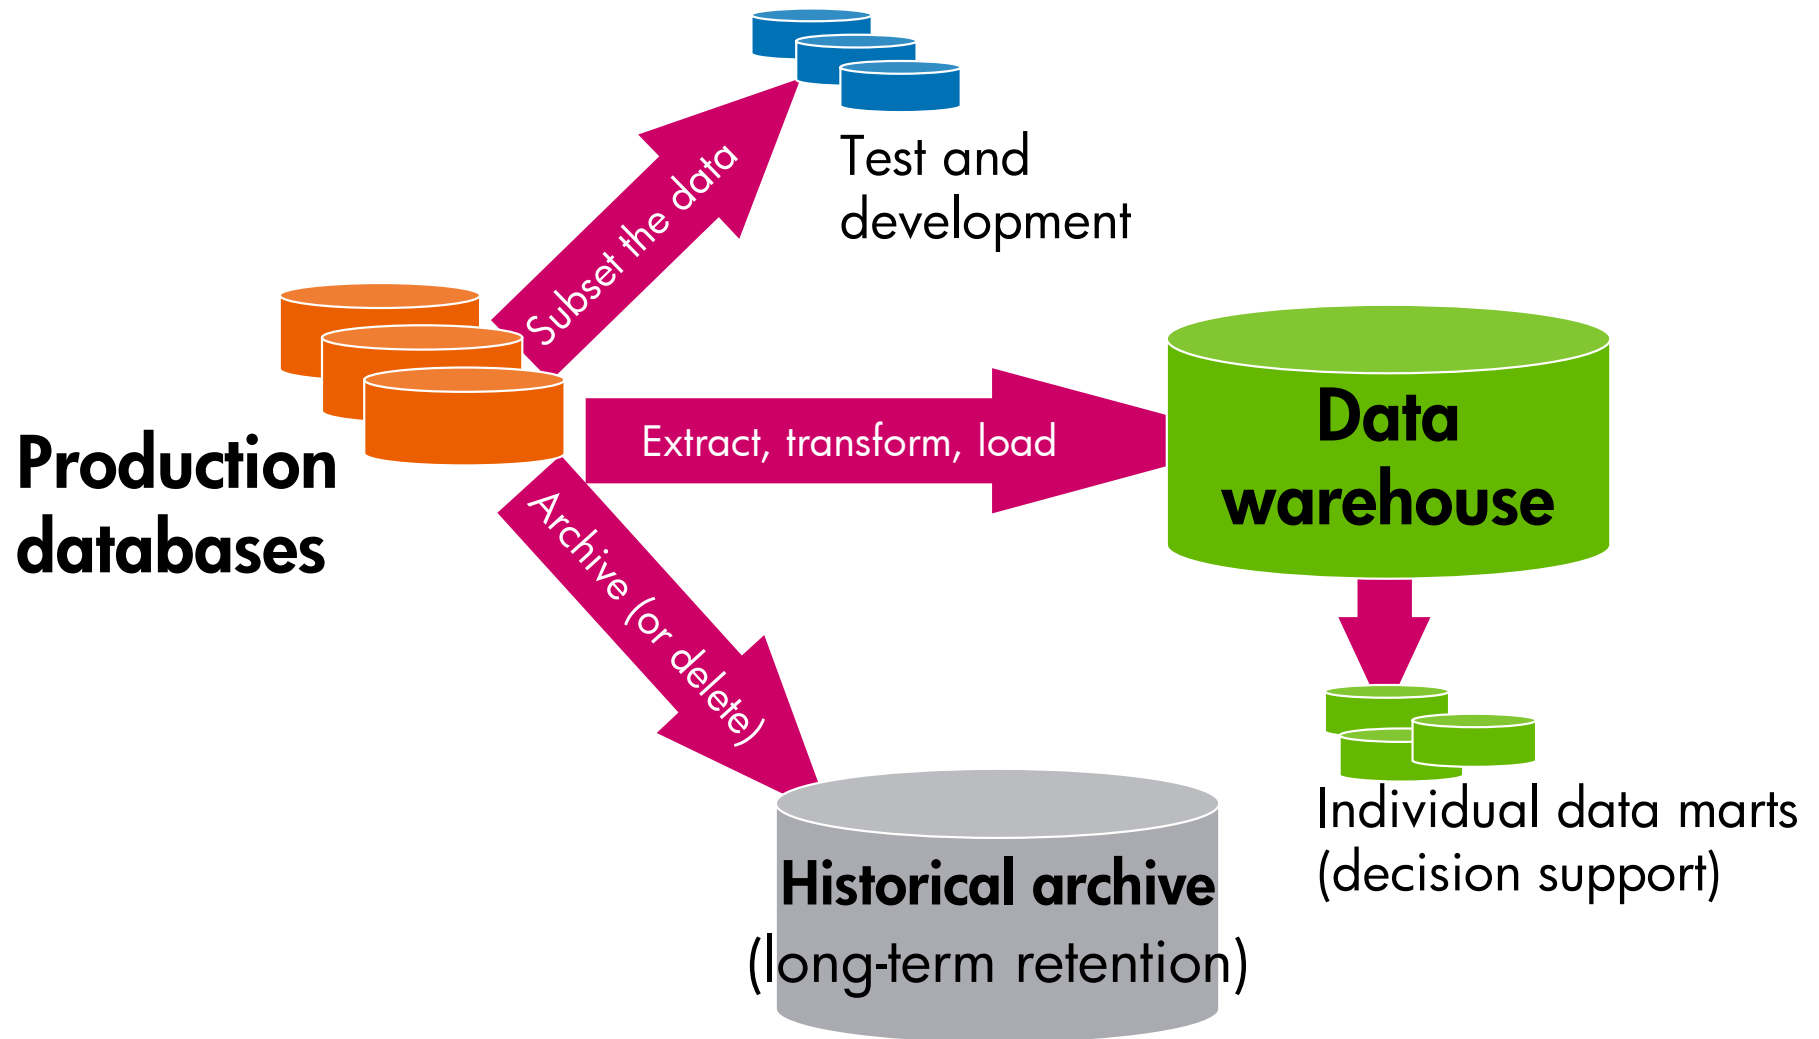
- infrequently updated
- converted to business record format

## Archival

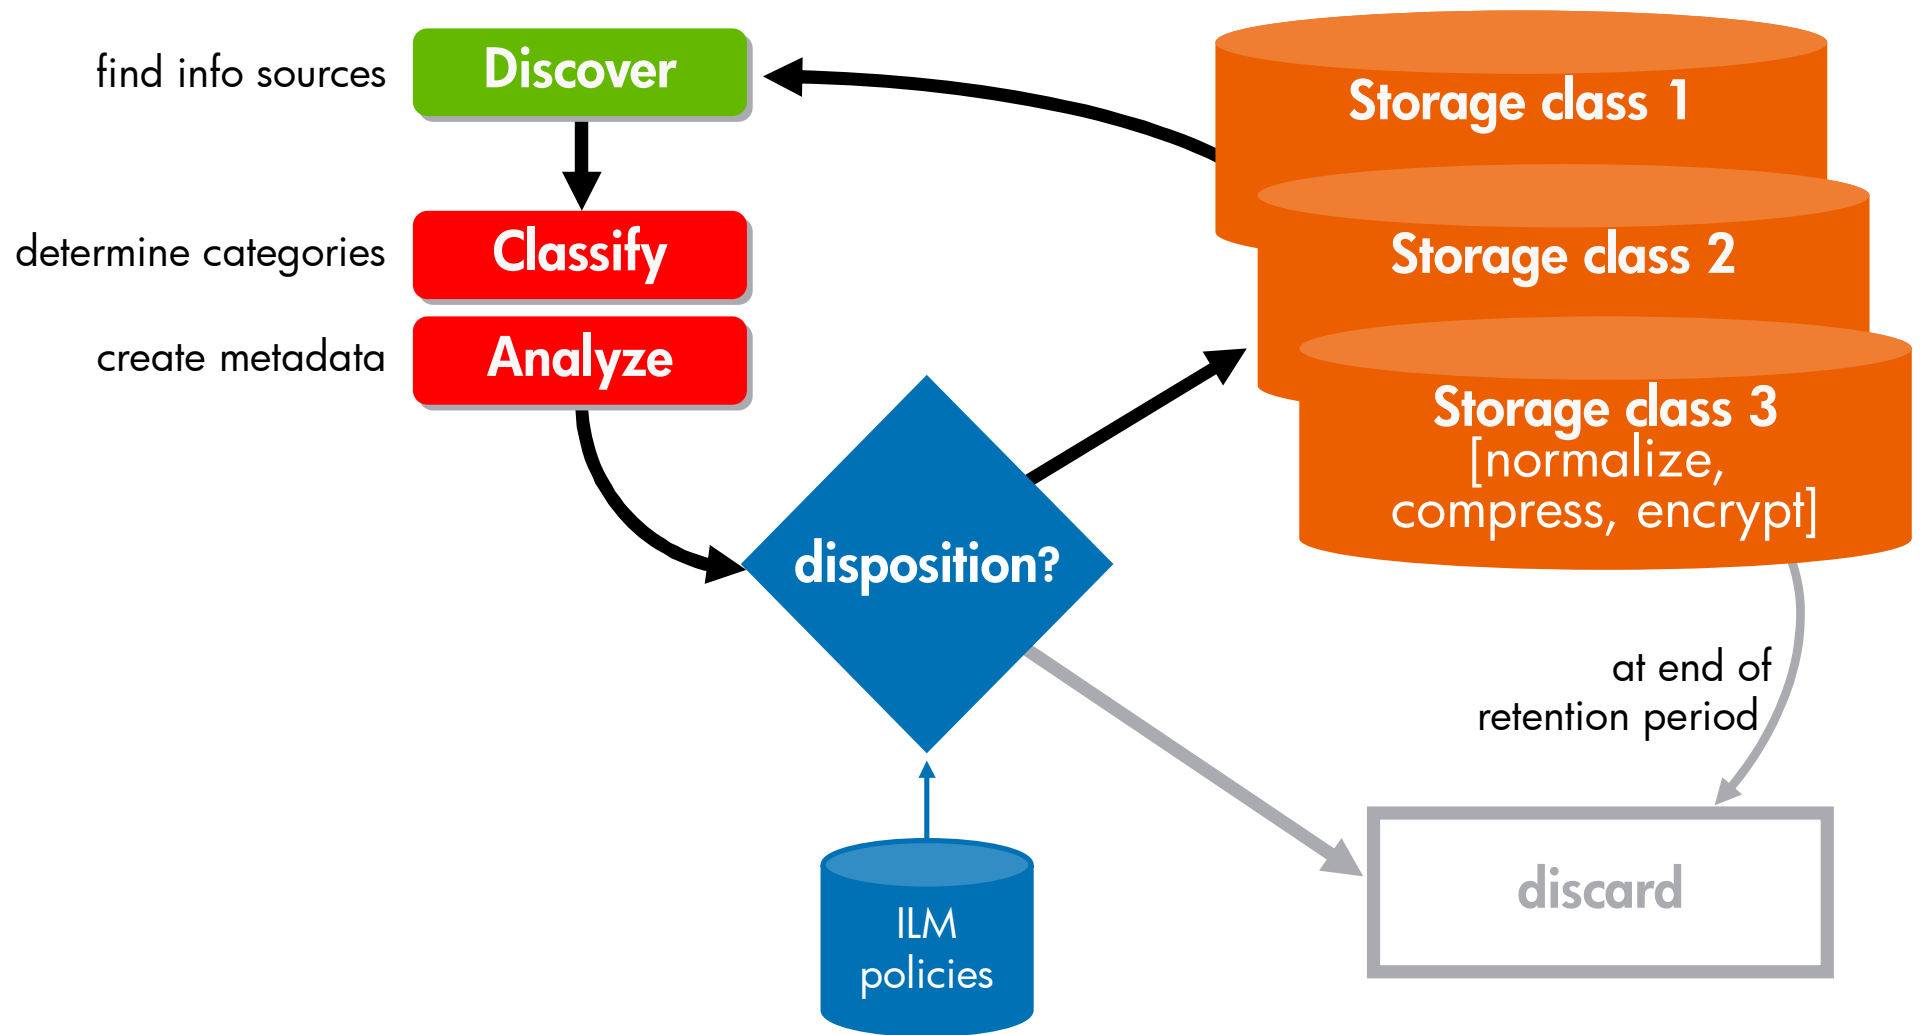
- static (rarely accessed)
- subject to long-term records management



# Lifecycles are not always linear ... the case of structured data



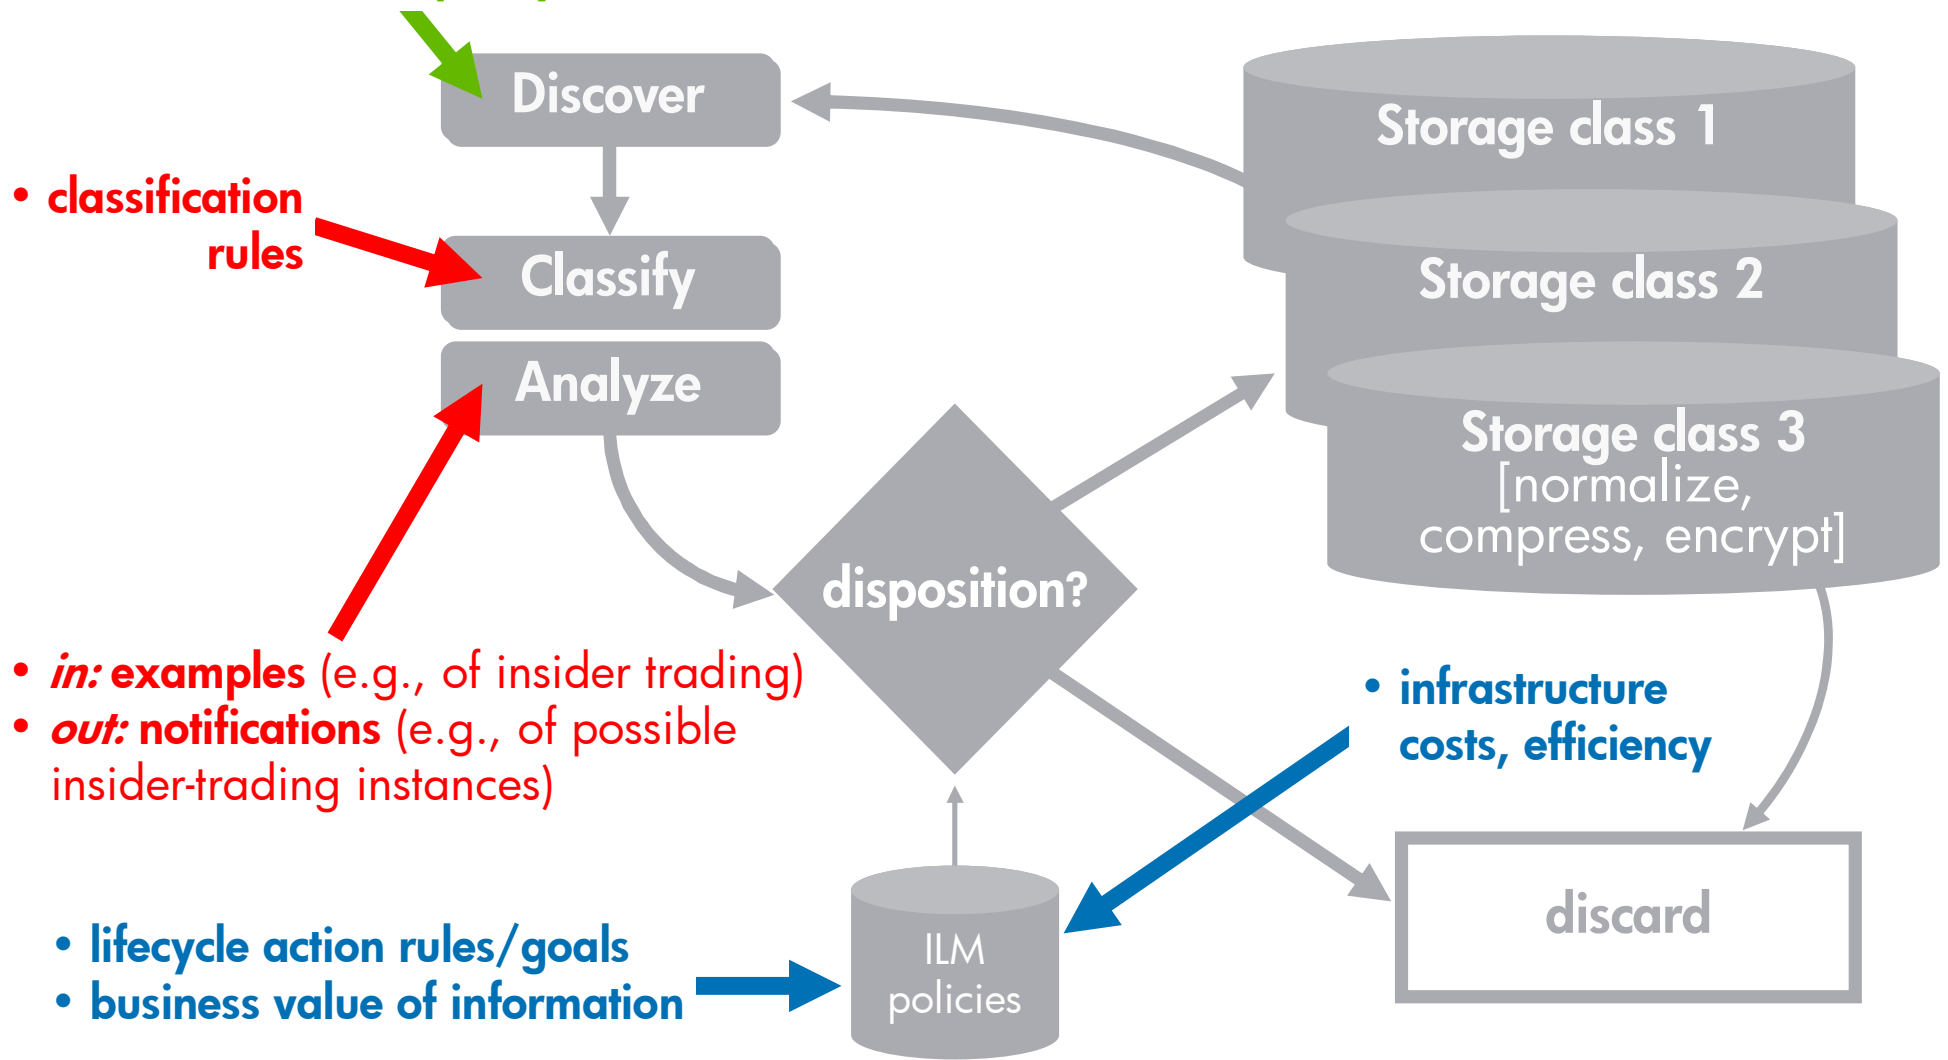
# ILM processing functions



# ILM processing functions



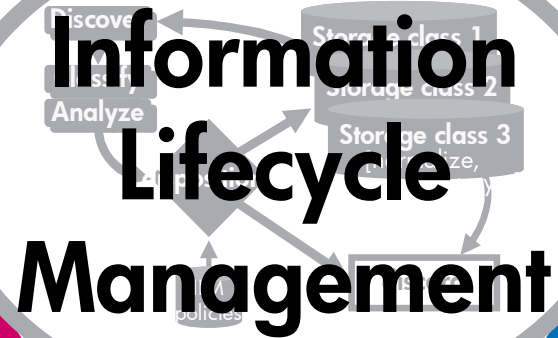
- *in*: business-process info needs
- *out*: information-quality metrics





# ILM adjacencies

**Business-process  
management**  
using information for  
business management



**Information  
semantics**  
what information  
means

**Information  
integration**  
across  
multiple sources





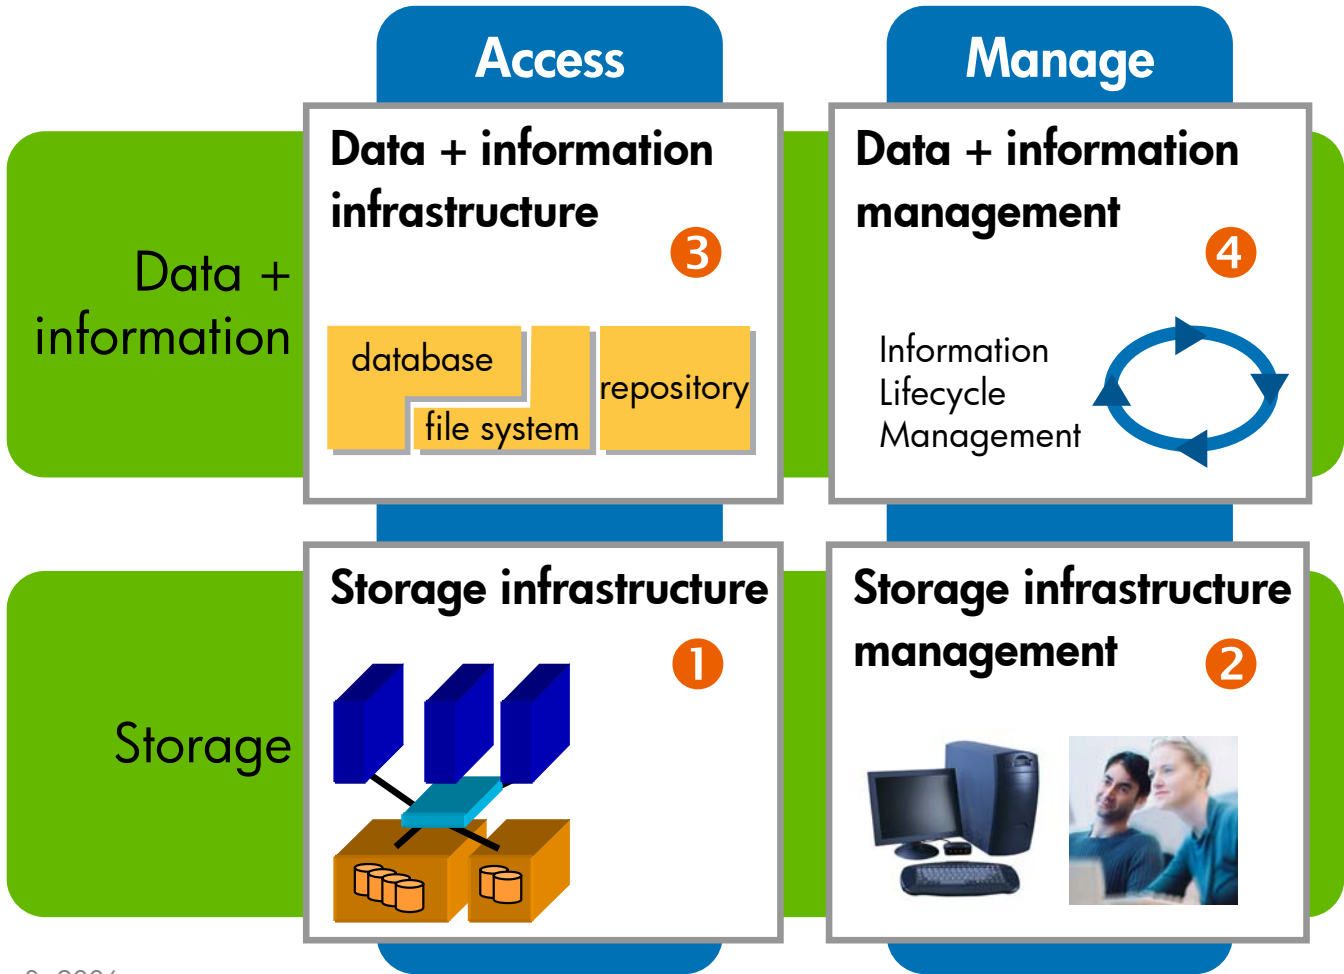
# Key messages about data/information management

- ILM helps align information management goals with business objectives
- today: a compliance-driven market today
- tomorrow: opened up by understanding of data contents

# Summary



Business logic, processes & applications







i n v e n t