

Virtualization for Storage Efficiency and Centralized Management

Genevieve Sullivan
Hewlett-Packard
gen_sullivan@hp.com

Agenda

Storage Virtualization



- what is it?
- who is it for and why?
- how does it work?
- where's the value?

Agenda

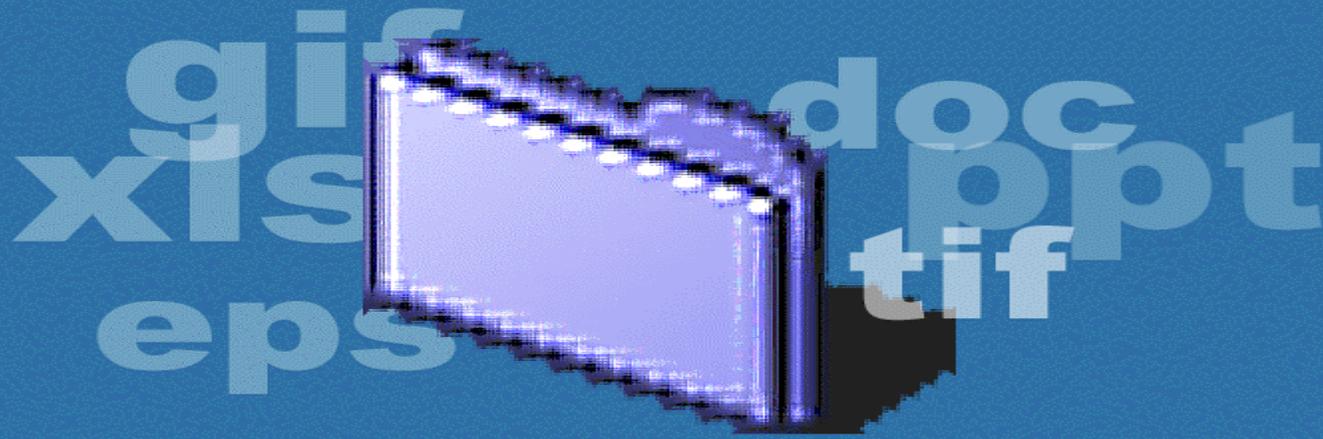
Storage Virtualization



- **what is it?**
- who is it for and why?
- how does it work?
- where's the value?

folders and files on your hard drive

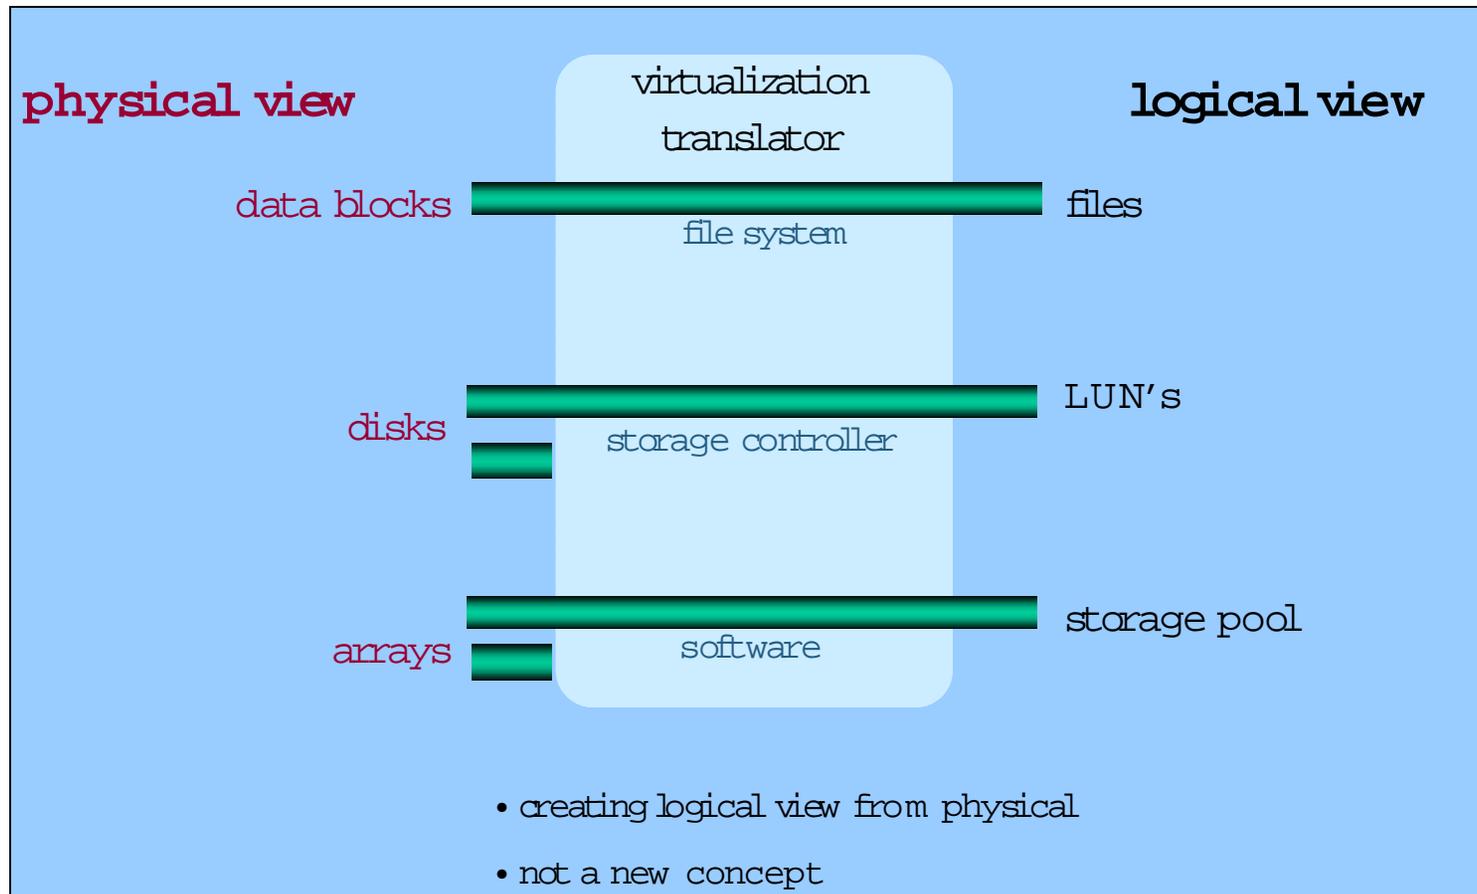
virtualization example



folders and files are virtual representation of data blocks

storage virtualization

a translator



defining storage virtualization

Storage Networking Industry (SNIA):

The act of abstracting, hiding, or isolating the internal function of a storage (sub) system or service from applications, computer servers or general network resources for the purpose of enabling application and network independent management of storage or data.

storage virtualization is:

- a logical view of physical resources
- a logical view of other logical resources
- a method for creating a common pool of storage from different disks or arrays
- a way to reduce complexity for the end user, simplifying storage maintenance activities
- an enabler of additional storage services such as replication and migration

storage virtualization is not:

- a new concept, although some implementations are
- a single defined technology
- defined the same way by all vendors
- always implemented the same way, even when different vendors use the same terms

Agenda

Storage Virtualization



- what is it?
- **who is it for and why?**
- how does it work?
- where's the value?

why virtualization?

water example

water users care about:

capacity

- household requirements
- ability to use more/less water as needed

availability

- there on the turn of a knob

they don't care about:

capacity

- size of reservoir
- number of reservoirs

availability

- neighborhood pipe construction
- number of water pumps

why virtualization?

virtualization enables storage to become more like a utility

storage users care
about:

capacity

- application requirements
- ability to use more space

availability

- failure resistance
- recovery time

they don't care about:

capacity

- disk size
- number of disks in array

availability

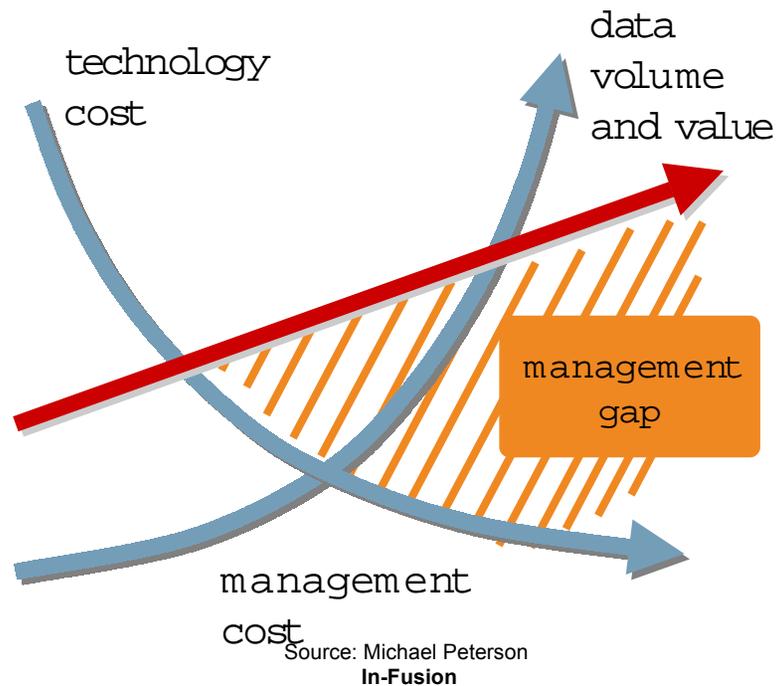
- mean time between failures (MTBF)
- path redundancy

storage virtualization – who is it for?



- storage networks with low storage utilization rates
- IT administrators spending too much time on disk management
- CxO's with imploding IT budgets
- applications with unpredictable storage demands
- heterogeneous environments that need data replication

What are their challenges?

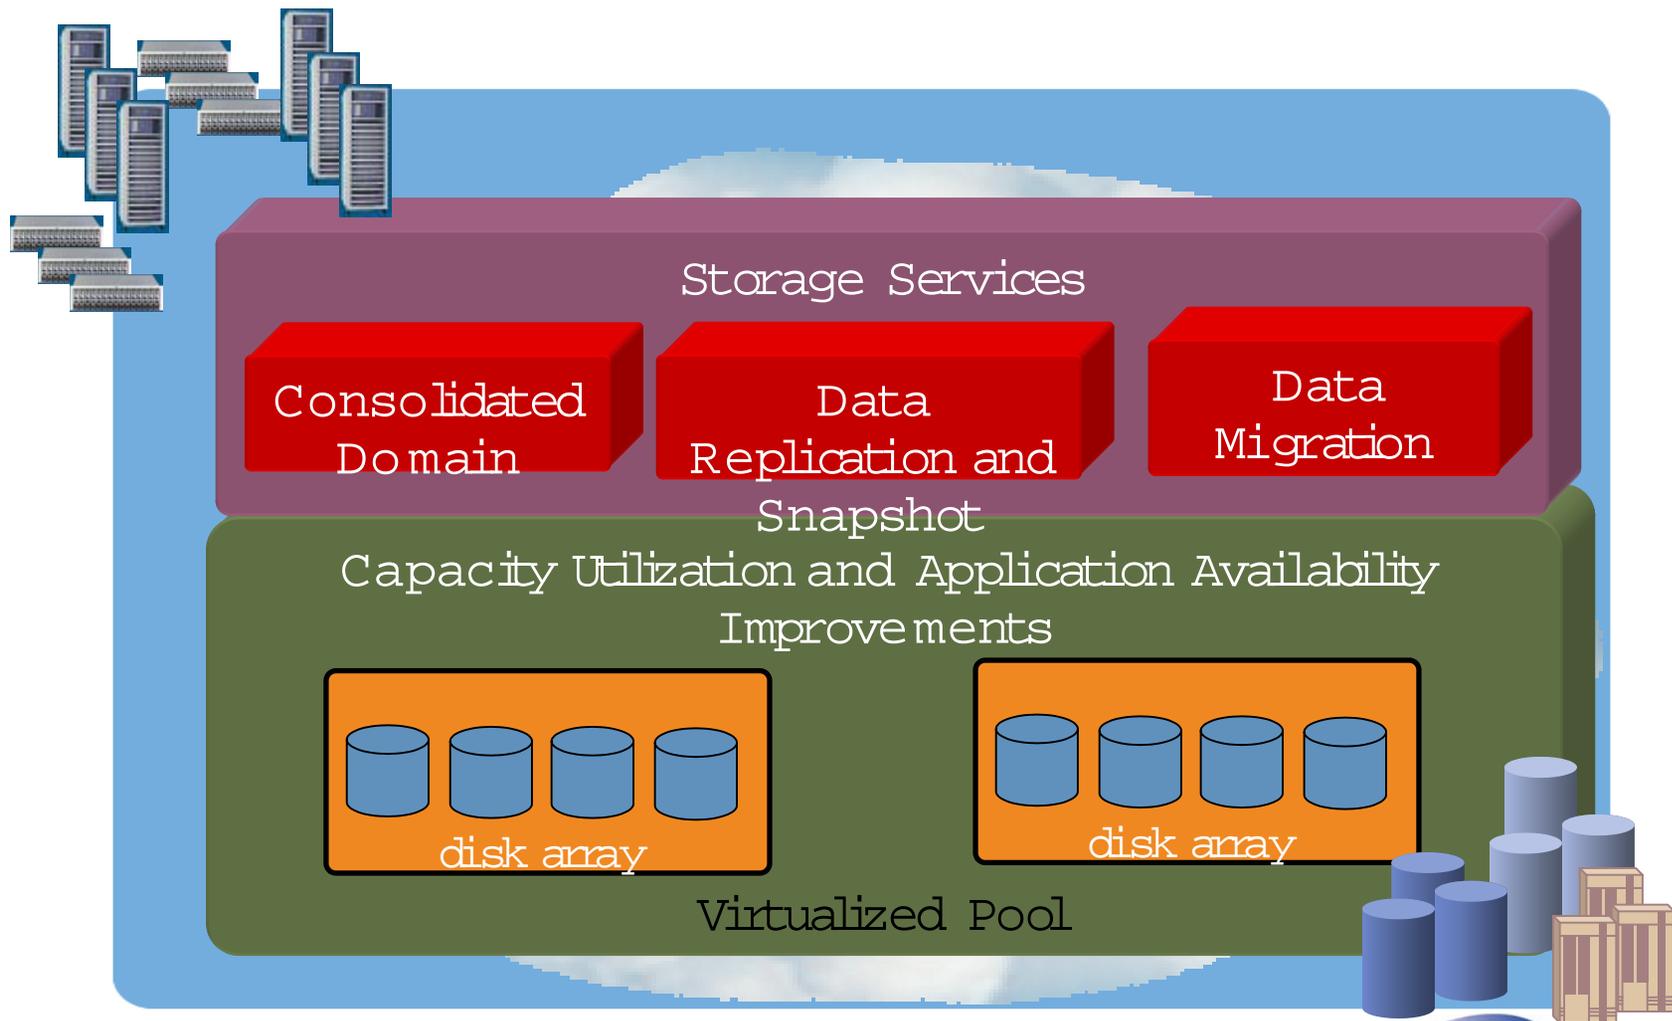


- business increasingly demands more data
 - demands better access
 - greater resource sharing
 - higher performance
 - less application downtime
- business reality requires greater efficiencies
 - capacity utilization
 - improved management
- **storage virtualization helps close the management gap**
 - storage virtualization
 - advanced storage services

Greater Efficiencies with Storage Virtualization

- Increased utilization of storage assets
 - Pooling of assets eliminates stranded capacity
- Reduced/ Eliminated application downtime resulting from storage maintenance activities
 - Storage administrators quickly perform storage maintenance activities without disrupting business applications
- Advanced storage services enabled across heterogeneous components
 - Data replication and data migration services are performed across and between unlike storage devices from multiple vendors
- Consolidated domain for centralized management
 - Management functions can be performed across the heterogeneous storage pool

Benefits of Virtualization



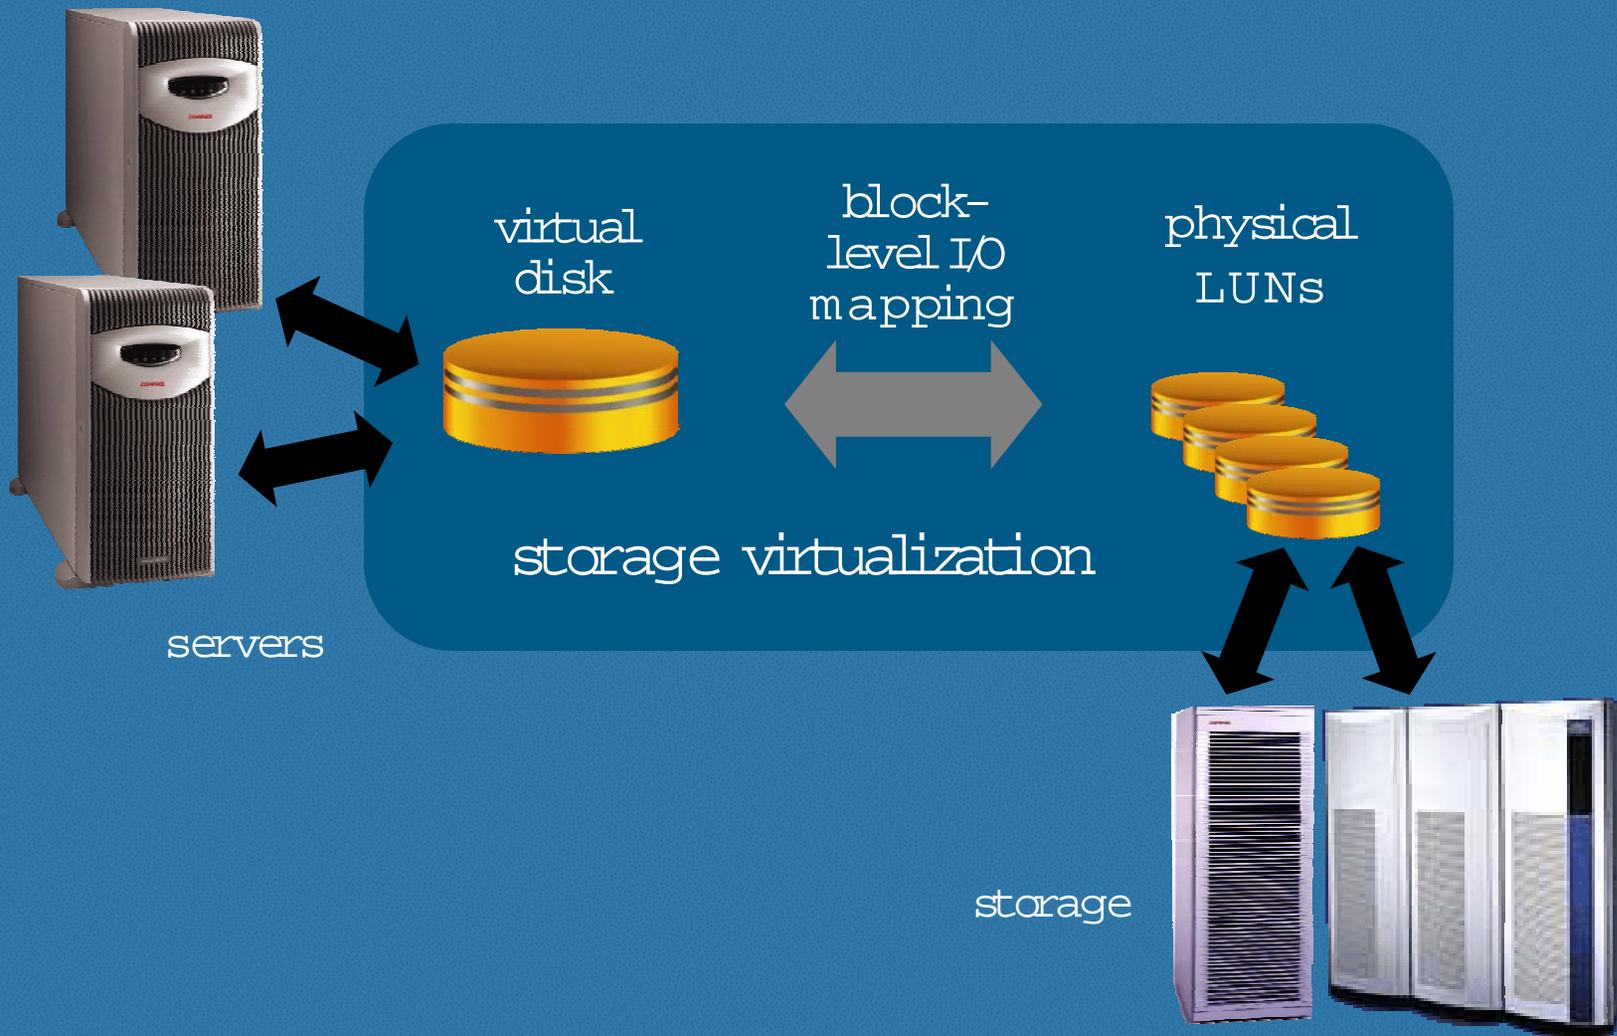
Agenda

Storage Virtualization

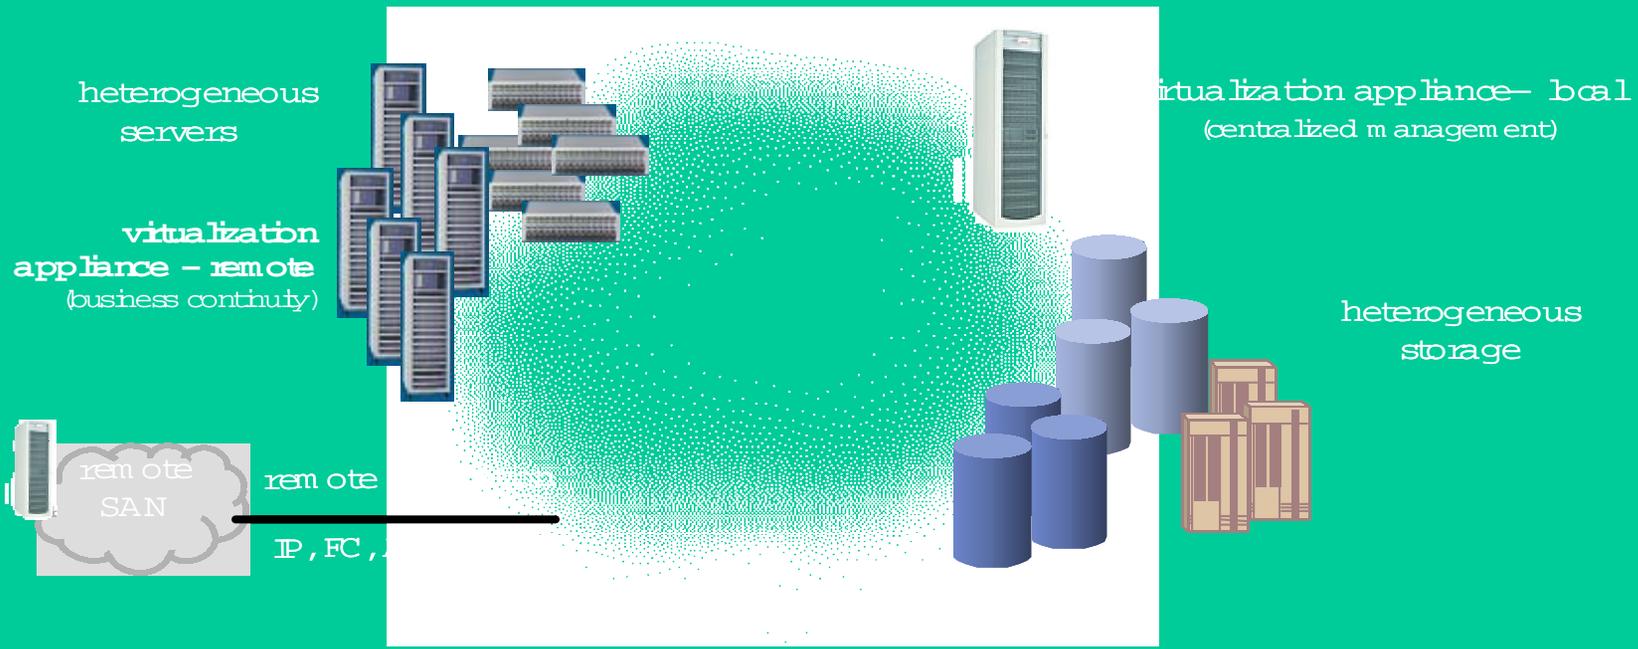


- what is it?
- who is it for and why?
- **how does it work?**
- where's the value?

how it works



virtualization solutions span many levels of IT



host based
Single host/cluster connected to many arrays

array based
SAN switch many, heterogeneous hosts connected to a single storage device

network based
Many hosts connected to many arrays

server-based virtualization

virtualization
intelligence



server



Storage
network



hp
500GB

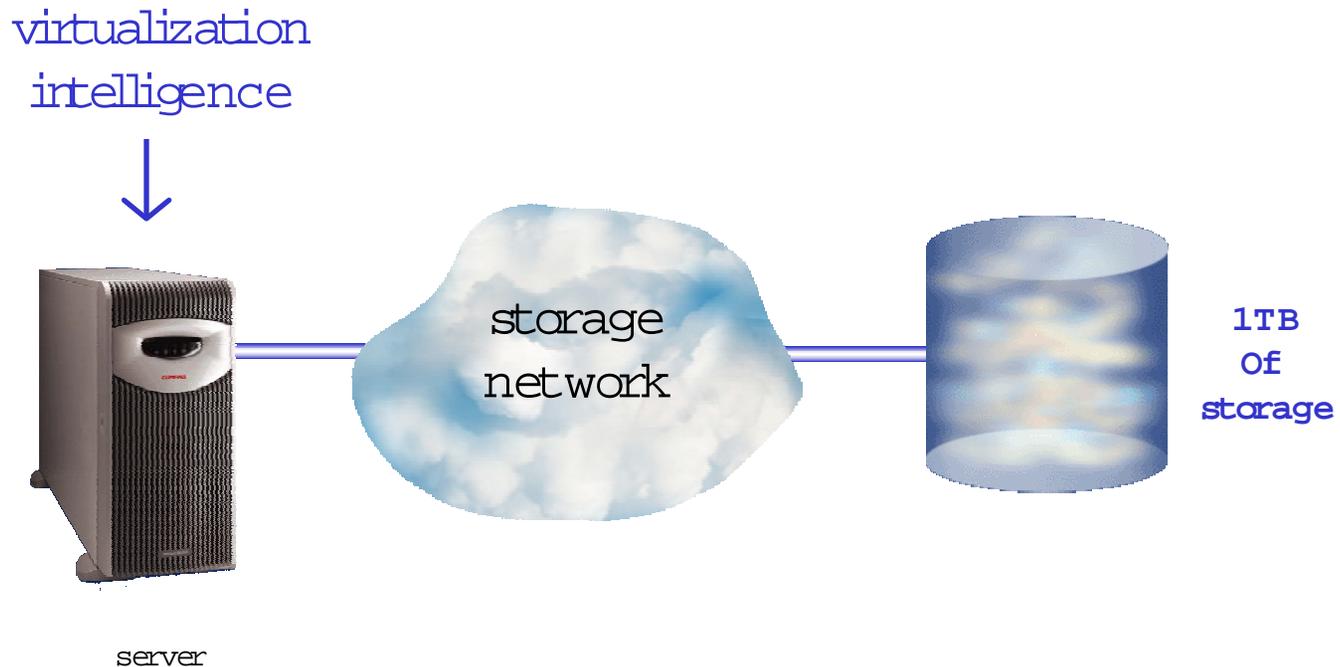
Array 1



Compaq
500GB

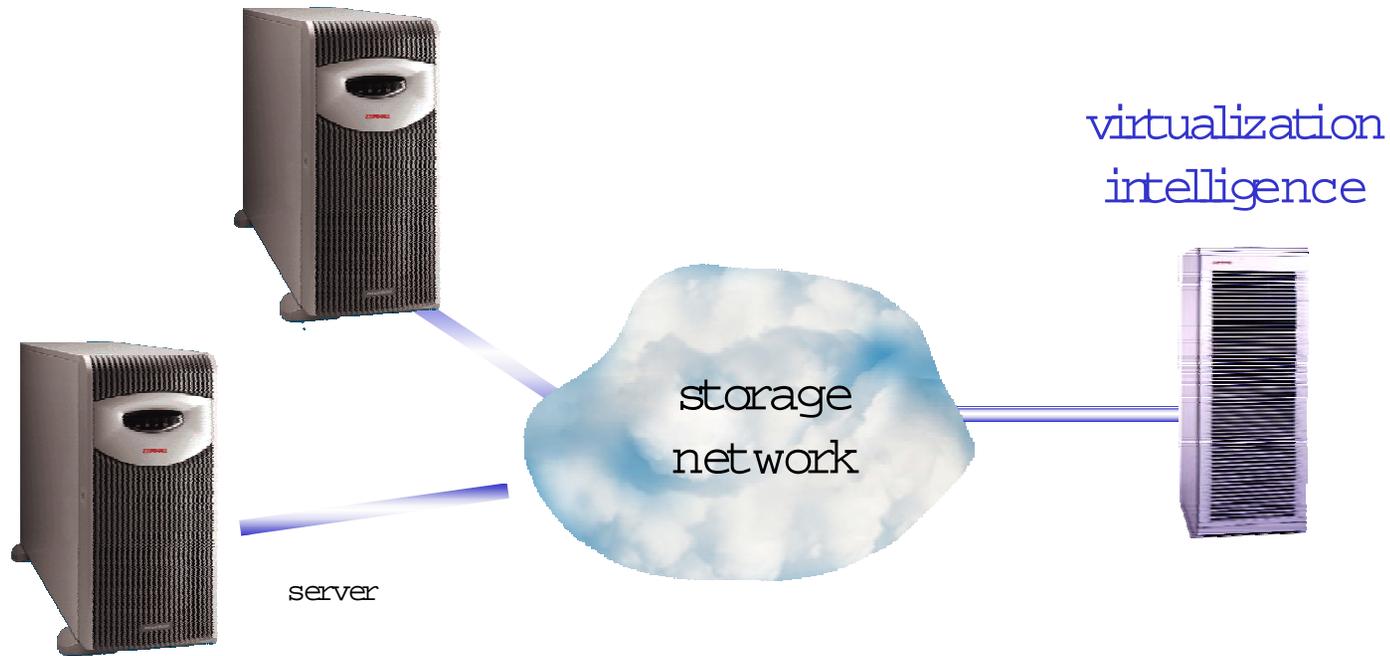
Array 2

server-based virtualization



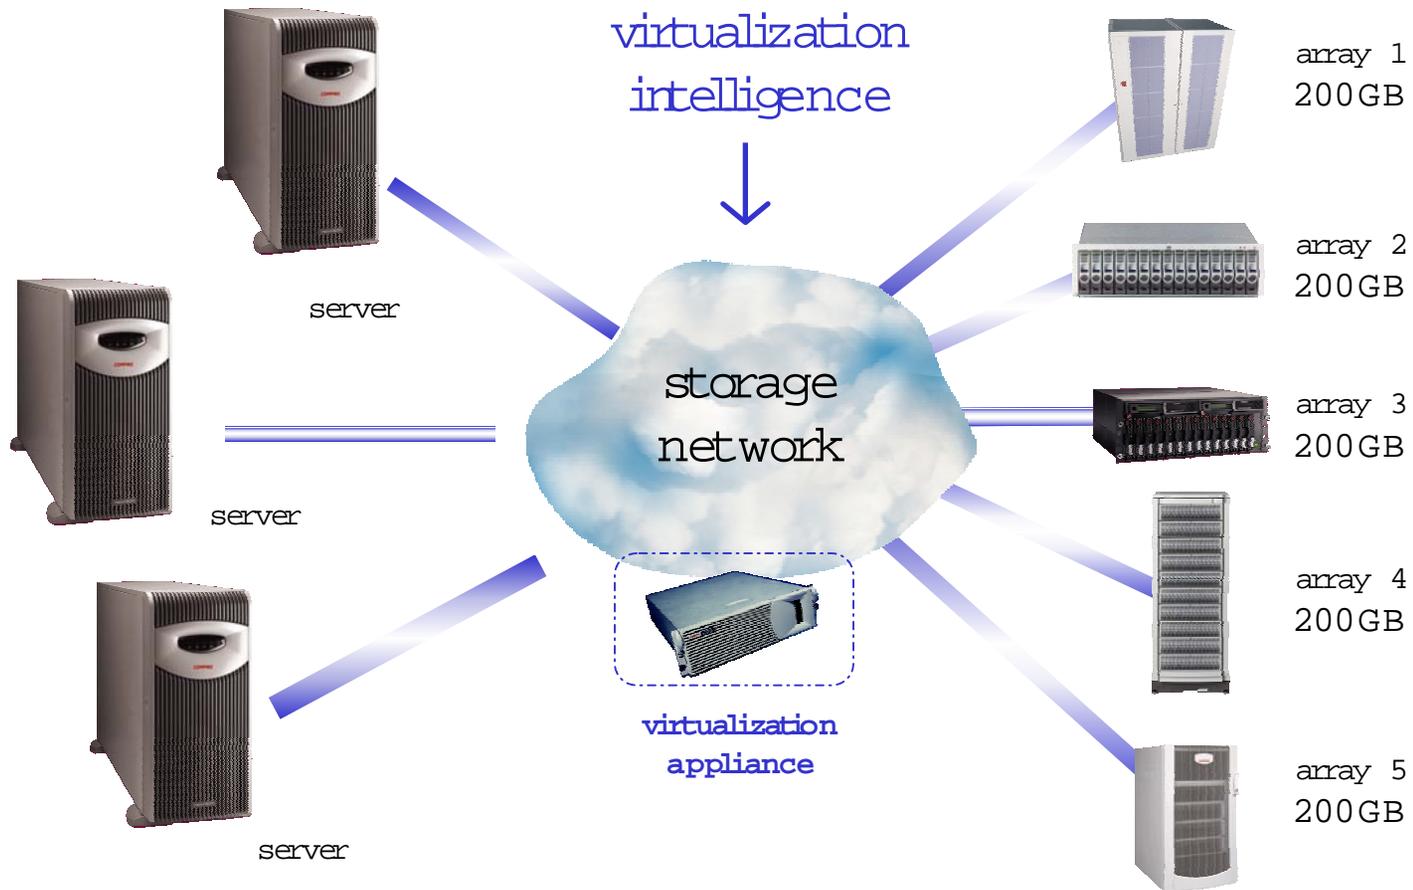
server sees 1TB of storage as one logical unit of volume through server-based software

array-based virtualization

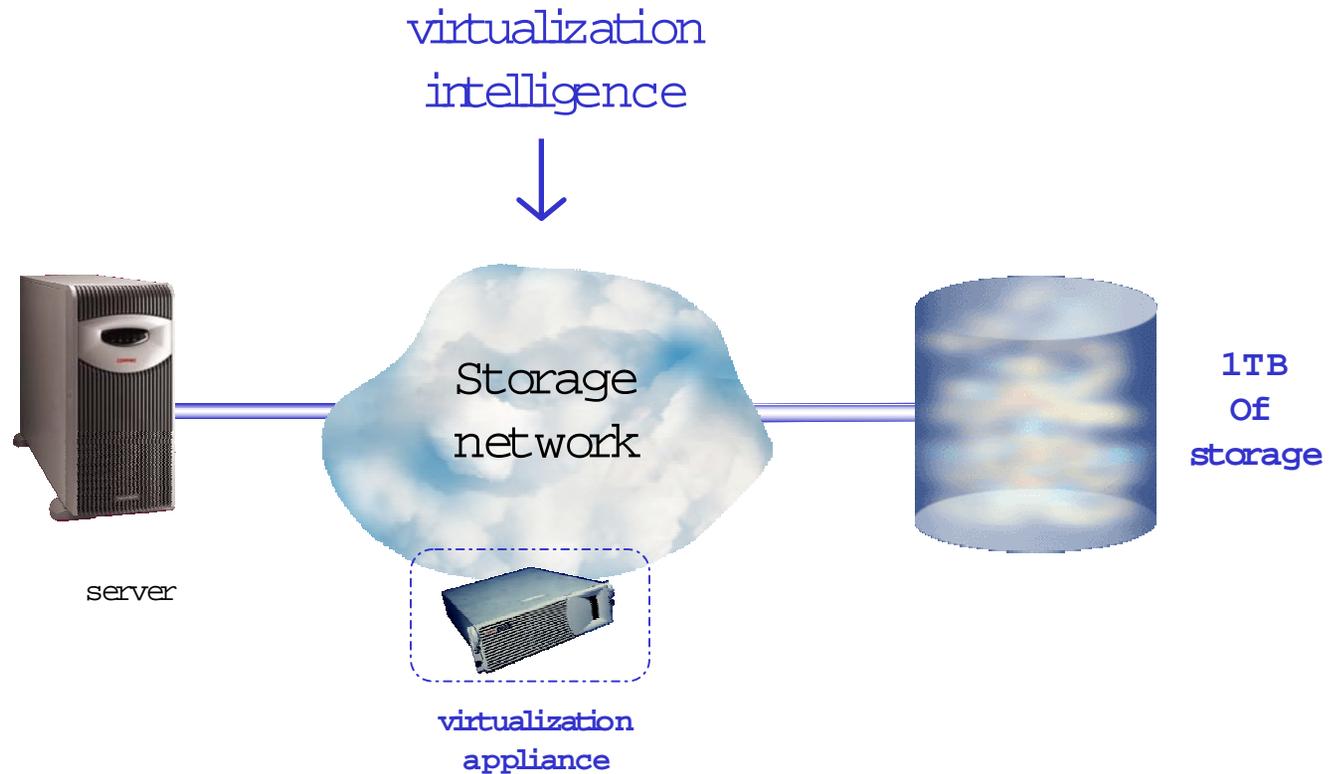


array automatically stripes data across disks and tunes for better performance and eliminate hot spots

network-based virtualization



network-based virtualization



server can see a 1TB LUN via network-virtualization appliance

Agenda

Storage Virtualization

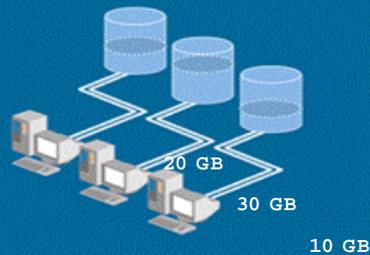


- what is it?
- who is it for and why?
- how does it work?
- **where's the value?**

value of virtualization
**storage
consolidation**

traditional distributed storage

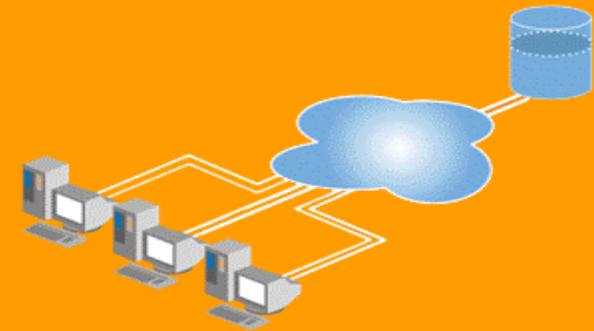
- difficult & expensive to manage
- 30-50% utilization



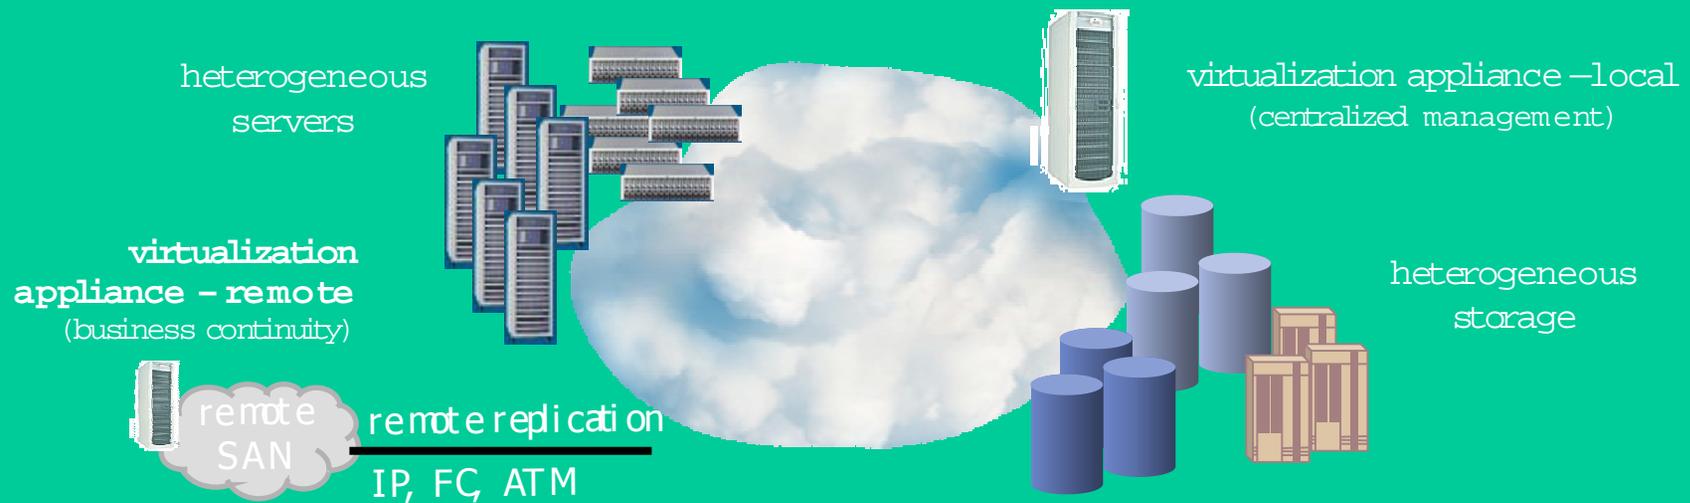
increase storage utilization
by creating virtual disks

bring arrays from different
vendors into one
manageable storage pool

lengthen useful life of older
storage by spanning
LUN's across arrays



virtualization solutions improve capacity utilization



host based

efficient storage utilization of capacity attached to a single host or cluster

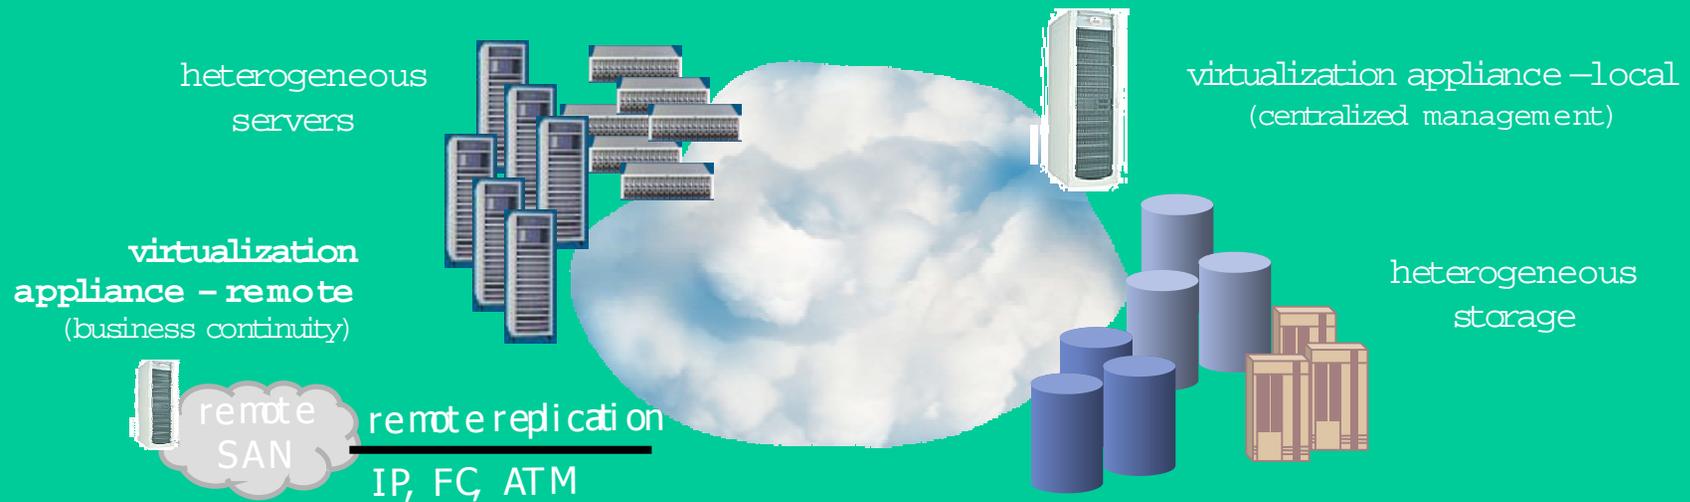
array based

Optimized storage utilization of capacity attached to multiple heterogeneous hosts

network based

Seamless SAN-wide storage capacity management

virtualization solutions reduce application downtime
resulting from storage maintenance activities



host based

array based

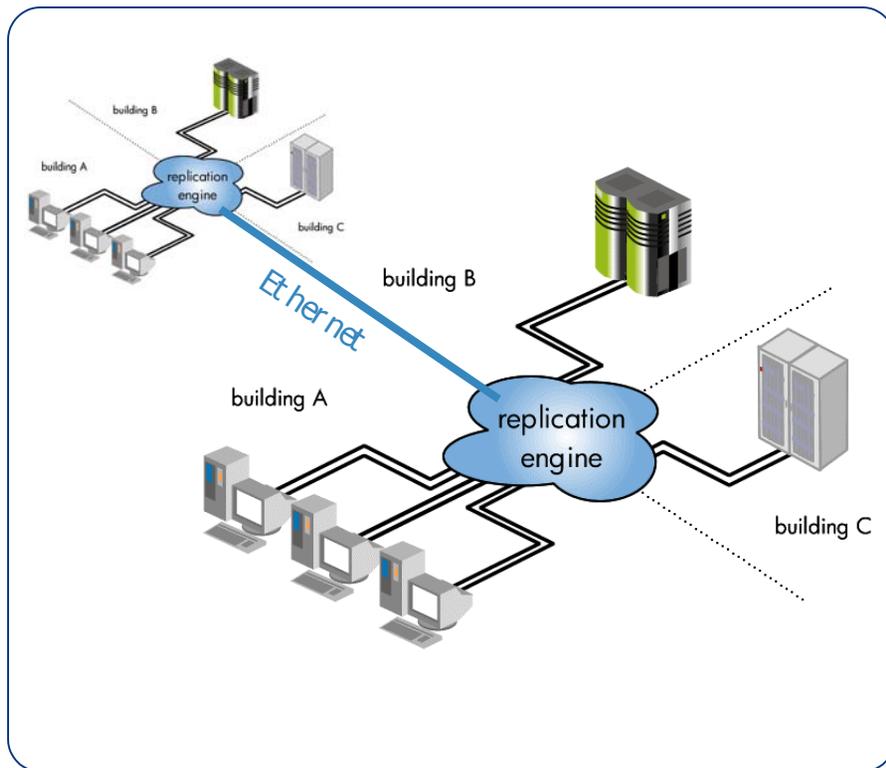
network based

Non disruptive capacity growth

Flexible allocation and reallocation of capacity

Increased administrator efficiency

value of virtualization business continuity



fibre channel mirroring

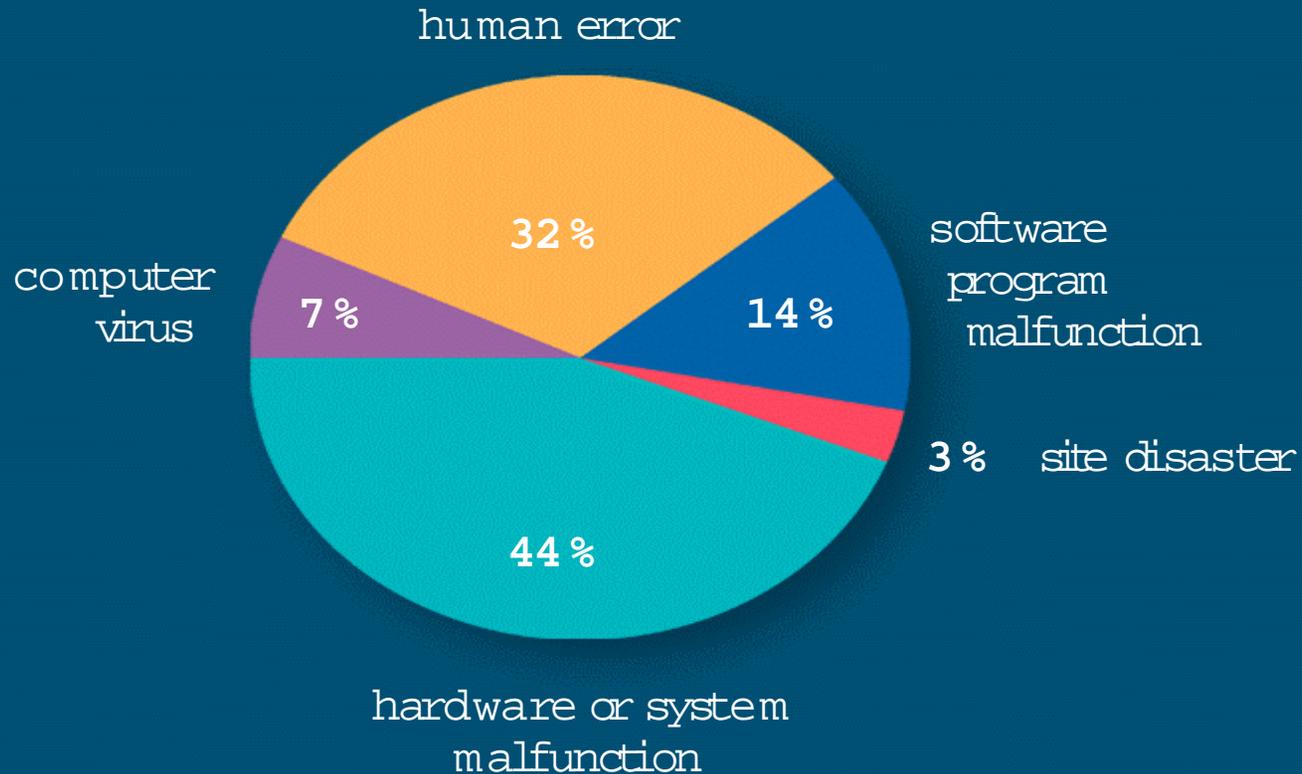
remote IP mirroring for DR

space-efficient snapshots

snapclones

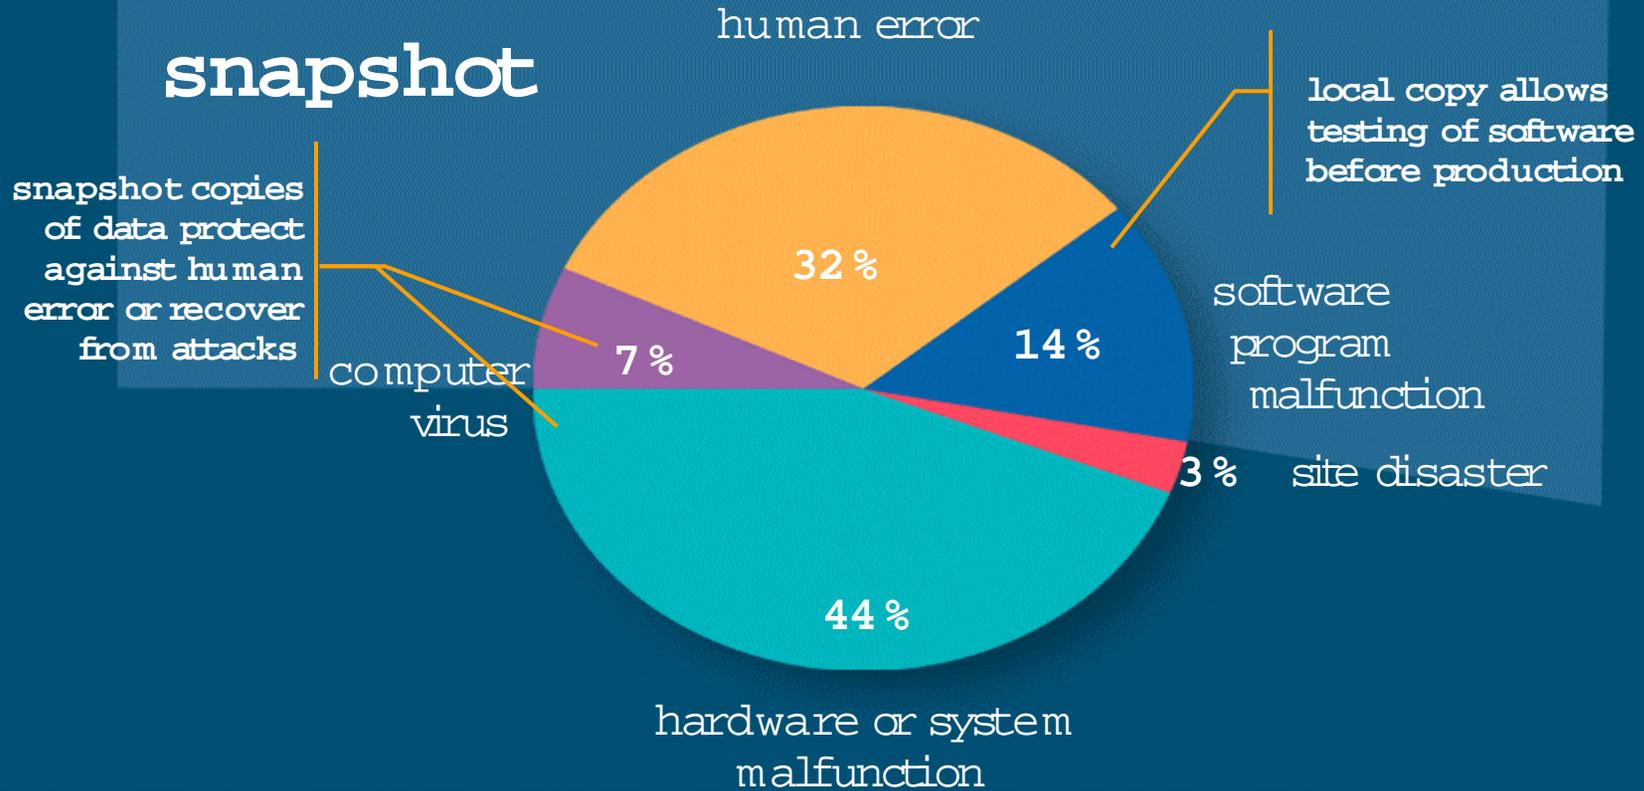
redundant/cluster

data loss



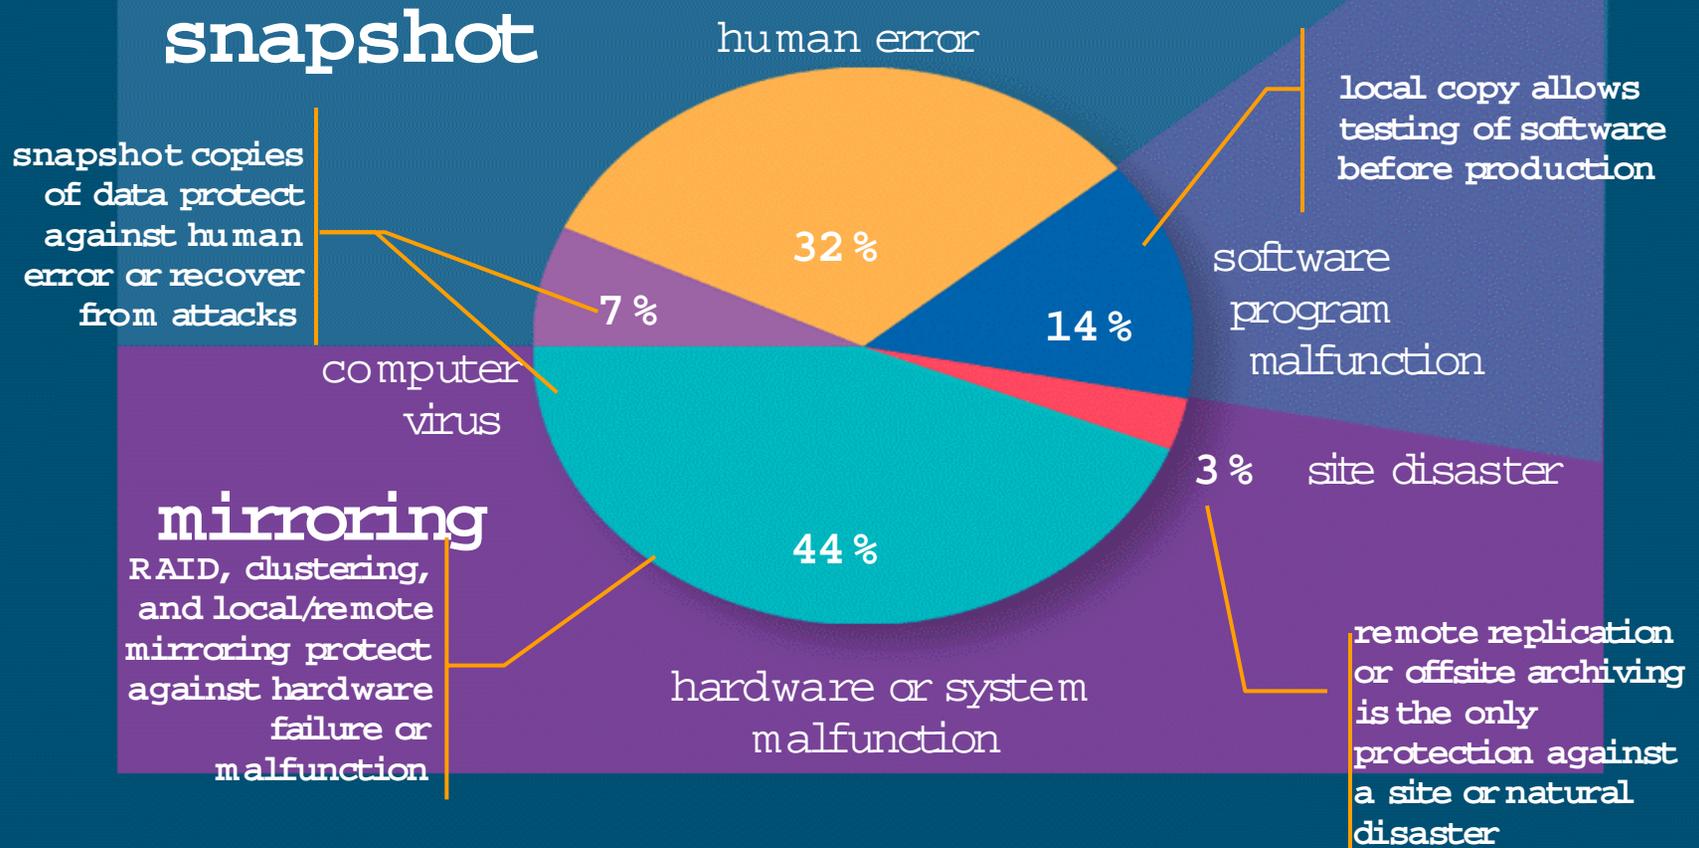
source: Ontrack, a data availability service provider

data loss

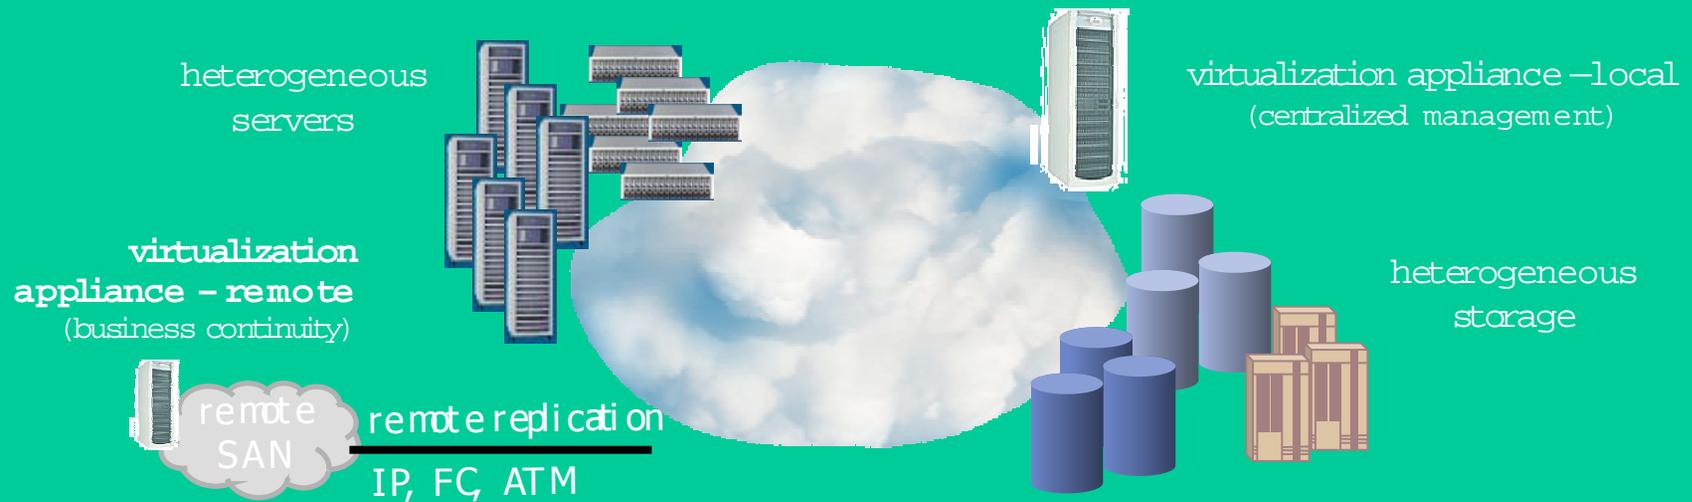


source: Ontrack, a data availability service provider

data loss



virtualization enables advanced storage services



host based

space efficient
snapshots for fast
backup and data
recovery

array based

Data replication for
business continuity

network based

Data replication
(snapshots and mirroring)
and data migration for
heterogeneous
environments

value of virtualization

Reduced Replication Costs

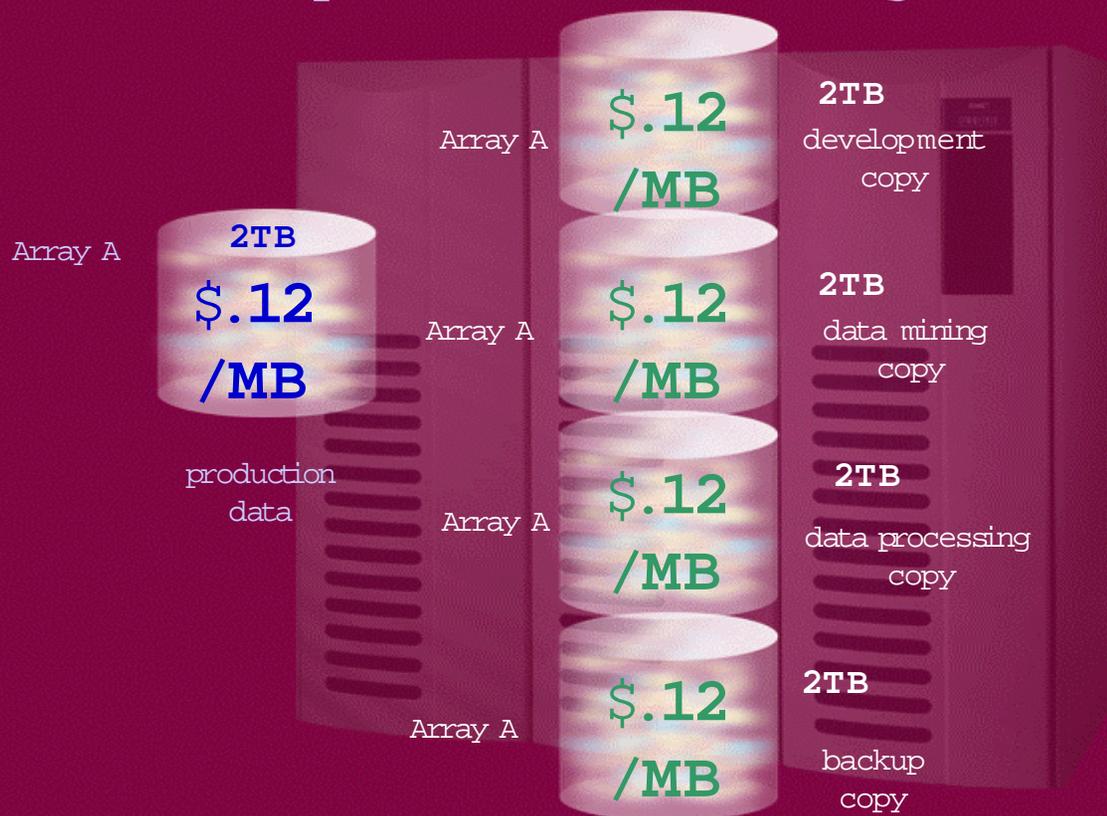
all copies of data should
NOT be treated equal!



- performance, scalability and availability requirements are different for each copy
- cost per MB for copies should reflect copy requirements, not production requirements
- many customers pay for high cost copies only because they are locked into creating copies on the same platform

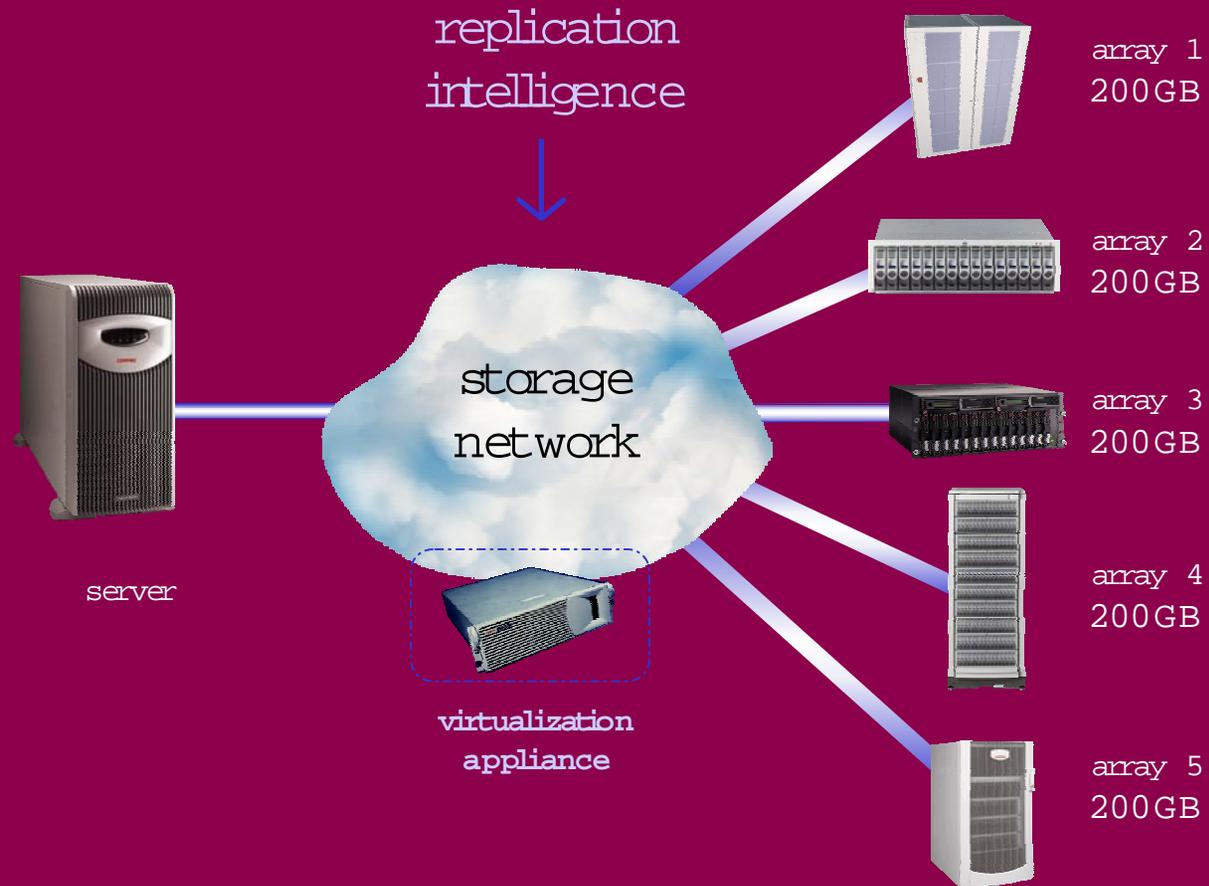
value of virtualization

copies via traditional homogeneous method

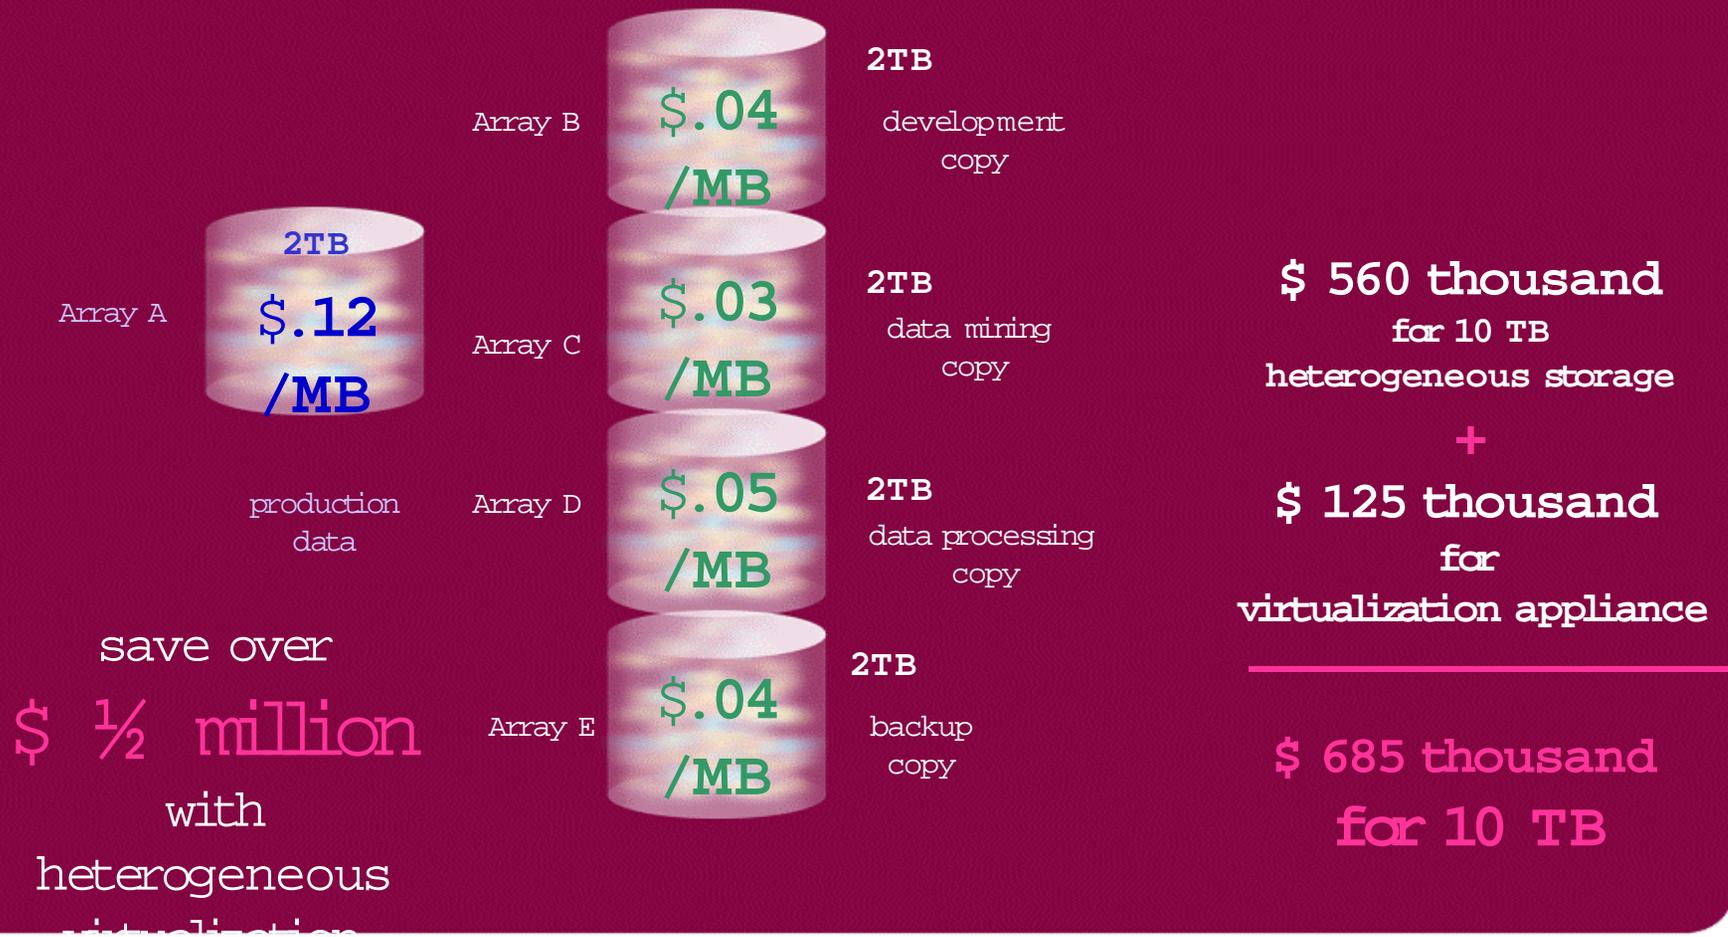


\$ 1.2 Million
for 10 TB

network-based heterogeneous virtualization



value of heterogeneous virtualization



value of virtualization lower storage management costs

The screenshot displays the HP StorageApps SV3000 management interface. At the top, there is a navigation menu with options: home, node conf, snapshot, LUN Map, LUN resize, hosts, IP mirror, FCP mirror, targets, utils, and LUNs. The HP logo and 'invent' tagline are visible. The main content area is titled 'view luns' and contains a table with the following data:

	lun	size	scsi id	physical lun	hba	bus
info	0	13.57 GB	0	0	2	0
info	1	13.57 GB	2	2	6	0
info	2	13.57 GB	3	3	6	0
info	3	11.57 GB	4	4	6	0
info	4	2.88 GB	5	5	6	0
info	5	6.88 GB	6	6	6	0

Below the table, a detailed view for 'lun 0' is shown, displaying various statistics and configuration options:

- bytes read: 1683056223
- bytes written: 384
- i/o read: 112063
- i/o written: 16
- bytes mirrored: 1682964576
- i/o mirrored: 111148
- partition: no
- partitioned: no
- expanded: no
- mapped: yes
- snapshot source: no
- snapshot target: no
- fc mirror source: no
- fc mirror target: no
- back...

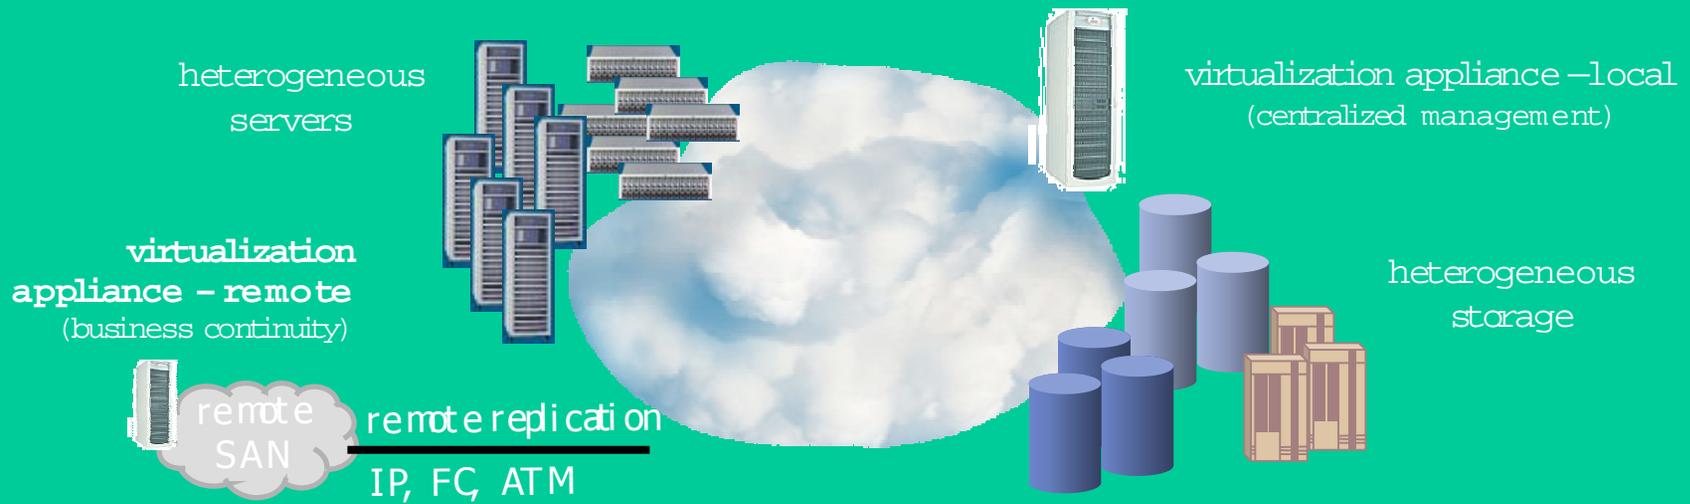
quickly allocate pooled
storage to hosts from
central console

on-the fly array tuning to
eliminate hot spots

online LUN expansion

automatic RAID level
migration

a consolidated domain for centralized management



host based

array based

network based

management functions are performed on a single pool of capacity

value of virtualization

storage virtualization addresses your storage challenges



explosive storage growth



dynamically add capacity without downtime



shortage of skilled IT staff



storage growth/allocation tasks take less time to accomplish



storage management complexity



manage storage pools centrally from simple GUI



manage technology lifecycle



extend useful life of arrays by adding storage intelligence via hosts or network appliance



keep business running while environment evolves



make copies of data via server, network, or array virtualization



total cost of ownership



lower management costs and higher storage utilization rates