

Obligatory rubric

- **Copyright © 2000 & 2001, Storage Networking Industry Association.**
- **SNIA-TC members may use this material freely; all others require permission from the SNIA-TC.**



SNIA shared-storage model

A work in progress ...



An architectural overview

April 2001

Contents

- **SNIA Introduction**
- **The SNIA storage model**
 - Purpose and benefits
 - Layers, functions, and services
- **Applying the SNIA storage model**
 - Common storage architectures
- **Conclusions**



SNIA Overview

- **Open forum**
 - **Work Groups open to all members**
 - **Public intellectual property policy**
- **Volunteer organization**
 - **Agenda determined by industry and member interests and participation**
 - **200+ members and growing**
- **Technical and market development efforts**
 - **Technical efforts: primarily a “feeder” organization to accredited standards organizations**



Standards Development Examples

- **Fibre Channel HBA-API**
- **MIBs and CIM extensions for storage**
 - **CIM Schemas for Disk Arrays, Media Libraries, NAS, and Fibre Channel Interconnect Components**
- **SCSI-3 Extended Copy command**
 - **Has been picked up by T10**
- **NDMP version 4**
 - **Work in progress**



Interoperability Efforts

- **Sponsored Interoperability Demonstrations**
 - HBA-API, Third party copy, FC Switches, etc.
 - Multiple venues, Multiple vendors
- **CIFS initiative**
 - Conference and interoperability demonstrations
- **SNIA Technology Center**
 - Permanent lab and test facilities, fixtures



Contents

- SNIA Introduction
- **The SNIA storage model**
 - Purpose and Benefits
 - Layers, functions, and services
- Applying the SNIA storage model
 - Common storage architectures
- Conclusions



Shared Storage Model: Purpose

- **Present a simple model for shared storage architectures**
- **Use it to describe common examples graphically**
- **Expose, for each one:**
 - **What services are provided, where**
 - **Where interoperability is required**
 - **[future] Pros and cons of the architecture**



Shared Storage Model: Benefits

- **A common “architecture vocabulary”**
- **Reference comparisons between common solutions**
- **Help to align the industry**
 - **Customers can better structure their choices**
 - **Vendors can better explain their differences**



What the model is, and what it is not

It is:


- ***A framework*** that captures the functional layers and properties of a storage system

It is not:

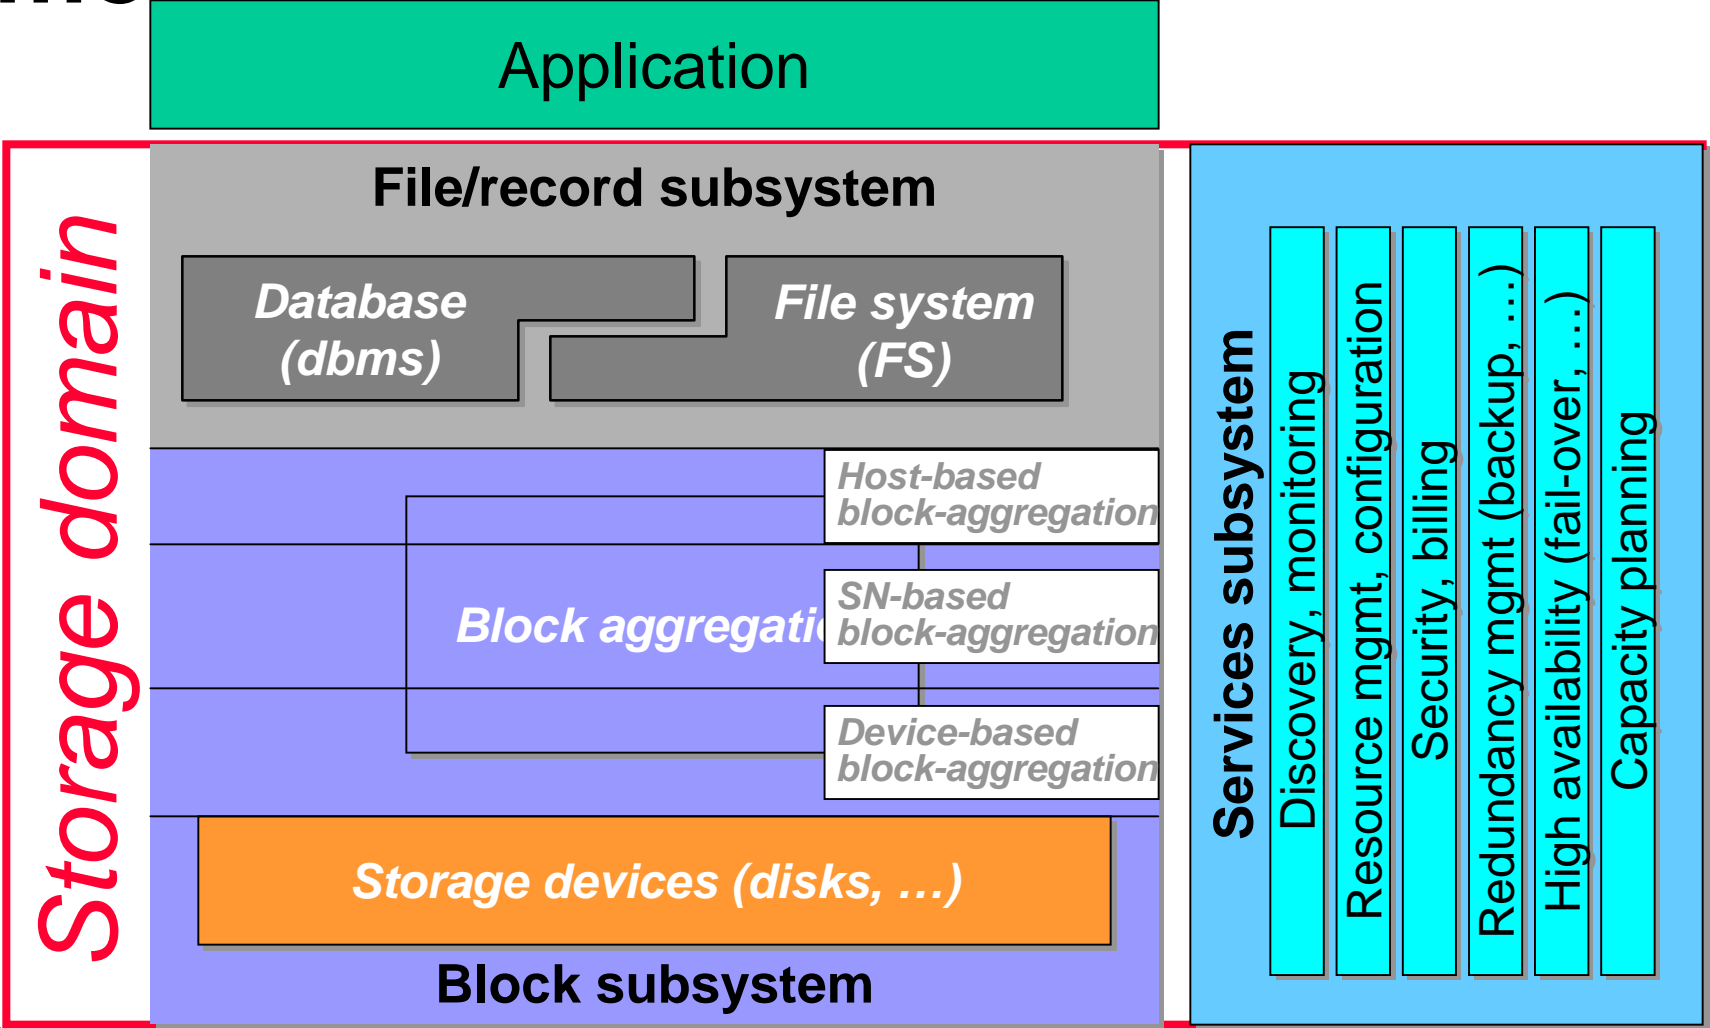
- ***A specification, an architecture, a design, a product, a recommendation, or an installation***



Classic storage model



The SNIA shared storage model



File/record subsystem

What can be done

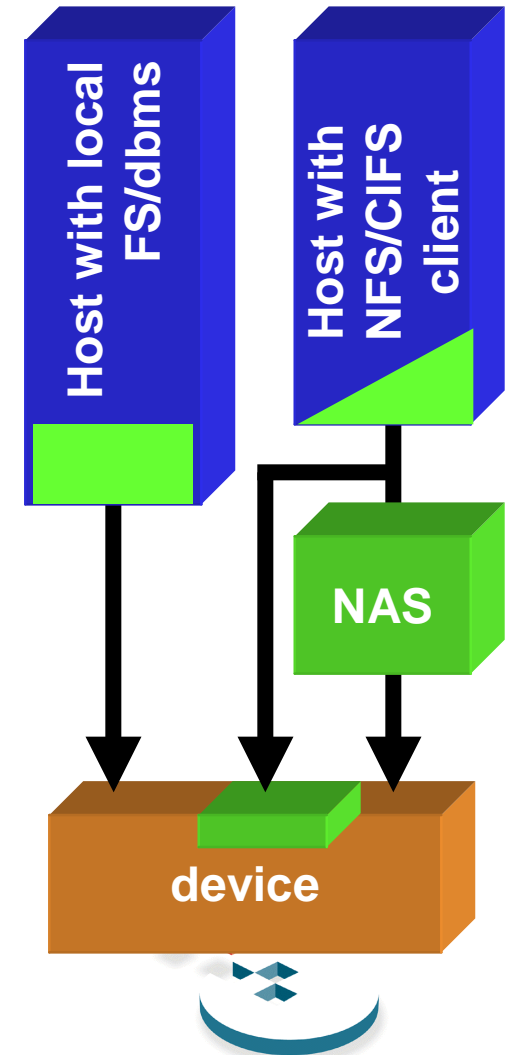
- **Database management systems**
 - tuples → tables
 - tables → table-spaces
 - table-spaces → volume
- **File systems**
 - files → volume
- **New types**
 - http caches: a kind of distributed file system?



File/record subsystem

Where it can be done

- **Host-side**
 - file systems and databases
 - NFS, CIFS, etc. are client-server splits inside the file system
- **SN-based**
 - NAS head
- **Device-based**
 - NAS functions in array box



Block subsystem

What can be done


- **Space management**
 - making a large store from many small ones
 - packing many small stores into one large one
- **Striping**
 - for performance (load balancing, throughput)
- **Redundancy**
 - full (local & remote mirroring, RAID-1, -10, ...)
 - partial (RAID-3, -4, -5, ...)
 - point-in-time copy



Block subsystem

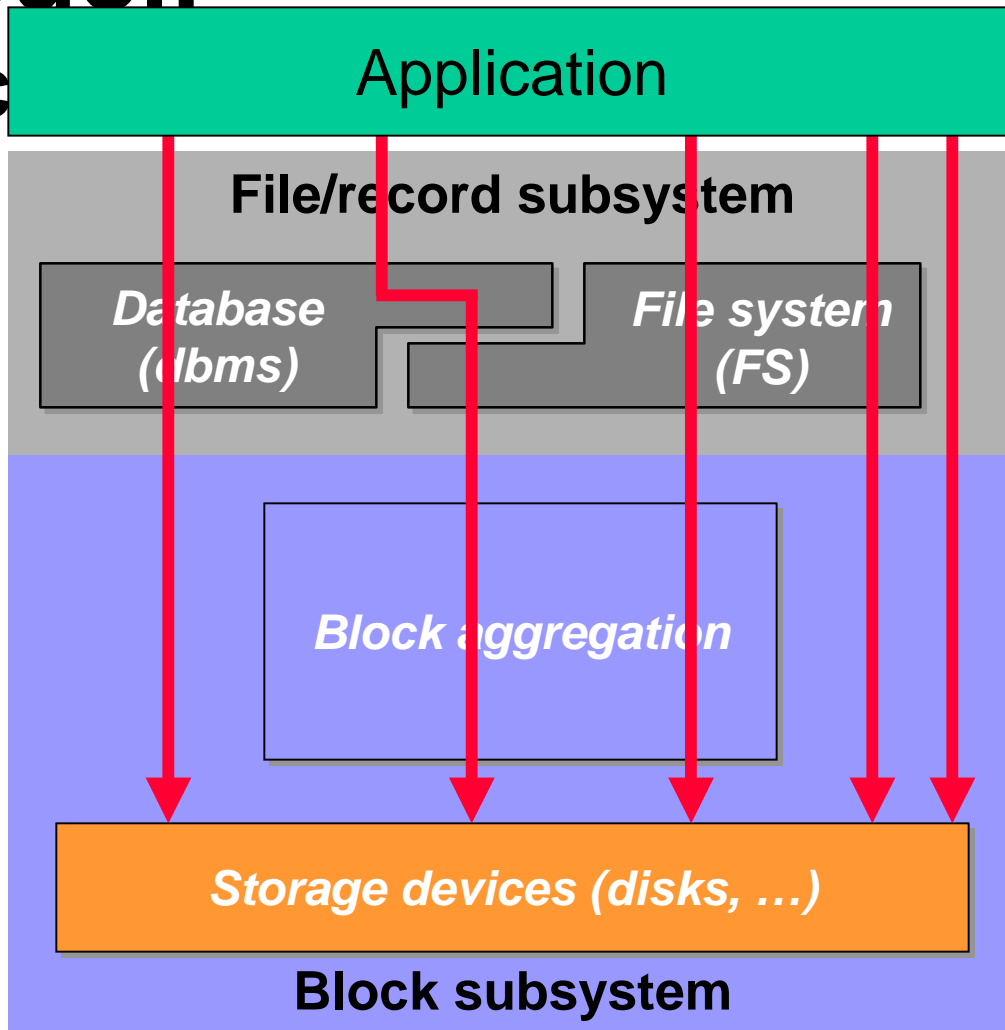
Where it can be done

- **Host-side**
 - logical volume managers
 - device drivers, HBAs
- **SN-based**
 - HBAs, specialized SN appliances
- **Device-based**
 - array controllers (e.g., RAID)
 - disk controllers (e.g., sparing)



The SNIA shared storage model:

Access

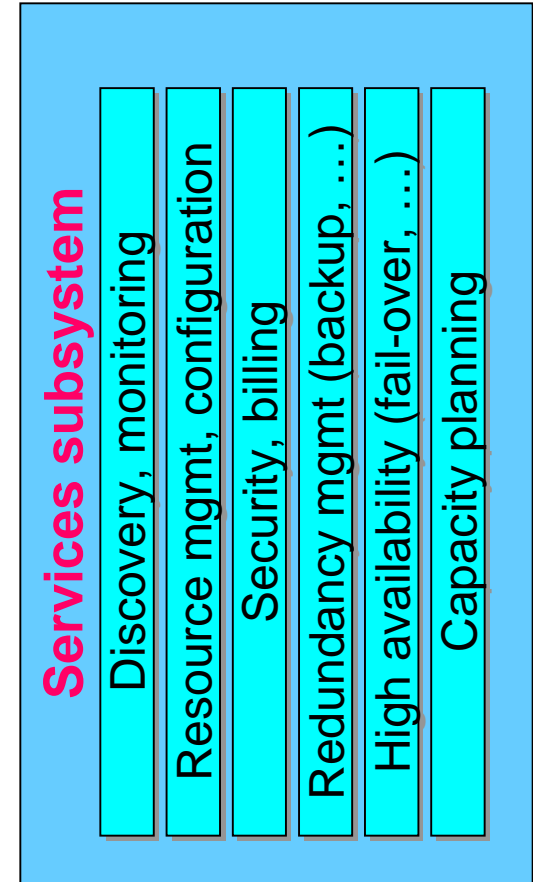
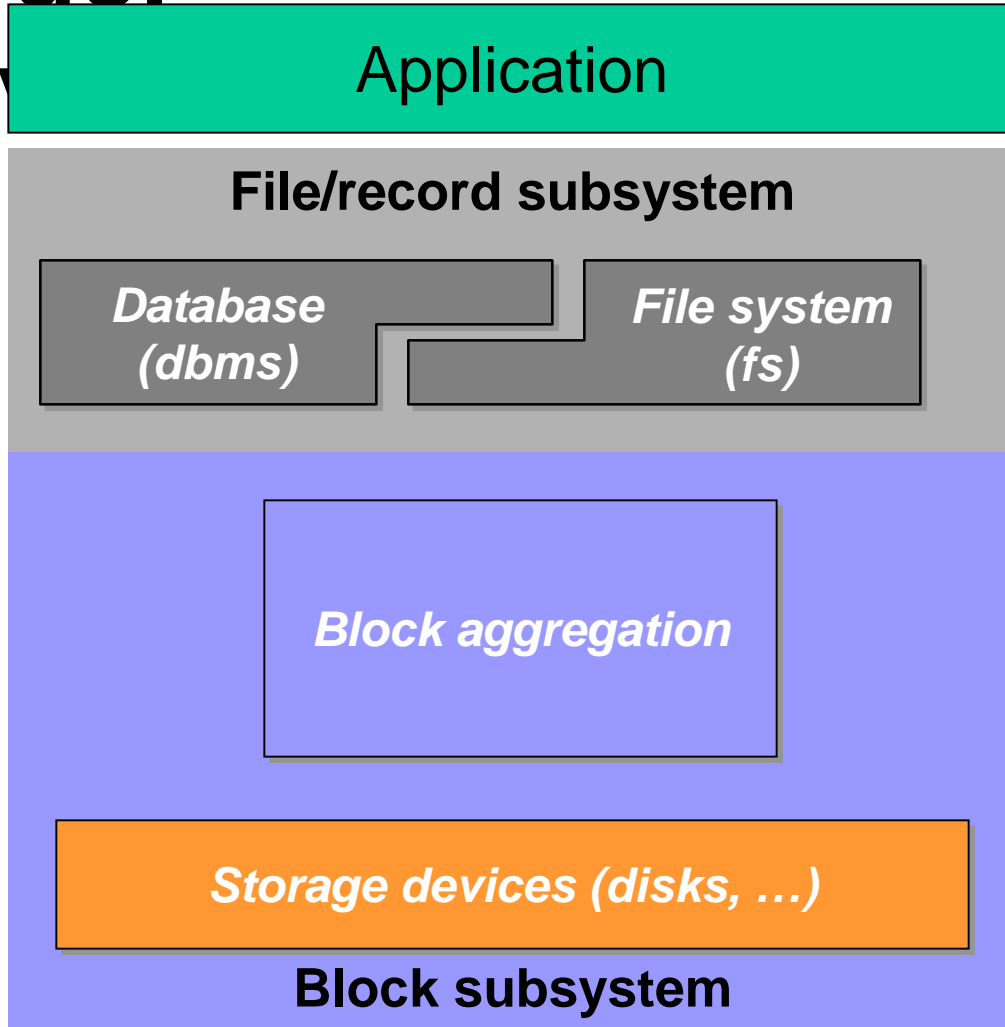


Note: all 8 possible paths can be used!



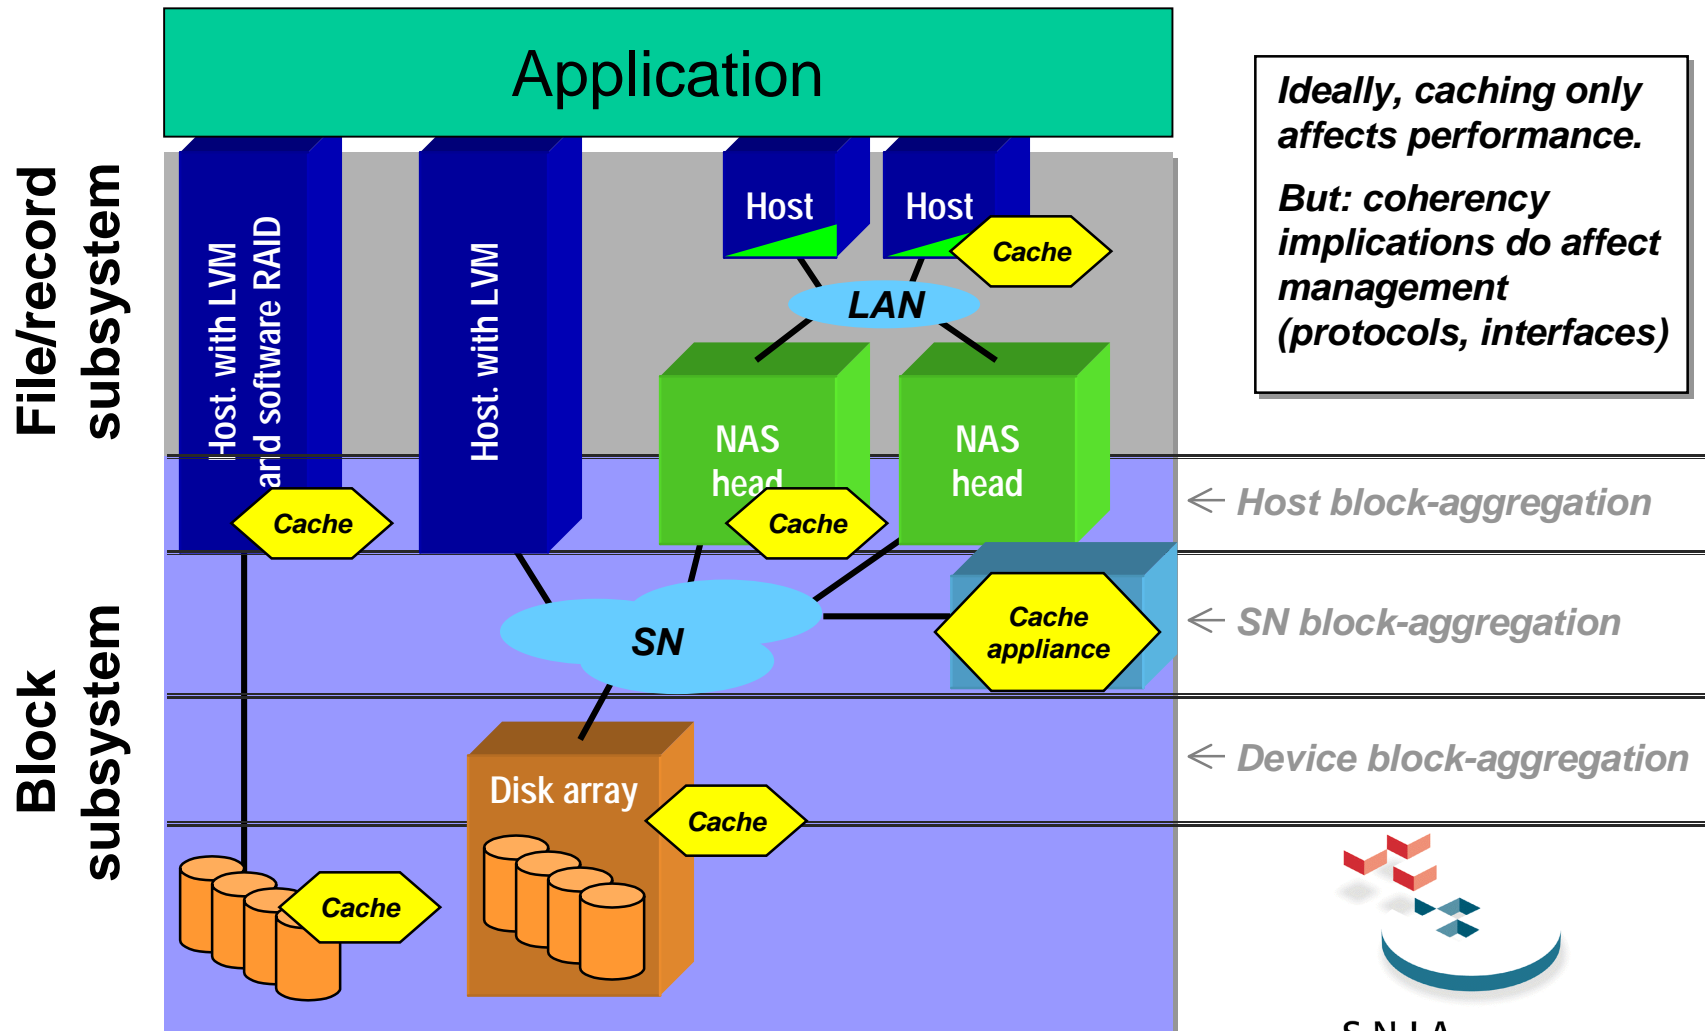
The SNIA shared storage model

Service



Caching

... can be added to almost any layer



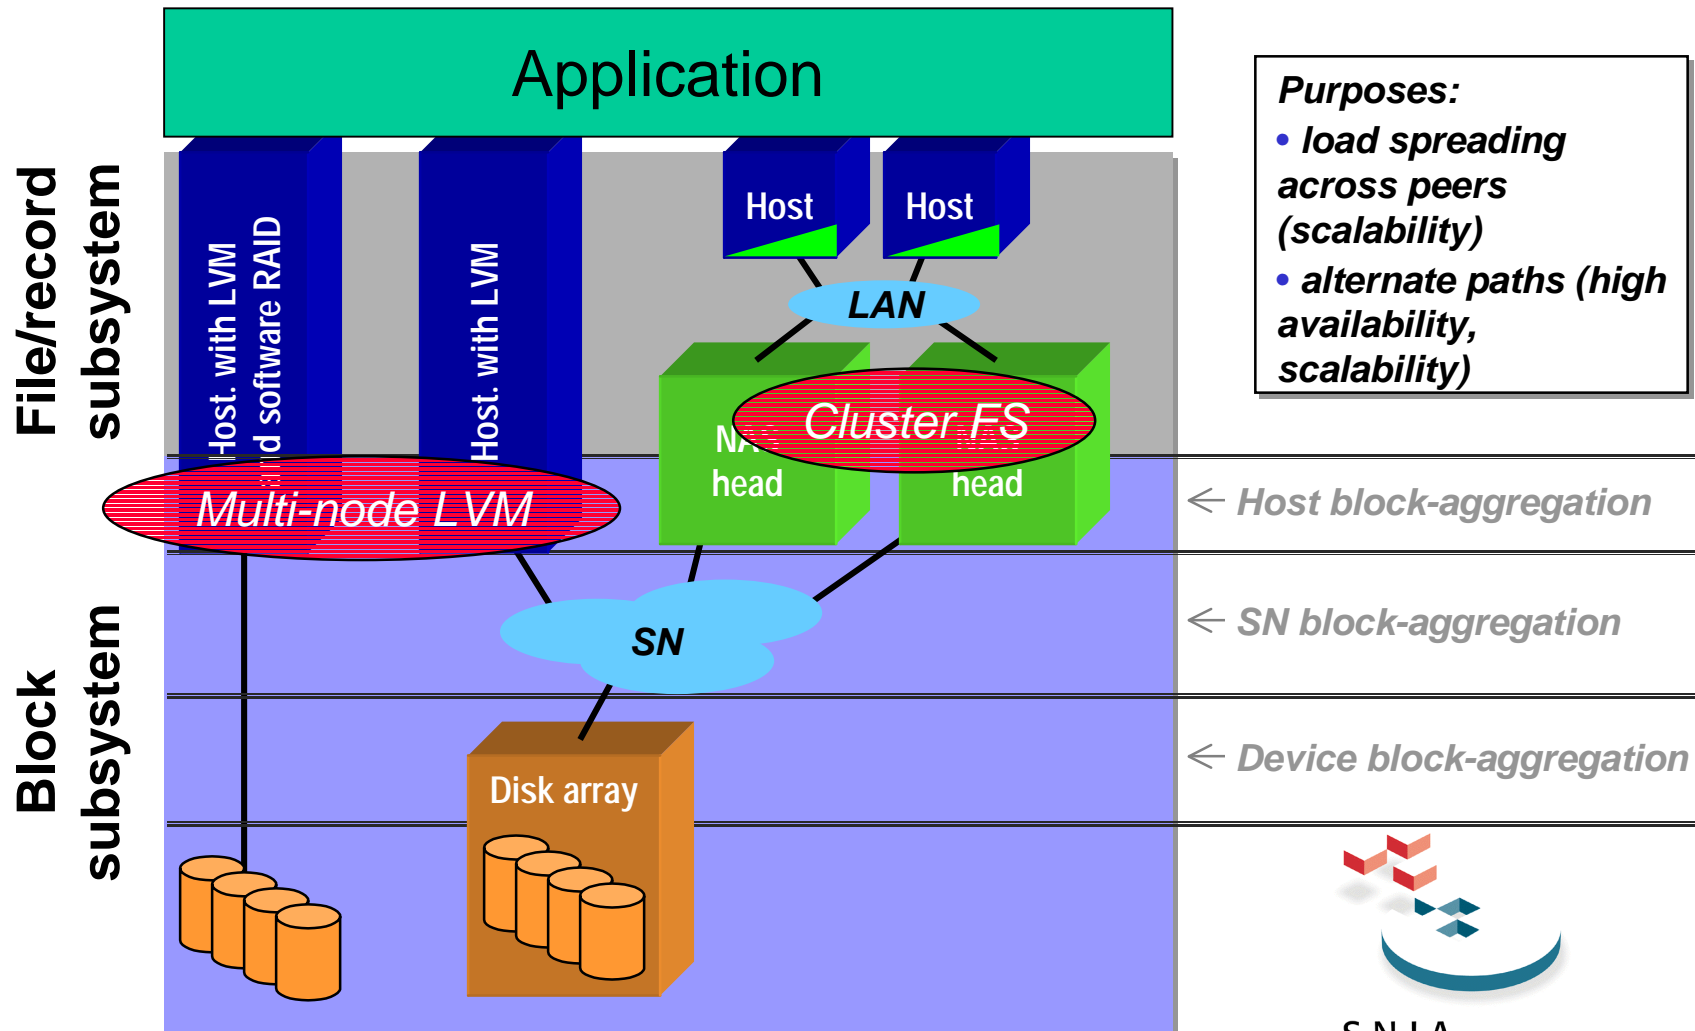
Ideally, caching only affects performance.

But: coherency implications do affect management (protocols, interfaces)



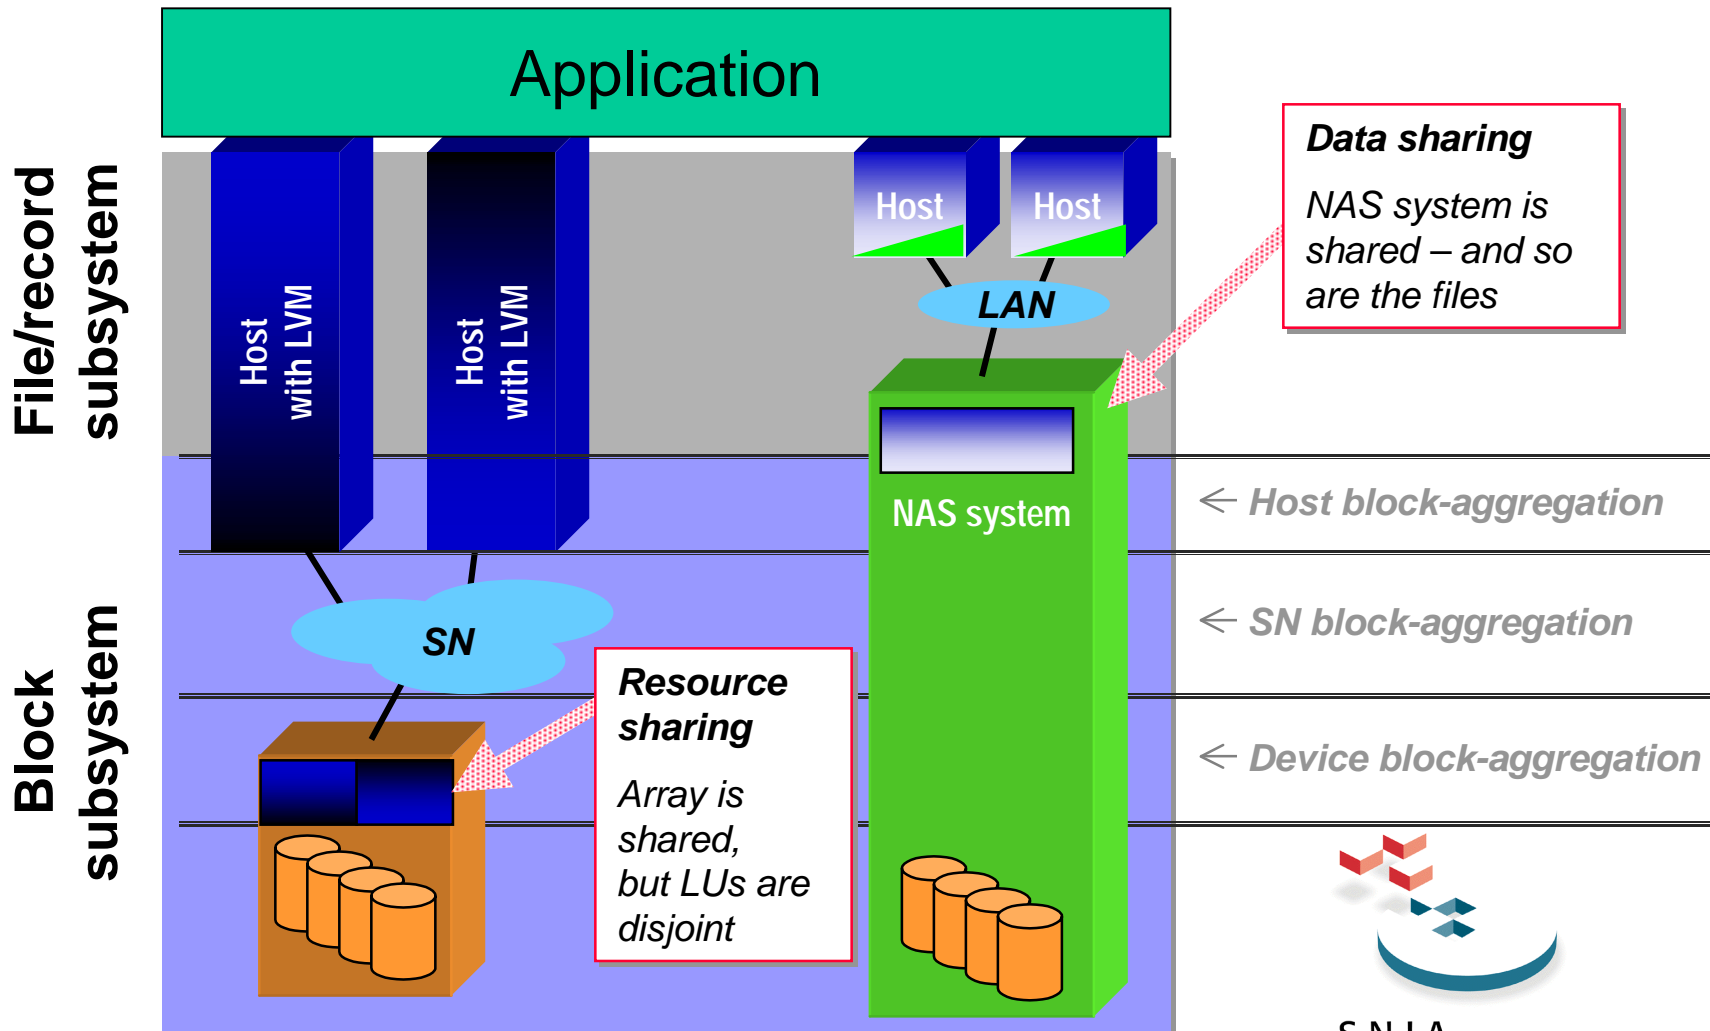
Clustering

Intra-box aggregation



Sharing


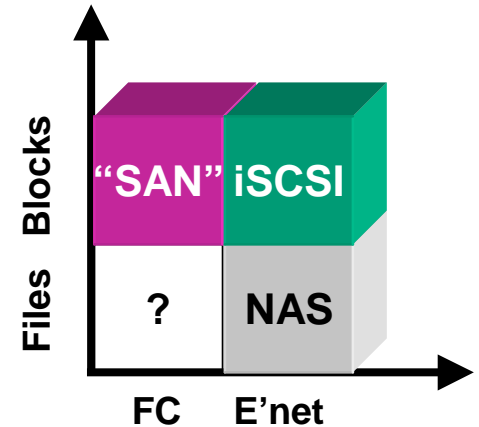
Content sharing and resource sharing



Q: “SAN” versus “NAS”?

A: a poorly-formed question

- Q: hardware: FibreChannel vs Ethernet vs InfiniBand?
- Q: API: blocks vs files (aka “NAS”) vs objects (OSD)?
- Q: protocol: FCP vs TCP/IP vs ... ?
- A: (to all the above) *it depends* ...
- **Storage network (SN):**
 - any (mostly) dedicated network, installed (mostly) for storage traffic
 - whatever the hardware, API, or protocol

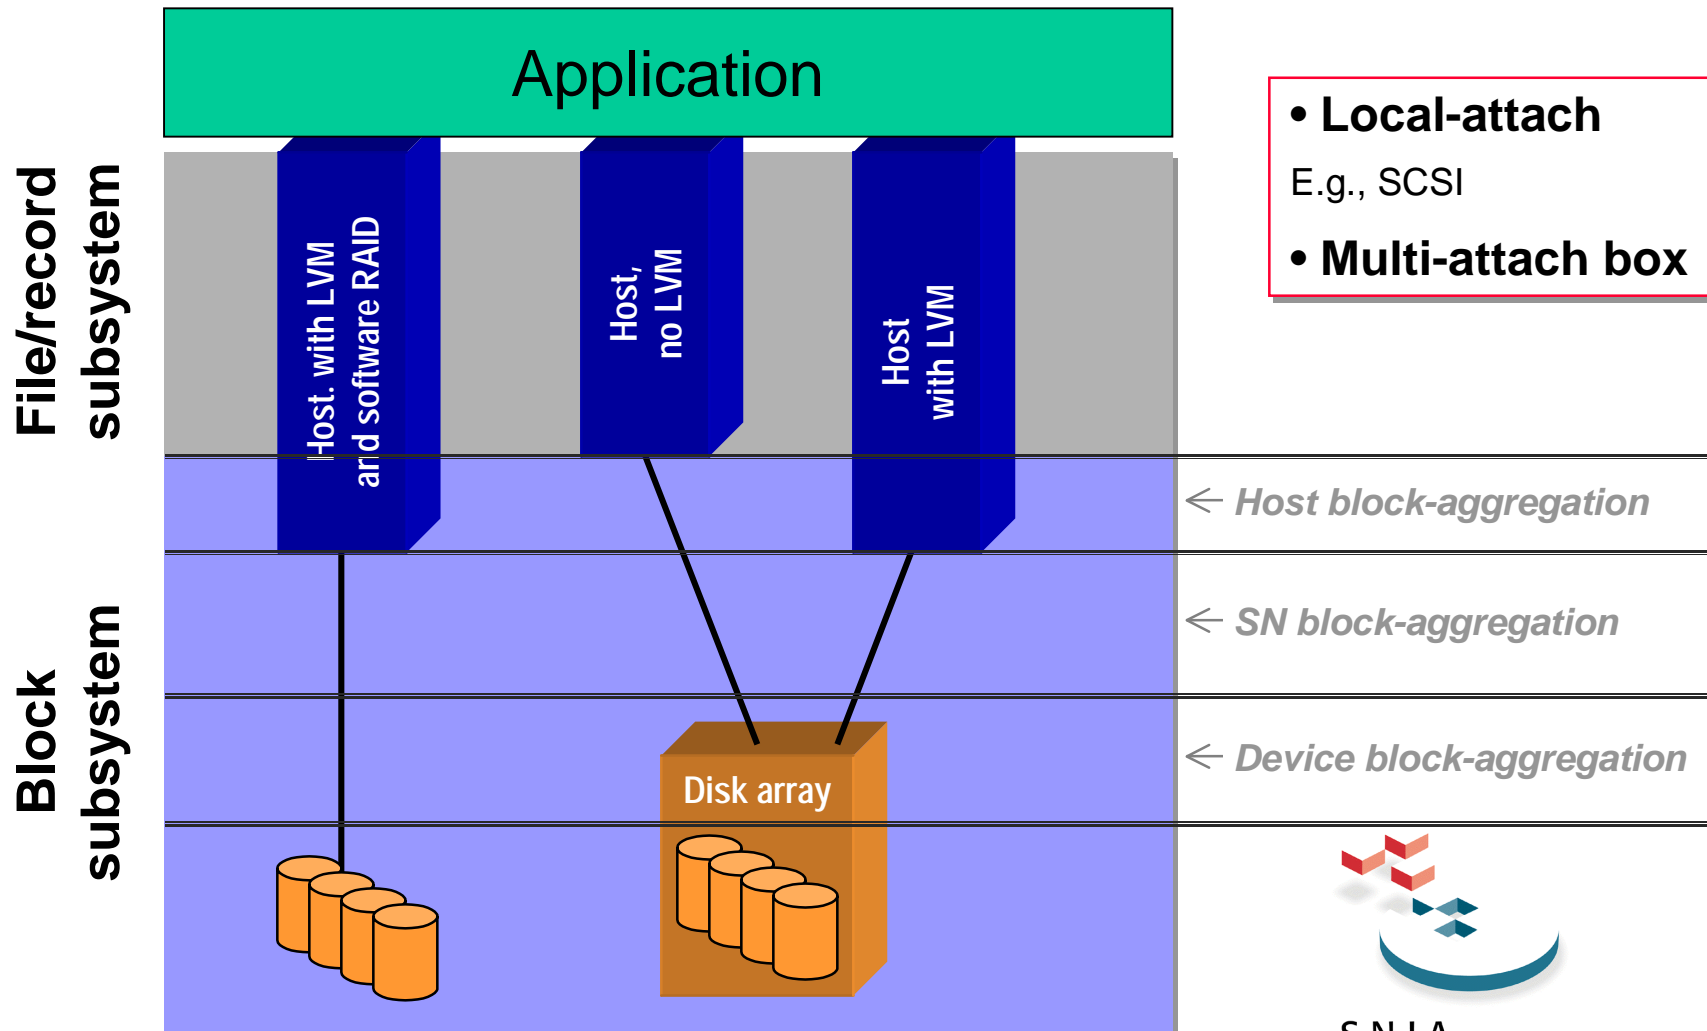


Some common storage architectures



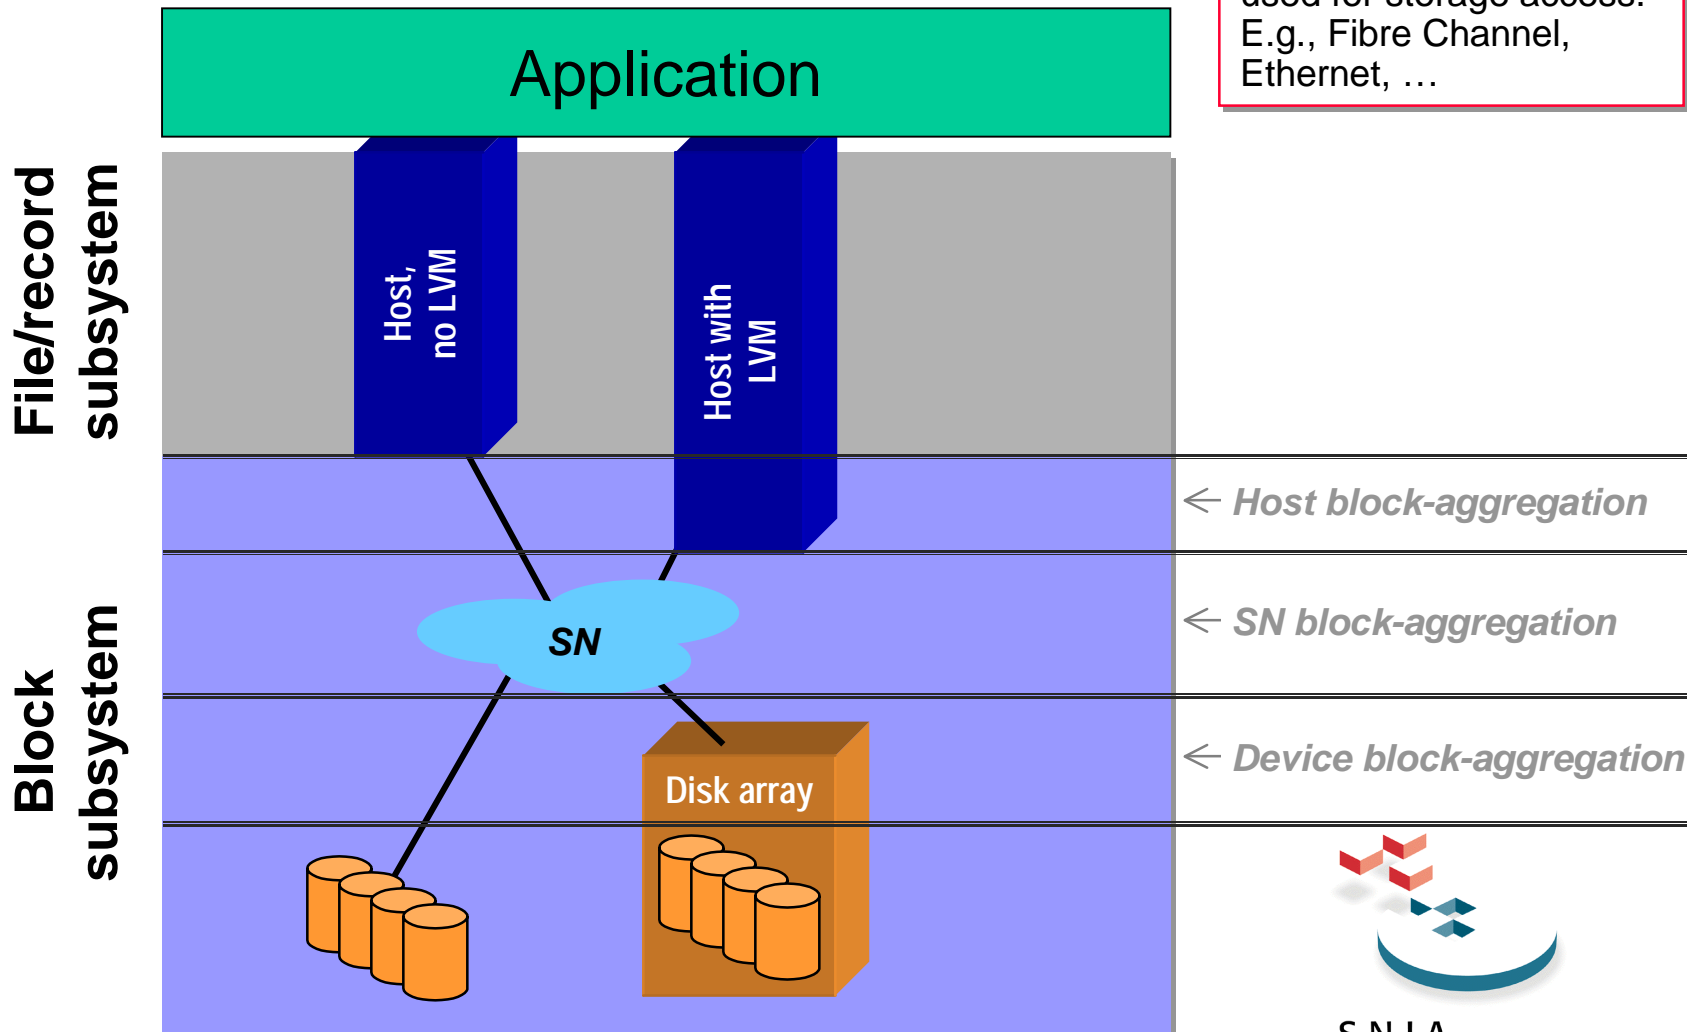
Mapping the SNIA model onto some current implementations

Direct-attached block storage

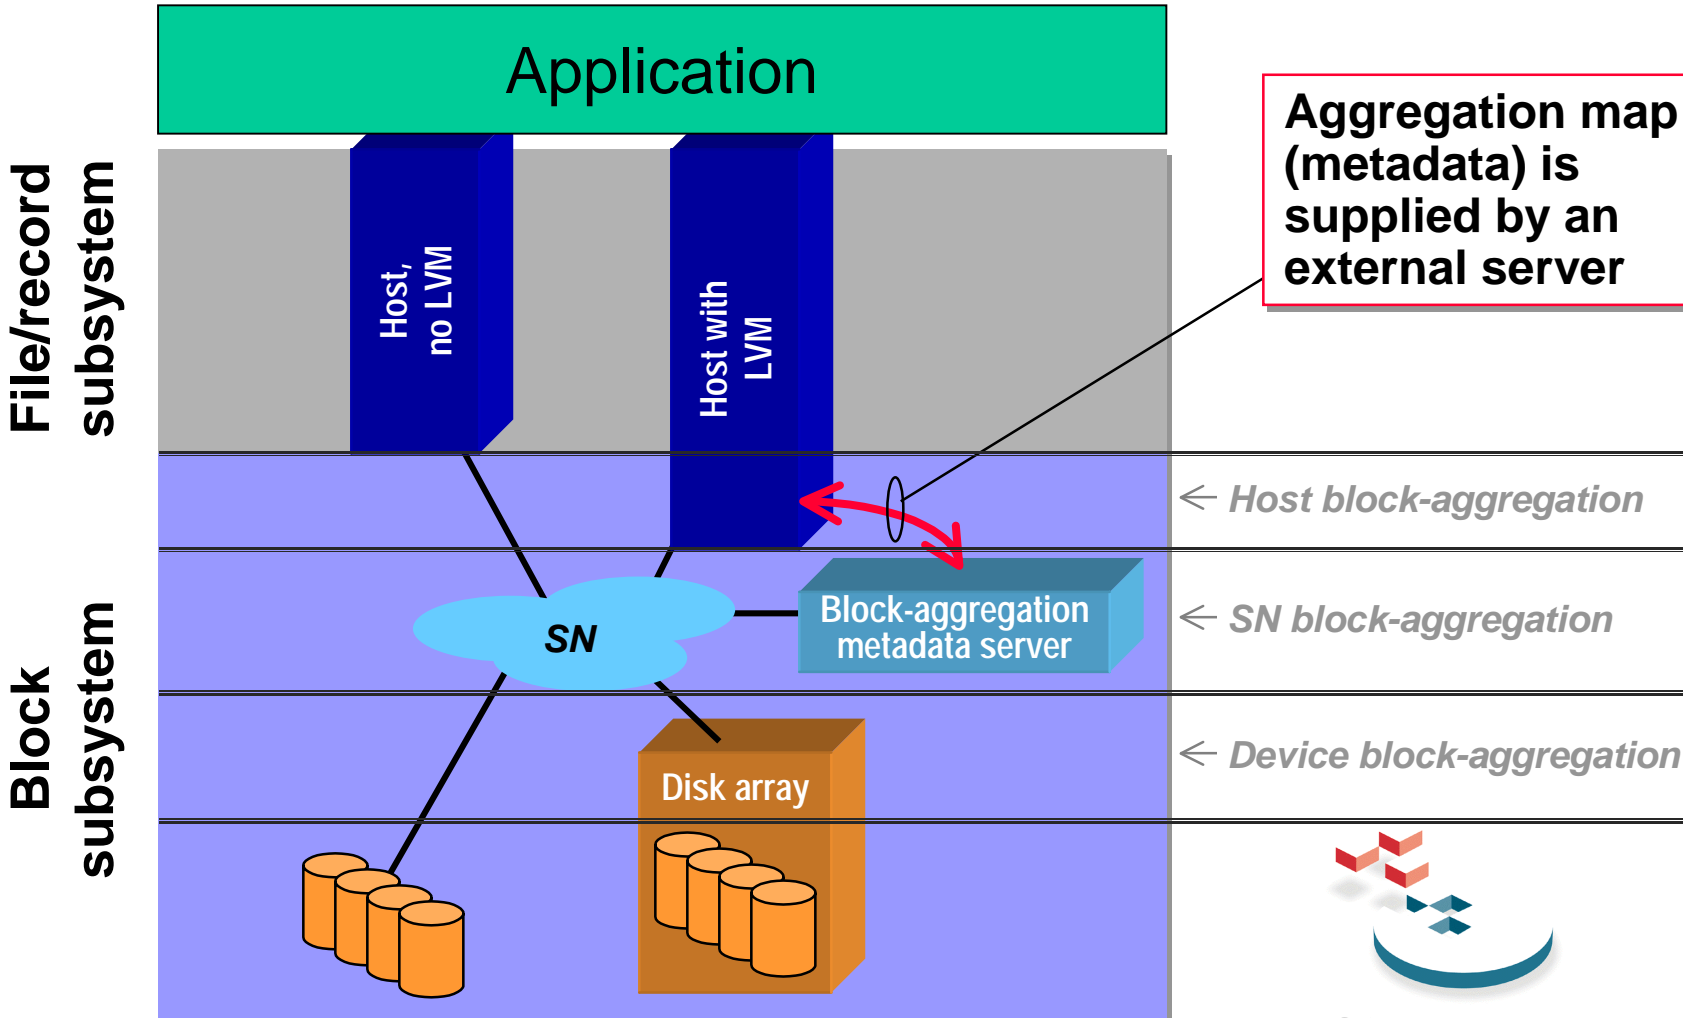


SN-attached block storage

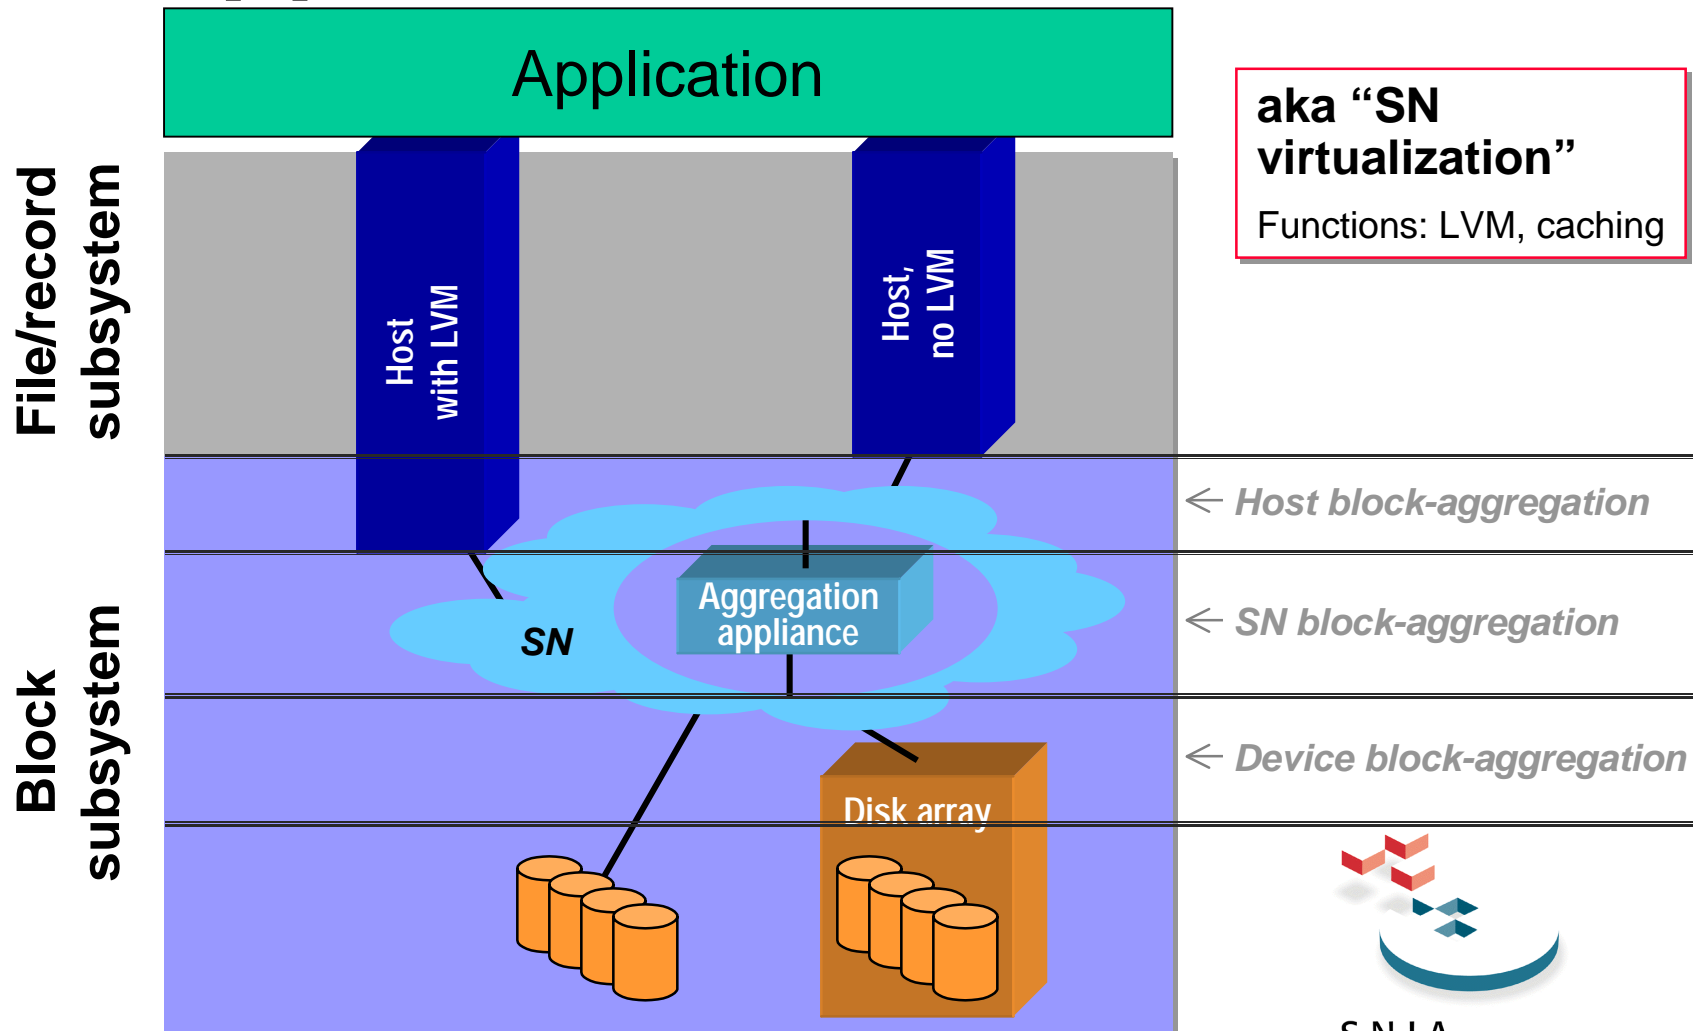
“SN” = any network used for storage access.
E.g., Fibre Channel,
Ethernet, ...



SN-attached block storage with metadata server

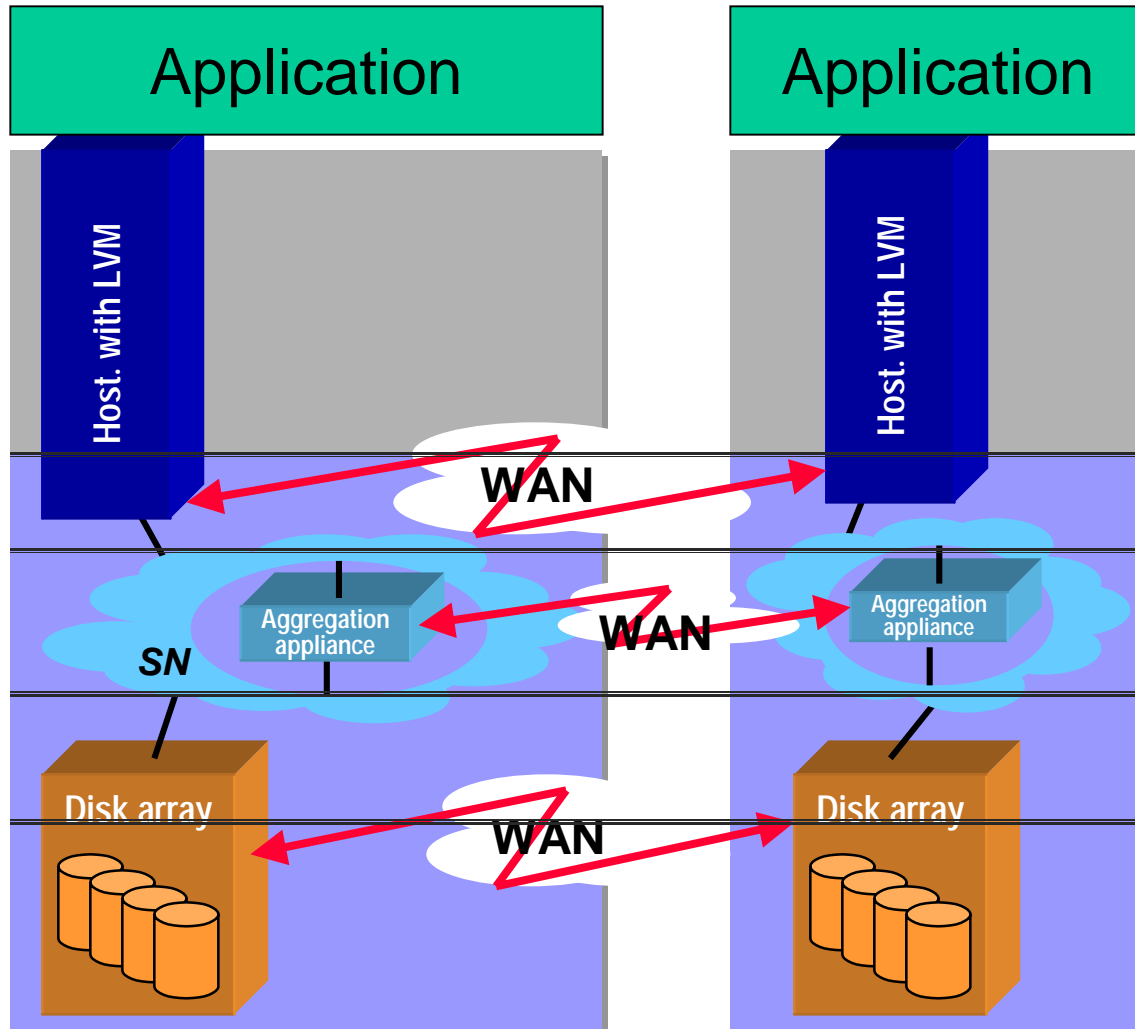


Block storage aggregation in a SN appliance



Multi-site block storage

Functions: point-in-time copy, caching, local & remote mirroring, ...



Host-to-host

SN appliance

Device-to-device

← *Host block-aggregation*

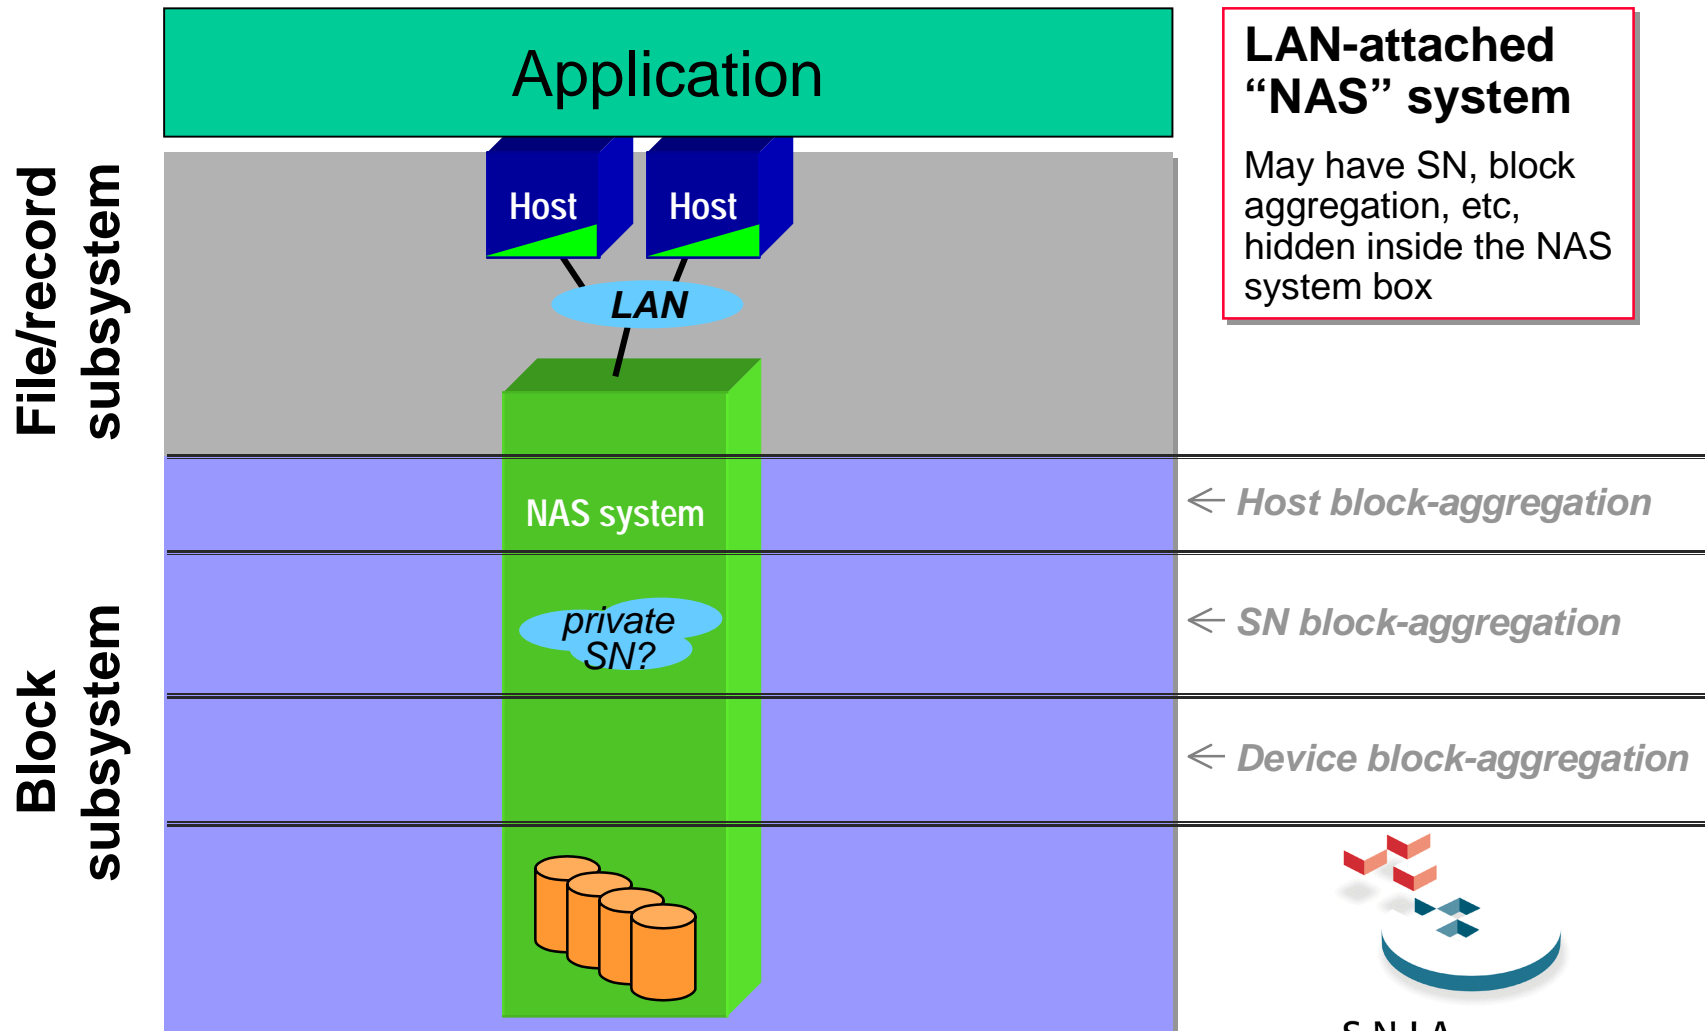
← *SN block-aggregation*

← *Device block-aggregation*



SNIA
Storage Networking Industry Association

File server

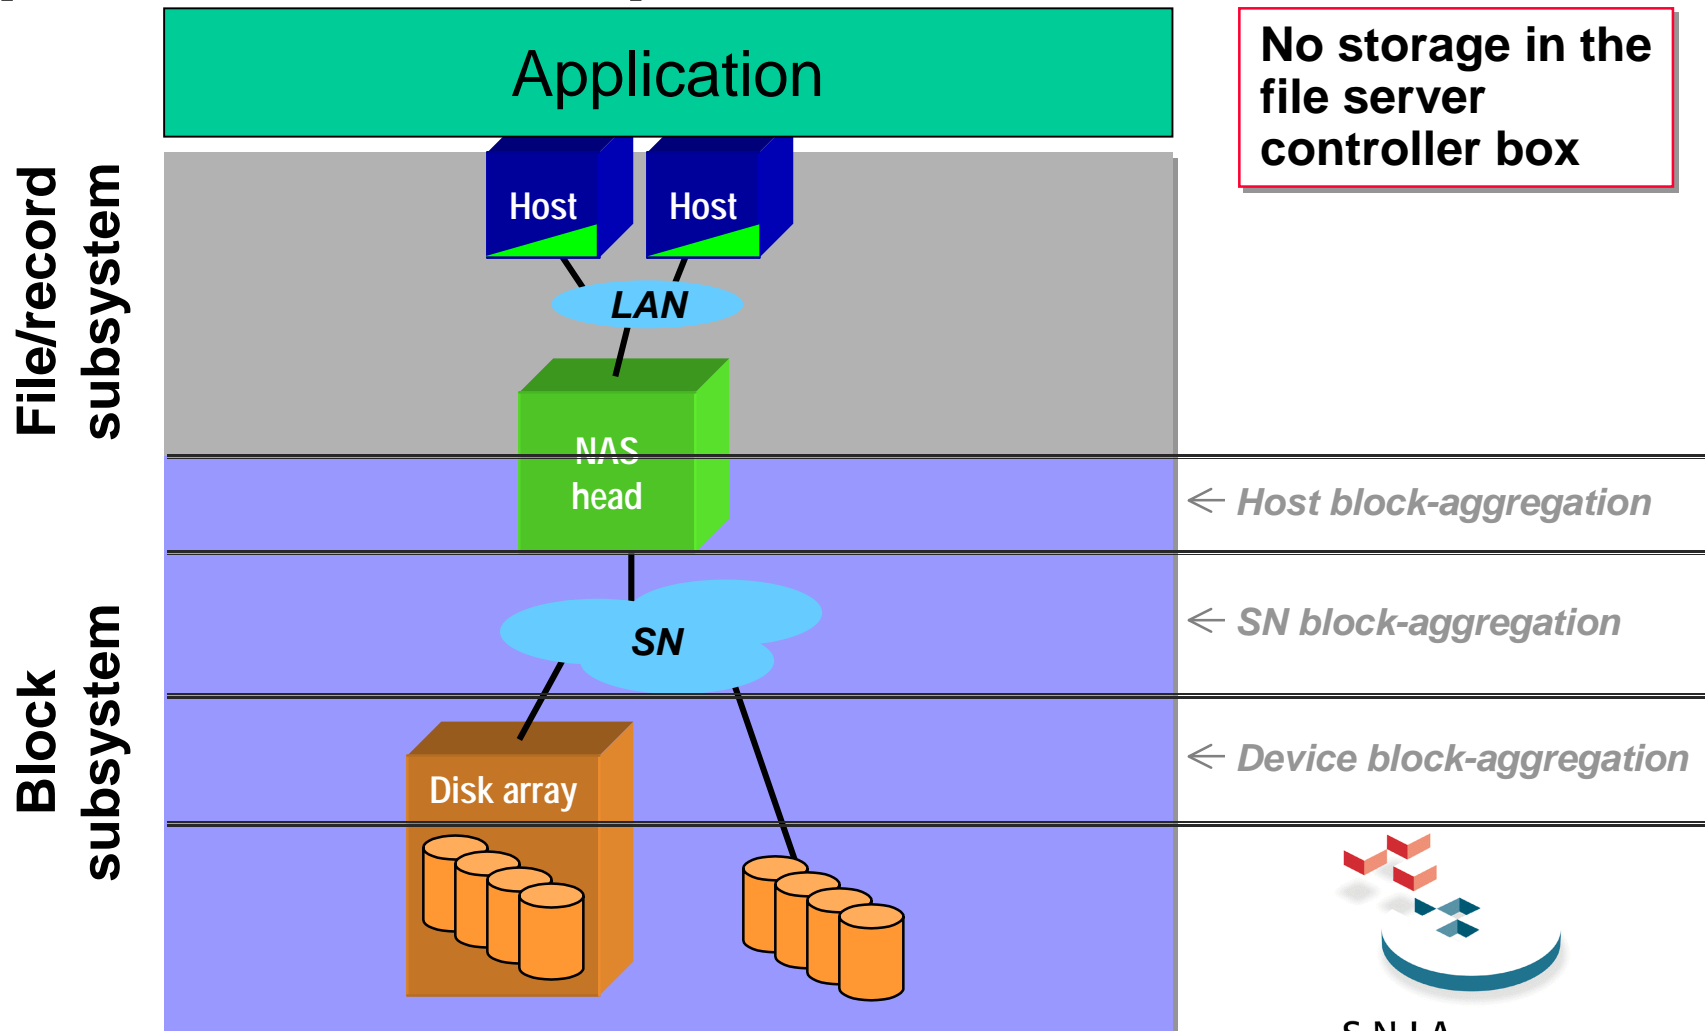


**LAN-attached
"NAS" system**

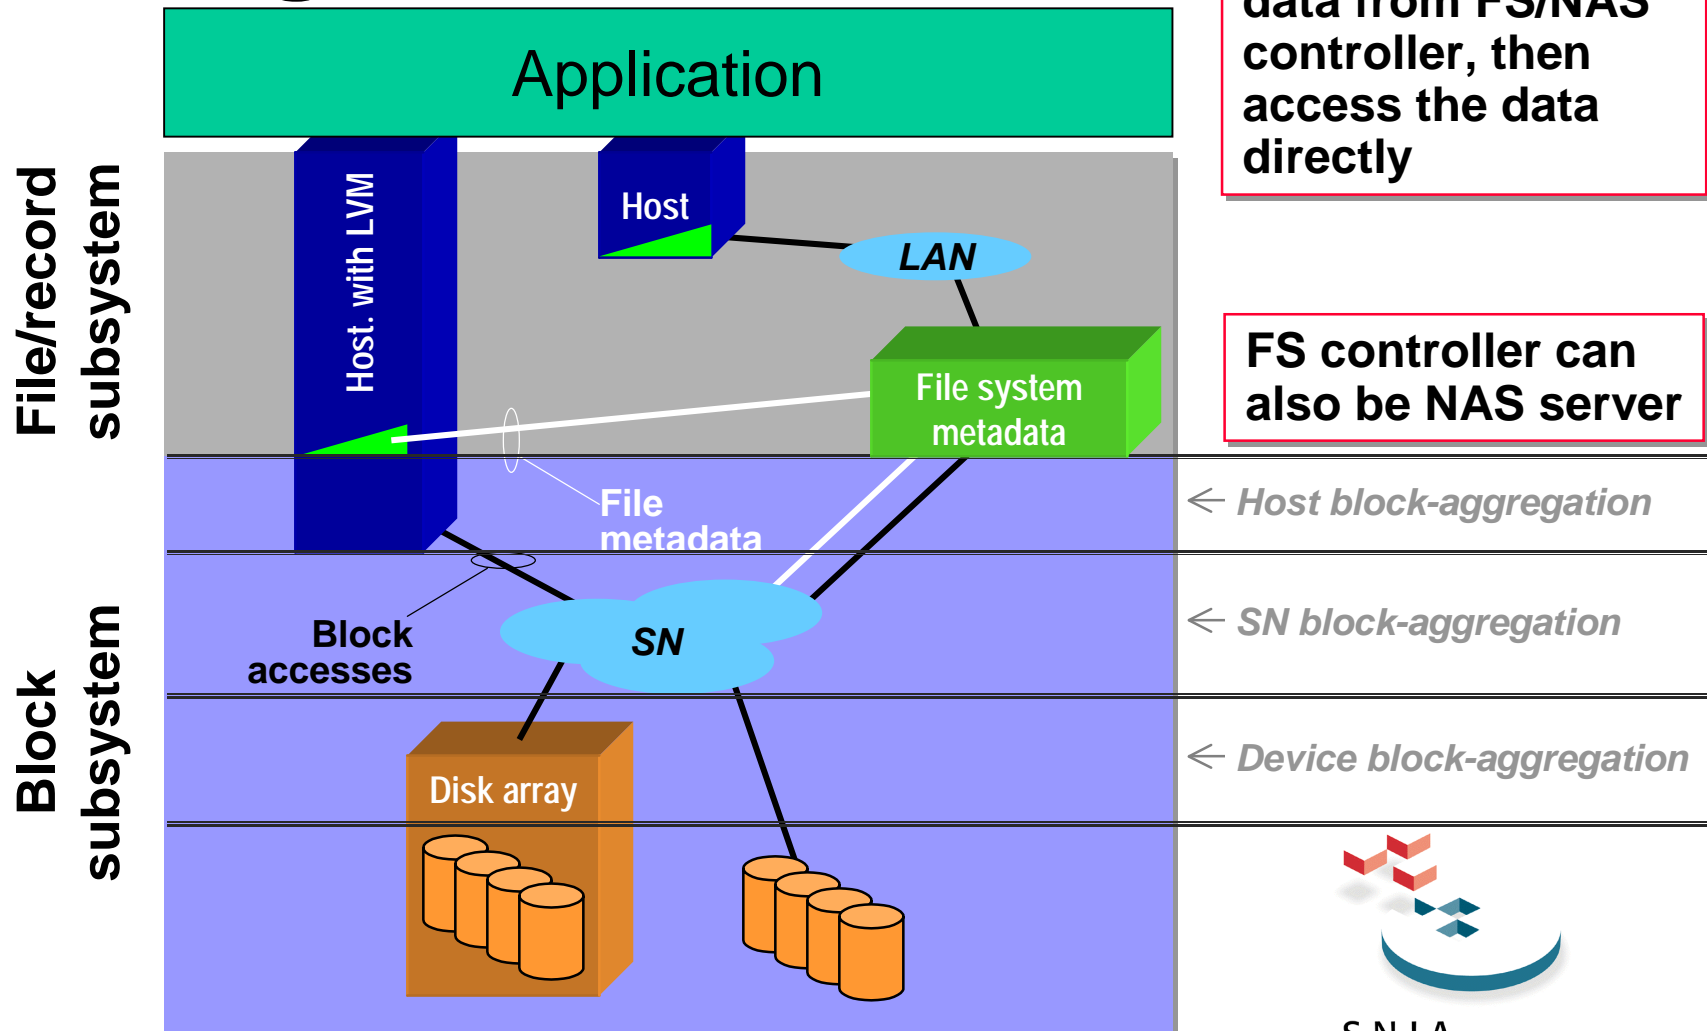
May have SN, block aggregation, etc, hidden inside the NAS system box



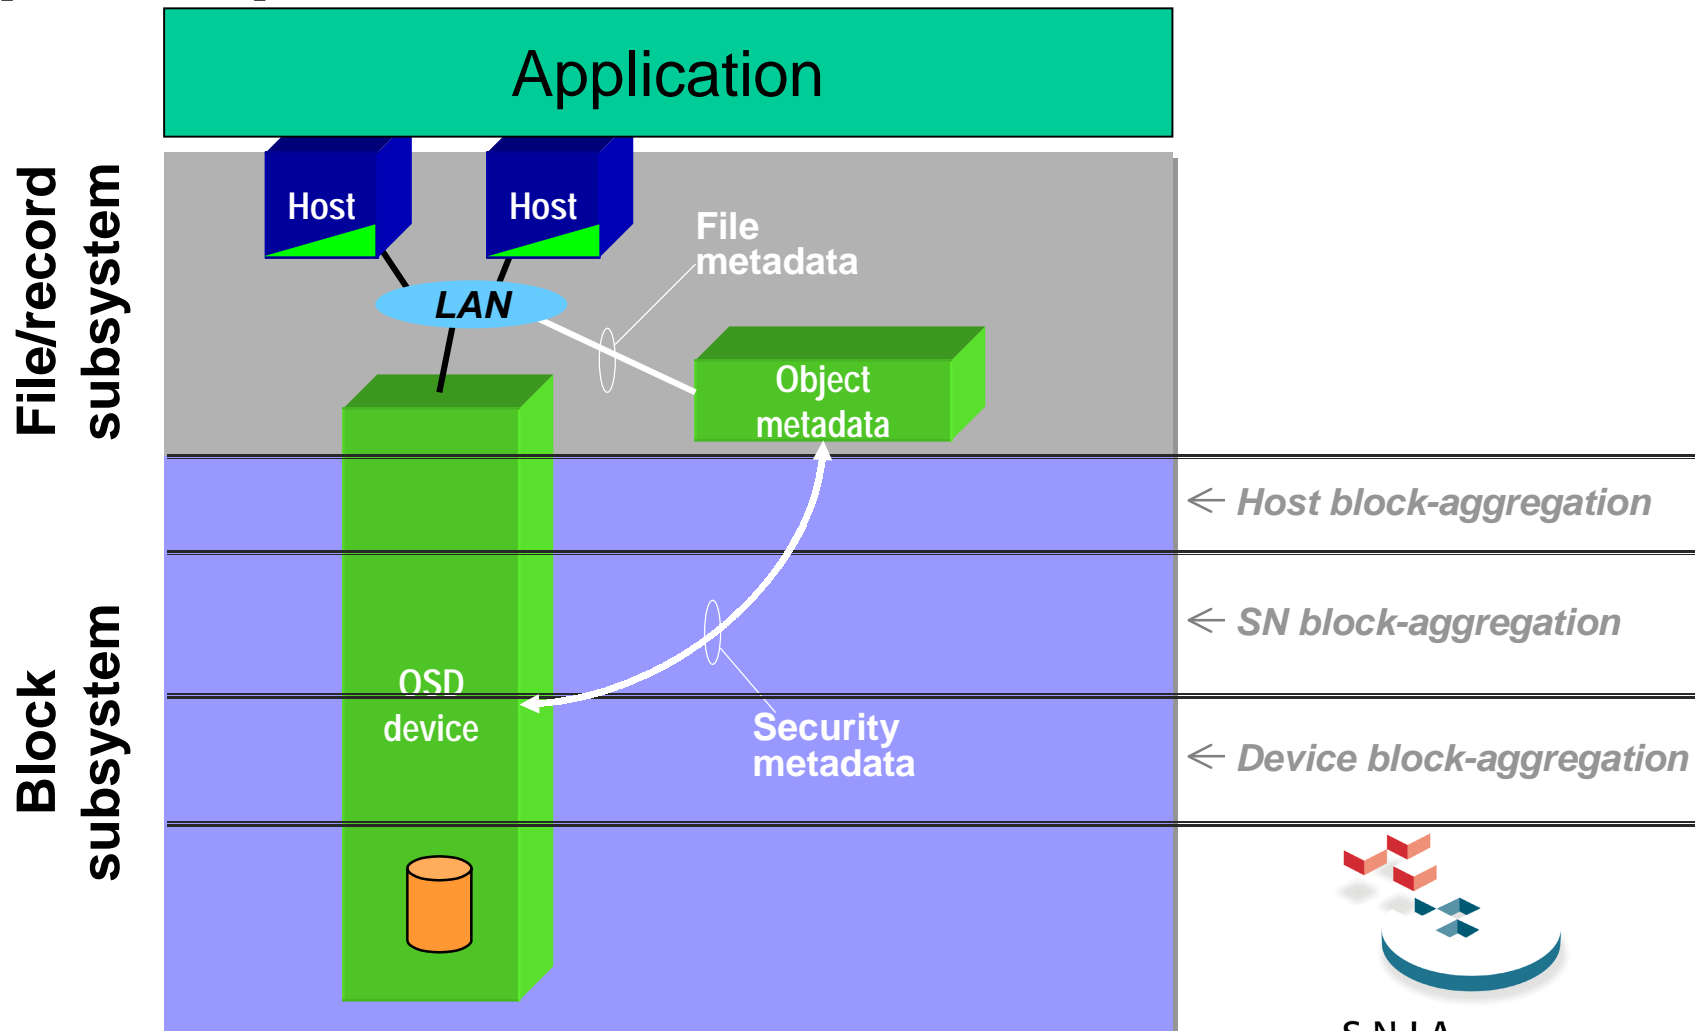
File server controller ("NAS head")




NAS/file server metadata manager (“asymmetric”)



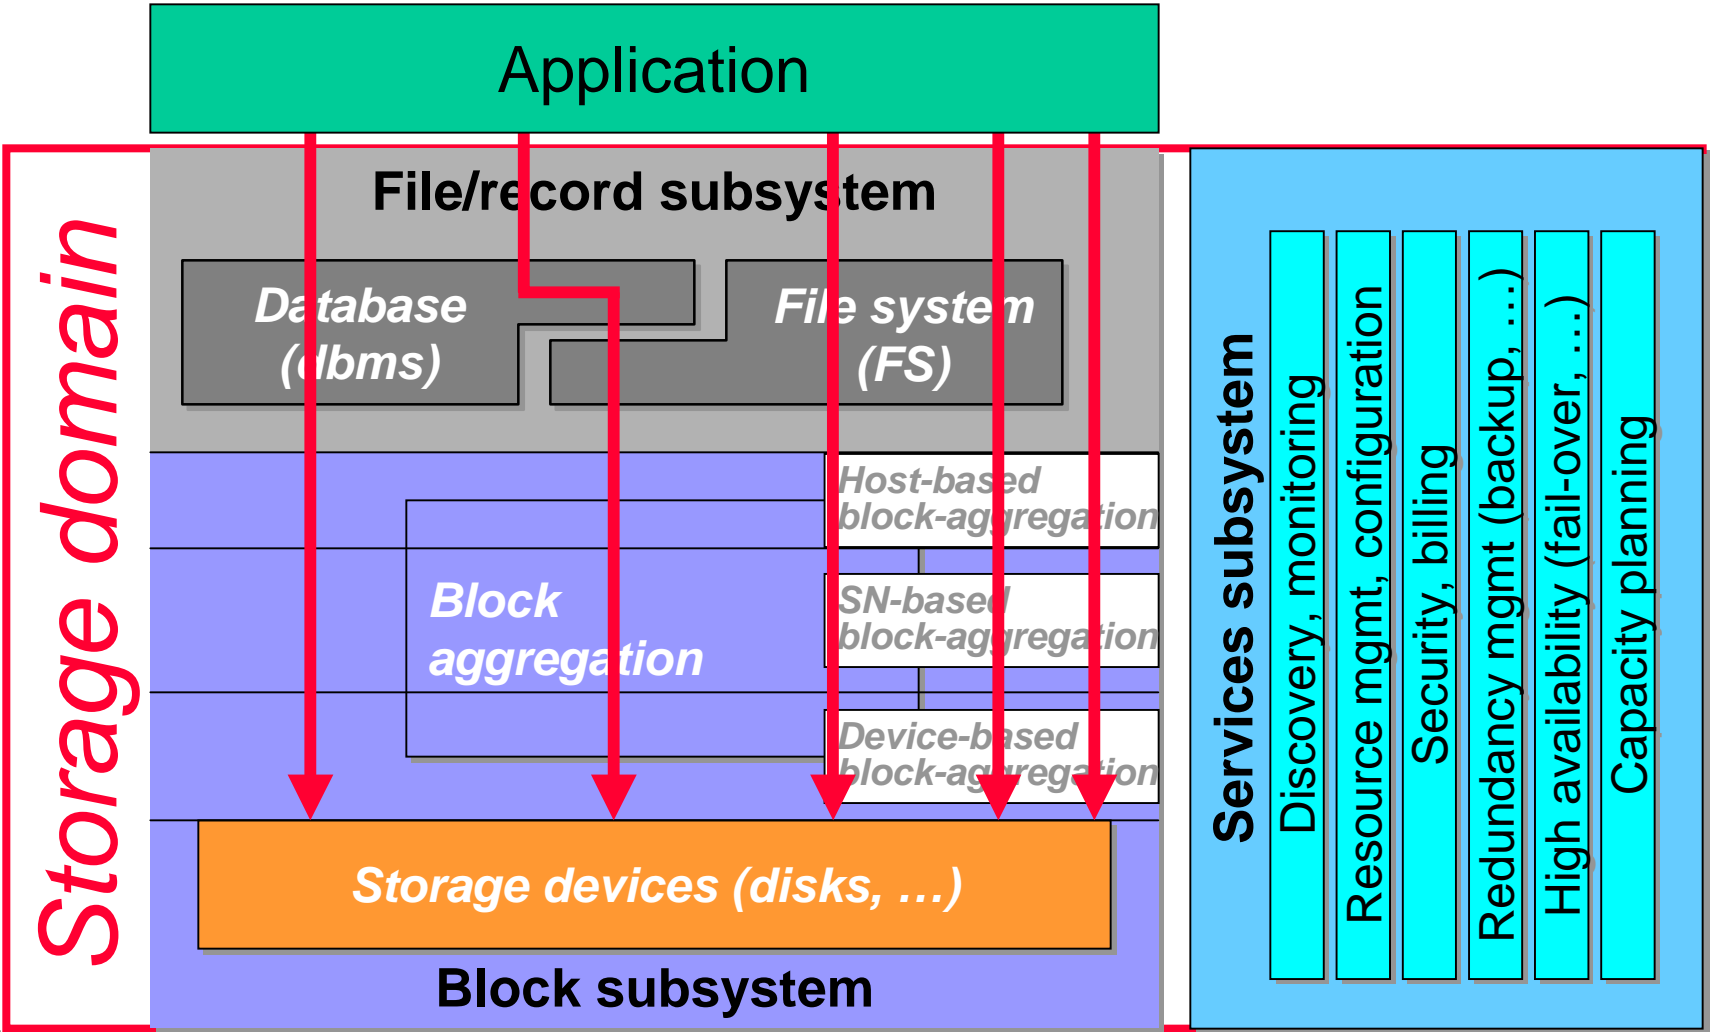
Object-based Storage Device (OSD), CMU NASD



Summary & conclusions

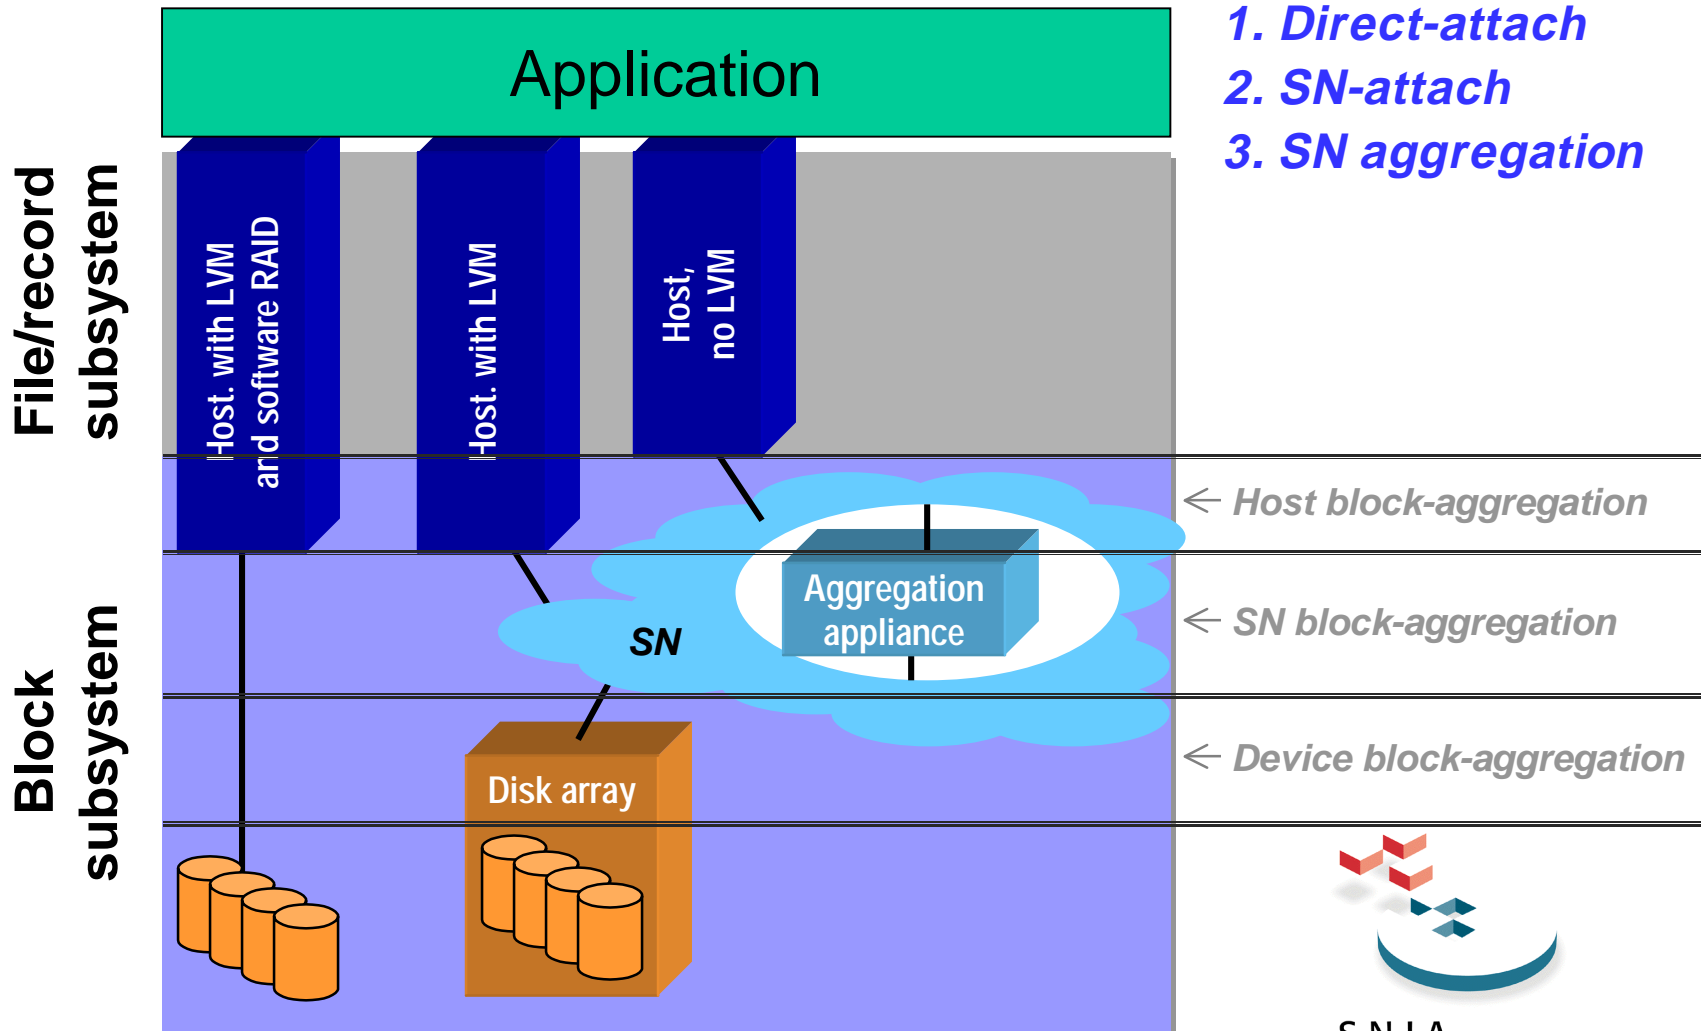


The SNIA storage model



Block subsystem

Sample architectures

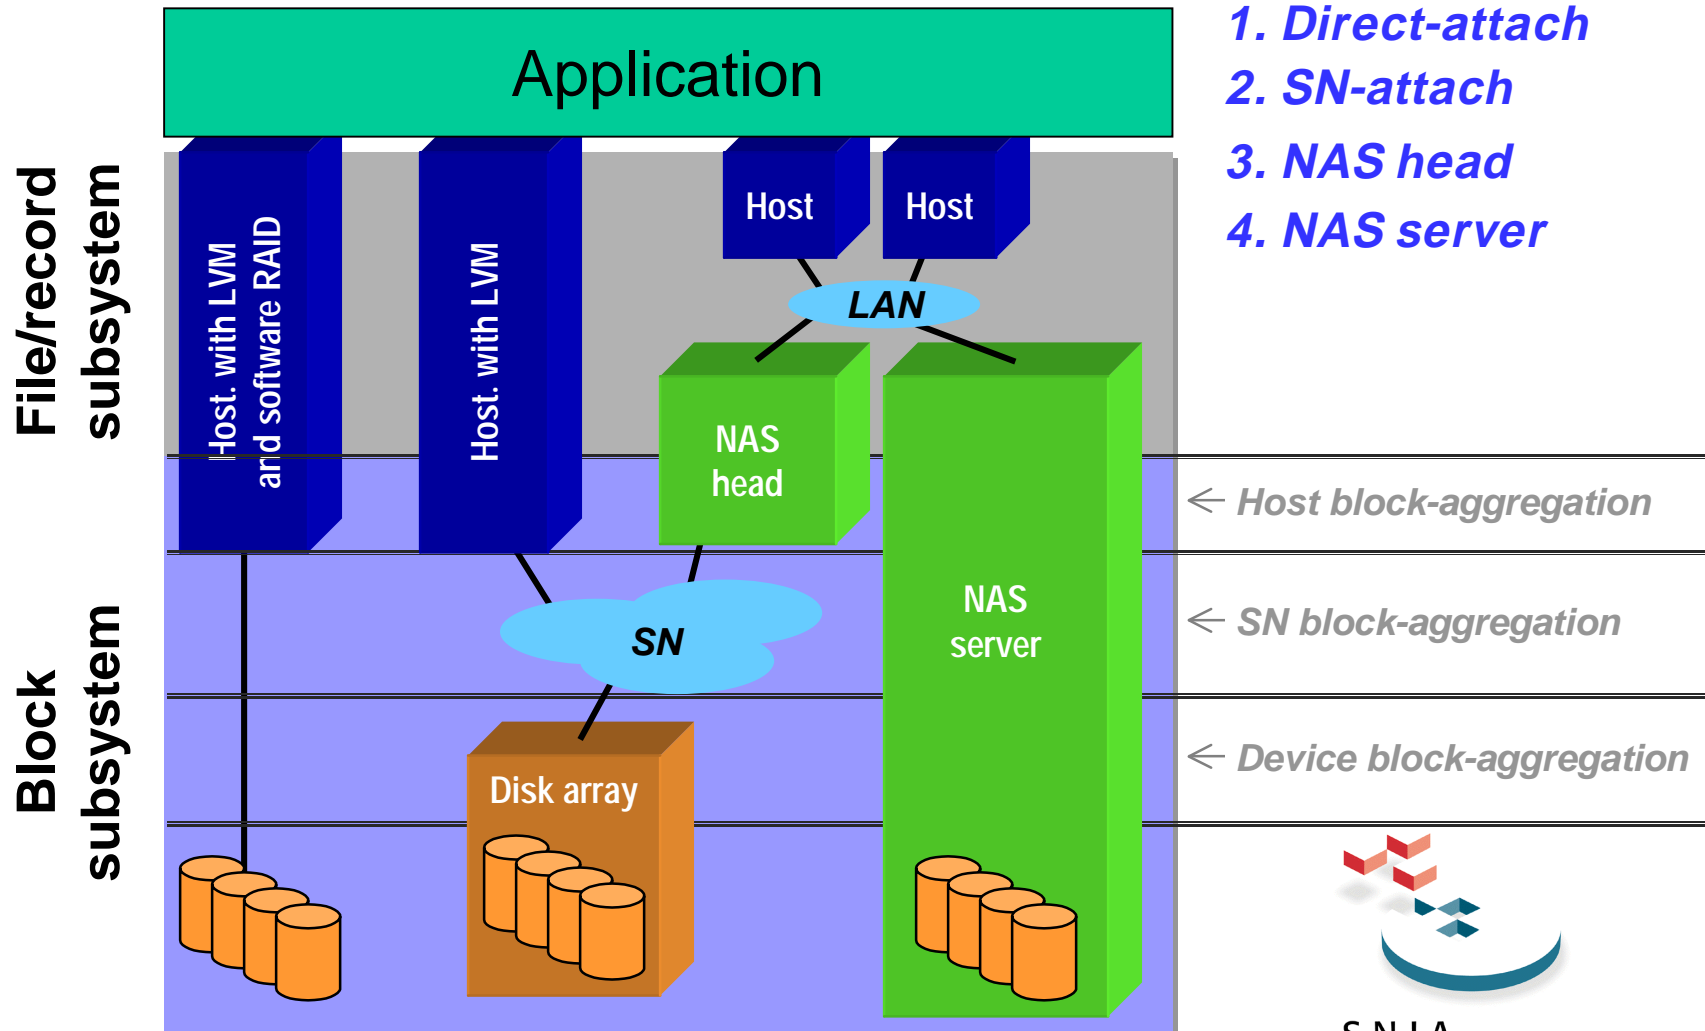


1. Direct-attach
2. SN-attach
3. SN aggregation



File/record subsystem

Sample architectures



1. *Direct-attach*
2. *SN-attach*
3. *NAS head*
4. *NAS server*



Uses for the model

- **Vendors**
 - place products in the space of architectures
 - clarify product differences
- **Customers**
 - understand vendor offerings better
- **The industry**
 - basis for common definitions, communication, understanding, interoperability



Conclusions

- **The SNIA shared storage model is both simple and useful**
 - to highlight similarities and differences
 - as a basis for comparisons
- **Still a work in progress**
 - data movers, tape drives, ...
 - better comparisons ...
 - suggestions?
- **The SNIA-TC welcomes input:**
 - <snia-tc@snia.org>



Acknowledgements and co-conspirators

- **John Wilkes, Hewlett-Packard Laboratories**
 - **Lead conspirator and tamer of PowerPoint**
- **Harald Skardal, Network Appliance**
- **David L. Black, EMC Corporation**
- **The rest of the SNIA Technical Council**
 - **Wayne Rickard (SNIA TC chairperson), Gadzoox**
 - **Dave Anderson, Seagate Technology**
 - **Jim Carlson, IBM**
 - **Garth Gibson, CMU**
 - **Kevin Phaup, HighGround Systems**
 - **David Thiel, Compaq**

