

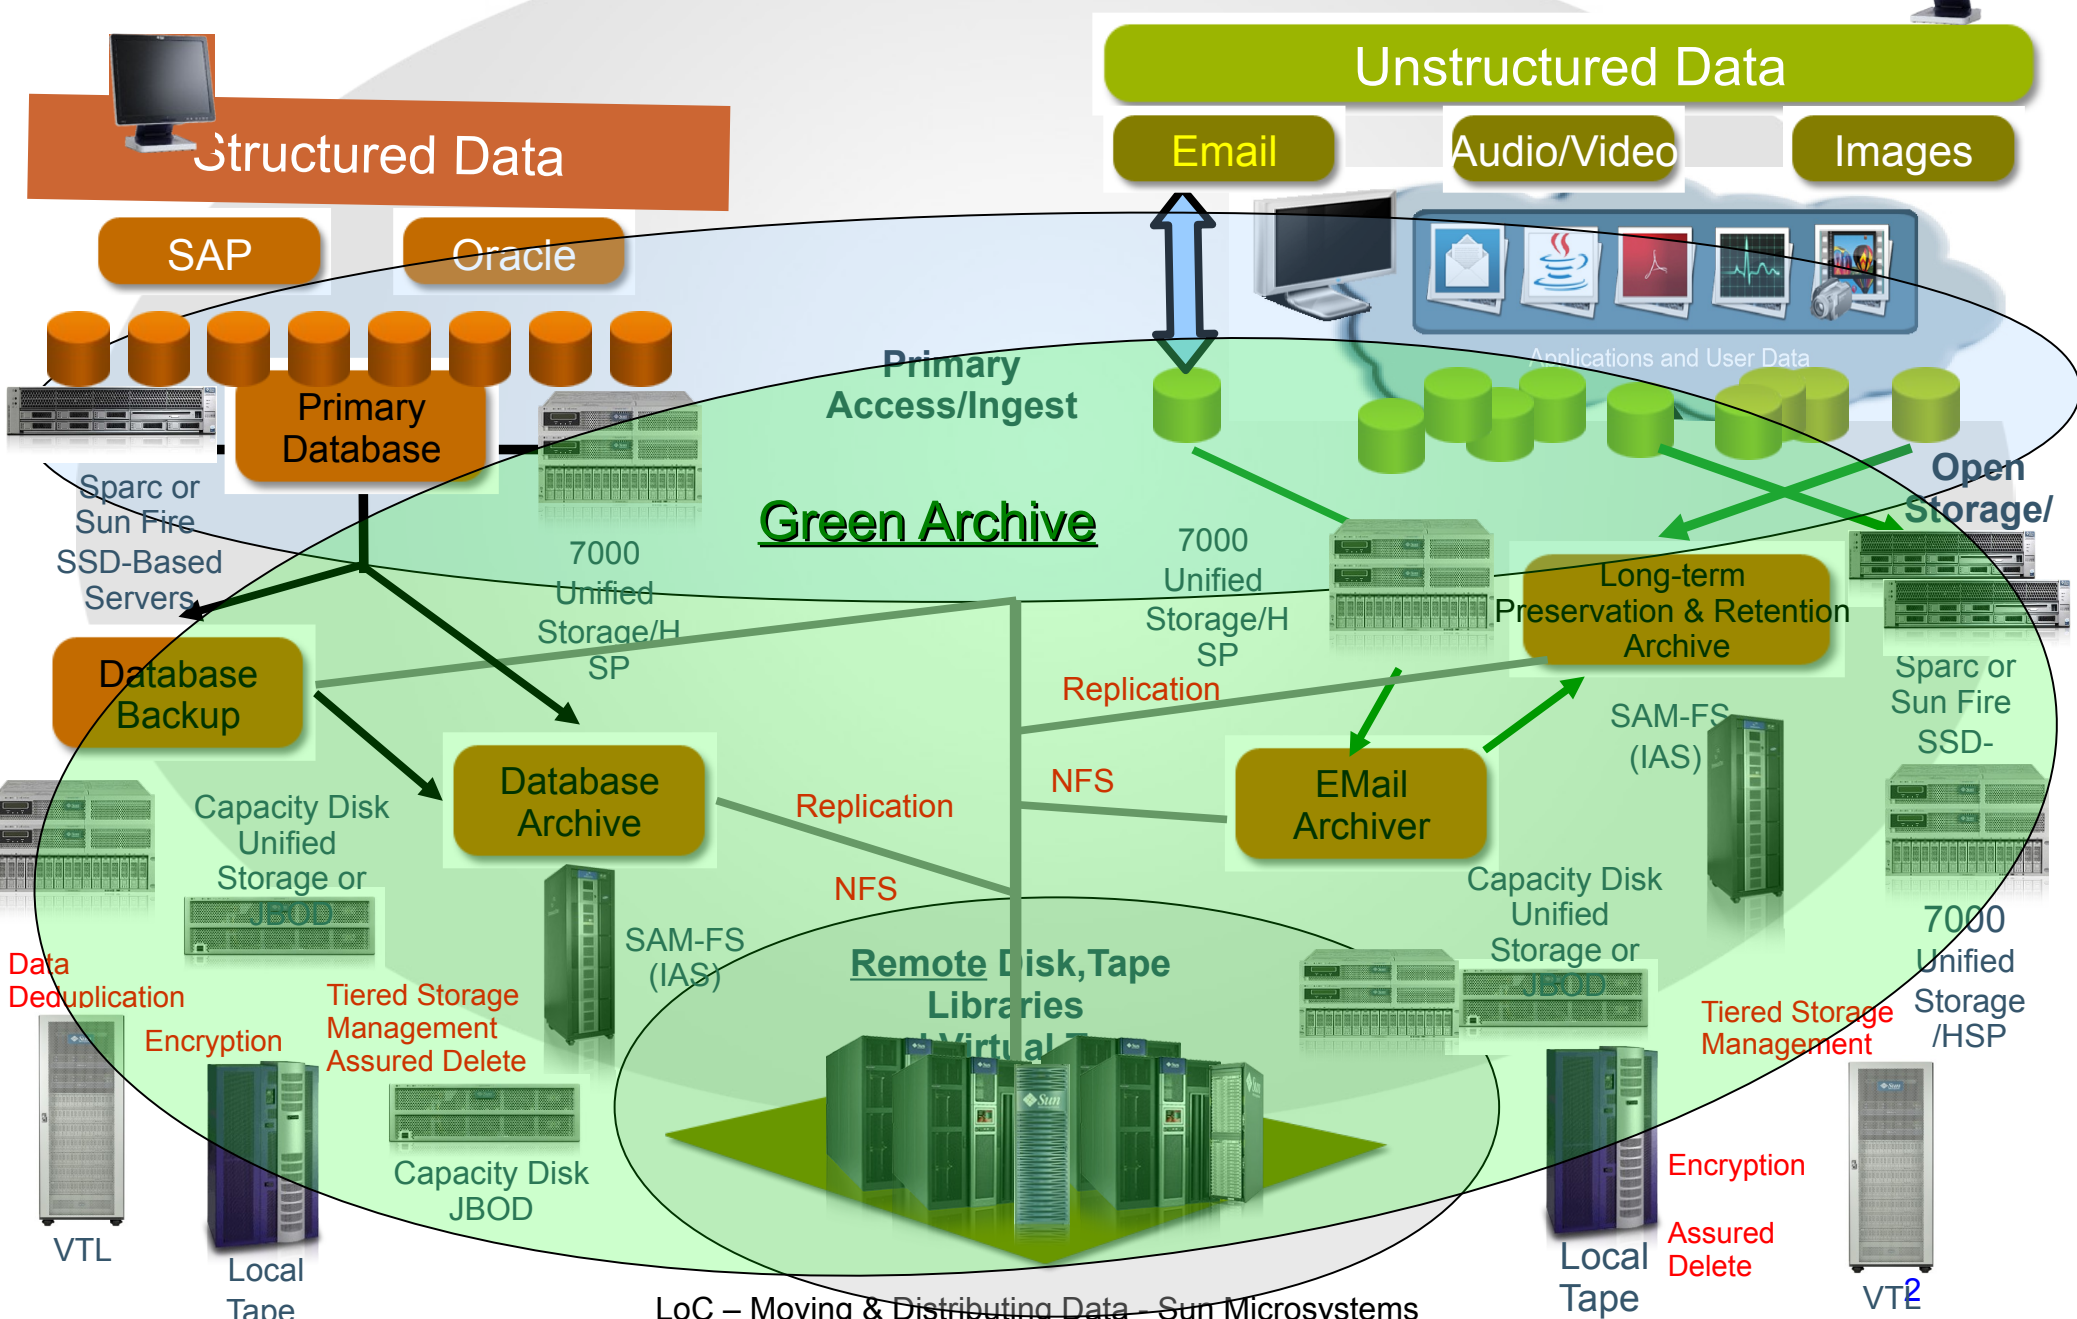


Moving Data and Distributing Data

Raymond A. Clarke

Sr. Enterprise Storage Solutions Specialist, Sun Microsystems - Archive & Backup Solutions
SNIA Data Management Forum, Board of Directors

Sun's Enterprise Archive Alternatives



CIS1 Slide from 2006 – SNIA DMF 100 Year Archive Task Force **Challenge: Manage Data for 75++ Years**

- HW typically only backward compatible N-1
- Yearly capacity increases
- Every 2-5 years HW becomes obsolete:
 - > Need to migrate current data to newer HW components
 - > Replace compute parcels
 - > Replace FC Parcel for performance and capacity
 - > Replace tape drives and media to current technology
 - > Replace SATA parcels for capacity / footprint
- Minimize vulnerability
- HW migration is inevitable; PLAN for it

What's SNIA Doing About All This?

- Educates, Defines and Taking Action to Address Industry Challenges
- Specific Activities
 - XAM – eXtensible Access Method
 - Self-Contained Information Retention Format(SIRF)
 - Rationale & Objectives
 - Requirements & Use Cases
 - Bridging Terminology
 - Green Storage Initiative
 - Cloud Storage Initiative(CSI)

What is XAM?

- **XAM is a SNIA Architecture**
 - > The XAM Architecture spec defines the normative semantics of the API for use by applications and implementation by storage systems and standardizes metadata and services across XAM compliant systems
- **XAM is an Application Programming Interface (API)**
 - > The XAM Java API spec defines the binding of the XAM Architecture to the Java Language
 - > The XAM C API spec defines the binding of the XAM Architecture to the C Language
- **XAM is SNIA Software**
 - > The XAM SDK provides a common library and reference implementation to promote widespread adoption of the standard

The need for MetaData Standards



- Which can contains corn?
- Open the cans.
- How much does it cost?
- Ask the clerk.
- How many calories does it have?
- Ask the vendor.
- How does the store automatically manage inventory?
- They can't.

MetaData Standards



Standardized labeling allows multiple vendors to consistently represent information to consumers

Nutritional Facts	
Serving Size 1/2 cup (130g)	
Servings per container about 3	
Amount per serving	
Calories 130	Fat Cal 5
	% Daily Value
Total Fat 0.5g	0%
Saturated Fat 0g	0%
Cholesterol 0mg	0%
Sodium 260mg	11%
Total Carbohydrates 22g	7%
Dietary Fiber 5g	22%
Sugars 0g	
Protein 10g	20%
Vitamin A 0%	Vitamin C 0%
Calcium 4%	Iron 10%

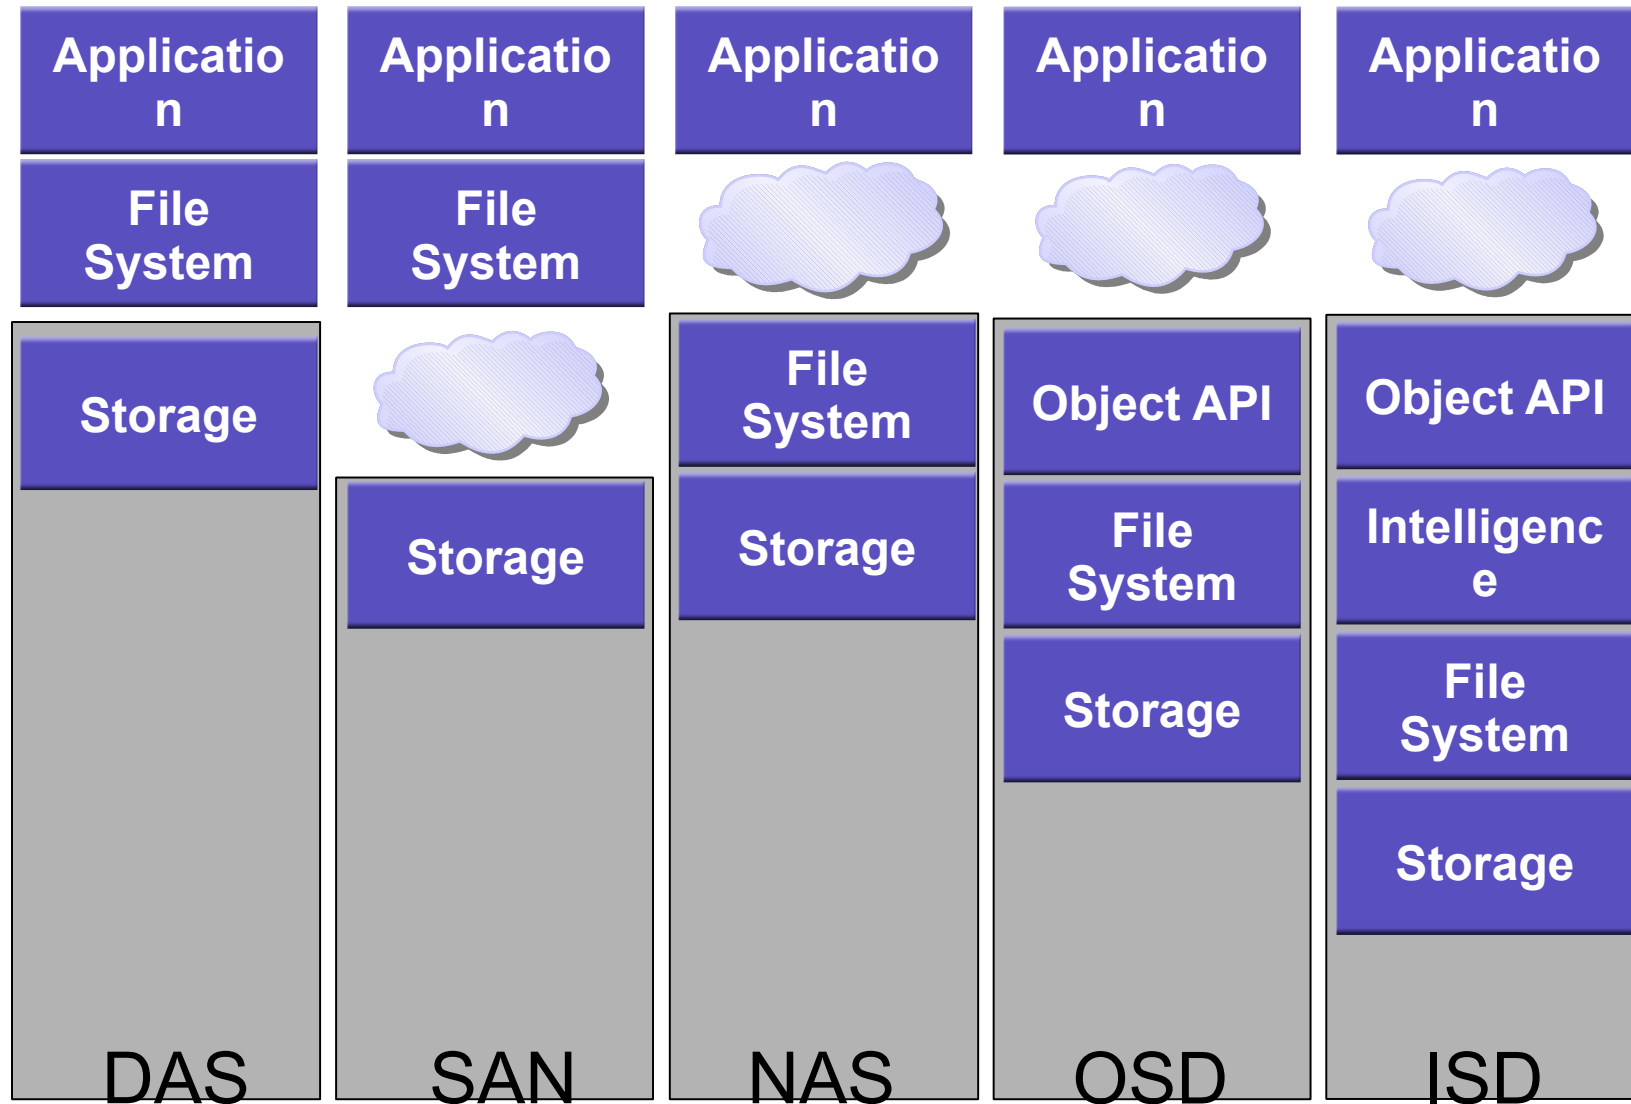
* Percent Daily Values are based on a 2,000 calorie diet

Extended labeling for LOB uses

What's Sun Doing About All This?

- Open Source/Standards Community Engagement
- Specific Activities
 - Open Solaris – ZFS(Hybrid Storage Pool)
 - Open Storage
 - Cloud Computing

The Evolution of Data Storage:



What is ZFS?

A new way to manage data

End-to End Data Integrity

With check-summing and
copy-on-write transactions

Easier Administration

A pooled storage model –
no volume manager



Immense Data Capacity

The world's
first 128-bit
file system

Data Services

Snapshots
Clones
Replication

FS/Volume Model vs. ZFS

Traditional Volumes


Abstraction: virtual disk

Partition/volume for each FS

Grow/shrink by hand

Each FS has limited bandwidth

Storage is fragmented, stranded



ZFS Pooled Storage

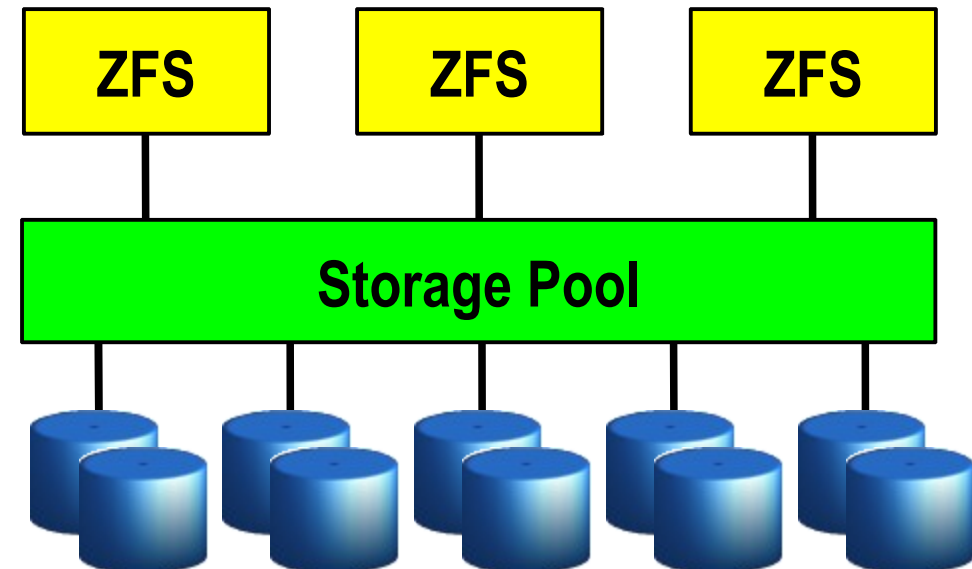
Abstraction: malloc/free

No partitions to manage

Grow/shrink automatically

All bandwidth always available

All storage in the pool is shared



What is ZFS?

A new way to manage data

End-to End
Data Integrity

With check-summing and
copy-on-write transactions

Easier
Administration

A pooled storage model –
no volume manager



Immense Data
Capacity

The world's
first 128-bit
file system

Integrated Data
Services

Snapshots
Clones
Replication


ZFS Snapshots

Provide a read-only point-in-time copy of file system



Copy-on-write makes them essentially “free”

Very space efficient – only changes are tracked

And instantaneous – just doesn't delete the copy



Open Storage/Open Archive Anatomy



- Replication
- Security
- Mirror/Snap
- Search
- Encryption
- De-duplication
- Migration
- Backup
- Compliance



Thank You
for
Your Time and
Attention

Raymond.Clarke@Sun.com

(212) 558-9321

