

A large red offshore oil rig structure is visible in the background, set against a blue sky and ocean. The rig's complex framework of pipes and beams is prominent. The text is overlaid on this image.

Big Storage for Big Problems

Petabyte File Systems on Tertiary Storage

Why Do We Care?

- **Data Explosion**
 - **Backup**
 - ▲ 1000 PCs = 2-3 TB Just for Backup
 - **Corporate Digital Assets**
 - ▲ Scanned Images
 - ▲ Video
 - ▲ Product
- **Long Term Storage Issues**
 - **Data Integrity**
 - **Vendor/Supplier Viability**

How Big is Big?



data storage solutions from



Why Not Use Disk?

Next question!

d a t a s t o r a g e s o l u t i o n s f r o m

 **EMASS**[®]
Storage Management Software

OK, So Now What?

- **Basic Requirements**

- **Low Operational Cost**

- ▲ Automated Library
- ▲ Minimal Fixed Requirements

- **Scalable**

- ▲ Predictable Access
- ▲ Manageable
- ▲ Capacity

- **Beneficial**

- ▲ Catalog Service/Asset Manager

What is Scalability?

- **Quantity**
 - Number of objects
 - Size of all objects
- **Quality**
 - **System Response**
 - ▲ Administrative
 - ▲ User
 - **Data Integrity**
 - ▲ Long Term Management

Project Goals

- **Understand Operational Characteristics**
 - Can it be managed?
 - How long do basic operations take?
- **Take the User View**
 - How does the system respond?
- **Get a Feel for How System Grows**
 - What resources run out?
- **Publish Good and Bad**

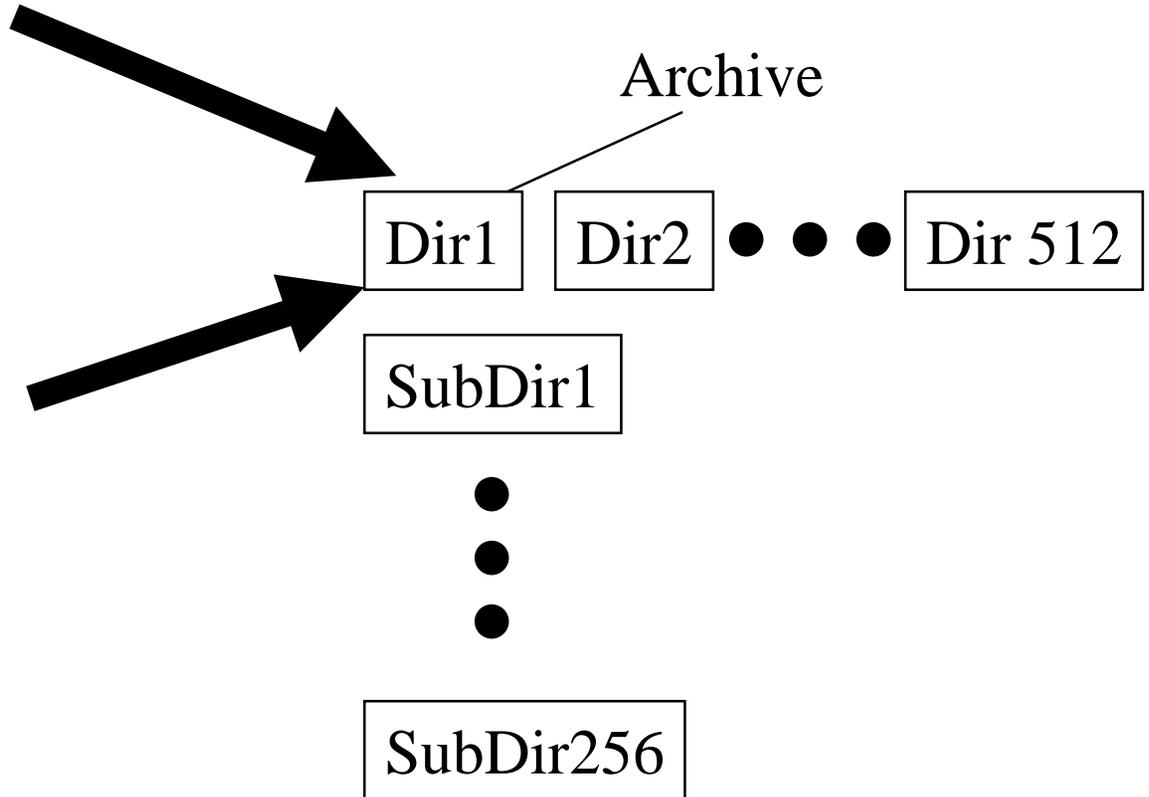
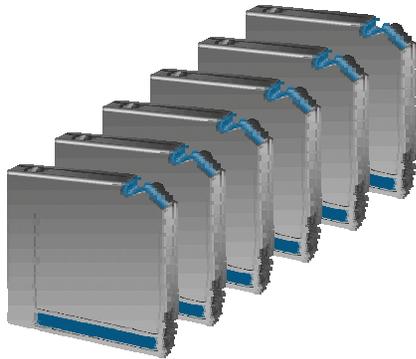
Test Case Parameters

- **1,000 TB of Data Represented**
 - **Metadata Access is the Issue**
 - **Could not Afford to write a PB**
- **Just Over 16 Million Files**
 - **XXXX levels of directories**
 - **512 Directories at the top**
- **Sun SPARC 50 MHz Classic**
- **8 GB of disk**

VGIMPORT

ASCII
version of
metadata

```
Dkkkd  
dkd s  
sks skj;kljdfkjadf  
klasdfkljkdjkj  
lkasdkkdjd  
  
asfdkl;kasd;  
a;lkasd  
a;laksdjflaskjdfkj  
kdkdkasdf;kjasdf  
akask;dd ;akdasd dlsa  
sdfkasdf;kl asdfkasf  
;aklsdf;ljkdf;kasdk  
;laksdfj;kl;sd;fks
```



Results

- **Measure**

- **User Level Commands**

- ▲ ls, cd, pwd

- **Size of Database**

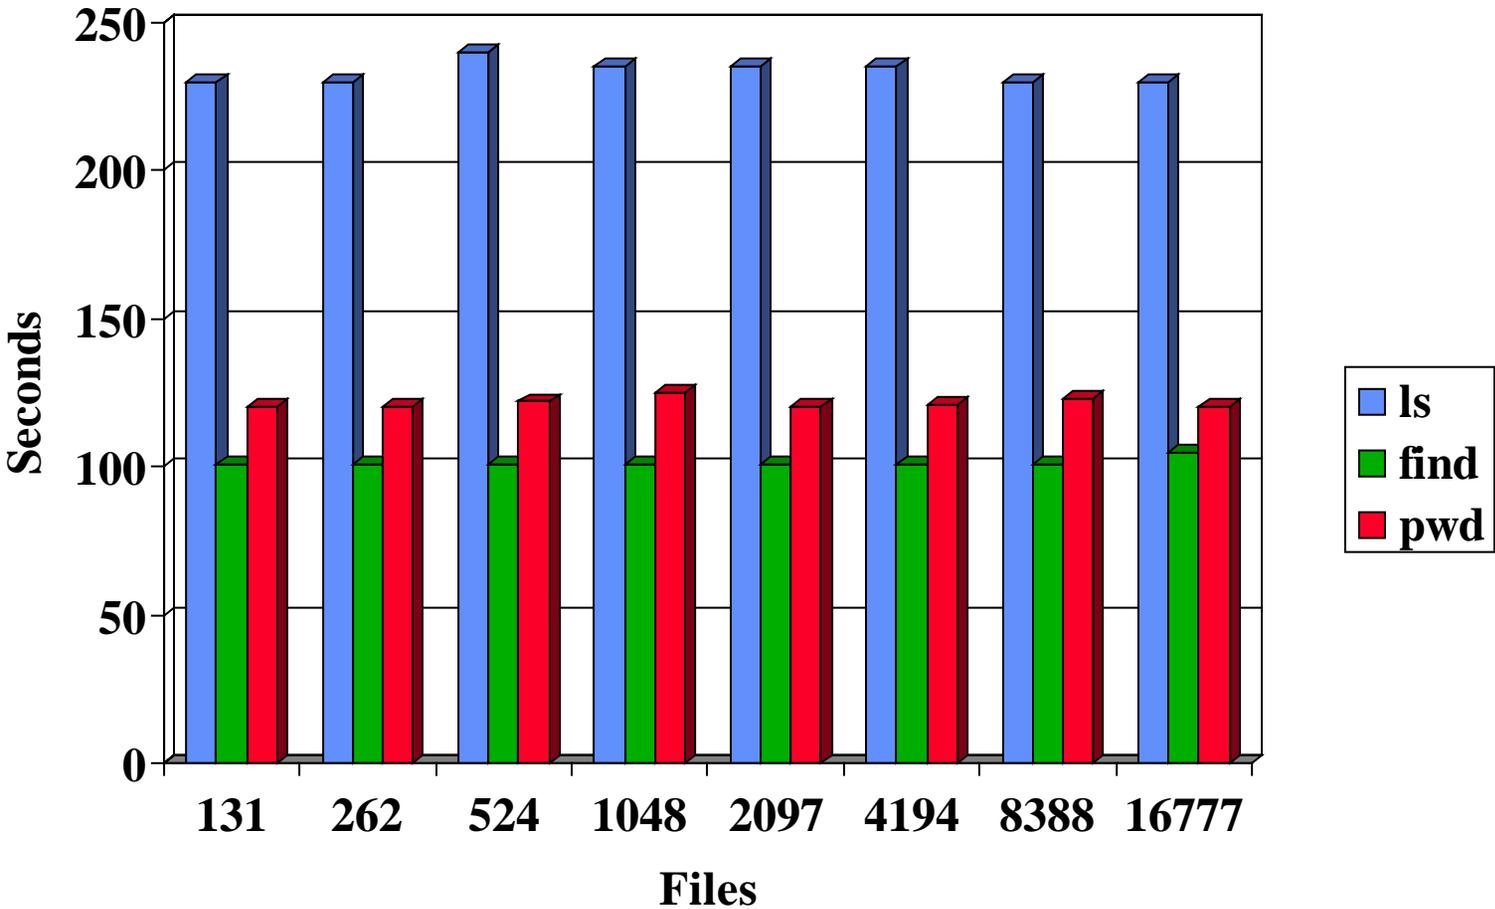
- ▲ Interested in Maximum Size

- **Administration Utilities**

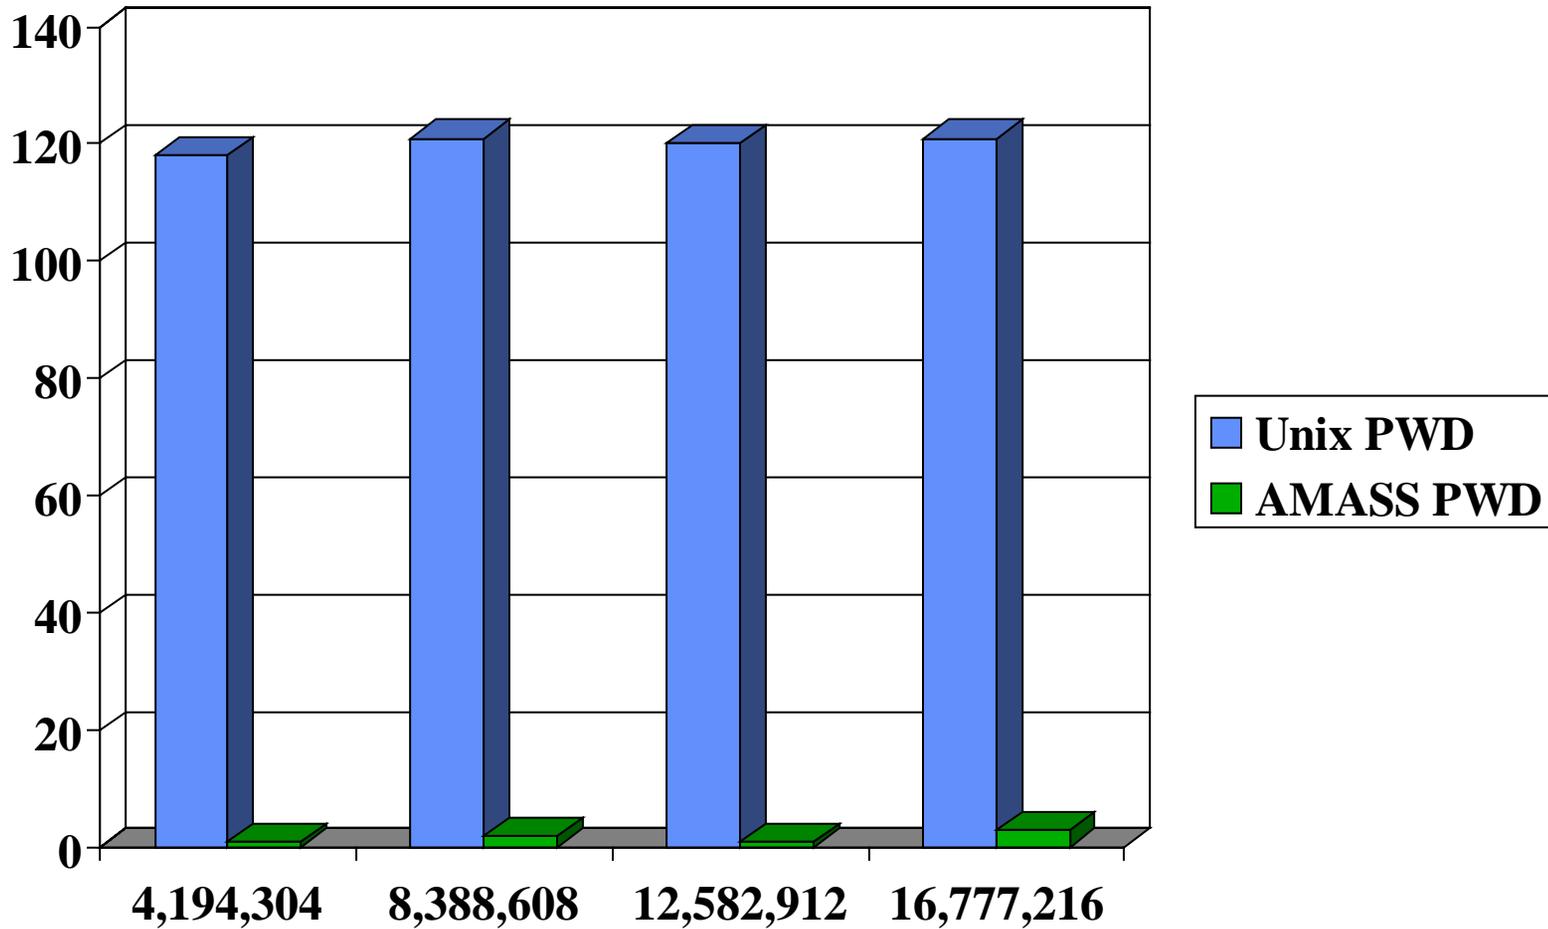
- ▲ Times for all required commands

- ▲ Include recovery processing

Metadata Access



Using the Database API



data storage solutions from

Database Size

- **AMASS**
 - Minimum of 128 bytes/file
 - 16 Million Files = 3.0 GB
 - 90 Million Files = 10 GB
- **File System + HSM**
 - 256 inode + 1024 HSM bytes/file
 - 20 GB for 16 Million Files

Phase II: Data Access

- **1 TB File**
- **100 IBM 3590**
- **Test**
 - **Single Reader**
 - **Multiple Reader**
- **Status**
 - **File was Written**
 - **Resources Moved to New Project**
 - **Will Resume this Summer**

Areas That Need Attention

- **Administration Commands**
 - Database Verification & Repair
 - Other Products Suffer Similar Fate
 - ▲ Backup, HSM, etc.
- **Some User Level Commands**
 - Implementations Need to Change
 - ▲ Memory and Algorithms, Browser Based Interfaces
- **Useability Features**
 - Catalogers
 - Long Term Media Management

Areas to Think About

- **Vendor Independence**
 - **Hardware and Software**
- **Storage and Access Paradigm**
 - **Transparent vs Purposeful**
- **Cost of Ownership**
 - **Flexibility vs Simplicity**