

# ***Web-Accessible File Storage***

***Audrey Bersot, UC Berkeley***

***abersot@cafe.berkeley.edu***

***Stephen D. Franklin, UC Irvine***

***franklin@uci.edu***

***Laura Kim, UC Berkeley***

***lkim@berkeley.edu***

# UC Berkeley: Objectives

---

- **Internet file storage and sharing**
- **Web publishing capability**
- **Windows and Macintosh support with WebDav functionality**
- **Kerberos/LDAP integration (CalNetID)**
- **Viable software company for technical and software support**

# UC Berkeley: Services

---

- **50MB file storage**
- **No cost (funded centrally – 5 years)**
- **Available to faculty, staff, and students (45K)**
- **Technical support provided (email only, no phone or walk-in consulting)**
- **On-line help**

# UC Berkeley: Selection Process

---

- **Historical Perspective**

- 1998-99 RFI/RFP formal process
- WebDisk v0, 1999-2002 idrive.com
  - RFI/RFP process
  - 3 year license
  - \$30K a year, covered all campus users
  - Outsourced
  - Company folded April 9, 2002
  - Contract extended services through June 30, 2002
- Little time for a back up plan to transition users to a new service: Fly by the seat of our pants.

# UC Berkeley: Selection Process

---

- WebDisk v1, 2002-2004 driveway.com
  - Met with the CTO of driveway.com
  - Negotiated an unlimited license
  - Supported in-house
  - 3 months to transition users from i-drive
  - Driveway folded before year-end 2002 leaving us with:
    - No technical support
    - User base of 8K+ to move someplace

# UC Berkeley: Selection Process

---

- **Lessons Learned (so far)**
  - Outsourcing a service has severe limitations
    - No control over downtimes, updates, slowness
  - Major economic downturns are unpredictable
    - Dot com bust
  - Contract renewals can be sketchy or even monetarily prohibitive
    - Try to get your money back when a company goes belly-up
  - Educational institutions beat to a different tune

# UC Berkeley: Selection Process

---

- **Vendor comparison (2002)**
  - Oracle Collaboration Suite
    - CorporateTime bought by Oracle
    - Web File Storage and Web Publishing separate entities
    - Company support difficult
  - Xythos
    - Met all of our objectives
    - Contracted July 2003

# UC Irvine's Objectives: File Storage

---

- **For UCI Faculty and Staff**
  - Permanent (backed up against system failures)
  - Substantial quantity (e.g., 1GB/person)
  - UCInetID (Kerberos) Authentication
- **“Ubiquitously” Accessible**
  - Via the Web
  - Via Folders on Windows and Mac desktops
- **Private or Easily Shared** (read-only or read-write)
  - With UCI and non-UCI collaborators
- **Security and Convenience**

# UC Irvine's Objectives: Applications

---

- **Institutional (vs personal) storage**
- **"Ubiquitous" access**
- **Alternative to large e-mail attachments**
- **Sharing/Collaboration**
  - Version control
  - Check-out/check-in
  - Locking and logging
- **(Limited) Web publishing**
- **Secure Electronic Information Resources**

# UC Irvine: Needs Not Addressed this Way

---

- **Course DropBoxes**

- Specialized Classes of Application
- Time-dependent change of access controls
- Access depends on fine-grain registration status
- Separate service provided as part of UC Irvine's  
EEE = Electronic Educational Environment

<http://eee.uci.edu/toolbox/dropbox/>

- **Individual/Personal (Student) File Store**

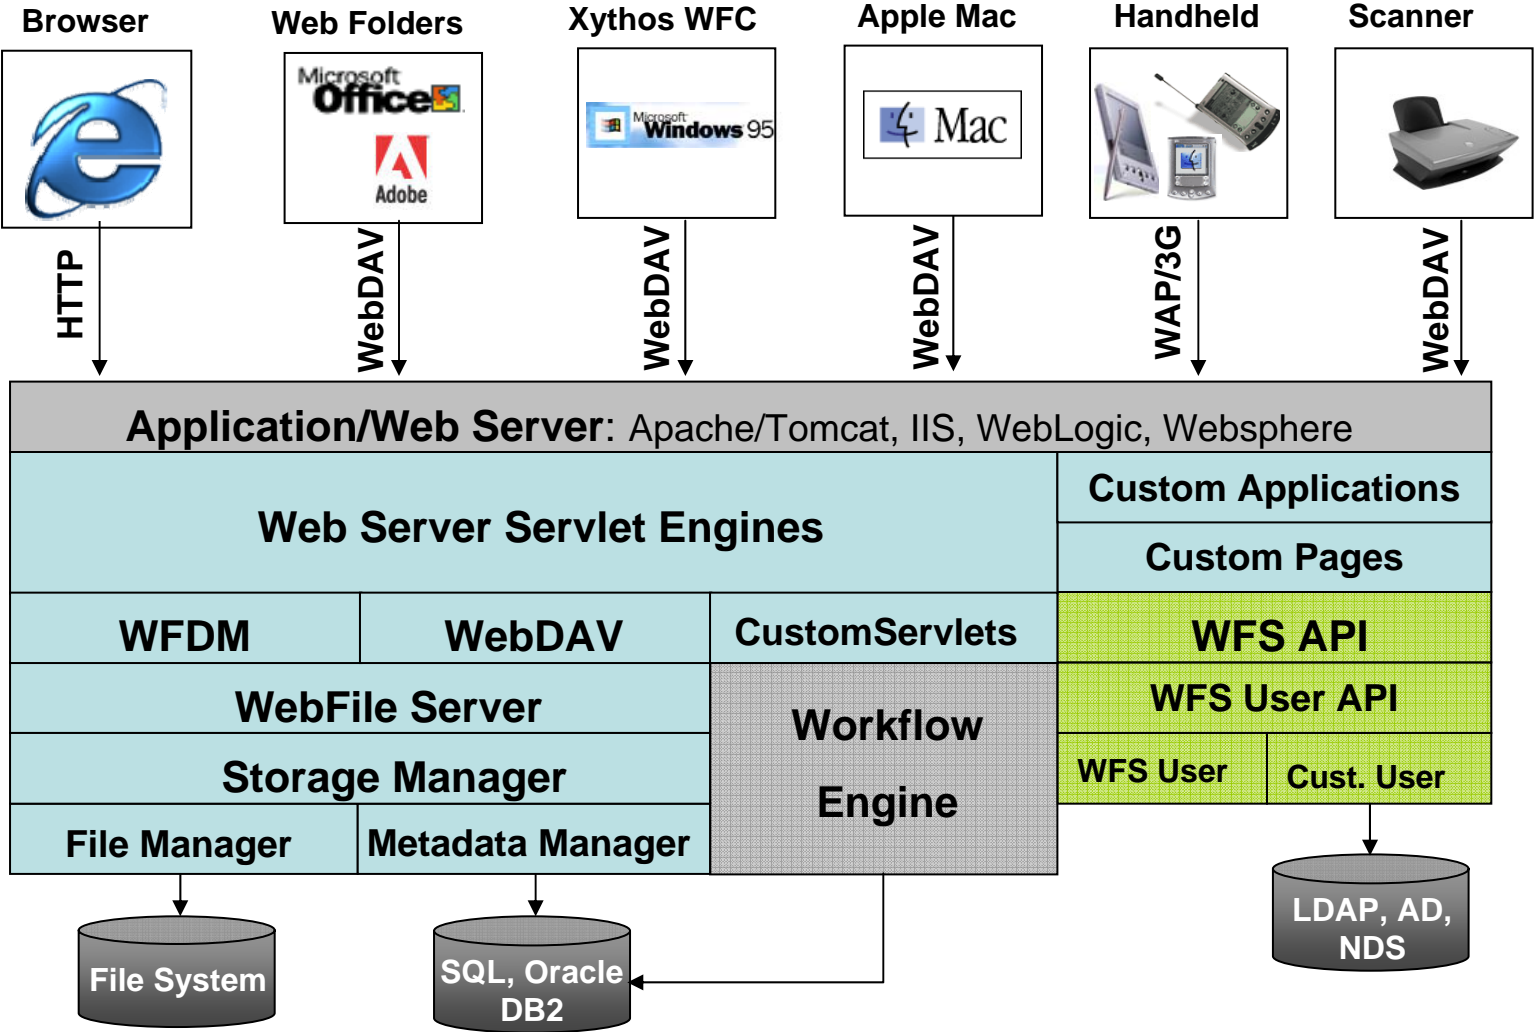
- Mobile personal memory devices  
USB connected, based on flash memory or  
mini-disks (e.g., iPod)

# WebDAV

---

- **Web Distributed Authoring and Versioning Access Control Protocol**
  - [IETF RFC 3744](#)
- **Extensions to HTTP protocol “allowing users to collaboratively edit and manage files on remote web servers”**
  - <http://webdav.org/>
- **Strong UC presence in its development**
  - Jim Whitehead while a grad student at UC Irvine and now faculty at UC Santa Cruz

# Xythos WebFile Server Architecture



# UC Irvine WebFile System Configuration

---

- **Front End Server (Western Scientific @ ~\$5K)**
  - User file access through this server via
    - user's web browser accessing web server on this system
    - other webDAV client programs on user's system
    - webDAV “speaking” parts of user's OS (e.g., “Web Folders” in Microsoft Windows)
  - Access control and logging
  - Uses database for metadata about user files
- **RAID Server (Western Scientific @ ~\$10K)**

# UC Irvine Webfiles Front End

---

- **Front End Server**
  - Dual 1.4 GHz Opteron 2GB RAM
  - 80 GB local disk, software only
  - 1 GigE connection (can go to 5) to network
- **Red Hat 9 Linux (moving to RHEL 3)**
- **Apache Tomcat Web Server and Java servlet container**
- **PostgreSQL database (moving to Oracle)**
- **Xythos WFS version 4.0.48 (current is 4.2)**

# UC Irvine WebFile File Store

---

- **RAID Server**
  - 160 MB/sec SCSI connection to front end server
  - 512MB RAM
  - 4 TB raw (16 disks @ 250 GB + 2 spares)
  - 1.7 TB useable fully mirrored:  
RAID 10 because reads predominate

# UC Berkeley: Selection Process

---

- **Xythos Contract**

- 3 year at \$15K for 6500 users license
  - \$2.31/user
  - Option to buy in increments of 1000 licenses
- Kerberos integration developed w/Xythos
- WFS 4.2
- Migration plan in process (user driven)
- Currently in testing phase
- Rollout to campus in mid-August

# UC Berkeley: Technical Specifications

---

- **Front End Server (Dell @ \$12K)**
  - Redhat Linux ES
  - 2 X 2.2 GHz CPU
  - 2GB RAM
  - 273GB Raid Array
- **Database (Dell @ \$16K)**
  - Windows Server 2003
  - 2 X 3.06 GHz CPU
  - 2GB RAM
  - 111GB Raid Array

# UC Irvine WebFiles Deployment

---

- <https://webfiles.uci.edu/>
- **1 GB for all faculty and staff**
- **Must register and then activate account**
- **Over 1,600 accounts (July 2004)**
- **Built-in on-line help**  
<http://webfiles.uci.edu/xythoswfs/static/en/HelpIndex.html>
- **UCI on-line help, tutorials, documentation**  
<http://www.nacs.uci.edu/webfiles/docs/documentation.html>
- **Presentations & Training Sessions**

# UC Berkeley: Future Ideas

---

- <http://webdisk.berkeley.edu>
- University consortium to collaborate, share code, and ideas
- Integrate Intellitach (email attachments) with email service
- Integrate with CourseWeb service
- Integrate with computer labs for transparent desktop storage
- Expand storage to 100MB per user

# webfile-hied@uci.edu

---

- **Discussion and sharing within Higher Education about Webfile (Xythos WebDAV) servers**
- **Primary “audience” (subscriptions lightly screened):**
  - Those supporting services based on Webfile servers.
  - Those considering offering such services
  - Some Xythos employees also.
- **Not directly for questions of individual end-users**
- **Xythos’s customer contact still primarily elsewhere**
- **Self-moderating: postings only from subscribers**
- **Subscription and (subscriber-only) archives:**  
**<https://maillists.uci.edu/mailman/admin/webfile-hied>**

# UC Berkeley: Credits

---

- **Audrey Bersot**
- **Laura Kim**
- **John Shiohita**
- **Rusty Wright**

# UC Irvine Credits

---

- **Webfiles**
  - Sylvia Bass
  - Joseph Farran
  - Kim Gerrard
  - John Mangrich
  - Allen Schiano
- **EEE DropBoxes**
  - Eric Carter
  - David Pritikin
  - Owen Shaw
  - Briandy Walden