

WebDAV and Apache

Greg Stein
gstein@lyra.org
<http://www.lyra.org/>

Agenda

- Overview
- Benefits
- Scenarios
- Setting up mod_dav
- DAV tools and clients
- Futures

What is WebDAV?

(1 of 2)

- **Web-based Distributed Authoring and Versioning**
 - “DAV” is the usual short form
- Goal: enable interoperability of tools for distributed web authoring
- Turns the Web into a *writable* medium

What is WebDAV?

(2 of 2)

- Applies to all kinds of content - not just HTML and images
- Based on extensions to HTTP
- Uses XML for properties, control, status
- RFC 2518

Benefits

- Benefits for all web users:
 - Users
 - Authors
 - Server administrators
- Technical benefits for developers, network administrators, and security personnel

User Benefits

- User: defined here as a web surfer
- Document metadata available
- More intelligent “directory” listings

Author Benefits

- Author: the person who writes the content
- Standard way to place content on server
- Move/copy the content around
- Tag the content with metadata
- Overwrite protection in group scenarios

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

7

Administrator Benefits

- Administrator: the person running the server
- DAV namespace does not have to mirror a specific directory structure (flexibility)
 - Could even be used for arbitrary repositories
- Authentication via HTTP; system accounts not required (security)

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

8

Technical Benefits

Overview

- Properties (“metadata”)
- Overwrite protection
- Namespace management
- Infrastructure: old and new
- Replacement protocol

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

9

Technical Benefits

Terminology

- Collection
 - A collection of resources
 - A collection is also a resource
- Resource
 - Generic name for collections or *member resources*
- Member Resource
 - “Leaves” in a URL namespace

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

10

Technical Benefits

Properties

- Properties are name/value pairs
 - Names are uniquely identified with URIs
 - Values are well-formed XML fragments
- All resources have properties
 - Files and directories
 - Server-defined/maintained, or client-defined
- Records metadata such as author, title, modification time, or size

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

11

Technical Benefits

Overwrite Protection

- Shared and exclusive locks
- Locks have characteristics such as timeouts, owners, and depth
- Identified by authentication and lock token
- Apply to whole resources, not portions

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

12

Technical Benefits

Namespace Management

- “Namespace” refers to the URL hierarchy
- DAV provides mechanisms to create, move, copy, and delete resources

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

13

Technical Benefits

Existing Infrastructure

- Receives benefits of HTTP infrastructure
 - Strong authentication
 - Encryption
 - Proxy/firewall navigation
 - Worldwide deployment
 - Huge talent pool; numerous tools, apps, etc

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

14

Technical Benefits

New Infrastructure

- DAV can provide infrastructure for:
 - Collaboration
 - Metadata
 - Namespace management
 - *Ordered collections*
 - *Versioning*
 - *Access control*
 - *Searching*

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

15

Technical Benefits

Replacement Protocol

- DAV providers read/write to the web server
- Can obsolete other mechanisms:
 - FTP
 - FrontPage and Fusion proprietary protocols
 - Custom or one-off solutions
- Robust enough for future enhancements

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

16

WebDAV Under the Covers

- A protocol layered on HTTP/1.1
- HTTP extensions
 - New HTTP headers
 - New HTTP methods
 - Additional semantics for existing methods
 - HTTP/1.1 clarifies the extension process

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

17

New HTTP Headers

- DAV:
- If:
- Depth:
- Overwrite:
- Destination:
- Lock-Token:
- Timeout:
- Status-URI:

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

18

New HTTP Methods

Overview

- COPY, MOVE
- MKCOL
- PROPPATCH, PROPFIND
- LOCK, UNLOCK
- Semantics for HTTP/1.1 methods:
 - GET, PUT, DELETE, OPTIONS

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

19

New HTTP Methods

COPY, MOVE

- Pretty obvious: copy or move resources
- Copying collections uses Depth: header
- Destination: header specifies target
- Also uses Overwrite: header
- Optional request body controls the handling of live properties

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

20

New HTTP Methods

MKCOL

- Create a new collection
- Avoids overloading PUT method

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

21

New HTTP Methods

PROPPATCH, PROPFIND

- PROPPATCH is used to set, change, or delete properties on a single resource
- PROPFIND fetches one or more properties for one or more resources

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

22

More on PROPFIND

- Using PROPFIND anonymously allows users to discover files
- Best to require authentication
- In the future:
 - Browsers will want it for “nice” directories
 - Clients will want PROPFIND for metadata
 - Server will have finer granularity to hide items

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

23

New HTTP Methods

LOCK, UNLOCK

- Add and remove locks on resources
- Both use the Lock-Token: header

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

24

Scenarios

- Collaborative authoring
- Network file system
- Unified repository-access protocol
- Remote software engineering
 - Minimal support until versioning arrives

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

25

Scenario: Departmental Server

(1 of 2)

- Department of 20 staff
- They operate a private web server
- Web server acts as a repository
 - File servers used to play this role
- Everybody needs to author documents
- Web server (vs file server) provides better navigation, overviews, and offsite links

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

26

Scenario: Departmental Server

(2 of 2)

- Web site is DAV-enabled
 - Allows remote authoring and maintenance
 - Allows tagging documents with metadata
- Security can be used to limit or partition areas for specific users
- Documents drop right onto the server
- New pages for summaries and overviews

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

27

Scenario: Web Hosting

(1 of 2)

- 5000 users
- `http://www.someisp.com/username/`
- No need to enter users into `/etc/passwd`
 - Use any Apache `mod_auth_*` module
- User directories can be distributed, shifted, updated as needed across the filesystem

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

28

Scenario: Web Hosting

(2 of 2)

- Apache's `httpd.conf` gets complicated
 - Need section for each user
 - Something like `UserDir` would be great
 - For now, include a generated file

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

29

Other Scenarios

- Not restricted to the Internet
- LAN environments
 - Departmental workgroups
 - Software development teams
- WAN/VPN environments
 - Remote workgroups, development
- Base protocol for client/server interactions

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

30

Setting up mod_dav

Overview

- Grab and install tarball
- One simple directive:
DAV On
 - Use within <Directory> or <Location>
- Need to change file/dir ownership and privs
- Enable locking
- Add security as appropriate

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

31

Setting up mod_dav

Installation

- Grab tarball:
 - http://www.webdav.org/mod_dav/
- Install in one of two ways:
 - Via APXS (easiest)
 - Build within the Apache source tree
- Expat (a subset) is part of Apache 1.3.9

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

32

Setting up mod_dav

Example Configuration

```
Alias /gstein /home/apache/davdirs/gstein
<Location /gstein>
  DAV On
</Location>
```

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

33

Setting up mod_dav

Filesystem Changes

- Assume Apache is run with UID “nobody” and GID “www”

```
% ls -la /home/apache/davdirs/gstein
total 3
drwxr-s--- 3 nobody www 1024 Jun 25 14:32 .
drwxr-s--- 3 nobody www 1024 Jun 28 17:26 ..
-rw-r--r-- 1 nobody www 424 Jun 26 16:36 index.html
drwxr-s--- 4 nobody www 1024 Jun 26 13:05 specs
```

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

34

Setting up mod_dav

Enable Locking

- Additional directive for the lock database
DAVLockDB /home/apache/davdirs/lock.db
- Lock databases are per-server

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

35

Setting up mod_dav

Security Considerations

- Prevent funny operations (CGI, includes, etc)
Options None
- Prevent .htaccess
AllowOverride None
- Limit method access
**<Limit PUT DELETE PROPFIND PROPPATCH \\
MKCOL COPY MOVE LOCK UNLOCK>**

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

36

Limiting PROPFIND

- Note that PROPFIND is in the <Limit> directive
 - Limits the use of PROPFIND to authorized users
 - Based on concerns mentioned earlier about “discoverability” of a web site

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

37

Example Configuration

```
<Location />
  AllowOverride None
  Options None
  DAV On
  AuthName "my web site"
  AuthType basic
  Auth_MySQL on
  Auth_MySQL http_auth
  <Limit PUT DELETE PROPFIND PROPPATCH MKCOL COPY \
    MOVE LOCK UNLOCK>
    Require user gstein
  </Limit>
</Location>
```

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

38

DAV Tools and Clients

Overview

- Open source
 - Joe Orton's sitecopy and cadaver
 - Nautilus
 - Python, Perl, C client APIs
- Commercial
 - Adobe: GoLive 5.0
 - Microsoft: Office 2000, IE5
- Lots of DAV servers

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

39

DAV Tools and Clients

Joe Orton's sitecopy

- Edit web site locally
 - Update remote web site
 - Operates via FTP or WebDAV
 - More/better functionality via WebDAV
 - Does not do two-way synchronization
-
- Screem, an HTML editor, embeds sitecopy

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

40

DAV Tools and Clients

Joe Orton's cadaver

- Interactive command-line tool
- Provides listing, moving, copying, and deleting of resources on the server

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

41

DAV Tools and Clients

Nautilus

- Nautilus is the file manager for GNOME
- Uses gnome-vfs
 - “Virtual File System”
 - Can target WebDAV repositories
- Provides full GUI-based management of a DAV repository

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

42

DAV Tools and Clients

Language APIs

- Good for experimentation
- Building new apps
- Most are layered onto existing HTTP APIs

- Python API from Greg Stein
- Perl API from Patrick Collins
- C API (Neon) from Joe Orton

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

43

DAV Tools and Clients

Adobe GoLive 5.0

- One of the first Web authoring tools to support the DAV protocol
- Page design, authoring, construction
- Uses locking to assist authoring teams
- Site management

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

44

DAV Tools and Clients

Microsoft Office 2000

- Broad distribution
- Word, Excel, etc are DAV-enabled
 - Open/save files directly from/to web server
 - Uses DAV locks for overwrite protection
- First round of Microsoft's move to DAV
 - Also: IIS5, Exchange 2000

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

45

DAV Tools and Clients

Internet Explorer 5.0

- Enabled with the "Web Folders" add-on
- Adds "Web Folders" section into Windows Explorer, under "My Computer"
 - Allows drag and drop of files
 - Standard move/copy/delete/rename of files

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

46

Future Clients (Wish List)

- KDE file manager
- Pharmacy (a GNOME frontend to CVS)
- Emacs EFS
- Mozilla
- CVS-like front-end
 - Subversion is going in this direction

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

47

Implementing mod_dav

- Apache has great extensibility
- But:
 - Hard to add new methods
 - Security: file ownership, SUID helpers, etc
 - Alternate access to repository
- Security issues led to private repository
- Module provides excellent speed

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

48

Futures: WebDAV

- Advanced Collections (idle; Q1 2001?)
 - Bindings, Ordering, References
- Searching (somewhat active; Q4 2000?)
- Access Control (active; Q4 2000?)
- Versioning (active; Q4 2000)

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

49

Futures: mod_dav

- mod_dav 1.0 was released on June 13, 2000
- Apache 2.0 includes core DAV features
 - fully integrated
 - better plug-in system
 - possibly: versioning
- Apache 2.1
 - Other DAV extensions

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

50

Review

- WebDAV can change the very nature of how people interact with the Web
- Great standard, replaces many protocols with a single protocol
- mod_dav brings DAV to Apache
- Tools and apps are starting to appear

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

51

Resources

- <http://www.webdav.org/>

Everything you need is on this web site, or linked from it.

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

52

Q&A

October 25, 2000

ApacheCon 2000 Europe: WebDAV and Apache

53