GridPP
UK Computing for Particle Physics

GridSite Web Servers for bulk file transfers & storage

Andrew McNab
Grid Security Research Fellow
University of Manchester, UK

MANCHESTER 1824

Outline

- Recent “bulk file” oriented additions to the GridSite (www.gridsite.org) system
  - GridSite overview
  - Security model
  - Read/write access via HTTP(S)
  - Onetime passcodes
  - Third party transfers
  - SiteCast file location

GridSite Overview

- GridSite has evolved from the GridPP website management system
- Now provides a Grid-oriented security toolkit (libgridsite) and extensions to the Apache webserver
- Supports Grid/Web services on Apache using CGI
  - C/C++, Perl, other scripting languages
- See GridSite Web Services poster for more details

Design philosophy

- Most Grid deployments (eg LCG + EGEE) are based on protocols and security technologies derived from the Web
- So we attempt to reuse high quality implementations like Apache from the mainstream
- This significantly reduces our support burden, since core Apache, mod_ssl, OpenSSL, ... is “RedHat's Problem” (or whoever does your distribution...)

Security model

- Authentication is done in Apache's mod_ssl using the client's X.509 certificate or GSI proxy
  - mod_gridsite dynamically modifies the OpenSSL callbacks to handle GSI proxies correctly
- VOMS attributes are extracted if present, and the server has access to a cache of any DN-Lists which have been fetched asynchronously.
- XML policy engine based on GACL or XACML languages decides whether access is permitted

Read / write access

- Almost all web traffic uses the GET method to fetch files, or POST to send the results of a form
  - But the HTTP/WebDAV RFCs also define PUT, DELETE and MOVE methods
- mod_gridsite adds support for these “write” methods, subject to the policy-based access model
  - So HTTP(S) servers act as read/write file stores
- Our htcp etc commands (cf scp) provide clients, but curl and many standard clients can be used too

Onetime passcodes

- For bulk files, may want an unencrypted data stream
  - cf GridFTP's use of an encrypted control channel and unencrypted data channel
- GridSite achieves this using an HTTPS GET/PUT to establish access rights
  - The server then issues an HTTP redirect to an HTTP URL
  - A onetime passcode is returned as a cookie
- This “GridHTTP” protocol works with unmodified versions of curl etc, and our htcp command

Third-party transfers

- WebDAV RFC defines a COPY method, which can be used for a client C to orchestrate a transfer of a file from remote server A to server B
- GridSite now implements this, both in the server (gridsite-copy.cgi) and client (htcp)
- We use onetime passcodes as a simple form of delegation from C to B, to give it the right to access the file
  - Supports both single stream and multistream HTTP

Third-party transfers

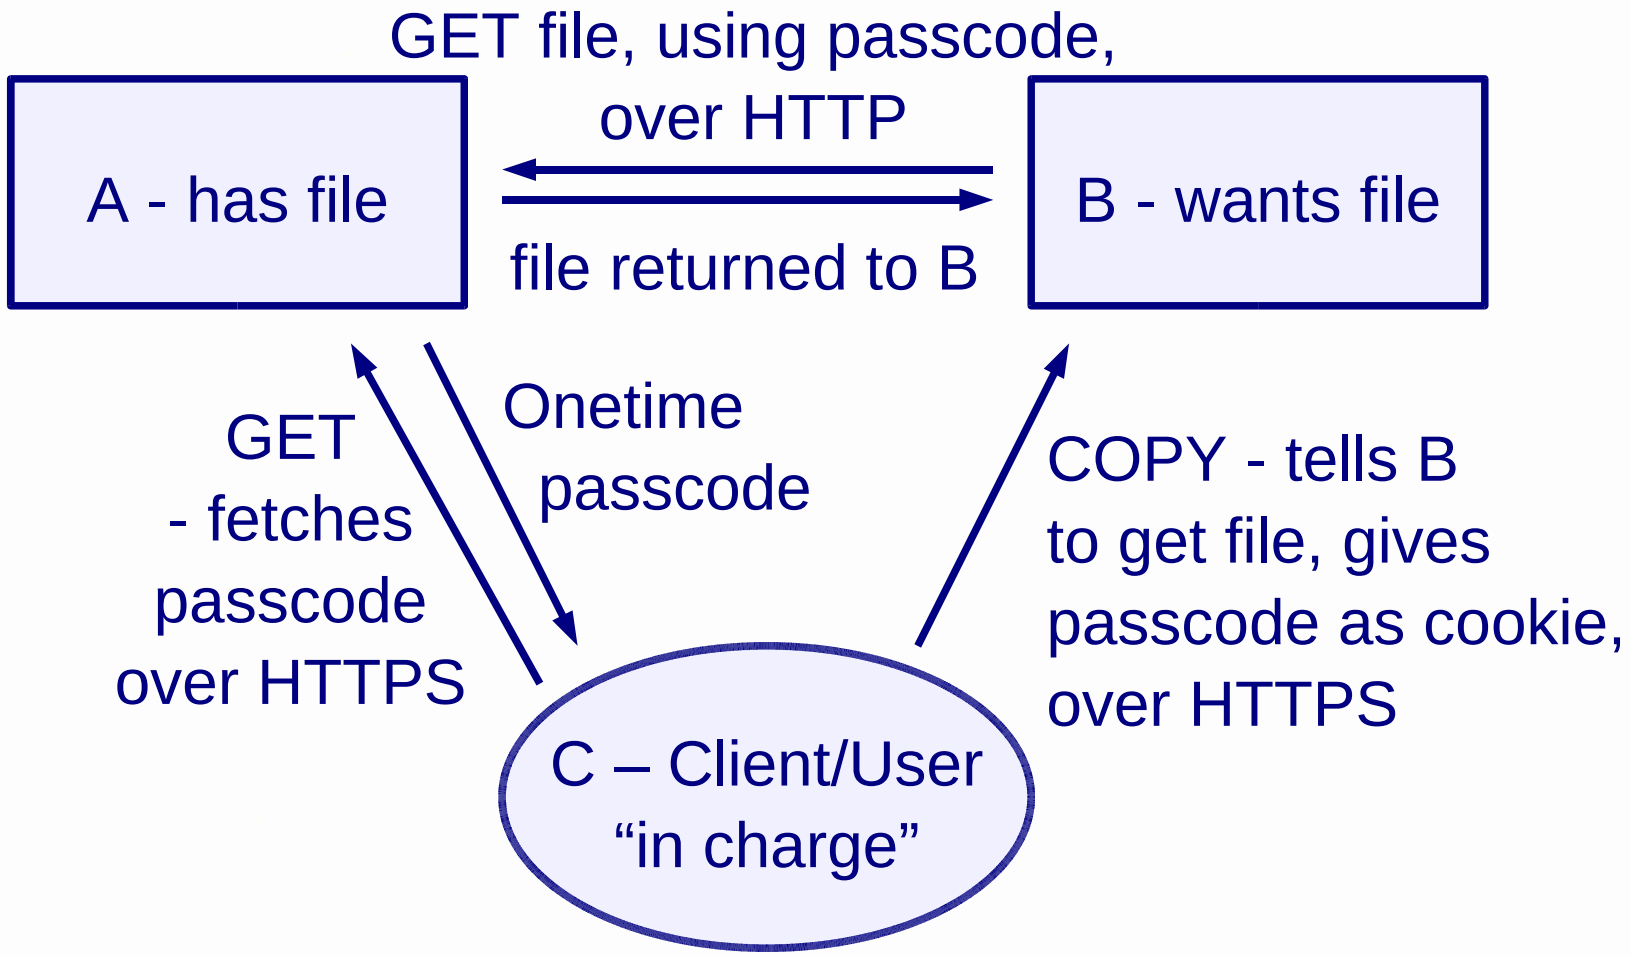

SiteCast

- Current work is looking at how to locate local replicas of files on GridSite HTTP(S) servers
- Have designed a simple replica location system for farms with many disks/hosts
  - Implemented in server-side (mod_gridsite) and htcp
  - Uses multicast of Hypertext Cache Protocol queries to find lists of replicas of a given file: looks at filesystem rather than any database
  - no database to keep in sync; automatically avoids replicas on dead machines; multicast can be filtered / routed by network hardware

Summary

- GridSite (www.gridsite.org) is already used for
  - Website/server management
  - Secured Web Services for grids, in C/C++/Scripts
- Now also has features for bulk file transfer
  - Fine grained, VOMS-aware access control
  - Secure Read/write using HTTP or HTTPS
  - Third party transfers using COPY
- Current work is on file location within a site
  - Using HTCP multicast to locate files
  - Very lightweight: no database needed