



Scaling ownCloud horizontally

by user redirection and proxying

Why don't we just use ...

Data grid technology

- ❖ LFC - high-performance file catalog
- ❖ Globus **gridftp** - parallel file transfer
- ❖ **dCache** - gridftp, SRM, xrootd, **HTTP/WebDAV**
- ❖ **iRODS/SRB** (Storage Resource Broker) - comprehensive data-management suite
- ❖ Etc.

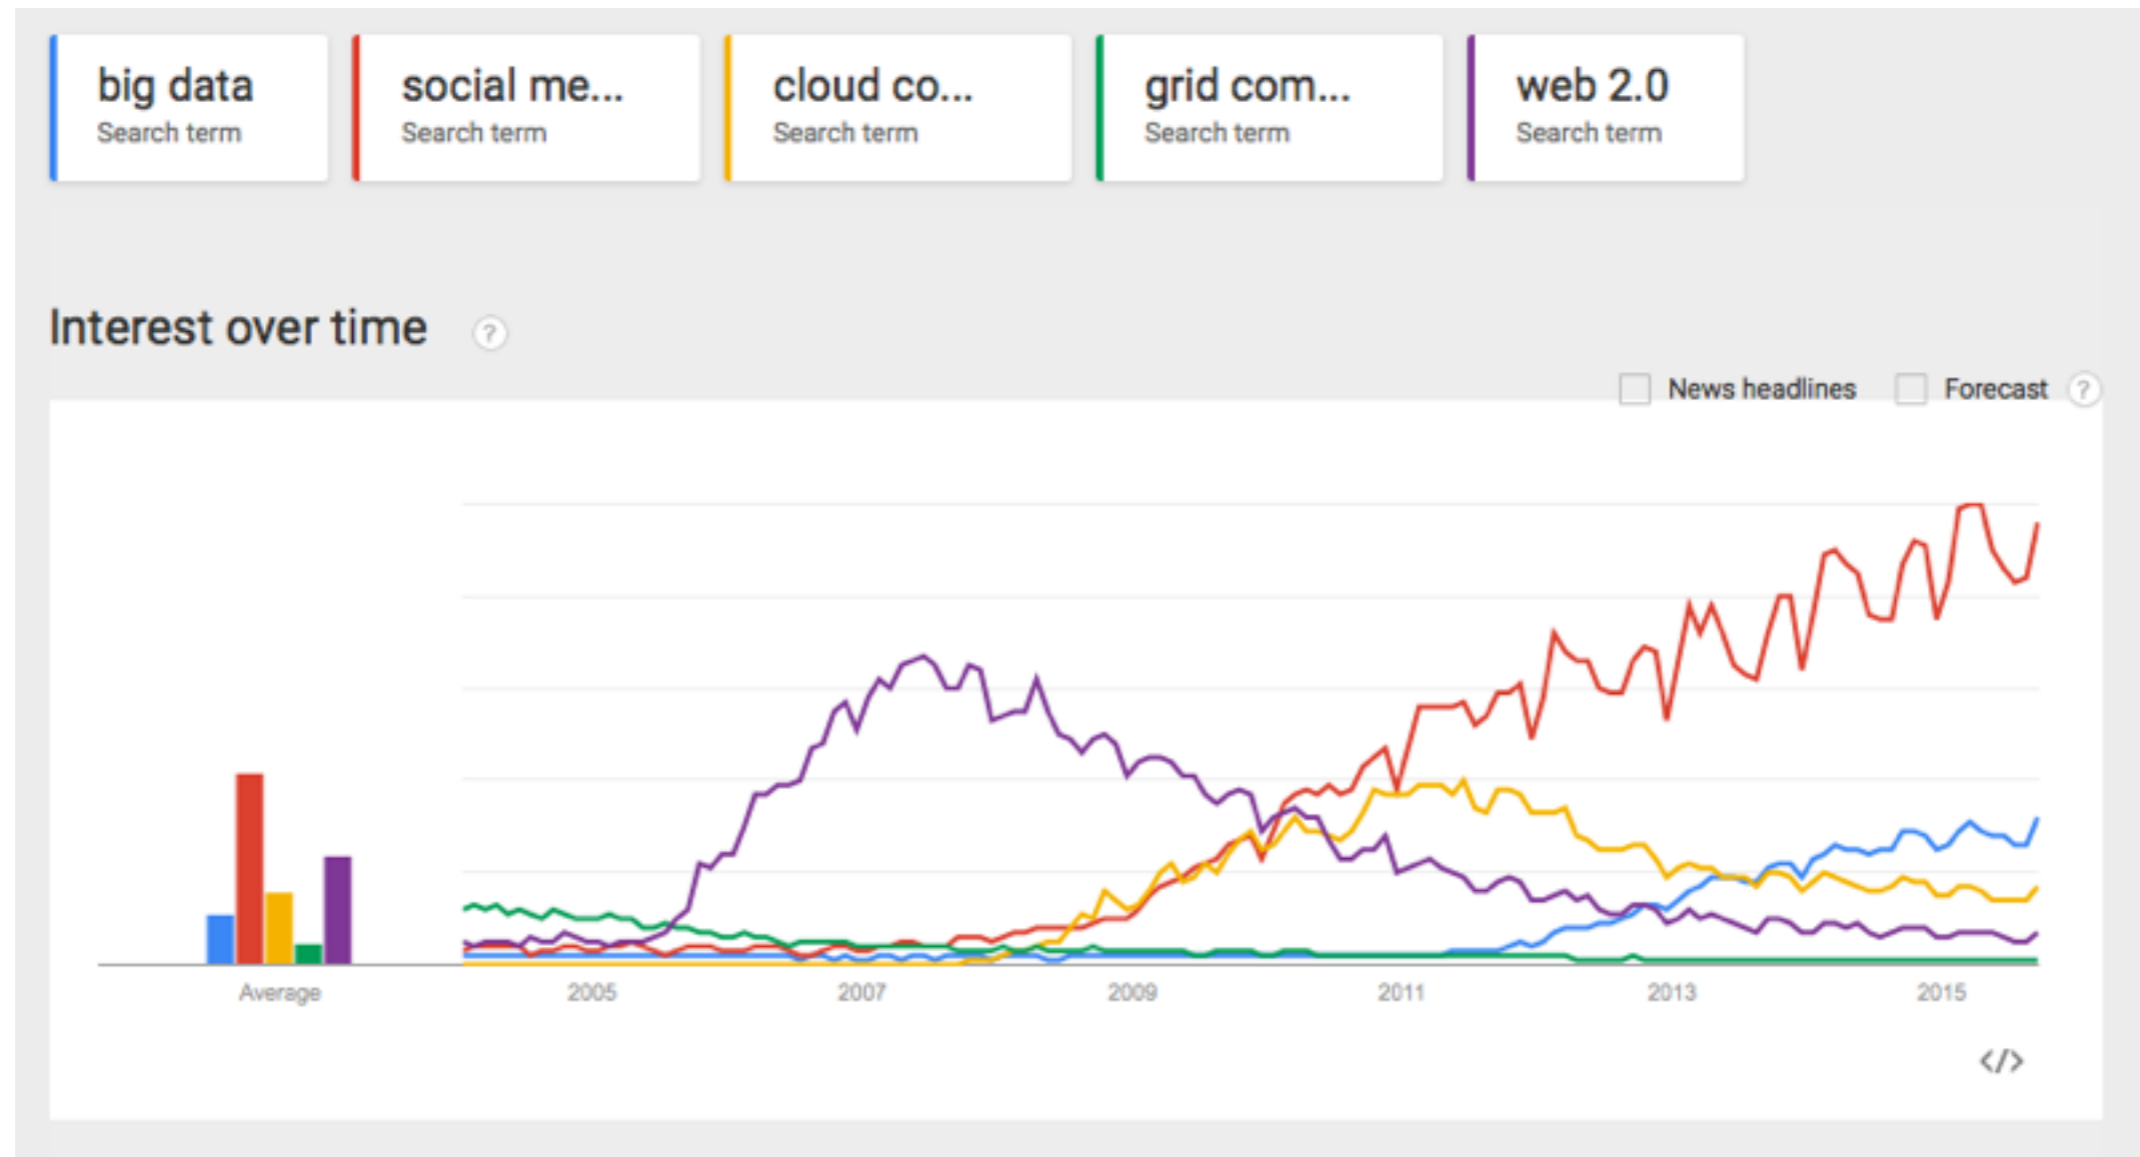
Public clouds

- ❖ Amazon S3/Dropbox
- ❖ Google Cloud Storage/Drive
- ❖ Microsoft Azure/OneDrive

Enterprise - private clouds

- ❖ Buy it all from EMC, CSC, IBM?

Technology changes



We're a national *research* data service provider

As such, we need:

- ❖ Long-term perspective
- ❖ Vendor independence
- ❖ Open standards
- ❖ *Customisability*

Places ▾

- Home
- documents
- Shared by me
- Shared with me
- Trashbin

Tags ▾

- mytag
- music

Apps ▾

- Metadata
- Teams
- Pictures
- Importer
- Uploader



[Upload](#) [New ▾](#)

<input type="checkbox"/>	Name ▲	Size	Modified
<input type="checkbox"/>	📁 Data	0 kB	3 months ago
<input type="checkbox"/>	📁 documents	6.3 MB	last month
<input type="checkbox"/>	📁 music	3.6 MB	9 months ago
<input type="checkbox"/>	📁 photos	663 kB	9 months ago
<input type="checkbox"/>	📁 testfolder	0 kB	6 days ago
<input type="checkbox"/>	📄 IPMIViewSuperBlade.pdf mytag	1.3 MB	last month
<input type="checkbox"/>	📄 msymdoc.pdf	1.8 MB	last month
<input type="checkbox"/>	📄 ownCloudUserManual.pdf	1.8 MB	9 months ago
5 folders and 3 files		15.4 MB	

Places ▾

- Home
- documents
- Shared by me
- Shared with me
- Trashbin

Tags

- mytag
- music

Apps

- Metadata**
- Teams
- Pictures
- Importer
- Uploader

New tag

Name	Color	On left menu	Public	Files
mytag		<input checked="" type="checkbox"/>	<input type="checkbox"/>	1
sound		<input type="checkbox"/>	<input type="checkbox"/>	0
				0
				2

Tag name:

Meta-data fields:

 ✕

✕

Add meta-data field

OK Cancel



Places

- Home
- documents
- Shared by me
- Shared with me
- Trashbin

Tags

- mytag
- music

Apps

- Metadata
- Teams
- Pictures
- Importer
- Uploader

Notify about my own actions


You need to set up your email address before you can receive notification emails.

Send emails:

WebDAV sccess control

Allow unauthenticated WebDAV access from compute nodes

Storage Use



2015

Daily History

Status	Month	Bill (DKK)	Balance (MB)	Invoice	Payment
Pending	August	0	0	2015-August-febd528a.pdf	Pay Now

Bottlenecks of a traditional ownCloud deployment

- ❖ Web server (Apache)
- ❖ Database (MySQL)
- ❖ File system
- ❖ Network card
- ❖ Internet uplink

Traditional scaling of LAMP sites

- ❖ Separation of web servers and DB server
- ❖ HTTP load balancing
- ❖ Reverse proxying
- ❖ MySQL clustering (sharding)

What's the problem?

- ❖ *Complexity*
- ❖ Non-static content on disk → performant shared file system → more *complexity*

Let's try something else

- ❖ Sharding/load balancing on HTTP level - redirect by user ID
- ❖ Each participating 'slave' server self-contained, i.e. can exist on its own (default single-machine LAMP)

A prototype - status

- ❖ SAML login on master - redirection to slave
- ❖ PHP session passed on to slave
- ❖ Shared files proxied in web interface
- ❖ WebDAV not redirected
- ❖ Shared files not (yet) available via WebDAV
- ❖ All done in app + theme

Places

- Home
- documents
- Shared by me
- Shared with me
- Trashbin

Tags

- mytag
- music

Apps

- Metadata
- Teams
- Pictures
- Importer
- Uploader

Personal settings

You have used **15.4 MB** of the available **64 TB**

Show First Run Wizard again

- Profile
- Password
- App settings
- Data settings**

Data location

Home site:

Backup site:

Server for sync clients: <https://silos1.data.deic.dk>

Backup server: <https://silos5.data.deic.dk>

Save

Is it simpler?

- ❖ From a deployment and operation point of view:
yes, a lot
- ❖ From a development point of view: no 😊
 - Lines of code: app: 7'000 + theme: 20'000 (a lot of duplication)