

Why a Storage Federation?

- Have a large number of storage endpoints
 - heterogeneous
 - loosely-coupled
 - geographically separated
- Want to unify them seamlessly
 - one namespace
 - one endpoint
- To solve the data locality problem for distributed cloud computing
 - and more

EMI Dynamic Federation

- High-performance

- aggressive metadata caching in RAM
- maximal concurrency
- scalable
 - $\sim 10^6$ hit/s per core
 - 24 GB of RAM for metadata cache is enough for ~ 100 PB of data in end-points

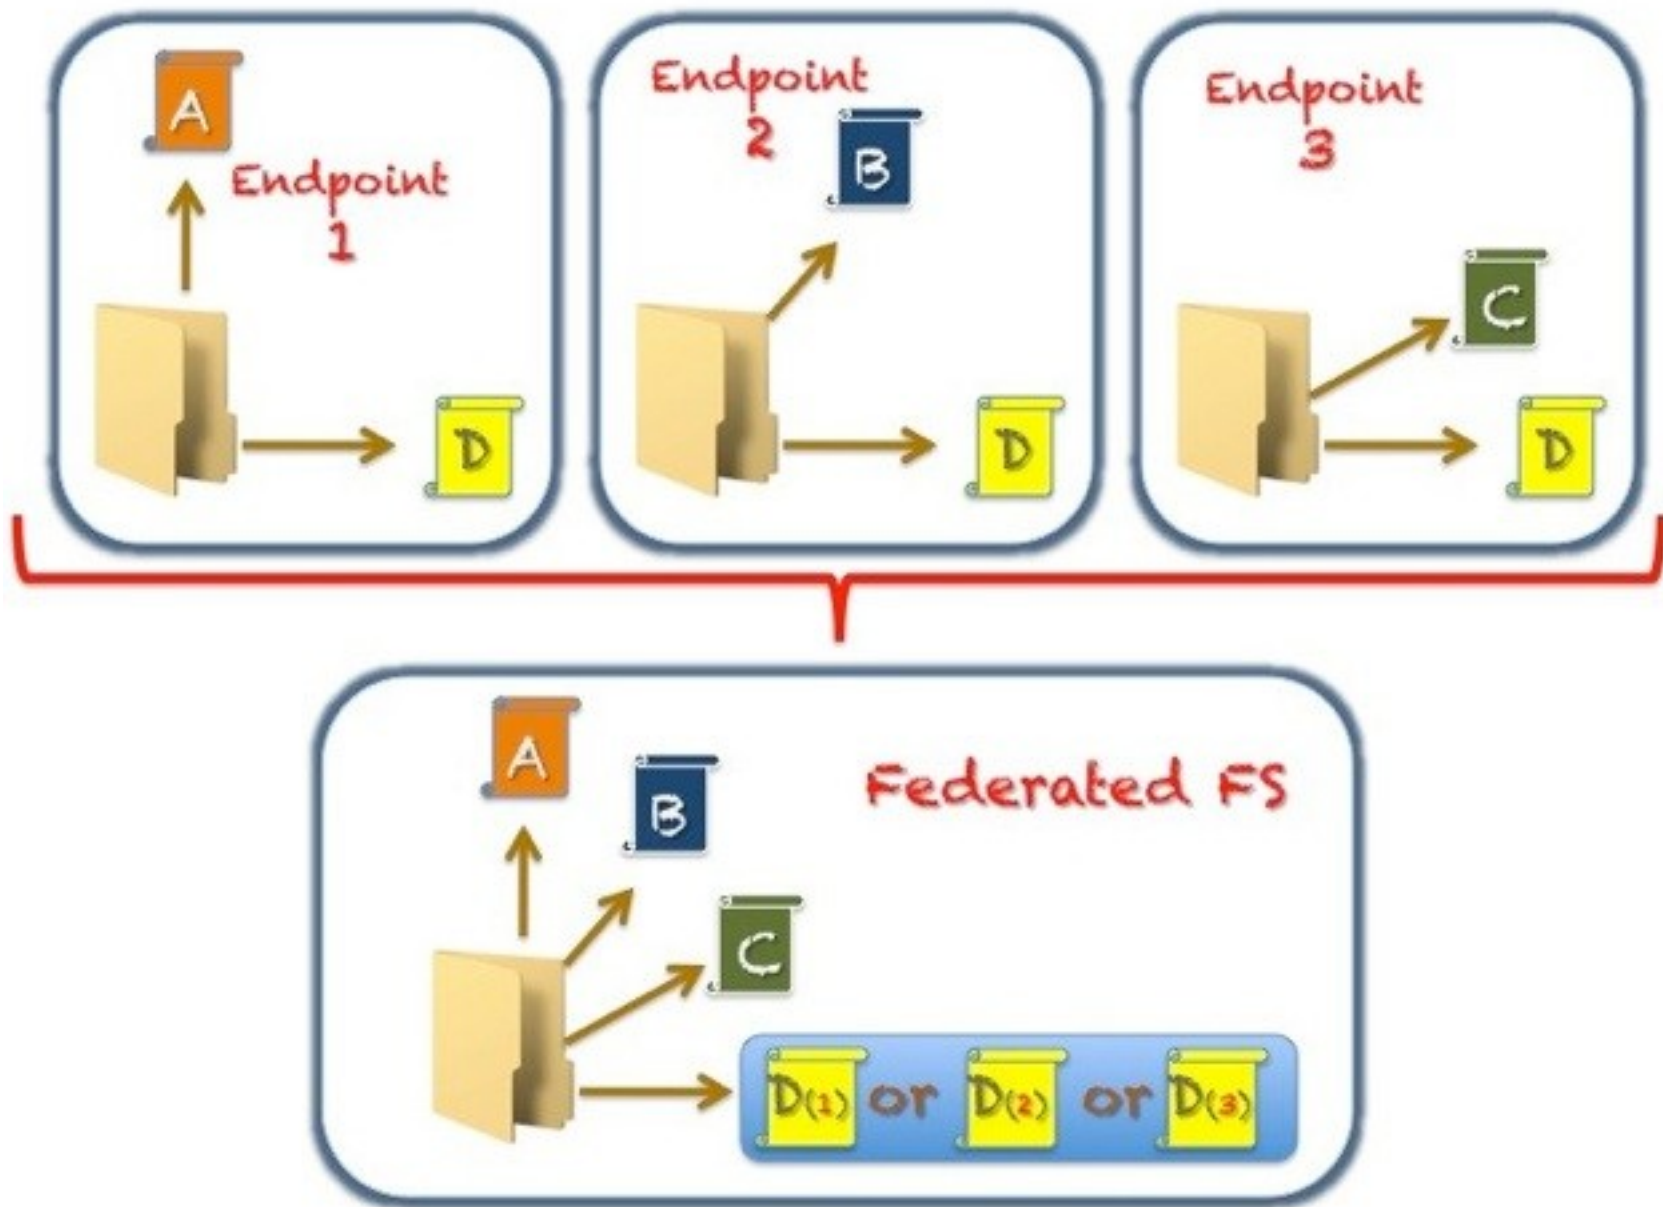


- Well-designed

- stateless, no persistency
- common standards-based components and protocols
- simple, modular
- trivial to integrate additional endpoints into the federation

- Dynamic

- automatically detect and avoid offline endpoints
- federating is done on-the-fly



The dynamic federation overlays the contents of the endpoints. If a file exists in only one of the endpoints, that location is selected. If a file exists in more than one endpoint, the GeoIP module determines which of the candidates is picked.

Software Components

















- Uniform Generic Redirector (UGR)
 - Core component containing all federation logic
 - Integrated as a plugin of Apache
- Apache HTTP server
 - Frontend to clients
 - Handles client redirection
- Memcached
 - for 2nd-level shared metadata caching in RAM
- DMLite
 - Name translations for unifying grid storage endpoints

Test Deployment

- Development SE at UVic
- HTTP server at NRC

⏪ ⏩ S ugr.heprc.uvic.ca/myfed/

/myfed/

Mode	UID	GID	Size	Modified	Name
drwxrwxrwx	0	0	0	Wed, 01 Aug 2012 16:29:41 GMT	 AOD/
drwxrwxrwx	0	0	0	Mon, 19 Aug 2013 20:01:17 GMT	 BaBarMovie/
drwxrwxrwx	0	0	0	Mon, 19 Aug 2013 20:01:28 GMT	 BaBarMovie_save/
drwxrwxrwx	0	0	0	Mon, 19 Aug 2013 20:01:39 GMT	 BaBarMovie_save2/
drwxrwxrwx	0	0	0	Thu, 23 May 2013 20:31:04 GMT	 HITS/
-rwxrwxrwx	0	0	2.2K	Fri, 15 Nov 2013 21:53:26 GMT	 MakeDataList.sh
drwxrwxrwx	0	0	0	Tue, 23 Apr 2013 17:21:31 GMT	 RDO/
drwxrwxrwx	0	0	0	Fri, 13 Jan 2012 09:48:30 GMT	 atlasdatadisk/
drwxrwxrwx	0	0	0	Fri, 13 Jan 2012 09:48:30 GMT	 atlasgroupdisk/
drwxrwxrwx	0	0	0	Fri, 13 Jan 2012 09:48:31 GMT	 atlashotdisk/
drwxrwxrwx	0	0	0	Fri, 02 Aug 2013 21:18:18 GMT	 atlaslocalgroupdisk/
drwxrwxrwx	0	0	0	Fri, 13 Jan 2012 09:48:33 GMT	 atlasproddisk/
drwxrwxrwx	0	0	0	Sun, 01 Dec 2013 02:16:12 GMT	 atlasscratchdisk/
drwxrwxrwx	0	0	0	Thu, 16 Jan 2014 10:32:32 GMT	 belle2/
-rwxrwxrwx	0	0	932.8M	Fri, 03 Jan 2014 21:47:56 GMT	 kk2f-1000.root
drwxrwxrwx	0	0	0	Fri, 24 May 2013 16:06:52 GMT	 new_RDOs/