



# Toward a Federated Global-Scale File System

Daniel Ellard, Craig Everhart - NetApp ADG  
Renu Tewari, Manoj Naik - IBM Almaden  
FAST'07 WiP

## What is a “Federated storage system”?

- ▶ **A federation is not simply a really big cluster.**
  - Federation members contribute/share resources.
  - Federation members retain local control:
    - How resources are allocated
    - Who is authorized to do what
    - Local replication / migration
  
- ▶ **Goal: enhance the user experience.**
  - Support for a global shared namespace.
  - Uniform abstraction.



**Currently, clusters of different types typically do not interact meaningfully**

Clusters of  
Server Brand A

Clusters of  
Server Brand B

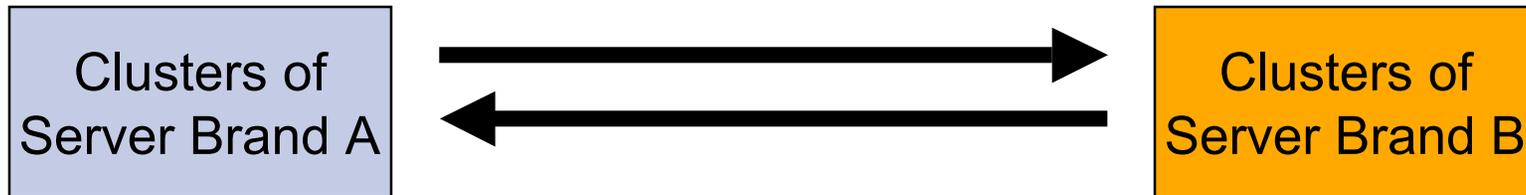
Islands of Data

## Let Brand A affiliate with Brand B



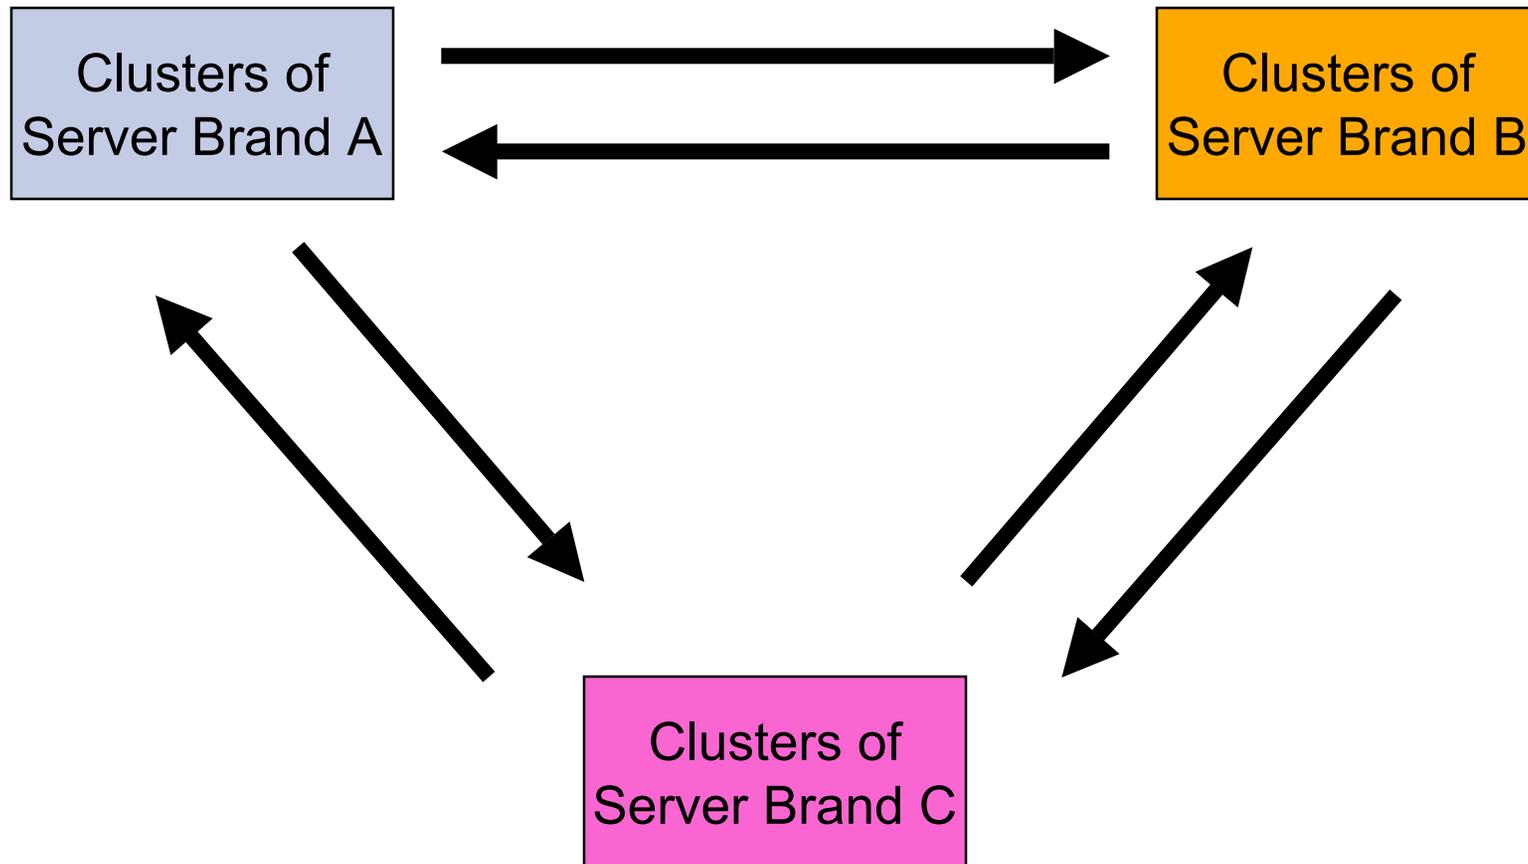
Cluster B can be a leaf in Cluster A's namespace

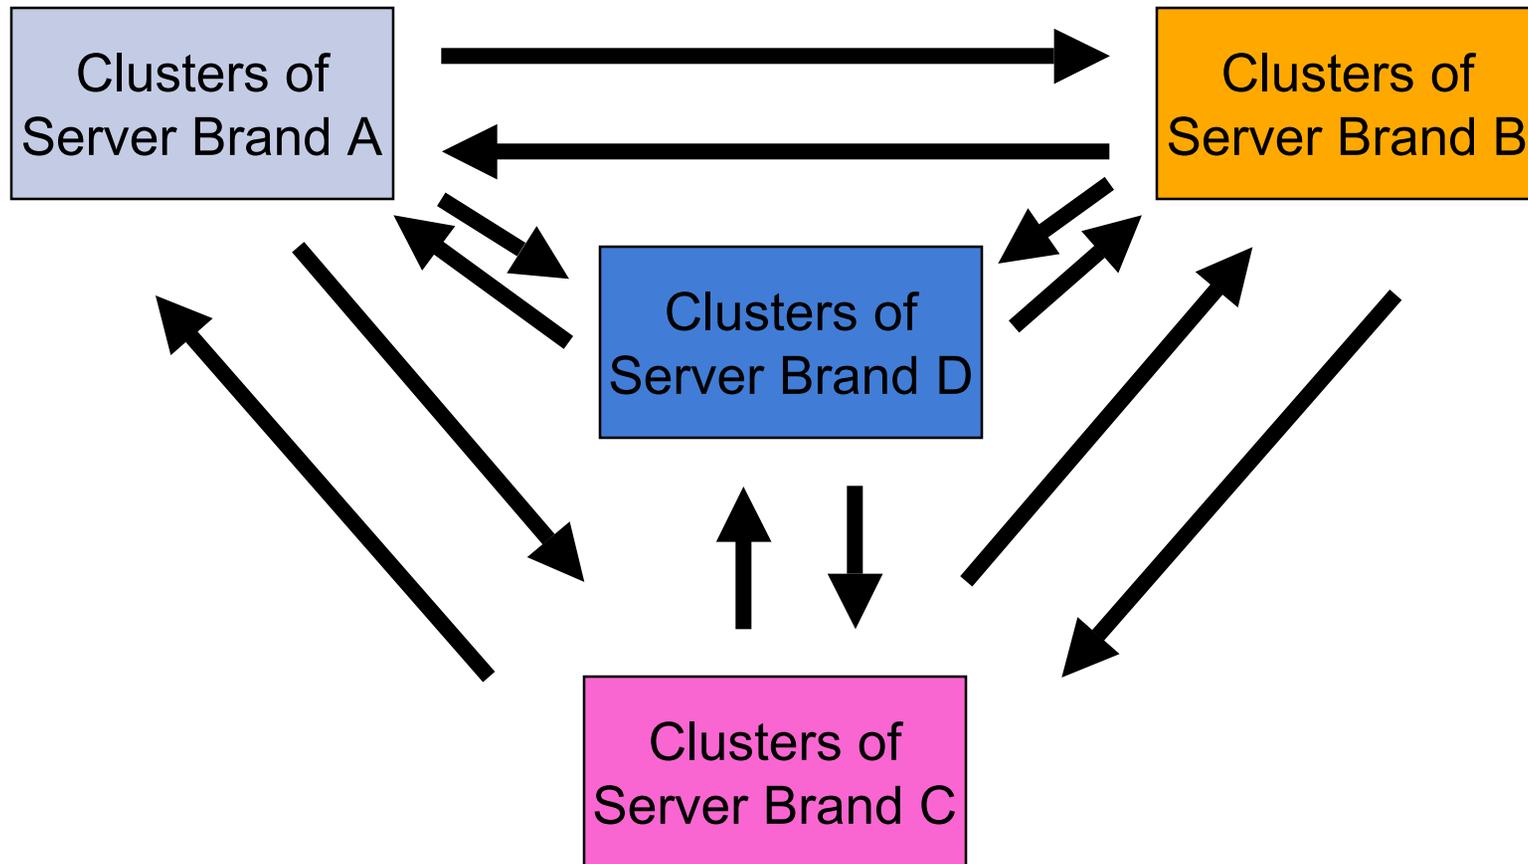
# Perhaps Brand B returns the favor



Clusters A and B are peers

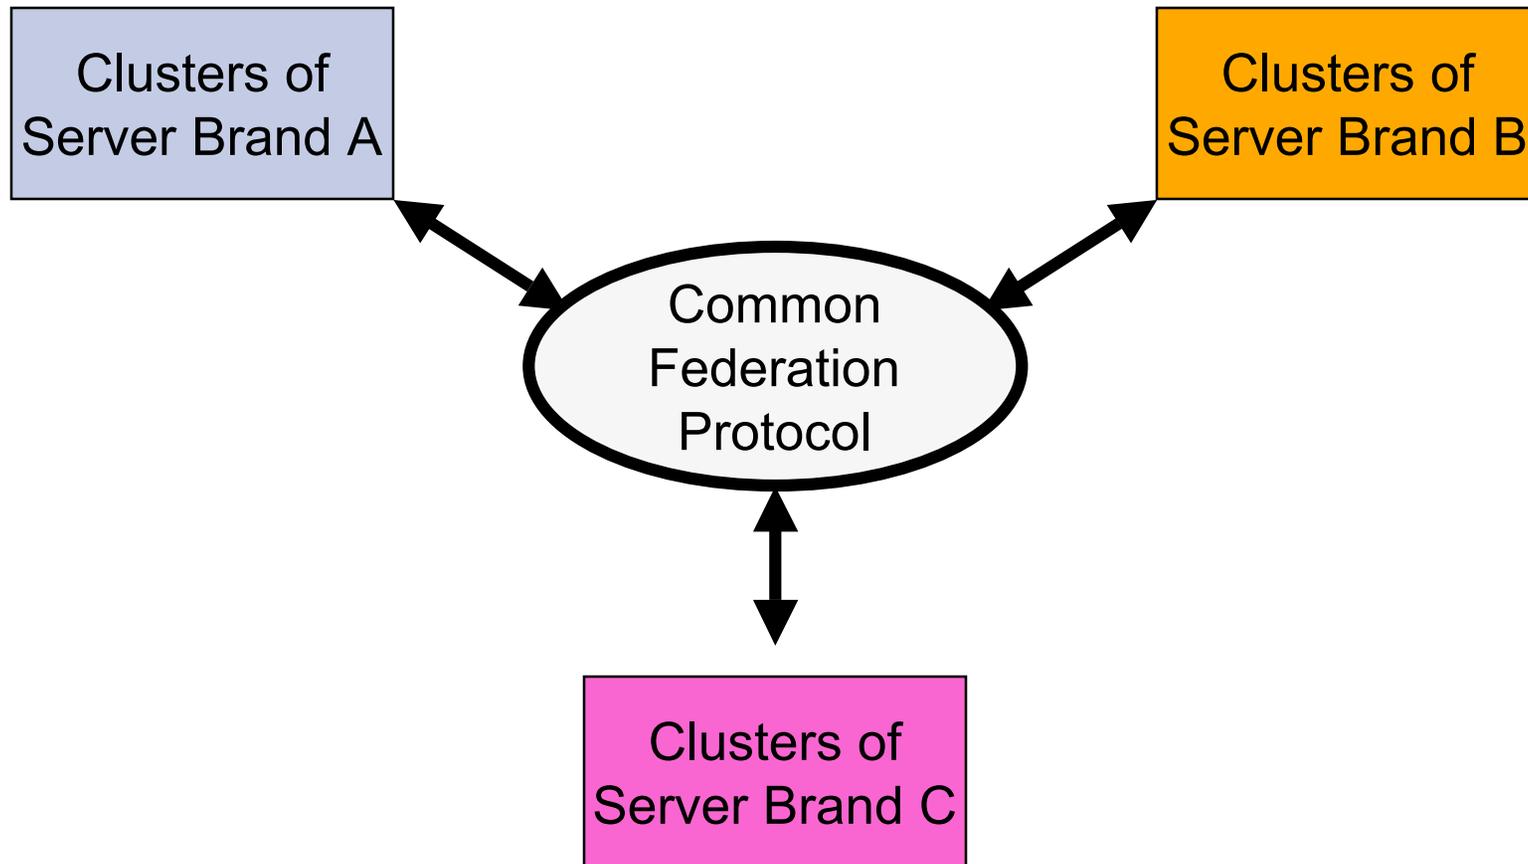
# But then Brand C comes along...





## This is no way to build a federation

- ▶ **For N cluster types, each cluster has N-1 peer relationships to spec, build, test, document, support...**
- ▶ **We don't have time for this.**
  - Not in this WiP
  - Not in our release cycles



- ▶ **Extends the standard protocol suite.**
  - Customers buy from their preferred vendor and use their preferred software.
  - Customers are not locked in.
  
- ▶ **Reduces the implementation effort.**
  - Interfaces are standard and well-defined.
  - Existing federation members do not need to be changed or upgraded in order to support new federation members.
  - Nobody is locked out.

- ▶ **Client issues: which clients to support?**
  - CIFS and NFSv4 simplify federation.
  - Many clients are still happily using NFSv3.
- ▶ **Admin issues: features versus commonality.**
  - Can one type of federation member, with permission from a second, administer the second?
  - Can we make things simpler via federation, or are we adding another tool that admins have to learn?
    - At the very least, creating a cross-cluster junction should be easy.

▶ **Current focus:**

- Mechanisms to build a platform-agnostic, global namespace.

[ellard@netapp.com](mailto:ellard@netapp.com)

**Come to the Federated File System BoF**

**Tonight, 8-9 in Ballroom V/VI**

**Enter to win an iPod nano 4GB**