



NetApp™

Go further, faster™

Referrals in Clustered ONTAP

Theresa Raj
NetApp



Agenda

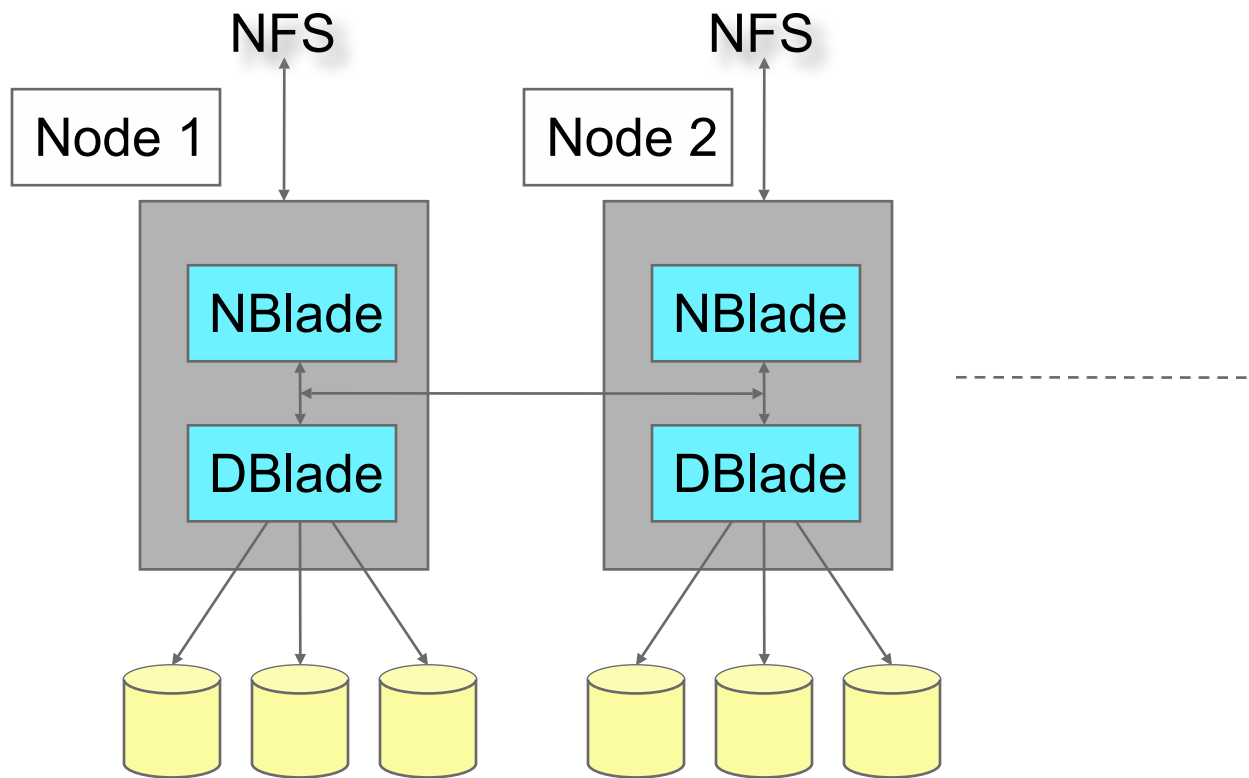
- Clustered ONTAP
- Scope of Referral
- Where does Referral occur
- How does Referral occur
- Client issues
- Federated FS in Clustered ONTAP



Clustered ONTAP

- Vserver
 - Provides a single file system namespace
 - Can span multiple physical nodes
 - Has one or more virtual interfaces
- Junction
 - Mount point of volume
 - Used to stitch the namespace
- NBlade
 - Network blade that handles all client facing protocols
- DBlade
 - Data blade that handles all file and block storage

Clustered ONTAP





Scope of Referral

- Intra-vserver only for now
 - NFSv4 client coming through one virtual interface to the vserver is redirected to another virtual interface within the vserver
- Criteria for intra-vserver referral
 - Existence of another virtual interface in the vserver that has greater proximity to the volume that the client is getting to
- Advantage of intra-vserver referral
 - no administrative setup for referrals is needed
 - we could defer all decisions on administrative setup for referrals and wait for Federated Filesystems protocols



Where does Referral occur

- At junction crossing; volume boundaries
- Client issues LOOKUP op
 - Looking up a junction
 - Target volume at the junction is closer to another virtual interface within the vserver
- Client issues READDIR op
 - Entry in directory is junction
 - Target volume at the junction is closer to another virtual interface within the vserver
 - Attributes requested beyond a limited set that includes `mounted_on_fileid`, `fsid`, `fs_locations`,



How does Referral occur

Client

=====

Compound ----->
PUTFH,LOOKUP,GETFH

Nblade

=====

Process PUTFH
(cfh is set)

Dblade

=====

Current Op is LOOKUP

component is junction

SpinNP lookup request ----->
SpinNP lookup response <--

JUNCTION-

Vldb Junction Lookup

Target Dblade is local? NO

Target Dblade has local VIF? YES.

cfh, compound tagged for Referral

Lookup Response NFS_OK

GETFH

cfh, compound tagged for Referral? YES



How does Referral occur (contd)

Client

=====

Compound ----->

PUTFH,LOOKUP,GETATTR(fs_locations) (cfh is set)

Nblade

=====

Process PUTFH

(cfh is set)

Dblade

=====

Current Op is LOOKUP
component is junction

SpinNP lookup request ----->
SpinNP lookup response <--

JUNCTION-

Vldb Junction Lookup

Target Dblade is local? NO.

Target Dblade has local VIF? YES.

cfh, compound tagged for Referral

Lookup Response NFS_OK

GETATTR (fs_locations)

Getattr Response with fs_locations attribute info

<-NFS4_OK-----
© 2008 NetApp. All rights reserved.



Client Issues

- Some clients see ERR_MOVED and fail with
 - EIO
 - EPROTO
- Need all clients to support referrals to enable Referrals
- Clients can't resolve fs_location.hostname
- Clients can't handle ERR_MOVED during mount



Federated FS in Clustered ONTAP

- Extend referrals outside the vserver into the Federated namespace
- Create junctions with junctionKey and nsdbName
- Lookup NSDB server for the FSL information of a FSN



References

- Implementation Guide for Referrals in NFSv4
 - <http://www.citi.umich.edu/projects/nfsv4/rfc/draft-ietf-nfsv4-referrals-00.txt>
- Data ONTAP GX: A Scalable Storage Cluster
<http://www.eisler.com/fast2007/>
- Requirements for Federated File Systems
 - <http://www.ietf.org/internet-drafts/draft-ellard-nfsv4-federated-fs-01.txt>
- NSDB Protocol for Federated Filesystems
 - <http://www.ietf.org/internet-drafts/draft-tewari-nfsv4-federated-fs-protocol-01.txt>