



ORACLE®

Solaris Referrals

Rob Thurlow
Solaris NFS Group



Starting Points

We had an old NFS referrals prototype

- Client support only

- Some bits were (gasp!) prototype quality

We had mirror mounts

- Always planned to drop in referrals

- Automatic mount & unmount worked

Biggest gap: server side

- Infrastructure to hand out referrals

- Basic administration model

Server Requirements

We wanted to support NFS and SMB

- Both have referrals

- Other uses, too e.g. an HSM system

Needed a persistent “junction” object

- Marked as needing special handling

- Carries purpose-specific data

Would prefer to not modify archivers

- So not a brand-new type of object

Server Solution

“Reparse point” - specially marked symlink

Kernel code detects with GETATTR

Archivers store as usual

Syntax to permit different types to coexist

Can do NFS and SMB referrals at same time

Syntax detected on creation

And also when archivers restore

Upcall to user-land with plugins

Location data direct from symlink or indirectly

Basic referrals now, FedFS later

Server Admin – nfsref utility

Needed a tool to manage NFS referrals

Share command felt clunky

SMB referrals usually managed from
Windows client

Came up with 'nfsref' utility

nfsref add path location [location ...]

nfsref remove path

nfsref lookup path

Client work

Upcall to handle hostnames

- Need IPv4/IPv6 address and transport

- Refactored some code from automounter

Mount triggering needed tweaks

- Same FSID crossing logic as mirror mounts

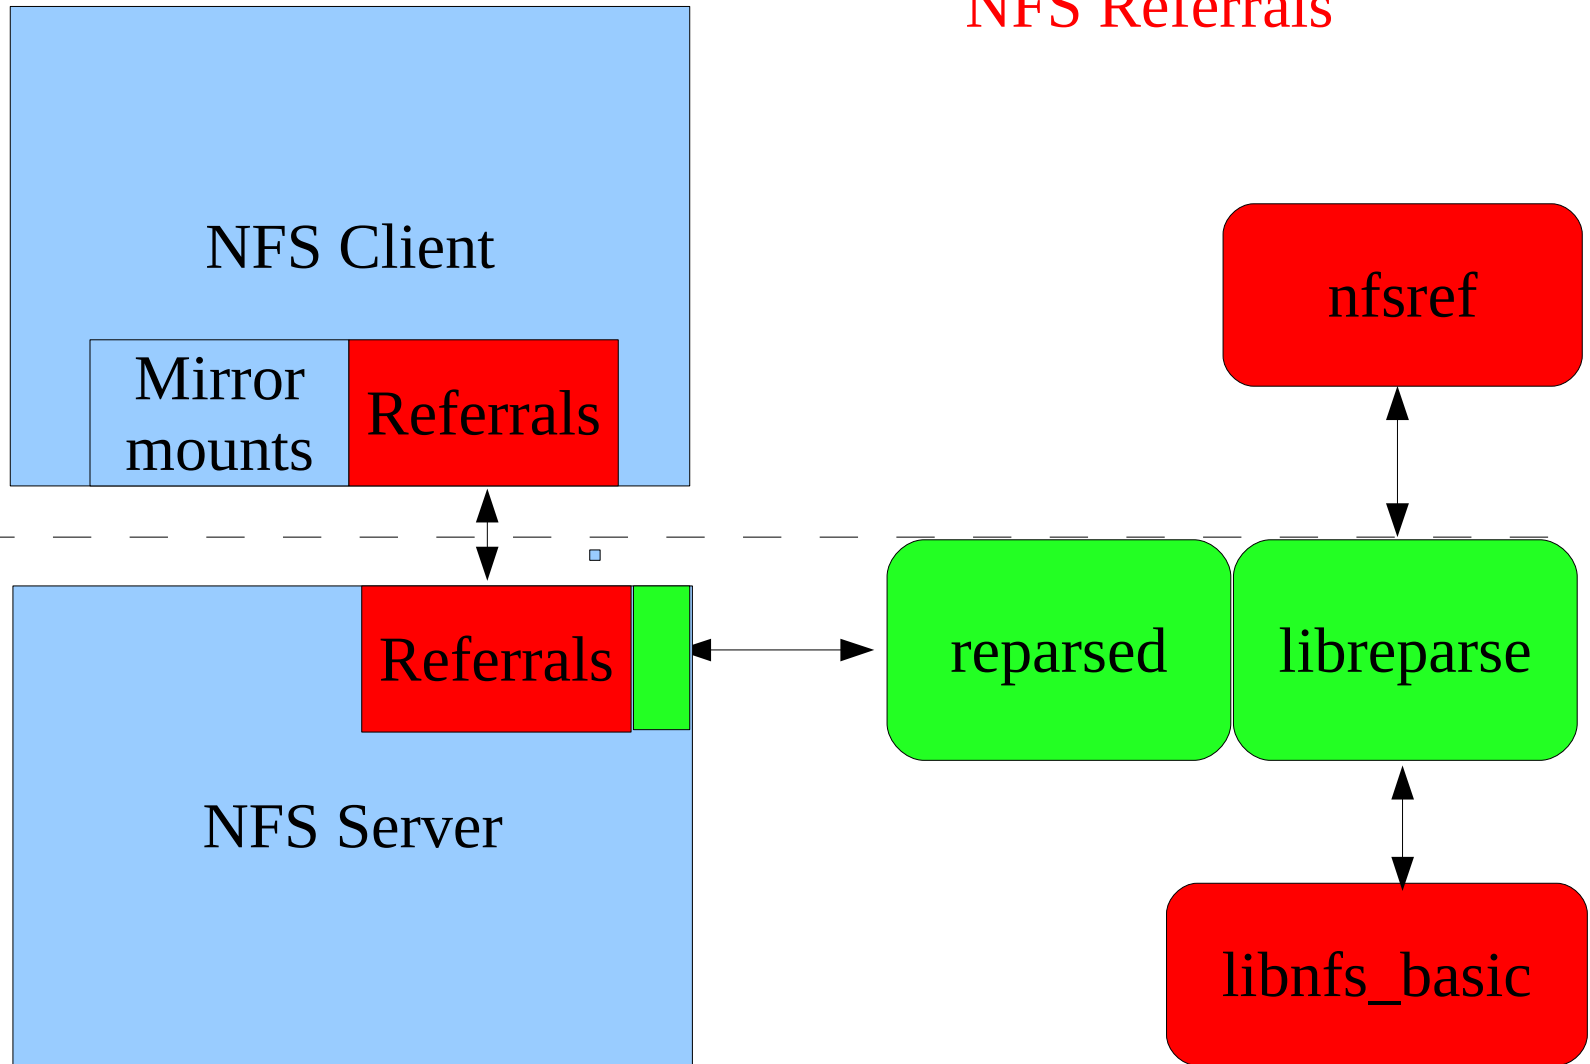
- Had to remove assumption about mounts coming from same server

- Had to teach GETATTR to fabricate info

Had to supply new referrals-aware code

New Functionality

Reparse points
NFS Referrals



Example

Creating a referral:

```
% nfsref add /test/ref jurassic:/home/thurlow
```

```
Created reparse point /test/ref
```

```
% ls -l /test/ref
```

```
lrwxrwxrwx 1 thurlow staff 45 Feb 5 15:28 /test/ref ->  
@{REPARSE@{nfs-basic:jurassic:/home/thurlow}}
```

Accessing a referral:

```
# mount trianon:/test /mnt
```

```
% cd /mnt/ref
```

```
% df .
```

```
Filesystem          kbytes  used  avail capacity  Mounted on  
jurassic:/home/thurlow
```

```
4977286087 9164965 4968121122 1% /mnt/ref
```


Referrals Status

Tested at October Bake-a-thon

NFS support integrated in December'09

OpenSolaris build 131

Tested here at Connectathon, too

Enjoy!