

NFSv4.2 Sparse files

Marc Eshel

11/11/09

eshel@almaden.ibm.com

NFSv4.2 Sparse files

- I propose to add Sparse file support to NFSv4.2
- Sparse file are common way to represent huge files
 - Database files
 - HPC applications
 - Virtual machine images
- Application are not aware of file organization
 - Read holes
 - Prefetch holes

NFSv4.2 Sparse files

- Simple change to NFS read RPC
- New return code when holes are read
 - Don't return data that is all zeros
- Return extra information with new return code
 - Next data offset
 - Length of data

Changes to the READ rpc:

Operation 25: READ - Read from File

...

RESULT

```
    struct READ4resok {
        bool        eof;
        opaque      data<>;
    };

+   struct READ4reshole {
+       offset4    data_offset;
+       count4     data_count;
+   };

    union READ4res switch (nfsstat4 status) {
        case NFS4_OK:
            READ4resok    resok4;
+       case NFS4ERR_HOLE:
+           READ4reshole  reshole4;
        default:
            void;
    };
```

DESCRIPTION

...

- + If return code NFS4ERR_HOLE is returned than entire read range can be
- + assumed to be zero.
- + The offset in READ4reshole might be set with the location of the next
- + chunk of data and count is the length of the chunk. If count is set to
- + zero just ignore the information in READ4reshole.