

**Subject:** Pathname Reparse Points  
**Submitted by:** Afshin Salek  
**File:** PSARC/2009/387/opinion.ms  
**Date:** August 5th, 2009  
**Committee:** Glenn Skinner (opinion written by Mark Martin), Garrett D'Amore, Mark Carlson, Richard Matthews, Sebastien Roy.

**Product Approval Committee:**

Solaris PAC  
solaris-pac-opinion@sun.com

**1. Summary**

This is one of a series of projects under the umbrella case: PSARC 2009/399 "Reparse Points and Referrals Umbrella Case".

There are situations where a mechanism is needed to reflect the concept that data is not present at a particular path, but can be found in some alternate location(s). Examples include "referrals" used to build unified name spaces in NFSv4.x and SMB, and data relocation in HSM systems. A "reparse point" is defined as the marker for a namespace redirection and a container for the metadata to specify where the target of this redirection is.

Reparse points are intended to be a general mechanism for location redirection and as such the file system that contains them is not cognizant of the reparse point format or content. Services that use reparse points know how to interpret and use the stored data.

**2. Decision & Precedence Information**

The project is approved as specified in reference [2].  
The project may be delivered in a minor release of Solaris.

**3. Interfaces**

The project exports the following interfaces.

Interfaces Exported		
Interface	Classification	Comments
VOP_LOOKUP, fop_lookup	Consolidation Private	Added new argument: vattr_t *vap
XAT_REPARSE	Consolidation Private	Reparse extensible attribute
symlink context syntax for Reparse points	Committed	
reparse_kderef reparse_init	Consolidation Private	Reparse kernel routines

Interfaces Exported		
Interface	Classification	Comments
reparse_parse reparse_free  reparse_init reparse_parse reparse_add reparse_remove reparse_unparse reparse_free reparse_create reparse_delete reparse_deref	Consolidation Private	Reparse library routines
reparse_plugin_ops_t	Consolidation Private	Reparse plugin ops
/usr/lib/reparse/reparsed	Consolidation Private	Reparse daemon
/usr/lib/libreparse.so	Consolidation Private	Reparse library
/var/run/reparse_door	Consolidation Private	User space door
svc: /system/filesystem/reparse	Uncommitted	SMF Service

## 4. Opinion

### 4.1. POSIX Conformance

The key issue in the case was the question of whether or not the presence of a reparse point affects POSIX conformance.

The problem is that a POSIX-conformant application could create a symlink whose contents match the reparse point syntax. Subsequently, that, or another POSIX-conformant, application could do a pathname traversal through that symlink, have it be treated as a reparse point (thereby inducing its interpretation, perhaps as an NFS referral), and experience normal symlink semantics (either a traversal failure or having the traversal resolve to a file whose name matches the reparse point's contents).

The project team noted that this situation can only arise when a non-POSIX conformant service (such as NFS or SMB) has control over the part of the file system name space containing the reparse point. The team further claimed that this involvement of a non-POSIX conformant service removes the traversal from the scope of activities subject to POSIX requirements.

After some discussion, the committee accepted this argument.

### 4.2. Follow-on Consumers

The question was raised whether this case should have a dependency on one of the follow-on cases mentioned in 2009/399 [Reparse Points and Referrals Umbrella Case] to ensure that a consumer of this project's deliverables is delivered in tandem with this project itself.

The project team argued that this case provides value in and of itself, and the committee members accepted this argument

### **4.3. Specification Changes**

#### **4.3.1. Release Binding**

The release binding changed from update to minor. This change was triggered by observing that the change to the signature of VOP\_LOOKUP() was not suitable for an update binding.

#### **4.3.2. Stability Level of the Reparse Format**

The stability level of the in-symlink reparse point format changed from Committed Private to Committed. This change was motivated by noting that there should be a way to know that one has not inadvertently created a reparse point when creating a symlink.

### **4.4. Reparse Point Enable/Disable**

The committee discussed whether adding some mechanism to enable and disable reparse points, perhaps on a file system-wide basis would help alleviate the compatibility concerns illustrated by the previous point.

The consensus among the committee and project team members was that this proposed cure was worse than the disease.

### **4.5. Other Minor Clarifications**

There were various other minor points of confusion (some relating to terminology). Where appropriate, the project team updated the specification to incorporate the clarifications.

### **5. Minority Opinion(s)**

None.

### **6. Advisory Information**

#### **6.1. Auditing in Follow-on Projects**

Any follow-on project that provides a service employing reparse points should verify that referral data provided in response to triggering a reparse point is properly audited on the client system that triggered the reparse point.

#### **6.2. Upcoming NFS Referrals Case**

As part of preparing the materials for the ARC review of the upcoming NFS referrals case, committee members advise the project team to provide a step by step account of what happens when pathname traversal encounters a referral.

### **7. Appendices**

#### **7.1. Appendix A: Technical Changes Required**

None.

#### **7.2. Appendix B: Technical Changes Advised**

None.

#### **7.3. Appendix C: Reference Material**

Unless stated otherwise, path names are relative to the case directory PSARC/2009/387.

1. Onepager.  
File: 20090709\_afshin.salek

2. Project Specification.  
File: final.materials/reparse\_spec.txt
3. Commitment Minutes.  
File: commitment-notes
4. PSARC 20 Questions.  
File: final.materials/reparse\_20q.txt
5. Global namespace: The future of distributed file server management.  
[http://findarticles.com/p/articles/mi\\_m0BRZ/is\\_2\\_23/ai\\_98709766/](http://findarticles.com/p/articles/mi_m0BRZ/is_2_23/ai_98709766/)
6. Implementation Guide for Referrals in NFSv4.  
<http://tools.ietf.org/html/draft-ietf-nfsv4-referrals-00>
7. DFS Technical Reference.  
<http://technet.microsoft.com/en-us/library/cc757042%28WS.10%29.aspx>