

Case Study: MS-DFS

Writing New Drivers to Extend Likewise-CIFS

Gerald Carter

<gcarter@likewise.com>

Director of Engineering

Likewise Software

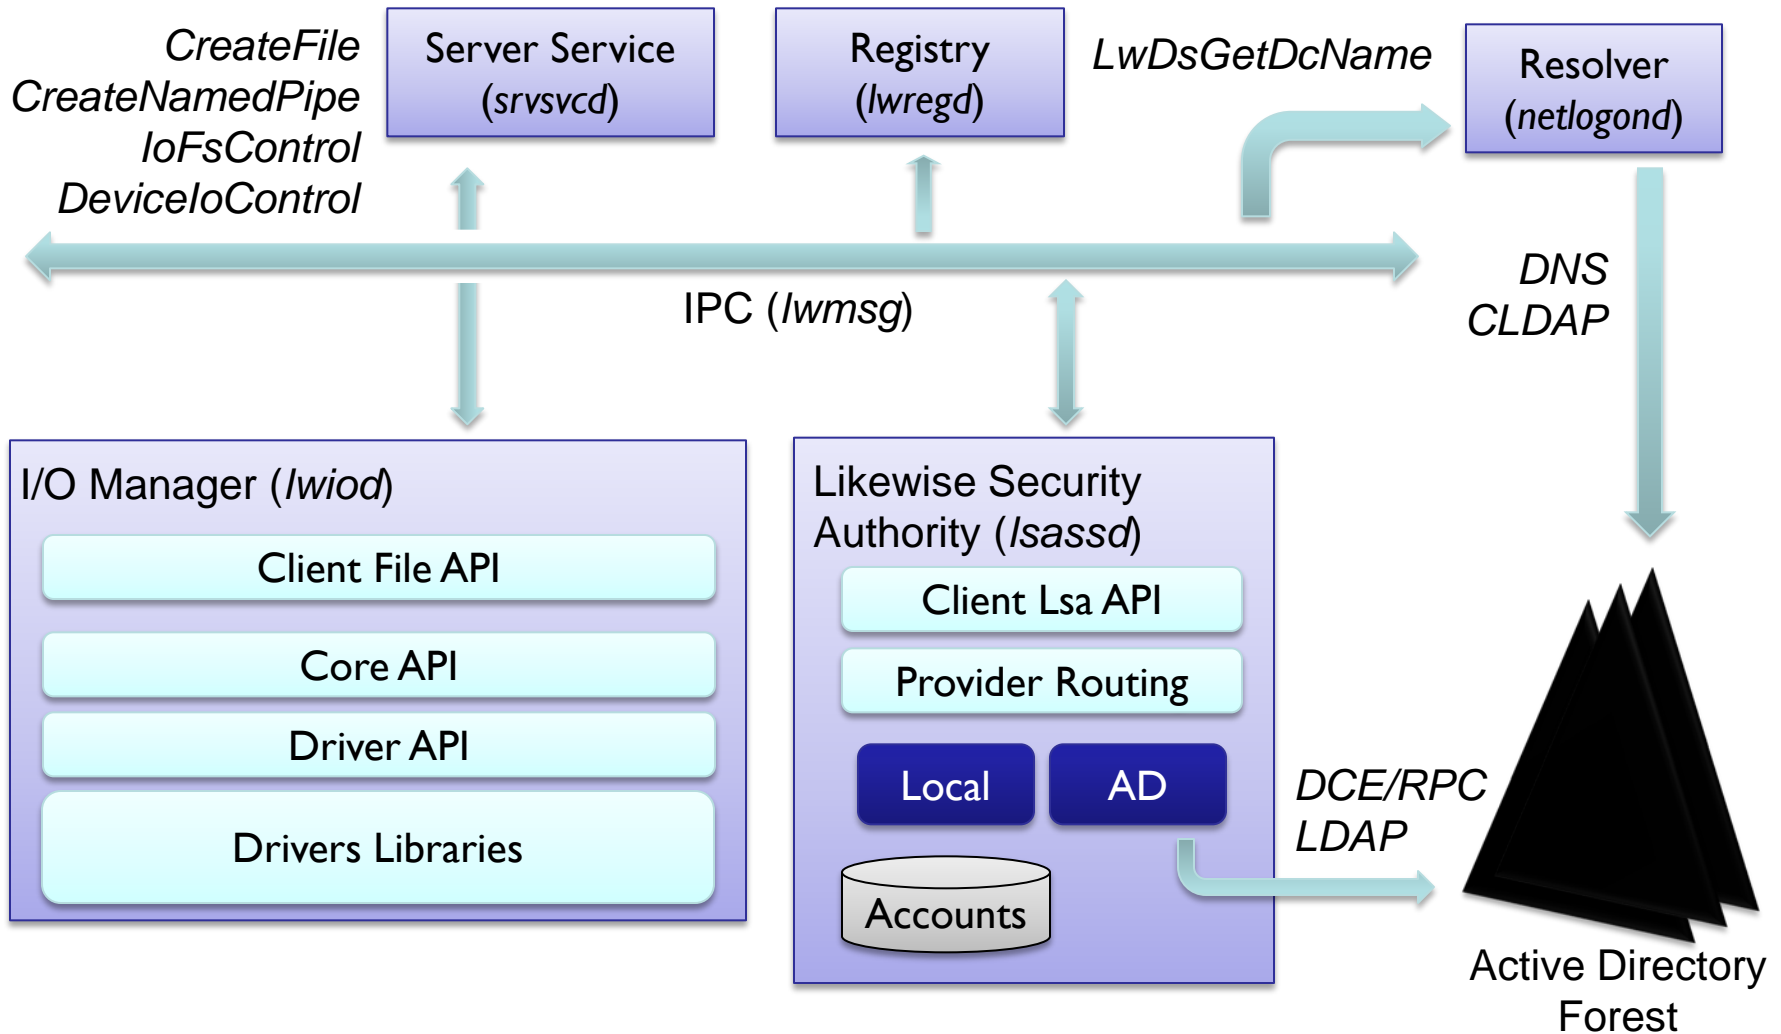
Likewise Open – Background

<http://www.likewiseopen.org/>

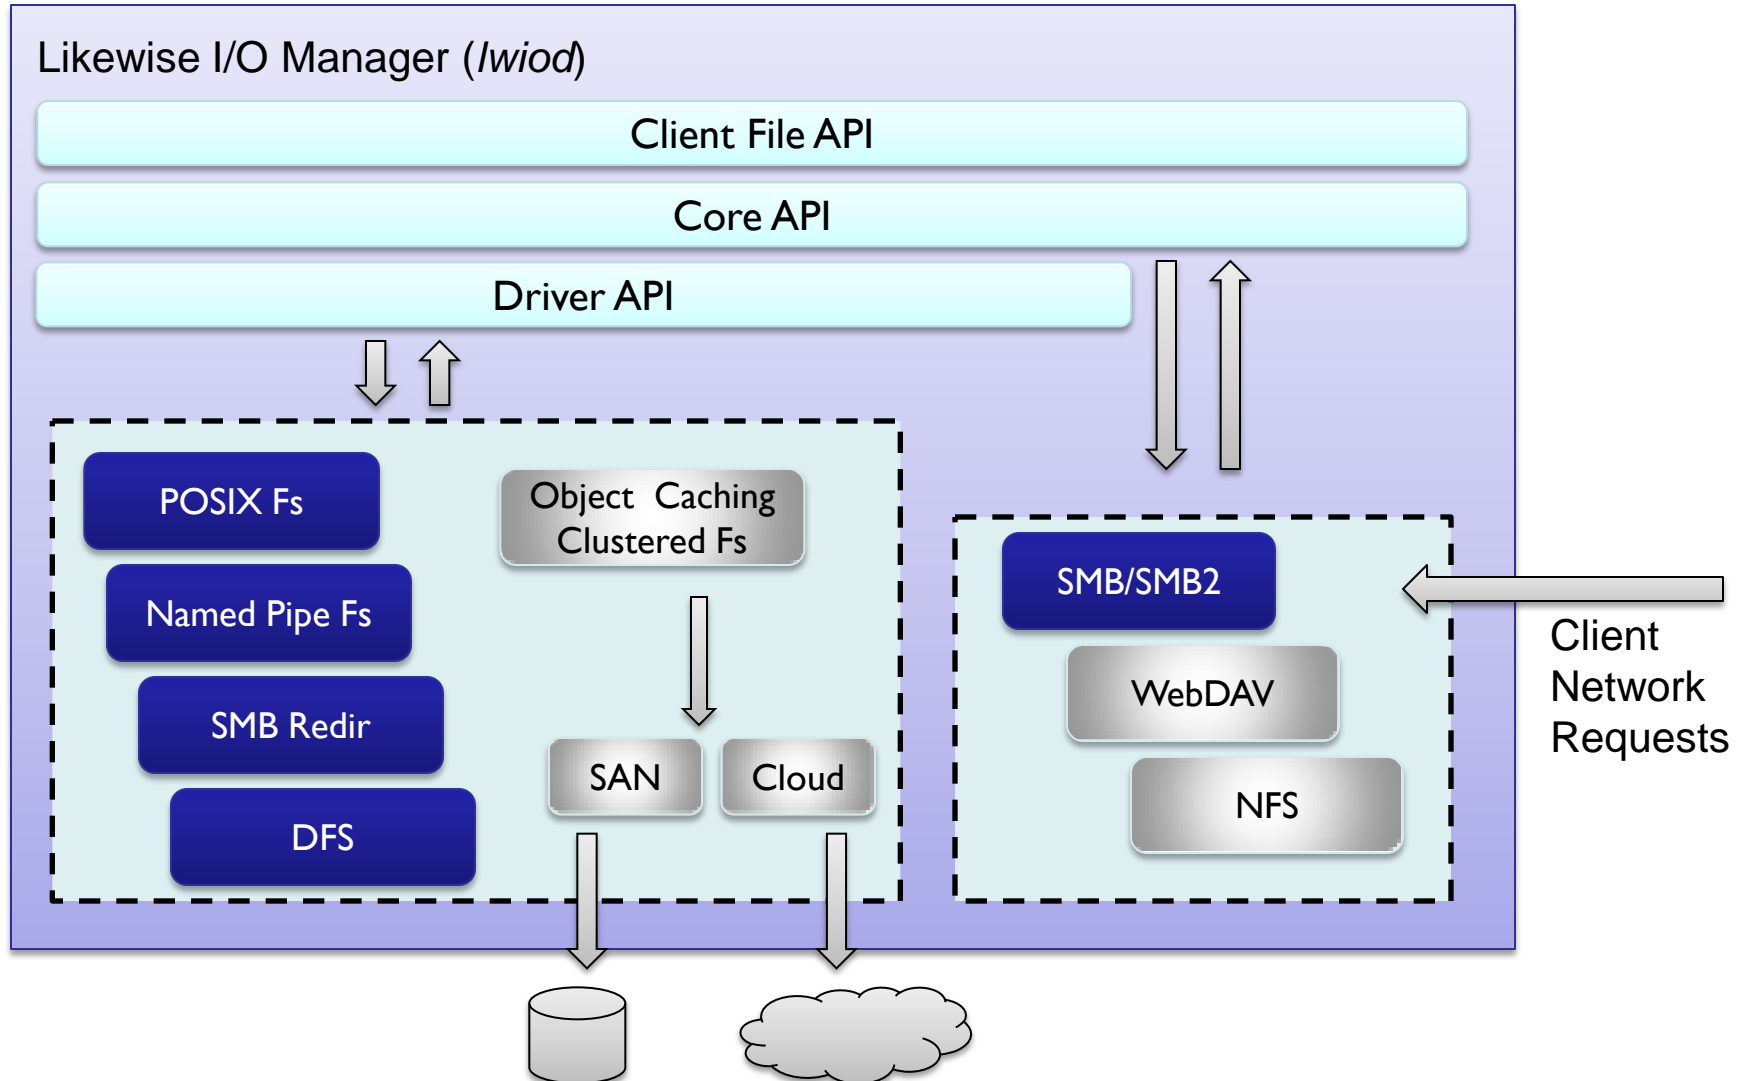
- ❑ Likewise Open Project is the umbrella project sponsored by Likewise Software designed to provide an interoperability platform for non-Microsoft clients and servers in Microsoft OS dominated networks.
- ❑ Likewise Open (product) refers to the open source authentication & Active Directory integration suite
- ❑ Likewise-CIFS is the file server software stack
 - ❑ Currently includes support for SMB and SMB2
- ❑ License
 - ❑ Choice: Commercial or GPLv2+/LGPLv2.1+
 - ❑ Single code base

- ❑ Describe the Likewise I/O Manager driver model
- ❑ Explain the SMB/SMB2 SRV driver representation of File Shares
- ❑ Show how SRV supports multiple File System Drivers (FSD)
- ❑ Understand a standalone MS-DFS Namespace
- ❑ Learn how to configure shares to use a specific FSD

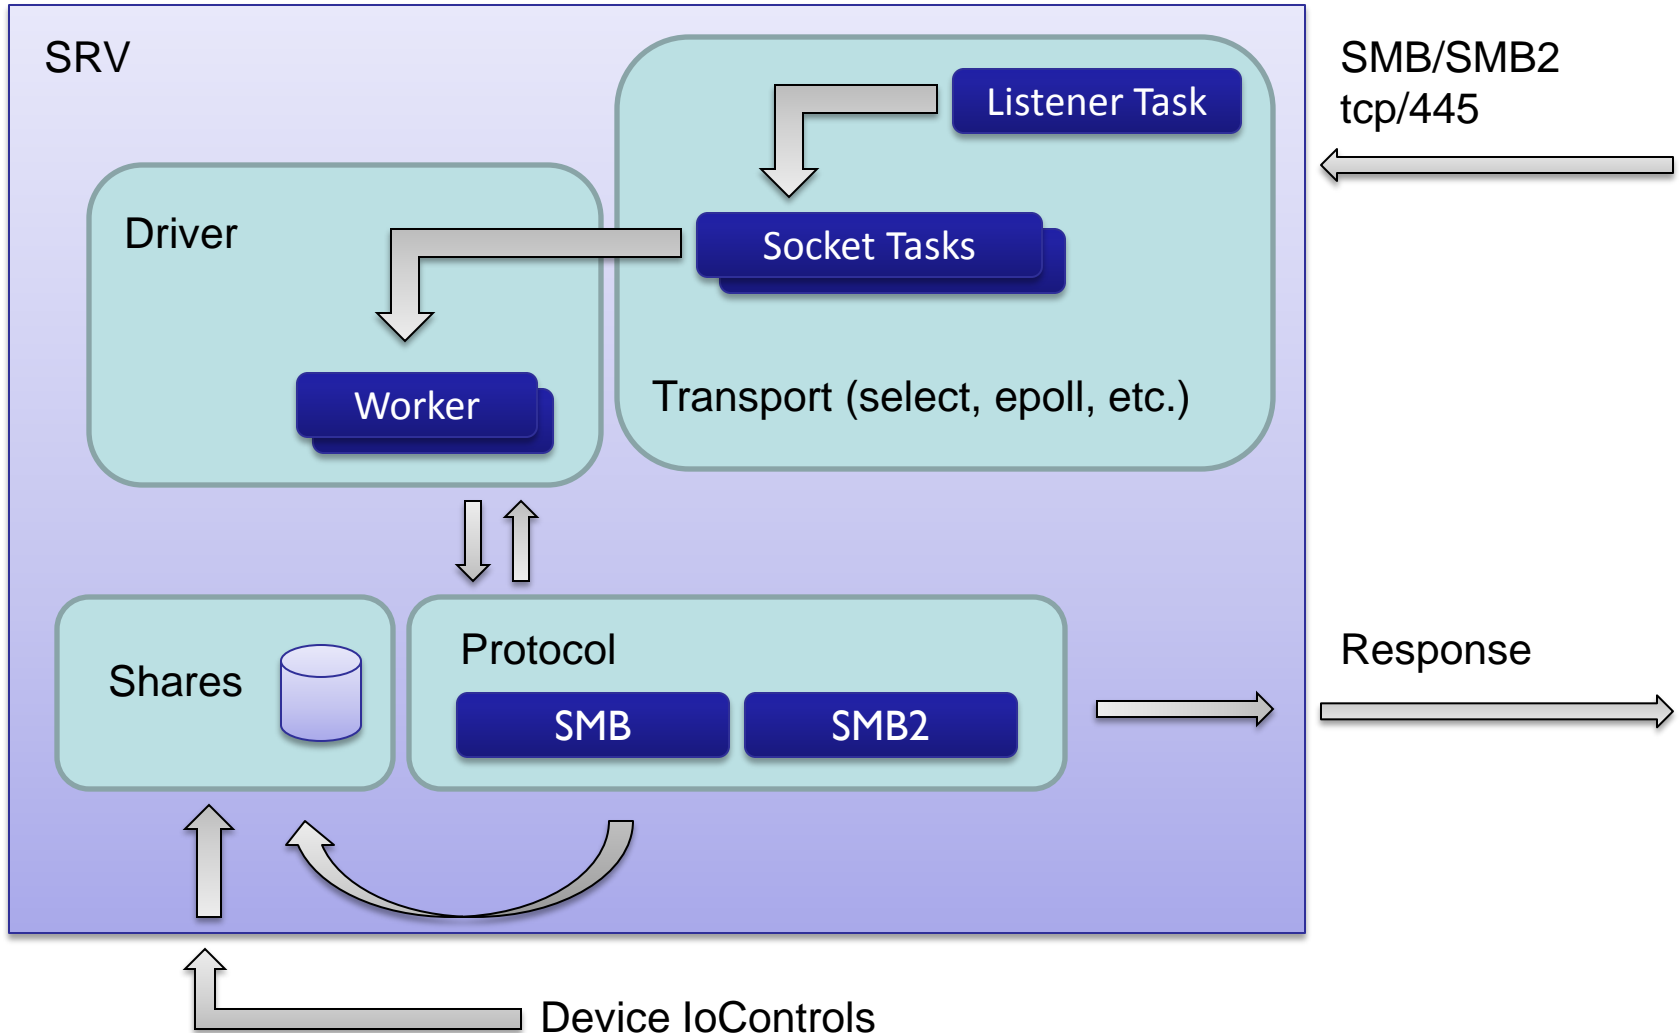
Likewise-CIFS Core Components



Likewise I/O Manager



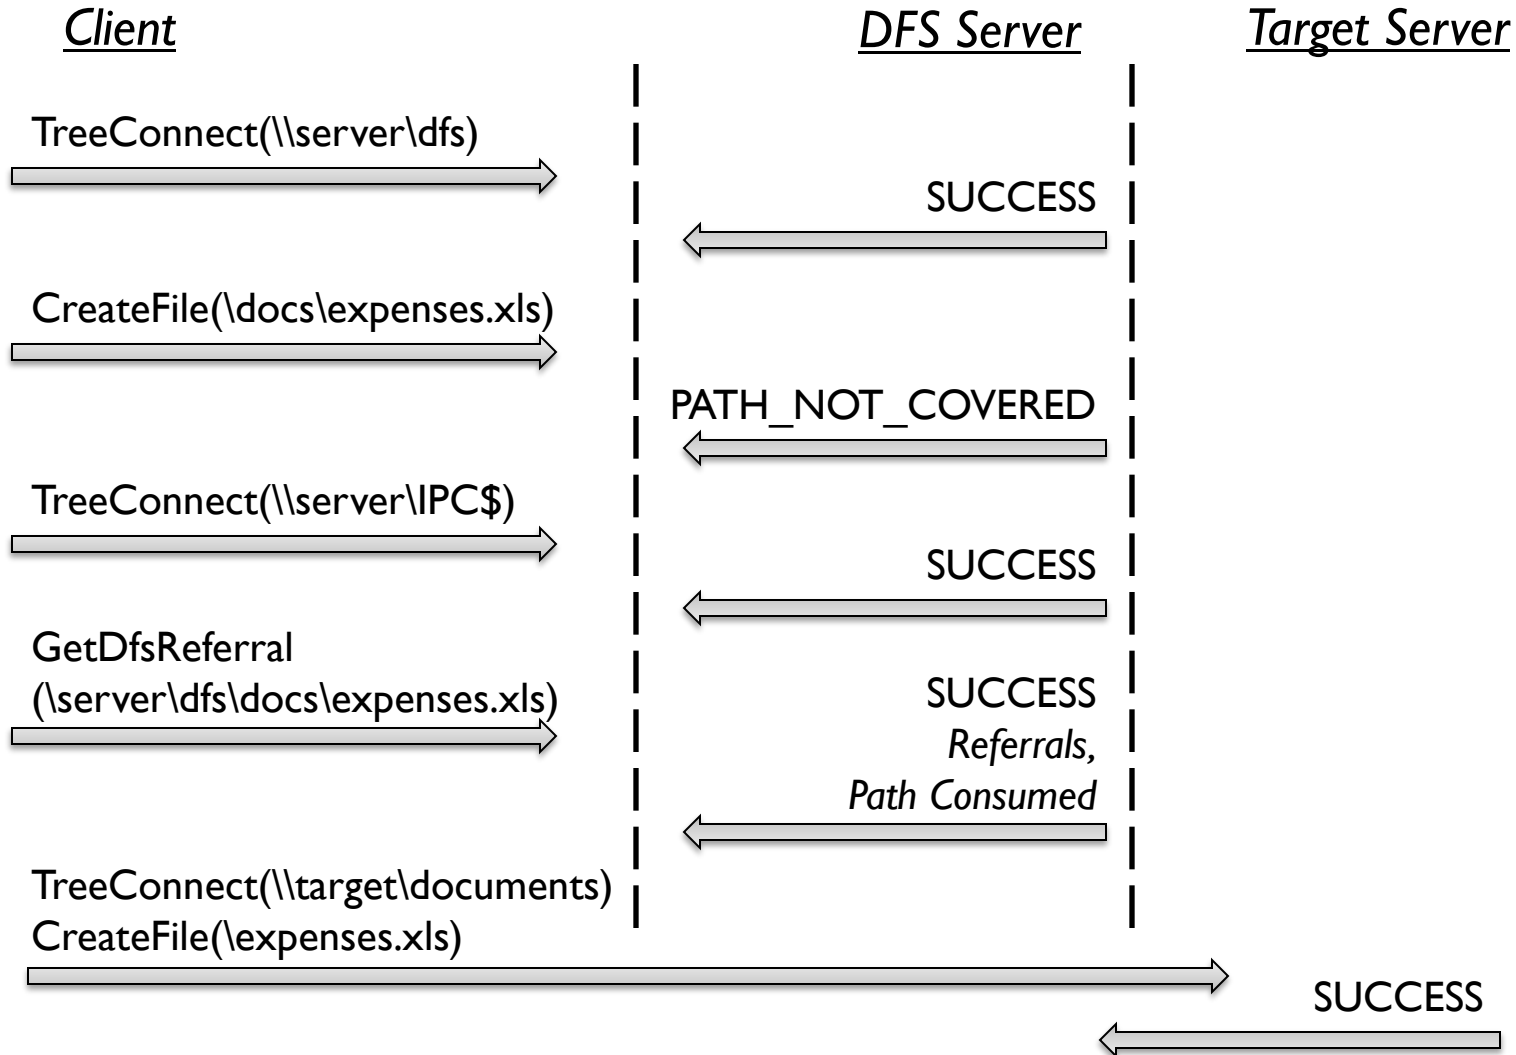
SRV (SMB/SMB2) Architecture



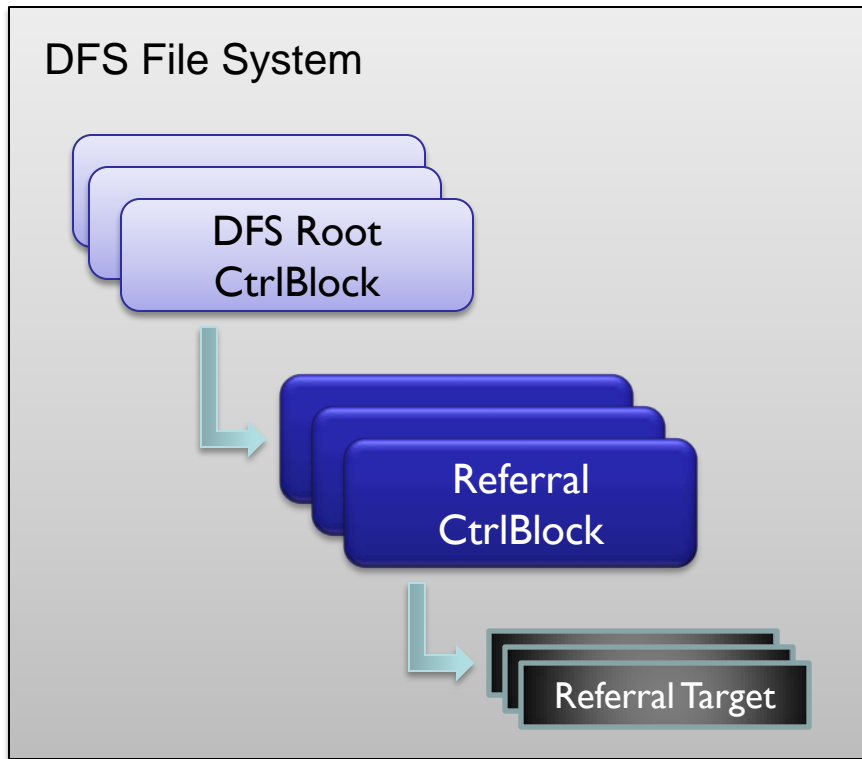
- ❑ Every driver registers a namespace
 - ❑ E.g. “/pvfs” “/npfs” “/dfs” “/rdr” “/srv”
- ❑ The SRV driver embeds the FSD namespace within the share’s absolute path
 - ❑ Binds a share to a given FSD
 - ❑ E.g. The share “Public” is bound to the I/O Manager path on disk as “/pvfs/data/public”
 - ❑ Client opens are relative to the Share root

- ❑ *libdfs.sys.so* – Driver library loaded by *lwiod*
- ❑ Peered with POSIX and Named Pipe FSDs
- ❑ A DFS Root Share is a share bound to the DFS FSD
 - ❑ All requests to the share are routed to the DFS driver
- ❑ Implements an ephemeral file system representing
 - ❑ Standalone DFS Root
 - ❑ A DFS Namespace held by a DFS Root
 - ❑ Link Referrals for each Namespace entry
- ❑ GetDfsReferral SMB/SMB2 operations on IPC\$ become DFS DeviceIoControls

MS-DFS Conversation



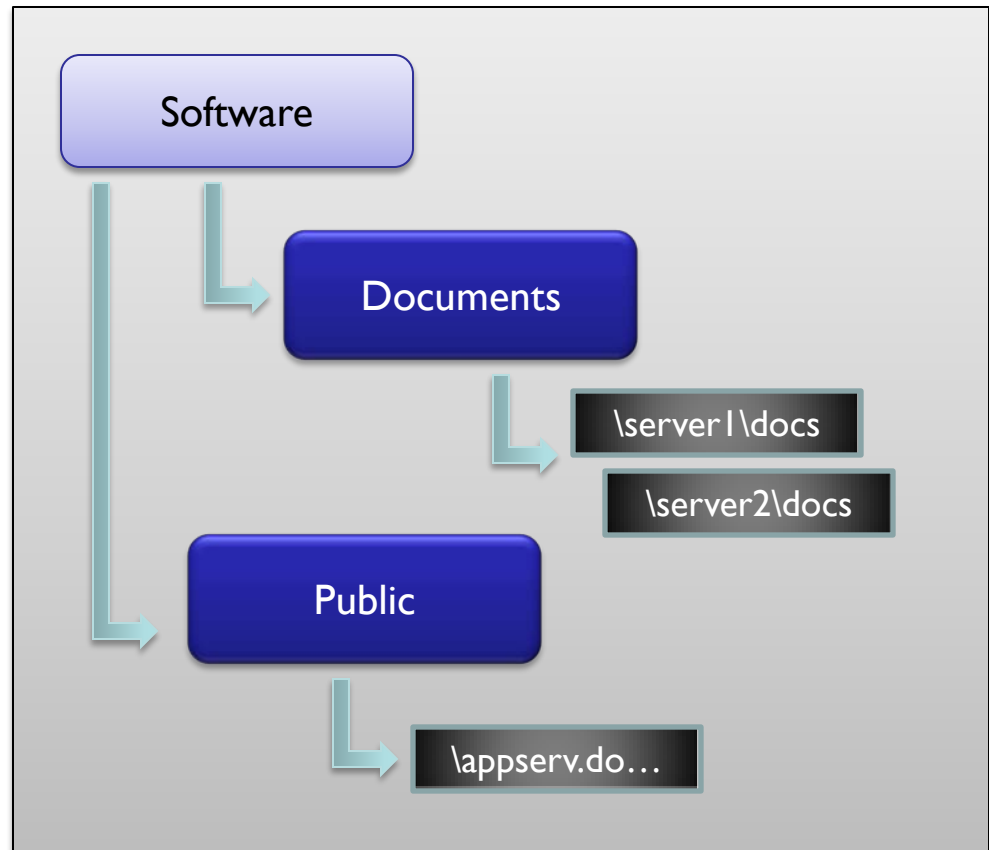
Likewise DFS – File System



- ❑ DFS Root Control Block holds a list of referrals
- ❑ Referrals are displayed as directories from `IoQueryDirectoryInfo()` and contain a list of referral targets
 - ❑ Registry value name
- ❑ Referral Targets are in the form `<UNC>:<TTL>`
 - ❑ Registry value data

DFS Registry Settings

```
[HKTM\Services\lwo\Parameters\Drivers\dfs\Roots\Standalone\Software]  
"Documents"=sza:"\\server1\docs:1800" "\\server2\docs:1800"  
"Public"=sza:"\\appserv.domain.com\public\software:1800"
```



❑ IRP_TYPE_CREATE

- ❑ Must match a DFS_CREATE_CTRL_BLOCK or a REFERRAL_CTRL_BLOCK and be a FILE_OPEN
- ❑ Partial namespace match returns PATH_NOT_COVERED
- ❑ Otherwise return OBJECT_PATH_NOT_FOUND

❑ IRP_TYPE_CLOSE

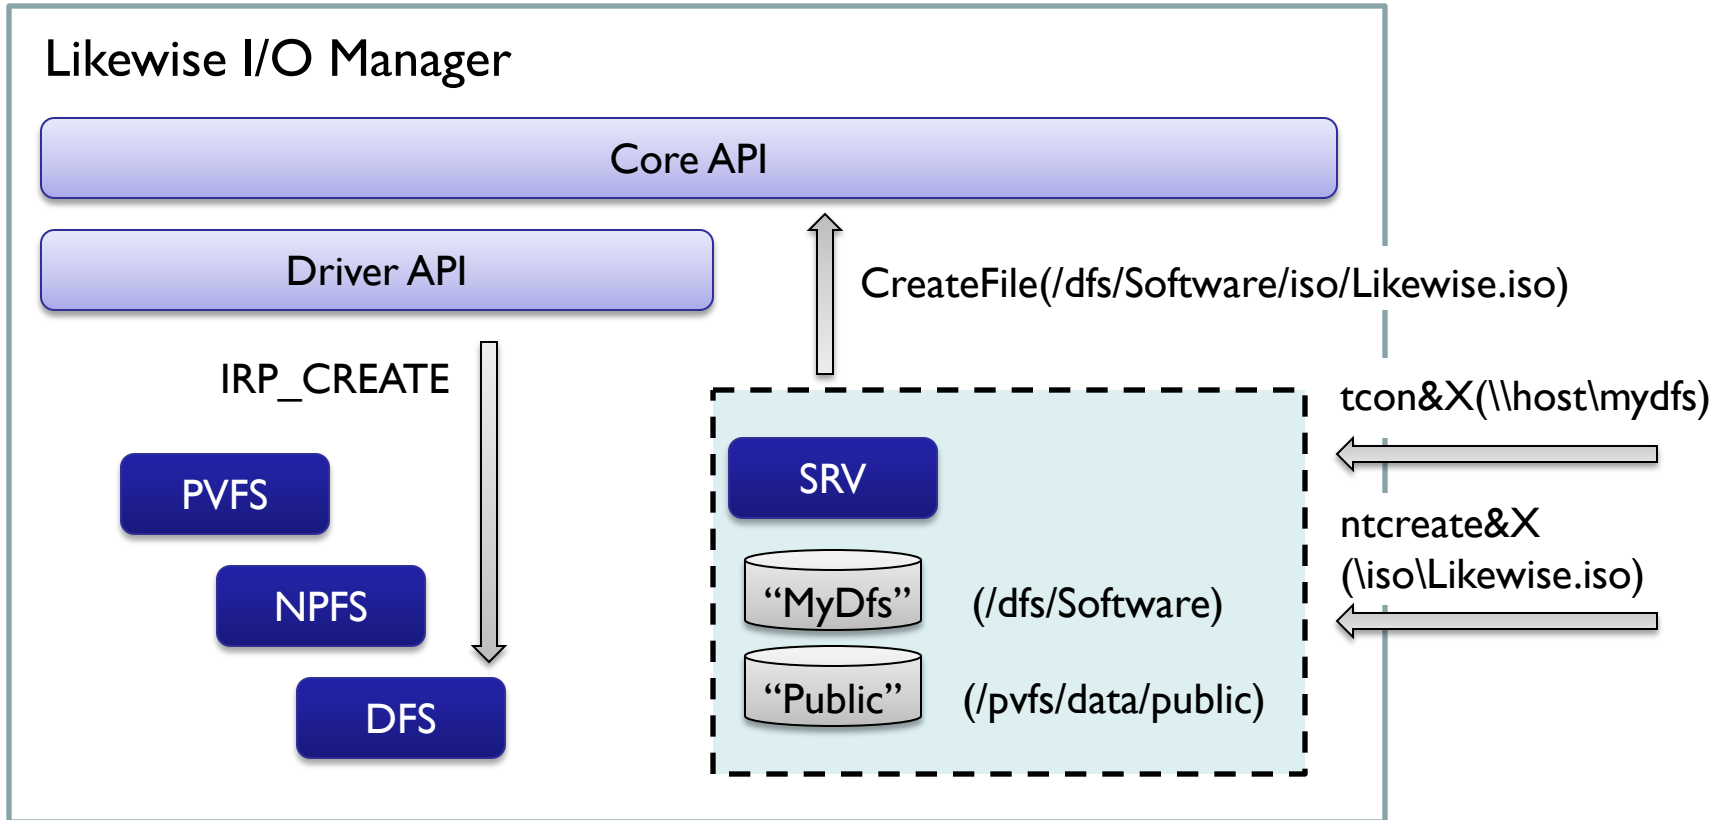
- ❑ Standard handle close

- ❑ **IRP_TYPE_QUERY_DIRECTORY**
 - ❑ If handle is to a **ROOT_CTRL_BLOCK**, then enumerate referrals control blocks

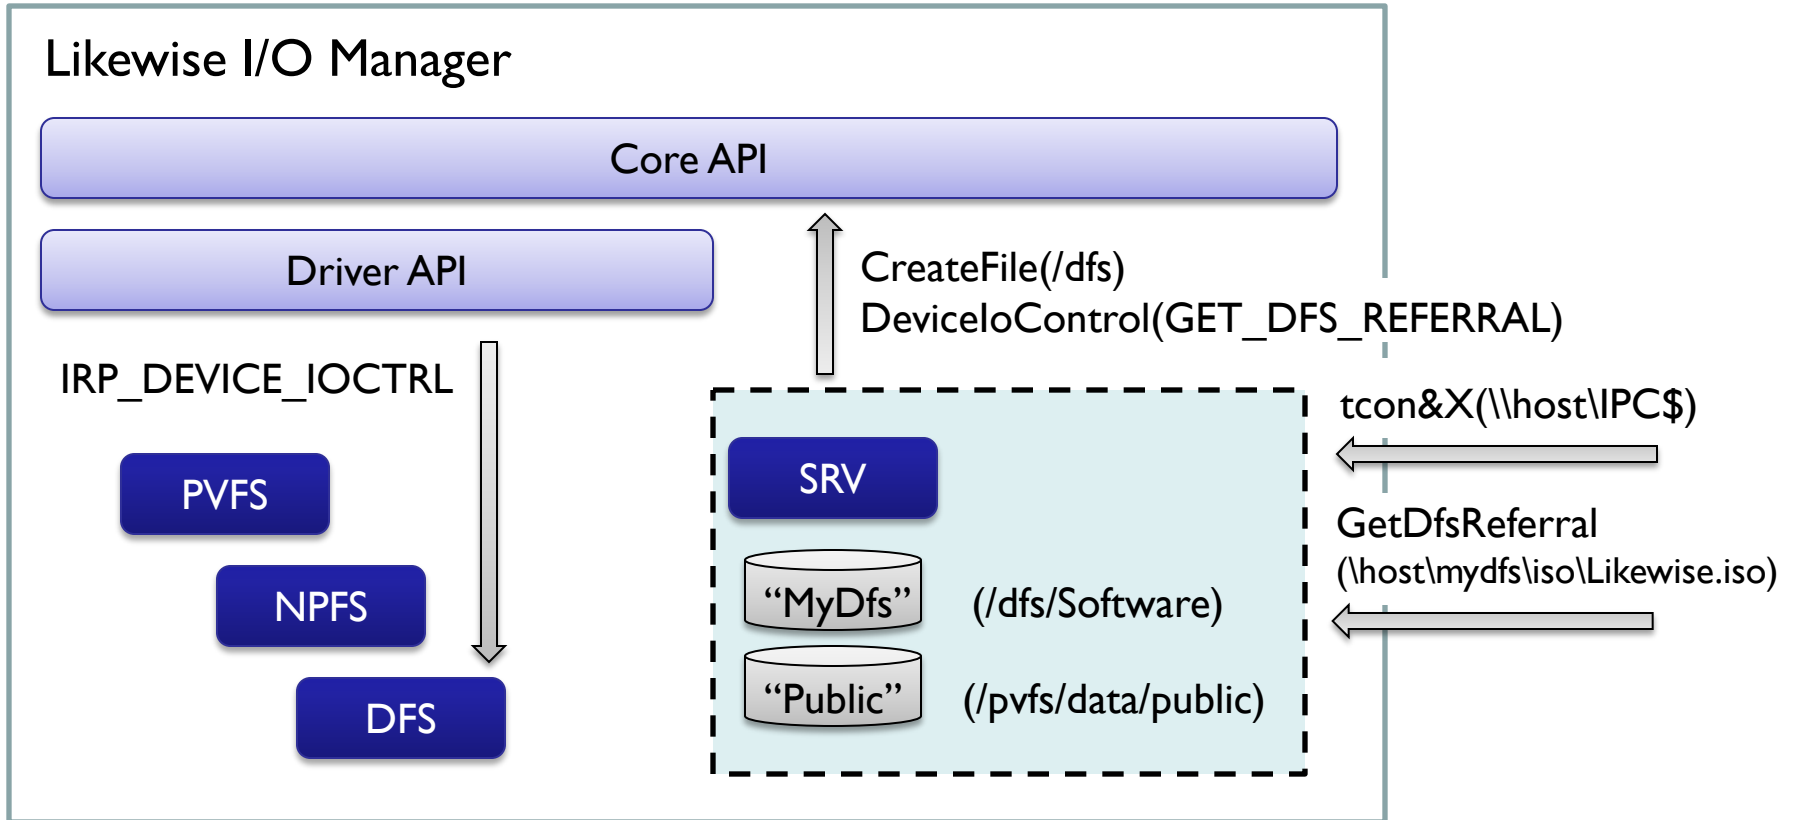
- ❑ **IRP_TYPE_QUERY_INFORMATION**
 - ❑ DFS Root Ctrl Blocks are reported as reparse points

- ❑ **IRP_TYPE_DEVICE_IO_CONTROL**
 - ❑ Interface to GET/SET **DFS_REFERRAL** info structures
 - ❑ **SMBTrans2GetDfsReferral**, **SMB2_IO_CTL**

DFS – Putting the Pieces Together



DFS – Putting the Pieces Together



Caveat – Windows Share Mgmt

- ❑ Driver namespace is not exposed through the NetrShareXXX() APIs
 - ❑ MMC “Shared Folders” plug-in
- ❑ “Default” FSD is used for new shares
 - ❑ Manual path configuration is possibly by directly provisioning the registry
 - ❑ [HKTM\Services\lwwio\Parameters\Drivers\srvc\shares]
- ❑ Administrative interface for standalone DFS Root shares provided by an upcoming Likewise DFS-N RPC service

- ❑ A multiple FSD driver model benefits developers by
 - ❑ Allowing a feature or functionality to be isolated and developed independently
 - ❑ “I’d like to write a database FSD”
 - ❑ Providing a “pivot point” against developers can verify a new driver against a known working configuration without recompiling or restarting
 - ❑ “Things work fine with PVFS but fail with GOOFS.”
- ❑ Drivers can be loaded only when required and therefore provide customized server builds
 - ❑ Not limited to FSDs

Building Likewise-CIFS

- ❑ Simple build system for Linux & FreeBSD
- ❑ Step 1: Download the source code
 - ❑ `$ git clone git://git.likewiseopen.org/likewise-open`
- ❑ Step 2: Build the likewise-open components
 - ❑ `$ build/mkcomp [--noincremental] [--debug] all`
 - ❑ Installs all pieces to “staging/install-root/”
- ❑ Step 3: Generate RPMs/DEBs (Linux only)
 - ❑ `$ build/mkpkg [--debug] all`
 - ❑ Creates packages in “staging/packages/”

Questions?

Gerald Carter

Director of Engineering

Likewise Software

gcarter@likewise.com

<http://www.likewise.com/>

<git://git.likewiseopen.org/likewise-open>