



LIKEWISE CIFS – A NEW GPL SMB/CIFS IMPLEMENTATION

LINUX PLUMBERS CONFERENCE

SEPTEMBER 25TH 2009



Likewise Open – Who are we?

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

- Goal – Likewise Open is the umbrella project sponsored by Likewise Software designed to provide a grounds-up Windows compatible distributed systems platform for Linux
 - Project officially launched Nov. '07
 - LWIS development began in Jan. '08
 - L-CIFS development began in Jan. '09
- License – Combination of GPLv2+ and LGPLv2.1+
 - Non-Likewise components (e.g. OpenLDAP and MIT Kerberos) remain under their original license.

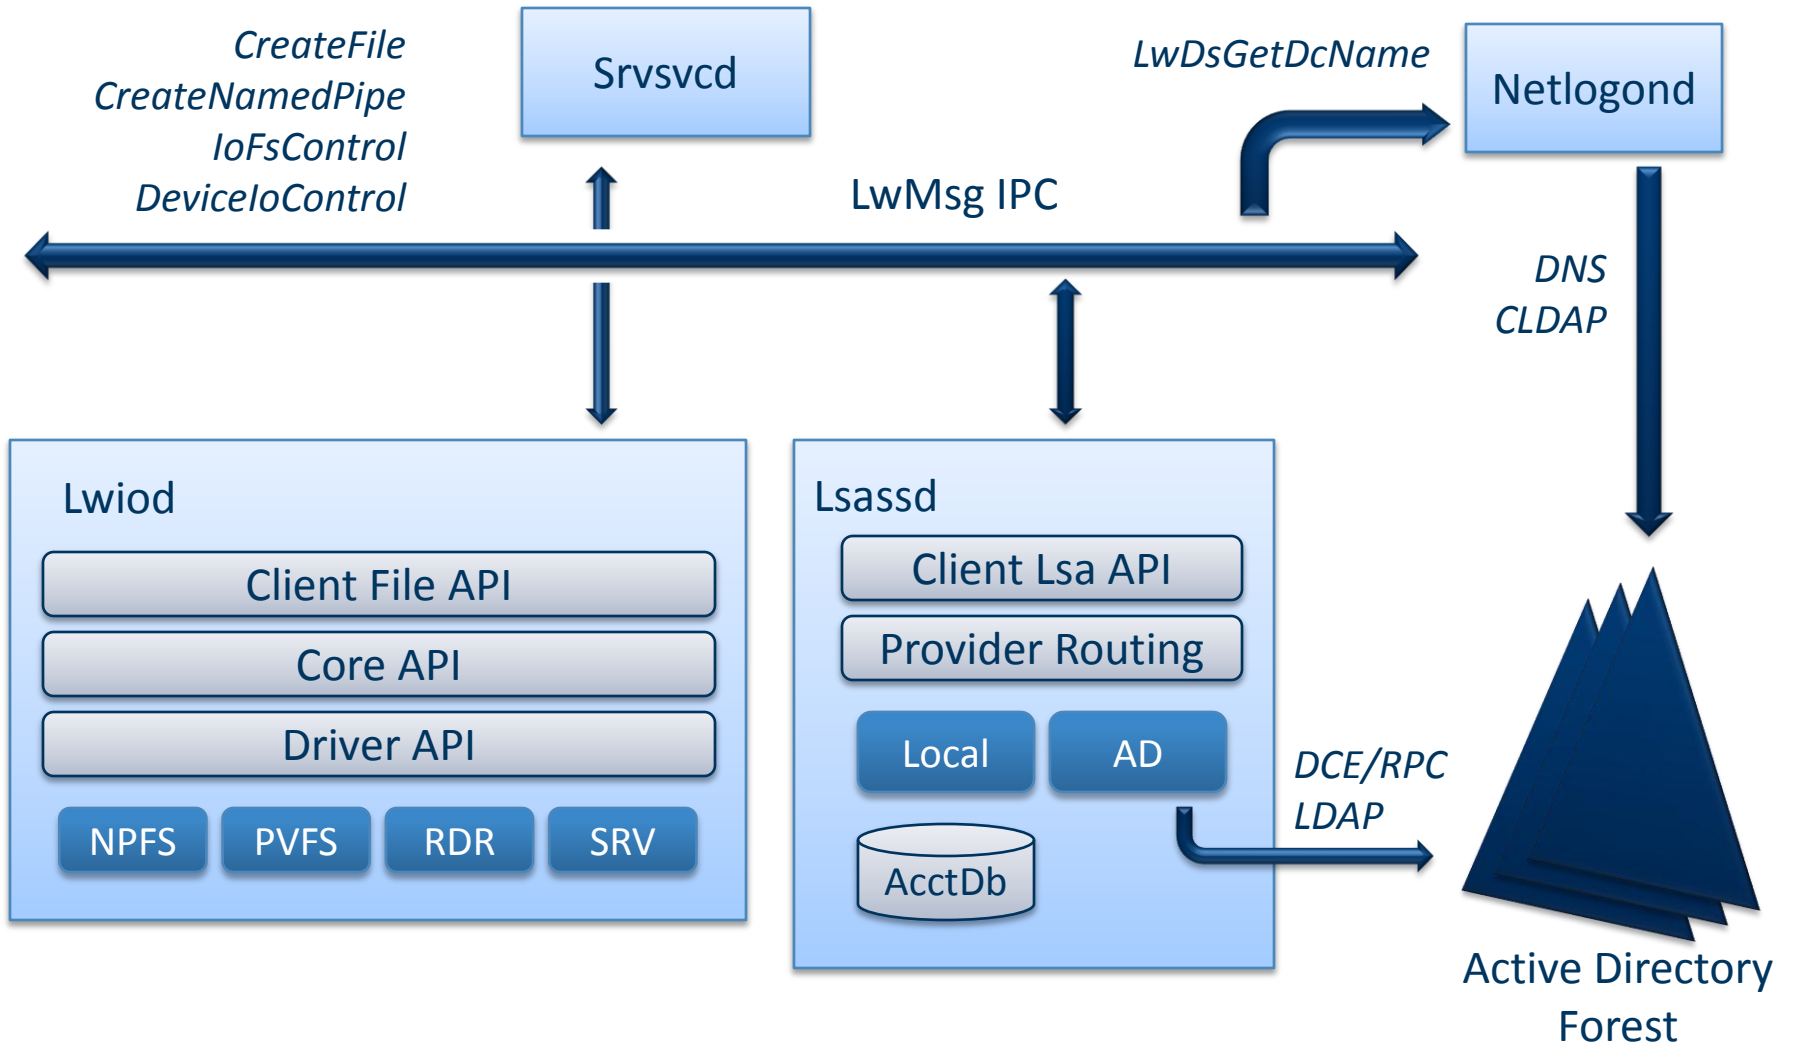
What does this mean?

- A new architecture for SMB/CIFS
 - User-mode I/O Framework and associated user-mode drivers
 - SMB redirector
 - SMB server
 - Named pipe server
 - Posix file system driver
- A new authentication daemon with plugins
 - local users,
 - Active Directory
 - LDAP (NIS
- A rehabilitated OSF DCE/RPC stack

Likewise Open Components

- All are single process, threaded services
- *lwiod* – Likewise I/O Manager
- *lsassd* – Likewise Security Authority
- *svsvcd* – Server and Workstation RPC Services
- netlogond – Domain Control locator
- dcerpcd – DCE/RPC endpoint-mapper
- eventlogd – Local/Remote logging service

Architectural Overview



Likewise I/O Manager

Likewise I/O Manager

- Provides an API inspired by the Windows ZwCreateFile(), et. al. interface
- Makes use of I/O request packets (IRPs) to communicate with drivers
- Drivers are loaded at run time by the I/O Mgr core
 - rdr.sys.so – SMB client file system
 - npfs.sys.so – Named pipe file system
 - pvfs.sys.so – POSIX compatible file system
 - srv.sys.so – SMBv1 & v2 server protocol head

I/O Mgr API

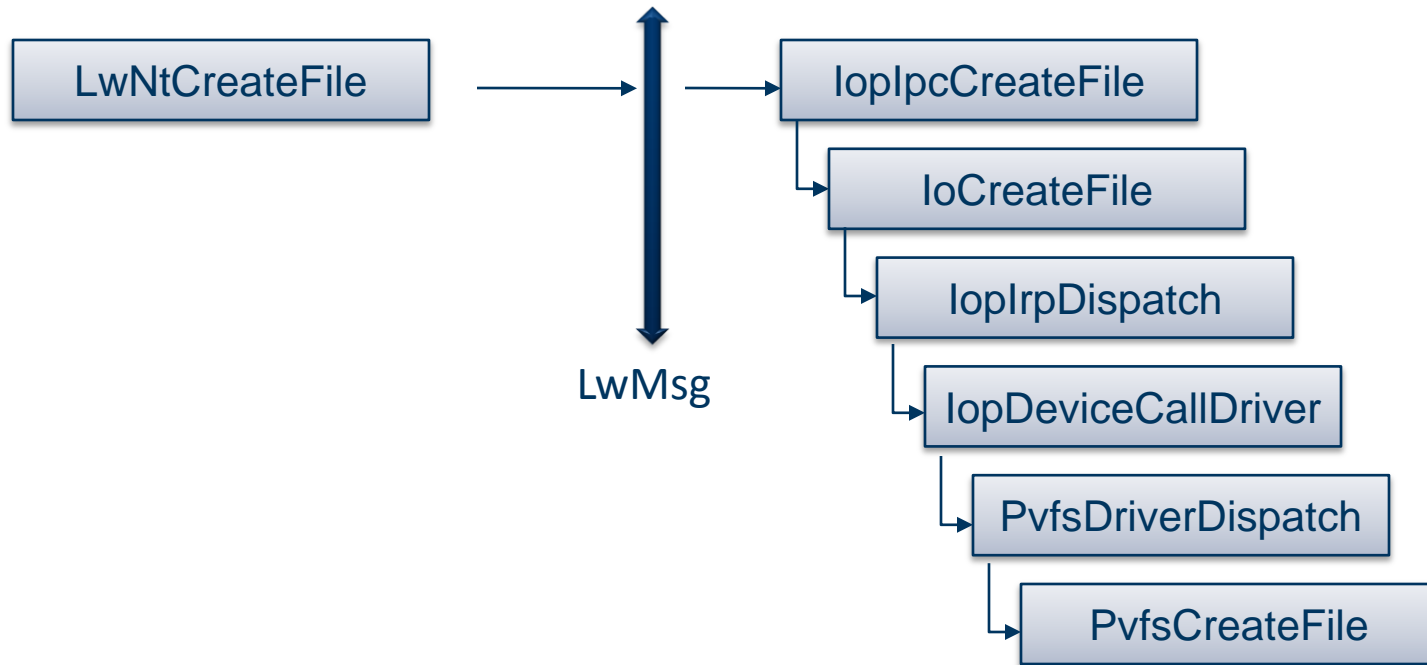
- IoCreateFile, IoCloseFile
- IoReadFile, IoWriteFile
- IoDeviceIoControlFile, IoFsControlFile
- IoQueryXXXInformation, IoSetXXXInformation
 - File, Directory, Volume
- IoLockFile, IoUnlockFile
- IoQuerySecurityFile, IoSetSecurityFile

Driver Namespace

- All drivers register a supported namespace
 - For example, “\pvfs” and “\npfs”
- The LwNtCreateFile() Client API call must include the driver namespace prefix in the filename.
 - Prefix is stripped by the I/O Mgr before sending the IRP to the correct driver
- A Win32 compatibility layer can be provided to insulate end-user applications
 - E.g. CreateFile(“\\server\share\file.txt”)

LwNtCreateFile Example

```
$> test_pvfs --cat /pvfs/etc/hosts
## /etc/hosts
127.0.0.1 localhost
127.0.1.1 sequoia.ad.plainjoe.org sequoia
```

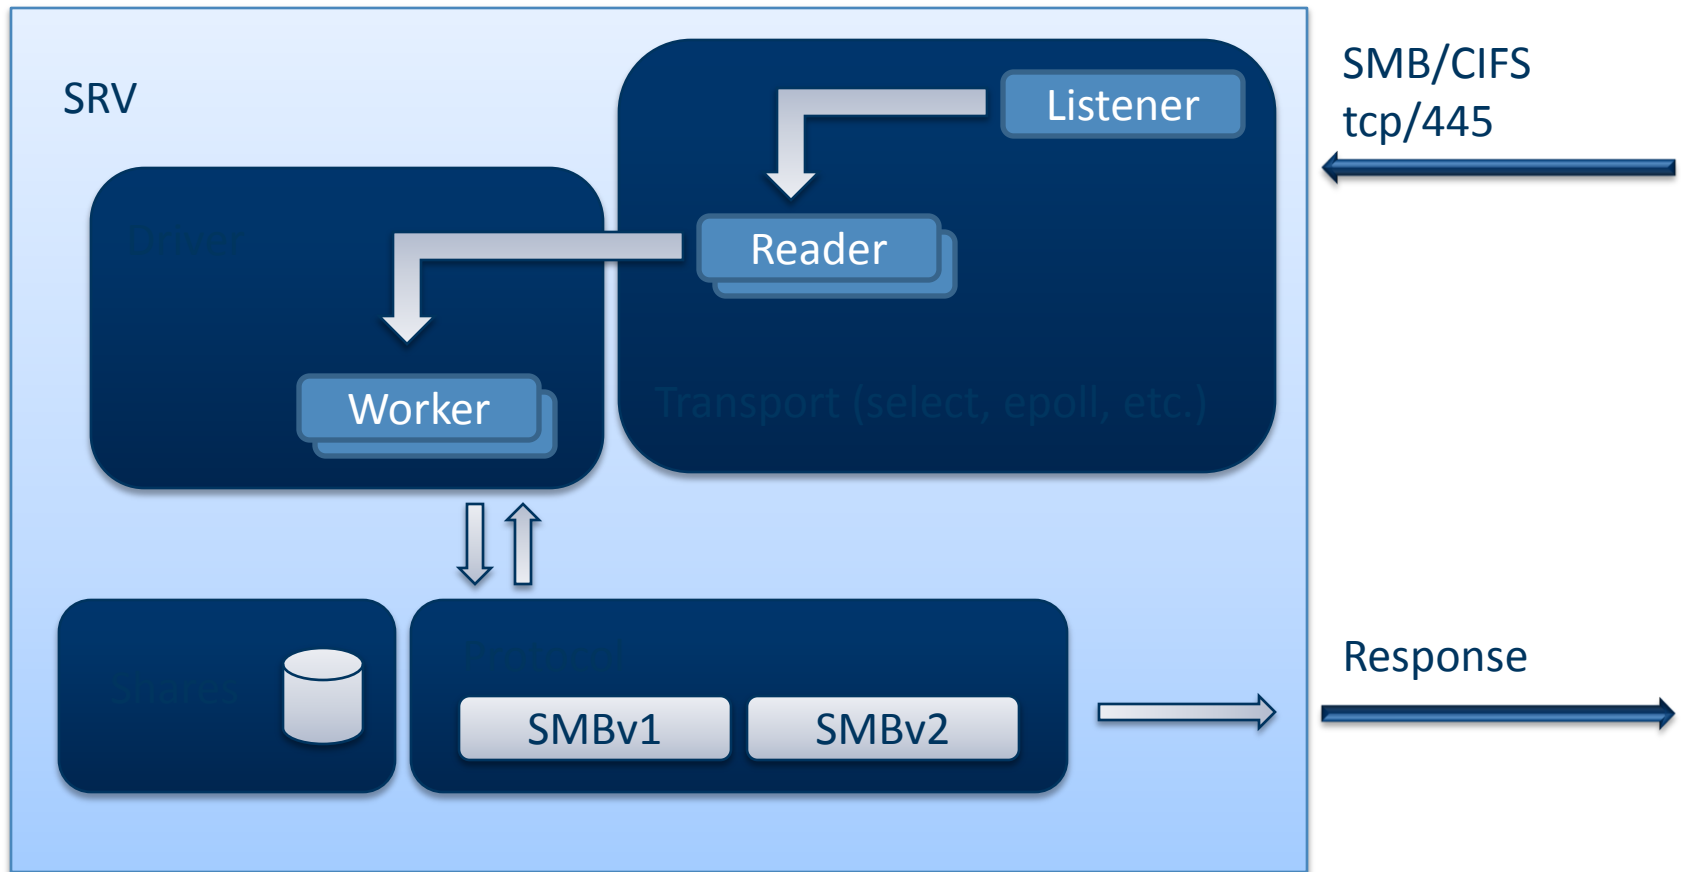


SRV & NPFS Drivers

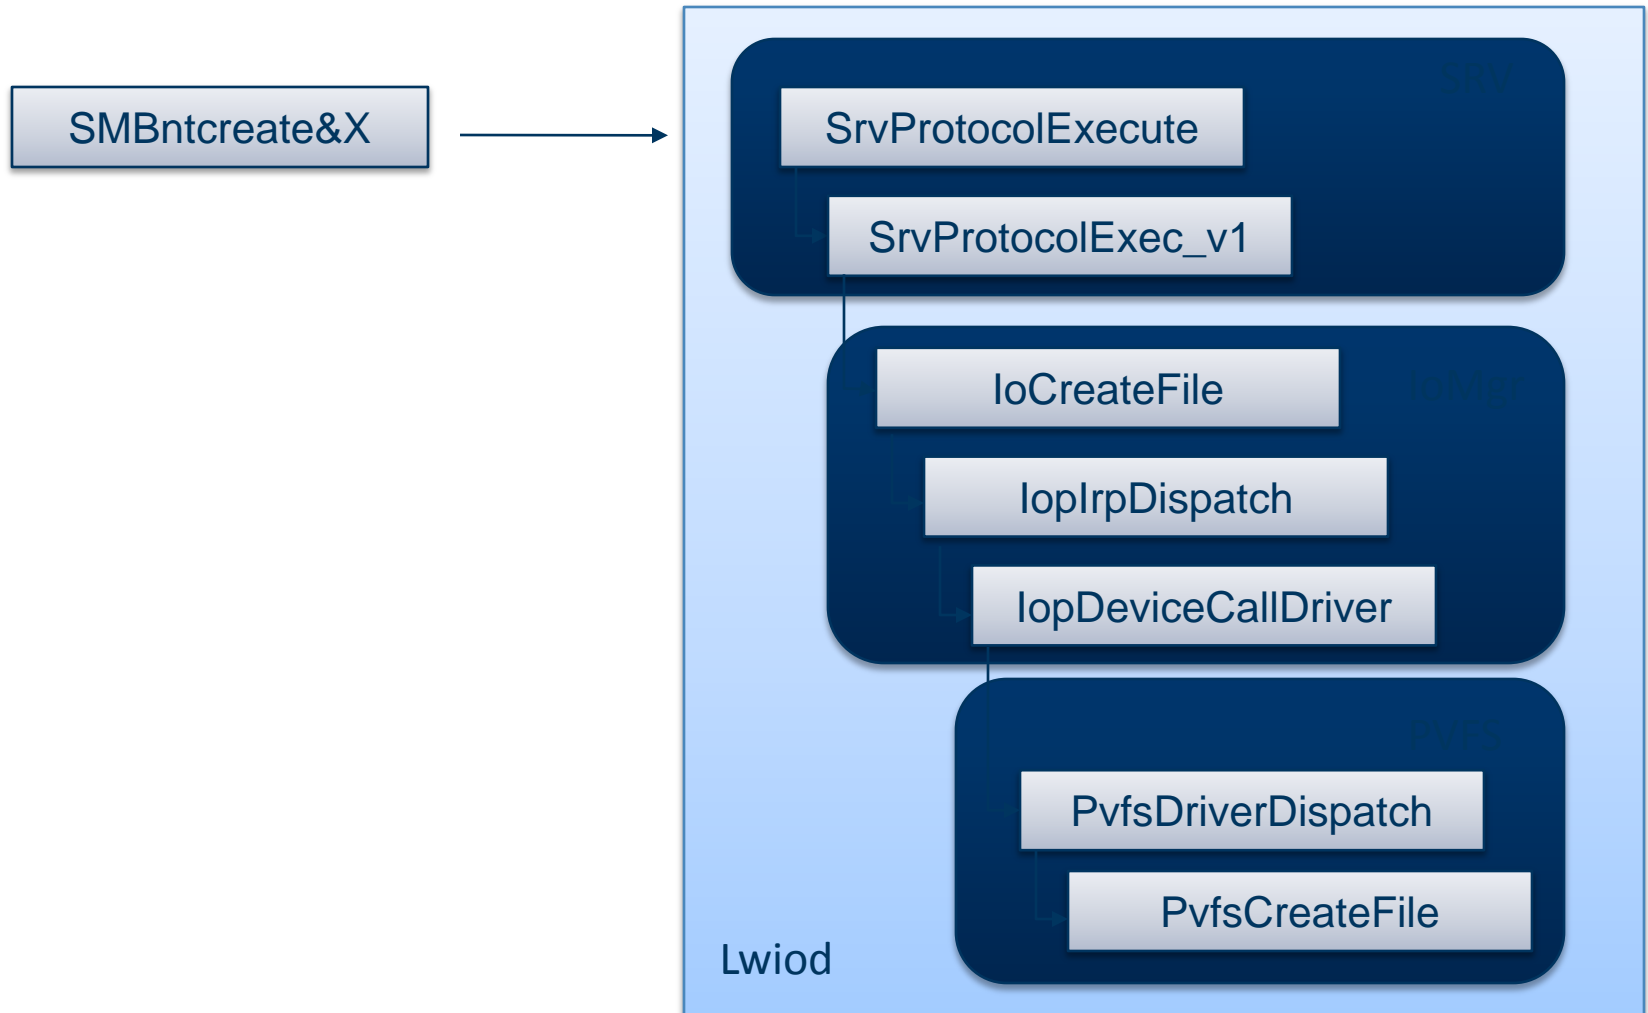
SRV.sys.so – SMB Protocol Head

- Support for SMBv1 and SMBv2
 - No NetBIOS support (only tcp/445)
 - NTLM 0.12 dialect or later
- Supported Clients
 - Windows XP/2003 and later
 - OS X and Linux clients forthcoming
- User mode security
 - Domain member and local authentication

SRV.sys.so - Architecture



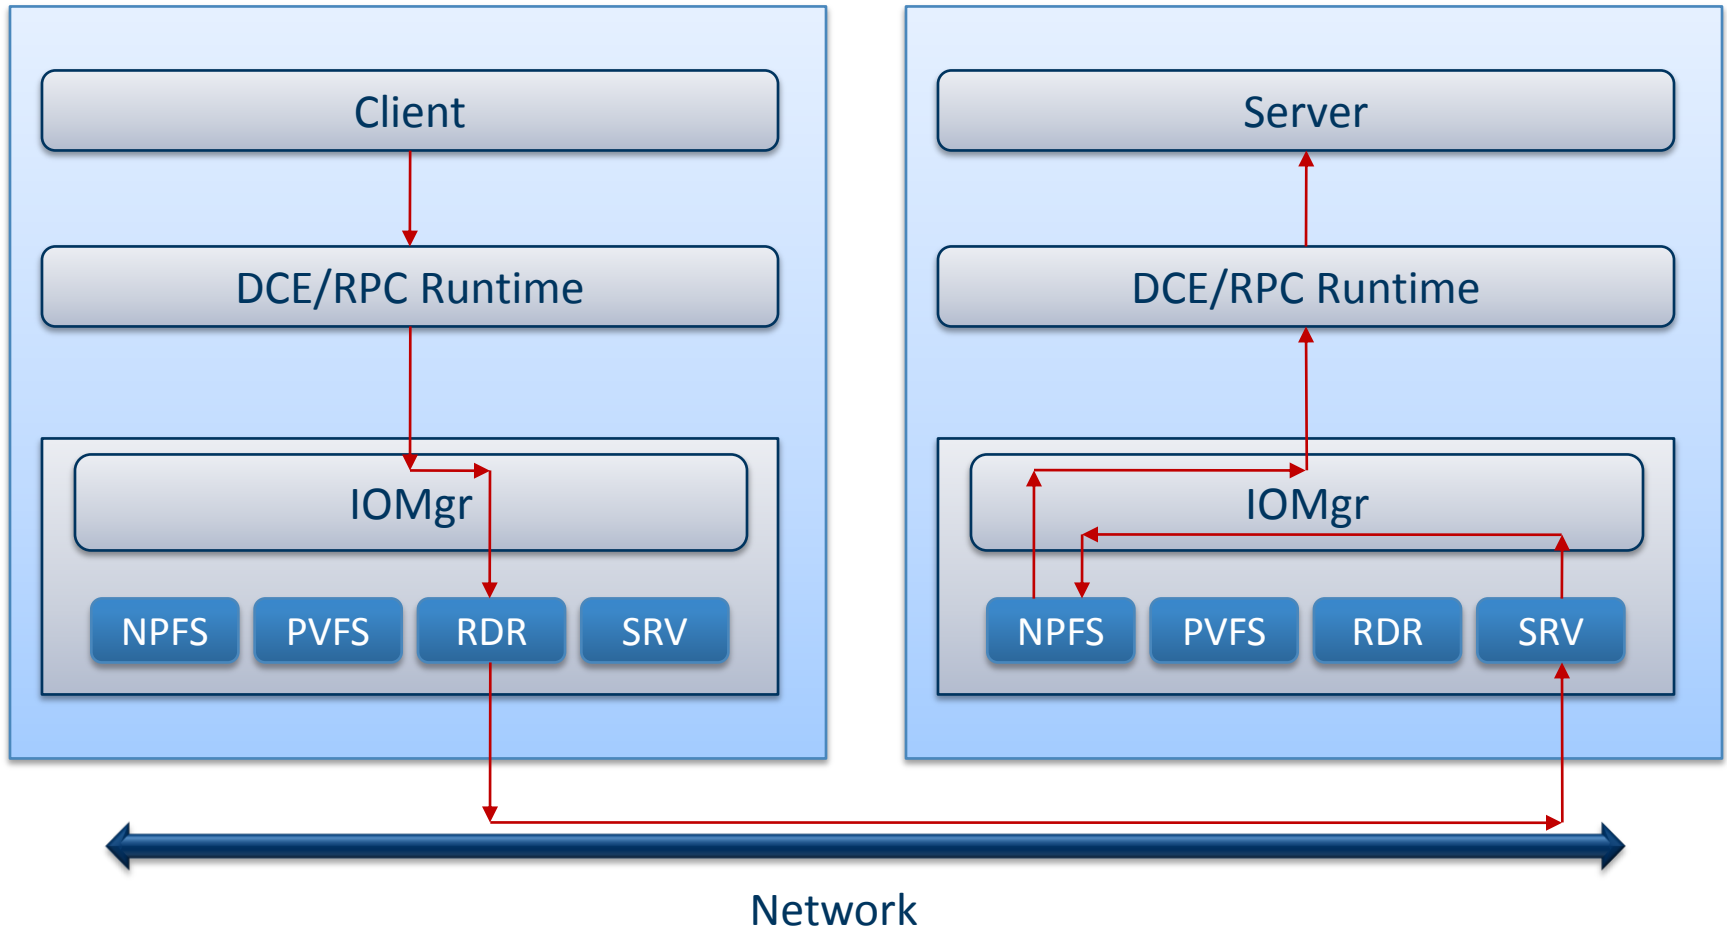
SMBntcreate&X Example



DCE/RPC & Named Pipes

- NPFS driver implements an in-memory named pipe file system
- DCE/RPC runtime supports clients and servers using the NPFS driver in `lwiod`
 - Registers an `ncacn_np` endpoint for server applications using `LwNtCreateNamedPipeFile()`
 - The client runtime calls `LwNtCreateFile()` to open a pipe on a remote host

DCE/RPC Clients & Servers

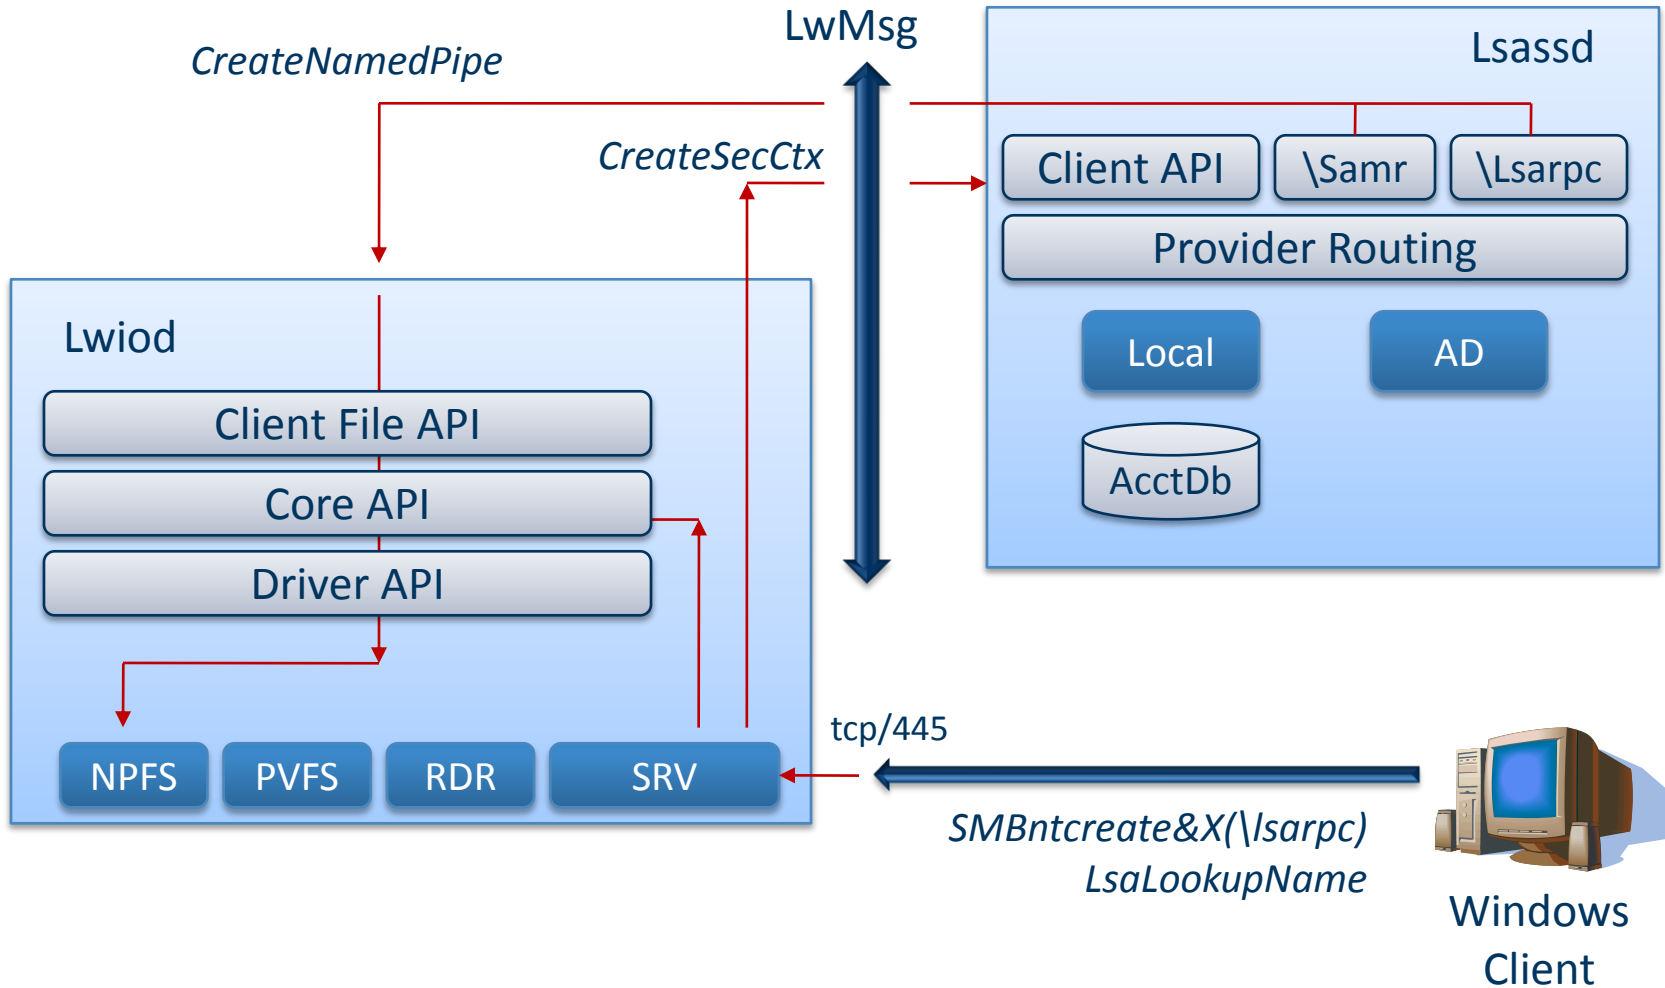


Likewise Security Authority

Likewise Security Authority

- User & Group Provider Routing
 - Local – Standalone account database
 - Privileged user management
 - Group nesting
 - MACHINE and BUILTIN domains
 - Active Directory – Member server functionality
 - Trust scenarios, Authentication, etc...
- Supplies session security contexts for Lwioid

Users Tokens and RPC Servers

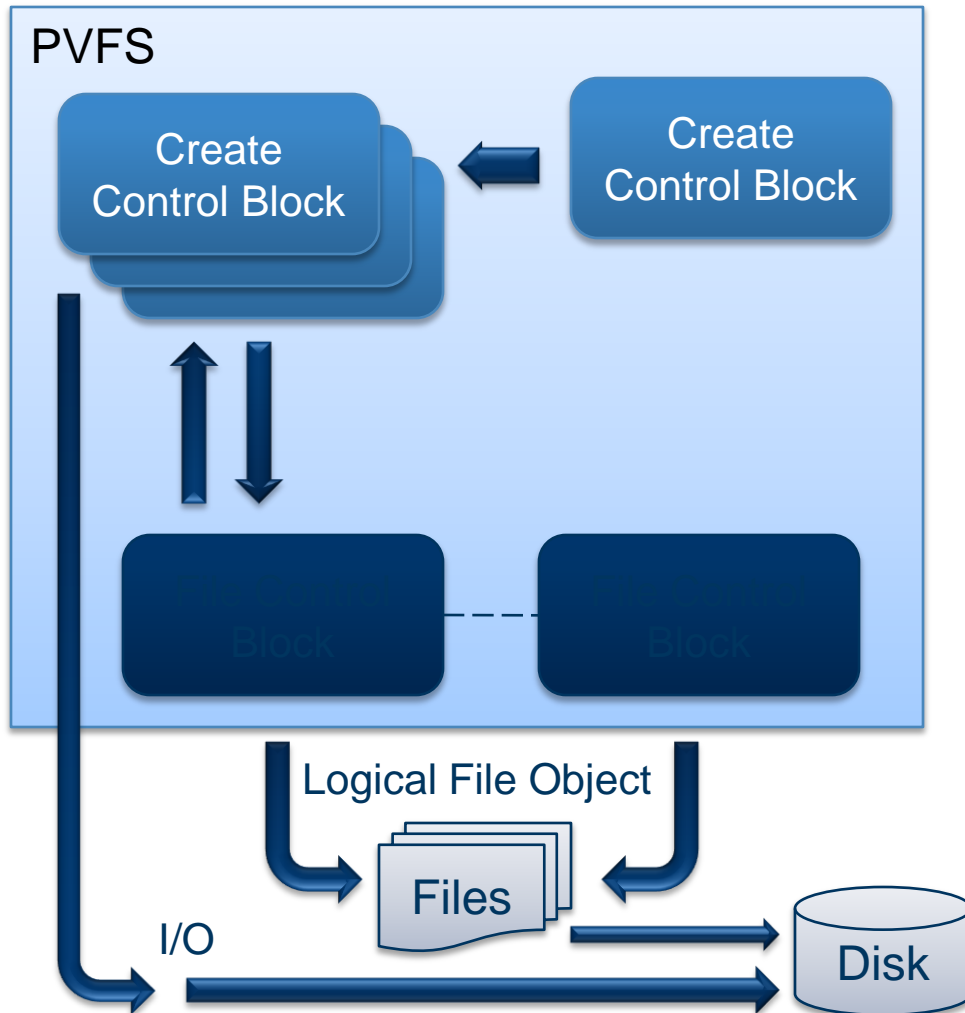


PVFS Driver

PVFS.sys.so

- Integration with POSIX file systems
 - Uses EAs for storing security descriptors, Attributes, etc..
 - Implements security and locking checks in process
- Provides a worker thread pool

PVFS – Data Structures



- FCB – File Object
 - Opcodes
- CCB – Open Handle
 - Pathname
 - Dev/Inode
 - BRL
 - Sharemode
 - File Descriptor

PVFS – Data Structures (cont)

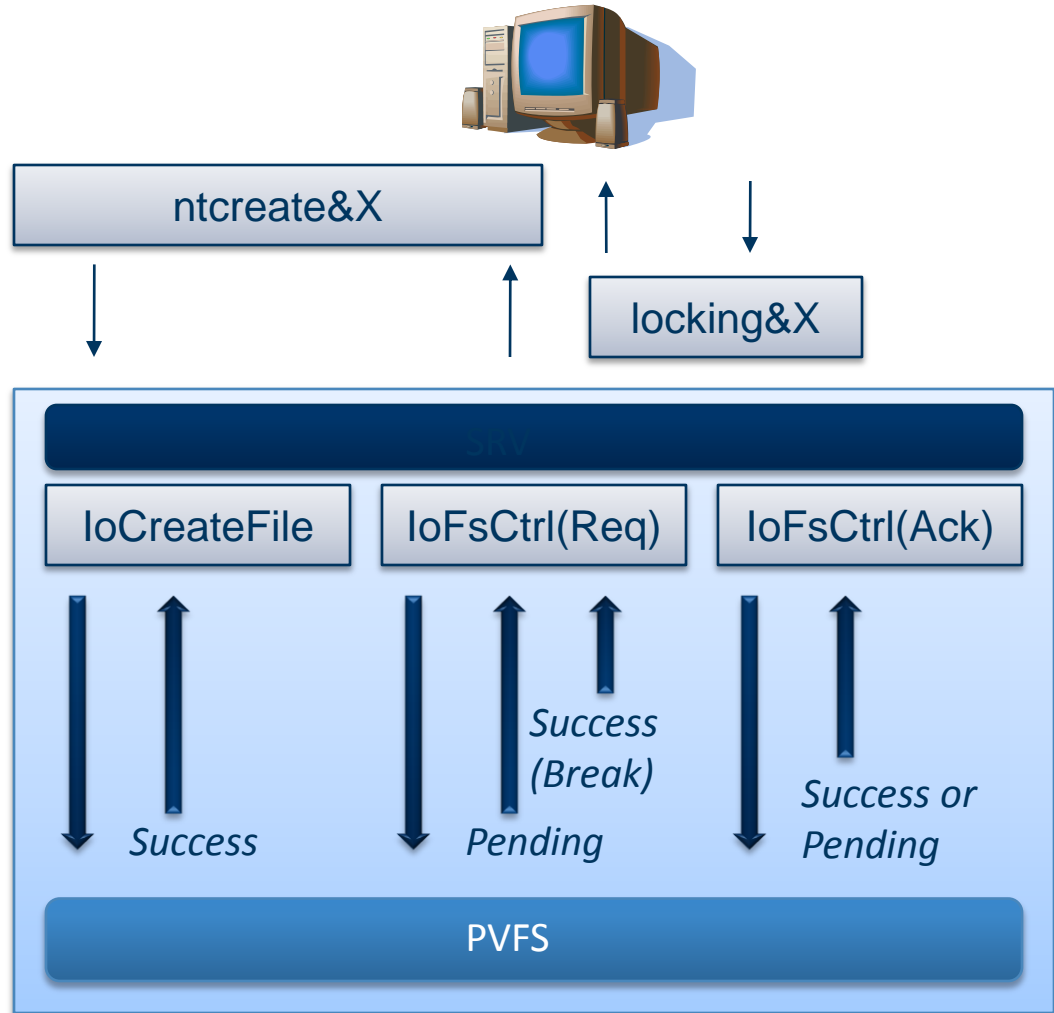
- File Control Block represents the file on disk
 - FCB is removed when last open handle is closed
- Create Control Block is open file handle
 - Stored in the `IO_FILE_HANDLE`
 - Lwlo API is handle based (i.e. All files and directories are processed first through `CreateFile`)
- CCB refers to its FCB; FCB owns a list of its CCBs

Share Modes and Byte Range Locks


- Share modes and byte range locking information is stored with the open handle in the CCB
- A share mode or BRL check checks all associated CCBs until a conflict is detected or success
 - PvfsEnforceShareMode(), PvfsCanLock(), PvfsAddLock()
- Pending locks are stored on the FCB
 - Backlink to the requesting CCB
 - Processed on any change to the lock table

Oplocks

- Legacy oplocks
- Requested using FsloCtrl on CCB
- Oplock list stored on the FCB
- Deferred ops stored in a queue on the FCB



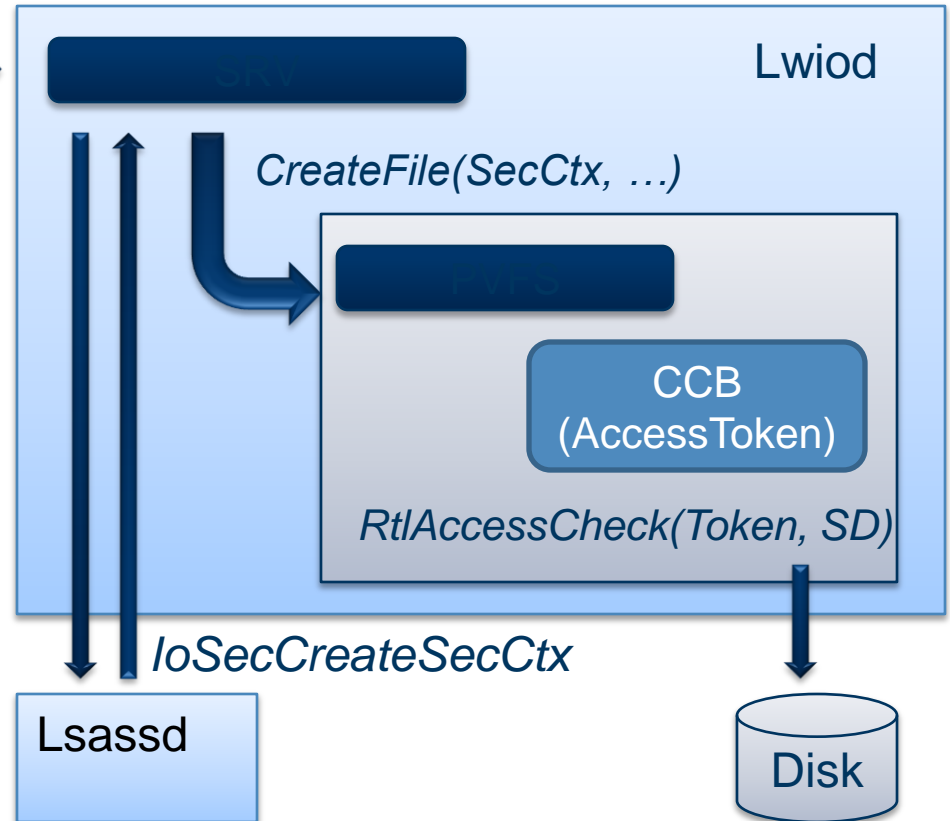
CREATE_SECURITY_CONTEXT



SessionSetup&X



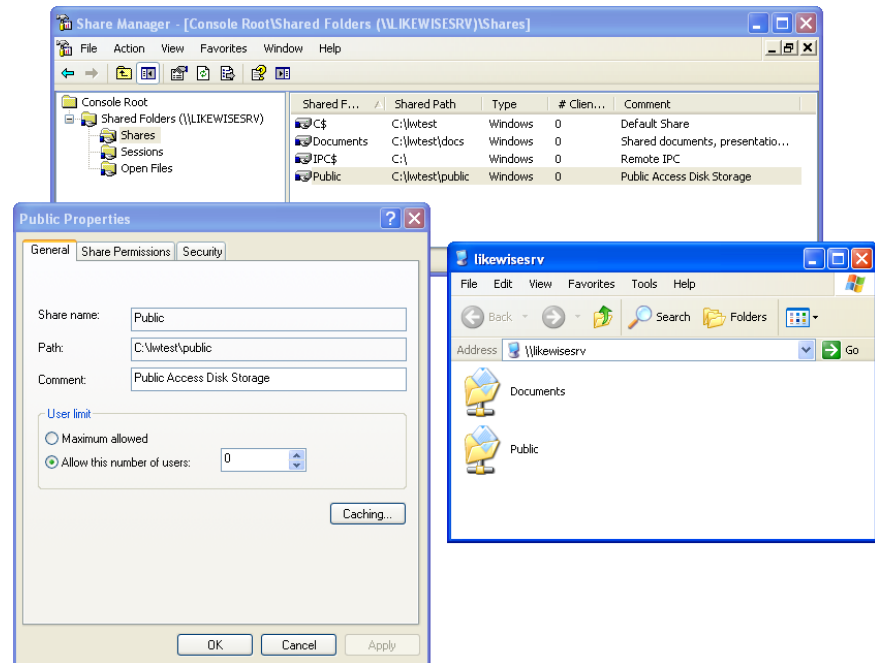
- Obtained from Lsassd during SessionSetup processing
- Passed to IoCreateFile()
- Contains user's Access Token



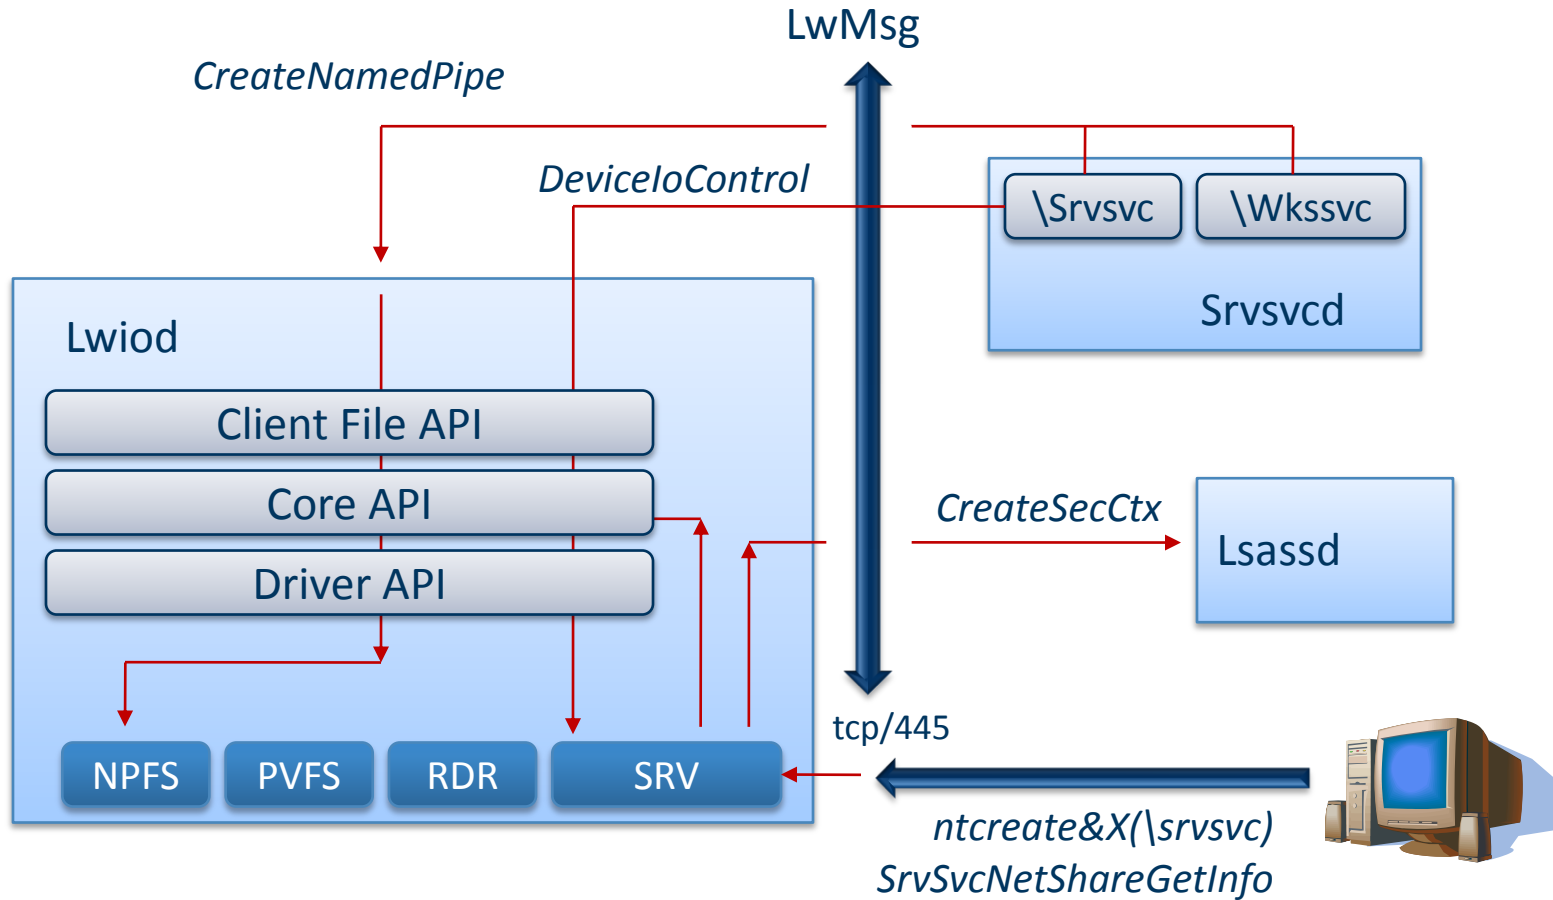
Server & Workstation Service

Server Service (srvsvcd)

- Implements the Srvsvc & Wkssvc RPC interfaces
- Retrieves information about file shares from `Lwiod/SRV LwNtDeviceIoControlFile()`



Server Service (cont)



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?

krishnag/gcarter@likewise.com

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

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