CIFS, ACLs and ZFS:
The One File System to Rule them all!

Aka: "Step-by-Step guide to get a CIFS server working for Windows clients"

Andrew Watkins
andrew@dcs.bbk.ac.uk
http://notallmicrosoft.blogspot.com
Birkbeck College
Department of Computer Science
and Information Systems

LOSUG
January 19th, 2011

Birkbeck
UNIVERSITY OF LONDON

Solaris CIFS Server: Background

- **"Seamless, ubiquitous, cross-protocol file sharing"**
  Alan Wright, Project Lead for CIFS Server

- **CIFS server is now a first class citizen in Solaris**
  - **Putback into Development/Nevada October 2007**
  - Available in Solaris Express and OpenSolaris 2008.03
  - **25+ ARC cases, 800 files, approximately 370,000 lines of code (including 180,000 lines of new code)**

- **Tight integration with NFS, ZFS, and Active Directory**
  - **Windows/CIFS concepts such as Security Identifiers and Access Tokens are now native to Solaris kernel**

Jarod Nash - LOSUG - September 2008

Solaris CIFS Service

"Seamless, ubiquitous, cross-protocol file sharing"

http://www.oug.org/files/presentations/cifs-losug.pdf

What you need?

- Solaris 11 Express or OpenIndiana
- Windows Domain or Windows Workgroup

Window Active Directory Domain

- ~£200 - £500 gets you a Supported version (Business/Standard)
  - with BUG FIXES and Patches!!
- Free for students with Bug Fixes and Patches!!
- CIFS is compatible with Windows 2003 and Windows 2008
- Windows 2008 needs some patches to work with CIFS:
  - NTLMv2 authentication problem: http://support.microsoft.com/kb/957441
  - Windows Server 2008 SP1 with hot fix KB951191
- Windows Server 2008 SP2 may have these fixes!

For Example

- Set AD Domain = test.int
- Windows 2008 server = windows = 192.168.56.3
- Solaris Server = openindiana = 192.168.56.5

- Setup DNS Server
- Check packages are installed

```
% pkg list smb
  NAME (PUBLISHER)          VERSION          STATE      UFOXI
  service/file-system/smb   0.5.11-0.148     installed  -----
  system/file-system/smb    0.5.11-0.148     installed  -----
```

- Sync clocks

CIFS Server

- Identity mapping of users and groups between systems

```
$ svcs \*idmap\*
STATE          STIME    FMRI
disabled       12:16:59 svc:/system/idmap:default
$ svcadm enable idmap
$ svcs \*idmap\*
STATE          STIME    FMRI
online         12:40:38 svc:/system/idmap:default

$ pfexec idmap add 'winuser:*@test.int' 'unixuser:*'
$ pfexec idmap add 'wingroup:*@test.int' 'unixgroup:*'
$ idmap list
add     winuser:*@test.int    unixuser:*
add     wingroup:*@test.int   unixgroup:*
```

idmap

- $ idmap dump

```
usid:S-1-5-21-275504925-2437894988-2844437058-500 ==  uid:2147483649
gsid:S-1-5-21-275504925-2437894988-2844437058-513 ==  gid:2147483650
gsid:S-1-5-21-275504925-2437894988-2844437058-512 ==  gid:501
gsid:S-1-5-21-275504925-2437894988-2844437058-519 ==  gid:2147483653
gsid:S-1-5-11          ==  gid:2147483656
gsid:S-1-5-32-544      ==  gid:2147483657
$ idmap dump –n
winuser:Administrator@test.int          ==  uid:2147483649
wingroup:Domain Users@test.int          ==  gid:2147483650
wingroup:Domain Admins@test.int         ==  unixgroup:winadmin
wingroup:Enterprise Admins@test.int     ==  gid:2147483653
wingroup:Authenticated Users            ==  gid:2147483656
wingroup:Administrators@BUILTIN         ==  gid:2147483657
```

Identity Mapping

- Unknown Windows identities are mapped to dynamically allocate UIDs / GIDs. windows SID => unix UID
- Unknown Unix identities are not mapped to Windows so they MUST exist in AD. ~~unix UID => windows SID~~

```
idmap[501]: [ID 523480 daemon.notice] AD lookup of winname
   root@test.int failed, error code -9961
idmap[501]: [ID 523480 daemon.notice] AD lookup of winname
   sys@test.int failed, error code -9961
idmap[501]: [ID 523480 daemon.notice] AD lookup of winname
   staff@test.int failed, error code -9961
```

- It is a good idea that well know accounts which may be used in ACL have a permanent mapping to a UNIX group

```
$ idmap add "wingroup:Domain Admins@test.int" unixgroup:winadmin
```

CIFS Server (Active Directory Domain)

- **Edit /etc/krb5/krb5.conf**

```
[libdefaults]
    default_realm = TEST.INT
[realms]
    TEST.INT = {
            kdc = windows.test.int
            admin_server = windows.test.int
            kpasswd_server = windows.test.int
            kpasswd_protocol = SET_CHANGE
    }
 [domain_realm]
     .test.int = TEST.INT
```

- **Start smb services**

```
$ pfexec svcadm enable -r smb/server
$ svcs \*smb\
STATE          STIME    FMRI
disabled       Nov_08   svc:/network/smb/client:default
online         15:21:01 svc:/network/smb/server:default
online         15:21:03 svc:/network/shares/group:smb
```

Join the Active Directory Domain

```
$ pfexec smbadm join -u Administrator test.int
After joining test.int the smb service will be restarted.
Would you like to continue? [no]: yes
Enter domain password:
Joining test.int ... this may take a minute ...
failed to find any domain controllers for test.int

$ tail /var/adm/messages
...openindiana …. smbd: failed locating domain controller
...openindiana …  smbd_dc_update: test.int: located windows
```

- **Set the LAN manager authentication level on your Solaris system smb(4)**

```
$ pfexec sharectl set -p lmauth_level=2 smb
$ pfexec smbadm join -u Administrator test.int
After joining test.int the smb service will be restarted
Would you like to continue? [no]: yes
Enter domain password:
Joining test.int ... this may take a minute ...
Successfully joined test.int
```

Join the WorkGroup (this what the manual says!)

```
$ pfexec smbadm join –w WorkGroup-Name
```

Need to setup Solaris server to handle authentication of users.
Edit /etc/pam.conf to support an encrypted SMB password

```
other    password required     pam_smb_passwd.so.1   nowarn

$ pfexec passwd andrew
```

Setup ZFS filesystem

- Enable Cross-Protocol Locking (nbmand)
  - SMB assumes mandatory locking
  - UNIX advisory locking
- Mixed case (casesensitivity)
- Enable SMB sharing on share (sharesmb)

```
$ pfexec zfs create -o nbmand=on -o casesensitivity=mixed
   rpool/export/homes
$ pfexec zfs create rpool/export/homes/andrew
$ pfexec zfs set sharesmb=name=andrew rpool/export/homes/andrew

$ zfs get nbmand,casesensitivity,sharesmb
   rpool/export/homes/andrew
NAME                PROPERTY         VALUE        SOURCE
../homes/andrew     nbmand           on           inherited from ..
../homes/andrew     casesensitivity  mixed        -
../homes/andrew     sharesmb         name=andrew  local

$ pfexec chown andrew:staff /export/homes/andrew
```

Check $path

```
# touch /export/homes/andrew/file
# echo $path
/usr/gnu/bin:/usr/bin:/usr/sbin:/sbin

$ ls –lv
-rw-r--r--   1 andrew   staff        0 Nov 18 18:42 file
$ /usr/bin/ls -lv
-rw-r--r--   1 andrew   staff        0 Nov 18 18:42 file
      0:owner@:read_data/write_data/append_data/read_xattr/write_xattr
        /read_attributes/write_attributes/read_acl/write_acl/write_owner
        /synchronize:allow
      1:group@:read_data/read_xattr/read_attributes/read_acl/synchronize:allow
      2:everyone@:read_data/read_xattr/read_attributes/read_acl/synchronize
        :allow

$ /usr/bin/ls -lV
-rw-r--r--   1 andrew   staff        0 Nov 18 18:42 file
             owner@:rw-p--aARWcCos:-------:allow
             group@:r-----a-R-c--s:-------:allow
          everyone@:r-----a-R-c--s:-------:allow
```

ACL properties on filesystem

discard - New objects, no ACL entries are inherited

noallow - New objects, only inheritable ACL entries that have access to type deny are inherited.

restricted - New objects, the write_owner and write_acl permissions are removed when ACL entry is inherited.

passthrough - New objects are created with a mode determined by the inheritable ACEs (Access Control Entries).

passthrough-x - As above, plus files are created with the execute (x) set.

So to get inheritance working from Windows:

```
$ pfexec zfs set  aclinherit=passthrough-x  rpool/export/homes
$ zfs get aclinherit  rpool/export/homes/andrew
NAME              PROPERTY    VALUE          SOURCE
…homes/andrew  aclinherit  passthrough-x  inherited from …/homes
```

ACLs

```
$ /bin/ls -ldv /export/homes/andrew
drwxr-xr-x   3 andrew   staff          3 Nov 30 12:40 /export/homes/andrew
     0:owner@:list_directory/read_data/add_file/write_data/add_subdirectory
         /append_data/read_xattr/write_xattr/execute/read_attributes
         /write_attributes/read_acl/write_acl/write_owner/synchronize:allow
     1:group@:list_directory/read_data/read_xattr/execute/read_attributes
         /read_acl/synchronize:allow
     2:everyone@:list_directory/read_data/read_xattr/execute/read_attributes
         /read_acl/synchronize:allow

$ /bin/ls -ldV /export/homes/andrew
drwxr-xr-x   3 andrew   staff          3 Nov 30 12:40 /export/homes/andrew
                owner@:rwxp--aARWcCos:-------:allow    (0)
                group@:r-x---a-R-c--s:-------:allow    (1)
             everyone@:r-x---a-R-c--s:-------:allow    (2)
```

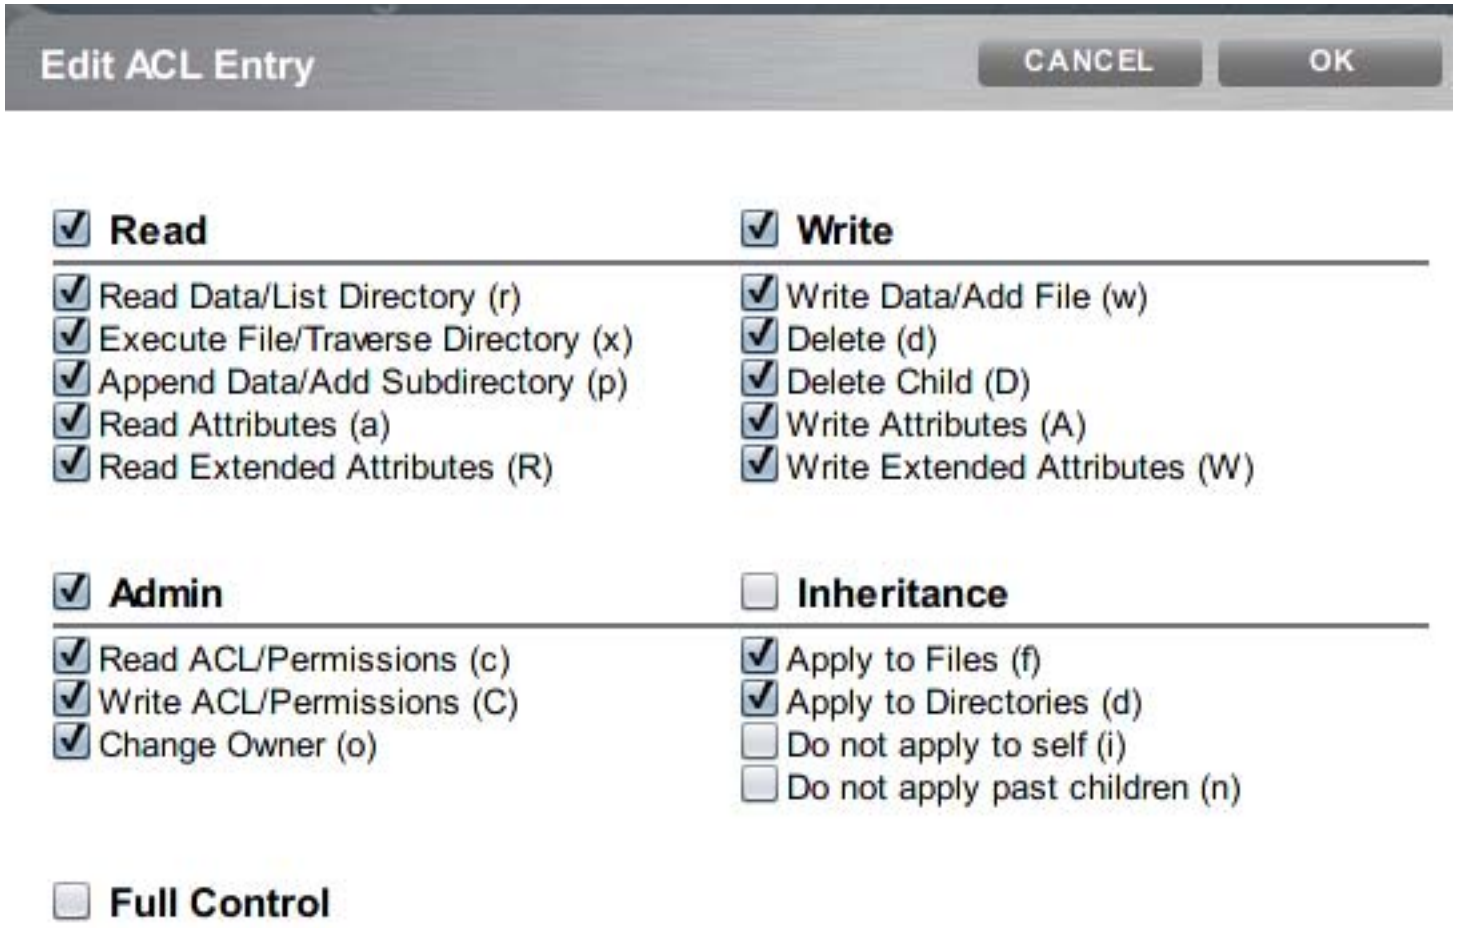


```
owner@:list_directory/read_data/add_file/write_data/add_subdirectory
/append_data/read_xattr/write_xattr/execute/delete_child/read_attributes
/write_attributes/delete/read_acl/write_acl/write_owner/synchronize:file_i
nherit/dir_inherit:allow

owner@:rwxpdDaARWcCos:fd-----:allow
```

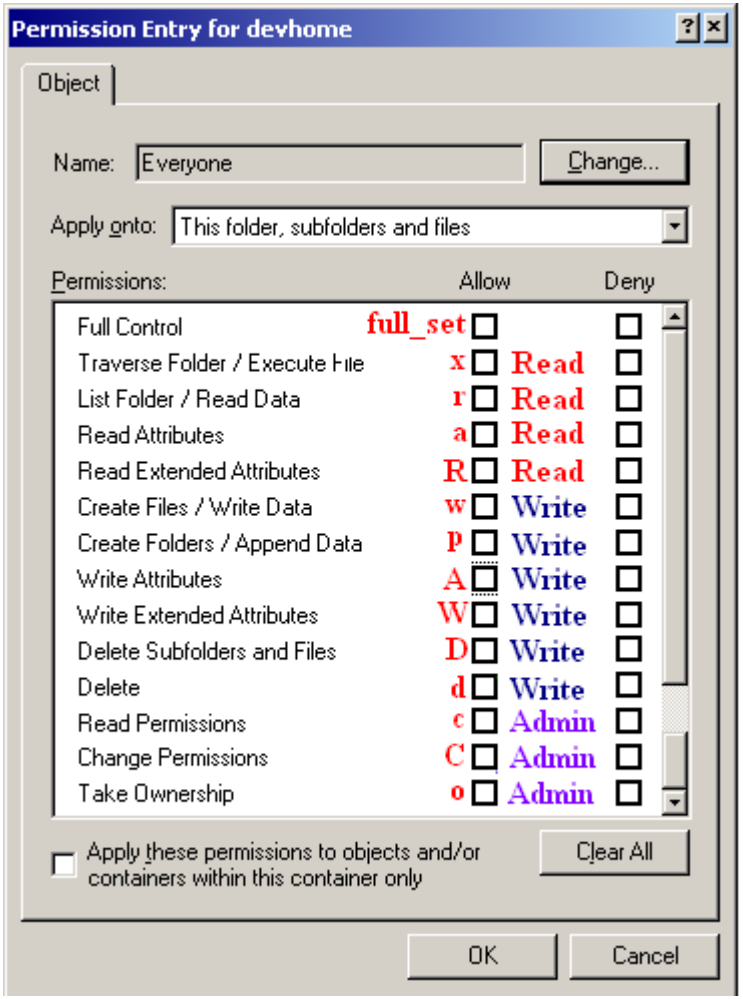


```
owner@:list_directory/read_data/add_file/write_data/add_subdirectory
/append_data/read_xattr/write_xattr/execute/delete_child/read_attributes
/write_attributes/delete/read_acl/write_acl/write_owner/synchronize:file_i
nherit/dir_inherit:allow

owner@:rwxpdDaARWcCos:fd-----:allow
```

ACL Entry Type

owner@	Specifies the access granted to the owner of the object.
group@	Specifies the access granted to the owning group of the object.
everyone@	Specifies the access granted to any user or group that does not match any other ACL entry.
user	With a user name, specifies the access granted to an additional user of the object. Must include the ACL-entry-ID, which contains a username or userID. If the value is not a valid numeric UID or username, the ACL entry type is invalid.
group	With a group name, specifies the access granted to an additional group of the object. Must include the ACL-entry-ID, which contains a groupname or groupID. If the value is not a valid numeric GID or groupname, the ACL entry type is invalid.

ZFS ACL Sets

ACL Set Name	Included ACL Permissions
full_set	All permissions :rwxpdDaARWcCos:-------:allow chmod "A+user:andrew:full_set:allow" file
modify_set	all permissions except write_acl and write_owner :rwxpdDaARWc--s:-------:allow chmod "A+user:andrew:modify_set:allow" file
read_set	read_data, read_attributes, read_xattr, and read_acl :r-----a-R-c---:-------:allow chmod "A+user:andrew:read_set:allow" file
write_set	write_data, append_data, write_attributes, and write_xattr :-w-p---A-W----:-------:allow chmod "A+user:andrew:write_set:allow" file

ACL Access Privileges

Access Privilege	Compact Access Privilege	
add_file	w	Permission to add a new file to a directory.
add_subdirectory	p	On a directory, permission to create a subdirectory.
append_data	p	Not currently implemented.
delete	d	Permission to delete a file.
delete_child	D	Permission to delete a file or directory within a directory.
execute	x	Permission to execute a file or search the contents of a directory.
list_directory	r	Permission to list the contents of a directory.
read_acl	c	Permission to read the ACL (ls).
read_attributes	a	Permission to read basic attributes (non-ACLs) of a file. Think of basic attributes as the stat level attributes. Allowing this access mask bit means the entity can execute ls(1) and stat(2).

ACL Access Privileges

Access Privilege	Compact Access Privilege	
read_data	r	Permission to read the contents of the file.
read_xattr	R	Permission to read the extended attributes of a file or perform a lookup in the file's extended attributes directory.
synchronize	s	Placeholder. Not currently implemented.
write_xattr	W	Permission to create extended attributes or write to the extended attributes directory. Granting this permission to a user means that the user can create an extended attribute directory for a file. The attribute file's permissions control the user's access to the attribute.
write_data	w	Permission to modify or replace the contents of a file.

ACL Access Privileges

Access Privilege	Compact Access Privilege	
write_attributes	A	Permission to change the times associated with a file or directory to an arbitrary value.
write_acl	C	Permission to write the ACL or the ability to modify the ACL by using the chmod command.
write_owner	o	Permission to change the file's owner or group. Or, the ability to execute the chown or chgrp commands on the file. Permission to take ownership of a file or permission to change the group ownership of the file to a group of which the user is a member. If you want to change the file or group ownership to an arbitrary user or group, then the PRIV_FILE_CHOWN privilege is required.

ACL Inheritance

Inheritance Flag	Inheritance Flag	
file_inherit	f	Only inherit the ACL from the parent directory to the directory's files
dir_inherit	d	Only inherit the ACL from the parent directory to the directory's subdirectories.
inherit_only	i	Inherit the ACL from the parent directory but applies only to newly created files or subdirectories and not the directory itself. This flag requires the file_inherit flag, the dir_inherit flag, or both, to indicate what to inherit.
no_propagate	n	Only inherit the ACL from the parent directory to the first-level contents of the directory, not the second-level or subsequent contents. This flag requires the file_inherit flag, the dir_inherit flag, or both, to indicate what to inherit.

ACL Inheritance

Inheritance Flag	Inheritance Flag	
Currently, the following flags are only applicable to a SMB client or server		
successful_access	S	Indicates whether an alarm or audit record should be initiated upon a successful access. This flag is used with audit or alarm ACE types.
failed_access	F	Indicates whether an alarm or audit record should be initiated when an access fails. This flag is used with audit or alarm ACE types.
inherited	I	Indicates that an ACE was inherited.

Birkbeck
UNIVERSITY OF LONDON

What's New since OpenSolaris

PSARC/2010/029 FastTrack

- deny ACL are not required in most cases now.
  - exceptions: 0705 (g-rwx), 0060 (u-rwx)
- aclmode has gone, which means that chmod will discard all ACLs
  - does not try to keep ACLs in place any more
- user and owner are treated together?
  - You no longer require both andrew: and owner:

An Interoperability Solution

```
$ /bin/ls -ldV /export/homes/andrew
drwxr-xr-x   3 andrew   staff          3 Nov 30 12:40 /export/homes/andrew
                  owner@:rwxp--aARWcCos:-------:allow     (0)
                  group@:r-x---a-R-c--s:-------:allow     (1)
                everyone@:r-x---a-R-c--s:-------:allow     (2)

$ cd /export/homes

$ chmod "A2=everyone@:r-x---a-R-c--s:fd-----:allow” andrew

$ chmod "A+user:andrew:rwxpd-aARWc--s:fd-----:allow" andrew

$ chmod "A+group:staff:r-x---a-R-c--s:fd:allow"  andrew

$ chmod "A+group:winadmin:full_set:file_inherit/dir_inherit:allow"
   andrew
```

An Interoperability Solution

```
drwxr-xr-x+  3 andrew   staff          3 Nov 30 12:40 andrew
         group:winadmin:rwxpdDaARWcCos:fd-----:allow          (0)
            group:staff:r-x---a-R-c--s:fd-----:allow          (1)
            user:andrew:rwxpd-aARWc--s:fd-----:allow          (2)
                owner@:rwxp--aARWcCos:-------:allow           (3)
                group@:r-x---a-R-c--s:-------:allow           (4)
              everyone@:r-x---a-R-c--s:fd-----:allow          (5)

$ chmod  "A4-" andrew
$ chmod  "A-owner@:rwxp--aARWcCos:-------:allow"   andrew

$ ls -ldV andrew
drwxr-xr-x+  3 andrew  staff 3 Nov 30 12:40 /export/homes/andrew
         group:winadmin:rwxpdDaARWcCos:fd-----:allow
            group:staff:r-x---a-R-c--s:fd-----:allow
            user:andrew:rwxpd-aARWc--s:fd-----:allow
              everyone@:r-x---a-R-c--s:fd-----:allow
```

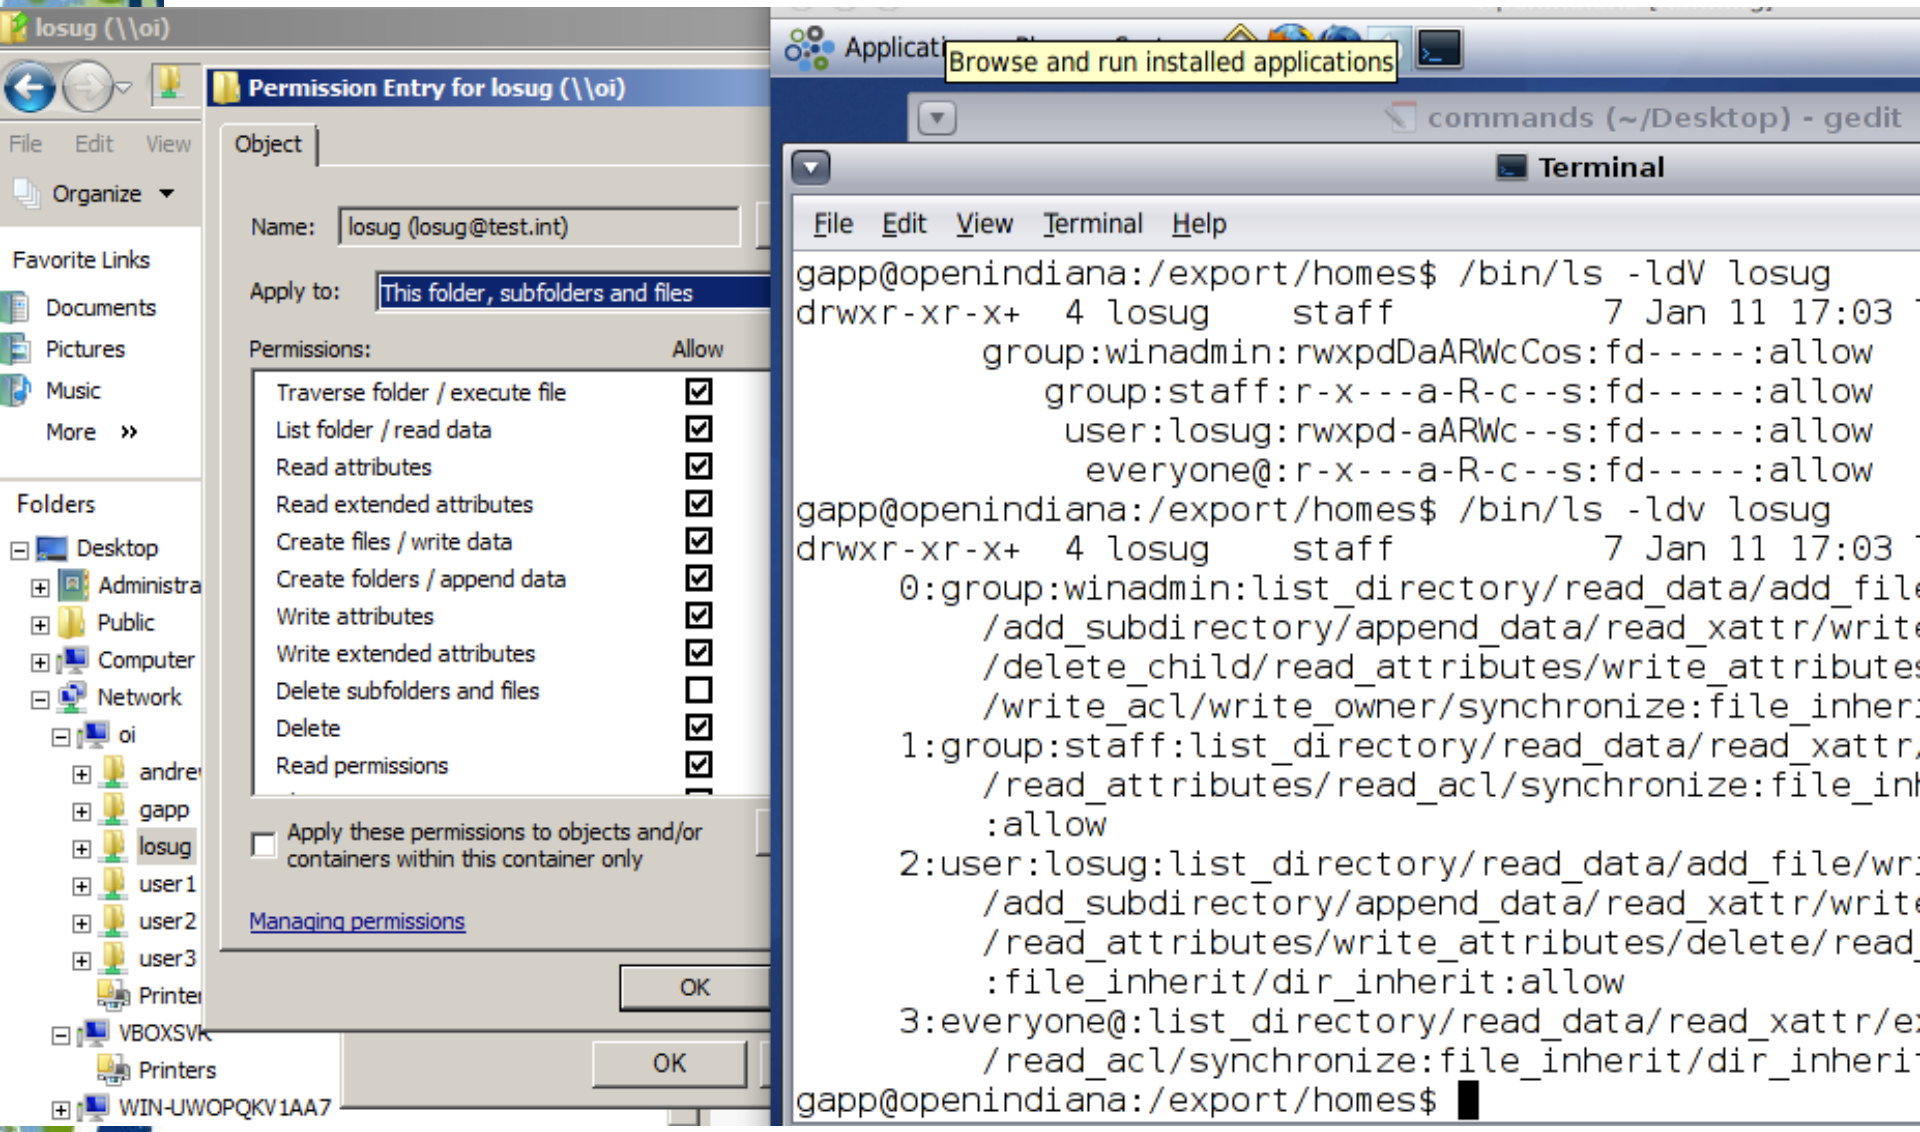
losug (\\oi)
File Edit View
Organize
Favorite Links
Documents
Pictures
Music
More
Folders
Desktop
Administra
Public
Computer
Network
oi
andre
gapp
losug
user1
user2
user3
Printer
VBOXSVR
Printers
WIN-UWOPQKV1AA7
Permission Entry for losug (\\oi)
Object
Name: losug (losug@test.int)
Apply to: This folder, subfolders and files
Permissions: Allow
Traverse folder / execute file
List folder / read data
Read attributes
Read extended attributes
Create files / write data
Create folders / append data
Write attributes
Write extended attributes
Delete subfolders and files
Delete
Read permissions
Apply these permissions to objects and/or containers within this container only
Managing permissions
OK
OK
Browse and run installed applications
commands (~/Desktop) - gedit
Terminal
File Edit View Terminal Help
gapp@openindiana:/export/homes$ /bin/ls -ldV losug
drwxr-xr-x+ 4 losug staff 7 Jan 11 17:03
group:winadmin:rwxpdDaARWcCos:fd-----:allow
group:staff:r-x---a-R-c--s:fd-----:allow
user:losug:rwxpd-aARWc--s:fd-----:allow
everyone@:r-x---a-R-c--s:fd-----:allow
gapp@openindiana:/export/homes$ /bin/ls -ldv losug
drwxr-xr-x+ 4 losug staff 7 Jan 11 17:03
0:group:winadmin:list_directory/read_data/add_file
/add_subdirectory/append_data/read_xattr/write
/delete_child/read_attributes/write_attributes
/write_acl/write_owner/synchronize:file_inher
1:group:staff:list_directory/read_data/read_xattr
/read_attributes/read_acl/synchronize:file_inh
:allow
2:user:losug:list_directory/read_data/add_file/wri
/add_subdirectory/append_data/read_xattr/write
/read_attributes/write_attributes/delete/read
:file_inherit/dir_inherit:allow
3:everyone@:list_directory/read_data/read_xattr/ex
/read_acl/synchronize:file_inherit/dir_inheri
gapp@openindiana:/export/homes$

Birkbeck
UNIVERSITY OF LONDON

References

CIFS Service Troubleshooting

http://wiki.genunix.org/wiki/index.php/CIFS_Service_Troubleshooting

Improved ACL interoperability

http://arc.opensolaris.org/caselog/PSARC/2010/029/20100126_mark.shellenbaum

Solaris 11 CIFS / ACLs

- http://download.oracle.com/docs/cd/E19963-01/821-1449/index.html

http://download.oracle.com/docs/cd/E19963-01/821-1448/ftyxi/index.html