Remote Name Mapping for Linux NFSv4

Andy Adamson
Center For Information Technology Integration
University of Michigan
August 2005

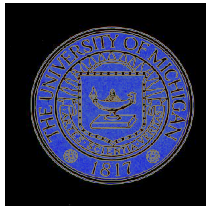

NFSv4 Administrative Domain

- NFSv4: names not numbers on the wire
- NFSv4 domain = unique UID/GID space
- Multiple Security Realms
  - Kerberos, PKI Certificate Authorities (SPKM3)
- Multiple DNS - NIS domains
- Pick one DNS domain to be the NFSv4 Domain Name <user@nfsv4domain>
  - nfsv4domain is used for ACL 'who' and GETTATTR owner and owner_group

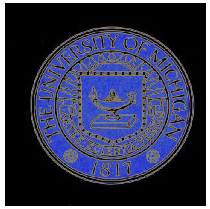

center for information technology integration

NFSv4 Domain

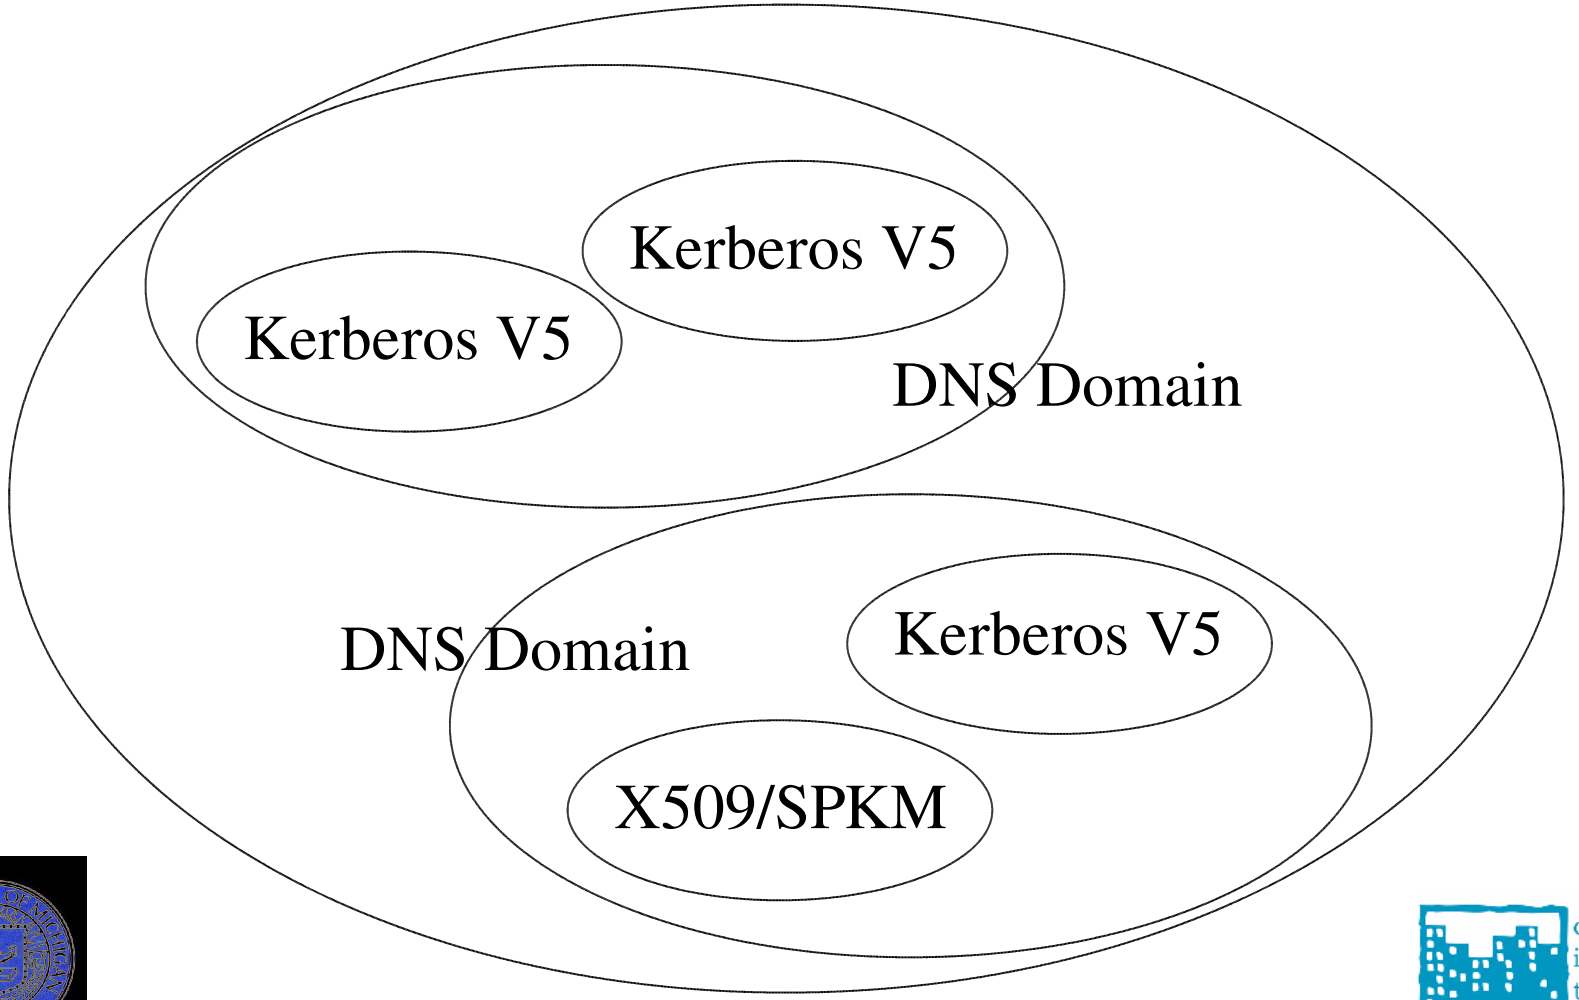

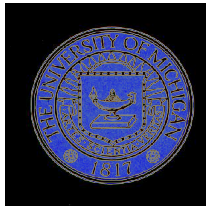

center for
information
technology
integration

Local NFSv4 Domain: Name to ID

- One to one correspondence between UID and NFSv4 domain name
  - joe@arbitrary.domain.org
- GSS Principal name will differ from NFSv4 domain name
  - Kerberos V: joe@ARBITRARY.DOMAIN.ORG
  - PKI: OU=US, OU=State, OU= Arbitrary Inc, CN = Joe User Email= joe@arbitrary.domain.org

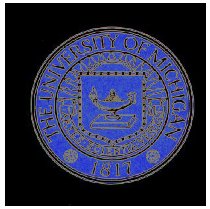

New LDAP Attributes

- We created a new LDAP object to hold two new LDAP attributes for NFSv4 id mapping
  - GSSAuthName
  - NFSv4Name
- We associate one NFSv4Name attribute with a RFC 2307 NSS-LDAP posixAccount to hold the users v4 domain name
- We associate multiple GSSAuthNames with a PosixAccount to hold the users multiple GSS principal names
- Attributes are configurable via /etc/idmap.conf

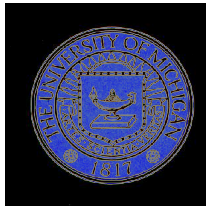

Local Mount: Kerberos V

v4 Domain: arbitrary.domain.org
K5 Realm: TANGENT.REALM
DNS Domain: citi.umich.edu

LDAP

Secure LDAP Call FAILS

nfs/host.citi.umich.edu@TANGENT.REALM

GSSD
/etc/krb5.keytab

NFSv4 Client

GSSD If machine name, map to nobody

NFSv4 Server

gss context call succeeds

nfs/host.citi.umich.edu@TANGENT.REALM

gss context creation

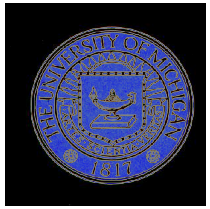

Local Mount: Kerberos V Issues

- Distribution of client keytabs? Linux: yes
  - With no keytab:
    - Allow AUTH_SYS for SETCLIENTID and mount of Kerberos export
    - User Kerberos credentials
- Server: maps machine credentials to nobody (mount)
- Client root user: UID 0?
  - Map to machine principal (no password)
  - Map to per server root principal (with password)

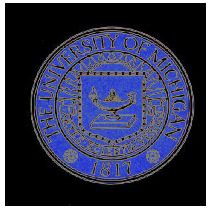

Local Principal: Kerberos V

- New Linux kernel keyring service enables kernel Kerberos credential storage, and PAG-like behaviour
- NSSwitch ID mapping (LDAP PosixAccount)
  - getpwid on principal portion assumes UNIX name (posixAccount uid) == K5 principal
- UMICH LDAP ID mapping
  - GSSAuthName attribute added to LDAP posixAccount to associate with uidNumber
- Server GSSD principal mapping failure = context creation failure

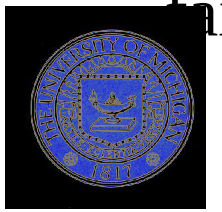

center for information technology integration

Local Principal: Kerberos V

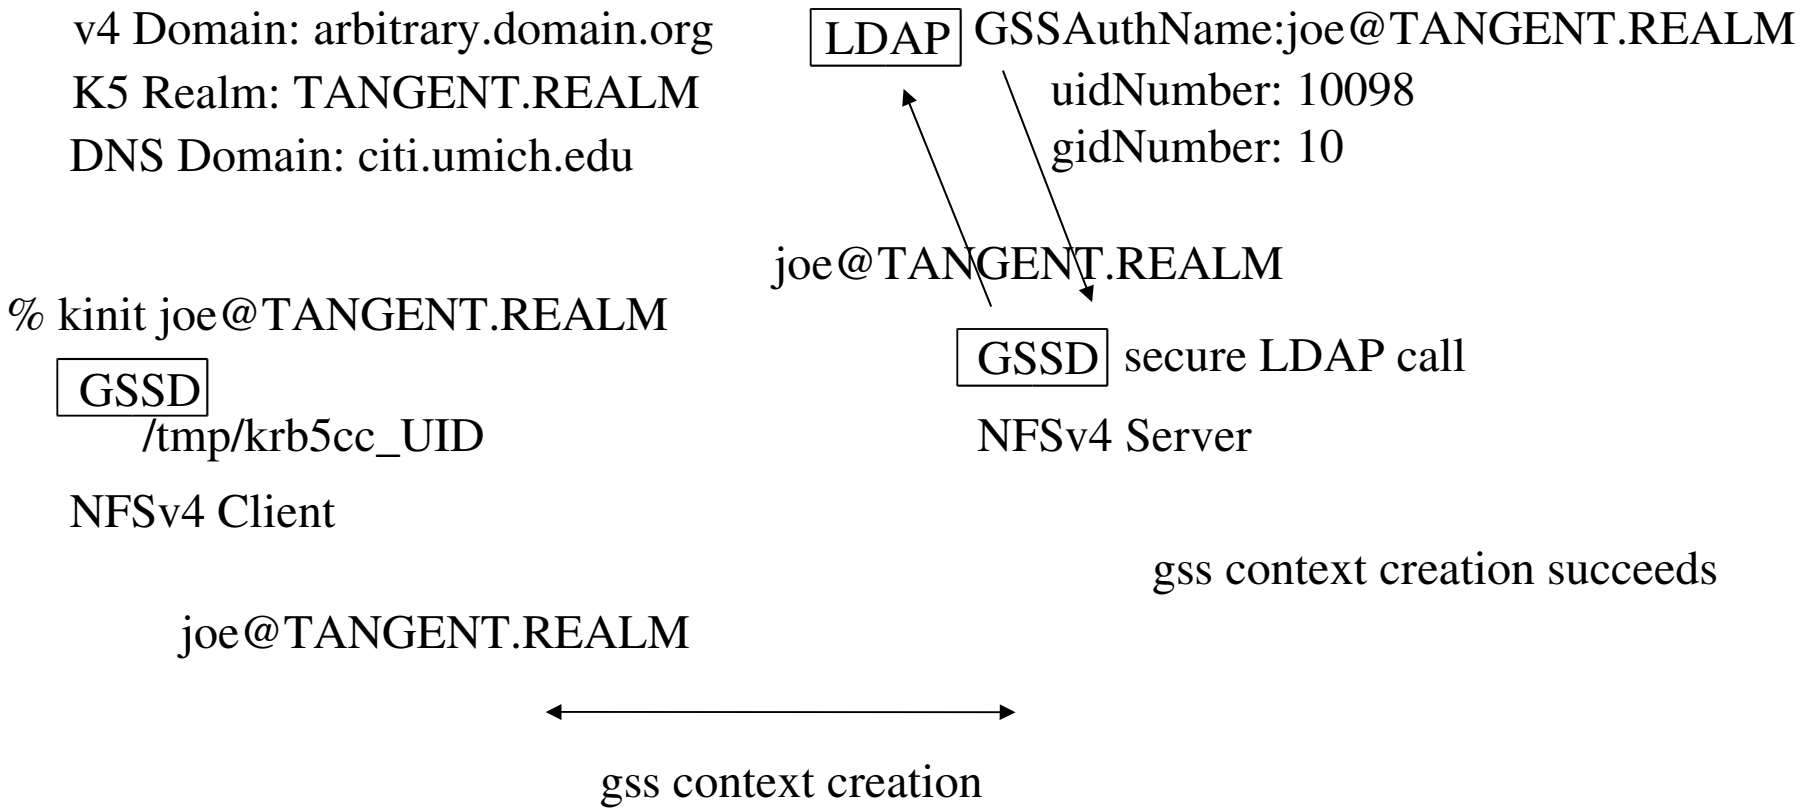

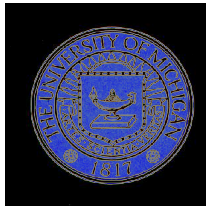

center for information technology integration

Local User: Set ACL issues

- Client setfacl POSIX interface uses UID/GID across kernel boundary (NS Switch)
  - Two name mapping calls
  - NSS posixAccount name (no @nfsv4domain)
  - NFSv4Name attribute added to LDAP posixAccount to associate full nfsv4 name with uidNumber
- New linux nfs4_setfacl interface passes string names across kernel boundary
  - No local name to ID mapping needed

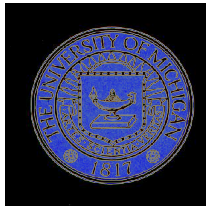

Local User: Set ACL

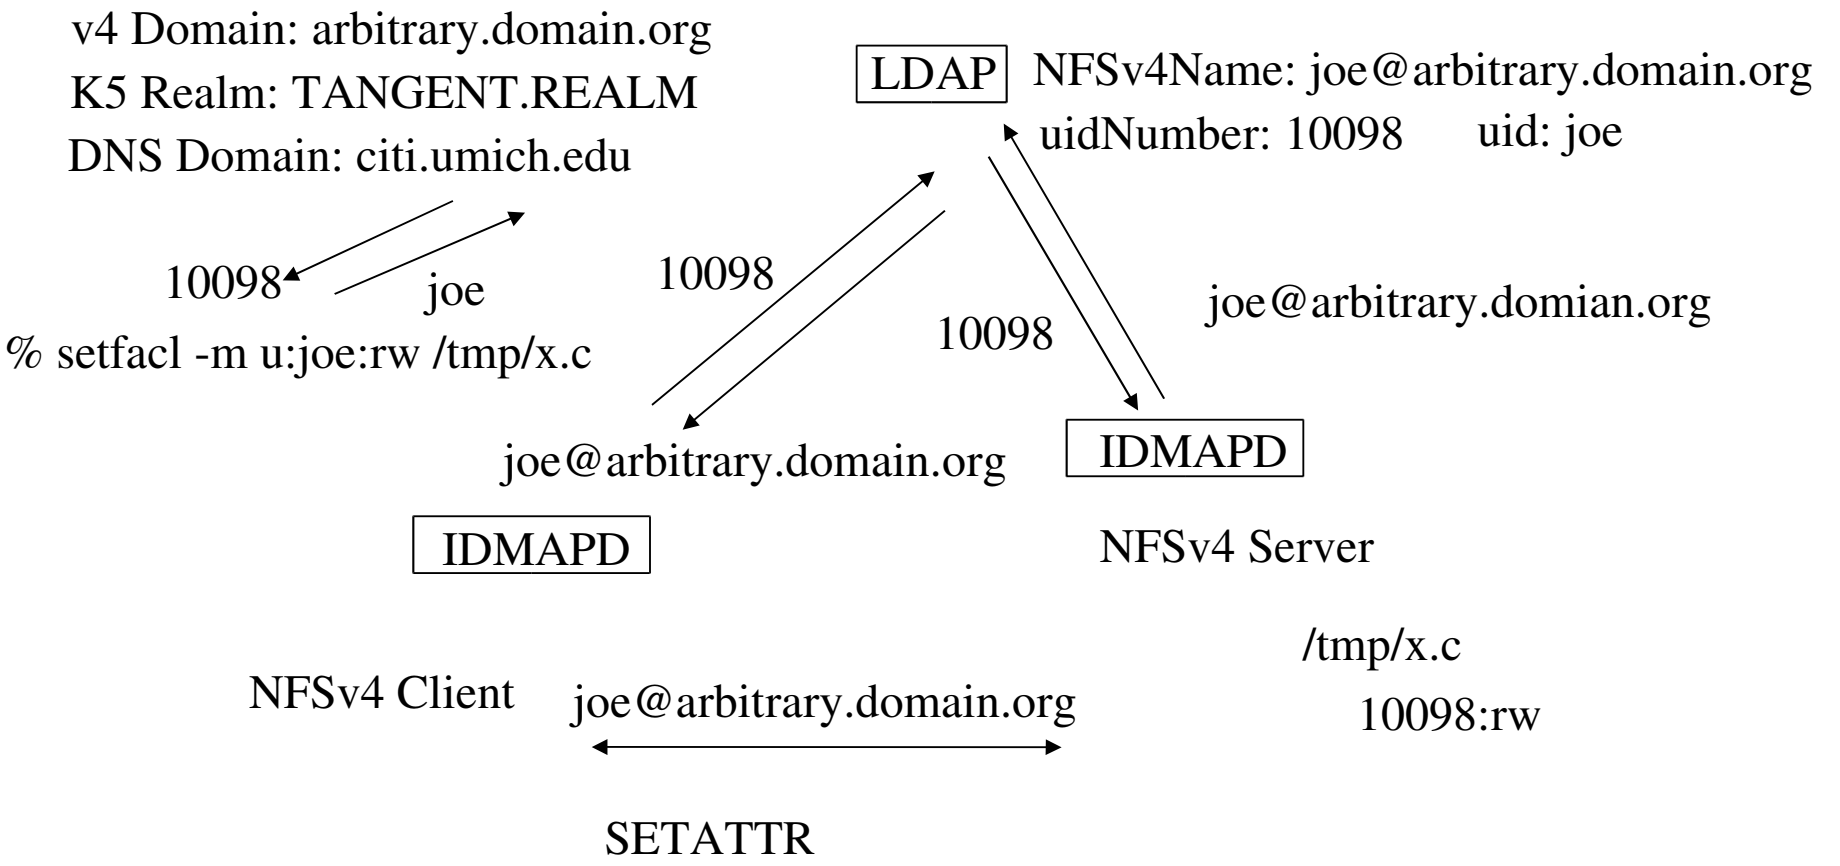

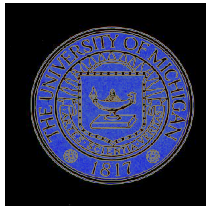

center for information technology integration

Local User: Get ACL issues

- Client getfacl POSIX interface uses UID/GID across kernel boundary (NS Switch)
  - NS Switch posixAccount: uid is displayed
  - Two name mapping calls
- New Linux nfs4_getfacl interface passes string names across kernel boundary

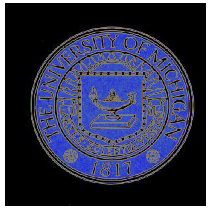

Local User: Get ACL

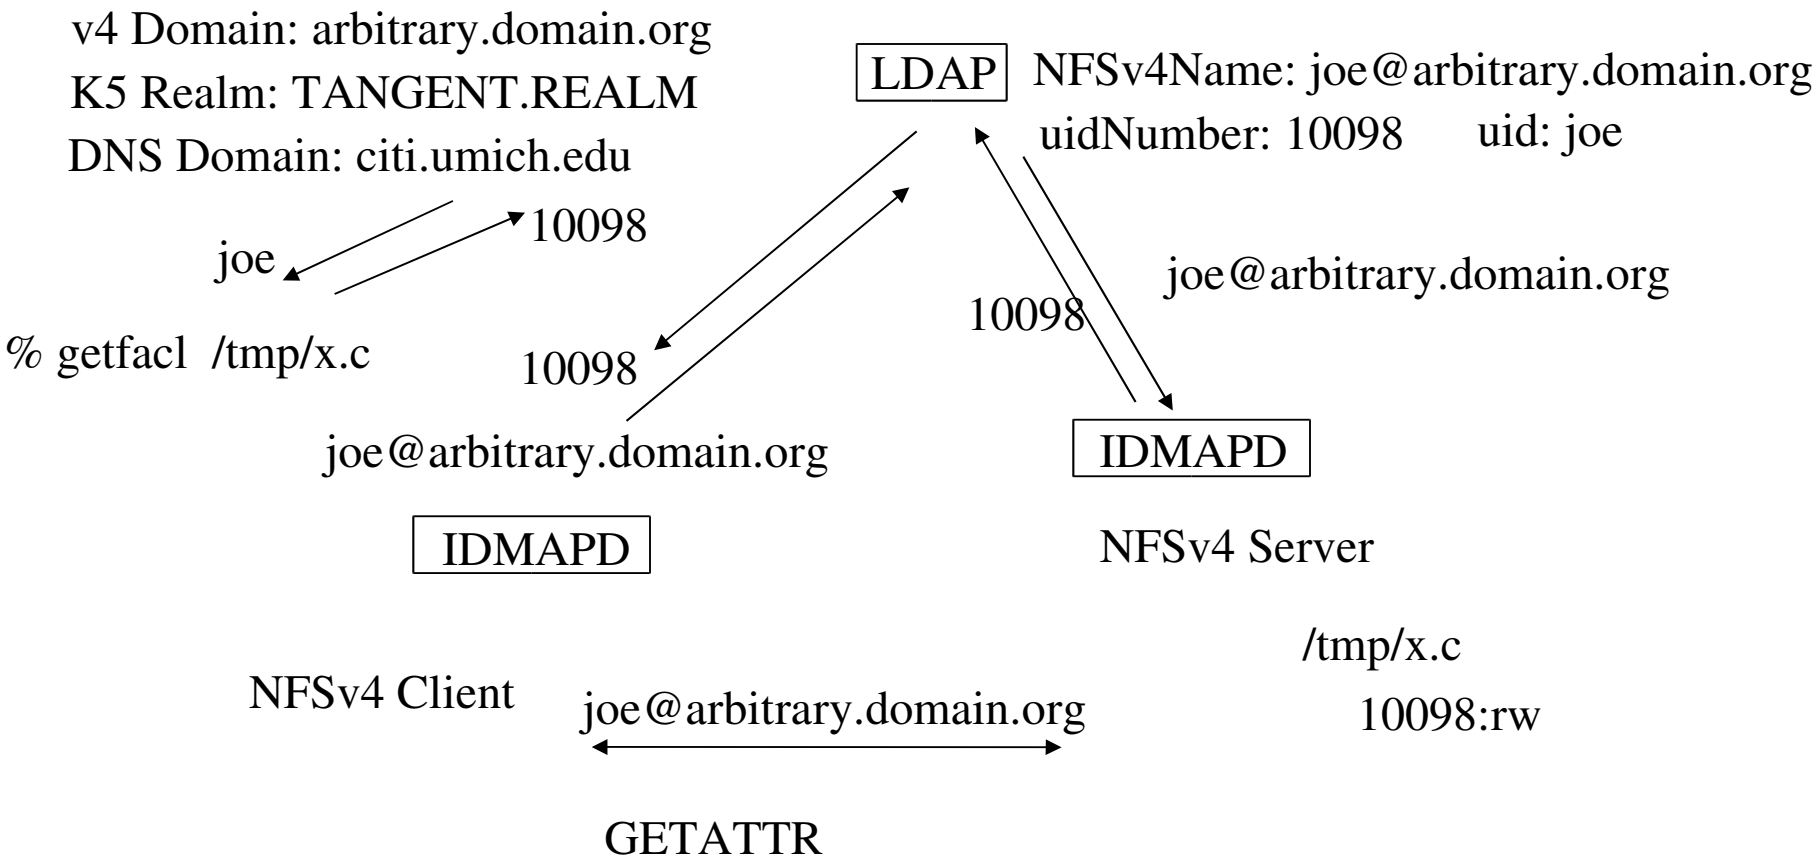

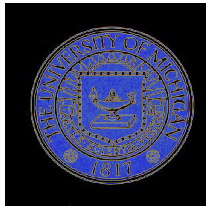

center for information technology integration

Kerberos V X-Realm and Linux NFSv4

- X-realm GSS context initialization just works
- GSSAuthName and NFSv4Name can hold remote user names.
- Need to add posix account with GSSAuthName for UID/GID mapping of remote user
- Set posixAccount shell to /dev/null for NFSv4 remote access without local machine access
- Secure LDAP communication required

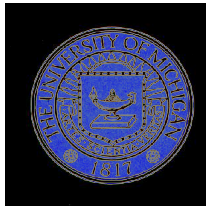

center for information technology integration

Remote Kerberos V Principal

v4 Domain: citi.umich.edu
K5 Realm: CITI.UMICH.EDU
DNS Domain: citi.umich.edu

v4 Domain: arbitrary.domain.org
K5 Realm: TANGENT.REALM
DNS Domain: citi.umich.edu

LDAP GSSAuthName:andros@CITI.UMICH.EDU
uidNumber: 10075
gidNumber: 10

% kinit andros@CITI.UMICH.EDU

GSSD
/tmp/krb5cc_UID

NFSv4 Client

andros@CITI.UMICH.EDU

GSSD secure LDAP call

NFSv4 Server

andros@CITI.UMICH.EDU

gss context creation succeeds

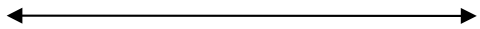

gss context creation

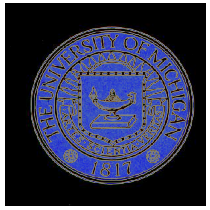

Remote User: Set ACL

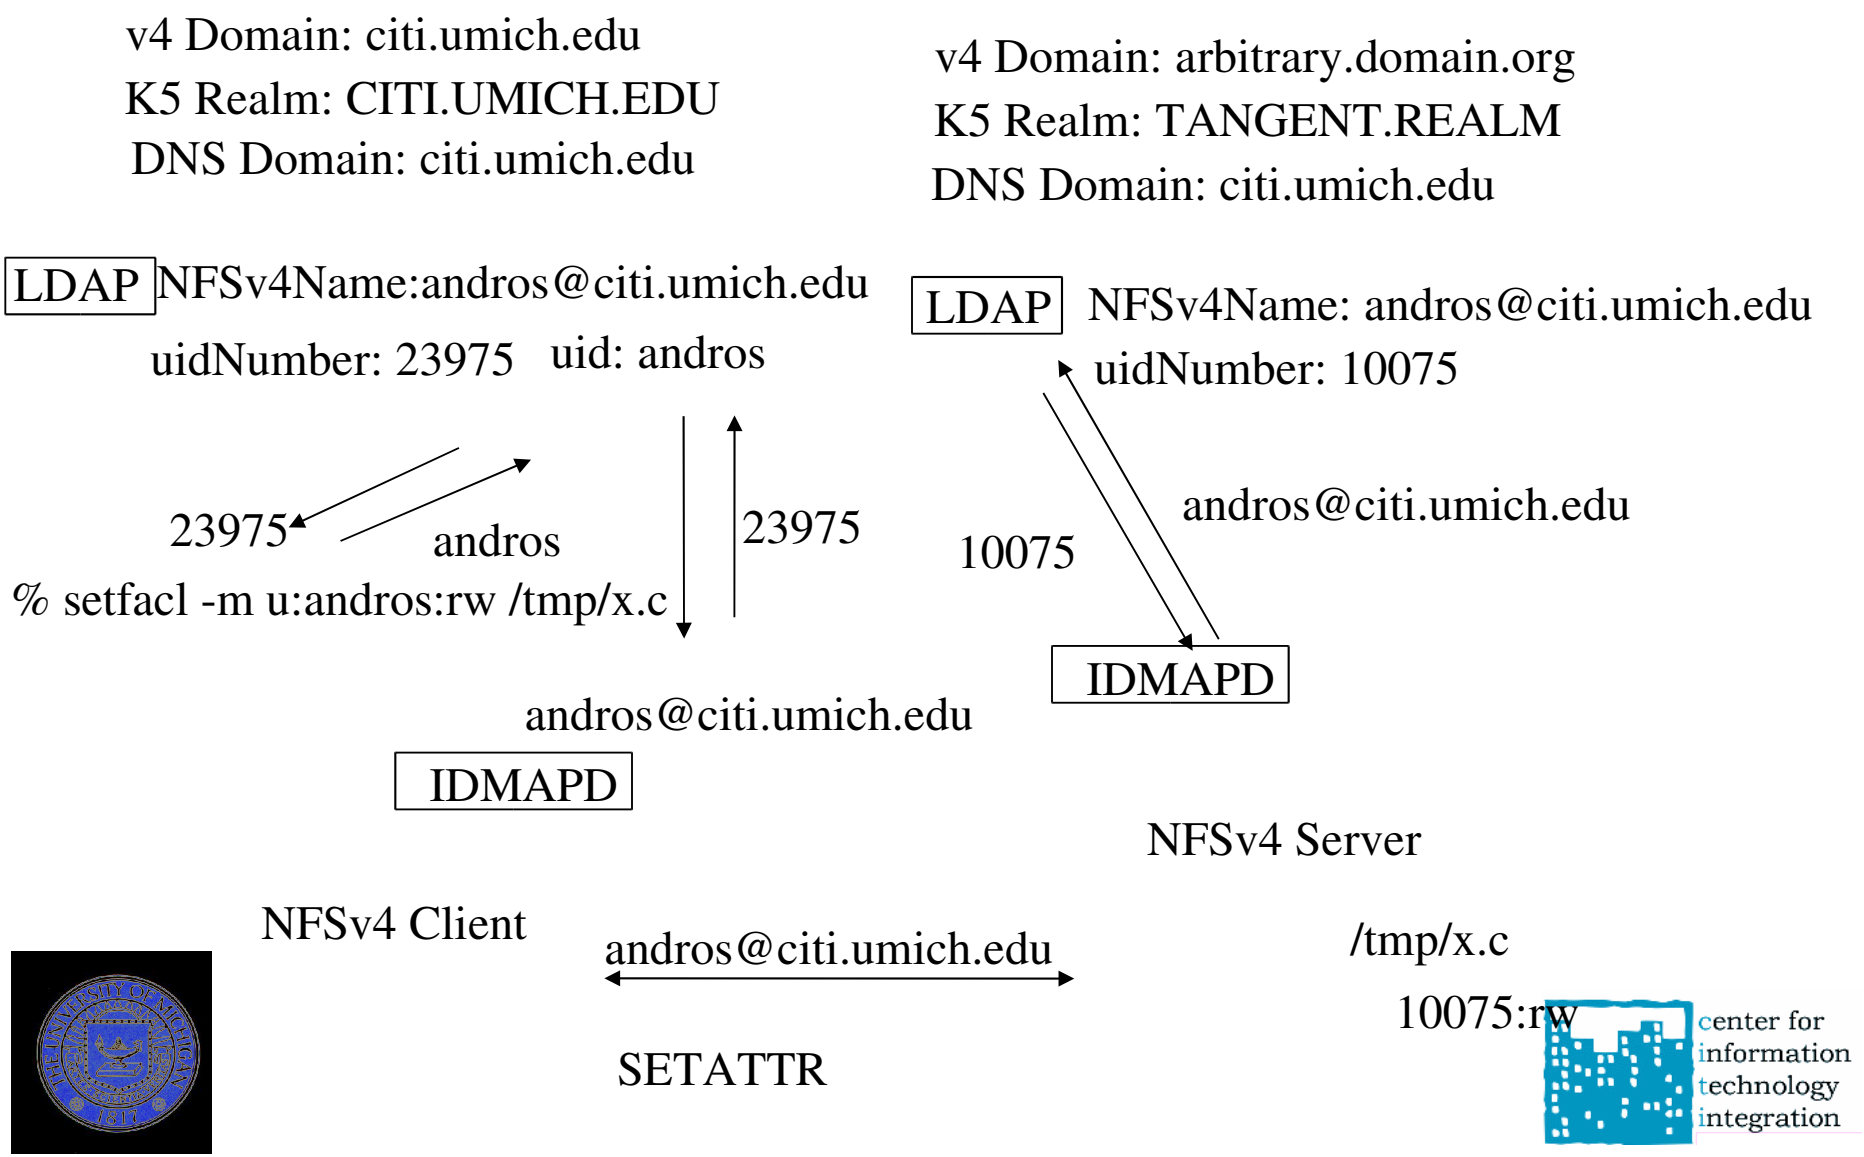

Remote User: Set ACL

- Remote realm: associate NFSv4Name with uidNumber, gidNumber, and GSSAuthName
  - NFSv4RemoteUser schema available
  - NFSv4domain name always used
- Secure LDAP communication required

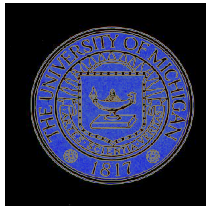

Remote User: Get ACL

v4 Domain: citi.umich.edu
K5 Realm: CITI.UMICH.EDU
DNS Domain: citi.umich.edu

v4 Domain: arbitrary.domain.org
K5 Realm: TANGENT.REALM
DNS Domain: citi.umich.edu

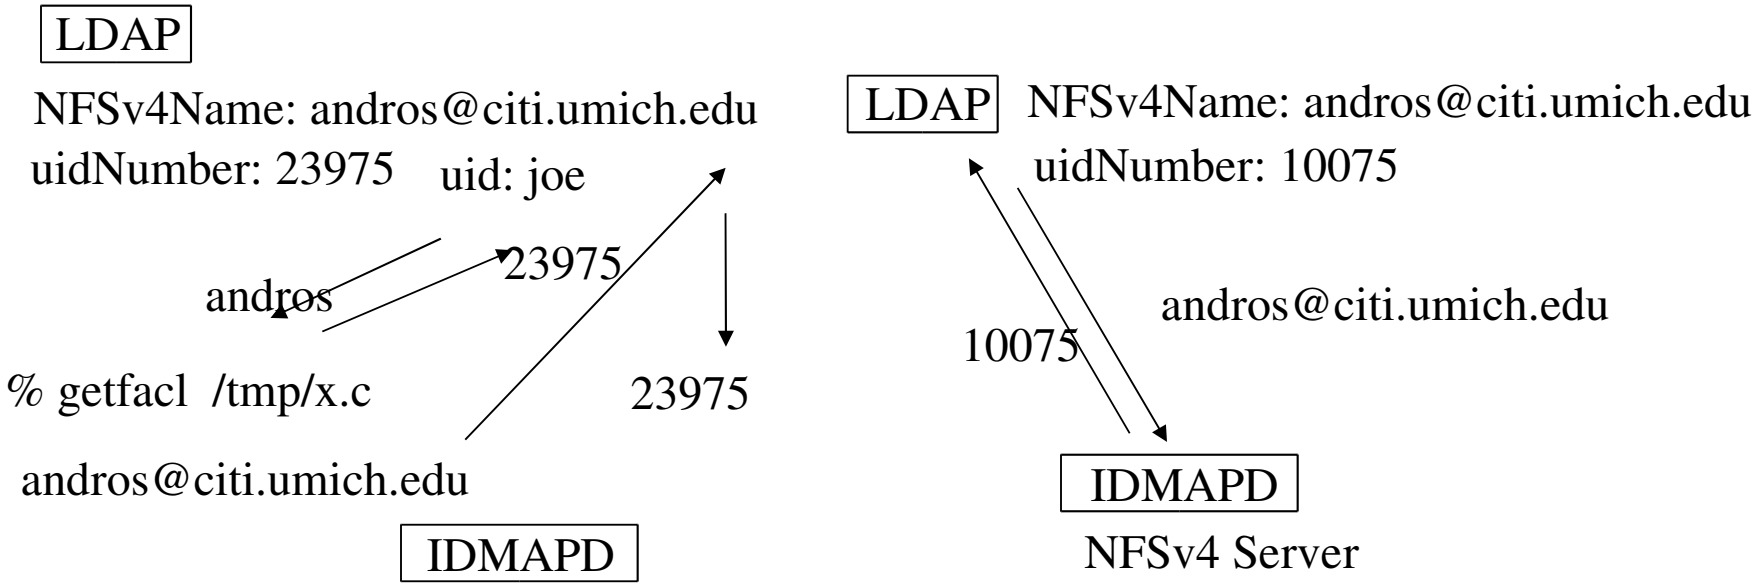

NFSv4 Client

andros@citi.umich.edu

GETATTR

/tmp/x.c
10075:rw

center for information technology integration

Remote User: Get ACL

- LDAP mappings required only for POSIX getfacl
  - NFSv4Name and uidNumber for remote user
  - uid (local user name) for remote user
- nfsv4_getfacl simply displays the on-the-wire ACL name
- Secure LDAP not required

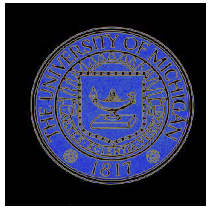

Foreign Groups

- Need design requirements
- Foreign group names could be assigned a local gid
- How does the server resolve foreign membership
  - Callback to foreign NFSv4 domain
  - Only resolve with local uid's (no callback)
  - Pass group list (names) in GSS initialization (a la EPAC)
  - Other?

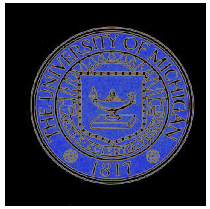

Cross Platform ID Mapping

- NS Switch which uses posixAccount is the common denominator
- Our cross realm mapping extends NS Switch
  - Not supported by other implementations
- IBM also has a cross realm solution

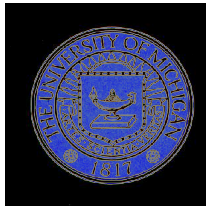

center for information technology integration

Any Questions?

http://www.citi.umich.edu/projects

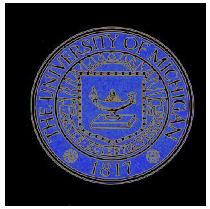

center for
information
technology
integration