







NFSv4 ACL Implementation in Solaris 10  
Lisa Week, Sam Falkner  
Sun Microsystems, Inc.  
Connectathon 2005



# Introduction

-  POSIX-draft ACLs
-  Translation between POSIX-draft and NFSv4
-  Issues with translation
-  Remaining NFSv4 ACL issues

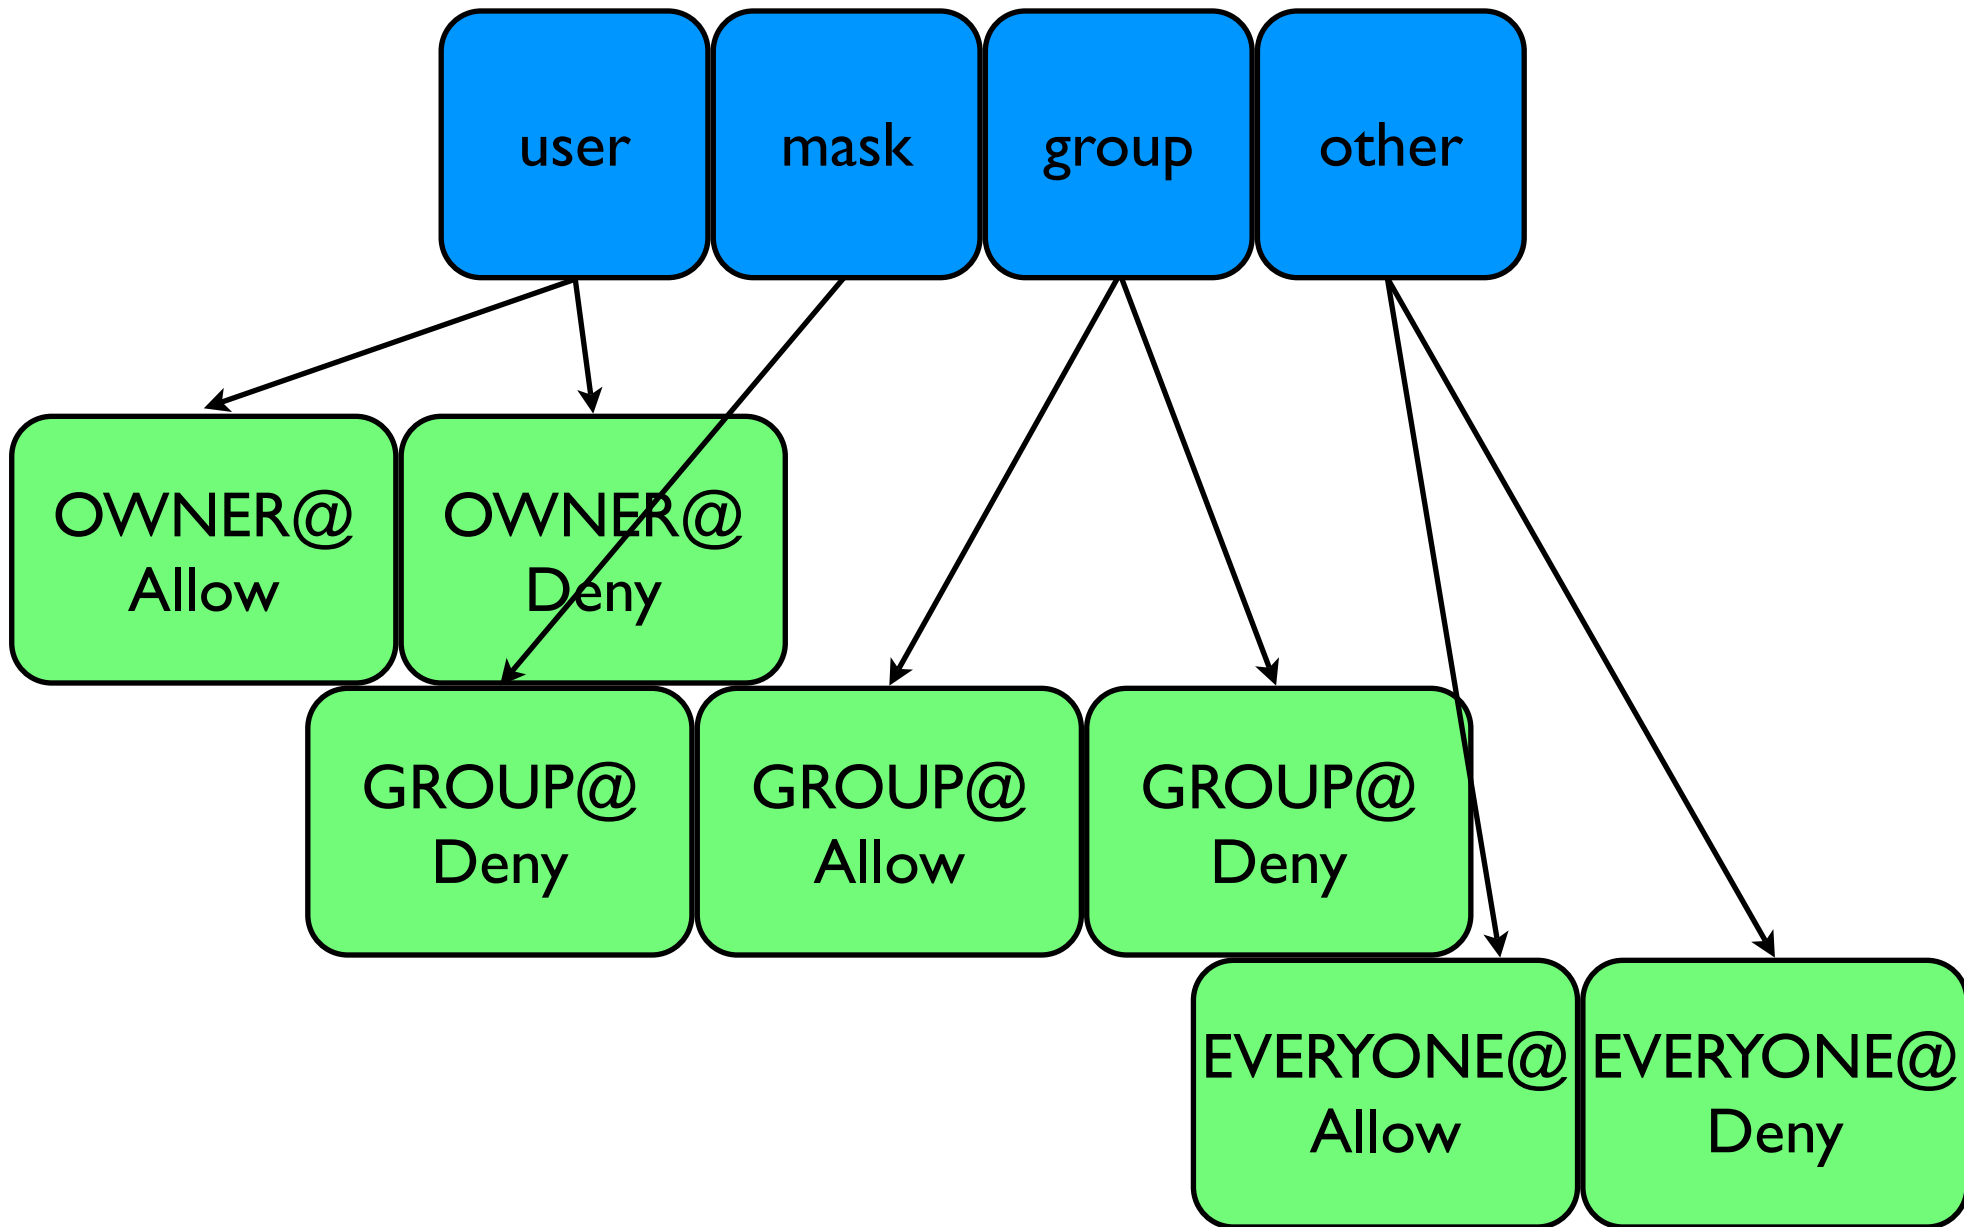
# POSIX-draft ACLs

- Never a finalized POSIX standard
- Full specification of POSIX-draft available here:
  - <http://wt.xpilot.org/publications/posix.1e/download.html>
- Implemented on Solaris Servers (UFS File System)

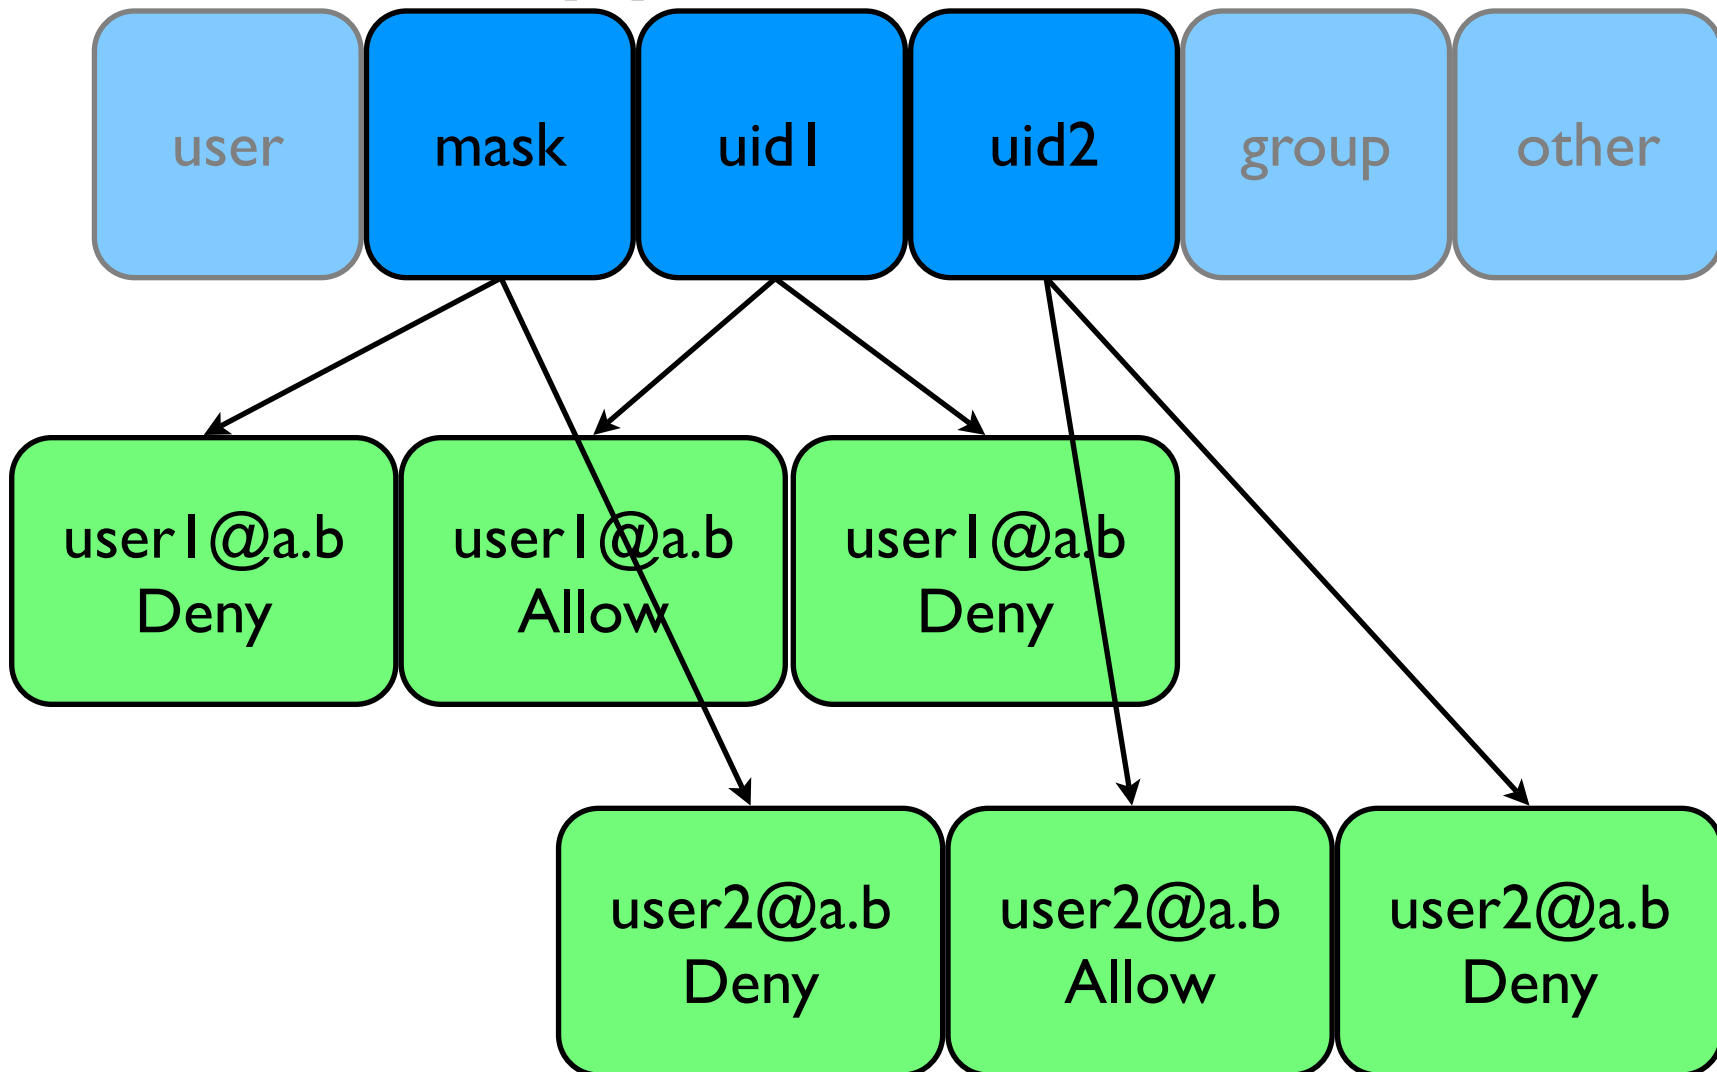
# ACL Translation

- Internet Draft by Bruce Fields and Marius Eriksen
- <http://www.ietf.org/internet-drafts/draft-ietf-nfsv4-acl-mapping-03.txt>
- This version resolves issues found at Bakeathon
- Philosophy: Don't make a file appear to be more secure than the server can enforce

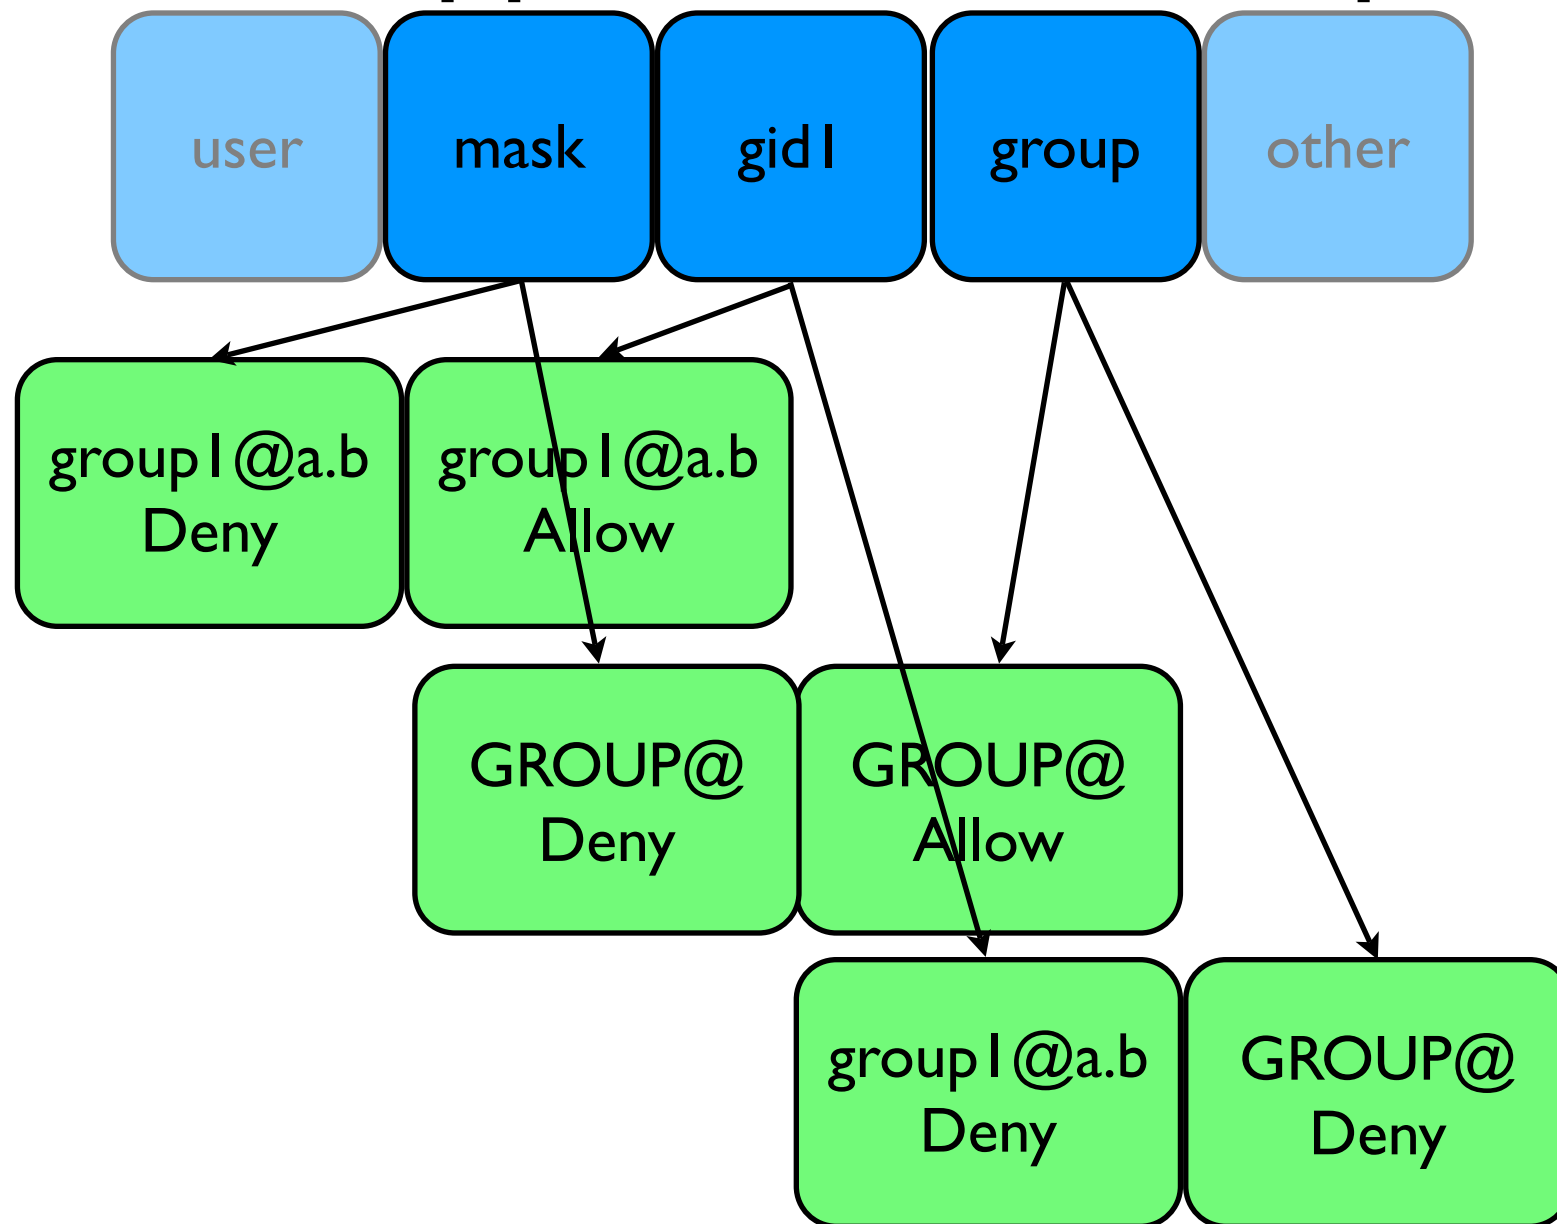
# POSIX-draft to NFSv4



# Supplemental Users



# Supplemental Groups




# Permissions to access\_mask

read	ACE4_READ_DATA
write	ACE4_WRITE_DATA   ACE4_APPEND_DATA
execute	ACE4_EXECUTE
writable directory?	ACE4_DELETE_CHILD
owner?	ACE4_WRITE_ATTRIBUTES   ACE4_WRITE_ACL
everything...	ACE4_READ_ACL   ACE4_READ_ATTRIBUTES   ACE4_SYNCHRONIZE



# Default ACLs

 POSIX-draft default ACLs will be translated to NFSv4 by turning on these flags:

 ACE4\_INHERIT\_ONLY\_ACE

 ACE4\_DIRECTORY\_INHERIT\_ACE

 ACE4\_FILE\_INHERIT\_ACE

# NFSv4 to POSIX-draft

- Reverse the POSIX-draft to NFSv4 Translation
- Not All NFSv4 ACLs can be Translated to POSIX-draft ACLs
  - ACEs in the wrong order
  - Allow/Deny pairs are not complements
  - Certain mask combinations not supported
  - Certain flags are not supported
  - User or group cannot be translated
- Some of the more controversial issues covered in later slides

# When Translation Fails...

- When an NFSv4 ACL cannot be translated to a POSIX-draft ACL
- When setting an ACL, Solaris server returns  
NFS4ERR\_ATTRNOTSUPP
- When a user on a Solaris client tries to receive a POSIX-draft ACL, user receives a fabricated ACL (based on the mode bits)

# Translation Issues

- Certain NFSv4 masks do not map to POSIX-draft ACLs (or even POSIX itself)
  - ACE4\_WRITE\_OWNER
  - ACE4\_DELETE
  - ACE4\_WRITE\_NAMED\_ATTRS
  - ACE4\_READ\_NAMED\_ATTRS
  - ACE4\_SYNCHRONIZE

# Write Attributes

- In ACE4\_WRITE\_ATTRIBUTES, what is included in “attributes”?
- mtime? Setting to an arbitrary time, or the current time?
- atime? (Same choices as mtime)
- mode? Does this go with ACE4\_WRITE\_ACL, or ACE4\_WRITE\_ATTRIBUTES?
- Primary group? Does this go with ACE4\_WRITE\_ATTRIBUTES, ACE4\_WRITE\_OWNER, or ACE4\_WRITE\_ACL?
- Decided to allow ACE4\_WRITE\_ATTRIBUTES only to OWNER@

# Remaining NFSv4 ACL Issues

- Even if we're not using POSIX-draft ACLs, POSIX itself isn't going away
- Solaris (and others?) needs to uphold POSIX access control semantics
- Create with inheritable ACEs: what to do with the mode passed to `creat()`?
- What to do with an existing NFSv4 ACL when doing a `chmod()`?

# Remaining NFSv4 ACL Issues

- What does ACE4\_READ/WRITE\_ATTRIBUTES cover; what should it explicitly not cover?
- What does ACE4\_READ/WRITE\_NAMED\_ATTRS cover?
  - Listing the named attribute directory, or
  - Reading/writing the named attributes themselves
- ACE4\_DELETE versus ACE4\_DELETE\_CHILD

# Questions?

 Email

 [sam.falkner@sun.com](mailto:sam.falkner@sun.com)

 [lisa.week@sun.com](mailto:lisa.week@sun.com)