Real world
MLS label translation in SELinux

Problem

- Current mcstransd does direct mapping
- Storage requirement scales linearly with number of translations supported
- Potential unrealistic storage requirement
  - 1024 Category bits
  - 10^308 possible combinations
  - 10^310 bytes to represent (10^292 Exabytes)

Scope the problem

- 268 Geographical entities in STANAG 1059
  - Still 10^80 combinations
- More than a few hundred combinations unlikely
  - But we don't always know which combinations in advance

Can we manage by exception?

- Manually add new translations when required
- Might be beyond the capabilities of the site
- Accreditation and test impact of editing security critical files
- Only feasible if new translations are rare and update time is available

Modify mcstransd

- Translate words to combinations of categories in raw label
- Support aliases (different words can translate to the same categories)
- Maintain word order (SECRET DOG CAT to s5:c11,c14 to SECRET DOG CAT, not SECRET CAT DOG)
- Support multiple domains of interpretation
- Support modularity

Implementation Concept

- Base translation
  - Level
  - Categories
- Modifier Groups
  - How words modify the base translation
  - Capture order dependencies
- Include files to improve modularity
- Fixed translation escape mechanism

New setrans.conf Syntax

Domain=Default
s0=SystemLow
s0=syslo
s15:c0.c1023=SystemHigh
s0-s15:c0.c1023=SystemLow-SystemHigh
...
Base=Sensitivity Levels
s1=UNCLASSIFIED
s1=UNCLAS
s1=U
...
s3:c0,c2,c11,c200.c511=CONFIDENTIAL
...
s4:c0,c2,c11,c200.c511=SECRET
...
Include=/etc/selinux/mls/mcstrans.d/rel.conf

Modifier Group Syntax

ModifierGroup=Inverse Releasable To
Whitespace=- ,/
Join=/
Prefix=RELEASABLE TO
Prefix=RELEASEABLE TO
Default=c200.c511

~c200.c511=EVERYBODY

Aruba - bit 201
~c200,~c201=ABW
~c200,~c201=AA
Afghanistan - bit 202
~c200,~c202=AFG
~c200,~c202=AF
…
Zimbabwe - bit 444
~c200,~c444=ZWE
~c200,~c444=ZI

Translation Approach

CONFIDENTIAL WORD1 WORD1

- Find 'Base' regexp that matches CONFIDENTIAL
- Walk modifier group tables in order looking for matching regexp
  - Make category bitmap changes
  - Iterate until nothing but whitespace left

Translation Approach

s1: c0,c2,c4.c9

- Find Base with smallest Hamming Distance
- Walk modifier group tables finding words that consume bits (shorten Hamming Distance)
  - Use the word that minimizes Hamming Distance
  - Iterate until all bits consumed
  - Emit in original table order

Status

- Prototype released
  - http://www.nsa.gov/selinuX/list-archive/0806/26366.cfm
- Supports prefixes, suffixes, join strings, whitespace definition and arbitrary combinations of words
- Translates a broad array of labels
  - S RELEASABLE TO AFG/CAN/ZWE
  - R HANDLE VIA SNEAKERNET CHANNELS ONLY

Issues

- No constraints
  - Allows translation of invalid labels
- Multi-domain support not exposed through API
- Breaks 'semanage translation'
- Can't handle embedded '-' in translated labels
- Implementation violates libsepol encapsulation

Encapsulation Issues

- Copied private mls_level_*_string routines out of libsepol
- Uses ebitmap routines from static libsepol
- Added several new ebitmap routines

Encapsulation Violation Discussion

- Hamming distance bit consumption calculation needs bitmap for adequate performance
- Silly to create new bitmap routines
- Bitmaps need to be converted to normal string
- Required code is private to libsepol