Persistent Reservation Problem & Enhancement Requirements

Roger Cummings
Symantec
roger_cummings@symantec.com

Background

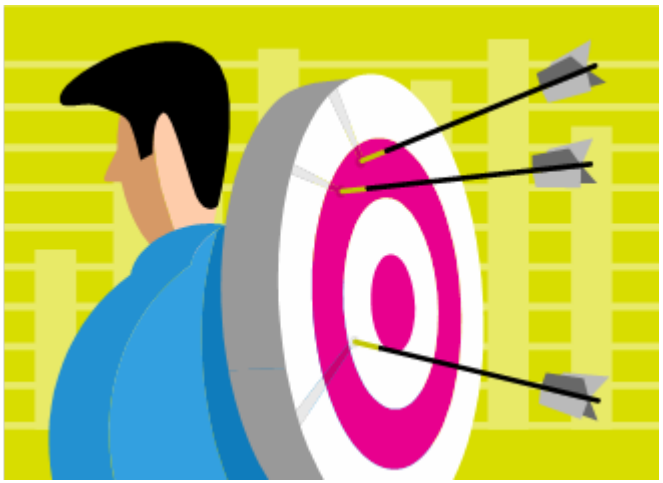

- On 12/10, Kevin Butt uploaded a proposal (08-025) & slides (08-024) for a new Persistent Reservation (PR) type
  - We thanked Kevin for providing a target!
  - And loosed a whole bunch of arrows

Problem

- Apps that deal with tapes have to date only used exclusive PR types
  - Major concerns about other servers accessing the media
    - Other PR types allow ANY server that registers to gain access
  - BUT….
- Most of those apps are now distributed across multiple servers
  - Using only exclusive types means all data to/from the media has to flow thru 1 server (aka bottleneck)

Need

- A way to use PR to provide access only to a designated set of Initiators
  - We’ve taken to calling this designated set a “Group”
    - Probably a bad term
    - Feel free to propose a better one but expect flames
- Note that today there are four different PR types (5-8h) that provide “shared” access
  - All are presently in use in real-world applications
  - We need “group” versions of all of them

Requirements

- One way of defining a Group for all PR types
- Don’t try to define Group only on a time basis
  - All members of the Group must NOT have to be registered at the time a reservation is created to get access
  - Or have to stay registered for the life of the reservation (they should be able to register again and gain access*)
- * Keep the same PR holder definition for the Group types as the equivalent non-Group type
- Minimize the impact on existing PR definitions, especially preempt & abort & All_TG_PT
- Must be able to use Group & non-Group types in same configuration
  - Don’t assume that all registering Initiators and reserving Initiator can be updated at the same time
  - But also must be no “leakage” i.e. an old type registering Initiator must not be able to fluke membership in a Group

3 ways to Define a Group?

- In Advance
  - By transport ID or some such scheme
  - Don’t let non Group members register
- At time of registration
  1. By specific key value, or
     - Impacts preempt rules & other existing functions etc.
  2. New REGISTER type include Group ID (GID)
     - Allow 1 register to indicate membership in N groups
- At time of reservation (Kevin’s proposal)
  - Using SPEC_I_PT bit & Transport IDs

Issues

- Reserving Initiator may not know the transport ID of all the Initiators it wants in the group
- Existing Initiator may be confused when Registering doesn’t enable access when a PR type other than 1h or 3h is in place