Persistent Reservation
What if the keys are not unique?

Tom Coughlan
Digital Equipment Corporation

To be as Specific as Necessary, and no Moreso

- The current definition of Persistent Reservation (intentionally) leaves the Device Server's behavior unspecified when multiple Initiators use the same Reservation Key.
- We should be more specific on this topic.
  - The advantage outweighs the cost.

PR Recap

- Each initiator (that is, each port on a host) registers a key
  - Device Server remembers the key-to-initiator mapping
- Initiator passes the key in each PR Out.
  - Device Server checks for a match, using key-to-initiator mapping.
  - Device Server remembers reservations, including the reserving initiator ID and/or the key.
- Initiator can request a list of all Reservations.
  - Device Server returns reservation descriptor, which includes the key of the reserving initiator (not its ID)

PR Recap (cont.)

- Initiator requests a preemption by issuing a PR Out specifying a key.
  - Device Server uses the key-to-initiator mapping to determine which initiator to preempt.

What if keys are not unique?

- Behavior is unpredictable.
  - Presumably, the Device Server shall be prepared to remember each initiator's key, even if not unique.
  - If Device Server locates Reservations by the initiator's key, then Reservations will be "shared".
  - Preempt may apply to all initiators with the same key, or Device Server may just pick one, or...?
- This gives implementation flexibility, but it may have more disadvantages than it is worth, for the typical applications...

Why Implement Non-Unique Keys?

- One key per node, same on all ports, is adequate for typical applications:
- Typically, Preemption is used to remove a *node* from a cluster.
  - The goal is to remove the state associated with *all* of a node's ports from the device server.
- The key-to-initiator map consumes scarce non-volatile storage in the Device Server
  - Less space may be required when there is just one key per node.

Proposal

- Change SPC-2 to specify the behavior of the Device Server when multiple initiators use the same key.
  - Reservations are identified according to the Initiator's ID, not it's key.
  - Preemption applies to all initiators with the key specified in the Service Action key.
- No change in behavior when all initiators register with the same key.