FCA Interfaces Changes to Support NPIV

Aaron Dailey
aaron.dailey@sun.com
Allan Ou
allan.ou@sun.com
Jack Meng
jack.meng@sun.com

Sun Microsystems
Version 0.2
30 July 2007

Theory of Operation
The Leadville device driver architecture is already well suited to supporting NPIV. Leadville is not a monolithic driver, but rather a collection of drivers with specific relationships. Today, there is currently a 1 to 1 relationship between FCA driver instances and fp driver instances. To support NPIV, we will change this relationship to 1 to many. So, for one FCA driver instance, representing one HBA port, we can have more than one fp instance. Each fp instance represents an NPIV Port ID or Permanent Port name.

We will create a command, `fcadm`, which will be responsible for creating/deleting and reporting status on NPIV ports. NPIV ports will be created and deleted dynamically i.e. without reboot. To implement this, we will modify HBA API. `fcadm` will actually be created as a hardlink to the existing `fcinfo` command.

`fp` instances will be created and deleted on demand using the ndi hotplug interfaces.

Glossary
Permanent Port Name: The port name that is sent with the first FLOGI, and remains constant. This is probably stored in non-volatile memory on the HBA
NPIV Port ID: The port WWN/Node WWN pair associated with an NPIV Port

Changed structures:

fca_tran_t
`existing fields:`

`int fca_version`
set to FCTL_FCA_MODREV_4, prior revisions unsupported

`new fields`

`int fca_num_npivports`
should be set to 0 for unsupported, !0 = number of NPIV ports supported

`la_wwn_t fca_perm_pwwn`
This is the permanent port wwn for the port. This is necessary so we can find related NPIV Port Ids in FibreChannel transport

fc_fca_bind_info_t

`new fields:`

`int port_npiv`
is this port a npiv port - 0== false, !0 = true

`la_wwn_t port_nwwn`
NPIV node wwn – valid if port_npiv != 0

`la_wwn_t port_pwwn`

NPIV port wwn – valid if port_npiv != 0

New Error Values

`FC_NPIV_FDISC_FAILED`

This error is returned if the FDISC command for the port failed. This implies the switch does not support NPIV.

`FC_NPIV_FDISC_WWN_INUSE`

This error is returned if the NPIV ID is already in use by a different permanent port name, as reported by the switch

`FC_NPIV_NOT_SUPPORTED`

This error is returned if the HBA does not support NPIV

`FC_NPIV_WRONG_TOPOLOGY`

This error is returned if the topology does not support NPIV (i.e. loop or point to point)

`FC_NPIV_NPIV_BOUND`

This error is returned when calling fca_unbind on a permanent port, and there are bound NPIV ports. No ports are unbound in this case.

fc_fca_port_info_t::fca_port_attrs_t::fca_hba_fru_details_t. This information is used for matching physical ports on an HBA. All ports, physical and NPIV, located on a given physical port, should use the same values for `high` and `low` and `port_index.`

Changed Driver/FC Transport Functions

fca's `_init` requirements are unchanged - must still call `fc_fca_init`

`fca's attach/detach` are unchanged - must still call `fc_fca_attach/fc_fca_detach.` The FCA driver does inform the driver if the HBA supports NPIV, and if so, how many ports are available. `detach` should fail by returning `DDI_FAILURE` if any NPIV or physical ports are bound.

Fibrechannel transport will now call `fca_bind_port` multiple times for each port to be enabled. The transport will first call bind for the permanent port name then for each NPIV port ID. The FCA driver must then enable NPIV instance, and should call back, as today, using fc_fca_bind::port_statec_cb(..., FC_STATE_ONLINE) when the port is ready.

New error returns:

`FC_NPIV_FDISC_FAILED`

`FC_NPIV_NOT_SUPPORTED`

`FC_NPIV_WRONG_TOPOLOGY`

After bind, the fc_fca_port_info.pi_login_params should contain the information for the new port

Similarly, `fca_unbind_port` will be called for the permanent port name only when all NPIV port ID have been unbound. FCA driver's `detach(9E)` entry point will only be called after all ports are unbound.

To be able to return an error, `fca_unbind_port` now returns an `int.` Return values are:

FC_SUCCESS successful return

FC_UNBOUND is already unbound.

FC_NPIV_NPIV_BOUND attempt to unbind a permanent port ID when NPIV ports are bound. No ports are unbound in this case.

Asynchronous Event Management

fc_fca_bind::port_statec_cb. For port state changes, the FCA driver will call back once for each NPIV port ID and once for each permanent name associated with the physical port.

fc_fca_bind::port_unsol_cb. The FCA driver need not do any special processing. In this case, the SAN will generate unsolicited requests for each NPIV port ID appropriately.

Unfinished:
fca_ub_alloc()