

Distributed FCF Requirements for High Availability

Steven Wilson

Dave Peterson

Howard Johnson

John Crandall

T11/11-162v0



What has been proposed

- In the proposed Distributed FCF model, FDFs cannot operate without an cFCF. cFCF Redundancy is therefore required to achieve a highly available environment
- Redundancy allows a secondary cFCF to take over if the primary becomes not operational
- The primary and secondary cFCFs must therefore remain synchronized. Augmented VE_Port links are used for the redundancy protocol used to keep the cFCFs synchronized.
- In a simple environment the redundancy protocols seem to work, but are they enough?



Redundancy Protocol Requirements

- For HA, two cFCFs shall be required
- Redundancy protocols between the controlling cFCFs should be able to use any transport network that is available.
 - These protocols should not be restricted to the Ethernet network in the Distributed FCF
 - Why not use the FC Fabric as a transport as well?
 - Good ole FC Fabric routing and operation is very reliable
- If the FC Fabric is used for the redundancy protocols, then relevant SW_ILSs need to be defined between Primary and Secondary cFCFs.
- Within the Distributed FCF at least two augmented VE_Port links between the cFCFs should be required



Heartbeat Requirements

- Currently the heartbeat details have just been proposed.
 - The payload contains the two cFCF switch names
- The heartbeat payload should contain information (e.g., identifiers, counters, timestamps) to manage heartbeat requests and responses
 - This is particularly important when multiple networks and/or links are used to transport heartbeats between the cFCFs
- The heartbeat interval needs to be short enough to ensure continued control and data path operation between the FC Fabric, Controlling FCFs, and FDFs in the FDF Set



Failover Requirements

- If the primary FCF becomes not operational, the secondary FCF takes over, becomes primary and a new secondary is selected
- To ensure availability, a new secondary FCF should be chosen as soon as possible
- Do we want to limit the redundancy model to only two cFCFs? Possible solution a pool of FCFs
- The bottom line is that dFCF needs to have the same resiliency and reliability as ANY FC Switch or Forwarder



Management Requirements

- The Distributed FCF is identified by the Switch_Names of the Controlling FCFs
- Also included in the Distributed FCF are the Switch_Names of the FDFs in the FDF Set
- This configuration needs to be propagated to the Fabric Configuration Server (topology) in the attached FC Fabric(s)
 - From an FC Fabric perspective the Domain IDs and Switch Names for the two cFCFs and the Distributed FCF and must be shown as they are visible to FSPF and other Fabric operations
 - To facilitate serviceability actions and performance analysis, the FDF Set in the Distributed FCF should have the capability to be shown as well





Thank You !