01/02/07 - DB2/Z on Parallel Sysplex clarifications

Question1 - I didn't understand how to depict the typically configuration of DB2 instances however. I have depicted having multiple DB2 instances "cloned" across multiple systems as with applications in parallel sysplex. Is this correct or should I have these DB2 instances names all be unique?

Every member of a DB2 Data Sharing Group must have its own unique DSNZPARM therefore, names in this sense must be unique (slide 10 looks to have mirrors of DB2 instances). A DB2 minimum instance is made up of a Services Address Space (MSTR), DataBase Address Space (DBMS) and an IRLM Address Space (3 regions in total). Data sharing requires that all DB2-related resources reside on shared disks (DB2 Catalog and Directory, MVS/ICF Catalog and user data). This is inclusive of CF member logs and Bootstrap datasets (I recommend workfiles as well). As far as locking, each DB2 in the data sharing group must have its own IRLM (its own startup procedure), therefore during initialzation, DB2 members join one XCF Group and the associated IRLMs join another XCF group (these connect to the lock structures in the CFRM Policy).
Note - there can be multiple DB2 instances within an LPAR or inter-LPAR all participating within the same Data Sharing Group. Slide 10 also may be misleading visually, although can be explained verbally to state the mainframe picture represents a separate DB2 instance (not a separate mainframe) .

Question 2 - a) I wanted to show both the normal operation of DB2 in conjunction with the sysplex distributor and WLM, as well as a similar failover scenario to what you showed with applications, only in this case for DB2. As with my previous question, I am not sure how to depict DB2 within the cluster.
b) Is it "cloned" like it is shown in the slide, or would it really be separate DB2 instances that could be spread across the systems, with multiples in separate LPARs. The chart is supposed to depict one data sharing group.

You are right in this slide displaying one data sharing group. I'm not sure of what you mean by "cloned". A data sharing group can have members at different release and version levels, therefore separate DB2s (as well as IRLMs). This slide is only a small sample of request spraying. Setting up the Network side can be complex especially with mixed protocol. The Diagram (12) displays one implementation using the Sysplex Distributor(SD). Note - this could be any edge-server, router, gateway, or external switch. The datashare group can be part of an SNA network, a Transmission Control Protocol/Internet Protocol (TCP/IP) network, or part of a network that uses both protocols. The data sharing group has a single-system image to requesting applications, whether those requests are coming in through TCP/IP or SNA network protocols. Requesting applications use the LOCATION NAME of the data sharing group to direct their SQL requests to that group. As an example - for VTAM one way is each DB2 instancane within the data sharing group has a unique DB2 NETID.LUNAME (LU = logical unit) for SNA connection(s). The requestor associates a given location name with a list of LU names at the requestor location called member-routing which defines a list of LUs for the data sharing group members. Anothre means is to use VTAM supplied "generic resources" which allows a generic LU name for the group serving as a single LU Name for the entire group (Sysplex Distributor).

For TCPIP, its supported only via DRDA and all members of the data sharing group

use the same DRDA port number to receive incoming SQL requests, although each member must have a resync port # unique within a sysplex so that if a failure occurs this unique number allows a client to reconnect to the correct member to resolve any inflight work (2PC). Also under TCPIP endusers connect to a data sharing group using a domain name and they'll receive a list of eligible members in the data shaing group and can connect to any eligible member.

NOTE – Yo can also use mixed network protocols (SNA & TCPIP). When sending a request DB2 will use the LINKNAME in SYSIBM.LOCATIONS table of the DB2 catalog to determine which protocol to use.

Keep your pitch simple. If you depict too much you will open more questions from the audience than anticipated unleashing the devil and requiring more details. Data Sharing is a subject in itself especially in light of remote networking. I'm not confident these two slide should be used or maybe should be redrawn for better clarity.