

## What I Learned This Month: DB2 Lock Duplexing Overhead

Scott Chapman

American Electric Power

Fair warning: if you aren't familiar with mainframe parallel sysplex technology and concepts, this column isn't for you. Check back next month when I may have a more general interest topic.

If you're still reading, you're probably familiar with the potential dual failure issue of losing both a Coupling Facility (CF) and z/OS LPAR that reside on the same physical machine. Mainframes rarely fail completely, but it's still possible for a variety of reasons. (Nobody has ever run the dual power feeds into the same PDU or UPS, have they?) Because the simultaneous loss of both a CF and a z/OS member of the same sysplex may result in loss of data that was stored in the CF, this situation has to be carefully considered. And if you're running ICFs (Internal Coupling Facilities), you undoubtedly have considered it.

One popular answer, and the one we chose when we started using ICFs in 2006, is to enable duplexing for the sensitive structures. This adds overhead because every CF update has to be done in two CFs, but it ensures that the loss of a single machine won't mean the loss of important CF data. At the time, we simply turned on duplexing for all the structures that IBM had listed as being important for availability purposes. I looked at the overhead, but it was deemed a cost of doing business—part of the cost of using the cheaper ICF solution instead of having a pair of stand-alone CFs.

While at CMG'11 I heard somebody (it might have been Robert Catterall) mention in passing that you didn't need to duplex DB2 lock structures. This struck me as odd since the recommendation for availability was to duplex the lock structures when there was a risk of dual failure. We'd been doing that for 5 years. But things change over time: DB2 has changed, CF code has changed, z/OS has changed, and the hardware has changed. So I filed the comment away as something interesting to follow up on. In January, I was evaluating mainframe performance reporting tools, two of which produced sysplex overhead reports. The numbers reported by those tools were fairly significant in some intervals. So I pulled out some of my old code that attempted to calculate the overhead and ran some of my reports as well, and my code also showed some potentially significant overhead. That, combined with my recollection of the comment from CMG, led me to search out some more information.

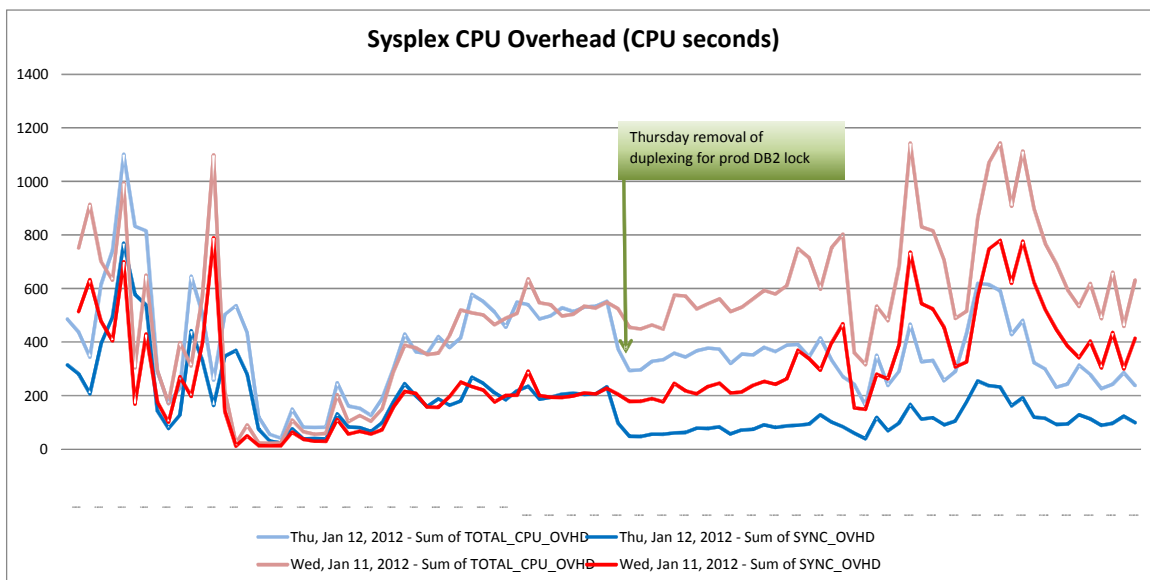
I came across a presentation<sup>1</sup> by Robert Catterall wherein he noted that today many organizations are not worrying about the dual failure scenario for DB2 lock and SCA structures because the overhead is so high and the probability of a failure of a physical system is so low. Also, if such a failure does occur, the DB2 group will fail, but can be restarted in a matter of minutes without the loss of any committed data. Recovery times have improved with newer versions of DB2, so

---

<sup>1</sup> See [http://www.mdug.org/Presentations/DB2\\_data\\_sharing\\_then\\_and\\_now.pdf](http://www.mdug.org/Presentations/DB2_data_sharing_then_and_now.pdf)

the total outage time at risk is less than it was in years past. I consulted with our DB2 Systems Programmer and we came to the conclusion that we should likely disable duplexing for the DB2 lock and SCA structures. In the highly unlikely event that we suffer the failure of an entire machine, the few extra minutes that it will take DB2 to restart on the surviving member is likely insignificant compared to the performance challenges we'll face from losing half our capacity with the machine that failed. Not that we couldn't fix the capacity situation with Capacity on Demand, but that too would take some time to get resolved.

As you can see below, the reduction in sysplex overhead was significant. The blue lines are the day I turned off duplexing, the red lines the day before. You can see before noon, the days tracked very closely to each other, and then when I made the change they diverged significantly.



I'm sure that my overhead calculations are not exactly correct. The lower lines that are brighter represent a minimum overhead calculation, while the lighter lines that are farther up the chart represent a potential total overhead figure. Regardless of the exact accuracy of the values, the relative change is obvious. In some intervals, at the high end, the change represents about 5% of our total installed capacity. That's a significant capacity savings! Because this is system overhead, this is the best kind of capacity savings to find—capacity that can be immediately put to use satisfying any workload that needs it.

Because we have more than sufficient latent demand to make up that 5%, total utilization hasn't gone down. But we have noticed some small performance improvements that may be at least partially related to the overhead reduction.

The amount of overhead incurred is related to a multitude of factors, so your mileage will vary. If you're duplexing your DB2 lock structures, you might want to take a closer look at whether it still makes sense today.

As always, if you think I've got it all wrong, please email me at [sachapman@aep.com](mailto:sachapman@aep.com) and let me know!