

Performance Improvements That Just Made Sense

We at Phoenix Software International are doing everything possible to make JES3^{plus} the obvious choice for existing IBM JES3 customers. It's the easiest path. It's the least risky path. It's a path that incurs no additional expense. And based upon results shown below from benchmark testing of enhancements we have made; it just makes sense.

What did we observe?

- Increased overall throughput and improved MVS dispatching efficiency
- More jobs completing with dramatically reduced channel utilization (a 92% reduction!)

What enhancements have we made?

- Support for Newer Channel Programming Techniques:
 - Advanced CKD to leverage modern RAID DASD I/O optimizations when zHPF is not available
 - High Performance FICON (zHPF) for significantly higher I/O start and data throughput rates
- Support for Parallel Access Volumes (PAV – V1R1 Only):
 - Multiple I/O launched for a single SPOOL extent instead of always chaining multiple requests
 - Customers control the level of parallelism vs chaining via the new MAXPAV= specification on the OPTIONS statement
- Reduced Local Lock Contention (V1R1 Only – In Testing):
 - Allows multiple SPOOL I/O operations to complete simultaneously
 - Reduces the total number of I/O-related SRBs scheduled

Benchmark Tests

Head-to-head comparison of IBM JES3 to JES3 ^{plus}		Experimental "What If?" Tests JES3 ^{plus} to IBM JES2	
IBM JES3 Single LPAR	JES3 ^{plus} Single LPAR	JES3 ^{plus} Three LPARs	IBM JES2 Three LPARs
1 ECKD I/O per volume (3 total) Single-threaded completion SRB	8 zHPF I/O per volume (24 total)	8x3=24 zHPF I/O per volume (72 total)	CF checkpoint for maximum efficiency

Performance Improvements That Just Made Sense

Benchmark Analysis Summary

Head-to-head comparison of
IBM JES3 to JES3^{plus}

Experimental "What If?" Tests
JES3^{plus} to IBM JES2

Metrics	IBM JES3 Single LPAR	JES3 ^{plus} Single LPAR	JES3 ^{plus} Three LPARs	IBM JES2 Three LPARs
Actual Job Counts	2,321	4,106	5,351	4,860
Average System CPU (in Cores)	1.1	2.1	2.85	4.18
Average I/O Rate (per sec)	381	2,499	3,178	17,314
Average Channel Utilization	12.68%	.99%	1.34%	74.87%

In addition:

- Average job CPU was the same in IBM JES3 and JES3^{plus}
- MVS Dispatching efficiency was superior under JES3^{plus}

Benchmark System Resources and Methodology

- System environment
 - z15 8562-B03: three CPs, one ICF, one IFL, and one zIIP
 - DS8882F DASD connected via four FICON 16S+ channels
 - SPOOL consists of three 3390 Mod-27 volumes – each on its own LCU
 - Fifty JES-managed initiators dedicated to running the benchmark
- Four [IEBDG](#) jobs with nine steps each writing 200,000 arbitrary SPOOL records
 - The tenth step of the last job (IEBDG4) resubmits the member twice (eight more jobs)
- CPU constraints
 - Vertical capacity constraint eliminated by using IPL boost (similar to 8562-Z04)
 - Four cores represents a horizontal capacity constraint compared to a typical large production environment
- I/O subsystem constraints
 - All DASD volumes are in an incremental Flashcopy relationship, which simulates the overhead of a DS8K MetroMirror (synchronous replication) configuration¹

¹ Significantly higher I/O rates should be achievable with all copy services relationships disabled