



STG - Systems Performance

## Intelligent Data Placement for Solid State Drives

Leslie Sutton

Senior Technical Staff Member, Systems Performance

IBM

## Intelligent Data Placement for Solid State Drives

---

- A quick overview of Solid State Drives
- Why data placement matters
- How System z I/O metrics can help
- z/OS tooling and performance studies
- z/TPF performance proof of concept
- Comparing use of software methodologies with hardware autonomics
- Additional SSD resources

## What are Solid State Drives?

*New Tier-0 drives for high priority, time-sensitive applications*

- Not Magnetic or Optical
- Semiconductor
  - Electronically erasable persistent medium
  - No mechanical read/write interface
  - No moving parts
- Plug compatible with HDDs
  - Form factors: 1.8", 2.5", 3.5"
  - Interfaces: SAS, FC, SATA, PCI-e
- **Recommended use:**
  - **High-intensity Random I/O**



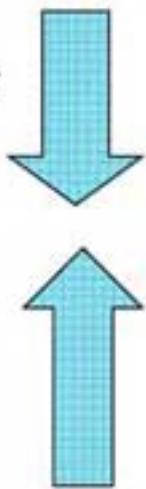
## SSDs compared to traditional spinning drives

### Advantages

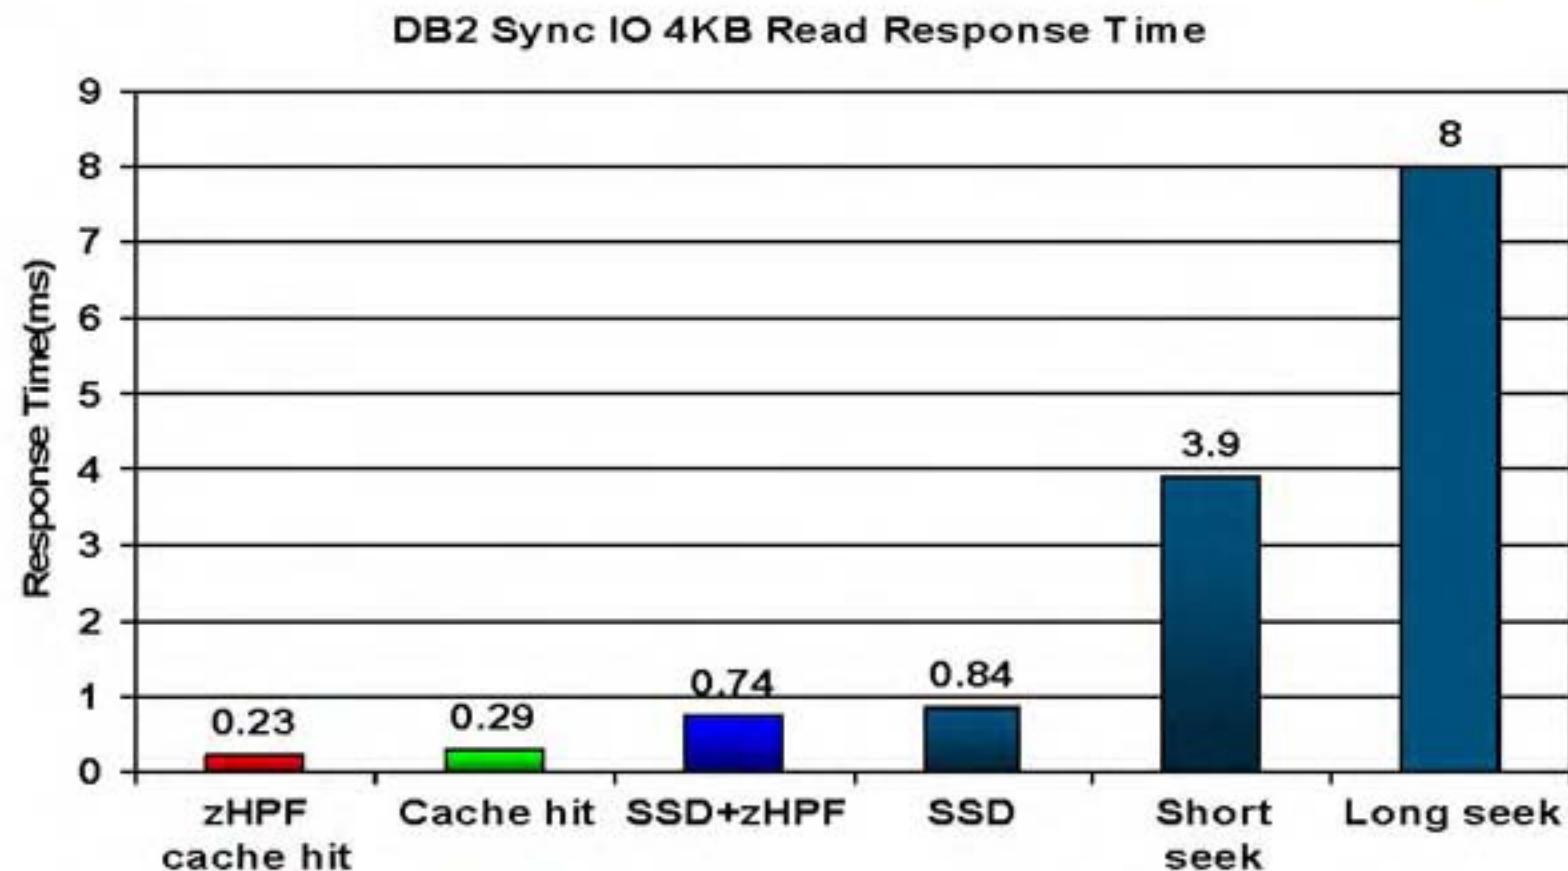
- Absence of mechanical moving parts
  - Significantly more reliable
- Green
  - Power - Fraction of energy consumed
  - Heat - Fraction of heat dissipated
- Performance
  - An order of magnitude improvement in IOPS as well as many times lower response time than 15K RPM spinning disks
  - Lower \$/IOPS than spinning disks

### Disadvantages

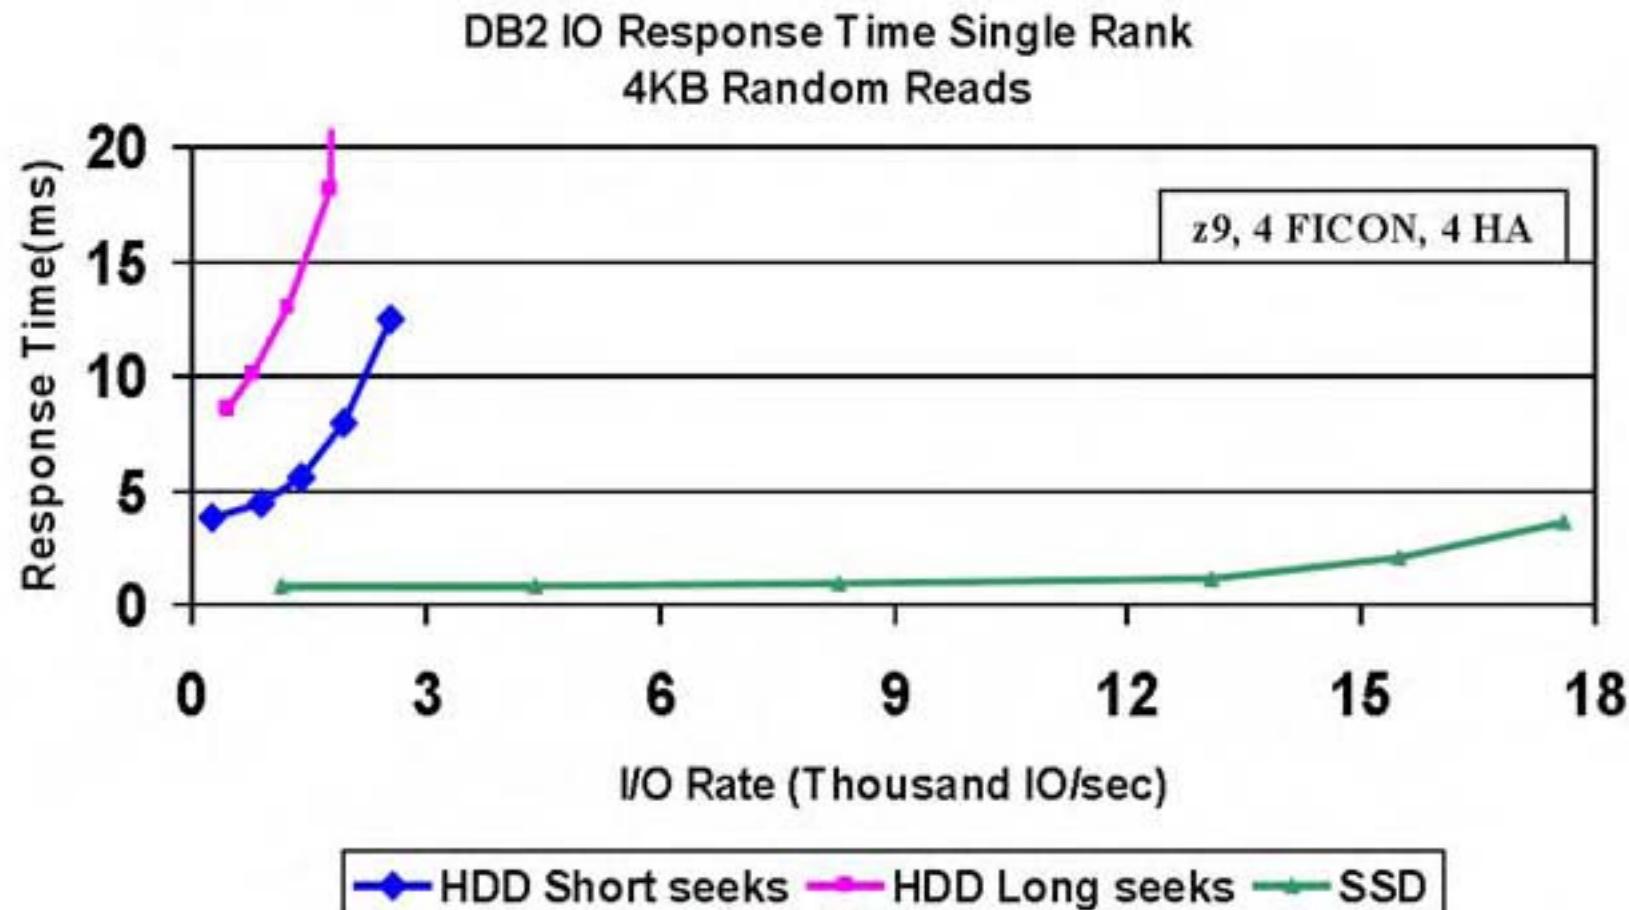
- Cost
  - About 10 - 30x more raw cost/GB today
  - Analysts foresee rapidly closing gap in pricing
- Capacity
  - Relatively small capacity per drive unit today compared to spinning disks
  - Industry moving to technologies that enable much higher unit capacities
- Wear issues
  - Flash cells wear out over time
  - SSD vendors have fixed this problem in enterprise class drives by providing overprovisioning and intelligent controllers with wear leveling.



## DB2 Sync (random) I/Os on DS8000 with z10 host

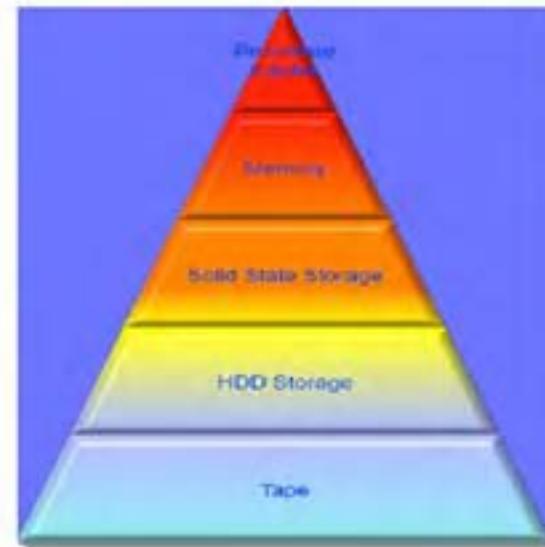


## HDD vs. SSD Response Time on DS8000 DB2IO Workload, one RAID-5 rank

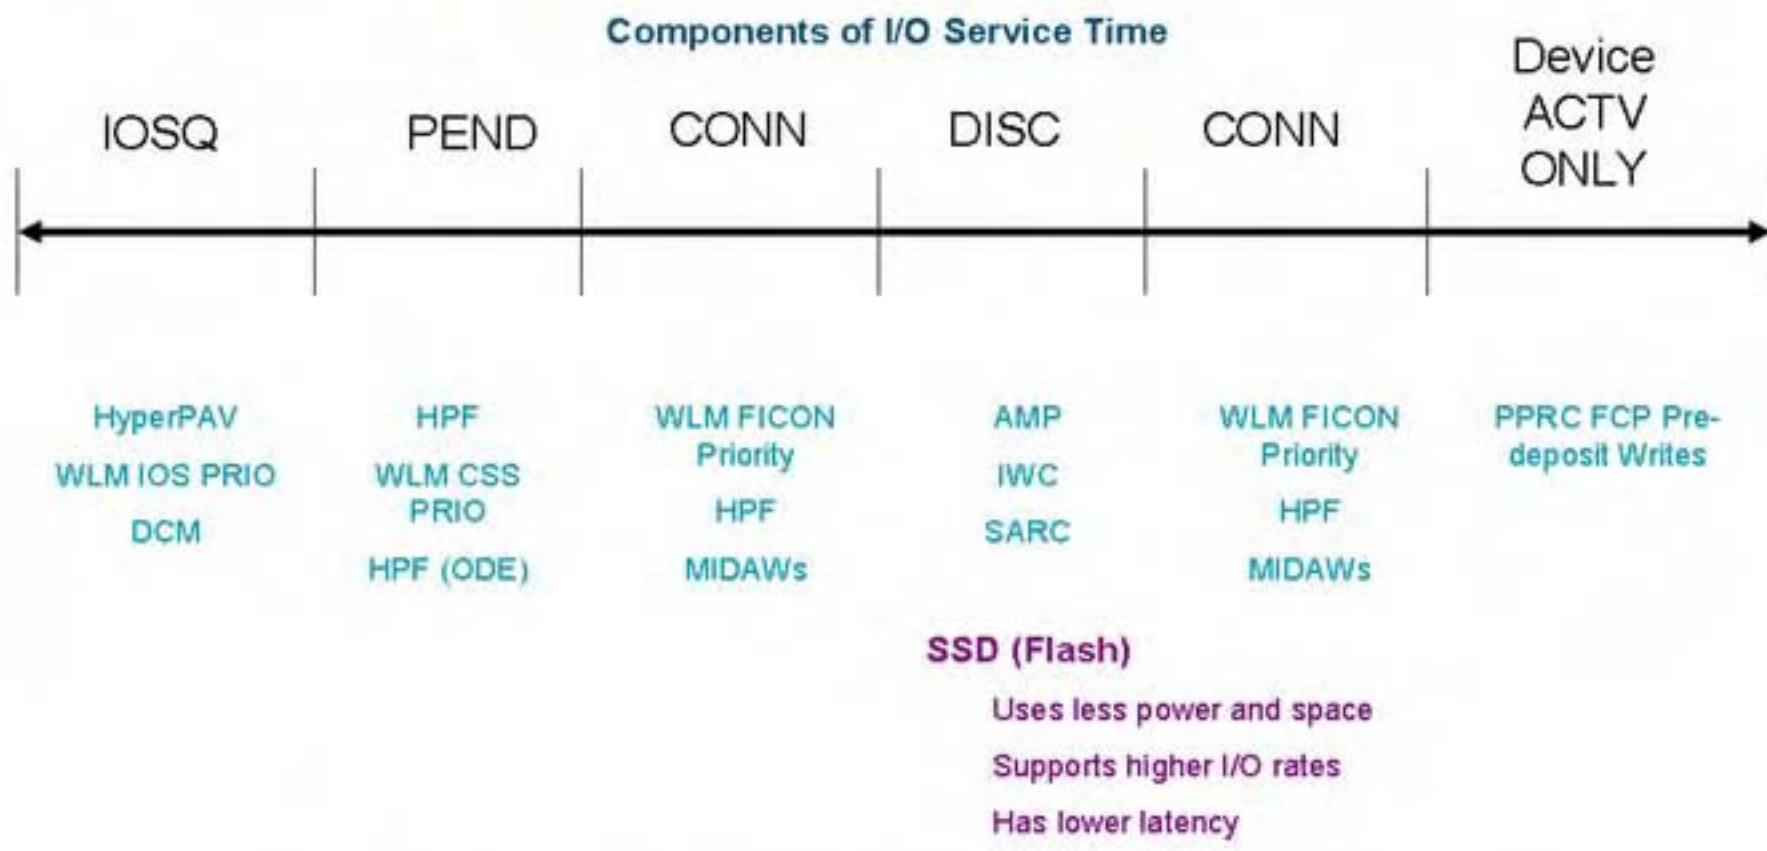


## Intelligent Data Placement for Solid State Drives

- Solid State Drive technology remains more expensive than traditional spinning disks, and the two technologies will coexist in hybrid configurations for several years
- During this transition, it's important to selectively migrate the data which can benefit most from the new technology.



## z/OS I/O Service Time Improvements



Another step in improving I/O performance for System z

## New Instrumentation for SSD

The advantages of SSD are greatest for random reads and writes.  
How can we identify the right data to move from HDD to SSD?

- Read-Only Disconnect Time
  - Disconnect time is accumulated on a device basis by the channel system as part of executing I/O operations
    - On a **READ** cache miss, includes the time required to stage data from disk into cache memory
    - On a **WRITE** operation, includes the time to propagate writes to the secondary DASD when using synchronous replication technologies
    - New instrumentation will separate write DISC time from read DISC time
- SMF type 42-6 record
  - z/OS accumulates channel measurement data by dataset
  - Read-Only Disconnect time is captured by dataset

General concept: mapping channel measurement data from each I/O to your logical view of the database

## New z/OS Tooling for SSD

---

- New Tooling based on SAS® software (FLASHDA) available as free download
  - Read-only disconnect time by dataset
    - Aggregate read DISC time for all data sets across all intervals from SMF data
    - Sort all datasets in descending order by total read DISC time
    - Datasets at the top (largest DISC time) will benefit the most from Solid State Drives
  - Device Cache Statistics to identify volumes that may be saturating RAID ranks
  - Datasets with highest write activity on these volumes are optionally displayed
  - Granular time periods to allow focus on market open, after lunch, batch window, etc...
  - Calculate "pseudo device load" to compare SSD vs. spinning
- Our SSD candidate selection algorithms have been adopted in additional tools from IBM Tivoli, Softek, and other vendors

leslie.zpet.demo.outr42 - Notepad

File Edit Format View Help

DEVICE NUMBER	VOLSER	DATA SET NAME	TOTAL DISCON.TIME (MILLISEC)	AVERAGE READ-ONLY DISCON.TIME (MILLISEC)	TOTAL DISCONNECT TIME (MILLISEC)	AVERAGE DISCONNECT TIME (MILLISEC)	AVERAGE RESPON TIME (MILLI)
831B	LB0101	DB2DBSG.DSNDBD.BOOK.TSTITLE.I0001.A001	597390.9	8.940	597390.9	8.940	13.09
5C06	SS0045	DMVSSPN.U2.ZFS.DATA	360126.4	3.836	365664.6	2.688	4.54
5B3A	NRLS21	RLSADSW.VF14D.BANKACCT.DATA	348377.5	9.521	346749.1	8.924	10.60
5925	PPRCR0	RLSADSW.VFA1D.BANKACCT.DATA	257148.3	2.923	267106.0	2.752	3.54
5B1A	DB3069	DB2DBWG.DSNDBD.CFCDB15.STAHIS15.I0001.A001	221873.8	9.788	221974.3	9.779	17.23
5B1E	DB3073	DB2DBWG.DSNDBD.CFCDB17.STAHIS17.I0001.A001	204584.3	6.591	204849.0	6.591	12.88
5B1B	DB3070	DB2DBWG.DSNDBD.CFCDB11.STAHIS11.I0001.A001	197327.9	5.308	197446.7	5.307	9.45
5F0E	NRLS14	RLSADSW.VF14D.HOTEL1.DATA	195790.9	6.838	194742.6	5.342	6.36
5E0F	NRLS19	RLSADSW.VF15D.BANKACCT.DATA	166745.5	5.755	164074.4	5.370	6.20
5D03	DB5D03	DB2DBWG.DSNDBD.CFCDB21.STAHIS21.I0001.A001	162632.8	7.219	162721.8	7.212	10.57
5F03	DB5F03	DB2DBWG.DSNDBD.CFCDB10.STAHIS10.I0001.A001	152836.2	7.281	152819.6	7.271	11.38
5C23	DB3084	DB2DBWG.DSNDBD.CFCDB20.STAHIS20.I0001.A001	147743.3	5.491	147860.4	5.486	8.99
5E12	DB3088	DB2DBWG.DSNDBD.CFCDB25.STAHIS25.I0001.A001	141717.4	5.057	141791.3	5.053	8.43
5F0E	NRLS14	RLSADSW.VF17D.HOTEL1.DATA	136608.9	5.693	137420.0	4.669	5.56
5A04	DB3077	DB2DBWG.DSNDBD.CFCDB18.STAHIS18.I0001.A001	128569.3	8.913	128517.1	8.890	13.03
5E13	DB1150	DB2DBWG.DSNDBD.CFCDB19.STAHIS19.I0001.A001	122697.3	3.761	122845.4	3.760	7.03
2E00	RLSA05	RLSADSW.VFA4D.DATAENDB.DATA	120246.8	9.407	122097.0	7.129	9.81
2E00	RLSA10	RLSADSW.VFA4D.CUSTOMER.DATA	112492.9	7.436	113983.9	5.505	7.83
2E00	RLSA10	RLSADSW.VFA2D.DATAENDB.DATA	106232.0	8.609	108186.6	6.261	8.95
5B23	DB3079	DB2DBWG.DSNDBD.CFCDB25.SUPHIS25.I0001.A001	104882.1	4.430	105145.8	4.430	6.57
2E00	RLSA05	RLSADSW.VFA3D.DATAENDB.DATA	100303.8	7.924	102726.2	5.567	7.89
5F06	DB5F06	DB2DBWG.DSNDBD.CFCDB20.SUPHIS20.I0001.A001	98206.3	3.535	98504.8	3.535	4.67
2E00	RLSA10	RLSADSW.VFA3D.CUSTOMER.DATA	93412.7	7.420	94616.8	5.444	7.82
2E00	RLSA10	RLSADSW.VFA2D.CUSTOMER.DATA	91395.8	7.435	92692.6	5.454	7.84
5B1D	DB3072	DB2DBWG.DSNDBD.CFCDB05.SUPHIS05.I0001.A001	85047.9	3.826	84731.8	3.804	4.97
2D80	RLSA0D	RLSADSW.VFA7D.DATAENDB.DATA	83812.0	8.476	85864.9	6.052	8.81
2D80	RLSA0D	RLSADSW.VFA9D.DATAENDB.DATA	83007.3	8.514	85204.4	6.197	9.00
5D03	DB5D03	DB2DBWG.DSNDBD.CFCDB08.SUPHIS08.I0001.A001	78871.7	2.659	79030.2	2.660	3.78
5905	DB0100	DB2DBWG.DSNDBD.CFCDB10.SUPHIS10.I0001.A001	76895.9	3.407	76288.8	3.371	4.49
5D03	DB5D03	DB2DBWG.DSNDBD.CFCDB12.SUPHIS12.I0001.A001	76754.0	2.604	76972.9	2.604	3.78
5E12	DB3088	DB2DBWG.DSNDBD.CFCDB13.SUPHIS13.I0001.A001	76367.7	3.046	76579.1	3.046	4.07
5A00	NRLS16	RLSADSW.VF15D.ITEMACT.DATA	76088.8	3.903	76588.4	3.456	4.67
5A04	DB3077	DB2DBWG.DSNDBD.CFCDB21.SUPHIS21.I0001.A001	75191.6	2.632	75349.6	2.632	3.67

leslie.zpet.demo.outr42 - Notepad

File Edit Format View Help

CT	AVERAGE RESPONSE TIME C)	AVERAGE I/O CONNECT TIME (MILLISEC)	AVERAGE I/O PENDING TIME (MILLISEC)	AVERAGE CONTROL UNIT QUEUE TIME (MILLISEC)	STORAGE CLASS	BLOCK SIZE	TOTAL READ I/O COUNT	TOTAL WRITE I/O COUNT
	13.099	3.730	0.128	0.000	DB2DAT	4096	66819	0
	4.543	1.280	0.128	0.000	SMSOE	4096	93880	42156
	10.603	1.039	0.128	0.000	DBNONR	512	36592	2264
	3.547	0.277	0.128	0.000	RLSA1	512	87969	9098
	17.239	6.990	0.147	0.000	DB2DAT	8192	22667	31
	12.885	5.827	0.184	0.000	DB2DAT	8192	31039	43
	9.452	3.600	0.158	0.000	DB2DAT	8192	37176	28
	6.366	0.428	0.128	0.000	DBNONR	4096	28632	7820
	6.204	0.256	0.128	0.000	DBNONR	512	28976	1576
	10.572	2.958	0.135	0.000	DB2DAT	8192	22527	37
	11.386	3.560	0.139	0.000	DB2DAT	8192	20992	25
	8.997	2.992	0.136	0.000	DB2DAT	8192	26906	44
	8.435	2.955	0.139	0.000	DB2DAT	8192	28024	36
	5.565	0.445	0.128	0.000	DBNONR	4096	23994	5436
	13.032	3.467	0.159	0.000	DB2DAT	8192	14425	31
	7.037	2.813	0.154	0.000	DB2DAT	8192	32627	42
	9.810	1.028	0.128	0.000	RLSA1	4096	12783	4343
	7.836	0.704	0.129	0.000	RLSA1	4096	15129	5575
	8.950	0.801	0.128	0.000	RLSA1	4096	12340	4940
	6.573	1.725	0.128	0.000	DB2DAT	16384	23677	57
	7.892	0.870	0.128	0.000	RLSA1	4096	12659	5795
	4.679	0.653	0.130	0.000	DB2DAT	16384	27782	80
	7.821	0.723	0.132	0.000	RLSA1	4096	12590	4789
	7.841	0.698	0.129	0.000	RLSA1	4096	12293	4701
	4.970	0.724	0.128	0.000	DB2DAT	16384	22229	44
	8.811	0.853	0.128	0.000	RLSA1	4096	9888	4299
	9.004	0.875	0.128	0.000	RLSA1	4096	9749	4000
	3.783	0.626	0.128	0.000	DB2DAT	16384	29659	57
	4.495	0.652	0.131	0.000	DB2DAT	16384	22573	55
	3.785	0.660	0.138	0.000	DB2DAT	16384	29474	82
	4.077	0.629	0.128	0.000	DB2DAT	16384	25071	68
	4.671	0.640	0.128	0.000	DBNONR	4096	19495	2666
	3.672	0.656	0.128	0.000	DB2DAT	16384	28567	62

Scrolling right  
in the report  
gives more  
data set level  
statistics

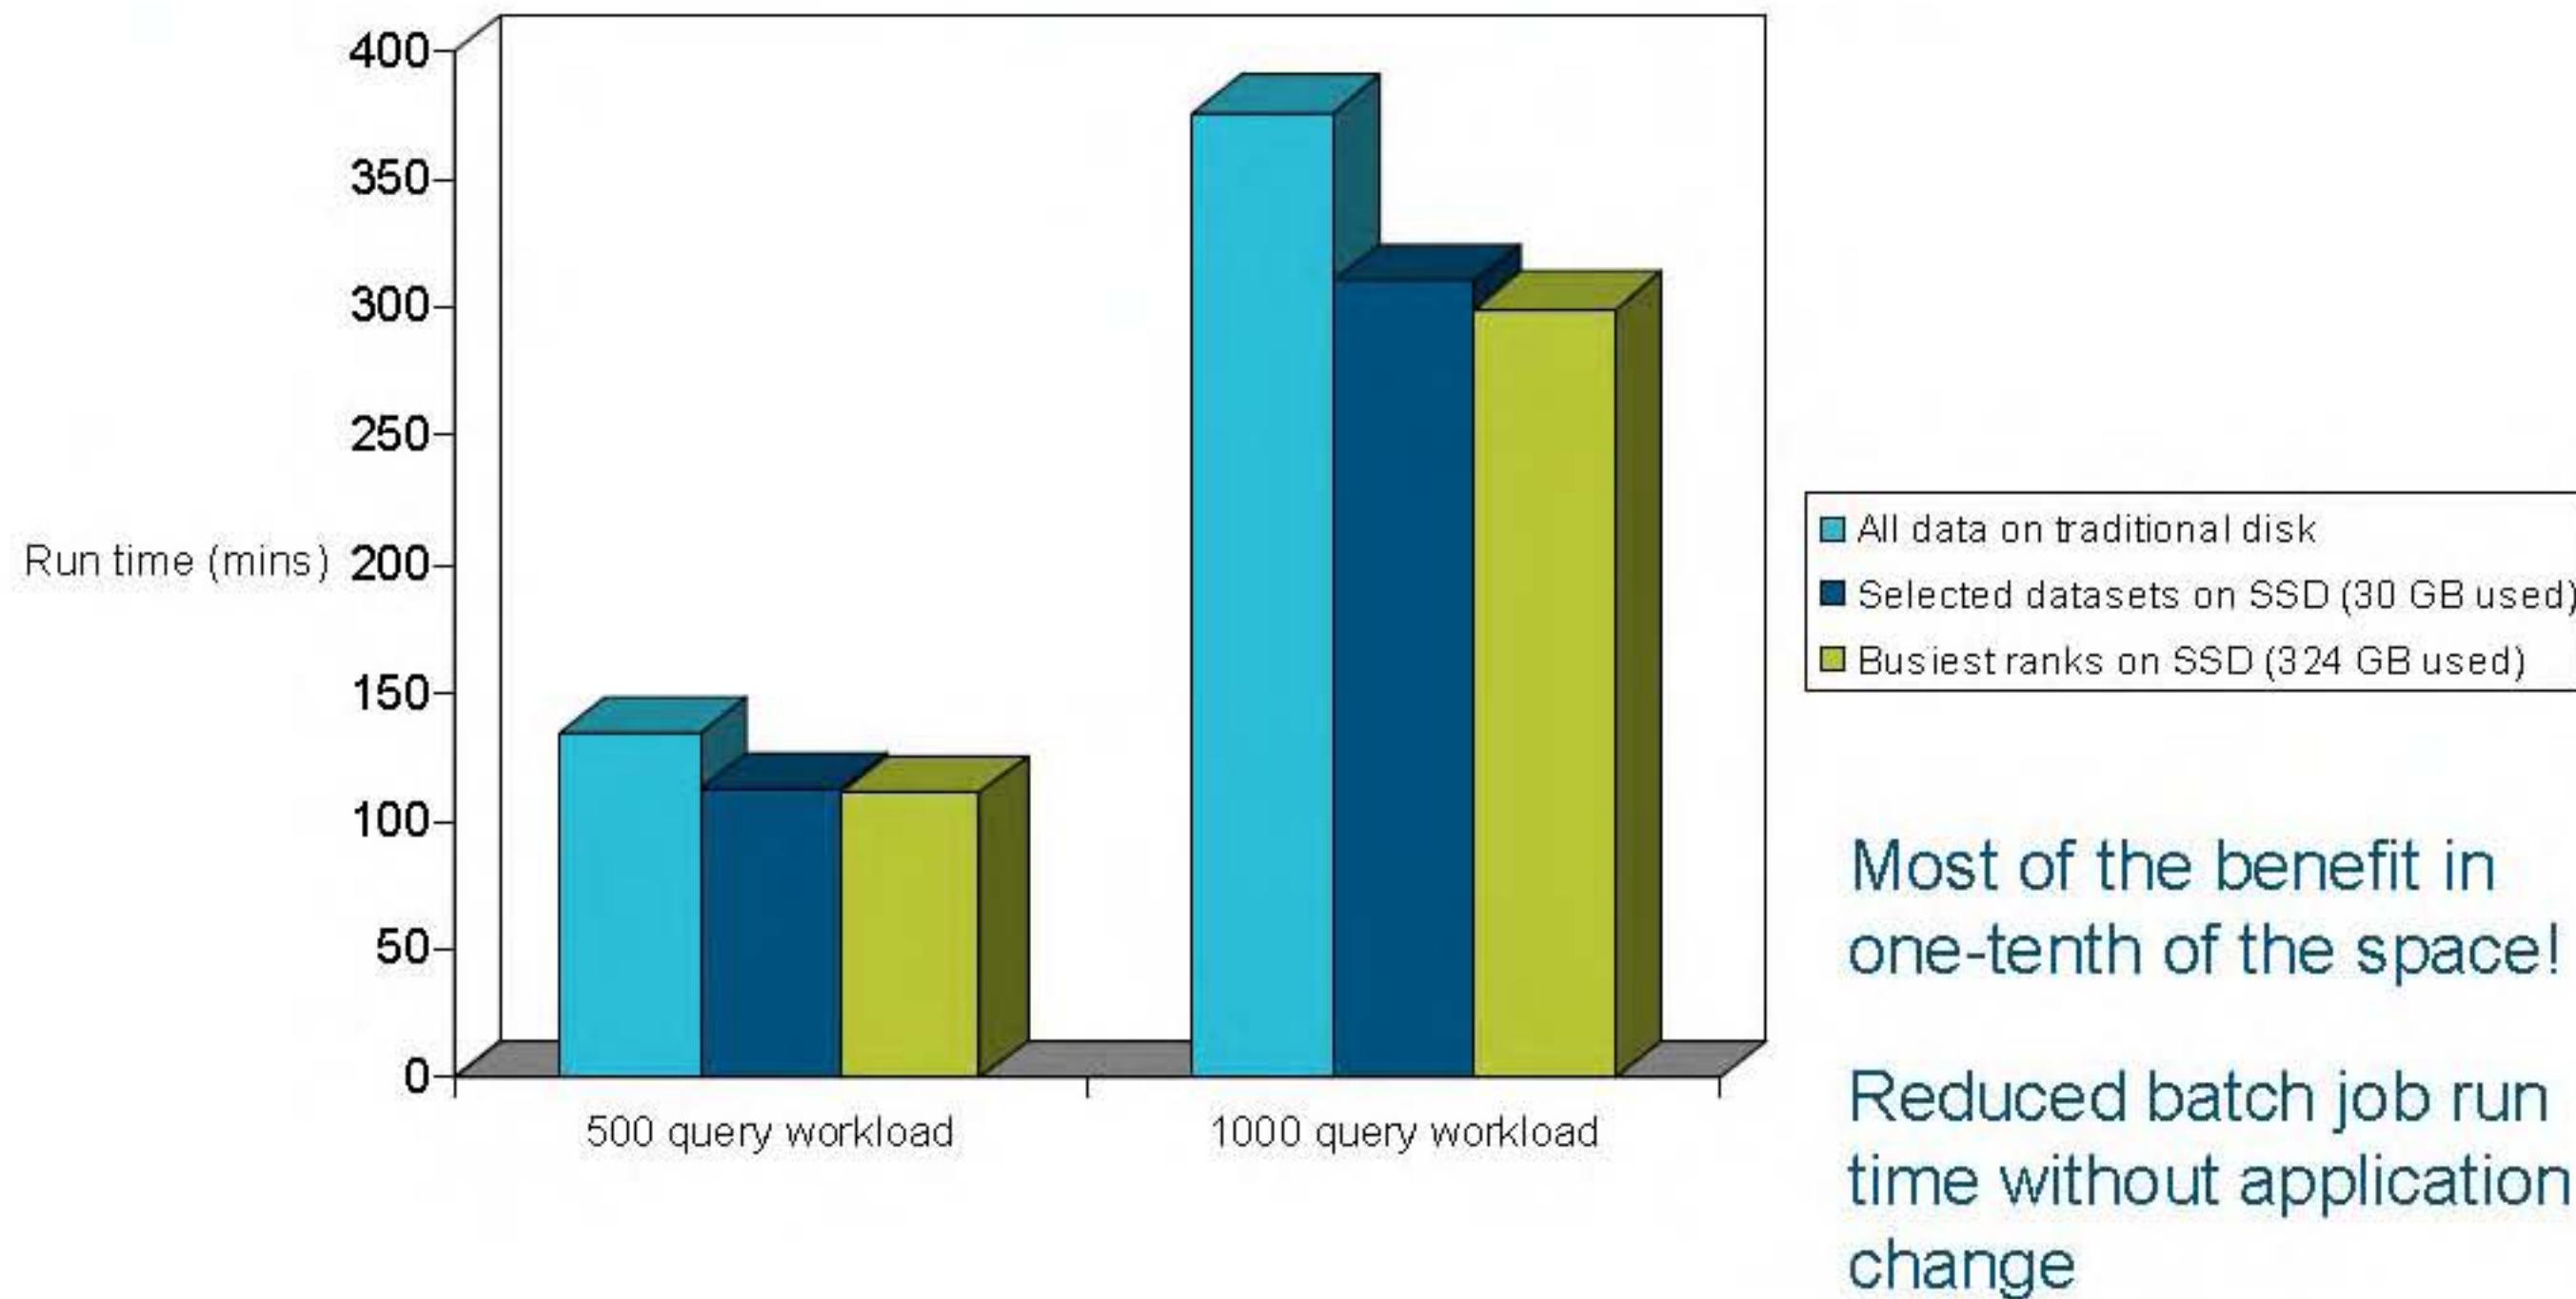
leslie.zpet.demo.outr74 - Notepad

File Edit Format View Help

DEVICE NUMBER	VOLSER	SPINNING PSEUDO DEVICE LOAD	SSD PSEUDO DEVICE LOAD	CACHE TO DEVICE	DASD TO DASD	DEVICE CACHE TRANSFERS	K BYTES WRITTEN	WRITE RESPONSE TIME (MILLISEC)	DEVICE MSECS PER K BYTE	DEVICE SEQUENTIAL REQUESTS	CACHE TO DASD/ WRITE RATIO	DEVICE VALIDITY TYPE	FLAG
3E0C	LOGR47	11.622	0.156	3493404	7983	166781184	4342720	0.026	878108	3.978	P	N	
2A80	DB2A80	6.780	0.071	1466404	129077	9534592	636496	0.067	40629	36.093	P	N	
3B01	LOGR44	6.379	0.084	1887390	4311	89294592	2234224	0.025	448203	4.211	P	N	
3B02	LOGR45	6.186	0.086	1921728	4405	91176576	2367776	0.026	524205	3.666	P	N	
5805	SS0079	6.112	0.073	1630112	17073	63692032	2070576	0.033	273989	5.950	O	Y	
3E21	LOGR46	5.621	0.122	2738106	4054	138920320	3958256	0.028	1466126	1.868	P	N	
831B	LB0101	5.491	0.055	189	1232815	0	0	0.000	0	0.000	P	N	
8432	XRST03	4.983	0.056	1253646	0	56536064	1468736	0.026	152679	8.211	P	N	
84CC	SS0076	4.713	0.073	1643076	549	73479680	2493568	0.034	585318	2.807	P	N	
880E	DB6359	4.414	0.047	999990	47925	7472000	381328	0.051	56835	17.595	P	N	
880C	DB6353	4.030	0.040	859755	45405	3536640	264160	0.075	288	2985.260	P	N	
8901	DB6356	3.657	0.039	844784	28026	7982976	454160	0.057	51759	16.321	P	N	
5D03	DB5D03	3.560	0.056	710943	538092	31081856	1028400	0.033	201869	3.522	O	Y	
5C06	SS0045	3.395	0.055	711451	525677	27609600	1030272	0.037	237768	2.992	O	Y	
8807	DB6348	3.374	0.035	770897	22310	6517376	274272	0.042	35649	21.625	P	N	
8437	XRST08	3.313	0.035	789579	0	33755904	1172128	0.035	45653	17.295	P	N	
A60B	D83IAB	3.231	0.045	1000911	0	33004416	1097008	0.033	275439	3.634	P	Y	
8813	DB6361	3.104	0.031	677735	19476	2271360	168864	0.074	144	4706.493	P	N	
8802	DB6343	2.705	0.028	602145	20945	3440512	191248	0.056	15714	38.319	P	N	
8433	XRST04	2.468	0.025	558135	0	22598144	623088	0.028	4059	137.506	P	N	
A60D	D83IAD	2.461	0.036	814761	2988	32534784	938592	0.029	257004	3.170	P	Y	
8436	XRST07	2.452	0.025	555939	0	22878336	629760	0.028	5383	103.277	P	N	
A600	D83IA0	2.340	0.028	628578	0	30101760	592560	0.020	103014	6.102	P	Y	
A609	D83IA9	2.269	0.028	629073	0	30594816	599184	0.020	119628	5.259	P	Y	
8902	DB6357	2.266	0.023	507423	15202	3383168	229504	0.068	13779	36.826	P	N	
5F12	LOGR25	2.206	0.074	1668383	855	78941184	2455968	0.031	1172733	1.423	O	Y	
5A17	LOGR23	2.041	0.126	2839802	558	135325952	4338656	0.032	2381499	1.192	O	Y	
5D0E	SS0046	2.005	0.021	440104	38205	11482624	522976	0.046	27468	16.022	R	Y	
8806	DB6347	1.859	0.020	425124	13128	3402240	149840	0.044	17865	23.796	P	N	
3A12	DB6007	1.705	0.043	958491	7281	74686976	2281616						
830F	IMS046	1.610	0.016	273600	87987	744192	58240						
3A00	DB3A00	1.382	0.033	739269	5184	57729152	1813072						
3A16	DB6011	1.278	0.029	638145	22212	47133824	1613184						

Additional reports merge  
volume and data set views

## Elapsed time savings with selective use of SSD for 8-way multistreaming batch query workloads



## Mixed workload example in 9-way sysplex

DEVICE	DATA SET	TOTAL READ-ONLY DISC.TIME	AVERAGE READ-ONLY DISC.TIME
2A00	SS0006 OMVSSPN.EAVZFS1.ZFS.DATA		
2A00	SS0006 OMVSSPN.EAVZFS2.ZFS.DATA		
5C06	SS0045 OMVSSPN.U2.ZFS.DATA		
5903	SS0078 OMVSSPN.SQUADZFS.ZFS.DATA		
8C00	MR0065 OMVSSPN.REEFZFS.ZFS.DATA		
5B3A	NRLS21 RLSADSW.VF14D.BANKACCT.DATA		
5E0F	NRLS19 RLSADSW.VF15D.BANKACCT.DATA		
5A00	NRLS16 RLSADSW.VF15D.ITEMACT.DATA		
5B1A	DB3069 DB2DBWG.DSNDBD.CFCDB15.STAHIS		
5925	PPRCR0 RLSADSW.VFA1D.BANKACCT.DATA		
5C23	DB3084 DB2DBWG.DSNDBD.CFCDB20.STAHIS		
5F03	DB5F03 DB2DBWG.DSNDBD.CFCDB10.STAHIS		
5A04	DB3077 DB2DBWG.DSNDBD.CFCDB18.STAHIS		
5B1E	DB3073 DB2DBWG.DSNDBD.CFCDB17.STAHIS		
5B1B	DB3070 DB2DBWG.DSNDBD.CFCDB11.STAHIS		
5A04	DB3077 DB2DBWG.DSNDBD.CFCDB21.SUPHIS		
5D03	DB5D03 DB2DBWG.DSNDBD.CFCDB21.STAHIS		
5A00	NRLS16 RLSADSW.VF15D.CUSTOMER.DATA	60057.6	12.800
5E12	DB3088 DB2DBWG.DSNDBD.CFCDB25.STAHIS25.I0001.A001	59973.5	26.643
5F0E	NRLS14 RLSADSW.VF17D.HOTEL1.DATA	54160.8	4.736
5F0E	NRLS14 RLSADSW.VF14D.HOTEL1.DATA	51604.2	5.376
2E00	RLSA05 RLSADSW.VFA4D.DATAENDB.DATA	49396.2	10.093
83A5	SS0075 OMVSSPN.SQUADZFS.ZFS.DATA	46428.6	0.384
5B1A	DB3069 DB2DBWG.DSNDBD.CFCDB15.STATUS15.I0001.A001	46218.7	1.217
5B23	DB3079 DB2DBWG.DSNDBD.CFCDB25.SUPHIS25.I0001.A001	45256.1	9.111
5B1D	DB3072 DB2DBWG.DSNDBD.CFCDB05.SUPHIS05.I0001.A001	44878.0	4.793
5C0D	DB3078 DB2DBWG.DSNDBD.CFCDB26.STATUS26.I0001.A001	43521.6	1.296

**z/OS File System workload**  
**3X I/O rate on SSD**

**VSAM RLS tran rate up 10%**  
**after moving 10 datasets**

**DB2 workload tran rate up 10%**  
**after moving 30 of almost**  
**50,000 datasets used**

## SSD for z/OS Mean Time To Recovery

- SYSRES on SSD to reduce IPL time
- Paging datasets on SSD to reduce dumping time
- Couple datasets on SSD to reduce restart time
- Observable improvement will vary for each customer's environment, but 10X reduction in random I/O response time can be helpful for these resources with long idle periods and very bursty I/O activity.

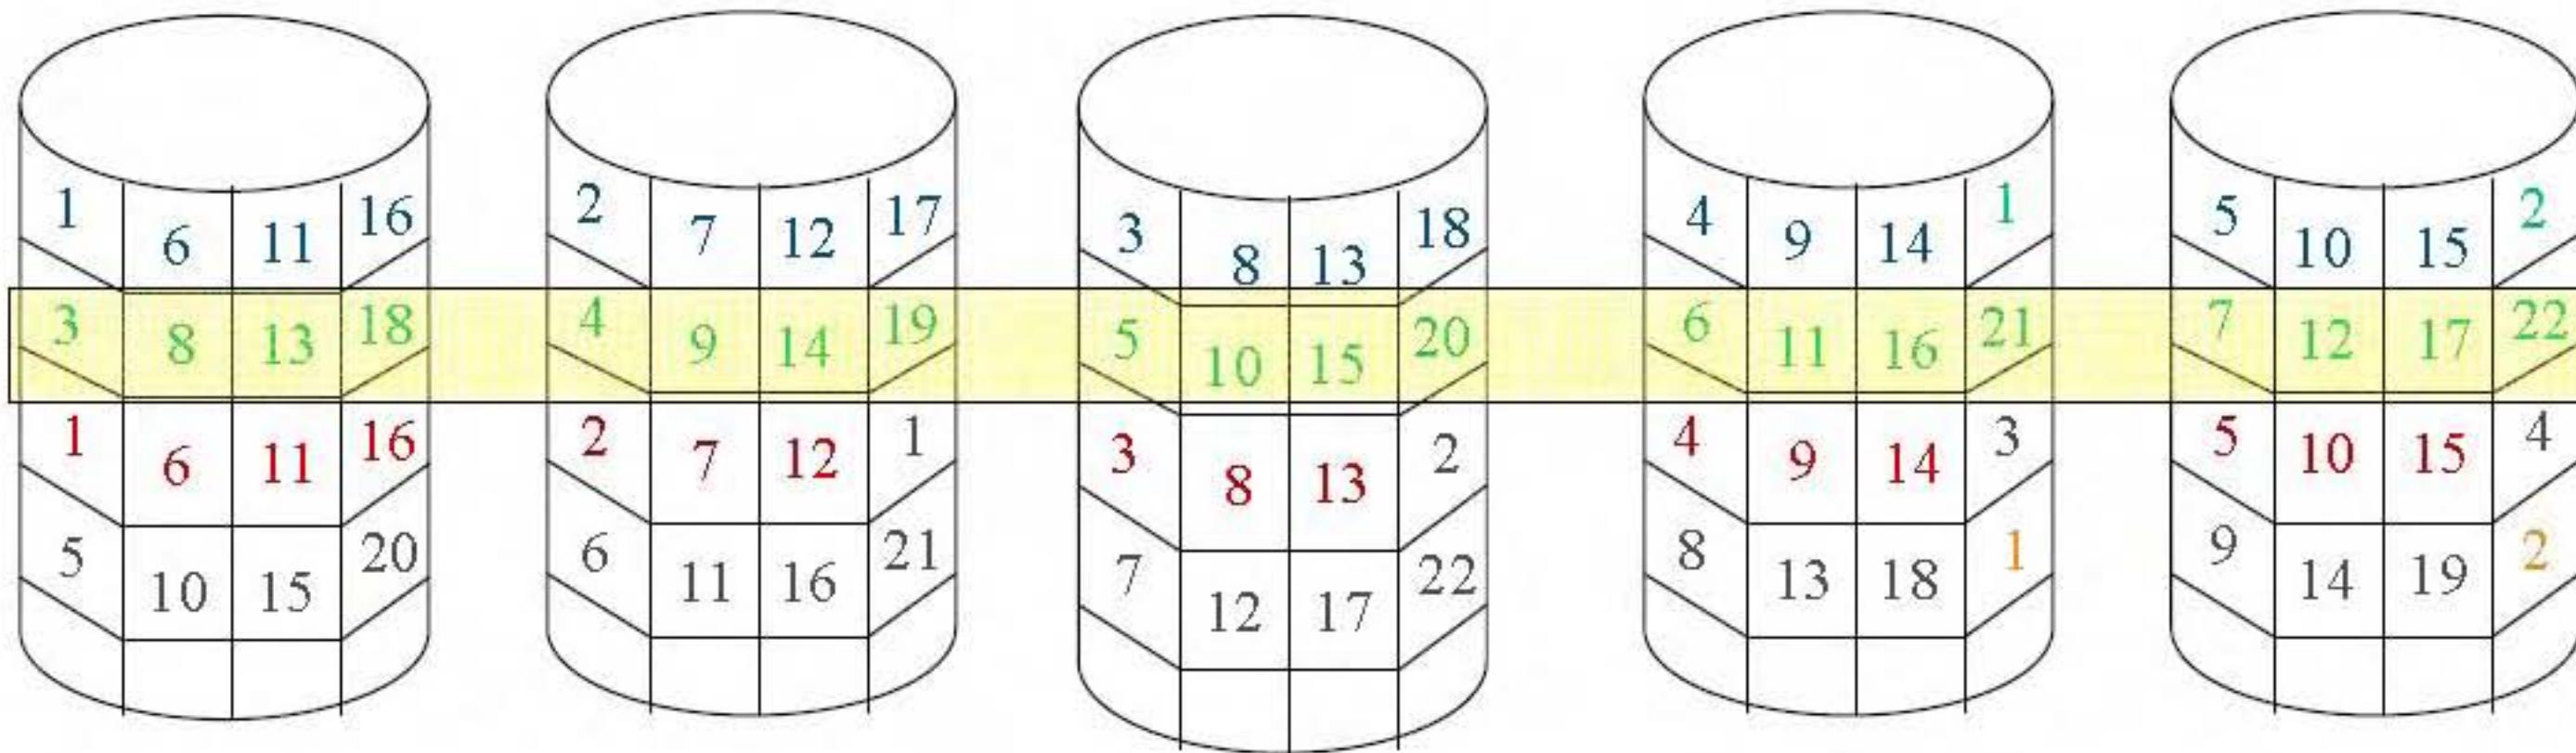
**General concept: think about your “cold” data that is critical during application restart/recovery processing**

## **z/TPF Intelligent Data Placement Study**

---

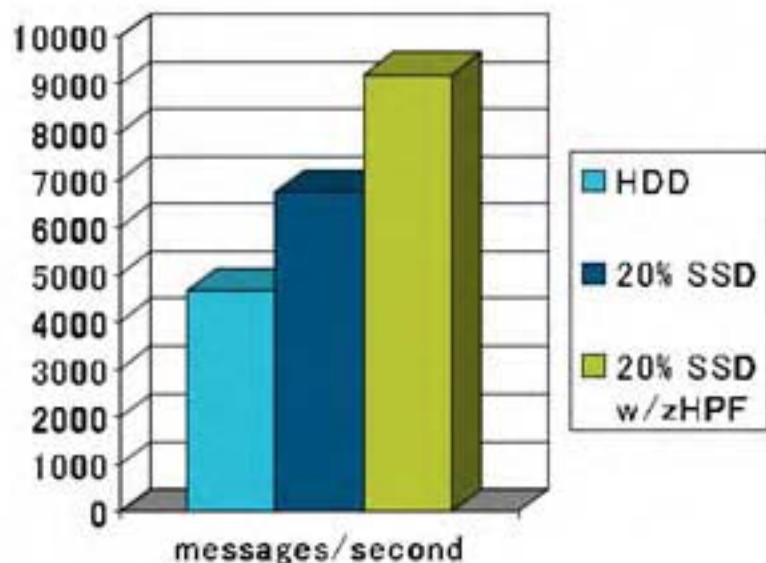
- Done in collaboration with z/TPF Development 4Q2009
- Focused on hybrid HDD/SSD configurations and intelligent data placement with customer-like workload
- Modified version of z/TPF AIR1 workload with localized hot spots for reads and writes
- “Before and After” study showing performance benefit of moving 10-20% of the database to SSD
  - Move data based on analysis of hot TPF “record IDs”
  - Use z/TPF zHPF prototype to show value of High Performance FICON with SSD
  - Prove that TPF customers can get benefit from SSD without migrating their entire databases

## TPF Horizontal Record Allocation and “Hot Bands”



- Logical view of TPF database (primary volumes only)
- Striped and mirrored in software
- Different record ID/ordinal ranges have different performance characteristics allowing identification of “hot bands” we can place on SSD

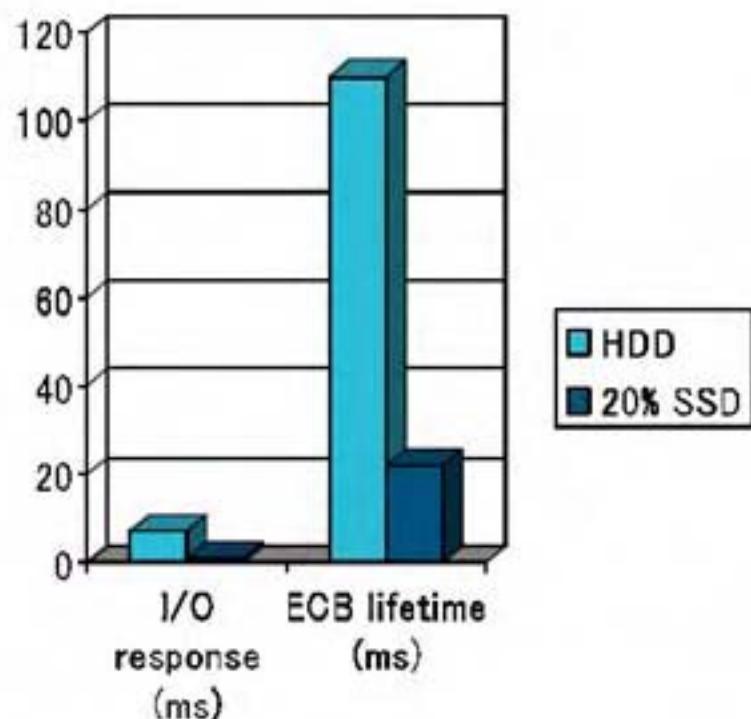
## AIR1 Workload – Maximum Throughput Study



- All database records are 4KB
- 85% read cache hit, 59% write hit
- Using 20% SSD alone gives a 45% increase in throughput
- Turning on zHPF raises this to a 98% increase in throughput
- I/O response time and ECB lifetime also improved 5X in 20% SSD case

- Baseline = 100% HDD
- 20% SSD = 20% of I/O directed to SSD
- 20% SSD w/ zHPF = 20% of I/O directed to SSD and z/TPF prototype code for zHPF support turned on

## AIR1 Workload – Constant Throughput Study



- Baseline = 100% HDD
- 20% SSD = 20% of I/O directed to SSD

- All database records are 4KB
- 85% read cache hit, 59% write hit
- Using 20% SSD gives 80% reduction in ECB lifetime which drives a corresponding reduction in working storage requirement
- Significant potential benefit to TPF 4.1 customers who are growing the number of TPF images in their loosely-coupled complexes today to compensate for working storage constraints until they can migrate to z/TPF.

## Software Methodologies vs. Hardware Autonomics

*One size does not fit all – choose what works best for you*

### Using host software metrics

- Requires skillful selection of SSD candidates, either through use of performance tooling or deep knowledge of applications
- Offers finest granularity of candidate selection and most opportunity for application of your business policy
- Needs to be revisited periodically as workload changes
- Data migration is initiated and managed by humans
- Compatible with using SSD for “cold” data critical for restart/recovery
- System-wide scope

### Using hardware autonomics

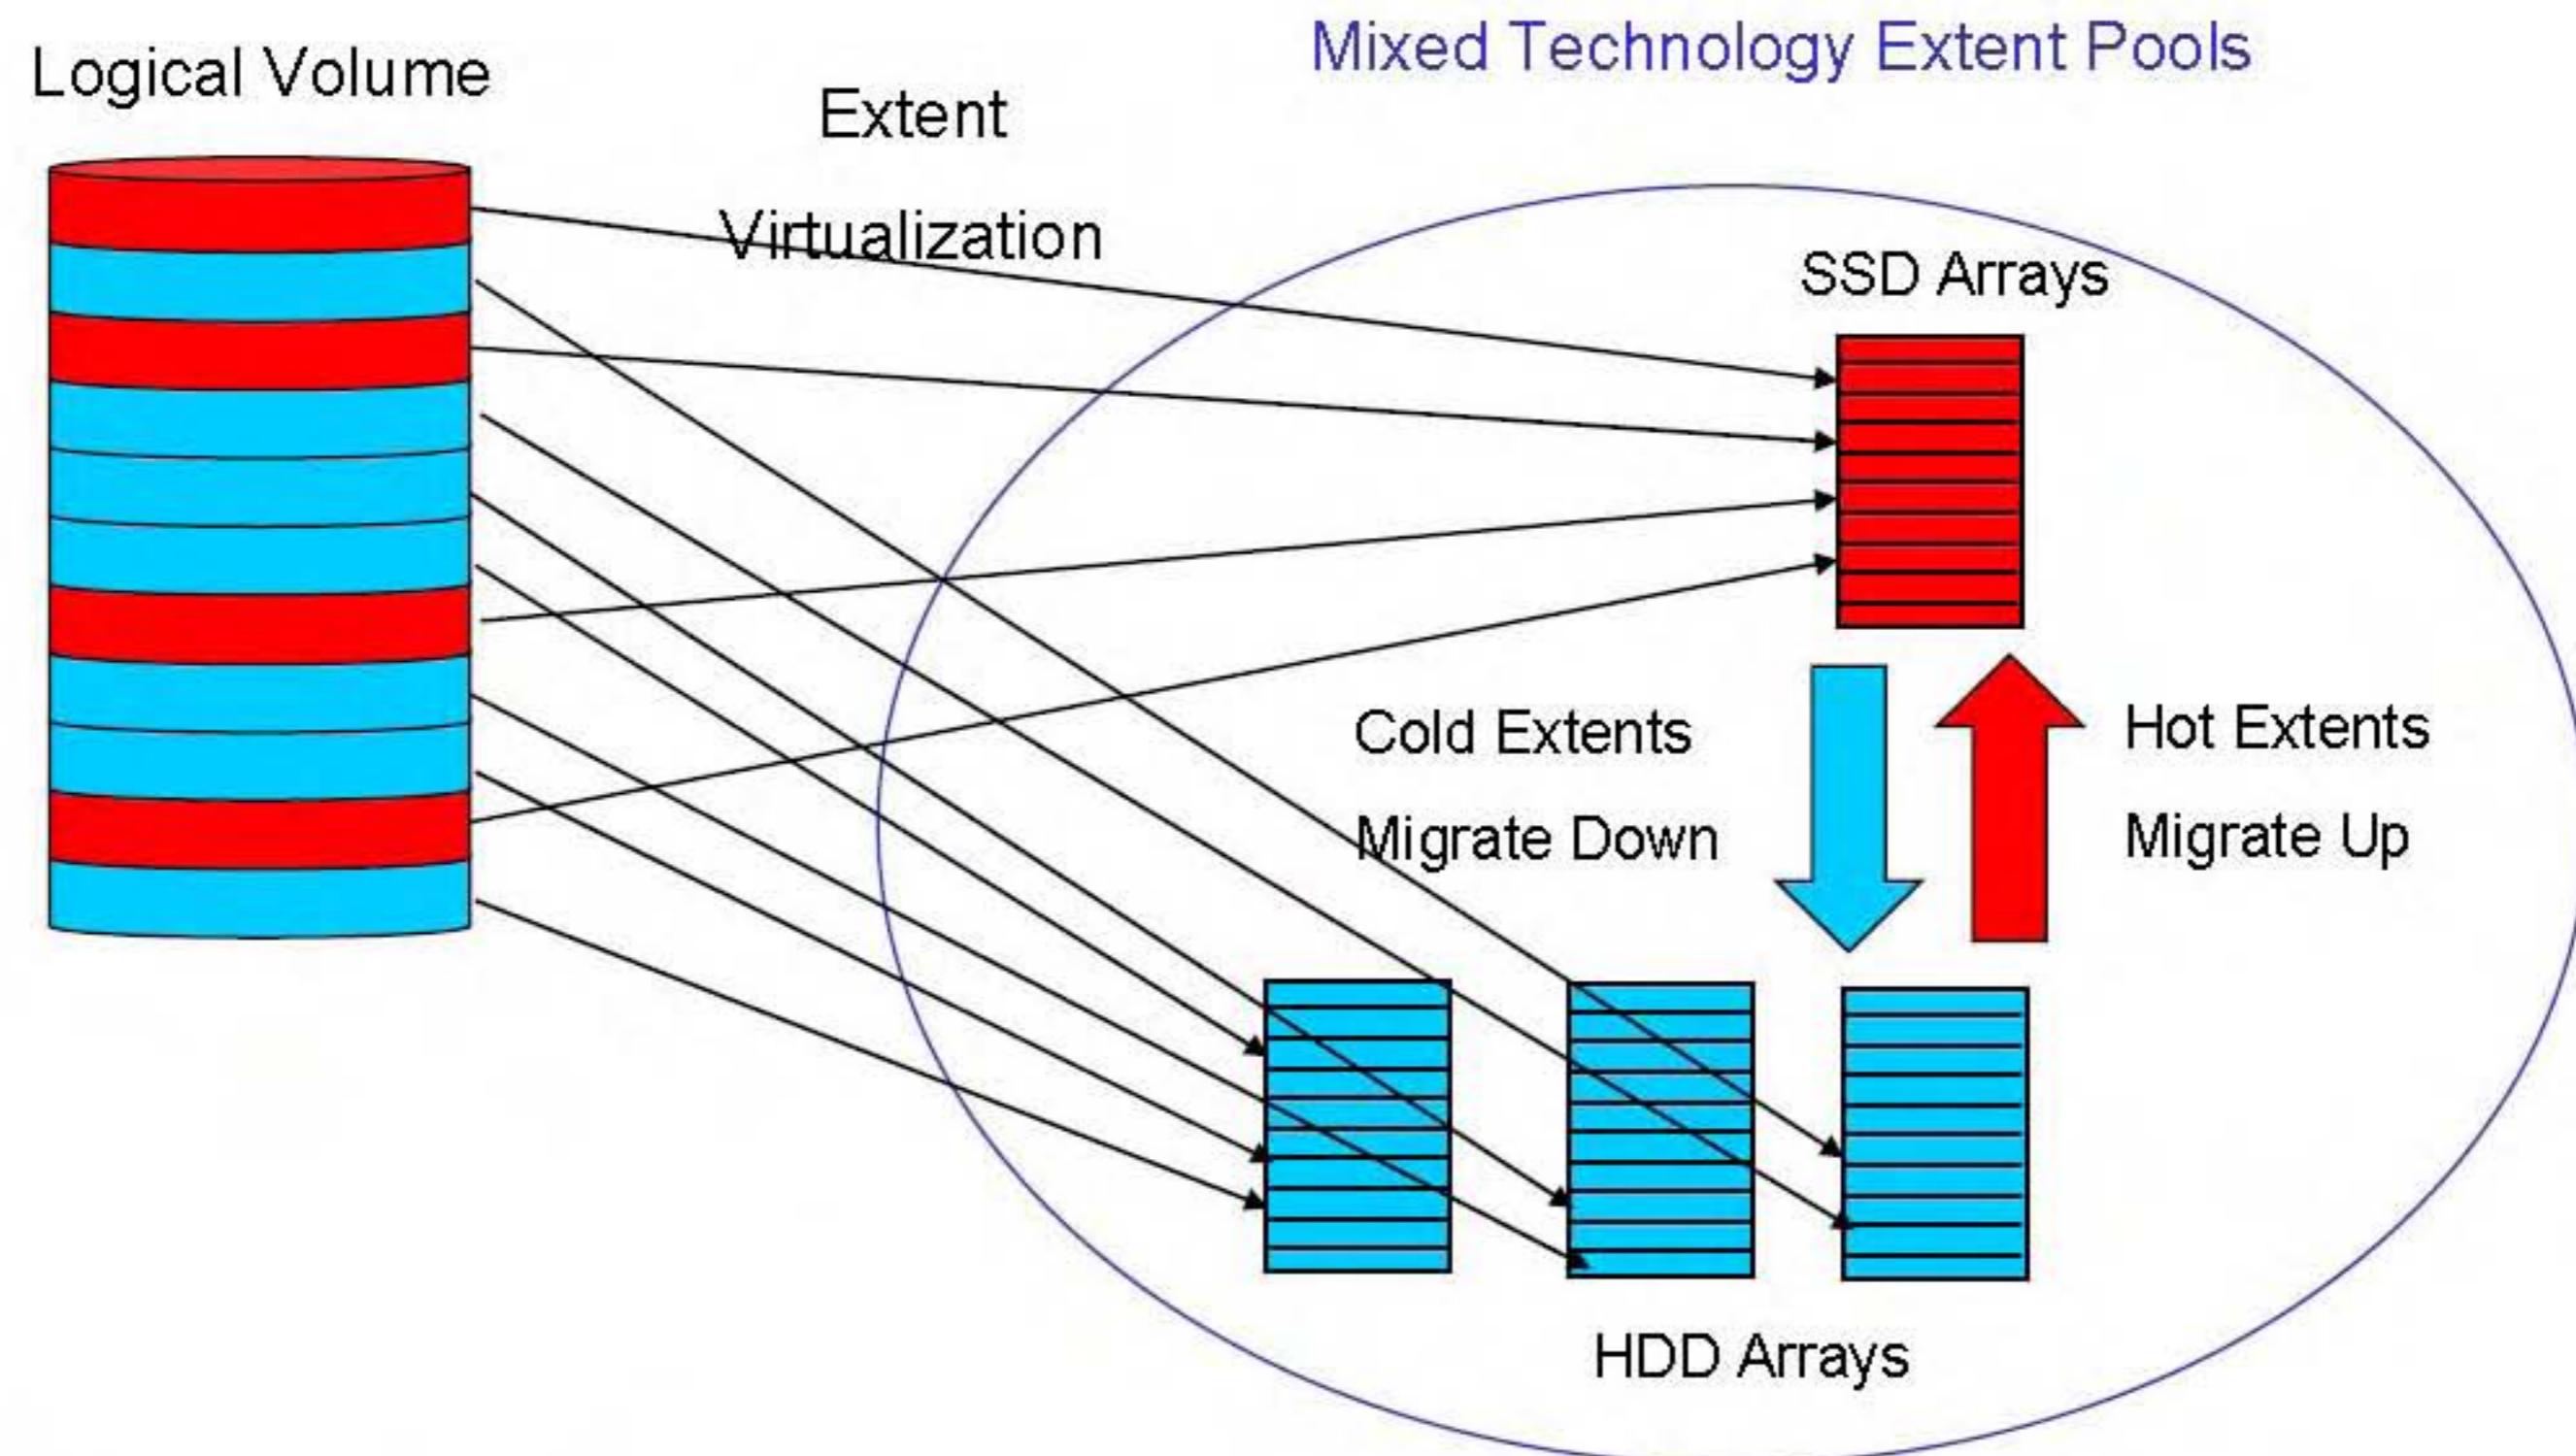
- Works from an internal view of storage performance metrics and has deep knowledge of resource utilizations not surfaced to software
- Selects SSD candidates at physical extent level (more granular than volume-based approaches, but much less granular than record-based approaches)
- Dynamically learns workload hotspots and automatically adjusts data placement over time in response to workload changes
- Limited integration of business policy today
- “Cold” data for restart/recovery would likely be demoted if placed in this pool
- Scope is a subset of a single storage system

## Advantages of Using Easy Tier Automatic Mode

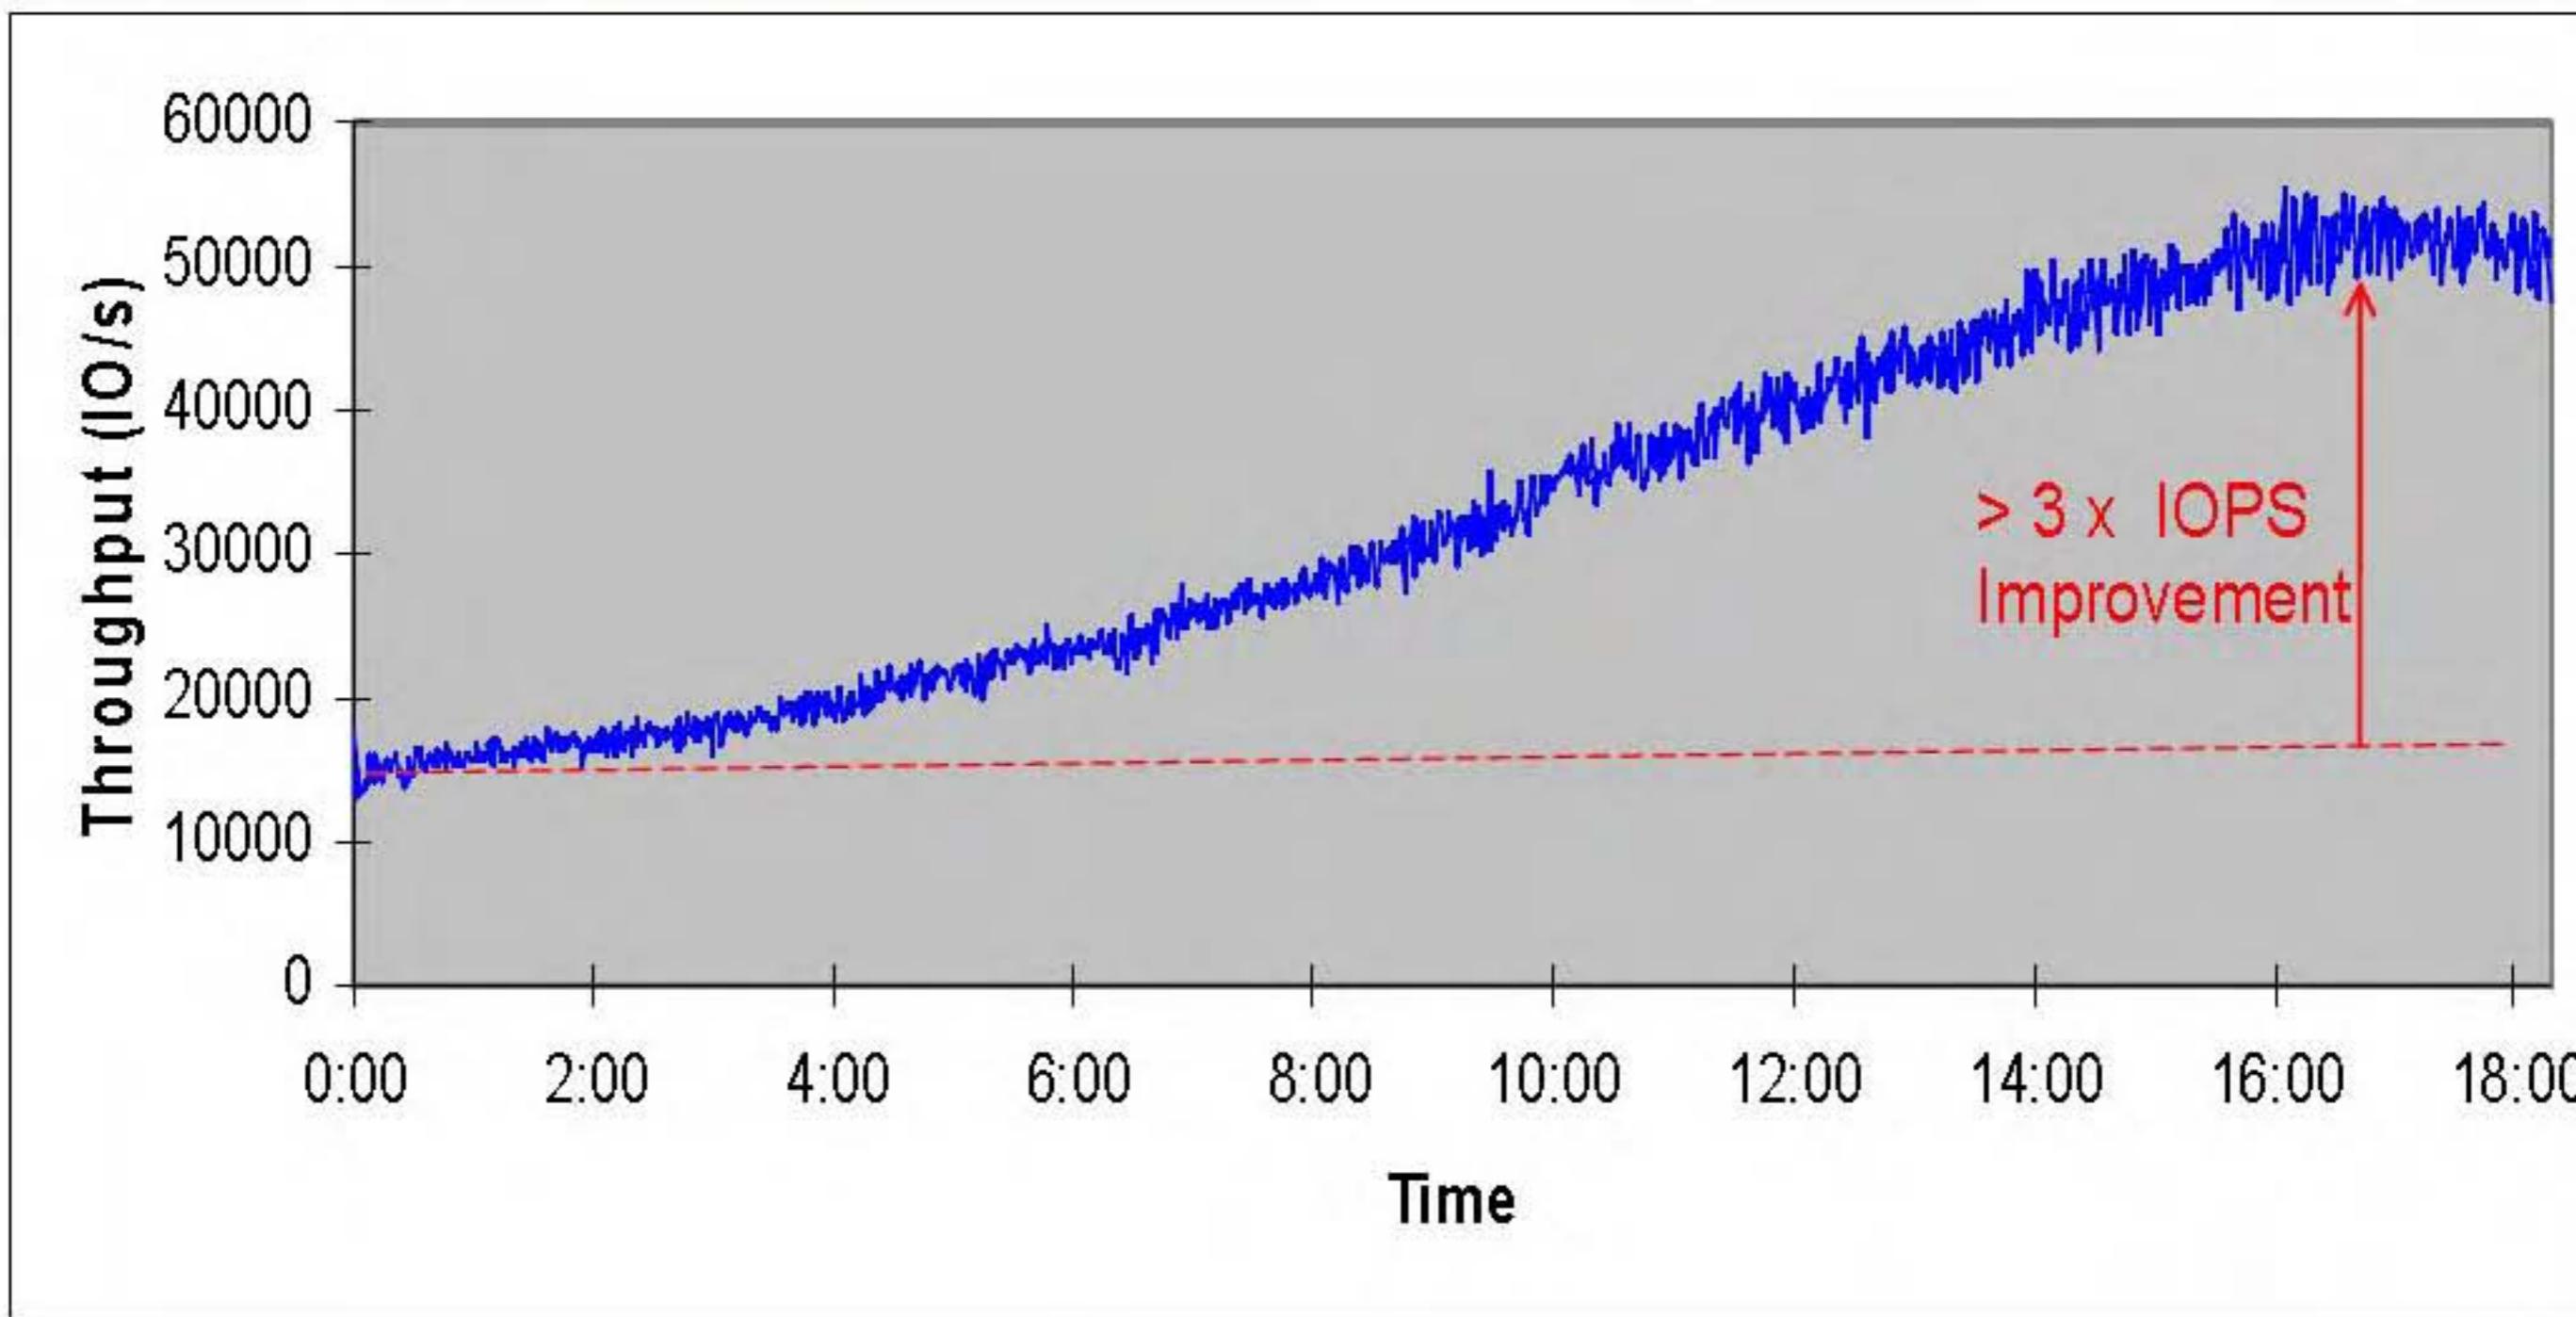
---

- **Designed to be Easy!**
  - The user is not required to make a lot of decisions or go through an extensive implementation process
- **Efficient Use of SSD Capacity**
  - Easy Tier moves 1 gigabyte data extents between storage tiers.
  - Less potential to waste SSD capacity than with volume based methods
- **Intelligence**
  - Easy Tier learns about the workload over a period of time (24 hours)
  - As workload patterns change, Easy Tier finds any new highly active ("hot") extents and exchanges them with extents residing on SSDs that may have "cooled off"
- **Negligible Performance Impact**
  - Easy Tier moves data gradually to avoid contention with host I/O activity
  - The overhead associated with Easy Tier management is nearly undetectable.
  - No need for storage administrators to worry about scheduling when migrations occur.

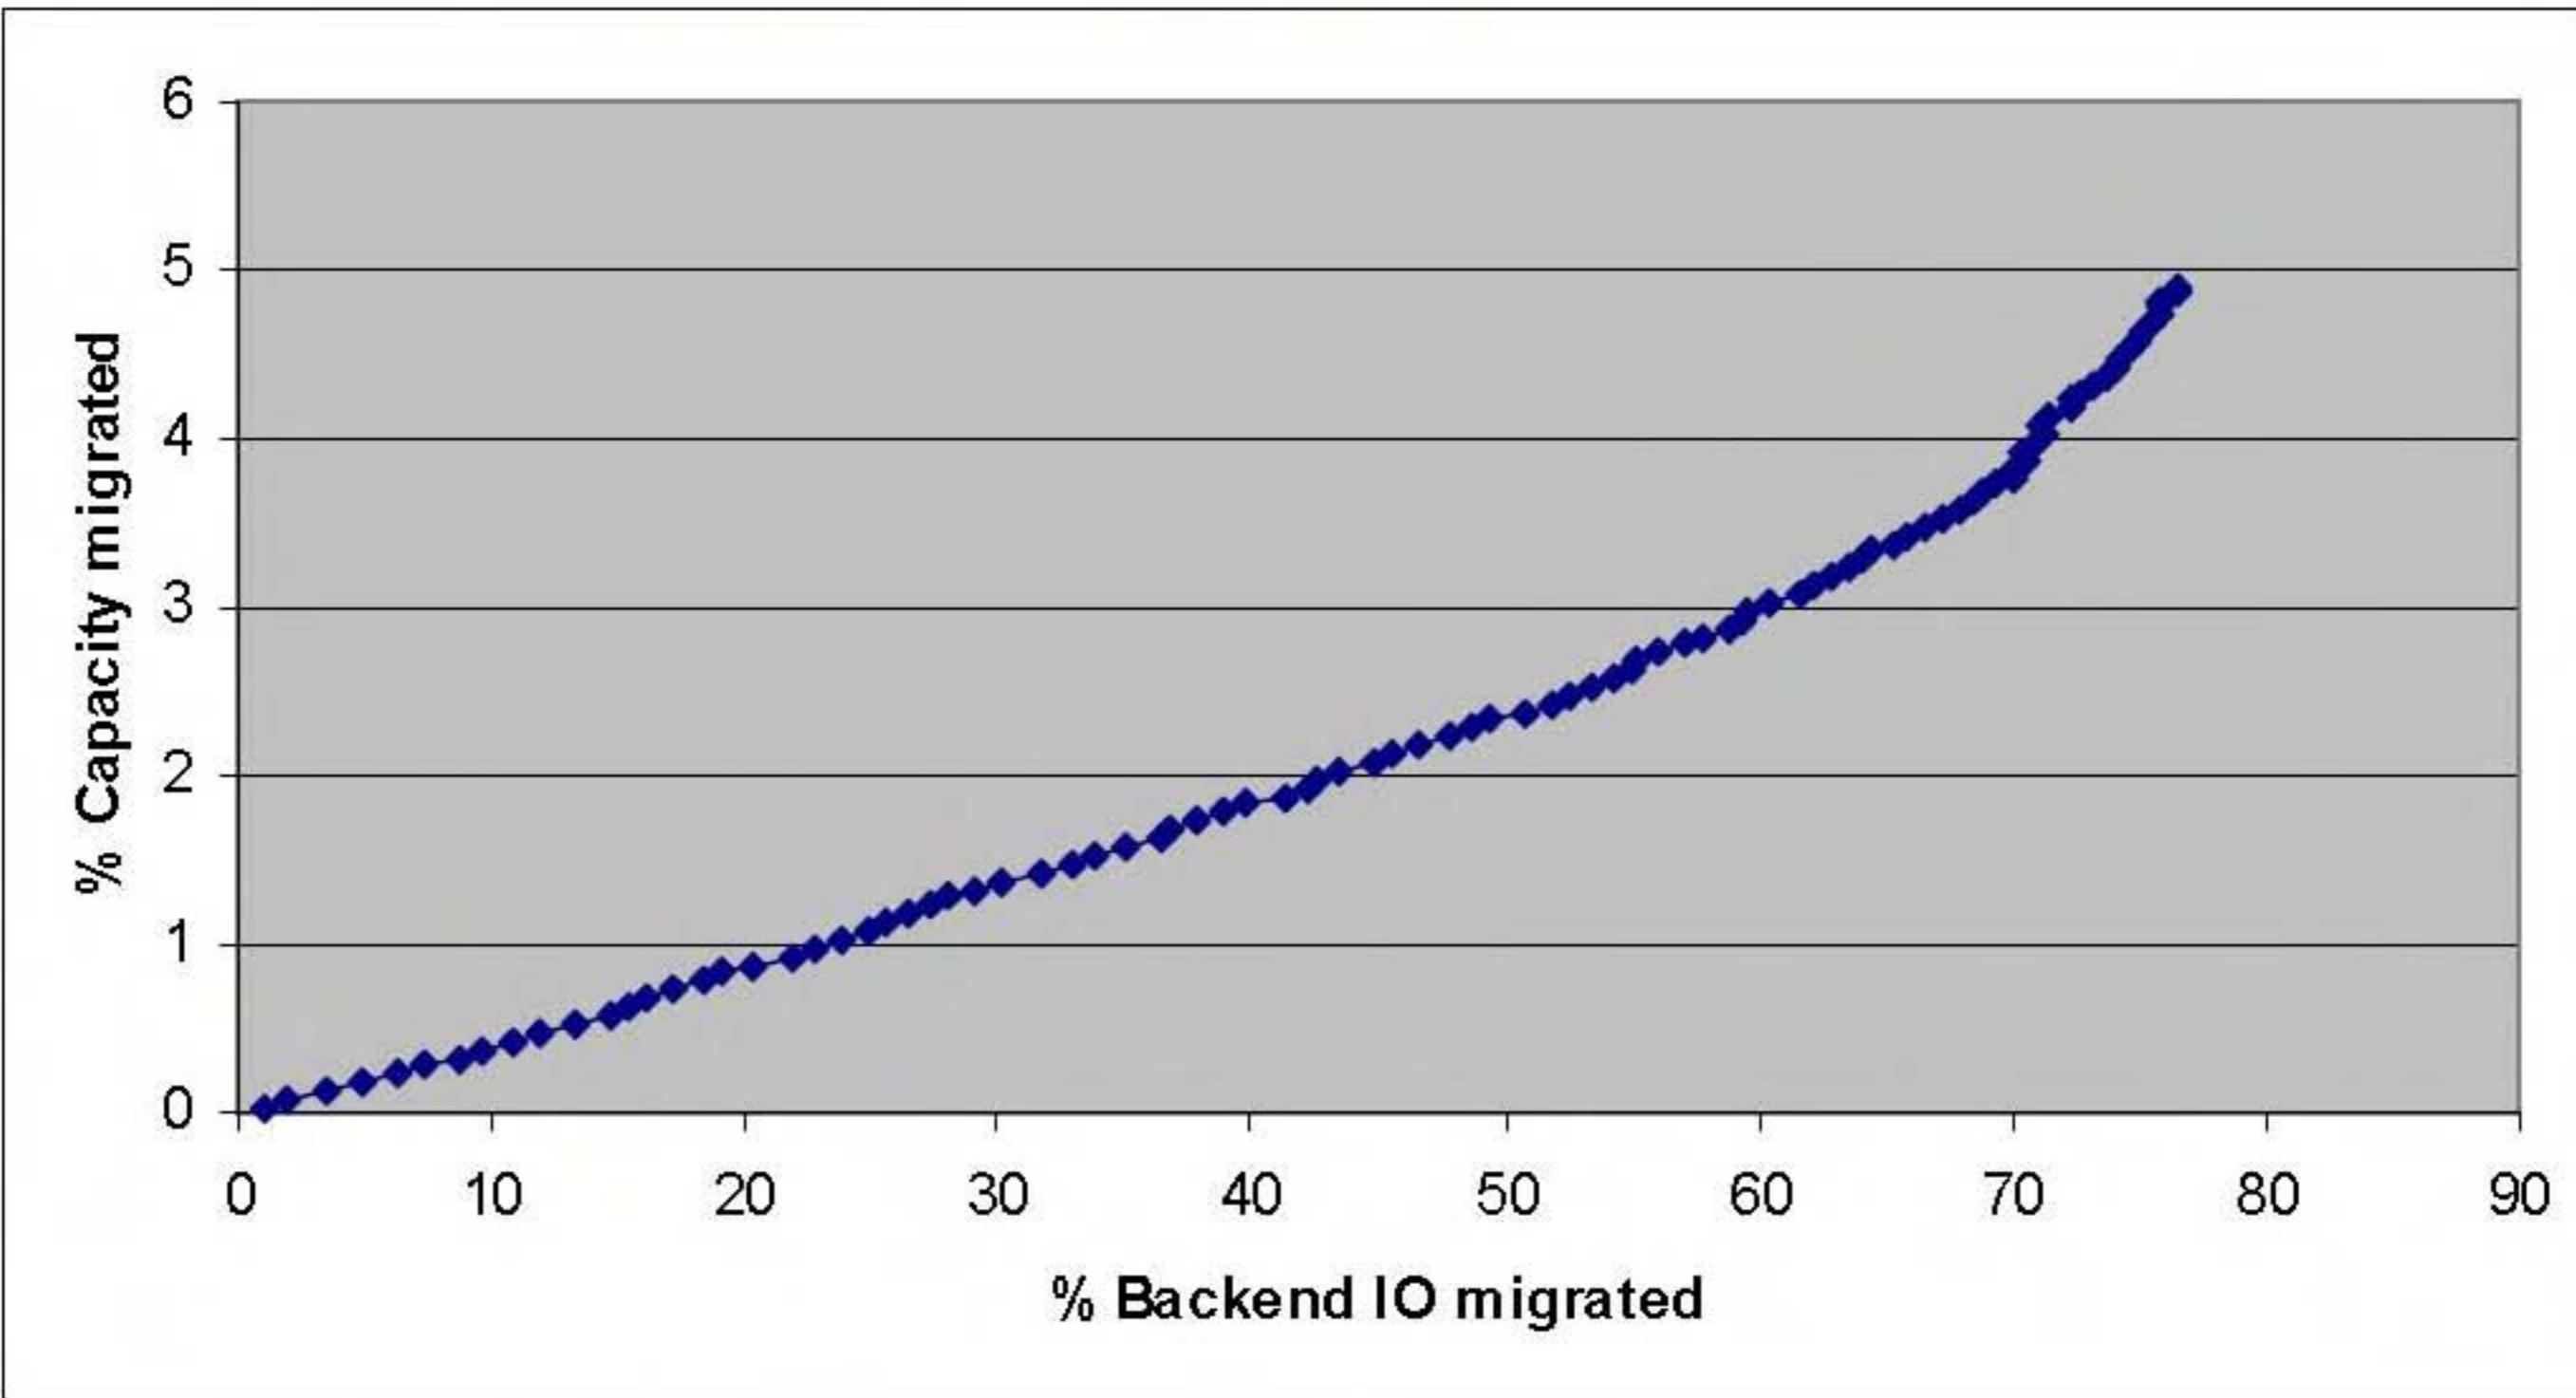
## How Easy Tier Automatic Mode Virtualizes Extents



## SPC-1 Throughput Improvement with Easy Tier



## SPC-1 Backend I/O Migration



## Summary

---

- TPF's combination of very high access density and random I/O is a natural match for Solid State Drives
- Significant improvements in throughput and response time can be achieved
- At the current price point for this technology, hybrid SSD and HDD configurations remain very attractive
- A variety of approaches are available to select what data should be migrated to SSD based on your specific requirements

## More SSD Resources

- An Assessment of SSD Performance in the IBM System Storage DS8000  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101495>
- IBM System Storage DS8000 with SSDs: An In-Depth Look at SSD Performance in the DS8000  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101466>
- z/OS Support for Solid State Drives in the DS8000 Storage Subsystem  
<http://publibz.boulder.ibm.com/epubs/pdf/e0z2n191.pdf>
- IBM System z® and System Storage DS8000: Accelerating the SAP® Deposits Management Workload With Solid State Drives  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101442>
- IBM® System Storage™ DS8700™ Performance with Easy Tier®  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101675>

## More SSD Resources

---

- The free tooling described in this presentation is available for download from the z/OS Tools and Tips web page at

<http://www.ibm.com/systems/z/os/zos/downloads/flashda.html>

- For additional information contact
  - Leslie Sutton, IBM STG Systems Performance
  - [leslie7@us.ibm.com](mailto:leslie7@us.ibm.com)

## Disclaimer

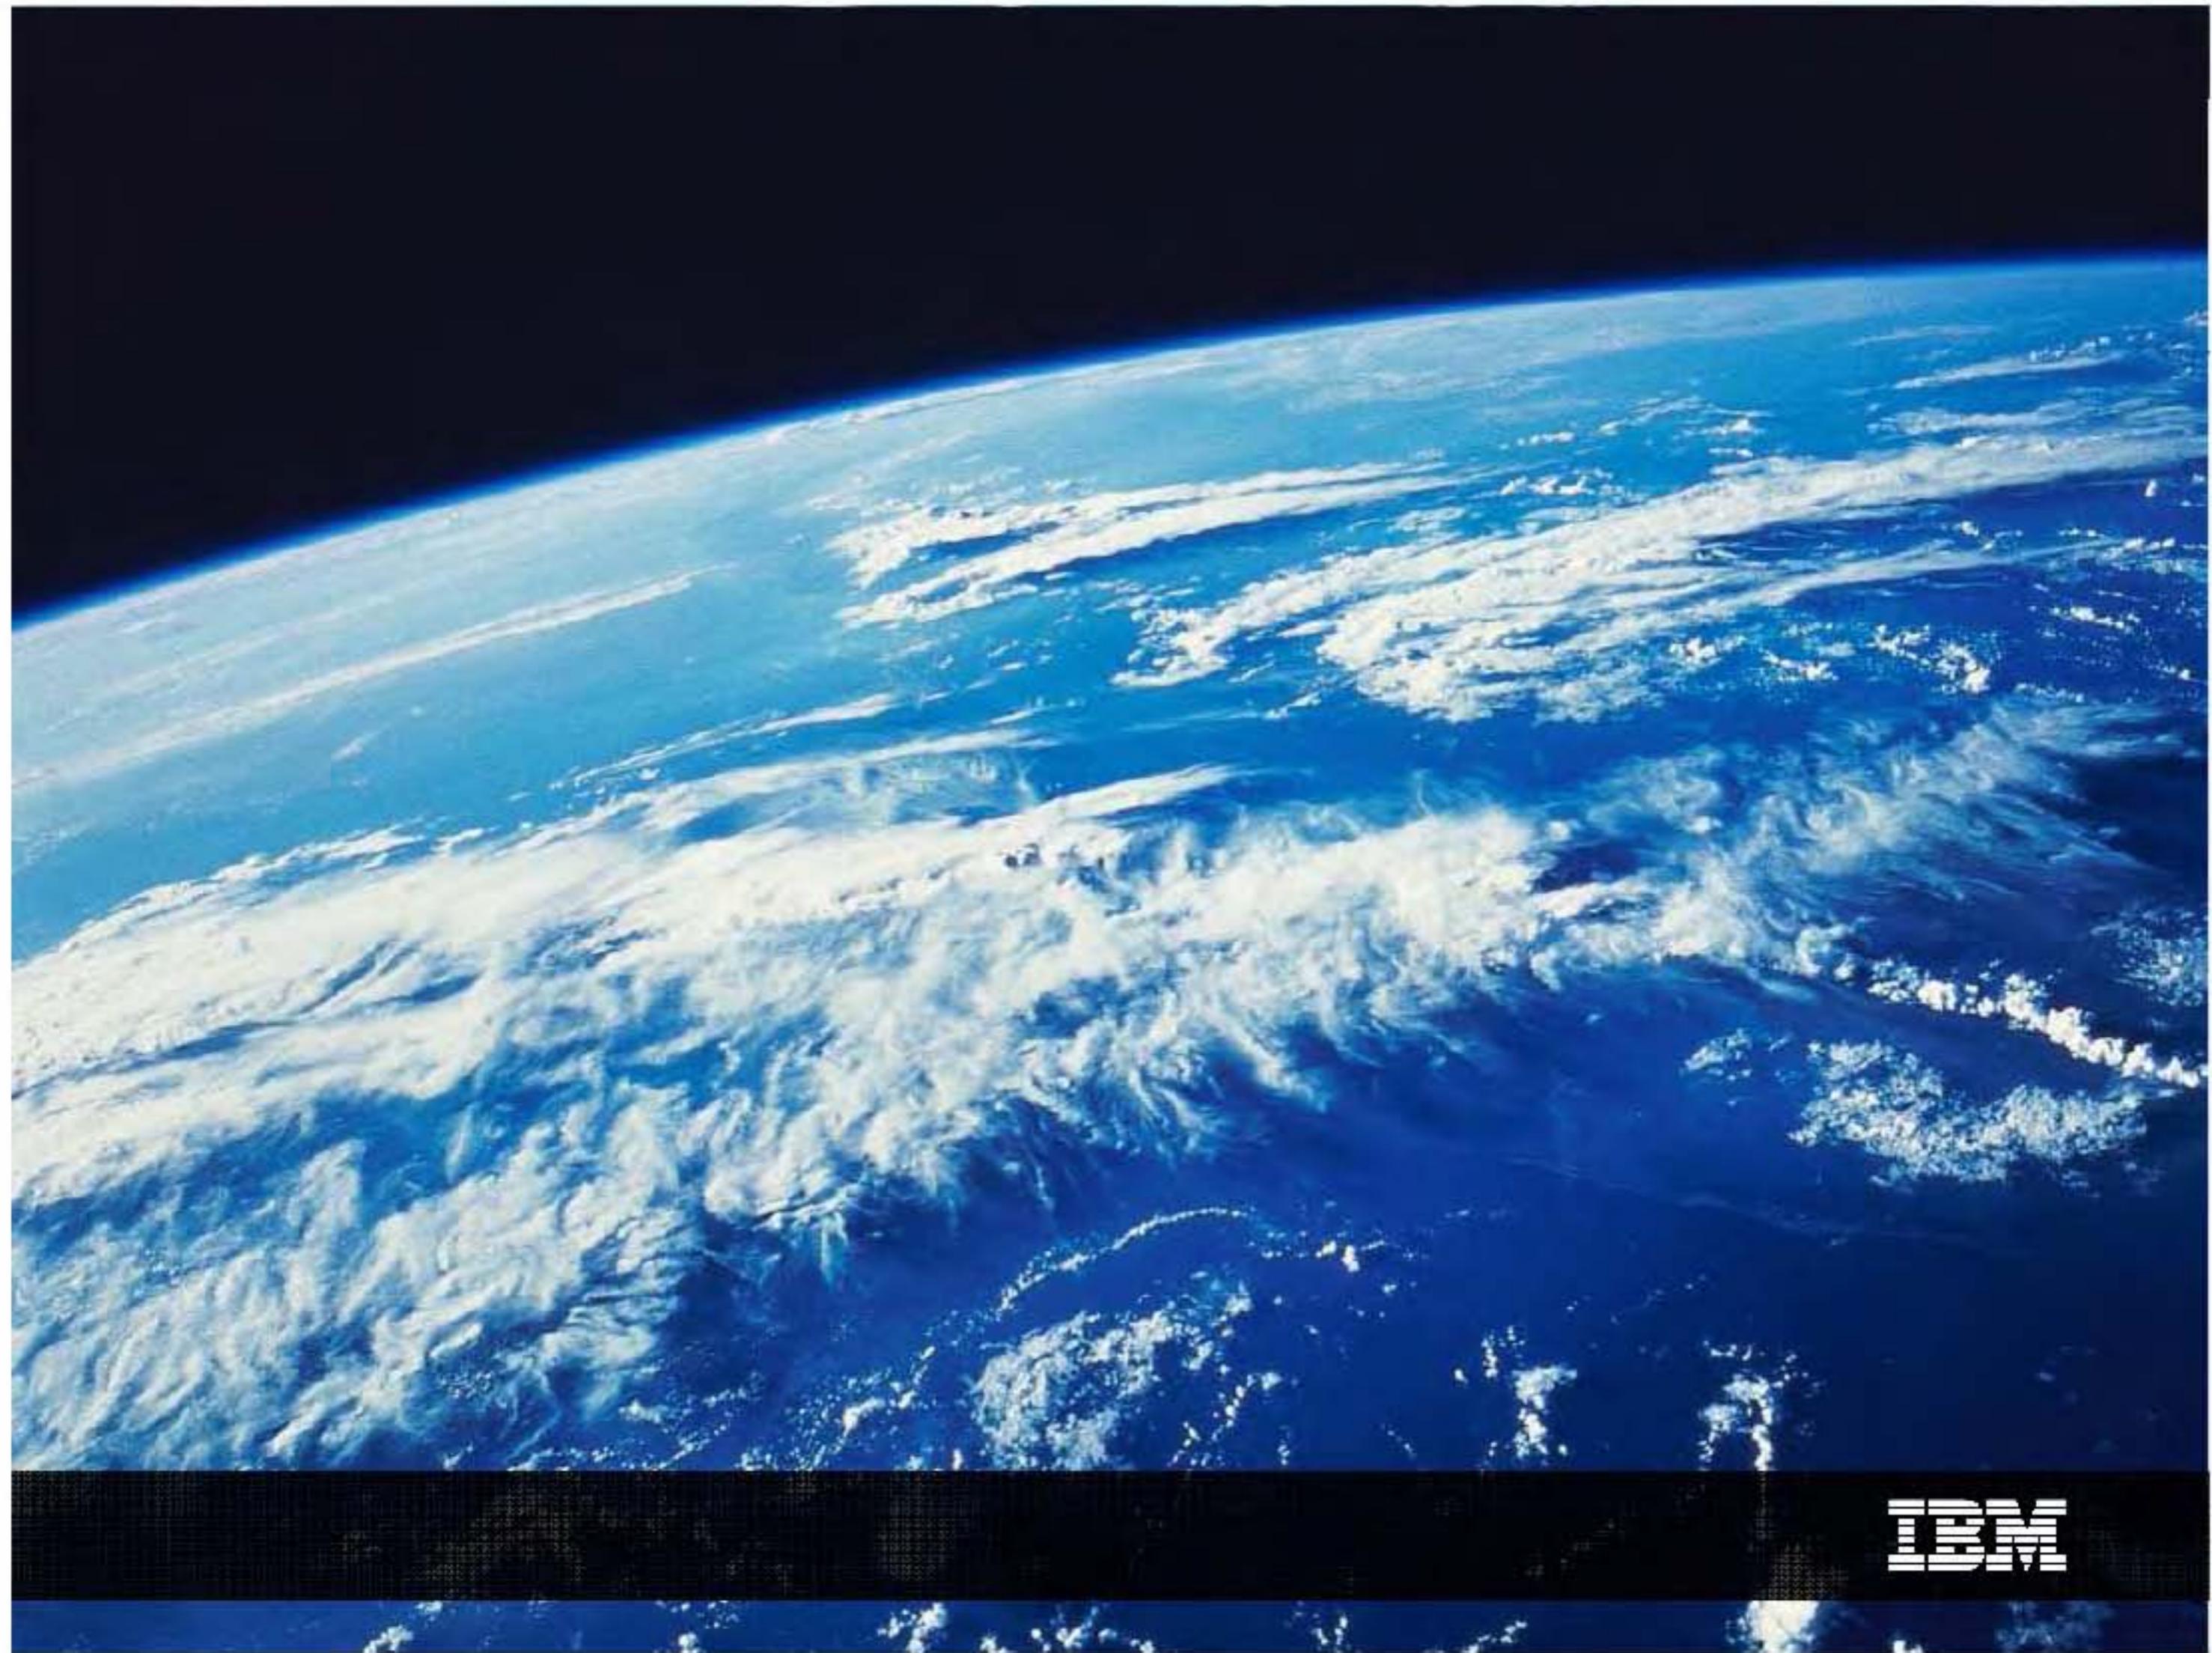
© Copyright IBM Corporation 2010. All rights reserved.

**U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.**

**THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.**

IBM, the IBM logo, ibm.com, System z, DFSMS, DFSMSdss, FICON, FlashCopy and DS8000 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.



IBM