

- Jana Jamsek
 - Advanced technical skills, Europe
-

Boost IBM i performance with IBM FlashSystem


Agenda

- IBM FlashSystem
- Benefits of FlashSystem for an IBM i customer
- IBM i workloads with IBM FlashSystem
- Tools for evaluating an IBM i workload for FlashSystem
- Customer cases
- Best practise for implementing IBM i with FlashSystem

FlashSystem - Target the weakest link in the IT stack...

Performance increased in the 10 last years ...

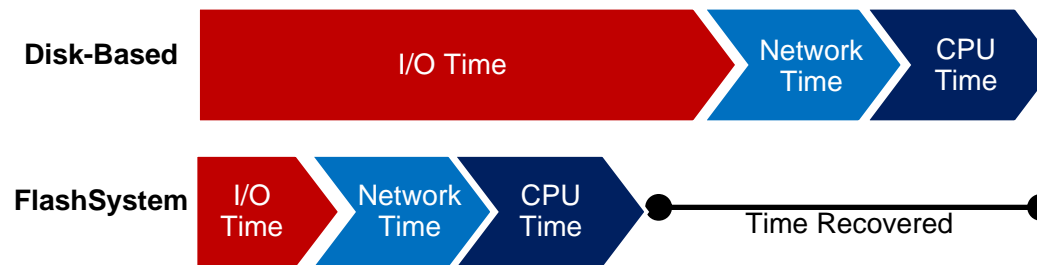
- CPU Speed: ~ 8-10x
- DRAM Speed: ~ 7-9x
- Network Speed: ~ 100x
- Bus Speed: ~ 20x
- Disk speed: > 30x (IBM FlashSystem)



The storage disk tier is the **BOTTLENECK** and is slowing everything else in the IT stack!

MicroLatency: How it Affects Your Business

CPU, Network, Memory and Bus are getting **faster**... so all of your applications **WAIT** on disk



Applications see time waiting, not IOPS

IBM FlashSystem

FlashSystem 900




FlashSystem V9000



IBM FlashCore technology

Flash chip

- Fast at reads
- Field Programmable Gate Array (FPGA)
- Fast writes / ECC, Wear Leveling and Overprovisioning
- Improves waer life




IBM MicroLatency Module

- IBM Patented Variable Stripe RAID™ (VSR)
- Preserves life, protection and performance
- Maximum level of system protection / Two-Dimensional (2D)

Flash RAID

- VSR protects from flash chip or sub-chip issues
- System-level RAID protects against abrupt module



Benefits of FlashSystem for IBM i customers

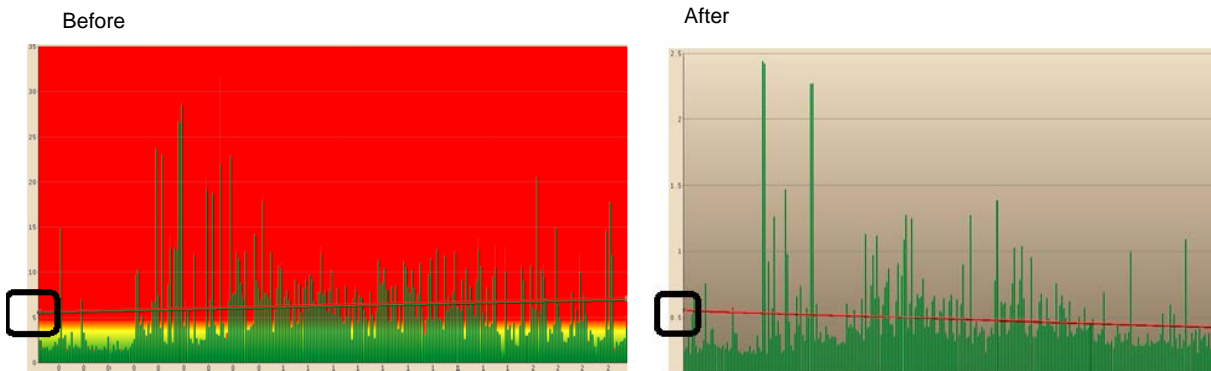
- Performance boost
- Sharing of Storage among several host systems
- Re-provisioning of storage capacity
- Employing Easy Tier (when virtualized with Spectrum Virtualize)

With V9000 or with virtualizing FlashSystem:

- Implementing Real-time Compression
- Implementing High Availability or Disaster recovery solutions with Copy services

Performance boost

- At many IBM i customers FlashSystem drastically reduces duration of jobs and improves response time
- Customer's example: Before implementing FlashSystem the job lasted 53 minutes, after FlashSystem the same job lasted 9.6 minutes.
- Another customer's example: Disk response time before and after employing FlashSystem show about 10 times improvement. The IBM i performance graphs are shown below



Sharing and re-provisioning

- FlashSystem is robust against different workload patterns, comparing to other Storage systems, so we recommend to share disk pools among several host systems. This enables very efficient usage of storage resources.
- Customer's example of re-provisioning from DS8000:
 - The ranks used for an LPAR became overloaded.
 - 4 additional ranks are added to the extent pool of the LPAR, Easy Tier is used to balance the data over new ranks


Real-time Compression

- IBM i can take advantage of Real-time Compression in either SVC, Storwize or FlashSystems V9000.
- Real-time Compression allows the use of less physical space on disk than is presented to the IBM i host. Capacity needed on the storage system is reduced due to both Compression and Thin provisioning.
- Customer's example: Real-time Compression of IBM i and VMWare with FlashSystem V840
- 73% of Compression savings
- Compression ratio is 3.7 : 1
- Response time is excellent, doesn't seem to be affected by compression

Implementing High Availability and Disaster recovery solutions


- V9000 Copy services integrated with PowerHA for i provide solutions for High Availability and Disaster Recovery
- Solutions are
 - Highly Automated
 - Use little resources
 - Require little maintenance
 - Low cost of ownership
 - Excellent RPO and RTO
- The customers who are not yet ready to implement applications in Independent Auxiliary Storage Pool (IASP) may use Full system HA and DR solutions with V9000

Example of DR solution with PowerHA for i: Scenario for Unplanned outages




IBM i workloads and FlashSystem

- **Good candidates for FlashSystem**
- Any jobs that are I/O intensive with higher disk read times than 500 microseconds = 0.5 ms per I/O
- The nightly batch workload and End of Day batch processes usually fall into this category
- Many jobs running during production hours can drive significant I/O too



IBM i workloads and FlashSystem

- **Good candidates for FlashSystem**
- Identify which jobs are driving many reads, these are good candidates for FlashSystem
- Can be identified with iDoctor



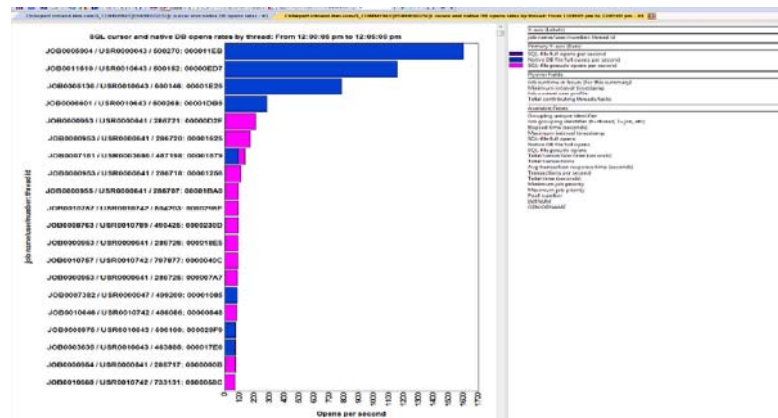
IBM i workloads and FlashSystem

- **Not good candidates for FlashSystem**
- Write intensive workloads that presently experience very good response time
- Jobs that perform a lot of IBM i journalling belong to this category
- Jobs that run presently on POWER internal SSD fall into this category
- Can be identified with iDoctor
- Customer's example from iDoctor:



IBM i workloads and FlashSystem

- **Not good candidates for FlashSystem**
- The workalods that drive excessive full-opens on the database files
- Can be identified by iDoctor
- Example of iDoctor: the jobs with more than 1000 full DB opens per second are not good candidates



Tools for evaluating FlashSystem for IBM i

We recommend the following tools to evaluate usage of FlashSystem with IBM i:

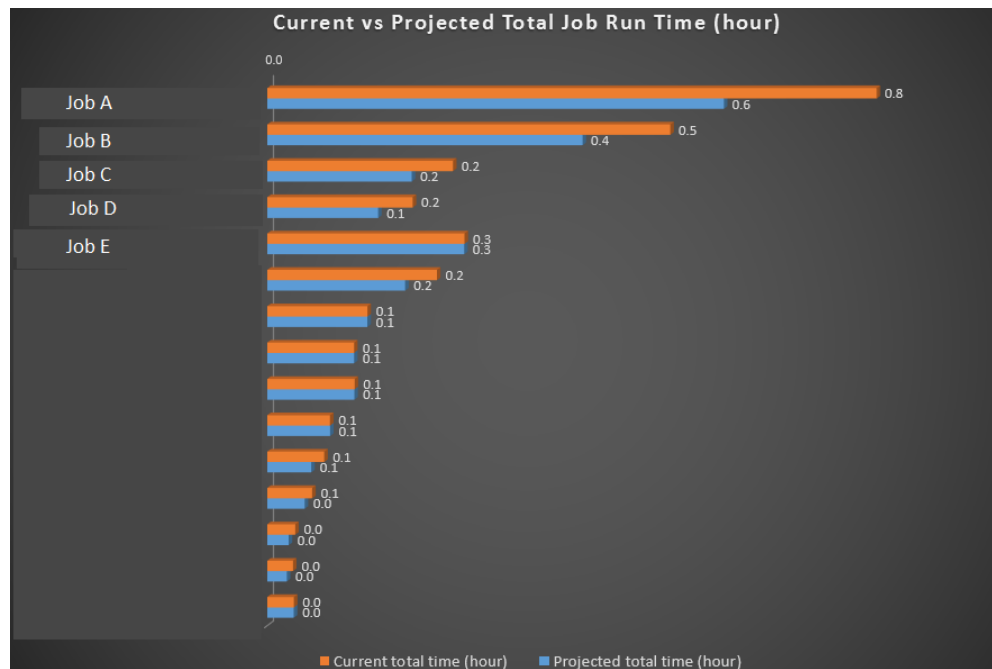
- Flash IBM i Performance tool (FLiP)
- Technical document *FlashSystems Evaluation Guide for IBM i Performance*
- Disk Magic

Flash IBM i Performance (FLiP) tool

- FLiP was created in IBM, can be downloaded from the following link:
<http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5291>
- Use FLiP tool for modelling performance improvement of IBM i **jobs** with FlashSystem
- Input to FLiP: IBM i Collection Services data
- In FLiP select the IBM i jobs you are interested in
- FLiP results:
 - Current job run time
 - Jobs sorted by Read wait time % - candidates for Flash
 - Total number of Read IOs per job
 - Average read disk response time per job
 - Current versus projected read and write latency for selected jobs
 - **Current versus projected job runtime**
 - Reduction and improvement factor of job run time

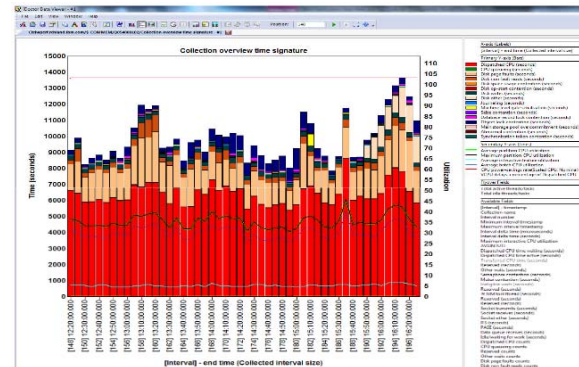
FLiP Example

- Following FLiP output shows current versus projected job runtime for selected jobs



FlashSystem Evaluation Guide for IBM i performance

- Available on the following link:
<http://w3-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD106347>
- Provides guidelines which types of IBM i workload profit and which don't profit from FlashSystem
- Describes the tools to identify potential Flash customers (iDoctor, Performance Data Investigator, etc)
- Describes which graphs should be used for evaluating FlashSystem for IBM i
 - Wait time signature
 - Page faults ranking
 - IO rates and KB per IO

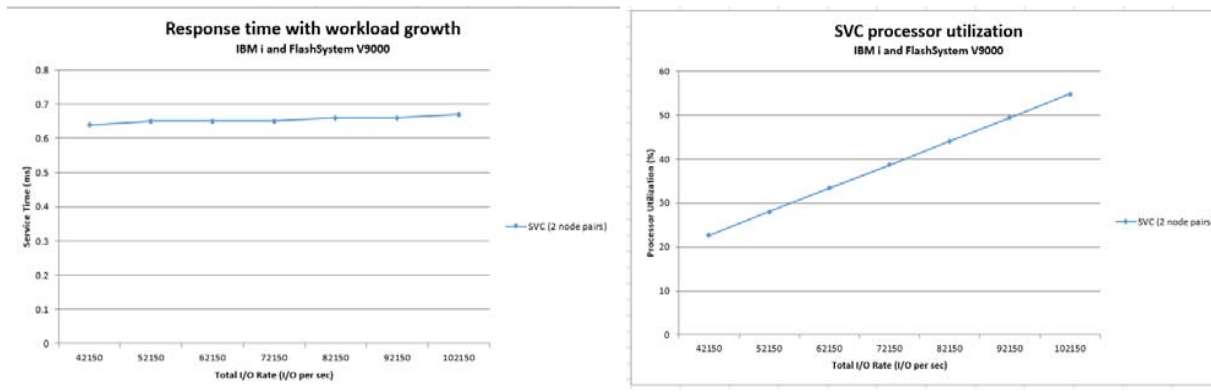


Disk Magic

- Modeling tool to help estimate performance of a Storage system
- Developed and maintained by the company IntelliMagic
- Models IBM, HP, HDS and EMC storage
- IntelliMagic closely collaborates with IBM development and performance teams
- Disk Magic can be obtained from the following web sites:
 - IBMers:
<https://w3-03.sso.ibm.com/sales/support/ShowDoc.wss?docid=SSPQ048068H83479I86>
 - Business partners:
<http://www.ibm.com/partnerworld/wps/servlet/ContentHandler/SSPQ048068H83479I86>
- Disk Magic currently doesn't support FlashSystem with IBM i
- Possible work-around: Model DS8870 / DS8884 / DS8886 with Flash to simulate FlashSystem
- Disk Magic support of FlashSystem for IBM i is planned for 4th quarter 2015

Disk Magic example

- Modelling FlashSystem V9000 for IBM i
- Insert IBM i reports of Collection Services data
- Save Disk Magic base with present storage configuration
- Model DS8870 with All Flash or DS8886 with HPF, virtualize with SVC
- Observe modelled response time, workload growth, utilizations




Customer case 1

- Customer – large transportation and logistic company for Europe
- Configuration before FlashSystem:
 - POWER7 740 with 128 GB memory
 - Production LPAR: 2 cores, IBM i 7.1
 - Testing and developing LPAR: IBM i 7.1
 - V7000 Gen 1 with 15 * HDD in Raid-10
- The customer's goal was to increase capacity and maintain / improve current performance
- Configuration with FlashSystem
 - POWER and IBM i LPARs the same as before
 - FlashSystem 840 with 12 TB capacity
 - FlashSystem is connected to V7000
 - Production LPAR is migrated with vdisk mirroring to FlashSystem

Customer case 1 - Before FlashSystem


From iDoctor Collection services investigator



- Read/Write ratio: $1200 / 4500 = 0.27$
- Max 8000 IO/sec with Read response time 5.2 ms
- Average Write resp. time = **0.7 ms**, Average Read resp. time = **2.7 ms**, Average resp. time = **1.2 ms**

Customer case 1 - After FlashSystem

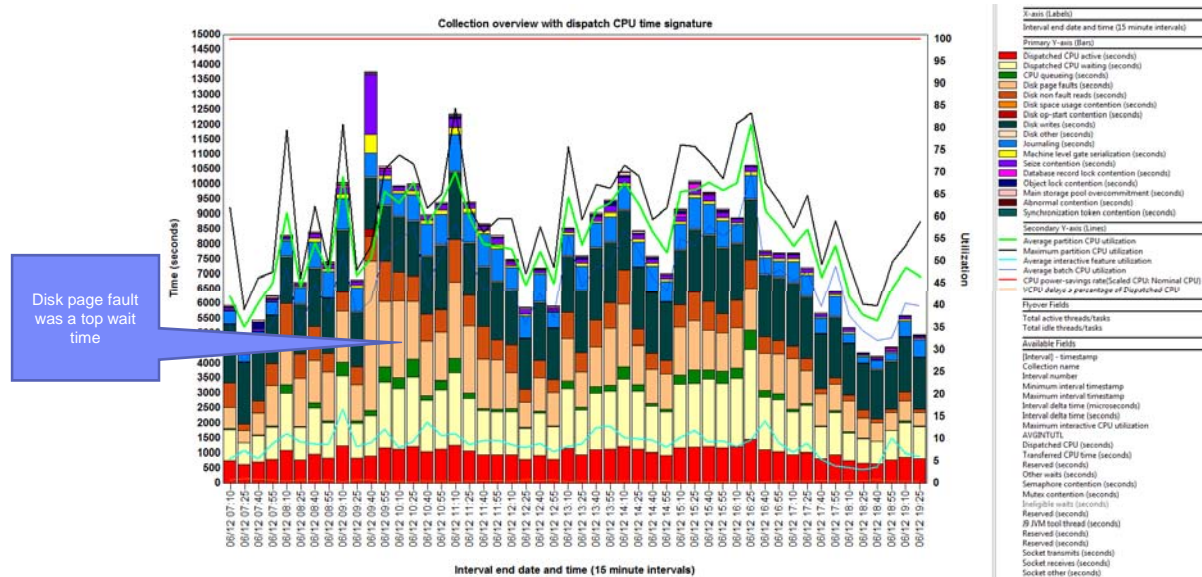
From iDoctor Collection Services investigator



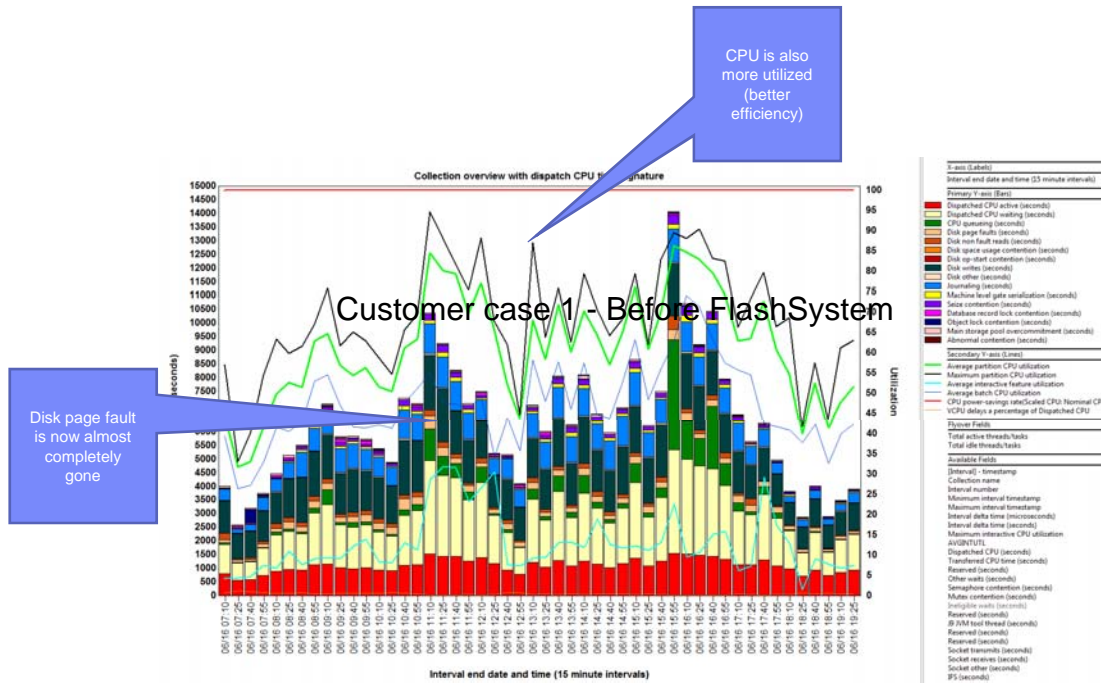
- Read write ratio: $2200 / 5500 = 0.4$
- Max 18000 IO/sec with Read response time 1.3 ms
- Average Write resp. time = **0.7 ms**, Average Read resp. time = **0.4 ms**, Average resp. time = **0.65 ms**

Customer case 1 - Before FlashSystem

From iDoctor Collection services investigator



Customer case 1 - After FlashSystem



Customer case 1 - Compare job runtimes before and after

- Example: Job A
- Before FlashSystem it ran 3152 sec (**53 min**) spending 44 min for waiting for Disk Reads. The time spent using the CPU was 120 sec (2 min)
- After FlashSystem the same job ran 576 sec (**9,6 min**) spending only 4,9 min waiting for Disk Reads and 134 sec (2,23 min) using the CPU
- Processing time for that job was reduced - **5,17 x faster**
- The CPU Efficiency for that job is **5,39 x better**

Customer case 1 - FLiP projections vs results

Status before FlashSystem and FLiP estimation

Job name / user / number - thread	Disk READ time	Disk WRITE time	Journal Wait time	Other Disk contention time	CPU time	Average Disk READ resp. time (msec)	Average Disk WRITE resp. time (msec)	Average Journal resp. time (msec)	Disk Read Time vs CPU time	Idle Time	Other Wait Time	Total Time
Job A / thread 1	2646.2	298.3	0.0	0.2	120.3	4,864	2,098	0.538	45.0	0.0	87.0	3152.5
Job B / thread 1	1136.5	76.9	16.5	1.6	324.3	1,639	4,355	3.022	9.0	0.0	4.0	1562.4
	986.8	0.7	0.0	0.2	33.5	5,397	1,046	0.532	58.0	0.0	0.1	1071.6
	889.6	0.8	0.1	0.4	51.0	3,905	0,811	2,266	46.0	0.0	0.1	942.0

Job name / user / number - thread	Projected Disk Write Time	Projected Journal Time	Projected Current total time (min)	Projected total time (min)	Improvement factor	Current total time (%)	Projected total time (%)	Total Job Runtime reduction (%)	Projected time saved (min)	Current CPU efficiency	Projected CPU efficiency
Job A / thread 1	46.43	0.1	457.15	52.5	7.6	15.9	15%	85%	44.9	1%	25%
Job B / thread 1	1.72	16.1	196.8	26.8	5.9	2.8	28%	80%	18.1	2%	32%
	2.21	17.0	17.0	17.0	1.1	6.8	19%	80%	16.3	1%	31%
	0.31	1.81	1.81	16.7	2.3	6.8	15%	85%	13.4	1%	37%

For Job A / thread 1 FLiP estimated a runtime improvement from **3152 sec to 457 sec**, in reality it is **417 sec**
 For Job B / thread 1 FLiP estimated a runtime improvement from **1562 sec to 596 sec**, in reality it is **596 sec**

Status after FlashSystem

Job name / user / number - thread	Disk READ time	Disk WRITE time	Journal Wait time	Other Disk contention time	CPU time	Average Disk READ resp. time (msec)	Average Disk WRITE resp. time (msec)	Average Journal resp. time (msec)	Disk Read Time vs CPU time	Idle Time	Other Wait Time	Total Time
Job B / thread 1	234.8	681.9	59.0	24.0	1403.0	0,299	0,589	0,841	0.0	44.1	1995.6	4442.5
	205.6	0.2	0.0	0.0	167.2	0,243	0,530	1,158	3.0	1167.2	0.0	1569.0
	183.5	102.8	13.0	1.9	290.1	0,312	4,271	3,006	2.0	0.0	5.3	596.5
	165.1	2752.4	160.0	2.1	8773.4	0,190	1,048	1,388	0.0	0.0	74212.8	74212.8
	151.5	11.8	6.6	0.3	481.3	0,342	2,970	4,213	1.0	0.0	24.0	675.6
	144.0	0.7	0.2	0.0	34.4	0,377	0,609	3,050	17.0	0.0	3143.3	3322.6
Job A / Thread 1	143.9	3944.6	721.4	0.3	5110.1	0,181	0,837	1,885	0.0	0.0	76143.3	76143.3
	134.5	131.9	0.0	0.2	106.7	0,364	0,910	0,532	3.0	0.0	44.3	417.6

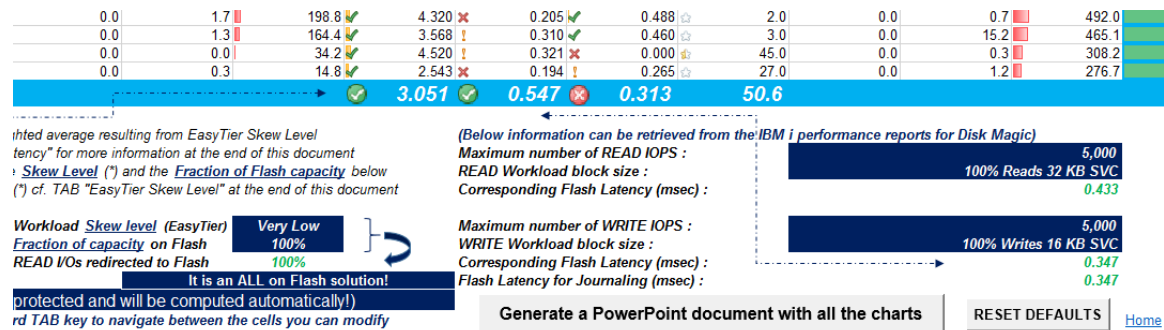
Customer case 2

- Customer – a bank in Europe
- Present configuration:
 - POWER, IBM i Rel 6.1
 - 72 * internal SAS 15 K RPM disk drives
- Proposed configuration:
 - POWER8, IBM i 7.2
 - VIOS_NPIV connection
 - FlashSystem V9000
- The customer's goal: Improving performance of EOD and EOM jobs with FlashSystem
- FLiP tool is used to evaluate proposed solution
- Proof of concept with POWER8, FlashSystem 900 and SVC is performed
- The case is on-going

Customer case 2 – Using FLiP

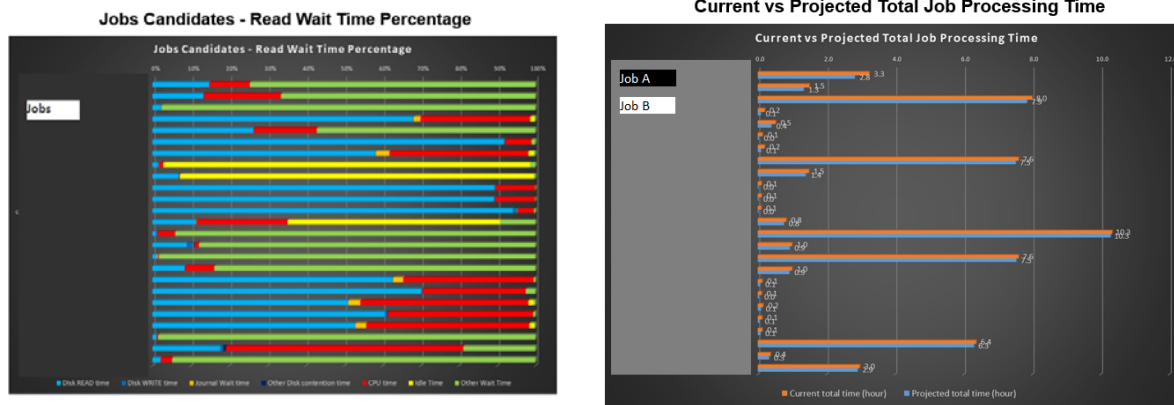
- Performance data are collected during 6 days
- FLiP was used for:
 - the jobs with highest accumulated read service time
 - the jobs of End-Of-Day

Workload characteristics obtained from Disk Magic were specified in FLiP
 100% capacity in Flash was specified in FLiP



Customer case 2 – Using FLiP

The jobs with highest accumulated read service time




The jobs with longer duration experience relatively small percentage of read wait time


Customer case 2 – Using FLiP

The jobs of EOD from two different days

Current vs Projected Total Job Processing Time



Current vs Projected Total Job Processing Time



Projected duration about 50% of present duration

Customer case 2 – Proof of Concept

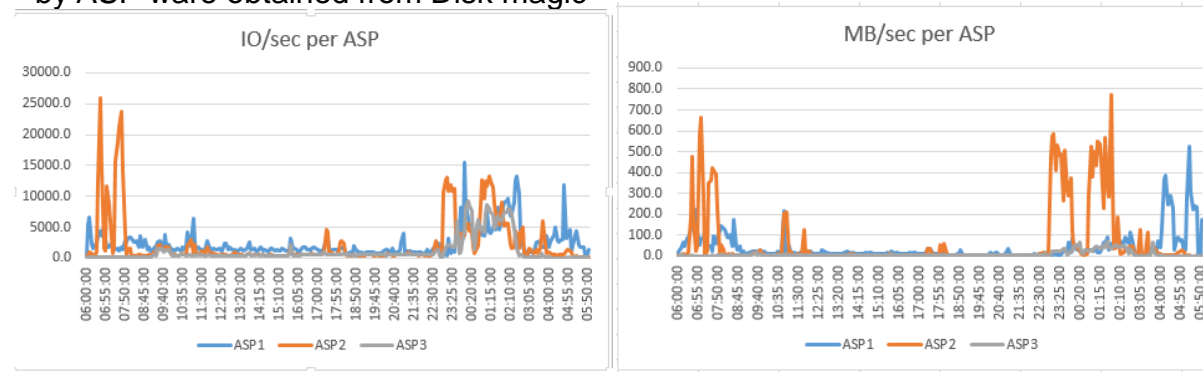
- Test of EOD:
- Previously: 4 hours : 40 min
- With FlashSystem: 2 hours : 20 min

Customer case 3

- Customer - Financial institution in Middle East
- Present configuration
 - POWER7 770, IBM i Rel 6.1
 - Internal disk drives in POWER for IBM i
 - ASP1: SSD in RAID-5
 - ASP2: SAS drives in RAID-5
 - ASP3 (Journal receivers): SSD in RAID-10
- The customer's goal: Improving performance of EOD and EOM jobs with FlashSystem
- Proposed configuration
 - Keep the same POWER, upgrade IBM i release to 7.2
 - FlashSystem 900
- To evaluate performance improvement we used Disk Magic and FLiP

Customer case 3 – Disk Magic modelling

- Performance data were collected during End-of-Month job, Workload characteristics by ASP were obtained from Disk magic



- Workload in ASP1
 - relatively write intensive with small block sizes,
 - random write cache efficiency is excellent (95%-100%).
 - Response time mostly below 0.1 ms, in the peaks below 0.5 ms.
- Workload in ASP2
 - shows high peaks in reads or writes with big block sizes
 - high percentage of sequential IO
 - Response time in the peaks can be as high as over 2 ms, otherwise it reaches 0.3 to 0.6 ms. .
- Workload in ASP3
 - Very write intensive, random oriented with small block size.
 - Write cache efficiency is excellent.
 - Response time is excellent, rarely goes over 0.1 ms

Customer case 3 – Disk Magic modelling

- Modelling of FlashSystem 900 was simulated with DS8870 All Flash 16 core with 256 GB cache
- Disk Magic default cache hits 20% read hit and 30% write efficiency were used
- Following peaks are modelled to obtain comprehensive picture of estimated response times:
 - by IO/sec and reads/sec
 - by writes/sec
 - by MB/sec
 - High response time in ASP2
 - Low response time in ASP2
 - by MB/sec in ASP1
- In most peaks the modelled response time is higher than present, in some peaks estimated resp time in ASP2 is lower

Customer case 3 – Disk Magic modelling

	IO/sec	Read %	Write cache eff % Random /Seq	Current Resp time (ms)	Estimated Resp time (ms)
Peak by IO/sec					
ASP1	4347	63	100 / 85	0.44	0.83
ASP2	26053	98	90 / 90	0.45	0.99
ASP3	198	18	100 / 94	0.11	0.34
Peak by MB/sec					
ASP1	6260	3	97 / 57	0.04	0.2
ASP2	9003	67	0 / 91	1.88	1.08
ASP3	5956	0	99 / 93	0	0.17

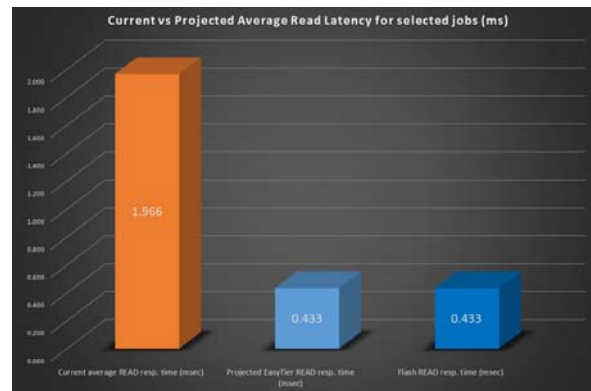
Customer case 3 – FLiP estimation

- Performance data are collected during 3 days
- Following jobs are selected for each day
 - The important jobs listed by the customer
 - The jobs with the longest duration
 - The jobs that consume the most read service time, on which we can expect the most improvement from FlashSystem
- FLiP provides estimation for jobs regardless in which ASP the job runs
- Workload characteristics from Disk Magic were inserted in FLiP

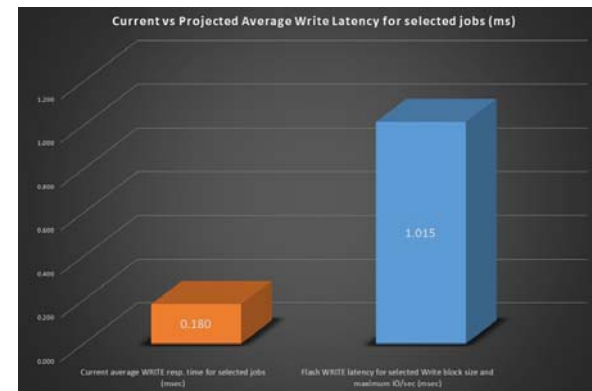
Customer case 3 – FLiP estimation

The important jobs listed by the customer

Current vs Projected Average Read Latency for selected jobs (ms)

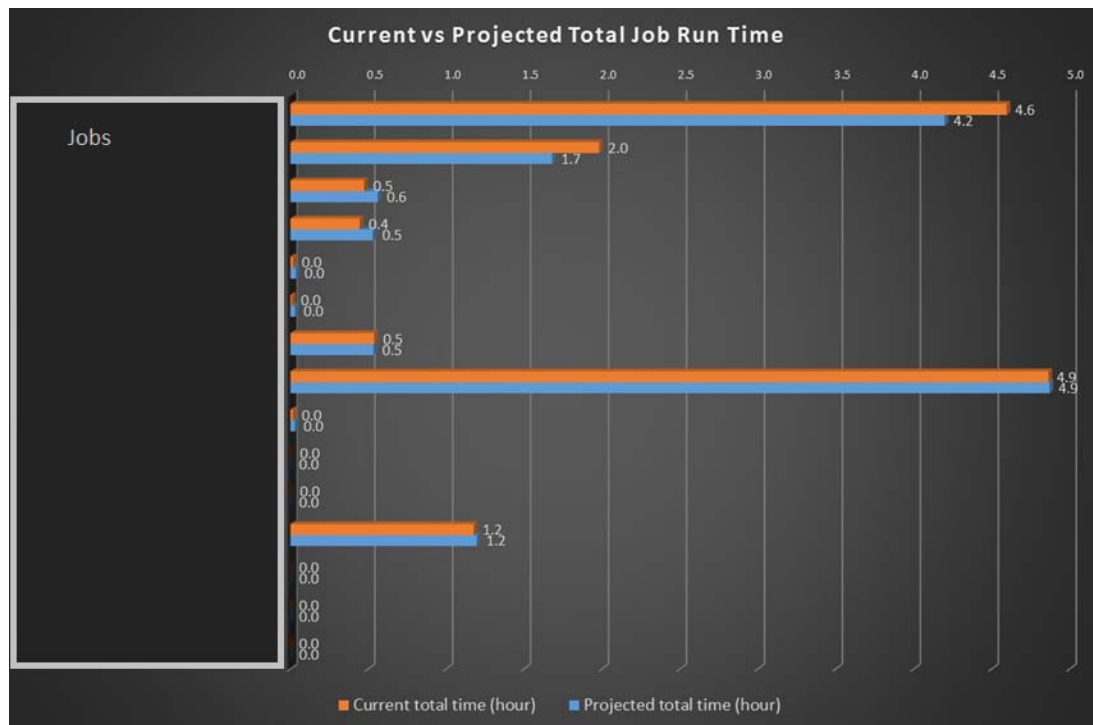


Current vs Projected Average Write Latency for selected jobs (ms)



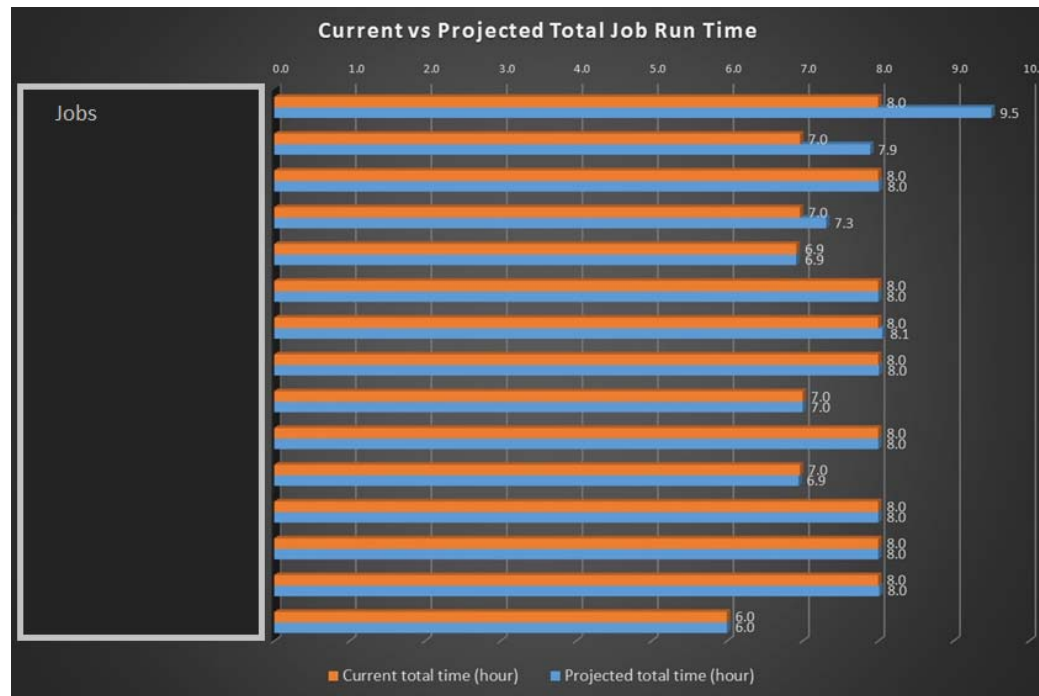
Customer case 3 – FLiP estimation

The important jobs listed by the customer



Customer case 3 – FLiP estimation

The jobs with the longest duration



Customer case 3 – Our Suggestion

Final decision

- After analysing with both Disk Magic and FLiP our suggestion is to stay with POWER internal disk and further invest in POWER resources
- FlashSystem wouldn't bring sufficient improvement since the performance on POWER internal disk are excellent

Customer case 4

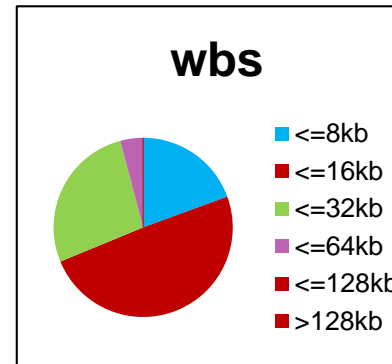
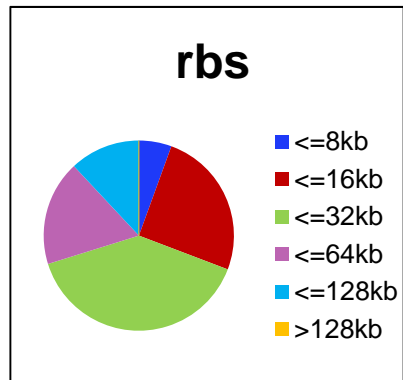
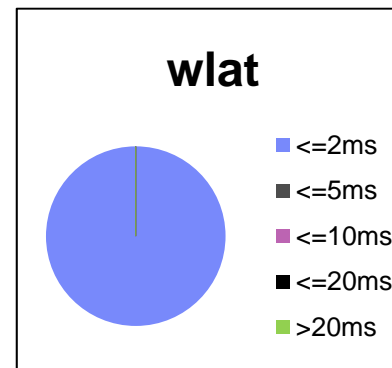

- Customer – car industry in Europe
- Present configuration
 - POWER8, 4 * LPARs with IBM i 7.1
 - Storage connection is VIOS_NPIV
 - V7000 Gen2 with
 - 42 * 15 K RPM disk drives, setup in 5 arrays .
 - FlashSystem 840 is connected as background storage and setup in 8 * 230 GB mdisks
 - Total capacity about 22.5 TB, capacity on Tier 0 (FlashSystem) is 1.84 TB - 8 % of total capacity
 - Easy Tier is used
 - Disk pool is shared among the IBM i LPARs
- Problem: performance don't meet customer's expectation

Customer case 4 – the Tools we used

- We used the following tools to investigate performance
- QPERF tool – provided information about performance of V7000
- IBM i Collection Services data
- IBM Storage Tier Advisory Tool (STAT)

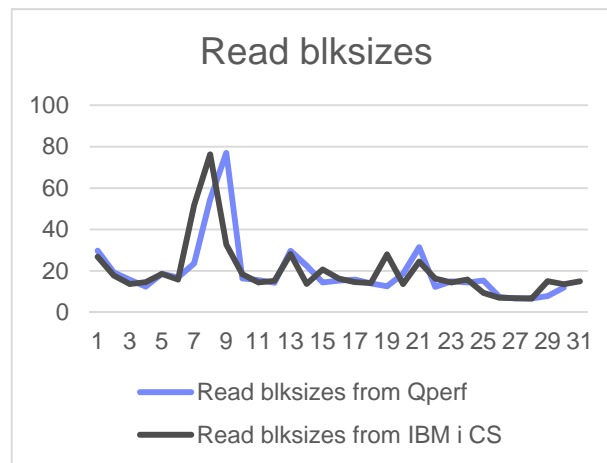
Customer case 4 – QPERF output

Latencies and Blocksizes

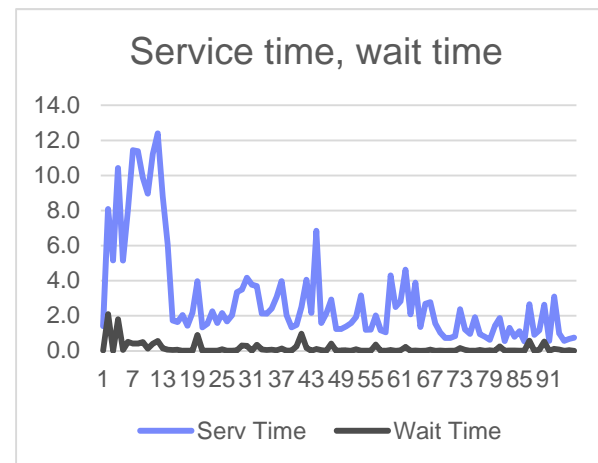


Customer case 4 – IBM i CS output

Read block sizes from one LPAR by QPERF and by IBM i CS data

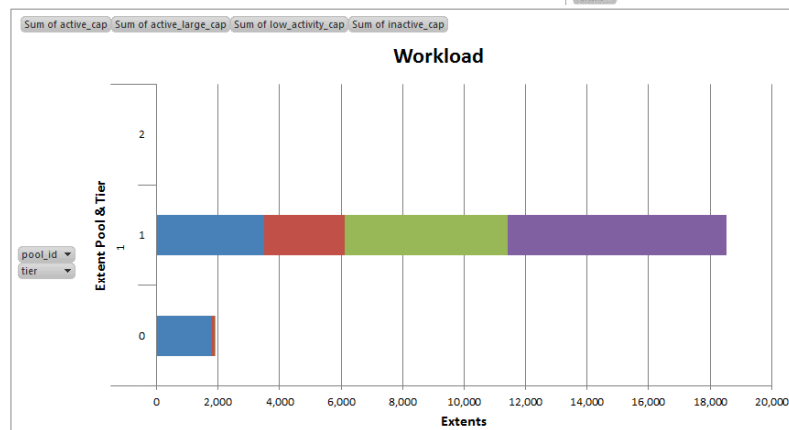
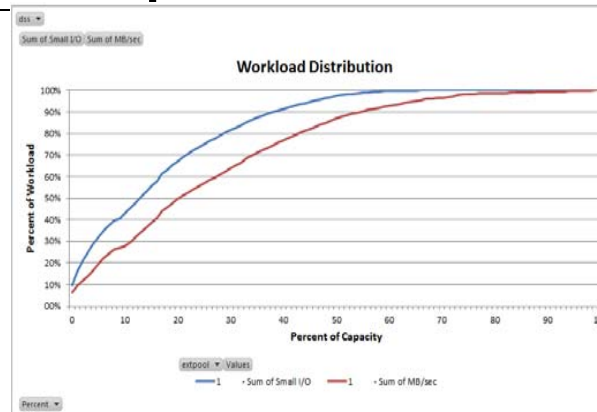



Service time and wait time of one IBM i LPAR



Customer case 4 – STAT output

EasyTierPool



- Tier 0 is small comparing to Tier 1
- Skew level is very flat
- Little movements of extents
- Small number of extents is considered as candidates (active capacity) due to big read blk sizes
- STAT recommendations are available

Customer case 4 – Recommendations and further actions



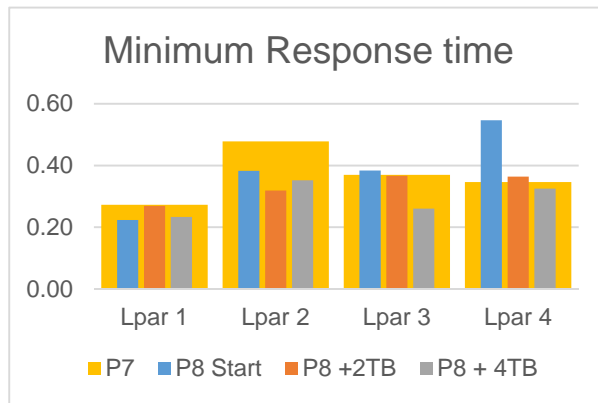
Our recommendations to improve performance

- Option 1: Add about 40% more capacity on FlashSystem, since the workload's skew level is flat
- Option 2: Add arrays of SSD, create a separate pool of SSD and migrate medium hot volumes to the pool
- Option 3: Add arrays of 15 K rpm disk drives to improve performance of data that are not relocated
- Option 4: Add NL drives for the cold data and let STAT optimize

Actions performed

- Test 1: add 2 TB of FlashSystem (4 TB altogether) as background storage
- Test 2: add 4 TB of FlashSystem (6 TB altogether) as background storage
- Optimum results (price performance) are achieved with Test 1: 4 TB on FlashSystem altogether
- Customer implemented 2 additional TB of FlashSystem – now about 15% of capacity resides on Tier 0
- Customer is satisfied with performance

Customer case 4 – results




Customer case 5

- Customer – big company in retail industry in Europe
- Configuration before FlashSystem
 - Power 7
 - 9 * LPARs with IBM i rel 7.1
 - 2 * DS8700 each with 128 ranks of 15 K RPM disk drives in RAID-5 and RAID-10
 - IBM i mirroring between the two DS8700 is implemented in each LPAR
- The customer decided for FlashSystem to boost performance of IBM i LPARs

Customer case 5 – configuration of FlashSystem

- POWER 7
- 9 * IBM i LPARs, release 7.1
- On each site:
- SVC with 2 I/O Groups
- Three FlashSystem 900, each with capacity app 52 TiB
- IBM i mirroring of each LPAR is implemented between the two SVC clusters



Customer case 5 – Issues to clarify before implementation

- How many LUNs to define on each FlashSystem
- How to connect FlashSystem to SVC cluster
- How to define disks pools for IBM i
- Which size of LUNs to use for IBM i
- How many path to implement for IBM i LUNs on each mirrored half
- Customer's request: 4 * important LPARs run each in a separate disk pool
- LUNs of the size 100 GB are planned
- FlashCopy will be implemented for backup purposes

Customer case 5 – Our recommendations

Recommendations for I/O groups and LUNs in FlashSystem

- Number of SVC I/O groups
 - Customer's IO rate reaches 205 000 IO/sec to one mirrored half in the peak, with blokisze about 18 KB.
 - One SVC I/O group with FlashSystem reaches performance knee at about 120 000 IO/sec
 - FlashCopy will pose some additional load to the configuration
 - We consider that SVC with two I/O ogroups and three FlashSystems in on performance edge. If possibe w recommned 3 SVC I/O groups
- Defining LUNs on FlashSystem 900
 - We agree with defining 32 LUNs from each FlashSystem 900

Customer case 5 – Our recommendations

How to define disk pools

- Each FlashSystem will provide capacity about 51 TiB
- Four important LPARs have capacities 78 TiB, 8 TiB, 9 TiB, 33 TiB
- In general we recommend to share a disk pool on FlashSystem among IBM i LPARs
- Option 1:
Create 4 Disk pools, and spread each disk pool across all three FlashSystems. This setup gives best average performance across all LPARs, however, it gives lowest resilience against failures - problems in one FlashSystem impact all LPARs, losing a single FlashSystem puts all workload on the other SVC cluster
- Option 2
Create one disk pool from each FlashSystem (Disk pool 1, Disk pool 2, Disk pool 3), Implement the biggest LPAR from Disk pool 1 and partially from Disk pool 2. Implement the second biggest LPAR from Disk pool 2. Share Disk pool 3 between the rest of the LPARs. This setup provides the best resiliency.

Customer case 5 – Our recommendations

Size of LUNs

- Generally, more LUNs better performance
- We presently recommend the size 40 GB – 200 GB
- The planned size of LUNs 100 GB is in-line with recommendations

Number of adapters


- Max IO rate on one 8Gb port at 70% utilization with IBM i performance PTFs installed, is 23100 IO/sec.
- Access Density of the biggest LPAR: 300000 IO/sec / 86240 GB = 3.5 IO/sec/GB
- Calculation of maximal capacity per port:
 - (IO/sec at 70% utilization / Access Density) * 40% for LUN utilization
 - (23100 IO/sec / 3.5 IO/sec/GB) * 0.4 = 2640 GB
- With 100 GB LUNs -> 26 LUNs per port
- If possible we recommend 4 path, so 64 LUNs per 4 ports
- *Further recommendation: Check that ether are enough ports available in SVC*

Customer case 5 – Performance after implementing FlashSystem

- Performance in the biggest IBM i LPAR:
- The jobs that run from midnight to 10:00 improved for about 30%
- The workload that runs from 10:00 to 14:00 experiences more than 90% of CPU usage, due to low disk latency enabling bigger workload. Consequently there is no performance improvement because of high CPU utilization.
- Future actions: enable more processor cores in this LPAR by RPQ or by upgrading to release 7.2


Customer case 5 – Performance after implementing FlashSystem

Performance data of one SVC with 2 I/O groups with 3 * FS900



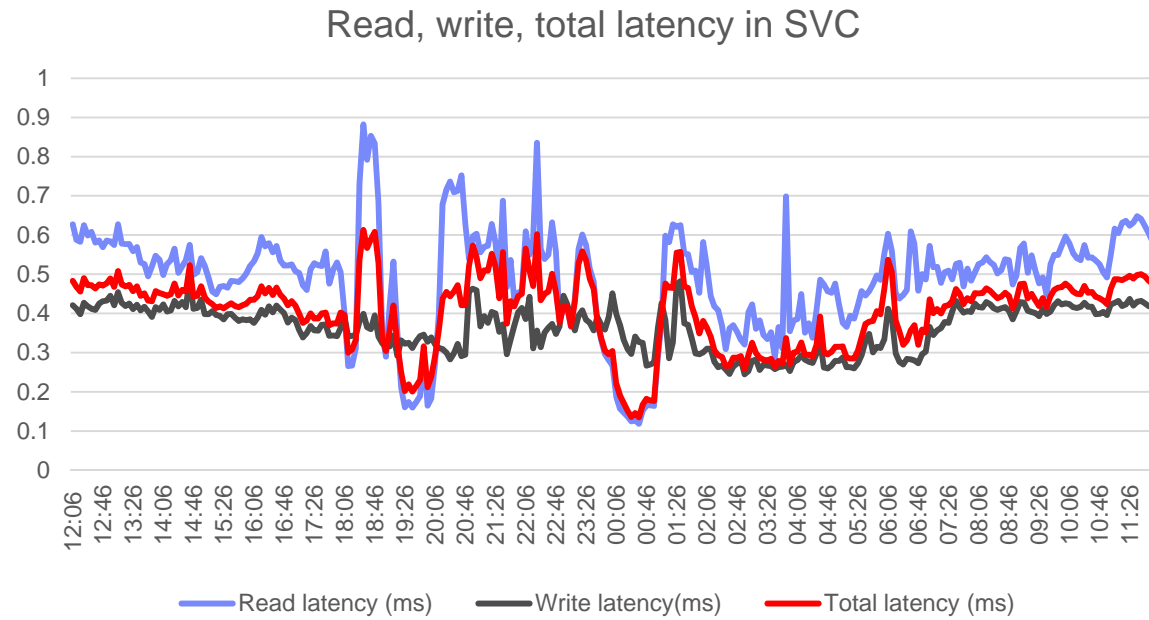
Customer case 5 – Performance after implementing FlashSystem

Performance data of one SVC with 2 I/O groups with 3 * FS900



Customer case 5 – Performance after implementing FlashSystem

Performance data of one SVC with 2 I/O groups with 3 * FS900



FlashSystem V9000 with IBM i – Best practise Tested scenarios

- Different LUN sizes covering the same capacity: 50GB, 150GB, 300GB, 600GB
 - **Best practise: size and numebr of LUNs**
- Disabled cache on the LUNs
 - **Best practise: enable or disable cache in FlashSystem**
- FlashCopy Snapshot influence on performance
 - **Best practise: FlashCopy of production IBM i**
- Metro Mirror influence on performance
 - **Best practise: Metro Mirror of production IBM i**
- Different number of path to IBM i: 2, 4, 6 path for a LUN
 - **Best practise: number of path for IBM i (adapters in IBM i)**

Testing environment in ESCC Lab

Pool V9000_Flash_Enclosure:

- 1 mdisk in the Disk pool,
- 11 * Flash modules of capacity 5.2 TiB in the mdisk
- Extent size 1 GiB
- 16* 150 GB LUNs
- Cache enabled on the LUNs
- 4 active path in VIOS NPIV

IBM i

- LPAR in Power7 770
- 7CPU, 8 GB memory
- IBM i 7.2 TR2, CUM level 15135
- Each of two VIOS: 1 CPU, 2 GB memory
- VIOS level: 2.2.3.52

IBM i workload for testing

- Workload consists of RPG programs with database files and journals
- **Writefile**
 - Sequential writes to 24 journaled files simultaneously
 - 70 million records to each file
 - Record: 75 packed decimal
 - OVRDBFILE FRCRATIO(1000)
 - Results in about 44000 writes/sec with blksize 11 KB
- **Readfile**
 - Sequential reads to 16 files simultaneously
 - 70 million records from each file
 - Record: 75 packed decimal
 - OVRDBFILE SEQONLY(*YES 2000)
 - Results in about 12000 reads/sec with blksize 80 KB

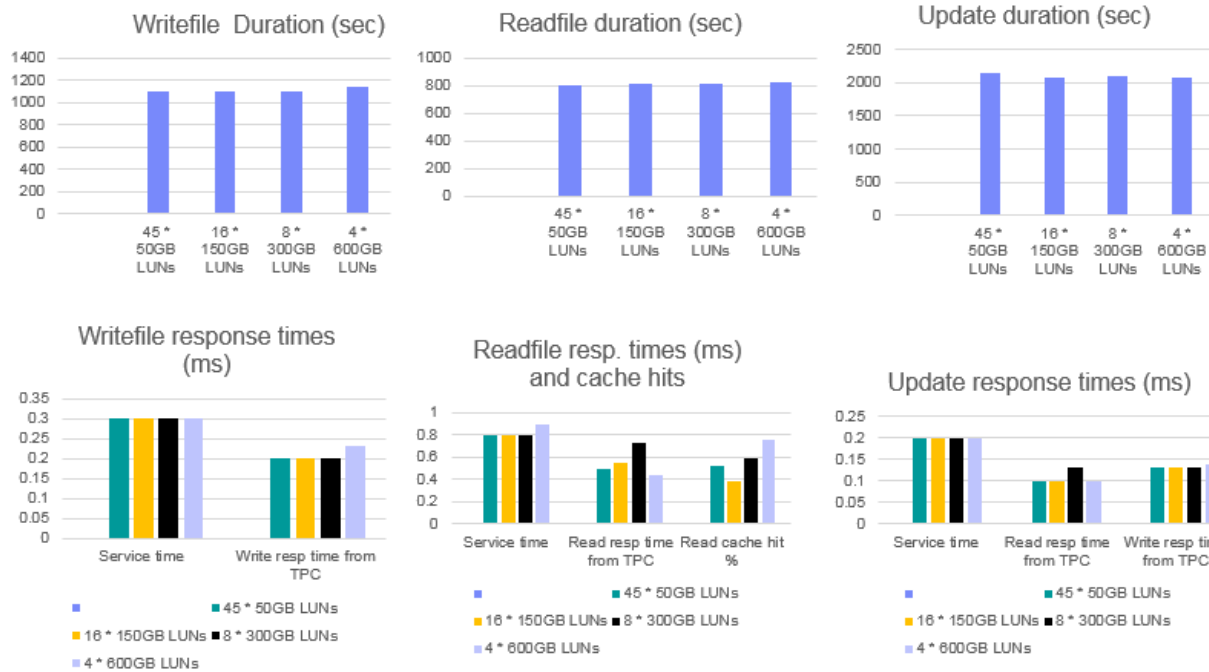
IBM i workload for testing

- **Update**
- Update1file
 - 50 iterations
 - Each iteration: 1 million times: Generates random number, retrieve the record with that number as key from file, and updates the record.
 - File is not journaled
 - Record: 20 packed decimal (key) and 15 character
 - OVRDBF with NBRRCD(30000) is used on the file
- Update2file
 - 20 iterations
 - Each iteration: Generates 5 million random numbers and writes them to a file1
 - Reads the record from file1 and updates the file1 by key with them
 - Files are not journaled
 - Record: 20 packed decimal (key) and 15 character
 - OVRDBF with NBRRCD(30000) is used on the file
- Update1file and Update2file run simultaneously
- Results in about 150 reads/sec and 23000 writes/sec with blocksize about 5 KB

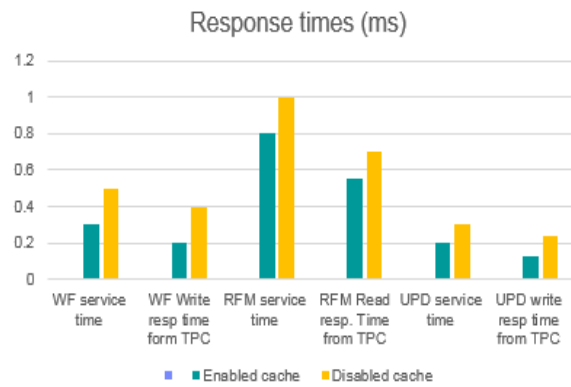

Test - Different LUN sizes

- Covering capacity 2.4 TB
- 45 * 50 GB LUNs, LoadSource 150 GB
- 16 * 150 GB LUNs
- 8 * 300 GB LUNs
- 4 * 600 GB LUNs
- Connected with 4 active path with 2 VIOS

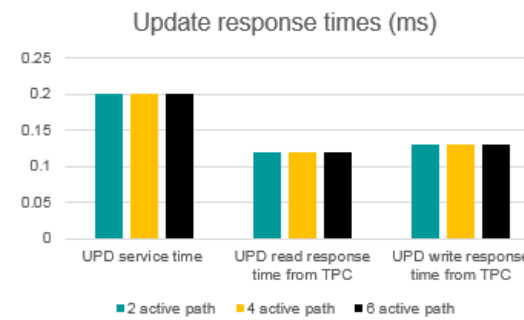
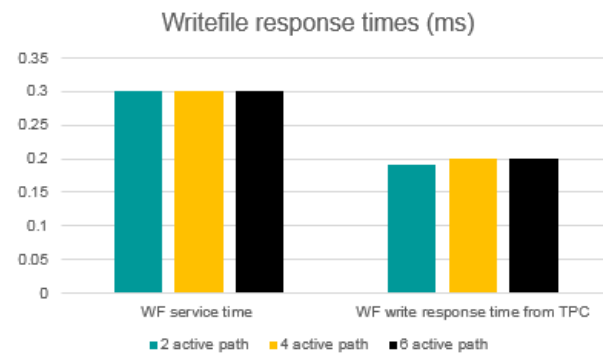


Different LUN sizes



Enabled / disabled cache




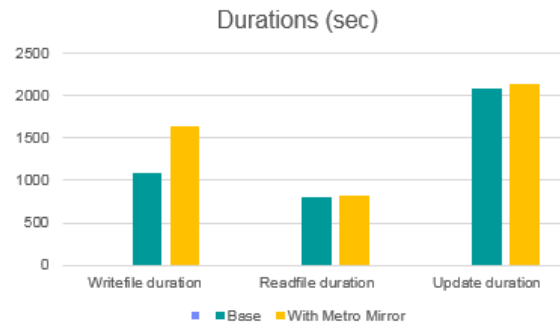
Different number of adapters



FlashCopy Snapshot influence



Metro Mirror influence



Conclusions: Best practise

- LUN capacity – we recommend from 50 GB to 250 GB
- Disabled cache – not recommended
- Number of ports – follow the guidelines in documentation, with big block sizes consider higher number of ports for good performance
- FlashCopy Snapshot influence - 50% - 60% with sequential writes, no influence with reads, insignificant influence with random writes
- Metro Mirror influence – 50% - 100% with sequential writes, no degradation with reads, small degradation with random writes