


# External Storage and IBM i Sizing and modelling

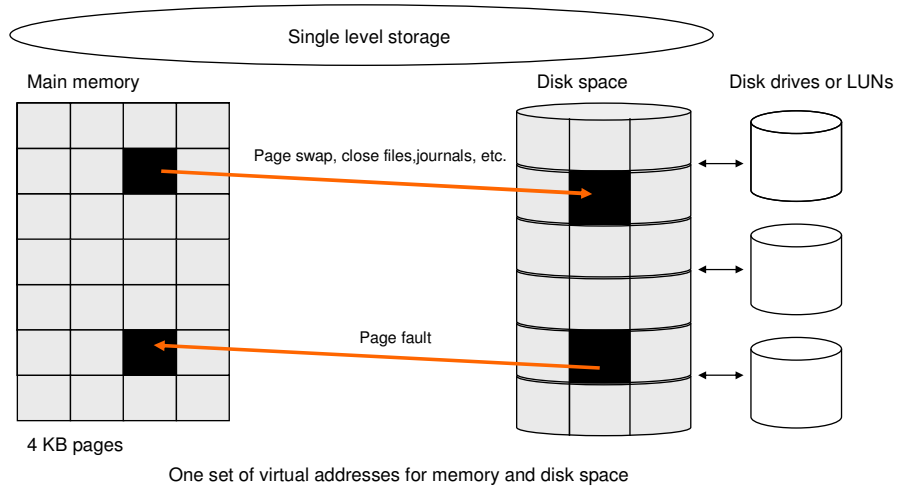
Jana Jamsek  
Advanced technical skills, Europe



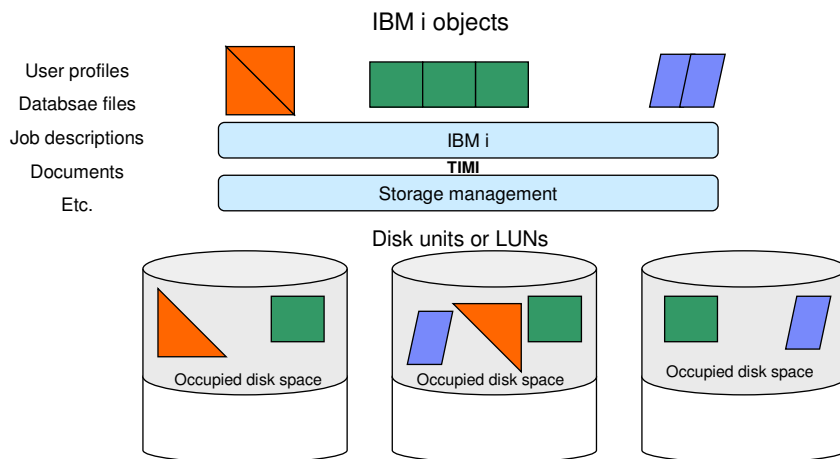
## Agenda

- IBM i architecture and external storage
- Storage systems that connect to IBM i
- Steps for sizing and modelling Storage systems for IBM i
- Disk Magic
- Collecting IBM i performance data
- Decide the peaks
- Sizing - decide the initial storage configuration for Disk Magic
- Modelling with Disk Magic
- Demo of Disk Magic

## IBM i Single level storage, Handling I/O operations

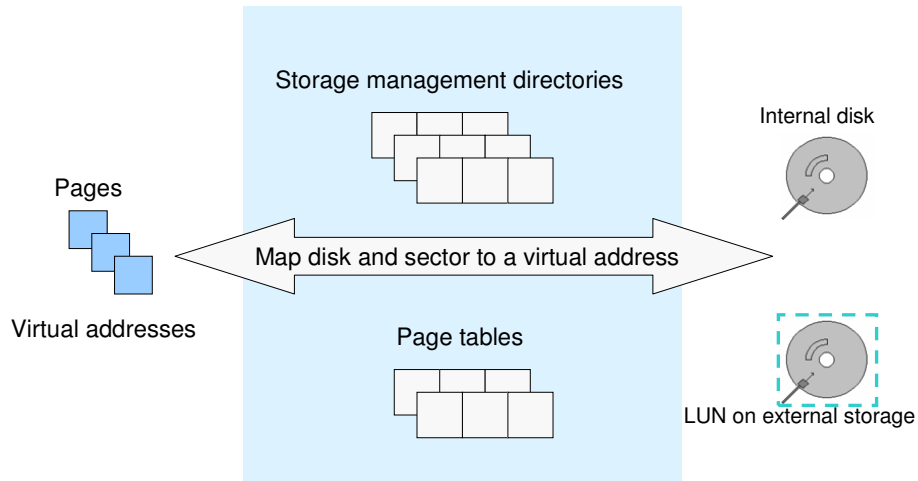


## Spreading IBM i objects across disk units - LUNs

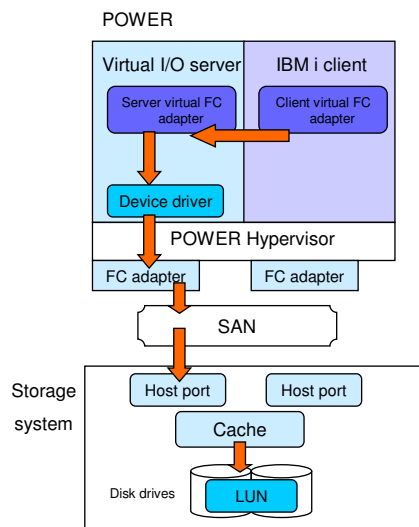


IBM i storage management spreads IBM i objects evenly across the available disk units or LUNs


## Storage Manager – Handling IO operations



## IO flow – Storage system connected with VIOS-NPIV

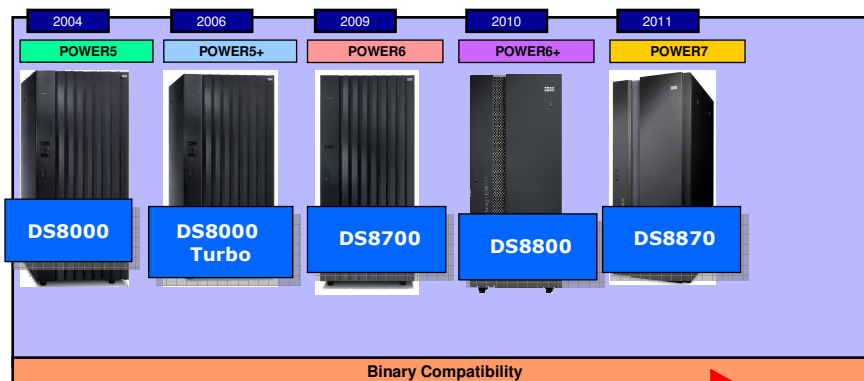


## IBM i and Storage systems



## DS8000 Enterprise disk system

The IBM POWER processor has been behind the success of IBM enterprise storage beginning with the Enterprise Storage Server in 1999



## DS8870 models

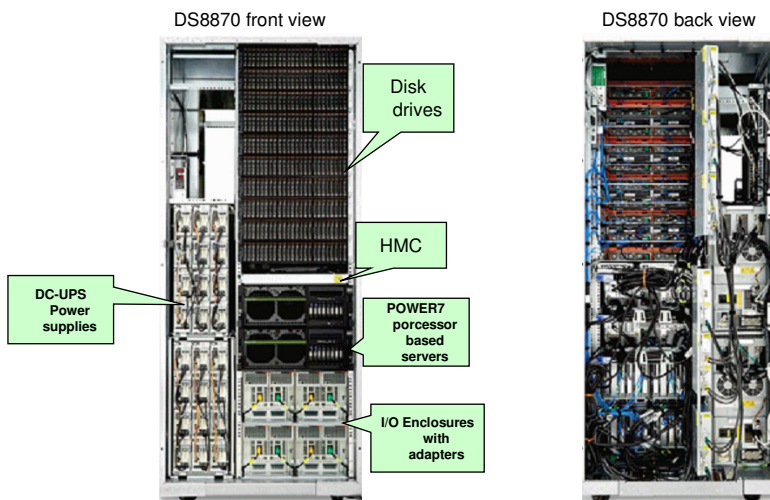


- **DS8870 Model 961 Enterprise model**
  - Dual 4-core, dual 8-core, or dual 16-core processor complex
  - 64 GB to 1 TB cache
  - 8 Gb/s host and device adapters
  - 2.5" Enterprise SAS-2 drives

- **DS8870 Model 961 Business Class model**
  - Dual 2-core processor complex
  - 16 GB or 32 GB cache
  - 8 Gb/s host and device adapters
  - 2.5" Enterprise SAS-2 drives

- **DS8870 Model 96E Expansion frame**
  - 8 Gb/s host and device adapters
  - 2.5" Enterprise SAS-2 drives

## DS8870 Components



## Disk Drives in DS8870

- 146 GB 15 K RPM – Small Form Factor (SFF)
- 300 GB 15 K RPM - SFF
- 600 GB 10 K RPM - SFF
- 900 GB 10 K RPM - SFF
- 3 TB 7.2 K RPM nearline SAS drives – Large Form Factor (LFF)
- 400 GB SSD
- RAID support:
  - SFF disk drives support RAID-5, RAID-6 and RAID-10
  - SSD support RAID-5
  - LFF disk drives ( nearline disk drives ) support RAID-6
  - RAID protection is enabled on Array site of 8 disk drives



SFF 24 disk drive enclosure – front view

## DS8870 Graphical User Interface (GUI)

### Overview

```

            graph LR
            Drives --> Extent Pools
            Extent Pools --> FB_Volumes[FB Volumes]
            Extent Pools --> OKD_LUNs[OKD LUNs and Volumes]
            FB_Volumes --> Volume_Groups[Volume Groups]
            OKD_LUNs --> Volume_Groups
            Volume_Groups --> Hosts
            
```

**Overview ( Version 5.6.1.1409 )**

**Watch e-Learnings:** [Overview](#)

The diagram above represents all of the objects that need to be configured. To learn more about each object, click the icon in the diagram. For some objects, e-Learning modules include a tutorial of the steps that are required to complete the task. To configure these objects, use the icons in the left navigation.

Before you begin storage configuration, please perform the initial setup tasks.


[Visit the Information Center](#)

**Initial Setup Tasks**

- **Activate your licenses**  
Go to **System Status** and select Apply Activation Codes in the action menu.
- **Create users and groups**  
Go to **Users** to set access.
- **Set up LDAP authentication**  
Go to **Remote Authentication** to set the authentication policy for LDAP user accounts.
- **Set up Encryption**

## IBM Storwize V7000 – Responding to Our Customers

- Requests for SAN Volume Controller
  - SVC with internal disk drives
    - Robust and reliable virtualized storage product with RAID integration and simplified system management (only one interface for everything)
  - Hide SVC virtualization complexity
  - Easy to use graphical user interface
- Requests for midrange disk
  - Easy system management
  - Enterprise attributes
  - Simplified data migration
  - Robustness, high availability
  - Strong interoperability
  - Application integration
  - Low cost



## IBM Storwize V7000 at a Glance



### Enterprise level software capabilities

- RAID 0,1,5,6,10
- Storage virtualization
- Non disruptive data migration
- New user interface
- Thin provisioning
- Easy Tier
- FlashCopy
- Remote copy
- Disaster recovery automation

### Enclosure front View

12 x 3.5" LFF Drive Bays Model 2076-112/212/312



24 x 2.5" SFF Drive Bays Model 2076-124/224/324




## Storwize V7000 Drive Options

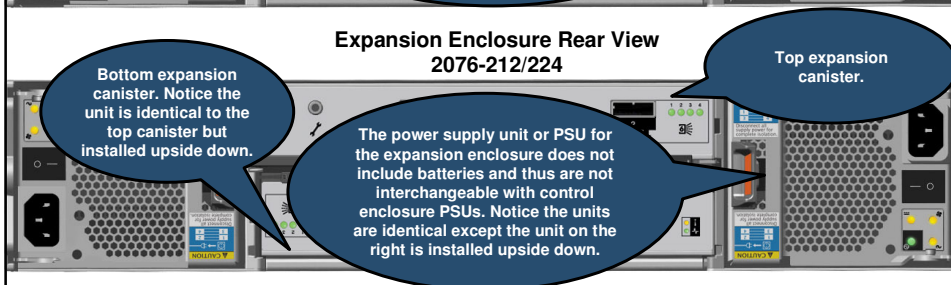
Drive Type		Drive Sizes
Model 2076 124 / 224 / 324:  2.5" Small Form Factor (SFF)	<b>SSD</b>	<b>300GB (MLC) 200/400GB (MLC)</b>
	<b>10K RPM SAS</b>	<b>300, 450, 600GB 900GB</b>
	<b>15K RPM SAS</b>	<b>146GB 300GB</b>
	<b>7.2K RPM NL-SAS</b>	<b>1TB 1.2TB</b>
No restrictions on mixing of drive types within the same enclosure at this time		
Model 2076 112 / 212 / 312:  3.5" Large Form Factor (LFF)	<b>7.2K RPM NL-SAS</b>	<b>2TB 3TB 4TB</b>

## Storwize V7000 Components

**Control Enclosure Rear View**  
2076-112/124/312/324




**Expansion Enclosure Rear View**  
2076-212/224





# Storwize V7000 Components

### Control Enclosure Rear View 2076-112/124/312




Two 24Gbps SAS chains for expansion enclosure attachment are available per node canister.

Four 8Gbps fibre channel ports are available per node canister.

Two gigE ports and two USB ports are available per node canister.

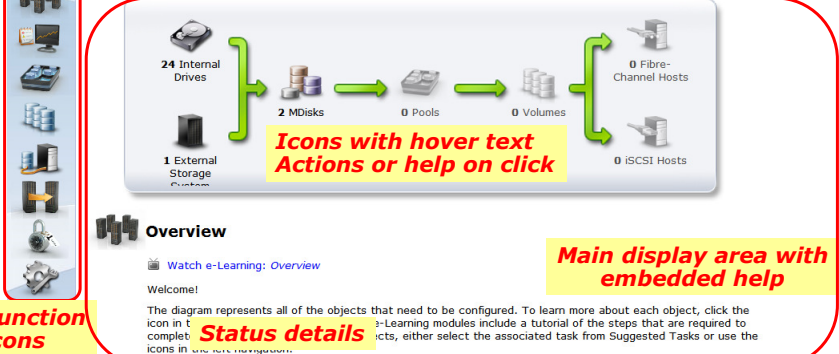
### Expansion Enclosure Rear View 2076-212/224



Two 24Gbps SAS ports are available per expansion canister for connection to other enclosures.

# Storwize V7000 GUI


Navigation, user indicator



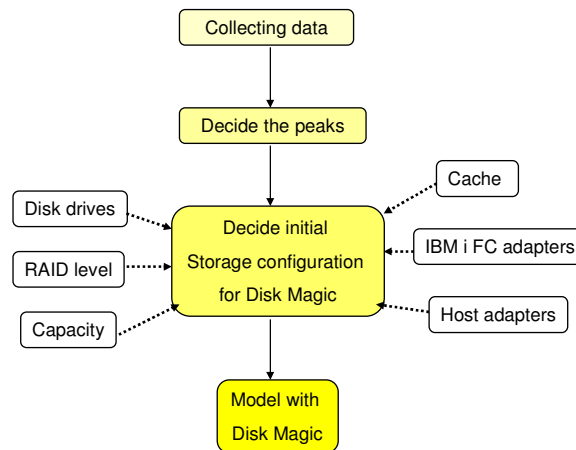
Function icons

Status details

Status indicators

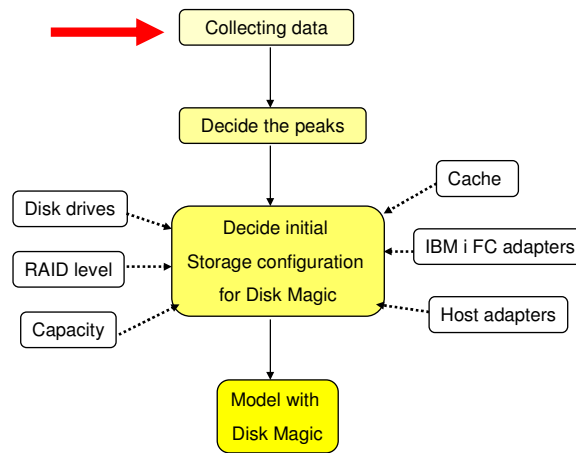


## Sizing and modelling External storage for IBM i



## Disk Magic

- Modelling tool to help estimate disk subsystem performance
- Developed and maintained by the company IntelliMagic
- Models IBM, HP, HDS and EMC storage
- IntelliMagic closely collaborates with IBM development and performance teams
- Standard tool for sizing and modelling IBM Storage systems
- Supports modeling of IBM i with Storage systems since 2003
- 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>
- Customers can purchase the Disk Magic license from Intellimagic



## IBM i performance data for sizing and Disk Magic modelling

- Collection Services data
  - Collected with 5 minutes intervals
  - Collected during 3 consecutive days
  - Collected also during intensive end-of-month jobs, if needed
- Needed Performance reports
  - System report / section Disk Utilization
  - Resource report / section Disk Utilization
  - Component report / section Disk Activity
- Performance reports must be in English language
- Performance reports must be transfered in .txt form to the PC with Disk Magic
- If a particular period of the day is important for the customer, create reports containing data of this period only



## Collecting IBM i Collection Services data with Performance Tools (5770-PT1)

The screenshots illustrate the process of collecting performance data:

- PERFORM**: Selecting '2. Collect performance data' from the main menu.
- Collect Performance Data**: Selecting '2. Configure Performance Collection' to set up the collection.
- Configure Perf Collection (CFGPFRCOL)**: Configuring parameters such as 'Default interval' (05:00), 'Default collection profile' (\*STANDARD), and 'Collection retention period' (00:00).
- Collect Performance Data**: Selecting '1. Start Performance Collection' to begin data collection.

© 2013 IBM Corporation



## Creating reports using Performance Tools

Example: Creating System report / section Disk utilization for a period 12:00 to 12:30

The screenshots show the steps to generate a report:

- Print Performance Report - Sample data**: Viewing a list of system members and their status.
- Select Sections for Report**: Choosing '1' for 'Disk Utilization' from the available sections.
- Select Categories for Report**: Selecting '1' for 'Time Interval' to specify the report's time range.
- Select Time Intervals**: Viewing a table of time intervals and selecting the 12:00 to 12:30 interval.

Date	Time	Transaction Count	Resp	CPU Util	Int	High	Pool
08/29	11:50	23	01	0	0	2	4002
08/29	11:55	11	00	0	0	1	4002
08/29	12:00	0	00	0	0	0	10001
08/29	12:05	0	00	0	0	0	10005
08/29	12:10	0	00	0	0	0	10003
08/29	12:15	0	00	0	0	0	10003
08/29	12:20	0	00	0	0	0	10003
08/29	12:25	0	00	0	0	0	10002
08/29	12:30	0	00	0	0	0	10002
08/29	12:35	0	00	0	0	0	10002

© 2013 IBM Corporation

## Collecting IBM i Collection Services data without Performance Tools

- CFGPFRCOL (Configure Performance Collection)
- STRPFRCOL (Start Performance Collection)
- ENDPFRCOL (End Performance Collection)

Example: Configuring performance collection with 5 minutes intervals

```

Configure Perf Collection (CFGPFRCOL)
Type choices, press Enter.
Default interval . . . . . 05_00_   *SAME, .25, .50, 1.0, 5.0...
Collection library . . . . . QPFRDATA  Name, *SAME
Default collection profile . . . . . *STANDARDP  *SAME, *MINIMUM, *STANDARD...
Cycle time . . . . . 000000         Time, *SAME
Cycle interval . . . . . 24          *SAME, 1-24 hours
Collection retention period:
  Number of units . . . . . 00120    *SAME, 1-720, *PERM
  Unit of time . . . . . *HOURS      *HOURS, *DAYS
Create database files . . . . . *YES_  *SAME, *YES, *NO

Bottom
F3=Exit  F4=Prompt  F5=Refresh  F10=Additional parameters  F12=Cancel
F13=How to use this display  F24=More keys
    
```

## Creating reports without Performance Tools

- PRSYSRPT ( Print System Report)
- PRTRSCRPT (Print Resource Report)
- PRTCPTRPT (Print Component Report)

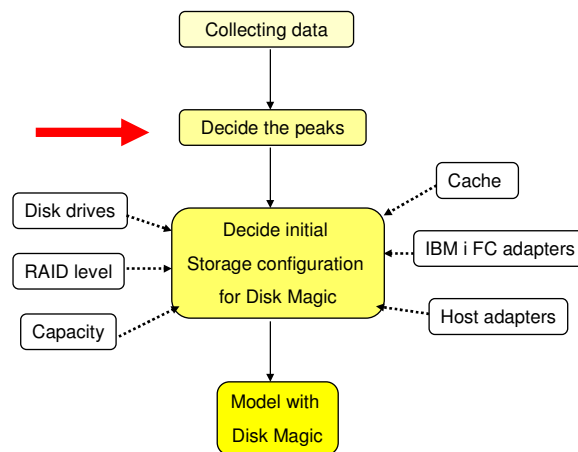
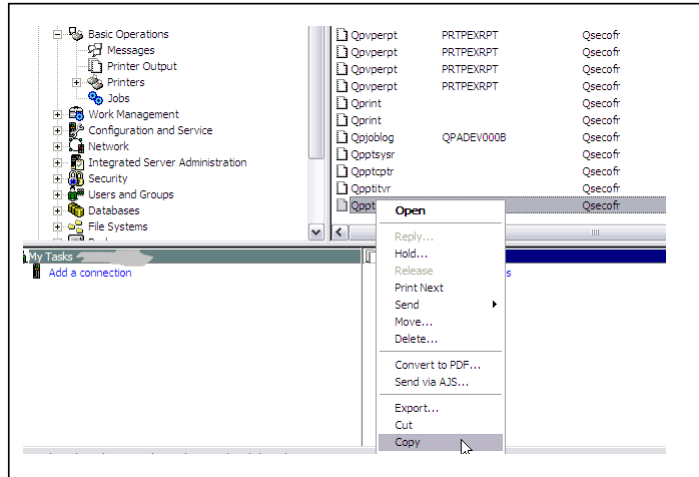
Example: Creating System report / Disk utilization from 4. November 00:00 to 8:00

```

Print System Report (PRSYSRPT)
Type choices, press Enter.
Member . . . . . > Q308000002  Name
Report title . . . . . Test_system_report_disk_utilization_
Time period for report:
Starting time . . . . . > 000000_   Time, *FIRST, *SELECT
Starting date . . . . . > 110413_   Date, *FIRST
Ending time . . . . . > 080000_   Time, *LAST
Ending date . . . . . > 110413_   Date, *LAST
Additional Parameters
Library . . . . . QPFRDATA_   Name
Type of information . . . . . > *DISK  *ALL, *WORKLOAD, *RSC...
+ for more values
F3=Exit  F4=Prompt  F5=Refresh  F12=Cancel  F13=How to use this display  More...
F24=More keys
    
```

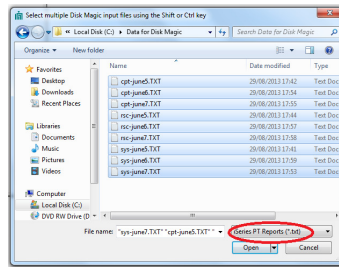


## Transfer the reports to .txt file on your PC

Use IBM i Operations Navigator

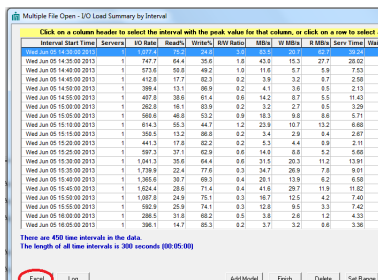






# Using Disk Magic spreadsheet - 1



# Using Disk Magic spreadsheet - 2







Interval	Start Time	I/O Rate	Read %	Write %	R/W Ratio	Total MB/s	Write MB/s	Read MB/s
2	Jun05:2013 06:40:00	142.6	7.5%	92.2%	0.1	1.8	1.7	0.1
3	Jun05:2013 06:45:00	118.7	11.6%	89.0%	0.1	1.9	1.7	0.2
4	Jun05:2013 06:50:00	120.7	3.6%	96.4%	0.0	1.2	1.2	0.0
5	Jun05:2013 06:55:00	176.5	16.2%	83.8%	0.2	2.1	1.7	0.3
6	Jun05:2013 07:00:00	209.1	28.1%	71.9%	0.4	2.6	1.9	0.7
7	Jun05:2013 07:05:00	223.3	30.9%	69.1%	0.4	2.4	1.7	0.7
8	Jun05:2013 07:10:00	145.0	7.6%	93.0%	0.1	1.7	1.6	0.1
9	Jun05:2013 07:15:00	161.0	9.9%	90.1%	0.1	1.9	1.6	0.2
10	Jun05:2013 07:20:00	154.5	10.2%	89.7%	0.3	1.7	1.3	0.4
11	Jun05:2013 07:25:00	152.5	8.5%	91.5%	0.1	1.8	1.6	0.2
12	Jun05:2013 07:30:00	141.4	7.8%	92.2%	0.1	1.7	1.6	0.1
13	Jun05:2013 07:35:00	149.2	7.3%	92.7%	0.1	1.4	1.3	0.1
14	Jun05:2013 07:40:00	141.4	7.8%	92.2%	0.1	1.7	1.6	0.1
15	Jun05:2013 07:45:00	162.0	10.5%	89.5%	0.1	1.6	1.4	0.2
16	Jun05:2013 07:50:00	137.8	11.7%	88.2%	0.1	1.4	1.3	0.2
17	Jun05:2013 07:55:00	219.5	24.6%	76.6%	0.3	2.6	2.0	0.6
18	Jun05:2013 08:00:00	344.8	47.5%	52.1%	0.9	7.5	3.9	3.6
19	Jun05:2013 08:05:00	341.9	21.2%	83.8%	0.3	4.7	3.7	1.0
20	Jun05:2013 08:10:00	260.1	26.5%	73.5%	0.4	4.9	3.6	1.3
21	Jun05:2013 08:15:00	202.2	25.2%	75.0%	0.3	3.9	2.9	1.0
22	Jun05:2013 08:20:00	318.3	26.6%	73.4%	0.4	2.4	1.8	0.6
23	Jun05:2013 08:25:00	398.9	36.2%	69.9%	0.6	4.6	2.9	1.6
24	Jun05:2013 08:30:00	453.1	49.5%	59.5%	1.0	19.0	9.6	9.4
25	Jun05:2013 08:35:00	564.0	54.4%	45.6%	1.2	22.6	16.3	16.3
26	Jun05:2013 08:40:00	286.4	9.2%	90.8%	0.1	4.5	4.1	0.4
27	Jun05:2013 08:45:00	332.1	29.9%	71.0%	0.4	6.3	4.5	1.8
28	Jun05:2013 08:50:00	407.8	46.1%	37.8%	1.6	17.4	6.6	10.8





## Decide for the peaks - 1



© 2013 IBM Corporation



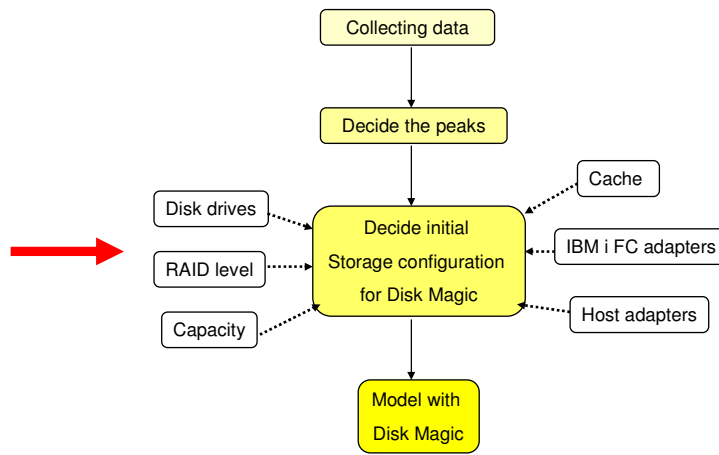
## Decide for the peaks - 2

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
Interval Date	Interval Start	IO Rate	Read %	Write %	R/W Ratio	Total MB/s	Write MB/s	Read MB/s	Serv Time	Wait Time	KB/Write	KB/Read	W SerTi	R SerTi	Reads/sec	Writes/sec	
Jun/05/2013	06:40:00	142.6	7.8%	92.2%	0.1	1.8	1.7	0.1	2.0	7.9	12.9	12.9	0.0	0.0	11.146	131.446	
Jun/05/2013	06:45:00	176.7	11.0%	89.0%	0.1	1.9	1.7	0.2	2.7	7.4	10.9	10.9	0.0	0.0	19.416	157.329	
Jun/05/2013	06:50:00	120.7	3.6%	96.4%	0.0	1.2	1.2	0.0	1.6	6.8	10.1	10.1	0.0	0.0	4.365	116.366	
Jun/05/2013	06:55:00	176.5	16.2%	83.8%	0.2	2.1	1.7	0.3	2.3	5.6	12.1	12.1	0.0	0.0	28.54	147.916	
Jun/05/2013	07:00:00	209.1	28.1%	71.9%	0.4	2.6	1.9	0.7	3.2	5.2	12.6	12.6	0.0	0.0	58.697	150.442	
Jun/05/2013	07:05:00	233.3	30.9%	69.1%	0.4	2.4	1.7	0.7	3.3	4.3	10.6	10.6	0.0	0.0	72.19499	161.104	
Jun/05/2013	07:10:00	145.0	7.0%	93.0%	0.1	1.7	1.6	0.1	1.8	7.9	12.3	12.3	0.0	0.0	10.221	134.776	
Jun/05/2013	14:25:00	610.8	42.9%	57.1%	0.8	14.6	8.3	6.2	5.8	2.5	24.4	24.4	0.0	0.0	262.124	348.726	
Jun/05/2013	14:30:00	1077.4	75.2%	24.8%	3.0	83.5	20.7	62.7	39.2	9.4	79.3	79.3	0.0	0.0	809.7151	267.728 Peak 2	
Jun/05/2013	14:35:00	747.7	64.4%	35.6%	1.8	43.0	15.3	27.7	28.0	7.3	58.9	58.9	0.0	0.0	461.3979	266.285	
Jun/07/2013	11:30:00	1457.8	22.3%	77.7%	0.3	29.4	22.9	6.5	5.8	4.4	20.7	20.7	0.0	0.0	324.439	1133.39	
Jun/07/2013	11:35:00	2200.0	11.7%	88.3%	0.1	29.9	26.4	3.5	3.7	1.4	13.9	13.9	0.0	0.0	257.613	1942.373 Peak 1	
Jun/07/2013	11:40:00	2473.3	21.9%	78.1%	0.3	41.7	32.6	9.1	5.2	1.7	17.3	17.3	0.0	0.0	540.625	1932.707	
Jun/07/2013	11:45:00	395.0	35.0%	65.0%	0.5	9.4	6.1	3.3	8.9	4.0	24.3	24.3	0.0	0.0	138.31	256.712	
Jun/07/2013	11:50:00	202.5	11.8%	88.2%	0.1	1.9	1.7	0.2	2.0	4.7	9.7	9.7	0.0	0.0	23.94	178.571	

**Recommend: Size for both peaks**


© 2013 IBM Corporation





## Hard Disk Drives

- Requirements for capacity
- Guidelines for disk arms





## Guidelines for RAID level

- RAID-10 provides better resiliency
- RAID-10 provides generally better performance:
  - RAID-5 results in 4 disk operations per write – higher penalty
  - RAID-10 results in 2 disk operations per write – lower penalty
- RAID-10 requires more capacity
- In DS8000 use RAID-10 when:
  - There are many random writes
  - Write cache efficiency is low
  - Huge workload
- In Midrange storage and Storwize V7000 use RAID-10 when:
  - There are more than 50% writes



## Guideline - Number of DS8000 ranks of 8 disk drives for an IBM i workload

<b>Max host IO/sec for HDD rank wit IOP-less adapters</b>	<i>Host iops at 70% Read</i>	<i>Host iops at 50% Read</i>
<b>SAS DDM</b>		
RAID-5 15 K RPM	940	731
RAID-10 15 K RPM	1253	1116
RAID-6 15 K RPM	723	526

### Example:

- IBM i workload with 10000 IO/sec, 70% reads, we are planning 15 K RPM disk drives in RAID-5
- Calculation for the ranks:  $10000 / 940 = \text{app } 10 \text{ ranks}$

## Guideline - Number of V7000 disk drives for an IBM i workload

Quick calculation for the number of disk drives for an IBM i workload

<b>Max hosts IO/ sec per disk drive</b>	<i>Host iops at 70% Read</i>	<i>Host iops at 50% Read</i>
<i>15 K RPM disk drive</i>		
RAID-1 or RAID-10	138	122
RAID-5	96	75
<i>10 K RPM disk drive</i>		
RAID-1 or RAID-10	92	82
RAID-5	64	50

### Example:

- IBM i workload with 3000 IO/sec, 50% reads, we are planning 15 K RPM disk drives in RAID-10
- Calculation for the disk drives:  $3000 / 122 = \text{app } 24$  disk drives in RAID-10

## Guidelines for cache size in external storage

- To understand the needed cache size determine the type of workload
  - Random or sequential
  - Read/write ratio
  - Is the workload cache friendly
- Modelling with Disk Magic
- Rough guidelines for DS8800
  - 10 to 20 TB capacity: 64GB cache
  - 20 to 50 TB: 128 GB cache
  - > 50 TB: 256 GB to 1TB cache
- V7000:
  - Fixed cache size of 16 GB per node pair

## LUNs – number and size

- The more LUNs the better performance
- Make sure that sufficient disk arms are provided to support the LUNs
- Guideline for sizes: 2 to 4 LUNs per disk drive
- Use one size of LUNs ( or minimum 2 sizes) for an IBM i ASP
- Consider minimal capacity of LoadSource
  - Minimal size for LoadSource in IBM i V7.1 is 17 GB
  - We expect bigger minimal size in future releases
- Best practice
  - Minimal size about 40 GB
  - Maximal size about 150-200 GB
  - Minimal number of LUNs per IBM i LPAR: 6
- Consider SCSI command tag queuing Queue depth (max. concurrent I/O operations to a LUN)

## Queue depth for IBM i LUNs

<b>SCSI command tag queuing - queue depth</b>	LUN in DS8000	LUN in SVC / V7000
Native connection with IOP-based adapters	1	n/a
Native connection with IOP-less adapters	6	16
VIOS_NPIV connection	6	16
VIOS vscsi connection	32	32

## Sizing FC adapters in IBM i or in VIOS

<b>Guideline for IO/sec per port / GB per port / MB/sec per port</b>	<b>4 Gb adapters, FC 5774/5276</b>	<b>8 Gb adapters, FC 5735/5273</b>
IO/sec at 70% utilization	10500	12250
GB per port for native connection Assumed: Access Density = 1.5	2800 With Dual path: 5600 per two ports	3266 With Dual path: 6532 per two ports
Sequential workload MB/sec at 70% utilization		122 per port, if one port in adapter is used 104 per port, if both ports are used
Transaction workload MB/sec at 70% utilization	175	

## Sizing host adapters in external storage

### *Rules of thumb for HBAs in DS8800*

- About 4 to 8 \* ports in IBM i per one HA card in DS8800
- For good performance: The number of HA cards should be the same or bigger than the number of device adapters in DS8800
- At least one HA card per IO enclosure in DS8800

### *Guideline for V7000 host ports*

- Maximum 50 000 IO/sec per host port
- Maximum 300 MB/sec per host port
- Native connection with SAN switches or VIOS\_NPIV connection:
  - Zone one IBM i port with two V7000 ports each from one node canister
  - For multipath use two or more IBM i ports each zoned this way Recommended: use ports from different adapters for multipath
- Native connection without SAN switches:
  - For resiliency: assign the LUNs to two IBM i ports each connected to different V7000 node canister
  - For Multipath: Use two IBM i adapters. Connect one port from each IBM i adapter to node canister 0, and the other port to canister 1. Assign the LUNs to all 4 IBM i ports.

## Data layout

- Spreading workloads across all components and resources of the storage system, maximizes the utilization of the hardware components. However, it is always possible when sharing resources that performance problems may arise due to contention on these resources.
- To protect critical workloads you should isolate them minimizing the chance that non-critical workloads can impact the performance of critical workloads. Isolation of workloads is most easily accomplished where each ASP or LPAR has its own managed storage pool. This ensures that you can place data where you intend.
- I/O activity should be balanced between the two DS8000 processor complexes, or the two nodes or controllers on the SVC / Storwize V7000.
- Use only IBM i LUNs on any storage pool (rather than mixed with non-IBM i).
- If you mix production and development workloads or many small workloads in storage pools make sure that the customer understands that this may impact production performance.

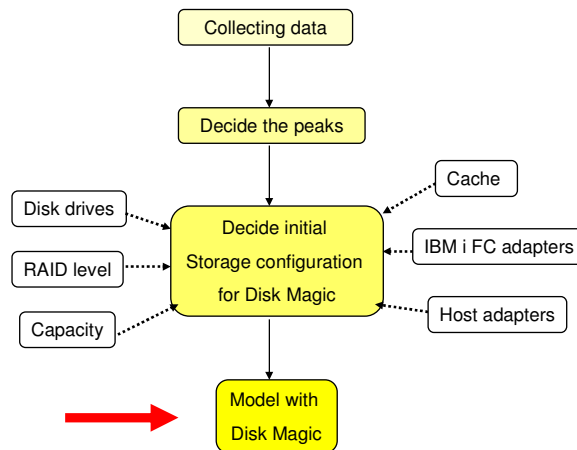

## Guidelines for VIOS resources

- Recommended: Use IBM i Workload Estimator to estimate needed VIOS resources
- Rule of thumb: 1 – 2 CPU per VIOS
- Rule of thumb: 8 GB memory in VIOS
- FC adapters in VIOS:
  - For 4Gb adapters calculate about 10500 IO/sec per port
  - For 8Gb adapters calculate about 12250 IO/sec per port
- With VIOS\_NPIV
  - Maximum one virtual FC adapter from one IBM i LPAR can be mapped to a port in adapter in VIOS
  - Up to 64 virtual FC adapters, each from different IBM i LPAR can be mapped to the same port in adapter in VIOS
- To calculate the number of adapters in VIOS you can also use the measurements of MB/sec listed in foil 41

## Solid State Drives with IBM i

- Hybrid disk pool and Easy Tier
- An IBM i ASP with application data resides on SSD
- The IBM i skew level is typically flat due to IBM i object orientated architecture
- Therefore IBM i might need some more SSD in the hybrid pool than other workloads
- The IBM i methods for data relocation can be used with DS8000, but can not be used with V7000

Typical IBM i Skew level and STAT

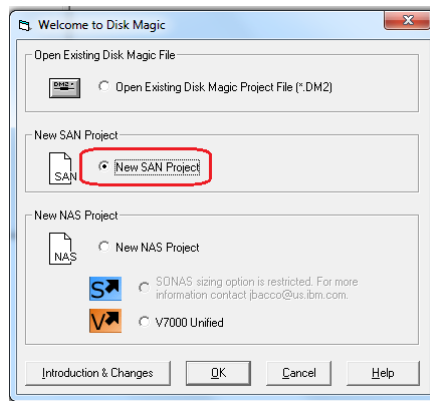


## Modelling with Disk Magic – inserting IBM i reports -1

Start Disk Magic




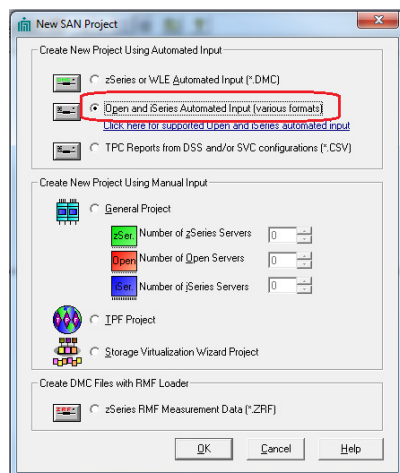
Select New SAN Project



## Modelling with Disk Magic – inserting IBM i reports -2

Select Open iSeries Automated ...

- Select the folder with IBM i reports
- Click Open

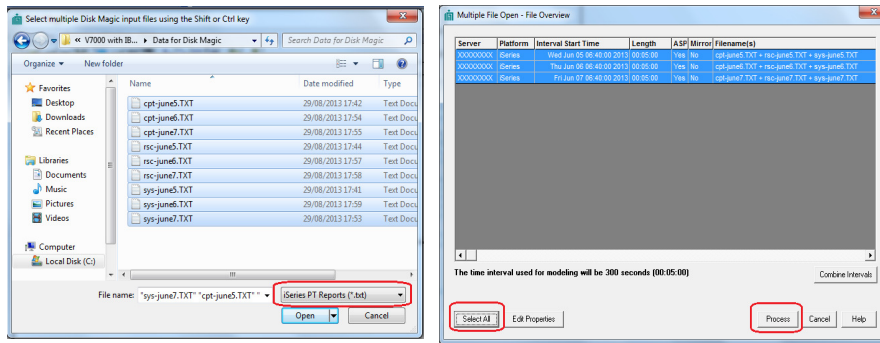




## Modelling with Disk Magic – inserting IBM i reports -3

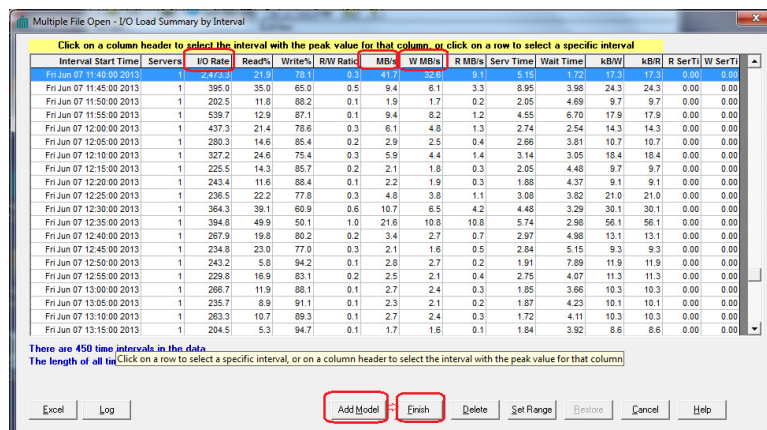
- Select “iSeries PT Reports”
- Select the reports to use for modelling

Use “Select all” and “Process”



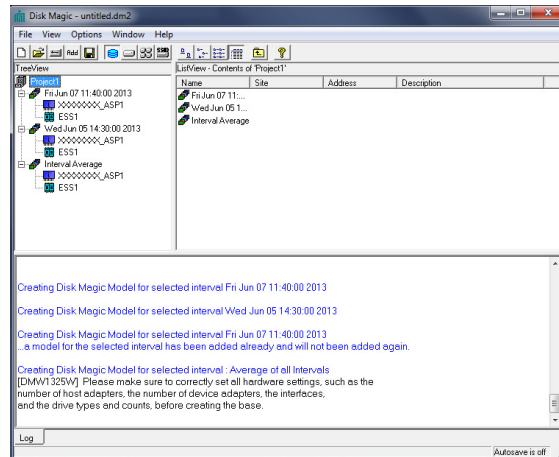
## Modelling with Disk Magic – selecting the peaks

- Select a peak by clicking the header of the column, click Add Model for each peak
- We recommend to select the peaks by IO/sec, MB/sec, Write MB/sec and Average
- After the models are added, click Finish



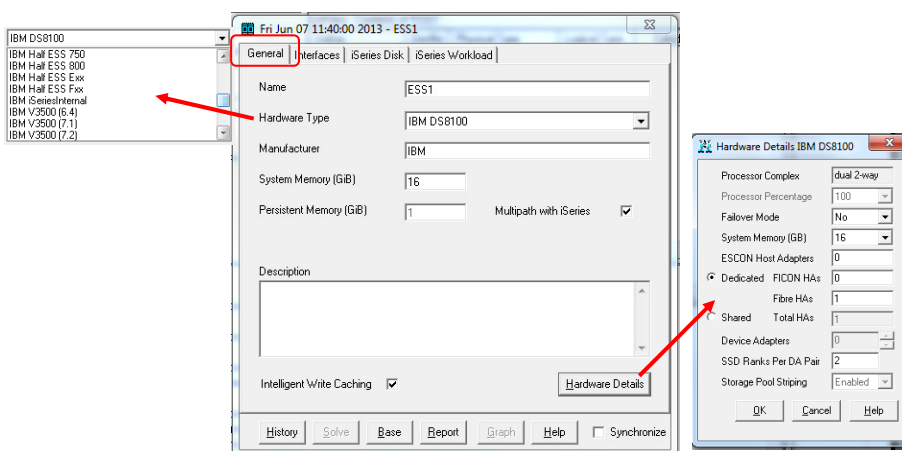
## Modelling with Disk Magic – start modelling ext. Storage

- Disk Magic creates a model for each peak
- Click on ESS1 to start modelling external storage for a particular peak



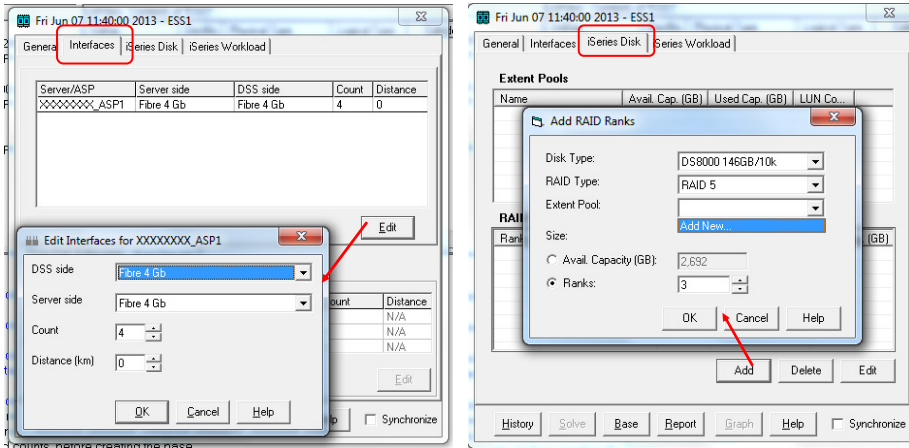
## Modelling with Disk Magic – Base (present conf.) -1

- Insert the current configuration of external storage or internal disks



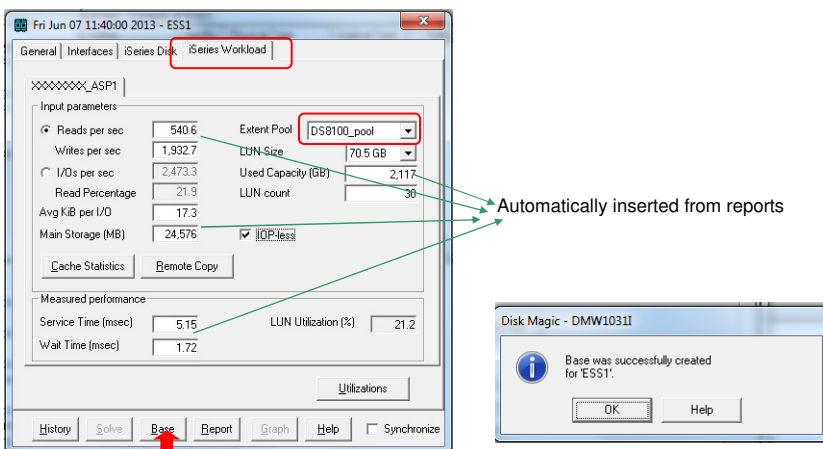
## Modelling with Disk Magic – Base - 2

- On interface tab insert the number and type of adapters used
- On iSeries Disk tab insert the number and type of presently used disk drives, create an extent pool



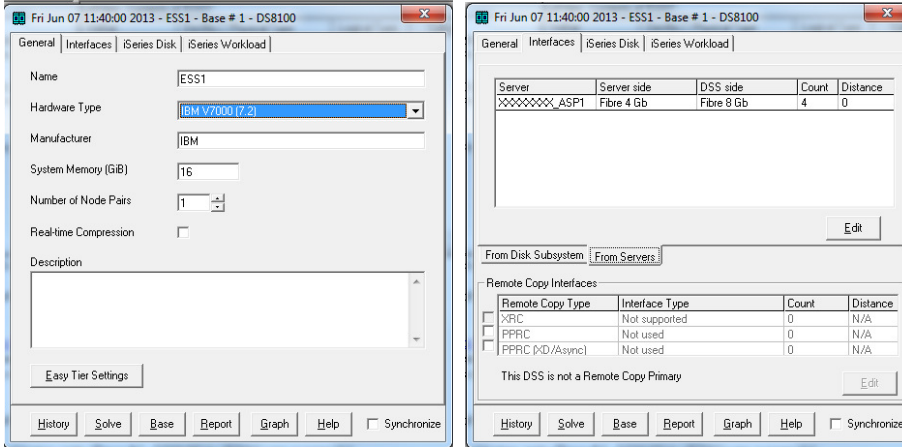
## Modelling with Disk Magic – Base - 3

- On iSeries Workload tab select the created extent pool, specify the type of adapters, Base
- Note: Service Time and Wait time for the base are automatically inserted from IBM i reports



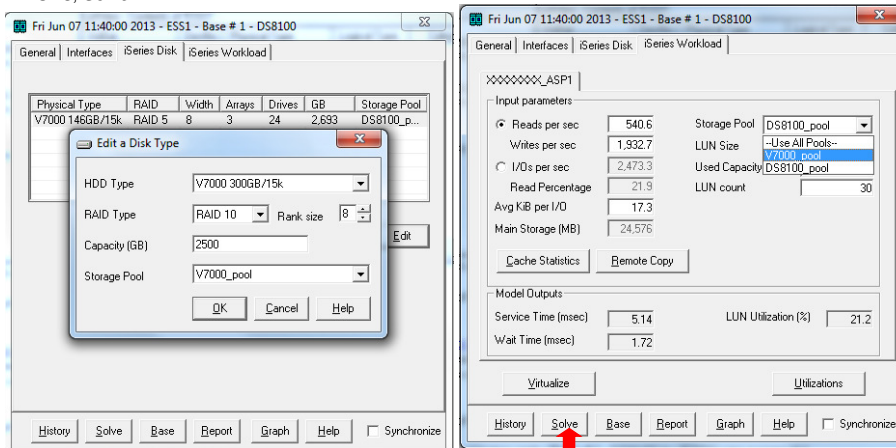
## Modelling with Disk Magic – new configuration - 1

- On General tab insert the storage system you want to model
- On Interface tab specify the number and type of adapters in IBM i and in storage system



## Modelling with Disk Magic – new configuration - 2

- On iSeries Disk tab choose the disk type, RAID level and capacity of the new configuration and add new disk pool
- On iSeries Workload tab select the pool for new configuration, if needed change the number or size of LUNs, **Solve**



## Modelling with Disk Magic – new configuration - 3

- At solving the model, Disk Magic calculates predicted response times.
- Compare present and predicted Service time and Wait time

© 2013 IBM Corporation

## Modelling with Disk Magic – utilizations, history

- Click Utilizations to show the utilizations of modelled configuration. Utilizations that exceed the threshold are marked in red
- Clicking History shows the modelled configurations. You may rename them with more descriptive names. Restore any configuration you want to look to, or to change and solve again.

© 2013 IBM Corporation


## Modelling with Disk Magic – modeling growth - 1

- Click *Graph*
- Choose the data to model growth for, graph type and the range for growth, *Plot*

The screenshot shows the Disk Magic software interface. The main window has several tabs: General, Interfaces, iSeries Disk, and iSeries Workload. The 'iSeries Workload' tab is active, showing input parameters for 'ASPI1'. The 'Graph Options' dialog box is open, allowing the user to configure the graph. The 'Graph Data' dropdown is set to 'Service Time in ms'. The 'Range Type' is 'I/O Rate' with a range from 2473.3 to 4000 by 500. The 'Graph' button in the main window is highlighted with a red arrow.

## Modelling with Disk Magic – modeling growth - 2

- Predicted service time and utilizations with IO growth



Utilization Overview (iSeries)	Total I/O Rate (I/Os per second)					
	Amber Threshold	Red Threshold	2473	2973	3473	3973
Processor (for I/O)	60%	80%	1.3%	1.5%	1.8%	2.0%
Processor (for Compression)	60%	80%	0.0%	0.0%	0.0%	0.0%
Average Bus	60%	80%	1.6%	1.9%	2.2%	2.6%
Highest HDD	50%	60%	12.7%	14.4%	16.0%	17.6%
Average HA	60%	80%	1.1%	1.3%	1.5%	1.8%
Back End Interface	60%	80%	2.9%	3.5%	4.1%	4.7%
Average Host Interface	60%	80%	1.1%	1.3%	1.5%	1.8%



## **DEMO**

# **Disk Magic modelling of V7000 for IBM i**



## **Backup charts**

## Response times in IBM i

- **Application response time:**
  - The response time of an application transaction. This time is usually critical for the customer.
- **Duration of batch job:**
  - Duration of batch job: Batch jobs usually run during the night; the duration of a batch job is critical for the customer, because it must be finished before regular daily transactions start.
- **Disk response time: service time + wait time**
  - **Service time:** processing I/O operation
  - **Wait time:** potential I/O queuing on the IBM i host

## Disk drives – additional considerations and best practise

- **Number of disk drives in native, VIOS\_NPIV and VIOS VSCSI connection**

The sizing guidelines and calculations for disk drives are the same regardless if the host connection is native, VIOS\_NPIV or VIOS VSCSI.
- **Sizing for big block sizes (transfer sizes)**

By experiences, big block sizes have significant impact on performance. Therefore we recommend to size with additional 25% - 30% of disk drives.
- **Implementation recommendations:**
  - create vdisks in Striped mode ( default )
  - extent size 256 MB ( default )



## Sizing for Metro Mirror or Global Mirror links

- Obtain the write rate in MB/sec
  - Use queries to IBM i collection services data
  - Or, use performance reports and calculate writes/sec \* blksize
  
- Based on highest write rate calculate needed bandwidth as follows:
  - Assume 10 bits per byte for network overhead
  - Assume a maximum 80% utilization of the network
  - Apply 10% uplift factor to the result to account for peaks in the 5 minutes intervals
  
- If the compression of devices for remote links is known you may apply it. If it is not known you may assume a 2:1 compression

## Sizing for FlashCopy

- GUI preset Snapshot – Flashcopy without background copying
- GUI preset Clone - Flashcopy with background copying
- Reasons for potential performance impact on production workload during FlashCopy relation:
  - At the beginning, every write operation to source volume or to target volume triggers copy from source to target; later copying usually decreases due to writes to the areas that have been already overwritten
  - A certain amount of read operations to target volume is done from source
- To minimize performance impact on production system it is important to provide sufficient disk drives to the FlashCopy target
- Rough guideline for the number of disk drives for FlashCopy target:
  - Calculate 150 writes/sec per 10 K RPM disk drive
  - Calculate 100 writes/sec per 15 K RPM disk drive
- With Thin provisioned FlashCopy target volumes, ensure sufficient capacity in the disk pool
  - The more write operations are done to either source or target LUNs, the more capacity is needed



# **DEMO**

## **Disk Magic modelling of V7000 for IBM i**