Introducing VIOS Investigator

Tom McBride, Ron McCargar
IBM i Global Support Center

What we heard from the LUG
We Want Flexibility

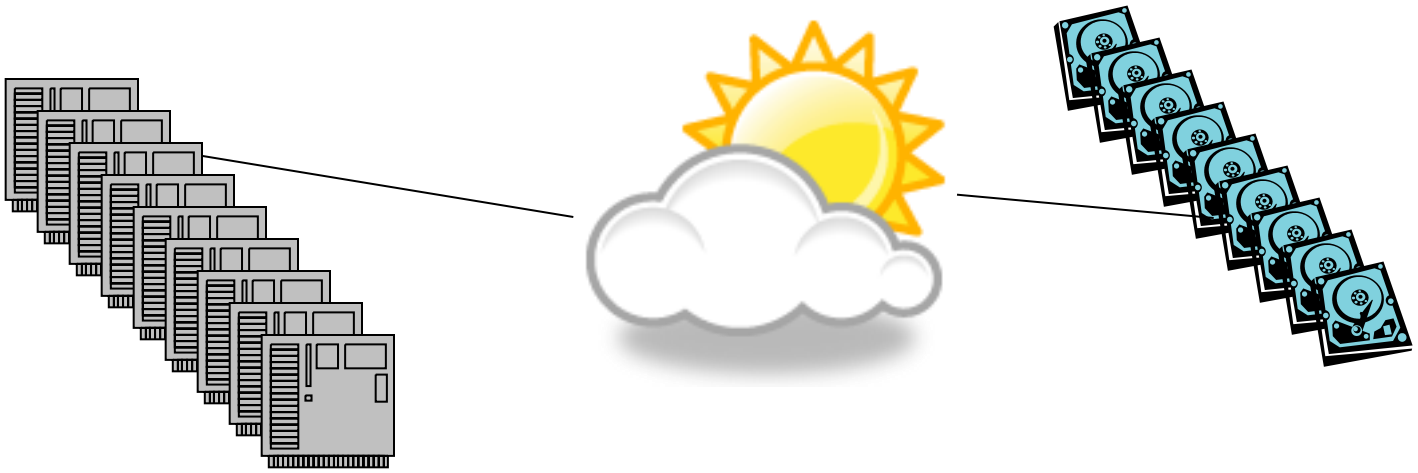

Now a problem occurs

Complexity leaves us wondering
Where is the problem? Which drive is which? What is going on?

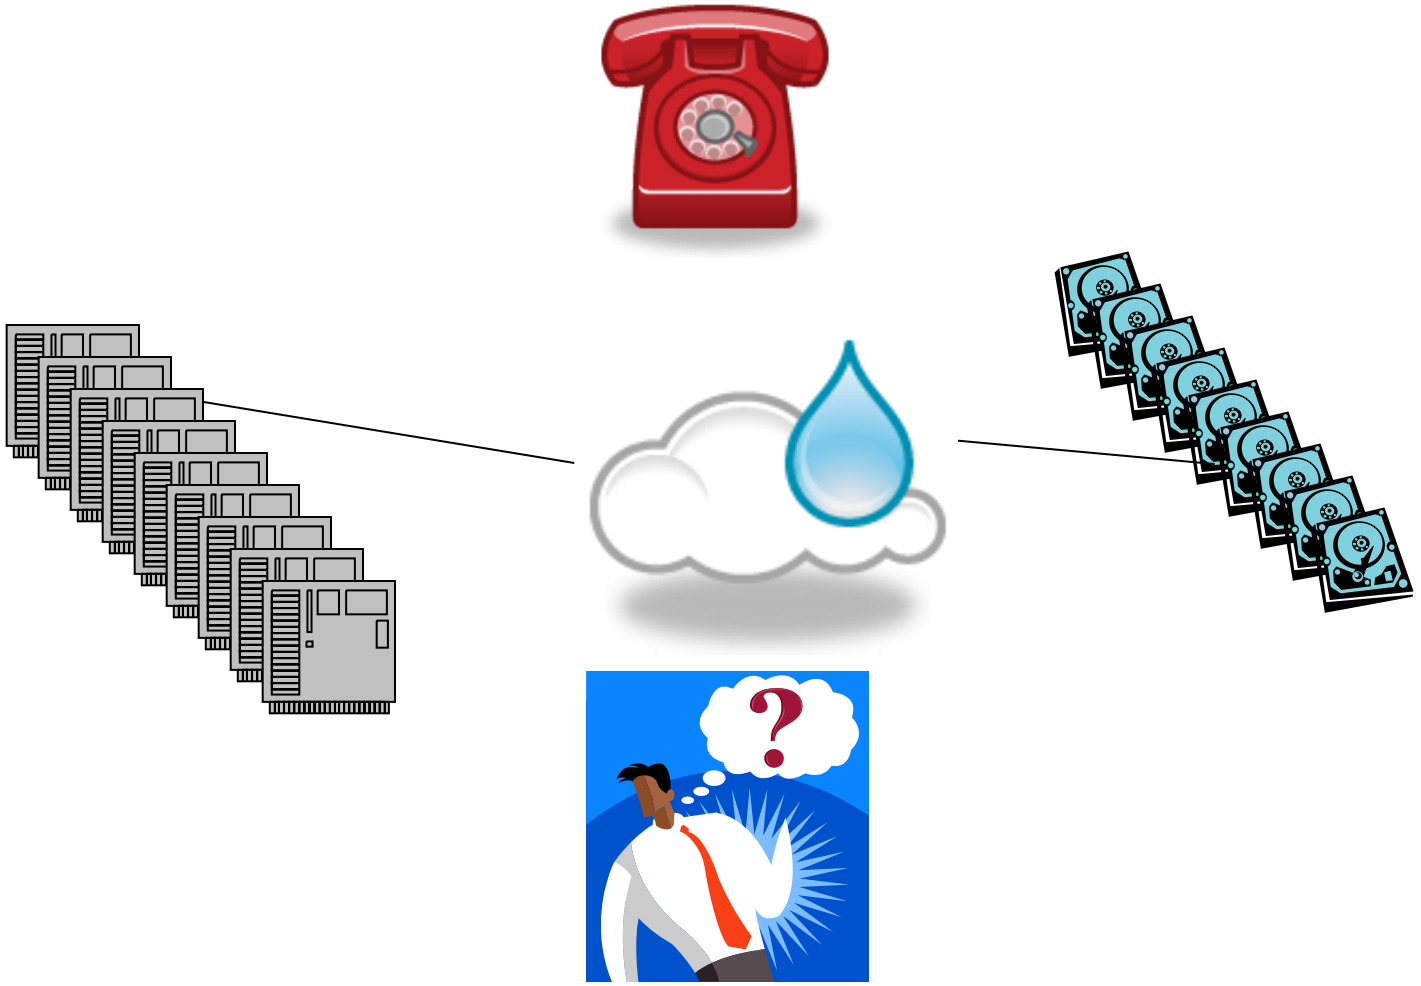

What we heard from LUG
We Want a Solution

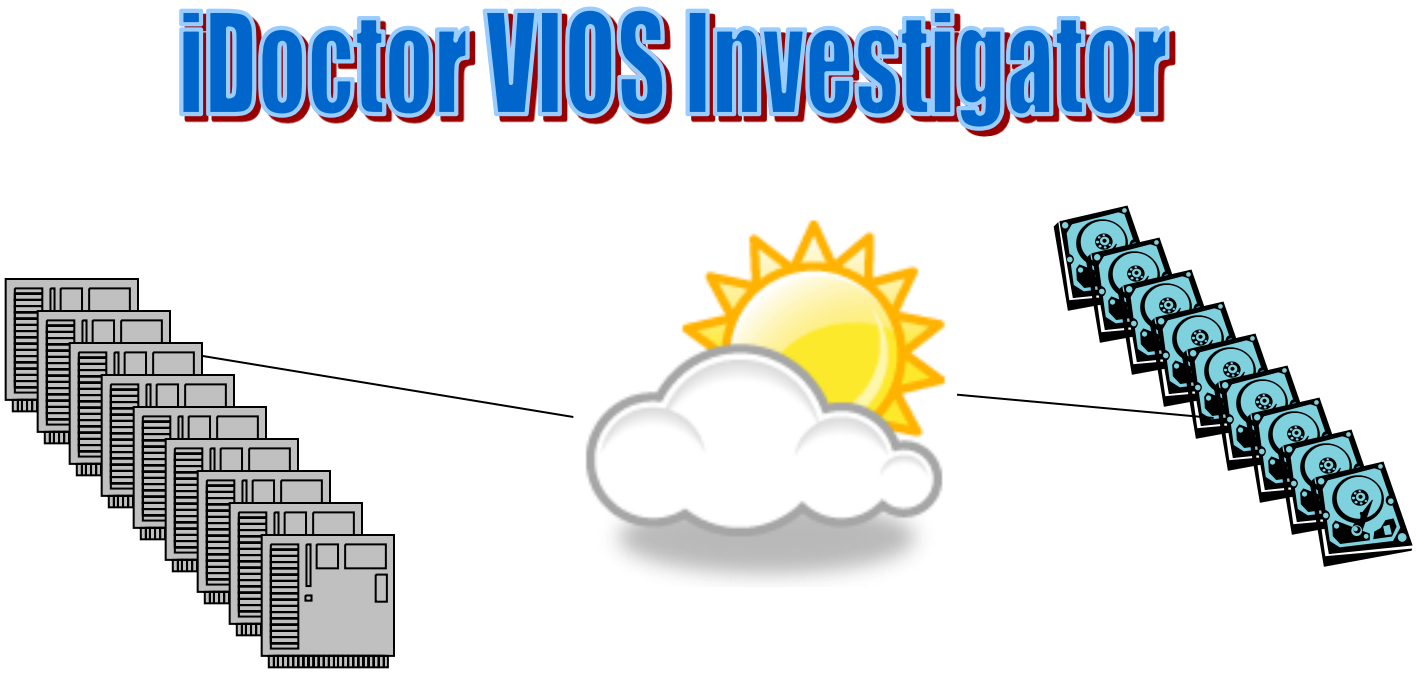

Pieces of VIOS Investigator

- **Disk Mappings (Tom)**
  - Provides mappings from IBM i disks to VIOS disks
- **topas_nmon data (Ron)**
  - Provides statistical data collection of VIOS disks
- **Analysis (Ron)**
  - Provides graphic analysis of statistical data mapped back to IBM i disks.

VIOS Investigator: Disk Mappings

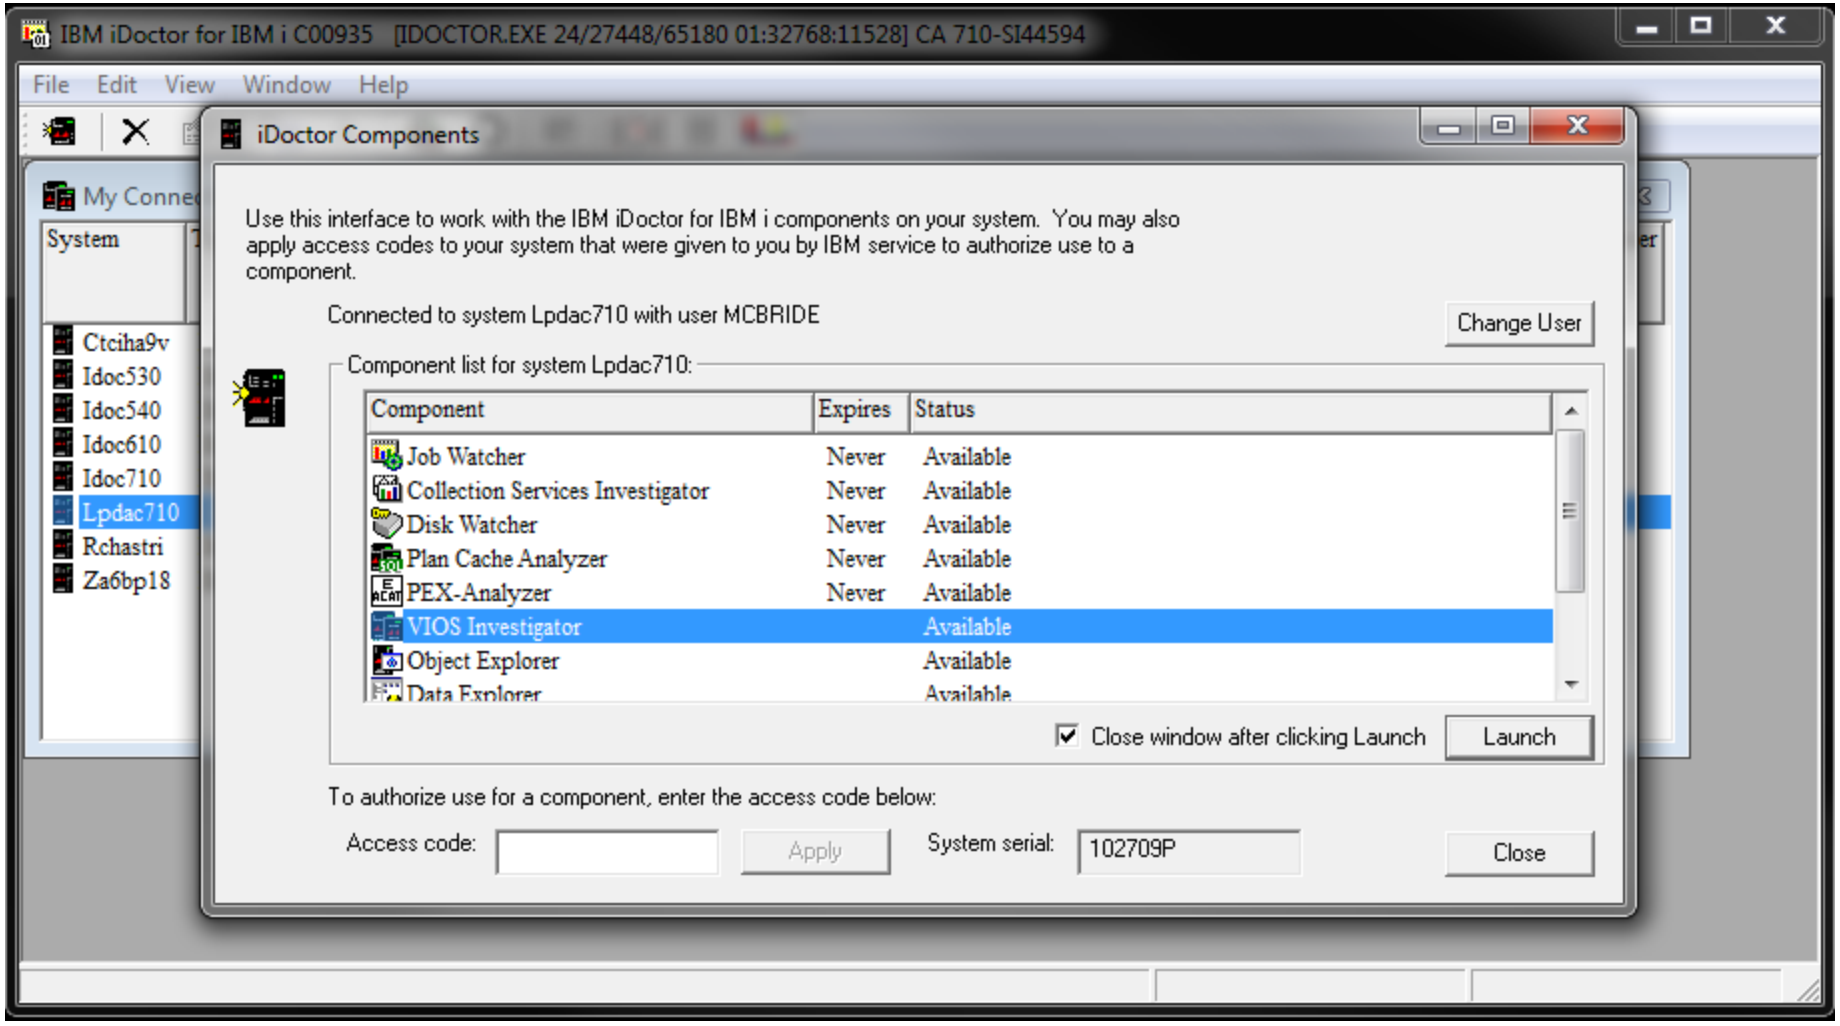

VIOS Investigator: Disk Mappings

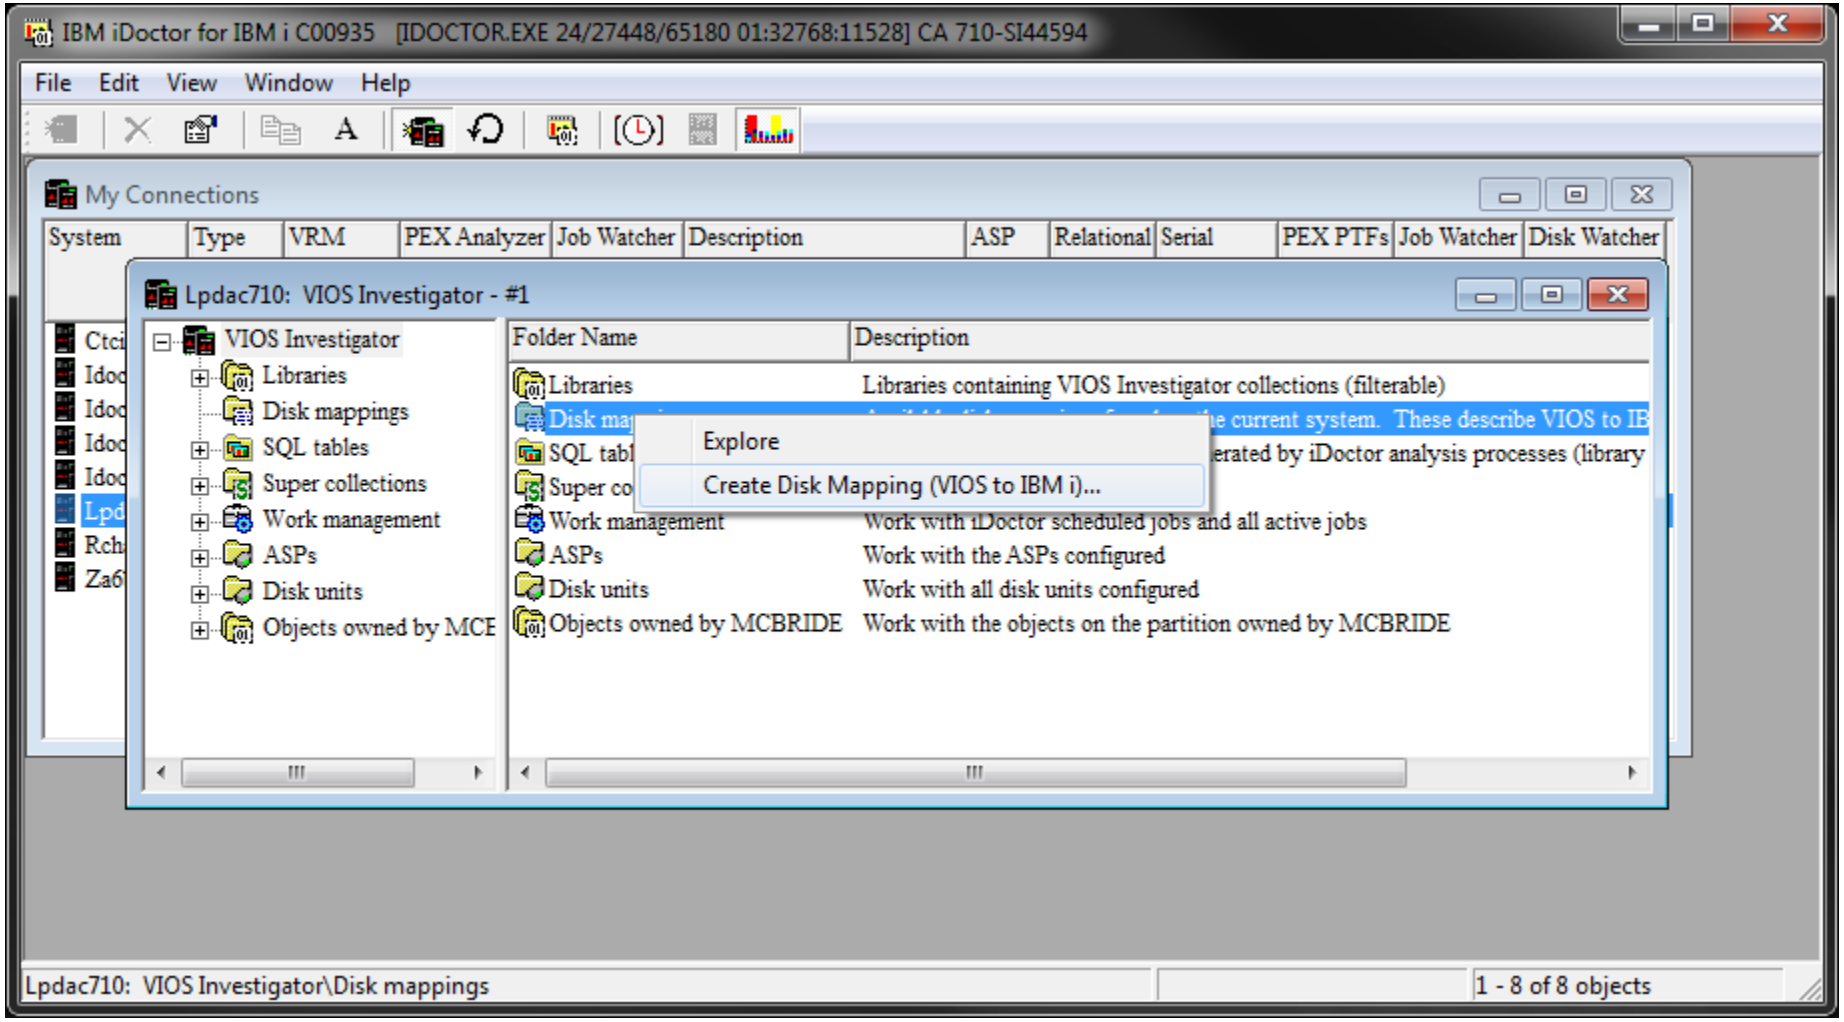

VIOS Investigator: Disk Mappings
Create Disk Mapping (VIOS to IBM i)…

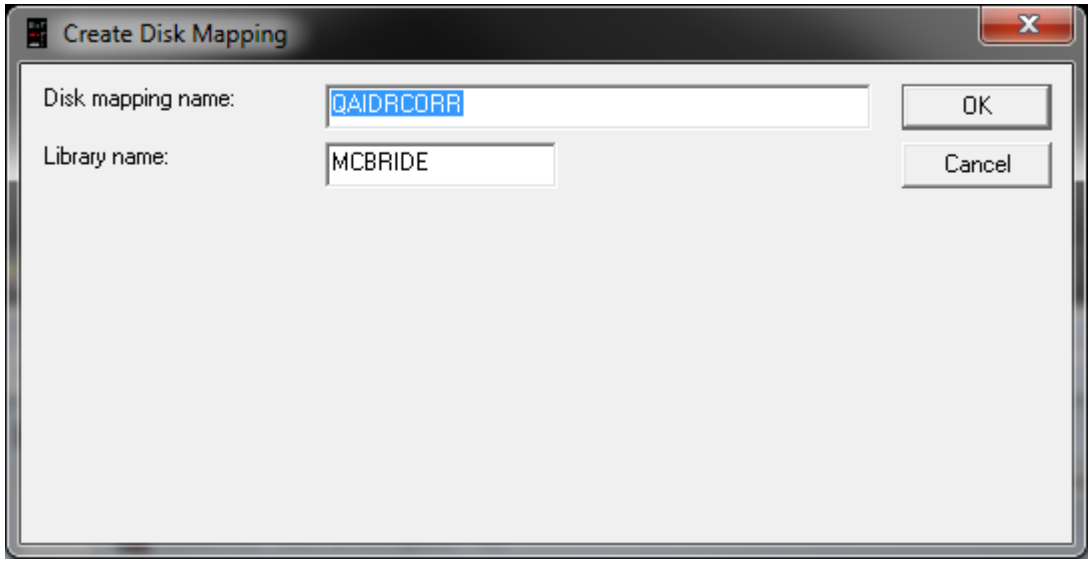

Where to store?

VIOS Investigator: Disk Mappings
Create Disk Mapping (VIOS to IBM i)…

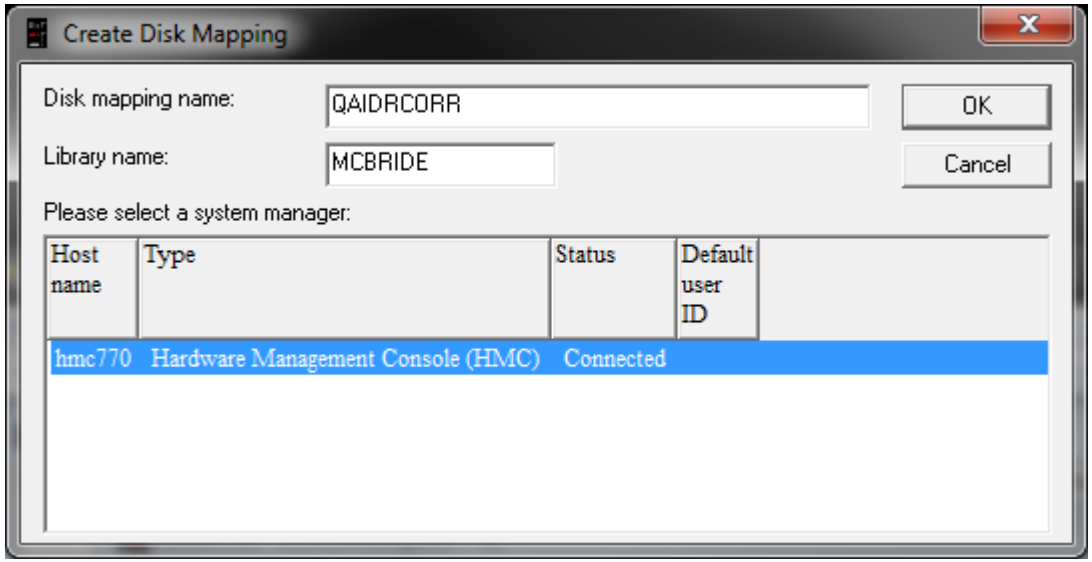

Which HMC to use?

VIOS Investigator: Disk Mappings
Create Disk Mapping (VIOS to IBM i)…

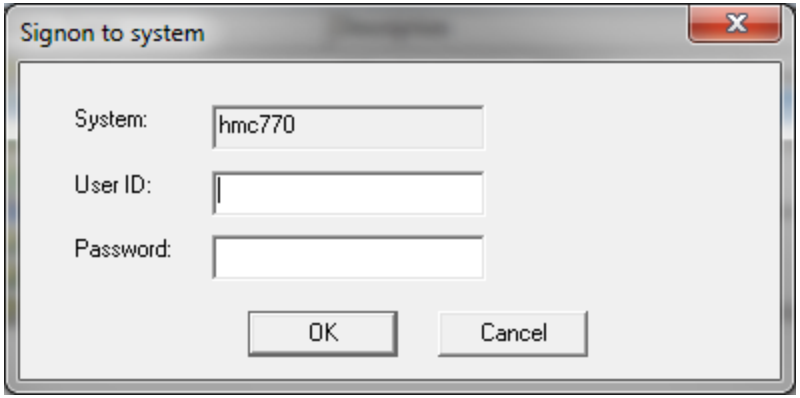

Credentials for HMC?

VIOS Investigator: Disk Mappings Explore

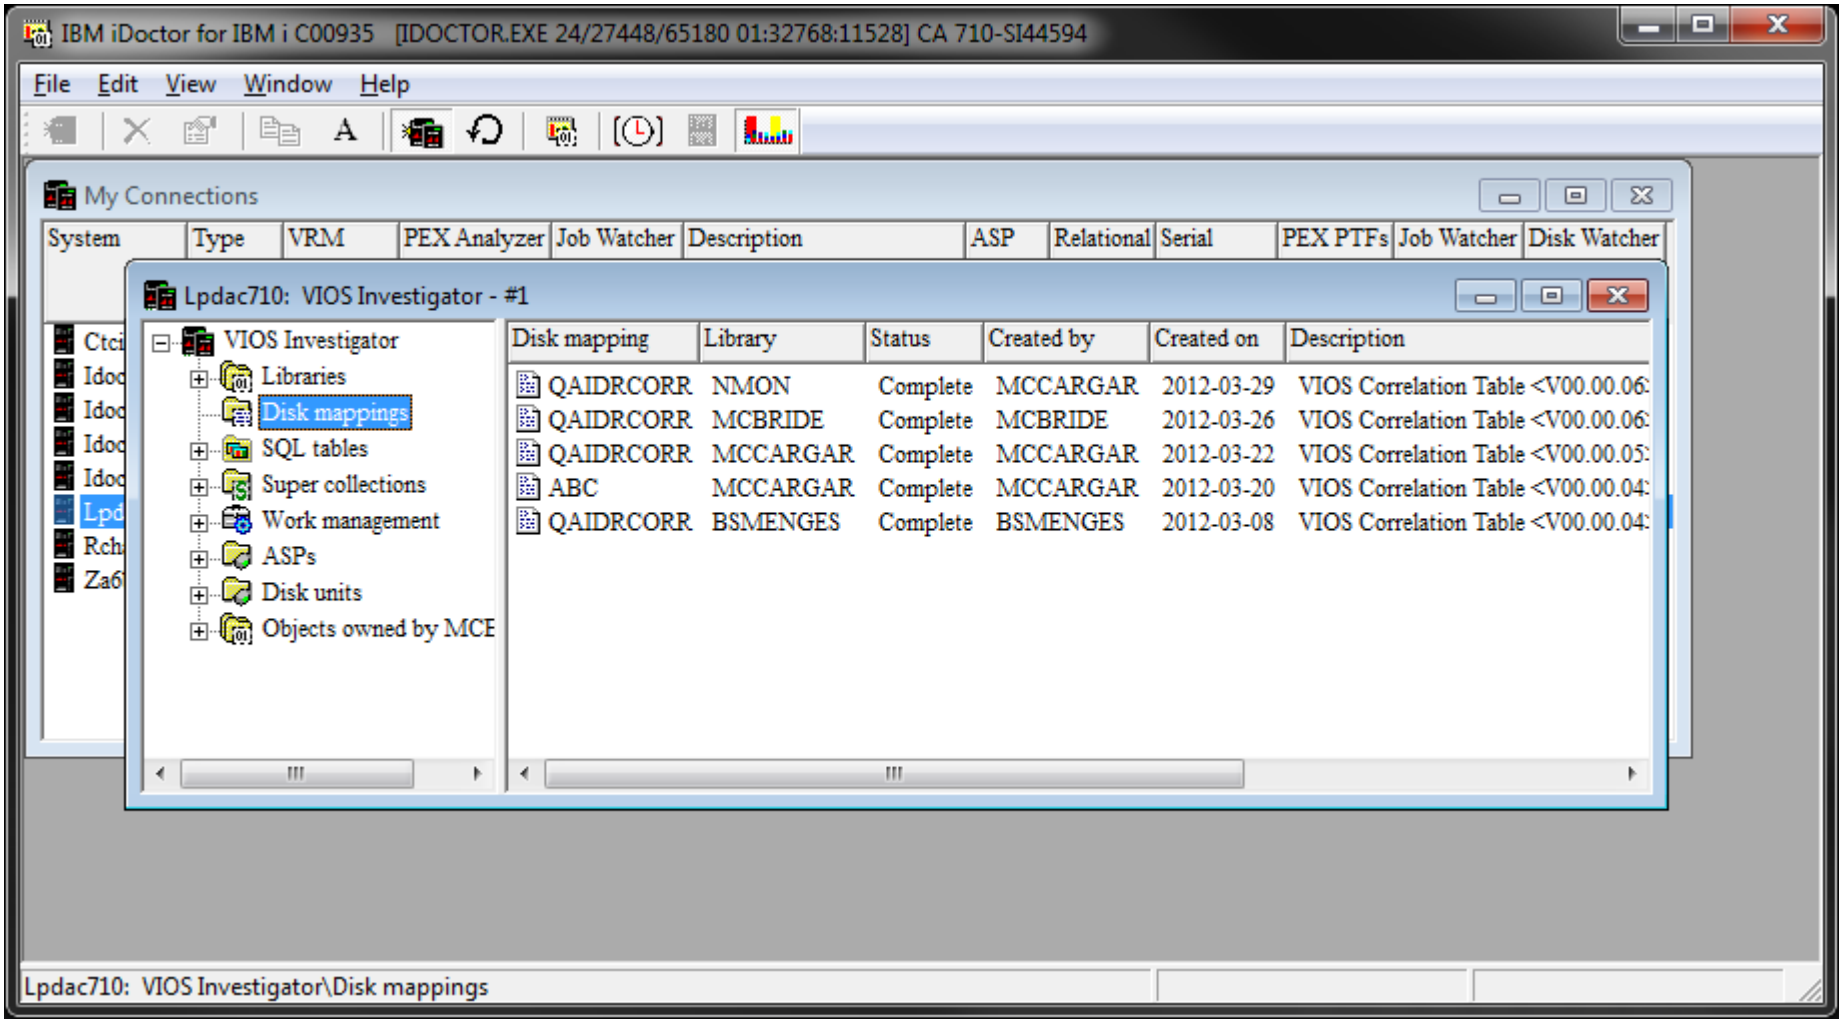

VIOS Investigator: Disk Mappings Results

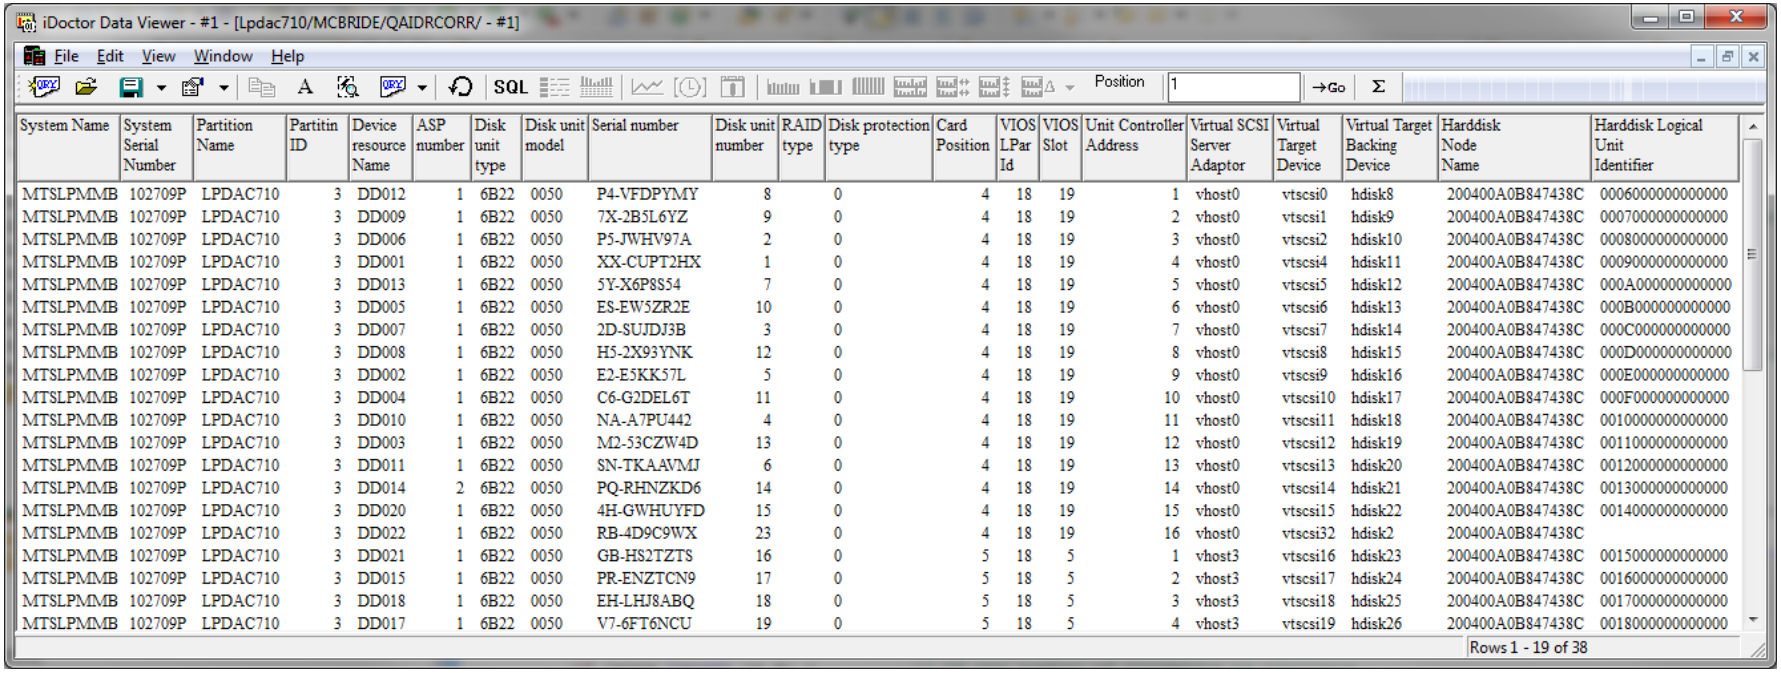

iDoctor Data Viewer - #1 - [Lpdac710/MCBRIDE/QAIDRCORR/ - #1]

File Edit View Window Help

Position 1 →Go Σ

System Name	System Serial Number	Partition Name	Partitin ID	Device resource Name	ASP number	Disk unit type	Disk unit model	Serial number	Disk unit number	RAID type	Disk protection type	Card Position	VIOS LPar Id	VIOS Slot	Unit Controller Address	Virtual SCSI Server Adaptor	Virtual Target Device	Virtual Target Backing Device	Harddisk Node Name	Harddisk Logical Unit Identifier
MTSLPMMB	102709P	LPDAC710	3	DD012	1	6B22	0050	P4-VFDPYMY	8	0		4	18	19	1	vhost0	vtscsi0	hdisk8	200400A0B847438C	0006000000000000
MTSLPMMB	102709P	LPDAC710	3	DD009	1	6B22	0050	7X-2B5L6YZ	9	0		4	18	19	2	vhost0	vtscsi1	hdisk9	200400A0B847438C	0007000000000000
MTSLPMMB	102709P	LPDAC710	3	DD006	1	6B22	0050	P5-JWHV97A	2	0		4	18	19	3	vhost0	vtscsi2	hdisk10	200400A0B847438C	0008000000000000
MTSLPMMB	102709P	LPDAC710	3	DD001	1	6B22	0050	XX-CUPT2HX	1	0		4	18	19	4	vhost0	vtscsi4	hdisk11	200400A0B847438C	0009000000000000
MTSLPMMB	102709P	LPDAC710	3	DD013	1	6B22	0050	5Y-X6P8S54	7	0		4	18	19	5	vhost0	vtscsi5	hdisk12	200400A0B847438C	000A000000000000
MTSLPMMB	102709P	LPDAC710	3	DD005	1	6B22	0050	ES-EW5ZR2E	10	0		4	18	19	6	vhost0	vtscsi6	hdisk13	200400A0B847438C	000B000000000000
MTSLPMMB	102709P	LPDAC710	3	DD007	1	6B22	0050	2D-SUJDJ3B	3	0		4	18	19	7	vhost0	vtscsi7	hdisk14	200400A0B847438C	000C000000000000
MTSLPMMB	102709P	LPDAC710	3	DD008	1	6B22	0050	H5-2X93YNK	12	0		4	18	19	8	vhost0	vtscsi8	hdisk15	200400A0B847438C	000D000000000000
MTSLPMMB	102709P	LPDAC710	3	DD002	1	6B22	0050	E2-E5KK57L	5	0		4	18	19	9	vhost0	vtscsi9	hdisk16	200400A0B847438C	000E000000000000
MTSLPMMB	102709P	LPDAC710	3	DD004	1	6B22	0050	C6-G2DEL6T	11	0		4	18	19	10	vhost0	vtscsi10	hdisk17	200400A0B847438C	000F000000000000
MTSLPMMB	102709P	LPDAC710	3	DD010	1	6B22	0050	NA-A7PU442	4	0		4	18	19	11	vhost0	vtscsi11	hdisk18	200400A0B847438C	0010000000000000
MTSLPMMB	102709P	LPDAC710	3	DD003	1	6B22	0050	M2-53CZW4D	13	0		4	18	19	12	vhost0	vtscsi12	hdisk19	200400A0B847438C	0011000000000000
MTSLPMMB	102709P	LPDAC710	3	DD011	1	6B22	0050	SN-TKAAVMJ	6	0		4	18	19	13	vhost0	vtscsi13	hdisk20	200400A0B847438C	0012000000000000
MTSLPMMB	102709P	LPDAC710	3	DD014	2	6B22	0050	PQ-RHNZKD6	14	0		4	18	19	14	vhost0	vtscsi14	hdisk21	200400A0B847438C	0013000000000000
MTSLPMMB	102709P	LPDAC710	3	DD020	1	6B22	0050	4H-GWHUYFD	15	0		4	18	19	15	vhost0	vtscsi15	hdisk22	200400A0B847438C	0014000000000000
MTSLPMMB	102709P	LPDAC710	3	DD022	1	6B22	0050	RB-4D9C9WX	23	0		4	18	19	16	vhost0	vtscsi32	hdisk2	200400A0B847438C	
MTSLPMMB	102709P	LPDAC710	3	DD021	1	6B22	0050	GB-HS2TZTS	16	0		5	18	5	1	vhost3	vtscsi16	hdisk23	200400A0B847438C	0015000000000000
MTSLPMMB	102709P	LPDAC710	3	DD015	1	6B22	0050	PR-ENZTCN9	17	0		5	18	5	2	vhost3	vtscsi17	hdisk24	200400A0B847438C	0016000000000000
MTSLPMMB	102709P	LPDAC710	3	DD018	1	6B22	0050	EH-LHJ8ABQ	18	0		5	18	5	3	vhost3	vtscsi18	hdisk25	200400A0B847438C	0017000000000000
MTSLPMMB	102709P	LPDAC710	3	DD017	1	6B22	0050	V7-6FT6NCU	19	0		5	18	5	4	vhost3	vtscsi19	hdisk26	200400A0B847438C	0018000000000000

Rows 1 - 19 of 38

VIOS Investigator GUI Overview

- **Collect data (NMON)**
- **Import data into DB2/IBM i**
- **Analyze**

Collect Data – Step 1 – Add VIOS Connection

- **Add your VIOS to iDoctor's Connection List**
- **Right-click Add Connection…**
- **Set the Connection Type to FTP**

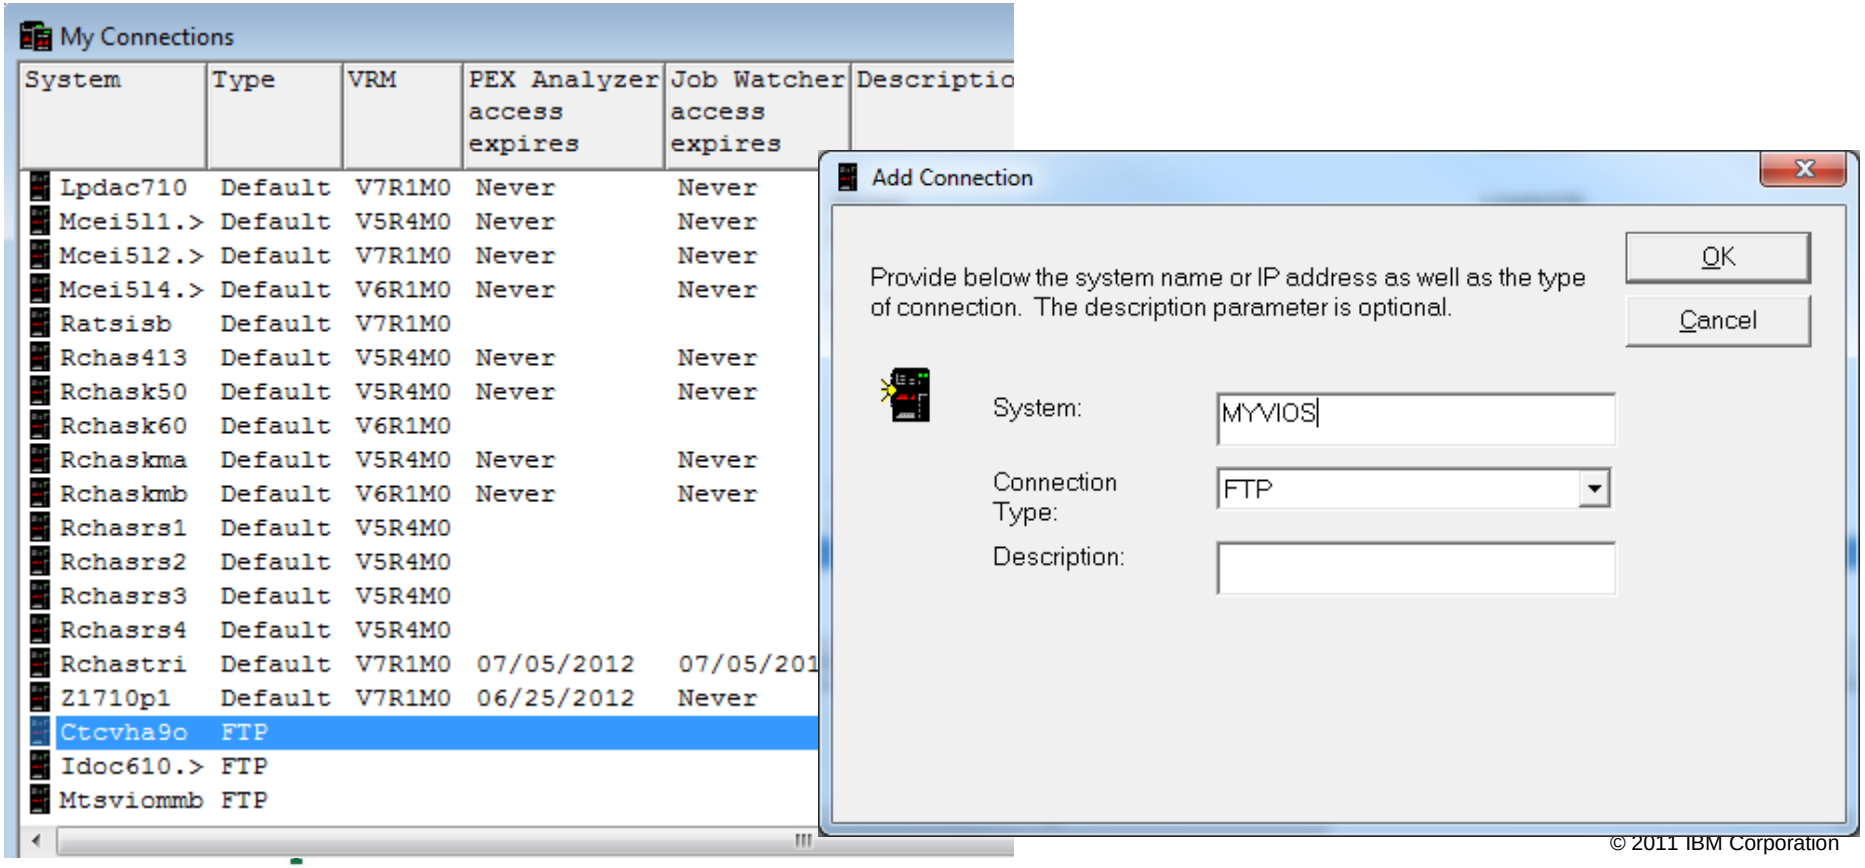

Collect Data – Step 2 - Start Collection menu in GUI

- **Use this to start an NMON collection via iDoctor GUI**
- **SSH 2.0+ must be installed on the VIOS.**
- **The user id must be authorized to run topas_nmon**

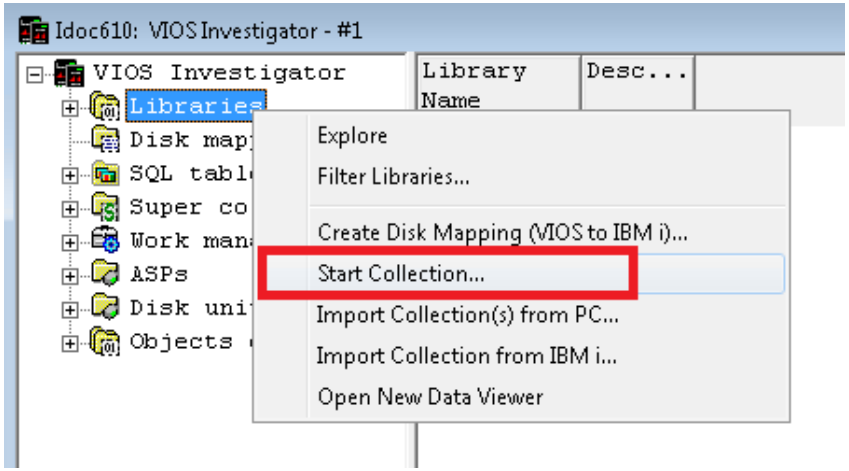

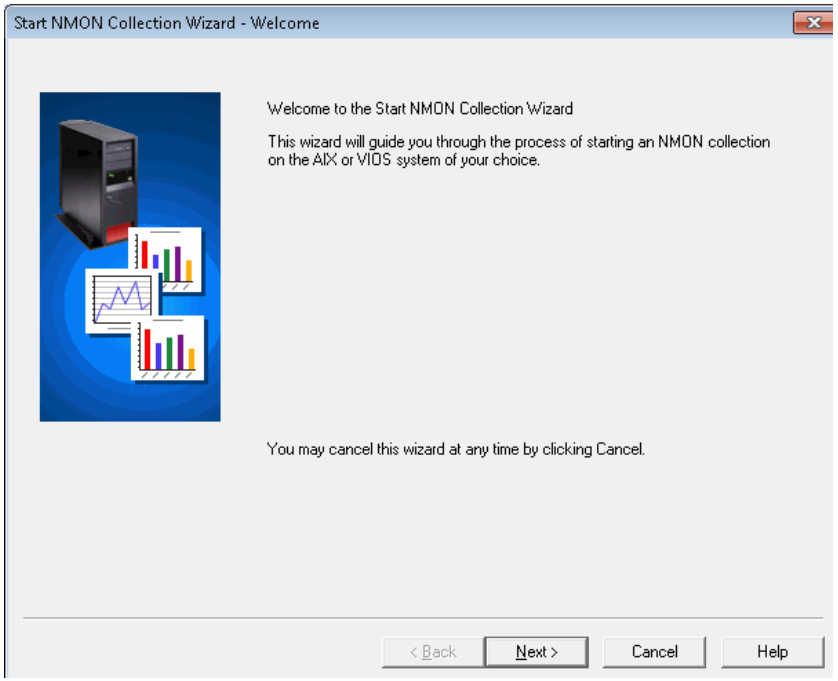

Collect Data – Step 3 – Select the VIOS

- **List of VIOS/AIX systems is the FTP type connections defined in iDoctor's connection list.**

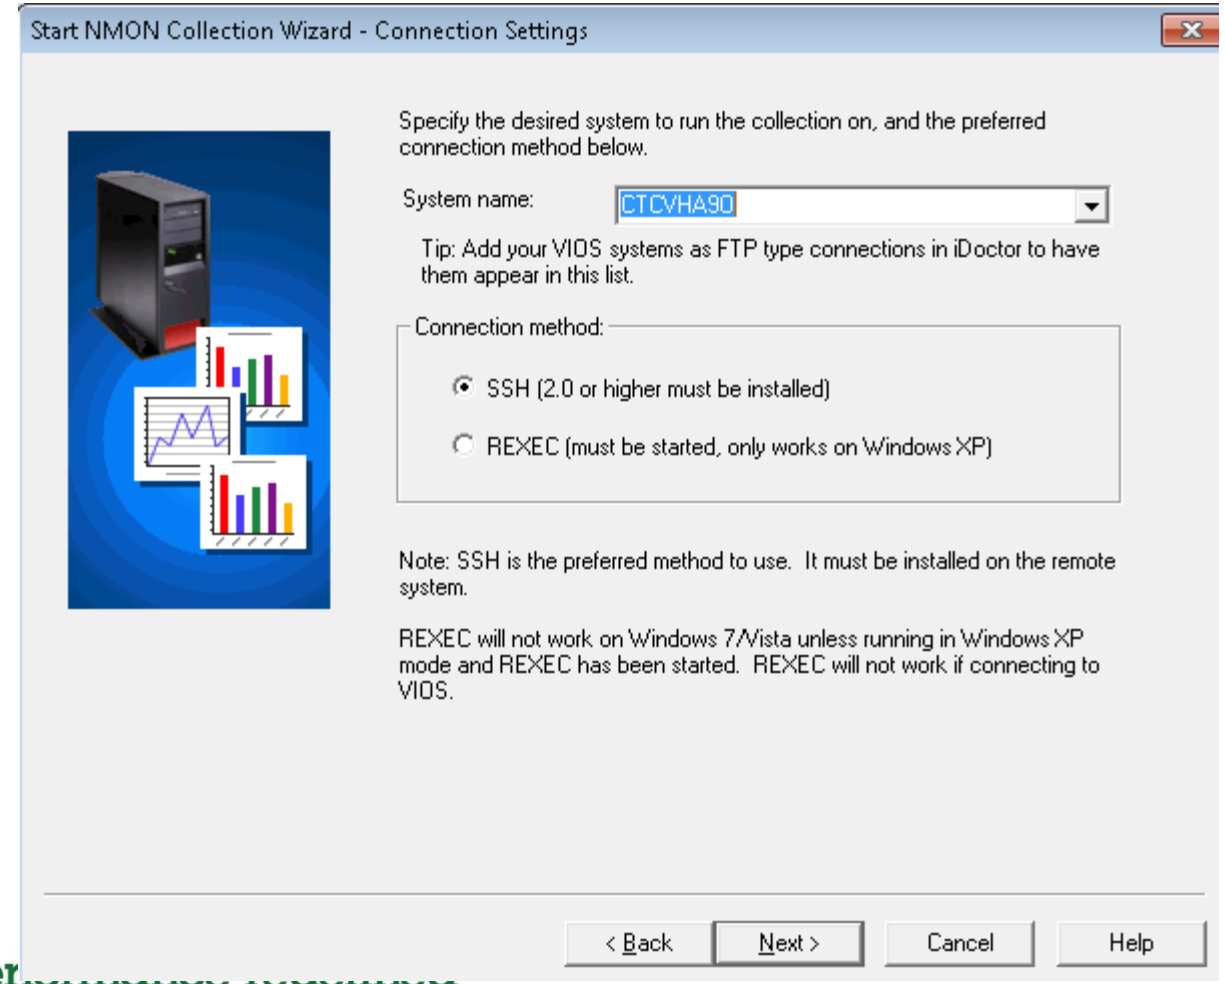

Collect Data – Step 4 – Configure collection options

- **Controls the key parameters on the command**

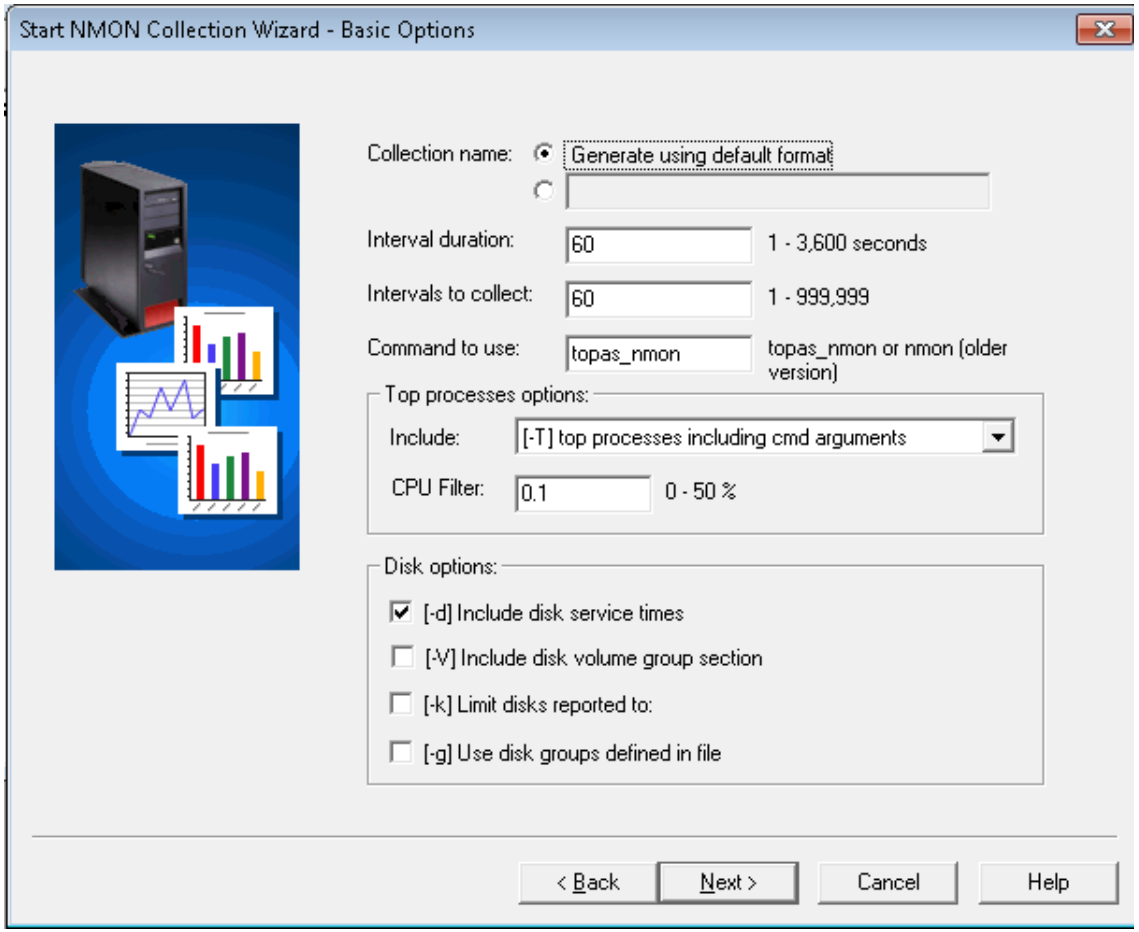

Collect Data – Step 5 – Submit the command

- **The final screen summarizes the parameters used and shows the command that will be executed on the VIOS.**

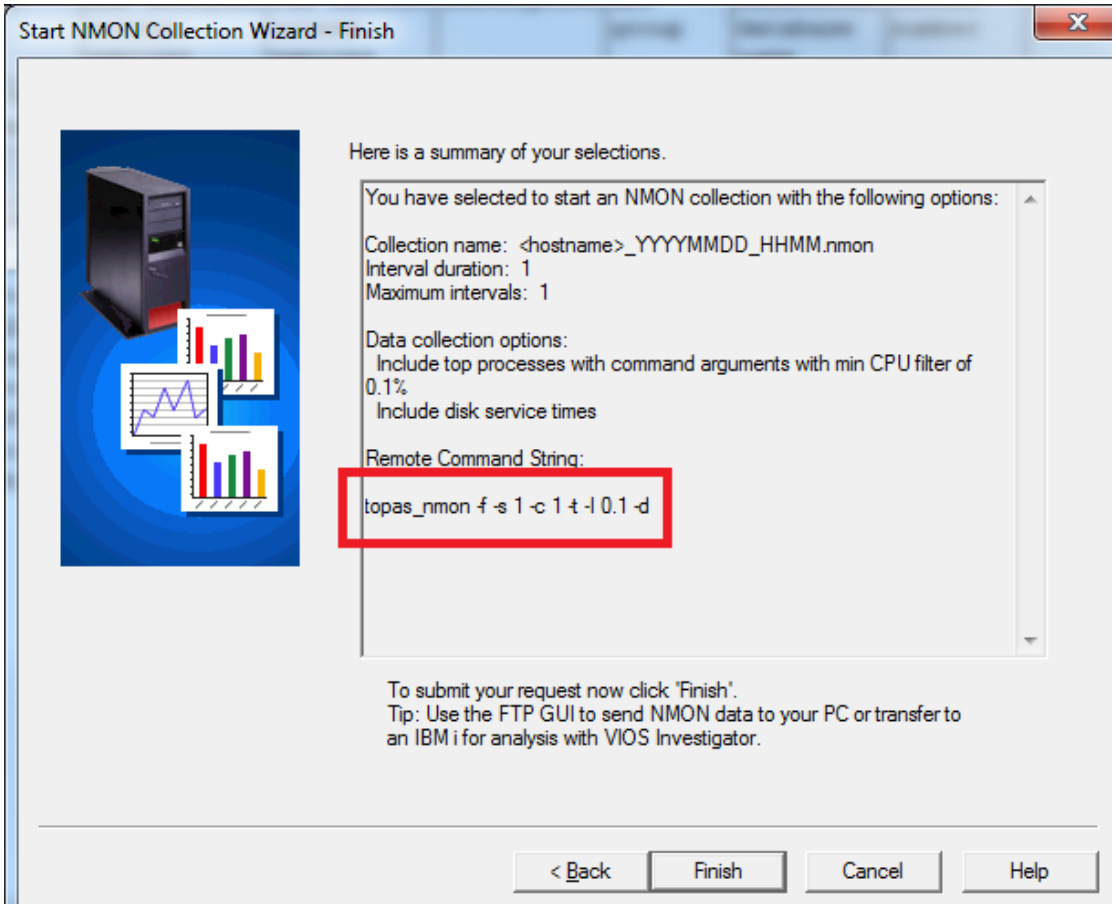

Import Data – Step 1 – Select the NMON files

- Start by creating a FTP type connection to the VIOS
- From My Connections (right-click Add Connection)
- Double-click the new connection to connect to VIOS using FTP
- Find and select desired .nmon files, right-click “Transfer to another system” menu

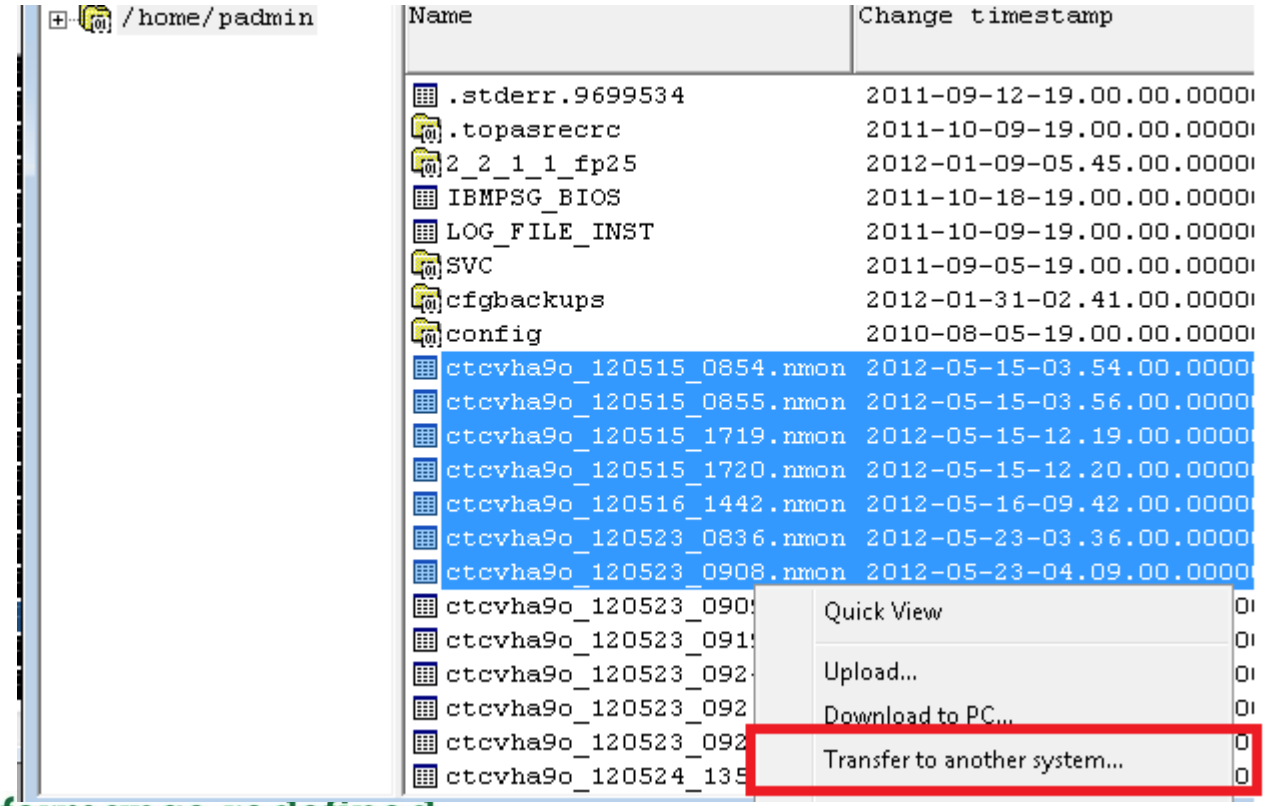

Import Data – Step 2 – Transfer File(s) to IBM i window

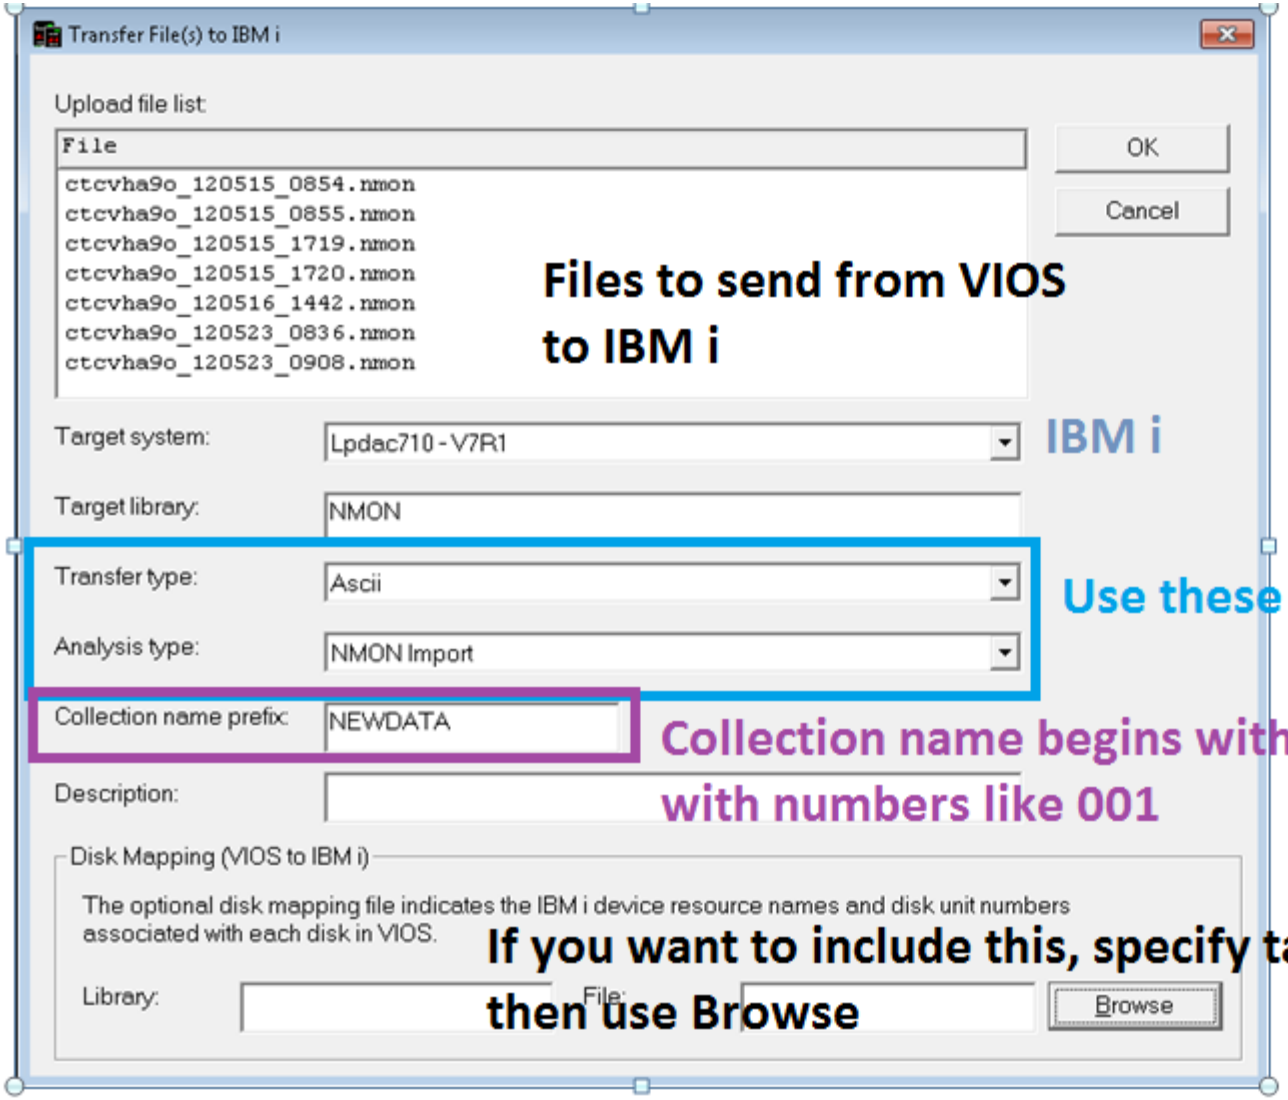

Import Data – Step 3 – Wait for analysis to complete

- Next the data is sent to the IBM i (via the PC)
- An SQL script is created and submitted to a batch job QIDRNMSUM that will run a stored procedure against each file to import the data.
  - Source file is QUSRSYS/QAIDRNMSQ1

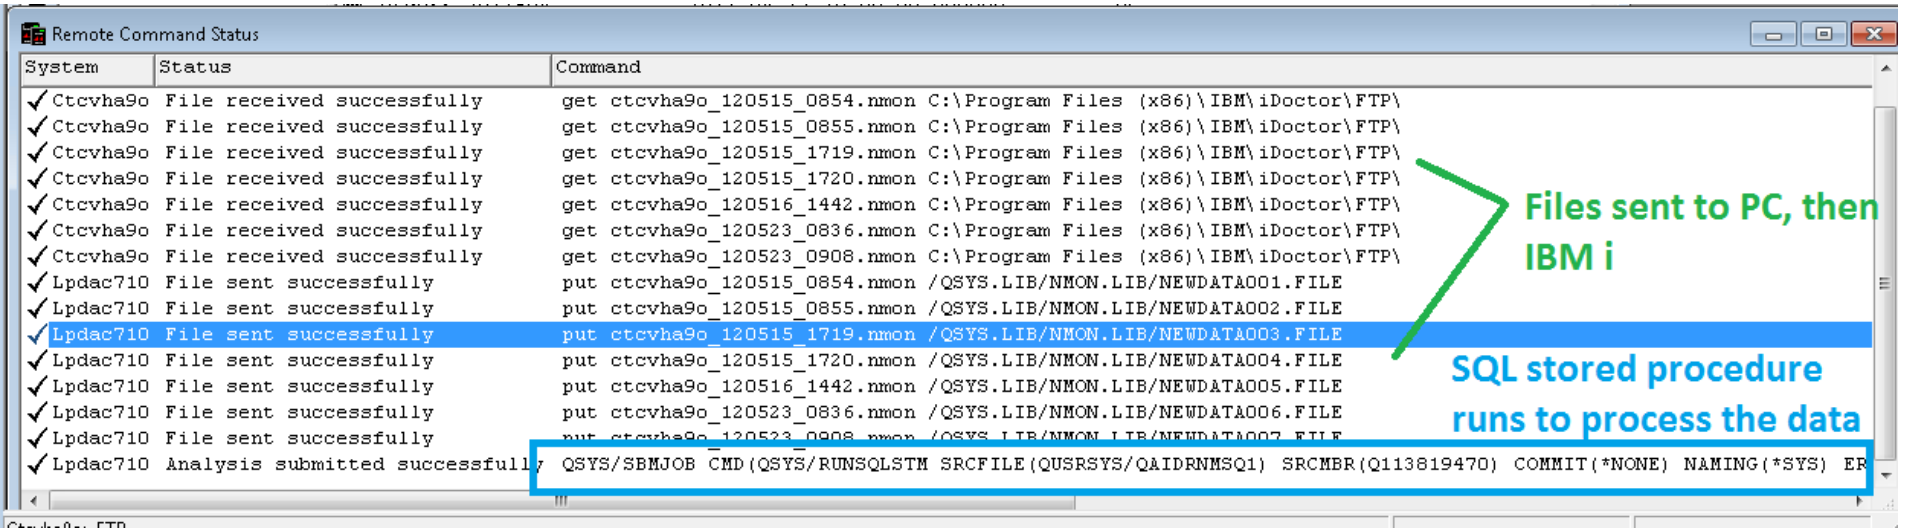

Analysis – Listing Collections in a Library

- Open VIOS Investigator on the IBM i to work with the data under the Libraries folder.
- **Note**: A status of “invalid disk mapping” or “No disk mapping provided” is normal if a disk mapping was not specified on the transfer/import window.

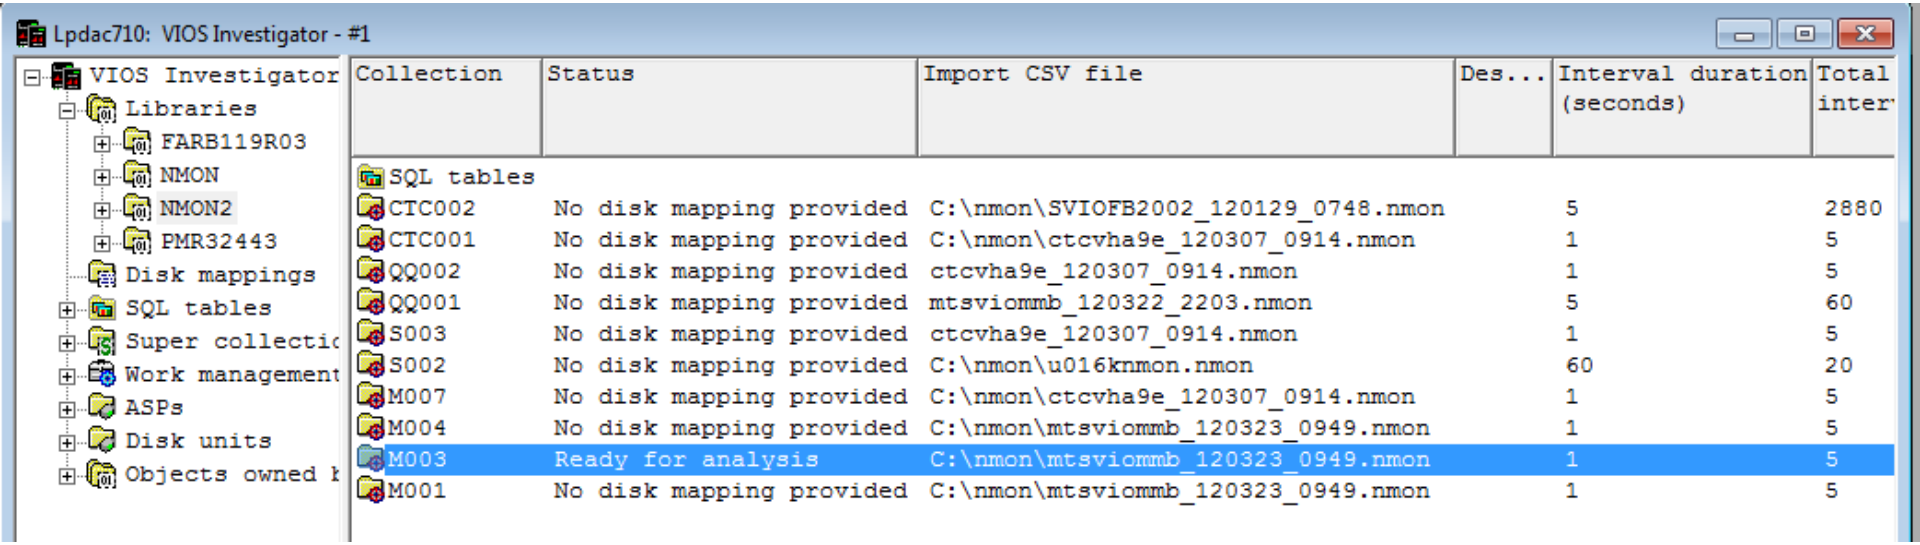

Analysis Options Overview

- **Graphs and reports are available under the Libraries folder in VIOS Investigator**
- Expand the desired library/collection to work with them

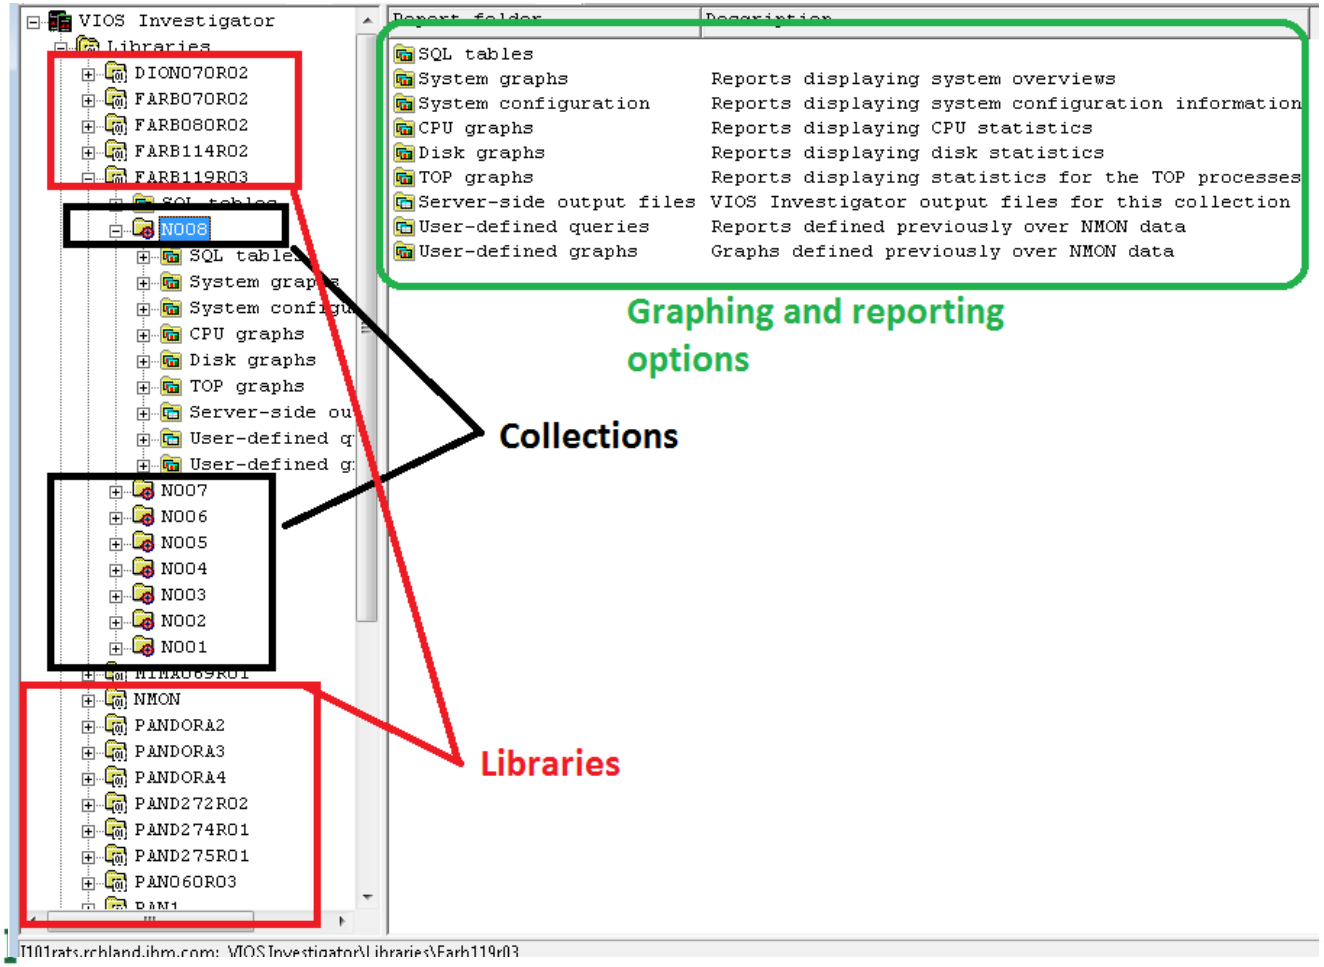

Analysis Options – Collection information

- This information is displayed in the library about each collection
- Right-click a collection -> Record Quick View to see in a vertical list

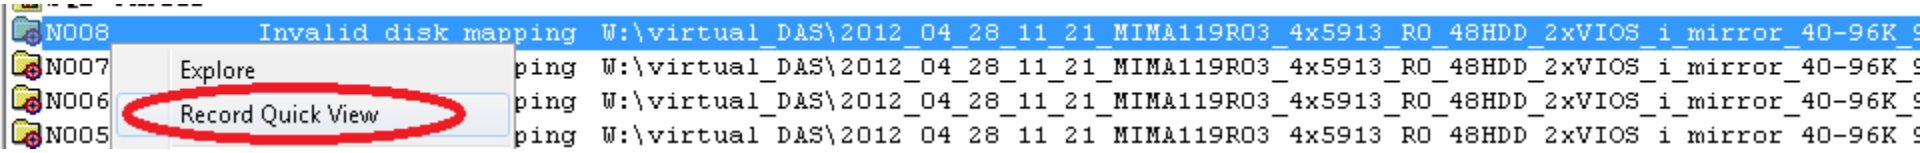

Record Quick View

Description	Record 1
System serial number	1026D2P
AIX	6.1.8.0
NMON command	/usr/bin/topas_nmon -F /tmp/nmon.nmon -s 60 -c 20 -t -I 0.1 -d -youtput_dir=/tmp/nmon.nmon -ystart_time=15:13:56
NMON build	AIX
User name	padmin
Node name	farbauti
Host name	farbauti
Start time	28-APR-2012 15:13:57
Import time	2012-04-30-07.11.13.560189
NMON version	TOPAS-NMON
Total intervals	20
Interval duration (seconds)	60
Description	
Import CSV file	W:\virtual_DAS\2012_04_28_11_21_MIMA119RO3_4x5913_RO_48HDD_2xVIOS_i_mirror_40-96K_9600WHS_write_cache_on\u096k\nm
Status	Invalid disk mapping
Collection	N008

Analysis Options – Open graph from collection menu

– You can also right-click a collection to launch graphs from the menu.

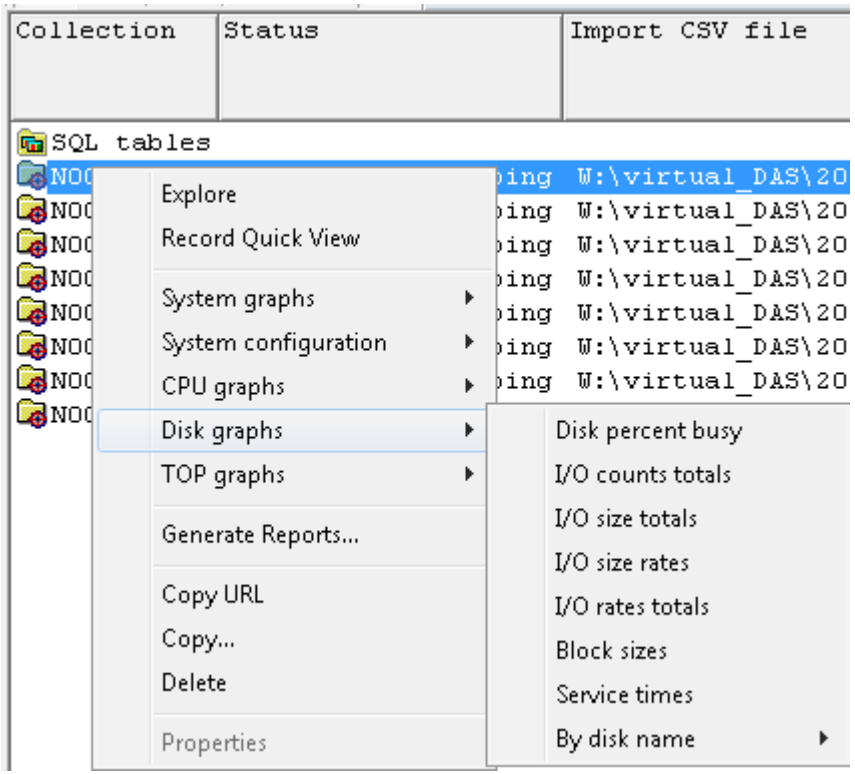

Analysis Options – Graphing multiple collections

– Select all desired collections, right-click and pick the desired graph

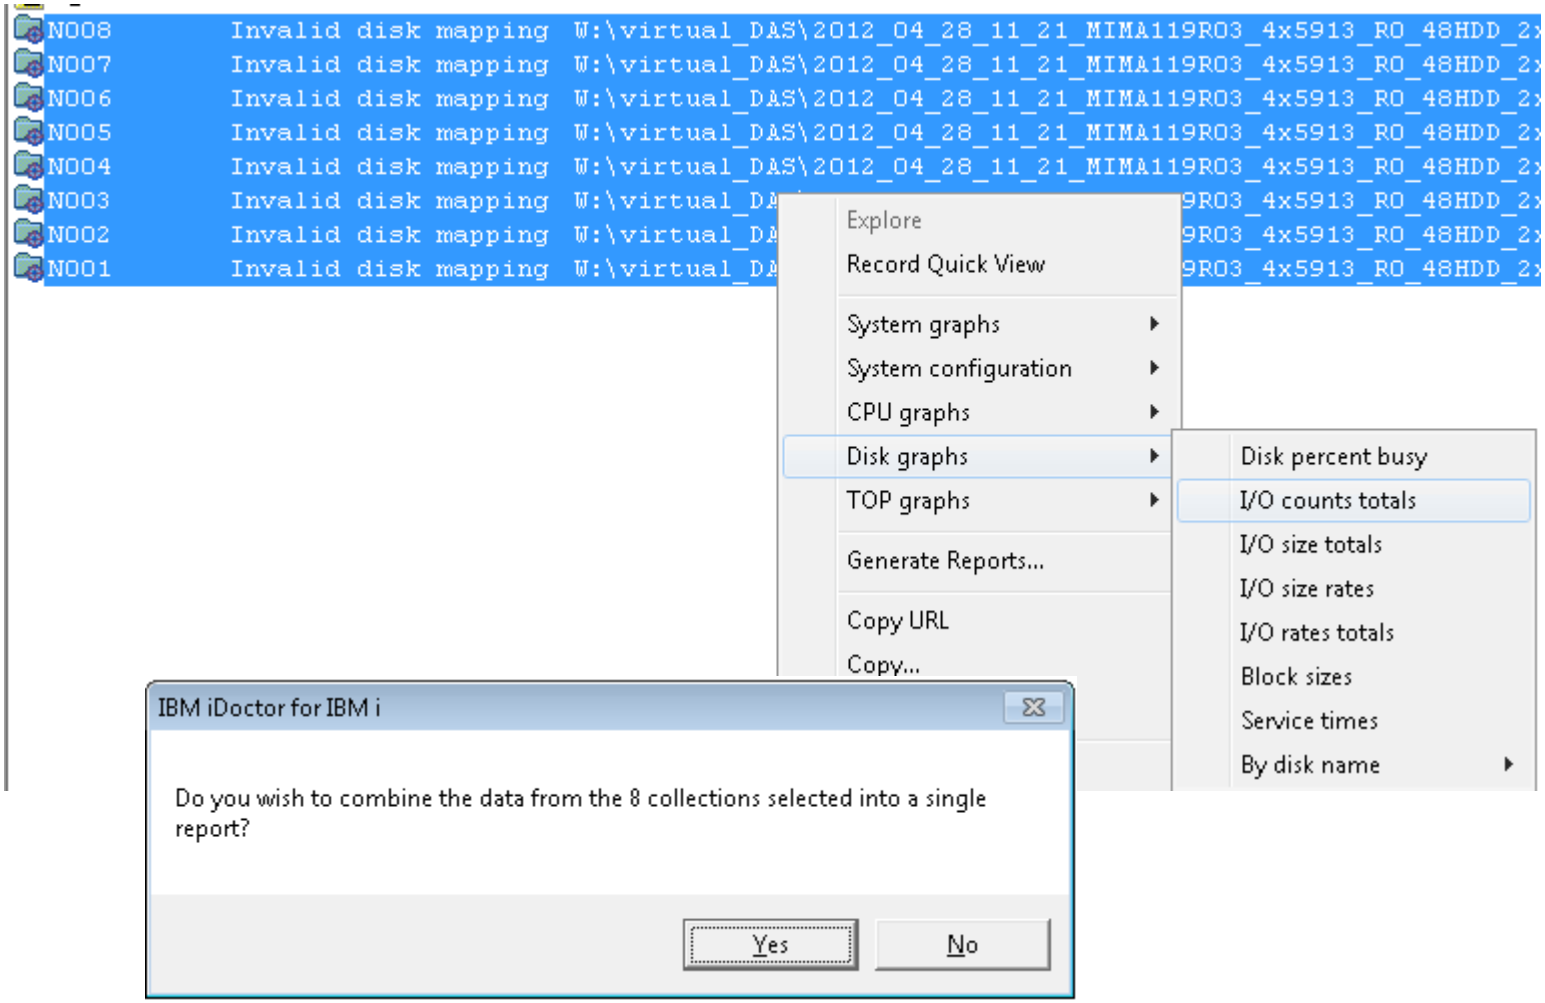

Analysis Options – Graphing multiple collections

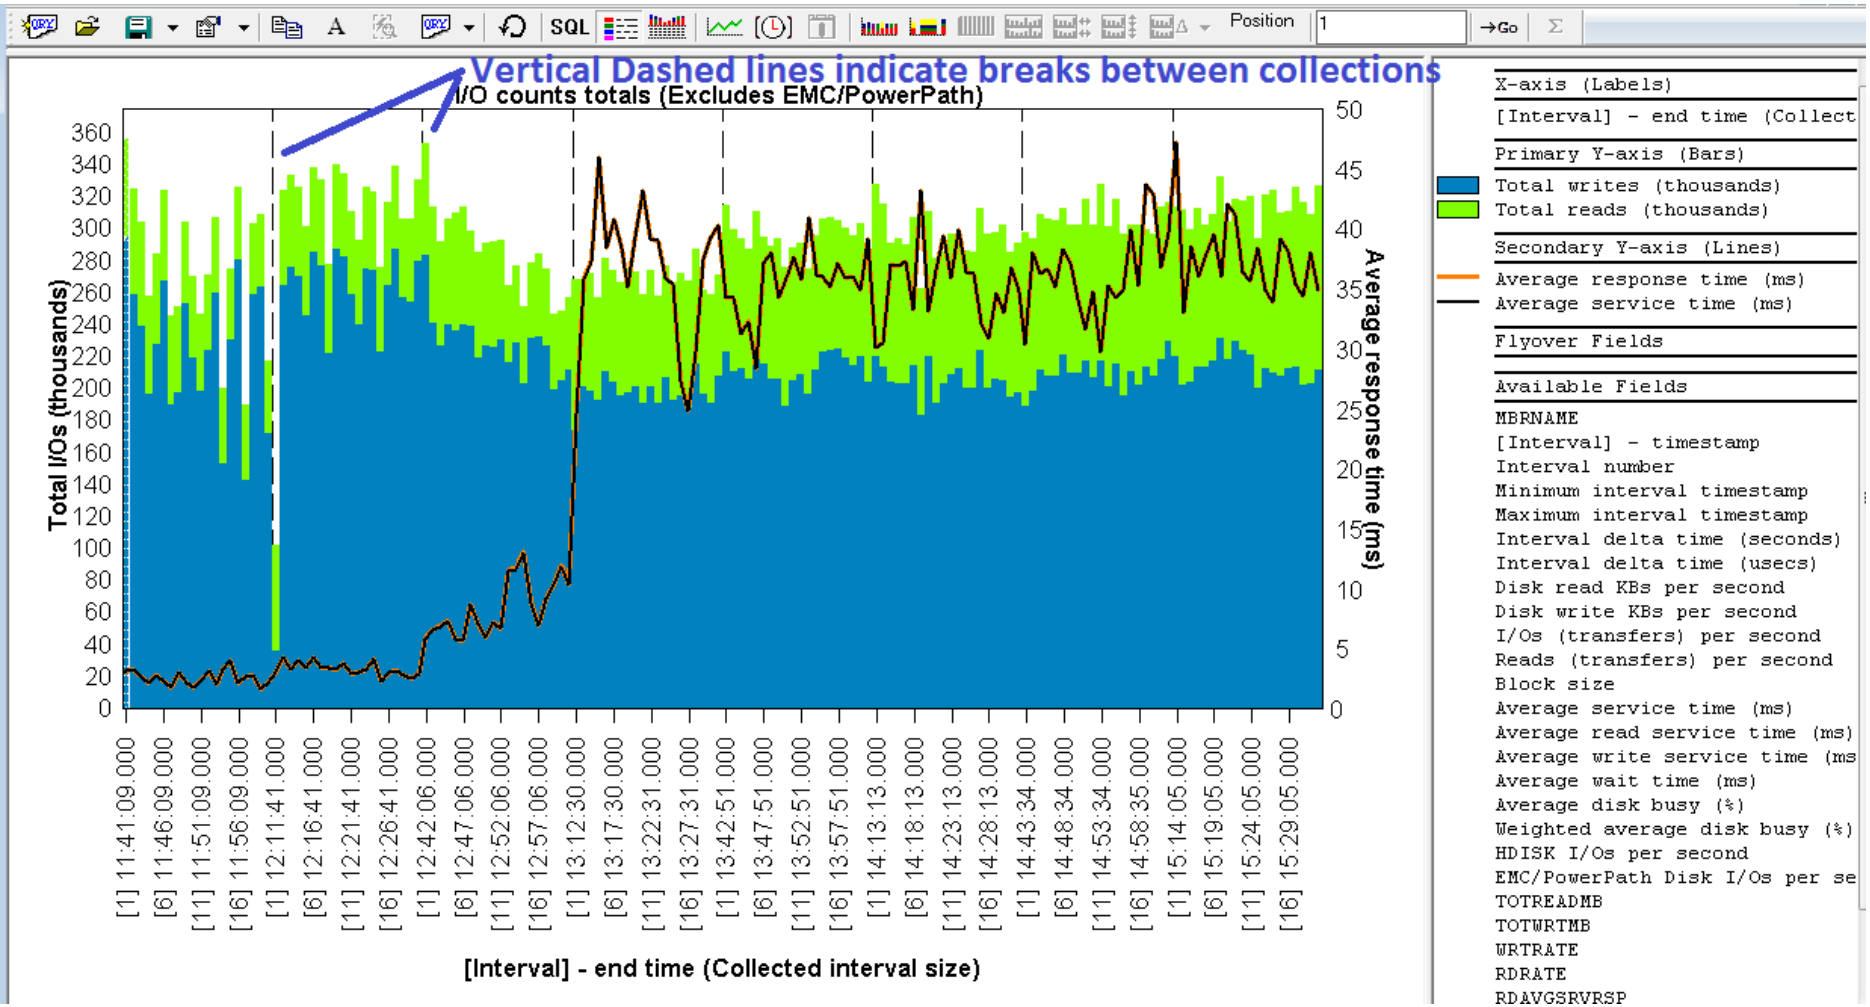

Analysis Options – Same graph shown differently

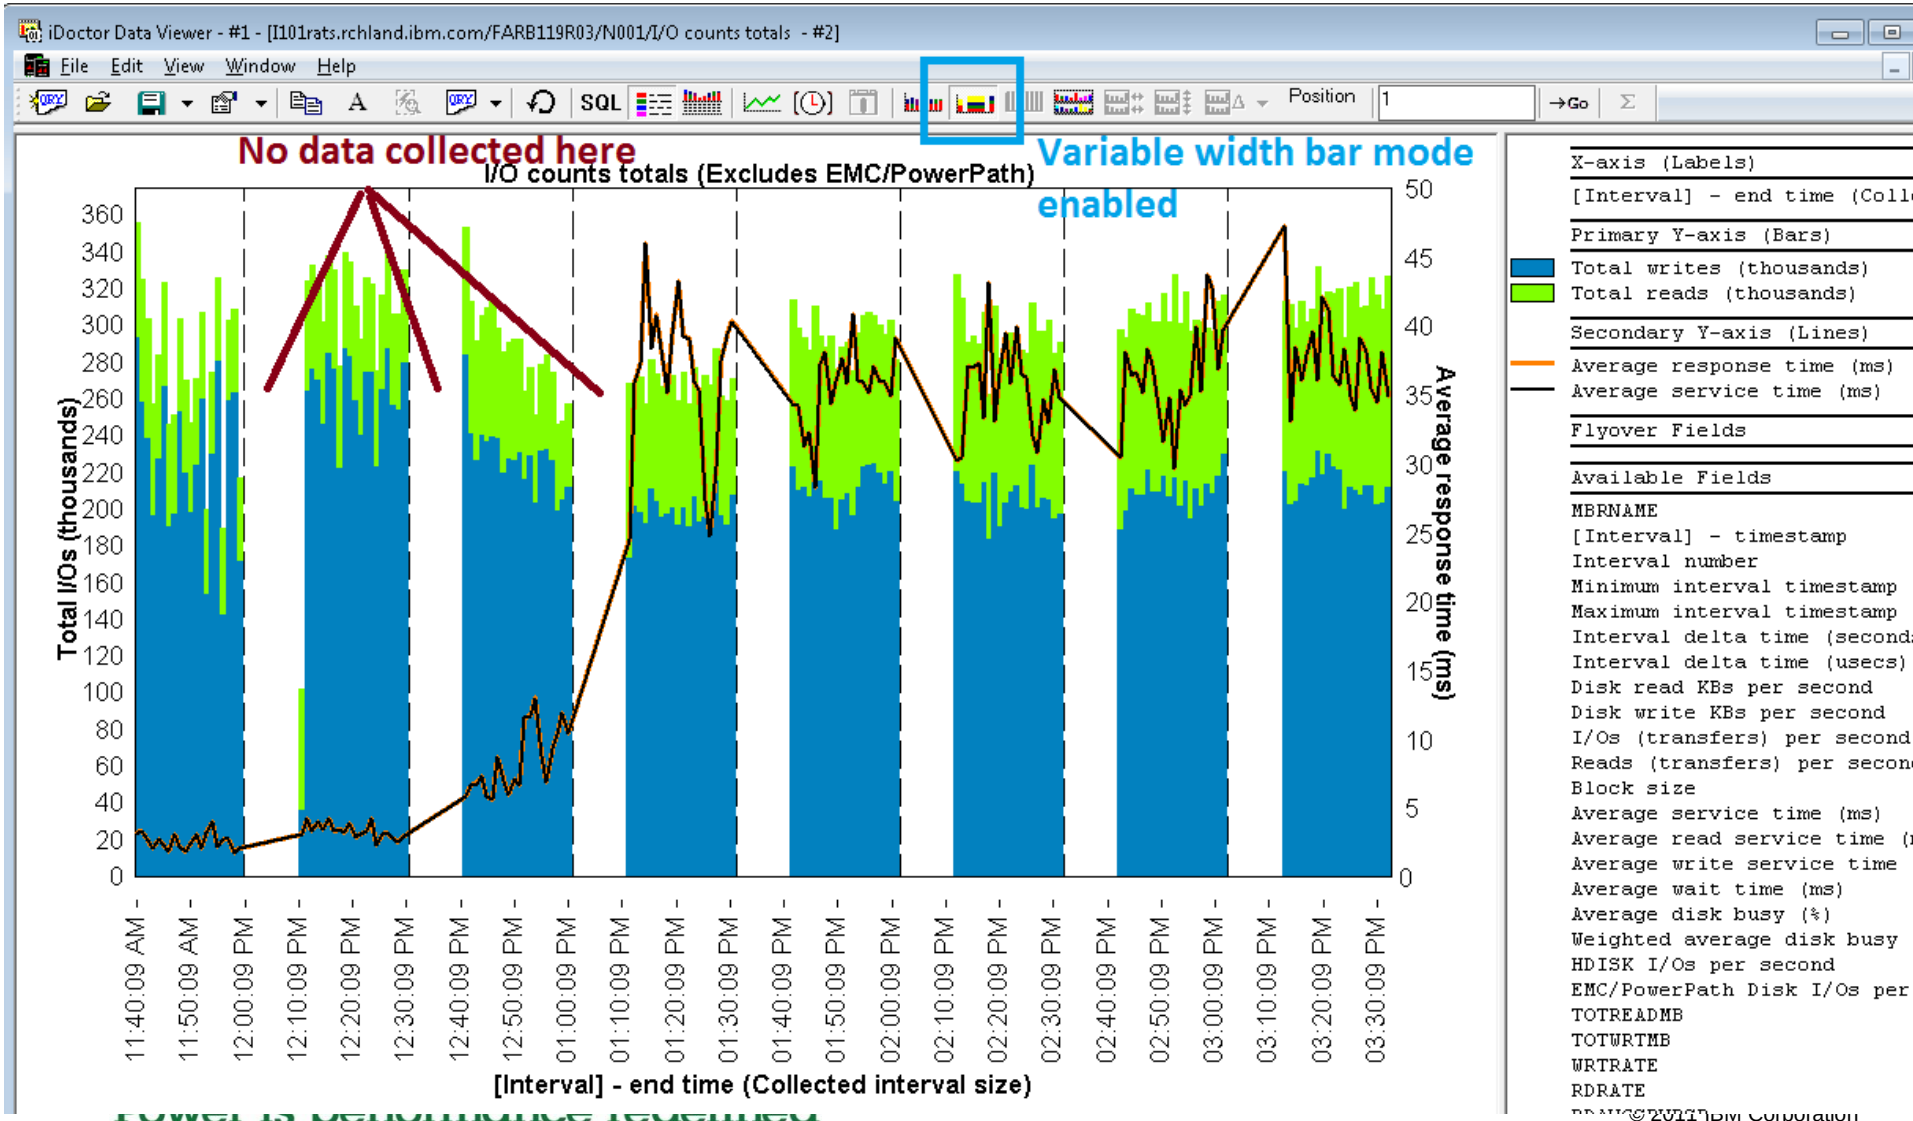

Analysis Options – Disk statistics provided by iDoctor

- **Under server-side output files -> QAIDRNMDISK_COLNAME**
- **The descriptions used where possible match the terms used in Collection Services Investigator's Disk graphs**

Field	Description	Record 4
INTERVAL	Interval number	1
DISKID	iDoctor disk ID	4
DISKNAME	Disk name	hdisk4
DISKBUSY	Percent busy	61.5000
DISKREAD	Read KBs per second	266.0100
DISKWRITE	Write KBs per second	1438.5000
DISKXFER	I/Os (transfers) per second	270.1300
DISKRXFER	Reads (transfers) per second	52.3400
DISKBSIZE	Block size	6.3300
DISKRIO	Reads per second	52.3000
DISKWIO	Writes per second	217.8000
DISKAVGRIO	DISKAVGRIO	3.6000
DISKAVGWIO	DISKAVGWIO	6.3800
DISKSERV	Average service time (ms)	1.5775
DISKREADSERV	Average read service time (ms)	1.9000
DISKWRITESERV	Average write service time (ms)	1.5000
DISKWAIT	Average wait time (ms)	0
INTENDSTR	Interval start time	2012-04-28-11.41.09.000000

Analysis Options – Disk graphs available without a disk mapping

- **Under the collection, expand the Disk Graphs folder**
- **The rankings graphs will include more options if a valid disk mapping has been used**

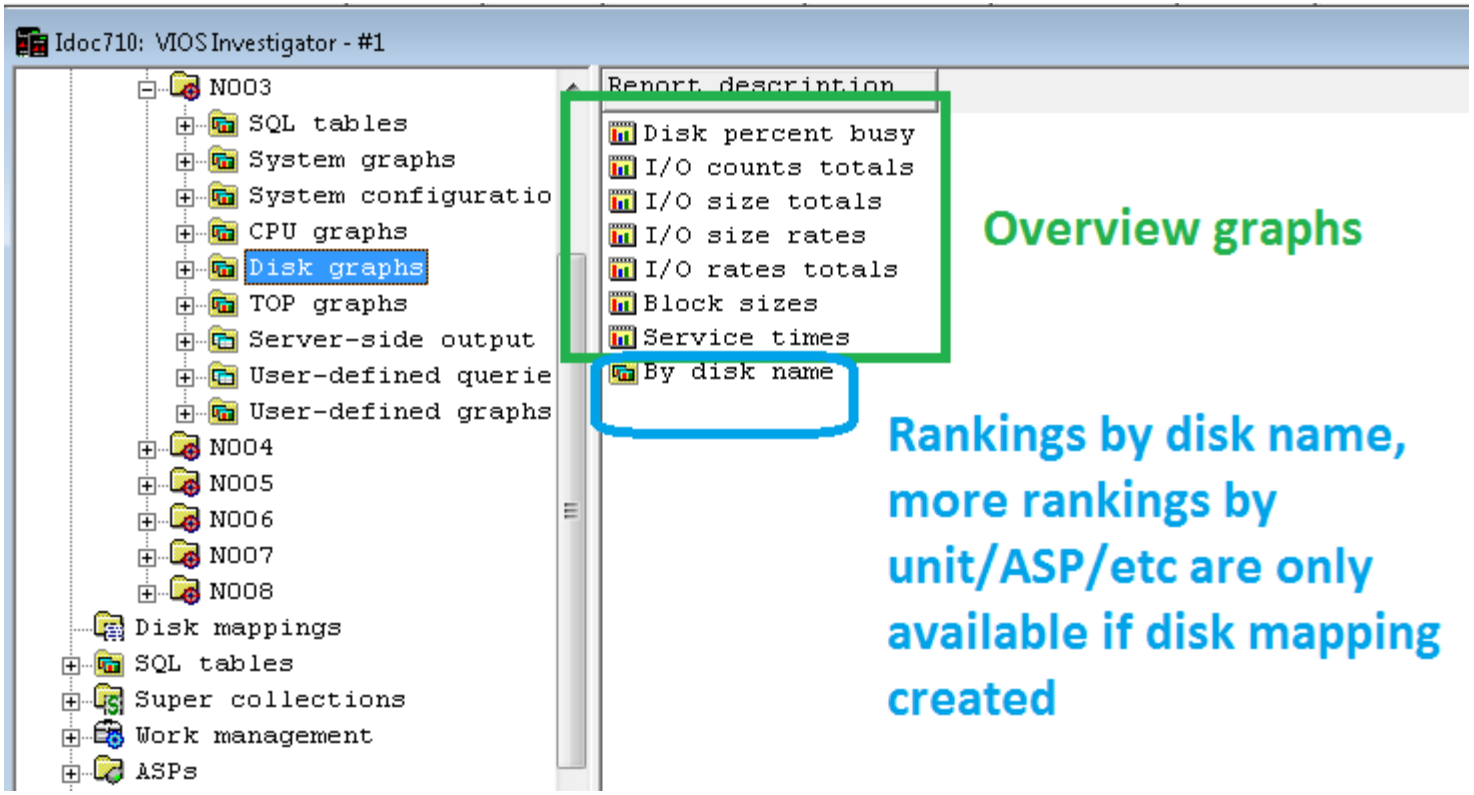

Analysis Options – Disk graphs available WITH a disk mapping

- **Under the collection, expand the Disk Graphs folder**

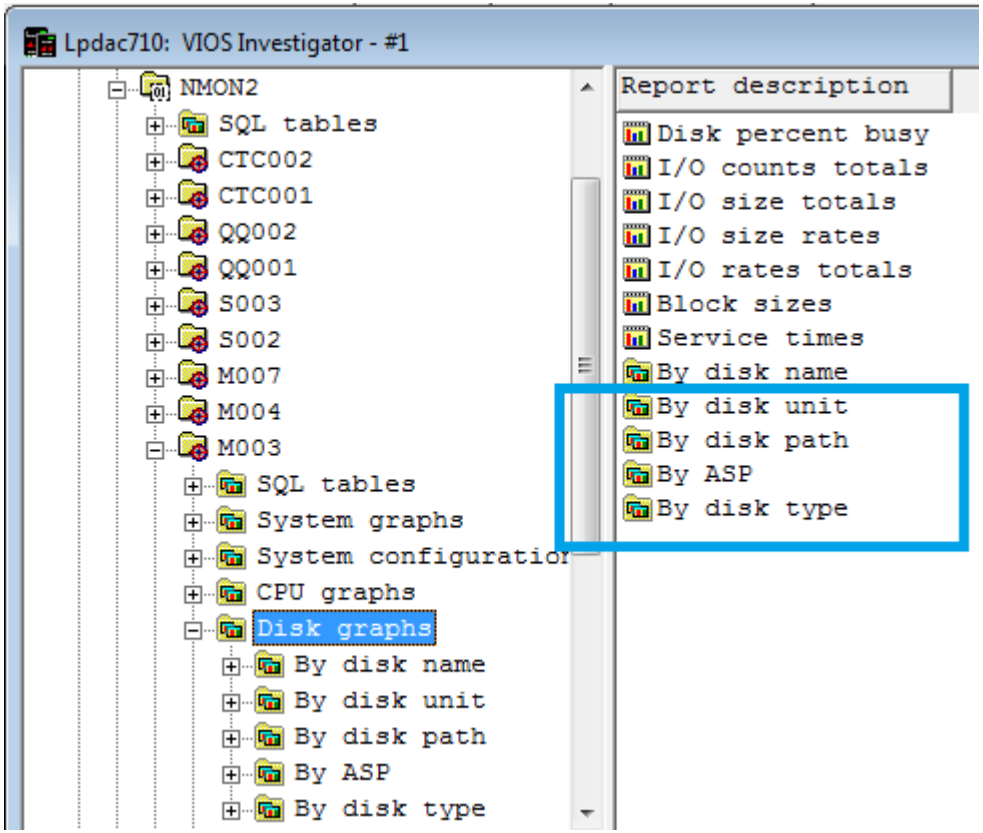

Analysis Options – Disk path rankings graph (requires a Disk Mapping)

- **Example showing I/O rates per IBM i unit/resource name.**

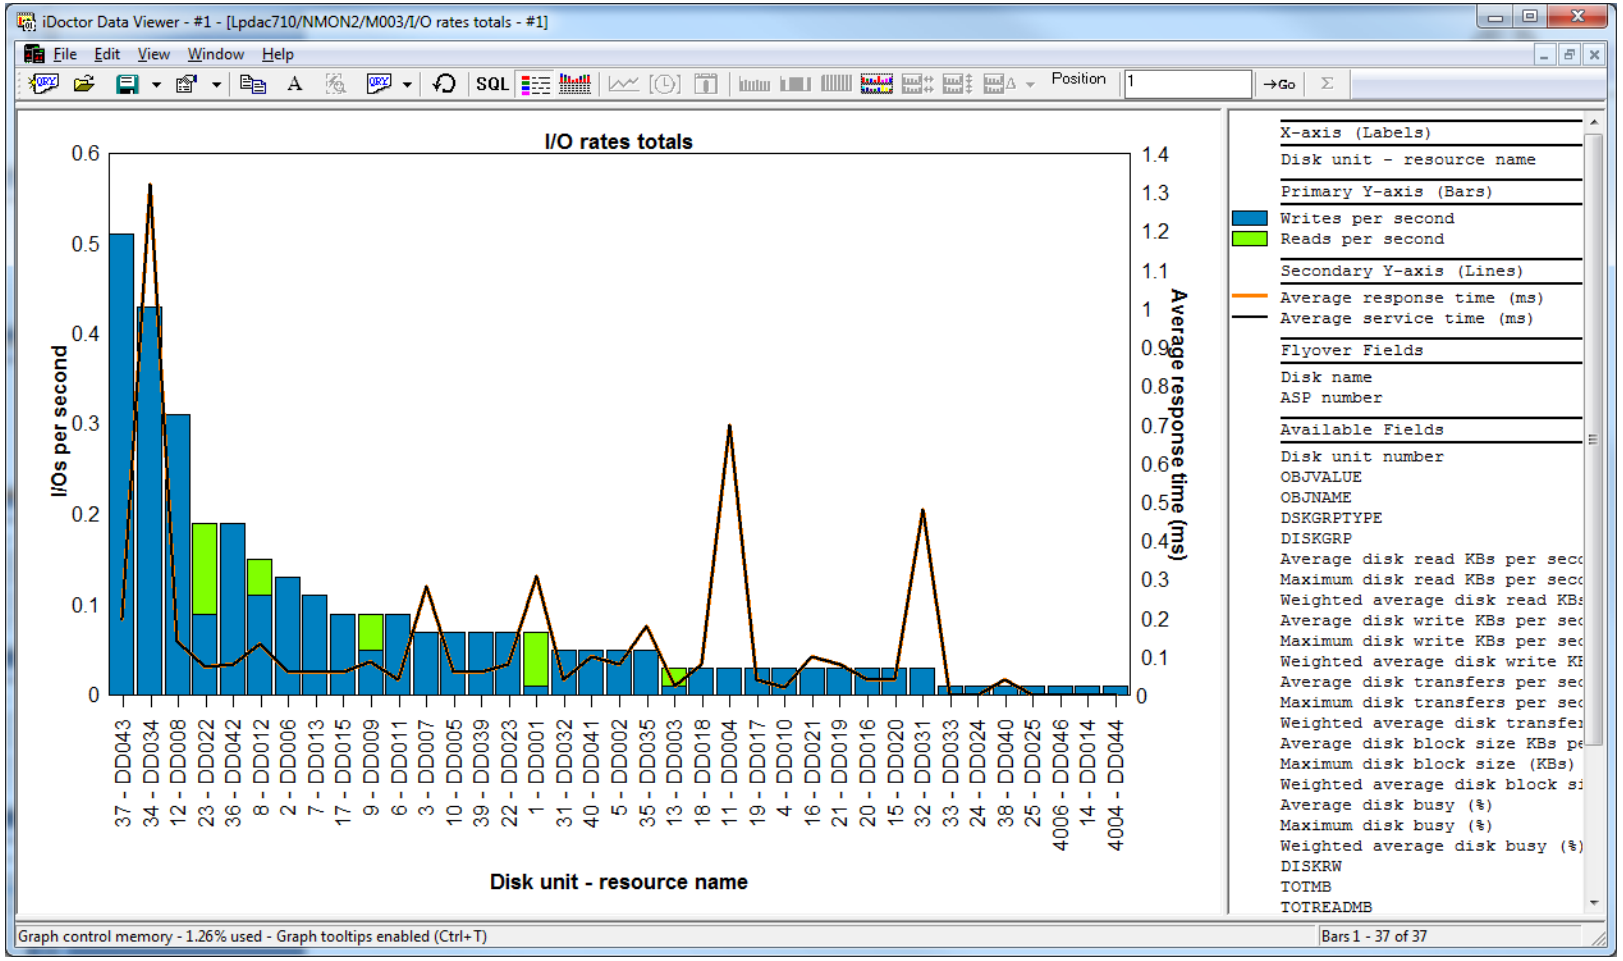

Analysis Options – ASP rankings graph (requires a Disk Mapping)

- **Example showing I/O size totals per ASP**

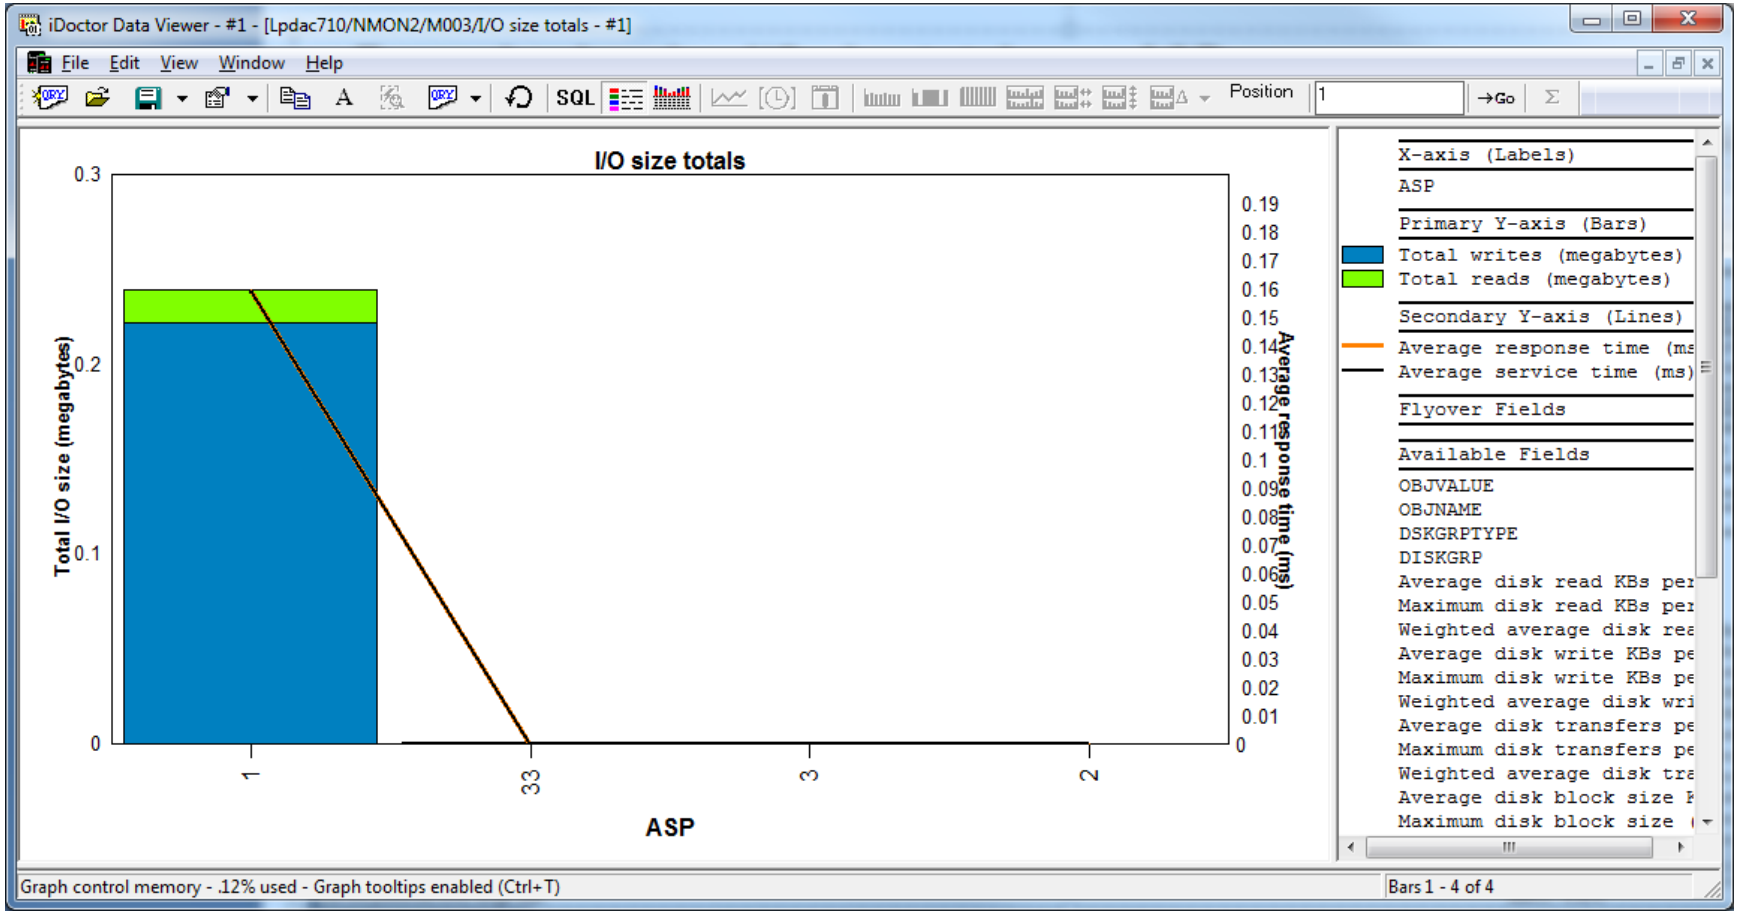

Analysis Options – Disk percent busy (overview)

- **This example graph shows "weighted average disk percent busy" and response times over 20 minutes (using 1 minute intervals.)**

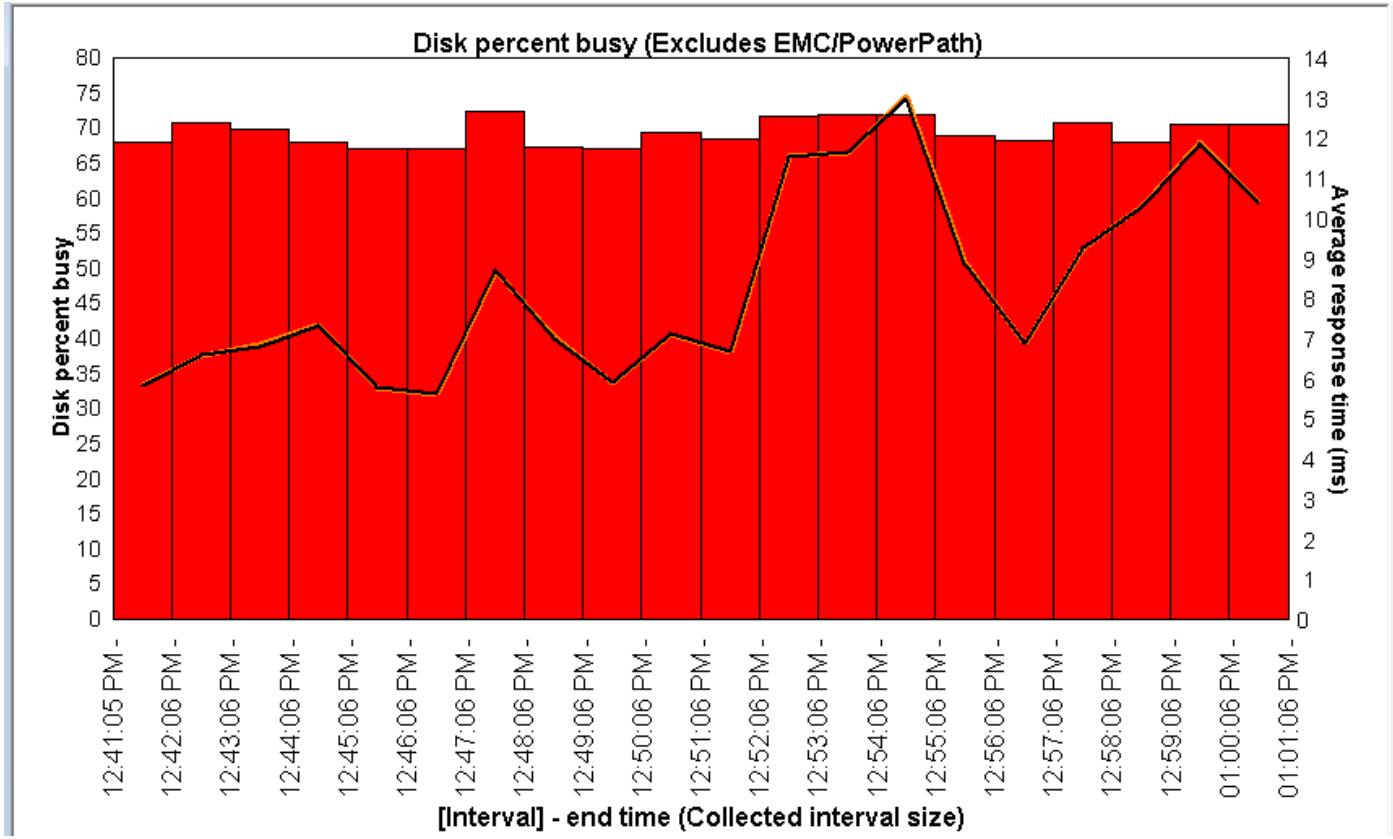

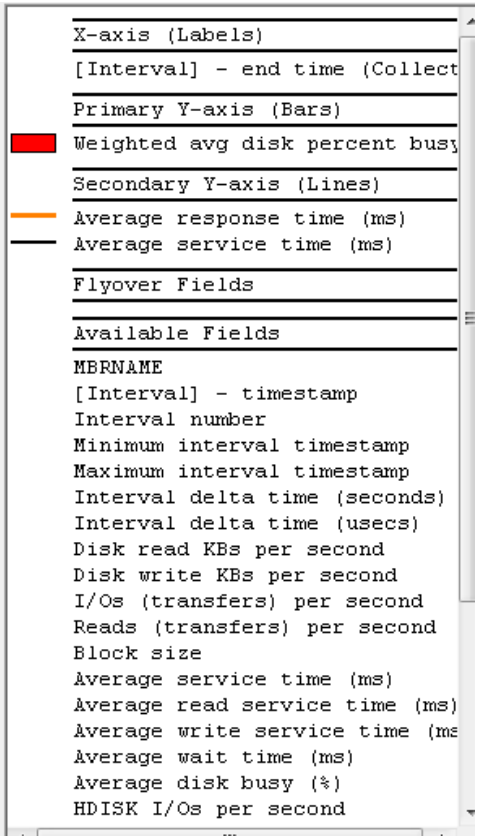

Analysis Options – I/O counts totals (overview)

- **This example graph shows total reads and writes with response times.**

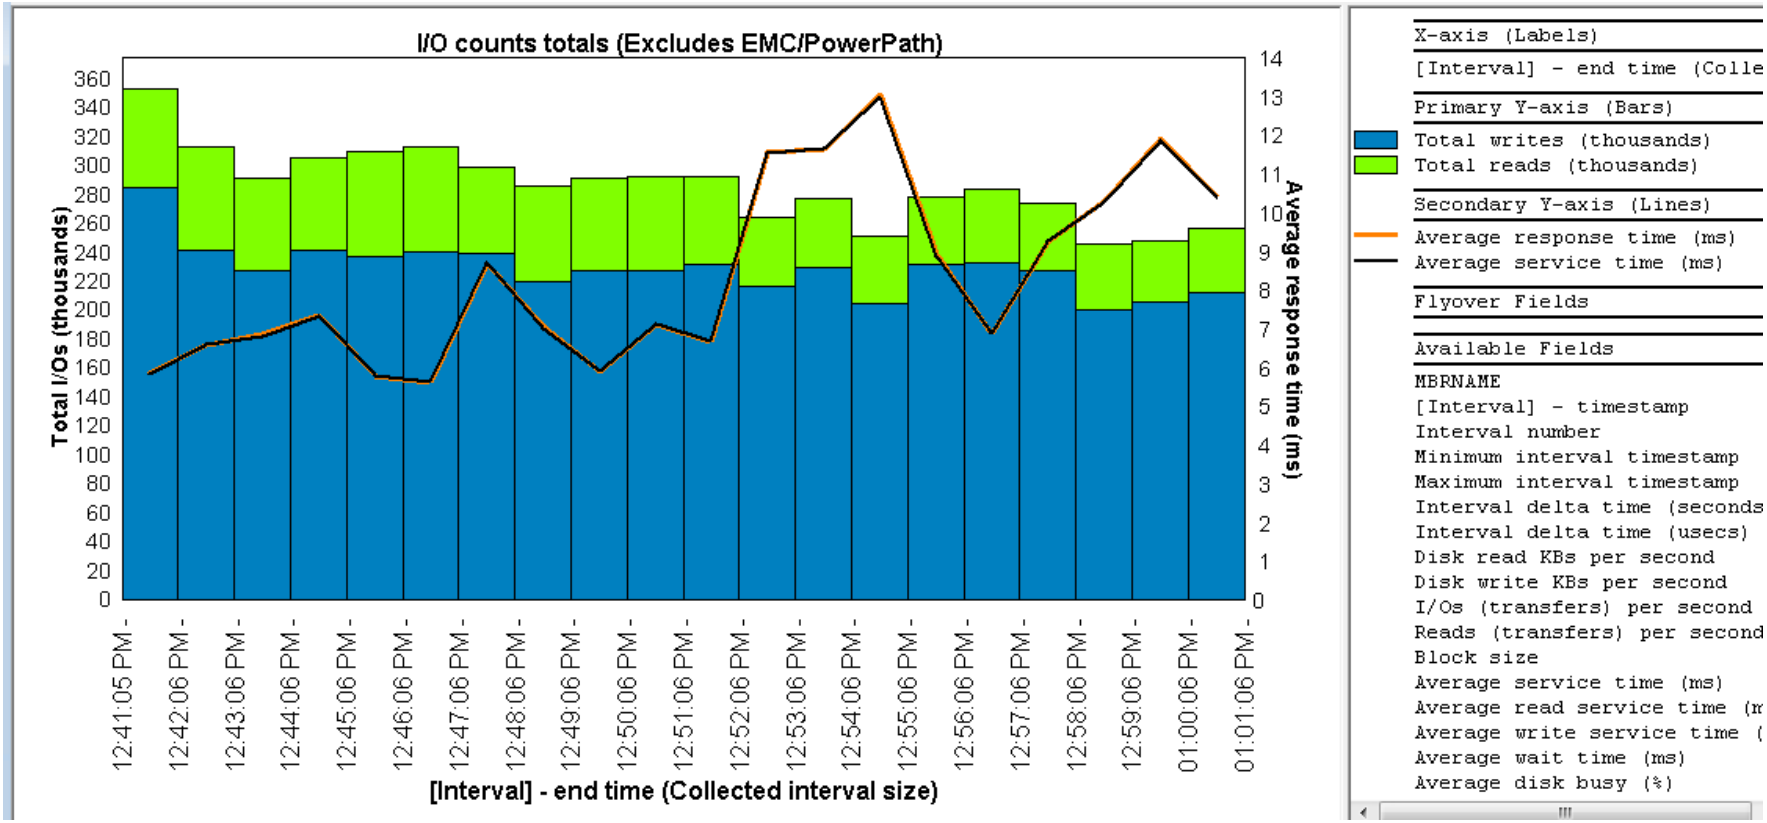

Analysis Options – I/O size totals (overview)

- **This example graph shows total I/O sizes for reads and writes with response times (this one shows 5, 20 minute collections)**

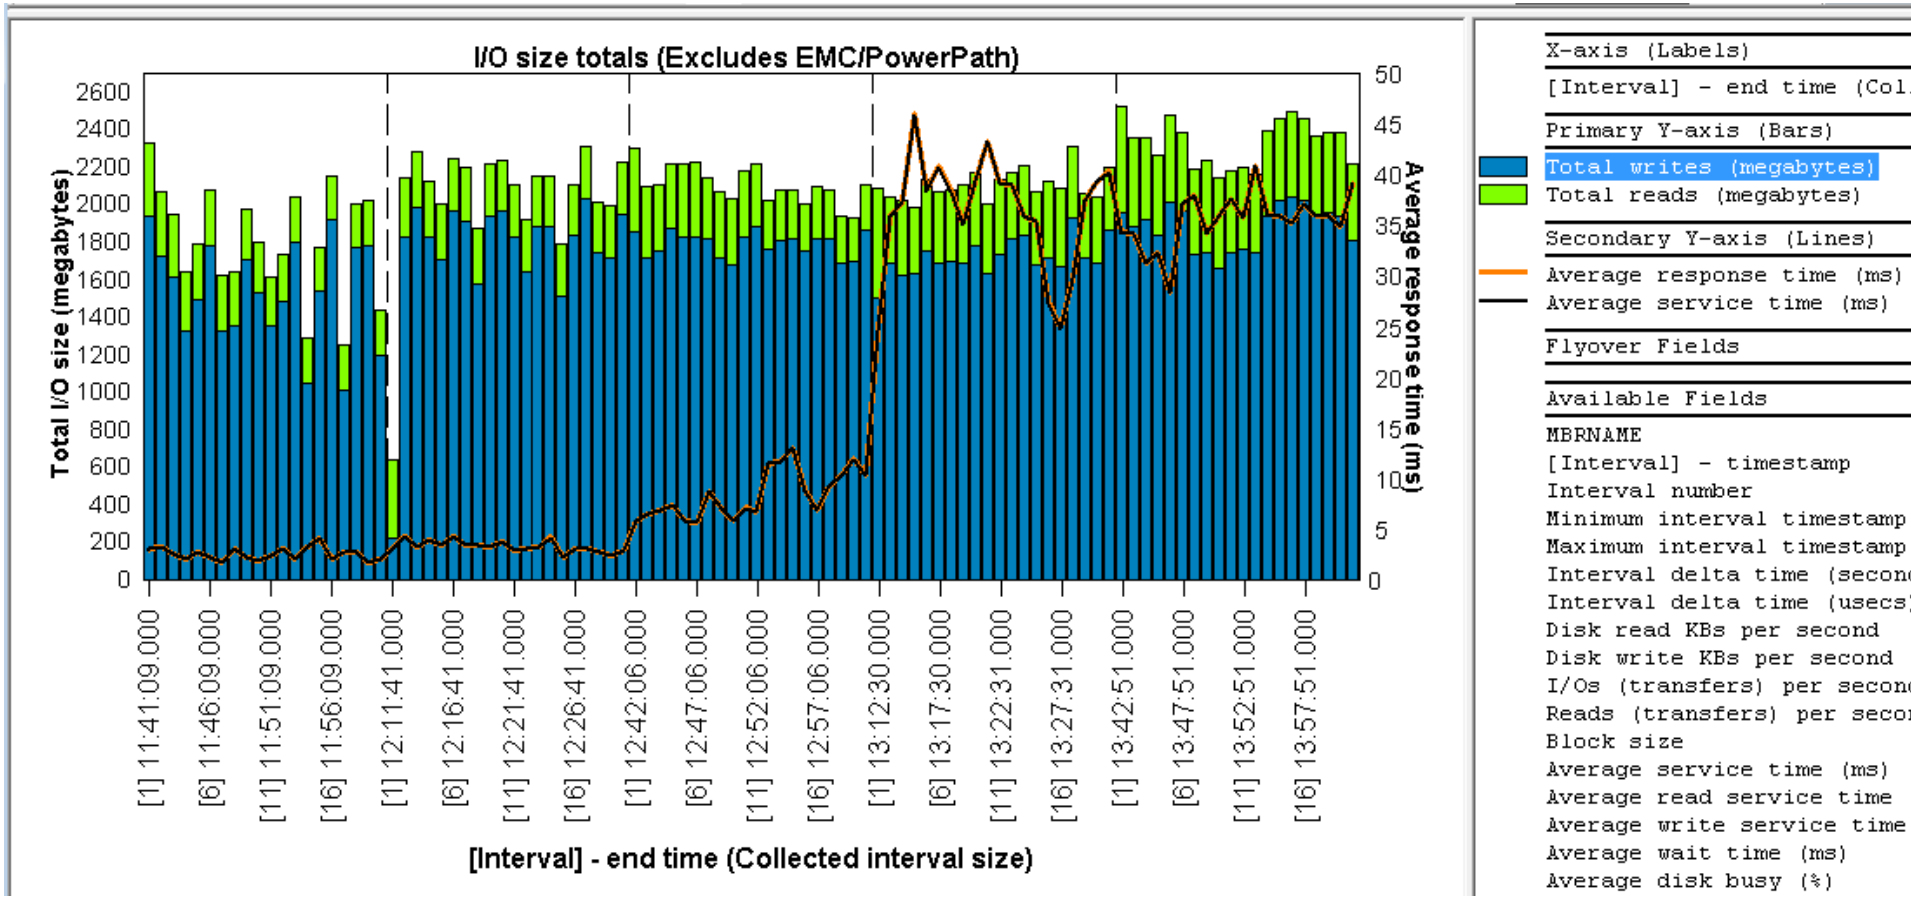

Analysis Options – I/O size rates (overview)

- **This example graph shows I/O size rates (in megabytes per second) and response times.**

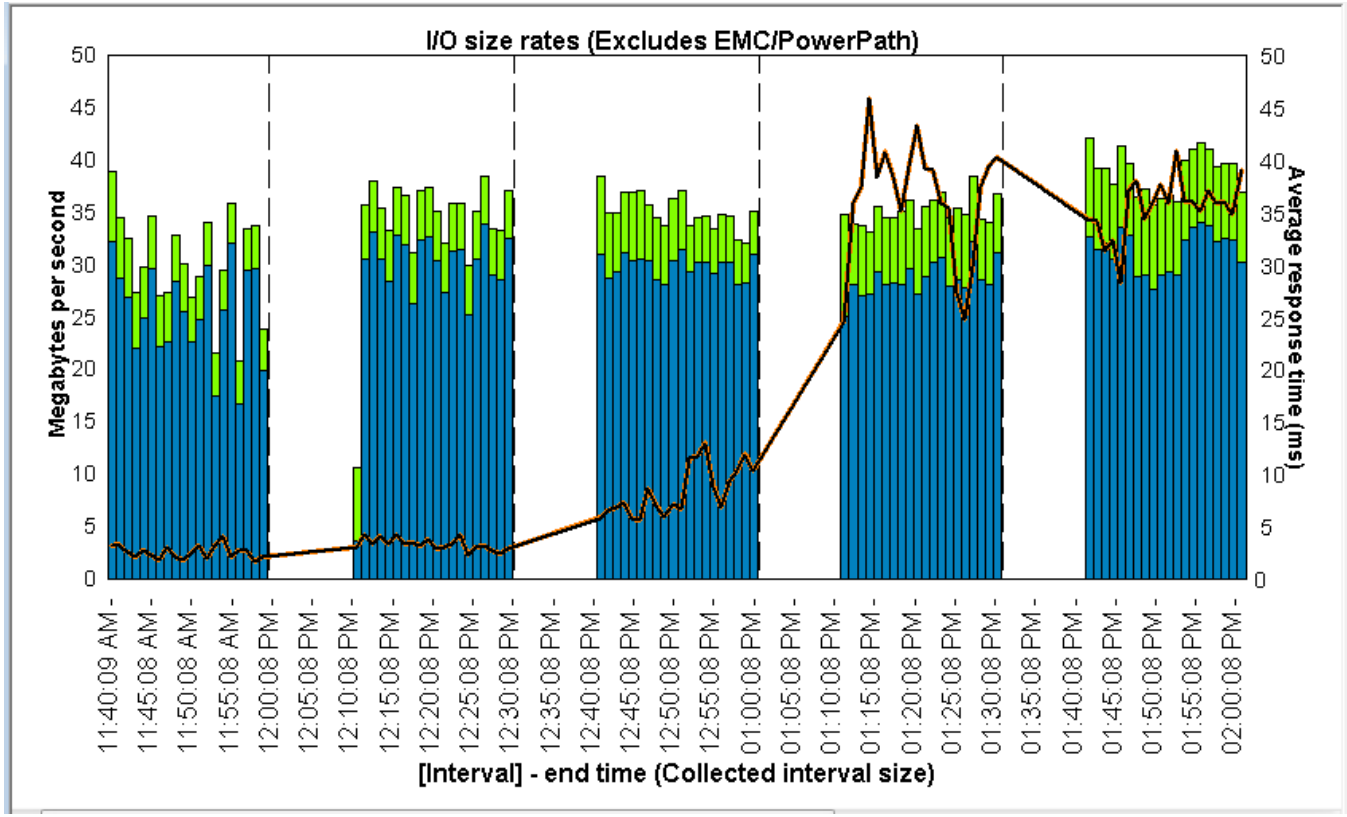

X-axis (Labels)
[Interval] - end time (Collect
Primary Y-axis (Bars)
Writes per second (megabytes)
Reads per second (megabytes)
Secondary Y-axis (Lines)
Average response time (ms)
Average service time (ms)
Flyover Fields
Available Fields
MBRNAME
[Interval] - timestamp
Interval number
Minimum interval timestamp
Maximum interval timestamp
Interval delta time (seconds)
Interval delta time (usecs)
Disk read KBs per second
Disk write KBs per second
I/Os (transfers) per second
Reads (transfers) per second
Block size
Average service time (ms)
Average read service time (ms)
Average write service time (ms
Average wait time (ms)
Average disk busy (%)

Analysis Options – I/O rates totals (overview)

- **The graph shows total reads and writes per second with average response times.**

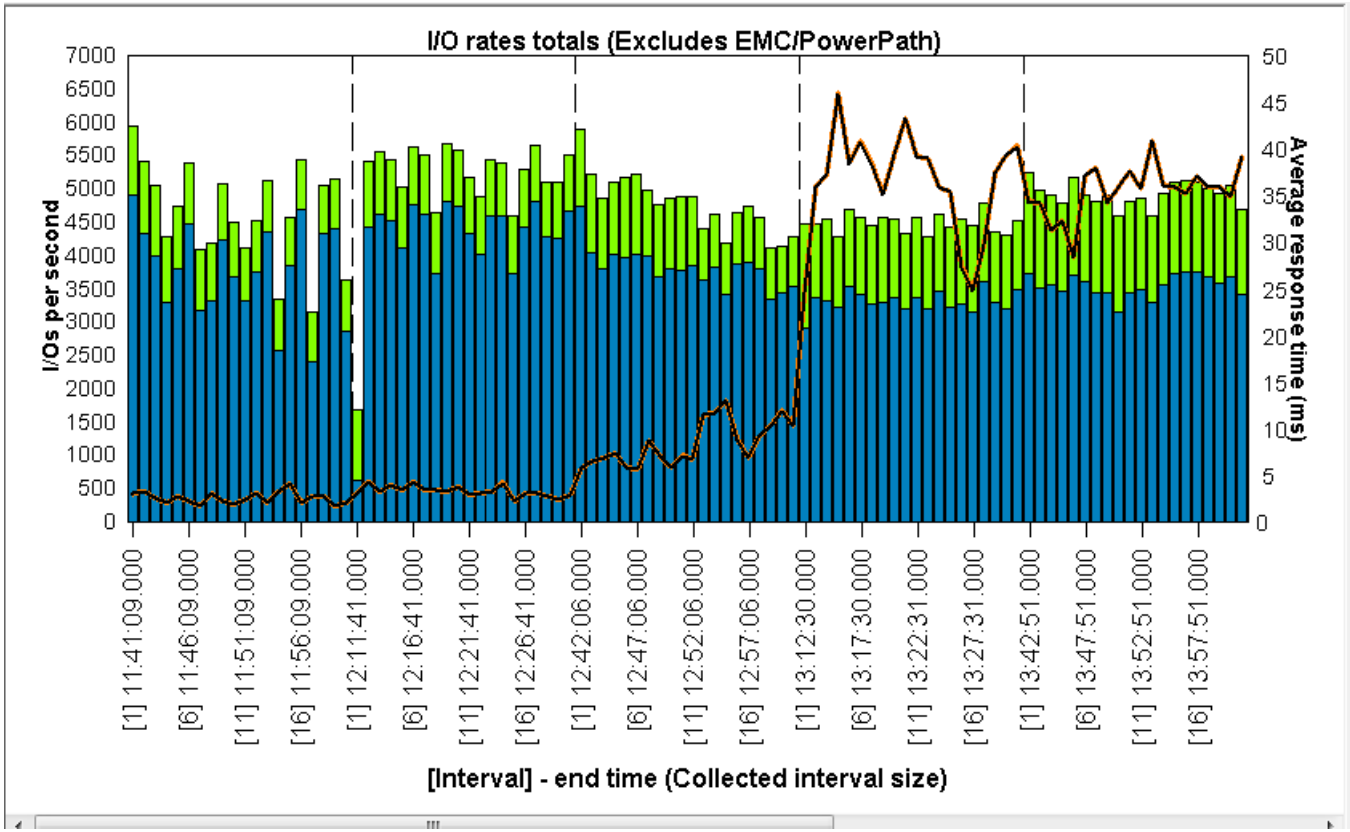

Analysis Options – Block sizes (overview)

- **This graph shows the average block size over time.**

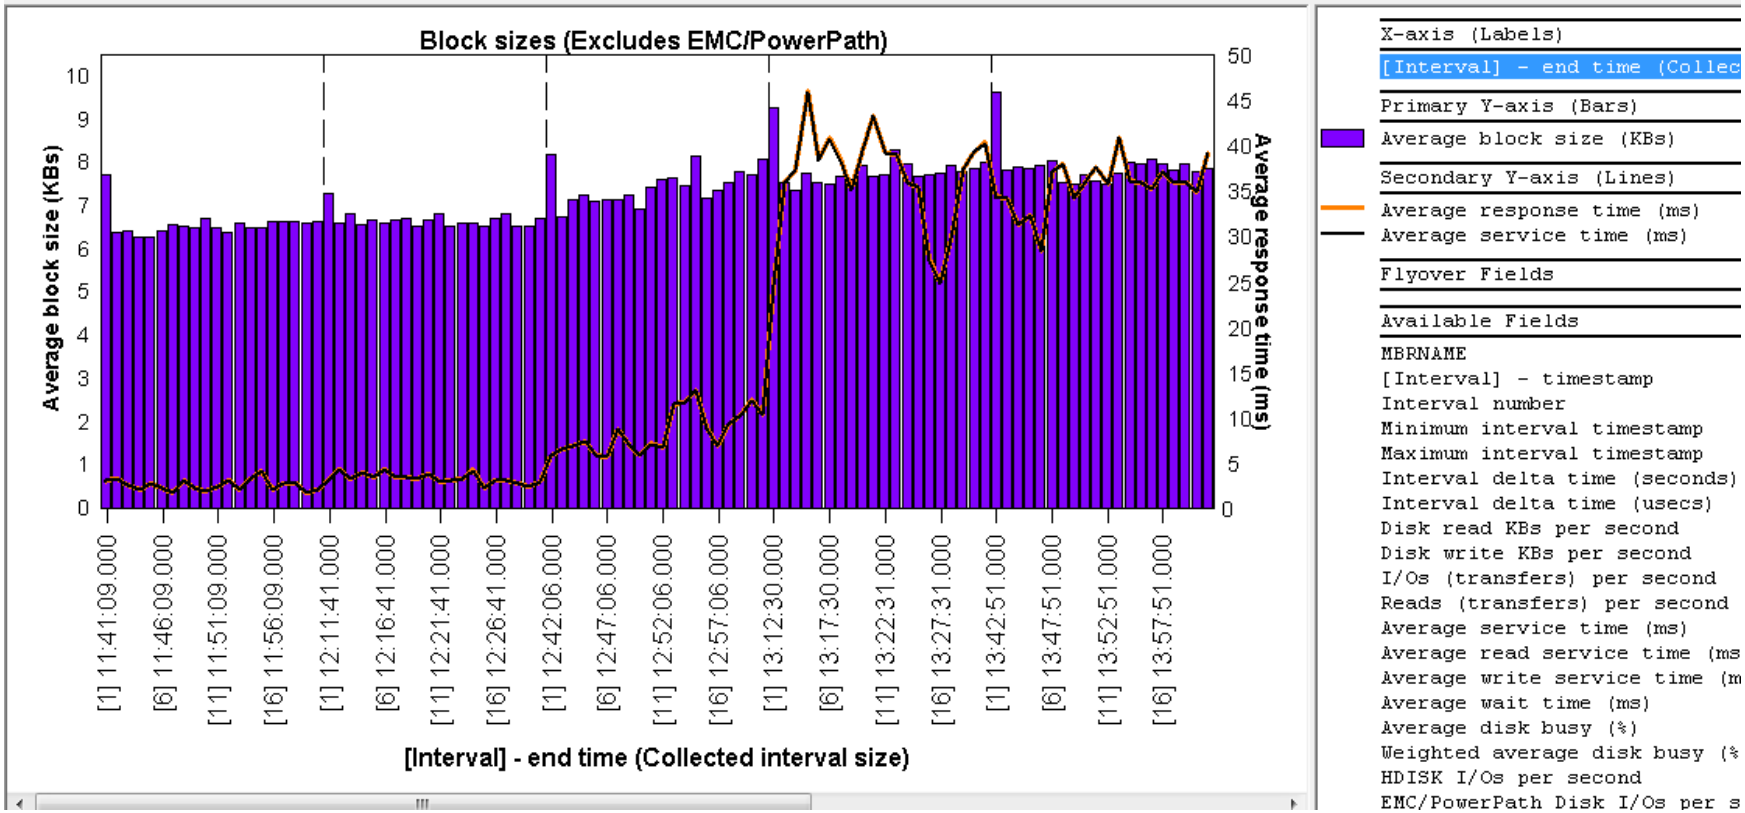

Analysis Options – Service times (overview)

- **This graph shows the average read and write service times**

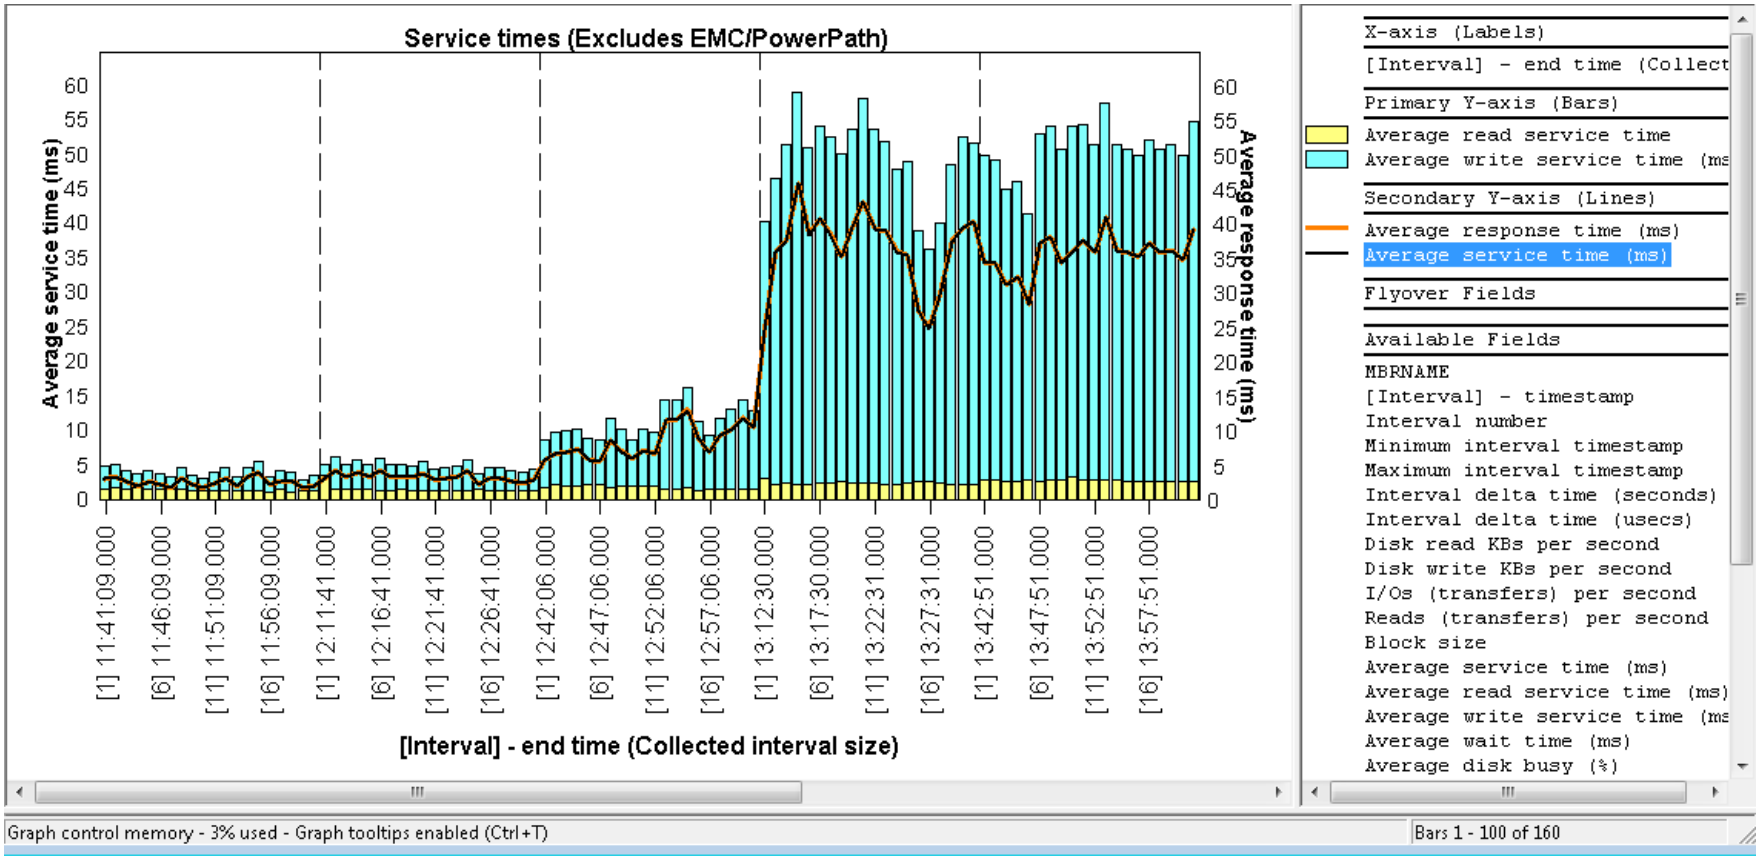

Analysis Options – Drilling down into rankings (from overview charts)

- **From the overview, select a time period, right-click and drill down using Disk graphs -> by disk name menu**

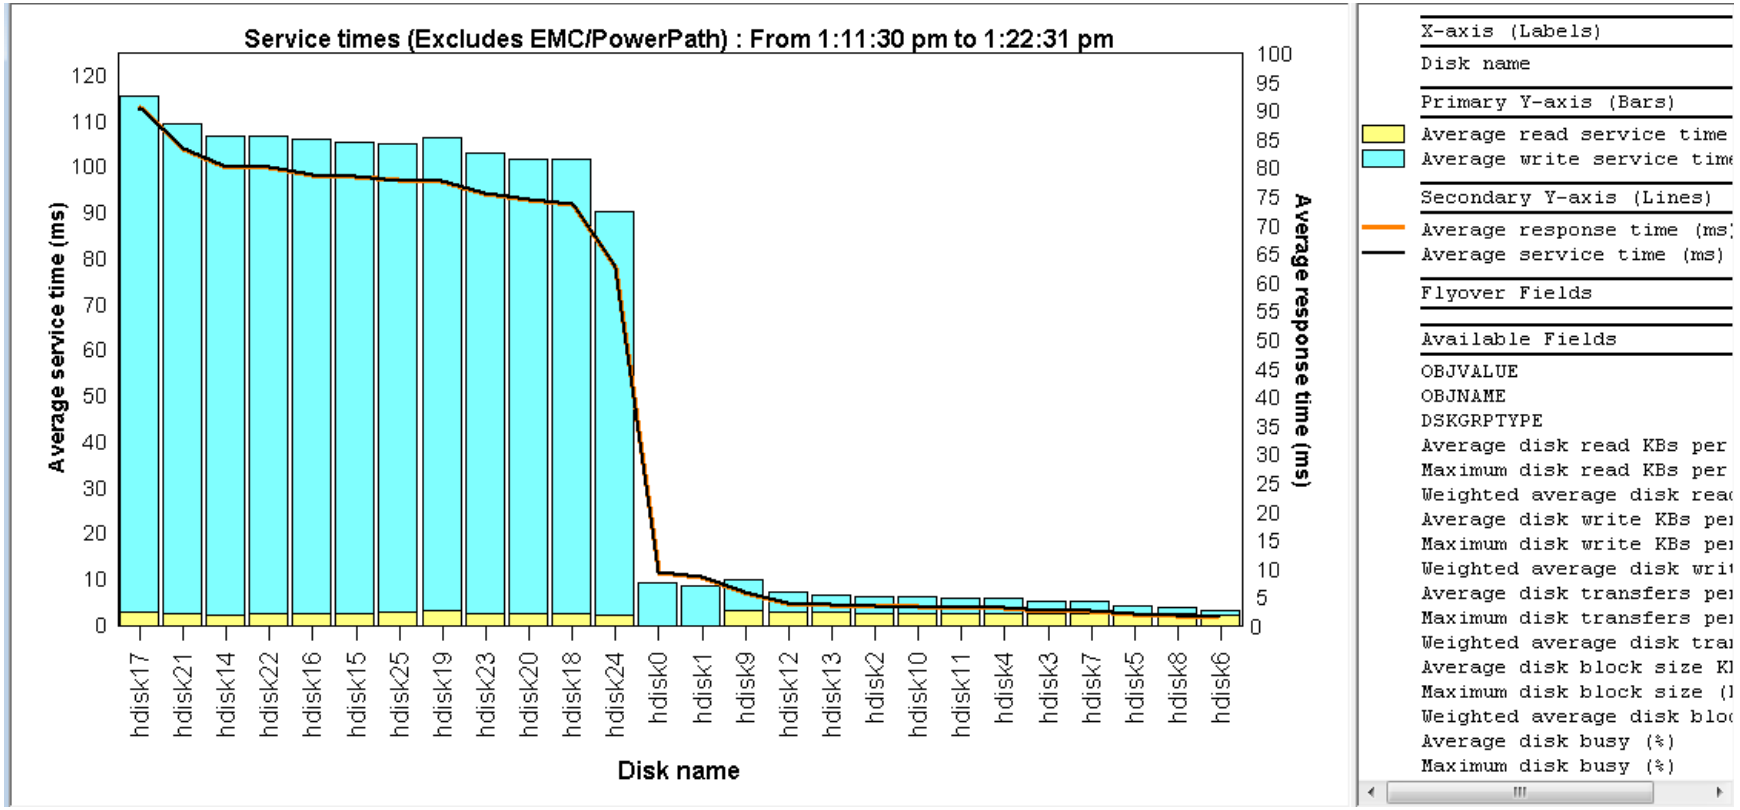

Analysis Options – Drilling down into rankings

- **Next if you wish to see a single disk over time, right-click it and drill down again.**

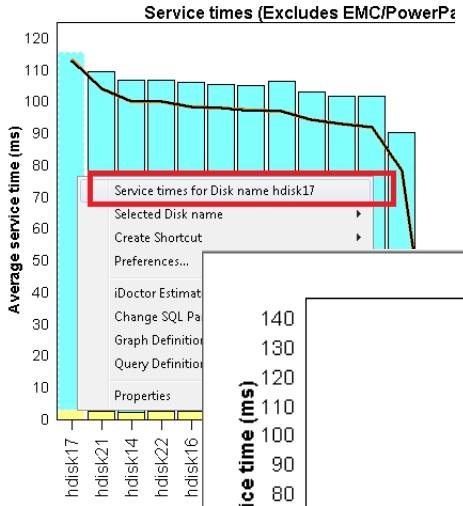

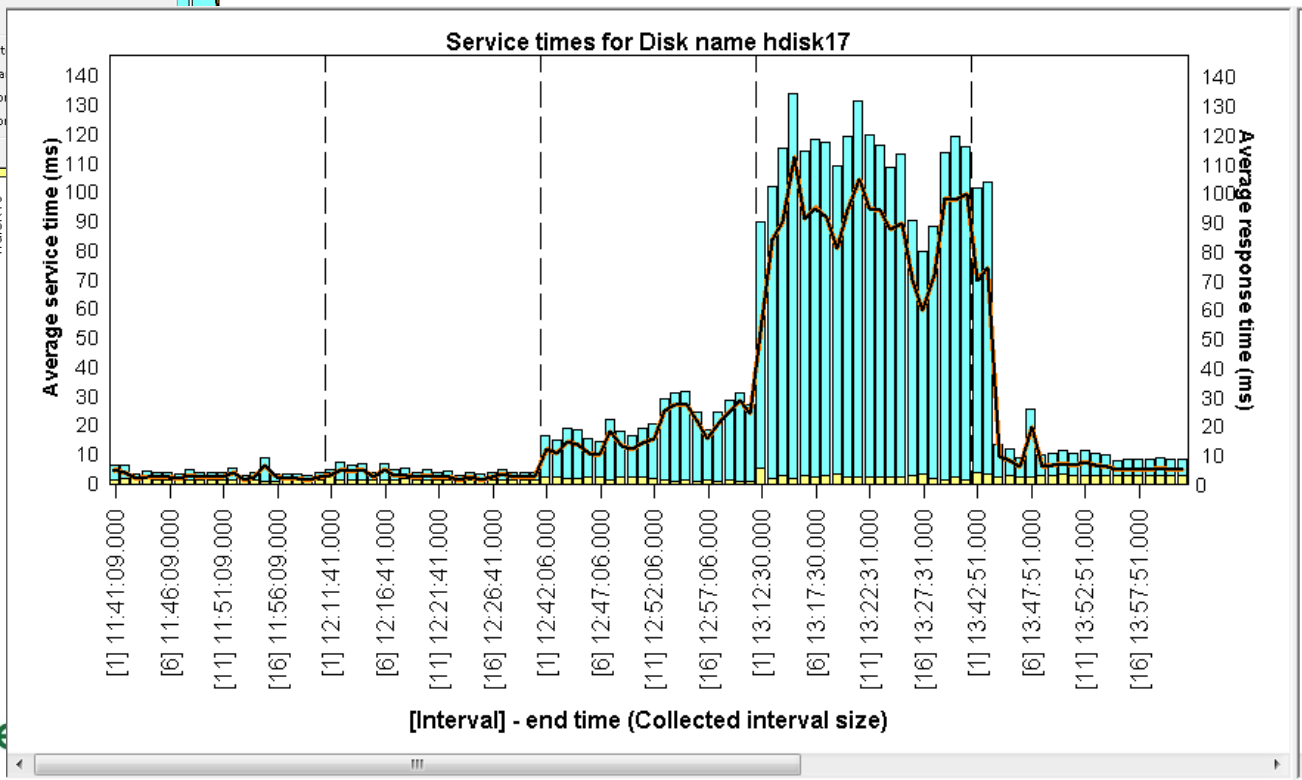

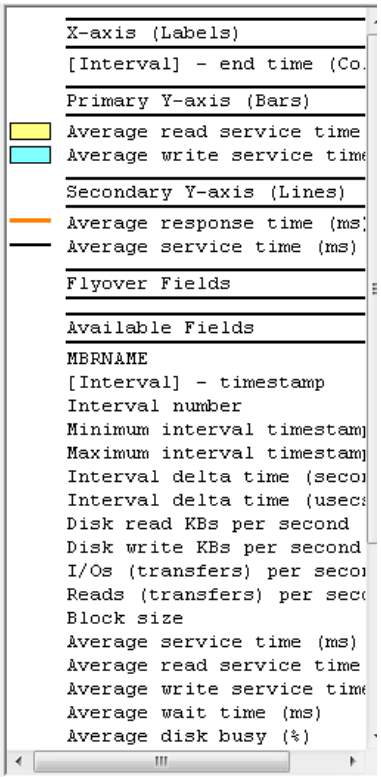

Analysis Options – Disk mappings benefits

- **But what device resource name / disk unit on IBM i is “hdisk17”?**
  - We can’t tell from this data alone. This is why we need the disk mapping discussed previously.
  - We also get these additional disk rankings graphs to group the data in additional ways (disk unit, path, ASP, disk type)

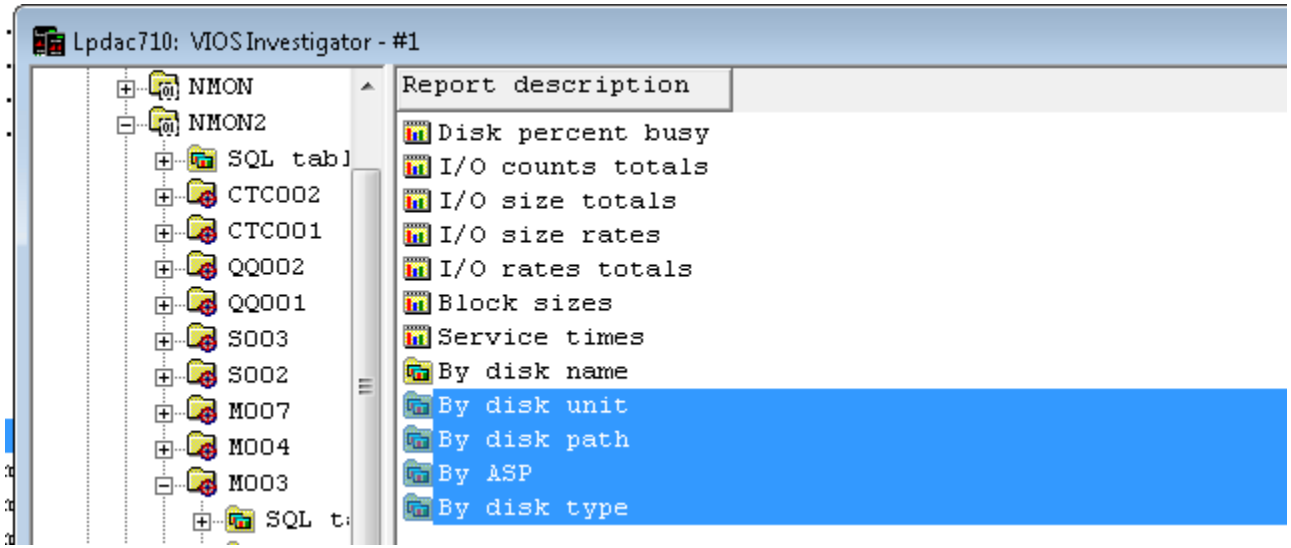

Analysis Options – Disk mappings benefits – using the flyover

- **If disk mapping is included, you can use the flyover to tell the disk correlated with the logical disk unit number on IBM i.**

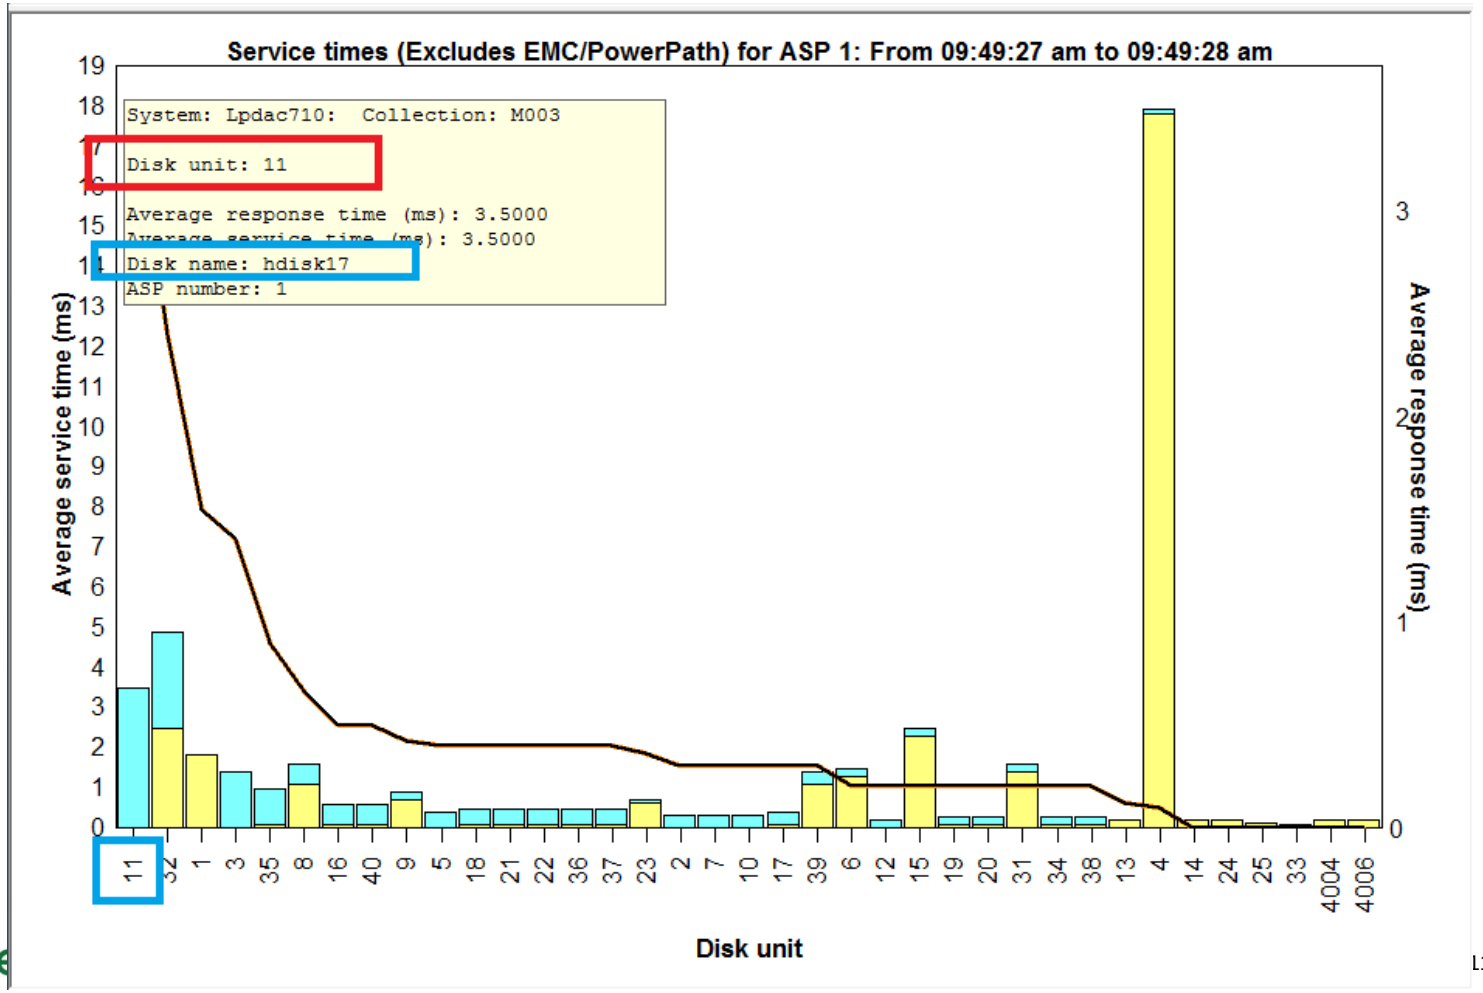

Analysis Options – VIOS Investigator Preferences

- **Under Preferences -> VIOS Investigator tab**

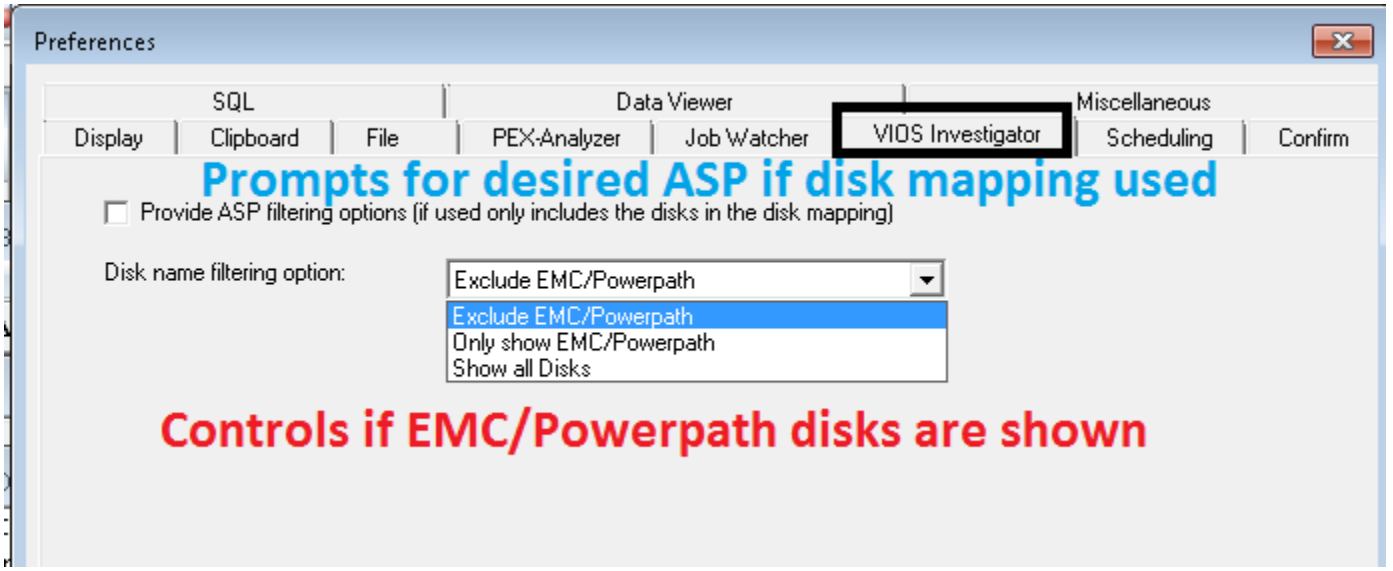

Analysis Options – Server-side output files

This contains the iDoctor generated SQL tables from the .NMON file.

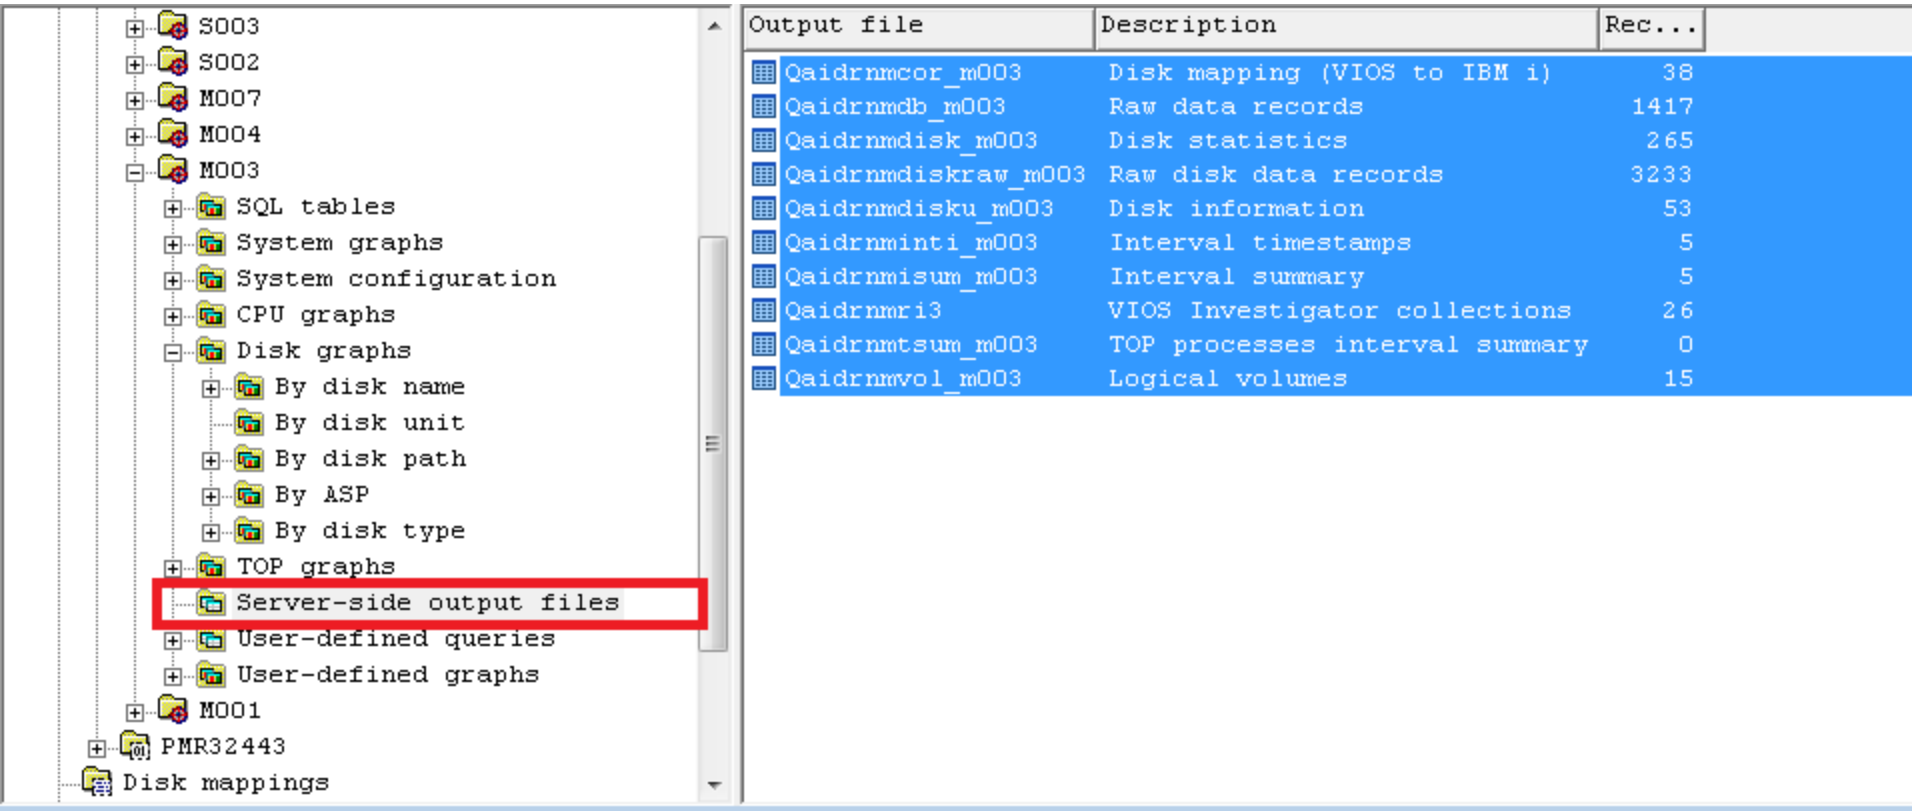

Analysis Options – System configuration folder

This contains reports over the NMON BBB* records

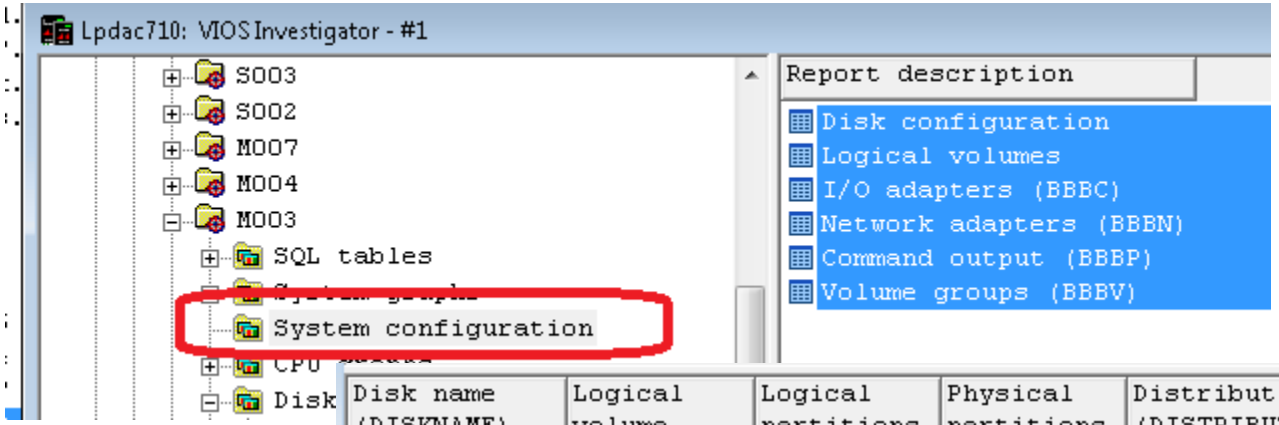

Disk name (DISKNAME)	Logical volume name (LVNAME)	Logical partitions (LPS)	Physical partitions (PPS)	Distributions (DISTRIBUTION)	Mount point (MOUNTPOINT)
hdisk0	hd1	20	20	00..00..20..00..00	/home
hdisk0	hd10opt	6	6	00..00..06..00..00	/opt
hdisk0	hd11admin	1	1	00..00..01..00..00	/admin
hdisk0	hd2	12	12	00..00..12..00..00	/usr
hdisk0	hd3	7	7	00..00..07..00..00	/tmp
hdisk0	hd4	1	1	00..00..01..00..00	/
hdisk0	hd5	1	1	01..00..00..00..00	N/A
hdisk0	hd6	1	1	00..01..00..00..00	N/A
hdisk0	hd8	1	1	00..00..01..00..00	N/A
hdisk0	hd9var	2	2	00..00..02..00..00	/var
hdisk0	lg_dumplv	2	2	00..02..00..00..00	N/A
hdisk0	livedump	1	1	00..01..00..00..00	/var/adm/ras/livedump
hdisk0	paging00	2	2	00..00..02..00..00	N/A
hdisk1	test	1	1	00..01..00..00..00	N/A
hdisk1	VMLibrary	40	40	00..40..00..00..00	/var/vio/VMLibrary

Analysis Options – More graphs available

- **Lots more graphs under System graphs, TOP graphs, CPU graphs**
- **Keep in mind you can use this tool for analyzing any system that can run NMON (AIX/Linux/VIOS)**

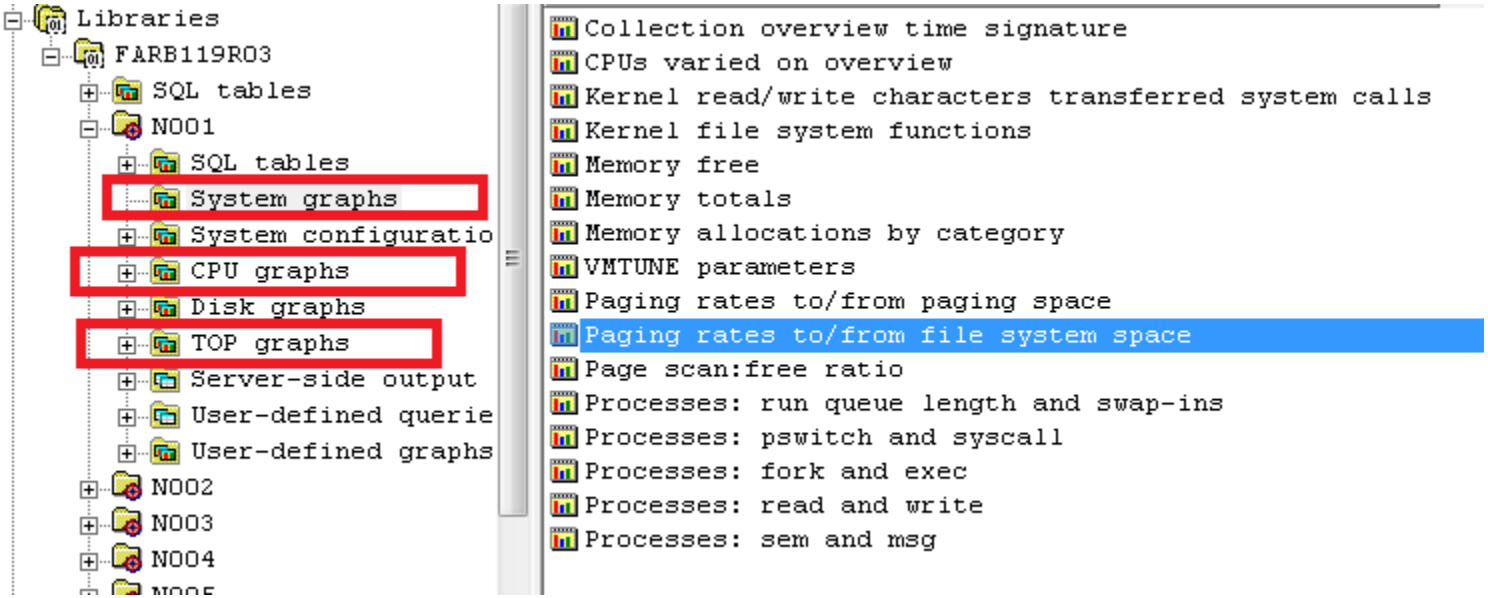

Introduction – Power is Performance Redefined

- For the past 10 years, through sustained investment in the Power Systems platform, IBM has gone head-to-head with our competitors in the UNIX market segment, and we won. Today, according to IDC, IBM commands a 47 percent share of the worldwide UNIX market segment.[1] The next 10 years, however, will be about helping our clients implement smarter computing. And for the Power Systems platform, that battle will center on our alternative value proposition to Linux and Microsoft Windows technology on x86 servers. To do that, we need to move beyond talking about pure system performance and industry benchmarks to placing a sharper focus on the performance of our clients' businesses and the business benefits of the IBM Power Systems platform.

- Industry benchmarks and our IBM POWER® processor technology are, and always will be, important. In the past they have enabled us to clearly and succinctly demonstrate our leadership position in terms of POWER processor performance versus our competitors. And we will continue to set those leadership benchmarks for the industry. But today the conversation must go beyond the performance of our systems and be framed in the broader context of smarter computing. *Power is performance redefined* sets out how we intend to shift the conversation with our clients. It defines how the Power Systems platform, and our associated software and services, can enable our clients to embrace smarter computing and derive business benefits from implementing big data, workload optimized infrastructure and cloud projects. In this messaging guide, you will learn that smarter computing isn't a product we sell; it isn't something clients can buy. Smarter computing is something our clients can implement through projects on the Power Systems platform to achieve better business outcomes. And it is smarter computing, enabled by IBM Power Systems servers, that will help our clients deliver services faster, with higher quality and with superior economics.

1 – IDC, "UNIX Server Rolling Four Quarter Average Revenue Share," Worldwide Quarterly Server Tracker, 2Q2011.

Smarter Computing – The Next Era of IT

- But this radical change is placing enormous pressure on businesses of practically every size, in just about every industry. The barriers of entry for competitors are lower. Companies can be blindsided by competitors that appear seemingly out of nowhere and seize market share by the handful. The need to be proactive, which requires an agile, flexible human and IT infrastructure, is critical. New ways of working, such as social media and mobile technologies, must be embraced ahead of the curve. Even customers are changing. Newly empowered by information, their expectations and the number of influencers that must be marketed to are rising. External forces such as compliance, regulations, privacy and security threats have to be addressed to survive. Ubiquitous mobile devices and instrumented, intelligent objects are creating unimaginable amounts of data volumes every day—data that must be analyzed to reveal systemic patterns, trends and insights that in turn inform the decisions businesses must make to stay competitive. And to deal with these changes, IT architectures must move from heterogeneous silos to flexible, workload optimized infrastructures. All of these forces must be dealt with in an era of tighter budgets and the directive to do more with less.
- But smarter companies are thinking differently about computing and how to deal with data that is growing exponentially and can become stagnant and unexploited simply because of its sheer volume. These smarter companies are breaking the vicious cycle of untrustworthy data, inflexibility and sprawl. They are reversing the always-guessing, reactive, costly IT conundrum by embracing what we call smarter computing. What smarter computing entails is the creation of an IT infrastructure that is **designed for data** and that harnesses enterprise information to unlock insights and make better, more informed choices. Organizations embracing smarter computing are creating IT infrastructures that are **tuned to the task** of the business, helping reduce costs by driving greater efficiency and performance for virtually every workload. And smarter computing is **managed with cloud technologies**, speeding delivery of services and creating an IT environment that has practically no boundaries, enabling the reinvention of processes and driving innovation.
- But to be clear, smarter computing isn't just a catch phrase or a lofty idea. It's not a metaphor, intro paragraph or headline. It's what the IBM Power Systems platform enables our clients to do. And this is the basis for our new brand identity *Power is performance redefined*. It's about how we believe clients measure IT performance – focusing less on processor performance and more on business performance. It's about our clients' ability to react more quickly to change, to innovate faster, and to seize new opportunities as they arise. It's about their ability to handle rapid growth and combat emerging competitors while responding to demands to meet increasingly higher service levels. And it's about doing more with less and delivering services within constrained IT budgets. We believe that with a new focus on business performance, we will enable our clients to deliver services faster, with higher quality and superior economics. Our message to clients is that, with Power Systems solutions, we can help them achieve these goals as they deploy smarter computing projects.

Power is Performance Redefined

- **In this new smarter computing era for business and IT, forward-thinking companies consider more than server performance, existing skills and ease of management when choosing a platform for new application workloads. They also evaluate how well the platform will help them achieve three core business objectives: delivering services faster, with higher quality and superior economics.**
- **By implementing smarter computing projects on an IBM Power Systems platform, businesses can outpace their competitors by delivering services faster. They can differentiate their offerings from the competition by delivering higher quality services. And they can turn operational cost into investment opportunity by delivering services with superior economics.**

Power is Performance Redefined

- **Deliver services faster**
  - A key measure of performance for IT today is around agility and the ability of IT to help the business gain a competitive edge and capitalize on emerging opportunities. Businesses need to simplify and integrate their IT infrastructure to *deliver services faster.*
  - The IBM Power Systems platform features deep integration and optimization across operating systems, databases and middleware for simpler, and more flexible, service delivery. Optimized with PowerVM virtualization for rapid cloud provisioning, clients can speed the delivery and deployment of new applications and processes to support their strategic business initiatives.

- **Deliver services with higher quality**
  - Today's IT departments are also measured on their ability to provide an infrastructure that can address demands for increased application service levels while at the same time balancing rapid change with managing business risk. Businesses need an integrated approach to managing security, resiliency and business risk to *deliver higher quality services.*
  - The IBM Power Systems platform, storage and software provide a highly secure and resilient infrastructure foundation for smarter computing. In addition to the built-in reliability, availability and serviceability (RAS) characteristics of Power Systems servers and blades, our IBM System Storage®DS8000® and IBM Storwize®V7000 Unified storage systems, and IBM PowerHA SystemMirror clustering software is tightly integrated with our operating systems to provide a system-wide solution for business resilience.

- **Deliver services with superior economics**
  - IT performance today is also measured on its ability to maintain existing services and deliver services within tight budget constraints. In order to do more with less, businesses need to *deliver services with superior economics.*
  - The Power Systems platform with PowerVM virtualization is central to our differentiation when compared to x86 servers. PowerVM technology is designed to offer more secure and scalable virtualization than VMware on x86, enabling cost-effective control of server and virtual image sprawl. PowerVM technology also is designed to help Power Systems servers deliver higher server utilization rates than VMware on x86. We believe that the superior economic model for workload consolidation on POWER7 servers with PowerVM software has been the key driver behind migrations from Oracle Sun and HP to Power Systems technology.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, Workload Partitions Manager and X-Architecture 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 full list of U.S. trademarks owned by IBM may be found at: http://www.**ibm.com**/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
AltiVec is a trademark of Freescale Semiconductor, Inc.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

Revised December 2, 2010

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 2, 2010

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm
	Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	**http://hmmer.janelia.org/** http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised December 2, 2010

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- **rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.**

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

==========

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007