IBM Systems & Technology Group

Collecting MONWRITE Data

Revision 2008-07-30 BKW

IBM z/VM Performance Evaluation
Brian Wade bkw@us.ibm.com
Bill Bitner bitnerb@us.ibm.com

Trademarks

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DBE, e-business logo, ESCO, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, VM/ESA, VSE/ESA, Websphere, xSeries, z/OS, zSeries, z/VM

The following are trademarks or registered trademarks of other companies

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
LINUX is a registered trademark of Linus Torvalds
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.
Intel is a registered trademark of Intel Corporation
* All other products may be trademarks or registered trademarks of their respective companies.

NOTES:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

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

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

Agenda

- **What is raw monitor data?**
- **How do I set up to collect it?**
- **When do I collect it?**
- **What tools are available to help me collect it?**
- **How do I package it for transmission?**
- **How do I study it myself?**
- **Summary**

What is Raw Monitor Data?

- **It is unformatted binary data describing system configuration or activity**
- **Logically, it is a sequence of *monitor records***
  - Each record comments on some specific aspect of system activity or performance
  - In aggregate they constitute a comprehensive, time-indexed record of system activity
- **There are three large classes of monitor records**
  - *Configuration records:* emitted when monitor starts, these describe system configuration
  - *Sample records:* emitted every so often, these comment on the accumulated activity of an entity (device, user, …)
  - *Event records:* emitted as needed, these comment on some specific phenomenon that just now occurred
- **Some records come from the Control Program and comment on its experience in running the system**
- **Other records come from guests and comment on their experiences in doing whatever it is they do**
- **We collect this data using an IBM-supplied utility program called MONWRITE**
- **During the rest of this presentation, we will call this data *MONWRITE data***

How Do I Collect MONWRITE Data?

- **By Default the z/VM system is set up with DCSS and user ID named MONWRITE**
- **If somehow skipped, then:**
  - You set up a DCSS where CP will buffer the monitor records it emits
    - CP DEFSEG and SAVESEG commands
  - You tell CP which kinds of records to emit, and how often to emit them, and in fact to begin emitting them
    - CP MONITOR command
  - You set up a guest that drains the DCSS to a disk or a tape via the MONWRITE utility
- **On some occasions, the default DCSS (named MONDCSS) is too small.**
  - See http://www.vm.ibm.com/perf/tips/mondcss.html
- **You run the guest**
- **You archive the resultant files or tapes, so that you have a long-term historical record of system activity and performance**

When Do I Collect MONWRITE Data?

- **Periodically, collect and archive some data during your peak periods, so that you have a historical record**
  - Every Tuesday at 10 AM for an hour
  - Month-end processing
  - Whenever you do that really big thing you do
- **When directed by IBM**
  - PMR, crit sit, ESP, whatever

Tool: Running MONWRITE By Hand

- **A great idea, assuming you are not running some other performance product**
  - If you know what you are doing, you can do both simultaneously
- **Create the DCSS to hold the buffered records**
- **Set up a guest to run our MONWRITE MODULE (collector)**
- **Issue some CP MONITOR commands to start CP emitting records**
  - Enable all samples
  - Enable all events except seeks and scheduler
  - Use a 1-minute sample interval and a 5-second HFS rate
- **In your guest, start MONWRITE to collect the data CP's emitting**
- **To stop collecting, type this: MONWSTOP**
- **You will end up with one MONWRITE file that you can:**
  - Archive for the historical record
  - Analyze yourself with z/VM Performance Toolkit
  - Send to IBM so we can look at it
- **There is an option for MONWRITE to close the file at regular times of day and a user exit to process the just-closed file.**
- **Good references:**
  - http://www.vm.ibm.com/perf/tips/collect.html - a good cheat sheet from Bill Bitner
  - z/VM Performance, chapter 9, "Monitoring Performance Using CP Monitor" – an excellent writeup of every last detail

Tool: Brian's LINMON Collector

- **At http://www.vm.ibm.com/devpages/bkw/monitor.html**
- **Based on a modified MONWRITE**
- **Sets up the DCSS, CP Monitor, etc. on its own, using certain assumptions that are probably safe for many systems**
- **Can be configured to:**
  - Account for presence of another performance product
  - Collect for a while then log off
  - Start a new file every so often
  - Keep only the last N files

Packaging MONWRITE Data For Transmission

- **MONWRITE files are binary CMS files, F 4096.**
- **Just attaching them to an e-mail is NOT recommended.**
- **We suggest you use the VMARC file archiver that runs on CMS for large files or when several files are being sent.**
  - Kind of like “zipping” on a PC (compresses, combines)
    - MONWRITE data is very compressible
    - The data compression and decompression steps assure data integrity
    - Sometimes you also want to send us a console spool, or some QUERY outputs, or whatever
    - You can package everything into one VMARC archive and send us only that
- **The standard z/VM Level 2 process for FTPing files calls for COPYFILE (PACK**
  - This is unnecessary for MONWRITE and VMARC files.
- **Always, always, always:**
  - Move the files in binary (ASCII is a sure showstopper)
  - Do not use FTP’s SITE FIXREC, QUOTE SITE FIXREC, or LOCSITE FIXREC features (error-prone)
- **You will FTP your VMARC archive file to IBM’s receiving server in Boulder, CO**
  - Testcase.boulder.ibm.com, anonymous, cd /toibm/vm, binary, put, quit
  - Name your file mnemonically and send us a note or update the PMR
  - In PMR and/or note be clear as to what is sent and how packaged
  - See http://techsupport.services.ibm.com/390/tcprocs.html for additional Testcase help
- **More on VMARC at http://www.vm.ibm.com/devpages/bkw/monitor.html**

Studying MONWRITE Data

- **z/VM Performance Toolkit**
- **Interactively – possible, but not so useful**
- **PERFKIT BATCH command – pretty useful**
  - Control files tell Perfkit which reports to produce
  - You can then inspect the reports by hand or programmatically
- **See z/VM Performance Toolkit Reference for information on how to use PERFKIT BATCH**

Other Types of Data Confused with MONWRITE Data

- **Asking for "raw VM monitor" data can be confusing.**
  - Velocity has its own form of raw monitor data and history files, and even a form that attempts to mimic MONWRITE.
  - "VM Monitor" sounds like the "VM:product" often associated with CA products.
  - Performance Toolkit's history, trend, and summary files do not have the same detail.
- **Be specific when asking for data.**

Summary

- **MONWRITE data is the single most comprehensive record of system activity**
- **It is invaluable in diagnosing performance concerns**
- **If you ask IBM for performance help, IBM will very likely ask you for MONWRITE data**
- **Practice collecting and transmitting MONWRITE data when you are not under duress**
- **Archive your MONWRITE data routinely so that you have a good record of your system's usual behavior**
- **Learn to use PERFKIT BATCH to generate reports, and get familiar with a few of the basic reports**