


Jim Fisher
Advanced Technical Skills
fisherja@us.ibm.com



Performance Analysis Tools for the TS7700 Virtualization Engine

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS4717>





....a series of *Customer directed* technically oriented 90 minute webinars on various **storage** topics

2011 and 2012 Customer Webinars

- Feb 29th, 2012 – ProtecTier Technical Update
- Feb 14th, 2012 – TS7700 R2.1 Technical Update
- Feb 7th, 2012 – Performance Analysis Tools for the TS7700
- Jan 31st, 2012 – Tape 101
- Jan 24th, 2012 – DS8000 DS Storage Manager
- Nov 15th, 2011 – Introducing IBM XIV Gen3
- Nov 8th, 2011 – DS8000 Technical Update
- Oct 18th, 2011 – IBM z/OS and DS8000 Synergy
- Oct 4th, 2011 – Crossroads RVA

For further information and session notification please **Subscribe** to the ATS blog <https://www.ibm.com/developerworks/mydeveloperworks/blogs/accelerate/?lang=en>

Send Ideas for future topics to Tony Abete

ARABETE@us.ibm.com

Google “Accelerate with ATS”

- **Introduction to the TS7700 Performance Analysis Tools**
 - Why, What and Where
- Data Collection Requirements
 - 90 Day Trending
 - 24 Hour Evaluation
- Populating the Spreadsheets
- Getting a Clear Picture of the Issue
- Dissecting the Cache Throughput Graph
- Examples of Types of Issues
 - Grid network configuration
 - Performance Increment Ceiling
 - Deferred Copy Queue and Deferred Copy Throttle
 - Pre-Migration Queue
 - Physical Backend Drive Usage
 - Data Center Migration



- There is a need for tools to simplify and speed up TS7700 performance analysis
 - Pro-active analysis
 - Critical situations and major issues
 - Requests for performance analysis are numerous
 - The field, support structure, development and ATS needs tools to help analysis
 - Requests for assistance from other geographies
- Each customer's needs and issues tend to be different
 - Analyzing the data is more of an art than a science
 - First step is to look at the data. How can this be sped up?
 - The amount of data available for analysis is overwhelming
 - The TS7700 provides over 125 pieces of data in 15 minute intervals for the past 90 days
- Get the customer familiar with the data and involved in the analysis
 - Customer gains a better understanding of their TS7700 environment
 - Customer now has the tools to monitor their own environment
 - Customer now has the tools to analyze their own data
 - Customer becomes a partner with IBM when an issue arises
 - Communicating with a knowledgeable customer during a crisis is significantly more productive

- Pro-active analysis of customer's environment
 - Improves customer's awareness of their environment
 - 90 day trending to look for changes and extremes in the workload
 - Focus on busier times such as month-end or quarter-end
 - Identifies opportunities to upgrade or enhance the subsystem before an issue occurs
 - Provide tuning recommendations
- Customer has an immediate or recent issue
 - Critical situation or major concern
 - Focus on a particular day or days when the event occurred
 - Provide tuning recommendations and recommend upgrades as necessary
- Provide the customer with a formal report containing the details of the analysis including tuning and additional recommendations.

- Available to IBMers, BPs and Customers on Techdocs
- Utilizes VEHSTATS data that the customer collects
- 90 Day Performance Trending
 - Data collection requirements document
 - Spreadsheets for 1, 2, 3, 4, 5, and 6 cluster grids that produce important graphs used for analysis
- 24 Hour Analysis
 - Data collection requirements document
 - Spreadsheets for 2, 3, 4, 5, and 6 cluster grids that produce important graphs used for analysis
- 7 Day Analysis – Under Construction
 - Will use 7 days worth of hourly data
- Working with various customers to refine the tools
- Tools posted to Techdocs
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS4717>
 - Demonstration showing how to locate tools on Techdocs
- Explanation of VEHSTATS Columns
 - White paper for BVIR statistics records
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100829>
 - VEHSTATS Decoder ring white paper
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD105477>



United States [change]

Search

Home Solutions ▾ Services ▾ Products ▾ Support & downloads ▾ My IBM ▾

Welcome James Fisher [Not you?] [IBM Sign in]

Techdocs Library
Flashes
Presentations & tools
Technotes & tips
FAQs
White papers
Solution scenario profiles
Customer support plans
Sizings
Auxiliary Material
Search Techdocs
Techdocs feedback

Related links

- Redbook publications
- IBM Software Support Handbook

Techdocs - the Technical Sales Library



This site provides access to the Technical Sales Support organization's technical information databases. It gives you access to the most current installation, planning and technical support information available from IBM pre-sales support, and is constantly updated. You can browse or search these databases by date, document number, product, platform, keywords, etc.

New to Techdocs? Take a look at our [detailed introduction](#), which describes the document categories available (those listed on the navigation area on the left side of this page).

Rather than browse these categories, as a convenience you may enter a search of the full ~~Techdocs~~ database, or of any category you wish, here:

Search:

for:

Hits: Order by:

Include docs updated:

☐ Allow word variants
☐ "Fuzzy" search

[Help for Search](#)

Also available: our [Advanced search](#), where you can select documents based on various assigned document attributes.

New to Techdocs?

Is this your first visit to **Techdocs** (the Technical Sales Library)?

→ [Learn more](#)

Returning to Techdocs?

Looking for what's new in the **Techdocs Library**?

→ [Latest updates](#)

Need Technical Support?

Looking for support resources or other documents and tools?

→ [Support & downloads](#)

United States [change]

Home Solutions Services Products Support & downloads My IBM

Welcome James Fisher [Not you?] [IBM Sign in]

Techdocs Library >

Searching Presentations & Tools

The results of your search are below. You may also refine & resubmit the search.

Search: Presentations & Tools
☐ Allow word variants
☐ "Fuzzy" search

for: ts7700

Hits: 50 Order by: relevance

Include docs updated: any time
[Help for Search](#)

Documents 1 through 12 (of 12 found)

Document Number	Document Title	Last Update	Doc Author or Owner
PRS2702	IBM Tape Storage & System z Synergy - Trifold	05/29/2007	bob kern
PRS3527	TS7700 Management Interface for Release 1.5	03/12/2009	Jim Fisher
PRS3568	TS7700 Release 1.5 Technical Overview Presentation	03/05/2009	Jim Fisher
PRS3906	TS7700 Release 1.6 Technical Overview	01/12/2010	Jim Fisher
PRS3581	TS7700 Installation and Planning Checklists (Customer Tasks)	06/22/2009	Youngjong Ham
PRS4643	TS7700 Release 2.0 Technical Update	11/17/2011	Jim Fisher
PRS4751	TS7700 Introduction Presentation	09/12/2011	Carl Bauske
PRS4402	TS7700 Release 1.7 Technical Overview	12/09/2010	Jim Fisher
PRS4717	TS7700 90 Day and 24 Hour Performance Evaluation Spreadsheets and data Collection Requirements	11/22/2011	Jim Fisher
PRS2844	Techdocs Library > Presentations & Tools > TS7700 HCD and IOCP Definitions With Logical Path Calculator	09/01/2011	Carl Bauske
PRS2717	IBM Tape Storage & System z Synergy	06/05/2007	bob kern
PRS2614	DB2 for z/OS - What's tape got to do with it?	04/17/2007	John Iczkovits

Documents 1 through 12 (of 12 found)

Techdocs Library

- Flashes
- Presentations & tools**
- Technotes & tips
- FAQs
- White papers
- Solution scenario profiles
- Customer support plans
- Sizings
- Auxiliary Material

Search Techdocs

- Most recent Techdocs updates
- Advanced search

Techdocs feedback

Related links

- Redbook publications
- IBM Software Support Handbook

Techdocs Library > Presentations & Tools >

TS7700 90 Day and 24 Hour Performance Evaluation Spreadsheets and data Collection Requirements

Document Author: Jim Fisher

Document ID: **PRS4717**

Doc. Organization: Advanced Technical Skills

Document Revised: 11/22/2011

Product(s) covered: # 3957; # 3957-V06; Virtualization Engine; Virtualization Engine TS7700

Abstract: Provided below are spreadsheets and data collection requirements to assist in the evaluation of the TS7700 performance. Spreadsheets for both a 90 day evaluation and a 24 hour evaluation are provided.

One 90 day evaluation spreadsheet can be used for 1, 2, 3 or 4 cluster grids and the other for 5 and 6 cluster grids. There is an accompanying data collection guide for each. The first worksheet in each spreadsheet has instructions for populating the data into the spreadsheet.

There are separate 24 hour spreadsheets for 2, 3, 4, 5, and 6 cluster grids. The spreadsheets use the 15 minute interval data to produce charts for a 24 hour period. There is also a data collection guide.

These spreadsheets are intended for experienced TS7700 users. A detailed knowledge of the TS7700 is expected as well as familiarity with using spreadsheets.

These are the 90 Day Evaluation Files:



TS7700 90 Day Perf Eval Data Collection Requirements 5 and 6 Cluster Grids V10.doc



TS7700 90 Day Perf Eval Data Collection Requirements 1-4 Cluster Grids V14.doc



90 Day Trending Up to 6 clusters with ORDER6CL V10.xls



90 Day Trending 1, 2, 3, 4 cluster V10.xls

These are the 24 Hour Evaluation Files:



TS7700 7 Day Perf Eval Data Collection Requirements 1-4 Cluster Grids V14.doc



TS7700 7 Day Perf Eval Data Collection Requirements 5-6 Cluster Grids V14.doc



Two Cluster One Day 4 Pools V14.xls



Three Cluster One Day 4 Pools V14.xls



Four Cluster One Day 4 Pools V14.xls



Five Cluster One Day 4 Pools V14.xls



Six Cluster One Day 4 Pools V14.xls

- Introduction to the TS7700 Performance Analysis Tools
 - Why, What and Where
- **Data Collection Requirements**
 - **90 Day Trending**
 - **24 Hour Evaluation**
- Populating the Spreadsheets
- Getting a Clear Picture of the Issue
- Dissecting the Cache Throughput Graph
- Examples of Types of Issues
 - Grid network configuration
 - Performance Increment Ceiling
 - Deferred Copy Queue and Deferred Copy Throttle
 - Pre-Migration Queue
 - Physical Backend Drive Usage
 - Data Center Migration



Data Collection Requirements – 90 Day and 24 Hour

- Reminder to obtain the latest set of tape tools from the free Tape Tools ftp site
- Describe TS7700 grid configuration
 - Type of TS7700s
 - Locations and role of each cluster in the grid (production, DR, etc.)
- Request for a detailed description of the issue and/or purpose of the analysis
- Time frame of the issue
- How to set up JCL for collecting the right data
- Grid link status (LI REQ, dlib_name, STATUS, GRIDLINK)
- Current settings in the TS7700s including:
 - Cluster settings (LI REQ, dlib_name, SETTING)
 - Inhibit reclaim schedule
 - Pool properties including:
 - Maximum number of pre-migration drives
 - Reclaim threshold
 - Borrow/No borrow, Return/Keep settings
 - Copy Policies
 - Retain Copy Mode setting
 - Copy Policy Overrides
 - Fast ready categories including expire and expire hold settings
 - Total number of logical volumes
 - Number of performance increments (FC5268 and FC9268)

- Ninety days of VEHSTATS Daily Summary data.
 - Daily summary (HDSUM) (DAYHSMRY file)
 - In ORDERV12 or ORDER6CL JCL make sure the active pools have their respective "ORDER=" lines uncommented.
 - For example, if they use pools 4, 5, 6 and 7, make sure the ORDER= lines for these pools are uncommented.
 - Make sure the statements for four pools are not commented, even if less than four pools are in use.
 - No more than four pools should be uncommented.
 - In VEHSTATS JCL:
 - REPORT= HDSUM;
 - ORDER = ORDERV12, (for one to four cluster grids)
 - ORDER = ORDER6CL, (for five or six cluster grids)
 - UTCMINUS or UTCPLUS set appropriately
 - Start and End dates set appropriately
 - Send the DAYHSUMRY file
- If a period interest is identified, send 7 to 14 days of 15 minute data per the 7 day evaluation instructions.

- One to two weeks of VEHSTATS data from a period with the issue
 - 15 minute intervals (QTR) (HOURFLAT file)
 - In ORDERV12 or ORDER6CL JCL make sure the active pools have their respective "ORDER=" lines uncommented.
 - For example, if they use pools 4, 5, 6 and 7, make sure the ORDER= lines for these pools are uncommented.
 - Make sure the statements for four pools are not commented, even if less than four pools are in use.
 - No more than four pools should be uncommented.
 - In VEHSTATS JCL:
 - REPORT= QTR;
 - ORDER = ORDERV12, (for one to four cluster grids)
 - ORDER = ORDER6CL, (for five or six cluster grids)
 - UTCMINUS or UTCPLUS set appropriately
 - Start and End dates set appropriately
 - Send the HOURFLAT file

- VEHSTATS – Use VEHSTxx JCL as appropriate

- ORDERV12 or ORDER6CL

```
//VEHSTATS PROC TOOLHLQ=TOOLID,      HLQ FOR LIBRARIES
//      USERHLQ=USERID,              FOR THE INPUT BVIR FILE
//      SITE=SITENAME,                2ND LEVEL QUALIFIER
//      ORDER=ORDERV12,  DEFAULT ORDER STATEMENTS FOR GRAPHING PACKAGE
//*      ORDER=ORDERALL,  ALL AVAILABLE ORDER STATEMENTS
//*      ORDER=ORDER6CL,  FOR 5 OR 6 CLUSTER GRIDS
//      RECL=260,  260 IS WIDE ENOUGH FOR REPORT=GRID WITH 4 CLUSTERS
//*      380 FOR 5 CLUSTERS,  530 FOR 6 CLUSTERS
//*      710 FOR 7 CLUSTERS,  916 FOR 8 CLUSTERS
//*      260 IS WIDE ENOUGH FOR 22 CLUSTERS ON COMPARE REPORT
```

- Report

```
QUEAGEMINUTES;  REPORT DEF & RUN QUEUE AGE AS MINUTES, NOT SECONDS
REPORT= QTR HDSUM;
*      = QTR      REQUEST 15 MINUTE REPORTING AS GENERATED BY TS7740
*      = HRS      REQUEST HOURLY ROLL-UP REPORTING
*      = GRID     SUMMARIZES ALL CLUSTERS BY GRID
*      = SHOP     SUMMARIZES ALL CLUSTERS WITHIN SHOP
*      = COMPARE  REQUEST SIDE BY SIDE CLUSTER COMPARISON
*      = HDSUM    DAILY SUMMARY FLAT FILE - HORIZONTAL 1 DAY/LINE
*      = DXFR     FOR DAILY ON DEMAND TRANSFER REPORTING
UTCMINUS= 04;    ADJUST UTC TO LOCAL TIME WEST OF GREENWICH
*UTCPLUS= 02;    ADJUST UTC TO LOCAL TIME EAST OF GREENWICH
```

ORDER='SECTION: POOL 04	' ;	SECTION HEADING
ORDER='POOL 04 DEVTXX' ;		POOL DEVICE TYPE
ORDER='POOL 04 ACT VV' ;		ACTIVE VIRTUAL VOLUMES IN POOL
ORDER='POOL 04 ACT GB' ;		ACTIVE GIGABYTES IN POOL
ORDER='POOL 04 # PRIV' ;		PHYSICAL PRIVATE VOLUMES IN POOL
ORDER='POOL 04 # SRCH' ;		PHYSICAL SCRATCH VOLUMES IN POOL
ORDER='POOL 04 MB WRT' ;		MB(CHGD TO GB FOR DAY+) WRITTEN
ORDER='POOL 04 MB RD ' ;		MB(CHGD TO GB FOR DAY+) READ
*ORDER='SECTION: POOL 05	' ;	SECTION HEADING
*ORDER='POOL 05 DEVTXX' ;		POOL DEVICE TYPE
*ORDER='POOL 05 ACT VV' ;		ACTIVE VIRTUAL VOLUMES IN POOL
*ORDER='POOL 05 ACT GB' ;		ACTIVE GIGABYTES IN POOL
*ORDER='POOL 05 # PRIV' ;		PHYSICAL PRIVATE VOLUMES IN POOL
*ORDER='POOL 05 # SRCH' ;		PHYSICAL SCRATCH VOLUMES IN POOL
*ORDER='POOL 05 MB WRT' ;		MB(CHGD TO GB FOR DAY+) WRITTEN
*ORDER='POOL 05 MB RD ' ;		MB(CHGD TO GB FOR DAY+) READ

SETTING V1

ALERTS

COPYLOW = 3000 COPYHIGH = 5000

PDRVLOW = 12 PDRVCRIT = 8

PSCRLOW = 50 PSCRCRIT = 20

RESLOW = 1000 RESDHIGH = 2000

CACHE CONTROLS

COPYFSC = ENABLED

RECLPG0 = DISABLED

PMPRIOR = 1500 PMTHLVL = 2500

REMOVE = ENABLED

THROTTLE CONTROLS

COPYFT = ENABLED

ICOPYT = ENABLED

DCOPYT = 125

DCTAVGTD = 100

RECLAIM CONTROLS

RCLMMAX = 0

DEVALLOC CONTROLS

SCRATCH = DISABLED

PRIVATE = ENABLED

COPY CONTROLS

CPYCNT RUN = 20

CPYCNT DEF = 20

RUN COPY SENSE = NONE

COPY TIMEOUT = 180

GRIDLINK STATUS V1

CAPTURE TIMESTAMP: 2008-02-08 12:45:32

LINK VIEW

LINK NUM	CFG	NEG	READ MB/S	WRITE MB/S	TOTAL MB/S	ERR	LINK STATE 01234567
0	1000	1000	87.2	102.4	189.6	0	-AA
1	1000	1000	74.9	104.6	179.5	0	-AA
2	0	0	0.0	0.0	0.0	0	
3	0	0	0.0	0.0	0.0	0	

LINK PATH LATENCY VIEW

LIBRARY	LINK 0	LINK 1	LINK 2	LINK 3
LATENCY IN MSEC				
TS001B	6	7	0	0
TS001C	19	20	0	0

CLUSTER VIEW


DATA PACKETS SENT:	103948956
DATA PACKETS RETRANSMITTED:	496782
PERCENT RETRANSMITTED:	0.4778

LOCAL LINK IP ADDRESS

LINK 0 IP ADDR:	9.11.200.60
LINK 1 IP ADDR:	9.11.200.61
LINK 2 IP ADDR:	
LINK 3 IP ADDR:	

- Introduction to the TS7700 Performance Analysis Tools
 - Why, What and Where
- Data Collection Requirements
 - 90 Day Trending
 - 24 Hour Evaluation
- **Populating the Spreadsheets**
- Getting a Clear Picture of the Issue
- Dissecting the Cache Throughput Graph
- Examples of Types of Issues
 - Grid network configuration
 - Performance Increment Ceiling
 - Deferred Copy Queue and Deferred Copy Throttle
 - Pre-Migration Queue
 - Physical Backend Drive Usage
 - Data Center Migration





1. The spreadsheet is set to read-only. As a first step, open the spreadsheet and save it using a different name, typically a name that describes the grid or dates, etc.
2. Do not change the names of the worksheet tabs. These are used by the spreadsheet to perform cross-worksheet calculations.
3. If you are analyzing a two, three, four, five or six cluster grid, use the appropriate two, three, four, five or six cluster spreadsheets.
4. Populate the cluster worksheets (Cluster 0 Daily, Cluster 1 Daily, etc)
 - With 90 days of daily summary data
 - With 24 hours worth of 15 minute interval data (96 samples)
 - Next slide has more detail for populating the spreadsheets
5. Modify the Y axis scales so that charts with the same type of data have the same magnitude of data. This makes interpreting the data easier by avoiding a misinterpretation due to not noticing the differences in the Y axis.
6. The charts can easily be copied and pasted into a Word document or a PowerPoint presentation.
7. Populating Spreadsheets for Oddball Configurations
 1. Three cluster grid, CL0, CL2, CL3 – Use 4 cluster spreadsheet and clear data from CL2 worksheet

1. Import or paste the data into cell A2 or A5 on each worksheet. Include the header line from the source data. You can use this to verify the data copied into the spreadsheet populates the expected columns.
 1. If you paste the data, you will need to tell the spreadsheet how to parse the data in the spreadsheet when you paste the first cluster's data. Subsequent pasting of data does not need the extra steps.
 2. Paste the data into cell A2 or A5. The data will appear as a long, multi-line string.
 3. Select "Data" -> "Text to Columns".
 4. On the first screen select the "Delimited" radio button followed by "Next".
 5. Make sure the "Tab" and "Space" boxes are checked then select "Finish".
 6. You will probably want to make the spreadsheet readable by making sure each column is wide enough. Do this by selecting all rows then selecting "Format"->"Column"->"Autofit Selection".
2. Check the columns across the worksheet to make sure they line up with the column headers in row 1. Here are columns to pay attention to:
 1. Columns CE through CP contain grid copy activity for six clusters (only data for 5 clusters is used)
 2. Columns CZ through EA contain 4 pools worth of data.
 3. Columns EB through EE contain Copy and Pre-migrate data.
3. Next make sure the new data has completely overwritten the template data by examining column A which is the Grid ID. Delete any extra rows of data.
4. Repeat for all clusters in the grid.

Pasting Data Into The Spreadsheet



The top screenshot shows a Microsoft Excel spreadsheet titled "Four Cluster One Day 4 Pools V14.xls [Read-Only]". The spreadsheet has columns A through Z. A red circle highlights the data in columns A through Z, starting from row 1. The data includes various metrics such as "Grid", "CLID", "Day", "Date", "Avg_Virt_Drvs", "Phy_Stg_Mnts", "Phy_Mig_Mnts", "Phy_Rcm_Mnts", "Tot_Phys_Mnts", "Max_Phys_Mtime", "Avg_Phys_Mtime", "Max_Phys_Mntd", "Avg_Phys_Mntd", "GiB_Read", "C", "PG1_GiB_to_CPY", "PG0_4Hr_VV_Mig", "PG1_4Hr_VV_Mig", "PG0_48H_VV_Mig", "PG1_48H_VV_Mig", "PG0_35D_VV_Mig", "PG1_35D_VV_Mig", "PG0_4Hr_Av_Min", "PG1_4Hr_Av_Min", "BlkSz_LE_2K", "BlkSz_LE_4K", "BlkSz_LE_8K", "BlkSz_LE_16K", "BlkSz_LE_32K", "BlkSz_LE_64K", "BlkSz_GT_64K", "Phys_Vols_Imp", "Phys_Vols_Exp", "Virt_Vols_Imp", "Virt_Vols_Exp", "EOI_VV_To_Recv", "EOI_MiBTo_Recv", "Max_Av_Def_Min", "Max_Av_Run_Min", "Max_VV_To_Recv", "Max_MiBTo_Recv", "MiB01_By_Copy", "Avg_01_MiB_s", "Max_01_MiB_s", "MiBRecv", "MiB03_By_Copy", "Avg_03_MiB_s", "Max_03_MiB_s", "CL0_Rmt_Wr_MiB", "CSPMED6_3592JB", "CSPMED2_3592JA", "CSPMED4_3592JJ", "PRIMED6_3592JB", "PRIMED2_3592JA", "PRIMED4_3592JJ", "POOL_01_DEVTXX", "POOL_01_ACT_VV", "POOL_01_ACTGiB", "POOL_01_Privat", "POOL_01_Scrch", "POOL_01_MiBWRT", "POOL_01_MiBRD", "POOL_02_DEVTXX", "POOL_02_ACT_VV", "POOL_02_ACTGiB", "POOL_02_Privat", "POOL_02_Scrch", "POOL_02_MiBWRT", "POOL_02_MiBRD", "POOL_03_DEVTXX", "POOL_03_ACT_VV", "POOL_03_ACTGiB", "POOL_03_Privat", "POOL_03_Scrch", "POOL_03_MiBWRT", "POOL_03_MiBRD", "POOL_04_DEVTXX", "POOL_04_ACT_VV", "POOL_04_ACTGiB", "POOL_04_Privat", "POOL_04_Scrch", "POOL_04_MiBWRT", "POOL_04_MiBRD", "EOI_GiB_to_Mig", "Max_GiB_to_Mig", "EOI_GiB_to_Cpy", "Max_GiB_to_Cpy".

The bottom screenshot shows the same spreadsheet with the "Data" menu open. The "Text to Columns..." option is highlighted with a red circle. The menu also includes options like "Sort...", "Filter", "Form...", "Subtotals...", "Validation...", "Table...", "Consolidate...", "Group and Outline", "PivotTable and PivotChart Report...", "Import External Data", "List", and "Refresh Data".

Pasting Data Into The Spreadsheet (continued)



Convert Text to Columns Wizard - Step 1 of 3

The Text Wizard has determined that your data is Fixed Width.
If this is correct, choose Next, or choose the data type that best describes your data.

Original data type
Choose the file type that best describes your data:

☒ Delimited - Characters such as commas or tabs separate each field.
☐ Fixed width - Fields are aligned in columns with spaces between each field.

Preview of selected data:

	Grid	CLIDMSER	Day	Date	Time	Code_Level	TVC_
2	AE1C0	CL0H1079	Sun	20NOV2011	00:15:00	17.00.0143	
3	AE1C0	CL0H1079	Sun	20NOV2011	00:30:00	17.00.0143	
4	AE1C0	CL0H1079	Sun	20NOV2011	00:45:00	17.00.0143	
5	AE1C0	CL0H1079	Sun	20NOV2011	01:00:00	17.00.0143	

Buttons: Cancel, < Back, Next >

Convert Text to Columns Wizard - Step 2 of 3

This screen lets you set the delimiters your data contains. You can see how your text is affected in the preview below.

Delimiters

☒ Tab ☐ Semicolon ☐ Comma
☒ Space ☐ Other:

☒ Treat consecutive delimiters as one

Text qualifier:

Data preview

Grid	CLIDMSER	Day	Date	Time	Code_Level	TVC_Size	Active
AE1C0	CL0H1079	Sun	20NOV2011	00:15:00	17.00.0143	9000	58491
AE1C0	CL0H1079	Sun	20NOV2011	00:30:00	17.00.0143	9000	58491
AE1C0	CL0H1079	Sun	20NOV2011	00:45:00	17.00.0143	9000	58491
AE1C0	CL0H1079	Sun	20NOV2011	01:00:00	17.00.0143	9000	58491


Buttons: Cancel, < Back, Next >, Finish

Microsoft Excel


Do you want to replace the contents of the destination cells?

Buttons: OK, Cancel



Pasting Data Into The Spreadsheet (continued)



Select first data cell in column F
(Code Level)



- Three Cluster Grid – HA/DR Configuration
- Populate 90 day data into canned spreadsheet
- Save the spreadsheet to its own filename
- Looking at the automatically generated charts
- Populate 1 to 2 weeks worth of data into blank spreadsheet
- Look for time period of interest
- Populate 24 hours worth of data into canned spreadsheet
- Save the spreadsheet to its own filename
- Looking at the automatically generated charts
- Do I need to create other charts specific to my issue?
- Changing graph titles
- Copying graphs to a Word document




- Introduction to the TS7700 Performance Analysis Tools
 - Why, What and Where
- Data Collection Requirements
 - 90 Day Trending
 - 24 Hour Evaluation
- Populating the Spreadsheets
- **Getting a Clear Picture of the Issue**
- Dissecting the Cache Throughput Graph
- Examples of Types of Issues
 - Grid network configuration
 - Performance Increment Ceiling
 - Deferred Copy Queue and Deferred Copy Throttle
 - Pre-Migration Queue
 - Physical Backend Drive Usage
 - Data Center Migration




- Be sure to get a clear definition of the issue.
- Time frame so you know where to look.
- VEHSTATS time vs wall clock time vs time issue reported vs GMT (25 or 6 to 4).
- Symptom – This will help determine what you will look at.
 - Host write throttling
 - Copy queue
 - Pre-migration
 - Job slowdowns
- Ninety Day Evaluation
 - Is there a specific issue being addressed?
 - Is the evaluation for customer awareness?
- 24 Hour Evaluation
 - What is the specific issue being addressed?
 - What time did the issue occur?
 - Is UTCMINUS/UTCPLUS set correctly in the VEHSTATS JCL?
 - Was there anything out of the ordinary occurring during the period with the issue?
 - Extra workload due to month-end, quarter-end, year-end?
 - Special, one time, processing?
- Look at the graphs for something that stands out.


Example 90 Day Charts




Example 90 Day Charts



24 Hour Analysis Examples




24 Hour Analysis Examples



- Introduction to the TS7700 Performance Analysis Tools
 - Why, What and Where
- Data Collection Requirements
 - 90 Day Trending
 - 24 Hour Evaluation
- Populating the Spreadsheets
- Getting a Clear Picture of the Issue
- **Dissecting the Cache Throughput Graph**
- Examples of Types of Issues
 - Grid network configuration
 - Performance Increment Ceiling
 - Deferred Copy Queue and Deferred Copy Throttle
 - Pre-Migration Queue
 - Physical Backend Drive Usage
 - Data Center Migration





Dissecting the Cache Throughput Graph



- Red line is uncompressed host IO rate
- Stacked bars show data moving through the disk cache
 - Lime green is compressed host IO rate
 - Cyan is copy out rate to the grid
 - Light blue is copy in rate from the grid
 - Yellow is pre-migration rate
 - Dark blue is recall rate from physical tape
 - Orange is remote reads by other clusters from this cluster's cache
 - Burnt orange is remote writes by other clusters into this cluster's cache
- Copy in/out rate can decrease when Deferred Copy Throttle (DCT) is being applied
- Pre-migration rate can decrease during periods of higher host IO


- Introduction to the TS7700 Performance Analysis Tools
 - Why, What and Where
- Data Collection Requirements
 - 90 Day Trending
 - 24 Hour Evaluation
- Populating the Spreadsheets
- Getting a Clear Picture of the Issue
- Dissecting the Cache Throughput Graph
- **Examples of Types of Issues**
 - **Grid network configuration**
 - **Performance Increment Ceiling**
 - **Deferred Copy Queue and Deferred Copy Throttle**
 - **Pre-Migration Queue**
 - **Physical Backend Drive Usage**
 - **Data Center Migration**





- Cluster 2 pulls logical volumes from clusters 0 and 1
- Cluster 2 has up to 20 active copy jobs
- Cluster 2 processes the copies in the order received (mostly)
- Cluster 2 may be pulling more copies from one cluster at any given time
 - Sometimes 10 copies from each
 - Other times up to 20 from one cluster and none from the other
- 148MB/s maximum copy rate when 10 copy jobs from cluster 0 and 10 from cluster 1
- 74MB/s maximum copy rate when 20 copy jobs from a single cluster
- Also, single grid path is a SPOF


Copy Rate Before and After Network Configuration Fixed



- Place one week's worth of data into a blank spreadsheet.
- One worksheet per cluster.
- Copy the copy-rate-to-CL2 columns from CL0 and CL1 worksheets to an additional worksheet and add them together.
 - Avg_02_MiB_s
 - Avg_12_MiB_s
- Copy the Time and CL2 MB-to-receive column to the additional worksheet.
- Plot Cluster 2 MB-to-receive and Copy-in-rate.

- FC5268 provides 100 MB/s Performance Increment
 - Maximum of 6 for 3957-V06
 - Maximum of 5 for 3957-VEA (One increment, FC9268, included with base model)
 - Maximum of 10 for 3957-V07
 - Maximum of 9 for 3957-VEB (One increment, FC9268, included with base model)
- Is the TS7700 actually limiting the Host IO due to performance increments?
 - Examine the uncompressed host IO rate
 - Sometimes it is clear that a plateau is being reached
 - Sometimes it doesn't appear to be limited, but it actually is.
 - VEHSTATS reports a 15 minute average of the host IO
 - The TS7700 evaluates whether to limit the IO rate every second
 - If the maximum host IO is within 50-75 MB/s of the limit, chances are the limit is being reached


- Each second the TS7700 zeros a counter.
- The counter is incremented with each byte of data written or read by the host. This is uncompressed data.
- If the Performance Increment Ceiling is reached within the second, the TS7700 will CCR any host writes until the counter is zeroed again.
- For example, with a 300MB/s ceiling, if 300 MB of data is written and read in 800ms, the TS7700 will CCR host write requests for the next 200ms.
- `rasutil -vma` used to collect this data



200MB/s Performance Increments Example





- Example of a clearly visible performance increment plateau
- Customer chose to not increase the number of performance increments since they were meeting their SLAs




- Next example is from a not-so-clear performance increment plateau
 - 300MB/s Performance Increments
 - Maximum IO rate 250-275MB/s

How Performance Increments Limit Performance – Before and After





- Customers have different expectations concerning the Deferred Copy Queue
 - Some need copies to occur within just a few hours
 - Others need copies to occur within 12-14 hours
 - Essential that the queue drop to nearly zero at least once per day to prevent run away copy queue
- Applying a Deferred Copy Throttle (DCT) is a trade-off of host IO versus copy time
 - With Power7 processor the need to limit deferred copies in deference to host IO is less than with the Power5 processor
- DCT affects copy rate in a non-linear fashion
- “Knee-of-the curve” is ~30ms delay
 - Below this value greatly effects host IO
 - Illustrated in the next set of charts

Effect of DCT – Theoretical versus Reality




Effect of DCT – One Way Latency of 10-15ms



Deferred Copy Queue Backlog




- Workload was increased
- Copy queue couldn't reach zero each day
- Job One was to get queue caught up
- Adjusted DCT to 0 to catch up
- Customer was still able to complete their workload with DCT set to zero
- In the process uncovered grid network that was not configured optimally (covered earlier)




DCT Set too High, Deferred Copy Queue Exceeds Size of Disk Cache



PG0 and PG1 in Cache at Midnight - Dec 2009 - Feb 2010





Deferred Copy Queue Depth at Midnight - Dec 2009 - Feb 2010



- PMPRIOR defaults to 1600 GB
 - Exceeding this threshold causes number of pre-migration tasks to be increased
- PMTHLVL defaults to 2000 GB
 - Exceeding this threshold causes host write and copy throttle to be applied
- Adjust to defer pre-migrations to a non-peak time
- Be sure there is time to catch up

Constant Pre-migration Limiting Throughput




- Back-end drives are used for Recalls, Pre-migration and Reclaim
- Recall mounts have priority
 - A pre-migration task will be stopped to allow a recall to occur after a 2-30GB chunk of data if pre-migrated
 - A Reclaim task will free up its two drives to allow a recall after the source tape is emptied
- Look for opportunities to move reclaim out of the high host IO period
 - Limit the number of reclaim tasks
 - Use the Inhibit Reclaim Schedule
 - Drive contention can often be reduced by effective use of reclaim tuning knobs
- The following charts show:
 - A cluster with perhaps not enough physical drives
 - A cluster with the right number of drives
 - A before and after example of moving reclaim to a better time


Not Enough Physical Drives vs. Enough Physical Drives



Cluster 0 Max and Avg Physical Drives Mounted - Feb 18, 2011




Cluster 0 Max and Avg Physical Drives Mounted




Physical Drive Usage Before and After Reclaim Moved to Weekends




Cluster 2 Max and Avg Physical Drives Mounted - October 11, 2011



Cluster 2 Max and Avg Physical Drives Mounted - 10/25/2011




- The cluster in question had a significant requirement for recall mounts



Reclaim Now on Weekends




- Weekends are a less busy time for the TS7700 Grid
- Customer is monitoring the physical scratch volume quantity to make sure reclaim has enough time to keep up



Family and Copy Tasks




- Each cluster defaults to 20 copy tasks
- DCT set to 5ms to insure timely copies.
- Two receiving clusters were added as part of a workload transfer
- Host IO suffered due to increased CPU demand



Family and Copy Tasks





- Two new clusters placed into a family
- Production cluster now only needs to supply 1 copy to the family
- However, there are still 60 copy tasks
- Next step was to reduce the number of copy tasks in each cluster in the family





- Put two new clusters into a family.
- Production cluster only needs to source two copies.

Without families Cluster 0 sourced each logical volume 3 times, once per cluster.





With families Cluster 0 only had to source each logical volume 2 times, once for cluster 1, and once for either cluster 2 or 3.


Copy Tasks – 20 per Cluster versus 10 per Cluster



Copy Tasks – 20 per Cluster versus 10 per Cluster



Copy Tasks – 20 per Cluster versus 10 per Cluster




In this environment reducing the copy tasks affected inter-family copies. Within a family first copies take precedence over second copies.


Copy Tasks set to Five and Eight




Grid 55555 Cluster 0 Cache Throughput - 1/19/2012



Grid 55555 Copy Rate to Clusters 2 and 3 - 1/19/2012



Questions?



Disclaimers and Trademarks



- Copyright© 2012 by International Business Machines Corporation.
- No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.
- The performance data contained herein were obtained in a controlled, isolated environment. Results obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. These values do not constitute a guarantee of performance. The use of this information or the implementation of any of the techniques discussed herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into their operating environment. Customers attempting to adapt these techniques to their own environments do so at their own risk.
- Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any on-IBM product, program or service.
- THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT.
- IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g. IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.
- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of 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.
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following terms are trademarks or registered trademarks of the IBM Corporation in either the United States, other countries or both.

- IBM, System Storage, TotalStorage, System i, System p, System x, System z, Virtualization Engine
 - z/OS, z/VM, VM/ESA, OS/390, AIX, DFSMS/MVS, OS/2, OS/400, i5, FICON, ESCON, Tivoli, VSE/ESA, TPF
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both
 - Other company, product, and service names mentioned may be trademarks or registered trademarks of their respective companies
 - LTO, Ultrium and Linear Tape Open are trademarks of HP, IBM and Quantum in the United States, other countries, or both
 - Linux is a registered trademark of Linus Torvalds in the United States, other countries or both