


System Storage

Tape Products Update



Lee Jesionowski
Lead Architect – Tape Automation Products



Agenda


- Tape trends
- Update on IBM tape drives
 - LTO 8
 - TS1155
 - Enterprise commonality and differentiators
- Update on IBM tape libraries
 - Integrated cooling



Areal Density Scaling

2015: IBM-FujiFilm demonstration of 123 Gb/in² on BaFe tape

2017: IBM-Sony demonstration of 201 Gb/in² on Sputtered Tape



IBM tape storage for hyperscale computing

The more things change, the more they stay the same...



TS1155 – Ethernet interface and 15 TB tape cartridge

How to store a zettabyte on a budget

- Aaron Ogus, Microsoft Azure
- Global IT Executive Summit - October, 2015
- <https://tapepower.fujifilmrmd.com/LA2015/video/id/presentation.5>



How Google Backs up the Internet

- Raymond Blum, Google Site Reliability
- NYC Tech Talk Series - October, 2013
- <https://www.youtube.com/watch?v=eNliOm9NtCM>

Lessons Learned Backing Up Google

- Ensuring durability and integrity of user data is job one
A lapse in availability can be ridden out, but data loss can be hard to recover from, if even possible
- Redundancy does not bring recoverability
Corruption and deletes can replicate quite nicely
- Distributed processing imposes data consolidation
You need to collect shards into one cohesive world view at some point
- The backup process has to scale with data volume
- If you haven't restored, you haven't backed up
- The Payoff: A case study



IBM tape storage for hyperscale computing

The more things change, the more they stay the same...



TS1155 – Ethernet interface and 15 TB tape cartridge

How to store a zettabyte on a budget

- Aaron Ogus, Microsoft Azure
- Global IT Executive Summit - October, 2015
- <https://tapepower.fujifilmrmd.com/LA2015/video/id/presentation.5>

[HPSS user testimonial here]

- [Your Name, Your Company]
- [Your testimonial URL]



How Google Backs up the Internet


- Raymond Blum, Google Site Reliability
- NYC Tech Talk Series - October, 2013
- <https://www.youtube.com/watch?v=eNliOm9NtCM>

Lessons Learned Backing Up Google

- Ensuring durability and integrity of user data is job one
A lapse in availability can be ridden out, but data loss can be hard to recover from, if even possible
- Redundancy does not bring recoverability
Corruption and deletes can replicate quite nicely
- Distributed processing imposes data consolidation
You need to collect shards into one cohesive world view at some point
- The backup process has to scale with data volume
- If you haven't restored, you haven't backed up
- The Payoff: A case study



IBM Tape Drive History and Roadmaps

LTO Generations	LTO-6	LTO-7	LTO-8	LTO-9
Max Format Capacity (Native) 	2.5 TB (L6)	6 TB (L7)	12 TB (L8)	Up to 25 TB (L9)
Other Format Capacities (Native)	1.5 TB (L5) (800 GB L4 R/O)	2.5 TB (L6) (1.5 TB L5 R/O)	6 TB (L7)	Up to 12 TB (L8) (6 TB L7 R/O)
Native Data Rate	160 MB/s	300 MB/s	360 MB/s	Up to 450 MB/s

2012


2015

2017

2011

2014

2017

TS1100 Generations	TS1140	TS1150	TS1155	Gen-6	Gen-7
Max Format Capacity (Native) 	4 TB (JC) 1.6 TB (JB)	10 TB (JD) 7 TB (JC)	15 TB (JD) 7TB (JC)	Up to 20 TB (JE) 15 TB (JD) 10 TB (JC)	Up to 50 TB (JF) Up to 30 TB (JE) 15 TB (JD)
Other Format Capacities (Native)	1 TB (JB) 700 GB (JB) (All JA R/O)	4 TB (JC)	10 TB (JD) 4 TB (JC, R/O)	10 TB (JD) 7 TB (JC) 4 TB (JC, R/O)	10 TB (JD)
Native Data Rate	250 MB/s	360 MB/s	360 MB/s	Up to 420 MB/s	Up to 1000 MB/s
Attachment	FC-8	FC-8	FC-8, 10 GigE (RoCEv2)	FC-16, 25 GigE (RoCEv2)	TBD

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

IBM LTO-8 Announcement on October 10

▪ New generation of Tape Drive and Libraries, LTO-8

- *New!* LTO-8 for TS4500 Tape Library
- *New!* LTO-8 for TS4300 Tape Library
- *New!* LTO-8 for TS2900 Tape Autoloader
- *New!* IBM® TS2280 Tape Drive
- *New!* LTO-8 for TS3100/3200/3310/3500 Tape Libraries
- *New!* Spectrum Archive EE/LE/SDE support for LTO-8 tape drive

GA:

TS3100 / 3200: 11/10/17

TS3500 / TS4500 / TS2280: 11/17/17

TS3310: 12/08/17

7226: 12/08/17

TS1155 Tape Drive integration into TS3500

Announce date: October 10

GA: Nov, 17



Some LTO-8 Topics of Discussion

- Time between LTO drive generations is shorter?
 - 2 years for LTO-8 vs history of 2 ½ to 2 ¾ years
- LTO-8 capacity increase vs data rate increase?
 - Native capacity of 12 TB vs 6 TB for LTO-7
 - Sustained data rate of 360 MB/s vs 300 MB/s for LTO-7
- Lack of LTO-6 media backward read compatibility?
 - TMR (Tunnel Magneto-resistive) LTO-8 head
 - MP (Metal Particle) LTO-6 media



Some LTO-8 Topics of Discussion (continued)



- What is the new “LTO 8 Type M” initialization option?

“New LTO generation 7 cartridges initialized as LTO-8 Type M media will be able to store up to 22.5TB* of data.”

Source: <https://www.lto.org/2017/10/lto-program-outlines-generation-8-specifications-extends-technology-roadmap-12th-generation/>

IBM plans to refer to this as “M8 media”

New barcode label – check with your tape library vendor for support

Media / Density Type	Barcode Label	Cartridge Packaging & Silkscreen	Native Capacity	Drive Compatibility
L8	xxxxxxL8	LTO Ultrium 8	12 TB	LTO-8
M8	xxxxxxM8	LTO Ultrium 7	9 TB	LTO-8
L7	xxxxxxL7	LTO Ultrium 7	6 TB	LTO-7, LTO-8







TS1155 Tape Drive – summary

- **Enhanced Capacity**
 - Up-formats existing JD media types
 - 15 TB JD/JZ media, 3 TB JL short media
 - Supports downward compatibility with JD and JC formats
 - No change in native data rate 360 MB/s
- **Supporting both fibre and new Ethernet interface capabilities**
 - Dual 10Gb optical attachment on new 55E model
 - RoCE v2 protocols support
 - Dual FC-8 interfaces supported on 55F model
- **Introducing longer life TMR read/write head**
- **Other features of TS1150 remain essentially unchanged**



Underlying Common Technology Description

- Integrated 32-channel head design
 - Advanced TMR sensor for robust signal output with protective overcoat
 - Three module W-R-W 3-bump head performs read-verification in both directions
- Advanced dd-NPML data detection channel
 - Trellis dd-NPML detection with integrated timing loop
 - Auto-adaptive FIR equalization and AGC
- Simple compact Tape Path
 - Low head wrap-angle enabled by flat-head contour – low friction
 - Flangeless rollers = no edge forces
- Advanced high performance track following actuator
 - Uses the servo pattern to precisely position for recording tape
 - Maintains lock to rejects lateral tape motion and tape runout errors
- Highly Integrated, low power electronics
 - Advanced IBM foundry ASIC technology
 - Low number of interconnects
- Common technology, but unique components for enterprise
 - Hardened mechanism, loader and higher speed motors, enterprise cartridge
 - Custom circuit card with larger buffer, faster data rate support, new interface processor
 - Microcode featurization for enterprise RAS, performance features



Enterprise Tape Drive Differentiators

- Capacity and Investment Protection
- Performance
- Access
- Interfaces / Attachments
- Reliability
- Crypto



Capacity and Investment Protection

Capacity and Media Up-format

- Enterprise has the highest single cartridge capacity and roadmap maintains the capacity advantage
- Enterprise supports media up-formatting, allowing previous generation media to be re-used at higher capacities and data rates on future drive models.
- Enterprise supports field MES drive model upgrades

- Advantages
 - Improved density / reduced library frames reducing overall system cost
 - Higher capacity and performance over time on existing media reducing media/storage costs
 - Maintaining footprint with future storage growth
 - Drive models may be upgraded at reduced cost and upgrades may be expensed




Performance

Native Data Rate

- Enterprise has the highest native data rate and roadmap increases differential over time. In addition, enterprise supports downward R/W of older formats at higher data rates improving migration
 - Jag5 360 MB/s, Jag6 420MB/s, Jag7 targeting 1000 MB/s (64 channel)
 - LTO7 300 MB/s, LTO8 360 MB/s, LTO9 400-450 MB/s (current outlook)
- Advantages
 - Fewer drives for given aggregate throughput, faster large file recalls
 - Significant performance benefit for higher single-threaded data rate applications
 - Some applications cannot multiplex data across many drives/tapes (seismic, etc)
 - Improved migration performance

FastSync write accelerator

- Enterprise supports the FastSync feature which is able to greatly improve performance for applications which write file marks (tape marks) with relatively small files (300MB or less).
- Advantages
 - Job times for some applications improve by a factor of 20x or greater
 - Many applications have at least some workload that writes relatively small files (under 200MB) such as audit logging where FastSync is highly effective
 - Reduces back hitches on media improving media and drive life



Access features

- **Search Speed**
 - Enterprise features highest search and rewind velocity about 30% faster than LTO
- **High Resolution Tape Directory (HRTD)**
 - Greater precision is maintained in the directory of the physical location of all blocks on tape
 - 1/64th of a tape wrap resolution per block versus ½ wrap for LTO (32x greater precision)
 - Physical motion in locating to a given file is reduced and more consistent
- **Recommended Access Ordering support**
 - Optimizes recall order for multi-file recalls
 - Improved performance for multi-file recall from a single volume

- **Advantages**
 - LTO average access to a file is 30% slower than enterprise and rewinds are 42% slower
 - Fewer enterprise drives are needed for a given recall rate
 - Applications that read or write a modest amount of data (recalls or appends) may be greatly improved due to locate/rewind operations inherent in this workload
 - If multiple files are often recalled from a single cartridge, this workload performance can be greatly improved with Recommended Access Ordering



Interfaces/ Attachment

▪ Ethernet Host Interfaces

- Enterprise drives have added support for Ethernet interfaces
- Supports dual port 10Ge physical attachment with optical SFP moving to 25Ge as available

▪ Enhanced Fibre Channel attachment

- Enterprise drives adding 16G FC support on Jag6 and 32G FC support on Jag7
- Supports higher host transfer rates

▪ Mainframe attachment

- Enterprise drives support mainframe attachment
- May be attached to z/OS attached TS7700
- Enables common technology in Mainframe and open system environments

Advantages

- Supports SAN elimination with move to converged Ethernet – cost and admin savings
- Supports higher Host and HBA transfer rates on faster fibre channel attachments – performance improvements and faster job execution




Figure 1 – iSCSI and iSER Protocol Layers

Reliability Features

- **Longer life loader**
 - supporting 300K cycles loader life, 3x higher than LTO
 - Consistent with faster access workload which results in more load/unloads per drive

- **Built-in, automatic in-drive Media and Drive quality management**
 - Pro-active alerts for usage and health issues (SARS)
 - Monitors the last 100 mounts for both the drive and the media
 - Automatically triggers cleaning if degradation is detected
 - Generates proactive alerts if signs of degradation persist

- **Media near end of life alert generation**
 - Application will be notified if media is near end of life condition for loads or meters processed
 - Enables pro-active media replacement


- **Enterprise media differentiation**
 - Constructed for greater physical durability - survives one meter drops without data loss
 - All media is end point tested in drives for qualify in the factory before shipment
 - Media is designed to support re-use at higher capacity and performance
 - Media supports a longer usage life (26M motion meters)

- **Advantages**
 - In general, Enterprise drives are designed to provide maximum data reliability and operational reliability reducing possibility of data loss or drive failure



3592 Robust Cartridge Design

- **Robust cartridge design**
 - Thicker plastics (vs. LTO)
 - Ribs to hold a reel
 - Five screws to tighten shells
 - Spec'd to withstand 1-m drops from all 6 axes without dataloss (not recommended!)
- **Dust-proof curved door design**
 - Effective dust-proof design for higher recording density
 - Passed an open/close test more than 50K cycles



Enterprise Tape Drive Differentiators


- Capacity and Media Up-format
- Performance
- Access
- Interfaces / Attachments
- Reliability
- Crypto



TS4500 Tape Library



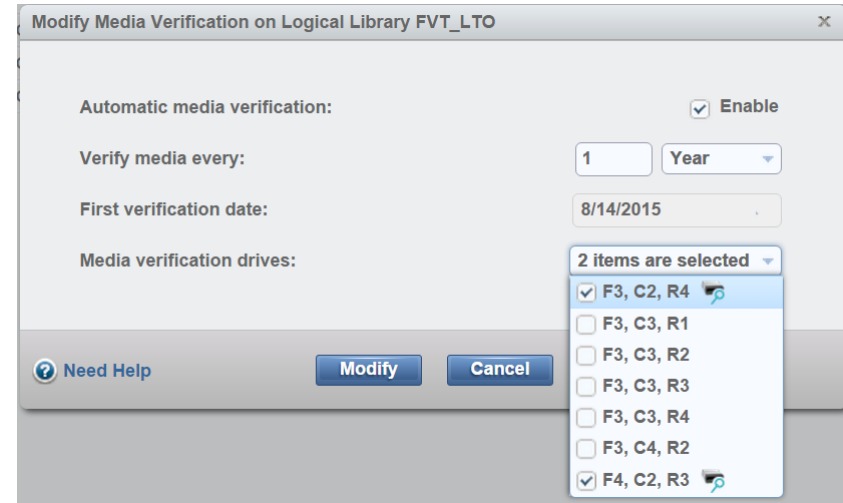
- Scales to over 278 PB native
- Supports up to 128 tape drives
- Next-generation storage density
 - 8.7 PB native in a single frame tape library
 - Scales at up to 1.5 PB per square foot
- Simplified management
 - Magazine I/O
 - Integrated management
 - Improved ease-of-use
 - Extensible platform



TS4500 Automatic Media Verification

Pro-actively verify the ability to read data stored on tape from any open systems application

- Simple policy designates policy period, start date, and media verification drives for the logical library
- Results are viewable in the Cartridges and Events pages of the TS4500 GUI
- Notification of failure sent via SNMP, email, or syslog
- Works with any open systems application*



Anacg4a > Cartridges > **Cartridges by Logical Library** IBM TS4500 Tape Library

VOLSER	State	Location	Last verification	Verification results
IMN427L5	Slot	F4,C6,R3,T0	8/14/15, 2:56 PM	✓ Passed
IM1392L6	Slot	F4,C8,R8,T0	8/14/15, 2:59 PM	✓ Passed
IM1376L6	Slot	F3,C1,R12,T2	8/14/15, 3:00 PM	✓ Passed
IM1374L6	Slot	F2,C6,R7,T0	8/14/15, 3:02 PM	✓ Passed
IM1375L6	Slot	F2,C10,R28,T0	8/14/15, 3:05 PM	✓ Passed
IM1373L6	Slot	F2,C4,R8,T0	8/14/15, 3:07 PM	✓ Passed
IM1349L6	Slot	F4,C8,R4,T0	8/14/15, 3:09 PM	✓ Passed



Why is there a need for tape library integrated cooling?

Table 1. Environment for operating, storing, and shipping the LTO tape cartridges

Environmental Factor	Environmental Specifications			
	Operating	Operational Storage ¹	Archival Storage ²	Shipping
Temperature	10 to 45°C (50 to 113°F)	16 to 32°C (61 to 90°F)	16 to 25°C (61 to 77°F)	-23 to 49°C (-9 to 120°F)
Relative humidity (noncondensing)	10 to 80%	20 to 80%	20 to 50%	5 to 80%
Maximum wet bulb temperature	26°C (79°F)	26°C (79°F)	26°C (79°F)	26°C (79°F)
Magnetic field	Stray magnetic field at any point on tape not to exceed 50 oersteds (4000 ampere/meter).			

Notes:

1. Operational storage equals less than 6 months.
2. Archival storage equals greater than 6 months.

6. MICROSOFT DATA CENTER, DUBLIN, 550,000 SF



An aerial view of Microsoft's huge Dublin data center.

Microsoft's new **Dublin Data Center** represents a milestone in data center design. The new facility will power much of Microsoft's global cloud computing operation, while using far less energy and water than typically consumed in other data centers of this scale. The Dublin center operates at a Power Usage Effectiveness (PUE) of 1.25, Microsoft says, compared to averages of about 2.0 for the industry. The design innovations driving its efficiency include a "free cooling" system that uses outside air to cool the data center, and

warmer operating temperatures inside the server space. Microsoft says it can run its server rooms at temperatures of up to 95 degrees F (35 degrees Celsius), much warmer than most data centers, which typically range between 68 and 72 degrees. For more, see [inside](#)

[Data Center.](#)

« Previous Article
Fjord-Cooled Data Center Gets Anchor Tenant

Next Article »
Cloud News: Microsoft, HP, Savvis, Nimsoft

Too Hot for Humans, But Google Servers Keep Humming

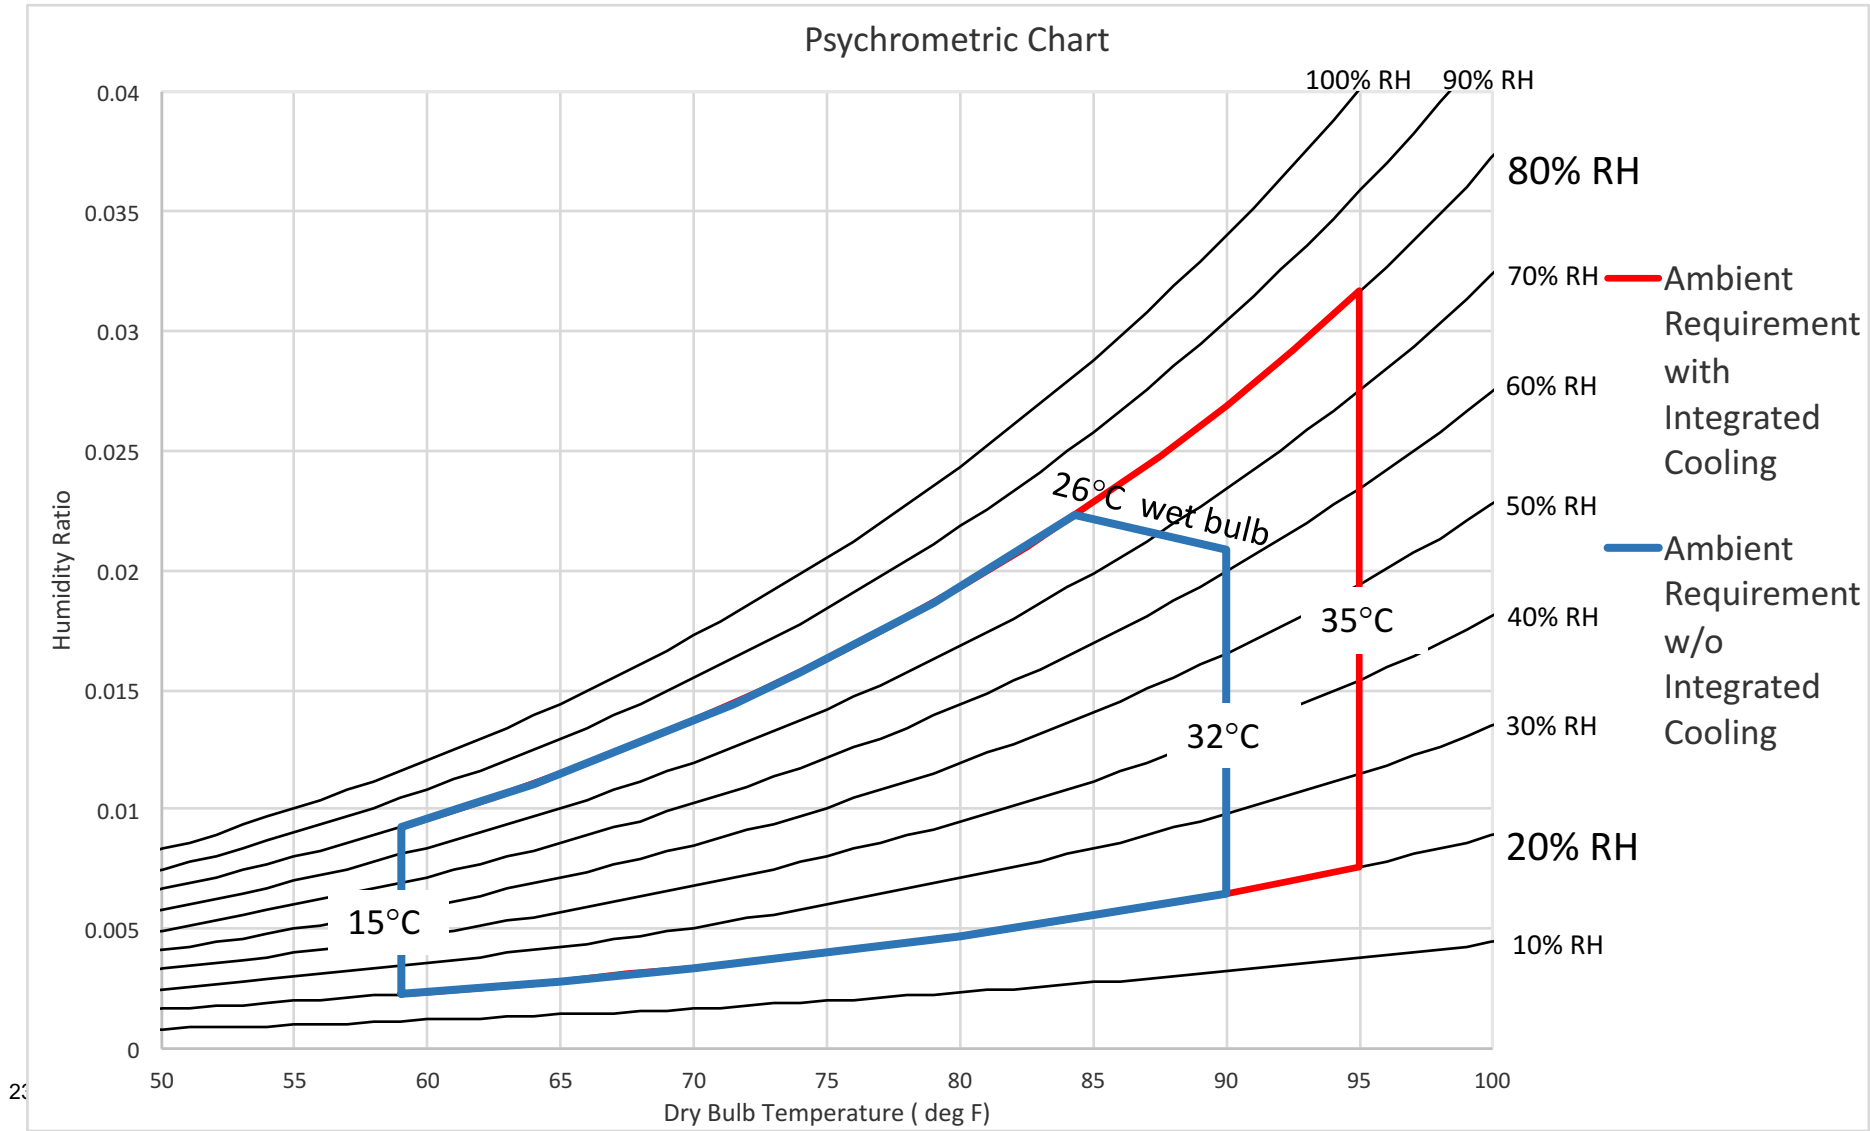
BY RICH MILLER ON MARCH 23, 2012 24 COMMENTS

140 299 91 41
Like Tweet Share +1






The Google data center in Belgium, which features no chillers and routes traffic to other facilities during hot spells. (Photo from Google)

Integrated Cooling Requirements



Integrated Cooling Air Flow



Thank you!

Contact:

Lee Jesionowski

ljesion@us.ibm.com

For additional information...

- Search the web for 'TS4500 Knowledge Center'



Disclaimers

- Copyright© 2007 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 intellectually 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.



Disclaimers (continued)

- 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.



Trademarks

- 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

