#SHAREorg

SHARE
Technology · Connections · Results

IBM Responses to Requirements Submitted by the SHARE MVSS Project

Barbara McDonald
IBM DFSMS Product Management
bawhite@us.ibm.com

August 6, 2012
Session 11916

SHARE in Anaheim 2012

SHARE
Technology · Connections · Results

Disclaimer

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.

Agenda

- IBM Responses to SHARE Requirements
  - Requirements Waiting for Responses
  - Requirements Open for Discussion
  - Existing requirements with updated responses
  - Requirements Available in R13

SHARE Requirements Summary

- Waiting for Response

Requirement #	Title	Status
SSMVSS11001	Force System Determined Blocksize (SDB)	**RJ – Rejected (Anaheim 2012)** Waiting for response (Atlanta 2012) Open for discussion (Anaheim 2011)
SSMVSS11013	Enhance DFSMSdss To Process An Offline DASD Volume	**AK – Acknowledged (Anaheim 2012)** Open for discussion (Atlanta 2012)

SSMVSS11001

- **Force System Determined Blocksize (SDB)**
  - **Response: RJ – Rejected (Anaheim 2012)**
    - If the application has hard coded the dcb blksize in the dcb then we shouldn't be ignoring it. The application could be dependent upon that and it could cause problems if we changed this 45+ year precedent. About the only time this would be acceptable is if the customer was writing a new application that understood this was happening or understood what their existing application required. But in the case of an existing application, most customers probably couldn't tell you what the application required and may not even have the source.
    - Concern is system programmers will get into trouble trying to force a blksize for those apps that hard code blksize in the DCB. Their first job will be to find those apps. IBM could change the tracker facility to help customers identify those apps that hardcode it.
- **Description**
  - The way the Force System Determined Blocksize option of the SMS Data Class as introduced in z/OS Release 10, when set to Y, creates the belief that SDB will be forced for a given data set. This is true during the creation of the data set only. Since most user's create and then use (open) the data set, the Force System Determined Blocksize option does not work when an application program provides a blocksize value in the DCB when issuing the Open macro.
- **Benefit:**
  - Since most files are allocated and opened within a step, this will provide a SDB as requested in the dataclas regardless of the DCB parm
- **Solution:**
  - This REQUEST is being submitted as a requirement to DFSMSdfp to provide the function to actually force SBD for data sets within a Data Class for the life of the data set regardless of the presence of any coded blksize value in the DCB JCL parm or program.

SSMVSS11013

- **Enhance DFSMSdss To Process An Offline DASD Volume**
  - **Response: AK – Acknowledged (Anaheim 2012)**
- **Description**
  - There are DASD volumes that need to be accessed by DFSMSdss that can not be brought online, such as duplicate VOLSER. These volumes need to be processed by functions such as DUMP, for backups, and COPY, for data promotion.
- **Benefit:**
  - To backup a DASD volume from a SYSPLEX that does not have tape drives.
  - To backup a DASD volume from a SYSPLEX with tape media that is stored offsite.
  - To copy a DASD volumes from a development SYSPLEX to promote data.
- **Impact:**
  - As the wall between production and development continues to get stronger, it becomes more difficult to provide work arounds.

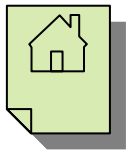

SHARE Requirements Summary

- Open for Discussion (previously reviewed)

Requirement #	Title	Status
SSMVSS10002	Excessive/invalid HSM RECALL commands generated by IKJEFT01	**More discussion needed (Anaheim 2012)** Returned for clarification (Atlanta 2012) AK – Acknowledged (Anaheim 2011)
SSMVSS11003	DFSMS: Change How System Managed Buffering Handles Empty Datasets	**More discussion needed (Anaheim 2012)** **Note:** Cannot find in FITS – need to resubmit. Open for discussion (Anaheim 2011)
SSMVSS11005	Increase the maximum VSAM Control Interval size (CISZ)	**More discussion needed (Anaheim 2012)** **Note:** Cannot find in FITS – may need to resubmit. Returned for clarification (Atlanta 2012) Open for discussion (Orlando 2011)

SSMVSS10002

- **Excessive/invalid HSM RECALL commands generated by IKJEFT01**
  - **Response: Returned for User Clarification**
    - TSO/E can not prescreen commands before passing the command string to the command processor. **However, it sounds like the customer simply wants HSM to screen the command before sending it to the coupling facility?** If so this requirement is for DFSMShsm.
- **Description:**
  - When IKJEFT01 passes HSEND RECALL commands to HSM using a dataset pattern, the catalog should be checked to ensure that HSM does not attempt to recall datasets that are already on L0. HSM currently issues a recall regardless of catalog status. Thousands of invalid recalls can flood the system, overwhelming the CRQ and its coupling facility structure.
- **Benefit:**
  - A dramatic decrease in overhead can be realized. Fewer MWEs in HSM, fewer HSM activity records will be written. More importantly, HSM recall users will not experience recall delays when excessive CRQ utilizatoin forces HSM recalls to fall back to LOCAL mode.
- **Solution:**
  - IKJEFT01 should do a catalog check before passing the HSEND RECALL request to HSM. Any recalls for files already on L0 should be dropped before enlisting HSM for the recall.

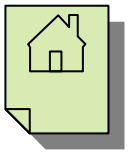

SSMVSS11003

- **DFSMS: Change How System Managed Buffering Handles Empty Datasets**
  - **Open for discussion**
    - VSAM is really WAD, however, its not the best design! There is a difference in the way NSR and LSR handle SEQ processing, LSR requires a POINT and NSR does not. COBOL is not issuing a POINT since its a NSR user. When SMB is used to convert it to use LSR under the covers, a POINT it now required.
    - **Note:** Cannot find in FITS – may need to be resubmitted. IBM recommendation is to resubmit the requirement to DFSMS, and we'll reassign the FITS requirement to COBOL.
- **Description**
  - Currently, if a program opens a VSAM data set with an ACB that specifies MACRF=DIR (which is what COBOL will specify if ACCESS IS RANDOM is coded on the File Definition) and the HURBA is zero, the open will fail if System Managed Buffering (SMB) is invoked. SMB selects ACCBIAS=DO, and that does not work for an empty data set. SMB should check the HURBA, and if it is zero ACCBIAS=DW should be selected instead.
- **Benefit:**
  - The application can open a data set that has a HURBA=0 if 1) the program is changed to specify MACRF=SEQ (or any combination other than MACRF=DIR) in Assembler or ACCESS IS DYNAMIC in COBOL, or 2) codes ACCBIAS=DW in the JCL on the DD statement for the data set. The first option is actually the correct fix, but in either case requires the application to make a change to a program or job that possibly has not changed in a number of years. If there are a lot of programs coded this way, the number of changes could be substantial. Making this change in System Managed Buffering would fix this coding deficiency without any adverse affects, and it would allow customers to continue rolling out SMB without fear of causing job failures.

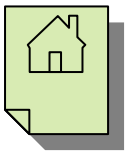

SSMVSS11005

- **Increase the maximum VSAM Control Interval size (CISZ)**
  - **Response: Returned for User Clarification**
    - VSAM Development will need more information to determine if the requirement is feasible.
    - Possible to support for LDS only?
- **Description**
  - CISZ(s) are currently limited to 32K, which limits the record size to 32-7 bytes. Increase the CISZs to 64K, 128K, 256K, 512K, 1024k or 1Meg. The 1Meg size would work well in conjunction with the z/OS Large page feature.
- **Benefit:**
  - Removing the 32k Limit on Application development will enable newer applications to be developed on z/OS, that would otherwise have to be done on "open" platforms where the 32K limit does not exist. To keep new application development on the z Platform.
- **Solution:**
  - z/OS does support greater than 32K I/O requests, but even if VSAM can't use them, physical records size could be used instead.
  - VSAM uses a Physical Record size, so a 64K CISZ, could be 4 physical records of 16K as an example.

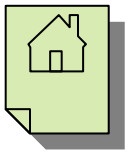

SHARE Requirements Summary

- Open for Discussion (new)

Requirement #	Title	Status
GO6SMG91037	DFP expand GDG limit beyond 255	Open for discussion (Anaheim 2012)
SSMVSS12001	Requirements should be accepted for storage 'hardware'	Open for discussion (Anaheim 2012)
SSMVSS12002	Document IECDDCE, the UCB DASD Class Extension Macro	Open for discussion (Anaheim 2012)
SSMVSS12003	REPRO MERGECAT Output should be condensed	Open for discussion (Anaheim 2012)
SSMVSS12004	Provide alternname MSGDD statement for DFDSS for warning or higher	Open for discussion (Anaheim 2012)
SSMVSS12005	Multi-Volume Catalogs	Open for discussion (Anaheim 2012)
SSMVSS12006	Tape Virtualization Engines should pre-cache during DR recovery	Open for discussion (Anaheim 2012)
SSMVSS12007	VTS systems need to supply mount statistics per logical volume	Open for discussion (Anaheim 2012)

GO6SMG91037

- **DFP expand GDG limit beyond 255**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - Many applications require GDS(s) for a year or more. The current limitation of 255 for GDG does not support this need. This causes confusion when attempting to go back more than the 255 generations to decide which GDS(s) would be required to input into a process.
- **Benefit:**
  - Considerable time and effort on the part of analysts and schedulers could be saved by reducing confusion and time in the attempt to decide which GDG's are required for processing. Additionaly, cost associated to reruns, related to this process would be eliminated.
- **Solution**
  - Raise the GDG limit to the max possible for the size of the field. This will provide the desired flexibility.
- **Discussion:**
  - I agree with this requirement, and will vote for it, but as for the 'year or more' I'd rather see implementation of something that allowed a date or time to be automatically substituted instead of a generalized generation. When I get some time I will try to write up that requirement.

SSMVSS12001

- **Requirements should be accepted for storage 'hardware'**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - Today's intelligent storage "devices" are really specialized computer systems in their own right. As such, IBM sells us a set of hardware with an operating system and various pieces of software which we interact with (one good example, the IBM TS77xx series of Virtualization Engines, which have a web management interface and BVIR for information retrieval). Customers need to be able, through the requirements process, to request improvements and enhancements to this software, just like we do with other software that we purchase from IBM. This requirement is being put in the Storage Management project because it is primarily DASD and tape equipment which has these features.
- **Benefit:**
  - Businesses benefit by not spending time and money on implementing workarounds and "surrounding software" for things which could be implemented via improvements to the software components of the storage hardware. Businesses also benefit by being able to improve their operational processes through needed enhancements to their intelligent storage systems. IBM may benefit by providing ease-of-use or other benefits which enhance sales.
- **Solution**
  - Suggestion is to take requirements for storage hardware software by product family (ex: IBM TS77xx) when appropriate, product name (ex: IBM TS7740), and component (for example BVIR or Web Management Interface).
- **Discussion:**
  - During disaster recovery testing, for an organization using a purely virtual tape setup, DASD volumes are recovered from backups on virtual tape. There currently is not a great method for making sure those virtual tapes are actually on the cache in the tape hardware. The current method requires retrieving the volume inventory via BVIR, and running jobs which reference the virtual volumes to force a recall in the back end. It would be far better if the current programming interface could be used to send a list of virtual volumes we would like recalled, in one request, _or_ if tape entry had an option which would cause the virtual volumes on a newly-entered stacked volume to automatically move to the cache. Both of these ideas would be changes to the _software_ which runs on these computer systems that the mainframe views as tape 'hardware'. SHARE members need a way to request improvements to this software.

  - I find this requirement peculiar in that during my 10 year plus tenure as IBM lead requirements rep, I always accepted in any form requirements for storage products software and hardware (in any form - paper napkins from SCIDS included) and worked directly with IBMers to provide a response within 6 months. What is stopping the customer from submitting a requirement? Even if the requirement was not for a storage device, I simply forwarded it to the appropriate project and/or IBM representative for handling. Keep in mind that customers are not limited to submitting requirements through SHARE. They may also submit requirements through their local branch/marketing team via a PASR or the RPQ process.

SSMVSS12002

- **Document IECDDCE, the UCB DASD Class Extension Macro**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - IECDDCE, otherwise known as the DCE, the UCB DASD Class Extension Mapping Macro, found in SYS1.MODGEN, remains undocumented in any publicly disclosed IBM publications. A search of SYS1.MODGEN and SYS1.MACLIB will surface numerous references to it. Moreover, the DCE is used prolifically by software and hardware vendors alike. Incorporate transparent and robust documentation in the DFSMSdfp Advanced Services or z/OS Data Areas manual for this control block.
- **Benefit:**
  - Whenever use of this macro is required, the only means by which a developer has at his/her disposal to surface fields and offsets within this macro is to research the macro - line-by-line. This is a grossly unproductive and inefficient use of a developer's time; additionally, a user has no idea if and when changes have been made to the DCE to add support for new z/OS software, hardware, or system control features. In addition, comments within the macro are poorly written and illuminate little with regard to use of many fields within the control block. Publishing this information will lift these restrictions upon software and hardware developers.
- **Solution**
  - Incorporate complete, robust, and transparent documentation in the DFSMSdfp Advanced Services or z/OS Data Areas manual for the IECDDCE UCB DASD Class Extension Mapping Macro.

SSMVSS12003

- **REPRO MERGECAT Output should be condensed**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - REPRO MERGECAT currently produces up to 6 lines for each entry that is processed: one blank line, one or two for the IDC0639I message(Sphere conversion started), a blank line, and one or two for the IDC01402I (Sphere conversion ended). The one or two lines depend it seems on the length of the dataset name. This should be condensed to one line per entry, with indicators of what worked (or didn't).
- **Benefit:**
  - Improve the usability of the report when searching for exceptions.
- **Solution**
  - Suggestion: the one line should contain a message ID, a blank, a dataset name, and indicators for 'conversion started' 'conversion ended' and 'error encountered'. Error messages to follow the entry (if any), possibly indented.
- **Discussion:**
  - Suggest rewriting this requirement to make the condensed output an option.
  - I do not agree with making it an option, unless someone has use case for the 6 lines per entry that makes business sense. If I were to ask for an option, I'd want the condensed output to go to a separate DD.

SSMVSS12004

- **Provide alternname MSGDD statement for DFDSS for warning or higher**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - When using DFDSS there can be thousands of messages produced. Looking through this listing for the odd WARNING, ERROR or SEVERE message type can be a challenge. It would be helpful to find a way to redirect these messages and leave the informational where they are.
- **Benefit:**
  - This change would reduce or eliminate the need to manually or programmatically scan the output to find the issues.
- **Solution**
  - A new parm entry like ERRMSG=ddname to redirect the error messages to a new DDNAME of my choice. If not coded, then normal behavior.
- **Discussion:**
  - This requirement resonates strongly and positively with me as a valuable productivity aid. I only wish I could reclaim the time I've spent plowing through mounds of DFSMSdss output to find the genesis of a failed job - typically, RC=8....
    - Having an alternate MSG DD statement to segregate output - much in the same vain as is currently implemented within compilers and SMP/E - will go a long way to resolving complex DFSMSdss backup, recovery, and replication job stream problems in a more timely, orderly, and less error-prone fashion.
    - Clearly, it will eliminate isolation of errors when visually inspecting DFSMSdss messages inter-leaved with valid job stream output.

SSMVSS12005

- **Multi-Volume Catalogs**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - Remove the 123 extent limit for Catalogs.
- **Benefit:**
  - Less Catalog reorgs/redefine due to reaching the 123 extent limit supporting for more 24x7 availability.
- **Solution**
  - Allow Catalogs to extend to more than 1 volume the same as other VSAM data sets.
- **Discussion:**
  - Yikes. I can only see catalog performance go down the drain. I don't even like seeing them in extents or large CI/CA splits. It's bad enough when a catalog volume is locked up by a reserve. Now you want to multiply that problem? Once a catalog gets too large for a volume it seems like it's time to split it. If necessary the implementation of multi-level alias names may be required to split a catalog that contains only one alias level's data set names. Can you explain the reason for this request? I just don't see the need.
  - I also think multi-level aliases provide enhanced availability in a better way; the main problem with multi-level alias is that as far as I've been able to determine it takes an IPL to change the number of levels (I may be wrong).

SSMVSS12006

- **Tape Virtualization Engines should pre-cache during DR recovery**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - When recovering a TSxxxx Virtualization Engine which uses physical cartridges, the customer should have the option of having all logical volumes on inserted physical cartridges recalled to the disk cache - to the limit of the disk cache threshold.
- **Benefit:**
  - Staging to disk cache at recovery time will eliminate a lot of physical cartridge contention during z/OS (or other OS) recovery. Doing recalls in parallel as described in the suggested solution will speed the necessary recalls as well.
- **Solution**
  - The interfaces pertaining to I/O station interface needs a UI item, button, or switch which acts in concert with the "Copy/Export Recovery" function to toggle "DR Insert" vs. "Normal Insert". During DR insert, all inserted physical tapes ought to be queued for mount by all drives but one, and each valid logical volume on the physical tapes recalled to disk cache until a threshold is reached. When the threshold is reached, recalls should quiesce (assuming those in progress had already allocated the necessary space for recall, the event logged, and cartridge insert return to "normal". The one drive omitted is there in case of a hardware error which might be correctable via swapping to a different drive.

SSMVSS12007

- **VTS systems need to supply mount statistics per logical volume**
  - **Open for Discussion (Anaheim 2012)**
- **Description**
  - While the performance of logical volume mounts in a Virtualization Engine is not necessarily controllable, nor is it necessarily desirable to control it, we still need the statistics related to mounting logical volumes - including, but not limited to, the time to mount, the physical cartridge mounted if a recall was necessary, the mount time of the physical cartridge if necessary, indicators of whether the physical volume was already mounted requiring a wait for this logical volume mount request, or whether action was necessary to make space in the disk cache for a mount request. I'm sure there may be others.
- **Benefit:**
  - These statistics are necessary to answer performance / elapsed time questions for systems using the VE - for example, increase in elapsed time of a critical batch job may be questioned; we currently can say 'the tape mount took some time' but we cannot directly answer why.
- **Solution**
  - Enhance Allocation or O/C/EOV (whichever applies) to record tape mount time, have VE send the statistics record as part of the 'drive is ready, tape is mounted' completion and add that information to the mount record.

SHARE Requirements Summary

- Existing requirements with updated responses

Requirement #	Title	Status
SSMVSS11014	ICKDSF: Automated Control Statement Generation for 1,000s of Vols	**RJ – Rejected (Anaheim 2012)** Open for discussion (Atlanta 2012)
SSMVSE10018	Provide creation date for alias	**AC – Accepted (Anaheim 2012)** RC – Recognized (Anaheim 2011)
GO6SMG90032	DFDSS copy cataloged data set name to unlike data set name	**AV – Available R6 (Anaheim 2012)** RC – Recognized (Jan 1996)

SSMVSS11014

- **ICKDSF: Automated Control Statement Generation for 1,000s of Vols**
  - **Response: RJ – Rejected (Anaheim 2012)**
    - ISMF is already used to build batch jobs for ranges of volumes. ISPF tailoring is also available that could be used.
- **Description**
  - A complementary tool, utility, or modification to ICKDSF is required to automate generation of its control statements. Such function shall allow for specification and incorporation of various formatting capabilities (e.g. creation of an indexed or non-indexed VTOC and OWNER) for various UCB ranges and consistently named volume serial numbers when executing ICKDSF with PARM=NOREPLY for very large numbers (in the thousands) of storage volumes.
- **Benefit:**
  - Given the historical context when ICKDSF was first written (1978), addressable numbers of storage volumes using high performance replication technology in the last 2 decades have transcended the

SSMVSE10018

- **Provide creation date for alias**
  - **Response: AC – Accepted (Anaheim 2012)**
- **Description**
  - We are attempting to cleanup obsolete HLQs (High Level Qualifiers). If an alias has no associated datasets, there is no easy way to determine whether this is a new alias and no datasets have been created or this is an obsolete alias that should be deleted. The alias creation date would also be useful for aliases that are related to datasets.
- **Benefit:**
  - Identifying the creation date would simplify alias cleanup.
- **Solution:**
  - LISTCAT needs to show the creation date. Creation date is currently not stored.

Impact:

- It is difficult to determine if HLQs are obsolete or current. Since the information isn't stored, there isn't an alternative solution.

GO6SMG90032

- **DFDSS copy cataloged data set name to unlike data set name**
  - **Response: AV – Available R6 (Anaheim 2012)**
- **Description**
  - DFDSS 2.5.0 Users Guide under "Rules for moving to pre-allocated target datasets" specifies "an existing dataset qualifies as a pre-allocated target for dataset COPY if the Format 1 DSCB names match." Under DFSMS, this is not possible. We would like the ability to move data from one catalogued dataset to another catalogued dataset.
- **Benefit:**
  - If DFDSS is to be the data mover, this function needs to be supported.

SHARE Requirements Summary

- **Requirements Available in R13**

Requirement #	Title	Status
SSMVSS09005	DFHSM command to release DASD recalls	AV – Available (Orlando 2011) **AC- Accepted (2009)**

SSMVSS09005

- **DFHSM command to release DASD recalls**
  - **Response: AV – Available (Orlando 2011)**
- **Description:**
  - Once a HOLD RECALL command has been issued, there is no way to release DASD recalls without also releasing tape recalls. The requirement asks for a way to specifically release DASD recalls.
- **Benefit:**
  - If an operator experiences a major problem with a tape subsystem s/he often responds with HOLD RECALL to avoid recall failures. However, once this is done, there is no way to release DASD recalls without issuing a general RELEASE RECALL. This reintoduces the tape issues (which can sometimes cause problems beyond a simple recall failure). This can be a serious problem if the tape subsystem will be down for an extended period of time due to maintenance (which is also more likely than having a DASD subsystem down for maintenance). Work that would otherwise be able to run is held up despite the fact all its datasets are only migrated to DASD.
- **Impact:**
  - Work is unintentionally held up for no reason. Once the error is made, it can be extremely difficult to undo if hardware work has begun in the meantime.
- **Solution:**
  - Introduce a RELEASE RECALL(DASD) command.

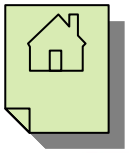

System z Social Media

- System z official Twitter handle:
  - @ibm_system_z
- Top Facebook pages related to System z:
  - Systemz Mainframe
  - IBM System z on Campus
  - IBM Mainframe Professionals
  - Millennial Mainframer
- Top LinkedIn Groups related to System z:
  - Mainframe Experts Network
  - Mainframe
  - IBM Mainframe
  - System z Advocates
  - Cloud Mainframe Computing
- YouTube
  - IBM System z

- Leading Blogs related to System z:
  - Evangelizing Mainframe (Destination z blog)
  - Mainframe Performance Topics
  - Common Sense
  - Enterprise Class Innovation: System z perspectives
  - Mainframe
  - MainframeZone
  - Smarter Computing Blog
  - Millennial Mainframer

Trademarks and Disclaimers

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:

CICS*
DB2*
DFSMS
DFSMSdfp
DFSMSdss
DFSMShsm
DFSMSrmm
DFSORT
DFSMS
DS4000
DS6000
DS8000
Enterprise Storage Server*
ESCON*
FICON*
FlashCopy*
GDDM*
GDPS*
geoManager*
HiperSockets
HyperSwap
IBM*
IBM logo*
ImagePlus*
IMS
Intelligent Miner
Language Environment*
Lotus*
MQSeries*
Multiprise*
OMEGAMON*
OS/390*
Parallel Sysplex*
PR/SM
QMF
RACF*
Rational*
RMF
System i
System z
System z9
System Storage
Tivoli*
TotalStorage*
Virtualization Engine
VisualAge*
VM/ESA*
VSE/ESA
VTAM*
WebSphere*
z/Architecture*
z/OS*
z/VM*
z/VSE
zSeries*
zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies**:**

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both

Microsoft, Windows,Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

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.

UNIX is a registered trademark of The Open Group in the United States and other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

NOTES:

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. Users of this document should verify the applicable data for their specific environment.

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

Information is provided “AS IS” without warranty of any kind.

Trademarks and Disclaimers (continued)

NOTES:

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 are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

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.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

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.