



IBM GPFS tuning guidelines for deploying SAS on IBM Power servers

Last updated: February 2013

Introduction

IBM® Power® servers along with IBM enterprise storage and IBM General Parallel File System (IBM GPFS™), provide a highly resilient and high performance infrastructure required for deploying SAS in a distributed grid environment.

This paper is a quick-reference guide for GPFS configuration and tuning parameters for deploying SAS on IBM Power servers with IBM AIX®.

Important note: The tunable given in the paper are based on a SAS mixed analytics workload run on a four-node SAS Grid, deployed on an IBM Power Systems™ server running AIX 7.1 with storage area network (SAN) attached IBM XIV Storage System. Refer to the paper *SAS 9.3 grid deployment on IBM Power servers with IBM XIV Storage System and IBM GPFS* published at ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102192 for detailed description of these tunable and tune them based on your workload and back-end storage system.

GPFS configuration and tunable parameters

SAS applications perform large sequential read and write operations. Some of the new SAS business intelligence (BI) applications do some random access, but for the most part, SAS workloads can be characterized as predominately large sequential I/O requests with high volumes of data.

Important GPFS tunable and their values include:

- Block size: 1 MB
 - Choose based on storage stripe size and keep it multiple of SAS BUFSIZE
- Block allocation type: Cluster
 - Cluster allocation works well for SAS I/O, which is predominantly sequential
- pagepool: At least 4 GB
- seqDiscardThreshold: 1 GB
 - Maximum size of the file for caching during sequential access. Change depending on the typical size of the files in workload
- prefetchPct: 40 (default 20)
 - Increasing to 40 can help SAS I/O, which is predominantly sequential
- maxMBpS: 5000
- maxFilesToCache: 20000
- GPFS file system mount options: *rw, mtime, atime, dev*
- Separate file system metadata from data while creating Network Shared Disks (NSDs)
 - For example, when storage has solid-state drives (SSDs), separate metadata to NSDs with SSD caching enabled
- Number of logical unit numbers (LUNs): 8 or more per file system

Other relevant tunable

- **SAS configuration:** The following SAS options were used to increase the I/O throughput of SAS applications.

SAS configuration	
MSIZE	2048 MB
SORTSIZE	256 MB
BUFSIZE	64 KB
UBUFSIZE	64 KB
IBUFSIZE	32767
BUFNO	10
UBUFNO	10
IBUFNO	10
ALIGNSASIOFILES	

- **AIX tunable:** Refer to *SAS AIX 5L, AIX 6, and AIX 7 Tuning Guide* published at ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101529
- **FC adapter and LUN tunable:** It is important to tune Fibre Channel (FC) adapters and LUNs along with GPFS tuning to achieve optimal performance. Refer to the *SAS 9.3 grid deployment on IBM Power servers with IBM XIV Storage System and IBM GPFS* paper published at ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102192.

Changing GPFS tunable: Changing just one tunable may not help in performance. It is important to understand how different tunable work with each other and find the right combination by running tests. For example, for the workload used in testing, setting this combination - *seqDiscardThreshold=1GB, pagepool=8GB, prefetchPct=40*, gave a slightly better performance.

GPFS licensing guidelines: A GPFS cluster requires a minimum of one node running a GPFS server license for cluster management. All other nodes can run with a GPFS client license. However, to ensure fault tolerance, licensing a minimum of two GPFS server nodes (or more, depending on customer availability requirements) per cluster is recommended. Refer to GPFS wiki for more details ibm.com/developerworks/wikis/display/hpccentral/GPFS+Licensing.

GPFS performance monitoring

Refer to the subsection *GPFS specific performance monitoring tools and techniques* under Appendix A of the *SAS 9.3 grid deployment on IBM Power servers with IBM XIV Storage System and IBM GPFS* paper published at ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102192.

Summary

The paper describes the summary of GPFS configuration and tuning parameters for SAS deployment on IBM Power servers with AIX OS. The tunable values might vary depending on your workload and back-end storage.

Trademarks and special notices

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

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Other company, product, or service names may be trademarks or service marks of others. Information is provided "AS IS" without warranty of any kind.