



# IBM Spectrum Scale general tuning guidelines for deploying SAS on IBM Power servers with AIX

Last updated: Oct 2016  
Narayana Pattipati (npattipa@in.ibm.com)

## Introduction

IBM® Power® servers along with IBM enterprise storage and IBM Spectrum Scale™ (formerly 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 general configuration and tuning guidelines for deploying IBM Spectrum Scale on IBM Power servers with IBM AIX® for SAS.

**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 or AIX 7.2 with storage area network (SAN) attached IBM XIV Storage System or IBM FlashSystem. 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/atmastr.nsf/WebIndex/WP102192](http://ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102192) for detailed description of these tunable and tune them based on your workload and back-end storage system.

## Spectrum Scale 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 Spectrum Scale tunable and their values include:

- Block size:
  - Choose based on storage stripe size and keep it multiple of SAS BUFSIZE
  - Use 1 MB block size for IBM XIV Storage System
  - User 256 or 512 KB for IBM FlashSystem
- Block allocation type: Cluster
  - Cluster allocation works well for SAS I/O, which is predominantly sequential
- pagepool: At least 4 GB (per node in the cluster)
- 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
- Spectrum Scale 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

*IBM Spectrum Scale tuning guidelines for deploying SAS on IBM Power servers*

© Copyright IBM Corporation, 2016

- Number of logical unit numbers (LUNs): 8 or more per file system
- scatterBuffers: Applicable for Spectrum Scale 3.5 version only. The scatterBuffers parameter affects how Spectrum Scale organizes file data in the pagepool. The default is scatterBuffers=yes (Starting in Spectrum Scale 3.5). If your workload is mostly sequential disable the scatterBuffers feature by running `mmchconfig scatterBuffers=no`. If you are not observing full blocksize IO's being sent to the storage during sequential IO operations disabling scatterBuffers or increasing scatterBufferSize may help (See scatterBufferSize below). This is an undocumented feature.
- scatterBufferSize: The scatterBufferSize parameter sets the size of the scatter buffer use by Spectrum Scale. The default is 32KB (Starting in Spectrum Scale 3.5). When tuning for sequential IO workloads it may help to increase scatterBufferSzie to be the same as the file system blocksize. If you are not observing full blocksize IO's being sent to the storage during sequential IO operations disabling scatterBuffers or increasing scatterBufferSize may help.
- Recommended version (as of April 2016): Spectrum Scale 4.1.1.3 or later version if you are using Spectrum Scale 4.1. Although Spectrum Scale version 3.5 is supported till end of life, it is recommended to use 4.1.

### Other relevant tunable

- **SAS configuration:** If you are using SAS release 9.4, there is no need to change any of the SAS options. If you are using SAS release 9.3, it is recommended to use the following SAS options.

SAS configuration	
MEMSIZE	2 GB
SORTSIZE	1 GB
CPUCOUNT	4
BUFSIZE	64 KB
UBUFSIZE	64 KB
IBUFSIZE	32767
BUFNO	10
UBUFNO	10
IBUFNO	10
ALIGNNSASIOFILES	

- **AIX tunable:** Refer to *SAS AIX 5L, AIX 6, and AIX 7 Tuning Guide* published at [ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101529](http://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 Spectrum Scale 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](http://ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102192).

**Changing Spectrum Scale 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.

**Spectrum Scale licensing guidelines:** A Spectrum Scale cluster requires a minimum of one node running a Spectrum Scale server license for cluster management. All other nodes can run with a Spectrum Scale client license. However, to ensure fault tolerance, licensing a minimum of two server nodes (or more, depending on customer availability requirements) per cluster is recommended. Refer to Spectrum Scale wiki for more details: <https://ibm.biz/Bd4MJb>

## Spectrum Scale 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](http://ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102192).

## Summary

The paper describes the summary of Spectrum Scale 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.

## Resources

Additional resources for Spectrum Scale tunable:

- Spectrum Scale wiki tuning parameters  
[ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General Parallel File System \(GPFS\)/page/Tuning Parameters](http://ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20(GPFS)/page/Tuning%20Parameters)
- Spectrum Scale tunable for SAS Software
- [ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General Parallel File System \(GPFS\)/page/SAS SoftwareSAS Business Analytics Deployment on IBM POWER8 Processor-Based Systems: With IBM XIV Storage System and IBM FlashSystem](http://ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20(GPFS)/page/SAS%20SoftwareSAS%20Business%20Analytics%20Deployment%20on%20IBM%20POWER8%20Processor-Based%20Systems%3A%20With%20IBM%20XIV%20Storage%20System%20and%20IBM%20FlashSystem)  
<http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5288.html>

## 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](http://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](http://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.