

IBM Storage Scale Frequently Asked Questions and Answers

Questions and answers

1. General questions:

- 1.1 Where can I find ordering information for IBM Storage Scale?
- 1.2 Where can I find the documentation for IBM Storage Scale?
- 1.3 What is there beyond the standard documentation that can help me learn more about and use IBM Storage Scale?
- 1.4 How can I ask a more specific question about IBM Storage Scale?
- 1.5 Does IBM Storage Scale participate in the IBM Academic Initiative Program?
- 1.6 Is IBM Storage Scale available in IBM PartnerWorld®?
- 1.7 Does IBM Storage Scale have a trial program?
- 1.8 Where can I find the documentation for IBM Storage Scale RAID?
- 1.9 Where can I find detailed information about stabilized, deprecated, and discontinued features of IBM Storage Scale?
- 1.10 Where can I find a list of the authorized program analysis reports (APARs) that are resolved for IBM Storage Scale?

2. Software questions:

- 2.1 What is supported on IBM Storage Scale for AIX, Linux, Power®, and Windows?
- 2.2 What is the IBM Storage Scale support position regarding *clone* Linux distributions (CentOS, Rocky Linux, ROCK, White box Linux, etc.) ?
- 2.3 What Linux distributions are supported by the integrated protocols access methods in IBM Storage Scale V4.1.1 and later?
- 2.4 What is the impact on the /dev, /proc/mounts, /etc/mtab directories and the mount command for IBM Storage Scale for Linux due to the recent changes in systemd? What happened to the block device in /dev? Why is the /dev/ prefix missing from the output of the mount command and also from /proc/mounts and /etc/mtab?
- 2.5 What are the limitations of IBM Storage Scale support for Windows?
- 2.6 What are the requirements for the use of OpenSSH on Windows nodes?
- 2.7 What levels of Cygwin are supported by IBM Storage Scale on Windows?
- 2.8 Can different IBM Storage Scale maintenance levels coexist?
- 2.9 Are there any requirements for Clustered NFS (CNFS) support in IBM Storage Scale?
- 2.10 Does IBM Storage Scale support NFS V4?
- 2.11 Are there any considerations for the use of the persistent reservation (PR) support in IBM Storage Scale?
- 2.12 What are the requirements/limitations for using native encryption in IBM Storage Scale Advanced Edition or Data Management Edition?
- 2.13 Are there any considerations when utilizing the Simple Network Management Protocol (SNMP)-based monitoring capability in IBM Storage Scale?
- 2.14 What are the current limitations and advisories for using the mmbackup command?
- 2.15 What are the current limitations and advisories for using Scale Out Backup and Restore (SOBAR)?
- 2.16 What are the current limitations for using the File Placement Optimizer (FPO) function?
- 2.17 What are the current limitations for using the Active File Management (AFM) Async DR function?
- 2.18 What are the current limitations for using the Active File Management (AFM) function?
- 2.19 What are the current limitations common to both the Active File Management (AFM) and AFM DR functions?
- 2.20 What is the currently recommended transport protocol for AFM and AFM DR data transfers?
- 2.21 What is the current limitation for using the snapshot restore function?
- 2.22 What are the current requirements when using local read-only cache?
- 2.23 What are the current requirements/limitations for using the Cluster Configuration Repository (CCR)?
- 2.24 What are the current requirements/limitations for using Ubuntu?
- 2.25 What are the current requirements/limitations for IBM Storage Scale for Linux on Z?
- 2.26 What are the current requirements/limitations for the deadlock amelioration function in IBM Storage Scale?
- 2.27 What are the requirements/limitations for using the IBM Storage Scale GUI?
- 2.28 What are the requirements/limitations for using the compression function?
- 2.29 What are the requirements/limitations for using the Quality of Service (QoS) function?
- 2.30 What considerations are there when running on SELinux?
- 2.31 What enhancements to quota management are available?
- 2.32 What are the current limitations of quota management?
- 2.33 Is file immutability supported with IBM Storage Scale?
- 2.34 Has file immutability been assessed for compliance?
- 2.35 Is there any guidance for RHEL 8 installations on IBM Storage Scale?

- 2.36 What are the requirements/limitations for using the file clone function?
- 2.37 What are the requirements/limitations for using the memory-mapped (mmap) function?
- 2.38 Are IBM Storage Scale packages signed by IBM?
- 2.39 What are the components required for GPUDirect Storage (GDS) support with IBM Storage Scale?
- 2.40 What are the requirements/limitations for manipulating the system.nfs4_acl extended attribute directly in IBM Storage Scale?
- 2.41 Does IBM Storage Scale support the secure boot specified by the Unified Extensible Firmware Interface (UEFI)?
- 2.42 Is IBM Storage Scale supported on Advanced RISC Machine (ARM) architectures?
- 2.43 What are the requirements and limitations for using the asymmetric data replication function?
- 2.44 What are the options that I can use with the mmsmb config command?

3. Machine questions:

- 3.1 What are the minimum hardware requirements for an IBM Storage Scale cluster?
- 3.2 On what servers is IBM Storage Scale supported?
- 3.3 What interconnects are supported for IBM Storage Scale daemon-to-daemon communication in my cluster?

4. Disk questions:

- 4.1 What disk hardware has IBM Storage Scale been tested with?
- 4.2 What Fibre Channel (FC) Switches are qualified for IBM Storage Scale usage and is there a FC Switch support chart available?
- 4.3 Can I concurrently access disks from both AIX and Linux nodes in my IBM Storage Scale cluster?
- 4.4 What devices does IBM Storage Scale support with SCSI-3 Persistent Reservations?
- 4.5 What considerations are there when setting up DM-MP multipath service?
- 4.6 Are there any steps that need to be taken before disks are used by IBM Storage Scale on AIX?
- 4.7 Does IBM Storage Scale support Logical Volumes (LVs) ?
- 4.8 Does IBM Storage Scale support AIX raw hdisks (rhdisk) ?
- 4.9 Does IBM Storage Scale for Linux on Z support Direct Attached Storage Devices (DASD)?
- 4.10 Does IBM Storage Scale support 4K disk sectors?
- 4.11 What are the considerations for using block storage systems that support thinly provisioned volumes (thin provisioning)?
- 4.12 What are the considerations for using IBM Storage Scale with block storage that supports data reduction features including compression or deduplication?
- 4.13 How is IBM Storage Scale capacity determined for licensing purposes with thinly provisioned volumes or data reduction?
- 4.14 Does IBM Storage Scale for Linux on Z work with disk hardware replication?
- 4.15 Does IBM Storage Scale for Linux on Z support HyperSwap® or SVC Stretch Clusters?
- 4.16 What disk connection technologies can be used for NSD storage?

5. Scaling questions:

- 5.1 What are the IBM Storage Scale cluster size limits?
- 5.2 What are some scaling considerations for the protocol (CES) nodes?
- 5.3 What is the current maximum tested limit for SMP scaling?
- 5.4 What is the current limit on the number of nodes that may concurrently join a cluster?
- 5.5 What is the limit of remote clusters that a client node can join?
- 5.6 What is the limit of remote clusters that can join a local cluster?
- 5.7 What are the current file system size limits?
- 5.8 What is the current limit on the number of mounted file systems in an IBM Storage Scale cluster?
- 5.9 What is the architectural limit of the number of files in a file system?
- 5.10 What is the architectural limit of the number of disks in a file system?
- 5.11 What are the limitations on IBM Storage Scale disk size?
- 5.12 What is the limit on the maximum number of groups a user can be a member of when accessing an IBM Storage Scale file system?
- 5.13 What are the current limits on the number of filesets in an IBM Storage Scale file system?
- 5.14 What is the maximum number of remote clusters for which an administrator can enable fileset access control?
- 5.15 What are the current limits on the number of snapshots in an IBM Storage Scale file system?
- 5.16 What is the current limit for the number of data and metadata replicas?
- 5.17 What is the limitation on pathname length?
- 5.18 What is the current maximum IBM Storage Scale file size limit?
- 5.19 What is the limitation on symbolic links?

6. Configuration and tuning questions:

- 6.1 What specific configuration and performance tuning suggestions are there?
- 6.2 What configuration and performance tuning suggestions are there for IBM Storage Scale when used primarily for Oracle databases?
- 6.3 What configuration requirements exist for utilizing Remote Direct Memory Access (RDMA)?

- 6.4 What configuration and performance tuning suggestions are there for the Active File Management function of IBM Storage Scale?
- 6.5 Sometimes IBM Storage Scale appears to be handling a heavy I/O load, for no apparent reason. What could be causing this?
- 6.6 What considerations are there when using IBM Storage Protect with IBM Storage Scale?
- 6.7 Does IBM Storage Scale use OpenSSL for RPC secure communication?
- 6.8 What ciphers are supported for use by IBM Storage Scale?
- 6.9 When I allow other clusters to mount my file systems, is there a way to restrict access permissions for the root user?
- 6.10 How do I determine the maximum size of the extended attributes allowed in my file system?
- 6.11 What are configuration considerations when using IPv6?
- 6.12 How should IBM Storage Scale Advanced Edition or Data Management Edition be configured to only use FIPS 140-2-certified cryptographic engines?
- 6.13 What are the configuration requirements for the use of the integrated protocols access methods?
- 6.14 What considerations are there when using IBM Spectrum Archive with IBM Storage Scale?

7. Virtualization questions:

- 7.1 How do I determine whether a server license or a client license is required when running IBM Storage Scale in VMs in a virtualized environment?
- 7.2 Is IBM Storage Scale on Linux (x86 and Power) supported in a virtualization environment?
- 7.3 Is IBM Storage Scale for Linux on Z supported in a virtualization environment?
- 7.4 Is IBM Storage Scale on Windows supported in a virtualization environment?
- 7.5 Can IBM Storage Scale run in a Workload Partitioning (WPAR) environment?
- 7.6 Does IBM Storage Scale support exploitation of the Virtual I/O Server (VIOS) features of POWER processors?
- 7.7 Does IBM Storage Scale support PowerVM® Shared Storage Pool (SSP)?
- 7.8 Are there any virtualization considerations when using an Oracle databases?
- 7.9 Can Linux on Z with IBM Storage Scale run in logical partition (LPAR) mode or on z/VM® as guest operating system?
- 7.10 Are there any limitations for using IBM Storage Scale in a virtualization environment?

8. Integrated Protocol Server questions

- 8.1 What are the requirements to use the protocol access methods integrated with IBM Storage Scale?
- 8.2 What is the minimum hardware requirement for a protocol node?
- 8.3 What are some configuration considerations when deploying the protocol functionality?
- 8.4 What is the guidance for NFS clients to be configured with the CES NFS function?
- 8.5 Can we configure protocol nodes with an IBM Storage Scale System/GSS?
- 8.6 Are there any limitations that I should be aware of before using the integrated CES Protocol function?
- 8.7 Where can I find additional information about protocols?
- 8.8 How do I reduce the logging level for performance monitoring of Swift?
- 8.9 What are the limitations when multiple file systems are exported by NFS?
- 8.10 How to effectively use VMWare with IBM Storage Scale?
- 8.11 What are the limitations when upgrading the SMB service?

9. Hadoop support questions:

- 9.1 What platforms are supported for Hadoop on IBM Storage Scale?
- 9.2 What versions of Hadoop are supported by IBM Storage Scale version?
- 9.3 What Hadoop distributions are supported by IBM Storage Scale?
- 9.4 What IBM Storage Scale licenses do I need to use the IBM Storage Scale Hadoop connector?
- 9.5 Can Hadoop connector be used for IBM Storage Scale that has shared storage (including IBM Storage Scale System)?
- 9.6 Can Hadoop connector be used for IBM Storage Scale that has shared storage (including IBM Storage Scale System) and internal storage (FPO pool) in the same file system?
- 9.7 What open-source Hadoop components are certified using IBM Storage Scale connector?
- 9.8 Why do you need special IBM Storage Scale connector instead of using Hadoop local file system (file:///) if not using FPO/internal storage?
- 9.9 What are the differences between the Hadoop connector and the HDFS Transparency connector ?
- 9.10 What are the requirements/limitations for using the IBM Storage Scale HDFS Transparency connector?
- 9.11 What are the current advisories for the IBM Storage Scale Hadoop connector?

10. Swift Object protocol support questions:

- 10.1 What considerations are there when using OpenStack Software with IBM Storage Scale?
- 10.2 What are the new features of IBM Storage Scale for Swift Object Storage?
- 10.3 How should I ensure that unauthorized users cannot access my object data when using IBM Storage Scale for Swift Object storage?
- 10.4 What is the level of compatibility between S3 API in IBM Storage Scale and Amazon S3?
- 10.5 How can I ensure secure data in flight between object client and IBM Storage Scale Swift Object storage?

- 10.6 What authentication schemes does IBM Storage Scale Swift Object Storage support?
- 10.7 When using Unified File and Object Access feature, when should I use the local_mode vs unified_mode for identity management?
- 10.8 Can I have SMB/NFS export over object data when using Unified File and Object Access feature?
- 10.9 Can I have read/write access from file interface as well as Swift Object interface on the same data, when using Unified File and Object Access feature?
- 10.10 What are the limitations for IBM Storage Scale Swift Object Storage with Unified File Access enabled?
- 10.11 What are the considerations when using HAProxy load balancer with IBM Storage Scale Swift Object Storage?
- 10.12 Are there known issues with OpenStack software?
- 10.13 Are there any additional limitations or restrictions when using IBM Storage Scale compression or encryption with Swift Object protocol?
- 10.14 What are the limitations when configuring objects to use CES Groups?
- 10.15 Does the Swift Object protocol support immutable objects?
- 10.16 Are OpenStack package repositories required to install the Swift Object protocol?
- 10.17 Can additional protocol nodes be added when the Swift Object protocol is in toleration mode?

11. CES S3 protocol support questions:

- 11.1 What are the considerations for using the CES S3 protocol with IBM Storage Scale?
- 11.2 What are the currently available features of CES S3 for IBM Storage Scale?
- 11.3 How can I ensure secure data in flight between a CES S3 client and IBM Storage Scale CES S3 protocol stack?
- 11.4 What are the additional limitations or restrictions when using IBM Storage Scale compression or encryption with CES S3 protocol?
- 11.5 Does the CES S3 protocol support immutable objects?
- 11.6 What are the configuration considerations when deploying the CES S3 protocol functionality?
- 11.7 Can the CES S3 (technology preview) that was introduced with 5.2.0.x be upgraded to the CES S3 that is included with 5.2.1.x onward?

12. Installation Toolkit questions:

- 12.1 Can I upgrade an IBM Storage Scale cluster with protocols directly from 4.1.1.x to 4.2.0.x or 4.2.1.x?
- 12.2 What are the limitations when I use the Installation Toolkit to install IBM Storage Scale 4.2.0.0 or upgrade from IBM Storage Scale 4.1.1.x to 4.2.0.0 or 4.2.1.0 with Swift Object Protocol?
- 12.3 Can GUI nodes be added during an upgrade to IBM Storage Scale V4.2.0.1?
- 12.4 What are the limitations when I use the Installation Toolkit to upgrade IBM Storage Scale from 4.1.1 to 4.2 with NFS or SMB Protocol?
- 12.5 How can I determine if the Installation Toolkit successfully upgrades from IBM Storage Scale V4.1.1 to V4.2?
- 12.6 Can I have EPEL repos enabled when using the spectrumscale installation toolkit for install or upgrade?
- 12.7 What are the functions that are not supported by the installation toolkit?
- 12.8 What are the potential limitations when using the Installation Toolkit with SLES12 SP1 or SP2 nodes?

13. Transparent cloud tiering questions:

- 13.1 What platforms is transparent cloud tiering supported on?
- 13.2 How many transparent cloud tiering bridge nodes are supported?
- 13.3 What are the networking requirements for transparent cloud tiering?
- 13.4 What Cloud Object Storage Providers are supported?
- 13.5 Can transparent cloud tiering be used with other IBM Storage Scale functions?
- 13.6 Is transparent cloud tiering supported with IBM Storage Scale System?
- 13.7 Are there restrictions on viewing or listing files?
- 13.8 What levels of operating systems are supported by transparent cloud tiering services in IBM Storage Scale?

14. Licensing and Pricing questions:

- 14.1 Where can I find detailed information about IBM Storage Scale and IBM Storage Scale System licensing and pricing?
- 14.2 How is IBM Storage Scale licensed?
- 14.3 How can I determine the number of licenses that I need? How is capacity measured?
- 14.4 If I have IBM Storage Scale Advanced Edition or IBM Storage Scale Standard Edition, can I move to storage capacity licensing?
- 14.5 What are the differences between the IBM Storage Scale editions?
- 14.6 How can I tell which edition of IBM Storage Scale I am running?
- 14.7 Can I migrate my GPFS V3.5 licenses to IBM Storage Scale?
- 14.8 Can I transfer my IBM Storage Scale licenses between machines?
- 14.9 Do I require more IBM Storage Scale licenses for my Disaster Recovery (DR) environment?
- 14.10 What are IBM Storage Scale OEMs, and how are their products licensed?
- 14.11 Can I get support for IBM Storage Scale from IBM if it is supplied as part of an OEM solution?

- 14.12 If I install IBM Storage Scale software and licenses from IBM on an OEM system from another vendor, is that software supported by IBM?
- 14.13 Can I mix products licensed from other vendors that embed IBM Storage Scale (OEMs) in the same cluster as IBM Storage Scale licenses from IBM?
- 14.14 Which platforms support IBM Storage Scale Backup?
- 14.15 What is the benefit of using IBM Storage Scale Backup?
- 14.16 Who are the potential users of IBM Storage Scale Backup?
- 14.17 Which IBM Storage Scale editions support IBM Storage Scale Backup?
- 14.18 May I use IBM Storage Scale Backup with existing deployments of IBM Storage Scale, IBM Storage Scale System, and IBM Storage Protect?
- 14.19 What is the benefit of using IBM Storage Fusion Data Cataloging Services?
- 14.20 Who are the potential users of IBM Storage Fusion Data Cataloging Services?
- 14.21 IBM Storage Scale editions support IBM Storage Fusion Data Cataloging Services?

15. Service questions:

- 15.1 What support services are available for IBM Storage Scale?
- 15.2 How do I download fixes for IBM Storage Scale?
- 15.3 What are the current advisories for all platforms supported by IBM Storage Scale?
- 15.4 What are the current advisories for IBM Storage Scale on AIX?
- 15.5 What are the current advisories for IBM Storage Scale?
- 15.6 Where can I find the IBM Storage Scale Software License Agreement?
- 15.7 Where can I find End of Market (EOM) and End of Service (EOS) information?
- 15.8 Where can I locate IBM Storage Scale code to upgrade from my current level ?
- 15.9 What are the release practices for IBM Storage Scale?

16. Scale Management Server (REST API) questions:

- 16.1 Does IBM Storage Scale support RESTful APIS for file system management?
- 16.2 What are the changes made to the API implementation in the 4.2.3 release?

17. File audit logging and clustered watch folder questions:

- 17.1 What are the requirements and limitations for file audit logging and clustered watch folder?
- 17.2 Is the message queue still supported in IBM Storage Scale 5.1.1 and later?

18. Support of Docker, Podman containers and use in Kubernetes and OpenShift® environments questions:

- 18.1 Can IBM Storage Scale storage be utilized inside of containerized environments?
- 18.2 Why do we need to use IBM Storage Scale Container Storage Interface? Can we bind mount IBM Storage Scale directly into docker containers?
- 18.3 What are the prerequisites and compatibility matrix for using IBM Storage Scale as a backend for containers?
- 18.4: Supported Upgrade Paths for IBM Storage Scale CSI and IBM Storage Scale Container Native
- 18.5 Can Kubernetes perform a health check on the underlying IBM Storage Scale cluster?
- 18.6 How can one identify which filesets in IBM Storage Scale are created by the IBM Storage Scale CSI driver?
- 18.7 Can I use IBM Storage Scale System storage to provision volumes into my remote cluster and then act upon these volumes with IBM Storage Scale CSI?
- 18.8 Is it possible to contribute to the IBM Storage Scale CSI driver or operator projects?
- 18.9 Where do I get IBM Storage Scale CSI and IBM Storage Scale Container Native?
- 18.10 What Cloud Paks and services does IBM Storage Scale currently support?
- 18.11 Where can I get more information about the subjects discussed in this section?
- 18.12 Does IBM Storage Scale Container Native support RDMA?
- 18.13 Can IBM Storage Scale storage be used for overlay files systems?

19. IBM Storage Scale Erasure Code Edition questions:

- 19.1 What is IBM Storage Scale Erasure Code Edition, and why should I consider it?
- 19.2 What are the limitations of IBM Storage Scale Erasure Code Edition?
- 19.3 Can IBM Storage Scale Erasure Code Edition exist with IBM Elastic Storage® Server or IBM Elastic Storage System 3000 in the same cluster and support the same file system?
- 19.4.1 What are the minimum hardware requirements for IBM Storage Scale Erasure Code Edition?
- 19.4.2 Can I use any vendor's server with IBM Storage Scale Erasure Code Edition?
- 19.4.3 Can I use x86, ARM 64, IBM z, and PowerPC architectures with IBM Storage Scale Erasure Code Edition?
- 19.4.4 What operating systems are supported for IBM Storage Scale Erasure Code Edition storage servers?
- 19.4.5 Is CentOS supported with IBM Storage Scale Erasure Code Edition?

- 19.4.6 Can I use SATA drives with IBM Storage Scale Erasure Code Edition?
- 19.4.7 Can I use SED drives with IBM Storage Scale Erasure Code Edition?
- 19.4.8 Can I use external enclosures with IBM Storage Scale Erasure Code Edition?
- 19.4.9 Can I run IBM Storage Scale Erasure Code Edition with heterogeneous servers?
- 19.4.10 Can I use a virtual machine as a storage node with IBM Storage Scale Erasure Code Edition?
- 19.4.11 Can I deploy IBM Storage Scale Erasure Code Edition in cloud environments?
- 19.5.1 What is the minimum and maximum number of nodes required for an IBM Storage Scale Erasure Code Edition recovery group? What is the maximum number of IBM Storage Scale Erasure Code Edition storage nodes in an IBM Storage Scale cluster?
- 19.5.2 How can I estimate the usable space in one recovery group with IBM Storage Scale Erasure Code Edition storage nodes in an IBM Storage Scale cluster?
- 19.6.1 What are the network requirements and considerations for IBM Storage Scale Erasure Code Edition?
- 19.6.2 Is IPV6 supported with IBM Storage Scale Erasure Code Edition?
- 19.7 How can I get the IBM Storage Scale Erasure Code Edition hardware, IBM Storage Scale network, and IBM Storage Scale storage precheck tools, and how do I execute them?
- 19.8 Can I use the IBM Storage Scale installation toolkit to install and upgrade IBM Storage Scale Erasure Code Edition?
- 19.9 Can I use sudo wrappers with IBM Storage Scale Erasure Code Edition?
- 19.10.1 Can I run CES protocol software on IBM Storage Scale Erasure Code Edition storage nodes?
- 19.10.2 Can I run application workloads on IBM Storage Scale Erasure Code Edition storage nodes?
- 19.10.3 Can I run the IBM Storage Scale GUI, AFM gateways, and perfmon collectors on IBM Storage Scale Erasure Code Edition storage nodes?
- 19.11 What workloads is IBM Storage Scale Erasure Code Edition storage recommended for? Can it be used for high performance workloads?
- 19.12 How can I migrate data from an existing IBM Storage Scale cluster or file system to IBM Storage Scale Erasure Code Edition?
- 19.13 When installing IBM Storage Scale Erasure Code Edition, IBM udev rules are installed. Also, when I configure servers using `mmvdisk`, the initial IBM Storage Scale configuration values are set. Is it okay to change these rules and values?
- 19.14 Are there any extra requirements for using or configuring NVMe drives for use with IBM Storage Scale Erasure Code Edition?
- 19.15 Can I deploy IBM Storage Scale Erasure Code Edition on a configuration that does not meet the requirements and recommendations detailed in the IBM Storage Scale Erasure Code Edition documentation?
- 19.16 What high-speed Ethernet networking adapters are tested with IBM Storage Scale Erasure Code Edition?

20. IBM Storage Scale Developer Edition questions:

- 20.1 What functions do I get with IBM Storage Scale Developer Edition?
- 20.2 Can I use IBM Storage Scale Developer Edition in a production environment?
- 20.3 How do I get IBM product support for IBM Storage Scale Developer Edition?
- 20.4 For IBM Storage Scale Developer Edition, can I upgrade my cluster to an IBM-supported IBM Storage Scale product offering like IBM Storage Scale Data Management Edition?
- 20.5 With IBM Storage Scale Developer Edition, how can I see the licensed storage usage?
- 20.6 Is the IBM Storage Scale installation toolkit supported with IBM Storage Scale Developer Edition?
- 20.7 What are the supported operating systems and architectures for IBM Storage Scale Developer Edition?
- 20.8 Can I run more than one IBM Storage Scale Developer Edition cluster in the same company, division, test lab, etc.?
- 20.9 Where can I download IBM Storage Scale Developer Edition?

21. Integrated protocol server authentication questions:

- 21.1 Is IBM Storage Scale affected by the Microsoft advisory ADV190023 regarding LDAP channel binding and LDAP signing?
- 21.2 Which Microsoft Active Directory versions does IBM Storage Scale support?
- 21.3 Which LDAP versions does IBM Storage Scale support?
- 21.4 What is the impact on IBM Storage Scale Protocol when you are using AD with RFC2307 or migrating to Windows 2016 AD server or later?

22. Cloud and cloudkit questions:

- 22.1 Can IBM Storage Scale be installed on the cloud?
- 22.2 What is the recommended way to install IBM Storage Scale on cloud?
- 22.3 What are the public clouds sported by cloudkit?
- 22.4 What IBM Storage Scale edition packages are supported by the cloudkit?
- 22.5 If I manually deploy IBM Storage Scale on the cloud without using cloudkit, are there any restrictions, functions, or environments that are not recommended for use?
- 22.6 What are the permissions needed to deploy cloudkit?
- 22.7 What IBM Storage Scale features are supported by cloudkit?
- 22.8 What operating systems are supported by cloudkit?
- 22.9 How to access the cluster that was created by cloudkit?

- 22.10 What are the limitation when using cloudkit?
- 22.11 What should one be aware of before deleting an IBM Storage Scale cluster?
- 22.12 Who should I contact for support?
- 22.13 How to collect debug data for cloudkit?
- 22.14 Why the file system capacity is relatively smaller or bigger than the exact size provided as input?

General Questions

Q1.1: Where can I find ordering information for IBM Storage Scale?

A1.1: You can view ordering information for IBM Storage Scale in the Announcement Letters at [IBM Announcements](#).

1. In the **keyword, letter or product number** bar, search for IBM Storage Scale in title, or search for <one of the product numbers> in product or document. Choose **Announcement letters, Software**.
2. For IBM Storage Scale V5, choose the corresponding product number to enter in the **Search for** field:
 - IBM Storage Scale Data Access Edition: **5737-I39** (Passport Advantage®), **5641-DA1**, **5641-DA3**, **5641-DA5** (eConfig/AAS)
 - IBM Storage Scale Data Management Edition: **5737-F34** (Passport Advantage), **5641-DM1**, **5641-DM3**, **5641-DM5** (eConfig/AAS)
 - IBM Storage Scale Erasure Code Edition: **5737-J34** (Passport Advantage)

Q1.2: Where can I find the documentation for IBM Storage Scale?

A1.2: Documentation is available on the IBM Knowledge Center.

- For IBM Storage Scale V4.1.1 or later, at http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html

Note: To change the documentation version, click the **Select** drop-down menu.

- For GPFS V3.5 and V4.1, at http://www-01.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html

Note: To change the documentation version, click the **Select** drop-down menu.

Q1.3: What is there beyond the standard documentation that can help me learn more about and use IBM Storage Scale?

A1.3: Additional resources include:

- The paper *On Making GPFS Truly General* at http://researcher.watson.ibm.com/researcher/files/us-dhildeb/login_hildebrand.pdf
- GPFS Technotes:
 - [GPFS Internal and External mount states](#)
- GPFS Web pages:
 - The [IBM Support Portal](#) (log in with your IBM ID)
 - The [IBM Storage Scale documentation site](#)
 - The GPFS Support Portal at [IBM Support > My Notifications > Product notifications](#)
- The IBM Systems Magazine site at <http://www.ibmsystemsmag.com/> and search on IBM Storage Scale and GPFS.
- The IBM Redbooks® and Redpapers site at www.redbooks.ibm.com and search on IBM Storage Scale and GPFS.
- The [IBM Technical Sales Library](#)
- For IBM Storage Scale for Linux on Z, see the white paper [Getting started with IBM Spectrum Scale for Linux on Z](#).
- Classes:

Table 13: Classes		
Course Code	Course Title	Course Type
H005G	IBM Storage Scale Basic Administration for Linux and AIX	System Administration classroom

IBM training provides education to support many IBM offerings. Descriptions of courses for IT professionals and managers are on the IBM training website <http://www.ibm.com/services/learning/>

Q1.4:
How can I ask a more specific question about IBM Storage Scale?

A1.4:
Depending upon the nature of your question, you may ask it in one of several ways.

- If you want to correspond with IBM regarding IBM Storage Scale:
 - If your question concerns a potential software error in IBM Storage Scale and you have an IBM software maintenance contract, please contact 1-800-IBM-SERV in the United States or your [local IBM Service Center in other countries](#).
 - If you have a question that can benefit other IBM Storage Scale users, you may post it to the GPFS technical discussion forum at [IBM Storage Community](#).
 - This FAQ is continually being enhanced. To contribute possible questions or answers, please send them to scale@us.ibm.com.
- If you want to interact with other GPFS users, refer to the [Storage Scale User Group](#). If you are interested in subscribing to the mailing list that is maintained by this group, fill in the [form](#).
- If you want to submit or view Request for Enhancements for IBM Storage Scale, go to [IBM System Storage Ideas Portal](#).

If your question does not fall into the above categories, you can send a note directly to the IBM Storage Scale development team at scale@us.ibm.com. However, this mailing list is informally monitored as time permits and should not be used for priority messages to the IBM Storage Scale team.

Q1.5:
Does IBM Storage Scale participate in the IBM Academic Initiative Program?

A1.5:
IBM Storage Scale Developer Edition is included in the Academic Initiative Program. Work with your IBM client representative to determine what educational discount may be available for IBM Storage Scale. See www.ibm.com/planetwide/index.html

Q1.6:
Is IBM Storage Scale available in IBM PartnerWorld?

A1.6:
IBM Storage Scale is available in IBM PartnerWorld. Search for "IBM Storage Scale", "General Parallel File System", or "GPFS" in the Software Access catalog <https://www-304.ibm.com/jct01004c/partnerworld/partnertools/eorderweb/ordersw.do>

Q1.7:
Does IBM Storage Scale have a trial program?

A1.7:
For trials:

- A free 90-day trial program is available to prospective customers of IBM Storage Scale. Contact your IBM sales representative to apply for this trial evaluation.
- The IBM Storage Scale Developer Edition is also available. For more information see, [IBM Storage Scale Developer Edition questions](#) on page 102.

ISVs can obtain the software and licenses from the IBM Partner World at <https://www-356.ibm.com/partnerworld/wps/servlet/ContentHandler/isv/sac>.

Q1.8:
Where can I find the documentation for IBM Storage Scale RAID?

A1.8:
The documentation for IBM Storage Scale RAID can be found:

- The IBM Storage Scale Erasure Code Edition documentation at <https://www.ibm.com/docs/en/spectrum-scale-ec>.
- The IBM Elastic Storage Server documentation <https://www.ibm.com/docs/en/ess-p8>.
- The FAQ at <http://www-01.ibm.com/support/knowledgecenter/api/content/nl/en-us/SSYSP8/gnfaq.html>.

Q1.9:
Where can I find detailed information about stabilized, deprecated, and discontinued features of IBM Storage Scale?

A1.9:
For information about the stabilized, deprecated, and discontinued features of IBM Storage Scale, see [Stabilized, deprecated, and discontinued features in IBM Storage Scale](#).

Q1.10: Where can I find a list of the authorized program analysis reports (APARs) that are resolved for IBM Storage Scale?

A1.10: A list of the resolved APARs for IBM Storage Scale 5.0.5.x and later releases is available on the following support page: [IBM Storage Scale APARs Resolved](#).

Software questions

Q2.1: What is supported on IBM Storage Scale for AIX, Linux, and Windows?

A2.1: **Supported OS and software versions**

The following tables detail the operating system and software package versions that are supported on IBM Storage Scale. Only the most recent versions of IBM Storage Scale releases are listed.

If you need information about earlier versions, see [Archived support information for IBM Storage Scale](#).

Interactive version: A dynamic table that lists the detailed support information for IBM Storage Scale releases is available at: <https://public.dhe.ibm.com/storage/spectrumscale/support/>.

This dynamic tool is still under development. The FAQ is the primary source of updated information.

IBM Storage Scale is not supported on OS versions that are out of standard or general support by the OS vendor. Use the following tables for the supported versions:

Important: Kernel errata can be applied to the current kernel version if the kernel is running on a supported operating system and IBM Storage Scale version. Before rolling out to production, you should always validate in a test environment any kernel changes that include errata with IBM Storage Scale. You must rebuild the portability layer after any kernel changes.

Table 14: IBM Storage Scale 6.0.0.2: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	RHEL ¹² 10 10.1	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP7
x86_64	RHEL 8.10: 4.18.0-553.117.1.el8_10 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 9.4 5.14.0-427.116.1.el9_4 RHEL 9.6 5.14.0-570.62.1.el9_6 RHEL 9.7 5.14.0-611.47.1.el9_7 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 10 6.12.0-55.66.1.el10_0.x86_64 RHEL 10.1 6.12.0-124.11.1.el10_1.x86_64 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	Ubuntu 22.04: 5.15.0-27-generic Ubuntu 22.04.4: 5.15.0-122-generic Ubuntu 22.04.5: 5.15.0-174-generic SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	Ubuntu 24.04.1: 6.8.0-51-generic Ubuntu 24.04.2: 6.8.0-54-generic Ubuntu 24.04.3 6.8.0-71-generic SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	SLES 15 SP7: 6.4.0-150700.53.34.1 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	RHEL ¹² 10 10.1	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP7
Power LE	RHEL 8.10: 4.18.0-553.115.1.el8_10 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 9.4 5.14.0-427.13.1 RHEL 9.6 5.14.0-570.12.1 RHEL 9.7 5.14.0-611.8.1 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 10 6.12.0-124.8.1.el10_0.ppc64le RHEL 10.1 6.12.0-124.8.1.el10_1.ppc64le SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)		No plan to support.	SLES 15 SP7: 6.4.0-150700.51-default
Linux on Z ³	RHEL 8.10: 4.18.0-553.117.1.el8_10 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 9.4 5.14.0-427.117.1.el9_4 RHEL 9.6 5.14.0-570.62.1.el9_6 RHEL 9.7 5.14.0-611.8.1 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 10 6.12.0-55.9.1 RHEL 10.1 6.12.0-124.8.1.el10_1.s390x SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	No plan to support.	No plan to support.	SLES 15 SP7: 6.4.0-150700.51-default SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)
ARM 64 ^{9, 10}	No plan to support.	RHEL 9.4 5.14.0-427.16.1.el9_4.aarch64 RHEL 9.6 5.14.0-570.13.1 RHEL 9.7 5.14.0-611.7.1	RHEL 10 6.12.0-55.22.1 RHEL 10.1 6.12.0-124.9.1.el10_1.aarch64	Ubuntu 22.04.x 5.15.0-97-generic linux-signatures-nvidia-6.2.0-1015-nvidia linux-signatures-nvidia-6.2.0-1015-nvidia-64k	Ubuntu 24.04.3 6.8.0-85-generic SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	No plan to support.

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 2004, OS Build 19041)	Windows 11 (>= OS Build 22000)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 2			
	TL 3			
	TL 4			

1. Red Hat® live kernel patching is not supported on IBM Storage Scale.

2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssh-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.
9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:
 - 4k and 64k for RHEL 9
 - Ubuntu 4k or 64k for Ubuntu 22.04.x
 - Nvidia 4k or 64k for Ubuntu 22.04.x
10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.
11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.
12. Cluster Export Services (CES) S3 protocol is supported on Red Hat Enterprise Linux (RHEL) 10.0 and RHEL 10.1 starting with IBM Storage Scale 6.0.0.2.

Table 16: IBM Storage Scale 6.0.0.1: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	RHEL ¹² 10 10.1	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
x86_64	RHEL 8.10: 4.18.0-553.105.1.el8_10 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 9.4 5.14.0-427.111.1.el9_4 RHEL 9.6 5.14.0-570.62.1.el9_6 RHEL 9.7 5.14.0-611.34.1.el9_7 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 10 6.12.0-55.60.1.el10_0.x86_64 RHEL 10.1 6.12.0-124.11.1.el10_1.x86_64 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	Ubuntu 22.04: 5.15.0-27-generic Ubuntu 22.04.4: 5.15.0-122-generic Ubuntu 22.04.5: 5.15.0-164-generic SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	Ubuntu 24.04.1: 6.8.0-51-generic Ubuntu 24.04.2: 6.8.0-54-generic Ubuntu 24.04.3 6.8.0-71-generic SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	SLES 15 SP6: 6.4.0-150600.23.65.1 SLES 15 SP7: 6.4.0-150700.51-default SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)
Power LE	RHEL 8.10: 4.18.0-553.109.1.el8_10 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 9.4 5.14.0-427.13.1 RHEL 9.6 5.14.0-570.12.1 RHEL 9.7 5.14.0-611.8.1 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 10 6.12.0-124.8.1.el10_0.ppc64le RHEL 10.1 6.12.0-124.8.1.el10_1.ppc64le SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)		No plan to support.	SLES 15 SP6: 6.4.0-150600.21-default SLES 15 SP7: 6.4.0-150700.51-default

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	RHEL ¹² 10 10.1	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
Linux on Z ³	RHEL 8.10: 4.18.0-553.109.1.el8_10 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 9.4 5.14.0-427.111.1.el9_4 RHEL 9.6 5.14.0-570.62.1.el9_6 RHEL 9.7 5.14.0-611.8.1 SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	RHEL 10 6.12.0-55.9.1 RHEL 10.1 6.12.0-124.8.1.el10_1.s390x SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	No plan to support.	No plan to support.	SLES 15 SP6: 6.4.0-150600.23.73.1 SLES 15 SP7: 6.4.0-150700.51-default SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)
ARM 64 ^{9, 10}	No plan to support.	RHEL 9.4 5.14.0-427.16.1.el9_4.aarch64 RHEL 9.6 5.14.0-570.13.1 RHEL 9.7 5.14.0-611.7.1	RHEL 10 6.12.0-55.22.1 RHEL 10.1 6.12.0-124.9.1.el10_1.aarch64	Ubuntu 22.04.x 5.15.0-97-generic linux-signatures-nvidia-6.2.0-1015-nvidia linux-signatures-nvidia-6.2.0-1015-nvidia-64k	Ubuntu 24.04.3 6.8.0-85-generic SMB: 4.22.4 6-1 NFS: 7.0 (ibm 001.18)	No plan to support.

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 2004, OS Build 19041)	Windows 11 (>= OS Build 22000)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 1			
	TL 2			
	TL 3			
	TL 4			

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmdaquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssh-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is

not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.

9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:

- 4k and 64k for RHEL 9
- Ubuntu 4k or 64k for Ubuntu 22.04.x
- Nvidia 4k or 64k for Ubuntu 22.04.x

10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.

11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

12. Cluster Export Services (CES) S3 protocol is not supported on Red Hat Enterprise on Linux (RHEL) 10 and 10.1 for IBM Storage Scale 6.0.0.1.

Table 18: IBM Storage Scale 6.0.0: Tested software versions and latest tested kernels for Linux ²					
Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
x86_64	<p>RHEL 8.10: 4.18.0-553.117.1.el8_10</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	<p>RHEL 9.4 5.14.0-427.116.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.62.1.el9_6</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	<p>Ubuntu 22.04: 5.15.0-27-generic</p> <p>Ubuntu 22.04.4: 5.15.0-122-generic</p> <p>Ubuntu 22.04.5: 5.15.0-173-generic</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	<p>Ubuntu 24.04.1: 6.8.0-51-generic</p> <p>Ubuntu 24.04.2: 6.8.0-54-generic</p> <p>Ubuntu 24.04.3 6.8.0-71-generic</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	<p>SLES 15 SP6: 6.4.0-150600.23.65.1</p> <p>SLES 15 SP7: 6.4.0-150700.53.34.1</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>
Power LE	<p>RHEL 8.10: 4.18.0-553.117.1.el8_10</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	<p>RHEL 9.4 5.14.0-427.13.1</p> <p>RHEL 9.6 5.14.0-570.12.1</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>		No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.21-default</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p>
Linux on Z ³	<p>RHEL 8.10: 4.18.0-553.115.1.el8_10</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	<p>RHEL 9.4 5.14.0-427.111.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.62.1.el9_6</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>	No plan to support.	No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.23.73.1</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.22.4 5-2</p> <p>NFS: 7.0 (ibm 001.15)</p>

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
ARM 64 ^{9, 10}	No plan to support.	RHEL 9.4 5.14.0-427.16.1.el9 _4.aarch64 RHEL 9.6 5.14.0-570.13.1	Ubuntu 22.04.x 5.15.0-97-generic linux-signatures- nvidia-6.2.0-1015- nvidia linux-signatures- nvidia-6.2.0-1015- nvidia-64k	Ubuntu 24.04.3 6.8.0-85-generic SMB: 4.22.4 5-2 NFS: 7.0 (ibm 001.15)	No plan to support.

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 2004, OS Build 19041)	Windows 11 (>= OS Build 22000)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 1			
	TL 2			
	TL 3			

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssh-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.
9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:
 - 4k and 64k for RHEL 9
 - Ubuntu 4k or 64k for Ubuntu 22.04.x
 - Nvidia 4k or 64k for Ubuntu 22.04.x
10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.
11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

Table 20: IBM Storage Scale 5.2.3.7: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.2 24.04.3 24.04.4	SLES 15 SP7
x86_64	<p>RHEL 8.10: 4.18.0-553.115.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	<p>RHEL 9.4 5.14.0-427.116.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.62.1.el9_6</p> <p>RHEL 9.7 5.14.0-611.45.1.el9_7</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	<p>Ubuntu 22.04: 5.15.0-27-generic</p> <p>Ubuntu 22.04.4: 5.15.0-122-generic</p> <p>Ubuntu 22.04.5: 5.15.0-174-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	<p>Ubuntu 24.04.2: 6.8.0-54-generic</p> <p>Ubuntu 24.04.3 6.8.0-84-generic</p> <p>Ubuntu 24.04.4 6.8.0-106-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	<p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>
Power LE	<p>RHEL 8.10: 4.18.0-553.115.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	<p>RHEL 9.4 5.14.0-427.13.1</p> <p>RHEL 9.6 5.14.0-570.12.1</p> <p>RHEL 9.7 5.14.0-611.8.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>		No plan to support.	<p>SLES 15 SP7: 6.4.0-150700.51-default</p>
Linux on Z ³	<p>RHEL 8.10: 4.18.0-553.117.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	<p>RHEL 9.4 5.14.0-427.117.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>RHEL 9.7 5.14.0-611.8.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>	No plan to support.	No plan to support.	<p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.21)</p>

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.2 24.04.3 24.04.4	SLES 15 SP7
ARM 64 ^{9, 10}	No plan to support.	RHEL 9.4 5.14.0-427.16.1.el9_4.aarch64 RHEL 9.6 5.14.0-570.13.1 RHEL 9.7 5.14.0-611.7.1	Ubuntu 22.04.x 5.15.0-97-generic linux-signatures-nvidia-6.2.0-1015-nvidia linux-signatures-nvidia-6.2.0-1015-nvidia-64k	Ubuntu 24.04.3 6.8.0-85-generic SMB: 4.20.8 10-1 NFS: 5.7 (ibm 028.21)	No plan to support.

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 1809, OS Build 17763)	Windows 11 (>= OS Build 22000)	Windows Server 2019 (>= version 1809, OS Build 17763)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 2				
	TL 3				
	TL 4				

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssl-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.
9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:
 - 4k and 64k for RHEL 9
 - Ubuntu 4k or 64k for Ubuntu 22.04.x
 - Nvidia 4k or 64k for Ubuntu 22.04.x
10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.
11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

Table 22: IBM Storage Scale 5.2.3.6: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP7
x86_64	<p>RHEL 8.10: 4.18.0-553.117.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	<p>RHEL 9.4 5.14.0-427.116.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.62.1.el9_6</p> <p>RHEL 9.7 5.14.0-611.47.1.el9_7</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	<p>Ubuntu 22.04: 5.15.0-27-generic</p> <p>Ubuntu 22.04.4: 5.15.0-122-generic</p> <p>Ubuntu 22.04.5: 5.15.0-171-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	<p>Ubuntu 24.04.1: 6.8.0-51-generic</p> <p>Ubuntu 24.04.2: 6.8.0-54-generic</p> <p>Ubuntu 24.04.3 6.8.0-84-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	<p>SLES 15 SP7: 6.4.0-150700.53.34.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>
Power LE	<p>RHEL 8.10: 4.18.0-553.117.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	<p>RHEL 9.4 5.14.0-427.13.1</p> <p>RHEL 9.6 5.14.0-570.12.1</p> <p>RHEL 9.7 5.14.0-611.8.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>		No plan to support.	<p>SLES 15 SP7: 6.4.0-150700.51-default</p>
Linux on Z ³	<p>RHEL 8.10: 4.18.0-553.115.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	<p>RHEL 9.4 5.14.0-427.116.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>RHEL 9.7 5.14.0-611.8.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>	No plan to support.	No plan to support.	<p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.20)</p>

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP7
ARM 64 ^{9, 10}	No plan to support.	RHEL 9.4 5.14.0-427.16.1.el9_4.aarch64 RHEL 9.6 5.14.0-570.13.1 RHEL 9.7 5.14.0-611.7.1	Ubuntu 22.04.x 5.15.0-97-generic linux-signatures-nvidia-6.2.0-1015-nvidia linux-signatures-nvidia-6.2.0-1015-nvidia-64k	Ubuntu 24.04.3 6.8.0-85-generic SMB: 4.20.8 10-1 NFS: 5.7 (ibm 028.20)	No plan to support.

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 1809, OS Build 17763)	Windows 11 (>= OS Build 22000)	Windows Server 2019 (>= version 1809, OS Build 17763)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 2				
	TL 3				
	TL 4				

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssl-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.
9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:
 - 4k and 64k for RHEL 9
 - Ubuntu 4k or 64k for Ubuntu 22.04.x
 - Nvidia 4k or 64k for Ubuntu 22.04.x
10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.
11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

Table 24: IBM Storage Scale 5.2.3.5: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
x86_64	<p>RHEL 8.10: 4.18.0-553.104.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>RHEL 9.4 5.14.0-427.109.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>RHEL 9.7 5.14.0-611.30.1.el9_7</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>Ubuntu 22.04: 5.15.0-27-generic</p> <p>Ubuntu 22.04.4: 5.15.0-122-generic</p> <p>Ubuntu 22.04.5: 5.15.0-164-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>Ubuntu 24.04.1: 6.8.0-51-generic</p> <p>Ubuntu 24.04.2: 6.8.0-54-generic</p> <p>Ubuntu 24.04.3 6.8.0-84-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>SLES 15 SP6: 6.4.0-150600.23.81.3</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>
Power LE	<p>RHEL 8.10: 4.18.0-553.104.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>RHEL 9.4 5.14.0-427.13.1</p> <p>RHEL 9.6 5.14.0-570.12.1</p> <p>RHEL 9.7 5.14.0-611.8.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>		No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.21-default</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p>
Linux on Z ³	<p>RHEL 8.10: 4.18.0-553.100.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>RHEL 9.4 5.14.0-427.106.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>RHEL 9.7 5.14.0-611.8.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	No plan to support.	No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.23.81.3</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.19)</p>

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6 9.7	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
ARM 64 ^{9, 10}	No plan to support.	RHEL 9.4 5.14.0-427.16.1.el9_4.aarch64 RHEL 9.6 5.14.0-570.13.1 RHEL 9.7 5.14.0-611.7.1	Ubuntu 22.04.x 5.15.0-97-generic linux-signatures-nvidia-6.2.0-1015-nvidia linux-signatures-nvidia-6.2.0-1015-nvidia-64k	Ubuntu 24.04.3 6.8.0-85-generic SMB: 4.20.8 10-1 NFS: 5.7 (ibm 028.19)	No plan to support.

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 1809, OS Build 17763)	Windows 11 (>= OS Build 22000)	Windows Server 2019 (>= version 1809, OS Build 17763)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 1				
	TL 2				
	TL 3				
	TL 4				

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssh-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.
9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:
 - 4k and 64k for RHEL 9
 - Ubuntu 4k or 64k for Ubuntu 22.04.x
 - Nvidia 4k or 64k for Ubuntu 22.04.x
10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.
11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

Table 26: IBM Storage Scale 5.2.3.4: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5} 9.4 9.6	Ubuntu ^{6, 7, 8} 22.04 22.04.4 22.04.5	Ubuntu ^{6, 7, 8} 24.04.1 24.04.2 24.04.3	SLES 15 SP6 SP7
x86_64	<p>RHEL 8.10: 4.18.0-553.97.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	<p>RHEL 9.4 5.14.0-427.100.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	<p>Ubuntu 22.04: 5.15.0-27-generic</p> <p>Ubuntu 22.04.4: 5.15.0-122-generic</p> <p>Ubuntu 22.04.5: 5.15.0-157-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	<p>Ubuntu 24.04.1: 6.8.0-51-generic</p> <p>Ubuntu 24.04.2: 6.8.0-54-generic</p> <p>Ubuntu 24.04.3 6.8.0-84-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	<p>SLES 15 SP6: 6.4.0-150600.23.65.1</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>
Power LE	<p>RHEL 8.10: 4.18.0-553.84.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	<p>RHEL 9.4 5.14.0-427.13.1</p> <p>RHEL 9.6 5.14.0-570.12.1</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>		No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.21-default</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p>
Linux on Z ³	<p>RHEL 8.10: 4.18.0-553.84.1.el8_10</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	<p>RHEL 9.4 5.14.0-427.94.1.el9_4</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	No plan to support.	No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.23.73.1</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>
ARM 64 ^{9, 10}	No plan to support.	<p>RHEL 9.4 5.14.0-427.16.1.el9_4.aarch64</p> <p>RHEL 9.6 5.14.0-570.13.1</p>	<p>Ubuntu 22.04.x 5.15.0-97-generic</p> <p>linux-signatures-nvidia-6.2.0-1015-nvidia</p> <p>linux-signatures-nvidia-6.2.0-1015-nvidia-64k</p>	<p>Ubuntu 24.04.3 6.8.0-85-generic</p> <p>SMB: 4.20.8 10-1</p> <p>NFS: 5.7 (ibm 028.18)</p>	No plan to support.

Table 27: IBM Storage Scale 5.2.3.4: Tested software versions and latest tested kernels for Windows and AIX

AIX 7.2 ¹¹	AIX 7.3 ¹¹	Windows 10 (>= version 1809, OS Build 17763)	Windows 11 (>= OS Build 22000)	Windows Server 2019 (>= version 1809, OS Build 17763)	Windows Server 2022 (>= OS Build 20348)
TL 5	TL 1				
	TL 2				
	TL 3				

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. For RHEL 9, openssh-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
5. NIS authentication is not supported in RHEL 9.
6. Only the Ubuntu Server edition is supported.
7. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
8. IBM Storage Scale packages are missing a dependency that ensures that the iputils-arping package is installed. Ubuntu users must manually install iputils-arping to make CES ip-failover work correctly. If the iputils-arping package is not installed manually, IBM Storage Scale packages are installed without error, but the CES ip-failover function is affected. CES protocol clients may not pick up the CES IP configuration after the CES IP move for many minutes.
9. The following page size kernels are supported for Advanced RISC Machine (ARM) architectures:
 - 4k and 64k for RHEL 9
 - Ubuntu 4k or 64k for Ubuntu 22.04.x
 - Nvidia 4k or 64k for Ubuntu 22.04.x
10. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.
11. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

Important: As of 30 September 2025 IBM Storage Scale 5.1.9 reached end of support (EOS). For more information about EOS policy and dates, see [15.7 Where can I find End of Market \(EOM\) and end of service \(EOS\) information?](#) For more information about EOS versions, see *Functional support matrices and feature exceptions* in [Archived support information for IBM Storage Scale](#).

Table 28: IBM Storage Scale 5.1.9.12: Tested software versions and latest tested kernels for Linux²

Architecture	RHEL ¹ 8.10	RHEL ^{1, 4, 5, 6} 9.4 9.5 9.6	Ubuntu 20.04 ^{7, 8} 20.04.5 20.04.6	Ubuntu 22.04 ^{7, 8} 22.04.4 22.04.5	SLES 15 SP6 SP7
x86_64	<p>RHEL 8.10: 4.18.0-553.80.1.el8_10</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>RHEL 9.4 5.14.0-427.94.1.el9_4</p> <p>RHEL 9.5 5.14.0-503.40.1.el9_5</p> <p>RHEL 9.6 5.14.0-570.58.1.el9_6</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>Ubuntu 20.04.5: 5.4.0-139-generic</p> <p>Ubuntu 20.04.6: 5.4.0-216-generic</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>Ubuntu 22.04.4: 5.15.0-118-generic</p> <p>Ubuntu 22.04.5: 5.15.0-157-generic</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>SLES 15 SP6: 6.4.0-150600.23.65.1</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>
Power LE	<p>RHEL 8.10: 4.18.0-553.80.1.el8_10</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>RHEL 9.4 5.14.0-427.13.1</p> <p>RHEL 9.5 5.14.0-503.11.1</p> <p>RHEL 9.6 5.14.0-570.12.1</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>Ubuntu 20.04.5: 5.4.0-146-generic</p> <p>Ubuntu 20.04.6: 5.4.0-216-generic</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>		<p>SLES 15 SP6: 6.4.0-150600.21-default</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p>
Linux on Z ³	<p>RHEL 8.10: 4.18.0-553.80.1.el8_10</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	<p>RHEL 9.4 5.14.0-427.94.1.el9_4</p> <p>RHEL 9.5 5.14.0-503.11.1</p> <p>RHEL 9.6 5.14.0-570.55.1.el9_6</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>	No plan to support.	No plan to support.	<p>SLES 15 SP6: 6.4.0-150600.23.73.1</p> <p>SLES 15 SP7: 6.4.0-150700.51-default</p> <p>SMB: 4.17.12 11-1</p> <p>NFS: 5.7 (ibm 028.19)</p>

Table 29: IBM Storage Scale 5.1.9.12: Tested software versions and latest tested kernels for Windows and AIX

AIX 7.2 ⁹	AIX 7.3 ⁹	Windows 10 (>= version 1809, OS Build 17763)	Windows 11 (>= OS Build 22000)	Windows Server 2019 (>= version 1809, OS Build 17763)	Windows Server 2022 (>= OS Build 20348)
TL 4	TL 1				
TL 5	TL 2				
	TL 3				

1. Red Hat live kernel patching is not supported on IBM Storage Scale.
2. If protocols and authentication are enabled, python3-ldap needs to be installed for mmadquery to run.
3. With IBM Storage Scale 5.1.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
4. Leap upgrade from RHEL 8 to 9 with IBM Storage Scale installed is not supported.
5. For RHEL 9, openssh-server-8.7p1-10.el9 or later is required ([RHBA-2022:6598 - Bug Fix Advisory](#)).
6. NIS authentication is not supported in RHEL 9.
7. Only the Ubuntu Server edition is supported.
8. Only the default, non-rolling generic kernel image for the Ubuntu Server edition is supported. Ubuntu Hardware Enablement (Ubuntu HWE) kernels and others are not supported with IBM Storage Scale.
9. Dynatrace cannot be used on IBM Storage Scale on AIX due to built-in memory tracking. Unstable behavior including crashes will occur.

IBM Storage Scale 5.1x through 5.2.x: Feature support exceptions

The following tables contain only the features that are not supported on AIX, Linux, Power, and Windows for the most recent versions of IBM Storage Scale. When an exception applies only to a specific release or PTF, it is listed in this [note](#).

To determine which edition of IBM Storage Scale includes specific features, see [14.1 Where can I find detailed information about IBM Storage Scale and IBM Storage Scale System licensing and pricing?](#)

Architecture	RHEL 8.x	RHEL 9.x	Ubuntu 20.04.x	Ubuntu 22.04.x	Ubuntu 24.04.x	SLES 15 SPx
x86_64	N/A	Authentication NIS method restricted and object protocol	ECE, HDFS, object protocol, and S3	ECE, HDFS, object protocol, and S3	ECE, HDFS, object protocol, and S3	ECE, HDFS, object protocol, and S3
Power LE ¹	ECE	ECE, Authentication NIS method restricted and object protocol	ECE, HDFS, installation toolkit, object protocol, S3, and IBM Storage Scale native REST API cannot run on Power LE nodes.	ECE, HDFS, installation toolkit, object protocol, SMB, NFS, S3, and IBM Storage Scale native REST API cannot run on Power LE nodes.	No plan to support.	ECE, HDFS, installation toolkit, SMB, NFS, S3, and clustered watch folder
Linux on Z ²	ECE, FPO, HDFS, object protocol, and S3	ECE, Authentication NIS method restricted and object protocol	No plan to support.	No plan to support.	No plan to support.	ECE, FPO, HDFS, object protocol, and S3
ARM 64 ³	No plan to support.	Protocol servers (SMB, NFS, CNFS, Object, and S3), NSD server, GNR, HSM, signed kernels, GUI, and IBM Storage Scale native REST API cannot run on ARM nodes.	No plan to support.	Protocol servers (SMB, NFS, CNFS, Object, and S3), NSD server, GNR, HSM, signed kernels, GUI, and IBM Storage Scale native REST API cannot run on ARM nodes.	No plan to support.	No plan to support.

Feature exceptions that apply to a specific PTF or release:

1. Performance monitoring ZIMON packages (gpfs.gss.pmsensors and gpfs.gss.pmcollector) are available from IBM Storage Scale 5.1.2.3 and later for SLES 15 with ppc64LE.
2. With IBM Storage Scale 5.1.x and 5.2.x, NFS and SMB protocols can be served from RHEL or SLES 15 on Linux on System Z servers. An RPQ would be required for IBM to review any requests for Integrated Protocol Server support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.
3. Starting with 5.2.3.2, ARM 64 can be used with ECE. In 5.2.0.x through 5.2.3.1, ARM 64 *cannot* be used with ECE.

Table 31: IBM Storage Scale - Feature exceptions for AIX and Windows

AIX	Windows
<ul style="list-style-type: none"> • AFM • AFM-DR • Call home • Clustered NFS • Clustered watch folder • ECE • File audit logging • FPO • GUI • HDFS • IBM Storage Scale Container Native deployment • IBM Storage Scale CSI Driver • Installation toolkit • Integrated protocols (CES) • Performance monitoring • RDMA • REST API • Transparent cloud tiering 	<ul style="list-style-type: none"> • AFM • AFM-DR • Backup of a file system • Call home • Clustered NFS • Clustered watch folder • ECE • Encryption of data at rest • File audit logging • File clones • FPO • GUI • HDFS • Health monitoring • Installation toolkit • Integrated protocols (CES) • Performance monitoring • Quotas • RDMA • REST API • Transparent cloud tiering • For more information, see <i>GPFS limitations in Windows</i> in the <i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i> and 2.6 What are the limitations of IBM Storage Scale support for Windows? in the IBM Storage Scale FAQ.

Q2.2:

What is the IBM Storage Scale support position regarding *clone* Linux distributions (CentOS, Rocky Linux, Oracle Linux RHCK, ROCKS, White box Linux, etc.) ?

A2.2:

There are many Linux distributions out there. It would not be practical to try to test and support them all. The IBM Storage Scale team focuses on testing Enterprise Linux distributions RHEL, SLES and Ubuntu. However, there are some popular distributions in the Linux community that are created by essentially building the code from source packages corresponding to one of the enterprise distributions, usually with some cosmetic changes. IBM Storage Scale code may be able to work correctly on such a distribution, since it very closely resembles a supported one. However, we do not test IBM Storage Scale explicitly on such *clone* distributions, and will not be able to provide support for any problems specific to the use of the latter. If a problem is reported in such an environment, we will investigate it, but if the problem is suspected to be related to the type of distribution used, we may request that the problem be recreated on a supported distribution. Note that other IBM products may have a different support policy. We recommend that a supported distribution is used on NSD servers and other nodes that have SAN connectivity, to make it possible to get support with storage-related issues.

Note: IBM Storage Scale for Linux on Z is only supported on the distributions and kernel levels as documented in the question [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)

Q2.3:

What Linux distributions are supported by the integrated protocols access methods in IBM Storage Scale V4.1.1 and later?

A2.3:

For more information about which Linux distributions are supported by the integrated protocols access methods in IBM Storage Scale, see [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)

Two NFS services cannot run at the same time. You will need to stop and mask the Linux kernel `nfs-server` that is shipped with the distribution in order to use the NFS-Ganesha server that is shipped with IBM Storage Scale.

For NFS:

Because the IBM Storage Scale version of the NFS server must be used, the NFS service must be run with the `mmces` command instead of `systemctl/service`.

1. Determine if the `nfs-server` is available on any CES node by running the command `mmdsh -N cesnodes systemctl status nfs-server`. If any of the nodes report that the service is not masked, you will need to stop and mask the `nfs-server`:
 - `mmdsh -N cesNodes systemctl stop nfs-server`
 - `mmdsh -N cesNodes systemctl mask nfs-server`
2. You can then restart the IBM Storage Scale NFS service:
 - `mmces service start nfs -a`

For Swift Object:

You will need to reload the `spectrum-scale-object-selinux` module if the CES nodes are in SELinux Enforcing mode. To check whether they are in SELinux Enforcing mode, run the command: `mmdsh -N cesnodes getenforce`. If the mode is returned as Enforcing or Permissive, then run the following commands:

- `mmces service stop obj -a`
- `mmdsh -N cesNodes semodule -i /usr/share/selinux/packages/spectrum-scale-object-selinux.pp`
- `mmces service start obj -a`

Note: Also, see the question [15.5 What are the current advisories for IBM Storage Scale?](#)

Q2.4:

What is the impact on the `/dev`, `/proc/mounts`, `/etc/mtab` directories and the `mount` command for IBM Storage Scale for Linux due to the recent changes in `systemd`? What happened to the block device in `/dev`? Why is the `/dev/` prefix missing from the output of the `mount` command and also from `/proc/mounts` and `/etc/mtab`?

A2.4:

Starting with IBM Storage Scale 4.2.1, GPFS on Linux no longer creates a block device in `/dev` for the corresponding GPFS file system. As the result, the prefix `/dev/` does not appear before the GPFS device name in the output of the Linux `mount` command and in files `/etc/fstab`, `/etc/mtab`, and `/proc/mounts`. On the other hand, commands that accept `/dev/file-system-name` as input will continue doing so. A few commands will still display the file system name as `/dev/file-system-name`.

For example:

```
c13c1apv7:~ # awk '$3 == "gpfs" { print }' /proc/mounts
/gpfs/automountdir/fs2 /gpfs/automountdir/fs2 gpfs rw,relatime 0 0
fs1 /gpfs/fs1 gpfs rw,relatime 0 0
fs5mpathd /fs5mpathd/a/few/level/mount/point gpfs rw,relatime 0 0
remote /remote gpfs rw,relatime 0 0
/gpfs/automountdir/fs3mpathb /gpfs/automountdir/fs3mpathb gpfs rw,relatime 0 0
/fs4mpathc /gpfs/automountdir/fs4mpathc gpfs rw,relatime 0 0
/gpfs/automountdir/autofs1 /gpfs/automountdir/autofs1 gpfs rw,relatime 0 0
/autofs2 /gpfs/automountdir/autofs2 gpfs rw,relatime 0 0
/autofs3 /gpfs/automountdir/autofs3 gpfs rw,relatime 0 0

c13c1apv7:~ # mount | awk '/type gpfs/ { print }'
/gpfs/automountdir/fs2 on /gpfs/automountdir/fs2 type gpfs (rw,relatime)
fs1 on /gpfs/fs1 type gpfs (rw,relatime)
fs5mpathd on /fs5mpathd/a/few/level/mount/point type gpfs (rw,relatime)
remote on /remote type gpfs (rw,relatime)
/gpfs/automountdir/fs3mpathb on /gpfs/automountdir/fs3mpathb type gpfs (rw,relatime)
/fs4mpathc on /gpfs/automountdir/fs4mpathc type gpfs (rw,relatime)
/gpfs/automountdir/autofs1 on /gpfs/automountdir/autofs1 type gpfs (rw,relatime)
/autofs2 on /gpfs/automountdir/autofs2 type gpfs (rw,relatime)
/autofs3 on /gpfs/automountdir/autofs3 type gpfs (rw,relatime)
```

Q2.5:

What are the limitations of IBM Storage Scale support for Windows ?

A2.5:

- GPFS for Windows supports most of the GPFS features that are available on AIX and Linux. Exceptions include certain GPFS commands to apply policies, administer quotas and administer ACLs, among others. These commands are thus unsupported in a Windows-only cluster. In a mixed (heterogeneous) cluster, these Windows-lacking commands can still be executed on Unix nodes without participation from the Windows nodes in that cluster.

For more information, see the *GPFS limitations on Windows* topic in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

For more information about GPFS features that are not supported on Windows nodes, see [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)

Other limitations include:

- Exporting IBM Storage Scale file systems as Server Message Block (SMB) shares (also known as CIFS shares) from IBM Storage Scale Windows nodes is not supported.
- NFS serving (any version of NFS) by GPFS Windows nodes is not supported.
- IBM Storage Scale for Windows is not supported in any environment where Citrix Provisioning Services are deployed.
- Desktop editions of Windows (such as Windows 10 and Windows 11) do not support direct attached disks and can only operate as NSD clients. Windows Server editions support direct attached disks and can operate as NSD servers.
- In a mixed cluster, it is recommended that most GPFS administrative commands be executed on non-Windows nodes.
- The only supported way to achieve Windows-Unix user-mapping between Windows and Unix compute nodes is via RFC 2307 attributes. These attributes can be administered via Identity Mapping for Unix (IMU) from Microsoft in Windows Server versions up to and including Windows Server 2012 R2. Beginning with Windows Server 2016, these RFC 2307 attributes can be specified via the Active Directory Users and Computers (ADUC) MMC Snap-in as follows: From Administrative Tools, launch Active Directory Users and Computers (ADUC). Under View, enable Advanced Features. Next, navigate to the desired User object under Users. Open Properties. Under Attribute Editor tab, edit uidNumber, gidNumber, primaryGroupID, loginShell, unixHomeDirectory, and so on. IBM Storage Scale primarily uses the uidNumber and gidNumber attributes for user-mapping.
- In IBM Storage Scale V4 and later, the following user commands require Administrative privileges. They can only be run by a user who is a member of the Administrators group: mmchfileset, mmcrsnapshot, mmdelsnapshot, mmdf, mmlsdisk, mmlsfileset, mmlsfs, mmlspolicy, mmlspool, mmlssnapshot, and mmsnapdir.
- Encryption is not supported on Windows. The encryption function in the Advanced Edition should be disabled if Windows nodes are present in the cluster.
- IPv4 subnets are not supported in a cluster that is defined with IPv6 primary addresses (hostname) that contains Windows nodes.
- For TSM V7.1.1, which is only supported with IBM Storage Scale V4.1, see:
 - IBM Tivoli® Storage Manager V7.1.1 Knowledge Center at http://www-01.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.tsm.doc/welcome.html.
 - TSM support page at https://www-947.ibm.com/support/entry/myportal/product/tivoli/tivoli_storage_manager?productContext=-2105539168.

For more information, see the following questions:

- [What are the requirements for the use of OpenSSH on Windows nodes?](#)
- [Is IBM Storage Scale on Windows supported in a virtualization environment?](#)

Q2.6:
What are the requirements for the use of OpenSSH on Windows nodes?

A2.6:
IBM Storage Scale requires the use of OpenSSH to support its administrative functions when the cluster includes Windows nodes and UNIX nodes. Install the Cygwin OpenSSH package as described in the *Installing IBM Storage Scale on Windows nodes* chapter of the [IBM Storage Scale: Concepts, Planning, and Installation Guide](#). If you are using an OpenSSH package from another vendor, make sure that it is compatible with the Cygwin namespace and environment.

OpenSSH 9.0 includes a change that is incompatible with IBM Storage Scale. Ensure that OpenSSH 9.0 is not used with IBM Storage Scale. Earlier OpenSSH packages from Cygwin work. It is expected that this issue will be resolved with OpenSSH 9.1.


Q2.7:
What levels of Cygwin are supported by IBM Storage Scale on Windows?

A2.7:
The latest level of Cygwin that has been tested with IBM Storage Scale is 3.5.x.

If you use a level of Cygwin that is more recent than 3.5.x and you encounter issues, use `downlevel Base->cygwin` to install Cygwin 3.5.x and retry.

Q2.8:
Can different IBM Storage Scale maintenance levels coexist?

A2.8:
Different releases of IBM Storage Scale can coexist, that is, be active in the same cluster and simultaneously access the same file system. For release co-existence, IBM Storage Scale follows the N-1 rule. According to this rule, a particular IBM Storage Scale release (N) can co-exist with the prior release of IBM Storage Scale (N-1). This allows IBM Storage Scale to support an online (rolling) upgrade, that is a node by node upgrade. As expected, any given release of IBM Storage Scale can coexist with the same release. To clarify, the term release here refers to an IBM Storage Scale release stream and the release streams are currently defined as 5.0.x > 5.1.x > 5.2.x > 6.0.x.



These coexistence rules also apply for remote cluster access (multi-cluster remote mount). A node running release N-2 cannot perform a remote mount from a cluster which has nodes running release N, and vice versa.

Recommended release level: To consult information about the recommended release levels for IBM Storage Scale, see <https://www.ibm.com/support/pages/ibm-storage-scale-software-version-recommendation-preventive-service-planning>.

Q2.9:
Are there any requirements for Clustered NFS (CNFS) support in IBM Storage Scale?

A2.9:
 IBM Storage Scale supports Clustered NFS (CNFS) on SLES12 (see [Archived support information for IBM Storage Scale](#) in the FAQ) and RHEL levels supported by your version of IBM Storage Scale (see [What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#) in the FAQ). However, there are limitations:

- Exporting using NFS V4 is supported starting with IBM Storage Scale V4.1 or later.
- CNFS over IPV6 is only supported with IBM Storage Scale V4.1 or later.
- It is very important to make all the nodes in the same group as identical as possible - from the hardware and software running on them to the configuration of IBM Storage Scale, NFS, and the network.
- CNFS is not supported on the Ubuntu distribution.
- CNFS is not supported on SLES 15 and later versions.
- CNFS is not able to export a remotely mounted filesystem.

Q2.10:
Does IBM Storage Scale support NFS V4?

A2.10:
Note: NFSv4 can be supported in the following ways:

- Clustered NFS
- Integrated protocols (CES and NFS), which support NFSv3, NFSv4.0 and NFSv4.1.

Note: Some features of NFSv4.1 are not supported. For example, pNFS.

Enhancements to the support of Network File System (NFS) V4 are available on:

- AIX V6.1 or AIX V7.1.
- The following Linux distributions:
 - RHEL 5.5 and later, 6.x, and 7x.
 - SLES 11 SP1 and later, and SLES 12

Restrictions include:

- To support NFSv4 ACLs, the package **nfs4-acl-tools** must be installed.
- Windows-based NFSv4 clients are not supported with Linux/NFSv4 servers because of their use of share modes.
- If a file system is to be exported through CNFS (the Linux kernel NFSserver), then it must be configured to support POSIX ACLs (with **-k all** or **-k posix** option). This is because NFSv4/Linux servers will only handle ACLs properly if they are stored in GPFS as posix ACLs. On the other hand, if a file system is to be exported through CES (Using Samba and NFS Ganesha), then it must be configured to only support NFSv4 ACLs (with **-k nfs4**). This is because the CES stack has only been qualified with NFSv4 ACLs.
- Starting with Linux kernel version 2.6, an fsid value must be specified for each GPFS file system that is exported on NFS. For example, the format of the entry in **/etc/exports** for the GPFS directory **/gpfs/dir1** might look like this:

```
/gpfs/dir1 cluster1(rw,fsid=745)
```

For further details see the Linux export considerations at [Linux export considerations](#)

- Concurrent AIX/NFSv4 servers, Samba servers and GPFS Windows nodes in the cluster are allowed. NFSv4 ACLs may be stored in GPFS filesystems via Samba exports, NFSv4/AIX servers, GPFS Windows nodes, ACL commands of Linux NFSv3 and ACL commands of GPFS. However, clients of Linux v4 servers will not be able to see these ACLs, just the permission from the mode.

For more information on the support of NFS V4, see the IBM Storage Scale documentation updates file at

- For IBM Storage Scale V4.1.1 and later, at http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html
- For GPFS V3.5 and V4.1, at http://www-01.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html

Q2.11:

Are there any considerations for the use of the persistent reservation (PR) support in IBM Storage Scale?

A2.11:

Considerations for the use of Persistent Reserve include:

- Starting with IBM Storage Scale 5.2.1.0, NVMe reservations are supported only for multi-attach volumes in AWS virtual machines. For more information, see [NVMe reservations](#) in AWS documentation.
- Support for Persistent Reserve requires:
 - For V3.5 support on AIX V6.1 requires APAR IZ57224
AIX 6.1 TL7 + Service Pack 4 is required to support Persistent Reserve without SDDPCM.
 - For V3.5 , V4.1 or V4.2 support on AIX V7.1, refer to the storage documentation to install the correct multipath driver.
- The use of Persistent Reserve is supported on GPFS tie-breaker disks with GPFS V3.5.0.21, or later, and IBM Storage Scale V4.1.0.4, or later.

Note: See the question [What are the current requirements/limitations for using the Cluster Configuration Repository \(CCR\) ?](#)

- For the Activate Persist Through Power Loss (APTPL) feature:
 - On Linux, if the storage is capable of supporting APTPL, GPFS V3.5.0.15, or later supports this feature.
- Starting with 3.5.0.16, it is possible to have a **descOnly** disk that resides on a device that does not support SCSI-3 Persistent Reserve while allowing Persistent Reserve to be used on other disks in the same file system. The lack of Persistent Reserve support for the **descOnly** disk will not result in fast failover being disabled.

Also see the question [What devices does GPFS support with SCSI-3 Persistent Reservations?](#)

Q2.12:

What are the requirements/limitations for using native encryption in IBM Storage Scale Advanced Edition or Data Management Edition?

A2.12:

Considerations for the use of native encryption (encryption of data at rest on GPFS disks) in IBM Storage Scale Advanced Edition include:

- For each node that acts as a key server, the installation and use of either of the following key servers is required.
 - [IBM Security Key Lifecycle Manager \(ISKLM\) V2.6](#) or later
 - [Thales CipherTrust Manager 2.5.x or 2.8](#) or later
 - [HashiCorp Vault Enterprise 1.12](#) or later
 - [IBM Key Protect with Key Management Interoperability Protocol \(KMIP\)](#)

Note: Key server hosts are not required to be members of the IBM Storage Scale clusters that use them.

Key server nodes are not required to be members of the IBM Storage Scale cluster(s) that use them.

For ISKLM:

- IBM Security Guardium Key Lifecycle Manager (GKLM) V5.0.0 is supported with the IBM Storage Scale Encryption feature.
- IBM Security Guardium Key Lifecycle Manager (GKLM) V4.2.1 is supported with the IBM Storage Scale Encryption feature.
- IBM Security Guardium Key Lifecycle Manager (GKLM) V4.1.1 is supported with the IBM Storage Scale Encryption feature.
- IBM Security Guardium Key Lifecycle Manager (GKLM) V4.1.0.1 (IF01) is supported with the IBM Storage Scale Encryption feature.
- ISKLM is not shipped with nor licensed with IBM Storage Scale and must be purchased separately.
- ISKLM V2.6, or later of the server software (D0887LL) must be installed on each node that acts as a key server.
- See the IBM Security Key Lifecycle Manager documentation for details about licensing that offering.

For Thales CipherTrust Manager:

- Thales CipherTrust Manager is not shipped with nor licensed with IBM Storage Scale and must be purchased separately. Contact Thales directly to purchase.

For HashiCorp Vault:

- HashiCorp Vault is not shipped with nor licensed with IBM Storage Scale and must be purchased separately. Contact HashiCorp directly to purchase.

- Every node that accesses the encrypted data must be running the Advanced Edition or Data Management Edition of IBM Storage Scale.
- Every node that accesses the encrypted data, and also nodes which play a management role in the file system (such as manager node, an NSD server, or a node which participates in the restripe of a file system), must have network connectivity to the key server.
- IBM Storage Scale client nodes do not require a key server license.
- For IBM Key Protect with KMIP:
 - IBM Key Protect with KMIP is available in IBM Cloud. For more information, see [IBM Key Protect for IBM Cloud](#).
 - To use IBM key Protect with KMIP in IBM Storage Scale, the regular setup is required. For more information, see [Regular setup: Accessing a remote file system](#).

Current limitations of the IBM Storage Scale encryption function include:

- Only user data is encrypted. The encryption of directories or other metadata is not supported.
- Extended attributes are not encrypted.
- Data, which is backed up, is in cleartext unless encryption is supported by the backup system.
- Data, which is migrated to tape using software such as IBM Storage Protect or IBM Spectrum Archive, is in cleartext unless the tape system and the connection between them (if Ethernet or InfiniBand) provide encryption.
- Encryption is not supported on Windows. The encryption function should be disabled when Windows nodes are in the cluster.
- The contents of encrypted files are placed into a local read-only cache (LROC) based on the settings of the `lrocEnableStoringClearText` configuration option. For more information, see the "Encryption and local read-only cache (LROC)" section.
- For more information, see [Encryption requirements and limitations](#) and [Q6.12 How should IBM Storage Scale Advanced Edition or Data Management Edition be configured to only use FIPS 140-2-certified cryptographic engines?](#)

Q2.13:

Are there any considerations when utilizing the Simple Network Management Protocol (SNMP)-based monitoring capability in IBM Storage Scale?

A2.13:

Considerations for the use of the SNMP-based monitoring capability include:

Important: SNMP events and queries that are directly received or sent by the `mmfsd` daemon in IBM Storage Scale have been discontinued with IBM Storage Scale 6.0.0. For more information, see [Table 3: Features discontinued in IBM Storage Scale in Stabilized, deprecated, and discontinued features in IBM Storage Scale](#)

- The SNMP collector node must be a Linux node in your GPFS cluster. GPFS utilizes Net-SNMP which GPFS does not support on AIX.
- Support for ppc64 requires the use of Net-SNMP 5.4.1. Binaries for Net-SNMP 5.4.1 on ppc64 are not available. You will need to download the source and build the binary. Go to <http://net-snmp.sourceforge.net/download.html>
- If the monitored cluster is relatively large, you need to increase the communication time-out between the SNMP master agent and the GPFS SNMP subagent. In this context, a cluster is considered to be large if the number of nodes is greater than 25, or the number of file systems is greater than 15, or the total number of disks in all file systems is greater than 50. For more information see [Configuring Net-SNMP](#) in the *IBM Storage Scale: Problem Determination Guide*.
- SNMP-based monitoring has not been tested in clusters composed of more than 127 nodes.

Q2.14:

What are the current limitations and advisories for using the `mmbackup` command?

A2.14:

Current limitations and advisories include:

- Beginning with IBM Storage Scale V4.2, in file systems that are managed by an HSM system, **mmbackup** will skip over candidates for backup that are migrated offline to avoid causing a recall storm. Instead, records for these files will be added to a file in the root of the fileset called **mmbackup.hsmMigFiles.name of server**. System managers should recall these changed files online to allow **mmbackup** to properly protect them in the next invocation.
- File systems with IBM Storage Protect for Space Management that have unlinked filesets, will be required to link all filesets when issuing the **mmbackup** command for the first time after you upgrade GPFS cluster from GPFS 3.5.0.10 or lower, to GPFS 3.5.0.11 or higher. If you have any concerns regarding this requirement, please contact GPFS service.
 - In the United States contact us toll free at 1-800-IBM-SERV (1-800-426-7378)
 - In other countries, contact your local [IBM Service Center](#)
- Use of the IBM Storage Protect Backup-Archive client option **SKIPACLUPDATECHECK** with the **mmbackup** command requires IBM Tivoli Storage Manager release 6.4.1.0 or later.

Note: Beginning with Version 7.1.3, IBM Tivoli Storage Manager is now IBM Storage Protect.

- The GPFS **mmbackup** command is not integrated with the IBM Storage Protect for Space Management-Multi HSM Server feature. See [Managing a file system with multiple Tivoli Storage Manager servers](#).

- Restoring a file via a node that has a different architecture than the one used to do the backup could cause the associated ACL to be corrupted.

For example, if a file was backed up using an x86_64 node, and then restored using a ppc64 node, this could cause its ACL to be corrupted (caused by differences in endianness of the architectures which is not supported by the GPFS APIs used in the restore operation). It is recommended that backup and restore operations be done on similar types of nodes.

- The **mmbackup** command supports backup of a whole file system from a global snapshot.

The **mmbackup -S** snapshot command option is supported with IBM Storage Scale V4.1.1 on either a global snapshot for the whole file system or for a fileset backup if the fileset was captured in that snapshot. It is also supported with a fileset snapshot for a fileset backup, providing the name of the snapshot is unique among all snapshot names. Do not use the same snapshot name for multiple snapshots.

The **mmbackup -S** snapshot command option is supported with GPFS V3.5.0.3 or later. GPFS V3.4 and GPFS V3.5.0.2 or lower do not support backup from a snapshot. If the snapshot directory for global snapshots and the directory for fileset level snapshots are different, then GPFS V3.5.0.4 or higher level is required.

- Doing backup from a snapshot in an IBM Storage Protect for Space Management managed file system could cause recall of migrated files.

In an HSM managed file system such as IBM Storage Protect for Space Management, using **mmbackup** to back up from a snapshot could cause the recall of migrated files if the migration was done after the snapshot was taken. This is due to the fact that a snapshot is a static view of the file system which does not reflect migration state changes. To avoid recalling data from migrated files, create the snapshot and complete the backup operation before migrating files or make sure that migration is done before the snapshot is taken for a backup operation. Backup operations will not recall files if the snapshot captured the files in their migrated state. Until and unless the migrated file stubs are removed from the live file system. In this case a recall will be required to populate the contents of the snapshot view of the files. If a snapshot exists, consider recalling files by using the IBM Storage Protect for Space Management "tape optimized recall" function before deleting migrated files from the active file system.

- Doing backup from a snapshot in an IBM Storage Protect for Space Management managed file system could cause failure to backup migrated files.

- The use of unsupported characters in the names of files or directories will cause failures.

mmbackup uses the IBM Storage Protect Backup Archive client to backup data to the IBM Storage Protect server. As IBM Storage Protect currently does not support all special characters in file or directory names, they cannot be supported by **mmbackup**. If special characters are used in the names of files or directories backed up by the **mmbackup** command, failures will result. Known special characters which can cause problems include: *, ?, ", ', control-X, control-Y, carriage return and the new line character. Use of the IBM Storage Protect options **QUOTESARELITERAL** and **WILDCARDSARELITERAL** along with the **--noquote** command line option to **mmbackup** will allow support for all special characters except carriage return, new line, control-X, and control-Y.

- Beginning with IBM Storage Scale V4.1.1, backup of either the entire file system or a selected fileset is supported. Nested fileset arrangements where one fileset is linked inside another are not supported by **mmbackup** on a fileset. Nesting remains supported for whole file system **mmbackup**. The first **mmbackup** of any fileset must be made using the option **-t full** to avoid causing accidental invalidation of existing backups that may exist of previously existing nested filesets.
- Differences in the way the **mmbackup** command and IBM Storage Protect process the **include** and **exclude** statements in the `dsm.sys` configuration file may cause files or directories to be included or excluded unexpectedly.

Known differences in processing include, but are not limited to:

- The **mmbackup** command does not support **exclude.archive**, **exclude.file.spacemgmt**, **exclude.spacemgmt**, **exclude.fs**.
- Whether or not there is a / at the end of an **exclude.dir** affects the way **mmbackup** decides what files or directories are excluded.
- exclude.file** may cause incorrect files to be backed up if the pattern presents a wildcard at the end.

Q2.15:

What are the current limitations and advisories for using Scale Out Backup and Restore (SOBAR) ?

A2.15:

Current limitations and advisories include:

- IBM Storage Scale Image Backup and Restore (SOBAR) has been tested in a standalone manner; but must be tested with Data Management/HSM products before deployment by customers with such products in production environments. Customers who are interested in making use of this function should contact scale@us.ibm.com.
- SOBAR does not support the backup or restore of a file system with Active File Management (AFM) filesets.
- SOBAR supports backup from a global snapshot. Independent fileset snapshots are not supported.

Q2.16:

What are the current limitations for using the File Placement Optimizer (FPO) function?

A2.16:

Important: File Placement Optimizer (FPO) was discontinued with IBM Storage Scale 6.0.0. For more information, see *Table 3: Features discontinued in IBM Storage Scale* in [Stabilized, deprecated, and discontinued features in IBM Storage Scale](#).

Current limitations include:

- The File Placement Optimizer (FPO) function is supported on IBM Storage Scale V5 on Linux and AIX operating systems.
- AFM ADR (primary/secondary filesets) is not supported on an FPO enabled file system.
- With the AFM function, if you want to maintain data locality on both home and cache, they must have the same Failure Group configuration. Additionally, block placement policy must be set via write-affinity-failure-group at both sites.
- Twin-tailed disks are supported in an FPO pool only when a single NSD server is defined for each disk.
-
- Contact scale@us.ibm.com if you plan to deploy a cluster with more than 32 nodes in a Shared Nothing Cluster, or SNC, in which no disks in the cluster are served by more than a single node. This includes FPO nodes. Shared Nothing Clusters that are larger than 32 nodes must be reviewed and approved by IBM before deployment. This limitation applies to clusters that have more than 32 nodes that have disks serving a file system with data replication enabled, and these disks are only accessible from a single node.

Note: You can determine if a given file system has data replication enabled by checking if the -R setting (the maximum number of data replicas) reported by the `/usr/lpp/mmfs/bin/mm1sfs` command is greater than 1.

- If a storage pool is FPO-enabled (allowWriteAffinity=yes), then layoutMap=cluster must also be specified.
-
- With GPFS V4.1, use of the `mmrestripefile -b, mmadddisk -r` and the `mmrestripefs -b` commands will break the original FPO file's placement.
- With GPFS V4.1, use of the `mmrestripefile -r` and the `mmrestripefs -r` commands is supported with locality awareness. Use of the commands with clones and snapshots will break the original FPO file's placement.
- On clusters with the FPO function enabled, in order to utilize the `mmrestorefs` command, you must specify the write-affinity-failure-group policy.
- If the size of a file is less than the value of the block size divided by 32, the write affinity depth policy and the write affinity failure group policy will not be followed. Data is widely striped instead.
- The setXattr function cannot set the FPO extended attributes writeAffinityDepth, write-affinity-failure-group, and BlockGroupFactor for a clone file or the policy MIGRATE rule. Respectively, setWAD, setWADFG, and setBGF should be used.
- The extended attributes writeAffinityDepth, write-affinity-failure-group, and BlockGroupFactor are for use only on an FPO pool.
- Starting in IBM Storage Scale 5.0.5, FPO and SNC remain available. However, it is recommended to limit the size of deployments to 32 nodes. There are no plans for significant new functionality in FPO nor increases in scalability. The strategic direction for storage using internal drives and storage rich servers is IBM Storage Scale Erasure Code Edition.
- The FPO configuration is not supported on IBM Storage Scale Erasure Code Edition. For more information on IBM Storage Scale Erasure Code Edition, see [IBM Storage Scale Erasure Code Edition documentation](#).

Q2.17:

What are the current limitations for using the Active File Management (AFM) Async DR function?

A2.17:

Limitations are added and deleted from time to time. For more information about the limitations that affect a particular release, see the AFM limitations section under **Product Overview > Active File Management** in the Knowledge Center or in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q2.18:

What are the current limitations for using the Active File Management (AFM) function?

A2.18:

Limitations are added and deleted from time to time. For more information about the limitations that affect a particular release, see the AFM limitations section under **Product Overview > Active File Management** in the Knowledge Center or in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q2.19:

What are the current limitations common to both the Active File Management (AFM) and AFM DR functions?

A2.19:

Limitations are added and deleted from time to time. For more information about the limitations that affect a particular release, see the AFM and AFM DR limitations section under **Product Overview > Active File Management** in the Knowledge Center or in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q2.20:

What is the currently recommended transport protocol for AFM and AFM DR data transfers?

A2.20:

For the current recommendations regarding the transport protocol for AFM and AFM DR data transfers, see *The backend protocol - NFS versus NSD* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q2.21:

What is the current limitation for using the snapshot restore function?

A2.21:

For mixed clusters containing GPFS V3.5, IBM Storage Scale 4.1.0, 4.1.0.4, 4.1.1, or V4.2 nodes, it is recommended that for both fileset and global snapshot restore, the **mmrestorefs** command is issued from a node running the most recent version of IBM Storage Scale. Version 4.1.1 and later, provides some new snapshot restore functionality so IBM Storage Scale attempts to intelligently use the latest features.

Note:

1. IBM Storage Scale V4.2 only interoperates with IBM Storage Scale V4.1

2. The **mmrestorefs** command is not supported on the Express Edition.

- To improve performance support for the **-N** parameter on the **mmrestorefs** command has been phased in over the past few releases.
 - In GPFS 3.5 and earlier there is no **-N** parameter
 - In GPFS 4.1 **-N** can be used for fileset snapshot restore only
 - In IBM Storage Scale 4.1.1 and later **-N** can be used for both fileset and global snapshot restores
- For global snapshot restore:
 - If the **mmrestorefs** command is issued from a pre-4.1.1 node:
 1. The file system must be unmounted.
 2. The file system manager performs the restore.
 - If the **mmrestorefs** command is issued from a 4.1.1 or later node:
 - ◆ The file system must be mounted.
 - ◆ By default the restore is performed on all nodes running the latest level of code.
- For fileset snapshot restore:
 - If the **mmrestorefs** command is issued from a V3.5 node:
 1. The file system must be unmounted.
 2. The file system manager performs the restore.
 - If the **mmrestorefs** command is issued from a 4.1 or later node:
 - ◆ The file system must be mounted.
 - ◆ By default the restore is performed on all nodes running the latest level of code.

Q2.22:

What are the current requirements when using local read-only cache?

A2.22:

The current requirements/limitations for using local read-only cache include:

- A minimum of IBM Storage Scale V4.1.0.1.
- Local read-only cache is only supported on Linux x86 and Power.
- The minimum size of a local read-only cache device is 4 GB.
- The local read-only cache requires memory equal to 1% of the local read-only device's capacity.

Note: Use of local read-only cache does not require a server license

Q2.23:

What are the current requirements/limitations for using the Cluster Configuration Repository (CCR) ?

A2.23:

The current requirements/limitations for using the Cluster Configuration Repository (CCR) include:

- IBM Storage Scale V4.1.0: The Disaster Recovery procedures described in the Advanced Administration Guide are not supported in a cluster with CCR enabled:
 - Do not run **mmchcluster --ccr-enable** for existing clusters
 - Use **mmcrcluster --ccr-disable** for new clusters
- IBM Storage Scale V4.1.1, or later, there are no limitations for using CCR.

Q2.24:

What are the current requirements/limitations for using Ubuntu?

A2.24:

The current requirements/limitations for using Ubuntu include:

- The minimum level of Ubuntu supported is 14.04.1.
- Only IBM Storage Scale 4.1.0.8 or later, is supported with 14.04.2.
- Only IBM Storage Scale 4.1.1.9/4.2.1.1 or later is supported with 14.04.4/16.04.
- The minimum kernel level supported is 3.13.0.32.
- Only IBM Storage Scale 4.2.3.10/5.0.1.2 or later is supported with 18.04.1.
- Only IBM Storage Scale for Linux on Z 4.2.3.10/5.0.1.2 or later is supported with 18.04.
- P8 is supported with Little Endian only and only with GPFS for Linux on System p base RPMs dated January 2015 (GPFS V4.1.0.5 or later).

If you have Software Maintenance Agreement (SWMA) for your products ordered through AAS/eConfig or IBM Subscription and Support (S&S) for orders placed through Passport Advantage, you may log into the respective systems and upgrade your level of GPFS:

- For products ordered through AAS/eConfig, please log into the Entitled Software page at: <https://www-05.ibm.com/servers/eserver/ess/OpenServlet.wss>
- For products ordered through Passport Advantage, please log into the site at: <http://www.ibm.com/software/lotus/passportadvantage/>
- GPFS V3.5.0.22 or later is only supported on x86_64 architecture
- When issuing **make World** for Ubuntu 14.04.1, this warning will appear but can be disregarded because **kdump-kern-dummy.ko** is not utilized by GPFS.

WARNING: ".TOC." [/usr/lpp/mmfs/src/gpl-linux/kdump-kern-dummy.ko] undefined!

- As Tivoli Storage Manager (TSM) does not support Ubuntu, GPFS commands that utilize TSM are not supported on Ubuntu.

Note: Beginning with Version 7.1.3, IBM Tivoli Storage Manager is now IBM Storage Protect.

- If you use CNFS or CES features under Ubuntu, verify that the iputils-arping package is installed. See the Software requirements page for more information.

Q2.25:

What are the current requirements/limitations for IBM Storage Scale for Linux on Z?

A2.25:

The current requirements and limitations for IBM Storage Scale for Linux on Z include:

- A leapp upgrade from Red Hat Enterprise 7.x to 8.x is not supported.
- File Placement Optimizer is not supported.
- For support of backup and restore functions with IBM Storage Scale for Linux on Z, see the following support matrices:

Note: Starting with IBM Storage Scale V5.0.0, the IBM Storage Protect for Space Management client is no longer supported on IBM Storage Scale for Linux on Z.

- For the IBM Storage Protect Backup Archive client, see [Hardware and software requirements for IBM Storage Protect Linux zSeries Backup-Archive and API Client](#).
- For the IBM Storage Protect for Space Management client, see [IBM Storage Protect for Space Management \(HSM\) requirements for Linux on IBM z Systems®](#).
- For supported storage, see the [What disk hardware has IBM Storage Scale been tested with?](#) and the [IBM Storage Scale for Linux on Z support Direct Attached Storage Devices \(DASD\)?](#) questions.
- Support for stretched cluster with synchronous mirroring utilizing block-level replication:
 - For IBM Storage Scale 5.0.0 and later, with distances up to 300 km.
- Kernel NFS (v3 and v4) is supported. Clustered NFS function (CNFS) is not supported.

Note:

1. Central Processor Assist for Cryptographic Function (CPACF) is supported. CPACF is IBM Z hardware encryption acceleration. It is incorporated in the central processors that are shipped with IBM Z. To benefit from the CPACF, you must install LIC internal feature 3863 (Crypto Enablement feature), which is available free of charge. By default, IBM Z is delivered to customers without this feature unless it is ordered explicitly by the customer. The installation of this feature at a future time is non-disruptive.
2. IBM z15 and z16 offer the Integrated Accelerator for zEnterprise® Data Compression (zEDC). This feature is enabled by default starting from IBM Storage Scale 5.1.0 when the CPU feature 'dflt' is listed in `/proc/cpuinfo`. Compression options

such as z, zfast, alphae and alphah benefit from this feature. The feature also depends on your installed zlib version. For more information, see <https://linux.mainframe.blog/zlib-acceleration/> for details.

Q2.26:
What are the current requirements/limitations for the deadlock amelioration function in IBM Storage Scale?

A2.26:
The current requirements/limitations for use of the deadlock amelioration function include:

- Deadlock amelioration functions are fully supported in IBM Storage Scale V4.
- In a cluster with **minReleaseLevel** below 4.1.0, that consists of all GPFS 4.1 nodes or a mixture of 4.1 and 3.5 nodes, the deadlock amelioration functions may still work partially. In order to avoid a problem of tracing not being turned off after GPFS code turns it on make sure to have 3.5.0.24 or later, or 4.1.0.7 or later, or have APAR IV69797 applied to all nodes. Running with tracing on could have performance implications.

Q2.27:
What are the requirements/limitations for using the IBM Storage Scale GUI?

A2.27:
Considerations for using the IBM Storage Scale GUI include:

- The GUI is available with IBM Storage Scale editions: Data Access, Data Management, Developer, Standard, and Advanced for Linux on x86, Linux on Z, and Power (Big Endian and Little Endian).
- The GUI is supported on RHEL 8.10 or later, SLES 15 SP6 and SP7, and Ubuntu 22.04 or later on Linux x86, Linux on POWER, and Linux on Z platforms. For more information, see [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)
- The maximum number of nodes supported is:
 - With V4.2.1 and later, 1000 nodes
 - With V4.2, 128 nodes
- The GUI works with either a client or server license.
- The GUI depends on a pre-installed PostgreSQL server usually installed already with the operating system installation but if this is not the case if you install the operating system from scratch you need to do this before you install the GUI or the installation fails.

Q2.28:
What are the requirements/limitations for using the compression function?

A2.28:
Compression support excludes the following:

- Compressing files in snapshots.
- Compressing clones and cloning compressed files.
- Small file compression (files consuming less than 2 sub-blocks, compressing small files into inode).
- Compression of non-regular files, such as directories.
- Compression of files in Windows hyper allocation mode.
- File compression does not compress a memory-mapped file.
- File compression does not compress a file that is opened for Direct I/O.

Additionally:

- Compression is supported in an FPO environment or horizontal storage pools with IBM Storage Scale V4.2.1 and later.
- Compression in this release is optimized for cold data or write-once objects and files. It uses the zlib data compression library and favors saving space over speed. Usage on other types of data may result in performance degradation.
- On Windows:
 - Compression of a file on Windows is only enabled via the `mmchattr` command.
 - The following Windows APIs are not supported:
 - ◆ `FSCTL_SET_COMPRESSION` to enable/disable compression on a file
 - ◆ `FSCTL_GET_COMPRESSION` to retrieve compression status of a file
 - In Windows Explorer, in the Advanced Attributes window, the compression feature is not supported.
- On IBM Z:
 - The IBM zEDC hardware compression feature is enabled by default starting from IBM Storage Scale 5.1.0 when the CPU feature 'dflt' is listed in `/proc/cpuinfo`. Compression options such as z, zfast, alphae and alphah benefit from this feature. The feature also depends on your installed zlib version. For more information, see <https://linux.mainframe.blog/zlib-acceleration/>.

Q2.29:

What are the requirements/limitations for using the Quality of Service (QoS) function?

A2.29:

Considerations for using the Quality of Service function include:

- Filesets can be associated only with user-defined classes; system classes are not supported.
- There is no fixed limit on the number of user-defined classes. However, limits might be imposed by the operating system or available hardware resources.
- Dynamic QoS Rebalancing (DQR) is enabled by default for user-defined classes.
- DQR is permanently disabled for system classes.
- The `mmqos` command is supported only on Linux.
- For Linux on Z, QoS is only supported with V4.2.1 and later.
- No throttling for direct I/O operations.
- Not supported on AFM cache and AFM-based asynchronous DR filesets.
- The following flash for QoS was issued:

Abstract:

In an IBM Storage Scale V4.2 file system with multiple storage pools, Quality of Service (QoS) settings should be set for all storage pools to avoid performance degradation for unspecified storage pools.

Problem Summary:

In an IBM Storage Scale V4.2 file system with multiple storage pools, if the user specifies Quality of Service for I/O operations (QoS) settings (for the **maintenance** and **other** classes) only for one storage pool then the I/O allocations for the unspecified pools will be set to a very low value, resulting in severe performance degradation when I/O is performed to the unspecified storage pool(s).

See the complete Flash at <http://www.ibm.com/support/docview.wss?uid=ssg1S1005464>

Q2.30:

What are the considerations when running on SELinux?

A2.30:

The following considerations apply to running SELinux:

- From the 5.0.5 release, IBM Storage Scale runs on Red Hat Enterprise Linux operating systems with Security-Enhanced Linux (SELinux). For more information, see the *Security-Enhanced Linux support* topic in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- When using the installation toolkit, the IBM Storage Scale Swift Object protocol functionality requires the following SELinux packages to be installed:
 - `selinux-policy-base` at 3.13.1-23 or higher
 - `selinux-policy-targeted` at 3.12.1-153 or higher
- When using the Swift Object protocol functionality, enabling SELinux after IBM Storage Scale has been installed is not supported. Contact IBM Storage Scale support at scale@us.ibm.com if you have questions about this restriction.

Q2.31:

What enhancements to quota management are available?

A2.31:

Enhancements to quota management available with IBM Storage Scale V4.1 and later, allow for quota clients to dynamically acquire and relinquish quotas based upon consuming rate; that is, the quota manager grants quota shares based upon global quota information such as the remaining quota limit and the number of mounted clients. Decisions based upon existing information provides for greater efficiency in managing quotas, however, this could result in an increase in the in-doubt values from earlier releases when running with a heavy IO workload. In order to get more accurate usage, issuing the `mm1squota` and `mmrepquota` commands with the `-e` option might help.

Note: Quota report accuracy is mainly affected by hardware errors, such as node/network failures. Having a large number of nodes could increase the chances of having these types of failures.

Q2.32:

What are the current limitations of quota management?

A2.32:

The aggregate total number of quota records in user, group, and fileset quota files is limited to 200K records per file system. This limitation is due to the maximum amount of data that can be exchanged between the quota manager and quota command client such as the `mmrepquota` command.

Note: A large number of quota records per file system can result from the following scenarios:

- There are a very large number of users, groups, or filesets.
- If the `--perfilesset-quota` option is enabled, the number of possible quota records is the number of filesets times number of users (and groups).

Q2.33:

Is file immutability supported with IBM Storage Scale?

A2.33:

Yes. For more information about IBM Storage Scale immutability functions, configuration, and operation, see *Immutability and appendOnly features* in the *IBM Storage Scale: Administration Guide* and the following Redpaper: <http://www.redbooks.ibm.com/abstracts/redp5507.html>.

Q2.34:

Has file immutability been assessed for compliance?

A2.34:

The immutability function of IBM Storage Scale 5.1.0 has been assessed for compliance in accordance to Securities and Exchange Commission (SEC) Rule 17a-4(f), Financial Industry Regulatory Authority (FINRA) Rule 4511(c) and the principles-based electronic records requirements of the Commodity Futures Trading Commission (CFTC) in 17 CFR § 1.31(c)-(d). To view the detailed assessment report, see [IBM Storage Scale Assessment Report](#).

The immutability function of IBM Storage Scale 5.0.0 has been assessed for compliance in accordance to US SEC17a-4f, EU GDPR Article 21 Section 1, German and Swiss laws and regulations by a recognized auditor. For more information, see the following links:

Assessment report: <http://www.kpmg.de/bescheinigungen/RequestReport.aspx?B290411BE1224F5A9B4D24663BCD3C5D>

Certificate: <http://www.kpmg.de/bescheinigungen/RequestReport.aspx?DE968667B47544FF83F6CCDCF37E5FB5>

Q2.35:

Is there any guidance for RHEL 8 installations on IBM Storage Scale?

A2.35:

Python 2 needs to be installed on new RHEL 8 installations because it is not installed by default. It is recommended that you create both the RHEL 8 BaseOS and AppStream repositories so that the package dependencies are met during installation.

If you intend to perform a leapp upgrade from RHEL 7.6 (versus a first-time installation) to RHEL 8, keep the following information in mind:

1. It is highly recommended that you use leapp-0.8.1-1 or higher. Make sure that the following requirements are met and the latest RHEL upgrade procedure is followed: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/upgrading_to_rhel_8/index
2. Consider running the boom utility to manage more boot loader entries on the system and to provide a path back to RHEL 7.6 if necessary: <https://www.redhat.com/en/blog/upgrading-rhel-7-rhel-8-leapp-and-boom>
3. As of the 5.0.4 release, if IBM Storage Scale 5.0.4 packages are installed and the leapp upgrade utility is run, leapp might remove some IBM Storage Scale packages and some dependencies. Manually reinstalling the removed packages will be required. Either proceed at your own risk, or consider not using the leapp upgrade utility and provisioning a new RHEL 8 node.
4. If protocols and authentication are enabled, python-ldap needs to be installed for mmadquery to run. A leapp upgrade from RHEL 7.x to RHEL 8.x removes the python-ldap package, which is a prerequisite for mmadquery. Ensure that you install python-ldap after the leapp upgrade if you intend to use mmadquery on RHEL 8.x.

Note:

- RHEL 8 is not supported by the Swift Object protocol in releases earlier than IBM Storage Scale 5.1.0.
- RHEL 8 is supported by Transparent cloud tiering starting with IBM Storage Scale 5.0.5.

For more information, see *Guidance for Red Hat Enterprise Linux 8.x on IBM Storage Scale nodes* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q2.36:

What are the requirements/limitations for using the file clone function?

A2.36:

The following requirements and limitations apply to using the file clone function:

- A compressed file cannot be cloned and a clone file cannot be compressed.
- mmap is not supported on clone child files.

Q2.37:

What are the requirements/limitations for using the memory-mapped (mmap) function?

A2.37:

mmap is not supported on clone child files.

Q2.38:**Are IBM Storage Scale packages signed by IBM?****A2.38:**

Starting with the IBM Storage Scale release 5.0.4, all IBM Storage Scale packages on Red Hat Enterprise Linux and SLES operating systems on supported architectures are signed by IBM with a GPG key. Starting with the IBM Storage Scale release 5.0.5.1, repository metadata is also signed by IBM.

You can use the available public key to verify the signatures on the packages and repository metadata.

The latest public key contents are as follows:

```
-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: EKM
mQENBF0tE4ABCADTU4imcpD1IHvcK/qWdMMrs72lL9EYDtA/JNL5YCPNeIa/54aIe3xXFJZbkjs
v+5INaxYv0DEQxXEFq8vA1pQGPIG1elb3fXgP7Iyfiy13KD1rVEB8AY/Cr/zTmHV8IJNMN8jcB16Z
vAED7fXE82Q4jQ3djbG00YBq2PeVS+wM5Y8n1+tmpVmcD9oLzYhJPeCsbFi6BAVgXBmyh4arrn15
OL5fD5jBnn0T926N2mpnsfubyGitQ1ywjJJuESnF9Ub9QMT7jNjGcg6frxHV0MUsIstmg01GBnvx
I/P/BvdiIqGj0TInka78+rYJpxZWP1bu/Xg/NXJ9sERjXuT30GCHABEBAAG0DVNwZWN0cnVtU2Nh
bGwJATkEEWIEIACMFA10tE4ACGy8HCwkIBwMCAQYVCAIJCgsEFgIDAQIeAQIXgAAKCRc9vnXD7+tu
654UCAC1AR99JhSDm47V2Jv0BYLxcxdHqoqY+MqgKxeuy11Tp/enpqoGigZAcPbz1Rv1JTOyh0Pa
PQC1y0oaKDROR5a0uCd0Cz3xbSQ92mWX0Fka7D9KN1AuxlGD6Ic58AvQ6RBv/mxblfH6gXH1c+Q0
+YFOY5Y1vgKLYJ+exGngzieZfxspyyTab7FZe06G/1Cm9U+m0fQ/70DH6AvNIRmsCCg5uUeAmQ0a
3+0RpWzN09nlkSY1kM1XvyZSWpwTEXLPtfdW0kzxs1k4IzgyFM0sw6oL05TVMZyL828M0uJtqU8
06rS6+/RIho7GhiQ8SklugSF1FnT9fx5TRJcCiJmhHeF
=v7fu
-----END PGP PUBLIC KEY BLOCK-----
```

Save these contents in a file and import it to verify the signature manually. The installation toolkit does the signature verification automatically before installation or upgrade. For more information, see https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.5/com.ibm.spectrum.scale.v5r05.doc/bl1ins_verifysignature.htm.

Q2.39:**What are the components required for GPUDirect Storage (GDS) support with IBM Storage Scale?****A2.39:**

The following software components are required for GPUDirect Storage (GDS) support with IBM Storage Scale:

- The required IBM Storage Scale version depends on the fabric type over which GPUDirect Storage is run:
 - Infiniband: IBM Storage Scale 5.1.2 or later.
 - RoCE: IBM Storage Scale 5.1.3 or later.
 - The recommended version is IBM Storage Scale 5.1.6 as it includes accelerated GDS writes.
- Linux distros:
 - For GDS clients: Ubuntu 20.04, RHEL 8.4, RHEL 8.6.
 - For storage servers: Any Linux distro that supports IBM Storage Scale 5.1.2 (Infiniband), 5.1.3 (Infiniband, RoCE) or later.
- Mellanox OFED: MOFED v5.4-1.0.3.0, MOFED v5.6-2.0.9.0
- Nvidia CUDA
 - Infiniband: CUDA 11.4.2, CUDA 11.5.1, CUDA 11.6.2, CUDA 11.7, CUDA 11.8
 - RoCE: CUDA 11.5.1, CUDA 11.6.2, CUDA 11.7, CUDA 11.8
 - CUDA 11.8 is required for accelerated GDS writes

Note:

- IBM Storage Scale 5.1.6 introduces accelerated GDS writes. For more information, see [GPUDirect Storage support for IBM Storage Scale](#).
- Asynchronous CUDA IO is not supported.
- The IBM Storage Scale 5.1.1 technical preview is not compatible with IBM Storage Scale 5.1.2 and later.

References:

- For more information and supported hardware, see *Planning for GPUDirectStorage* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

- For the base configuration of IBM Storage Scale on RoCE fabrics, see [Highly Efficient Data Access with RoCE on IBM Elastic Storage Systems and IBM Storage Scale](#).

Q2.40:

What are the requirements/limitations for manipulating the system.nfs4_acl extended attribute directly in IBM Storage Scale?

A2.40:

The following are the requirements/limitations for manipulating the system.nfs4_acl extended attribute directly in IBM Storage Scale:

- Manipulating the system.nfs4_acl extended attribute is supported only on Linux.
- It is recommended to use version 0.4.2 or later when installing the nfs4-acl-tools package.
- nfs4-acl-tools 0.4.2 and later versions are expected to work without problems.
- nfs4-acl-tools 0.3.6, 0.3.7, and 0.4.1 have the problem of showing the error message `Invalid filename` for all the errors that are returned from `stat()`. For example, if a directory does not have execute permission for user A and user A invokes `nfs4_getfact` on a file within the directory, `Invalid filename` is returned instead of `Permission denied`. Version 0.3.7 is included in Ubuntu 22.04.
- nfs4-acl-tools 0.3.5 has the problem of showing undocumented flags O, G, and E, which can be ignored. For example, if the parent directory has ACE `A:d:gpfsuser:rtncy`, a subdirectory created under it will inherit the ACE with the additional O flag `A:do:gpfsuser:rtncy`, the O flag is not part of the protocol and can be ignored. This version is included in RHEL 8 and RHEL 9.
- nfs4-acl-tools 0.3.3 can run into segmentation faults. This version is not supported. This is included in RHEL 7, Ubuntu 20.04, SLES15 SP3, and SLES15 SP4.
- Listing the extended attributes through `listxattr` does not include the system.nfs4_acl attribute. This is to avoid redundant work in case of preserving xattrs with `cp (cp --preserve=xattr)`, as the system.gpfs_nfs4_acl extended attribute already exists. The system.nfs4_acl attribute can still be retrieved with `getxattr`.
- In addition to strings, numeric IDs can also be accepted for the principal in a NFSv4 access control entry.
- To add new ACL entries to inodes that do not have NFSv4 ACL set and for inodes that have only mode bits, an equivalent NFSv4 ACL is created from the mode bits and returned when retrieving the attribute.
- When the file system allows only NFSv4 authorization (`-k nfs4`), and the inode has a POSIX ACL, a conversion from POSIX to NFSv4 for the ACL is performed when retrieving the attribute.
- When the file system allows both POSIX and NFSv4 authorizations (`-k all`) and the inode has a POSIX ACL, no conversion from POSIX to NFSv4 for the ACL is done when retrieving the attribute, and the query is denied. In this case, change the permission to mode bits or a NFSv4 ACL first before attempting to use the tools in the nfs4-acl-tools package.

Q2.41:

Does IBM Storage Scale support the secure boot specified by the Unified Extensible Firmware Interface (UEFI)?

A2.41:

IBM Storage Scale supports secure boot in the following ways:

1. For RHEL, the following scenarios are supported:
 - a. IBM provides signed kernel modules together with a validation key starting with RHEL 9.2. on x86_64. For more information, see [Signed kernel modules for UEFI secure boot](#).
 - b. The kernel modules are signed manually by the customers with their key.
 - c. Secure boot is disabled.
2. Secure boot is not supported with SLES or Ubuntu, which means that it has to be disabled in the BIOS.

Q2.42:

Is IBM Storage Scale supported on Advanced RISC Machine (ARM) architectures?

A2.42:

Yes. Starting with IBM Storage Scale 5.2.0, IBM Storage Scale is supported on ARM architectures. Before 5.2.0 was released, a technology preview was made available with IBM Storage Scale 5.1.9.0 through 5.1.9.4 (for more information, see [ARM technology preview for IBM Storage Scale 5.1.9](#)). IBM Storage Scale has been developed for ARM architectures with an instruction set of at least ARMv8.2-A.

IBM Storage Scale has been tested on systems based on Ampere Altra and Nvidia Grace Hopper processors. IBM Storage Scale has been successfully installed on AWS Graviton 2, AWS Graviton 3, and systems based on Fujitsu A64FX.

For information about the supported operating systems and kernels, see [What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)

The following features are **not** supported by release 5.2.0.x:

- NSD server (support starts with 5.2.1)
- IBM Storage Scale GUI (ARM nodes can be managed in the GUI but the GUI cannot run on an ARM node)

- Offline files (HSM)
- Kernel signing
- IBM Storage Scale GPFS native RAID (GMR); that is, Erasure Code Edition (ECE).
- Protocol servers: SMB, NFS, CNFS, Object, BDA

The packages for ARM can be recognized by the platform token “aarch64” or “arm64”. For example, `gpfs.base-5.1.9-0.aarch64.rpm` or `gpfs.base_5.2.0-0_arm64.deb`.

ARM technology preview for IBM Storage Scale 5.1.9

With IBM Storage Scale 5.1.9.0 through 5.1.9.4, IBM Storage Scale was supported on ARM architectures as a technology preview. Meaning, it was supported only for non-production environments. The ARM technology preview is superseded by the support made available with the 5.2.0 release; thus, no packages for ARM technology preview are available for 5.1.9.5 onward. Information about this technology preview is still available on the next IBM Support page: [IBM Storage Scale Technical Preview: Support for ARM based processors](#).

Q2.43:

What are the requirements and limitations for using the asymmetric data replication function?

A2.43:

Currently, the use of asymmetric data replication for file systems is not supported.

The feature will be supported for specific configurations over time. The first use of asymmetric data replication for file systems will be supported as hyper store in an upcoming IBM Storage Scale System release. As new configurations become supported, this section of the FAQ will be updated to provide details.

Q2.44:

What are the options that I can use with the `mmsmb config` command?

A2.44:

The following options can be used with the `mmsmb config` command:

- `deadtime`
- `gpfs:dfreequota`
- `min domain uid`
- `restrict anonymous`
- `server signing`
- `server string`
- `server smb encrypt`

For more information about the `mmsmb config` command, see [mmsmb command](#).

Machine questions

Q3.1:

What are the minimum hardware requirements for an IBM Storage Scale cluster?

A3.1:

The minimum hardware requirements are:

- IBM zSystems:
 - 2 vcpu with 2 GB memory.
 - IBM Storage Scale V5.1.2 or later : Family z13 or later.
 - IBM Storage Scale V5.1.1 or earlier : Family z196 or later.
- IBM Storage Scale V5 on Power:
 - AIX on Power is supported by IBM POWER8 processors, or higher processors , supported by your level of AIX, with a minimum of 2 GB of system memory.
 - Linux on Power is supported on IBM POWER8, or higher processors, with a minimum of 2 GB of system memory.
 - POWER9 and Power10 CPUs support a new Radix MMU mode. The Linux kernel can use this MMU mode to implement additional restrictions for memory access, called Kernel Userspace Access Prevention (KUAP). IBM Storage Scale releases 5.1.3 and 5.1.2.5 have proper support for this feature. Earlier IBM Storage Scale releases require a workaround. The only affected Linux

distribution by earlier IBM Storage Scale releases is Ubuntu 20.04. To verify whether the POWER9 and POWER10 feature is active, run the following command:

```
grep Radix /proc/cpuinfo
```

To check whether KUAP is active, run `dmesg` or check the syslog for the following message:

```
radix-mmu: Activating Kernel Userspace Access Prevention
```

If KUAP is active, modify the boot loader of your Linux distribution to pass the `nosmap` parameter to the Linux kernel. Then reboot and again run the above checks. Only when KUAP is no longer active, attempt to start IBM Storage Scale.

Note: For more information about specific Power processor and operating system version requirements, see [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#).

- IBM Storage Scale V4/IBM GPFS V3 or later on x86 architecture is supported on:
 - Intel 64 processors, with 2 GB of memory.
 - AMD Opteron processors, with 2 GB of memory. Other AMD x86-64 processors are supported as long as they are completely compatible with AMD Opteron, and as long as the SMP scaling limit is not exceeded. For more information, see [5.3 What is the current maximum tested limit for SMP scaling?](#).

Additionally, it is highly suggested that a sufficiently large amount of swap space is configured. While the actual configuration decisions should be made taking into account the memory requirements of other applications, it is suggested to configure at least as much swap space as there is physical memory on a given node.

IBM Storage Scale is supported on systems which are listed in, or compatible with, the IBM hardware specified in the *Hardware requirements* section of the Sales Manual for IBM Storage Scale.

To access the Sales Manual:

1. Go to the **Recent sales manual** section in the [IBM Announcements](#) site.
2. On Information Type, choose **HW&SW Desc (sales manual,RPQ)**.
3. For IBM Storage Scale V5, choose the corresponding product number to enter in the **Search for** field:
 - IBM Storage Scale Data Access Edition: **5737-I39** (Passport Advantage), **5641-DA1**, **5641-DA3**, **5641-DA5** (eConfig/AAS)
 - IBM Storage Scale Data Management Edition: **5737-F34** (Passport Advantage), **5641-DM1**, **5641-DM3**, **5641-DM5** (eConfig/AAS)
 - IBM Storage Scale Erasure Code Edition: **5737J34** (Passport Advantage)
4. The Hardware Requirements section, which is part of the Technical Description section.

Q3.2: On what servers is IBM Storage Scale supported?

A3.2:

1. IBM Storage Scale for Linux on Z is supported:
 - a. with the distributions and kernel levels as listed in the question [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)
 - b. on servers that meet the minimum hardware model requirements as listed in the question [What are the minimum hardware requirements for a IBM Storage Scale cluster?](#)
2. IBM Storage Scale for AIX is supported:
 - a. with levels of AIX as listed in the question [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)
 - b. on servers that meet the minimum hardware model requirements as listed in the question [What are the minimum hardware requirements for a IBM Storage Scale cluster?](#)

Note: IBM Storage Scale runs on the POWER8 processor in default, POWER8, and POWER7 compatibility modes.

3. IBM Storage Scale for Linux on Power is supported:
 - a. with the distributions and kernel levels as listed in the question [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)
 - b. on servers that meet the minimum hardware model requirements as listed in the question [What are the minimum hardware requirements for a IBM Storage Scale cluster?](#)

Note: IBM Storage Scale runs on the POWER8 processor in default, POWER8, and POWER7 compatibility modes.

4. IBM Storage Scale for Linux on x86 Architecture is supported:
 - a. with the distributions and kernel levels as listed in the question [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)

- b. on servers that meet the minimum hardware model requirements as listed in the question [What are the minimum hardware requirements for a IBM Storage Scale cluster?](#)
5. IBM Storage Scale for Windows on x86 Architecture is supported:
- a. with the levels of Windows Server as listed in the question [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)
 - b. on servers that meet the minimum hardware model requirements as listed in the question [What are the minimum hardware requirements for a IBM Storage Scale cluster?](#)

Q3.3:
What interconnects are supported for GPFS daemon-to-daemon communication in a GPFS cluster?

A3.3:
 The interconnect for GPFS daemon-to-daemon communication depends upon the types of nodes in your cluster.

Note: This table provides the list of communication interconnects which have been tested by IBM and are known to work with GPFS. Other interconnects may work with GPFS but they have not been tested by IBM. The IBM Storage Scale support team will help customers who are using interconnects that have not been tested to solve problems directly related to GPFS, but will not be responsible for solving problems deemed to be issues with the underlying communication interconnect's behavior including any performance issues exhibited on untested interconnects.

<i>Table 32: GPFS daemon -to-daemon communication interconnects</i>		
Nodes in your cluster	Supported interconnect	Supported environments
Linux(x86 , Power)/AIX/Windows	Ethernet	All supported IBM Storage Scale environments
	10-Gigabit Ethernet 40-Gigabit Ethernet	All supported IBM Storage Scale environments
	InfiniBand	All supported IBM Storage Scale environments IP only
Linux (x86 and Power)	Ethernet	All supported IBM Storage Scale environments
	10-Gigabit Ethernet 40-Gigabit Ethernet	All supported IBM Storage Scale environments
	InfiniBand	IP and optionally VERBS RDMA Note: See the question Are there any considerations when utilizing the Remote Direct Memory Access (RDMA) offered by InfiniBand?
Linux on Z	Ethernet 10-Gigabit Ethernet 25-Gigabit Ethernet Hipersockets Vswitch (z/VM only) GuestLAN (z/VM only) ISM devices (SMC-D)	All supported IBM Storage Scale environments

Table 32: GPFS daemon -to-daemon communication interconnects (continued)

Nodes in your cluster	Supported interconnect	Supported environments
AIX	Ethernet	All supported IBM Storage Scale environments
	10-Gigabit Ethernet 40-Gigabit Ethernet	All supported IBM Storage Scale environments
	InfiniBand	All supported IBM Storage Scale environments IP only
Windows	Ethernet	All supported IBM Storage Scale environments
	10-Gigabit Ethernet 40-Gigabit Ethernet	All supported IBM Storage Scale environments
	InfiniBand	All supported IBM Storage Scale environments IP only

Disk questions

Q4.1:

What disk hardware has IBM Storage Scale been tested with?

A4.1:

This set of tables displays the set of disk hardware which has been tested by IBM and known to work with IBM Storage Scale. Other disk devices may work with IBM Storage Scale using NSD disk leasing, though they have not been tested by IBM. The IBM Storage Scale support team will help customers who are using devices outside of this list of tested devices, using NSD disk leasing only, to solve problems directly related to IBM Storage Scale, but will not be responsible for solving problems deemed to be issues with the underlying device's behavior including any performance issues exhibited on untested hardware. Untested devices should not be used with GPFS assuming SCSI-3 PR as the fencing mechanism, since our experience has shown that devices cannot, in general, be assumed to support the SCSI-3 Persistent Reserve modes required by GPFS.

These test statements apply to all current releases of IBM Storage Scale unless specified otherwise.

It is important to note that:

- Each individual disk subsystem requires a specific set of device drivers for proper operation while attached to a host running GPFS. The prerequisite levels of device drivers are not documented in this GPFS-specific FAQ. Refer to the disk subsystem's web page to determine the currency of the device driver stack for the host's operating system level and attachment configuration.

For information on IBM disk storage subsystems and their related device drivers levels and Operating System support guidelines, go to www.ibm.com/servers/storage/support/disk/index.html

- Microcode levels should be at the latest levels available for your specific disk hardware.

For the IBM System Storage, go to www.ibm.com/servers/storage/support/allproducts/downloading.html

DS4000 customers: Please also see

- The IBM TotalStorage DS4000 Best Practices and Performance Tuning Guide at publib-b.boulder.ibm.com/abstracts/sg246363.html?Open
- For the latest firmware and device driver support for DS4100 and DS4100 Express Midrange Disk System, go to <http://www.ibm.com/systems/support/supportsite.wss/selectproduct?brandind=5000028&familyind=5329597&osind=0&oldbrand=5000028&oldfamily=5345919&oldtype=0&taskind=2&matrix=Y&psid=dm>
- For the latest storage subsystem controller firmware support for DS4200, DS4700, DS4800, go to:
 - <https://www.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-5075581&brandind=5000028>
 - <https://www.ibm.com/systems/support/supportsite.wss/docdisplay?Indocid=MIGR-5073716&brandind=5000028>

Table 33: Disk hardware tested for AIX on Power

XtremIO 4.0.10

VMAX 3 Hudson 5977.810.784 and Trinity 5977.932.887

AIX 6.1 TL9 SP6, AIX 7.1 TL4 and AIX 7.2 or later, with IBM Storage Scale V4.1.0.0 or later

IBM FlashSystem® 900 Minimal Firmware Level: 1.2.0.11

This storage subsystem has been tested on
AIX 7.1.3.16 with IBM Storage Scale V4.1.0.8 or later

IBM FlashSystem 840

Minimal Firmware Level: 1.1.1.2
AIX 6.1(6100-09) and AIX 7.1(7100-02-03-1334) with GPFS V3.5.0.19 or later,
and IBM Storage Scale V4.1 or later

IBM FlashSystem 820

Minimal Firmware Level: 6.3.0.6

AIX 6.1(6100-06) and AIX 7.1(7100-01) with GPFS V3.5.0.11 or later, and IBM Storage Scale V4.1 or later

IBM Storwize® V7000/V3500/V3700/SVC

Note: Placing GPFS metadata on thinly provisioned or compressed volumes is not supported.

AIX 6.1 and 7.1 with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later

IBM XIV® 2810

Minimum Firmware Levels: 10.1, 10.2
AIX 6.1 and 7.1 with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later
For more information, directions and recommended settings for attachment please refer
to the latest Host Attach Guide for Linux located at the IBM XIV Storage System Knowledge Center go to
<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

IBM System Storage DS8000® using SDDPCM

IBM System Storage DCS3700

AIX 6.1 and 7.1 with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later

IBM System Storage DS5000 all supported expansion drawer and disk types including SSD

This includes models: DS5100, DS5300 and DS5020 Express.

on AIX V7.1 with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later

on AIX V6.1 with a minimum level of TL2 with SP2 and APAR IZ49639 GPFS V3.5 or later,
and IBM Storage Scale V4.1 or later

Firmware levels:

7.60.28.00

7.83.22.00

7.77.34.00

IBM System Storage DS3400 (1726-HC4)

See <http://www-07.ibm.com/storage/includes/pdf/TSD03007USEN.pdf>.

IBM TotalStorage ESS (2105-F20 or 2105-800 with SDD)

IBM System Storage Storage Area Network (SAN) Volume Controller (SVC) V2.1 and V3.1

See www.ibm.com/support/docview.wss?rs=591&uid=ssg1S1002471 for specific advice on SAN Volume Controller recommended software levels.

Table 33: Disk hardware tested for AIX on Power (continued)

Hitachi Virtual Storage Platform (VSP G200, G400, G600, G800, F400, F600, F800)
Hitachi Virtual Storage Platform (VSP G350, G370, G700, G900, F350, F370, F700, F900)
Hitachi Virtual Storage Platform (VSP E590, E790, E990, E1090, E590H, E790H, E1090H)
Hitachi Virtual Storage Platform (VSP G1000, G1500, F1500)
Hitachi Virtual Storage Platform (VSP 5100, 5200, 5500, 5600, 5100H, 5200H, 5500H, 5600H)

Note: These tests were conducted on:

- AIX 7.2 TL03 or later, AIX 7.3 or later
- IBM Storage Scale V5.0.2.3 or later
- HDLM V8.7.0 or later, built-in multipath MPIO

Hitachi Universal Storage Platform (USP V)
Hitachi Adaptable Modular Storage (AMS)
- AMS Series (includes 2100, 2300 and 2500 models)

Note:

1. In all cases Hitachi Dynamic Link Manager™ (HDLM) (multipath software) or MPIO is required.
2. Hitachi Vantara (HV) was previously known as Hitachi Data Systems (HDS).
3. MPIO-ODM packages for AIX supplied by HV are required for all the listed USP V and AMS devices.
4. Customers should consult with HV to verify that their proposed combination of the listed components is supported by HV.

EMC Symmetrix VMAX and DMX Storage Subsystems (FC attach only)

Device driver support for Symmetrix includes both MPIO and PowerPath.

Selected models of EMC CLARiiON CX/CX-3 family including CX300, CX400, CX500 CX600, CX700 and CX3-20, CX3-40 and CX3-80

Note: CX/CX-3 requires PowerPath.

See <http://www.emc.com>.

Customers should consult with EMC to verify that their proposed combination of the above components is supported by EMC.

HP XP 128/1024
HP StorageWorks Enterprise Virtual Arrays (EVA) 4000/6000/8000
and 3000/5000 models that have been upgraded to active-active
configurations

Note: HDLM multipath software is required

HPE 3PAR OS 3.3.1
Minimum Firmware Level: HPE 3PAR OS 3.3.1
RHEL 6.7 with GPFS V4.2.3.4 or later

Table 34: Disk hardware tested with Linux on x86 servers

Hitachi Virtual Storage Platform (VSP)
Hitachi Storage Platform (VSP G200,G400,G600,G800, microcode 83-03-24-00/00)
Hitachi Virtual Storage Platform (VSP) (F400, F600, F800, microcode 83-03-24-00/00)
Hitachi Storage Platform (VSP G1000, VSP G1500, microcode 80-04-22-00/00 or later)
Hitachi Unified Storage VM

RHEL 7.1 or later, with IBM Storage Scale V4.2.1.1 or later

[IBM FlashSystem A9000/A9000R](#) Firmware version 12.0.1

RHEL 7.2 or later, with IBM Storage Scale V4.2.1.0 or later

See [Q4.12 What are the considerations for using thinly provisioned or compressed volumes with GPFS?](#) for considerations on thin provisioning and compression on these storage subsystems.

[IBM FlashSystem 900](#) Minimal Firmware Level: 1.2.0.11

RHEL 6.5, or later, and SLES 11 SP3, or later, with IBM Storage Scale V4.1.0.8 or later

Table 34: Disk hardware tested with Linux on x86 servers (continued)

IBM FlashSystem 840

Minimum Firmware Level: 1.1.1.2

RHEL 6.5 and SLES 11 SP with GPFS V3.5.0.20 or later, and IBM Storage Scale V4.1.0.2 or later

IBM FlashSystem 820

Infiniband and FC attach

Minimal Firmware Level: 6.3.1.0

RHEL 6.3 or later and SLES 11 SP 1 or later, with GPFS V3.5.0.11 or later

RHEL 6.3 or later and SLES 11 SP 2 with IBM Storage Scale V4.1 or later

IBM Flex Systems Storwize V7000

Firmware Level: SVC 6.4.1.4

RHEL 6.4 or later, and SLES 11 SP2 or later, using multipath Device Mapper

GPFS V3.5.0.9 or later, and IBM Storage Scale V4.1 or later

RHEL 5.9 and SLES 10 SP4

GPFS V3.5.0.9 or later

IBM Storwize V7000/V3500/V3700/SVC

RHEL 6.x and 5.x with levels of GPFS that support the distribution

SLES11 SP2 or later with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later

SLES 10 SP1 and SP2 with GPFS V3.5 or later

IBM XIV 2810

Minimum Firmware Level: 10.0.1

RHEL5.1 and greater with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later

SLES 10.2 with GPFS V3.5 or later, and IBM Storage Scale V4.1 or later

For more information, directions and recommended settings for attachment please refer to the latest

Host Attach Guide for Linux located at the IBM XIV

Storage System Knowledge Center go to

<http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp>

IBM System Storage DS8000

IBM System Storage DS3500/DS5300

RHEL 6.2 or later and SLES 11 SP2 or later with IBM Storage Scale V4.1 or later

RHEL 6.x, 5.x with GPFS V3.5 or later

SLES 11, 10 with GPFS V3.5 or later

Firmware level 7.84.44.00

IBM System Storage DCS3700

RHEL 6.0, 5.6 and 5.5

SLES 11.1, 10.4 and 10.3

IBM System Storage DS5000 all supported expansion drawer and disk types including SSD

This include:

Models: DS5100, DS5300 and DS5020 Express.

Firmware levels:

7.60.28.00

7.83.22.00

7.77.34.00

IBM System Storage DS3400 (1726-HC4)

Table 34: Disk hardware tested with Linux on x86 servers (continued)

IBM TotalStorage Enterprise Storage Server® (ESS) models 2105-F20 and 2105-800, with [Subsystem Device Driver \(SDD\)](#)

EMC ScaleIO V2.0 without persistent reserve

RHEL 7.1, or later, with IBM Storage Scale V4.2.0.1, or later

EMC Symmetrix Direct Matrix Architecture (DMX) Storage Subsystems 1000 with PowerPath v 3.06 and v 3.07

IBM System Storage Storage Area Network (SAN) Volume Controller (SVC) V2.1 and V3.1

See www.ibm.com/support/docview.wss?rs=591&uid=s59151002471 for specific advice on SAN Volume Controller recommended software levels.

IBM DCS9550 (either FC or SATA drives)

FC attach only

minimum firmware 3.08b

QLogic drivers at 8.01.07 or newer and

IBM SAN Surfer V5.0.0 or newer

IBM DCS9900 (either FC or SATA drives)

FC attach only

Table 35: Disk hardware tested with Linux on Power

IBM FlashSystem™ 900

Minimal Firmware Level: 1.2.0.11

RHEL 6.5 or later, and SLES 11 SP3 or later, with IBM Storage Scale V4.1.0.8 or later

IBM FlashSystem 840

Minimal Firmware Level: 1.1.1.2 using default DM-MP

RHEL 6.4 or later, and SLES 11 SP2 or later

GPFS V3.5.0.18 or later, and IBM Storage Scale V4.1 or later

IBM Flex System V7000

Firmware Level: SVC 6.4.1.4, 7.1, 7.2

RHEL 6.4 or later and SLES 11 SP2 or later using multipath Device Mapper

IBM Storage Scale V4.1 or later

RHEL 5.9 or later, RHEL 6.4 or later, SLES 10 SP4 or later, SLES 11 SP2 or later, using multipath Device Mapper with GPFS V3.5.0.9 or later

IBM Storwize V7000/V3500/V3700/SVC

Note: Placing GPFS metadata on thinly provisioned or compressed volumes is not supported.

RHEL 6.x and 5.x with levels of GPFS that support the distribution

SLES11 SP2 or later, with IBM Storage Scale V4.1 or later

SLES 10 SP1 or later, and SLES11 SP1 or later with GPFS V3.5 or later

IBM System Storage DS3500/DS5300

RHEL 6.2 or later and SLES 11 SP2 or later with IBM Storage Scale V4.1 or later

RHEL 6.x, 5.x with GPFS V3.5 or later

SLES 11, 10 with GPFS V3.5 or later

Firmware level 7.84.44.00

IBM System Storage DCS3700

RHEL 6.0, 5.6 and 5.5

SLES 11.1, 10.4 and 10.3

Table 35: Disk hardware tested with Linux on Power (continued)

IBM System Storage DS5000 all supported expansion drawer and disk types including SSD

This include:

Models: DS5100, DS5300 and DS5020 Express.

Firmware levels:

7.60.28.00

7.83.22.00

7.77.34.00

IBM System Storage DS8000

Table 36: Disk hardware tested with IBM Storage Scale for Linux on Z

HDS USPV/VSP/VSP G1000 without persistent reserve

EMC without persistent reserve

IBM DS8000 Series

IBM FlashSystem

IBM Storwize V7000

IBM XIV

IBM SVC

IBM Elastic Storage Server (ESS) V4.5.0

Note:

1. Ensure you are running with the latest firmware levels available.
2. DS4000 is *not* supported
3. See the question [Does IBM Storage Scale for Linux on Z support Direct Attached Storage Devices \(DASD\)?](#)
4. IBM Storage Scale for Linux on Z is supported with EMC without persistent reserve. Customers should consult with EMC to verify if their proposed solution is supported by [EMC](#).

Q4.2:

What Fibre Channel Switches are qualified for IBM Storage Scale usage and is there a FC Switch support chart available?

A4.2:

There are no special requirements for FC switches used by IBM Storage Scale other than the switch must be supported by AIX or Linux or Windows. For further information see www.storage.ibm.com/ibmsan/index.html

Q4.3:

Can I concurrently access SAN-attached disks from both AIX and Linux (x86 and Power) nodes in my IBM Storage Scale cluster?

A4.3:

While the architecture of IBM Storage Scale would generally allow LUNs to be shared between different operating systems (Linux (x86 and Power), AIX, and Windows), the actual implementation of various OS specific features preclude this from being exploited at the current time. There are differences in how disks are labeled, how partitions are created and managed, and how multi-pathing managers react to error conditions between the various OS such that this support isn't offered in IBM Storage Scale today.

For IBM Storage Scale for Linux on Z, SAN-attached disks can only be accessed from Linux on Z cluster nodes.

Q4.4:

What devices does IBM Storage Scale support with SCSI-3 Persistent Reservations?

A4.4:

The following devices are supported with SCSI-3 Persistent Reservations:

- EMC XtremIO 4.0.10, VMAX 3 Hudson 5977.810.784 and Trinity 5977.932.887 using Native MPIO running AIX 6.1 TL9 SP6, AIX 7.1 TL4 and AIX 7.2 or later, through Fiber Channel connection (IBM Storage Scale V4.1.0.0 or later).
- Hitachi Storage Platform (VSP G1000, G1500, F1500) (microcode 80-04-22-00/00 or later) using default DM-MP on x86 Linux running RHEL 7.1 or later, through Fiber Channel connection (IBM Storage Scale V4.2.1.1 or later).
- Hitachi Storage Platform (VSP 5100, 5200, 5500, 5600, 5100H, 5200H, 5500H, 5600H) using default DM-MP on Power Linux running RHEL 8.1 (Power) or later, through Fiber Channel connection (IBM Storage Scale V5.1.0 or later).

- Hitachi Virtual Storage Platform (VSP G200, G400, G600, G800, F400, F600, F800, G350, G370, G700, G900, F350, F370, F700, F900, E590, E790, E990, E1090, E590H, E790H, E1090H, VSP G1000, G1500, F1500, 5100, 5200, 5500, 5600, 5100H, 5200H, 5500H, 5600H) (for the microcode, refer back to Hitachi Vantara for supported microcode levels), using MPIO and HDLM V8.7.0 or later on AIX 7.2 TL03, AIX 7.3 or later, through Fiber Channel connection (IBM Storage Scale V5.0.2.3 or later).
- **DCS3700** (firmware 08.20.12.00) using IBM RDAC driver or the default DM-MP on x86 Linux running RHEL7.2, SLES 11 SP3 (IBM Storage Scale V4.2.0.3 or later)
- **IBM FlashSystem A9000/A9000R** (firmware version 12.0.1) using default DM-MP on x86 Linux running RHEL7.2, or later through Fiber Channel connection (IBM Storage Scale V4.2.1.0 or later)
- **IBM FlashSystem 900** (firmware 1.2.0.11) on x86 Linux running RHEL 6.5, or later, or SLES 11 SP3, or later, through Infiniband or Fiber Channel connection
- **IBM FlashSystem 900** (firmware 1.2.0.11) on Power Linux running RHEL 6.5, or later, or SLES 11 SP3, or later, through Infiniband or Fiber Channel connection
- **IBM FlashSystem 900** (firmware 1.2.0.11) using default AIX PCM on AIX 7.1.3.16 through Fiber Channel and Infiniband connections (GPFS V4.1.0.8 or later)
- **IBM FlashSystem 840** (firmware 1.1.1.2) on x86 Linux running RHEL 6.5, or later, or SLES 11 SP3, or later, through Infiniband or Fiber Channel connection
- **IBM FlashSystem 840** (firmware 1.1.1.2) on Power Linux running RHEL 6.5, or later, or SLES 11 SP3, or later, through Infiniband or Fiber Channel connection
- **IBM FlashSystem 840** (firmware 1.1.1.2) using default AIX PCM on AIX 6.1.9.0 and AIX 7.1.2.3 through Fiber Channel and Infiniband connections (GPFS V3.5.0.19 or later and IBM Storage Scale V4.1.0.0 or later)
- **IBM Storwize V7000** (firmware SVC 7.1.0.3) using default DM-MP on Power Linux running RHEL 6.4, or later, or SLES 11 SP2 (GPFS V3.5.0.16 or later and GPFS V3.4.0.20 or later)
- **IBMIB Flex V7000** (firmware SVC 6.4.1.4) using SDDPCM (2.6.3.2) on AIX 6.1.8 or AIX 7.1.2 (GPFS V3.5.0.16 or later and GPFS V3.4.0.20 or later)
- **IBMIB Flex V7000** (firmware SVC 6.4.1.4) using default DM-MP on Power Linux running RHEL 5.9/6.4 or SLES 10.4/11.2 (GPFS V3.5.0.16 or later and GPFS V3.4.0.20 or later)
- **IBMIB Flex V7000** (firmware SVC 7.1.0.0) using default DM-MO on x86 Linux running RHEL 5.9/6.4 or SLES 10.4/11.2 (GPFS V3.5.0.16 or later and GPFS V3.4.0.20 or later)
- **IBM Storwize V7000** (firmware SVC 7.1.0.3) using SDDPCM on AIX 6.1.0 and AIX 7.1.0 through Fiber Channel connection
- **IBM FlashSystem 820** (firmware 6.3.1 SP1) using default DM-MP on Power Linux running RHEL 6.3 or SLES 11 SP2 through Infiniband or Fiber Channel connections (GPFS V3.5.0.21 or later, and IBM Storage Scale V4.1.0.4 or later)
- **IBM FlashSystem 820** (firmware 6.3.1 SP1) on x86 Linux running RHEL 6.3 or SLES 11 SP2 through Infiniband or Fiber Channel connection (GPFS V3.5.0.21 or later, and IBM Storage Scale V4.1.0.4 or later)
- **IBM FlashSystem 820** using default AIX PCM on AIX 6.1.0 and AIX 7.1.0 through Fiber Channel connection (GPFS V3.5.0.21 or later, and IBM Storage Scale V4.1.0.4 or later)
- **IBM Storwize V7000/V3500/V3700** (SVC firmware) on x86 Linux running SLES 10 SP4, SLES 11 SP2, RHEL 5.8, or RHEL 6.2
- DS5000 using SDDPCM or the default AIX PCM on AIX
- DS8000 (all 2105 and 2107 models) using SDDPCM or the default AIX PCM on AIX
- DS4000 subsystems using the IBM RDAC driver and AIX MPIO on AIX. (devices.fcp.disk.array.rte or MPIO)
- DS3500 using IBM RDAC driver or the default DM-MP on Linux
- DS4800 using IBM RDAC driver or the default DM-MP on Linux
- DS5020 using IBM RDAC driver or the default DM-MP on Linux
- DS5300 using IBM RDAC driver or the default DM-MP on Linux
- DS8000(2107 models) using IBM SDD driver or the default DM-MP on Linux
- EMC VMAX using EMC PowerPath 5.5 P04 B003 and EMC AIX ODM Package 5.3.0.6.

Please check EMC PowerLink for support details, and consult EMC to verify that the proposed configuration is supported by EMC.

Note: The use of AIX MPIO is also supported in this environment.

- HPE 3PAR (HPE 3PAR OS 3.3.1) using the default DM-MP on x86 Linux running RHEL6.7 (IBM Storage Scale V4.2.3.4 or later)

The most recent versions of the device drivers are always recommended to avoid problems that have been addressed.

Note: For a device to properly offer SCSI-3 Persistent Reservation support for GPFS, it must support **SCSI-3 PERSISTENT RESERVE IN** with a service action of **REPORT CAPABILITIES**. The **REPORT CAPABILITIES** must indicate support for a reservation type of **Write Exclusive All Registrants**. Contact the disk vendor to determine these capabilities.

Also see the question [Are there any requirements for Persistent Reserve support in GPFS ?](#)

Q4.5:

What considerations are there when setting up DM-MP multipath service?

A4.5:

To setup the DM-MP multipath service, depending on node distribution and storage controller firmware level, you may need to modify the `/etc/multipath.conf` file to fit your individual storage requirements. A default copy of the `multipath.conf` file can be copied from the `/usr/share/doc` directory.

As an example, the following attributes are tested with the IBM products DS3500 (1746), DS5020 (1814), DS4800 (1815), and DS5300 (1818) :

```
device {
    vendor                "IBM"
    product               "1746"
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout          "/sbin/mpath_prio_rdac /dev/%n"
    features               "0"
    hardware_handler      "1 rdac"
    path_selector         "round-robin 0"
    path_grouping_policy  group_by_prio
    failback              immediate
    rr_weight              uniform
    no_path_retry         fail
    rr_min_io             1000
    path_checker          rdac
}
```

Note: In order for GPFS failover to take place, the following steps must be taken:

1. The following parameters must be set:

- features "0"
- failback immediate
- no_path_retry fail

2. The `mmnsddiscover -a` command must be issued in order for the NSD server to rediscover the disks.

Additionally, see

- The IBM Storage Scale documentation at <https://www.ibm.com/docs/en/spectrum-scale>
- Please refer to each distribution's multipath document for details. For instance:
 - https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/6/html/DM_Multipath/
 - https://www.suse.com/documentation/sles11/stor_admin/?page=/documentation/sles11/stor_admin/data/bookinfo.html

Q4.6:

Are there any steps that need to be taken before disks are used by IBM Storage Scale on AIX?

A4.6:

Yes. Most of the following specifics are for IBM disks. If you have non-IBM disks, comments below help explain how you would need to adjust the commands that are shown. The `lsattr` and `chdev` commands are used for all disk types:

- Set all disks that will be used as NSDs to the `no_reserve` reservation policy.
 - This may require `chdev` commands to modify the attributes if they are found not set to `no_reserve` as the default.
 - The `reserve_policy` needs to be checked on all nodes with access to the disks even if they are not specially coded as NSD servers.
 - ◆ The `lsattr -El hdiskX -a reserve_policy` command should show `no_reserve`
 - ◆ Issue the `chdev -El hdiskx -a reserve_policy=no_reserve` command to update if necessary.
- Please note this advice pertains to all disks being used by IBM Storage Scale on AIX, not just disks that will be used in a persistent reserve model.
- If there are any issues with the `chdev` commands contact AIX support for assistance.

Disks from other manufacturers should have their own unique attributes that similarly controls disk reserves. A specific example is the `reserve_lock` attribute, which needs to have the value `no`. The `lsattr -El device` command would show all the attributes that the device supports, along with the current value of each. If it is not obvious which attribute controls reserves on the disk, contact the manufacturer for that information. The `lsattr -R -l device -a attribute` command can be used to find out all the legal values for the specified attribute. For example `lsattr -R -l hdiskx -a reserve_policy`.

Q4.7:

Does IBM Storage Scale support Logical Volumes (LVs) ?

A4.7:

Logical Volumes (LVs) are minimally supported under these conditions:

- The customer must maintain LV availability. IBM Storage Scale does not support the management of the export/import or varying on/off of LVs between nodes.
- Conventional LVs can be used when only attached to a single node as a **descOnly** disk.
- Starting with GPFS 3.5.0.16, Concurrent Mirrored LVs can be used as **descOnly** disks.

Q4.8:

Does IBM Storage Scale support AIX raw hdisks (rhdisk) ?

A4.8:

Input to the **mmcrnsd** command requires the use of the hdisks format, not rhdisks. Internal logic converts the hdisk format to rhdisk.

Q4.9:

Does IBM Storage Scale for Linux on Z support Direct Attached Storage Devices (DASD)?

A4.9:

The DASD device driver provides access to real or emulated Direct Access Storage Devices (DASD) that can be attached to the channel subsystem of an IBM Z. This device driver supports the ECKD (Extended Count Key Data) and FBA (Fixed Block Access) devices.

Note: Prior to IBM Storage Scale V4.2.1, an ECKD device has to have all the same Bus-ID on all NSD server nodes.

To enable the usage of FBA devices, the cluster needs to run IBM Storage Scale V4.2.2 or later.

It is recommended to set the failfast parameter of the DASD device so the device driver immediately returns "failed" for an I/O operation when the last path to a DASD is lost. If the failfast parameter is not set, GPFS might hang until the path to the DASD is restored. See the Device Drivers, Features, and Commands documentation for your Linux distribution.

For more information about the device drivers, features and commands of your Linux platform, see [Device Drivers, Features, and Commands](#).

Q4.10:

Does IBM Storage Scale support 4K disk sectors?

A4.10:

Yes, 4K disk sector support requires IBM Storage Scale V4.1.0.5 or later. The following disk subsystems with 4K sector size have been tested by IBM:

- ECKD disk devices (Linux on Z only)
- [IBM FlashSystem 820](#).
- [IBM FlashSystem 840](#)
- [IBM FlashSystem 900](#)

Note: Other disk devices may work with IBM Storage Scale, though they have not been tested by IBM. See the question [What disk hardware has IBM Storage Scale been tested with?](#)

Q4.11:

What are the considerations for using block storage systems that support thinly provisioned volumes (thin provisioning)?

A4.11:

Block storage uses the following terminology:

Thin provisioning

Thin provisioning is the ability to create a volume without immediately allocating the requested space from the pool of usable space. Blocks are allocated to the volume only as required.

Over-provisioning

Over-provisioning occurs when more space is provisioned from a pool than the total amount of usable space in the pool. This is permissible with thin provisioning as a volume's space is not reserved from the usable space until it is needed, so the total amount of space that is provisioned to all the volumes can exceed the usable amount.

Note: When a pool is over-provisioned, a volume might not be allocated all of the space that is nominally provisioned if all of the usable space has already been allocated to other volumes.

Full allocation

Full allocation instructs the storage system to immediately provision all of the space that is requested for a volume. Fully allocated volumes eliminate the risk that is associated with over-provisioning such as being unable to allocate provisioned space when requested. This is sometimes called fully provisioned.

Note: Not all storage systems support full allocation.

Thin-provisioned disks are supported only through the IBM RPQ or SCORE process. Additionally, when using a storage system or storage pool that supports thin provisioning, the following conditions need to be satisfied:

- All nodes mounting or playing a management role in the file system should be at least at version 5.0.4, and the file system must be upgraded to file system format version 5.0.4 or later.
- The stanza file must include the following line to add thin disks into the file system:

```
thinDiskType={scsi | nvme}
```

- Thin-provisioned disks must be connected to nodes that are running the Linux operating system.

For more information, see the topic *IBM Storage Scale with data reduction storage devices* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q4.12:
What are the considerations for using IBM Storage Scale with block storage that supports data reduction features including compression or deduplication?

A4.12:
Note: IBM Storage Scale 5.0.4 has introduced support for data reduction storage devices, so configurations that operate with such devices should upgrade IBM Storage Scale to at least that release. As with thin provisioning, the support needs to go through the RPQ or SCORE process. For more information, see the topic *IBM Storage Scale with data reduction storage devices* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Metadata

It is critical that IBM Storage Scale does not unexpectedly run out of space to write or rewrite metadata due to block-level data reduction features. The system or storage administrator is responsible for allocating volumes in such a way that this cannot happen.

- Volumes that are used for metadata must be fully allocated and not use deduplication.
- Compression is permitted; however, you cannot rely on compression to accommodate more metadata than the usable capacity of the volume (for example, you should assume a compression ratio of 1:1 and allow for any overhead that is incurred by the storage system in managing compression).

Data

Deduplication and compression might be supported after a review by IBM. Ask your sales representative to contact IBM Storage Scale development through the RPQ process.

Q4.13:
How is IBM Storage Scale capacity determined for licensing purposes with thinly provisioned volumes or data reduction?

A4.13:
When using thinly provisioned volumes, the capacity to be licensed is the provisioned capacity presented as NSDs to IBM Storage Scale.

Note: If the storage pool is over-provisioned, the capacity to be licensed might be more than the usable capacity of the pool.

For fully allocated volumes, the provisioned capacity and the usable capacity of the NSD are equal. Data reduction does not affect IBM Storage Scale licensing. While it might increase the effective capacity of the storage system, it does not change the provisioned capacity.

Q4.14:
Does IBM Storage Scale for Linux on Z work with disk hardware replication?

A4.14:
Storage subsystems, such as IBM DS8000, offer data replication mechanisms. For example, Metro Mirror provides synchronous data replication; whereas, Global Mirror provides asynchronous data replication over distance. IBM Storage Scale supports disk hardware replication provided that the disks (FCP LUNs or ECKD volumes) are within a consistency group. The IBM Storage Scale configuration needs to be set up in a way that handles the different device addresses used for primary and secondary devices.

Disk replication can be managed by products such as IBM GDPS® (Geographically Dispersed Parallel Sysplex®) or IBM CSM (Copy Services Manager).

Q4.15:
Does IBM Storage Scale for Linux on Z support HyperSwap or SVC Stretch Clusters?

A4.15:
IBM Storage Scale supports HyperSwap, SVC Stretch Clusters, or similar technologies on IBM Z beginning in IBM Storage Scale V4.2.3.1. The actual swap from one device to another results in a pause to IO while the necessary storage level actions are taken. The `failureDetectionTime` and `leaseRecoveryWait` tunables need to be set accordingly. It is suggested to use a value of 1.5X the

expected pause time. The user should contact their provider to discuss expected IO pause time for their particular configuration as the actual values depend on the HyperSwap/stretch cluster technology, as well as the storage subsystem and configuration. For more information about HyperSwap with IBM DS8000, see the following white paper: [Hyperswaping IBM Spectrum Scale for Linux on z under GDPS Virtual Appliance management control](#).

Q4.16:
What disk connection technologies can be used for NSD storage?

A4.16:
 IBM Storage Scale NSDs can be created on arrays accessed through many different connection technologies including Fibre Channel SAN, SAS, InfiniBand SRP, iSCSI, NVMeOF, and other disk connection technologies.

Scaling questions

Q5.1:
What are the IBM Storage Scale cluster size limits?

A5.1:
 The current maximum tested IBM Storage Scale cluster size limits are:

<i>Table 37: IBM Storage Scale maximum tested cluster sizes</i>	
IBM Storage Scale for Linux (x86, Power, and IBM Z)	9620 nodes
IBM Storage Scale for AIX	1530 nodes
IBM Storage Scale for Windows on x86_64 Architecture	64 Windows nodes
FPO-enabled	732 nodes
IBM Storage Scale for Linux and IBM Storage Scale for AIX	3906 (3794 Linux nodes and 112 AIX nodes)
<p>Note: Contact scale@us.ibm.com if you intend to exceed:</p> <ul style="list-style-type: none"> • Configurations with Linux nodes exceeding 512 nodes. • Configurations with AIX nodes exceeding 128 nodes. • Configurations with Windows nodes exceeding 64 nodes. • FPO-enabled configurations exceeding 32 nodes, see 2.16 What are the current limitations for using the File Placement Optimizer (FPO) function? 	

Although IBM Storage Scale is typically targeted for a cluster with multiple nodes, it can also provide high performance benefit for a single node so there is no lower limit. For a given I/O configuration, typically multiple nodes are required to saturate the aggregate file system performance capability. If the aggregate performance of the I/O subsystem is the bottleneck, then IBM Storage Scale can help achieve the aggregate performance even on a single node.

Q5.2:
What are some scaling considerations for the protocol (CES) nodes?

A5.2:
 When planning for scaling the protocols function, consider the maximum supported or maximum recommended number of protocol nodes and client connections. For detailed information, see [Scaling considerations](#).

Q5.3:
What is the current maximum tested limit for SMP scaling?

A5.3:
 The largest SMP scale tested to date is 192 cores. The largest vCPU (hardware thread) count tested to date is 1536 total vCPUs. The largest NUMA Complexity metric tested to date is 3. The NUMA Complexity metric is the number of different **node distances** values as reported by `numactl --hardware` on Linux or **REF1** numbers as reported by `lssrad -av` on AIX. **This 1536 vCPU limit is a hard-coded and enforced GPFS limit.**

For the following Linux example, the `numactl --hardware` distinct **node distances** values are {10; 11}. The NUMA Complexity metric is therefore 2.

```
node distances:
node 0 1
0:   10 11
1:   11 10
```

Q5.4:

What is the current limit on the number of nodes that may concurrently join a cluster?

A5.4:

As of GPFS V3.4.0.18 and GPFS V3.5.0.5, the total number of nodes that may concurrently join a cluster is limited to a maximum of 16384 nodes.

A node joins a given cluster if it is:

- A member of the local GPFS cluster (the **mmlscluster** command output displays the local cluster nodes).
- A node in a different GPFS cluster that is mounting a file system from the local cluster.

For example:

- GPFS **clusterA** has 2100 member nodes as listed in the **mmlscluster** command.
- 500 nodes from **clusterB** are mounting a file system owned by **clusterA**.

In this example **clusterA** therefore has 2600 concurrent nodes.

Q5.5:

What is the limit of remote clusters that a client node can join?

A5.5:

The maximum number of remote clusters that a client node can join is 31 (32 when counting the local cluster).

Q5.6:

What is the limit of remote clusters that can join a local cluster?

A5.6:

There is not really a limit. The smallest cluster possible is a single node cluster, which means that 16,383 clusters can join a local cluster (16384 - 1).

Q5.7:

What are the current file system size limits?

A5.7:

The current file system size limits are:

<i>Table 38: Current file system size limits</i>	
GPFS 2.3 or later, file system architectural limit	2 ⁹⁹ bytes
GPFS 2.2 file system architectural limit	2 ⁵¹ bytes (2 Petabytes)

Q5.8:

What is the current limit on the number of mounted file systems in an IBM Storage Scale cluster?

A5.8:

The current limit on the number of mounted file systems in an IBM Storage Scale cluster is 256 on all supported OSs except for Windows. On Windows, the limit is the number of unused drives in the range A-Z.

Q5.9:

What is the architectural limit of the number of files in a file system?

A5.9:

The architectural limit of the number of files in a file system is determined by the file system format:

- For file systems created with GPFS V3.4 or later, the architectural limit is 2⁶⁴.
The current tested limit is 9,000,000,000.
- For file systems created with GPFS V2.3 or later, the limit is 2,147,483,648.

- For file systems created prior to GPFS V2.3, the limit is 268,435,456.

Please note that the effective limit on the number of files in a file system is usually lower than the architectural limit, and could be adjusted using the `mmchfs` command (GPFS V3.4 and later use the `--inode-limit` option).

Q5.10:
What is the architectural limit of the number of disks in a file system?

A5.10:
 The architectural limit of the number of disks in a file system is 2048.

Q5.11:
What are the limitations on IBM Storage Scale disk size?

A5.11:
 The maximum disk size is only limited by the OS kernel and device driver support

<i>Table 39: Maximum disk size supported</i>	
OS kernel	Maximum supported GPFS disk size
AIX, 64-bit kernel	>2TB, up to the device driver limit
Linux 2.6 64-bit kernels	>2TB, up to the device driver limit
Windows	>2TB, up to the device driver limit

- Note:**
1. On systems running with the Linux kernel 3.0, both `processor.max_cstate` and `intel_idle.max_cstate` should be set to zero.
 2. IBM Storage Scale supports 16MB file system block size.

Q5.12:
What is the limit on the maximum number of groups a user can be a member of when accessing a GPFS file system?

A5.12:
 Each user may be a member of one or more groups, and the list of group IDs (GIDs) that the current user belongs to is a part of the process environment. This list is used when performing access checking during I/O operations. Due to architectural constraints, GPFS code does not access the GID list directly from the process environment (kernel memory), and instead makes a copy of the list, and imposes a limit on the maximum number of GIDs that may be smaller than the corresponding limit in the host operating system. The maximum number of GIDs supported by GPFS depends on the platform and the version of GPFS code. Note that the GID list includes the user primary group and supplemental groups.

<i>Table 40: Maximum number of GIDs supported</i>	
Platform	Maximum number of GIDs supported
AIX	2,048 ¹
Linux with 4K page size (all supported platforms except the two below)	1,020
Linux with 64K page size (PPC64/RHEL5/RHEL6/RHEL 7 platforms)	16,380
Windows	Windows OS limit (no limit in GPFS code)

- Note:**
1. The table reflects the maximum value that can be achieved with IBM Storage Scale 5.0.0 or later and AIX 7.1 or later. On earlier versions of IBM Storage Scale or AIX, the limit is 128. For more information about configuring the Number of Groups allowed, see https://www.ibm.com/support/knowledgecenter/ssw_aix_71/com.ibm.aix.security/number_groups_allowed.htm.

Q5.13:
What are the current limits on the number of filesets in an IBM Storage Scale file system?

- A5.13:**
- Maximum number of independent filesets per file system: 3,000
 - Maximum number of AFM filesets per file system: 1,500

- Maximum number of filesets (dependent + independent) per file system: 10,000

Note: The listed numbers include the root fileset. The root fileset is an independent fileset.

Q5.14:
What is the maximum number of remote clusters for which an administrator can enable fileset access control?

A5.14:
 Per file system, fileset access control can be enabled as follows:

- For file systems which format level is 35.00 or earlier (IBM Storage Scale 5.2.1 or earlier), a maximum of 15 remote clusters.
- For file systems which format level is 36.00 or later (IBM Storage Scale 5.2.2 or later), a maximum of 31 remote clusters.

For more information, see [Fileset access control for remote clusters](#).

Q5.15:
What are the current limits on the number of snapshots in an IBM Storage Scale file system?

A5.15:

<i>Table 41: Maximum number of snapshots</i>	
Global snapshots	Maximum number of snapshots of each independent fileset
256	256

Q5.16:
What is the current limit for the number of data and metadata replicas?

A5.16:
 The maximum supported number of data and metadata replicas is 3 for GPFS V3.5.0.7 and later, and 2 for older versions.

Q5.17:
What is the limitation on pathname length?

A5.17:
 The maximum supported pathname length (directory tree path) must not exceed 4096 bytes.

Q5.18:
What is the current maximum IBM Storage Scale file size limit?

A5.18:
 The maximum single IBM Storage Scale file size architectural limit is ~ 9 EB (9,223,372,036,854,775,807).

Q5.19:
What is the limitation on symbolic links?

A5.19:
 The maximum allowed length of symbolic link targets is 1023 bytes.

Configuration and tuning questions

Please also see the questions:

- [What are the current advisories for all platforms supported by IBM Storage Scale?](#)
- [What are the current advisories for IBM Storage Scale on AIX?](#)
- [What are the current advisories for IBM Storage Scale on Linux?](#)
- [What are the current advisories for IBM Storage Scale on Windows?](#)

Q6.1:
What specific configuration and performance tuning suggestions are there?

A6.1:
 In addition to the configuration and performance tuning suggestions in the [IBM Storage Scale: Concepts, Planning, and Installation Guide](#) for your version of IBM Storage Scale:

- With IBM Storage Scale V4.2.0.3 and later, the `mmchconfig` command supports the `workerThreads` attribute to control the maximum number of concurrent file operations at any one instant, as well as the degree of concurrency for flushing dirty data and metadata in the background and for prefetching data and metadata. This attribute may be used instead of `worker1Threads` and `prefetchThreads` as a simpler and more comprehensive way to tune the file system for systems capable of handling higher sequential as well as random read/write workloads and small file activity.

See the `mmchconfig` command as documented in the [IBM Storage Scale: Administration and Programming Reference](#)

- On systems running with the Linux kernel 3.0, both `processor.max_cstate` and `intel_idle.max_cstate` should be set to zero.
- IBM Storage Scale supports 16MB block size.

Note:

1. For support of 8MB block size with SLES, the minimum level of support is SLES 10 SP4 with this patch <http://download.novell.com/Download?buildid=IOqqokqjuQQ>
2. For support of 8MB block size with RHEL, the minimum level of RHEL 6.0 shipped with `coreutils-8.4` is needed to take full advantage of block sizes larger than 4MB. Using RHEL5 with `coreutils-5.97` is supported, but can result in degraded performance from basic operations including but not limited to the `cp` command.
 - If your IBM Storage Scale cluster is configured to use SSH/SCP, it is suggested that you increase the value of `MaxStartups` in `sshd_config` to at least 1024.
 - You must ensure that when you are designating nodes for use by IBM Storage Scale you specify a non-aliased interface. Utilization of aliased interfaces may produce undesired results. When creating or adding nodes to your cluster, the specified hostname or IP address must refer to the communications adapter over which the GPFS daemons communicate. When specifying servers for your NSDs, the output of the `mm1sc1uster` command lists the hostname and IP address combinations recognized by IBM Storage Scale. Utilizing an aliased hostname not listed in the `mm1sc1uster` command output may produce undesired results.
 - On Linux systems it is recommended you adjust the `vm.min_free_kbytes` kernel tunable. This tunable controls the amount of free memory that Linux kernel keeps available (i.e. not used in any kernel caches). When `vm.min_free_kbytes` is set to its default value, on some configurations it is possible to encounter memory exhaustion symptoms when free memory should in fact be available. Setting `vm.min_free_kbytes` to a higher value (Linux `sysctl` utility could be used for this purpose), on the order of magnitude of 5-6% of the total amount of physical memory, but no more than 2GB, should help to avoid such a situation.

Also, see the following GPFS Redpapers:

- *GPFS Sequential Input/Output Performance on IBM pSeries 690* at www.redbooks.ibm.com/redpapers/pdfs/redp3945.pdf

Q6.2: What configuration and performance tuning suggestions are there for IBM Storage Scale when used primarily for Oracle databases?

A6.2: In addition to the performance tuning suggestions within the *Configuring and Tuning your system for GPFS* section of the [IBM Storage Scale: Administration Guide](#), the following recommendations are provided.

IBM Storage Scale, previously known as General Parallel File System or GPFS, is a high-performance clustered filesystem that is a complimentary solution when deploying the Oracle Database Real Application Cluster (RAC) configurations. IBM Storage Scale has the following certified uses:

- `ORACLE_HOME` directory for shared Oracle RAC database installation
- Database files for tablespaces and other general database object containers
- Oracle Clusterware registry and membership files including Oracle Cluster Registry (OCR) and Vote Disks, as well as the Grid Infrastructure Management Repository (GIMR)
- `ORACLE_BASE` for a common repository of alert logs and diag traces for the RAC cluster

After the initial IBM Storage Scale installation, the following configuration considerations and tuning parameters are suggested:

IBM Storage Scale Performance Tuning

By default, the Oracle databases open and access the IBM Storage Scale files in the correct manner. Do not use any special mount options (for example, `DIO`) for IBM Storage Scale. The Oracle database instance parameter `filesystemio_options` should remain at the default value of `SETALL`.

When configuring the Network Shared Disk (NSD) devices, there will be a one-to-one relation of a storage LUN for each IBM Storage Scale NSD. One or more LUNs/NSDs can be used for a single IBM Storage Scale filesystem. Storage LUNs for a single filesystem should use the same RAID type (for example, RAID-5 or RAID-10). It is not recommended to mix RAID types within the same IBM Storage Scale filesystem. However, different filesystems can use different RAID types. As an example, one might use RAID-10 arrays for Oracle REDO logs which need good sequential write performance and RAID-5 for data and index table spaces which may be accessed in a random manner.

Storage LUNS / GPFS NSDs can be created from different arrays (different HDDs, controllers etc) within the storage subsystem. In this manner, when the IBM Storage Scale filesystems are created, multiple NSDs would be used to produce the desirable effect

of spreading I/O across the various controllers and cache regions in the storage subsystem. This method achieves the general objective of the commonly-used Stripe And Mirror Everything (SAME) strategy.

IBM Storage Scale provides the option to set the blocksize for each filesystem individually. The Oracle-specific recommendations are as follows:

- 512KB is generally suggested
- 1MB is suggested for filesystems that are 100TB or larger

The mount options to suppress atime (-S) and mtime (-E) on the data filesystems may be helpful in reducing the overhead for the filesystem management and increasing the performance. If any operating system utility, like backup software, is using file modification time ensure that it is not suppressed.

To suppress atime and mtime (either or both), set the parameters as follows:

- To disable exact mtime tracking - `mmchfs <device> -E no`
- To suppress atime tracking - `mmchfs <device> -S yes`

For IBM Storage Scale versions earlier than 4.2.3, I/O thread tuning parameters are recommended to be initially set as follows:

- `prefetchThreads = 150`
- `worker1Threads = 450`

For IBM Storage Scale versions 4.2.3, 5.0 and later, the I/O thread tuning is controlled by a single parameter (`workerThreads`). The recommended initial value should be set as follows:

- `workerThreads=512` (or `1024`)

• **IBM Storage Scale Resilience and Availability**

- **Quorum:**

Availability of the IBM Storage Scale cluster is paramount for production or mission-critical databases. As such, the parameter `minQuorumNodes` may be set to decrease the possibility of losing cluster quorum and incurring unplanned downtime. Quorum loss or loss of connectivity occurs if a node goes down or becomes isolated from its peers by a network failure. Quorum is typically defined as one + half of the explicitly defined quorum nodes in the IBM Storage Scale cluster.

In small clusters it may be desirable to have the IBM Storage Scale cluster remain online with only one surviving node. In that case, tiebreaker disks must be used. The following parameter values are an example of this configuration option (the names of the tiebreaker disks will be different):

- ◆ `minQuorumNodes=1`
- ◆ `tiebreakerDisks = tiebreakerdisk1;tiebreakerdisk2;tiebreakerdisk3`

- **IBM Storage Scale administration and file manager network:**

As stated previously, availability of the IBM Storage Scale cluster network is an important consideration for production or mission-critical environments. As such, the IBM Storage Scale cluster network may be protected using link aggregation methods such as IEEE 802.3ad or Etherchannel.

The IBM Storage Scale network has modest bandwidth requirements as it does not transfer large sets of data from node to node. Although not a hard requirement, the administration and file manager network may be dedicated as it is in the certification tests.

- **Storage failure detection:**

IBM Storage Scale makes use of storage subsystems that employ SCSI-3 persistent reservations to control multi-node access to the shared storage. Failover times can be significantly reduced when this parameter is enabled in the filesystem cluster. IBM tests and certifies storage subsystems for use of this feature. To confirm the currently supported storage subsystems and for further considerations for implementation, see [4.1 What disk hardware has IBM Storage Scale been tested with?](#) To enable this feature, the following parameters should be set:

- ◆ `usePersistentReserve = yes`
- ◆ `failureDetectionTime = 10`

For a device to properly offer SCSI-3 Persistent Reservation support for IBM Storage Scale, it must support SCSI-3 PERSISTENT RESERVE IN with a service action of REPORT CAPABILITIES. The REPORT CAPABILITIES must indicate support for a reservation type of Write Exclusive All Registrants. Contact the disk system vendor to verify if these capabilities are provided.

Note:

- Only a subset of releases are certified for use in the Oracle environments. To confirm the certified versions log into the Oracle support (<https://support.oracle.com/>) and search on the certify tab for IBM Storage Scale product and note the target version to be used.
- For AIX, see [IBM Storage Scale and Oracle RDBMS RAC \(Doc ID 2587696.1\)](#).
- For Linux, see [RAC Technologies Matrix for Linux Platforms](#).
- Oracle certification is for storing RDBMS files in the IBM Storage Scale direct access model. Configuring an Oracle database to access through Protocol Nodes (NFS, SMB) is not certified.

- There is no plan to certify Oracle DB versions prior to 19c on IBM Storage Scale 5.1 as those versions are out of support.
- There are currently no supported levels of IBM Storage Scale qualified with Linux on Power.
- Oracle has not been certified with IBM Storage Scale on Linux on Intel and there are no current plans to do so.
- For the list of virtualization and partitioning technologies supported by Oracle, see [Certified Virtualization and Partitioning Technologies for Oracle Database and RAC Product Releases](#)

Q6.3:

What configuration requirements exist for utilizing Remote Direct Memory Access (RDMA)?

A6.3:

IBM Storage Scale supports RDMA on Linux only. IBM Storage Scale uses the VERBS programming interface to provide RDMA support. While IBM Storage Scale uses the VERBS programming interface for RDMA support, the underlying implementation of RDMA is vendor-specific.

IBM Storage Scale supports RDMA in the following configurations:

- RDMA over Infiniband fabrics is supported on the following Linux RDMA stacks, provided that the Distribution version and kernel are supported by IBM Storage Scale:
 - Mellanox RDMA stacks on ppc64le, x86_64, and arm64 or aarch64 provided that the Mellanox HCA, Distribution version, and kernel are supported by the Mellanox RDMA stack.
 - Linux Distro RDMA stacks on ppc64le, x86_64, and arm64 or aarch64 provided that the Mellanox HCA, Distribution version, and kernel are supported by Mellanox.
- RDMA over Omni-Path fabrics is supported on the following Linux RDMA stacks, provided that the Distribution version and kernel are supported by IBM Storage Scale:
 - Cornelis Networks RDMA stacks on x86_64, provided that:
 - ◆ The Cornelis Networks HFI, Distribution version, and kernel are supported by the Cornelis Networks RDMA stack.
 - ◆ IBM Storage Scale V5.1.0, or later, is required to enable Omni-Path 8K path MTU support.
 - Omni-Path 8K MTU support is enabled with the `mmchconfig verbsRdmaQpRtrPathMtu=8192` command
 - Linux Distro RDMA stacks on x86_64, provided that the Cornelis Networks HFI, Distribution version, and kernel are supported by Cornelis Networks.
 - RDMA over Omni-Path fabrics is not supported when Dynamic Page Pool is enabled.
- RDMA over Converged Ethernet (RoCE) is supported on the following Linux RDMA stacks provided that the Distribution version and kernel are supported by IBM Storage Scale:
 - Mellanox RDMA stacks on ppc64le and x86_64, provided that:
 - ◆ The Mellanox HCA, Distribution version, and kernel are supported by the Mellanox RDMA stack.
 - ◆ RDMA Connection Manager (RDMA-CM) must be enabled with the `mmchconfig verbsRdmaCm=enable` command.

The following restrictions apply for IBM Storage Scale RDMA support:

- The protocols export over CES does not utilize RDMA.
- A single Infiniband subnet is supported.
 - Clusters that use multiple fabrics that are not connected should use the `mmchconfig verbsPorts=Device/Port/Fabric` option to ensure that proper RDMA connections are created.
- Mixing of Infiniband and RoCE in a single cluster.
 - Clusters that use Infiniband and RoCE must set different fabric numbers to Infiniband and RoCE ports. To make this differentiation, use the `mmchconfig verbsPorts=Device/Port/Fabric` command. For example, set fabric to 1 for all Infiniband ports and 2 for all RoCE ports.
- The Mellanox MOFED levels MOFED 5.4.2.x and MOFED 5.5.x cannot be used with IBM Storage Scale and ESS. For more information see the following:
 - IBM Storage Scale: <https://www.ibm.com/support/pages/node/6552842>
 - IBM ESS: <https://www.ibm.com/support/pages/node/6554496>
- RDMA over Converged Ethernet (RoCE) restrictions:
 - All nodes must use IBM Storage Scale V5.1.0 or later.
 - If a node is using multiple ports for RoCE, in a single subnet IP subnet, special routes need to be created because the Linux kernel does not specifically tie IP addresses to MAC addresses by default. For more information, see [Highly Efficient Data Access with RoCE on IBM Elastic Storage Systems and IBM Spectrum Scale](#), section 3.4.7.
 - IPv6 must be enabled to use RoCE, if interfaces are selected using the port name.
- RDMA stacks based on OFED 1.1 and OFED 1.2 are not supported.
- RDMA is not supported on a node when both Mellanox HCAs and Cornelis Networks Omni-Path HFIs are enabled for RDMA.

In IBM Storage Scale 5.0.4 and later, the GPFS daemon startup service waits for a specified time period for the RDMA ports on a node to become active. You can adjust the length of the timeout period and choose the action that the startup service takes if the timeout expires. For more information, see the descriptions of the `verbsPortsWaitTimeout` attribute and the `verbsRdmaFailBackTCPIfNotAvailable` attribute in the topic *mmchconfig command*.

Note:

1. Ensure you are at the latest firmware level for both your switch and adapter.
2. When enabling Infiniband on AMD64 hardware, `iommu=soft` may be required in grub boot options to permit allocations greater than 1GB to the VERBS RDMA device. This may impact performance and CPU utilization.
3. See the question [What are the current advisories for IBM Storage Scale on Linux?](#)

Q6.4:
What configuration and performance tuning suggestions are there for the Active File Management function of GPFS?

A6.4:
In addition to the performance tuning suggestions in the *IBM Storage Scale: Advanced Administration Guide*:

- There is a known TCP performance issue with the NFS server in certain kernel releases. It is suggested for best performance to use RHEL 6.1 (or later) or SLES 11 SP2 (or later) for the NFS server in a cache relationship.

Q6.5:
Sometimes GPFS appears to be handling a heavy I/O load, for no apparent reason. What could be causing this?

A6.5:
On some Linux distributions the system is configured by default to run the file system indexing utility `updatedb` through the `cron` daemon on a periodic basis (usually daily). This utility traverses the file hierarchy and generates a rather extensive amount of I/O load. For this reason, it is configured by default to skip certain file system types and nonessential file systems. However, the default configuration does not prevent `updatedb` from traversing GPFS file systems.

In a cluster this results in multiple instances of `updatedb` traversing the same GPFS file system simultaneously. This causes general file system activity and lock contention in proportion to the number of nodes in the cluster. On smaller clusters, this may result in a relatively short-lived spike of activity, while on larger clusters, depending on the overall system throughput capability, the period of heavy load may last longer. Usually the file system manager node will be the busiest, and GPFS would appear sluggish on all nodes. Re-configuring the system to either make `updatedb` skip all GPFS file systems or only index GPFS files on one node in the cluster is necessary to avoid this problem.

Q6.6:
What considerations are there when using IBM Storage Protect with IBM Storage Scale?

A6.6:
Considerations when using IBM Storage Protect with IBM Storage Scale include:

- When using [IBM Storage Protect](#) with IBM Storage Scale, verify the supported environments:
 - IBM Storage Protect for Space Management technotes:
 - ◆ For Linux x86 at <http://www.ibm.com/support/docview.wss?uid=swg21248771>
 - ◆ For AIX at <http://www.ibm.com/support/docview.wss?uid=swg21248419>
 - ◆ For Linux on Z at <http://www.ibm.com/support/docview.wss?uid=swg21966164>
 - General overview on the integration between IBM Storage Scale and IBM Storage Protect: <https://www.ibm.com/support/pages/ibm-spectrum-protect%E2%84%A2-ibm-spectrum-scale%E2%84%A2-introduction>
 - Tivoli Field Guide for TSM for Space Management for UNIX-GPFS Integration at <http://www-01.ibm.com/support/docview.wss?uid=swg27018848>
 - IBM Storage Protect Requirements for IBM AIX Client at <http://www.ibm.com/support/docview.wss?uid=swg21052226>
 - IBM Storage Protect Linux x86 Client Requirements at <http://www.ibm.com/support/docview.wss?uid=swg21052223>
 - IBM Storage Protect Linux on Z at <http://www-01.ibm.com/support/docview.wss?rs=663&context=SSGSG7&q1=clientrequirements&uid=swg21066436>
 - To search IBM Storage Protect support information go to www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html and enter **GPFS** as the search term
- When configuring IBM Storage Scale Active File Management, see <https://www.ibm.com/support/pages/configuring-ibm-spectrum-protect-ibm-spectrum-scale-active-file-management>
- Quota limits are not enforced when files are recalled from the backup using IBM Storage Protect. This is because `dsmrecall` is invoked by the root user who has no allocation restrictions according to the UNIX semantics.
- IBM Storage Protect Backup Archive 7.1.3 client is the only supported version to work with IBM Storage Scale 4.1.1
- IBM Storage Protect Backup Archive 7.1.1 client is only verified to work with IBM Storage Scale V4.1 or later.

- IBM Storage Protect Backup Archive 7.1.0 client is only verified to work with GPFS V3.5 or later. IBM Storage Scale V4.1 is not supported at this level.
- IBM Storage Protect Backup Archive 6.3 client is only verified to work with GPFS V3.4.0.4 or later and V3.5. IBM Storage Scale V4.1 is not supported at this level.
- A DMAPI-enabled file system may be mounted on a Windows node, with certain restrictions. For more information, refer to the following resources:
 - IBM Storage Protect version 6.3 Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/index.jsp>
 - IBM Storage Protect support page at <http://www-01.ibm.com/support/docview.wss?rs=663&tc=SSGSG7&uid=swg21248771>

Q6.7:

Does IBM Storage Scale use OpenSSL for RPC secure communication?

A6.7:

IBM Storage Scale no longer uses OpenSSL for secure communication across nodes. Instead, it uses the GSKit toolkit, which is shipped in all Editions of IBM Storage Scale, as of V4.1 and later, as `gpfs.gskit`.

Q6.8:

What ciphers are supported for transport security by IBM Storage Scale?

A6.8:

TLS has an inherent limitation in that the protocol does not periodically refresh the key material that is used to protect the data that is exchanged. Depending on the cipher suite that is used, and the amount of data that is transmitted between two nodes, if the key is not updated after the threshold for that cipher is reached, the data might be at risk to an attack that would compromise the confidentiality or integrity of the transmitted data. The AES-GCM cipher suite (only available on TLS 1.2) is affected by this issue because the number of bytes that can be safely exchanged on a single TLS session is the lowest among cipher suites. This limit is on the order of hundreds of GiB. For more information, see the following links:

- <http://dx.doi.org/10.6028/NIST.SP.800-38D>
- <http://www.isg.rhul.ac.uk/~kp/TLS-AEbounds.pdf>

This issue could be exploited by an attacker with the capability to collect large quantities of network traffic exchanged between two nodes and then perform sophisticated cryptanalysis to decrypt part of the traffic exchanged, or be able to inject messages in the encrypted communications between two nodes (with partial control over their content). IBM Storage Scale uses long-lived TLS connections and when using the AES-GCM cipher suite it might exchange enough data to increase the risk of this type of weakness being exploited.

- For environments where `nistCompliance=off`
 - AES128-SHA
 - AES256-SHA
- For environments where `nistCompliance=SP800-131A`
 - AES128-SHA
 - AES128-SHA256
 - AES256-SHA
 - AES256-SHA256

Note:

- When a cluster contains both GPFS V3 (use of OpenSSL) and IBM Storage Scale V4 (use of GSKit) nodes, ensure the use of a cipher that is supported by all nodes:
 - AES128-SHA
 - AES256-SHA
- IBM Storage Scale also supports the keywords **DEFAULT**, **EMPTY**, and **AUTHONLY** in place of a cipher list. **DEFAULT**, **EMPTY**, and **AUTHONLY** are not affected by this issue. The default security mode is **EMPTY** in IBM Storage Scale V4.1 or earlier and is **AUTHONLY** in IBM Storage Scale V4.2 or later. When **EMPTY** is specified, IBM Storage Scale does not authenticate or check authorization for network connections, or encrypt transmitted data. When **AUTHONLY** is specified, IBM Storage Scale checks network connection authorization, but data that is sent over the connection is not encrypted, therefore not protected.

Q6.9:

When I allow other clusters to mount my file systems, is there a way to restrict access permissions for the root user?

A6.9:

Yes. A *root squash* option is available when making a file system available for mounting by other clusters using the `mmauth` command. This option is similar to the NFS root squash option. When enabled, it causes GPFS to *squash* superuser authority on accesses to the affected file system on nodes in remote clusters.

This is accomplished by remapping the credentials: user id (UID) and group id (GID) of the root user, to a UID and GID specified by the system administrator on the home cluster, for example, the UID and GID of the user `nobody`. In effect, root squashing makes the root user on remote nodes access the file system as a non-privileged user.

Although enabling root squash is similar in spirit to setting up UID remapping, there are two important differences:

1. While enabling UID remapping on remote nodes is an option available to the remote system administrator, root squashing need only be enabled on the local cluster, and it will be enforced on remote nodes.
2. While UID remapping requires having an external infrastructure for mapping between local names and globally unique names, no such infrastructure is necessary for enabling root squashing.

When both UID remapping and root squashing are enabled, root squashing overrides the normal UID remapping mechanism for the root user. See the `mmauth` command man page for further details.

Note: Administrators who use UID remapping to configure users with many group memberships are advised to ensure that ID remapping helper functions (IRHF) scale appropriately. For more information about UID remapping, see https://www.ibm.com/support/knowledgecenter/SSFKCN/com.ibm.cluster.gpfs.doc/gpfs_uid/uid_gpfs.html.

Q6.10:

How do I determine the maximum size of the extended attributes allowed in my file system?

A6.10:

As of GPFS 3.4, the space allowed for extended attributes for each file was increased and the performance to get and set the extended attributes was improved. To determine which version of extended attribute your file system uses, issue the `mmfsfs - -fastea` command. If the new *fast* extended attributes are enabled, **yes** will be displayed on the command output. In this case, the total space for user-specified extended attributes has a limit of 50K out of 64K and the size of each extended attribute has a limit of 16K, otherwise the total space limit is 8K out of 16K and the size of each extended attribute has a limit of 1022 bytes.

Q6.11:

What are configuration considerations when using IPv6?

A6.11:

These considerations apply when you use IPv6:

- IPv4 subnets are not supported on clusters that use IPv6 primary addresses (hostnames) and include Windows nodes.
- IBM Storage Scale does not support IPv6 for these components:
 - Clustered NFS
 - Swift Object protocols
 - HDFS protocol
 - Transparent cloud tiering
- IPv6 support for the Cluster Export Services (CES) S3 protocol is available starting with IBM Storage Scale 6.0.0.2.

For more information, see the IBM Storage Scale documentation.

Q6.12:

How should IBM Storage Scale Advanced Edition or Data Management Edition be configured to only use FIPS 140-2-certified cryptographic engines?

A6.12:

Important: IBM Storage Scale uses IBM Global Security Kit (GSKit) as the underlying cryptographic engine. In July 2022, the GSKit FIPS 140-2 certificate status was changed to *historical*.

To only use FIPS 140-2-certified cryptographic engines, you need to perform the following two steps *before* performing any configuration steps to enable encryption:

1. On IBM Security Key Lifecycle Manager (ISKLM), turn the FIPS configuration parameter **on**. See the ISKLM installation guide for more information at [Configuring compliance for FIPS in IBM Security Guardium Key Lifecycle Manager](#).
2. Issue the command `mmchconfig FIPS1402Mode=yes`.

Restrictions:

- File system operations can work with FIPS-certified cryptographic engines.

- Integrated protocol components use other cryptographic libraries and are currently not ensured to utilize FIPS- certified cryptographic engines.
- You are strongly advised to contact IBM before enabling FIPS mode in your IBM Storage Scale cluster.

Q6.13:
What are the configuration and tuning considerations for using the IBM Storage Scale integrated protocols access methods?

A6.13:
 Configuration considerations for using the integrated protocols access methods include:

- Client nodes can access integrated protocols with integrated NFS, SMB, and Swift Object services with the CES infrastructure by using IPV4.

Q6.14:
What considerations are there when using IBM Spectrum Archive with IBM Storage Scale?

A6.14:
 For the latest support information about using IBM Spectrum Archive with IBM Storage Scale, see the [Required software for Linux systems](#) topic in the IBM Spectrum Archive Enterprise Edition (EE) Knowledge Center.

Note: To see support information about a specific release of IBM Spectrum Archive Enterprise Edition, select from the options in the drop-down menu on the upper left corner of the page.

Virtualization questions

Q7.1:
How do I determine whether a server license or a client license is required when running IBM Storage Scale in VMs in a virtualized environment?

A7.1:
 Whether you need a server license or a client license is determined by the function of the virtual node. Virtual hosts that perform management functions such as cluster configuration manager, quorum node, manager node, Network Shared Disk (NSD) server, and protocol node require a server license. Virtual hosts/servers only providing a disk image to a local virtual machine (a guest) may be licensed via a client license as no management functions are performed.

With a client license, IBM Storage Scale can also execute in the hypervisor or in a VM and then export data to VMs or daemons executing on the same physical server via a protocol such as NFS as long as the client license covers all sockets available to all the VMs on that physical server.

Q7.2:
Is IBM Storage Scale on Linux (x86 and Power) supported in a virtualization environment?

A7.2:
 In a virtualization environment, the level of support depends on whether an individual node is an NSD server (has direct-attached or SAN-attached disks) or an NSD client.

Note: IBM Storage Scale is only supported as an NSD client on a XEN guest.

The following tables contain the support information for running GPFS in a virtualization environment:

<i>Table 42: KVM support matrix on Virtual Machine (VM) Guest</i>				
Configuration	KVM Version	OS Distribution	Supported Configurations	Known Limitations
GPFS nodes with no direct disk access	RHEL 7.6 or higher	Linux distributions are supported by both KVM and GPFS	N/A	<ul style="list-style-type: none"> • Live migration is not supported • KVM high availability is not supported • Local read-only cache is not supported

Configuration	KVM Version	OS Distribution	Supported Configurations	Known Limitations
GPFS nodes with direct disk access	RHEL 7.6 or higher	Linux distributions are supported by both KVM and GPFS	<ul style="list-style-type: none"> • PCI passthrough • Devices supported by GPFS and KVM • Virtio SCSI is supported with RHEL 6.4 or higher 	<ul style="list-style-type: none"> • Live migration is not supported • KVM high availability is not supported • SCSI-3 PR on Virtio is not supported • Local read-only cache is not supported

Configuration	pKVM Version	OS Distribution	Supported Configurations	Known Limitations
GPFS nodes with no direct disk access	IBM PowerKVM releases 2.1.0, 2.1.1 3.1.x	Linux distributions are supported by both pKVM and GPFS	N/A	<ul style="list-style-type: none"> • Live migration is not supported • pKVM high availability is not supported • Local read-only cache is not supported
GPFS nodes with direct disk access	IBM PowerKVM release 2.1.1 3.1.x	Linux distributions are supported by both pKVM and GPFS	<ul style="list-style-type: none"> • PCI passthrough only • Devices supported by GPFS and pKVM 	<ul style="list-style-type: none"> • Live migration is not supported • Virtio SCSI is not supported • pKVM high availability is not supported

Configuration	IBM PowerKVM Host Version	Supported Configurations
GPFS nodes with no disk access	IBM PowerKVM Release 3.1	N/A
GPFS nodes with disk access	IBM PowerKVM Release 3.1	Disk hardware supported by both pKVM and GPFS.

Configuration	VMware Version	OS Distribution	Supported Configurations	Known Limitations
GPFS nodes with no direct disk access	VMware ESX 7.x and 8.x	Linux x86_64 distributions are supported by both VMware and IBM Storage Scale. For more information, see 2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows? Refer to VMware documentation for supported Linux Distributions.	vSphere vMotion is supported	<ul style="list-style-type: none"> • vSphere Fault Tolerance (FT) is not supported • Local read-only cache is not supported

Table 45: VMware support matrix on VM guest (continued)

Configuration	VMware Version	OS Distribution	Supported Configurations	Known Limitations
GPFS nodes with direct disk access	VMware ESX 7.x and 8.x	Linux x86_64 distributions are supported by both VMware and IBM Storage Scale. For more information, see 2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows? Refer to VMware documentation for supported Linux Distributions.	Pass-through Raw Device Mapping (RDM) with physical compatibility mode are supported. VMDK disks are supported on IBM Storage Scale 5.1.4 and later. For details on configuring devices with Virtual Machine Clusters, see VMware vSphere Documentation	<ul style="list-style-type: none"> vSphere vMotion is not supported vSphere Fault Tolerance (FT) is not supported Use of Persistent Reserve is not supported Local read-only cache is not supported

Note: For current generally supported versions, check the VMware Lifecycle Policies.

Q7.3:

Is IBM Storage Scale for Linux on Z supported in a virtualization environment?

A7.3:

For IBM Storage Scale for Linux on Z, the level of support depends on the virtualization technology in use and whether an individual node has direct disk access or has no direct disk access.

Table 46: Virtual Machine support matrix for Linux on Z

Configuration	Hypervisor	OS distribution	Supported configurations	Known limitations
GPFS nodes with no direct disk access	NONE / LPAR	All supported Linux distributions	N/A	
GPFS nodes with direct disk access	NONE / LPAR	All supported Linux distributions	N/A	
GPFS nodes with no direct disk access	z/VM	All supported Linux distributions	N/A	
GPFS nodes with direct disk access	z/VM	All supported Linux distributions	N/A	
GPFS nodes with no direct disk access	KVM qemu 2.11 or higher	Linux distributions supported by both KVM and GPFS	N/A	<ul style="list-style-type: none"> Live migration is not supported. KVM high availability is not supported. Local read-only cache is not supported.
GPFS nodes with direct disk access	KVM qemu 2.11 or higher	Linux distributions supported by both KVM and GPFS	<ul style="list-style-type: none"> Devices supported by GPFS and KVM. Virtio-blk Virtio-scsi is supported with qemu 2.12 or higher. 	<ul style="list-style-type: none"> Live migration is not supported. KVM high availability is not supported. Local read-only cache is not supported. Sharing of ECKD-type DASD between KVM and non-KVM nodes is not supported. The share-rw property on scsi-block and scsi-generic requires KVM qemu 2.12.

For more information, see the following questions:

- [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)
- [What are the minimum hardware requirements for an IBM Storage Scale cluster?](#)
- [On what servers is IBM Storage Scale supported?](#)

Q7.4:

Is IBM Storage Scale on Windows supported in a virtualization environment?

A7.4:

In a virtualization environment, the level of support depends on whether an individual GPFS node uses direct-attached disks or SAN-attached disks. IBM Storage Scale on Windows is supported in the following virtualization environment:

Table 47: Hyper-V support matrix

	Hyper-V Host OS Versions	Hyper-V Guest OS Versions	Supported Configurations	Known Limitations
GPFS nodes with no direct disk access	Windows Server 2022 Windows Server 2019	RHEL 9 and all versions of Windows that IBM Storage Scale supports. For more information, see 2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?		<ul style="list-style-type: none"> • Hyper-V Live migration is not supported • Local read-only cache is not supported
GPFS nodes with direct disk access	None	None	None	

Table 48: VMware support matrix

	VMware Versions	VMware Guest OS Versions	Supported Configurations	Known Limitations
GPFS nodes with no direct disk access	VMware ESX 7.x and 8.x	All versions of Windows that IBM Storage Scale supports. For more information, see 2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?		<ul style="list-style-type: none"> • vSphere vMotion is not supported • vSphere Fault Tolerance (FT) is not supported • Local read-only cache is not supported
GPFS nodes with direct disk access	None	None	None	

IBM Storage Scale for Windows does not support any raw disk I/O when it is running as a VM guest.

Note: For current supported versions, check the VMware Lifecycle Policies.

Q7.5:

Can GPFS run in a Workload Partitioning (WPAR) environment?

A7.5:

GPFS can only be run in the global environment. It is not possible to run the GPFS subsystem or mount a GPFS file system in a WPAR. A GPFS file system can be made available to a WPAR using namefs.

By definition, a global instance is each AIX operating system that is running. The instance consists of all the program and services that compose AIX. If WPARs are inside of an instance of AIX, the parent AIX is referred to as the *global instance*. The global instance can share resources with the WPARs, but WPARs cannot directly share resources with other WPARs (http://public.dhe.ibm.com/software/passportadvantage/SubCapacity/Scenarios_Power_Systems_AIX_System_WPARs.pdf).

As GPFS does not run in a WPAR, licenses are not required for the WPAR, only the global instance. The type of GPFS license required by the global instance depends on what functions the instance is performing. See the *Licensing and Pricing* section of this FAQ for more information.

Q7.6:

Does IBM Storage Scale support exploitation of the Virtual I/O Server (VIOS) features of Power processors?

A7.6:

Yes, IBM Storage Scale allows exploitation of Power VIOS configurations. N_Port ID Virtualization (NPIV), Virtual SCSI (VSCSI), LPM (Live Partition Mobility) and Shared Ethernet Adapter (SEA) are supported in single and multiple Central Electronics Complex (CEC)

configurations. This support is limited to IBM Storage Scale nodes that are using the AIX V7.1 or 7.2 operating system or a Linux distribution that is supported by both VIOS (see www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html) and IBM Storage Scale (see [2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)).

There is no IBM Storage Scale fix level requirement for this support, but it is recommended that you be at the latest IBM Storage Scale level available. For information on the latest levels, go to the IBM Storage Scale page on [Fix Central](#)

For further information on Power VIOS go to www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html

For VIOS documentation, go to www14.software.ibm.com/support/customer/sas/f/vios/home.html

Q7.7:
Does IBM Storage Scale support PowerVM Shared Storage Pool (SSP)?

A7.7:
Yes, it is supported on PowerVM 2.2.5 and later, IBM Storage Scale 4.2.2 and later, and RedHat 7.2 and later on ppc64 NSD server.

Q7.8:
Are there any virtualization considerations when using an Oracle databases?

A7.8:
Ensure the virtualization or partitioning technology you are utilizing is supported by both GPFS and Oracle. For the list of virtualization and partitioning technologies supported by Oracle, go to <http://www.oracle.com/technetwork/database/virtualizationmatrix-172995.html>

Q7.9:
Can Linux on Z with GPFS run in logical partition (LPAR) mode or on z/VM as guest operating system?

A7.9:
Yes, GPFS cluster nodes on Linux on Z can be configured either as a NSD server or as a NSD client. All of the cluster nodes can run on LPAR directly or on z/VM as a guest operating system.

Q7.10:
Are there any limitations for using IBM Storage Scale in a virtualization environment?

A7.10:
The following limitations apply to using IBM Storage Scale in a virtualization environment:

- The same file system cannot be used on the KVM host and the virtual machine.
- The same GPFS cluster cannot be used on the KVM host and the virtual machine.

Integrated Protocol Server questions

Note: Additionally, see the questions:

- [What are the scaling considerations for the integrated protocols access methods?](#)
- [What Linux distributions are supported by the integrated protocols access methods in IBM Storage Scale V4.1.1?](#)
- [What are the configuration requirements for the use of the integrated protocols access methods?](#)

Q8.1:
What are the requirements to use the protocol access methods integrated with IBM Storage Scale?

A8.1:
Requirements to use the integrated protocols access methods include the following points:

- All protocol nodes in a cluster must use the same IBM Storage Scale version. The only scenario where an exception for this requirement can apply is if the SMB service is disabled and this service is not needed for authentication or ID mapping.
- All protocol nodes in a cluster must be based on the same CPU architecture and must be running the same operating system distribution and release. Minor releases of the same version can be mixed; you cannot mix minor releases that belong to different versions.

RHEL versions are indicated by the number that appears right after "RHEL", like in RHEL **X**; the RHEL minor releases of a version are indicated by the number that appears right after the decimal point in a version, like in RHEL **X.x**.
- Other nodes in the cluster can use a different CPU architecture and operating system release.
- Verify that your edition of IBM Storage Scale includes CES support for your operating system
- Nodes configured as a protocol node must have an IBM Storage Scale server license designation.
- On Linux on System Z servers, an RPQ would be required for IBM to review any requests for Integrated Protocol Server support.

Q8.2:

What is the minimum hardware requirement for a protocol node?

A8.2:

The protocol functionality is delivered in software only, so capability and performance depend on the configuration you choose.

- If you enable only one protocol (NFS or Swift Object), use a server with a minimum of one CPU socket of Power or Intel processor and at least 64 GB of memory, or a minimum of two IBM z14 vCPUs with at least 64 GB of memory.
- If you enable multiple protocols or if you enable SMB, use a server with a minimum of two CPU sockets of Power or Intel processor and at least 128 GB of memory.
- Ensure that the operating system disk has sufficient space for logs and memory dumps. For example, for a node with 64 GB of memory, allocate at least 64 GB of free space in the `/var` file system.
- Network configuration is important. Use at least a 10 Gb Ethernet connection for protocol client access.

Q8.3:

What are some configuration considerations when deploying the protocol functionality?

A8.3:

Configuration considerations include:

- Protocol and NSD serving functions can coexist on the same systems if the hardware can handle both workloads in terms of network, CPU, and memory. In larger scale deployments, it is advised to separate the functions on separate hardware.
- If you have an FPO configuration and you want to use the integrated protocol function, the protocol nodes must be nodes that are not FPO disk servers.
- As it is the case with basic IBM Storage Scale functions, the protocol function also relies on the administrator of the cluster to appropriately setup the network. This networking setup includes helping to ensure that the following configurations are in place.
 - The appropriate firewall ports are opened.
 - The Domain name service (DNS) is configured for hostname lookups.
 - The DNS is configured for reverse hostname lookups.
- DNS caching on CES nodes is recommended to avoid latency.
- IPV6 protocol support does not extend to the HDFS protocol.
- Some GPFS configuration aspects have not been explicitly tested with the protocol function:
 - Local Read Only Cache
- The protocol software includes open source components of NFS server (Ganesha), SMB (Samba), and S3 server (Noobaa). You must use only the versions of these components that are provided by IBM for use in the integration with CES (failover/monitoring).

CNFS aspects:

- IBM Storage Scale Clustered NFS (CNFS) and integrated protocol support by using CES are not available on the same cluster.
- If you want to integrate the SMB and S3 object functions with CES, you must migrate from CNFS to CES NFS. As you plan for that migration, consider that the failover group function of CES NFS is not completely equivalent to CNFS. If CNFS serves IBM Storage Scale AFM home, any remote cluster caching data needs to clear cache and have it repopulated.

Q8.4:

What is the guidance for NFS clients to be configured with the CES NFS function?

A8.4:

It is generally recommended that the NFS client mount with the options `mount -o hard,intr`. Mounting with `-o soft` is strongly discouraged because of the risk of data loss or corruption. Hard mounts and `intr` (interruptible) enables the application to be sure of a successful write. In addition, it is advised that the GPFS filesystem used for the NFS export have the `syncfs` option. Use the `mmfsfs` command to display if the `sync` option is set:

```
mmfsfs gpfs_file_system -o
```

and the `mmchfs` command to change the `syncfs` option:

```
mmchfs gpfs_file_system -o syncfs
```

The AIX NFS client is unable to reestablish a connection with the NFS server after NFS server failover. To resolve this problem, complete the following steps:

1. Ensure that the AIX version of the NFS client is 6.1 or 7.1.
2. Install any of the following IBM APARs from the IBM AIX support site:

- AIX 6.1: AIX 6.1 TL7 SP4 or earlier versions up to AIX 6.1 TL6 SP0 with an ifix for IV07784 and IV07918
- AIX 7.1: AIX 7.1 TL0 or later versions with ifix for IV04555, IV08311, and IV08310

Use the NFS mount options `hard,intr,timeo=1000,dio,noac` on the AIX client. For example:

```
mount -o hard,intr,timeo=1000,dio,noac spectrumScaleCESIP:/path/to/exportedDirectory
/localMountPoint
```

NFSv4.0 server uses 64-bit cookies for `readdir` and the IBM AIX NFS client truncates them into 32-bits. This causes the `readdir` from the AIX NFS clients to fail. To resolve this problem, install the following APARs from the IBM AIX support site:

- AIX 6.1 TL6 SP10 : IV28464
- AIX 6.1 TL7 SP6 : IV28372
- AIX 6.1 TL8 : IV25166
- AIX 7.1 TL0 SP8 : IV26554
- AIX 7.1 TL1 SP6 : IV28894
- AIX 7.1 TL2 : IV24863

Q8.5:

Can we configure protocol nodes with an ESS/GSS?

A8.5:

Configuration requirements to run with an ESS/GSS with V4.2 and later:

- Prior to running the 4.2.1.x Installation Toolkit for protocol deployment on a cluster containing an ESS, all servers in a cluster need to have IBM Storage Scale at V4.2.0.0 or later, code level to use the protocol function. To allow the use of protocols in this special case of using ESS NSD servers (since the IBM Storage Scale version on ESS is not entirely under the control of the user) along with other servers running V4.2.0.0 or later, we have additional configuration requirements:
 1. The ESS/GSS has to run latest ESS/GSS level that supports V4.2.0.0 or later.
 2. Install IBM Storage Scale manually on the protocol nodes using rpms from the `/usr/lpp/mmfs/4.2.x.x/gpfs_rpms` directory?
 3. Join the protocol nodes to the existing ESS cluster using the `mmaddnode` command
 4. The cluster should have CCR enabled. Issue the `mmlscluster` command to determine if CCR is enabled. Issue the `mmchcluster -ccr-enable` command to enable CCR if needed.
 5. Any nodes designated as quorum or manager nodes must be running 4.2.0.0 or later code. Depending upon the configuration, this may mean movement of quorum and/or manager function to higher level nodes within the cluster.
 6. ESS nodes need to be in the same cluster as the protocol nodes that export ESS file systems.
 7. We expect a node class called `gss` or `gss_ppc64` (`mmlsnodeclass --all`)
 8. Input the protocol nodes into the Installation Toolkit (do not input the ESS IO nodes nor EMS node).
 9. Configure the protocols using the Installation Toolkit.
 10. Proceed with a protocol deployment using the Installation Toolkit .
 11. Run protocol CLI commands from a protocol node if other nodes in the cluster are at a lower level.

Configuration requirements to run with an ESS/GSS at pre V4.2.0.0 levels:

- Prior to running the Installation Toolkit for protocol deployment on a cluster containing an ESS, all servers in a cluster need to have IBM Storage Scale at V4.1.1 or later, code level to use the protocol function. To allow the use of protocols in this special case of using ESS NSD servers (since the IBM Storage Scale version on ESS is not entirely under the control of the user) along with other servers running V4.1.1 or later, we have additional configuration requirements:
 1. The ESS/GSS has to run latest ESS/GSS level that supports V4.1.1.
 2. Install IBM Storage Scale manually on the protocol nodes using rpms from the `/usr/lpp/mmfs/4.x.0.0/gpfs_rpms` directory
 3. Join the protocol nodes to the existing ESS cluster using the `mmaddnode` command
 4. The cluster should have CCR enabled. Issue the `mmlscluster` command to determine if CCR is enabled. Issue the `mmchcluster -ccr-enable` command to enable CCR if needed.
 5. None of the V4.1.0.8 nodes can have quorum or management – any nodes that are designated as quorum or manager nodes should be running V4.1.1 or later code (note that the ESS management function and GUI can be on a node that runs V4.1.0.8).
 6. ESS nodes need to be in the same cluster as the protocol nodes that export ESS filesystems.
 7. We expect a node class called `gss` or `gss_ppc64` (`mmlsnodeclass --all`)
 8. Input the protocol nodes into the Installation Toolkit (do not input the ESS nodes) .
 9. Configure the protocols using the Installation Toolkit.
 10. Proceed with a protocol deployment using the Installation Toolkit
 11. No protocol CLI will be run from ESS nodes. The CLI only runs on nodes that are at V4.1.1 or later.

Note:

1. Asynchronous Disaster Recovery function of V4.1.1or later, requires a V4.1.1 filesystem format and therefore cannot be used with ESS 3.0 (and by implication ESS+Protocols) ESS does not allow any function to reside on ESS nodes including protocol node functionality.
2. Expected sequence for configuring Protocols with ESS
 - Install and configure an ESS using standard procedures.
 - Perform a standard install of IBM Storage Scale on additional nodes that will be added to the cluster with ESS nodes (include any Protocol nodes).
 - Add these nodes (including protocol nodes) to the cluster created during ESS installation and configuration.
 - Change node roles to ensure none of the ESS NSD servers are designated manager or quorum.
 - Create any prerequisite file systems (including shared-root).
 - Configure protocol nodes for CES use and enable protocols.

Q8.6:
Are there any limitations that I should be aware of before using the integrated CES Protocol function?

A8.6:
 For more information, see the *SMB limitations* topic in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q8.7:
Where can I find additional information about protocols?

A8.7:
 Additional information can be found at:

- The IBM Storage Scale documentation (<https://www.ibm.com/docs/en/storage-scale>) introduces the function, provides guidance on the set-up, administration and management of protocols including information on logs.

Q8.8:
How do I reduce the logging level for performance monitoring of Swift?

A8.8:
 To reducing the logging level for performance monitoring of Swift, run the following commands on each of the Swift Object protocol node.

- To clear them immediately issue the commands:

```
perl -p -i -e "s/PMS\_LOG\_LEVELS\[\'DEBUG\'\\]/PMS\_LOG\_LEVELS\[\'ERROR\'\\]/g" /usr/local/pmswift/pmswiftparams.py
rm -f /var/log/pmswift/pmswift*
systemctl restart pmswiftd.service
```

- To automatically clear the logs after seven days issue the commands:

```
perl -p -i -e "s/PMS\_LOG\_LEVELS\[\'DEBUG\'\\]/PMS\_LOG\_LEVELS\[\'ERROR\'\\]/g" /usr/local/pmswift/pmswiftparams.py
systemctl restart pmswiftd.service
```

Q8.9:
What are the limitations when multiple file systems are exported by NFS?

A8.9:
 If a file system that was previously exported successfully by NFS on a CES node becomes unavailable, the NFS daemon exits and the CES node becomes unhealthy. If the file system becomes unavailable on all CES nodes, the whole CES cluster becomes unhealthy. In this case, the NFS daemons need to be restarted on all nodes.

Q8.10:
How to effectively use VMware with IBM Storage Scale?

A8.10:
 IBM Storage Scale 4.2.2 does not support NFS v4.1. Therefore, it is recommended that NFSv3 be used with VMWare ESX with IBM Storage Scale 4.2.2. For more information, see [VMware support matrix on VM guest](#) and [VMware support matrix](#).

Q8.11:
What are the limitations when upgrading the SMB service?

A8.11:

All protocol nodes that are running the SMB service must have the same version of `gpfs.smb` installed at any time. Upgrading the SMB service also requires an outage. For a manual upgrade, it is recommended that you upgrade all of the other parts of the system first before taking an outage to upgrade the `gpfs.smb` package on the protocol nodes. For more information, see the procedure at https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.3/com.ibm.spectrum.scale.v4r23.doc/bl1ins_updatingsmb.htm. If you use the toolkit for the upgrade, a similar process is followed to ensure a proper SMB upgrade.

Hadoop Support questions

Q9.1:

What platforms are supported for Hadoop on IBM Storage Scale?

A9.1:

Hadoop can run with IBM Storage Scale on x86_64 systems. The verified operating system for HDFS Transparency is Red Hat Enterprise Linux (RHEL).

Important:

- IBM Storage Scale does not support Hadoop on AIX, Linux on Z, or Windows.
- Hadoop support on Power LE system is deprecated for HDFS Transparency 3.1.1 and is not supported in versions 3.2.2 and 3.3.6.

Q9.2:

What versions of Hadoop are supported by IBM Storage Scale?

A9.2:

IBM Storage Scale Hadoop support is aligned with the Hadoop versions supported by Cloudera. For the particular versions that are supported, see [Hadoop distribution support](#).

Note:

- For more information, see [Limitations and differences from native HDFS](#).
- If there are multiple connector versions that support the Hadoop version that you are using, it is recommended to use the latest connector.

Q9.3:

What Hadoop distributions are supported by IBM Storage Scale?

A9.3:

Cloudera distributions and open source Apache Hadoop are supported. For more information, see [Hadoop distribution support](#).

Q9.4:

What IBM Storage Scale licenses do I need to use the IBM Storage Scale Hadoop connector?

A9.4:

The IBM Storage Scale Hadoop connector can be used with all license types and all editions. There is no additional license requirement for Hadoop access. The IBM Storage Scale Erasure Code Edition can be used for the centralized storage to connect to HDFS Transparency. If you want to use IBM Storage Scale Erasure Code Edition to run in hyper converged mode, consult with your IBM storage seller on how to do this properly.

Q9.5:

Can Hadoop connector be used for IBM Storage Scale that has shared storage (including IBM Storage Scale System)?

A9.5:

If you want to utilize the HDFS Transparency connector, you need to install the `gpfs.hdfs-protocol` image. For more information, see the [HDFS Transparency download](#) section.

Q9.6:

Can Hadoop connector be used for IBM Storage Scale that has shared storage (including IBM Storage Scale System) and internal storage (FPO pool) in the same file system?

A9.6:

Yes. The HDFS Transparency connector (`gpfs.hdfs-protocol`) works with both the types of storage and leverage locality when the files are stored on the FPO pool. You can also use the ILM feature of IBM Storage Scale to move data between the FPO and shared/IBM Storage Scale System pool.

Q9.7:**What open-source Hadoop components are certified using IBM Storage Scale connector?****A9.7:**

IBM Storage Scale HDFS Transparency is tested and certified with Cloudera HDP and Cloudera Private Cloud Base Hadoop distributions. Open Source Apache Hadoop components that are also part of Cloudera HDP and Cloudera CDP distributions and have the same major and minor release numbers are also supported. Ambari is only supported with Cloudera HDP distribution. If you have specific questions about other components, send an email to scale@us.ibm.com.

Q9.8:**Why do you need special IBM Storage Scale connector instead of using Hadoop local file system (file:///) if not using FPO/internal storage?****A9.8:**

As compared with **file:///**, the IBM Storage Scale connector does more than handle I/O from Hadoop applications. This includes reporting data block location for data awareness scheduling in scheduling engineer; large data chunk size support for scheduling; optimization for workloads, such as HBase and Hive; supporting both FPO and shared storage within the same cluster etc. Additionally, using a cluster file system as a local - per node file system will introduce challenges in Yarn scheduling as a map reduce split will map to the entire file versus creating multiple splits per file which is handled in the IBM Storage Scale connector.

Q9.9:**What are the differences between the Hadoop connector and the HDFS Transparency connector?****A9.9:**

The Hadoop connector (gpfs.hadoop-connector) implements the Hadoop File System API. It does not support Kerberos authentication or the WebHDFS REST API. The Hadoop connector is no longer supported.

The HDFS Transparency connector implements the Hadoop HDFS RPC and supports wider Hadoop workloads including full Kerberos, WebHDFS, and distcp.

Q9.10:**What are the requirements/limitations for using the IBM Storage Scale HDFS Transparency connector?****A9.10:**

Considerations for using the IBM Storage Scale HDFS Transparency connector support include:

- The Hadoop Distributed File System (HDFS) Transparency connector supports both FPO mode and shared storage including IBM Storage Scale System since the HDFS Transparency connector 2.7.0-1.
- The HDFS Transparency connector is available in the IBM Storage Scale self-extracting package. The HDFS Transparency connector 2.7 and 3.1.0 are also available from IBM on Fix Central.
- The HDFS Transparency connector 3.1.0 and earlier have no dependencies on the level of IBM Storage Scale. However, it is fully tested with IBM Storage Scale V4.1.1 or later and it is recommended at these levels.
- The CES HDFS Transparency 3.1.1 and later are depended on for the BDA integration toolkit and IBM Storage Scale versions. If using the IBM Storage Scale installation toolkit, you can only install from the packages from the self-extracting package directory.
- Linux on Z is not supported.

Features Discontinued

- From December 31, 2021, IBM will discontinue support for Hadoop Distributed File System (HDFS) Transparency Connector 3.1.0 for Hortonworks Data Platform (HDP) from the IBM Storage Scale offerings.
- From October 8, 2021, IBM will discontinue support for Hadoop Distributed File System (HDFS) Transparency Connector 2.7 for Hortonworks Data Platform (HDP) from the IBM Storage Scale offerings.

Q9.11:**What are the current advisories for the IBM Storage Scale Hadoop connector?****A9.11:**

The current advisory is:

Abstract:

IBM Storage Scale (GPFS) Hadoop connector is affected by a security vulnerability (CVE-2015-7430)

Summary:

A security vulnerability has been identified in the IBM Storage Scale (GPFS) Hadoop connector which could allow an unprivileged user the ability to read, write, modify, or delete any data in a GPFS file system (CVE-2015-7430)

See the complete bulletin at either <http://www-01.ibm.com/support/docview.wss?uid=isg3T1022979> **or** <http://www.ibm.com/support/docview.wss?uid=ssg1S1005461>

Swift Object protocol support questions

Q10.1:

What considerations are there when using OpenStack Software with IBM Storage Scale?

A10.1:

Important:

- CES Swift Object protocol feature is not supported from IBM Storage Scale 5.1.9 onwards.
- IBM Storage Scale 5.1.8 is the last release that has CES Swift Object protocol.
- IBM Storage Scale 5.1.9 will tolerate the update of a CES node from IBM Storage Scale 5.1.8.
 - *Tolerate* means:
 - ◆ The CES node will be updated to 5.1.9.
 - ◆ Swift Object support will not be updated as part of the 5.1.9 update.
 - ◆ You may continue to use the version of Swift Object protocol that was provided in IBM Storage Scale 5.1.8 on the CES 5.1.9 node.
 - ◆ IBM will provide usage and known defect support for the version of Swift Object that was provided in IBM Storage Scale 5.1.8 until you migrate to a supported object solution that IBM Storage Scale provides.
- Please contact IBM for further details and migration planning.

IBM Storage Scale includes support for Swift Object protocol:

- For IBM Storage Scale 4.2.3 and 5.0.0, this support is built with the Mitaka release of OpenStack. The OpenStack Mitaka Swift details can be found at <https://releases.openstack.org/mitaka/#mitaka-swift>.
- For IBM Storage Scale 5.0.1 to 5.0.5, this support is built with the Pike release of OpenStack. The OpenStack Pike Swift details can be found at <https://docs.openstack.org/releasenotes/swift/pike.html>.
- For IBM Storage Scale 5.1.0, this support is built with the Train release of OpenStack. The OpenStack Train Swift details can be found at <https://releases.openstack.org/train/#train-swift>.

Note: Swift Object protocol in release 5.1.0.0 requires APAR IJ26961 to be installed.

IBM Storage Scale includes the Swift and Keystone components to provide a complete object storage capability that is tightly integrated with IBM Storage Scale.

Other resources for using OpenStack software with IBM Storage Scale include:

- IBM Storage Scale in an OpenStack Environment at <http://www.redbooks.ibm.com/abstracts/redp5331.html>.
- The OpenStack Cinder documentation for configuring the IBM Storage Scale volume driver at <http://docs.openstack.org/mitaka/config-reference/block-storage/drivers/ibm-gpfs-volume-driver.html>.
- The GPFS support in OpenStack Cinder Havana release blog at <http://www.redbooks.ibm.com/redpapers/pdfs/redp5331.pdf>.
- A Deployment Guide for Elastic Storage Object at <http://www.redbooks.ibm.com/redpieces/abstracts/redp5113.html>.
- IBM Private, Public, and Hybrid Cloud Storage Solutions at <http://www.redbooks.ibm.com/abstracts/redp4873.html>.
- Customers who are interested in learning more about IBM Storage Scale and OpenStack Software should contact scale@us.ibm.com for additional guidance.

Q10.2:

What are the new features of IBM Storage Scale for Swift Object Storage?

A10.2:

The new features of Swift Object Storage in 5.1.0 include:

- OpenStack Train Release, including Swift 2.23.1 and Keystone 16.

The new features of Swift Object Storage in 5.0.1 include:

- OpenStack Pike Release, including Swift 2.15.1 and Keystone 12.0.1.
- Swift3 release 1.12, including minimum segment size for multi-part upload in S3 protocol.

The new features of Swift Object Storage in 5.0.0 include:

- Support on Ubuntu 16.04.
- Cumulative upgrades from older releases directly to 5.0.0.

The new features of Swift Object Storage in 4.2.3 include:

- OpenStack Mitaka Release, including Swift 2.7.3 and Keystone 9.3.1.

The new features of Swift Object Storage in 4.2.1 include:

- OpenStack Liberty Release, including Swift 2.5.0 and Keystone 8.0.0.
- Storage policy support for encryption. Storage policies allow encryption to be enabled on a per container basis.
- Support for issuing `mmobj` commands on GPFS client nodes.
- Improved problem determination documentation.
- Improved documentation for configurations using an external Keystone identity service.
- Simplified enablement of S3 API support.
- Simplified enablement of Unified File and Swift Object access support.
- Monitoring of AD and LDAP services used with Keystone.
- Support for object configuration using CES groups.

The new features of Swift Object Storage in 4.2 include:

- Storage policy support for compression, uniform file and object access and multi-region active object storage. Storage policies allow these features to be enabled on a per container basis.
- Compression allows object data to be compressed in the background after being committed to storage.
- Unified file and object access allows object data to be ingested from the object interface and then be accessed (read/update/delete) from the file interface, as well as data to be ingested from the file interface and then accessed from the object interface.
- Multiregion active-active object storage allows you to configure containers that have data replicated between multiple sites.
- S3 emulation support has added support for S3 ACLs on the object interface, and support for S3 multi part uploads.

Q10.3:

How should I ensure that unauthorized users cannot access my object data when using IBM Storage Scale for Swift Object storage?

A10.3:

To ensure against unauthorized access to your object data:

- It is extremely important to set up firewall rules to limit access to the ports used by Swift Object storage services.
- Shell access by non-root users must be restricted on IBM Storage Scale protocol nodes where the Swift Object services are running to prevent unauthorized access to object data.
- See the *IBM Storage Scale: Advanced Administration Guide* for your level of code at <https://www.ibm.com/docs/en/storage-scale>. Refer to the section on *Object port configuration*.

Q10.4:

What is the level of compatibility between S3 API in IBM Storage Scale and Amazon S3?

A10.4:

IBM Storage Scale uses the OpenStack Swift3 code to implement the S3 API. The compatibility level is documented at https://review.openstack.org/#/c/504281/11/doc/source/s3_compat.rst.

Q10.5:

How can I ensure secure data in flight between object client and IBM Storage Scale Swift Object storage?

A10.5:

IBM Storage Scale Swift Object storage should be configured with suitable load balancer (for example, HAProxy) enabled in SSL mode to ensure secure data in flight between object client and the object storage system. It is the customer's responsibility to provide and configure the load balancer.

Q10.6:

What authentication schemes does IBM Storage Scale Swift Object Storage support?

A10.6:

IBM Storage Scale for Swift Object storage must be configured with OpenStack Keystone (either installed on IBM Storage Scale protocol nodes or using an external Keystone instance). Keystone supports integration with Microsoft Active Directory, LDAP or can use a local Postgres repository for user data. The Unified File and Object Access feature can be configured to use either local or unified identity management mode. When using unified mode, the object identity back end must be the same as that used for file. If using local mode, these can be different. See the support Authentication Matrix for Swift Object for your version of IBM Storage Scale at http://www-01.ibm.com/support/knowledgecenter/STXKQY_4.2.0/com.ibm.spectrum.scale.v4r2.ins.doc/bl1ins_authconcept.htm

Q10.7:

When using Unified File and Object Access feature, when should I use the `local_mode` vs `unified_mode` for identity management?

A10.7:

The use of identity management mode depends upon your use case. Generally:

- `local_mode`: Suitable when authentication schemes for file and object are different and file access is required for applications and file ownership of data ingested via object interface is not of interest.
- `unified_mode`: Suitable for unified file and object access for end users. File ownership for data ingested via object interface is required and one can leverage features like common ILM policies for file and object data based on data ownership. In this mode, Object and File are expected to use a common authentication back end coming from the same directory service (AD+RFC 2307 or LDAP)

Q10.8:

Can I have SMB/NFS export over object data when using Unified File and Object Access feature?

A10.8:

Yes, but it is important to note that the file authorization is independent from object authorization, and that changing a file ACL does not impact the object access of the data and vice versa. Also, the ownership of the data seen on file interface depends upon the ID management mode (`local_mode` or `unified_mode`) being used. We recommend creating exports at the container/bucket level in the Unified File and Object directory hierarchy.

Q10.9:

Can I have read/write access from file interface as well as Swift Object interface on the same data, when using Unified File and Object Access feature?

A10.9:

Yes, but it is important to consider the use case for this data. Swift Object semantics do not support locking of objects. From the Swift Object interface, every PUT operation creates a new object atomically. From the file interface, files can be created and then updated many times. Even though simultaneous read/write access is possible between file interface and Swift Object interface, doing this will lead to unpredictable results. We recommend a serial work flow where any object or any file is only accessed from one interface at any point in time. One way to achieve this is to have either of the interfaces (file interface or Swift Object interface) to be read only and the other to be read/write at any point in time.

Q10.10:

What are the limitations for IBM Storage Scale Swift Object Storage with Unified File Access enabled?

A10.10:

Limitations for IBM Storage Scale Swift Object Storage with Unified File Access enabled include:

- We see acceptable results in tests with up to 10 containers and 400,000 files/objects per container when running the objectizer at its default interval of 30 minutes. If you require a larger number of containers, we recommend using an objectizer interval of 120 minutes or longer. The interval can be changed as shown here, for example, setting it to 2 hours (in seconds):

```
mmdsh config change --ccrfile spectrum-scale-objectizer.conf --section DEFAULT \
    --property objectization_interval --value 7200
```

- In some situations, the objectizer process can complete before all of the new files are added to the container listing but are still queued as asynchronous operations. In this case, the files are visible and can be accessed from the Swift Object interface, but they do not show up in the container listing for some time. These objects eventually show up in a container listing when this asynchronous queue is processed.
- It is possible when stopping GPFS that the **ibmobjectizer** service may not be stopped automatically. You can verify if this is the case and force it to stop using the `systemctl` command:

```
mmdsh -N cesnodes systemctl status ibmobjectizer -n 0
mmdsh -N cesnodes systemctl stop ibmobjectizer
```

- Additional limitations for Unified File and Object Access are documented in the Knowledge Center. See the *IBM Storage Scale: Administration and Programming Reference, Managing Object Storage* section for your level of IBM Storage Scale at <https://www.ibm.com/docs/en/storage-scale>.

Q10.11:

What are the considerations when using HAProxy load balancer with IBM Storage Scale Swift Object Storage?

A10.11:

When using HAProxy as a load balancer to distribute Swift Object workloads across multiple protocol nodes, users need to be aware:

- The default HAProxy timer values may interfere with communication between the Swift client and protocol node. With HAProxy default timer values that are typically lower than either the default object server or default client settings, HAProxy can timeout and terminate a transaction before either the client or the server timers expire. This may result in the server logging an "unexpected client disconnect" indicating to the object server administrator that the client disconnected when actually HAProxy terminated the connection.

- The recommended debug method for HAProxy environments is to either remove HAProxy from the configuration (preferred) or disable the timers in `/etc/haproxy.cfg` (or set the HAProxy timers to very large values, ex. 60 minutes), investigate the problem between the client and server, and then re-instate HAProxy as the load balancer.

Q10.12:

Are there known issues with OpenStack Software?

A10.12:

The following are the known issues with documented workarounds:

In Release 4.2.1 and earlier:

- Authentication fails when openstackclient prompts for a password: see <https://bugs.launchpad.net/python-openstackclient/+bug/1473862>. The workaround is to include the password in your `openrc` or environment settings, or pass on the command line.
- The IBM Storage Scale Swift Object protocols functionality on the Linux (standard and advanced) platform is affected by security vulnerabilities in the TLS and SSL protocols. For the workaround, see <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1009336>

In Release 4.2.2 and later:

- Credentials required for s3 access are created and stored locally in Keystone even when Keystone is configured to use an external identity source. Using an external identity source for credentials required for s3 access is not supported.

Q10.13:

Are there any additional limitations or restrictions when using IBM Storage Scale compression or encryption with Swift Object protocol?

A10.13:

No, there are no additional limitations or restrictions for these with Swift Object Protocol. Any limitations or restrictions that exist with IBM Storage Scale compression or encryption also apply to Swift Object protocol.

Q10.14:

What are the limitations when configuring objects to use CES groups?

A10.14:

Limitations when configuring objects to use CES groups include:

- Configuring objects to use CES groups is supported in IBM Storage Scale V4.2.1 or later.
- If the CES group feature is used as described in the *Configuration of object for isolated node and network groups* section of the [Advanced Administration Guide](#) the following limitation needs to be considered:

If a CES IP address that has the database or singleton attribute assigned is changed by either of the following, it needs to be ensured that the selected CES IP address is within the object group:

- Removed via the command `mmces address remove`
- Any of the attributes are changed to a different CES IP via the command `mmces address change`

Use the `mmces address list` command to list the current group and attribute assignment.

If the assignment is incorrect so that an attribute is assigned to a CES IP that is not part of the object group, use the `mmces address change --ces-ip IP --attribute Attribute` command to change the attribute to CES IP address assignment. For example, for `ces_ip8` within the object group

```
mmces address change --ces-ip ces_ip8 --attribute object_database_node
```

Q10.15:

Does the Swift Object protocol support immutable objects?

A10.15:

No, immutable objects are not supported. Although immutability is supported by IBM Storage Scale filesets, the eventual consistency model of OpenStack Swift means immutability of object data cannot be guaranteed.

Q10.16:

Are OpenStack package repositories required to install the Swift Object protocol?

A10.16:

Configuration of OpenStack repositories is needed in certain release streams. Release streams 5.1.0.0 through 5.1.2.0 require the configuration of OpenStack repositories. Release streams 5.1.2.1 and higher and 5.0.x and lower do not require OpenStack repositories.

If installing a release stream where an OpenStack repository is necessary, refer to the documentation associated with the specific release for relevant setup instructions.

Q10.17:

Can additional protocol nodes be added when the Swift Object protocol is in toleration mode?

A10.17:

The Swift Object protocol was discontinued starting with IBM Storage Scale 5.1.9. However, the Swift Object protocol may continue to be enabled in the CES environment in a toleration mode.

In this toleration mode, attempting to add additional protocol nodes might fail because the necessary Swift Object packages are no longer included with the installation media, which means that the necessary Swift Object packages are not automatically installed on the new node.

Before any attempt to enable a new protocol node as a CES node, it is necessary to first install the pre-5.1.9 Swift Object packages manually on the new protocol node.

To manually install the Swift Object packages, complete the next procedure.

1. To determine the version of Swift Object packages that are installed, use the following command on one of the active CES nodes:

```
rpm -qi spectrum-scale-object
```

2. From the Spectrum Scale installation media that corresponds to the active spectrum-scale-object package, copy the object_rpms directory to the new protocol node.

3. On the new protocol node, create the file /etc/yum.repos.d/ssobjectprotocol.repo with the following contents:

```
[objectrpms]
name=Spectrum Scale Object Protocol
baseurl=file:///NEWPATH/object_rpms/rhel8/
enabled=1
```

Where NEWPATH is the path to the object_rpms directory on the new protocol node.

4. Install the spectrum-scale-object package on the new protocol node by using the next command:

```
dnf install spectrum-scale-object
```

After the spectrum-scale-object package and all its dependencies are installed, the new node can be enabled as a CES protocol node with the mmchnode or spectrumscale commands.

For more information about toleration mode for the Swift Object protocol, see [Stabilized, deprecated, and discontinued features in IBM Spectrum Scale](#).

CES S3 protocol support

Q11.1:

What are the considerations for using the CES S3 protocol with IBM Storage Scale?

A11.1:

From IBM Storage Scale 5.2.1.0 onward, Cluster Export Services (CES) S3 protocol is available.

Q11.2:

What are the currently available features of CES S3 for IBM Storage Scale?

A11.2:

The CES S3 provides the following features:

- Optimized for multiprotocol data access to enable workflows that access the same instance of data by using S3 object protocol and other file protocols like POSIX, NFS, SMB and CSI.
- High-performing, highly available, and scalable S3 object access to data that is stored in IBM Storage Scale file systems.
- Support for S3 API calls that are required to process data that is stored in IBM Storage Scale file systems.
- Files and directories in IBM Storage Scale file systems are represented 1:1 as S3 objects and S3 buckets.
- Installation and deployment support with IBM Storage Scale installation toolkit.
- Support for S3 service management with the existing mmces commands.
- For administrators, management of S3 configuration, S3 accounts, and S3 buckets by using the new mms3 command.

- Backup and restore of S3 configuration data.
- Health monitoring of S3 protocol stack.
- Secure storage of S3 keys with encrypted S3 secret keys.
- Support for ILM including tiering to tape (RPQ).
- Support for S3 access to AFM managed data and AFM S3 to CES S3.
- Stretch cluster support.

Q11.3:
How can I ensure secure data in flight between a CES S3 client and IBM Storage Scale CES S3 protocol stack?

A11.3:
 CES S3 must be configured with self-signed SSL/TLS certificates to ensure secure data in flight between S3 client and the CES S3 protocol stack. It is the customer's responsibility to provide and configure the self-signed SSL/TLS certificates.

Q11.4:
What are the additional limitations or restrictions when using IBM Storage Scale compression or encryption with CES S3 protocol?

A11.4:
 No, no additional limitations or restrictions should be considered for S3 Protocol. Any limitations or restrictions that exist with IBM Storage Scale compression or encryption also apply to S3 protocol.

Q11.5:
Does the CES S3 protocol support immutable objects?

A11.5:
 No, immutable objects are not supported.

Q11.6:
What are the configuration considerations when deploying the CES S3 protocol functionality?

A11.6:
 Configuration considerations include:

- Operating system: RHEL8 or RHEL9.
- Architectures supported: x86_64, Power (ppc64le), Linux on Z (s390x for RHEL9 only).
- IPv6 is not supported.
- While the IBM Storage Scale cluster uses RDMA, the CES S3 protocol does not utilize RDMA.
- When using the CES S3 service, depending on the workload requirements, the number of S3 protocol stack endpoint processes can be increased or decreased manually. For handling heavy workloads, consider assigning a higher number of endpoint processes. For information about configuring endpoints for the S3 protocol stack, read the [S3 protocol quick reference](#) in the IBM Storage Scale documentation.

Q11.7:
Can the CES S3 (technology preview) that was introduced with 5.2.0.x be upgraded to the CES S3 that is included with 5.2.1.x onward?

No. Upgrading the 5.2.0.x S3 protocol (technology preview) to the S3 protocol that is included with 5.2.1.x and higher versions is not supported.

Installation Toolkit questions

Q12.1:
Can I upgrade an IBM Storage Scale cluster with protocols directly from 4.1.1.x to 4.2.0.x or 4.2.1.x?

A12.1:
 IBM Storage Scale V4.1.1.x clusters with protocols must first upgrade to V4.2.0.0 and then to V4.2.0.x or 4.2.1.x. See question [What are the limitations when I use the Installation Toolkit to upgrade IBM Storage Scale from 4.1.1 to 4.2 with NFS or SMB Protocol?](#) for help with the first step of this dual upgrade.

Q12.2:
What are the limitations when I use the Installation Toolkit to install IBM Storage Scale 4.2.0.0 or upgrade from IBM Storage Scale 4.1.1.x to 4.2.0.0 or 4.2.1.0 with Swift Object protocol?

A12.2:
 Current limitations for the Installation Toolkit include:

- The Installation Toolkit does not support installing the IBM Storage Scale GUI during an upgrade of a cluster which has the Swift Object protocol enabled. If you happen to do this, the GUI will not start properly after the upgrade completes. Before installing the IBM Storage Scale GUI, (and after completing the cluster upgrade) you will need to first remove or rename the postgresql.service file from /etc/systemd/system/postgresql.service. After that, install the IBM Storage Scale GUI by running the **spectrumscale install** command. It is generally advisable to check for this file if the GUI reports database issues after install or upgrade. For example:

```
mv /etc/systemd/system/postgresql.service /etc/systemd/system/postgresql.service.sav.4.1.1
```

- The Installation Toolkit does not support installing or upgrading a configuration that uses an external keystone. This limitation will be corrected in an upcoming refresh of Installation Toolkit. The install or upgrade can be accomplished by using the IBM Storage Scale CLI. See the *IBM Storage Scale: Administration and Programming Reference* section on *Configuring object authentication with external Keystone server* at http://www-01.ibm.com/support/knowledgecenter/STXKQY_4.2.0/ibmspectrumscale42_welcome.html
- In some cases, the Installation Toolkit may fail during upgrade of the performance monitoring component. The reason for this is that the Installation Toolkit fails to stop **pmswiftd.service** before upgrade, leaving the **pmswiftserver** running and the corresponding ports open. This causes a failure when starting **pmswiftd.service** after upgrade. A workaround for this problem is to stop all running pmswiftserver processes on ALL protocol nodes and then manually start **pmswiftd.service**. To do this, use the following commands on each protocol node:
 - Check **pmswiftserver** process, this step is optional

```
$ ps aux|grep pmswiftserver
```

- Stop **pmswiftserver** process

```
$ kill -9 $(pgrep pmswiftserver)
```

- Start **pmswiftd.service**

```
$ systemctl start pmswiftd.service
```

- Check the status of **pmswiftd.service**

```
$ systemctl status pmswiftd.service
```

- You may be required to restart **pmsensors.service** if **systemctl status** shows FAILED or corresponding **pmswiftproxy** is not active (this step may or may not be required). You should look for something like **/usr/bin/python2.7 /usr/local/pmswift/pmswiftproxy** in the output.

```
$ systemctl restart pmsensors.service
```

- Check the status of **pmsensors.service** to make sure that corresponding **pmswiftproxy** is active. You should look for something like **/usr/bin/python2.7 /usr/local/pmswift/pmswiftproxy** in the output.

```
$ systemctl status pmsensors.service
```

Q12.3:

Can GUI nodes be added during an upgrade to IBM Storage Scale V4.2.0.1?

A12.3:

When planning to add GUI nodes with the Installation Toolkit, add them via **spectrumscale install** or **spectrumscale deploy**, either before performing an upgrade to 4.2.0.1 or afterwards. Attempting to add GUI nodes during the upgrade itself may result in a failure during the *Upgrading Performance Monitoring* step.

Q12.4:

What are the limitations when I use the Installation Toolkit to upgrade IBM Storage Scale from 4.1.1 to 4.2 with NFS or SMB Protocol?

A12.4:

Current limitations for using the Installation Toolkit to upgrade include:

- An intermittent failure may occur during Deploy while 'enabling gpfs fileset quota sensors'. This failure may result in SMB being down and CTDB being unhealthy. The following solution will recover the cluster if this issue is hit:

1. This problem will create core files on the CES nodes. Typically the core files are in and may have even filled up the entire root file system. The files will be named similar to /core.1734. Remove all core files:

```
rm -rf /core.*
```

2. Stop SMB services on all CES nodes:

```
/usr/lpp/mmfs/bin/mmces service stop SMB -a
```

3. Stop pmsensors on all CES nodes. ssh to each CES node and issue the following command:

```
systemctl stop pmsensors
```

4. Locate **SMBStats.cfg** and **SMBGlobalStats.cfg** on the installer node in **/usr/lpp/mmfs/4.2.0.0/installer/cookbooks/zimon_on_gpfs/files/default** Copy these files to **/opt/zimon** on all CES nodes. If the installer node is a CES node, also copy these files to **/opt/zimon** on the installer node.
5. Check GPFS cluster state on all nodes (double check the CES nodes since they may be down)

```
/usr/lpp/mmfs/bin/mmgetstate -a
```

6. Start GPFS on any down nodes

```
/usr/lpp/mmfs/bin/mmstartup -a
```

7. Verify GPFS becomes active on all nodes:

```
/usr/lpp/mmfs/bin/mmgetstate -a
```

8. Restart SMB on all CES nodes

```
/usr/lpp/mmfs/bin/mmces service start SMB -a
```

9. Start the **pmsensors** service on all CES nodes. ssh to each CES node and execute:

```
systemctl start pmsensors
```

10. Verify cluster state and service state on all nodes

```
/usr/lpp/mmfs/bin/mmgetstate -a  
/usr/lpp/mmfs/bin/mmces service list -a
```

11. Resume the deploy

```
/usr/lpp/mmfs/4.2.0.0/installer/spectrumscale deploy
```

12. Verify CES service states on all nodes after the deploy is successful.

```
/usr/lpp/mmfs/bin/mmces service list -a  
/usr/lpp/mmfs/bin/mmces state cluster
```

- Occasionally the cluster-wide knowledge of the state of the protocol nodes (viewable with the **mmces state cluster** command) may become out of sync with the local state (viewable with the **mmces state show** command) for some nodes. The most common case where this may occur is upgrading a system that has SMB enabled and uses Active Directory authentication. In order to bring the nodes back to the correct state the monitors on the affected nodes need to be restarted. This can be done by running the **mmcesmoncontrol restart** command on the nodes that have inconsistent state information.
- There is also a known issue that it is possible to hit on clusters when performing an upgrade with SMB enabled. This issue occurs very rarely but is most frequent if using Active Directory based authentication. What this issue looks like and the steps to fix it are:

1. If the upgrade fails during SMB upgrade the system state will look something like this:

```
$ mmlscluster --ces
Node Daemon node name IP address CES IP address list
-----
3 node01 172.31.132.1 node failed
4 node02 172.31.132.2 node failed
5 node03 172.31.132.3 Node suspended
6 node04 172.31.132.4 node failed
7 node05 172.31.132.5 Node suspended
8 node06 172.31.132.6 node failed
```

2. Recovery from this failure requires completing the SMB upgrade manually. The first step is to stop SMB on all nodes

```
$ mmdsh -N cesNodes /usr/lpp/mmfs/bin/mmces service stop SMB
```

3. Check which nodes have already been upgraded to the newer gpfs-smb level

```
$ mmdsh -N cesNodes "rpm -qa | grep grep gpfs.smb"
node01: gpfs.smb-4.3.0_gpfs_8-1.el7.x86_64
node02: gpfs.smb-4.2.2_gpfs_31-1.el7.x86_64
node03: gpfs.smb-4.3.0_gpfs_8-1.el7.x86_64
node04: gpfs.smb-4.2.2_gpfs_31-1.el7.x86_64
node05: gpfs.smb-4.3.0_gpfs_8-1.el7.x86_64
node06: gpfs.smb-4.2.2_gpfs_31-1.el7.x86_64
```

4. Manually upgrade the rpm on each of the nodes that are down-level.
 - a. Copy the newer rpm to each node (replace /usr/lpp/mmfs/4.2.0.0/ with the directory you extracted the GPFS self-extracting package to if you chose a different location)

```
$ scp /usr/lpp/mmfs/4.2.0.0/smb_rpms/gpfs.smb-4.3.0_gpfs_8-1.el7.x86_64.rpm node02:/tmp/
```

- b. Use rpm to upgrade the package on each node

```
$ rpm -U /tmp/gpfs.smb-4.3.0_gpfs_8-1.el7.x86_64.rpm
```

- c. Check that the level is as expected using the rpm query above. If all nodes are now at the same **gpfs.smb** version then you can restart SMB on all nodes

```
$ mmdsh -N cesNodes /usr/lpp/mmfs/bin/mmces service start SMB
```

- d. Determine which are the suspended nodes

```
$ mmlscluster --ces
Node Daemon node name IP address CES IP address list
-----
3 node01 172.31.132.1 node failed
4 node02 172.31.132.2 node failed
5 node03 172.31.132.3 Node suspended
6 node04 172.31.132.4 node failed
7 node05 172.31.132.5 Node suspended
8 node06 172.31.132.6 node failed
```

- e. Resume the suspended nodes.

```
$ mmces node resume -N node03,node05
```

Q12.5:
How can I determine if the Installation Toolkit successfully upgrades from IBM Storage Scale V4.1.1 to V4.2?

A12.5:

To check if the Installation Toolkit successfully upgraded from IBM Storage Scale V4.1.1 to V4.2, issue the following command:

```
./spectrumscale upgrade -po
```

If the only error condition displayed is the following, the upgrade completed successfully

```
"TypeError: sequence item 0: expected string, Node found"
```

Q12.6:

Can I have EPEL repos enabled when using the spectrumscale installation toolkit for install or upgrade?

A12.6:

EPEL repos must be disabled on all nodes that have been added to the IBM Storage Scale installation toolkit when attempting to install, deploy or upgrade.

See the Flash at <http://www-01.ibm.com/support/docview.wss?uid=ssg1S1009275>.

Q12.7:

What are the functions that are not supported by the installation toolkit?

A12.7:

For detailed information on functions that are not supported by the installation toolkit, see: [Limitations of the spectrumscale installation toolkit](#).

Q12.8:

What are the potential limitations when using the Installation Toolkit with SLES12 SP1 or SP2 nodes?

A12.8:

There is a potential issue with IBM Storage Scale packaged with SMB and SLES version of samba-winbind that can lead to an installation failure. If this occurs, remove the samba-winbind rpms and continue the installation and deployment. For more information, see [Package conflict on SLES 12 SP1 and SP2 nodes while doing installation, deployment, or upgrade using installation toolkit](#).

Transparent cloud tiering (discontinued) questions

Important:

TCT is now discontinued. Our strategic direction for cloud object storage tiering is IBM Storage Scale Active File Management (AFM). If you are currently using TCT, you must migrate to AFM before you upgrade to 5.2.0; otherwise, the preexisting TCT environment will be lost. For information about how to migrate from TCT to AFM, see the documented processes for [filesets](#) and for [file systems](#).

Q13.1:

What platforms is transparent cloud tiering supported on?

A13.1:

The transparent cloud services must be installed on CES protocol nodes or NSD nodes that are running RHEL 7 or RHEL 8. Both x86 and IBM POWER8 and later servers are supported. In order to transparently recall a file that has been migrated via transparent cloud tiering, a node must be running: RHEL, SLES, or Ubuntu on x86 or RHEL on POWER8 Little Endian. The IBM Storage Scale cluster might include nodes with other platforms or operating systems, but these nodes will not be able to migrate or recall files directly.

Note: To enable transparent cloud tiering nodes, you must first enable the transparent cloud tiering feature. These nodes must have GPFS server licenses enabled. This feature provides a new level of storage tiering capability to the IBM Storage Scale customer. Please contact your IBM Client Technical Specialist (or send an email to scale@us.ibm.com) to review your use case of the transparent cloud tiering feature and to obtain the instructions to enable the feature in your environment.

Q13.2

How many transparent cloud tiering bridge nodes are supported?

A13.2:

Up to four cloud services nodes can be configured per file system. Up to four file systems can be managed in a cluster.

Q13.3:

What are the networking requirements for transparent cloud tiering?

A13.3:

Transparent cloud tiering utilizes the cloud services node in order to communicate with an external storage cloud. Typically, this communication will utilize standard HTTP or HTTPS TCP ports (port 80 or 443). However, some storage cloud providers may use other TCP ports. Please check with your cloud provider for details. The bridge node must be able to communicate to the storage cloud. Prior to migrating files to a cloud provider, ensure that there is sufficient bandwidth to both send and receive files as needed. The bandwidth required will vary based on workload and user requirements.

Q13.4:**What Cloud Object Storage Providers are supported?****A13.4:**

For information on the supported Cloud Object Storage Providers, see [Supported cloud providers](#).

Q13.5:**Can transparent cloud tiering be used with other IBM Storage Scale functions?****A13.5:**

Considerations for using transparent cloud tiering with IBM Storage Scale include:

- IBM Spectrum Archive - Linear Tape File Systems (LTFS) and IBM Storage Protect for Space Management (HSM)

Running IBM Spectrum Archive and transparent cloud tiering on the same file system is not supported. However, both HSM and transparent cloud tiering can coexist on the same systems (as long as they are configured with different file systems)

- AFM
 - Running transparent cloud tiering service on the AFM gateway nodes is not supported.
 - Data from the AFM and AFM DR filesets must not be accessed by transparent cloud tiering.
- FPO

Transparent cloud tiering is supported on an FPO cluster.

- IBM Storage Scale Swift Object

Transparent cloud tiering can be configured on IBM Storage Scale Swift Object fileset(s) only. Support for native object storage is not provided.

- Snapshots
 - Transparent cloud tiering cannot be used to migrate/recall snapshots.
 - Space contained in snapshots will not be freed if files are migrated to cloud object storage.
- Sparse Files

Transparent cloud tiering can be used to migrate and recall sparse files, but sparseness will not be retained. Full blocks will be allocated.

- Native Encryption

Transparent cloud tiering can be used with IBM Storage Scale built-in encryption. All data migrated to Cloud Object Storage will be migrated with the encryption key configured in transparent cloud tiering. When the file is read from the filesystem, the data will be unencrypted and re-encrypted using transparent cloud tiering encryption algorithms prior to being sent to cloud storage.

- Compression

Transparent cloud tiering can be used along with Scale file system level compression capability. When the file is read from file system, the file will be uncompressed, transparent cloud tiering will transfer the uncompressed file to cloud storage. Recalled files will be uncompressed on the file system.

- CES nodes (Protocol Services)

Transparent cloud tiering can coexist along with NFS, SMB or Swift Object Services on the CES nodes.

- IBM Storage Protect.

Note: Beginning with Version 7.1.3, IBM Tivoli Storage Manager is now IBM Storage Protect.

Files should be backed up prior to transferring them to cloud storage via transparent cloud tiering. Failure to do so will cause files to be recalled in order to perform the back up.

Q13.6:**Is transparent cloud tiering supported with IBM Storage Scale System?****A13.6:**

Transparent cloud tiering cannot be deployed directly on IBM Storage Scale System nodes, however it can be deployed on other nodes in the IBM Storage Scale cluster that meet the hardware and software requirements.

Q13.7:**Are there restrictions on viewing or listing files?****A13.7:**

Standard Unix tools and windows utilities such as **ls** or **dir** can be used to view files that have been migrated to the cloud. Some file viewers, such as Windows Explorer and GNOME File viewer utilize *preview* functions which will open files in order to generate a preview. These functions may result in files being unintentionally recalled from the cloud.

Q13.8:**What levels of operating systems are supported by transparent cloud tiering services in IBM Storage Scale?****A13.8:**

For more information about the OS matrix, see *Software requirements for Cloud services* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Licensing and Pricing Questions

Q14.1:**Where can I find detailed information about IBM Storage Scale and IBM Storage Scale System licensing and pricing?****A14.1:**

For more information about IBM Storage Scale and IBM Storage Scale System licensing and pricing, see https://www.ibm.com/docs/en/spectrum-scale?topic=STXKQY/IBMScale_ESS_Licensing.pdf.

Q14.2:**How is IBM Storage Scale licensed?****A14.2:**

IBM Storage Scale is available in the following editions, which are licensed individually:

- The following editions are currently available:
 - IBM Storage Scale **Data Access Edition** (DAE)
 - IBM Storage Scale **Data Management Edition** (DME)
 - IBM Storage Scale **Erasure Code Edition** (ECE)

Note: These IBM Storage Scale editions are licensed by capacity: per terabyte (TiB).

- The following editions are no longer available:
 - IBM Storage Scale **Express Edition**
 - IBM Storage Scale **Standard Edition**
 - IBM Storage Scale **Advanced Edition**

Note:

- These IBM Storage Scale editions are licensed per socket with options for client, server, and FPO servers.
- Existing licensees with active entitlement can renew and add licenses. You can renew existing socket-based licenses through normal renewal channels. To add licenses, contact your IBM representative or Business Partner because parts are not available by normal ordering processes.
- Entitlement to purchase more socket-based licenses is determined by the IBM Customer ID.
- Earlier versions of GPFS (General Parallel File System, the previous generation of IBM Storage Scale) were licensed by processor core. These licenses are expired. If you have these older licenses and want to extend entitlement, contact your IBM representative.

Q14.3:**How can I determine the number of licenses that I need? How is capacity measured?****A14.3:**

Licensing of current versions and editions of IBM Storage Scale is capacity-based only.

- For capacity-based licenses:
 - Per-TB, which for IBM Storage Scale licensing is defined as binary, where 1 TB is $2^{40}=2^40$.
 - Per-PB, which for IBM Storage Scale licensing is defined as binary, where 1 PB is $2^{50}=2^50$.
 - Per-drive (if purchased with an IBM Elastic Storage System):

- ◆ Different licenses for disk storage and solid-state storage (NVMe, SSD, etc.).
- ◆ Per-drive is considered a capacity-based license for intermixing with per-TB and per-PB licenses within the same clusters.
- ◆ If thin provisioning is supported, entitlement is based on the provisioned capacity.
- ◆ For specific licensing requirements, see https://www.ibm.com/docs/en/spectrum-scale?topic=STXKQY/IBMScale_ESS_Licensing.pdf.

Q14.4:

If I have IBM Storage Scale Advanced Edition or IBM Storage Scale Standard Edition, can I move to storage capacity licensing?

A14.4:

Yes, contact your IBM representative or Business Partner for more information. Each situation needs to be uniquely evaluated.

Q14.5:

What are the differences between the IBM Storage Scale editions?

A14.5:

- **IBM Storage Scale Data Access Edition:** Includes the base IBM Storage Scale functions, including ILM between online storage tiers or tape (IBM Spectrum Archive or IBM Storage Protect licenses are required for ILM to tape), AFM, multi-cluster mount, synchronous replication, and integrated protocol access methods (NFS, SMB, S3, and HDFS).
- **IBM Storage Scale Data Management Edition:** Includes all of the features in the Data Access Edition, plus multisite replication, native encryption for secure storage and secure deletion, Asynchronous Disaster Recovery (AFM DR), file audit logging and clustered watch folder to track user access to the file system and events across all nodes and all protocols, ILM tiering to and from onsite and Cloud-based Swift Object storage; and exports to and from Swift Object storage (sync to Cloud).
- **IBM Storage Scale Erasure Code Edition:** Includes all of the features in the Data Management Edition, plus enterprise-grade durability on commodity storage rich server hardware.

For more information about withdrawn editions, see https://www.ibm.com/support/knowledgecenter/STXKQY/IBMScale_ESS_Licensing.pdf.

Q14.6:

How can I tell which edition of IBM Storage Scale I am running?

A14.6:

To determine which edition of IBM Storage Scale you are running, you can execute the `mm1slicense` command. It will tell you which edition is running on the local node.

Q14.7:

Can I migrate my GPFS V3.5 licenses to IBM Storage Scale?

A14.7:

If you have valid entitlement and active subscription and support, your licenses can be manually migrated to a capacity license. While the base license can be migrated to IBM Storage Scale, subscription and support must be reinstated for capacity of the entire cluster. Contact your IBM representative for more information.

Q14.8:

Can I transfer my IBM Storage Scale licenses between machines?

A14.8:

Yes, licenses are not tied to any particular machine. For more information, see https://www.ibm.com/support/knowledgecenter/STXKQY/IBMScale_ESS_Licensing.pdf.

PVU and socket-based licenses can be registered to a particular server in IBM systems. This does not affect your entitlement.

Q14.9:

Do I require more IBM Storage Scale licenses for my Disaster Recovery (DR) environment?

A14.9:

Yes. Regardless of whether your IBM Storage Scale is licensed based on PVUs, sockets, or capacity, the determination on whether more licenses are needed depends on whether "work" is being done and whether the DR location is "cold," "warm," or "hot". Since that answer is specific to your situation, it is important to see how IBM's policy aligns with your relevant dynamics.

Q14.10:

What are IBM Storage Scale OEMs, and how are their products licensed?

A14.10:
IBM Storage Scale licenses can be used with storage systems and servers from any vendor provided those systems comply with the supported hardware definitions and supported operating environments for IBM Storage Scale as documented in the IBM Knowledge Center. In addition, some IBM solution providers incorporate IBM Storage Scale with specific software and hardware to create complete system solutions. These partners are known as OEMs, and their offerings are known as OEM solutions. The solution includes support for IBM Storage Scale, which is provided by the OEM and embedded in the solution.

Q14.11:
Can I get support for IBM Storage Scale from IBM if it is supplied as part of an OEM solution?

A14.11:
An OEM solution is supported completely by the partner (OEM) that supplies it, which includes support for IBM Storage Scale when used in that solution. The included IBM Storage Scale licenses do not entitle you to support directly from IBM.

If you obtain separate IBM Storage Scale software licenses from IBM for use with an OEM solution, be aware that the OEM solution might include hardware, software, or configuration that is not supported by IBM itself. IBM Support might ask you to recreate reported problems on a supported system or configuration. IBM will not be able to provide support for the solution as a whole; only the OEM can do that.

Q14.12:
If I install IBM Storage Scale software and licenses from IBM on an OEM system from another vendor, is that software supported by IBM?

A14.12:
IBM tests and supports IBM Storage Scale software on supported platforms as documented in the IBM Knowledge Center. OEM systems from other vendors often contain customized components or unique additions, including operating system releases or distributions, that are not tested or supported by IBM. If that is the case for a given OEM system, it is considered an unsupported platform by IBM.

Q14.13:
Can I mix products licensed from other vendors that embed IBM Storage Scale (OEMs) in the same cluster as IBM Storage Scale licenses from IBM?

A14.13:
No, systems from OEM vendors cannot be part of the same cluster as systems licensed under IBM Storage Scale licenses. Each cluster must be supported by one vendor only. If you wish to integrate OEM systems and IBM systems in the same environment, use the multi-cluster capabilities of IBM Storage Scale to combine separate clusters for each vendor.

Q14.14:
Which platforms support IBM Storage Scale Backup?

A14.14:
IBM Storage Scale Backup is supported on the following platforms:

- IBM Power Systems (little endian)
- x86-64 servers
- IBM zSystems
- IBM LinuxONE enterprise servers

Q14.15:
What is the benefit of using IBM Storage Scale Backup?

A14.15:
IBM Storage Scale Backup simplifies and modernizes the existing licensing models by providing IBM Storage Scale and IBM Storage Scale System users with the ability to easily add data protection services under a front-end and capacity-based license model for new and existing clients.

Q14.16:
Who are the potential users of IBM Storage Scale Backup?

A14.16:
IBM Storage Scale Backup is useful for IBM Storage Scale clients who are looking to protect and manage space in their IBM Storage Scale environments. The new licensing model creates a simple licensing option based on front-end terabytes for IBM Storage Scale and IBM Storage Scale System solutions.

Q14.17:**Which IBM Storage Scale editions support IBM Storage Scale Backup?****A14.17:**

The following IBM Storage Scale Editions support IBM Storage Scale Backup:

- IBM Storage Scale Data Access Edition
- IBM Storage Scale Data Management Edition
- IBM Storage Scale Erasure Code Edition
- IBM Storage Scale Data Management Edition for IBM Storage Scale System
- IBM Storage Scale Data Access Edition for IBM Storage Scale System

Q14.18:**May I use IBM Storage Scale Backup with existing deployments of IBM Storage Scale, IBM Storage Scale System and IBM Storage Protect?****A14.18:**

Yes. IBM Storage Scale Backup provides a simple licensing model to allow IBM Storage Scale and IBM Storage Scale System customers to purchase IBM Storage Protect Extended Edition and IBM Storage Protect for Space Management. There are no restrictions in using IBM Storage Scale Backup licenses with existing IBM Storage Protect deployments backing up IBM Storage Scale and IBM Storage Scale.

Q14.19:**What is the benefit of using IBM Storage Fusion Data Cataloging Services?****A14.19:**

IBM Storage Fusion Data Cataloging Services provides unified metadata management and insights for heterogeneous unstructured data, on-premises and in the cloud. Data Cataloging Services delivers the following key capabilities:

- **Discover:** Automatically ingest and index system metadata from multiple file and object storage systems, on-premises and in the cloud.
- **Classify:** Automatically identify and classify data, including sensitive and personal identifiable information.
- **Label:** Enrich data with system and custom metadata tags that increase the value of that data.
- **Find:** Find data quickly and easily by searching catalogs of system and custom metadata.

Q14.20:**Who are the potential users of IBM Storage Fusion Data Cataloging Services?****A14.20:**

IBM Storage Scale Data Management Edition and IBM Storage Scale Erasure Code Edition clients are the potential users who are entitled to use IBM Storage Fusion cataloging services to enable a comprehensive unstructured data catalog to simplify AI data organization for machine learning and business analytics initiatives.

Q14.21:**Which IBM Storage Scale editions support IBM Storage Fusion Data Cataloging Services?****A14.21:**

The following IBM Storage Scale editions support IBM Storage Fusion Data Cataloging Services from the 5.1.8 release:

- IBM Storage Scale Data Management Edition
- IBM Storage Scale Erasure Code Edition
- IBM Storage Scale Data Management Edition for IBM Storage Scale System

Service questions

Q15.1:**What support services are available for IBM Storage Scale?****A15.1:**

The following support services are included:

- IBM Support Guide: <https://www.ibm.com/support/pages/ibm-support-guide>.
- Forums:
 - Technical discussion forum: [IBM Storage Community](#).
 - For the latest announcements and news, subscribe to the IBM community: <https://community.ibm.com/community/user/home>.

- Notifications:

Customize your support portal: <http://www-01.ibm.com/software/support/einfo.html>.

- IBM Global Services - Support Line for Linux

A 24x7 enterprise-level remote support for problem resolution and defect support for major distributions of the Linux operating system. Go to www.ibm.com/services/us/index.wss/so/its/a1000030.

- IBM Systems Lab Services

IBM Systems Lab Services can help you optimize the utilization of your data center and system solutions.

Lab Services has the knowledge and deep skills to support you through the entire information technology race. Focused on the delivery of new technologies and niche offerings, Lab Services collaborates with IBM Global Services and IBM Business Partners to provide complementary services that will help lead through the turns and curves to keep your business running at top speed.

Go to <http://www.ibm.com/systems/services/labservices/>.

- Software maintenance

Defect resolution for current holders of IBM software maintenance contracts:

- In the United States contact us toll free at 1-800-IBM-SERV (1-800-426-7378)
- In other countries, contact your local [IBM Service Center](#)

You may also visit the Support Portal online at https://www-947.ibm.com/support/entry/portal/Open_service_request/Software/Other_Software/General_Parallel_File_System?category=2&locale=en_us&lnk=mhsd

Contact scale@us.ibm.com for all other services or consultation on what service is best for your situation.

Q15.2:

How do I download fixes for IBM Storage Scale?

A15.2:

To download fixes, go to Fix Central: <https://www.ibm.com/support/fixcentral/>.

- Search for IBM Storage Scale.
- For earlier releases, search for General Parallel File System.

Q15.3:

What are the current advisories for all platforms supported by IBM Storage Scale?

A15.3:

For more information about the current advisories for all platforms, see [IBM Storage Scale advisories](#).

Q15.4:

What are the current advisories for IBM Storage Scale on AIX?

A15.4:

For more information about the current AIX advisories, see [IBM Storage Scale advisories](#).

Q15.5:

What are the current advisories for IBM Storage Scale?

A15.5:

To consult a list of the latest security bulletins for IBM Storage Scale, go to [IBM Product Security Central](#).

For a more comprehensive and customized list of critical updates of your IBM products, subscribe to [My Notifications](#). The sign up instructions and a demonstration are available on [Stay up to date with My Notifications](#).

Q15.6:

Where can I find the IBM Storage Scale Software License Agreement?

A15.6:

For more information about license information documents, see <http://www.ibm.com/software/sla/sladb.nsf>. To search for a specific program license agreement, search for IBM Storage Scale.

Q15.7:

Where can I find End of Market (EOM) and End of Service (EOS) information?

A15.7:

IBM Storage Scale EOS dates can be found at the IBM support lifecycle page: <https://www.ibm.com/software/support/lifecycle/lc-policy.html>. IBM Storage Scale follows the *Standard* IBM Support Lifecycle Policy.

EOM and EOS information can also be found in the IBM Sales Manual pages on the IBM Offering Information site for the program:

1. Go to <https://www.ibm.com/support/pages/lifecycle>.
2. Enter **IBM Storage Scale** in the search box.
3. Click the search button.

Q15.8:

Where can I locate IBM Storage Scale code to upgrade from my current level?

A15.8:

If you have active entitlement, you may log into Passport advantage or Fix Central and upgrade your level of IBM Storage Scale.

Q15.9:

What are the release practices for IBM Storage Scale?

A15.9:

For information about the IBM Storage Scale release practices, see the following IBM Support page: [IBM Storage Scale Software Version Recommendation Preventive Service Planning](#).

Scale Management Server (REST API) questions

Q15.1:

Does IBM Storage Scale support RESTful APIs for file system management?

A16.1:

In v4.2.2 and later, IBM Storage Scale supports a REST API for configuring, managing, and monitoring various components of IBM Storage Scale system. The IBM Storage Scale REST API is an HTTP programming interface for performing IBM Storage Scale management tasks. With the REST API, you can automate storage management operations and integrate IBM Storage Scale capabilities into your applications. The APIs are installed on the GUI stack of the IBM Storage Scale cluster. The GUI installation and setup takes care of the API installation. You do not need to perform any additional steps to set up APIs. Clients communicate using HTTPS protocol and JSON syntax is used to frame data inside HTTP requests and responses.

Q15.2:

What are the changes made to the API implementation in the 4.2.3 release?

A16.2:

The API Version 1 was introduced with the IBM Storage Scale 4.2.2 release. The implementation was based on Python and the deployment was limited only to the manager nodes that run on RHEL7. The API Version 2 is introduced in the 4.2.3 release. The current implementation is based on GUI stack. That is, GUI server manages and processes the API requests and commands. Version 2 has the following features:

- Reuses the GUI deployment's backend infrastructure, which makes introduction of new API commands easier.
- Uses the same role-based access feature that is available to authenticate and authorize the GUI users. No additional configuration is required for the API users.
- Makes deployment easier as the GUI installation takes care of the basic deployment.
- Supports filtering of objects and paging if several thousand objects are retrieved.
- Highly scalable and can support large clusters with thousands of nodes.
- The APIs are driven by the same lightweight WebSphere® Liberty server and object cache that is used by the IBM Storage Scale GUI.

Note: Although the REST API delivered with IBM Storage Scale V4.2.3 still supports version 1 requests, it is highly recommended that you switch to REST API version 2 requests at your earliest convenience since version 1 is deprecated and will not be enhanced.

File audit logging and clustered watch questions

Q16.1:

What are the requirements and limitations for file audit logging and clustered watch folder?

A17.1:

For more information, see *Requirements, limitations, and support for file audit logging* and *Requirements, limitations, and support for clustered watch folder* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q16.2:

Is the message queue still supported in IBM Storage Scale 5.1.1 and later?

A17.2:

No. From IBM Storage Scale 5.1.1 and later we no longer support the use of the message queue with file audit logging and clustered watch folder. If you are upgrading to IBM Storage Scale 5.1.2 from a code level that still has the message queue, you must disable file audit logging and clustered watch folder and then run `mmmsgqueue config --remove` prior to the upgrade, or upgrade to IBM Storage Scale 5.1.0.x, run `mmmsgqueue config --remove-msgqueue`, and then proceed with the upgrade to IBM Storage Scale 5.1.2. If upgrading from IBM Storage Scale 5.1.1, you can upgrade to IBM Storage Scale 5.1.2 without removing the message queue and disabling file audit logging and clustered watch folder as the message queue should already be removed.

Support of Docker, Podman containers and use in Kubernetes and Red Hat OpenShift environments questions

Q18.1:

Can IBM Storage Scale storage be utilized inside containerized environments?

A18.1:

Yes. There are multiple ways to utilize IBM Storage Scale storage inside containerized environments. Two of the ways are described as follows:

1. RHEL or Ubuntu worker nodes under Kubernetes or Red Hat OpenShift Container Platform:

IBM Storage Scale Container Storage Interface (CSI) driver and operator is a free open-source download, enabling provisioning of volumes that are filesets or directory paths in a preconfigured file system, to the containers. Volumes can be dynamically provisioned, or existing filesets or directory paths can be used.

In this CSI configuration, classic non-containerized IBM Storage Scale is loaded directly upon RHEL or Ubuntu worker nodes. These worker nodes will belong to a Kubernetes or OpenShift cluster. The IBM Storage Scale CSI driver and operator will run as containers to allow provisioning of the underlying IBM Storage Scale storage up and into the container applications.

For more information, see the [IBM Storage Scale CSI Documentation](#).

2. Red Hat CoreOS worker nodes under Red Hat OpenShift Container Platform:

Coupling IBM Storage Scale Container Native with IBM Storage Scale CSI allows for a fully containerized deployment of IBM Storage Scale on Red Hat CoreOS worker nodes, where the classic non-containerized IBM Storage Scale packages cannot be installed.

In this Container Native + CSI configuration, an OpenShift cluster consisting of Red Hat CoreOS worker nodes pre-exists, and some or all the worker nodes are designated to host the IBM Storage Scale containers. IBM Storage Scale Container Native is a containerized version of IBM Storage Scale in which the components such as the IBM Storage Scale daemon, the GUI Rest-API, Performance Monitoring, Health monitoring run as pods and sidecars, spread across designated worker nodes of an OpenShift cluster. The IBM Storage Scale Container Native operator controls the deployment, cluster configuration, and overall cluster activities. This containerized IBM Storage Scale cluster is a client cluster (Container Native Storage Access or CNSA) which will remote mount storage from a non-containerized IBM Storage Scale or IBM Storage Scale System storage cluster. The storage is dynamically provisioned for use with container applications by IBM Storage Scale CSI. The IBM Storage Scale GUI node is required on the storage cluster to act as a REST-API server for all the operator actions.

For more information, see [IBM Storage Scale Container Native Documentation](#).

Q18.2:

Why do we need to use IBM Storage Scale Container Storage Interface? Can we bind mount IBM Storage Scale directly into docker containers?

A18.2:

The IBM Storage Scale CSI driver offers the ability for container applications to provision dynamically created volumes that map to filesets, which can be existing filesets or a dynamically created fileset, into containers. If your use case does not need the flexibility that is offered by the IBM Storage Scale CSI driver, an alternative method is to bind mount the file system path inside of the container.

Q18.3:

What are the prerequisites and compatibility matrix for using IBM Storage Scale as a backend for containers?

A18.3:

IBM Storage Scale CSI Compatibility Matrix

Use the following matrix with CSI and non-containerized IBM Storage Scale to determine which level of IBM Storage Scale, OpenShift/Kubernetes, and OS/arch to use with each release:

Table 49: IBM Storage Scale CSI Compatibility Matrix

CSI	Architecture ²	non-containerized IBM Storage Scale level for worker nodes	IBM Storage Scale level if remote cluster is used	OCP level
CSI 2.12.x	x86, ppc64le	5.1.9.x or later	5.1.9.x or later	4.14, 4.15
CSI 2.13.x	x86, ppc64le	5.1.9.x or later	5.1.9.x or later	N/A
CSI 2.14.x	x86, ppc64le	5.1.9.x or later	5.1.9.x or later	N/A
CSI 3.0.x	x86, ppc64le	5.2.3.x or later	5.2.3.x or later	N/A

Notes:

- ¹ OpenShift levels previous to 4.16 and Kubernetes levels previous 1.30 are no longer maintained; and their status is considered as end of support (EOS).
- ² The minimum supported ppc64le architecture is Power9.
- ³ Ubuntu 24.04 is supported by IBM Storage Scale Container Storage Interface driver 2.13.1 and later versions.
- ⁴ Support for OpenShift with RHEL worker nodes was discontinued with IBM Storage Scale Container Storage Interface 2.13.x.
- ⁵ Kubernetes 1.33 supports CSI 2.14.4 and later.
- ⁶ RHEL 10 and RHEL 10.1 are supported by IBM Storage Scale Container Storage Interface driver 3.0.1 and later.

For more information about CSI prerequisites, see [IBM Storage Scale CSI Documentation](#).

IBM Storage Scale Container Native and CSI compatibility matrix

Use the following matrix with IBM Storage Scale Container Native and CSI, to determine which level of CSI, IBM Storage Scale remote cluster, and OpenShift to use with each release:

Table 50: IBM Storage Scale Container Native and CSI compatibility matrix

Container Native	CSI	Architecture	Containerized IBM Storage Scale level	Remote non-containerized storage cluster IBM Storage Scale level	OCP level	UBI level*	RHCOS	vanilla K8s level	IBM Cloud Satellite	RHEL	Ubuntu
CNSA 5.1.9.12	CSI 2.10.10	x86, ppc64le, s390x	5.1.9.12	5.1.3.0+	4.12, 4.13, 4.14	9.6	4.12, 4.13, 4.14	N/A	N/A	N/A	N/A
CNSA 5.2.0.1	CSI 2.11.1	x86, ppc64le, s390x	5.2.0.1	5.1.3.0+	4.13, 4.14, 4.15	9.3	4.13, 4.14, 4.15	N/A	N/A	N/A	N/A
CNSA 5.2.1.1	CSI 2.12.1	x86, ppc64le, s390x	5.2.1.1	5.1.9.0+	4.14, 4.15, 4.16	9.4	4.14, 4.15, 4.16	N/A	N/A	N/A	N/A
CNSA 5.2.2.1	CSI 2.13.1	x86, ppc64le, s390x	5.2.2.1	5.1.9.0+	4.15, 4.16, 4.17, 4.18	9.5	4.15, 4.16, 4.17, 4.18	N/A	N/A	N/A	N/A

Table 50: IBM Storage Scale Container Native and CSI compatibility matrix (continued)

Container Native	CSI	Architecture	Containerized IBM Storage Scale level	Remote non-containerized storage cluster IBM Storage Scale level	OCP level	UBI level*	RHCOS	vanilla K8s level	IBM Cloud Satellite	RHEL	Ubuntu
CNSA 5.2.3.4	CSI 2.14.3	x86, ppc64le, s390x	5.2.3.4	5.1.9.0+	4.16, 4.17, 4.18, 4.19	9.6	4.16, 4.17, 4.18, 4.19	1.29, 1.30, 1.31, 1.32	N/A	8.10, 9.4, 9.5**9.6*	22.04, 24.04**
CNSA 5.2.3.5	CSI 2.14.4	x86, ppc64le, s390x	5.2.3.5	5.1.9.0+	4.16, 4.17, 4.18, 4.19, 4.20	9.7	4.16, 4.17, 4.18, 4.19, 4.20	1.29, 1.30, 1.31, 1.32, 1.33	N/A	8.10, 9.4, 9.5**9.6* *9.7**	22.04, 24.04**
CNSA 5.2.3.6	CSI 2.14.5	x86, ppc64le, s390x	5.2.3.6	5.1.9.0+	4.16, 4.17, 4.18, 4.19, 4.20	9.7	4.16, 4.17, 4.18, 4.19, 4.20	1.29, 1.30, 1.31, 1.32, 1.33	N/A	8.10, 9.4, 9.5**9.6* *9.7**	22.04, 24.04**
CNSA 5.2.3.7	CSI 2.14.6	x86, ppc64le, s390x	5.2.3.7	5.1.9.0+	4.16, 4.17, 4.18, 4.19, 4.20	9.7	4.16, 4.17, 4.18, 4.19, 4.20	1.29, 1.30, 1.31, 1.32, 1.33	N/A	8.10, 9.4, 9.5**9.6* *9.7**	22.04, 24.04**
CNSA 6.0.0.0	CSI 3.0.0	x86, ppc64le, s390x, aarch64**	6.0.0.0	5.2.0.0+	4.17, 4.18, 4.19, 4.20	9.6	4.17, 4.18, 4.19, 4.20	1.29, 1.30, 1.31, 1.32, 1.33	N/A	8.10, 9.4, 9.5**9.6* *9.7**	22.04, 24.04**
CNSA 6.0.0.1	CSI 3.0.1	x86, ppc64le, s390x, aarch64**	6.0.0.1	5.2.0.0+	4.17, 4.18, 4.19, 4.20	9.7	4.17, 4.18, 4.19, 4.20	1.29, 1.30, 1.31, 1.32, 1.33	N/A	8.10, 9.4, 9.5**9.6* *9.7**	22.04, 24.04**
CNSA 6.0.0.2	CSI 3.0.2	x86, ppc64le, s390x, aarch64**	6.0.0.2	5.2.0.0+	4.17, 4.18, 4.19, 4.20, 4.21	9.7	4.17, 4.18, 4.19, 4.20, 4.21	1.29, 1.30, 1.31, 1.32, 1.33, 1.34	N/A	8.10, 9.4, 9.5**9.6* *9.7**	22.04, 24.04**

- * UBI (universal base image) level reflects the internal packaging and build of each IBM Storage Scale Container Native level. This cannot be changed. It is shown to understand and track future compatibility.
- ** IBM Storage Scale Container Native with Vanilla Kubernetes on RHEL and Ubuntu is available only for the x86 architecture.
- *** Red Hat OpenShift Container Platform (OCP) is not supported on AArch64 (ARM64) architecture.

For more information about IBM Storage Scale Container Native prerequisites, see [IBM Storage Scale Container Native Documentation](#).

Q18.4:

Supported Upgrade Paths for IBM Storage Scale CSI and IBM Storage Scale Container Native

A18.4:

When planning an upgrade, it is important to understand the following:

- Come-from and go-to possibilities of each component involved.

- Recommended order of upgrade for each component.
- Overall support statements for each component.

The combination of this information will determine overall support of the solution and the upgrade strategy.

Come-from and go-to possibilities of each component involved

Table 51: IBM Storage Scale Container Storage Interface (CSI) upgrade paths

Upgrading from:	To: CSI 2.10.x	To: CSI 2.11.x	To: CSI 2.12.x	To: CSI 2.13.x	To: CSI 2.14.x	To: CSI 3.0.x
CSI 2.10.x	✓	✓	✓	✓	--	--
CSI 2.11.x	--	✓	✓	✓	--	--
CSI 2.12.x	--	--	✓	✓	✓	--
CSI 2.13.x	--	--	--	✓	✓	✓
CSI 2.14.x	--	--	--	--	✓	✓
CSI 3.0.x	--	--	--	--	--	✓

Table 52: IBM Storage Scale Container Native upgrade paths

Upgrading from:	To: CNSA 5.1.9.x	To: CNSA 5.2.0.x	To: CNSA 5.2.1.x	To: CNSA 5.2.2.x	To: CNSA 5.2.3.x	To: CNSA 6.0.0.x
CNSA 5.1.9.x (December 2023)	✓	✓	X	X	X	X
CNSA 5.2.0.x (April 2024)	--	✓	✓	✓	X	X
CNSA 5.2.1.x (August 2024)	--	--	✓	✓	✓	X
CNSA 5.2.2.x (August 2024)	--	--	--	✓	✓	✓
CNSA 5.2.3.x (May 2025)	--	--	--	--	✓	✓
CNSA 6.0.0.x (October 2025)	--	--	--	--	--	✓

Recommended order of upgrade

When considering an upgrade of OpenShift or Kubernetes, first check the IBM Storage Scale Container Native or IBM Storage Scale Container Storage Interface driver CSI compatibility matrix tables to ensure compatibility.

Likewise, when considering an upgrade of either IBM Storage Scale Container Native or IBM Storage Scale Container Storage Interface driver, check the respective compatibility matrix tables to ensure compatibility with the installed OpenShift or Kubernetes versions.

Support Statements for each component:

- IBM Storage Scale recommends keeping to the latest supported versions of all the dependencies. Fixes for CSI and CNSA will be included with the latest release only.
- IBM Storage Scale CSI driver/operator is directly dependent upon the Kubernetes/OpenShift and IBM Storage Scale/CNSA levels.
- [Red Hat's support policy for OpenShift reflects continued support and fixes.](#)
- [Red Hat's OpenShift Container Platform 4.x Tested Integrations list.](#)
- [Kubernetes version skew support policy reflects continued support and fixes for the latest 3 releases of Kubernetes.](#)
- IBM Storage Scale will continue to support both the CSI and CNSA previous versions as long as their dependent K8s and/or OCP levels have not reached end of support, and so long as the dependent IBM Storage Scale levels have not reached end of support.
- If opening a support ticket against out of support levels, be aware that a recreate against a currently supported level, may be requested.

Q18.5:

Can Kubernetes perform a health check on the underlying IBM Storage Scale cluster?

A18.5:

On Kubernetes clusters that have classic IBM Storage Scale RPM or DEB packages that are installed upon RHEL or Ubuntu worker nodes, the monitoring is done like on any other IBM Storage Scale deployment and is not integrated into Kubernetes.

On IBM Storage Scale Container Native clusters that run on Red Hat OpenShift Container Platform or on a supported Kubernetes distribution, the IBM Storage Scale pods are fully integrated with Kubernetes and managed through a Kubernetes operator.

Q18.6:

How can one identify which filesets in IBM Storage Scale are created by the IBM Storage Scale CSI driver?

A18.6:

Filesets that are created by the IBM Storage Scale Container Storage Interface driver have the following comment tagged to them: "*Fileset created by IBM Container Storage Interface driver*". This can be viewed using the `mmfsfileset` command with `-Y` option.

Q18.7:

Can I use IBM Storage Scale System storage to provision volumes into my remote cluster and then act upon these volumes with IBM Storage Scale CSI?

A18.7:

For standalone CSI, the required IBM Storage Scale System version is 6.1.9 or higher. For more information, see the [IBM Storage Scale CSI Knowledge Center](#).

For IBM Storage Scale Container Native configurations, the IBM Storage Scale storage cluster must be running IBM Storage Scale 5.1.9.0 or higher. An IBM Storage Scale System can be utilized as the storage cluster as long as the IBM Storage Scale System code level is 6.1.9 or higher.

Q18.8:

Is it possible to contribute to the IBM Storage Scale CSI driver and operator?

A18.8:

Yes, a public open-source GitHub is available and can be accessed at the following link:

IBM Storage Scale CSI GitHub: <https://github.com/IBM/ibm-spectrum-scale-csi>

Q18.9

Where do I get IBM Storage Scale CSI and IBM Storage Scale Container Native?

A18.9

IBM Storage Scale Container Native

IBM Storage Scale Container Native images are released through IBM Cloud Container Image Registry and *not through IBM Support Fix Central*. Customers entitled to IBM Storage Scale Data Access Edition or IBM Storage Scale Data Management Edition gain entitlement to the container images.

Check [My IBM > ... > Entitlement keys](#) for an entitlement key that allows access to the images.

Documentation that is relevant to use this entitlement key include the following pages:

- Must read: [IBM Storage Scale Container Native documentation for full installation guidance as well as all prerequisites](#)
- [Red Hat OpenShift Container Platform / IBM Cloud Container Registry \(ICR\) pull secrets](#)
- [Planning / Air-gapped installs](#)
- [Planning / Container images](#)

Installation of IBM Storage Scale Container Native automatically pulls IBM Storage Scale CSI images. Therefore, it is not necessary to follow installation instructions or entitlement procedures that are specific to IBM Storage Scale CSI; access to any container image registries other than the IBM Cloud Container Registry is not necessary either.

If IBM Storage Scale was purchased through IBM Entitled Systems Support, an additional step is necessary before the entitled image bundles can be viewed with My IBM and accessed by using an entitled key during installation.

1. Go to [My IBM dashboard](#) > [Container Software & Entitlement Keys](#) > [View Library](#) > [Entitlement keys](#).
2. Register the customer number of your software license.

After this registration, the entitlement key generated via My IBM can be used to authorize access to the images. For more information, see [IBM Entitled Systems Support for more information](#).

IBM Storage Scale CSI

IBM Storage Scale CSI is available in multiple image registries. But it is *not available via IBM Support Fix Central*. For full installation guidance, read the [IBM Storage Scale CSI documentation](#). IBM Storage Scale CSI pulls the prerequisite images from multiple image registries. Before installation, make sure to understand the dependencies explained in [Deployment considerations](#).

Q18.10

What Cloud Paks and services does IBM Storage Scale currently support?

A18.10

IBM Storage Scale Container Native is now supported with/for the following six IBM Cloud Paks:

- Cloud Pak for Data (CP4D): <https://www.ibm.com/docs/en/software-hub/5.2.x?topic=planning-storage-considerations>
- Cloud Pak for Security (CP4S): <https://www.ibm.com/docs/en/cloud-paks/cp-security/1.11.0?topic=planning-storage-requirements>
- Cloud Pak for Network Automation (CP4NA): <https://www.ibm.com/docs/en/cloud-paks/cp-network-auto/2.7.x?topic=planning-storage-requirements>
- Cloud Pak for Business Automation (CP4BA): <https://www.ibm.com/docs/en/cloud-paks/cp-biz-automation/25.0.0?topic=considerations-storage-options>
- Cloud Pak for Integration (CP4I): <https://www.ibm.com/docs/en/cloud-paks/cp-integration/16.1.2?topic=requirements-storage-considerations>
- Cloud Pak for Watson AIOps (CP4WAIOPs): <https://www.ibm.com/docs/en/cloud-paks/cloud-pak-aiops/4.11.1?topic=planning-storage-requirements#scale>

For IBM Storage Fusion Cloud Pak support, see: <https://www.ibm.com/docs/en/fusion-software>

For IBM Storage Fusion HCI Cloud Pak support, see <https://www.ibm.com/docs/en/fusion-hci-systems>

Q18.11:

Where can I get more information about the subjects discussed in this section?

A18.11:

IBM Storage Scale Container Storage Interface (CSI) documentation:

<https://www.ibm.com/docs/en/spectrum-scale-csi>

IBM Storage Scale Container Native documentation:

<https://www.ibm.com/docs/en/scalecontainernative>

Q18.12:

Does IBM Storage Scale Container Native support RDMA?

A18.12

Yes, Remote Direct Memory Access (RDMA) is supported with both InfiniBand and RDMA over Converged Ethernet (RoCE). For more information, see [NVIDIA Networking](#) in the IBM Storage Scale Container Native documentation.

Q18.13

Can IBM Storage Scale storage be used for overlay filesystems?

A18.13

Using IBM Storage Scale as a backend for overlay filesystem is not supported. Overlay/Union filesystem is used by container runtime to build and run container images. These container images are usually stored under `/var/lib/kubelet` or `/var/lib/docker` directories. Using IBM Storage Scale as a storage for such directories can result in issues during container operations.

IBM Storage Scale Erasure Code Edition questions

Q19.1:

What is IBM Storage Scale Erasure Code Edition, and why should I consider it?

A19.1:

IBM Storage Scale Erasure Code Edition provides all the functionality, reliability, scalability, and performance of IBM Storage Scale on the customer's own choice of commodity hardware with the added benefit of network-dispersed IBM Storage Scale RAID, and all of its features providing data protection, storage efficiency, and the ability to manage storage in hyperscale environments.

Q19.2:

What are the limitations of IBM Storage Scale Erasure Code Edition?

A19.2:

In addition to the items that are listed in the *IBM Storage Scale Erasure Code Edition limitations* section of the IBM Documentation, also see the following sections of the IBM Documentation:

- *IBM Storage Scale Erasure Code Edition Hardware requirements*
- *IBM Storage Scale Erasure Code Edition installation prerequisites*

Q19.3:

Can IBM Storage Scale Erasure Code Edition exist with IBM Elastic Storage Server or IBM Elastic Storage System in the same cluster and support the same file system?

A19.3:

Yes, with the following limitations:

- Adding IBM Storage Scale Erasure Code Edition into IBM Elastic Storage Server (ESS) cluster:
 - IBM Storage Scale Erasure Code Edition must be at version 5.0.3.1 or later, and IBM Elastic Storage Server must be at version 5.3.4 or later. For information, see *Incorporating IBM Storage Scale Erasure Code Edition in an IBM Documentation Elastic Storage Server (ESS) cluster* section in IBM.
- Adding IBM Elastic Storage Server (ESS) building block into IBM Storage Scale Erasure Code Edition cluster:
 - IBM Storage Scale Erasure Code Edition must be at version 5.1.5.1 or later, and IBM Elastic Storage Server must be at version 6.1.5 or later. For information, see *Incorporating IBM Elastic Storage Server (ESS) building block in an IBM Documentation IBM Storage Scale Erasure Code Edition cluster* section in IBM.

Q19.4.1:

What are the minimum hardware requirements for IBM Storage Scale Erasure Code Edition?

A19.4.1:

For more information, see *Minimum hardware requirements* in the IBM Storage Scale Erasure Code Edition Documentation.

Q19.4.2:

Can I use any vendor's server with IBM Storage Scale Erasure Code Edition?

A19.4.2:

Yes, if the hardware meets the minimum hardware requirements. You can verify this using the hardware precheck tool. For more information about this tool, see [Q19.7 How can I get the IBM Storage Scale Erasure Code Edition hardware and IBM Storage Scale network precheck tools, and how do I execute them?](#)

Q19.4.3:

Can I use x86, ARM 64, IBM z, and PowerPC architectures with IBM Storage Scale Erasure Code Edition?

A19.4.3:

At this time, only servers with x86 or aarch64 (ARM 64) architectures are supported.

Q19.4.4:

What operating systems are supported for IBM Storage Scale Erasure Code Edition storage servers?

A19.4.4:

Currently, only RHEL is supported. For information on the IBM Storage Scale Erasure Code Edition releases with supported operating systems, see [Q2.1 What is supported on IBM Storage Scale for AIX, Linux, Power, and Windows?](#)

Q19.4.5:

Is CentOS supported with IBM Storage Scale Erasure Code Edition?

A19.4.5:

Only RHEL is supported at this time. For more information about the use of unsupported distributions with IBM Storage Scale, see [Q2.3 What is the IBM Storage Scale support position regarding clone Linux distributions \(CentOS, ROCK, White box Linux, etc.\)?](#)

Q19.4.6:

Can I use SATA drives with IBM Storage Scale Erasure Code Edition?

A19.4.6:

No, only SAS, NL-SAS, and NVMe drives are supported at this time.

Q19.4.7:

Can I use SED drives with IBM Storage Scale Erasure Code Edition?

A19.4.7:

Self-encrypting drives are only allowed if they have never been enrolled into SED (locked) and do not require a key to unlock after power on. Starting with release 5.1.9, IBM Storage Scale Erasure Code Edition support migrates an existing recovery group to enable the SED. To learn about limitations and the procedure, see [Self-encrypting drive support](#) in the IBM Storage Scale Erasure Code Edition documentation.

Q19.4.8:

Can I use external enclosures with IBM Storage Scale Erasure Code Edition?

A19.4.8:

No, a typical IBM Storage Scale Erasure Code Edition configuration uses direct attached storage devices. An RPQ would be required for IBM to review any requests for external enclosure support. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.

Q19.4.9:

Can I run IBM Storage Scale Erasure Code Edition with heterogeneous servers?

A19.4.9:

No, all servers in an IBM Storage Scale Erasure Code Edition recovery group must have the same CPU, memory, and storage configuration with consistent adapter hardware and firmware levels. If you plan to introduce a new server or a new storage topology into your cluster, it must be done with servers in a separate recovery group.

Q19.4.10:

Can I use a virtual machine as a storage node with IBM Storage Scale Erasure Code Edition?

A19.4.10:

Constrained VMWare virtual machine is supported and for more information see the *Deploying IBM Storage Scale Erasure Code Edition on VMware infrastructure* topic in the IBM Storage Scale Erasure Code Edition Documentation.

Any other use of a virtual machine as a storage node in a production environment must be reviewed by IBM. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process. For more information, see the *Hardware checklist* topic in the IBM Storage Scale Erasure Code Edition Documentation and [Q7.2 Is IBM Storage Scale on Linux \(x86 and Power\) supported in a virtualization environment?](#)

Q19.4.11:

Can I deploy IBM Storage Scale Erasure Code Edition in cloud environments?

A19.4.11:

IBM Storage Scale Erasure Code Edition is supported on IBM Cloud and Oracle Cloud infrastructure.

To deploy IBM Storage Scale Erasure Code Edition in IBM Cloud Infrastructure, an RPQ is required. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.

To deploy IBM Storage Scale Erasure Code Edition in Oracle Cloud Infrastructure, please contact OCI HPC infrastructure team for the hardware and provisioning.

Q19.5.1:

**What is the minimum and maximum number of nodes required for an IBM Storage Scale Erasure Code Edition recovery group?
What is the maximum number of IBM Storage Scale Erasure Code Edition storage nodes in an IBM Storage Scale cluster?**

A19.5.1:

Each IBM Storage Scale Erasure Code Edition recovery group can have 3 - 32 storage nodes from 5.1.4 release. Releases earlier than 5.1.4 can have 4 - 32 nodes. Starting with the 5.2.0 release, there can be up to 256 storage nodes in an IBM Storage Scale cluster that is using IBM Storage Scale Erasure Code Edition. Releases previous to 5.2.0 can have up to 128 storage nodes. For more information, see *Planning for erasure code selection* in the IBM Storage Scale Erasure Code Edition Documentation.

Q19.5.2:

How can I estimate the usable space in one recovery group with IBM Storage Scale Erasure Code Edition storage nodes in an IBM Storage Scale cluster?

A19.5.2:

You can calculate usage capacity using the capacity estimator tool. You can download the tool from <https://github.com/IBM/SpectrumScaleTools> under [ece_capacity_estimator](#) directory.

Q19.6.1:

What are the network requirements and considerations for IBM Storage Scale Erasure Code Edition?

A19.6.1:

The network-dispersed IBM Storage Scale erasure coding in IBM Storage Scale Erasure Code Edition makes heavy use of network resources. For this reason, it is critical that every IBM Storage Scale Erasure Code Edition installation have a fast and low latency network for best results. For this reason, a minimum of one 25 Gbps interface for file system and IBM Storage Scale Erasure Code Edition data traffic is required. If CES IPs are used, they should be defined on a separate network. For more information, see *Network requirements and precheck* in the IBM Storage Scale Erasure Code Edition Documentation.

Q19.6.2:

Is IPV6 supported with IBM Storage Scale Erasure Code Edition?

A19.6.2:

An IBM Storage Scale Erasure Code Edition cluster can be configured to use IPV6 for file system and data traffic. There are some IBM Storage Scale services that do not support IPV6. For more information, see [Q6.11 What are configuration considerations when using IPV6?](#)

Q19.7:

How can I get the IBM Storage Scale Erasure Code Edition hardware, IBM Storage Scale network, and IBM Storage Scale storage precheck tools, and how do I execute them?

A19.7:

The precheck tools can be downloaded from the following link:

<https://github.com/IBM/SpectrumScaleTools>

- Hardware precheck: under [ece_os_readiness](#) directory
- Check that a collection of nodes meets the IBM Storage Scale Erasure Code Edition building block requirements: under [ece_os_overview](#) directory
- Network precheck: under [ece_network_readiness](#) directory
- Storage precheck: under [ece_storage_readiness](#) directory

The installation requirements and execution instructions are documented in the README files for each repository. These can be run at any time to check compliance with IBM Storage Scale Key Performance Indicators. The network and storage precheck tools should not be run when other critical workloads are running to avoid a heavy load.

Q19.8:

Can I use the IBM Storage Scale installation toolkit to install and upgrade IBM Storage Scale Erasure Code Edition?

A19.8:

Yes, this is the recommended method to install and upgrade IBM Storage Scale Erasure Code Edition. For more information, see *Installing IBM Storage Scale Erasure Code Edition* and *Upgrading IBM Storage Scale Erasure Code Edition* in the IBM Storage Scale Erasure Code Edition Documentation. To upgrade from version 5.0.4.3 or later to version 5.0.5 or later, you can use the installation toolkit for an online upgrade. To upgrade from a version earlier than 5.0.4.3 to a later version (including version 5.0.4.3 to version 5.0.4.4), you can use the installation toolkit for offline upgrades only, or you can use the manual upgrade process.

Q19.9:

Can I use sudo wrappers with IBM Storage Scale Erasure Code Edition?

A19.9:

Yes, IBM Storage Scale Erasure Code Edition can be configured in a cluster with sudo wrappers enabled. At this time, the installation toolkit does not support sudo wrappers. If you use the installation toolkit, you must first install and configure your cluster with the root login enabled and then change the configuration to use sudo wrappers.

Q19.10.1:

Can I run CES protocol software on IBM Storage Scale Erasure Code Edition storage nodes?

A19.10.1:

In this release, CES protocol software should be configured on separate nodes for protocol workloads with high performance requirements. An RPQ is required if you want to run CES protocol software on IBM Storage Scale Erasure Code Edition storage nodes. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.

Q19.10.2:

Can I run application workloads on IBM Storage Scale Erasure Code Edition storage nodes?

A19.10.2:

Yes, application workloads must be deployed in an environment where network, CPU, and memory utilization can be constrained (for example, with Linux cgroups or containers). The IBM Storage Scale Erasure Code Edition storage servers must be sized with enough resources to support IBM Storage Scale Erasure Code Edition requirements and the added requirements of the application workload.

Q19.10.3:

Can I run the IBM Storage Scale GUI, AFM gateways, and perfmon collectors on IBM Storage Scale Erasure Code Edition storage nodes?

A19.10.3:

These services should be configured to run on separate nodes. For more information, see *Planning for node roles* in the IBM Storage Scale Erasure Code Edition Documentation.

Q19.11:

What workloads is IBM Storage Scale Erasure Code Edition storage recommended for? Can it be used for high performance workloads?

A19.11:
IBM Storage Scale Erasure Code Edition is expected to work best for workloads that require high bandwidth and low latency. This includes but is not limited to data analytic, AI, and other unstructured data processing workloads. Any workload that is planned for IBM Storage Scale Erasure Code Edition should be thoroughly tested prior to deploying in a production environment.

Q19.12:
How can I migrate data from an existing IBM Storage Scale cluster or file system to IBM Storage Scale Erasure Code Edition?

A19.12:
There are several strategies that can be used for data migration between storage pools in an existing cluster or between clusters. There is no support for in place migration of data on an existing cluster to IBM Storage Scale Erasure Code Edition storage using existing hardware. Contact IBM to discuss what will work best for your specific requirements.

Q19.13:
When installing IBM Storage Scale Erasure Code Edition, IBM udev rules are installed. Also, when I configure servers using `mmvdisk`, the initial IBM Storage Scale configuration values are set. Is it okay to change these rules and values?

A19.13:
Yes, both udev rules and IBM Storage Scale configuration values are meant to be a good starting point for typical hardware and typical workloads, but you might need to adjust both of these for your configuration and workload. In particular, pagepool might need to be adjusted for optimal performance.

Q19.14:
Are there any extra requirements for using or configuring NVMe drives for use with IBM Storage Scale Erasure Code Edition?

A19.14:
Enterprise class NVMe drives with U.2 form factor are required. When deploying IBM Storage Scale Erasure Code Edition, you must define the mapping of NVMe drive location to PCI bus. For more information, see *Setting up IBM Storage Scale Erasure Code Edition for disk slot location* in the IBM Storage Scale Erasure Code Edition Documentaion.

NVMe drives that are used by IBM Storage Scale Erasure Code Edition must be formatted with a metadata size of zero and the protection information disabled. All NVMe drives in the same declustered array should be formatted with the same LBA size. For more information, see the *Hardware checklist* in the IBM Storage Scale Erasure Code Edition Documentation.

Q19.15:
Can I deploy IBM Storage Scale Erasure Code Edition on a configuration that does not meet the requirements and recommendations detailed in the IBM Storage Scale Erasure Code Edition documentation?

A19.15:
Configurations that do not meet our minimum requirements must be reviewed by IBM using the RPQ (SCORE) process. A proof of concept or additional testing might be required as part of this process. Ask your sales representative to contact IBM Storage Scale development about the RPQ process.

Q19.16:
What high-speed Ethernet networking adapters are tested with IBM Storage Scale Erasure Code Edition?

A19.16:
IBM Storage Scale Erasure Code Edition is tested with NVIDIA Mellanox (ConnectX-5 and later) and Broadcom (400 GbE and 200 GbE) adapters. Both Ethernet adapters are supported with RDMA over Converged Ethernet (RoCE) as well.

The network-dispersed erasure coding in IBM Storage Scale Erasure Code Edition makes heavy use of network resources. Therefore, always run the Network Precheck located in the [ece_network_readiness](#) directory to ensure your configuration meets the network requirements. Configurations that do not meet the minimum requirements must be reviewed by IBM through the RPQ (SCORE) process.

IBM Storage Scale Developer Edition questions

Q20.1:
What functions do I get with IBM Storage Scale Developer Edition?

A20.1:
This edition provides all of the features of the IBM Storage Scale Data Management Edition, but it is limited to 12 TB per cluster.

Q20.2:
Can I use IBM Storage Scale Developer Edition in a production environment?

A20.2:
No, use of IBM Storage Scale Developer Edition in a production environment is prohibited.

Q20.3:
How do I get IBM product support for IBM Storage Scale Developer Edition?

A20.3:
There is no support from IBM for IBM Storage Scale Developer Edition.

Q20.4:

For IBM Storage Scale Developer Edition, can I upgrade my cluster to an IBM-supported IBM Storage Scale product offering like IBM Storage Scale Data Management Edition?

A20.4:

No, an upgrade to any other IBM Storage Scale offering is not supported.

Q20.5:

With IBM Storage Scale Developer Edition, how can I see the licensed storage usage?

A20.5:

Run the `mmfsllicense` command with the `--licensed-usage` option.

Q20.6:

Is the IBM Storage Scale installation toolkit supported with IBM Storage Scale Developer Edition?

A20.6:

Yes, it is supported.

Q20.7:

What are the supported operating systems and architectures for IBM Storage Scale Developer Edition?

A20.7:

The IBM Storage Scale Developer Edition supports RHEL at levels identical to IBM Storage Scale Data Management Edition. The only supported architecture is x86.

Q20.8:

Can I run more than one IBM Storage Scale Developer Edition cluster in the same company, division, test lab, etc.?

A20.8:

Yes, you can if each cluster is limited to 12 TB. You can also cross mount clusters if each cluster is limited to 12 TB.

Q20.9:

Where can I download IBM Storage Scale Developer Edition?

A20.9:

IBM Storage Scale Developer Edition is available for download at the following link: <https://www.ibm.com/products/spectrum-scale/pricing>.

Integrated protocol server authentication questions

Q21.1

Is IBM Storage Scale affected by the Microsoft advisory ADV190023 regarding LDAP channel binding and LDAP signing?

A21.1:

No, IBM Storage Scale is not affected by the Microsoft advisory ADV190023 regarding LDAP channel binding and LDAP signing. For more information about that advisory, see <https://portal.msrc.microsoft.com/en-us/security-guidance/advisory/adv190023>.

The advisory recommends activating the LDAP channel binding and LDAP sealing on the Active Directory Domain Controllers. The LDAP channel binding setting is related to LDAP authentication over SSL/TLS. The LDAP signing setting is related to simple LDAP binds or SASL (Simple Authentication and Security Layer) LDAP binds over an encrypted or unencrypted channel.

The IBM Storage Scale CES stack supports integrating with Active Directory Domain Controllers as one of the supported authentication mechanisms for FILE protocols. In such a configuration, the FILE protocol stack communicates over LDAP with Active Directory Domain Controllers. It binds with the domain controllers over SASL (Simple Authentication and Security Layer) using Kerberos authentication. The Samba configuration setting `client ldap sasl wrapping` defines whether these SASL binds are signed or signed and sealed. The value for this setting is by default `sign`. Thus, the FILE protocol stack works seamlessly after the setting that is recommended in the advisory has been applied.

Q21.2:

Which Microsoft Active Directory versions does IBM Storage Scale support?

A21.2:

IBM Storage Scale supports all versions of the Active Directory servers that Microsoft officially supports. Tests were run on Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, Windows Server 2022, and Windows Server 2025.

Q21.3:

Which LDAP versions does IBM Storage Scale support?

A21.3:

LDAP server hosting RFC2307 schema compliant user and group entries are supported for integration with IBM Storage Scale. Only such users and groups are recognized when accessing IBM Storage Scale over NFS and SMB. Performing SMB access requires additional attributes that are defined on user and group entries, which are available through the Samba schema. For SMB access, the Samba schema must be imported in the LDAP server and the user and group entries should be updated for the relevant Samba attributes.

Q21.4:

What is the impact on IBM Storage Scale Protocol when you are using AD with RFC2307 or migrating to Windows 2016 AD server or later?

A21.4

When you are using AD with RFC2307 authentication scheme, IBM Storage Scale requires certain attributes of the user and group identities (for example, uidNumber for user, gidNumber for primary group and secondary groups of the user) on Active Directory server to be populated based on RFC2307 schema. From Windows 2016 AD server, Microsoft is removing the Identity Management for Unix and the plugin for management of the RFC2307 attributes. The attributes are going to stay, only the ability to manage the attributes using the IDMU plugin has been removed. There are multiple ways to manage the attributes. For information on managing the attributes, see the following link from Microsoft:

[Clarification regarding the status of Identity Management for Unix \(IDMU\) & NIS Server Role in Windows Server 2016 Technical Preview and beyond](#)

Note: There is no impact on IBM Storage Scale.

Cloud and cloudkit questions

Q22.1

Can IBM Storage Scale be installed on the cloud?

A.22.1:

Yes. IBM Storage Scale deployments are supported on a variety of clouds, directly through IBM, but also through Business Partners who have crafted specific solutions that involve IBM Storage Scale.

Cloud resources can be provisioned and IBM Storage Scale can be installed upon them by using the included cloudkit. Alternatively, IBM Storage Scale can be installed upon manually provisioned cloud resources.

Q22.2

What is the recommended way to install IBM Storage Scale on cloud?

A22.2:

Use of the IBM Storage Scale cloudkit for deployments is highly recommended because each option and configuration parameter has been vetted against best practices; which results in a known and workable cloud configuration.

Q22.3

What are the public clouds sported by cloudkit?

A22.3:

- Amazon Web Services (AWS)
- Google Cloud Platform (GCP)
- Microsoft Azure

Q22.4

What IBM Storage Scale edition packages are supported by the cloudkit?

A22.4:

The following IBM Storage Scale edition packages are supported by the cloudkit:

- IBM Storage Scale Data Management
- IBM Storage Scale Data Access
- IBM Storage Scale Developer

Q22.5

If I manually deploy IBM Storage Scale on the cloud without using cloudkit, are there any restrictions, functions, or environments that are not recommended for use?

A22.5:

IBM Storage Scale consists of many features, some of which are highly useful and relevant for cloud environments, while other features are not recommended in a cloud environment. A few examples:

- SMB, NFS, Object protocols using Cluster Export Services (CES stack). The CES stack relies upon existence of an assigned IP pool (non-DHCP) that can freely rotate across nodes. The IBM Storage Scale CES stack controls assignment and rotation in cases of node failure. In a cloud environment, IPs can only be assigned and moved by using cloud-specific APIs; which means that the IBM Storage Scale CES stack is currently incompatible with cloud APIs without further development.
- AFM home (using NFS protocol). See the previous CES stack explanation, because of the CES NFS protocol is used for AFM home.
- Cloud block disk services are a feature of every cloud and provide block access to storage. IBM Storage Scale uses disk services to create NSDs for building the IBM Storage Scale file system. For compatibility to be ensured, these disk services must be fully vetted with IBM Storage Scale. An inclusion list of services specifically compatible with IBM Storage Scale is shown in the supported features matrix for each cloud.
- Custom Kubernetes services are a feature of every cloud. IBM Storage Scale supports manually installed Vanilla Kubernetes with its IBM Storage Scale Container Storage Interface driver (CSI driver). Custom services such as EKS and GKS are not supported by the CSI driver.
- OpenShift is available in every cloud as either a self-managed instantiation or a bundled service. Refer to the IBM Storage Scale Container Native (CNSA) documentation for a list of supported clouds for OpenShift.
- Virtualization technologies on the cloud: VMware or KVM. The IBM Storage Scale FAQ contains a compatibility matrix of features and functions supported in VMware or KVM. This same matrix applies to VMware or KVM on the cloud.
- For other exceptions, as well as lists of supported features, see the supported features matrix of IBM Storage Scale for each cloud.

Q22.6

What are the permissions needed to deploy cloudkit?

A22.6:

For information about the permissions needed to deploy cloudkit, see the validation included in [Understanding the cloudkit installation options](#) in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q22.7

What IBM Storage Scale features are supported by cloudkit?

A22.7:

For information about the IBM Storage Scale features supported by cloudkit, see the *IBM Storage Scale features* table in [Supported features of cloudkit](#) of the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q22.8

What operating systems are supported by cloudkit?

A22.8:

- The cloudkit binary is only supported on Red Hat Enterprise Linux (RHEL) releases. Therefore, the cloudkit can only be executed from a machine running the mentioned RHEL versions. For more information, see the *Operating systems that are supported by cloudkit installer nodes* table in [Supported features of cloudkit](#) of the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- The cloudkit can create IBM Storage Scale clusters that run on the RHEL and kernel version that are supported by IBM Storage Scale (for more information, see the [support information tool](#)). However, by default, the cloudkit supports RHEL 8.10 and 9.4 on the AWS and GCP clouds.

Q22.9

How to access the cluster that was created by cloudkit?

A22.9:

For information about accessing the cluster that was created by cloudkit, see [Accessing the cluster on the cloud](#) of the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Q22.10

What are the limitation when using cloudkit?

A22.10:

For information about current limitations of the cloudkit, see [Limitations of cloudkit](#).

Q22.11

What should one be aware of before deleting an IBM Storage Scale cluster?

A22.11:

- Before deleting the IBM Storage Scale cluster on the cloud, ensure that any data that is required is backed up. All data stored in the IBM Storage Scale file systems will be permanently removed during deletion.
- Deleting a cluster may lead to failure if there are other resources that reference the resources created by cloudkit; or if there are additional resources that are not recognized by cloudkit but share the network resources created by cloudkit.

Q22.12

Who should I contact for support?

A22.12:

- For any issues related to the cloud infrastructure, including any cloud resources that are used by the IBM Storage Scale cluster, contact the support team of the hosting cloud.
- For any issues related to the IBM Storage Scale cluster, contact the IBM Storage Scale support team.

Q22.13

How to collect debug data for cloudkit?

A22.13

Collect the following information for debugging a cloudkit issue:

- `cloudkit log`. You can find the cloudkit log location displayed during every run of the `cloudkit` command.

Example:

```
# ./cloudkit create cls
I: Logging at /root/scale-cloudkit/logs/cloudkit-10-10-
2023_6-9-22.log
```

The default log location is: `${Home}/scale-cloudkit`.

- `gpfs.snap`

Q22.14

Why the file system capacity is relatively smaller or bigger than the exact size provided as input?

A22.14

During deployment, the NSD size is calculated internally within the cloudkit based on the profile. This calculation rounds the size, which could lead to increase or decrease in the overall capacity. For an instance with store-based profile, the entire disk size shown by the cloud vendor is not fully utilizable for the file system.

In few cases, the block to subblock calculations, the inode configurations, and the file system capacity could be shown as relatively smaller than the accumulated attached disk sizes.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10594-1785
USA

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This 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 may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Intellectual Property Law
2455 South Road,P386
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

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

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom

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.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

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

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

April 2026

© Copyright International Business Machines Corporation 2004, 2025.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.