

GHI Installation Guide

Jae R. Kerr

[<jrkerr@us.ibm.com>](mailto:jrkerr@us.ibm.com)

Table of Contents

- [1. Preface](#)
- [2. Conventions Used in This Document](#)
- [3. Preparing for GHI Install](#)
 - [3.1. Prerequisites](#)
 - [3.2. Operating System](#)
 - [3.3. ISHTAR](#)
- [4. Spectrum Scale](#)
 - [4.1. Install Spectrum Scale](#)
 - [4.2. Configure Spectrum Scale](#)
- [5. DB2](#)
 - [5.1. Users and Groups](#)
 - [5.2. ACL Account Validation](#)
 - [5.3. Setup GHI Tablespace](#)
 - [5.4. Install DB2 Client](#)
 - [5.5. Create DB2 Permanent License](#)
- [6. HPSS](#)
 - [6.1. HPSS Users and Groups](#)
 - [6.2. Install HPSS](#)
 - [6.3. Configure HPSS Client](#)
- [7. GHI Installation & Configuration](#)
 - [7.1. Install GHI](#)
 - [7.2. GHI Users and Groups](#)
 - [7.3. Configure GHI](#)
 - [7.4. Create GHI Cluster](#)
 - [7.5. Create Spectrum Scale file systems GHI will manage](#)
 - [7.6. Create IOMs for each GHI Managed File systems](#)
 - [7.7. Modify xinetd.conf for number of IOMs](#)
 - [7.8. \(Optional\) Configure Syslog Facilities](#)
 - [7.9. Information Lifecycle Management \(ILM\) Policies](#)
- [8. Backup & Restore](#)
 - [8.1. Backups](#)
 - [8.2. Restore](#)
- [9. Appendix A - GHI Conversions](#)
 - [9.1. Conversion steps for 2.5.0.1 to 3.0](#)

Spectrum Scale HPSS Interface Release 2.5

September 2016

(c) 1992-2016 International Business Machines Corporation, the Regents of the University of California, Los Alamos National Security, LLC, Sandia Corporation, and UT-Battelle.

All rights reserved.

Portions of this work were produced by Lawrence Livermore National Security, LLC, and Lawrence Livermore National Laboratory (LLNL) under Contract No. DE-AC52-07NA27344 with the U.S. Department of Energy (DOE); by the University of California, Lawrence Berkeley National Laboratory (LBNL) under Contract No. DE-AC02-05CH11231 with DOE; by Los Alamos National Security, LLC, Los Alamos National Laboratory (LANL) under Contract No. DE-AC52-06NA25396 with DOE; by Sandia Corporation, Sandia National Laboratories (SNL) under Contract No. DE-AC04-94AL85000 with DOE; and by UT-Battelle, Oak Ridge National Laboratory (ORNL) under Contract No. DE-AC05-00OR22725 with DOE. The U.S. Government has certain reserved rights under its prime contracts with the Laboratories.

DISCLAIMER

Portions of this software were sponsored by an agency of the United States Government. Neither the United States, DOE, The Regents of the University of California, Los Alamos National Security, LLC, Lawrence Livermore National Security, LLC, Sandia Corporation, UT-Battelle, nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Printed in the United States of America.

GHI 2.5 September 2016

High Performance Storage System is a trademark of International Business Machines Corporation. IBM is a registered trademark of International Business Machines Corporation. IBM, DB2, DB2 Universal Database, AIX, pSeries, and xSeries are trademarks or registered trademarks of International Business Machines Corporation. AIX is a trademark of International Business Machines Corporation. UNIX is a registered trademark of the Open Group. Linux is a registered trademark of Linus Torvalds in the United States and other countries. Kerberos is a trademark of the Massachusetts Institute of Technology. Microsoft Windows is a registered trademark of Microsoft Corporation. NFS is trademark of Sun Microsystems, Inc. Other brands and product names

appearing herein may be trademarks or registered trademarks of third parties.

1. Preface

The Spectrum Scale (formerly GPFS) to HPSS Interface (GHI) Installation Guide is for use at system installation time.

It outlines the steps required to install and configure a GHI system.

2. Conventions Used in This Document

Example commands that should be typed at a command line will be preceded by a percent sign (%) and be presented in a **bold** face font:
% sample command

Example command output and example contents of ASCII files are presented in *italic* font:

sample file line 1

sample file line 2

Any text preceded by a pound sign ('#') is considered comment lines:
This is a comment

3. Preparing for GHI Install

Prior to installing GHI, a series of planning meetings should have been conducted to discuss the customer's requirements in order to properly size and configure the GHI system. The planning process is not a subject of this Installation Guide and is done during the Proposal and/or System Engineering phase of the project. In addition, refer to the GHI Management Guide for planning considerations.

IBM recommends you install GHI on a Spectrum Scale cluster that does not have another Hierarchical Storage Management (HSM) application running, for example Tivoli Storage Manager (TSM). If another HSM managed file system is required, it must run on a separate cluster and remotely mounted on the GHI managed cluster. GHI is dependent on timely Data Management Application Programming Interface (DMAPI) events from Spectrum Scale, therefore there should not be two applications competing for events.

For systems installed with High Availability (HA) Core Server, it is critical

to ensure that the required GHI components are installed on the backup or stand-by Core Server. These components include DB2 accounts creation and configuration, DB2 Server configuration and Independent Standalone HPSS TAR (ISHTAR).

GHI installation requires root or root-equivalent privileges, except where noted otherwise. Privileges are normally acquired by logging into the system as root or another privileged user, or by issuing the command “*su -*”.

3.1. Prerequisites

Before installing GHI, review the GHI Release Notes on the HPSS Admin wiki for prerequisites, special notes, and possible known issues for the version you plan to install. The Release Notes define the software version of each prerequisite software:

- Operating system
 - Setting up ulimits
 - Enabled core dump
 - rsyslog configuring
 - Install memcached
 - Configure memcached
- Python - not covered in this document
- IBM_db egg (python support for DB2) - not covered in this document
- ISHTAR
- Spectrum Scale
- DB2
- HPSS

3.2. Operating System

3.2.1. Set ulimits



In GHI 2.5, the default inode scan bucket size is 1000. Increase the max open file descriptors limit to 65536 in /etc/security/limits.conf on all systems that will run GHI. Reboot each node to validate each change is correct and

persists after a reboot.

Example:

```
% vi /etc/security/limits.conf
#*                soft    core    0
#*                hard    rss     10000
#@student        hard    nproc   20
#@faculty        soft    nproc   20
#@faculty        hard    nproc   50
#ftp             hard    nproc   0
#@student        -      maxlogins 4
*                -      nofile  65536          <- (add to
the end of block.)
```

3.2.2. Enable core dump at the operating system level

Verify that operating system core dump functionality is configured correctly in the case that GHI needs to create a core file for debug purposes.

3.2.3. rsyslog configuration



Recommendations for setting up rsyslog for tuning rate limiting off and suppressing repeat messages.

Information on log rate limiting: :website: <http://www.rsyslog.com/tag/rate-limiting/>

Information on repeat messages: :website: http://www.rsyslog.com/doc/rsconf1_repeatedmsgreduction.html

3.2.4. Install memcached and libmemcached

GHI uses memcached to reduce the load on the HPSS metadata. It is not required for GHI but it improves performance of GHI full backups, GHI image backup verification, **ghiverifyfs** and **ghi_ls**. Install the memcached and libmemcached-devel RPMs from the RHEL software distribution on each machine you want memcached to run to improve data transfers. You must install the RPMs even though you may not use memcache.

3.2.5. Configure memcached

1. Start memcached daemon on each node.

```
% memcached -d -u root -m 1024
```

This will start the server with 1024MB of memory. Adjust to the desired cache size. The rule of thumb is 10,000 items in the cache will use ~16MB.

2. Create /var/hpss/etc/memcached.conf and copy it to each GHI node.

The file format is:

```
--POOL-MIN=4  
--POOL-MAX=32  
--SERVER=<ghi node # name>:11211  
--SERVER=<ghi node # name>:11211
```

3. Automatically start memcached.

If using RHEL 7, follow the steps below.

- a. Create /usr/lib/systemd/system/memcached.service
- b. Add the lines below to the file.

```
[Unit]  
Description=Memcached  
Before=httpd.service  
After=network.target  
  
[Service]  
Type=simple  
EnvironmentFile=-/etc/sysconfig/memcached  
Restart=always  
ExecStart=/usr/bin/memcached -u $USER -p $PORT -m  
$CACHESIZE -c $MAXCONN $OPTIONS  
  
[Install]  
WantedBy=multi-user.target
```

- c. cat /etc/sysconfig/memcached

```
PORT="11211"  
USER="memcached"  
MAXCONN="1024"  
CACHESIZE="64"  
OPTIONS=""
```

- d. Run commands to start.

```
% systemctl enable memcached.service  
% systemctl daemon-reload  
% systemctl start memcached.service
```

If using RHEL 6, follow the steps below.

- a. Create `/etc/init/memcached.conf`.
- b. Add the lines below to the file.

```
stop on runlevel [012]  
start on starting rc RUNLEVEL=[345]  
respawn  
exec /usr/bin/memcached -d -u root -m 1024
```

- c. Run command to start.

```
% memcached -d -u root -m 1024
```

3.3. ISHTAR

Independent Standalone HPSS TAR (ISHTAR) uses the same authentication mechanism as those used by GHI.

3.3.1. Install ISHTAR

ISHTAR must be installed on all GHI IOM nodes. The ISHTAR RPM is compatible with GHI and is provided by IBM. Only the provided version should be used with GHI. Check the GHI Release Notes for the version of ISHTAR to install.

3.3.2. Configure ISHTAR

1. Copy the `htar.ksh` wrapper script to `"/var/hpss/hsi/bin"`.

```
% cd /var/hpss/hsi/bin  
% cp htar.ksh.template htar.ksh  
% edit htar.ksh - Variables to modify are described in step 2  
below  
% /bin/chmod 755 htar.ksh
```

2. Modify the `htar.ksh` script to provide correct values for the following information:

```
TMPDIR:                Location of the temporary files. The  
amount of space required is based on the size of an  
aggregate, plus temporary files created for the data files.  
It is recommended to set this  
variable to "<Spectrum Scale_mount_point>/scratch/.ghi/".  
  
DEFAULT_REALM:         Realm name for location of the HPSS  
Core server. This name must match what is set for "site  
name" in /var/hpss/etc/site.conf from CORE server.
```

HPSS_AUTH_METHOD: Set this variable for desired authentication type. < Unix or Kerberos > This variable will determine your keytab file to be used.

HPSS_KEYTAB_PATH: Location of keytab. Set this variable when using UNIX authentication (ex. /var/hpss/etc/hpssdmg.unix.keytab).

HPSS_HOSTNAME: Interface to be used for the data path.

HPSS_NET_FAMILY: Must be set to the same value as the /var/hpss/etc/env.conf value.

* Verify that ISHTAR verification and CRCs are enabled in the htar.ksh file. Do not disable this feature.

* Copy all files in /var/hpss/hsi/bin/* from the session node to all GHI I/O Manager nodes.

4. Spectrum Scale

4.1. Install Spectrum Scale

Contact your IBM Spectrum Scale customer support representative to obtain the Spectrum Scale software. Install Spectrum Scale per instructions from the IBM Spectrum Scale support representative.

4.2. Configure Spectrum Scale

1. After Spectrum Scale is installed, make sure ssh or rsh is working between nodes in the cluster. If using ssh, be certain to complete additional configuration steps to allow for password-less command execution (steps are covered in the Spectrum Scale documentation).

2. Enable threshold processing.

% mmlsconfig - to check to see if it is set.

% mmchconfig enablelowspaceevents=yes

3. Configure NSD Multipath.

If using multipath, please follow the steps below to create NSDs.

- a. Create a /etc/multipath/bindings file. The file needs to match on

all nodes using the NSD.

b. **% cp /usr/lpp/mmfs/samples/nsddevices.sample /var/mmfs/etc/nsddevices**

c. Edit /var/mmfs/etc/nsddevices to look like the example below:

```
osName=$(/bin/uname -s)
if [[ $osName = Linux ]]
then
CONTROLLER_REGEX='mpath[a-z]+'
for dev in $( /bin/ls /dev/mapper | egrep
$CONTROLLER_REGEX )
do
#
# dmm vs. generic is used by Spectrum Scale to prioritize
internal order of searching through available disks,
# then later Spectrum Scale discards other disk device names
that it finds that match as the same NSD device by a different
path.
# For this reason, dmm vs. generic is an important distinction if
you are not explicitly producing the entire and exclusive set
# of disks that Spectrum Scale should use, as output from this
script (nsddevices) _and exiting this script with a "return 0".
-Brian Finley_
#
echo mapper/$dev dmm
#echo mapper/$dev generic
done
fi

if [[ $osName = AIX ]]
then:
# Add function to discover disks in the AIX environment.
fi
# To bypass the Spectrum Scale disk discovery (/usr/lpp
/mmfs/bin/mmdevdiscover),
return 0
# To continue with the Spectrum Scale disk discovery steps,
# return 1
```

d. **% chmod +x /var/mmfs/etc/nsddevices**

e. Execute /var/mmfs/etc/nsddevices.

Example:

```
# /var/mmfs/etc/nsdevices
```

```
mapper/mpatha dmm
mapper/mpathb dmm
mapper/mpathc dmm
```

f. Create NSD stanzas file that uses the multipath aliases.

Example:

Edit /var/hpss/ghi/gpfs_config/nsd.stanzafile and insert the lines below.

```
%nsd: device=/dev/mapper/mpatha
nsd=nsd1
servers=kenosha,sundance
usage=dataAndMetadata
```

g. Continue with Creating NSDs.

4. Enable DMAPI on Spectrum Scale file system.

```
% mmchfs <file system> -z Yes
```

5. DB2

5.1. Users and Groups

The DB2 accounts should be created when HPSS is installed. Verify the following UserIDs and GroupIDs exist on the HPSS Core Server and GHI Nodes:

Table 1. User & Group IDs

UserID:	Group ID:
hpss	hpsssrvr
hpssdmg	hpssdb
hpssdb	

5.2. ACL Account Validation



This is a sample configuration from the CORE Server. The steps below are examples. Verify the ACL menu on your system before selecting any option.

% /opt/hpss/bin/hpss_server_acl

hsa> acl -t CORE

PVL Mount Notification Interface (v1) 007ff347-e533-1cc6-b22d-02608c2cedf4

Client Interface (v1) 32ba9692-4667-11d6-aa3a-0004ac49692b

Account Validation Interface (v1) 647f22a8-a1e9-11d3-a739-000001341966

Realtime Monitor Interface (v1) 80c9a256-2f13-11d3-a0c8-000001341966

Select an interface

Choose an item by number (RET to cancel):

> 3

hsa> show

```
perms - type - ID (name) - realm ID (realm)
=====
rw-c-d- - user - 303 (hpsssm) - 80960 (RS37.LOC.GOV)
rw-c-dt - user - 307 (hpssmps) - 80960 (RS37.LOC.GOV)
r--c--- - user - 308 (hpssfs) - 80960 (RS37.LOC.GOV)
r--c--- - user - 314 (hpssftp) - 80960 (RS37.LOC.GOV)
-----t - any_other
```

hsa> add user hpssdmg rwc

hsa> show

```
perms - type - ID (name) - realm ID (realm)
=====
rw-c-d- - user - 303 (hpsssm) - 80960 (RS37.LOC.GOV)
rw-c-dt - user - 307 (hpssmps) - 80960 (RS37.LOC.GOV)
r--c--- - user - 308 (hpssfs) - 80960 (RS37.LOC.GOV)
rw-c--- - user - 309 (hpssdmg) - 80960 (RS37.LOC.GOV)
r--c--- - user - 314 (hpssftp) - 80960 (RS37.LOC.GOV)
-----t - any_other
```

hsa> quit

5.3. Setup GHI Tablespace

GHI should be configured to use the same DB2 storage group that is used in HPSS.



This is a sample configuration from the CORE Server. The actual configuration for DB2 should be determined during the system engineering/planning phase of the deployment. The GHI DB2 mapping table has the potential to become very large and care should be taken in configuring DB2 to handle it.



Repeat this section to setup the GHI tablespace on the HA Backup Core Server for proper failover operations.

1. Source the database profile.

```
% source ~hpssdb/sqllib/db2profile
```

2. Create the database.

```
% db2 "CREATE DATABASE HGHI ON /hpss1_tbspc1/hpssdb  
/hghi, /hpss1_tbspc2/hpssdb/hghei, /hpss1_tbspcX/hpssdb/hghei"
```

Where:

X = the number of table spaces you have on your system.

3. Connect to the HGHI database.

```
% db2 CONNECT TO HGHI
```

4. Create the bufferpool used for GHI DB table space.

```
% db2 "CREATE BUFFERPOOL "SMALLTABLES" SIZE  
AUTOMATIC PAGESIZE 4K"
```

5. Create the bufferpool used for GHI mapping tablespace.

```
% db2 "CREATE BUFFERPOOL bp32k SIZE 10000 PAGESIZE  
32K"
```

6. Create the storage group.

```
% db2 "CREATE STOGROUP "IBMSTORGROUP" \  
ON '/hpss1_tbspc1', '/hpss1_tbspc2', '/hpss1_tbspcX' \  
OVERHEAD 6.725000 \  
DEVICE READ RATE 100.000000 \  
DATA TAG NONE \  
SET AS DEFAULT"
```

7. Create DB2 tablespaces.

- a. Create DB2 tablespace for GHIDB.

```
% db2 "CREATE LARGE TABLESPACE GHIDB PAGESIZE
```

```
4K \
MANAGED BY AUTOMATIC STORAGE \
USING STOGROUP "IBMSTORGROUP" \
AUTORESIZE YES \
INITIALSIZE 32M \
MAXSIZE NONE \
EXTENTSIZE 128 \
PREFETCHSIZE AUTOMATIC \
BUFFERPOOL "SMALLTABLES" \
OVERHEAD 7.500000 \
TRANSFERRATE 0.060000 \
NO FILE SYSTEM CACHING \
DROPPED TABLE RECOVERY ON \
DATA TAG NONE"
```

- b. Create DB2 tablespace for GHI MAPPING.

```
% db2 "CREATE LARGE TABLESPACE GHIMAPPING
PAGE SIZE 32K \
MANAGED BY AUTOMATIC STORAGE \
USING STOGROUP "IBMSTORGROUP" \
AUTORESIZE YES \
EXTENTSIZE 128 \
PREFETCHSIZE AUTOMATIC \
BUFFERPOOL "BP32K" \
DATA TAG NONE \
OVERHEAD 7.500000 \
TRANSFERRATE 0.060000 \
MAXSIZE NONE \
NO FILE SYSTEM CACHING \
DROPPED TABLE RECOVERY ON"
```

8. Grant connect to the database.

```
% db2 "grant connect on database to user hpss"
% db2 "grant createtab on database to user hpss"
% db2 "grant dbadm on database to user hpss"
```

9. Configure the primary, secondary logs, log archives, log file size, and number of logs similar to the standard of the HPSS databases.

```
% db2 "update db cfg for hghi using NEWLOGPATH
<primary_log_path> hghi"
% db2 "update db cfg for hghi using MIRRORLOGPATH
<secondary_log_path> hghi"
% db2 "update db cfg for hghi using AUTO_MAINT off"
% db2 "update db cfg for hghi using AUTO_RUNSTATS off"
```

```

% db2 "update db cfg for hghi using AUTO_TBL_MAINT off"
% db2 "update db cfg for hghi using LOGARCHMETH1 \
DISK:/ <primary_log_archive_path>/hghi/"
% db2 "update db cfg for hghi using LOGARCHMETH2 \
DISK:/<secondary_log_archive_path>/hghi/"
% db2 "update db cfg for hghi using LOGFILSIZ 25000"
% db2 "update db cfg for hghi using LOGPRIMARY 10"
% db2 "update db cfg for hghi using LOGSECOND -1"

```

Table 2. LOGBUFSZ

Machine Memory	LOGBUFSZ < Table Value
< 16 GB RAM	4096
16 - 64 GB RAM	8192
> 64 GB RAM	16384

```

% db2 "update db cfg for hghi using LOGBUFSZ <table
value>"

```

```

% db2 "update db cfg for hghi using DFT_QUERYOPT 2"

```

10. Disconnect from the database.

```

% db2 disconnect all

```

11. Restart the database.

```

% rc.hpss stop
% rc.hpss start

```

5.4. Install DB2 Client

Install the DB2 client on each Spectrum Scale quorum node (all nodes which include "quorum" in the "Designation" column from the **mmlscluster** command). Follow the DB2 Command Reference document to install the client.

5.5. Create DB2 Permanent License

Create a permanent license on each Spectrum Scale quorum node that has the DB2 client installed.

1. Create License.

```
% cd /opt/ibm/db2/V10.5/adm (on AIX, substitute 'IBM' for 'ibm')  
% ./db2licm -a <path name to DB2 generic license file>
```



The generic DB2 license file (“*/db2/license/db2ese.lic”) can be found on the DB2 Installation CD or image. It can also be obtained by contacting your HPSS customer support representative.

2. Update the license with the number of processors your quorum node possesses.

```
% ./db2licm -n db2ese <number of processors>
```



Refer to the DB2 Command Reference document for more information on how to use the db2licm utility to manage the DB2 license. Create the DB2 database connection on the GHI Session nodes, which should already have the DB2 client installed per the prerequisites.

a. Create an instance.

```
% /opt/ibm/db2/V10.5/instance/db2icrt -a SERVER -s  
client -u hpssdb hpssdb
```

Verify the hpssdb \$HOME directory. If the \$HOME directory for hpssdb is not set to “/var/hpss/hpssdb”, issue the following command:

```
% ln -s <hpssdb_home> /var/hpss/hpssdb
```

Where:

<hpssdb_home> = the assigned \$HOME directory for hpssdb.

b. Catalog the database profile.

```
% source ~hpssdb/sqllib/db2profile  
% db2 catalog tcpip node $NODE remote $HPSS_CORE  
server $PORT  
% db2 catalog db hghi as hghi at node $NODE
```

Where:

```
$NODE = unique name; recommend using short host name
of current machine. +
$HPSS_CORE = hostname of the HPSS Core server. +
$PORT = port number acquired from the Core server
/etc/services - +
Steps to check hpssdb port on core server:
1) Source the database profile.
% Source ~hpssdb/sqllib/db2profile
2) Run the command
% db2 get dbm cfg | grep SVCENAME
3) Look at the value for the SVCENAME
TCP/IP Service name          (SVCENAME) =
db2_hpssdb
4) Cat the /etc/services file and grep for the
SVCENAME from above.
% cat /etc/services | grep db2_hpssdb
5) Use the port number found from the grep of the
/etc/services file for $PORT.
```

- c. Cycle DB2 on the HPSS Core Server to pick up the changes.

```
% rc.hpss stop
% rc.hpss start
```

6. HPSS

HPSS Core server must also be able to connect to the network configured for the Spectrum Scale configuration. For example, if the Spectrum Scale cluster is configured exclusively on a data network, HPSS must be able to connect to that data network, even if the Spectrum Scale nodes also have and additional network to connect to the HPSS Core Server.

6.1. HPSS Users and Groups

GHI needs three users (hpss, hpssdb, hpssdmg) and two groups (hpss, hpsssrvr) on all GHI nodes that will have the HPSS client installed. The user and group ID numbers created on the GHI nodes must match the corresponding user and group ID numbers on the HPSS Core Server.

Use the system command `id` to verify the required users and groups exist.

```
% id <user>
```

Table 3. User & Group IDs

User	Group	Home Directory:
HPSS	HPSS	/opt/hpss
hpssdb	hpss	/var/hpss/hpssdb
hpssdmg	hpsssrvr	/home/hpssdmg



If any of the above users or groups do not exist, use the `useradd` system command to add them.

```
% useradd -d <home directory> -g <group> -p hpss  
<user>
```

6.2. Install HPSS

Install the RPMs on all the GHI nodes:

- * `hpss-clnt-<version>.el<level>.<type>.rpm`
- * `hpss-lib-<version>.el<level>.<type>.rpm`
- * `hpss-lib-devel-<version>.el<level>.<type>.rpm`

Where:

```
<level> = the RedHat OS level, 6 or 7  
<type> = the architecture of the GHI node machines (e.g.  
x86_64, PPC, etc.).
```

6.3. Configure HPSS Client

1. Set up `/var/hpss/etc` on GHI client machines.
 - a. Copy `/var/hpss/etc/*` from the HPSS Core Server to GHI client.
 - b. Add `HPSS_API_HOSTNAME=<long hostname>` to `/var/hpss/env.conf`
 - c. Add `HPSS_PTHREAD_STACK=524288` to `/var/hpss/env.conf`

2. Set up authentication.

Copy the HPSS PAM module (`/etc/pam.d/hpss`) from the HPSS Core Server to `/etc/pam.d/hpss` on all GHI nodes.

If using Kerberos authentication, copy `/etc/krb5.conf` from the HPSS Core Server to all GHI nodes.

1. Specify the *HPSS_NET_FAMILY*. Ensure that the HPSS client configuration has the correct `HPSS_NET_FAMILY` in `/var/hpss/etc/env.conf`.

Examples:

ipv6_only

ipv4_only

ipv6

7. GHI Installation & Configuration

7.1. Install GHI

The following RPMs must be installed on all GHI nodes:

* `ghi-lib-2.5.0.1-0.<architecture>.rpm`

* `ghi-2.5.0.1-0.<architecture>.rpm`

These directories will be created when the RPMs are installed:

* `/opt/hpss/src/ghi`

* `/opt/hpss/src/ghi/man`

* `/opt/hpss/bin`

* `/opt/hpss/lib`

* `/var/hpss/ghi`

* `/var/hpss/ghi/policy`

* `/var/hpss/ghi/config`

* `/var/hpss/ghi/config/templates`

7.2. GHI Users and Groups

All authentication and authorization are done using the `hpssdmg` principal. The numeric IDs must match those on the HPSS Core Server, which may be obtained from the `"/etc/passwd"` file on your HPSS Core Server.

1. Verify `hpssdmg` User ID exists on each GHI node.
2. Verify Group ID `hpssrsvr` is set for `hpssdmg`.

If User ID and Group ID do not exist, create them.

7.3. Configure GHI

GHI is configured using command-line tools. All of the GHI commands discussed in this section are fully documented in the GHI Management Guide.

These are the steps to configure GHI:

1. Create GHI cluster from the Spectrum Scale configuration.
2. Add Spectrum Scale file system for GHI to manage.
3. Add IOMs for each GHI managed Spectrum Scale file system.
4. (Optional) Configure Syslog Facilities.

7.4. Create GHI Cluster

Define the overall cluster configuration, including the nodes which will be known to GHI (not necessarily all nodes known to Spectrum Scale). This is accomplished via **ghicrcluster** command. **ghicrcluster** must run on the session node that is designated as "Cluster manager node". Use **mmlsmgr** command to determine which node is the cluster manager. In addition, Spectrum Scale must be running when defining the cluster configuration to GHI.

All nodes that are designated as quorum with **mmlscluster** command must be listed after the cluster manager. This will allow GHI to assign them as a manager node in the case of a failover.

```
% ghicrcluster [-v] {--aix | - -ppc_64 | --x86_64 <node>} ...  
<GHI_node> ..  
% ghicrcluster -N [-v] {--aix | --ppc_64 | --x86_64 <node>} ...  
<nodelist_file>
```

Where:

```
<Architecture> = System hardware type configuration.  
<node> = This is the node the source code was compiled on  
(*ghilscluster*).  
<GHI_node> | <nodelist_file> = Node list of machines from  
*mmlscluster* which will have the designation of "manager"  
in the command *mmlscluster*.
```



Command below is an example.

```
% ghicrcluster -v --ppc_64 firefly firefly falcon
```

After **ghicrcluster** returns "Done.", restart Spectrum Scale and GHI.
% ghishutdown -G
% ghistartup -G



If **ghicrcluster** fails during the configuration, retry the configuration step with the "-r" option after the errors from the failure are resolved.
% ghicrcluster -r [-v]

7.5. Create Spectrum Scale file systems GHI will manage

Use the command **ghiaddfs** for each file system to be created, which may be issued from any node in the cluster. File systems to be defined must NOT be mounted in Spectrum Scale when the **ghiaddfs** command is issued. **ghiaddfs** will supply default values for the file system which can be updated or changed with the command **ghichfs**.

For each file system, the name and mount point are to be supplied by the user. The ports to be used by the associated SD and ED may also be user-supplied or left to their default values.

% ghiaddfs [-v] <FS_Name> [-c "# <comment>"] <Mount_Point> [<SD_Port> <ED_Port>]

Where:

<FS_Name> = The same as the Spectrum Scale configuration name.
<Mount_Point> = The same as the Spectrum Scale configuration mount point.



Beginning with version 2.4, GHI introduced the capability to configure a GHI-managed file system to be read-only with respect to HPSS. The preceding discussion of **ghiaddfs** applied to creation of a normal, full-access file system. Please refer to the GHI Management Guide for a complete discussion of GHI read-only file systems and how they should be configured.



Command below is an example.

```
% ghiaddfs firef1 /firef1
```

7.6. Create IOMs for each GHI Managed File systems

The default port selected for an IOM is 80*x*2, where **x** is the order in which the file system was configured (8012 for the first configured file system, 8022 for the second configured file system, etc.). For more details about ports, see the GHI Management Guide.

Refer to the GHI Management Guide for more details about **ghiaddiom**.

```
% ghiaddiom [-vd|D] <File System name> [-c "# <comment>"]  
<IOM Node> \ <active_on_session_node>  
<estimate_transfer_rate>
```

Where:

```
<File System name> = Name of the file system added with the  
ghiaddfs command.  
<IOM Node> = Name of node the IOM will run.  
<active_on_session_node> = Active state of the IOM on the  
manager session node.  
<estimate_transfer_rate> = Estimated data transfer rate.
```



Command below is an example.

```
% ghiaddiom -v firef1 firefly:8012 TRUE 1GB
```

7.7. Modify xinetd.conf for number of IOMs



This setting should be adjusted whether using RHEL 6.X or RHEL 7.X

```
% vi /etc/xinetd.conf
```

change *cps = 50 10* to *cps = <IOM Thread Pool Size * number of IOMs>*

7.8. (Optional) Configure Syslog Facilities



Configuring syslog facilities is different than just configuring GHI syslog. Please contact GHI support representative with any questions on this optional configuration.

Syslog facilities can be configured using the FACILITY_CODE macro in Makefile.macros.

To configure syslog facilities, install the following RPM:

* ghi-src-2.5.0.1-0.el7.x86_64.rpm

The FACILITY_CODE macro can be set to any of the facility codes defined by the syslog standard, but GHI defaults to LOG_USER.

Table 4. Facility Code

Facility Code	Keyword	Description
0	LOG_KERN	kernel messages
1	LOG_USER	user-level messages
2	LOG_MAIL	mail system
3	LOG_DAEMON	system daemons
4	LOG_AUTH	security/authorization messages
5	LOG_SYSLOG	security generated internally by syslogd
6	LOG_LPR	line printer subsystem
7	LOG_NEWS	network news subsystem
8	LOG_UUCP	UUCP subsystem
9	LOG_CRON	clock daemon

Facility Code	Keyword	Description
10	LOG_AUTHPRIV	security/authorization messages
11	LOG_FTP FTP	daemon
15	LOG_DAEMON	scheduling daemon
16	LOG_LOCAL0	local use 0 (local0)
17	LOG_LOCAL1	local use 1 (local1)
18	LOG_LOCAL2	local use 2 (local2)
19	LOG_LOCAL3	local use 3 (local3)
20	LOG_LOCAL4	local use 4 (local4)
21	LOG_LOCAL5	local use 5 (local5)
22	LOG_LOCAL6	local use 6 (local6)
23	LOG_LOCAL7	local use 7 (local7)

7.9. Information Lifecycle Management (ILM) Policies

GHI makes use of Spectrum Scale ILM policies. A policy is a plain-text file that describes files and directories to be included or excluded from processing. IBM provides templates which you may use as a starting point to configure custom policies. These templates can be found in the `/var/hpss/ghi/policy` directory. Below is a list of policy templates.

migrate.policy

This file can be placed in any directory in the system. The policy should have separate rules for aggregates and non-aggregates. The script, `ghi_migrate`, that gets invoked from the policy engine requires a “-a” option to process aggregates.

reset_incomplete_migration.policy

Use this policy to reset files for which a migration was started but never completed. Such files will show as “[incompletely-migrated]” when listed with “`ghi_ls -h`”. They are “migrated enough” such that Spectrum Scale will not select them to be re-migrated, and the migration-reset process will result in the files being set back to “un-migrated” so that Spectrum

Scale will select them in the next applicable migration policy run. This file can be placed in any directory in the system.

recall.policy

The recall policy does not use a bulk size. The policy generates on list. That list is parsed into aggregates and non-aggregates. The recall.policy file can be placed in any directory in the system.

tape_smart_migration.policy

This is an example used to migrate files in a tape smart manner. Files are migrated by HPSS file families and by path name. This policy can be used in combination with the `-split-filelists-by-weight` option for `mmapplypolicy` to generate file lists that contain elements with the same WEIGHT value.

backup_migration.policy

The migration policy will run a full Spectrum Scale scan and will attempt to migrate any files that are not stored currently in HPSS. The policy file should be updated to reflect the migration rules used for this file system. The policy should be able to select every file that has not been migrated to HPSS and exclude any file which should not be migrated. Verify that the backup migration policy matches what is being backed up in the `backup_metadata.policy` to ensure that files which have not been migrated are included in the metadata backup.

backup_metadata.policy

This policy is used by the Spectrum Scale SOBAR `mmimgbackup` command. Spectrum Scale file system namespace and file metadata are sent to GHI and HPSS.

[WARNING] Do not change the `backup_metadata` policy without contacting IBM GHI support.

backup_error.policy

The backup error policy contains the rules that are used to validate the capture of a file system's metadata.

[WARNING] Do not change the `backup_error` policy without contacting IBM GHI support.

threshold.policy

The Spectrum Scale ILM threshold policy provides the capability for GHI to space manage the Spectrum Scale file system. New and modified Spectrum Scale files are copied to HPSS on a periodic basis. When the Spectrum Scale file system reaches a pre-defined space threshold, the Spectrum Scale ILM threshold policy is executed to identify file candidates whose data can be removed from the file system. This file must be copied from the `/var/hpss/ghi/policy` directory to `/var/hpss/ghi/policy/<file system>` and modified to be file system specific. The

script, `ghi_migrate`, is invoked from the policy engine requires a “-p” option to punch holes in the file system.

To apply a threshold policy to a file system, issue the following command:
% `mmchpolicy <file system> <full path of threshold.policy>`

8. Backup & Restore

8.1. Backups

To backup a Spectrum Scale file system, use the GHI `ghi_backup` command line interface. The backup interface uses the Spectrum Scale `mmimgbackup` command, which uses the ILM policy management engine. Each file system to be backed-up uses its own copy of each of the following backup policy templates that resides in the `/var/hpss/ghi/policy` directory:

GHI backups use the Spectrum Scale snapshot feature to take a point-in-time image of the file system.

When running a backup:

1. A snapshot of the Spectrum Scale namespace is saved after the backup migration policy and any other running migration policies completed.
2. The state of each of the files is saved.

backup_migration.policy:

The backup migration policy contains the migration rules for the Spectrum Scale file system to be backed up. The rules can migrate files as an aggregate or non-aggregates. The rules must select all the files to be backed up.

backup_metadata.policy:

The backup metadata policy contains the rules that previous GHI versions need to capture a file system’s metadata. The new image backup feature does not require a metadata policy for metadata backup. The metadata is contained in the image generated by Spectrum Scale as part of the backup process.

backup_error.policy:

The backup error policy contains the rules that are used to validate the capture of the file system’s metadata.



IBM recommends running a daily backup. Checking the backup logs daily to correct any errors is a good practice to ensure successful backups. GHI backup options are image, non-image (full), and incremental. Details of each type of backups are provided in the GHI Management Guide.

Most sites create a crontab entry to run a daily backup.

```
ghi_backup <filesys> <type> [-E <Fsets> [-U <user_args>]] | [-U  
[<user_args>]]
```

```
ghi_backup <filesys> <type> [-E -F <List of Fsets> [-U  
<user_args>]]
```

Where:

```
type = one of the backup type:  
image = image backup using Spectrum Scale SOBAR capability.  
full = full backup using GHI-native backup.  
incr = incremental changes since the last backup using  
GHI-native backup.  
Fsets = filesets which need not be linked-in, either space-  
separated list of filesets  
or  
-F space-separated list of files containing list of filesets  
user_args = args to be passed to policy.
```



Commands below are examples.

```
% ghi_backup firef1 image
```

```
% ghi_backup firef1 full
```

```
% ghi_backup firef1 incr
```

8.2. Restore

```
ghi_restore [-v] [-c config_file] [-R <BUIdx>] -g <temporary space>  
<file system>
```

Where:

```
-v = verbose.
```

```
-c <config_path>          = specifies an alternate
configuration path.
-R <BUIdx>                = resume restoration of specified
backup.
<file system>            = file system (string).
-g <temporary space>     = a separate Spectrum Scale mounted
file system temporary space in addition to the restoring
file system where to place temporary backup files.
```

9. Appendix A - GHI Conversions

9.1. Conversion steps for 2.5.0.1 to 3.0



Doing conversion process will use rpms for GHI 3.0.

1. Shutdown GHI.

```
% ghishutdown -g
```

2. Verify all GHI processes are not in process list

```
% ps -ef | grep ghi
```



If any "defunct" or left GHI processes are still in the list use `kill -9 <pid>` to shut them down.

3. Remove old GHI rpms

```
% rpm -qa | grep ghi - List installed rpms.
```

```
% rpm -e <ghi_rpms installed from above>
```

```
% rpm -qa | grep ghi - Verify no ghi rpms are installed
```

4. Install new GHI 3.0 RPMs



RPMs deployed will vary get with GHI support rep before installing

```
% rpm -ivh <path to GHI 3.0 rpms>/ghi-3.0.0.0-0.<arch>,
ghi-lib-3.0.0.0-0.<arch>, ghi-ishtar-5.1.1.0-0.<arch>
```

5. Run ghiupdate

```
% ghiupdate -vT --<OS arch> <all nodes listed in
ghilscluster> - -T option is Test recommended to run this option
first.
```

```
% ghiupdate -v --<OS arch> <all nodes listed in ghilscluster> -
This will perform actions.
```

6. Start new GHI

```
% ghistartup -g
```

7. Verify PM, CM, MD running

```
% ps -ef | grep ghi
```

8. Mount GHI managed File Systems

```
% mmmount <file system> -a
```

9. Verify SD, ED, and then IOMs are active

```
% tail /<FS>/scratch/mon/mon_sd.out<mon_iom.out>
```



It could take up to 10 minutes for ED to connect to SD. It could take up to 15 minutes for the IOMs to connect to the ED. If after 20 minutes things are still not connecting contact your GHI support rep.

Last updated 2017-12-04 15:27:56 CST