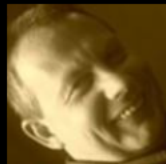
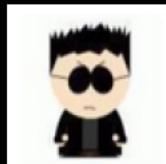


IBM Spectrum LSF

Best Practices for Upgrading your LSF Clusters



Bill.McMillan@uk.ibm.com
Principal Product Manager
Spectrum Computing



Larry Adams
Systems Expert Labs

Upcoming Seminars

Details, Registration and Replays on www.ibm.biz/LSFCommunitySeminars

August 19, 2021

IBM Spectrum LSF: New Features and Futures

Speaker: Bill McMillan, Principal Product Manager

September 2, 2021

What's New in LSF Service Pack 12

Speaker: LSF Support Team

September 16, 2021

Best Practices for Upgrading your LSF Clusters

Speaker: Larry Adams, Expert Labs

September 30, 2021

What if....? Using the LSF Simulator to answer those hard questions

Speakers: John Welch & Renita Leung, LSF SME

October 14, 2021

Simplifying HPC – Just push the button! Tips & Tricks

Speaker: Gabor Samu, Product Management

October 25-28, 2021

IBM TechU: <https://www.ibm.com/training/events>

- Deploying LSF with OCP: Tips and Tricks
- HPC Cloud Bursting on IBM Cloud: Tips and Tricks

November 11, 2021

Expediting PMR's with the LSF Support Tools

Speaker: LSF Support Team

November 29, 2021

High Performance Computing - Health Check Services from Lab Services

Upcoming Seminars

Details, Registration and Replays on www.ibm.biz/LSFCommunitySeminars

August 19, 2021

IBM Spectrum LSF: New Features and Futures

Speaker: Bill McMillan, Principal Product Manager

September 2, 2021

What's New in LSF Service Pack 12

Speaker: LSF Support Team

September 16, 2021

Best Practices for Upgrading your LSF Clusters

Speaker: Larry Adams, Expert Labs

September 30, 2021

What if....? Using the LSF Simulator to answer those hard questions

Speakers: John Welch & Renita Leung, LSF SME

October 14, 2021

Simplifying HPC – Just push the button! Tips & Tricks

Speaker: Gabor Samu, Product Management

October 25-28, 2021

IBM TechU: <https://www.ibm.com/training/events>

- Deploying LSF with OCP: Tips and Tricks
- HPC Cloud Bursting on IBM Cloud: Tips and Tricks

November 11, 2021

Expediting PMR's with the LSF Support Tools

Speaker: LSF Support Team

November 29, 2021

High Performance Computing - Health Check Services from Lab Services

Agenda

IBM Spectrum LSF Packaging

The LSF Install Alternatives

Install/Upgrade Best Practices

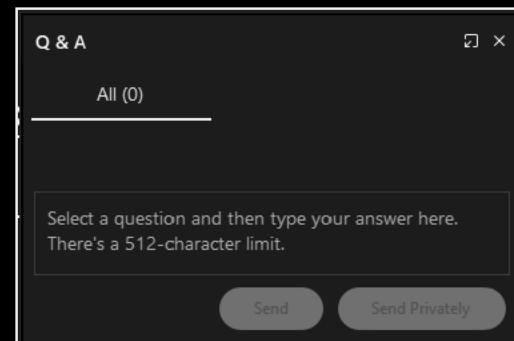
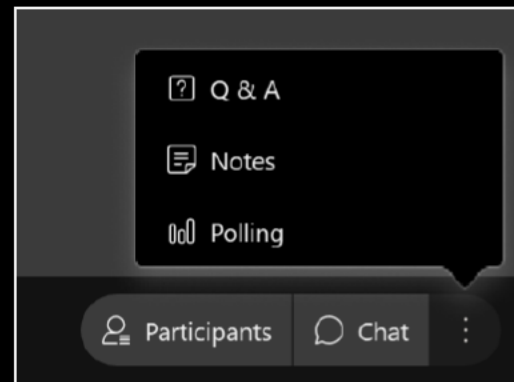
Upgrading LSF Standard

Upgrading the LSF Suites

Recent Security Changes Post LSF 10.1.0.9

What's Next?

Q&A: You can ask questions at any time in the Q&A panel



IBM Spectrum LSF Packaging

Over the last 30 years LSF has evolved from a single scheduler to a full family of products.

Proven track record in delivering customer value

- 10 Major Releases
- 30+ Minor Releases
- 1000+ customer driven features

Historically, odd numbered major releases have delivered significant architectural changes. The next release of LSF will follow this pattern.



The LSF Install Alternatives

LSF Standard

Pros:

- Tried and true. In use by most customers since as early as 2006 (my start date)
- Leverages NFS as a common binary, library, and database location as well as for configuration management.
- Has an established patching process that is, in general, non-service affecting.

Cons:

- Not well suited to distributed environments where NFS is not practical or avoided
- LSF Standard installer does not cover other aspects of system configuration that are critical to HPC operations.

LSF Suites

Pros:

- Leverages Ansible to perform installation tasks in parallel
- Installs all LSF Suite components from a single installer
- Allows Single point of management for deploying node types
- Supports NFS based installs
- Performs additional system tuning as required for optimal system health (/etc/sysctl.d, /etc/security, /etc/profile.d, etc)

Cons:

- Not all architecture types are supported
- Forces non-RPM Linux variants to utilize RPM's (aka SLES, Debian)
- Non-NFS installs beg for a local configuration management tool such as puppet, chef, salt etc. to manage LSF plugins. However, starting in LSF 10.1.0.10 it is possible to force a shared LSF plugin location even when using local installs.

Install/Upgrade Best Practices

- Use an `$LSF_TOP` that does not include a version number (ie `/apps/lsf` and not `/apps/lsf10`)
- Audit your LSF plugins prior to upgrade you will have to copy them during major version upgrades
 - ELIM – System/Cluster Resource Monitoring
 - ESLIM – Static string Resources
 - ESUB – Submission Control
 - EAUTH – Authentication verification
 - EEXEC – Job monitoring
- Don't forget to audit your `/etc/lsf.sudoers` file if in place
- If monitoring Licenses, use either License Scheduler Standard or Basic vs. the legacy ELIM approach
- Rerun `hostsetup` on all server hosts after upgrading or applying service packs to obtain important `systemd` updates
- Have a Parallel Distributed Shell Tool (preferably `pdsh`)
- Remove clutter from `lsf.shared` and `lsf.cluster` files
- Use a configuration management database for your LSF configuration and plugins (`git`, `svn`, etc.)
- Periodically update your LSF security keys, noting that this may be a service impacting change, so it will happen less frequently
- Always open a proactive ticket with IBM prior to major upgrades to make support aware of your intentions and to receive upgrade advice.
- Always review the LSF release notes to verify if any prior bugs that you reported are not included in a planned service pack and to capture any LSF security changes that may be service impacting!
- Engage IBM Systems Expert Lab Services if you are unsure of what steps should be performed and when
- You may have to recompile MPI frameworks such as OpenMPI and any API driven libraries after an upgrade

Recently Introduced LSF Commands

battach – Allow you to get a shell in your docker container

battr – Allow you to assign/remove attributes from hosts for affinity scheduling

bctrld – Used to start & stop daemons by authorized users

bgpinfo – Global Fairshare inquiries

bimages – Find out what containers are distributed where

bsubmit – Allow authorized users to submit a job as a service account

bwait – Blocking wait on a job to finish that does not continually query the LSF batch system

lsportcheck – Show ports in use by an LSF server

lsreghost – Register a host with an LSF cluster

Also note that there are many new “badmin” options and most commands are now supporting JSON output now! Additionally, “bsub” now supports both JSON and YAML based job decks.

Upgrading LSF Standard



LSF Standard Upgrade/Patch Release Notes

What's new in IBM Spectrum LSF

Review the new and changed behavior for each version of LSF.

System requirements and compatibility

The following sections describe the system requirements and compatibility information for version 10.1 of IBM Spectrum LSF.

IBM Spectrum LSF product packages

The IBM Spectrum LSF product consists of distribution packages for supported operating systems, installation packages, and entitlement files.

Getting fixes from IBM Fix Central

After you install or upgrade LSF, use IBM Fix Central to find and download the fixes that are recommended by IBM Support for LSF products. From Fix Central, you can search, select, order, and download fix packs and interim fixes for your system with a choice of delivery options.

Bugs fixed

LSF Version 10.1 releases and Fix Packs contain bugs that were fixed since the general availability of LSF.

Known issues

LSF 10.1 has the following known issues.

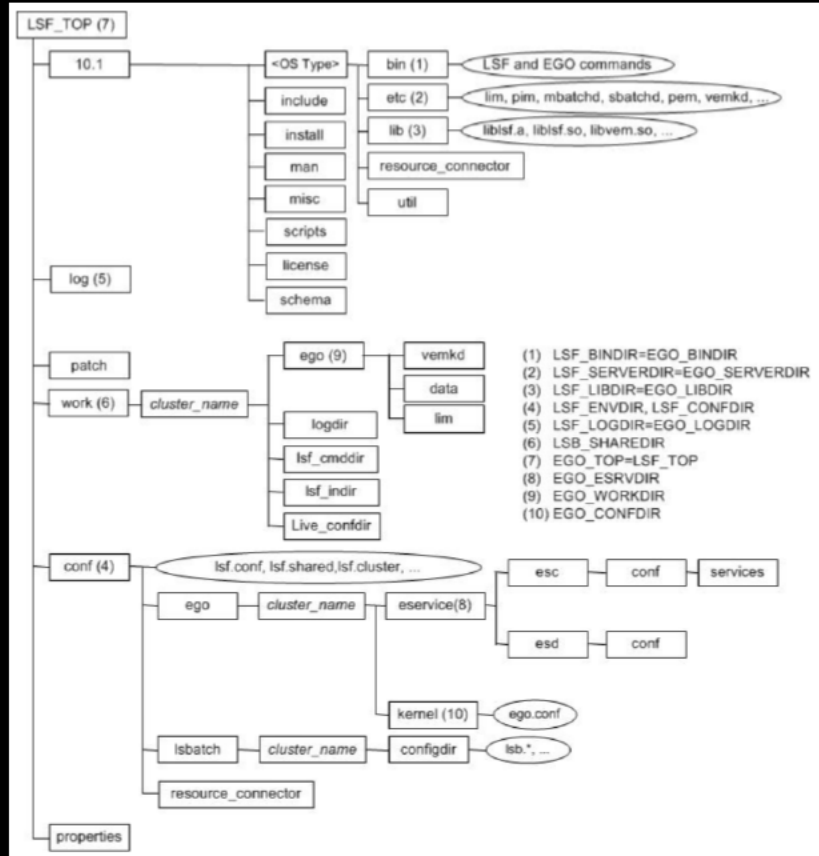
Limitations

LSF 10.1 has the following limitations.

Detailed changes

The following topics describe new and changed commands, options, output, configuration parameters, environment variables, accounting, and job event fields in LSF 10.1.

LSF Directory Structure



← Plugins (`elim`, `esub`, `eauth`)
← Library files (for dynamically linked binaries)

← Critical Operational Database

← Critical Configuration Database

The LSF README Fixed Bugs

IBM Spectrum LSF 10.1 Fix Pack 12 (600488) Readme File

Abstract

LSF Version 10.1 Fix Pack 12. This Fix Pack includes new issues and solutions resolved between **12 November 2020 and 10 June 2021**. For detailed descriptions of the issues and solutions in this Fix Pack, refer to the LSF 10.1 Fix Pack 12 Fixed Bugs List (lsf10.1.0.12_fixed_bugs.pdf), which can be downloaded from Fix central via fix ID lsf-10.1.0.12-spk-2021-Jun-build600488.

Description

Readme documentation for IBM Spectrum LSF 10.1 Fix Pack 12 (600488) including installation-related instructions, prerequisites and co-requisites, and list of fixes.

The new issues addressed in LSF Version 10.1 Fix Pack 12:

ID	Fixed Date	Description
P104243	2021/5/31	This fix ensures that jobs are forwarded to remote clusters, which also enables the LSF resource connector.
P104239	2021/5/26	This fix prevents rusage errors when users submit jobs that have spaces in the rusage string.
P104203	2021/4/25	This fix prevents jobs on spot instances from being wrongly terminated when the spot instance is reclaimed, even though the jobs are requeued and dispatched to run on a new instance.
P104194	2021/4/16	This fix prevents pending jobs from showing the pending reason of "Template host only used for demand calculation" in the LSF resource connector.

The LSF README New Features

The new solutions in LSF Version 10.1 Fix Pack 12:

ID	Fixed Date	Description
1280	2021/6/10	<p>This solution enhances the functionality of the LSF resource connector OpenShift plugin including</p> <ul style="list-style-type: none">- The LSF cluster name can be configured by LSF_OPERATOR_CLUSTER in openshiftprov_config.json.- The timeout value of the OpenShift API request can be configured by TIME_IN_SEC in openshiftprov_config.json.- The status sync mechanism of new Pod can be changed by WAIT_POD_IP in openshiftprov_config.json.- The mount point of each persistent volume claim can be configured by mountPaths field for each template in openshiftprov_templates.json.- The change of the management Pod's IP address does not affect the functionality of cluster.
941	2021/6/10	<p>This solution enforces the LSF security mechanism by default by enabling LSF_STRICT_CHECKING=ENHANCED, LSF_AUTH_QUERY_COMMANDS=Y, LSF_ENV_OVERRIDE=N, and also introduces the "badmin security view" command to view the cluster's security mechanism configuration. Meanwhile, the cluster administrator can set the grace period of the previous LSF_EAUTH_KEY by specifying the LSF_EAUTH_OLDKEY and LSF_EAUTH_OLDKEY_EXPIRY parameters in the /etc/lsf.sudoers file.</p>
159 166	2021/5/31	<p>This solution correctly handles new VM instances with the same hostname and IP address joining the cluster in the cloud environment.</p>
P104238	2020/5/26	<p>This solution enhances the LSF resource connector to support new Google Cloud features, including instance templates, local SSDs, preemptible VMs, and bulk APIs.</p>

Binary Availability Matrix

Component	linux2.6- glibc2.3- x86_64 / linux3.10- glibc2.17- x86_64	linux3.10- glibc2.17- ppc64le	linux3.12- glibc2.17- armv8	aix- 64	hpuxia64	macos- x86_64	sparc- sol10- 64	x86- 64- sol10
LSF/TaskStarter	√	√	√	√	√	√	√	√
LSF/augmentstarter	√	√	√	√	√	√	√	√
LSF/bacct	√	√	√	√	√	√	√	√
LSF/badmin	√	√	√	√	√	√	√	√
LSF/bapp	√	√	√	√	√	√	√	√
LSF/batch-acct	√	√	√	√	√	√	√	√
LSF/battach	√	√	-	-	-	-	-	-
LSF/battr	√	√	√	√	√	√	√	√
LSF/bbot	√	√	√	√	√	√	√	√
LSF/bchkpnt	√	√	√	√	√	√	√	√
LSF/bclusters	√	√	√	√	√	√	√	√
LSF/bconf	√	√	√	√	√	√	√	√

General Process for LSF Standard Patches or Fix Packs

1. Download Fix Pack Binary Packages from Fix Central
2. Review README for Custom Instructions/Breaking Changes
3. Inactivate Queues (badadmin qinact all)
4. Run “patchinstall” for each binary type
5. Rerun “hostsetup” as root on each LSF Server using “pdsh” or equivalent
6. Implement changes to maintain non-breaking state per the README
7. Run “pdsh -w nodes -f 20 systemctl restart lsfd” to leverage “pdsh” to restart key LSF daemons
8. Reactivate Queues (badadmin qact all)

With only few exceptions, this change is NOT a service impacting upgrade. Note that “lsfrestart” is currently broken post FP9. A ticket has been logged to fix the script.

General Process for LSF Standard Major Upgrades

<https://www.ibm.com/docs/en/spectrum-lsf/10.1.0?topic=linux-upgrade-lsf-unix>

1. Download LSF Installer, Entitlement, and Binary Type Packages from Passport Advantage
2. Download the latest LSF Fix Pack from Fix Central for each Binary Type
3. Untar the LSF Installer Package
4. Inactivate Queues (badadmin qinact all)
5. Edit install.conf as if it were a new cluster
6. Failover the LSF Master to the Secondary Master
7. Run the LSF Installer from the primary master
8. Accept License Agreement
9. Select Binary Type Packages to Install
10. Wait for Upgrade to Complete
11. Copy any critical LSF plugins to the new binary directories
11. Review README for Version Upgrade for Breaking Changes
12. Review README for Fix Pack for Breaking Changes
13. Apply the Fix Packs using “patchinstall”
14. Rerun “hostsetup” on the Primary Master
15. Apply all changes required to ensure that there are no breaking changes in the lsf.conf file
16. Start the LSF Daemons on the Primary Master using “systemctl restart lsfd”
17. Wait for the Primary Master to retake control (lsid command)
18. Use “pdsh” to rerun “hostsetup” on all remaining LSF servers
19. Restart the remainder of the cluster using “pdsh -w nodes -f 20 systemctl restart lsfd”
20. Reactivate Queues (badadmin qact all)

With only few exceptions, this change is NOT service impacting.

Important Notes/Exceptions

- To Apply the Latest LSF Service Pack, you can bypass all prior Fix Packs unless otherwise notified in the Fix Pack README.
- Previously running jobs will continue to utilize OLD RES binaries, so do not remove the previous binary directories until those jobs clear from the system.
- To implement some new security features and settings available in Fix Packs 10 through 12, you can introduce breaking changes, which require a draining of the cluster and full restart.
- Some new Security Features, if enabled, force the use of the `/etc/lsf.sudoers` files and a new `eauth` binary.
- Carefully check the Known Issues for your LSF release, especially if you run parallel jobs. Some upgrades are best executed if certain settings in `lsf.conf` are in place prior to upgrading
- Starting with Fix Pack 10, by default, some LSF commands will not longer run as root, so ensure that you review any custom integrations, esubs, elims for “batch” command use before upgrading.
- Some LSF Changes, will require RTM to be upgraded prior to LSF. Therefore, upgrade RTM prior to making major LSF changes across clusters and leverage the latest available RTM Data Collector.
- Regarding RTM and LSF 9 clusters, prior RTM Data Collectors should continue to work with your LSF 9 Clusters after upgrading RTM. However, ensure you backup these RTM Binaries before upgrading and restore them if required.
- It is recommended that you upgrade License Scheduler at the same time as you upgrade LSF. Though those steps are not covered here, you can find them in the License Scheduler README file.

Important Notes Relative to Legacy GPU Integrations

- Prior to LSF 10.1.0.6, all GPU integrations relied on LSF host-based resources that had to be created in the `lsf.shared` and `lsf.cluster` file. Starting in LSF 10.1.0.6, these shared resources are no longer required and should be removed after upgrade.
- Nvidia DCGM has version dependency issues and does not provide good backward compatibility. If you are using Nvidia GPU's you must plan to upgrade your DCGM to one that is compatible with the LSF Fix Pack that you are upgrading to, otherwise, the GPU resource reporting will potentially break. If you have any questions about how to upgrade, contact your IBM Support representative.
- If after upgrading from a legacy LSF version, if there are pending jobs that relied on these legacy resources, you will be required to "bmod" those jobs to update their GPU requirements for them to dispatch correctly.
- There were scalability changes in LSF starting with LSF 10.1.0.11 that if you use RTM, you should ensure that you upgrade RTM to a least Fix Pack 11 before upgrading RTM so that your GPU Graphs can continue to update.

Upgrading LSF Suites



LSF Suite for Workgroups Upgrade/Patch Release Notes

Getting fixes from IBM Fix Central

After you install or upgrade IBM Spectrum LSF Suite for Workgroups, use IBM Fix Central to find and download the fixes that are recommended by IBM Support for LSF products. From Fix Central, you can search, select, order, and download fix packs and interim fixes for your system with a choice of delivery options.

Determining the cluster configuration

There are different update procedures depending on whether the master host is installed locally and whether there is a shared configuration between hosts. Determine the configuration of your cluster to determine the specific update procedure to use.

Updating IBM Spectrum LSF Suite for Workgroups on a shared file system

Update IBM Spectrum LSF Suite for Workgroups on a shared file system, which affects all hosts in the cluster.

Updating IBM Spectrum LSF Suite for Workgroups on a single-host POC installation

Update IBM Spectrum LSF Suite for Workgroups that is locally installed on a single host.

Updating IBM Spectrum LSF Suite for Workgroups on a basic or large installation

Update IBM Spectrum LSF Suite for Workgroups on a basic or large installation. IBM Spectrum LSF Suite for Workgroups is locally installed, but with a shared configuration directory. This means that the binary files are local, but the configuration files are on a shared file system.

Updating the SMPI fix pack manually

As of version 10.2 Fix Pack 9, the bundled SMPI will not be refreshed for each Fix Pack. To refresh the SMPI manually, it can be downloaded directly from Fix Central and applied to the IBM Spectrum LSF Suite for Workgroups host.

LSF Suite Documentation


- LSF Suite Documentation is separated by Workgroup, HPC, Enterprise, and Enterprise Plus
- LSF Suite Documentation Explains what is new by point release.
- Describes new features and breaking changes, if any
- Generally, the LSF Fix Pack released with the suites aligns with the LSF binaries included (aka 10.2.0.11 will include LSF 10.1.0.11)
- LSF Suite users and administrators should follow the LSF standard documentation when looking for LSF usage tips and techniques.
- Spectrum MPI must be updated separately in all future releases. (Available for the ppc64le architecture only).

IBM Spectrum LSF Suite for Workgroup

Change version

10.2.0

☒ Show full table of contents

 Filter on titles

IBM Spectrum LSF Suite for Workgroups
V10.2 documentation

Overview

Getting started

Quick start guides

Release notes

What's in Version 10.2 Fix Pack 11

What's new in 10.2 Fix Pack 11

Software packages

What's in Version 10.2 Fix Pack 10

What's in Version 10.2 Fix Pack 9

What's in Version 10.2 Fix Pack 8

What's in Version 10.2 Fix Pack 7

What's in Version 10.2 Fix Pack 6

Learn more

General Process for Upgrading LSF Suites Fix Packs

1. Download Fix Pack Binary Packages from Fix Central
2. Backup your shared LSF directory in case of a potential backout scenario
3. From the “deployer” node, run the binary which will provide additional instructions on upgrading
4. Inactivate Queues (badmin qinact all)
5. From the “deployer” node switch directories to “/opt/ibm/lsf_installer/playbook”
6. From the “deployer” node as root, run “ansible-playbook -i lsf-inventory lsf.upgrade.yml”
7. Reactivate Queues (badmin qact all)

With only few exceptions, this change is NOT a service impacting upgrade though some LSF commands may not work properly during some aspects of the upgrade.

For LSF Suite Patches, you will simply change the ansible syntax to run the “lsf-apply-fix.yml” playbook

Special Notes Related to LSF Suites

- If you are running LSF Suites prior to 10.2.0.9 it is recommended that you migrate your ElasticSearch services to ElasticSearch 7.14+ prior to upgrading the Suites packages.
- If you are impacted by this upgrade scenario, contact IBM Support or IBM Systems Expert Lab services to request help migrating your old data or on what steps will be required to upgrade.
- If the old data is not critical, you should consider a complete re-install of the LSF Suites which will be service impacting change.
- Ansible is great, you can selectively upgrade hosts in your cluster by simply using the “--limit” option to the ansible-playbook command. For example, to upgrade some hosts before the masters, you can use the “--limit” option to setup testing of the new build.
- Outside of ElasticSearch, LSF Suite upgrades, should be non-service impacting and strait forward.
- You may run the Ansible ‘pre-deploy.yml’ playbook will check for any compatibility issues prior to executing the actual upgrade
- If you are running without NFS, you may be required to backup and then re-push your custom LSF plugins to the various \$LSF_SERVDIR locations after the playbooks have run.
- If this is required, you may have to restart the LSF daemons using pdsh from the deployer node after re-installing your plugins.

Important Recent Security Changes in LSF Post 10.1.0.9

- Starting in LSF 10.1.0.10 through LSF 10.1.0.12, the Spectrum LSF development team has been incrementally hardening LSF from a security perspective.
- Several changes that are potentially breaking changes were introduced that must be put into your upgrade plans.
- Some of these changes, if you want to apply them, require a hard cluster shutdown. However, most changes are optional if proper care is taken.
- Starting with LSF 10.1.0.10, it is critical that you thoroughly review the README files to avoid cluster downtime if you plan to bypass those settings.
- Starting in LSF 10.1.0.12, we have introduced a “badadmin security view” to report on the security hardening level of your cluster.

Important Recent Security Changes in LSF Post 10.1.0.9

- Changes in `lsf.conf` include:
 - Breaking/Deprecating `LSF_ROOT_REX=“...”`
 - Introducing `LSF_ROOT_USER=Y|N`
 - Introducing `LSF_ADDON_HOSTS=“...”`
 - Introducing `LSF_ENV_OVERRIDE=Y|N`
 - Introducing `LSF_EXT_SERVERDIR=“...”`
 - Introducing `LSF_STARTUP_USERS=“...”`
 - Forcing `LSF_STRICT_CHECKING=“...”` on new installs
 - Forcing `LSF_AUTH_QUERY_COMMANDS=Y` on new installs
- General security hardening includes:
 - Requiring a static encoding key required in the `/etc/lsf.sudoers` file
 - Providing a revised encryption algorithm for LSF batch command checksums to mitigate MIM attacks
 - The ability to block bypassing ESUB through forcing a common LSF plugin directory
 - Preventing “root” from running batch commands by default
 - Introducing the “bctrld” command to control who can manage LSF daemons
 - Introducing a “--setuid” option when using the “hostsetup” binary
 - No log file will be owned by “root” by default
 - Blocking `ldpreload` from overriding system commands

What's Coming or Going?

- I'm riding my 40-year-old touring bike weekly losing weight and feeling great and currently COVID free
- I'm moving on to a new role outside of IBM
- My last day @ IBM is 9/24
- I'll still be wrapped around both the LSF and RTM axles in my new role
- I hope to be able to remain an advocate for both LSF and RTM moving forward, but from the customer perspective
- Look me up on LinkedIn if you want to stay in touch
- Stay safe and healthy. Don't forget daily "vitamin d", and "zinc" are essential to regulate a healthy immune system.



