Concertim

HPC Health Observations

An Open Reference Guide

Generated 24 June 2026, 19:37 UTC

47 observations across Vitals (36), Security (10), Performance (1)

Source: OBSERVATIONS_AND_ACTIONS*.csv

A point-in-time companion to the live, filterable online guide. Content is added over time — for the current version, refer to the online guide rather than this snapshot. This document describes how the Concertim team begins investigating each observation; it is a reference, not a prescription of what anyone must do.

About this guide

This is an open, living reference describing how the Concertim team begins investigating HPC health observations. It's shared so customers and peers can compare our practices against their own — **not as a prescription** of what anyone must do. Each entry captures an observation, what it can mean for the service, and where we'd start looking into it.

> **Why no priority or severity labels?** Because what's urgent depends entirely on the specific customer, site, and the assets an observation touches. A "critical" filesystem fill on one cluster may be routine headroom on another. We've deliberately left severity out of this guide to avoid implying a fixed ranking — work with us to triage against your environment.

Observations are grouped into topics: **Vitals** (essential health indicators), **Security**, and **Performance**. Entries still being written are badged *Draft* — they're shown so the gaps are visible, not hidden.

WHAT'S INSIDE

Vitals	36
Security	10
Performance	1

Vitals 36 observations

Unreachable

Host

This indicates hosts which we were unable to contact with our checks. This could be due to the host being off, network issues or other faults.

POTENTIAL SERVICE IMPACT

When a node is unavailable/unreachable it is likely not functioning as part of the resource pool and therefore reducing the availability of the HPC Environment to end-users. Depending on the systems that are unreachable this could impact more than workload resource availability.

WHERE WE'D START LOOKING

1. For each of these unreachable systems we would investigate the cause for them being down
2. Attempt to bring them back up
3. Investigate the systems for any issues that may have caused them to be down
4. Perform burn testing and benchmarks to ensure system stability
5. In the event that issues are reproduced we would follow our hardware troubleshooting methodology *(Appendix A: Concertim Hardware Troubleshooting Methodology)* to narrow down the cause, arrange replacement and bring the system back into service

Uncorrectable Memory Error (GPU)

Host

A GPU has reported a number of uncorrectable memory errors, these can occasionally occur sporadically or be indicative of early hardware failure.

POTENTIAL SERVICE IMPACT

The GPU may perform incorrect calculations during workloads and potentially end up falling out of service.

WHERE WE'D START LOOKING

1. Identify the gpu with the uncorrectable fault
2. Make sure gpu server is taken off job scheduler
3. Perform power-cycle on server
4. Investigate the system for any issues that may have caused the uncorrectable memory errors
5. Perform GPU benchmarks to ensure the gpu has recovered properly
6. In the event that issues are reproduced we would follow our hardware troubleshooting methodology *(Appendix A: Concertim Hardware Troubleshooting Methodology)* to narrow down the cause, arrange replacement and bring the system back into service

GPU Reset Report

Host

A GPU has reported it needs to be reset in order to recover from an error state or restore responsiveness

POTENTIAL SERVICE IMPACT

The GPU may be unavailable for user workloads and could negatively impact performance of workloads due to the degraded availability

WHERE WE'D START LOOKING

1. Identify the gpu which requires reset
2. Make sure gpu server is taken off Job Scheduler
3. Perform power-cycle
4. Review gpu was reset successfully
5. Resume back on job scheduler

Multipath Errors

Storage Array

A storage array multipath error typically occurs when one or more paths to the storage device become unavailable. These can occasionally happen sporadically, or can be indicative of a hardware failure in the storage system.

POTENTIAL SERVICE IMPACT

- Loss of redunancy in storage system - Potential for loss of access should further failures occur

WHERE WE'D START LOOKING

1. Identify the systems and multipaths affected.
2. Establish if the number of available paths has recovered redundancy for each affected path.
3. In the event that the number of paths has not recovered, investigate the systems and storage arrays for hardware failures.
4. In the event that the number of paths has recovered, no further action is required.

High Memory Utilisation

Host

A host has high memory utilisation (above 85%)

POTENTIAL SERVICE IMPACT

The workload(s) on the affected hosts may be pushing the resource to its limits. In the event that memory peaks further than current usage it could lead to workload slowdown or failure.

WHERE WE'D START LOOKING

1. Review current memory usage on the server
2. Identify the source of the high memory
3. If related to a current job running on the scheduler to monitor and make sure usage has dropped after completion
4. If this is due to leftover of a historic job to make sure this is tidied up correctly

High Filesystem Utilisation

Filesystem

Storage utilisation is above 85%

POTENTIAL SERVICE IMPACT

Workloads may be affected if storage utilisation continues to increase, and performance may begin to degrade on certain filesystem types.

WHERE WE'D START LOOKING

For the increase of filesystem usage we would provide clarification on the following points:

1. Confirm the filesystems current utilisation
2. Generate report of usage of filesystem
3. Inform customer of the current utilisation
4. Take measures to inform relevant end-users to be conscious of data storage and to reduce usage if possible

Critical Filesystem Utilisation

Filesystem

Storage utilisation is above 95%

POTENTIAL SERVICE IMPACT

The filesystem is close to 100% utilisation which may impact workloads should this continue to increase. Performance is likely to be degraded on certain filesystems, and workloads may begin to fail if utilisation reaches capacity.

WHERE WE'D START LOOKING

For the increase of filesystem usage which would now be set to critical usage we would provide further information on:

1. Confirm the filesystems current utilisation is at critical utilisation
2. Generate report of usage of filesystem
3. Inform customer of the critical utilisation and the effects this can have on their cluster such as user access or performance on workloads

Missing Filesystem Mounts

Host

The affected hosts are missing defined filesystem mounts

POTENTIAL SERVICE IMPACT

Depending on the missing mounts and the type of hosts affected this could render the host unusable within the service as critical data needed for workloads could be unavailable.

WHERE WE'D START LOOKING

1. Review all currently mounted and active filesystems
2. Attempt to mount missing filesystem manually
3. Review hardware to make sure their is no fault to cause/continue the missing filesystem

High Load

Host

Load average per CPU is above 1.5

POTENTIAL SERVICE IMPACT

High load can be indicative of various workload processes running as expected (and can be affected by I/O and other factors). If load is exceedingly high it could lead to a host becoming slow or completely unresponsive, or be indicative of a misconfigured workload or a workload utilising more resources than allocated.

WHERE WE'D START LOOKING

1. Review current CPU usage on the server
2. Identify the source of the high CPU load
3. If related to a current job running on the scheduler to monitor and make sure usage has dropped after completion
4. If this is due to leftover of a historic job to make sure this is tidied up correctly

Unable to Collect Hardware Storage Data

Host

Due to missing software dependencies, checks of the storage hardware could not be completed.

POTENTIAL SERVICE IMPACT

State of storage hardware on these hosts is potentially unknown which means that storage hardware failures could take these hosts out of service without forewarning, potentially impacting user workloads or service access.

WHERE WE'D START LOOKING

1. Install MegaCLI on the relevant host. This can be obtained from the Broadcom support section online.

Infiniband Devices Inactive

Host

Infiniband devices are not in the ACTIVE state

POTENTIAL SERVICE IMPACT

If the identified Infiniband devices are configured for use, then the infiniband will likely not be functional for parallel workloads or high speed connections to networked storage (if HPC environment configured for that) or have degraded performance.

WHERE WE'D START LOOKING

1. Make sure server is taken off the job scheduler
2. Power off the server before powering on the server to see if link is re-established
3. If link is still reporting to be down, attempt to reseat cable
4. If issue still hasn't been resolved attempt hardware troubleshooting steps of swapping cables and network ports
5. Raise support call with hardware vendor to replace affected part

Down/Drained on Scheduler

Host

Hosts are not available for running workloads

POTENTIAL SERVICE IMPACT

When a node is down/drained it is not functioning as part of the resource pool and therefore reducing the availability of the HPC Environment to end-users.

WHERE WE'D START LOOKING

1. Identify issue as to why server is removed from the scheduler
2. Perform power-cycle on server
3. Investigate the system for any issues that may have caused the server to be down/drained on the scheduler (if none then to be resumed back onto the scheduler)
4. In the event that issues are reproduced we would follow our hardware troubleshooting methodology *(Appendix A: Concertim Hardware Troubleshooting Methodology)* to narrow down the cause, arrange replacement and bring the system back into service

Physical Disk Not OK

Storage Array

The affected arrays are reporting that disks are not okay.

POTENTIAL SERVICE IMPACT

Depending on the configuration and state of the array this may not impact storage performance, however it is likely there is reduced redundancy and further failures may lead to either performance degredation or data loss.

WHERE WE'D START LOOKING

1. Identify the disk reporting an issue
2. Gather support logs and raise a support call with the hardware vendor
3. Schedule the replacement of the faulty disk
4. Monitor copyback of the drive after the replacement

Overall System Health Not OK

Storage Array

The affected arrays are reporting a system health issue.

POTENTIAL SERVICE IMPACT

Depending on the failure, the array may continue to operate but in a degraded state, and further failures may lead to either performance degradation or data loss.

WHERE WE'D START LOOKING

1. Identify the specific issue for the array to be reporting degraded system health
2. Collect support logs and raise a support call with the hardware vendor
3. Depending on the failure this may require a complete outage of the cluster to conduct hardware call
4. Once the scheduled hardware call is completed to review the overall health of the storage array and review any data loss

Unable to Collect Enclosure Data

Storage Array

This issue can occur when the management interface of the disk array fails to return data to queries.

POTENTIAL SERVICE IMPACT

The state of the array was unable to be determined via one of the redundant controllers. Should the other controller also stop responding, any hardware failures may unknowingly cause degredation to performance or data loss.

WHERE WE'D START LOOKING

1. Review both management controllers to make sure both are pingable and host access is reachable
2. Make sure both of these controllers have SNMP enabled and can communicate via SNMP
3. May require the affected controller or both controllers to be reset
4. If issue is still persistent after these steps then it will require a support call raised with the hardware vendor
5. Depending on the failure this may require a complete outage of the cluster to conduct the hardware call
6. Once the schedule hardware call is completed to then review the overall health of the array and review any potential data loss

Unable to Collect Management Interface Data

Host

This issue can occur when the management interface of the system fails to return data to queries.

POTENTIAL SERVICE IMPACT

The state of the system hardware is unknown, which means hardware failures may then unknowingly cause performance degredation or take the host out of service without forewarning.

WHERE WE'D START LOOKING

1. Check to see if you are able to ping/access the management IP address or gui
2. If unable to reach IP or gui to reseat the management cable from host end/switch port end
3. If able to reach IP address/gui to make sure all SNMP settings are enabled
4. Review overall hardware health on the host server to make sure nothing could be affecting the data gathering
5. Test to make sure SNMP is now communicating as expected

Overall System Health Not OK

Host

The affected systems are reporting a system health issue.

POTENTIAL SERVICE IMPACT

Depending on the issue, the performance of the system may be degraded or the system may be entirely unavailable, impacting availablity of HPC resource to end users.

WHERE WE'D START LOOKING

1. Identify the specific issue on the host to be reporting degraded system health
2. If required to make sure the server is taken off the job scheduler
3. Conduct a power-cycle on the server to see if that updates the system health
4. Collect support logs and raise a support call with the hardware vendor
5. Depending on the failure this may require a complete outage of the cluster to conduct hardware call
6. Once the schedule hardware call is completed to review the overall health of the host server before returning back into service

Outdated Hardware

Host

This issue relates to hardware in use which is over 5 years old

POTENTIAL SERVICE IMPACT

Outdated hardware can have a number of impacts on the service. It may perform tasks a lot slower than newer hardware, lack hardware-level optimisations for workload execution or be more prone to instantanious component failure. Warranty for outdated hardware can be difficult to obtain, as can replacement parts in general. Additionally it may cost more in power to execute a workload than more modern kit.

WHERE WE'D START LOOKING

To minimise the service impact of outdated hardware we would:

1. Ensure that outdated hosts are flagged to end-users (e.g. separate queue on the scheduler)
2. Configure monitoring solutions to consistently check for indications of hardware issues
3. Seek warranty for the affected hosts Beyond establishing risk management for the outdated hardware we would also seek upgrades with newer kit coming in to retire the existing ones.

Unknown Warranty Status

Host

We are unable to determine the warranty status of these hosts

POTENTIAL SERVICE IMPACT

Warranty for hardware ensures an agreement to replace faulty components. Being unable to establish the status of warranty means it is unknown whether faults with this hardware can be handled through an existing agreement and may lead to affected hosts being out of service for a long time as warranty status is renewed.

WHERE WE'D START LOOKING

At Concertim we would:

1. Identify the warranty status of these hosts (this could be via supplier's website or through purchase agreement details)
2. Track expiry for all hosts
3. Trigger extesion of warranty 3 months before expiry (to ensure time to allocate funds and purchase renewal)

Host Powered Off

Host

The system is not reporting as powered on.

POTENTIAL SERVICE IMPACT

When a host is not reporting as powered on, this can be indicative of a critical hardware or power failure unless it is an expected condition. Depending on the affected host, this could reduce available compute or render critical parts of the cluster unavailable.

WHERE WE'D START LOOKING

1. Attempt to power on the server to see if it boots
2. If boots successfully then will review the overall health of the server before returning to service
3. If it fails to boot, identify the specific issue on the host causing it to not successfully boot through the management interface
4. Collect support logs and raise a support call with the hardware vendor
5. Depending on the failure this may require a complete outage of the cluster to conduct hardware call
6. Once the schedule hardware call is completed to review the overall health of the host server before returning back into service

Reported GPUs Changed

Host

When a host reports the number of GPUs has changed, this typically indicates one or more of the expected GPUs is no longer recognised by the system.

POTENTIAL SERVICE IMPACT

The missing GPU(s) will be unavailable to workloads scheduled to the host. This is likely to affect performance and also potentially cause job failures.

WHERE WE'D START LOOKING

1. Review the host to identify which GPU is no longer being recognised
2. Make sure the gpu host is removed off the job scheduler
3. Attempt power-cycle on the server to see if this corrects the numbers of GPU's
4. Perform burn testing and benchmarks to ensure system stability
5. If the issue is still persistent contact your gpu hardware vendor and provide logs of the affected gpu's
6. Once issue is resolved with vendor perform benchmark testing to ensure system stability

Expected Service Not Running

Host

A service which is expected to be running on this host is not reported as running. This could be due to a service failure, or be indicative of a larger issue.

POTENTIAL SERVICE IMPACT

Depending on the affected host and service, this may affect user workloads, compute availability or overall HPC cluster availability.

WHERE WE'D START LOOKING

1. If required make sure the hosted server is off the job scheduler
2. Restart down service
3. Review server to make sure it is now running as expected
4. Resume host back onto the job scheduler

Excessive Number of Zombie Processes

Host

An excessive number of zombie processes are present on the node. This can be indicative of a failed workload, system service or other processes have not exited cleanly.

POTENTIAL SERVICE IMPACT

Impact on workloads is likely to be minimal, however routine housekeeping is important in maintaining system health, preventing exhaustion of process tables which may affect future workloads and assist in early identification of problems.

WHERE WE'D START LOOKING

1. Review the current zombie processes being reported on the server
2. If these are related to a current workload on the server, continue to monitor and see if these processes tidy up once the workload is completed
3. If these are defunct process of a previous workload then these would be cleared.
4. If unable to clear processes this will require the server to be power-cycle
5. If there are a very high number of these processes found on a login host then we would schedule a power-cycle with the primary contact of the cluster to clear these Zombie processes

System Time Not Synced

Host

The system time is not reporting correctly synced to a central NTP server.

POTENTIAL SERVICE IMPACT

Unsynchronised system time across nodes can result in job failures, scheduling inconsistencies and complicate troubleshooting due to inconsistent logging.

WHERE WE'D START LOOKING

1. Review system configuration to ensure it matches the expected configuration.
2. If no NTP server is configured, configure a suitable NTP server.
3. Ensure the system is able to reach the specified NTP server.
4. If the system is unable to reach the NTP server, continue further diagnosis to determine the issue.
5. Once resolved, attempt to sync time to the NTP server.
6. Verify time has been synced correctly.

PSU Not Reporting OK

Host

The system is reporting a PSU issue. This can be indicative of an upstream power failure or a hardware failure in the system.

POTENTIAL SERVICE IMPACT

Immediate impact on most systems is likely to be low if equipped with redundant power, however certain systems or environments may be power limited when not functioning with all expected PSUs which may affect workload or cluster performance. Further failures could also lead to unexpected loss of systems or the HPC cluster as a whole.

WHERE WE'D START LOOKING

1. Identify which PSU is reporting the issue
2. If required make sure the server is taken off the job scheduler
3. Raise support call with hardware vendor
4. These may require further troubleshooting steps such as swapping PSU cables which the hardware vendor will advise on support ticket
5. Arrange hardware replacement
6. Review overall health of the host and bring the system back into service

Fan Not Reporting OK

Host

The system is reporting a fan issue. This is typically indicative of a hardware failure, however these can also occur sparodically under very high load conditions.

POTENTIAL SERVICE IMPACT

The system may be unable to continue to function or may function with degraded performance. Depending on the system, this may reduce available HPC compute or affect the wider HPC cluster service.

WHERE WE'D START LOOKING

1. Identify which fan is reporting the issue
2. If required make sure the server is taken off the job scheduler
3. Raise support call with hardware vendor
4. These may require further troubleshooting steps such as swapping fans which the hardware vendor will advise on support ticket
5. Arrange hardware replacement
6. Review overall health of the host and bring the system back into service

VNC session older than 30d.

User

This VNC process has been running for at least 30 days, and may no longer be in use or be required by the user.

POTENTIAL SERVICE IMPACT

Too many unused VNC sessions can lead to resource exhaustion on your cluster login nodes, which may lead to degraded user experience or even affect wider cluster access.

WHERE WE'D START LOOKING

1. Identify which VNC's sessions are over 30 days old
2. Review to see if the user has connected into the VNC or is still running a process actively
3. Review if these stale VNC's are using unnecessary resources for the login node or could have any potential security risks for being left open for so long
4. Possible termination of the VNC session

Recent HA warnings in log messages.

Host

The host has recently logged HA warning messages.

POTENTIAL SERVICE IMPACT

Repeated warning logs from HA services can often be an indicator for wider issues with the HA service.

WHERE WE'D START LOOKING

1. Review the current state of the HA services and ensure everything is reporting ok.
2. Review the log messages reported. Typically these are due to monitor time outs caused by a temporary high load on the server.
3. If the load has reduced and the service is operating as expected, clear the alert from the HA service.
4. Further investigation may be required if the time out has resulted in a service failure, the load has not subsequently reduced or the issue reoccurs frequently.

User/group quotas are not enabled.

Filesystem

The system is not reporting user or group quotas as enabled for this filesystem.

POTENTIAL SERVICE IMPACT

Storage usage quotas will not be enforced for this filesystem, so users may be able to consume disproportionate storage, leading to disk exhaustion or degraded performance in extreme cases.

WHERE WE'D START LOOKING

1. Is this a filesystem quotas should be enabled on?
2. If so, run a quota check and re-enable the quotas on the filesystem export
3. If not, ensure usage is routinely reviewed to prevent filesystem exhaustion

User is above their set soft file/storage quota.

User

This user has gone over the allocated storage quota for this filesystem and is currently in a grace period.

POTENTIAL SERVICE IMPACT

Once the grace period has expired, the user will no longer be able to read or write data.

WHERE WE'D START LOOKING

1. Identify the user which has exceeded their soft file/storage quota
2. Review current quota utilisation
3. Notify the user for being in their grace period

Users grace period has expired for their soft file/storage limit.

User

This user has gone over the allocated storage quota for this filesystem and their grace period has now expired.

POTENTIAL SERVICE IMPACT

The user will no longer be able to write data to this filesystem, and may have trouble accessing the HPC cluster.

WHERE WE'D START LOOKING

1. Identify the user which has exceeded their soft file/storage quota
2. Review current quota utilisation
3. Possible further notifications to advise user can no longer write data to this filesystems and have potentially failures in any jobs currently running

User is above their set hard file/storage quota.

User

This user has gone over their allocated storage quota for this filesystem by a substantial amount, and has therefore been prevented from writing any further data to the filesystem.

POTENTIAL SERVICE IMPACT

The user will no longer be able to write data to this filesystem, and may have trouble accessing the HPC cluster.

WHERE WE'D START LOOKING

1. Identify the user which has exceeded their soft file/storage quota
2. Review current quota utilisation
3. Possible further notifications to advise user can no longer write data to this filesystems and have potentially failures in any jobs currently running

HA Corosync service is not reporting ok.

Host

The HA services on this node are reporting an issue.

POTENTIAL SERVICE IMPACT

Depending on the affected host and service, this may affect user workloads, compute availability or overall HPC cluster availability.

WHERE WE'D START LOOKING

1. Review the current state of the HA services.
2. If necessary, fail over any services that have not automatically failed over to restore cluster service.
3. Identify the underlying cause of the failure, typically this is due a partial or total host failure.
4. Once the underlying cause has been resolved and all testing has been completed on the failed host, arrange a suitable period to return it to service and fail services back over to restore high availability.

Unable to collect Linux health data

Host

This issue can occur when remote access to the host is not possible either due to network issues or access requirements not being satisfied

POTENTIAL SERVICE IMPACT

The state of the Linux operating system on these hosts is potentially unknown which means that errors or misconfigurations could take these hosts out of service without forewarning, potentially impacting user workloads or service access.

WHERE WE'D START LOOKING

1. Check host is up
2. Review network and access configuration to ensure requirements met
3. Attempt to reach the host via SSH and other protocols

Uncorrectable Memory Error (DIMM)

Host

A host has reported a number of uncorrectable memory errors, these can occasionally occur sporadically or be indicative of early hardware failure.

POTENTIAL SERVICE IMPACT

The host may perform incorrect calculations during workloads and potentially end up falling out of service.

WHERE WE'D START LOOKING

1. Identify the DIMM with the uncorrectable fault
2. Make sure host is taken off job scheduler
3. Perform power-cycle on server
4. Investigate the system for any issues that may have caused the uncorrectable memory errors
5. Perform memory benchmarks to ensure the host has recovered properly
6. In the event that issues are reproduced we would follow our hardware troubleshooting methodology *(Appendix A: Concertim Hardware Troubleshooting Methodology)* to narrow down the cause, arrange replacement and bring the system back into service

Unscheduled Processes

Host

End-user processes have been identified on the affected hosts which are running outside of the scheduler

POTENTIAL SERVICE IMPACT

This can have a negative impact on the environment's performance for other users by reducing available resources.

WHERE WE'D START LOOKING

1. Check if the server is currently idle/free/available or is running a job on the scheduler
2. If the server is idle/free/available then the process will be terminated to free up any used resource.
3. Review the job that is running on the scheduler and check if the unscheduled process is related to the current running job.
4. If it is related such as the same user id or to the current jobid then continue to monitor until the job is completed to see if the unscheduled processes have cleared.
5. However if the unscheduled process is not related to the current running job then the unscheduled process will be terminated to free up any available resources.

Security 10 observations

Enable Fail2Ban on Access Hosts

Environment Perimeter | Draft

Fail2Ban improves security by blocking IP addresses with repeated authentication failure attempts. Enabling it on access hosts provides an additional threat deterrent.

POTENTIAL SERVICE IMPACT

Without fail2ban, anyone with access to an access host can repeatedly attempt to authenticate, opening the door to potential brute force attacks.

WHERE WE'D START LOOKING

No steps documented yet.

Review Exposed Services on External Network

Environment Perimeter | Draft

Over time it is possible for services to be exposed to external networks that no longer are required.

Authoring note:

- SOME NOTE ABOUT WHICH NETWORKS - MAKE IT SPECIFIC TO THEM

POTENTIAL SERVICE IMPACT

Reviewing services exposed to external networks provides the opportunity to update and reclarify access requirements, ensuring that no unnecessary access pathways are open.

WHERE WE'D START LOOKING

No steps documented yet.

Tighten the scope of administrative permissions

Privileged Access | Draft

It is common practice for administrative users to have privileged access to run any commands in a system. For the purpose of auditing and hardening it can be beneficial to limit the commands particular administrative users have privileged access to.

POTENTIAL SERVICE IMPACT

- administrators could perform destructive actions

WHERE WE'D START LOOKING

No steps documented yet.

Audit the commands and movements of administrative accounts

Privileged Access | Draft

Not yet documented.

POTENTIAL SERVICE IMPACT

Not yet documented.

WHERE WE'D START LOOKING

No steps documented yet.

Periodically review administrative user accounts

Privileged Access | Draft

Not yet documented.

POTENTIAL SERVICE IMPACT

Not yet documented.

WHERE WE'D START LOOKING

No steps documented yet.

Periodically review administrative user permissions

Privileged Access | Draft

Not yet documented.

POTENTIAL SERVICE IMPACT

Not yet documented.

WHERE WE'D START LOOKING

No steps documented yet.

User storage locations external to environment

Storage | Draft

Not yet documented.

POTENTIAL SERVICE IMPACT

This is not inherently insecure but could pose a security risk as sensitive user information for the HPC Environment is exposed more than it may need to be.

WHERE WE'D START LOOKING

No steps documented yet.

Only mount external storage to access hosts to enable users to copy workload information to internal storage

Storage | Draft

Not yet documented.

POTENTIAL SERVICE IMPACT

Not yet documented.

WHERE WE'D START LOOKING

No steps documented yet.

Minor CVE Patching

Software and Applications | Draft

These CVEs are of mild concern and are worth repairing at the next maintenance opportunity

POTENTIAL SERVICE IMPACT

Not yet documented.

WHERE WE'D START LOOKING

No steps documented yet.

Major CVE Patching

Software and Applications | Draft

These CVEs are of high concern and are recommended to be repaired as soon as possible

POTENTIAL SERVICE IMPACT

Not yet documented.

WHERE WE'D START LOOKING

No steps documented yet.

Performance 1 observation

Memory performance reductions across affected Blocks

Individual Integrity

May indicate performance degradation

POTENTIAL SERVICE IMPACT

Bad performance = worse throughput

WHERE WE'D START LOOKING

Review the results Run further investigations (longer, more intense checks) If problem persists - hardware replacement?

No steps documented yet.