

# EQUALLOGIC

## PS Series Firmware Version V5.0.2

### Issues Corrected in This Version

---

Issues that have been corrected in PS Series Firmware Version 5.0.2 are described below.

- On PS6000, PS6010, PS6500 and PS6510 systems, the CPU fan may be intermittently incorrectly reported as failed when it is operating properly.
- Cancelling a replication on a volume may cause a member in the primary group to restart continuously.
- Properly functioning SATA drives may be incorrectly marked as failed.
- A lack of internal resources may cause a controller failover or restart.
- An update to V5.0.0 or V5.0.1 may cause a temperature sensor on PS6000, PS6010, PS6500 and PS6510 arrays to be incorrectly marked as failed.
- Updating the firmware to V5.0.0 or V5.0.1 may cause volumes with more than 4 ACLs to respond slowly to iSCSI initiator login requests. This may cause logins to fail as the number of ACLs increases.
- After an update to 5.0.0 or 5.0.1, arrays may respond slowly and virtual machines running on vSphere 4.1 may experience timeouts.
- Disabling replication, followed by enabling replication does not restart replication as expected.
- After an update to V5.0.0 or V5.0.1, arrays may go offline.
- After upgrading the firmware to V5.0.0 or V5.0.1, volumes distributed across multiple arrays, moved between pools, or involved in member-move operations may not be replicated correctly even if the user interface indicates that the replica completed successfully.
- Immediately following an update to V5.0.0 or V5.0.1, some volumes on PS5500, PS6500 or PS6510 arrays may not come online.
- During a firmware update, the primary controller may be unable to failover to complete the update process.
- An array may require assistance from customer support to be brought back online after a failover or power outage.
- Firmware updates may time out when updating the passive controller on Type 1 arrays.
- iSCSI Connections may be redirected to Ethernet ports without valid network links.

### Issues Corrected in Version 5.0.1

---

Issues that have been corrected in PS Series Firmware Version 5.0.1 are described below.

- In some circumstances, pressing the “restart” option in the GUI update utility would not restart the array due to the utility presenting the restart option before a restart could be performed. Pressing restart would work properly once the controller was actually ready to restart.
- PS Series arrays models PS50 – PS5000 running V5.0.0 Firmware may experience a controller restart if the network transmits Ethernet packets containing VLAN tags to the array.
- In some circumstances, replication over low bandwidth network links is unable to complete successfully.
- Occasionally, an upgrade to V5.0.0 firmware would cause the eth1 port on a PS6010 or PS6510 to stop functioning after the upgrade.

**NOTE:** If you have an array with a malfunctioning Ethernet port, updating the PS Series Firmware to version 5.0.1 will not resolve the issue. If you are running version 5.0.0 firmware and experience this issue, contact your EqualLogic customer support provider.

## **Issues Corrected in Version 5.0**

---

Issues that have been corrected in PS Series Firmware Version 5.0 are described below.

**Note:** All group members must be running version 5.0 firmware for these fixes to take effect.

### **Management Interface**

- Many long volume names could cause the GUI to be unable to display all volumes properly.
- An invalid replica collection could prevent the GUI from displaying valid replica collections.
- Occasionally, the GUI would present an invalid warning about mixing disk types when adding a new member to a pool.
- Non-English characters may cause the GUI to react slowly.
- Sorting of entries in the replication information portion of the GUI has been corrected.
- Customers using Belgian or Dutch keyboards could not enter the “@” symbol into the GUI using the keyboard.
- Up and down arrows do not work in Date and Time Panel.
- The CLI and GUI use different units of measurement to report the size of a storage pool.

## Initiator Connections

- Double clicking on a volume ACL to modify it could cause an ACL element to be de-selected in the Modify ACL dialog box.
- The restriction on using a replication partner group name as a local CHAP user name has been removed.
- Following a controller failover, an internal process may not relinquish resources causing iSCSI connections to be unable to reconnect automatically.
- When configured to use IPv6 with the TCP timestamp option, initiators would be able to discover volumes, but would be unable to log in to them.
- In some cases, additional MPIO connections initiated by the MPIO DSM were unable to login automatically.
- Repeated loss of link by an array Ethernet port could cause an array to become unresponsive to Ethernet traffic.

## Replication

- Cancelling a replication may leave fewer replicas available at the target site for possible recovery use than defined in the replica max-keep policy.
- Delegated space usage for the target replication group was incorrectly calculated on the source replication group. This could result in a negative amount of delegated space being reported on the source group.
- After changing replication partners for a volume or volume collection, and performing replication to the new partner, deleting old replica sets for the old partner removes replica sets for both partners.
- After changing replication target groups for a volume and performing the initial full volume replication to the new target group, deleting the replica from the original replication target group causes a second full replication to new target replication group.
- In rare circumstances, attempting to cancel a replication that was already in the process of being canceled could cause an internal process to restart.
- Cancelling movement of delegated space from one pool to another while the movement was in progress might not result in the operation being cancelled.
- A volume with 16 ACLs may be unable to fail back. This can mean the replication may have to transfer the entire contents of the volume to the original group, instead of only the changes that occurred to the recovery volume.

## Snapshots

- In rare circumstances, a snapshot schedule for a deleted volume collection cannot be deleted.
- Snapshots may be unable to be created, generating an error message, while a volume is being moved to a new member as part of a load balancing or pool migration operation.

## Array Restart

- After an online RAID conversion, a PS5500, PS6500 or PS6510 requires a restart to display the correct new member capacity.
- Restarting an array during a load-balancing event may result in a “clear-lost-data” error message, which requires support assistance to resolve.
- In some circumstances, an erroneous message would appear in the event log during a controller startup indicating that a panic had occurred when the startup was actually a normal startup.
- Restarting an array with a paused replication could cause an internal process to be unable to restart properly, resulting in the array becoming unresponsive.

## Control Modules and Group Members

- Some controller component failures did not result in a controller failover.
- Updating the password for the group administrator account is not reflected on the secondary controller until a controller failover or array restart.
- In rare circumstances, an unexpected failover could occur while allocating space for a volume.
- In rare circumstances, an internal process may require a restart. This now results in a controller failover.
- PS6010 and PS6510 controllers using 5m or longer copper cables may occasionally bounce the network link, causing some 10Gb switches to disable the switch port.
- In a PS5500, PS6500, or PS6510, the secondary controller can generate a log event with a value of -1 for the ambient temperature, if the sensor is out of the normal operating range.
- In some cases, controller failovers and restarts could take longer than the recommended timeout settings.
- In rare circumstances, the following issues may occur with PS6000 and PS6500 arrays:
  - The secondary controller may unexpectedly restart.
  - A member of the group can be spontaneously removed, setting volumes offline.
  - A member of the group can lose connectivity to the rest of the group.
  - An Ethernet port on the array may become unresponsive, but not go down, causing volumes to disconnect.
  - An internal process may restart unexpectedly.
  - The array may unexpectedly restart.
  - In a multi-member group, a member may take longer than the timeout settings to fail over or restart.

## Firmware Updates

- On arrays with very low free space, firmware updates are not able to be copied to the array.
- After upgrading firmware, a group could report that the replication partner is down when it is not, causing replication to be unable to restart.

## Volumes

- Modifications made to a volume's properties within one second of the creation of the volume can result in the volume property modifications not being implemented.
- In rare circumstances, a volume could be offline, with the status "offline-member-down" when all group members are operating normally.
- Volumes in pools with extremely low free space may be unable to move to other pools.
- A thin-provisioned volume may not be able to allocate additional capacity while being moved, resulting in the volume going offline.
- When many volumes without ACLs are present in a Group, an internal process may restart due to a lack of resources.
- Restarting an internal process on multiple arrays in a large group may cause volumes to go offline.

## Storage Pools

- An internal process may restart continuously when attempting to add, remove or move members to or from a pool with low free space.

## Disks

- Normal disk error handling messages generated messages in the group event log. This now only generates an event when the frequency is at a level that is a concern for possible future disk failure.
- An unexpected increase in disk failures may occur on some EqualLogic chassis systems (pre-Dell) with 750GB drives (model PS400E) after upgrading to v4.3 firmware.
- Multiple drives could be incorrectly marked as "failed" simultaneously, requiring support to clear the condition and bring the array back on line.
- A failed spare drive was marked as failed, but did not generate a health condition. A failed spare now generates a health condition indicating that it should be replaced.

## Other Issues

- In rare circumstances, an array that is performing the management functions for the group will not relinquish the group IP address after a new member is elected to assume the management role.
- The Host Scripting Toolkit Perl module modifies the paging setting in the CLI but did not restore it to its previous setting. The module now restores the paging setting in the CLI after a script is run.
- When sending email to an administrator, the group would send a message with the email subject in the format of: "n events from group group\_name", which is difficult to sort. The format of the email subject has been changed to: "group\_name: n events" for better clarity and ease of sorting.
- Security of information in the support DIAG command report has been enhanced.
- MIBS have been enhanced to provide clearer system information.
- When using the EqualLogic MPIO DSM, all members of a Windows cluster are not able to log into cluster volumes.
- In rare circumstances, an SNMP community string could become corrupt, causing management utilities such as SANHQ to be unable to poll the group.

- In some cases, performance statistics could over count IOs when volumes were distributed on multiple members.
- In rare circumstances, an internal process may consume an excessive amount of system resources, causing the system to respond slowly.
- Microsoft clusters using SCSI-3 persistent reservations may be unable to pass the cluster validation test, or may be unable to properly negotiate file locks on CSVs.