

# RFE Questionnaire Results

Jez Tucker, Group Chair

Claire Robson, Group Secretary

## Foreword

The GPFS User Group was formed in the UK in November 2010 with the aims:

- Provide a 'localised' community for support and knowledge sharing
- Empower the GPFS community to use the product to its fullest potential
- Provide a centralised mechanism for development feedback
- Raise awareness of GPFS as a capable data management platform

Users meet to discuss relevant GPFS related topics and hear the latest product updates from key GPFS speakers. The User Group has over 200 members with representatives from a variety of industries including academia, media, defence, automotive, research and finance.

At the 10th meeting of the group in April 2014 users were asked to rate their their importance to a series of Request For Enhancement (RFE) questions generated from discussion points at previous meetings and through the GPFS User Group mailing list; from very unimportant to very important.

Of 92 attendees, 32 feedback forms were completed and returned.

Each attendee did not have to answer all questions, therefore some questions have a reduced number of responses.

The responses have been compiled anonymously and are presented in the following section.

In addition the RFE questions have been entered into the RFE portal for users to additionally vote on. RFE items were entered with the priority represented via the results returned. An IBM RFE URL follows each question.

This document has been shared with IBM GPFS Development as part of the aim of providing feedback and interaction between IBM and the Users.

On behalf of the User Group,

Jez Tucker & Claire Robson

## Compiled Results of RFE Questions


### Q1 Cache control via runtime policy engine

If caching could be controlled by the policy engine, the user could intelligently control that valuable resource with respect to their known workloads / data types. Example:

NO\_CACHE WHERE NAME LIKE %.bigfile (never cache any files of type bigfile)

PIN\_TO\_CACHE WHERE PATH\_NAME = '/path/to/my/data/%' (never expire this file from cache)


PRELOAD\_TO\_CACHE WHERE NAME = '/path/file.data, /path/file2' (load these files entirely into cache)



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56765](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56765)

### Q2 Ability to see what is in the cache


mmlscache gives directory, file, size/blocks, % used of cache, cache hits, last access time etc. -Y option



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56766](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56766)

### Q3 Ability to flush / clear the cache (not to disk)

The cache on a node should be able to be cleared without unmounting and remounting the file system




RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56767](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56767)

#### Q4 Globally accessible and coherent cache

A cached block is cached on all nodes and hence accessible from all nodes.


Globally coherent cache would give a performance advantage to multiple nodes performing the same task. If a file is cached on one node, then other client nodes reading the same file would not have to take a cache hit. Similarly, multiple blocks could be read in parallel from all nodes where the file is globally cached.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56769](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56769)

### Q5 mmfsck -P poolname

Checking a file system with multiple pools means results in a check of the entire metadata for the file system when each pool is checked. Not optimal. Check individual pools instead.




RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56769](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56769)

### Q6 Pool placement on path


Ability to place data into a pool based on the directory location of the data. Example:

RULE 'example' SET POOL 'fastpool' WHERE PATH\_NAME LIKE '/fastdata/'



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56769](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56769)


**Q7 Use the running policy engine to SetXattr for all files and folders under a PATH\_NAME**  
Ability to automatically tag the Xattr of data with defined attributes such as a project name



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56772](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56772)

**Q8 AFM queue control**

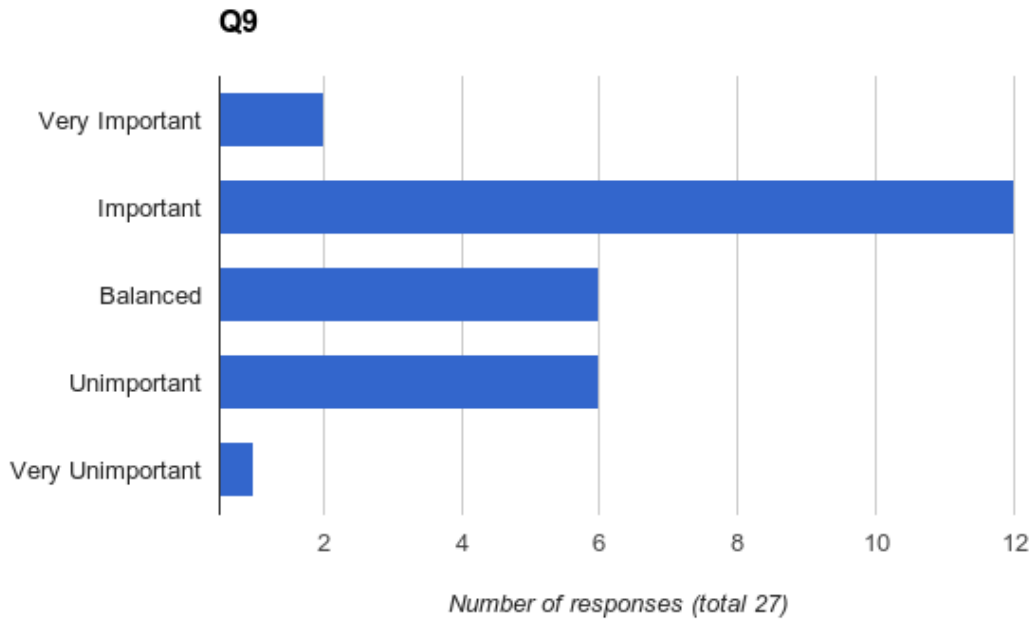
Complete control over the order of each prefetch (analogy: priority WAN transfer queue for GPFS)



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56772](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56772)

### Q9 SNMP 64 bit counters + CIM support


Proper 64 bit counters. Obtaining correct SNMP values (E.G. for pool sizes) is not as simple as it should be.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56774](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56774)

### Q10 Metablocks to require less context switching (I.E. increased performance)

At small (64k) block size there's too much overhead / load / cpu / context switches




RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56775](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56775)



**Q11 Multiple block sizes at the pool level.**


Pools could be configured with different block sizes (not metablocks) based on workload requirement.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56776](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56776)

**Q12 Treat all independent filesets identically to the global namespace operations**


E.G. Independent filesets do not raise callbacks on lowSpace or noSpace



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56779](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56779)

**Q13 Increase the supported number of independent and dependent filesets x 10**


1000 independent filesets does not offer enough granularity for AFM.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56779](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56779)

**Q14 Fast approximate df for filesets**

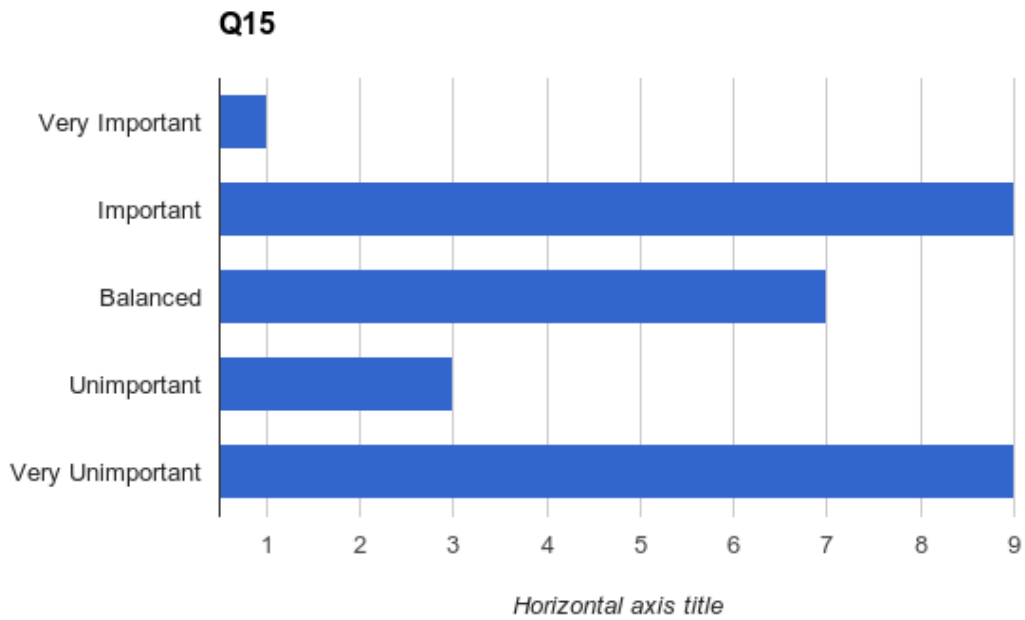
mmlsfileset df calculations are slow. A fast approximate filesetdf (alike mmlsquota) is needed.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56781](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56781)

**Q15 Support GPFS on Ubuntu 14.04 LTS and future LTS releases**


Ubuntu is now mainstream. Support Ubuntu LTS releases in addition to Debian.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56783](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56783)

**Q16 OSX 10.8+ client : GPFS protocol over IP**


Performance improvements over NFS3/4 and SMB2



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56786](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56786)

### Q17 NDMP support in native GPFS


Required for supporting fast block streaming to 3rd party (non-IBM) backup applications



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56786](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56786)

### Q18 Override the sort routine within policy processing


Ability to specify extra args to sort (E.G. for natural sort for numerical sequences) or to use a user specified program to perform the sort using a different algorithm. Sort is frequently the longest stage of execution.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56786](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56786)

**Q19 Ability to add and remove parallel threads / nodes to a currently executing policy**

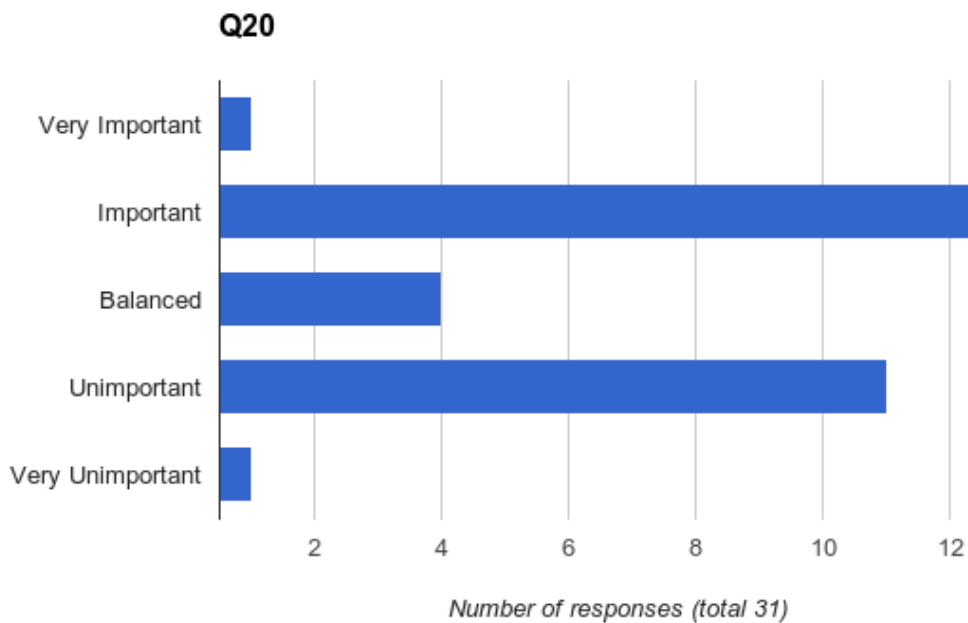
Allows the administrator to make decisions and changes to the performance impact of a running policy



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56786](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56786)

**Q20 Mirror snapshots to another cluster as a destination snapshot**


Ability to mirror a snapshot to a remote cluster. The source snapshot would appear as a snapshot on the remote cluster, not as live data.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56791](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56791)

**Q21 Option to mount the file system with the root being as a subdirectory**


NFS clients can mount /gpfs/mydata as /mydata. GPFS clients would mount /gpfs with a subdir of mydata. Ideally selected GPFS clients could mount /gpfs/mydata as /mydata without the use of symlinks.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56792](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56792)

**Q22 Make GPFS more atomic**

At present you cannot delete or create snapshots or filesets if a migration is occurring.



RFE URL: [http://www.ibm.com/developerworks/rfe/execute?use\\_case=viewRfe&CR\\_ID=56793](http://www.ibm.com/developerworks/rfe/execute?use_case=viewRfe&CR_ID=56793)

## Summary

Firstly, we extend our thanks to our members and the IBM GPFS team for the support and level of engagement at User Group 10.

This involvement enabled the User Group to collate data and personalised feedback from entirely relevant, cross industry parties at all levels of experience.

Users can use the RFE URLs in this document to vote for the RFEs covered - or submit additional RFEs, informing the User Group of their creation via the mailing list.

We welcome all constructive feedback and actively encourage Users and IBM alike to engage and debate at all possible levels for the furtherment of GPFS.

## Anonymised Raw Data

Sector	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22
N/A (anon)		3	1	3	3	1	2	5	1	2	5	5	5	2	5	5	5	5	2	5	5	2
Academia														4	5	5						
Academia	2	2	3	4	3	2	2	4	3	2	2	2	5	2	5	5	4	5	4	2	4	5
Research	2	3	3	4	2	5	2	3	2	2	4	4	4	2	3	4	3	2	2	2	2	3
Academia	3	2	2	3	3	2	3	2	2	3	2	3	4	3	2	2	3	4	3	2	4	5
Academia	4	2	1	3		4	5		2	5	5	5	5	2	5	5	5	3	2	5	5	3
Research	3	2	2	3	4	4	4	5	2	2	2	3	3	1	3	2	3	3	2	3	3	2
Storage Vendor	1	2	2	1	3	1	3	2	2	3	1	3	1	2	2	1	1	1	2	3	2	2
Solution Provider	1	1	3	3	3	1	5		1	2	1	2	3	3	5	5	3	3	2	2	3	3
Media	3	2	3	2	2	2	2	5	2	1	2	3	4	2	2	1	2	2	3	2	3	2
Media	2	2	2	3	2	1	3	2	4	4	3	3	3	1	5	3	1	3	2	2	3	3
Media	1	2	1	4		1	2	2	3	3	2	4	5		2	2	5				4	4
Solution Provider	2	2	4	4	4	2	2	2	2	4	2	4	4	4	2	2	5	3	4	4	2	4
Solution Provider	1	1	2	2	1	1	5	3	5	2	2	4	2	5	5	4						
IBM	2	1	1	2	1	1	1	3	3	3	1	3	3	1	4	2	4	2	2	4	4	2
Storage Vendor	1	3	2	4	2	1	1	3	4	3	1	2	4	2	2	2	3	3	2	4	3	2
Government	1	1	2	1	1	1	1	1	2	1	1	1	2	1	4	5	3	2	2	2	2	1
Research	4	5	3	3	5	4	3	5	3	3	4	4	3	4	3	2	2	4	5	4	4	5
Research	1	1	1	2	1	1	1	2	1	1	1	4	4	4	3	4	4	2	5	4	3	
Academia	2	1	1	3	2	2	2	2		2	2			2	4	4	2	4	3	2	3	2
Academia	4	2	5	4	1	3	4	3	2	1	1	4	1	3	2	4	3	4	2	2	4	1
Research	3	3	3	3		1	2		2	2	2		1	2	5	5	3	2	2	3	2	
Research	1	2	3	3	2	1	3	1	3	2	2		4	2	5	5	5	3	3	2	4	4
Finance	2	2	4	1						3	5	4	4	3	5	4	4	3	2	3	3	2
Research	2	1	2	3	3	2	3	1		1	2	4	4	1	5	1	3	2	3	4	1	3
Academia																						
Finance	4	3	3	3	3	2	3	2	4	2	2		3	2	5	3	5	3	2	2	3	2
Academia	2	2	4	4	4	4	2	2	4	2	3	3	3	1	2	3	4	4	4	4	4	4
Academia	2	2	3	2	2	2			2	2	2	3		2	3	3	3		2	4	3	
N/A (anon)	2	4	3	2	1				1	1					2	1			2	2		
Distributor	4	5	2	4	4	2	3	3	4	5	4	3	1	4	3	4	3	3	4	4	3	3
Academia	3	4	1	4	4	2	2	4	3	2	2	4	2	1	5	4	2	3	4	3	4	4
N/A (anon)	4	4	3	4	3	4	2	5	2	3	3	2	1	2	3	2	2	4	2	1	1	2
Academia	4	4	4	4	3	3		4					3	4			4			4	4	4