

Ganesha statistics overview

Existing stats gathering and extraction

With Ganesha 2.5 code, below performance stats were gathered and extracted.

Global stats (no option / global) – Global stats are enabled by default. It keeps track of how many requests are handled per protocol. Here protocols are NFSv3, NFSv4, NLM, MNT and RQUOTA. When extracted global stats, one get count for all protocols which got exercised by the clients. For holding the stats count, global_st structure is being used. Typical o/p for global stats is as below:

```
# ganesha_stats
Timestamp: Tue Dec 10 10:31:13 2019580388834 nsecs
Total NFSv3 ops: 10909
Total NFSv4.0 ops: 49638
Total NFSv4.1 ops: 0
Total NFSv4.2 ops: 0
```

Client listing (list_clients) – This option provides the list of NFS clients, which have accessed NFS server. Along with client IP address, the o/p also provides an indication whether a protocol specific stats are present for that client or not. “list_clients” provide below kind of o/p:

```
# ganesha_stats list_clients
Timestamp: Tue Dec 10 10:32:01 2019197071186 nsecs
Client List:
Address: ::ffff:10.0.100.1
  NFSv3 stats available: 1
  MNT stats available: 1
  NLM4 stats available: 1
  RQUOTA stats available: 0
  NFSv4.0 stats available 1
  NFSv4.1 stats available: 0
  NFSv4.2 stats available: 0
  9P stats available: 0
```

Delegations (deleg <client-ip-address>) – This option provides delegations related stats. I have not used/tested it.

Caching related (inode) – This option provides meta-data caching related statistics. A typical o/p is as below:

```
# ganesha_stats inode
Timestamp: Tue Dec 10 10:33:01 2019197072289 nsecs
Cache Requests: 0
Cache Hits: 107844
Cache Misses: 0
Cache Conflicts: 0
Cache Adds: 3134
Cache Mapping: 6946
```

NFSv3 I/O stats (iov3 [export-id]) – This option provides information on NFSv3 related IO stats like how much data is transferred (in bytes), how many such requests came in, related errors and latency. The output provides information on READ and WRITE requests. This information is for all exports by default, else it will be export specific. A typical o/p is as below:

```
# ganesha_stats iov3
EXPORT 32: Export does not have any NFSv3 activity
EXPORT 33: Export does not have any NFSv3 activity
EXPORT 34:
      requested transferred          total          errors
latency
READv3:  31686656  31457280          170           0  263847420          0
WRITEv3:  74424634  74424634          530           0  19056953192         0
EXPORT 0: Export does not have any NFSv3 activity
```

```
# ganesha_stats iov3 34
EXPORT 34:
      requested transferred          total          errors
latency
READv3:  31686656  31457280          170           0  263847420          0
WRITEv3:  74424634  74424634          530           0  19056953192         0
```

NFSv4 I/O stats (iov4 [export-id]) - This option provides information on NFSv4.0 related IO stats like how much data is transferred (in bytes), how many such requests came in, related errors and latency. The output provides information on READ and WRITE requests. This information is for all exports by default, else it will be export specific. A typical o/p is as below:

```
# ganesha_stats iov4
EXPORT 32: Export does not have any NFSv4.0 activity
EXPORT 33: Export does not have any NFSv4.0 activity
EXPORT 34:
      requested transferred          total          errors
latency
READv4:  31705943  31476567          182           0  328462795          0
WRITEv4:  84904207  84904207          1182          0  10477924722         0
EXPORT 0:
      requested transferred          total          errors
latency
READv4:           0           0           0           0           0           0
WRITEv4:          0           0           0           0           0           0
```

Export stats (export) – This option provides hint about whether protocol specific stats are available for an export or not. It provides information on all exports. A typical o/p is as below:

```
# ganesha_stats export
Timestamp: Tue Dec 10 10:43:01 2019167897842 nsecs
Export id: 32
  Path: /ibm/gpfs0/exp15
  NFSv3 stats available: 0
  NFSv4.0 stats available: 0
  NFSv4.1 stats available: 0
  NFSv4.2 stats available: 0
  MNT stats available: 0
  NLMv4 stats available: 0
  RQUOTA stats available: 0
  9p stats available: 0
```

Export id: 33
Path: /ibm/gpfs0/exp17
NFSv3 stats available: 0
NFSv4.0 stats available: 0
NFSv4.1 stats available: 0
NFSv4.2 stats available: 0
MNT stats available: 0
NLMv4 stats available: 0
RQUOTA stats available: 0
9p stats available: 0

Export id: 34
Path: /ibm/gpfs0
NFSv3 stats available: 1
NFSv4.0 stats available: 1
NFSv4.1 stats available: 0
NFSv4.2 stats available: 0
MNT stats available: 0
NLMv4 stats available: 1
RQUOTA stats available: 0
9p stats available: 0

Export id: 0
Path: /
NFSv3 stats available: 0
NFSv4.0 stats available: 1
NFSv4.1 stats available: 0
NFSv4.2 stats available: 0
MNT stats available: 0
NLMv4 stats available: 0
RQUOTA stats available: 0
9p stats available: 0

Total stats (total [export-id]) – This option provides operation count for NFSv3, NFSv4.0, NFSv4.1 and NFSv4.2 protocols. If an export ID provided then information for that export only is provided. A typical o/p is as below:

```
# ganesha_stats total
Timestamp: Tue Dec 10 10:43:28 2019167897842 nsecs
Export id: 32
  NFSv3: 0
  NFSv40: 0
  NFSv41: 0
  NFSv42: 0
Export id: 33
  NFSv3: 0
  NFSv40: 0
  NFSv41: 0
  NFSv42: 0
Export id: 34
  NFSv3: 10559
  NFSv40: 64229
  NFSv41: 0
  NFSv42: 0
Export id: 0
  NFSv3: 0
```

```

    NFSv40: 40
    NFSv41: 0
    NFSv42: 0
# ganesha_stats total 34
Timestamp: Tue Dec 10 10:43:31 2019167897842 nsecs
Export id: 34
    NFSv3: 10559
    NFSv40: 64229
    NFSv41: 0
    NFSv42: 0

```

Fast stats (fast) – This option provides count for different ops in the protocol. The o/p will list out all the operations served by Ganesha in that protocol. A typical o/p is as below:

```

# ganesha_stats fast
Timestamp: Tue Dec 10 10:43:51 2019167897842 nsecs
Global ops:

```

```

NFSv3:
NULL                :           1
GETATTR             :          1433
SETATTR             :          1195
LOOKUP              :          1911
ACCESS              :           451
READ                :            85
WRITE               :           265
CREATE              :           867
MKDIR               :           175
SYMLINK             :           250
REMOVE              :          1365
RMDIR              :           174
RENAME              :           358
LINK                :           250
REaddirPLUS        :           625
FSSTAT              :          1500
FSINFO              :             2
PATHCONF            :             1
COMMIT              :             2

```

```

NFSv4:
ACCESS              :          1332
CLOSE               :          1069
COMMIT              :             2
CREATE              :           425
GETATTR            :          9632
GETFH               :          3441
LINK                :           250
LOCK                :          2052
LOCKT               :             50
LOCKU               :          2056
LOOKUP              :          2862
OPEN                :          1069
OPEN_CONFIRM        :             1
PUTFH              :         17988
PUTROOTFH           :             1
READ                :             91
REaddir            :           621
REMOVE              :          1566

```

```

RENAME          :          376
RESTOREFH       :          250
SAVEFH          :          626
SETATTR         :         1237
SETCLIENTID     :           1
SETCLIENTID_CONFIRM :         1
WRITE           :          591
RELEASE_LOCKOWNER :        2048

NLM:
TEST            :           8
LOCK            :           1

MNT:
NULL            :           2
MNT             :           1

```

pNFS stats (pnfs [export-id]) – This option provides pNFS related stats. To be honest never used / explored by me.

Additions done by IBM (Sachin)

The Ganesha code can be divided into 3 layers: RPC layer (ntirpc), the server code (handling of protocol specific tasks, and FSAL specific code. Any performance issue can be at any of these layers. But most of the time its due to underneath file-system latency or there is contention in server code.

To detect this, firstly FSAL specific stats were added. This feature helps in getting performance statistics when FSAL specific code starts its execution. Currently it is fully available for GPFS, and for other FSAL basic infrastructure is available. As this functionality requires lot of book keeping, it is not enabled by default. Hence a functionality has been added to enable it dynamically and disable it. In addition one can reset these stats counter. Later the functionality to enable/disable has extended to other stats counting as well and reset action works for all stats.

FSAL specific stats (fsal <fsal_name>) – This option provides details on what all FSAL specific operations are invoked and number of those calls and response time (average, min & max). This option can be enabled via config option enable_FSALSTATS in NFS_Core_Param section. A typical o/p is as below:

```

# ganesha_stats fsal gpfs
FSAL stats for - GPFS
Stats collected since: Tue Dec 10 12:53:59 201992522344 nsecs
Duration: 155.4605736732 seconds
FSAL Stats (response time in milliseconds):
  Op-Name          Total      Res:Avg      Min      Max
NAME_TO_HANDLE    47709      0.003126    0.001434  0.151229
OPEN_BY_HANDLE    52473      0.008643    0.004173  16.698437
INODE_UPDATE      529        128.639432  0.001193  29794.744530
GET_XSTAT         60427      0.004458    0.000913  0.198555
SET_XSTAT         2508       0.850394    0.003859  82.463500
GET_LOCK          58         0.007511    0.005643  0.017185
SET_LOCK          4108       0.008528    0.006346  0.073093
FSYNC             4          520.241355  63.859949  1066.216852
CLOSE_FILE        52473      0.003139    0.001991  0.441536

```

LINK_BY_FH	500	0.565615	0.334298	7.363347
RENAME_BY_FH	734	1.046008	0.345693	239.168672
STAT_BY_NAME	3841	0.004900	0.002042	0.074003
UNLINK_BY_NAME	3107	0.754904	0.356306	68.213809
CREATE_BY_NAME	2110	1.016229	0.364189	598.707250
READ_BY_FD	176	3.782423	0.006406	31.717223
WRITE_BY_FD	842	16.791472	0.384152	544.854531
STATFS_BY_FH	3000	0.004872	0.004324	0.025681
TRACE_ME	2	0.001882	0.001441	0.002324

To detect server code contention, next functionality added was to get statistics for all operations in NFSv3 and NFSv4. Again these features can be enabled/disabled dynamically.

Stats for NFSv3 all operations (v3_full) – This option provides details on all NFSv3 related operations carried out by Ganesha server. Details includes number of operations, related errors, any dup requests and operation latency (average, min & max). A typical o/p is as below:

```
# ganesha_stats v3_full
NFSv3 Detailed statistics
Stats collected since: Tue Dec 10 12:53:59 201992522344 nsecs
Duration: 1210.2047095299 seconds
```

Operation Details				Operation Latency (in milliseconds)		
Name	Total	Error	Dups	Avg	Min	Max
GETATTR	1427	0	0	0.012034	0.006625	0.094338
SETATTR	1195	0	0	0.703558	0.022861	10.729440
LOOKUP	1902	0	0	0.041535	0.011379	0.171812
ACCESS	450	0	0	0.020110	0.007937	0.091183
READ	85	0	0	4.895197	0.020499	31.739047
WRITE	255	0	0	44.750014	0.432070	544.887935
CREATE	867	0	0	1.725369	0.426161	598.829033
MKDIR	175	0	0	0.725980	0.438336	7.636447
SYMLINK	250	0	0	0.658445	0.436082	1.342323
REMOVE	1365	0	0	0.859535	0.403199	68.317329
RMDIR	175	0	0	0.900281	0.455928	38.541144
RENAME	358	0	0	1.256913	0.399444	239.275290
LINK	250	0	0	0.601190	0.363632	6.728508
REaddirPLUS	625	0	0	0.837069	0.012343	3.361598
FSSTAT	1500	0	0	0.013503	0.012366	0.063472
FSINFO	2	0	0	0.029196	0.018103	0.040289
PATHCONF	1	0	0	0.016904	0.016904	0.016904
COMMIT	2	0	0	115.714986	63.911308	167.518664

Stats for NFSv4 all operations (v4_full) - This option provides details on all NFSv4 related operations carried out by Ganesha server. Details includes number of operations, related errors, any dup requests and operation latency (average, min & max). A typical o/p is as below:

```
# ganesha_stats v4_full
NFSv4 Detailed statistics
Stats collected since: Tue Dec 10 12:53:59 201992522344 nsecs
Duration: 1480.1579866409 seconds
```

Operation Details				Operation Latency (in milliseconds)		
Name	Total	Error	Dups	Avg	Min	Max
ACCESS	1332	0	0	0.619486	0.011668	25.192819
CLOSE	1069	0	0	0.040280	0.023921	0.521888
COMMIT	2	0	0	924.888101	783.496666	1066.279536
CREATE	425	0	0	0.617226	0.423520	2.122538
GETATTR	9633	0	0	0.576492	0.015825	321.532584
GETFH	3442	0	0	0.350488	0.014266	25.189050
LINK	250	0	0	0.596659	0.403118	7.395720
LOCK	2052	0	0	0.052119	0.028803	7.083470
LOCKT	50	0	0	0.057371	0.044383	0.101646
LOCKU	2056	0	0	0.027167	0.024745	0.119471

LOOKUP	2862	915		0.035787	0.012751	0.266933
OPEN	1069	0		0.805058	0.044318	25.182468
OPEN_CONFIRM	1	0		0.027805	0.027805	0.027805
PUTFH	17990	0		0.009840	0.006355	0.160252
PUTROOTFH	1	0		9.515081	9.515081	9.515081
READ	91	0		2.793285	0.026140	17.611686
REaddir	622	0		0.934111	0.019732	3.466086
REMOVE	1567	0		0.764680	0.402491	19.677058
RENAME	376	0		0.968536	0.405123	44.514241
RESTOREFH	250	0		0.599618	0.405953	7.398653
SAVEFH	626	0		0.011843	0.009684	0.058360
SETATTR	1237	0		1.017108	0.026840	43.474325
SETCLIENTID	1	0		8.197145	8.197145	8.197145
SETCLIENTID_CONFIRM	1	0		5.590990	5.590990	5.590990
WRITE	587	0		4.695734	0.421245	321.503810
RELEASE_LOCKOWNER	2048	0		0.022850	0.020414	0.425051

All above stats helps in checking server side issues. But at times customer do report that a specific client or set of client face performance issues. This raises need to have client specific stats, which will provide details on what operations are being requested by the client.

By default client specific stats collection was being made for READ, WRITE, & “rest all” operations. This stat collection used to gather latency info as well. But as the latency can not be client specific, made code change for not to calculate the latency. As we have these stats already gathered, provided an option to extract those.

Stats for a clients IO related operations (client_io_ops <ip address>) - This option provides information on how many READ, WRITE and rest all operations requested by the specific client related to NFSv3 and NFSv4.x. It also provides info on bytes transferred in or out. A typical o/p is as below:

```
# ganesha_stats client_io_ops ::ffff:10.0.100.1
Client last active at: Tue Dec 10 12:55:15 2019778872927 nsecs
      NFSv3:
      total      errors      transferred
READ :          170           0      31457280
WRITE:          510           0      74415901
Other:        10544           0
      NFSv4.0:
      total      errors      transferred
READ :          182           0      31476567
WRITE:         1174           0      84901811
Other:        68377        1830
      No NFSv4.1 activity
      No NFSv4.2 activity
```

Stats for an exports latency for IO related operations (export_details <export_id>) - This option provides information on export specific latency. May be this is redundant as we can get the similar info via iov3 & iov4 combined. The typical o/p is as below:

```
# ganesha_stats export_details 34
Export last accessed at: Tue Dec 10 12:55:15 2019778872927 nsecs
Latency is in milliseconds for Read/Write/Other Operations
      NFSv3:
      total      errors      latency/delays      transferred
READ :          170           0           2.447598      31457280
WRITE:          510           0          22.375007      74415901
Other:        10544           0           0.515889
      NFSv4.0:
      total      errors      latency/delays      transferred
READ :          182           0           1.391416      31476567
WRITE:         1174           0           2.341099      84901811
```

```

Other:          64237          1830          0.262847
      No NFSv4.1 activity
      No NFSv4.2 activity

```

Stats for a client all ops (client_all_ops <client_ip_address>) - This option provides information all operations in all protocols (NFSv3, NFSv4.x, NLM). The typical o/p is as below:

```

# ganessa_stats client_all_ops ::ffff:10.0.100.1
Client last active at: Tue Dec 10 12:55:15 2019778872927 nsecs

```

NFSv3 Operations

Op Name	total	errors	dups
GETATTR	1427	0	0
SETATTR	1195	0	0
LOOKUP	1902	0	0
ACCESS	450	0	0
READ	85	0	0
WRITE	255	0	0
CREATE	867	0	0
MKDIR	175	0	0
SYMLINK	250	0	0
REMOVE	1365	0	0
RMDIR	175	0	0
RENAME	358	0	0
LINK	250	0	0
REaddirPLUS	625	0	0
FSSTAT	1500	0	0
FSINFO	2	0	0
PATHCONF	1	0	0
COMMIT	2	0	0

NLMv4 Operations

Op Name	total	errors	dups
TEST	8	0	0
LOCK	1	0	0

NFSv4 Operations

Op Name	total	errors
ACCESS	1332	0
CLOSE	1069	0
COMMIT	2	0
CREATE	425	0
GETATTR	9633	0
GETFH	3442	0
LINK	250	0
LOCK	2052	0
LOCKT	50	0
LOCKU	2056	0
LOOKUP	2862	915
OPEN	1069	0
OPEN_CONFIRM	1	0
PUTFH	17990	0
PUTROOTFH	1	0
READ	91	0
REaddir	622	0
REMOVE	1567	0
RENAME	376	0
RESTOREFH	250	0
SAVEFH	626	0
SETATTR	1237	0

SETCLIENTID	1	0
SETCLIENTID_CONFIRM	1	0
WRITE	587	0
RELEASE_LOCKOWNER	2048	0

NFSv4 Compound Operations		
total	errors	Ops in compond
68377	1830	51470