

AFPLOOK

The command sampler (AFPLOOK) can be used to determine where the best results may be expected from Fastpath by reporting on the command constructs that qualify for Fastpath.

The sampler can be controlled and viewed online using SYSAFP but may also be controlled without SYSAFP using the defaults module AFPLUKD. In either case, the sampler will print a report in DDPRINT at ADAEND time if a sample is active at that time.

- Enabling AFPLOOK
 - Operational Defaults
 - Adjusting the Defaults
 - AFPLOOK Parameters
 - AFPLOOK Report
-

Enabling AFPLOOK

AFPLOOK is enabled using the following ADARUN command:

```
ADARUN FASTPATH=YES
```

Use the SYSAFP administration center to activate the sample.

Operational Defaults

AFPLOOK is set up with certain operational defaults that control the amount of memory used during command analysis by restricting

- the maximum number of files sampled; and
- the number of concurrent users.

If any parameter is exceeded, AFPLOOK tries to ignore the excess while still reporting maximum information. In this way, AFPLOOK audits a general sampling of the database command workload to determine Fastpath optimization parameters. The operational defaults can be modified for site requirements as described in the section *AFPLOOK Parameters*.

Adjusting the Defaults

AFPLOOK is designed for dynamic use. It is inactive by default. You use the SYSAFP administration center to activate and target it. However, you can control it independently of the administration center.

It is possible to configure AFPLOOK independently of SYSAFP by creating a defaults module AFPLUKD. Please refer to the sample job SAGLUKD which is distributed as part of the Adabas release tape.

Each of the parameters described in the section APLOOK Parameters. can be pre-configured using AFPLUKD.

AFPLOOK Parameters

This section describes the AFPLOOK parameters, which are used to define the boundaries of the sample and limit the amount of memory required.

- Maximum Files
- Command/Descriptors per File
- Maximum Concurrent Users
- Maximum CIDs per User
- Maximum Commands Processed
- Job Name
- Selected Files
- Demo Sampling (AFPLUKD Only)
- Real Sampling (AFPLUKD Only)
- Report Title (AFPLUKD only)

Maximum Files

The maximum number of files to be sampled.

Once the maximum number of files is put in the analysis table, no additional files are sampled; however, additional files show in a command count so that it can be determined whether or not this parameter should be increased for subsequent executions.

Default: 64

Command/Descriptors per File

The maximum number of command/descriptor entries per file.

In conjunction with the Maximum Files parameter, this parameter restricts the amount of memory used. If the maximum entries is reached for a file, the last entry is converted into a general accumulator. Note that only one entry is required for the Adabas command types L1,L2,S8, and S9.

Default: 32

Maximum Concurrent Users

The maximum size for the table of concurrent users.

If all the user areas are being used at one time, a new request is satisfied by releasing the 'oldest usage' user area. The number of times this reusage occurs is noted and printed in the summary. When a sample contains a high percentage of reusage, this parameter should be adjusted.

Default: 100

Maximum CIDs per User

The maximum concurrent Adabas Command IDs (CIDs) sampled for each user processed.

In conjunction with the Maximum Concurrent Users parameter, this parameter restricts the amount of memory used. Commands for Command IDs that exceed this maximum are ignored and reported as rejected. When a sample contains a high percentage of rejections, this parameter should be adjusted.

Default: 10

Maximum Commands Processed

The maximum number of commands to be sampled.

Default: No limit

Job Name

Used to restrict sampling to a particular job name. One or more asterisks (*) can be used in the job name as a wild card character so that the sample can select all jobs that match the name ignoring the character positions occupied by an asterisk (*).

Default: None

Selected Files

Used to restrict sampling to specific files.

This option may be useful where the maximum files overflowed, or file activity is known and detailed analysis is required.

Default: All files

Demo Sampling (AFPLUKD Only)

For use when pre-configuring AFPLOOK with the defaults module AFPLUKD and licensed Fastpath is not present.

DEMO= sets whether demo AFPLOOK samples by default or not.

Settings are: ON / OFF

Default: ON

Real Sampling (AFPLUKD Only)

For use when pre-configuring AFPLOOK with the defaults module AFPLUKD and licensed Fastpath is present.

REAL= sets whether real AFPLOOK samples by default or not.

Settings are: ON / OFF

Default: OFF

Report Title (AFPLUKD only)

For use when pre-configuring AFPLOOK with the defaults module AFPLUKD.

This parameter sets the title that will be seen on the output report. This can be a maximum of 30 characters.

Default: AFPLOOK report summary

AFPLOOK Report

This section describes the types of information available on the AFPLOOK report:

- File Summary
- Potential Optimization Summary
- Sample Command Analysis
- Report Parameters

File Summary

This section of the report provides a summary of the file commands.

```

-----
FNR CC DESC DIRECT ACC RC SEQUENTIAL SEQUENCES
-----
 20 L1 -- 1
    L2 -- 4 4
    L3 CC 1
    L9 AA 1
    L9 BB 2
    L9 CC 2 1 1
    S1 AA 3 1
    TOTALS 7 8 21 (18%)
                EXCLUDED COMMANDS: 2
                ALREADY PREFETCHED: 3
(UPDATES 2,INSERTS 1,DELETES 1) (MAX.RBL DA 0,SEQ 32)
-----

```

| Column | Explanation |
|------------|---|
| FNR | Adabas file number. |
| CC | Adabas command code. |
| DESC | Primary descriptor. |
| DIRECT ACC | Maximum number of direct access commands that can be optimized. |
| RC | Maximum number of RC commands that can be optimized. |
| SEQUENTIAL | Maximum number of sequence commands that can be optimized. |
| SEQUENCES | Number of sequences that caused the number of sequential commands. The sequence factor for optimization may be calculated from these two numbers. |

The rightmost number shows the total sampled commands for the file together with the percentage relative to all the sampled file commands. On a large report, this number can be used to determine quickly which files should be considered for optimization.

Commands that have been ignored for the file are also listed along with the reason for exclusion.

The final line shows the update commands as well as the maximum record buffer lengths found for direct access or sequential commands that can be optimized.

Potential Optimization Summary

This section of the report summarizes the total commands sampled for all files and expresses this as a percentage of all commands seen. Excluded commands are similarly reported.

```

-----
POTENTIAL OPTIMIZATION SUMMARY -----
SAMPLED COMMANDS          MAXIMUM OPTIMIZATION
SAMPLED FILE COMMANDS    116 ( 77%) <----- SEQUENTIAL:          55 (47%)
                                DIRECT ACCESS:       32 (27%)
                                RCS:                 4 ( 3%)
EXCLUDED COMMANDS        33 ( 22%)
TOTALS                   149 (100%)
                                91 (61%)
-----

```

The maximum optimization numbers are an estimation of potential optimization. The sequential commands, direct access, and RC totals are expressed as a percentage of the total sampled file commands. The total is expressed as a percentage of all commands.

These numbers indicate the estimated *potential* optimization using Fastpath. The actual optimization will depend on various factors unique to each user site. Contact Software AG for assistance when interpreting samples.

Sample Command Analysis

This section of the report provides command analysis information.

| ----- | | | |
|-------------------------|---|------------|--|
| COMMAND ANALYSIS | | | |
| REJECTED COMMANDS | | | |
| MAX. USERS EXCEEDED: | 0 | | |
| MAX. CIDS EXCEEDED: | 0 | | |
| MAX. FILES EXCEEDED: | 0 | 0 (0%) | |
| EXCLUDED COMMANDS | | | |
| BAD COMMANDS: | 4 | | |
| NON-FILE COMMANDS: | 7 | | |
| NON-FILE RCS: | 2 | | |
| EXCLUDED FILE COMMANDS: | 8 | | |
| UPDATE COMMANDS: | 4 | | |
| ALREADY PREFETCHED: | 8 | 33 (22%) | |
| SAMPLED FILE COMMANDS | | 116 (77%) | |
| ALL COMMANDS SEEN | | 149 (100%) | |
| ----- | | | |

The numbers shown:

- illustrate the type of commands processed, and
- put the previous section into perspective.

Rejected commands are categorized by users, CIDs, and files exceeded. If the total percentage is high, estimates reported elsewhere may not give an accurate assessment.

Excluded commands are split into the following categories:

| Category | Explanation |
|------------------------|---|
| Bad commands | Unexpected Adabas response codes. |
| Non-file commands | Commands that cannot be attributed to a file; for example, OP,CL,ET,C1,RE. Plus file commands HI,LF,RI. |
| Non-file RCs | All RC commands plus any RC for which the CID is not stored by AFPLOOK. |
| Excluded file commands | L4,L5,L6,S4,S5. |
| Update commands | A1,A4,E1,E4,N1. |
| Already prefetched | Any command that could qualify for sequential optimization that has prefetch or multifetch already set. |

Report Parameters

This section of the report

- shows the important parameters used to produce the report; and
- gives an indication of the parameters needed.

| | | | |
|---------------------------------|----------|---------------|----|
| ----- | | | |
| PARAMETERS USED | | | |
| MAX. FILES: | 64 | FILES NEEDED: | 5 |
| ..MAX. DE: | 32 | OVERFLOWS: | 0 |
| MAX. USERS: | 100 | HIGH USERS: | 15 |
| ..MAX.CID: | 10 | HIGH CIDS: | 4 |
| * REUSED USER AREA OCCURRENCES: | | | 0 |
| MAX.RECORDS: | NO LIMIT | | |
| ----- | | | |