

Natural Roll Server Operation

This section covers the following topics:

- Roll Server System Requirements
- Formatting the Roll File
- Starting the Roll Server
- Roll Server Messages, Condition Codes and Abend Codes
- Return Codes and Reason Codes of the Roll Server Request
- Operating the Roll Server
- Roll Server Performance Tuning
- Roll Server User Exits

OS/390 Environment - Other Topics

Natural under OS/390 | Authorized-Services Manager | Shared Natural Nucleus | Natural Roll Server Functionality

Roll Server System Requirements

This section describes the roll server system requirements. The following topics are covered:

- APF Authorization
- System Linkage Index
- Virtual Storage
- CF Structure
- XCF Signalling Paths

APF Authorization

Copy module NATRSM41 to an APF-authorized library. NATRSM41 is shipped with AC=1.

System Linkage Index

As the Roll Server reserves one system linkage index (System LX), check whether there is a high enough value of NSYSLX in member IEASYSxx of library SYS1.PARMLIB.

When the Roll Server terminates, its address space ID is no longer available because a System LX has been used. It becomes available again with the next IPL.

To avoid this, deactivate the Roll Server with the DEAL operator command and restart it afterwards.

Once a System LX has been reserved, it is reused with every restart of the Roll Server until the next IPL.

Virtual Storage

| | |
|--------------------------------------|-------------------------------------|
| ECSA | 84 bytes |
| Private program storage | 30 KB above |
| Fixed subpool storage (incl. ELSQA): | 10 KB below, 50 KB above |
| LRB directory: | |
| 100 slots per roll file | 4 KB above |
| Every additional roll file | 30 KB above |
| Working storage | depending on load, about 1 MB above |

CF Structure

The space required is calculated using the following formula:

$$24 \text{ KB} + (\text{RFN} + 1) * 1 \text{ KB} + (\text{RFS} + 8) * 160 \text{ bytes}$$

where RFN denotes the number of roll files and RFS denotes the total number of roll-file slots in all roll files.

Example:

There are five roll files with 1000 slots each.

$$24 \text{ KB} + 6\text{KB} + 5008 * 160 \text{ bytes} = 24 \text{ KB} + 6 \text{ KB} + 783 \text{ KB} = 813 \text{ KB}$$

The CF structure should thus be defined with 820 KB:

```
STRUCTURE NAME(NATROLLS) SIZE(820) PREFLIST(CF1)
```

XCF Signalling Paths

In a Parallel Sysplex, the Roll Servers communicate via the XCF Signalling Services. As the default XCF group name, the leftmost eight bytes of the CF structure name are used.

If you want to specify your own XCF groupname, use the NATRSU24 user exit. For more information on this user exit, see NATRSU24 User Exit.

Formatting the Roll File

To format the roll file, proceed as follows:

1. Allocate it as a physical, sequential dataset with a fixed-record format.
2. Format it using module NATRSRFI.

During formatting, the roll file is converted to BDAM format with a device-dependent block size.

To format, enter the following parameter string under the DD name RFIPARMS, or as PARM on the JCL EXEC statement:

function, dd-name, slot-size, number-of-slots

All parameters are positional; they are explained in the table below:

| Parameter | Description |
|-----------------|--|
| function | <p>FORMAT ☺ Format the roll file.</p> <p>RESET - All roll file slots are reset (marked as free). You can only use this parameter value if the roll file has already been formatted. The only other parameter allowed is dd-name.</p> <p>LIST ☺ Print a list of session-IDs contained in the roll file and their last activity. The only other parameter allowed is DD-name.</p> |
| dd-name | The name of the DD statement under which the roll file has been specified. |
| slot-size | <p>The size of a roll file slot in bytes. This size is rounded to the next higher multiple of the block size used.</p> <p>It is recommended to initially use a slot size equal to the size of the Natural thread. Then look at the Roll Server statistics. They also show the largest occurrence of a thread size. Use this value to reduce the slot size, if necessary.</p> |
| number-of-slots | The number of roll file slots to be allocated. This number is the maximum number of concurrently active users. This parameter is optional. If omitted, the entire roll file, as allocated, will be formatted. |

Note: For compatibility with Version 3.1, the old formatting function CINIT is still supported. Please refer to the Natural Version 3.1.x for Mainframes documentation.

To calculate the required disk space in cylinders (SPACE parameter of the DD statement), use the following formula:

number-of-cylinders = ceiling (slot-size / 30*block-size)

or

number-of-tracks = ceiling (slot-size / 2*block-size)

The block size used is:

23476 for 3380 DASD

27998 for 3390 DASD

22928 for 9345 DASD

In addition, space is needed for the roll file directory header (40 bytes) and one directory entry for each roll file slot (24 bytes). Thus, one additional block is needed for roughly 976 slots on 3380, 1164 slots on 3390, or 953 slots on 9345 DASD.

NATRSRFI Output

If DD-name RFIPRINT is specified, NATRSRFI directs its output to this dataset. When RFIPRINT is omitted, output is written to JESMSGLG using the WTO macro (ROUTDCE=11). Note that RFIPRINT must be specified for the LIST function.

NATRSRFI Conditon and Abend Codes

The following condition codes are used:

| | |
|----|---|
| 0 | Normal completion. |
| 4 | Number of slots formatted is less than requested. |
| 20 | Parameter error. |

The following user abend codes are used:

| Abend Code | Cause |
|------------|---------------------------------------|
| U0100 | Open for RFIPARMS or RFIPRINT failed. |
| U0101 | Open for Roll File failed. |

Example 1:

```
//FBRUNRFI JOB (FB,218),FB,CLASS=K,MSGCLASS=X,NOTIFY=FB
//FORMAT EXEC PGM=NATRSRFI
//STEPLIB DD DISP=SHR,DSN=NATURAL.NAT41.LOAD
//RF1 DD DISP=SHR,DSN=FB.SYSF.ROLLF1
//RF2 DD DISP=SHR,DSN=FB.SYSF.ROLLF2
//RFIPARMS DD *
FORMAT,RF1,200000,1000
FORMAT,RF2,200000
/*
```

Exerpt from resulting JESMSG LG:

```
+FBRUNRFI: FORMAT,RF1,200000,1000
+FBRUNRFI: RF1: FB.SYSF.ROLLF1
+FBRUNRFI: Creation date: 2001/06/13 Volume: ADA002(3390)
IEC031I D37-04,IFG0554P,FBRUNRFI,FORMAT,RF1,305B,ADA002,FB.SYSF.ROLLF1
+FBRUNRFI: Not enough space for 1000 slots.
+FBRUNRFI: 60 Blocks written. Block size is 27998.
+FBRUNRFI: 1 Directory block.
+FBRUNRFI: 8 Blocks per slot. Slot size is 223984.
+FBRUNRFI: 7 Slots initialized. Roll file version 411.
+FBRUNRFI: 3 Blocks unused.
+FBRUNRFI: FORMAT,RF2,200000
+FBRUNRFI: RF2: FB.SYSF.ROLLF2
+FBRUNRFI: Creation date: 2001/06/08 Volume: USRF08(3380)
IEC031I D37-04,IFG0554P,FBRUNRFI,FORMAT,RF2,020F,USRF08,FB.SYSF.ROLLF2
+FBRUNRFI: 60 Blocks written. Block size is 23476.
+FBRUNRFI: 1 Directory block.
+FBRUNRFI: 9 Blocks per slot. Slot size is 211284.
+FBRUNRFI: 6 Slots initialized. Roll file version 411.
+FBRUNRFI: 5 Blocks unused.
```

Example 2:

```
//FBRUNRFI JOB (FB,218),FB,CLASS=K,MSGCLASS=X,NOTIFY=FB
//FORMAT EXEC PGM=NATRSRFI,PARM='FORMAT,RF1,200000'
//STEPLIB DD DISP=SHR,DSN=NATURAL.NAT41.LOAD
//RF1 DD DISP=SHR,DSN=FB.SYSF.ROLLF1
//RFIPRINT DD SYSOUT=X
```

Resulting RFIPRINT:

```
Natural Roll Server - Roll File Utility Version 411
FORMAT,RF1,200000
RF1: FB.SYSF.ROLLF1
Creation date: 2001/06/13 Volume: ADA002(3390)
60 Blocks written. Block size is 27998.
1 Directory block.
8 Blocks per slot. Slot size is 223984.
7 Slots initialized. Roll file version 411.
3 Blocks unused.
```

Notes Concerning the Formatting or Resetting of Roll Files

- You can format or reset several roll files at once by specifying several parameter lines in RFIPARMS.
- You cannot format or reset a roll file while the roll server is active.
- When the roll file is formatted in a Parallel Sysplex, the roll server Coupling Facility structure must also be cleared using the SETXCF operator command, for example:
SETXCF FORCE,STR,STRNAME=NATROLL1

Starting the Roll Server

You start the Roll Server either as a batch job or as a started task by executing module NATRSM41. The roll file(s) must be defined as DD-names ROLLF1 to ROLLF5.

On the JCL EXEC statement, specify as PARM the following parameters:

| |
|--|
| <p><i>subsystem-id, number-of-roll-files, number-of-LRB-slots, LRB-slot-size, CF-structure-name, low-water-mark, high-water-mark</i></p> |
|--|

All parameters are positional and must be separated by a comma. They are explained in the table below:

| Parameter | Possible Values | Default | Comment |
|-----------------------------|--------------------------|---------------------|--|
| <i>subsystem-id</i> | 4-byte non-blank string | NAT4 | The specified value must match the value of the Natural profile parameter SUBSID. Note: With Natural under CICS, refer to the ROLLSRV parameter in the NCMDIR macro for setting the appropriate subsystem ID. |
| <i>number-of-roll-files</i> | 0 - 5 | 1 | In a non-Parallel Sysplex, the Roll Server can operate without a roll file, using only the in-storage Local Roll Buffer. |
| <i>number-of-LRB-slots</i> | 1 - 32767 | none | The number of LRB slots multiplied by the slot size must not exceed 2 GB. The same number of LRB slots is assigned for each LRB, i.e. for each roll file used. The total number of LRB slots is calculated by the formula: <i>number-of-roll-files * number-of-LRB-slots</i> |
| <i>LRB-slot-size</i> | any numeric value | roll file slot size | Value in number of bytes. This parameter must be specified if no roll file is used. |
| <i>CF-structure-name</i> | any valid structure name | none | If you specify less than 16 characters, blanks are appended. Only specify this parameter if you use the Coupling Facility (with Parallel Sysplex). |
| <i>low-water-mark</i> | 1 - 9 | 7 | Specifies the low water mark in steps of ten percent of the number of LRB slots. Only specify this parameter if you do not use the CF. |
| <i>high-water-mark</i> | 1 - 10 | 8 | Analogous to low-water-mark parameter. Value "10" means that the staging task will never be activated. It is only recommended to specify "10" if the LRB is large enough to serve all simultaneously active Natural sessions. |

Note: From z/OS V1.2 on, the Local Roll Buffer resides in a Memory Object "above the bar". Use the MEMLIMIT parameter on the EXEC statement to ensure enough memory can be allocated "above the bar".

Examples for Starting the Roll Server as a Batch Job

```
// EXEC PGM=NATRSM41,PARM=NA41,,1000
//ROLLF1 DD DSN=SYSF.ROLLFILE
```

The subsystem ID is NA41, one roll file is used (default), and the Local Roll Buffer has 1000 slots. The slot size used is identical with the roll file's slot size. The low water mark is 70% (default), the high water mark is 80% (default).

```
// EXEC PGM=NATRSM41,PARM=,5,1000,150000,NATROLL1,MEMLIMIT=800M
//ROLLF1 DD DSN=DASD1.ROLLFILE
//ROLLF2 DD DSN=DASD2.ROLLFILE
//ROLLF3 DD DSN=DASD3.ROLLFILE
//ROLLF4 DD DSN=DASD4.ROLLFILE
//ROLLF5 DD DSN=DASD5.ROLLFILE
```

The subsystem ID is NAT4 (default), five roll files are used, and each of the five Local Roll Buffers has 1000 slots. The LRB slot size is 150000 bytes. The roll file directory resides in the Coupling Facility structure NATROLL1. Low and high water marks are ignored, because every thread is written to the roll file (see Natural Roll Server Functionality). Since this job is intended for z/OS, the MEMLIMIT option specifies 800 Megabytes for the Local Roll Buffers.

Note: The Roll Server will not start if:

- Another Roll Server is running with the same *subsystem-id*.
- Another Roll Server is accessing a Roll File specified in its JCL
- A Roll File has been reformatted without resetting the CF structure, using the SETXCF FORCE command.

Roll Server Messages, Condition Codes and Abend Codes

The Roll Server writes informational and error messages to JESMSGLG using the WTO macro (ROUTCDE=11). The messages are preceded by a message identifier and the Roll Server's job name, for example:

RSM0019 FBRSM41 : Roll Server Version 411 is active

The following condition codes are used:

| | |
|----------------|-----------------------|
| 0 | Normal completion |
| 12 | Wrong parameter input |
| 16 | Runtime error |
| 20 | Abend has occurred |
| >100 | Initialization error |

User Abend Codes

When an unexpected return code is issued by an XCF or XES Service Call, an abend with a dump is forced. Register 14 of the abend register contains the reason code. To find a description of the reason, refer to *Programming: Sysplex Services Reference* (IBM documentation). If the error was not environment-specific, send the dump to Software AG support.

The following user abend codes are used:

| Abend Code | Cause |
|-------------------|-----------------|
| U0200 | IXLCONN failed |
| U0201 | IXLFORCE failed |
| U0202 | IXLLIST failed |
| U0203 | IXLDISC failed |
| U0204 | IXCLEAVE failed |
| U0301 | IXLLIST failed |
| U0302 | IXCMMSGO failed |
| U0401 | IXLLIST failed |
| U0501 | IXLLIST failed |

Return Codes and Reason Codes of the Roll Server Request

These are codes that Natural may receive from the Roll Server's PC services routines. They are reported by the respective teleprocessing interfaces (Natural CICS or Natural IMS interface). For a list of these codes, refer to the Return Codes and Reason Codes of the Roll Server Request (in the Messages and Codes documentation).

Operating the Roll Server

The following commands can be passed to the Roll Server via the MODIFY operator command:

| Command | Description |
|-------------------------|--|
| TERM | Stops the Roll Server. The roll file directory and all modified LRB slots are written to the roll file and the address space is terminated. The address space ID is no longer available until the next IPL. |
| DEAL | The Roll Server is stopped, but the address space is not terminated. The roll file directory and all modified LRB slots are written to the roll file. In a de-allocated status, the Roll Server can be restarted with new parameters and the old address space ID. The roll files can be reformatted in de-allocated status. If you do that, however, currently active Natural sessions are no longer restartable. |
| SNAP | Debugging function. The Roll Server's address space is dumped to SYSUDUMP. |
| TRSTART | Debugging function, only to be used at Software AG's advice. Activates the Trace Task. If the General Trace Facility (GTF) is started and enabled for user records of Type 200, trace records are written to the GTF. |
| TRSTOP | Deactivates the Trace Task. |
| START,parmstring | Reactivates the Roll Server with the specified parameters. You can only use this command in deactivated status. |

Roll Server Performance Tuning

As a general rule for Roll Server performance tuning, give the Roll Server a higher dispatching priority than the address spaces where Natural runs.

To find out where the weaknesses in performance are, analyze system performance with the Roll-Server Statistics function of the SYSTP utility.

When looking at Roll-Server Statistics, keep an eye especially on the following values:

- The number of direct writes.
"Direct write" means that the Natural thread that was received was written to the roll file directly. There are two possible reasons:
 1. No LRB slot available. Increase the LRB.
 2. The compressed thread was larger than a single LRB slot. Increase the LRB slot size.
- The number of direct reads.
"Direct read" means that the requested thread was no longer in the LRB and had to be read directly from the roll file.
If the ratio of direct reads to the total number of reads is very high in a single OS/390 system, the LRB is too small (increase it).
If the ratio of direct reads to the total number of reads is very high in a Parallel Sysplex OS/390 system, this may also mean that there are many inter-system activities, which in turn means that a Natural session changes OS/390 images quite frequently during its lifetime.
- The number of staging waits (in a single OS/390 environment).
A "staging wait" is a situation where a write request had to wait until the Staging Task had written the LRB slot to the roll file. If the ratio of staging waits to the total number of write requests is very high, this indicates that the high and low water marks are set inappropriately or that there is a bottleneck on the roll file device/roll file channel.
Based on experience with stress tests, the following is recommended:

If the ratio of maximal number of active users to number of LRB slots is very small, increase the high water mark. If not, decrease the high water mark.

The difference between high water mark and low water mark should not be larger than three (30%).

Ideally, if the number of LRB slots is definitely larger than the maximum number of concurrent users, the high water mark should be set to 10.

Roll Server User Exits

The roll server has two user exits.

- NATRSU14
- NATRSU24

Sample source modules are delivered for these.

NATRSU14 User Exit

Specifies the roll file number to be used.

Entry calling conventions:

- Register 1 addresses the parameter list that is described by the following DSECT:

```

PLIST      DSECT
PLRSVER    DS CL4      Roll server version (= '411')
PLNRF      DS H        Number of roll files
PLUID      DS CL16     Userid
PLTSNUM1   DS H        Total number of slots Roll file 1
PLUSNUM1   DS H        Number of slots in use Roll file 1
PLTSNUM2   DS H        Total number of slots Roll file 2
PLUSNUM2   DS H        Number of slots in use Roll file 2
PLTSNUM3   DS H        Total number of slots Roll file 3
PLUSNUM3   DS H        Number of slots in use Roll file 3
PLTSNUM4   DS H        Total number of slots Roll file 4
PLUSNUM4   DS H        Number of slots in use Roll file 4
PLTSNUM5   DS H        Total number of slots Roll file 5
PLUSNUM5   DS H        Number of slots in use Roll file 5
PLISTL    EQU *-PLIST

```

- Register 13 points to a 36-fullword save area.
- Register 14 contains the return address.
- Register 15 contains the entry address of NATRSU14.

Return calling convention:

- Register 15 contains the number of the roll file in binary format.

Note:

If access registers are modified within this user exit, these access registers must be saved and restored on return. This user exit is called in primary addressing mode with PSW Key 8. Since it runs in cross-memory mode, no SVC except SVC 13 may be used.

NATRSU24 User Exit

Specifies the XCF group name to be used.

Entry calling conventions:

- Register 1 points to an 8-byte area in which the group name must be generated.
- Register 13 points to an 18-fullword save area.
- Register 14 contains the return address.
- Register 15 contains the entry address of NATRSU24.

As a group name default, the Roll Server will use the leftmost 8 bytes of the CF structure name.

This user exit is called in primary mode, PSW key 8 and in task mode.