IBM Software Group Enterprise Networking Solutions
z/OS® V1R11 Communications Server

z/OS V1R11 Communications Server Simplification and usability – syslogd enhancements

z/OS Communications Server Development, Raleigh, North Carolina

This presentation will give you an overview of the enhancements to the Communications Server in z/OS V1R11 for simplification and usability. The simplification and usability theme covers enhancements that in some way or another make it easier to deploy or to use certain functions of the Communications Server.

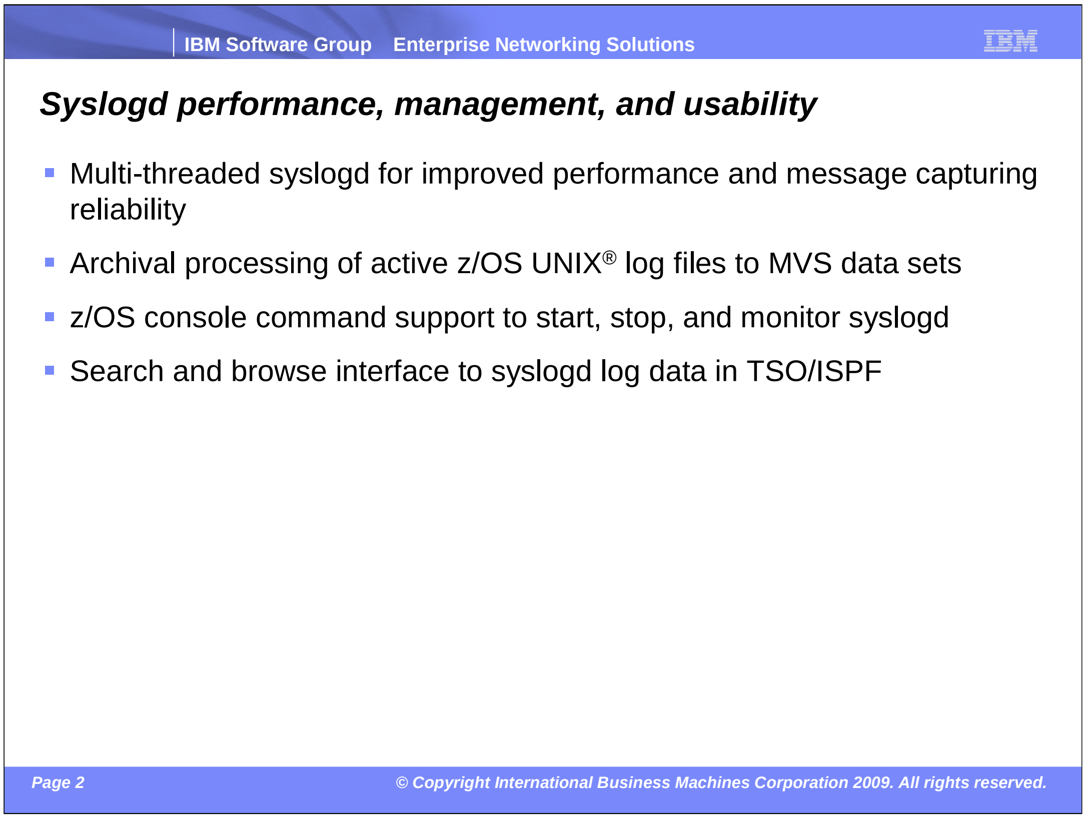

Syslogd performance, management, and usability

- Multi-threaded syslogd for improved performance and message capturing reliability
- Archival processing of active z/OS UNIX® log files to MVS data sets
- z/OS console command support to start, stop, and monitor syslogd
- Search and browse interface to syslogd log data in TSO/ISPF

Syslog daemon logs messages from z/OS UNIX applications. Syslogd is enhanced in V1R11 to perform better, thus reducing the likelihood of losing messages during peak periods. Syslogd is also enhanced to archive active z/OS UNIX log files based on various criteria, and syslogd is finally enhanced to support a set of MVS console commands for operational purposes. To improve accessibility to the logged messages, a new ISPF-based syslogd browser function is part of z/OS V1R11 Communications Server.

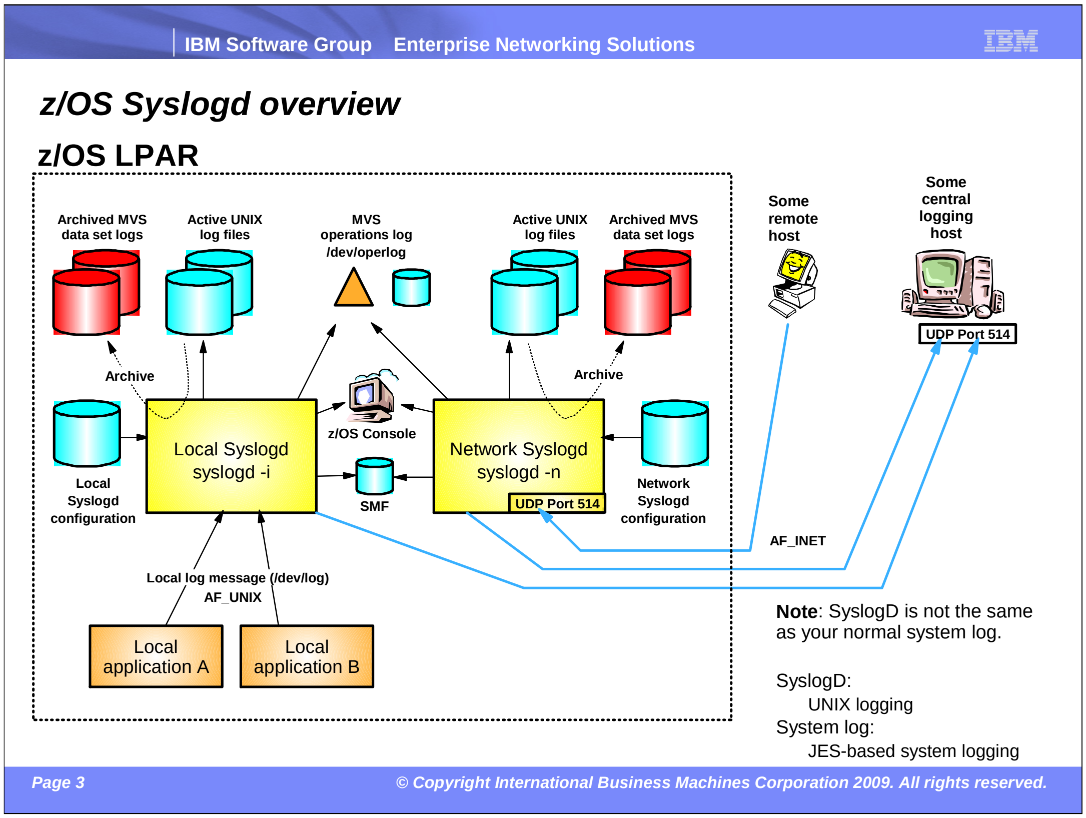

IBM Software Group Enterprise Networking Solutions

z/OS Syslogd overview

z/OS LPAR
Archived MVS data set logs
Active UNIX log files
MVS operations log /dev/operlog
Active UNIX log files
Archived MVS data set logs
Some remote host
Some central logging host
UDP Port 514
Archive
Archive
Local Syslogd syslogd -i
z/OS Console
Network Syslogd syslogd -n
UDP Port 514
Local Syslogd configuration
SMF
Network Syslogd configuration
AF_INET
Local log message (/dev/log)
AF_UNIX
Local application A
Local application B

Note: SyslogD is not the same as your normal system log.

SyslogD:
UNIX logging
System log:
JES-based system logging

Page 3

© Copyright International Business Machines Corporation 2009. All rights reserved.

z/OS supports running two syslogd instances:

One instance is used by local applications on the z/OS system where the syslogd instance is running. You start syslogd with a –i flag to indicate it is a local instance. Such a syslogd instance does not open UDP port 514 and is not subject to remote attacks on that UDP port. Such a local syslogd instance can be set up to send log messages to other syslogd instances, but it cannot receive any such messages.

Another instance is used by remote syslogd instances that are configured to send their log messages to a local consolidated syslogd on z/OS. You start syslogd with a –n flag to indicate it is a network syslogd instance. Such a syslogd instance opens UDP port 514 and receives messages from remote Syslogd instances. It does not receive messages from any local applications.

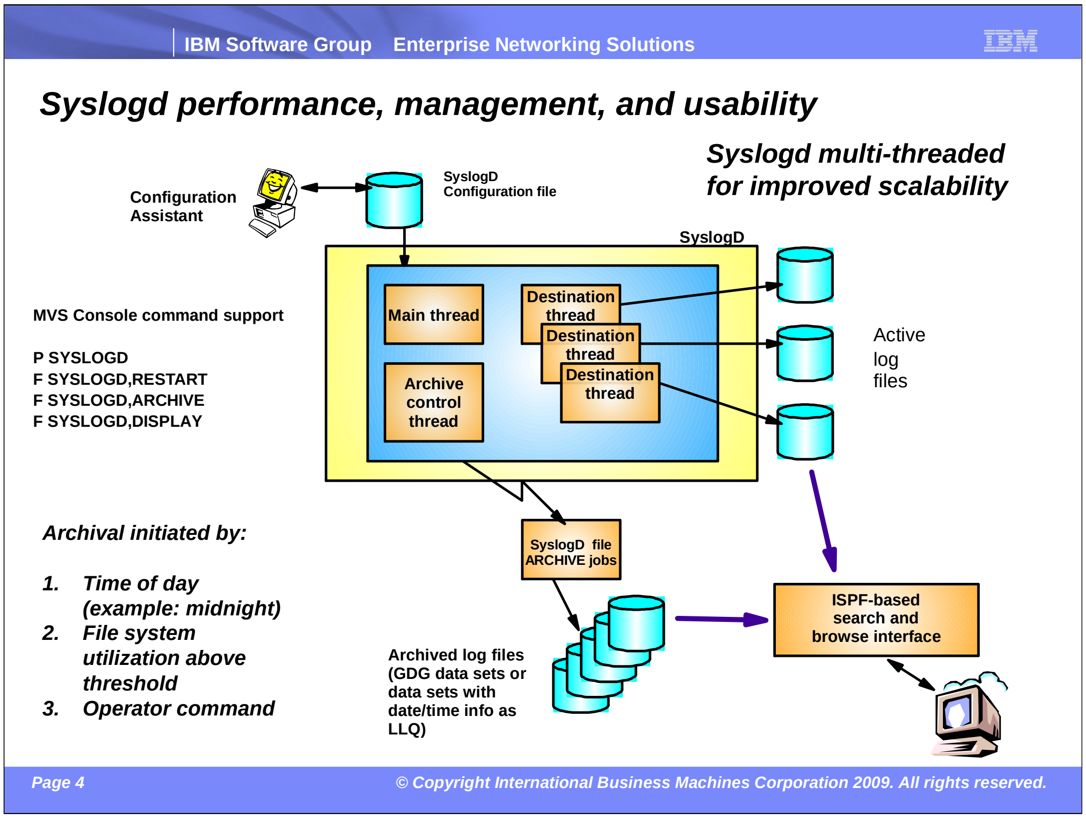

IBM Software Group Enterprise Networking Solutions

Syslogd performance, management, and usability

Syslogd multi-threaded for improved scalability

Configuration Assistant

SyslogD Configuration file

SyslogD

Main thread

Archive control thread

Destination thread

Destination thread

Destination thread

Active log files

MVS Console command support

P SYSLOGD
F SYSLOGD,RESTART
F SYSLOGD,ARCHIVE
F SYSLOGD,DISPLAY

SyslogD file ARCHIVE jobs

ISPF-based search and browse interface

Archival initiated by:

1. *Time of day (example: midnight)*
2. *File system utilization above threshold*
3. *Operator command*

Archived log files (GDG data sets or data sets with date/time info as LLQ)

Page 4

This slide shows a high-level view of the new and improved syslogd components.

Syslogd is now a multi-threaded implementation allowing for more parallel processing in peak periods. Syslogd continues to write log messages to z/OS UNIX files. A new archive function archives the content of a z/OS UNIX log file to an MVS data set. The MVS data set can either be a sequential data set (low level qualifiers specify date and time) or a new generation of a generation data group (GDG). The archive operation can be initiated by an operator at a specific time or when the utilization of a file system to which the z/OS UNIX log files are written exceeds a configurable threshold.

Command support includes the ability to shut syslogd down using a P (stop) command. Syslogd will not change its address space name after it has started. If you start a procedure by the name of SYSLOGD, the resulting address space name remains SYSLOGD.

The ISPF browser starts by reading the syslogd configuration file and locates the active z/OS UNIX files and all available MVS archives. It supports browsing individual files or data sets, in addition to performing extensive searches in one or a series of files or data sets.

The solution adds an automatic archival mechanism to syslogd which also supports on demand archival. You can archive once per day using a configurable time of day, or archive when any of the UNIX file systems reaches a configurable percentage full. You can also archive using an operator command. Syslogd archives UNIX files to either sequential or generation data group (GDG) data sets, and you can include system symbols in parts of the target data set names. You do not need to determine the space requirements of the target data sets: syslogd takes care of that. Syslogd automatically retries previously failed archives at the next archive event. You can monitor a console message for failed archives to correct any problems, and syslogd will eventually successfully archive all previously failed files. You can also use a new operator command to display the utilization of the syslogd UNIX file systems.

The default for all UNIX files is not to perform an automatic archival. If you want to use this new function you must explicitly configure it.

You have three choices for each rule that contains a UNIX file log destination. You can archive the file by using the new -N parameter. You can reinitialize the file (delete its contents) when an archive occurs. Use this option with care, because the contents of the file are lost. Or you can do nothing with the file by not using either the -N or -X parameters.

All eligible files are archived for the time of day and operator command triggers. But for the file system threshold trigger, syslogd attempts to reduce the space used by archiving files until half of the configured threshold is reached. For example, if you configure 80% as the threshold, syslogd archives files until the file system reaches 40% utilization. A console message is issued if all eligible files are archived but the file system utilization was not able to be reduced to half the configured threshold. This can happen if the file system contains files that are not managed by syslogd.

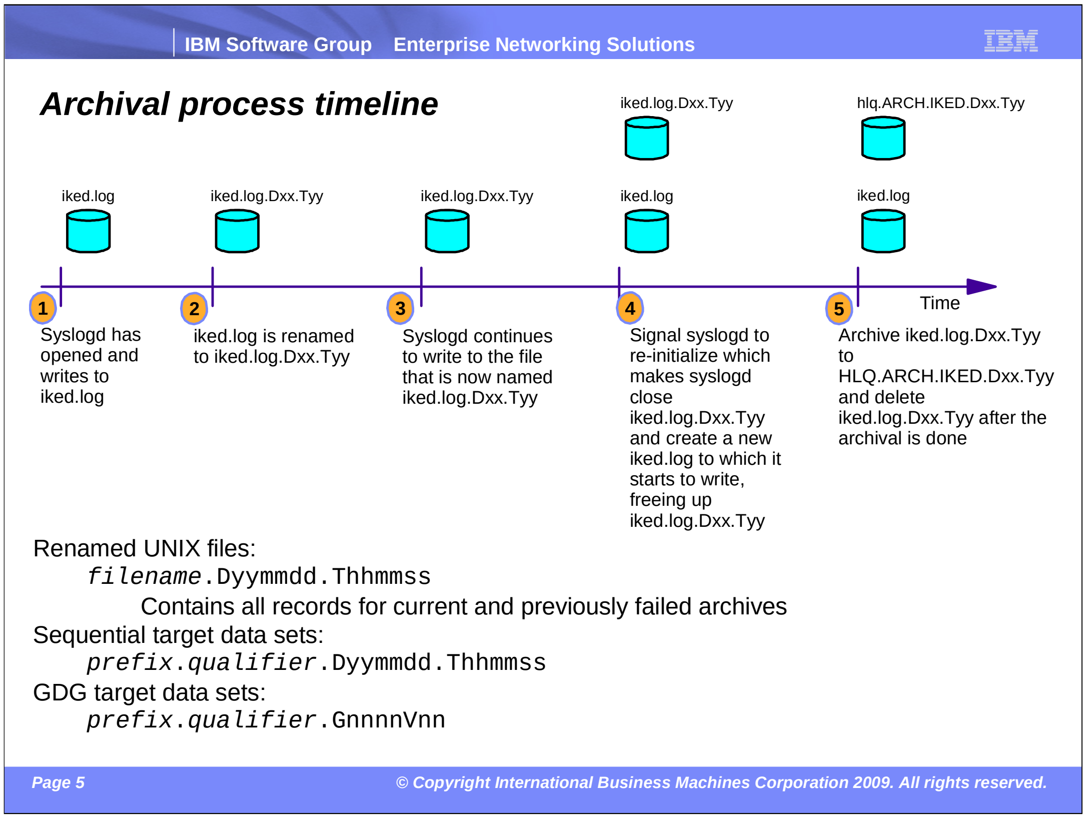

IBM Software Group | Enterprise Networking Solutions

Archival process timeline

iked.log.Dxx.Tyy

hlq.ARCH.IKED.Dxx.Tyy

iked.log | iked.log.Dxx.Tyy | iked.log.Dxx.Tyy | iked.log | iked.log

Time

1 Syslogd has opened and writes to iked.log

2 iked.log is renamed to iked.log.Dxx.Tyy

3 Syslogd continues to write to the file that is now named iked.log.Dxx.Tyy

4 Signal syslogd to re-initialize which makes syslogd close iked.log.Dxx.Tyy and create a new iked.log to which it starts to write, freeing up iked.log.Dxx.Tyy

5 Archive iked.log.Dxx.Tyy to HLQ.ARCH.IKED.Dxx.Tyy and delete iked.log.Dxx.Tyy after the archival is done

Renamed UNIX files:

`*filename*.Dyymmdd.Thhmmss`

Contains all records for current and previously failed archives

Sequential target data sets:

`*prefix.qualifier*.Dyymmdd.Thhmmss`

GDG target data sets:

`*prefix.qualifier*.GnnnnVnn`

Page 5 © Copyright International Business Machines Corporation 2009. All rights reserved.

This diagram shows the timeline of an archive event for a single file. The log file is named iked.log for this example.

At step one, the iked.log file is open and syslogd is writing records to it.

At step two, syslogd renames the open log file with a unique date and time suffix.

At step three, the file is still open so syslogd can continue to write records to it.

At step four, the renamed file is closed and the original iked.log file is re-created. Syslogd now writes to the open iked.log file.

At step five, the renamed file has been archived into a target data set, and the renamed file has been deleted.

The archive process consists of this sequence of events.

Syslogd allocates the target data set, opens the source file and target data set, and copies the file. Then it closes the file and data set, de-allocates the data set, and then deletes the renamed source UNIX file.

When syslogd renames UNIX files to prepare them for archival, it adds a unique suffix that identifies the current date and time. You might see these files in your file system if an automatic archival fails. If successive archive failures occur, this file contains the file contents from all previous failures, and is named with the most recent date and time stamp.

The format of the target data set name depends on whether you are using sequential data sets or GDG data sets. For sequential data sets, syslogd appends unique date and time values to the configured prefix and qualifier values. For GDG data sets, the system creates a unique suffix value according to how the GDG base data set was defined.

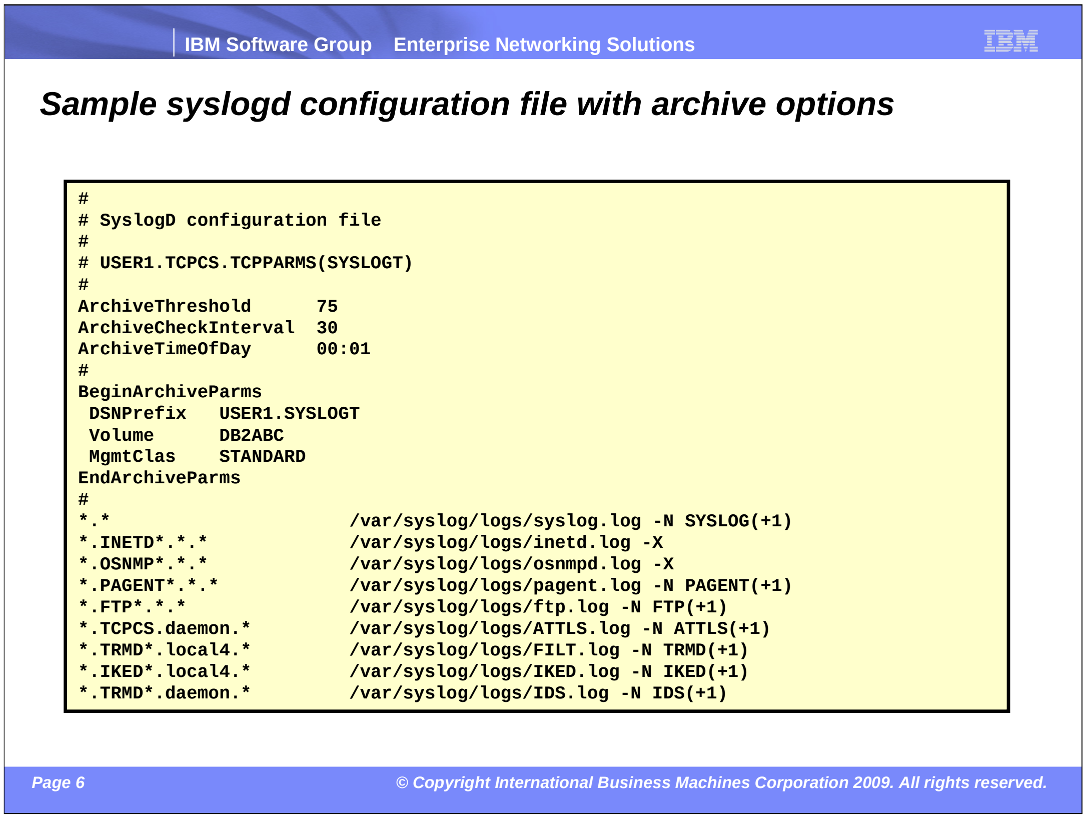

IBM Software Group Enterprise Networking Solutions

Sample syslogd configuration file with archive options

```
#
# SyslogD configuration file
#
# USER1.TCPCS.TCPPARMS(SYSLOGT)
#
ArchiveThreshold      75
ArchiveCheckInterval  30
ArchiveTimeOfDay      00:01
#
BeginArchiveParms
 DSNPrefix   USER1.SYSLOGT
 Volume      DB2ABC
 MgmtClas    STANDARD
EndArchiveParms
#
*.*                      /var/syslog/logs/syslog.log -N SYSLOG(+1)
*.INETD*.*.*             /var/syslog/logs/inetd.log -X
*.OSNMP*.*.*             /var/syslog/logs/osnmpd.log -X
*.PAGENT*.*.*            /var/syslog/logs/pagent.log -N PAGENT(+1)
*.FTP*.*.*               /var/syslog/logs/ftp.log -N FTP(+1)
*.TCPCS.daemon.*         /var/syslog/logs/ATTLS.log -N ATTLS(+1)
*.TRMD*.local4.*         /var/syslog/logs/FILT.log -N TRMD(+1)
*.IKED*.local4.*         /var/syslog/logs/IKED.log -N IKED(+1)
*.TRMD*.daemon.*         /var/syslog/logs/IDS.log -N IDS(+1)
```

Page 6

The **ArchiveThreshold** statement configures the percentage of UNIX file system utilization that triggers an archive. The utilization is checked at the interval specified with the **ArchiveCheckInterval** statement .

The **ArchiveTimeOfDay** statement configures the time of day for an automatic archival

The **BeginArchiveParms** statement configures the data set name prefix for the target data set. You must configure a data set name prefix before using the -N parameter on any syslogd rule. You can repeat the **BeginArchiveParms** statement multiple times for different groups of syslogd rules, or you can use a single instance of the statement to apply to all rules. In addition to the data set name prefix, you can configure several allocation parameters for the archive data set. These parameters have the same names, syntax, and meaning as the corresponding parameters on the DD JCL statement.

Use the **-N** parameter on a syslogd rule to indicate that the rule is eligible for automatic archival. Specify a data set name qualifier with the **-N** parameter.

Use the **-X** parameter on a syslogd rule to indicate that the contents of the file should be deleted when an archive event occurs. You should only use this parameter if you do not need to keep the contents of the file.

If you do not use the **-N** or **-X** parameter on a syslogd rule that specifies a UNIX file destination, then the file does not participate in automatic archival processing.

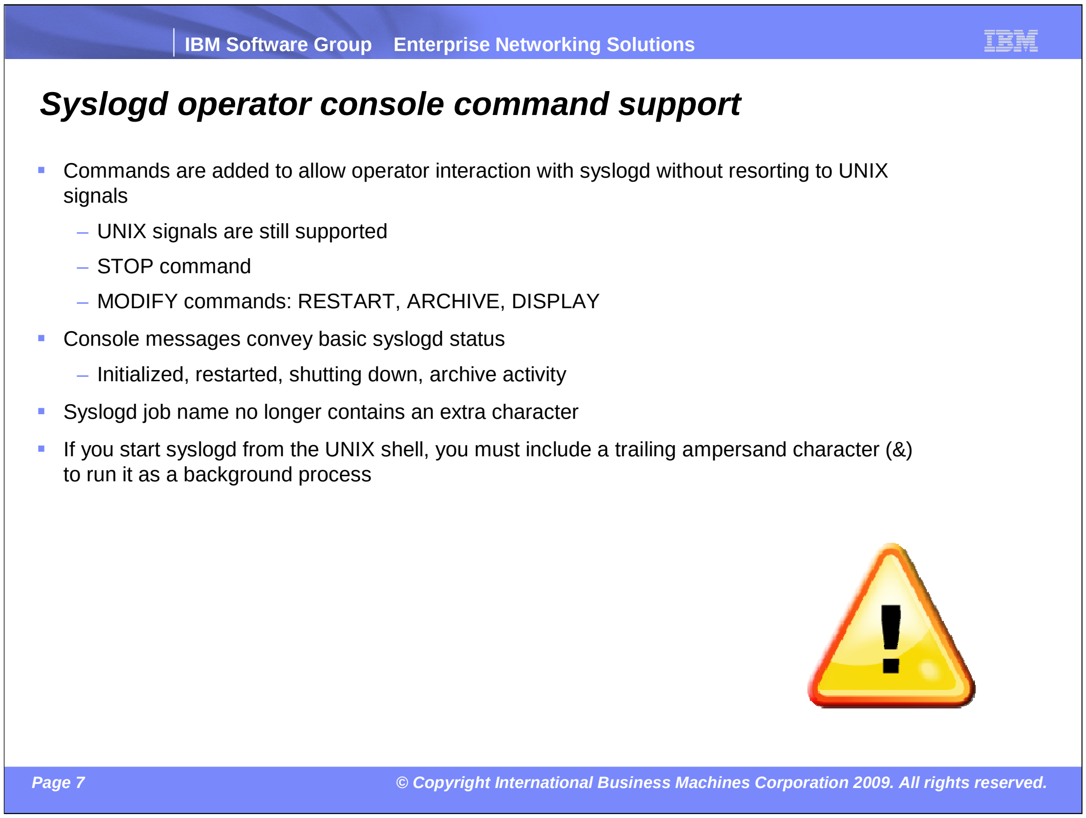

IBM Software Group Enterprise Networking Solutions

Syslogd operator console command support

- Commands are added to allow operator interaction with syslogd without resorting to UNIX signals
  - UNIX signals are still supported
  - STOP command
  - MODIFY commands: RESTART, ARCHIVE, DISPLAY
- Console messages convey basic syslogd status
  - Initialized, restarted, shutting down, archive activity
- Syslogd job name no longer contains an extra character
- If you start syslogd from the UNIX shell, you must include a trailing ampersand character (&) to run it as a background process

Page 7

Syslogd adds both operator command and console message support. The STOP and MODIFY operator commands are supported. The STOP command is equivalent to sending a SIGTERM signal to syslogd. Several MODIFY commands are supported. The MODIFY RESTART command is equivalent to sending a SIGHUP signal to syslogd. The MODIFY ARCHIVE and MODIFY DISPLAY commands are used to perform an on demand archival and display the archive status. Console messages now indicate basic syslogd status, such as initialization complete, restarted due to a MODIFY RESTART command or SIGHUP signal, shutting down due to a STOP command or SIGTERM signal, and archive activity.

In prior releases, an additional character (typically “1”) was added to the syslogd program name if syslogd was started from the UNIX shell. The same was true for the cataloged procedure name if the procedure name was less than eight characters long. Now, the syslogd job name is typically equal to the program name or procedure name, although it can be controlled from the UNIX shell using the _BPX_JOBNAME environment variable. This makes it easier to issue operator commands, because the job name is predictable.

Before V1R11, syslogd automatically ran as a background process when started from the UNIX shell. It no longer does this in V1R11. If you start syslogd from the UNIX shell, you must include a trailing ampersand character to force it to run as a background process. This is especially true if you start syslogd from a shell script such as /etc/rc. If you do not include the trailing ampersand, control will not return to the shell session until syslogd ends.

This is not a concern if you start syslogd from a cataloged procedure.

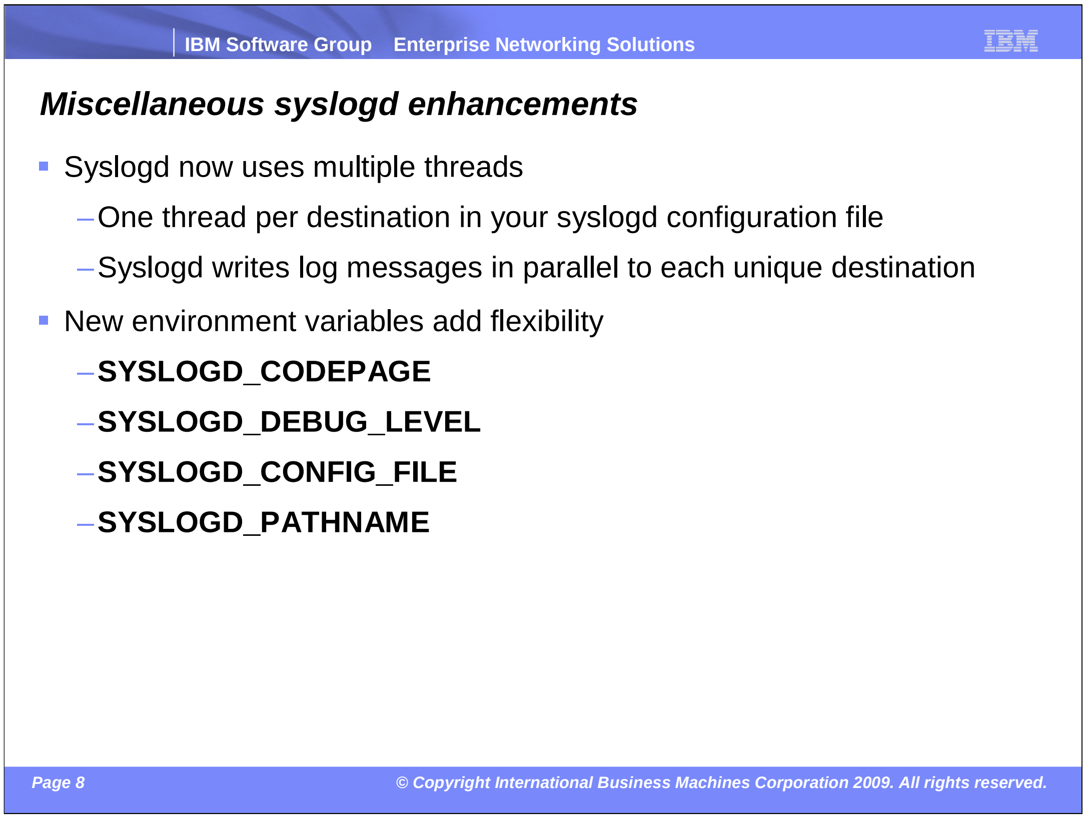

Miscellaneous syslogd enhancements

- Syslogd now uses multiple threads
  - One thread per destination in your syslogd configuration file
  - Syslogd writes log messages in parallel to each unique destination
- New environment variables add flexibility
  - **SYSLOGD_CODEPAGE**
  - **SYSLOGD_DEBUG_LEVEL**
  - **SYSLOGD_CONFIG_FILE**
  - **SYSLOGD_PATHNAME**

The solution to the performance problem is to change syslogd to a multi-threaded application, where log messages for each unique destination are written on a separate thread.

The solution to the remaining problems is to add support for a set of environment variables.

The first environment variable is for you to specify the code page for reading the configuration file. A finite set of single byte EBCDIC code pages is supported. See the ***IP Configuration Guide*** or ***IP Configuration Reference*** for the list of supported code pages. The default code page is IBM-1047.

Another environment variable lets you limit the amount of debug output produced when you specify the -d start option. See the ***IP Configuration Reference*** for the complete list of debug levels. Of particular note is that you can exclude all debug output produced for each individual log message being written - normally a large amount of debug output is produced for each message written to syslogd. With a high volume of logging, excluding this debug output can significantly reduce the amount of debug output produced.

The final two environment variables allow you to specify the values for the -f and -p start options. The -f start option specifies the name of the configuration file. You probably do not need to use the -p start option. When these values are supplied to syslogd using environment variables instead of start options, the length of the start options is significantly reduced. This might help you specify the start options in JCL without running into the 100 character PARM length restriction.

You can specify any of several single byte EBCDIC code page values using the SYSLOGD_CODEPAGE environment variable. The default code page is IBM-1047.

You can specify a debug level with the SYSLOGD_DEBUG_LEVEL environment variable. This is used with the -d start option. The range of debug levels is from 0 (no debugging) to 127 (all debug output is produced), and the default debug level is 127.

You can also use the SYSLOGD_CONFIG_FILE environment variable instead of the -f start option to specify the configuration file name. You should rarely need to use the SYSLOGD_PATH_NAME environment variable, but it can be used instead of the -p start option.

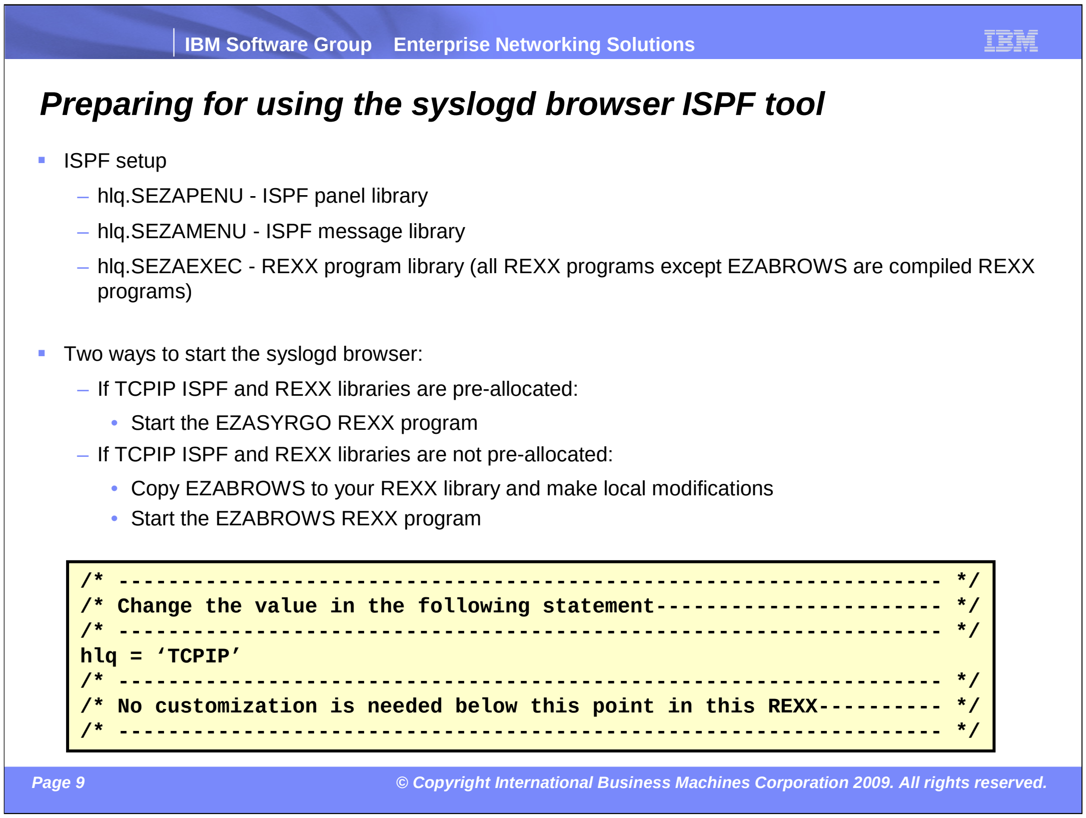

Preparing for using the syslogd browser ISPF tool

- ISPF setup
  - hlq.SEZAPENU - ISPF panel library
  - hlq.SEZAMENU - ISPF message library
  - hlq.SEZAEXEC - REXX program library (all REXX programs except EZABROWS are compiled REXX programs)
- Two ways to start the syslogd browser:
  - If TCPIP ISPF and REXX libraries are pre-allocated:
    - Start the EZASYRGO REXX program
  - If TCPIP ISPF and REXX libraries are not pre-allocated:
    - Copy EZABROWS to your REXX library and make local modifications
    - Start the EZABROWS REXX program

```
/* ------------------------------------------------------------------ */
/* Change the value in the following statement----------------------- */
/* ------------------------------------------------------------------ */
hlq = ’TCPIP’
/* ------------------------------------------------------------------ */
/* No customization is needed below this point in this REXX---------- */
/* ------------------------------------------------------------------ */
```

All components of the syslogd browser have member names that start with EZASYxxx.

z/OS Communications Server delivers ISPF components for panels, messages, and REXX programs. ISPF panels are in hlq.SEZAPENU. ISPF messages are in hlq.SEZAMENU. REXX programs for TSO are in hlq.SEZAEXEC.

You can pre-allocate ISPF and REXX libraries using DD names in your TSO LOGON procedure or TSO LOGON CLIST. hlq.SEZAEXEC is a new z/OS Communications Server system library in z/OS V1R11. It is an FB, 80 library.

If you use the EZABROWS REXX to start the browser, you can copy EZABROWS to a REXX library that is pre-allocated to your TSO environment. You must customize the copied EZABROWS to identify high level qualifier of your z/OS Communications Server ISPF libraries.

EZABROWS uses ISPF LIBDEF commands to add the z/OS Communications Server ISPF Libraries to ISPF (ISPPLIB and ISPMLIB). It uses the TSO ALTLIB command to add the hlq.SEZAEXEC library to TSO. EZABROWS finally starts the syslogd browser (EZASYRGO) using an ISPF SELECT with NEWPOOL, PASSLIB, and NEWAPPL(EZAS).

Syslogd browser entry panel

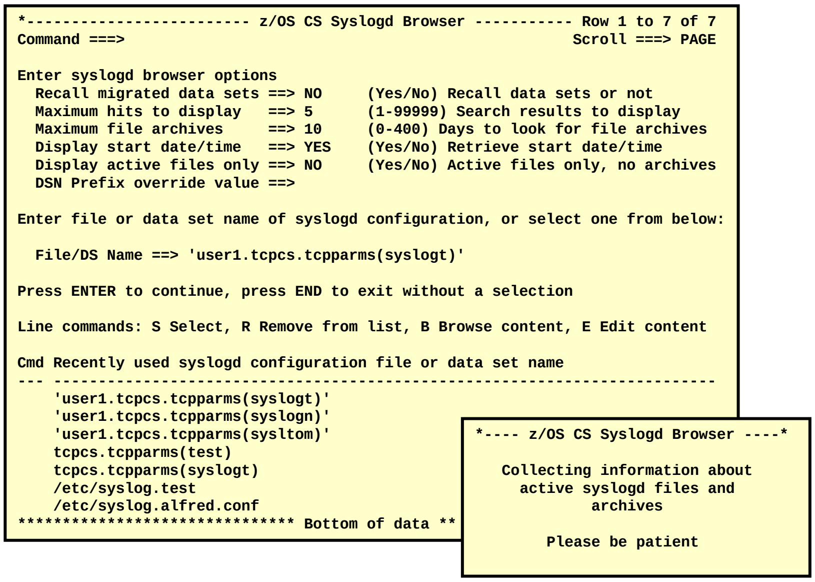

```
*------------------------ z/OS CS Syslogd Browser ----------- Row 1 to 7 of 7
Command ===>                                                Scroll ===> PAGE

Enter syslogd browser options
  Recall migrated data sets ==> NO     (Yes/No) Recall data sets or not
  Maximum hits to display   ==> 5      (1-99999) Search results to display
  Maximum file archives     ==> 10     (0-400) Days to look for file archives
  Display start date/time   ==> YES    (Yes/No) Retrieve start date/time
  Display active files only ==> NO     (Yes/No) Active files only, no archives
  DSN Prefix override value ==>

Enter file or data set name of syslogd configuration, or select one from below:

  File/DS Name ==> 'user1.tcpcs.tcpparms(syslogt)'

Press ENTER to continue, press END to exit without a selection

Line commands: S Select, R Remove from list, B Browse content, E Edit content

Cmd Recently used syslogd configuration file or data set name
--- ------------------------------------------------------------------------
    'user1.tcpcs.tcpparms(syslogt)'
    'user1.tcpcs.tcpparms(syslogn)'
    'user1.tcpcs.tcpparms(sysltom)'
    tcpcs.tcpparms(test)
    tcpcs.tcpparms(syslogt)
    /etc/syslog.test
    /etc/syslog.alfred.conf
******************************* Bottom of data **
```

```
*---- z/OS CS Syslogd Browser ----*

   Collecting information about
     active syslogd files and
          archives

        Please be patient
```

This is the first panel you will see when starting the syslogd browser

Syslogd browser options:

Do you want the browser to access MVS data sets that have been migrated? If you specify NO, you are not able to browse migrated archive data sets.

Specify the maximum number of hits you want displayed as the result of a search operation.

If you specify an archive file name with %-symbols (for day, month, and year), the syslogd browser will look for archives in the same directory as the active z/OS UNIX file. The browser will look for such archives day by day. You use this option to specify the maximum number of days you want to look for such archives.

The display start date/time option is used to control the display of start date and time for each active file and each archive. Set this to NO if you don't need it and want to improve the performance of the syslogd browser initialization.

The Display active files only option controls if the syslogd browser is to be used for browsing the currently active syslogd files only, or if it is to be used for browsing both active syslogd files and archives. Set this to NO if you know you're only going to browse the active syslogd files. It will improve the performance of the syslogd browser initialization.

The DSN Prefix override value overrides the DSNPREFIX keyword in your syslogd configuration file. This option is especially useful if you use system symbols in your DSNPREFIX and want to browse the syslogd files of another LPAR than the one you are logged into.

The browser will save the last 10 syslogd configuration files you have used. For each of those, you can edit, browse, remove from the list, or select the configuration file for use by the browser.

If you have many syslogd UNIX files and archived MVS data sets, it will take a little while for the browser to collect information about all those files and data sets. You can speed the initialization up by either answering NO to display start date and time, or answering YES to display active files only. If you know you are going to look into the active UNIX files only, then there is no need to collect information about archives.

Syslogd destination view

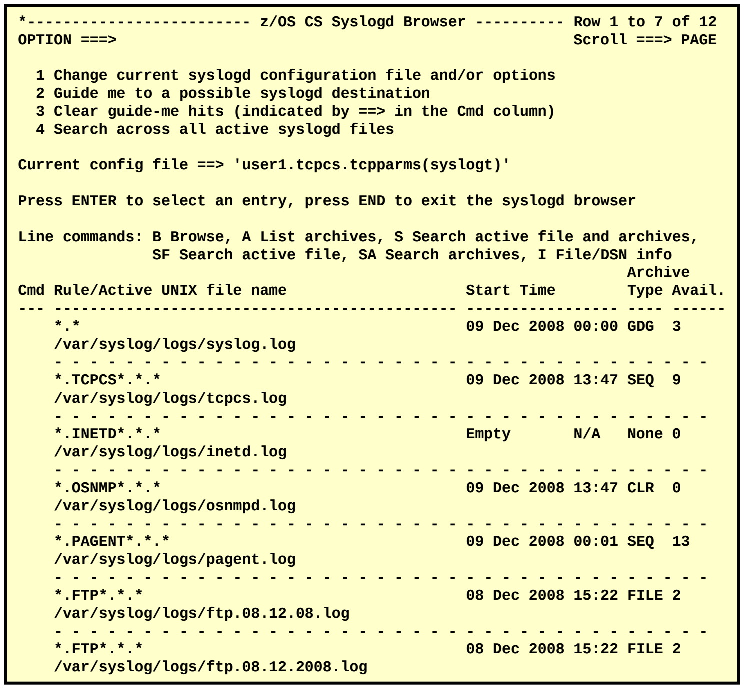

```
*------------------------- z/OS CS Syslogd Browser ---------- Row 1 to 7 of 12
OPTION ===>                                                   Scroll ===> PAGE

  1 Change current syslogd configuration file and/or options
  2 Guide me to a possible syslogd destination
  3 Clear guide-me hits (indicated by ==> in the Cmd column)
  4 Search across all active syslogd files

Current config file ==> 'user1.tcpcs.tcpparms(syslogt)'

Press ENTER to select an entry, press END to exit the syslogd browser

Line commands: B Browse, A List archives, S Search active file and archives,
               SF Search active file, SA Search archives, I File/DSN info
                                                                    Archive
Cmd Rule/Active UNIX file name                        Start Time        Type Avail.
--- --------------------------------------------- ---------------- ---- ------
    *.*                                           09 Dec 2008 00:00 GDG  3
    /var/syslog/logs/syslog.log
    - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
    *.TCPCS*.*.*                                  09 Dec 2008 13:47 SEQ  9
    /var/syslog/logs/tcpcs.log
    - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
    *.INETD*.*.*                                  Empty       N/A   None 0
    /var/syslog/logs/inetd.log
    - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
    *.OSNMP*.*.*                                  09 Dec 2008 13:47 CLR  0
    /var/syslog/logs/osnmpd.log
    - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
    *.PAGENT*.*.*                                 09 Dec 2008 00:01 SEQ  13
    /var/syslog/logs/pagent.log
    - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
    *.FTP*.*.*                                    08 Dec 2008 15:22 FILE 2
    /var/syslog/logs/ftp.08.12.08.log
    - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
    *.FTP*.*.*                                    08 Dec 2008 15:22 FILE 2
    /var/syslog/logs/ftp.08.12.2008.log
```

This panel shows the syslogd destination rules that direct messages to z/OS UNIX files

The files on this panel are the active log files

Page 11

After having parsed a syslogd configuration file, all z/OS UNIX file destinations are selected and all associated available archives are located. This syslogd destination view is the main panel of the syslogd browser interface from which other functions are selected.

Main options:

1: To change which syslogd configuration you're using. This option can also be used to re-initialize with the current syslogd configuration file. This can be useful if an archive occurs while you are using the browser. The new archive file or data sets is not accessible until you re-Initialize.

2: To invoke the guide-me function (help me to find which destination my log messages go to)

3: To clear the indicators that were returned from an invocation of the guide-me function

4: To start a search operation across all the active syslogd destination files (such a search can take some time if the active files are large)

Scrollable section:

The display includes one entry per z/OS UNIX file destination for which the active file can be found. Each entry includes rule, active file name, date and time of the first logged message in the active file, archive type, and number of available archives. For MVS archives that means archives that were found online in the z/OS UNIX file system or in a z/OS catalog. For each entry, several line commands are available to browse the active file, search at various levels, and so on.

Archive types:

None - No archive processing for this file

GDG - Archive done to an MVS generation data set group

SEQ - Archive done to a sequential MVS data set

CLR - No archive. The z/OS UNIX file is cleared during archive processing (the -X option used in the syslogd configuration file)

FILE - z/OS UNIX files based on use of %-symbols

Browse an active syslogd file

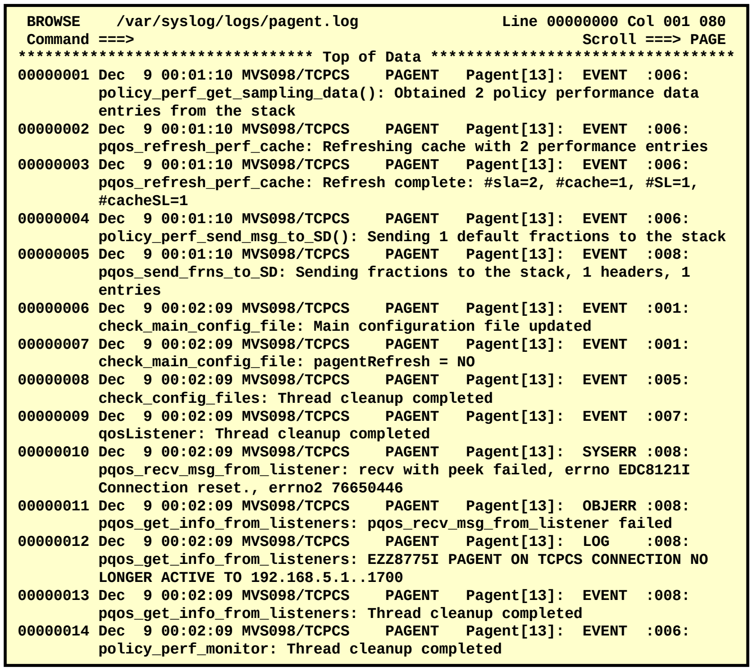

```
 BROWSE    /var/syslog/logs/pagent.log                Line 00000000 Col 001 080
 Command ===>                                                  Scroll ===> PAGE
********************************* Top of Data **********************************
00000001 Dec  9 00:01:10 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :006:
         policy_perf_get_sampling_data(): Obtained 2 policy performance data
         entries from the stack
00000002 Dec  9 00:01:10 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :006:
         pqos_refresh_perf_cache: Refreshing cache with 2 performance entries
00000003 Dec  9 00:01:10 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :006:
         pqos_refresh_perf_cache: Refresh complete: #sla=2, #cache=1, #SL=1,
         #cacheSL=1
00000004 Dec  9 00:01:10 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :006:
         policy_perf_send_msg_to_SD(): Sending 1 default fractions to the stack
00000005 Dec  9 00:01:10 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :008:
         pqos_send_frns_to_SD: Sending fractions to the stack, 1 headers, 1
         entries
00000006 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :001:
         check_main_config_file: Main configuration file updated
00000007 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :001:
         check_main_config_file: pagentRefresh = NO
00000008 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :005:
         check_config_files: Thread cleanup completed
00000009 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :007:
         qosListener: Thread cleanup completed
00000010 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  SYSERR :008:
         pqos_recv_msg_from_listener: recv with peek failed, errno EDC8121I
         Connection reset., errno2 76650446
00000011 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  OBJERR :008:
         pqos_get_info_from_listeners: pqos_recv_msg_from_listener failed
00000012 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  LOG    :008:
         pqos_get_info_from_listeners: EZZ8775I PAGENT ON TCPCS CONNECTION NO
         LONGER ACTIVE TO 192.168.5.1..1700
00000013 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :008:
         pqos_get_info_from_listeners: Thread cleanup completed
00000014 Dec  9 00:02:09 MVS098/TCPCS    PAGENT   Pagent[13]:  EVENT  :006:
         policy_perf_monitor: Thread cleanup completed
```

A normal ISPF browser interface

By entering a ‘B’ for an entry at the destination view, you will see a display of the active UNIX file.

The actual browse window is built using the ISPF BRIF interface, which allows the browser to read only portions of a file or data set into storage at a time.

Long messages are folded into lines that fit the current ISPF screen width.

Normal ISPF FIND command support is available and can be used for simple searches in the file that is being browsed.

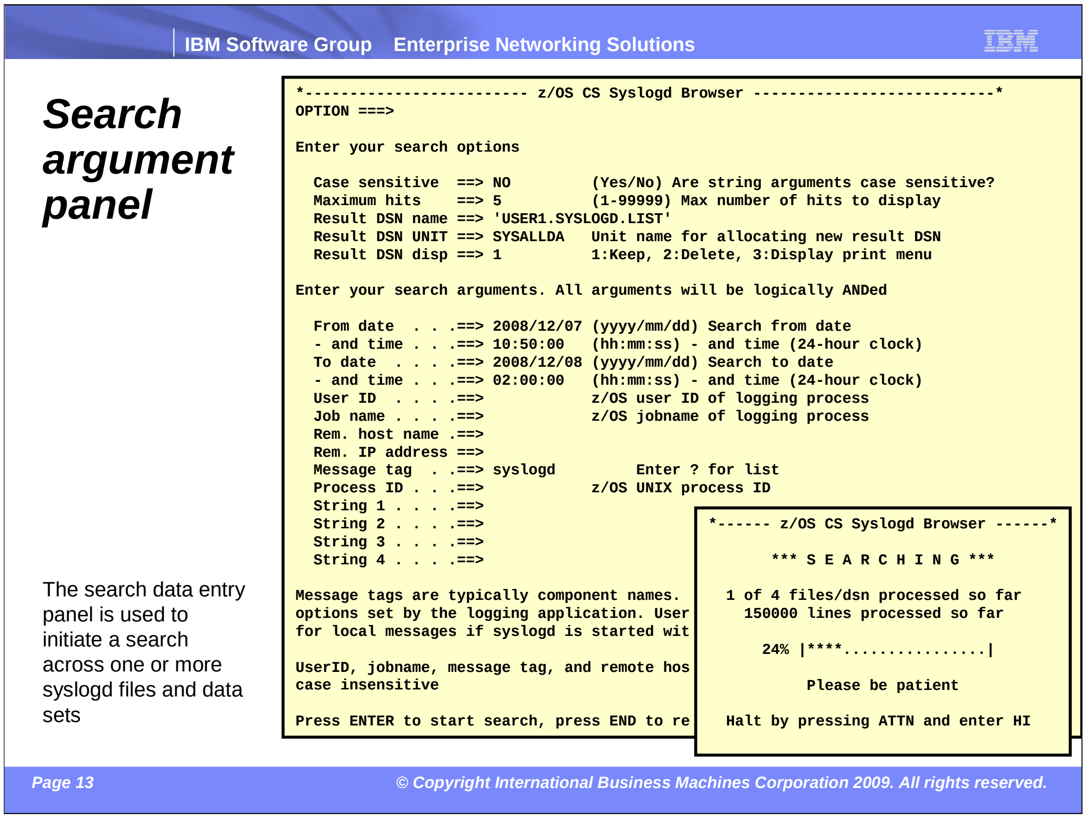

IBM Software Group   Enterprise Networking Solutions

Search argument panel

```
*------------------------- z/OS CS Syslogd Browser ---------------------------*
OPTION ===>

Enter your search options

  Case sensitive  ==> NO          (Yes/No) Are string arguments case sensitive?
  Maximum hits    ==> 5           (1-99999) Max number of hits to display
  Result DSN name ==> 'USER1.SYSLOGD.LIST'
  Result DSN UNIT ==> SYSALLDA    Unit name for allocating new result DSN
  Result DSN disp ==> 1           1:Keep, 2:Delete, 3:Display print menu

Enter your search arguments. All arguments will be logically ANDed

  From date  . . .==> 2008/12/07 (yyyy/mm/dd) Search from date
  - and time . . .==> 10:50:00   (hh:mm:ss) - and time (24-hour clock)
  To date  . . . .==> 2008/12/08 (yyyy/mm/dd) Search to date
  - and time . . .==> 02:00:00   (hh:mm:ss) - and time (24-hour clock)
  User ID  . . . .==>            z/OS user ID of logging process
  Job name . . . .==>            z/OS jobname of logging process
  Rem. host name .==>
  Rem. IP address ==>
  Message tag  . .==> syslogd         Enter ? for list
  Process ID . . .==>            z/OS UNIX process ID
  String 1 . . . .==>
  String 2 . . . .==>                        *------ z/OS CS Syslogd Browser ------*
  String 3 . . . .==>
  String 4 . . . .==>                               *** S E A R C H I N G ***

Message tags are typically component names.     1 of 4 files/dsn processed so far
options set by the logging application. User      150000 lines processed so far
for local messages if syslogd is started wit
                                                       24% |****................|
UserID, jobname, message tag, and remote hos
case insensitive                                            Please be patient

Press ENTER to start search, press END to re    Halt by pressing ATTN and enter HI
```

The search data entry panel is used to initiate a search across one or more syslogd files and data sets

Page 13

© Copyright International Business Machines Corporation 2009. All rights reserved.

By entering one of the S-commands for an entry in the destination view, you will get to the search interface.

Search options govern the search operation. The result data set name can be an existing data set or it is allocated. After the search, you can keep the result data set, delete it, or have a standard ISPF print dialog displayed.

All search arguments are optional. All specified search arguments are logically ANDed together.

> If you specify a specific search criteria and a message has no value for that criteria, the message is considered a non-hit. Example: if you specify a user ID, but a message has no user ID, such a message is not considered a hit. This can be the case if syslogd has not been started with the –u option.
>
> If you say NO to ‘Case sensitivity’ and search for a string of ‘abc’, messages with ‘ABC’, ‘abc’, ‘Abc’ are considered hits. If you specify YES to ‘Case sensitivity’ and search for a string of ‘abc’, then only messages with the exact matching case ‘abc’ are considered hits. Note the case sensitivity option only applies to the four free-form string fields.

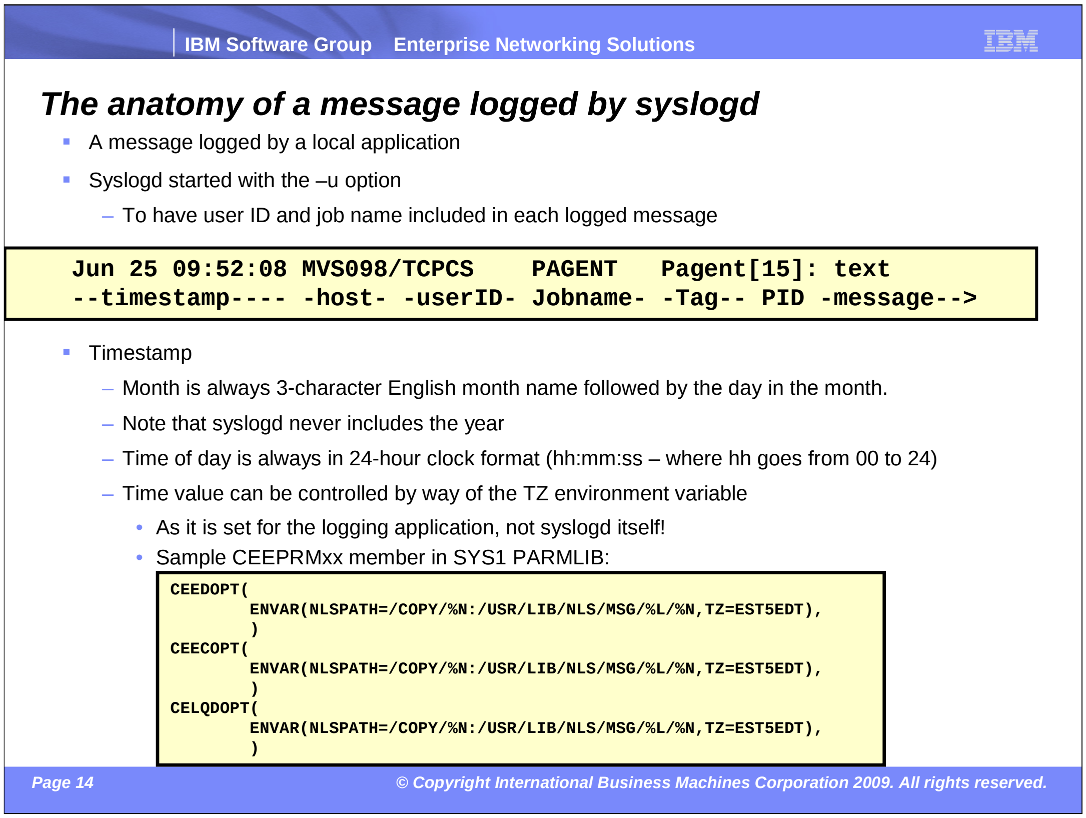

The anatomy of a message logged by syslogd

- A message logged by a local application
- Syslogd started with the –u option
  - To have user ID and job name included in each logged message

```
Jun 25 09:52:08 MVS098/TCPCS    PAGENT   Pagent[15]: text
--timestamp---- -host- -userID- Jobname- -Tag-- PID -message-->
```

- Timestamp
  - Month is always 3-character English month name followed by the day in the month.
  - Note that syslogd never includes the year
  - Time of day is always in 24-hour clock format (hh:mm:ss – where hh goes from 00 to 24)
  - Time value can be controlled by way of the TZ environment variable
    - As it is set for the logging application, not syslogd itself!
    - Sample CEEPRMxx member in SYS1 PARMLIB:

```
CEEDOPT(
        ENVAR(NLSPATH=/COPY/%N:/USR/LIB/NLS/MSG/%L/%N,TZ=EST5EDT),
        )
CEECOPT(
        ENVAR(NLSPATH=/COPY/%N:/USR/LIB/NLS/MSG/%L/%N,TZ=EST5EDT),
        )
CELQDOPT(
        ENVAR(NLSPATH=/COPY/%N:/USR/LIB/NLS/MSG/%L/%N,TZ=EST5EDT),
        )
```

To have all messages logged with your local time, set the TZ environment variable in the CEEPRMxx PARMLIB member. You need to define the TZ environment variable for all three LE option sets (CEEDOPT, CEECOPT, and CELQDOPT).

To support search across new year, the browser applies this logic to all time stamps. If message month.date is later than the current month.date, the year is assumed to be the previous year, otherwise the year is assumed to be the current year. Example: if today is Jan 4 2009, and a message with a date of Dec 30 is processed, the year of the message is set to 2008. Another message with the date of Jan 1 is processed, the year of the message is set to 2009. This logic allows browsing a year back in time across a new year, but not more than one year.

For local messages, host name is the host name that is configured in TCPIP.DATA.

For remote messages, host name is the DNS name of the remote host or the IP address of the remote host where the IP address is included in parenthesis: (10.1.2.3). Syslogd will resolve remote IP addresses to host names only when you start syslogd with the –x option.

User ID and job name are available for local messages when syslogd has been started with the –u flag. The message tag is an optional character string that can be passed by the logging application and generally identifies the application or component that created this log message.

The process ID is included if the logging application specifies the LOG_PID option on its open_log call. The PID is always enclosed in square brackets and those square brackets are always encoded according to IBM-1047 (the square brackets in the logged messages are not subject to any locale configured by the installation).

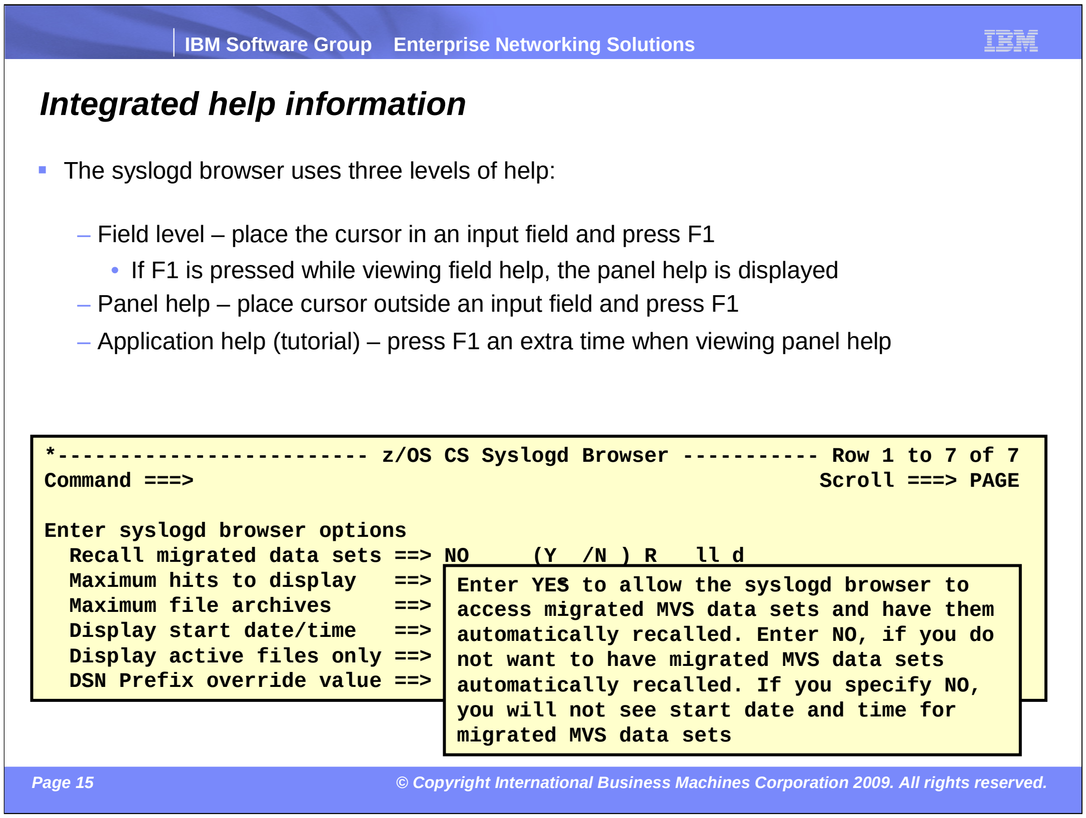

IBM Software Group | Enterprise Networking Solutions

Integrated help information

- The syslogd browser uses three levels of help:
  - Field level – place the cursor in an input field and press F1
    - If F1 is pressed while viewing field help, the panel help is displayed
  - Panel help – place cursor outside an input field and press F1
  - Application help (tutorial) – press F1 an extra time when viewing panel help

```
*------------------------ z/OS CS Syslogd Browser ----------- Row 1 to 7 of 7
Command ===>                                               Scroll ===> PAGE

Enter syslogd browser options
  Recall migrated data sets ==> NO     (Y  /N ) R   ll d
  Maximum hits to display   ==> Enter YES to allow the syslogd browser to
  Maximum file archives     ==> access migrated MVS data sets and have them
  Display start date/time   ==> automatically recalled. Enter NO, if you do
  Display active files only ==> not want to have migrated MVS data sets
  DSN Prefix override value ==> automatically recalled. If you specify NO,
                                you will not see start date and time for
                                migrated MVS data sets
```

Page 15

There is extensive help built into the syslogd browser. All input fields have field-level help associated with them. If in doubt about an input field, press F1. If more help is needed press F1 again and you will see the help for the panel in which the input field occurs. If yet more help is needed, press F1 a third time, and the tutorial is presented.

The tutorial consists of several sections in which you can select directly from the tutorial index panel, or you can browse through all of the tutorial panels by repeatedly hitting the ENTER key.

Use of the syslogd browser is optional.

The main objective is to display and search syslogd messages that are generated by z/OS.

Most capabilities will work with log messages received from other platforms.

Some inconsistencies might exist as to the format of the actual messages, which can confuse the browser when searching.

When browsing syslogd files from another system than where syslogd is running, you need to consider:

Make sure that this system has access to the z/OS UNIX file system of the other system.

If system symbols are used in the data set prefix option in the syslogd configuration file, consider using the override option on the initial browser panel.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

z/OS

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.