IBM eServer™

Systems management

syslogd and netstat

ebusiness on demand software

Agenda - Systems management

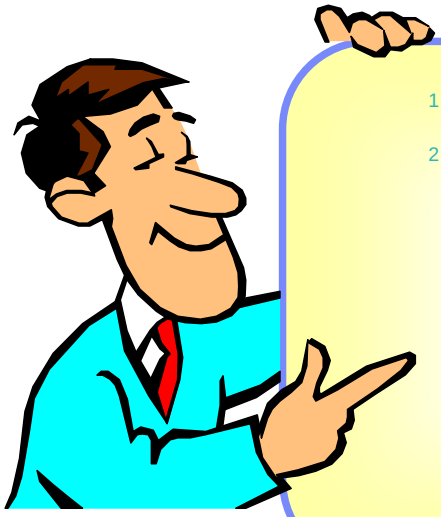

1 **Linux® message integration (syslogd)**

2 **New netstat options**

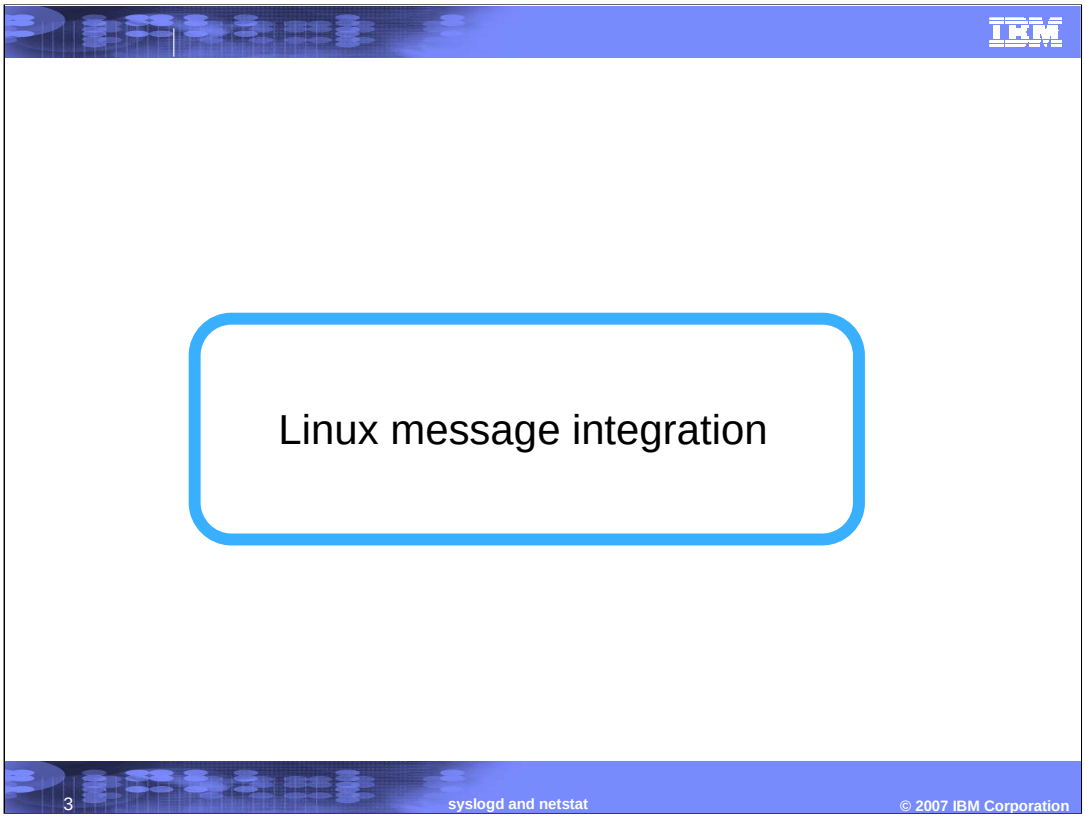
IBM
Linux message integration
3
syslogd and netstat
© 2007 IBM Corporation

Background information - syslogd

- **syslogd is the syslog daemon**
  - Used by UNIX® System Services processes to write log messages to files and other destinations
  - The default configuration file for syslogd is /etc/syslog.conf
  - The configuration file contains filter rules and message destinations
- **syslogd rules**
  - Rule conditions select messages based on:
    - message facility (for example: daemon, user, cron)
    - message priority (for example: warn, error, info)
    - userid of the USS process generating the message
    - jobname of the USS process generating the message
    - multiple combinations of the above
  - Rule destination may be:
    - unix file system file
    - the MVS™ operator console (/dev/console)
    - network destination specified as a host name or IPv4 or IPv6 address (UDP port 514)
    - unix socket
  - syslogd may be started in "normal" mode
    - processes messages from local applications and network applications
  - syslogd may be started in "local only" mode with the -i option
    - process only messages from local application and ignores network messages

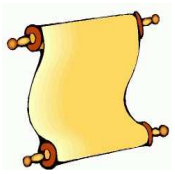

Sample syslogd environment

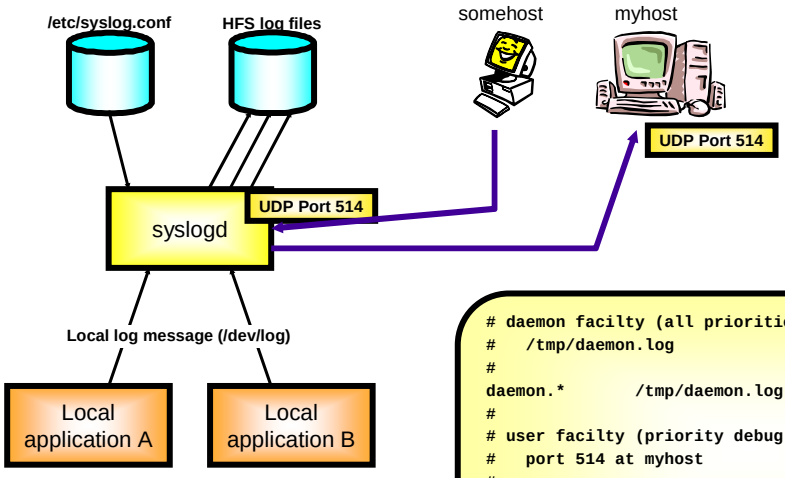


```
# daemon facilty (all priorities) go to
#   /tmp/daemon.log
#
daemon.*       /tmp/daemon.log
#
# user facilty (priority debug and higher) go to
#   port 514 at myhost
#
user.debug     @myhost
#
# messages from syslogd itself go to console
syslogd.*.daemon.*    /dev/console
```

Linux message integration

- **System z customers are deploying functions/workloads on Linux on System z**
  - Some of these functions/workloads may interact with z/OS® images or perform functions on behalf of z/OS. For example:
    - Middle-tier servers such as Web Servers or Application Servers
    - Communications Controller for Linux (CCL)
  - In these scenarios it is often desirable to be able to integrate certain key messages from these Linux hosts into the z/OS environment
    - Allow z/OS operator to monitor key events that occur on the Linux hosts that may have an effect on z/OS operations as well
- **syslogd log messages can now be collected from numerous network sources including Linux hosts and can be filtered to log to the intended destination based on the source IP address or host name**
- **log messages from network hosts can be written to the MVS operations log (operlog)**
  - operlog can be used in place of or in addition to MVS syslog (console log)
  - in a sysplex environment operlog can be configured as a log stream in the coupling facility
    - provides a single sysplex-wide consolidated message log that contains z/OS generated messages and syslogd messages
  - better performance than writing to /dev/console
    - You may still want to write some messages to /dev/console for automation purposes
- **Performance of syslogd is improved**
  - a local-only and a network-only instance may be run concurrently
  - new command option (-x) improves performance by avoiding IP address-to-host name resolution for network log messages

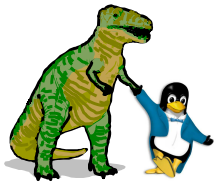

syslogd - start command options

- **syslogd [-f conffile] [-m markinterval] [-p logpath] [-c] [-d] [-i] [-n] [-x] [-u] [-?]**
  - -i option starts syslogd in local-only mode (option exists before V1R8)
  - -n option starts syslogd in network-only mode (new option in V1R8)
  - -x option causes syslogd to avoid resolver calls for converting IP addresses to host names.
- **It is now possible to start two instances of syslogd**
  - One instance in local only mode (-i option)
  - One instance in network only mode (-n option)
- **If you use both local and network logging, IBM recommends that you use two instances of syslogd**
  - helps ensure that local syslogd logging is not adversely affected by the amount of remote messages being forwarded to z/OS

Sample dual-syslogd environment

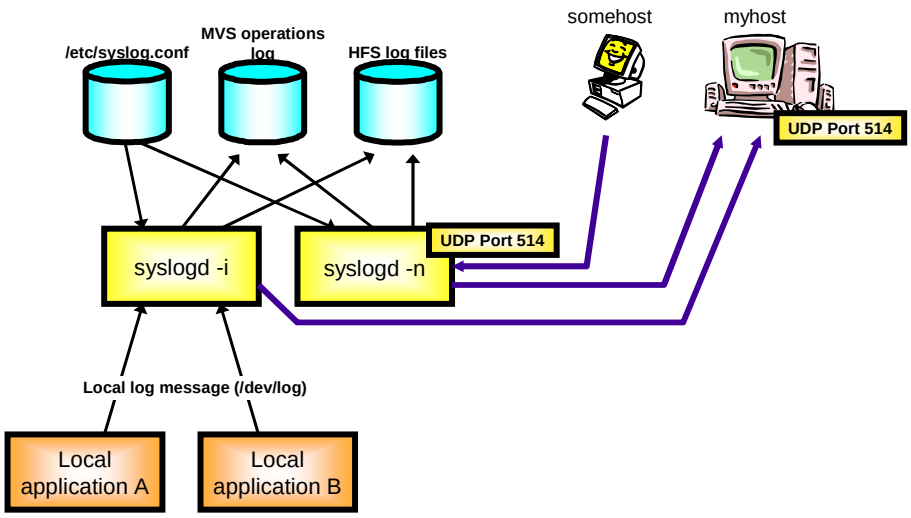

new options in the syslogd configuration file

- **syslogd configuration file is by default /etc/syslog.conf**
- **New options for filters in the condition part of rules in the syslogd configuration file**
  - An IP address (IPv4 or IPv6) or host name may now be used
  - If an IP address is used, an optional prefix length may be specified

```
(host1.xyz.com).*.*                 /tmp/syslogd/host1.log
(192.168.0.1/24).daemon.info        /tmp/syslogd/host2.log
```

- **New option for destination in syslogd rules**
  - MVS operations log (operlog) can now be specified as a destination

```
user1.job1.daemon.*                 /dev/operlog
(host1.xyz.com).daemon.debug        /dev/operlog
```

Things to think about

- **At most, you may run two instances of syslogd**
  - one with -i option
  - one with -n option

- **The -x option can provide better performance when syslogd is processing log messages from remote syslog daemons**
  - when -x is not used, syslogd makes resolver calls to attempt to determine the host name associated with the source IP address of the network syslogd message
    - If found, the host name is added to the prefix information when the message is logged
    - If not found, the IP address is added to the prefix information when the message is logged
    - the host name lookup can be slow - especially if resolver has to search external DNS servers
  - when -x is used, syslogd does not make resolver calls to determine the host name
    - if syslogd "knows" the host name associated with the IP address (from the syslogd rule) then the host name will be added to the prefix information
    - if syslogd does not know the host name, the IP address is added to the prefix information

- **The use of a Dynamic VIPA address on z/OS has many advantages**
  - configure remote syslog daemons to send to the DVIPA address
  - if one image in a sysplex fails (or must be taken down for maintenance), the DVIPA can be moved to another image
    - remote syslog daemons continue to forward messages to the same address
    - when properly configured, the z/OS syslogd on the second image can continue to log to the same sysplex-wide MVS operations log stream

A few more things to think about

- **Network syslogd messages are delivered over UDP**
  - UDP does not guarantee the delivery of messages
  - under some conditions messages may not be delivered to z/OS
  - use of z/OS automation on these messages has limitations

- **IPSEC should be considered for protecting the syslog daemon's UDP port 514 when running in normal or network-only mode**
  - remote syslogd daemon hosts must also use IPSEC
  - ensures that remote hosts are authorized to send messages to the z/OS syslog daemon
  - can ensure data privacy (using encryption) if needed


IBM
New netstat options
12
syslogd and netstat
© 2007 IBM Corporation

Netstat changes in z/OS V1R8 - overview

- **To easier identify listening sockets, support for a SERVER modifier to the Netstat ALL/-A and Netstat COnn/-c reports has been added**
  - When SERVER is specified only TCP connections in listen state are displayed
- **To increase ability to better limit a netstat report to a specific connection, support for a combined IP address and port filter (the IPPort/-B filter) has been added to the the Netstat ALL/-A, ALLConn/-a, COnn/-c, SOCKets/-s, TELnet/-t, VCRT/-V, and VDPT/-O reports**
  - Display connections that match the IP address and port number
  - Allow the IP address to match either the local or remote IP address
  - For the VCRT/-V report the IP address can be a source, destination IP, or destination XCF address
- **To avoid having to browse the TCP/IP profile to see which servers were autologged, the Netstat CONFIG/-f report has been changed to:**
  - Display the AUTOLOG list specified in the z/OS Communications Server profile
- **To reduce the number of output lines from a netstat port list report when many ports have been reserved, the Netstat PORTList/-o report has been changed to:**
  - Display the reserved ports specified in the PORTRANGE statement more concisely
- **To limit a Netstat SLAP report to only list the currently active policies, as new ACTIVE modifier to the Netstat SLAP/-J report has been added**
  - Displays only the policies that are active

The SERVER modifier

- **Server modifier is supported on the Netstst ALL and the Netstat CONN reports.**
- **Netstat report will only include those socket end points that are in listen state (the TCP server sockets)**

```
netstat CONN SERVER

MVS TCP/IP NETSTAT CS V1R8       TCPIP Name: TCPCS          08:20:47
User Id  Conn     State
-------  ----     -----
BPXOINIT 0000001E Listen
  Local Socket:   0.0.0.0..10007
  Foreign Socket: 0.0.0.0..0
FTPD1    0000001F Listen
  Local Socket:   ::..21
  Foreign Socket: ::..0
SYSLOGD5 00000028 Listen
  Local Socket:   0.0.0.0..7
  Foreign Socket: 0.0.0.0..0
```

The IPPORT filter

- **Available on the Netstat ALL, ALLConn, COnn, SOCKets, TELnet, VCRT, and VDPT**
  - For ALL, ALLConn, COnn, TELnet the filter can be either local or remote IP address
  - For SOCKets the filter can the be IP address to which the socket is bound
  - For VCRT the filter can be source, destination, or destination XCF IP address
  - For VDPT the filter can be destination or destination XCF IP address
  - Allow up to six IP address and port filter pairs
  - No wildcards (*) or masks allowed for the IP address
- **IP address and port number separated by plus ("+") sign**

```
netstat -A -B 9.37.218.154+21

MVS TCP/IP NETSTAT CS V1R8       TCPIP Name: TCPCS           08:49:40
Client Name: FTPD1                    Client Id: 00000074
  Local Socket: ::ffff:9.42.105.47..21
  Foreign Socket: ::ffff:9.37.218.154..1395
....
```

Netstat CONFIG to include autolog information

```
netstat CONFIG

MVS TCP/IP NETSTAT CS V1R8       TCPIP Name: TCPCS           10:48:06
TCP Configuration Table:
...
UDP Configuration Table:
...
IP Configuration Table:
...
IPv6 Configuration Table:
...
SMF Parameters:
...
Global Configuration Information:
...
Network Monitor Configuration Information:
...
Autolog Configuration Information: Wait Time: 0060
ProcName: FTPD      JobName: FTPD
  ParmString:
ProcName: RPCBIND   JobName: RPCBIND
  ParmString: TRACE='-d'
```

PORTLIST report to better handle port ranges

- **Display z/OS Communications Server profile PORTRANGE statement in a single statement**

```
netstat -o

MVS TCP/IP NETSTAT CS V1R8       TCPIP Name: TCPCS          10:36:06
Port# Prot User     Flags    Range
----- ---- ----     -----    -----
...
00021 TCP  FTPD1    DAU
...
05000 TCP  TESTSYS  DAR      05000-05199
00007 UDP  MISCSERV DAB
```

Things to think about

- **Only the ususal:**
  - There are migration concerns for those Miscellaneous changes if customers have programs that collect data from the Netstat report output. The programs may need to be updated to handle the changes.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM MVS z/OS

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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.

Information is provided "AS IS" without warranty of any kind. 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.