

An Introduction to Syslog

Rainer Gerhards
Adiscon

What is Syslog?

- The “heterogeneous network logging workhorse”
- a system to emit/store/process meaningful log messages
- both a communications protocol as well as a set of actual programs and libraries
- designed for system admins to keep machines up and running
- Great way to get heterogeneous data into a single data repository

What is Syslog being used for?

- Troubleshooting Routers/Firewalls/Devices
 - during installation
 - in problem situations
- Intrusion Detection
- Operations Management
- (Long Term) Auditing
- tracking user and admin activity

Where is Syslog Available?

- Native to all flavors of UNIX and Linux
- Third parties on Windows and other OSs
- Syslog Data is sent by almost all network equipment:
 - Routers
 - Firewalls
 - Switches
 - And other “active” boxes

Syslog Roles

- Device – generates message (may be a program)
- Collector – receives and optionally stores messages. Commonly known as syslog daemon or server.
- Relay – receives and forwards message
- Sender – anyone who sends syslog messages (device & relay)
- Receiver – anyone who receives syslog messages (relay & collector)

Syslog – Protocol vs. Application

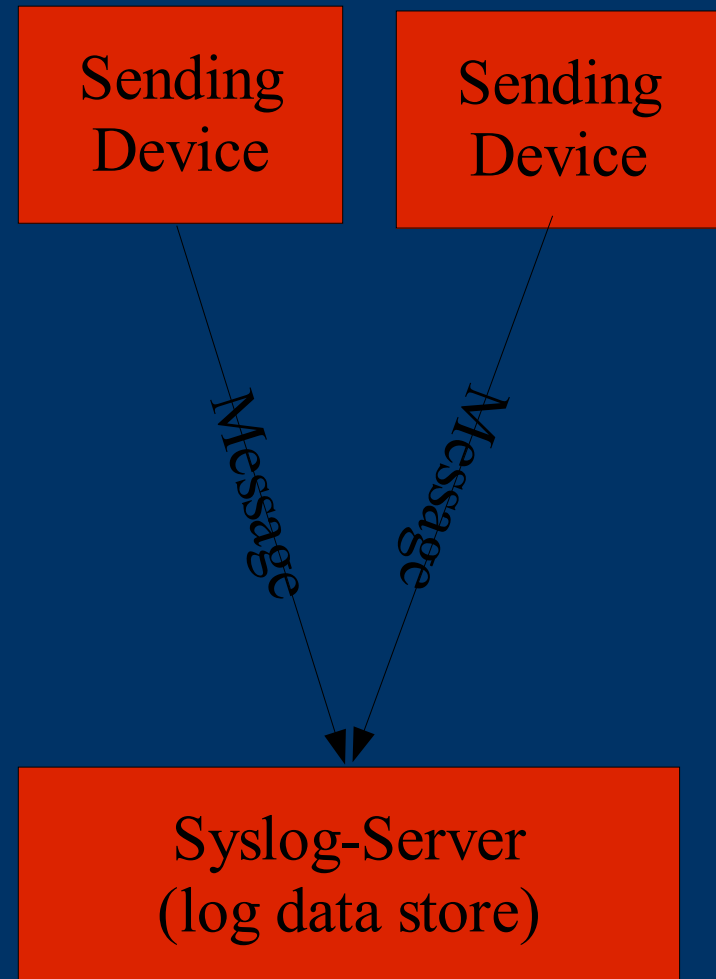
- Some confusion stems from the fact that the name “syslog” is often used both for the protocol and several “syslog applications”.
- The actual syslog applications vary greatly (performance, stability, security... - as usual)
- This presentation talks about protocol issues, which should be common across all applications.

Passive Nature of Syslogd

- “just like TV” - only records what is sent to it
- The syslog receiver is passive and waits for incoming input – it does **not** actively gather it (but it needs to be configured to accept network messages)
- The senders must be pointed to send data to the syslog receivers

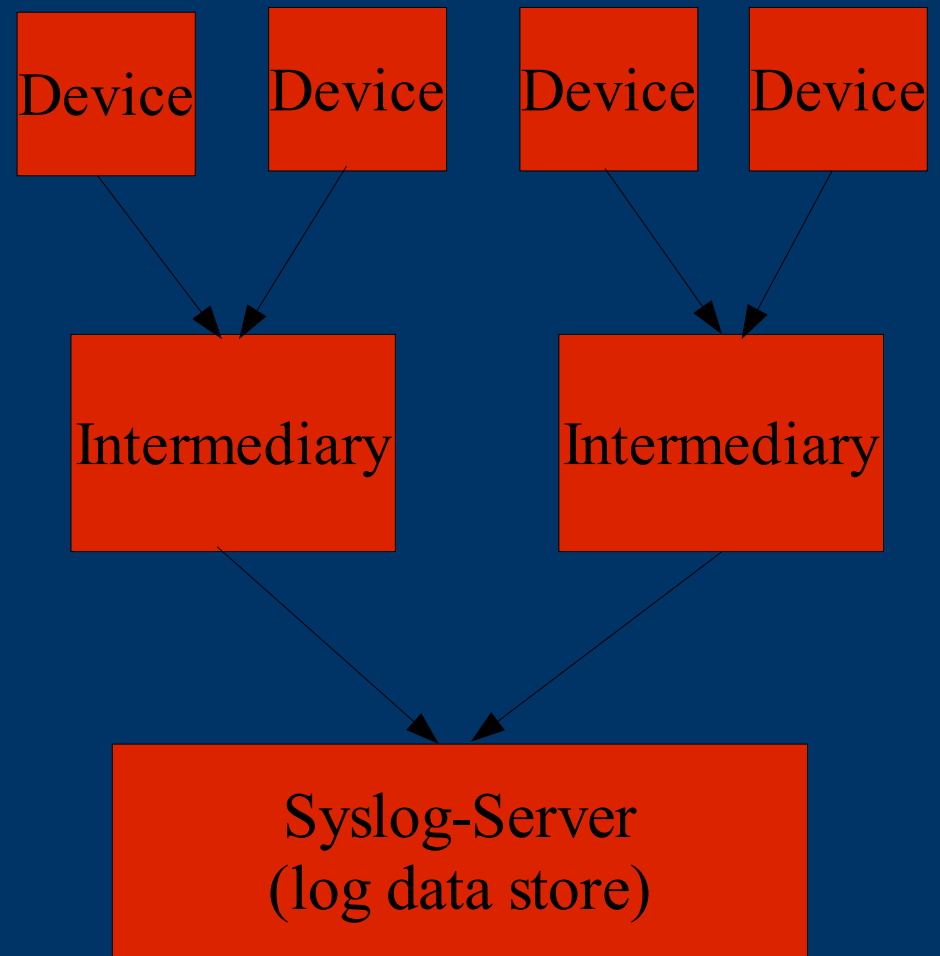
Typical Syslog Setup - Small

- One or multiple senders send data to a single, central syslog server
- Server stores (e.g. analysis) and/or may act on received data
- Typically all on one LAN.



Typical Syslog Setup - Larger

- Senders send data to intermediary syslog servers
- Intermediary receives, may filter and forward message to final destination
- Typically found in larger enterprises.



Fields you Should Know

- Fields commonly found
 - Facility
 - Severity
 - Timestamp
 - Host
 - Tag
 - Message
- Not all of them always present – largely depending on implementation and configuration
- Fields can be used for filtering and processing

Syslog - Facility

- numerical indicator (0..23) of sender component/application that sent the message
- Each of the values is assigned a friendly name
- For example, on Unix systems, facility 0 is traditionally used for kernel messages, 2 for mail subsystem messages, and values 16-23 are reserved for local customization.
- Other implementations may use it differently.
- In general: a more-or-less user-assigned value to be used for filtering at the server side.

Syslog - Severity

- Another integer value with a range from 0 to 7
- Indicates the severity / importance of a message (thus the name ;-)) - but some implementations do not properly support this
- The lower the number, the more important the message is deemed to be (e.g. 0 is “emergency” and 7 “debug”)
- This too is most often be used to filter incoming messages (e.g. Where to store, page users, ...)

Syslog - Timestamp

- If present, most often a timestamp with just the date and day of month, hour, minutes and seconds
- Most often no time zone, year or better-than-second resolution
- Often **wrong!** ... due to out-of-sync internal device clocks (e.g. Clock always starts at Jan, 1st 1997 after power up)
 - If supported (by device), plan for NTP or similar mechanism to solve this.
- Improved in upcoming standards

Syslog - Host

- Name or IP-Address of the sender
- Sometimes missing, sometimes present, sometimes meaningless or invalid (depending on configuration)
- Often duplicate if multiple networks are being monitored (e.g. a service provider monitoring customer networks)
- Intention is to provide the name of the original sender when passing through syslog relays.

Syslog - Tag

- A short ID made up of printable characters
- Most often identifies the process/device sending the message
- Sometimes contains a static process name and a dynamic process ID (changing after each reboot)
- Actual format is **very** different between implementations

Syslog - Message

- Confusing – a “syslog message” contains a “message”
- The textual part after all the “header” fields (facility, severity, ..., tag)
- Often non-structured clear-text intended for human recipient
- Sometime better machine-parsable form (but then less human readable)

Syslog Limitations

- 1K message size
- Not the most secure protocol in this world...
- Reliability issues
- As the content format is not standardized, there is a large variety of message contents (but almost all are human readable).
- Side note: “human readable” does not necessarily mean that the message makes sense in all cases...

Can you Trust Syslog?

- Good enough for many cases
- BUT
 - Not 100% reliable delivery (UDP based)
 - Sender address can be easily faked
 - Open to replay attacks
 - ... and more
- If you deploy syslog, spend some time thinking on how to do it securely!
- Again, newer standards provide better protection

Syslog Reliability

- “traditional” syslog works OK over a non-congested LAN
- Expect some packet loss when the network is congested, the sender or receiver is busy, the network is slow (WAN) or there is a large burst of syslog traffic
- Plan for some minor packet loss in your analyzer
- Experience shows it is reliable enough to work
- Newer developments promise better reliability

What if I need more Reliability?

- Think about non-standard extensions, e.g. raw tcp syslog (some vendors do this)
- Try to implement the new standards (RFC3195!) as soon as possible
- Be sure to think about the “big picture” in a larger network
- Make sure the central syslog receiver is monitored and backed up! Use a redundant configuration if it is mission-critical. (reliability issue #1!)

Syslog Standardization

- Been around for long, but only recently standardized (beginning in 1999)
- To learn about standardization process, visit IETF syslog-sec WG homepage at <http://www.employees.org/~lonvick/index.shtml>
- NO standardization on the actual content so far.
- Still a large variety of “interpretations”
- So far, few if any implementations of the new standards. They will appear over time...

Questions?

- My page with further links and information to syslog:

<http://www.monitorware.com/en/topics/syslog/>

- Just mail me ;-)

rgerhards-at-adiscon.com

(I am sometimes swamped, so please bear a little with me if a reply takes time – but re-send when you do not hear back within 1 week – this is an indication your message ended up in the SPAM filter...)

Credits and Copyright

- **Special thanks go to Tina Bird of www.loganalysis.org for reviewing this presentation.**
- Copyright © 2004 by Rainer Gerhards and Adiscon. All rights reserved.
- Reprint and Distribution require a prior written permission from either Adiscon or Rainer Gerhards.
- This work is provided “as is” without warranties of any kind. Use at your own risk.