

 **ONE IDENTITY™**

Syslog-ng 101, part 2: Basic concepts

Peter Czanik

What is syslog-ng?

- Logging: recording events, for example:

```
Jan 14 11:38:48 linux-0jbu sshd[7716]: Accepted  
publickey for root from 127.0.0.1 port 48806 ssh2
```

- syslog-ng: an enhanced logging daemon with a focus on portability and high-performance central log collection. Originally developed in C.

Why central logging?

Ease of use

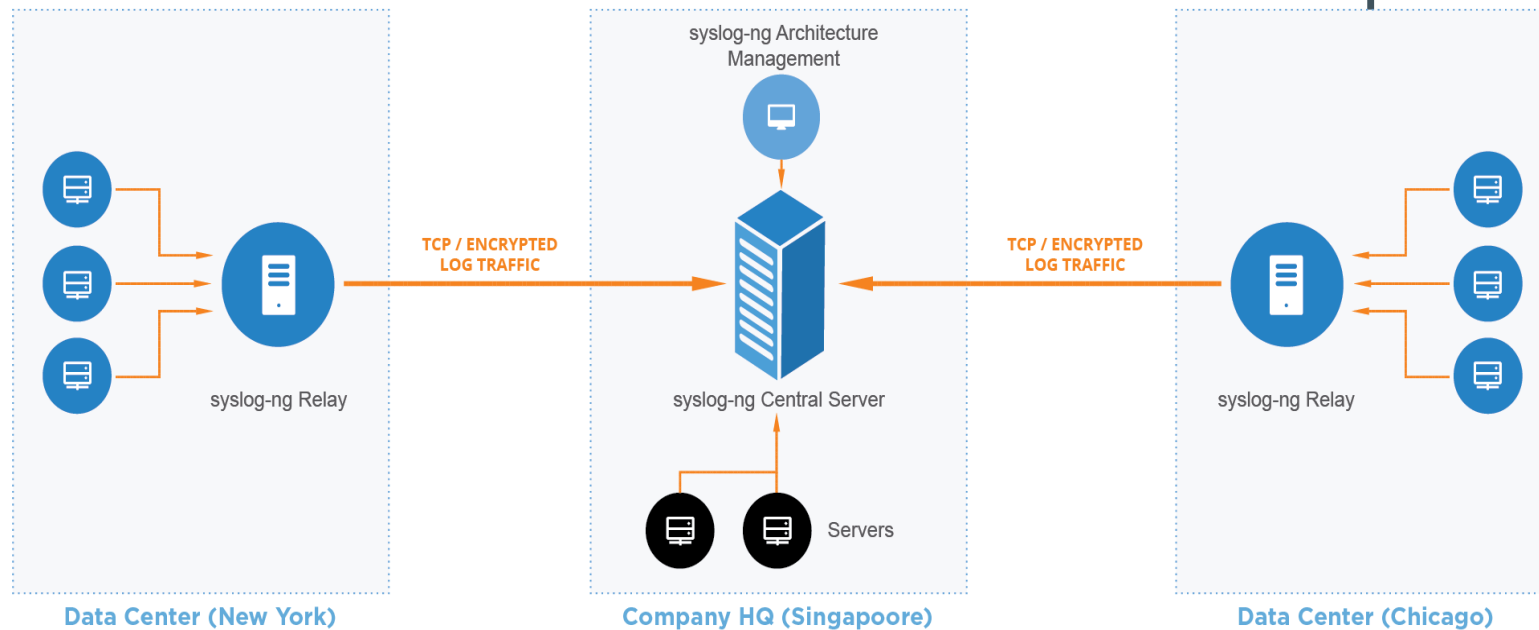
One place to check instead of many

Availability

Even if the sender machine is down

Security

Logs are available even if sender is compromised



The four major roles of syslog-ng

- Collect
- Process
- Filter
- Store (or forward)

Role: data collector

Collect system and application logs together: contextual data for either side

- A wide variety of platform-specific sources:
 - /dev/log & Co., Journal, Sun streams
- Receive syslog messages over the network:
 - Legacy or BSD (RFC3164) and new (RFC5424), UDP/TCP/TLS
- Logs or any kind of text data from applications:
 - Through files, sockets, pipes, application output, etc.
- Python source: Jolly Joker
 - HTTP server, Kafka source, etc.

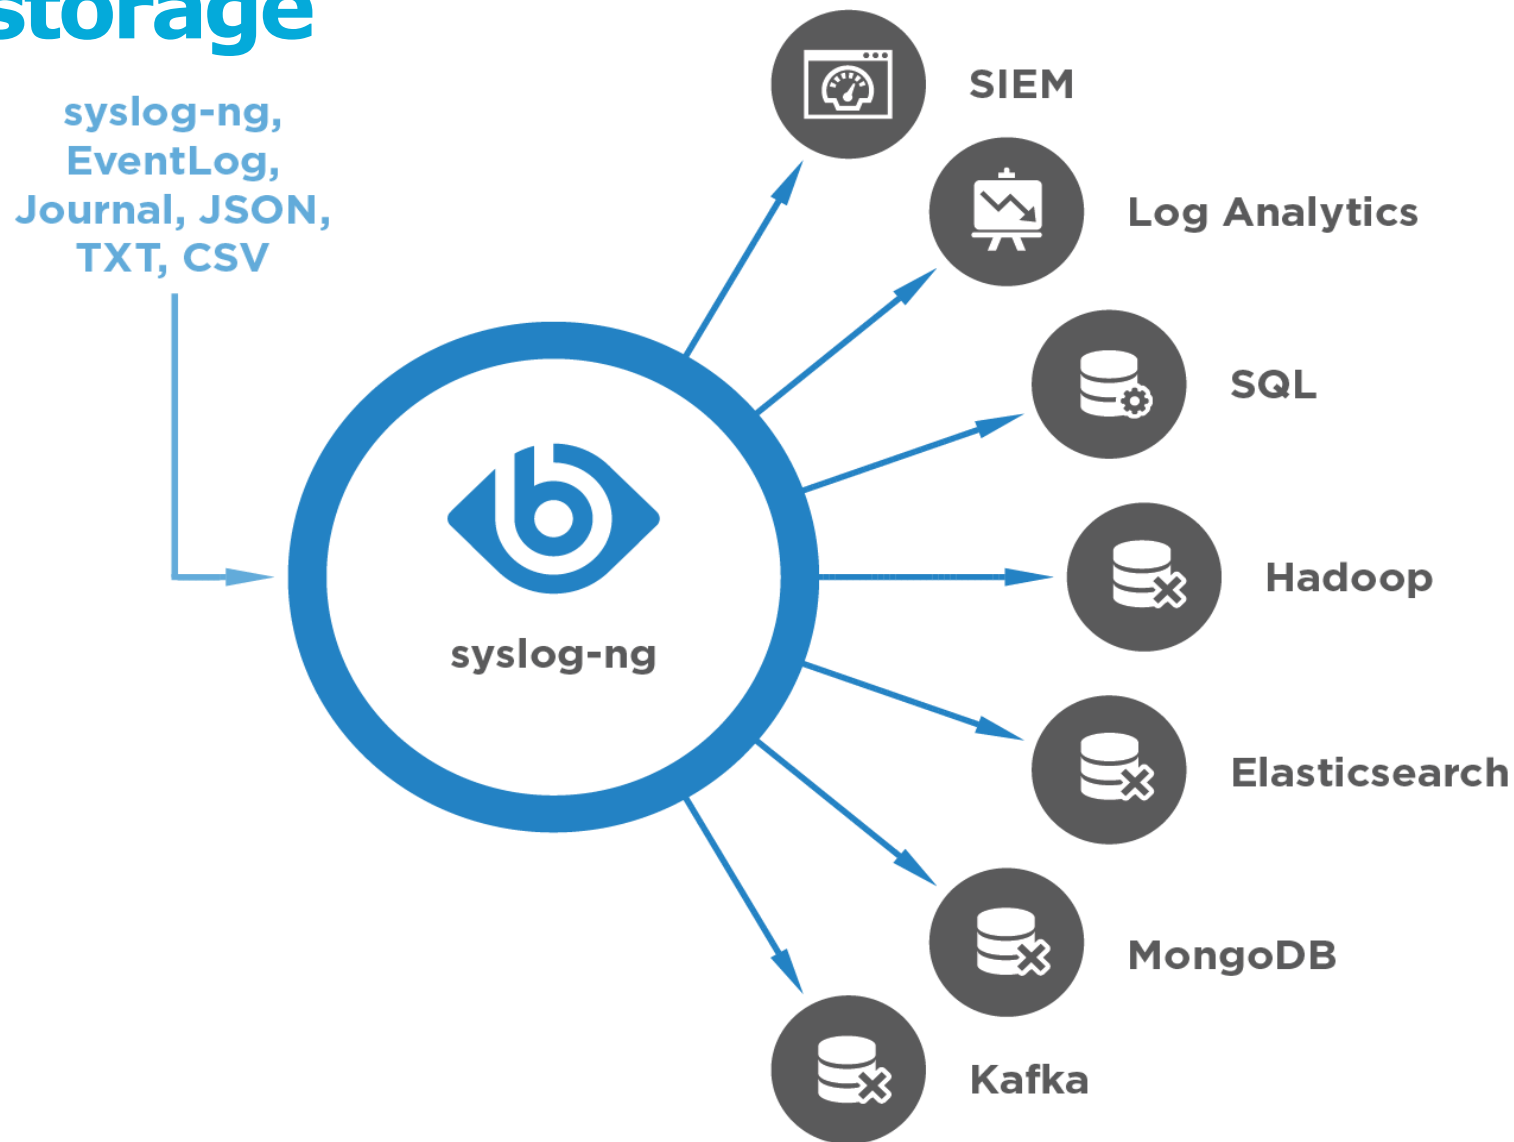
Role: processing

- Classify, normalize, and structure logs with built-in parsers:
 - CSV-parser, PatternDB, JSON parser, key=value parser
- Rewrite messages:
 - For example: anonymization
- Enrich data:
 - GeoIP
 - Additional fields based on message content
- Reformatting messages using templates:
 - Destination might need a specific format (ISO date, JSON, etc.)
- Python parser:
 - all of above, enrich logs from databases and also filtering

Role: data filtering

- Main uses:
 - Discarding surplus logs (not storing debug-level messages)
 - Message routing (login events to SIEM)
- Many possibilities:
 - Based on message content, parameters, or macros
 - Using comparisons, wildcards, regular expressions, and functions
 - Combining all of these with Boolean operators

Role: storage



Freeform log messages

- Most log messages are: **date + hostname + text**

Mar 11 13:37:56 linux-6965 sshd[4547]: Accepted keyboard-interactive/pam for root from 127.0.0.1 port 46048 ssh2

- Text = English sentence with some variable parts
- Easy to read by a human
- Difficult to create alerts or reports

Solution: structured logging

- Events represented as name-value pairs. For example, an ssh login:
app=sshd user=root source_ip=192.168.123.45
- syslog-ng: name-value pairs inside
 - Date, facility, priority, program name, pid, etc.
- Parsers in syslog-ng can turn unstructured and some structured data (CSV, JSON) into name-value pairs

Which is the most used syslog-ng version?

Some hints:

- Project started in 1998
- RHEL EPEL 8 has version 3.23
- Latest stable version is 4.2

 **ONE IDENTITY™**