Making Logs Understandable

The Missing Step in Log Automation

Jiri Setinsky

CESNET

TNC Conference | 2026


100,000+
logs per second
At scale, volume is no longer the main bottleneck.

Processing keeps improving.

Meaning is the bottleneck.

Structure without meaning cannot be automated.

Parsing finds structure

Raw log

User alice
logged in from
10.0.0.8

Parser output

User <VAR1>
logged in from <VAR2>

Current parsers identify patterns and variable positions.

Structure is not meaning

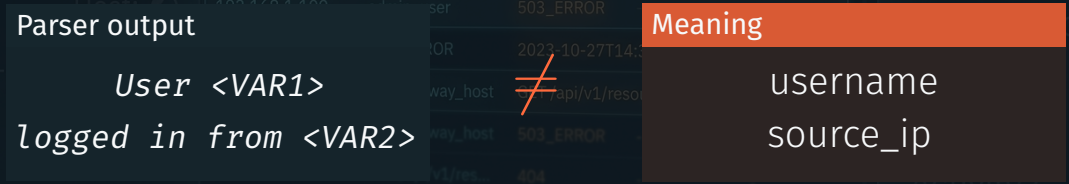

Templates show where values are, not what they mean.

Semantics makes parsing useful

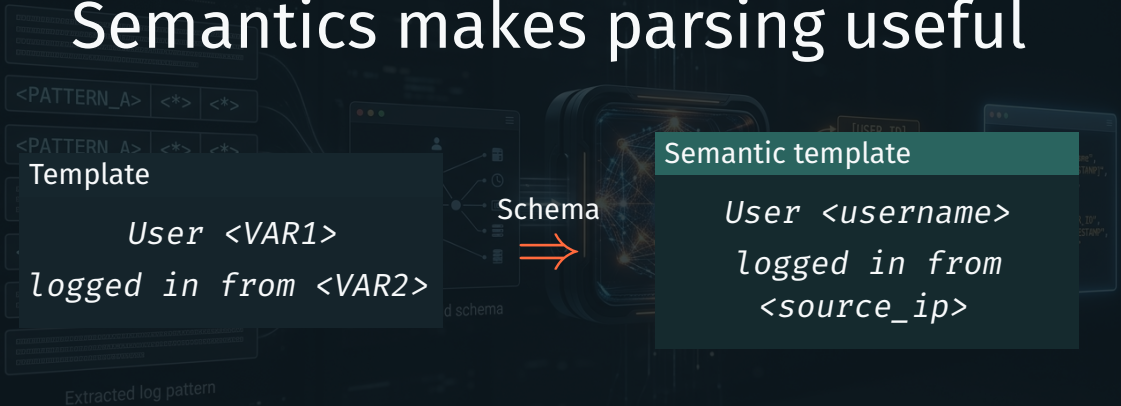

Semantic labels turn extracted variables into fields people can search, trust, and automate on.

Why this matters

New sources

Faster onboarding

Operations

Less expert effort

Pipelines

Earlier action

The goal is not parsed logs. The goal is useful logs.

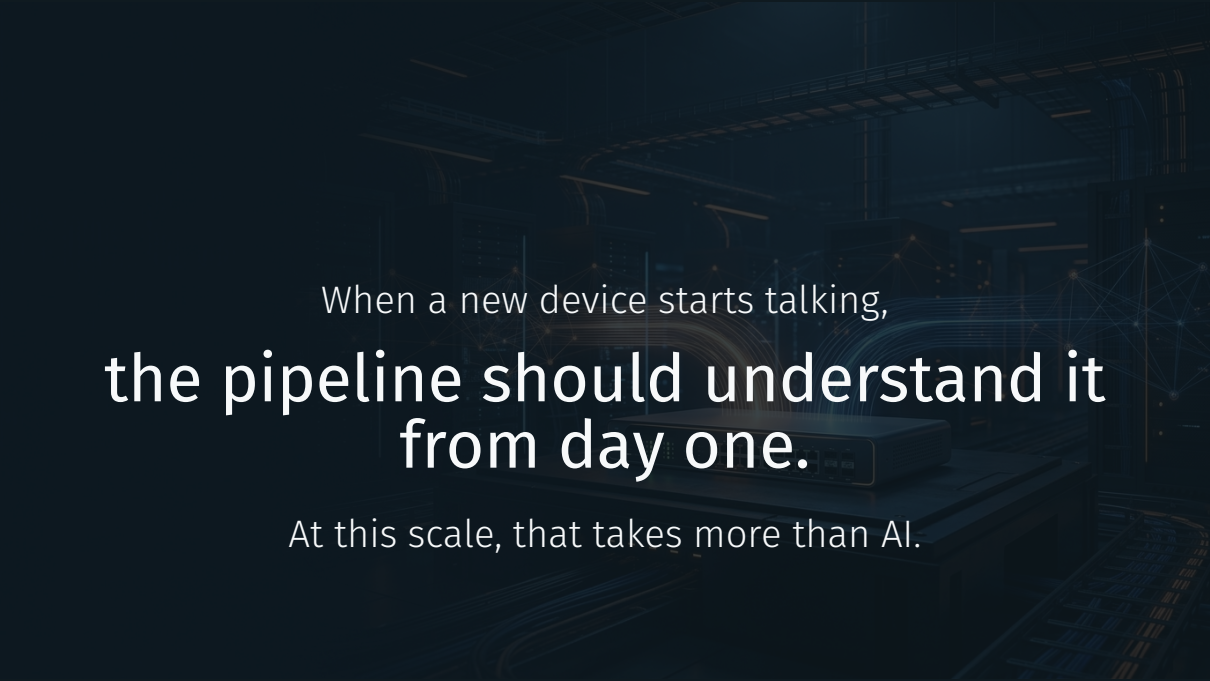
When a new device starts talking,
the pipeline should understand it from day one.
At this scale, that takes more than AI.

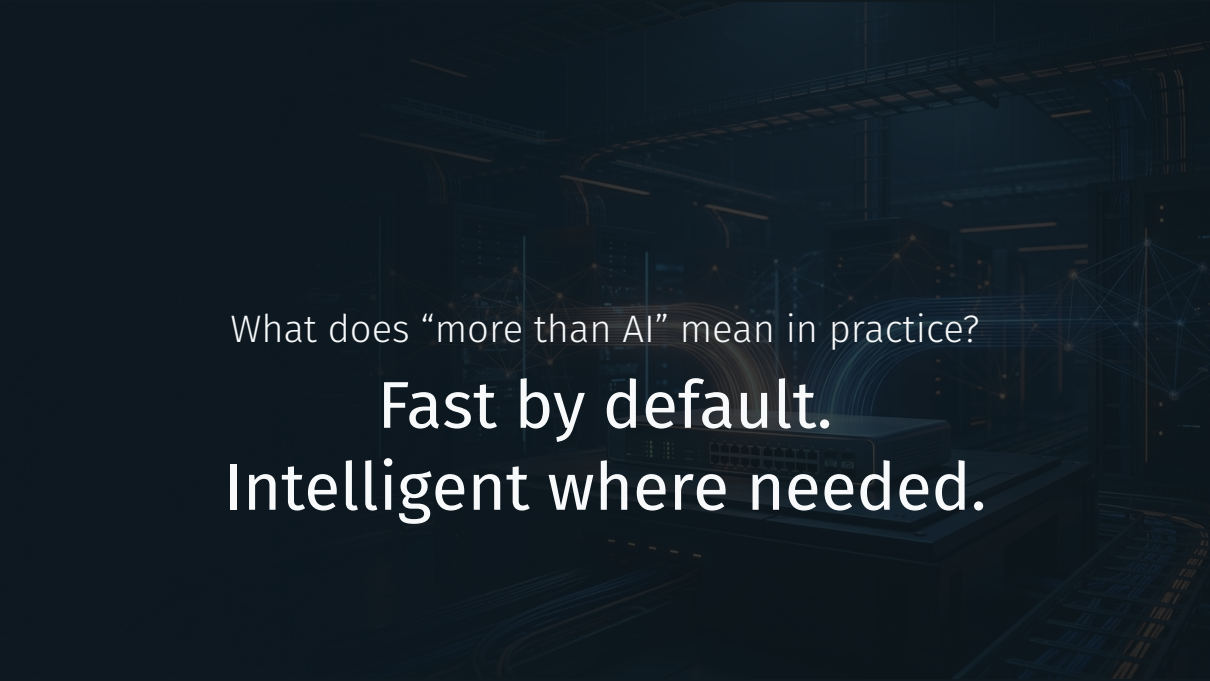
What does “more than AI” mean in practice?
Fast by default.
Intelligent where needed.

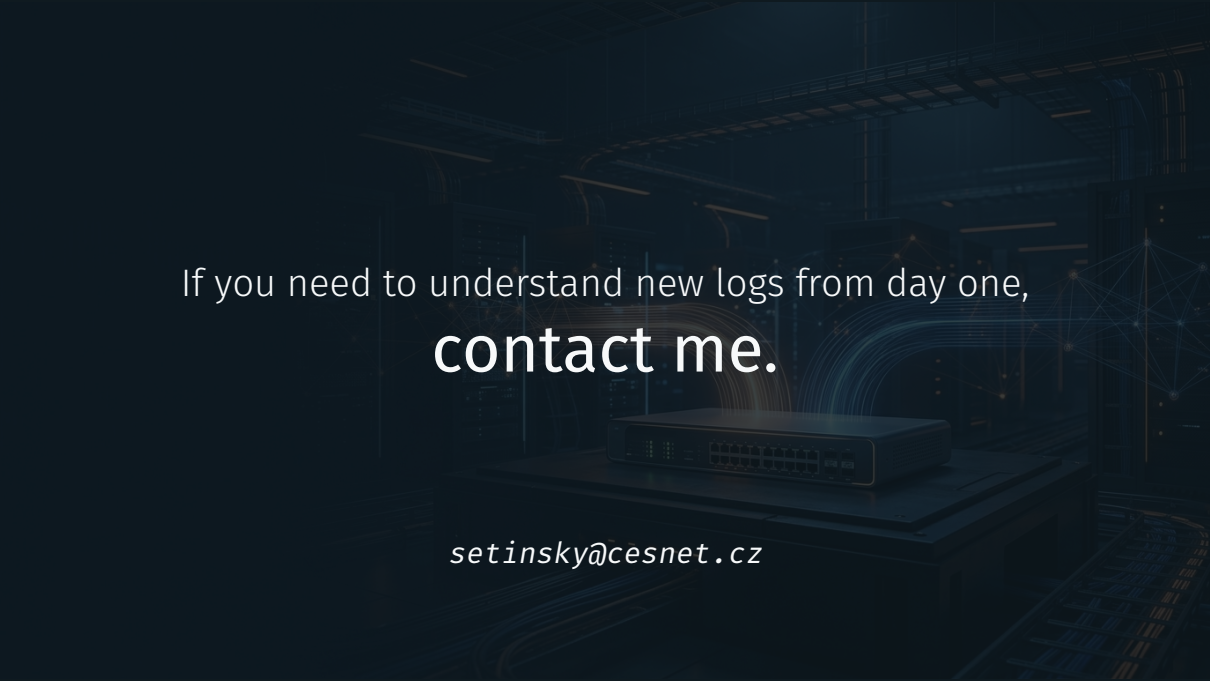

If you need to understand new logs from day one,

contact me.

setinsky@cesnet.cz