DATA FINDINGS

SPL2 & The Modern Splunk Platform

Helsinki Splunk User Group | 3 June 2026

Session Agenda

Helsinki Splunk User Group · SPL2 & Modern Platform

- Welcome & Paradigm Framing — Juha Tamminen (EN)
- Boost Productivity with SPL2: The Next-Gen Language for Splunk — Jay Rakhe (EN)
- Cisco Data Fabric
  - SPL2 & Demo
  - Federation
  - EP & IP
  - Q&A
- Regulatory Context — GDPR, NIS2, DORA etc — Ismo Soutamo (FI)
- Practitioner Checklist — Ismo Soutamo (FI)
- Closing — Ismo + Juha (FI)
- Open Floor (FI/EN)

The Platform Has Changed

"A rolling stone gathers no moss"

- **The platform you bought is not the same platform you are running today**
  - That is not a problem — it is an opportunity
- **You have been running Splunk for years**
  - Bought it for security, operations, compliance, visibility — and it worked
- **But the platform has changed structurally**
  - Architectural assumptions have shifted — affects data, licensing, your team
- **The good news: the direction is the right one**
  - But you need to understand it to take advantage of it

The Old Contract

How Splunk worked — and the pain it created

- **"Give us all your data. We index it. We store it. We make it searchable."**
- **Licensing pressure**
  - More data = more cost. Battles about what was worth indexing.
- **Data gravity**
  - Some data too large, sensitive, or legally constrained to centralise.
- **Architectural rigidity**
  - Single central platform became a bottleneck as data sources multiplied.
- **Skill fragmentation**
  - Ingest pipelines needed completely different expertise from writing searches.

The New Contract

A fundamental reversal

- **”Your data can live anywhere. We bring the analytics to where the data is.”**

And now accelerated by Cisco

- **Why this is irreversible for now**
  - Cisco's world is distributed by definition — networks, endpoints, cloud edges
  - Cisco's bet: AI era requires distributed, observable and secure infrastructure
  - Chuck Robbins at Cisco Live 2025: "The private data center is back for AI"
  - Splunk is the data and analytics layer of that bet

How the Pieces Fit Together

SPL2 — the language the fabric speaks

- **Edge Processor**
  - Analytics at the source — on-prem or cloud
- **Ingest Processor**
  - Analytics at cloud ingest — Splunk Cloud only
- **Federated Search**
  - Analytics across data stores — Splunk, S3, Snowflake, Databricks…
- **Search & Reporting**
  - Analytics at the index — Enterprise + Cloud

- **SPL2 — one language for all of it**
  - Each product uses a profile of SPL2 — same syntax, product-specific feature set
  - The PM will explain exactly what this means in practice

SPL2 as a Programming Model

What is actually new — beyond a query language update

- **Modules**
  - Named, reusable .spl2 files inside a standard Splunk app
  - Version-controlled in Git — same CI/CD pipeline as your TAs
- **Views**
  - Schema definitions enforced at query time
  - Analyst gets the view — never the raw index
- **Custom Types**
  - Schema validation at ingest via EP or IP
  - Catch non-compliant data before it reaches your indexes
- **Lot of other useful features**

- SPL: search language — everything is a pipeline, no reuse above the search bar
- **SPL2: programming model — functions, modules, views, types you can name and version**
- SPL remains fully supported — migration is incremental, no deadline

The Official Story

Straight from Splunk Product Management

- **Jay Rakhe**
  - Product Manager, SPL2 — Splunk
  - Joining remotely

- Cisco Data Fabric
- SPL2 in depth — what it is, why it was built, how it spans the platform
- Federation — S3, Snowflake, Databricks, Azure Data Lakes…
- DMX — EP and IP
- Q&A

- After the PM presentation we return with Finnish-language practitioner content
  - Slides remain in English throughout

Käytännön osuus

Short technical break?

We now switch to Finnish · Regulatory context · Practitioner checklist · Closing

Regulatory Landscape by Sector

Why the timing of SPL2 is excellent

- Universal baseline — all organisations
  - GDPR
  - Kyberturvallisuuslaki 124/2025 (NIS2, in force April 2025)
- **Finance**
  - DORA · MiFID III · MiCA · EMIR Refit · PSD3/PSR · FiDA · AMLA · EBA ESG
- **Healthcare**
  - Cybersecurity Act · EHDS (in force March 2025)
- **Energy · Transport · Telecoms · Public sector**
  - Cybersecurity Act — essential services
- **All sectors**
  - EU AI Act — documentation of AI processes + human oversight

Common core across all frameworks

Who can see what data —
and can you prove it?

Are your processes documented, consistent and auditable —
and can you prove that too?

Modules: Documented Procedures

Repeatable, auditable, version-controlled

- **The requirement**
  - NIS2 / Cybersecurity Act: systematic and repeatable detection
  - Not dependent on who is on shift
  - DORA: consistent detection · MiFID III: audit trails · EU AI Act: documented logic
- **SPL2 module = procedure documentation**
  - Analyst opens the module and runs it — the module is the procedure
  - You can show it to an auditor. You can put it in Git.
  - You can demonstrate your process is repeatable.
- **Module in Git = process audit trail**
  - No longer 'we trust the right person is on shift today'

Views: Who Sees What

Access control and data governance — enforced at the platform level

- **Old model — index-level role permissions**
  - Coarse-grained, hard to audit
  - Raw searches can expose more than intended

- **SPL2 view — completely different model**
  - Write a view definition — e.g. mask username, remove workstation names
  - Assign to a role — analyst gets the view, never the raw index
  - Enforced at platform level · Documented in version control
  - Defensible to any regulatory authority

- **One mechanism — covers the access control requirement across all sectors**
  - GDPR · DORA · EHDS · MiCA · NIS2

Custom Types: Ingest-Time Data Quality

Schema validation at ingest — GDPR Art. 5(1)(d) · DORA data integrity

- Jay showed you custom data type definitions
  - Schema definitions enforced at ingest via EP or IP
- If data arrives malformed — or contains a personal identifier in the wrong field
  - Custom types catch and route it at ingest
  - Non-compliant data flagged or quarantined before it reaches your indexes
- Before: data quality problems discovered during investigation or audit
- **With custom types: caught the moment data arrives**
- **Ingest-time control > post-hoc detection**

Three SPL2 Capabilities

Together they cover the compliance requirement across the entire regulatory landscape

- **Modules — documented, auditable procedures**
  - Investigation logic in Git = process audit trail
  - NIS2 · DORA · MiFID III · EU AI Act
- **Views — access control & data governance**
  - Compliance team gets the view — never raw operational logs
  - GDPR · DORA · EHDS · MiCA · NIS2
- **Custom Types — data quality at ingest**
  - Non-compliant data caught before indexing
  - GDPR Art. 5(1)(d) · DORA

- Bonus: federated search works between Splunk instances too
  - Security + ops stacks · post-acquisition environments — data stays local, query spans all

Version Check — What You Actually Have

Do this on Monday

- **SPL2 in any visible app which has search box requires**
  - Splunk Enterprise 10.2+ (Linux) or Splunk Cloud Platform 10.2.2510+
  - No separate install — language toggle above the search bar
  - Ad hoc SPL2 works from any app with a search bar (ES, ITSI, your own apps)
  - Full module editor: Any visible app with normal Search GUI
- **On Enterprise 9.x? Edge Processor is your SPL2 path now**
- **First steps on 10.2+ / Cloud 10.2.2510+**
  - Open Search & Reporting (or your own app) → switch language toggle to SPL2
  - Try Convert to SPL2 on a search you already know well
  - Open Modules tab → create your first module (Search in SPL2 Module)
- **What to defer**
  - Full SPL migration — no deadline, do it incrementally
  - Federated search — wait until you have a specific use case

Modules Are Files or not?: The Developer Story

Same pipeline as everything else in Splunk

- SPL2 modules are stored into side car PostgreSql
  - Those can get out from there at least enterprise and put then into Git
- **VS Code**
  - Official Splunk extension — full syntax support
- **Git**
  - Version-control the app exactly like any TA
  - Same CI/CD pipeline you already use — no new tooling
- **Deployment**
  - ACS to Splunk Cloud · Deployment server or SHC deployer to Enterprise
- **REST API extraction**
  - Pull modules from a running instance — bridge for GUI-first teams

```
[soutamo@baga]  /opt/docker/splunk/10.4/etc>
(0) $ cat /tmp/spl2/df_splunk_health_check/default/data/spl2/filtered_internal.spl2
@run_as_owner;

import _internal from ../../../../indexes

$internal_errors =
  from _internal
  where sourcetype="splunkd"
    AND log_level IN ("ERROR", "WARN")
  | fields _time, host, component, message, log_level

export $internal_errors
[soutamo@baga]  /opt/docker/splunk/10.4/etc>
(0) $ echo "==="
===
[soutamo@baga]  /opt/docker/splunk/10.4/etc>
(0) $ cat /tmp/spl2/df_splunk_health_check/default/data/spl2/idx_internal_component_exec.spl2
import _internal from ~indexes

$srch_main = search index=_internal source=*splunkd.log
| where component = "ExecProcessor"

export $srch_main
[soutamo@baga]  /opt/docker/splunk/10.4/etc>
(0) $
```

App Promotion: Plan Before You Build

Most teams skip this — and regret it

- **Three questions to answer before building at scale**
  - Source of truth — Git or Splunk? Pick one explicitly.
  - Promotion mechanism — SCP: ACS · Enterprise: deployment server or SHC deployer
  - Environment-specific config — local/ layer per env, or separate branches
- **Minimum viable approach for a GUI-first SCP team**
  - REST API extraction script pulling modules + KOs to Git on a schedule?
  - Or some other method for extracting those from SCP?
  - Run it before people start creating content
  - ACS export/import as the formal promotion gate between environments
- This applies to all Splunk content — not just SPL2
  - SPL2 modules make it visible because they are files or rows in sql? — use that as the hook

Where to Learn — Beyond the Docs

Best sources for practical SPL2 learning

- **help.splunk.com**
  - SPL2 Overview · SPL2 Search Manual · SPL2 Search Reference
- **lantern.splunk.com**
  - Docs tell you what you can do — Lantern tells you what you should do
  - Use-case guidance by industry and data source
- **.conf24 & .conf25 session recordings**
  - Splunk YouTube — real walkthroughs from the people who built it
- **community.splunk.com**
  - Splunk Answers · Community Slack · Recurring office hours live sessions
- **This UG — future sessions**
  - Edge Processor · Federated search architecture · SPL2 migration · App CI/CD

Back to Where We Started

The platform you bought is not the same platform you are running today

- The old model is not going away overnight — your investments are not wasted
- **But the direction is clear and it is not reversible**

- The regulatory angle is not a separate story — it is the same story
  - Views, modules, custom types pay off every time a new requirement lands

- The question is not whether to engage with this
  - It is when and how — version, data challenges, regulatory obligations, team capacity
  - There is no universal right answer

These Conversations Happen Here

Not in a vendor webinar — not in a sales call

- **Potential future sessions — a commitment, not a teaser**
  - Edge Processor — architecture, configuration, use cases
  - Federated search architecture — storage strategy, latency, governance
  - SPL2 migration patterns — incremental paths from SPL1
  - App development & CI/CD — pipeline design, KO governance, env promotion
- **Practitioners. Honesty. Experience — not marketing**

Thank You

Thank you — Jay Rakhe, Splunk Product Management

- Ismo Soutamo
  - ismo.soutamo@data-findings.com · Data Findings Oy
- Juha Tamminen
  - juha.tamminen@mode2.fi · Mode2 Oy

- Ask about app promotion after the session
- **See you at the next session**