



From Packets to Prompts: AI-Assisted Insights and Root Cause Analysis for Trading Operations

Matt Davey, MD, Corvil Product

Deutsche Boerse Open Day, September 24 2026

The content provided in this presentation or in any other information, data or content (whether written or oral), notice or document supplied or otherwise made available to you or your agents and/or advisors in connection with this presentation contains confidential and proprietary information of Pico. Please see the Legal Notice at the end of this presentation for more information.

© Pico 2026. All rights reserved. The Pico symbols and word marks are among the registered and unregistered trademarks of Pico. All other trademarks are the property of their respective owners.

PICO What is Corvil Network Data and Analytics Platform?

SOLUTION: Capture, decode and normalize network data to deliver independent visibility and intelligence across complex trading environments

3 Comprehensive Business Observability

- Business logic to track and join orders across complex trading environments
- Connect data directly to LLM-powered chat
- **Outcome** delivers value for broad set of audiences and use cases


2 Decode, Analyze & Stream Data

- Decode market data, order-entry and enterprise protocols
- Stream data to Kafka, Elastic, KX, JDBC and Pico
- **Outcome** trusted record of every order, trade and transaction

1 Start with Trusted, Golden Data Source

- Network Data capture - 100G/200G throughput
- Lossless capture, Granular timestamps
- No logs, agents or application changes
- **Outcome** independent and accurate data


CONFIDENTIAL



Internal


PICO Evolution of Accessing Large Data Sets with Corvil

Innovating Access to Large Datasets with AI Enablement



PICO Many Consumers, One Gatekeeper

Trading, Risk, Quants and Compliance rely on Operations for access





Goal: Democratize access to Corvil Data using Natural Language Queries

Operations

"Which feeds are gapping today?"


"Compare microburst, 1-sec peak, and packet count for these feeds. Analyze the data and provide a visual representation."



Trade Support

"Client complained about latency of order ABC123."

"Show me how we handled this order. Any issues?"



Quants

"Give me a breakdown of latency by Algo, Order-type and Market."

"Anything surprise you?"



Trade Desk


"Which clients are sending the most orders today?"

How are fill-rates and cancel-ratios trending?"



Accessing Corvil data via AI/LLM

PICO AI Tool Ecosystem Overview



PICO MCP enables access; skills guide action

MCP | Access

What can the agent access and do?

A standard interface to external systems.

Servers expose tools and data resources.

Connects agents to APIs, databases and apps.

Agent Skills | Know-how

How should the agent do the task?

Reusable instructions loaded into context.

Define steps, constraints and output formats.

Guide consistent use of the available tools.

Together: access + know-how

Database example: MCP exposes a query tool; a skill guides safe, efficient SQL for the schema. The agent applies the skill and calls the tool.

Complementary, not inseparable: skills can also guide direct API calls without MCP.

PICO Creating a Corvil MCP Server

1. Use python 'FastMCP' framework

Framework provides server, schema generation, validation. Supports authentication, elicitation and many other features.

```
from fastmcp import FastMCP

mcp = FastMCP("Demo 🌟")

@mcp.tool
def add(a: int, b: int) -> int:
    """Add two numbers"""
    return a + b

if __name__ == "__main__":
    mcp.run()
```

gofastmcp.com

2. Create MCP tools

Declare MCP 'tools' to proxy to Corvil API. **Use plain English to describe the tool usage.** Use self-explanatory parameter names.

```
@mcp.tool
async def get_lens_data(reporting_period: str,
                        stats: Optional[list[str]],
                        is_request: bool,
                        group_by: Optional[list[str]],
                        group_by_sessions: bool,
                        include_overall_values: bool,
                        filter_by_tag_type: str,
                        filter_by_tag: str,
                        context: Context): # -> LensData

    """
    This tool can be used to retrieve statistics via
    a range of "session" objects defined on Corvil
    into tags from one or more tag types (hierarchy)
    be filtered by a tag from a tag type.
    Args:
        reporting_period: A string with one of the
            "Last 1 hour", "Last 24 hours",
```


3. Integrate with Chat UI

Use a System Prompt to provide any further local information.

Helps the LLM to figure out *when* it should use the MCP tool.

- What is being monitored by Corvil?
- What session naming / tagging conventions are in use?
- Locally defined metrics (or expose by adding tools to the MCP server)
- Default prompts

PICO Accessing Corvil data from Chat Interface via MCP



PICO MCP trade-offs: the case for skills

MCP: strengths

- Discover exposed tools, data and how to access them.
- Standard interface across compatible clients and vendor servers.
- A place to implement authentication, access controls, audit and logging.

MCP: trade-offs

- Large, unfiltered results can exhaust the LLM context window.
- Bulk data needs filtering or aggregation—not raw delivery to the model.
- Server deployment and governance still need to be managed.

Skills + agents: go direct to APIs when it fits

Create a playbook for the Agent; the agent writes code to call APIs, filter and aggregate data, then return concise evidence to the LLM.

“MCP gives agents abilities; skills teach agents how to use those abilities well” – [goose-docs.ai](#)

Skills package know-how so agents can repeat it reliably

WHAT A SKILL IS

A reusable playbook that teaches an agent how to perform a task—loading only the guidance it needs.

DESCRIPTION always visible

Full SKILL.md file loaded when invoked

FILES · SCRIPTS · ASSETS

loaded only when needed

Progressive disclosure keeps the model context focused while deeper guidance stays available on demand.

CREATING A SKILL

01 DO THE TASK

Guide the agent until the workflow succeeds.

02 CAPTURE THE METHOD

Ask it to write a reusable SKILL.md and any supporting files.

03 TEST FROM CLEAN CONTEXT

Clear context, run the skill from scratch, fix gaps and repeat.

Alternatives: write it manually or generate it from code / documentation.

Best starting point: preserve a workflow that already works, then prove it is repeatable.



Creating a simple Corvil Skill

Step 1: Interactively using the LLM

 User

I want you to help me find what sessions are defined on a Corvil CNE. Run the “./CorvilApiStatsClient.py” script and read the usage instructions

Use CorvilApiStatsClient.py to retrieve 'summary' stats from ny4cne01. Filter for sessions containing the string Bergman. Summarise the metrics for me.

[...]

that's not how you filter. don't use

'Session/subject:contains:' - just use a glob syntax

Bergman

[...]

 Claude Code (10+ minutes interactive)

```
* Bash(python3 CorvilApiStatsClient.py
summary ny4cne01 -f "Bergman" -r "1-
hour")
* #client version:
3.2.0.202206301037-GA+273102
#summary request generated at
2026-09-08 11:21:29
#CNE: ny4cne01:5101
#Reporting period: 1-hour
[...]
```

? that's not how you filter. don't use 'Session/subject:contains:' - just use a glob syntax '*Bergman*'

Commands used by Claude:

```
CorvilApiStatsClient.py --help
```

```
CorvilApiStatsClient.py ny4cne01 summary
```

```
CorvilApiStatsClient.py ny4cne01 summary -f "AlgoSpeed"
```

```
CorvilApiStatsClient.py ny4cne01 summary -f "*AlgoSpeed"
```

[...]



Creating a simple Corvil Skill

Step 2: Preserve the Findings

User

Now that you know how to answer this question, create a reusable skill to accomplish this task, in the root directory of this project.

Claude Code

I have created a 'search-sessions' skill to document the process for finding sessions in the **SKILL.md** file at the root of the project. The output is a compact Markdown table that lists only the rt-class entries (the API name and the human-readable display name) together with the most useful summary columns..

Resulting SKILL.md file:

```
**Skill name:** `search-sessions`

**Description:** Search for Corvil sessions whose name (or description)
contains a given substring.

**Usage:**
...
/?search-sessions <host> <substring>
...

The skill runs the `CorvilApiStatsClient.py summary <host>` command
with a glob filter matching the provided substring (e.g. `*
<substring>*`). It then prints a compact table of the matching sessions
(name, display name, total bytes, average utilisation, etc.).

**Implementation notes:**
- The skill invokes:
  ```bash
 python3 CorvilApiStatsClient.py summary <host> -f "*"<substring>*"
  ```
- It parses the CSV output and formats a tidy Markdown table with one
row per rtclass
- It ignores the 'channel' endpoints and just reports the 'rtclass'
entries
- It outputs the 'rtclass' API name as well as the display name
- Errors (e.g., no matches) are reported to the user.
...

```



Creating a simple Corvil Skill

Step 3: Retry from Scratch

 User

What are the summary metrics for AlgoSpeed, from the probe2216 CNE for the last hour?

Commands used:

```
CorvilApiStatsClient.py probe2216 summary -f  
"*AlgoSpeed*"
```

 Claude Code (after 3s)

```
what are the summary metrics for AlgoSpeed from probe2216 CNE for the last hour?
```

```
Summary metrics for AlgoSpeed (host probe2216, last 1 hour)  
  
Direction: Forward (-)  
rt-class (API): rt-class/CUST AlgoSpeed//NY4-CNE/CustomGateway//class-default  
Config Cap (kbps): -  
Eff Cap (kbps): 100 000  
Total Bytes: 2 058 659  
Avg Util %: 0  
NSI: 0.1  
Measure Msgs:   
Mon Mech: pnm  
1-sec Peak (msg/s): 23  
Max Microburst (msg/s): 800  
Packet Microburst Avail.: false  
  
Direction: Reverse (-)  
rt-class (API): rt-class/CustomGateway/NY4-CNE//CUST AlgoSpeed//class-default  
Config Cap (kbps): -  
Eff Cap (kbps): 100 000  
Total Bytes: 4 443 181  
Avg Util %: 0  
NSI: -  
Measure Msgs:   
Mon Mech: none  
1-sec Peak (msg/s): 444  
Max Microburst (msg/s): 8 250  
Packet Microburst Avail.: false
```


After creating a Skill, the agent is faster, more accurate, using fewer tokens

Accessing Corvil via Chat Interfaces

PICO Corvil MCP – natural language queries direct to Corvil data


1. CNE Direct Data

CNE Lens API metrics with grouping and aggregation



Natural Language UI

- Which Clients are sending the most orders today
- What happened to Order X – where was the latency?
- Which market-data feeds are gapping today?



PICO Example Interactions

“Bergman has pretty high latency. Please take a look and see if it looks like an application layer issue or a network layer issue. Give me your analysis.”

metrics

reported (likely not configured), but loss % is low.

delay is likely happening inside the application stack (e.g., processing queues, downstream system latency, thread contention, GC pauses, etc.).

Conclusion – Application-layer issue

- The network is delivering packets promptly (low RTT, no retransmissions, ample head-room in bandwidth).
- The huge gap between network RTT (≤ 134 ms) and observed end-to-end latency (≥ 1.66 s p99) points to processing time inside the order-entry gateway or downstream systems (matching engine, risk checks, database look-ups, etc.).
- The modest loss % and absence of zero-window events further support that the bottleneck is not at the transport layer.


Recommendations for a deeper dive

1. **Instrument the gateway code** – add timestamps at key stages (receive, decode, risk-check, order-routing, send) and correlate with the Corvil latency

“What 5 clients have the highest cancel ratio today? Compare to order and trade volumes, and visualize.”

Top 5 clients by cancel-ratio (today – UTC)

| CLIENT | ORDERS (EG-IN) | CANCELS (EG-IN) | CANCEL-RATIO % | FILLED EVENTS | FILLED QUANTITY ☐ |
|----------------|----------------|-----------------|----------------|---------------|--|
| HugeTrade | 32 067 | 10 800 | 33.68 % | 1926 339 | 87 859 524 |
| ABCXYZ | 31948 | 10751 | 33.65 % | 1945 311 | 87 139 541 |
| FergalOTrading | 32 075 | 10 785 | 33.62 % | 1933 148 | 87 303 646 |
| GoGoStock | 31808 | 10 692 | 33.61 % | 1916 735 | 87 373 347 |
| BigMoneyInc | 32 120 | 10 785 | 33.58 % | 1942 464 | 87 771 698 |



PICO Example Interactions

“What is the CIOrdID of the child order with highest parent-to-first-child latency in the past hour? Follow that CIOrdID end to end. On which component was the latency incurred? ”

Answer

The biggest parent-to-first-child latency for the child order CIOrdID = `204491.ScenarioState@75e8a146.3379` was measured at the `SOR-OUT` capture point.

| TIER | CLIENT NODE | SERVER NODE | ISSUER | CUSTODIAN | LATENCY (NS) | EVENT | GROUP NAME |  |  |
|---------|-------------|-------------|--------|-----------|------------------------------------|-------|----------------------|---|---|
| SOR-OUT | SOR20 | SOR-BUS-MAV | SOR20 | NAPA | 89 597 905 297 ns (≈ 89.6 seconds) | order | SOR-OUT_MA-IN_MA-OUT | | |

The later `MA-IN` and `MA-OUT` events for the same order show only a few micro-seconds of parent-to-first-child latency (1 211 ns and 9 539 ns), confirming that the substantial delay happened **as the order left the SOR (Smart Order Router) and entered the market-access path**. In other words, the latency was incurred on the SOR-OUT side—between the SOR20 instance (client node) and the downstream market-access gateway (the first child in the chain).

investigate the latency across SOR20. How does it compare to all the other SOR instances? Is there an issue across all latency percentiles, or just outliers?



DSA Algo/SOR Plant Monitoring

Thought for 17 seconds ▾

View Result from `tool_run_select_query_post` ▾



Using Agents and Corvil Skills for AI-Assisted operations workflows

PICO Corvil skills for automated and interactive AIOps

Build reusable skills once—combine them for different operational questions.

REUSABLE CORVIL SKILLS

01 SEARCH

Find sessions
List metrics
Retrieve tags

02 METRICS

Summary, grouping
Timeseries
Forensics

03 MESSAGES

Decoded messages
Retrieve packets
Filtering

04 SYSTEM

Configuration
Dashboarding
System Health

COMPLEMENT WITH BUILT-IN SKILLS

Report creation • Automated scheduling • tshark packet analysis • Visualizations • ...

ONE TOOLKIT, TWO OPERATING MODES

AUTOMATED REPORTS

ROOT CAUSE ANALYSIS FOR ALERTS

Event trigger → collect evidence → create PDF report

INTERACTIVE INVESTIGATIONS

AI-ACCELERATED WORKFLOWS

Agentic CLI → analyse interactively → inspect packets → create reports → optionally create skills

PICO Customer Use Case: Automated WAN drop root cause

AUTOMATED WAN DROP ROOT CAUSE

1 Trigger

WAN switch-drop alert with timestamp and link.

2 Corvil evidence via API skills

Find CNE session; 5-second microburst; top conversations, talkers and listeners.

3 Agent reasoning

Test threshold, rank contributors and separate single-flow from aggregate congestion.

4 Automated report

Report likely culprit flows and hosts, with evidence window and confidence.

Outcome

Faster triage, consistent evidence collection and an auditable hand-off.

PICO Example automation skill: investigate a WAN drop

A skill makes expert practice invocable, efficient and consistent across investigations.

SKILL.md

```
---
name: investigate-wan-drop
description: Investigate WAN drops.
---
# Inputs
drop alert: link | timestamp | threshold
# Procedure
1. Resolve monitoring CNE session.
2. Call getAnalytics() for microburst  $\pm 5s$ 
3. Compare peak with threshold.
4. If exceeded, request top-N:
   conversations, talkers, listeners.
5. Rank flows and hosts.
# Output
PDF report + evidence.json + summary.md
# Guardrails
Read-only; timezone explicit;
attribute only with evidence.
```

CONFIDENTIAL

WHAT THIS STANDARDIZES

Trigger contract

Requires a named link, alert timestamp and threshold before the agent starts.

Tool sequence

Resolves the right CNE session before issuing time-bounded Corvil queries.

Evidence discipline

Expands only above threshold and separates observed contributors from attribution.

Reusable outcome


A concise conclusion and structured evidence another analyst or agent can reproduce.

PICO WAN Drop: example report output

Corvil provides the microburst drill down showing the moment the WAN drop occurred

Fast summaries of talkers, listeners, conversations.

Use of Microburst is key to identifying culprits *that contributed to drop*



PICO Exchange case study: Fragmented Orders

INTERACTIVE ORDER FRAGMENTATION INVESTIGATION

1 Trigger

Corvil dashboard highlighted a small number of market participants with > 30% out-of-order TCP
Corvil identified every Order completed by an out-of-order segment using runtime decoder analytics

2 Using Agents to gather evidence

Retrieve list of Orders with out-of-order TCP – Filter by presence of “**VirtIsolatedSmallFragment**” field
Retrieve all packets contributing to the Order (bounded by timestamp of previously decoded message)

3 Agent reasoning

Reconstruct TCP sequence and identify timestamps of fragments comprising order
Using protocol specification, Agent can map fragments to message fields – **What was sent late?**

4 Generate Report

Compile report with analysis, showing participant / session details, specific transmission pattern, which fields were transmitted per segment, packet evidence and a neutral explanation.

Outcome

Faster triage, consistent evidence collection and an auditable hand-off.

PICO Exchange case study: Fragmented Orders

A skill makes expert practice invocable, efficient and consistent across investigations.

SKILL.md

```
---
name: analyse-fragmented-order
description: Explain flagged messages.
---
# Inputs
fragmented-message-alert | flow |
timestamp
# Procedure
1. Retrieve message triggering alert.
2. Find timestamp of preceding message.
3. Retrieve bounded packet window.
4. Reconstruct TCP reassembly.
5. Map late bytes to protocol fields.
6. Cite evidence and gaps.
# Output
rca.json + customer_report.md
# Guardrails
Redact payloads; preserve artefacts;
do not infer intent.
```

CONFIDENTIAL

WHAT THIS STANDARDIZES

Trigger contract

Starts from identified orders with timestamp
— not a broad packet search.

Evidence window

Uses message-completion boundaries to
retrieve the smallest useful PCAP.

Reporting rule

Combines TCP reconstruction with field
mapping; states evidence and uncertainty.

Reusable outcome

A concise conclusion and structured
evidence another analyst or agent can
reproduce.

PICO Exchange case study: Fragmented Orders

Completed report details the Operations investigation.

Backed by precise packet and message data captured by Corvil and retrieved by the Agent.

Agent uses neutral tone to document the findings

“The combination of message CSV, decoder virt fields and pcap from Corvil is just so powerful!”

**- Operations Engineer
Tier 1 EMEA Exchange**

CORVIL AGENTIC INVESTIGATION REPORT


SYNTHETIC

A 145 ms TCP GAP HOLDS BACK THE COMPLETE ORDER


Finding: Fragment 1 starts the order. Fragment 2 follows just $0.3\ \mu\text{s}$ later, but its sequence begins beyond a missing byte range. The absent bytes—Symbol, Side, Price and Quantity—arrive in Fragment 3 after 145 ms and close the gap.

1 OBSERVED ARRIVAL ORDER

NOT TO TIME SCALE



2 TCP BYTE ORDER AFTER FRAGMENTS 1 + 2



3 GAP FILLED

+145 ms · Symbol · Side · Price · Quantity

DELIVERY IMPACT

TCP holds Fragment 2 until Fragment 3 closes the gap. Only then can the decoder deliver the complete Order to the matching-engine application. The 145 ms delay occurs before application delivery.

Illustrative byte ranges and identifiers. No customer artefact used.

PICO Corvil API / LLM Roadmap

REST API: Metrics, Data, Configuration

- Agent friendly, self-describing OpenAPI spec.
- Supports MCP and direct agentic usage
- AI assistance to manage config, build dashboards

MCP Servers



- Modules to add MCP functionality to existing CNEs

Connectors and Observability

- Throughput and functionality for streaming flows, metrics, events, packets into Elastic, Kafka, SQL...

Corvil curated Skills library

- Skills library to create structured agentic capabilities
- Repeatable, efficient, and informed by experts



**Open Weight models
are viable when
control matters**

Great option when data sensitivity / infosec is critical
Fast for prototyping – just stand up a local GPU
Good domain knowledge (trading, netops, root-causes)

**Use Frontier Models
for more complex APIs
and workflows**

Smaller context windows become a constraint for data tasks
Frontier models are more capable at API use and reasoning
Can complete more complex workflows and create rich reports

**MCP plus local skills
turns access into
repeatable practice**

MCP ideal for lower-volume data and plugging directly to Chat
Augment vendor-provided skills with locally creations to create
powerful, repeatable, valuable agentic workflows

SEE THE DEMO!

Scan the QR code for a follow up.



Trade Desk

“Which clients are sending the most orders today?”



Trade Support

“Client complained about latency of order ABC123. Show me how we handled this order. Any issues?”



NetOps

“Which feeds are gapping today? Please compare microburst, 1-sec peak, and packet count for these feeds. Analyze the data and provide a visual representation.”



Quants

“Give me a breakdown of latency by Algo, Order-type and Market. Anything surprise you?”



FINANCIAL TECHNOLOGY SERVICES

THANK YOU



LEGAL NOTICE

The content provided in this presentation or in any other information, data or content (whether written or oral), notice or document supplied or otherwise made available to you or your agents and/or advisors in connection with this presentation (the "Information") contains confidential and proprietary information of Pico. The Information may not be reproduced or redistributed or publicized or disclosed to any person outside the recipient's organization, in whole or in part, for any purpose, without the prior written consent of Pico. The Information is provided solely for informational purposes, and is not intended to provide legal, tax, investment or accounting advice, nor should it be construed as a recommendation, or an offer to buy or sell, or the solicitation of an offer to buy or sell any security, financial product or to participate in any trading strategy. No determination of your suitability for, or the appropriateness of, any service provided by Pico has been made. Any performance data provided in the Information represents past performance and is not necessarily indicative of future performance. The Information may contain forward-looking statements, which by their nature involve risks and uncertainties. Future results may vary from what is presented in the Information and in no event are guaranteed. Pico undertakes no obligation to update the Information. The Information is provided as is, with no representation or warranty made or implied, and Pico assumes no responsibility for the accuracy, completeness or reliability of the Information, and Pico disclaims all liability for any use you or any other party may make of the Information. The Information, in any form, is not intended for distribution to, or use by, any person or entity in any jurisdiction or country where such distribution or use would be contrary to local law or regulation.

©2026 Pico Quantitative Trading LLC. All rights reserved. The Pico symbols and word marks are among the registered and unregistered trademarks of Pico. All other trademarks are the property of their respective owners.