

precisely


IBM i & AI: Friend or Foe?


9th September — 14:00 - 15:00

Stephan LEISSE
stephan.leisse@precisely.com



For over thirty years, this was a defence



The background features a collection of 3D-rendered geometric shapes in various shades of purple. These include a large horizontal rectangle in the top left, a sphere to its right, a tall vertical rectangle on the right side, a small vertical rectangle in the center, a sphere below it, a sphere in the bottom left, and a square below that. The shapes are arranged in a way that creates a sense of depth and balance.

The Foe: what AI actually changed

AI did not invent a new vulnerability.

It industrialised the ones we already tolerate: authority creep, clear-text protocols, open IFS directories, and a platform the SOC cannot see.

What changed in eighteen months

82.6%

Share of phishing email showing meaningful signs of AI involvement by early 2026. In November 2025 it was 4%.

+44%

Increase in attacks that begin by exploiting a public-facing application, driven largely by AI-enabled vulnerability discovery. (IBM X-Force Threat Index, February 2026)

400%

Growth in successful phishing attributed to AI tooling across 2025. Reported click rates around 54%, credential capture around 34%.

80–90%

Share of tactical operations executed by the model itself in the first publicly reported AI-orchestrated espionage campaign, across roughly 30 targets. (November 2025)

42%

Of global phishing breaches now driven by autonomous agents rather than a person at a keyboard. (SentinelOne, 2026)

90+

Organisations where adversaries turned the defender's own generative AI tooling against it through prompt injection. (CrowdStrike, 2026)

How AI sharpens the attack

Hyper-real phishing

Generative models draft flawless spear-phishing aimed at named IBM i administrators. Deepfaked voice backs up the call requesting a password reset.

Machine-speed exfiltration

Scripted ODBC and JDBC sessions crawl DB2 tables and the IFS in minutes — well inside any manual review cycle.

Automated exploit generation

PTFs and OS updates reverse-engineered at scale to find zero-days, then payloads tailored to a specific IBM i release and configuration.

IFS ransomware

One infected PC with a mapped drive is enough. Automated tooling prioritises the highest-value IFS directories first.

Low-and-slow credential stuffing

Tooling learns your profile naming and password conventions, then paces attempts to stay under QMAXSIGN and threshold-based lockout.

Insider aggregation

Valid credentials plus an AI query assistant assemble a complete sensitive dataset out of individually unremarkable reads.

Anatomy of an AI-assisted breach



Recon

Minutes, not
weeks



Entry

Valid credentials,
background port




Escalate

Hunt for high
privileges
(*ALLOBJ)



Exfiltrate

Paged reads over
communication
Interfaces
(ODBC)



IBM I is not the most secure
system. It is the most
securable.

— which puts the outcome entirely in your hands

The background is a solid dark purple. Scattered across it are several 3D-rendered geometric shapes in a lighter shade of purple. These include a horizontal rectangle in the top left, a sphere to its right, a vertical rectangle in the top right, a sphere in the center, a vertical rectangle in the bottom right, a sphere in the bottom left, and a horizontal rectangle in the bottom center. The shapes have soft shadows, giving them a three-dimensional appearance.

**The Friend :
putting AI on our side of the wall**

AI on defence: what it changes

Predictive detection

- Models baseline normal behavior across millions of QAUDJRN events.
- Surfaces the subtle case: legitimate authority, abnormal volume, abnormal hour.

Automated response

- Detection triggers containment without waiting for a human to read a log.
- End a rogue job, deny a source address, revoke an elevation, isolate a session.

Continuous assessment

- Constant discovery of *PUBLIC authority leaks and misconfigured profiles.
- Flags drifted system values and inherited authority far faster than an annual audit.

Plain-language evidence


- Turns journal entries into sentences an operations manager or auditor can act on.
- Drafts the compliance narrative from the actual data, and names the gaps.

Assure Security

Complete IBM i Security and Compliance

Assure Security delivers market-leading IBM i security capabilities that help organizations successfully comply with cybersecurity regulations and reduce security vulnerabilities

Six layers. Assure Security covers four.



Concept to module: the Assure Security map

SECURITY CONCEPT	ASSURE MODULE	WHAT IT STOPS
Identity verification	Assure Multi-Factor Authentication	A phished or stuffed password on its own
Port and protocol defence	Assure System Access Manager	Scripted ODBC, JDBC, FTP and IFS access
Privilege containment	Assure Elevated Authority Manager	Permanent *ALLOBJ and lateral movement
Unified visibility	Assure Monitoring and Reporting	The blind spot between IBM i and the SIEM

The background features an abstract composition of various geometric shapes in shades of purple and blue. These include a large horizontal rectangle at the top left, a sphere to its right, a vertical rectangle on the right side, a sphere below it, a small vertical rectangle in the center, a sphere to its right, a sphere at the bottom left, and a square below it. The shapes are arranged in a way that creates a sense of depth and balance.

Assure Multi-Factor Authentication

Assure MFA: the capability

Factors and services

Two or more factors: something the user knows, has, or is.

Codes delivered by mobile device, email or hardware token.

Free Precisely Authenticator, plus TOTP (RFC 6238) and RADIUS-based servers.

RSA SecurID, Duo, Okta, Entrust and Microsoft Entra ID.

RSA certified. One-step authentication available, as PCI DSS requires.

Where and when it applies

Rules engine decides when a second factor is required — not a global switch.

By user, group profile, special authority, IP address, device or time of day.

5250 sign-on, on demand, or via FTP and ODBC/ JDBC in combination with SAM.

Self-service profile re-enablement and self-service password change.

Supports the four-eyes principle for supervised changes.

What it delivers

A layer above a memorised, written down or phished password.

A compromised credential stops being a compromised system.

Addresses PCI DSS, NYDFS 23 NYCRR 500, SWIFT Alliance Access and GLBA/ FFIEC.

Failure messages neutral and configurable; no trace left in the job log.

Where AI helps: Multi-Factor Authentication

Risk-scored step-up

The rules engine already decides when a second factor applies. Feed it a risk score built from AMR history rather than a static list, and the challenge lands on the sessions that deserve it.

Push-fatigue detection

Repeated authenticator prompts that are eventually accepted are a known attack pattern. A model spots the burst; policy stops approving.

Impossible travel

A 5250 sign-on from Montpellier and an ODBC bind on the same profile from another continent nine minutes later. Correlated in the SIEM, returned as a hard challenge or a deny.

Deepfake-resistant enrolment

The weak point is now the help desk, not the algorithm. Synthetic voice makes 'reset my password' calls unreliable — route re-enablement to MFA self-service and the four-eyes principle instead.

Device and hour baselines

Two years of connections from the same two subnets between 08:00 and 18:00 is a baseline. The 03:00 login from an unknown device is the exception worth interrupting.

Exception review that keeps up

Every bypass and every exempted profile summarised and ranked each month, so the exception list stops growing quietly.

The background is a solid dark purple. Scattered across it are several 3D-rendered geometric shapes: a horizontal rectangle in the top left, a sphere to its right, a vertical rectangle in the top right, a sphere in the center, a vertical rectangle in the bottom right, a sphere in the bottom left, and a horizontal rectangle in the bottom center. All shapes have soft shadows, giving them a three-dimensional appearance.

Assure System Access Manager

System Access Manager: the capability

Access it controls

Network access: FTP, ODBC, JDBC, OLE DB, DDM, DRDA and NetServer.

Communication ports: IP addresses and sockets, covering SSH, SFTP and SMTP.

Database access over open-source stacks: JSON, Node.js, Python and Ruby.

Command access: including remote command execution.

One rule set across every exit point instead of scattered custom programs.

How it is managed

Graphical interface with a standard configuration for out-of-the-box deployment.

Rules keyed on date and time, user profile attributes, object and IP address.

Simulation mode tests rules with no impact on users.

Alerts and reports; logs access data for SIEM integration.

No application modification required.

What it delivers

Supports SOX, GDPR, PCI DSS and HIPAA requirements.

Cuts the time and cost of demonstrating regulatory compliance.

Detects incidents quickly enough to remediate them.

Low impact on system performance.

Removes the hand-written exit programs nobody wants to own in five years.

Where AI helps: System Access Manager

Rules written from evidence

Point a model at ninety days of SAM logs and it will tell you which profiles actually use ODBC, from which addresses, at which hours. Write the rule from that, not from a guess.

Anomaly-driven containment

A profile that has never issued a DDM request suddenly issues four hundred. The model raises it; the containing rule is drafted in seconds instead of days.

Predicting what a rule would break

Simulation mode already shows what a rule would have blocked. AI turns that output into a plain-language impact statement your change board can read.

Orphan and shadow rule cleanup

Rule sets accumulate. AI flags the rules that have not matched a connection in a year, and the ones silently shadowed by a broader rule above them.

Natural-language rule drafting

'Deny FTP outside the two batch profiles between 18:00 and 06:00' becomes a draft rule, staged in simulation, never applied automatically.

Enrichment before the decision

Source addresses enriched against threat intelligence before the allow decision, so a known-bad range never reaches the exit point at all.

An abstract background featuring various 3D geometric shapes in shades of purple. There are several spheres of different sizes, some rectangular blocks, and a thin vertical bar. The shapes are arranged in a way that creates a sense of depth and balance. The text 'Assure Elevated Authority Manager' is centered in a bold, white, sans-serif font.

Assure Elevated Authority Manager

Elevated Authority Manager: the capability

Controlling elevation

Requests granted manually by an administrator, or automatically by rule.

Rules on source and target profiles, group profiles, supplemental groups and user lists.

Context conditions such as time of day, job name and IP address.

*SWAP or *ADOPT methods supported.

Handles processes connecting over ODBC, JDBC, DRDA and FTP.

Monitoring elevated profiles

Elevated users and remaining duration visible from the GUI or 5250 displays.

Audit trail built from job logs, screen captures, exit points and journals.

Log-only option records activity without changing authority.

Alerts on events such as exceeding the authorised time window.

Reports in a range of formats; ticketing system integration.

What it delivers

An audit trail of privileged actions that satisfies an auditor.

Enforced segregation of duties.

Far fewer powerful profiles standing on the system.

The user never learns the privileged password — EAM swaps rather than shares.

Less exposure from human error as well as from deliberate misuse.

Where AI helps: Elevated Authority Manager

Right-sizing from log-only mode

A month in log-only mode gives you evidence rather than opinion. A model reads it and proposes who needs elevation, for what, and for how long.

Session summary in plain language

Instead of reading a job log, the approver gets three sentences: what was elevated, what was run, and whether anything sat outside the norm.

Risk-scored approval

The same request at 10:00 from a managed laptop and at 03:00 from an unmanaged one are not the same request. Score it before a human is asked to approve it.

Early revoke on divergence

The window has not expired but the behaviour has. Ending the elevation early is a policy decision AI can trigger and a human can review afterwards.

A command vocabulary per administrator

Every admin has a normal command set. A CHGSYSVAL or a DLTLIB from someone who has never run one is the signal worth acting on.

Evidence the auditor accepts

Every elevation correlated to the ticket that justified it, with the unmatched ones surfaced rather than buried in ten thousand rows.

The background features a collection of 3D geometric shapes in various shades of purple. These include a large horizontal rectangle in the top left, a sphere to its right, another sphere in the top right, a small vertical rectangle in the center, a sphere below it, a sphere in the bottom left, a square below that, and a tall vertical rectangle on the right side. The shapes are arranged in a way that creates a sense of depth and balance.

Assure Monitoring and Reporting

Monitoring and Reporting: the capability

What it monitors

Simplifies the analysis of complex journals — QAUDJRN and database journals.

System and database changes, message queues, QHST and job logs.

Query engine with extensive filtering to find deviations from best practice.

Static security information too: users, groups, system values, configuration.

Customisable models supplied for common ERP applications and for GDPR.

Alerts and reports

Event alerts by email, popup or syslog.

Reports continuous, scheduled or on demand — down to the second.

PDF, XLS, CSV and PF formats, distributed by SMTP, FTP or IFS.

SIEM add-on feeds QRadar (certified), Splunk, Sentinel, ArcSight, LogRhythm and ELK.

Formats include *LEEF, *CEF, *RFC3164, *RFC5424 and user-defined.

What it delivers

Fast identification of incidents and compliance deviations.

A continuous check that the controls you implemented are still in place.

Meets GDPR, SOX, PCI DSS and HIPAA logging requirements.

A journal-based audit trail with real segregation of duties.

No application modification required.

Where AI helps: Monitoring and Reporting

Correlation, not more alerts

AF, PW and CP entries that mean nothing individually become one ranked incident with a beginning, a middle and a named profile.

Triage that ranks

The value is not detecting more. It is putting the six events that matter above the six thousand that do not.

Plain-language incident narrative

'At 02:14 profile BATCH01 read 412,000 rows from CUSTMAST over ODBC.' A sentence an operations manager can act on.

Compliance evidence drafted for you

The PCI DSS or GDPR narrative assembled from the actual journal data, with the gaps named — leaving the auditor conversation to the humans.






Baselines instead of thresholds

A fixed threshold is either noisy or blind. A learned baseline knows what a Tuesday night normally looks like on your partition.

Ask the journal a question

'Who read the payroll file outside business hours last quarter?' answered against the AMR feed and DB2 journal services, without anyone writing SQL.

Making Assure Security agent-callable



Why an MCP server, and not just another API

The knowledge problem

IBM i operational knowledge lives in runbooks and in a pool of people that is shrinking every year.

'Check the system status' is the easy part. Knowing that it is *WRKSYSSTS* is the part that is disappearing.

The integration problem

If every product ships its own connector, that is another entry point to firewall, certificate-manage and audit.


A typed tool layer plugs into whatever AI tooling the customer already runs, instead of a bespoke integration per product.

The trust problem

No security team is going to hand an agent direct access to a production partition on trust alone.

The answer is not a smarter model. It is a narrow, allow-listed, least-privilege surface with a human gate on every write.

How it fits together



Future Concept

Built on interfaces
Assure Security
already exposes.
(Not a committed roadmap,
and no shipping date.)

The Precisely (MCP) Gateway

Where it runs

- One hardened front door. No IBM i MCP server is ever directly exposed.
- Runs on IBM i PASE, on Linux, or in cloud.
- Aggregates products, so a single conversation can span Assure Security, MIMIX and QEDD.
- Python 3.11 and the pure-Python stack are verified on IBM i — no Rust, no compiling.

What it enforces

- mTLS between the agent and the gateway.
- RBAC per tool, per user, per product.
- Rate limiting, and a full audit trail into the customer's own SIEM.
- Security tiers are opt-in flags on the same build.

What it means for you

- Product knowledge becomes agent-callable rather than tribal.
- One integration to review and certify, not one per product.
- You keep the enforcement decision — the agent proposes, your process disposes.
- Forward-looking: direction of travel, not a commitment. Features and timing may change.



Defence in depth, with one solution

No single control answers an AI-driven attack. MFA stops the stolen password. Elevated authority management contains the session that gets through anyway. Exit-point rules close the background paths. Monitoring makes all of it visible to the people whose job it is to respond.

A woman with dark hair, wearing a purple long-sleeved shirt, is smiling and looking at a laptop screen. Her hands are clasped near the screen. To the left, the side of a man's face with glasses is partially visible. The background is a bright, out-of-focus office setting with large windows. The entire image has a light purple overlay.

precisely

**Questions?
Answers!**

precisely

IBM i & AI: Friend or Foe?

9th September — 14:00 - 15:00

Stephan LEISSE
stephan.leisse@precisely.com

