



**National Institute on
Retirement Security**
Reliable Research. Sensible Solutions.

From AI Readiness to Results: Implementing AI in Public Pension Systems

Webinar | July 22, 2026

Agenda

- Logistics
- Speaker Introductions
- From AI Readiness to Results Conversation
- Q&A



Logistics

- Attendees in listen-only mode.
- Submit questions via Q&A function throughout presentation.
- Recording in progress.
- Replay will be available at www.nirsonline.org.events



Today's Speakers



Barbara Butrica

NIRS



Nate Haws

Linea Solutions



Korey McConachie

Linea Solutions

Linea Solutions - Who we are

Business Focus – Pension, Workers' Compensation, and Insurance Consulting

- Provide **strategy, business and technical assessments, project management and change specialists**, assisting in all phases of business and technology transformation
- AI, Cybersecurity & Data as key focus areas (50+ clients)



Linea Solutions AI Practice

Practical AI for Pension and Benefits Organizations

What We Focus On

- Visible, safe, and governed AI Adoption
- Practical use cases tied to real business problems
- AI-enabled modernization knowledge access (*requirements, testing, training, hypercare, etc.*)
- Knowledge assistants built on approved source material
- Human review, decision rights, accountability, and trust

How We Help

- AI strategy, governance, and policy
- Use case discovery, prioritization, and roadmap development
- AI solution implementations
- AI-enabled modernization support
- AI vendor selection and implementation oversight



A Culture of AI, Three Years On

Background

In 2023, we stood here and talked about creating a Culture of AI - organizations where staff use AI tools **every day they work.**



A Culture of AI, Three Years On

Background

Three years later, AI has stopped being just a chat window. It works towards a goal, plans, uses tools, manages memory, and carries out whole pieces of work — **as agents.**



A Culture of AI, Three Years On


Background

And the question has shifted from
“should we allow AI?” to “**how do we
make its work visible, governed,
and useful?**”



What to Expect Today

- The AI vocabulary of 2026: agents, orchestration, skills, MCP and connectors
- What Microsoft's 2026 Work Trend Index says is changing about knowledge work
- A live Public Pension AI demo: a human directed manager of agents works a contact center's morning queue
- The 2023 Culture of AI steps, revisited for the agentic era
- Q&A




The 2026 Work Trend Index

Microsoft's 2026 Work Trend Index describes the rise of Frontier Firms - organizations redesigned around humans directing AI agents, and a new role it says everyone will eventually hold: the agent boss

Work Trend Index Report

- 58% of AI users say they now produce work they could not have produced a year ago
- Only 19% of workers sit in the “frontier zone” where individual capability and organizational readiness align



Source: Microsoft 2026 Work Trend Index



The vocabulary of 2026

AGENTS

- Plan and execute multi-step work, not just answer questions
- Range from a knowledge retrieval chatbot to systems working across applications via orchestration
- Act at human direction, with human review of anything produced
- Grounded in curated, approved sources only

SKILLS & OPERATING INSTRUCTIONS

- Operating Instructions provide the ground rules and goals of an agent – included in every prompt
- Skills are packaged task instructions an agent follows: triggers, roles, hard rules, output formats
- Both written once, reused every run, the asset is the instruction file, not the prompt

MCP SERVERS, CONNECTORS & APIs

- Model Context Protocol (MCP), an open standard connecting AI to tools and API's/data
- How agents reach your systems - scoped, least-privilege, logged and monitored
- Same governance questions as staff access: who, what, and how much

Other Observations:

The skill has moved up a level: from writing the perfect prompt to directing a team of agents.

In 2023...

- You wrote the perfect prompt, one question at a time
- AI was a tool you operated - you reviewed every answer as it arrived



In 2026...

- You write the operating instructions & goals, orchestrate the agent team and delegate tasks in a specific sequence
- AI is a small team you manage, you review the finished, reconciled work



In 2023, we asked what AI could do...

In 2026, we ask what agents can get done.

Next: a supervisor handles a team of agents one January morning at a pension contact center - and reviews the finished, reconciled work.

⚡ Model: Default (GPT-3.5)

BH What can AI do?



AI, or artificial intelligence, refers to the development of computer systems that can perform tasks that typically require human intelligence. AI has a wide range of applications and capabilities. Here are some things that AI can do:

1. Natural Language Processing (NLP): AI can understand and generate human language, allowing for tasks like language translation, sentiment analysis, chatbots, and voice assistants.
2. Machine Learning (ML): AI can learn from data and improve its performance over time. ML algorithms can be used for tasks such as image recognition, speech recognition, predictive analytics, and recommendation systems.
3. Computer Vision: AI can interpret and understand visual information from images and videos. This enables applications like object recognition, facial recognition, image and video analysis, autonomous vehicles, and surveillance systems.
4. Robotics: AI can be integrated into robots to enable them to perform tasks autonomously or with minimal human intervention. This includes tasks like assembly line work, logistics, healthcare assistance, and exploration in hazardous environments.
5. Expert Systems: AI can simulate human expertise in specific domains. Expert systems use knowledge and rules to make decisions or provide recommendations in areas such as medicine, law, finance, and engineering.
6. Data Analysis: AI can process and analyze vast amounts of data to identify patterns, extract insights, and make predictions. This is particularly useful in fields like finance, marketing, healthcare, and cybersecurity.
7. Automation: AI can automate repetitive or mundane tasks, freeing up human resources for more complex and creative work. This includes tasks like data entry, customer support, content generation, and process optimization.
8. Gaming: AI can play games at a high level and compete against human players. Deep reinforcement learning has enabled AI agents to achieve remarkable performance in games like chess, Go, poker, and video games.
9. Personalization: AI can personalize user experiences by analyzing individual preferences and behavior. This is commonly seen in recommendation systems for streaming platforms, online shopping, and targeted advertising.
10. Creativity: AI can generate new content, such as music, art, and writing. It can assist in

Reality check before the demo


Agentic workflows are still leading edge for public pension systems.

Some organizations are finding early wins, but many are still building governance, approved sources, and staff fluency.

The demo is a destination, not the average current state

Important Note: Claude Cowork (the tool being demonstrated) still requires significant security review & we are demonstrating in an isolated environment

A smaller group is experimenting here now



The demo:

GVPERF



Golden Valley Public Employees' Retirement Fund - a fictional pension fund in a common scenario: the peak of 1099-R season.

Note – All demo data is synthetic, and **nothing is ever sent to a member.**

The scenario:



GVPERF is overwhelmed

- 18 contacts arrived overnight - about 3× normal - plus yesterday's 8 calls waiting for QA
- Why? Late January stacks two events: 1099-R tax forms just hit mailboxes, and the monthly pension payment goes out on the 30th
- The forms bring missing 1099-Rs, Box 7 confusion, and address changes; and payment week brings anxious “where’s my deposit?” calls
- And it draws predators: two account-takeover attempts timed to the payment date – plus a widow reporting a death and an attorney’s QDRO deadline

The agent team



Manager (Claude Cowork) sequences the agents, reconciles every item exactly once, and compiles the report.



Triage Agent

Categorizes all 18 items, cross-checks the roster, sets urgency + one action each.



QA Agent

Scores yesterday's 8 calls against the rubric; flags skipped ID checks and tax advice.



Response Agent

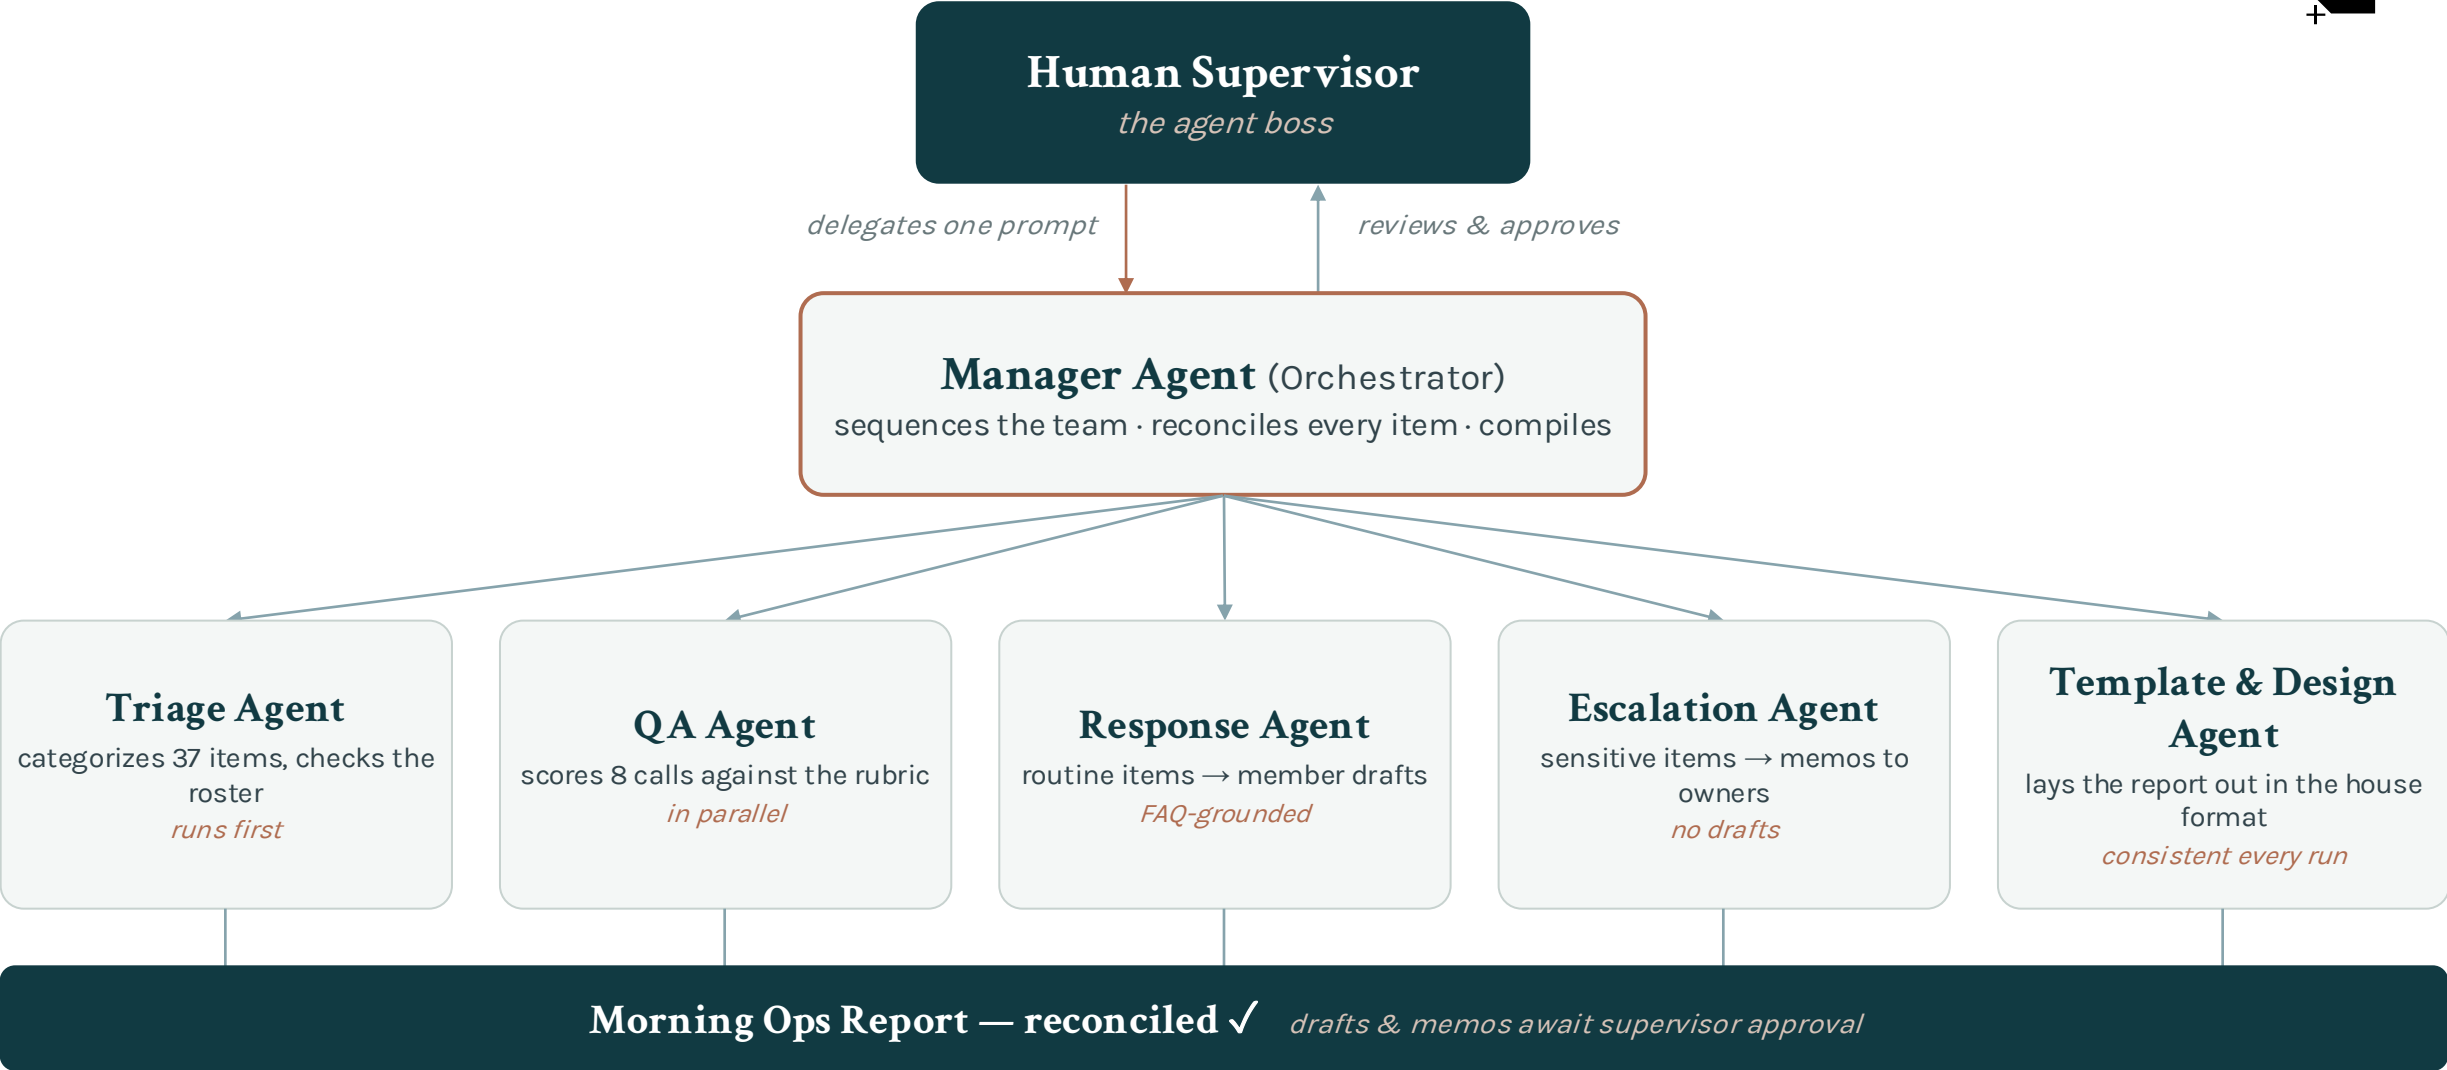
Drafts routine replies only, grounded in the FAQ. No tax advice, always the contact number.



Escalation Agent

Pulls deceased-member, fraud, and QDRO items; writes a memo to a human owner – no member reply.

One prompt in, one reviewed report out



Drafts only – nothing is sent to a member without human approval

What one run produced



18 = 13 + 4 + 1

every item reconciled exactly once: 13 drafts to approve, 4 escalation memos, 1 duplicate merged.

8 calls QA'd

5 commendable, 3 coaching flags — including tax advice caught before the member acted on it.




The payoff: about two hours of manual triage became a short supervisor review — with fraud holds in place before the Jan 30 payment run.

Why a team beats one giant prompt



Specialization

Each agent has one narrow job and tight rules, so each decision is more accurate than a single sprawling prompt.



Parallelism

Triage and QA run at the same time on different inputs – the morning's work compresses instead of queuing.



Auditability

You can see which agent decided what, and the counts prove nothing slipped through the cracks.



How to build a Culture where AI Agents thrive?



OUR WAY IS NOT THE ONLY WAY, BUT WE BELIEVE IT...

- Is systematic
- Protects against risks
- Leverages the expertise and skills within the organization
- Is nimble enough to adapt to the rapid rate of innovation
- Lives on through Continuous Process Improvement

Three Years On

Step 1



Govern Agents Like Staff



- Extend your governance charter to agents: every agent gets an identity, scoped access, and a named owner
- Decide decision rights - what an agent may do alone, and what needs human approval
- Same team as 2023 - leaders, early adopters, IT, training and security/ISO
- Rule of thumb: give an agent the protections and constraints you would give a new hire



Three Years On

Step 2



Curate Sources & Skills



Agents are only as good as what they may read - approve the FAQs, policies, and rosters they cite.

Write the operating instructions once - roles, hard rules, report formats. The reusable asset is the instruction file, not the prompt.

Messy or missing sources produce confident, wrong output.




Three Years On

Step 3



Connect Systems Safely



MCP servers, connectors, and APIs let agents reach CRM, telephony, and admin systems - scope access narrowly, read-only where possible.

Treat every connector like a third-party integration: **vet it**, log it, monitor it.

Start with folders and documents; graduate to systems integrations as trust builds and security is proven.

Three Years On

Step 4



Pilot One Orchestrated Workflow



Pick one recurring, painful process - a morning queue, an inbox, a reconciliation

One manager agent, a few narrow specialists, one reviewed deliverable

Dry-run repeatedly: sensitive items must always escalate, and counts must reconcile

Prove it on synthetic data before it touches member data

Testing is still very important to build trust in the systems and outputs - automate regression testing where possible



Three Years On

Step 5



Keep Humans in the Loop



Drafts only: nothing reaches a member without appropriate approval

Escalations go to people, not auto-replies - even when the reply looks easy


Watch the reviewers: time saved generating must not become burden shifted downstream

Three Years On

Step 6



Train Agent Bosses



The 2026 skill is delegation: writing instructions, reviewing output, spotting failure

Move AI fluency training from only prompting to directing and verifying

Subject-matter expertise matters more than ever - experts catch what agents miss

Three Years On

Step 7



Iterate, Measure & Scale



Reconcile every run: items in = items out – trust is built on counts that prove it

Feed QA findings and coaching themes back into the instruction refinements

Scale by adding agents, not headcount - then revisit steps 1-6 as tools evolve





1. Govern Agents Like Staff
2. Curate Sources & Skills
3. Connect Systems Safely
4. Pilot One Orchestrated Workflow
5. Keep Humans in the Loop
6. Train Agent Bosses
7. Iterate, Measure & Scale

A Culture of AI, Three Years On

Summary of Steps



From exploration to implementation

ASSESS THE OPPORTUNITIES

- Define Governance, Policy & Strategy
- Start from pain, not hype: queues, backlogs, reconciliations
- Score candidates on impact/value, risk, data readiness, ROI
- Keep a prioritized backlog
 - governance approves each use case

EVALUATE & PROCURE

- Decide build vs. buy vs. configure, use case by use case
- Vendor diligence: data use/residency, IP ownership, liability, security attestations
- Prove it on synthetic/obfuscated data before a use case scales
- Security is paramount

IMPLEMENT

- Varying levels of technical expertise required based on the agent platform (Pro Code / Low Code / No Code)
- Operating instructions are plain English - written by your SMEs, not necessarily engineers for business tasks
- More integration -> more technical expertise & security
- Leaders govern outcomes: counts, escalation triggers, approval gates
- Your role is the agent boss: delegate, review, decide




Please feel free to reach out to:

Nathan Haws

Senior Director of AI, Linea Solutions

nhaws@lineasolutions.com

Questions?



Check out NEW NIRS Podcast



<https://www.nirsonline.org/research/retirementinamericapodcast-episode3/>