



Data Readiness and Knowledge Readiness for AI

Opening Note

In nearly every NAED PowerShift® AI conversation, the same reality eventually emerges. A distributor can identify a compelling use case, select the right AI platform, and launch a well-designed pilot, yet the outcome still depends on information the organization may not fully trust, cannot find quickly, does not own clearly, or has never organized for reuse.

For electrical distributors, that information is spread across ERP transactions, item and customer masters, supplier portals, product files, quotes, emails, spreadsheets, branch practices, project notes, operating procedures, and the judgment of experienced employees. AI can make this information easier to access, connect, and use, but it can also expose inconsistency, staleness, duplication, and hidden dependencies far more quickly than traditional systems.

As AI adoption accelerates across the electrical distribution industry, conversations are shifting from experimentation to measurable business outcomes. Organizations are looking for ways to improve customer responsiveness, increase productivity, strengthen decision making, preserve institutional expertise, and reduce manual work. Many are discovering that the limiting factor is no longer the AI platform itself. It is the readiness of the information environment that AI depends upon. The distributors that realize the greatest value from AI will not necessarily be those with the largest technology budgets. They will be those that can consistently provide trusted information, clear business context, and governed operational knowledge.

Historically, experienced employees bridged gaps in systems, documentation, and process through years of accumulated judgment. AI does not eliminate the need for that expertise. It creates an opportunity to preserve more of it, make it reusable across the organization, and combine it with trusted data to improve how decisions are made. Organizations that strengthen both their data readiness and their knowledge readiness will be better positioned to implement AI successfully while building more resilient operations over time.

This whitepaper does not argue that distributors should clean every record, standardize every process, or build a perfect enterprise data platform before they begin. That approach is unrealistic and often becomes a reason to delay. Instead, it advocates use-case-specific readiness:

- Define the business problem
- Identify the information required
- Determine whether that information is trusted enough for the intended purpose
- Correct the few gaps most likely to undermine value

The objective is practical. Rather than treating data readiness and knowledge readiness as standalone initiatives, this paper presents a framework for improving the information foundation while solving real business problems. Every successful AI initiative should not only deliver business value today, but also leave the organization with stronger data, better knowledge, clearer governance, and a more capable foundation for tomorrow.

For additional insight or a deeper conversation about these topics, including how to apply the readiness frameworks presented in this document, please feel free to reach out.

Thank you,

Scott A. Wagner | Director of Industry Transformation | NAED
C 205.542.0402 | www.naed.org/dcoe | swagner@naed.org

Executive Summary

AI does not create reliable business information. It depends on it. When an AI-assisted workflow produces weak output, the AI platform may be part of the problem, but the deeper cause is often incomplete data, conflicting definitions, inaccessible source material, outdated documents, inconsistent process execution, or knowledge that exists only in the minds of experienced employees.

Electrical distributors are especially exposed because value is created through a combination of structured records and operational judgment. ERP and related systems may hold order, customer, item, pricing, inventory, purchasing, branch, and rebate information. The meaning required to act on those records often lives elsewhere. A buyer knows which supplier date is believable. An inside salesperson understands the customer's terminology. A pricing analyst knows why a special pricing agreement or rebate claim will fail. A warehouse supervisor knows why available inventory and physical inventory do not appear to agree. AI cannot apply that context responsibly until enough of it is visible, governed, and reusable.

As AI becomes more accessible, competitive advantage will shift toward distributors that can consistently provide trusted information, governed knowledge, and meaningful business context. Most organizations will have access to similar AI capabilities. Far fewer will have the information foundation needed to produce reliable, scalable business outcomes.

The practical standard is not perfect data. It is information that is ready enough for a defined use case. That requires two closely connected disciplines:

- **Data readiness.** Structured and transactional information is accurate, complete, consistent, timely, owned, accessible, and understood well enough to support a specific workflow.
- **Knowledge readiness.** Documents, procedures, supplier and product guidance, operating rules, communications, and human expertise are findable, trusted, current, access-controlled, and maintained well enough for people and AI systems to use reliably.

The two disciplines must work together.

- An order-status assistant can have clean order records and still make a weak customer commitment if it cannot interpret supplier acknowledgements or project-release notes.
- A guided-selling tool can have product attributes and still recommend the wrong item if it does not understand units of measure, customer terminology, or approved substitution rules.
- A pricing assistant can have current cost and price data and still create margin leakage if agreement terms, effective dates, or rebate logic are missing.

Data provides the facts. Knowledge provides the context needed to apply those facts correctly.

This whitepaper presents a practical readiness framework built around seven dimensions: accuracy, completeness, consistency, timeliness, accessibility and permission, ownership and stewardship, and context and interpretability. Together, these dimensions help leaders determine whether a use case is ready now, needs foundation, or not now, while focusing improvement efforts where they deliver the greatest business value.

Perhaps the greatest opportunity AI creates is the ability to preserve and scale operational expertise that has traditionally lived only in the experience of trusted employees. Capturing more of that knowledge, alongside trusted data, creates a stronger foundation for better decisions, faster onboarding, greater consistency, and future AI initiatives.

The leadership message is straightforward. Do not launch a broad data-cleanup program simply because AI has arrived. Start with one meaningful workflow, identify the data and knowledge it requires, determine what is trusted enough for the intended purpose, assign clear ownership, and correct the few gaps most likely to undermine value. AI should become a catalyst for improving the organization's information foundation while delivering measurable business value.

Contents

Opening Note 1

Executive Summary 2

AI Is a Mirror of Your Operating Environment..... 4

What Data Readiness and Knowledge Readiness Mean 4

Human Bridges: Value, Risk, and the Knowledge Capture Opportunity 5

A Use-Case-Specific Readiness Standard 6

The Seven Dimensions of Readiness 6

Authority, Ownership, and Stewardship..... 8

Assessing Readiness Without Launching a Data Program 9

Turning Readiness Gaps into a Controlled Pilot 10

How Readiness Supports NAED PowerShift® 11

The 30, 60, 90-Day Readiness Roadmap 11

Conclusion: Build the Foundation One Use Case at a Time 12

Appendix A: Data and Knowledge Readiness Scorecard..... 13

Appendix B: Source Authority Matrix 13

Appendix C: Human Bridge Discovery Guide 14

Appendix D: Readiness Remediation Plan 14

AI Is a Mirror of Your Operating Environment

AI can search, summarize, classify, compare, predict, generate, and orchestrate work at a speed humans cannot match. That capability is real, but it does not operate outside the business. It reflects the quality of the information, process, ownership, and decision standards underneath it.

For an electrical distributor, that reflection can be uncomfortable:

- If leadership has three versions of gross margin, fill rate, or on-time delivery, AI does not know which one the business trusts unless authority is defined.
- If a supplier portal shows one shipment date and the ERP shows another, AI can repeat the conflict faster without resolving it.
- If item descriptions, manufacturer part numbers, pack quantities, or units of measure are inconsistent, guided selling can become faster without becoming safer.
- If special pricing agreements, rebate rules, or ship-and-debit requirements live in email and spreadsheets, a pricing recommendation can look precise while missing the economics that protect margin.
- If counter and inside sales teams rely on customer shorthand, project history, or application context that is not documented, product matching can miss what an experienced employee immediately understands.
- If branches use local workarounds because the standard process does not fit every operating condition, AI may learn uncontrolled variation as if it were policy.

Weak AI output is therefore often diagnostic. It reveals conditions that existed before the tool arrived: inconsistent master data, unreliable supplier dates, undocumented judgment, fragmented repositories, weak version control, unclear ownership, or process variation the organization has learned to work around.

The wrong response is to blame the model and move on. The equally weak response is to declare that every data issue must be solved before AI can be used. The mature response is to identify which condition affected the result, determine whether it can be bounded or corrected within the use case, and decide whether the pilot should continue, narrow, or pause.

What Data Readiness and Knowledge Readiness Mean

Data Readiness

Data readiness concerns the structured and semi-structured information used to describe transactions, entities, status, quantities, dates, relationships, and outcomes. In distribution, this may include item, customer, and supplier masters; inventory and available-to-promise positions; sales orders and purchase orders; branch and warehouse transactions; costs and prices; special pricing agreements; rebate and claim records; contracts; delivery events; and historical performance.

Data can exist and still be unready. A field may be populated but not trusted. An inventory balance may ignore an allocation or transfer. A promised date may reflect the last update rather than the customer's original commitment. A unit of measure may be correct in one system and interpreted differently in another. A cost may be current while the related agreement or rebate treatment is not. Accessibility may depend on a manual export that only one analyst knows how to produce.

Data readiness is not simply whether data is present. It is whether the data is fit for the intended decision.

Knowledge Readiness

Knowledge readiness concerns the information and context people use to interpret data and perform work. It includes procedures, product documentation, supplier instructions, customer and contractor history, project notes, quote notes, training materials, decision rules, escalation paths, exception handling, lessons learned, and the expertise developed through years of operational experience.

Knowledge can be valuable and still be unusable at scale. The approved procedure may be buried in a shared drive. A supplier rule may exist only in a buyer's email. A branch may have developed a local workaround that nobody else understands. A new employee may interrupt three people to answer a routine customer question. A subject-matter expert may know why a report is wrong but have no practical way to transfer that understanding.

Knowledge readiness means the organization can find, trust, interpret, control access to, and maintain the information required to act.

Explicit, Tacit, and Institutional Knowledge

Distributor knowledge takes different forms. Understanding those forms matters because each requires a different readiness response.

Knowledge type	Typical distributor examples	Readiness requirement
Explicit knowledge	SOPs, product and technical documentation, approved policies, pricing rules, ERP instructions, role guides, and training materials.	Govern source authority, ownership, version control, permissions, retrieval, and review cadence.
Tacit knowledge	Operational judgment, exception handling, supplier experience, customer familiarity, product application knowledge, and workflow expertise.	Observe decisions, capture recurring drivers and exceptions, validate with experts, and preserve human review where judgment remains necessary.
Institutional knowledge	Shared branch practices, cross-functional assumptions, informal escalation paths, project-handling conventions, and accumulated operating history.	Identify where shared understanding influences work, then decide what should be standardized, governed, locally preserved, or left judgment-based.

Explicit knowledge provides consistency and scalability. Tacit knowledge provides judgment, interpretation, and context. Institutional knowledge preserves the shared understanding of how work is actually accomplished.

Data tells the organization what is recorded. Knowledge explains how to interpret it and what to do next.

Human Bridges: Value, Risk, and the Knowledge Capture Opportunity

Most distributors can name at least one **go-to person** who connects gaps between systems, departments, information sources, and real operating conditions. These people are often not senior leaders. They may be a buyer, counter or inside salesperson, branch operations manager, warehouse supervisor, pricing analyst, ERP power user, or process champion. They know how to get things done when the formal process, the system answer, and the customer need do not line up.

These employees function as **human bridges**. They know which supplier commitments are dependable, which item descriptions map to a customer's terminology, why a transfer temporarily distorts inventory, which project release instructions matter, how a rebate or special pricing claim must be documented, and when a local exception protects customer service without creating unacceptable risk.

The value is obvious: customers are served and operations continue. The risk is less visible. Expertise becomes concentrated in a small number of people. Onboarding takes longer. Decisions vary by branch or employee. Work slows when the expert is unavailable. The organization may also automate the documented workflow without understanding the real one.

The problem is not that expert judgment exists. The problem is that the organization has not converted enough of the repeatable, high-value portion into reusable organizational capability.

Signal	What it may reveal	Readiness action
Side spreadsheets	The system of record does not support the decision or is not trusted.	Document purpose, owner, inputs, logic, and whether to integrate, govern, or retire.
Repeated interruptions	Knowledge is hard to find or concentrated in a few people.	Capture recurring questions, approved answers, sources, and escalation.
Manual reconciliation	Sources disagree, identifiers do not align, or timing differs.	Define source authority and the minimum matching logic.
Unofficial reports	Official reporting does not reflect operating reality or accepted definitions.	Compare definitions, lineage, and decision use before choosing a source.
Local branch workarounds	The standard process does not fit all operating conditions.	Separate legitimate local exceptions from uncontrolled variation.
Expert-only judgment	Decision criteria are tacit or inconsistently applied.	Capture drivers and exceptions, then test decision support before automation.

AI as a Catalyst for Knowledge Capture

Historically, preserving tacit knowledge required time-consuming interviews, procedure projects, and formal training development. Those efforts often captured only a small portion of what experienced employees actually used to make decisions.

AI changes the economics of that work. It can help structure expert interviews, summarize recurring questions, classify corrections and overrides, draft decision trees, turn conversations into searchable knowledge, and identify where different experts apply different rules. It can make more of the organization's operating reasoning visible for review.

That does not make AI the authority. Experienced employees and business owners remain the designers, validators, and teachers. They determine whether captured knowledge reflects operating reality, which parts can be reused, which exceptions need governance, and which decisions must remain judgment-based.

The objective is not to document everything a person knows. That is neither possible nor useful. The objective is to make the repeatable, high-value portion available in the workflow while preserving escalation to experienced people when the situation still requires judgment.

A Use-Case-Specific Readiness Standard

The most damaging readiness mistake is making the standard too broad. Leaders hear that AI needs clean data and conclude that the company must solve enterprise data quality before doing meaningful work. That approach is usually too expensive, too slow, and too disconnected from business value.

Readiness should be evaluated against a defined use case. The information required to summarize a meeting is different from the information required to recommend a product at the counter, commit to a project delivery date, predict a stockout, initiate supplier follow-up, or recommend a price. The acceptable error rate, freshness requirement, approval threshold, and risk exposure also differ.

A use case is **ready enough** when the team can define the business problem, identify the required sources, explain which source is authoritative, access the information at the necessary cadence, understand the major limitations, assign owners, and establish human review proportional to risk. Ready enough does not mean perfect. It means the weaknesses are known, bounded, and manageable.

This standard can be expressed as **minimum viable readiness (MVR)**: the smallest set of trusted data, documented rules, source authority, permissions, ownership, and human controls required to learn safely from a pilot.

Classification	Meaning	Leadership response
Ready now	The problem, sources, ownership, scope, and review model are clear enough to learn safely.	Launch a narrow pilot and use the pilot to improve the foundation.
Needs foundation	The opportunity is valuable, but one or more information gaps would make the test misleading or unsafe.	Define the minimum corrective work, assign owners, and set a readiness date.
Not now	Value is vague, risk is excessive, ownership is absent, or the information environment is too weak for a useful test.	Deprioritize without killing the idea. Document what would need to change.

The Seven Dimensions of Readiness

The same seven questions can be applied to both data and knowledge. The required score is not universal. It should reflect the consequence of the decision and the strength of the human control around it.

1. Accuracy

Accuracy asks whether the information reflects operating reality closely enough for the intended use. A customer-facing delivery commitment requires a higher standard than internal analysis. An inventory quantity can be numerically correct but still misleading if allocations, holds, transfers, or damaged stock are excluded. A cost can be current while the agreement or rebate context required to interpret margin is not.

Test accuracy against known outcomes and experienced users, not just field-population statistics. A populated date can be wrong. A complete description can still point to the wrong product. A mathematically correct score can be operationally useless.

2. Completeness

Completeness asks whether the use case has the fields, documents, rules, and context required to produce a useful result. Missing information forces the AI or the user to infer. For guided selling, a product description may be insufficient without voltage, amperage, poles, interrupting rating, dimensions, material, unit of measure, or approved substitution rules. For project status, order data may be insufficient without release instructions and customer commitments.

Distinguish critical completeness from desirable completeness. The pilot does not need every attribute or historical record. It does need the elements without which the output cannot be judged or safely used.

3. Consistency

Consistency asks whether the same entity, event, status, or rule is represented the same way across sources. Common distribution examples include customer and supplier identifiers, manufacturer part numbers, units of measure, pack quantities, branch codes, product categories, promised dates, active-customer definitions, and margin measures.

AI can help reconcile differences, but the matching and precedence rules must be owned. Otherwise the system may hide disagreement rather than resolve it.

4. Timeliness

Timeliness asks whether the information is current enough for the decision. A historical analysis may tolerate a weekly refresh. An order-status workflow may require near-real-time supplier or carrier events. A pricing or rebate workflow may need current effective dates, program terms, and cost changes before a recommendation is released.

Define the required refresh cadence and the maximum acceptable age of information before the pilot begins. If the source cannot meet that requirement, the output should show the age and the resulting uncertainty.

5. Accessibility and Permission

Accessibility asks whether the right employees, systems, and AI tools can retrieve the information when needed. Permission asks whether they should. Those are separate questions. Information may be technically available but restricted by customer, supplier agreement, role, location, confidentiality, or commercial sensitivity.

A mature design does not solve access by copying ERP, CRM, PIM, e-commerce, supplier portal, and document content into one uncontrolled repository. It aligns access to the workflow, preserves source permissions where possible, and records how information is used.

6. Ownership and Stewardship

Ownership asks who is accountable for the source, definition, correction path, and update cadence. Stewardship is the recurring work required to keep that source usable. Without ownership, quality becomes everyone's concern and no one's responsibility. Without stewardship, a one-time cleanup decays.

The owner does not need to fix every record or document personally. The owner does need authority to define the standard, resolve conflict, prioritize correction, and ensure the process stops recreating the same defect.

7. Context and Interpretability

Context asks whether users and systems understand what the information means, where it came from, what limitations apply, and how it should influence action. Interpretability is especially important when AI produces a product match, risk score, pricing recommendation, customer response, or summary.

A recommendation without drivers is hard to trust. A retrieved answer without source traceability is hard to verify. A metric without a definition invites conflict. Context converts information from raw material into usable business judgment.

A low readiness score does not automatically stop a use case. It identifies where scope, remediation, source traceability, or human review must be stronger.

Authority, Ownership, and Stewardship

AI-assisted search and decision support require more than access to information. They require **source authority**. When several systems, files, portals, and people hold related information, the business must define which source controls the answer, under what conditions, and who resolves disagreement.

No single system is authoritative for every fact. The ERP may own posted order quantity while a supplier portal or carrier event owns the latest shipment status. A PIM or manufacturer source may own technical attributes while a customer-specific cross-reference lives elsewhere. A pricing system may own an approved special pricing agreement, but a sales note cannot silently override it. Rebate eligibility may depend on manufacturer program terms and claim rules that are not represented in the sales transaction itself.

Information need	Typical authoritative source	Supporting context	Conflict owner
Posted order quantity	ERP or order management system	Customer or project notes	Order process owner
Current supplier shipment event	Supplier portal, EDI/API event, or carrier source	Buyer notes and exception history	Purchasing owner
Technical product attribute	PIM or approved manufacturer content	Cross-references and application notes	Product data owner
Customer-specific price or SPA	Approved pricing or agreement system	Quote history and sales context	Pricing or sales operations owner
Rebate or ship-and-debit eligibility	Manufacturer program terms and claim system	Analyst notes and prior exceptions	Rebate or finance owner
Safety, HR, or operating policy	Approved policy repository	Training guide and job aid	Named policy owner

An authority model should be explicit enough that the AI retrieves from approved sources, the user can open the underlying evidence, permissions are preserved, staleness is visible, and disagreements route to a named owner. Retrieval-augmented generation (RAG) can reduce unsupported answers, but it cannot settle governance questions the business has never answered.

Business-Led, IT-Partnered Ownership

Data and knowledge readiness must be **business-led and IT-partnered**. The business owns meaning, workflow, value, decision rules, risk thresholds, and adoption. IT and data teams enable access, integration, architecture, security, permissions, monitoring, and scalable controls.

This distinction matters because many readiness problems are not technical. Purchasing must define which supplier dates matter. Sales must define what makes a counter, quote, or project response acceptable. Operations must define exception categories. Finance and pricing must define margin, agreement, and rebate logic. IT cannot invent those rules without creating a technically organized version of business ambiguity.

IT and data teams should not be treated as order takers. They bring the reality of data lineage, integration effort, vendor architecture, access controls, supportability, and security. Where a manufacturer, supplier, or rep controls a source, an internal owner must still define how the information is consumed, validated, and escalated.

Role	Primary responsibility
Executive sponsor	Connect readiness work to strategy, resolve cross-functional tradeoffs, and protect focus.
Business process owner	Define the workflow, decision, value, source meaning, exceptions, risk threshold, and adoption standard.
Data owner or steward	Define quality rules, correction paths, lineage, and maintenance cadence.
Knowledge owner	Approve content, identify authoritative sources, manage versioning, and retire outdated material.
IT / data partner	Enable secure access, integration, architecture, monitoring, permissions, and scalability.
Frontline manager	Validate operating reality, reinforce correct use, and surface recurring gaps.
Subject-matter expert	Explain tacit judgment, test outputs, and help define exceptions and quality thresholds.

Assessing Readiness Without Launching a Data Program

A readiness assessment should be fast enough to support a decision and deep enough to prevent self-deception. It is not an enterprise maturity assessment unless the organization explicitly needs one. For most early AI work, the right unit of assessment is the use case.

Begin with the workflow, not the data inventory. Define the business decision or task, the user, the input, the output, the action that follows, and the measure of improvement. Then identify the information required at each step. This prevents the assessment from becoming a broad cataloging exercise with no connection to value.

A practical assessment can be completed through five moves:

1. **Map the workflow.** Show where information enters, changes, is interpreted, and becomes a commitment or action.
2. **Inventory required sources.** List systems, files, documents, communications, supplier feeds, and human expertise used today.
3. **Score the seven readiness dimensions.** Use evidence, not optimism. Record limitations and uncertainty.
4. **Identify critical blockers.** Separate issues that would invalidate the pilot from issues that can be managed through scope or human review.
5. **Assign a readiness decision.** Classify the use case as ready now, needs foundation, or not now, then name owners and timing.

The assessment should produce decisions, not a large report. The most useful outputs are a workflow map, a source map, a readiness scorecard, a short list of blockers, named owners, and a launch or remediation date.

Minimum Viable Readiness by Distributor Use Case

Use case	Critical data and knowledge	Minimum viable readiness	Human control
Internal knowledge assistant	Approved procedures, policies, product documents, training materials, metadata, and permissions.	Authoritative repository defined; outdated content marked or removed; source citations available; access rules preserved.	Users verify high-impact answers and escalate conflicts.
Customer order and project-status support	Sales orders, purchase orders, supplier acknowledgements, shipment events, customer commitments, project-release instructions, and exception notes.	Order, PO, and shipment linkage is reliable; timestamps and status definitions are consistent; project instructions are findable.	Employee approves external commitments and communicates uncertainty.
Counter sales, product matching, and guided selling	Item attributes, manufacturer numbers, units of measure, cross-references, customer terminology, application context, and substitution rules.	Priority categories have sufficient attributes; matching confidence is visible; approved substitutes and application limits are governed.	Sales or product expert confirms ambiguous matches and substitutions.
Inventory and stockout-risk support	Branch inventory, allocations, transfers, open demand, inbound supply, lead-time history, and criticality rules.	Core quantities and dates are trusted for pilot scope; transfer, allocation, and demand definitions are fixed; refresh cadence is known.	Planner or branch leader reviews risk drivers and selects intervention.
Supplier follow-up and vendor alignment	Open PO status, acknowledgement rules, supplier contacts, portal or EDI events, templates, and escalation policy.	Triggers are accurate enough to avoid false outreach; contacts and templates are current; ownership and escalation are clear.	Buyer reviews early messages and owns exceptions.
Pricing, SPA, rebate, and margin guidance	Current cost and price, customer and product segments, agreements, rebate or ship-and-debit terms, effective dates, margin rules, and exception logic.	Definitions are controlled; agreement and program data are current; ownership and approval thresholds are explicit.	Authorized employee approves recommendations and customer-facing action.
Onboarding and branch training support	SOPs, role guides, product content, FAQs, escalation paths, local exceptions, and examples of good judgment.	Approved content is findable, current, role-specific, and access-controlled.	Manager reinforces judgment, exceptions, and real-world application.

The pattern is consistent: the more a use case affects a customer commitment, product selection, margin, safety, or automated action, the stronger the source authority, traceability, and human control must be.

Turning Readiness Gaps into a Controlled Pilot

Prioritize the Gaps That Matter

Once gaps are visible, the organization needs discipline in deciding what to fix. Not every defect has equal value. A missing attribute that affects a high-volume counter or guided-selling workflow may matter more than thousands of low-use records. An outdated safety procedure deserves faster attention than a formatting inconsistency. A supplier mapping problem that blocks inbound visibility may justify immediate work, while a historical field unrelated to the pilot can wait.

The organization must also distinguish one-time cleanup from operating control. If the process continues creating bad data or outdated knowledge, cleanup without control only delays recurrence.

Priority factor	Leadership question
Business impact	What customer, margin, service, capacity, cash, or risk outcome is blocked?
Frequency	How often does the gap affect work or AI output?
Decision criticality	Could the gap create a wrong commitment, product recommendation, unsafe action, or material financial effect?
Dependency	Does this issue block several use cases or only one narrow scenario?
Correction effort	Can the issue be corrected through a rule, mapping, targeted cleanup, content decision, or process change?
Sustainability	What prevents the issue from returning after cleanup?

A strong remediation plan is narrow, owned, and time-bound. It states what will be corrected, what standard will apply, who owns the source, how the fix will be tested, and what evidence will indicate that the use case is ready.

Run the Pilot with Imperfect Information

No distributor should expect perfect information before piloting AI. The goal is controlled learning. A pilot can proceed with known weaknesses when the scope is narrow, consequences are limited, human review is strong, and the team is explicit about what the pilot is testing.

- A product-matching pilot may begin with one product family.
- A stockout model may begin with selected branches and suppliers.
- An order-status workflow may begin with a limited customer and supplier set.
- A pricing or rebate workflow may begin with one agreement type.

Narrow scope creates a cleaner learning environment and prevents unrelated enterprise problems from overwhelming the test.

Do not allow **silent compensation**. If experienced employees repair every weak output without recording why, the pilot may appear successful while remaining dependent on the same human bridge. Capture whether the correction came from data, knowledge, source conflict, prompt design, workflow definition, model behavior, or user judgment.

Pilot control	Purpose
Bounded scope	Limits exposure and isolates the information environment being tested.
Known-limitations log	Prevents the team from rediscovering the same weakness repeatedly.
Source traceability	Allows users to verify where an answer, recommendation, or summary came from.
Human review standard	Matches review depth and approval authority to the business risk of the output.
Correction and override capture	Turns expert judgment into reusable learning and exposes recurring gaps.
Weekly readiness review	Separates tool issues from information, process, and workflow issues.
Stop, refine, or scale decision	Prevents indefinite pilots and forces evidence-based action.

Measure Readiness Through Business Usefulness

Readiness should be measured through business usefulness, not only technical quality. Field completeness, duplicate rates, retrieval accuracy, and document freshness matter, but they should connect to the workflow the organization is trying to improve. Early measures should show whether information is easier to find, trust, and verify. Later measures should show whether the improved foundation changes operating performance.

Measure category	Examples
Source quality	Completeness of critical fields; invalid units of measure; duplicate rate; date accuracy; obsolete-document rate.
Findability	Time to locate approved information; search success; zero-result rate; reduction in repeated interruptions.
Trust and traceability	Percent of outputs with source citations; disagreement rate; verification effort; confidence by defined user group.
Workflow	Quote or order-status response time; rework; inventory-exception resolution; claim exception rate; order or product-match accuracy.
Adoption	Use of approved sources; manager reinforcement; active users in the defined workflow; participation in corrections and overrides.
Sustainability	Owner review completion; stale-content rate; recurring-defect rate; time to correct priority issues.

The strongest measure is not, 'our data is cleaner.' It is, 'the organization can now make this decision faster, more consistently, and with less risk because the required information is more usable.'

How Readiness Supports NAED PowerShift®

The NAED PowerShift® FORGE maturity path helps distributors move from AI urgency to measurable business value. Readiness is not a gate that must be completed before that path begins. It is an operating discipline that matures through every stage.

FORGE stage	Readiness responsibility
Fear / Friction	Identify where weak data, inaccessible knowledge, manual reconciliation, side spreadsheets, or expert dependency is creating drag.
Opportunities	Define the information required for each use case and classify the opportunity as ready now, needs foundation, or not now.
Run Pilots	Use bounded tests to validate source quality, reveal missing context, and capture corrections from experienced users.
Govern & Scale	Formalize source authority, ownership, permissions, quality thresholds, review rules, and update cadence.
Embed Adoption	Place trusted information inside the workflow and reinforce how employees verify, correct, and improve it.

A readiness gap discovered during a controlled pilot is useful evidence. The same gap discovered after a customer commitment, product recommendation, pricing action, or automated message is exposure. The purpose of readiness is not to stop learning. It is to move the learning into a controlled environment.

The 30, 60, 90-Day Readiness Roadmap

First 30 Days: Define and Diagnose

- Select one meaningful distributor use case and define the decision or workflow it should improve.
- Name the business owner, intended users, measurable outcome, and risk level.
- Map the current workflow, including manual handoffs, side spreadsheets, and human bridges.
- Inventory the required data, documents, supplier sources, and expert knowledge.
- Define source authority and minimum acceptable accuracy, completeness, timeliness, and permission rules.
- Score the seven readiness dimensions and classify the use case.
- Baseline one workflow measure and one readiness measure.

Days 31 to 60: Correct the Highest-Value Blockers

- Complete targeted cleanup, mapping, document retirement, metadata, permission, or integration work.
- Document critical decision rules and recurring exceptions that currently live with experienced employees.
- Establish owners, correction paths, conflict rules, and maintenance cadence for priority sources.
- Design the pilot with source traceability, human review, and correction capture.
- Test the information flow before testing the full AI workflow.
- Set the launch criteria, stop criteria, and readiness review cadence.

Days 61 to 90: Pilot and Decide

- Launch a narrow pilot with defined users, scope, cadence, and exit criteria.
- Review corrections and classify root causes weekly.
- Measure both information quality and workflow improvement.
- Decide whether to stop, refine, or scale.
- Before scaling, formalize ownership, controls, training, support, and monitoring.

Conclusion: Build the Foundation One Use Case at a Time

Data readiness and knowledge readiness are not reasons to delay AI indefinitely. They are reasons to move with discipline. Distributors do not need perfect enterprise information before they begin. They do need enough clarity to know what a use case depends on, where the information comes from, who owns it, how it will be verified, and what weaknesses could undermine the result.

The practical danger is not simply bad data. It is **hidden uncertainty**. When teams do not know which source is authoritative, when experienced employees silently repair weak outputs, when outdated documents remain searchable, or when managers cannot explain why reports disagree, the organization creates confidence without control. AI can make that confidence look more polished and spread it more quickly.

The opportunity is larger than avoiding failure. Readiness work can improve customer commitments, reduce manual reconciliation, strengthen inventory and supplier decisions, protect margin, improve rebate recovery, accelerate onboarding, and preserve the expertise of the people who keep the business moving when systems and processes fall short.

The mature path is use-case-specific, business-led, and iterative.

- Start with a real friction point.
- Identify the information required.
- Make uncertainty visible.
- Fix the highest-value blockers.
- Pilot with human judgment.
- Capture corrections.
- Assign ownership.
- Scale only when the information foundation and the workflow can support it.

Start with one use case. Put ownership around the information. Capture expert judgment. Scale only what the organization can support. Crawl, Walk, Run, Fly!

Appendix A: Data and Knowledge Readiness Scorecard

Score each dimension from 1 to 5 for the specific use case. A low score does not automatically stop the work. It identifies where controls, remediation, or narrower scope are required.

Dimension	1: Weak	3: Usable with controls	5: Strong
Accuracy	Frequent errors; users routinely correct or ignore the source.	Known limitations; acceptable for bounded use with review.	Validated against outcomes; trusted for the intended decision.
Completeness	Critical fields, documents, rules, or context are missing.	Most required elements exist; gaps are visible and manageable.	Required information is complete for the workflow.
Consistency	Identifiers, definitions, units, or rules conflict across sources.	Mappings or authority rules handle common conflicts.	Sources align or differences are explicitly governed.
Timeliness	Information is stale or refresh cadence is unknown.	Cadence is defined but may require manual steps.	Information is current at the level required by the workflow.
Accessibility	Users or systems cannot retrieve information reliably or permissions are unclear.	Access exists for pilot users with some manual support.	Authorized access is reliable, secure, and embedded in the workflow.
Ownership	No clear owner or correction path.	Owner is named but stewardship is inconsistent.	Owner has authority, standards, review cadence, and correction process.
Context	Meaning, limitations, source authority, and exceptions are unclear.	Key definitions and caveats are documented.	Sources are interpretable, traceable, and supported by decision rules.

Suggested decision guidance:

- **Ready now.** No critical dimension is below 3, and high-risk outputs have defined human review.
- **Needs foundation.** One or more critical dimensions are below 3, but targeted remediation is achievable.
- **Not now.** The use case lacks ownership, source authority, measurable value, or a safe way to bound uncertainty.

Appendix B: Source Authority Matrix

Information need	Authoritative source	Supporting sources	Owner	Refresh or review cadence	Conflict rule
Order quantity					
Supplier promised date					
Approved substitution					
Customer-specific pricing or SPA					
Rebate or ship-and-debit rule					
Safety or policy guidance					

Use this matrix to distinguish the source that controls the answer from sources that provide context. The conflict rule should explain what happens when sources disagree and who has authority to resolve the issue.

Appendix C: Human Bridge Discovery Guide

Use these questions with experienced employees, managers, and new hires to identify knowledge that is carrying the workflow outside formal systems.

- What questions do people repeatedly bring to you?
- Which reports, fields, dates, costs, or system messages do you not fully trust?
- What spreadsheets, notes, bookmarks, email folders, or personal files do you rely on?
- Where does the official process differ from what actually happens?
- What exceptions require judgment, and what signals influence your decision?
- Which customer, contractor, supplier, product, project, or branch situations require special handling?
- Which pricing, agreement, rebate, or claim rules are understood by experience but not clearly documented?
- What mistakes do new employees make because the real rule is not documented?
- What would stop working if you were unavailable for two weeks?
- What knowledge should be standardized, and what must remain judgment-based?
- What source should an AI tool use, and what should it never assume?

Discovery should not feel like an attempt to extract or replace expertise. The purpose is to reduce avoidable dependency, preserve institutional knowledge, and involve experts in designing better work.

Appendix D: Readiness Remediation Plan

Readiness gap	Business impact	Minimum corrective action	Owner	Due date	Evidence of readiness

Keep the remediation plan limited to issues that materially affect the selected use case. Broad cleanup work should be justified separately through business value, risk, or shared dependency across several use cases.