SAMPLE CHAPTER · FREE DOWNLOAD

THE FORWARD DEPLOYED INVENTORY SYSTEM

You Cannot Optimize What You Cannot See

Chapter 9 · Part III · Level 1: Foundation

> *Do not spend millions optimizing*
> *what you have not first learned to trust.*
>
> FROM THE FDIS MATURITY MODEL

Glenn Caudill

COMPLIMENTARY PREVIEW · NOT FOR RESALE

PART III · LEVEL 1: FOUNDATION

Chapter 9

You Cannot Optimize What You Cannot See

> *“Many intelligence reports in war are contradictory; even more are false, and most are uncertain.”*
>
> — Carl von Clausewitz

> *“The farther information travels from the point where reality occurs, the more certainty it seems to acquire.”*
>
> — Glenn Caudill

Trust in the Information

In Chapter 4, we saw that trust is not merely a virtue between individuals; it is a foundational principle of efficient organizations, especially when work must move across departments, functions, and organizational boundaries. Processes that cross these boundaries depend on people being willing to act on one another’s commitments, judgments, and decisions without constantly stopping to verify every step.

The same is true of information. Trust between people can carry an organization only so far if those same people are working from data they do not believe. In fact, uneven confidence in the data can begin to erode the personal trust that took time to build. One group may believe the numbers are authoritative, while another knows from daily experience that they are not. When that happens, disagreements that appear to be about judgment or competence are often really disagreements about which version of reality can be trusted.

That distinction matters because people do not make decisions from relationships alone; they make them from relationships combined with information. Supply chain may trust field service, and field service may trust supply chain, but if each group is relying on a different understanding of what inventory actually exists, where it is located, or whether it has been consumed, the process will still break down. Effective organizations therefore need both forms of trust: trust in the people involved and trust in the information they use together.

Trust in information also has to be maintained. Forward-deployed inventory is constantly moving—received, consumed, replenished, transferred, returned, and sometimes misplaced—so even accurate data can drift away from reality surprisingly quickly. For that reason, data accuracy cannot be treated as a periodic cleanup exercise or something established once during a system implementation. It has to be sustained by an operating process that continually compares what the system says with what is actually happening in the field, identifies discrepancies, and corrects the behaviors that produced them.

The People Closest to the Work Know

The first signs that information cannot be trusted usually appear at the point where the work is actually being done. Field Service Personnel discover that a part shown in inventory is not on the shelf. Local managers learn which locations are routinely inaccurate, which transactions are entered late, and which movements of inventory never make it into the system at all. Over time, the people closest to the operation develop a remarkably detailed understanding of where the official record departs from reality.

They also learn how to work around it. A technician calls another technician before driving across town for a part. A manager keeps a spreadsheet alongside the corporate system because experience has taught her that the two will not always agree. Someone checks a shelf manually before trusting an inventory balance, while another person holds a little extra stock because replenishment has failed often enough to make caution seem sensible. These practices may look inefficient from a distance, but from the perspective of the people doing the work, they are rational responses to information that has repeatedly proved unreliable.

The danger is that these workarounds can become so normal that the organization stops recognizing them as evidence of a deeper problem. The operation continues because experienced people compensate for the weaknesses in the information system. They remember exceptions, make phone calls, maintain unofficial records, and use judgment to bridge the gap between what the system says and what they know to be true. In effect, they are keeping the process together through personal knowledge and constant intervention.

For a time, this can create the appearance that the system is working better than it actually is. Service continues, customers are supported, and problems are resolved, but much of that performance is being produced by people compensating for unreliable information rather than by a process that can be trusted on its own. The organization is not truly in control; the people closest to the work are simply holding it together.

The Farther Data Travels, the More Certain It Appears

As information moves upward through an organization, something curious often happens: the uncertainty surrounding the data begins to disappear. At the point where the work is being performed, people may know that a number is approximate, that a transaction is missing, or that a location has not been reconciled recently. By the time that same information reaches a dashboard or an executive report, those qualifications are often gone.

The field may say, “The system shows three, but I would verify that before sending someone.” A regional manager may understand that one of the three is probably no longer there. A report prepared several levels higher simply shows an available quantity of three. What began as a number surrounded by experience, context, and doubt has become a fact on a screen.

This creates an important asymmetry inside the organization. The people closest to the work may have less confidence in the data precisely because they have more knowledge of how it was created. They see the missed transactions, the workarounds, the delayed updates, and the exceptions. People farther from the operation see a cleaner representation of the same reality and can therefore become more

confident in it than the people who actually produced it.

That leads to a principle worth remembering:

"The farther information travels from the point where reality occurs, the more certainty it seems to acquire." — Glenn Caudill

The problem is not that senior leaders are careless or that people in the field possess some special monopoly on truth. The problem is that uncertainty is rarely carried through an organization as carefully as the data itself. Numbers are easy to aggregate. Context is harder. As a result, management can become increasingly confident in information whose limitations are best understood by the people farthest from the decision-making table.

Data Is Not Automatically Information

We often use the words data and information as though they mean the same thing. They do not. Data is a record of what a system captured. Information is data that has enough context, accuracy, and reliability to support a decision. The distinction matters because organizations can accumulate enormous amounts of data and still have a poor understanding of what is actually happening.

For forward-deployed inventory, the quality of that understanding depends on questions that are often left unasked. How current is the inventory record? How complete is it? How was it created? Which movements may never have been captured? How often do system balances disagree with physical reality? Most importantly, what is the error rate of the data being used to make decisions?

That last question is critical because the absence of an error rate does not mean the absence of error. It simply means the organization does not know how wrong the data may be. A dashboard can show a precise inventory balance without showing the uncertainty behind it. A report can be internally consistent and still be materially different from what exists in the field. Without some measure of accuracy, the organization is not really managing information; it is managing an assumption.

There is a certain comfort in that assumption, at least for a while. The reports arrive on time, the dashboards look complete, and decisions can be made with confidence. But confidence is not the same as knowledge. Eventually the results begin to fall short of expectations, and the organization is left trying to explain why decisions that looked sound on paper did not produce the expected outcome. The problem may have begun much earlier, when data was accepted as information without first understanding how closely it reflected reality.

Optimization on Untrusted Data Is Wasteful

Once data has been accepted as information, organizations naturally begin trying to improve the decisions made from it. This is where increasingly sophisticated tools enter the picture. Optimization engines, simulation models, artificial intelligence, predictive analytics, and advanced planning systems all promise to help companies extract more value from their operations. Used on a sound foundation, they can do exactly that.

The problem arises when the sophistication of the analysis advances faster than the quality of the underlying information. An organization may spend millions of dollars building a model expected to reduce inventory by 10 percent, improve utilization by 8 percent, or lower premium freight by 15 percent. Those targets may be perfectly reasonable. But if the field inventory data feeding the model contains an error rate of 25 or 35 percent, the organization is pursuing a relatively small improvement inside a much larger field of uncertainty.

This is more than the familiar warning of “garbage in, garbage out.” The greater danger is that sophisticated tools can make uncertain data appear more trustworthy than it really is. A rough spreadsheet invites questions. A mathematically elegant model tends to command confidence. The output may be presented to several decimal places, supported by charts, scenarios, and statistical analysis, but none of that precision changes the quality of the reality represented by the inputs.

The distinction between precision and accuracy is critical. Precision tells us how finely a result can be calculated. Accuracy tells us how closely that result corresponds to reality. An organization can become extraordinarily precise while remaining materially wrong. In that situation, greater sophistication does not reduce uncertainty; it disguises it.

This helps explain why ambitious technology programs sometimes produce disappointing business results. The problem may not be that the promised improvement was impossible. The organization may simply have attempted to optimize a version of the operation that did not exist.

The Tool May Not Be the Problem

When those expected results fail to materialize, attention usually turns to the tool. Leaders question the software, the consultants, the algorithm, or the assumptions in the model. A new vendor may be considered. Another implementation may be proposed. Additional analytics may be added in an attempt to explain why the original analytics did not work.

Sometimes the tool is, in fact, the problem. No technology should be immune from scrutiny. But before concluding that the model has failed, the organization should look upstream at the processes that produced the data the model was given.

A planning system cannot know that an outbound transaction was never recorded. An optimization engine cannot know that a part shown at one location was informally transferred to another. A predictive model cannot infer that a replenishment was delivered but never properly received into inventory unless some other reliable signal tells it so. The tool sees the digital representation of the operation. The people in the field live inside the physical operation.

If those two realities have drifted apart, the analytical system is being asked to solve the wrong problem with considerable sophistication.

That distinction changes the corrective action. If the tool is at fault, improve or replace the tool. If the underlying processes are continuously producing inaccurate data, replacing the tool merely places a new layer of technology on the same unstable foundation. The organization may spend more money, obtain a more modern interface, and generate more sophisticated recommendations while preserving the

conditions that caused the original disappointment.

Before asking why the tool failed, therefore, it is worth asking a more fundamental question: What did we give the tool to work with?

Episodic Cleanup Creates an Illusion of Improvement

Most organizations eventually become aware of serious gaps between their records and physical reality. The traditional response is an event. A full physical inventory is scheduled. A major cycle count is performed. A special reconciliation team is assembled. Consultants are brought in to clean the data before a system conversion. For several days or weeks, the organization concentrates enormous effort on restoring accuracy.

These efforts can be valuable. At the end of them, the inventory records may be substantially better than they were at the beginning. Management sees the accuracy number rise, discrepancies are resolved, and there is a genuine sense that the problem has been addressed. But the more important question is what happens after the event ends.

If the same processes remain in place, the same processes begin producing the same errors. Parts continue to move without being recorded. Transactions are delayed or missed. Employees continue to rely on workarounds that make sense locally but weaken the integrity of the system. Replenishments are received incorrectly. Inventory is transferred informally. The organization has corrected the accumulated errors without changing the mechanisms that created them.

The result is a predictable cycle. Accuracy is restored, slowly deteriorates, eventually becomes painful enough to attract management attention, and is restored again. Each cleanup creates a temporary high point, followed by another period of drift.

For the sophisticated tools depending on this information, the problem is even more subtle. Immediately after a major reconciliation, their inputs may be relatively reliable. Months later, the same models are operating on a much less accurate picture of the field. The model has not changed, but the quality of its inputs has. When the results become inconsistent, the organization may once again question the model without recognizing that the foundation underneath it has been steadily deteriorating.

Replace Episodic Cleanup with a Continuous Corrective Process

The alternative is to stop treating data accuracy as something that must occasionally be restored and begin treating it as something that must continuously be maintained. That requires more than better counting. It requires a process that addresses both the incorrect record and the reason the record became incorrect.

When the physical inventory does not agree with the system, correcting the balance is necessary. But it should be the beginning of the investigation, not the end of it. The organization also needs to ask why the two diverged in the first place.

Was an inbound transaction missed? Was a part removed without an outbound transaction? Was inventory moved between Field Service Personnel or locations outside the normal process? Did the technology make the correct behavior too difficult? Was training inadequate? Was ownership unclear? Did an emergency situation cause someone to bypass the normal procedure? Is there a recurring system defect that encourages people to develop their own workaround?

Each discrepancy contains information about the process that produced it. If the organization merely adjusts the inventory balance, it corrects the symptom while discarding much of the diagnostic value contained in the error.

A discrepancy should therefore be understood as more than a bad record. It is evidence that something in the operating process did not work as intended. Repeated discrepancies of the same kind reveal patterns, and those patterns tell management where the process needs to change.

The goal is not simply to make the inventory right today. The goal is to reduce the probability that it will be wrong tomorrow.

Know the Error Rate in Real Time

This leads to a different way of thinking about data quality. The objective is not to declare that the data is accurate and move on. Nor is it reasonable to demand perfect information at every moment. In a dynamic forward-deployed inventory network, parts are continually moving and some amount of error will always need to be detected and corrected.

What matters is knowing how much uncertainty exists and whether it is increasing or decreasing.

An organization that understands its current error rate can make much better decisions than one that simply assumes its records are correct. A leader considering an optimization recommendation should know whether the underlying inventory data is 99 percent accurate, 95 percent accurate, or perhaps only 70 percent accurate. Those conditions call for very different levels of confidence in the resulting recommendation.

The important change is that accuracy becomes an operating metric rather than the result of an occasional audit. Instead of learning once or twice a year that the system has drifted substantially from physical reality, the organization develops mechanisms that continuously reveal discrepancies and provide an ongoing picture of the quality of its information.

This does not eliminate uncertainty. It makes uncertainty visible.

And visible uncertainty is manageable. Hidden uncertainty is not.

Once leaders understand the error present in the information, they can decide how much confidence to place in an analytical result, where additional verification is required, and which processes are creating the greatest deterioration in accuracy. The organization can then improve the quality of the information at the same time that it improves the decisions made from it.

Frustration Is Reasonable, but Leadership Is the Answer

Discovering that important business decisions have been made from unreliable information can be infuriating. It is especially frustrating when significant investments have already been made in systems, analytics, consulting, and optimization efforts that were expected to produce measurable results.

That frustration is understandable.

People may ask why the problem was not identified earlier, why those closest to the operation did not raise the alarm more forcefully, or why expensive technology failed to expose the weakness in the data. There will be a natural temptation to identify the person, department, vendor, or system responsible.

But blame rarely improves the foundation.

In many cases, no single person created the problem. The inaccuracies accumulated through hundreds or thousands of individually understandable actions inside a process that did not reliably capture or correct them. A technician solved an urgent customer problem by moving a part. A manager created a spreadsheet because the system was difficult to use. A planner made an assumption because a transaction had not yet arrived. Each decision may have been rational in isolation, while the combined effect gradually weakened the integrity of the information.

Correcting that condition requires leadership rather than anger. Leadership establishes that accurate information is important, creates a process for maintaining it, assigns responsibility for discrepancies, removes obstacles that encourage workarounds, and insists that the organization learn from recurring errors.

Most important, leadership creates the patience to repair the foundation before demanding more sophistication from the layers built on top of it.

Why Foundation Comes First

This is why Foundation is the first level of the Forward Deployed Inventory System Maturity Model.

It is tempting to look past foundational capabilities because the technologies at the higher levels are more exciting. Optimization, simulation, artificial intelligence, demand sensing, reliability engineering, and predictive placement all promise impressive improvements. They are also easier to present to executives as transformational initiatives.

Tracking, reconciliation, discrepancy management, and transaction discipline can appear much less glamorous.

Yet everything above them depends on their success.

The purpose of the Foundation level is not simply to install scanning technology or conduct more cycle counts. Its purpose is to create a reliable connection between the physical world and the information used to manage it. Inventory movement must be captured. Expected and actual events must be reconciled. Discrepancies must become visible. Their causes must be understood. The resulting level of accuracy must be measured and maintained rather than assumed.

Only then does higher-order optimization have a stable surface on which to operate.

An organization should certainly aspire to become more sophisticated. It should use better models, artificial intelligence, simulation, and predictive technologies where they create value. The mistake is believing that sophistication can substitute for foundation.

It cannot.

The proper sequence is not to avoid advanced technology. It is to earn the right to use it effectively by first establishing information that deserves to be trusted.

That is the logic of the FDIS Maturity Model. Foundation comes first not because it is elementary, but because every improvement that follows depends upon it.

If the process keeps recreating the error, cleaning the data only resets the clock.

Do not spend millions optimizing what you have not first learned to trust.

Looking Ahead: Chapter 10 — Closed-Loop Starts With Accountability

If trustworthy information must be continuously maintained, the next question becomes unavoidable: How do we do it?

The answer cannot be another annual physical inventory or another emergency reconciliation project. An organization needs a process that connects the movement of inventory to the information that represents that movement, detects when the two diverge, and creates accountability for closing the gap.

That is the purpose of the closed loop.

In Chapter 10, we will move from the principle of trusted information to the operating discipline required to create it. We will follow inventory from expected arrival through receipt, storage, consumption, replenishment, reconciliation, and discrepancy resolution, and examine why every handoff matters.

Because in forward-deployed inventory, accountability is not simply knowing who touched the part. Accountability is knowing whether the physical event and the information describing that event remained connected from beginning to end.

ABOUT THIS SAMPLE

The Forward Deployed Inventory System

A Maturity Model for Seeing, Trusting, and Improving Field Inventory

This complimentary chapter is offered so leaders in field service, spare parts, aftermarket operations, and supply chain can test the argument before committing to the full book. Chapter 9 sits at the Foundation level of the FDIS Maturity Model: the point at which an organization either earns the right to optimize — or spends millions optimizing a version of the operation that does not exist.

The full book develops the complete maturity path: how to capture movement, close the loop between physical events and system records, assign accountability at every handoff, and only then apply planning, optimization, and predictive tools that can actually be trusted.

If this chapter named a problem you already feel in the field — unofficial spreadsheets, phantom parts, dashboards that look certain, models that disappoint — the rest of the book is written for the work of fixing it.

© Glenn Caudill. Sample chapter provided for evaluation. All rights reserved.