




What Operations Leaders Need to Know About Data Integrity in Distributed Environments

Keeping distributed operations running on data
you can trust

SUCCESS
SUCCESS
SUCCESS



When the Data Your Teams Trust Stops Being True



Now two versions of reality are running side by side. The maintenance schedule the platform is following isn't quite the one head office sees. The inventory count doesn't match. And every decision made from that moment on is built on data that quietly fell out of step.

This isn't a worst-case scenario, but the default outcome of assuming your data is in sync without ever verifying that it is.

This guide is about closing the gap between assuming your data is consistent and knowing every site is working from the same operational truth. The challenges it describes apply to distributed operations of any kind, but the focus here is on organizations running distributed IFS environments, where replication is what holds that consistency together.

Know every site is running on the same operational truth.

See it in action

The Foundation of Distributed Operations:

Maintaining Data Consistency at Scale

A distributed operation only works when every location sees the same underlying picture. The moment that picture diverges, good decisions get made on bad data.

Underneath every site sits the same shared record:



Work orders



Inventory



Maintenance schedules



Financial records

Each one has to mean the same thing at the distributed site as it does at the hub.

That’s harder than it sounds, because of where these operations run. Offshore platforms, vessels, and field facilities sit at the far end of an unreliable connectivity link. And here’s what’s easy to underestimate: connectivity loss in these environments isn’t the exception. It’s routine. Weather, distance, satellite handoffs, equipment limits – the link will drop. The only real question is what happens to your data when it does.

When a site goes offline, even briefly, data stops flowing. When it comes back, the gap it left behind may be invisible. But the decisions made in that window were based on incomplete information, and those small divergences don’t always heal when the connection is restored. They persist, quietly shaping the next decision, and the next.

The stakes are measurable.

\$12.9M

The average cost of poor data quality to an organization per year.¹

In asset-intensive operations, a single misaligned work order or inventory figure can ripple into downtime, a safety exposure, or a missed maintenance window. The true cost is usually felt long before it’s ever counted.

Know every site is running on the same operational truth.

See it in action

¹Gartner: Data Quality: Best Practices for Accurate Insights



The Ways Distributed Data Goes Silent

When a site goes offline, it lets data fall out of sync in three distinct ways.

1. Stoppages

Data stops moving between the site and the hub entirely. When connectivity returns, both sides have already carried on, working from data that diverged the second the link dropped. *Nothing about the restored connection reconciles them.*

2. Queued data that silently backs up

Data is queued to sync, but nothing is processing. The site looks connected. The dashboards look normal. Underneath, the backlog grows. The system looks healthy; the data is not.

3. Connectivity failures

Intermittent signal loss – common in offshore and distributed environments – interrupts the flow again and again without ever triggering a full outage. No single drop is big enough to notice. *Small, cumulative gaps add up to real divergence over time.*

The root cause is always the same: distributed sites in low-connectivity conditions are exposed the moment the link drops. When it returns, the gap isn't obvious – and the operation has already moved on.

The failure isn't the real risk. Not noticing is.

By the time the consequences surface, they're already embedded in the operation.

Know every site is running on the same operational truth.

See it in action

Why Most Teams Don't Know There's a Problem Until It's Too Late

Most teams already watch their data somehow. The problem is that the three most common methods were never built for silent, connectivity-driven gaps. Each has a structural blind spot.

Manual queue checks tell you what was true when you looked, not what's true now.

A site drops overnight and recovers by morning. Your 8 am check shows it connected. What it won't show is the gap that piled up at 10 pm while everyone was asleep. Periodic sampling gives you confidence without coverage.

Email alerts wait for a human to notice, interpret, and act.

They fire to a distribution list. Sometimes at 2 am. They get filtered, missed, or ignored. Where drops are brief and frequent, the threshold may never trip, even as divergence grows. And they don't scale: what works for five sites is noise across fifty. Your IT team can spot a problem this way. A notification alone doesn't let them own it.

In-house scripts get built once, then forgotten.

The script breaks. Replication starts failing. Nobody connects the two, because the very thing meant to catch the problem is the thing that stopped working. This is the flaw at the heart of self-managed monitoring: the monitor itself has no monitor.

The gap is structural, not a failure of effort.

These approaches simply weren't built for the connectivity realities of distributed operations.

And the effort is enormous:

82% of teams spend a day or more each week fixing data quality issues.


66% still rely on manual review to catch them.²

That's skilled people spending real time chasing problems after the fact with exactly the methods least able to catch them early.

² McKinsey: Master data management: The key to getting more from your data

What Effective Data Oversight Looks Like

If the common methods share one blind spot, effective oversight has to be built differently. Four principles separate oversight that works from oversight that only looks like it does.



Continuous, not periodic.

Always-on, including during and after connectivity events. No scheduled window for a failure to hide in.



Detection-based, not threshold-based

It catches data divergence the moment it happens, including the subtle gaps that build when connectivity flickers rather than fully fails.

Know every site is running on the same operational truth.

See it in action



Owned response, not just flagged.

A notification tells you something's wrong; it doesn't make sure something's being done. The right people should already be on it before the gap is hours old.



Scaling, not reconfiguring.

Every new site, vessel, or facility is covered automatically, with no fresh setup project each time the footprint grows.

STEP UP
STEP UP
STEP UP



How IFS Success Keeps Distributed Operations Running on Trusted Data

Look again at those four principles. None of them describes software. They describe a service – and the team accountable for it.

Replication Monitoring

Always-on, is a service IFS operates, not a tool your team runs. It continuously monitors replication health across your distributed IFS environments, detecting the stoppages, stalled queues, and connectivity failures that occur when sites go offline. When it finds one, it opens an IFS-owned case – and IFS is already working the problem before your team knows there is one.

Detection is only half of it.

A case is only as good as the team that picks it up. Replication Monitoring is delivered as part of IFS Success, so when an issue surfaces, it reaches people who already know your environment and what normal looks like across your sites. Detection without that context is just a faster alert.

[Learn more about IFS Success](#)

What IFS Success brings

A long-term partnership across your whole IFS lifecycle, built to help you accelerate time to value, achieve operational excellence, and get more from your IFS investment across its lifetime.

Backed by over 2,500 IFS experts with deep experience in your industry, it works across four areas:

- **Business initiatives and the path to innovation** – turning your strategic goals into a joint plan with measurable outcomes
- **Solution best practices** – product expertise applied to keep your system aligned to the business, efficient, and performing
- **Operational best practices** – optimizing processes and guiding release management to keep your environment evergreen
- **Business adoption** – helping your teams get full value from what you've already invested in

Is Your Operation Running on Data You Can Trust?

Every operation is different. The scale of your distributed environment, the connectivity each site faces, the way you watch your data today – all of it varies. So, the right next step isn't a checklist. It's a conversation.

We can help you assess your current environment and find the right fit, so replication stops being something you hope is working and becomes something you know is, with every site and every team acting on data they can trust.



Get in touch to start the conversation.

About IFS

IFS is the world’s leading provider of Industrial AI and enterprise software for hardcore businesses that make, service, and power our planet. Our technology enables businesses which manufacture goods, maintain complex assets, and manage service-focused operations to unlock the transformative power of Industrial AI™ to enhance productivity, efficiency, and sustainability.

IFS Cloud is a fully composable AI-powered platform, designed for ultimate flexibility and adaptability to our customers’ specific requirements and business evolution. It spans the needs of Enterprise Resource Planning (ERP), Enterprise Asset Management (EAM), Supply Chain Management (SCM), and Field Service Management (FSM). IFS technology leverages AI, machine learning, real-time data and analytics to empower our customers to make informed strategic decisions and excel at their Moment of Service™.

IFS was founded in 1983 by five university friends who pitched a tent outside our first customer’s site to ensure they would be available 24/7 and the needs of the customer would come first. Since then, IFS has grown into a global leader with over 7,000 employees in 80 countries. Driven by those foundational values of agility, customer-centricity, and trust, IFS is recognized worldwide for delivering value and supporting strategic transformations. We are the most recommended supplier in our sector. Visit ifs.com to learn why.

