The Metrics That Actually Matter

- % *critical apps with SBOM*
- # *AI models with training-data lineage*
- *time-to-verify authenticity in "is this real?" events*

Gartner flags provenance as a response to code tampering, abandoned open-source risk, and deepfake-driven disinformation, plus regulatory pressure (e.g., watermarking/provenance tracking) with the goal of reduced fraud/impersonation, better incident forensics and, cleaner audits.

ACTION BOX: Build a Trust Infrastructure

A 90-Day "Stop Guessing" Plan for Hybrid IT

1 Days 0–30: start where damage would hurt most

Focus on:

- Crown-jewel apps (customer portals, payments, identity)
- Externally exposed brand channels (exec comms, status pages)

Do this:

- Generate SBOMs for crown jewels (and refresh per release)
- Define minimum AI lineage metadata (if you deploy models)
- Map "must-trust artifacts" (images, binaries, model versions, public statements)

2 Days 31–60: Add attestation and verification gates

Do this:

- Sign builds and container images
- Attach provenance attestations to releases
- Enforce verification in at least one critical runtime (e.g., Kubernetes admission controls)

With AI-based solutions, track:

- Establish model registry + dataset versioning workflows
- Require "model card + data card" basics before production promotion

3 Days 61–90: Operationalize it (because paper controls don't work)

Do this:

- Write a deepfake/disinformation response playbook with Legal + Comms + Security
- Run "Is this real?" tabletop drills
- Wire SBOM + exposure management into change/release workflows

The Day 90 finish line:
You're not "done." But you've stopped relying on trust-by-default for your most important systems and communications.

DIGITAL ERA