HOUSEBLEND

HOUSEBLEND / RESEARCH & PRACTICAL GUIDANCE

NetSuite Safety Stock Optimization: Item-Location Review

> Make NetSuite work better for your finance and operations teams with Houseblend. We help design, implement, integrate and improve ERP systems, with practical support for the people who use them every day.

Published by Houseblend · 2026-11-25 · netsuite safety stock optimization · netsuite safety stock · netsuite reorder point · netsuite inventory optimization · safety stock by location · inventory optimization results workbook · supply analysis interval · service level planning · netsuite 2026.2

Original article: https://www.houseblend.io/articles/netsuite-safety-stock-diagnostics-item-location



In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes: What the Workbook and the September 2026.2 Additions Store
4. Prerequisite Checks: Features, Permissions, and Item Eligibility
5. Demand History, Lead-Time History, and the Service Level Decision
6. Evaluating a Recommendation: Field-by-Field Interpretation
7. Triage: Blanks, Zeros, Stale Runs, and Missing Data
8. Implementation Considerations and Process Changes
9. Data Analysis and Evidence
10. Implications and Future Directions
11. Frequently Asked Questions (FAQs)
12. Conclusion

Executive Summary

NetSuite's Inventory Optimization feature writes a safety stock, a reorder point, and a preferred stock level for each eligible item-location. The September release of NetSuite 2026.2 added a **Supply Analysis Interval** preference that sets the lead-time history window separately from demand history, plus projected measures and two new workbook views, **Inventory Optimization KPIs** and **Inventory Optimization Diagnostics** (docs.oracle.com). This report treats those additions as an approval toolkit rather than a formula tutorial. The decision it supports: for each recommendation, approve it, override it, or hold it for investigation.

The workbook holds stored results written by inventory optimization processing. Changed inputs appear only after a scheduled or manual calculation processes the item-location, and the standard workbook has no current item-location quantity fields (docs.oracle.com) (netsuitechangelog.com). A recommendation is therefore a dated calculation output, not a live position and not a forecast accuracy guarantee. NetSuite computes safety stock as the lead-time-period demand standard deviation times a z-score from the service level, rounds up, and assumes normally distributed lead-time demand (docs.oracle.com).

The new diagnostics need careful reading. The target service level, the previous projected service level, and the optimized projected service level are three different numbers, and the projected fill rate is a fourth that "differs from Service Level" (docs.oracle.com). Blank and zero also differ: numeric zero is a valid result, while a blank means required source data, an input, or a result is unavailable (docs.oracle.com). A missing service level makes processing fail for that item-location, and fewer than three qualifying lead-time values leave safety stock and reorder point uncalculated (docs.oracle.com).

Census Bureau data for September 2026, released in November 2026, put the US total business inventories-to-sales ratio at 1.33, manufacturers' inventories-to-shipments at 1.46, and the merchant wholesaler ratio at 1.21 (www.census.gov). These aggregates describe the macro setting, not any item, and no public benchmark can validate an individual reorder point. The ISM Supplier Deliveries Index read 59.0 in September 2026, its tenth consecutive month of slower deliveries (www.ismworld.org). The report therefore proposes a seven-gate

approval workflow, a sign-off register, and a user acceptance testing (UAT) checklist, all anchored on the account's own demand and lead-time history, with measured post-change monitoring before any claim of better service or lower working capital.

Introduction and Background

Safety stock and reorder points are often discussed as if they were one setting. They are not. The reorder point is the quantity at which a replenishment order is triggered, and in textbook form it is average demand during the lead time plus safety stock (en.wikipedia.org). Safety stock is the buffer inside that reorder point, sized to cover variability in demand and lead time up to a chosen service level (learn.microsoft.com). Oracle's educational material adds that reorder points govern the timing of an order, not its size (www.netsuite.com). This answers the common query "reorder point vs safety stock": the reorder point is the trigger, safety stock is the part of the trigger that protects against variability, and the preferred stock level is the order-up-to target built on top of safety stock.

NetSuite Inventory Optimization, available starting with 2026.2 according to NetSuite Changelog's summary of Oracle's release notes, recalculates lead time, safety stock, reorder point, and preferred stock level for eligible item-locations, while Reorder Point items not included in the feature keep using standard calculations (netsuitechangelog.com). The segmentation layer assigns each included item (or item-location) an ABC category based on annual consumption value and a 123 category based on demand variability, and each of the nine resulting segments carries a default service level (netsuitechangelog.com).

The decision problem begins after the calculation. Planners receive a table of recommended values for hundreds or thousands of item-locations, and a controller or chief operating officer (COO) must decide whether those numbers are ready to apply. Three realities shape that decision: the workbook shows stored results that can lag the inputs, the calculation rests on a normal-distribution assumption that Oracle flags as approximate for intermittent, promotional, strongly seasonal, or skewed demand (arxiv.org), and the quality of the answer depends on the business's own demand and lead-time history, which no public benchmark can replace (knowledge.wharton.upenn.edu).

This report is written for COOs, supply planners, inventory finance managers, and NetSuite implementation leads at multi-location distributors and manufacturers. It uses Oracle's documentation for field definitions, supplements it with inventory theory from university and peer-reviewed sources, and anchors the macro context in Census Bureau, New York Fed, and Institute for Supply Management (ISM) data. It does not claim that the algorithm improves service or reduces working capital, since those outcomes must be measured in each account. Organizations that want implementation help with NetSuite configuration, SuiteAnalytics workbooks, or managed administration can engage a NetSuite consultancy; the publisher of this report, Houseblend, is one such provider, and nothing in the analysis depends on that.

Key Changes: What the Workbook and the September 2026.2 Additions Store

The Results Workbook Is a View of Stored Results

Oracle describes the **Inventory Optimization Results Workbook** as a set of table views over active stored results, one row per item-location combination. Its source dataset joins Item, Location, Item Location Configuration, and Inventory Optimization Segment records. The standard workbook shows result records whose Inactive field is false, and inventory optimization processing writes the stored values (netsuitechangelog.com).

Two omissions shape interpretation. The standard workbook has no current item-location quantity fields, pivot table, chart, or dashboard page (netsuitechangelog.com), so on-hand stock and open purchase orders must come from live transaction searches. And because it holds stored output, users who customize it should use Save As and not overwrite the standard workbook. One third-party summary lists dataset fields hidden in the standard view differently from Oracle's own table, so teams should confirm the Data Grid in their own account (netsuitechangelog.com).

The Supply Analysis Interval Separates Lead-Time History From Demand History

Before the September release, one history window drove both demand and lead-time statistics. Oracle's release note says users can now "set the supply history period for lead-time calculations independently of demand history," entered in months on the Optimization Preferences tab of Inventory Management Preferences. Demand calculations continue to use the Order Analysis Interval. The preference accepts 1 through 60 months with a default of 6, applies account-wide, and cannot be overridden for individual items or locations (docs.oracle.com).

Two behaviors matter for approvals. The interval governs only lead-time inputs. And a change affects only future results: "Changing this preference affects subsequent calculations," and previously created results keep their stored supply analysis dates (docs.oracle.com). A planner who changes the interval and then reads yesterday's results is reading numbers built on the old window, so the run must be repeated and compared with the prior one.

Projected Measures and the KPIs View

The September update added stored projections for shortages, inventory levels and values, quantity fill rate, turnover, and service levels, intended to help users "assess inventory value and expected shortages under the planning parameters used in the calculation". The **Inventory Optimization KPIs** view surfaces the financial subset. Oracle defines expected units short as a per-replenishment-cycle estimate and stresses that projected results "don't describe actual transaction results". The projection describes the planning parameters, not next quarter's shortages.

The Diagnostics View Compares Previous and Optimized Safety Stock

The **Inventory Optimization Diagnostics** view compares the projected service level of the previous safety stock with that of the newly optimized value, and both projections use the current calculation's inputs. The previous safety stock is therefore re-scored against today's variability: the comparison asks how the old buffer would perform under today's statistics. Oracle adds that these diagnostics "don't represent Expected Quantity Fill Rate or achieved customer service".

Prerequisite Checks: Features, Permissions, and Item Eligibility

Feature and Permission Checklist

Oracle instructs users to enable Advanced Inventory Management and Inventory Optimization on the Enable Features page, and to enable Advanced Item Location Configuration to segment inventory by item-location (netsuitechangelog.com). Enabling the feature alone does not change planning values, because preferences, item inclusion, and calculations must follow (netsuitechangelog.com).

Permissions are the next gate. A role needs SuiteAnalytics Workbook at the Edit level and Inventory Optimization at the View level or higher, and users see only the rows available to their roles (netsuitechangelog.com). Running Calculate Inventory Levels also needs the Items permission (netsuitechangelog.com). In UAT, a reviewer who sees fewer rows than the planner is signaling role scope, not missing results.

Which Items Use Reorder Point Replenishment

Only some items are in scope. The Include in Inventory Optimization setting is editable only when the item uses the Reorder Point replenishment method (docs.oracle.com). Planning values are calculated at the item-location level for active, included Reorder Point items at locations where Make Inventory Available is checked (netsuitechangelog.com), and Reorder Point items outside the feature keep the standard calculation (netsuitechangelog.com).

Before any review meeting, list which active items use Reorder Point replenishment, which of those are included, and which locations make inventory available. An item-location absent from the workbook may be excluded, inactive, unprocessed, or failed, and each cause has a different owner.

Item-Location Segmentation Versus Item-Level Segmentation

The **Calculate Inventory Segmentation per Location** preference is cleared by default, in which case NetSuite calculates one segment per item instead of one per item-location. Planning calculations still run at the item-location level. This answers the query "NetSuite safety stock by location": the safety stock figure is always location-specific, but the service level feeding it may be shared across locations under item-level segmentation.

Demand History, Lead-Time History, and the Service Level Decision

Defining Demand History

For nonseasonal items, demand uses the Order Analysis Interval, and the calculation uses average daily demand and daily demand standard deviation drawn from demand history. Days with no demand are included as zeros, which lowers the mean and raises relative variability for slow items. Pending Approval, Cancelled, and drop-ship sales orders are excluded from demand under the orders setting. Each result stores whether its demand source was ORDERS or ACTUALS, so a reviewer comparing two runs should confirm both used the same source.

Three data-quality tests apply before accepting any recommendation. Does the demand window contain one-time events, such as a large project order, that inflate the standard deviation? Does it reflect the current customer base? Are the underlying inventory records trustworthy? A peer-reviewed study of nearly 370,000 inventory records at one retailer found 65% inaccurate, which is one retailer's experience and not a distributor benchmark, but it shows why record accuracy belongs in the review (ideas.repec.org).

Defining Lead-Time History

Lead-time statistics come from inbound supply history over the Supply Analysis Interval: purchase orders in Pending Billing, Billed, or Closed status, received transfer orders, and built or closed work orders. NetSuite requires at least three qualifying lead-time values. That floor is a computational minimum, not a statistical comfort level. Academic guidance warns that treating estimated parameters as true values leads to insufficient safety stock, stock-outs, and low service (ideas.repec.org), and a tutorial from North Carolina State University illustrates the calculation from only 10 weeks of demand and 8 prior orders, which is acceptable for teaching and risky in production (scm.ncsu.edu).

Lead-time variability deserves special attention because it enters the formula through a squared term. A longer but perfectly reliable lead time can need less safety stock than a shorter one that varies widely (user.engineering.uiowa.edu). Reviewers should look at lead-time standard deviation as closely as the mean, and should ask whether receipts are recorded when goods arrive, since batched or late receipt entry makes recorded lead times differ from physical ones.

Service Level: Segment Default, Override, and Target

Each segment has a default service level, which can be overridden on a segment, item, or item-location (netsuitechangelog.com). Overrides must be from 50% through 99.99% (docs.oracle.com). When segmentation runs it updates the Service Level field unless Override Service Level is checked, in which case NetSuite keeps the override. The dataset's Override formula field shows Manual Override or Segment Default, although Override Service Level itself is hidden in the standard views (netsuitechangelog.com). Planners should therefore read the Override column and not assume a service level reflects segment logic.

Higher service levels generally produce higher safety stock, and the relationship is nonlinear near 100% (ocw.mit.edu). The Association for Supply Chain Management (ASCM) states that "Statistically, a 100% service level is impossible" and that typical goals fall in a 90%-98% range (www.ascm.org). The National Institute of Standards and Technology (NIST) publishes the normal critical values used to convert a service

level to a z-score (itl.nist.gov), and NetSuite's educational material gives approximately 1.645 for 95% and 2.326 for 99% (www.netsuite.com). Oracle's calculation page uses a rounded z-table, so a hand check should allow for rounding.

Evaluating a Recommendation: Field-by-Field Interpretation

The Four Service Numbers

Four quantities are easily confused. The **target service level** is the segment default or override. The **previous projected service level** is the probability implied by the previous safety stock under the current run's variability, calculated as the standard normal cumulative distribution of previous safety stock divided by lead-time demand standard deviation. The **optimized projected service level** applies the same function to the new safety stock. The **expected quantity fill rate** is a different concept: Oracle says it estimates the percentage of units filled during a replenishment cycle and "differs from Service Level".

This mirrors textbook theory. Cycle service level is the share of replenishment cycles without a stock-out, fill rate is the fraction of demand served from stock, and two policies can share a cycle service level yet have quite different fill rates (user.engineering.uiowa.edu) (www.ncbi.nlm.nih.gov). A peer-reviewed simulation also found that a common fill-rate formula can overestimate the true fill rate, so a projected fill rate is a screening metric and not a promise (www.ncbi.nlm.nih.gov).

Field Table

Table 1 below lists the fields a reviewer needs, the view in which each appears, and the reading discipline each requires. Definitions follow Oracle's workbook and calculation pages.

Field	View	What it holds	How to read it
Result Created / Result Updated	Results, KPIs (created)	Timestamps of the stored result	Compare with the last input change; a stale run means the displayed figure predates the change
Service Level, Override	Results	Target service level and whether it is Manual Override or Segment Default	Confirm the override has an owner and a rationale; segment defaults come from the segment table
Demand Analysis Period Start/End	Results	Window used for demand mean and standard deviation	Check it spans a representative period and excludes one-off events
Supply Analysis Period Start/End	Results	Window used for lead-time mean and standard deviation	After a Supply Analysis Interval change, old rows keep old dates
Lead Time Mean / Standard Deviation	Results	Statistics from inbound supply history	Needs at least three qualifying values; check variability, not only the mean
Daily Demand Mean / Standard Deviation	Results	Statistics from demand history, zero days included	Zero days lower the mean and raise relative variability

Field	View	What it holds	How to read it
Lead Time Demand Mean / Standard Deviation	Results, Diagnostics (SD)	Demand over the lead time	The standard deviation is the denominator of both projected service levels
Safety Stock, Reorder Point, Preferred Stock Level	Results	Rounded-up recommendations	Reorder point equals safety stock plus lead-time demand mean
Previous Safety Stock	Diagnostics	Copied from the latest prior result for the item-location	Prior result may be inactive; blank if no prior result exists
Previous / Optimized Projected Service Level	Diagnostics	Service level implied by each safety stock under current variability	Mathematical projection, not achieved service
Distance From Target Before / After	Diagnostics	Absolute gap between target and each projection	Smaller after than before indicates alignment with the target, not value
Expected Units Short / Value	Results, KPIs	Shortage per replenishment cycle and its cost value	Uses the final rounded reorder point
Expected Average Inventory Value, Expected Inventory Turnover	Results, KPIs	Projected inventory value and turnover	Blank if Average Cost is missing
Expected Quantity Fill Rate	Results, KPIs	Estimated percentage of units filled per cycle	Differs from service level

Read dates and override status first, then variability, then projections. On a first run with no prior result, the previous values, the delta, and Distance From Target Before are blank.

A Fictional Item-Location Before and After (Hypothetical Example)

The following example is entirely fictional and uses no benchmark. It applies the formulas Oracle documents for safety stock, reorder point, projected service level, and the KPI measures. Assume item SKU-X at location DC-East with average daily demand of 40 units, daily demand standard deviation of 18 units, average lead time of 12 days, preferred stock level of 30 days of demand, and unit cost of $12. The target service level is 95%, using the z-value of 1.65 (scm.ncsu.edu), and the previous safety stock was 150 units.

Lead-time demand mean is 40 times 12, or 480 units. Lead-time demand variance is 12 times 18 squared plus 40 squared times lead-time standard deviation squared. With a lead-time standard deviation of 3 days, that is 3,888 plus 14,400, or 18,288, a standard deviation of about 135.2 units. Safety stock is 1.65 times 135.2, or 223.1, rounded up to 224, and the reorder point is 224 plus 480, or 704. *Table 2* compares this run (A) with a second run (B) in which lead-time standard deviation rises from 3 to 5 days, as could happen if a longer Supply Analysis Interval pulled in a few late receipts.

Measure (Hypothetical Example)	Run A: lead-time SD 3 days	Run B: lead-time SD 5 days
Lead-time demand SD	135.2 units	209.5 units
Optimized safety stock	224 units	346 units

Measure (Hypothetical Example)	Run A: lead-time SD 3 days	Run B: lead-time SD 5 days
Reorder point	704 units	826 units
Previous safety stock	150 units	224 units (Run A result)
Previous projected service level	about 86.6%	about 85.8%
Optimized projected service level	about 95.1%	about 95.1%
Projected delta	about +8.5 points	about +9.3 points
Expected units short per cycle	about 2.7 units	about 4.3 units
Expected units short value	about $33	about $52
Expected average inventory value	$7,872	$8,238
Expected quantity fill rate	about 99.6%	about 99.4%

The same demand history produces a 122-unit higher reorder point solely because lead-time variability rose, and the earlier safety stock of 224 would now project only about 85.8%. Moving safety stock from 224 to 346 adds 122 units, or about $1,464 at $12 per unit, to restore the target projection. Approval depends on whether the new variability is real (a supplier reliability change, which matters because a longer but reliable lead time can need less safety stock than a shorter, highly variable one (user.engineering.uiowa.edu)) or an artifact (a pattern in receipt entry), and the workbook cannot answer that.

Two cautions apply. Oracle defines projected turnover as lead-time demand mean times unit cost over expected average inventory value, so it is a planning-cycle ratio and not an annual cost-of-goods-sold turnover as in accounting texts (openstax.org); compare it across runs and not against an annual target. And the normal-distribution assumption means the 95.1% figure is only as good as the shape of real lead-time demand at that location.

Triage: Blanks, Zeros, Stale Runs, and Missing Data

Blank Versus Zero

Oracle's rule is explicit: numeric zero is a valid result, and a blank indicates that required source data, an input, or a calculation result is unavailable. The calculation page's condition table yields this triage.

- **Zero safety stock, reorder point, and preferred stock level:** average daily demand is zero, and a warning is logged. This is a finding about demand history, not a decision to stop stocking; check whether the window missed a seasonal period.
- **Blank safety stock and reorder point:** fewer than three qualifying lead-time values exist. The owner is whoever maintains purchase receipt data.
- **Blank unit cost, inventory value, shortage value, and turnover:** Average Cost is missing. The owner is inventory accounting.
- **Blank previous values and delta:** expected on a first run for the item-location, not an error.

- **Failed processing:** a missing service level makes processing fail for the item-location, so segmentation should be current before the levels calculation unless service levels are overridden.

Stale Runs and Inactive Results

A row can be wrong because it is old. Changed inputs appear only after a scheduled or manual calculation processes the item-location, and the levels calculation schedule is separate from the segmentation schedule. Only one Inventory Optimization calculation can run at a time, so a manual run may be blocked. The standard workbook also hides inactive result records, so a reviewer who sees no row should not infer that no recommendation exists. Segment records retain only the latest completed run's results, so segment-level history must be captured by the reviewer.

Missing Current On-Hand

Because the workbook has no current quantity fields, a reorder point can be approved without anyone looking at the position it will be compared with. The register below therefore asks for an on-hand and on-order snapshot from a live search on the sign-off date. A reorder point of 704 means little without knowing whether the location holds 1,500 units or 90.

Items the Normal Model Handles Poorly

Oracle's note on the normal-distribution assumption is the formal trigger for a second look at intermittent, promotional, strongly seasonal, or skewed items. The literature agrees: a 2025 preprint says real demand is often non-normal, intermittent, and highly skewed, and a peer-reviewed paper calls the normal lead-time demand assumption invalid for spare parts (arxiv.org) (orca.cardiff.ac.uk). Wikipedia's summary lists further assumptions, including independent and identically distributed demand and statistically constant supply, and notes that one mean and standard deviation can produce stock-outs in peak season and excess off-season (en.wikipedia.org). A 1997 study of fast-moving items with low variability found the normal approximation robust, so the assumption is a risk to test per item class and not a universal defect (ideas.repec.org). SAP documentation likewise treats Poisson models as suited to intermittent demand (help.sap.com).

Implementation Considerations and Process Changes

How to Calculate Safety Stock in NetSuite, End to End

The steps below answer the query "how to calculate safety stock in NetSuite" in the order a planner experiences them.

1. **Enable the features** and confirm roles, as described above.
2. **Include eligible Reorder Point items** in Inventory Optimization and set preferences, including the Order Analysis Interval and the new Supply Analysis Interval.
3. **Run segmentation** so that each included item or item-location has a segment and a service level.
4. **Set or review overrides** with an owner for each.
5. **Run Calculate Inventory Levels**, manually or on schedule, which writes purchase lead time, safety stock, reorder point, and preferred stock level for eligible item-locations.

6. **Open the Results, KPIs, and Diagnostics views** in a saved copy of the workbook.
7. **Apply the approval workflow** below before treating any number as a decision.

An Approval Workflow in Seven Gates

Each recommended change is a gated decision, with an exception owner for every failed gate.

- **Gate 1, currency:** the result was created after the last input change and after any Supply Analysis Interval change.
- **Gate 2, eligibility:** the item uses Reorder Point replenishment, is included, and the location makes inventory available.
- **Gate 3, data quality:** demand and lead-time windows pass the tests above, and the three-value lead-time floor is comfortably exceeded.
- **Gate 4, service level ownership:** the service level is a segment default or a documented override within the permitted range.
- **Gate 5, diagnostics:** previous projected service level, optimized projected service level, and distance from target are consistent with the proposed change.
- **Gate 6, economics:** expected shortage value, expected average inventory value, and turnover are reviewed together, and the change in safety stock times unit cost is recorded.
- **Gate 7, monitoring plan:** a named owner and review dates are set before new values are applied.

The economics gate is where finance participates. Oracle's educational material says holding costs often represent 20% or more of an inventory's total cost, a vendor-stated figure with no method disclosed, and academic work has long cautioned that carrying cost estimates vary and are often ignored (www.netsuite.com) (emerald.com). The safer approach is to use the company's own cost of capital and storage cost.

Sign-Off Register Template

Table 3 shows a register to be completed for each item-location approved, overridden, or held. Its fields mirror what the workbook stores plus the elements it lacks (live position and owner).

Item	Location	Demand window (start to end)	Supply window (start to end)	Target service level and source	Current reorder point	Proposed reorder point	Expected shortage value and inventory value change	Live on-hand and on-order (date)	Decision	Exception owner	Review date
SKU-X (Hypothetical Example)	DC-East	2026-04-01 to 2026-09-30	2026-04-01 to 2026-09-30	95%, Segment Default	630	704	Expected shortage value $33; safety stock value change +$888 (74 units at $12)	Snapshot from live search, dated	Approve, override, or investigate	Planner or buyer named here	30 and 90 days

The decision column accepts only three values, and every non-approval has an exception owner: the buyer for lead-time anomalies, inventory accounting for cost anomalies, and the planner for demand anomalies. Recording baseline dates also freezes what was known on the approval date, which is the basis for measuring results later.

UAT Checklist

User acceptance testing should run in a sandbox or on a representative subset before production use.

- **Features and roles:** the planner role opens the workbook and sees the same rows as the approver role.
- **Eligibility reconciliation:** active Reorder Point items equal included items plus documented exclusions.
- **Recalculation by hand:** for five item-locations, a spreadsheet recomputation from stored means and standard deviations matches the stored safety stock and reorder point after rounding up.
- **Interval change test:** change the Supply Analysis Interval in a sandbox, rerun, and confirm new rows carry new supply dates while old rows keep theirs.
- **Blank and zero test:** include an item with no demand, an item with fewer than three receipts, and an item with missing Average Cost, and confirm results match Oracle's documented behavior.
- **Override test:** apply an override and confirm the Override column shows Manual Override and that a later segmentation run keeps it.
- **Customization test:** save a workbook copy, edit it, and confirm the standard workbook is unchanged.

Measured Post-Change Monitoring

Monitoring belongs in the register, not in later improvisation. For each approved item-location, record the baseline window's actual stock-out events, backorder quantity, and average on-hand value, then compare the same measures over the review period, adjusting for demand changes. Projections are no substitute, since they do not describe actual transactions.

Data Analysis and Evidence

Macro Inventory Conditions as of November 2026

As of the Census Bureau releases for September 2026, total business inventories were estimated at an end-of-month level of $2,873.4 billion, with a total inventories-to-sales ratio of 1.33 (www.census.gov). Manufacturers' inventories-to-shipments stood at 1.46, down from 1.47 in August (www.census.gov), while the merchant wholesaler ratio was 1.21 against 1.34 a year earlier (www.census.gov). The St. Louis Fed's FRED service explains the wholesale ratio as months of sales on hand (fred.stlouisfed.org).

Lead-Time Signals

Lead-time variability is the input most exposed to external conditions. The New York Fed's Global Supply Chain Pressure Index (GSCPI) integrates transportation cost data and manufacturing indicators into a gauge of global supply chain conditions (www.newyorkfed.org). A May 2026 New York Fed analysis reported that the index rose by 1.3 points to 1.8 standard deviations above its average value, driven predominantly by

lengthening delivery times and growing order backlogs (libertystreeteconomics.newyorkfed.org). ISM's September 2026 report showed the Supplier Deliveries Index at 59.0, slowing for the 10th month in a row, and ISM notes that a reading above 50 percent indicates slower deliveries (www.ismworld.org). The October 2026 Empire State Manufacturing Survey said delivery times lengthened substantially and supply availability worsened further (www.newyorkfed.org). These indicators are national and cannot replace a supplier-level review, but they support a rule: with a long Supply Analysis Interval, the lead-time mean may lag a recent deterioration, and with a short one, the standard deviation may rest on few receipts.

What Drives Inventory Levels, and What Is Known About Cost

A panel study of 722 public companies summarized by the Wharton School identifies average demand, demand uncertainty, lead times, and cost of capital as the biggest determinants of inventory levels, consistent with the structure of the safety stock formula (knowledge.wharton.upenn.edu). On cost the evidence is untidy. A 1979 academic study stated that many firms do not even consider inventory carrying costs (emerald.com); the Council of Supply Chain Management Professionals' 2026 report, as relayed by a trade publication, put inventory carrying costs at 30.4% of total US business logistics costs in 2025, which is a share of logistics spend and not a percentage of inventory value (www.scdigest.com); and Oracle's vendor material cites holding costs of 20% or more (www.netsuite.com). No single authoritative carrying-cost rate emerged, so the register asks finance to supply the company's own.

Statistical Caveats

Forecast and demand error interact with the formula in ways stored results cannot show. A peer-reviewed paper notes that forecast errors across the periods of a lead time are positively correlated, which makes scaling per-period variance by lead time flawed, and a Lancaster University preprint reaches a similar conclusion about overlooked autocorrelations (ideas.repec.org) (eprints.lancs.ac.uk). A 2025 preprint adds that safety stock based on forecast-error variance is biased when errors are autocorrelated (arxiv.org). These cautions apply most strongly when daily demand has trend, seasonality, or clustering. Other platforms face the same issue: Microsoft's Dynamics 365 safety stock journal proposes a minimum from an item's historical usage (learn.microsoft.com), and Business Central says safety stock should absorb fluctuations in demand up to a targeted service level (learn.microsoft.com). Every platform's safety stock is a statistical proposal that a planner must validate.

Implications and Future Directions

For finance and operations leaders, the workbook shifts the review from "what number did the system produce" to "what evidence supports applying it". The September views expose projected shortage value, inventory value, turnover, and service-level diagnostics beside the recommendation, but they do not make the decision. With supplier deliveries slowing for ten consecutive months in ISM's September 2026 reading (www.ismworld.org), a recommended reorder point can be internally consistent and still wrong if the lead-time window includes late-recorded receipts, if a segment service level fits a location poorly, or if demand has changed since the window closed.

For implementation teams, the Supply Analysis Interval adds a configuration decision with change-management cost. Because it is account-wide and cannot be set per item or location, a distributor with both domestic and overseas suppliers must pick one window for both. The six-month default is a starting point and not a recommendation, so teams should test two or three values in a sandbox before changing production.

For multi-location networks, the item-location grain allows different protection at a hub and a satellite, and NetSuite supports replenishing a location by transfer order for both optimization-managed and standard Reorder Point items. Oracle's product pages describe multi-location visibility in marketing terms, a vendor claim to test in the account (www.netsuite.com). More rows mean more review, which argues for exception-based review: pass rows that clear every gate, and spend human time on the rest.

Best practices for NetSuite inventory management follow from this framing. Maintain record accuracy through cycle counting, since inaccurate records distort both demand history and the on-hand comparison (ideas.repec.org). Review ABC assignments periodically, because Oracle's educational material notes that up to 50% of items can change category every quarter or year (www.netsuite.com). Document overrides, keep the approval register, and treat the first quarter after a change as a measurement period. Re-check field names and view structure after each release.

Frequently Asked Questions (FAQs)

What does NetSuite inventory optimization do?

It recalculates lead time, safety stock, reorder point, and preferred stock level for eligible item-locations, and Reorder Point items outside the feature keep standard calculations (netsuitechangelog.com).

How does demand planning safety stock differ from Inventory Optimization?

In standard Advanced Inventory Management, the Default Safety Stock preference is entered in days (default 7), and NetSuite uses it to auto-calculate the item's reorder point (docs.oracle.com). When Demand Planning is used, NetSuite includes the safety stock amount in demand plan calculations (docs.oracle.com). Inventory Optimization instead derives safety stock from a service level and measured demand and lead-time variability.

What is the difference between reorder point and safety stock in NetSuite?

Safety stock is the buffer sized from lead-time-period demand variability and a service level z-score, while reorder point adds the average lead-time demand to that safety stock (en.wikipedia.org). The preferred stock level is a separate target that adds preferred days of demand to safety stock.

Does the workbook show current inventory?

No. The standard workbook has no current item-location quantity fields, and its values are stored results.

Why is a value blank instead of zero?

Zero is a valid result, while a blank indicates that required source data, an input, or a calculation result is unavailable.

Does changing the Supply Analysis Interval change old results?

No. It affects subsequent calculations and does not change dates on previously created results.

How does NetSuite optimize inventory across multiple locations?

Calculations run at the item-location level, with the service level potentially shared by item-level segmentation unless the per-location segmentation preference is checked.

Can NetSuite's recommendation guarantee a service level?

No. The diagnostics describe mathematical planning projections and do not represent achieved customer service, and the method assumes normally distributed lead-time demand.

Which service level should be used?

The segment default is the starting point, and overrides must be from 50% through 99.99%. ASCM describes typical goals of 90% to 98%, but the right level balances the cost of stocking out against the cost of holding extra inventory (www.ascm.org) (scm.ncsu.edu).

Conclusion

The September 2026.2 additions make NetSuite's safety stock and reorder point recommendations reviewable in a way earlier releases did not allow. The Results, KPIs, and Diagnostics views expose the inputs, the projected value and shortage, and a comparison of previous and optimized service levels. The same documentation defines what these views are not: stored results, without live on-hand, built on a normal-distribution assumption, with projections that do not describe achieved service.

The decision framework converts those facts into practice. Confirm features, permissions, and item eligibility. Check that the result is current and that its demand and supply windows are representative. Identify who owns the service level. Read target, previous projected, and optimized projected service levels as three distinct numbers, and fill rate as a fourth. Compare expected shortage value with expected inventory value before approving, treat blanks and zeros as different signals, record each decision with an exception owner, and measure actual stock-outs and inventory value afterward. The evidence reviewed here supports no claim that the algorithm improves service or reduces working capital in any account. Only the account's own measured results can establish that.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_42174309908.html
2. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_0807050821.html
3. https://netsuitechangelog.com/articles/inventory-optimization-results-workbook
4. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_0702045810.html
5. https://www.census.gov/mtis/current/index.html

6. https://www.ismworld.org/supply-management-news-and-reports/reports/ism-pmi-reports/pmi/september
7. https://en.wikipedia.org/wiki/Safety_stock
8. https://learn.microsoft.com/en-us/dynamics365/business-central/design-details-handling-reordering-policies
9. https://www.netsuite.com/portal/resource/articles/inventory-management/reorder-point-rop.shtml
10. https://netsuitechangelog.com/articles/inventory-optimization-reorder-point-netsuite
11. https://netsuitechangelog.com/articles/inventory-segmentation-location-2026-2
12. https://arxiv.org/html/2511.04616v1
13. https://knowledge.wharton.upenn.edu/article/inventory-management-a-way-to-give-it-a-grade
14. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_0702084957.html
15. https://netsuitechangelog.com/articles/inventory-optimization-results-dataset
16. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_58201401242.html
17. https://ideas.repec.org/a/inm/ormnsc/v54y2008i4p627-641.html
18. https://ideas.repec.org/a/eee/intfor/v35y2019i1p224-238.html
19. https://scm.ncsu.edu/scm-articles/article/safety-stock-analysis-inventory-management-models-a-tutorial
20. https://user.engineering.uiowa.edu/~dbricker/Lecture_Slides/Reorder_Point.pdf
21. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_0702044406.html
22. https://ocw.mit.edu/courses/15-772j-d-lab-supply-chains-fall-2014/0d50c5c77382852102ee30b98f1d4657_MIT15_772JF14_Lec14.pdf
23. https://www.ascm.org/ascm-insights/safety-stock-a-contingency-plan-to-keep-supply-chains-flying-high
24. https://itl.nist.gov/div898/handbook/eda/section3/eda3671.htm
25. https://www.netsuite.com/portal/resource/articles/inventory-management/safety-stock.shtml
26. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8853522
27. https://openstax.org/books/principles-financial-accounting/pages/10-5-examine-the-efficiency-of-inventory-management-using-financial-ratios
28. https://orca.cardiff.ac.uk/65447/1/Lengu%20et%20al%20%282014%20-%20EJOR%29.pdf
29. https://ideas.repec.org/a/wly/navres/v44y1997i2p165-186.html
30. https://help.sap.com/saphelp_scm700_ehp01/helpdata/en/8e/2bc95360267614e10000000a174cb4/content.htm
31. https://emerald.com/insight/content/doi/10.1108/eb014449/full/html
32. https://www.census.gov/manufacturing/m3/current/index.html
33. https://www.census.gov/wholesale/current/index.html
34. https://fred.stlouisfed.org/series/WHLSLRIRSA
35. https://www.newyorkfed.org/research/policy/gscpi
36. https://libertystreeteconomics.newyorkfed.org/2026/05/will-mounting-supply-chain-strains-hamstring-the-ai-investment-boom
37. https://www.newyorkfed.org/medialibrary/media/survey/empire/empire2026/2026_10_empire-state-manufacturing-survey.pdf
38. https://www.scdigest.com/firstthoughts/26-06-26.php?CID=23144
39. https://ideas.repec.org/a/eee/ejores/v256y2017i2p454-461.html
40. https://eprints.lancs.ac.uk/169200/1/Approximations_for_the_Lead_Time_Variance_a_Forecasting_and_Inventory_Evaluation_6_.pdf

41. https://learn.microsoft.com/en-us/dynamics365/supply-chain/master-planning/safety-stock-journal
42. https://www.netsuite.com/portal/products/erp/demand-planning.shtml
43. https://www.netsuite.com/portal/resource/articles/inventory-management/abc-inventory-analysis.shtml
44. https://docs.oracle.com/en/cloud/saas/netsuite/ns-online-help/article_162488535703.html

THE PUBLISHER

About Houseblend

Houseblend is a NetSuite consulting firm serving finance and operations teams. We help organizations implement ERP systems, connect business applications, improve existing configurations and maintain the systems that support everyday work. Our audience includes finance leaders, controllers, operations managers, NetSuite administrators and implementation teams.

Implementation and architecture

Houseblend provides NetSuite implementation, architecture and data migration services. We help teams evaluate how business processes, reporting requirements and existing data should fit together in an ERP environment. Training supports the people responsible for adopting and operating the resulting system.

Integrations, customization and AI

Our services include NetSuite integrations and customization, as well as AI integrations and AI transformation work. These engagements connect ERP data and workflows with the broader application landscape. The right design depends on the organization's systems, controls and operating needs.

Improve and support an existing system

Houseblend offers NetSuite health checks, optimization, managed support and project rescue services. We also provide expertise for analytics and specialist workflows, including NetSuite Analytics Warehouse, warehouse management and field service management. Published educational material helps teams investigate options and prepare informed questions for their implementation or support work.

Work with Houseblend

Explore NetSuite implementation, integrations, managed support and AI integrations. Contact Houseblend to discuss your current system and priorities.

Article examples explain concepts rather than promising a particular license, product capability, delivery schedule or outcome. Engagement scope is confirmed with the Houseblend team.

Website: https://www.houseblend.io

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.