

How to calculate **safety stock** for inventory management

Your guide to calculating safety stock using a range of formulas



Contents

The basics of safety stock	3
The importance of safety stock	4
Challenges of calculating safety stock	5
3 common methods for calculating safety stock	6
Statistical calculations for safety stock	10
Examples of statistical safety stock formulas	11
Dynamic safety stock	20
Summary	21
Appendix	22

The basics of **safety stock**

Determining appropriate stock levels is one of the most important and challenging tasks faced by inventory managers. If you carry too much, you'll tie up working capital. Carry too little and you'll risk potential stock-outs and unhappy customers.

There needs to be a balance between investing the right amount of capital in inventory and achieving service level targets (a measure of stock availability). One key challenge – and an important piece in the puzzle – is calculating safety stock levels.


Safety stock, also called buffer stock, is important at all stages of the supply chain. Safety stock is the extra layer of stock (raw materials or finished goods) that businesses carry to mitigate the risk of run-out due to uncertainties in supply or demand. Put simply, safety stock is used to prevent service levels from being negatively impacted by unforeseen changes in demand or supplier lead times.

In this guide we'll look at various ways to calculate safety stock, starting simple and then moving to more advanced, statistical formulas.



The importance of safety stock

Safety stock is used to manage any demand or lead time uncertainty and guard against supply chain or fulfillment disruptions. Deficiency situations can arise because of fluctuating demand, forecast errors and variation in suppliers' lead times.



Safety stock is intended to cover any shortfall in cycle stock during the lead time period. It is an important element of the **reorder point (ROP) formula**:

$$\text{ROP} = (\text{avg. consumption} \times \text{lead time}) + \text{safety stock}$$

The challenges of calculating safety stock

The goal of safety stock is to minimize disruption to order fulfillment, while investing the lowest possible amount of capital in inventory.

When demand for inventory items is consistent and lead times are reliable, it's fairly easy to set safety stock levels that will achieve this. But, when demand and supply fluctuate, many inventory planners find it much tougher to calculate safety stock accurately. All too often they chose to use simple safety stock formulas, which are insufficient to deal with the supply and demand challenges they face.

Why carry safety stock?

- To protect against demand and lead time variance
- To compensate for forecast errors (when demand exceeds forecast)
- To fulfill demand if there are disruptions in production or deliveries
- To avoid stock-outs and back orders that can lead to lost sales and poor customer service





3 common methods for calculating safety stock

Here are three simple methods used by inventory managers to calculate safety stock.

01 | Fixed safety stock

Many companies set a fixed level of safety stock for their inventory items e.g. they add a 'best-guess' quantity to the reorder point to allow for any issues.

This number is often set at item group level and based on the judgment or assumptions of the inventory management team – there are no formal calculations used. For example, an inventory planner may decide to carry a week's worth of safety stock and use last month's highest week of sales as the safety stock figure.

Pros/Cons

While easy to set up and manage, this basic approach often leads to stock imbalances. This could mean an unnecessarily large amount of capital invested in excess stock, while stock-outs could be frequent on other items.



02 | Time-based calculations

A time-based safety stock calculation finds the average sales/demand over a fixed period of time and uses this value as the safety stock level. For example, it's simple to calculate average demand based on last month's consumption; in the case below, it would be 200 units.

Period	Consumption
Week one	300 units
Week two	150 units
Week three	250 units
Week four	100 units

Therefore, safety stock for the upcoming month would be 200 units.

Pros/Cons

Both fixed and time-based calculations take a “one size fits all” approach. They assume that the forecasted demand will be accurate and that lead times will remain consistent. In reality, of course, this rarely happens.

For starters calculating demand based on historical sales alone cannot account for demand variance, particularly for items with erratic or intermittent demand.

At the same time, lead times can be inconsistent due to a range of issues, such as production downtime or delivery delays.

Using a basic safety stock model will no doubt result in getting the right levels of stock for some items, but others will see levels too high and others will be too low.

03 | Average/max. calculations

A more prudent safety stock calculation is to use an average/max. formula, which accounts for when lead times rise and sales max out.

In our previous example, average sales were 200 units a week but rose to 300, and let's say that the lead time is 1 week but can be as high as 1.5 weeks. Using the formula below, we can work out safety stock:


$$\begin{aligned} &(\text{max. sales} \times \text{max. lead time}) - (\text{avg. sales} \times \text{avg. lead time}) \\ &(300 \times 1.5) - (200 \times 1) \\ &= 250 \text{ units} \end{aligned}$$

Pros/Cons

Problems arise with this formula if the maximum lead time and sales are considerably higher than the average, which results in the levels of safety stock being significantly inflated.

While simple to use, a key drawback of all the formulas so far is that they fail to link back to service levels. Service levels are important as they are connected to how well you can serve your customers with on-time deliveries.

Any out-of-stock item, even if it's just one SKU, will lead to an incomplete order which can be detrimental to customer satisfaction. Therefore the service level KPI closely correlates with customer service, acquisition, loyalty and retention.



Statistical calculations for safety stock

Statistical safety stock calculations overcome many drawbacks we've just discussed. Despite being more complicated to use, they are much more accurate. This is because they use probability distributions to model demand and account for variance.

A probabilistic approach accepts that there is uncertainty when predicting future events, such as demand volume and frequency, and accommodates for this by covering for a percentage of all possible inventory requirements.

Statistical safety stock formulas are based on achieving a desired service level. Service level is the expected probability of being able to satisfy all possible demand scenarios within a particular period of time.

For example, if you set a service level target of 99%, this means your safety stock levels will cover for 99% of all probable requests; in other words, you can give your customers what they want, when they want it, 99% of the time.



Examples of statistical safety stock formulas

Safety stock formula with demand uncertainty

Let's start with a basic statistical safety stock calculation. This formula assumes that there are no supplier lead time variations. It only takes into account variation in demand.

$$SS = Z * \sigma_d * \sqrt{LT}$$

Key:

- SS = Safety stock level
- Z = Z-score
- σ_d = Standard deviation of demand
- LT = Average lead time

This formula introduces two statistical elements – the Z-score and standard deviation. Let's break down each element in turn.




Z-score

The service factor, or Z-score, is based on your service level target – balancing inventory costs with the risk of a stock-out. The higher the desired service level, the more safety stock is required.

Once you set your service level target, you can work out your Z-score using a normal distribution table.


Normal distribution is used to allow you to predict the probability of selling a certain quantity of stock. With normal distribution, the probability of an occurrence, if graphed, is a bell shaped curve, where all data is equally spread out around a central value with no bias left or right.



For example, if you sell an average of 1,000 units, you are as likely to sell more than 1,000 next month as you are to sell less than 1,000. The probability decreases the further from the center of the curve e.g. the probability of selling 2,000 is lower than selling 1,500, as is the probability of selling 500 over 750.


As the chart on the next page shows, if you decide to set your service level target at 50% you don't need safety stock because you have a 50/50 chance of selling more or less than the average next month.

However, if you want a service level of 95%, the normal distribution will give you a Z-score (a multiplier coefficient) of 1.64.



The Z-score can easily be determined using a service factor table (as shown on page 22), or you can use Excel function NORMSINV to convert a service level percentage to a service factor.

The graph below shows the relationship between target service level and the Z-score. The relationship is non-linear e.g. higher service levels require disproportionately higher Z-scores and therefore disproportionately higher safety stock levels.



Standard deviation of demand

Standard deviation is used in statistics to measure how spread out a certain set of numbers is from the average. This allows the safety stock calculation to take into account demand variance.

Standard deviation of demand is calculated like this:

1. Determine the mean (average) of a set of numbers
2. Determine the difference between each number and the mean
3. Square each difference
4. Calculate the average of the squares
5. Calculate the square root of the average

Let's take a working example. Mary works at The Plumbing Superstore. Below you can see her forecasted demand for a line of showerheads. The average weekly demand is 525 units.

Period	Forecasted demand: shower head A	Difference between actual and forecasted demand	Difference squared
Week one	550	25	625
Week two	500	-25	625
Week three	650	125	15,625
Week four	400	-125	15,625

The total of the squared differences is 32,500 units.

The average of 32,500/four periods is 8,125 units.

The square root of 8,125 is 90.

90 units is the standard deviation of demand per week.

You can also use Excel function STDEVPA to calculate standard deviation. In safety stock calculations, the forecast quantity is often used instead of the mean in determining standard deviation.

Always ensure that you use the same unit of measurement for both the demand and lead time calculations e.g. if your standard deviation of demand is calculated in days, then calculate your average lead time in days too!

Let's now put all these elements together, using The Plumbing Superstore as an example. Mary wants to achieve a 95% service level for her showerheads. In the service factor table on page 22, this gives a Z-score of 1.64. The workings on the previous page show the standard deviation of demand is 90 units per week and average lead time is 3.5 weeks.

$$SS = Z * \sigma_d * \sqrt{LT}$$

$$SS = 1.64 * 90 * 1.87$$

$$SS = 276 \text{ units}$$

Pros/Cons

If your supplier lead times are fairly stable and you rarely have delivery delays, this formula will work well. But in practice, this is often unlikely and you need to add factors to compensate for variance in your supply chain.

Safety stock formula with demand uncertainty

In the previous equation, safety stock was only used to mitigate demand variability. When variation in lead time is also a concern, the equation must be expanded. The formula below does just that:

$$SS = Z * \sqrt{\sigma_d^2 * LT + x^2 * \sigma_{LT}^2}$$

Key:

- SS = Safety stock level
- Z = Z-score
- σ_d = Standard deviation of demand
- LT = Average supplier lead time
- x = Average demand
- σ_{LT} = Standard deviation of supplier lead time

Let's take a look at a working example of this formula, using The Plumbing Superstore.

Mary still wants to provide a 95% service level for the showerheads.

- Average demand is 525 units per week
- The standard deviation from demand is 90 units
- The average supplier lead time is 3.5 weeks, but is as fast as 2.5 weeks and as slow as 4.5 weeks
- The standard deviation in lead time is 0.79 weeks (approx. 5.5 days stock)

$$SS = Z * \sqrt{\sigma_d^2 * LT + x^2 * \sigma_{LT}^2}$$

$$SS = 1.64 * \sqrt{90^2 * 3.5 + 525^2 * 0.79^2}$$

$$SS = 1.64 * \sqrt{28,350 + 172,018}$$

$$SS = 1.64 * 447.6$$

$$SS = 734 \text{ units}$$

As you can see, the safety stock level is now much higher, as lead time variance is also being taken into account.

Pros/Cons

This formula is based on the assumption that the difference between the forecast and the actual demand follows a normal probability distribution. If the forecast error does not have a normal distribution, the quality of the result will be diminished.

For items that have intermittent demand patterns, such as slow and lumpy demand, you'll achieve higher accuracy using Poisson distribution or negative binomial distribution.





Choosing the right **safety stock model**

Whichever safety stock calculation you use, it's very important to carefully test it prior to final implementation. This is so you can ensure that it works correctly and to analyze what impact it will have on inventory levels and cashflow.

While more sophisticated models lead to more accurate safety stock levels, they also take time and resources to implement. Manual calculations are also frozen in time and virtually impossible to implement at SKU level.

A solution is to use software that will automatically do this for you. A good inventory optimization tool will reduce the time involved in undertaking manual calculations and avoid the risk of human error and the associated costs.

Dynamic safety stock

An inventory optimization tool like EazyStock takes statistical safety stock calculations to the next level. Safety stock levels are set at SKU level and consider a number of key elements:

Demand type

EazyStock uses a range of statistical algorithms and probability distributions e.g. normal distribution, Poisson distribution, empirical distribution and negative binomial distribution. The most appropriate algorithm is chosen based on each item's demand type. So whether an item's demand is fast, slow, lumpy or erratic, EazyStock will calculate safety stock as accurately as possible.

Supplier lead times

EazyStock's dynamic lead time feature reviews purchase order history and alerts you when supplier lead times deviate from the norm. You can then request the system automatically increases safety stock levels, to help achieve target service levels.

Safety stock calculations are automatically reviewed and updated daily (when required) based on service level targets, demand variance, forecast accuracy, supplier lead times and order frequency. When any of these factors change, EazyStock dynamically updates safety stock algorithms and adjusts final safety stock levels accordingly.

With a tool such as EazyStock, you no longer need to calculate and update safety stock manually and enter this into your ERP. Instead EazyStock does all the math for you.

Summary

Today's marketplaces and supply chains are more dynamic and uncertain than ever. Demand and supply volatility is almost inevitable, making safety stock a necessity for many businesses.

It's important to calculate safety stock levels as accurately as possible, so you have the right balance between inventory investment and stock availability e.g. hitting your service levels.

Statistical safety stock calculations are much more accurate than fixed safety stock or time-based methods, as they allow for more variance in demand and lead times. However, using these more complex formulas in manual spreadsheets is time-consuming and can lead to errors – especially if you manage a large number of SKUs.

Companies who really want to optimize their safety stock and inventory levels should consider an inventory optimization software solution.

This will plug into existing ERP or stock control systems and automatically calculate forecasts, safety stock, service levels and reorder points, giving operations teams the data and the time to make informed, strategic supply chain decisions.



Target service level	Service factor (Z-score)
50.00	0
55.00	0.13
60.00	0.25
65.00	0.39
70.00	0.52
75.00	0.67
80.00	0.84
81.00	0.88
82.00	0.92
83.00	0.95
84.00	0.99
85.00	1.04
86.00	1.08
87.00	1.13
88.00	1.17
89.00	1.23

Target service level	Service factor (Z-score)
90.00	1.28
91.00	1.34
92.00	1.41
93.00	1.48
94.00	1.55
95.00	1.64
96.00	1.75
97.00	1.88
98.00	2.05
99.00	2.33
99.50	2.58
99.60	2.65
99.70	2.75
99.80	2.88
99.90	3.09
99.99	3.72



**Learn more about
inventory optimization**

Book a demo