Operating the Precision-Flow 5000

You are the Lead Quality Engineer at **Statypus Water**. To ensure your production line is running efficiently, you have programmed your new bottling machine to a target mean (μ) of **500.5 mL**.

While you can control the target, you are bound by the machine's physical hardware. The manufacturer guarantees that the machine is accurate "up to" a specific tolerance: an **Ethereal Population Standard Deviation** (σ) of 0.2 mL.

The **Central Limit Theorem (CLT)** tells us that while the machine's 0.2 mL noise is constant for every bottle, the distribution of our **sample average** ($\bar{x}$) becomes significantly more concentrated as our sample size (n) grows. Our goal today is to audit the line and see how often our production stays within our preferred **Target Zone** of **500.2 mL to 500.8 mL**.

Phase 1: The Lone Wolf ($n = 1$)

If we pull just one bottle of Statypus Water off the line, our "ruler" for its precision is simply the machine's inherent tolerance: $\sigma = 0.2$. At this stage, we are measuring raw, unadulterated machine noise.

Calculation 1: Individual z-scores

Calculate the z-scores for our 0.3 mL Target Zone (500.5 ± 0.3):

Reflection: The Individual

On your shading sheet, locate these z-scores on the first curve and shade the area. Looking at your shaded region, how much of the total production (the area under the curve) is actually falling within our 0.3 mL target? Is this machine as precise as you expected for a single bottle?

Phase 2: The Four-Pack ($n = 4$)

Now, imagine we take a sample of 4 bottles and average them. According to the **CLT**, the machine hasn't changed. However, the random variations of the individual bottles (some slightly over, some slightly under) begin to cancel each other out.

To measure the consistency of this *average*, we need a new, more refined ruler: the standard deviation of the sample mean ($\sigma_{\bar{x}}$).

Calculation 2: The Shrinking Ruler

1. Calculate the new ruler for the average ($\sigma_{\bar{x}}$):

2. Calculate the z-scores for the **same 0.3 mL target**:

Reflection: The Power of Averaging

Shade the area on the second curve of your shading sheet. Even though the physical target is still 0.3 mL, your shaded "net" has widened significantly. Why is the average of 4 bottles so much more reliable than a single "Lone Wolf" bottle?

Phase 3: The Full Case ($n = 16$)

Finally, we look at the average of a full 16-pack. In the world of Statypus Water, this is where the **CLT** turns our "Shotgun" into a "Sniper Rifle." The variation is now so compressed that our 0.3 mL target should be nearly impenetrable.

Calculation 3: The Sniper Rifle

1. Calculate the final ruler for the average ($\sigma_{\bar{x}}$):

2. Calculate the z-scores for the **same 0.3 mL target**:

Reflection: Statistical Certainty

Attempt to shade the third curve on your sheet. If the machine is working perfectly at its 0.2 mL tolerance, is it likely that a 16-pack average would ever fall outside this 0.3 mL zone? Explain your reasoning using the z-scores you just calculated.

The Statypus Audit

Every bottle of Statypus Water is labeled as **500 mL**. While you set the machine to 500.5 mL to provide a safety buffer, you pull a random 16-pack off the line and find an average volume of **500.1 mL**.

Final Engineering Assessment

1. Your z-score for this 16-pack average (relative to our 500.5 mL setting) is: ____________

2. Based on your Phase 3 calculations, is it *physically possible* to get this result if the machine is actually operating at the promised 0.2 mL tolerance?

3. Even though 500.1 mL is above the 500 mL label, why should this result alarm you as a Quality Engineer? What does it suggest about the machine's actual performance?

Phase 1: Average of n = 1

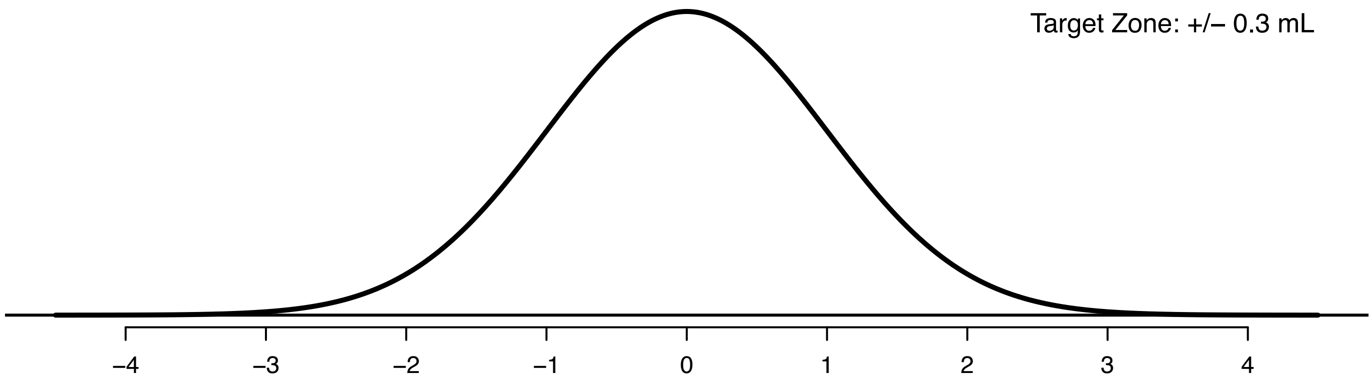

Phase 2: Average of n = 4

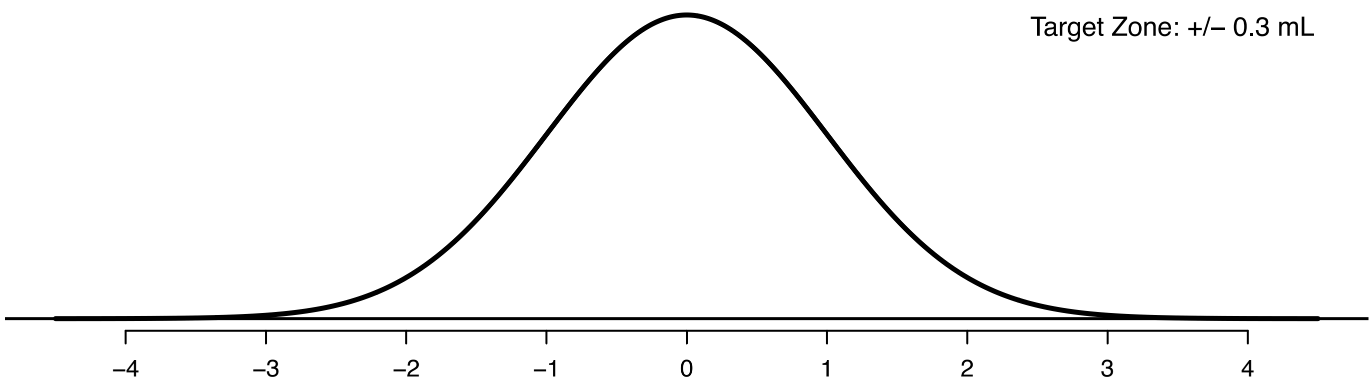

Phase 3: Average of n = 16

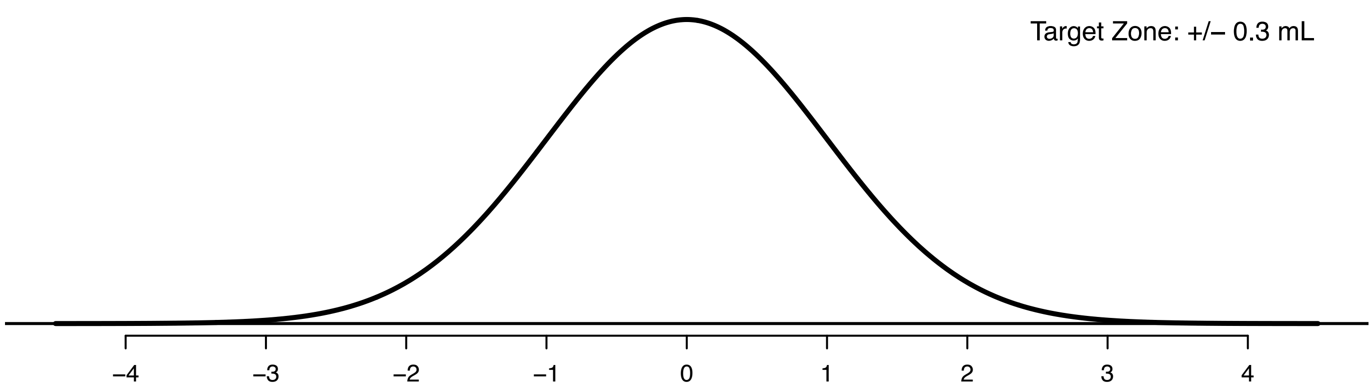

Final Audit: The 500 mL Legal Limit (Individual Bottle)