


# System Performance Simulation

Product backlog 2019-03-18

1. [SPRINT 1] As an engineer, I want to run the simulation with a particular day's data rates, so that I can compare actual system performance with the simulation result.
2. [SPRINT 1] As an engineer, I want to know average message latency and throughput, so that I can evaluate if my design is acceptable.
3. [SPRINT 1] As an engineer, I want to try monolithic design in simulation, so that I can evaluate the size of the buffer needed to process all data with no loss.
4. [SPRINT 2] As an engineer, I want confidence that the simulation results are correct, so that I can trust the simulation.
5. [SPRINT 2] As an engineer, I want to use more realistic input data rate model, so that the simulation produces more realistic results.
  - a. Instead of queuing all messages once per second, they need to arrive every micro-second.
  - b. Example: if the data rate specified in a file is 1000000 messages per second, then 1 message is added to the queue every microsecond (thus resulting in 1000000 messages a second).
6. [SPRINT 2] As an engineer, I want to know the distribution of latency as 50<sup>th</sup>, 75<sup>th</sup>, 90<sup>th</sup>, 99<sup>th</sup>, and 99.9<sup>th</sup> percentiles.

An example distribution is shown below. In this chart, 50% of messages had latency of 1000 nanoseconds or less, and 99.9% of messages had latency of 9000 nanoseconds or less. Note, that providing the data for the chart is sufficient (no need to make the chart).



7. As an engineer, I want to try "task parallel" design in simulation, so that I can evaluate the advantage of splitting monolithic system into an "assembly line".

8. As an engineer, I want to see the latency distribution of the simulation, so that I can tell the 50<sup>th</sup>, 75<sup>th</sup>, 90<sup>th</sup>, 99<sup>th</sup>, and 99.9<sup>th</sup> percentiles of message latency.
9. As an engineer, I want to combine "task parallelism" with "data parallelism" in a design simulation, to determine the minimal resources needed to meet latency and throughput expectations.
10. As a product owner, I want to test various data rates scenarios, so that I can plan future product improvements.