TidalScale SOLUTION BRIEF | ANYLOGIC

Tackle Any Model, Any Simulation

When companies across an array of industries need to tackle complex business challenges, they turn to AnyLogic, the standard in multimethod simulation modeling.

But even with a powerful software platform like AnyLogic, analyzing large, complex models poses unique computing challenges.

Data growth makes it worse. Those challenges only intensify as data sets grow—requiring not just powerful software, but lots of memory and plenty of CPU cores.

Limited options.
In the past, when large problems overwhelmed your largest server, your options were limited.

- Incur costs by purchasing expensive specialized computers
- Lose time by rewriting applications using complex distributed algorithms to run across clusters
- Delay insights by shrinking the problem to fit the limits of your server

A server sized to solve your largest problems.
With TidalScale, you can push past the limits of your existing hardware and get even more value from AnyLogic. No matter how big your model or detailed your simulation, TidalScale's Software-Defined Servers deliver all the memory and cores you need, on demand.

Size Servers to Fit Any Workload

With TidalScale, you can create a Software-Defined Server to run models and simulations across dozens of commodity servers. TidalScale's technology combines those servers (along with all their associated resources, including memory, cores, storage and bandwitdth) into a single aggregated system.

Let your model or simulation decide your server size.
Incorporate up to hundreds of processor cores, tens of terabytes of memory, and all the storage and networking I/O you need.

Scale as needed, no coding necessary.
Just select the ideal combination of resources from available systems in your datacenter or cloud. TidalScale makes your Software-Defined Server available to your operating system and applications without a single modification.

Create a new server in minutes.
Using TidalScale WaveRunner, our point-and-click control panel, you can configure and deploy a new Software-Defined Server in as little as five minutes. Reconfigure just as quickly.

Get More Value from AnyLogic

With TidalScale, modeling with granularity is easy at any scale.

Deploy Your Largest Workloads
Stop limiting the size of your model to the memory, cores or other resources available on a single server. TidalScale pools the resources of multiple commodity servers into one or more virtual servers—available on demand.

Engage in Exploratory Analysis
Scale without changing your models, tools operating environment--without the cost or burden of traditional scale-up or scale-out solutions.

Get to Work in 5 Minutes
Configure and deploy a Software-Defined Server in minutes.

Scale on the Fly
Match your server instance to the size of your AnyLogic workload. No more limiting inputs or granularity.

Keep Up with Your Data
Just as quickly, reconfigure your Software-Defined Server to accommodate new workloads.

Achieve Results Faster
Cut the time required to validate models.

Drive Down Costs
Tackle your biggest data challenges using already available commodity hardware, either on-premise or via cloud IaaS providers.

USE CASE

Modeling Customer Behavior in AnyLogic

Customer:
A large financial institution running an AnyLogic model of customer behavior.
Challenge: Previously, models were too large and detailed to achieve the granularity and sensitivity analysis the institution wanted. This delayed results.

Solution:
Needing more than 3TB of memory to properly run its model, the customer used TidalScale to combine five physical servers into a single system with 3.5TB. Analysts ran the AnyLogic model on unmodified Linux 7.2.

Results:
More concurrent runs, faster time to results. For the first time, the institution was able to run its simulation at a 1:1 granularity level and perform sensitivity analysis on the outcomes.

3 Breakthroughs on TidalScale

TidalScale's ability to size servers to the model had a dramatic impact on three aspects of AnyLogic modeling.

Model granularity

Increasing the agent count improves confidence in the result, but it requires more memory.

TidalScale Impact: Ability to scale agent count into the tens of millions for more accuracy in high-impact business decisions.

Sensitivity analysis

Analysts vary input parameters as they repeatedly invoke the model. This vets the sensitivity of the model to input assumption. The agent count, concurrent runs, and iterations can scale in both memory and cores.

TidalScale Impact: Eliminated limitations on the number of concurrent runs and scale. Decreased the time needed to validate a model.

Optimization

Analysts repeatedly invoke a model to locate the inputs needed to achieve the desired outcome. The scaling factor, number of inputs optimized for, and iterations all scale in both memory and CPU resources.

TidalScale Impact: TidalScale allowed optimization across all of the inputs rather than some.

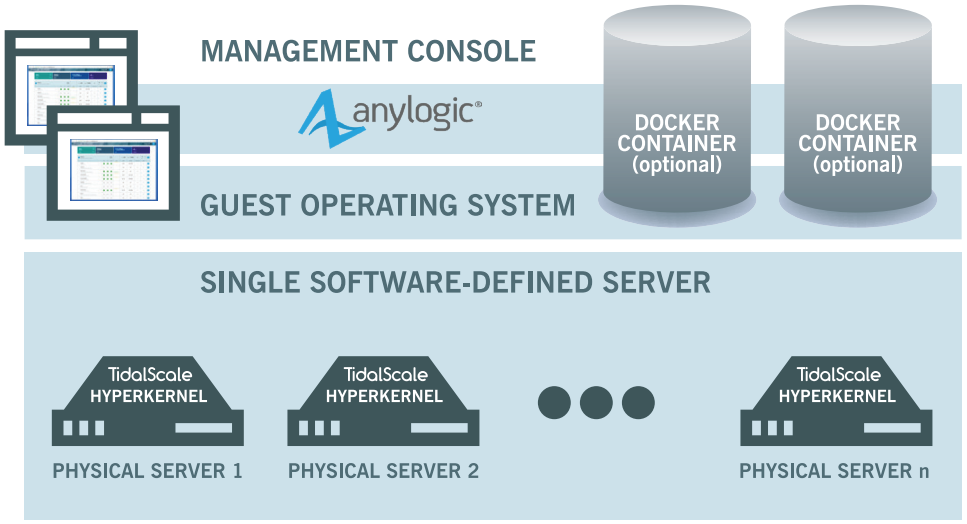

anylogic®

FIND OUT MORE:
www.tidalscale.com

REV 2018-01-22.001