Решения и технологии Intel

Технологии Intel для PostgreSQL

Владимир Слинько, Intel

Правовая информация

Использование технологий Intel может потребовать соответствующего оборудования, программного обеспечения или активации обслуживания. Никакая продукция или компоненты не являются абсолютно безопасными. Ваши расходы и результаты могут варьироваться. Оптимизация Intel для компиляторов и другой продукции может не осуществляться в той же мере для продукции других производителей.

Intel не контролирует содержание и не проводит аудит информации, предоставленной партнерами. Для оценки достоверности такой информации вам следует проверять другие источники.

ТЕХНОЛОГИИ INTEL

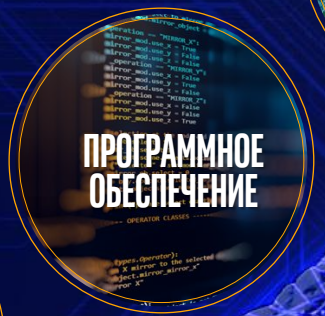

intel®

Архитектуры общего назначения

СКАЛЯР

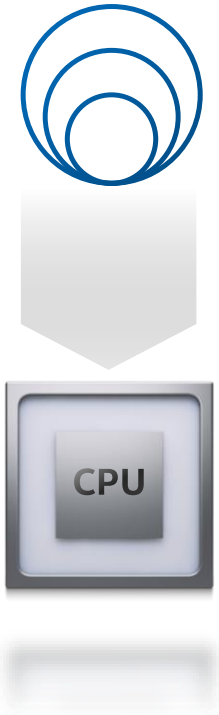

ВЕКТОР

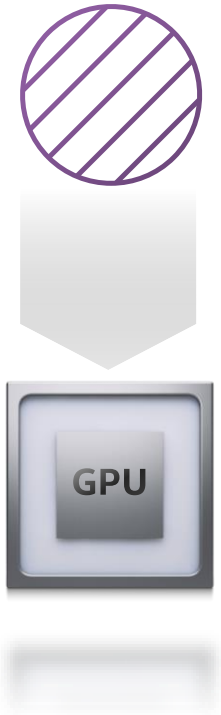

МАТРИЦА

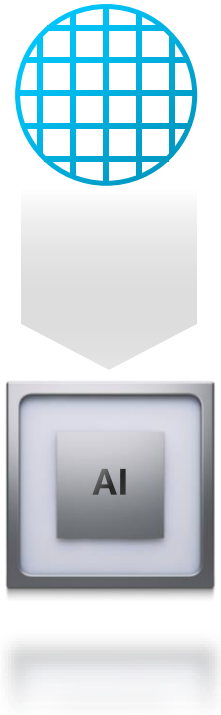

ПРОСТРАНСТВА

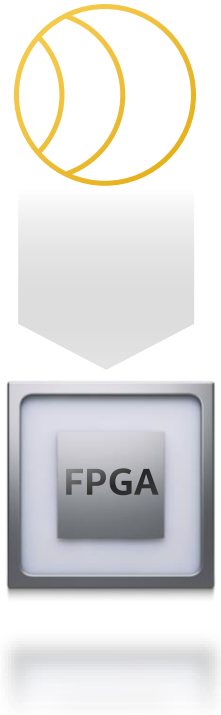

АРХИТЕКТУРНЫЙ ВЫБОР

Баланс скорости и универсальности

Иллюстрация. Масштаб не соблюдается

intel®

3-е поколение Intel Xeon® Scalabale

созданы для современных сервисов на основе баз данных

1.9X
СРЕДНИЙ ПРИРОСТ ПРОИЗВОДИТЕЛЬНОСТИ
К 5-ЛЕТНЕЙ ПЛАТФОРМЕ

ВСТРОЕННОЕ УСКОРЕНИЕ ИИ

Intel Deep Learning Boost
HOBOE: bfloat16*

до **1.98X**
ВЫШЕ ПРОИЗВОДИТЕЛЬНОСТЬ БАЗ ДАННЫХ
к 5-ЛЕТНЕЙ ПЛАТФОРМЕ

НОВАЯ ПАМЯТЬ

Intel Optane Persistent Memory
200 Серии

ДЛЯ 4S-8S СИСТЕМ

ГИБКОСТЬ

Улучшенная технология
Intel Speed Select

Alibaba Cloud · AsiaInfo 亚信科技 5G 原力进化 · Baidu 百度 · FACEBOOK · FUJITSU · GIGABYTE · 紫光集团 H3C 核心企业 数字化解决方案领导者 · 海鑫科金 HISIGN TECHNOLOGY · Hewlett Packard Enterprise · HITACHI · HUAWEI · hyve solutions · inspur 浪潮 · Inventec Inventec Data Center Solutions · Lenovo · Neusoft · Quanta Computer · SAP · sas · SUPERMICRO · Tencent Cloud · wiwynn · ZTE

intel®

*Available on select 3rd Gen Intel Xeon Scalable processors

Intel Optane™ Persistent Memory

Энергонезависимая память

Развитие среды обитания Баз Данных

СПРОС НА ВЫЧИСЛЕНИЯ РАСТЕТ

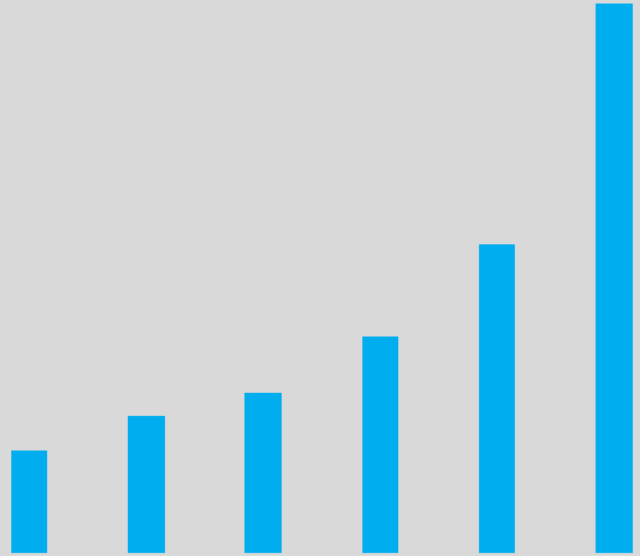

CPU Core Growth Projection Over Time

РОСТ DRAM ЗАМЕДЛЯЕТСЯ

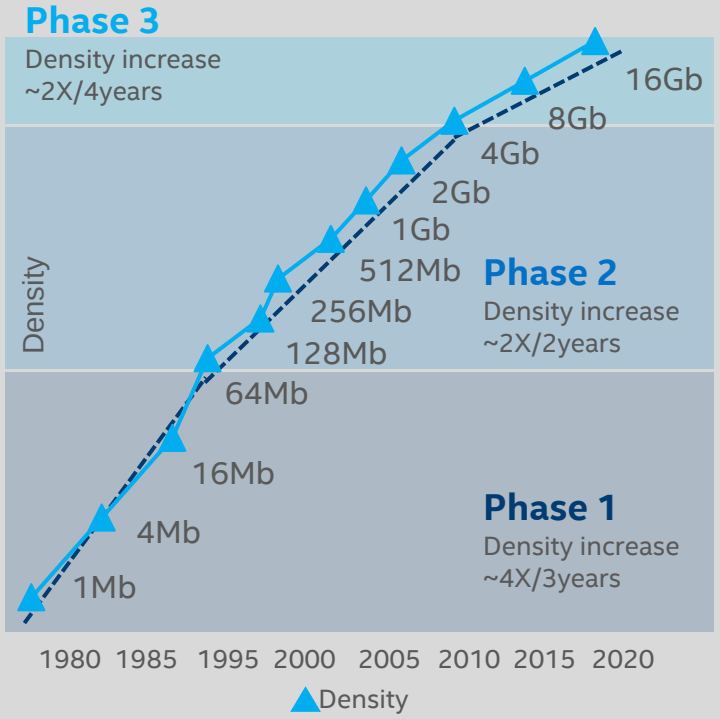

ПАМЯТИ НАДО ВСЕ БОЛЬШЕ

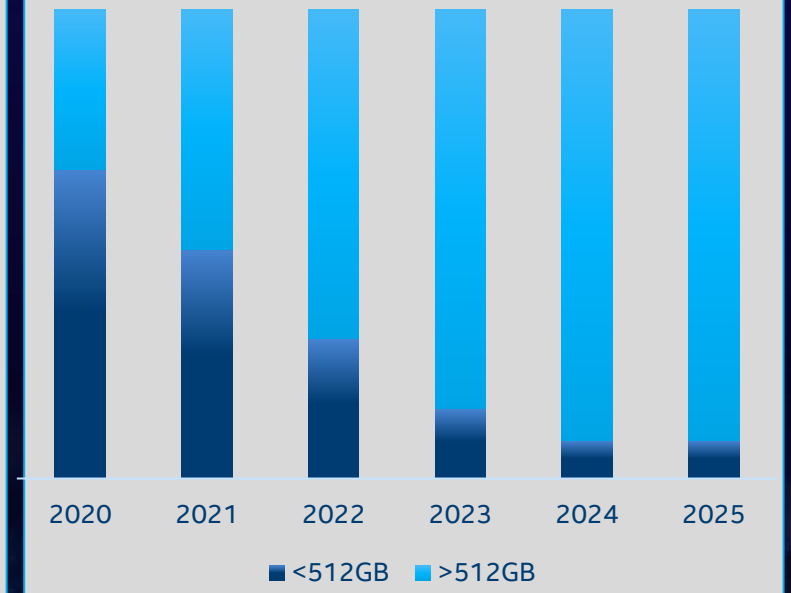

Источник “Compute Demand Accelerating and Large Memory Systems Growing”: Intel Market Intelligence Group - DRAM Market Monitor Q1-20" by Yole Development
Источник «DRAM is Not Scaling»: ”3D NAND Technology – Implications for Enterprise Storage Applications” by J.Yoon (IBM), 2015 Flash Memory Summit

Полная иерархия памяти и хранения

90/10 "ПРАВИЛО" ЛОКАЛЬНОСТИ

10X CAPACITY

1/10TH PERFORMANCE

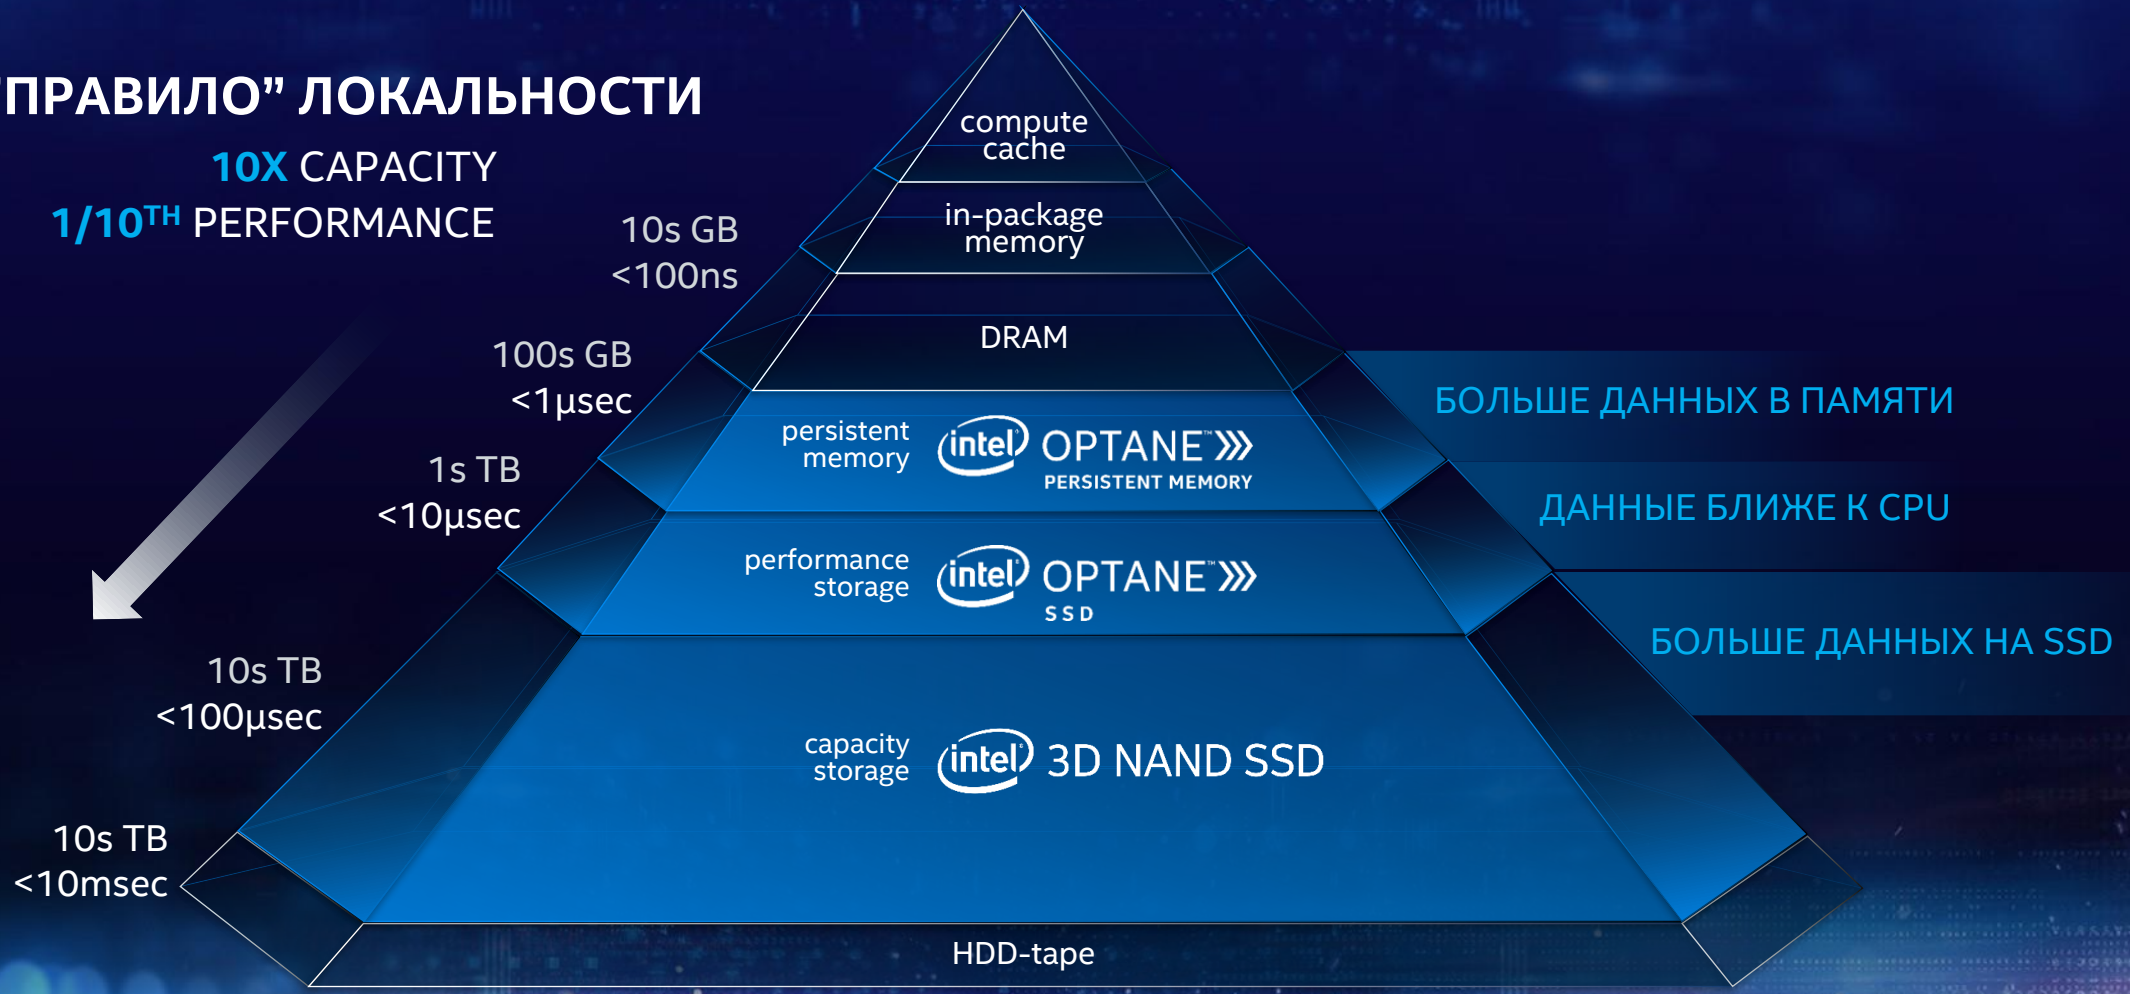

intel

Технология Intel Optane™ уникальна

ПОСТОЯННАЯ

информация сохраняется даже при потере питания

ПРЯМАЯ ЗАПИСЬ

Данные перезаписываются сразу, **очистка носителя не требуется**

ПОБАЙТНАЯ АДРЕСАЦИЯ

Каждая ячейка может быть адресована **индивидуально**

НИЗКИЕ ЗАДЕРЖКИ

Быстрый носитель и интерфейсы

ЛУЧШЕЕ ОТ ПАМЯТИ И ДИСКОВ

intel®

Источник: “The Challenge of Keeping Up with Data” https://www.intel.com/content/www/us/en/products/memory-storage/optane-dc-persistent-memory.html

Обзор Intel Optane Persistent Memory (PMem)

Intel Optane PMem 100 series

2nd Gen Intel® Xeon® Scalable processors

Available on 2S/4S/8S platforms

up to		
BASELINE	**18W**	**128, 256 & 512GB**
bandwidth	TDP	capacity

Supports up to **4.5TB** total memory per socket

Intel Optane PMem 200 series

3rd Gen Intel® Xeon® Scalable processors

Available on 4S platform

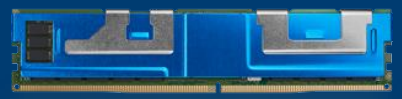

average		
25%	**15W**	**128, 256 & 512GB**
more bandwidth vs 100 Series PMem	TDP	capacity

Supports up to **4.5TB** total memory per socket

УВЕЛИЧЕН ОБЪЕМ ПАМЯТИ ДЛЯ РАЗЛИЧНЫХ ПРИЛОЖЕНИЙ

Memory Mode

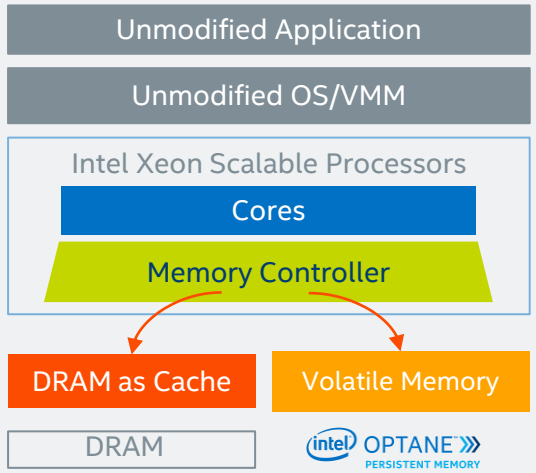

ДОБАВЛЕН НОВЫЙ УРОВЕНЬ ХРАНЕНИЯ PERSISTENT MEMORY

App Direct Mode

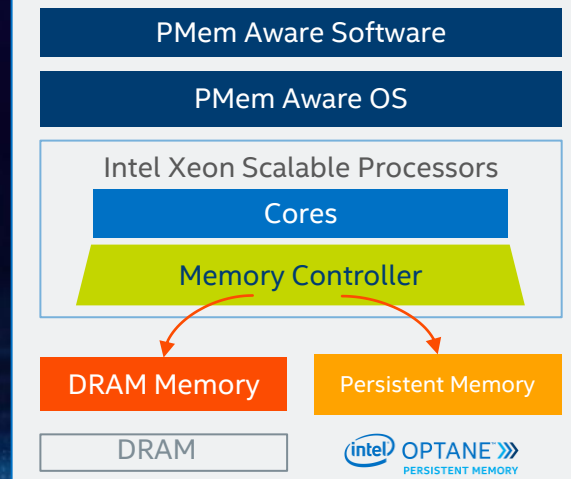

intel®

Примеры использования

Customer Success Story

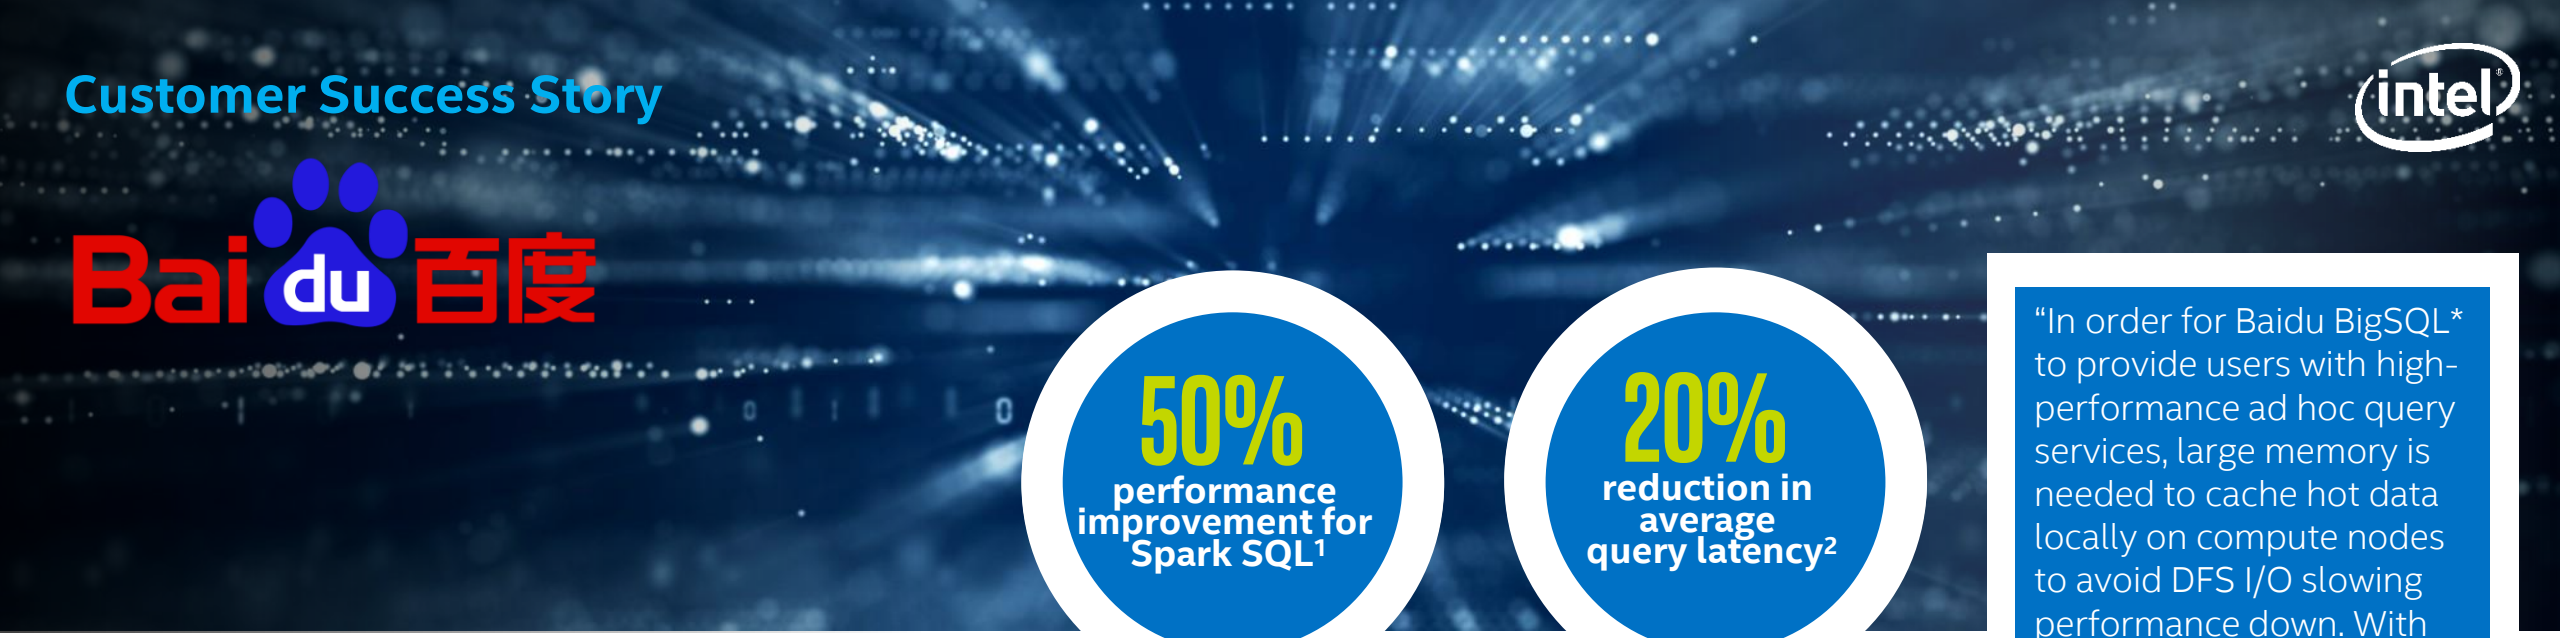

Baidu BigSQL Delivers Faster Spark Interactive Queries with Lower TCO

Baidu's BigSQL data processing platform is based on Spark SQL* and has many features and performance enhancements that improve on it. As Baidu's business expands, the scale of hot data grows rapidly. Memory scaling is needed to deliver the same level of performance that users demand. However, the high cost of Dynamic Random-Access Memory (DRAM) adds increasing pressure to the total cost of ownership (TCO). To lower TCO while ensuring satisfactory performance, Baidu deployed Intel® Optane™ persistent memory and used it to optimize its ad hoc query service – Tuling. Supported by Intel Optane PMem, the cluster offloaded more than 30% of the workload from Tuling[3]. Additionally, the average query latency reduced by 20%[2].

"In order for Baidu BigSQL* to provide users with high-performance ad hoc query services, large memory is needed to cache hot data locally on compute nodes to avoid DFS I/O slowing performance down. With Intel Optane persistent memory, we managed to ensure outstanding cache performance, while at the same time greatly improving cluster processing and achieving significant TCO benefits."

LI Shiyong, Senior System Engineer, Baidu

Products and Solutions
2nd Gen Intel® Xeon® Scalable processors
Intel® Optane™ persistent memory

Industry
Internet Software, IT Services

Organization Size
10,001+

Country
People's Republic of China

Learn more
Case Study

1,2 For more complete information about performance and benchmark results, visit https://www.intel.com/content/www/us/en/customer-spotlight/stories/baidu-bigsql-customer-story-html

Customer Success Story

Durham University

Durham University is Simulating the Universe with Intel® Xeon® Scalable processors

Durham University's Institute for Computational Cosmology builds universes. The ambitious EAGLE (Evolution and Assembly of GaLaxies and their Environments) project simulated a 300 million light-year section of space in 2015. Building upon the success of the 2015 EAGLE simulation, Durham University is pushing forward with a significant software and hardware upgrade for the next iteration. With rewritten code and a 12,000-core on-premise supercomputing solution based on the latest Intel® Xeon® Scalable processors, the cosmology team will be able to simulate a universe that's 30 times bigger. Durham University is also testing Intel Optane persistent memory for different aspects in their simulations which gives them a chance to create universes that are large enough to compare in detail with the universe we see.

"More processing power will allow for more details. Secondly, more processing power means we can do more simulations with more variations. Instead of only being able to do one experiment a month, we want to be able to do one each week."

Matthieu Schaller, Research Associate, Durham University Institute for Computational Cosmology

Products and Solutions
2nd Gen Intel® Xeon® Scalable processors
Intel® Optane™ persistent memory

Industry
Higher Education

Organization Size
1,001-5,000

Country
United Kingdom

Learn more
Article
Video

1 For more complete information about performance and benchmark results, visit https://www.intel.co.uk/content/www/uk/en/it-management/cloud-analytic-hub/simulating-the-universe.html

Customer Success Story: - крупный телеком оператор Азии

Customer: global technology and business solutions provider delivering fully integrated services, including global networks, cybersecurity, managed IT and applications, cloud and datacenter services.

Challenge: Provider was interested in optimizing their data center and workloads for more scalable database management to support AI and big data analytics. NTT evaluated a solution that included Intel® Optane™ persistent memory in App Direct mode testing PostgreSQL, Apache Kafka and Redis.

Solution: With 2nd Generation Intel® Xeon® and Intel Optane persistent memory, saw results that showed overall better TCO when compared to a DRAM only solution.

Пример - Redis Enterprise

Redis Enterprise Use Case Overview

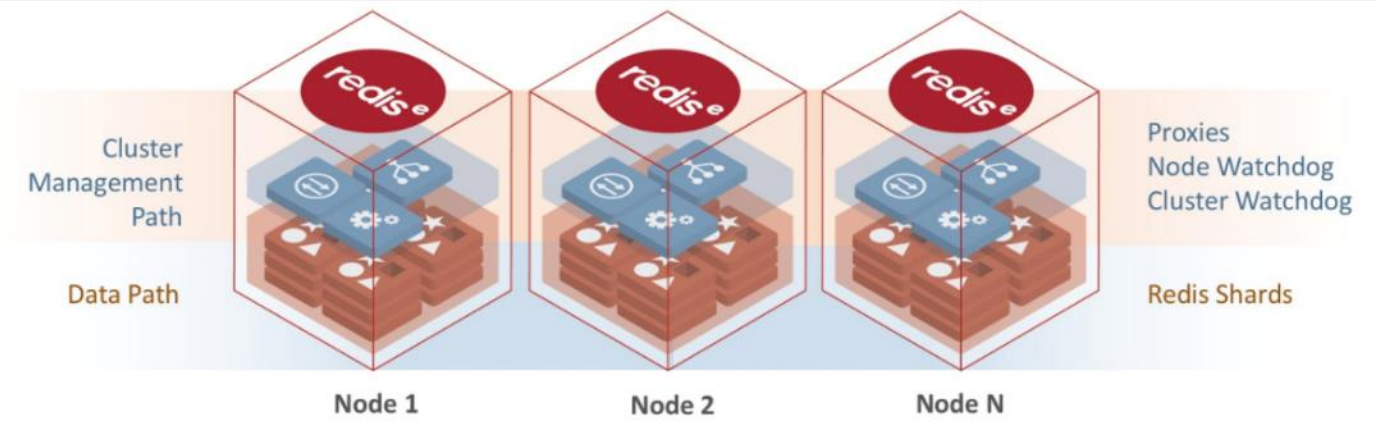

Customer pain points

- Large-capacity DRAM is cost-prohibitive
- Customers respond by splitting databases and adding complexity

Solution

- Add more memory with Intel® Optane™ persistent memory in Memory Mode
- Full dataset can be deployed in a single in-memory database

Value proposition

- **Maintain performance** SLAs at **reduced hardware costs**[1]
- Full dataset can be deployed in a **single in-memory database**
- **Sub-millisecond response** with dramatically reduced cost

[1] Performance results are based on testing or projections as of Jan 15, 2019 and may not reflect all publicly available security updates. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.
Configuration: See Redis Enterprise Consolidation config slide

Redis Enterprise Сравнение

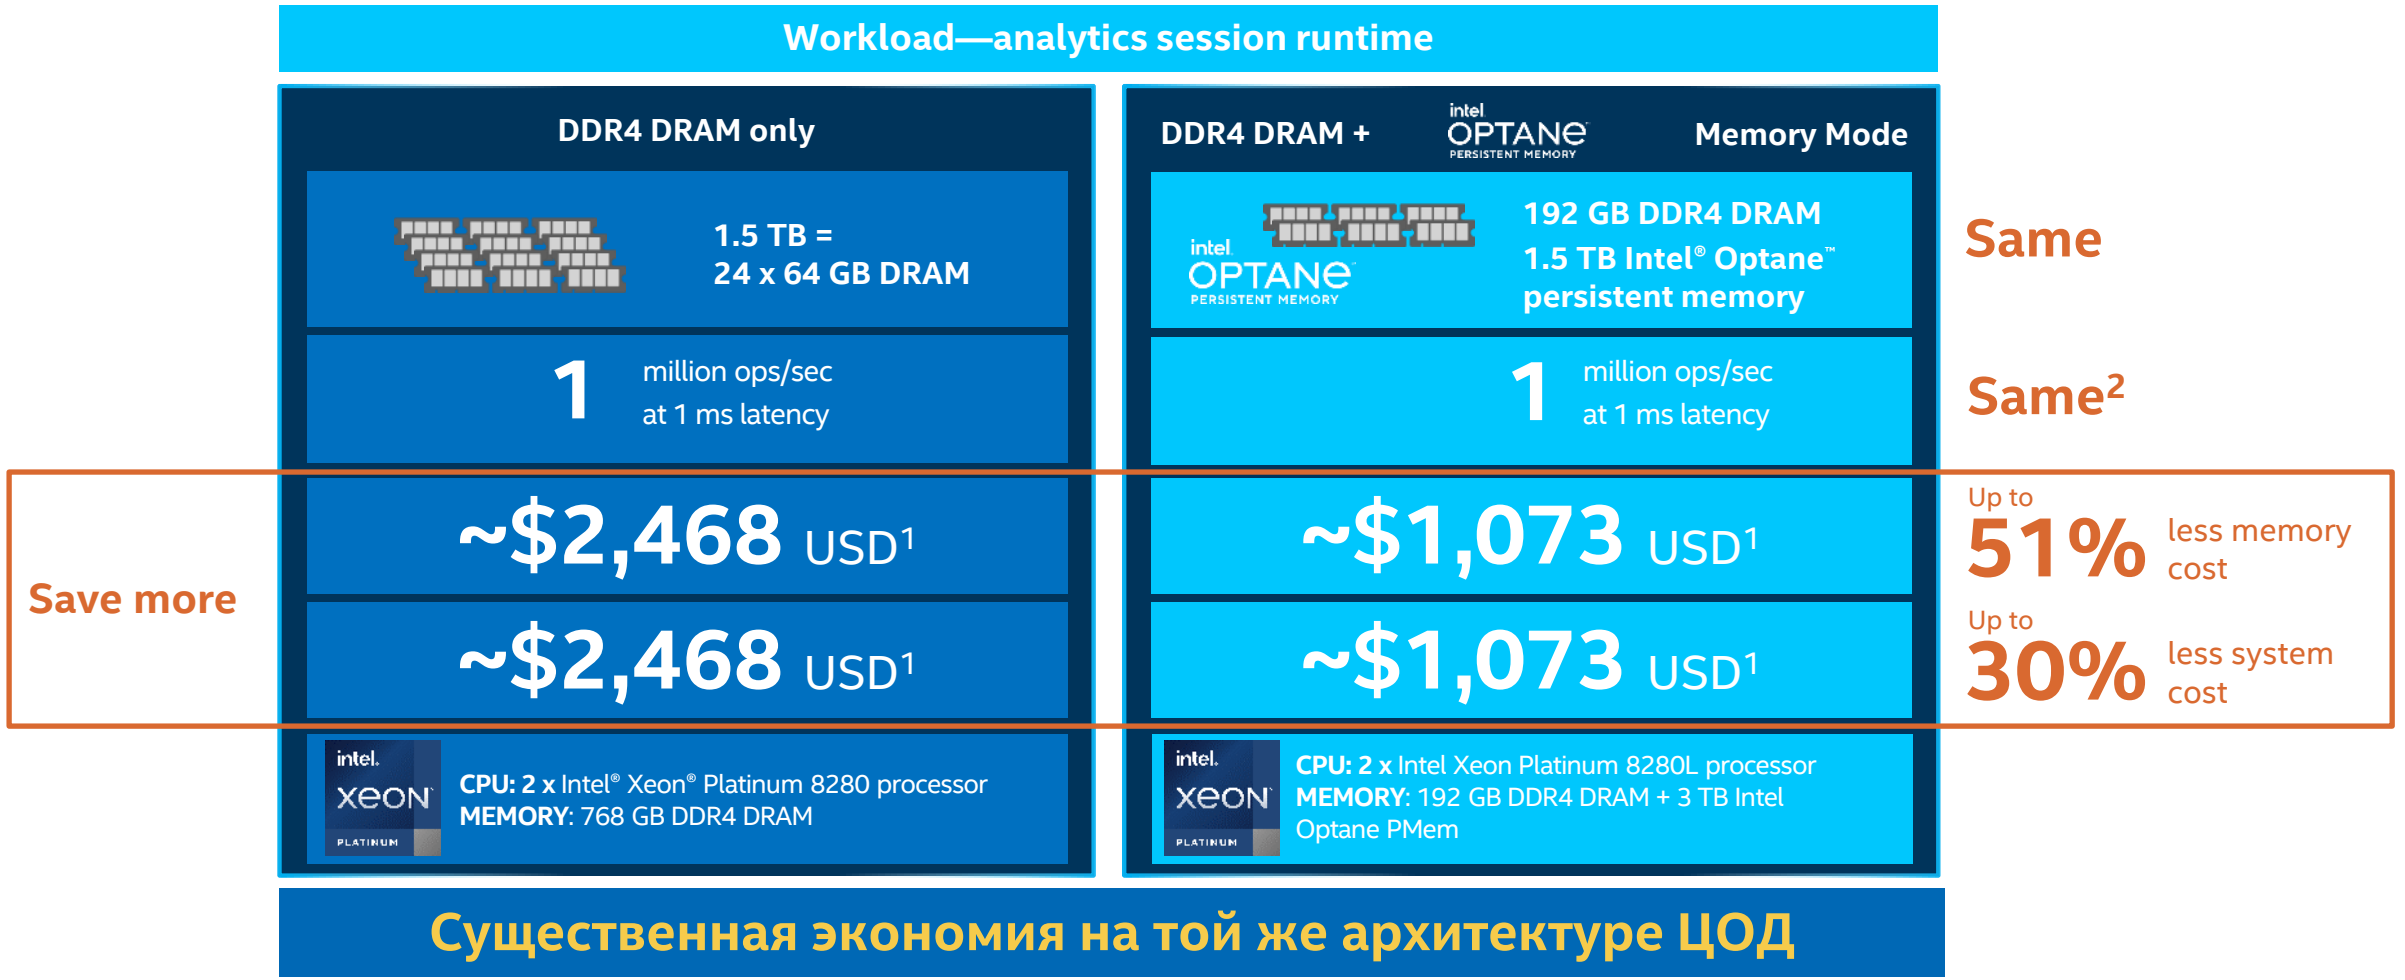

[1] Performance results are based on testing or projections as of January 15, 2019, and may not reflect all publicly available security updates. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.
Configuration: See Redis Enterprise Consolidation config slide
[2] For typical use-cases (1 million ops/sec)

Когда стоит использовать Intel Optane PMEM для Redis ?

Основные факторы: размер БД и производительность

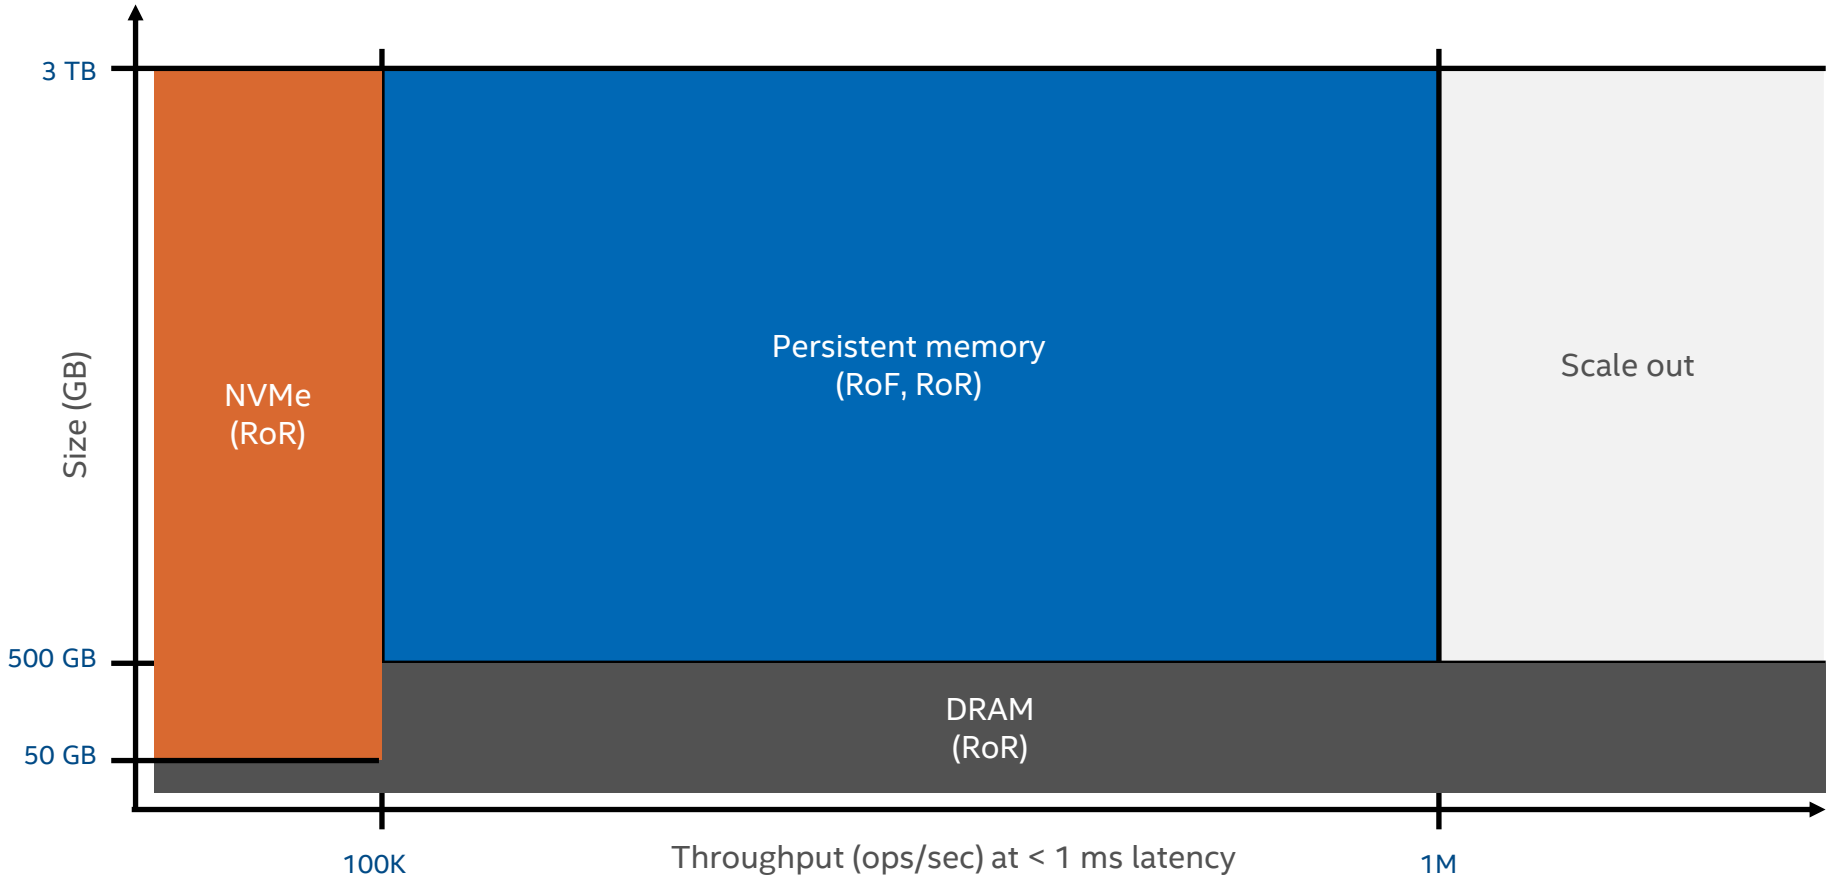

RoR: Redis on RAM (Redis Enterprise)
RoF: Redis on Flash

Предлагаем подготовить аналогичный сайзинг для PostgreSQL

intel OPTANE PERSISTENT MEMORY

Пригодится ли это для ваших проектов ?

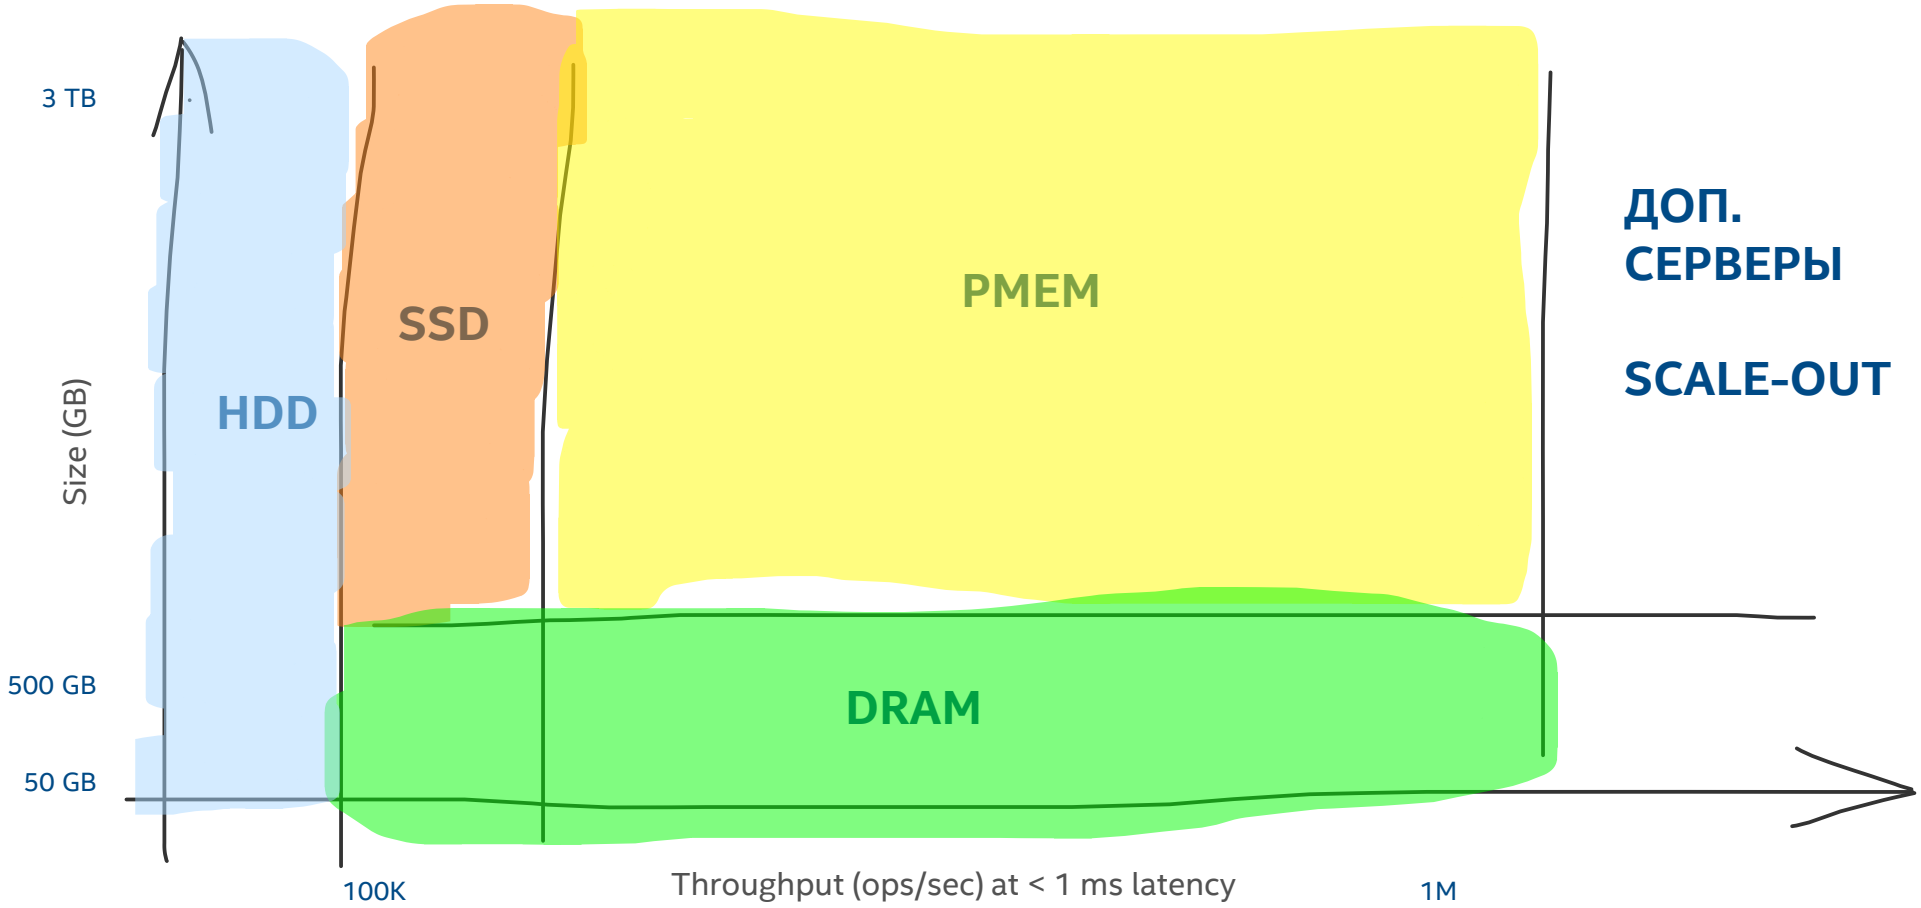

Какие факторы важны для вас ? Размер базы , Скорость работы, задержки ?


intel