Excelero

LEVERAGING DISTRIBUTED NVME FOR DATABASE AS A SERVICE

USE CASE

INTRODUCTION

The world of databases has gone through massive transition over the past decades. Databases used to be an enterprise game, but today they play a crucial role in a wide range of applications, across all industry verticals. Virtualization, cloud and container innovations have had a big impact on database deployment methods: many databases run on hosted servers, leaving the "bare metal" world and leveraging virtualization or containers. As end-users are seeking to leverage cloud economics, Database as a Service (DBaaS) is becoming one of the most popular cloud computing services, both in private and public cloud environments. Providers of DBaaS are deploying infrastructures using a variety of technologies including multi-tenant database systems, SQL and NoSQL technologies, scale-out architectures, shared-nothing and shared-all mechanisms, etc., following the examples set by tech giants like Google and Amazon. Instead of using proprietary, hardware-based storage solutions, they leverage intelligent software and standard servers to achieve their flexibility, efficiency and scalability goals.

„Our original DBaaS infrastructure used NVMe for hot tables and flash for cold tables. Excelero enabled us to increase efficiency to the extent that we were able to switch to a full NVMe infrastructure, which simplifies and optimizes operations even more."

FAST-GROWING MARKET

Database as a Service (DBaaS) refers to the on-demand delivery of database management software to be consumed by end users as a service, without the need to first install any hardware or software, or configure the environment. Services often include database administration, allowing users to focus solely on their applications. DBaaS offers many benefits, including reduced licensing and infrastructure costs, faster time to application development and reduced administration overhead. Market research by The451 Research predicts that the DBaaS market will witness a CAGR of more than 46% between 2015 and 2020, to a valuation of $19bn. Therefore, providers of DBaaS are constantly seeking ways to achieve better efficiency and more flexibility at scale.

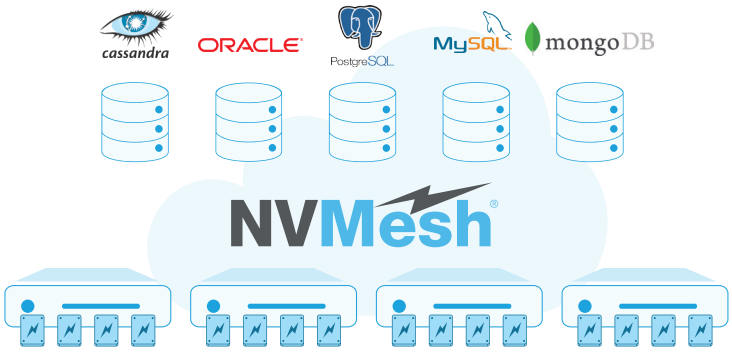

DBAAS REQUIREMENTS: SCALE FLEXIBLY AND EFFICIENTLY

Scalability and Flexibility are the primary requirements for DBaaS, but must not impact efficiency. The purpose of DBaaS is to reduce costs for customers by offering shared infrastructure resources for database services. This can only be achieved by running the underlying hardware as efficiently as possible. There is a common understanding that, for an optimal ROI, you need to achieve 70% hardware utilization: if your utilization is lower than that, your costs will be too high, if it is higher, you risk not being able to handle bursts.

DBaaS infrastructure needs to scale (up and down) in many ways: think of capacity and performance, but also number of instances and users. Compute resources need to be grown, or shrunk on demand. Storage resources need to "move" with the database application for true mobility and flexibility. These requirements put a severe strain on traditional hardware architectures in that to readily achieve both objectives the only perceived solution is to over-provision the compute and local flash storage resources. This of course makes it near impossible to achieve that 70%+ hardware utilization and hence exemplifies the dilemma. To ensure flexible scalability, DBaaS designs need intelligent software that is **completely decoupled** from the underlying hardware. Data volumes need to scale beyond individual disk sizes and even beyond the flash that can fit in an individual node. Volumes should be accessible from any database node to allow migrating a database instance from one server to another without impacting the service while providing operational flexibility. At the same time, performance can't be impacted – access speeds must be fast and consistent. DBaaS infrastructure should be capable of efficiently handling a large number of smaller, rarely used database instances while also enabling other instances to scale both in terms of dataset size and in terms of processing power.

To meet their scalability, flexibility and efficiency goals, DBaaS providers are deploying infrastructures with scale-out, shared-nothing architectures and Software-defined Storage on standard servers. As mentioned above, a specific challenge to achieve high utilization efficiency in such environments is balancing out storage IO, capacity, and CPU power. Especially with standard flash storage, DBaaS providers will typically have an abundance of CPU power but never enough IO. Architectural flexibility and the option to mix converged and disaggregated topologies are essential to maximizing CPU & IO utilization. A software defined storage (SDS) product that can logically disaggregate compute and storage without having to do so physically enables simultaneously achieving the efficiency, flexibility and scalability goals.

BENEFITS AND CHALLENGES OF NVMe

A game changer for DBaaS infrastructures is NVMe flash, which is technically just flash storage, but with an optimized controller and protocol (that supersedes AHCI and SCSI by leveraging PCIe technology). Today's NVMe controllers can handle four times more parallel IO commands than SAS/SATA SSD controllers. NVMe delivers better performance and reduced latency - you need fewer drives to achieve the concurrent performance levels required for your workload. This makes it a lot easier to maintain a balanced CPU/IO ratio. But there is an important caveat: to enjoy the performance benefits of NVMe flash, the storage needs to be used by the application locally, in-server.

The inability to leave the confines of the server is a serious limitation that directly affects all requirements of DBaaS designs. You are limited by the capacity and performance of the NVMe drives within each server and it's virtually impossible to level out utilization across your entire infrastructure. Without a management layer to orchestrate beyond these boundaries, you will waste capacity and performance, e.g. for RAID setups, or use of cumbersome volume management tools that were not designed for modern SDS requirements.

GETTING TO THE NEXT LEVEL WITH DISTRIBUTED NVME: EXCELERO NVMesh

With Excelero's NVMesh, providers of **DBaaS can enjoy all the benefits of NVMe** in a distributed fashion while meeting all their storage requirements. NVMesh is a Software-Defined Storage platform that enables customers to design scale-out block storage infrastructures for the most demanding database applications. NVMesh supports legacy enterprise SQL databases, but also more modern NoSQL and distributed database platforms.

NVMesh is a Software-Defined Block Storage solution that features Elastic NVMe, a distributed block layer that allows any database application to utilize pooled NVMe storage devices across a network at local speeds and latencies. Distributed NVMe storage resources are pooled with the ability to create arbitrary, dynamic block volumes that can be utilized by any host running the NVMesh block client. These virtual volumes can be striped, mirrored or both while enjoying centralized management, monitoring and administration. In short, database applications can enjoy the latency, throughput and IOPS of a local NVMe device while at the same time getting the benefits of centralized, redundant storage. NVMesh provides the ability to attach volumes ubiquitously,

enabling users to mount databases on any server at any time. NVMesh is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures, giving customers full freedom in their architectural design.

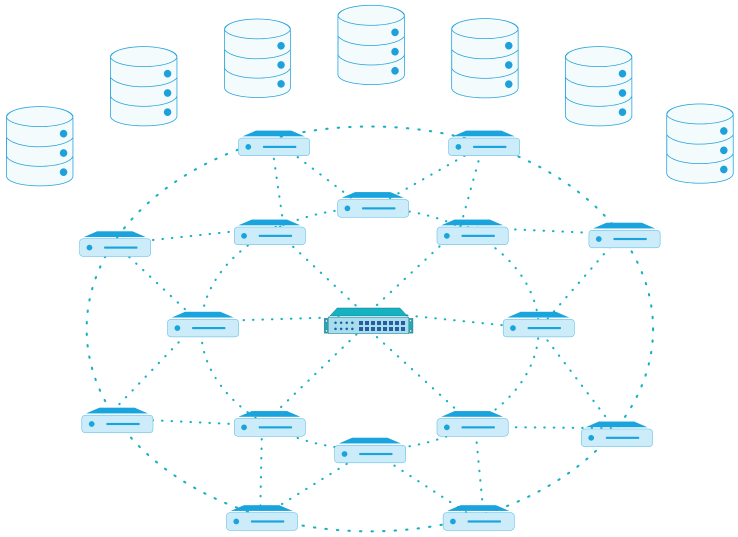

NVMesh BENEFITS FOR DBAAS

SCALE & PERFORMANCE

- Leverage the full performance of your NVMe flash at any scale, over the network
- Predictable database performance – ensure that storage is not a database bottleneck
- Scale your performance and capacity linearly
- Leverage high IOPS, high bandwidth or mixed

EFFICIENCY

- Maximize the utilization of your NVMe flash devices
- Choose hardware from any server, storage and network vendor
- Easy to manage & monitor, reduces the maintenance TCO
- Balance CPU and storage resources

FLEXIBILITY

- Scale your databases up and down
- Move databases seamlessly from server to server
- Run RDBMS or NoSQL solutions within containers or virtual machines to conserve compute resources when they are not used without sacrificing performance
- Choice of architecture: converged, disaggregated or mixed
- Mix different storage media types to optimize for cost, scale or performance
- Scale storage and compute separately, as needed