Database Workloads

Faster seek time and increased storage capabilities to run any database workload

Databases are a key tool inside any IT infrastructure and are considered one of the most important criteria for server and storage infrastructure selection. These serve to centralize where data is recorded, tracked and updated. Data availability is essential for any organization, from departmental databases to critical workloads. Having an effective database management system in place will guarantee that operations run smoothly and properly inside any company.

Challenge

Database seek time - the time it takes a program or device to find a specific piece of data has often been reported as one of the main challenges faced when deploying and managing a database. The higher the seek time, the slower a database returns a piece of data, which in-turn affects how quickly applications can process the data and return on the requested operation to the end-user. Disk subsystems are a typical bottleneck and often the culprit behind sluggish database response times.

Performance degradation of seek time can often be caused by legacy disk media, such as overloaded shared storage platforms or local spinning disks in a server.

Cost and Performance

The performance requirements of a database platform are often the most challenging to meet and also the most expensive. In traditional setups where shared storage (even an All Flash Arrays) are leveraged for this highly available and performant platform, costs associated with the implementation are prohibitive.

In addition to the cost of implementing such an infrastructure other shortcomings typically faced by traditional database architectures are:

Scalability

Handling performance and data growth can be hard with a traditional approach where individual compute, memory and storage planning need to be defined upfront and considered for the life of the deployment. Deploying a new SAN for more storage performance and availability can be very costly exercise.

Management

Managing a distributed architecture can be challenging, requiring technical specialists from a wide range of backgrounds – virtualization, storage, networking. In addition, it's complex to integrate into monitoring and management platforms and DR has to be architected separately. All but the largest of enterprises cannot afford to duplicate their primary architecture into a DR environment.

Why Hive Fabric™ for Database Workloads

Running databases on Hive Fabric provides better utilization of hardware resources and simplified management through a single platform running all workloads.

Hive Fabric has a number of technical benefits that provide database workloads with the performance and flexibility they required. A unique acceleration layer in front of the disk subsystem is just the start.

Hive Fabric also has the ability to add a RAM based acceleration layer on top of the physical storage which will allow up to 90% of reads to come from RAM. As a result, seek time is accelerated regardless of the type of media being used to store at-rest data. This acceleration layer also speeds up deduplication, which increases disk capacity as redundant copies of data are eliminated. As the data is written to disk it is also compressed further increasing the capacity available from the platform.

This results in a performance transformation by simply using RAM on top of the spinning media to process database workloads. The cost of delivering this is significantly cheaper than a traditional architecture.

Additional Hive Fabric advantages:

Hive Fabric offers a disruptive and comprehensive solution for server virtualization. This includes shared storage (hyperconverged) that transforms the traditional 3 tier approach associated with database workloads into a hyperconverged architecture, simplifying the infrastructure supporting the database.

Scalability: Hive Fabric allows compute, memory and storage to scale out as needed. More instances of the database can be added in the same platform and scaled out as data grows or more performance is required. More compute memory or storage can be provided by simply adding another server to the cluster.

Availability: Hive Fabric allows a database workload to be live migrated between hosts. This free movement of workload results in better infrastructure usage since the database is unlikely to need the same CPU and memory resources all the time. Hive Fabric's Cluster Resource Scheduling allows for automatic movement of database workloads through the cluster without input from the admin to optimize the resources available for the database.

Simplified Management: Hive Fabric brings more flexibility and easier management to the IT infrastructure with its unique Message Bus architecture. Hive Fabrics' Zero-Layer management simplifies management operations to the point of requiring far less time from expensive resources allowing IT to focus on delivering more for the business.

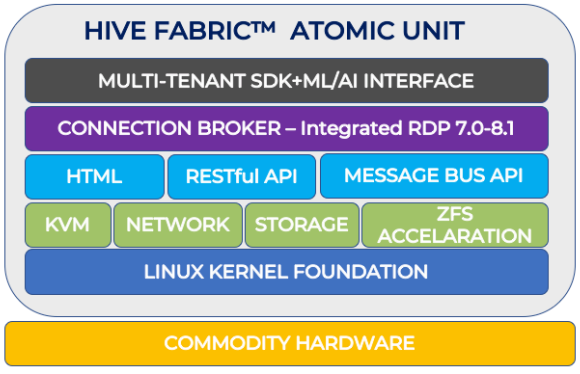

When considering how to best utilize the resources you have in your datacenter, add additional capacity to run more workload, or even entirely transform how IT delivers on business requirements, question the level of complexity that you have in your datacenter and whether you are increasing or decreasing the overhead on your IT team with the changes that are being made. It's important that your datacenter works with you, not against you, leverages your resources as you scale, and is capable of intelligently powering outcomes for the business. Hive Fabric can remove this complexity and cost, install on your existing infrastructure, and ease the transformation to a software-defined datacenter.

Getting started with Hive Fabric is easy. We do not ask for a support and maintenance agreement, this is included as part of the license, no need for multi-year renewals pay as you go and your budget allows. We do not require complex hardware configurations to get you to the starting line.

Why HiveIO

Complexity can be interesting. It can add richness; it can inject ambiguity and variability. It can lead to redundancy. It can be short-lived. Simple is Powerful.

At HiveIO we believe that the real journey should begin immediately, and that the promise can be realized. We do not want you to go from the data center team, to the networking team, to the storage team, to the virtualization team, to the Windows or Unix team. We want you to go to a single team that can see you through your journey.

Realize the promise.

www.hiveio.com | info@hiveio.com | @HiveioInc
T. +1 415 340 2089 | F. +1 415 715 9028