NetApp MAX Data

PRODUCT BRIEF

The Architect's View

MAX Data is a software solution from NetApp, Inc that addresses the challenge of delivering single-digit microsecond latency to applications. The technology works with persistent media such as storage-class memory and persistent memory to create a local file system within a server that is backed by external NetApp storage. MAX Data can be used to accelerate existing workloads by making use of SCM/PM without any application rewrite.

Background

As the slowest component in traditional computer architectures, persistent storage is always looking at ways to become faster. The difference in time between a single I/O and accessing data in memory is effectively eons if we scaled up memory access time to the equivalent of a single second.

However, we need persistent storage for a number of reasons. There are limits to accessibility for main memory, which isn't persistent and isn't scalable. Cost/benefit analysis shows that it isn't practical to store all of our data in memory all of the time. Add in the issues of resiliency and availability and we get a balancing act where data needs to be as close to the processor as possible, while still being persistently retained on the most cost-effective media.

Solutions already exist to cache data or even tier within shared storage appliances. These have worked well, but still have significant latency challenges compared to addressing media that's either on the server memory or I/O bus. New technologies like Storage-Class Memory offer super-low latencies with high throughput and can be byte, rather than block addressable.

Devices can exist on the PCI Express bus within a server or even emulate DRAM and sit in a memory slot. The hardware technology exists to get persistent I/O into the single-digit microseconds. We just need the software to exploit it.

Latency is King

With persistent media latencies reaching single-digit microseconds, accessing storage from a centralised appliance is starting to show weakness in terms of performance. Centralisation provided the benefits we've seen in storage area networks for the past 20 years, including greater resiliency, efficiency and reduced administrative costs. But for many applications, the latency penalty introduced by storage networking is simply too high a price to pay.

Persistent storage in the server is clearly the answer, but there are challenges. The storage hierarchy in operating systems like Linux means I/O time is wasted with many levels of indirection. Having persistent media within the host without any means of protecting against failure represents a return to simple DAS design. This isn't practical or efficient.

The challenge then, is to implement a scalable storage layer within the server that can exploit local media without the compromises of DAS.

What is MAX Data?

Memory Accelerated or MAX Data is a new software product from NetApp, Inc. The solution uses persistent memory technology such as NVDIMMs and Intel Optane to create a file system that runs locally within a server. I/O is serviced from the MAX Data tier, providing both high performance and low latency. MAX Data has been

developed from the intellectual property acquired by the purchase of Plexistor in 2017.

Ordinarily, accessing data locally within a server is simply directly-attached storage, or DAS. However, MAX Data operates as an extension to NetApp ONTAP storage appliances and can tier data between local storage and an ONTAP system.

In order to protect against hardware failure, data is both backed off to the ONTAP appliance and can be synchronously replicated to another fall-back server. If the primary server has a hardware issue, then data can be recovered from the fall-back once the primary server has been fixed.

MAX Data is not a caching solution but a tier of storage where data only exists either in the server (tier 0) or on the backing storage appliance (tier 1). This is an important distinction because caching can introduce latency time variability if the working set isn't fully in memory or if the data itself is relatively random (and so unpredictable).

Caveats

The MAX Data solution provides I/O latency quoted as low as 3 microseconds for workloads including NoSQL databases. The benefits extend past just I/O performance and can include more efficient server utilisation and lower licensing costs.

However, the trade-off in moving away from shared or SAN-type storage designs is in the increased complexity MAX Data can introduce. Resiliency, for example, is achieved through an additional local server that acts as nothing more than a failover cache.

Another consideration is the workload mix that MAX Data is suited for. Customers that need HA or high levels of resiliency (generally achieved with synchronous replication) will find MAX Data does not suit their requirements. Instead, the solution is more appropriate for read-intensive applications or those where asynchronous or periodic replication of data via snapshots can provide sufficient recovery capability.

Reference Information

Further details on MAX Data can be found in the following Architecting IT blog posts:

- MAX Data: Optimising Application Performance with SCM (12 February 2019)
- Accelerating Workloads with NetApp Plexistor (20 November 2017)
- What are Storage Class Memory (SCM) and Persistent Memory (PM) (6 February 2018)
- NetApp, Storage Class Memory and Hyperconvergence (27 September 2016)

Further details on NetApp MAX Data can be found with the following Storage Unpacked podcasts:

- #83 – Introduction to NetApp MAX Data (Published 18 January 2019)

We recommend the following Tech Field Day videos that provide additional background on MAX Data and Plexistor:

- NetApp MAX Data Presentation at Tech Field Day Extra at NetApp Insight 2018
- Plexistor Presents at Storage Field Day 9
- Plexistor Presents at Tech Field Day 11