**UNICOM Engineering White Paper - How Intel's 3D Xpoint, Omnipath, and NVMe are Revolutionizing Tiered Storage**

**Big Data's Storage Woes**

We are flooded with data, and the levels are rising fast. Only 5 years ago, 90 percent of all of the digital data ever created had come into being within the previous two years.[1] Using its own research, IBM states that 2.5 quintillion bytes of data are created every single day. [2] Clearly, there is a growing data problem affecting enterprise and storage, and it's just going to get larger.

Increasing expectations from customers require more data processing, while at the same time, data chunks are growing in size. Drivers of this trend include burgeoning social media use, increasing use of Internet-connected apps, ballooning data-capture volumes for big data enterprises, and an influx of data from The Internet of Things (IoT) sensors. All have caused data usage, and storage needs to proliferate exponentially, compounding the problem. And the progression is only moving upwards, with no signs of leveling off.

To make things even more complex, storage performance is growing linearly, while data and data demands are growing exponentially. There is a light at the end of the tunnel, however, with new storage tech that will help engineers solve these pressing issues by advancing storage performance dramatically to keep pace with demand.

It's not just about increasing storage, either. Fortunately, technologies have been developed that not only store more data but reduce latency and increase the amount of data that can be processed in a period of time. This allows for more efficient use of processor cores and allows for more processing actions to be taken on the data. In other words, this new tech essentially has both increased hard drive capacity along with better memory performance, allowing the solution to do more, faster.

So, what does all this mean for the end user? Technologies such as Intel's 3D XPoint can process big chunks of data in a shorter time frame, which can mean immediate responses on even large-volume data requests. This, in turn, encourages greater development and usage of products which can take full advantage of and do amazing things with the sea of data at our fingertips.

New non-volatile memory technologies such as Intel's 3D XPoint anticipate growing needs in the form of ever-evolving storage media. These ultra-dense and economical storage mediums with ultra-low latency protocol options give enterprises more choices than ever. Whether they are choosing traditional HDDs handling bulk storage functions to cutting-edge, low-latency solutions spooling up live data for real-time analytics, enterprise computing operations will be able to wield precise control over how the rising tide of data is both managed and stored.

**Non-Volatile Storage Media**

As with most things in the technology world, capacity and performance tend to get better while materials get cheaper. In the case of Solid State Drives (SSDs), multi-level cells have given rise to increased density, consistently lowering the cost per gigabyte month after month. Improved wear levelling mixed with breakthroughs in caching, backups, and speed-enhancing controllers have only

[1] https://www.sciencedaily.com/releases/2013/05/130522085217.htm

[2] http://www-01.ibm.com/software/data/bigdata/what-is-big-data.html

recently allowed SSDs to outpace the abilities of (Hard Disk Drives) HDDs—although the latter still maintain a preferable cost per gigabyte when looking at slower data throughput archiving scenarios. In general, although HDDs will most likely maintain their function as high-capacity, archived storage media, at the current pace of development, SSDs are on the cusp of being both viable and affordable on a much larger scale for enterprise computing entities.

**Tiering Storage Solutions with SSDs and HDDs**

Even though they have a comparatively high price-point, the read/write speed of SSDs makes them extremely desirable in systems that need rapid data throughput with low latency. To control costs while leveraging the performance benefits of SSDs, many enterprise computing companies use the following storage tiering architecture:

- High-performance SSDs
    - Immediate data intake and throughput
    - Temporary storage of data for real-time or low-latency operations
- Mid-Range SSDs/High-Performance HDDs
    - Caching
    - Extended storage
- Mid-Range or Low-End HDDs
    - Long-term, archival storage

Across the enterprise computing world, Solid State Drives (SSDs) have shown much promise when it comes to the performance point/component price tradeoff. And they have become cheaper over time, with today's offerings much less expensive than the first commercially-available SSDs, which appeared back in 1974 and cost $1.8 million per gigabyte. [3] However, their expense and relatively low storage capacity have slowed their wide adoption across industries bottom line: current SSDs have certainly helped improve performance, but they aren't cheap or fast enough. This shortcoming has left an opportunity wide open for further development and innovation along different lines.

**Building the Future of Storage: NVMe, 3D Xpoint, Omnipath and More**

One of the issues with the tier storage arrangement outlined above is that investing in advanced SSDs can have diminishing returns. While NAND-based flash memory has tremendous potential over physical media like hard disks, until recently those capabilities were restrained by the adherence to older protocols. In other words, treating SSDs like HDDs has underutilized their full potential and capabilities, undermining their true advantages; however, three emerging technologies are set to change that.

**NVMe**

NVMe (also called Non-Volatile Memory Host Controller Interface Specification, or NVMHCI) replaces interface protocols like SATA and SAS, which were designed for HDDs. NVMe allows SSDs to plug directly into PCIe lanes, boosting data speeds considerably.

- SATA ≈ 500 Mbps
- SAS ≈ 1.5 Gbps
- NVMe ≈ 1 Gbps per PCIe Gen 3 lane (4Gbps in four-lane slots)

[3] http://www.pcworld.com/article/246617/storage/evolution-of-the-solid-state-drive.html#slide2

The dramatic increase in speed available via NVMe will revolutionize tiered storage architecture as extremely low-latency data access becomes possible. [4]

**3D XPoint**

3D XPoint, a non-volatile memory (NVM) technology produced by Intel and Micron, promises to rewrite what SSDs are capable of and offer unheard-of performance. Write durability is allegedly 1,000 times higher than NAND-based flash memory, and latency promises to be around 10 times lower than a NAND SSD over NVMe.

3D XPoint is also bit-addressable, allowing it to function similarly to DRAM in addition to enabling more efficient wear levelling. Key specs include (Input/Output per Second) IOPS of 5x 78,000 and a latency of 7µs. In 2015, Intel announced its Optane brand of storage products designed to take advantage of 3D XPoint technology. New chipsets such as Intel's Xeon Scalable Processor Family, released in mid-2017, also work seamlessly with Optane and 3D XPoint technology. [5]

**Omnipath Architecture**

Intel Omnipath Architecture (OPA) replaces True Scale as Intel's fabric-based computing system and is included with Xeon Scalable Processor Family, Intel's just-released storage server chipset. Like the InfiniBand EDR, Omnipath fabric can transfer data at a rate of 100 Gbps; however, it also promises more user-friendly features in its ability to scale, customize, and adjust to complex architecture. Its features include:

- Traffic flow optimization — Enables prioritization and urgency tiering within data
- Dispersive routing — Maximizes efficiency with simultaneous packet distribution
- Adaptive routing — Picks optimal pathing for packet distribution
- Packet integrity protection — Allows for rapid recovery of transmission errors
- Dynamic lane scaling — Prevents process disruptions if a single lane drops.[6]

**A Revolution in Tiered Storage Architecture**

A true revolution in tiered storage is at hand. With blindingly-fast I/O speeds and minimal latency now available, some computing applications can finally deliver on their promise of real-time results. Many data storage and enterprise computing applications are seeing marked improvement in their performance as they wean their SSD storage components away from SATA and start utilizing NVMe instead.

These new storage and data transmission capabilities mean that enterprise computing companies have more price points and performance considerations than ever before. SSD capabilities can all be turbocharged through NVMe connections. 3D XPoint memory can replace flash-based memory or even DRAM. Intel Omnipath fabric makes connections faster, more stable, and more customizable than ever before.

In light of these sweeping changes, what does the future hold for tiered storage architecture? Most likely, a rise in highly environment-specific use applications, with many options available for designers to choose from for their specific needs. For example, the rise of the Internet of Things (IoT) means that the

---

[4] http://www.computerweekly.com/feature/Storage-briefing-NVMe-vs-SATA-and-SAS

[5] https://www.krollontrack.co.uk/blog/the-world-of-data/intel-optane-memory-what-you-need-to-know/

[6]http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-architecture-fabric-overview.html

increasing numbers of device applications will need to access various data pools over multiple connections at extremely high volumes with minimal latency. Storage gateways will increasingly manage the task of making information accessible to the cloud and the Edge, while customizable solutions like Dell Networking switches and Compellent can optimize automation to fit the needs of the specific enterprise.[7] Controllers may also be needed to allow memory technology like 3D XPoint to switch dynamically from DRAM-type functionality to SSD-storage and back on the fly. All of these demands will place increasing complexity and responsibility on independent software vendors and other technology architecture designers attempting to develop solutions for a broad market of enterprise clients.

Adjusting to this new environment involves not just creating and planning for new types of storage tiers, but also how each of those storage tiers' interfaces and uses may alter based on each use case.

**A Hypothetical Example of a Tiered Storage Architecture Configuration**

Fictional intrusion detection company "Ft. Noxx" uses a sophisticated pattern recognition algorithm along with thousands of ID metrics to fingerprint each incoming access request for a client's system. This access data must be rapidly compared to "Ft. Noxx's" home server data to identify access instances that have a high probability of being a threat. Threats are therefore detected in near-instant real time.
To accomplish this task, "Ft. Noxx" needs to maintain:

- "Fresh" data on NVMe, 3D XPoint and Omnipath systems. These systems can rapidly capture, process, and transmit raw data as it arrives for near-instantaneous response.
- "Cool" Data on high-capacity SDDs. Standard NAND SDs can handle this priority of data for a competitive dollar per GB price point, storing relatively recent data from the past few days or weeks in order to detect patterns of access requests.
- "Cold" historical data on large capacity HDDs. High-volume, low dollar per GB hard disks can handle archival duties through their economic and cost-effective price points and extended read/write lifespans. These systems can also be accessed periodically to study previous patterns and draw comparisons between the latest models as well as past successful models.

"Ft. Noxx" may also need their clients to deploy a similar architecture in order to achieve low-latency data relay between the field-based appliance and the company's own proprietary data center architecture.

**New Challenges Designers Face in the Increasingly Complex Tiered Storage Field**

With the emergence of tech protocols such as NVMe, SSDs suddenly present a drastic night-and-day comparison to HDDs. Moving forward, "fresh" data will be completely dependent on SSD systems with HDDs gradually being relegated to a legacy-based archiving system where low cost-per-terabyte is valued above metrics like latency levels or IOPS.

Storage controller applications will need to respond in kind, knowing exactly when, why and how to bump data to a lower tier or transfer it to a more readily-available one through an automated process. Any kinks in this process could mean more than inefficienc;, they can mean many hours of manual coding or data-dumping in order to rectify the disorganization.

---

[7]http://en.community.dell.com/dell-blogs/direct2dell/b/direct2dell/archive/2016/05/17/the-tiered-storage-revolution

All of this complexity may be completely new territory for many software and hardware designers and architects. Security solutions providers, for instance, will need to know which processes to delegate to low-latency NVMe systems and when to dump “stale” data into archiving, mass-storage HDDs. In the midst of the complexity will be a host of issues that will likely require an integration partner, in order to cut through the confusion, optimize solutions for various price points, and bring products to market quicker while anticipating and responding to setbacks and challenges.

Possible challenges could include:

- Keeping up with all the advancing interfaces while enabling the appliance product to interact uninhibitedly with elements like storage gateways and controllers
- Enabling quick, painless integration of the appliance into a network environment
- Optimization based on price points, and client storage preferences
- The general research and development-intensive process of creating solutions that are compatible and in compliance with current expected standards

Additional challenges faced by enterprise computing developers include logistics, shipping, long-distance support, warranty replacements, and a host of other complexities that only add risk as operations scale up.

**UNICOM Engineering Makes Navigating New Storage Solutions Easier**

Overcoming challenges on a product’s way to market necessitates having hardware and compliance expertise in your corner. Rather than investing in employees or training that lies outside their core competency, solutions developers can look to Global Value Integrator partners like UNICOM Engineering.

We make it our business to understand the latest standards, protocols, and technologies in order to anticipate new emerging challenges to integration. We then work tirelessly to discover new standards or workarounds that allow for environment-specific appliance builds perfectly suited to a client’s demands.

Capabilities UNICOM Engineering offer include:

- Revealing solutions for low-cost ways to incorporate the latest technology protocols or standards within your appliance build
- Optimizing appliance builds to specific verticals or use cases, including accounting for:
    - Heat output
    - Power consumption
    - Mechanical and environmental stressors
- Applying best practices to enable seamless integration within intended environment
- Sourcing hardware components from trustworthy OEMs who partner with Dell and Intel
- Build-testing appliance prototypes to create consistent benchmarks
- Developing a range of products based on our clients’ performance or price demands
- Logistics, global shipping, support and warranty services
- Lifecycle management to help your product stay current, then prepare for eventual replacement
- Navigating complex compliance demands based on industry standards and government regulations

In short, a partner like UNICOM Engineering can reduce your risk, help you achieve required efficiencies, ensure needed oversight, and enable you to scale globally as your business grows.

Don’t let the increasingly nuanced field of tiered data storage minimize your ability to develop amazing solutions that could impact the market. With UNICOM Engineering at your side, utilize these data storage technologies for success.

The demand for data is only increasing, and the need for solutions for storing and processing this data is creating many new challenges. New hardware solutions make finding results far easier, but only for enterprise computing companies that develop sophisticated enough architectures to stratify tiered storage standards and develop protocols for controlled and automated data transfer.

Solutions developers caught up in the mix are likely to face more challenges than ever before when deciding how to approach storage protocols as well as their products’ own onboard storage needs.

Having a partner like UNICOM Engineering in your corner can provide answers to your challenges while reducing risk and enabling your operation to scale rapidly. UNICOM Engineering is a leading provider of server-based application platforms and lifecycle support services for software developers and OEMs worldwide. Through its expertise and comprehensive suite of solution design, system integration, global logistics, trade compliance, support, and business analytics services, UNICOM Engineering is redefining application deployment solutions to provide customers with a sustainable competitive advantage. More than a decade of appliance innovation and strong technology partnerships make UNICOM Engineering one of the most trusted and capable software deployment partners in the industry.

Founded in 1997, UNICOM Engineering has facilities in Canton, Massachusetts; Plano, Texas; and Galway, Ireland. For more information, visit www.unicomengineering.com.

**Contact Us**

For more information on our platforms and deployment services, please contact us by telephone at (800) 977-1010 or by email at info@unicomengineering.com.