Dell EMC's 2020 Server Trends & Observations

By Robert W Hormuth, Technology Fellow, Infrastructure Solutions Group Office of the CTO
& the Senior Technical Community in Server and Infrastructure Systems at Dell EMC

While many of our 2019 Server Trends and Observations came to fruition, and some are still ongoing, our technical leadership team has collaborated to bring you the top twelve trends and observations that will most greatly affect server technologies and adoption for 2020.

As the global leader in server technology, Dell EMC has attracted some of the brightest minds in the industry. Sharing a small glimpse into our mind trust – with deep roots in listening to our customers and leaders around the industry – each of these twelve trends and observations is authored by one of our Senior Fellows, Fellows, or Senior Distinguished Engineers.

Contents

Data is King

Jimmy Pike – Senior Fellow & SVP, Server Infrastructure Solutions office of CTO
Shawn Dube – Senior Distinguished Engineer, Server Infrastructure Solutions, Advanced Engineering

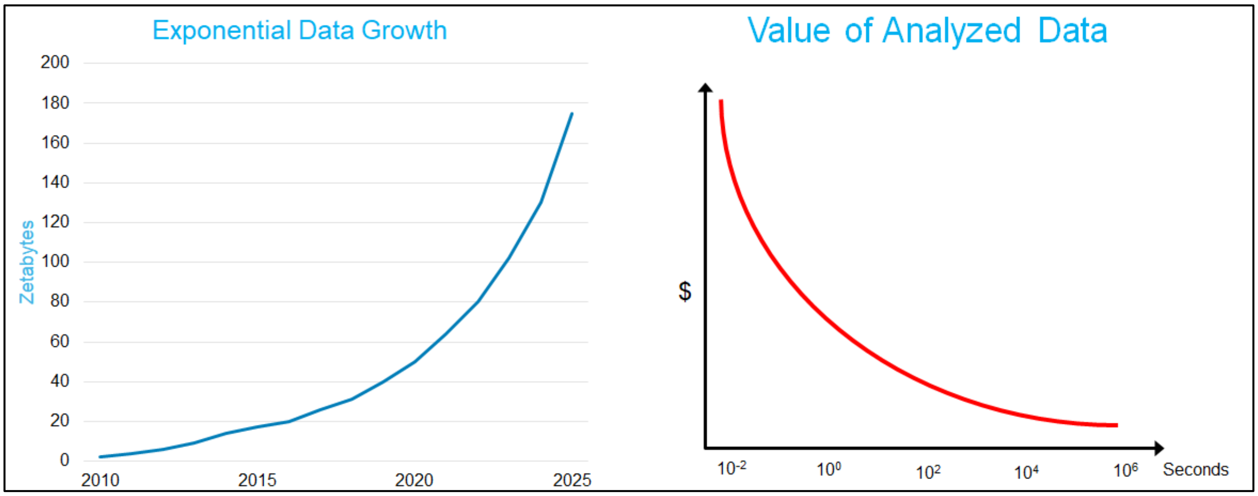

Data is the key to the future as data-driven insights are changing how business is run and presenting new challenges at all points from the cloud to the core to the edge. Perhaps the greatest challenge is from the 50 billion intelligent devices and the deluge of edge data they are generating. In the modern information technology world, among the essential practices are data (and metadata) management and governance. As business goals and technology capabilities converge, more business users will find themselves in a new era of problem modeling using data which must be validated and qualified, transformed and normalized. Existing data practices are in part automated and in part manual. With the sheer increase in volume both management and governance are an increasingly important part of data science and offer more power and control to both citizen data scientists and expert business analysts.

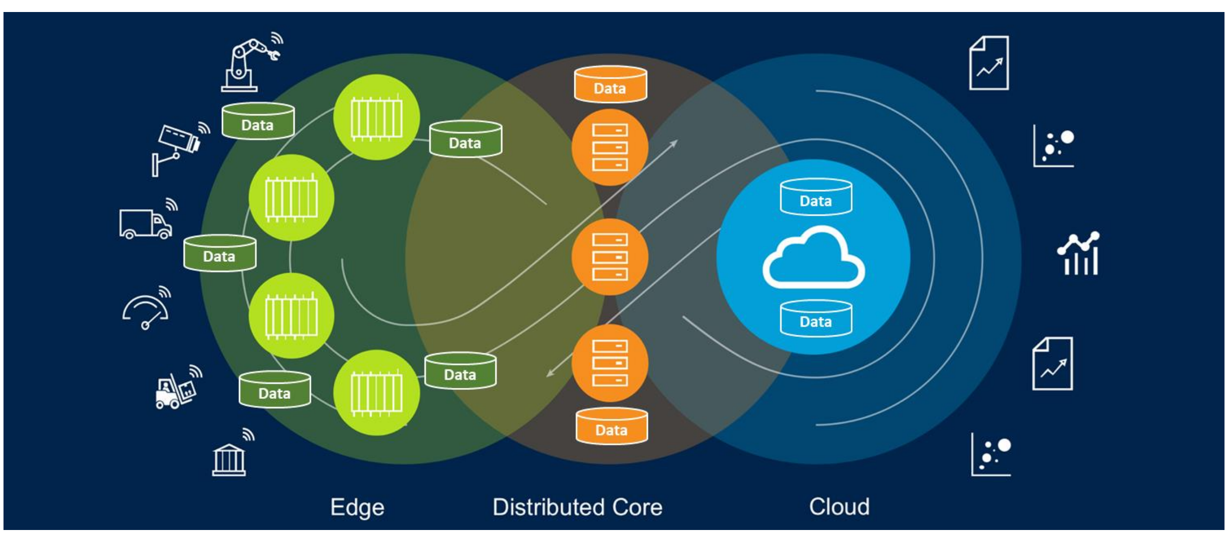

Enter machine learning (ML), deep learning (DL), and artificial intelligence (AI), all forms of data analytics that are used interchangeably but have distinct meanings. ML and DL are primarily patterned recognition techniques where ML is algorithmic in nature, while DL requires extensive sets of training data. AI represents actions or responses that are taught (or trained) in response to the pattern detection mentioned above. These are no longer futuristic concepts, but rather practices that are here today and being integrated with and deployed into a variety of business. One can easily see the importance of data in both the preparation and execution of each of these disciplines. In fact, the majority of most project time is consumed in the preparation and cleansing for the data, in a process generally referred to as data curation. The impact of data curation is especially keen as inaccurate or incorrect data can result in erroneous conclusions from AI engines.

The first wave of digital transformation has increased reach and customization, improved processes, and boosted productivity while also systematically incorporating big data into business processes. As data increasingly drives business, manual processes aren't sufficient to power the next wave of evolution. The need to be customer-driven and the requirement to deliver relevant services and products in a timely fashion is taking data science to new levels. Using artificial intelligence—driven by data in an automated manner—is the only way to tackle emerging problems. Automation is required to accommodate the vast quantity of data that is continuously changing and evolving in context. These challenges mandate automatic data management and governance as new use cases drive capabilities and adoption.

Modern business requires that data engines exist in an elastic, scalable, and resilient manner, including the ability to interoperate with existing data sources such as data lakes, distributed data sources, and now real-time data which can't be duplicated. Finally, from an architectural perspective, the use of data in the public cloud, conventional enterprise, on-premise platforms, and now at the edge must abstract the infrastructure with ecosystems that will autonomously manage data pipelines, independent of where the data resides. While data science and advanced analytics continue to make up the broad field of algorithms and statistical routines, the integration of data-driven ML/DL/AI, including data from the edge, will be the focal point providing significant business impact as new complex businesses evolve. Data is King! He who masters the data will rule the world.

All Hail Domain-Specific Architectures

Onur Celebioglu - Senior Distinguished Engineer, Server Infrastructure Solutions, HPC
Bhyrav Mutnury - Senior Distinguished Engineer, Server Infrastructure Solution, Advanced Engineering

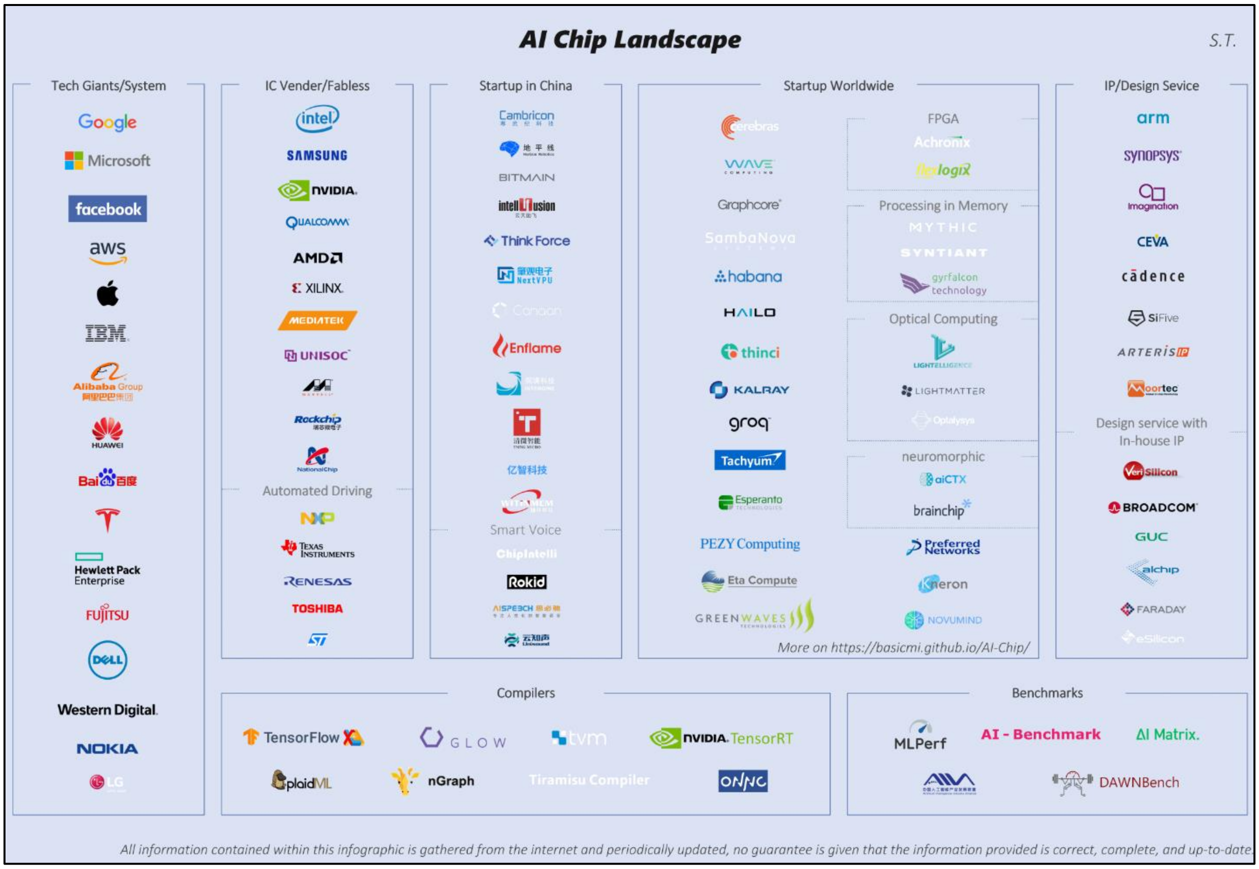

Artificial Intelligence (AI)/Machine Learning (ML) and Deep Learning (DL) usages are becoming more and more pervasive with each passing year. With Internet of Things (IoT) to smart homes to autonomous cars, the explosion of AI/ML usage is here.

Not long ago, the CPU compute was based on RISC or X86 architectures (IBM, Intel and AMD). In the last decade, we have moved to GPUs and FPGAs (Nvidia, AMD, Intel, Xilinx, etc.) to propel AI/ML. In the last few years, tensor processing units (TPUs) and neural processing units (NPUs) are gathering more traction. The next wave of revolution will be to have AI chips that will thrive and drive in the cloud to core to edge space. AI chips would help in performing parallel computing and faster execution of AI related jobs by orders of magnitude compared to traditional CPU architectures. Compared to traditional GPUs and FPGAs, these chips would be more application specific covering areas such as computer vision, speech recognition, robotics and autonomous vehicles. These chips are optimized for inferencing, natural language processing and perform ML and DL jobs.

On the other hand, processor vendors are adding specialized instructions (VNNI, Bfloat16, etc.) to be able to handle mixed precision arithmetic better and optimize performance for AI workloads. Compared to the specialized accelerators, traditional CPUs provide a more general-purpose platform for users to run a wide variety of workloads and perform data manipulation / data prep tasks that are essential in AI/ML/DL workflows. Although a general-purpose CPU will not be as efficient as a purpose-built domain specific accelerator in running AI workloads, consider a CPU-heavy system architecture if the system is to be used for a wide variety of workloads and use cases.

This new wave of AI chips has given rise to a lot of companies trying to make their mark. In addition to big players like Nvidia, Intel, Apple and Alphabet, there is a new line of start-ups who are making their presence felt through AI chips, including Graphcore, Groq, Hailo technologies, Wave computing and Quadric. With the convergence of the AI/ML/DL algorithms and GPUs/accelerators enabling the adoption rate, standardization of AI platforms is happening. AI/ML benchmarking software like Mlperf, DawnBench and ReQuest show the advantage or value of one AI chip and software framework versus the other.

A key enabler for this variety is the availability of software tools and middleware that abstracts the underlying hardware. Although most of these new AI chips have low level APIs and interfaces for programming, it would be a heavy lift for developers to port their software to accelerator-specific programming models. Higher level AI frameworks like Tensorflow, Pytorch, MXnet, etc promise to make the AI models more portable across different hardware accelerators. Availability of cross-framework formats such as Open Neural Network Exchange (ONNX) promise to provide even more flexibility by increase portability across different frameworks.

This market is expected to grow to tens of billions of dollars in the coming 2-3 years and with that, there will be a lot of players (big, medium and small) who would try to differentiate their chips, platforms and applications and prove their merit. The numerous use cases of AI/ML in the market is fueling this race. The usage of AI chips for autonomous cars would be different from usage on image processing and would be different on AI usage on robotics. What would be key in surviving would be to enable usage of AI chips that can help differentiate in speed and agility with which the models run. The next big game changer would be who can make these AI chips faster and conducive to the application. Big companies will have to strive to be nimble to survive in this race, as this is going to be a race where one size will not fit all. Whoever happens to win at the end, it is clear that users will have a lot more choices. The systems of the future will be more hybrid with a mix of general-purpose CPUs and domain specific accelerators. As solution vendors, our challenge is going to be to recommend the right architecture for the right workload and to make the operation and use of such hybrid systems easier for our customers.

It’s All About Transparency

Mukund Khatri - Fellow & VP, Server Infrastructure Solutions, Advanced Engineering

As we kick off 2020 – the year that marks the beginning of what we at Dell Technologies are calling the Next Data Decade - we’ll see organizations accelerate their digital transformation by simplifying and automating their IT infrastructure. Data - massive amounts of data - will be accessible everywhere, requiring data to also be secured everywhere – at the edge, in the data center, and in the cloud.

2020 also marks the beginning of a new era in privacy regulations with the new California Consumer Privacy Act (CCPA) taking effect on Jan 01, 2020. CCPA is set to be the toughest privacy law in United States by broadly expanding the rights of the consumers and requiring businesses within scope to be more transparent in how they collect, use and disclose personal information.

The exponential rise in discovery of new security threats as well as data breaches across IT sectors and devices will unfortunately not cease in the foreseeable future. IT administrators will look for innovative and easy ways to stay current with patch management tools. These tools will be key to minimizing the impact to their business, resource allocations, and other business disruptions. Increased VC funding and innovative startups in the area of cybersecurity will continue to emerge as they try to tap the breadth of cyber-related opportunities. We will also see innovations leveraging new technologies like AI, ML and blockchain across various security solution spaces, including supply chain risk management, advanced threat monitoring solutions, and enhanced access and identity management (AIM).

There will be broad recognition of the critical need and increasing demand for built-in resilience in computing devices against advanced persistent threats, enabled by Silicon Root of Trusts and trustworthy boot flows along with conformance to industry standards like the Platform Resiliency Guidelines, NIST SP800-193, and others.

Newer technologies like Storage Class Memory (SCM), Field Programmable Gate Arrays (FPGA), Smart NICs, all of these while critical for Digital transformation, will require innovative security considerations. Activities tied to quantum safe crypto algorithms will also intensify towards protection of critical assets and data as we enter 2020.

And finally, in 2020 and through the Next Data decade, supply chain security considerations will be one of the top criteria driving purchase decisions for IT devices. Transparency in supply chain inclusive of hardware and software – around how products are developed, delivered, deployed and managed – will be critical for trusted platforms.

Transparency of the state of IT infrastructure will be the mandate of the Next Data Decade.

Game Over –The Server Wins

Gaurav Chawla - Fellow & VP, Server Infrastructure Solutions, Office of CTO
Alan Brumley – Senior Distinguished Engineer, Server Infrastructure Solutions, OEM Engineering

The drive towards software defined architectures started around 10 years ago with SDN (Software Defined Networking) and SDS (Software Define Storage). This gave rise to new scale-out storage architectures for block, file, and object that run on standard x86 based servers. VMware vSAN, Dell EMC VxFlex OS, ECS are just a few examples. The networking architectures also evolved from proprietary networks to Open-flow based programmable switches and new distributed virtual switching for VM and container networking. Dell EMC OS10 Network OS, Microsoft SONiC, VMware NSX, Linux OVS (Open Virtual Switch) and Network Service Mesh (NSM) coupled with L4-7 network services are just some examples. Most hyperscale clouds and large data centers are built leveraging software defined architectures.

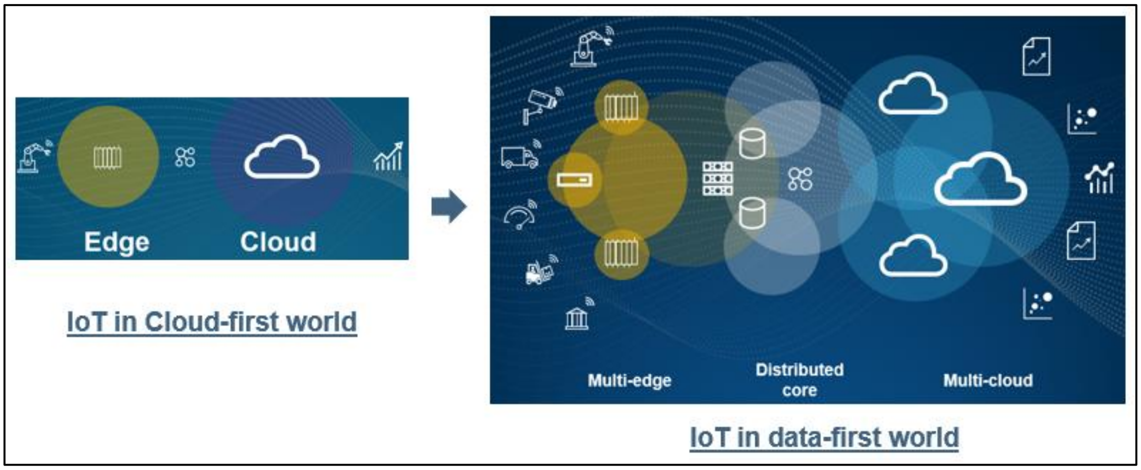

This movement towards software defined is further accelerating as customers move from a "cloud-first" strategy to "data-first" strategy, with both IoT and 5G driving this transition. As customers connect IoT devices to network, there is the need to process the data either on-prem or at the edge. These edge (or distributed core) locations often require devices that are ruggedized, compact, and adaptable to extended temperature ranges. Where data is processed depends on the cost of moving data upstream over the network. In order to move the data processing applications to the edge, the underlying infrastructure services (networking and storage) also need to move to the edge to ensure security, privacy of data, and to enable data analytics at low-latency. This is leading to x86 server systems coupled with offload accelerators as the underlying platforms to host edge workloads. The x86 servers provide an edge environment that connects to other edges and centralized public cloud to enable distributed processing of data.

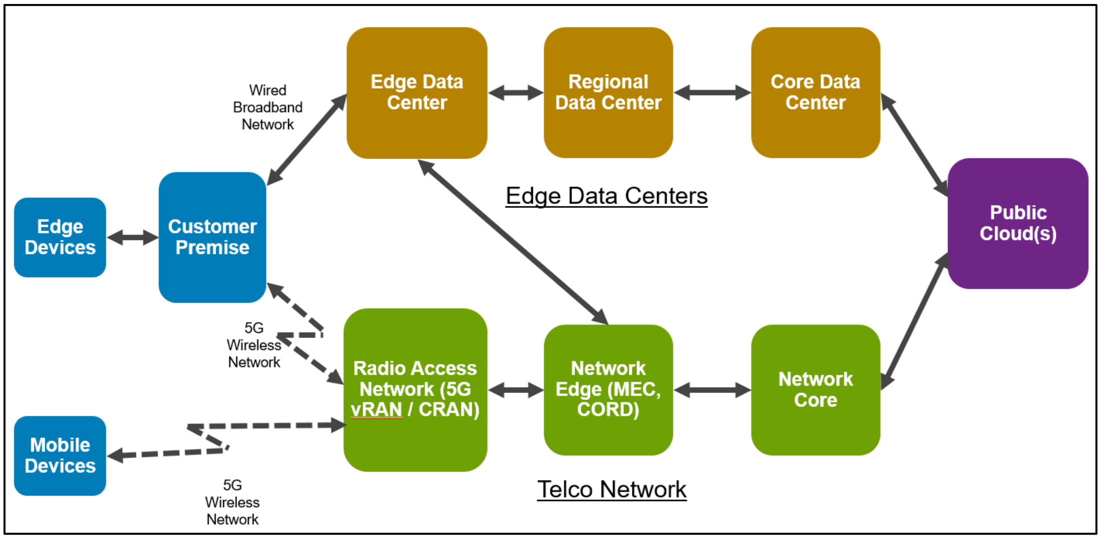

For the connectivity of mobile users (mobile devices and connected/autonomous vehicles) Telcos are moving towards 5G enablement. The 5G infrastructure is causing a transition of Telco edge to become software defined. Telco companies were already on a journey to transition their network core to software defined with NFV (Network Function Virtualization). High speed, low-latency and distance limitations of 5G spectrum is leading to densification of cell sites. This densification is resulting in evolution of proprietary RAN (Radio Access Network) architecture to CRAN/vRAN (Centralized RAN / Virtualized RAN) which leverages x86 servers combined with offloading of RAN and network processing to Smart NICs and FPGAs. In order to process mobile data at the edge, Telcos are also working on MEC (Multi-Access edge Computing) platforms. MEC will enable moving NFV services from network core to edge and hosting 3rd party applications on these MEC servers for data processing. To do this effectively, the Telco edge Cloud must evolve to create a dynamic and well-orchestrated hybrid / multi-cloud environment that can reconfigure and re-scale itself as needed and connect to these new applications made possible by 5G. Examples of some of the applications are AR/VR, gaming, connected/autonomous vehicles and content distribution. MEC platforms use x86 servers as the underlying infrastructure to deliver a cloud like user experience with integrated network, storage, security and Telco virtual network services. The new Telco edge will also create an environment and market where other solutions can deploy to the edge and take advantage of the increased connectivity. For example, this software defined Telco edge will enable enterprises to lease out private 5G as a logical slice on the shared 5G infrastructure. This can be coupled with a leased edge presence in the same Telco MEC infrastructure.

To summarize, scalability, performance and multiple use cases for edge will further accelerate “Software Defined Infrastructure” with all workloads running on the servers. It will also result in Domain Specific Architectures where in certain aspects of infrastructure services and user applications will be accelerated or offloaded to purpose-built co-processors like GPUs, FPGAs and SMART-NICs. The server systems will consist of x86 processors and memory combined with domain specific accelerators.

Infrastructure Automation & AI/Telemetry Are the Keys to the Future

Elie Jreij - Fellow & VP, server Infrastructure Solutions, Advanced Engineering
Jon Hass - Senior Distinguished Engineer, server Infrastructure Solutions office of CTO

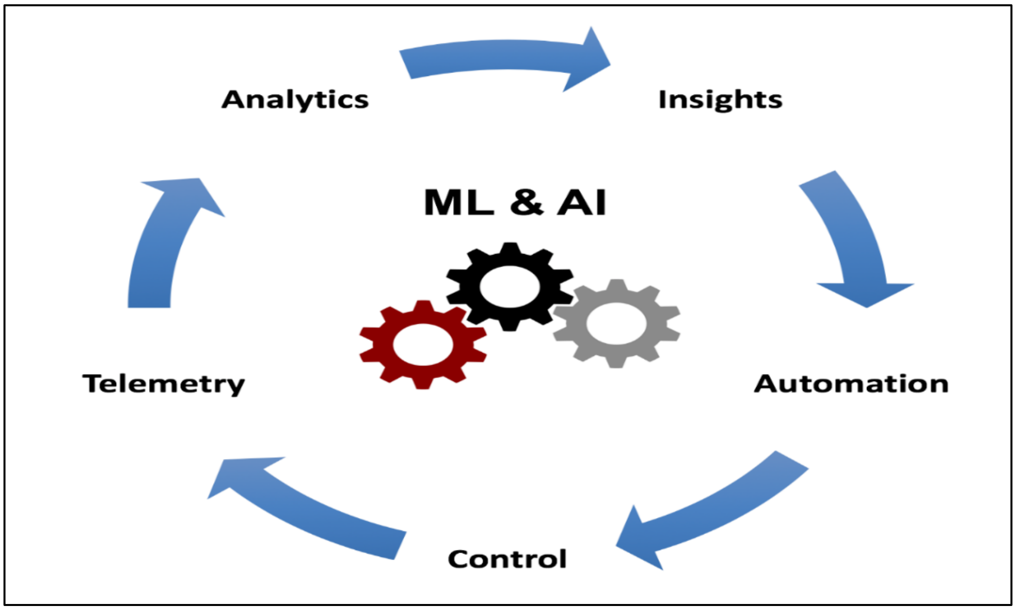

Software defined data centers require flexible, autonomous and self-managing characteristics that will need advancement in several technology areas including zero-touch deployment, composability, advanced telemetry and analytics.

Zero-touch deployment – Security/Discovery

Servers are increasingly being deployed at many remote and edge locations. Having onsite administrators becomes impractical and there is need for complete automation of deployment and lifecycle management. Dell has a long history of providing embedded manageability features that support automation. Recent versions of PowerEdge server iDRAC service processor augment these capabilities with additional security including the expansion of Hardware Root-Of-Trust to include platform BIOS and support for OpenID Connect. Additionally, support for multicast DNS and node-based discovery was added to enhance zero-touch network deployment.

Composability – Abstraction and Outcome Paradigms

As technologies for disaggregating compute architectures to address resource stranding become more prevalent and spans compute, storage, and networking, the complexity of composing systems greatly increases. There is a need to simplify system composition and abstract the administrator and orchestration solutions from needing deep understanding and granular manipulation of the underlying infrastructure. Dell EMC recognizes the importance of outcome-oriented management paradigms and provides some solutions in this area with more sophistication and maturity to come. The focus is on providing declarative techniques that programmatically operationalize management details based on a desired end state definition.

Major increase in compute platform telemetry – new iDRAC capabilities

Asynchronous events have historically been generated by equipment management instrumentation when systems and components needed attention (errors, warnings). However, there are use cases for the continuous streaming of telemetry data that can provide advanced insight into system behavior and early discovery of undesirable operational patterns. The focus on providing the telemetry necessary to perform big data analysis on the operations of the infrastructure is resulting in new and optimized capabilities in Dell EMC product offerings. An example is the newly released PowerEdge server iDRAC9 Telemetry feature. The breadth of available telemetry data has expanded with in-depth power, thermal, network/storage/CPU/memory/accelerator component performance along with historical statistics. The quality of the telemetry data is improved by including timestamp and meta-data for each metric sample. Additionally, the data is marshaled into easily usable JSON documents and delivered via push and streaming protocols to consuming applications and databases.

Applying ML Analytic techniques to N-Dimensional metric data

The amount of data being generated is becoming too large to be managed by historical tools and data bases which require human interaction and interpretation. Increasing sources of telemetry from IT and infrastructure equipment are fueling the ability and need to apply AI and Machine Learning analytics to further the automation of Data Center management, with the goal of attaining full autonomic infrastructure control. Pattern recognition analytics applied to multi-dimensional metrics provide the basis for a breadth of automation from failure prediction and abnormality detection to the optimized execution of workloads and utilization of equipment. Multi-variate visualization techniques help IT administrators make sense of the big data and enable subject matter experts to capture and transfer domain knowledge to general users and automation frameworks.

To Play the Infinite Game, IT Must Embrace edge to Core to Multi-Cloud

Ty Schmitt – Fellow & VP, Server Infrastructure Solutions, Extreme Scale Infrastructure
Joe Vivio - Fellow & VP, Server Infrastructure Solutions, Extreme Scale Infrastructure

Businesses are rushing to transform their business, services and operations using digital technologies like AI, cloud-native and multi-cloud “to win.” The challenge with digital transformation though is to what end? What defines winning? Is being #1, winning? Being #1 at what? Revenue, margin, market share, units… who knows? Understanding the game being played is vital to digital transformation

Business, like many things in life, is an infinite game – there are no winners or losers, just players in game – some known and unknown – with no formal set of rules. You can be in the lead, but you can’t ever win because the game never ends unlike a finite game like baseball that have known rules, known players, known ending state – time, innings, sets, periods, quarters. The goal of the infinite player is to keep the game going.

For businesses to survive in the digital age, digital transformation must include a strategy that embraces the edge to core to multi-cloud era. Whether or not businesses need all 3 elements during phase 1 of digital transformation is not critical, but what is critical is they need to ensure the architectural elements and decisions made include all 3 elements, otherwise another transformation could be required to re-architect the other elements later – adding risk to business.

For example, consider the case of smart cities and thousands of IoT devices and sensors. While it may be possible to start small and have all the IoT devices communicate back to city hall, as the city grows and more intelligence is added, city hall will soon be consumed with a deluge of data – impacting any chance at real time decision making. If you then continue that expansion to the suburbs and larger at the state level, you will quickly realize the need for a tiered approach.

Another example may be a large national retailer or transportation company wanting to utilize captured in-store/vessel data to improve operational effectiveness and marketing communication, ultimately resulting in lower operational expenses, higher margins, and a better overall customer experience. This starts with adding sensor capabilities and on prem capture, compute and storage capabilities. This may be complimented with AR/VR further increasing the amount of data captured and transmitted. The amount of data collected, the timeframe by which something productive is done with this data, the security of this data, and the storage requirements drive costs and what is done on prem vs at the core vs in the cloud.

With these approaches, you will need localized decision making much closer to the action and aggregation points along the way from store/vessel/suburbs to city/state to availability zones. AKA, edge to core to cloud.

As the global leader in IT infrastructure, Dell Technologies continues to build and architect the necessary ingredients for the edge to Core to Cloud era . From Dell edge Gateways for IoT, to SD-WAN edge, to Micro Modular Data Centers, to Rugged servers, to the #1 server Portfolio and #1 Storage Portfolio, leading the way with Open Networking, and finally the software glue to tie it all together with VMware.

Think Outside the Box … And about the Container

Stephen Rousset – Senior Distinguished Engineer, server Infrastructure Solutions office of CTO
Glen Campbell – Member Technical Staff, server Infrastructure Solutions office of CTO

From a server design standpoint, we have always had to consider the software that will be running on the platform to ensure it incorporates enough processing, memory and I/O – sometimes resulting in hardware taxed to support a superset of workloads or hardware specialized to specific functions. Now with the onset of composable and kinetic hardware designs there is the ability for software to dynamically take advantage of the hardware on which it is deployed. We see this with the maturation of containers and the container ecosystem. The meteoric rise of containers and cloud native applications saw initial implementations of abstracting the software coupling to specific hardware capabilities, but as container development has matured, we now see containers being able to use hardware capabilities through application-infrastructure awareness. This is allowing developers the ability to marry application behaviors to specific platform functionality and maximize outcomes.

There is, however, another step of optimizations that can be had in the container ecosystem, one that is facilitated by the telemetry data being provided by the latest Dell PowerEdge servers via Redfish. As containerized applications are deployed within a fleet, telemetry data of the workload and the infrastructure utilization can be collected and fed into the scheduling and dispatching decisions to ensure the deployed jobs are run on an increasingly more optimized platform. We see the making of better container scheduling decisions based on dynamic hardware utilization results through collected telemetry as a collaboration of several Dell technologies across servers, infrastructure management, orchestration and the community. This collaborative effort will manifest itself in new scheduling algorithms and composition of kinetic hardware pieces where in the future state, the platform foundation, the platform resources, the workloads and the composition engine will be intertwined with schedulers via telemetry metrics.

The Rack’s the Limit

Mark Bailey – Senior Distinguished Engineer, Extreme Scale Infrastructure
Robert Hormuth - Technology Fellow, Infrastructure Solutions Group Office of the CTO

Rack power density has remained relatively unchanged for the past decade. Most experts would agree that a target rack density of 5-10kW is a conservative assessment of existing and future hardware requirements. In fact, a world-wide survey of customers suggests that most IT and facility managers don’t view increased rack density as a major inhibitor to their future data center growth and execution metrics. This rear-view perspective is an inaccurate forecast to which both the IT and facility teams will be held accountable looking into the future.

The appetite for compute solutions associated with the global digital transformation, coupled with the rise of AI/ML/DL, suggests that future rack densities will far exceed those of today. New CPU/GPUs will exceed 300Watts each in the near future. DDR5 power and channel number increases, PCIe Gen4/5 power and lane increases, 100G+ Ethernet, and increasing NVMe adoption throws next-generation rack power into the new territory. Customers will soon be faced with critical decisions if they wish to remain in their existing environment: (1) adopt future server solutions but cut the number of servers per rack in order to maintain total rack density, (2) pivot to lower power servers in order to maintain server node count and total rack density, or (3) adopt future server solutions and therefore increase rack power and data center power and cooling accordingly . Some customers believe that they could solve these challenges by moving to a colocation facility or to the public cloud. The reality is, however, that the same challenges exist in those locations as well.

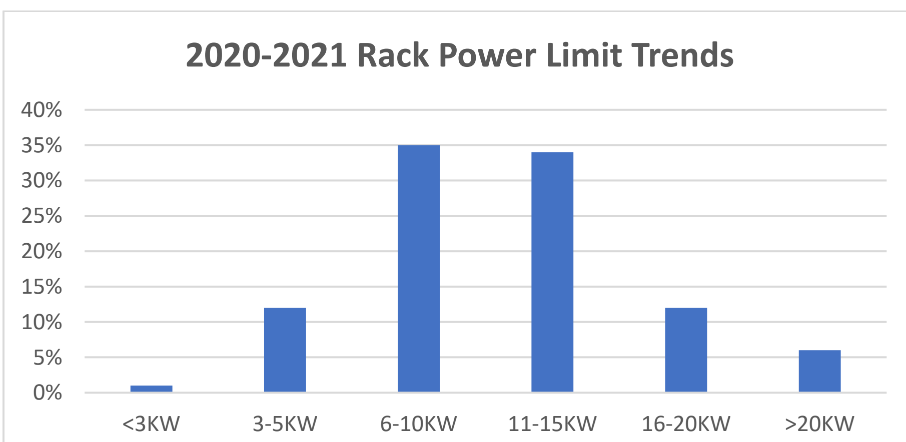

The chart above is a conservative view of rack power densities within the next 18 months. The average projected rack power is more than double what it is today. At Dell EMC, we are working to help customers by innovating across all facets of the data center, allowing real solutions to these emerging challenges. Dell EMC now offers a wide array of 1U and 2U single socket servers that can greatly reduce the per node power and resulting rack density. We also offer a wide variety of 480/277V and 415/240V rack distribution units and PSUs, complimenting the evolution of data center distribution voltage for higher density cabinets.

Dell EMC systems are configured with liquid-to-the-chip and rack-level rear door heat exchangers to more effectively remove heat exchangers (see below).

The data center is just the beginning for next-generation power density challenges. As more customer applications and associated server, IoT/IIoT hardware move outside of the data center, the edge presents significant future challenges for both IT and facility managers. The edge server may differ from its data center counterpart. Edge locations may not always have access to the large amount of power the data center rack density will require. High power cooling may not be available in existing edge locations such as retail stores and office buildings. These edge solutions will evolve to utilize rich, energy efficient, single socket servers. These servers will provide the performance required to handle the massive amount of real-time data generation, in addition to supporting Domain-Specific-Architectures (DSA) such as GPUs, FPGAs and AI accelerators.

The environment in which edge solutions reside will vary greatly, promoting a wide variety of options for deployment. In some cases, edge solutions will require ruggedization to operate in expanded environmental conditions. Alternatively, modular edge data centers – data centers optimized to deploy IT workloads at the edge - will protect edge hardware, allowing more general-purpose options to be deployed. Dell EMC has many innovative and award-winning solutions to accommodate a multitude of deployment opportunities.

The advent of the era of data, AI, ML, DL is upon us all. Rack densities will more than double. The edge will present a multitude of new opportunities and challenges. Dell EMC has the experience, solutions and capabilities to guide customer partners during this industry transition.

Money for Nothing, Chips for Free

Stuart Berke - Fellow & VP, server Infrastructure Solutions, Advanced Engineering

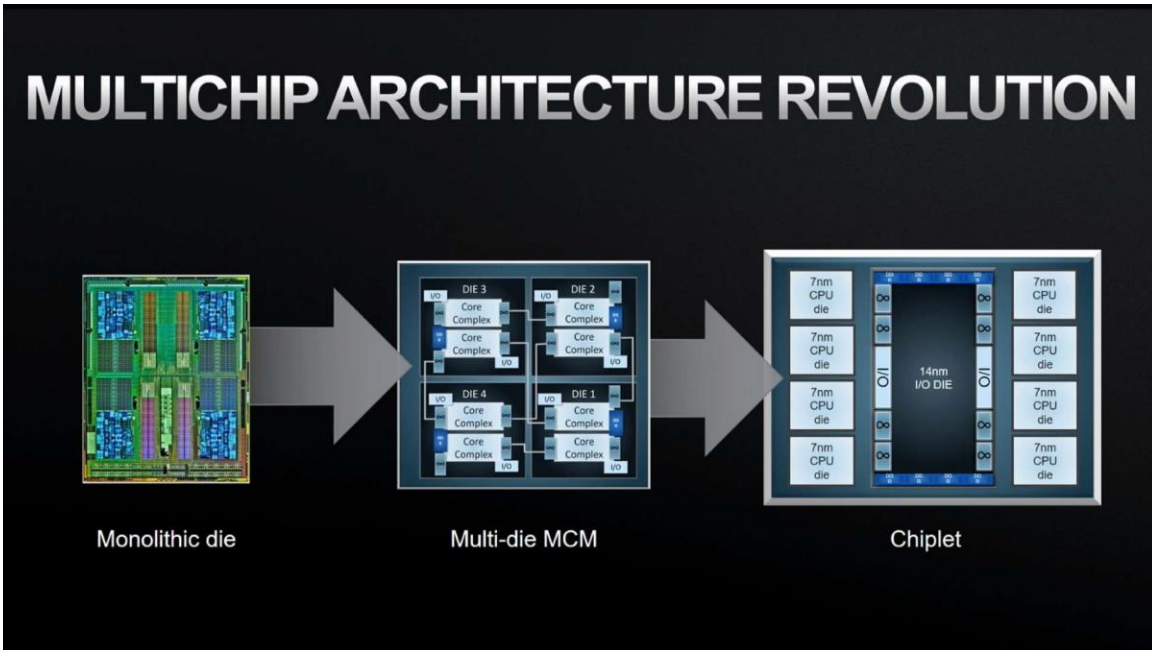

Well, not literally “free”, but the era of multi-die scalable server CPUs is here, enabling more linear pricing and feature set consistency of CPUs as core counts scale to 64 and beyond.

CPU core counts, high speed IO lanes, DDR channels, embedded memory, and other features are growing faster than historically, fueled by silicon process geometries at 10nmm and below. Using traditional monolithic CPU die architecture and implementation, we would be in ‘Dire Straits’ from a yield and cost stand point. However, new approaches to silicon and substrate packaging will increasingly be used to ensure CPUs can scale up affordably.

While Multi-chip modules (MCMs) have been around for generations, they typically have been used to tie together multiple discrete device die within a package to save space or improve signal integrity. The die-die interconnects were limited in width and required high power I/O drivers or SERDES. But improvements in die and packaging now allow a device or element, such as a server CPU “Socket” or “SOC”, to be comprised of multiple sub-CPU silicon dies, commonly called chiplets, tiles, or slices without noticeable performance or power disadvantages.

One example of this new approach is AMD’s second gen EPYC server CPU, in which an IO Die (IOD) and up to eight 8-core CPU Cache Die (CCD) are arranged on a CPU package substrate in a manner that the interconnection of the CCDs and IOD do not impose any significant latency or bandwidth impact vs. if the CPU was built out of a single monolithic die . As the individual CCDs and IOD die are a small fraction of the size of what an equivalent monolithic die would be, large improvements in die yield are realized. Further, this “chiplet” approach allows the CCDs to use a more advanced silicon process than the IOD, saving even more cost.

Another example is Intel's recently announced EMIB (Embedded Multi-die Interconnect Bridge) and Foveros (2D and 3D die packaging and stacking) technologies. EMIB requires just a small embedded silicon connection, tying two chiplets together with high bandwidth and small distance. EMIB is now utilized in Intel's FPGAs, and in the Kaby Lake-G, connecting the GPU to on-package high bandwidth memory. Foveros is a silicon stacking technique that uses TSVs (Through Silicon Vias) whereby for example, separate IO, cores, and memories die are efficiently connected. In some instances, a die at the bottom of the silicon stack acts as an active 'interposer' or substrate.

As CPUs, as well as GPUs, FPGAs, AI/ML chips continue to incorporate larger numbers of computing elements, heterogenous computing elements, on-package memory such as DRAM, HBM, SCM, etc., and IO technologies, the need for "chiplet" style packaging is becoming imperative to improve yield, time to market, development risk, and more linear cost of computing that will drive customer adoption.

The Bus Wars Are Over

Bill Dawkins – Fellow & VP, server Infrastructure Solutions office of CTO

The exponential growth of data, the increased amount of computation done on data due to the larger application of AI/ML/DL and other analytical processes, and the increased value of real-time results have led to the proliferation of hardware techniques to satisfy these trends. GPUs have become mainstream. FPGAs and AI accelerators will become more common in servers. Storage Class Memory (SCM), like Intel Optane, will allow larger data sets to reside in the server memory address space. These HW technologies have something in common, they need to move data in and out of the CPU and between themselves with very low-latency and high-bandwidth. Indeed, we believe that to efficiently use SCM and many accelerators, the latencies between these devices and the CPU need to be less than one microsecond. Furthermore, much like SAN, NAS and software defined storage free data contained in HDDs and SSDs to be shared across multiple servers, there is great benefit for not having data in and functions of these sub-microsecond devices confined to the sheet metal of an individual server.

Many interconnect technologies (or buses – to use the term in a less than technically exact way) have been expanded or developed to address the increased adoption of sub-microsecond devices. Most of these bus technologies have limitations. Some expand existing I/O technologies through RDMA, but these buses do not allow SCM to be accessed as native memory through a CPU's Memory Management Unit (MMU). Other technologies are proprietary point-to-point interconnects that do not expand beyond the chassis or are limited to supporting devices from one or two vendors. Gen-Z and Compute Express Link (CXL) are emerging as the winners in this battle by providing interconnects that are vendor agnostic and support low-latency memory devices and accelerators.

Both Gen-Z and CXL have garnered broad industry support and their respective consortiums both have over 70 members (see genzconsortium.org/about-us/membership/members/ and computeexpresslink.org/members). Both allow memory semantic access to SCM directly through the CPU's MMU with native load/store instructions. Both have the capability to provide sub-microsecond latencies. We believe CXL and Gen-Z are complementary. CXL, as defined in the CXL 1.0 specification, is a point-to-point interface that allows a CPU to directly access memory in a CXL memory device or allow a CXL accelerator to directly access the CPU's memory and do this in a manner that is coherent to the CPU's cache. Gen-Z is a memory semantic fabric that includes multipathing, built-in security and supports large numbers of devices (e.g., up to 4096 in a single subnet). As such, CXL is well suited to support SCM or accelerators inside the server chassis that require cache coherency with the CPU. Gen-Z provides rack scale expansion beyond the server chassis.

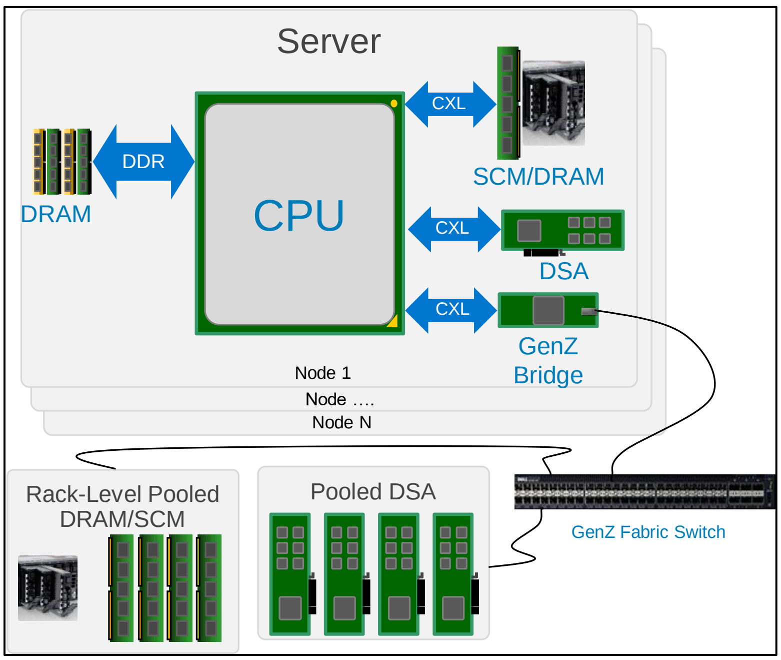

[Note: DSA stands for Domain Specific Architecture. DSAs can be any GPU, FPGA or any other HW accelerator or offload]

The combination of CXL and Gen-Z will usher in new data center innovations. CXL will allow servers to have very large persistent memory address spaces, addressing workloads that have large real-time datasets. Gen-Z will allow these persistent memory devices to be pooled and shared across servers when redundancy and high-availability are required. Gen-Z and CXL also provide a path to full server composability by allowing the disaggregation of sub-microsecond devices at low-latency. CXL will connect to the Gen-Z fabric and the memory centric nature of these interconnect technologies will allow sub-microsecond devices to be accessed directly as memory with CPU native load/store instructions and without the overhead of software I/O protocol stacks. CXL and Gen-Z win the bus wars because they are open, address the needs of sub-microsecond devices and allow a fully composable infrastructure.

It's Time to Make Persistent Memory Boring

Stuart Berke - Fellow & VP, Server Infrastructure Solutions, Advanced Engineering
Gary Kotzur - Senior Distinguished Engineer, Server Infrastructure Solutions, Advanced Engineering

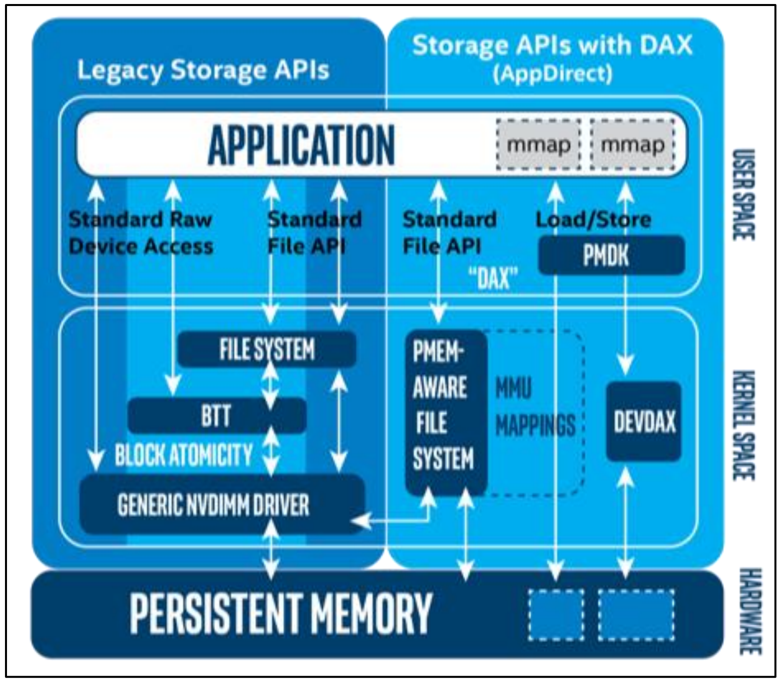

It's been a journey, but all the necessary ingredients to enable ubiquitous persistent memory are falling into place. Proprietary Non-volatile DIMMs are disappearing as industry standard NVDIMMs gain multi-sources and widespread CPU support. Near-DRAM class Storage Class Memory (SCM) media such as Intel Optane DCPMM DIMMs have arrived, enabling new natively non-volatile persistent memory usage models, as well as higher capacities and lower $/GB. And critically, an open industry standard software stack ecosystem, from BIOS and BMC to OS/Hypervisor to Application layer has been embraced to ensure current and new persistent memory types are easily adopted into a common framework.

Improvements in CPU hardware and firmware support for persistent memory will reduce or remove many of today's challenges/barriers, including support for Memory Encryption to safeguard the persistent data while in operation or after power off, architectural optimizations that eliminate the need for software and drivers to manage low-level write cache flushing and ordering consistency, enhancements in hardware memory tiering / caching of persistent memory to improve latency and bandwidth. And natively persistent media will eliminate or lessen the need for batteries to provide hold-up energy to ensure data is not lost during system power loss scenarios.

Industry standard form factors for persistent memory will grow and span multiple standards organizations, including JEDEC DIMMs (NVDIMM-N, NVDIMM-P and others), Compute Express Link (CXL) persistent and storage class cards and modules, and GenZ modules and fabrics, allowing greater support flexibility across server segments within the node, within the chassis, and at rack scale.

Additional SCM media types and variations suitable for server persistent memory are on the horizon, including Phase Change Memory (PCM), Spin Torque Magnetic Memory (ST-MRAM), Carbon Nanotube (CNT or NRAM), and Resistive Memory (ReRAM), which will offer greater choice in capacity, endurance, cost, and performance.

As we look out at 2020 and beyond, the industry will move from primarily early adopter targeted markets today, such as storage, databases, and server appliances, to more general purposes computing deployments where persistent memory improves overall system cost, performance, capacity, boot and recovery time, power, or other criteria.

The World as a Service (is not enough)

Stephen Rousset – Senior Distinguished Engineer, Server Infrastructure Solutions office of CTO

With apologies to Ian Fleming, this is not the title of a James Bond sequel, but more of a statement of where the IT industry has landed and a note of caution about this model. Obviously, with the cloud XaaS (anything as a service) model also comes the associated subscription; paying at some interval for used services. The subscription model is a change from the larger, one-time capital dollars to recurring, smaller OpEx dollars, and provides some short-term balance sheet benefits. So, at a minimum, a long-term financial view should certainly be used when implementing an XaaS model. One must ensure a full understanding of the long-term costs outlay, the price complexities of different service tiers and your utilization of the service. Beyond the financial, there are other aspects to consider in the subscription-based services model to ensure your company is getting the most from the services you are paying.

As has been said, data is the new gold, and with that, everyone wants to stake a claim on your data. As companies embark on more service-oriented offerings, be sure you are aware of how your data is being used within the service, where your data is located, and how securely it is housed. Is your data encrypted at rest and/or in-flight? Are you adhering to your company and countries governances around data? Just a change in the purchase model doesn't affect these items, but several subscription-based services can be cloud-based, requiring this additional scrutiny as your data becomes a desired nugget of value.

Additionally, as you evaluate any XaaS offering, understand the service reliability and availability of the service to your business. Many times, service offerings are monitored by the providing vendor to ensure they are running, but the impact of the systems using the service is not monitored. Such an impact was seen earlier this year when Amazon's S3 had a DDoS attack, which showed the service up, but customers using the service were unable to access their data. The take-away here is, no one is more invested in the success of your business than you, and you need to have a watchful eye on the services you depend and the ones you provide.

And finally, to get the best of the services you use, you also need to have top technical support when you encounter problems. As you look at your subscription services, also look at the service and support history of the company offering those services. At Dell Technologies, we are helping our customers transition to service-oriented offerings for their on-prem and multi-cloud architectures. We bring the breadth of the Dell Technologies portfolio with the security, reliability and support that has made us the premier provider in servers, storage, data protection, HCI, virtualization and security.