OrionX.net

Event

The OrionX Constellation: Data Center

OrionX 2016 Data Center Issues and Predictions

Here at OrionX.net, we are fortunate to work with tech leaders across several industries and geographies, serving markets in Mobile, Social, Cloud, and Big Data (including Analytics, Cognitive Computing, IoT, Machine Learning, Semantic Web, etc.), and focused on pretty much every part of the "stack", from chips to apps and everything in between. Doing this for several years has given us a privileged perspective. We spent some time to discuss what we are seeing and to capture some of the trends in a set of reports.

1- Cloud computing drives further consolidation in hardware

In any value chain, a vendor must decide what value it offers and to whom. With cloud computing, the IT value chain has been disrupted. What used to be a system is now a piece of some cloud somewhere. As the real growth moves to "as-a-service" procurements, there will be fewer but bigger buyers of raw technology who drive hardware design towards scale and commoditization.

2- Composable infrastructure matures, leading to "Data Center as a System"

The computing industry was down the path of hardware partitioning when virtualization took over, and dynamic reconfiguration of hardware resources took a backseat to manipulating software containers. Infrastructure-as-code, composable infrastructure, converged infrastructure, and rack-optimized designs expand that concept. But container reconfiguration is insufficient at scale, and what is needed is hardware reconfiguration across the data center. That is the next frontier and the technologies to enable it are coming.

3- Technical Debt continues to accumulate, raising the cost of eventual modernization

Legacy modernization will get more attention as micro-services, data-flow, and scale-out elasticity become established. But long-term, software engineering is in dire need of the predictability and maintainability that is associated with other engineering disciplines. That need is not going away and may very well require a wholesale new methodology for programming. In the meantime, technologies that help automate software modernization, or enable modular maintainability, will gain traction.

4- Tools emerge to relieve the DB-DevOps squeeze

The technical and operational burden on developers has been growing. It is not sustainable. NoSQL databases removed the time-delay and complexity of a data schema at the expense of more complex codes, pulling developers closer to data management and persistence issues. DevOps, on the other

hand, has pulled developers closer to the actual deployment and operation of apps with the associated networking, resource allocation, and quality-of-service (QoS) issues. This is another "rubber band" that cannot stretch much more. As cloud adoption continues, development, deployment, and operations will become more synchronized enabling more automation.

5- In-memory computing redefines straight-through apps

The idea of a "memory-only architecture" dates back several decades. New super-large memory systems are finally making it possible to hold entire data sets in memory. Combine this with Flash (and other emerging storage-class memory technologies) and you have the recipe for entirely new ways of achieving near-real-time/straight-through processing.

6- Multi-cloud will be used as a single cloud

Small and mid-size public cloud providers will form coalitions around a large market leader to offer enterprise customers the flexibility of cloud without the lock-in and the risk of having a single supplier for a given app. This opens the door for transparently running a single app across multiple public clouds at the same time.

7- Binary compatibility cracks

It's been years since most app developers needed to know what CPU their app runs on, since they work on the higher levels of a tall software stack. Porting code sill requires time and effort but for elastic/stateless cloud apps, the work is to make sure the software stack is there and works as expected. But the emergence of large cloud providers is changing the dynamics. They have the wherewithal to port any system software to any CPU thus assuring a rich software infrastructure. And they need to differentiate and cut costs. We are already seeing GPUs in cloud offerings and FPGAs embedded in CPUs. We will also see the first examples of special cloud regions based on one or more of ARM, OpenPower, MIPS, and SPARC. Large providers can now offer a usable cloud infrastructure using any hardware that is differentiated and economically viable, free from the requirement of binary compatibility.