2017 Reflects the Tipping Point in IT Infrastructure Design and Protection

At the end of the year people naturally reflect on the events of the past year and look forward to the new. I am no different. It is as I reflect on the past year and look ahead on how IT infrastructures within organizations have changed and will change, 2017 has been as transformative as any year in the past decade if not the past 50 years. While that may sound presumptuous, 2017 seems to be the year that reflects the tipping point in how organizations will build out and protect their infrastructures going forward.

Over the last few years technologies have been coming to market that challenge two long standing assumptions regarding the build out of IT infrastructures and the protection of the data stored in that infrastructure.

1. The IT infrastructure stack consists of a server with its own CPU, memory, networking, and storage stack (or derivations thereof) to support it
2. The best means of protecting data stored in that stack is done at the file level

Over the last two decades, organizations of all sizes have been grappling with how best to accommodate and manage the

introduction of applications into their environment that automate everything. They have been particularly stressed on the IT infrastructure side with each application needing its own supporting server stack. While managing one or even a few (less than 5) applications may be adequately achieved using the original physical server stack, more than that starts to break the stack and create new inefficiencies.

These inefficiencies gave rise to virtualization at the server, networking, and storage levels which helped to somewhat alleviate these inefficiencies. However, at the end of day, one still had multiple physical servers, storage arrays, and networking switches that now hosted virtual servers, storage arrays, and fabrics. This virtualization solved some problems but created its own set of complexities that made managing these virtualized infrastructures even worse if one did proactively put in place or have in place frameworks to automate the management of these virtualized infrastructures.

Further aggravating this situation, organizations also needed to protect the data residing on this IT infrastructure. In protecting it, one of the underlying assumptions made by both providers of data protection software and those who acquired it was that data was best protected at the file level. While this premise largely worked well when applications resided on physical servers, it begins to break down in virtualized environments and almost completely falls apart in virtualized environments with hundreds or thousands of virtual machines (VMs).

These inefficiencies associated with very large (and even not so large) virtualized environments have resulted in the following two trends coming to the forefront and transforming how organizations manage their IT infrastructures going forward.

1. Hyper-converged infrastructures will become the

predominant way that organizations will deploy, host, and manage applications going forward
2. Data protection will predominantly occur at the volume level as opposed to the host level

I call out hyper-converged infrastructures as this architecture provides organizations the means to successfully manage and scale their IT infrastructure. It does so with minimal to no compromise on any of the features that organizations want their IT infrastructure to provide: affordability, availability, manageability, reliability, scalability, or any of the other abilities I mentioned in my blog entry from last week.

The same holds true with protecting applications at the volume level. By primarily creating copies of data at the volume level (aka virtual machine level) instead of the file level, organizations get the level of recoverability that they need with the ease and speed at which they need it.

I call out 2017 as a tipping point in the deployment of IT infrastructures in large part because the combination of hyper-converged infrastructures and the protection of data at the volume level enables the IT infrastructure to finally get out of the way of organizations easily and quickly deploying more applications. Too often organizations hit a wall of sorts that precluded them from adopting new applications as quickly, easily, and cost-effectively as they wanted because the existing IT infrastructures only scaled up to a point. Thanks to the availability and broad acceptance of hyper-converged infrastructures and volume level data protection, it appears the internal IT infrastructure wall that prevented the rapid adoption of new technologies has finally fallen.