White Paper

Lenovo Delivers Seamless Rack-Scale Hyperconvergence with ThinkAgile SX for Nutanix

Sponsored by: Lenovo

Eric Sheppard
June 2017

SITUATION OVERVIEW

The use of converged infrastructure as an alternative to traditional procurement and delivery models has been one of the most important developments in the evolution of the datacenter infrastructure market. IT departments and service providers increasingly look to technology suppliers for converged infrastructure that encapsulates server, storage, networking, and management software into a tightly integrated solution that is supported and managed more efficiently than disparate solutions.

The pain points that have driven companies toward converged systems have evolved but have remained broadly consistent over the years. The most common challenges IT departments look to address when deploying converged infrastructure are:

- High capital costs associated with datacenter resources
- Length of time required to deploy new applications and business services (i.e., time to value)
- Amount of time spent troubleshooting and maintaining datacenter infrastructure
- Difficulties associated with scaling and refreshing datacenter infrastructure
- Application outages due to planned and unplanned infrastructure downtime
- Inability to provide, or test, disaster recovery (DR) capabilities due to costs and complexity
- Operational inefficiencies related to silos of infrastructure requiring silos of experts to manage
- Low infrastructure utilization rates due to sprawl of infrastructure islands throughout the organization

Converged infrastructure solutions are designed to address these pain points by eliminating infrastructure silos, enabling centralized management, leveraging software-defined infrastructure, offering the choice of scale-up or scale-out architectures, leveraging automation to increase predictable repeatability and reduce human errors, and reducing costs associated with DR.

In return, organizations that have deployed converged infrastructure tell IDC that they are able to:

- Rightsize capital costs
- Reduce complexity and risk
- Reduce errors
- Increase reliable repeatability
- Increase predictability

- Reduce operational expenses
- Increase operational efficiency
- Increase staff agility
- Reduce time needed to roll out business services

With such benefits, it should not be surprising to learn that converged infrastructure adoption has grown quickly over the years to become a multibillion-dollar market. Total spending on integrated systems exceeded $12.5 billion in 2016 and is expanding rapidly. An important reason for the sustained market growth of converged infrastructure can be attributed to the constant evolution of solutions brought to market. Today's converged infrastructure solutions offer far greater levels of automation in many areas such as deployment and life-cycle management. This has helped to reduce errors, increase the speed of such tasks, and ensure reliable repeatability of many critical tasks. These solutions are also increasingly managed by centralized teams with common skills and goals rather than silos of infrastructure-centric experts that often point fingers when problems arise.

The next stage of development in converged infrastructure comes from hyperconverged infrastructure, which IDC considers a subset of the broader converged infrastructure market. The global hyperconverged market was worth $2.3 billion in 2016. This represented 18.5% of total 2016 converged infrastructure spending. IDC expects hyperconverged solutions to represent almost all of the market growth through 2021, when it will represent 41% of global converged infrastructure spending. Hyperconverged infrastructure is helping deliver all of the proven benefits reviewed previously (e.g., reduced capex, reduced complexity, reduced risk, reduced opex, and improved time to services through staff agility) but does so through a clustered, scale-out architecture that is built on highly virtualized x86-based servers. A key characteristic of hyperconverged systems that differentiate these solutions from traditional converged infrastructure is their ability to provide compute, storage, and networking functionalities through the same server-based resources (or nodes). Tightly integrating these technologies into a simple, easy-to-manage, single resource pool also eliminates the specialized skills and certifications typically needed in traditional infrastructure siloes of separate compute, storage, and virtualization products. Nodes are usually deployed as clusters, with each node in the cluster contributing all of its resources to an abstracted (i.e., virtualized) pool of capacity, memory, and compute resources. This pool of virtualized resources provides the foundation for all server-centric workloads (e.g., VM, hypervisor, and application) as well as storage-centric workloads (e.g., data persistence, data management, replication, snapshots, and deduplication).

Evolution of Hyperconvergence

Hyperconverged systems found early success within midsize environments, individual business units in larger organizations, and targeted workloads such as VDI. Broader acceptance of hyperconvergence and increased awareness of the benefits these solutions bring to the table (agility, flexibility, scale, web-scale economics, and ease of use) have driven rapid innovation and growth in this market. Today, the breadth of workloads running on hyperconverged solutions is expanding rapidly. Highly virtualized and mission-critical business applications are more common than people may realize. In fact, hyperconverged vendors such as Lenovo and Nutanix report that business-critical applications such as SAP and Microsoft SQL Server are the most frequently deployed applications on their platforms. Growth of new hyperconverged deployments and expansion of workloads running on these systems have helped to drive total 2016 hyperconverged sales (including hardware and software) up 110% to more than $2 billion in global sales. Indeed, the hyperconverged market has reached a new phase of maturity.

While the first phase of hyperconverged maturity saw deployments of relatively small clusters of loosely coupled appliances, this new phase of market maturity will see increased deployments of hyperconverged solutions that are able to support further consolidation of virtualized enterprise workloads without increasing performance volatility. Since the workloads being consolidated onto these systems will include a wide range of enterprise applications like databases, collaboration and messaging systems, web-based and file-based applications, virtual desktop infrastructure (VDI), and other clone-intensive environments, modern hyperconverged solutions need to accommodate a wide range of workload profiles. This means consistently delivering submillisecond response times and supporting hundreds of thousands to millions of IOPS. Thus users will increasingly expect hyperconverged solutions to provide a mature quality of service (QoS) engine to address "noisy neighbor" issues and ensure that performance can be consistently delivered as required, no matter what else is going on in the system.

Demand for increased workload density also means modern hyperconverged solutions must scale easily without increasing complexity or management burdens. Consolidating virtualized workloads onto ever-larger hyperconverged solutions will also drive the need for solutions that can scale into at least the petabyte range, with a road map that extends the range to tens of petabytes. The scale-out nature of hyperconverged architectures has an advantage in that hyperconverged architectures can start smaller, add performance and/or capacity resources only as needed, and do not impose disruptive forklift upgrades as environments grow. With increased scale, however, comes a commensurate need for simplified management. This will partly come from a reduction of management interfaces that can span multiple systems and locations. Simplified management will also result from increased levels of automation, better use of APIs, and introduction of software-defined technologies within the networking of the hyperconverged solution.

Emergence of Rack-Scale Hyperconverged Solutions

IDC believes this new phase of hyperconverged market maturity will also see increased deployment of solutions referred to as rack-scale hyperconverged, which combines the best aspects of converged and hyperconverged.

Rack-scale hyperconverged solutions are loosely coupled hyperconverged solutions that provide all of same x86-based virtualized compute (hypervisor) and software-defined storage (SDS) resources as other hyperconverged solutions but are optimized for larger deployments that are intended to scale multiple datacenter racks. To achieve this, rack-scale hyperconverged solutions include the following additional software and hardware:

- **Software-defined networking:** Rack-scale hyperconverged systems ship with software-defined networking to provide network virtualization and automation of networking management tasks that can be triggered via API calls when changes are made within a virtual machine's environment.
- **Top-of-rack switching:** In addition to the internode switches commonly leveraged for hyperconverged clusters, rack-scale hyperconverged will leverage top-of-rack switching for inter-rack scaling.

Although there is overlap between rack-scale hyperconverged solutions and traditional appliance-based hyperconverged solutions, the benefits provided by rack-scale solutions come from the following attributes:

- They are designed to scale beyond multiple datacenter racks as a single system.

- They ship as preconfigured, full systems that can be deployed with ease and reduce onsite services time.
- They provide predictable network performance with QoS features and extensive use of flash-based storage media (SSDs and/or NVMe), thus addressing noisy neighbors and increased system I/O traffic.
- They eliminate network complexity issues through increased use of automation and SDN features.
- They offer network management, which had traditionally been limited in early hyperconverged offerings.

To be clear, the future of hyperconvergence will be one in which traditional, appliance-based solutions will see new deployments continue to grow very rapidly. Increased adoption of rack-scale hyperconverged solutions is expected to occur in tandem as use of hyperconverged solutions expands within midsize and larger enterprises with a need for larger-scale deployments. Both are considered important to the future of the hyperconverged market.

Overview of Lenovo's Hyperconverged Solutions

Lenovo offers a broad portfolio of converged systems solutions that are intended to drive improved levels of IT staff agility, reduce capital and operational costs, and offer comprehensive services and support covering all aspects of its solutions. Lenovo's family of converged systems solutions includes two core hyperconverged offerings: the ThinkAgile HX Series and the ThinkAgile SX for Nutanix (SXN).

Lenovo's ThinkAgile HX Series is the company's appliance-based hyperconverged solution that fully integrates Lenovo x86 servers with core Nutanix software to provide application mobility, a distributed storage fabric and a hypervisor (all part of Nutanix Acropolis), and system management software (Nutanix Prism to manage the cluster and virtual machines and Lenovo XClarity to manage the hardware resources). Lenovo's hyperconverged solutions leverage virtualization and software-defined storage to natively collapse core storage, compute, and networking functions into a single pool of resources that are deployed as a scale-out cluster.

Lenovo's ThinkAgile SX for Nutanix is a fully configured and validated rack-scale hyperconverged solution that leverages ThinkAgile HX Series nodes as the core building blocks in conjunction with Lenovo networking and infrastructure software. All ThinkAgile SXN solutions are delivered as a fully integrated rack-scale hyperconverged solution with all infrastructure software (e.g., hypervisor, management software, SDS, SDN) pre-installed so that it can be running production workloads hours after arriving onsite. Specifically, Lenovo's ThinkAgile SX for Nutanix includes the following elements:

- **Lenovo ThinkAgile HX Series.** These hyperconverged appliances and software include Nutanix Acropolis OS and Prism Manager for virtual machine and cluster management (Pro or Ultimate).
- **Lenovo ThinkSystem and RackSwitch top-of- rack switches** run Cloud Network Operating System (CNOS).
- **Lenovo's ThinkAgile Network Orchestrator** provides automated network provisioning and dynamic synchronization between the virtual machine layer and the physical switch. ThinkAgile Network Orchestrator is software based on Lenovo's CNOS Ethernet switches that fully integrate with Prism management software so that any changes to the virtual machines impacting the virtual networks will trigger an API call to the physical Lenovo switches, which

then dynamically (and automatically) reconfigure VLAN settings to accommodate the required changes. Many operations leave unused VLANs configured to reduce ongoing opex and downtime, which can impact security and performance. With ThinkAgile Network Orchestrator, that is not needed. The result is that ThinkAgile Network Orchestrator eliminates human error, reduces the need for scheduled maintenance windows, reduces maintenance/administration time, increases security, improves performance, and enhances overall operational efficiency through reduced dependency on network administrators.

- **Lenovo's XClarity Administrator or Lenovo XClarity Integrator (LXCI) for Prism** allows users to choose between two software suites to manage system resources:
  - **Lenovo's XClarity Administrator** is a suite of centralized resource management software that allows IT managers to automate a number of lower-level tasks including hardware discovery, hardware monitoring, and firmware updates. XClarity also automates system configuration and the deployment of bare metal operating systems and hypervisors. This, along with a dashboard-driven interface, allows IT staff to spend less time running/managing their IT and more time innovating and improving their business.
  - **Lenovo XClarity Integrator for Prism** integrates Lenovo XClarity directly into Nutanix Prism management software to provide an administrator with the ability to monitor and manage both the physical and the virtual resources from a single pane of glass. An important part of this integration will be the ability to leverage Nutanix's Prism Life Cycle Manager for updating firmware, which allows administrators to monitor the firmware of all cluster resources and display any required updates. If updates are available, an administrator can easily update firmware in the cluster through a simple one-click option that initiates an automated process deploying updates throughout the cluster, nondisruptively.
- **Lenovo premier services and ThinkAgile Advantage support** provide customers with a single point of support for all aspects of the solution.

With the introduction of ThinkAgile SXN, Lenovo offers a truly modern rack-scale hyperconverged solution that provides users with flexibility of choice (choose the best combination of ThinkAgile HX nodes and management software for each environment), the power of systemwide automation, and true enterprise scale. Regardless of the configurations chosen, those deploying ThinkAgile SXN can expect a rack-scale hyperconverged solution that will expand the core benefits of the company's ThinkAgile HX Series appliances (e.g., IT infrastructure cost reductions, IT staff productivity and agility increases, risk mitigation and reduced downtime, faster time to new business services) to provide additional benefits related to network automation and rack-scale management features (e.g., elimination of human error, reduced need for scheduled maintenance windows, reduced maintenance/administration time, enhanced overall operational efficiency through reduced dependency on network administrators).

It should be noted that the ThinkAgile HX Series appliances used within ThinkAgile SXN are built on Lenovo's ThinkSystem servers. Lenovo servers have a long history of deployments within datacenters where quality, reliability, and security have been high priorities. While ThinkAgile SXN is built with Lenovo's ThinkAgile HX Series appliances, there are important differences between ThinkAgile HX Series and ThinkAgile SXN. ThinkAgile HX deployments are at the node level, typically 1-5 nodes at a time. ThinkAgile SXN solutions are targeted at environments needing to start with 8-16 nodes at a time and intended for customers that want the fully integrated, automated, configured bundle with the premier services at the rack level. Lenovo's use of its ThinkAgile HX Series appliances as core building blocks for its ThinkAgile SXN ensures a consistent, seamless experience for customers based on a mature and truly tested hyperconverged technology. ThinkAgile SXN provides true enterprise scale without increased complexity or risk by ensuring everything is pre-validated, fully integrated,

preconfigured, delivered, deployed, and supported as a holistic/single system. The company has done a considerable amount of work to ensure ThinkAgile SXN can be up and running much faster than what it would be if one were to build a comparable system with individual components. Lenovo plans to automate as much of the validation and configuration of ThinkAgile SXN in manufacturing as possible. Last, Lenovo bundles its global support and services capabilities to ensure a seamless, end-to-end customer experience and issue resolution.

As an early example of a rack-scale hyperconverged solution, Lenovo has shown the value of incorporating both manufacturing automation and software-automated networking when designing hyperconverged solutions for this new phase of hyperconvergence.

CHALLENGES/OPPORTUNITIES

Total spending on converged systems has recently surpassed $12.5 billion in annual spending. While hyperconvergence investments represent a relatively small $2.2 billion portion of this market, sales are expanding at triple-digit rates and expect to account for 41% of total spending on converged infrastructure by 2021. This rapid market expansion can be attributed to the ability of hyperconverged solutions to improve long-standing datacenter metrics such as capital costs, utilization rates, time to service, operational efficiency, and levels of risk. The breadth of workloads running on hyperconverged solutions is also expanding rapidly, with an increasing percentage of these workloads being mission-critical business applications. As we move forward in time, rack-scale hyperconverged solutions will only help push hyperconverged deployments deeper into larger organizations looking to deploy the technology within large-scale, mixed-workload environments. IDC recommends viewing appliance-based hyperconvergence and rack-scale hyperconverged solutions as complementary solutions that can help drive new levels of capital and operational efficiencies in datacenters that are large or small. With the relatively new addition of rack-scale hyperconverged solutions, the hyperconverged market in general appears to be better aligned with contemporary datacenter challenges than any other time in the past.

IDC believes the most pressing challenge for Lenovo will be to leverage the growing awareness of ThinkAgile HX Series to help drive adoption of ThinkAgile SXN and ensure its portfolio of offerings is on the short list of companies considering hyperconvergence. As noted previously, Lenovo's use of its ThinkAgile HX Series appliances as core building blocks for its ThinkAgile SXN ensures a consistent, seamless experience for customers based on a mature and truly tested hyperconverged technology. The ThinkAgile SXN's ability to provide automated network provisioning and dynamic synchronization between the virtual machine layer and the physical switch is very much a forward-looking design that helps drive true rack-scale hyperconvergence without increasing complexity, management burdens, or risk to the datacenter. This should help provide considerable customer traction, as users begin to learn more about the company's converged systems offerings.

CONCLUSION

The use of converged infrastructure as an alternative to traditional procurement and delivery models will continue to be one of the most important developments in the evolution of the IT infrastructure market. After nearly a decade of real-world deployments, the benefits of convergence have proven to be very real. These benefits of infrastructure convergence are expected to remain tied to faster time to service/market, increased cost advantages, greater operational efficiencies, and improved resiliency. The core technology and architectural underpinnings of converged infrastructure, however, have

begun to change dramatically. Converged infrastructure is increasingly software defined, server based, and highly automated. Appliance-based hyperconvergence has represented a large part of this architectural shift to date and is expected to drive considerable market expansion moving forward. That said, the emergence of rack-scale hyperconverged solutions like Lenovo's ThinkAgile SXN should be viewed as an important milestone for this market and one that will help support further consolidation of virtualized enterprise workloads without increasing complexity, management burdens, or performance volatility. IDC views appliance-based hyperconverged solutions and rack-scale hyperconverged solutions as complementary technologies that will coexist and drive market expansion within the diverse global datacenter market. Both types of hyperconverged solutions are also expected to help customers shift datacenter investments toward infrastructure supporting a more agile IT department that is better able to support increasing investments in developers and programmers.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com