

Creating cloud revenue streams for Managed Service Providers

WHITE PAPER

A guide for MSPs to solving legacy cloud problems and creating new revenue streams whilst offering clients a tangible competitive advantage.

Table of Contents

Executive Summary	3
The Inevitable Cloud	3
Current Industry Technology Trends	4
I. SDDC, Cloud and Hyperconvergence.....	4
II. Hyperconverged Appliances vs. The Hyperconverged Platform	5
III. Cloud Computing	5
Cloud Landscape	6
Legacy Approaches	6
Disruptive Innovation - Hyperconverged IaaS	6
Cloud 2.0 Hyperconverged Architecture.....	8
Cloud Orchestration And Virtualization	9
Integrated Storage	9
High-Availability By Design.....	10
Software-Defined Networking (SDN).....	10
Core Cloud 2.0 Revenue Streams	11
Conclusion	12
References For Action.....	13

Executive Summary

The Inevitable Cloud 2.0

In this whitepaper we discuss current industry trends, the traditional cloud landscape and the new and emerging cloud 2.0 built on hyperconvergence platform. We also evaluate the Build Your Own (BYO) Cloud vs. the Retail IaaS cloud offered by technology giants such as Microsoft Azure and Amazon Web Services. We then compare these options and evaluate how they stack up against the new hyperconvergent platform that offers many times the performance and scalability against such providers.

Businesses ranging from large corporations through to SME are today competing on two levels: through the quality of their actual product and through the technology they employ to communicate and deliver it. As a result of a more strategic emphasis on technology as a differentiator, it is incumbent on executives to focus on viable business tools and also provide their company with a technology edge over competitors. Furthermore, corporations are becoming increasingly aware that these solutions must be agile enough to deal with changing business requirements and disruptive technological innovation.

In response to this changing landscape enterprise IT has moved away from siloed solutions and is embracing converged, virtualized and cloud-based technologies. Those who have not yet done so, will. It is only a matter of time.

However, not all cloud-based solutions are created equal. The traditional build your own cloud (BYO) is a viable option only to this with vast financial and human resources, and even then the advantages of control and reliability have virtually been eliminated. Large initial and ongoing capital expenditures combine with technical complexity and data center overheads to make for a very risky proposition. Due to the rapidly changing technology environment the investment of time and money runs the risk of being outdated well before it has been completely installed or fully amortized.

Today's infrastructure environments are typically

comprised of 8 to 12 hardware and software products from as many vendors, with each product offering a different management interface and requiring different training. Furthermore, each product in this type of legacy stack is grossly overprovisioned, using its own resources (CPU, DRAM, storage, and so on) to address the intermittent peak workloads of the resident applications. The value of a single shared resource pool, offered by server virtualization, is also limited to the server layer. All other products are islands of overprovisioned resources that aren't shared. Low utilization of the overall stack therefore results in the ripple effects of high acquisition, space, and power costs. Simply put, too much money and too many resources are wasted in today's legacy environments.

Recently emerged technology stands in stark contrast to the aforementioned BYO approach, however. Hyperconverged technology is everything that BYO is not. Although it is commonly referred to as Cloud 2.0, hyperconverged cloud is the first service that realizes the full potential of cloud-based technology. To IT departments, Cloud 2.0 delivers on the cloud's promise of turnkey infrastructure competency that can be deployed in a few days and does not require extensive teams of specially trained engineers for maintenance and scaling. Cloud software reduces the business costs of implementing, configuring, and maintaining software. By lowering the "friction" associated with software change, the Cloud helps business keep a flexible computing environment and adapt more quickly to changes. Given the increasing importance of technology to business operations of every kind, this adaptability is becoming increasingly critical.

Unsurprisingly, the shift to cloud computing is growing at an exponential rate and is a journey that businesses simply must make sooner or later. Business executives can now effectively manage their technology risk, as well as dramatically reduce their capital expenditures for IT.

Current Industry Technology Trends

In order to better understand the current industry technology trends, it is essential to have an in-depth understanding of some of the technical terms and the technology used. The most significant transformation in the IT infrastructure is reflected by three major trends:

- I. Software-defined data centers (SDCCs)
- II. Hyperconvergence
- III. Cloud computing

The emergence of these trends is a result of the current IT infrastructure clutter, complexity and high cost. These trends are attempts to simplify IT and reduce the overall cost of infrastructure ownership.

Companies like Google and Facebook have been pioneers in creating sheer scalability at reasonable costs, but a lot of these cloud principles used by these companies can be adapted for smaller environments using hyperconverged technology for any size company. The factors driving us toward such hyperconverged technology include:

I. Software-defined data centers:

Companies like Google discovered the potential of software-centric design of software-defined data center (SDCC) years ago and created efficiencies in its hardware platform by wrapping it inside software layers. This is an integral part of software-centric design.

Commodity hardware and seamless changes:

Change is constant in the cloud. A change in data center hardware should not require a reconfiguration of all the virtual machines and policies. Using commodity hardware and resisting expensive proprietary components can facilitate the need for scaling your cloud environment as needed. A hyperconverged environment reduces the risk of using commodity hardware as its software layer is built with the understanding that hardware can eventually fail. Commodity hardware that is readily interchangeable removes the need to recode the entire system which in-turn results in no disruption of service.

Bite-sized scalability:

Customers no longer need to expand just one component or hardware rack at a time; they can simply add another node to a homogenous environment.

VM-Centricity:

The individual VMs are the workload that takes center

stage in the cloud. You can now apply policies directly to the VMs without having to worry about the LUN, shares, data stores or any other constructs. Mission-critical applications are designed expecting high-availability infrastructure, disaster recovery, backup and recovery along with a range of other necessary services. A hyperconverged IaaS platform must deliver on these requirements. Hence, hyperconvergence brings cloud-type consumption-based infrastructure economics and flexibility without compromising on performance, reliability and scalability.

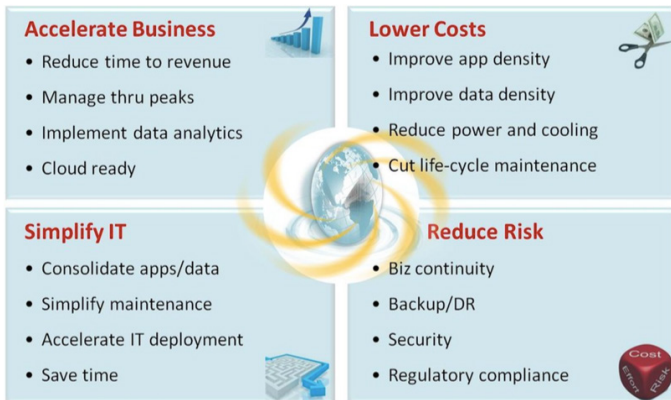
II. Hyperconvergence:

Hyperconvergence takes away the need to buy huge servers every few years by simply allowing you to add building blocks of infrastructure as needed. It delivers simplification and savings by consolidating all required functionality into a single infrastructure stack by completely reinventing the underlying data structure required by Software-Defined Data centers (SDCC).

By using convergence and a single shared resource pool, hyperconvergence makes many legacy services obsolete (backup, replication, duplication appliances, solid-state drive (SSD) arrays, SSD cache arrays, public cloud gateways, and replication appliances or software).

A recent study from IDC (a premier global market intelligence firm) shows the most common drivers for businesses adopting converged infrastructure:

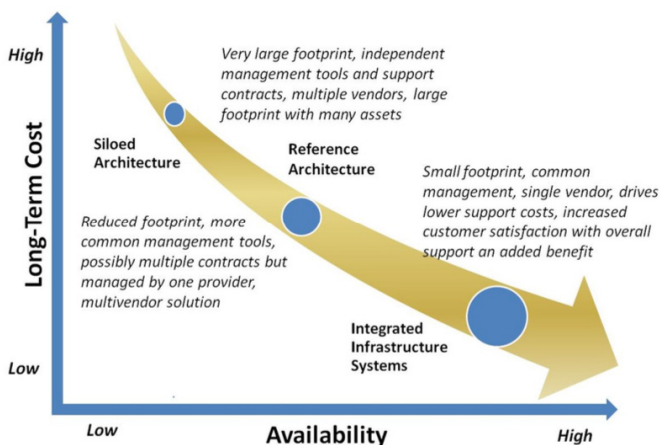
Integrated and Converged Systems Drivers



Advantages to hyperconvergence include a relatively simplified design and architecture and administration model. Instead of having different teams to manage software, hardware, storage arrays, virtualization, etc.; one team can manage the complete hyperconverged stack.

Another IDC research study indicates that an IT operations team typically spends 70% of their time on day-to-day operations leaving little time for them to add value to the business or support innovation. Hence, data center managers and Managed Service Providers (MSPs) are looking for robust solutions to help streamline operational inefficiencies.

Total Cost of Support by Architecture Type



Hyperconverged appliances vs. hyperconverged platform:

It is important to state that hyperconverged appliances are distinctly different from the hyperconverged platform. Hyperconverged appliances have a number of disadvantages:

- Inability to make granular upgrades or tweaks
- Storage growth and performance tuning is not easy to do
- Inefficiency – If a cluster runs low on storage but not on computer, you end up adding more computer along with storage by adding another appliance and vice versa.

A full featured hyperconverged platform on the other hand is the most recent addition to the cloud landscape and is capable of resolving all the aforementioned issues.

III. Cloud computing:

Cloud computing is simply the delivery of computing services such as servers, storage, applications and network infrastructure delivered over the Internet. It helps lower your operating costs, provides scalability and allows companies to create an efficient infrastructure. Large cloud computing services run on secure data centers which are constantly upgrading their resources with the latest generation of fast and efficient computing hardware.

The cloud landscape has evolved considerably over the years. Cloud service providers now offer public, private and hybrid cloud solutions based on company's needs and requirements. Companies can build their own clouds or purchase cloud services through retail IaaS (Infrastructure as a Service) providers.

More recently, a new technology has emerged as a disruptive technology known as hyperconvergence. We take an in-depth look at the overall cloud landscape and hyperconvergence in the following section.

Cloud Landscape

Now that we have talked about the current industry trends and identified key elements of software defined data centers, hyperconvergence and cloud computing, let's look at the cloud landscape in greater detail. It is extremely important to understand that different cloud platforms can either make or break your strategy. ***Not all clouds are created equal.***

Legacy Approaches

Build Your Own Cloud (BYO):

Some MSPs choose to build their own cloud using either open-source or proprietary solutions. In this case, the service provider is responsible for building the infrastructure and buying or building a cloud management platform. The issues with this approach are: 1) On-going capital expenditure, lengthy time to market, technical complexity, high risk and data center management overhead. Additionally, delayed revenue-generating services and profit margins of a meager 5-30% makes it a costly option.

Retail cloud:

Public cloud solutions are highly appealing to Managed Service Providers (MSPs). These instant-on services are elastic and only cost a few cents an hour, convincing some MSPs to resell retail cloud products. They can make a small margin by simply adding a markup building services by adding value through customer support and consulting services. But there are a number of other hidden costs and technological shortcomings that one should be aware of. Aside from the typical low margins (5-10%) and spotty support, the bigger issues are with loss of your own branding and the inherently weak blocking architecture of Cloud 1.0 with single points of failure and costly horizontal scaling. This results in an unreliable solution that incurs escalating operational and capital expenditures.

Upon due consideration, you will find that the true cost of public cloud appears dramatically high when you start adding the cost of predictable storage performance, high availability, backup, disaster recovery, private networking and more.

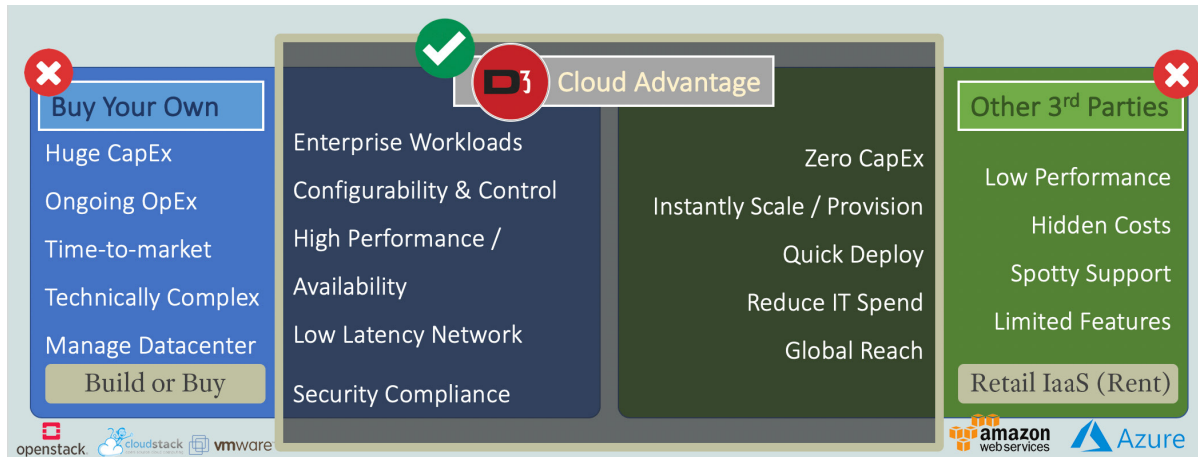
Disruptive Innovation - Hyperconverged IaaS:

Recently, a unique approach to the cloud has emerged that allows MSPs to deliver enterprise-grade infrastructure competency without the need to engineer, deploy and scale the solution yourself. Products and services are turnkey, can be white labelled and are easily integrated without heavy capital and staffing investments dramatically reducing the cost of ownership.

The biggest advantage to such an approach is to deliver enterprise-grade cloud solutions internally within their own-organization and the ability to brand it as their own (if so desired). For MSPs, it offers lucrative margins often between 30-60% which is considerable higher than what's possible with BYO or retail IaaS solutions.

- A turnkey business model, incorporating operational processes as well as supporting systems and tools.
- Clients acquire business know-how wholesaler has accumulated, includes valuable lessons learned & pitfalls to avoid.
- Immediate access to latest innovations.
- MSPs are responsible for updating the technology that drives the cloud services being sold. Helps clients reduce in-house R&D costs and leverage latest/greatest tech advancements in a proven, market-tested & market-ready platform.
- Global marketing support.
- MSPs provides sample marketing collateral templates, shares promotional strategy and market positioning advice, which clients can adapt for their own local market conditions and specific competitive environments.

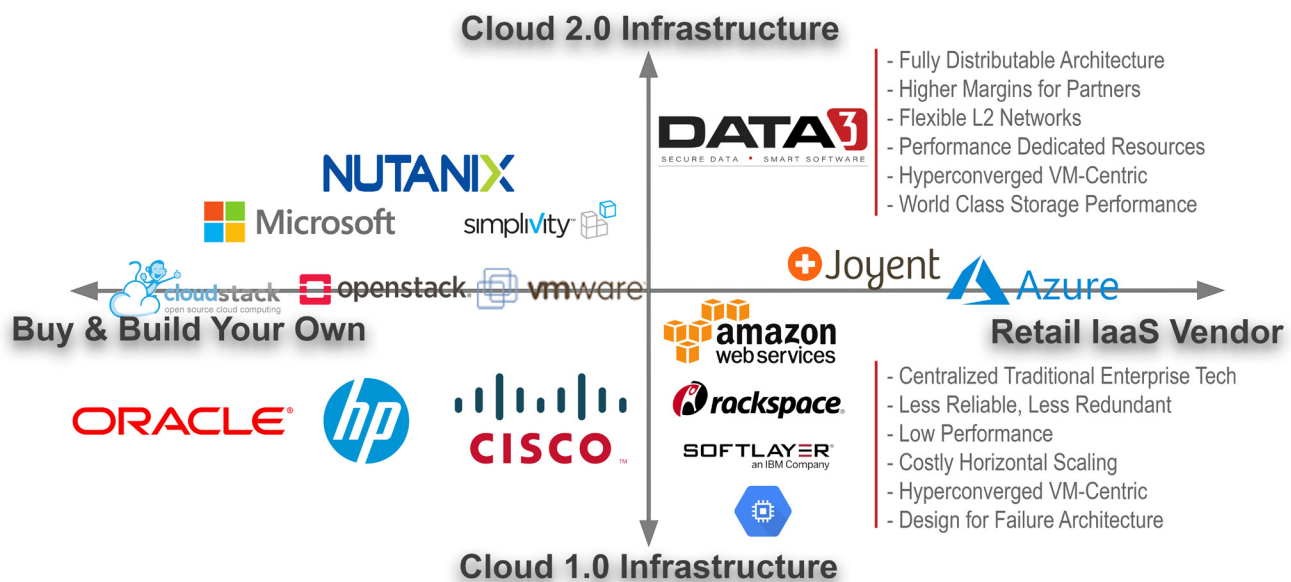
The below graphic shows how hyperconverged cloud for MSPs, as a new third option, combines the advantages of both BYO cloud and retail IaaS while eliminating the disadvantages:




The DATA3 cloud is powered by industry-leading technology and gives MSPs the ability to price, brand and customize the platform to suit their business needs. It comes with out-of-the-box usage tracking and rate-card support along with industry-first N-tiered branding and pricing functionality allowing for multiple levels of resellers and end users; companies even have the ability to use this feature for internal on-cost billing or use tracking.

There are large differences between platforms in terms of how complex, time consuming and expensive it is to

create core cloud services. DATA3's over 20 years of experience and its ability to recognize industry trends and best available technology has allowed it to build the next-gen "The D3 Cloud". Its 2.0 cloud infrastructure is based on state-of-the-art hyperconverged platform. This offers the fastest, easiest and least expensive path to producing these cloud services. MSPs success will not simply be determined by the features or capabilities of a cloud platform but also the set of services that can be directly monetized. It puts DATA3 in the top-right of the quadrant (see below).





Cloud 2.0 - Hyperconverged Architecture



The whole concept behind Cloud 2.0 is that there are challengers emerging in the space who are building breakthrough hyperconverged technology. Incumbents of Cloud 1.0 have architectures that are wedded to less efficient business models. Even if they wanted to innovate, these companies would have great difficulty in doing so,

and the effort would jeopardize current revenue streams. Cloud 1.0, especially the retail model, relies on centralized, traditional enterprise technology. This "design for failure" architecture uses costly horizontal scaling, and performance suffers due to limited redundancy and lesser reliability.

The below graph illustrates the architectural design differences between cloud 1.0 and cloud 2.0.



Cloud orchestration and virtualization

DATA3's hyperconverged cloud uses the industry standard Xen Hypervisor to run clients' virtual machine instances. Xen is a microkernel hypervisor that allows you to run multiple instances on the same physical server, all with complete isolation and security. On most instances, we take advantage of a specific mode in Xen called PVHVM – which stands for Paravirtualised Drivers on HVM. This delivers the best of both worlds – the operating system support and CPU performance of HVM (Hardware Virtual Machine) with the Network and Disk performance of a PV (Paravirtualised) driver. Utilizing Xen, we are able to run almost any operating system within our environment, including Linux, Windows Server and FreeBSD.

Rather than work around the issues relating to traditional IaaS systems, or retrofit existing infrastructure to fit as others have done, we redesigned the whole system from scratch, entirely removing these issues from the equation.

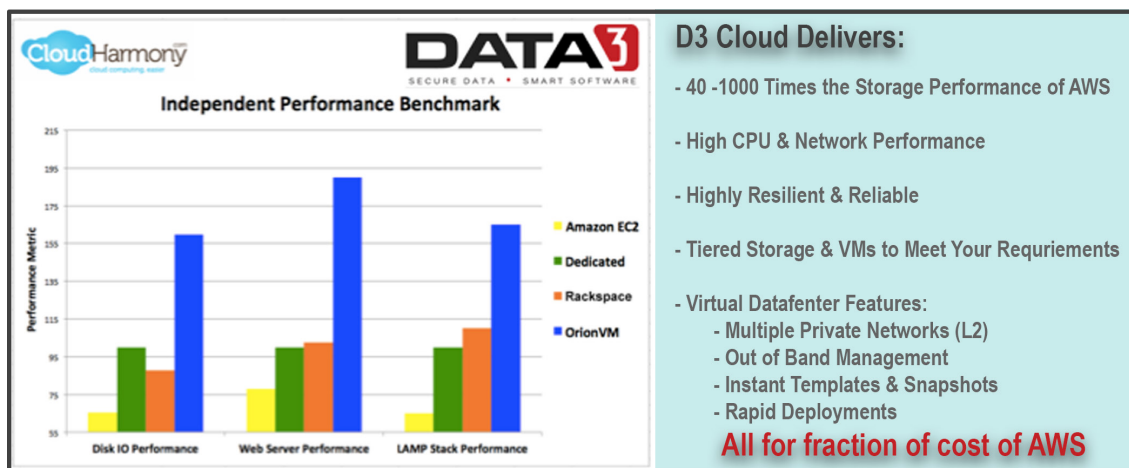
Firstly, our cloud infrastructure platform is built around a decentralized, distributed and hyperconverged architecture. We utilize a homogeneous node architecture,

where each node has compute, memory and storage assets. By utilizing homogeneous servers we are able to deliver greater performance at a reduced cost.

Integrated storage

Traditionally, using direct attached storage has introduced issues such as reduced flexibility and redundancy, but DATA3 has overcome this issue by utilizing a distributed storage architecture over InfiniBand, creating a virtual SAN across our cluster. This architecture reduces latency, is more cost effective, and mitigates the points of contention and failure that Centralized SANs can cause. The virtual disk of each server can be mounted on any physical server (allowing us to move the instances around the cluster) and is also replicated across multiple servers for redundancy.

InfiniBand was designed specifically for low latency and sustained data transfer, making it perfect for a high performance cloud to be built on. Combined with multiple storage performance tiers, customers are able to mix and match SSD with Spinning Disk to get the perfect performance/price mix for their requirements. Benchmarks and figures aside, it has also allowed us to design and develop some great features that regular Ethernet and Xen would not have afforded.



High-availability by design

With this redundant, decentralized InfiniBand fabric, we are also able to continuously replicate storage across several nodes while maintaining performance. It allows us to distribute storage efficiently across our cluster, as well as remove hotspots without impacting high performance for customers. As for your instances, having network-backed storage allows us to automatically migrate your instances to different nodes in the event of a physical server failure, rather than having to wait for a technician to enter a datacenter and perform recovery operations on a physical node. This design drastically reduces downtime, simplifies the deployment of instances, and provides more immediate feedback to customers.

Software-Defined Networking (SDN)

By default, DATA3 deploys two networks for the customer environment:

Public Network


This is a network connected to the public Internet, either by DATA3 or partner supplied connectivity. Using DATA3 IP addresses is the easiest way to get started. We utilize a multi-homed mix of providers and run our own BGP network (Routers/Switches/AS). Our upstreams are a mix of Tier 1 and Tier 2 carriers, with at least 10 gigabit connections. We utilize a fully HA network with router, switch and gateway redundancy from the VM through to the Internet. A public network is not mandatory and can be turned off for private cloud deployments.

Private Network

All accounts are by default set up with one private network. A private network is a Layer 2 network that connects all VM's within your account together to allow for East-West Traffic. This is a fully segregated and private network that isn't visible to any other customer. It is just like having your own switch with your customers connected to it. You are able to provision as many private networks within an account as required to fit your network architecture. As this is your cloud, you can define it as either Public/Private or Hybrid.

Additional DATA3 capabilities

- Connect private networks to cross connects within a datacenter at L2, or interconnect with a MPLS or Metro-E provider to deliver a private tale to your existing office/infrastructure.
- A network appliance can be utilized to provide a NAT gateway for internal infrastructure, as well as to perform VPN termination, Load Balancing and Firewall duties.
- Hybrid Cloud - we are able to place our platform on the same network as an existing colo/private cloud deployment to extend existing capacity.
- DATA3's composable platform also allows you to leverage existing network connectivity and data center capacity, our networks are very flexible and can be used to fit most network architectures and deployments.
- Service providers don't need to worry about coding, integrating modules, and completing complex configurations, DATA3 makes it easy to integrate with existing billing systems, so providers can rapidly start selling cloud services.
- DATA3 is easy to deploy and provides extensive support throughout the process. Using a combination of turnkey features and DATA3 support, some customers have achieved a production-ready cloud in just two weeks.



Core Cloud 2.0 - Revenue

By offering cloud resources to your customers new business opportunities will emerge, enabling you to add additional revenue to your existing business and prevent customer loss. You will be able to concentrate on your core competencies, confident in knowing that your infrastructure is powered by the best performing and most reliable platform on the market. Rapid time-to-market is also key in retaining existing customers and winning new ones.

Perhaps most importantly, you will be able to capitalize on the vastly improved margins - why give away revenue by moving your customers onto retail cloud platforms that provide low margins and de-value your brand? Despite what some analysts and CMOs might have you think, the dawn of Cloud 2.0 brings great opportunity and margins for IaaS resellers. MSP can look forward to creating the following revenue streams:

COMPUTE

- **Multitenant Virtual** - Servers are divided into multiple, isolated compute instances using server virtualization. Separate cloud customers (tenants) use compute instances from the same underlying server, without access to each other's applications or data.
- **Bare Metal Servers** - Customers provision servers that have not been virtualized, and can use the hardware any way they choose.

NETWORK

- **Load Balancing** - Customers can specify a group of servers that will share the processing of incoming requests. This helps eliminate performance bottlenecks, and is good for large workloads with fluctuating performance requirements.
- **Firewalls** - Firewalls provide customers with protection by monitoring network activity and intercepting security threats and unapproved traffic

STORAGE

- **Block Storage** - Block level storage lets customers create raw storage volumes that can be used for almost any type of application. Servers connect to these raw storage volumes and use them the same way they would an individual hard drive
- **Archival Block Storage** - Using magnetic drives as opposed to SSD, sacrificing performance for cost savings, perfect for archival storage.
- **CDN** - Content delivery networks allow customers to meet performance and availability objectives by distributing applications and workloads across multiple, geographically dispersed servers that reside closer to content consumers.
- **DNS Management** - This allows customers to establish and manage a namespace for their domains, or the domains of their customers.

Conclusion

DATA3 is 100% channel focused powering organizations to differentiate their services and out-compete public cloud and BYO solutions on everything from the base compute, networking and storage resources to the total end-to-end solution delivery.

In this whitepaper we have explained why moving workloads to the cloud is inevitable and how MSPs can capitalize on this shift in infrastructure consumption by offering their own cloud services. We also explained why MSPs are perfectly placed to add cloud services to their offerings and deliver greater value to their customers.

The first section of the white paper focused on how MSPs can create cloud revenue streams, along with how choosing the right platform will enable the fastest and most effective route to profitability. We analysed the two most common paths to the cloud and recognised the need for a third option, the reason behind DATA3's hyperconverged cloud platform. The specific revenue streams that MSPs have the potential to create have been called out, along with the effects they will have for the MSPs business.

In the second section we looked at the industry technology trends that have forged the way for DATA3's hyperconverged cloud, these included the SDDC, Cloud and its predecessor virtualization, and convergence/hyperconvergence. Organizations are now wanting to take advantage of the economics and scalability characteristics of webscale environments like Facebook and Google. This whitepaper has shown how MSPs can make the most of these differences with the DATA3's hyperconverged cloud platform.

The third and last section was a technical deep dive into how the DATA3 cloud platform works. Key differentiating

Availability by design and Software Defined Networking to deliver industry leading performance/price ratios and sustainable cost advantages.

With the demand for these solutions expected to rise significantly in the coming years it seems likely that MSPs will dramatically alter the landscape of cloud providers, redefining how and by whom cloud services are sold. As IT departments are continually being directed more toward the business and less toward the technology, the need and desire to turn to solutions like cloud and hyperconverged infrastructure will cause MSP firms to simplify their datacenter and redistribute workloads in new and different ways.

About DATA3

DATA3 is a Tulsa based company at CityPlex Towers that has been in business since 2001. The company provides the following services:

- Datacenter
- Business Continuity Planning
- Software Development

The company operates three datacenters with a combined datacenter space of over 18,000 sq. ft. with another 35,000 sq. ft. available for business continuity services. The company also offers custom software and mobile development using its proprietary Enterprise Focused Software (EFS) strategy. Its web portals and Single Sign-On solutions are accessed by customers from over 35 different countries.

Some of its customers include American Airlines, ONEOK Energy, Pop TV (previously known as TV Guide), AAR Corporation and BAE SYSTEMS.

www.datathree.com

References for action

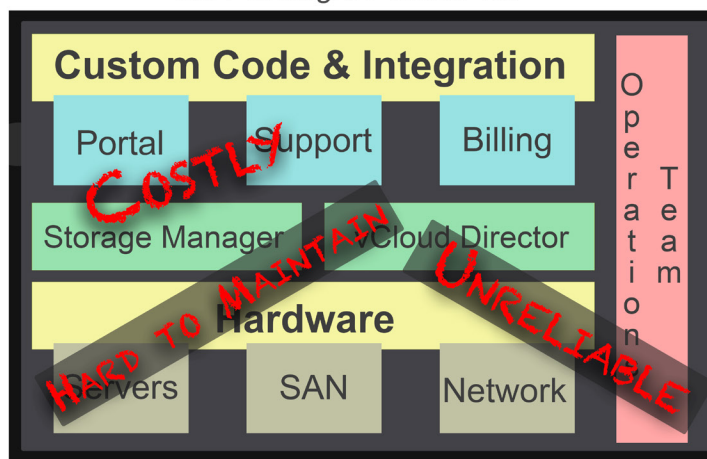
If you would like to further evaluate the DATA3 platform you can:

a) Contact our sales team via www.datathree.com/msp-partners-program. They will be happy to discuss any particular details, give you a demonstration of the platform and provide a free 30 day trial account; or

b) Register for a free technical consultation at www.datathree.com/contact-us, and we will assign an engineer and implementation consultant to work directly with your technical and business line managers. In addition to aiding in the evaluation of the platform, they will complete a review of your current technology stack and business processes. Objectives will be identified and agreed upon so that a detailed report including tasks, costs and timeline can be given for achieving the objectives.

Traditional Enterprise

Old Paradigm / Cloud 1.0



Hyperconverged Cloud

Next Gen Cloud 2.0 - IaaS

