

Creating cloud revenue streams for system integrators

WHITE PAPER

A system integrator's guide 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
Historical background.....	4
Current industry technology trends.....	5
SDDC, Cloud and Hyperconvergence.....	5
Hyperconverged appliances vs. the hyperconverged platform.....	7
Cloud landscape	7
Legacy approaches	7
Disruptive innovation	8
Hyperconverged Wholesale IaaS	8
Cloud 2.0 hyperconverged architecture	10
Cloud orchestration and virtualization	11
Integrated storage.....	11
High-availability by design	12
Software-defined networking (SDN).....	12
Core Cloud 2.0 revenue streams	13
Conclusion.....	14
References for action.....	15

Executive Summary

The inevitable cloud

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. They are consequently expecting system integrators to provide solutions that not only are viable business tools but also give them a technological edge over competitors. Furthermore, and even more importantly, corporations are becoming increasingly aware that these solutions have to be agile enough to deal with changing business requirements and disruptive technological innovation, which is steadily gaining momentum.

Enterprise IT has therefore 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 two legacy approaches are in fact failing to resolve some of the fundamental issues that supposedly should have been eliminated by cloud-based technology.

Building your own (BYO) cloud is really only an option available to those with vast resources, and even for such corporations, the costs outweigh the advantages. Large initial and ongoing capital expenditures combine with technical complexity and data center overheads to make for a very risky proposition that due to the lengthy time to market might well be dated even before it has been fully implemented. It is very far from the agile, turnkey solution cloud-based services are heralded as, and it offers system integrator profit margins of only 5-20%.

The traditional alternative has been to purchase retail Infrastructure as a Service (IaaS), but this also comes with a number of disadvantages. Among the more obvious ones are very low reseller margins and spotty support. Worse still is the sacrifice of branding power and the limited scope for technical differentiation as well as the inherently weak blocking architecture of Cloud 1.0, with single points of failure and costly horizontal scaling. The

result is an unreliable solution that incurs escalating operational and capital expenditures without providing neither the system integrator nor the client with the desired competitive edge, agility or profit margins. For incumbents of Cloud 1.0, this is particularly worrisome since they cannot innovate without jeopardizing their current revenue streams.

Recently emerged technology stands in stark contrast to the aforementioned two legacy approaches, however. Hyperconverged Wholesale IaaS is everything that BYO and retail IaaS is not. It utilizes the web-scale power of a Facebook-like cluster to enable system administrators to rebrand and resell infrastructure with lucrative 30-60%+ profit margins whilst safeguarding the efficiency, agility and internal compliance needs of their clients.

Although it is commonly referred to as Cloud 2.0, hyperconverged wholesale IaaS is actually the first service that realizes the full potential of cloud-based technology. Including features such as per-second billing, live snapshotting, and multi-tier reseller and org management, its software-centric design has been built from the ground up rather than working around the issues in traditional IaaS solutions or retrofitting existing infrastructure. This has resulted in far superior scope for differentiation, scalability, control- all at less than 80% of the cost of a legacy Amazon Web Services (AWS) solution.

To IT departments, it 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. To system integrators, it presents a novel white-label opportunity to add value over and above management of servers and databases. Enabling seamless adding of other on-demand IT services all the way up the stack, Cloud 2.0/wholesale IaaS is unsurprisingly starting to become known by yet another name.

IT as a service.

Historical background

System integrators have traditionally developed or deployed solutions onto dedicated hardware in a data center. Usually, they have procured hardware for a customer and made additional profits by reselling off-the-shelf components. In a legacy data center environment, growing can be expensive due to the proprietary nature of each individual piece of hardware. The more diverse the environment, the more difficult it is to maintain.

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 many resources are wasted in today's legacy environments.

Prompted by the above issues, enterprise IT has moved away from siloed servers, storage, information, and processes and is now embracing converged, virtualized, and cloud-based technologies. This transformation and the core technologies behind it are driving changes in how businesses consume software and infrastructure.

Virtualization revolutionized the enterprise's use of infrastructure in the previous decade. Most organizations ultimately moved to the cloud for the same reasons that they turned to virtualization, enjoying additional benefits like paying only for what you use, eliminating

capital expenditure in favor of operational expenditure, on-demand capacity expansion and spending less time on IT and more time on business. The cloud represents layers of new technology combined with a revolutionary business model.

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 businesses simply must make sooner or later. However, as we shall see, legacy cloud platforms have until recently only offered systems integrators partial capitalization on the technological potential of the cloud and little or nothing in the way of profit margins.

Current industry technology trends

The IT infrastructure market is undergoing unprecedented transformation. The most significant transformation is reflected by three major trends: Software-defined data centers (SDDCs), hyperconvergence and cloud computing. These trends are responses to the IT realities of infrastructure clutter, complexity and high cost; they represent attempts to simplify IT and reduce the overall cost of infrastructure ownership.

SDDC, Cloud and Hyperconvergence

The hallmarks of Google's and Facebook's environments are, among other things, sheer scalability and reasonable economics. Many of these cloud principles have been adapted for use in smaller environments and packaged in hyperconverged products that any company can buy. Factors driving toward such hyperconverged cloud technology include:

Software-centric design

Companies like Google discovered the potential of the software-defined data center (SDDC) years ago and tamed its hardware beast by wrapping it inside software layers.

Commodity hardware and seamless changes

Change is constant in the cloud, so ensuring that changes are made without disruption is critical. In the world of system integration, things should work the same way. A change in data center hardware shouldn't necessitate reconfiguration of all your virtual machines and policies. Google and Facebook are good examples of this, scaling their environments without relying on expensive proprietary components. Instead, they leverage commodity hardware.

To some people the word commodity is a synonym for cheap or unreliable, and that is correct to a point. When you consider the role of commodity hardware in a hyperconverged environment, you should however keep in mind that the hardware takes a back seat to the software layer, which is built with the understanding that the hardware eventually will fail. Software-based architecture is designed to anticipate and handle any

hardware failure, and by using commodity hardware that is readily interchangeable with other components, a hyperconvergence vendor can switch its hardware platform without recoding the entire system. This ensures that their customers can utilize commodity hardware without risking disruption to services.

Bite-sized scalability

Hyperconverged infrastructure takes a bite-sized approach to data-center scalability. Customers no longer need to expand just one component or hardware rack at a time; they simply add another node to a homogenous environment.

VM-centricity

The workload takes center stage in the cloud. In the case of system integration, these workloads are individual VMs. When it comes to policies in cloud-based environments, the VM is the center of the world. It's all about applying policies to VMs — not to LUNs, shares, data stores, or any other constructs. Bear in mind the plight of the VM administrator, who is VM-centric. Why wouldn't the administrator assign backup, quality-of-service, and replication policies to a VM?

SI's applications are usually designed expecting high-availability infrastructure, disaster recovery, backup, recovery and a range of other necessary services that place a different set of demands on the infrastructure. Therefore, any hyperconverged IaaS platform must deliver on these requirements. Hyperconvergence brings cloud-type consumption-based infrastructure economics and flexibility to system integrators without compromising on performance,

reliability, and availability. Rather than making huge buys every few years, IT simply adds building blocks of infrastructure to the data center as needed. This approach gives the business much faster time-to-value for the expanded environment.

Hyperconvergence delivers simplification and savings by consolidating all required functionality into a single infrastructure stack running on an efficient, elastic pool of x86 resources. The underlying data architecture has been completely reinvented, allowing data management to be radically simplified. As a result, hyperconverged infrastructure delivers on the promise of the Software Defined Data Center (SDDC) at the technological level. It also carries forward the benefits of convergence, including a single shared resource pool. Hyperconvergence goes far beyond servers and storage, bringing into the convergence fold many services that make some legacy services obsolete, including data protection products (backup, replication), deduplication appliances, wide-area network (WAN) optimization appliances, solid-state drive (SSD) arrays, SSD cache arrays, public cloud gateways, and replication appliances or software.

A recent study from IDC shows the most common drivers for business adopting converged infrastructure:

Integrated and Converged Systems Drivers

<p>Accelerate Business</p> <ul style="list-style-type: none"> • Reduce time to revenue • Manage thru peaks • Implement data analytics • Cloud ready 	<p>Lower Costs</p> <ul style="list-style-type: none"> • Improve app density • Improve data density • Reduce power and cooling • Cut life-cycle maintenance
<p>Simplify IT</p> <ul style="list-style-type: none"> • Consolidate app/data • Simplify maintenance • Accelerate IT deployment • Save time 	<p>Reduce Risk</p> <ul style="list-style-type: none"> • Biz continuity • Backup/DR • Security • Regulatory compliance

Source: IDC, 2014

Advantages to these solutions are the relatively simple design for a complex infrastructure environment (of course, normally, it could take months to design an infrastructure, using a mix of best-in-class technologies). In addition to the simplified architecture, there's also a simplified administration model. The hyperconverged systems are managed via "a single pane of glass." Instead of having a set of applications and a team to manage your storage array, a team to manage virtualization and a team to manage the server hardware, one team (or, in some environments, a single person) can manage the complete hyperconverged stack.

Another IDC research study indicates that an IT operations team typically spends over 70% of their time on day-to-day IT management operations — monitoring, troubleshooting, patching, updating, and configuring resources. This leaves little time for them to add value to the business or support new innovation. Hence, technical leaders and data center managers are looking for robust solutions to help streamline these operational inefficiencies.

Total Cost of Support by Architecture Type



Source: IDC, 2014

Hyperconverged appliances vs. the hyperconverged platform

Since the popular industry term hyperconvergence hitherto most commonly has been used when referring to hyperconverged appliances, it should be underlined that these appliances are distinctly different from the hyperconverged platform.

Hyperconverged appliances are characterized by a number of disadvantages, including the inability to make granular upgrades or tweaks to the system. Storage growth and performance tuning are also pain points for most organizations. If a cluster runs low on storage but not compute, you must still upgrade overall compute capability by adding another appliance. Likewise, if you want to tune storage disk configuration for a particular application, this can be a challenge.

In the past the remedy for these drawbacks was to build-your-own platform, for many the risk, lead-time and cost of doing so was inhibitive, causing most to wait for solutions to mature and cover their platform requirements.

In contrast to the hyperconverged appliances, there is today a full-featured hyperconverged platform capable of resolving all of the aforementioned issues: the most recent addition to the cloud landscape.

Cloud landscape

The right choice of cloud management platform can do much more than help system integrators quickly get to market with a core set of competitive services; it can also provide flexibility to differentiate and layer additional services as needed. ***It is however 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 SI 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. Common issues with this strategy include large initial

and ongoing capital expenditures, lengthy time to market, technical complexity, high risk and data center management overhead.

Depending on your capabilities, resources and timeline, there is a good chance that you'll struggle with the initial stages of deploying, stabilizing, and integrating it with your current systems before you can finally start delivering revenue-generating services, at which point the meager profit margins will be in the 5-30% range.

Retail Cloud

Public cloud solutions are highly appealing to system integrators. These instant-on services are elastic and only cost a few cents an hour, convincing some SI to resell retail cloud products. They can make a small margin by simply adding a markup, bundling services in unique ways or adding value through customer support and consulting services.

However, to gain substantial market shares and achieve healthy margins, you must balance revenue streams, time to market, service differentiation, pricing strategy and performance factors. Upon due consideration, you will then find that the true cost of public cloud appears dramatically high when compared against the cost of predictable storage performance, high availability, backup, disaster recovery, private networking, and more. IT ends up paying for a server that is running at 15 percent utilization, and the cloud provider is benefiting from packing those VMs onto a single host.

You will also realize that retail Cloud comes with a number of other hidden costs and technological shortcomings. Among the more obvious ones are very low reseller margins (typically between 5-10%) and spotty support. Worse still is the sacrifice of branding power and the limited scope for technical differentiation as well as the inherently weak blocking architecture of Cloud 1.0, with single points of failure and costly horizontal scaling. The result is an unreliable solution that incurs escalating operational and capital expenditures without providing neither the system integrator nor the client with the desired competitive edge, agility or profit margins. For incumbents of Cloud 1.0, this is particularly worrisome since they cannot innovate without jeopardizing their current revenue streams.

Disruptive innovation

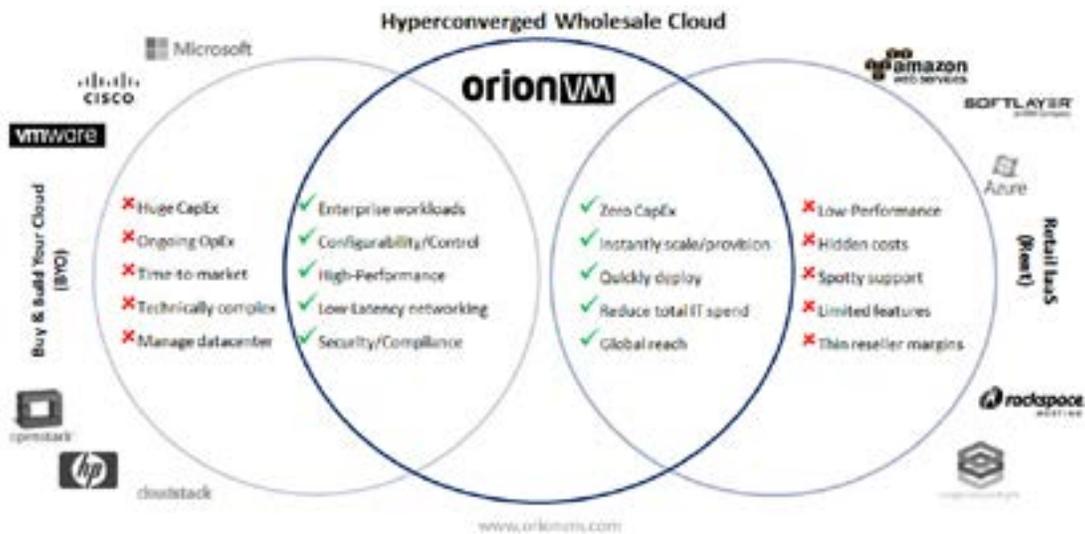
Hyperconverged Wholesale IaaS

Recently, a unique approach to the cloud has emerged that allows for a wholesale business model. Wholesale cloud delivers enterprise-grade infrastructure competency without needing 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, all of which dramatically reduce total cost of ownership.

A wholesale cloud platform enables companies to resell their own branded enterprise-grade cloud solutions or deliver cloud services internally within their own organization. Most attractively, wholesale cloud delivers lucrative margins often between 30-60%, which is considerably 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.**
 - Wholesaler is 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.**
 - Wholesaler 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.

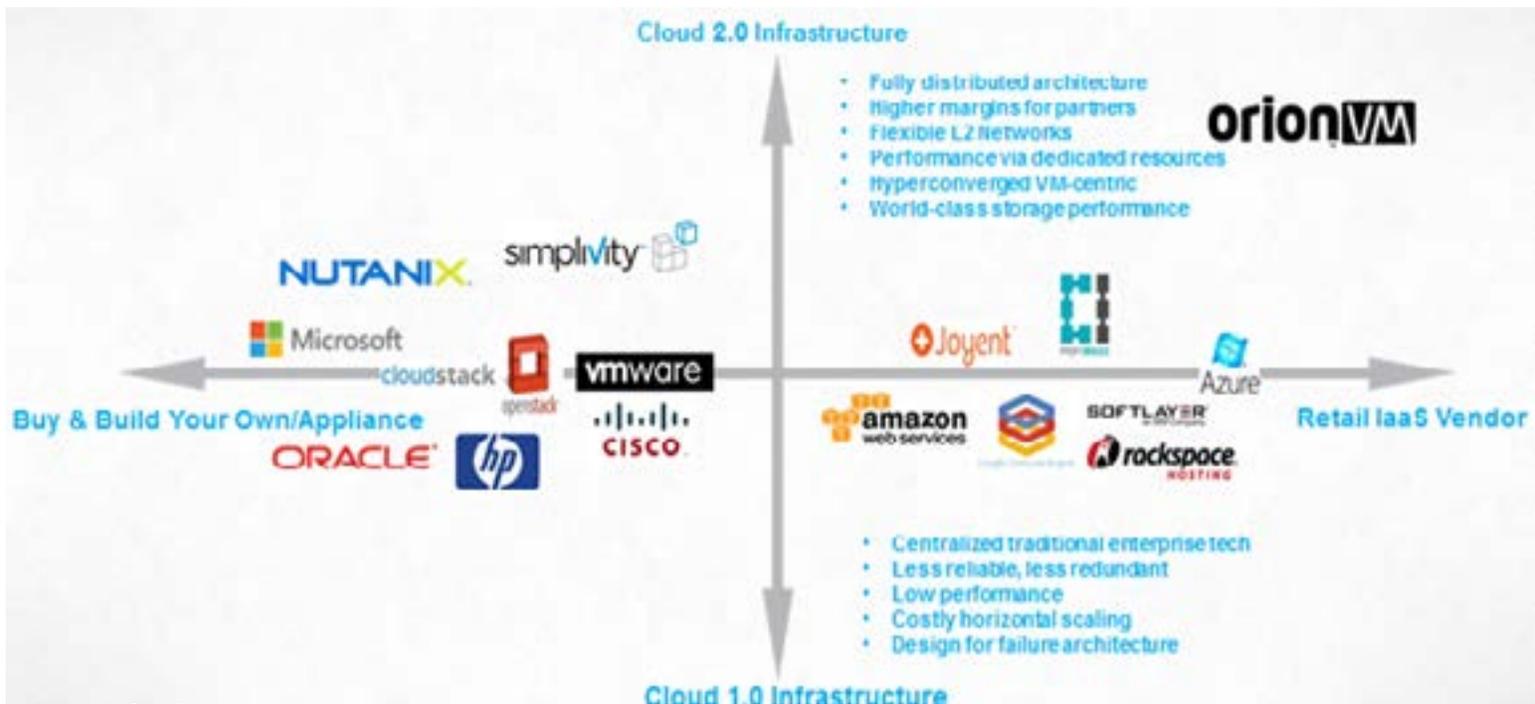
The below graphic shows how hyperconverged wholesale cloud, as a new third option, combines the advantages of both BYO cloud and retail IaaS while eliminating the disadvantages:



OrionVM's wholesale cloud is powered by industry-leading technology and gives SIs 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. OrionVM offers the fastest, easiest and least expensive path to producing these cloud services. SI 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.

Below is an industry map showing where some of the common solutions sit with reference to common paths and the technology stack powering them:



Source OrionVM

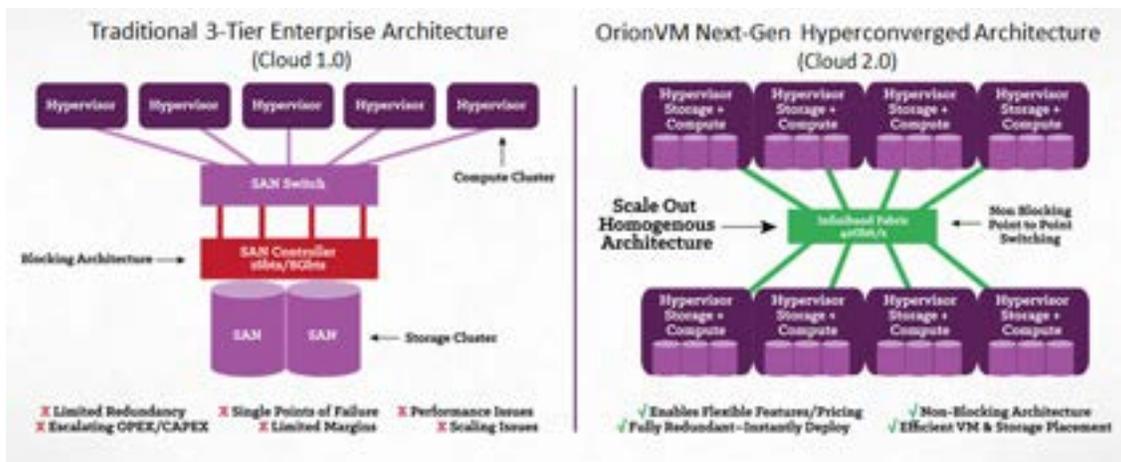
Cloud 2.0 Ring of Innovation



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

OrionVM's hyperconverged cloud use 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.

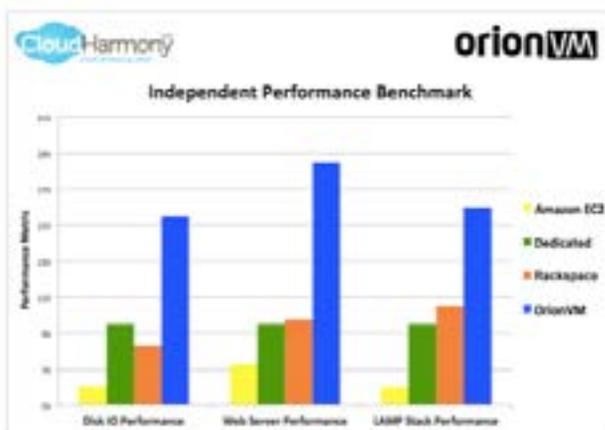
Firstly, our cloud infrastructure platform is built around a decentralized, distributed and hyperconverged architecture. We utilize a homogeneous node architecture,

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 OrionVM 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.



OrionVM Delivers:

- 40 – 1000 times the storage performance of AWS
- High CPU & Network Performance
- Highly resilient & reliable
- Tiered Storage & VMs to meet your requirements
- Virtual Datacentre Features:
 - Multiple Private Networks (V2)
 - Out of Band Management
 - Instant Templating & Snapshotting
 - Rapid deployment
- All for the fraction of the cost of AWS.

www.orionvm.com

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, OrionVM deploys two networks for the customer environment:

Public Network

This is a network connected to the public Internet, either by OrionVM or partner supplied connectivity. Using OrionVM 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 OrionVM 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.
- OrionVM'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, OrionVM makes it easy to integrate with existing billing systems, so providers can rapidly start selling cloud services.
- OrionVM is easy to deploy and provides extensive support throughout the process. Using a combination of turnkey features and OrionVM support, some customers have achieved a production-ready cloud in just two weeks.



Core Cloud 2.0 revenue streams

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 cloud tech 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. System integrators 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

OrionVM 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 SIs can capitalize on this shift in infrastructure consumption by offering their own cloud services. We also explained why SIs 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 SIs can create cloud revenue streams, along with how choosing the right platform will enable the fastest and most effective route to profitability. We analyzed the two most common paths to the cloud and recognised the need for a third option, the reason behind OrionVM's hyperconverged wholesale cloud platform. The specific revenue streams that SIs have the potential to create have been called out, along with the effects they will have for the SI's business.

In the second section we looked at the industry technology trends that have forged the way for OrionVM's hyperconverged wholesale IaaS cloud (Cloud 2.0), these included the SDDC, Cloud and its predecessor virtualization, and convergence/hyperconvergence. We explored the CAPEX/OPEX benefits of utilizing hyperconverged infrastructure as an IaaS platform vs. purchasing it as an appliance. 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 SIs can make the most of these innovations with the OrionVM hyperconverged wholesale cloud platform.

The third and last section was a technical deep dive into how the OrionVM platform works. Key differentiating technical features were explained including how we manage cloud orchestration and virtualization, integrated storage, High Availability by design and Software Defined Networking to deliver industry leading performance/price ratios, enterprise-grade stability and resilience, and sustainable cost advantages.

With the demand for these solutions expected to rise significantly in the coming years it seems likely that system integrators 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 SI firms to simplify their datacenter and redistribute workloads with maximum efficiency, on-demand flexibility, and unprecedented cost-savings and ROI.

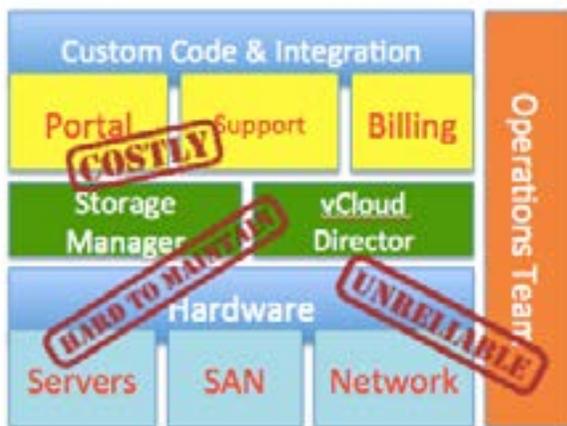
References for action

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

a) Contact our sales team via www.orionvm.com/partner-with-us/. They will be happy to discuss any particular details, pricing info, give you a free live demonstration of the platform. and provide access to a free introductory trial account; or

b) Register for a free technical consultation at www.orionvm.com/consulting-offer, 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

Traditional Enterprise (old paradigm/cloud 1.0)



OrionVM Hyperconverged cloud (next-gen cloud 2.0 IaaS)



Replace multiple services/vendors with 1 Integrated system.

www.orionvm.com