



THE PROFESSIONAL- GRADE CLOUD.

**Zunicore: A Purpose-Built Cloud
for Agile Businesses**

*BY GREG RUSU (GENERAL MANAGER) AND
DAVE GEADA (DIRECTOR OF PRODUCT MARKETING)*

TABLE OF CONTENTS

Executive Summary	1
Technological Advances Force a Business Paradigm Shift	2
Cloud 1.0 Has Not Responded to Business Needs	2
The Causes of Cloud 1.0's Inadequacy for Mainstream Businesses	3
Inflexibility	3
Unreliable Scaling	4
Faulty Security	4
Poor Visibility and Cost Controls	4
Enter Professional-Grade Cloud, Improving Cloud Dramatically	4
Professional-Grade Cloud is Flexible	5
Professional-Grade Cloud is Cost Efficient	5
Professional-Grade Cloud is Reliable and Secure	5
Professional-Grade Cloud Enables Control	5
Zunicore: A Professional-Grade Cloud Backed by Superior Delivery	5
Professional Grade with Unmatched Network Infrastructure	6
Conclusion: Switching Has Never Been Easier	7
Bios	8

EXECUTIVE SUMMARY

While cloud computing has shown strong market growth in recent years, it has served a limited set of business customers that could benefit from its one-size-fits-all approach. Many businesses across a diverse range of industries, such as marketing agencies, SMB product and financial services companies, professional services firms, and gaming companies have very specific needs when it comes to the cloud. As such, they have not been able to benefit as much from the cloud's promise of providing highly scalable, instant-on, and cost-effective computing resources to serve a global audience.

For these businesses to gain from the cloud, they require a cloud service offering that allows them to work with their existing people, processes, and technology infrastructure as they deliver services to their own customers over the Web.

A new type of cloud service has emerged that provides businesses with the four critical elements they seek today: control, flexibility, reliability, and cost efficiency. This new type of cloud, **professional-grade cloud**, has been developed specifically for businesses with unique and changing technology requirements as they bring varying types of campaigns to the Web.

This paper will discuss how Cloud 1.0 has fallen short, and how **professional-grade cloud** offers cloud benefits to a wider group of business types today. It will detail the services provided by a professional-grade cloud leader and pioneer, Zunicore, a division of PEER 1 Hosting.

Technological Advances Force a Business Paradigm Shift

Businesses have had to dramatically adjust their operational models to keep pace with technology changes, especially those sparked by the Internet and related capabilities. With the near ubiquity of social media sites like Facebook and Twitter, many business sites can rapidly attract more traffic than they can handle. Rich media sites with the latest “must-see” videos; financial sites during stock market fluctuations; news sites delivering breaking news; and ecommerce sites selling much-awaited products all can unexpectedly lose control over their demand flow. Business professionals at these companies focus on the business, not on the technology that enables them. Yet many find themselves responsible for both generating demand and fulfilling it, suddenly having to ensure that their site doesn't crash during changes in demand flow.

Unfortunately, cloud 1.0, Managed Hosting, Dedicated Hosting, Colocation, and In-House Hosting were not designed to address this issue. They have also failed to help businesses address the risks of unpredictable costs resulting from greater than expected demand. As the levers of demand generation become harder to control on today's Internet, businesses need solutions to help them prepare and respond to varying traffic volumes in an effective and profitable manner.

The next generation of cloud – or Cloud 2.0 – allows businesses to address unexpected spikes in demand by more effectively managing corresponding traffic volumes. Businesses can leverage a cloud or hybrid cloud environment to run targeted applications or functions, and then decide for themselves how their solutions will scale in response to varying demand levels. Up-to-the-minute visibility also shows the effects that different levels of scaling have on the bottom line. In this way, Cloud 2.0 puts control and flexibility back into the hands of businesses and enables them to realize significant budget and resource savings. In this new paradigm, both business and IT leaders can now focus their time and talent on revenue-generating projects. They become far more nimble than competitors who must waste excess time, money, and talent using antiquated IT models.

Through Cloud 2.0, true competitive advantage comes from the IT infrastructure itself. With the ability to scale resources seamlessly and instantaneously, businesses can proactively pursue market opportunities rather than reacting to environmental and business conditions. This provides critical time-to-market advantage over competitors. The cost savings from Cloud 2.0 enable businesses to invest more in other areas of the business such as R&D, marketing, and sales to maintain market leadership. In this way, Cloud 2.0 fulfills the promise of cloud which cloud 1.0 could not deliver for all but a select few business models.

Cloud 1.0 Has Not Responded to Business Needs

For many organizations, the advent of the early cloud technologies (known as cloud 1.0), signified an opportunity for business and IT professionals to easily manage all types of demand-generation scenarios. Gone were the days (or so they thought) that their Web infrastructure could no longer manage demand. However, for many businesses, leveraging cloud services feels like trying to fit a square peg into a round hole. They want to obtain the highly touted benefits of cloud but cannot seem to make it work within their current process frameworks. Cloud 1.0 service providers such as Amazon Elastic Cloud Computer (Amazon EC2), Salesforce.com, and Rackspace have largely ignored business complexities such as data reliability, cloud resource scaling, business flexibility, and cost control.

Business users have found improving cloud performance difficult. Their cloud 1.0 service providers maintain a one-size-fits-all cloud infrastructure that offers little visibility into the resources that customers are using and how they are using them. Without transparency, businesses must guess at architectural changes when targeting areas for application performance improvements. Often, the process of fine-tuning cloud performance can take months. Yet, companies that leverage the Web for revenue and lead generation simply cannot afford to wait months while they optimize their cloud architecture. This lack of business flexibility and responsiveness inherent in cloud 1.0 services causes many businesses to distrust the cloud.

The Causes of Cloud 1.0's Inadequacy for Mainstream Businesses

Understanding why cloud 1.0 service providers entered the market can illuminate the reasons behind cloud deficiencies. Service providers never meant for their cloud architecture to appeal to most businesses. Instead, cloud providers originally developed their offerings and underlying infrastructure to suit a narrow set of customers with like needs. For the most part, these customers include Web developers, small business entrepreneurs, and pure Web companies. Such businesses have very little of their own existing business processes and IT infrastructure to match with their cloud services. Also, the majority of these businesses have technologically savvy teams with the knowledge and resources to navigate cloud's technical complexities, making them ideal early adopters.

To serve the early cloud audience, Amazon EC2, Rackspace, Salesforce.com, and other cloud providers did not need to concern themselves with customer technology integration at a great level of detail. They merely had to provide application programming interfaces (APIs) to their cloud infrastructure. With just a few application developers, any small or purely Web startup company could create applications to run on cloud providers' technology building blocks (i.e., databases, storage, servers, etc.).

Despite the success of early cloud 1.0 for a limited market, mainstream businesses with existing infrastructure or existing solutions simply cannot use it. For the most part, their back-end technologies were not written to run against a configured, inflexible, structured, and vast cloud infrastructure. Not wanting to entirely re-engineer their companies, businesses struggle to make their existing technologies interact with the cloud. For most, success occurs infrequently or for a limited number of applications. In fact, according to The 2011 Future of Cloud Computing survey conducted by North Bridge, GigaOM Pro, and The 451 Group, "interoperability and vendor lock-in were the second and third biggest inhibitors to adoption of cloud computing."¹

To fully understand the vendor lock-in and interoperability issues of cloud that The 451 Group references, one must first recognize the inherent faults in the cloud's original approach. They fall into four main categories: inflexibility, unreliable scaling, faulty security, and poor visibility and cost controls.

Inflexibility

Cloud 1.0 customers often spend great effort performing the most basic business tasks, such as storing a single file. One of cloud's signature flagship offers, cloud storage, imposes an interface on customers that requires developers to write code to use it. The cloud was meant to simplify operations. Very few customers have personnel that would want to prescribe inputs and outputs just to store a file. This type of arcane inflexibility has held back user adoption.

¹ Network World, The cost of bad cloud-based application performance, March 2, 2011

Cloud is an Imperative, But Cloud 1.0 Needs Improvement

According to IDC, public cloud services will grow five times faster than traditional IT products, from over \$16 billion in 2009 to \$55.5 billion by 2014. However, many CIOs have serious concerns about cost monitoring/management.⁴

Unreliable Scaling

When writing applications that need to scale to cloud's potential, businesses need to have insight into the underlying technology architecture. Developers must have specific information, such as when to cache, when to retrieve data sets from the associated database, and when to write to the database. This type of information changes for each specific application implementation. Unfortunately, cloud 1.0 abstracts this information and keeps it hidden from application developers. As a result, they cannot easily write applications that scale well. Such poor performance costs businesses dearly. According to Network World's review of a recent Compuware study, "IT directors at 378 large enterprises in North America reported their organizations lost almost \$1 million annually due to poorly performing cloud-based applications."²

Faulty Security

Cloud 1.0 architectures do not decouple back-end services such as Web servers and database servers from the front end. This exposes cloud customers to the many security threats that can penetrate the front end. At the same time, businesses cannot use their own security measures and must settle for what the cloud provider offers. This drastically increases risk exposure, because application developers typically have less security expertise than the system administrators who would normally secure a business's back-end technologies. Even if system administrators are involved, cloud 1.0 security tools likely differ from the tools to which they have grown accustomed to using.

Poor Visibility and Cost Controls

Operating blindly against cloud 1.0's back-end systems, developers must continuously re-work their application code to optimize performance and the end-user experience. This represents a significant and unforeseen cost to cloud, particularly for small and medium-sized businesses with limited IT staff. SMB developers do not control cloud resources and have little administrator-level insight into system configurations and the ways in which systems run.

Even in scenarios where cloud 1.0 helps scalability, it frequently does so at extremely high costs. As traffic bursts and drops, cloud customers strain to control their costs. Considering the massive traffic volumes that crowd-sourcing tactics can produce (for example, Groupon promotions), successful marketing campaigns can quickly turn into disastrous events that break the bank. Moreover, marketers cannot easily predict their costs or match costs to each individual campaign. Businesses need the ability to turn resources on and off instantly. Equally important, they need reporting that is accurate, detailed, and timely so that they can plan effectively. Analyst Dennis Callaghan of The 451 Group characterizes the situation this way: "Organizations will want to know what kind of performance they're getting from cloud resources and how much that's costing them versus their internal infrastructures."³

Enter Professional-Grade Cloud, Improving Cloud Dramatically

Today's new professional-grade cloud architectures offer businesses far greater control over their cloud-based services. That's because professional-grade clouds provide transparency across the entire platform. As a result, business customers can interact with key back-end resources and set parameters as they see fit. This information availability underpins the major competitive advantage that professional-grade cloud offers over cloud 1.0 architectures.

³ Network Computing, Uptime First To Solve Cloud Cost Complexities, June 22, 2011

⁴ IDC, Worldwide and Regional Public IT Cloud Services 2010-2014 Forecast, June 2010

Intel's View: the Added Security of Professional-Grade Cloud

Recently, Intel highlighted the value of professional-grade cloud architectures. Intel stated, "Intel VT-d can help end users improve security and reliability of the systems and also improve performance of I/O devices in a virtualized environment. These help IT managers reduce the overall total cost of ownership by reducing potential down time and increasing productive throughput by better utilization of the datacenter resources."⁵

⁵ Intel, Intel® Virtualization Technology for Directed I/O (VT-d): Enhancing Intel platforms for efficient virtualization of I/O devices, 2009



Professional-Grade Cloud is Flexible

Unlike the one-size-fits-all approach behind cloud 1.0, professional-grade cloud business customers receive customizable resource pools, with no need to work their applications and processes around pre-defined virtual machines. As such, they can piece together virtual images that form one or more of their own unique clouds. These independent virtual clouds run over a public cloud architecture, but the ability to parse them from one another gives businesses the ability to treat each uniquely based upon business rules.

Professional-Grade Cloud is Cost Efficient

While difficult to find, select professional-grade clouds provide for user-defined scaling, wherein users define traffic thresholds that they do not want to exceed. This helps businesses to both predict and minimize costs, especially for broad marketing campaigns. With many businesses currently using a hybrid approach to cloud, measuring and predicting costs become more important.

Professional-Grade Cloud is Reliable and Secure

Dedicating virtual machines clusters to specific business needs, professional-grade cloud customers can add more capacity to a cluster at any time. This completely eliminates any scaling concerns. Reliability also comes in the form of greatly enhanced security. Customers provide their own servers to run in the cloud. In this way, customers can own and configure each cloud instance on their own. Not only does this make cloud more reliable, it greatly simplifies the security model. The main reason is the decoupling of the front end from back-end services. Professional-grade clouds also allow customers to maintain whatever security structure they see fit. This means they can use the same security tools that they have always used.

Professional-Grade Cloud Enables Control

Professional-grade cloud provides system administrator rights to the back end of the cloud. As a result, customers gain the same level of control as if it were their own infrastructure. This type of control is critical for fast-paced businesses with continuously evolving needs, especially those with multiple high-volume marketing campaigns.

Zunicore: A Professional-Grade Cloud Backed by Superior Delivery

Zunicore Cloud Hosting is the first professional-grade cloud that's purpose-built for business professionals. The Zunicore solution offers business professionals the control to manage their cloud hosting environments with simplified yet robust tools that correspond to the ways in which fast-changing businesses work today. In essence, businesses have their own Web server in their own Web instance. By de-coupling back-end services from the cloud's front end, each user can own, configure, control, scale and secure one or more cloud instances simultaneously. Zunicore ensures cloud success by delivering in six critical areas:

- 1 Control:** Zunicore provisions a pool of computing resources and allows users to decide how they want to use those instances. As needs change for the business users, they have the control to modify their resources and add more virtual machines to their pool, and increase or decrease the size of their pool, for example.
- 2 Reliability:** Most clouds use local hard drive storage and offer extra storage at an additional cost. Zunicore offers persistent SAN storage, which is more reliable, much less susceptible to data loss, and higher performing. SAN storage within Zunicore

Cloud Architecture Security Tip

A best practice to optimize a business's cloud security is to leverage a Web server and a database, and have all incoming traffic enter a reverse proxy that passes traffic to a virus scanner. The rest of the back-end resources communicate with each other on a virtual private network that only that cloud can access. In this way, the upfront virtual image can have virus scan, email scan, SSL access, VPN, etc. Additionally, system administrators within a company know what kinds of attack vectors they are likely to encounter (spam, DDoS, HTML injection, etc.) and that skillset is perfectly transferable to Zunicore's model.

provides dual data paths, hot spare drives, and an operations staff that monitors storage health on a 24x7 basis, all at no additional cost.

- 3 Security:** Every Zunicore customer has direct, unscanned access to the Internet and can leverage that access as they see fit. They can take data streams and deploy any packet inspection, packet software, encryption, or other security technology directly. Most cloud providers do not put such control in the hands of their customers.
- 4 Scalability:** Zunicore's Auto-scaling feature allows users to specify how much to increase or decrease resources. Recognizing that business drivers often underpin cloud initiatives, users can manage Auto-scaling based upon business rules or technical resource rules. For example, they can set scaling thresholds based upon cost, CPU and RAM usage, or bandwidth or any combination of cost and technical resource usage. For businesses with small staff size or unpredictable traffic demand, Zunicore offers Hands-Free Autoscaling, which lets users define thresholds. This cloud form of "overdraft protection" makes additional Auto-scaling available for every resource pool. Plus, a simplified dashboard interface makes it easy for administrators to make changes quickly.
- 5 Cost Efficient:** Zunicore's FairPlay Pricing™ reflects the way in which businesses want to work in partnership with their professional-grade cloud provider. Rather than the traditional antagonistic pricing models of the cloud with their inherent hidden costs, surprise penalties, and overage charges, Zunicore enables businesses to only pay for the resources they actually use, without making up-front commitments. Since Zunicore's architecture allows for instantaneous changes, performance will never suffer. High levels of architectural transparency and detailed resource usage reporting help customers to understand campaign costs. Equally important, cost information comes in advance, with costs broken out by individual campaign. For ultimate flexibility, customers can scale different machines and different clusters at varying levels so that they can respond to varying business needs.
- 6 Simple:** The complexities of managing cloud have driven many businesses away from the public cloud. Multiple levels of administration lead to mistakes, wasted time, and wasted resources within internal IT departments. Zunicore developed its platform with a simple dashboard interface to show the resources being used, the amount of resource remaining, and options to increase or decrease resources. Users can leverage these tools on the fly or set thresholds for unforeseen traffic bursting.

Professional Grade with Unmatched Network Infrastructure

Zunicore leverages the established infrastructure of PEER 1 Hosting, one of the world's fastest-expanding hosting companies, with access to multiple connected datacenters and its wholly owned FastFiber Network™. As a result, not only do professional grade cloud customers of Zunicore obtain all the benefits of business-friendly cloud infrastructure, their cloud instances run better and faster. Additionally, Zunicore has extended its network infrastructure to the last mile so that Zunicore professional-grade cloud customers can reach any Internet user, anywhere.

Professional-Grade Cloud on Professional- Grade Network Architecture

1- Network Reach: Zunicore makes its services available in 3 geographically dispersed datacenters on two continents, extending to 18 datacenters and 21 network points-of-presence (POPs) around the world.

2- Content Delivery: Zunicore partners with established and proven third-party content delivery networks (CDNs). Zunicore has also enhanced traditional CDN. With traditional CDNs, Web-only content follows the traffic load. With Zunicore, computing resources follow the traffic load to ensure that everything in the architecture can respond as fast as possible.

3- Enhanced DNS: Having AnyCast DNS implemented across the FastFiber Network™ means a faster, more responsive DNS lookup. This ensures customers with load-balanced cluster nodes will go live in different datacenters as time zones change and traffic loads move across them.

4- Reliable Storage Backup: Zunicore makes persistent SAN available by default to deliver superior performance and reliability.

5- Service Assurance: Zunicore backs its cloud services with a 100% network uptime SLA.

Conclusion: Switching Has Never Been Easier

Business professionals today are searching for a more economical cloud solution that delivers high levels of flexibility and performance, at a reasonable price. Zunicore Cloud Hosting offers these businesses an unprecedented level of technical capability while enabling users to focus on their business, not the IT technical delivery architecture. In addition to its appeal for business leaders, Zunicore offers the first cloud that appeals to system administrators and application developers. They can greatly benefit from a cloud that takes their perspective into account and helps businesses gain more from their extensive architectural knowledge.

Convinced? Zunicore is partnering with developer-centric companies to offer easy migration for cloud 1.0 customers to make the switch to professional-grade cloud.

Bios



Dave Geada, Director of Product Marketing

Dave Geada is the Director of Product Marketing for PEER 1 Hosting, responsible for the global marketing of PEER 1 Hosting's product and service offerings, which include Managed Hosting, Dedicated Hosting, Cloud Computing and Colocation. With over 15 years of experience in high tech marketing, Dave has worked at name brand companies like Rackspace, Quest Software, Computer Associates, Network Solutions and VeriSign, as well as cloud start-ups like Buzzoink and StrateScale.



Greg Rusu, General Manager

Greg Rusu is the General Manager of Zunicore, PEER 1 Hosting's cloud computing division. Greg's career spans over 20 years, recently having worked with start-ups leveraging virtualization and cloud technologies in Europe's banking, finance and insurance sectors, as well as the end-user consumer and mobile phone data service markets. His experience extends to establishing software product lines for Dell, managing strategic partnerships with enterprise ISVs, and originating software products supporting AMD's technology in the high-end x86 markets.