

A Whitepaper



Cyclozzo: Platform done right


The beginning of sleepless nights

Running a high traffic web application is a black art or almost. There are many organizations in the annals of cyber history that have gone down, unable to scale gracefully. There are three very important aspects of any high traffic web application that are the most difficult to get right and are even more difficult to maintain well. These are scalability, high availability and performance. And all three of these aspects are unrelated to each other. However under certain circumstances one may affect the other. Each of these aspects have to be dealt with at different levels: servers, networking infrastructure, databases and also the application itself. Every choice of a component in the system has to be thought about. You will agree, it is a black art or almost.

Ask system admins of high traffic web applications, which situations they fear the most. They have one word to describe it: downtime. It is worse for B2B portals like Payment Gateways. They take down several businesses with them. So, system admins are always required to build in redundancy. So there are solutions but not all are graceful solutions. As the complexity of the infrastructure that supports high traffic web applications increases, it tends to get more and more inelegant.

The usual solution

The HTTP protocol is notorious for its simplicity and its inability to support complex applications. Even a simple concept of a user's logged-in session is emulated using one of the several available ugly techniques. But simplicity is good and engineers like it. The HTTP request originates at the browser, travels to the web server and a response goes back from it. But if there is only one web server it will soon run out of capacity, if there are too many users accessing it. The usual approach is to use a better architecture that can scale fairly well.



A Fairly Scalable Architecture





As the diagram explains, the web portal or web application will be made scalable just by adding in more and more application servers. Since the load balancer does not do a lot of work, one or two of those would easily suffice. This is called “Horizontal Scaling” as opposed to “Vertical Scaling”, where the crunching power of the machine is increased in response to more and more traffic. While it is clear that the horizontal scaling approach is better than the vertical one, the solution reaches its limits easily. The reason for the failure to scale rests with database. The traditional database is usually the most difficult to scale due to inherent design restrictions which makes it hard to scale horizontally. Some sites use a technique called sharding where information is spread across servers in a pre-determined scheme. Other sites use separate servers for reads and writes. There are almost endless possibilities. The greatest challenge here is maintaining the consistency of data across the cluster without complicating the application too much. There is no simple scheme that works for high traffic web applications without changing the application code.

A New Paradigm

A good solution is one where one does not make any compromises. Tackling the database problem has been done in several ways by different web applications, but with compromises - sometimes in the design, sometimes for performance and sometimes in the need to modify application code. The latest approach is radical and a very elegant one too. The new approach is to completely eliminate the old relational database and embrace the NOSQL approach. The new approach helps in building applications which will scale very elegantly even with huge amounts of data and web traffic.

It's already been field tested in production environment by companies like Google and Facebook to satisfy their huge scalability demands. Rather than using a database server residing on a single sever or multiple sharded servers, a distributed, fault tolerant cluster of a NOSQL datastore will help satisfy rigorous demands. In fact even the application servers themselves can also participate in hosting of the datastore. All access to the datastore is made possible using APIs familiar to web application programmers. Losing servers due to hardware or other issues has no impact on the availability or the integrity of data, since it is redundantly distributed across the cluster.

Welcome to The Platform

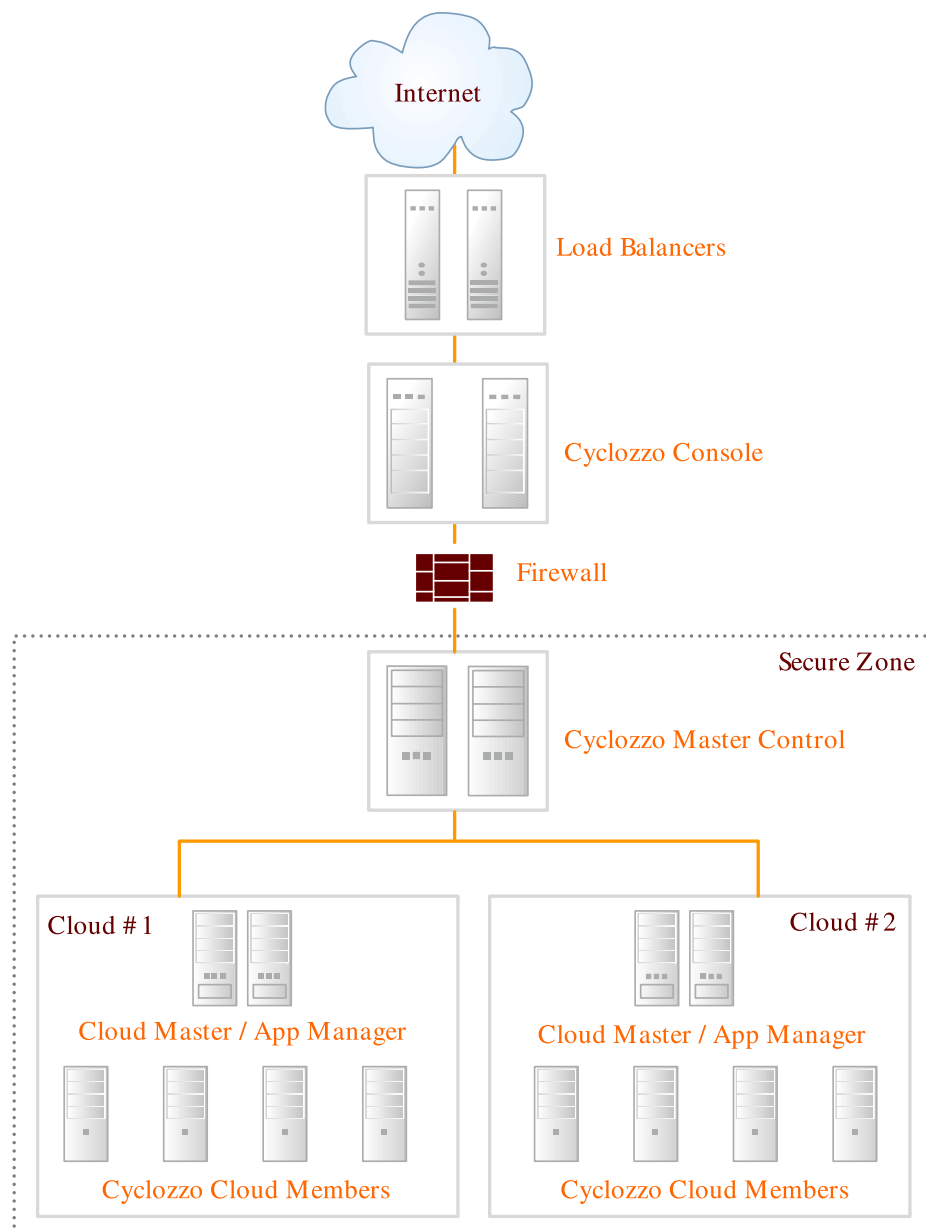
Cyclozzo is not just a web hosting product. It is a Platform that makes it very simple to run highly available and scalable web applications without hitting performance problems. And while doing so, it makes the life of the developers easy. Cyclozzo provides a very flexible API available for different, supported languages that makes it possible to develop web applications that scale well and are highly available. The API is the interface to the Cyclozzo platform as far as the application is concerned. There are some very simple rules to follow and when the application is hosted on the platform, it gains some amazing capabilities, which would have been very difficult to achieve otherwise.





Cyclozzo is much more than just the API supporting several languages. It is a complete system where workloads are managed from the physical infrastructure to the application level. From the Cyclozzo Admin Panel you can administer and monitor your physical infrastructure, look at application loads or manage users. Cyclozzo also continuously monitors applications and scales them up or down as necessary. And, applications can define maximum resource budgets too. Cyclozzo will take them into account.

All it takes is a bunch of physical servers. Once Cyclozzo is installed, it will form a cluster of cooperating systems that will make it possible to run multiple highly available and highly scalable web applications. Cyclozzo does not use virtualization. This results in excellent performance. Cyclozzo is based on a fully multi-tenant architecture. It can host several web application in a way that one does not affect the other and also in such a way that one cannot access the other's resources.





Who will want Cyclozzo

There are two main audiences for Cyclozzo. The first audience are hosting providers who want to provide the new business of the “Platform as a Service”, popularly known as PaaS. Deployment of PaaS also presents new upselling opportunities to existing customers on dedicated hosting plans.

The second main audience for Cyclozzo are organizations with web applications that bring in huge traffic and are growing or plan to grow.

Google App Engine Compatibility

Cyclozzo is 100% compatible with Google App Engine (GAE) Application Programming Interfaces. This lets organizations move their code without changes from GAE to their own infrastructure, for example. Cyclozzo currently supports the Java and Python programming languages supported by GAE.

Contact

Web: <http://www.k7cloud.com>

E-mail: info@k7cloud.com

Phone:

Shuveb Hussain / +91 98403-80386

Navin Sylvester / +91 95000-87617

