The Hybrid Cloud: Benefits and Challenges of a Mixed Sourcing Model

International Systems Group (ISG), Inc. White Paper in association with The Outsourcing Institute

The hybrid cloud has attracted a lot of interest in recent years. Surveys show that two thirds of organizations that are using a private cloud are considering to implement a hybrid cloud. So what is a hybrid cloud? Here is a popular definition:

"A hybrid cloud is a cloud computing infrastructure composed of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability."

From a practical perspective it typically means that an organization has implemented a private cloud and is using one (or several) public clouds. These two types of clouds serve different purposes and the workloads put on one or the other are often unrelated. For example, a public cloud could be used to outsource commodity services like email, or we can set up a temporary testing environment whenever we are at a point in the lifecycle of an application development project when we need it. The private cloud is used for applications that have characteristics that make it prohibitive to run them in a public cloud. Those could be regulatory requirements, restrictions in where company data can be physically located, or performance guarantees that could not be met by a public cloud provider's SLA, among other things.

In a hybrid cloud configuration we have a workload that has one part running in the private cloud, another in a public cloud, and there is a relationship between them. There is no strict definition what constitutes such a relationship. It could be organizational, operational, technical, or a mix of all, as illustrated in Figure 1.

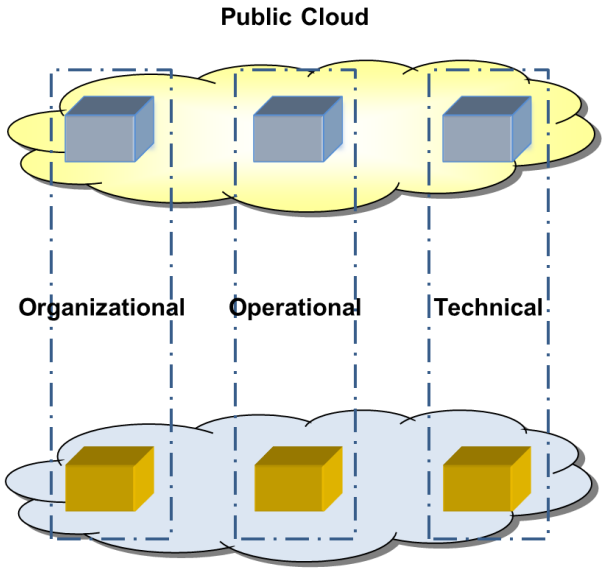

Figure 1 Private / Hybrid Cloud Relationships

Let's look at some hybrid cloud use cases to see what exactly this means. An example for an organizational relationship is building an application in a public cloud and utilizing the private cloud for production deployment. A public cloud is well suited for the development and test activities, since those require temporary

resources that could be released once the project is completed. There are no limiting factors like strict security or performance requirements, and there is no need to move production data to the public cloud - a sample of artificial test data does suffice. However, exactly these requirements may apply to running the application in production, so that's where the private cloud comes into the picture. The two environments have no technical dependencies; it's all about organizing the transition through the various stages from development to production.

This use case is shown in Figure 2, along with three other use cases that we are going to discuss next. Disaster recovery is an example that constitutes for the most part an operational relationship, but it also has a technical component. In this case we run a workload in the private cloud and use a public cloud as a backup environment that we can switch to if our data center that hosts our private cloud has a major outage. The only active system in the backup environment is typically the database; other components like load balancers and Application Servers are configured but not actively running. This setup requires that IT solves two key issues: first, how does data get moved from the private cloud to the public cloud and how are both systems kept in sync; second, how to activate the components of the backup system that are not running and switching users over in case of a private cloud outage. The largest part of this hybrid cloud use case is setting up operations procedures, but there are also infrastructure and middleware that need to be put in place for the data synchronization, hence the technical component.

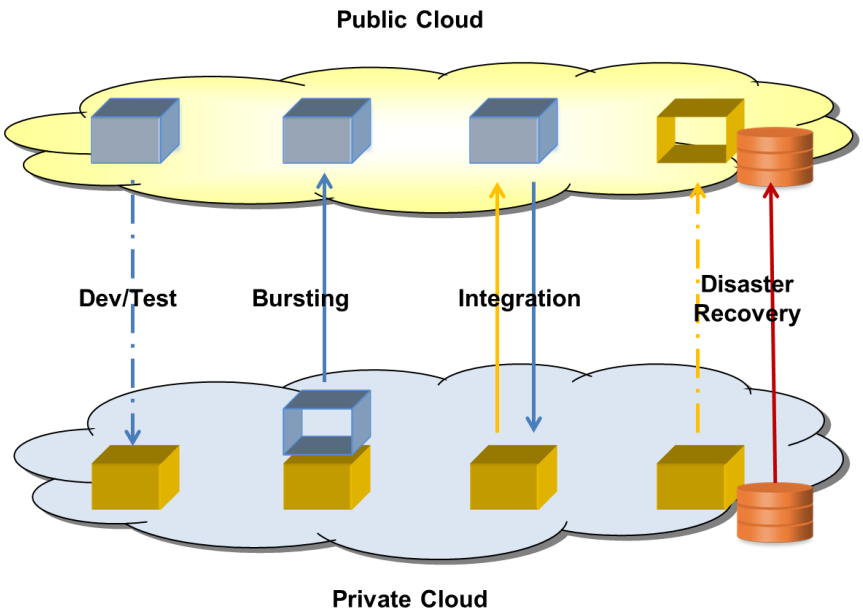

Figure 2 Hybrid Cloud Use Cases

Cloud bursting has some similarities with disaster recovery, but it is more complicated since it involves workloads that are distributed in real time and execute in parallel. The idea is to run a workload in the private cloud and utilize a public cloud as an additional resource during periods when demand exceeds the capacity of the private cloud. This setup requires careful planning and identification of applications that can take advantage of cloud bursting, since not all applications are feasible, and some require changes that could be cost prohibitive. The biggest hurdle is getting applications that are bursted to the public cloud access to the data that they need to run. There are basically two options to accomplish this: first, one could send a portion of the data to the public cloud that is needed for the on-demand workload that is running there. However, this approach is problematic, since it incurs transport costs and latency, and the big question is how to segregate the data between what stays in the private cloud and what goes into the public cloud.

The second option is to have the application run in the public cloud while accessing data stored on-premise in the private loud. Is adding a small delay to each transaction acceptable? It is for the gaming company Zynga (the provider of the FarmVille game), they chose the second option. They put front-end game applications into a public cloud and let them access common services that run in their private cloud. The data flows over the network back and forth between the public and private cloud. This approach provides a number of advantages that show why the hybrid cloud is sometimes called “best of both worlds”. First, the components that are run in the public cloud are highly scalable, and this scalability is available on demand. Examples include load balancers, Web servers, Application Servers, and front-end applications. Second, data at rest can be highly secured, since it is kept in the private cloud. As we can see, this is an example of a mostly technical relationship between the two clouds.

Our last hybrid cloud use case example demonstrates a purely technical relationship. When we move an existing application into the public cloud, or replace its functionality with a SaaS application, we often have to deal with a multitude of upstream and downstream dependencies with other applications. For existing applications these dependencies have been implemented through some form of integration, be it as custom code or by utilizing integration middleware. However, since the application is now running in a public cloud, the integration has to cross boundaries like firewalls and NAT (Network Address Translation) devices. This is not a problem when calling outbound (i.e. from an on-premise system to a cloud based system), but the other way requires to either change the custom integration code (not advisable) or adopt middleware for cloud integration, like an Internet Service Bus. The latter is also what we will have to do in order to integrate a SaaS application, since there is no integration out of the box with our internal applications.

The four hybrid cloud uses cases we discussed show the potential benefits of a mixed cloud sourcing model, but we also highlighted some of its challenges. However, the process of migrating applications from an on-premise data center to a public cloud comes with an entire set of additional issues that need to be considered. There are technical questions, for example what infrastructure requirements does the application have and how can the application be configured in a target cloud, and what tools can be utilized to automate some (or most) of this process. There is also a financial optimization problem to be solved: there are many public clouds to choose from, and they offer a multitude of server sizes and pricing models, which can result in up to 50 times variance in cost for running a particular application. Stay tuned for future articles where we will examine these issues.

About the Author

Gerhard Bayer is a Senior Consultant with International Systems Group (ISG), Inc. He has more than 25 years of industry experience, working for software vendors as well as for consulting firms. He is currently focusing on the migration to the Cloud and integration architectures for the Internet of Things. In addition, he is responsible for maintaining ISG's training curriculum and teaches a variety of IT seminars. He is a contributing author at Cutter Group, teaches seminars for Technology Transfer Institute, SOA Institute, and EITA Global, and presents at international conferences.

International Systems Group (ISG), Inc is a provider of strategic consulting, project management, architecture, and development/integration services to Fortune 2000 companies and government agencies. ISG helps organizations with implementation of large-scale applications, application integration, Service Oriented Architectures, and the migration to Cloud Computing.

ISG has been serving its clients for over 20 years and has extensive public- and private-sector experience. Requirements, technologies, and solutions keep changing, but one constant keeps proving itself: there is no substitute for experience. We are proud to look back on a long track record of successful projects, and we are looking forward to assist our clients to navigate the challenges they are facing today. Every new project we engage in builds on best practices that have been established over years, and at the same time we strive to remain a thought leader in IT – our key personnel is well published and regularly speaks at major industry events.