Knowledge Transfer Network

ICT

Case Study: Cloudsoft

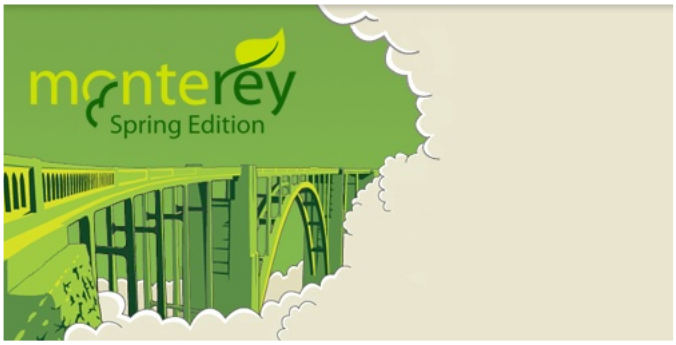

By Gillian Law
19 May 2011

Cloud offers great flexibility but that flexibility has to be controlled if enterprises are going to entrust their core business applications to the cloud, says Duncan Johnston-Watt, founder and chief executive officer of Cloudsoft.

Cloudsoft has developed an 'intelligent application mobility platform' to manage the increasing complexity of cloud computing, recognising that as cloud gets more ubiquitous and as larger companies begin to use it, they are demanding more control over how applications are scaled and distributed within and across clouds. In addition, as applications running on the cloud become more complex there is a need for more fine-grained management to keep things running as smoothly and quickly as possible. This is true whether an application is run in a private, public, or hybrid cloud environment.

With that in mind, Cloudsoft has developed Monterey. Based on IP that was initially developed in Johnston-Watt's previous company, Enigmatec , Monterey is designed to augment cloud application platforms such as VMWare's vFabric. Within these platforms, Monterey automates how applications are scaled and distributed via user-defined policies. These policies continually optimise an application based on its behaviour, where data is located, where user-demand is coming from, resource availability, costs, and SLAs. This is particularly powerful for applications that run across wide area networks, since the location where a specific piece of processing is executed can have a major impact on performance.

While companies like Amazon and Rackspace have developed enormous cloud businesses based on making computing infrastructure very easy to provision, Johnston-Watt says there is "an emerging requirement for automated control at the application level as well as simply at the infrastructure level. Especially for business critical applications, it's not enough to just have the right number of servers allocated to the application as a whole. It's equally or more important to look within the application and ensure that the various application components have the specific resources they need, and are located where they can deliver the best performance. The only way to achieve this, while ensuring that business controls and jurisdictional constraints are enforced, is through application policies. These policies are user-defined, and guide the cloud provider as to how best to meet the non-functional business requirements."

The future development of cloud is likely to be in this direction, Johnston-Watt says.

"We pretty much know how to do Infrastructure as a Service, that's really just utility computing. The focus now is very much on moving beyond utility computing to providing a platform that delivers a range of services directly to the application – what's known as Platform as a Service. Without a platform that transparently controls the infrastructure on behalf of the application, it's very hard for people to develop new services that can fully exploit the underlying flexibility and elasticity of cloud computing," he says.

Monterey isn't a Platform as a Service in itself, but rather is the essential middleware that provides its foundation. "Our middleware can be used directly, or else can be embedded in a framework that makes it all but invisible to the developer so that it can be used without needing to learn anything new – it's simply a part of vFabric or whatever other platform it's built into".

"We're careful to be complimentary to the big players in the application platform space, and to work with them. Within their over-arching platforms we provide specific services that make it easy to scale and distribute transactional business applications much more efficiently and flexibly than has previously been possible. Monterey does this by creating an elastic processing fabric that spans machines, data centres and clouds, and by activating transactional applications as highly mobile segments that can flow across the fabric as needed. The size and shape of the processing fabric created by Monterey, and where any given segment finds itself in the fabric at any given point in time is automatically determined by your user-defined policies. It all adds up to an application that is continuously self-optimising and which adapts near-instantly to changes in operating conditions. We call this Intelligent Application Mobility".

"There are many reasons why application segments should move around machines, locations and clouds. For example, segments may need to move in order to scale-out an application; or to move closer to specific users in order to improve response times; or to move closer to where the data is located; or to implement "follow-the-sun" processing; or to use the most cost-effective computing resource; or to escape from failing infrastructure; or to move between cloud providers. Whatever the reason, Monterey moves the segments while they are still running so that there is no interruption to service. It's completely dynamic and transparent to the user".

Part of the theory underpinning Monterey is that only the critical data needs to move in order for processing to move. As an example, Johnston-Watt describes a share-trading service. "If the centre of activity for trading a specific share suddenly switches from, say, Singapore to London then let's move the trade-processing service for that individual share to London. All we need to do is move the order-book data for that specific share – that's the core information – we can then use this to completely re-constitute the trading service in London. It's quite simply the fastest way to move things around, and ensures we can keep the service running uninterrupted. By co-locating the work with the demand in London response times are dramatically reduced, which in this case provides valuable competitive edge for securing the best possible trades".

Another benefit of keeping critical data with the processing segments is that it can dramatically reduce the workload on databases: instead of referring every request for data to a database, the required data is already active in the segment ready to be processed. "If the ticketing system for the 2012 Olympics had been architected in this way it wouldn't have ground to a halt" he contends.

"It ran out of capacity because of the way bookings were done. Each of the sports, each event, each round, could have been dealt with as a separate segment. If you're interested in show jumping, you don't want to be competing for resources with someone who's interested in swimming. So you need the ability to split these

things out and still guarantee that people are dealt with in the correct order and that bookings aren't lost. Once booked, then and only then, the data needs to be sent to a database".

Middleware has been used for decades to simplify and enhance enterprise application architectures Johnston-Watt says. "The advent of the cloud computing model has created the need for a new type of middleware, one that can fully exploit cloud's elasticity and pay-as-you-go benefits while maintaining control through automatic enforcement of policies even, or especially, where your services are moving around a public cloud. That's the argument we're making, and people are beginning to buy into it."