

Virtualization: The need to know what is going on

Tim Chaffe¹

¹University of Auckland, New Zealand, t.chaffe@auckland.ac.nz

ABSTRACT

The University of Auckland began virtualizing its infrastructure around five years ago. In 2010 we have virtualized our storage and our compute layers to such a degree that greater than 90% of our application solutions and services are delivered using some form of compute or storage virtualization. We are currently sitting at one thousand server instances and are still growing. With the next wave of private and public cloud offerings are around the corner. But are we or they really delivering a quality service?

As applications and services are aggregated onto massively shared platforms, management issues come to the fore. Virtualization succeeds by sweating aggregated resources. However the simultaneous behavior of those applications can change dynamically causing actual end user experience to be highly variable.

Without proper management the performance of an application can swing between excellent to unacceptable and back again within a small timeframe. Such micro-outages will not be visible on the system charts or even traceable using standard systems monitoring techniques. The IT department will deliver its aggregated statistics showing perfect uptime and averaged CPU/memory consumption. However from a service delivery perspective, no adequate measurement of service performance is being provided.

In the past the technical view is that it is impossible to provide such information. Writing synthetic transactions against every business interface is time consuming and wasteful. Instrumenting every piece of application software to track user performance is prohibitive for software vendors and lacks a standard approach across different vendors. (and you probably have to pay extra for it!)

A solution is at hand however in the domain of Application Performance monitoring. This monitoring works by watching transactions between clients and servers as they traverse the network. This allows precise measurement of the conversation between a client and the server. If a client is on a slow link such as a dial-up modem or wireless handset, this can be measured and recorded. If the virtual infrastructure slows down due to a micro outage, it will be recorded.

This allows - applications to be measured against SLA - problem web pages to be identified - systems tuning to be tested and proven ahead of the proposed change (no more 'trust me' black magic)

In the presentation I will show what Application Performance monitoring is and give real examples of how it has helped the University of Auckland. This will include commentary on the evolution of the way IT folk run systems with observations on why having the users actually know what the performance is a scary thing for IT folk. In simple terms, if an IT Director is virtualizing or Cloud sourcing without using Application Performance monitoring they're flying blind.

<http://creativecommons.org/licenses/by/4.0/>



Attribution 4.0 International