

# Cloud Computing, Economics and a Universal Truth



Robert Haynes, 2011-26-08

---

I was reading an interesting article in ZDNet about the [economics of cloud computing](#) and was struck by a universal truth - to deliver a service for the lowest cost you need to make maximum use of your resources *for the maximum amount of time*. This is the principle that drives a wider range of designs than you'd maybe imagine, e.g:

- Storage subsystem designs
- IT outsourcing
- Cloud computing

All of these are designed in such a way as to smooth the peaks of demand to reduce the resource needed to supply the service. Whilst the RAM cache of a storage controller might be somewhat different to a team of network engineers in a [NOC](#), they are all there to service a workload.

Just as it's more economical to provide memory cache to deal with the peak I/O workloads, rather than hundreds of extra disk spindles, the cost of running a 24x7x365 NOC with enough staff to cover sickness, holidays and training is better shared by multiple organisations than each providing its own.

The same works for cloud computing, where cloud service providers will rely on providing enough compute resource to meet their average requirements, rather than the theoretical maximum if every customer used their resources all at once. Whilst all this was running around my head (I prefer to believe that it's this high-value deep thinking time that results in my somewhat modest productivity rather than my easily distracted nature), F5 quietly announced availability of [Version 11 of BIG-IP](#).

There are loads of new and pretty cool things in this release, which we promise to bombard you about over the coming months, but one that stood out was our new Scale<sup>N</sup> architecture, which lets you break away from traditional two node clustering into a world where *workloads* can migrate between members of a pool of application delivery controllers.

So if, for example, your downloads site hits a huge peak, you could migrate off your Outlook Web Access workload to one application delivery controller and your E-commerce site to another. If one device fails, its workload can be spread around multiple peers. Suddenly we seem to have a way to smooth those peaks and meet even the toughest SLAs.

---

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | [f5.com](http://f5.com)

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)