031

Living on the edge with HA-clusters

Adam Maulis maulis@ludens.elte.hu

ELTE IT center

A fine piece of hardware can be run with 99% availability. (Let's assume that. If it is really fine.) At times, this is just not enough, so, they say, let's make a cluster. There are many types of clusters, even if only the ones that provide higher availability are considered. Some operating systems and clustering solutions only know one kind; administrators of such systems tend to think this is "The" Cluster. Many types of clusters may be implemented, ranging from the shrink-wrapped, pre-configured, howto-backed solutions (which may or may not fit the needs of a particular company) to the completely custom-built hardware-network-software combination (for which one may or may not have enough in-house expertise).

Common pitfalls, told in a somewhat informal way. Why is it good to have a week of scheduled downtime to avoid two minutes of unscheduled downtime? With a service already being online and used by the public, what good is it to differentiate between the two types? What will your cluster do, if the VAC-repairman faints into the rack cabinet, tearing your heartbeat apart? How happy will you be to recognize that even if your web shop continued to operate after a hardware failure, your MySQL server does not know whether there have been any orders? What do you think, having developed an application for 50 grand without considering any clustering, how much will it cost to `cluster' it? Would you like to know your data safe, ten kilometers away? Being exposed to heaps of primitive DoS-attacks for that?

This talk will show several popular, known-by-many customary clustering solutions, which are perfect to bring availability down to 95%. Or even further...