

Clustering Environments

In today's fast paced world, the importance of having an e-mail service needs no further explanation. We all use it on a daily basis, relying on it to keep in touch with the surrounding world. However, when the throughput of information becomes so great that a simple system cannot keep up with it, new methods must be deployed to cope with the growing requirements of networks and their users. Clusters fill up the gap between the need for faster, more powerful systems, and the safety, reliability and easy management of the already existing networks.

Introduction

Any seasoned system administrator knows very well the implications and testing requirements of any new service that needs to be added to an already existing design. It must fit just like a hand in a glove and there is very little room for mistakes. This is the main concern, because severe repercussions can cause down-time to other services previously running at normal parameters. Because of their complicated structure, clusters are somewhat feared and seldom included into complex setups. However, a well structured and designed cluster is in fact a very sturdy and resilient entity that adds up to the overall network stability.

Cluster advantages

Computing clusters, as their name implies, are usually made up of a group of computer systems that work together to achieve a common objective. The systems that make up a cluster are called either "nodes", or "cluster nodes". The purpose of a cluster differs from setup to setup. The methods used to achieve the preset goals dictate the type of the cluster that needs to be used and all the specific details related to the actual configuration.

Clusters are mainly used for the following reasons:

Scalability: Large networks, with a large number of users and processing or providing large amounts of information, grow very fast. To keep up with the growth factor, provided services must be extended to cover the new requirements. A cluster can provide a very quick and stable resolution to this issue by adding new elements to help process the surplus, either permanently or temporarily. This is a very valuable ability of the cluster as most thriving networks reach a stage when their systems and services can no longer keep up with daily demands. Being able to add or remove systems from the setup can solve a lot of problems very fast and with no complications entailed.

Redundancy: Networks usually have weak spots, like central information nodes that are vital to their normal operation. These weak spots can cause important losses if they fail for any reason, either hardware or software. Redundancy provides the means to prevent or solve such situations. If a system fails for any given reason, it is automatically replaced by another one standing by, or by a previously designated system in the network. This way, down-time is reduced to a minimum, usually within a limit of only a few seconds. The process is completely automatic and the network administrators can be notified right away regarding the problem. Complete redundancy for all of the services provided by a network can be achieved by using clusters.

Load Balancing & Speed: Quality of service often depends on speed. No matter how tuned up a system is, it can process only so much. No more processing power translates into a loss in quality. Using more than one system to share the processing burden or the number of connections served can be easily achieved in a cluster. This ability, along with the above discussed scalability, turns the cluster into a hot-swap processing plant with an unlimited amount of processors available.

Ease of Maintenance: In a medium or large network, managing individual systems is a very tedious task. Clusters have the unique ability of being a centralized resource that can be managed entirely from one location, given the right tools. All of the systems within the cluster can be examined and supervised with ease, while services deployed within the environment can also be kept under control. This greatly enhances the network administrator's ability to troubleshoot and solve any issues that might occur, in a time-effective manner.

Conclusion

For each of the above mentioned points of discussion, a different approach must be taken. The important thing to notice here is that all clusters rely on more than one node to carry out their tasks. As such, this is the common starting point for all cluster setups. Clusters are usually being considered when one type of resource is no longer sufficient, or when a single system can no longer keep up with the demands. This means that it is often too late to prevent the shortage situation with the aid of a cluster, a very complex technical undertaking that requires careful planning and even more careful testing before it becomes ready and stable enough to append to an already crippled network. This is the most frightening side of clustering. Implementing a very complex system on short notice with high expectations can be frustrating.

Note: Clusters are really complex; there is no doubt about that. However, if the components that make up the cluster are understood correctly, only minor issues should come between the initial project and the resulting solution. Nothing can compensate for the cluster engineer's experience. It is imperative to correctly estimate and manage available resources, as a cluster has the bad habit of consuming as much as it can out of every available

resource.

About the Author

For the original story, please check: [Clustering Environments](#)

Source: <http://www.articletrader.com>