A Common Database Replication and Backup System for Network and Application Services in the Internet.

Motivation

The majority of the Internet network services (DHCP, DNS, E-Mail, etc) have already some form of redundancy mechanism that sets their resilience to a certain level. The resilience level for these services can be satisfactory only if the resilience of the underlying network and supporting computer resources is high enough so the frequency and duration of episodes that make these services unavailable (or performing deficiently) have no discernible effect on the overall quality of service as it is perceived by the users.

Server substitution is usually not considered but these redundancy mechanisms and some services don't rely on secondary servers (like the routing service, for example, which relies on replication of the routing database to all running servers without the use of secondary routing servers).

Sometimes, load balancing is used has an indirect method to implement redundancy but it will also mandate the need of, at least, two running servers at all times, and they will not share exactly the same service runtime databases.

Finally, some services can use complete server replications, although this is usually achieved using a front-end interface (virtual server) that replicates all the interactions between clients and servers.

Another issue of service redundancy is service recovery after failure, which is most often, directly related with the runtime database replication and backup mechanism.

One of the most important responsibilities for the network services management staff is to insure an acceptable level of resilience for each network or application service installed and maintained. This task mandates significant knowledge about each installed service and its methods of implementing redundancy through mirroring, load balancing or replication strategies. Additionally, each service could have several alternative software implementations with their specific features, configuration languages and runtime database formats.

All this accounts for the great complexity associated with such an endeavor. Some of this complexity could be diminished if these services could use similar strategies and common configuration languages and runtime database formats. Even for implementation of different strategies, it should be possible to use common components of these redundancy systems.

A Common Redundancy and Backup System for Network and Application Services on the Internet would have an important impact on the current network management activities and practices since it would provide a unique methodology for implementation of redundancy and a backup feature for all these services, eliminating the need for knowledge of several systems and the use of a different software implementation for each service.

Also, this would permit separate development and implementation of the productive functions of the services and the redundancy and replication aspects.

Objectives

Runtime database format representation and replication is one of the most important aspects of all the Internet network and application services redundancy mechanisms. Each service uses its own replication technique to try to maintain all server databases coherent. This is a complex task and some of the simplest services cannot rely on runtime database replication strategies.

Furthermore, database replication techniques can also be dependent on implementation packages and the most advanced techniques are not standard methods.

Finally, the majority of these services do not define any precise language or syntax for representation of the runtime data to be passed to the database sub-system. Although this brings flexibility to the standards it creates additional difficulties when defining a common runtime database redundancy and replication system to be shared by the majority of the services.

The solution should be in the use of a generic redundancy and replication strategy and a unique meta-language for representation for the runtime data to be stored on the database system. This meta-language would permit the use of the same database system to be used for storing and managing data for all services.

So, the first goal for this work is to study and analyze the redundancy mechanisms of all the major TCP/IP network services and to discover common strategies, features and techniques used specifically for runtime database redundancy, replication and backup. Special attention should be given to advanced state of the art replication mechanisms that are used on experimental implementation packages of DNS and DHCP services.

The next research goal would be the definition of a unique, generic strategy supporting redundancy, secure replication and backup requisites of all major TCP/IP network and application services and a meta-language for data representation off all standard services. This would permit the use of only one database system for all standard services on a domain or group of networks.

To complement the previous research goals a prototype redundancy and replication system should be developed that would be compliant with the new strategy and use the newly defined meta-language for runtime data representation. A well known open source database technology should be chosen that permits a simple implementation of the newly defined methodology. This system should be applied to a well known network service and its efficiency and features should be compared with the same system using present redundancy and replication strategy solutions.

Supervisor: Bruno Dias, <bruno.dias@di.uminho.pt>
Research Unit: CCTC