

A Generic Hot-Failover State Machine for Gateway Services and its Application to a Linux Firewall

H. Roelle
Munich Network Management Team
University of Munich
Oettingenstr. 67
D-80538 Munich, Germany
roelle@informatik.uni-muenchen.de

Abstract

Nowadays, companies of any size rely on their IT-infrastructure since it provides connectivity to the outside world. Services like firewalls, being positioned between the own domain and a foreign one, form a premises for higher level services. Therefore, such gateway services must be considered as especially mission-critical. While there exist high availability solutions for special service types, a generic solution which can be applied to arbitrary gateway services, especially for smaller sized scenarios, is missing. Fault tolerance in terms of high availability is addressed by this paper through the concept of redundancy. Presenting a generic state machine for monitoring and takeover processes, it leads to an universally applicable logic. The state machine's basis is derived from requirements posed by the generic scenario of gateway services. Furthermore, our solution's practical applicability is shown by presenting an implementation carried out for a Linux-based firewall system.

Keywords

High Availability, Gateway Services, Hot-Standby, Fault Management

1 Introduction

Today, IT-equipment and provided services represent mission-critical parts in a company's business infrastructure. Serving as a fundament for advanced services, this especially applies to those involved in connectivity to foreign domains. Against the background of increasing security threats even small and mid-range companies cannot abstain from complex gateway services like firewalls and application-level proxies. Hence, for such services, fault tolerance in terms of availability is gaining rising importance in environments of any size. While large companies already afford adequate solutions, in smaller scenarios the problem is often disregarded.

Especially in lower budget environments, nowadays off-the-shelf hard- and software components (like industry-standard PCs running multi-purpose operating systems like Solaris, Linux or Windows) are used to realize gateway services. But with a rising number of mission-critical components the probability of service failures increases equally. As a common and widespread solution, this problem can be solved by hot-standby redundancy, which means to supply additional systems with the ability to immediately take over service provisioning in case of a failure. While this approach is often feasible from a financial viewpoint, it poses problems in management: To realize seamless service availability, failures have to be detected and service provisioning has to be shifted, preferable automatically, to a different system. Furthermore service functionality has to be replicated on the redundant systems.

As gateway services like firewalls and application-level proxies often are not dependent on a persistent state, service replication is fairly straightforward by spreading static configuration data among systems in the redundancy cluster. Therefore this paper concentrates on failure detection and hot service takeover. It presents a reasonable generic solution for realizing fault tolerance by redundant systems, particularly for smaller environments based on off-the-shelf hard- and software components. For being universally applicable, the whole logic regarding the processes of monitoring and takeover is depicted by a state machine. Thereby an eye is kept on being directly implementable and easily customizable for concrete services with minimal additional effort.

In the following section an abstracted scenario is used to derive requirements for the solution. Afterwards Section 3 presents related work and compares it to these requirements. Section 4 presents our solution and the following Section 5 shows an implementation of our solution for a packet filtering firewall. Finally Section 6 concludes the paper und indicates further work.

2 Scenario Analysis and Requirements for a Generic Solution

For developing a widely applicable solution an abstract scenario is introduced in this section. By means of this scenario, important terms will be defined and requirements for a generic solution will be identified afterwards.

Fig. 1 shows a logical view of our abstract scenario, which is divided into two domains. One provides the gateway service that is linked by a communication network both inside the own domain (**downstream side**) and to the foreign domain (**upstream side**). The particular network links ending at the **reference points** are referred as **uplink** and **downlink**. Systems accessing service functionality from either sides are called **clients**. A **gateway service** in our context now is defined as a service being linked to different domains by discrete network links, while being completely located inside a single domain. In this scenario it is furthermore assumed that the communication protocol used to access the service is not genuinely capable of load balancing, otherwise fault-tolerance would already be present by service redundancy via the communication protocol. Examples for a scenario in our sense are IP edge routers, firewalls or application-proxies without a dynamic routing protocol.

Now that the system hosting the service is identified as a classical single-point-of-failure, adding extra systems to the scenario introduces fault-tolerance in terms of high availability (Fig. 2). The system which currently provides the service is called **master system**. The additional hosts, which are ready to take over service provisioning are called **backup systems**. Having backup systems raises the question how failure detection and handover in the case of a failure could be realized. For specifying a universally applicable answer, we have identified specific requirements by performing an in-depth analysis of a variety of concrete scenarios:

Requirements

Separation of logic and actions: In designing failover solutions two major layers of abstraction can be differentiated. The logic is responsible of steering and coordinating the processes during monitoring and handover. In contrast, actions describe the execution of concrete activities. For gaining a generic solution, it is necessary to separate the logic (which is identical for any service the solution is designed for) from service specific actions.

Distinction of service monitoring from communication link monitoring: To design a flexible solution applicable to a wide range of scenarios in a modular way, monitoring of service functionality must be strictly separated from checking of communication links. This enables partial reuse of realizations, when only the service type changes but the communication technology remains the same, and vice versa.

Service monitoring from client's perspective: In the sense of customer orientation, service monitoring must not only be carried out locally [HaRa 01]. From a client's point of view the service is provided properly when service functionality and communication links are available. Therefore, in addition to testing service functionality itself, monitoring also must include service accessibility through the communication links. In case of gateway services both links to the upstream and downstream side must be taken into account.

Independence from specific services and communication technology: A generic solution must neither make any assumption on the specific service nor the communication technology being used. Particularly in smaller scenarios, where off-the-shelf communication services (like IP mass-services) are deployed, special management and communication protocols are neither available by default nor can they be ordered optionally.

Independence from technology-specific communication primitives: For gaining wide applicability, special communication primitives like multicasting must not be used for the purpose of coordination the takeover. Only unicast communication can be assumed to be available in general.

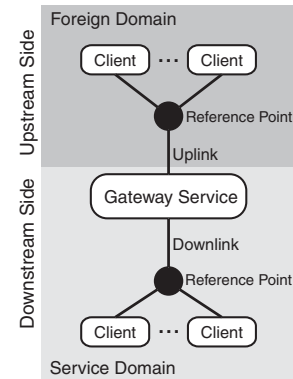


Figure 1: Logical View on Abstract Scenario

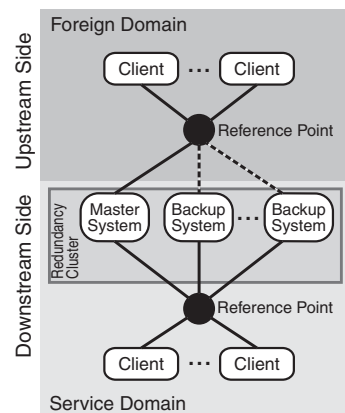


Figure 2: Logical View with Redundant Systems

Minimal active links to foreign domain: Regarding security, connections to foreign domains should be considered to be harmful. Therefore the number of active links to the foreign domain has to be kept to a minimum. This implies that upstream links should be kept down unless they are absolutely needed.

No need for extra hardware: For gaining flexibility in the usage of machines and achieving a short setup time of backup hosts, no hardware should be needed in addition to the one required to deliver the service anyway.

The next section will explain why existing solutions to the given problem need to be rethought and afterwards our own solution will be presented.

3 Related Work

In the field of fault-tolerance and high availability a great amount of valuable work has been published. A whole class of solutions deals with specific services and is mostly aimed for large-scale scenarios. In addition, they use specific properties and address special requirements of the particular service type, especially dealing with questions of service replication. High availability solutions for database management systems are a typical example. As they are aiming for a specific scenario, they are not considered further.

Another class of fault tolerance solutions proposes special communication protocols, like Stream Control Transmission Protocol [RFC 2960], to manage the takeover process. As these are specific protocols, the requirement of independence from specific communication architectures is violated.

The IETF Working Group for Virtual Router Redundancy Protocol (VRRP) [IETFWG-VRRP] has developed a protocol [RFC 2338, HMMH+ 02] delivering hot-standby redundancy for routing devices. Major limitations regarding our requirements result from being specific in communication technologies: On layer 2 it assumes IEEE 802, on layer 3 IPv4 is required. This limitation in generality is underlined by the fact that even for IPv6 a new revision [Hind 02] had to be prepared. In specifying a state machine consisting of only three states, specification is imprecise. Nevertheless VRRP was a good source of inspiration in general. Furthermore in [RFC 2787] it addresses management explicitly.

Closely related to VRRP is the Cisco Hot Standby Router Protocol (HSRP) [RFC 2281]. It is also designed specifically for routing services. In addition to the prerequisites of VRRP, it requires a dynamic routing protocol on the upstream side. Although not being directly applicable as a solution, it served as a valuable inspiration, namely its in-detailed specified state machine.

Many solutions try to overcome availability limitations by adding load balancers, whose purpose is to spread service requests to a cluster of hosts. While this approach remedies shortcomings on service hosting machines, it simply shifts the single-point-of-failure problem to the load balancer itself. Consequently they don't solve the problem, but the load balancing service as a gateway service itself is subject for redundancy and considerations made so far apply for it.

The IETF Working Group for Reliable Server Pooling (rserpool) [IETFWG-RSP] proposes an architecture [TXSS+ 02] which can be used to provide high availability for various services. In fact two different solutions are suggested. The one introduces a special protocol [SXS 02] for servers and clients and therefore violates the requirement of being independent from specific communication architectures. The other solution suggests to install a proxy which is equivalent to a load balancer. Hence, this solution also must be disregarded. Anyway, other documents of the working group helped in developing our own solution. For identifying requirements [RFC 3237] made suggestions and [LSXS 02] contributes by reviewing further related work.

Finally two prominent projects of the open source community should also be reviewed: High-Availability Linux (HA-Linux) Project [HALP] and Linux Virtual Server (LVS) Project [LVS, Zhan 00]. Both projects' goal is to provide high availability to the Linux operating system by redundancy in specifying clustering solutions. In the context of HA-Linux the takeover daemon "heartbeat" [Robe 00] was developed. It fully meets the requirement of independence from concrete communication technologies and does not require extra communication channels for message and heartbeat signaling. Unfortunately it is strictly limited to a two host setup (one master / one backup) and does not include any options for service monitoring facilities at all. As LVS's failover daemon "keepalived" [Cass 01] implements VRRP, it suffers directly from VRRP's limitations. Furthermore, LVS focuses in general on services solely connected to the domain it is hosted in and therefore does not address problems related to gateway services.

4 Generic State Machine

In this chapter our solution to the problem of designing generic redundancy for gateway services is presented. It starts with showing a design overview (Sec. 4.1), followed by an in-detail explanation of our

generic state machine (Sec. 4.2). Section 4.3 explains messages and timers used in the state machine. Afterwards customizable local procedures for applying the generic state machine to concrete scenarios are presented in Section 4.4. Finally, Section 4.5 summarizes our design.

4.1 Design Overview

The overall goal of our solution is to keep the gateway service operational in case of a failure. This is realized by setting up a redundancy cluster, where one single host acts as the master and an arbitrary number of hosts are serving as hot-standby backups, ready to immediately take over service provisioning in case of a failure (see Fig.2 in Sec 2). Hereby the following basic failure types are considered:

- Failure of **service functionality** of any cause on the current master host (e.g. internal service errors, hardware failures, ...)
- Failures in **communication links** of any cause on any machine in the redundancy cluster, either on upstream and downstream side.

Our design incorporates three main components: a **local state machine** executed on each host in the redundancy cluster, **message exchange** between hosts and a **status table** which lists all hosts in the redundancy cluster.

The state machine is responsible for monitoring service operation and coordinates service takeover in case of a failure. Customizable procedures are used to describe the necessary activities. They are noted as actions inside the states and are executed sequentially on entry into a state.

Message exchange: For coordinating a service takeover, messages are exchanged between hosts in the cluster and thereby our design assures that exactly one master host is present in the whole cluster. As the following explanations will show, the state machine's design is robust against loss of messages and assures that communication is operational before a backup becomes the new master host.

Status table: To designate which machine is to become the new master host, the status table is used. It prioritizes all hosts in the redundancy cluster in the manner of a total order, with the current master owning the highest priority. All hosts in the cluster are listed in this status table together with their current priority. Every host keeps a local copy of the table. Maintenance and distribution of the table is managed by the state machine, including new hosts being dynamically added to the cluster and others removed from it.

Furthermore, our solution relies on the following reasonable assumptions:

- State machine is implemented as specified and executed correctly as long as hardware is operational.
- Monitoring and testing procedures deliver positive-definite results. This means, whereas an erroneous negative result is acceptable (claiming that the object under test has failed, while it is still operational), a positive result must be reliably valid.

In the following sections our generic hot-failover state machine is explained in detail.

4.2 Hot-Failover State Machine

The UML statechart diagram in Fig. 3 shows the Main state machine. It gives an overview over the state machine as it is executed on each host participating in the redundancy cluster for a certain service.

The process starts with an initialization phase (init). Next, based on the initial priority received during init main state, a decision on the initial state of the host is taken (decide). When the host's own priority is lower than the one of the current master host, backup main state is entered. In case of the own priority being equal to the current master priority this host is the active master, consequently service main state is the next one. If the current master's priority is lower than the one of this host, a service takeover is necessary. This is accomplished by moving to handover main state.

In case a service takeover being necessary, handover main state is entered. If the takeover succeeds, service main state is reached, else the machine falls back to backup. When a host gives up service main state for

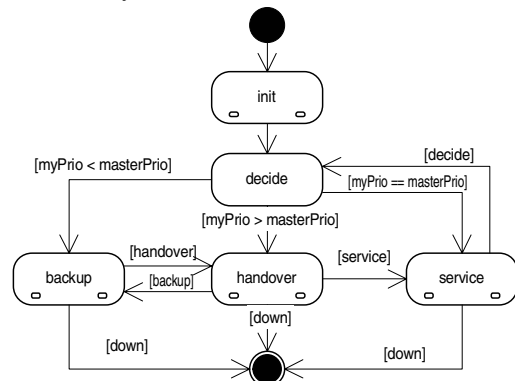


Figure 3: State Diagram of Main State Machine

any reason it falls back to the initial decision process. Additionally, from each main state a transition to the final state is possible.

In the following sections each of the mentioned main states will be explained in detail by showing sub-states inside and depicting conditions triggering transitions between the main states.

4.2.1 Init Main State

Init main state is passed only once in the life-cycle of the state machine and detects the initial priority of the host.

Entering check service present state (Fig. 4) the presence of a master host is detected by the customizable procedure `remoteSvcCheck()` (see also Sec. 4.2.2). When it succeeds a `SVC_CHECK` message is sent. When receiving a `SVC_CHECK_REPLY` message, the status table is initialized by the message's contents (see Sec. 4.3). When no `SVC_CHECK_REPLY` message arrives the cycle is restarted.

In case of an initial bootstrap of the redundancy cluster, no master service will be detected. Therefore the status table has to be initialized from a statically predefined initial table.

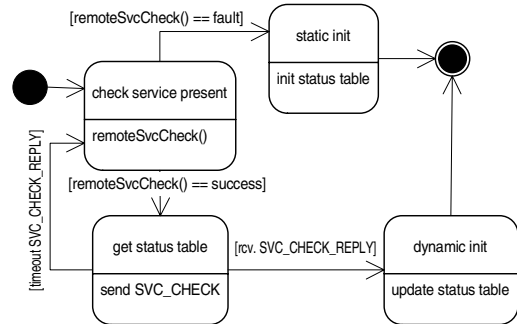


Figure 4: State Diagram of Init Main State

4.2.2 Backup Main State

The backup main state describes the behavior of a host acting as a backup for the service, remaining in this state until a service failure is detected. To detect such a failure (fulfilling one of the main requirements) an active service monitoring from the client's perspective is performed here.

Entering in backup idle state the `SVC_CHECK_TIMER` is started. Being started on each entry, this timer is responsible for triggering periodical remote checks of the service. In case there is a `SVC_CHECK_TIMER` alarm, the service functionally is first checked from the client's point of view using `remoteSvcCheck()`. In case of failure, backup main state is left and handover main state is entered (Fig. 3). On success, a local service check on the master host is triggered by sending a `SVC_CHECK` message. If a `SVC_CHECK_REPLY` message with fault status is received or such a message does not arrive at all, the service is considered to have failed and a takeover is initialized. When a success indicating `SVC_CHECK_REPLY` message was received, in learn state the local status table is updated from the message's data. Depending whether the own priority in the status table is the highest, learn state is left. If it is the highest, this host is designated to become the new master host and therefore a handover is necessary. Otherwise the system returns to backup idle state and the so far described process restarts.

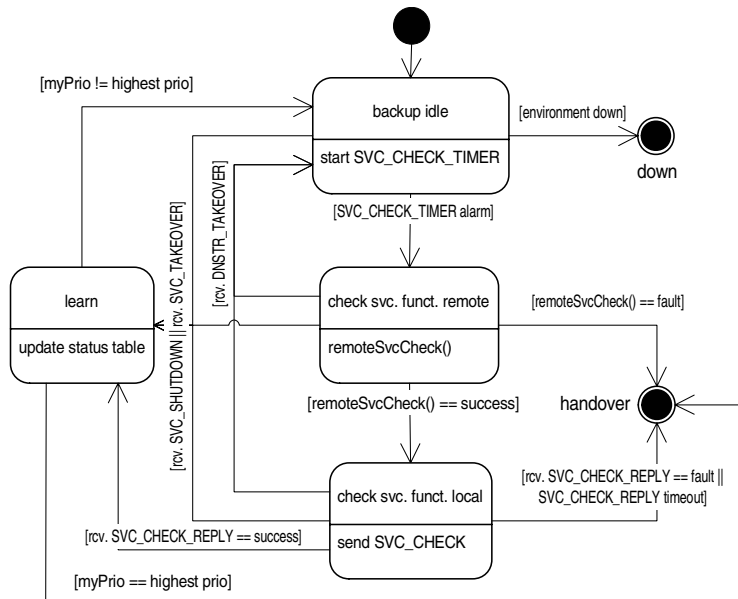


Figure 5: State Diagram of Backup Main State

Both checking processes are interrupted immediately when either a `DNSTR_TAKEOVER`, `SVC_SHUTDOWN` or `SVC_TAKEOVER` message (see Sec. 4.2.3 and Sec. 4.2.4) is received. As they all indicate a change of the master host, further service checking would be useless. In case of a `SVC_SHUTDOWN`

Both checking processes are interrupted immediately when either a `DNSTR_TAKEOVER`, `SVC_SHUTDOWN` or `SVC_TAKEOVER` message (see Sec. 4.2.3 and Sec. 4.2.4) is received. As they all indicate a change of the master host, further service checking would be useless. In case of a `SVC_SHUTDOWN`

or a *SVC_TAKEOVER* message, the status table is updated in learn state and the above described decision on the following state is taken.

4.2.3 Handover Main State

When the service monitoring procedures in backup main state indicate a service failure, a backup host needs to become the new master host. As the failure indication might not only be provoked by a real service failure, but also by a failure in the backup host's connectivity to the network, these both cases have to be separated before the service takeover finally can be completed. This distinction is the main purpose of the handover main state shown in Fig. 6. Additionally the current master host is informed about the ongoing takeover.

In handover start state the downlink is checked by `downlinkCheck()` to ensure communication to the other hosts is functional. On success, in handover downstream takeover state a *DNSTR_TAKEOVER* message is sent to indicate the beginning takeover to other hosts in the redundancy cluster, especially to the master host. The master may now reply a *DNSTR_TAKEOVER_DENY* message (see Sec. 4.2.4), leading to fall back to backup main state. If none is received, in handover uplink check state, first the uplink is activated by `activateUplink()` and gets tested by `uplinkCheck()`. If it succeeds, both links are proven to be functional. Therefore taking over the upstream is signaled to the other hosts by sending a *UPSTR_TAKEOVER* message in handover uplink takeover state. If the current master host raises its veto by sending a *UPSTR_TAKEOVER_DENY* message (see Sec. 4.2.4), the uplink is shut down by `deactivateUplink()` and the host remains in backup main state. Otherwise the takeover is completed and service main state is reached.

Takeover is interrupted if either a *DNSTR_TAKEOVER*, *UPSTR_TAKEOVER* or *SVC_TAKEOVER* message with a higher priority than the own is received, since this indicates that another host is going to take over the service. When any of the link checks fail, the host is considered to be inoperable for backup use and the state machine quits.

4.2.4 Service Main State

A host being in service main state designates it as the master host on which the service currently is executed. Furthermore it is responsible for maintaining the status table. To avoid inconsistencies it is changed only on the master host. Besides, a local monitoring of the service is accomplished and on-demand checks are carried out.

When entering service main state, first the status table is updated to reflect the new situation which is announced to other hosts in the redundancy cluster by sending a *SVC_TAKEOVER* message. Afterwards the service is launched (`activateService()`) and then the takeover is made public to clients on both interfaces (`takeoverDownstream()`, `takeoverUpstream()`). Before moving to service active state

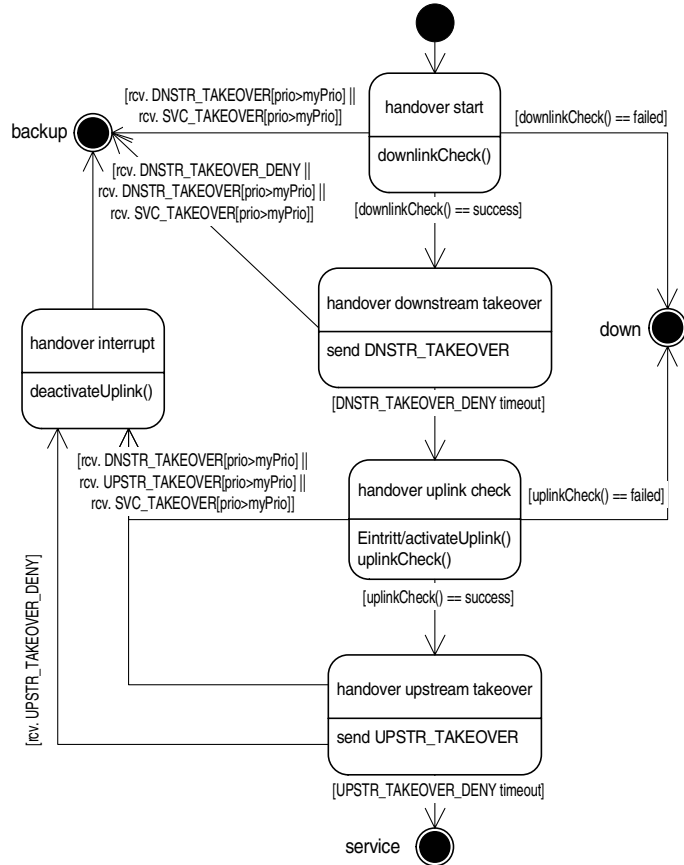


Figure 6: State Diagram of Handover Main State

the `SELF_CHECK_TIMER` is started. It is responsible for periodically triggering the self monitoring process similar to `SVC_CHECK_TIMER` for remote monitoring in backup main state (see Sec. 4.2.2). When a `SELF_CHECK_TIMER` alarm occurs, a full service check performed in service full check A state. Its sub-states are described at the end of this section.

Service full check B state is entered on arrival of a `DNSTR_TAKEOVER` message from a host with a lower priority than the own. As this indicates a takeover attempt on the downstream side by a backup host, a complete check is performed. In case all steps were successful, the takeover is considered to be unnecessary and consequently is vetoed by replying a `DNSTR_TAKEOVER-DENY` message. This process is only executed in case of a takeover attempt by a lower prioritized host, as this leaves the possibility of direct takeovers by new higher prioritized candidates entering the redundancy cluster. A similar process starts for the upstream side in case of a `UPSTR_TAKEOVER` message arriving.

As the master host is responsible for maintaining the status table, for each host in the redundancy cluster a `SVC_REQCHECK_TIMER` is present. An alarm of such a timer indicates that no heartbeat was recognized from the associated host. Therefore this host is considered to have left the redundancy cluster and the status table is updated accordingly.

A simple service check is carried out by `localSvcCheck()` on arrival of a `SVC_CHECK` message from a backup host. As this also indicates that this host is alive, it serves as a heartbeat signal. In consequence, the associated `SVC_REQCHECK_TIMER` is reset. When the incoming `SVC_CHECK` message has been sent by a new host in init main state (see Sec. 4.2.1), the new host is added to the status table. The result of the test is sent back by a `SVC_CHECK_REPLY` message.

Service main state may be left for several reasons. First, in case of failure detection (leaving to shutdown of the system) and second, when an explicit service shutdown is requested (returning to decide state in Fig. 3). In both cases service main state is left carrying out the same actions but with different succeeding states. A shutdown might result from one of the full service check states being left in failed state or a manual shutdown of the service. Explicit service takeover is signaled by another host from entering service main state, therefore it is detected by receiving a `SVC_TAKEOVER` message. Before leaving service main state, first the status table is updated and a `SVC_SHUTDOWN` message

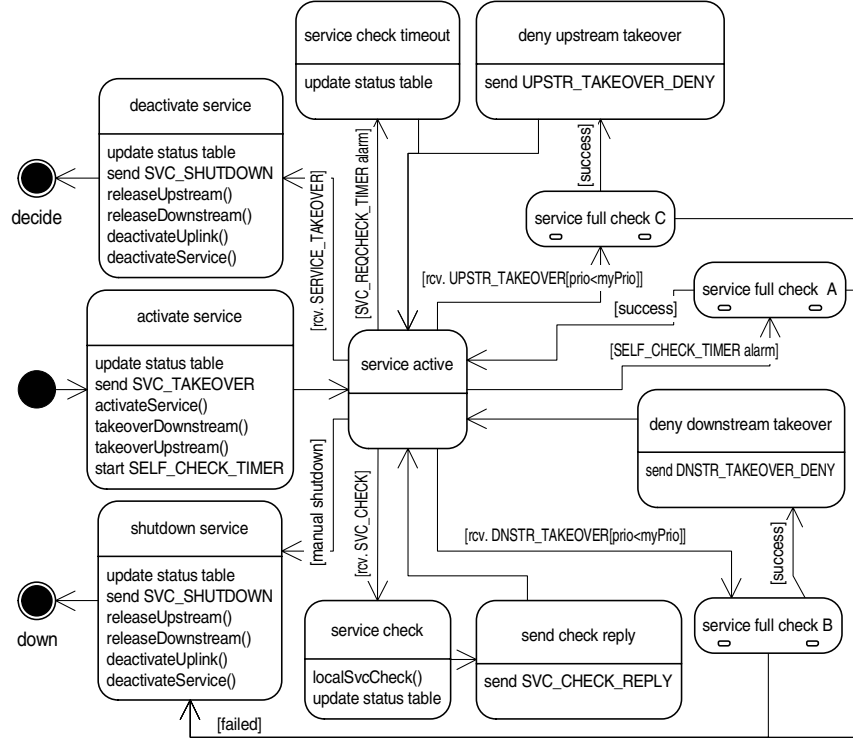


Figure 7: State Diagram of Service Main State

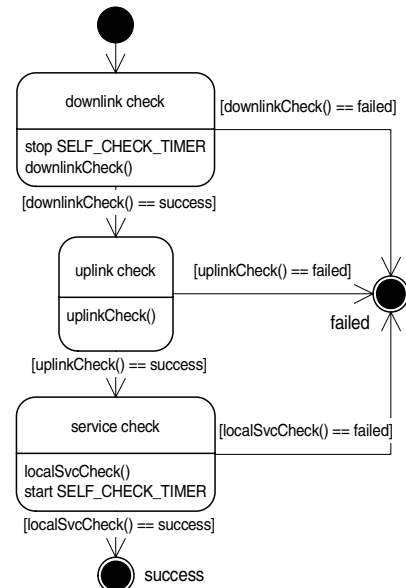


Figure 8: State Diagram of Service Full Check A|B|C Sub-States

Before leaving service main state, first the status table is updated and a `SVC_SHUTDOWN` message

is sent to all hosts. Clients are informed on both interfaces by `releaseDownstream()` and `releaseUpstream()`. Finally the uplink and the service itself are shut down by `deactivateUplink()` and `deactivateService()`.

Service Full Check A|B|C Sub-States: Although the three service full check states have different preceding / succeeding states, they have the same internal behavior in common. Therefore they are all specified in common in Fig. 8. First, in downlink check state the `SELF_CHECK_TIMER` is stopped and the downlink is tested via the customizable procedure `downlinkCheck()`. On success the uplink check state is entered, where the uplink is tested by `uplinkCheck()`. When the test was successful, too, the service is checked locally using `localSvcCheck()`. If this test succeeds the whole subprocess was successful. In case any of the checks failed, the subprocess is left immediately in failed state.

4.3 Messages and Timers

While semantics of messages exchanged between hosts in the redundancy cluster already became clear in the previous sections, Table 1 specifies data carried with the messages, whether they are uni- or multicast, the communication side they are exchanged on, and their principle source and destination. Note that although some messages are multicast ones in principle, this does not imply that communication services must be capable of true multicasting. As every host explicitly knows the other participants in the redundancy cluster through its local status table, true multicasting can be replaced by multiple unicast messages. Therefore the requirement of only having unicast communication primitives is not violated.

message	data	type	side	source	destination
<i>SVC_CHECK</i>	own prio.	UC	DS	BH	MH
<i>SVC_CHECK_REPLY</i>	status; status table	UC	DS	MH	host which prior sent <i>SVC_CHECK</i>
<i>SVC_SHUTDOWN</i>	status table	MC	DS	MH	all hosts in table
<i>DNSTR_TAKEOVER</i>	own prio.	MC	DS	BH	all hosts in table
<i>DNSTR_TAKEOVER_DENY</i>	-	UC	DS	MH	host which prior sent <i>DNSTR_TAKEOVER</i>
<i>UPSTR_TAKEOVER</i>	own prio.	MC	US	BH	all hosts in table
<i>UPSTR_TAKEOVER_DENY</i>	-	UC	US	MH	host which prior sent <i>UPSTR_TAKEOVER</i>
<i>SVC_TAKEOVER</i>	status table	UC	DS+ US	BH	all hosts in table
UC=unicast MC=multicast; DS=downstream US=upstream; BH=backup host MH=master host					

Table 1: Messages exchanged between hosts

To ensure proper operations regarding the logic of the distributed system, queuing of asynchronously arriving messages is necessary to preserve them for processing. E. g. in service main state, while processing a *SVC_CHECK* from one backup host, another *SVC_CHECK* message arriving from a different host must not be lost. Otherwise the second host might be removed from the redundancy cluster as its associated `SVC_REQCHECK_TIMER` might alarm, indicating that this backup host is not available anymore.

Table 2 shows a list containing the main state and the messages to queue there. It should be stressed that this does not include messages which are awaited synchronously as an answer to a prior request message, like the *SVC_CHECK* / *SVC_CHECK_REPLY* pair in backup main state.

main state	queued messages
backup	<i>DNSTR_TAKEOVER</i> , <i>SVC_SHUTDOWN</i> , <i>SVC_TAKEOVER</i>
handover	<i>DNSTR_TAKEOVER</i> , <i>UPSTR_TAKEOVER</i> , <i>SVC_TAKEOVER</i>
service	<i>DNSTR_TAKEOVER</i> , <i>UPSTR_TAKEOVER</i> , <i>SVC_TAKEOVER</i> , <i>SVC_CHECK</i>

Table 2: Messages to be queued

For messages possibly triggering a direct reply, timeout values have to be specified. Assuming that the roundtrip delay is some magnitudes smaller than the execution time of customizable monitoring procedures, lower bounds for their timeout read as follows:

$$\begin{aligned}
\text{timeout}_{\text{SVC_CHECK_DENY}} &> \text{tmax}_{\text{localSvcCheck}}() && (1) \\
\text{timeout}_{\text{DNSTR_TAKEOVER_DENY}} &> \text{tmax}_{\text{downstreamCheck}}() + \text{tmax}_{\text{upstreamCheck}}() + && (2) \\
&&& \text{tmax}_{\text{localSvcCheck}}() \\
\text{timeout}_{\text{UPSTR_TAKEOVER_DENY}} &> \text{tmax}_{\text{downstreamCheck}}() + \text{tmax}_{\text{upstreamCheck}}() + && (3) \\
&&& \text{tmax}_{\text{localSvcCheck}}()
\end{aligned}$$

Regarding local alarm timers, the values of `SELF_CHECK_TIMER` (see Sec. 4.2.4) and `SVC_CHECK_TIMER` (see Sec. 4.2.2) directly influence how promptly a service failure is detected. Their concrete values depend on complexity, load and duration of the respective checks and the local policy on acceptable service downtime. Therefore they must be determined individually for the concrete service in its particular environment. To avoid unnecessary load, `SVC_CHECK_TIMER`, which triggers remote monitoring in backup main state, should be set indirect proportional to the backup hosts priority. Doing so, a lower ranking host will check service functionality less often than the master's direct successor. Additionally it should not be smaller than `SELF_CHECK_TIMER` in any case.

4.4 Customizable Procedures

To adopt the generic state machine for a concrete service, several procedures are available for customization. While the redundancy logic remains untouched and is given by the generic state machine as of Section 4.2, only the customizable procedures are specific to a certain service.

Monitoring and Testing Procedures: The procedures showed in the following are all responsible for monitoring the various parts of the gateway services environment. They all return a boolean status, indicating whether the check was successful or not. All monitoring procedures must be implemented to deliver positive-definite results, meaning that positive answers must be valid, whereas erroneous negative results are acceptable. By specifying distinct procedures for service and communication monitoring an important requirement from Section 2 is fulfilled.

The `downlinkCheck()` and `uplinkCheck()` procedures are used in several states in the generic state machine. Their aim is to solely check network connectivity relative to respective reference points (Fig. 2).

In backup main state the `remoteSvcCheck()` procedure checks the services functionality from the remote side. Checks should be implemented from a clients side of view. This means that tests should act like a client, checking service functionality and connectivity as a whole.

In contrast `localSvcCheck()` solely is used to test the services functionality on the host currently providing the service. Hence additional service-internal data can be taken into account leading to the possibility of proactive monitoring.

(De-)Activation Procedures: The two pairs of `(de-)activateUplink()` and `(de-)activateService()` are responsible for activating and deactivating the network interface on the uplink side and the service itself. No monitoring/checking actions are carried out within these procedures

Client Related Procedures: To announce a change in the host providing the service to clients two procedures, each for upstream and downstream, are available (`takeoverDownstream()`, `takeoverUpstream()`, `releaseDownstream()`, `releaseUpstream()`). The takeover procedures should implement actions necessary to inform clients of the fact that a new master host is starting service provisioning. The release procedures can be used for activities towards clients necessary on shutdown of the service functionality.

4.5 Summary and Fulfillment of Requirements

In the previous sections we presented a robust and sound solution for management of redundancy in terms of high availability. Hereby, the requirements identified in Section 2 are fulfilled as follows.

The separation of logic from actions is accomplished by specifying the logic as a generic state machine (Sec. 4.2), while actions are noted as customizable procedures (Sec. 4.4). Consequently the state machine design makes no assumption whether a later implementation is done as an integrated part of a (new) service application or a standalone addition to an already existing application. Section 5 gives an example for implementing the state machine as additional component. Leaving concrete actions to be carried out in customizable procedures, the requirement for independence from specific services and communication techniques is also fulfilled.

Messages used to coordinate takeover — to ensure that only one single master host is present in the redundancy cluster and to define a total prioritization order of hosts — are exchanged on communication links present already for service provisioning. Thus no extra heartbeat link is presumed.

Furthermore the requirement of having only a minimal number of active upstream links is fulfilled. During normal service operation, all messages between hosts are exchanged via the downstream side (Sec. 4.3) and uplinks of a backup host are just activated shortly before becoming the master host (Sec. 4.2.3).

To gain service monitoring from the client's perspective, a host serving as a backup system actively tests the status of the service (Sec. 4.2.2). By this, monitoring of connectivity and service functionality as a whole is achieved. On the other hand, to enable proactive actions, the active master host additionally does local service monitoring (Sec. 4.2.4), taking internal conditions of the service into account. Especially in case of an integrated implementation, where internal data are more easily observable, monitoring capabilities are substantially extended compared to external monitoring only.

By specifying different customizable procedures for service monitoring and communication link monitoring the distinction of monitoring types is realized (Sec. 4.4).

Independence from technology-specific communication primitives are addressed in Section 4.3 and Section 4.4.

5 Application to a Linux-based Firewall

Applying the generic state machine to solve the single point of failure problem induced by the Linux-based packet filtering firewall of our institute's infrastructure resulted in the implementation described in this section.

Implementing the generic state machine for use in a concrete service consists of three main steps:

1. Implementing the state machine
2. Implementing the message communication mechanism according to the scenario's specific circumstances
3. Implementing the customizable procedures

The following sections show how these steps were carried out in our scenario.

5.1 Universal IP-Service Daemon

Performing steps 1 and 2 from above lead to a universal application for IP-based services. It was implemented in C as a background daemon on Linux. For message implementation simply UDP packets were used with an additional symmetric encryption. All hosts in the redundancy cluster share the same key. Configuration is done on file basis. The configuration files need to be distributed manually among the hosts.

Knowing the fact that IP-communications is used, parts of step 3 also can be realized already. Transparency to clients is solved by using a roving pair of IP-addresses, each one for up- and downstream side. Therefore all hosts in the redundancy cluster own a host-specific pair of IP-addresses for message exchange and testing in backup and handover state. In service state the roving pair of IP-addresses is additionally bound to the interfaces (this capability is often referred as "single link multihoming" or "IP-aliasing") by `takeoverDownstream()` and `takeoverUpstream()`. Additionally, to announce the change of the host which owns the roving IP-addresses, a broadcast ping on both sides is sent, giving clients an opportunity to notice the change and re-adapt their IP / Layer-2 translation mechanism (e.g. change ARP table entries in case of Ethernet). One of the main advantages of the roving IP-addresses scheme is that IP-related service configuration remains independent of the actual host, permitting simple configuration replication among backup hosts.

The `downlinkCheck()` is implemented as a broadcast ICMP ping on the downstream IP-network, claiming the connection to be healthy when at least one foreign reply arrives. The upstream side is checked in `uplinkCheck()` by sending an ICMP ping on the upstream interface to the next hop router.

Having only made the assumption of an IP-based infrastructure the pair of `activateUplink()` / `deactivateUplink()` cannot be determined as they are specific for the underlying layer 2 technology. Their implementation is delegated to external programs (e.g. shell scripts) configurable in the configuration file. The same applies for service related checking procedures (`remoteSvcCheck()` and `localSvcCheck()`) and service (de-)activation (`((de-)activateService())`).

The next section will show how those delegated procedures were implemented for a firewall service.

5.2 Tailoring for Firewall

For use in our institute's firewall scenario the above described daemon was enhanced by some shell scripts. As Ethernet is used on the upstream side, the scripts for `(de-)activateUplink()` (`un-)load` the kernel

driver modules for the Ethernet-card, realizing a total link deadness on deactivation of the link. Additionally IP-configuration of the interface is done here. Service activation is done by enabling IP forwarding, adapting the routing table and setting up the firewall rules via `iptables`. Deactivation is done by the respective inverse actions. Remote service checking in `remoteSvcCheck()` is done by an ICMP ping from the downlink interface to the next hop router on the current firewall host's upstream side. On the master host `localSvcCheck()` is realized by comparing currently active filter rules to the expected ones.

6 Conclusion and Further Work

Services which are provided and connected to one domain, but also are linked to a foreign domain, so called gateway services, demand special requirements to fault tolerance in terms of availability. As current solutions are restricted to specific services and/or violate one of the requirements for gateway services, no generic solution for high availability is available.

The introduced generic state machine gives a solution for any gateway services by separating monitoring and failover logic from individual actions specific for concrete services. Its applicability has been proved by implementing a universal daemon for IP-based gateway services and its deployment for a Linux-based firewall.

Moreover this paper could serve as a first starting point for standardization. By reusing the logic described by the generic state machine in conjunction with an in-depth specification of the customizable procedures for a concrete service type, a vendor-independent and interoperable failover mechanism for this service type could be specified.

Although an applicable solution is provided, there is room for further improvements. First, security considerations regarding authorization and authentication in message delivery between redundant hosts will be specified at the level of the generic state machine. For this purpose [RFC 2338, HMHH+ 02] and [Robe 00] will provide a valuable basis. By now message security topics are left to the responsibility of a concrete implementation. Furthermore, to setup redundancy clusters where hosts serve as a backup for multiple services, the generic state machine will be extended to incorporate the notion of a service type, enabling to distinct between different services. Included are questions regarding load balancing and active feedback of backup hosts to influence their priority ranking in the redundancy cluster.

Acknowledgment The author wishes to thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of this paper. The MNM Team directed by Prof. Dr. Heinz-Gerd Hegering is a group of researchers of the University of Munich, the Munich University of Technology, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. Its web-server is located at <http://www.mnmteam.informatik.uni-muenchen.de/>.

References

- [Cass 01] CASSEN, A.: *Linux Virtual Server High Availability using VRRPv2*. Technical Report, nov 2001, <http://www.keepalived.org/pdf/LVS-HA-using-VRRPv2.pdf> .
- [HALP] *High-Availability Linux Project*, <http://linux-ha.org/> .
- [HAN 99] HEGERING, H.-G., S. ABECK and B. NEUMAIR: *Integrated Management of Networked Systems – Concepts, Architectures and their Operational Application*. Morgan Kaufmann Publishers, ISBN 1-55860-571-1, 1999.
- [HaRa 01] HAUCK, R. and I. RADISIC: *Service Oriented Application Management — Do Current Techniques Meet the Requirements?* In *New Developments in Distributed Applications and Interoperable Systems, Proceedings of the 3rd IFIP International Working Conference (DAIS 2001)*, Krakow, Poland, September 2001. Kluwer Academic Publishers, <http://www.nm.informatik.uni-muenchen.de/Literatur/MNMPub/Publikationen/hara01/hara01.shtml> .
- [Hind 02] HINDEN, R.: *Virtual Router Redundancy Protocol for IPv6*. Internet Draft, IETF, February 2002, <http://www.ietf.org/internet-drafts/draft-ietf-vrrp-ipv6-spec-02.txt> .
- [HMHH+ 02] HINDEN, R., D. MITZEL, P. HUNT, P. HIGGINSON, M. SHAND, A. LINDEM, S. KNIGHT, D. WEAVER and D. WHIPPLE: *Virtual Router Redundancy Protocol*. Internet Draft, IETF, February 2002, <http://www.ietf.org/internet-drafts/draft-ietf-vrrp-spec-v2-06.txt> .
- [IETFWG-RSP] *Reliable Server Pooling (rserpool)*, <http://www.ietf.org/html.charters/rserpool-charter.html> . IETF Working Group Charter.
- [IETFWG-VRRP] *Virtual Router Redundancy Protocol (vrrp)*, <http://www.ietf.org/html.charters/vrrp-charter.html> . IETF Working Group Charter.
- [LSXS 02] (ED.), J. LOUGHNEY, M. STILLMAN, Q. XIE and R. STEWART: *Comparison of Protocols for Reliable Server Pooling*. Internet Draft, IETF, March 2002, <http://www.ietf.org/internet-drafts/draft-ietf-rserpool-comp-03.txt> .
- [LVS] *Linux Virtual Server Project*, <http://www.linuxvirtualserver.org/> .

- [RFC 2281] LI, T., B. COLE, P. MORTON and D. LI: *RFC 2281: Cisco Hot Standby Router Protocol (HSRP)*. RFC, IETF, March 1998, <http://www.ietf.org/rfc/rfc2281.txt> .
- [RFC 2338] KNIGHT, S., D. WEAVER, D. WHIPPLE, R. HINDEN, D. MITZEL, P. HUNT, P. HIGGINSON, M. SHAND and A. LINDEM: *RFC 2338: Virtual Router Redundancy Protocol*. RFC, IETF, April 1998, <http://www.ietf.org/rfc/rfc2338.txt> .
- [RFC 2787] JEWELL, B. and D. CHUANG: *RFC 2787: Definitions of Managed Objects for the Virtual Router Redundancy Protocol*. RFC, IETF, March 2000, <http://www.ietf.org/rfc/rfc2787.txt> .
- [RFC 2960] STEWART, R., Q. XIE, K. MORNEAULT, C. SHARP, H. SCHWARZBAUER, T. TAYLOR, I. RYTINA, M. KALLA, L. ZHANG and V. PAXSON: *RFC 2960: Stream Control Transmission Protocol*. RFC, IETF, October 2000, <http://www.ietf.org/rfc/rfc2960.txt> .
- [RFC 3237] TUEXEN, M., Q. XIE, R. STEWART, M. SHORE, L. ONG, J. LOUGHNEY and M. STILLMAN: *RFC 3237: Requirements for Reliable Server Pooling*. RFC, IETF, January 2002, <http://www.ietf.org/rfc/rfc3237.txt> .
- [Robe 00] ROBERTSON, A.: *Linux-HA Heartbeat System Design*. In *4th Annual Linux Showcase & Conference (ALS 2000)*, Atlanta, USA, October 2000. , <http://www.linuxshowcase.org/2000/2000papers/> .
- [SXS 02] STEWART, R. R., Q. XIE and M. STILLMAN: *Aggregate Server Access Protocol (ASAP)*. Internet Draft, IETF, May 2002, <http://www.ietf.org/internet-drafts/draft-ietf-rserpool-asap-03.txt> .
- [TXSS+ 02] TUEXEN, M., Q. XIE, R. STEWART, M. SHORE, L. ONG, J. LOUGHNEY and M. STILLMAN: *Architecture for Reliable Server Pooling*. Internet Draft, IETF, April 2002, <http://www.ietf.org/internet-drafts/draft-ietf-rserpool-arch-02.txt> .
- [Zhan 00] ZHANG, W.: *Linux Virtual Server for Scalable Network Services*. In *Proceedings of Ottawa Linux Symposium 2000*, July 2000, <http://www.linuxvirtualserver.org/ols/lvs.ps.gz> .