CHAPTER 2

Introduction to High Availability

This chapter describes the high availability (redundancy) and protection options available for units and gateways:

- High Availability Overview, page 2-1—Provides an overview of high availability in the Cisco ANA fabric.
- Watchdog Protocol, page 2-2—Describes the Watchdog protocol that monitors the processes on the units.
- Unit N+m High Availability, page 2-2—Describes the clustered N+m high availability mechanism within the Cisco ANA fabric designed to handle the failure of units.
- Limitations and Restrictions, page 2-3—Describes the restrictions and limitations relating to high availability.

High Availability Overview

High availability is the provision of multiple interchangeable components to perform a single function to cope with failures and errors.

The high availability architecture is designed to ensure continuous availability of assurance and fulfillment functionality, by detecting, and recovering from a wide range of hardware and software failures, such as failures in the server machines, connectivity, software breakdowns and so on.

The distributed design of the system enables the “impact radius” caused by a single fault to be confined. This prevents all types of fault from setting into motion the “domino” effect, which can lead to the meltdown of all the management services.

The high availability of the server backbone is achieved at several complementing levels, namely:

- NEBS-3 compliant carrier-class server hardware.
- Internal watchdog within each unit, in charge of monitoring (and if necessary automatically reloading) failed processes. For more information see Watchdog Protocol, page 2-2.
- N+m warm standby protection for units clusters. For more information see Unit N+m High Availability, page 2-2.

Watchdog Protocol

Each unit executes several processes: one control process and several Agent Virtual Machine (AVM) processes that execute Virtual Network Elements (VNEs). Each process within the unit is completely independent. The isolation concept is tailored throughout the design: a failure of a single process does not affect other processes on the same machine. The exact number of processes on each unit depends on the capacity and computation power of the unit.

The control process executes a Watchdog protocol, which continuously monitors all other processes on the unit. This Watchdog protocol requires each AVM process to continuously handshake with the Control process. A process that fails to handshake with the control process after a number of times (namely, is "stuck") will be automatically killed and reloaded. All the Watchdog protocol parameters are configurable by the operator.

The dynamic design of the control process implements runtime adaptation and escalation. The escalation procedure moves the AVM to suspended mode, namely, the process is suspended. An example of an escalation procedure is to stop reloading a process that has crashed more than **N** times within a given period, as it is suspected of having a recurring software problem.

The reload process is local to the unit, and thus very rapid, with a minimal amount of downtime. Since the process can use its previous cache information (temporary persistency used to improve performance), once the stuck process is detected, reloading the process takes only a few seconds with no data loss.

All Watchdog activity is logged, and an alarm is generated and sent when the watchdog reloads a process.

Unit N+m High Availability

The clustered N+m high availability mechanism within the Cisco ANA fabric is designed to handle the failure of a unit. Such failures include hardware failures, operating system failures, power failures, or network failures, which disconnect a unit from the Cisco ANA fabric.

Unit availability is established in the gateway, running a Protection Manager process, which continuously monitors all the units in the network. Once the Protection Manager detects a unit that is malfunctioning, it automatically signals one of the m servers in its cluster to load the configuration of the faulty unit (from the system registry), taking over all its managed network elements. This design provides many possibilities for trading off protection and resources. These possibilities range from just segmenting the network into clusters without any extra machines, up to having a warm-swappable empty unit for each and every unit in the setup. It is recommended that units are clustered according to geography and that an additional empty unit is added to heavily loaded clusters.

The switchover of the redundant standby unit does not result in any loss of information in the system, as all the information is auto-discovered from the network, and no persistent storage synchronization is required. Hence, the redundant standby unit relearns all the information from the network elements, with no danger of persistent information corruption. Furthermore, where there is cluster saturation (namely, more than one unit in a cluster fails at the same time and there are no extra machines), the remaining units will continue to operate and manage their network scope normally.

When a unit is configured it can be designated as being an active or standby unit. The active units (excluding the standby unit) that are connected to the gateway are known as a protection group. The standby unit that is configured for the gateway is linked to that protection group. The administrator can define more than a single protection group. Each protection group defined has a set of protected units and a protecting standby unit.

The following example shows a protection group (cluster) of units, controlled by a gateway with one unit configured as the standby for the protection group.

Figure 2-1 Cisco ANA Architecture

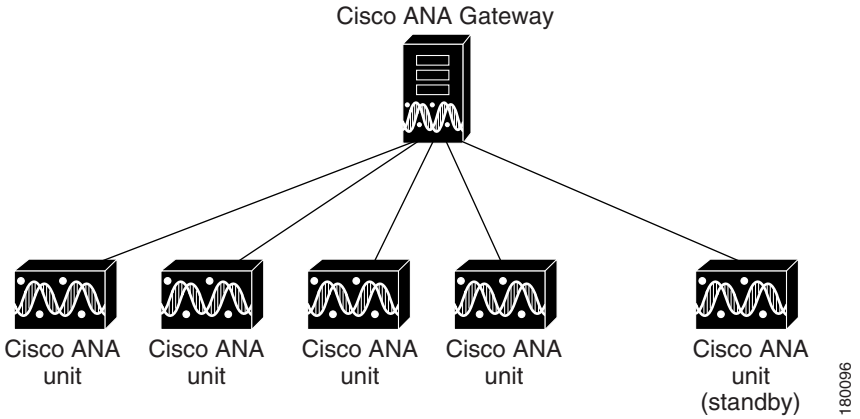

In the above configuration, when the gateway determines that one of the units in the protection group has failed, it notifies the protection group's standby unit to immediately load the configuration of the failed unit. The standby unit loads the configuration of the failed unit, including all its AVMs and VNEs, and functions as the failed unit.

These events are all recorded in the EventVision system log, which enables the user to take the necessary action to bring the failed unit up again. When the failed unit becomes operational, the user can decide whether to configure it as the new standby unit or to reinstate it to the protection group and configure another unit as the standby unit.

Limitations and Restrictions

The high availability mechanism will attempt to load an AVM after it crashes (whether the AVM comes up or not), a maximum of seven times. Thereafter, the high availability mechanism will not try to reload this AVM again.

Related Documentation

For more detailed information see the following publications:

- Cisco Active Network Abstraction Administrator Guide
- Cisco Active Network Abstraction NetworkVision User Guide
- Cisco Active Network Abstraction EventVision User Guide

Note Changes to the registry should only be carried out with the support of Cisco Professional Services.