



TwinStar™: How it Works

Document version: 1.0

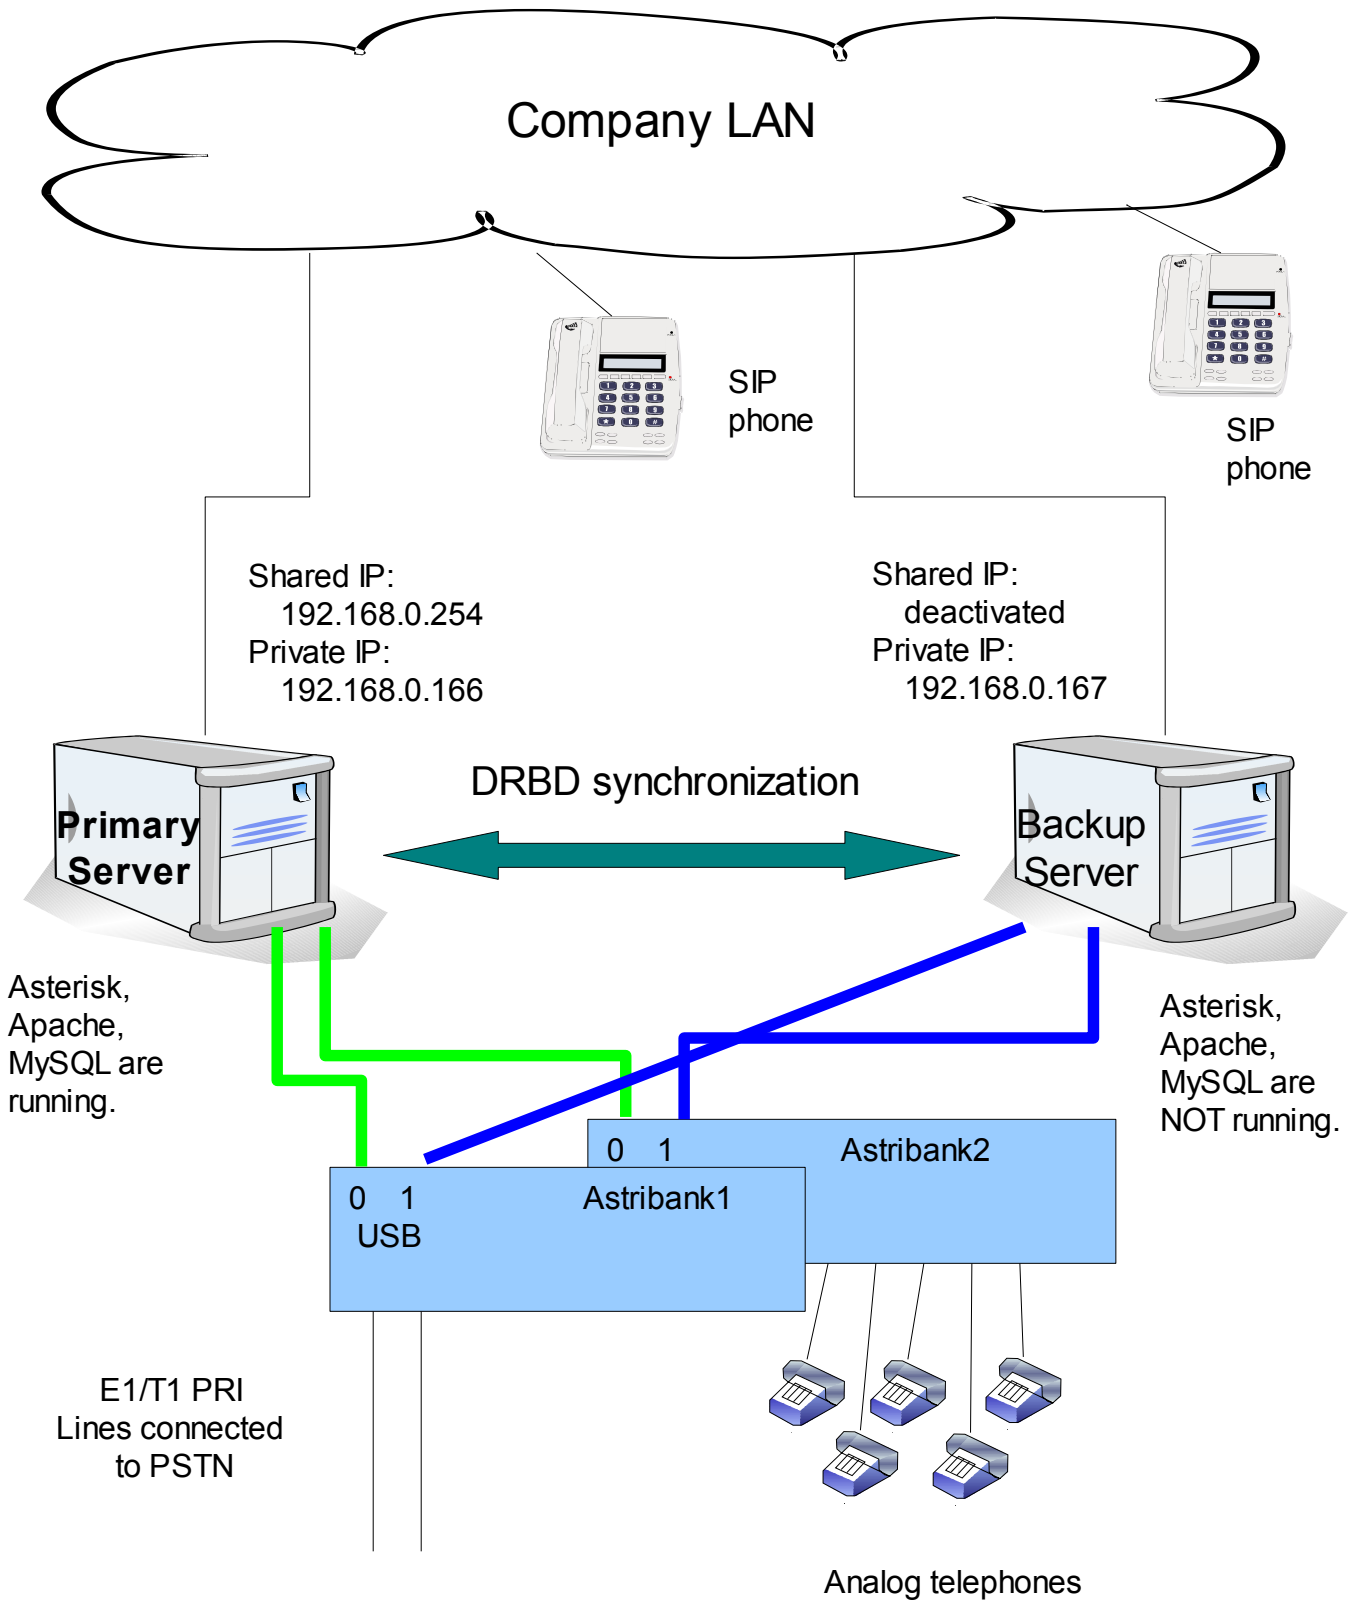
Modification: Aug 18, 2009

Preface

This document describes the Atribank TwinStar operation and configuration principles. It does not describe the specific step-by-step procedure for Asterisk high availability cluster building.

An Example of an Asterisk High Availability Cluster

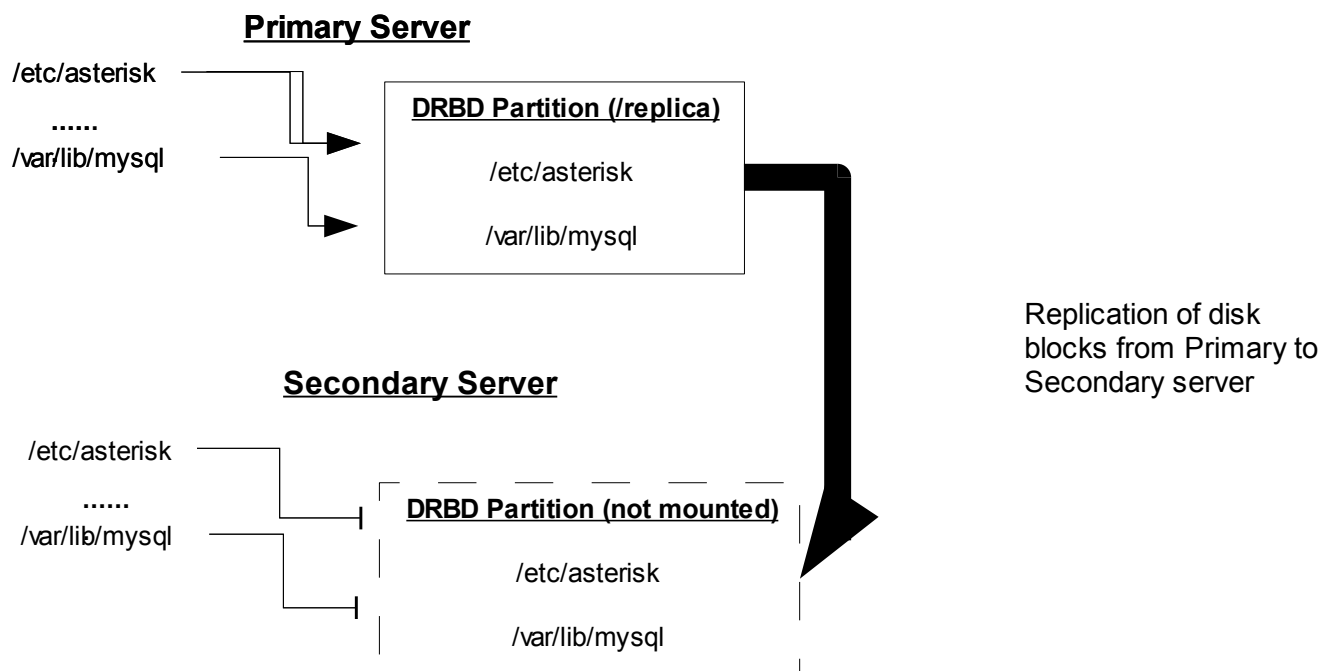
Two Asterisk servers can comprise an Asterisk cluster. One of the servers acts as the Primary server and the other server acts as the Secondary (hot-backup) server. Each server has its own private IP address. This allows server management and configuration replication to be performed between the two servers via the IP network. The VoIP PBX extensions communicate with the PBX cluster via a special public (or cluster) IP address, which is activated only on the current Primary server. When the servers swap roles, for example due to the Primary server failure, then that cluster IP address will be activated automatically on the Secondary server. As a result, the IP phones will continue to use the same PBX IP address.



The replication of the servers' configuration files may be performed by using DRBD (“Distributed Replicated Block Device”). Visit the <http://www.drbd.org> Web site for further details. Nowadays, the DRBD is widely used for building many different high availability clusters. The DRBD concept assumes that there are special hard disk partitions that are replicated on the physical level via the IP network. Thus a changed hard disk block on the Primary server will be replicated to the Secondary server.

All files that need to be replicated must be allocated in the DRBD partition. For example, usually it is necessary to replicate the Asterisk configuration files, MySQL database files, etc. Therefore, folders such as `/etc/asterisk` and `/var/lib/mysql` must be allocated in a DRBD partition and soft links need to be created instead of the regular file names. For example, the `/etc/asterisk` will be a link that points to `/replica/etc/asterisk`, where “`/replica`” is the default mount name of the DRBD partition.

Please note that the DRBD partition (`/replica`) is mounted on the Primary server only. On the Secondary server the DRBD partition is not mounted and therefore the `/etc/asterisk` files as well as other replicated files are not available. Services that depend on those files do not run (Asterisk, MySQL server etc). This ensures that the replicated files cannot be changed simultaneously on both servers.





In addition to the data replication function, a tool that detects the Primary server failure and forces the Secondary server to takeover the cluster services is required. In many cases, High Availability (HA) software is used. Visit the <http://www.linux-ha.org/> Web site for further details. The HA uses its own IP watchdog protocol to detect the active server failure and activate one of the backup servers. Atribank TwinStar technology allows the implementation of similar cluster management tools.

Each Atribank device has two USB interfaces (0 and 1). Once powered up, the Atribank will always use the USB interface 0, which should be connected to the Primary server. After all of the services have started it is possible to activate the TwinStar watch dog mechanism that monitors (a) voltage availability on the USB port and (b) continuity of communication between the Atribank device and the Atribank DAHDI driver.

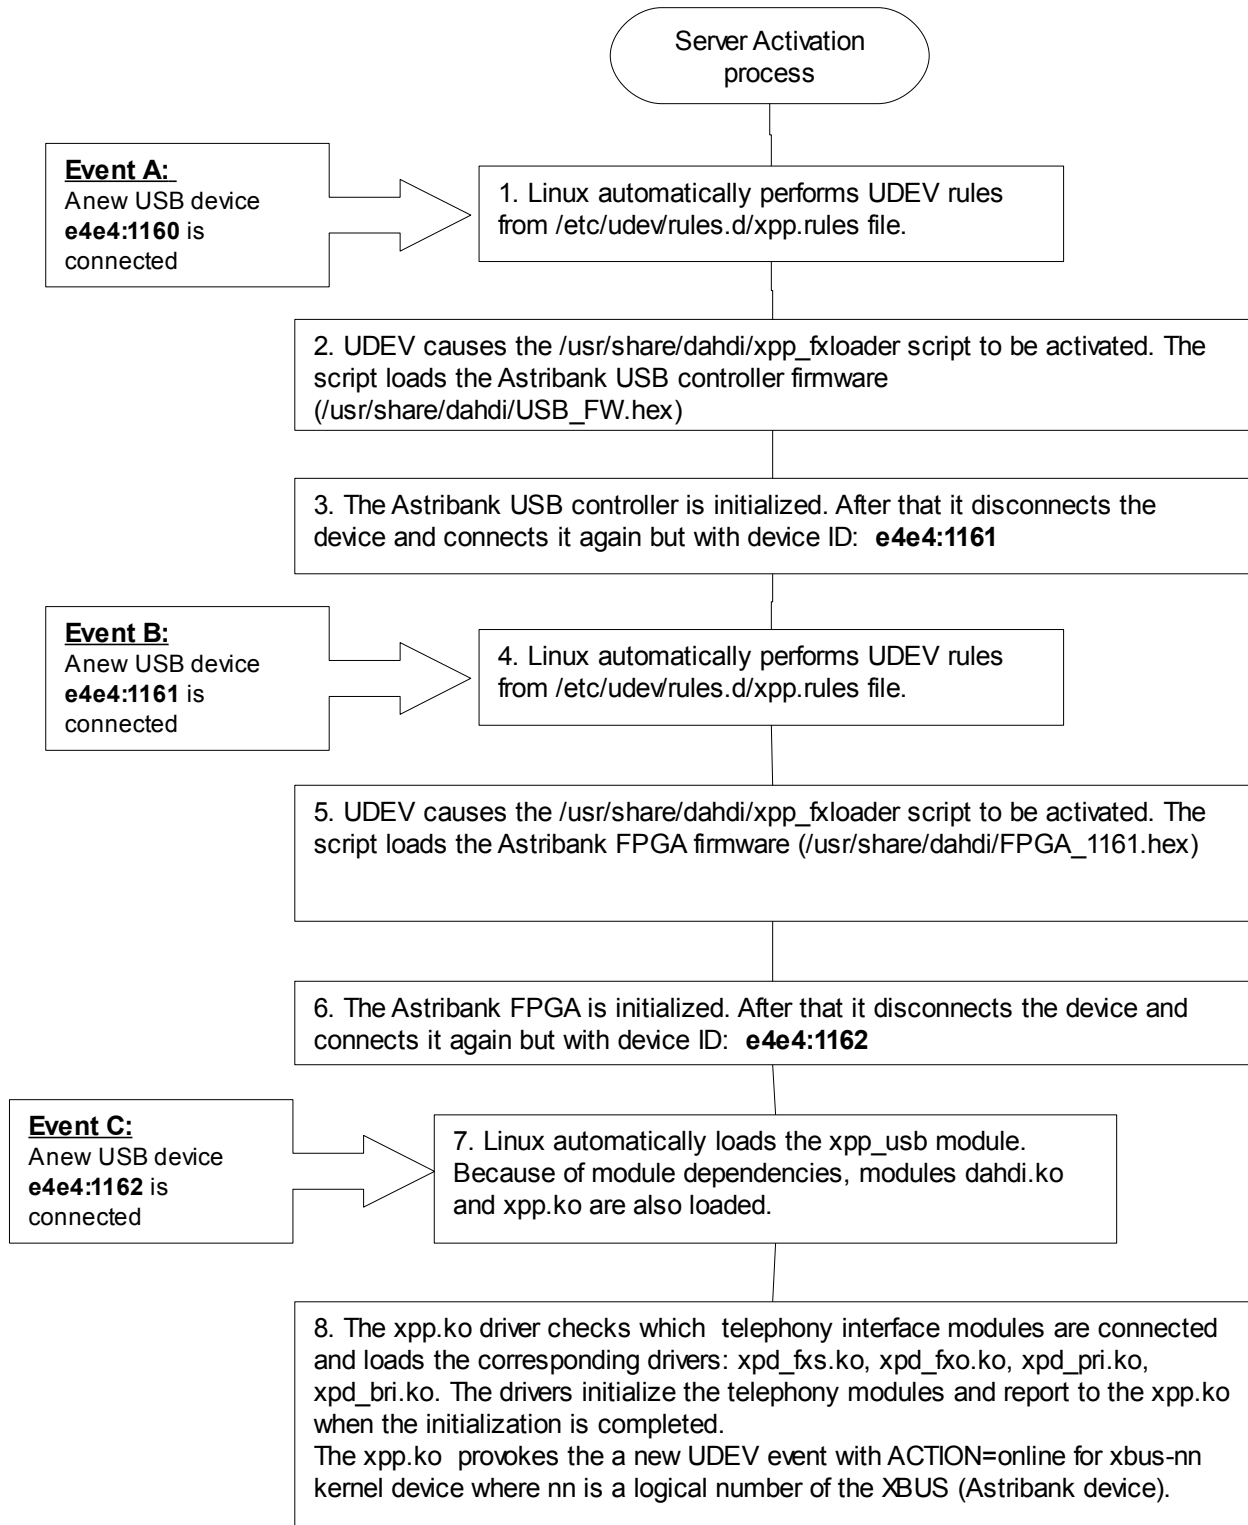
When at least one of the conditions is not met, the Atribank will deactivate the USB interface 0 and automatically activate the USB interface 1. As a result, the Secondary server receives a plug-and-play event that a new USB device is connected to the server. The plug-and-play system calls special scripts that check that all of the configured Atribank devices have switched to the Secondary server. After that, a special script is called and the following occurs:

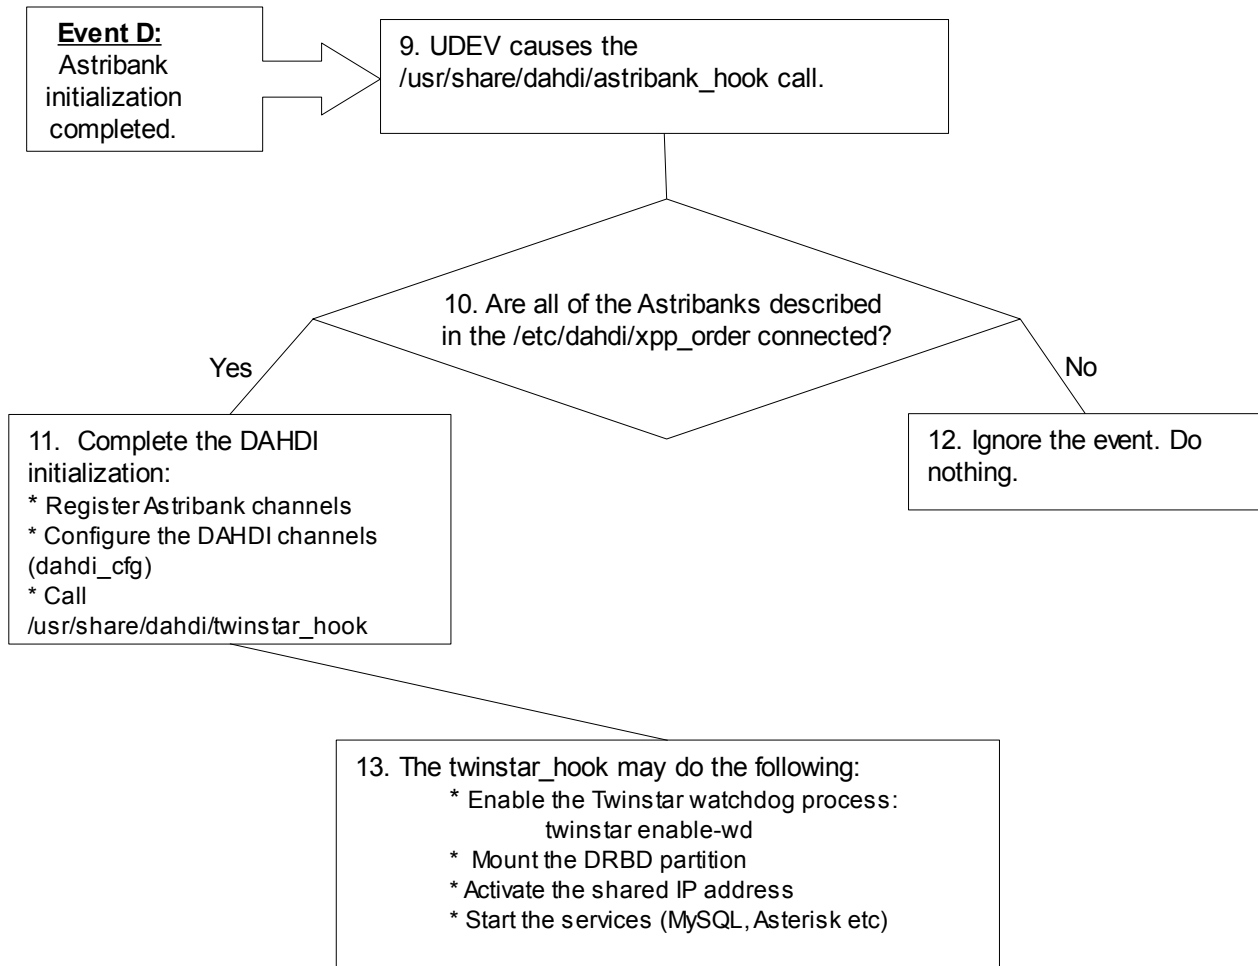
1. The script registers the Atribank channels on DAHDI and configures them.
2. The script mounts the DRBD partition (/replica). As a result, the /etc/asterisk, /var/lib/mysql etc. folders become available.
3. The shared IP is activated.
4. MySQL, Asterisk, and Apache services are started.

TwinStar Implementation Details¹

The Atribank device is plug-and-play by nature. When an Atribank device is discovered by the Linux plug-and-play system and, based on the predefined UDEV rules (/etc/udev/rules.d/xpp), the /usr/share/dahdi/xpp_fxloader script is called. This utility is responsible for loading the Atribank firmware. Once the firmware has been loaded Linux loads the Atribank drivers that initialize the device. After that, a special UDEV event is generated and as a result, the /usr/share/dahdi/atribank_hook script is called. The atribank_hook may check if all of the required Atribank devices have been activated. In order to do so the script inspects the /etc/dahdi/xpp_order file which contains a sorted list of configured Atribank devices. The script also checks the /sys file system for existence and status of connected Atribanks. If all of the devices listed in the xpp_order file are active, then the atribank_hook may complete the DAHDI start-up procedure and start all of the necessary services. For example, to mount the DRBD partition, activate the shared IP, start Asterisk, etc. Those services actually may be started from a separate twinstar_hook script that is called by the atribank_hook. The flow chart below provides a detailed explanation of the process.

1 The technical details described here are subject to change.





As you can see, the proposed solution assumes that the services such as Asterisk, MySQL, Apache, etc. -- which are usually started automatically during the Linux start-up--must be started only on the Primary server and only when DAHDI has been successfully started (when all of the Atribanks are active) and their configuration files are available (the DRBD partition is mounted).

Which Failure Events May be Caught

As was mentioned above, the Astribank TwinStar monitors the voltage of the active USB port and controls the regularity of the watchdog messages exchanged between the device and the Astribank driver. Therefore, the Astribank will switch over to the second USB interface when one of the following happens:

- the server is switched off
- the Astribank driver stopped working, for any reason. For example, the Linux kernel crashed, the driver was unloaded, etc.
- the USB cable is disconnected.

Switching of the servers can be made conditional upon other events as well. For example, another server monitoring system like Monit (<http://mmonit.com/monit>), Runit (<http://smarden.sunsite.dk/runit>) or even the Linux Heart Beat (<http://www.linux-ha.org>) can be installed and configured in such a way that Astribank switching will take place when the monitoring system discovers a problem.

How to Control the TwinStar Status

You can use the 'twinstar' utility that is located in the /usr/sbin folder.

1. In order to check status:

```
#twinstar status
DEVICE          PORT      WATCHDOG      POWER0      POWER1
usb:001/006     0         on             yes         yes
usb:001/007     0         on             yes         yes
```

2. In order to enforce switching to another USB interface:

```
twinstar jump
```

3. In order to start USB voltage monitoring and exchange of watchdog messages:

```
twinstar enable-wd
```



4. In order to stop USB voltage monitoring and exchange of watchdog messages:

```
twinstar disable-wd
```

The watchdog should be disabled before stopping the DAHDI for maintenance purposes.

How to Install TwinStar

Please note that some TwinStar-related utilities and scripts are not installed by default during regular DAHDI installation procedure and need to be installed manually (this will be changed in future).

Notes:

Xorcom Ltd. does not provide technical support for DRBD or any server monitoring software/hardware packages.

Xorcom provides pre-configured DRBD/TwinStar-ready XR2000/XR3000 Asterisk appliances in its product offering.

Trademarks

Asterisk is a registered trademark of Digium, Inc. TwinStar, Atribank, XR1000, XR2000, and XR3000 are trademarks of Xorcom, Ltd. All other marks are the property of their respective holders.