Product Description

Convolo™ Cluster DataGuard Edition Version 2.0

August 2001

Welcome to Linux's leading high availability cluster software product – Convolo Cluster DataGuard Edition. DataGuard was designed and built specifically for Lunix by the most industry proven group of cluster developers. If you have just purchased the product, thank you. We are confident that you will be completely satisfied with the software and appreciate the benefits it brings to your Linux environment. If you are considering purchasing DataGuard cluster software, we hope that this document will provide the answers to any technical questions you may have.

This product description provides a technical overview of DataGuard cluster software and its features. It also outlines supported hardware, prerequisite software, and other information that will be useful as you configure your Linux cluster. Information regarding the product Warranty and Support Services is also included.

Sincerely,

Mission Critical Linux, Inc.
100 Foot of John Street
Lowell, Massachusetts 01852
`www.missioncriticallinux.com`
(+1) 978.606.0200
(+1) 877.62LINUX

Description

DataGuard cluster software provides a high availability computing environment for multiple Linux-based systems running static (read-only) and dynamic update (read/write) applications. It provides support for two systems accessing shared mass storage (not including boot disks) under a single application management environment. Within this highly integrated environment, systems retain their independence because they use private copies of the Linux operating system. Thus, the Linux systems can boot and shut down independently while benefiting from common storage resources. Applications running on servers in a DataGuard cluster configuration access shared disk resources in a fully coordinated manner, guaranteeing the highest levels of data integrity.

Because disk storage resources are shared, DataGuard cluster configurations offer higher availability than standalone Linux systems. Properly configured DataGuard cluster configurations can withstand the shutdown or failure of various components. For example, if one server in a DataGuard cluster is shut down for planned maintenance, applications and network clients can continue operation after restarting on or reconnecting to the remaining cluster member. Because mass storage can be shared cluster wide, restarted applications and clients are able to access their original data. Applications and network clients can be designed to survive these events automatically. For example, if clients, such as NFS, are designed correctly, they do not have to reconnect on failover.

Hardware configurations can be designed to have no single-point-of-failure, and will provide very high levels of application availability. DataGuard cluster software provides comprehensive integrity guarantees for applications and data through multiple polling mechanisms and programmable power management.

Features

DataGuard cluster software includes the following features:

- **Heterogeneous File Sharing with NFS and Samba –** DataGuard leads the industry with the only complete NFS V2 and V3 high-availability solution, allowing client access through NFS for UNIX or Linux clients. Only DataGuard maintains full access permissions and client lock preservation through failover.

  DataGuard is the most flexible file server available today with NFS and SMB/CIFS (Samba) support completely integrated. DataGuard delivers Samba network file sharing and print serving capabilities to UNIX and Windows clients.

  By combining DataGuard with NFS or Samba, you can configure a scalable and highly available client/server environment that maintains data integrity while sharing data. Optionally, you can use the ReiserFS journalled file system to minimize failover times of disk based services.

- **High availability for Commercial Applications** – DataGuard enables you to make many data-intensive commercial applications highly available. You can set up cluster services for Oracle, DB2 EE, PostgreSQL, Informix, Sybase, and MySQL databases, and unlike other Linux clustering products, there is never an extra charge.

- **Support for Major Linux Distributions** – DataGuard supports all the major Linux distributions, including both RPM and DEB-based distributions. Although not required, it is recommended that you run the same Linux distribution on each member system.

- **Commodity Hardware Configurations** – DataGuard supports commodity hardware in a variety of cluster configurations, ranging from a low-cost minimum configuration to a no-single-point-of-failure configuration. Both SCSI and Fibre Channel storage can be used in a cluster. You can choose the cluster

hardware that is right for the availability and data integrity needs of your business.

- **Easy Cluster Software Installation and Configuration** – Mission Critical Linux provides the DataGuard cluster software in a binary distribution that makes time-consuming and error-prone manual builds unnecessary. To accommodate popular Linux distributions, both RPM and DEB-based packages are available. DataGuard also provides a comprehensive cluster configuration verification utility and commands that check quorum partition accessibility and power switch operation.

  The cluster software includes configuration scripts and user interfaces that guide you through the process of setting up the cluster and services, and then add the site-specific information to the cluster database on the shared quorum partitions. Whenever you modify the cluster configuration, the database is updated and the changes are automatically propagated to all the member systems, elminating the need to perform duplicate operations on each member system.

  In addition, cluster event logging enables you to detect and resolve cluster problems before they affect service availability. The cluster daemons log messages by using the `syslog` subsystem. You can customize the severity level of the logged messages and also the log file location.

- **Comprehensive Documentation** – To facilitate setting up and managing a DataGuard cluster, Mission Critical Linux provides comprehensive documentation in the printed DataGuard *Installation and Administration* manual. The manual provides step-by-step installation and configuration instructions, complete with diagrams and detailed information about shared storage configurations. An online version of the manual is located on the DataGuard CD-ROM. In addition, the latest release notes are location on the `www.missioncriticallinux.com` Web site.

- **Multiple User Interfaces for Cluster Management** – DataGuard includes multiple user interfaces for easy service and cluster administration. The `cluadmin` utility provides an interactive text based user interface and also has a command line interface that is suitable for shell scripts. The Web-based graphical user interface (GUI) enables you to perform cluster operations from any browser. Cluadmin's command line interface facilitates usage in scripts as well as makes DataGuard well suited for deployment in embedded configurations.

- **Service Configuration Framework** – DataGuard provides an easy-to-use framework for creating services for heterogeneous file sharing and database applications.

  Service failover is accomplished by scripts that start and stop an application or, in the case of an NFS or Samba service, redirect an IP address to a functioning member system. DataGuard includes kernel patches which are required in order to provide NFS failover operation which is transparent to clients. DataGuard also provides all the scripts necessary for failing over NFS and Samba services. For database services, you can use sample service scripts and templates. In addition, the service configuration framework can easily be extended to other applications, such as mail and print applications.

- **Network Interface Failover** – DataGuard supports network interface failover on a member system. If a network interface fails, DataGuard can fail over the network I/O and an IP address to another interface, with no disruption to the cluster, clients, or applications. In addition, you can configure

DataGuard so that it fails over a medium access control (MAC) address.

- **Robust Cluster Communication Mechanisms** – To enable failover to occur quickly when a problem occurs, but also prevent premature or erroneous failover, DataGuard uses various monitoring mechanisms in order to accurately assess the status of a member system. Periodically, each member system writes system status information ("up" or "down") and a timestamp to the two quorum partitions located on shared disk storage. A system can join the cluster only if it is able to write to both quorum partitions. In addition, each member system periodically monitors the information written by the other member system to determine whether to fail over its services.
- **Service Failover** – To ensure proper cluster operation and maintain data integrity, if a member system cannot reliably run services or if the state of a member system cannot be determined, its services will automatically fail over to a functioning member system. Services already running on this system are not disrupted.
- **Service Relocation and Load Balancing** – DataGuard enables administrators to manually stop one or more services and relocate them to a functioning member system. This enables you to balance the service workload across the member systems or perform maintenance on a system while keeping services available. For example, if you need to upgrade hardware, you can relocate the services running on a member system, upgrade the hardware on that system, and relocate the services to the upgraded system.
- **Active-Active Configuration** – Leverage your hardware investment by setting up multiple separate instances of cluster services in active-active configurations. No need to have idle hardware standing by in a back up mode.

Configuration Support

The following hardware, software, and application infrastructures are supported by DataGuard cluster software.

Application Service Support

Cluster Service startup and shutdown scripts are provided for the following applications:

- Samba
- NFS Server V2 and V3
- IBM DB2 EE
- Oracle 8i and 9i
- PostgreSQL
- Informix
- Sybase
- MySQL
- Apache Web Server
- Zeus Web Server

Template Cluster Services are provided to permit creation of scripts to control any site-specific application.

Hardware Requirements

Minimum hardware requirements are as follows:

- Two Linux servers configured with one Ethernet interface, one RS-232 serial interface, and a local boot disk.
- One shared SCSI or Fibre Channel storage subsystem for cluster services, applications, and data.
- Optional: two programmable power controllers. Use of programmable power controllers is strongly recommended in environments that require stringent data integrity guarantees.

Computer System

DataGuard cluster software supports Intel® 32-bit (IA-32) hardware platforms. DataGuard includes initial support for IA64-based (eg. Itanium) processors. Please contact our Sales organization at Mission Critical Linux for more information. DataGuard is supported on the most popular Linux distributions. Refer to the distribution support list in the Software Requirements section of this document for more information.

Cluster nodes do not need to be identical. They can be of different speed, memory size, and SMP configuration, and can be supplied by different vendors. Selecting the appropriate configuration for each server depends on the applications and failover strategy to be deployed. The shared disk storage namespace must be identical for both cluster nodes, so it is recommended that the nodes have symmetric storage subsystems.

Additionally, it is recommended that each node in the cluster run the same Linux distribution and kernel version.

Storage Subsystem

The Linux boot device may be configured using any commodity hard disk, including IDE and SCSI. RAID support for the boot device may be provided by means of host software, adapter or controller based methods.

The shared storage subsystem must be carefully configured to meet the availability requirements of the overall DataGuard cluster configuration. The following guidelines should be observed:

- DataGuard cluster software supports SCSI and Fibre Channel for the shared storage subsystem.
- Production environment guidelines: For robustness, ease of maintenance, and correct handling of storage bus reset conditions, it is strongly recommended that the shared storage be configured using a multi-port storage controller. A multi-port controller provides a separate physical bus for each cluster node. The multi-port controller must support concurrent access to all logical units by all cluster nodes. These controllers are easier to configure and terminate correctly, and greatly simplify removal of a node from the configuration while the cluster remains operational. Because of the bus isolation provided by multi-port controllers, it is possible to use any modern, good quality host adapter in these configurations.
- Development environments guidelines: In development environments, it is often acceptable to use a lower cost, single-port storage controller. In these configurations, all cluster nodes as well as the storage controller are connected to the same storage bus. Note that depending on the host adapter and Linux device driver used, these configurations may not reliably recover from cluster node failures and reboots. Mission Critical Linux, Inc. is actively evaluating adapters and Linux device drivers that provide reliable recovery in these configurations.

  Please refer to the `www.missioncriticallinux.com/products/convolo` Web site for the latest adapter and device driver information.
- DataGuard cluster software does not use SCSI Reserve/Release commands or SCSI Target Mode, so support for these functions is not required.
- It is not possible to use host/server-based RAID software or adapter-based RAID firmware for the shared storage subsystem with DataGuard cluster software. RAID capabilities, if required, must be provided by a storage controller.

The following recommended SCSI RAID array products have multiple host channels and provide simultaneous access to all the logical units on the host channels:

- Winchester Systems FlashDisk RAID Disk Array, available from www.winsys.com
- CMD Technology CRD-7040, available from www.cmd.com (integrated RAID arrays based on the CRD-7040 are available from www.synetexinc.com)

Recommended Fibre Channel RAID controllers that have multiple host channels and provide simultaneous access to all the logical units on the host channels are the CMD Technology CRD-7220 and CRD-7240. Dual-redundant configurations are also available. Integrated RAID arrays based on the CMD CRD-7220 and CRD-7240 are available from www.synetexinc.com.

Network Hardware

DataGuard cluster software supports Local Area Network and serial line hardware supported by the specific Linux distribution on which DataGuard is installed. Refer to the network hardware support list associated with the Linux distribution for additional information.

Power Controller

Server systems in DataGuard cluster configurations can optionally be connected to programmable power controllers. The controllers are cross-coupled between the systems so that each system can control the power supplied to the other system. This feature permits DataGuard cluster software to ensure data integrity under certain rare failure scenarios, such as a system hang. It is strongly recommended that power controllers are included in production environments.

Development environments with less stringent data integrity requirements may be configured without programmable power controllers. For more information on configuring DataGuard cluster solutions without power controllers, refer to the DataGuard *Installation and Administration* manual.

DataGuard cluster software provides support for the following serial-based power switches:

- WTI RPS-10 (model M/HD in the US, and model M/EC in Europe), available from www.wti.com/rps-10.htm
- APC Serial On/Off Switch (part number AP9214), available from www.apc.com (Note that DataGuard cannot obtain status from this power switch.)

If you are using IBM xSeries servers, models x230 or higher, ending in 0 (zero), or IBM Netfinity servers, models 4500 and higher, you can set up the systems to utilize the built-in IBM Serial Management Console port for power cycling member systems. An external power switch is not needed in this configuration.

If you are using HP Netserver systems, models beginning with "L," you may be able to use the built-in HP Integrated Remote Assistant management port for power cycling member systems. An external power switch is not needed in this configuration.

Recommended network-based power switches include:

- BayTech RPC-3 or RPC-5, available from www.baytech.net (best performer of network-based switches on busy networks)
- WTI NPS-115, available from www.wti.com
- APC MasterSwitch (part number AP9211), available from www.apc.com

Software Requirements

DataGuard cluster software supports the following distributions:

- Red Hat Linux 6.2 or 7.1
- SuSE Linux 7.1
- Debian GNU/Linux 2.2r3

Note: When using other distributions not listed above, DataGuard cluster software requires either Linux Kernel 2.2.19 or 2.4.6. The DataGuard cluster software distribution CD contains a complete 2.2.19 and 2.4.6 kernel source pools complete with a corresponding set

of patches that are required to ensure correct cluster behavior. These patches include:

- NFS Server and utilities
- Raw I/O and bigmem support (2.2.19 only)
- ReiserFS file system support
- I/O fairness algorithm enhancement

Documentation

Comprehensive configuration, installation, and system administration information for DataGuard cluster software is provided in the DataGuard *Installation and Administration* manual. The hardcopy manual is supplied with the software product kit. The manual is also provided electronically in PDF format on the DataGuard cluster software distribution CD and on the `www.missioncriticallinux.com/products/convolo` Web site.

Distribution

DataGuard cluster software is provided as a boxed software product including the following items:

- DataGuard Product Description
- Distribution CD-ROM with two license numbers
- Hardcopy DataGuard *Installation and Administration* manual

There is no requirement to build DataGuard cluster software prior to installation. The software is provided as fully built images in RPM and DEB formats. However, it is necessary to rebuild the Linux kernel if you do not choose one provided in the product to include mandatory features (for example, support for Raw I/O operations and NFS failover).

Licensing

DataGuard cluster software is provided with two server licenses. *Both servers in the cluster must be licensed separately.* A valid license is required during software installation and to obtain Warranty services.

Warranty

DataGuard cluster software is provided with 90 days of warranty and support. Refer to the license agreement at the front of the DataGuard *Installation and Administration* manual for warranty details.

Support Services

Support for DataGuard cluster software is provided by Mission Critical Linux, Inc. Support Services are provided using our Web browser-based LifeLine customer relationship management system, email, and telephone. Service is provided during business hours (8:00 a.m. – 8:00 p.m. EST) Monday through Friday. Warranty services include assistance with DataGuard cluster configuration, installation and setup, and problem resolution. Critical problems receive a response within an hour.

To receive warranty services, please have your DataGuard cluster software license number available. Please call or send email to enable your LifeLine account.

Additionally, Mission Critical Linux, Inc. offers a range of support service contracts to meet any support requirement, including 24 x 7 coverage as well as long-term support for DataGuard cluster software and the complete Linux software environment. For additional information, please contact the Sales organization at Mission Critical Linux, Inc.

Training

Training courses are available for DataGuard cluster software and cover configuration, installation, management, and theory of operation. For additional information, please contact the Sales organization at Mission Critical Linux, Inc.

Ordering

DataGuard cluster software is orderable as follows:

- On the Internet at:
  `www.missioncriticallinux.com/products/convolo`
- Directly from the Mission Critical Linux, Inc. Sales organization.

DataGuard cluster software boxed package with:

Licenses	Part Number
2 Server Licenses	CCS-200-002

Contact Information

Phone: (+1) 978.606.0200

FAX: (+1) 978.446.9470

Toll Free: (+1) 877.62LINUX

Email:

`sales@missioncriticallinux.com`

`support@missioncriticallinux.com`

Internet:

`www.missioncriticallinux.com`