

Running TeleForm Enterprise in a High Availability Environment

Verity White Paper



Running TeleForm Enterprise in a High Availability Environment

August 2005

Table of Contents:

Objective	2
Background	2
Infrastructure Risk Areas	2
Servers	2
Workstations	2
TeleForm Support for High Availability	2
TeleForm Internal Database (MS SQL Server)	3
Typical Clustered Server	3
Linking TeleForm to the Cluster	3
Failover Process	3
TeleForm Internal File Storage (MS File Server)	3
TeleForm License Services	4
Summary	4
Definitions	4
Resources	4
High Availability Information	4
MS SQL Clusters	4
MS File Clusters	4
About Verity	5

Publisher's Note: Information contained in this document is intended for guideline purposes only. Verity product documentation supersedes information contained in this document. The situations described in this document are offered as examples; actual configurations and results will vary from system to system.

Running TeleForm Enterprise in a High Availability Environment

Objective

The purpose of this document is to provide an overview of what “High Availability” (HA) means, how these systems are built, and how Verity supports IT professionals in building TeleForm systems with high availability

Background

A common goal for all IT departments is to provide maximum availability of their systems to end users. The availability of a multi-user, network-based application is influenced by many factors that include multiple hardware and software components, all of which have the potential to fail. In the pursuit of high availability, all of these factors must be examined and addressed in an appropriate order based on a cost/benefit analysis for each component.

Below is a list of some key hardware components, examples of how they might fail, and some common responses to minimize the risks associated with component failure.

- Infrastructure Risk Areas
- Servers
- Workstations

While not comprehensive, the following tables give the reader an introduction to the types of things that must be addressed when building a high availability system.

Infrastructure Risk Areas

Risk Area	Potential Failures	Possible Responses
Electrical Power	Failure of externally-supplied power from utility companies.	Use of internal power generators for primary or backup power.
Cooling Systems	Water pipes in the A/C system can break, resulting in complete or partial failure of server room equipment.	Use of redundant and non-destructive cooling systems.
Fire	Electrical equipment can potentially catch fire and ignite other equipment.	Use of non-destructive fire suppression technologies.
Internal Network	Failure of routers, switches, network cables and other equipment essential for internal network communications.	Setup redundant systems that can compensate for failed components.
External Network	Failure of internal or external equipment required to provide access between the internal network and external networks, such as remote offices or the Internet.	Concurrent use of multiple vendors and different technologies to provide connectivity between internal and external networks.

Servers

Risk Area	Potential Failures	Possible Responses
Buildings	Natural and unnatural disasters such as earthquakes, hurricanes or terrorist attacks have the potential to disable entire buildings or cities.	Complete or partial duplication of mission critical operations at a secondary location.
Database Server	The database server application may stop working due to software bugs in the database application.	Use “clustered” servers so that if one server fails another server is able to take its place.
File Server Storage	One or more hard drives in the file server may crash.	Use of RAID-based file servers and/or clustered file servers.
Server Hardware	Any hardware component in a server computer, such as CPU, RAM, motherboard, NIC, etc., has the potential to fail.	Use “clustered” servers so that if one server fails another server is able to take its place. High-end servers also include redundancy for many key components, allowing them to continue operating even if one component fails.
Server-Based Application Services	Some network-based applications have additional hardware or software server-based components that are potential points of failure.	Generally speaking, the application vendor must build-in features specifically designed to address high availability concerns.

Workstations

Risk Area	Potential Failures	Possible Responses
Workstation Hardware	Similar to Server hardware, workstations can suffer from hardware failures.	Typically, redundancy is accomplished by obtaining duplicate hardware for any workstation components that are unique to a particular task.
Workstation Software	The workstation OS or a workstation-based application has the potential to stop working due to problems within the application or due to other applications running on the same computer.	In many cases, there are multiple workstations that are interchangeable. If there are special-purpose workstations that are unique, duplicate workstations should be available.
Additional Hardware	Some applications require specific hardware, such as high-performance scanners.	Typically, redundancy is accomplished by obtaining duplicate hardware for any workstation components that are unique to a particular task.

As seen in the tables above, there are many possible points of failure. While action can be taken to increase reliability in each area, there is always the chance of downtime if a failure occurs. Therefore the goal is to find the right balance between costs, data security and availability.

TeleForm Support for High Availability

The most basic level of support for high availability systems is for an application to be able to function reliably in an HA environment. This means that the product is certified and tested to

be compatible with system configurations typical to HA environments. However, compatibility does not mean the application leverages the HA features within the environment.

New features in TeleForm Enterprise Version 9.1 go beyond compatibility and allow TeleForm to be “aware” of clustered Microsoft SQL Servers and clustered Microsoft File Servers. This awareness allows TeleForm to leverage the clustered environment, resulting in increased uptime and reliability during potential failovers of the clustered servers.

Another functional area of TeleForm that supports high availability is the area of licensing. These changes are discussed in greater detail below:

- TeleForm Internal Database (MS SQL Server)
- TeleForm Internal File Storage (MS File Server)
- TeleForm License Services (tied to a specific computer)

TeleForm Internal Database (MS SQL Server)


TeleForm uses MS SQL Server for internal data storage. MS SQL Server can be configured to run in a clustered configuration where two identical servers are configured to run SQL Server. One of the servers becomes the primary server while the other is assigned to be the backup server in case the primary server fails.

See the following sections for more information on this topic:

- Typical Clustered Server
- Linking TeleForm to the Cluster
- Failover Process

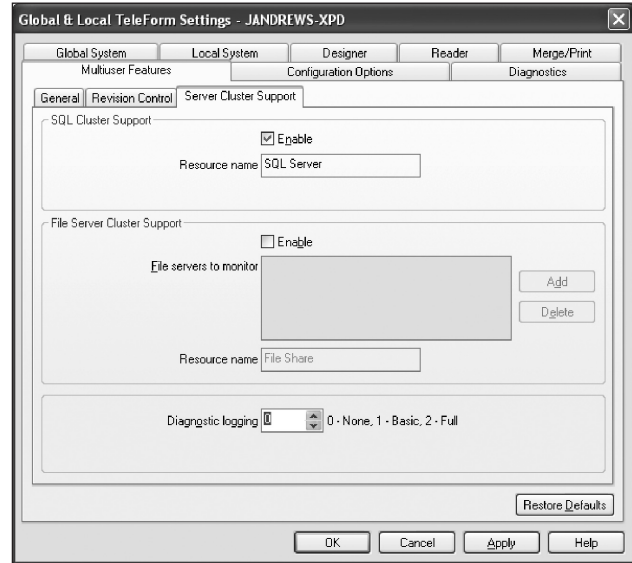
Typical Clustered Server

A conventional clustered SQL Server arrangement is represented in the following graphic:



Linking TeleForm to the Cluster

Once your clustered MS SQL Server is configured, you can make TeleForm aware of the clustered configuration by enabling the SQL Cluster Support in the Server Cluster Support tab of the Global & Local TeleForm Settings dialog box.



When you enable this option, TeleForm contacts the clustered SQL server and registers with the server so that the TeleForm applications will be notified if the clustered server fails.

Failover Process

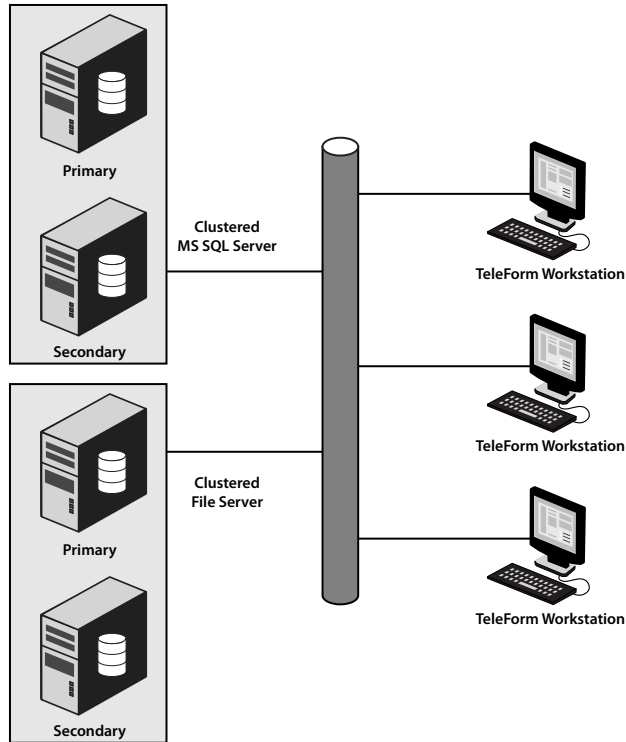
It is important to realize that the failover process, meaning the process of the clustered servers switching from the primary to the backup server, is not instantaneous and may actually take several seconds or even longer. It is also not transparent to the client applications because some in-process information, such as an active connection to the database, may become invalid during the failover.

As long as the primary SQL Server is running, TeleForm will run normally. However, if the primary SQL Server fails, TeleForm will receive a message from the backup server as soon as the failure is detected by the clustered server environment. When TeleForm receives this message it will immediately stop what it is doing and wait for another notification message that the backup server is on-line. When this second message is received, TeleForm will attempt to resume processing from where it stopped. If a SQL-related error occurs, TeleForm will attempt to recover by re-attempting the operation and in some cases re-initializing the SQL connection.

TeleForm Internal File Storage (MS File Server)

TeleForm supports clustered Microsoft File Servers in the same manner that it supports Microsoft SQL Servers. When you enable File Server Cluster support, TeleForm contacts the specified file servers and registers for notification when the file

server fails over to a backup server. A clustered file server and SQL server might be arranged in the following manner:



As with SQL Server clusters, it is important to realize that the failover process is not instantaneous and may actually take several seconds or even longer. It is also not transparent to the client applications because some in-process information, such as an open file handles, may become invalid during the failover.

As in the SQL Server failover situation, TeleForm will stop processing when it is notified of a File Server failover situation and will resume processing when it receives notification that the backup server is on-line. If errors related to the File Server are encountered, those operations will be reattempted and in some cases file handles will be re-opened.

TeleForm License Services

TeleForm uses activation-based licensing, where the product license is tied to a specific computer. This can be a potential point of failure if the computer used to activate the product becomes inoperative because it can no longer authenticate the license information.

Prior to TeleForm Enterprise Version 9.1, if the computer running the TeleForm License Service stopped working, the rest of the TeleForm modules would shut down within a few minutes or less. But with the changes made to TeleForm Enterprise in Version 9.1, this “grace period” has been extended up to two weeks after the last time the license was authenticated. This provides more than enough time for you to contact Verity for assistance in reactivating your product.

Summary

While building a high availability infrastructure may not be a simple task, it is easy to integrate TeleForm into a high availability environment that is using clustered Microsoft SQL Servers and clustered Microsoft File Servers. With the latest “cluster aware” features added to TeleForm Enterprise Version 9.1, the benefits of running TeleForm in a clustered environment have greatly increased beyond being simply “cluster compatible.”

Definitions

The following definitions are from the IEEE Task Force on Cluster Computing discussion on high availability:

Continuous Availability: This implies non-stop service, with no lapse in service. This represents an ideal state, and is generally used to indicate a high level of availability in which only a very small quantity of downtime is allowed. High availability does not imply continuous availability.

Fault Tolerance: This is a means to achieve very high levels of availability. A fault tolerant system has the ability to continue service despite a hardware or a software failure, and is characterized by redundancy in hardware, including CPU, memory, and I/O subsystems. High availability does not imply fault tolerance.

Single Point of Failure (SPOF): A hardware or software component whose loss results in the loss of service; such components are not backed up by redundant components.

Failover: When a component in an HA system fails resulting in a loss of service, the service is started by the HA system on another component in the system. This transfer of a service following a failure in the system is termed failover.

Resources

High Availability Information

- <http://www.ieeetfcc.org/high-availability.html>

MS SQL Clusters

- http://msdn.microsoft.com/library/default.asp?url=/library/en-us/architec/8_ar_cs_7ab7.asp
- <http://search.microsoft.com/search/results.aspx?st=b&na=88&View=en-US&qu=sql+server+cluster>

MS File Clusters

- <http://search.microsoft.com/search/results.aspx?view=en-us&st=b&na=82&qu=file+server+cluster>
- <http://www.microsoft.com/technet/prodtechnol/windowsserver2003/library/ServerHelp/36fbf586-99a4-4d57-b973-00c5a183a98a.mspx>

About Verity

Headquartered in Sunnyvale, California, Verity is recognized by industry analysts such as Gartner, IDC and Forrester as the market-leader in enterprise search software that enables organizations to discover, analyze and process all the information within their enterprises.

Verity's intelligent content services provide integrated search, classification, recommendation, monitoring, concept extraction and analytics across the real-time flow of enterprise information, along with question and answer interfaces for effective online self-service. When deployed in business critical applications, these capabilities let employees, customers and partners locate information, uncover and evaluate trends, and respond to questions quickly and accurately.

Verity content capture and business process management solutions deliver integrated paper and electronic forms capture and process automation. These solutions make the content on paper and electronic forms actionable, accelerate business processes from content capture to completion, connect people, processes and information across departments, enterprises and geographies, and enforce best practices and policies across processes.

Besides solving real business problems in the enterprise, Verity technology also serves as a core component of more than 260 applications from leading independent software vendors. Integrating Verity increases the value of the applications by both allowing them to display, extract information from and convert close to 300 file formats, and giving end-users the ability to search, classify and monitor content in the file formats they use every day.

Verity, Inc.

894 Ross Drive
Sunnyvale, CA 94089
t. 408.541.1500
f. 408.541.1600

Verity San Diego

3220 Executive Ridge Drive
Vista, CA 92081
t. 800.659.8755
f. 760.936.4800

Verity New York

230 W. 41st Street
12th Floor
New York, NY 10036
t. 646.366.9500
f. 646.366.9540

Verity Washington, D.C.

2941 Fairview Park Drive
8th Floor
Falls Church, VA 22042
t. 703.289.8800
f. 703.289.8840

Verity Chicago

200 South Wacker Drive
31st Floor
Chicago, IL 60606
t. 312.674.4525
f. 312.674.4571

Verity Boston

275 Grove Street, 4th Floor
Newton, MA 02466
t. 617.663.5716
f. 617.663.5382

Verity Canada

400 – 4th Avenue SW
Suite 2500
Calgary, AB T2P 0J4
Canada
t. 403.750.4000
f. 403.750.4100

Verity Australia

115 North Oatlands Road
Yarrambat, VIC 3091
Australia
t. 61.3.9436.2024
f. 61.3.9436.2078

Verity Benelux

Coltbaan 31
3439 NG Nieuwegein
The Netherlands
t. 31.30.669.2120
f. 31.30.662.2094

Verity China

18F, China Merchants Tower
No. 118 Jian Guo Road
Chao Yang District, Beijing 100022
P. R. China
t. +86-10-65675899
ext. 6091

Verity France

Les Collines de l'Arche
Immeuble Madeleine
D76 route de la Demi-Lune
92057 Paris La Défense Cedex
France
t. 33.1.41.49.0450
f. 33.1.40.89.0981

Verity Germany

Robert-Bosch-Strasse 9 „Europahaus“
D-64293 Darmstadt
Germany
t. 49.6151.86082.0
f. 49.6151.86082.20

Verity Italia

S.R.L.c/o Ernst & Young Via Cornaggia,
n°1020123 Milano
Italia
t. +39 050 770 157z
f. +39 050 770 157z

Verity Japan

15th floor, Cerulean Tower,
26-1 Sakuragaoka-cho, Shibuya-ku, Tokyo,
150-8512, Japan
t: 81.3.5456.5792
f: 81.3.5456.5733

Verity United Kingdom

Verity House
2 Heath Road
Weybridge KT13 8AP
United Kingdom
t. 44.1932.844200
f. 44.1932.843823

© 2005 Verity, Inc. All rights reserved. Verity®, the Verity logo, Ultraseek®, TOPIC®, KeyView®, TeleForm®, and Knowledge Organizer® are registered trademarks. Verity® Portal One™, Verity® LiquidOffice™, Verity® LiquidCapture™, Verity Profiler™ are trademarks of Verity, Inc., in the United States and numerous other countries. All other trademarks or symbols are those of their respective owners.

**Sales and Product Information**

software-sales@verity.com or 888.328.7335

Partnerships

partners@verity.com

Technical Support for Existing Customers

tech-support@verity.com

Verity, Inc.

894 Ross Drive, Sunnyvale, CA 94089 | t. 408.541.1500 | f. 408.541.1600 | www.verity.com