

IBM Security Secret Server
Version 10.4

High Availability and Disaster Recovery Guide

Contents

Overview	2
Suggestions	3
Architecture Planning.....	3
Virtualization	3
Frequent Backups	3
Uptime Monitoring.....	4
Secret Server HADR Features.....	5
Features by Edition	5
Secret Server Standard and Enterprise Editions	5
Secret Server High Availability Add On.....	5
Maintaining Secret Server in a “Disaster”	6
Simple Installation & Architecture	6
Restoring from Backup	6
HADR Solutions by Architecture	7
Basic Configurations.....	7
High Availability Configurations	9
Appendix	11
Automated and Manual Backups	11
Offline Caching and API	11
Cleartext Exports.....	12
Database HADR	12
Application HADR.....	12

Last updated: August 23, 2018

Overview

High availability and disaster recovery (HADR) is a primary concern of nearly every system administrator. Establishing a set of procedures to recover data from (or maintain operations during) failing hardware, natural disasters, or unforeseen circumstances is a requirement in most businesses and institutions. Secret Server can help System Administrators by storing and securing critical infrastructure data including admin-level credentials and important documents or files. Secret Server has an extensive list of features specifically designed to aid restore processes as well as methods for achieving business continuity during disaster scenarios.

This document is designed to assist system admins to integrate Secret Server with their disaster recovery plan. Electrical, connectivity, and similar concerns are out of the scope of this document. The assumptions made below are not intended to fit every situation, merely to serve as a guide for Secret Server administrators.

Suggestions

There are many factors to a solid disaster recovery plan. Physical hardware, network connectivity, power requirements, operating system configuration, and similar concerns are out of the scope of Secret Server. However, Secret Server can help with the management of and access to the Secrets saved in the application database. Available Disaster Recovery features in Secret Server depend on the edition of Secret Server. It is worth noting that Secret Server performs well in virtual environments, and integrating Secret Server into your virtualization solution is recommended.

ARCHITECTURE PLANNING

Consider adopting geographic redundancy when and where economically feasible. Multiple sites are not required by any means, but work to ensure network, power and hardware replication between Secret Server instances. Complex configurations for Secret Server can involve 3+ IIS Servers, Mirrored or AlwaysOn Microsoft SQL Server Clusters, and geographically-diverse work sites connected via WAN connections. The design of Secret Server lends itself to multi-site WAN-based use. However, when using Secret Server from multiple geographically-diverse sites, the issue of High Availability is typically one of the largest concerns.

See the [Secret Server Architecture and Sizing](#) guide for information about the most common configurations of Secret Server and Distributed Engine that can facilitate better preparation and planning for your highly-available and disaster-recovery ready environment.

VIRTUALIZATION

Secret Server fully supports being run on Windows in a virtualized environment. Using a virtual server may be preferable alternative when using the single server approach. Virtual servers allow the System Administrator to make a full backup of the server in addition to backups within Secret Server. They also allow for the server to be transferred to different hardware quickly.

If virtualization redundancy is not available or not sufficient in a multi-site environment, it may be beneficial to install Active Nodes on a continental or regional plan. These Nodes still communicate with the same database, but localized users will not suffer from WAN latency.

FREQUENT BACKUPS

When using a configuration without multiple servers, making frequent (daily) backups is the next best solution. All editions of Secret Server support creating backups manually and Secret Server Professional, and Premium Editions support scheduling these backups. In single server configurations, it is highly recommended to save the backup files for Secret Server on a different device. This will isolate sensitive files from a hardware failure on that server.

UPTIME MONITORING

Secret Server is a standard IIS web application and can be monitored for uptime like any other website. Consider using a web uptime management tool to verify that the front end website is still available. Using a load balancer to automatically redirect users between active nodes will increase uptime and can help prevent any outages for users.

Secret Server HADR Features

FEATURES BY EDITION

This is a complete list of HADR features built into the Secret Server application. Features are listed by their availability in the editions. Note that the Active Web Clustering is included for customers with Secret Server Enterprise Plus licenses, which are no longer sold.

Secret Server Standard and Enterprise Editions

- ✓ Web Services
- ✓ Automated and Manual Web Application Backup
- ✓ Automated and Manual Application Database Backup
- ✓ Offline Caching
- ✓ Microsoft SQL AlwaysOn
- ✓ Passive Web Clustering

Secret Server High Availability Add On

- ✓ Active Web Clustering

For details about the features and links to further resources for each, see “Appendix” on page 11.

Maintaining Secret Server in a “Disaster”

The framework of a solid Secret Server Disaster Recovery Plan should follow these methods of maintaining operations.

SIMPLE INSTALLATION & ARCHITECTURE

Secret Server can operate on typical modern server and even workstations in the simple configurations without requiring high-end hardware.

[Secret Server System Requirements](#)

By design, Secret Server’s installation is a quick and easy process. Keeping this process as quick and easy to install was a goal from the outset of Secret Server. This serves as a viable fallback option should redundancy plans fail. In a worst-case scenario where the host server fails, a cluster/mirror fails, and the other backup plans fail, Secret Server can be installed from scratch quickly and data imported from various methods. Users familiar with Microsoft SQL and IIS can typically install Secret Server in about 30-45 minutes on a prepared server.

See the [Installation Guide](https://ibm.biz/BdYBMS) (https://ibm.biz/BdYBMS).

RESTORING FROM BACKUP

Restoring a backup of Secret Server’s web application folder is as simple as copying the contents of the last available zipped backup file into place. Microsoft SQL database restores are simple as well, but require several steps, depending on the backup scenario.

Start by preparing a server for installation according to the Installation Guide. When the server is prepared, restore the application and database. This guide below will explain the procedure of restoring the database. Some specific web configurations may be needed to match the previous IIS settings.

See the following documentation for steps to restore from backup:

[Restoring Secret Server from a Backup](https://ibm.biz/BdYBgn) (https://ibm.biz/BdYBgn)

When restoring from backup in the single-server configurations, be certain to make copies of the Secret Server backup files on a different device or media. If you are unable to restore Secret Server after following these steps, please contact [Technical Support](#).

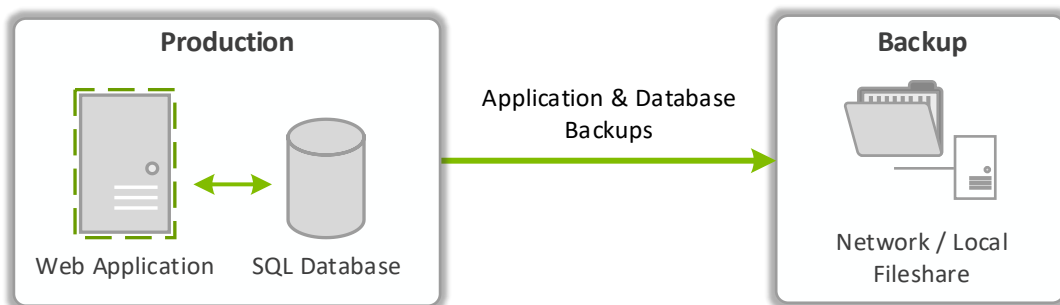
HADR Solutions by Architecture

The next four examples reference commonly-used architectures of Secret Server implementations and provide detail around how disaster recovery and high availability tie into each. For more information about planning Secret Server architecture, system requirements, and sizing, see the [Architecture and Sizing Guide](#).

Basic Configurations

Single Site, Single Server

Microsoft SQL Server Express, Secret Server Free, Professional Edition or higher

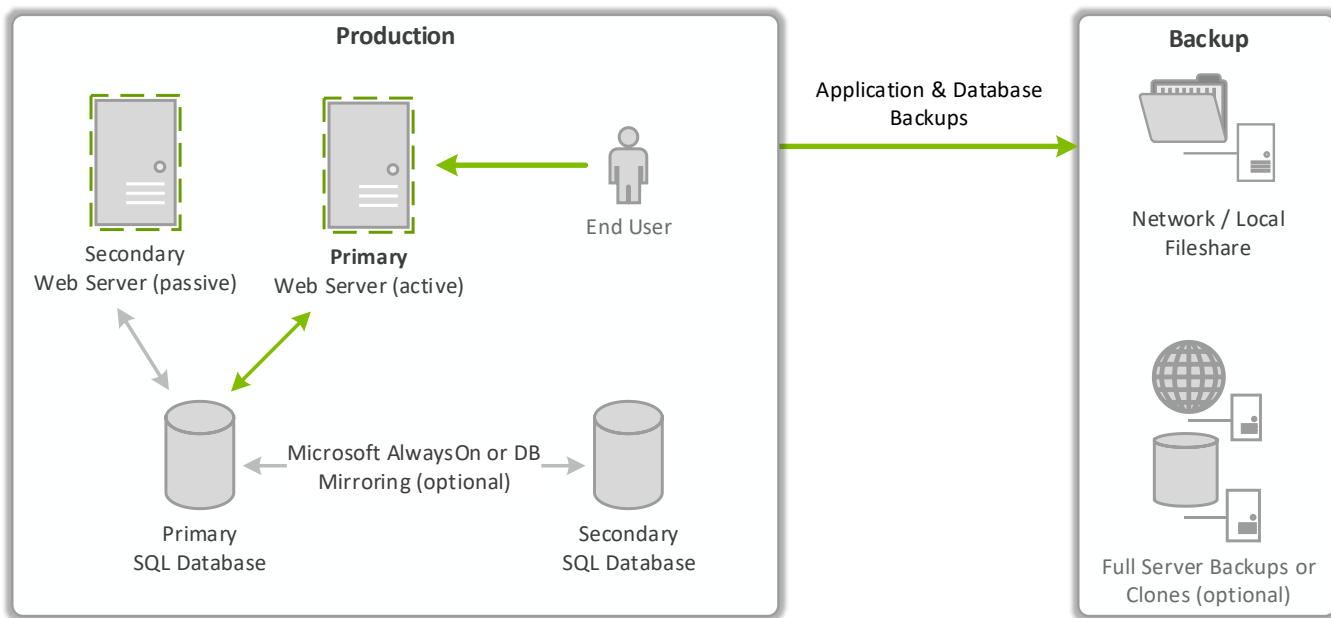


This is an example of a basic Secret Server installation. Secret Server and Microsoft SQL Server Express are installed on a single Windows Server.

A Disaster Recovery Plan for this configuration would consist of automatic or manual application and database backup procedures. Backups can be stored on another server or a fileshare. If the Windows Server is virtualized, leveraging strategies such as making scheduled snapshots or having a hot/cold site may add additional layers of redundancy.

Single Site, Active-Passive

Microsoft SQL Server Standard, Secret Server Professional Edition or higher



In this example, there are more than one front-end web servers, but only one active node. A web server node being active means that end users will hit this site when browsing to Secret Server at any given time. This server handles background processes (such as Remote Password Changing and Heartbeat) as well, making it the Primary server in addition to the fact that it is active.

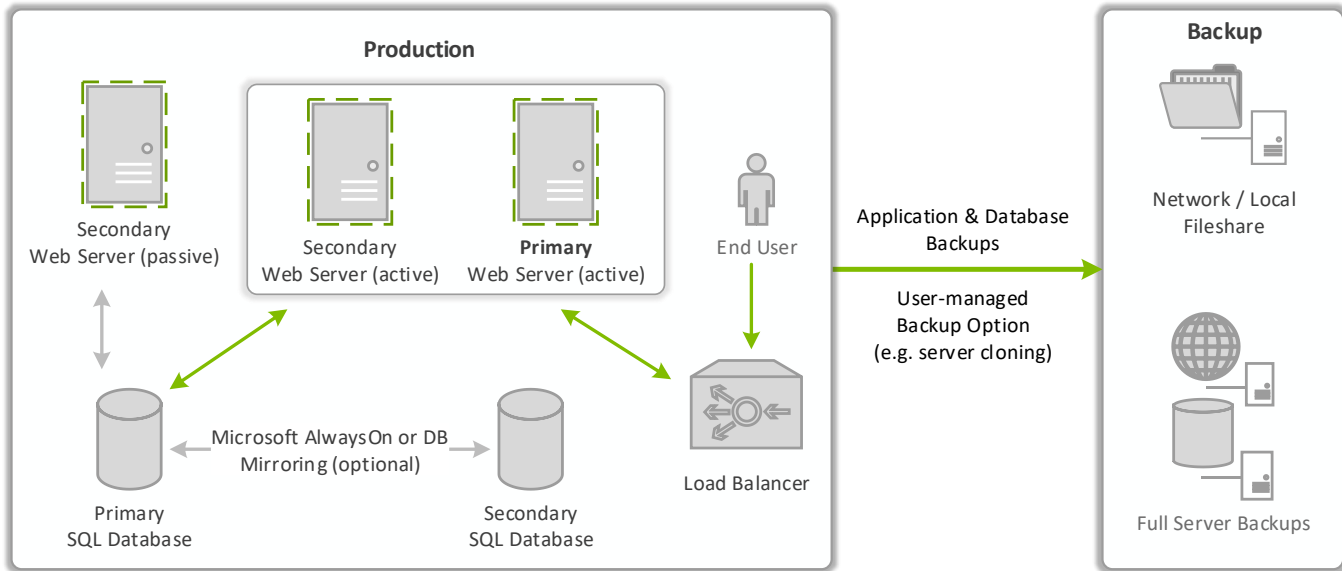
The passive web server also points to the backend SQL database, but informs users browsing to it that they must use the active node to use Secret Server. If the active node is unavailable, the “Primary” status can be transferred to the passive node, and users can resume using the application. There can be more than one passive server node (no limit), depending on the needs of the organization.

A Disaster Recovery Plan for above configuration would consist of failover for Web Server or Microsoft SQL Server issues. If the failover members were to themselves fail, then Automated Web Application Backups and Automated Application Database Backups can be used to restore functionality. If these Servers are virtualized, leveraging strategies such as making scheduled Snapshots or having a hot/cold Site may add additional layers of redundancy.

High Availability Configurations

Single Site, Active-Active

Microsoft SQL Server Standard or higher, Secret Server HA Add On

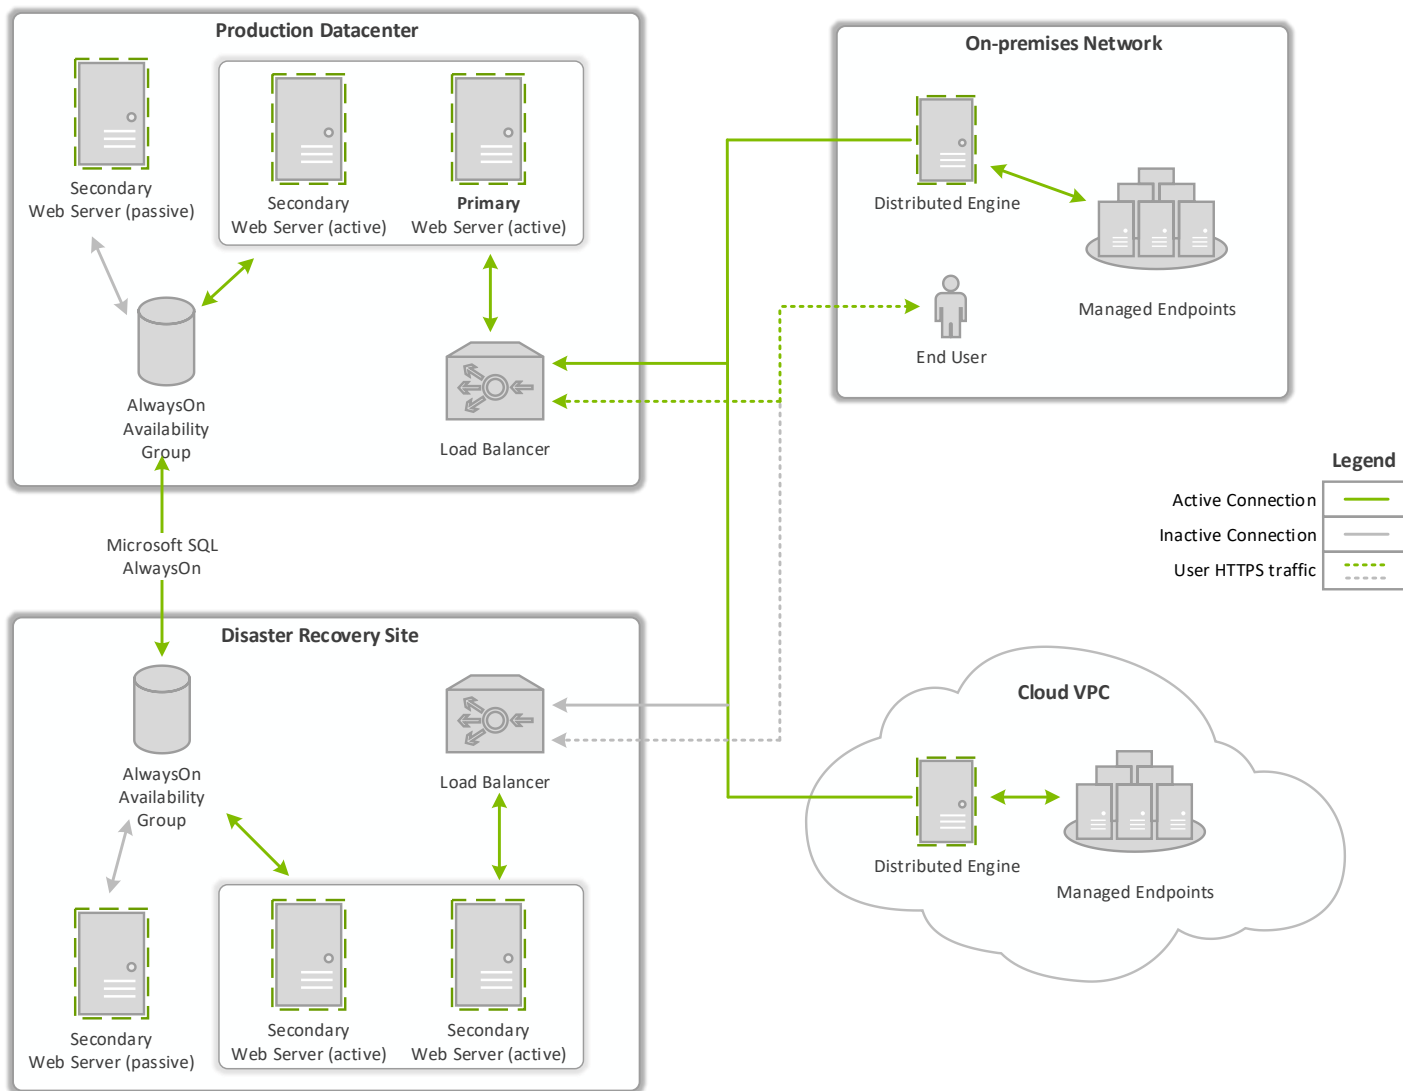


In this example, there are more than one front-end web servers, and more than one active node. Allowing users to use Secret Server through more than one active node simultaneously requires enabling clustering within the application. Only one server handles background processes (such as Remote Password Changing and Heartbeat), meaning that one of the active nodes will be designated as the Primary server at any given time (this can be changed manually, if necessary, in the application). If the Primary active node becomes unavailable, the "Primary" status will be transferred to one of the other active nodes and users can continue using the application without interruption. There can be more than one active and passive server nodes (no limit), depending on the needs of the organization.

A Disaster Recovery Plan for above configuration would consist of failover for Web Server or Microsoft SQL Server issues. If the failover members were to themselves fail, then Automated Web Application Backups and Automated Application Database Backups can be used to restore functionality. If these Servers are virtualized, leveraging strategies such as making scheduled Snapshots or having a hot/cold Site may add additional layers of redundancy.

Multi-Site, Active-Active

Microsoft SQL Server Enterprise, Secret Server HA Add On



This is an example of an implementation that would be used by a larger enterprise. Not only are there more than one front-end web servers, more than one active node, and SQL AlwaysOn in use – there is a setup of this configuration at each of two sites. Users can browse to the application through one load balancer at the production site, and if anything happens on the production site that cannot be handled by the Active-Active setup at that location, administrators can have users directed to the load balancer at the second site and an active server node there will become Primary. Still, only one server handles background processes, but it's recommended to offload Remote Password Changing, Heartbeat, Discovery, etc. to Distributed Engines. This means that regardless of which server is Primary, Distributed Engines will be able to retrieve workload tasks from Secret Server and connect to managed endpoints per usual.

Appendix

AUTOMATED AND MANUAL BACKUPS

Secret Server natively supports local and network backups. By configuring locations for the application folder and Microsoft SQL database, Secret Server backs up this data based on a highly-configurable user-defined schedule with detailed logging.

Restoring just a database backup and reconnecting the web application to the database typically requires less than 30 minutes (depending on the size of the database being restored).

Restoring just a web application backup in conjunction with a valid Secret Server database, functionality can typically restore functionality in less than 30 minutes.

Automatic Backups

This is a configurable “one-click” process that creates backups of the Secret Server Application and Database. Please refer to the following documentation to configure automated backups:

[Backing up Secret Server to a Network Share](#)

[Backup Configuration File Path Settings](#)

Manual Backups

This process involves manually making a copy of the Secret Server application directory from the web server and backing up the Secret Server SQL database from SQL Server Management Studio. Please refer to the following documentation for Manual Backup procedures:

[Manually Backup Secret Server](#)

OFFLINE CACHING AND API

Offline Caching (Mobile App)

Secret Server allows users to connect using applications (iOS, Android, Windows desktop Chrome app). These applications allow users to cache a copy of their available Secrets on their desktop or mobile device. This cache is protected and has a customizable time limit. For more information about the Thycotic PAM mobile app, see the [Thycotic PAM Guide](#).

Web Services API

Secret Server provides web services to allow for custom development (SOAP v1.2) to interact with Secret Server while maintaining security. Customers can develop software for interacting with Secret Server or to script specific tasks. See the [Web Services API Guide](#) for details and the [Thycotic Knowledge Base](#) for examples.

CLEARTEXT EXPORTS

Secret Server supports two formats for exporting data into clear text for printing or electronic storage.

DATABASE HADR

Secret Server supports pointing to an AlwaysOn listener to connect to the database. In this setup, AlwaysOn handles HADR of the database. Secret Server also supports synchronous and asynchronous SQL database mirroring in Microsoft SQL Server, however this setup is becoming deprecated, and is therefore not recommended. Adding geographic redundancy to this plan is recommended for customers with multiple sites, for that added layer of protection. Log File Shipping is not supported in Secret Server.

Please refer to the following documentation for setup:

<https://msdn.microsoft.com/en-us/library/ff878265.aspx>

APPLICATION HADR

Active and Passive Web Server Clustering

Secret Server allows users to install multiple Web Applications for Disaster Recovery purposes. In the case of an active-passive configuration, users can manually switch between nodes to designate a single active server that will be writing to the database.

With the High Availability add on licensing, Secret Server supports High Availability web server clustering (active-active configuration). Customers can configure Secret Server to be highly-available and geographically-redundant quickly. Users can rapidly deploy Secret Server active instances by first cloning pre-configured web servers, then copying the website “application” folder, and lastly running the database connection procedure. Background threads are still only written by the “Primary” node. This feature is useful for users who have geographically diverse environment with multiple Secret Server instances.

See the following documentation for detailed steps and potential issues with Active Web Clustering-enabled web servers:

[Setting up Clustering in Secret Server](#)