

Protecting Exchange 2010

Enhancing the Availability of Microsoft Exchange Server 2010

Native backup can provide higher levels of data and system protection, but the need for 3rd-party software and hardware is still apparent. This document describes protection for Exchange Server 2010 and its component roles and illustrates how a combination of native solutions and Double-Take Software offerings can provide a more complete solution for both high availability (HA) and disaster recovery (DR).

Exchange 2010 Roles and Structures

The Exchange platform is broken up into distinct roles which can be installed on one or more servers in an environment. Most enterprise-level implementations will have several servers for each role, using a combination of technologies to provide redundancy and availability. There are five key roles :

Mailbox (MBX) - Required: Responsible for storage of all email data for all mailboxes and public folders within the organization, including shared contacts, user and utility calendars and tasks and files. Databases that contain this information are organized into Information Stores, which are part of Storage Groups (SG's). SG's are logical relational databases built on the Extensible Storage Engine (ESE) database system. While not tied to specific MBX servers, each SG must be hosted on one or more MBX servers in the Organization in order to function.

Hub/Transport (HT) – Required: HT servers function as the primary internal gateway for all Exchange communication between servers. They may also function as the bridgehead servers to external email systems. HT servers also function as the main point of message hygiene (anti-virus, anti-spam, etc) for all internal communications and may serve that purpose for inbound and outbound external communications as well.

Client Access Services (CAS) – Required: CAS servers provide a connection point for MAPI and other forms of client-side communication software. In some cases the

requests are ferried to HT or MBX servers, in others the CAS server will act as the single communication point for the client. Outlook Web Access (OWA) and Outlook Anywhere (OA – formerly ActiveSync) are two examples of technologies enabled by CAS systems.

Edge Transport (ET) – Optional: ET servers are external gateway hosts for Exchange systems. They lie outside both the firewall and the Domain, and contain only a limited subset of Domain information to confirm that a user exists, but nothing else. ET servers route external mail that has cleared first-stage message hygiene onward to HT servers for continued processing. They also act as SMTP communication systems to external mail servers, moving mail into and out of the enterprise.

Unified Messaging (UM) – Optional: UM servers function as a bridgehead between the Exchange system and the VOIP or PBX systems used by the enterprise for telecommunication. This allows Exchange to handle storage and routing of voicemail as well as email, and enables Outlook Anywhere to be accessed via telephone service by converting text email to voice information and accepting voice and keypad input.

Additional Servers in an Exchange Environment

Along with these roles, Exchange servers often leverage other technologies as part of an overall messaging and collaboration solution. Some examples of major forms of external systems are:

Mobile Messaging – Example: Blackberry Enterprise Server (BES): Though automatically enabled with Windows Mobile interface technologies through Outlook Anywhere, many enterprises choose to support additional mobile messaging platforms. BES is an example of a highly-popular, non-Outlook-Anywhere enabled mobile solution platform. The BES server communicates with the Exchange server via MAPI technology (the main communication system for Outlook clients) and shuttles messages and other data to and from handheld devices running on many mobile networks.

Protecting Exchange 2010

Extended Collaboration Tools – Example: Microsoft SharePoint Services (SPS): While Exchange has continued to offer basic support for collaboration through public folders in Exchange Server, that technology as a whole has been depreciated by Microsoft and is subject to be discontinued in some future Exchange version. Instead, Microsoft recommends more robust and flexible collaboration solutions such as SPS for sharing documents, project information and other team-related or Intranet information.

Utility Servers – Example: Microsoft SQL Server (SQL): Though the Exchange Server system is self-contained, in most cases there will be auxiliary servers and services that depend on databases or other resources that exist in the environment. Both SPS and BES require SQL (or SQL Express) instances in order to function, and at the enterprise-level these databases are typically held on their own server systems. UM servers also leverage external SQL databases for storage and management purposes, creating a need for these utility servers even in cases where no other applications outside the Exchange 2010 solution set are used.

Legacy Systems in an Exchange 2010 Environment

While the primary focus of effort in an Exchange Server 2010 environment will be the Exchange 2010 servers themselves, many environments will not be able to immediately decommission earlier versions of the platform. During migrations, there will be at least a limited period of co-existence with previous versions of Exchange Server, and in most cases there will be an extended time before the organization can completely remove the legacy systems. During the co-existence periods, legacy Exchange servers still contain critical data for the organization, and therefore must provide some level of redundancy, availability and recovery.

Resilience: The total solution for Exchange 2010

Two Tiers for Uncompromising Protection

Tier One: Native Solutions

With Exchange Server 2010, Microsoft has leveraged existing Windows Server technologies and combined them with new solutions native to Exchange 2010. A cooperative solution set that uses both sets of technologies allows you to provide basic availability for the core components of an Exchange Server 2010 system.

Database Availability Groups (DAG) are your first line of defense for mailbox databases and MBX Role servers within the environment. DAG allows each Storage Group to be replicated using log-shipping technologies to up to 15 other MBX servers in the same Domain. This effectively creates up to 16 total copies of each database, any of which can be made active. Only one copy of each SG may be active at any one time, but servers acting as passive nodes for one SG may be the active node for other SG's in the organization.

Configuration of DAG solutions requires multiple MBX servers be installed in one or more physical locations. For Exchange 2010 Server Standard Edition, you may not have more than five SG's within the domain, and therefore you can protect no more than two SG's via DAG without exceeding the total SG limit in this version of Exchange. With Exchange 2010 Server Enterprise Edition, the SG limit is raised to 100 total (active or passive) SG's per server, allowing for much more flexibility for DAG.

Once the MBX servers are installed and configured, and the SG's each server will host as primary are established, a series of PowerShell commands and GUI interactions will establish DAG replication between all nodes specified for each SG. DAG may be configured to replicate SG information to one, more than one, or all MBX nodes in the Domain, at the Administrator's discretion. If a DAG node or DAG protected SG fails, another node within the DAG system will automatically take over, redirecting all Outlook 2007 or higher clients and re-establishing mail flow. CAS servers will re-direct most other SMTP, POP3 and IMAP clients as well as Outlook Web Access and Outlook Anywhere connections. Legacy Outlook versions may need to be either manually re-homed or wait for DNS updates to occur to allow for Profile re-direction. The failover process may also be executed on-command.

Protecting Exchange 2010

Network Load Balancing (NLB) allows the organization to host multiple CAS servers within a single site. This technology is not configured automatically, but is available for the Exchange Server 2010 system, as long as the underlying version of Windows Server is capable of supporting NLB. The technology allows for incoming communications for CAS servers to be routed to more than one physical server by using a virtualized machine identity. Typically this is done by addressing all connections to a virtual IP address managed by Active Directory and associated DNS entries.

HT servers are natively redundant, with each HT server able to accept communications from all Exchange servers within the same Active Directory Site they are installed to. No NLB is required for this Role within Exchange 2010.

Across multiple physical and/or Active Directory Sites, HT and CAS servers are independently managed. As long as each Active Directory Site has at least one HT and one CAS Role installed (possibly both Roles on the same server), then Exchange communications for each site will be correctly handled in the event of a loss of the primary site. Note that this refers to internal communication only, and would require that DNS systems correctly re-direct incoming SMTP communications to the appropriate HT server if the organization does not leverage the Edge Transport Role.

The combination of these native Windows and Exchange technologies can provide for basic database, server and site resilience. However, this is often not sufficient to meet enterprise-class organizations' needs with regards to uptime, compliance and Disaster Recovery Planning (DRP) requirements.

Tier Two: Double-Take Software Solutions for Extended Availability

There may be many instances where DAG systems are not the most effective solution, especially in single-server (combined role) environments. Some environments may not have the required networking infrastructure to properly support a DAG system between locations, while others may not have the appropriate knowledge-base within their staff to manage a DAG solution. In these instances, Double-Take Availability offers a failover solution for single-server environments. As Availability replicates not only data, but system state information and system/application binaries; a secondary single-server can be configured at a remote site to provide all-in-one failover.

As discussed above, the native solutions for Exchange 2010 can provide a starting point for total messaging and collaboration protection in more complex configurations, but do not complete the picture. Adding Double-Take Software products can provide availability for all components of an enterprise-class platform, and these products are designed to be used along-side the native tools.

Edge Transport servers cannot leverage NLB or other hardware or software load balancing solutions due to the problems these would cause with aspects of the ET Role functions (such as blacklist/whitelist push). Each Active Directory Site can independently control ET communication, and round-robin DNS configurations can offer some level of protection against single-server failure. This will, however, cause performance issues if one or more ET servers go completely offline within a single site using round-robin DNS routing.

Double-Take Availability allows another physical or virtual machine to assume the role of the failed ET server within minutes. This minimizes the amount of time that the organization must work at limited capacity due to a single-server failure, and eliminates the time and effort required to rebuild the server from the ground up during an emergency. This effort should not be underestimated.

While the Exchange Server components are relatively easy to reinstall, remember that ET servers are responsible for first-line message hygiene. This means that you will also have to manually install and re-configure all Microsoft and 3rd-Party tools that perform anti-virus, anti-spam, blacklist/whitelist processing, etc. Availability ensures that all system-state components and application binaries are mirrored to another server, removing the need to manually configure these systems even if the hardware is completely different.

Protecting Exchange 2010

Unified Messaging servers also suffer from an inability to leverage most load balancing systems, as UM servers are tied to a specific VOIP or PBX system and therefore must be uniquely configured to interface with those devices. Double-Take Availability brings another UM server up and online in minutes, provided another server (physical or virtual) has the required interface hardware already integrated. The server hardware may be completely different, but the standby server must be capable of communication with the VOIP or PBX system in question. As the entire voicemail and possibly the call-routing system for the enterprise relies on this server Role, the time required to manually restore and configure a failed UM server could cripple the organization for an extended period.

Availability allows a different server to take over for the failed system quickly, and give the IT staff time to properly repair the failed hardware without interrupting communications to do so. Protection of UM servers across multiple sites is typically not a useful solution as the VOIP or PBX systems are generally bound to one physical location. However, for enterprises that have created stretched telecommunications systems via VOIP bridges and/or contracts with traditional telco companies, a UM server could be failed over between sites by leveraging FSFO technologies within Double-Take Availability.

Mobile Messaging platforms like BES often include some rudimentary form of failover solutions to provide basic protection and cutover in the event of a loss of a single server. However, these solutions cannot provide robust multi-server and multi-site failover platforms. Most native tools do not permit for multi-path failover, and do not easily permit multi-subnet failover. Technology sets within Double-Take Availability can provide for both eventualities; allowing you to have multiple target devices which can take over in the event of either single-server or multi-server (site-wide) failures. Having a unified tool set that can handle multiple eventualities will simplify management and enhance overall availability.

In the case of Extended Collaboration Tools like SharePoint Services, these often come with the ability to produce redundant systems within the same physical location. The loss of a single SPS web server (where at least one other SPS web server is available) will allow for continued access to SharePoint data by end-users.

The problem comes into play when a site-wide failure occurs. In these instances, Double-Take Availability can move roles, responsibilities and end user connectivity to another site and to different server platforms, physical or virtual. This allows business to resume quickly without manual reinstallation and reconfiguration of an entire SharePoint Site – a process that is typically rated in “days to completion” and not the hours or minutes required by most enterprise organizations.

Utility Servers, such as SQL back-end systems for BES and SPS, are also vital links to enterprise data and would require resiliency. SQL does have native tools that can allow another SQL server to store a copy of database information, and even perform limited failover operations. These tools, however, require that they be configured, maintained and operated during failover by an experienced SQL DBA, which may not be an available option during an emergency. Other Utility Servers may not have any native tools at all, leaving them unable to provide resiliency even with experienced personnel available.

Since the systems that these databases and other servers enable cannot function if the Utility Server itself is lost, using a tool set such as Double-Take Availability will allow these systems to resume their functions within minutes either within the same site or to an entirely different location. Hardware for these target devices does not have to match, so long as the target hardware (physical or virtual) is capable of running the applications in question.

Exchange 2003 and 2007 Servers that are running as legacy systems in an Exchange 2010 organization will also need to be considered for availability, as they will contain critical data within SG's homed on these servers until such time as the migration is complete. Until that time, leveraging Double-Take Availability tools for these Exchange platforms will allow you to provide local and remote failover for all Exchange servers within the enterprise while you remain in co-existence periods. Double-Take Software has whitepapers and information on how to protect both of these legacy platforms available from the <http://www.doubletake.com> website.

Protecting Exchange 2010

Summary

While Microsoft has made great strides in resiliency for the Exchange Server platform with the introduction of the native tools in Exchange 2010, these tools are still limited to their intended purposes. This leaves gaps in both High Availability planning and Disaster Recovery operations. Leveraging both native technologies and the Double-Take Software product family can allow enterprise organizations to ensure that all critical services and solutions will survive in the event of single-server fault or site-wide emergencies. It also ensures that recovery of failed servers back to their original or new hardware will be able to move forward without concern for hardware incompatibility or manually intensive application by application re-installation.



QUESTIONS?

1-888-674-9495

info@doubletake.com

www.doubletake.com

© Double-Take Software, Inc. All rights reserved. Double-Take, Balance Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.