

Backup Strategies for Microsoft Exchange Server 2003

Microsoft® Exchange Server 2003 installations require a solid backup and recovery architecture. This article presents common topologies and operational strategies for backing up Exchange in a variety of enterprise scenarios, including stand-alone servers, LANs, and storage area networks.

BY SUMAN KUMAR SINGH AND QUOC DAT NGUYEN

Related Categories:

Backup

Dell PowerEdge servers

Dell PowerVault storage

Dell/EMC storage

Disaster recovery

Microsoft Exchange

Visit www.dell.com/powersolutions for the complete category index.

Enterprises typically translate e-mail downtime into lost revenue, lost productivity, or both. The size of an organization and the number of applications and service-level agreements that the data center must support are essential considerations when determining a suitable e-mail backup architecture. Enterprises must assess cost/benefits trade-offs that include the impact of the backup infrastructure on the availability and performance of business-critical applications—defining an acceptable duration for the backup and restore window, for example.

To protect e-mail data from potential disaster, the first line of defense is usually to back up critical information using tape or disk. To help administrators determine the backup method that is most appropriate for their specific enterprise requirements, this article describes three models for backing up the e-mail infrastructure: stand-alone server backup, LAN-based backup, and storage area network (SAN)-based backup.

Stand-alone server backup model

A stand-alone backup and restore scenario can be appropriate for small Microsoft Exchange environments that are hosted on a single stand-alone server. This approach typically locates the Exchange database on the server's internal storage or on direct attach SCSI or Fibre Channel storage. Even if the Exchange database is fairly large, the high storage capacities of advanced tape technologies may allow an organization to back up data onto a single tape. For example, the tape backup unit in the stand-alone server backup model can be a dedicated tape drive, such as the Dell™ PowerVault™ 110T Ultrium 3 Linear Tape-Open (LTO-3) tape drive, or an autoloader tape backup library, such as the Dell PowerVault 132T tape library (see Figure 1).

The advantages of this model are that it is simple, easy to deploy and implement, and cost-effective. However, there are a few limitations. For example, the management

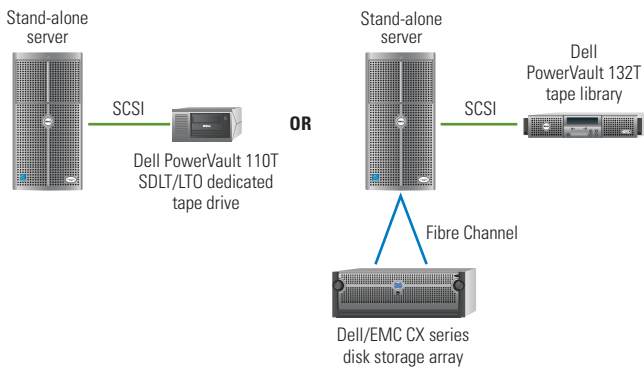


Figure 1. Stand-alone server backup model

of a stand-alone server backup model can be difficult if the data center is running multiple applications that must be backed up separately. In addition, this backup model offers limited scalability compared to LAN-based and SAN-based approaches. Because this configuration does not use a separate backup server, the application server’s system resources must be shared and dedicated among different applications, including the backup task. As a result, backups may affect the performance of production applications.

LAN-based backup model

In the LAN-based scenario, the tape library attaches to a separate server known as the backup server or media server. The backup server connects to the application servers over the LAN, as shown in Figure 2. The LAN-based model may be suitable for environments that support multiple servers running multiple applications—for example, Oracle® database, Microsoft SQL Server™, and

Compared to the stand-alone server backup model, the LAN-based model offers enhanced scalability to meet future needs for expanded storage and backup capacity.

Microsoft Exchange applications. Configured either on stand-alone servers or on clustered servers, the applications can share the same backup tape library across the LAN. The master backup server initiates backup tasks and provides a centralized location housing the catalog database as well as the logical/physical management tree for the entire backup organization.

The backup application agent must be installed on all application client nodes to facilitate application backups via the backup server. The backup server has a physical SCSI- or Fibre Channel-based interface to the tape library. The data from the Exchange server or other applications is sent over the LAN to the backup server and then to the tape library. The data flow is shown in Figure 2.

LAN-based backup has several advantages over the stand-alone server backup model. In this model, different application servers can share a single tape library over the LAN. This centralized approach helps simplify backup administration. Compared to the stand-alone server backup model, the LAN-based model offers enhanced scalability to meet future needs for expanded storage and backup capacity. However, the LAN approach can also incur significant performance penalties because backup is performed over the network. To help avoid contention between the application traffic and backups, best practices recommend configuring a separate, isolated subnet dedicated to backup traffic. This limitation can be addressed in a SAN-based backup model.

SAN-based backup model

The SAN-based scenario is similar to the LAN-based approach—it also streamlines administration through centralized backups. The network topology for SAN-based backups is designed to improve application performance because it is routed over a high-speed Fibre Channel network. As a result, SAN-based backups enable the following benefits:

- Enhanced performance of transaction-intense applications using a high-bandwidth network interface
- High reliability and availability
- Ability of heterogeneous servers and operating systems to coexist and share the same tape library across the network

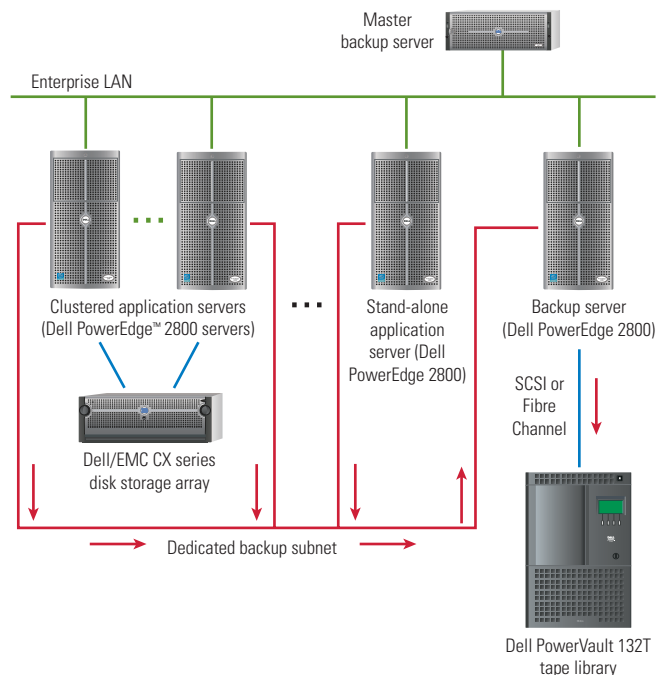


Figure 2. LAN-based backup model

As shown in Figure 3, application servers, a master backup server, a storage system, and a tape library are connected across the Fibre Channel fabric.

The SAN approach has the potential to drastically reduce backup and restore windows compared to LAN and stand-alone server models—thereby helping to improve performance for application service-level agreements.

As in the LAN-based model, data traffic can be routed through a high-speed Fibre Channel switch and written directly to the tape library. As in the LAN-based model, the master backup server controls the backup tasks.

At press time, the front-end Fibre Channel interface for Dell-based SANs is designed to support data transfer rates of up to 4 Gbps (up to 800 MB/sec full duplex). Moreover, the latest multi-drive autoloader tape libraries or modular tape libraries—such as the Dell PowerVault 136T and PowerVault ML6000 series, respectively—enable administrators to back up several applications concurrently.

Of the three backup topologies described in this article, the SAN-based model is designed to provide the highest I/O throughput for backup. Consequently, the SAN approach has the potential to drastically reduce backup and restore windows compared to LAN and stand-alone server models—thereby helping to improve performance for application service-level agreements.

In addition to these benefits, the SAN-based model can take advantage of storage software and backup techniques such as EMC® SnapView™ snapshot software, EMC MirrorView™ SAN-based mirroring, and the Microsoft Windows®-based Volume Shadow Copy Service (VSS). The SAN-based approach also enables cluster-aware backups, in which a backup job can be failed over to another available node in the event of a failure during backup operations.

Backup strategies

For all three scenarios described in this article, administrators must ensure that backup software running on the backup infrastructure is compatible with Microsoft Exchange Server 2003. To take advantage of online Exchange backup, backup software must support the Exchange Server 2003 Backup and Restore application programming interface (API) or the Windows VSS writer.

It is equally important to back up all the data required to restore applications running on an organization’s server to a previous known good state. Along with the applications, support software and management scripts must be backed up. For Exchange Server 2003, backing up the contents of mailboxes, public folders,

and requisite configuration data for the Exchange environment is critical. In addition, best practices recommend that Exchange data be backed up separately—not together with Windows or with the full server backup operation.

Administrators should ensure that backups include the following:

- Microsoft Windows OS
- Backup and systems management software
- Management scripts
- Microsoft Active Directory® data
- System state, including the Microsoft Internet Information Services (IIS) metabase
- Cluster quorum (if Exchange uses clusters)
- Certification services (if applicable)
- Exchange databases and log files
- Exchange message-tracking logs

After hardware and software components are configured properly in the backup infrastructure and critical backup data is identified, administrators must implement a backup strategy. Exchange works with one or a combination of the following methods: full backup, differential backup, incremental backup, and mirror backup.

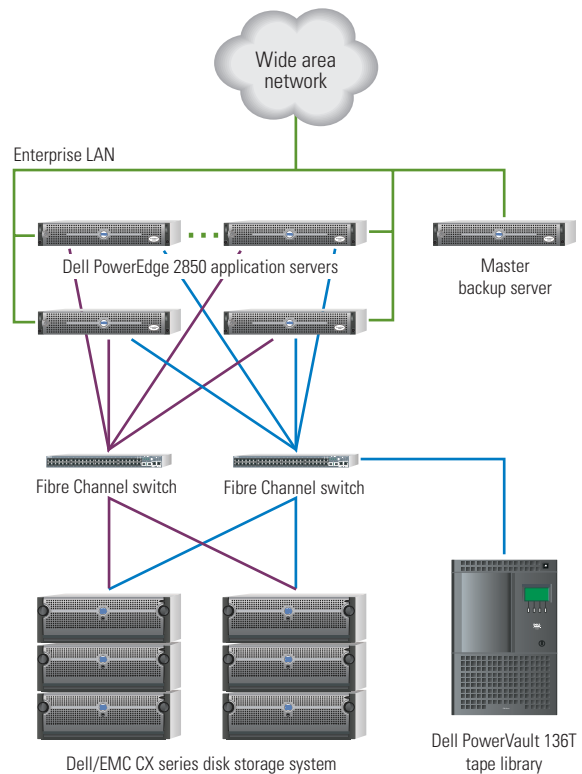


Figure 3. SAN-based backup model

Full backup. A full backup is designed to store all data, including Exchange database files and transaction logs. This approach helps simplify the recovery process because it saves all the data files and transaction log files in a single backup session. However, a full backup operation consumes the most network bandwidth and requires the most storage space compared to differential, incremental, and mirror backups. For that reason, best practices recommend a full backup operation be performed at regular intervals, in rotation with other backup strategies.

Differential backup. A differential backup contains only the Exchange transaction log files that have changed since the last full backup; the database files are not copied. Because all the transaction logs since the last full backup are required for a restore operation, circular logging cannot be enabled during a differential backup. Recovery requires both the last full backup and the last differential backup. Best practices recommend that

Key considerations when determining the e-mail infrastructure include the time necessary to perform the backup, the number of tape backup sets required for the restore, the time necessary to complete the restore, and the system resources available to perform the restore.

a full backup be performed at regular intervals and supplemented with daily differential backups.

Incremental backup. An incremental backup contains the Exchange transaction log files that have changed since the last full, differential, or incremental backup. Of these three types, incremental is the fastest backup method and may be suitable for large Exchange databases with a high volume of daily activity. The drawback to the incremental approach is that recovery requires the last full backup and all subsequent incremental backups. Best practices recommend that a full backup be performed at regular intervals and supplemented with daily incremental backups.

Mirror backup. A mirror backup is similar to a full backup except that no file marking is performed. Mirror backup is not ordinarily used for recovery purposes. This method can be used to make a full copy of the Exchange database without disrupting any incremental or differential backup procedures.

Microsoft Exchange Server 2003 supports online backup, which can be performed using a program such as Microsoft Windows NT® Backup (Ntbackup) or a third-party backup utility that supports the Exchange backup API. During the online backup process, Exchange services typically continue to run

normally and users experience no downtime. Online backups can be performed for full, differential, incremental, and mirror backup strategies.

Practical considerations

For enterprises of all sizes, e-mail has become a mission-critical application. When determining a suitable backup architecture for the e-mail infrastructure, administrators must factor in the number and variety of applications and service-level agreements that the data center must support. Of particular concern is the impact of the e-mail infrastructure on the availability and performance of business-critical applications. Key considerations

when determining the e-mail infrastructure include the time necessary to perform the backup, the number of tape backup sets required for the restore, the time necessary to complete the restore, and the system resources available to perform the restore. Last but not least, administrators must ensure that backup media is stored in a secure location.

The three strategies described in this article for backing up Microsoft Exchange Server 2003—stand-alone server backup, LAN-based backup, and SAN-based backup—offer general guidelines for Exchange backup operations. Actual implementations may differ based on the specific requirements of individual Exchange organizations. ☞

Suman Kumar Singh is a systems engineer in the High-Availability Systems Group at Dell. He specializes in messaging systems architecture and sizing. His other interests include SANs, virtualization, and security. Suman has published several papers and presented at enterprise computing-related industry conferences.

QuocDat Nguyen is a systems engineer in the High-Availability Cluster Development Group at Dell. His responsibilities include developing SAN-based, high-availability clustering products that comprise Dell servers and Dell/EMC Fibre Channel storage systems. QuocDat has a B.S. in Electrical Engineering from the University of Houston.

FOR MORE INFORMATION

Microsoft Exchange Server 2003 upgrades on Dell platforms:

www.dell.com/exchange

Dell storage:

www.dell.com/storage