

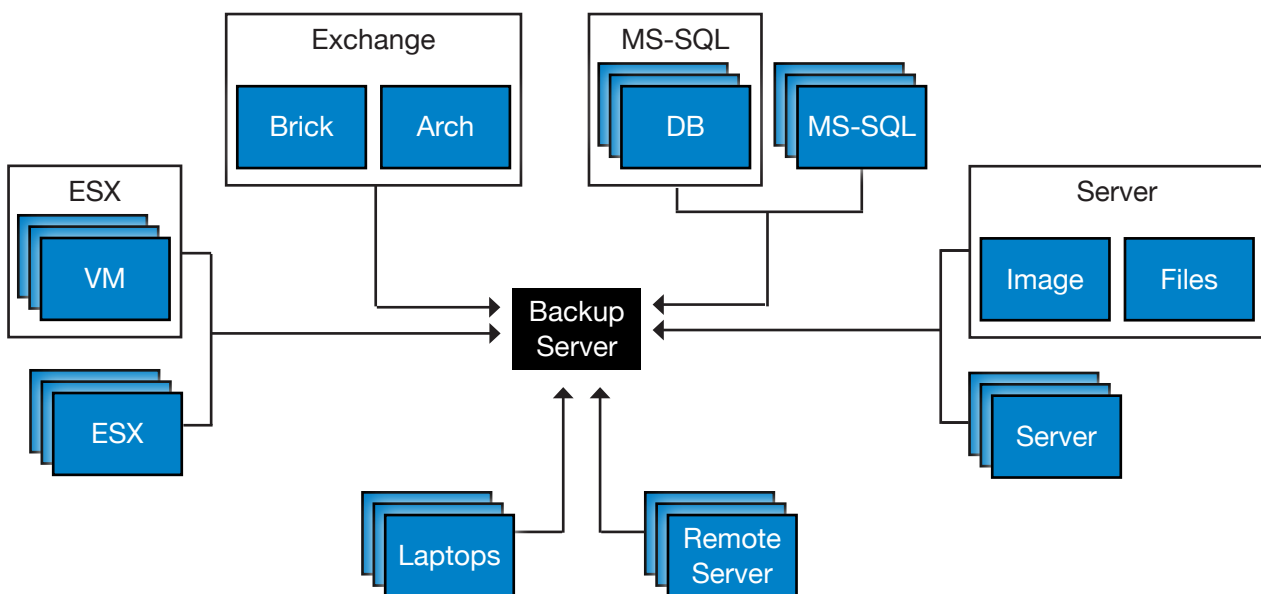
This document describes the functionality of a backup server and how multiple servers can be networked to create a backup infrastructure. The emphasis is on technical and design issues for organisations with large data volumes spread across multiple sites.

THE BACKUP SERVER

The basic unit used to build a backup infrastructure is the Backup Server. This stand alone server hosts all the necessary software. It resides on a particular site and will be configured to carry out all backups on that site. It provides scheduling, reporting and has a web based interface. It will connect to the organisation's servers to carry out the following types of backup:

- Windows file based backup (including locked files)
- Imaging of the Windows system disk in VHD / VMDK format (also called bare metal)
- Exchange "brick level" backup of message stores
- Exchange email level archiving
- Databases (Oracle, MS-SQL, MySQL)
- Virtual Machines (ESX, Hyper-V, Xen)
- Other OS file based backups, including Linux, Unix, Apple

The Backup Server can also be the receiver for laptops not on the company network, these laptops require the installation of an agent. Remote servers on a WAN or VPN can also be backed up, providing there is sufficient bandwidth. After capturing data, the system creates a version history of data using point-in-time snapshots.



DISASTER RECOVERY

To protect the data against a full scale disaster, data should be replicated to a second geographically distinct location. The Backup Server can be configured to replicate data to one or more secondary Backup Servers within the organisation.

A Backup Server, while being the primary for its site, could function as the secondary for other sites. This approach requires less hardware but has the drawback of providing less redundancy.

TECHNIQUES FOR EFFICIENTLY STORING AND REPLICATING DATA

The main issues when implementing a backup solution for larger multi-site organisations are storage and bandwidth. The system provides the following functionality that can help with this:

- De-duplicated version histories – Although point-in-time snapshots are presented as complete file trees, only the changed files are actually stored.
- Block level disk de-duplication – Where two files are part identical, the identical data is only held once.
- Bit differenced replication (also called delta blocking) – Only the changes within files are replicated, not the whole file, and these “deltas” are compressed in transit.
- Throttling – restricting the amount of bandwidth replication uses. By running replication at full speed during the night but throttling during the day, user impact is limited.

DESIGN

The overall system design needs to reflect the data volumes, connection speeds between sites and the organisation’s disaster recovery strategy.

Careful planning is required to achieve the correct design, know the following:

- Total number of servers and data at each site.
 - List of servers, operating systems, applications, backups required.
 - The criticality of the various data.
 - A schedule for the backups.
- Connection speed between sites.
 - When and for how long the connections can be utilised for backups.
- How the system will be used in the event of various DR scenarios.


Once this information is known it is relatively easy to size the required hardware and design a system.

EXAMPLE DESIGN – FIVE SITES

The following diagram is for a company with five sites. London, Bristol and Paris have a significant number of servers and data, requiring an on-site Backup Server. Brussels and Lyon have their servers backed up to Paris over a WAN. All the data on Paris is replicated nightly to London, the London data is replicated to Bristol, Bristol is replicated to Paris.

Every design is a trade off between cost and redundancy. The implications of this architecture are as follows:

- Cost effective - the entire organisation is backed up using just 3 servers.
- Limited redundancy – a site disaster at Bristol would require data to be downloaded from Paris, or the Paris server would need to be couriered to Bristol, temporarily leaving Paris without a primary.
- Restores from Brussels and Lyon are dependent on the speed of the WAN.



PARTITIONING SCHEME FOR A LARGER ORGANISATION

The following diagram is for a company with a large number of sites. Most sites are branch sites, with one or two computers, regional offices have a significant number of servers and data. The Backup Servers are located in the regional offices, branch sites are backed up remotely. Each Backup Server replicates to another regional office.

This design scales. Small sites with few servers are backed up remotely to a designated larger site which has a Backup Server. Each Backup Server replicates to a partner. The entire organisation can be partitioned in this way. There is no centralised backup, instead it is distributed throughout the organisation.

