FORTRA

Maximizing Uptime
For Mission-Critical File Transfers

How To Choose The Most Effective Solution
For Your File Transfer Needs

Managed File Transfer (MFT) solutions are essential to a wide range of business processes. Secure file transfer solutions support business-critical operations, such as supply chain management or logistics, to provide a highly available environment. These file transfer processes are often part of a service provided to customers and partners, with "uptime" guaranteed in negotiated service-level agreements.

When choosing an MFT solution, you must first decide whether you want to deploy the solution as onsite software or use cloud services, and whether you want to manage the solution yourself or outsource that service.

If you have decided to deploy and manage your own MFT solution, the next task is identifying availability and scalability requirements and choosing the appropriate deployment architecture. This paper provides guidance on common high availability and scale-out deployment architectures, and discusses the factors to consider for your specific business environment.

Three Deployment Architectures for Availability and Scalability

At a high level, there are three basic models for deploying the MFT solution on-site: as a single server, in a two-node active/passive failover cluster, or in a horizontal scale-out model with load balancing among multiple servers. This section describes the basic attributes of each option depicted below.

1. Single Server, Stand-Alone Configuration

The simplest deployment model has the MFT solution running on a single server, but does not provide any clustering or failover capabilities.

PROS

This solution has the lowest initial deployment cost. It requires the least amount of hardware, power, networking, and licensing resources, and is relatively simple to install and manage.

CONS

You will need maintenance windows to keep the core server up-to-date (OS patches, software updates, etc.). An unplanned outage due to hardware or system failure may cause significant downtime while you either repair or replace the hardware or move the application to another system.

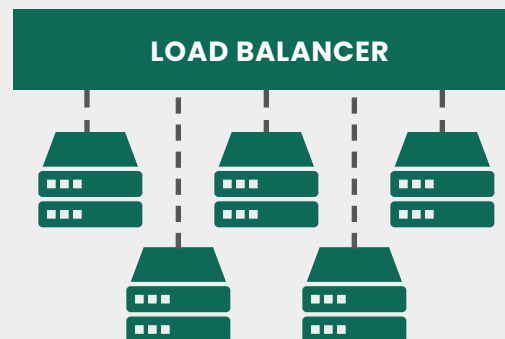## Deployment Configuration Options

Single Server   Active/Passive HA cluster   Active/Active HA Deployment

LOAD BALANCER

2. Active / Passive HA Cluster

You can improve application availability by using an active/ passive failover configuration with a clustering service such as Microsoft Windows Server failover clustering or Symantec Veritas Cluster Server. Each server accesses shared storage, and the cluster server can be set up to automatically failover during maintenance or if a problem arises. Only one MFT server at a time is active.

PROS

In case of an unplanned outage, automated failover to the passive node limits downtime. Maintenance windows are minimized, as you can do a planned "failover" for maintenance purposes. This model imposes moderate additional costs for hardware, power, networking, and OS licenses; the MFT vendor should discount the extra license for the passive node. If you use industry-standard clustering technologies, you do not need specialized expertise to deploy or manage the solution.

CONS

During a hardware failure, failover is not instantaneous, and active sessions will be terminated and need to restart. This model adds complexity to the initial installation and ongoing maintenance—with two or more servers to maintain instead of one, and clustering software in addition to the MFT software.

3. Active / Active HA Cluster

A scale-out configuration uses multiple MFT servers accessing shared storage and a load balancer to distribute requests among them. This architecture offers similar scalability and availability features as a web server farm. It works well for high-volume, around-the-clock, and globally-dispersed MFT environments with mission critical file transfer operations.

When the load balancer detects that a server is not responding, it stops sending requests to that server. The current connections to that server will be lost, but all other nodes remain operational. There is no failover window to accommodate. And if you need more capacity for file transfer operations, you can simply add another MFT server to the architecture.

PROS

This solution offers superior scalability, as you can easily add more processing power when needed. It also offers the best uptime during unplanned outages; when one node fails, alternate nodes are ready with no failover window. It is well-suited to mission-critical, high-volume file transfer operations.

CONS

A scale-out deployment imposes the highest TCO, with higher costs for initial deployment and more servers to manage. Creating and maintaining the load-balanced environment requires network infrastructure and administrator expertise.

Mapping Deployment Architecture to Business Considerations

Each of the deployment architectures described above offers a different balance of cost, complexity, availability, and scalability. Before choosing an option, you should carefully assess your business requirements to determine the relative importance of the different factors.

Availability and Downtime

The importance of downtime depends on the specific business environment in which managed file transfer is operating:

- How much downtime can you tolerate? 5-10 minutes? A two-node failover cluster may work well for you.

- Do you need to offer "dial-tone" reliability to customers? A scale-out cluster may be necessary.

- Can you negotiate a maintenance window, and if so, how long can it be? Global operations often have no good time windows for maintenance.

- Do you offer service-level agreements to customers? If so, what are the costs of downtime exceeding negotiated levels?

You need to balance the costs of SLA violations with the additional configuration and maintenance costs of the different availability models.

Scalability

What kind of growth in file transfer traffic do you expect? Can you handle it by scaling up (more powerful hardware)? If so, a failover cluster or stand-alone server may work.

Do you have extremely high or unpredictable volumes of file transfers? Many MFT solutions are highly efficient; one adequately provisioned server can handle a lot of traffic. However, if you have an extremely high volume of file transfers, the scale-out configuration is the better choice.

Costs

Your selection should balance costs against availability and scalability. As you consider costs, account for the total cost of ownership of the solution, including deployment, management, and maintenance costs.

Against the total costs, factor the potential costs of planned and unplanned downtime. For organizations with high volume or mission-critical MFT processes, the cost of downtime easily outweighs the investment in availability architectures.

Summary

MFT solutions play a critical role in many businesses. No matter how reliable and robust the solution itself, every MFT deployment is subject to outages or problems due to failures outside the application, such as OS problems, hardware failures, and networking failures.

You can protect your business from downtime by choosing a deployment architecture that limits downtime through clustering strategies. A two-node, active/passive cluster is relatively simple to deploy, while a scale-out, load-balanced cluster offers the significant benefit of nearly limitless scalability using many smaller (less expensive) servers.

Each of these models is tested and proven in thousands of installations around the globe.

To choose the right deployment architecture for your MFT solution, you should understand the potential cost of downtime, and balance that against the cost of deploying and managing additional hardware and software for availability.

Please contact our sales team if you would like to see a demo or have specific questions you would like answered.

For more information please go to:
www.globalscape.com

Fortra.com