



A trading name of Vitanium systems Limited

VSL Business Continuity Strategy

Author: Paul Houselander (Technical Director)

Date: 01/03/2010

Description: Published

Contents

Purpose	2
Overview of Disaster Recovery Infrastructure	2
Scope	2
Overview of Disaster Recovery Methodology	3
Infrastructure	3
Disaster Recovery Invocation	4
Points of Contact and Escalation Paths	4

Purpose

The aim of this document is to provide an overview of the VSL technical architecture to deliver Business Continuity. Business Continuity is provided by eliminating single points of failure from the physical infrastructure, by utilising industry leading hardware and diverse networks to provide high availability, backups and geographically separated mirrors of data and infrastructure.

Overview of Disaster Recovery Infrastructure

The VSL business continuity solution was designed so that there are no single points of failure. To achieve this two diversely situated `high availability` server clusters was implemented. The server clusters are split between two physical sites, and have multiple power supplies, multiple network cards, and mirrored internal hard drives. The storage for both the clusters is in a RAID 5 configuration with two hot spare disks. The arrays are asynchronously mirrored at the block level between the sites.

There are multiple links between sites, taking different routes and with different presentation points to each building. Networks within sites are fully resilient, utilising multiple links between servers, switching and storage. Internet connectivity is provided via six (6) 100mbit/sec (burst able to 1Gb) Internet connections, supplied by different ISPs. The networks are protected with multiple firewalls.

Data centres in both sites are physically secured, and have fire detection, fire prevention, flood detection, redundant air conditioning, and UPS with diesel generator backup. Servers and services are monitored by VSL staff on a 24x7 basis. For a service failure, there is a response time of 15 minutes during core hours, or two hours at other times.

Scope

Business continuity plans will provide protection against single points of failure. Multiple concurrent failures may not be covered. For example, losing multiple storage arrays at different sites or losing both sites simultaneously is not covered.

Losing an individual piece of hardware, such as a server, and losing an entire site will be covered by the DR plans. In addition, recovery from database corruption will also be provided, by utilising mirrored and timed backups to recover the database to any point in time.

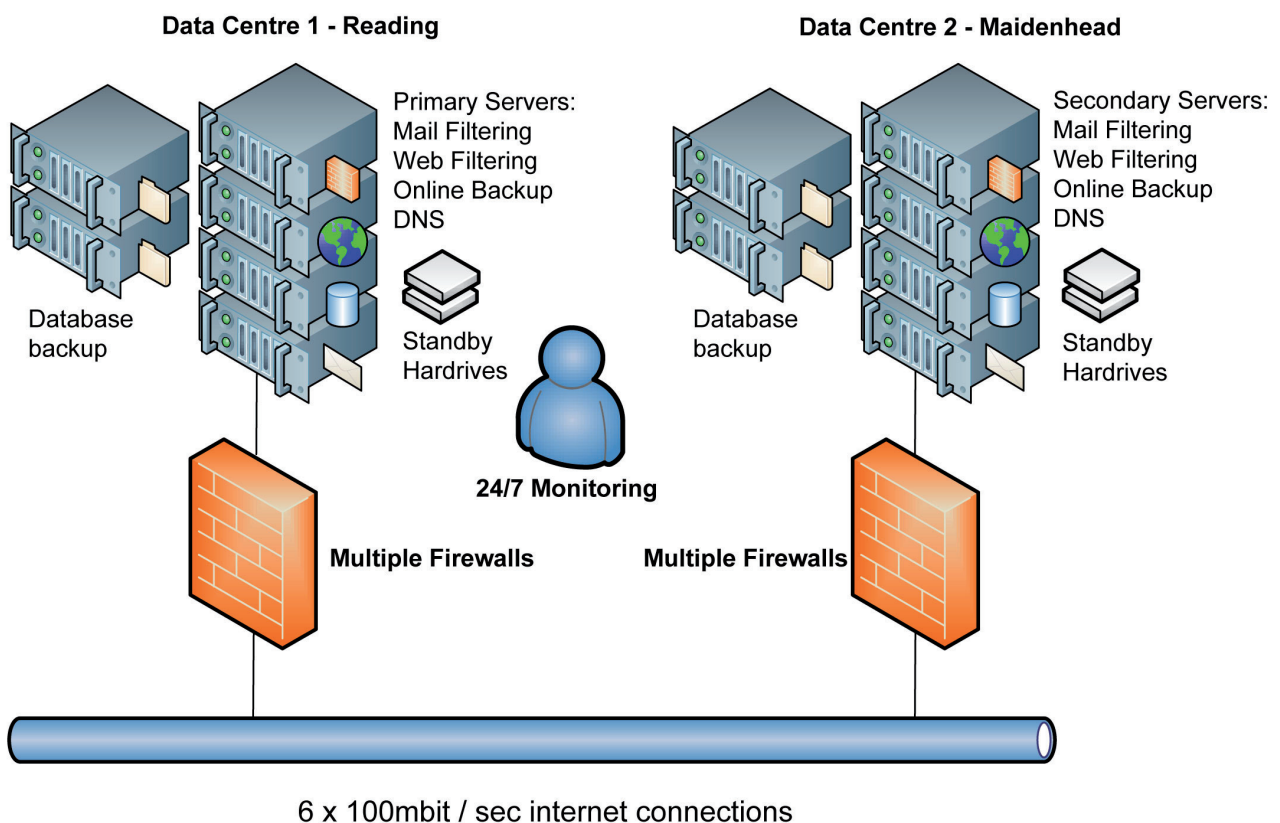
In addition, business continuity is provided during maintenance and also for predictive failure of hardware. These are predicted by the use of management software and detailed monitoring. The two server clusters are designed to also provide load balancing.

Overview of Disaster Recovery Methodology

High availability will require a restart of services on an alternate server. This form of failover requires manual intervention by a system administrator to check database consistency, restart services, and unit test the system before the service can be accessed by customers. In the event of loss of the live storage or of the live site, recovery will utilise the mirrored file systems. The mirrored disks are inaccessible, and need to be promoted to a read/write state and presented to the surviving server hosts. Once this is complete, the servers and services can be started. It is important to note that the mirroring is asynchronous; the remote mirror can be up to 10 minutes behind the live system.

This process does not guarantee that the database (DB) is in a consistent state. Nightly backups to additional hardware will be used to provide this guarantee. If the DB has been corrupted during the failure and will not restart from the mirrored disk, then the database files will be restored from the last backup.

The infrastructure is depicted below:



Disaster Recovery Invocation

Every member of VSL has a copy of the VSL Disaster Recovery Procedure document. The initial trigger to invoke DR is to contact the VSL technical department about the possibility of a disaster. The information is immediately escalated to the on duty Emergency co-ordinator. The Emergency co-ordinator reports to the monitoring technicians and the decision team company directors, who take responsibility for deciding to invoke the DR procedure. The invocation target timeframe is 24 hours from confirmation of a business impacting disaster and the VSL Management Team agreement to invoke DR platforms and processes. Contact details and methods are documented and maintained in the VSL Recovery Procedure document.

Points of Contact and Escalation Paths

Once DR invocation has begun, VSL directors will become the point of contact for customers. The Technical implementation teams will present progress updates to the Emergency Co-ordinator on an hourly basis, or at intervals agreed during the invocation. Information will be fed out to VSL Directors via the Emergency Co-ordinator.

