

# Storage Solutions

---

## Planning the fall of Sprawl



## Addressing Microsoft SQL Database Sprawl

Consolidation. It continues to be the most prolific data centre infrastructure project for all the usual reasons. Ease of management, reduce capital spend, reduce operational costs, running out of data centre space and Going Green.

And we are consolidating everything. Whether we are consolidating our server estate by introducing server virtualisation technology from the likes of VMware, or application consolidation, by introducing a single vendor's application suite to replace multiple applications from a variety of vendors.

Consolidation is what IT departments are doing, because they have to.

One of the *overlooked* areas (and I'll explain the italics later) is database consolidation. Most organisations have grown an IT estate with many database servers in it, from various vendors, at various versions with various levels of support and availability (and are maintained at varying degrees of strictness!) We use different types of database vendor providers for different applications. It might be that the application vendor has dictated the database required, or it might be that it is believed that only one vendor's offering is suitable for the type of application, or, unbelievable as it may sometimes seem, you may have a strategy around your database platform of choice and that is the database platform that is most common in your environment.

Most large IT departments have seen a rapid increase in the use of Microsoft SQL Server. The abundance of SQL Server and the cost-effective servers on which it runs has shown significant benefits. But the success of SQL server has also created issues for IT administrators who are responsible for MS SQL deployment and management.

MS SQL Server does provide a high degree of functionality and performance at a low cost with it running on commodity server hardware providing unmatched price performance. But what has resulted in this is **MS SQL Server Sprawl** which has introduced the problems being observed by organisations today.

Here are the few main issues of SQL Sprawl:

#### Low Utilisation and Poor Technology Efficiency.

Historically, a one database to one server approach to database deployment has been adopted in many organisations, perhaps due to departmental investment or to simplify capacity planning where servers are provisioned with capacity for peaks now and in the future. This approach has produced very low levels of server utilisation. In Computacenter's experience working with our customers, it is not uncommon to find most database servers running at 5-15% CPU utilisation, particularly with modern server hardware. Of course, low levels of utilisation imply high capital costs or that there is much capacity in the data centre that is underutilised but unavailable to utilise elsewhere.

#### Deficient Server Availability

The reality in many data centres today is that few MS SQL Server databases are protected by some form of high availability. Typically, this comes from an understanding that high availability options are costly and complex. Therefore only those deemed 'critical' are protected with some form of high availability, but more and more, just about all production databases (and associated applications) are viewed as business critical.

#### Management Overhead

The higher number of server instances the higher the management costs, that much is obvious. The cost of acquiring server hardware and associated software actually represent a small part of the Total Cost of Ownership (TCO) of servers. A larger proportion of the TCO is made of the cost of managing and maintaining the server.

With operating system patches and upgrades server in the organisations environment are an end point that requires multiple discovery and execution of various tasks when carrying out upgrades and patching.

So why are these MS SQL databases *overlooked* for consolidation? (I said I would explain the *italics*) Because IT administrators don't believe there is the solution to address it. Perhaps using server hardware virtualisation technology is seen as inappropriate for high performing databases? Perhaps vendor support is the issue? Some clustering technologies are considered too complex and expensive to implement for all the organisations databases.

But Computacenter has been implementing a solution that addresses these issues and it stands head and shoulders above anything else available. The

deployment of an **HP PolyServe** based solution can offer today's organisation significant advantages and cost benefits over other clustering solutions.

The HP PolyServe Database Utility for SQL Server is based on PolyServe's Matrix Server software, a shared data clustering architecture that enables all servers to see all data within the SAN. This solution architecture is quite unique in the Wintel space.

The fundamental idea of the PolyServe solution is to treat the hardware resources (servers and storage) that support the implementations of MS SQL Server databases as a single pool that can be directed, as required, toward any database.

The key point here is that the files containing the data for all the databases are stored in a single place that can be accessed by all of the servers simultaneously.

With PolyServe's cluster file system architecture, each node in the cluster mounts the shared storage on the SAN. This means that during a failover, there are no delays whilst the new active node in the 'farm' mounts the physical disk resource.

The major advantage to this approach is that cluster failovers don't have to wait for disk mounts to be completed. This can result in a failover being accomplished in seconds as opposed to several minutes.

Thus, the adoption of a Polyserve design offers the following fundamental benefits:

**A single pool of host servers:** Using HP Polyserve, we no longer think of installing a database on any particular host. Instead, we install into the cluster, and the database can then run on any host server in the cluster. This allows a DBA or system administrator to move a database from one server to another in order to adjust server load and maintain an optimum balance of utilisation and performance (without any need to copy or migrate data).

**Flexibility:** The shared data cluster can be formed from Windows servers in the existing environment. There is no requirement or guideline to buy matching servers, these can be assorted servers from different vendors, with different processor types and speeds, different CPU numbers and different memory.

**Simple scalability:** It is easy to add servers to a cluster if overall utilisation demand grows, and databases can be relocated in a matter of seconds onto newly added servers with no need for data copying. Equally, if workloads drop, you can relocate databases off some of the servers and remove them for repurposing. Databases receive complete native performance with no virtualisation overhead, so if a database needs the full capacity of a larger server, it can have it.

**Uncomplicated High Availability:** Because all servers have access to all databases, high availability becomes simple to implement with no requirement for replicating hardware. Simply by indicating where a database should be restarted in the event of failure, from among any of the servers in the cluster, we can ensure the database will remain available. The shared storage pool requires none of the complex configuration on a server-pair-by-server-pair basis that other failover clustering can entail.

Automated failover can be configured via a simple management interface that makes the configuration of the failover path particularly easy to define.

'Point and Click' type manual failover is available for ease of management of the farm to enable node maintenance with One-click utility-wide install for software installation and updates

Typically, sub-30 second failover time from service disruption can be achieved.

### **In Summary**

There are many good technology solutions that can address individual areas such as high availability, disaster recovery, business continuity or management automation with regards to an enterprise scale deployment of MS SQL servers.

But what we have in HP Polyserve Database Utility for MS SQL is a single, flexible, enterprise-class solution available to benefit any organisation suffering from MS SQL Sprawl and it can really make significant impact on Capital and Operational expenditure and service improvement.

*Only one HP Partner in the UK has extensive production design and implementation experience of HP Polyserve Database Utility for SQL Server. If you believe your organisation suffers from MS SQL Sprawl, Computacenter can look at your current infrastructure and show where the HP Polyserve technology can be leverage to implement operational efficiencies provide fast return on investment.*