Server Consolidation – A Software Perspective

By

John D. Peedle

January 2004

1 Server Consolidation – a Software Perspective

1.1 Abstract

This paper considers the drivers towards Intel-based server (also known as Industry Standard Architecture or ISA) consolidation. It then looks in more depth at a software product suite, *VMware,* which offers solutions in this area.

1.2 Acknowledgements

Figure 1 is copyright VMWare inc.

1.3 The need

Server consolidation is an increasing priority for today's IT departments. The proliferation of ISA servers that has taken place over the last decade has left many organisations with large numbers of servers, most of which are not used to best advantage and are typically 1-2 CPU systems running at much less than 50% CPU utilisation. This has led to heavy over-provision when the organisation is viewed as a whole.

The lion's share of the ISA server TCO is the ongoing administration overhead, which accounts for 60% - 70% of total costs[1].

Recent figures from Gartner[2] demonstrate the level of interest in server consolidation. 61% of their interviewees already had a server consolidation project underway, with a further 25% considering the problem. When asked what kind of consolidation was being considered, 58% were looking at Physical consolidation.

Gartner Group has made the following high probability prediction:

> *By 2008, enterprises that do not leverage virtualization technologies will spend 25 percent more annually for hardware, software, labor and space for Intel servers, and 15 percent more on the same for RISC servers[3].*

There is a growing market out there.

1.4 The problem defined

Until the advent of the Intel 80386, no memory protection was available to individual programs. This substantially limited the opportunity for high availability on the Intel platform and it was usually relegated to non-critical operations.

There is a limit to the number of individual applications that can be run concurrently on any single Windows™ instance. Not only are there compatibility issues between applications, inherent platform instabilities also quickly increase the risk of application unavailability above an acceptable level.

In order to allow a more efficient use of system resources, a means of partitioning an individual server is required. Each partition must exist as a distinct operating system entity in its own right, but remain subservient to the controlling system, only having access to those resources allotted to it. Of critical importance is the principle of separation; whereby no errant application can affect another.

---

[1] The Next Virtual Reality: Server Consolidation - http://www.battery.com
[2] Server Consolidation: An Updated Look – Gartner Inc. May 2003
[3] Gartner Note Number: SPA-21-5502

IBM aficionados will instantly recognise the concept of LPARs. Introduced in 1988 for the S/370 and 1999 for AS/400, for some time this had been the only option for dynamic partitioning. This changed with the emergence of software alternatives into the market.

1.5 A new alternative

With roots tracing back as far as Turing's Universal Machine, the idea of emulating hardware in software is not new. One of the first viable solutions was the UCSD Pascal *p*-system. This is in many respects the forerunner of Java and allows the same executable to run unchanged on many different platforms, both big and little endian.

The commoditisation of the IBM PC clone brought more software development opportunities to the home enthusiast. Previously, talented programmers made their money and reputations creating games for home computers such as the Commodore 64 and its many competitors. Once operating systems such as the various BSD flavours and GNU/Linux appeared, with the powerful and free development tools available for those platforms, interest blossomed and people started to consider what else could be achieved on the Intel platform. Many software engineers had also worked either for IBM or with its systems and were aware of the advantages of partitioning. The question started to be asked: 'Can we emulate some kind of LPAR functionality on an Intel system?' The commonly adopted answer was to create a wrapper layer around the actual hardware and represent whatever interfaces were required to the child operating system in software. The idea of the *Virtual Machine* was born.

The earliest true virtual machine appeared in the late 1990s. Open source projects such as Bochs[1] were followed by the first commercial offering from a small Palo Alto based company, VMware, in 1999.

1.6 From little acorns

Formed in 1998, with its first product launched in 1999, VMware now brings much IBM LPAR functionality to the Intel space. Individual virtual machines are instantiated as *guests* within the host environment under the dynamic control of the parent operating system.

The original product, *VMware Workstation,* operated as an application in user space. A number of guest instances could be run, subject to the capacity of the host system hardware. The host operating system could be either Windows or Linux and the guest could be virtually any operating system capable of running on an Intel x86 CPU. The drawback was that, as a user-space application, an individual instance could quickly appropriate more than its fair share of CPU resources to the detriment of other instances. Dynamic allocation of resources was also unavailable.

[1] http:// Bochs.sourceforge.net

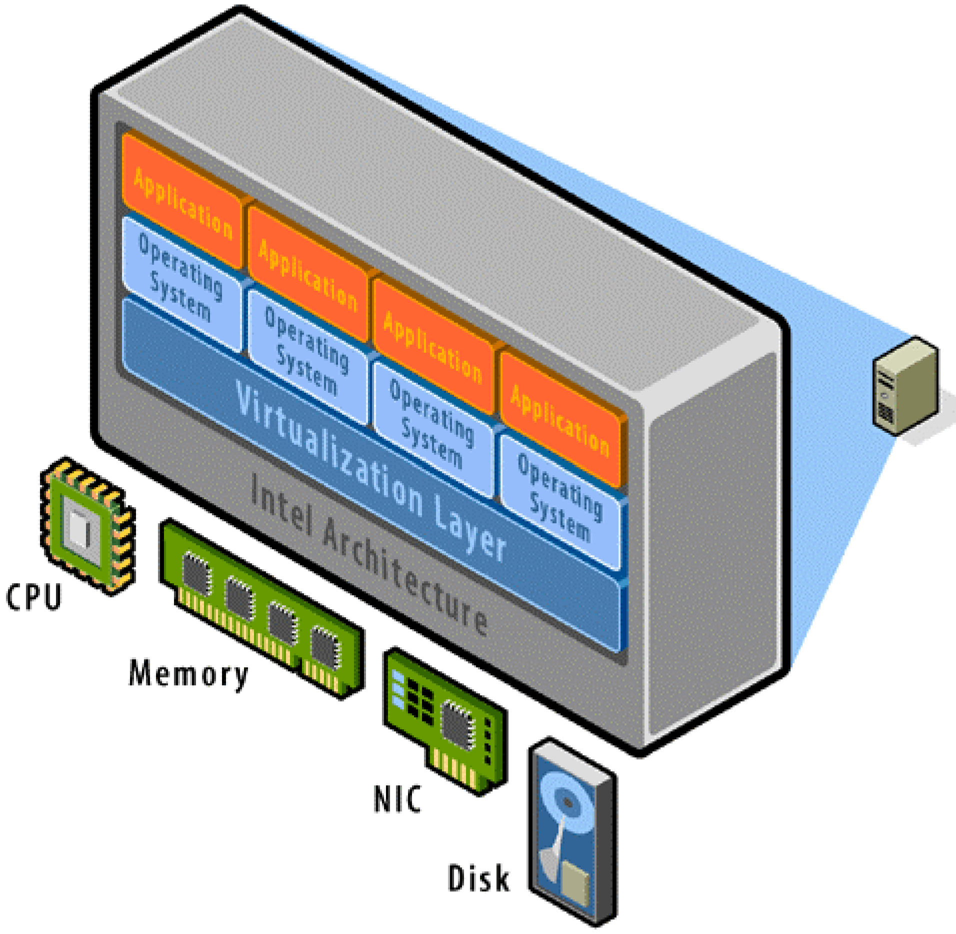

Figure 1 - VMware ESX/GSX server architecture

In 2001, VMware launched a new product aimed squarely at the server consolidation space. ESX Server uses a customised distribution of Linux (based heavily on Redhat Linux) and offers considerable control over each virtual machine. Launched at the same time, GSX server offers similar functionality, but aimed at smaller, less critical deployments. Figure 1 shows the architecture of the current product suite that now encompasses the three variants (including VMware Workstation).

Feature	GSX Server 2.5	ESX Server 2
Environment	Workgroup	Enterprise
Target System	2-8 CPU	2-16 CPU
Resource Controls	Static memory	Dynamic memory, CPU, disk, network
# of Virtual Machines	4 per CPU (16 Max.)	8 per CPU (80 Max.)
System RAM	Up to 64GB	Up to 64GB

Table 1 - ESX/GSX server comparison

1.7 ESX Server

The flagship product, ESX Server is designed to function in the Enterprise space. Typically platformed on 8-way and above servers, it allows every last processor cycle to be used to best advantage. Dynamic control of host resources allows for fluctuating needs within the modern business. Coupled with the recent advantages in hot-plug support for ISA servers it allows for very high levels of availability.

The controlling operating system is a customised distribution of Linux. A limited number of device drivers are supplied and supported, allowing a high level of compatibility and reliability to be assured.

Each ESX Server machine can host up to 80 virtual CPUs in virtual machines (and up to 200 registered virtual machines) or up to 8 virtual machines for each CPU.

The maximum physical storage assignable to a single ESX server is only limited by the following factors:

- 16 host bus adapters per ESX Server system
- 128 logical unit numbers (LUNs) per storage array
- 128 LUNs per ESX Server system

Each virtual machine can have up to two virtual CPUs (the host system must be multi-CPU to support this and the SMP add-on must be purchased). Up to 3.6GB RAM can be assigned. Storage capabilities for each virtual machine are as follows:

- 4 host bus adapters per virtual machine
- 15 targets per host bus adapter
- 60 targets per virtual machine; 256 targets concurrently in all virtual machines

Server provisioning is reduced from many hours to a few minutes. All that is required is to select an appropriate system image, copy it to the desired storage destination, allocate hardware resources and launch the new server, complete with built operating system. As the virtual hardware is based on standard devices, no additional manufacturer's device drivers are required over and above those supplied with the operating system. Once the image is deployed, specific server configuration can be carried out as normal.

1.8 GSX Server

This product is designed to fulfil the needs of an organisation needing to deploy large numbers of server instances where the fine-grain resource control available within ESX server is not required. This might include training departments where each student requires an individual server instance or large test or demonstration facilities. In addition, GSX server runs under a host operating system that can be either Windows or Linux and thereby supports a wider range of host hardware configurations. This allows for greater hardware compatibility, but at the potential cost of reliability.

1.9 Virtual Center

This is a management console that allows visual control of all an organisation's ESX servers. Dynamic control of each ESX server and its subordinate virtual servers is available from a single point. A library of server templates can be created and managed to allow for fast server provisioning. A running server can be used to create a template for future re-use, thereby allowing best practice methods to persist. Threshold and event alerts are configurable to assist with administration. A comprehensive Application Programming Interface (API) allows many functions to be automated thereby further reducing the administration overhead.

1.10 VMotion

This is a powerful feature that allows a *running* virtual machine to be migrated from one physical ESX server to another with virtually no interruption in processing. If a particular service or application requires more physical resources than are available on the current host, the virtual machine can be moved to a different host with the required capacity. A host exhibiting hardware problems can be quickly removed for repair with little or no impact on the services it delivers. The running state is transferred across to the new host and the load taken on by the new virtual machine. The period of outage is typically 2 seconds, which in an IP environment will be usually unnoticed by the user.

1.11 Migration Support

Migrating physical servers to virtual machines can be a complex and lengthy task. VMware has created a tool that is designed to assist in this process. The source system is automatically analysed and a new virtual server created with all the applications configured ready to run. This tool is also useful for diagnosis of production systems; the system can be migrated to a virtual server for detailed analysis without impacting on the live system and its users. Disaster recovery is also made simpler as snapshots can be taken regularly.

1.12 Conclusion

The need for physical server consolidation will impact most businesses in the near future. The costs of not consolidating are anticipated to be considerable and will impact on a business' potential to grow. The opportunities for service improvement and reduction in administration overhead cannot be overstated. Products such as VMware will grow in strategic importance – indeed this is amply demonstrated by the recent acquisition of VMware by the storage manufacturer EMC.