Server Consolidation through Virtualization with Quad-Core Intel® Xeon® processors

Abstract: Server Consolidation can be adopted in both homogeneous and heterogeneous environments, wherein the servers cater to different types of applications. This paper showcases the benefits of virtualization using Quad-Core Intel® Xeon® processors to consolidate various physical servers and reduce total cost of ownership.

Infosys® POWERED BY INTELLECT DRIVEN BY VALUES

intel®

Table of Contents

Executive Summary 3
Server Consolidation Using VMware ESX Server 3.0.1 4
Proof of Concept (PoC) 4
Architecture 5
Approach 8
Application and Workload Selection 9
  File Server 9
    Workload 9
    Tool 10
  Web Server 11
    Workload 11
    Tool 11
  Active Directory Server 12
    Workload 12
    Tool 12
Test Results 12
  Homogeneous Experiments 12
    File Server Consolidation 13
    Web server Consolidation 18
    Directory Server Consolidation 19
  Heterogeneous Experiments 21
  Stability Test 24
Recommendations 25
TCO 25
Conclusion 28

Jake Smith

Intel Corporation

Systems Integration Group

Infosys Technologies

January 2008

Executive Summary

Today's IT enterprises offer an array of services to support their increasing business growth and customer needs. Data centers consist of a complex mix of underutilized server platforms, other hardware and applications. Moreover, mergers and acquisitions have brought in their share of complexity to the existing environment, resulting in the following consequences:

- Increased number of servers in data centers
- Increased need for reliability and availability issues
- Increased complications in contingency planning
- Increased server administration and maintenance costs
- Increased TCO (Total Cost of Ownership)

In search of solution, Enterprises look into server consolidation as viable solution to their problems. Traditional server consolidation promises cost reduction through consolidation of hardware and software; while costs are reduced, performance, availability and agility can also be compromised. Server consolidation using virtualization may not only reduce the costs of hardware and software licensing while retaining performance but also offers simplified architecture, improved quality of datacenter management, reduced power consumption and cooling requirements, simplified backup and recovery activities, and significantly enhances the enterprise agility, as described in Figure 1. The business and technical rationale for server consolidation using virtualization technology is clear: bottom line benefit to the IT budget as well as improvement in the productivity of the enterprise workforce.

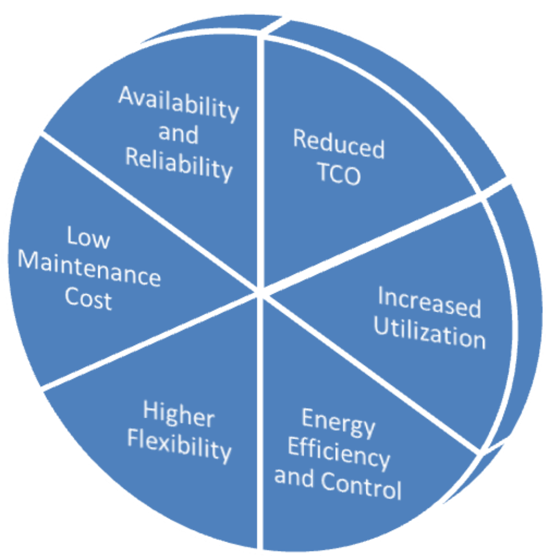

Figure 1: Benefits of Virtualization

Infosys, a global system integrator with proven expertise and a track record in the server consolidation arena, has teamed with Intel to demonstrate the scalability of Intel's Quad-Core Intel® Xeon® processor based server and VMware's* server virtualization products.

This paper highlights following key points

- Server Consolidation through use of Virtualization technology is a key strategy for improving flexibility, reducing TCO of IT infrastructure services and increasing enterprise agility
- Sizing and scalability of resources and allocation of workload across virtual machines is a project risk to be managed in a server consolidation project
- Presents and validates a methodology for baselining scalability of multiple VM workloads on Intel's Quad-Core Intel® Xeon® processor based servers
- Extrapolating the projected capacity requirements to size resources and allocate VMs for practical homogenous and heterogeneous server consolidation scenarios

Server Consolidation Using VMware ESX Server 3.0.1

Virtualization has emerged as a compelling technology for server platforms, presenting data center managers with an ability to consolidate multiple workloads on one physical server. Such consolidation offers lower hardware acquisition costs as well as improved data center maintenance, performance, and power efficiency.

Intel® Virtualization Technology (Intel® VT) consists of a set of processor enhancements to traditional software-based virtualization solutions. These integrated features give virtualization software the ability to take advantage of offloading certain workloads to the system hardware, enabling more streamlined virtualization software stacks and "near native" hardware performance characteristics.

Server consolidation using VMware Server virtualization technology revolutionizes the way IT infrastructure is managed. It streamlines the computing environment; offers the ability to run multiple instances of one or different operating systems on single physical server, simplify the server provisioning, ease the management of solution life cycle from test and development to production, enhance high availability and eliminate the need for porting the legacy application to new platform.

Proof of Concept (PoC)

The purpose of this PoC is to showcase the server consolidation benefits using real world end user scenarios and by calculating the TCO for given use case.

To further examine application and server behavior in a virtual environment we employ a dual testing strategy, first run test using homogeneous environment (File, Web or Directory server) and repeat the test using heterogeneous environment (mix of Web, Directory and File servers). The aim is to stress the single server with maximum number of virtual machines (VMs) until the defined exit criteria is met. The exit criteria determined based on real world customers' requirements. The server threshold parameters are as follows.

- Average CPU Utilization - 65%
- Average Memory Utilization - 75%
- Average Disk Utilization – 75%
- Network Utilization was not considered in the virtualized environment.

The PoC was conducted by joint team from Infosys and Intel at the Infosys labs. The Quad-Core Intel® Xeon® processor server loaded with VMware's hypervisor (VMware® ESX Server version 3.0.1) combined with Fiber Channel SAN from NetApp*.

Architecture

The architecture was designed to represent a real world scenario of a server consolidation. The diagram below shows a server setup in a consolidation scenario. Clients are connected to a typical high-end server hosting the VMs, which communicates to the SAN through a fiber channel.

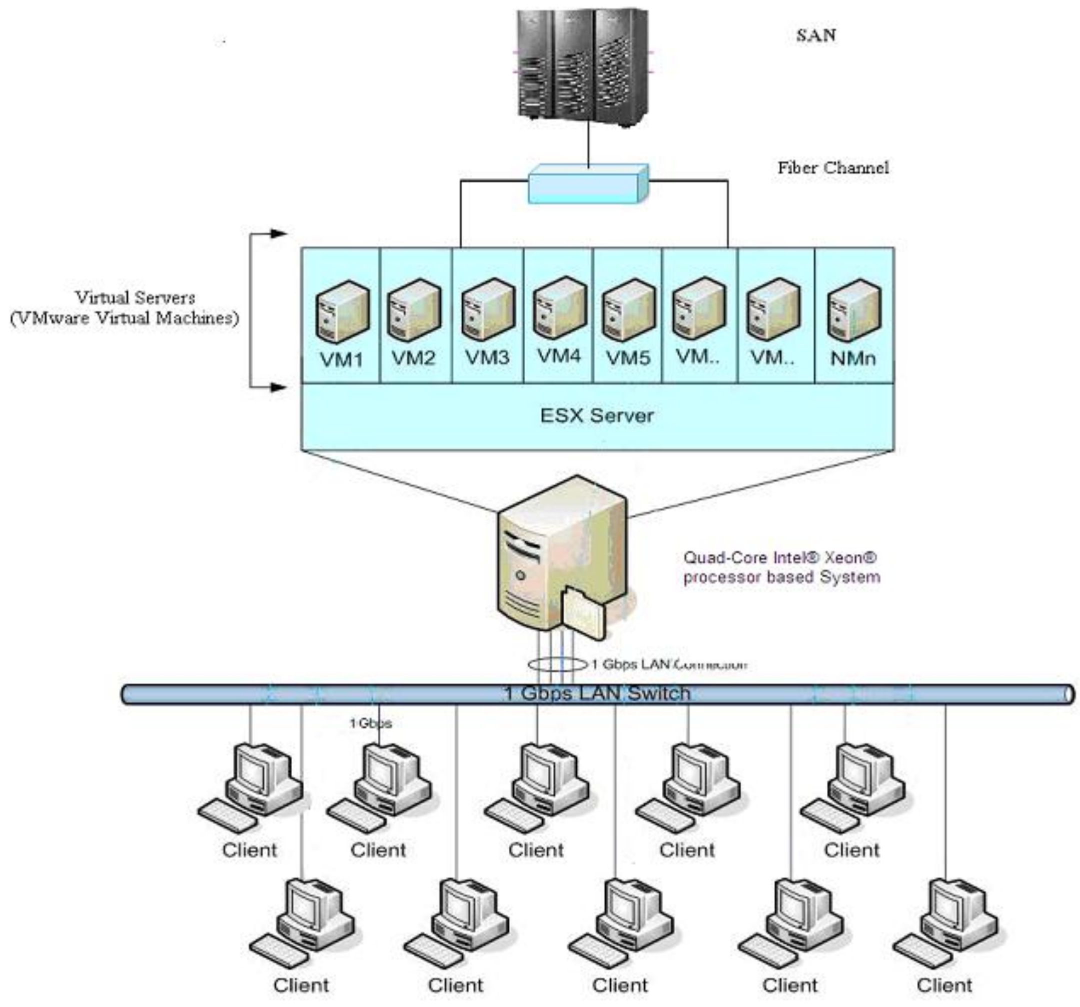

Figure 2: Architecture for Server Consolidation

The main components of the PoC infrastructure are as follows:

Server – System under test

Client PCs – Traffic generators

Network – Switches and cables for interconnection

SAN – VM storage

Server
Intel's EPSD Dual-processor Server with two Quad-Core Intel® Xeon® processors (X5340) represents a mid-high performance server to consolidate multiple file servers, web servers, and directory servers. The detailed server configuration of each blade is shown in table 1.

Table 1: Server Details

Make and Model	**Intel EPSD System**
CPU	Two CPU (Quad-Core Intel® Xeon® processor X5340 (4*2.40GHz) 8MB L2, Cache, 1.06GHz FSB
RAM	2GB x 8 memory (667 MHz)
Storage Internal	4x 72 GB Internal Disc SAS 15K RPM
Network	4 x 10/100/1000 Mbps Ethernet Ports
Hypervisor	VMware ESX Server 3.0.1
HBA	Qlogic* QLE2462 (Dual Channel – 4Gbps, Fiber Channel)

Clients
The Client machines loaded with appropriate tools were used to generate the workload on the server. Total of 15 PCs were used in the PoC, with client system configuration as shown in Table 2.

Table 2: Client PC Details

Make and Model	**Intel® Pentium® 4**
CPU	2.8 GHz
RAM	512 MB
Storage Internal	37GB
Network	Connected to 1Gbps switch
OS	Microsoft* Windows* XP

Network

The Server is configured with a total of four Gigabit Ethernet NICs. Best practice is to dedicate one Ethernet card to the VMware ESX Server service console. The virtual machines are distributed across the remaining three NICs to enable load balancing of the network traffic during data transfer between the VMs and clients. Figure 3 illustrates VM connectivity and the virtual network topology:

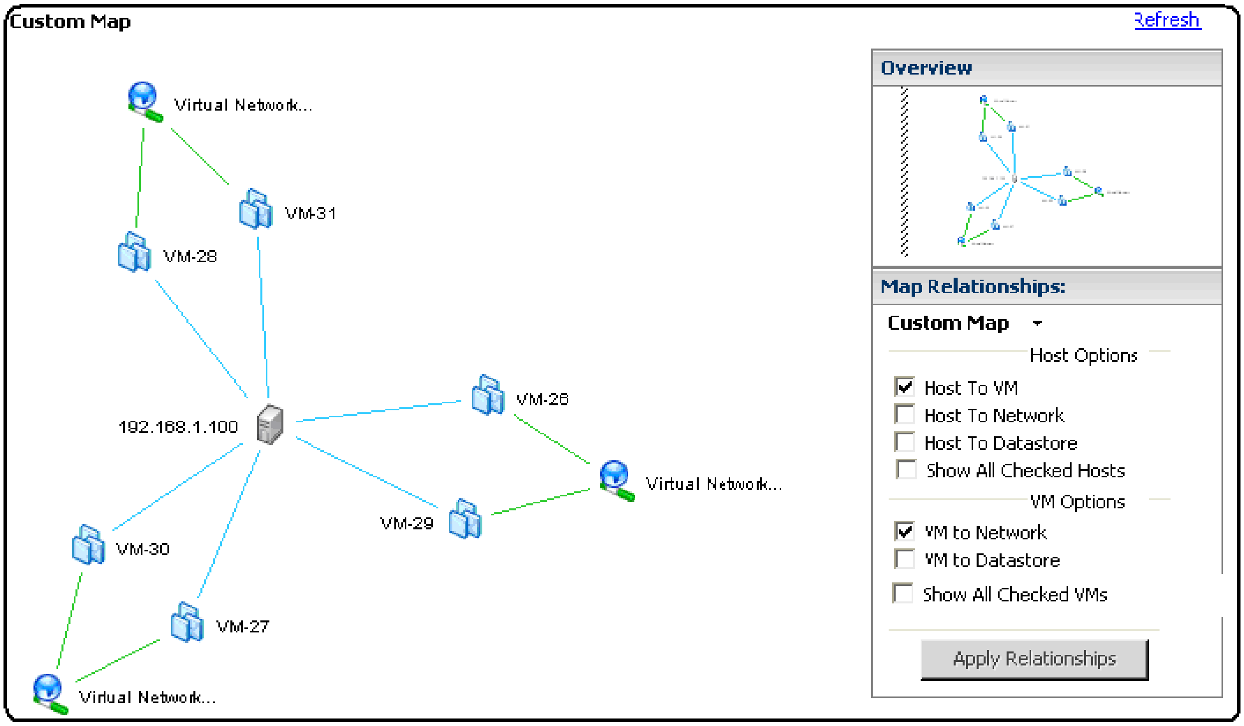

Figure 3: Virtual Machine to Network Association

A Dell* PowerConnect switch was used to connect client systems to server, switch configuration shown in table 3.

Note: The best practice is to dedicate one Ethernet NIC card for VMotion feature, since the tests conducted did not use the live migration feature we have dedicated one NIC for the service console and the other 3 NICs were distributed among the Virtual Machines (VMs) .

Table 3: Network Switch Details

Model	**DELL POWERCONNECT 5324**	
Ports	Ports	24
	Port Speed	1Gbps

The Server is connected through gigabit links while the clients are connected through 100 Mbps links to the Switch. The Host Bus Adapter (HBA) QLogic QLE2462, Dual Channel, 4 Gbps, Optical, Fiber Channel card is used (PCIe) for connecting server to SAN.

SAN

A NetApp FAS3040 FC-SAN was chosen for this PoC. Generally, ESX Server hosts are connected via fiber channel to boot and swap volumes hosted on shared data storage as required to use features like vMotion, DRS and HA. VM image volumes and log files are stored on LUNs (Logical Units) on the FC-SAN as shown in Table 4.

Table 4: NetApp SAN

Total Usable Space	1TB (14 disks of 144GB each)
LUN Space	Each LUN is of 20 GB
RAID	RAID-DP
Connectivity to Server	Fiber Channel

NetApp Unified Storage supports FC SAN, IP SAN and NFS protocols on the same storage system; this configuration is supported by VMware. RAID-DP on NetApp SAN provides protection against dual disk failure in the same RAID group without any performance overhead. NetApp FlexVol helps rapid disk provisioning, improves disk utilization and increases performance by striping data across multiple disks spindle.

NetApp SnapShot can be implemented to help provide faster VMware backup and restoration of data with little performance impact on servers. NetApp storage systems help provide instantaneous and space efficient cloning of VMware data with little impact on server performance.

Approach

The PoC follows a two phase approach to test the real life scenarios. The consolidation can be realized as either a homogeneous consolidation or a heterogeneous consolidation. In a homogeneous consolidation all the consolidated servers (onto a high end server) are of a similar type, say fileservers. The first phase will stress and measure the scalability of the individual workloads in homogeneous (single application type) configurations. In the second phase, capacity planning for distribution of heterogeneous workload VMs will be projected based on the first phase data, and the ultimate scalability, performance, and reliability of the solution will be verified.

The homogeneous consolidation has the following characteristics:

- The average resource utilization for a particular Application type
- Resource Consumption of a particular server resource

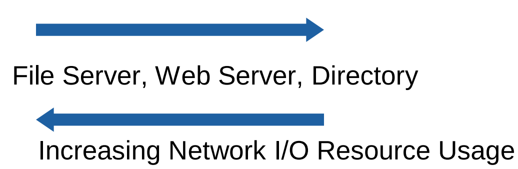

The heterogeneous consolidation proffers the following benefits:

- Leads to an optimal server resources (CPU, Memory, I/O. etc) use rather than predominantly utilizing one server resource while the others remain underutilized. This is possible because some kinds of workloads are CPU intensive while others require I/O bandwidth etc.

Tests were conducted on both homogenous and heterogeneous environment.

In the homogenous case for fileserver two sets of tests were conducted. Test-I entails increasing the number of users per server (load) for a fixed number of servers. Test-II keeps increasing the number of virtual servers, keeping the number of users per server constant.

In the case of Web and Directory servers (homogeneous case), the load was kept constant and the number of VMs was increased. For web servers, the number of users per server typically corresponds with the number of requests per second; similarly, for Directory, operations/second is the parameter that better captures users' activity. For the heterogeneous testing combination of File, Web and Directory servers were chosen which are detailed in the sections below.

In addition, a set of stability tests was undertaken to test the servers' ability to sustain moderate loads over long periods.

Application and Workload Selection

File Server

It is estimated that 20 percent of Microsoft* Windows Server*-based servers used in enterprises are file and/or print servers. The demand for file servers has grown with the increase in storage requirement. The existing infrastructure in many organizations is insufficient to maintain the expected quality of service with lower cost. Through consolidation, IT department can save on the TCO and improve the overall efficiency, high availability, and data protection of the critical file-based infrastructure.

Workload

The size of files transferred in the test workload is distributed in the ratio of 80:20, where 80% of the clients transfer small files and 20% transfer large files. To simulate typical usage, 60% of operations per VM are READs and 40 % operations are WRITEs.

READ and WRITE signifies the following operations

READ – Reading from the server and writing on to client.

WRITE – Reading from the client and writing on to the server.

There are two variables in case of file servers - the number of users and the number of virtual machines. In the first scenario, we kept the number of VMs constant (9 VMs), and increased the number of users until the defined exit threshold is reached, i.e. until we reached the target CPU, Network I/O, or memory utilization. In the second scenario, keeping users constant (e.g.: 60 users/VM) we increase the number of VMs until the performance threshold is hit.

The **Consolidated Workload (CW)** for File Servers is of two types – THIN and THICK.

THIN CW loads the system with small file sizes while the THICK CW uses large file sizes to load the system. This means if we run a THIN CW on a client system then file sizes ranging from the size 0.10MB to 5MB would be read and written on the server under test.

Consolidate Workload (CW) File Size Details

Thin CW (MB)	0.10, 0.15, 0.25, 0.50, 1, 2, 5
Thick CW (MB)	50

This load was chosen to simulate a real world scenario, which combines small as well as large file transfers. More details on the CW are given below.

Tool

Perl* scripts were used to generate the file server load. Scripts performed read and write of thin and thick CWs from the clients to the server. To simulate multiple users with limited client workstations, the scripts used thread based methodology for transfer; and each thread simulated the load generated by a single user. Each CW consists of 9 threads (users) doing a pattern of 'read' and 'write' on a particular VM and each thread communicates to one particular VM.

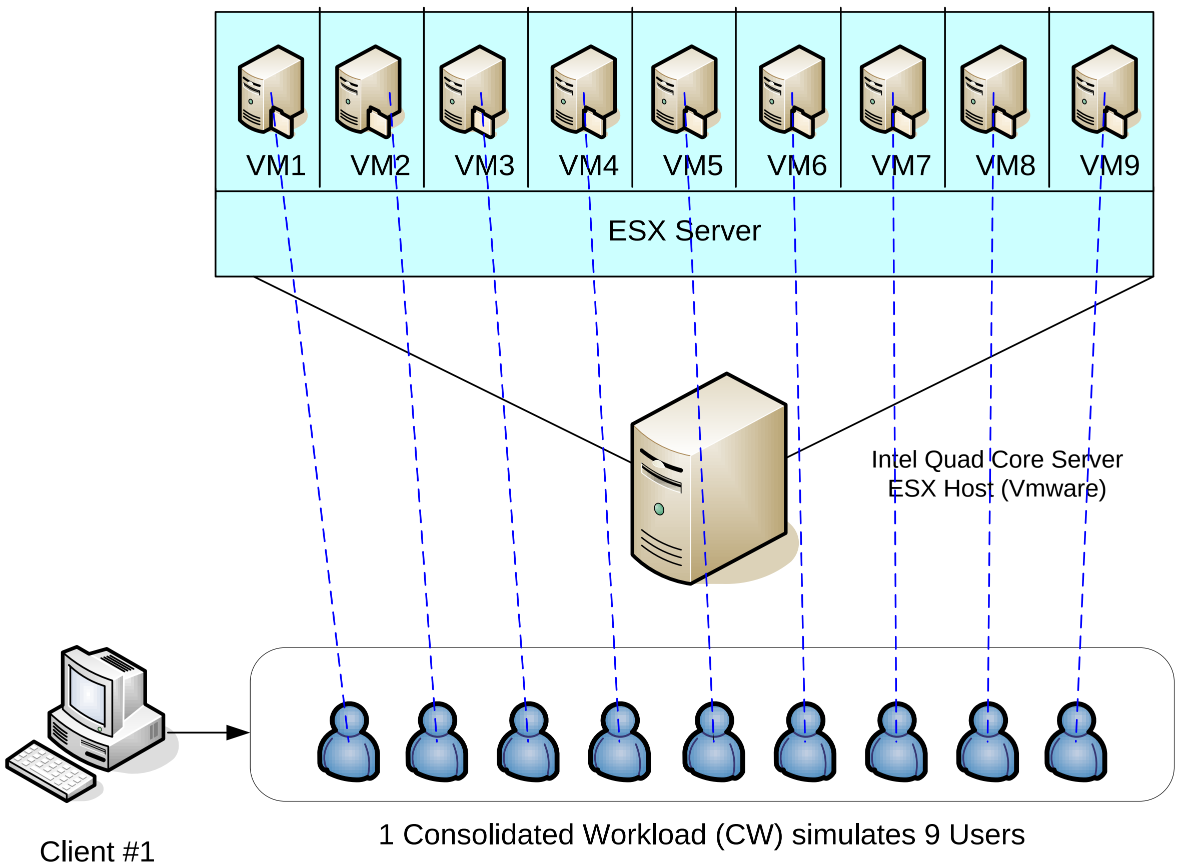

Figure 4: Consolidated Workload (CW) depiction on a physical client

The following table explains the mechanics of file transfer involved when a single CW is executed from a client. Both THIN CW and THICK CW are depicted below.

Table 5: Consolidate Workload Distribution

Thin CW									
Operation	File Size (MB)								
▼	Thread 1	Thread 2	Thread 3	Thread 4	Thread 5	Thread 6	Thread 7	Thread 8	Thread 9
R	0.10	0.15	0.25	0.50	1	2	5	0.10	0.15

R	0.15	0.25	0.50	1	2	5	0.10	0.15	0.25
W	0.25	0.50	1	2	5	0.10	0.15	0.25	0.50
R	0.50	1	2	5	0.10	0.15	0.25	0.50	1
R	1	2	5	0.10	0.15	0.25	0.50	1	2
W	2	5	0.10	0.15	0.25	0.50	1	2	5
R	5	0.10	0.15	0.25	0.50	1	2	5	0.10
Thick CW									
File Size	50	50	50	50	50	50	50	50	50
Operation	R	R	W	R	R	W	R	R	W

Read and write operations are executed in the ratio of 2:1 iteratively on the set of Thin CW files transfer.

The Perl scripts were run on the client machines to initiate the file transfers and log the response time throughput for each file transfer. To simulate real-world user "think time", delays ranging from 1 to 10 seconds between file transfers and delays of 100 seconds between loops were introduced.

Web Server

Web applications are currently large, complex and are used to perform online processing. Online processing is impacted when the operations can't be done in real time. Extension of the application is expected to exceed the capacity of the current IT infrastructure. Thus server consolidation with Web Servers is an attractive way to support the upcoming growth.

Web server consolidation will greatly reduce the power used while improving the performance and gaining better control over the entire system.

One of the benefits of consolidating web servers is the support for larger numbers of users accessing the same application by running on a larger server that may be cost justified in a distributed environment. It also opens up the opportunity to provide new services by enabling or improving interaction between different applications, or between new applications and data in existing systems.

It can be used in Web serving/Electronic commerce based on secure access to data in the existing production systems.

Workload

The web server test methodology involved simulating user-hits on a web server through an application. To generate the appropriate number of hits with limited clients, a typical workflow of customer transaction with the web application was recorded as a script with Microsoft Web Application Stress Tool* (WAST) which is available from Microsoft TechNet*. The think time between user hits was removed and the script was looped to stress the server. A set of parameters, namely requests per second, response time, CPU, memory and network utilization, was monitored to arrive at the results.

Tool

Following are the tools used:

Web Server: Xitami* (version 2.5c2) open source web server by iMatix Corporation*, which is implemented in C and implements LWRP and virtual multithreading. This simple web server is portable to various flavors of Windows, Linux, UNIX*, OS/2, and other systems.

Web Application: Perl, html, JavaScript* based Online Booking Web Application.

Web Application Stress Tool (WAST): Free utility available from Microsoft TechNet to perform stress testing of web servers. Has an elaborate set of parameters that can be configured to suit individual needs. Any workflow of a web application can be recorded as a script and the script can be run to load the server. WAST generates comprehensive reports which include standard performance parameters for a web server.

Active Directory Server

Microsoft Active Directory* (AD) is a distributed directory server component of Microsoft's Windows Server operating system. Multiple AD servers in an enterprise cause multiple redundant copies of directory data to be placed on individual systems. By consolidating multiple AD server instances onto a single physical server we can benefit by largely reducing the operational cost while increasing the informational security.

Workload

A set of logon, search and modify operations were performed on an AD instance running in each VM. Each AD workload consisted of 40,000 simulated users distributed across 20 OUs (Organizational Units). Script based stress methodology was used with a custom script written and integrated with Microsoft ADTest tool.

Tool

The following tool was used:

Microsoft ADTest: ADTest.exe is an Active Directory load generation and performance analysis tool that simulates client transactions for performance and scalability testing. It has features that assist in testing the whole AD directory lifecycle, from creating users and organizational units in an enterprise Active Directory to performing stress test by simulating logon, search, modify and other typical operations. It also provides a set of scripts used to simulate these typical operations, which can be used as-is or can be customized to specific requirements.

Test Results

The tests using servers based on Quad-Core Intel® Xeon® processors show the benefits of using quad-core Intel® processors along with VMware's ESX Server and NetApp SAN for server consolidation. The supporting data from these tests is summarized in this section.

Homogeneous Experiments

Experiments targeting consolidation of the same type of servers are classified under homogeneous experiments. The server was tested for performance and scalability with fileserver, web server and directory server workloads separately until the exit criteria were reached. The results are summarized and analyzed below:

File Server Consolidation

Following are the results that have been collated for the File Server consolidation exercise which was undertaken at the Infosys labs. Two experiments were conducted for the Homogeneous file server consolidation.

Experiment 1: Virtual Machines Constant, Increasing Load

Approach:
In the first set of tests, 9 VMs were created, each running a file server. We arrived at 9 VMs to evenly distribute the load across 3 network ports in the server. Each VM has the following configuration

Table 6: Virtual Machine Configuration Details

RAM Size	1GB
Disk Space	20GB on SAN (separate LUN)
Vcpu	1
Operating System	Windows 2003 Server – 32bit
VMware Tools installed	Yes

For details on CW refer to File Server section.

The number of physical Clients/Computers used to load the server was 15. If in each of the Client we simulate 1CW then the total load on the server is 135 users (15 Physical PCs * 9 simulated users).

We started with one CW running on 15 PCs and scaled to 20 CWs/PC and hence the read/write load was increased in steps. **The graph shows that the defined threshold was exceeded at around 15CWs/PC**.

Hence with the given workload definition, we can scale up to 15CWs/PC which sums to a total of 225 CWS on the server simultaneously accessing the files. (15 CWs/PC = 15 * 9 = 135 Users simulated per PC; we have 15 PCs this means 2025 users in total (135 *15)).

The graph shows the server performance and user experience during the experiments. It indicates that the user experience, throughput in MB/sec is related to CPU utilization plotted against the increasing load.

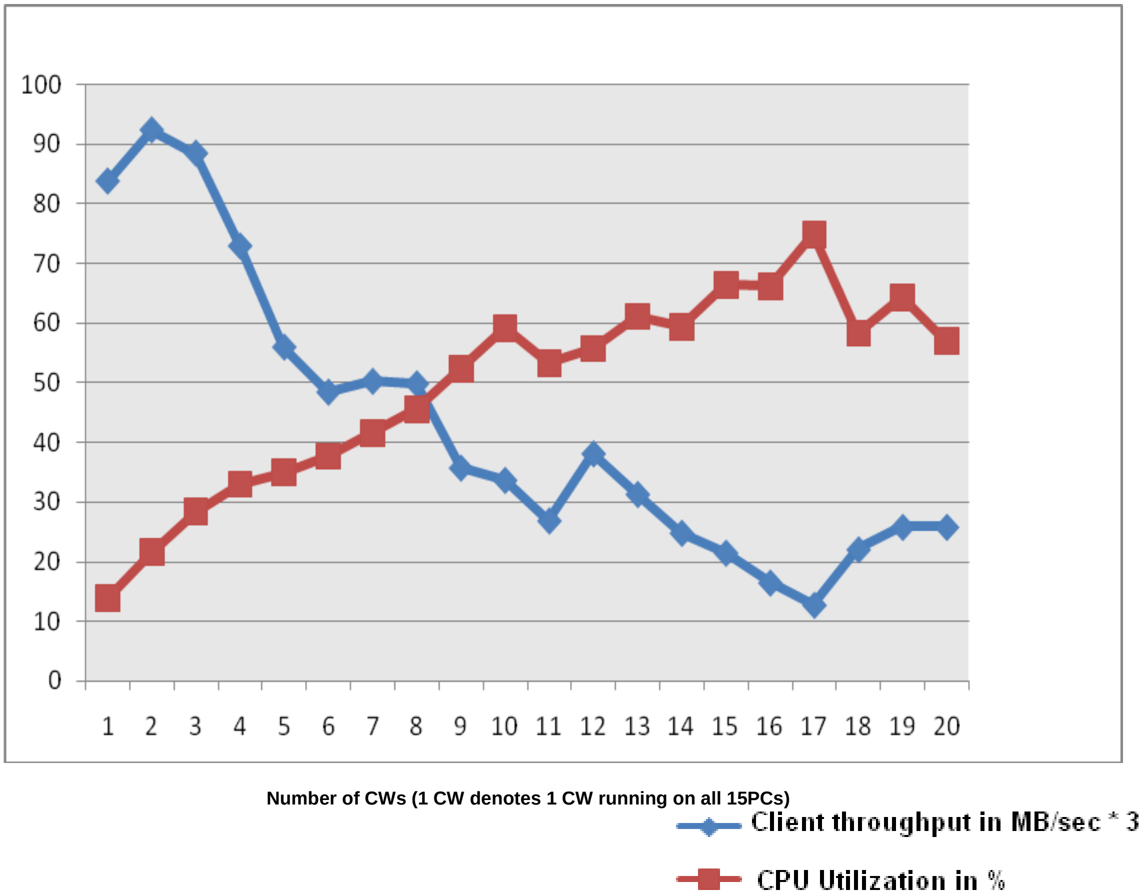

Figure 5: Experiment 1- CPU Utilization Vs Client Throughput

We can observe from the graph that the defined exit threshold for the CPU is reached at 15 VMs. The succeeding graph shows the server performance during the experiments.

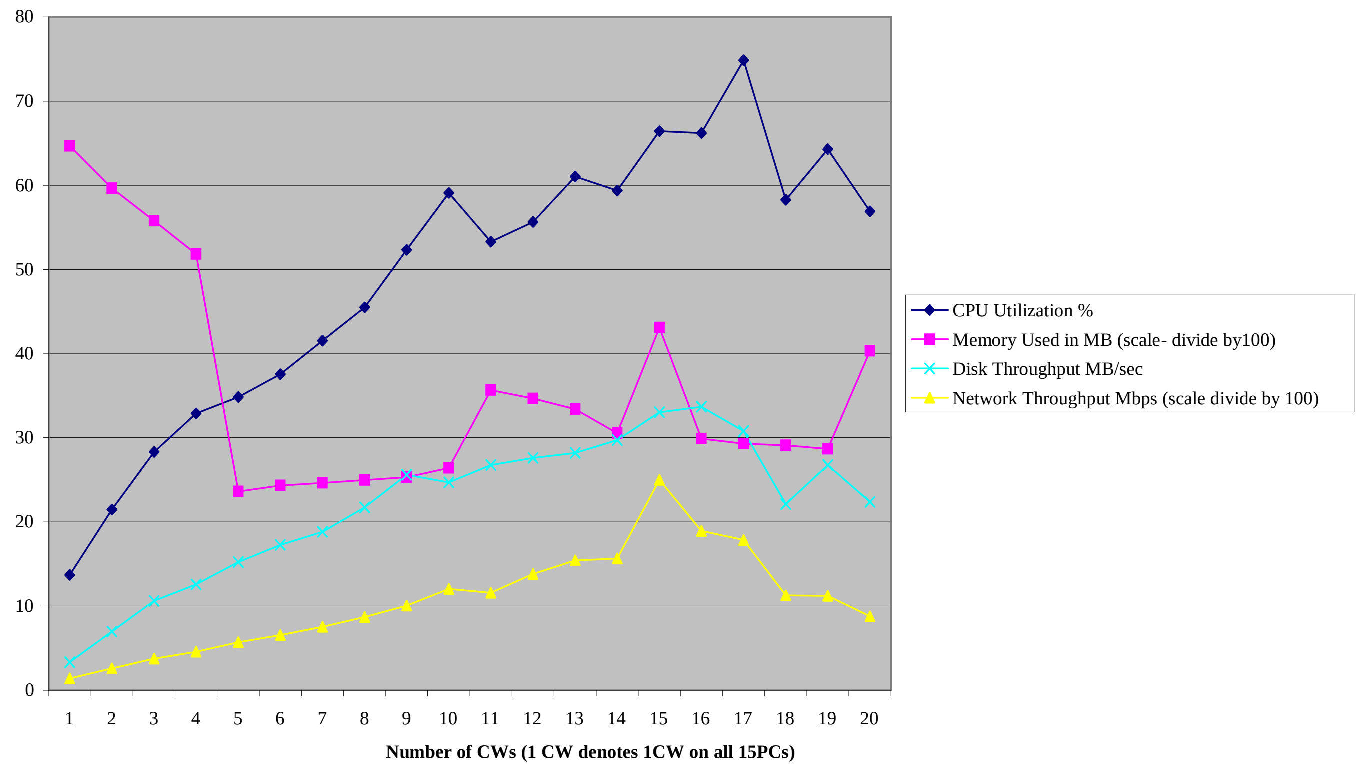

Figure 6: Server Performance 9VMs constant but user load increased

Metrics	Threshold	Threshold reached
CPU	65	Yes
Memory	75	No
Disk	75	No

This shows that CPU Utilization alone exceeded the defined exit criteria.

Experiment 2: Load Constant, Increasing the number of VMs

Approach

In the second set of experiments in file server consolidation, the load per VM was kept constant but VMs were increased consistently. The configurations of the VMs are shown in the following table.

RAM Size	512MB
Disk Space	20GB on SAN (separate LUN)
vCPU	1
Operating System	Windows 2003 Server – 32bit
VMware Tools installed	Yes

The number of users per VM was fixed at 60. We selected 60 users based on the previous experiment which denoted that for a 9 VM case the CPU Utilization for 60 Users/VM was around 30%. The VMs were increased in steps of 3, so that the load on the NIC is evenly distributed, i.e. each NIC was assigned a new VM at the same time.

The following graph indicates the user experience in MB/sec plotted against the server CPU utilization. The time taken in seconds to complete the entire operation and the size of the file is taken to calculate the throughput.

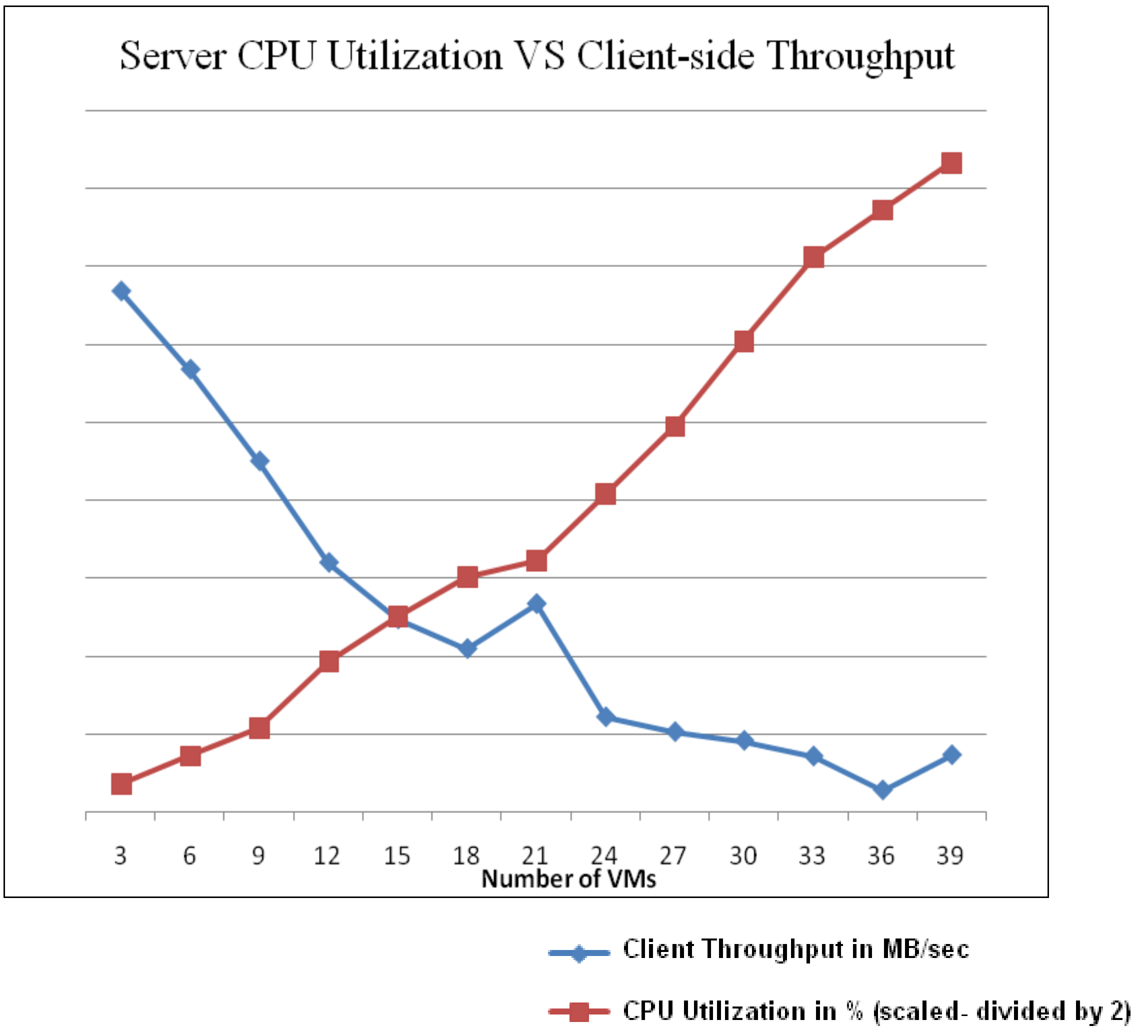

Figure 7:Experiment 2- CPU Utlization Vs Client Throughput

The server performance characteristics are shown in the following graph.

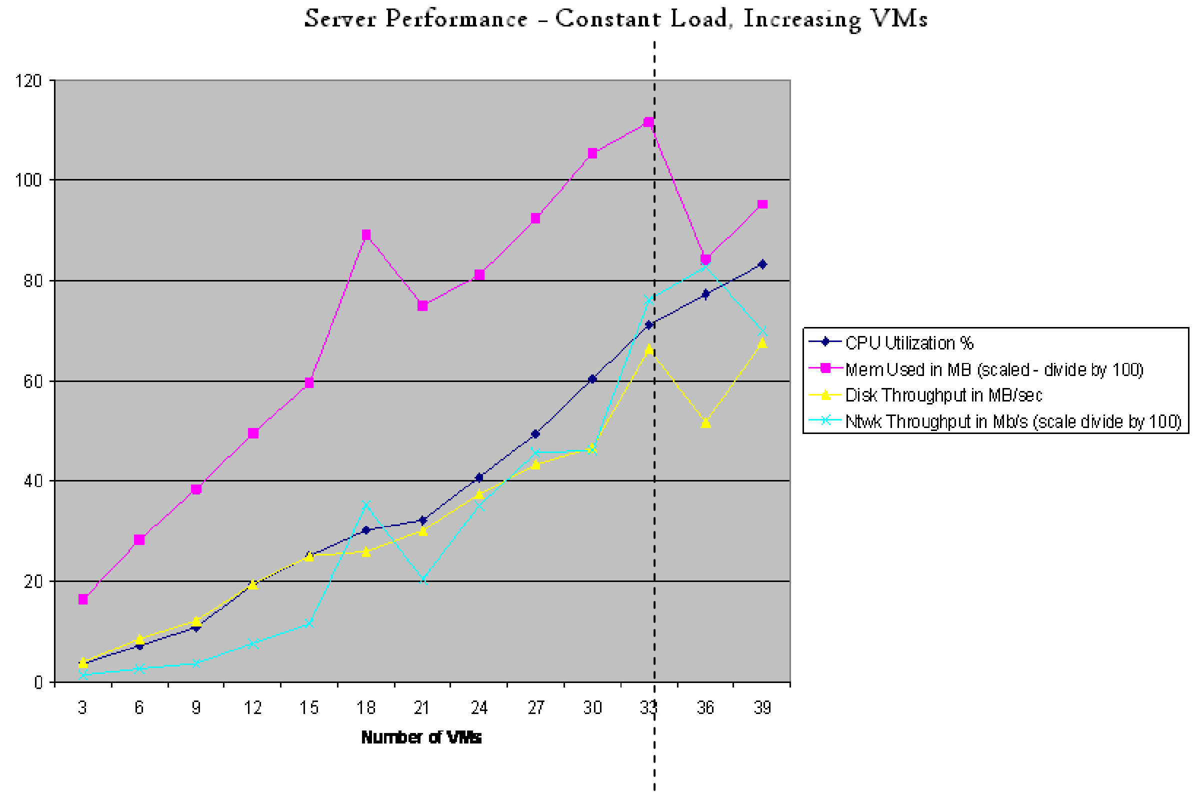

Figure 8: Server Performance- load constant/VM and Number of VMs increased

Metrics	Threshold	Threshold reached
CPU	65	Yes
Memory	75	Yes
Disk	75	Yes

For the given workload when the number of VMs was increased beyond 30, the threshold was exceeded. **Thus 30 traditional file-servers can be consolidated to a single machine for the given workload.** This clearly points to the fact that the ESX Server uses the CPU resources efficiently in spite of the oversubscription (8 CPU cores Vs 30VMs) and the workload throughput increases linearly until the physical CPU resources are used up.

Web server Consolidation

The objective of this experiment is to showcase the maximum number of web servers that can be consolidated on the same server within the defined threshold limits.

Web server and a web application are configured on a VM. Client machine with Windows XP and WAST tool installed on it will stress the web server by increasing the number of concurrent connections and thus the http requests/second (RPS) processed by the web server.

This experiment was carried out by applying a standard load of 47 RPS per web server and increasing the number of servers one by one until we reach the exit criteria.

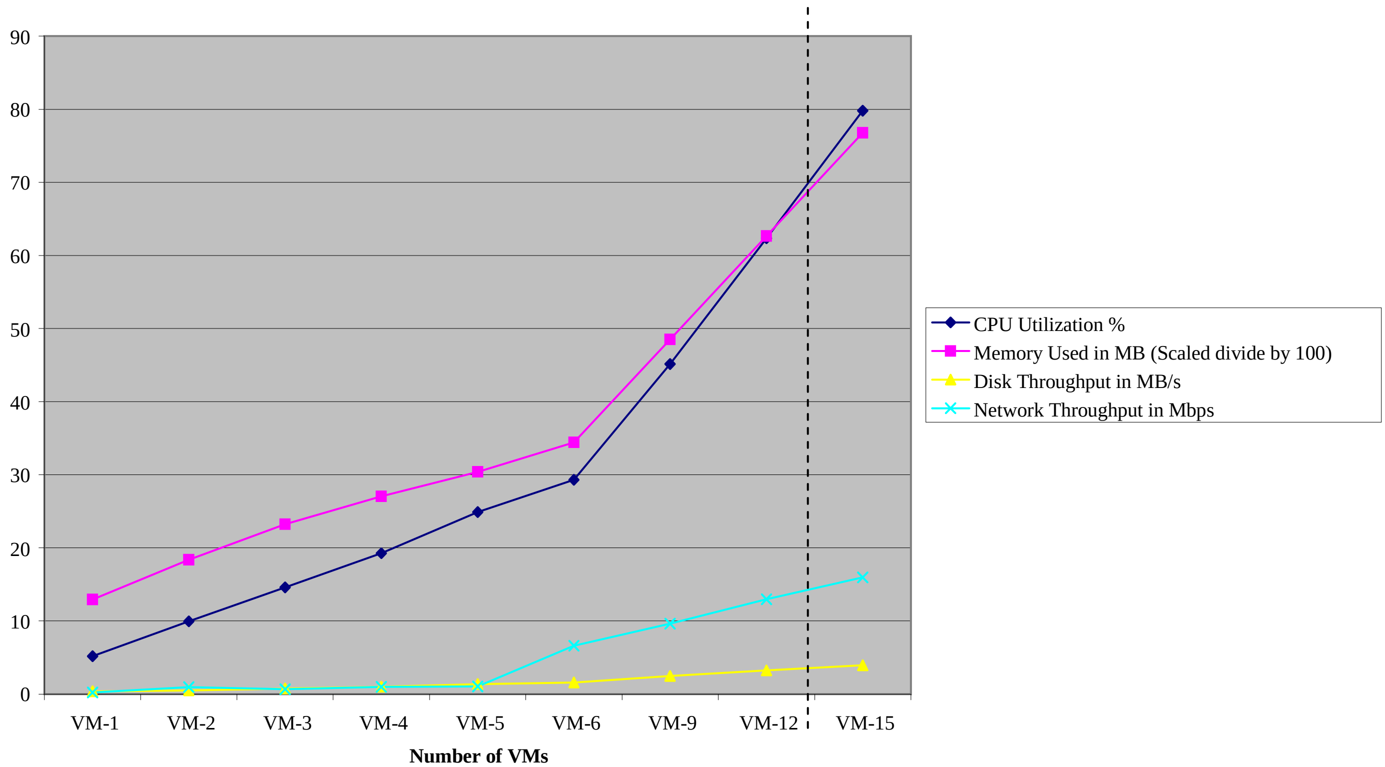

Figure 9: Server Performance- load constant/VM and Number of VMs increased

Parameters	Threshold	Threshold reached
CPU	65	Yes
Memory	75	Yes
Disk	75	No

We observed that the threshold CPU Utilization of 65% is reached when we increase more than 12 VMs for the specified load. The details of the user experience are tabulated below:

Sl. No	Time run (sec)	No of VMs	Requests/ Web Server	Requests/ Sec	Avg Response Time(ms)	CPU Utilization in percentage
1	300	1	14151	47.17	125.70	5.20
2	300	2	14316	47.72	123.71	9.96
3	300	3	14331	47.77	124.10	14.58
4	300	4	14283	47.61	124.40	19.26
5	300	5	14256	47.52	124.69	24.89
6	300	6	14238	47.46	124.75	29.32
7	300	9	14067	46.89	126.50	45.14
8	300	12	13881	46.27	128.67	62.33
9	300	15	13452	44.84	134.15	79.80

The above data can be interpreted to show the maximum number of web servers supported within the threshold limit, which adds up to 12 VMs with around 45 Requests/Sec.

Directory Server Consolidation

Depending upon the nature of business, certain applications may be more critical than others. AD server is one such application which was evaluated and in this section, we showcase the number of AD servers that can be consolidated on quad-core Intel Xeon processor-based systems.

The Microsoft AD server for a domain was configured on the VM. Client machines with Windows XP were added to the domain, and ADTest tool running on the client was used to stress the corresponding domain by performing logon, LDAP search and LDAP modification operations depending on the tests given.

This experiment applied a standard load of 400 Operations/sec (mix of logon, LDAP search and LDAP modification operations) on each AD server instance, and the experiment increased the number of servers one by one until the exit criterion was reached.

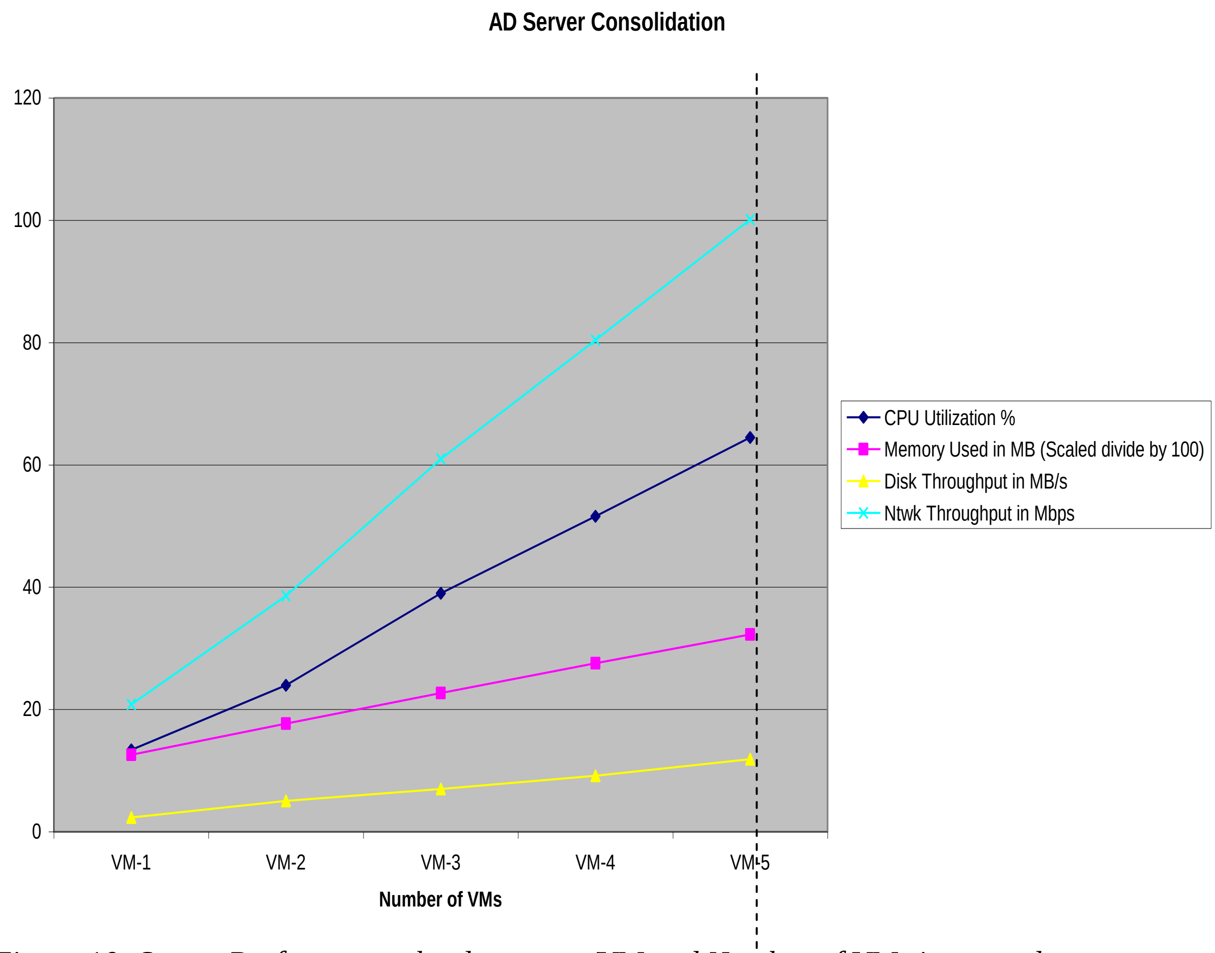

Figure 10: Server Performance- load constant/VM and Number of VMs increased

Parameters	Threshold	Threshold reached
CPU	65	Yes
Memory	75	No
Disk	75	No

From the above graph, we can see that the CPU threshold was reached with 5 AD Servers, while Memory and Disk utilization remained relatively lower. The user experience for the AD Server is detailed below:

Sl. No.	Time run (sec)	Number of VMs	No of operations / AD Server	CPU Utilization in Percentage
1	600	1	398	13.38
2	600	2	409	23.95
3	600	3	407	39.00
4	600	4	391	51.58
5	600	5	410	64.45

From the above table, we can infer that a maximum of 5 AD Servers can be consolidated within the CPU utilization threshold limit of 65%.

Heterogeneous Experiments

Heterogeneous experiments consist of a mix of File, Web and Directory servers hosted on the consolidation server. Three different mixes were identified to correspond to three different scenarios in which one of File, Web and Directory servers predominates in resource utilization. The number of servers of each type in a mix was formulated with inputs obtained from homogeneous tests so as to almost reach the exit criteria and yet remain beneath it. Also, the load chosen for each of File, Web and Directory servers was the same as that in the homogeneous experiments.

The objective of the three scenarios is to discover the VM allocation mix that enables hosting of the maximum number of File, Web and Directory servers respectively and the details are given below.

Scenario 1:
Predominant File Servers

Scenario	No of File Servers	No of Web Servers	No of AD Servers
1	**15**	4	1

The user experience of the Web, AD and file servers are tabulated below:

FILE SERVER		
Scenario	**Number of File Servers**	**Throughput in MB/sec**
1	15	14.99

WEB SERVER					
Scenarios	**Time run (sec)**	**Number of Servers**	**Requests/WebServer**	**Requests/Sec**	**Avg Response Time(ms)**
1	649	4	29203	44.97	133. 91

AD SERVER		
Scenario	**Number of Ads**	**Operations/sec**
1	1	362.79

The individual results of Web, AD and File Server for scenario 1 is specified in the above table. The 15 File Servers were stressed with 60 Users/Server and maintained an average throughput of 14.9 MB/sec. The test was run with an average of 45 Requests/Sec and Response time of 133.9 for the Web Server, while the AD completed 362.8 Operations/Sec. The CPU threshold nearly reached as the utilization of 58% getting closer to the defined exit criteria.

Scenario 2:

Predominant Web Servers

Scenario	No of File Servers	No of Web Servers	No of AD Servers
2	3	**7**	1

The user experience of the Web, AD and file servers are tabulated below:

FILE SERVER		
Scenario	Number of File Servers	Throughput in MB/sec
2	3	27.88

WEB SERVER					
Scenarios	Time run (sec)	Number of Servers	Requests/WebServer	Requests/Sec	Avg Response Time(ms)
2	625	7	29007	46.41	127.53

DIRECTORY SERVER		
Scenario	Number of Ads	Operations/sec
2	1	407.49

The results of Web, AD and File Server for scenario 2 can be analyzed as follows. Only 3 File Servers were used with 60 Users/Server and had an average throughput of 27.8 MB/sec. For the Web Server, the average Requests/Sec is 46 and Response time of 127.5. The number of Operations/Sec for the AD is 407. The CPU utilization of 55% was reached with the above setup.

Scenario 3:

Predominant Directory Servers

Scenario	No of File Servers	No of Web Servers	No of AD Servers
3	6	5	**2**

The user experience of the Web, Directory and file servers are tabulated below:

FILE SERVER		
Scenario	Number of File Servers	Throughput in MB/sec
3	6	25.68

WEB SERVER					
Scenarios	Time run (sec)	Number of Servers	Requests/WebServer	Requests/Sec	Avg Response Time(ms)
3	541	5	24967	46.23	128.33

DIRECTORY SERVER		
Scenario	Number of Ads	Operations/sec
3	2	386.93

Scenario 3 hosting more Directory servers had the following output: The average throughput for the File Server is 25.6 and each server had 60 users as the previous scenarios. The average Requests/Sec and Response time for the Web Server is 46 and

128 respectively. The average number of operations for the 2 AD is 386.9. The CPU threshold of 60% was crossed, but maintained within permissible limits with this mix of Servers.

The number of Directory servers in this scenario has been limited to 2 because each one uses up a major portion of the CPU resource (~30 %).

The overall result for the set of Heterogeneous experiments is summarized in the following graph:

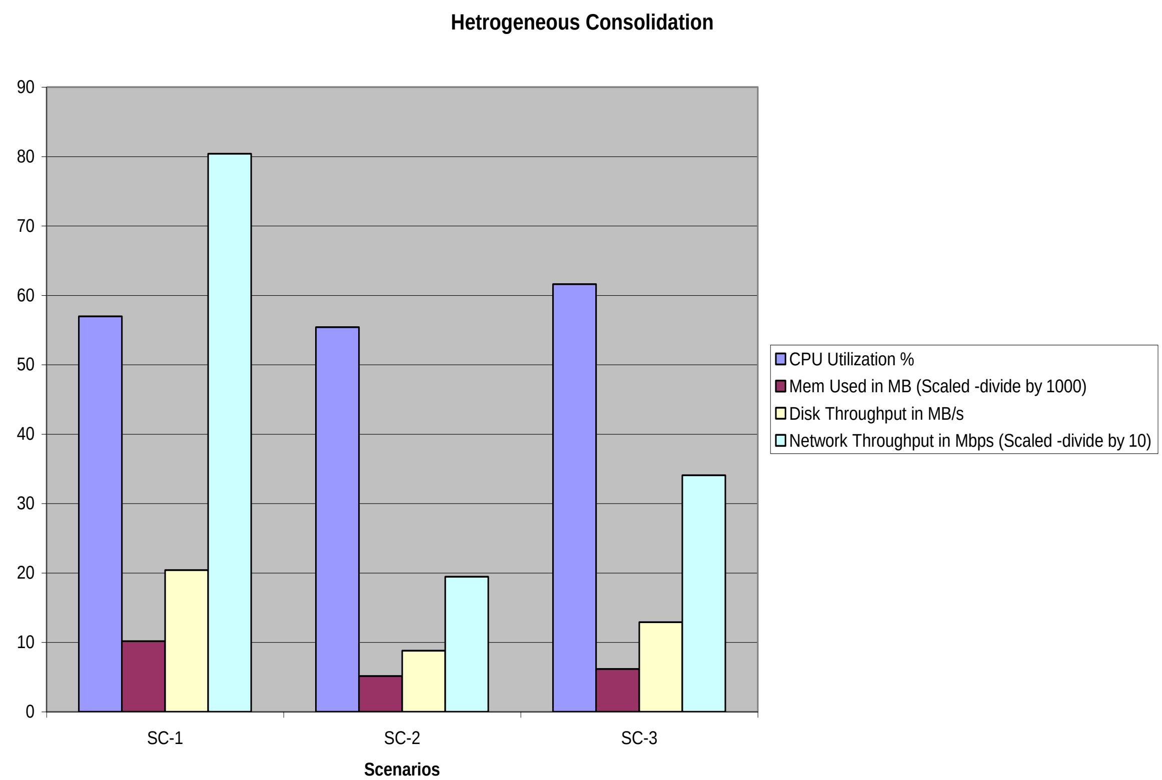

Figure 11 – Comparison of Scenarios 1, 2 and 3

Metrics	**Threshold Limit**	**Threshold reached**		
		Scenario		
		1	**2**	**3**
CPU	60	Yes	Yes	Yes
Memory	75	No	No	No
Disk	75	No	No	No

Conclusion:
The Heterogeneous consolidation graph shows the CPU, memory, disk and throughput for each of the three scenarios. The CPU threshold is reached for all the three scenarios while the memory and disk exit criteria is not reached. The contribution to CPU

utilization for similar numbers of Web, Directory and File Servers in heterogeneous setup is around the same as in the Homogeneous experiments implying a direct and undistorted contribution to resource use while translating from homogeneous to heterogeneous setup. This illustrates the scalability of Quad-Core Intel Xeon processor-based servers for heterogeneous server consolidation.

Stability Test

Stability tests were run to make sure that results obtained are reliable and stable. A scenario with Heterogeneous consolidation of File and Directory Servers was chosen. A mix of 8 File Servers and 1 Directory Servers were used to serve the desired load. Each File Server is accessed by 60 users and a medium load is sent to the Directory Server. The experiments ran for 75 hours over a period of 3 days to evaluate stability.

Logs were captured for the entire test run and the analysis shows that there were no irregularities or unexpected issues in the output. No performance degradation was observed over the period of time. The results of the server performance characteristics are as follows:

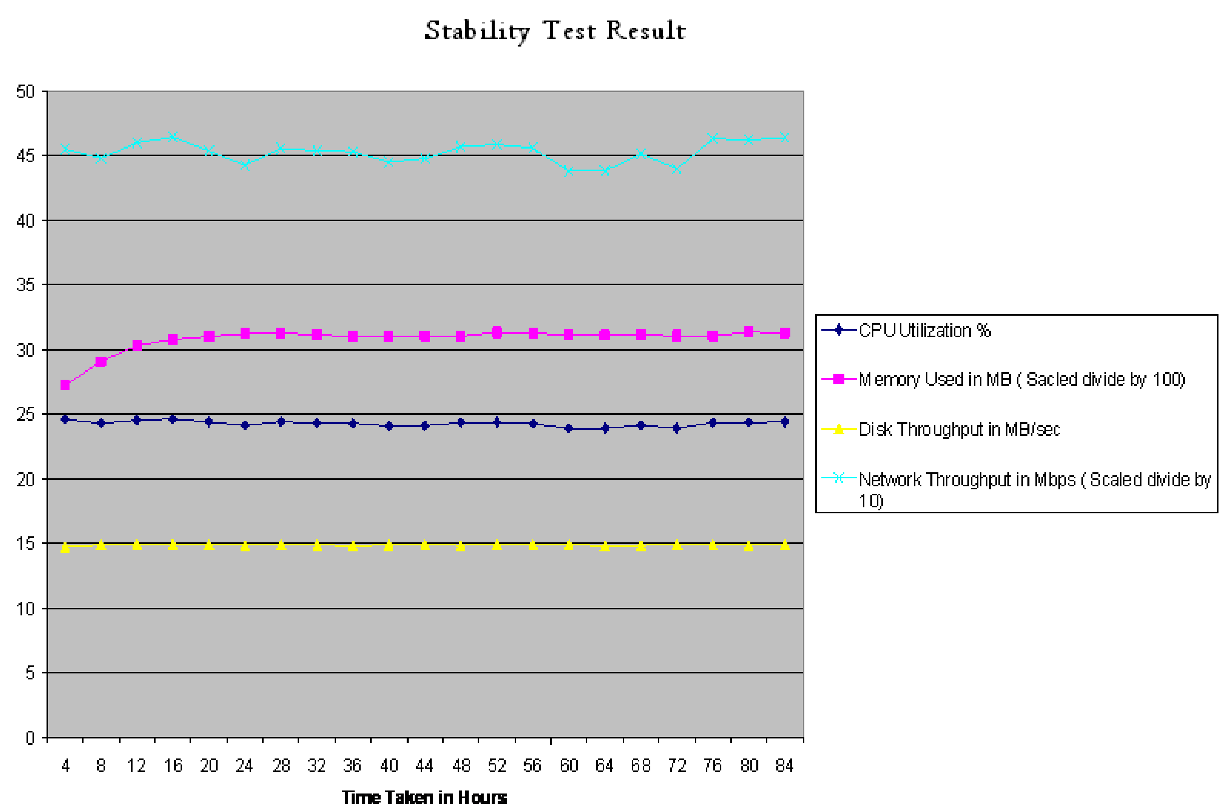

Figure 12 – Stability Test Characteristics

	Threshold	Threshold reached
CPU	65	No

Memory	75	No
Disk	75	No

Conclusion:
As observed in the graph, none of the thresholds were reached and the output is steady with no deviations from the expected result. This ensured that real life scenarios where a server may be stressed for a long duration can be supported in the given setup.

Recommendations

Listed here are a set of recommendations:

- Quad-Core Intel® Xeon® processors are found to be very suitable for homogeneous and heterogeneous server consolidation
- Infrastructure assessment needs to be undertaken at first to characterize the utilization of CPU, memory and network i/o resources for the servers and workloads to be virtualized and consolidated.
- File and print servers are recommended as the first level for server consolidation rather than mission critical servers.
- These test results can be used as a reference for planning datacenter Server Consolidation and the test methodology can be used to size the deployment based on actual workloads.
- Stability and stress testing has to be performed before deploying the server consolidation in production environment.
- Microsoft utility tools can be used to test the Web and Active Directory based Directory Servers.
- Heterogeneous server consolidation efficiently maximizes the server hardware utilization, as File server workloads are network intensive, and Web and Directory Servers are found to be CPU intensive. This guideline can be used to consolidate more application servers per physical server.
- Server resource utilization guidelines can be used for server consolidation capacity planning. Such guidelines can be based on performance measurements with the workloads to be consolidated with server virtualization technology.
- Use of a SAN is recommended to prevent disk bottleneck when numerous file servers are consolidated as opposed to using internal disks.
- It is highly recommend that each VM be mounted on separate LUN, if you have small number of VMs. This configuration is not practical for large number of VMs as overhead of management significantly will increase by increase in number of VMs.

TCO

TCO calculation is done based upon 9:1 consolidation. 9 physical servers were converted into 9 VM's on one physical server. Following is the TCO calculation for 3 years. We have assumed that there are 900 servers and we are consolidating them onto 100 servers.

The parameters considered are hardware, software, maintenance and utility cost. The following table depicts the actual cost savings in dollars if a 9:1 consolidation is undertaken. The savings (difference in dollars) is highlighted part of every step.

Note: $0 indicates that the same amount is required in a non-virtual and virtual environment and hence the cost is not considered.

		Non Virtualized Environment			Virtualized environment			Savings
		Numbers	Rate	Cost	Numbers	Rate	Cost	
Hardware	Servers[0]	900	$5,464.00	$4,917,600.00	100	$8,363.00	$836,300.00	**$4,081,300.00**
	HBA	1800	$1,800.00	$3,240,000.00	200	$1,800	$360,000.00	**$2,880,000.00**
	Racks	75	$5,000.00	$375,000.00	9	$5,000.00	$45,000.00	**$330,000.00**
	Fiber Cable	1800	$90.00	$162,000.00	400	$90.00	$36,000.00	**$126,000.00**
	SAN Switch	37.5	$9,500.00	$356,250.00	5	$9,500.00	$47,500.00	**$308,750.00**
	Network Switch	56.25	$2,499.00	$140,568.75	9	$2,499.00	$22,491.00	**$118,077.75**
	UPS Cost	100	$1,366.00	$136,600.00	44	$1,366	$60,104	**$76,496**
Utility	**Power**							
	Total watt used by server[1]			300			447	
	Total KWH			0.3			0.447	
	Total power used in 3 years			7884			11747.16	
	Typical KwH cost[2]			$0.098			$0.098	
	Total cost of running 1 server			$772.63			$1,151.22	
	Total cost of running all	900 server		$695,368.80	100 servers		**$115,122.17**	**$580,246.63**
	Cooling							
	Total number of Racks			75			9	
	Typical Square feet / square meters per rack[3]			7			7	
	Total SQ foot for racks			525			63	
	Typical cooling watt /sq foot[a]			80			80	
	Total watt for racks for AC			42000			5040	
	Total KWH			42			5.04	
	Typical KWH Cost			$0.098			$0.098	
	Total KWH for Racks			$4.12			$0.49	
	Total for 3 years			**$108,168.48**			**$12,980.22**	**$95,188.26**
	Real Estate[3]	40 USD/Sq Foot per year						
	Total SQ foot for racks			525			63	
	Total cost of real estate for 1 year			$21,000.00			$2,520.00	

	Total cost of real estate for 3 years			$63,000.00			$7,560.00	**$55,440.00**
Software								
				$0.00			0	**$0.00**
	Windows 2003 EE	No Change						**$0.00**
	VMware 3.0.1¯			$0.00	100	$5,750.00	$575,000.00	**($575,000.00)**
Total Savings								**$8,076,498.64**
Others	Business continuity	Loss of application/availability			Application can be migrated to another running server			
	Server buying/server/provisioning	typical wait time is 3 months			Can be done immediately			**3 months**
	Ease of new deployment	Typically 3-5 days			Couple of Hours			

ºMore expensive hardware compared to legacy systems.
Source of all hardware costs: www.hp.com, as of 30th January, 2008.

The cost quoted in the TCO calculation above is for the following hardware configuration:

Servers	HP* ProLiant DL380 series
HBA	HP StorageWorks FC1243 4Gb PCI-X 2.0 Dual Channel HBA
Racks & Rack accessories	HP 10642 G2 Crated Universal Racks
Fiber Cable	HP Fiber Optic cable multi mode type
SAN Switch	HP Brocade 4/24 SAN Switch
Network Switch	ProCurve Switch xl 16-port 10/100/1000 Module
UPS	HP R3000 2U L5-30 NA UPS (2700W)

[1]Average Power consumption in watt by the platform:

Non Virtual environment approximately 300W

Intel®5000 chipset	29W
Intel® Xeon® processor	80W (2 Sockets)
FBD DIMMs	44W (11W/DIMM)
I/O Bridge	13W
Other devices	54W

Virtual environment approximately 447W

Intel®5000 chipset	29W
Quad-Core Intel® Xeon® processor	120W (2 Sockets)
FBD DIMMs	88W (11W/DIMM)
I/O Bridge	13W
Other devices	77W

Source: Intel, November 28, 2007. Intel specifications for CPU, chipset, LAN component power (TDP). Memory power represents typical DIMM power (as of October 07) based on survey of leading DIMM providers. Memory power may vary depending on vendor and memory speed.

Total Power used in 3 years by one server is calculated as follows. If 0.3KWH is the Power consumption of the server per hour, then the total power consumed for 3 years is 3years*365 days*24 hours*0.3 KWH = 7884KWH.

[2]Typical Power KWH Cost is taken as $0.098

Default Energy Cost: 9.8 cents/KWh, which is the U.S. average for commercial customers as of September, 2007. Source: U.S. Department of Energy, http://www.eia.doe.gov/cneaf/electricity/epm/table5_6_a.html#_ftnref1

[3] Datacenter rack space and floor space cost was estimated using VMware TCO/ROI calculator. http://www.vmware.com/files/pdf/tco_roi_methdology.pdf as of 3rd February 2008.

[a]Typical cooling watt /sq foot is taken as 80.

http://www.abrconsulting.com/Custom_Code_Pages/calc4.php as of 3rd February 2008.

¯VMware VI3 infrastructure cost is $5750 for VI enterprise edition dated February 1, 2008.

By using virtualization technology for server consolidation productivity is greatly improved and man power utilized is reduced due to the reduction in the number of physical servers.

Conclusion

The PoC undertaken by Infosys at its SI labs along with support from Intel Corporation demonstrates the benefits of server consolidation using VMware ESX Server VI 3.0.1, Net App SAN FAS3040 and Quad-Core Intel® Xeon® processor-based servers, by consolidating several applications on a single physical server, thereby significantly reducing the HW, SW and Maintenance costs. Furthermore, we showed that the user expectation of performance and quality of service can be met in a virtual environment using Intel's quad-core servers. In a typical example we showed a saving of nearly $8M over a period of 3 years when 900 physical servers are reduced by a 9:1 ratio.

Acronyms

PoC	Proof of Concept
Intel® VT	Intel Virtualization Technology
AD	Active Directory

TCO	Total Cost of Ownership
LUN	Logical Unit Number
SAN	Storage Area Network
VM	Virtual Machine
Intel EPSD	Enterprise Platforms & Solutions Division
FSB	Front Side Bus
RAID	Redundant Array of Independent Drives
PCIe	Peripheral Component Interconnect – Express
NIC	Network Interface Card
HBA	Host Bus Adapter
DRS	Distributed Resource Scheduler
HA	High Availability
TDP	Thermal Design Power

Authors

Hema Magesh
Senior System Engineer in Digital Enterprise Group/End-User Platform Integration group

Systems Integration Group
Infosys Technologies

Reviewers from Intel Corporation

Naresh Sehgal, Lead Software Architect in Digital Enterprise Group/Enterprise Platforms and Services Division
Parviz Peiravi, Principal Architect in Software Solutions Group/ Enterprise Solution Sales
Rob Sullivan, Lead Architect Virtualization in Digital Enterprise Group/End-User Platform Integration group

INFOSYS TECHNOLOGIES	INTEL CORPORATION
Infosys Technologies Ltd (NASDAQ: INFY) defines, designs and delivers IT-enabled business solutions that help Global 2000 companies win in a flat world. These solutions focus on providing strategic differentiation and operational superiority to clients. Infosys creates these solutions for its	Intel, the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live.

clients by leveraging its domain and business expertise along with a complete range of services. With Infosys, clients are assured of a transparent business partner, world-class processes, speed of execution and the power to stretch their IT budget by leveraging the Global Delivery Model that Infosys pioneered.	
More information about Infosys Technologies is available at www.infosys.com.	For more information visit www.intel.com/technology/platform-technology/virtualization/

Infosys®
POWERED BY INTELLECT
DRIVEN BY VALUES

intel®