Server Virtualization and Consolidation - A case study

Ravi G Singh
Consulting IT Specialist – System p
rsingh@us.ibm.com

Special Notices

This document discusses Server Virtualization and Consolidation designed and implemented by the author for an IBM Customer. This document is presented "As-Is" without any warranty, guarantee or assurances of any kind, either express or implied. IBM does not assume responsibility for the statements expressed herein and it reflects the opinions of the author. If you have questions about the contents of this document, please direct them to the author at rsingh@us.ibm.com.

Author is not responsible for errors in this document that may result in any kind of inaccuracies.

Acknowledgements

Thanks to John R Hock, IBM Certified IT Specialist – System p - Advanced Technical Support Americas (ATS) for reviewing this White Paper.

Thanks to the customer and IBM team for their contribution and support to this project.

Trademarks

Copyright

URLs

This document can be found on the web under the category of Whitepapers in IBM TechDocs Library.

Internet http://www.ibm.com/support/techdocs
Intranet http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

Release Date

Apr 2007

Table of Contents

Abstract

This Whitepaper discusses a case study of server consolidation and virtualization project using the Advanced POWER Virtualization (APV) feature of IBM System p Server for a Business Intelligence application environment of a large enterprise in Industrial Sector. The discussion is focused on business requirement, server architecture, APV features used, capacity planning and resource planning, methodology and tools, and benefits to the customer.

The scope of this project was to enable a technology refresh of old hardware, Virtualize and consolidate twenty seven stand alone servers, upgrade the OS and middleware to the supported levels, improve utilization and availability of hardware, and adopt a industry standard network infrastructure to communicate and share data between servers. This paper highlights how the Server Virtualization enabled the customer to reduce the TCO by optimizing the utilization of computing resources.

Introduction

Server Virtualization has emerged as one of the top IT objectives for corporate and IT Management. The major factors driving Server Virtualization as one of the top IT goals are

- Consolidation of servers
- Real time response to changing business needs
- OnDemand computing model
- Sharing of computing resources to improve utilization
- Improving Flexibility, High Availability and Performance
- Reducing TCO

IBM's Advanced POWER Virtualization (APV) feature of System p Server and AIX provide Server Virtualization capability to run multiple workloads in a shared and/or dedicated resource environment. APV can be used very effectively to balance resource utilization between LPARs on a server running commercial and engineering applications such as business critical, ERP, product development, Web/App/Database, datawarehouse and BI applications. Businesses are migrating standalone servers running these types of applications to LPAR environment using APV on IBM System p enterprise servers. Benefits of APV and Server Virtualization are many-fold and presented in the diagram below. Synergy of these benefits leads to a significant business advantage in IT operations.

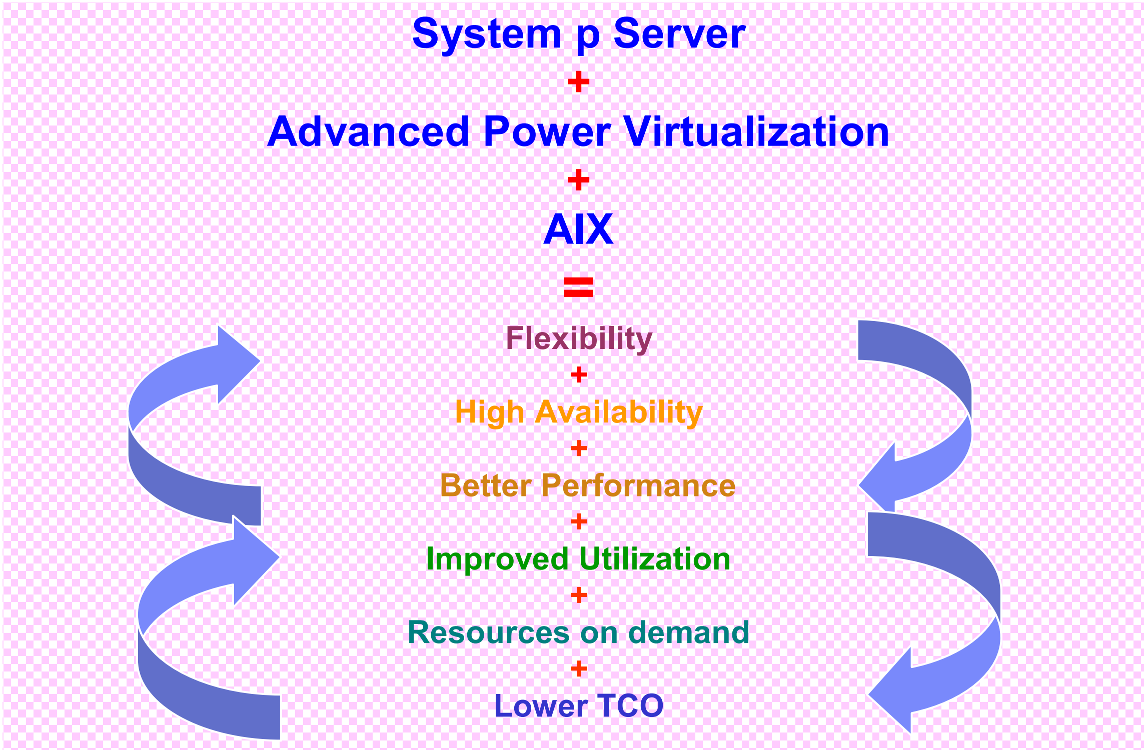

This paper discusses the architecture, design, and methodology of Server Virtualization for a multinational organization in industrial sector. The virtualization architecture deployed highlights APV features used and its Best Practices.

Background

Three Business Intelligence (BI) application environments using IBM UNIX servers provided critical warranty and quality information to customer's management and plants. These business critical applications ran in a complex environment involving multiple servers, applications, middleware stack, dependency on other servers for data and applications, network and SAN connectivity, and infrastructure requirements. These BI applications with multiple TB of data running on clustered IBM pSeries servers were implemented in a Massive Parallel Processing (MPP) shared nothing architecture for DB2 queries. The parallel queries running on these servers used a private high speed network. The workload of these applications involved running complex parallel queries on databases using both real-time and historical data, and could take anywhere from hours to days to complete. These datawarehouse and datamart environments exchanged and updated information between them.

The IBM UNIX servers deployed for this project went through several technology refreshes in the past. The first two generation of servers were RS/6000 Scalable Parallel (SP) nodes with SP Switch network and the next generation consisted of POWER4 pseries servers with a SP2 switch network. The customer preferred the SP clustering technology with a high speed SP private network due to the excellent application performance. The discussion in this paper is related to the implementation of the latest technology refresh of these servers using virtualization.

Existing configuration

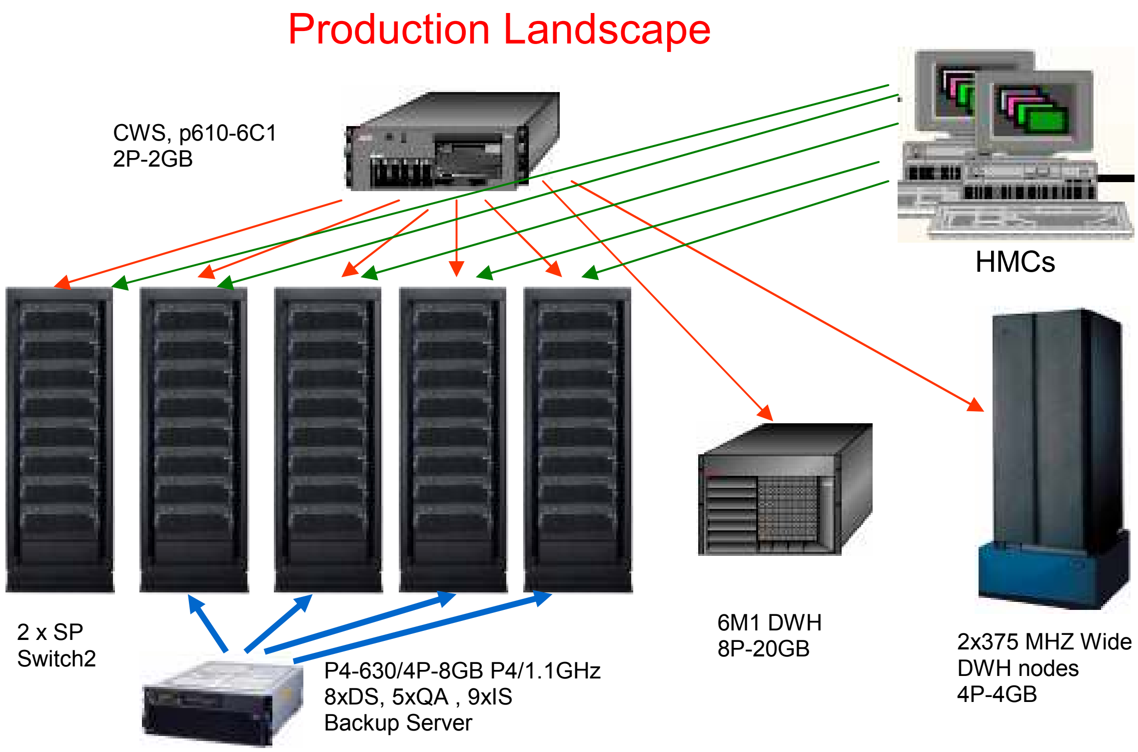

Production

- Total Processors-Memory 112P-222GB
- 22 x P630s (4P-8GB)
- 7026-6M1 (8P-20GB)
- 2 x 375 MHz wide node (4P-4GB)
- Backup Node (2P-4GB)
- 7028-6C1 Control Work Station (CWS) (2P-2GB)
- 2x7315-C01 HMC
- 2x9076-558 SP Switch2

Development

- Total Processors-Memory 16P-16GB
- 4 x 375 MHz wide node (4P-4GB)
- 7028-6C1 CWS (2P-2GB)
- 1x9076-558 SP Switch

Software Stack

- Operating System – AIX 5.2
- Database - UDB 8.2 (8.1FP11)
- Websphere – 5.0
- IHS (IBM HTTP Server) – 2.0.47.1
- Cognos – 7.3

- JDK – 1.4.2 & 1.3.1
- Perl – 5.8.2.30
- SiteMinder – 5.5QMR7
- Crystal Reports

During the technology refresh of these servers, several server solution alternatives were considered before selecting the virtual architecture. The following sections discuss the solutions architected and presented, architecture of the virtualization solution, tools used, and post-implementation benefits.

Discussion of APV functionality and features, application requirements and implementation/customization details are beyond the scope of this White Paper. The APV architecture designed and implemented for this project is documented as best practices in the IBM Redbooks and papers mentioned in the References Section.

Methodology

This technology refresh project was architected and built using a three phase methodology. Output from the preceding phase was used as the input in the next phase along with several other factors discussed below.

1. Capacity Planning: During this phase, replacement servers/LPARs were sized using a customized tool. Factors such as current utilization, desired target utilization, growth factor and spare pool were used to size each LPAR. APV best practices and recommendations, and customer configuration standards and practices were used in this sizing.

2. Resource Planning: In this phase, standards and guidelines were built for resource allocation based on the existing practices, type of LPAR (Web, App. Or Database), applications to be run, bandwidth requirement and APV best practices. The resources allocated to each LPAR were either dedicated or shared depending on these guidelines.

   LPAR sizing information from capacity planning was used in this phase. Processor, memory, boot disks and adapter requirements finalized in this phase were used in eConfig for server configuration.

3. Resource Management: Based on the Resource Planning, adapters, disks and HBAs were assigned to LPARs. Multiple adapters of same type required for a LPAR were assigned from multiple RIO units to eliminate RIO draw as a single point of failure. The information from this tool was used by Server Architects and System Administrators during LPAR creation and then later for LPAR management.

Capacity Planning and Technology refresh alternatives

The customer's corporate standard performance monitoring tool was used to collect utilization data from all SP nodes and servers in the SP cluster. Processor utilization and projected workload growth during the next lease period for each server was used to compute and size replacement servers/LPARs and multiple server solutions were proposed. Customized scripts were used to collect network (intranet and SP2 Switch) and SAN bandwidth requirements.

Two alternatives proposed to replace SP Switch2 network were 10 Gb/sec Ethernet and the Virtual Ethernet network feature of System p server to provide the required bandwidth for high speed data transfer between the LPARs to run parallel queries.

Four server architecture solutions proposed for production environment were:

Option 1 - Consolidation and virtualization using p5-595

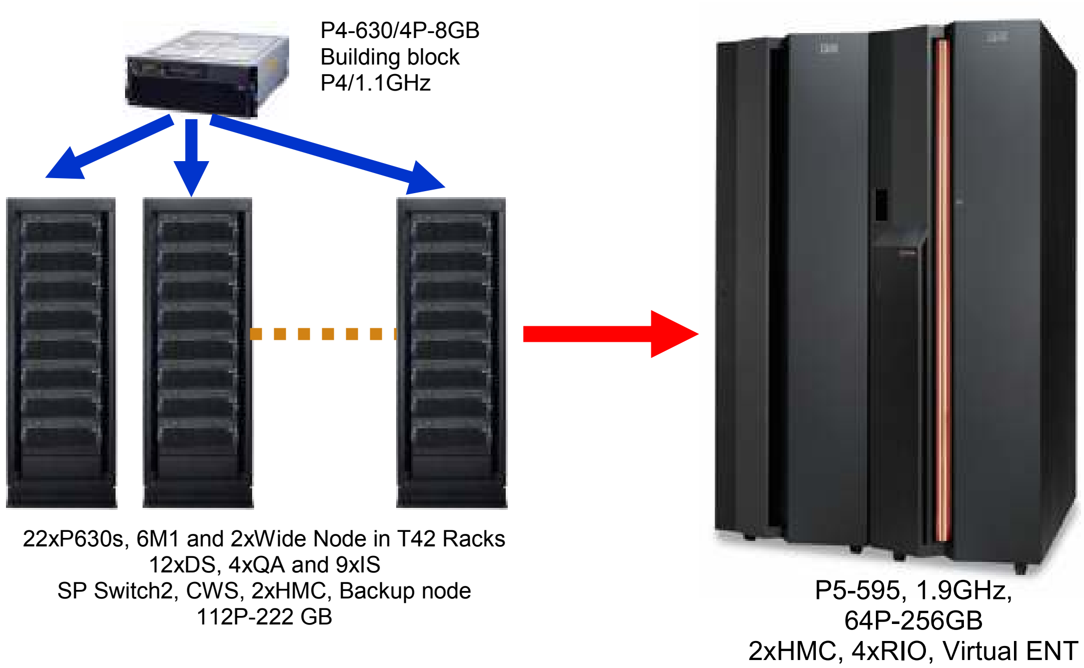

Option 2 – One for one replacement using entry server

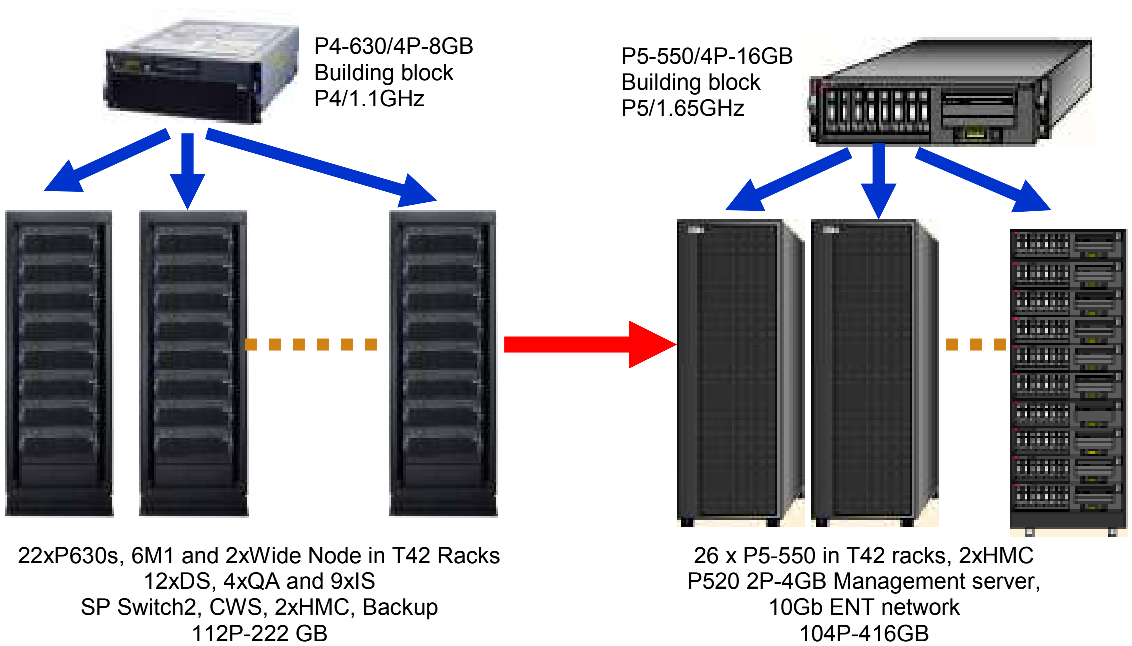

Option 3 – Consolidation and virtualization using midrange server

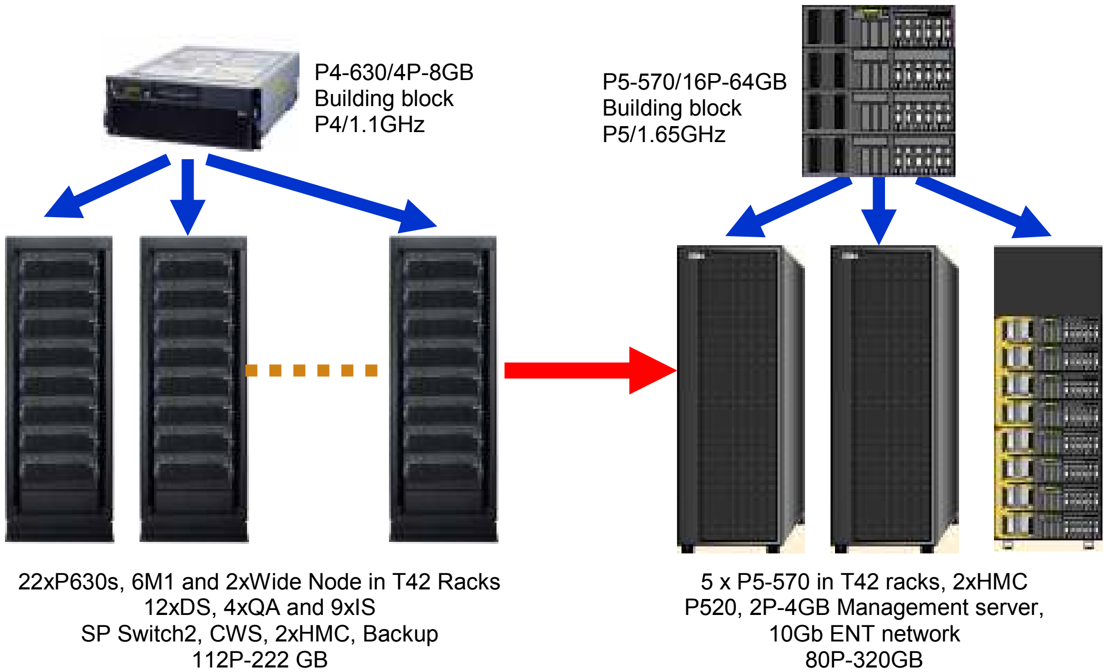

Option 4 – Consolidation and virtualization using BI server

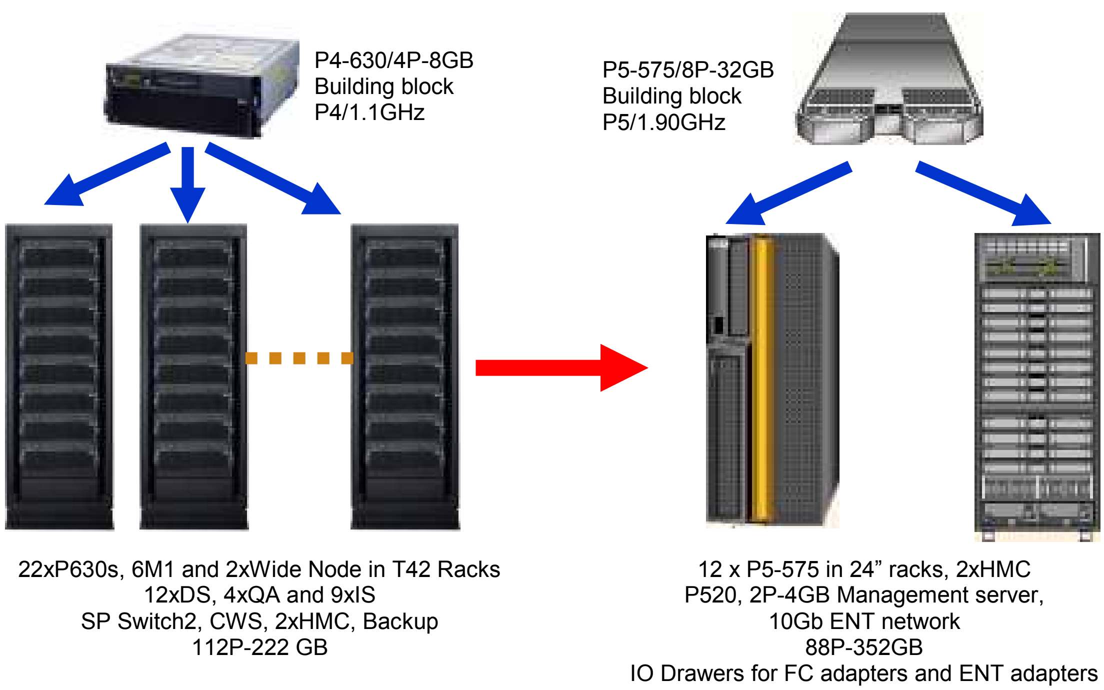

Virtualized Server Architecture

Out of the four server solution alternatives architected and presented, the customer selected the p5-595 architecture to virtualize and consolidate servers running MPP datawarehouse applications. Virtualization and consolidation on the p5-595 server was preferred due to the flexibility in sharing resources, higher availability, and higher utilization allowing server provisioning for new workload. One of the key decision factors favoring this architecture was the opportunity to reduce the overall TCO and not just Total Cost of Acquisition (TCA) of the server. These are discussed in the Results section of this White Paper.

This flexible server architecture was built using APV Best Practices to enable resource sharing for optimum utilization and a highly available environment to eliminate single points of hardware failure.

The diagrams below show LPARs with the recommended processor and memory configuration based on the utilization and projected workload for the production and development environments.

Consolidation and virtualization using p5-595 - Production

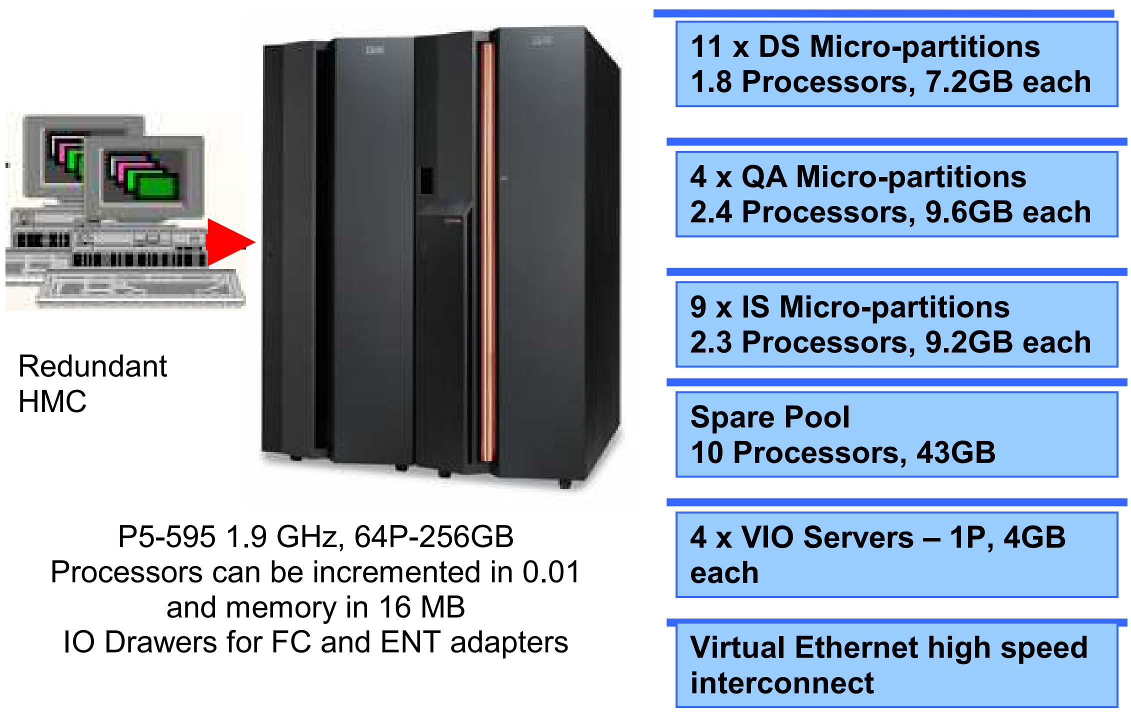

Six LPARs for test and development environment of these applications were consolidated with other workload onto a second p5-595 server as shown in the diagram below. This provided a development environment identical to the production environment using the same APV features such as micro-partitioning, dual VIO Servers, shared boot disks, Ethernet and SAN adapters, spare processor pool, and private high speed Virtual Ethernet network to run parallel queries and share data between LPARs.

Consolidation and virtualization using p5-595 - Development

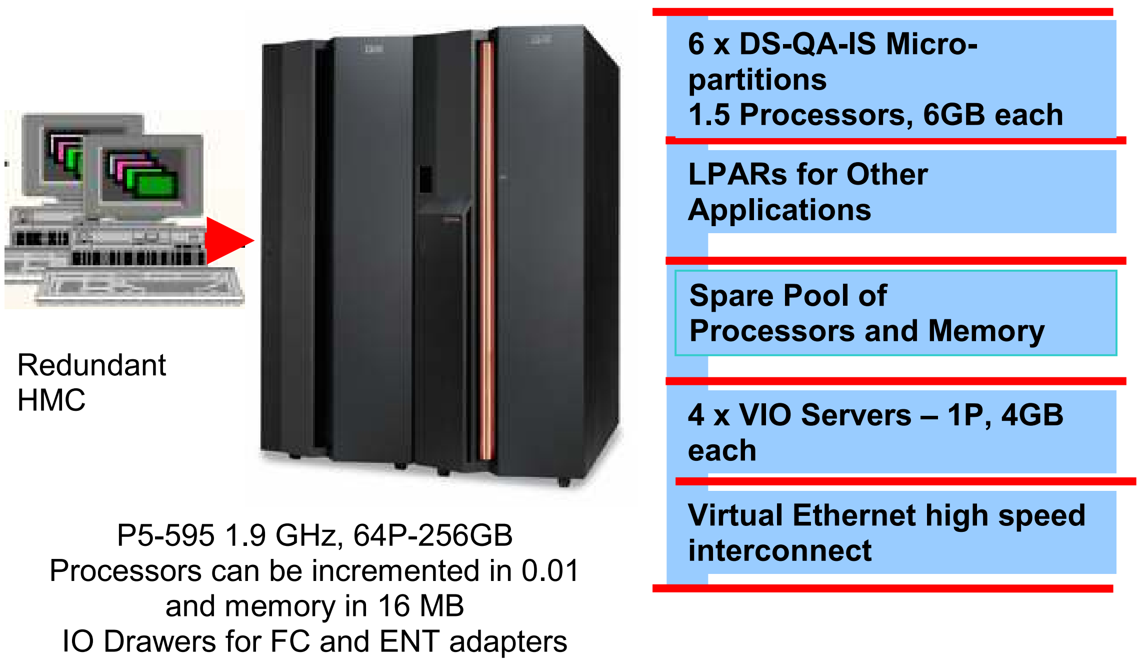

APV and High Availability features

APV and AIX offer several features and options to design a flexible server with high availability features to eliminate single points of failure. The APV features used in this server solution architecture are listed below.

- Micro-partitioning – Processing units are allocated in 0.1 increments, and can be added and removed using DLPAR without a reboot.
- Max. processor utilization - Virtual Processors are used to limit the maximum utilization of processing units.
- Shared processor pool - Uncapped LPARs can dynamically add more processing units from the pool On Demand for the changing workload.
- Memory – Can be added and removed from the LPARs without a reboot.
- Dual Virtual I/O Servers - Enables sharing of resources with dual paths to LPARs for SAN and network connectivity.

- Shared Boot Disks – 146GB Ultra 320 SCSI disks are shared by four LPARs. Each LPAR gets three VSCSI disks for OS: boot, mirror and alt.
- Shared Ethernet Adapters – Two adapters for intranet network and two adapters for Backup network are shared by multiple LPARs. Network Interface Backup (NIB) is configured at VIOS and client LPARs for high availability.
- Shared SAN Adapters – Each pair of VIOS using EMC Powerpath drivers present SAN disks as VSCSI disks to client LPARs for path availability and load balancing.
- Virtual Ethernet – An internal high speed private network connects LPARs to share data.

- Multiple OS levels - LPARs can run either AIX 5.2 or AIX 5.3 level of the operating system.
- Security - 100% safe, secure and isolated LPARs.
- High Availability - Single Points of Failures (SPOF) for boot disks, Ethernet and SAN adapters eliminated.

LPAR Types

APV offers options to create shared and/or dedicated LPARs using shared and/or dedicated resources. Both the production and development p5-595 servers were designed and configured for the following type of LPARs.

- **VIO Servers –** Dual LPARs to enable sharing of I/O resources.
- **Database LPARs –** LPARs with shared resources for parallel database.
- **Web and App. LPARs –** LPARs with dedicated resources to run web server and applications.
- **NIM LPAR –** Management and administration LPAR with dedicated I/O resources to install and manage AIX on client LPARs, to run scheduled DLPAR operations for other LPARs, and to collect performance data from all the LPARs.

Virtualization of I/O Resources

Two pairs of VIO servers were configured and built to host groups of LPARs serving the three BI applications.

- VIO Server 1 and 2 hosting the first BI application using nine client LPARs.
- VIO Server 3 and 4 hosting the second and third BI application running on twelve client LPARs.
- Physical I/O resources (adapters and disks) were allocated to VIO Servers.
- Each pair of VIOS presented the following shared resources to client LPARs
    - Shared boot disks
    - Shared adapters for intranet
    - Shared adapters for Backup network
    - SAN disks with multiple paths
- High speed private Virtual Ethernet network was configured for LPARs to communicate with each other.

Shared Processor pool and Capped/Uncapped LPARs
The shared processor pool was configured to enable LPARs to dynamically add more processing units in two situations: fluctuating and peak workload demands, and provision for the estimated growth in workload. This requirement was sized during the capacity planning phase. All the AIX 5.3 LPARs were uncapped micro-partitions and capable of allocating additional processing units beyond the entitled capacity as and when required from the shared pool. The Power Hypervisor released the additional processing units back to the shared processor pool when no longer required by the LPARs.

Three AIX 5.2 LPARs were built with dedicated resources (processors, memory and I/O devices) due to application requirements.

Virtual Processors (VP) and Weightage
The maximum processing units that could be used by the uncapped LPARs was controlled by the Virtual Processors. The desired Virtual Processors (VP) was set to twice the quantity of the desired Processing units, rounded-off to the nearest integer value.

Users ran complex interactive and batch database queries and these could take anywhere from hours to days to complete. To accommodate these spikes and fluctuating workload, additional processor capacity beyond the entitled capacity for each LPAR was provided from the shared pool as and when required. Virtual Processors feature of APV allocated physical processors for multi-threading and additional capacity to LPARs on demand. The priority for allocating additional processing units to LPARs was controlled by the Weightage factor and its assignment to LPARs was as follows.

- VIOS – 255
- Database LPARs – 196
- Web and Application LPARs - 128
- NIM LPAR - 64

Virtual SCSI
The SCSI disks from dual VIO Servers were presented as Virtual SCSI boot disks to AIX 5.3 client LPARs. Two 36GB Logical Volumes presented from the first VIOS to each LPAR were used for rootvg and alt_rootvg. One 36GB Logical Volume presented from the second VIOS was used to mirror rootvg. This eliminated a LPAR failure either due to a SCSI disk failure on a VIOS &/or a VIOS failure.

Virtualization of Boot Disks

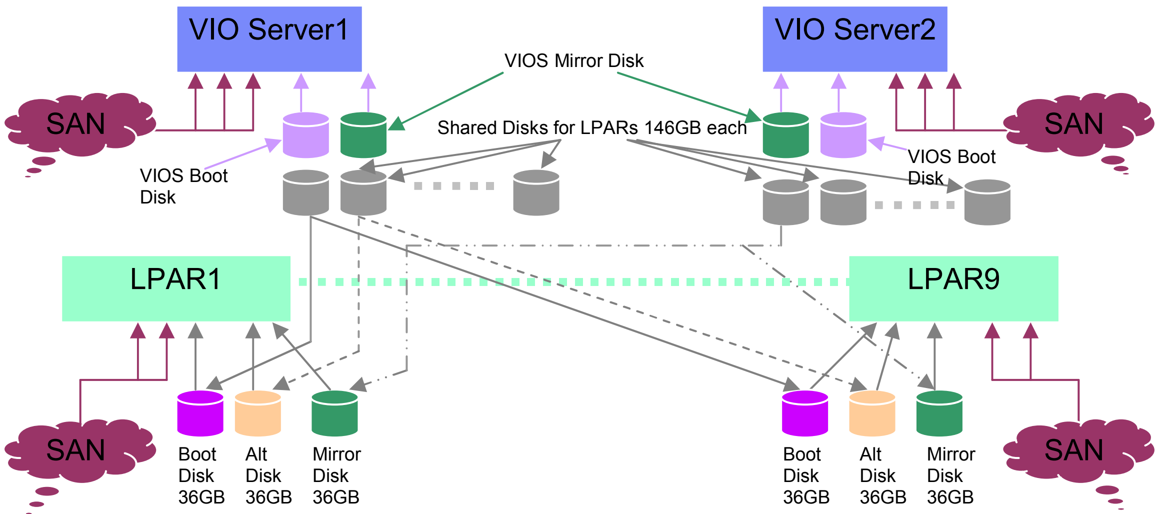

- 146GB Disk from a VIO Server allocated to 4 LPARs, each 36GB in size.
- Boot disk failure and/or a VIO Server failure will not crash a LPAR

Two 73GB disks were allocated to each VIO Server and NIM LPAR for boot and mirror and three 73GB disks were assigned to AIX 5.2 LPARs for boot, mirror and alt disks. alt_disk_install clones rootvg and altinst_rootvg disk can be used to boot when the OS on the boot and mirror disk is corrupted or when both the boot and mirror disks fail.

Virtual Ethernet
Four options available to replace the SP2 high speed private network for LPARs to communicate and share UDB EEE data with each other were

- 10Gb network
- High Performance Switch (HPS)
- Infiniband (IB)
- Virtual Ethernet

The 10Gb and IB options required a new network infrastructure to be built while the HPS option required redesigning the database architecture from MPP to dB logical partitions on a single large SMP server. Since all the LPARs were configured and built on one p5-595, Virtual Ethernet was the best alternative and it eliminated the need for building a 10Gb Ethernet or HPS or IB network.

Virtual Ethernet adapters were created for all LPARs and a private non-routable 192.x.y.z network was configured.

Shared Ethernet Adapters (SEA)
Two network connections required to all LPARs were: corporate intranet for user access and Backup network to backup the data. Dual VIO Servers serving a group of LPARs were assigned either a two single port or one dual port Ethernet adapter.

These adapters/ports were configured as the primary and standby adapter using Network Interface Backup (NIB) feature on each VIOS. Shared Ethernet Adapters (SEA) were created to all the LPARs using this NIB adapter from each VIOS. Two SEAs for each network on the client LPARs presented from dual VIOS were configured in NIB mode. This configuration provided high availability to both network connections and eliminated network failures due to an adapter, cable, switch port and/or VIOS failure.

The network traffic from the client LPARs was distributed between the two VIO Servers by changing the primary adapter in NIB configuration. One half of the LPARs used the SEA from VIOS1 as the primary adapter and the second half of LPARs used the SEA from VIOS2 as the primary adapter thereby distributing the network traffic between a pair of VIOS.

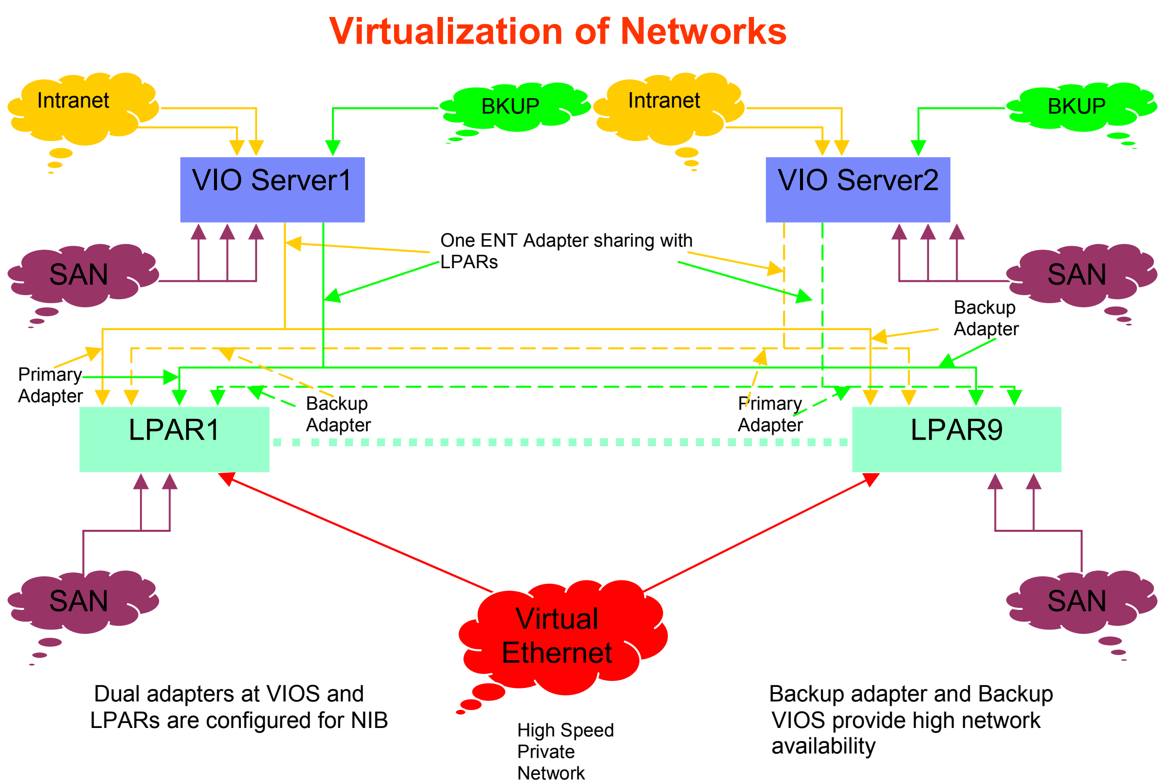

Dedicated SAN adapters
Two dedicated HBAs were allocated to database LPARs for SAN connectivity and to provide the required I/O bandwidth. HBAs were also dedicated to AIX 5.2 LPARs running Web and App. Servers as AIX 5.2 does not support sharing of resources.

Shared SAN adapters
SAN disks for data storage were configured and presented from dual VIOS to LPARs using shared HBAs. Three HBAs were assigned to each VIOS and the same SAN disks or LUNs were presented to a pair of VIOS. EMC Powerpath was installed on all the VIO Servers to provide multiple paths to SAN and balance the I/O between the HBAs. Dual VIOS presented the same power disks as VSCSI disks to client LPARs. The MPIO

feature of AIX on client LPARs configured the same LUN presented from dual VIOS as a single disk.

The SAN traffic was distributed between a pair of VIOS by changing the primary or active path on client LPARs. The SAN adapters on VIOS1 provided primary path to one half of the LPARs and HBAs on VIOS2 provided primary path to the second half of the LPARs. Sharing of HBAs on both VIOS enhanced the availability and increased the bandwidth by effectively using all the six paths. This configuration on client LPARs made the SAN disks highly available and eliminated path failures due to a HBA, SAN cable, switch port and/or VIOS failure.

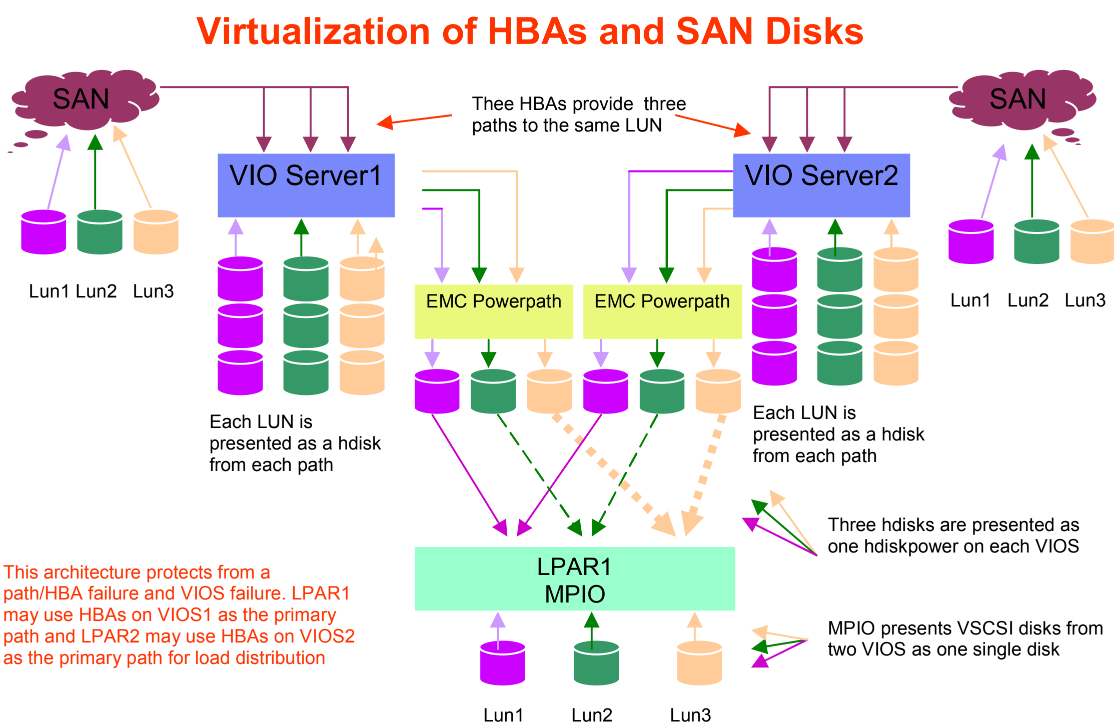

Resource Planning and Management tools

1. Resource Planning: Resource allocation standards and guidelines for LPARs were built based on the existing practices, type of LPAR, applications to be run, bandwidth requirement and APV Best Practices. I/O resources allocated to each LPAR were either dedicated or shared depending on these guidelines. Typically a Web and App. LPAR needed large network bandwidth while a database LPAR required large bandwidth to SAN. The standards defined were:

    a. Web and App. LPARs: Dedicated Ethernet adapter for intranet and shared HBAs using a pair of VIOS.
    b. Database LPARs: Dedicated HBAs for SAN connectivity and shared Ethernet adapters for intranet via a pair of VIOS.
    c. Backup network: LPARs to use shared network adapters presented from a pair of VIOS to connect to Backup network.
    d. Private Gb network: To share data between LPARs and other servers in the same datacenters and transfer mainframe data to LPARs, shared ethernet adapters presented from a pair of VIOS to LPARs provided connectivity to a private Gb network.
    e. Multiple DMZs: Best Practices and guidelines were built to connect LPARs to multiple DMZs on a single p5 server.
    f. Boot disks: Two 73 Gb disks were allocated to each VIOS and the NIM LPAR, three 73GB disks were allocated to AIX 5.2 LPARs and three 36GB LVs created from multiple 146GB disks were allocated to each AIX 5.3 LPAR. The total no. of disks required were consolidated and distributed between disk packs. As each disk pack with four slots could be allocated to only one LPAR, disks packs were populated with the required no. of disks.
    g. Dual loops were used to connect each RIO draw to the p5-595 server CEC.

    The sizing information for LPARs from the capacity planning was used in the Resource Planning. Processor, memory, boot disk and adapter requirements from the Resource Planning step were used in eConfig to configure the p5-595 server with dual HMCs.

2. Resource Management: Based on the resources allocated in Resource Planning, adapters, disks and HBAs were assigned to LPARs. Multiple adapters of same type required for a LPAR were assigned from multiple RIO units to eliminate RIO draw as a single point of failure. This information was used by the Server Architects and System Administrators during LPAR creation and then later for LPAR management.

Sample output from these tools is given in the Appendix.

Server Virtualization and Consolidation results

The desired Server Virtualization benefits are very distinct and different at each management level in a corporation. This section highlights the results and benefits to the Senior Management, IT Management and Business Unit after migrating stand-alone servers to a virtualized enterprise p5-595 server.

1. Management perspective

- ✔ Optimum utilization of computing resources.
- ✔ Reduced cost (TCO).
- ✔ Simplification of server operations and management.
- ✔ Standardized Infrastructure.
- ✔ Improved business agility/flexibility.
- ✔ Real time response to dynamically changing computing needs.
- ✔ Product differentiation and competitive edge in the market place due to innovation and collaboration in IT Infrastructure.

2. IT perspective

- ✔ Flexible to allocate resources – On Demand allocation of processors and memory to accommodate changing workload.
- ✔ Consolidation - No. of servers reduced from 27 to one.
- ✔ Sharing of resources - No. of network adapters and drops reduced from 54 to 12. No. of SAN adapters and drops reduced.
- ✔ High Speed network - Virtual Ethernet eliminated the need to build a new high speed private network (HPS or 10Gb Ethernet).
- ✔ New workload – Unutilized processor and memory resource could be used for additional &/or new workload to further reduce TCO.
- ✔ High Reliability and Availability – Mainframe like features and redundancy provides better server reliability.
- ✔ Better utilization of resources due to sharing.

2.1. Improved Resource utilization

- A customized script was used to monitor the CPU and Memory utilization on all the LPARs and p5-595 server.
- DLPAR was used to change the Processor and Memory allocation to LPARs.
- Server Processor utilization was increased from 25% to 65%.

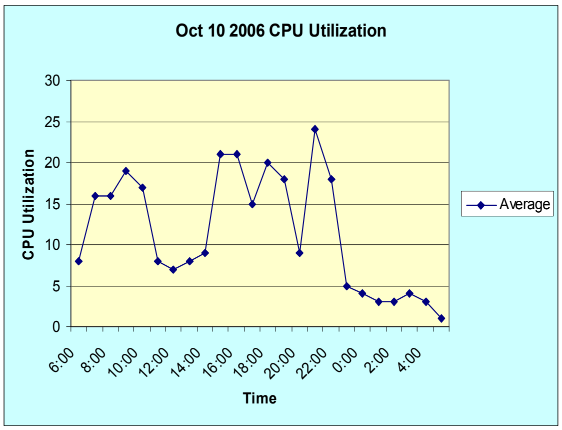

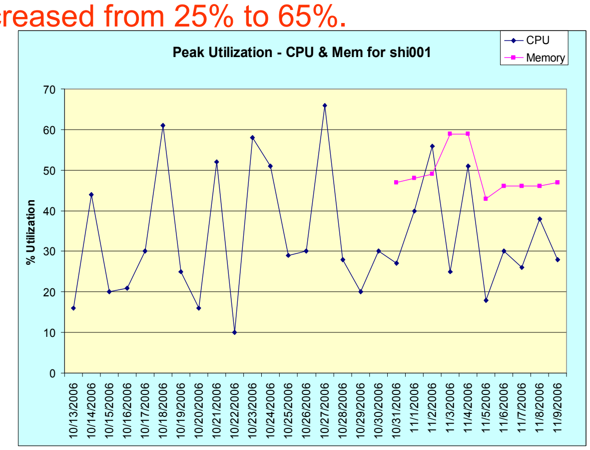

2.2. Flexibility in Resource Allocation

Dynamic changes to Resource Allocation

11 x DS Micro-partitions 1.8 Processors, 7.2GB each	**11 x DS Micro-partitions 1.0 Processors, 7.2GB each**
4 x QA Micro-partitions 2.4 Processors, 9.6GB each	**4 x QA Micro-partitions 1.0 Processors, 9.6GB each**
9 x IS Micro-partitions 2.3 Processors, 9.2GB each	**9 x IS Micro-partitions 1.3 Processors, 9.2GB each**
Spare Pool 10 Processors, 43GB	**Spare Pool 36.3 Processors, 43GB**
4 x VIO Servers 1P, 4GB each	**4 x VIO Servers 0.5P, 4GB each**
Total Allocated Processors 54	**Total Allocated Processors 28.7**

The performance data from all the LPARs was collected and utilization graphs for LPAR and server were plotted using the scripts developed by the customer. The average processor utilization for all the LPARs was increased from 25% to 65% by reducing the desired processing units from each LPAR using DLPAR operation. As all the AIX 5.3 LPARs were uncapped and configured with adequate VPs to use additional processing units as and when required from the spare pool, there was no performance impact on the applications. This resulted in a bigger size spare pool and the no. of unused processors increased from 10 to 36.

2.3. Additional &/or new Workload

- ✔ 26 additional spare processors are available in the shared processor pool.
- ✔ New workload can be deployed on the production server without procuring additional computing resources.
- ✔ This flexible environment enables provisioning of computing resources.

3. Business Unit Perspective

- ✔ Application and database Queries run faster in the virtual environment.
- ✔ Database Backup from all LPARs takes less than half the time compared to the old environment.
- ✔ Benchmark tests show performance improvement ranging from 40% to 70%.
- ✔ New environment has better reliability and stability due to built-in high availability characteristics.
- ✔ Virtual Ethernet is faster and eliminates hardware failures compared to the SP Switch2 private network in the old environment.
- ✔ LPARs can scale-up to accommodate the workload growth.
- ✔ LPARs provide a safe and secure environment as per the Business Unit requirements.
- ✔ AIX and Middleware stack was upgraded to supported levels and eliminated unsupported PSSP software, SP2 Switch and adapters.
- ✔ Reduced Software licensing cost (IHS, WAS, UDB, Cognos, SiteMinder) due to smaller no. of processors allocated to each LPAR.

Conclusions

Virtualizing and consolidating twenty seven stand alone servers to one single enterprise server provided a standardized server architecture, reliable, flexible and cost effective server solution to the customer. This implementation reduced the overall TCO due to cumulative savings in the infrastructure (power, cooling, and floor space), reduced network and SAN connections, reduced boot storage, Virtual Ethernet eliminating the need to build a new high speed private network, reduced Software Licensing, reduced server maintenance cost, and provisioning of computing resources due to higher utilization. Provisioning of processor, memory and shared I/O resources helped to quickly deploy a new workload on the same server by reducing the time required to plan and provide the datacenter infrastructure.

The APV architecture and the methodology used to design and implement this project can be considered as Best Practices for projects involving similar server virtualization configuration and implementation.

Appendix – Methodology and Tools

1. Methodology Flowchart

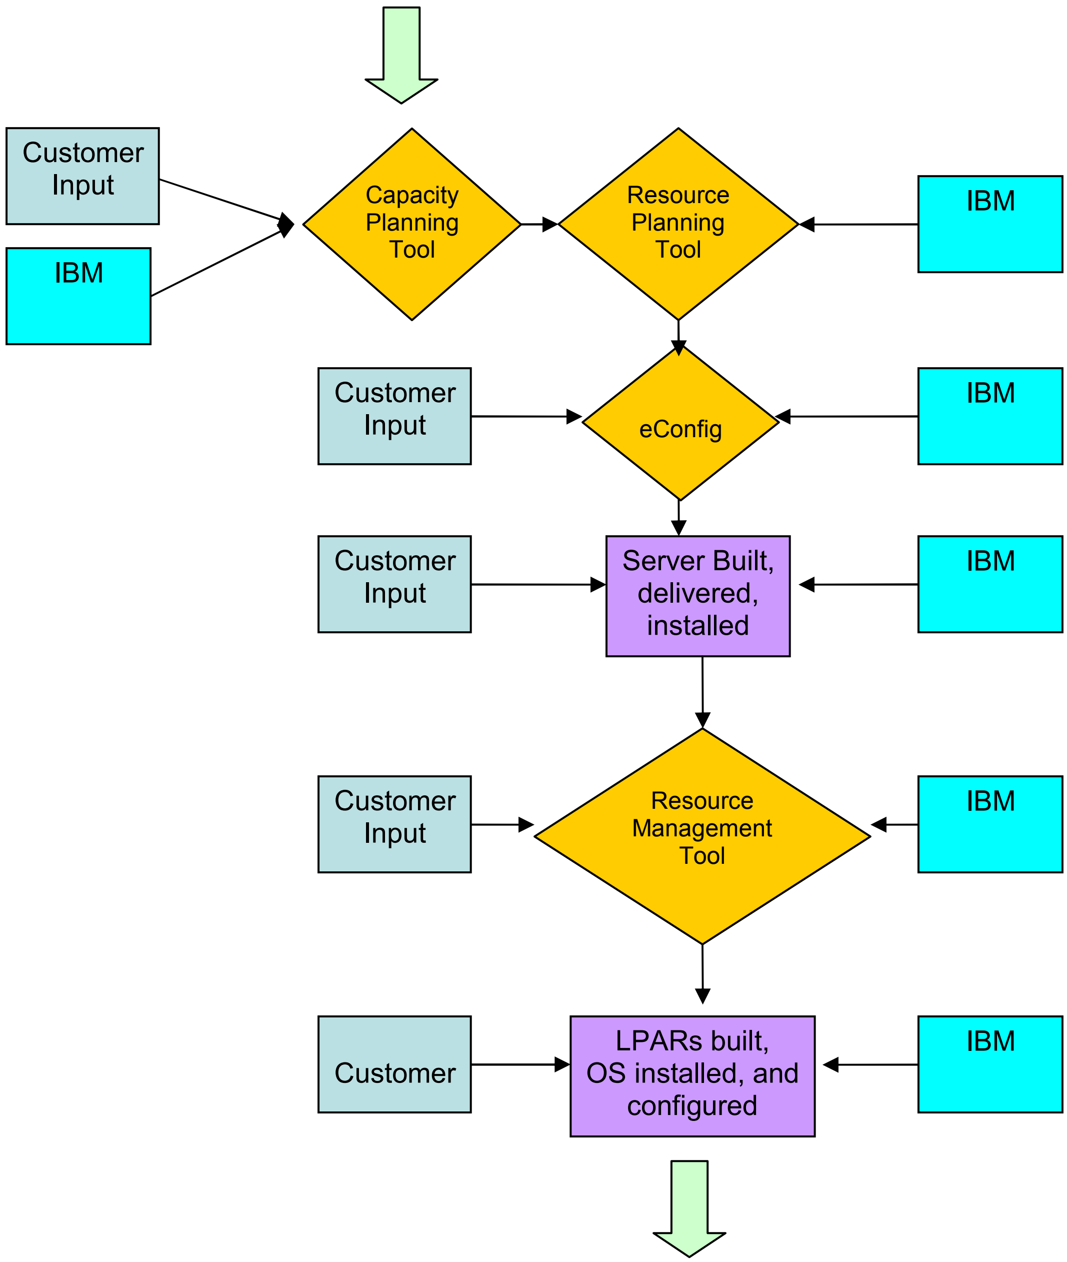

2. Capacity Planning and Server sizing

A customized capacity planning tool was used to compute the processor and memory required for the replacement servers. A sample screen shot of this tool below shows the details.

This tool was used to size LPARs, desired no. of VIO Servers and the spare pool. Variables used in sizing these LPARs were utilization, projected workload growth and maximum utilization factor.

1. Present configuration									
Server	Type-Model Serial #	Clock Speed & Type	Processors	Memory in GB	Max CPU utilized %	Max. rPerf of the config.	Utilized rPerf		Applications
Production									
DS									
dsc01	p630	1GHz, p4	4	8	75	7.12	5.34		UDB
dsc03	p630	1GHz, P4	4	8	75	7.12	5.34		UDB
dsc05	p630	1GHz, P4	4	8	90	7.12	6.41		UDB
dsc07	p630	1GHz, P4	4	8	75	7.12	5.34		UDB
dsc09	p630	1GHz, P4	4	8	80	7.12	5.70		UDB
dsc11	p630	1GHz, P4	4	8	95	7.12	6.76		UDB
dsc13	p630	1GHz, P4	4	8	100	7.12	7.12		UDB
dsc15	p630	1GHz, P4	4	8	100	7.12	7.12		UDB
wp01	p630	1GHz, P4	4	4	50	3.59	1.80		UDB
isp01	p630	1GHz, P4	4	8	60	7.12	4.27		WAS
wp02	p630	1GHz, P4	4	4	60	3.59	2.15		WAS
Backup	p630	1GHz, P4	2	4	0	4	0.00		
CWS	p610	450MHz	2	2	0	2.27	0.00		
Other H/W - SP Switch2, HMC									
Totals	**13**		**48**	**86**		**77.53**	**57.35**		

Dev									
DS-IS	3xSP Nodes	375 MHz	12	12	100	10.77	10.77		UDB
QA									IHS, WAS
tdv6	1xSP Node	375 MHz	4	4	100	3.59	3.59		
Totals	**3**		**16**	**16**		**14.36**	**14.36**		

2. Capacity Planning and server sizing									
Server	Max. rPerf of the present config.	Utilized rPerf	Growth Factor	Desired Max CPU utilization	Sizing Factor	Required rPerf	Reqd. Processors	Memory in GB	Model/ Type
Production									
11xDS	77.53	57.35	1.3	0.8	1.63	93.19	19.5	78	
4xQA	28.48	28.05	1.3	0.8	1.63	45.59	9.5	38	
9xIS	64.08	61.09	1.3	0.8	1.63	99.27	20.7	83	
Totals for Prod						**238.05**	**49.75**	**199**	
4xVIO Servers						19.14	4.00	16.00	
Shared pool @ 20%						51.44	10.75	43	
Totals						**308.62**	**64.50**	**258**	
Recommended Configuration						306.21	64	256	16x4-way P5-595 block, 1.9GHz

Dev									
6xDS-IS-QA	14.36	14.36	2.15	0.8	2.69	38.59	9.0	36	
1xWeb AIX52	3.59	3.59	1.90	0.8	2.38	8.54	2	8	
Totals						**47.13**	**11.04**	**44**	
2xVIO Servers						8.54	2.00	8	
Shared pool @20%						11.13	2.61	10	
Totals						**66.81**	**15.65**	**63**	
Recommended Configuration						68.32	16	64	16x4-way P5-595 block, 1.65GHz

Note:

1. In a LPAR environment, rPerf is estimated/CPU (rPerf of 64P/p595-1.65GHz is 273.1 with a rPerf/CPU of 4.27 and for 1.9 GHz is 306.21 and rPerf/CPU is 4.78).
2. RPE2 is the Relative Performance Matrix published by Ideas International.
3. Memory has been configured 4 GB/processor.
4. P5 servers offer several options to consolidate and configure servers, the option recommended here is one of them.
5. For Servers with no, performance data, it is assumed to be 75% utilized.
6. Growth rate is estimated to be 30% for servers/LPARs during the next three years lease period.
7. The performance of sized servers/LPARs may vary and depend on many factors including varying peak load, system hardware configuration, software installed and tuned, applications installed and configured, I/O and network performance, tuning of H/W, OS and application for optimum performance. The suggested configuration is a best estimate to run the projected workload and does not provide a guaranteed system performance.
8. Performance in a virtualized environment depends on VIO servers resources, shared resources such as Virtual SCSI (internal disks and LAN disks), Virtual Ethernet, Shared Ethernet and Shared processor pool.
9. rPerf (Relative Performance) is an estimate of commercial processing performance. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks and does not represent any specific benchmark results. Although rPerf

may be used to compare estimated IBM UNIX commercial processing performance, actual system performance may vary and dependent upon many factors including Hardware configuration and Software design and configuration.

3. Resource Planning

The resources required for each LPAR were estimated and allocated using the Resource Planning tool. The table below shows the I/O resources – boot disks, Ethernet adapters and SAN adapters required and if these resources were either shared or dedicated. Based on this information, eConfig tool was used to configure and order p5-595 servers.

Production									
LPAR	**Procs**	**Mem**	**Drives**	**Intranet**	**BKUP**	**DMZ**	**Pvt. GB**	**V-ENT**	**SAN**
			(Bt+Alt+Mir)						
9xDS	16.1	64	9x3x36	S11	S12		S14	S15	18
2x Web & App	4	16	6x72	2x4P	<==Use 2x2P		2x2P		4
4xQA	9.5	38	4x3x36	S11	S12		S14	S15	8
8xIS	18.4	74	8x3x36	S11	S12		S14	S15	16
Web/App	2	8	3x72	1x4P	<==Use 1x2P		1x2P		2
4 x VIOS	4	16	8x72	4x2P	4x4P		4x2P	S15	12
Total	54	216	17x72 24x146	4x2P 3x4P	4x4P		7x2P		60
Spare pool @ 20%	10.42	40							
Total required	64	256	17x72 24x146	4x2P 3x4P	4x4P		7x2P 0x4P		60
NIM	1	2	2 x 36	S11				S15	S16
P5-595 Config.	64	256	17x72GB 24x146GB	11x2P 7x4P					60
Development									
LPAR	**Procs**	**Mem**	**Drives**	**Network1**	**BKUP**	**DMZ**	**Pvt. GB**	**V-ENT**	**SAN**
			(Bt+Alt+Mir)						
6xDS-IS-QA	9	36	6x3x36	S11	S12			S15	S16
1x Web AIX5.2	2	8	3x72	1x2P	1x2P				2
2 x VIOS	2	8	4x72	2x2P 1x4P	2x2P 1x4P			S15	6
Total	13	52	7x72GB 6x146GB	3x2P 1x4P	3x2P 1x4P				8
Spare pool @ 20%	2.6	10							

Total required	16	62	7x72GB 6x146GB	3x2P 1x4P	3x2P 1x4P				8
4x4x16/1.9GHz Blocks	16	64	12	3x2P 1x4P	3x2P 1x4P				10
NIM	1	2	2 x 36	S11				S15	S16
4 Way BB	4	16	3xDASD	2x4P					2.5
P5-595 Config.	64	256	20x72GB 24x146GB	24x2P 8x4P					40
Guidelines									
App. Server	2	16	3 x 36	1x2P	S				S
DB	3	32	3 x 36	S	S				3
2x VIOS	2	8	4 x 72	2x2P	2x2P				6
NIM	1	2	2 x 36	S	S			S	S

Notes:

1. Data traffic between
1.1. web/app server (AIX5.2) to dB LPARs - Use Pvt. Gb network or intranet
1.2. dB LPARs - Use VENT
1.3. Mainframe and dB LPARs - Use intranet or Pvt. Gb.

2. SAN
2.1. VIOS will have 3 FCs.
2.2. Prod-Both AIX5.2 and AIX5.3 LPARs have dedicated FCs.
2.3. Dev-AIX5.2 has dedicated FCs, all AIX5.3 LPARs share FCs assigned to VIOS.

3. Network
3.1. Each 4P adapter - On AIX52 LPARs, 2Ports for Intranet (NIB), 1 Port for Backup network.
3.3. Backup network - Shared on VIOS for AIX5.3 LPARs, dedicated for AIX5.2 LPARs
3.4. Prod - Each VIOS has a 4P adapter for the backup network.
Atleast two adapters to be shared with VIPs to back-up TBs of data.
3.5. VNET - Virtual Ethernet to share data between AIX 5.3 LPARs, change MTU to 64K.
3.6 Explore with Network group to enable Jumbo Frame and Etherchannel for GB networks.

4. VIOS
4.1. VIOS1 and 2 will serve DS LPARs, VIOS3 and 4 will serve IS and QA LPARs.
4.2. VIOS1 & 2 and VIOS3 & 4 will serve both VSCSI (boot disks and SAN) and SEA to client LPARs.

5. Disks
5.1. VIOS will have 1Pack x 2 x 73 GB Drives
5.2. VIOS1 and 3 will have 2Packs x 4 x 146 GB Drives for LPARs (Boot and Alt)
5.3. VIOS2 and 4 will have 1Pack x 4x 146 GB drives for LPARs (Mirror).
5.4. AIX5.2 LPARs will have 1Pack x 3x 72 GB drives (Boot, Alt, Mirror).
5.5. Will have either 2 or 3 spare disks packs for additional drives, if required.

6. Spare pool accommodates spikes and growth requirements.

7. Shared Resource naming convention

	Intranet	Bkup	DMZ	Pvt. Gb	VENT	SAN
Intranet	S11	S12	-	S14	S15	S16
DMZ Web	-	S22	S23	-	S25	S16
DMZ App.	-	S32	S33	-	S35	S16
DMZ Data	-	S42	S43	-	S45	S16

4. Resource Allocation and Management

This tool was used to allocate resources to LPARs as planned in the Resource Planning tool. While allocating the resources, following guidelines/techniques were used.

1. Boot disks: Boot disk and alt_disk were allocated from one VIOS and rootvg mirror disk was allocated from the second VIOS to each LPAR.
2. Ethernet adapters: All the network adapters were configured using the NIB feature of AIX. Multiple adapters of same type were allocated from different Remote I/O Drawers (RIO) to eliminate SPOF.
3. SAN adapters: Multiple SAN adapters to a LPAR were allocated from different RIOs to eliminate SPOF.

This document titled "p5-595 Resource Management Worksheets" can be downloaded from IBM TechDocs Library using one of the URLs below.

IBM intranet
http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1986

Internet
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1986

Business Partners
http://partners.boulder.ibm.com/src/atsmastr.nsf/WebIndex/PRS1986

This tool has several worksheets for Resource allocation and management. These were used during LPAR profile creation and then later for managing the resources.

Worksheet1 – System Profile: Here Processor and memory allocation to LPARs are managed. This also shows the total I/O resources – boot disks, Ethernet and SAN adapters available and free, and spare pool of processors and memory.

Worksheet2 – Resource Allocation of Boot disks, Ethernet and SAN adapters: In this worksheet boot disks (dedicated or shared), Ethernet adapters (dedicated or shared) and SAN adapters allocated with WWN (dedicated or shared) are listed for each LPAR. Slot inf. is also provided for each adapter allocated to a LPAR.

Worksheet3 – Adapter Information: This worksheet gives the adapter view of the RIO draw with slot nos., locations codes, adapter type and the LPAR information.

Worksheet4 – Boot disk information: This shows the disks view of the RIO draw with slot nos., locations codes, disk type and the LPAR information.

References

1. Advanced POWER Virtualization on IBM eServer p5 Servers: Architecture and Performance Considerations, SG24-5768-01.

2. Advanced POWER Virtualization on IBM System p5, SG24-7940-01.

3. Partitioning Implementations for IBM eServer p5 Servers, SG24-7039-02.

4. AIX 5.3 and System p Documentation.

http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/aix43.htm?lang=en_US#sysmgmt
http://publib.boulder.ibm.com/eserver/
http://www-304.ibm.com/jct01004c/systems/support/

5. Several Whitepapers and Redpapers on Advanced Power Virtualization.

http://www-03.ibm.com/systems/p/library/index_lit.html
http://publib-b.boulder.ibm.com/Redbooks.nsf/redbooks/

6. IBM System p Advanced POWER Virtualization Best Practices, Redpaper 4194.