IBM Server Technology Group

IBM Virtualization – SAP and Power Technologies

Jan Kristian Nielsen
jankn@dk.ibm.com

IBM Systems

Table of contents

IBM Systems Agenda

- *A strategic design for delivering innovative technology, resources and skills*

IBM Systems Agenda

Working together with clients and IBM Business Partners, IBM provides the breadth of expertise and resources to design, build and implement IT solutions that can drive innovation and help transform your business.

IBM is committed to sharing technology with the industry and providing the industry's richest portfolio of interoperable server and storage systems.

Optimize IT operations and dynamically respond to the priorities of the business by managing the IT environment more efficiently with proven IBM cross-platform virtualization capabilities.

IBM Power Architecture™

From consumer electronics to supercomputers

A common architecture . . . the most scalable technology

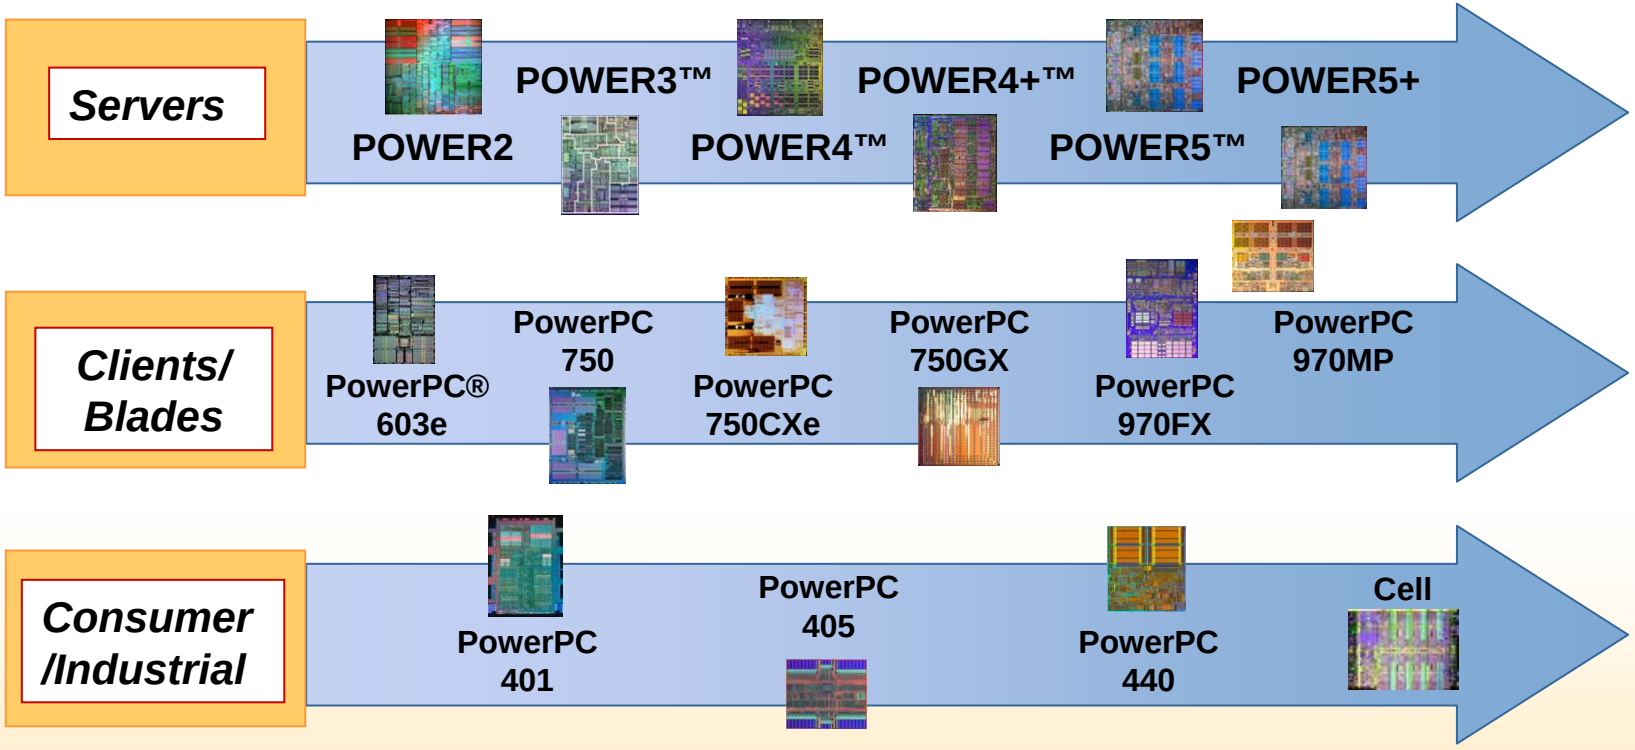

Source: http://www.ibm.com/chips/power/aboutpower/

POWER5+ packaging

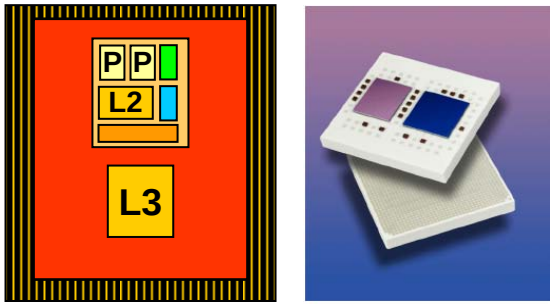

Dual-Core Module
POWER5+ Dual-Core chip + L3 cache chip
(505, 510, 520, 550, 570)
Two processor cores

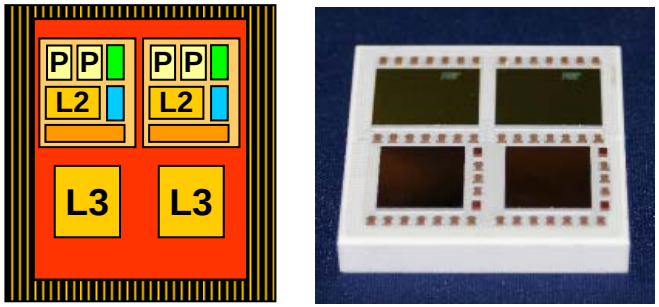

Quad-Core Module
Two POWER5+ Dual-Core chips + two L3 cache chips (505Q, 510Q, 520Q, 550Q, 560Q)
Four processor cores

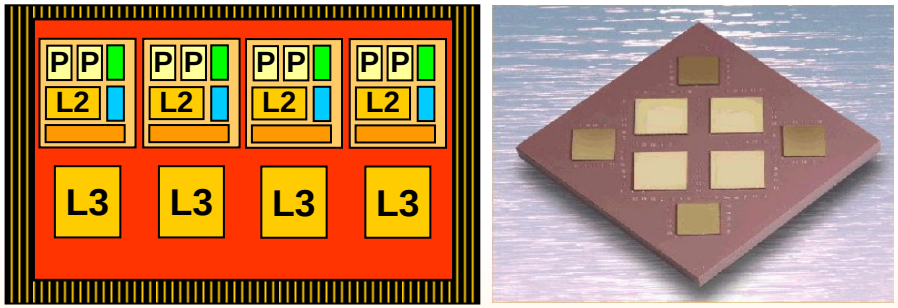

Multi-Chip Module
Four Dual-Core POWER5+ chips + four L3 cache chips
(590, 595)
Eight processor cores

IBM System p5: Simultaneous multithreading

POWER5+ (simultaneous multithreading)

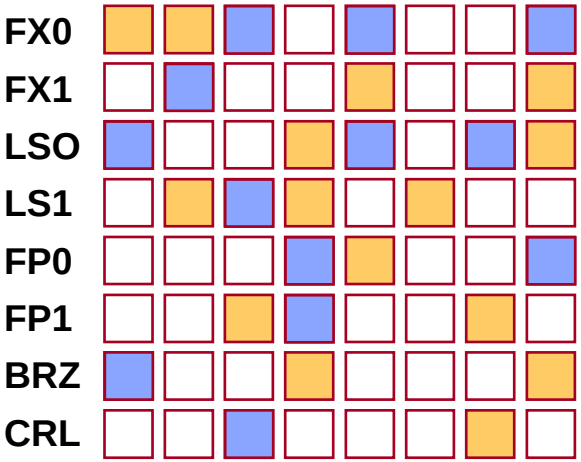

Appears as four CPUs per chip to the operating system (AIX 5L V5.3 and Linux)

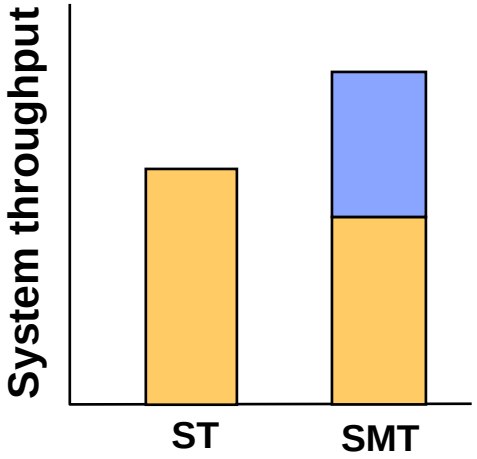

- Utilizes unused execution unit cycles
- Presents symmetric multiprocessing (SMP) programming model to software
- Natural fit with superscalar out-of-order execution core
- Dispatch two threads per processor: *"It's like doubling the number of processors."*
- Net result:
  - ***Better performance***
  - ***Better processor utilization***

IBM's 39-year history of leadership in virtualization

1967	1973	1987	1997	2001	2004
IBM develops **hypervisor** that would become VM on the mainframe	IBM announces first machines to do **physical partitioning**	IBM announces **LPAR on the mainframe**	**POWER LPAR design** begins	IBM introduces **LPAR in POWER4™ based systems with AIX 5L™**	**Advanced POWER™ Virtualization** ships

"In our opinion, they [System p servers] bring mainframe-quality virtualization capabilities to the world of AIX®."

- Ulrich Klenke, CIO, rku.it

January 2006

Timeline reference http://www.levenez.com/unix/history.html#01
Customer quote source: rku.it case study published at http://www.ibm.com/software/success/cssdb.nsf/CS/JSTS-6KXPPG?OpenDocument&Site=eserverpseries

Advanced POWER Virtualization on IBM Power System

IBM APV Benefits

- ☑ ***Can help lower the cost** of existing infrastructure **by up to 62%**[4]*
- ☑ *Can increase **business flexibility** allowing you to meet anticipated and unanticipated needs*
- ☑ *Can **reduce** the **complexity to grow your infrastructure***

Advanced POWER Virtualization[1]

Virtual I/O Server

- Integrated Virtualization Manager[3]
- Share Ethernet, SCSI and Fibre Channel disks

Partition Load Manager[2]

- Automatically balance processor and memory requests

Micro-Partitioning™

- Share processors across multiple partitions
- Minimum 1/10 th processor

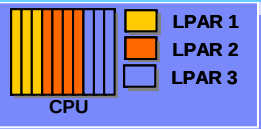

Linux on POWER

AIX 5L

1) Advanced POWER Virtualization (APV) is an optionally orderable feature on IBM System p, 2) Partition Load Manager (PLM) is not supported on OpenPower / Linux Partitions, 3) Only available on select models, 4) "Business Case for IBM System p5 Virtualization," Economic Benefits of IT Simplification. International Technology Group, 02/10/2006. Study methodology: Companies in financial services, manufacturing and retail with $15 Billion+ revenues and total 200,000+ employees focusing on UNIX® large enterprise environments with multiple, broad-ranging applications. Study compared the cost of the company's workload running on multiple vendor servers and employing minimal virtualization to the cost of the company's workload running on the p5-510, 550, 570, 590 and 595 – all using Advanced POWER Virtualization [APV]. APV is standard on System p5 590 and 595. Other System p servers have the option to add APV except the System p5 185. This cost analysis was performed for financial services, manufacturing and retail example environments with an overall average savings of up to 62% in TCO savings by virtualizing and consolidating on the System p servers. For further information, see the white paper at: http://www-03.ibm.com/systems/p/library /consult/ itg_p5virtualization.pdf Total Cost of Ownership may not be reduced in each consolidation case. TCO depends on the specific customer environment, the existing environments and staff, and the consolidation potential.

POWER5 / POWER5+ virtualization architecture

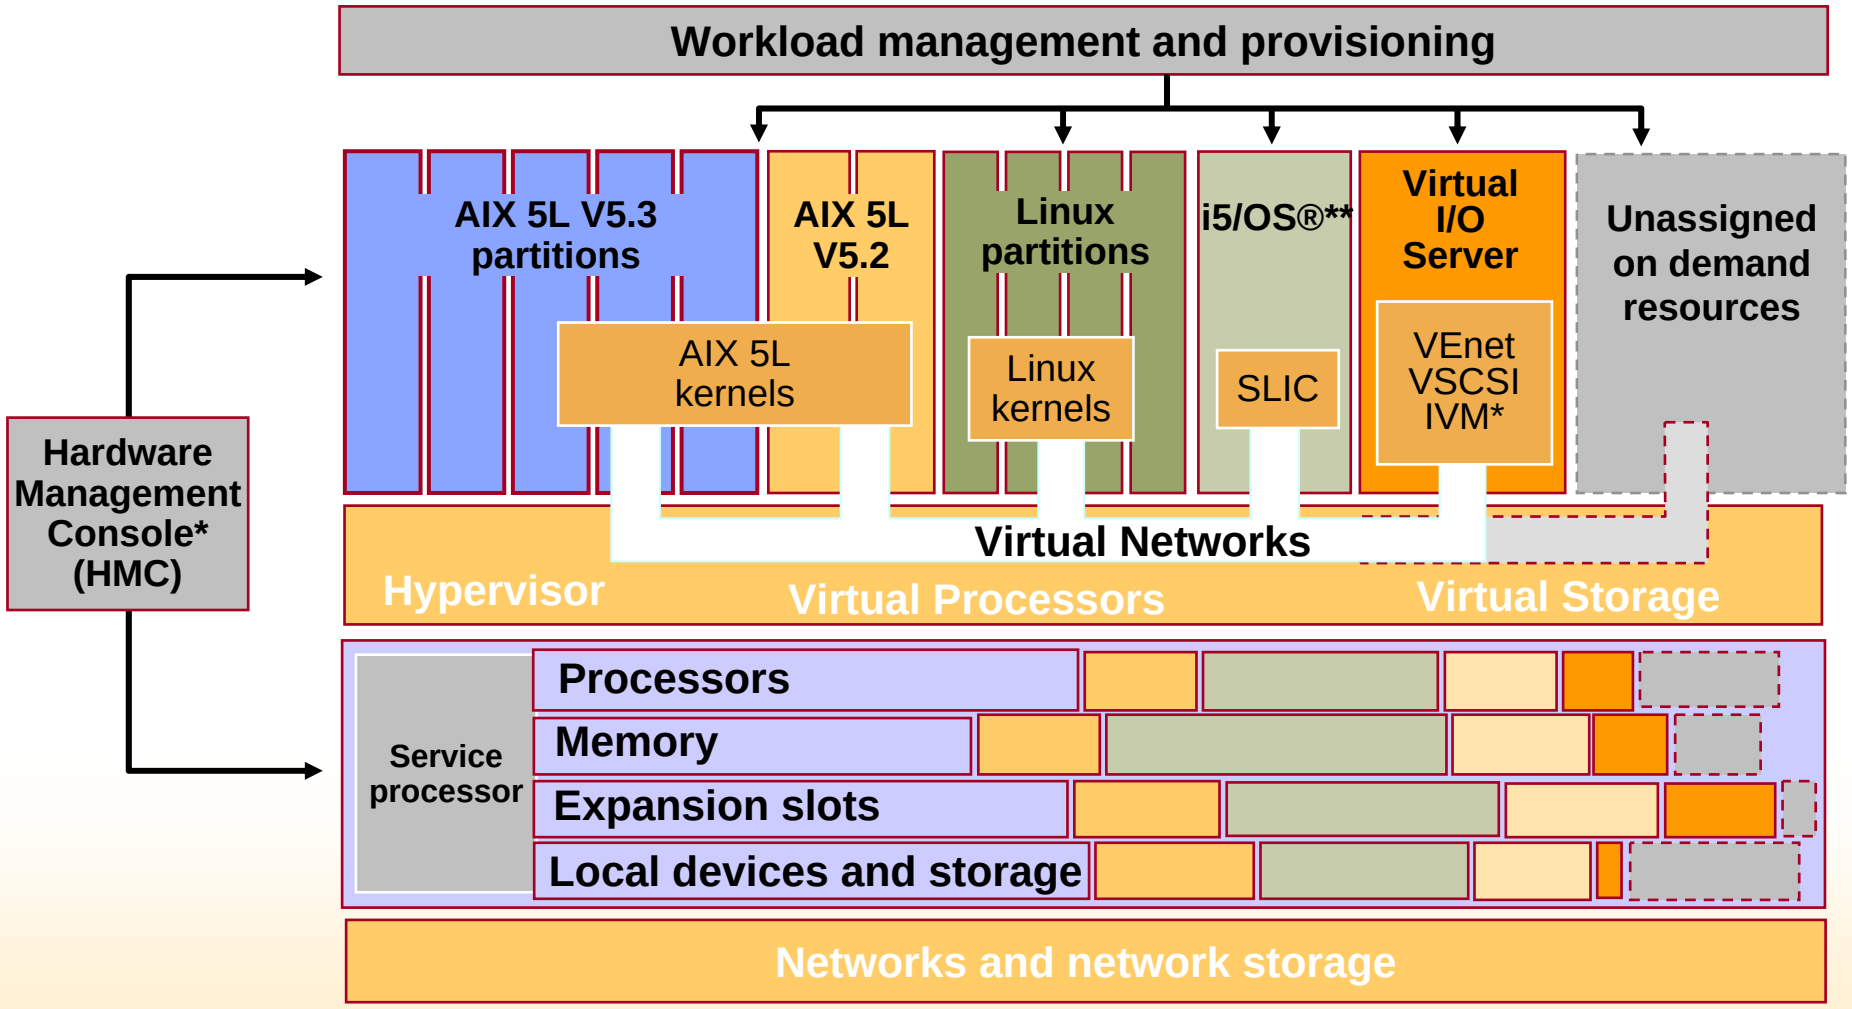

* Integrated Virtualization Manager (IVM) is disabled if HMC attached

**Available on System p5 560 and below as well as the BladeCenter® JS21

Virtual Ethernet connections

- **VLAN technology implementation**
  - Partitions can only access data directed to them.
- **Virtual Ethernet switch provided by the POWER Hypervisor**
- **Virtual LAN adapters appears to the OS as physical adapters**
  - MAC-Address is generated by the HMC.
- **1-3 Gb/s transmission speed**
  - Support for large MTUs (~64K) on AIX.
- **Up to 256 virtual Ethernet adapters**
  - Up to 18 VLANs.
- **Bootable device support for NIM OS installations**

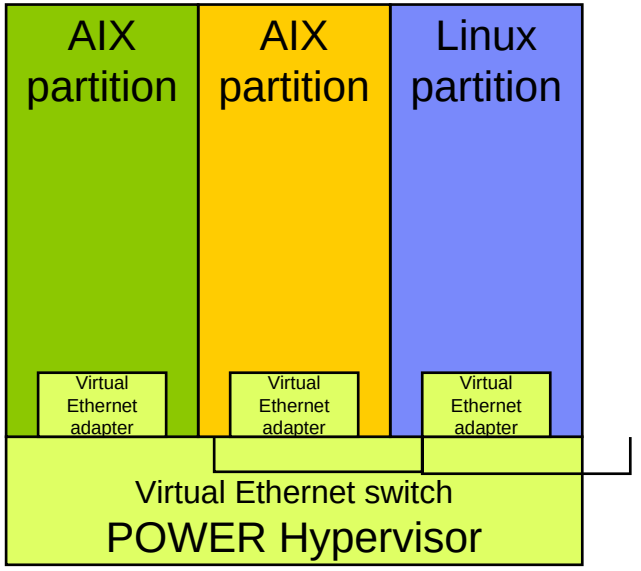

LVM mirroring

- **This configuration protects virtual disks in a client partition against failure of:**
  - One physical disk
  - One physical adapter
  - One virtual I/O server
- **Many possibilities exist to exploit this great function!**

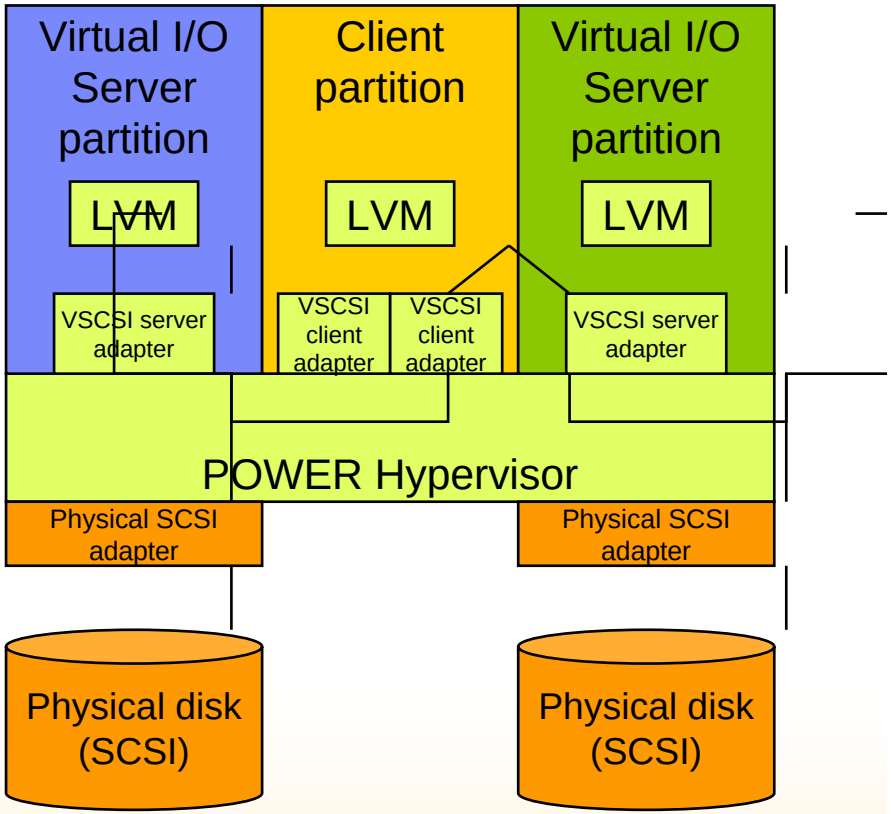

Advanced POWER Virtualization option

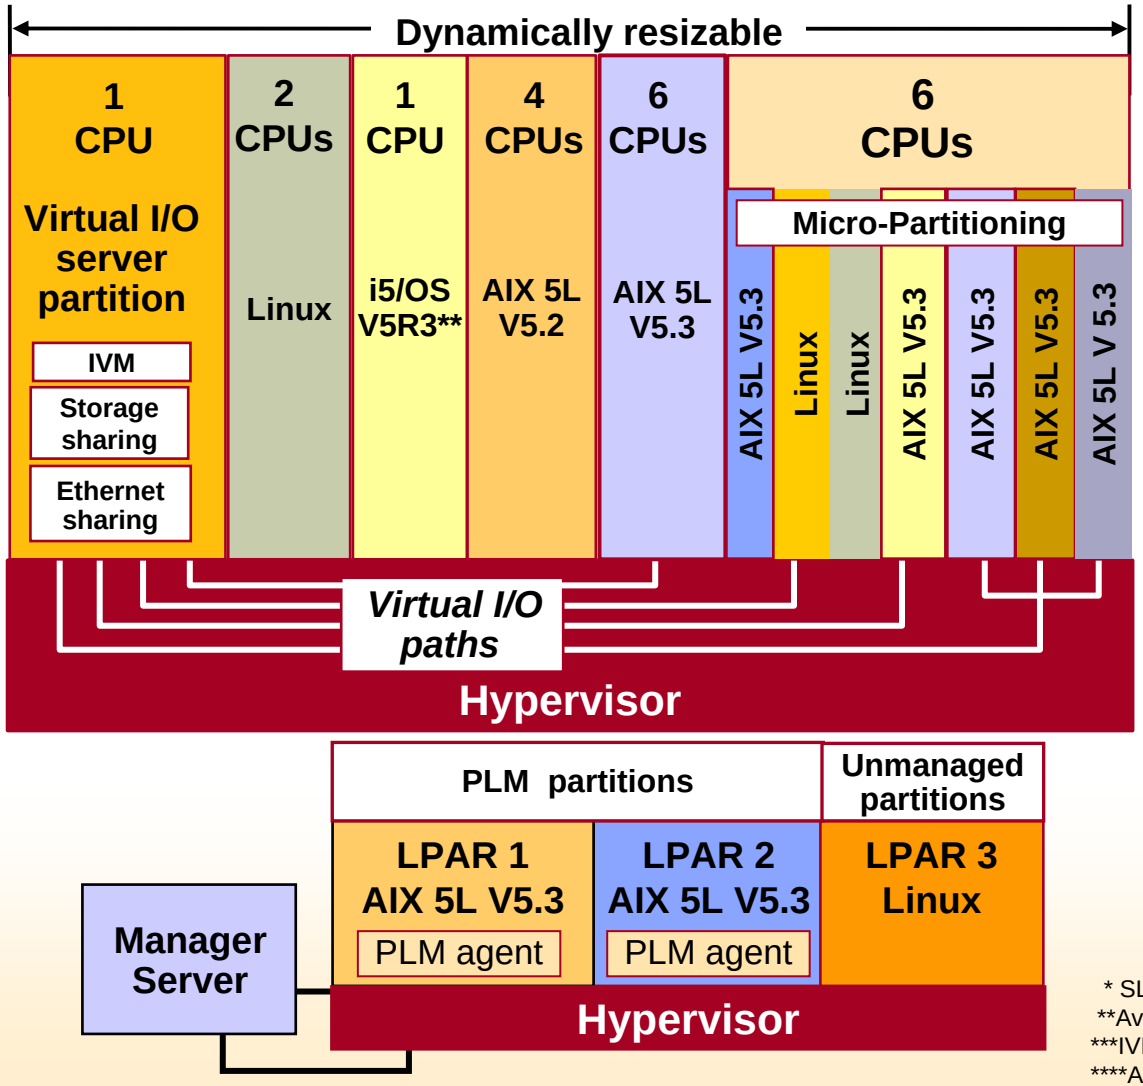

Virtual I/O Server

- Shared Ethernet
- Shared SCSI and Fibre Channel-attached disk subsystems
- Supports AIX 5L V5.3 and Linux* partitions

Micro-Partitioning

- Share processors across multiple partitions
- Minimum partition 1/10th processor
- AIX 5L V5.3, Linux*, or i5/OS**

Partition Load Manager****

- Balances processor and memory request

Managed via HMC or IVM***

* SLES 9 or RHEL AS 4 and above
**Available on selected p5-570, p5-590 and p5-595 models
***IVM on p5-560Q and below
****Available for AIX 5L V5.2 or above (RPQ required on POWER4)

Micro-Partitioning technology

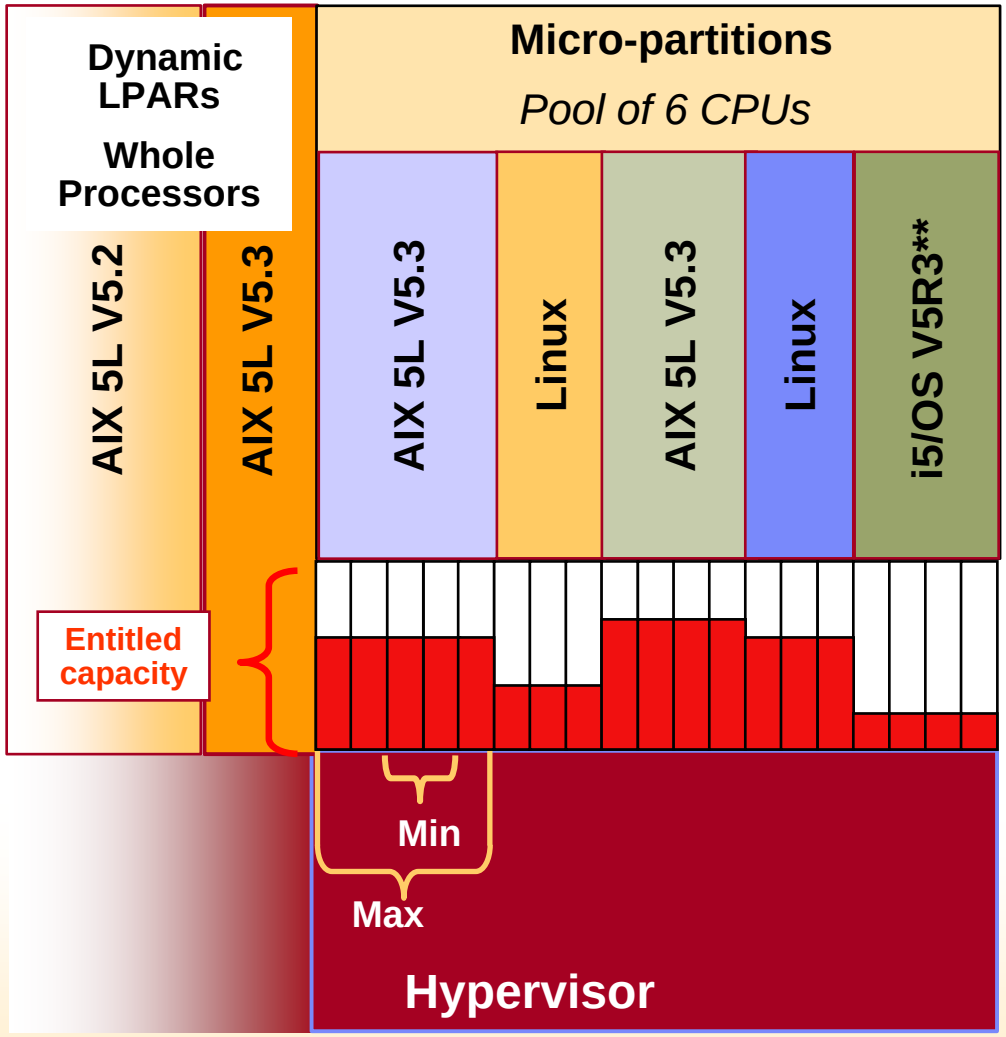

Note: Micro-partitions are available via optional Advanced POWER Virtualization or POWER Hypervisor and VIOS features.

Micro-Partitioning technology allows each processor to be subdivided into as many as 10 "virtual servers", helping to consolidate UNIX® and Linux applications.

Partitioning options

– Micro-partitions: Up to 254*

Configured via the HMC

Number of logical processors

– Minimum/maximum

Entitled capacity

– In units of 1/100 of a CPU
– Minimum 1/10 of a CPU

Variable weight

– % share (priority) of surplus capacity

Capped or uncapped partitions

*on p5-595
** on p5-570, p5-590, and p5-595

Partition Load Manager for AIX 5L

- **Policy-based, automatic partition resource tuning**
- **Dynamically adjust CPU and memory allocation**
- **Included with purchase of optional Advanced POWER Virtualization***

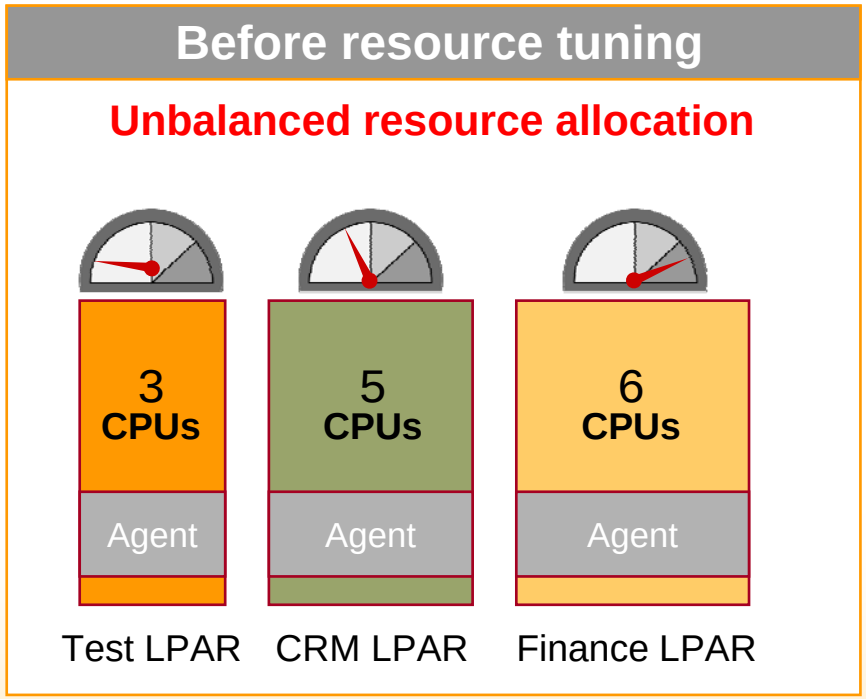

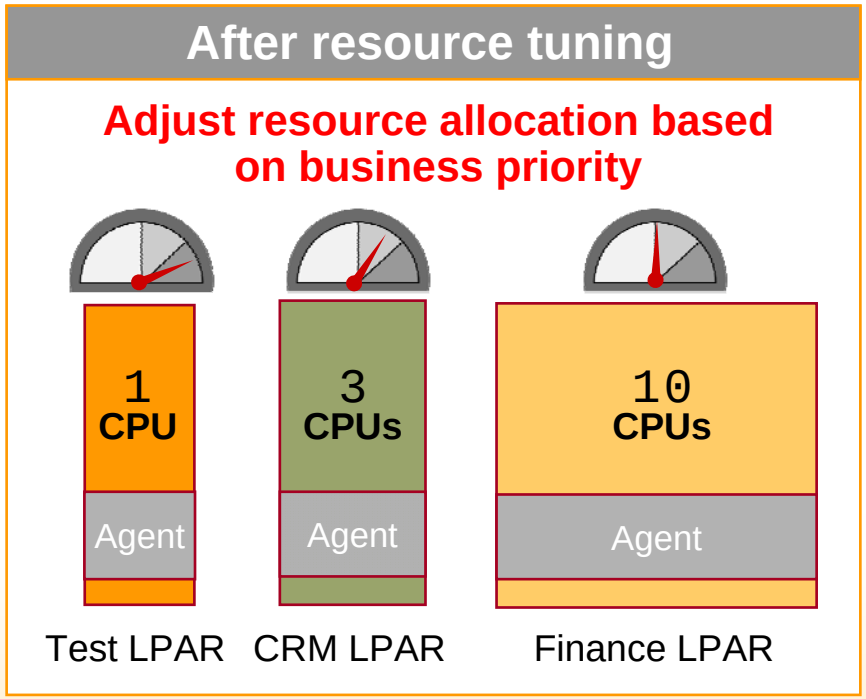

Note: Micro-partitions are available via optional Advanced POWER Virtualization or POWER Hypervisor and VIOS features. AIX 5L V5.2 supports LPAR only.

* Standard on p5-590 and 595; not available on p5-185

New!

APV enhancements for August 2006

Virtual I/O Server (VIOS) version 1.3

– VIOS Partition Monitoring **through Topas and PTX**

Topas (part of AIX 5L V5.3)

```
Topas CEC Monitor            Interval:  10
Partitions     Memory (GB)                Proce
Shr:  3        Mon:24.6  InUse: 2.7       Shr:1
Ded:  3        Avl:   -                   Ded:

Host         OS  M Mem InU Lp  Us Sy Wa Id  Ph
-------------------------------------shared---
ptoolsl3     A53 c 4.1 0.4  2  20 13  5 62   0
```

PTX (AIX LPP)

– Performance Enhancements **for Virtual SCSI and Virtual Ethernet**

Integrated Virtualization Manager (IVM)

– Decrease downtime **- Resize and modify partitions without server disruption with support for DLPAR**

– Save time – **With new simplified GUI for IP configuration support**

What problem does it solve?

Inability to allocate IT costs, usage, and value

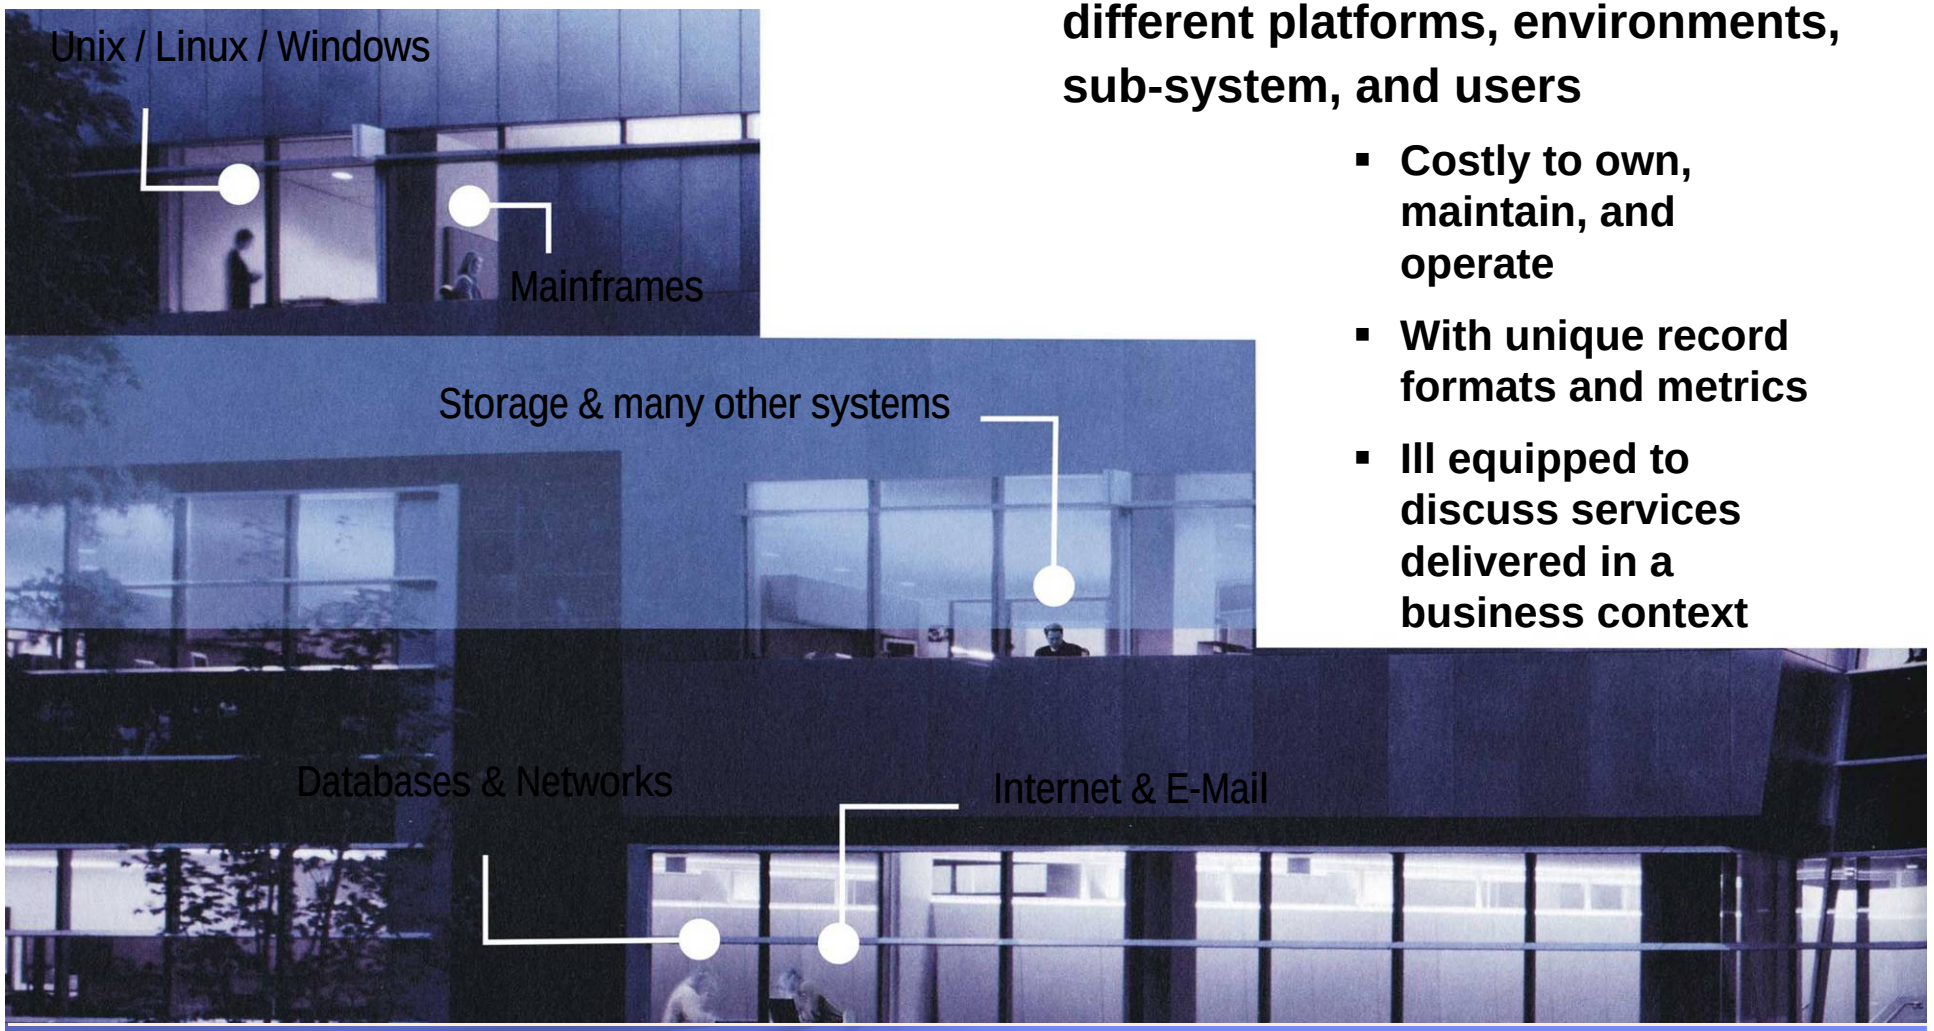

IT Departments contain many different platforms, environments, sub-system, and users

- **Costly to own, maintain, and operate**
- **With unique record formats and metrics**
- **Ill equipped to discuss services delivered in a business context**

Know what it costs – The IBM UAM funnel

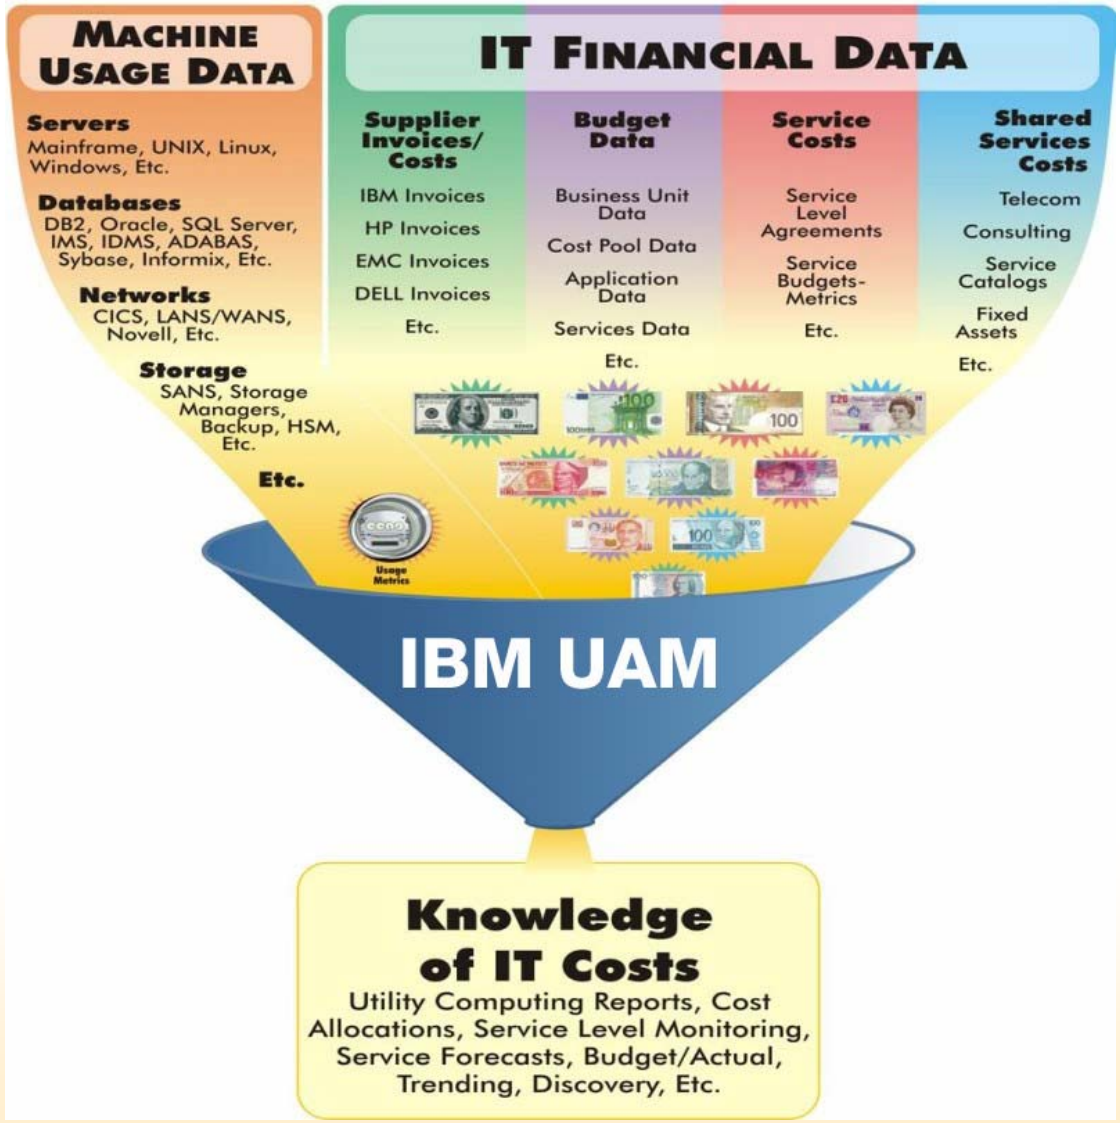

Objectives of an IT Usage and Accounting Management System…

- **Allocate/Distribute or Charge IT Costs to the Users, Cost Centers, Applications and other organizations that consumed them in a . . .**
  - Fair
  - Understandable
  - Auditable/Reproducible, and
  - Easy to administer manner
- **Optimize IT costs through…**
  - Costs trend identification
  - Real time analysis

IBM UAM features

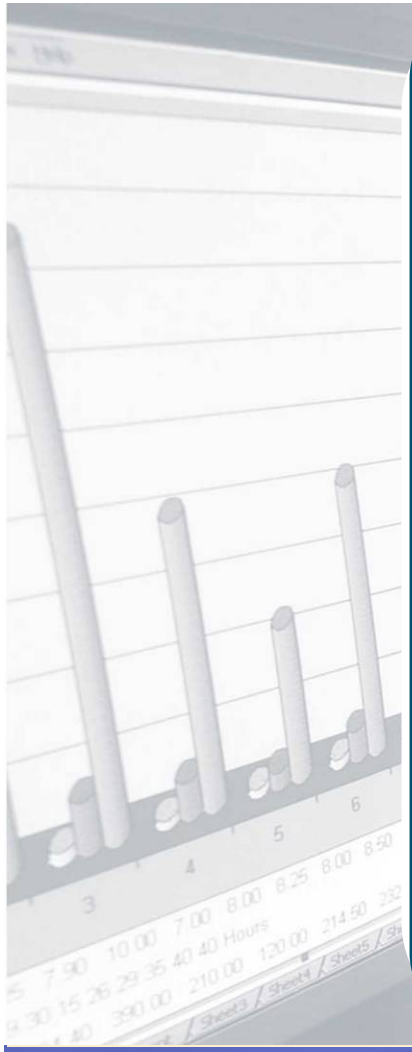

- Account Code Editing
- Account Code Validation
- Automatic E-Mail Reporting
- Automatic Web Reporting
- Billing Equation
- Budget / Actual Reporting
- Business Rules Engine
- Contract Pricing
- Conversion Engine
- Cost & Resource Analysis Reports
- Creates GL Transactions
- Disk Space Accounting
- Discounts
- Efficient Daily Processing
- Efficient EOM Processing
- External Billing
- Flexible Account Code Generation
- Flexible Account Code Reporting
- Full Time Administrator Not Required
- Integrated Rate Table/Service Catalogue

- Miscellaneous and Recurring Transactions
- Multiple CPU's of Differing Speeds Supported (Normalization)
- Multiple Rate Tables by Acct.
- Multiple Reporting Levels with drilldown
- Paper & Form Chargeback
- Proration
- Rate Modeling
- Sales Tax
- Security Authentication/LDAP
- Server Based Reporting System
- Server Based Stand-Alone System
- Shift, Class, Priority Surcharge
- Tiered Pricing
- Usage Discovery
- Web Enabled & Automatic HTML Creation
- Work Shift Reporting
- Year To Date Reporting
- Zero Based Budget Support