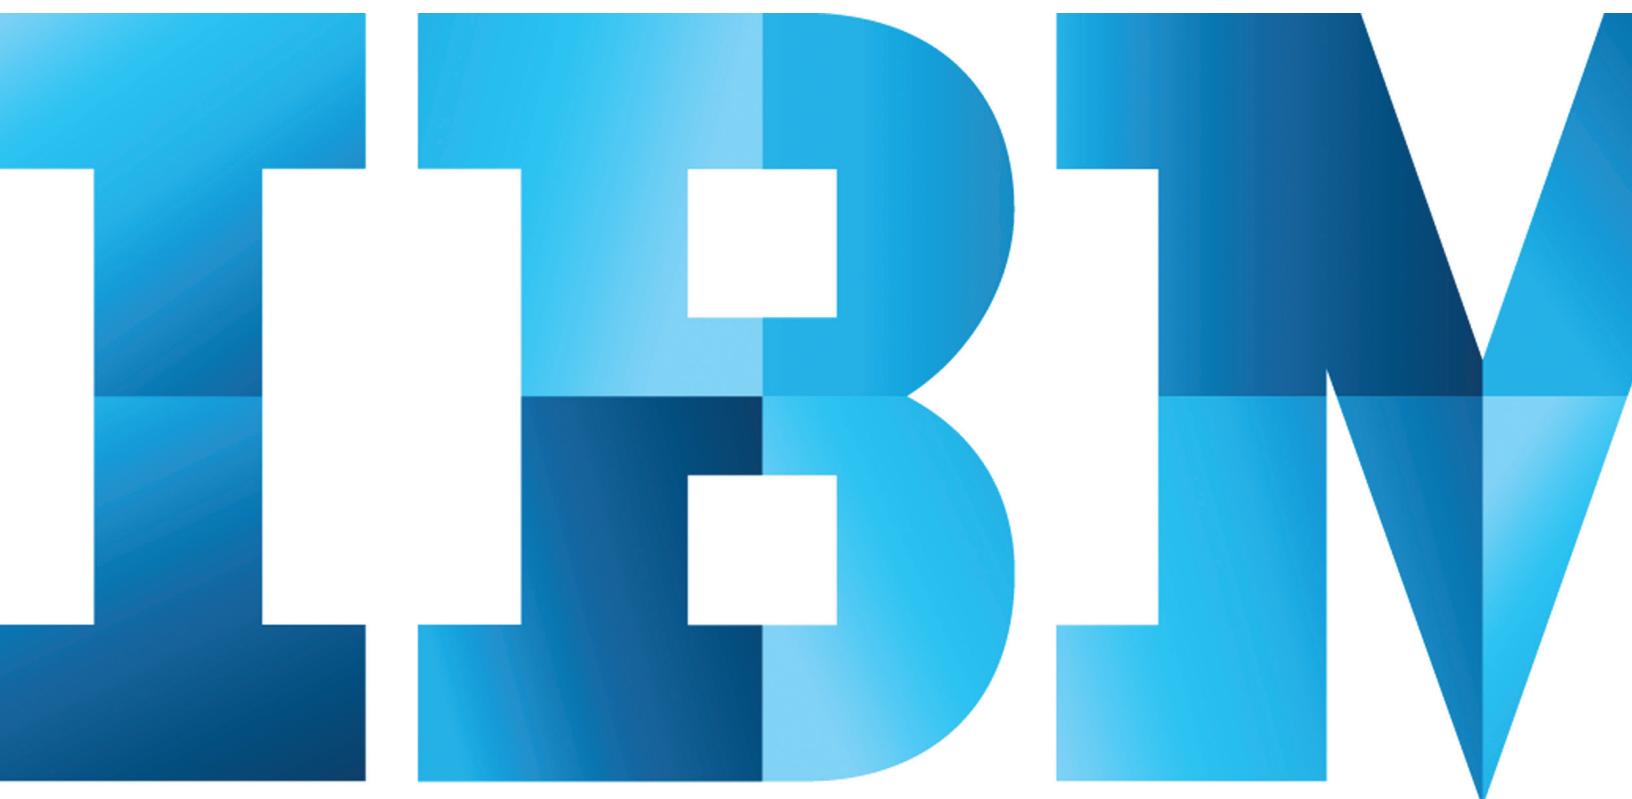


PureSystems

IBM PureFlex System

Integrated Solution for Greater Business Value



IBM

Introduction

The main concern for customers today is the lack of solutions to manage an exponential growth in the volume of structured and unstructured data. Another concern is with shrinking product cycles and focus on top-line growth. Customers are looking for infrastructure solutions that can be deployed rapidly, managed easily and that can run optimally without compromising on reliability.

As part of its mission to deliver IT solutions for a smarter planet, IBM has introduced IBM® PureSystems™, a family of expert integrated systems, that combine the flexibility of general-purpose systems, the elasticity of cloud computing and the simplicity of an appliance that is tuned to the workload. These system offerings are designed to deliver value in the following ways:

- Built-in expertise that helps address complex business and operational tasks automatically.
- Integration by design that helps tune systems for optimal performance and efficiency.
- Rapid deployment capabilities of these systems that provide a simplified experience where individual elements of an IT infrastructure need not be separately procured thereby helping customers focus on their core product deliverables.

In this paper, we first discuss the technology view and composition of IBM PureFlex System™ (one of the IBM PureSystems offerings) where we present the various compute nodes, networking and storage configurations that are available. Next we present the platform and virtualization management aspects of the system. Subsequently, we enumerate the challenges and pain points gathered from customer feedback and discuss the value proposition of IBM PureFlex System as a potential solution.

Technology View

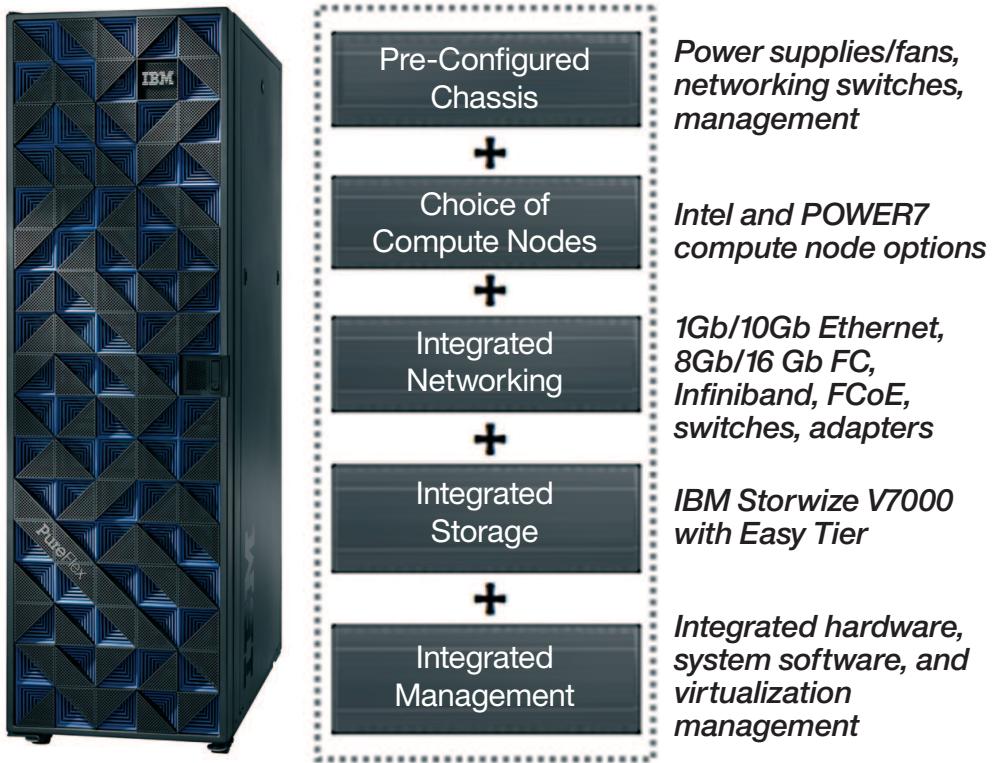
IBM PureFlex System is an expert integrated system by design. The server, storage, network, and software components are designed to work together as a coherent system, which offers the flexibility to choose any platform and virtualization capabilities such as:

- IBM PowerVM® on Power Systems™
- KVM/Hyper-V/VMware on Intel processor-based systems
- Variety of network options
- High-performance storage subsystems

This allows for rapid deployment since individual elements of an IT infrastructure need not be separately procured, thereby helping customers focus on their core product deliverables.

This section will provide an overview of all the building blocks that comprise the PureFlex System, including the IBM Flex System™ Enterprise Chassis, which houses various IBM Flex System compute nodes, storage nodes, the Flex System Manager appliance, and networking adapters and switches.

IBM PureFlex™ System



Management

IBM Flex System Manager™ provides a single pane of glass to manage hardware, software, virtualization, workloads, networks and storage elements seamlessly. The management solutions come with integrated views, intuitive GUI, and can intelligently manage and deploy workloads based on resource availability and predefined policies, and manage events and alerts to increase system availability and reduce downtime while reducing operational costs.

Flex System Manager is provided as an appliance in a dedicated compute node. The key features of Flex System Manager include:

- Multiple view overlays to track system health, firmware inventory, and environmental metrics
- Configuration management for setup of compute, network, and storage devices
- Remote presence application for remote access to compute nodes with single sign-on

Flex System Manager enables virtualization and workload optimization through various policies such as:

- **Resource utilization:** Congestion detection, customizable notification policies, and relocation of physical and virtual machines that include storage and network configurations.

- **Resource pooling:** Pooled network switching with placement advisors that consider VM compatibility, processor, availability, and energy.
- **Intelligent automation:** Automated and dynamic VM placement based on utilization, energy, predictive hardware failure alerts, and host failures.

IBM Flex System Chassis

The major components of the IBM Flex System Enterprise Chassis are:

- Fourteen 1-bay compute node bays which supports nodes based on both IBM POWER® and x86 processors.
- Six 2500-watt power modules with redundant power supply.
- Ten fan modules.
- Four physical I/O modules.
- Networking solutions that include Ethernet, Fibre Channel, FCoE, and InfiniBand.
- Two Flex System Manager management appliances for redundancy. Flex System Manager manages up to 4 chassis.
- Two IBM Chassis Management Modules (CMMs). CMM provides single-chassis management support.

Compute Nodes

IBM PureFlex System offers two major categories of compute nodes, based on x86 processors or IBM POWER processors.

IBM Flex System product offerings

Product	Value Proposition	Features
Flex System x220 (standard node)	Ideal for infrastructure workloads and entry virtualization.	Two Intel Xeon® E5-2400 processors, 16 cores, 12 DIMMs
Flex System x240 (standard node)	Ideal for mainstream virtualization and enterprise workloads.	Two Intel Xeon E5-2600 processors, 16 cores, 24 DIMMs, two hot-swap drives, two PCI Express I/O adapter slots, and an option for two internal USB connectors
Flex System x440 (double-wide node)	Ideal for high-end virtualization, mainstream databases, and memory-intensive high-performance workloads	Two Intel Xeon E5-4600 processors, 32 cores, 48 DIMMs, four I/O adapter slots, and an option for up to two internal drives for local storage

IBM Power product offerings

Product	Value Proposition	Features
System p260 (standard node)	Ideal for infrastructure workloads and entry virtualization.	Two POWER7 processor sockets, 16 memory slots, two I/O adapter slots, and an option for up to two internal drives for local storage.
System p24L (standard node)	Ideal for mainstream virtualization and enterprise workloads.	Two POWER7 processor sockets, 16 memory slots, two I/O adapter slots, and an option for up to two internal drives for local storage (The IBM PureFlex System p24L Compute Node is optimized for lower-cost Linux installations).
System p460 (double-wide node)	Ideal for high-end virtualization, mainstream databases and memory-intensive high-performance workloads.	Four POWER7 processor sockets, 32 memory slots, four I/O adapter slots, and an option for up to two internal drives for local storage.

Storage

IBM Storwize® V7000 storage systems integrated into the PureFlex System are designed to be easy to use and enable rapid deployment. Storwize V7000 systems provide extraordinary performance and flexibility through built-in solid-state drives (SSD), and optimization and thin provisioning technologies. Advanced storage features like automated tiering, storage virtualization, clustering, replication and multi-protocol support are designed to improve efficiency.

Network

Ethernet Adapters and Switches

IBM Flex System provides high-performance Ethernet offerings coupled with on-demand scalability to scale as IT requirements grow.

- IBM Flex System Fabric EN4093 10 Gb Scalable Switch supports 10G and 1G Ethernet
- IBM Flex System EN2092 1 GB Ethernet Scalable Switch along with upgrade options for future expandability.
- IBM Flex System EN4091 10 GB Ethernet Pass-Thru Module offers easy connectivity of the Flex System Enterprise Chassis to any external network infrastructure.
- IBM Flex System CN4054 10 GB Virtual Fabric Adapter based on Emulex BE3 ASIC supports multiple advanced protocols for Intel processor-based IBM Flex System compute nodes.

- IBM Flex System EN4054 4-port 10 Gb Ethernet Adapter Emulex BE3 ASIC is designed to run on IBM Flex System POWER7® processor-based compute nodes
- IBM Flex System EN4132 2-port 10 Gb Ethernet Adapter based on Mellanox Connect X3 ASIC supports RDMA and RoCE technologies, helping provide application acceleration and low latency for specialized applications.
- IBM Flex System EN2024 4-port 1 Gb Ethernet Adapter Supports IO virtualization features like VMware, NetQueue and Microsoft VMQ technologies.

Fiber Channel Adapters and Switches

The IBM Flex System portfolio offers both 8 GB and 16 Gb Fiber Channel SAN for storage connectivity.

- IBM Flex System FC5022 8/16 Gb SAN Scalable Switch from Brocade
- IBM Flex System FC3171 8 Gb SAN Switch from QLogic
- IBM Flex System FC5022 2-port 16 Gb Fibre Channel Adapter From Brocade offers 4G/8G and 16G Fiber Channel
- IBM Flex System FC3052 2-port 8 Gb Fibre Channel Adapter from Brocade works with any 8 Gb or 16 Gb Flex System Fibre Channel switch modules
- IBM Flex System FC3172 2-port 8 Gb Fibre Channel Adapter from QLogic works with any of the 8 Gb or 16 Gb Flex System Fibre Channel switch modules. When compared to the previous generation of 4 Gb adapters, the new 8 Gb adapters double the throughput speeds for Fibre Channel traffic.

InfiniBand Adapters and Switches

IBM Flex System portfolio includes next generation InfiniBand offering for high performance and low latency required for HPC and financial services applications.

- IBM Flex System IB6131 InfiniBand Switch based on Mellanox architecture designed to offer the performance to support clustered databases, transactional services and high-performance embedded I/O applications.
- IBM Flex System IB6132 2-port FDR InfiniBand Adapter providing 56 Gbs bandwidth based on Mellanox architecture
- IBM Flex System IB6132 2-port QDR InfiniBand Adapter based on Mellanox architecture delivers 40 Gbps per port (4x10 Gbps), providing the necessary bandwidth for high-throughput applications.

Security

IBM PureSystems leverages the Trusted Computing Group standards to create a secure infrastructure across the compute, network and storage elements. Security is guaranteed at various levels:

- **Secure Boot**—Ensures that only trusted firmware, hypervisor, and workloads get executed on the IBM hardware
- **Hardware SPI headers**—Methods to read the flash device headers in the field to validate firmware authenticity and thereby ensure that the system is not a victim of supply chain attacks
- **Secure Management traffic**—The traffic between the Chassis Management Modules, Flex System Manager and the various compute nodes are encrypted and secure.
- **Network isolation** between servers is implemented using VLANs

Business Value Challenges

Businesses today face several critical challenges that impact profit margins, productivity, revenue and customer satisfaction.

- Complex IT infrastructure
 - Aging infrastructure with little or no flexibility
 - Server and storage capacity saturation
- Low utilization and efficiency: un-optimized applications
 - Significant performance bottlenecks
- High operational cost
 - Increasing management and administration costs
 - Increasing energy footprint
 - High workload deployment delay

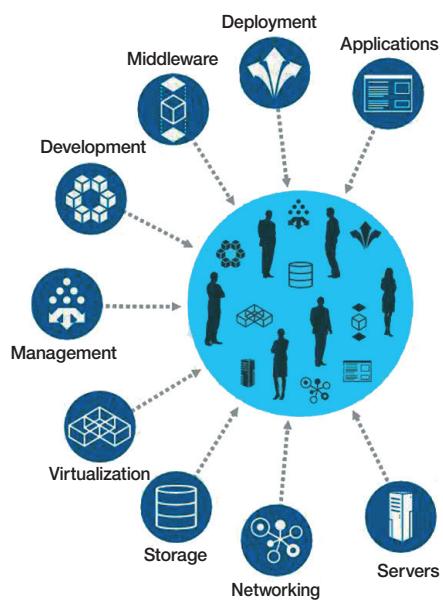
What Customers Want

In order to overcome these challenges companies are looking at end-to-end solutions which offer:

- Seamless HW integration across server, storage and networking
- Centralized management on the compute, storage and network layers
- High-performance compute nodes, which are ideal for high-end virtualization, mainstream databases and memory-intensive high-performance workloads
- Enterprise-class storage solution with Fibre Channel, which provides reliability along with virtualization support, easy-tiering and thin provisioning features
- Enterprise-class networking solution to support growing businesses' bandwidth needs along with support for open standards, such as 802.1qbg with virtual machine-aware switches
- Improved workload optimization and deployment solutions

Value Proposition

IBM PureSystems integration by design



Optimizes the complete solution stack:

- All hardware and software components factory integrated and optimized
- Born virtualized and ready for cloud
- Storage tuned to data needs
- Hardware directly tuned to the software
- System resource allocation uniquely optimized per selected pattern for each application workload



IBM PureFlex System is designed to enable companies to improve the way they deploy and manage their IT environments to achieve key benefits:

- **Consolidation:**

- Integrated infrastructure supporting compute, network, and storage to reduce physical space, HW/SW costs, maintenance, and energy use
- High flexibility of choice for compute, OS, and hypervisor technology, yielding low cost in moving legacy applications and adding new applications
- Integrated storage management of IBM Storwize V7000

- **Optimize:**

- Single pane of glass management of all resources with automated processes, yielding a reduction in the skill requirement and time required to manage and deploy a system
- Dynamic System Scaling: Workload Optimization, deployment under five minutes, faster storage deployment and provisioning
- Dynamic Workload provisioning and deployment
- Tightly integrated server, network and storage, which leads to improved setup time and energy savings of up to 42 percent.

- **Innovate:**

- Rapidly deploying new applications with the built-in patterns of expertise.
- Patterns are applied across compute, storage, and network resources of the IBM PureFlex System, providing a foundation for simplifying deployment of tuned partner applications.
- Faster and simpler application deployment. ISVs can provide virtual appliances that are designed to be downloaded and run as is, complete and ready to go—similar to an ‘app Store’, or ‘app service policies’.

BPTP Case Study

BPTP Limited is a leading real estate company in India that has adopted IBM PureFlex System to streamline its IT infrastructure and improve the overall home buying experience for its customers. BPTP utilizes business analytics to match the best properties with its customers.

BPTP’s storage, processing and overall IT infrastructure requirements have grown considerably owing to the country’s urban growth. The key requirement for BPTP is to enable consolidation, streamlined management, and improve performance while building a highly scalable infrastructure.

BPTP had been looking at a hardware refresh, and the key driving factors were:

- Aging infrastructure
- Growing business needs beyond available storage, compute, and network capacity
- End of life of current storage solutions

This section discusses the challenges faced by BPTP and the value-add solution provided by IBM PureFlex System.

Existing Infrastructure

Chassis/Compute Nodes

BPTP’s previous architecture was based on an IBM BladeCenter® E chassis with 4xLS20, 4xLS21, 2xHS21 and 3xHS22V hosting SAP ECC 6.0 (Windows Server/SQL). The infrastructure supported SAP applications for both development and production environment, authentication servers, business applications, and exchange servers hosted in VMware ESXi 5.0-based virtual machines. These applications were hosted on approximately 120 virtual machines in a clustered fail-over environment.

Storage

BPTP's existing storage architecture comprised EMC and IBM DS4700 hardware, hosting all the data from the SAP application, databases, Microsoft Exchange server and Virtual Machine data.

Storage connectivity

Connectivity to EMC and DS4700 storage is provided using the redundant SAN switch modules of the BladeCenter E chassis.

Network

BPTP's previous network architecture comprised a Juniper L3 EX4500 10G-capable switch connecting to the BladeCenter E chassis, running at 1 Gbps capacity.

Management

Management of the chassis was through the Advance Management Module on the BladeCenter E chassis. Individual subsystems were managed through independent management entities.

Solution Provided

Chassis/Compute Nodes

The proposed solution involved the IBM Flex System chassis with x86-based compute nodes. The chassis consists of two redundant CMMs, one Flex System Manager and four Flex System x240 Compute Nodes. The entire IT workload of BPTP including all the business applications, and the SAP development and production environment were consolidated using 120 virtual machines hosted on x240 compute nodes running embedded VMware ESXi 5.1.

Storage

IBM Storwize V7000 storage is proposed to improve application flexibility, responsiveness, ease of management and simplify storage-based migrations of data. The storage consisting of SAP, business applications and virtual machines data were migrated to IBM Storwize V7000 through the storage virtualization feature. IBM Storwize V7000 Easy Tier® is configured to improve the responsiveness of business critical workloads like SAP.

Storage connectivity

IBM Storwize V7000 and the existing BladeCenter E is integrated along with the two SAN switches as part of the IBM Flex System Enterprise chassis to provide a redundant configuration. A round robin-multipathing policy is configured for effective distribution of load across the configured paths.

Network

IBM Flex System Fabric EN4093 10Gb scalable switch modules which are integrated with the Flex System Enterprise Chassis are used to improve the network bandwidth for BPTP's growing infrastructure needs. EN4093 delivers the scalability and performance needed, with forty-two 10Gbps ports for internal chassis data movement and ten 10Gbps ports for connectivity to the customer's backbone. BPTP's solution uses the first partition (10 ports) and two external ports on each EN4093 switch to upgrade the existing infrastructure to end-to-end 10G networking. The VMReady feature of the EN4093 switches is configured to ensure that the virtual machine's network policies (VLAN, ACL, QoS) are not affected by relocation of the VM from one host to another. BPTP has integrated VMReady to configure four VLANs to support a DMZ connection, internal business communication, R&D communication and management network data movement. These switches are also capable of providing 40Gb uplink connections. The external ports connect to the existing Juniper L3 switch in a redundant configuration. Each compute node has two Ethernet ports operating at 10Gbps configured in NIC teaming mode to provide a redundant and load-balanced configuration.

Management

The IBM PureFlex System provides two levels of management options. At the chassis level is a redundant Chassis Management Module to monitor and manage all the chassis components such as compute nodes, fans, and power modules. Flex System Manager is included as a management appliance to provide advanced management capabilities for compute nodes, storage, and network. Flex System Manager features like virtualization

management (IBM VMControl™), chassis management, automation management, remote control, storage management, and service and support manager, are configured at the BPTP environment.

BPTP Value Proposition

- Smooth migration of the 120+ VMs hosted on thirteen servers (4xLS20, 4xLS21, 2xHS21 and 3xHS22V) to three Flex System x240 compute nodes. The VMs were upgraded to VMware ESXi 5.1
- Eight Flex System Enterprise Chassis slots are available to expand the compute nodes. BPTP can integrate additional x86 nodes or Power nodes depending on the business requirements.
- The SAP development and production environment was successfully migrated from a physical to a virtual environment running in an HA cluster mode.
- Exchange Server and SAP application performance increased considerably, as reported by end users and administrators.
- BPTP's VMs are all based on VMware BPTP has acquired an additional x240 node to research KVM as an option for the future.
- IBM Storwize V7000's virtualization feature enabled a smooth migration of all the data from existing EMC and IBM System Storage® DS4300 storage.
- One of the key highlights of the BPTP solution was Storwize V7000's Easy Tier functionality, which provided incredible SAP application performance (10,000-15,000 IOPS in 1-2ms).
- IBM Storwize V7000's management console is extremely user friendly and easy to learn.
- IBM Flex System Fabric Switch EN4093 was used to increase the network bandwidth from 1G to 10G. The switch allows for 40G support for future extensibility.
 - Flex System Fabric EN4093 switch increases network performance and flexibility, optimized network virtualization with virtual NICs, and simplified the network infrastructure.

- Flex System Fabric EN4093 switch has been configured for VMReady function where in network policies migrate automatically along with virtual machines (VMs) to enable that security, performance, and access to remain intact as VMs move from server to server.
- The BPTP solution aims for redundancy between compute nodes and chassis switches, chassis switches to the backbone network, redundancy in storage access, and allows for redundancy at the VM level.
- The Service and Support Manager feature in Flex System Manager has allowed auto logging of support calls, and the automation manager sends e-mail alerts for critical events. The customer can also use a mobile application on smart phones to monitor the chassis status.
- One Flex System Manager is able to manage up to four Flex System Enterprise Chassis

Conclusion

IBM PureFlex System offers a unique holistic solution for resource consolidation, workload optimization and a simplified deployment strategy to help clients deliver resources and infrastructure efficiently and securely.

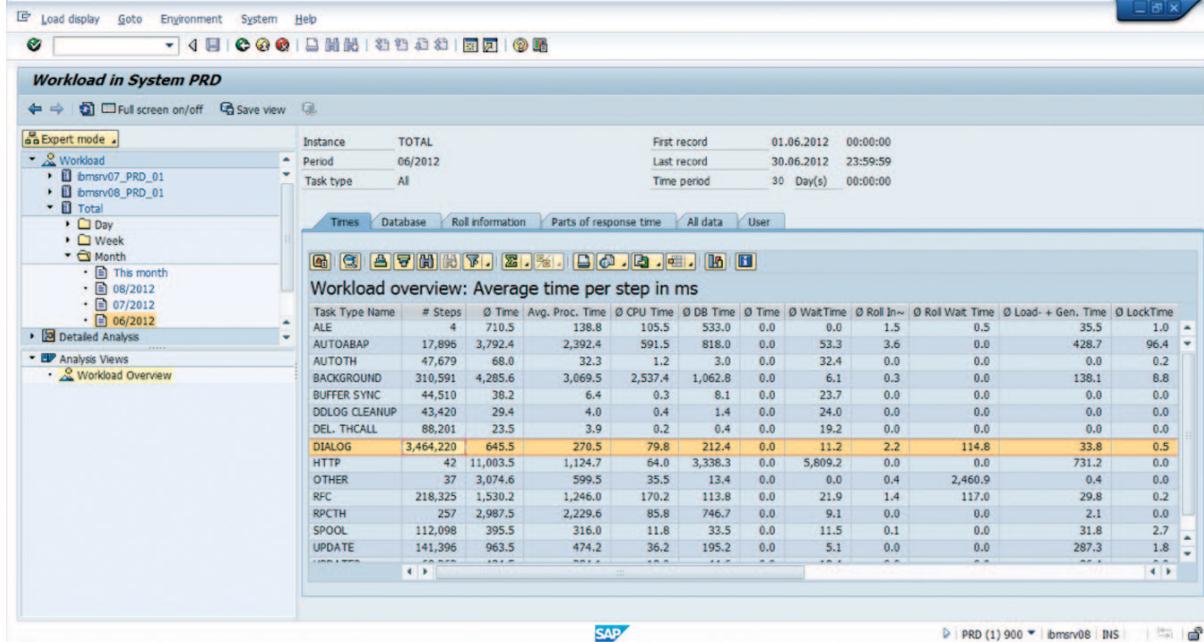
Appendix A: Supporting Materials

Illustration 1:

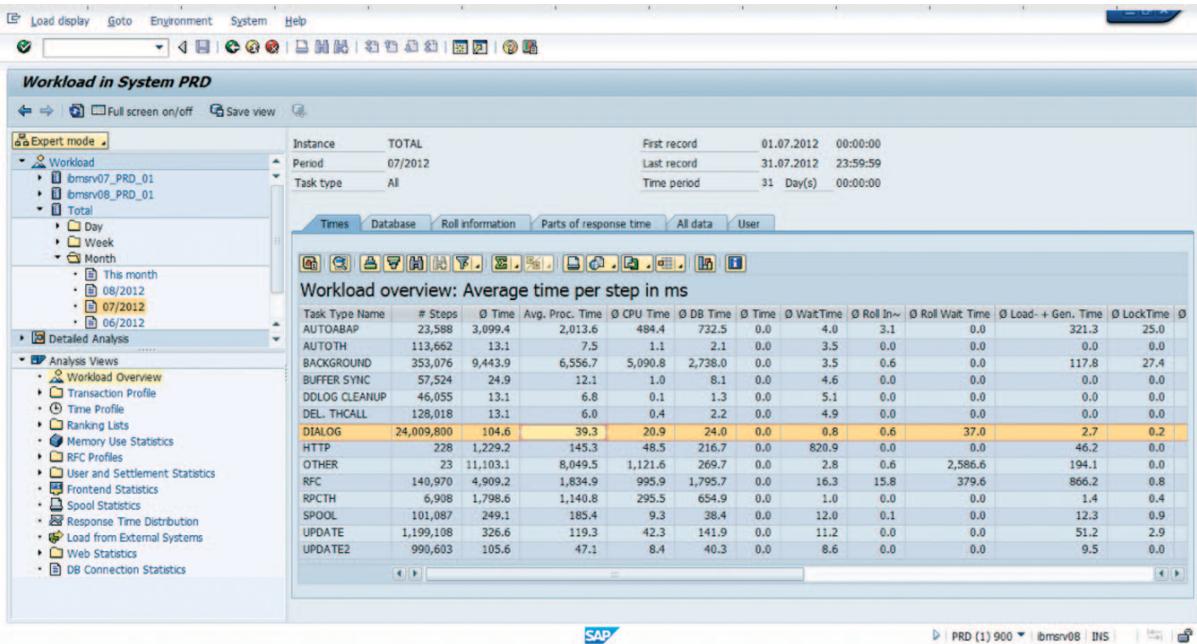
The following illustration shows the significant improvement averaged across a month in Dialog response time and SAP DB updates in the BPTP environment.

Summary of Performance metrics as shown below (less is better):

- Dialog Avg. processing time spread across a whole month: 270.5ms – >39.3ms (**85% decrease in processing time**)
- SAP DB update change over a month 474.2ms – >119.3ms (**75% reduction in DB update time**)



Data 1: BPTP performance metrics collected in month of June 2012



Data 2: BPTP performance metrics collected in month of July 2012

Illustration 2:

The following illustration shows significant improvement in BPTP workload on migrating to IBM PureFlex compute nodes. SAP reports have indicated 48% performance improvement compared to the previous BladeCenter infrastructure.

T-code	Report Name	Mode	Duration (22.06.2012) in Sec	Duration (25.06.2012) in Sec	Performance jump %
ZSDR031	TCODE for CPR Report	Background	1,225	626	48.90
T-code	Report Name	Mode	Duration (22.06.2012)	Duration (25.06.2012)	Performance jump %
ZOCR	Other Charges Register	Background	1,211	768	36.58
T-code	Report Name	Mode	Duration (22.06.2012)	Duration (25.06.2012)	Performance jump %
ZSDR007	New Booking Report	Background	862	453	47.45
T-code	Report Name	Mode	Duration (22.06.2012)	Duration (25.06.2012)	Performance jump %
ZFIR002_CHARGE	Charge Collection	Background	2,964	2,243	24.33
T-code	Report Name	Mode	Duration (22.06.2012)	Duration (25.06.2012)	Performance jump %
FBL5N	Customer Line Items	Background	8,841	5,446	38.40

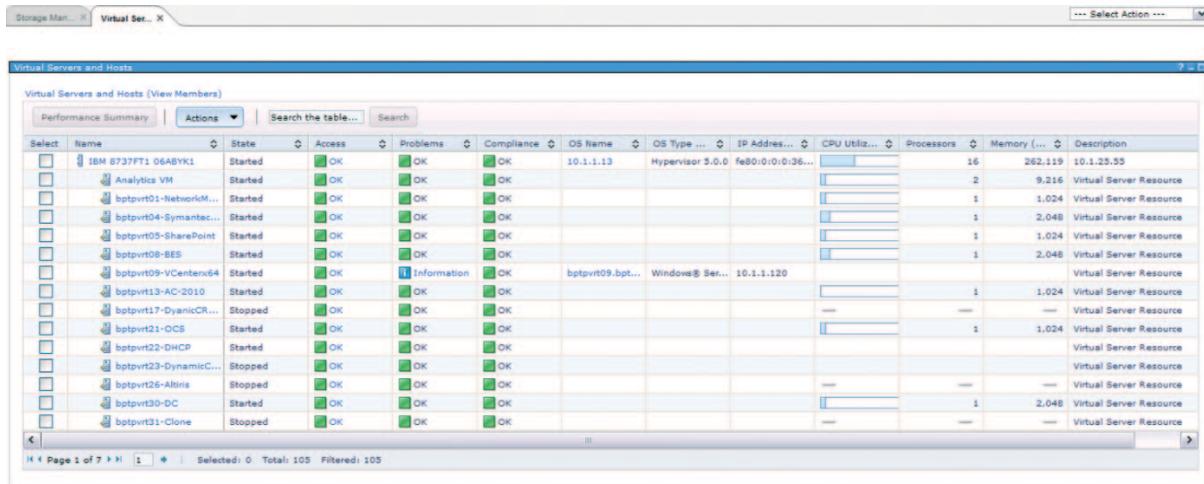
Illustration 3:

The IBM Flex System Manager web interface provides several ways to view chassis properties, configure chassis settings, and view chassis events and problems. The Chassis Manager view in the management software is the primary interface for selecting managed resources and performing management tasks on those resources. A selected BPTP chassis management is shown in the following illustration.



Illustration 4:

BPTP has deployed multiple VMs in their environment. The illustration below shows a tabular view of the chassis manager which provides management of all physical and virtual servers in BPTP's environment. This view provides performance monitoring and VM life-cycle management along with relocation of VMs and automation plan creation.



The screenshot shows a software interface titled "Virtual Servers and Hosts". The main area is a table with columns: Select, Name, State, Access, Problems, Compliance, OS Name, OS Type, IP Address, CPU Utiliz..., Processors, Memory, and Description. The table lists various VMs and hosts, such as "IBM 8737PT1 06ABYK1", "Analytics VM", and "bptprt01-NetworkM...". Most entries show "OK" status in the Access and Problems columns. The "Description" column indicates they are "Virtual Server Resource". The table has 105 rows and 13 columns. At the bottom, there are navigation buttons for page 1 of 7, and a status bar showing "Selected: 0 Total: 105 Filtered: 105".

Appendix B: Resources

These websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
ibm.com/partnerworld/systems
- IBM Publications Center ibm.com/e-business/linkweb/publications/servlet/pbi.wss?CTY=US
- IBM Storwize V7000 ibm.com/storage/storwizev7000
- IBM Redbooks® ibm.com/redbooks
- IBM TechDocs Library
ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

For more information

To learn more about the IBM PureFlex System, please contact your IBM representative or IBM Business Partner, or visit the following website: ibm.com/systems/pureflex/

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. We'll partner with credit-qualified clients to customize an IT financing solution to suit your business goals, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward.

For more information, visit: ibm.com/financing

About the authors

Rashmi Narasimhan—*Architect Security & Networking, Service Processor (IBM)*

Vikas Bhardwaj—*Solution Architect, Lab Services (IBM)*

Manish Misra—*Program Director Systems Firmware (IBM)*

Vilakshan Jakhu (Will)—*Sr. Vice President & CIO (BPTP)*

Acknowledgements

Special thanks to Rene Neumann—Market Segment Manager, PureFlex System Tim Alpers—Product Manager, PureFlex System, Rajesh Sehgal—India Sales Leader, PureSystems, Venkatesh Sainath, IBM STG, Pradipta Banerjee, IBM STG and Vishal Kulkarni, IBM STG for reviewing this White Paper. Thanks to the BPTP team and the IBM team for their contribution and support to this project.



© Copyright IBM Corporation 2012

IBM Corporation
Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States of America
October 2012

IBM, the IBM logo, ibm.com, BladeCenter, Easy Tier, Flex System Manager, IBM Flex System, POWER7, Power Systems, Redbooks, Storwize, VMControl, VMReady, POWER, PureSystems, PureFlex, and PowerVM are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

Microsoft, Windows, Hyper-V, and the Windows logo are registered trademarks of Microsoft Corporation in the United States and/or other countries.

LINUX is a registered trademark of Linus Torvalds.

Intel and Xeon are registered trademarks and MMX, Pentium II Xeon and Pentium III Xeon are trademarks of Intel Corporation in the United States and/or other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.



Please Recycle