Electric Green

An IBM *Project Big Green* Solution

Authors: High Performance On Demand Solutions (HiPODS) team
Technical contacts: Nimesh Bhatia, Stewart Edelman, Chin Huang, Marty Yarnall
Management contacts: Jinzy Zhu, Tom Pavela, Michael Azzi
Web address: www.ibm.com/websphere/developer/zones/hipods

Date: May 25, 2007

Status: Version 1.0

Abstract: This paper describes the High Performance On Demand Solutions (HiPODS) request driven provisioning (RDP) solution. RDP is based on a composite service-oriented architecture pattern and provides IT infrastructure management services for multiple, globally distributed data centers. With its power saving services, RDP helps enterprises increase the level of energy efficiency in IT.

To receive information about future workshops and seminars pertaining to this solution, please send an email to **hipods@us.ibm.com**.

Executive summary

Technology analyst IDC reports that energy costs consume roughly half of every dollar spent on computer hardware, a number that's expected to increase 54 percent, to 71 cents, over the next four years[1].

IBM, who currently runs one of the world's largest commercial technology infrastructures with more than eight million square feet of data centers in six continents, is battle testing a new initiative it calls *Project Big Green*. Using new approaches, IBM expects to double the computing capacity of its data centers within the next three years without increasing power consumption or its carbon footprint. Based on greenhouse gas production, the size of a carbon footprint is determined by the amount of carbon dioxide that comes from human activities. Compared to doubling the size of its data centers by building out new space, IBM expects this could help save more than five billion kilowatt hours of energy per year.

The IBM High Performance On Demand Solutions (HiPODS) team developed request driven provisioning (RDP), a solution that automates a subset of IBM service management processes. HiPODS is aggressively developing an RDP extension called *Electric Green* to support Project Big Green. Automation provisioning software such as Electric Green can reduce power consumption by 80 percent on servers by automatically putting them on standby mode when they are not needed. If this software was deployed in all the estimated U.S. data centers, the country could save 5.4 billion kilowatt hours per year, enough electricity to heat 370,000 homes for a winter.

This paper introduces Electric Green, primarily RDP and its planned power conservation techniques. RDP is an integration of assets and products based on a service-oriented architecture and Web services interfaces. It supports the full life cycle of a service delivery from order placement through final reporting.

RDP highlights the coordination of the different roles involved in this full cycle scenario. It automates capacity planning based on business objectives, the deployment topology determination and provisioning plan creation; it enforces SLA terms through orchestration and provides monitoring, reporting, and power saving services.

RDP can provide automated infrastructure management, thereby offering these benefits:

- Customers can focus on their core business instead of IT
- Customers have fast access to computing capacity and IT services
- Customers incur costs based on usage and do not have to invest in IT infrastructure, thus allowing them to respond quickly to varying economic environments
- Data centers can save energy costs by monitoring the system usage data and apply appropriate power saving techniques

More information can be found at:
www.ibm.com/developerworks/websphere/zones/hipods/library.html.

---

[1] Source: IDC, *Worldwide Server Power and Cooling Expense 2006–2010 Forecast*, Doc #203598, September 2006

Contents

Introduction

For seven years, the IBM High Performance On Demand Solutions (HiPODS) team has been engaging with customers seeking solutions to IT infrastructure challenges. Our success with these projects depended on a high degree of innovation, fast turnaround, and containment of technical risk due to the new technology components in the solutions. HiPODS ensured innovative solutions and approaches to the solutions by collaborating closely with customers and by building innovative teams consisting of the best subject matter experts from IBM research, hardware, software, and services groups. These teams leveraged the state-of-the-art HiPODS hardware and software lab at IBM Silicon Valley developing, testing, and ensuring deployment readiness of the solutions. The successes of projects have led to growth in the number of projects, including a dramatic increase in projects outside of North America. This growth has led HiPODS itself to experience challenges in IT infrastructure management.

Request driven provisioning

Request driven provisioning (RDP) automates a subset of the IBM service management process; in other words, it's a pattern based on IBM service management. RDP is an end-to-end service delivery solution that supports the on demand operating environment and integrates existing assets and products based on SOA. In general, SOA is critical for enterprises to execute the on demand vision and prepare for the incremental changes over time. RDP is an instantiation of a service delivery solution based on SOA implemented in Web services, providing a flexible and efficient solution for offering business services that operate in a service-oriented environment. It supports the full life cycle of a service delivery from offering creation through order placement, capacity estimation, contract fulfillment, monitoring, and reporting.

Engaging with customers on a daily basis, HiPODS is able to leverage the value of RDP in every step of the process. RDP is an IBM service management pattern based on information technology infrastructure library (ITIL) standards helping organizations better manage their IT infrastructure and more effectively and efficiently deliver IT services. RDP is also capable of leveraging the IBM service management change management process manager to deliver infrastructure provisioning.

The RDP solution is designed to help define services more conveniently and manage IT infrastructures such as data centers more efficiently and cost effectively. It generalizes service definition, virtualizes the infrastructure, and automates infrastructure management, thereby offering these benefits:

- Shorter time to market for making new services available
- Streamlined business process to deliver a service
- Cost effective sizing of the infrastructure needed to run business applications
- Maximized resource utilization with resource virtualization and reservation
- On demand computing capacity and IT services to customers
- Power savings by tracking system usage and employing power saving techniques

Shown conceptually in Figure 1, RDP is composed of a suite of loosely coupled components (core components described below) that are integrated through predefined data and process models. These components can be selectively customized, composed, and deployed to provide greater modularity and consumability to satisfy various business requirements.

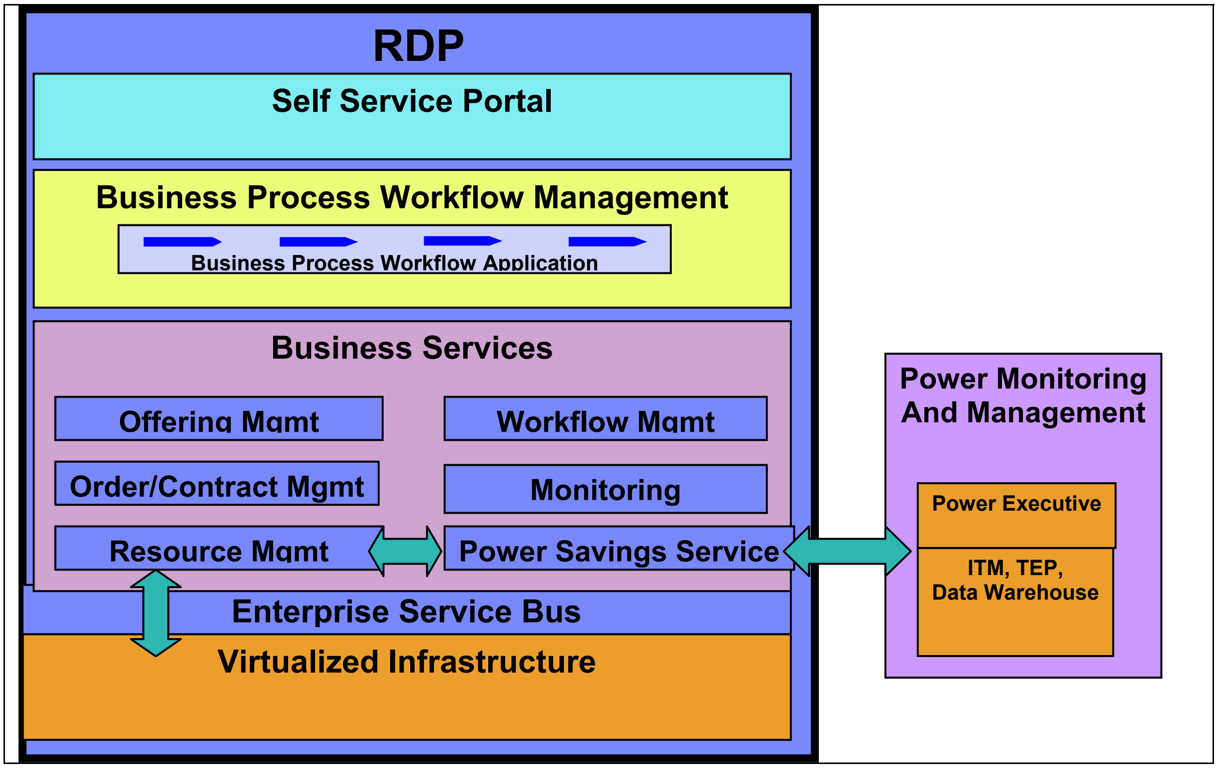

Figure 1. RDP components

The **resource management** component manages hardware allocation and software deployment. The hardware allocation functions include provisioning of servers, storage, and network bandwidth, as well as the configuration set up of supporting equipment such as switches and routers. The software deployment functions include operating system builds, BIOS updates, application installs, and application configuration. It also has a reservation and scheduling system that supports ordering of infrastructure resources for future consumption.

The **offering management** component consists of the repository of the available services, information about subscriptions of those resources, and business processes to manage offerings. It maintains a taxonomy of services that enables efficient searching of services for subscription and composition.

The **order/contract management** component consists of the repository of the service requests and contracts, a state machine to ensure relationship and integrity between service requests and contracts, and searching capability for quick access to those objects.

The **process workflow management component** manages communication between components at the process level and enforces an overall choreography.

The **service monitoring component** collects and provides system and application performance data for the individual services and their composition. Both real time and historical data are provided to help IT management to understand and analyze system usage, and adjust system allocation when desired. Monitoring provides critical data for other IT management components such as metering, SLA management, and problem determination tools.

The **power savings services component** is the data mining and analysis tool to find out average server usage, based on different filters and criteria and taking some power saving actions. It uses the IBM® Tivoli® Monitoring to collect the data and apply the filtering capabilities to find the servers least used and their average utilization during the day or over a week. A data center administrator can select the servers and collaborate with project teams to return the unused servers to the active inventory by de-provisioning them or schedule the workflows to place the servers on standby for off-peak hours using IBM ®Tivoli® Provisioning Manager.

Figure 2 illustrates the flow of the typical service delivery process.

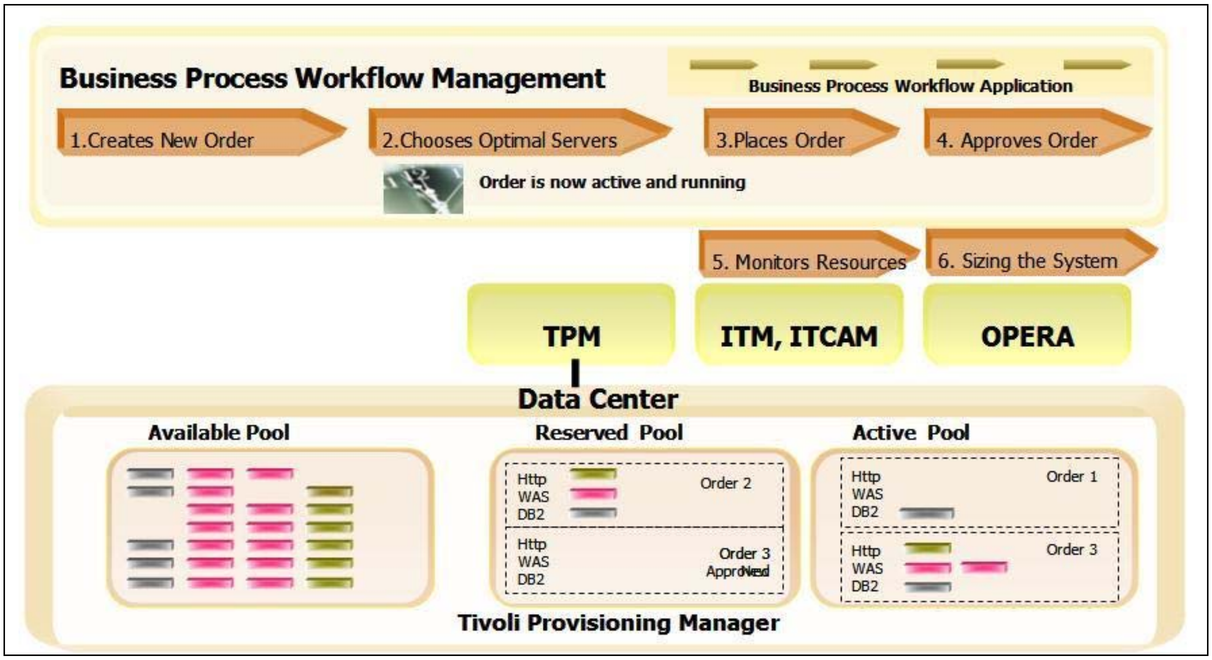

Figure 2. Flow of typical service delivery

Techniques for saving power

Considerations during capacity planning

Power consumption should be an important consideration during the initial phase of capacity planning. Different sets of servers can be used during the phases of development; the design, development, and functional test phases may require less heavy computing capabilities than performance testing and production environments. Server selection for each phase should include a consideration of power consumption.

Capacity planning tools can give information about power consumption when returning recommended infrastructure topologies based of different service level agreements. Figure 3 shows sample results from such a tool.

Vertical Solution

Total Power Consumption **1041-1073 Watts**

	Brand	Model	OS	# of Nodes	Disk Access(ms)	# of disks
Tier 1	xSeries	x206 2-way 2800	Linux	1	13.0	4
Tier 2	xSeries	x306 1-way 3200	Linux	1	13.0	16

Vertical Solution Details (Click to expand)

Horizontal Solution

Total Power Consumption **711-742 Watts**

	Brand	Model	OS	# of Nodes	Disk Access(ms)	# of disks
Tier 1	xSeries	x3800 1-way 3660	Linux	2	13.0	4
Tier 2	xSeries	x306 1-way 3200	Linux	1	13.0	16

Horizontal Solution Details (Click to expand)

Minimal Solution

Total Power Consumption **1086-1132 Watts**

	Brand	Model	OS	# of Nodes	Disk Access(ms)	# of disks
Tier 1	xSeries	x306 1-way 3200	Linux	1	13.0	4
Tier 2	xSeries	x306 1-way 3200	Linux	1	13.0	16

Minimal Solution Details (Click to expand)

Figure 3. Sample capacity planning results displaying power consumption

Continuous power monitoring and analysis during project lifetime

Any project goes through different cycles during its life time. System resources go from being fully utilized during some phases, or time of the day or part of a week, to minimally utilized at various instances.

IBM Tivoli Monitoring can be used to track various parameters and activities of server utilization. Data collected should be analyzed and the trend should be identified. When the time period and the underutilized servers are determined, action can be taken to conserve energy.

Electric Green offers an add-on service option to perform trend analysis and filtering on various parameters on the data collected by Monitoring Service. Figures 4 and 5 show output samples from trend analysis and filtering.

Filtered sets of servers can be tagged for further monitoring and Tivoli Provisioning Manager workflows can be scheduled to put the selected servers on standby mode during off-peak hours.

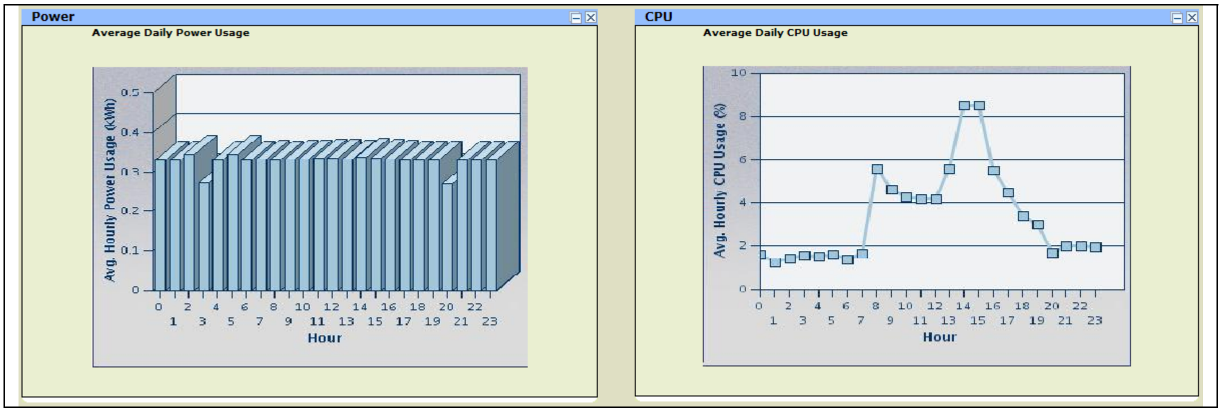

Figure 4. Sample trend charts for CPU utilization and power consumption on average day

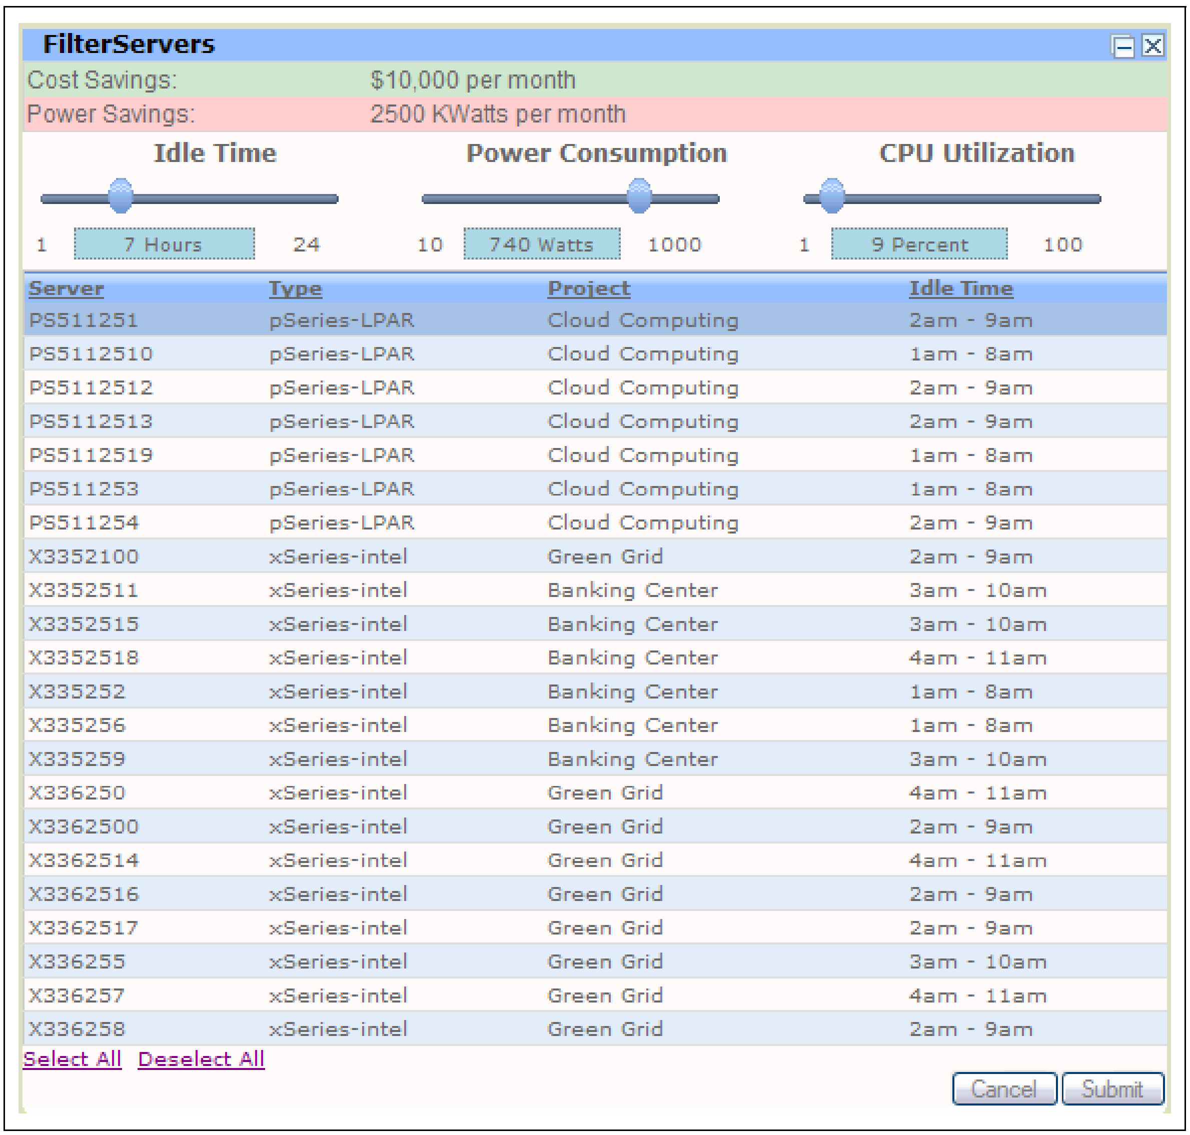

FilterServers

Cost Savings: $10,000 per month

Power Savings: 2500 KWatts per month

Idle Time	Power Consumption	CPU Utilization
1 — 7 Hours — 24	10 — 740 Watts — 1000	1 — 9 Percent — 100

Server	Type	Project	Idle Time
PS511251	pSeries-LPAR	Cloud Computing	2am - 9am
PS5112510	pSeries-LPAR	Cloud Computing	1am - 8am
PS5112512	pSeries-LPAR	Cloud Computing	2am - 9am
PS5112513	pSeries-LPAR	Cloud Computing	2am - 9am
PS5112519	pSeries-LPAR	Cloud Computing	1am - 8am
PS511253	pSeries-LPAR	Cloud Computing	1am - 8am
PS511254	pSeries-LPAR	Cloud Computing	2am - 9am
X3352100	xSeries-intel	Green Grid	2am - 9am
X3352511	xSeries-intel	Banking Center	3am - 10am
X3352515	xSeries-intel	Banking Center	3am - 10am
X3352518	xSeries-intel	Banking Center	4am - 11am
X335252	xSeries-intel	Banking Center	1am - 8am
X335256	xSeries-intel	Banking Center	1am - 8am
X335259	xSeries-intel	Banking Center	3am - 10am
X336250	xSeries-intel	Green Grid	4am - 11am
X3362500	xSeries-intel	Green Grid	2am - 9am
X3362514	xSeries-intel	Green Grid	4am - 11am
X3362516	xSeries-intel	Green Grid	2am - 9am
X3362517	xSeries-intel	Green Grid	2am - 9am
X336255	xSeries-intel	Green Grid	3am - 10am
X336257	xSeries-intel	Green Grid	4am - 11am
X336258	xSeries-intel	Green Grid	2am - 9am

Select All Deselect All

Cancel Submit

Figure 5. Sample list of servers based on a specific set of filter criteria

Power conservation with deprovisioning

When a project is over, returning resources to the active inventory and shutting down servers until they are ready to be allocated for new projects can save significant power.

Towards the end of the cycle, projects go though a phase typically known as the *long tail*. The resources used in this cycle have low utilization. Virtualizing and consolidating those resources can save capital and energy.

Using some technologies from IBM Research to track power consumption, our studies determined that a powered-on blade server consumed about 230W/hr, a measurement that decreased to 40W/hr when we put the server on standby power consumption. After applying the above mentioned techniques in the HiPODS data center, which has over 1000 servers, we estimated that we would save $32.7K or 327,000 kilowatt hours per year.

Summary

Electric Green is a highly automated IT infrastructure provisioning service that can quickly deliver applications on standardized hardware, software, and applications at a low cost.

The service is composed of:

- Order management GUI: The service portlets provide the user-friendly experience to identify available infrastructure resource, specify hardware, software, and configuration parameters, and view service requests and contracts.
- Order processes: The service business processes allow the users to go through standard process flow for order approval, resource reservation, provisioning, and message notification.
- Infrastructure provisioning service: operating system provisioning (for example, Linux®, Microsoft Windows®, IBM® AIX®), software and middleware installation and configuration (for example, IBM® WebSphere®, IBM®DB2®)
- Virtualization support: IBM® System P5 LPAR resource virtualization (CPU, memory, storage)
- Storage provisioning: IBM® TotalStorage® SAN Volume Controller, IBM System P5 storage
- Security module: System security/patch management
- Power saving services: Power conservation by tracking utilization and applying power policies

With its unique ability to monitor and track system utilization and power consumption, set power saving policies, and put servers into standby mode when not in use, Electric Green provides power saving techniques to reduce power consumption for corporate data centers.

HiPODS is excited about contributing RDP provisioning and power savings technology to help IBM customers combat the data center energy crisis, and looks forward to deploying Electric Green to transform existing enterprise infrastructures in to green data centers.

References

To learn more about IBM Project Big Green, see www.ibm.com/press/us/en/presskit/21440.wss

See all the HiPODS white papers at www.software.ibm.com/wsdd/zones/hipods/library.html
Of special interest:

Innovation and collaboration	*Innovation Factory: An integrated solution for accelerating innovation* *Introducing HiGIG: The HiPODS Global Innovation Grid*
Network delivered Services	*Sonoma: Capacity & Performance Estimation Web Service for Service-Oriented Architecture (SOA) Workloads*
Business-driven IT	*On Demand Service Delivery for SOA Environments* *Case Study: SOA-based On Demand Service Delivery Solution at Star Technology Services Ltd.* *Automating to Become an On Demand Business*
Provisioning	*Techniques to Configure IBM Tivoli Intelligent Orchestrator to Manage IBM pSeries Servers* *Provisioning for On Demand Data Centers* *Provisioning Best Practices for On Demand Data Centers* *Provisioning Case Study: pSeries and AIX LPARs* *Provisioning Case Study: WebSphere Application Server Version 5*
Orchestration	*Orchestration for On Demand Data Centers* *Orchestration Case Study: Integrating WebSphere and DB2* *Orchestration with Enterprise Workload Manager* *Prepare your WebSphere Web Site for e-business on demand , An Update*
Virtualization	*Virtualization and Automation for a WebSphere Application Server Environment / An Update* *Using VMware ESX Server with IBM WebSphere Application Server*

Acknowledgements

The HiPODS Architecture Board (Lance Amundsen, Dave Coleman, Jeff Coveyduc, Catherine Cuong-Diep, Harold Hall, Chin Huang, Eugene Hung, Rahul Jain, Linda Legregni, Luis Ostdiek), Qi He, Larry Lange, Dikran Meliksetian, Noshir Wadia, Susana Meza, Mike Tuso

Notices

Trademarks
The following are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM
IBM logo
AIX
DB2
pSeries
Tivoli
TotalStorage
xSeries
WebSphere

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Other company, product, and service names may be trademarks or service marks of others.

Special notices
The information contained in this document has not been submitted to any formal testing by or for IBM, and the document is therefore distributed "AS IS" by IBM to the fullest extent allowed by applicable law. Any use of the information or implementation of any of the techniques contained in this document is a responsibility of the user. The results of any such use will depend on the user's ability to evaluate and integrate the information and/or techniques contained in this document into the user's operational environment.

While IBM may have reviewed each item for accuracy in a specific situation, IBM offers no guarantee or warranty to any user that the same or similar results will be obtained elsewhere. Any person attempting to adapt the techniques contained in this document to their own environment(s) does so at their own risk. Any performance data contained in this document were determined in various controlled laboratory environments and are for reference purposes only. Customers should not adapt these performance numbers to their own environments as system performance standards. The results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.