sageza

Strategic Snapshot

The Workload Optimized Approach
Why Do Systems Matter?

By Clay Ryder

The Sageza Group
November 2009

The Sageza Group
931 Campbell Avenue
Los Altos CA 94024
510·675·0700 fax 501·640·8105

sageza.com
info@sageza.com

The Workload Optimized Approach

ABSTRACT

The world is getting smarter as it is more instrumented, interconnected, and intelligent: we are living in a smarter planet. The data center of today is being called upon to deliver a dynamic and real-time IT infrastructure that can support business transformation as the smarter planet increases the scale, complexity, and diversity of workloads. Pervasive instrumentation will create vast amounts of data and drive new types of applications that require real-time data analysis and predictive intelligence.

Historic approaches bent business process around the limitations of technologies; today the expectation is for the technology to align with business process. The one-size-fits-all approach is no longer sufficient for organizations that are seeking to maximize their competitive advantage and ROI on IT investments. Organizations' strategic approach to IT must evolve to meet the corporate need not only of today, but of tomorrow as well.

In this paper, we examine the importance of IT's ability to deliver a dynamic, scalable, and optimized infrastructure that is tailored to an organization's workloads. We will review how IBM is embracing the smarter planet by delivering a Dynamic Infrastructure upon which workload-optimized solutions can be deployed to effectively meld business process with IT infrastructure. We will also highlight an example workload-optimized solution that addresses the vexing challenge of traffic congestion management in the inner city. With its workload-optimized approach, IBM has positioned itself to address organizations' search for an IT infrastructure that not only meets the present need, but can act as a catalyst to drive business transformation and enhanced competitive positioning in the marketplace.

The Workload Optimized Approach

TABLE OF CONTENTS

Introduction 1

The Need for Responsive and Realtime IT Infrastructure 1

Workloads: Old, New, and Ever Changing 2

- Four Categories of Everyday Workloads 2

IBM's Dynamic Infrastructure 3

IBM Architecture for Workload Optimization 4

New Workload Optimization Product and Service Offerings 5

- Enhanced Virtualization Management 5
- New Datacenter Networking Hardware 5
- New Workload Optimized Systems 5
  - Industry Solutions 6
  - Smart Business Systems 6

Case Study: Smart Traffic 6

What It All Means 8

Introduction

The traditional approach to IT infrastructure has been to meld business processes and IT resources at point-in-time, with the delivery of supporting IT services not varying greatly over the lifetime of the IT solution. In this scenario, a widening disconnect between current business practices and the IT infrastructure is inevitable and organizations typically would attempt to narrow this gap when IT infrastructures were refreshed. Such an approach, while existing within the technological constraints of the era, does not meet the needs of organizations today.

> The world is getting smarter as it is more instrumented, interconnected, and intelligent.

The world is getting smarter as it is more instrumented, interconnected, and intelligent. As a result, there is a significant shift in organizations' assumptions about IT and the strategic approach to IT investments. The data center of today is being called upon to deliver a dynamic, realtime, IT infrastructure that can keep up with the demands for information and speed of delivery. Business is transforming as a smarter planet increases the scale, complexity, and diversity of workloads that most any organization will be called upon to support. Pervasive instrumentation will create vast amounts of data and drive new types of applications that require realtime data analysis and predictive intelligence.

Organizations' strategic approach to IT infrastructure must evolve to meet the corporate need not only of today, but of tomorrow as well. Historically, workloads often have been deployed in a siloed fashion. While there may have been organizational and financial biases within business that favored this approach, the resultant infrastructure suffered from an inherent inefficiency in operations combined with high demands on floor space, energy consumption, and cooling capacity. In addition, the sharing of applications and data across the corporation was often difficult in such a scenario, which leads to underutilized assets.

Rigid technological constraints combined with complexity made change difficult to achieve which resulted in increased operational expense and a competitive disadvantage for organizations. Further, this has challenged regulatory compliance and security initiatives as well as limiting organizations' efforts to implement industry best practices. In a nutshell, businesses are unnecessarily being exposed to risk.

Historic approaches bent business process around the limitations of technologies; today the expectation is for the technology to align with business process. The one-size-fits-all approach no longer supports what is needed. As a result, the current IT infrastructure deployed in most organizations was not designed to meet their current needs let alone meet the future needs of business.

The Need for Responsive and Realtime IT Infrastructure

To meet their current and future business requirements, organizations must change their IT strategy to an expectation of a dynamic IT infrastructure that seamlessly blends business process with IT resources to enable the making of informed, realtime decisions at the point of business impact. Key capabilities are requisite for organizations seeking to implement a dynamic infrastructure:

- It is necessary to create Workload Optimized Systems that more effectively align technology capability with the business need. This is in sharp contrast with the general-purpose or one-size-fits-all approach of the past.

- Organizations will be called upon to manage service delivery across a larger number of more instrumented and interconnected business and IT processes and assets. It is essential that the IT department be able to take a more integrated approach to Service Management.
- It is essential to fundamentally change storage economics to enable information-led growth without scaling the cost of storage management out of reach. To this end, organizations will need to be able to leverage emerging cloud-enabled information services and innovative storage technologies.

Long-term corporate viability and competitive success dictate that IT solutions will need to seamlessly integrate visibility, control, and automation across all business and IT assets. Computing solutions must be highly optimized to do more with less. Integrated service management is necessary to provide visibility, control, and automation across business and IT services to ensure consistent, high-quality delivery.

Workloads: Old, New, and Ever Changing

There are many kinds of workloads: old, new, and those yet to be implemented, each with its own unique IT requirements. Modern demands being placed on old workloads as well as the expectations for emerging new workloads require smarter and more versatile solutions. The degree of instrumentation now available is fundamentally changing the scope, scale, and complexity of data being collected and processed by organizations. For example, Business Analytics is one burgeoning new workload that organizations are increasingly deploying to take advantage of the wealth of instrumentation now available in order to create a competitive advantage.

The demand for data storage continues at a seemingly relentless pace. Beyond the relatively simple matter of acquiring sufficient storage capacity looms the much larger challenge of effectively managing increase in data volume, understanding its content and context, and ensuring compliance with regulatory and industry best practices. The growth of unstructured data such as images, email, audio, video, etc. is geometric and its effective management and seamless inclusion into business processes is essential for corporate success.

In today's global economy, there are no "off-hours"…

Another way in which workloads are changing is that historically batch-focused tasks are now being called upon in a realtime environment. For example, account settlement used to be considered a once-a-day activity typically completed after business hours. However, in today's global economy, there are no "off-hours," and more importantly, account settlement may be just one part of an answer to larger business question. The intertwining and leverage of information is by nature realtime and dynamic; the enterprise can no longer afford to have information that is minutes, let alone hours or days removed from reality.

Four Categories of Everyday Workloads

No two workloads are exactly the same. Each has a unique requirement for IT whether it is computation, storage, bandwidth, indexing, or any number of other resources. Nevertheless, at a higher level we note that there are four broad categories of workloads, each of which has a given resource affinity. These four categories and their requirements are as follows:

- Transaction Processing and Database workloads are typified by their need for scalability, high transaction rates, high quality of service, the ability to handle peak workloads, and the highest levels of resiliency and security.

- Analytics and High Performance Computing workloads are generally compute-intensive, require high I/O and memory bandwidth, and are floating-point calculation intensive.
- Business Applications workloads require scalability, high quality of service, a large memory footprint, and a responsive infrastructure.
- Web, Collaboration, and Infrastructure workloads are typically highly multi-threaded and throughput-oriented, assume rapid scaling, and can operate successfully at a lower quality-of-service threshold.

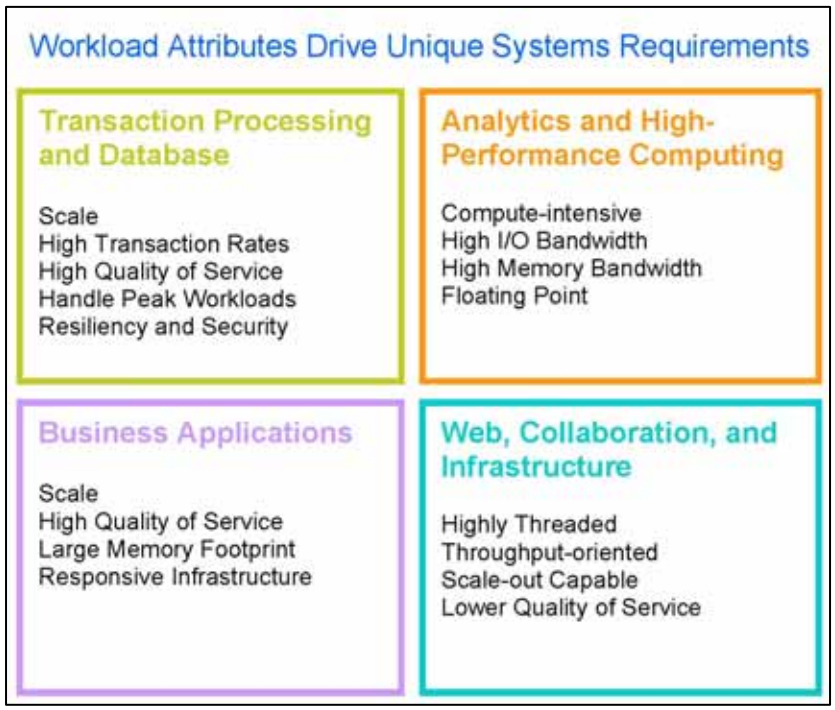

Figure 1: Categories of Everyday Workloads

It is clear from these simple categorizations that a one-size-fits-all approach to infrastructure would result in far too many compromises, thus artificially stymieing the success of any workload. Although infrastructure standardization is a laudable goal, the value of standardization is manifest best at the workload level, not at the physical server. The reality of new workloads and ever more competitive business environments dictate that the IT infrastructures of the past are no longer up to the task. In essence, that which was "good enough" is no longer "good enough."

IBM's Dynamic Infrastructure

To address the needs of organizations to implement IT solution that mirror their business objectives, IBM articulated its Dynamic Infrastructure vision for the next generation of dynamic and real time business focused computing. The Dynamic Infrastructure strategy includes a high-level overview of the IT challenge while providing a ground-level response for day-to-day operations delivering a responsive and realtime IT environment that meets organizations' current and future computing needs. Dynamic Infrastructure addresses corporate IT challenges holistically:

- It is designed for a smarter planet to help organizations address the increasing cost and complexity of IT infrastructure.
- It links and manages all IT and business assets.
- It makes business and IT infrastructure as dynamic as the organization demands.

Dynamic Infrastructure addresses the new information challenge while proactively managing risk to the enterprise from loss of information integrity, security, and regulatory compliance. Solutions can be delivered in a flexible fashion using a variety of platforms in any combination of inhouse solutions, private clouds, and the public cloud.

Within the Dynamic Infrastructure framework, IBM's Workload Optimized approach allows organizations to transform all of their new IT infrastructure investments into a state-of-the-art dynamic and business-focused solution. By providing systems for a smarter planet, the organization's business demands come first, not the technology's limitations. With a sharp focus on Service Management IBM customers can benefit from new and enhanced capabilities that provide greater visibility, control, and automation of their IT environment which in turn enhances organizations' ability to decrease time-to-market for new competitive services. Collectively this delivers the opportunity for businesses to deploy a new information infrastructure focused on cost-reduction, information-led IT, and business transformation.

IBM Architecture for Workload Optimization

Workload Optimization by definition is not a one-size-fits-all approach; rather it seeks to create the best fit of IT infrastructure with business process. As a result, a workload-optimized approach can deliver greater performance, scale, and efficiency to corporate operations. As we see in figure 2, workload optimization is an iterative process whereby the workload drives the system requirements which in turn impact system design which dictates further optimizations that can influence the workload attributes, and the cycle continues.

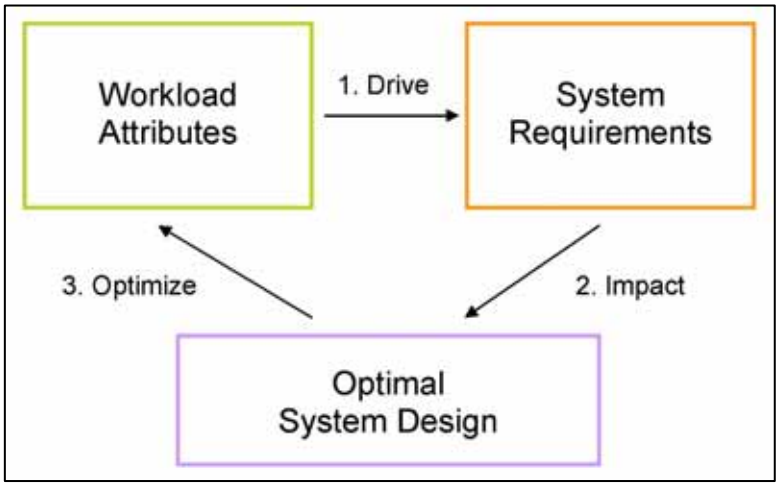

Figure 2: Workload Optimization Cycle

IBM delivers Workload Optimization choice to its customers through three complementary approaches:

- Systems Choice provides organizations a variety of hardware platforms and software solutions through which to best match specific business workloads.
- Purpose Built Solutions are deeply tuned solutions that are designed to meet specific requirements and deliver a closely tailored solution for specific workloads.

- Integrated Workloads are achieved through integrated heterogeneous systems that optimize multiple workloads across the entirety of the IT infrastructure.

IBM is uniquely positioned to create workload-optimized solutions as the company has extensive vertical expertise of business processes and IT workloads combined with a deep technical knowledge of hardware and software components along with a choice in hardware platforms that can best match customers' needs. As part of IBM's Dynamic Infrastructure strategy, organizations have the opportunity to rethink how they approach IT service delivery and undertake a simplified, flexible, and future-ready approach for delivering IT services with better ROI and lower TCO.

New Workload Optimization Product and Service Offerings

On October 20, IBM announced several new product and service offerings related to Workload Optimization. Each of these new offerings delivers an important component supporting a holistic, optimized, approach to IT.

Enhanced Virtualization Management

One of the most important aspects of modern day IT is the virtualization and management of resources through the enterprise. As IBM Systems Director begins delivering VMControl capability in 2009 and 2010 it will increasingly condense over seventy controls that will allow simpler policy-based automatic control of physical and virtual servers. This should benefit organizations that to date have had to discretely manage their virtual servers from their physical instantiations with improved efficiency through a single management interface. Management simplification is also achieved through the integration of IBM Systems Director Network Manager and the Tivoli networking products. Virtual System Pools, which are enabled through VMControl Enterprise Edition, seek to ease system platform and resource management through interoperation with the existing service management platform to enable organizations to be able to reconfigure and scale for dynamic workloads especially within a cloud computing context.

New Datacenter Networking Hardware

Networking technology is a key enabler of modern IT.

Networking technology is a key enabler of modern IT. IBM's most recent network offerings feature Ethernet and Fibre Channel over Ethernet (FCoE) capabilities in addition to traditional Fibre Channel. The new Ethernet offerings, result from partnerships that IBM has established with network equipment providers Juniper and Brocade, and the FCoE offerings are from Brocade and Cisco. There is also a new Fibre Channel router as well as new router and FCoE blades available for IBM Fibre Channel directors. In addition, there is a new IBM BladeCenter Virtual Fabric that enhances the value proposition of Ethernet, Fibre Channel, and Fibre Channel over Ethernet. This new solution can reduce hardware acquisition costs while delivering Virtual Network Interface Card (NIC) and I/O convergence.

New Workload Optimized Systems

Workload optimization implies that solutions will tailor their configuration around specific resource utilization patterns, which can be expressed either in technological terms or industry vertical emphasis. What follows are a few examples of the workload optimized systems available from IBM.

Industry Solutions

There are several segments within the retail marketplace each with its own IT infrastructure needs. IBM delivers workload-optimized retail solutions through its Green Store infrastructure that integrates servers, storage, and networking within IBM BladeCenter systems to help retailers reduce complexity and improve systems management. These retail store solutions can support a variety of commercial establishments such as quick-service restaurants and major grocery retailers. In addition, IBM Retail Systems Management integrates Remote Management Agents with the IBM Systems Director console to provide a set of tools designed specifically for retail that can provide front-of-store and back-of-store management from one central location.

IBM System Cluster 1350 is a family of integrated solutions based on the IBM System x rack, BladeCenter, and iDataPlex server that can scale from the departmental cluster level up to super-computer implementations. These end-to-end solutions integrate IBM server, system clustering, and Tivoli technologies. System Cluster 1350 is available with optimized applications for industrial design and manufacturing, financial services, life sciences, government, and educational institutions. In addition, the platform is well suited for applications that require horizontal scaling, such as Web serving and collaboration.

Smart Business Systems

IBM Information Archive is a next-generation information retention solution that is designed as a universal repository for structured and/or unstructured content to help organizations address complete information retention needs as required by business, legal, or regulatory requirements. This mixed-media storage repository provides up to three customizable information collections with three information protection levels in a storage optimized environment that supports data deduplication and compression. Data integrity is maintained until deletion is permitted by retention policy and offers enhanced security through a data encryption option as well as a patent-pending Enhanced Tamper Protection feature.

IBM CloudBurst 1.2 is the latest version of a prepackaged and self-contained integrated service delivery platform of server, storage, networking, virtualization, service management software, and Quickstart services designed to enable private cloud, self-service environments. Built on the IBM BladeCenter platform, CloudBurst is a fault-tolerant solution that scales based upon changing business needs, supports the management of both physical and virtual workloads and systems, and features integrated power management to actively manage server power consumption. Building upon the function of version 1.1, CloudBurst now extends its recommended deployment in test/development environments to include operation in full production environments. With the addition of Tivoli Service Automation Manager 7.2, the solution can now manage systems outside of the integrated BladeCenter platform, including System x, Power Systems, and System z platforms. CloudBurst is workload-optimized to deliver customer solutions based on the appropriate architecture dictated by an organization's workload profile.

Case Study: Smart Traffic

Traffic congestion is a growing concern in most urban areas. While at one level the solution, decreasing the number of vehicles using the road system at a given time, may seem like a simple matter of reduction, it is in fact a complex undertaking that spans many different planning, execution, and management actions. The city of Stockholm, Sweden embarked on a gridlock reduction initiative that exemplifies the need for workload-optimized solutions. In the following sections we review how very different IT business processes and IT workloads combined to deliver the solution to the city's very vexing transportation challenges.

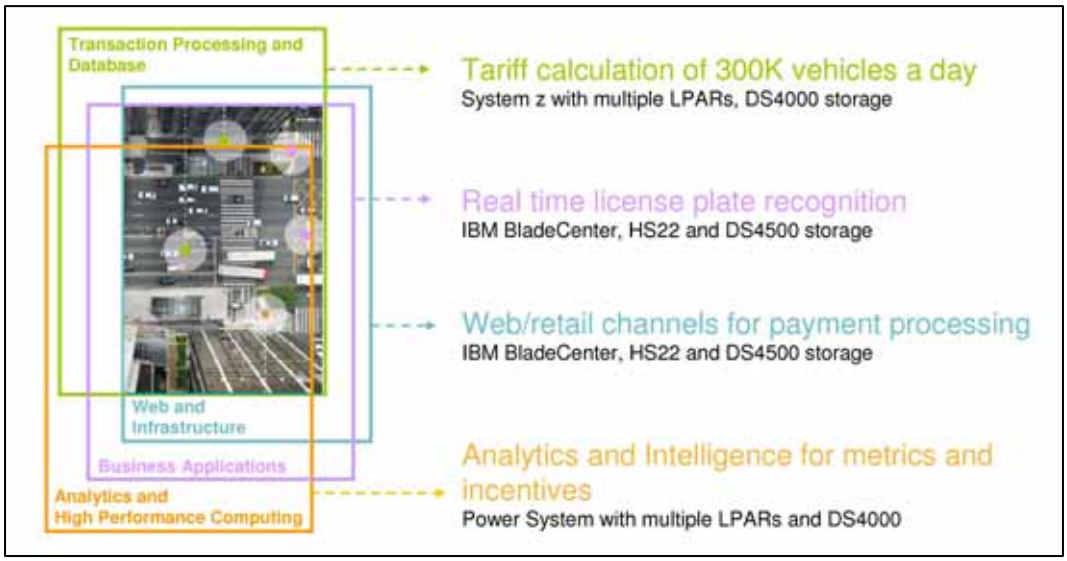

Figure 3: Smart Traffic Solution

The Stockholm City Council, with the assistance of the Swedish Road Administration (SRA), embraced a market-driven approach to reducing the number of vehicles trips (~450,000) through the 24 sq. km. central business district (CBD) each day. Varying fees are assessed during periods of the highest congestion to discourage unnecessary travel by private vehicle. In order to achieve effective congestion management, the city would need to be able to track each vehicle as it entered and exited the CBD, identify its owner, and ultimately collect the toll. As illustrated in figure 3, the Smart Traffic solution weaves together information from multiple sources, each with a unique workload profile:

- Traffic Information, the collection and processing of realtime traffic flow, represents a *Web-based, multimedia, and collaborative* workload that is continuously garnering information from traffic cameras, news reports and government agency information, as well as private citizens reporting first hand.
- Traffic Flow Prediction is an example of a *business analytics* workload in that it is analyzing historical data in conjunction with realtime road, weather, and related conditions.
- Transit Management, the management of all public transit assets and their maintenance and operational scheduling, is very much a *transaction processing*-oriented workload.
- Toll Collection represents a traditional *business application* workload focused on account management through charges and payments derived from the implementation of the Smart Traffic solution.

Each time a vehicle crosses the CBD boundary it automatically triggers one of two sensing mechanisms. If equipped with an SRA-issued transponder a roadside gateway detects the vehicle's passage and sends the information to the central processing system. Otherwise, the vehicle's front and rear license plates are photographed by gateway cameras and the images are forwarded to the centralized processing facility where a license plate recognition application verifies ownership against vehicle registration data. The road usage fee is calculated and a transaction posted against the owner's account. Account payments can be received through automatic bank withdrawal, Web-based self-service payment, or at kiosks located at various retailers throughout the city.

As can be seen, there are many different combinations of workloads necessary to make this Smart Traffic implementation a reality. The level of IT efficiency and resilience required to process the average 2.5 million daily transactions would not be cost-effectively achieved without tailoring the solution to meet the multiple disparate workloads. Smart Traffic requires a high level of instrumentation, real time collection and processing of structured and unstructured data on a massive scale, high levels of transaction processing, traditional business applications, and service management across multiple disparate workloads and infrastructure components. The confluence of data and transactions evident in a smart traffic solution is manifest in many different commercial scenarios every day. Wherever there is data and interactive behavior (buying and selling), there is a dynamic opportunity for business. As organizations seek competitive differentiation through business transformation, the efficiency gains afforded by workload optimization become all the more apparent.

Software	Servers	Services
• IBM DB2 Universal Database • IBM WebSphere MQ • IBM WebSphere Message Broker • IBM WebSphere Application Server • IBM WebSphere Portal	• IBM System p Server • IBM System z Server • IBM Blade Center HS22 • IBM DS4500 Storage • IBM DS4000 Storage	• IBM Global Business Services • IBM Global Services e-business Hosting • IBM Business Transformation Outsourcing

Table 1: Key IBM Technology Components in the Smart Traffic Solution

What It All Means

Business is transforming as a smarter planet in