

# Enhancing Business Intelligence in a Smarter Computing Environment through Cost Analysis

Saurabh Kacker, Vandana Choudhary, Tanupriya Choudhury, Vasudha Vashisht  
Department of Computer Science  
Lingaya's University  
Faridabad, India

**Abstract**-The paper aims at improving Business Intelligence in a Smarter Computing Environment through Cost Analysis. Smarter Computing is a new approach to designing IT infrastructures to create new opportunities like creating new business models, find new ways of delivering technology-based services, and generate new insights from IT to fuel innovation and dramatically improve the economics of IT. The paper looks at various performance metrics to lower the cost of implementing Business intelligence in a smarter computing environment, to generate a cost efficient system. To ensure it, smarter services are deployed with business strategy. The working principle is based on workloads optimizations and their corresponding performance metrics like value metrics, Advanced Data Capabilities and Virtualizations so as to decrease the total IT cost.

**Keywords**- Smarter Computing; Business Intelligence; Cost Analysis; Virtualizations; Advanced Data Capabilities; Value Metrics.

## I. INTRODUCTION

We can see dramatic shifts as our planet becomes smarter. These shifts are changing the way the world works. Cities are becoming smarter by transforming traffic systems, water systems, security—every possible form of municipal infrastructure. Business process is evolving across every industry—banking, trading, manufacturing. And we're seeing changes in the way people live, enjoying advancements ranging from reduced congestion and pollution to new ways to communicate and collaborate. Every aspect of life is benefited from the instrumentation, interconnection and infusion of intelligence into the systems of the world [11].

Nothing is changing more than information technology: the way it's accessed, the way it's applied, and the way it's architected [11]. Change is the name of the game – particularly in Information Technology, it starts with the ever-changing expectations of everyone in the chain: the providers of IT, the partakers and the consumers. To a large extent this is driven by evolving business landscape, stiffer competition and the threat of becoming too outdated. Customer demands for top value products and services are escalating – with which IT infrastructure and systems simply need to keep up. The opportunities for innovation have never been greater.

By thinking differently about computing, the leaders have addressed the IT conundrum—meeting exploding demand for service on a flat budget. This conundrum traps IT organizations

in a vicious cycle in which a rigid infrastructure and lack of trusted data lead to reactive decision-making. Meanwhile, attempts to overcome these challenges through additional IT investments result in a more sprawling and costly infrastructure [8].

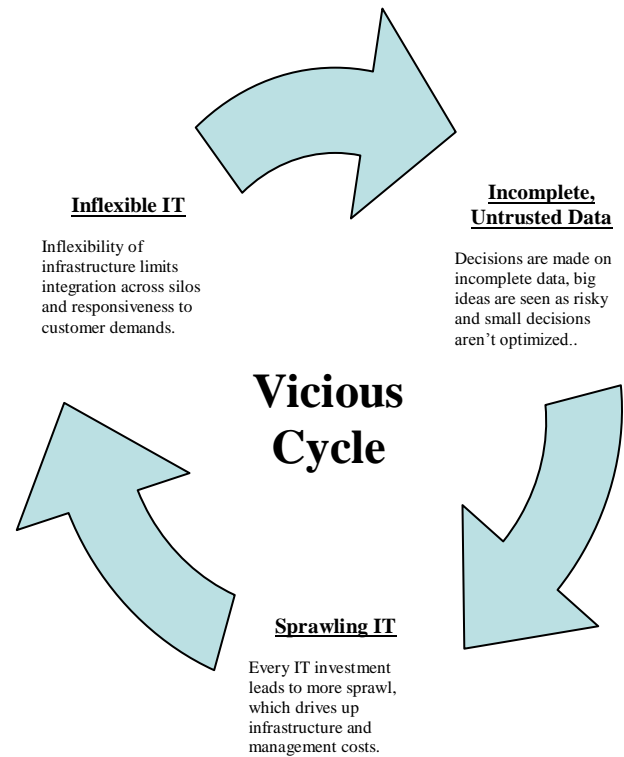


Figure 1. Current IT Business Situation

Leaders in this new era understand that the traditional approach of just adding more servers, storage, networking and other system components to meet increasing demand is no longer a sustainable model. This traditional approach ends in hardware sprawl, bloated software and labour costs, higher environmental costs, all contributing to higher total IT costs. This makes us to enter in a new era of computing known as Smarter Computing—the era of insight for discovery. Smarter Computing is taking a new approach to designing IT infrastructures to create new opportunities like creating new

business models, find new ways of delivering technology-based services, and generate new insights from IT to fuel innovation and dramatically improve the economics of IT. A smarter approach to computing makes it possible to meet increasing demand and to support innovation while managing to stay within nearly flat IT budgets. Optimized Systems for superior economics and Cloud to reinvent business processes and drive innovation. Any enterprise can enter this new era by architecting an IT infrastructure that is designed for data, tuned to the task and managed in the cloud. Thus Smarter Computing will improve the conundrum that is trapping us into a vicious cycle and making a virtuous environment for IT leaders [8].

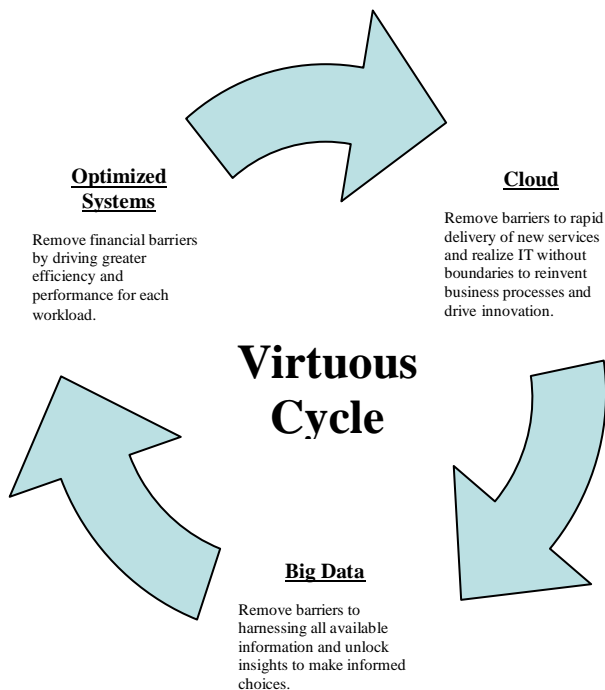


Figure 2. Reversing the Current IT Business situation through Smarter Computing

## II. PROBLEM FORMULATION

Recently the research community has believed that Business Intelligence has become choice of midsize and large size companies that want to improve business processes, discover new opportunities, and edge out the competition. Business Intelligence gives you the information when you need it, in the format you need. By integrating data from across your enterprise and delivering self-service reporting and analysis, IT spends less time responding to requests and business users spend less time looking for information. Business Intelligence tries to improve the productivity of business analysts and preserve information consistency throughout an organization. Smarter Computing is a promising environment which provides computation and storage resources as services along with efficient method to reduce efforts. Smarter computing infrastructure and related mechanisms allows the stability, equilibrium, efficient resource use and sustainability of a Business Intelligence System [4].

IT serves the strategy of the business, but keeping business objectives and IT operations in alignment is not always easy. We may have a clear business strategy mapped to detailed business processes that are ready to implement but still the execution stumbles. Why? One reason is that the information systems needed to execute the strategy are insufficient or poorly matched to the requirements. Smarter computing and service oriented architectures can mitigate the risk of such misalignments, assuming they are used in ways supportive of business strategy. Aligning business strategy and IT services is a several step process, at least at the most coarse level:

- Identifying key business objectives
- Identifying IT services needed to support those objectives
- Assessing the current state of IT services and identifying gaps between the existing set and the needed set of IT services.
- Developing a plan for reducing the gap between the existing and needed set of information services

Key business objectives may include controlling and reducing costs, enabling more rapid response to changing market conditions, improving governance of the organization, or improving the resiliency of IT operations to adverse events, such as hardware failures, loss of power, or natural disaster. Many of the services needed to support business objectives can be readily identified once the business objectives are known. Cost controls and cost reduction come with more efficient server utilization, more self-service in systems management, and reduced overhead associated with infrastructure services such as backups, high availability, and disaster recovery. The gap analysis process should take into account both technical and organizational considerations. For example, will existing hardware readily deploy in cloud architecture or will new hardware be required? Are service management practices mature enough to implement in self-service delivery systems? Is a billing or chargeback mechanism in place if a private cloud is under consideration? The first steps in creating a plan to move from the existing to the needed systems are to prioritize the gaps and identify dependencies in the process. The next step is creation of Dynamic Infrastructure. A Dynamic Infrastructure is required for delivery of smarter products and services. A dynamic infrastructure is one that is designed to enable today's instrumented and interconnected world, helping clients integrate new smarter components of their business with a flexible, secure and seamlessly managed IT infrastructure. Organizations supported by a dynamic infrastructure can deliver superior business and IT services with agility and speed, while addressing the day-to-day operational needs to improve service, reduce cost and manage risk. A first step towards enabling a dynamic infrastructure is to break down the traditional barriers that separate software development and IT operations. In most organizations today, a separation exists between software development and IT operations. The two groups use different tools, speak different languages, maintain proprietary data stores and expend few resources on integration with external entities.

Generally, this separation is born out of the requirement to enable the two functions to specialize and focus on their individual goals. For development the goal is application creation, for IT operations, the goal is service assurance. Given this distinction, specialization within development and operations makes sense. As with many undertakings, specialization underpins organizational effectiveness, creating value by allowing teams to focus on a subset of well-defined objectives.

Specialization, however, only works up to the point where work flows are contained within organizational boundaries. When focus is put on the on the larger goal of end-to-end service delivery where workflows must be coupled across teams, this specialization of tools, skills, and processes can expose the business to unwarranted risk [1].

### III. ARCHITECTURAL STRATEGY

Current changes in IT and User expectations make architecting for a smarter strategic computing plan a necessity for businesses. Alignment with past technology initiative like Service Oriented Architecture (SOA) and delivery processes like Agile is important. After all, many organizations are deep into their SOA and data centre consolidation & transformation programs so any low hanging fruit that they can obtain towards a smarter computing path needs to be exposed.

However, a strong architectural governance and management process is by no means the only criteria for a smarter computing plan. Not all organizations or businesses will face the same integration, architectural or evolutionary challenges. That is why there is a resurgence of Service Level Agreements (SLA) which provides business units with performance metrics that they can monitor as part of their justification for using internal rather than external services. Business managers may look at cloud services and find the lower costs, greater control, and potential for scaling business processes compelling reasons to use cloud services. These reasons are often not enough, though. It is not sufficient for a cloud to work well today. This is why we need SLAs. SLAs are standard in IT, and it is no surprise that they are used with cloud services. Rather than focus just on the availability of a specific application, cloud SLAs may be more general and apply to capacity commitments, network infrastructure, storage infrastructure, and availability and recovery management. These SLAs are closely coupled to the infrastructure of the cloud, but the primary concern is on the business commitments cloud providers make to their customers.

Over the years, Organizations have implemented numerous architectural strategies and business driven IT systems to create an integrated business oriented computing strategy. Business processes have been reengineered over and over again to meet each new technology model of architecture change. For different industries and organizations sizes there may be a broad range of concerns and issues to consider ranging from regulatory or statutory requirements, to technology maturity and management commitment and culture.

Ability to Integrate, Automate and apply secure design and delivery practices are amongst the necessary criteria for delivering a framework that is able to:

1. Respond and deliver value quickly and pragmatically.
2. Make use of existing investments and incorporates a broad range of appropriate facilities and resources.
3. Address and protect against key security concerns and requirements.

Making decisions on the above factors help in designing a smarter computing plan for Business valued IT. The IT estate for many businesses and organizational data centres has expanded to include a wide range of infrastructure platforms from mainframes to distributed servers. Advances in software infrastructure, application technology and the arrival of service based delivery models have resulted in businesses facing new architectural models and challenges:

1. *Infrastructure:* Virtualized, hosted/service based, distributed, complex, hybrid on demands.
2. *Application:* Virtualized, distributed, packaged complex, custom, on demands, mobile.
3. *Data:* Federated, private, public, big, structured, unstructured, transient.
4. *IT Process:* Mixed (agile, iterative, formal), automated.
5. *Mobile:* Smartphone, Tablets, hotel business centres, internet cafes [7].

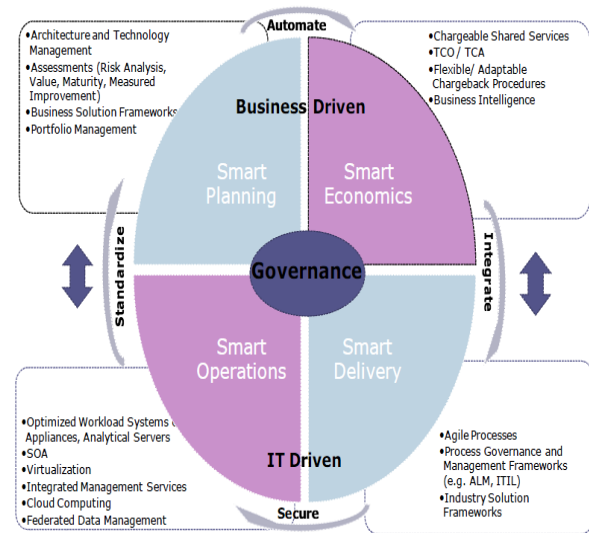


Figure 3. Integrated Strategy for Business valued IT

#### A Workload Optimizations

A large server farm is indistinguishable from a set of cloud servers when looking at the hardware. Servers, switches, routers, power supplies, and other components are the same.

The difference lies in how these resources are used. The servers in a typical corporate data centre prior to the advent of cloud computing were assigned to a particular department or application use. The configuration was relatively fixed and changed only when the server was upgraded, reassigned, or decommissioned. These servers were configured to do one type of operation. This makes for a reliable compute resource, but not an efficient one.

Servers with fixed configurations are less likely to have high- utilization rates. Unless there is a steady stream of jobs that fits the machine's configuration, there will be idle periods. Without proper infrastructure for rapidly deploying virtual machines, the cost of reconfiguring a server is so high that it is done only for significant long- term changes. In the cloud, the cost of switching virtual machines is low enough that idle servers can be reconfigured with different virtual machine images allowing other applications to run on the same physical server that had just been running other types of jobs.

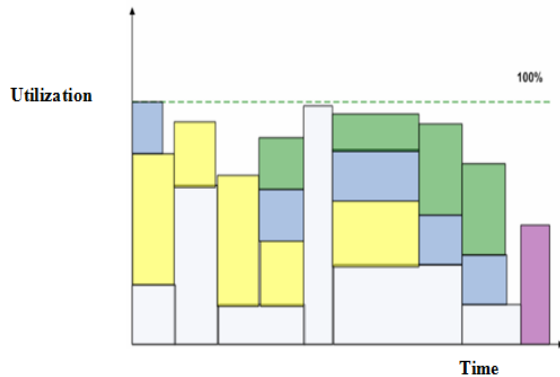


Figure 4. Cloud Server Utilization

The figure shows that in the cloud, server utilization can be significantly higher when workloads are distributed and optimized over available servers [4].

#### B. Virtualizations

Currently, Virtualization revolves around use of commodity server to better utilize resources. For general purpose computing such as email server, file or print server and many application server hardware utilization hardly goes above 10%. Using Virtualizations to get better utilization has reduced the physical estate substantially.

Virtualization is not just about servers. Network and storage virtualization are also used in many organizations as they look to improve throughput and get the most use out of the assets. Storage Virtualization is critical to the ability to support the huge volume of data available today. There are others benefits to the wider use of Virtualization such as backup, replication and the ability to clone mission critical servers on demand [7].

Setting up a set of virtual machines on a single server is straightforward: install a hypervisor and create virtual machine instances based on OS(s) of choice. Scaling virtualization to a large number of servers requires management software that can manage multiple hypervisor clients from a single console. Storage services also need to be virtualized so that they appear to cloud consumers to be a single storage device. Virtual machine instances in the cloud, for example, should be able to address storage space on the cloud SAN(s) without having to manage implementation Details. Ideally, the same management console that is used to control servers in the cloud will support management and administration of storage resources. Computing and storage clouds hide many of the implementation details that go into Building and maintaining a large IT infrastructure. By standardizing services, streamlining service

management, and virtualizing physical resources, cloud providers enable the technical resources needed by users to leverage cloud services. Those same users, however, also require attention to business considerations [4].

#### IV. COST ANALYSIS

The current system in which Business Intelligence is implemented is not flexible enough to provide result at a low cost as well as there is no load management between multiple work flows. Smarter Computing aims at transforming the economics of IT while freeing teams to focus on new innovation. To ensure smarter services are deployed with business strategy, we should focus on workloads optimization and their corresponding performance metrics so as to decrease the total IT cost [3].

##### A. Workload Analysis

Right now in business there are hundreds, thousands, or even more applications executing business processes. Some of these are transaction- processing systems that provide high- volume, rapid processing of orders, inquiries, reservations, or a broad array of other narrowly focused business activities. Other applications are performing batch operations, such as generating invoices, reviewing inventory levels, or performing data quality control checks on databases. Still others are extracting data from one application, transforming the data into a format suitable for analysis, and moving it into a data warehouse. There is a wide array of different types of applications that are needed to keep an enterprise functioning.

These different types of applications have different requirements and constraints that must be considered when moving them to the cloud. For example, they might need:

1. To start and finish executing within a particular time period.
2. To wait for another job to complete before it can begin.
3. To limit the functionality of some services, for example, write locking a file to perform a backup.
4. To provision a significant number of servers for a short period of time for a compute intensive operation.

Every Business will have finite resources. As part of the planning process, we need to understand the true cost of delivering a set of workloads. Often, when people compare the cost of deployment options, they limit the comparison to the cost of hardware acquisition. This can be quite misleading. It is important to consider all of the key elements of cost. Even for a Total Cost of Acquisition (TCA) calculation alone, the different elements of cost include software acquisition costs, software S&S (support & subscription) costs, hardware maintenance costs etc., in addition to the base hardware acquisition costs. A Total Cost of Ownership (TCO) calculation typically is much broader and includes many other relevant elements of cost— some being administration/labor costs, systems management software costs, power and cooling costs, facilities costs, refresh costs etc. Another important aspect to consider is the amount of work being done on the two systems. If two systems being compared have been sized upfront so as to guarantee the same amount of work is being done on both, a direct total cost

comparison is valid. In scenarios where we compare systems delivering different amounts of work, it is imperative that we reduce it to a cost per workload comparison to understand true value.

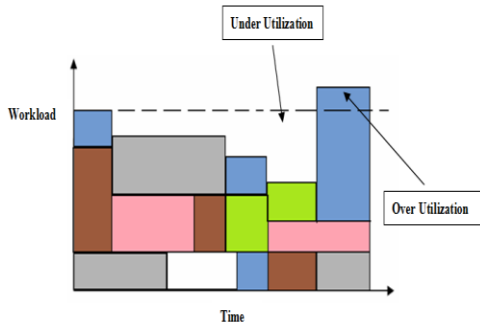


Figure 5. Combinations of workload determining the Total Utilization at a point of time.

The Total Utilization can be increased by minimizing under and Over Utilization. To do so, we can look at existing workloads and try to answer following questions:

1. How often jobs execute on dedicated servers.
2. The level of utilization of those servers.
3. Time constraints on when those jobs execute.

Answering these questions is not an easy task as problem of load balancing arises. For this we want to deploy sufficient cloud infrastructure to avoid periods when cloud consumers want to run more jobs than there is capacity for (over- utilization) at the same time we do not want extended periods of time when servers are idle (under- utilization). This brings us to the second aspect of business alignment: value metrics [4].

### B. Value Metrics

The last few years have seen a tremendous increase in server hardware capabilities, especially in the number of processor cores and hardware threads available per server. For example, IBM® POWER7® can support 4 simultaneous threads per core and can scale to 256 cores, making available a massive 1024 hardware threads in a single frame for meeting the ever increasing demand for processing. However, having this kind of hardware capability is not as useful if the software running on that machine cannot exploit those hardware capabilities. The IBM Software suite of products is designed to take better advantage of available hardware computing capabilities. Together with IBM hardware and their differentiated features, IBM Software exploits the higher number of cores and threads, increasing application utilization efficiency, and delivers better price/performance for many workloads.

The figure 6 shows that that the IBM solution is able to deliver an online banking workload at a much lower cost per workload than a competitive option. Thus adding optimization systems with specialized, purpose-built appliances such IBM WebSphere DataPower® for SOA, Netezza for data warehousing, and Cloudburst™ for simplified cloud administration for running specific types of workloads provide focused capabilities with very quick return on investment.

IBM WebSphere Application Server Designed to Leverage Threads in POWER 7 Systems		
<b>WebSphere Application Server v7.0</b> 1 Instance AIX v6.1 64 bit	<b>IBM BladeCenter PS701</b> 8 Cores/3.0 GHz 32 Threads	<b>5,009 Transactions per Second</b>  <b>\$32 per transaction per second</b>
<b>Competitive Application Server</b> 1 Instance Solaris 10 64 bit	<b>Oracle Sparc T3-1B</b> 8 Cores/1.65 GHz 64 Threads	<b>746 transactions per second</b>  <b>\$137 per transaction per second</b>

Figure 6. How IBM Web Sphere Application Server exploits the hardware capabilities of an IBM POWER7 server.

These appliances are generally comprised of a base operating system, the necessary middleware, and the application in question into a single stack, pre-configured to work out-of-the-box. They typically require lesser skills and labor to configure and operate. Install time is often mentioned in hours instead of days or weeks. Focused delivery capabilities make appliances ideal for certain workload types. Integrating hardware and software to create optimized systems, lowers cost per workload dramatically [5].

### C. Advanced Data Capabilities

We are seeing an explosion in data – the volume of information being generated has increased exponentially in the last few years. This data needs to be stored, managed and used efficiently. Storage systems need to be able to handle this growth. Software for managing data needs to leverage hardware capabilities. Business analytics and deriving intelligence from information is becoming an increasingly important factor that decides how competitive a company will be.

IBM software, servers and storage systems are meeting this challenge. IBM offers a broad portfolio of storage systems, ranging from small to mid-range to enterprise class. There are numerous innovations in this space. For example, IBM's Easy Tier® capability can automatically allocate optimum amounts of Solid State Drives (SSD), thus maximizing SSD performance gains while minimizing costs. Another example of innovation in this space is the IBM Smart Analytics Optimizer—a purpose-built appliance for optimizing business analytics on zEnterprise.

Real world Business Analytics workloads need to support concurrent performance. A typical Business Analytics solution will need to support multiple users executing a wide variety of queries and reports concurrently. IBM offers packaged solutions like the IBM Smart Analytics System (ISAS) that have been built as pre-packaged, pre-configured solutions to support this kind of real world usage patterns

Solid state disk drives are revolutionizing storage performance—they support much higher Input/Output Operations (IOPS) than traditional hard disk drives (HDDs). However, SSDs are also more expensive than HDDs. Over provisioning storage systems with SSDs in the hope of getting the best performance often results in huge storage cost increases. We find that, often a small fraction of SSDs will yield the majority of the performance gain possible for a given workload. The ideal price/performance point is reached by having a judicious mix of SSDs and HDDs. Doing this manually is quite

cumbersome and inefficient. IBM's Easy Tier can optimize the amount of SSD allocated. It dynamically moves data to SSD, based on hot spots detected. Further, Easy Tier can dynamically share the available SSDs across many workloads, efficiently allocating the SSDs to the hottest spots.

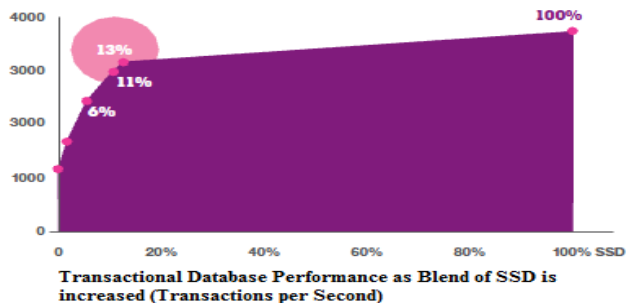


Figure 7. Easy Tier in IBM Storage Systems Automates Optimum Use of SSD

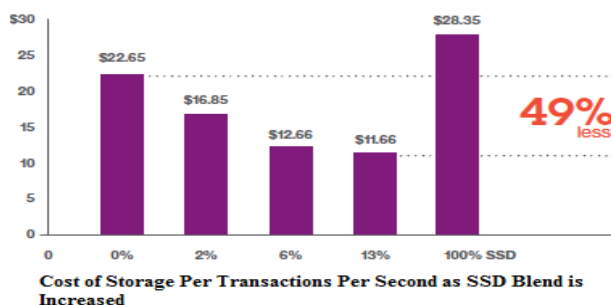


Figure 8. Easy Tier in IBM Storage Systems Automates Optimum Use of SSD (Continued)

Figure 7 & 8 demonstrates how Easy Tier helps improve IT economics when dealing with data workloads. Figure 7 shows just 13% blend of SSD to HDD achieves 171% performance gain and Figure 8 shows Easy Tier achieves 78% of the maximum SSD performance potential with a blend of just 13% SSD.

Automatic allocation of expensive SSD resources in an optimum way results in lowering overall storage costs. Automation also reduces labor costs. IBM offers a comprehensive suite of software products aimed at efficiently automating various data center tasks. Some examples are software that automatically provisions systems, software that synchronizes the start and stop of sequenced applications, software that automatically adjusts system resources available to workloads to meet varying demand etc. Automation not just reduces labor costs; it also improves the quality of service delivery. This is especially true as automating mundane repetitive tasks eliminates the risk of manual operations error [5].

#### D. Improve IT Economics with Private Clouds and New Service Delivery Models

Cloud computing is poised to transform the service delivery model. At an abstract level, it is about delivering hosted services with some key distinguishing attributes:

1. *Elastic Scaling:* Resources scale up and down by large factors as the demand changes.

2. *Flexible Pricing:* Utility pricing, variable payments, pay-by-consumption and subscription models make pricing of IT services more flexible.
3. *Ease of Use:* End user often just needs a PC with Internet access to request services, without IT skills or knowledge of the system.

Many businesses are moving to a Cloud Computing model. Line-of-business units within organizations are going to public cloud providers as their low cost, pay-per-use model is seen as more cost effective. Furthermore, faster provisioning of resources offered by them can enable users to respond with agility. In some cases, the public cloud model makes the most economic sense [5].

There are three broad modes of delivery for cloud services: private, public, and hybrid. A private cloud is deployed and managed by an organization for its own internal use. The organization controls all aspects of cloud implementation, management, and governance. One of the most significant advantages of this approach is that data never leaves the control of its owner. This reduces the risk that an outside party will gain access to private or confidential data. Depending on the implementation and management details, private clouds may be more cost effective as well. For example, a business may have significant investment in servers that can be redeployed in the cloud, lowering the initial costs.

A public cloud is one that is managed by a third party that provides services to its customers. The primary advantage is low start-up costs on the part of customers and minimal management overhead, at least with respect to basic cloud services. Businesses will still need to manage their workloads, allocate chargeback, and so on. Choosing between public and private cloud implementations is not an all-or-nothing proposition. Hybrid clouds, or the combination of private and public implementations to run business services, have emerged as a third alternative. Consider the economic benefits. There may be a point, however, at which the benefit of adding servers to a private cloud is not sufficient to offset the costs of adding them. For example, the distribution of workloads may entail a number of peak periods where demand exceeds the capacity of the private cloud. These peaks may be regular short periods (for example, at the end of the month when accounts are closed and data warehouses and data marts are updated and many reports are generated) or they may be more unpredictable periods of high demand.

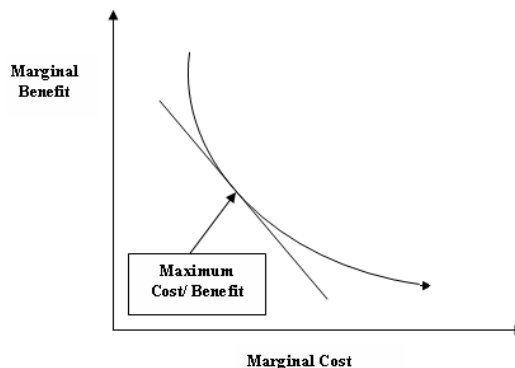


Figure 9. Cost/Benefit of additional Investment in a Private Cloud.

The figure demonstrates the cost of adding and maintaining additional cloud resources eventually reaches a point where the costs outweigh the benefits. At this point, a hybrid cloud approach may be the most cost effective option.

In general, private clouds are a model for delivering IT services in a pay-as-you-go fashion similar to what public clouds can offer. The difference is that a private cloud is built from resources inside an organization and is typically hosted within the data center to be used by line of business and other end users within the organization. Workloads run on large scale virtualization platforms. This approach may significantly reduce hardware, software and labor costs. Users request services via a self-service portal and virtual machines are quickly provisioned. Labor costs may be significantly reduced via standardization and automation. Furthermore, users may be given tools to manage their applications on the virtual machines that are running their services and they only pay for what they consume [4].

For example, IBM provides large scale virtualization environments in POWER7 and zEnterprise that are ideal to host private clouds. IBM provides software that manages the entire lifecycle of virtual servers—everything from self-service automated provisioning to metering and billing based on usage. IBM also offers industry specific Cloud Service Platforms. An example is the recently announced IBM Cloud Service Provider Platform, a comprehensive set of hardware, software and services to help providers rapidly deliver cloud computing on their own.

The zEnterprise system provides the broadest architectural support for building a private cloud. Different environments in zEnterprise may be used to run workloads without requiring a port or rewrite. Where there is an option, workloads may be best fit to an environment to run at the lowest cost per workload. A Fit-for-Purpose deployment strategy aims to assign a workload to the environment that best satisfies the particular requirements of that workload. For example, workloads with heavy IO demand may be best fit on Linux on z/VM on the z196 portion of zEnterprise. Workloads that have high CPU demand and that can exploit multithreading may be best fit on the POWER7 blades in the zBX. Large scale virtualization on z/VM drive down acquisition costs. The private data network as well as the private management network between the z196 and the zBX plus network access control mechanisms ensure a secure network environment. Managing this environment with zManager and with Integrated Service Management (ISM) software, resulting in reduced operational costs [5].

### I. CONCLUSION

We started with a question: How can IT better serve the business? There is no one simple answer to this question but it is clear that aligning development and operations processes driven by today's business and technological maturation has exposed new challenges. There are three main challenges when it comes to implementing a Business Intelligence solution: staying within budget, meeting the needs of the business, and efficiently maintaining the solution going forward.

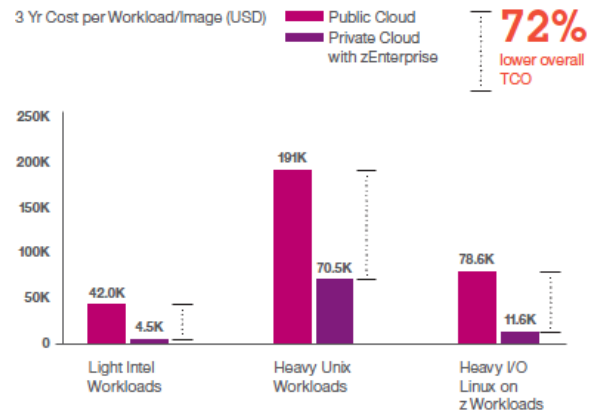


Figure 10. Private Cloud on Z Enterprise Dramatically Reduces Costs using virtualization servers

However, the approach discussed above to reduce the economics of IT and thereby improving the business has produced cost-effective solutions to maintain enterprise-wide workloads, integrate within the existing IT environment, and adapt to changing business needs. The graph below shows that if Business Intelligence is implemented in a Smarter Computing Environment can minimize costs and complexity by leveraging existing IT investments and increase IT productivity by reducing Application Development time, Total cost per workload, Server acquisition cost, Database cost, Power Consumption, New Application Development Time, Floor Space and thereby maximizing organizational effectiveness and helping the Business to grow in a Smarter way as the planet is becoming Smarter.

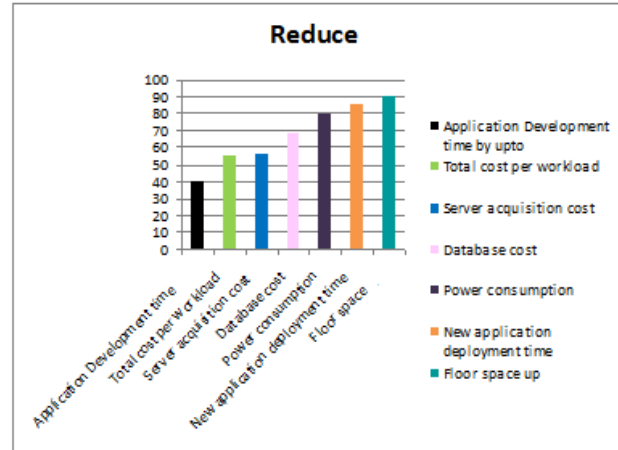


Figure 11. Smarter Computing improving Business efficiency.

### II. FUTURE SCOPE

The system accuracy can be improved by adding Data Governance process in place in order to maintain data integrity and the quality of data being used across the enterprise, i.e., being able to control where the data comes from, who controls data values, consistent data types, etc. Having a tight process in place can also help ensure that business data is not exposed to unauthorized personnel, especially sensitive information that can be exploited.

In the past, database administrators had exclusive access to the data and would provide canned reports to the business users. Now business users are being given greater access to data so that they can slice and dice the information and organize the data in a way that best meets their needs. Making your Data Governance processes more robust can not only allow for broader usage of your data to improve insight and decision making capabilities, but it will also help ensure that data quality is maintained in a safe and secure way and using Predictive Analytics to find patterns in historical and transactional data to identify opportunities and risks for the business has become increasingly sophisticated over the last few years [10].

#### REFERENCES

- [1] IBM, "Deliver smarter products and services by unifying software development and IT operations," Issue Date : Sep. 2009, ISBN RAW14175-USEN-00.
- [2] Hien Nguyen Van, Frederic Dang Tran, "Autonomic virtual resource management for service hosting platforms," Issue Date: May 23, 2009, ISBN 978-1-4244-3713-9.
- [3] Derrick Kondo, Bahman Javadi, Paul Malecot, Franck Cappello, David P. Anderson, "Cost Benefit Analysis of Cloud Computing versus Desktop Grids," Issue Date : Sep. 2009, ISBN 978-1-4244-3750-4.
- [4] Dan Sullivan, "The Definitive Guide to Cloud Computing," Realtime Publishers.
- [5] IBM, "Smarter Computing & Breakthrough IT Economics," Issue Date : Apr. 2011, ISBN XBL03007-USEN-01.
- [6] Gary Barnett, "Mainframe and the Cloud," Issue Date : Nov. 2010, ISBN BC-INF-W-00068UK-EN-00.
- [7] Bola Rotibi, Ian Murphy, "Business Benefits of an Architectural Strategy for Smarter Computing," Research Analyst in Creative Intellect Consulting, Issue Date : Jun. 2011.
- [8] Sam Palmisano, "Smarter Computing -The Next Era of IT," IBM, Issue Date: April 2011.
- [9] G. Khanna, K. Beaty, G. Kar and A. Kochut, "Application Performance Management in Virtualized Server Environments," Network Operations and Management Symposium. 2006.
- [10] Raj Gill, "Top Ten Data Management Trends," Scalability Experts, Issue Date: June 2011.
- [11] "www-03.ibm.com/systems/data/flash/ae/smartercomputing".

#### AUTHORS PROFILE

**Saurabh Kacker** received his bachelor's degree in Computer Science from Rajasthan Technical University, Kota, India and currently pursuing Master's Degree from Lingaya's University, Faridabad, India. His areas of interests include Cloud Computing, Network Security, Natural Language Processing etc.

**Vandana Choudhary** received her bachelor's degree in Computer Science from Kurukshetra University, Kurukshetra, India and currently pursuing Master's Degree from Lingaya's University, Faridabad, India. Her areas of interests include Cloud Computing, Natural Language Processing etc.

**Tanupriya Choudhury** received his bachelor's degree in CSE from West Bengal University of Technology, Kolkata, India, master's Degree in CSE from Dr. M.G.R University, Chennai, India and currently pursuing his Doctoral Degree. He has two year experience in teaching. Currently he is working as Asst. Professor in dept. of CSE at Lingaya's University, Faridabad, India. His areas of interests include Cloud Computing, Network Security, Data mining and Warehousing, Image processing etc.

**Vasudha Vashisht** received her bachelor's and master's degree in Computer Science from M.D. University, Haryana, India. She has 6 years of experience in teaching. Currently, she is working as Assistant Professor in the Dept. of Computer Sc. & Engg. at Lingaya's University, Faridabad, Haryana, India. She has authored 10 papers and her areas of interests include artificial intelligence, Cognitive Science, Brain Computer Interface, Image & Signal Processing. Currently she is pursuing her doctoral degree in Computer Science & Engg. She is a member of reputed bodies like IEEE, International Association of Engineers, International Neural Network Society, etc.