

Batch Processing with WebSphere Compute Grid: Delivering Business Value to the Enterprise



Redguides
for Business Leaders

Snehal Antani



- Reduce costs
- Improve time to market
- Address aggressive non-functional requirements



Executive overview

Batch processing is a mission critical workload for the enterprise. End of day/month/year reporting, bulk account processing for credit scores and assessing interest, and reconciling banking activities are just a few examples of tasks carried out with batch processing. Whether they like it or not, or even are cognizant of its role and importance, enterprises depend on batch processing. Yet, in this context, enterprises are suffering. Over the past decade, their online-transactional processing systems (OLTP) have evolved, where application servers such as WebSphere® Application Server served as the foundation for that evolution. Standards for Web services and other OLTP technologies emerged, programming models like JEE were standardized, and service-oriented architecture (SOA) was pursued. Throughout this evolution however, batch systems were forgotten. Enterprises are suffering because these batch systems and new batch requirements are causing them pain in the following ways:

- ▶ There are pressures to reduce costs, where operations must be more efficient and application development and maintenance costs must be reduced.
- ▶ Time-to-market must be improved, where the solution is agile and can quickly deliver new business value.
- ▶ New non-functional requirements, where batch workloads must run faster or handle larger volumes of data, must be addressed.

These pain points cannot be ignored because with time they will just get worse. WebSphere Compute Grid is not only the painkiller, but also an end-all therapy that ensures enterprises will remain agile, scalable, and cost efficient. WebSphere Compute Grid delivers a complete, cloud-enabled batch-processing platform for the enterprise that includes the following features:

- ▶ Comprehensive development and management tools for building and deploying Java™-based, cloud-enabled batch applications
- ▶ A resilient, highly available, secure, and scalable runtime with container-managed services for batch applications
- ▶ A platform capable of supporting 24x7 batch and OLTP processing, and parallel computing on highly virtualized and cloud-based runtimes
- ▶ Integration capability with existing infrastructure processes like enterprise schedulers (Tivoli® Workload Scheduler, and so forth), archiving/auditing technologies (Beta 92 from Beta Systems), and other technologies typically deployed in an enterprise batch solution

- ▶ Integration capability with the overall SOA strategy of reuse by enabling services to be shared across multiple execution domains: batch, OLTP, real-time, and so forth
- ▶ Ensures high-performance batch on the mainframe by integrating with z/OS®, making use of workload management and performance optimizations gained by running in close proximity of its data
- ▶ Platform and vendor-neutral batch applications that allow the location of the application's data to dictate its deployment platform.

Compute Grid delivers the much-needed painkillers to alleviate the pain in the short-term, but also a therapy that ensures such pains are not endured again. The relief can be summarized as the following:

- ▶ Cost reductions
 - Reduce code development and maintenance costs in the following ways:
 - Sharing business services across multiple execution domains (for example, batch, OLTP, real-time). Approximately 20–40% of the application code can be eliminated through the use of shared-services design patterns.
 - Reducing application infrastructure code through the use of container-managed services (for example, checkpoint/restart). Approximately 25–50% of home-grown batch code can be eliminated with WebSphere Compute Grid.
 - Reducing operational costs via reduced complexity in the following ways:
 - Standardizing operational procedures (archiving, auditing, scheduling) across the enterprise.
 - Eliminating redundant infrastructure processes (for example, testing, deployment, disaster recovery, scheduling).
 - Integrating with existing processes (for example, archiving, auditing, scheduling).
 - Reducing operational costs by making use of specialty compute engines on the mainframe (for example, zAAP) and reducing physical resource consumption in the following ways:
 - Implement Java-based batch to make use of zAAP processor execution environment on the mainframe (for example, Java batch jobs can be more than 90% zAAP eligible, and the performance is comparable to native workloads due to just-in-time (JIT) compiler optimizations in Java on z/OS).
 - Use the efficiencies of processing data in bulk, which leads to reduced resource consumption such as CPU, as well as improved throughput. JDBC batching, pre-fetching of data, and processing multiple records within a transaction are some examples of bulk processing efficiencies.
- ▶ Improve time-to-market
 - Improve application agility in the following ways:
 - Application design patterns that promote reuse of shared services across multiple execution domains (for example, batch, OLTP, real-time).
 - Development frameworks, tooling and methodologies to improve development speed, quality and maintainability.
 - Platform neutral applications, where the batch container (and therefore applications) can run on multiple application servers (for example, WebSphere, WebLogic, JBoss, and so forth).
 - Integration with existing application assets (heterogeneous applications by integrating Java and COBOL)

- Improve infrastructure agility in the following ways:
 - End-to-end batch processing platform that runs on multiple platforms (z/OS, z/Linux®, AIX®, UNIX/linux/wintel, and multiple middleware vendors (WebSphere, WebLogic, JBoss, and so forth).
 - Location of the application data dictates application placement, coupled with flexible deployment options ensures that batch will execute within close proximity of the data
- ▶ Addressing new and aggressive non-functional requirements
 - Reduce/eliminate batch windows in the following ways:
 - Run batch and OLTP systems 24x7, where OLTP SLA's and batch deadlines are achieved by manipulating physical (CPU, and so forth), and logical (db locks, and so forth) resources.
 - Achieve higher throughput to address larger data volumes in the following ways:
 - Efficiencies in bulk reading and writing of sequential data
 - Parallel processing across jobs, steps, and application components

Each of these value propositions can be quantified financially, in real dollar savings terms that can effect the bottom line of the business. A business value assessment (BVA) can be conducted where the enterprise works closely with WebSphere Compute Grid subject-matter experts (SMEs) to quantify the financial benefits accurately. A sample BVA, which is based on experiences delivering the product to numerous customers, is included with this paper.

The batch platform

WebSphere Compute Grid delivers a batch platform capable of solving today's business pains, while ensuring the enterprise is agile and able to tackle tomorrow's problems. As Figure 1 illustrates, batch job submitters are not burdened with the details of the batch platform. They instead interact with a job management tier and view the rest of the platform as a cloud.

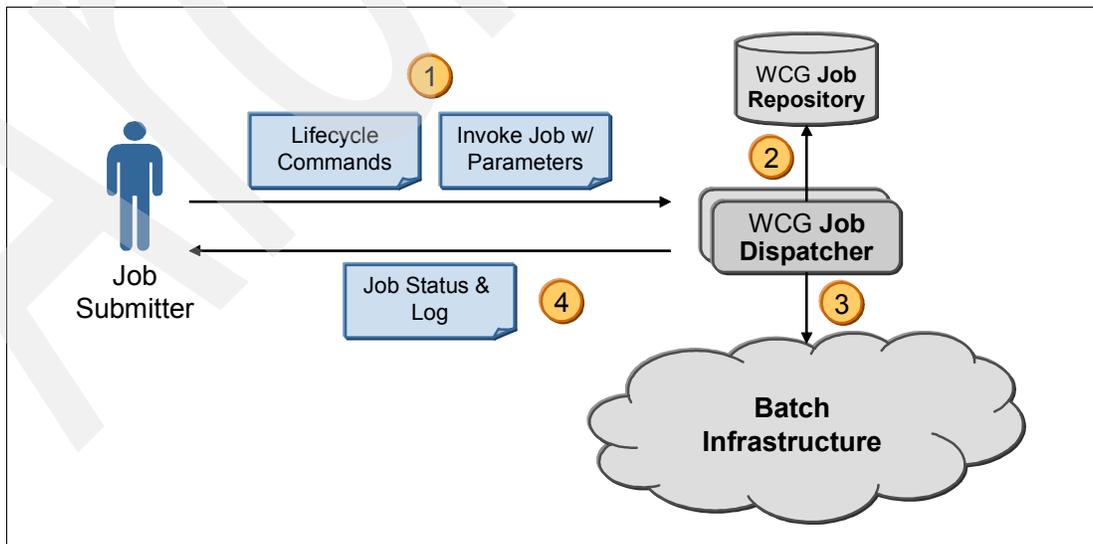


Figure 1 WebSphere Compute Grid's Batch Platform from the perspective of the job submitter

The experience of a job submitter can be described as the following:

- ▶ Job submitters submit batch job definitions to the job management tier. These definitions reference the business logic to be run, parameters describing the input and output data locations, any job-specific qualities of service, and so on. Job submitters are also capable of submitting operational commands like stop/start/cancel/restart on their job instances. Job submitters can be enterprise schedulers such as Tivoli Workload Scheduler, JES on z/OS, users making use of a graphical user interface (GUI), programs using Web Services, Java Messaging Service (JMS), or EJBs, and shell commands.
- ▶ Batch job definitions can be stored in a job repository, which enables the life-cycle of the job definitions to be managed. Invoking batch jobs stored in the repository is synonymous to making a remote-procedure call in other distributed computing paradigms: the name of the job coupled with instance-specific parameters are passed to the system for execution
- ▶ The output of the job execution, which is typically archived for auditing purposes, can be streamed back to the job submitter.

WebSphere Compute Grid is comprised of two primary components:

- ▶ Job dispatcher tier
This component is responsible for managing the execution of batch jobs across a collection of resources, and
- ▶ Batch container tier
This component is responsible for executing the jobs themselves.

The batch containers are multi-tenant in nature, and are capable of running on both WebSphere and non-WebSphere application servers. The strategy for the batch containers is ubiquity. Compute Grid batch containers should transcend application servers and run everywhere. Ubiquitous batch containers translate to portable batch applications. Figure 2 illustrates the infrastructure components of the technology.

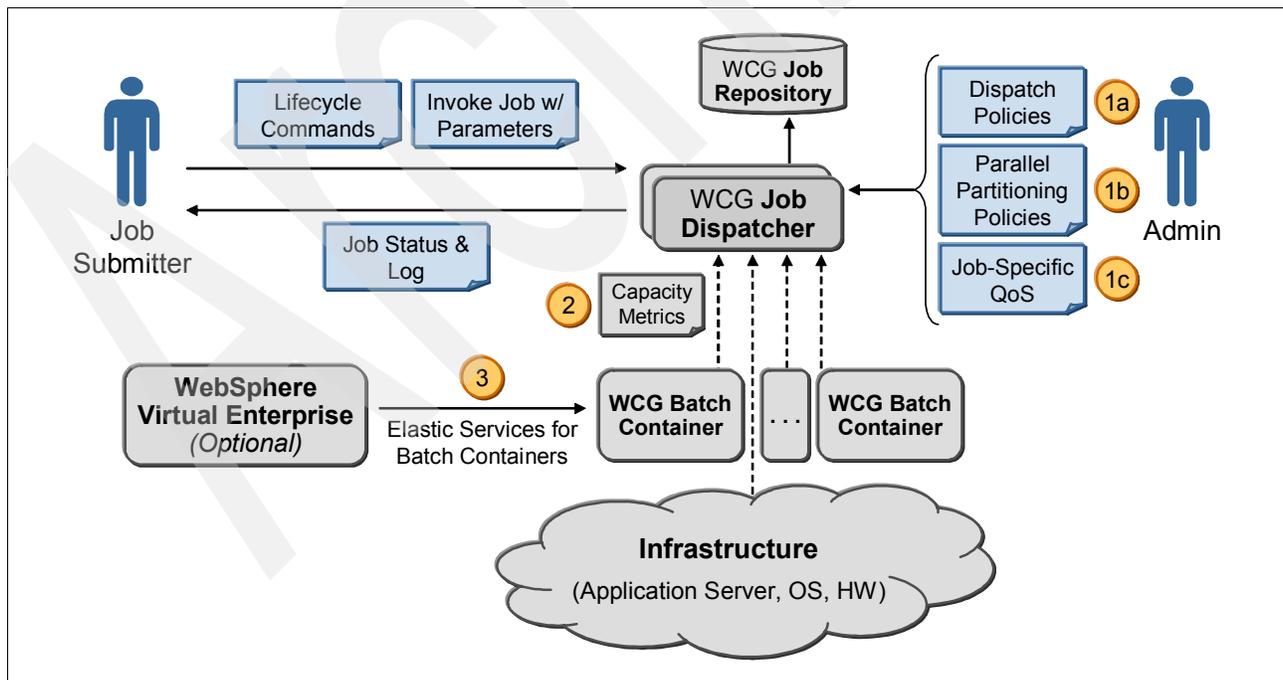


Figure 2 A detailed view of the WebSphere Compute Grid infrastructure

WebSphere Compute Grid administrators define policies that influence how batch jobs are executed. These policies, coupled with autonomics built into the infrastructure, serve as the foundation for cloud-enabled batch. Figure 2 depicts these elements:

1. Policies can be defined by WebSphere Compute Grid administrators that govern how jobs are executed. The life cycles of these policies can be managed through standard life-cycle management technologies, and have a life cycle that is independent of the applications.
 - a. Dispatch policies can be configured that influence where batch jobs are executed. For example, a dispatch policy can be defined where all jobs of a certain type must run in a 64-bit JVM.
 - b. Partitioning policies can be configured that define how batch jobs should be broken into parallel processing elements. These policies typically match how the data used by the batch applications have been partitioned. Dispatch policies can be defined in conjunction with the parallel partitioning policies to create a highly parallel solution with data-aware routing.
 - c. Job-specific qualities of service (QoS) can be configured that influence how jobs are executed within the batch container.
2. Both the batch container tier, as well as the underlying infrastructure cloud, conveys capacity and execution metrics to the job-dispatching tier. The metrics are used by the job-dispatching tier to determine the best endpoint to execute the jobs on. This system uses autonomics algorithms to ensure that jobs are load-balanced across the system.
3. As spikes in batch jobs occur, elasticity services, where the batch container and infrastructure cloud must scale up or down to meet demand, may be necessary. The z/OS platform has these services built in, and WebSphere Compute Grid takes full advantage of this service. On distributed platforms, WebSphere Compute Grid can be combined with WebSphere Virtual Enterprise to create an elastic infrastructure.

WebSphere Compute Grid has been designed to deliver a sophisticated batch-processing platform that is capable of executing highly parallel workloads within the cloud. There are two critical elements within the design that enable this:

- ▶ The policy-based dispatching, coupled with the autonomics, ensure the system continually optimizes the execution;
- ▶ The degree of parallelism is a point-in-time decision, where the deadlines for the jobs, relative job priority, SLA's for competing workloads, and available system resources influence how many parallel instances will execute.

Figure 3 illustrates the view of executing highly parallel workloads within the cloud.

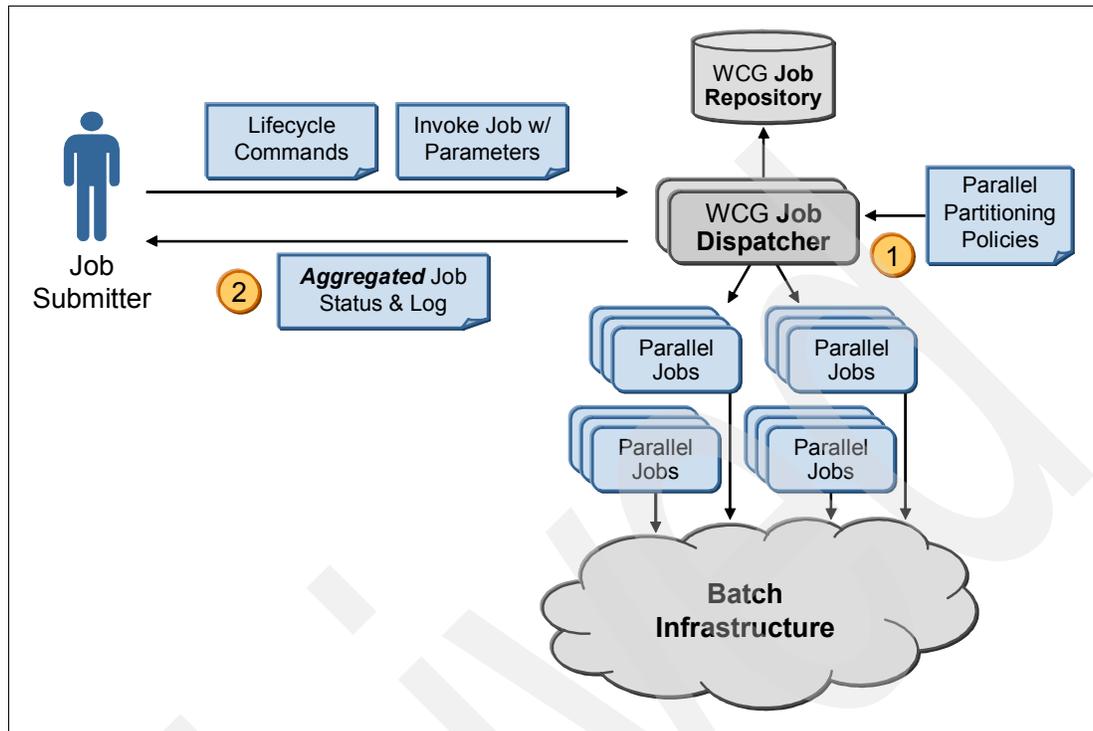


Figure 3 WebSphere Compute Grid can execute highly parallel batch workloads within the cloud

As the use of cloud computing evolves, WebSphere Compute Grid, WebSphere Virtual Enterprise, WebSphere Application Server for z/OS, WebSphere Cloudburst, WebSphere eXtreme Scale, and other cloud-enabling technologies will serve as the foundation. Mixed application workloads (OLTP, batch, message-driven, and others) will run within highly virtualized infrastructures. Physical resources, such as CPU and memory, as well as logical resources like database locks, will be manipulated to ensure that SLA's across different workloads are met by the overall system. Dynamic provisioning coupled with elastic applications will ensure the system can tolerate spikes in application demand.

The batch vision with WebSphere Compute Grid

Batch workloads are not unique to the mainframe. Batch processing exists on many platforms, in numerous applications within the enterprise. The location of the data that a batch application needs should dictate where that batch application should run. Proximity of data in batch processing is critical for performance. Enterprises today have portable OLTP applications, where standards such as J2EE and infrastructure components like WebSphere Application Server facilitate the portability. Batch workloads, however, were never standardized, nor were platform-neutral batch platforms delivered to the market prior to WebSphere Compute Grid. As a result, OLTP infrastructures and applications are standardized across an enterprise, while batch applications were built in silos, as illustrated in Figure 4 on page 7. These silos impact time-to-market where the infrastructure is not agile. Business processes are duplicated across the enterprise and consolidation and workload relocation efforts are impeded. To solve these problems, WebSphere Compute Grid enables enterprises to achieve a Unified Batch Architecture (UBA).

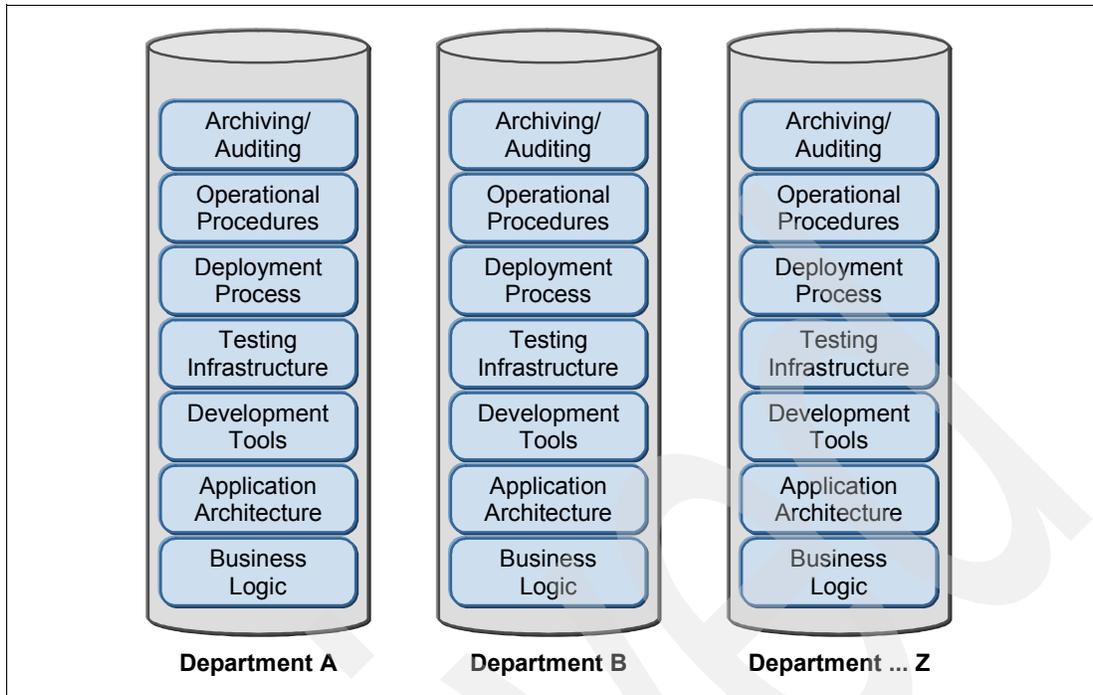


Figure 4 Batch application infrastructures exist in silos within enterprises today

The UBA strategy is simple: the standardization of business processes and technology across the enterprise ensure that enterprises remain agile and efficient. As Figure 5 illustrates, pursuing a UBA strategy eliminates numerous redundancies. Application libraries and functions can be shared across the enterprise. Common disaster recovery, archiving or auditing, and operational procedures ensure that resources such as people and hardware are not wasted.

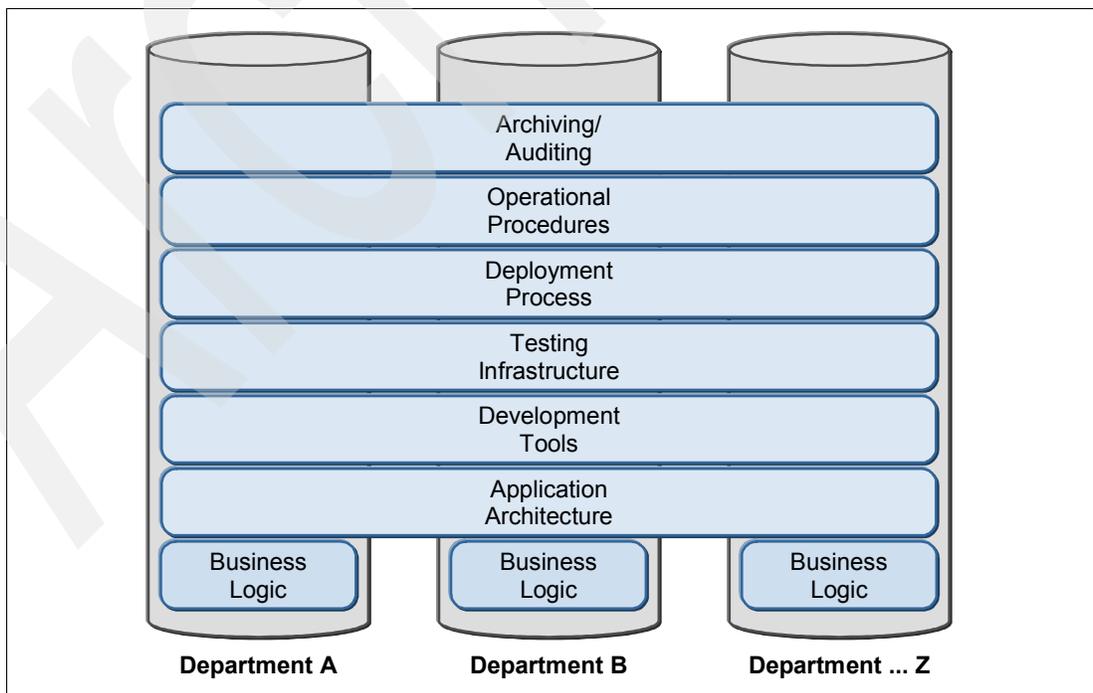


Figure 5 The goal of the UBA strategy is to standardize as much of the batch infrastructure as possible across the enterprise, increase agility, and reduce costs

The batch vision, as illustrated in Figure 6, is simple:

- ▶ A common batch processing runtime within the enterprise should exist.
 - This enterprise batch platform should integrate with existing business processes, such as archiving or auditing systems, security infrastructures, enterprise schedulers, and so forth.
- ▶ Batch applications should be portable, and run where its data is located.
- ▶ To maximize performance, the batch platform must integrate where possible (for example with z/OS i/o optimizations) to maximize performance on the mainframe.
- ▶ The location of batch application data should dictate where that application should run. Since most enterprise data resides on the mainframe, first-class mainframe integration is crucial.

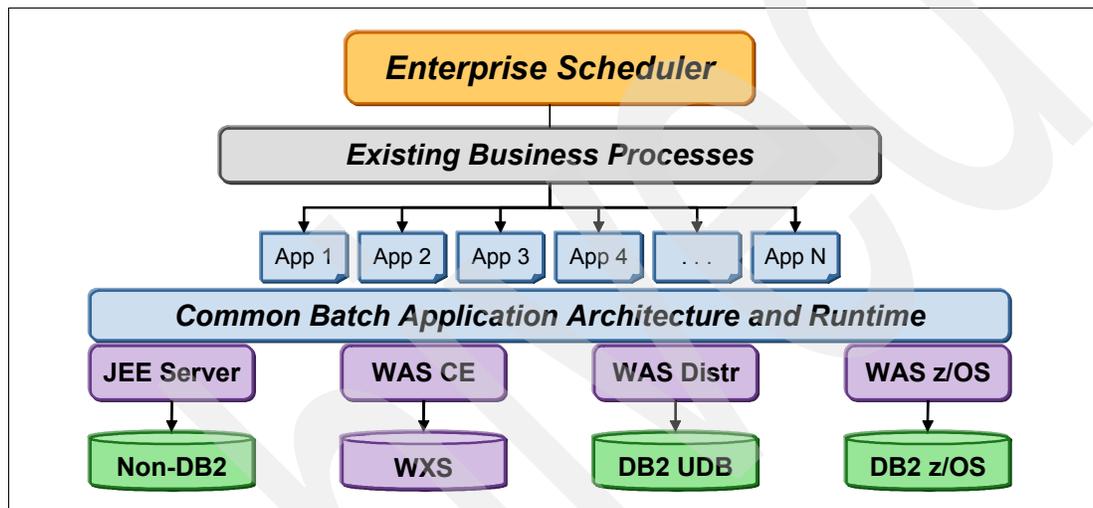


Figure 6 The batch vision

A sample BVA

In this section, we examine an example BVA that was conducted for a fictional customer.

Customer business case

The business case for Compute Grid has a few major elements:

- ▶ Reduce home-grown batch infrastructures, where we reduce development costs by replacing customer-developed infrastructure code with Compute Grid. Think of this as replacing a home-grown application server with WebSphere Application Server. The customer no longer must deal with developing and maintaining server code. Instead they can focus on developing business logic and delivering real value.
- ▶ Reduce general CP usage charges on the mainframe by shifting workloads to zAAP processors. zAAP processors are specialty engines on the IBM® System z® mainframes dedicated to processing Java workloads. zAAP workloads fall under a different licensing and pricing schemes than non-Java workloads (Cobol, Assembler, PLX, and so forth).
- ▶ Reduce redundant development and infrastructure costs by sharing services across batch and OLTP.

Scope and goals

The purpose of the BVA is to document the WebSphere Compute Grid solution investment streams and associated benefits pertaining to a valued IBM customer's batch modernization and cost reduction initiatives. The focus of this assessment is to analyze the customer's current batch challenges and conservatively to quantify benefits associated with a modernized batch environment through WebSphere Compute Grid.

The following business and/or IT goals were defined by the customer SMEs:

- ▶ Shift from general-purpose MIPS on z/OS to IBM System z Application Assist Processor (zAAP)-eligible MIPS, which reduce costs
- ▶ A desire to improve time to market by using modern design patterns, tooling, open source libraries, and container-managed infrastructure services to speed development.
- ▶ A move towards a shared-services infrastructure run on WebSphere Application Server, where the group has common batch and OLTP components, such as business services, application development process, testing and deployment infrastructure, operational management, disaster recovery, and security.
- ▶ Alleviate any staffing issues due to the diminishing availability of COBOL skills, which could impact the group's ability to deliver new functions

BVA methodology

The BVA methodology uses a five step process to weigh project costs versus tangible benefits, intangible benefits and project risk, extending traditional ROI to better measure and manage line-of-business and IT related value. The methodology helps uncover hidden project costs, identifies key business-orientated benefits, as well as IT benefits, and allows quick assessment of different scenarios. Benefits are grouped into the following categories:

- ▶ **IT Benefits**
This includes the savings in IT spending such as hardware, software, professional services, and infrastructure, plus improvements in IT productivity.
- ▶ **Business Operating Efficiencies**
This includes the savings to business operating expenses such as the benefits of capital cost avoidance, business productivity improvements, asset use improvements and reduced application or project backlogs.
- ▶ **Business Strategic Advantages**
This includes the incremental improvements in revenue and profitability that are associated with the solution.

Within each category there may be both direct and indirect benefits:

- ▶ **Direct benefits** have a first order effect based on implementing the solution. The benefits are hard savings or improvements, which are quantifiable and can be directly associated with a recommendation or action.
- ▶ **Indirect benefits** are secondary benefits that are harder to quantify but are real and need to be accounted for. Indirect benefits may be discounted to reflect both risk and assumption. For example, if the solution will result in decreased system downtime, workers or customers may benefit from the increased availability. Productivity, revenue, and satisfaction may all improve.

The five steps that are performed during the assessment are depicted in Figure 7.

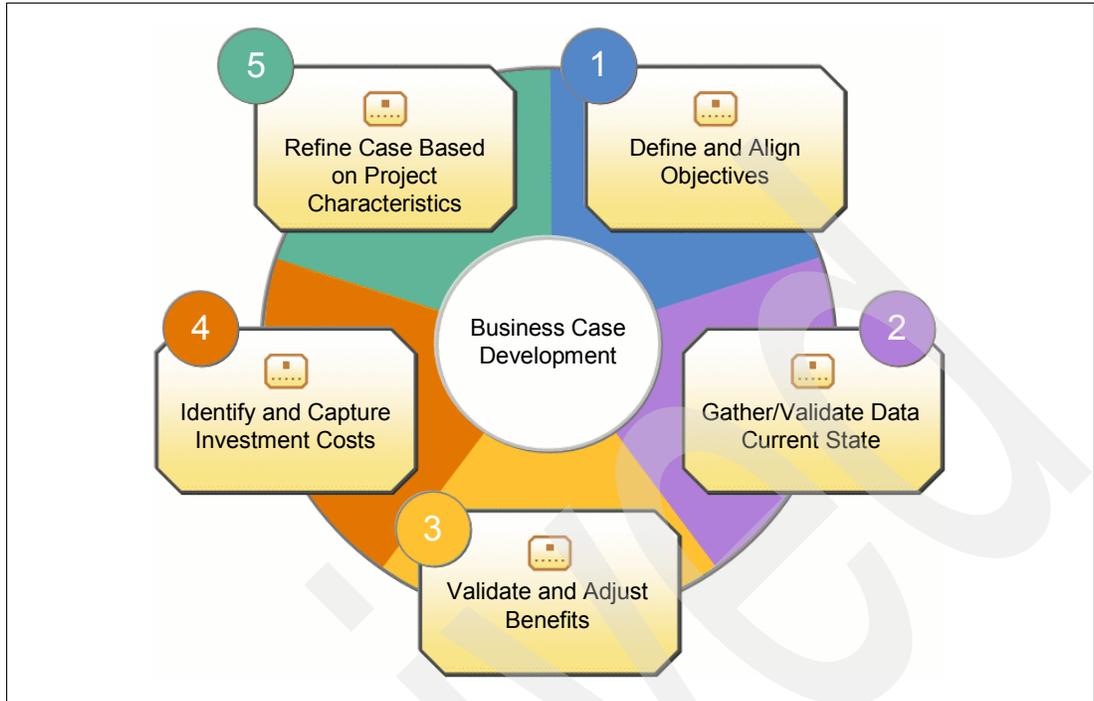


Figure 7 BVA methodology steps

In step 1, the customer's industry profile, business objectives, high-level use cases, and scope are verified and clarified to help define and align the solution and supporting components.

During step 2, information is gathered on the customer's current environment to establish the baseline from which the benefits will be derived. This information includes both line-of-business and IT resources plus associated metrics as appropriate.

In step 3, the benefits that have been selected are validated and refined, with the potential for new benefits to be added and others dropped based upon facilitated discussions with the customer's SMEs. At this time the implementation schedule and realization or adoption curve are taken into consideration using the customer's projected deployment roadmap.

During step 4, the estimated investments associated with the solution are captured.

In step 5, specific project risks, as well as any accelerators are taken into consideration. Also, specific benefits or costs may be tuned more granularly, such as defining unique realization schedules for benefits based on implementation plans or allocating particular costs with financing options in mind.

Steps 3–5 may be performed iteratively as the return on investment results are analyzed and additional considerations arise.

Risk adjusted analysis and cumulative benefits

A risk-adjusted analysis of the proposed solution's impact was conducted and it was projected that implementing the proposed WebSphere Compute Grid solution would result in \$14,078,654 worth of five year cumulative benefits. Of these projected benefits, \$13,183,738 are direct benefits and \$894,916 are indirect benefits.

Top cumulative benefits for the project include:

- ▶ Reducing home-grown batch development costs through Compute Grid container-managed services: \$7,320,230
- ▶ Reducing general CPU usage charges through zAAP offload, MSU cost reductions: \$2,420,682
- ▶ Reducing home-grown batch on-going code maintenance related charges: \$1,702,711
- ▶ Reducing service level agreement penalties due to competing OLTP service levels and batch deadlines: \$994,351
- ▶ Improving and accelerating innovation: \$894,916
- ▶ Reducing redundant development costs through shared services across execution domains (OLTP and batch): \$745,763

These benefits can be grouped regarding business impact as follows:

- ▶ \$12,189,387 in IT cost reductions
- ▶ \$994,351 in business operating efficiency improvements
- ▶ \$894,916 in business strategic advantage benefits

The proposed project is expected to help the company meet the following goals and gain the following benefits:

- ▶ Reduce application development deployment and maintenance cost: \$9,768,705
- ▶ Improve IT resource use: \$2,420,682
- ▶ Improve IT system availability and service levels: \$994,351
- ▶ Improve time to market for new offerings: \$894,916

The proposed project is expected to deliver the following benefits to specified stakeholders:

- ▶ Application development: \$9,768,705
- ▶ Information technology: IT \$2,420,682
- ▶ Strategic and executive management: \$994,351
- ▶ Finance and accounting: \$894,916

To implement the proposed project will require a five-year cumulative investment of \$1,850,000 including:

- ▶ \$1,100,000 in capital expenditures
- ▶ \$750,000 in operating expenditures

Comparing the costs and benefits of the proposed solution using discounted cash flow analysis and factoring in a risk-adjusted discount rate of 6.5%, the proposed business case predicts:

- ▶ Risk Adjusted Return on Investment (RA ROI) of 570%
- ▶ Return on Investment (ROI) of 661%
- ▶ Net Present Value (NPV) savings of \$9,731,116
- ▶ Internal Rate of Return (IRR) of 176%
- ▶ Payback period of 10.0 month(s)

The solution has been risk-adjusted for an overall deployment schedule of three months, realized benefits to include 100% of direct benefits and 60% of indirect benefits, plus a benefit adoption curve of 75% for the first year and 100% over each successive year of the analysis.

In summary, the customer can derive substantial value (for example, cost reductions, time-to-market improvements, and so forth) by modernizing the current batch environment through the use of WebSphere Compute Grid.

Benefit summary

As shown in Table 1, the proposed solution is expected to deliver \$14,078,654 over the 5 year analysis period, with \$13,183,738 in direct (hard) benefits, and \$894,916 in indirect (soft) benefits.

Table 1 Benefits summary

Benefits Summary	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Total Benefits	\$1,470,825	\$2,627,610	\$3,042,645	\$3,318,041	\$3,619,533	\$14,078,654
IT Cost Reductions						
Reduce home-grown batch development costs via Compute Grid container managed services	\$750,150	\$1,350,270	\$1,576,982	\$1,734,680	\$1,908,148	\$7,320,230
Reduce general CP usage charges via zAAP offload - MSU cost reductions	\$248,063	\$446,513	\$521,483	\$573,631	\$630,994	\$2,420,682
Reduce home-grown batch ongoing code maintenance related charges	\$174,488	\$314,078	\$366,812	\$403,493	\$443,842	\$1,702,711
Reduce redundant development costs - via shared services across execution domains - OLTP and batch	\$84,375	\$146,250	\$163,406	\$171,577	\$180,155	\$745,763
Total IT Cost Reductions	\$1,257,075	\$2,257,110	\$2,628,682	\$2,883,380	\$3,163,139	\$12,189,387
Business Operating Efficiency						
Reduce Service Level Agreement Penalties - due to competing OLTP service levels and batch deadlines	\$112,500	\$195,000	\$217,875	\$228,769	\$240,207	\$994,351
Total Business Operating Efficiency	\$112,500	\$195,000	\$217,875	\$228,769	\$240,207	\$994,351

Business Strategic Advantage						
Improve and Accelerate Innovation (indirect)	\$101,250	\$175,500	\$196,088	\$205,892	\$216,187	\$894,916
Total Business Strategic Advantage	\$101,250	\$175,500	\$196,088	\$205,892	\$216,187	\$894,916
Direct Benefits	\$1,369,575	\$2,452,110	\$2,846,557	\$3,112,149	\$3,403,346	\$13,183,738
Indirect Benefits	\$101,250	\$175,500	\$196,088	\$205,892	\$216,187	\$894,916

Figure 8 illustrates these benefits in a chart.

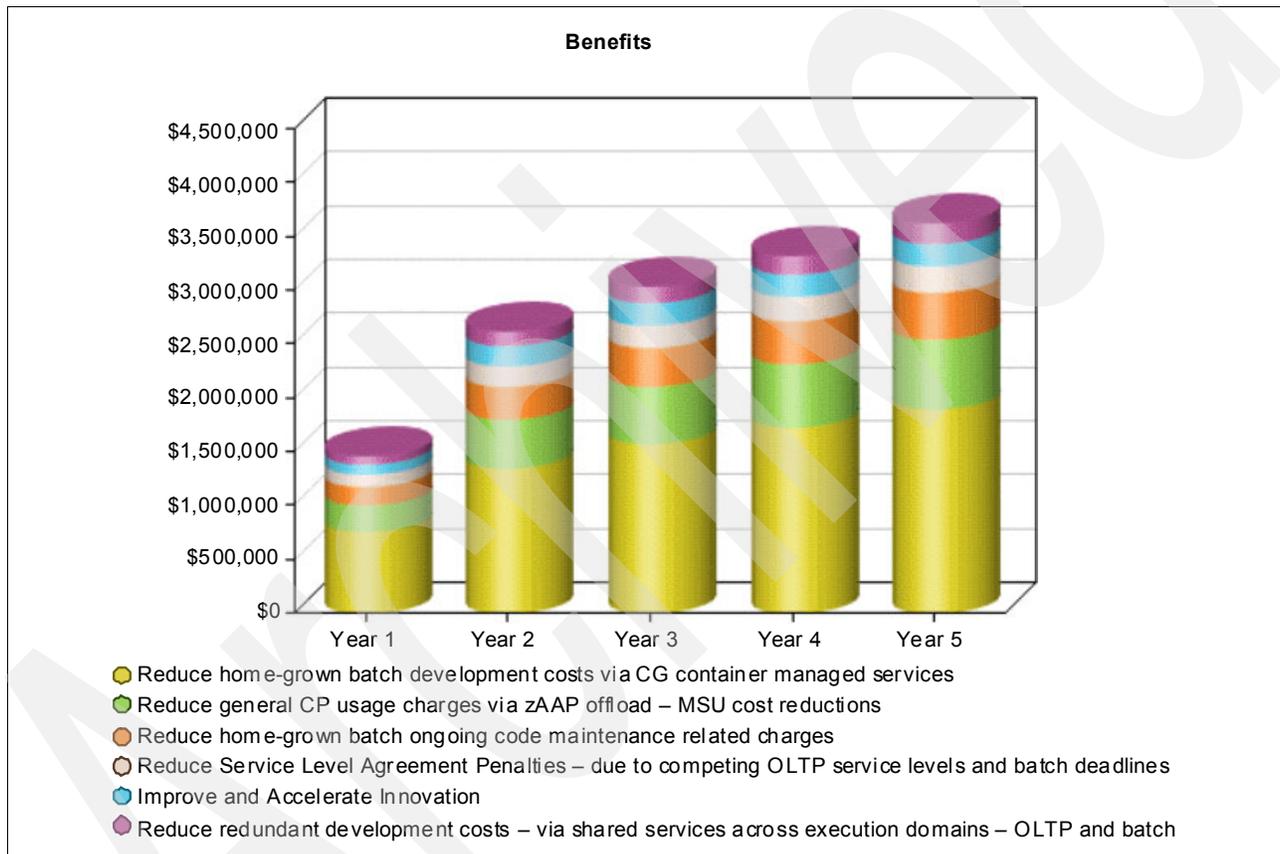


Figure 8 Graphical view of the benefit summary

The pie chart in Figure 9 helps show where the top benefits lie.

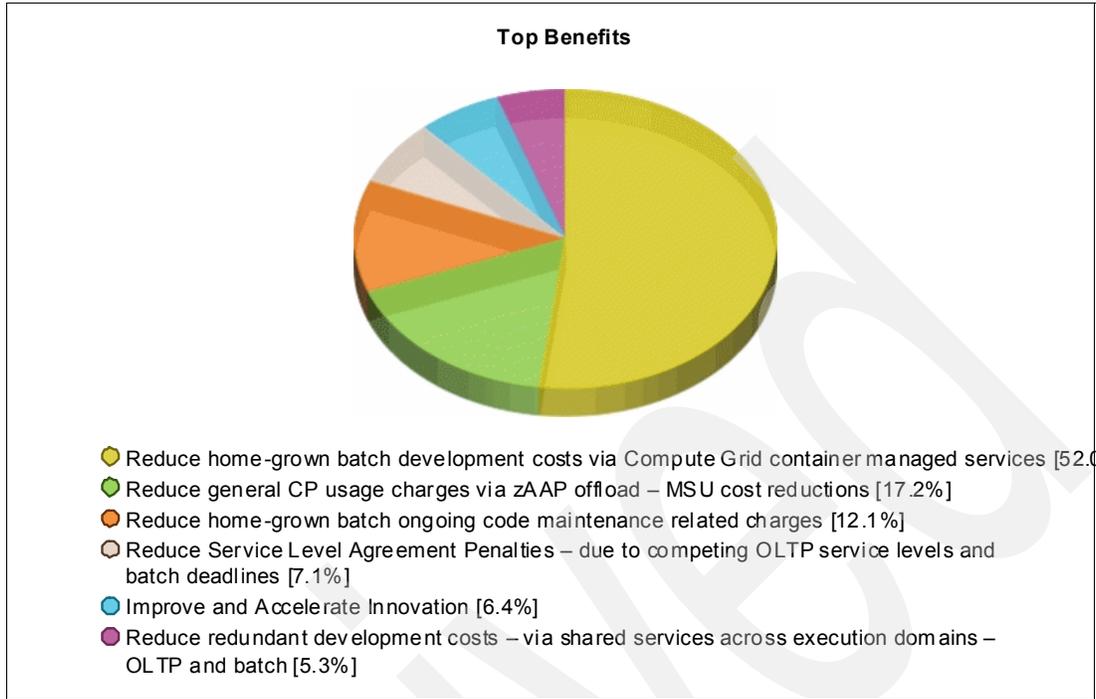


Figure 9 Summary of top benefits

Figure 10 shows a view of the benefits by category and the value of the direct versus the indirect benefits.

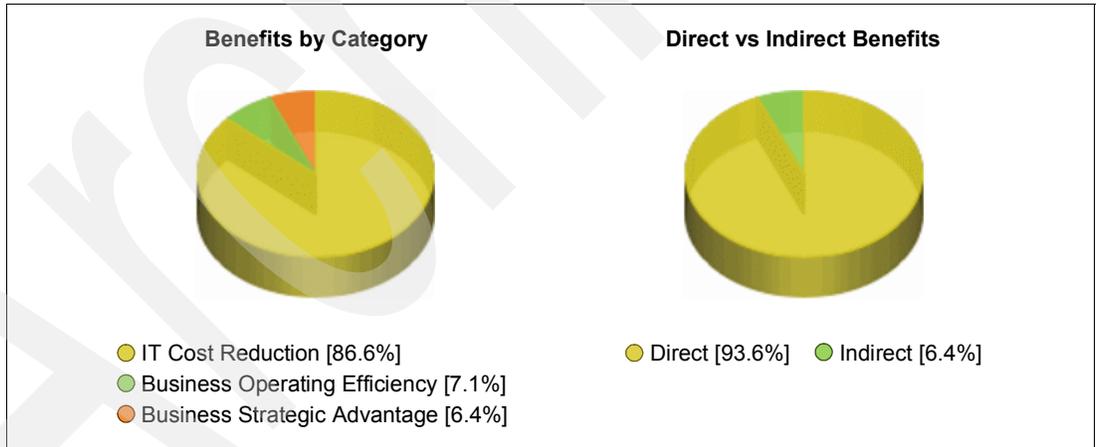


Figure 10 Benefits by category and direct versus indirect benefits

Figure 11 shows the benefits by goal and the benefits with respect to the stakeholders.

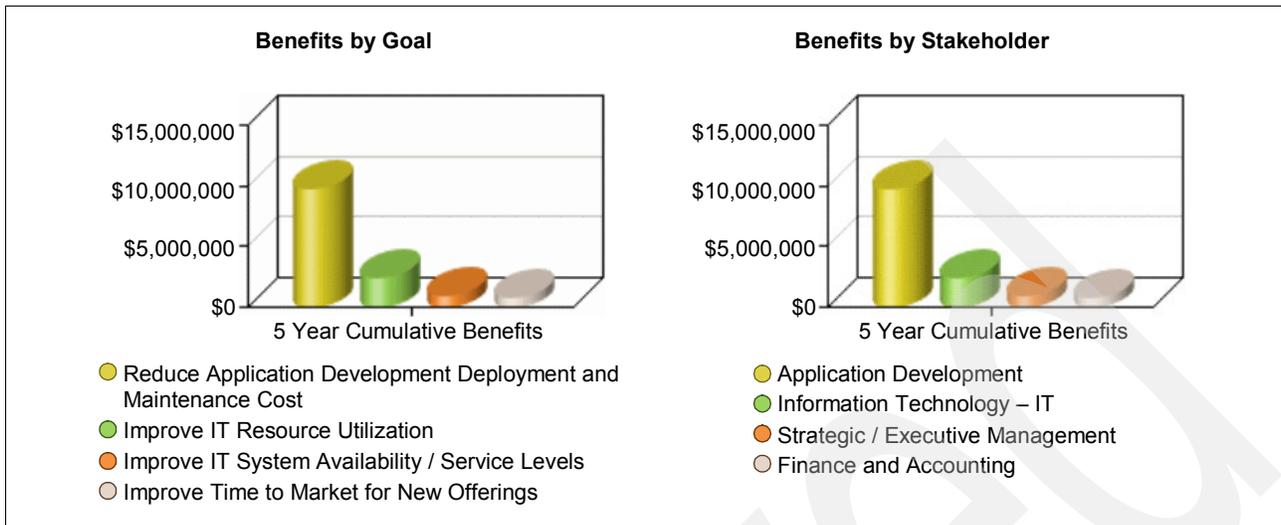


Figure 11 Benefits by goal and by stakeholder

Investment summary

To implement the proposed project will require a five year cumulative investment of \$1,850,000 (Table 2 on page 16) including:

- ▶ \$1,100,000 in capital expenditures
- ▶ \$750,000 in operating expenditures

Table 2 Investment summary by year

Investment Summary	Initial	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Total Investment	\$1,000,000	\$170,000	\$170,000	\$170,000	\$170,000	\$170,000	\$1,850,000
Capital Investment							
WebSphere Compute Grid Licensing Cost (IT)	\$750,000	\$0	\$0	\$0	\$0	\$0	\$750,000
Solution Hardware Cost - zAAP related acquisitions (IT)	\$250,000	\$20,000	\$20,000	\$20,000	\$20,000	\$20,000	\$350,000
Software and Operating Systems Cost (IT)	\$0	\$0	\$0	\$0	\$0	\$0	\$0
Network and Communications Costs (IT)	\$0	\$0	\$0	\$0	\$0	\$0	\$0
Total Capital Investment	\$1,000,000	\$20,000	\$20,000	\$20,000	\$20,000	\$20,000	\$1,100,000
Operating Investment							
WebSphere Support and Maintenance Cost (IT)	\$0	\$0	\$150,000	\$150,000	\$150,000	\$150,000	\$600,000
IBM Professional Services Cost (IT)	\$0	\$75,000	\$0	\$0	\$0	\$0	\$75,000
Project Implementation and Migration Costs (IT)	\$0	\$60,000	\$0	\$0	\$0	\$0	\$60,000
Training Cost (IT)	\$0	\$15,000	\$0	\$0	\$0	\$0	\$15,000
Ongoing WebSphere Administration Cost (IT)	\$0	\$0	\$0	\$0	\$0	\$0	\$0
Internal Implementation Labor (IT)	\$0	\$0	\$0	\$0	\$0	\$0	\$0
Independent Learning Costs (IT)	\$0	\$0	\$0	\$0	\$0	\$0	\$0
Total Operating Investment	\$0	\$150,000	\$150,000	\$150,000	\$150,000	\$150,000	\$750,000

Figure 12 shows the investment summary in a chart.

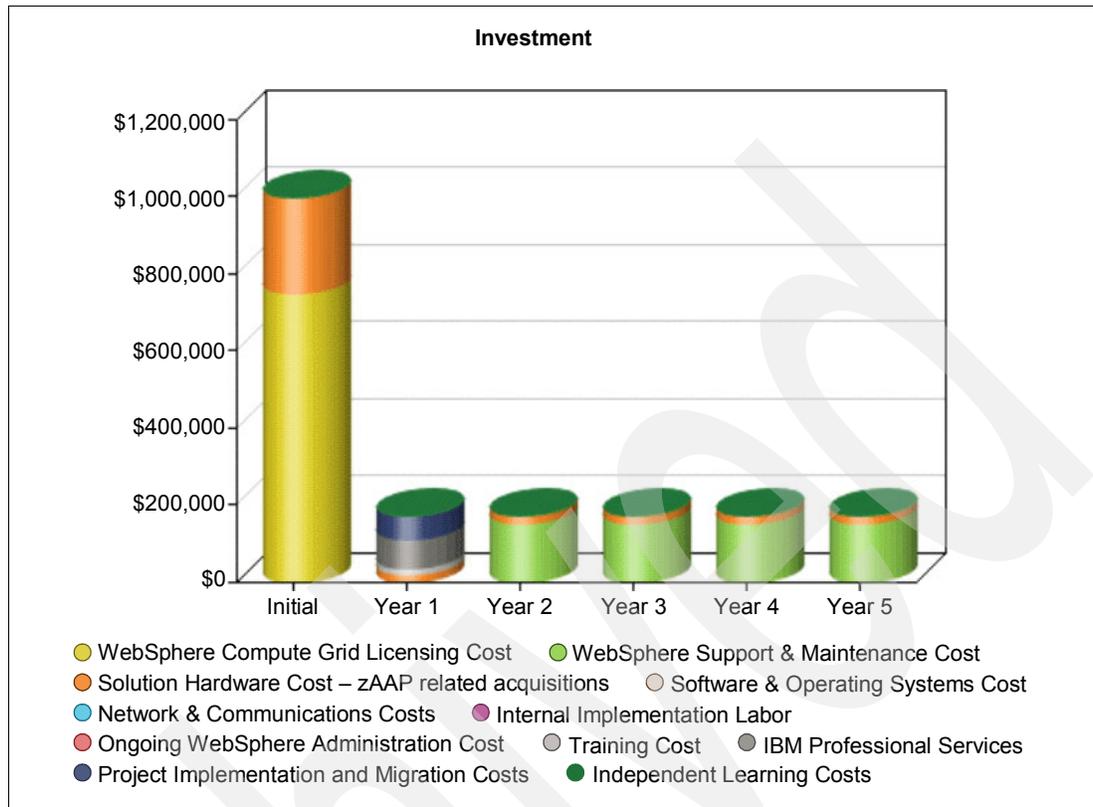


Figure 12 Investment summary - graphical view

Figure 13 shows a graphical view of the investment summary by category and expense type.

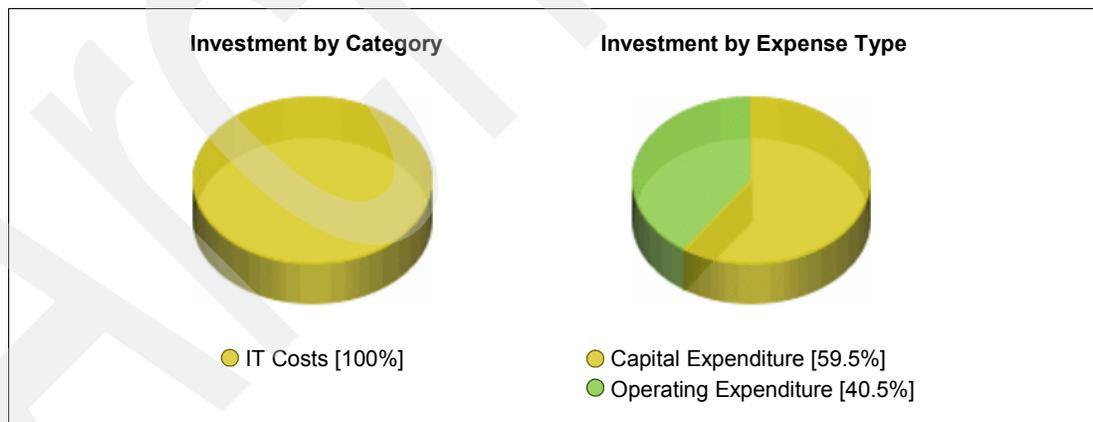


Figure 13 Investment summary by category and expense type

ROI analysis

Analyzing the WebSphere Extended Deployment Compute Grid investment and benefit streams - cash flow and key financial metrics were calculated as follows (See Table 3):

- ▶ Risk Adjusted Return on Investment (RA ROI) of 570%
- ▶ Return on Investment (ROI) of 661%
- ▶ Net Present Value (NPV) savings of \$9,731,116

- ▶ Internal Rate of Return (IRR) of 176%
- ▶ Payback period of 10.0 month(s)

Table 3 ROI analysis

ROI Analysis (Probable Case)	Initial	Year 1	Year 2	Year 3	Year 4	Year 5
Benefits	\$0	\$1,470,825	\$2,627,610	\$3,042,645	\$3,318,041	\$3,619,533
Cumulative Benefits		\$1,470,825	\$4,098,435	\$7,141,080	\$10,459,121	\$14,078,654
Investment	\$1,000,000	\$170,000	\$170,000	\$170,000	\$170,000	\$170,000
Cumulative Investment	\$1,000,000	\$1,170,000	\$1,340,000	\$1,510,000	\$1,680,000	\$1,850,000
Cash Flow	(\$1,000,000)	\$1,300,825	\$2,457,610	\$2,872,645	\$3,148,041	\$3,449,533
Cumulative Cash Flow	(\$1,000,000)	\$300,825	\$2,758,435	\$5,631,080	\$8,779,121	\$12,228,654
ROI	661%					
Risk Adjusted ROI	570%					
NPV Savings	\$9,731,116					
IRR	176%					
Payback period (including deployment period)	10 month(s)					
Risk Adjusted Discount Rate	6.5%					

Figure 14 shows the ROI analysis in a chart.

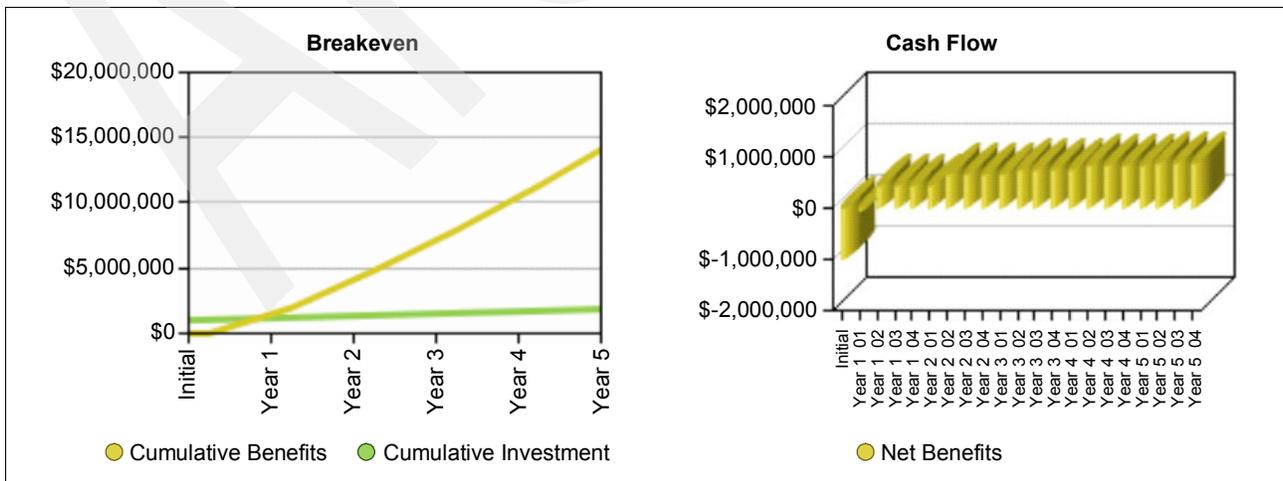


Figure 14 ROI analysis - graphical view

Additional resources

The following resources can provide additional information.

- ▶ Customer reference:
 - Swiss Re and IBM collaborate on modern batch processing for z/OS
<http://www-01.ibm.com/software/tivoli/features/ccr2/ccr2-2008-12/swissre-websphere-compute-grid-zos.html>
- ▶ Batch Application Design:
 - Designing Batch Applications
<http://www.ibm.com/developerworks/forums/thread.jspa?threadID=275702&tstart=0>
 - Introduction to batch programming using WebSphere Extended Deployment Compute Grid
http://www.ibm.com/developerworks/websphere/techjournal/0801_vignola/0801_vignola.html
- ▶ Batch Infrastructure:
 - Emerging technologies make WebSphere Extended Deployment Compute Grid ideal for handling mission-critical batch workloads
http://www.ibm.com/developerworks/websphere/techjournal/0804_antani/0804_antani.html
 - WebSphere Compute Grid Wiki
<https://www.ibm.com/developerworks/wikis/display/xdcomputegrid/Home>
 - *Batch Modernization on z/OS, SG24-7779*
<http://www.redbooks.ibm.com/abstracts/sg247779.html>
- ▶ Get Started:
 - How to Create a WebSphere Compute Grid Test Server
<http://www.ibm.com/developerworks/forums/thread.jspa?threadID=227667&tstart=0>
 - New and Improved: Eclipse Workspace for Batch Development
<http://www.ibm.com/developerworks/forums/thread.jspa?threadID=228339&tstart=0>
 - Ask WebSphere Compute Grid experts a question
<http://www.ibm.com/developerworks/forums/forum.jspa?forumID=1240&start=0>

About the author



Snehal Antani works for WebSphere Product Management, where he leads the strategy for WebSphere's cloud computing platforms and runtimes. Prior to this role, he worked for IBM Software Services for WebSphere (ISSW) as a senior managing consultant. He led the consulting practice for the WebSphere XD suite of products, which includes: WebSphere Virtual Enterprise, WebSphere Compute Grid, and WebSphere eXtreme Scale. His focus is on grid/cloud/HPC/middleware architectures and design. He serves as a technical advisor to key customers, driving clients towards production, and speaking at industry conferences worldwide. Prior to joining ISSW, Snehal worked in product development for both WebSphere Application Server for z/OS and WebSphere

Extended Deployment. He has worked closely with many large distributed and z/OS clients around the world, helping them achieve production with WebSphere products as well as influence their technical strategies. Snehal has disclosed numerous patents and technical publications in the domains of enterprise application infrastructure and high-volume transaction processing. He earned a BS in computer science from Purdue University, and an MS in computer science from Rensselaer Polytechnic Institute (RPI) in Troy, NY with a thesis in the area of quantifying and improving the resiliency of middleware infrastructures. You can contact Snehal at <http://snehalantani.googlepages.com>

Contributors

The BVA example has been provided by Glen Ritche, Business Value Assessment Team, IBM.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Archived

Contact an IBM Software Services Sales Specialist



Start SMALL, Start BIG, ... **JUST START** architectural knowledge, skills, research and development . . . **that's IBM Software Services for WebSphere.**

Our highly skilled consultants make it easy for you to design, build, test and deploy solutions, helping you build a smarter and more efficient business. **Our worldwide network of services specialists wants you to have it all!** Implementation, migration, architecture and design services: IBM Software Services has the right fit for you. We also deliver just-in-time, customized workshops and education tailored for your business needs. You have the knowledge, now reach out to the experts who can help you extend and realize the value.

For a WebSphere services solution that fits your needs, contact an IBM Software Services Sales Specialist:
ibm.com/developerworks/websphere/services/contacts.html

This document, REDP-4566-00, was created or updated on January 11, 2010.



Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>



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

AIX®	System z®	z/OS®
IBM®	Tivoli®	
Redbooks (logo)  ®	WebSphere®	

The following terms are trademarks of other companies:

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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