May 2007

MANAGEMENT BRIEF

VALUE PROPOSITION FOR IBM SYSTEM i

Cost/Benefit Case for JD Edwards EnterpriseOne Deployment in Midsize Businesses

International Technology Group

4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net

TABLE OF CONTENTS

List of Figures

EXECUTIVE SUMMARY

What does the future hold for Oracle's JD Edwards EnterpriseOne users? In a demand-driven world, the dynamics of success for midsize businesses are changing.

No longer is it sufficient for a company to operate efficiently. It must also respond rapidly and flexibly to changing market conditions. Strategies and processes must be adapted to handle faster cycle times. Closer, more responsive relationships must be developed with customers and partners. Lean manufacturing and distribution models and real-time business practices are increasingly mandated.

These shifts are driving a fundamental transformation of enterprise resource planning (ERP) systems. Core transaction processing applications are joined by tools that allow organizations to use information as a new resource for competitive advantage, and enable high value-added interaction with customers, suppliers, partners and co-workers through the Internet and intranets.

Latest-generation JD Edwards EnterpriseOne solutions offer the potential for significant business gains. Their effectiveness, however, will be determined not only by application-level functionality, but also by underlying IT infrastructures.

These must handle software environments that are more complex and technologically sophisticated than most midsize organizations have dealt with in the past. They must also manage increasingly diverse, interdependent workloads and maintain the around-the-clock continuity of service that is mandated for businesses operating in real time in globalized, Internet-connected markets.

How will these challenges be met? This document deals with one potential solution: the IBM System i platform. Specifically, it looks at the business case for deploying EnterpriseOne applications on this platform in midsize companies.

This case is based on the following:

- ***Distinctive capabilities***. Companies that have deployed EnterpriseOne solutions on the System i report that its ability to maintain extremely high levels of availability – often on a near-24x365 basis – is a critical strength. Among 52 users surveyed for this report, 50 (96 percent) cited this as a principal benefit of employing this platform.

  Other System i capabilities valued by users included ease of management, reflected in low IT staffing; high levels of stability and system integration; partitioning and virtualization capabilities optimized for mixed workload environments; and distinctive System i strengths in disaster recovery, security, virus protection and other areas that reduced organizational risk exposure.

- ***Lower IT costs***. If overall IT expenditures for implementing, supporting and administering servers over time are compared, costs for EnterpriseOne deployment on System i may be significantly lower than for other platforms.

  This is illustrated by three composite profile-based comparisons of manufacturing companies presented in this report. In these comparisons, three-year IT costs for System i-based scenarios averaged 28.5 percent less than those for Microsoft Windows server equivalents; 38.0 percent less than those for Hewlett-Packard (HP) Integrity server equivalents; and 41.6 percent less than those for Sun Solaris server equivalents.

  Figure 1 summarizes these results.

Figure 1
Profile Companies: Average Three-year IT Costs for IBM System i, HP Integrity, Sun Solaris and Windows Server Scenarios

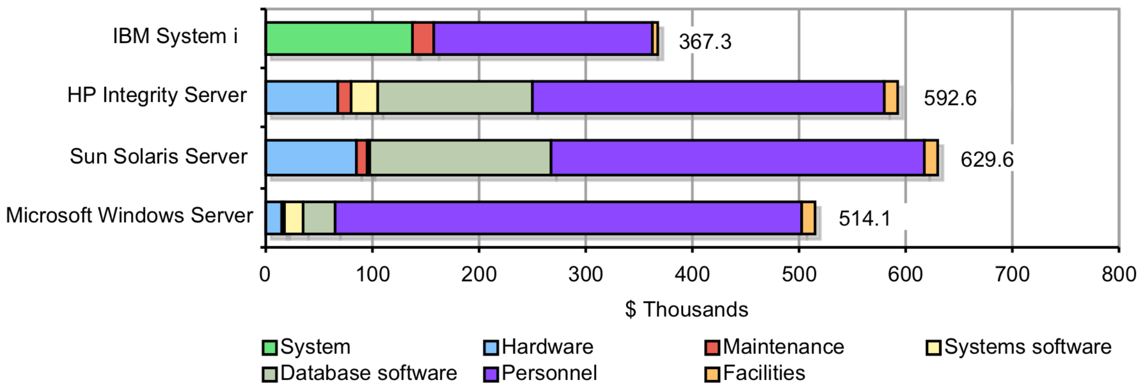

Higher IT costs for HP Integrity, Sun Solaris and Windows server scenarios are due in large part to higher full time equivalent (FTE) staffing levels and personnel costs for system administration, database administration and related functions; and (for HP Integrity and Sun Solaris scenarios) costs of Oracle database software.

In System i scenarios, 525 and 550 platforms handle application and Web serving, as well as database serving functions. Advanced partitioning, virtualization and management capabilities enable these platforms to handle three-tier workloads more efficiently than conventional server architectures, and provide additional performance and availability benefits.

Lower IT costs for System i scenarios also reflect comparatively aggressive IBM pricing for the new System i 525, as well as reductions in i5/OS costs for 550 platforms due to use of the IBM application server licensing option. Introduced in April 2007, this offers lower-cost license fees for processor-based partitions that do not require use of the i5/OS database.

Profiles are based on automotive parts, food and beverage, and electronics manufacturers with between $100 million and $650 million in sales, and 800 to 4,000 employees. Detailed profiles and calculations are presented in the Detailed Data section of this report.

- ***Lower costs of downtime***. More than a decade of experience with ERP systems has shown that outages impact bottom lines. Operations may be disrupted, personnel and capacity idled, orders and shipments delayed, and a wide range of other business activities affected. Customers may also be alienated, and their business lost or reduced.

  Recent trends have increased vulnerability. A tightly integrated supply chain operating in real time, with minimal inventories, is a great deal more vulnerable to outages than a looser, less efficient structure. In such environments, disruptions tend to generate "cascading effects" that spread rapidly across the entire organization, and extend to customers and suppliers.

  These effects, moreover, may continue to be felt long after service has been restored. Even brief outages may continue to affect businesses for days or even weeks. If orders or customers are lost as a result, bottom-line impacts may be significantly greater and of longer duration.

As organizations move toward 24x7 business operations, they are increasingly exposed to the effects not only of unplanned (i.e., accidental) outages, but also of planned downtime for such tasks as hardware changes, software upgrades and routine maintenance. In the past, such tasks might be scheduled to occur outside normal business hours. In a growing number of businesses, however, "normal business hours" are now 24 hours per day, 365 days per year.

There is thus a direct correlation between system availability and the effects of bottom-line business disruption. For the three profile companies described above, differences in availability levels for System i, HP Integrity, Sun Solaris and Windows server scenarios translate into major variations in costs of downtime over a three-year period. Figure 2 summarizes these costs.

Figure 2
Profile Companies: Three-year Costs of Downtime by Scenario

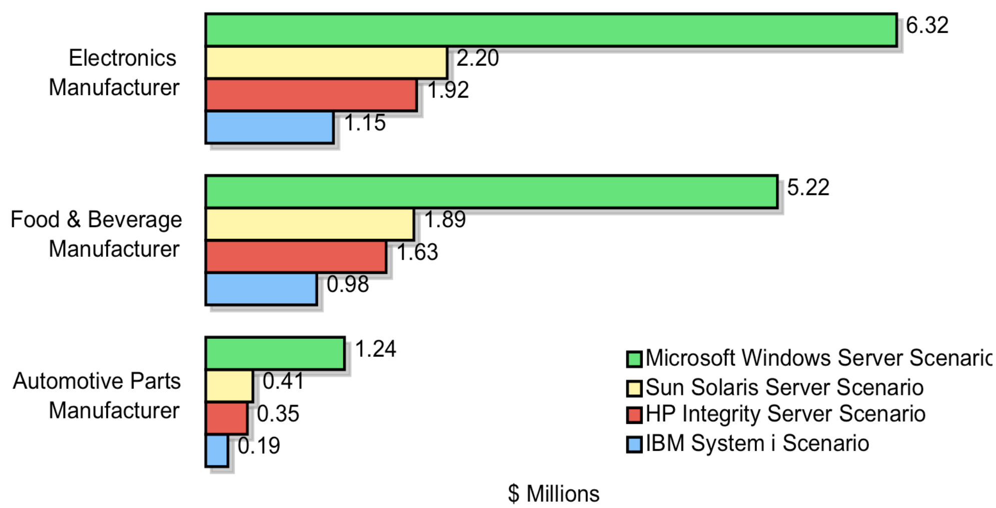

Overall three-year costs of downtime for System i scenarios average 81.8 percent less than those for Windows server scenarios, and 40.4 percent and 48.4 percent less than those for HP Integrity and Sun Solaris scenarios respectively.

Costs of downtime in this presentation include costs of idle and underutilized capacity and personnel for production and logistics operations; costs of scheduling and production setup changes; costs of order, shipment and payment delays; additional inventory carrying costs; customer-related costs including late delivery and imperfect order fees; and other components. The basis of these calculations is presented in the Detailed Data section of this report.

User experiences with EnterpriseOne deployment on the System i also highlight a broader benefit. The distinctive capabilities of this platform materially reduce the technical complexities with which organizations must deal.

In the past, the effects of IT complexity have debilitated the ERP strategies of more than a few large corporations. In midsize companies with small IT staffs and limited technical and financial resources, the damage may be greater and longer lasting.

It is from this perspective that the business case for the System i platform should be viewed. It enables midsize users to realize the business benefits of EnterpriseOne solutions, while minimizing the downsides that use of latest-generation technologies might otherwise represent for them.

BUSINESS CASE

Introduction

Since the JD Edwards EnterpriseOne ERP system (originally known as OneWorld) was introduced in 1996, it has been widely deployed on the System i and its predecessors. Worldwide, more than 2,000 EnterpriseOne users run their core systems on this platform.

The continued appeal of the System i reflects multiple factors. One is that this platform has evolved to implement new technologies. There have been no significant lags in System i technological currency compared to UNIX or Windows servers.

A second factor has been that the core strengths of the System i platform in such areas as stability, integration, availability, manageability, recoverability and security have remained relevant – indeed, have become increasingly relevant – through successive generations of ERP systems and business challenges.

The following sections present additional information on and analysis of this business case in three areas – complexity, availability and IT costs. Key trends affecting users are first outlined, and the relevance of System i strengths to these is discussed and illustrated by user survey results.

System i technology differentiators are then detailed. The report concludes with an in-depth presentation of profiles, IT costs and costs of downtime comparisons, and the methodologies used to develop these.

Complexity

A decade ago, an ERP system was primarily a transactional system handling such processes as order management, manufacturing, logistics, purchasing, accounting, payroll and financial management. The focus of ERP deployment was on increasing the efficiency of operational processes.

This picture has changed. Both the structure and technology base of the EnterpriseOne environment, and the business challenges it addresses have undergone major shifts.

Companies have moved aggressively to adopt lean operating models that minimize inventories and cut cycle times across the entire range of supply chain processes. The ability to identify and exploit market trends, along with effective forecasting and planning, organizational agility and responsiveness to customers have increasingly become the vehicles through which competitive advantage is achieved.

Shifts have occurred in two key areas:

1. ***Use of information***. The role of the ERP system has evolved beyond simple transaction processing. It now also collects information from all stages of procurement, production and distribution as well as from all points of contact with customers, suppliers and other partners. The manner in which this information is used is becoming a principal determinant of competitiveness.

   Tools for assembling and using information have become more sophisticated, and more closely integrated into the EnterpriseOne environment. Even relatively small users now routinely employ solutions for advanced forecasting, production and logistics planning, network optimization, product data management, business intelligence and other “informational” functions.

   Customer relationship management (CRM) processes, increasingly, have also become information-driven. Visibility into current inventory levels and production plans enables salespeople to respond more quickly to customer needs and to quote more accurate delivery dates. Real-time order tracking enables better handling of customer queries. Many other such examples might be cited.

2. ***Internet interaction***. In all industries, the Internet has emerged as a pervasive vehicle for interaction with customers, suppliers and partners. New systems allow orders to be placed online, and provide 24x7 access to key information. The ERP environment now extends to a new range of e-commerce, self-service and portal applications.

   The Internet also, increasingly, enables internal collaborative and informational processes. Financial as well as HR self-service systems permit convenient and timely access to information for managers and employees while reducing administrative costs. Portals and dashboards provide interfaces to these and other sources of information for a wide range of users.

The combined result is that today's EnterpriseOne environment, illustrated in figure 3, is functionally much broader, technologically more diverse, and structurally more complex than in the past.

Figure 3
JD Edwards EnterpriseOne Environment

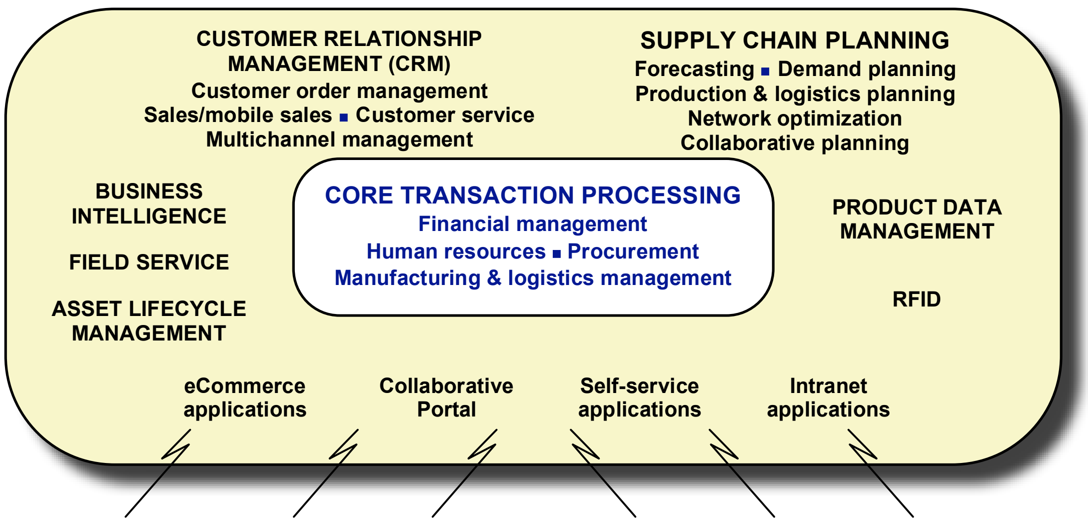

This environment is characterized by application processes and workloads that are not only heterogeneous, but also highly interdependent.

Challenges are magnified by reduced cycle times and "real-time" practices. Traditional batch windows are becoming shorter or disappearing altogether. Application processing, database refreshes and updates, and posting of information for Internet and intranet access now occur around-the-clock.

There are two important implications. First, major new demands are placed on underlying platforms. To support such environments effectively, these must be equipped with sophisticated capabilities for concurrent execution of diverse processes, mixed workload management, and dynamic allocation and re-allocation of resources.

Second, growth in technical complexity easily translates into more IT staff, enhanced technical skills and increased IT expenditure. In complex IT environments, integration and optimization become more challenging processes, and it becomes more difficult to maintain high levels of availability and to protect critical systems and data.

There is also a risk that IT agendas will become dominated by the minutiae of managing and maintaining system infrastructures, to the extent that focus on meeting business objectives is diluted. An organization preoccupied with the basic tasks of making systems work properly will tend to devote less attention and less effort to anticipating and responding to business needs.

Availability

The importance of maintaining high levels of availability for business-critical manufacturing and supply chain operations is widely recognized. The experiences of ERP users in a wide range of industries have demonstrated that even brief outages may have a significant bottom-line impact in terms of lost sales, customer dissatisfaction, organizational disruption and other negative effects.

Moreover, recent shifts in ERP deployment are increasing organizational vulnerability to the effects of outages. Four factors are driving this trend:

1. ***Integration***. For organizations that have deployed ERP systems, replacement of fragmented legacy applications, automation of manual processes, and creation of data transparency across organizations have been primary sources of business benefit. However, there has been a side effect: they have become dependent on their systems.

   EnterpriseOne users surveyed for this report repeatedly commented on this effect. One company noted, for example, that its business was "now almost totally dependent on (EnterpriseOne)" and that any outage would "bring us to a complete halt." Another observed that the company was "fundamentally dependent on EnterpriseOne. If it goes down, everything stops."

   A third company, which estimated the bottom-line cost of downtime at over $1 million per hour, noted that any (system) interruption "would bring us to a complete standstill."

   Vulnerability to disruptions was increased when organizations had consolidated ERP systems. Whereas outages in the past would have disabled only certain segments of the business, they were now more likely to shut down the entire company.

2. ***Supply chain strategies***. Pervasive adoption of lean manufacturing and distribution models and just-in-time (JIT) practices has also had important implications. As inventory buffers and cycle times are reduced, the potential impact of disruptions increases.

   Again, this was clearly recognized by organizations surveyed for this report. An automotive parts manufacturer noted, for example, that the company now received customer orders and queries about product availability on a daily or even hourly basis, and that these required "real-time" response. Outages that impaired response times could "hit hard."

   An electronics distributor made a firm commitment to customers that orders placed by 6 a.m. local time would be shipped the same day. More than 50 percent of the average day's volume was picked, packed, labeled, loaded and expedited during the three hours immediately before the company's shipping deadline. An outage during this period would mean disappointing (and potentially losing) a significant number of customers.

   A fresh foods supplier reported that its produce had to be delivered to retailers within hours of its arrival at their warehouses. Outages affecting EnterpriseOne systems could affect sales, result in spoilage of tons of product, and undermine relationships with customers. Downtime could also disrupt quality control and impair traceability, with potentially even more serious implications.

   A retailer noted that an outage affecting EnterpriseOne-based supply chain operations could delay store deliveries, potentially resulting in stockouts. The bottom-line impact would be particularly serious if they occurred during the company's peak Christmas selling period. Many other such examples might be cited.

Users recognized the significance of cascading effects. A consumer products manufacturer noted, for example, that if an outage disrupted distribution center operations, "it might take us days to catch up." The company's manufacturing plants might also be affected. If changes to production schedules, delays in setups and teardowns, and incoming supplier deliveries were required, it might again take days to overcome backlogs.

The implications of cascading may be simply illustrated. Even a basic manufacturing supply chain will typically involve most or all of the processes summarized in figure 4. This presentation is based on selected segments of the Supply Chain Operations Reference (SCOR) model developed by the Supply Chain Council.

Figure 4
Basic Manufacturing Supply Chain Processes

SOURCE MATERIALS/ COMPONENTS	MAKE PRODUCT	DELIVER PRODUCT	
Identify sources of supply	Schedule production	Process inquiry & quote	Rate shipments
Select supplier(s)	Issue product	Receive & configure order	Receive product
Negotiate with supplier(s)	Produce product	Enter & validate order	Pick product
Schedule product deliveries	Inspect/test	Reserve delivery resources	Pack product
Receive product	Package	Determine delivery date	Load product
Verify product	Stage product	Consolidate orders	Generate shipping docs
Transfer product	Release product	Build loads	Ship product
Authorize supplier payment		Route shipments	Verify customer receipt
		Select carrier(s)	Invoice customer

Even in small businesses, these processes may be replicated hundreds or thousands of times every day for different products, customers, production lines and distribution centers. A disruption at any point may affect the entire sequence of processes, and the combined impact may continue to be felt days or weeks after service has been restored.

Companies often estimate the cost of downtime per hour based on sales or production value; e.g., the cost of an hour of downtime is the same as an hour's worth of sales or production. However, this may seriously understate actual damage. If cascading effects are included in the calculation, it may be necessary to multiply this by, say, three or four times to allow for the cumulative impact of disruptions over longer periods of time.

There is another significant effect. Disruptions tend to increase error rates. These may affect all phases of the supply chain, resulting in order processing, product, packing and shipping errors that may generate additional resolution costs, and further alienate customers.

The impact of disruptions may extend beyond operational activities. System outages may delay forecasting and planning cycles; impair the ability to respond to customer queries about delivery times, inventory availability and order status; and affect other informational processes across the entire organization.

In a business operating in real time, damage may be caused not only by failing to deliver a part to a manufacturing line, or a finished product to a customer, but also by delays in delivering the right information to the right person at the right time.

It can be expected that technology advances such as radio frequency identification (RFID) will reinforce these trends by speeding movement of goods and further reducing inventory stocks across all stages of the supply chain where they are employed. They will also tend to make organizations even more dependent on real-time information availability.

3. ***Globalization***. Most of the companies surveyed for this report operated internationally, or employed foreign suppliers, channel partners, or both. This was the case, for example, for 21 out of 24 manufacturers (the exceptions were food and beverage companies dealing in perishable goods), and for all 10 retailers and distributors.

   This reflects a broader trend – the supply chains of most companies are now regional or global. Large segments of manufacturing production, in particular, have moved to China and other offshore bases. As result, certain processes – including those related to procurement, logistics and, in many cases, sales, order management and customer service – now occur on a 24x7 basis.

   Companies surveyed for this report were aware (in some cases, acutely aware) of the implications of this shift. An electrical equipment company, for example, manufactured most of its products in Asia. A system outage could cause the company – literally – to "miss the boat" for a critical customer order. Even if alternative shipping arrangements could be made, it might be obliged to use air rather than ocean freight. This would be a great deal more expensive.

4. ***Internet applications***. The trend across most of the industries in which EnterpriseOne is employed is toward Internet-based customer and supplier self-service systems. These increasingly handle business-critical processes such as inventory availability queries, order placement and customer service response.

   The Internet is, almost by definition, a 24x7 medium, and the normal expectation is that online systems should be accessible at any time. Internet systems often experience some level of activity at all hours of the day and night, and on weekends and public holidays. Any outage at any time may affect customers.

   Figure 5, for example, shows the frequency of online orders placed with a health and personal care products distributor over a 24-hour period.

Figure 5
24-hour Online Order Activity: Distributor Example

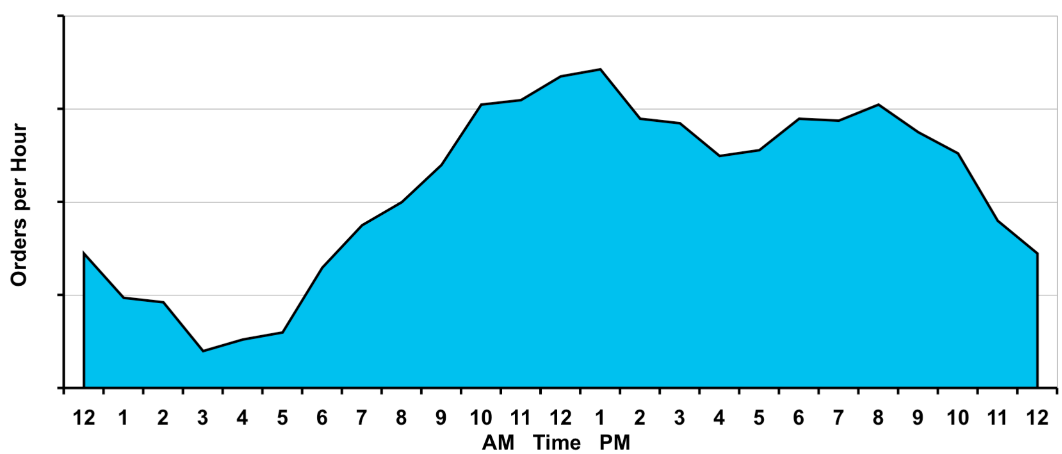

Although many customers ordered on a real-time basis, smaller outlets often did not have time to check inventories and place orders until evening or early morning. Inability to access the distributor's online system at this time, at best, would be inconvenient, and could easily result in lost sales. If the experience was repeated, customers might defect.

This example highlights a broader implication of failing to maintain adequate system availability. For many companies, the largest potential damage resulting from an outage affecting any major system is not internal disruption, but customer loss.

A customer who is impacted directly (e.g., because an online self-service system is down) or indirectly (e.g., because supplier order management, production or delivery operations are disrupted) by an outage will inevitably be dissatisfied, and may defect. Long-term lost business may be substantial.

Even if customers are not lost, there can be a number of negative bottom-line effects. For business-to-business suppliers, these might include late delivery and imperfect order penalties as well as the need to offer special discounts or terms and conditions in order to win back the customer's business.

A less visible, but potentially more significant erosion of confidence might also occur. This could cause the customer to hedge by diverting some future purchases to other suppliers in order to reduce dependence. In addition, the customer might be reluctant to rely upon the company for future strategic orders, particularly where these were time-sensitive.

System outages may thus have a wide range of potential cost impacts. Figure 6 contains a representative list of these for manufacturing companies. Customer-related and operational costs items shown here were employed in calculating costs of downtime for profiles and scenarios presented in this report.

Figure 6

Potential Costs of Outages: Manufacturing Companies

Strategic Costs		
Charge against earnings Financial metrics/ratios Share price decline Share price volatility Cost of capital (discount rate) Increased risk provision Reduced brand value Insurance premiums	Damaged reputation - Financial markets - Customers/prospects - Banks - Business partners - M&A candidates	Legal exposure - Customers - Third parties - Shareholders Compliance exposure - Regulatory reporting - Impaired inspection - Impaired traceability
Customer-related Costs		
Lost Sales & Profit	**Other Costs**	
Lost short-term sales Lost short-term profit Lost future sales Lost future profit	Late delivery penalties Imperfect order penalties Product defect penalties	Customer rebates Buyback pricing/concessions Additional customer service cost
Operational Costs		
Idle Capacity Overall supply chain Procurement Plant operations Logistics/distribution Transportation Warehouses Third-party services **Personnel Costs** Idleness/underutilization Reduced productivity Additional work required Overtime/shift premiums Additional T&E costs	**Finance Processes** Delayed billing/receivables Inventory carrying cost Cash flow cost Delayed close **Costs of Change** Procurement change Revised order processing Special order cost Production schedule change Line change cost Costs of logistics change Supplier premiums Expedited transportation Additional handling cost Additional inventory cost Additional checking cost	**Error-related Costs** Order processing errors Product defect Specification error Manufacturing error Quality failure Shipment error Damaged product Wrong packaging Routing error Wrong delivery time **Other Costs** Lost promotional expenditure Lost marketing expenditure IT costs Administrative costs Overhead

An additional set of "strategic" costs, however, may also be incurred. These will tend to occur if outages are severe, protracted, or both. Share prices may be impacted. Other effects such as reduced brand value, increased risk provision, higher insurance premiums, and a variety of reputational, legal and compliance problems may be experienced.

The potential significance of such effects was highlighted by a recent study co-authored by Kevin Hendricks of the University of Western Ontario and Vinod Singhal of the Georgia Institute of Technology. After reviewing the financial results of more than 800 public companies that had experienced severe supply chain disruptions, the authors concluded that company stocks experienced 33 to 40 percent lower returns relative to industry benchmarks over a three-year period because of these.

The study also reported declines of 7 percent in sales growth, 107 percent in operating income, 114 percent in return on sales, 93 percent in return on assets, and increases in cost of sales, selling, general and administrative (SG&A) expenses and inventory levels.

A clear conclusion emerges. In a modern supply chain, whether outages result in operational disruption, customer-related costs, strategic costs, or combinations of these, they have direct and significant bottom-line impacts. For organizations preparing to deploy a new ERP system, maintenance of the highest possible level of availability should be a core goal from the earliest stages of the planning process.

IT Costs

What does a server cost? Since the appearance of the first total cost of ownership (TCO) studies in the 1980s, it has been clear that initial acquisition costs may be a relatively small component of overall costs over time.

This has notably been the case for "commodity" server platforms. Inexpensive hardware and operating systems for these platforms have often been outweighed by costs for third-party software such as database and system management tools, and by high levels of FTE administrative staffing.

Other factors may also affect costs, including:

- ***Scalability***. If an organization must support high-volume workloads or experiences significant growth in workloads over time, this may be an important cost variable. Inadequate scalability may mean that multiple servers must be employed, or that servers must be periodically replaced with more powerful and expense models as they reach their capacity limits, or both.
- ***Efficiency***. A server capable of operating at, say, 80 percent of capacity may offer superior economics compared to one operating at 30 percent or 40 percent. The level of capacity utilization achieved will be largely determined by the efficiency of the overall system design, and by the effectiveness of system and workload management capabilities.

Such factors may be familiar to many ERP users. High staffing levels reflect inefficient process structures as well as lack of automation. The impact of scalability limitations may be compared to the effects of a production or distribution bottleneck. Equally, the cost structure of a server platform, like that of a supply chain, is materially affected by the efficiency with which capacity is used.

A further lesson may be drawn from ERP experiences. The acquisition cost of a server platform provides as much insight into overall IT costs as the cost of purchasing a new machine tool, fork lift or truck provides into overall production and supply chain costs. In both cases, it is necessary to understand – and address – the entire cost structure.

USER VIEW

User Survey

What is the appeal of the System i for EnterpriseOne deployment? A recent International Technology Group (ITG) survey of companies that ran EnterpriseOne solutions on this platform offers answers.

The survey population included manufacturing (24), distribution (8), business services (7), transportation (3), retail (2) and energy, gaming, and health care (1 each) companies, along with government organizations (3) and utilities (2). These were located in North America (39), Europe (9), the Asia/Pacific region (3) and Latin America (1).

The following benefits were cited.

High Availability

Although different terms were used to describe this capability, the ability of the System i to operate for long periods without outages was cited by most organizations in all industries as a critical benefit.

The System i was variously described as "very reliable…very stable…extremely reliable…completely reliable…highly available…highly dependable" and cited "high reliability…extremely high availability…rock-solid reliability." Others referred to "99.999 percent uptime," "virtually 100 percent uptime," "continuous availability" and equivalents.

The ability of the System i to deliver high levels of reliability and availability, as figure 7 indicates, was the most highly valued strength of this platform.

Figure 7

Benefits of JD Edwards Deployment on IBM System i: User View

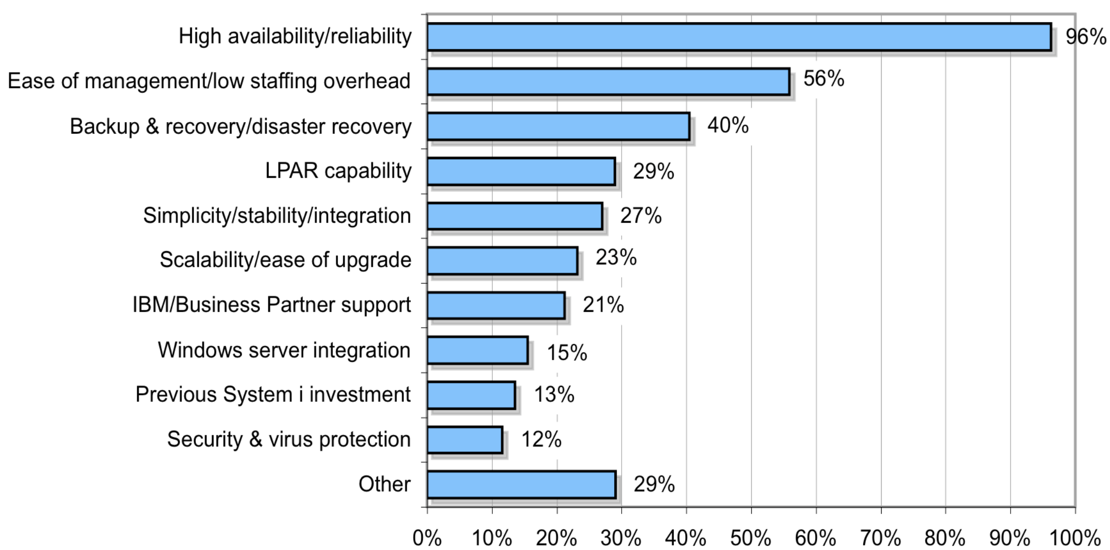

Base: 52 organizations

Source: International Technology Group

Three companies reported that no unplanned downtime had occurred since they had deployed EnterpriseOne on System i platforms to meet Y2K deadlines. Others reported similar experiences.

Low Staffing Levels

The second most commonly reported benefit of employing the System i was reported to be the ease with which this platform could be managed, which was reflected in the comparatively small number of personnel required to handle system administration, database administration and related tasks. This was cited by 29 users (56 percent).

Among organizations that also employed UNIX servers, staffing levels for these tasks were reported to be significantly higher. Two organizations that employed Oracle databases noted that database administrator staffing was generally lower for the DB2 environment.

Comparisons between System i and Windows server staffing levels were more common (most of the organizations surveyed employed the latter in various roles), and differences were reported to be larger than for UNIX servers.

One industrial equipment manufacturer with approximately 1,500 employees reported, for example, that a single System i administrator handled EnterpriseOne and other core business systems, while five server administrators and one DBA handled close to 70 Windows servers supporting other applications. The System i administrator also acted as a DB2 DBA.

Another manufacturing company, with approximately 700 employees, reported that a single administrator handled the company's System i-based EnterpriseOne and Domino e-commerce systems while performing other tasks. Two full-time administrators were required to handle the company's Windows-based e-mail, end-user and departmental applications. Comparable ratios were cited by others.

Lower levels of System i FTE administrative staffing was cited as a benefit valued by organizations of all sizes, in most industries. In three large companies, high levels of System i automation were said to enable "lights out" or unattended operations for core EnterpriseOne systems.

Other Benefits

Other benefits were reported to include:

- ***Backup and recovery/disaster recovery***. System i capabilities in this area (the two terms were often used interchangeably) were cited by 21 organizations. These included manufacturing, distribution, retail, transportation, health care, business services and other companies, along with government organizations and utilities.

  There was widespread agreement that disaster recovery (DR) coverage was becoming increasingly important. Two users, a freight shipper and an automotive parts manufacturer, reported that they had been obliged to put effective "hot site" failover arrangements in place to meet customer demands.

  Others – including manufacturers, distributors and transportation companies – reported that such arrangements were not only essential to avoid the bottom-line effects of protracted outages, but had also become a de facto competitive requirement in their industries. Customers were reluctant to entrust their business to companies that did not have adequate DR coverage.

  In addition, a health care company cited U.S. federal government regulations requiring protection of sensitive records as a key driver of its DR strategy, while government organizations and utilities cited similar mandates.

Users cited i5/OS Remote Journaling and Save While Active, along with third-party clustering solutions for the System i platform as enabling highly effective DR capability. These are discussed in more detail in the Technology Differentiators section.

- ***Logical partitioning***. The ability to deploy multiple instances of EnterpriseOne and other systems on System i platforms using logical partitions (LPARs) was cited by 15 organizations.

  The most common applications were deployment of multiple EnterpriseOne production, development and test instances on a single System i platform; hosting of i5/OS-based Domino e-mail and e-commerce systems; and use of LPARs to deploy Linux and, in one case, AIX computer aided engineering (CAE) applications on System i platforms.

  One key benefit of LPAR capability was reported to be that it enabled consolidation of platforms. For example, one company had consolidated four regional EnterpriseOne production instances onto a single System i. Another had adopted a similar approach for three divisional systems.

  A third had consolidated three production instances of EnterpriseOne as well as other i5/OS applications onto a single System i platform. A second System i hosted development and test instances, and acted as a failover for the primary system. This configuration is shown in figure 8.

Figure 8

IBM System i EnterpriseOne Configuration Using LPARs: Example

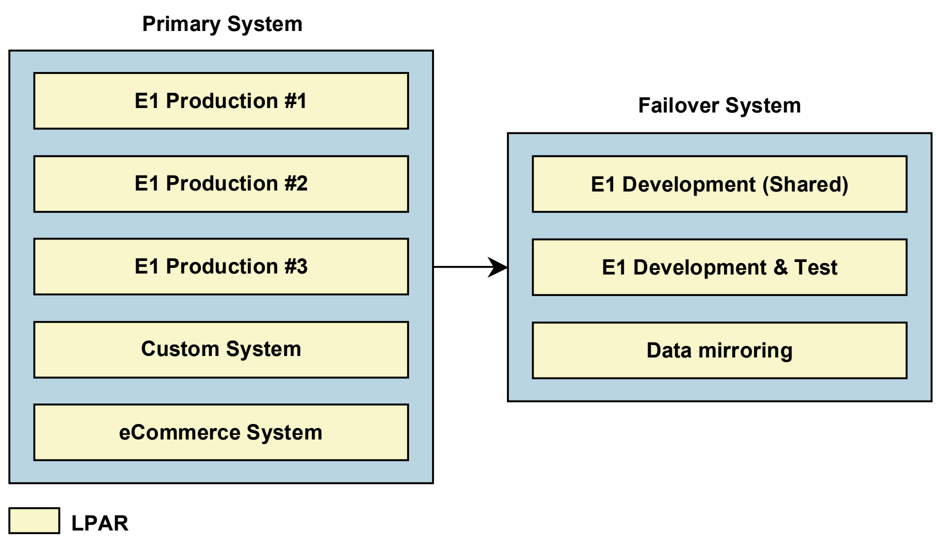

Another benefit was said to be that LPARs enabled overall capacity to be used more efficiently. Several companies reported, for example, that they were able to re-allocate capacity dynamically between LPARs to handle workload peaks that varied between countries and time zones, or were better able to handle short-term volume spikes, or both.

LPAR capability, it was also noted, could improve availability by allowing partitions to be used for data staging, backups, version upgrades, application development and testing, and other functions without taking systems offline.

System i DR solutions fully supported LPAR-based environments, with the result that consolidation of instances did not necessarily increase vulnerability to disruptions. System i configurations with multiple LPAR-based systems were routinely backed up failover servers, with data mirroring processes occurring in some cases in near real time.

- ***Simplicity, stability and integration***. These characteristics of the System i platform and i5/OS environment were cited by 14 organizations. The general view was that they created less work for system and database administrators as well as for application, operations and other IT personnel.

  In these areas, it was reported, the System i contrasted with UNIX and Windows server environments. In the case of the former, it would be necessary, at a minimum, to deal with different systems software and database vendors. In the case of the latter, it was often necessary to deal with an even more diverse range of components.

  In either case, users were obliged to configure, install and support products from multiple vendors, and to handle comparatively frequent, often-uncoordinated vendor upgrades and modifications to these, including issue of patches. A great deal of IT staff time might thus be spent on comparatively minor software tasks.

- ***Scalability***. The comparative ease with which System i capacity could be upgraded to handle short-term workload spikes, long-term business growth, or both was cited by 12 organizations.

  One distributor who had expanded rapidly through organic growth and a series of regional acquisitions reported, for example, that its system capacity requirements had more than quadrupled since 2000. These requirements had been met by progressively upgrading the original System i platform. Similar experiences were cited by other companies.

  The speed and ease with which System i capacity could be upgraded was also cited as a benefit by organizations that experienced workload fluctuations due to seasonal variations, unpredictable demand patterns (e.g., customer "rush orders") and use of e-commerce systems. In all of these cases, the scalability strengths of the System i platform were seen as an important benefit.

  IBM Capacity Upgrade on Demand (CUoD) offerings, which enabled companies to activate additional capacity when needed and pay for it only while it was in use, received favorable comment. One company noted that it could activate CUoD capacity "within seconds" in the event of demand surges. This avoided the need to "pay overhead for underutilized systems."

- ***IBM and Business Partner support***. The quality and responsiveness of support available from IBM and IBM Business Partners were cited by 11 organizations as a benefit of EnterpriseOne deployment on System i platforms.

  Key advantages were said to include high levels of industry-specific EnterpriseOne expertise (one respondent noted that "IBM really understood our business…we were able to get up the learning curve much faster because of that") as well as the ability of the IBM Global Services organization to provide specialists in a wide range of business and technical disciplines.

  Others cited high levels of cooperation between IBM, IBM business partners and the Oracle JD Edwards organization as an important benefit.

- ***Windows server integration***. The ability to use older IBM technologies such as the Integrated xSeries Adapter (IXA) and Integrated xSeries Server (IXS) as well as new capabilities for iSCSI attachment of IBM BladeCenter servers was cited by eight organizations.

  In most cases, these solutions were employed to host Windows-based applications. These included Windows file and print serving; Microsoft Outlook and Exchange; e-commerce; Internet and intranet infrastructure, including Web servers and firewalls; and decision support, departmental and other non-core business applications.

Use of System i Windows server integration solutions was said to result in lower administrative costs, improved security and better storage management capability. The ability to use System i storage management services to improve disk capacity utilization, facilitate backup to tape, and extend System i DR coverage to Windows-based applications was notably valued.

These capabilities particularly appealed to smaller organizations that required few Windows servers (in these installations, System i administrators could often handle Windows as well as i5/OS tasks) and to organizations of all sizes where Windows-based applications were an essential complement to i5/OS-based EnterpriseOne systems.

In the latter case, System i Windows server integration solutions were particularly used to support e-commerce and collaborative extensions to the EnterpriseOne environment, and decision support and planning tools that exploited data generated by EnterpriseOne. Both sets of applications benefited from higher availability levels as well as reduced latencies for interaction with i5/OS-based core EnterpriseOne systems.

- ***Previous System i investments***. The ability to leverage previous investments in System i infrastructures and skill bases was cited by seven organizations as a benefit of deploying EnterpriseOne solutions in this platform.

  Part of the appeal of EnterpriseOne solutions to this group was that they were able to implement latest-generation ERP capabilities without radical changes in key skill sets. Moreover, they could continue to employ embedded System i-based infrastructures for data center operations, backup and recovery, and other functions.

  Three organizations reported that they continued to employ some legacy RPG-based applications alongside EnterpriseOne. The ability to host these on the same platform was reported to facilitate interoperability, and to enable better utilization of staff time and skills than would be the case if multiple server platforms, operating systems and databases were employed.

- ***Security and malware resistance***. System i strengths in these areas were cited by six organizations. These included a U.S. health care company that valued the role the System i played in enabling compliance with the privacy and records retention provisions of the Health Insurance Portability and Accountability Act (HIPAA), a government organization, and manufacturing and retail companies.

  Among manufacturers and retailers, key concerns included security of global supply chains (one company noted, for example, that dealing with approximately 250 suppliers in more than 20 countries inevitably increased risks of online security violations and malware damage) and potential vulnerabilities caused by large-scale use of e-commerce systems.

Other benefits of EnterpriseOne deployment on System i platforms were reported to include high levels of performance, as well as System i strengths in database functionality, and in workload management and storage management.

TECHNOLOGY DIFFERENTIATORS

System Architecture

From a system architecture perspective, the System i platform incorporates multiple elements. These include a distinctive basic system design; a highly integrated operating environment incorporating database management and other components; reliability, availability and serviceability (RAS) features drawn from mainframe systems and other IBM sources; advanced virtualization and system and workload management facilities; and industry-leading autonomic (artificial intelligence) features.

The System i also benefits from the performance strengths of the IBM POWER5+ reduced instruction set computing (RISC) technology. POWER5+ is a recognized industry performance leader that employs dual- and quad-core processors, chip symmetric multithreading and other advanced capabilities.

While the System i platform benefits from the performance levels that are characteristic of latest-generation RISC processor technology, its architecture has evolved in a manner that is significantly different from that of UNIX as well as Windows servers.

The advantages of System i described in this report are not merely a function of effective hardware design and software engineering. The unique characteristics of the System i compared to competitive platforms reflect fundamental technical differences that are highly relevant to the requirements for effective use of EnterpriseOne systems.

This is particularly the case in the following areas:

- ***System design***. The core System i design, which is derived from the earlier IBM System/38 and AS/400, is built around an object-based kernel in which all system resources are defined and managed as objects.

  The kernel also incorporates single-level storage capability – meaning that the system treats all storage resources, including main memory and disks, as a single logical entity. These may be exploited transparently by applications, regardless of physical type and location. Management functions for these, as figure 9 shows, are built into the core system design.

Figure 9
IBM System i Single-level Storage Structure

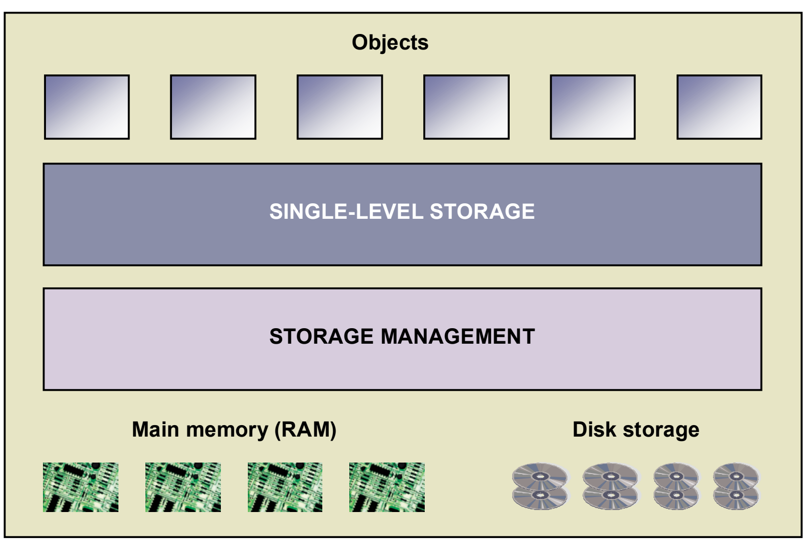

In addition, key portions of the i5/OS operating environment, including system storage management functions, are implemented in microcode.

These features enable high levels of configuration flexibility, and materially improve the efficiency with which processor and storage resources are used, with corresponding benefits in performance and capacity utilization.

- ***System integration***. The i5/OS environment includes not only core operating system functions, but also the DB2 database management system, WebSphere Application Server, and tools for system, storage, backup and recovery, communications, security, operations and other management tasks. These are closely integrated, and share common administrator interfaces.

  Equivalent functionality in UNIX and Windows server environments typically requires that users acquire, install, configure and administer multiple software products, often sourced from different vendors. This increases deployment complexity, and tends to create integration and administration challenges that are greater than those faced by System i users.

- ***Partitioning and virtualization***. System i LPAR capability enables organizations to host multiple operating system instances and applications on a single physical platform, and to better manage and balance diverse workloads across the system. Up to 254 LPARs are supported using the System i firmware-based Hypervisor.

  Resources may be allocated and re-allocated dynamically between partitions, enabling organizations to achieve significantly higher levels of capacity utilization than may be realized with monolithic server environments.

  LPARs are commonly used to host development, test and production instances on the same system, and to perform software updates and other tasks without taking systems offline. LPAR capability also offers potential to consolidate multiple production system instances. Organizations surveyed for this report had adopted both approaches.

  In addition to multiple copies of i5/OS, AIX 5L Version 5 (the current IBM version of UNIX) and Linux operating systems may also run in LPARs.

  System efficiency is augmented by three further virtualization capabilities. ***Virtual I/O*** and ***Virtual Disk*** enable multiple LPARs to share adapters and physical disk drives respectively – i.e., fewer physical adapters and less disk capacity are required.

  Virtual LAN supports high-speed linkages between LPARs and enables creation of virtual Ethernet switches and adapters, which may result in additional hardware economies and reduced configuration complexity. A further benefit is that Virtual LANs are more secure than external LAN connections. These capabilities are illustrated in figure 10.

- ***System and workload management***. System i partitioning and virtualization capabilities are reinforced by sophisticated i5/OS system and workload management facilities.

  These are designed into the core System i architecture and embedded into i5/OS rather than, as is the case for Windows, HP-UX and Solaris server environments, implemented through software add-ons and overlays. Key facilities such as the Enterprise Workload Manager (EWLM) and IBM Director provide uniform management support for all system-level and LPAR resources.

  A key point should be emphasized here. With any platform, it is not sufficient simply to create partitions. The overall efficiency with which a partitioned system operates also depends on how effectively workloads are distributed and how resources are managed across these. System i capabilities in these areas are among the most advanced in today's IT world.

Figure 10
IBM System i Partitioning and Virtualization Capabilities

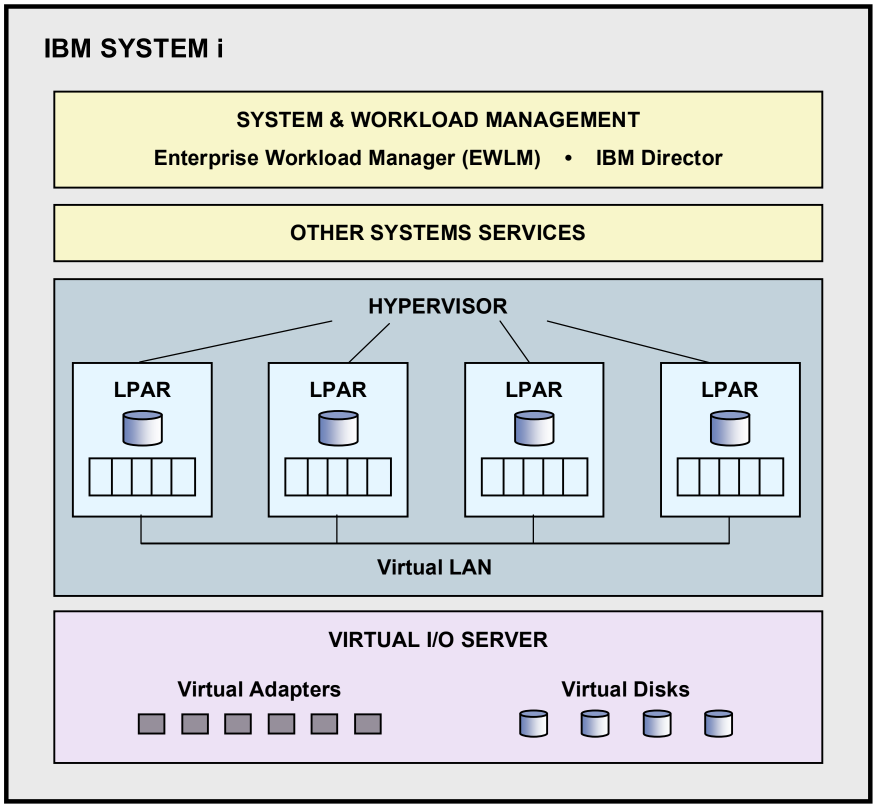

A broader characteristic of the System i environment is that its different components are implemented in a highly synergistic manner. DB2, for example, exploits the underlying System i object-based structure and single level storage capabilities, and these also assist in enabling LPARs to utilize system resources with high levels of efficiency. System-level and Virtual I/O Server capabilities as well as software- and microcode-based management facilities are closely integrated and optimized.

A key benefit of the System i's distinctive capabilities is that it enables organizations to run the application, Web and database serving components of the EnterpriseOne environment on a single physical platform, while realizing higher capacity utilization than less sophisticated platforms.

This approach also delivers advantages in overall performance (latencies for cross-server interaction are reduced), availability (Web and application servers benefit from System i high availability strengths) and reduced system administration overhead.

High Availability

Maintenance of high levels of system availability over long periods can be a challenging process. Organizations must deal with risks of unplanned outages caused by hardware and software failures, administrator and operator errors, workload spikes that overload systems, hacker attacks, virus damage and other factors. Vulnerability to these may vary widely between platforms.

Unplanned outages can be highly disruptive, particularly if they occur during periods of peak business activity. The majority of downtime, however, will be due to planned outages for such functions as hardware, operating system, database and applications software upgrades; database reorganization; and scheduled hardware and software maintenance.

Like unplanned outages, planned downtime tends to be platform-dependent. In a Windows server environment, for example, frequent planned outages will typically be required for hardware and software upgrades, including patching, preventative maintenance and other functions. Figure 11 illustrates these.

Figure 11
Representative Causes of Planned Outages: Windows Servers

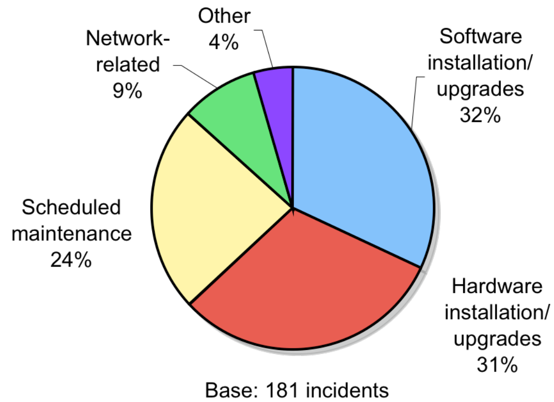

Source: International Technology Group

Avoidance of both types of outage is a central design parameter of System i architecture. Distinctive capabilities are implemented at multiple levels. At the most basic level, these include industry-leading embedded RAS features. High levels of redundancy, along with monitoring, diagnostic, and fault isolation and resolution facilities are built into all major components and subsystems.

Operating system-level capabilities include Remote Journaling (file and system changes may be automatically copied to a second server); Save While Active (backups may be performed without taking systems offline); Independent Auxiliary Storage Pools (data may be mirrored to remote site systems); and Clustering Management Services, which are embedded into i5/OS and leveraged by third-party solutions.

LPARs further contribute to reduction of planned outages. Software modifications may be made and new versions installed and assured without disrupting operations. Backups may also be performed, and batch workloads executed concurrently with mainstream processes.

Concurrent maintenance ("hot plugging"), dynamic firmware upgrades, LPARs and other capabilities also reduce requirements for planned outages.

Functionally similar capabilities may be implemented for UNIX and Windows servers. However, a key principle should be noted. Achieving high levels of availability and recoverability inevitably increases configuration complexity. If underlying environments are already complex, the overall solution will tend to be less stable, or will require greater expenditure, higher staffing levels, or both.

i5/OS integration also facilitates availability. Risks of software failure are typically less than for complex multivendor environments. Simplified, high-productivity management interfaces mean that the potential for administrator and operator errors causing outages, data loss, or both is reduced.

Autonomic Functions

The System i platform benefits from the one of the most advanced implementations of autonomic technologies within the IBM product line. Autonomic capability – meaning the application of artificial intelligence technologies to IT administration and optimization tasks – has been a major IBM development focus since the late 1990s, and the company is the clear leader in this area.

There are four categories of System i autonomic functions: self-configuring; self-optimizing; self-protecting; and self-optimizing. These functions, summarized in figure 12, extend across i5/OS and DB2 and are embedded into hardware components and subsystems.

Figure 12
IBM System i Autonomic Functions

SYSTEM	
Self-configuring	**Self-protecting**
Connect automated services CPU capacity upgrade on demand Enterprise Identity Mapping EZSetup Wizards Hot plug disk & I/O Linux & Windows Virtual I/O RAID subsystem Switchable auxiliary storage pools Windows NetServer file/print support Windows dynamic storage addition Wireless system management access	Automatic virus removal Chipkill Memory Digital certificates Digital object tagging Enterprise Identity Mapping Integrated Kerberos support Integrated SSL support IP takeover RAID subsystem Self-protecting kernel Tagged storage
Self-optimizing	**Self-healing**
Adaptive e-transaction services Automatic performance management Automatic workload balancing Dynamic disk load balancing Dynamic LPAR for i5/OS & Linux Expert Cache Global resource manager Heterogeneous workload manager Quality of service optimization Single-level storage	ABLE problem management engine Auto-fix defective PTFs Automatic performance adjuster Chipkill Memory, dynamic bit steering Concurrent maintenance Domino auto restart, clustering Dynamic IP takeover, clustering Electronic Service Agent (“call home”) First-failure data capture & alerts Service director
DATABASE	
Self-configuring	**Self-protecting**
Automatic collection of object relationships Automatic data striping & disk balancing Automatic distributed access configuration Automatic object placement Automatic self-balancing indexes Automatic TCP/IP startup Graphical database monitor	Automatic enforcement of user query limits Automatic enforcement of user storage limits Digital object signing Object auditing OS-controlled resource management
Self-optimizing	**Self-healing**
Automatic Index Advisor Automatic memory pool tuning Automatic query plan adjustment Automatic rebind & reoptimization Automatic statistics Caching of open data paths & statements Cost-based Optimizer	Automatic access path protection Automatic backup/restore of database objects Automatic database object extents Automatic database restart Automatic system-managed journaling First-failure data capture & alerts

The most obvious benefit of System i autonomic capabilities is that they improve the productivity of system and database administrators. Automation of administrative tasks also, however, reduces the potential for human errors leading to performance bottlenecks, outages, data loss or corruption and other negative effects.

Security and Malware Resistance

There are also significant differences between UNIX and Windows servers, and i5/OS environments in security capability as well as in resistance to viruses and other forms of malicious code ("malware") such as worms, Trojans and zombies.

i5/OS users are generally less vulnerable to security violations. The U.S. National Vulnerability Database, operated by the National Institute of Standards and Technology (NIST), for example, has recorded no vulnerabilities for i5/OS or its predecessor, OS/400, since collection of statistics began in 1992. Although there are variations in definition of vulnerabilities between different databases, other security sources confirm the i5/OS position in this area.

There are, similarly, no known native viruses for i5/OS, and malware incidents involving System i platforms are virtually unknown.

To some extent, this is because i5/OS is less well-known and less attractive to virus initiators than Windows or UNIX environments. Technical characteristics also, however, result in a level of vulnerability that is lower than for UNIX and (to an event greater extent) Windows servers.

Security functions are designed into the core system architecture, and closely embedded into the operating environment. They are tightly integrated with the system's compiler structure, directory server and object-based file system. Programs that do not meet highly demanding validity criteria cannot execute.

A further benefit is that high levels of system integration and administrator productivity make it easier to implement and maintain security policies than in a more complex IT environment. This is particularly valuable for companies with few or no security professionals and limited security resources.

These strengths do not mean that System i users do not face security or malware threats. But the probability that these will materialize is significantly less than for other platforms. The staff time and funds that must be expended to guard against them is correspondingly less.

DETAILED DATA

Comparison Profiles

Methodology

The IT cost and cost of downtime comparisons presented in this report were calculated using data supplied by 21 companies in the same industries and approximate size ranges, and employed most or all of the major components of the EnterpriseOne suite.

Data on two subjects were obtained from these companies:

1. ***EnterpriseOne installations***. Data on application suites and modules, numbers of instances and users, workloads, server configurations, staffing and other variables was collected for use in IT cost calculations. Data on overall availability levels and on the frequency and duration of outages for different platforms was also collected as input for cost of downtime calculations.
2. ***Business operations***. Data on manufacturing and supply chain operations and costs, customer and supplier bases, key cycle times, inventory levels and potential fees for delivery shortfalls, along with details of order management, procurement, production, distribution, billing and other processes potentially affected by outages was collected for use in cost of downtime calculations.

Using these inputs, three composite profiles were constructed. A best practices approach drawing upon the experiences of multiple organizations was employed. For example, the experiences of one automotive parts manufacturer with core EnterpriseOne ERP applications were combined with the experiences of a second with CRM solutions, and a third with online procurement. The food and beverage, and electronics manufacturer profiles were constructed in a similar manner.

Profile Summaries

The three profile companies may be summarized as follows:

1. ***Automotive parts manufacturer*** is a $100 million company producing automotive suspension and transmission components for cars, SUVs and light trucks. It employs more than 800 people at manufacturing and distribution centers in the United States and Mexico, and at branch offices worldwide. Its customers are primarily automakers. Its business is mainly build-to-order, with some build-to-stock volume.

   Faced with weak automotive demand, fierce competition from local and (increasingly) offshore suppliers, and highly demanding customers, the company has moved to a demand-driven manufacturing model. To support this shift, a broad suite of EnterpriseOne applications has been deployed. There are approximately 150 internal users of these. Customer and supplier self-service systems are also heavily used.

2. ***Food and beverage manufacturer*** is a $500 million company offering a range of branded consumer products. The company has approximately 1,250 employees at five manufacturing and distribution centers, and eight sales offices. It employs a mix of direct sales to major retail and institutional accounts, and third-party distributors serving smaller customers.

   The company's EnterpriseOne ERP suite supports the entire business, including 24x7 manufacturing and logistics operations. Its CRM system is employed for interactions with clients through direct sales representatives as well as call center and online channels. Collaborative forecasting and planning arrangements have been put in place with leading retail customers.

There are approximately 250 ERP and CRM users. Most employees and managers access the company's HR self-service system, and traveling salespeople, managers and professionals handle expense processing through its intranet.

The company deals with worldwide customer and supplier bases. Orders are received and processed, and most production and distribution facilities operate on a 24x7 basis.

3. ***Electronics manufacturer*** is a $650 million company specializing in computer peripherals and subassemblies. The company sells to computer vendors as well as through distributors. It employs more than 4,000 people at four manufacturing plants in Asia, assembly and distribution centers in North America and Europe, and branch offices worldwide.

   EnterpriseOne plays a central role in the company's strategy to deal with market volatility, shortened product life cycles, price erosion and increasing customer pressure for reduced order-to-delivery times. The company operates in near real time on a 24x7 basis. Forecasts and plans are updated daily.

   The company employs three production EnterpriseOne ERP systems, supporting North American, European and North American operations, and a single global CRM instance. These are hosted at a single data center, which is equipped for hot site failover and recovery in the event of a disaster. There are approximately 400 EnterpriseOne users worldwide.

Additional details of business profiles, application modules employed, numbers of users, server hardware and software configurations and FTE staffing for system and database administration for EnterpriseOne installations are shown in Figure 13.

In this presentation, numbers of processors and cores are shown for HP Integrity, Sun Solaris and Windows servers – e.g., "X4600 8/16 x 2.6 GHz" refers to a Sun Solaris X4600 server with 8 Advanced Micro Devices (AMD) Opteron 2.6 GHz dual-core processors for a total of 16 cores, while "4/8 x Xeon 3.4 GHz" refers to a server with 4 Intel Xeon 3.4 GHz dual-core processors for a total of 8 cores.

Numbers of cores only are shown for System i platforms. All three configurations employ dual-core 1.9 GHz POWER5+ processors. The 3-core configuration for the Food and Beverage Manufacturer profile is a 4-way (2x2) model with one inactive core.

In System i scenarios, EnterpriseOne Web serving is implemented on 525 and 550 models. In other platform scenarios, this function is handled by dedicated servers equipped with two dual-core Intel Xeon 2.0 GHz processors and 2 GB of RAM.

Configurations for all platforms are based on user-reported data. In some cases, users supplied configuration data for older iSeries, HP Integrity, Sun Solaris, and Windows server models. Where this was the case, configurations were updated to current-technology models using vendor comparative performance guidelines and sizing estimates, and other inputs as appropriate.

All configurations are for EnterpriseOne 8.12 suites.

Figure 13
Profiles Summary

Automotive Parts Manufacturer	Food & Beverage Manufacturer	Electronics Manufacturer
Business Profile		
$100 million manufacturer of transmission & suspension products 800 employees 2 manufacturing & distribution centers + 12 branch offices	$500 million manufacturer of food & beverage products 1,250 employees 5 manufacturing & distribution centers + 8 sales offices.	$650 million manufacturer of computer peripherals & subassemblies 4,000 employees 4 manufacturing plants & distribution centers + 25 branch offices
EnterpriseOne Applications		
Customer Order Management/ Customer Relationship Management Customer Self Service, Mobile Sales Sales Force Automation Sales Order Management ***Financial Management*** Accounts Payable, Accounts Receivable Fixed Asset Accounting, General Ledger ***Human Capital Management*** Human Resources Management, Payroll ***Logistics*** Advanced Stock Valuation Inventory Management Warehouse Management ***Manufacturing*** Engineer to Order Product Data Management ***Supply Management*** Buyer Workspace Procurement & Subcontract Management Supplier Self Service	***Customer Order Management/ Customer Relationship Management*** Advanced Pricing Customer Self Service, Mobile Sales Sales Force Automation Sales Order Management ***Financial Management*** Accounts Payable, Accounts Receivable Expense Management, General Ledger ***Human Capital Management*** Employee Self Service Human Resources Management Manager Self Service Payroll, Time & Labor ***Logistics*** Inventory Management Warehouse Management ***Manufacturing*** Quality Management Requirements Planning, Shop Floor ***Supply Management*** Buyer Workspace Procurement Management Supplier Self Service	***Customer Order Management/ Customer Relationship Management*** Advanced Pricing, Configurator Customer Self Service Sales Force Automation Sales Order Management ***Financial Management*** Accounts Payable, Accounts Receivable Expense Management, General Ledger ***Human Capital Management*** Human Resources Management Payroll, Time & Labor ***Logistics*** Inventory Management ***Manufacturing*** Quality Management Requirements Planning, Shop Floor ***Supply Management*** Buyer Workspace Procurement Management Requisition Self Service Supplier Self Service
Number of Users		
150	250	400
SCENARIOS		
IBM System i		
525 2 x 1.9 GHz 12 GB RAM i5/OS V5R4 0.5 FTEs	550 3 x 1.9 GHz 20 GB RAM i5/OS V5R4 0.6 FTEs	550 4 x 1.9 GHz 28 GB RAM i5/OS V5R4 0.8 FTEs
HP Integrity Servers		
rx2660 2/4 x 1.6 GHz 8 GB RAM HP-UX 11i V3, Oracle 10g 0.85 FTEs	rx6600 4/8 x 1.6 GHz 16 GB RAM HP-UX 11i V3, Oracle 10g 1.0 FTEs	rx7640 6/12 x 1.6 GHz 24 GB RAM HP-UX 11i V3, Oracle 10g 1.4 FTEs
Sun Solaris Servers		
V40Z 4/8 x 2.8 GHz 8 GB RAM Solaris 10, Oracle 10g 0.9 FTEs	X4600 8/16 x 2.8 GHz 16 GB RAM Solaris 10, Oracle 10g 1.05 FTEs	V890 8/16 x 1.8 GHz 24 GB RAM Solaris 10, Oracle 10g 1.5 FTEs
Windows Servers		
2/4 x Xeon 3.0 GHz 8 GB RAM Windows Server 2003, SQL Server 2005 1.1 FTEs	4/8 x Xeon 3.4 GHz 16 GB RAM each Windows Server 2003, SQL Server 2005 1.3 FTEs	2 x (4/8 x Xeon 3.4 GHz) 16 GB RAM each Windows Server 2003, SQL Server 2005 2.3 FTEs

IT Cost Calculations

IT costs were calculated as follows:

- ***Server costs***. Costs for System i configurations include acquisition of hardware and i5/OS licenses, along with three-year 24x7 hardware and software maintenance coverage for these. IBM Application Server processor entitlement licensing arrangements are employed for two cores each in the System i scenarios for the food and beverage, and electronics manufacturer profiles.

  Costs for HP Integrity and Sun Solaris configurations include server hardware, operating system licenses, three-year 24x7 hardware and software maintenance coverage, and licenses and three-year update and support coverage for Oracle 10g, including Diagnostic and Tuning Packs.

  Costs for Windows server configurations include acquisition and three-year 24x7 maintenance coverage for hardware, along with licenses and three-year Microsoft Software Assurance coverage for Windows Server 2003 and SQL Server 2005 software.

  Costs for HP Integrity, Sun Solaris and Windows server scenarios include hardware acquisition, systems software and maintenance for dedicated Web servers.

  Costs for all items were calculated using discounted published vendor U.S. list prices as of April 2007. All server configurations are rack-mounted.

- ***Personnel costs***. These were calculated for the FTE values shown in figure 13. Calculations were based on annual salaries of $83,909 and $70,399 per year for i5/OS and Windows server system administrators respectively, and $75,093 for UNIX system administrators supporting HP Integrity as well as Sun Solaris servers; and annual salaries of $82,806 and $90,714 for SQL Server and Oracle DBAs for Windows server scenarios, and HP Integrity and Sun Solaris scenarios respectively. DBA tasks in System i scenarios are handled by i5/OS administrators.

  Annual salaries for all personnel in all scenarios were increased by 28.2 percent to allow for bonuses, benefits and other non-cash compensation, training and other items.

- ***Facilities costs***. These were calculated using vendor specifications for electricity consumption and footprints for server configurations included in scenarios. Costs include data center occupancy – calculations were based on standard rack mount units and service clearances for these, plus allowance for inactive areas – and electricity consumption by servers.

  Facilities costs also include those for power and cooling equipment acquisition, maintenance, occupancy and electricity consumption. Configurations for this equipment are appropriate for electricity consumption and heat generation by servers in scenarios.

  Cost calculations for power and cooling equipment were based on U.S. specifications and discounted list prices for equipment from leading vendors.

  Data center occupancy costs for servers as well as for power and cooling equipment were calculated using a conservative assumption for annual average cost per square foot for existing facilities (i.e., costs do not include new facilities construction). Electricity costs were calculated based on specific utilization levels and hours of operation for each profile, using a conservative assumption for average price per kilowatt/hour. All assumptions are for U.S. costs.

Calculations do not include costs for applications software, external storage, peripherals, networks, PCs and other client devices, personnel other than for system and database administration, and other resources not identified above.

Costs of Downtime

Costs of downtime were calculated using a two-phase process.

First, average costs per hour of downtime were calculated for each composite profile using appropriate industry- and company-specific values. "Average" in this context means that costs are based on overall annual volumes of business activity divided by hours of operation (in all three cases, 24x365 = 8,760).

Results are summarized in figure 14.

Figure 14
Profile Companies: Average Costs per Hour of Downtime

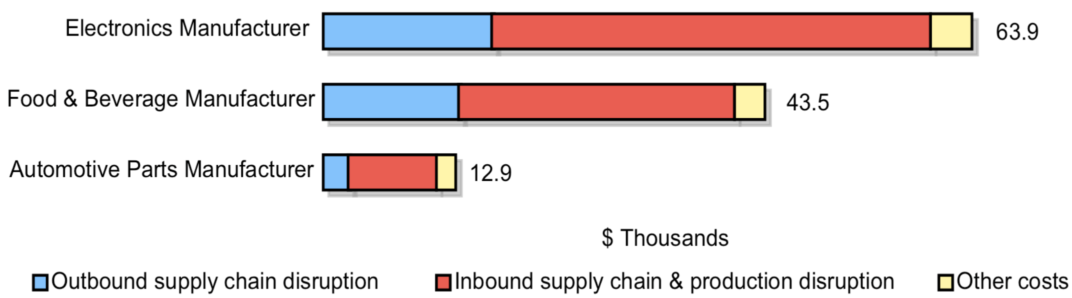

Three main cost categories were employed:

1. ***Outbound supply chain disruption*** consists of costs caused by disruption of activities between factory release and customer delivery. This category includes costs of idle and underutilized capacity, including personnel costs; handling of delivery delays (including distribution center and transportation costs); inventory carrying costs; and costs of change for affected processes.
2. ***Inbound supply chain and production disruption*** consists of costs incurred for activities between initial supplier queries and factory release. This category includes costs of idle and under-utilized capacity, including personnel costs, for inbound logistics and production operations; handling of delivery delays (including transportation costs); inventory carrying costs; and costs of supplier order, production scheduling and other changes.
3. ***Other costs*** include customer penalties for late delivery and imperfect orders, along with buyback costs such as rebates; costs of customer billing delays; reduced productivity of and additional work performed by sales, customer service and administrative personnel; and other items.

Second, average costs per hour of downtime were multiplied by numbers of hours of downtime per year for each server scenario, which were calculated based on user inputs. The focus was placed on downtime for underlying hardware and software platforms, rather than downtime on for EnterpriseOne version upgrades, application maintenance and other functions that would be similar for all scenarios. Annual costs of downtime for each scenario were then multiplied by three for three-year totals.

Results are specific to the profiles and calculation methods described above. Actual costs may vary in individual cases. Companies should develop their own calculations using organization-specific data.

Additional Tables

Detailed breakdowns of three-year IT costs, and of availability levels and three-year costs of downtime for profiles and scenarios are presented in figures 15 and 16 respectively.

Figure 15
Three-year IT Costs Breakdown

Company	Automotive Parts	Food & Beverage	Electronics
SYSTEM i			
System	83,124	140,498	191,178
Maintenance	4,498	25,398	28,458
Personnel	161,357	193,628	258,171
Facilities	4,319	5,208	6,155
TOTAL ($)	**253,298**	**364,732**	**483,962**
HP INTEGRITY SERVER			
Hardware	13,806	39,067	151,684
Maintenance	2,595	4,267	27,930
Systems software	5,986	20,079	46,641
Database software	23,217	97,608	320,712
Personnel	257,502	303,827	428,362
Facilities	6,367	10,836	17,311
TOTAL ($)	**309,473**	**475,684**	**992,640**
SUN SOLARIS SERVER			
Hardware	17,452	40,868	198,157
Maintenance	3,929	4,818	19,303
Systems software	3,183	3,183	3,183
Database software	73,206	267,260	170,016
Personnel	271,943	318,268	457,243
Facilities	5,717	8,453	22,144
TOTAL ($)	**375,430**	**642,850**	**870,046**
MICROSOFT WINDOWS SERVER			
Hardware	8,230	12,753	22,646
Maintenance	2,859	2,840	4,348
Systems software	10,984	15,056	23,084
Database software	13,648	27,295	54,591
Personnel	304,988	361,524	641,822
Facilities	4,984	10,768	19,840
TOTAL ($)	**345,693**	**430,236**	**766,331**

Figure 16
Availability Levels and Three-year Costs of Downtime Breakdown

SCENARIOS	Automotive Parts	Food & Beverage	Electronics
IBM SYSTEM i			
Availability Level	99.94%	99.91%	99.93%
Three-year costs ($000)	**193.3**	**978.8**	**1,149.5**
HP INTEGRITY SERVER			
Availability level	99.86%	99.90%	99.89%
Three-year costs ($000)	**348.0**	**1,631.3**	**1,915.8**
SUN SOLARIS SERVER			
Availability level	99.88%	99.83%	99.87%
Three-year costs ($000)	**406.0**	**1,892.3**	**2,203.2**
MICROSOFT WINDOWS SERVER			
Availability level	99.63%	99.54%	99.62%
Three-year costs ($000)	**1,237.2**	**5,220.0**	**6,322.1**

ABOUT THE INTERNATIONAL TECHNOLOGY GROUP

ITG sharpens your awareness of what's happening and your competitive edge
. . . this could affect your future growth and profit prospects

The International Technology Group (ITG), established in 1983, is an independent research and management consulting firm specializing in information technology (IT) investment strategy, cost/ benefit metrics, infrastructure studies, deployment tactics, business alignment and financial analysis.

ITG was an early innovator and pioneer in developing total cost of ownership (TCO) and return on investment (ROI) processes and methodologies. In 2004, the firm received a Decade of Education Award from the Information Technology Financial Management Association (ITFMA), the leading professional association dedicated to education and advancement of financial management practices in end-user IT organizations.

The firm has undertaken more than 100 major consulting projects, released approximately 160 management reports and white papers, and delivered nearly 1,800 briefs and presentations to individual clients, user groups, industry conferences and seminars throughout the world.

Client services are designed to provide factual data and reliable documentation to assist in the decision-making process. Information provided establishes the basis for developing tactical and strategic plans. Important developments are analyzed and practical guidance is offered on the most effective ways to respond to changes that may impact or shape complex IT deployment agendas.

A broad range of services is offered, furnishing clients with the information necessary to complement their internal capabilities and resources. Customized client programs involve various combinations of the following deliverables:

Status Reports	In-depth studies of important issues
Management Briefs	Detailed analysis of significant developments
Management Briefings	Periodic interactive meetings with management
Executive Presentations	Scheduled strategic presentations for decision-makers
Email Communications	Timely replies to informational requests
Telephone Consultation	Immediate response to informational needs

Clients include a cross section of IT end users in the private and public sectors representing multinational corporations, industrial companies, financial institutions, service organizations, educational institutions, federal and state government agencies as well as IT system suppliers, software vendors and service firms. Federal government clients have included agencies within the Department of Defense (e.g. DISA), Department of Transportation (e.g. FAA) and Department of Treasury (e.g. US Mint).

International Technology Group
4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net