Management Report

September 2014

IBM i on Power Systems for Midsize Businesses

Minimizing Costs and Risks for Cloud, Analytics and Mobile Environments

International Technology Group
609 Pacific Avenue, Suite 102
Santa Cruz, California 95060-4406
Telephone: 831-427-9260
Email: Contact@ITGforInfo.com
Website: ITGforInfo.com

Table of Contents

List of Figures

Executive Summary

Challenges and Solutions

The challenges faced by midsize businesses remain daunting. Yet as signs of economic recovery emerge, IT spending by these businesses is increasing. Technology continues to offer the potential for greater competitiveness, improved efficiency and higher productivity.

The technology value equation is, however, becoming broader. In most industries, organizations continue to upgrade core business systems and expand in e-commerce, while seeking to exploit new waves of analytics, mobile, cloud and social media technology.

Adoption of next-generation technologies has been driven by fundamental business trends. More complex business environments have led to unprecedented growth in analytics. Demand from customers, partners and in-house staff has made use of mobile devices routine. Growth in social media has obliged businesses to follow their customers into new channels. Cloud computing offers new opportunities to reduce IT costs, accelerate solution delivery and increase flexibility of deployment.

Certain things, however, do not change. Companies still require systems that run core business processes securely and predictably. Enterprise resource planning (ERP) and equivalent core systems remain the backbone of IT infrastructures – and as new technologies are deployed, their role becomes more, not less significant.

Analytics tools draw heavily upon customer, operational and financial data generated by core systems. Mobile solutions interact with them, and this is also the case when social media support such functions as ordering, inventory availability queries and account services. Most cloud deployments in midsize businesses complement and interface to core systems.

The quality and functionality of core systems – and of the platforms they run on – have *ripple effects* that may extend across entire application portfolios. If quality of service is impaired, processes across organizations may be impacted. If core systems are expensive to deploy and operate, fewer resources will be available to meet new technology challenges.

This report deals with these issues. Specifically, it compares the IBM i operating system deployed on POWER8-based systems with two alternatives: Microsoft Windows Server 2012 and SQL Server 2014, and x86 Linux servers with Oracle Database 12c.

There are sharp distinctions between IBM i 7.2 and Power Systems, and these alternatives. Architectures and software environments are significantly different. IBM i 7.2 and Power Systems are optimized to deliver levels of availability and security that are – by wide margins – higher than those of Windows and x86 Linux servers. Risk exposure is correspondingly less.

Such capabilities would justify a cost premium. In practice, however, overall IT costs for use of IBM i 7.2 and Power Systems may be significantly lower.

IT Costs

In representative installations in midsize manufacturing, distribution and retail companies, three-year IT costs for use of IBM i 7.2 and Power Systems average 45 percent less than for use of Microsoft Windows Server 2012 and SQL Server 2014, and 51 percent less than for x86 Linux servers and Oracle databases.

Figure 1 summarizes these results.

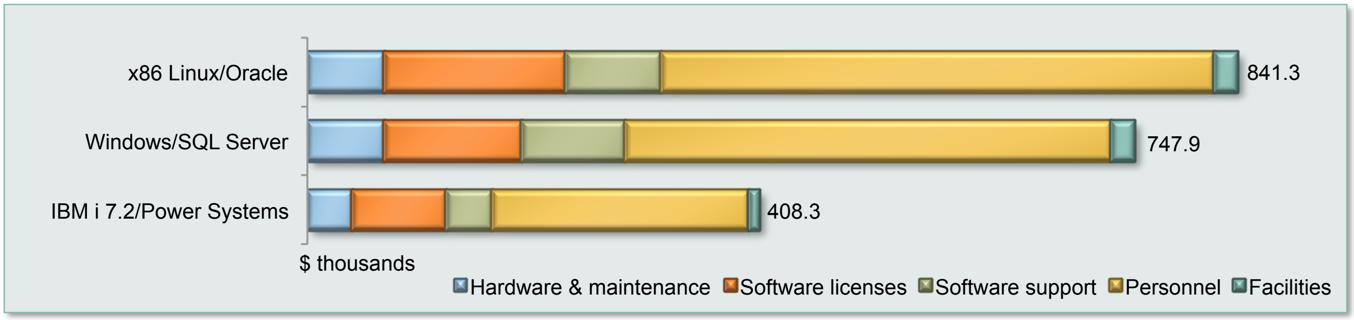

Figure 1: *Overall Three-year Costs by Platform – Averages for All Installations*

Costs included hardware acquisition and maintenance; license and support costs for operating systems, databases and other systems software; personnel costs for system and database administration; and facilities (primarily energy) costs. Hardware, maintenance and software license and support costs are based on discounted prices reported by users.

Installations are composites based on input from 42 midsize businesses employing IBM i on Power Systems, Windows servers or x86 Linux servers in manufacturing, distribution and retail companies. Companies ranged from $300 million to $1.65 billion in revenues, and employed 500 to more than 5,000 people.

Comparisons are between IBM Power S814 and S824 models equipped with POWER8 and dual- and four-socket x86 servers equipped with Intel E5 and E7 processors. IBM PowerVM, Microsoft Hyper-V and a comparable x86 Linux hypervisor are employed for virtualization.

Costs for use of IBM i 7.2 and Power Systems are lower across the board. For example, initial acquisition costs for hardware and software licenses average 35 percent less than for use of Windows and SQL Server, and 46 percent less than for x86 Linux servers with Oracle. Ongoing costs average 49 percent and 53 percent less respectively. Figure 2 summarizes these results.

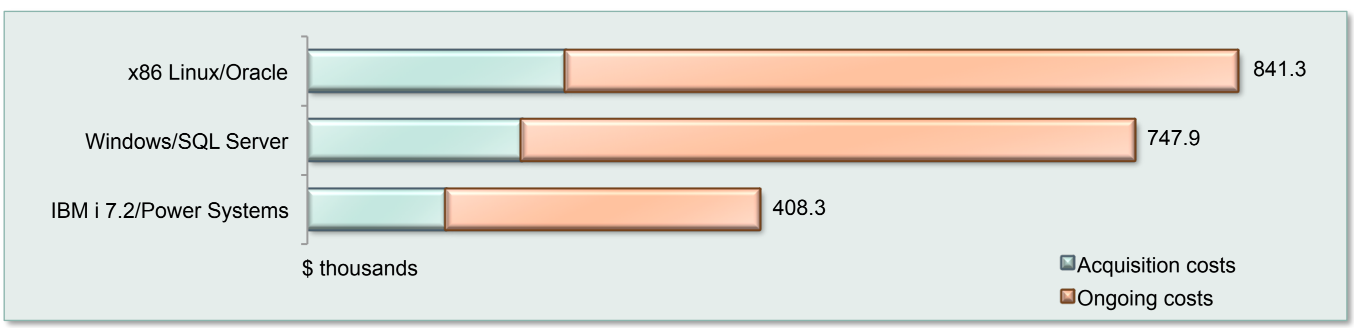

Figure 2: *Three-year Acquisition and Ongoing Costs by Platform – Averages for All Installations*

A number of differences contribute to cost disparities. More granular partitioning and real-time workload management mean that greater *workload density* may be achieved with IBM i 7.2 and Power Systems.

In Windows and x86 Linux environments, separate servers are typically deployed to handle database, application and Web serving, and to support test and development systems. Hypervisors allow sharing of some physical platforms, primarily those supporting non-production and light-duty instances. The effects on overall configurations are, however, incremental.

In smaller installations, between three and five physical x86 servers are required for workloads handled by single Power System. In others, between 6 and 11 physical servers are required for workloads handled by pairs of Power Systems duplexed for redundancy. Hardware, software and support costs are multiplied accordingly. Greater administrative complexity also increases personnel costs.

The pricing picture has recently changed in favor of IBM i and Power Systems. New POWER8-based models offer better price/performance levels than their POWER7+-based predecessors. This is particularly the case for the company's new four-core S814 model. Equally, Microsoft's per core pricing model, first employed for SQL Server 2012, has pushed up Windows database server costs.

Moreover, per core costs for SQL Server and Oracle have been affected by the trends toward use of denser multi-core processors. For example, where two-processor, eight-core servers might once have hosted core databases, the norm is now likely to be 12 or 18 cores, and high-end Xeon E7 processors embed 10 to 15 cores each.

Personnel costs are – by wide margins – lower for use of IBM i than for Windows and Linux servers. IBM i is the simplest, most automated, most tightly integrated operating system available today. It incorporates unique features such as an object-based architecture, integrated SQL database (DB2 for i), single-level storage and integrated workload management that minimize administrative overhead.

Similarly, because of the tight integration of DB2 for i, the same individual can typically handle database as well as system administration; i.e., separate database administrators (DBAs) are not be required.

Details of installations, along with methodology and assumptions employed, and cost breakdowns may be found in the Detailed Data section of this report.

Risk Exposure

Costs of Downtime

IT costs are only part of the picture. The ability to maintain high levels of availability and security also has major bottom-line implications.

It is a truism that downtime costs money. In world of globalization, Internet and mobile commerce, and social media interaction, 24/7 availability has become the norm for a growing number of systems.

The impact not only of unplanned (i.e., accidental) outages, but also of repeated planned outages for tasks such as software updates and scheduled maintenance may be substantial. Operations may be disrupted, orders and shipments delayed, and a wide range of other activities affected. Customers may be alienated and business lost.

Costs of downtime have been widely documented for core business as well as e-commerce systems. In some industries, they are clearly increasing. Among businesses that operate tightly integrated, *lean* supply chains, for example, there is growing evidence that disruptions at any point may *cascade* rapidly through the entire supply chain. The effects may continue to be felt long after service has been restored.

E-commerce companies have also learned that their customers expect 24/7 access, and that even short outages can impact sales and customer relationships. Users are only a few clicks away from competitors, and once they divert to these, they may not return.

Next-generation applications are proving to be equally if not more sensitive to uptime. For example:

- *Mobile users* are more demanding than their desktop and laptop counterparts, and tend to lose interest if they cannot access information or services within three to five seconds. In mature economies, mobile users now account for a quarter to a third of e-commerce sales, and the proportion is typically higher in developing geographies. Everywhere, numbers are increasing.

- *Social media users* expect the same level of access as through conventional Internet channels. The effects of downtime in lost business and dissatisfied customers are proving to be equally severe.

- *Cloud networks* are subject to the same dynamics. Core business systems delivered through clouds remain as sensitive to availability as on-premises equivalents. Among companies that contributed to this report, clouds had been widely adopted to interact with customers, suppliers and business partners spread across multiple geographies and time zones.

  In other cases, clouds delivered key sales, customer relationship management, financial, human resources, collaboration and other applications. Among all users, interaction with and/or use of current data generated by core business systems was the norm. Not only unplanned, but also planned outages could disrupt service to key users.

  A further point should be noted. Quality of service delivered by cloud service providers as well as private clouds was affected by underlying platforms. Clouds based on Windows and x86 Linux systems are subject to the same availability limitations as in non-cloud environments.

- *Analytics* increasingly require continuous uptime. As decision-making cycles accelerate, delays in obtaining and acting upon information may have a wide range of negative impacts. In state-of-the-art supply chains, for example, decisions about ordering, stocking and deliveries may now be made in real-time. If key systems are down, users will at best be working with stale data.

  Time sensitivity increases as companies move to solutions that embed analytics into transactional systems, enabling delivery of information and decision-making in *real time*. Most major ERP, CRM and supply chain management (SCM) vendors – including leading suppliers of IBM i-based systems – have adopted this approach, which is rapidly gaining traction among best practice users.

Where next-generation applications depend upon data supplied by, or interoperability with core business systems, the bottom-line impact of outages affecting these is magnified.

There are marked differences between platforms in this area. The availability strengths of IBM i and Power Systems have been widely demonstrated. Industry surveys, as well as user experiences, have consistently shown higher levels of uptime than for any other platform employed by midsize businesses. Planned outages are shorter and less frequent, and unplanned outages less common.

These differences are reflected in significantly lower costs of downtime; i.e., bottom-line business costs due to outages. In the same companies that form the basis of IT cost calculations, costs of downtime averaged 72 percent less than for use of Windows and SQL Server, and 79 percent less than for use of x86 Linux servers with Oracle. Figure 3 illustrates these disparities.

Calculations for all companies include costs of supply chain disruption caused by core business system outages. Costs for manufacturers and distributors also include related costs such as late delivery and imperfect order fees. Retail company costs also include costs of lost sales and in-store disruption.

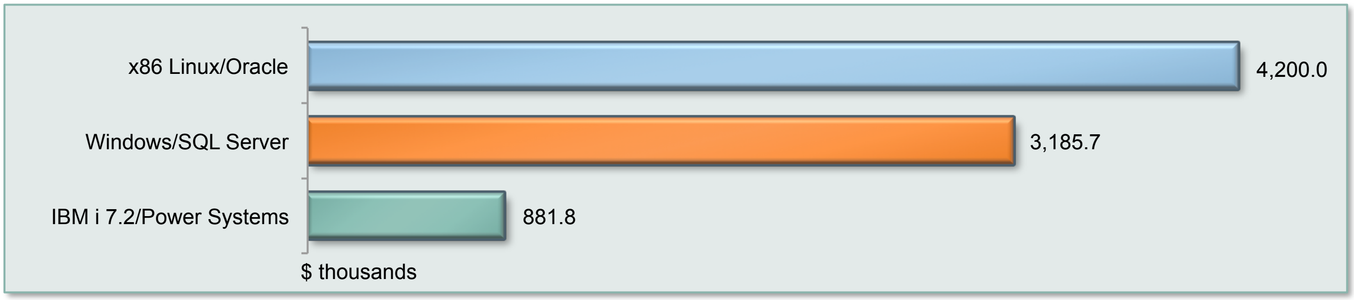

Figure 3: *Three-year Costs of Downtime – Averages for All Installations*

Allowance is also made for costs of downtime affecting next-generation analytics, mobile, cloud, social media and collaborative (e.g., IBM Connections) applications that draw data from and/or interface to core systems.

The impact of core system outages on next-generation applications was striking. The ability to process customer transactions was interrupted. Queries went unanswered because current information was not available. Salespeople were unable to quote inventory availability and delivery schedules. Service departments were unable to resolve problems. Processes supported by cloud applications were disrupted.

Comparisons allow for use of Microsoft AlwaysOn, a new SQL Server 2014 high availability (HA) feature built upon Windows Server Failover Clustering (WSFC); a comparable Linux-based solution; and the IBM equivalent, Independent Auxiliary Storage Pools (IASPs). The basis of these calculations is again described in the Detailed Data section of this report.

Security and Malware Protection

Hacking and infection by malware (malicious code) remain ubiquitous threats for organizations of all sizes. Most midsize businesses experience both on a regular basis. Many intrusions are not detected for long periods, or not detected at all.

There are, again, bottom-line impacts. Businesses that experience customer data breaches may incur fines and other penalties, along with costs of remedial actions such as notifications, credit monitoring subscriptions, query handling and technical fixes. Risks of customer loss and reputational damage may be even more significant.

As incidents such as the recent data breach affecting U.S. retailer Target have demonstrated, prevention of data loss is moving to the forefront of IT strategy. Even if customer data is not compromised, other types of sensitive information may be stolen, and malicious damage to systems and software may occur.

In these areas, differences between IBM i and competitive platforms are not merely significant – they are dramatic. IBM i security incidents are rare, and malware infection is virtually unknown. These strengths reflect the system's object-based architecture. Objects are encapsulated in a manner that places strict controls on data as well as system code, making it extremely difficult for unauthorized instructions to execute.

Capability differences are reflected in data compiled by Secunia, one of the industry's leading security and malware authorities.

Figure 4 summarizes numbers of advisory notices issued by the company between the beginning of 2008 and the end of June 2014 for the most recent versions of IBM i, the two principal Linux distributions – Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) – and the Windows Server operating system.

Operating System	Windows Server 2012	Windows Server 2008	RHEL Server 5	RHEL Server 6	SLES 10	SLES 11	i5/OS 6.x	IBM i 7.1
Release Date	*October 2012*	*February 2008*	*March 2007*	*November 2010*	*July 2006*	*March 2009*	*January 2008*	*April 2010*
Extremely critical	0	3	1	0	0	0	0	0
Highly critical	36	89	152	144	176	139	3	3
Moderately critical	10	44	225	136	90	90	6	0
Less critical	25	101	212	176	74	122	5	0
Not critical	3	11	78	55	18	23	1	1
TOTAL ADVISORIES	74	248	668	511	358	374	15	4

Source: Secunia

Figure 4: *Comparative Advisory Data – January 2008 through June 2014*

Data was not available for the latest version 7.2 of IBM i, which was introduced in April 2014.

The importance of malware resistance should be highlighted. For some time, the trend among cybercriminals has been toward use of malware that harvests information over time. Malware is now so prevalent that even minor security lapses may lead to infections.

The significance of these strengths is reinforced by two factors. One is that *costs of protecting data* are lower than those for competitive platforms. The time and effort that must be spent on routine security and malware protection, and in patching, auditing and other tasks is a great deal less.

Companies have invested in information security for decades. The sophistication of cybercriminals continues, however, to evolve, as do the techniques and technologies they employ. In response, expenditures on security tools, personnel and services continue to escalate. IBM i provides an opportunity to break this cycle.

A second factor is that, most security authorities recognize, perimeter defenses are no longer sufficient. Penetration of these has become increasingly common, and they do not prevent insider abuse. The trend is toward creation of *data firewalls*, which provide a further level of protection for the most sensitive data resources within enterprises. IBM i provides such protection, with no additional effort or cost.

Conclusions

IBM i originated with the OS/400 operating system in 1988. It was designed to provide a simple, reliable, secure and easy-to-administer platform for core business systems, and user experiences confirm that it has retained these characteristics. It is currently employed by more than 150,000 businesses and at least 80 cloud service providers worldwide, and enjoys the highest customer loyalty of any major platform.

IBM i users routinely characterize it as *highly stable...extremely robust…completely dependable…rock-solid*. Such terms are not commonly applied to Windows or x86 Linux environments. At a time when the IT world trends toward ever-greater complexity, IBM i and Power Systems offers a simple and cost-effective platform upon which to anchor increasingly diverse organizational IT infrastructure.

The significance of IBM i strengths increases as midsize organizations move to next-generation technologies. All channels through which businesses interact with customers and trading partners, and all mechanisms through which companies make decisions and apply operational processes to compete, increasingly require 24/7 availability, credible security and high levels of cost-effectiveness.

It is that simple.

Risk Trends

Overview

Key industry trends mean that the significance of IBM i strengths in availability and disaster recovery, and in security and malware resistance are increasing over time. These trends are discussed in more detail in this section.

The following section, Platform Differentiators, deals with differences in architecture and technology between IBM i 7.2 and Power Systems, and Windows and x86 Linux environments that affect comparative costs and risks.

Supply Chain Disruption

Risk Demographics

IT failures also rank as one of the most common causes of supply chain disruption. A survey of supply chain experts published by the World Economic Forum (WEF) in January 2012, found that information and communication disruption was cited by 30 percent of respondents.

Although information and communication disruption ranked as the fifth most significant trigger overall, it was ranked second among controllable risks. The distinction is important. Few supply chain or risk management executives would be penalized for failing to predict a tsunami. But potential exposure to IT-related disruptions is a great deal easier to measure, and to mitigate.

The prevalence of IT-related disruptions is confirmed by other sources. The 2013 Supply Chain Resilience Survey by the Business Continuity Institute, for example, found that IT outages were the most commonly reported cause of supply chain interruptions, cited by 57 percent of respondents. Severe weather effects ranked first, cited by 59 percent. Figure 5 summarizes these results.

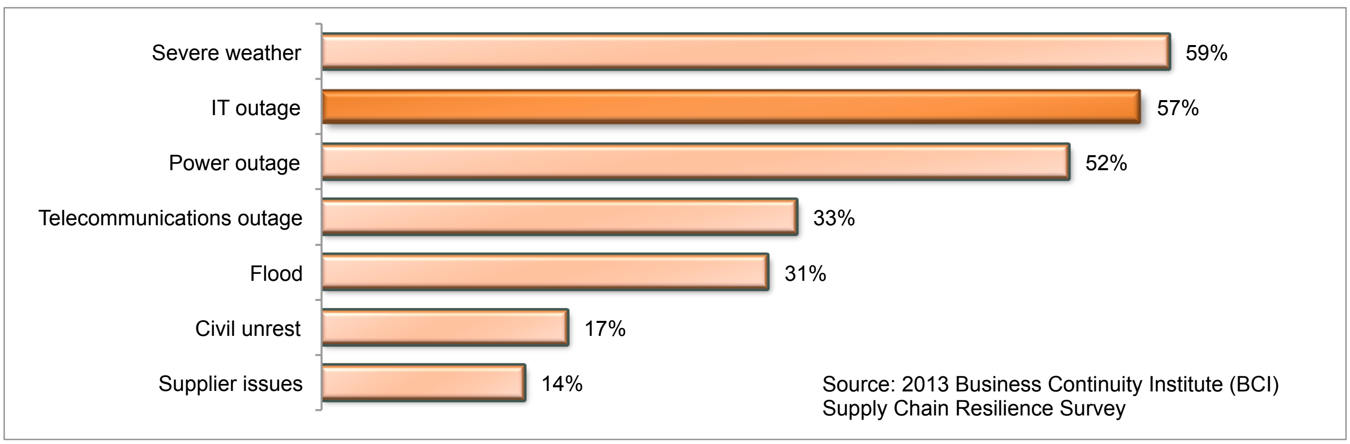

Figure 5: *Most Common Causes of Supply Chain Disruptions – Percent of Companies Reporting*

A breakdown of causes of IT outages from the same study shows that security and malware incidents, as well as production system shutdowns, were ranked highest by respondents who included 519 executives and managers worldwide. Figure 6 summarizes these results.

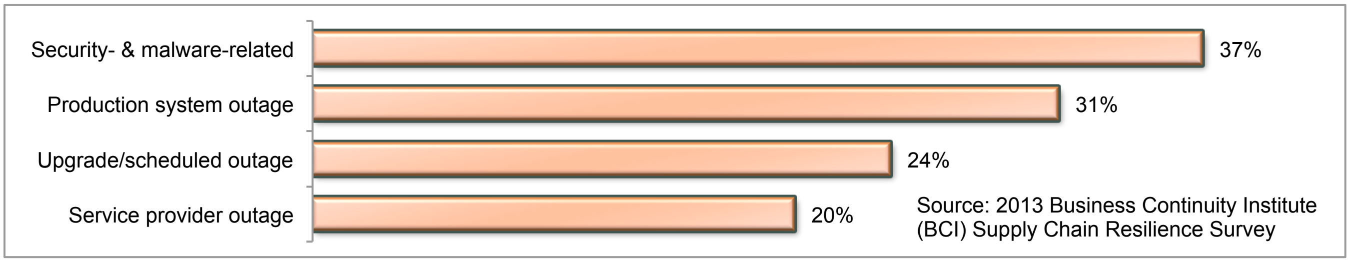

Figure 6: *Causes of IT Outages – Percent of Companies Reporting*

The results of other studies of this subject are consistent.

Industry Trends

The effects of supply chain disruption have been the subject of greatly increased attention since the mid 2000s. This shift been driven by a number of trends, including the following:

- *Integration*. Core business systems in most industries have progressively expanded to integrate a broader range of transactional processes, as well as new analytical and collaborative functions. This evolution has been particularly apparent for ERP, CRM and SCM systems, but has also affected other core systems employed in a wide range of industries.

  Examples include core merchandising systems in retail; core banking systems; policy management revenue and service delivery systems in insurance companies; customer information and billing systems in telecommunications and utilities; reservation systems in travel and hospitality; and revenue and service delivery systems in government and others.

  Interfaces between these systems and next-generation applications have increasingly become the norm. Numerous distributors and manufacturers that contributed to this report, for example, had deployed mobile and/or cloud applications to support field sales and service teams. Others had extended e-commerce and CRM systems to support customers and trading partners who preferred to interact with them using tablets or smartphones.

  Companies had also deployed cloud solutions (e.g., for supplier interaction, workforce management, and specialized ERP functions) to supplement their principal application portfolios, and to support development and test activities; and use of analytics tools was pervasive. Users ranged from CXO-level executives to customer service representatives.

  Even many small companies had adopted the use of social media for marketing, sales and account service. Manufacturers and retailers serving consumer markets have begun to move to broader customer engagement strategies.

  In virtually all cases, these applications interfaced to core business systems, either directly or (via data warehouses) indirectly. The result was that these systems increasingly became the backbone of interdependent multi-channel, multi-media complexes through which companies compete.

- *Globalization*. Most large manufacturers, retailers and distributors operate internationally, or employ foreign suppliers, channel partners or both. Key processes – including procurement, logistics and, in many cases, sales, order processing and customer interaction – now routinely occur around the clock.

  The impact of disruptions tends to be greater for global supply chains. A delay in shipping from a local plant to a nearby distribution center, for example, may mean waiting for another truck. A delay in shipping from a Chinese plant to North America or Europe may mean waiting for the next ship.

- *E-commerce and m-commerce*. The trend across many industries is toward Internet-based customer and supplier self-service systems that handle processes such as inventory availability queries, order placement and customer service.

  The Internet is, almost by definition, a 24/7 medium, and the expectation is that online systems should be accessible at any time. Figure 7, for example, illustrates the frequency of online orders placed with a wholesale distributor over a 24-hour period.

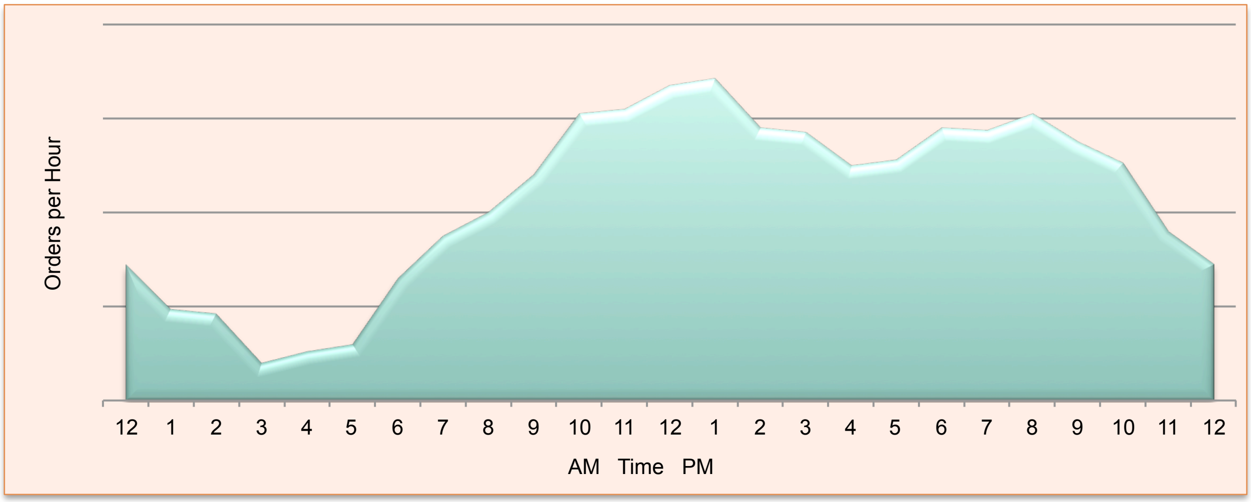

Figure 7: *24-hour Online Order Activity – Distributor Example*

  M-commerce – meaning use of mobile devices for key customer interactions – also places a premium on uptime. According to recent market research surveys, for example, more than 80 percent of U.S. tablet and smartphone owners use these for comparison-shopping from home, or while visiting stores. Even storefront visitors are now *only a few clicks away from competitors*.

  Similar trends are occurring in other geographies. In developed as well as developing economies, m-commerce represents the fastest growth segment of online retailing.

- *Supply chain strategies*. For years, in manufacturing, distribution, retail and other industries, best practice supply chain strategies have focused on lean operating models and streamlined process structures. There is an important implication: as inventory buffers removed or reduced, and process delays are eliminated, the potential impact of disruptions is magnified.

  Cycle times are declining. At the corporate or business unit level, forecasting cycles may be reduced from weeks to days, or to less than 24 hours. At the other end of the spectrum, cross docking is reducing safety stocks at distribution centers. Other best practice techniques have similar results.

  Among retailers and their suppliers, shorter planning cycles have become the norm. A decade ago, for example, sales for new products and/or promotions were typically monitored for the first 60 to 90 days. The norm is now the first two weeks, and many companies have reduced this to days or even hours. Inventories may be monitored in real time.

  In storefront retailing, the industry *rule of thumb* is 40 to 80 percent of stockouts result in lost sales rather than purchases of an alternative in-store product. Additional costs may be incurred for changes to store displays, backorders, restocking, markdowns and other actions. Online retail loss rates are even higher.

Further pressures have been added by moves toward same-day and weekend delivery. In the United States, for example, Amazon.com has driven this trend, and it can be expected that others will follow suit. In such conditions, system downtime becomes a matter of *zero tolerance*.

- *Cascading effects*. There has been a growing body of research on this subject. The basic principle is that in lean, tightly integrated supply chains operating in real time, the effects of a disruption may cascade rapidly, affecting other processes and extending to customers and trading partners.

  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 8.

SOURCE		
▪ Identify sources of supply ▪ Select supplier(s) ▪ Negotiate with supplier(s)	▪ Schedule product deliveries ▪ Receive product ▪ Verify product	▪ Transfer product ▪ Authorize supplier payment
MAKE		
▪ Schedule production ▪ Set up production ▪ Issue product	▪ Produce ▪ Inspect/test product ▪ Package product	▪ Stage product ▪ Release to delivery
DELIVER		
▪ Process inquiry & quote ▪ Receive, enter & validate order ▪ Reserve inventory resources ▪ Reserve delivery resources ▪ Determine delivery date ▪ Consolidate orders	▪ Build loads ▪ Route shipments ▪ Select carrier(s)/rate(s) ▪ Receive product ▪ Pick product ▪ Pack product	▪ Load product ▪ Generate shipping docs ▪ Ship product ▪ Customer receipt & verify ▪ Install product ▪ Invoice customer

Figure 8: *Basic Manufacturing Supply Chain Processes – SCOR Model*

  This presentation is based on selected segments of the Supply Chain Operations Reference (SCOR) model developed by the Supply Chain Council. 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.

  A delay in delivering components to a plant, for example, might cause finished product deadlines to slip. This may in turn impact transportation schedules and warehouse operations, resulting in further delays and causing disruption to spread. The effects are cumulative.

- *New technologies*. Experience has shown that adoption of *RFID* reduces cycle times for production, logistics and retail processes.

  The impact of *3D printing* may be even greater. Early adopters in such industries as apparel and accessories, consumer products and industrial manufacturing have reported that cycle times for design, production and distribution may be slashed to a matter of hours. It can be expected that similar effects will occur in other sectors as costs decline and new applications emerge.

- *Customer expectations*. Competitive pressures, acceleration of supply chains, and *on demand* use of electronic media have made customers less tolerant of supplier failures.

  A customer who is affected 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. Dissatisfaction may translate into immediate lost sales and longer-term customer attrition.

Even if customers are not lost, business-to-business suppliers may be subject to late delivery and imperfect order penalties. It may be necessary to offer rebates, special discounts and/or other incentives to win back business.

A less visible, but potentially more significant erosion of confidence might also occur. Customers might hedge by diverting some future purchases to other suppliers, and might be reluctant to rely upon the company for future time-sensitive orders. For example, no manufacturer wants to hear that customers consider them a *high-risk supplier*.

There is a further implication. Disruptions tend to raise error rates across any or all supply chain stages. This is particularly likely if there is a rush to catch up with backlogs.

An additional set of *strategic* costs may be incurred if outages are severe, protracted or both. Share prices may be affected. 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.

Outages may thus have a wide range of potential cost impacts. Figure 9, for example, shows a representative list for manufacturing companies.

STRATEGIC COSTS		
Charge against earnings **Financial metrics/ratios** **Share price decline** **Share price volatility** **Cost of capital** **Increased risk provision** **Reduced brand value** **Insurance premiums**	**Damaged reputation** ▪ Financial markets ▪ Customers/prospects ▪ Banks ▪ Business partners ▪ M&A candidates **Impaired credit** **Liquidity exposure**	**Legal exposure** ▪ Customers ▪ Third parties ▪ Shareholders **Compliance exposure** ▪ Regulatory reporting ▪ Impaired inspection ▪ Impaired traceability
CUSTOMER-RELATED COSTS		
Lost short-term sales **Lost short-term profit** **Lost future sales/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

Figure 9: *Potential Costs of Outages – Manufacturing Companies*

The potential significance of such effects was highlighted by a study 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.

Other researchers have confirmed this picture. A 2013 study by William Schmidt and Ananth Raman of Harvard Business School cited research on 517 supply chain disruptions in 412 public U.S. companies.

The authors concluded that company stock prices were reduced by an average by 3.8 percent by internal disruptions, and 1.1 percent by external disruptions caused by such factors as earthquakes, storms and supplier failures. The disparity is striking. Financial markets tend to be tolerant if incidents are judged not to have been the fault of management, but a great deal less so in the case of *self-inflicted wounds*.

Hacking and Malware

Threat Matrix

Security and malware attacks are now so common that security authorities have largely abandoned efforts to quantify their frequency. There is, however, general agreement that both the frequency and severity of incidents continues to increase.

The 2014 U.S. State of Cybercrime Survey, for example, reported that 77 percent of respondents had detected a significant security event within the previous 12 months. Among these, 34 percent reported that the number of security incidents detected had increased over the previous year, and more than 59 percent reported that they were more concerned about security threats during 2014 than they had been 2013.[1]

Comparable trends are occurring worldwide. Most estimates of the number of malware variants – including viruses, Trojans, worms, spyware, rootkits, backdoors and assorted hybrids of these – circulating on the Internet is between 600 million and 800 million. The number is expected to exceed one billion within the next two years.

Traditional distinctions between hacking (i.e., penetration of systems by individuals or organized groups) and malware are eroding. There has been consistent growth in *gateway* attacks (which use malware to create covert breaches that can be exploited over time), and in the prevalence of *spyware* (which collects and forwards information from computers without the knowledge of users).

There is general agreement on another subject: that mobile devices and social media represent the fastest-growing areas of malware activity. Businesses employing these channels must expect that security challenges will expand.

Data Breaches

Midsize users increasingly risk intrusions that expose sensitive customer data. The bottom-line impact of such incidents may be substantial.

In most countries, privacy laws expose businesses to penalties in the event of data breaches. The Data Protection Directive (a.k.a. Directive 95/46/EC) has notably affected businesses throughout the European Union, and comparable regulations are in place or expected in other geographies.

The combined total of fines, penalties and other costs may be a great deal larger than is generally realized. Figure 10 lists examples from U.S. experience.

[1]The survey, which drew responses from 500 executives of U.S. businesses, law enforcement services and government agencies, was sponsored by PricewaterhouseCoopers (PwC), CSO Magazine, the Software Engineering Institute of Carnegie Mellon University and the U.S. Secret Service.

ACTIVITY	COSTS
Investigation & technical fixes	Weeks to months using specialists at $1,000-$5,000 per person/day
Customer notification	$0.20 to $6.50 per customer, depending on medium
Query-handling	$10 to $25 per customer (call center)
Credit/identity monitoring	$65 to $250 per customer per year, depending on quality of service
Other customer remedial actions (e.g., special offers, coupons)	$15 to $500+ per customer
Reissue payment card	$3 to $15 per card
Legal costs*	Average legal defense cost: $574,984 Average legal settlement: $258,099
Fines & penalties	$600 to $200,000+
Management, PR costs	Variable
Customer attrition, brand damage	Variable
Sources: *NetDiligence, Cyber Liability & Data Breach Insurance Claims, A Study of Actual Payouts for Covered Data Breaches, 2013 All others: International Technology Group, based on 2013 and 2014 incidents	

Figure 10: *Data Breach Costs – U.S. Examples*

In the United States and Europe, most estimates put the average cost of a data breach in the range of $150 to $300 per record exposed. A leading industry authority, the Ponemon Institute, put the overall average cost of U.S. breaches in 2013 at more than $5.85 million each. In European countries, averages were in the $2.69 million to $4.74 million range, varying by country.

These averages include direct costs, including notification; remedial action such as security fixes, subscriptions to credit monitoring services and costs of marketing initiatives to retain disaffected customers; as well as indirect costs due to lost business, customer defections and other effects. In individual cases, costs may be higher.

Platform Differentiators

Overview

IBM i and Power Systems represent the convergence of two major technology streams:

1. *IBM i*, according to the company, is employed by more than 150,000 businesses worldwide. Although the installed base has decreased since the early 2000s, most of this has been due to consolidation – many organizations that had deployed AS/400s to remote sites later consolidated these to centralized systems.

   IBM i is supported by 850 software vendors, including leading suppliers of ERP systems as well as core systems for banking, insurance, retail and other industries. Its ecosystem extends to business partners, user groups, online communities, service providers and consultants worldwide.

   Although developed markets such as North America, Western Europe and Japan account for the majority of IBM i installations, this platform enjoys strong growth in emerging markets such as Latin America, Eastern Europe and the former Soviet Union, and the Asia/Pacific region.

   IBM has maintained the *technological currency* of the i environment. It currently incorporates the full function SQL-compliant DB2 database, along with full suites of Internet and mobile standards. A wide range of development languages, including C/C++, Java (including JTOpenLite for mobile devices), PHP, XML, IBM Rational Enterprise Generation Language (EGL) and others may be employed.

   Numerous open source products – including Apache Web Server, Zend, MySQL and SugarCRM – are supported on IBM i. The company has also continued to invest in established technologies such as the RPG II, COBOL and CL languages, and a wide range of application modernization tools are offered by ISV partners.

   IBM policy on i *technology upgrades* is distinctive. As a general principle, IBM introduces new i releases every two years and Technology Refreshes (TRs) – which may be applied in a simple and non-disruptive manner – every six months. This approach, which was widely requested by customers, avoids the disruptions caused by frequent version migrations.

   The IBM policy contributes to higher availability levels. Downtime, as well as risks of severe outages, tends to highest after such migrations.

   The company recently reaffirmed its *commitment to IBM i*. Announcing the POWER8 generation of technology in April 2014, IBM indicated that it would provide concurrent support for IBM i as well as AIX (the IBM version of UNIX) and Linux for all new Power Systems hardware. Enhancements to the IBM i environment are planned until at least 2026.

2. *Power Systems* are built upon IBM reduced instruction set computing (RISC) architecture. Power Systems routinely handle enterprise-class workloads requiring high levels of performance and scalability and offer highly granular, real-time virtualization even in demanding, high-volume production conditions.

   New *POWER8-based systems* deliver significant advances over previous generations of technology. Processor performance is accelerated, up to eight threads per core are supported (compared to four on POWER7+-based systems and two on x86 servers), and memory and I/O features are upgraded to support faster throughput.

According to Commercial Processing Workload (CPW) metrics, POWER8-based systems offer from 20 percent to more than two times higher performance than POWER7+-based models with the same number of cores. CPW is the standard performance metric for systems running IBM i.

Current 4- to 32-core POWER8-based systems correspond approximately to POWER7+-based 710, 720, 730, 740 and 750 models – which, according to IBM, accounted for 85 percent of Power Systems unit shipments in 2013. They provide sufficient performance for the vast majority of midsize users.

In terms of price/performance, IBM i on POWER8-based systems are more aggressively positioned than was the case for POWER7+-based systems. This is notably the case for the new IBM 4-core S814 model, which qualifies for the least expensive P05 IBM i pricing bracket.

IBM offers a variety of pricing options for i on Power Systems. Solution Editions, which are customized for more than 160 industry-specific and cross-industry ISV offerings, offer packages of hardware, software and services that reduce overall costs by up to 25 percent. Pre-configuration and testing for individual customer requirements also enables more rapid and cost-effective deployment.

Other programs allow users of larger Power Systems to license Power cores employed for application serving at lower costs than for database servers; and offer capacity-based pricing for PowerVM partitions and IBM i-defined workloads. In addition, Managed Server Provider (MSP) utility pricing offers special terms for service providers.

IBM i

Distinctive characteristics include the following:

- *Core design.* The core IBM i design is built around a unique object-based kernel in which all system resources are defined and managed as objects.

  The kernel incorporates single-level storage capability – meaning that the system treats all storage resources, including main memory and disks, as a single logical entity. This capability, as figure 11 illustrates, is built into the core system design.

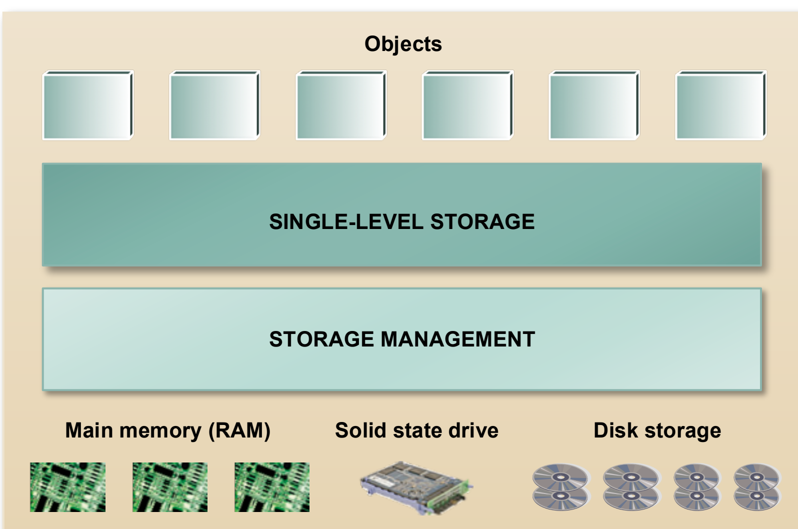

Figure 11: *IBM i Single-level Storage Structure*

Placement and management of data on all resources is handled automatically by the system, minimizing tasks that must be handled by administrators. This approach also enables high levels of configuration flexibility; and materially improves the efficiency with which processor and storage resources are used. A further benefit is that integration and management of solid state drives (SSDs) is comparatively simple.

Another kernel component, *Technology Independent Machine Interface (TIMI)* acts as a *virtual* instruction set with which applications interact, regardless of the instruction set of underlying processor hardware. TIMI allows users to update underlying hardware platforms without recompiling applications software.

There are no Windows or Linux equivalents to single level storage or TIMI.

- *System integration.* In addition to core operating system functions, IBM i includes the DB2 for i relational database, file systems, WebSphere Application Server (WAS), Tivoli Directory Server, Java Virtual Machine (JVM) environments, Zend Server and more than 300 tools handling system, database, storage, backup and recovery, communications, security, operations and other management tasks.

  Software components are not simply bundled. They are implemented in a highly synergistic manner, and engineered to interact with each other simply and efficiently. For example, DB2 for i exploits the underlying object-based structure and single-level storage capabilities of the operating system.

  Equivalent functionality in Windows and Linux server environments requires that users acquire, install, configure and administer hardware and software products from multiple vendors. Integration and testing is less coordinated, and upgrades may not follow the same schedules. Deployment complexity and management challenges are increased.

  In addition to increasing FTE administration staffing less integrated environments are more likely to degrade performance. Maintenance of availability, recoverability and security also become a great deal more problematic.

- *Workload management.* Since its inception, IBM i has incorporated industry-leading workload management (in IBM i terminology, *work management*) capabilities.

  IBM i subsystems leverage object-based architecture. Individual workloads or applications (e.g., ERP, CRM, e-mail, e-commerce) are described and managed independently. The system allocates memory, limits consumption of resources by individual workloads, and manages scheduling, tuning and other tasks automatically, based on user-defined priorities.

  Subsystems are integral to the IBM i design, and may be employed independently of or in conjunction with PowerVM virtualization.

- *Security and malware resistance.* The strengths of IBM i's object-based design are reinforced by tight integration of security functions with compiler, directory server and object-based file system structures. In contrast, security in Windows- and Linux-based environments is implemented in separate software tools and subsystems. The level of integration is significantly less.

  IBM i also contains a full IP security suite, including support for the principal industry security standards and encryption techniques, and extensive access control and audit facilities. Single sign-on is enabled using an industry-leading IBM autonomic technology, Enterprise Identity Mapping (EIM), which maps user IDs across all middleware and application components.

- *Simplification and automation.* IBM i was designed to automatically handle a wide range of functions – including configuration, tuning, software updates, availability and security optimization and other common operational tasks – for which most other systems require extensive manual intervention.

High levels of automation are combined with streamlined, high-productivity interfaces. These allow administrators to perform system and database management tasks using fewer, simpler actions, in less time than for Windows- and Linux-based environments.

The most visible effect of automation is that it reduces FTE staffing. Other benefits may also be realized. For example, a system that can determine workload requirements and reallocate system resources in a matter of milliseconds will use capacity more efficiently than one that is dependent on intervention by administrators. Automation also reduces the potential for human errors.

Core automation features have been reinforced by IBM autonomic technologies. Autonomic computing – meaning the application of artificial intelligence technologies to IT administration and optimization tasks – has been a major IBM development focus since the 1990s.

Four categories of autonomic functions – self-configuring, self-optimizing, self-protecting and self-healing – are implemented in IBM i. These functions, which represent one of the most advanced implementations of autonomic technologies within the IBM product line, are summarized in figure 12.

Additional features further *minimize risks of data loss* in the event of an unplanned outage. These include kernel-based IASPs, Remote Journaling (file and system changes may be automatically copied to a second server) and Save While Active (backups may be performed without taking systems offline).

A further capability, Live Partition Mobility, allows movement of active Logical Partitions (LPARs) between systems without disrupting operations. Service interruptions of one or two seconds may occur due to network latency, but are rarely noticeable to users.

IBM PowerHA SystemMirror for i enables *live failover clustering*. The best practice norm for this solution is that operations may be fully restored within two hours with no data loss. Users have achieved mainframe-class failover and recovery even for complex large-scale workloads. This solution, previously named High Availability Clustered Multiprocessing (HACMP), has been in widespread use since the 1990s.

In comparison, Microsoft's *SQL Server 2014 AlwaysOn* is a newer, less mature offering that corresponds to IASPs rather than PowerHA SystemMirror for i. Experiences with its base component, WSFC, have shown that this is complex, requires a great deal of administrator intervention, and is limited in scalability. Failover and recovery processes may effectively handle simple, low-volume environments, but tend to be slower and less reliable for larger workloads and databases.

Experiences with Linux-based failover and clustering solutions have generally been similar. This has particularly been the case for open source HA tools.

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 file/print services 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, Dynamic SMT Dynamic System Optimizer, Expert Cache Global resource manager Heterogeneous workload manager Processor-memory affinity processing Quality of service optimization Single-level storage Workload Group Support	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 spreading & disk allocation Automatic data striping & disk balancing Automatic disk space allocation Automatic distributed access configuration Automatic object placement Automatic self-balancing indexes Automatic tablespace allocation Automatic TCP/IP startup Graphical database monitor	Automatic Encryption management Automatic enforcement of user query & storage limits Automatic synchronization of user security Digital object signing Object auditing OS-controlled resource management
Self-optimizing	**Self-healing**
Adaptive Query Processing Automatic Index Advisor Automatic memory pool tuning Automatic query plan adjustment Automatic rebind & reoptimization Automatic statistics collection Auto Tuner, Cost-based Query Optimizer Caching of open data paths & statements On Demand Performance Center Performance monitoring & analysis	Automatic object backup/restore Automatic database object extents Automatic database restart Automatic index rebalancing Automatic journaling of indexes & objects Automatic rebuild of catalog views Automatic restart of journal processing Self managed database logging Self-managed journal receivers Systems managed access path protection

Figure 12: *IBM i 7.2 and IBM Power Systems Autonomic Functions*

Power Systems

IBM has progressively enhanced the Power Systems platform since its introduction in 1990. Recent processor generations have included POWER7 (February 2010), POWER7+ (February 2013) and POWER8 (April 2014). IBM has also put in place industry-leading capabilities in the following areas:

- *Performance optimization*. Power Systems performance is a function not only of POWER processors, but also of close optimization at all levels of design and implementation. Key capabilities include highly effective compiler acceleration; chip symmetric multithreading (SMT); low levels of symmetric multiprocessing (SMP) overhead; and on-chip memory acceleration and compression technologies.

  Active Memory Expansion, for example, enables compression rates of up to 50 percent for data in memory; i.e., usable main memory may be up to double physical memory. In addition to lower memory costs, system throughput is improved.

  POWER8-based systems incorporate additional performance boosts. The maximum number of threads per core, for example, increases to eight. A new hardware-based transactional memory feature accelerates high-volume parallel applications, and a Coherent Accelerator Processor Interface (CAPI) enabling higher-bandwidth CPU access for specialized co-processors.

  A new on-chip accelerator also reduces overhead for encryption and decryption of data conforming to Advanced Encryption Standard (AES) specifications.

- *PowerVM virtualization*. This is one of the industry's most sophisticated virtualization architectures. PowerVM offers highly granular combinations of hardware- and software-based partitioning.

  Hardware-based LPARs isolate workloads more effectively than software-based partitioning techniques, and provide additional security functions. Hardware-based partitioning is not supported on Windows or x86 Linux servers.

  PowerVM also allows AIX, the IBM variant of UNIX, and Linux to run in partitions on the same physical system as IBM i. Linux support has allowed IBM i users to deploy Internet and intranet infrastructures, along with open source applications on Power Systems that also host core business systems.

- *Mixed workload management* is another key strength. Features such as intelligent threading and intelligent cache adjust numbers of threads and cache configurations respectively to more efficiently execute workloads. Parameters may be set automatically or by system administrators.

  The overall architecture, which is illustrated in figure 13, combines IBM i and Power Systems features to manipulate a wider range of variables – including threads, processors, cache, main memory and I/O, multiple types of partition, multiple threads and dedicated or pooled processors – with higher levels of granularity and flexibility than any competitive platform.

  VMware, Hyper-V and other x86-based hypervisors also provide certain of these capabilities. There are, however, significant differences in the efficiency with which these operate.

  Partitioning creates the potential for high levels of capacity utilization. However, the extent to which this occurs in practice depends on mechanisms that allocate system resources between, and monitor and control workload execution processes across partitions. If these mechanisms are ineffective, a high proportion of system capacity may be idle over time.

  Administrators will, moreover, tend to operate systems at well below capacity to avoid risks of bottlenecks. These effects explain why, in many organizations, use of x86 hypervisors falls short of their architectural potential. Power Systems enable higher workload densities.

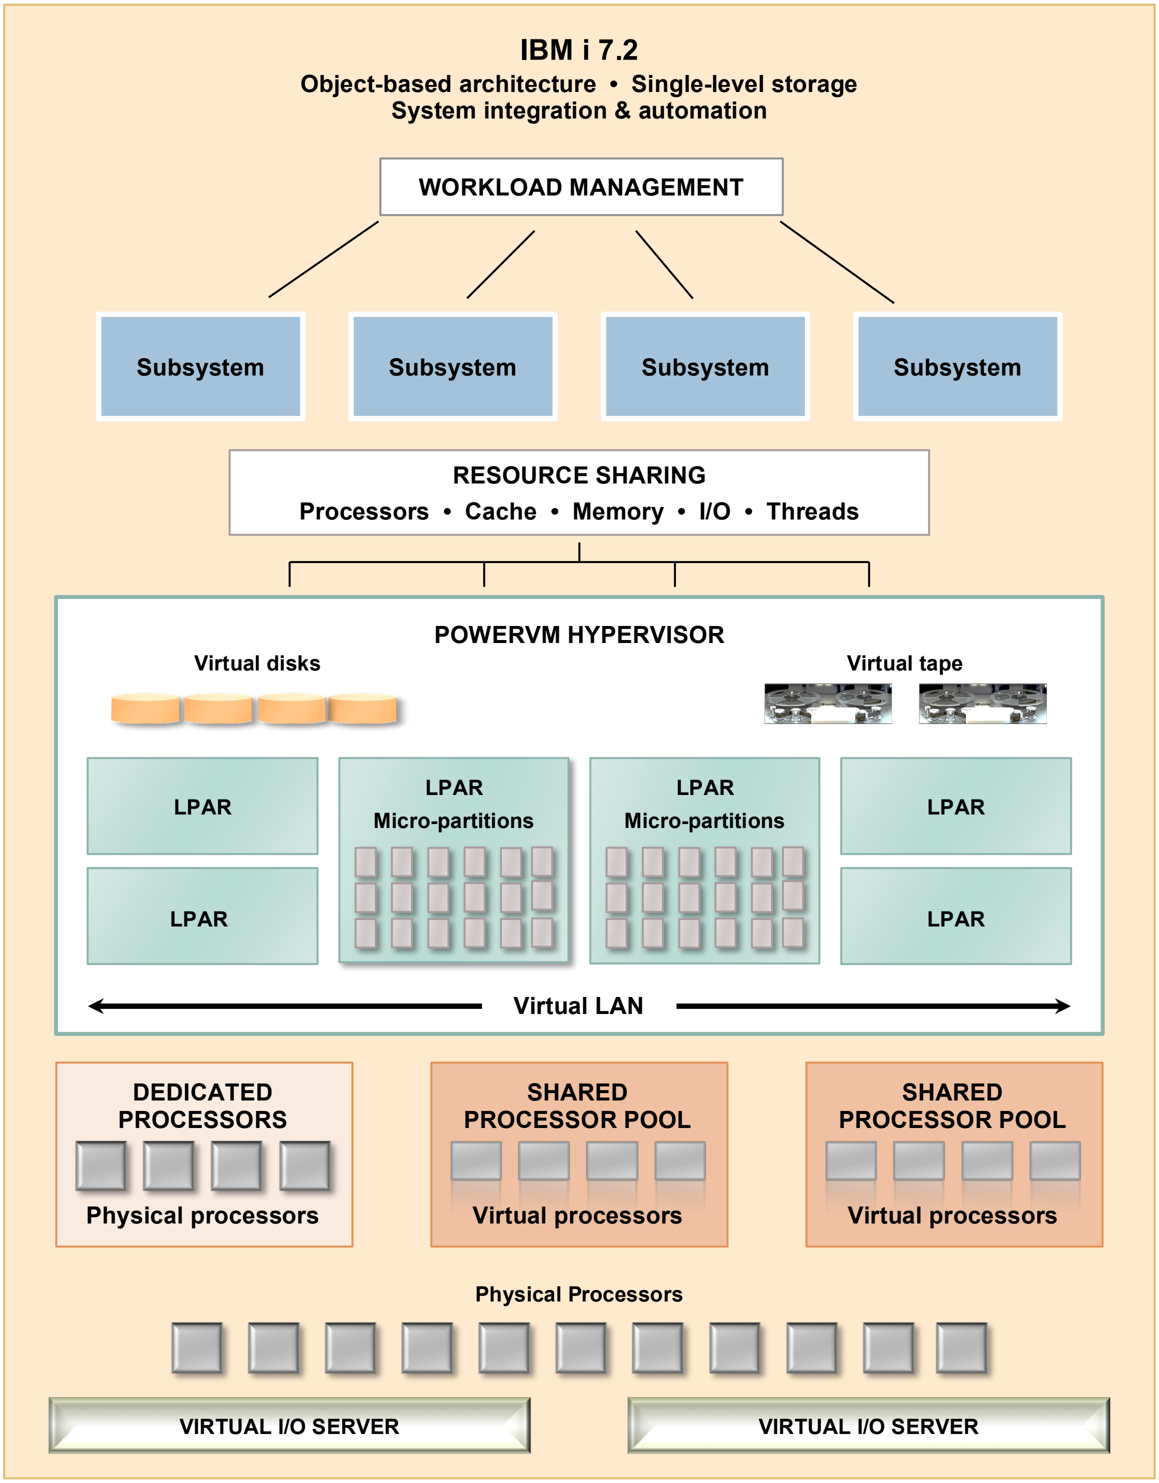

Figure 13: *Overall IBM i and Power Systems Architecture*

In addition, reliability, availability and serviceability (RAS) features in Power Systems hardware are among the most sophisticated in the industry today. Comparable features may be found in x86 servers in some cases. However, the microelectronics technology in Power Systems is more advanced, and systems have longer track records of stable and effective operation.

Detailed Data

Installations

The cost comparisons presented in this report are based on the installations, server configurations and FTE staffing levels summarized in figure 14. Next-generation applications are highlighted in blue.

CONSUMER PRODUCTS DISTRIBUTOR	INDUSTRIAL DISTRIBUTOR	SPECIALTY RETAILER
Business Profile		
Health & cosmetic products distributor $500 million sales 500 employees 3 distribution centers	Specialty industrial distributor $300 million sales 650 employees 10 distribution centers	$300 million sales 1,500 employees 250 stores 2 distribution centers
Applications		
ERP/CRM, query & reporting Data warehouse: product performance, customer relationship, sales & supplier analysis; sales & operations planning, revenue planning, promotions management Mobile sales & customer service Social media merchandising & marketing	CRM, order management, finance, HR, supply chain, e-commerce Distribution analytics system Sales & talent management clouds Mobile sales Social media marketing	Core retail/merchandising, supply chain, finance, HR, Web & wireless applications Retail analytics, workforce management & supplier clouds Mobile point of sale Social media customer engagement, merchandising & marketing, online product/inventory data & sales
Number of Users		
300	350	400
PLATFORM SCENARIOS		
IBM i 7.2/Power Systems		
Power S814 4-core 3 GHz IBM i 7.2, PowerVM *0.3 FTE*	Power S814 4-core 3 GHz IBM i 7.2, PowerVM *0.3 FTE*	2 x Power S814 4-core 3 GHz IBM i 7.2, IASPs, PowerVM *0.45 FTE*
Windows/SQL Server		
3 x 2/8 x 2 GHz Windows Server 2012 with Hyper-V, SQL Server 2014 AlwaysOn *0.5 FTE*	1 x 2/8 x 2.5 GHz 3 x 2/8 x 2 GHz Windows Server 2012 with Hyper-V, SQL Server 2014 AlwaysOn *0.55 FTE*	2 x 2/12 x 2 GHz 3 x 2/8 x 2.3 GHz 1 x 2/8 x 2 GHz Windows Server 2012 with Hyper-V, SQL Server 2014 AlwaysOn *0.95 FTE*
Linux/Oracle		
1 x 2/12 x 2 GHz 2 x 2/8 x 2.2 GHz Linux, Oracle Database 12c *0.55 FTE*	2 x 2/12 x 2 GHz 3 x 2/8 x 2.2 GHz Linux, Oracle Database 12c *0.6 FTE*	2 x 2/12 x 2.5 GHz 7 x 2/8 x 2.2 GHz Linux, Oracle Database 12c *1.05 FTEs*

Figure 14: *Installations and Scenarios Summary*

DISCRETE MANUFACTURER	PROCESS MANUFACTURER	AGRIBUSINESS COMPANY
Business Profile		
Industrial machinery & components manufacturer $600 million sales 2,500 employees 5 manufacturing & distribution centers	Food & beverage manufacturer $1 billion sales 2,000 employees 6 manufacturing plants	$1.65 billion sales 5,000 employees 10 production & distribution centers
Applications		
ERP, supply chain, e-commerce Analytics: demand forecasting & planning; customer, sales, inventory, procurement, cost & productivity analysis, quality management Sales & CRM cloud Mobile sales Social media marketing	ERP, CRM, supply chain, e-commerce, departmental Analytics: demand forecasting, revenue planning; customer, product sales & merchandising analysis; financial analysis; traceability & other applications Marketing, sales & other collaborative applications (IBM Connections network) Mobile/online consumer engagement, promotions and coupons Social media marketing & research	ERP, procurement, e-commerce, EDI Analytics: 30+ business & compliance applications CRM, procurement, talent management & warehouse management clouds Mobile e-commerce extensions
Number of Users		
600	500	1,200
PLATFORM SCENARIOS		
IBM i 7.2/Power Systems		
Power S814 1/6 & 1/4 x 3GHz IBM i 7.2, IASPs, PowerVM *0.65 FTE*	2 x Power S824 1/6 x 3.89GHz IBM i 7.2, IASPs, PowerVM *0.6 FTE*	2 x Power S824 1/8 x 4.15GHz IBM i 7.2, IASPs, PowerVM *1.0 FTE*
Windows/SQL Server		
2 x 2/16 x 2GHz 1 x 2/12 x 2.3GHz 5 x 2/12 x 2GHz 1 x 2/8 x 2GHz Windows Server 2012 with Hyper-V, SQL Server 2014 AlwaysOn *1.25 FTEs*	2 x 2/24 x 2.7GHz 5 x 2/12 x 2GHz 1 x 2/12 x 2GHz Windows Server 2012 with Hyper-V, SQL Server 2014 AlwaysOn *1.2 FTEs*	2 x 4/40 x 2GHz 7 x 2/16 x 2.6GHz 2 x 2/12 x 2 GHz Windows Server 2012 with Hyper-V, SQL Server 2014 AlwaysOn *2.3 FTEs*
Linux/Oracle		
2 x 2/16 x 2 GHz 1 x 2/12 x 2.3 GHz 5 x 2/12 x 2 GHz 1 x 2/8 x 2 GHz Linux, Oracle Database 12c *1.2 FTEs*	2 x 2/24 x 2.7 GHz 5 x 2/12 x 2 GHz 1 x 2/12 x 2 GHz Linux, Oracle Database 12c *1.3 FTEs*	2 x 4/40 x 2 GHz 7 x 2/16 x 2.6 GHz 2 x 2/12 x 2 GHz Linux, Oracle Database 12c *2.65 FTEs*

Figure 14 (cont.): *Installations and Scenarios Summary*

In these tables, numbers of processors and cores are shown for all platforms; e.g., 1 x Power S814 1/4 x 3 GHz refers to a Power S814 with one quad-core 3 GHz POWER8 processor, while 2/12 x 2 GHz refers to an x86 server with two six-core Intel E5 processors.

Installations were constructed using data on applications and workloads, server configurations, database and system administrator staffing and other variables supplied by 42 companies employing IBM i on Power Systems, Windows servers or x86 Linux servers to support ERP and equivalent core systems. Companies were in the same industries and approximate size ranges, with similar business profiles.

Based on these inputs, six composite company profiles were created. IT costs and costs of downtime per hour for all affected applications were calculated for each company.

IT Costs

These were calculated as follows:

- *Server costs* include hardware and software licenses, along with three-year maintenance and software support coverage. Maintenance and support costs are for 24/7 coverage. Calculations were based on vendor list prices discounted to reflect prevailing street prices.

  Power Systems were POWER8-based models with IBM i 7.2 and PowerVM Standard or Enterprise Edition. Branded x86 servers were equipped with current-generation Intel Xeon E5 or E7 processors.

  Windows servers were configured with Windows Server 2012 R2 Standard or Datacenter Editions with Hyper-V and, for database serving, SQL Server 2014 Enterprise Edition with AlwaysOn. License costs for Microsoft software include Client Access Licenses (CALs). Software support costs are for Microsoft Software Assurance.

  x86 Linux servers were configured with a major commercial Linux distribution including hypervisor and HA components, along with Oracle Database 12c Enterprise Edition. Oracle configurations also include Diagnostics and Tuning Packs providing functionality equivalent to that incorporated by IBM in i 7.2, and by Microsoft in SQL Server 2014 at no additional charge.

- *Personnel costs* were calculated based on annual average salaries of $89,745 for IBM i 7.2 and Power Systems administrators handling database as well as system administration tasks; $77,820 and $80,022 for Windows and x86 Linux system administrators respectively; and $93,986 and $98,070 for SQL Server and Oracle DBAs respectively.

  Salaries were increased by 56.7 percent to allow for benefits, bonuses, training and other personnel-related items. Overall costs were calculated for a three-year period.

- *Facilities costs* include data center occupancy and energy consumption by servers as well as by power, cooling and other data center infrastructure equipment supporting these.

  Energy costs were calculated using vendor electricity consumption values for server configurations and data center equipment. Specific utilization levels and hours of operation for the installation were then applied, and a conservative assumption for average price per kilowatt/hour was employed to determine three-year costs.

IT cost breakdowns are presented in figure 15. Costs for Linux subscriptions are included in software licenses.

Company	Health Products Distributor	Industrial Distributor	Specialty Retailer	Discrete Manufacturer	Process Manufacturer	Agribusiness Company
IBM i 7.2/POWER SYSTEMS						
Hardware & maintenance	14,973	17,470	34,294	38,973	51,546	86,709
Software license	12,260	12,260	24,521	93,684	158,862	207,917
Software support	9,010	9,010	18,020	44,965	72,425	96,220
Personnel	126,567	126,567	189,851	274,229	253,135	421,891
Facilities	4,263	5,645	9,932	9,786	11,532	13,121
TOTAL ($)	**167,073**	**253,246**	**387,879**	**514,374**	**625,891**	**917,310**
WINDOWS/SQL SERVER						
Hardware & maintenance	31,308	30,028	53,134	87,718	84,869	130,509
Software license	44,883	43,014	96,788	130,380	159,340	272,952
Software support	33,663	32,260	72,591	97,785	119,505	204,714
Personnel	194,315	216,407	370,339	487,688	469,397	898,410
Facilities	7,860	9,746	10,845	22,893	30,493	43,536
TOTAL ($)	**368,288**	**544,849**	**669,311**	**998,286**	**1,098,480**	**1,494,093**
x86 LINUX/ORACLE						
Hardware & maintenance	31,308	30,028	53,134	87,718	84,869	130,509
Software license	85,420	93,420	128,130	192,195	162,840	322,808
Software support	42,240	47,520	63,360	95,040	79,200	190,080
Personnel	223,870	246,921	428,930	493,842	527,218	1,081,730
Facilities	7,860	9,746	10,845	22,893	30,493	43,536
TOTAL ($)	**498,292**	**711,745**	**858,080**	**1,344,230**	**1,401,115**	**1,896,290**

Figure 15: *Three-year IT Costs Breakdown*

Costs of Downtime

Downtime costs were calculated using a two-phase process. First, average costs per hour of downtime were calculated for all companies using appropriate industry- and organization-specific values.

Averages were calculated for two sets of costs:

1. *Direct operational costs for downtime due to core business system outages.*

   Calculations for all companies include costs of supply chain disruption, including idle and underutilized capacity, including personnel; handling of delivery delays (including distribution center and transportation costs); additional inventory carrying costs; costs of customer billing delays; customer penalties; and remedial costs such as buyback rebates.

   For manufacturing companies, calculations also allow for production disruption, including costs of supplier order, production scheduling and setup and other changes.

   For the retail company, calculations included lost storefront sales due to stockouts; costs of storefront disruption, including increased markdowns, set-up delays, handling and administrative costs; and remedial actions such as reordering, display changes and restocking.

2. *Costs of downtime due to core business system outages affecting next-generation applications*.

   For all companies, costs of downtime include lost sales, supply chain disruption and additional overheads caused by outages affecting customer interaction through Internet, mobile and social media channels; and interaction with direct sales personnel though clouds (e.g.; Salesforce.com), mobile applications and, in the case of the process manufacturer, IBM Connections.

   For supplier clouds operated by the retailer and agribusiness, costs of downtime were for supply chain disruption, and were calculated in the same manner as for direct operational costs.

   For analytics, workforce management clouds and non-sales IBM Connections applications, costs of downtime are for lost productivity by individuals unable to obtain key information or initiate actions due to core business system outages. Calculations were based on conservative assumptions for time lost, and average annual salary and benefits per hour for users affected.

Average costs of downtime per hour were then multiplied by numbers of hours of downtime per year for each platform. These were calculated based on user input. The focus was placed on downtime for underlying hardware and software platforms, rather than application downtime. Annual costs of downtime for each platform were then multiplied for three-year totals.

All values for IT costs and costs of downtime were for the United States.

International Technology Group

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

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.

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 complex IT deployment agendas. A broad range of services is offered, furnishing clients with the information necessary to complement their internal capabilities and resources.

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).