ITG | International Technology Group

Management Report

September 2014

Value Proposition for Infor ERP Solutions Deployment on IBM i

Comparing Capabilities and Costs with Microsoft Windows Servers

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

The Landscape Changes

Manufacturing is among the world's most challenging industries. Pressures to cut costs, improve efficiency and strengthen customer relationships remain pervasive. Companies must also grapple with increasingly volatile business environments that require them to plan and execute faster, and with greater flexibility.

As the effects of recession diminish, IT spending is rising. Companies are investing in next-generation enterprise resource planning (ERP) systems that address not only transaction processing, but also such functions as planning and forecasting, supply chain optimization and customer relationship management (CRM).

Manufacturers are also beginning to exploit the potential of cloud computing, analytics, mobile, social media and other new technologies. The new ERP landscape is significantly different, and its potential for business transformation is far-reaching.

Infor has been at the forefront of these trends. The *Infor10* solution set combines state-of-the-art core ERP capabilities, one of the industry's most sophisticated integration suites, and tools that address the full range of new application opportunities. Focused vertical capabilities address all of the parameters through which manufacturing companies compete in today's markets.

Core ERP systems become increasingly critical to the overall Infor10 environment. They determine the efficiency and flexibility of transaction processes. They also, increasingly, act as the backbone of supply chain structures, and assume new importance as sources of real-time data for analytical and customer-facing solutions.

Selection of a server platform to host core ERP systems becomes a business-critical decision. Key criteria will include performance, as well as the ability to maintain availability for 24/7 operations, and – an increasingly important issue – security and resistance to malicious code (malware). In a time of continued pressures on expenditure, platforms must also be cost-effective.

This report deals with the option of deploying Infor core ERP systems – specifically, *Infor10 LX, XA and System21* – on IBM i on Power Systems. This combination meets these criteria in a manner that is – by wide margins – superior to Windows and SQL Server platforms implementing the latest Microsoft technologies.

Cost-effectiveness

Three-year IT costs for Infor ERP installations on Windows and SQL Server-based systems averaged 62 percent more than installations on IBM i on Power Systems. This result, which is based on representative manufacturing company installations, is illustrated in figure 1.

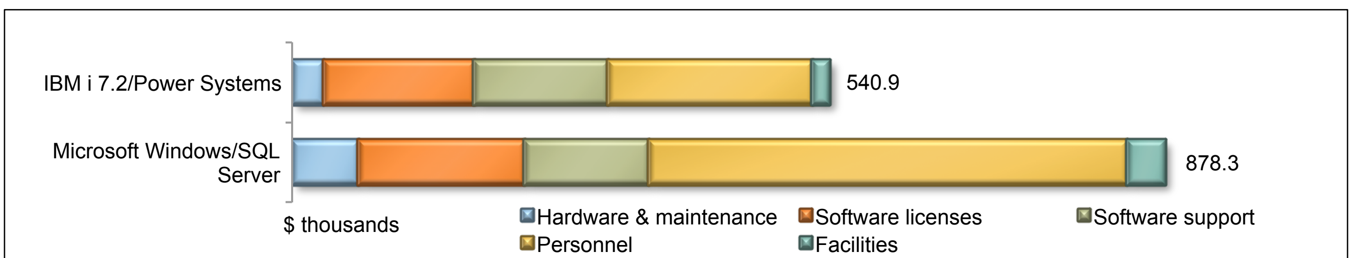

Figure 1: *Three-year IT Costs for IBM i on Power Systems and Microsoft Windows Servers Supporting Core ERP Systems – Averages for All Installations*

These three-year IT costs were, in fact, 50 percent higher for Infor LX installations, 58 percent higher for Infor XA installations and 98 percent higher for Infor System21 installations on Windows and SQL Server-based systems than on IBM i on Power Systems. The specific breakdowns for individual applications are presented later in this report.

IT costs include hardware acquisition and maintenance, software licenses and support, full time equivalent (FTE) personnel for system and database administration and facilities, including data center occupancy and energy. Hardware, maintenance and software costs were calculated based on discounted prices reported by users. Costs of applications software are not included.

Installations were constructed using data on applications and workloads, server configurations and database, system and storage administrator staffing supplied by 44 companies employing Infor10 or comparable Windows-based ERP systems. Users were in the same industries and approximate size ranges, with similar business profiles.

Comparisons were based on use of IBM i 7.2 on Power Systems with POWER8 and, in smaller installations, POWER7+ processors; and current-generation Intel E5-based x86 servers with Windows Server 2012, including Hyper-V, and SQL Server 2014. Failover and recovery was provided by use of independent auxiliary storage pools (IASPs) in IBM i 7.2 and by Microsoft AlwaysOn respectively.

Availability Impacts

Higher availability levels enabled by IBM i and Power Systems have important bottom-line impacts. It is a truism that downtime costs money. In manufacturing companies, operations may be disrupted, personnel and capacity idled, orders and shipments delayed, and a wide range of other activities affected. Customers may be alienated and business lost.

Globalization and Internet commerce require 24/7 availability. Moreover, not only is there growing evidence that the effects of downtime may be greater than is generally realized, but also that effects may continue to felt long after service has been restored.

In lean structures operating with minimal inventory buffers, the effects of an outage may generate *cascading* effects that spread rapidly across the entire supply chain. Vulnerabilities are reinforced by integration of systems, cycle time reductions and growing interdependency of processes.

Lower Windows and SQL Server availability is reflected in higher three-year *costs of downtime*; i.e., business costs due to outages. In the same companies employed for IT cost calculations, costs of downtime for use of Windows and SQL Server-based systems averaged 4.6 times more than for use of Infor10 systems on IBM i and Power Systems. This result is illustrated in figure 2.

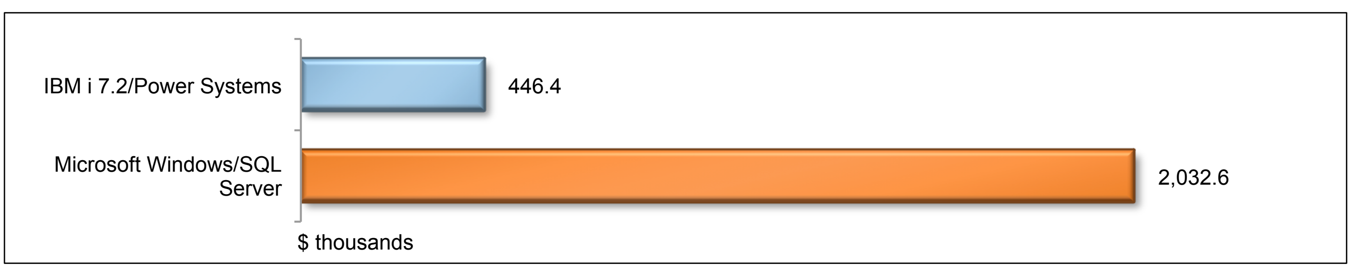

Figure 2: *Three-year Costs of Downtime for IBM i on Power Systems and Microsoft Windows Servers Supporting Core ERP Systems – Averages for All Installations*

Costs of downtime for use of Windows and SQL Server-based systems averaged approximately four, six and five times more than for use of Infor10 LX, XA and System21 respectively on IBM i and Power Systems.

Costs include idle and underutilized capacity and personnel; handling of delivery delays; additional inventory carrying costs; costs of change for procurement, production and logistics processes; costs of customer billing and payments delays; late delivery and imperfect order penalties; and costs of remedial actions. The methodology employed is described in the Detailed Data section of this report.

Technology Differentiators

Lower IT costs and costs of downtime compared to Microsoft Windows and SQL Server on x86 hardware reflect fundamental differences in architecture and technology. There are sharp differences in overall designs, and in hardware and software environments between these and IBM i on Power Systems.

Differences include the following:

- *IBM i* is the most highly integrated and automated operating environment in the industry today. Core features – including a unique object-based kernel and single-level storage – are tightly coupled with the DB2 for i, a full-function SQL relational database; multiple file systems; Web application and services servers; and more than 300 management tools.

  IBM i was designed to automate a wide range of functions – including configuration, tuning, subsystem management, software updates, availability and security optimization and other operational tasks – which in most other systems require extensive manual intervention.

  Because of the tight integration of DB2 for i, moreover, administrators can typically handle database as well as system administration. Separate database administrators (DBAs) are rarely required.

  As a result, *FTE staffing* for system and database administration is – by wide margins – lower than for Windows and SQL Server. In ERP deployments, the industry norm is that Windows FTE staffing is two to four times higher than for IBM i-based systems.

  In the comparisons presented in this report, numbers of FTE administrators for use of Windows and SQL Server averaged over 2.5 times more than for use of IBM i and Power Systems. Three-year personnel costs averaged 2.3 times more. Ratios were largely consistent among Infor LX, XA and System 21 users.

  IBM i also offers industry-leading *workload management* built around *IBM i subsystems* that are designed to handle diverse applications and workloads in a highly efficient manner. One implication is that, in contrast to Windows-based deployments, it is not necessary to employ separate database, application, batch and Web servers.

  Most IBM i users are able to run all core ERP system components on a single physical system. This results in higher capacity utilization, reduced management complexity and lower hardware and software costs. A second system is typically employed for failover. In Windows and SQL Server environments, multiple servers are required in this role.

  IBM i strengths in *security and malware resistance* should be highlighted – it is virtually immune to hacking and malware. The system's object-based architecture places strict controls on data as well as system code, making it extremely difficult for unauthorized instructions to execute. Security violations are rare, and there are no known native IBM i viruses.

  These differences are reflected in data compiled by Secunia, one of the industry's leading authorities on security and malware exposure. Figure 3 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 and Microsoft Windows Server.

Severity	Windows Server 2012	Windows Server 2008	IBM i 7.1	i5/OS 6.x
Extremely critical	0	3	0	0
Highly critical	36	89	3	3
Moderately critical	10	44	0	6
Less critical	25	101	0	5
Not critical	3	11	1	1
TOTAL ADVISORIES	74	248	4	15

Source: Secunia

Figure 3: *Comparative Advisory Data: January 2008 Through June 2014*

IBM i strengths in this area not only reduce the risks to which companies are exposed, but also contribute to lower FTE staffing costs. The time and effort that must be spent on routine security and malware protection tasks, and in patching and auditing is a great deal less than for Windows-based environments.

- *Power Systems* have been recognized server performance leaders since the mid-2000s. This is function not only of POWER processors, but also of performance optimization at all levels of design and implementation.

  Key capabilities include highly effective compiler performance acceleration; chip symmetric multithreading (SMT); low levels of symmetric multiprocessing (SMP) overhead; and unique on-chip memory acceleration and compression technologies.

  Power Systems also employ one of the industry' most sophisticated virtualization architectures. *PowerVM* offers highly granular combinations of hardware- and software-based partitioning techniques that are closely integrated with IBM i workload management. Higher levels of concentration may be realized than with x86 hypervisors such as VMware and Hyper-V.

  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 a wide range of open source applications on Power Systems that also host ERP systems.

New *POWER8-based systems*, introduced in April 2014, 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 memory and I/O features are upgraded to support faster throughput.

According to Commercial Processing Workload (CPW) benchmarks, POWER8-based systems offer 20 percent to more than two times higher performance than POWER7+-based models with the same number of cores. The average is over 57 percent. CPW is the standard performance metric for systems running IBM i.

High levels of *availability* reflect features built into the IBM i kernel, and into Power Systems hardware and firmware. IBM i has a long track record of stability and robustness. Unplanned outages are rare, and planned outages for such tasks as software updates and modifications, and scheduled maintenance are less frequent and shorter than for Windows-based systems.

The overall simplicity, integration and automation of the IBM i environment minimize the potential for downtime. Specialized features include *Independent Auxiliary Storage Pools*, which allow data to be mirrored to and recovered from alternate systems, and others.

Further protection is offered by IBM or third-party high availability (HA) clusters. *IBM PowerHA SystemMirror for i*, for example, builds upon IASPs to provide more advanced database mirroring, failover and recovery.

Operating system features, as well as IBM PowerHA SystemMirror for i, have been deployed by tens of thousands of users worldwide and used in some cases for decades.

Microsoft AlwaysOn, which is implemented in SQL Server 2012 and 2014, is a more rudimentary, less mature solution built upon Windows Server Failover Clustering (WSFC). Both are less reliable than IBM i equivalents, and require significantly more personnel time to deploy and administer.

Power Systems incorporate reliability, availability and serviceability (RAS) features that 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.

Conclusions: Dealing with Complexity

Most areas of manufacturing, production and supply chain operations are becoming more complex. This trend has contributed to widespread replacements of and upgrades to ERP systems that are unable to meet these challenges.

The new generation of ERP systems – which includes Infor LX, XA and System 21, along with the middleware and solutions that accompany them – provides new tools to deal with complexity. More effective use of information and collaboration extends across all processes within the enterprise to enable greater responsiveness to customers and faster, more effective realization of business opportunities.

Complexity is not, however, only a business issue. Next-generation ERP systems incorporate a broader, more rapidly changing portfolio of applications and technologies than their legacy counterparts. Properly exploited, these may open new dimensions of competitiveness. Carelessly exploited, these may create unmanageable complexities within IT infrastructures.

IBM i and Power Systems represent – by a wide margin – the simplest, most stable, most cost-effective platform available today for deployment of ERP systems. They deliver industry-leading performance, availability and security; require fewer staff and cost less overall than Windows-based environments; and enable companies to avoid the disruptions and distractions of overly complex IT infrastructures.

This is the case regardless of whether systems are operated by customers or provided through a cloud – to which Infor as well as IBM are heavily committed and well positioned. Business-critical systems delivered through a cloud require the same capabilities as those operated by customers.

It is also worth highlighting the close technical relationship between IBM and Infor. The two companies have worked closely together for decades and this collaboration has resulted in exceptional integration and optimization across the hardware and software solution stack. This insures that the advanced capabilities of Power Systems and IBM i are fully realized by the Infor ERP applications and tools.

The Infor and IBM Power Solution Center of Excellence (COE) acts as a focal point for this work across the two companies, leading technical collaboration and supporting joint-customer activity. This collaboration has been most evident during the rollout of IBM's POWER8-based systems and has resulted in one of the strongest IT industry partner relationships in existence today.

As the challenges that face manufacturers continue to mount, the manner in which next-generation ERP systems are employed will, in no small measure, determine how well these are met. There is no reason to make deployment and operation of business-critical systems more difficult – or expensive – than needs to be the case.

Risk Trends

Availability

The importance of maintaining high levels of uptime for business-critical ERP systems has been recognized for decades. Companies have planned and invested heavily to guard against the effects of unplanned outages and data center disasters.

Avoidance of planned outages has also become a priority. Businesses operating on a 24/7 basis are finding it difficult to find windows during which software updates, scheduled maintenance and other processes can be performed without disrupting production operations.

Most industries are vulnerable to outages. In manufacturing, however, availability is a particularly critical requirement. There are a number of reasons for this, including:

- *Integration*. Core ERP systems have progressively expanded to integrate a broader range of transactional processes, as well as new informational and collaborative functions.

  The Infor10 portfolio, for example, has expanded to include CRM, enterprise asset management (EAM) product lifecycle management (PLM), along with forecasting, planning, analytics, mobile and cloud computing as well as other tools. While these may run on other platforms, they interact with and depend upon data supplied by core ERP systems.

  A core ERP outage would, at a minimum, oblige users to work with stale data. Salespeople could not access current inventory availability data, production and logistics scheduling would become problematic, and customer service systems and personnel would be reduced to using historical information. At worst, an outage could grind the entire company to a halt.

- *Globalization*. Even relatively small manufacturers operate internationally, or employ foreign suppliers, channel partners or both. Certain processes – including procurement, logistics and, in many cases, sales, order processing and customer service – now routinely occur around the clock.

  The impact of disruptions tends to be greater for regional and global supply chains. For example, rescheduling shipments may be a significantly more demanding process for offshore suppliers and global logistics contractors.

- *Supply chain strategies*. For years, best practice strategies have focused on lean operating models and streamlined process structures. As inventory buffers are removed or reduced, and process delays are eliminated, the potential impact of disruptions increases.

  At the corporate or business unit level, forecasting and planning cycles may be reduced from weeks to days, or to 24 hours or less. A growing number of manufacturers receive continuous demand signals from their customers, and initiate procurement, production and logistics actions on a daily or even hourly basis.

  At the other end of the spectrum, cross docking (i.e., the immediate transshipment of goods between arriving and departing vehicles, without intermediate storage) in distribution centers may increase not only efficiency, but also vulnerability to disruption.

  In such environments, the effects of an outage may cascade through the entire supply chain. 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.

- *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 on a 24/7 basis.

  M-commerce – meaning use of mobile devices such as tablets and smartphones for key business interactions – also places a premium on uptime. Interaction within organizations, and with customers and partners, may now occur continuously. The effects of disruption are further magnified.

- *RFID and 3D printing* will also tend to reduce cycle times for production, distribution center and logistics processes, and design and production processes respectively. Early adopters of 3D printing have included companies in fashion, consumer products, aerospace and engineering, industrial manufacturing and other industries where Infor ERP systems are widely used.

- *Customer impacts*. 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. The long-term impact may be significantly greater.

  Even if customers are not lost, companies may be subject to late delivery and imperfect order penalties. It may be necessary to offer special discounts or terms and conditions to win back the customer's 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. No manufacturer wants to hear that customers consider them a *high-risk supplier*.

An additional set of *strategic* costs may be incurred 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.

System outages may thus have a wide range of potential cost impacts. Figure 4 shows examples.

There is another implication. Disruptions tend to raise error rates. This may occur across multiple stages of the supply chain, and may cause further customer dissatisfaction, penalties and remedial costs.

The potential significance of supply chain disruptions was highlighted by a *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.

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 publicly traded U.S. companies. The study focused on the impact of publicly reported disruptions since the passage of the Sarbanes-Oxley Act in 2002.

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*.

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

Figure 4: *Potential Costs of Outages for Manufacturing Companies*

Security and Malware

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 2013. [1]

The number of malware variants – including viruses, Trojans, worms, spyware, rootkits, backdoors and assorted hybrids of these – circulating on the Internet continues to expand. Security firms estimate that there are currently between 600 million and 800 million unique variants, and the number is expected to exceed one billion within the next two years.

[1] The survey, which drew responses from 500 executives of US 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.

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).

Use of bots and zombies – which allow attackers to gain control over a computer for illicit purposes – has also begun to evolve in new directions. Security analysts report that attackers are increasingly targeting servers, rather than PCs, because of the greater bandwidth that these offer for spam and malware propagation. New malware variants also increasingly target mobile devices.

Although manufacturers are usually not targeted as often as financial institutions, retailers and e-commerce companies, they remain exposed to a growing momentum of threats. In any industry, a security weakness may be rapidly and opportunistically exploited by any number of attackers worldwide.

In most countries, privacy laws expose companies to penalties in the event of data breaches. Other costs –for investigation, customer notification, credit monitoring, problem fixes and other items – may also be substantial.

In the United States and Europe, for example, 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.

Companies have invested in information security for more than a decade. The extent of improvement is, however, questionable. The sophistication of cybercriminals continues to evolve, as do the techniques and technologies they employ. The general industry perception is that companies are at best *breaking even*, while expenditures on security tools, personnel and services continue to escalate.

Use of IBM i provides an opportunity, at least for core ERP data, to break this cycle.

A final point should be noted. A growing number of analysts have noted that perimeter defenses are no longer sufficient. Penetration of these is common and perimeter defenses do not protect against 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 already provides this capability.

Technology View

IBM i

Status Report

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. Although IBM has upgraded its technology base since that time, user experiences confirm that it has retained these characteristics. It enjoys an exceptional record of customer satisfaction and loyalty.

According to IBM, more than 150,000 companies worldwide currently use the operating system. It benefits from an ecosystem that includes more than 2,300 solutions from over 850 independent software vendors (ISVs), along with a global network of business partners, user groups, online communities, service providers and consultants.

Although developed markets such as North America, Western Europe and Japan account for the majority of IBM i installations, IBM i and Power Systems are enjoying 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 IBM i environment. It has been progressively enhanced to incorporate latest-generation SQL technology, C/C++, Java and Eclipse, the PHP Web enablement language, XML, the MySQL database, Apache Web server, IBM Rational Enterprise Generation Language (EGL) and others. IBM i is also fully supported by IBM mobility solutions.

The company and its ISV partners have also continued to invest in established IBM i technologies such as the RPG II, COBOL and CL languages.

Announcing the POWER8 generation of technology in April 2014, IBM made it clear that the company's policy was to provide concurrent support for IBM i as well as AIX (the IBM version of UNIX) and Linux for all new Power Systems platforms.

The most recent IBM i versions 7.1 and 7.2 were introduced in April 2010 and April 2014 respectively. The company has indicated that it plans for new releases and Technology Refreshes (TRs) until at least 2026.

Technology Refreshes are periodic software upgrades – typically introduced every six months – that provide new features and functions in an incremental manner. This approach was introduced at the request of customers who wished to avoid the disruptions caused by frequent version upgrades.

Infor Commitment

With 2013 sales of $2.7 billion, Infor is the third largest ERP software vendor after SAP and Oracle. The company expanded rapidly during the 2000s through a series of more than 30 acquisitions. More recently, it has stabilized its operational structure, rationalized its solution portfolio, and begun to invest aggressively in new product development, support and services.

In 2011, Infor announced its Infor10 strategy. The company began to emphasize *micro-vertical* solutions focused on specific industry segments, and began to deliver new XML-based ION middleware designed to integrate these, and to support new cross-platform offerings. ION solutions incorporate Microsoft SharePoint, SQL Server and Windows Server, along with the Silverlight development framework and Microsoft security technologies.

Seven of the company's principal ERP systems – LX, XA, System21, Lawson M3 and S3, Infor Distribution A+ and the Infinium suite of financial and human capital management (HCM) applications – support IBM i.

According to Infor, its IBM i customers include more than 15,000 manufacturers and around 4,000 companies in other industries. Although the company reports more than 60,000 customers overall, IBM i users are typically larger, and contribute disproportionately to Infor revenues and profit.

Infor management has repeatedly emphasized that it will continue to enhance and support its IBM i-based Infor10 solutions indefinitely. The company operates a dedicated IBM i business unit, and development and support of IBM i-based solutions also occurs elsewhere in the Infor organization.

Key Differentiators

Key IBM i differentiators include the following:

- *Core design*. The core IBM i design is built around an 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 5 illustrates, is built into the core system design.

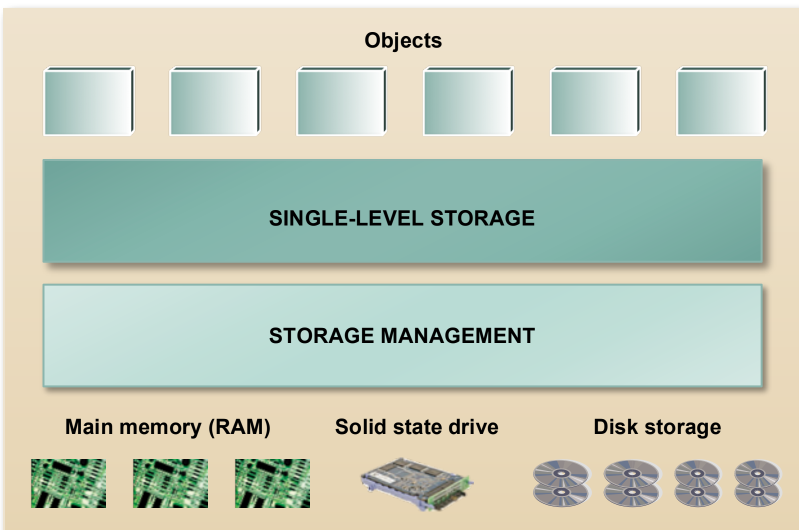

Figure 5: *IBM i Single-level Storage Structure*

  Placement and management of data on all resources is handled automatically by the system, minimizing tasks that must be handled by administrators.

  A further benefit of single-level storage is that integration and management of solid state drives (SSDs) is comparatively simple. The operating system automatically places the *hottest* (most frequently accessed) data on SSDs, reallocates data to SSDs or hard drives as workloads evolve, and optimizes performance on an ongoing basis. No application changes are required.

  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. The TIMI allows users to update underlying hardware platforms without recompiling applications software.

- *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 and tools handling a wide range of management tasks.

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

This approach contributes to low FTE staffing for system and database administration. Improvements in performance (efficient software structures generate lower system overhead) and quality of service (tightly integrated, tested systems are less likely to experience outages) may also be realized.

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

  The backbone of these capabilities is provided by *IBM i subsystems*, which leverage the IBM i object-based architecture. Individual workloads or applications (e.g., ERP, CRM, e-mail, Web serving) are described and managed independently. The system allocates memory, limits consumption of resources by individual workloads, and manages scheduling, tuning and other tasks automatically, or based on priorities set by users.

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

- *Security and malware resistance*. In addition to native object-based security, described earlier, IBM i benefits from features built into DB2 for i. These have been progressively expanded in recent Technology Refreshes and in IBM 7.2, which includes a new IBM technology, Row- and Column-level Access Control (RCAC).

  RCAC, which also forms part of IBM DB2 for Linux, UNIX (LUW), enables exceptionally fine-grained security through a combination of row- and column-based techniques. It is an emerging standard among the most demanding corporate and government users.

  IBM i also contains a comprehensive IP security suite. This includes encryption support, and extensive access control and audit facilities. Single sign-on is enabled using an IBM autonomic technology, Enterprise Identity Mapping (EIM).

- *Simplification and automation*. High levels of IBM i integration are accompanied by simplified interfaces that allow administrators to perform system and database management tasks using fewer, simpler actions, in less time than for Windows and SQL Server environments.

  Management of IASPs and other IBM i availability features is also a great deal simpler than for Microsoft AlwaysOn and WSFC.

  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. The company is the recognized industry leader in this area.

  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 6.

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 6: *IBM i 7.2 and IBM Power Systems Autonomic Functions*

In addition, *avoidance of planned as well as unplanned outages* is a central IBM i design parameter. High levels of stability, integration and automation minimize risks of unplanned outages caused by software failures and human error, and reduce the frequency and duration of planned outages.

Specialized features further minimize risks of data loss in the event of an unplanned outage. In addition to IASPs, key features include 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.

Power Systems

Overview

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 performance optimization, mixed workload management, virtualization and other areas.

POWER8-based systems are offered by IBM in a different models designed for use with Linux, AIX and/or IBM i. Those supporting IBM i include the single-socket 2U model S814 and the dual-socket 4U S824. These employ 4-, 6-, 8- and 12-core processors rated at 3.02 GHz to 4.15 GHz. S814 and S824 systems may be configured with up to 8 and 24 cores respectively.

The new 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. The company continues to market the POWER7+- based 760 (up to 48 cores), 770 (64 cores) and 780 (128 cores), and the POWER7-based 795 (up to 256 cores).

POWER8-based systems incorporate additional performance boosts. The maximum number of threads per core, for example, increases from four to eight, and cache and main memory optimization is improved. 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.

POWER8-based systems also employ PCIe Gen3 I/O, replacing IBM GX++ technology employed in earlier Power Systems.

The latest systems continue to implement key POWER7-based mixed workload management features such as intelligent threading and intelligent cache, which adjust numbers of threads and cache configurations to more efficiently execute individual workloads. Parameters may be set automatically or by system administrators.

Another retained feature, Active Memory Expansion, enables system-managed compression and decompression of data in memory. Compression rates of up to 50 percent are supported; i.e., usable main memory may be up to double physical memory.

PowerVM Virtualization

PowerVM virtualization supports three types of partitioning:

1. *Logical partitions* are microcode-based partitions that may be configured in increments as small as 1/10th core. The technology was originally developed for IBM mainframes.

   This approach (a.k.a. hard partitioning) isolates workloads more effectively than software-based techniques. Workloads running in different partitions are less likely to interfere with each other, enabling higher levels of concentration. LPARs also provide additional security functions.

   System resources used by LPARs may be dedicated (Static LPARs), or shared according to application priorities (Dynamic LPARs). Static LPARs are typically employed for applications with high levels of business criticality.

   Hard partitioning is not supported on Windows servers.

2. *Micro-partitions* are software-based partitions. They are typically employed to support instances requiring limited system resources, and to improve load balancing for large, complex workloads. Micro-partitions may be configured in initial increments of 1/10th core, and subsequent increments as small as 1/100th core.

   LPARs and micro-partitions are supported by mechanisms that allow processor, memory and I/O resources to be pooled and re-allocated in an extremely granular manner. The system monitors resource utilization every 10 milliseconds, and may change allocations as rapidly.

   Business-critical workloads may run in dedicated LPARs, using dedicated physical processors. Other workloads may be executed based on assigned priorities using combinations of threads, partitions and shared processor pools.

3. *Virtual I/O Servers* allow operating system instances running in multiple LPARs to share a common pool of LAN adapters as well as FC, SCSI and RAID devices; i.e., it is not necessary to dedicate adapters to individual partitions. Hardware, maintenance and energy cost savings may be realized. Virtual I/O Servers are duplicated to provide redundancy.

The overall IBM i and Power Systems architecture, which is illustrated in figure 7, combines IBM i and Power Systems capabilities 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. Workload management for IBM i on Power Systems is a great deal more effective.

The implications are important. 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.

Most workloads experience fluctuations, and processes (e.g., online, batch, collaborative) may vary. Unexpected spikes may also occur. When multiple applications are concentrated on a single physical platform – particularly if these generate mixed workloads – highly granular, real-time monitoring and resource assignment will be required.

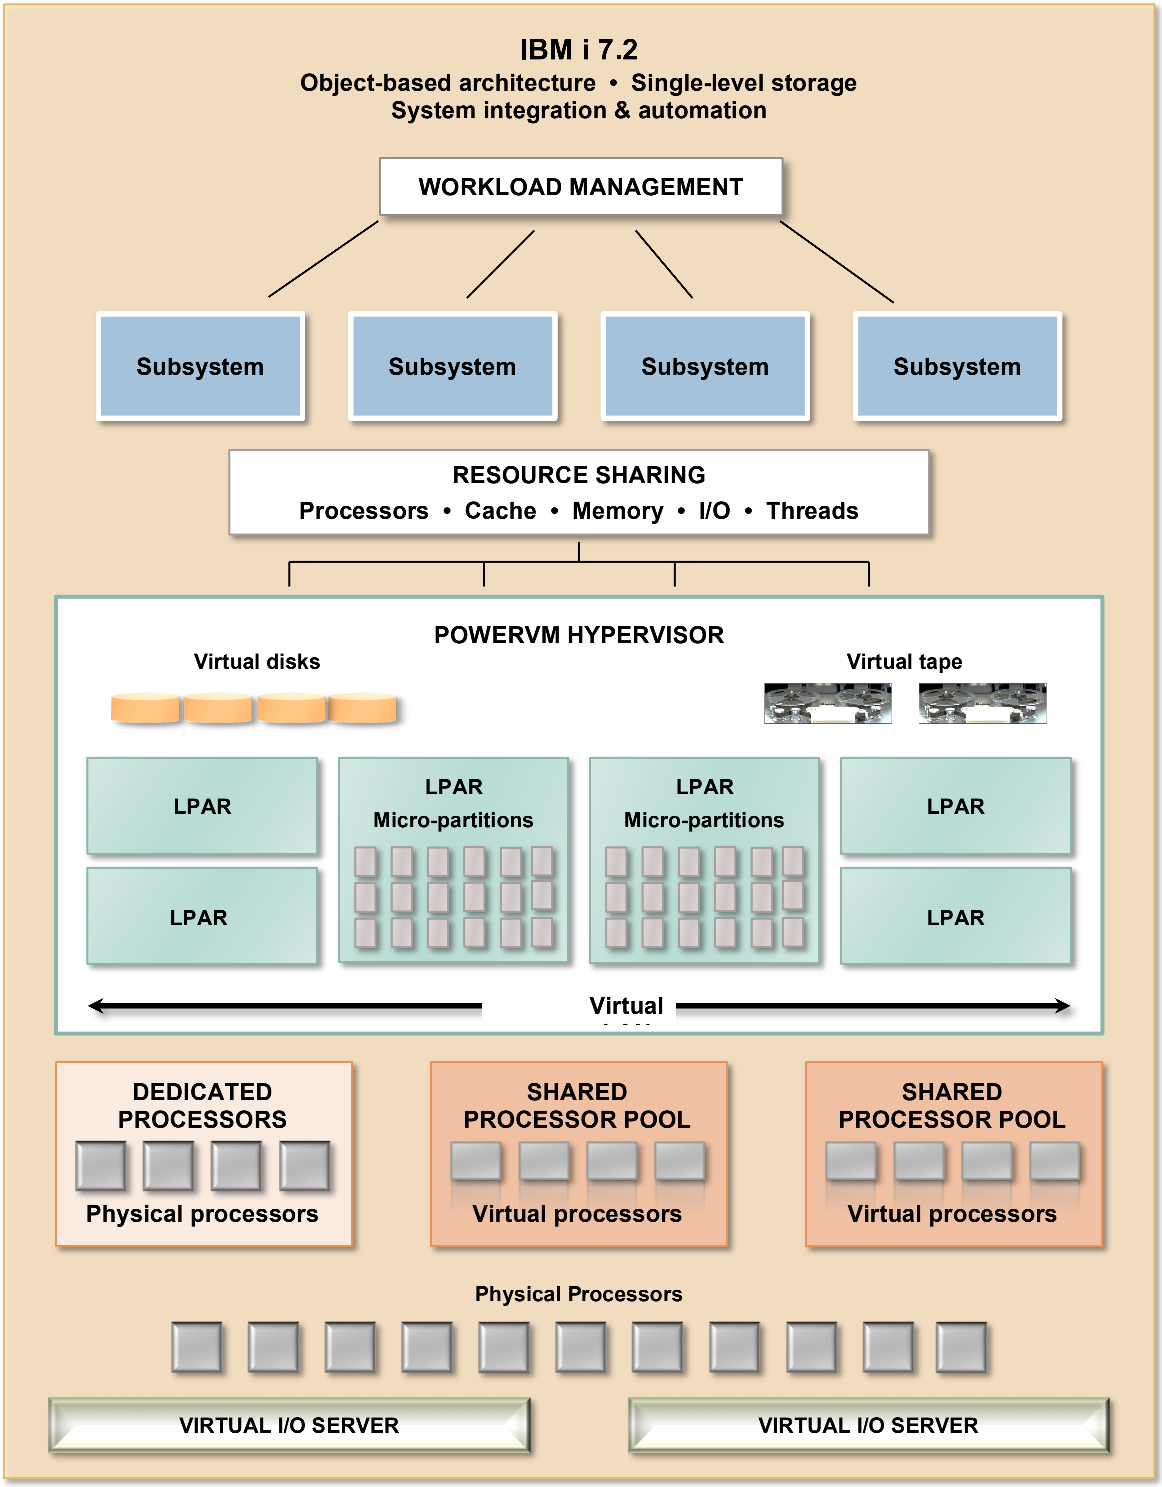

Figure 7: *Overall IBM i and Power Systems Architecture*

If systems cannot provide such capabilities, administrators will tend to limit the number and size of partitions to prevent workloads interfering with each other. This is one of the weaknesses of such tools as VMware and Hyper-V, and helps explain why most installations of these realize only a fraction of their architectural potential.

RAS Features

Key RAS features in Power Systems include the following:

- *Basic capabilities* include high levels of component reliability and redundancy, along with hot swap capabilities enabling devices to be replaced without taking systems offline. Redundant and hot swap components include disk drives, PCI adapters, fans, blowers, power supplies and, on high-end models, system clocks, service processors and power regulators.

- *Monitoring, diagnostic and fault isolation and resolution facilities* are built into all major Power Systems components, including processors, main memory, cache and packaging modules, as well as adapters, power supplies, cooling and other devices. Multiple layers of protection and self-test are implemented.

  IBM-developed *Chipkill* and *First Failure Data Capture (FFDC)* technologies are significantly more reliable than conventional error correction code (ECC) techniques. FFDC employs embedded sensors that report failures to a separately powered *Service Processor*, which also monitors environmental conditions.

- *Fault masking capabilities* prevent outages in case failures do occur. For example, in the event an instruction fails to execute due to a hardware or software fault, the system will automatically repeat the operation. If the failure persists, the operation will be repeated on a different processor and, if this does not succeed, the failed processor will be taken out of service.

  In addition, *memory sparing* enables alternate memory modules to be activated in the event of failures, and *enhanced memory subsystem* enables memory controller and cache sparing.

These features are summarized in figure 8. Additional capabilities are provided for high-end Power 770, 780 and 795 models.

Basic capabilities	
Redundancy, hot-swap & related	Redundant/hot-swap disks, PCI adapters, GX buses, fans & blowers, power supplies, power regulators & other components. Redundant disk controllers. I/O paths & oscillators. Concurrent system clock repair.
Concurrent firmware update	Server microcode may be updated without taking systems offline.
Concurrent maintenance	Allows processors, memory cards & adapters to be replaced, upgraded or serviced without taking systems offline.
Monitoring, diagnostics & fault isolation/resolution	
Hardware-assisted memory scrubbing	Automatic daily test of all system memory. Detects & reports developing memory errors before they cause problems.
Chipkill error checking	Employs RAID-like striping of data across memory devices to provide redundancy & enable reinstatement of original data. Significantly more reliable than conventional error correction code (ECC) technology.
First Failure Data Capture (FFDC)	Employs 1,000+ embedded sensors that identify errors in any component. Causes of errors are determined without the need to recreate problems or run tracing or diagnostics programs.
Fault masking	
Processor instruction retry Alternate processor recovery Processor-contained checkstop	If an instruction fails to execute due to a hardware or software fault, the system automatically retries the operation. If the failure persists, the operation is repeated on a different processor &, if this does not succeed, the failed processor is taken out of service (checkstopped). Only LPARs supported by failed processor are affected.
Dynamic processor sparing	Allows idle Capacity Upgrade on Demand (CUoD) processors to be automatically activated as replacements for failed processors.
Partition availability priority	In the event of a processor failure, maintains LPAR-based workloads based on assigned priorities; i.e., remaining processor capacity is assigned to the highest-priority workloads.
Memory sparing	Enables redundant memory to be activated in the event of failure.
Enhanced memory subsystem	Enables memory controller & cache sparing.
Enhanced cache recovery	Detects & purges processor & cache errors. Recovers original data.
Dynamic I/O line bit repair (eRepair)	Detects & bypasses failed memory pins.
PCI bus parity error retry	Retries an I/O operation if an error occurs.

Figure 8: *Key Power Systems Reliability, Availability and Serviceability Features*

Detailed Data

Installations

The cost comparisons presented in this report are for the profile installations summarized in figure 9.

Infor LX			
Specialty Vehicles Manufacturer	**Industrial & Automotive Components Manufacturer**	**Industrial Equipment Manufacturer**	**Cable & Wire Products Manufacturer**
Business Profile			
$400 million sales 1,700+ employees 5 manufacturing plants & engineering centers	$600 million sales 3,000 employees 7 manufacturing & distribution centers, technical/R&D center	$1.2 billion sales 7,000 employees 3 manufacturing plants	$3.5 billion sales 8,000 employees 35 manufacturing locations
Number of Users			
400	600	1,250	2,000+
CONFIGURATIONS & STAFFING			
IBM i 7.2 + Power Systems			
720 1/6 x 3.6 GHz 720 1/4 x 3.6 GHz IBM i 7.2 0.3 FTE	S814 1/8 x 3.6 GHz 720 1/6 x 3.6 GHz IBM i 7.2 0.4 FTE	S824 2/12 x 3.8 GHz 720 1/8 x 3.6 GHz IBM i 7.2 0.6 FTE	S824 2/24 x 3.6 GHz 740 2/16 x 3.6 GHz IBM i 7.2 0.75 FTE
Microsoft Windows Server 2012 + SQL Server 2014			
2 x 2/8 Hyper-V servers SQL Server 2014 AlwaysOn, Windows Server 2012 0.7 FTE	2 x 2/12 database servers 3 x 2/12 application & Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 0.9 FTE	2 x 2/16 database servers 6 x 2/12 application servers 2 x 1/4 Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 1.6 FTEs	2 x 4/24 database servers 10 x 2/12 application servers 2 x 2/8 Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 2.0 FTEs

Infor XA			
Pipes & Tubing Manufacturer	**Automotive Parts Manufacturer**	**Avionics Components & Subsystems Manufacturer**	**Industrial Packaging Equipment Manufacturer**
Business Profile			
$75 million sales 800 employees 2 manufacturing & distribution sites, 3 distribution centers	$350 million sales 2,500 employees 5 manufacturing & distribution facilities	$1 billion sales 7,000 employees 20+ manufacturing & development centers	$750 million 4,000 employees 6 manufacturing plants
Number of Users			
250	725	1,400	2,500
CONFIGURATIONS & STAFFING			
IBM i 7.2 + Power Systems			
720 1/6 x 3.6 GHz 720 1/4 x 3.6 GHz IBM i 7.2 0.25 FTE	S814 1/6 x 3 GHz 720 1/4 x 3.6 GHz IBM i 7.2 0.45 FTE	S824 2/12 x 3.8 GHz 720 1/8 x 3.6 GHz IBM i 7.2 0.75 FTE	2/24 x 3.8 GHz 730 2/12 x 4.2 GHz IBM i 7.2 0.9 FTE
Microsoft Windows Server 2012 + SQL Server 2014			
2 x 2/12 Hyper-V servers SQL Server 2014 AlwaysOn, Windows Server 2012 0.7 FTE	2 x 2/12 database servers 3 x 2/8 application & Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 1.15 FTEs	2 x 2/12 database servers 5 x 2/12 application servers 1/4 x Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 1.8 FTEs	2 x 4/24 database servers 8 x 2/12 application servers 1 x 1/4, 1 x 2/8 Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 2.45 FTEs

Figure 9: *Installations Summary*

Infor System21			
Specialty Paper Products Manufacturer	**Apparel & Accessories Manufacturer**	**Food & Beverage Manufacturer**	**Agribusiness Company**
Business Profile			
$70 million sales 200 employees Single manufacturing plant	$150 million sales 1,200 employees 3 manufacturing plants & 2 distribution centers	$400 million sales 1,500 employees 3 manufacturing plants & distribution centers	$750 million sales 5,000+ employees 5 manufacturing plants
Number of Users			
80	350	500	1,000+
CONFIGURATIONS & STAFFING			
IBM i 7.2 + Power Systems			
2 x 710 1/4 x 3.6 GHz IBM i 7.2 0.2 FTE	720 1/6 x 3.6 GHz 720 1/4 x 3.6 GHz IBM i 7.2 0.3 FTE	S814 1/6 x 3 GHz 720 1/4 x 3.6 GHz IBM i 7.2 0.4 FTE	S814 1/8 x 4.1 GHz 720 1/6 x 3.6 GHz IBM i 7.2 0.55 FTE
Microsoft Windows Server 2012 + SQL Server 2014			
2 x 2/12 Hyper-V SQL Server 2014 AlwaysOn, Windows Server 2012 0.5 FTE	2 x 2/8 database server 2 x 2/8 application & Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 0.65 FTE	2 x 2/12 database server 3 x 2/8 application & Web servers 1 x 1/4 Web server SQL Server 2014 AlwaysOn, Windows Server 2012 0.95 FTE	2 x 2/12 database servers 5 x 2/12 application servers 2 x 1/4 Web servers SQL Server 2014 AlwaysOn, Windows Server 2012 1.35 FTEs

Figure 9 (cont.): *Installations Summary*

Configurations and staffing levels are for core ERP systems, and do not include servers and personnel supporting other Infor10 applications or ION middleware.

IT Costs

IT costs were calculated as follows:

- *Hardware and maintenance, and software license and support* costs were calculated based on street prices; i.e., discounted vendor prices.

  IBM i software license and support costs were for numbers of used Power cores only, and allowed for IBM Solution Edition pricing for Infor ERP systems. Windows servers were configured with Windows Server 2012 and SQL Server 2014 AlwaysOn. Support costs were for Microsoft Software Assurance.

- *Personnel costs* were calculated based on annual average salaries of $89,745 for IBM i and Power Systems administrators; $93,986 for SQL Server and AlwaysOn administrators; and $77,820 for Windows system administrators. Salaries were increased by 56.7 percent to allow for benefits, bonuses and related items.

- *Facilities costs* include data center occupancy and energy consumption by IBM Power Systems and x86 servers, and for infrastructure equipment supporting these.

  Energy costs were calculated using vendor electricity consumption values. 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.

  Data center occupancy costs were calculated based on a conservative assumption of cost per square foot for existing facilities (i.e., costs do not include new construction).

Detailed breakdowns of IT costs are presented in figure 10.

Infor LX				
	Specialty Vehicles Manufacturer	Industrial Components Manufacturer	Industrial Equipment Manufacturer	Cable & Wire Products Manufacturer
IBM I 7.2/Power Systems				
Hardware & maintenance	10,215	24,620	45,821	87,866
Software licenses	112,245	129,696	188,625	377,250
Software support	97,478	114,419	163,455	326,910
Personnel	126,567	168,756	253,135	316,418
Facilities	11,255	16,784	22,986	26,439
Total ($)	**357,760**	**454,275**	**674,022**	**1,134,883**
Microsoft Windows/SQL Server				
Hardware & maintenance	47,706	49,238	81,715	160,043
Software licenses	106,368	153,611	196,379	325,958
Software support	79,776	115,208	147,284	244,468
Personnel	275,081	352,048	623,329	781,061
Facilities	11,648	29,414	54,506	87,681
Total ($)	**520,579**	**699,517**	**1,103,212**	**1,599,210**

Infor XA				
	Pipes & Tubing Manufacturer	Automotive Parts Manufacturer	Avionics Components & Subsystems Manufacturer	Industrial Packaging Equipment Manufacturer
IBM I 7.2/Power Systems				
Hardware & maintenance	13,523	21,852	45,864	59,129
Software licenses	112,245	112,245	188,625	335,805
Software support	97,478	97,478	163,455	291,837
Personnel	105,473	189,851	316,418	379,702
Facilities	12,434	15,409	22,760	24,541
Total ($)	**341,153**	**436,835**	**737,122**	**1,091,014**
Microsoft Windows/SQL Server				
Hardware & maintenance	27,650	45,938	49,622	133,196
Software licenses	100,544	158,022	151,272	338,642
Software support	75,408	118,517	113,454	253,982
Personnel	271,282	451,105	707,895	960,885
Facilities	11,931	29,414	42,185	71,262
Total ($)	**486,813**	**802,996**	**1,064,427**	**1,757,966**

Infor System21				
	Specialty Paper Products Manufacturer	Apparel & Accessories Manufacturer	Food & Beverage Manufacturer	Agribusiness Company
IBM I 7.2/Power Systems				
Hardware & maintenance	9,338	11,624	15,594	34,309
Software licenses	15,240	58,628	58,628	132,458
Software support	25,782	61,928	61,928	114,419
Personnel	84,378	126,567	168,756	232,040
Facilities	7,578	12,434	15,212	16,980
Total ($)	**142,316**	**271,181**	**320,117**	**530,205**
Microsoft Windows/SQL Server				
Hardware & maintenance	26,144	49,565	60,347	71,818
Software licenses	51,044	103,544	150,611	167,597
Software support	38,283	77,658	112,958	125,697
Personnel	198,115	256,790	370,339	528,071
Facilities	11,844	23,295	32,585	48,779
Total ($)	**325,429**	**510,850**	**726,840**	**941,962**

Figure 10: *IT Costs Breakdowns*

Costs of Downtime

Costs of downtime involved two sets of calculations:

1. *Supply chain disruption* consisted of costs incurred for operational and planning processes between initial supplier contacts and final customer delivery. This category generally corresponds to the *Source, Make and Deliver* segments of the Supply Chain Operations Reference (SCOR) model developed by the Supply Chain Council.

   Costs include idle and underutilized capacity and personnel; handling of delivery delays (including distribution center and transportation costs); additional inventory carrying costs; costs of procurement and logistics changes; effects of delays on production operations, including costs of scheduling and setup changes; and costs of customer billing and payments processing delays.

2. *Customer penalties and remedial costs* included late delivery and imperfect order fees, along with buyback costs such as additional discounts and rebates. Costs of potential customer loss and hedging were not included.

Cascading effects were also allowed for. Calculations were based on industry norms applied to the business profiles of individual companies. User survey input, financial documents such as annual and quarterly reports, and other sources were employed to quantify cost of downtime variables.

All IT and downtime costs 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).