




CA Resource Management Suites Address the Challenges of Optimizing the Online Mainframe



Billions of eBusiness and supply-chain dollars depend on highly available, high performance mainframe applications and subsystems. Millions of people access and analyze petabytes of structured data stored in mainframe-based databases. Yet datacenter experts still rely on older, home-grown tools to aggregate and deliver critical system data to key management applications. Without timely and accurate system information, datacenter experts are hard pressed to optimize capacity utilization of these critical systems, resolve system issues and improve system cost effectiveness. This paper discusses why home-grown solutions have difficulty scaling to meet the new management challenges of online, highly-connected mainframes and outlines how CA's Resource Management Suites address these challenges to deliver specific business and IT benefits.

Mainframe Innovation Drives New Datacenter Management Challenges

Enterprises have long valued the high performance, availability and scalability aspects of mainframe systems able to process billions of dollars worth of transactions and efficiently provide access to petabytes of structured data. As enterprises are finding new ways to leverage those aspects towards fuelling continued growth and increased competitiveness, their mainframe workloads have grown in scale and complexity at rates that are raising the eyebrows of even the most ardent mainframe advocates.

This growth is also creating new datacenter management challenges. As one operations manager at a community hospital put it, “As the volume and complexity of our mainframe workloads keeps growing, we need management tools that ensure our ability to keep pace—despite the resource constraints we face in IT.”

A key aspect of keeping pace with the growing mainframe workload volumes and complexity involves transforming the sea of raw data captured by the operating system's System Management Facility (SMF) into information usable by datacenter experts through their management solutions. Simply put, datacenter staff cannot effectively manage enterprise resources without rapid access to accurate management data about the mainframe environment. However, the growing scale and complexity of enterprise mainframe systems and workloads has dramatically:

- Changed the volume of raw SMF data captured,
- Increased data processing efficiency requirements,
- Shrunk information availability timeframes, and
- Increased difficulties in managing the accuracy and integrity of archived SMF data.

Let us examine how these changes came about and the impact they have on datacenter staff's ability to manage mainframe resources to cost effectively attain business goals.

Workload volumes outpace in-house developed data processing solutions

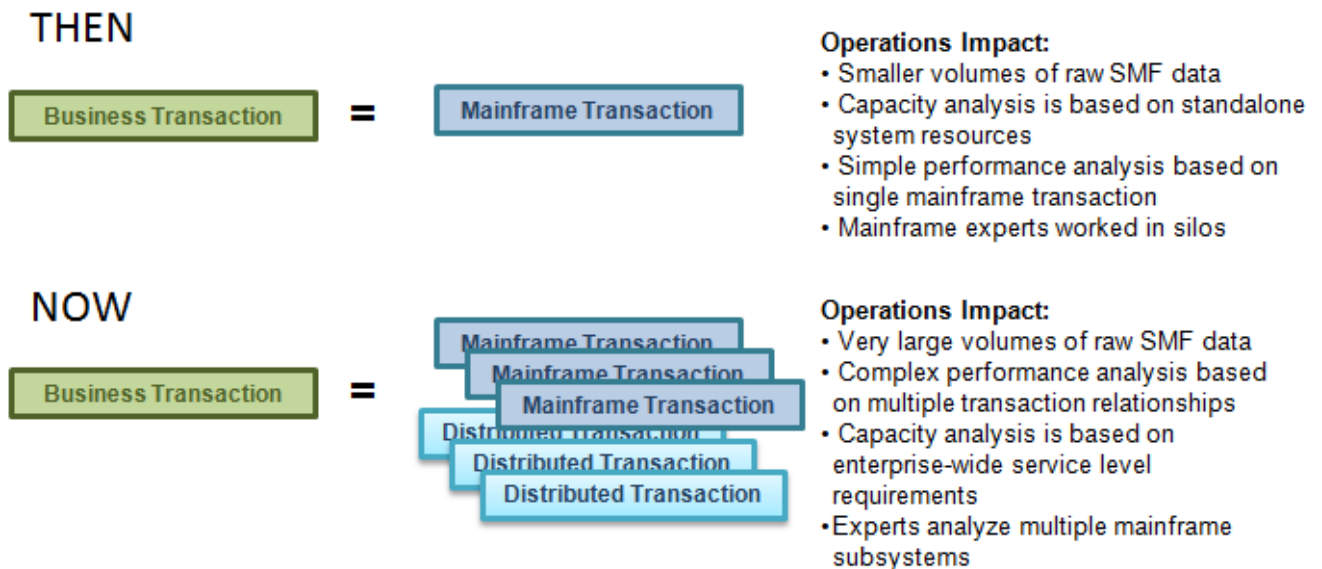
Datacenter consolidation is a key driver of mainframe workload increases. Global enterprises that once operated several smaller datacenters, each with a relatively small mainframe, are increasingly consolidating workloads into larger systems in fewer locations. Corporate acquisitions and mergers have had the same effect. For example, a large financial institution consolidated an additional 80 million per-day business transactions onto their mainframe environment over the course of several corporate acquisitions.

Web-based and SOA-based applications have also increased mainframe workloads dramatically. One factor contributing to this growth is that web-based applications have tiered architectures, where distributed servers interact with mainframe transactional and database systems. Every new web-based application deployed on distributed application servers generates additional mainframe workloads. For example, one institution saw its mainframe transaction volumes grow at 30%-40% annually as every year brought deployments of new web-based applications. Another example is

Sol Meliá, Spain’s largest hotel chain owner, which saw online sales increase by more than 200 percent within a year of launching an online application targeting new customers.¹

Another contributing factor is that business transactions are becoming more complex. As recently as ten years ago, a business transaction was a fairly simple entity often encapsulated as a single mainframe transaction. A 1% increase in business transactions, therefore, would result in a 1% increase in mainframe transactions. With SOA-based applications, a single business transaction is often supported by a collection of orchestrated SOA services, where each SOA service can involve one or many mainframe transactions (Figure 1). Therefore a 1% increase in business transactions, supported by multiple mainframe SOA services, could easily result in a tenfold increase in mainframe workload. As a result, many large organizations are seeing their mainframe transaction volumes grow faster than their business transaction volumes.

Figure 1: Operational impact of innovative business use of mainframe assets



Additionally, this simple example does not factor in the reuse of SOA services. As more mainframe capabilities are exposed as SOA services, or through other standards-based database protocols, modern developers will leverage these capabilities within their current development platforms and tools. As a result, Ptak, Noel & Associates (PNA) believes that enterprises are only at the beginning of SOA-related workload increases in mainframe environments.

This dramatic increase in workload volume in turn generates similar dramatic increases in SMF data.² The SMF data growth means that management applications which analyze this data must

¹ “Spanish hotel chain boosts web sales by 200 percent with flexible online IT solution,” CA Customer Success Story.
² Indeed, IBM had to completely redesign its system logging capabilities a few years ago to keep pace with increasingly large volumes of data generated by larger mainframe installations that are processing more workloads as global enterprises consolidate their datacenter operations.

sift through millions of log records to find the critical subset of data needed to solve problems and conduct decision support analysis. In-house developed data processing solutions that were effective when mainframe systems were smaller and simpler typically have problems scaling to current mainframe usage. Datacenter staff should migrate to solutions designed to process the much larger volumes of raw SMF data into information usable by mainframe management applications.

Problem resolution necessitates a new approach to managing SMF records

The need for higher levels of end-to-end transaction availability and performance grows as more mainframe workloads are directly driven by customer and partner transactions. This puts enormous pressure on datacenter staff to rapidly solve availability and performance problems. Yet, many enterprises actually are experiencing longer problem resolution times as their home-grown SMF data extraction processes struggle to keep pace with data volumes. The extraction overhead may produce negligible delays in providing data to management applications when processing a million SMF records. But when that extraction process is applied against a hundred million SMF records, the overhead may create serious delays which prevent staff from starting any problem resolution activities. "There is nothing more frustrating than having business managers upset with you about a service problem that is costing the company millions of dollars and there is nothing you can do about it because you are waiting for the system to finish nightly processing before you can get started," said a systems analyst at a large financial institution.

One cause of the rising data availability problems during troubleshooting is the increasing number of individuals requiring access to SMF data. For example, to solve a problem with a single business transaction may require a team of database, networking, and application experts to perform root cause analysis, each using specific diagnostic tools that analyze varying subsets of SMF records. This scenario results in three simultaneous ad-hoc requests to extract different records from an enormous pool of raw SMF data. The increasing volume of SMF data produced by the multiple mainframe transactions supporting that single business transaction exacerbates the problem as millions of SMF records must be sifted through over and over again to extract the relevant data for each request. This time intensive, highly error prone data extraction process produces a backlog of work which must be completed *before* problem resolution analysis can begin. Therefore, this data processing backlog hampers datacenter staff's ability to resolve service problems quickly or prevent performance slowdowns.

24x7 business and shrinking staff size require higher data processing efficiencies

Traditionally, the data management housekeeping needed to create a reliable, complete, non-duplicated SMF data source was carried out as a batch job during maintenance windows or during lulls in CPU and I/O utilization. It was more common to have spare CPU and I/O capacity to complete data management tasks when mainframes were standalone enterprise systems. However online business models that generate business transactions around the clock, coupled with enterprise globalization that consolidated multiple smaller mainframes into larger more complex systems, have dramatically shrunk maintenance windows and created high demand for CPU and I/O resources at all times. Spare resources have become a thing of the past. Thus SMF

data processing tasks must operate at the same time as business transactions. This co-existence dictates SMF-related jobs be more resource efficient than ever before.

Even as mainframe environments grow in complexity and size, the pool of internal mainframe expertise continues to shrink. This combination of trends makes staff productivity increases extremely important as enterprises often rely on a smaller, centralized staff to monitor and manage multiple datacenters. This condensed staff requires high efficiency aggregation of SMF data from multiple systems across multiple locations.

Regulatory compliance places a spotlight on data integrity

Global enterprises must comply with a growing number of country-specific and industry-specific regulations, even as existing regulations undergo near constant change. This has led to an increasing demand for processes to manage datacenter logs centrally, audit datacenter controls once, and report results in multiple, diverse ways. Most governance risk and compliance solutions focus on consolidating the number of auditing and reporting aspects.³ However, most enterprises still struggle to centrally manage the accuracy and integrity of mainframe logs. Manually-created log aggregation and archiving schemes are stressed by the size and complexity of current mainframe environments and the dramatic increases in SMF data volumes. In addition, many enterprises are struggling to adapt these schemes to work with IBM's redesigned log streaming capabilities. The resulting data overlap and accidental data elimination threaten data integrity at a time when international focus on corporate compliance is highest.

Meeting These Challenges with CA Resource Management Suites

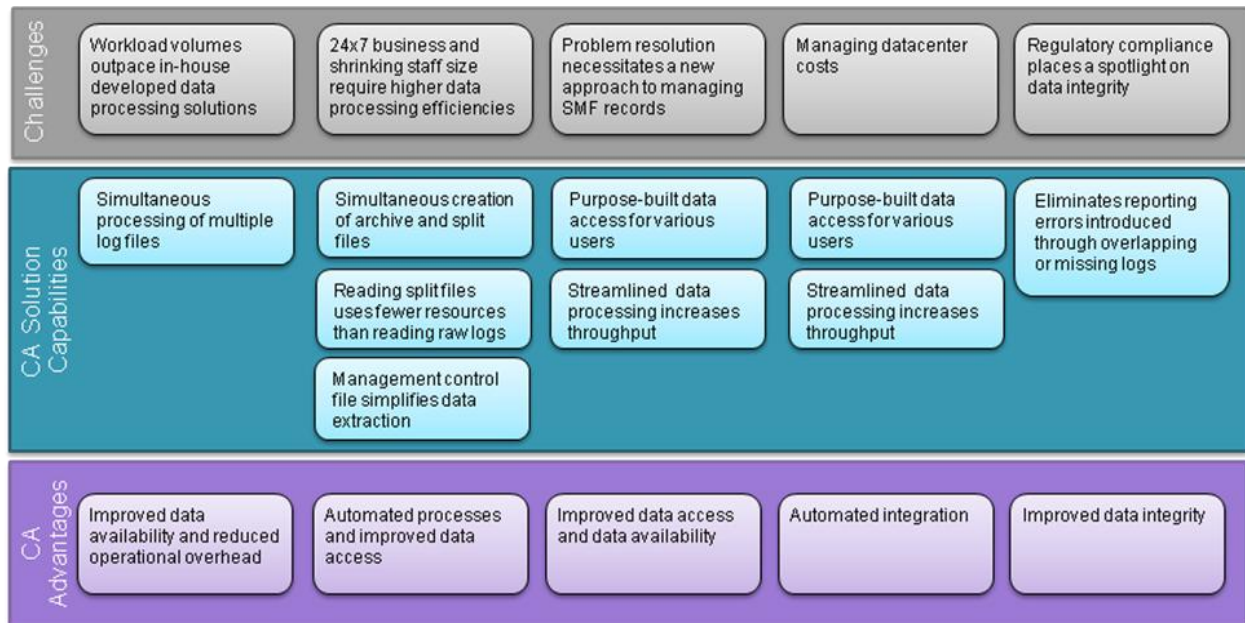
These new operational realities demand new approaches and new solutions to providing reliable, complete and non-duplicate SMF data. Home-grown and disparate SMF data management solutions, coupled with the shrinking pool of mainframe expertise, are no longer equipped to cope with demands created by mainframe systems supporting innovative business services and processes. The CA Resource Management Suites, comprised of CA SMF Director and either CA MICS[®] Resource Management or CA JARS[®] Resource Accounting, were created specifically to address these new datacenter management challenges (Figure 2).

The CA Resource Management Suites address the demand for increased operational efficiencies in several ways. CA SMF Director processes multiple raw log files from multiple systems simultaneously, rather than serially. This eliminates potential bottlenecks and achieves throughput levels that can keep pace with multiple LPARs and hefty workloads of the largest data centers. Similarly, the solutions are architected to read raw log files once and create both the historical archive and the application-specific split files simultaneously, rather than serially. By increasing the throughput of SMF processing and reducing the number of reads of raw data, CA SMF Director significantly reduces the operational overhead associated with providing accurate and accessible management data. Enterprises can redeploy the saved system resources to more mission-critical and profitable workloads. The suites also streamline capacity analysis considerably, allowing

³ "CA: Innovation in Governance, Risk and Compliance," Ptak, Noel & Associates, Oct 2007.

CA MICS® Resource Management or CA JARS® Resource Accounting experts to run multiple models with wider ranges of system and business variables in shorter timeframes.

Figure 2: Mapping how CA's solutions address new challenges in mainframe management



CA's solution suites organize vast amounts of SMF data into multiple split files⁴ for application-specific purposes. This capability makes data subsets rapidly accessible to a wide variety of stakeholders without requiring repeated extractions from raw SMF data. For example, capacity management tools like CA MICS or CA JARS can access a split file with resource utilization and cost records while other management tools, such as CA Performance Management Suite or CA Mainframe Database Performance Management Suite, can access split files containing data related to DB2, CICS or IBM WebSphere applications. By splitting SMF data more effectively, CA's solutions empower the technical teamwork required to manage and troubleshoot increasingly complex mainframe environments supporting business transactions. CA's suite packaging also addresses the data accessibility challenge by providing many data analyzer integration options, specified by the customer's infrastructure profile, as part of the basic solution suite. These options speed the integration with existing management tools used by various stakeholders.

The CA Resource Management Suites use standardized, automated processes to catalog all SMF data from multiple systems and subsystems, create multiple split files, and maintain a management control file to simplify data extraction. One benefit of these standardized, automated processes is productivity improvements across all mainframe management activities by allowing datacenter professionals to apply their technical expertise in more strategic ways. For example, capacity

⁴ Split files are a container for a specific subset of SMF data. A split control statement is used to identify the SMF records that should be written to each split file. Records can be selected based on system ID, subsystem name, application ID, record types, record subtypes or within a specific time period.

planning requires significant technical knowledge but the real business value is obtained when that skill is applied to interpreting the business implications of system utilization, service performance and user experience information. However, most capacity planners must focus their technical knowledge to determine how to obtain and aggregate utilization, performance and experience information from raw SMF data. They then spend much of their time either waiting for various home-grown solutions to complete their ad-hoc requests against raw SMF data or verifying the accuracy of the aggregated data. This use of talent and time is hardly a recipe for quick, decisive, and effective decision-making.

The standardized, automated processes embedded within CA's solution suites eliminate much of the time and effort taken to collect and extract the information. Datacenter experts can then focus on using the suites' analysis and interpretation capabilities to provide the insight required for better decision making in shorter timeframes.

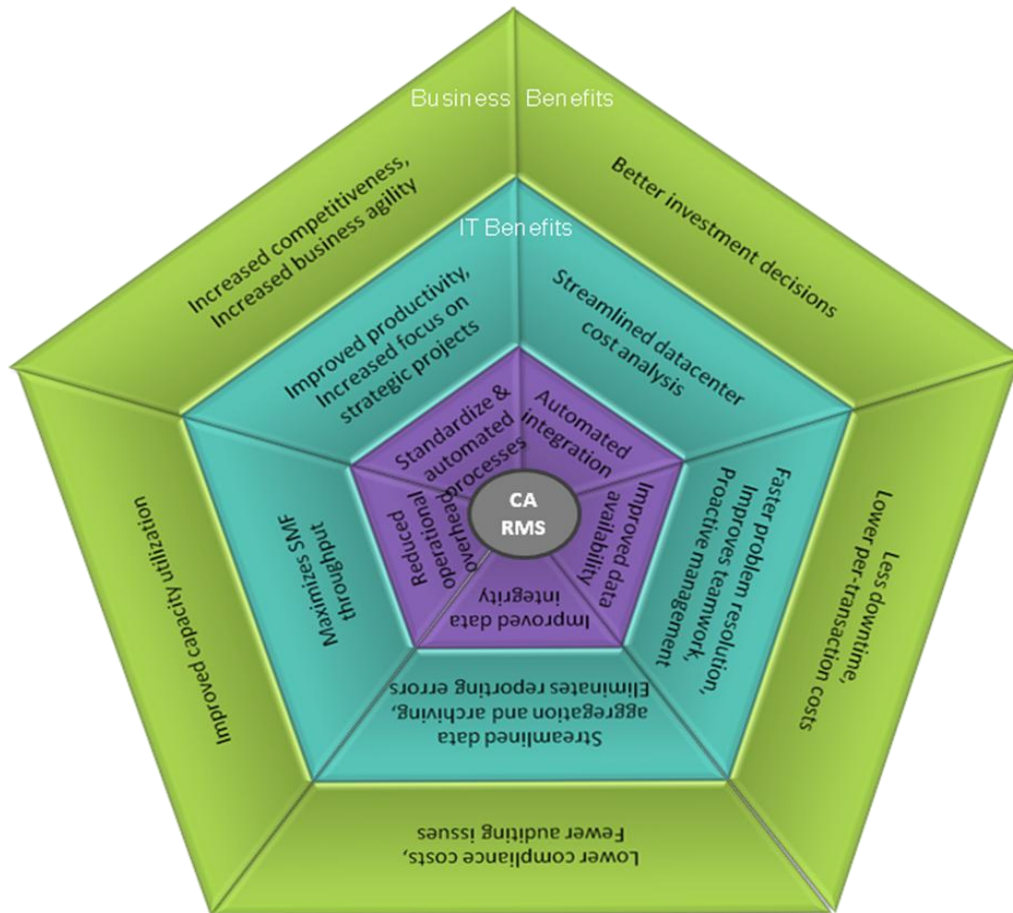
CA's solution suites also ensure data integrity, even with the inherent difficulties in aggregating SMF data coming from multiple systems and subsystems possibly in different data centers across multiple geographies and time zones. For example, dynamic date parameters are difficult to maintain manually and can produce overlapping and duplicate logs. Similarly, using multiple archiving utilities can introduce duplicate data or eliminate data coming from multiple sources. The CA Resource Management Suites streamline and standardize the aggregation and archiving processes to capture all SMF data seamlessly. This creates a non-duplicated, non-overlapping data store regardless of the recording method⁵, which eliminates reporting errors often introduced through overlapping or missing logs.

Understanding the Business Value of Resource Management

There are several business benefits that 'trickle-up' from the use of CA Resource Management Suites. The health of IT organizations and the productivity of datacenter staff have a growing impact on business agility and profitability as online business models and services play an increasingly important role in overall enterprise strategies. IT-business alignment is really about clarifying and quantifying the inherent links between IT management solutions and business benefits (Figure 3).

⁵ The solution can work with MAN datasets, System Logger feeds and mixed environments.

Figure 3: Mapping the benefits of CA's Resource Management Suites



This visibility helps IT organizations provide both a clearer understanding of the business unit costs and more effective advice on right-sizing operational investments. For example, the profitability of a business model can only be determined if both the per-transaction income and costs are well known. With online business models, quickly executed analysis SMF data provides knowledge about resource relationships, usage and cost, which provides insight into how to increase transaction completion rates at lower operational costs. Both top and bottom line improvement possibilities become visible when SMF data is automatically integrated with these analysis capabilities.

From a top-line perspective, the business impact of problem resolution processes streamlined by CA Resource Management Suites is often directly measurable in online sales situations where application availability and performance directly impact revenue. Average revenue per hour, employee productivity, and business transaction throughput can be used to estimate the immediate financial returns of CA's suites.

Similarly, strategic increases in business competitiveness and agility often hinge on making better decision making in shorter timeframes. CA's standard processes for managing SMF data facilitate

these decision making procedures. For example, changes in business forecasts significantly change system capacity needs. Capacity planners now have the luxury of carrying out a more complete analysis for multiple economic scenarios in shorter timeframes. Business managers can now have access to more information on which to create contingency plans, allowing more flexibility in dealing with rapid shifts in the market or economy.

From a bottom-line perspective, CA's suites reduce the operational overhead of processing the increasingly large volumes of management data. This minimizes the impact on mission-critical business transactions, enabling enterprises to utilize their mainframe capacity more effectively. The time saved and improved teamwork of an enterprise's technical organization can have a far-reaching impact on total operational costs as well. The efficiencies gained by the IT organization provide more time for staff to develop proactive and automated problem management strategies to handle technical issues before they grow into service-level problems and empower existing staff to efficiently manage more complex business transactions and increasingly diverse business services and priorities. Simply put, proactive datacenter management can dramatically lower per-transaction costs.

The integrity of archived SMF data is of increasing concern for compliance purposes. Global enterprises must comply with a growing number of country-specific regulations. The regulations are changing, which places an enormous burden on datacenter staff to generate accurate reports that support the various control proof-points that auditors require. By eliminating the risks of overlapping and missing log data during the aggregation and archiving process, CA dramatically lowers the cost of compliance and the risk of report errors which may lead to failed compliance audits.

Final Word

CA's Resource Management Suites provide the means for mainframe experts to understand their increasingly complex mainframe environments and to adapt those environments to how enterprises are innovating and integrating their mainframes. The solutions' architecture addresses some fundamental changes, including the vast increases in SMF volumes, the complexities of accurately aggregating and consolidating data from multiple disparate sources, and improving data availability and accessibility by diverse datacenter experts.

Tackling these fundamental issues now is critical because PNA believes datacenter complexity will only increase as more organizations become more sophisticated in their use of SOA-powered business process modeling and complex event processing solutions to facilitate business agility. Additionally, the need to efficiently integrate SMF data with multiple enterprise solutions will only grow as more enterprises incorporate system data sources to fuel business intelligence analysis of online services.

Equally important, the solutions' architecture allows enterprises to build upon and leverage their existing management investments. For example, the included data analyzer integration options allow customers to leverage CA's efficient data processing across existing management investments. CA MICS and CA JARS users can use familiar capabilities to forecast resource

demand in relation to both system and business variables, but will have dramatic productivity gains as SMF data is available faster and ad-hoc reports are easier to generate and have faster completion rates because of CA SMF Director.

It is these productivity gains which impress CA's current clients. As one senior systems specialist commented, "With CA, we have been very successful in automating our SMF management tasks, allowing us to deliver a level of service and data availability that was previously impossible. The solution manages the data, and we don't have to worry how it is kept or which 'Tape' or what 'File' it is on to be able to produce reports. We can now easily and automatically deliver SMF data that is specific to our customer's application and ad-hoc needs, enabling us to better meet all of their requirements."



This paper was sponsored by:

CA, Inc.

One CA Plaza

Islandia, NY 11749

1 800 225-5224

www.ca.com

No part of this document can be used in any medium without the written permission of Ptak, Noel & Associates LLC.

All trademarks, trade names, service marks and logos referenced herein belong to their respective companies.

While every care has been taken during the preparation of this document to ensure accurate information, the publishers cannot accept responsibility for any errors or omissions. PNA makes no assumptions as to the potential return on investment that other organizations will receive. PNA opinions reflect judgment at the time and are subject to change.

About Ptak, Noel & Associates LLC

With a belief that business success and IT success are inseparable, Ptak, Noel & Associates works with clients to identify, understand and respond to the implications of today's trends and innovations on the future of IT Operations.

www.ptaknoelassociates.com

September 2008