IBM Smarter Planet Storage Solutions

John Webster
February, 2011

Technology Insight Series

Evaluator Group

enabling you to make the best storage decisions

Executive Summary

IBM's "Let's Build a Smarter Planet" TV ads portray a vision of interconnected systems producing actionable information in real time. The ad campaign is very much in keeping with core values IBM has honed for decades.

In the beginning IBM manufactured computers for companies that, initially, didn't know what to do with them. So IBM had to first show them how to compute. IBM had to sell computing first before it could sell computers. Smarter Planet builds on that legacy to draw attention to a fundamental way to use computing in a business context for the 21st century—one that encompasses a myriad of sensory devices and compute nodes all working within an "orchestrated" system of systems.

It's a big picture message that IBM is fundamentally delivering in these ads: systems that harvest data from a variety of wired and wireless sources are capable of producing new types of information and solving some of humanity's most important problems. And while these may feel futuristic even for seasoned IT professionals, the technologies needed to turn the vision into reality (RFID, GPS, pattern recognition, Complex Event Processing, and so on) are here and now. Pervasive Analytics is one term we have seen that has been used to describe the integration of pervasive data sources with existing business intelligence systems to produce useable information for a wide range of decision makers in real or near-real time[1].

Here we examine a fundamental question for storage professionals: what possible role does storage play on a "smarter planet?" In this paper, we show that pervasive data-generating technologies are now available, that emerging applications are harnessing them, and that there is a desire among business executives across a broad range of industry segments to mine these applications for competitive advantage in some cases, and for the well-being of humanity in others. But, where does storage fit and what roles do storage technologies play?

How real is Smarter Planet?

In a book by Chris Stakutis and John Webster entitled "Inescapable Data – Harnessing the Power of Convergence," published in 2005, the coauthors speak of developing new information sources – possibly in real time – from the convergence of networked systems with pervasive, digital sensory and wireless technologies including radio frequency identification (RFID), global positioning system (GPS), charge-coupled device (CCD) imaging, cellular telephony, and others. In 2004, most senior executives interviewed for the book said they were aware of and understood the underlying technologies. Some were moving ahead with initial projects. Now, there are many examples of pervasive information systems built upon the successes of initial deployments.

Some examples from five years ago and today include:

[1] From a paper entitled "The Analytics Revolution: Optimizing Reporting and analytics to Make Actionable Intelligence Pervasive," by David Loshin, Knowledge Integrity, Inc.

Healthcare

2004

The CEO of a major metropolitan teaching hospital described an ongoing pilot project that used RFID in an operating room setting. Plastic patient bracelets contained RFID tags that were detected when the patient was wheeled into the OR. The clinical system in the OR would immediately identify the patient, verify the scheduled procedure, provide the clinicians with a checklist of equipment and medications that needed to be in the OR at the beginning of the procedure, and have the patient's medical history available for immediate recall. RFID was also being used in another project where medication containers were tagged so as to verify that a drug being administered to a patient was in fact one that had been prescribed. The CEO reported that the reduction in the hospital's insurance rates would more than pay for these projects in the first year.

2010

Today healthcare providers are investing in a broad range of sensory devices and interconnected systems to improve patient outcomes and do so cost effectively. With the median age of the US population rising, the number of people at risk of chronic illness is the highest ever—a contributing factor to rising healthcare costs. The healthcare industry is now beginning to use a variety of sensory devices to empower self-care, reduce the number of episodes that result in doctor visits and hospitalization, and improve a patient's sense of well-being and morale[2]. These devices communicate with a healthcare network, often via a home PC or laptop, or sometimes via a robot, to provide a continuous communications link between patient and healthcare provider. Devices include:

- Wearable, wireless sensors and wireless-enabled stationary devices (wireless-connected bathroom scales for example) that communicate with a laptop or robotic device in a home environment that report blood pressure, weight, respiration, pulse, and other vital signs, daily and sometimes hourly.
- Mobile and stationary robotic devices that can query patients every day confirm that medications are taken when prescribed for example, and monitor the current state of a patient's condition.
- The wireless ePen – a pen used by the patient to complete a daily checklist. The results from check-off items (meals, exercise routines, emotional state, pain intensity, etc.) are transmitted to a local data collection station (PC, laptop) connected to the healthcare network.

These devices and others are the endpoints of a network of systems that provides:

- Patient data aggregation points communicating with centralized healthcare systems
- Patient tracking, reporting, and alerting for clinicians and healthcare professionals in real time if required
- Bidirectional communication between patients and providers
- Electronic patient records and historical patient data

[2] Healthcare 2.0, a presentation made by Chris Stakutis at the Spring 2010 Storage Networking World conference, Orlando ,FL.

As a result, patient outcomes can be dramatically improved simply by keeping them out of hospitals. Practitioners can also analyze the resulting data across a broad patient population to see the effectiveness of one treatment protocol compared with another.

Retail

2004

The CEO of a large metropolitan retail chain described a project that used wireless digital video cameras to track the how shoppers moved around on a retail floor. The system's primary application was to track customer movements and count the number of customers as they entered and left a particular sales area. Using this data, the company could better understand in-store traffic patterns and optimize product placement, optimize staffing levels throughout the day, and measure the impact of advertising and special promotions.

2010

A business analytics system at Rooms to Go, an innovative nationwide furniture retailer, uncovers market trends and customer buying behaviors to improve up-sell and cross-sell results. The system from IBM integrates several regional transactional databases into a single enterprise data warehouse that delivers comprehensive sales reports to help staff identify buying patterns, among other things. The system increases sales by helping staff optimize floor planning, and reduces advertising costs by helping staff target promotions to the most likely buyers. The system also helps the retailer's buyers make inventory purchase decisions based on relationships among the data rather than gut feel.

A Small Window on 2015

As pervasive information systems move forward, a new compute paradigm will be established—one that leverages and places maximum value in the delivery of information in real time. It will not be enough to know what has happened in the past or what may happen based on past events. Both organizations and individuals will want to know what's happening now, and to have the ability to make predictions based on both what is happening now and what just took place a moment ago. The more the mechanics of real time information delivery are understood, the more we will all want to use it.

Additionally, corporate executives will want to use these systems to eliminate inefficiency and waste to the greatest extent possible. This will come as the result of the system integration work now being driven by the emergence of cloud computing.

The technologies that will foster real-time information delivery and increased efficiency are:

1. The continued delivery of systems performance increases driven by advances we can already see in silicon chip processing and the possible commercialization of nano technology

2. Systems integration on a scale that is massive when compared to what is going on today. The word we use to describe this integration effort is "cloud." We may well be calling it something different five years from now.

By 2015 we could well see the following:

In healthcare, the use of electronic patient records among healthcare providers is commonplace. Massive system integration within the healthcare industry will give researchers the ability to perform therapeutic studies in a fraction of the time now required. Advances in genomics will have reduced the cost of recording an individual human genome to an affordable level. Treatment therapies will be individualized and structured around a patient's genome.

A growing list of retailers will likely offer patrons the ability to customize and streamline the shopping experience. Customers can opt-in to using a retailers system that understands their preferences, knows their shopping histories, makes specially discounted offers based on preferences and previous selections, knows what they have selected as they move around a retail space, and automatically deducts funds from their accounts as they bypass the lines at the checkout counter. Customer can be offered these incentives because of the underlying efficiencies that these new systems enable.

Most major metropolitan cities will have public safety systems similar the ones currently in use in Chicago and London. In addition, these public safety systems will be integrated with law enforcement systems at the state and federal government level. In spite of the fact that more that 50% of the US population will live in a city, by 2015 cities will actually be safer places to live than suburban or rural areas.

IBM's Smarter Planet and Storage

For IT administrators that will be tasked with resolving storage architectural and management issues that could arise, the question is: what impact will these projects and the technologies deployed to support them have on the storage environment, both now and in the future?

One way to address this question is to first determine what attributes storage must present to these converged data applications and then match these attributes to actual storage devices and systems. We believe this approach is smarter for buyers than starting with a particular technology in mind and then attempting to fit it to various application requirements. Regardless of the approach, the discussion quickly gets to technology. Here, we use IBM's storage product portfolio as a reference, since it supports the important performance and scalability requirements and is used successfully in several Smarter Planet deployments.

Pervasive information systems and applications are likely to require the following attributes of the storage infrastructure:

Scalability

Depending on the application and the desired platform, IBM offers storage system scalability in multiple dimensions. IBM's DS8700 takes a more traditional scale-up storage controller-based approach. It scales to 2 PB of tiered storage, and up to 32 internal controller processors to maximize I/O performance as the number of external, server-facing I/O ports increases on the front-end and disk storage capacity is added on the back-end. IBM's XIV system is a virtualized storage array that achieves scalability through the addition of nodes within a grid framework, which adds more I/O performance and fault resilience capabilities as capacity grows.

Given the sheer size of the data volumes in play for some applications, accessing data using block-based protocols is not efficient. Storage subsystems based on larger constructs than blocks such as data objects the size of many blocks may be a better match for these "big data" applications. IBM's Scale-out NAS (SONAS) can answer this requirement. Like XIV, SONAS also features a grid-based architecture, but divides the workload between interface nodes for I/O processing and storage nodes for managing access to disk. Both interface and storage nodes can be scaled independently. SONAS supports up to 7.2 PB of raw storage using 1 TB drives, and will scale to greater capacities as larger capacity drives are introduced and incorporated.

Performance

The first RFID reference architecture proposed in 2003 prescribed a computational system supported by a real-time, in memory event database (RIED). Rotating disk is considered to be too slow for these types of applications. However, server-based memory is expensive and for some situations—a database that is too large to fit into memory for example—flash-based solid-state disk (SSD) can outperform the highest performance rotating disk by a wide margin. SSD is now generally available as a discrete storage tier within a high-performance disk system.

Storage systems are emerging that offer differing levels of performance to match the demands of specific applications. Today, most storage systems integrate flash memory based SSD as a high-performance tier but, like server-based memory, SSD can be an expensive alternative.

Currently, IBM offers two storage systems-based approaches to SSD optimization. IBM Easy Tier automatically migrates data based on the performance requirements of the application environment as sensed by Easy Tier software residing in the storage system, and works at the sub-LUN, or sub-volume level. With IBM's SAN Volume Controller (SVC), configurations that support SSD and are used in conjunction with IBM's Tivoli Storage Productivity Center (TPC) to help IT administrators identify "hot spots" within the disk systems attached to SVC and migrate data to SSD based on recommendations made by TPC. TPC further assists this process by generating scripts to automate non-disruptive data movement under SVC.

Change Data Propagation using Data Copy Functions

Copies of data are often required. For example, a copy of a production database used with OLTP systems may be needed for use in analytic processes. Storage-based data copy functions can be used here. IBM offers a number of these including:

> SVC and DS8700 FlashCopy Copy-on-Write – This function creates the appearance of a copy of data but improves efficiency by copying only data that is changed, usually a fraction of the total data volume involved.

XIV Snapshots – A virtually unlimited number of point-in-time "pictures" of XIV-based volumes can be created using minimal additional storage for each "picture".

SVC FlashCopy to different devices – A copy of data can be created on a different disk system enabling storage used for a copy to be matched to the application using the copy. For example, analytic processes may have very different performance requirements from OLTP applications

Efficiency

It is precisely because pervasive data sources can generate huge data volumes that storage infrastructures need to be far more efficient than they are now. There are two ways to measure efficiency in the context of Smarter Planet: management practices that reduce complexity and maximize administrative effectiveness, and storage devices that optimize internal efficiency.

Disk storage virtualization, offered by IBM SVC and XIV, groups available storage into 'pools' that helps address both measures: configuration of storage is less complex and, because applications share storage in a pool, storage use is also better optimized.

SVC and XIV support "thin provisioning", which automatically allocates disk space to applications only as needed. Thin provisioning eliminates the need to pre-allocate disk for future growth and enables acquisition of additional disk space as growth occurs.

SVC also makes it easier to pool together new and existing storage to support Smarter Planet applications. Configuration and utilization of storage are improved, acquisition of new or replacement storage can be reduced or delayed, and migration to new storage can be done without disruption to applications.

Process improvements can also drive efficiency. IBM Storage Optimization and Integration Services uses best practices and software to improve storage strategy, architecture, reporting, and operations.

Internal efficiency for backup storage targets can be optimized by implementing data deduplication processes that eliminate redundancies among stored data objects (volumes, files, and blocks). IBM ProtecTIER offers high throughput deduplication for both mainframe and open systems environments. Tivoli Storage Manager includes client and server deduplication for disk storage pools.

Data protection

Smarter Planet applications can be particularly stressful on existing data protection processes. Data volumes get larger given the sheer amounts of data being generated while traditional backup windows get smaller and may in fact disappear. In addition, many of these applications are generating relatively large data volumes at distributed sites (RFID-enabled retail outlets for example). Storage administrators have traditionally had difficulty protecting data at remote sites. Pervasive data applications only exacerbate the problem. The situation calls for very new ways to approach data protection.

Two very different technologies—mirroring for continuous change data capture and data deduplication, can enable the efficient protection of these data volumes. IBM Metro Mirror and Global Mirror, and TSM FastBack technologies continuously capture only changed data, supporting near continuous operations for systems that can't be down for as long as it takes to perform a traditional recovery.

Data deduplication removes redundancies in the data to further reduce the amount of data managed. When a backup is needed for recovery, it is "rehydrated" automatically upon restore to preserve the original content. Data deduplication is able to dramatically decrease the amount of disk space required for backup data when disk is used as a backup target, while retaining the significant performance improvements that disk based backup devices have over tape.

Archiving, retention, and compliance

While data archiving solutions may seem to be the farthest thing away from delivering the new types of information envisioned by Smarter Planet, they in fact are critical as well. It is not uncommon to find organizations where "save everything forever" is the default policy for digital archiving. For Smarter Planet applications, this practice is irrational and unsustainable. Needed here are functions, best practices, and automated processes that make data retention manageable and can potentially turn archival data into a meaningful source of information for operational business intelligence applications. These include:

Active archiving software for email, file systems, SAP and other key applications that works by using policies to identify older or less-frequently used information and redirecting data to a more efficient tiered storage environment.

Automated Tiering, which moves data from primary storage to archive storage, is implemented in IBM Information Archive and Tivoli Storage Manager. Archive storage tiers can include high capacity SATA disk, tape, non-erasable non-rewritable storage for compliance, and cascaded storage arrays that no longer meet primary storage requirements.

Tape File Systems add new archival use cases for tape technology. Tape File Systems, such as IBM's Long Term File System for LTO-5, allow users to access data on tape as if it were on disk, so data and the applications that use it can stay in synch over time.

IBM Information Lifecycle Management Services uses best practices for data classification, tiered storage implementation, content management, archive and retention to help customers better manage information based on its business value.

Acquisition Alternatives

When budgeting for Smarter Planet projects, we believe that it is important to consider alternatives to outright purchase such as equipment leasing, managed services and OPEX-based project financing. Why? These projects may require a significant commitment of both capital and human resources, but typically reduce operating expenses sufficiently to justify the investment. Applying off-balance sheet financing alternatives to IT projects allows IT administrators to move projects forward as an operational expense when the CFO wants to conserve capital. We also note that IT provisioning in general is moving in the direction of new services delivery models as exemplified by the cloud computing structures that are now emerging. Services are paid for as they are consumed, usually on a monthly basis. Payment for the underlying infrastructure can also be made on a monthly basis and IBM Global Finance can be used to structure the acquisition alternative for Smarter Planet projects.

Conclusion

Traditional data warehousing and business intelligence applications typically provide enterprises with an analysis of the past because they are built upon data that represents past activity. And while these applications have served and continue to serve decision makers, there no doubt exists a desire to integrate and combine wired and wireless data sources into the data analytics-based decision making process such that actionable information is delivered in as close to real time as possible. We believe this to be a fundamental objective of IBM's smarter planet.

For customers to succeed with new implementations of pervasive information systems, we have very briefly outlined what IBM storage technologies and services may be needed in the context of the new data analytics models now emerging and how IBM delivers them. As one can see, some overlapping capability does exist at a general level. For instance, the decision to use SSD for scale-up application acceleration or virtualization to avoid hot spots will depend on a number of factors, not the least of which is what fits best with the existing environment. What IBM demonstrates is an ability to bring these capabilities to bear on Smarter Planet projects for a wide range of use cases and customer deployment scenarios.