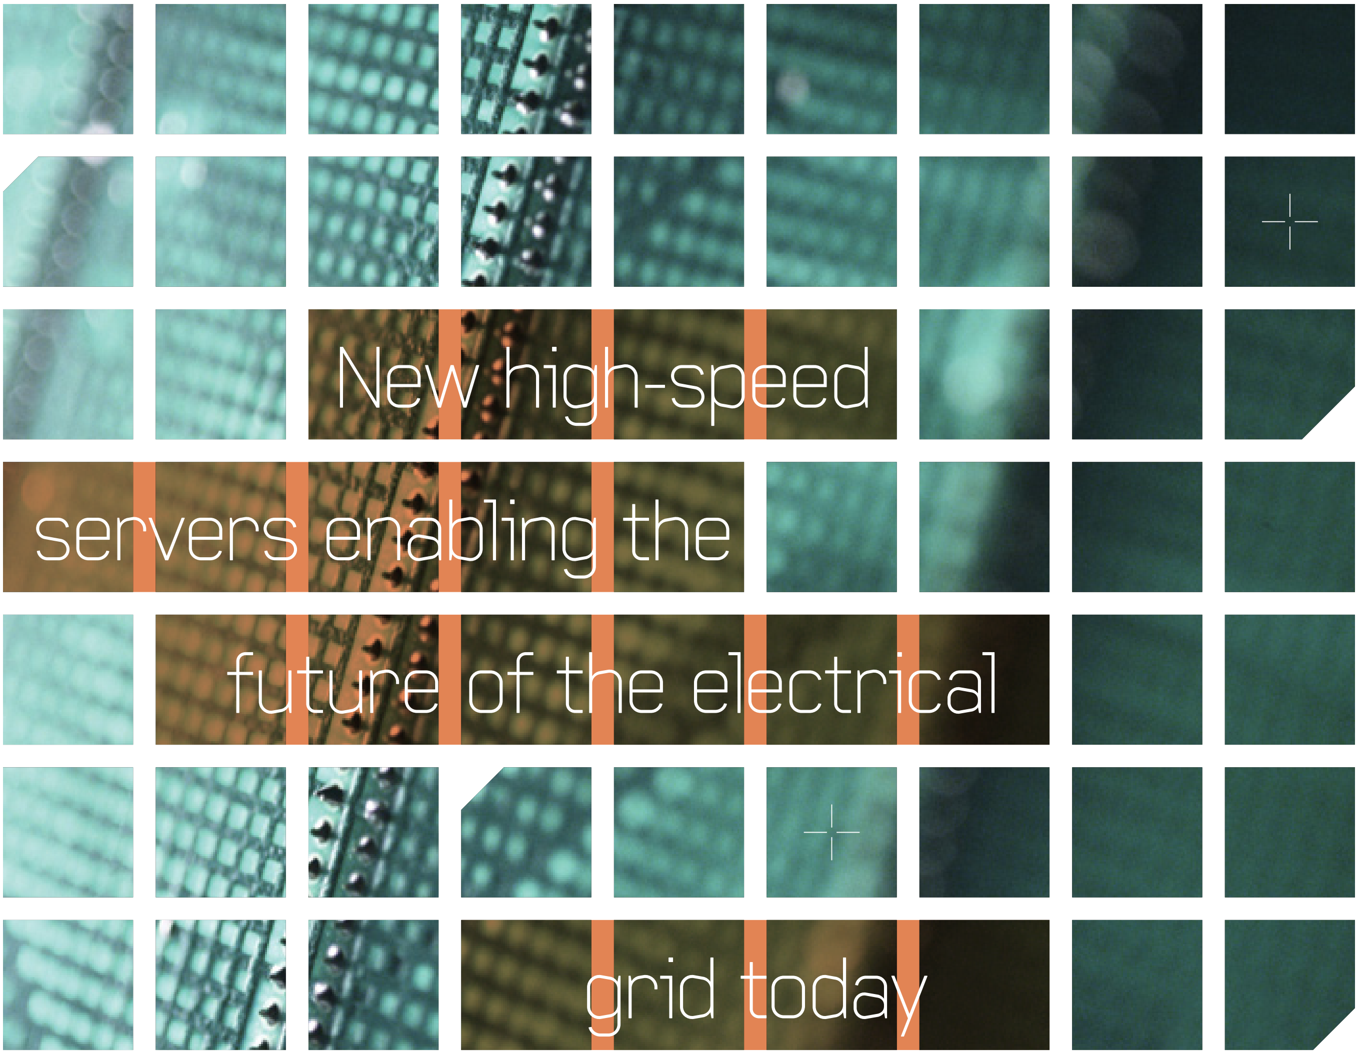

New high-speed servers enabling the future of the electrical grid today

The worldwide utility industry has seen the future and it is high-speed, integrated, flexible and intelligent. The distribution of electricity is a technological marvel that now is taken for granted in the developed world, and is rapidly being built out around the globe. Sometimes it's hard to believe that the entire technology is only a little more than 100 years old and was essentially invented by one man, Thomas Alva Edison. Edison not only invented the light bulb, he designed and developed the fundamental generation, transmission and distribution systems that made ubiquitous electricity possible.

What most people don't realize is that the fundamental principles and practices of electric generation and distribution have not changed dramatically since Edison invented them. However, with the advent of high-speed electrically powered computers only a little more than 50 years ago, this overlaid digital technology opened up new vistas for development and refinement of Edison's invention.

Today, as computer systems become faster, more efficient in their use of electricity, and more intelligent software is developed, the two systems are merging. Eventually they will become one—a high speed electric power system that is generated, distributed, transmitted, repaired and managed almost entirely by high-speed computers that will adjust to weather and generation changes, react to outages much more rapidly than human beings can, and make electric distribution essentially self-healing.

As computer systems become married with robotics, which are becoming more sophisticated daily, human intervention in the electric distribution system gradually will evolve to only top-level decision-making. Such a system actually is in experimental development in several places, particularly in DV2010, a project led by WE Energies, Milwaukee, WI, and involving a consortium of utilities and vendors in the Midwest. Numerous other advances are under way elsewhere.

The main thing the electric system of the future will require is ever faster and more capable processors, computers, and servers. The distribution system of today produces staggering amounts of data—a lot of which is not even yet being used, but will be in the future. By their very nature, utilities and the newer business entities that are involved in the restructured utility environment worldwide, generate and use tremendous quantities of data. This is particularly true on the "operations" side where the actual process of delivering electricity, natural gas and water to homes and businesses occurs. For more than 20 years now, utilities have used SCADA (Supervisory Control and Data Acquisition) and DA (Distribution Automation) to monitor and control the millions upon millions of miles of wire and pipes necessary to deliver energy and water.

In recent years, regional transmission organizations (RTOs) and independent system operators (ISOs) have been created in the United States and similar organizations elsewhere to regulate the distribution of power, natural gas and water and enable the development of wholesale and retail markets. These newer organizations, which control distribution over much larger areas, require even more massive amounts of data to perform their functions.

Since electricity flows in a continuous process from generation through transmission to distribution and consumption (unless, of course, there is an outage), the only way of counting what is used actually is at the retail level—businesses and homes that use it to perform work and have meters to determine how much they consume. In order to ensure that each of the players in that chain (generators, transmission firms or divisions, and distributors) are paid for their services, it is necessary to measure consumption at the retail level. Since there will be relatively few generators, transmission firms and distributors in a particular distribution chain, but many retail points of consumption, it is necessary to track and measure what is consumed and assign it to the proper entities back up the chain.

Combining these measurements and assigning them properly is called "aggregation," and is one of the primary responsibilities of the RTOs and ISOs. However, all of them don't do this aggregation themselves. Some rely upon transmission/distribution organizations to aggregate and provide the ISO with the compiled numbers, putting increasing data handling demands on utilities themselves.

Aggregation is an extremely complex computing problem involving massive amounts of data. As an example, in Texas the ERCOT does between three and nine "resettlements" per day, processing about a terabyte of data each time, or up to nine terabytes of data each day.

Local utilities generate large quantities of data themselves with SCADA, DA and the newer EMS (energy management systems) constantly generating automated reports on the status of the network including voltages at hundreds or even thousands of points, heating of wires, operational status and maintenance alerts on different capacitors, transformers, switches, etc.

FIGURE 1

Moore's Law

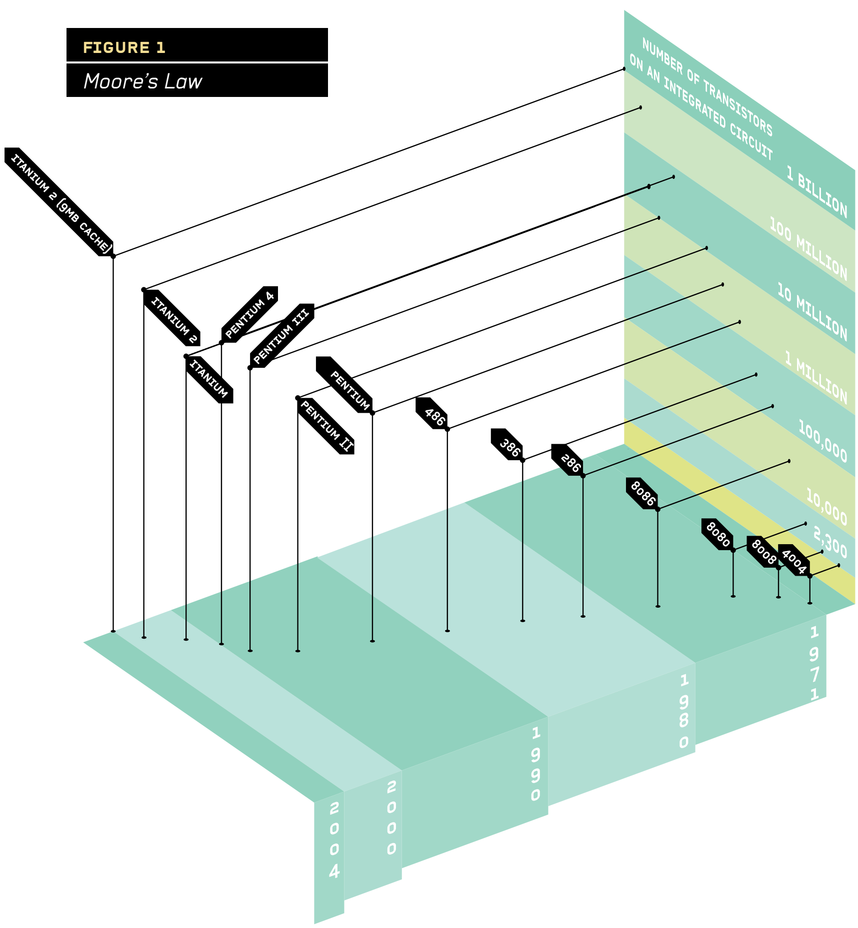

THE NEED FOR SPEED

All of this data must be collected, analyzed, reacted to, and stored. That requires tremendous computing power with as much processing speed as possible.

Over the past 20 years, processing speed and power has generally kept up with the exponential increase of data, although the increase has been taxing existing processors and servers. The normal evolution of computing power, generally is expressed in an axiom known as Moore's Law. Moore's Law is the empirical observation that the transistor density of integrated circuits, with respect to minimum component cost, doubles every 24 months. It is attributed to Gordon

E. Moore, a co-founder of Intel, Santa Clara, CA. Moore's law is illustrated in *Figure 1*

There also is the capability to place processors in tandem, which Intel now is doing in many of its processors with two processors residing on the same chip. Recently, Intel announced that by November 2006 it will begin releasing "quad-core" chips that will contain four processors in dual-core configurations.

The density of transistors in processors is not the only determinant of processing speed and power, however. There also are the issue of architecture and how the processors are addressed by compilers and operating systems. The traditional architecture is known as RISC (reduced instruction set computing), which has been around for about 20 years.

Because Hewlett-Packard Co. (HP), Palo Alto, CA, servers are in use at all but one of the ISOs in the United States, are used by most of the major EMS, SCADA and DMS software vendors, and in a large proportion of utility control centers (90% of all transmission and distribution control centers in China), HP servers have to be very fast. Now, they are becoming even faster.

EPIC and the Evolution of the Itanium Processor

For nearly 20 years, HP servers have been the mainstay of utility control rooms and operations departments, hosting the data-intensive SCADA, DA, Outage Management and DMS (Distribution Management Systems). These were primarily "Alpha" servers running HP's Unix operating system (Tru64 Unix),.or OpenVMS. Newton-Evans Research, a respected automation market research firm that specializes in "real-time" utility systems estimates that more than 60% of the world's energy management systems run on HP hardware.

Alpha servers were based upon processors designed and built by Digital Equipment Corp, which was acquired by Compaq in 1997. Then in 2001, HP acquired Compaq.

The Alpha servers were true 64-bit machines using Digital's proprietary chip. That chip was available only from HP. However, as the need for speed and the processing of even more massive quantities of data, particularly at RTOs and ISOs continued to escalate, HP began working with its energy management solution partners to move to the "next generation" of processors. Since manufacturing chips containing processors is expensive and time-consuming, HP developed a partnership in 1997 with Intel, the world's leading chip-maker. This partnership led to the development of the Itanium processor, now known as the Itanium2.

In addition to offering dual-core processing, the Itanium2 also makes use of more sophisticated compilers and a new-generation architecture known as EPIC (explicitly parallel instruction computing). This new architecture enables the processor to "look ahead" and "learn" as it handles data. The goal of the partnership, which included several other vendors besides HP, was to develop a high-end, mainframe-class computing platform with PC volume economics HP's implementation of the Itanium2 processor is the Integrity line of servers which HP now is making available to its clients.

As the utility industry continues to evolve, operators of electric power Generation, Transmission and Distribution systems need to deploy more modern enterprise systems to provide the scalability, adaptability, security and compliance that the energy business demands and will increasingly need in the future. HP's Integrity servers, using Intel™ Dual-Core Itanium2 micro-architecture, are designed to meet these comprehensive needs. And HP and its software partners can ensure a clear migration path, transition services, and protection of one of the utility industry's most important and increasingly scarce investments—people.

The utility industry is facing an increasingly serious threat as a result of the aging of its workforce. Some utilities estimate that more than 50% of their existing workforce will be eligible for retirement within the next five to 10

FIGURE 2

IOUs Ranking of Major Issues

(5-point scale of importance)

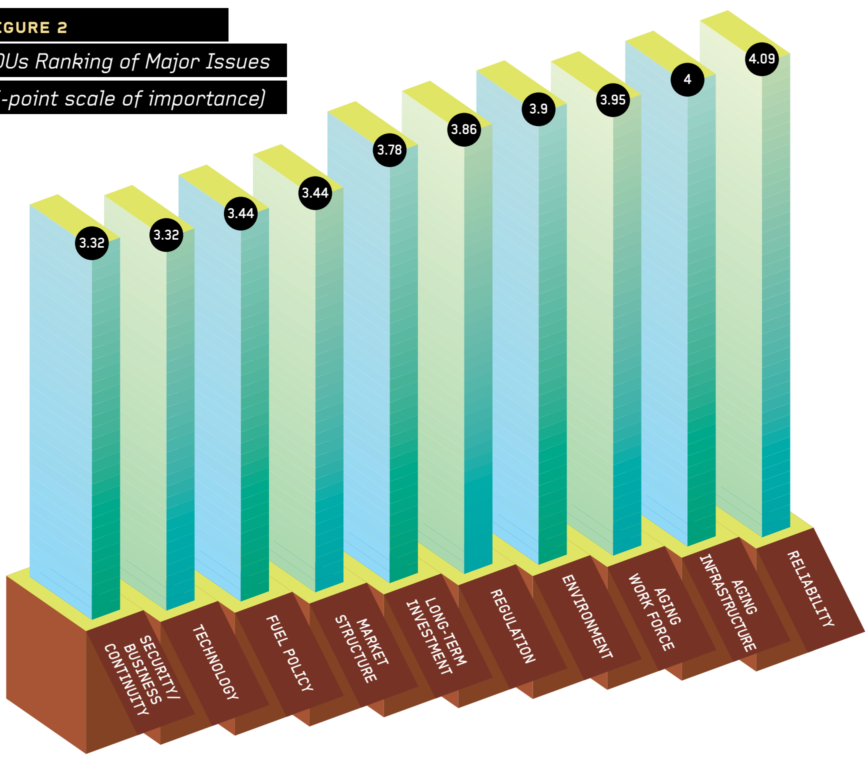

From the Annual Black & Veatch Strategic Directions Electric Utility Industry Survey (used with permission) years. Utility industry executives and managers are aware of the impending problem, as shown in ***Figure 2*** from a recent Sierra Energy Group survey.

The aging workforce is a major driver of the movement toward more intelligent, self-managing and self-healing systems not only because it represents the best forward-thinking, it also deals with a current and impending problem.

Revolutionary improvements in automation such as the self healing grid will require breakthroughs in processing performance. Fortunately development is underway on new architectures such as the Intel®Tera-scale Computing Research Program-Intel's overarching effort to shape the future of Intel processors and platforms.

According to a recent Intel company white paper:

> *"Intel researchers are already working on over 100 R&D projects worldwide to address the hardware and software challenges of building and programming systems with dozens of energy-efficient cores with a sophisticated memory hierarchy. Currently, projects in this program span circuit technologies, micro-architecture, interconnects, memory, and software technologies.*
>
> *Some may wonder why Intel is putting so much emphasis on tera-scale research. The reason is simple: tera-scale computing is a twofold revolution, both in the capabilities that devices will have, and in the amount of innovation that will be required to handle tomorrow's advanced applications. The term itself—tera-scale—refers to the terabytes of data that must be handled by platforms capable of teraflops of computing performance. That's a thousand times more computing capability than is available in today's giga-scale devices.*
>
> *Why such a leap forward? Because incremental improvements in performance and capabilities simply won't support real-time data mining across teraflops of data (ie. Required for self healing a wide area power system); artificial intelligence (AI) for smarter cars and appliances; virtual reality (VR) for modeling, visualization, physics simulation, and medical training; and other applications that are still on the edge of being science fiction."*

High-speed servers that can handle the massive quantities of data being generated today, and enable the development of the control systems of the future, are coming on line in the HP Integrity servers. Additional developments including the placing of even more processors on chips to enable even faster and higher-quantity data already obviously are on the drawing board at HP and Intel. As utility executives and managers look toward the future, they already are envisioning systems that will require this next generation of computing power and beyond.

Because HP provides the hardware, but not the software, to operate EMS, DA, and other high-volume utility systems, the servers also must provide a clear and easy migration path for the vendors that do provide the software. These vendors include virtually all, and certainly a who's who of major software vendors: Areva, ABB, Itron, Advanced Control Systems, Telvent, Cognicase, GE Smallworld, SPL WorldGroup, etc.

The HP Integrity tools integrated with the servers simplify deployment, system tuning and infrastructure management to make the servers quicker to deploy and easier to configure and manage. Some of these tools include:

→ *The Intel Itanium2 processors for faster throughput, more concurrent users and lower power consumption;*

→ *Exceptional system reliability, availability and future-readiness;*

→ *Multi-level HP-UX 11i v2 security features including intrusion detection, host and network security tools, centralized authorization with single log-in and security containment;*

→ *Tools to simplify and streamline infrastructure management;*

→ *Lower total cost of ownership; reduced capital expenditures and predictable maintenance costs;*

→ *Server virtualization options for consolidation of selected Network Manager modules*

Handling today's needs, tomorrow's vision

Thus having faster, more capable servers represents more than just meeting the need for speed. Such servers also help meet other current critical utility industry concerns, such as doing more with fewer, less-experienced people and helping meet the vital needs for security and rapid expansion of computing power and capabilities. And they open the possibilities now being envisioned at utilities for an ever more "intelligent" electrical system that require ever less human intervention to continue to operate efficiently and effectively.

Edison's intent was to make life easier. The electrical system he invented has done that. Now computer systems overlaid on that electrical system will make the operation of that system easier, faster and less labor-intensive. Thomas Edison would be proud! ***-Warren Causey***