

Smart Systems Manifesto

How networked computing
and the physical world are
merging to transform
civilization



Smart
Systems
Design

Harbor
Research

The bit, the byte, and later the packet made possible the entire enterprise of digital computing and global networking. Until the world agreed upon these basic concepts, it was not possible to move forward. The next great step in Information and communications technologies—completely fluid information and fully interoperating Machines and devices—requires an equally simple, flexible, and universal abstraction schema that will make information itself truly portable in both physical and VIRTUAL space, and among any conceivable information devices and systems. For the internet of things to really take off, Radical new thinking about information technology must begin at the most basic levels, with new conceptions about how devices, information, people and systems will interact.

TABLE OF CONTENTS

01: Introduction	Page 04
02: Designing The Future of Information	Page 06
03: What's Required?	Page 08
04: Enabling Universal Connectivity	Page 09
05: Automating Design and Development	Page 11
06: Achieving Truly Flexible, Scalable Systems	Page 13
07: Leveraging Collective Intelligence	Page 15
08: Smart Systems Impact On Business Models	Page 16
09: How Should We View Future Architecture?	Page 17
10: Smart Systems Is A Radical Departure	Page 19
11: Getting There First; But To Where?	Page 20

M2M & IT TECHNOLOGY ONLY TELL PART OF THE STORY

01

Introduction

For quite a few years now, Harbor Research has focused most of its research and consulting on what we call “Smart Systems”—the convergence of pervasive or embedded computing with the packet-switching “network of networks” called the Internet.

These days, many people refer to this phenomenon as “the Internet of Things.” We prefer “Smart Systems” over other terms in common use—notably “M2M,” which usually stands for “machine-to-machine”—because it captures the profound enormity of the phenomenon - something much greater in scope than just machine connectivity.

Whatever we chose to call it -- “Smart Systems” or “Pervasive Computing” or “The Internet of Things” — we are referring to digital microprocessors and sensors embedded in everyday objects. But even this makes too many assumptions about what the smart systems phenomenon will be. Encoded information in physical objects is also smart—even without intrinsic computing ability. Seen in this way, a printed bar code, a house key, or even the pages of a technical manual can have the status of an “information device” on a network. For that matter, all of these characterizations do not even begin to address the human-machine dimension of collaboration.

But very few people are thinking about smart connected systems on that level. Current IT and telecom technologists are operating with outdated models of data, networking and information management that were conceived in the mainframe and client-server eras and cannot serve the needs of a truly connected world. “Smart Systems” should automatically be understood as “real-time networked information and computation,” but it isn’t.

We have now entered the age when everyday objects will communicate with, and control, other objects over a global data network—24/7/365. It’s not “the future,” it’s now—this year, next year—and thus it is vitally important that business leaders understand this phenomenon, its effects on their business, and what they should do right now to position themselves for opportunities that are literally just around the corner.

Before delving into the new thinking that makes this story possible, let’s talk about why it’s necessary at all. The IT and telecom sectors have failed to re-evaluate their relationship to advancing technology and to their constituents. The business and technology paradigms to which these industries cling today are far too limiting, too cumbersome and too expensive to foster and sustain new growth.

From a Telco perspective, today’s discussions of M2M systems focus almost exclusively on communications -- the “pipe” -- and very little on the information value. In other

WHY DO WE NEED A MANIFESTO?

words, on things that look good to the carriers. There are many popular visions about wireless monitoring and wireless control. Such as it is, wireless is a fantastic new advance -- no question. But, focusing on the communication element alone as 'first-order' business value amounts to grabbing the wrong end of the technology stick. Wireless communications alone steals the limelight and potentially eclipses the real revolution -- utilizing new networking technologies and processes to liberate information from sensors and intelligent devices to leverage collective awareness and intelligence.

From an IT perspective, today's corporate IT function is a direct descendent of the company mainframe, and works on the same "batched computing" model—an archival model, yielding a historian's perspective. Information about events is collected, stored, queried, analyzed, and reported upon. But all after the fact.

That's a very different thing from feeding the real-time inputs of billions of tiny "state machines" into systems that continually compare machine-state to sets of rules and then do something on that basis. In short, for connected devices to mean anything in business, the prevailing corporate IT model has to change.

The next cycle of technology and systems development in the smart connected systems arena is supposed to be setting the stage for a multi-year wave of growth based on the convergence of innovations in software architectures; back-room data center operations; wireless and broadband communications; and smaller, more powerful client devices connected to personal, local and wide-area networks. But is it?

This paper is about an important new perspective on smart systems and services. This perspective does not just come from our own thinking. It is from the diverse group of people who are thinking about the scope and on the scale that Smart Systems deserves. It is because this community of technologists and business developers is so fragmented that we believe a new approach is required. This is why we believe a new "manifesto" is required to begin to integrate and bind these disparate views of the future of information technology.

What do we plan to address in this Smart Systems Manifesto?

the required architecture and technologies to inform a radically new view of information services; and,

the corresponding emerging business models these technologies will inform.

Most importantly, we plan to demonstrate that for the first time in the evolution of networked businesses, these two "architectures" must be viewed in close proximity. The two thrusts need to be mutually supportive without inhibiting one or the other.

CAN THE INTERNET OF THINGS BE DESIGNED?

However, trying to coordinate and leverage the respective roles of technology architecture and business architecture often creates contention. Many of the participants in this emerging arena we speak with are coming to see the continuously evolving relationship between these two dimensions as fertile ground for innovation. They need to be interwoven and mutually supportive. In fact, from our own direct consulting experiences, we believe success in either increasingly goes to the company that effectively utilizes the combined potential of both.

In our years of work on the Internet of Things phenomenon and its real-world effects on business, we have not encountered very many compelling visions about the complete integration of things, people, systems and real-time real-world inputs. The world today, for the most part, lacks a group of coordinated innovators that understand that the tools we are working with to make products “smart” on networks were not designed to handle the scope and diversity of interactions they are being enabled to accomplish.

Since the beginning of computing there have essentially been three generations of technology and architecture: mainframe computing, personal computing, and network computing. Each generation of technology has had significant impacts on productivity and efficiencies: mainframes standardized transactions; personal computing placed processing power into the hands of professionals; and, networked systems enabled business process automation.

What is important about the next generation of smart systems is the combined impact of these legacy cycles of innovation. While there is standalone value in each discipline - embedded software, servers, network infrastructure, and client devices - it is the combination of these innovations that will inform smarter systems.

“Smart Connected Systems” really means the future of information, and that means the future of civilization. It will require a remarkably agile global network that could comfortably scale to trillions of nodes—some of them hardware, some software, some purely data, many of them coming into and out of existence or changing location constantly. Obviously, such a network system cannot be “designed” in any ordinary sense. Certainly, it cannot be designed “top-down.”

And yet the Internet of Things must be designed in some sense. Such a network will easily be the biggest technical achievement in the history of humanity. Its closest predecessor is the global financial economy—with which, in fact, it will share vital characteristics.

Some basic design principles must be put in place to guide the growth of a vast, distributed technological organism that must remain organized as it evolves

02

Designing the Future of Information

SIMPLE INCREMENTAL IMPROVEMENTS IN TODAY'S TECHNOLOGIES WILL NOT WORK FOR THE INTERNET OF THINGS

according to a logic all its own. It demands that we design not only devices and networks but also information itself in ways not addressed by current IT.

The Internet of Interactions -- between and among "Things" and "People" -- requires much more than simple incremental improvements in today's technologies to be fully realized. The challenge is much more than a simple patch, Band-Aid, or new flavor of what we already do.

What's required is a true shift in thinking about how devices, people and physical systems will be integrated and how they will interact. We need an approach that is not about leveraging aging IT technology into a new application context; its about looking forward to a single, unified architecture for the nearly infinite interactions to which any PERSON or any THING can contribute.

What will this require?

Its about the ability to create more universal standards for connectivity and integration; one simple standard, based upon the Internet Protocol itself for networking any information device (even non-electronic ones), rather than the multiplicity of so-called standards we have today.

A vision that looks toward the creation of more universal means to integrate and manage disparate data, to which anyone can contribute, and which liberates information by abandoning traditional relational databasing and the client-server computing models that have us so deceptively "trapped" today.

These broader requirements are driving several important development trends, including:

The rapidly rising requirements for smart products to be interoperable with a growing array of applications, systems and people.

End customers are beginning to recognize the requirement for a unified family of software and tools to develop smart connected devices; today's diverse software offerings and fragmented supplier community will not meet their needs.

The need for new development protocols to address the growing complexity of new devices and systems.

The ability to aggregate, fuse and manage structured, unstructured and time series data.

The maturation of radically new information models and architectures that can more readily integrate classical IT capabilities with real-time, state-based devices and systems.

THE PRACTICES OF THE MAINFRAME AND PC ERAS ARE NOW OLD AND NOT SUITABLE FOR THE PERVASIVE COMPUTING ERA

Taking these initiatives seriously does not mean junking all current IT practice in one fell swoop. The pillars of present-day information technology will not crumble overnight, nor has the great existing investment in them suddenly lost all value. There are reasonable, fiscally sane paths for migrating to the future. But migrate we must.

The assumptions and practices of the mainframe and PC eras are now decades old and not suitable for the pervasive computing era that informs the internet of things and people. The nature and behavior of a truly distributed global information system are concerns that have yet to really take center stage—not only in business communities, but in most technology communities, too.

When it comes to preparing for the global information economy of the 21st century, most people assume that “the IT and telco technologists are taking care of it.” They take it on faith that the best possible designs for the future of connected things, people, systems and information will emerge from large corporations and centralized authorities. But those are big, unfounded assumptions. In fact, most of today’s entrenched players are showing little appetite for radical departures from current practice. Yet current practice will not serve the needs of a genuinely connected world. What are the major obstacles that need to be overcome?

Automated Development: When telephones first came into existence, all calls were routed through switchboards and had to be connected by a live operator. It was long ago forecast that if telephone traffic continued to grow in this way, soon everybody in the world would have to be a switchboard operator. Of course that has not happened, because automation was built into the systems to handle common tasks like connecting calls. We are quickly approaching analogous circumstances with the proliferation of smart connected devices. Each new device requires too much customization and maintenance just to perform the same basic tasks. We must develop software and methods to automate development and facilitate re-use, or risk constraining the growth of this market.

Flexible, Scalable Systems: IT professionals rarely talk these days about the need for ever-evolving information services that can be made available anywhere, anytime, for any kind of information. Instead, they talk about web services, enterprise apps and now cloud computing. The Web stores information in one of two basic ways: utterly unstructured, or far too rigidly structured. The unstructured way gives us typical static Web pages, blog postings, etc., in which the basic unit of information is large, free-form, and lacking any fundamental identity. The overly structured way involves the use of relational database tables that impose rigid, pre-ordained schemas on stored information. These schemas, designed by database administrators in advance, are not at all agile or easily extensible. Making even trivial changes to these schemas

03

What’s Required

WE DON'T NEED MANY STANDARDS; WE WANT ONE

is a cumbersome, expensive process that affects all the data inside them. Just as importantly, they make deep, inflexible assumptions about the meaning and context of the data they store. Both of these approaches to data-structure enforce severe limitations on the functions you want most in a global, pervasive-era information system: scalability, interoperability and seamless integration of real-time or event-driven data. The client-server model underlying the Web greatly compounds the problem.

Optimizing All Assets - Tangible And Intangible: New software technologies and applications need to help organizations address the key challenge of optimizing the value of their balance sheets, allowing them to move beyond just financial assets and liabilities to their physical assets and liabilities (like electric grids or hospitals) and then to their intangible assets and liabilities (like a skilled workforce). The task of optimizing the value of financial assets, physical assets and people assets requires new technologies that will integrate diverse asset information in unprecedented ways to solve more complex business problems.

Leveraging Collective Intelligence: For all its sophistication, many of today's M2M systems are a direct descendent of the traditional cellular telephony model where each device acts in a "hub and spoke" mode. The inability of today's popular enterprise systems to interoperate and perform well with distributed heterogeneous device environments is a significant obstacle. The many "nodes" of a network may not be very "smart" in themselves, but if they are networked in a way that allows them to connect effortlessly and interoperate seamlessly, they begin to give rise to complex, system-wide behavior. This allows an entirely new order of intelligence to emerge from the system as a whole—an intelligence that could not have been predicted by looking at any of the nodes individually. What's required is to shift the focus from simple device monitoring to a model where device data is aggregated into new applications to achieve true systems intelligence.

Smart Systems holds the potential for what Harbor Research calls "the networking of every manufactured thing." but it also presupposes the existence of "a zero-infrastructure, ad-hoc network" that makes seamless physical connection possible.

How will the physical aspects of this networking actually take place? How will virtually any device, right down to a lowly light bulb, become a peer that connects at will to the global data network?

Obviously, billions of devices of wildly varying types cannot each receive individual attention and configuration, or conform to elaborate a priori specifications. If it

04

Enabling Universal Connectivity

WE ARE AT THE CUSP OF A PERFECT STORM

literally takes a network engineer to screw in a smart light bulb, pervasive computing is never going to work.

Many schemes and so-called “standards” for device connectivity already exist. But of course, all those “solutions” add up to one big problem. We don’t want many standards; we want one.

Of course, an elegant answer is offered by IP which is predicated upon the idea that we already have an excellent way to get every conceivable device on the Internet—the original design principles and protocols of the Internet itself.

“Internet” is short for “internetworking.” The Internet was designed in the 1960s to allow the incompatible data networks and computing systems of the time to share information—to “talk to each other,” as people like to say. The Internet is literally a “network of networks.” That’s what the “inter” means.

The public Internet as we know it today is a worldwide embodiment of those original data communications protocols—which are, by design, extremely simple. For this discussion, their key attribute is that they make very few assumptions about the data they are sending and about the devices connecting to the network to send and receive data.

It is this extensible, technology-neutral basis of the Internet that has allowed it to scale so dramatically (and gracefully) since its inception, with minimal central administration.

There have been previous high profile examples of industries trying to create niche standards for special purpose networks only to eventually evolve to open IP standards years later. This circumstance is typified by the evolution of fieldbus networks in the industrial arena in the 1990s and early part of this decade. Several larger players slowed down their own developments, plowed billions of dollars into supporting closed and then ‘pseudo open’ standards and fought over the spoils for nearly 20 years only to find the market eventually came to understand the benefits of open IP based network standards. We see no reason that this lesson from history shouldn’t be applied now.

Looking forward, Harbor Research believes there is a major opportunity for select market participants to act as an ‘IoT architect or catalyst’ and harmonize the network standards for short range WLAN and WSN applications around an existing, new or adapted suite of interoperable standards.

SOME BASIC DESIGN PRINCIPLES MUST BE PUT IN PLACE TO GUIDE THE DEVELOPMENT OF SMART DEVICES

05

Automating Design and Development

Our society is at the cusp of a “perfect storm” of network connectivity. This phenomenon is not just about the dichotomy between people communicating with people or machines communicating with machines: it also includes people communicating with machines (e.g. a networked ATM), and machines communicating with people (e.g. automated stock ticker alerts on your smartphone). The concept of network effects states that the value of a network grows exponentially with the number of nodes connected to it. Along with the value, however, so too grows the complexity of the software integrated into devices and the reliance on people and organizations developing and supporting these systems.

With the rapid growth of wireless networks, connecting these devices to the Internet has never been easier. The growth of devices on the Internet today is chiefly occurring in two distinct ways. The first is that previously separate networks – such as video, voice, cellular, etc., - are all migrating toward shared IP. This trend requires the Internet to absorb wholesale transitions of full-scale networks into its existing framework.

At the same time, new classes of devices are becoming network enabled. The types of devices being connected today extend far beyond the laptops and cell phones we have become so accustomed to. Today, virtually all products that use electricity - from toys and coffee makers to cars and medical diagnostic machines - possess inherent data processing capability and the potential to be networked.

The fact that a rapidly expanding range of devices have the capability to automatically transmit information about status, performance and usage and can interact with people and other devices anywhere in real time points to the increasing complexity of these devices. For example, software for the average mobile phone contains over 10 million lines of code. Over the last five years, software in automobiles has grown from an average of 35 million lines of code to over 100 million. The astronomical growth of features and functions within and among connected devices pushes the bounds of what designers of software for products ever had in mind.

Some basic design principles must be put in place to guide the development of smart connected devices. It demands that we design not only devices and networks but also information management systems in ways not addressed by current IT technologies. The reader may ask, don't we already have a vast number of software products to address these challenges? Don't we have real-time operating systems and development tools? And aren't these technologies working together today quite handsomely?

Almost everyone will answer with a resounding “Yes!” But consider this analogy from Buckminster Fuller: Suppose you are traveling on an ocean liner that suddenly begins

TECHNOLOGY DOES NOT HAVE TO BREED COMPLEXITY

to sink. If you rip the lid off the grand piano in the ballroom, throw it overboard, and jump on it, the floating piano lid may well save your life. But if, under normal circumstances, you set about to design the best possible life preserver, are you going to come up with the lid of a grand piano?

The growing scale of interactions between feature-rich devices and the antiquated tools available today to develop them are like that piano lid. In a period of great change and tumult, it worked—in the sense that it kept us afloat. But that does not make it the best possible design, or qualify it to be something that we should plan to live with forever.

The tools we are working with today to make products “smart” were not designed to handle the scope of new capabilities, the diversity of devices and the massive volume of data-points generated from device interactions. These challenges are diluting the ability of engineering organizations to efficiently and effectively manage development. The fragmented nature of software offerings available today make it extremely difficult, if not impossible, to leverage design and engineering work across different platforms and devices.

What is needed is a common means of development that can leverage tools, tasks and code across families of interrelated devices. Customers expect evolving software tools to be functional, ubiquitous, and easy-to-use. Within this construct, however, the first two expectations run counter to the third. In order to achieve all three, a new approach is required -- a unified development framework for smart connected devices.

What would this entail:

Software tools and development protocols that work together seamlessly, securely and safely across diverse devices - historically software development has been very distinct for different types of products such as computers, phones and other intelligent [non-IT] equipment - in the future diverse device segments will be served by increasingly common software and tools.

Software that supports interoperability over long product life cycles - the ability to support products for 15-20 year life cycles and bridge complete technology shifts will be required. Devices will require multiple code upgrades to remain compatible and interoperable.

Software and development tools to address a broad range of application requirements - increasingly, product manufacturers will need to develop both GUI-based/data centric applications as well as low end “headless” non-GUI-based systems where boot loaders, schedulers, comms stacks and other OS components are used. OEMs need a single unified framework to design and build solutions that can be certified to interoperate across all the uP and OS classes they use.

WE ARE REACHING A CRITICAL JUNCTURE WHERE ORGANIZATIONS WILL BE CRYING OUT FOR A COMPLETELY NEW APPROACH

Software and tools that leverage re-use - given the scale of the Internet of Things it will simply not be humanly possible to write all the code required without large scale re-use - the ability, for example, to develop a security module or comms stack that can be designed and certified once but used multiple times.

We are reaching a critical juncture in market development where engineering organizations will soon be crying out for a completely new approach - one where the operating system and its components can be utilized again and again across an ever broader spectrum of devices. We must develop software and methods to automate development and facilitate re-use, or risk constraining the growth of this market.

IT professionals talk these days about the need for ever-evolving information services that can be made available anywhere, anytime, for any kind of information. Most of this talk surrounds the advent of cloud computing and services. All too often, IT people speak about cloud and web services with “the Internet” interchangeably without giving it a thought.

But the Web is not the Internet. The Internet itself is a simple, elegant, extensible, scalable, technology-neutral networking system that will do exactly what it was designed to do for the indefinite future. The same cannot be said of the Web, which is essentially an application running on top of the Internet. It is hardly the only possible Internet application, nor is it the most profound one conceivable.

The Web’s Achilles heel does not originate in its browsing software, or markup languages, or the other superficial aspects that most users touch directly. Those inventions are not necessarily ideal, but they are useful enough today, and they can be replaced over time with better alternatives. Rather, the Web’s weakness lies with its basic enabling technologies—in particular, relational databasing and the client-server model—and the restrictions they place upon structuring, storing, and retrieving data.

The Web, and now cloud services, stores information in one of two basic ways: utterly unstructured, or far too rigidly structured. The unstructured way gives us typical static Web pages, blog postings, etc., in which the basic unit of information is large, free-form, and lacking any fundamental identity.

The overly structured way involves the use of relational database tables that impose rigid, pre-ordained schemas on stored information. These schemas, designed by database administrators in advance, are not at all agile or easily extensible. Making even trivial changes to these schemas is a cumbersome, expensive process that affects all the data inside them. Just as importantly, they make deep, inflexible

06

Achieving Truly Flexible, Scalable Systems

TODAY'S CLOUD AND IT WORLD ARE STILL DISCONNECTED ISLANDS

assumptions about the meaning and context of the data they store. Both of these approaches to data-structure enforce severe limitations on the two things you want most in a global, pervasive-era information system: scalability and interoperability.

The client-server model underlying the Web greatly compounds the problem. Regardless of data-structure, when you put information on the Web you put it in a specific physical location—i.e., you upload your HTML or middleware files, or your database tables, to a specific directory on a specific server. Thus, a Web URL does not point to information per se; it points to the information's physical location, which is a very different thing.

This means that the Web is machine-centric, not information-centric. Information on the Web is not free to move, and because its life is tied to the life of a physical machine, information on the Web can easily become extinct. If a particular server goes offline, temporarily or forever, all its URLs are dead and the information is unavailable. You can “mirror” your data so that copies reside on multiple servers, but this does not change the fact that Web-based information is always tethered to specific locations on specific machines.

All of this adds up to a huge collection of information-islands called the World Wide Web. Assuming the islands remain in existence reliably, they are still fundamentally incapable of truly interoperating with other information-islands. We can create bridges between them, but islands they remain. That's what they were designed to be.

The latest bridges are called “cloud brokering services,” by which islands automatically “feed” information to each other. In terms of convenience for human beings, Web services are a distinct advance over manually navigating Web pages. You can now set up a browser page or a desktop applet to “consume” (receive and display) a Web-feed of news headlines, stock quotes, weather information, package tracking, and so on. And that's a good thing—a pleasant, time-saving, incremental improvement on yesterday's way of doing things.

But many people talk about Web and cloud services as “the next big thing,” and if that's true we can only say: We were hoping for something bigger. Consuming a cloud service is exactly the same thing as loading a Web page.

All of these new services are still prey to the client-server rigidity and vulnerability we've described, making them particularly unsuitable for vital public information services. And in order to concatenate information from diverse sources, these services depend upon data-tagging conventions that require yet more layers of prior agreement, schemas, markup, and manual human administration.

WHAT WOULD TRULY LIBERATED INFORMATION BE LIKE?

With each additional layer of such engineering and administration, the Web comes closer and closer to resembling a fantastically jury-rigged Rube Goldberg contraption. The reason is simple. These new computing schemes were not designed for a world driven by pervasive information flow.

In today's world, information is not free (and that's free as in "freedom," not free as in "free of charge"). In fact, thanks to present information architectures, it's not free to easily merge with other information and enable any kind of search-based intelligence.

What would truly liberated information be like? It might help to think of the atoms and molecules of the physical world. They have distinct identities, of course, but they are also capable of bonding with other atoms and molecules to create entirely different kinds of matter. Often this bonding requires special circumstances, such as extreme heat or pressure, but not always.

In the world of information, such bonding is not all that easy. Today's software platforms focus on execution processes that generate one of three types of data - unstructured, transactional or time series. For each of these data types, a specific set of intelligence tools have evolved to provide "insight" but, in most cases, these tools limit the questions that can be answered to those known in advance. So for a user attempting to do something as simple as asking a certain multi-dimensional question, creating new information from multiple data types that is an easily perceivable, manipulable, or mappable "model" of the answer to that question is a significant challenge.

The Internet of Things requires new platforms to fundamentally changes this paradigm, treating data from things, people, systems and the physical world as augmented representations. In other words, treating diverse data types equally. This enables processes connecting diverse data in any combination to be rapidly built and deployed.

The traditional approaches to data discovery and systems intelligence have two failings: they can't provide a holistic view of these diverse data types and, the types of intelligence tools available to users are, at best, arcane and typically limited in use to "specialists."

Users need new ways to find information and discover patterns on their own, without requiring specialists or IT support. This allows users to determine where deeper analytics or the creation of an ad hoc business process can add value.

Given the immature state of today's real-world systems, most people have trouble grasping the power and importance these capabilities enable. The ability to detect

07

Leveraging Collective Intelligence

TRADITIONAL APPROACHES TO SYSTEMS INTELLIGENCE WON'T WORK IN THE SMART SYTEMS WORLD

patterns in data is the holy grail of smart systems and The Internet of Things because it allows not only patterns but a whole higher order of intelligence to emerge from large collections of ordinary data. The implications are obviously immense.

We need an entirely new approach that avoids the confinements and limitations of the today's differing data types and tools. It allows data to maintain their fundamental identity while bonding freely with other data. Facilitating discovery, based on data and information accessibility, fusion and cumulative systems intelligence, is a fundamental goal for information systems. The world needs systems designed for a genuinely connected world in which there are no artificial barriers between pieces and types of information.

Smart connected devices are a global and economic phenomenon of unprecedented scale - potentially billions if not trillions of nodes will be connected if the right technology and business architectures are in place. Today the relationship between the evolving technologies driving the Internet of Things and the business models that these systems could inform is "disconnected" and problematic at best.

Apple and some of its peers in consumer space present an interesting case for how Smart Business models are developing. Apple is primarily a consumer-focused tech vendor, with no need for the vertical industry capabilities required in B2B applications.

Still, Apple provides a model for creating a Smart System solution that pulls together technologies from multiple domains and packages that solution in a way that wins buyer acceptance. Looking beyond Apple to

Google, Amazon, Facebook and other players coming from their roots in the evolving consumer mobile internet arena, there are a variety of new business models emerging from cloud and related services platform players that are the embodiment of Smart Systems. Add to this the momentum from these players are creating via collaboration with their app communities -- Amazon has over 100,000 developers building applications and businesses; Apple's App Store has created a phenomenal lead position -- they are all driving entirely new forms of collaboration and peer product development.

The traditional notion of M2M applications has largely grown up in a B2B context -- equipment manufacturers developing remote services and support automation tied closely to their equipment services contracts.

These models are focused almost exclusively on customer support and automation -- not on new Smart Services value beyond support. As these two classes of business models inch closer to each other in the marketplace it is increasingly evident that consumer smart business models provide many lessons for the "cloistered" equipment manufacturers in B2B arenas. The business benefits of large

08

Smart Systems Impact on Evolving Business Models

NEW BUSINESS MODELS AND NEW BUSINESS DESIGN PRACTICES ARE REQUIRED TO REALIZE THE POTENTIAL

scale collaboration and social networking tools in the B2B arena are finally being recognized.

Executives in many product manufacturing companies are only just beginning to understand the opportunities driven by the complicity between collaboration tools, networked device intelligence and service business design and it is this set of relationships, not the technological shift, that will benefit but also challenge many product manufacturers.

What product manufacturers looking to leverage collaboration or benefit from connecting “smart products” to the Internet need to understand is we have entered a phase in the marketplace where ideas can emerge from anywhere in the world; new network and IT tools have dramatically reduced the cost of utilizing them. The bottom line: no single company should look to innovate on their own.

Collaborative communities both inform and express the strategy. Built to pursue multiple aims simultaneously, a dynamic network of connected products and people drive new information values which, in turn, create new influences in the marketplace. Power in human and device collaboration structures falls to those who best understand how to use this information and influence to get and keep a key position.

Forging collaborative communities means managing uncertainty. A product OEM needs a clear understanding of the forces at work between and among devices and people. They must try to identify those few “interactions” that make a difference. From this understanding provided via real time interactions, early indicators of the true direction of customer needs and behaviors can be understood long before others.

Some things that look easy turn out to be hard. That’s part of the strange saga of the Internet of Things and its perpetual attempts to get itself off the ground. But some things that should be kept simple are allowed to get unnecessarily complex, and that’s the other part of the story. The drive to develop technology can inspire grandiose visions that make simple thinking seem somehow embarrassing or not worthwhile. That’s not a good thing when defining and deploying real-world technology to deliver innovation. This is where the new values of a true smart systems architecture really come into focus.

We’ve articulated a “smart systems” vision and road map that emphasizes several points, including:

The architecture required to inform an Internet of Things looks nothing like today’s kludge diagrams of smart devices connected to clouds; it will inevitably evolve to a radically different peer-to-peer architecture.

09

How Should We View Future Smart Systems Architecture?

WE ARE JUST BEGINNING TO UNDERSTAND THE OPPORTUNITIES

As the architecture evolves, the definitions of its most basic elements that everyone is so anxious to quantify today (servers, devices, routers, hubs, etc.) will change radically.

What most people who do not understand this evolution are missing is that you cannot accurately quantify the future state of the Internet of Things by merely counting today's mess of architecture and systems elements.

Today the world of smart communicating devices is mostly organized in hierarchies with smart user interface devices at the top and the dumb devices [often analog or serial sensors and actuators] at the bottom. Within this structure, there are typically various types of "middle box" supervisory devices forming a point of connectivity and control for the sensors and actuators as well as the infrastructure for the network. From our perspective, this description of today's M2M, remote services and connected systems architectures looks very familiar and is largely organized much like client-server based computer systems. no surprise given they were designed in the 1980's.

As the Internet of Things opportunity expands, the sensor and actuator devices will all become smart themselves and the connectivity between them (devices, for the most part, that have never been connected) will become more and more complex. As the numbers of smart devices grow, the existing client-server hierarchy and the related "middle boxes" acting as hubs, controllers and interfaces will quickly start to blur. In this future-state, the need for any kind of traditional client-server architecture will become superfluous. In a future Internet of Things, the days of hierarchical models are numbered.

Now, imagine a future smart systems world where sensors and devices that were once connected by twisted pair, current loops or were hardwired become networked with all devices integrated onto one IP-based network (wired or wireless). In this new world, the "middle boxes" don't need traditional input/output (I/O) hardware or interfaces. They begin to look just like network computers running applications designed to interact with peer devices and carry out functions with their "herd" or "clusters" of smart sensors and devices – let's call these new middle boxes, "aggregators".

We can readily imagine an application environment where there may be several aggregators running applications which overlap sharing their sensors and actuators, some even 'sharing' a whole herd – a smart building application, for example, where the processor in an occupancy sensor is used to turn the lights on, change the heating or cooling profile or alert security.

In this evolving architecture, the network essentially flattens until the end-point devices are merely peers and a variety of applications reside on one or more network

THE NOTION THAT ALL THESE “THINGS” WILL PRODUCE STREAMING DATA TO BE PROCESSED IN SOME CLOUD WON’T WORK

controllers which look for all intents and purposes like today’s cellular router/modems, industrial PCs or small “headless” high availability distributed servers.

In a smart systems application world designed to capture, log and analyze large volumes of data from sensors, the aggregators, such as we are describing here, will carry out the process of taking raw data and distilling it into information “locally.” Local processing is required to reduce the otherwise untenable Internet traffic challenges that arise from connecting billions of devices. The notion that all these “things” and devices will produce streaming data that has to be processed in some cloud will simply not work. It makes more sense structurally and economically to execute these interactions in a more distributed architecture near the sensors and actuators where the application-context prevails.

This is the move we’ve all been waiting for to a truly distributed architecture because today’s systems will not be able to scale and interact effectively where there are billions of nodes involved. In this portrayal, the aggregators will become peers in a network along with the sensors, actuators and controllers whether or not they have GUIs (headed) or not (headless). The aggregators will also become unified application platforms from which to provide services to devices and users where the applications run, where the data is turned into information, where storage takes place, and where the browsing of information ultimately takes place too – not in some server farm in a cloud data center. Even the mobile handsets we admire so much today are but a tiny class of user interface and communications devices in an Internet of Things world where there will be 100 times more “things” than humans.

In this future Smart Systems arena, the architecture and relationship structures between and among smart devices will become more or less flat. Hierarchy will disappear. Peer-to-peer will become the “norm” for interactions (with its attendant benefits and pitfalls). From our view the movement towards peer-to-peer, and the view that many people hold that this is somehow novel, is ironic given that the Internet was originally designed for peer-to-peer interactions.

This is a vast opportunity of untold scale. If we hypothesize that there will be 2-5 “headed” user interface devices for every person on the planet – say 12-30 billion in the future. Of these, if 4-5 billion will be portable and/or low power, there will likely be 30-40 billion controller/aggregator devices, potentially managing around 60-100 billion end-point sensor/actuator devices – a ratio of 20:1 or so...and that’s just the relatively near term. When you think there are only 1 billion PCs on the planet and 3 billion cell phones, this all starts to come into focus and in many ways, aggregators will be the next “PC”...but with 30X more of them.

10

**Smart Systems Is A
Radical Departure From
Current Thinking**

THIS CREATES A NEW MARKET AND TECHNOLOGY CATEGORY

Today we can only see the beginnings of this structurally: Industrial PC's, communications hubs and all manner of gateway-like boxes all morphing to aggregators with handsets and tablets morphing to universal user interface devices. In this vision all we can hear in the today's marketplace is the siren song of client-server and Cloud stories.

In our years of experience, we have all too often seen the unfortunate scenarios that managers create when uncertainty and complexity force them to rely on selective attention. Unfortunately, when this happens, selective attention naturally gravitates toward what's readily available: past experience and uncertain assumptions. Today's IT and telco infrastructure players are doing just this. By ignoring important trends simply because it's difficult to perceive an alternative future, these managers are certainly leaving the door open for competition that will lead to their eventual obsolescence...which will make for a very interesting world to live in...

Though their business models are intermingling today, all of the major categories of suppliers in the "traditional" so-called M2M software arena have historically operated within well-established assumptions about product scope and business models. No one would characterize the existing players of being technology or business model innovators or disruptive in nature.

Radical new thinking about information technology must begin at the most basic levels, with new conceptions about the interactions of information with people, systems and devices. We think more about future proofing innovations by making the fewest possible assumptions about the nature of networked objects and the data they produce, carry or process - we need a much broader, all-encompassing view of information. Ultimately, this type of smart systems architecture will alter traditional business models and how new applications are realized.

Since all of this that we are describing is a radical departure from current offerings and business practices, and is driven by a very unique set of needs, it stands to reason that these types of solutions do not fall within the narrow specialties of the existing players. In fact, the architecture being described is probably best viewed as an entirely new market category. This is particularly true given the disjointed patchwork of solutions presently in place and the apparent lack of vision from existing players of what's required in the future. The opportunity to lead in developing and shaping this market looks wide open.

11

Getting There First; But To Where?

