Architecting in the Age of the Industrial Internet

Amine Chigani
Joseph Salvo
Benjamin Beckmann
Thomas Citriniti

Business Integration Technologies Laboratory
Software Sciences & Analytics
GE Global Research
Niskayuna, NY
E-mail: chigani@ge.com

9th Software Architecture Technology User Network (SATURN) Conference
April 29 – May 3, 2013 in Minneapolis, MN

imagination at work

Public

Overview

- Background
  - Cyber Physical Systems
  - CPS Systems Coming Online
  - The Industrial Internet
  - Convergence-driven Savings

- Key Architectural Challenges
  - Abstraction – Standards - Big Data – Cloud – Engineering

Cyber Physical Systems

Cyber Physical Systems (CPS) are large-scale, smart networked systems with embedded sensors, processors, and actuators designed to sense, control, and interact with the physical world and people, and support real-time performance in safety-critical environments.

imagination at work

CPS Systems Coming Online

imagination at work

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

The Industrial Internet

The industrial internet refers to the *convergence* of the global industrial ecosystem, advanced computing and manufacturing, pervasive sensing, and ubiquitous network connectivity that enables complex, complete cyber physical systems to come online.

Convergence-driven Savings

Minds

Machines

Technology

Industry

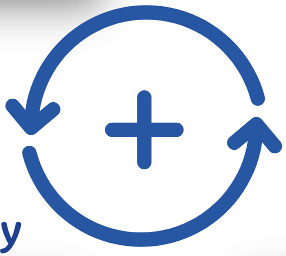

Source: GE 2012 Annual Report. http://www.ge.com/ar2012/

INDUSTRIAL INTERNET = CUSTOMER OUTCOMES

THE POWER OF 1%

Through efficiencies enabled by the Industrial Internet, a 1% change can deliver tremendous value to customers.

		15-YEAR SAVINGS
AVIATION	1% Fuel Savings	$30B
POWER	1% Fuel Savings	$66B
RAIL	1% Reduction in System Inefficiency	$27B
HEALTHCARE	1% Reduction in System Inefficiency	$63B
OIL & GAS	1% Reduction in Capital Expenditures	$90B

Potential of adding $15 Trillion to global GDP by 2030

imagination at work

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

What does this new context mean for the software architecture practice?

Key Architecting Challenges

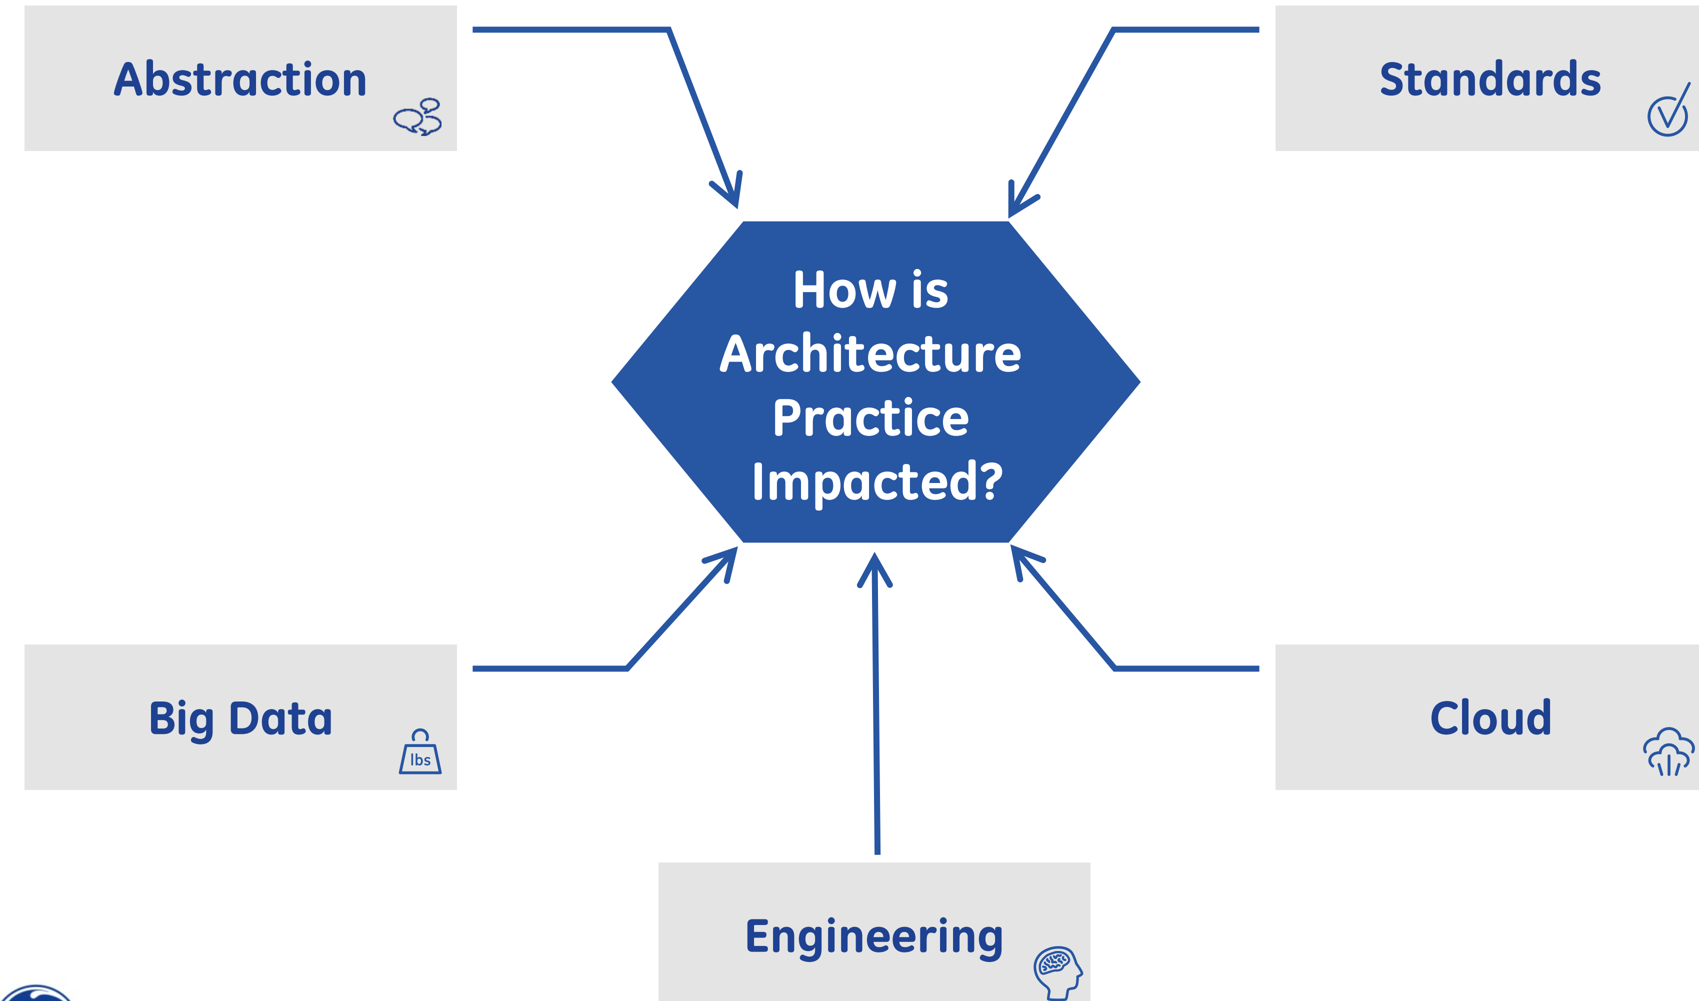

imagination at work

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Abstraction

The scale of CPS systems and interdependency among their elements will mandate a greater emphasis on systems-level, end-to-end thinking about solution architectures that can be used by stakeholders of different organizations, disciplines, and expertise. Architecting the software backbone of a network of CPS systems will require skills beyond those related to the software craft. Multi-discipline knowledge of the basics of machinery, analytics, mechanics, supply chain, manufacturing, and ultimately business value will become critical to enable the architect to create an abstraction of architecture decisions and choices suitable for a broad audience.

Architecture Considerations

- Multi-disciplinary domain knowledge
  - Architects' skills should include basic knowledge of machinery, analytics, supply chain, value, etc.
- Cross-discipline language and iconology
  - Model, component, resource, workflow, architecture, etc.: terms that mean different things in different disciplines
- Architecture views catering to broader audience
  - Various types of engineers not just software
  - Mapping of architecture choices to business value

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Focus on Shared Abstraction

Develop architecture views in collaboration with stakeholders from targeted audience

- **Example:** Architect develops the system's conceptual data model in collaboration with a materials engineer
- Inspired by Agile's customer involvement principle
- Enables a shared abstraction level

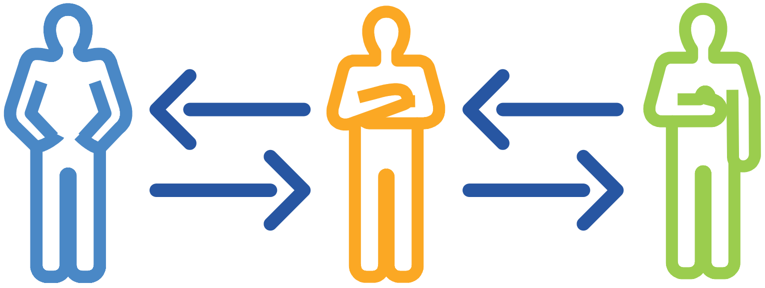

imagination at work

Standards

Corollary to abstraction, enabling communication and collaboration among a wider community of stakeholders will require standardization beyond the software architecture community. Standardized architecture tools and nomenclature should include other engineering disciplines such as mechanical engineering, physics, natural sciences, mathematics, manufacturing, and many others.

Architecture Considerations

+	-
It is software that enables us to compose, connect, and take advantage of a network of CPS systems	However, poor software practice is key to system failures (84% in 1996 and 66% in 2003)

- Standards that hold software same accountability as other engineering disciplines
  - Regulatory, industry, and organization compliance
  - **example:** Congress' attempt to pass cyber security law (defeated)
- Standards that can be understood and applied across disciplines
  - **example:** DODAF 2.0 standardizes how architecture knowledge is developed, documented, and consumed across DOD

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research

Big Data

The sheer number of machines (e.g., engines, medical devices, turbines, etc.) expected to come online and the volume of data expected to be generated and transmitted through the industrial internet as a result will bring about big data challenges beyond what we know today. One major architecture challenge will be to decide what gets thrown away, processed at the edge (i.e., point of contact of the CPS system with the physical world), or transmitted and processed away from the point of generation (i.e., the cyber world). In other words, intelligent, resource-conscious analytics will be used at the points of data generation and aggregation to decide what gets passed from the sensor to the industrial internet and into the backend big data system for storage and processing. In this effort, new data architecture tactics and modeling approaches will become paramount.

lbs

The Issue is at the Edge

- The challenge is the decision to be made at the edge – where the physical meets the cyber
- Often, computing at the edge is resource scarce
- Environments at the edge are often challenging
  - i.e., in the sky; under water; 100 mph; etc.

What do you do when one element (i.e., engine) of your network of CPS systems generates 1T of data a day, and you have thousands of them?

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research

Architecture Options

1	Discard – ignore – throw away
2	Process at the edge (point of generation)
3	Transmit via industrial internet and process away from the edge
4	Hybrid option

- It depends on architectural goals:
  - Insights into historical data: Diagnostics
  - Foresight about future behavior: Prognostics

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research

Context-driven Approach

- Identify requirements related to edge processing and long-term analytics
- Conduct tradeoff analysis and weigh in options
  - Consider distributed computing principle supported by cloudlet architecture: processing should follow data.
- Decide on architecture solution taking into consideration:
  - Environment & Infrastructure at point of generation
  - Evolution space of the system
    - Will analytics that informs processing at the edge evolve?

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research

Cloud

Cloud-based computing enables scale and elasticity – two essential elements of the expending and evolving nature of CPS systems. To address the computing and networking needs of many engineering companies, manufacturers, suppliers, and other stakeholders in this context, cloud computing offers an affordable, efficient strategy to come aboard the industrial internet early to ensure a continued competitive edge. However, privacy issues related to export control, intellectual property, corporate identity, governance, and ownership among others must be addressed.

The Scalability Problem

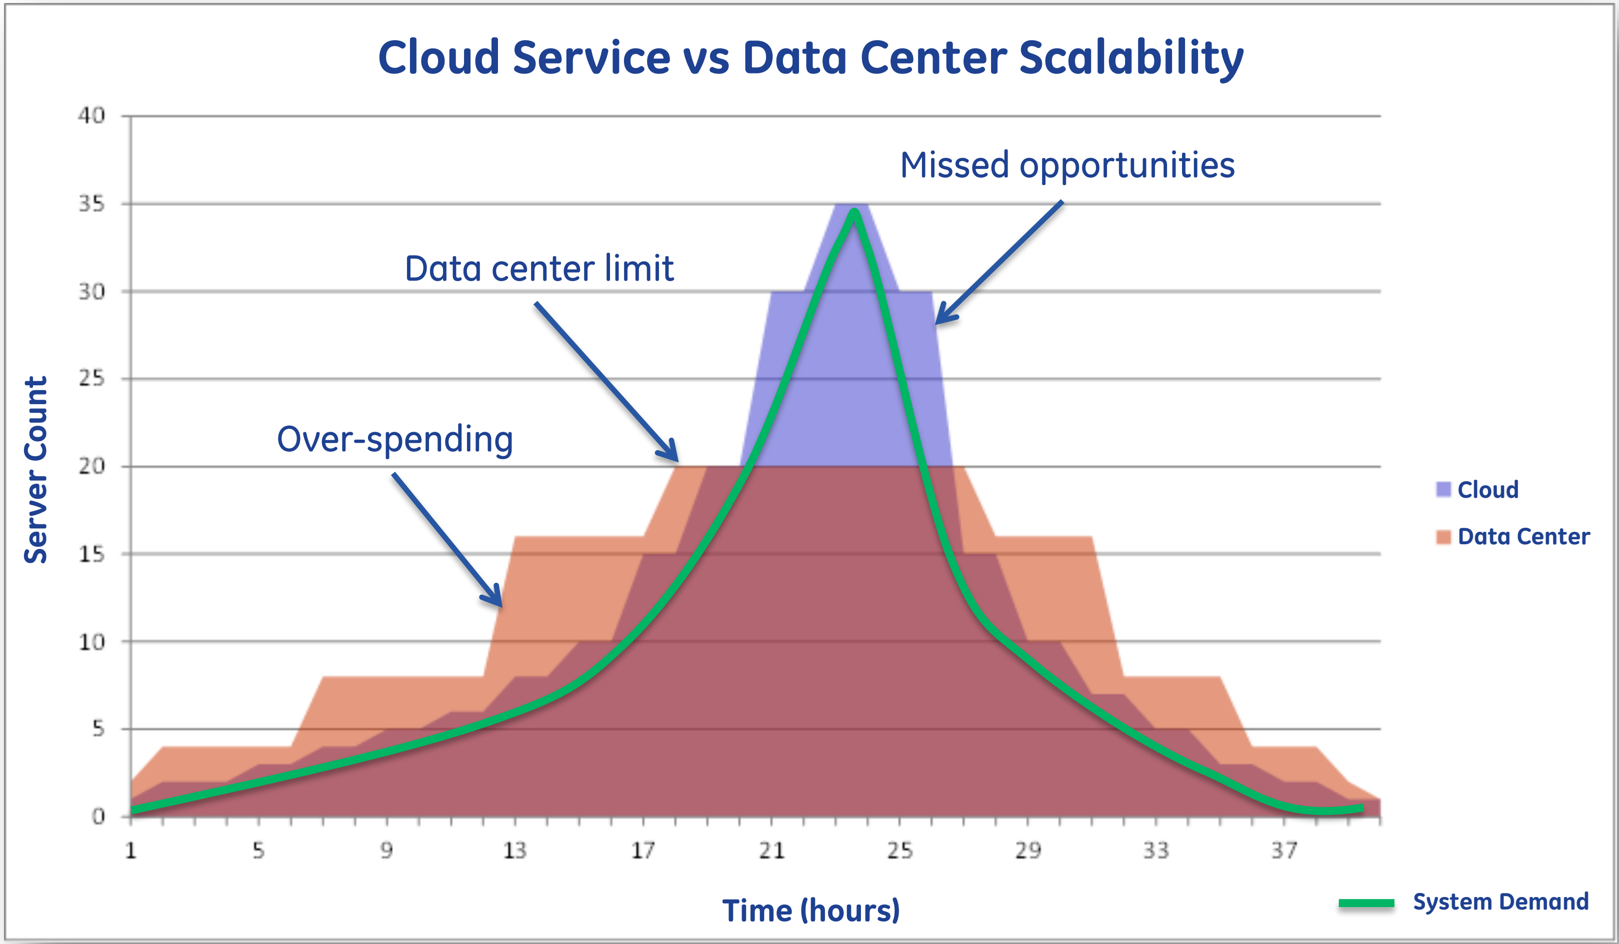

Example of performance tests conducted for a platform development effort at GE Global Research.

imagination at work

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Cloud Considerations

Benefits

Upfront cost:

- None compared to large NPI investments

Scale:

- Delivery and sustainable support mechanism for global deployment

Elasticity:

- Need-based growth

Issues

Governance issues:

- Intellectual property
- Security and privacy
- Corporate identify
- Ownership and access

Compliance issues:

- Export control and related regulations

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Example: GE CEED on the Cloud

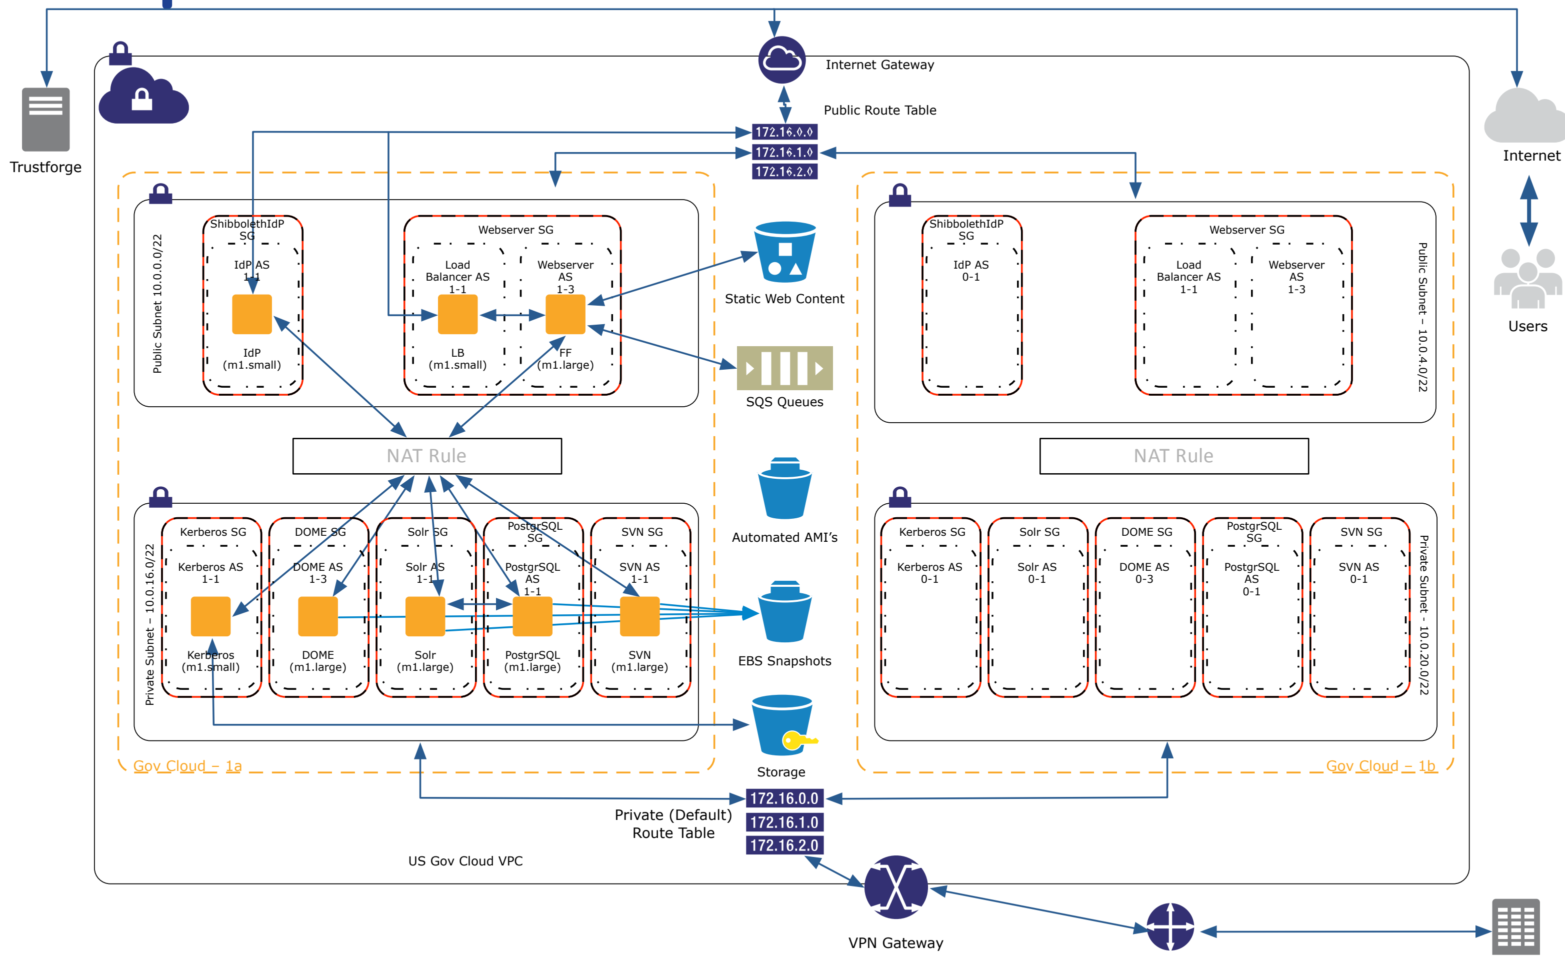

Collaborative Ecosystem for Evolutionary Design (CEED) Platform

First Commercial Use of the Amazon's GovCloud

imagination at work

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Cloud Industry Needs

- **Certifications and Accreditations**
  - ITAR-compliant environment
  - Federal Information Processing Standard (FIPS) 140-2 (cryptography)
  - Federal Information Security Management Act (FISMA)
  - Payment Card Industry Data Security Standard (PCI DSS) Level 1
  - Health Insurance Portability and Accountability Act (HIPAA)
  - … and more!

- **Infrastructure Quality Attribute Support**
  - Redundant availability zones supports failover
  - Public & private subnets protects sensitive data
  - Security & auto scaling groups limit connectivity and support elasticity
  - Data & content stored using scalable, reliable, fast, inexpensive service

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Engineering

Stove-piped, single-discipline-focused engineering of products and services that form the components of CPS systems will no longer fit within an industrial internet enabled environment. Time to market, cost, complexity, and competitiveness will require a much more robust engineering design methodology. The potential to transform how engineering is conducted through the adoption of a collaborative, crowdsourcing driven approach to engineering and manufacturing is becoming a reality.

Unsustainable Approach

APPROVED FOR PUBLIC RELEASE. DISTRIBUTION UNLIMITED.

DARPA Historical schedule trends with complexity

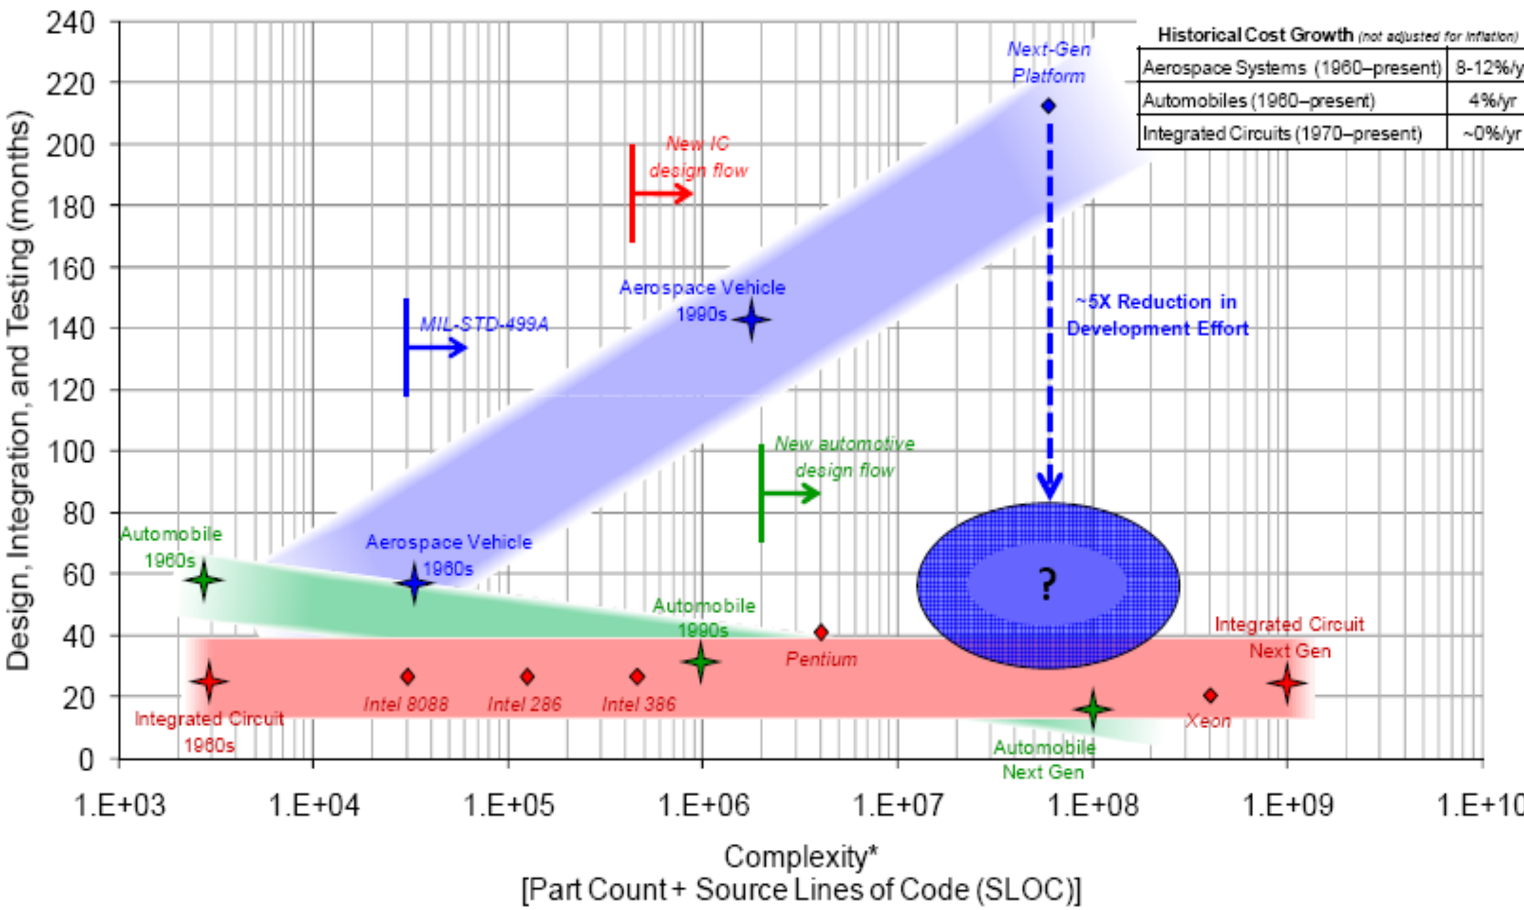

Historical Cost Growth *(not adjusted for inflation)*

Aerospace Systems (1960–present)	8-12%/yr
Automobiles (1960–present)	4%/yr
Integrated Circuits (1970–present)	~0%/yr

Note (*): Not a great metric. But that's what we have today. META will come up with better metrics.

Radical Paradigm Change

Goal: 5x development time reduction

- Disrupt the traditional design build paradigm
- Blur traditional talent and knowledge boundaries through enablement of virtual and physical collaboration
- Attract orders of magnitude more talent through SME-focused collaboration

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research
4/26/13

Solution Elements

From:	Towards:
• Isolated Expertise	• Integrated Knowledge
• Separate Tool & Data Repositories	• Unified & Searchable
• Manual Workflow & Data Movement	• Automated Workflow & Data Movement
• Task Centric	• Project/Process Centric

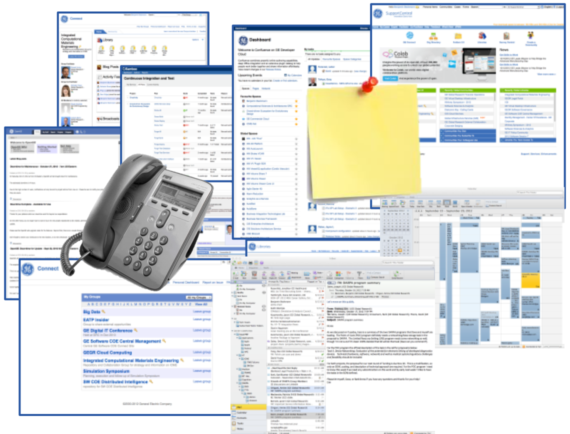

→

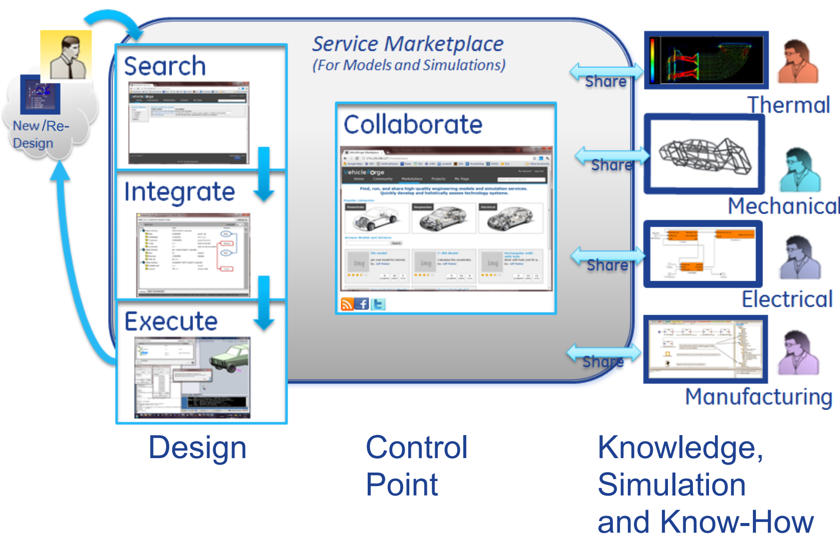

One unified environment

Streamlining Productivity of $Billion+ Engineering Community

imagination at work

A. Chigani, J. Salvo, B. Beckmann, T. Citriniti
GE Global Research

4/26/13

About GE Global Research

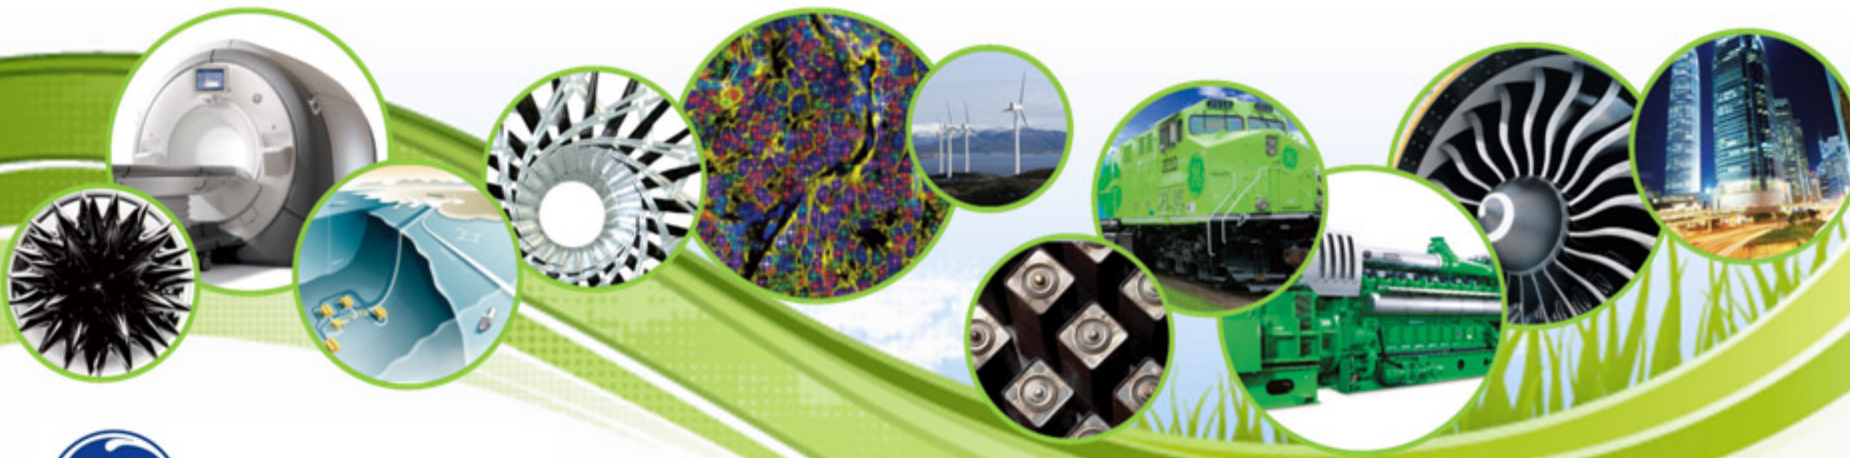

imagination at work

GE Overview

Founded in 1892, 330,000 employees worldwide, $173B annual revenues, only company in Dow Jones index originally listed in 1896.

Energy

- Power & Water
- Energy Management
- Oil & Gas

Technology Infrastructure

- Aviation
- Healthcare
- Transportation

GE Capital

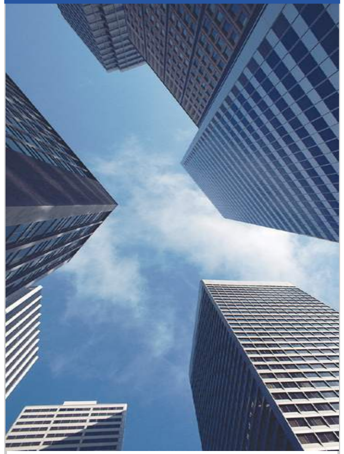

- Aviation Financial
- Commercial Finance
- Energy Financial
- GE Money
- Treasury

Home & Business Solutions

- Appliances & Lighting
- Intelligent Platforms

GE Global Research: Market-focused R&D

- First U.S. industrial lab
- Began 1900 in Schenectady, NY
- Founding principle … improve businesses through technology
- One of the world's most diverse industrial labs

A Tradition of Innovation

1909 Ductile tungsten
1913 Medical X-ray
1927 First television broadcast reception
1932 Langmuir Nobel Prize in chemistry
1938 Invisible/glareless glass
1942 First US jet engine
1953 LEXAN™ polycarbonate
1955 Man-made diamonds
1962 Semi-conductor laser
1973 Giaever Nobel Prize in physics
1984 Magnetic resonance imaging
1994 GE90® composite fan blade
1999 Digital X-ray
2004 Lightspeed VCT
2009 Wide Bore 1.5T MR System
2010 Energy Smart® LED




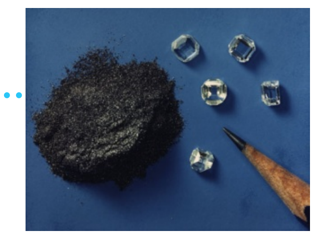



imagination at work

Amine Chigani, Ph.D.

Dr. Amine Chigani is a Computer Scientist in the Business Integration Technologies Lab at GE Global Research. His basic research focuses on software architecture, system of systems engineering, process modeling, software quality, and service-oriented computing. He has experience primarily in applying systems-level analysis and architecting to large-scale, enterprise-wide ecosystems. His current work focuses on the development of a GE-based, crowd-driven, engineering design ecosystem that transforms how engineers collaborate on designing products, running analysis models, and interacting with data. Prior to joining GE Global Research, Dr. Chigani worked as a Visiting Scientist at Carnegie Mellon's Software Engineering Institute (SEI) helping the Department of Homeland Security develop integration strategies to guide the adoption of the Commercial Mobile Alert System (CMAS) by emergency alert originators nationwide. He led the development of the Integration Strategy Framework currently used to develop these strategies.

Dr. Chigani has publications in the International Journal of System of Systems Engineering, IEEE Software Engineering Workshop, SEI Software Architecture User Network Conference, IEEE Conference on Software Engineering Education and Training, and International Conference on Software Reuse. Professionally, he is a member of ACM, IEEE, SEI, and ASEE. He holds a BS (2003) in Computer Science from Radford University, and MS (2007) and Ph.D. (2011) in Computer Science from Virginia Tech. He also holds the Software Architecture Professional Certificate (2010) from the SEI.

Joseph Salvo, Ph.D.

Dr. Joseph Salvo and his group have developed a series of large-scale internet-based sensing arrays to manage and oversee business systems and deliver value-added services. Their most recent business releases include a number of complex decision ecosystems (e.g. Electric Vehicle Charging Systems, GE Veriwise™, GE Railwise™, Global Vendor Managed Inventory, Ener.GE™, and E-Materials Management) that deliver near real-time customer value through system transparency and knowledge-based computational algorithms.

Dr. Salvo's work focuses on pervasive networked sensors and knowledge repositories that deliver time-critical, high fidelity data to enable information analysis across traditional business process boundaries; crowdsourcing technology that enables the democratization of the design, model, manufacture paradigm to collapse the time to produce complex cyber physical systems; and total supply chain, energy management and financial services that can be integrated to create a virtual enterprise environment that encourages discovery and process improvement on a global basis. Commercial business implementations of this work are currently active in Europe, Asia as well as in North and South America.

He is a member of the board at the M.I.T. Forum for Supply Chain Innovation.

Benjamin E. Beckman, Ph.D.

Dr. Benjamin Beckmann is a Computer Scientist in the Business Integration Technologies Lab at the General Electric Global Research. He received B.S. (2002), M.S. (2004) degrees in computer science from Western Michigan University and his Ph.D. (2010) degree in computer science from the Michigan State University. His research in the areas of evolutionary computation, cloud computing, crowdsourcing, self-organizations, and artificial life are motivated by his desire to understand interactions among individual entities in a complex virtual, cyber-physical, and social systems.

Dr. Beckmann has established a strong track record as an independent researcher with a steady flow of publications ranging from the evolution of quorum sensing in self-replicating computer programs to adaptive logic for balancing of non-functional tradeoffs during system reconfiguration. His research has focused on evolutionary pressures the form communities and mold behavior. He has been an active contributor to open-source software platforms (including Avida, PECL, and vehicleforge) that support his research. His research has led to twenty refereed publications, and has been the focus of articles in The New York Times, New Scientist, and the Wall Street Journal.

Thomas Citriniti

Mr. Thomas Citriniti is a Computer Scientist in the Software Sciences and Analytics organization at GE Global Research. Mr. Citriniti has over 20 years of experience building software products in the commercial market prior to joining GE. This work includes initial research, architecting, and delivering a new and existing software products. His focus over the last 10 years has been working with large fortune 500 customers to help architect their enterprise applications to ensure uptime and scalability.

Mr. Citriniti has authored publications and presented at conferences including SIGGRAPH, The Where in Business, Directions, Insights, and MapWorld conferences. Mr. Citriniti also authored and taught a graduate-level course at Rensselaer Polytechnic Institute on 3D Computer Graphics and Scientific Visualization. This work focused on the presentation of real world phenomena using 3D simulation and multiple visualization methods. Mr. Citriniti has a B.S from SUNY Cortland and a M.S. from Rensselaer Polytechnic Institute.