Architecture of Federated Enterprise

Jay S. Bayne

Senior Member IEEE, ACM, ISA
Echelon 4 LLC
jbayne@echelon4.com

Abstract

The capabilities of the Internet and its associated World-wide Web have reached a level of maturity in both technology and application that they now provide the platform of choice for many services of public and private enterprise. Web services, as key elements of this platform, provide a framework and tools for transaction-oriented (primarily stateless) applications. With the emergence of the global information grid, we are moving into an era in which growth through innovation and value generation is encouraging, if not forcing, individual enterprises to aggregate into more responsive and tightly coupled collaborative alliances, or virtual organizations. In open free-market economies such federated enterprises require new web-based services to support process-oriented (primarily stateful) applications. This is especially true of applications that provide automation and control of value production processes. We introduce here a framework for discussing grid-based distributed real-time automation and control of federated enterprise.

Key Words

Distributed real-time systems, cybernetics, federated systems, intelligent automation

Introduction

The next generation Internet and its supporting web services will support new classes of applications. One such class provides distributed real-time[1] automation and control services. For this class we are interested in the problems of orchestration (coordination and regulation) of performance in federated enterprise[2] – teams of [semi-] autonomous value-producing entities that are bound together for higher purpose.

Following Jeffersonian principles, members of a federated enterprise are required to 1) be viable and identifiable members of a community, 2) be governed by collectively (and democratically) determined laws, and 3) contribute their individual efforts to coherent ensemble behaviors that characterize the outcomes of the federation as a whole. As such, there is a higher order logic that defines a federation's mission. As defined in modern democracies federation objectives include securing the freedom of individuals (enterprises) to pursue their individual objectives while establishing and securing the common good. As such, federation rules enumerate the powers granted the federation by its members for their mutual well being. Thus, the federation establishes the intellectual and commercial "commons" on which "acceptable" individual and group behavior is defined, carried out and evaluated. We assume that such democratic principles apply no less to federated automation systems acting on behalf of free-market enterprise.

The three principles enumerated above have direct implications on the architecture and design of distributed real-time automation (aka, command and control applications) applied to grid-connected enterprise. Lawrence Lessig, Stanford professor of constitutional law, has eloquently argued that enterprise architectures *allow* [9] – meaning that the code that implements enterprise automation is a determining factor in establishing laws of cyberspace. Next generation Web services will contribute to the establishment of such laws, especially through support of command and control protocols and services.

The term *federation* as used here has two connotations. The first is in reference to the internal organization of a legal entity. The second refers to loosely coupled alliances among independent legal entities. In multi-domestic corporations, a single commercial entity (e.g., division or business unit) may comprise several domestic legal entities. In this case one or more legal entities in a multi-domestic corporation may participate in federations with one or more entities from another corporation. For this discussion we shall consider only a single federated corporate structure.

---

[1] A system is *real-time* to the extent that meeting time constraints is an explicit requirement of its correct behavior.
[2] *Enterprise* is an arbitrary interactive and systematic unit of production of a quantifiable value proposition.

Figure 1 depicts a multilevel enterprise command and control (C2) hierarchy. The structure depicts the accountability hierarchy of two divisions (business areas) and their governing corporate control level. The divisions are each organized into business units, plants, production areas and production units. At the corporate level the structure includes operations and development functions.

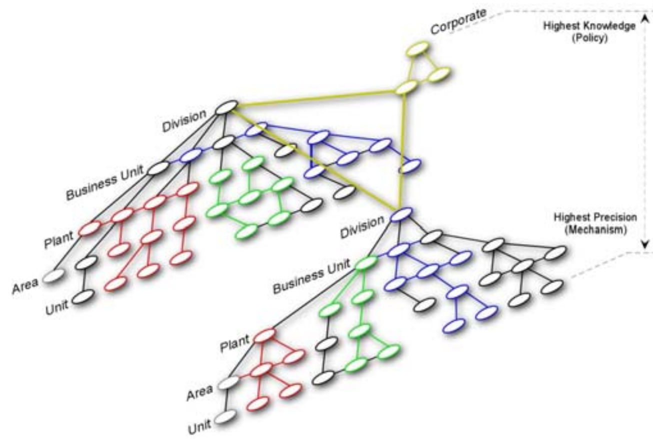

Figure 1: Command & Control Hierarchy

Automation and control services exist at each level. For example, at the unit operations level they provide sequence and modulating controls for motor, power, steam, start-up and shut-down, safety and production applications. At the intermediate levels planning, scheduling, maintenance management and supervisory activities are performed. At the business unit and divisional levels we find order processing, fulfillment, financial controls and transportation logistics. Nodes in figure 1 define the *spans of a control* of management services that populate the enterprise infrastructure. As corporations, seeking growth in revenues and market presence, continue to expand through mergers, acquisitions and alliances, the problems of coordination and control are compounding.

Extended Enterprise – Emergence of Virtual Organizations

In the 20 years following the January 1983 deployment of TCP/IP within the ARPANET, scientific and commercial enterprises have relied on the Internet to provide transport services for transactional data management applications. As a result, eBusiness has blossomed into broad services providing supply chain, electronic banking, information management (e.g., search engines) and “etailing” applications. With the advent of the Global Information Grid [8] and its emphasis on distributed data and supercomputer support for eScience applications, new services are emerging that provide computational resource scheduling, security, and batch control for large-scale collaborative computing activities. These new services are being applied to medical research, leading to the emergence of clinical eMedicine applications. The Department of Defense has adopted the grid as part of its “network-centric warfare” initiatives, and in support of its objectives to harmonize and rationalize the command and control of defense and homeland security (aka, eDefense) [7].

In support of internet-hosted computing processes an entire industry has emerged to take advantage of scale economics and provide consolidated and shared communications and computing (ISP) and applications (ASP) infrastructure. This has led to the need to remotely manage network-connected data centers through “network operating centers” (NOC) that are operationally similar to telecommunications and broadcast television network operations centers. These centers are focused on low-level C2 applications related to infrastructure support. They are precursors to higher-level enterprise applications in the same way that autonomic controls (e.g., the sympathetic nervous system) in man and animals are precursors to consciousness.

With the advent of standards and technologies sufficient to host higher-level distributed applications, the electronic frontier is poised to address the challenges of defining protocols and services to govern *extended enterprises* – virtual organizations whose “rules of engagement” are fashioned on, and defined within, the intellectual and commercial (i.e., legal, contractual, management) *commons* offered by services of the next generation internet and its enhanced web. The foundation of such a commons must certainly include a shared concept of value – a means of establishing relative worth between parties engaged in the trading of information, goods and services.

Value Production

Enterprises, whether public or private, for-profit or not-for-profit, military or commercial, exist to promote and sustain their unique *value propositions* within evolving financial (equity) and product market conditions. Value production is therefore a continuous and typically time-critical computation of those value propositions. In federations of real-time enterprise, as well as within the legal boundaries of its members, value is often difficult to define and measure. It is difficult within enterprises since value propositions (and associated value creation processes) are typically spread across organizational boundaries and involve weakly defined business processes. It is difficult across federated enterprise boundaries since joint (mutual) value propositions must be expressed in terms that transcend participants' own unique and often conflicting definitions of asset and supply chain objectives and resources. Hence, any effort at defining policies and mechanisms for the management and control of virtual organizations requires development of a common lexicon, and an agreed-upon set of interface specifications to a core set of joint value management services. This paper offers a framework for discussing value production in grid-connected systems that recognizes individual objectives while at the same time enabling management of shared objectives.

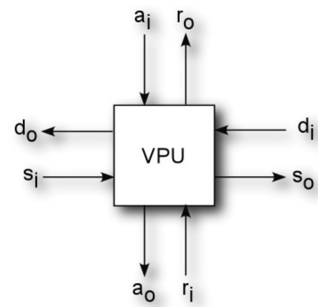

Figure 2 – VPU

The C2 model used here is one based on the notion of "rational agents" [2, 11], semi-autonomous computations that are self-serving (goal seeking) but that participate in alliances as a means of remaining viable[3], of proactively sustaining their existence. In abstract terms, a viable enterprise is one whose value production units (VPU, figure 2), while executing their respective value production processes, are able to cooperate through well-defined interfaces. As such, VPUs define distributed virtual machines. For this paper we shall discuss VPUs associated with eBusiness activities, but we could equally well discuss eScience, eMedicine, eDefense or other grid-enabled federations.

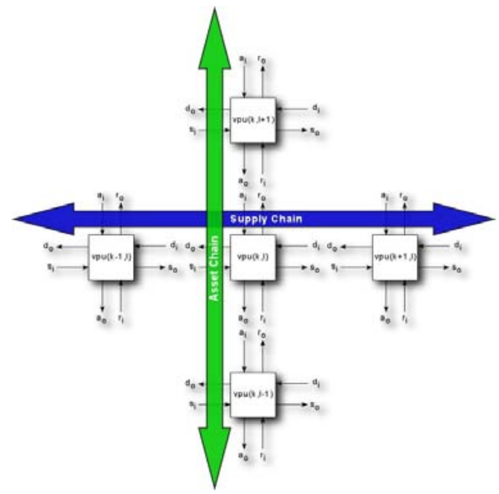

Figure 3 - VPU Production Web

The essential character of value creation is defined in traditional cost-benefit terms as in income statements and balance sheets. In simplified terms, a VPU is a computation that simultaneously satisfies the requirement to produce positive returns on invested capital assets in a vertical *asset chain* while profitably providing goods or services in a horizontal *supply chain*. Value production, viewed as a black box, is a process with two clients – *investors* supplying assets, and *customers* supplying demands for goods and services. A given enterprise may include one or more VPUs. A federation comprises two or more enterprises, each with one or more VPUs.

At a more detailed level, the two-axis model of a VPU supports its asset and supply chains through eight communications ports. Investors provide capital assets at port a_i (assets-in) that subsequently yield investment returns on port r_o (returns-out). Customers provide demands for goods or services on port d_i (demand-in) that are subsequently fulfilled on port s_o (supply-out). A VPU supports two subsidiary channels, one for subordinate (i.e., embedded) VPUs and one for supplier VPUs. Subordinate VPUs are allocated investment assets on port a_o (assets-out) that generate returns on port r_i (returns-in). Supplier VPUs receive their demands on port d_o (demand-out) and return their production on port s_i (supply-in). The horizontal flows may be expressed in terms of cost per unit ordered and delivered. The vertical flows may be expressed in terms of cost per asset deployed or returned. Costs are typically measured in net present value dollars.

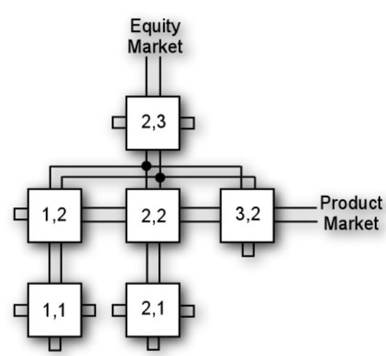

Figure 4 – Example Enterprise

This model allows VPUs to participate in the *production web* of a federated enterprise, the two principle threads of which are shown in figure 3. Each VPU is uniquely identified (i.e., named) by its indexed location vertically and horizontally in the web. Thus, `VPU[k,l]` is subservient in the investment chain to `VPU[k,l+1]`, and is a supplier to `VPU[k+1,l]` in the supply chain. Figure 4 depicts an example enterprise comprising a federation of six VPUs whose aggregation serves the equity and product markets. Indexing may be

[3] A system is viable to the extent that it is able to remain productive while adapting to changing environmental conditions.

interpreted as follows. Business unit `VPU[2,3]` operates three production plants, `VPU[1,2]`, `VPU[2,2]` and `VPU[3,2]`, that together act as a horizontal supply chain to product market customers. Vertically, production areas `VPU[1,1]` and `VPU[2,1]` are subservient to their respective plants, with the entire enterprise collectively serving the equity market through `VPU[2,3]`. An important question is one of governance – what is an efficient set of management controls to *allow* this enterprise to effectively and simultaneously serve these two evolving markets. VPU performance models are discussed in [4].

Enterprises today measure value production through interpretation of operating data produced via transactions in various financial management application subsystems (e.g., ERP) and stored in various databases. Figure 5 identifies major applications found in today's corporations. The five rings represent the level ("`l`") dimension in `VPU[k,l]`, where `l=0` represents production devices (people and machines); `l=1` represents production units; `l=2` production areas; `l=3` business units; `l=4` divisions or business areas; and `l=5` corporations. The radials are representative of primary domains of management C2 focus. The four quadrants represent the major application domains of production, infrastructure, conformance and asset management. And finally, the colored overlay regions represent application *suites* provided by vendors of enterprise software subsystems that are increasingly web-based and subject to adaptation to GIG standards [5, 10].

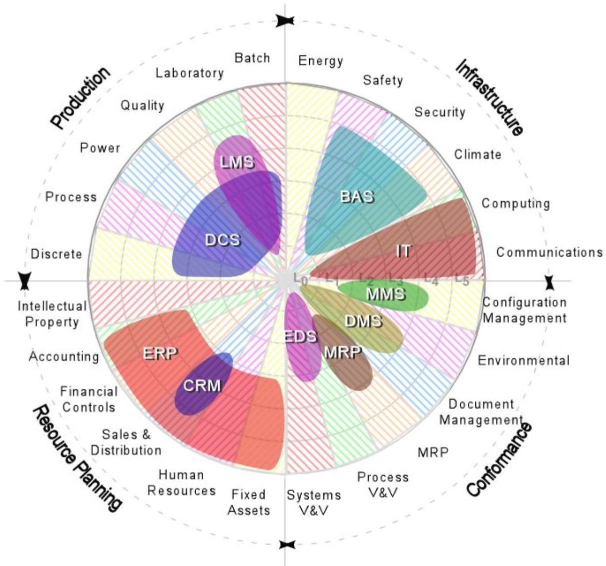

Figure 5 – Application Infrastructure

Internet standards and web services entered this world at rings 4 and 5, being initially external to core operations. In addition to inter-enterprise applications, the development of the grid is encouraging its intra-enterprise use within the inner automation rings where proprietary or more traditional communications and computing systems reside. Within this *inner sanctum* are found distributed control systems for manufacturing (DCS); building automation systems (BAS) for climate, safety and security; document management systems (DMS); maintenance management systems (MMS); laboratory management systems (LMS); and myriad incarnations of financial management systems (ERP), to name a few. Our particular interest is how to view these evolving applications as formal components in higher-level real-time command and control of inter- and intra-enterprise value production processes. For this we rely on the semantics of cybernetics, the science of communication and control in adaptive systems [1, 2, 3].

Federation Command and Control Services

Figure 6 is a model for *intelligent objects* capable of providing reactive and proactive command and control services. Such *adaptive controllers* have been studied and implemented in robotics, discrete, batch and continuous control systems for many years [2, 11]. In this context, the process to be controlled is one or more VPUs.

The lower loop of the adaptive controller represents the reactive (autonomic) management function, capable of sensing the state of a process as measured through one or more *sensors*. This input is presented to an internally maintained VPU *process model* through sensory signal processing or *sensory perception* logic. The process model qualifies the signal processing activity and accepts updates to the database(s) representing the current state and history of the process. *Behavior generation* logic produces appropriate next state outputs that are delivered to the VPU through the object's final controls, or *actuators*.

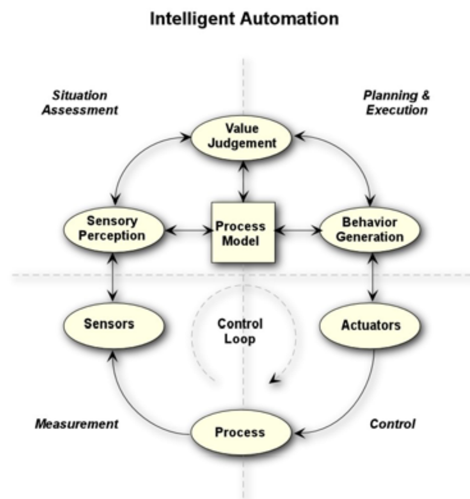

Figure 6 – An Adaptive Control

The upper loop of figure 6 represents the object's internal proactive management function, supporting volition and the adaptability and intelligence exhibited by the VPU. In addition to having access to sensory input, this *value judgment* function is able to communicate with other federated controls to declare its intentions and capabilities. It provides methods capable of making qualitative decisions and appropriately biasing the activities of both sensory perception and behavior generation. Taken together, these two interconnected and interoperable loops provide for specific degrees of *intelligent automation*.

Conclusions

Intelligent enterprise automation, especially adaptive C2, requires specification of and adherence to time constraints. As such, C2 regimes are inherently *real-time*, both within a given enterprise controller as well among cooperative controls spread horizontally and vertically within a federation. Timeliness, and it correlated assumptions about predictability of performance, underlie a system's ability to resolve current states and move to new states in timeframes required to maintain homeostasis – dynamic stability within changing contexts. As a consequence, imbuing the Internet and its web services with real-time semantics is at the core of the technical challenge.

This view of federated enterprise, and more importantly the "intellectual commons" upon which freely associative, high integrity and open commerce may be supported, requires the resolution of a number of legal, commercial and technical issues. These issues speak directly to requirements for next generation Internet and web services. Several infrastructural issues are currently being worked in community standards activities such as WC3, the Global Grid Forum, IETF and OMG. Of the activities in these working groups, the key elements leading to the success of distributed real-time control of federated enterprise include the following topics.

- Distributed real-time programming models [6]
- Distributed resource scheduling [10]
- Distributed system security through authentication, access control and accounting [8]

Important open issues remain in the areas of

- Standardized enterprise operating performance metrics
- Standardized enterprise binding services defining and monitoring "rules of engagement"
- Standardized enterprise alarm and event services for real-time control
- Standardized enterprise operating services for collaborative (coalition) C2 applications

References

1. Ashby, R., *Introduction to Cybernetics*, Chapman & Hall, Ltd, London, 1957, PDF version available by permission at www.echelon4.com.
2. Bayne, J., *Automation & Control in Large-scale Interactive Systems*, Proc IEEE International Symposium on Object-Oriented Real-time Computer Systems (ISORC), 2002.
3. Bayne, J., *^{Meta}Systems*, Proceedings of the Instrumentation, Systems and Automation (ISA) Society, 1996.
4. Bayne, J., *Value Production in Distributed Real-time Enterprise*, publication pending, December 2002.
5. Chinnici, R., et al, *Web Services Definition Language (WSDL) - Version 1.2*, WC3 Working Group Draft, July 2002, www.w3.org/2002/07/wsdl.
6. Clark, R., Jensen, E., Wellings, A., Wells, D., *The Distributed Real-time Specification for Java: A Status Report*, www.real-time.org, Revised January 2002.
7. DOD, *Joint Tactical Architecture – Version 4.0*, April 2001.
8. Foster, I., et al, *The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration (Draft 2.9)*, June 2002.
9. Lessig, L., *Code and Other Laws of Cyberspace*, Basic Books, 1999. See also www.cyberlaw.stanford.edu/lessig/.
10. Schwiegelshohn, U. and Yahyapour, R., *Grid Scheduling Architecture*, Working Draft 1, Global Grid Forum, October 2002.
11. Wooldridge, M., Reasoning About Rational Agents, MIT Press, 2000.