Three Big Challenges of Smart Grid Applications

And How the Blue Vector Platform Can Address Them

The time is now

Despite more than a decade of starts and stops, the Smart Grid shows more promise today than ever before. Market forces have combined to create a perfect storm of opportunity: Stimulus money, investment in green technology, and a drive toward energy independence have contributed to a business climate that is ready for Smart Grid deployments. Virtually every large IT vendor or appliance manufacturer has placed Smart Grid on its strategic agenda, and utilities have begun to discuss their plans publicly as well as place bets on technology providers and commission pilots.

Smart Grid investments need to enable multiple use cases and accommodate future growth and change.

The industry has contributed use cases ranging from automated metering to incorporating, storing, and managing renewables to managing home and building energy consumption, all under the umbrella of the Smart Grid. Any entity managing an energy grid – from large utilities to building complex managers – recognizes the importance of enabling multiple use cases with the same infrastructure, as well as ensuring that today's Smart Grid investments will accommodate future growth and change. Energy utilities and industry experts have called for standards-based platform technology from which to launch and distribute applications. Utilities aren't the only ones who require a platform; whether out of a desire to be more "green" or to save costs, businesses and consumers are becoming more aware of their energy consumption patterns and becoming more participative in the conservation process.

The three challenges of Smart Grids

Smart Grid applications must address the scale, mix of new and legacy infrastructure, and unification requirements common to all grids.

We believe that these factors – multiple audiences, many use cases, and market growth and change – point not just to the need for platform technology, but technology that addresses three significant characteristics of large, complex systems like Smart Grids. All implementers of Smart Grid applications will need to address the following three things common to grids:

1. **Scale.** All implementations need to deal with the scale at which grids operate. With each new wave of "smartening" the grid emerge control system problems that must span millions of end-points.

2. **Legacy.** No grid implementation can expect to completely replace existing infrastructure. The commercial feasibility of all solutions must encompass insertion and expansion of the solution into an existing framework. This requires flexibility in integration with existing software systems and the millions of devices that are already deployed and functional in the grid.

3. **Unifiability.** National grid infrastructures are a conglomeration of numerous different networks and sub-networks operated by many different companies, operating under many different types of public policy and business realities. Just as the Internet was designed to integrate across many types of networks, the next generation Smart Grid must have inherent architectural principles that will unite the diverse sub-systems that form nations' grid systems.

Addressing each of these characteristics is a challenge in itself. Addressing them simultaneously in a single platform requires a new architectural paradigm that enables distributed processing with centralized control and visibility. We at Blue Vector call this platform a **distributed control plane**.

The Blue Vector platform

Blue Vector offers an application development platform that serves as a distributed control plane for large, complex systems ranging from large video networks to global supply chain automation solutions. Deployed in mission-critical environments in Fortune 500 companies, our platform is based upon a unique architecture called Aspect Oriented Programming. It is this architecture and key technological capabilities that allow Blue Vector to simultaneously address the above three characteristics that make Smart Grid applications so challenging.

Blue Vector was born out of research and development performed at Xerox PARC for the Smart Dust Initiative

Born out of research and development performed at Xerox Palo Alto Research Center in the 1990s for the Department of Defense Smart Dust Initiative, Blue Vector's platform harnesses Aspect Oriented Programming, an architectural paradigm that increases modularity by allowing the separation of cross-cutting concerns and enabling them to be expressed separately and automatically unified into working systems[1].

Because Blue Vector's platform was originally developed specifically for large, highly distributed, heterogeneous, and ever-changing military sensor networks, it is a tremendous asset in the development of today's Smart Grid applications. Below is a simplified view and description of how our platform works.

A simplified view of Blue Vector's platform

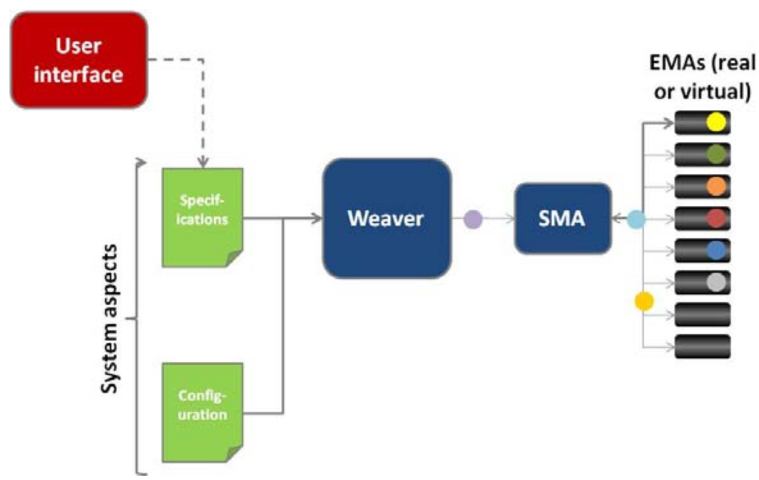

Blue Vector generates and distributes code automatically for fast, simple, and error-free applications.

In deploying the Blue Vector platform, users indicate system specifications separately from each other via a **user interface**. Those specifications are combined with configuration information in Blue Vector's **Weaver**. The Weaver uses a recursive process to solve all constraints, automatically generates code for each compute unit in the system, and provides the **Blue Vector SiteManager Appliance** (SMA) with instructions for where to distribute each code stub. The SMA then distributes each code stub to its intended location (to each **EdgeManager Appliance** (EMA) or other compute unit, as appropriate) and keeps track of where it is

[1] Wikipedia

and what it contains. When the Weaver process is used for re-configurations, the platform only updates those code stubs that are impacted by the change, leaving the rest of the system unchanged.

In Smart Grids, Blue Vector's platform can be inserted at multiple layers in a communications network not only to control or communicate with central- and end-points in, but it is flexible enough to be as loosely- or tightly-coupled as desired with both communications network as well as grid devices themselves based on performance requirements and device intelligence and connectivity.

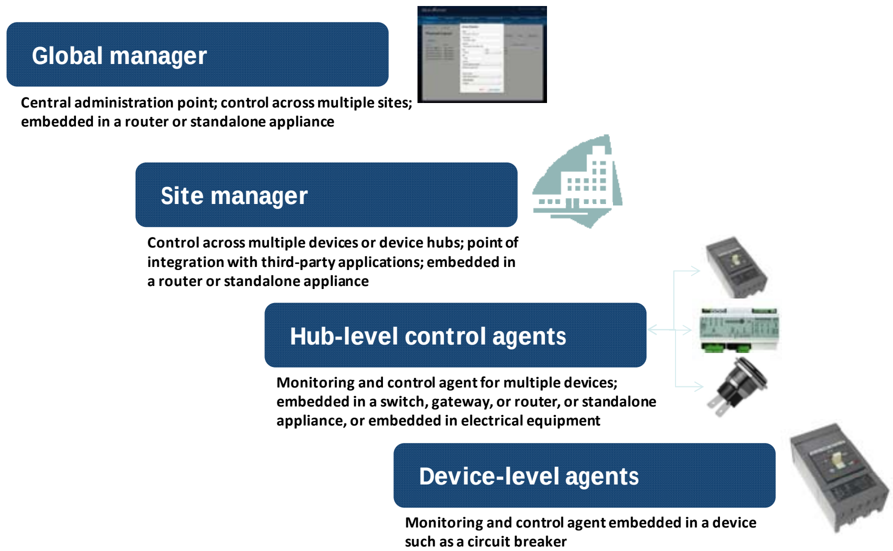

Blue Vector in micro-generation: a use case

Below, we illustrate a micro-generation use case because we believe it is representative of many Smart Grid applications and carries with it many common Smart Grid requirements (such as closed-loop device control and integration with third-party applications). It also represents a powerful illustration of the three challenges that must be addressed in Smart Grid applications – scale, legacy, and unification.

In the use case, there is a micro-grid that is coupled to a main grid by way of a point of coupling controller (PCC). In the micro-grid, there is a conventional power generator as well as a photovoltaic solar power system. There are inverters that convert the current after it is generated or before it moves into storage. There are relays that control the flow of energy to loads, or consumers of energy within the micro-grid (see architecture diagram below).

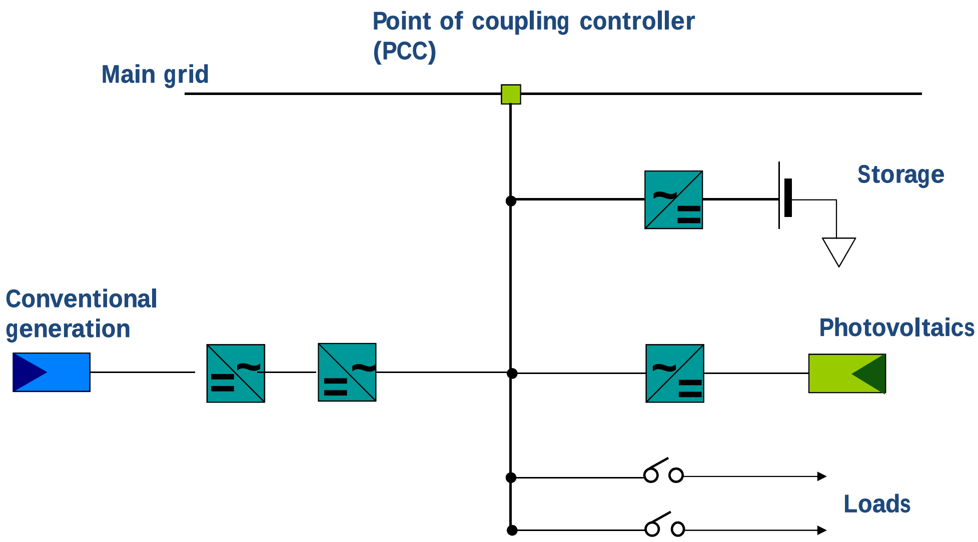

Integrating distributed generation into the main grid. Micro-grid operation typically relies on a hierarchical control algorithm. It requires supervisory control from the main electrical grid, and requires local control to effect changes internally. At every point, the micro-grid must be informed by supervisory control systems about various set-points and reference points, and it must act upon this information to suitably control the power sources within the micro-grid.

In general, the parameters being controlled are voltage, frequency, and reactive-power. This is essential for maintaining the stability of the grid, and optimizing power flow. In addition, the micro-grid must be able to operate standalone (in island mode).

In this scenario, the supervisory control system relies on a local controller to implement the control policy. The supervisory control system might typically reside at the sub-station or at the utility's control center and integrated with various legacy systems within the utility such as the SCADA repositories, the GIS, and other information sources. It communicates with the local controller (assumed here to be the PCC, which might be the home automation gateway or a building automation gateway) through a command set that will request changes to set-points, or provide reference values which would influence the local control algorithms.

Demand response: Optimization of energy cost. Within the micro-grid, there is a need to make automatic adjustments to energy supply in response to demand changes that are signaled to the micro-grid by the supervisory control through pricing signals or other types of commands.

The system can operate in different states wherein a particular set of energy supply policies is in effect. For example, if the micro-grid is servicing a corporate campus, during "normal" operations the policies in effect might enable user-driven energy consumption such as allowing users to turn on and off lights when they wish or run the air conditioning at the temperature they wish. However, when the state changes to "high-demand" or "emergency" operations, the PCC needs to switch to a different set of policies (prioritized and restrictive), which supersede those in effect during the "normal" state. The micro-grid may need to shed certain loads such as air conditioning in office buildings, while maintaining others, such as air conditioning in data centers.

In such a scenario, the PCC (which, again might be a home or building automation gateway), must be able to implement the policy locally by modifying different control parameters on the different loads on the micro-grid, acting as the load control transponder. In larger networks, there would be hierarchies of load control transponders that would similarly implement load control policies.

The solution needs to be able to understand a state change and send the necessary events to the system loads to have them change the set of policies under which they are consuming energy. It needs to distribute fine-grained policies to devices which control the flow of energy and even to the loads themselves in some cases. The solution needs to allow those policies to be cross-cut and prioritized, as well as integrate-able with utilities' information systems such as billing and reporting.

Local micro-control and islanded operation. One of the best means of maintaining power quality within a micro-grid is to implement a multi-agent PQ control algorithm that changes set-points at the inverters within the micro-grid. In particular, in an islanded operation, the micro-grid must be able to maintain stability internally by implementing both load-shedding as well as fine-grained control over the inverters in the grid.

Many inverters implement droop control as part of maintaining stability when no centralized control is available. So, for example, if the PCC were to go down during islanded operation, an ideal backup control strategy is to have each inverter switch to a droop-control policy for maintaining stability. This would require the controllers at each inverter to be smart enough to switch control policies when the environment around them changes.

A scenario like this requires for control policies to cover both grid-wide control and local micro-control. The solution needs to be able to cohesively integrate inverters that can actively participate in implementing a global control strategy.

How Blue Vector addresses the use case. Below is an architectural diagram indicating how Blue Vector solves the above examples. The shaded area indicates our distributed control plane, the span of control enabled by the distribution of our agent code (depicted by the colorful circles), which is generated automatically and centrally and contains logic that is specific to the local node.

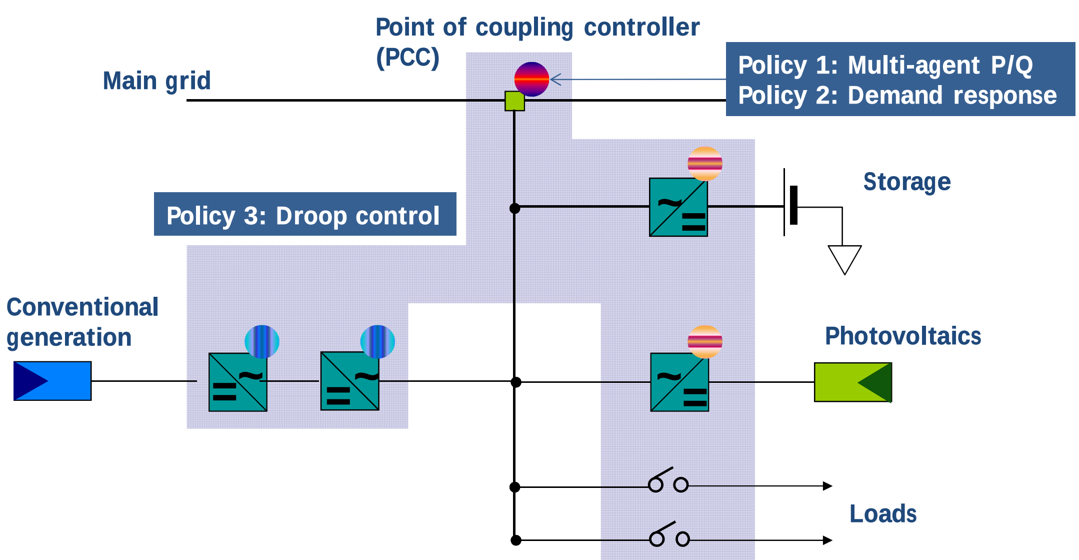

> Blue Vector agents can reside on the point of coupling control or on the inverters themselves and carry out system-wide or local policies.

Blue Vector's Weaver automatically generates agent code based on separately-specified inputs, including policies and processes for system and individual component control parameters, and demand response policies. The Blue Vector SMA distributes the agent code to the appropriate nodes in the system. In the micro-grid, the agents can reside on the PCC only or on the PCC and the inverters themselves, depending on system requirements and constraints. In the integrating distributed generation example, the code-images distributed to the PCC will incorporate all the policies necessary to integrate this micro-grid into the main grid. These policies will respond to supervisory control as well as implement local stability policies. In the demand response example, the code image at the PCC would include policies determined by utilities, consumers, and localities for conditions under which energy is delivered to various loads in the system, processes for doing so, and inverter and relay device drivers. In the local micro-control example, agents on the PCC and inverters would respond appropriately based on a coordinated or islanded operation policy.

Summary

We believe that creating and deploying Smart Grid applications represent a tremendous opportunity for business growth, environmental progress, and energy independence. But pulling it off is a big challenge given the nature of grids – their sheer scale, mix of legacy and new infrastructure, and the fact that they are amalgamation of interconnected networks.

We believe Blue Vector brings an important ingredient to the table in simultaneously addressing each of these challenges by providing an inherently scalable distributed control plane for enabling a multitude of Smart Grid use cases. We believe our unique architecture, solid platform purpose built for distributed systems, and solution competency will enable our partners and customers to build powerful and sustainable Smart Grids.