

Convergence of IoT, Edge and Cloud Computing for Smart Cities

Editor:

Mazin Yousif,
mazin@computer.org

Editor-in-Chief Mazin Yousif discusses what it will take to make smart cities a reality.

Cities require extensive underlying infrastructure to make them smart. But what does the concept of a “smart city” really mean? Does it mean a city designed for just one goal using standalone infrastructure—for example, an optimal transportation system, healthcare system, or safety system? Or, to qualify as a smart city, do we mean that all such services have been optimized to communicate and work together through fully-instrumented infrastructure? Of course, it all depends on who you talk to. I am more of the mind that all city services need to be optimized for a city to become truly “smart.”

Fully instrumenting any city to become smart is a monumental effort that is likely to take a very long time. If we are building a new city from scratch according to a comprehensive planned design, then design may take longer but execution will no doubt be easier and more likely to keep to a planned timeframe. Existing cities, on the other hand, are a completely different story, as augmenting an existing city to become smart is likely to disturb normal city life, likely to frustrate its residents, and likely to cause disruption to normal city services such as transportation, police and healthcare. After all, consider the turmoil and angst during a simple lane closure.

However, regardless of such disruption, smart cities and their interacting systems and ecosystems are not going to be forgotten. They are inexorable because of their many benefits, such as improved safety, reduced costs, and externalities such as air pollution from idling cars due to suboptimal traffic flows, enhanced social welfare, etc. However, inexorability does not imply simplicity; the infrastructure, applications, and security concerns of such a vision will require the state of the art in information technology and engineering design.

The basic components of any smart city are the devices/sensors/cameras and other information sources that collect data from physical assets such as roads, buildings, cars, buses, and other resources; distributed computing infrastructure (fog and edge) to locally analyze the collected data to impact local decisions; and a cloud, somewhere, to host the city data for further analysis, decision-making, and decision enactment—for example, examining current traffic and resetting traffic lights to smooth flow. These components cannot work together without comprehensive networking, including state-of-the-art cellular networks and distributed WiFi to support data transfers among the digital infrastructure, services, apps, businesses, and individuals in the city.

On the software side, there is a need for an extensible and scalable software framework that will be able to host all the software components—applications, middleware, analytics, intelligence and dashboards—that are basically the brains making the city smart. Note that I am not talking about one single installation of the framework, but rather multiple installations, depending on the complexity of the smart city’s infrastructure and services.

On the data side, a data architecture is expected to be substantially complex because there is a need to bring together and possibly harmonize data from the city’s many locations and services. This means many data connectors: message brokers to direct the right message to the right queue for processing; possibly hierarchical common data buses with a common data bus for one location in the city or one type of service (e.g., transportation and healthcare) in the city and then one parent common bus that brings all together; and more. I also mentioned earlier that the cloud needs to be “somewhere.” The reason is it needs to be able to maintain the compliance and privacy profiles of the data. In addition to architecture, policies for data ownership, rights of use, monetization, access, and trust will be necessary.

There is also a need for an extensive and comprehensive management infrastructure that will manage every device and digital aspect of the smart city. This management will maintain the integrity and security of digital infrastructure in a smart city and maintains the integrity of the data pertinent to that city. In particular, “fail-safe” systems are needed with warm failover to ensure the smart city systems operate without interruption. Additionally, due to a smart city’s extensive reliance on information systems, it will be a target for hackers, criminals, terrorists, nation-states, and sociopaths who want to either profit from disruption or the threat of disruption, or extract data in an unauthorized fashion.

Clearly it is necessary to have edge and cloud services operate jointly in any smart city deployment. There will also be many such edge deployments, depending on the size, diversity, and complexity of the city. There are likely to be multiple compute layers of “fog” between the edge and the cloud, for the aforementioned reasons. The cloud will be the place where all the data will be stored and further analyzed in combination with other types of data for different purposes. The recent story in New Haven of 100 overdoses in a park underscores the need for various smart city ecosystems such as healthcare, law enforcement, and transportation to interoperate, most likely via the cloud.

With all the above considerations in mind, we can definitely think of a smart city as one complex, enormous IoT deployment, with edge and cloud as the underlying building blocks. But more importantly and due to the complexity and diversity of services in a smart city, it will likely include software and hardware components from very different providers and technology vendors. The same is true for all the devices that will be used to sense, collect, and act on data. As a result, interoperability and open framework architectures—both hardware and software—will be essential for any successful smart city deployment. Such loose coupling, modularity, and heterogeneity will enable intelligence while also admitting incremental deployment and upgrades, best in class components, and cost optimization through market competition.

These topics are the theme of this issue. What makes this complex topic interesting is the need to satisfy the unique requirements of each city as it offers services based on its country, region, weather, culture, crime rates, laws, vision, and investment flexibility. I am thus confident that there is much research and development that needs to take place before we can really appreciate the full meaning and potential of smart cities. I would like to thank the guest editors of this special issue, namely Maria Fazio, Rajiv Ranjan, Michele Girolami, Javid Taheri, Schahram Dustdar, and Massimo Villari for all their efforts in producing this issue. I also urge the readers to read their guest editor introduction.

ABOUT THE AUTHOR

Mazin Younis is the editor in chief of *IEEE Cloud Computing* magazine. Contact him at mazin@computer.org.