White Paper

Driving the Internet of Space Things through Adoption of Commercial Technologies

Introduction

The adoption and use of cutting-edge commercial technologies in space applications have enabled the design of smaller, lighter, and more-powerful vehicles, lowering launch costs and providing greater mission capability than possible a decade ago. Smaller mission-capable devices called CubeSats have emerged, bringing space experiments and research into classrooms, smaller corporations, and institutions. CubeSats also broaden the types of missions performed for commercial and defense markets.

This Internet of Space Things (IoST) will continue to grow and expand as adoption of commercial systems and network technologies provides increased processing power, sensor capability, and interconnect fabrics that link collaborative space systems into massive orbital micro–data centers. Taking advantage of the modularity in commercial-off-the-shelf (COTS) systems, these IoST devices will replace many monolithic, proprietary, and special-purpose space systems to:

- further reduce costs,
- increase scalability and flexibility in mission design,
- and reduce development time needed to put platforms into orbit.

Ultimately, many space systems may consist of networked, modular functional units responsible for various tasks and functions that are docked into required mission-specific configurations. After missions are completed, the systems are broken down into modular elements again and are ready for reuse.

IoST devices have rigorous design requirements to be able to exist in space and will need to operate autonomously while collaborating with other systems in orbit. Scalable, orbital processor and data storage clusters will provide always-available analytics capability and advanced situational awareness for many commercial, federal, and defense use cases. This orbital "dynamic edge" will require a software fabric capable of supporting commercial hardware platforms and providing performant, scalable, and resilient node networking; advanced distributed computing capabilities; and seamless upgrading to support new use cases and mission profiles. Intel's Situational Awareness at the Dynamic Edge (SADE) is a software framework designed to meet these challenges and is ideally suited for the rapidly emerging IoST applications segment.

The use of COTS technology in space applications

Historically, there was very little COTS technology in space. With the drive toward putting more, smaller satellites of an IoST system into orbit, companies are looking to move away from proprietary to commercial technologies. They're opting for faster technology development at a lower price, with better availability and compatibility. They want to have a major technology partner with breadth and capability that understands modular open systems and can help them take advantage of solutions available in the commercial world.

Intel has demonstrated that it can provide the functionality required for space applications with COTS components. The first use of a commercial 32-bit processor in a NASA spacecraft was the Intel® 80386 and its peripheral support ICs, deployed in NASA's Small Explorer (SMEX) spacecraft that launched in 1992. The explorer had a five-year life-span goal but actually operated for nearly 20 years.[1]

Intel® products include components that are appropriate for mission operators and autonomous systems as well as technology that makes satellites more resilient to failure and capable of operating under extremely challenging conditions. Intel radiation disclosure information can be provided on a case by case basis under CNDA. Intel® FPGAs and Intel® Xeon® D processors can process radio signals received at the ground with high speeds. From processing and storage in LEO satellites to receiving ground stations, Intel has a family of products that are COTS based with open source tools for lower cost and exceptional performance.

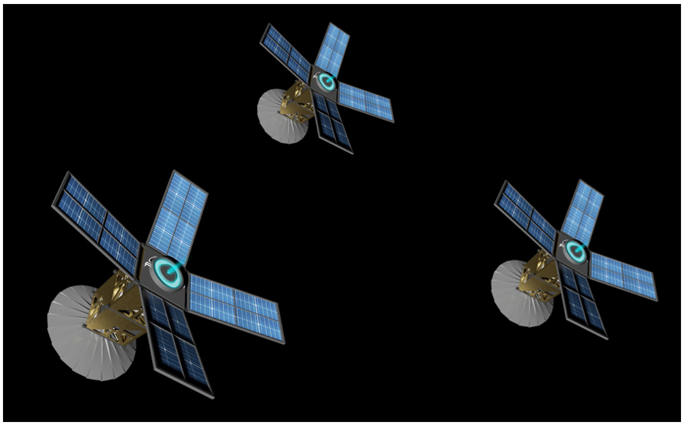

Figure 1: Heterogeneous CubeSat swarm.

IoST and the dynamic edge

IoST devices have numerous potential applications. For example, IoST devices operating in low Earth orbit can be deployed like drones to multiple locations around the globe to capture near-real-time multimodal sensor data for use by ground control and mission operators. They can be tasked to behave like a cluster of vehicles, providing customized capabilities with resilience to operational loss if one or more of the individual units are lost or damaged. They can also perform simple mechanical operations in space such as repairs of external systems on spacecraft. Or, through modular design, these IoST devices may physically displace a failed system by docking to a larger vehicle and becoming a replacement wireless sensor or data processing array.

The dynamic edge is where multiple clustered nodes may be vulnerable to communications denial back to ground stations. Connectivity to the cloud or central data center could be denied for numerous reasons—for example, due to a remote location limiting line-of-sight communications or active signal denial. To function efficiently, IoST systems need to be capable of operating autonomously at the dynamic edge, performing time-critical processing locally. They shouldn't have to rely solely on communicating with their ground-control system for computational support. The terabytes of information that would have to be sent and received would strain communications, and intermittent connectivity to the control system would hamper performance and introduce instability in the overall system. Rather, IoST devices need the ability to operate independently and to be configured, dynamically, into orbital micro data centers to process computationally significant tasks.

IoST design requirements

With the proper software infrastructure, COTS-based devices can be developed and deployed more rapidly and less expensively than they can be using today's custom devices. Moreover, the adaptation of commercial technology to space use by third parties could enable new features and capabilities for IoST devices to meet rapidly growing challenges.

Multiple low-cost IoST devices operating in coordinated, modular fashion can be used in a variety of exploration and defense/intelligence operations, removing the need to deploy costly monolithic satellite assets. Several key systems are required:

- IoST technology must have a software fabric for machine-to-machine and human-to-autonomous-system collaboration that extends independent mission operation beyond line of sight. This can reduce dependency on vehicle to ground-control communication to meet mission objectives.
- A networked multinode platform architecture could support redundancy and scalability in operational scenarios. This provides the ability to relocate relevant software workloads to the most-capable and -available platforms while providing for scalability of the cluster of nodes as needs demand. Properly configured and with synchronized communications and operations, a cluster can operate like a micro data center in space.
- Vehicles need near-real-time situational awareness and scene intelligence capabilities to ensure that they have the necessary environmental information to navigate safely and avoid potential threats. Scene intelligence provides a high-fidelity and contextually detailed view of any interesting volume of space for simulation and evaluation. Situational awareness and scene intelligence are cornerstones to scalable collaboration of thousands of heterogeneous Internet of Space Things devices.
- Finally, the IoST devices will require algorithms and logic enabling them to potentially operate in proximity and be capable of coordinated maneuvering using low-impulse propulsion systems. This will require a combination of multimodal sensor data, heuristic decision-making, and analytics bridged by low-latency communications. Vehicles will be capable of adapting and learning from operations in varying conditions and react appropriately to emergency scenarios.

As the scope of operation of IoST grows and greater distances separate command and control from the IoST nodes, autonomy and self-reliance will be increasingly important, along with a local communications fabric that provides resilience of operation. Allowing teams of smaller devices to coordinate workloads and share capabilities provides critical redundancy, supporting mutual vehicle survival and overall mission success.

Situational Awareness at the Dynamic Edge

A dynamic edge framework can interconnect and coordinate the computational and sensor data tasks of multiple IoST nodes. Intel's SADE is a software framework based on a commercial approach using a lightweight, low-latency messaging backbone. This backbone interconnects the multiple nodes with a common fabric for application deployment and intelligent execution. Upon the backbone, all microservices and applications are completely modular and containerized to provide isolation and portability from one node to another. This is beneficial for recovery and scalability purposes. All elements may exist on a single platform or may be fully distributed across a network-connected set of nodes.

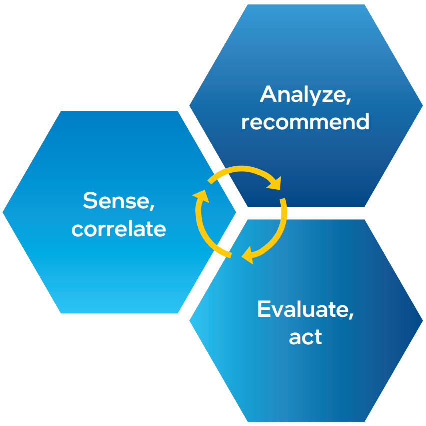

Figure 2: Situational awareness is derived from a continuous loop of operations.

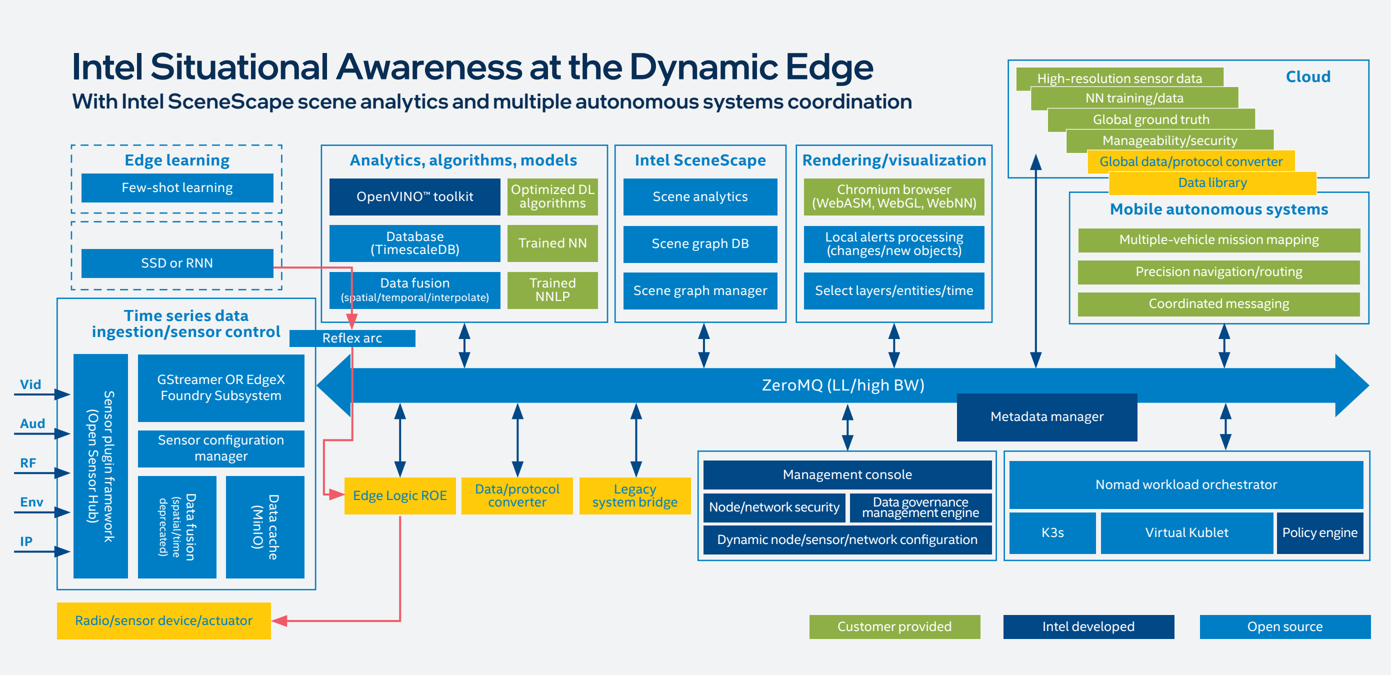

Figure 3: Proposed SADE functional block diagram.

SADE core functional blocks

Time-series data ingestion/sensor control and remote sensors identifies and manages multiple sensors and is responsible for initial data correlations. This may be used locally and/or transferred to other database layers for further analysis, consolidation with additional data, and generation of recommendations.

Analytics, algorithm, models module is a computationally intensive module that fuses sensor and situational data from multiple platforms and performs inference operations on the data stream. Like other containers, there can be several such modules anywhere on the network that provide services as needed and are retired when not required. Intel uses a robust Timescale database for the temporal-synchronization of multiple sensor sources. This creates ground-truth layers that show the evolution of situational awareness events. Updated ground-truth information and classified objects may be pushed toward the endpoints for immediate response and upstreamed to ground-truth databases in the command and control center or the cloud.

4D scene analytics module (see Figure 4) uses multiple sensor inputs and simultaneous location and mapping technology to synthesize a digital space. This digital space is modeled from the real world and positions detected objects within the volume. It provides an efficient means to evaluate near-real-time situational awareness using 3D and the ability to change perspectives and points of view as events unfold. In addition, the 4D scene graph database caches recent events, allowing for the "playback" of events within that region, the creation of hotspots and tripwires to trap incidents or intrusions, and the compositing of multiple scene graph databases to create larger volumes of reference. It's a powerful tool for simulating and understanding the context of an environment and comprehending objects and events within that space.

Rendering and visualization (RV) module is the end-user experience capable of displaying situational awareness information and managing application operations. The RV module is designed as a lightweight and portable web application with connected mode and offline cached data modes. Based on widely available browser technology and web standards, it allows for easy use and secure access from PCs, tablets, and smartphones. It also supports local hardware acceleration through standards such as WebGL, WebNN, and WebASM. With it, a central command and control location can monitor and manage the entire infrastructure. It also allows the apportioning of systems and assets to mission operators, providing them critical local control of situational awareness resources and data.

The cloud provides the backbone for the overall system, doing the heaviest of computational and data analytics tasks. The cloud also provides a protocol and translation bridge to interconnect the other branches and agencies needing access to the system or data. Commercially hardened Docker Swarm and Kubernetes allow containers of services and functionality to be moved from the cloud—on ground or in orbit—right to the dynamic edge. This may be useful to colocate computational and analytics functionality to where near-real-time sensor data is being captured, reducing latency and improving performance.

Mobile autonomous systems coordinate the behavior of multiple IoST nodes. An example is the intelligent swarming of vehicles in flight. This allows each node not only to evade or navigate independently as obstacles are encountered, but also to coordinate as a group toward a specific goal or to intelligently quarter and search a region.

Distributed application management and management console are two modules that constitute the control plane for the overall system. They are responsible for configuration and initialization, as well as for maintaining security and permissions, including reconfiguring to add or deprecate satellites, drones, other nodes, and networks. Application management is responsible for gathering system telemetry and, with the data orchestrator, ensuring that containers are targeted to the appropriate platforms for load balancing; it also ensures proper data locality for performance. The command center and cloud are the logical points for overall network and node control. However, it would be expected that a group of local platforms that are operating disconnected from the larger network would be resilient and capable of establishing an independent mesh. Being functional and independent, they would be able to continue executing on original mission directives and policies even when the system is disconnected from the cloud or command and control resources. This small mesh could reconnect and synchronize data and operations with the larger infrastructure at some later time or continue autonomous operation for extended periods.

Metadata manager ensures the efficient tagging and updating of information based on established definitions. It ensures efficient data access without necessarily moving physical data across locations. Acting as a fast index of metadata-processed information on the network, it minimizes data traffic. It does this by working as a switchboard pointing to new, relevant, and urgent data. It correlates data streams and files to identify relevance that assists in situational awareness analysis and supports the scene analytics module.

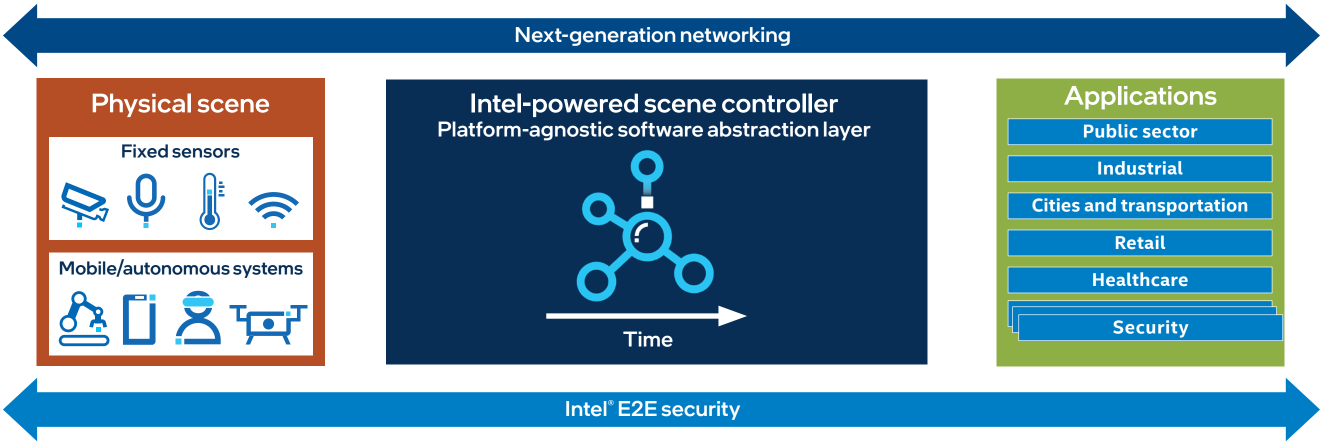

Figure 4: Multimodal scene intelligence creates customer value through a standards-based scene graph abstraction layer.

Summary

The emergence of the Internet of Space Things is still in its early stages, and use cases and architectural approaches from ground systems to space are evolving rapidly. Of key importance is how US commercial, defense, and aerospace organizations will evolve and deploy new generations of IoST devices rapidly enough to support new missions and use cases. The commercial sector helps ensure that innovation will occur quickly, while economics will tend to drive down costs of updating existing systems to make them more powerful and efficient. This dual benefit makes commercial off-the-shelf technology and modular open systems critical to IoST success.

A high degree of autonomous coordination and collaboration between IoST systems with other platforms and humans will require heuristic decision-making software. Intel's SADE software framework is designed to support the development and deployment of this. It enables rapid, modular application development and tasking along with operational redundancy and resilience across networked systems.

Learn more

For more information about the Internet of Space Things discussed in this white paper, please contact your Intel account executive, or email us at **IOTG-PublicSector@intel.com**.

---

1. LaBel, Kenneth A, NASA, "NASA Past, Present, and Future: The Use of Commercial Off the Shelf (COTS) Electronics in Space," May 2017. https://nepp.nasa.gov/files/28843/NEPP-CP-2017-LaBel-SEE-MAPLD-COTS-Presentation-TN42792.pdf

1221/LM/CMD/PDF