Taking the Tactical Cloud with You

Small Form Factor, Modular Data Centers at the Edge of the Battlefield

In order to achieve and maintain warfighting overmatch, coordinate deployed forces and enable new capabilities, the US Army, Air Force, and Navy are actively looking to new programs such as Joint All Domain Command and Control (JADC2) to ensure warfighters have maximum situational awareness. These programs will deliver a variety of compute and bandwidth intensive technologies, increasing the use of big data analytics, artificial intelligence/machine learning, and video for example, using common technical standards, APIs and data formats to deliver the command and control information that warfighters need to coordinate their activities.

The software needed to run these new capabilities is increasingly being developed to rely on the cloud, which itself might reside in a variety of data centers, ranging from large commercial services, such as Amazon Web Services (AWS) GovCloud and Microsoft Azure Government, to the DoD's Regional Hub Nodes (RHN), which are located in five separate strategic regions and used by deployed Marine Corps and Army units to access transport information from theater tactical networks around the world.

As forces develop ever greater dependence on cloud-based services, the impact of denied wide area access, (known as Disconnected, Intermittent, and Limited (DIL) environments) to the cloud, due to contested communications through electronic warfare (EW), or diminished access due to reduced bandwidth must be addressed. The key to delivering advanced new tactical capabilities and ensuring their continued availability when there's no access to the cloud is to vastly increase the capability, speed, and mobility of networking and compute at the tactical edge of the battlefield. This will make possible the replication of critical data and services in mobile, distributed clouds, moving information resources closer to the edge of the network.

To keep up with the proliferation of compute and bandwidth hungry applications relying on video, sensor data and AI/ML, to support battlefield operations, processing needs to happen locally at maximum speeds. To enable the next generation of situational awareness capabilities, the DoD is looking to deploy distributed processing and data replication between remote computing nodes and the central cloud This will provide continuity of operations in the case of network outages or low bandwidth.

Examples of current cloud replication programs include the US Army Cross Functional Team's CPI-2, which is currently developing prototypes of its Command Post Support Vehicle (CPSV). Essentially a mobile data center, CPSV is a rugged truck that carries a small data center's worth of servers, that in future iterations could support a local cloud.

Frame Stock Footage/Shutterstock.com

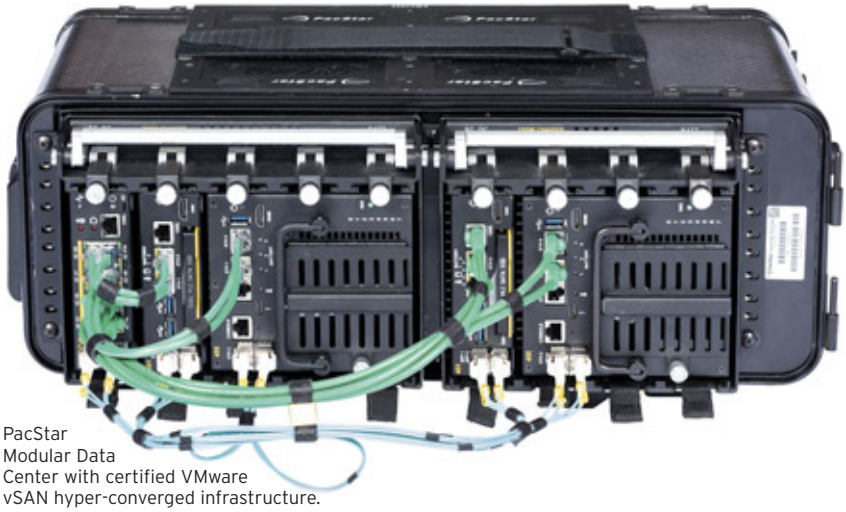

PacStar Modular Data Center with certified VMware vSAN hyper-converged infrastructure.

Likewise, the Army Future Command has published a roadmap that specifies requirements for robust cloud capabilities. The US Air Force is developing the Advanced Battle Management System (ABMS), a federated cloud system that will provide secure processing from its CloudONE security cloud already running in Amazon AWS and Microsoft Azure government clouds. ABMS also defines a local cloud, EdgeONE, to provide continued security in the event that communications with CloudONE are disconnected. And the Navy recently issued an RFP to industry to support Manned-Unmanned Air Vehicle Team Tactical Cloud analysis. The RFP includes an approach for providing remotely deployed cloud or processing services in case the tactical unit becomes temporarily disconnected from the tactical network.

To realize the capabilities that the mobile cloud can bring to the tactical edge, industry will require a new class of rugged, fieldable, network processing and data storage solutions that can provide tactical and expeditionary teams access to all of the data and compute resources they need. These systems will need to be able to provide cloud-like services, maintaining high operational availability of applications and data, even in DIL environments where WAN connectivity is not assured. What's more, mobile cloud systems will need to integrate seamlessly with leading public, government and private cloud providers.

Earlier efforts to bring these levels of performance and service to the battlefield involved trying to deploy huge sets of standard datacenter equipment, installed in shipping containers such as the Army "Container Express" or CONEX trials. These systems proved to be too large and hard to transport, significantly limiting their mobility.

In addition, these standard solutions are power hungry, which adds and additional burden on mobility. At the edge of the battlefield, fuel is both difficult and dangerous to transport, not to mention, very expensive. Higher performance, small form factor servers, combined with new approaches to distributed processing, offer the potential to move smaller slices of processing even closer to the edge, improving mobility.

Ruggedness is also a major factor when it comes to supporting mobile applications. A mobile data center should be designed to meet military environmental standards, such as MIL-STD-810, to ensure that the system can operate optimally when exposed to the harsh temperatures, vibration, shock and EMI typical of battlefield conditions. Standard 19" rack mount datacenter equipment is not designed to withstand harsh environments and not designed to operate on the move – leading to concerns about equipment failure and loss of availability of mission critical communications.

To meet the requirements for mobility and integration into the types of platforms on which they might be mounted, a mobile cloud solution will need to provide high reliability compute and storage in a small size, weight, and power (SWaP) optimized form factor, that is designed for mission-critical applications and able to support large volumes of data. The good news is that solutions that can deliver the high-density compute, storage and networking infrastructure needed to handle such large loads of data are available today. These modular systems can be optimized for program needs, so that the number of CPU cores, GPU cores or size of solid-state storage, can be maximized depending on program needs.

When these modular datacenters are based on industry standard processors and Nvidia GPUs, they will be compatible with a wide variety of applications and can meet the needs of a vast array of C5ISR use cases including data gathering, analytics/AI, and situational awareness. That means that this new class of modular datacenter can also support emerging distributed processing, storage/replication and hyperconvergence infrastructure software needed to deliver mobile cloud services.

An example of a modular datacenter available today that can support mobile cloud applications at the tactical edge is Curtiss-Wright's PacStar Modular Data Center (MDC), a COTS-based, modular, tactical and expeditionary, rugged data center capable of hosting cloud/storage, AI, and analytics applications. It uses proven small form factor modules for compute, storage, and networking functions with industry leading reduction in SWaP. This system can be deployed dismounted, in FOBs, command posts, ground vehicles, and aircraft, as well as in upper echelons – for military, intelligence, law enforcement, and Homeland Security use. Depending on the specific use case, PacStar MDC configurations can include a mix of compute modules, storage modules, and GPU modules, along with the company's switching/routing modules.

With the advancements such as this in high capacity, rugged, small form-factor hardware, it's now possible to deploy emerging *software infrastructures* to automate distributed processing and storage,

communications, and cloud replication required to ensure situational awareness is maintained to the edge of the battlefield. Large enterprise software companies are rapidly developing new technologies to ensure that data and applications can move seamlessly from the cloud to the edge and back, with little operator intervention. Several key technologies enabling this include application virtualization and containerization, network virtualization, hyperconverged storage infrastructure.

Application virtualization and containerization are technologies that decouple applications from the underlying hardware, allowing multiple applications to run securely on a single server, optimizing SWaP by reducing the number of servers required to deliver needed processing. Both technologies also enable applications to be replicated or moved from server to server, and from cloud to edge, balancing the availability of computing resources or minimizing latency over network connections. For tactical organizations, this offers the potential to move applications close to the warfighter, reducing processing delays and providing processing even when disconnected.

Hyper-converged storage infrastructure (HCI) technologies decouple the storage of application data from hardware, while also eliminating reliance on legacy Network Attached Storage (NAS) or Storage Area Network (SAN) architectures. These new technologies are foundational for replication of data between cloud and edge processing – enabling warfighters to take copies of cloud data into theater, and replicate changes to data over DIL connections when appropriate.

HCI also provides local data replication, for deployed organizations with a need for high reliability, ensuring that data is available in-theater even in the event of a server or disk failure. New advances in these technologies are rapidly improving the ability for tactical organizations to replicate data between any of the major cloud providers, and moving applications between the edge and cloud, providing maximum mobility and maneuver options for our warfighters.

The benefits of a converged compute/storage/networking modular datacenter with modern software infrastructure at the edge of the network are numerous. It can support a diverse array of use cases when operating in DIL environments, including hosting situational awareness, mission command and C2 applications, supporting SIGINT, HUMINT and IMINT data gathering and analytics workloads, as well as emerging IoT and sensor fusion-based applications. This is ideally suited to ensure the new all-domain situational awareness-driven warfighting doctrine will succeed while addressing the DIL reality of communications at the tactical edge.

PacStar
Modular Data Center with
high-capacity storage and Nvidia GPU-enhanced servers.

PacStar 400-Series Network Package on Joint Light Tactical Vehicle (JLTV)

This article was written by Charlie Kawasaki, CTO, and Dominic Perez, VP of Systems Engineering, Curtiss-Wright Defense Solutions, PacStar (Portland, OR). For more information, visit ***http://info.hotims.com/79418-503****.*