

Protect Sensitive Data with a Confidential Computing Solution

In collaboration with:



Tech Note by

Todd Mottershead

Todd.mottershead@dell.com

Seamus Jones

Seamus.jones@dell.com

Brian Porter

Brian.porter@intel.com

Krzysztof Cieplucha

krzysztof.cieplucha@intel.com

Summary

There are multiple considerations to take into account when deploying artificial intelligence and machine learning environments. This paper serves as a discussion and suggestion as to the possible hardware configurations to achieve a server infrastructure deployment that is secure and can grow with your increased need.

Enterprises in most industries are applying artificial intelligence/machine learning (AI/ML) to data. However, data privacy and sensitivity issues are preventing the use of AI/ML in health and financial sectors. This data cannot be shared, and it is limited to on-premises usage. Although this data must be protected from exposure to unauthorized parties, it is a valuable resource that could lead to groundbreaking discoveries and innovation in areas such as pandemic response, anti-money-laundering tactics and human trafficking response and prevention.

Confidential computing offers a way to expand the utility of such data while also keeping sensitive details sequestered and private. Dell EMC™ PowerEdge™ servers, built on 3rd Generation Intel® Xeon® Scalable processors, are available for the first time with confidential computing. A key feature is Intel Software Guard Extensions (Intel SGX), which provides an extra layer of hardware-based encryption in memory that helps protect data while it is being accessed. With Intel SGX, organizations can access and use multiple expansive datasets for AI applications, leading to greater insights. Intel SGX also helps ensure the integrity of the AI app to protect against intrusion, and it provides increased integrity to the platform whilst helping satisfy sovereignty requirements.

Key Considerations

- **Decrease attack surfaces.** Dell EMC PowerEdge servers based on Intel technologies expose fewer attack surfaces to hackers by making use of Intel SGX. The entire solution is built on an architecture that provides multiple layers of hardware- and software-based security. Security settings can be tailored to sequester only the most sensitive, privacy-protected data on Intel SGX to optimize performance.
- **Accelerate inferencing.** Intel performance optimizations that are built into Dell EMC PowerEdge servers can help speed AI inferencing. These optimizations include Intel and open-source software and hardware technologies, such as Intel Deep Learning Boost (Intel DL Boost) with Vector Neural Network Instructions (VNNI) for AI acceleration, Intel oneAPI Deep Neural Network Library (Intel oneDNN), the OpenVINO™ toolkit and optimized versions of TensorFlow™ and PyTorch®.
- **Secure ecosystem.** By creating a confidential computing environment with Intel SGX, ecosystem partners can improve the security of data collaboration between organizations that might have a need to keep data confidential and to protect intellectual property, even if a lack of trust exists between the parties. Secure ecosystem partners enable code and system integrity with Intel SGX for a growing list of AI solutions.

Available Configurations

	Base Configuration	Plus Configuration (More Memory for Larger Workloads)
Platform	Dell EMC™ PowerEdge™ R650 servers, supporting 10 NVM Express® (NVMe®) drives (direct connection with no Dell™ PowerEdge RAID Controller [PERC]), 1 RU	
CPU	2 x Intel® Xeon® Gold 6348 processor (28 cores at 2.6 GHz) with 64 GB/CPU Intel® SGX enclave capacity	2 x Intel® Xeon® Platinum 8368 processor (38 cores at 2.4 GHz) with 512 GB/CPU Intel® SGX enclave capacity
DRAM	256 GB (16 x 16 GB DDR4-3200)	512 GB (16 x 32 GB DDR4-3200)
Boot device	Dell EMC™ Boot Optimized Server Storage (BOSS)-S2 with 2 x 480 GB Intel® SSD S4510 M.2 Serial ATA (SATA) (RAID1)	
Storage adapter	Dell PERC H755N front NVMe RAID adapter ⁱ	
Cache storage (optional)	1 x 400 GB Intel® Optane™ SSD P5800X (PCIe Gen4) or 1 x 375 GB Intel® Optane SSD DC P4800X (PCIe Gen3) ⁱⁱ	
Capacity storage	1 x (up to 9 x) 3.84 TB Intel® SSD P5500 (PCIe Gen4, read intensive)	
Network interface controller (NIC)	Intel® Ethernet Network Adapter E810-XXV for OCP3 (dual-port 25 Gb)	

Learn More

Learn more **about** secure AI inferencing:

- [“Intel and Consilient Join Forces to Fight Financial Fraud with AI”](#)
- [“Intel SGX Enables Magnit to Create a Trusted Computing Environment”](#)
- [“Spain’s largest hospitals connect for federated learning”](#)

Contact your Dell or Intel account team. [1-877-289-3355](tel:1-877-289-3355)

- Read the Principled Technologies report: [“Reap better SQL Server OLTP performance with next-generation Dell EMC PowerEdge MX servers.”](#)
- View the Principled Technologies vSAN [infographic](#).

ⁱ An NVM Express® (NVMe®) RAID adapter is optional, but it is recommended for configurations with a large number of capacity drives.

ⁱⁱ Cache storage is optional. Intel® Optane™ SSD P5800X drives are recommended when available, but the previous-generation Intel® Optane SSD DC P4800X can be used otherwise.

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.