Futurum®

DELL Technologies

Responsible Artificial Intelligence: Dell Cyber Resilience for AI Applications

OCTOBER 2025

Introduction

The introduction of generative AI applications, such as ChatGPT, has led to rapid acceleration of AI application development, both generative and otherwise. Many organizations have initially focused on understanding the basic hardware and software requirements of AI, and initial deployments are often limited in functionality and scope. As AI development continues, however, AI applications will become mission-critical. As a result, IT organizations must be aware of the associated cyber resilience considerations.

AI provides a powerful tool that can solve complex challenges, optimize workflows, and drive impactful decisions across a wide range of industries, including healthcare, finance, transportation, and manufacturing. These industries will increasingly leverage AI for mission-critical use cases that can have a direct impact on business operations and financial outcomes. In some cases, AI's impact can extend to individuals' wellbeing, such as AI-assisted medical diagnoses, AI-powered workplace hazard detection, and self-driving vehicles.

The importance of cyber resilience for mission-critical workloads is not a new concept, and IT organizations are familiar with the approaches and best practices required. While many existing cyber resilience practices still apply, AI applications have specific data protection and security requirements to consider. IT organizations must recognize the increasingly critical nature of these new AI applications and understand the unique requirements for protecting them.

Protecting Data Sources

As with other critical workloads, a fundamental aspect of protecting AI applications is protecting the underlying data sources. While advances in AI technology are leading to powerful AI tools, AI is ultimately reliant on the underlying data that is used to train and augment models. Achieving powerful AI models capable of powering mission-critical applications often requires significant capacities of data, often from various data sources. To protect AI applications, there are several data-related aspects for IT operations to consider.

AI models rely on data used for model training, which ultimately leads to the accuracy and the overall usefulness of the model. To create powerful AI applications that fit their individual requirements, organizations will utilize their own private data to train or fine-tune models, often from many different data sources. IT organizations should understand which data sources are being utilized to train their AI models and verify that these sources have backup and recovery policies in place to ensure data availability for the training process. When considering an organization's existing data, it is important to check that existing cyber resilience policies meet the required service levels for recovery point objectives ([RPOs] or the amount of data loss that can be tolerated) and recovery time objectives ([RTOs] or the length of time that it takes to recover) for AI model training.

Alongside the need for training data availability, securing access to training data is also crucial. The outcomes of AI models are a direct result of the data they have been trained on. Modifications to training data, known as data poisoning, changes the resulting model, leading to inaccurate or unwanted results. Training data changes may be caused intentionally by malicious actors or accidentally. In either scenario, modifying the data can significantly disrupt an AI model, and as a result, significantly disrupt key decisions and operations. To protect against data poisoning, organizations should secure and limit access to data used in model training with role-based access control (RBAC), multi-factor authentication (MFA), and other data privacy policies.

While actual model training happens over a limited period of time, the training data used may require protection for an extended period. Organizations may have requirements to ensure that data used in training is maintained for use in future model training or for auditing purposes. For long-term storage requirements of inactive training data, organizations should leverage data archival processes to maintain data cost effectively while ensuring data is secure and available for retrieval if needed.

While protecting training data should be a key priority in the implementation of AI applications, there may be additional data sources that also require protection. A major trend in AI is the use of Retrieval Augmented Generation (RAG) to supply AI models with additional context during inferencing. RAG is highly beneficial as it allows AI models to leverage additional data to generate more accurate outcomes, including those that may be specific to a particular industry or organization, without additional training. Utilizing RAG, however, involves additional data that also requires protection. Unlike model training, inferencing is an ongoing process across an AI model's lifespan. Just as with training data, RAG data must be kept available and secure. If data is unavailable to be queried during inference, the model will be unable to provide the additional context and accuracy that RAG provides. Similar to training data, unwanted modifications to this data can lead to inaccurate or undesirable results to critical AI applications. Organizations should ensure data used for RAG is backed up and recoverable and has proper access control mechanisms in place.

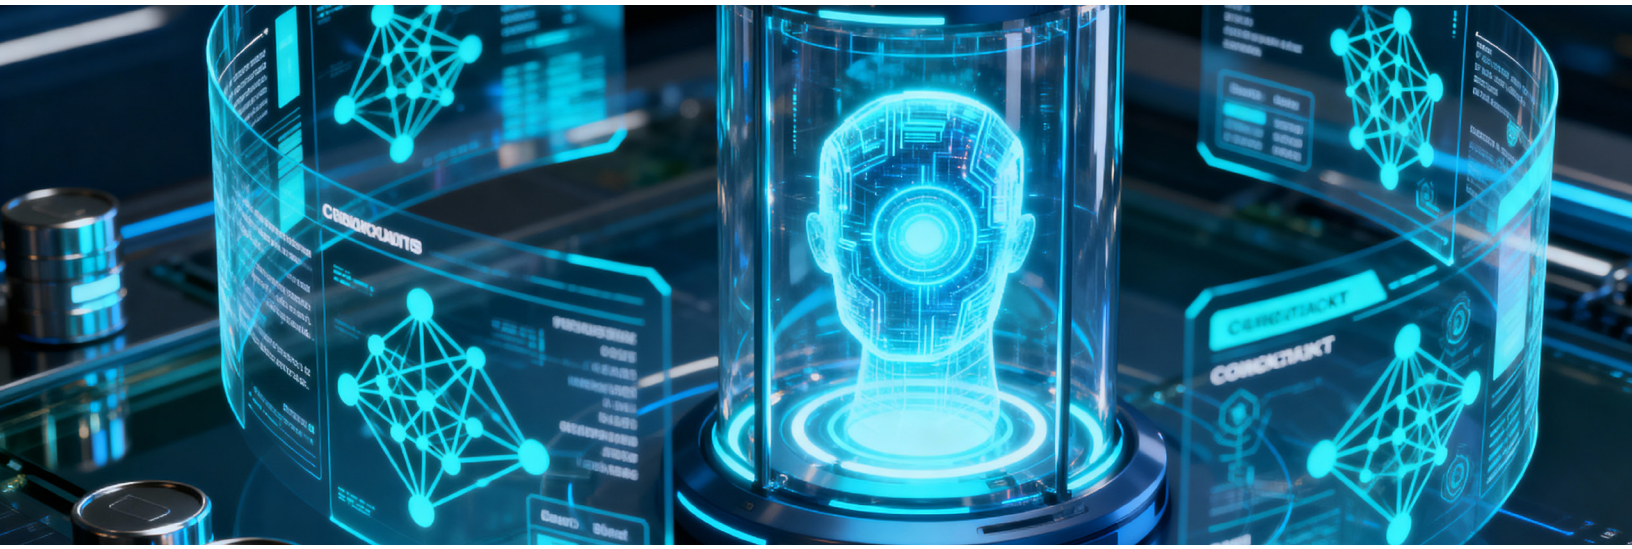

Protecting AI Models

Along with the protection of data that underpins AI applications, it is crucial for IT organizations to recognize the importance of securing and protecting their AI models. Trained models are the core of AI applications, and as these applications become tasked with driving operations and decisions that are increasingly sophisticated, the protection of the underlying models will become a growing priority for IT. Just as with data sources, IT organizations should consider the protection and security of their models both during training and deployment.

The training process is a crucial step in developing AI applications, as it is responsible for determining model weights that ultimately drive AI's ability to solve problems. The time and complexity of the training process can vary significantly, depending on several factors including the model architecture, the amount of training data, the target solution, the hardware utilized, and whether the model is undergoing full training or a simpler fine-tuning process. Time spent training models can be significant, ranging from a few hours to multiple weeks, and any disruptions to this process can result in a significant setback. During training, the intermediate state of a model can be saved and stored as a checkpoint. To avoid any significant disruptions, checkpoints of AI models should be created at regular intervals to ensure progress can be resumed if any errors do occur. Creating and retaining secondary copies of model checkpoints can be useful in certain situations including future AI model training, or if the model needs to be rolled back to a previous state.

Protection of AI models, however, is not limited to training. Deployed AI models and the associated application should be protected as any other mission-critical application. To do so, organizations should have backup and recovery policies in place for these applications and ensure that the RPOs and RTOs meet business requirements.

While, in many cases, protecting critical AI applications is similar to protecting any other business-critical application, the critical role of AI models introduces unique considerations for cyber resilience and cybersecurity. The output of a neural networkbased AI model is determined by a large number of weights that are configured during training. Unlike traditional programs, which operate according to coded logic, AI models rely on model weights that operate as a black box in which the inner workings are not easily decipherable. As discussed previously, modifying the training data is one vector that attackers could take to influence the output of an AI model. Similarly, the model could be compromised by changing the model's weights, even without access to training data, if a malicious actor has access to the model itself.

By targeting model weights, as with targeting training data, malicious actors could severely impact the results of a model, rendering a critical AI application useless or even altering it to perform an opposite task. Due to the black box nature of AI models, changes to model weights may be more difficult to detect and understand compared with changes in a traditional application in which the developers are familiar with the code and intended logic. To defend against such attacks, IT organizations should ensure that access to AI models is limited and that any alterations or anomalies can be detected. Additionally, it is crucial that a model can be safely and quickly restored if an issue does occur.

Data Privacy, Compliance, and Regulations

An additional consideration is that AI applications must meet regulatory and compliance requirements surrounding the use of private and sensitive data, both for model training and inference. IT organizations should be aware if personal data is used in their training data sets. Personal data should be identified as such and stored in compliance with any regulations. The requirements to identify and secure personal data are not unique to AI applications, and IT organizations should be familiar with the techniques and processes involved. A differentiator for AI, however, is recognizing that this data has been used to train an AI model and ensuring that the model will not return such data as an output. IT organizations should ensure that there are controls in place to limit an AI application's output to include only data to which users have access.

Data privacy concerns additionally extend beyond model training. Sensitive data, such as personally identifiable information (PII), may additionally be used during inferencing as inputs. This may include healthcare-related applications, facial recognition models, or AI-powered job screening solutions. Applications taking in such data must be capable of detecting such sensitive data and storing it securely.

Certain models, along with their inputs and outputs, may also be subject to heavy regulation in which copies must be maintained and retrieved if needed. Organizations with these types of regulatory requirements must be capable of securely storing, managing, and retrieving their AI models and data for an extended period of time.

Protecting AI Solutions with Dell Technologies

Dell PowerProtect Data Manager software and PowerProtect Data Domain systems are well-suited to help organizations to meet the cyber resilience needs of their AI applications.

Dell PowerProtect Data Manager is a cyber resilience software offering that provides backup, operational recovery, and disaster recovery for a variety of sources. It offers a number of capabilities that are well-suited to the requirements of AI applications specifically. These include the ability to create application-consistent backups, which capture the application state along with data, including factors such as ongoing training progress or various model versions, to facilitate faster and more reliable restoration in case of failures. This functionality is important considering the unique data and models that underlie the AI application's functionality. Backup orchestration can also be automated to ensure that AI data and its current state are captured regularly, avoiding the loss of data and time and effort spent developing underlying models.

The solution additionally offers Kubernetes and container support, making it well-suited for the many new AI applications that will be containerized. In some instances, AI data and models may be spread across on-premises, cloud, and edge environments. For these scenarios, PowerProtect Data Manager provides cloud-agnostic backups, allowing organizations to back up these distributed environments with management from a central console.

Complementing PowerProtect Data Manger, Dell PowerProtect Data Domain systems can be leveraged to provide secure and scalable storage infrastructure for AI workloads. Notably, Data Domain key functionality is the use of software and general-purpose CPUs to deduplicate data inline during transfer to a disk storage target that is connected to the Data Domain controller. This feature allows for compelling deduplication rates (up to 65x, according to Dell) while minimizing impact to backup and recovery performance. This approach goes far in allowing the large datasets that are common to many AI applications to be stored cost-effectively. Data Domain deduplication capabilities not only save storage space but also accelerate data recovery times. When restoring backups for AI applications, faster retrieval times minimize downtime and allow the organization to resume training or development processes quicker. Data Domain also includes security features, such as encryption at rest and in-flight, safeguarding an organization's sensitive AI training data and models from unauthorized access. Additionally, its redundant architecture ensures data availability even during hardware failures. This reliability is crucial for maintaining the integrity of AI assets. Data Domain tiering and archive capabilities can also be leveraged to assist organizations in maintaining AI training data, AI models, and associated application data with requirements for long-term retention.

In terms of procurement, it is material to note that PowerProtect Data Domain is a target storage system that can be used with PowerProtect Data Manager software or a number of third-party backup software offerings. Additionally, PowerProtect Data Manager can also be deployed as part of a fully integrated appliance consisting of both PowerProtect Data Manager software and PowerProtect Data Domain.

Final Thoughts

Advances in AI technology, spurred by new generative AI applications, is leading to a wave of AI application development across enterprise organizations. While initial AI applications have largely been experimental, organizations are increasingly leveraging AI in ways that are vital to business operations and outcomes. As the status of AI applications evolves from low stakes proof-of-concepts to mission-critical applications, IT organizations must address the associated cyber resilience considerations. In many cases, the requirements for protecting critical AI applications are similar to protection of other critical workloads; however, there are characteristics unique to AI that require additional consideration. IT organizations must recognize the growing importance of AI applications, understand their protection requirements, and implement the appropriate cyber resilience solutions.

Important Information About This Report

AUTHORS

Mitch Lewis
Research Analyst | The Futurum Group

Krista Case
Research Director | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DELL Technologies

ABOUT DELL TECHNOLOGIES INC

Dell Technologies Inc engages in designing, developing, manufacturing, marketing, selling, and providing support for information technology infrastructure such as laptops, desktops, mobile devices, workstations, storage devices, software, cloud solutions, and notebooks.

Futurum®

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum®
CONTACT INFORMATION: The Futurum Group LLC | futurumgroup.com | (833) 722-5337