



# **NetApp AI Pod Mini for ERAG - Infra Readiness Checklist**

NetApp artificial intelligence solutions

NetApp  
January 29, 2026

# Table of Contents

NetApp AI Pod Mini for ERAG - Infra Readiness Checklist . . . . .	1
Business & Use Case Readiness . . . . .	1
Technical & Infrastructure Requirements . . . . .	1
Software & Ecosystem Alignment . . . . .	1
Data Governance & Security . . . . .	1
Operational & Support Considerations . . . . .	2
Commercial & GTM Alignment . . . . .	2

# NetApp AI Pod Mini for ERAG - Infra Readiness Checklist

This document outlines a comprehensive infrastructure readiness checklist for NetApp AI Pod Mini for Enterprise RAG, serving as a pre-deployment reference.

## Business & Use Case Readiness

- Is the solution aligned to line-of-business outcomes (e.g., productivity, customer service, legal, healthcare, manufacturing, public sector)?
- Have you estimated Time to First Token (TTFT) and latency needs for your LLM workloads?
- Do you know the expected concurrency/user load (e.g., 32 concurrent users per 2-worker node for RAG)?
- Have you identified the primary AI/GenAI workloads (RAG, inferencing, fine-tuning, departmental LLMs, vector DB integration)?
- Are you evaluating CPU-based AI options (OPEA, Intel Xeon) versus GPU alternatives for cost/performance balance?

## Technical & Infrastructure Requirements

- Is your data pipeline ready (data prep, ETL, secure ingest into vector DB)?
- Do you require high availability, redundancy, or DR capabilities?
- Are you leveraging Ubuntu Linux / Kubernetes / Red Hat OpenShift AI stack support?
- Have you validated network performance (25–100GbE depending on workload)?
- Is storage provisioned with NetApp ONTAP + Trident CSI driver for Kubernetes persistence?
- Minimum 3 compute nodes (2 worker, 1 control plane) sized correctly?

## Software & Ecosystem Alignment

- Are your containerized apps compatible with Kubernetes & Helm charts provided?
- Which vector database(s) (e.g., Milvus, pgvector) are planned for deployment?
- Do you need OPEA (Open Platform for Enterprise AI) pre-integration for retrieval-augmented generation (RAG)?
- Are you leveraging hybrid-cloud options (Cloud Volumes ONTAP, FSxN, Anthos, Azure Arc)?
- Do you need partner ISV integrations (ESRI, healthcare PACS, financial/legal ISVs)?

## Data Governance & Security

- Have you enabled role-based access control (RBAC) in Kubernetes?
- Is there a data protection & backup plan (SnapMirror, SnapCenter, ransomware protection)?
- Have you mapped data compliance needs (HIPAA, GDPR, FedRAMP, CJIS)?
- Do you require private AI deployment (air-gapped, on-premises, secure enclave)?

## Operational & Support Considerations

- Are admins trained/enabled on Kubernetes, Trident CSI, and OPEA stack deployment?
- Do you need support for multi-tenancy (departments, SLED agencies, business units)?
- Is there a plan for monitoring & observability (ONTAP System Manager, Cloud Insights, Prometheus/Grafana)?
- Who will own day-2 operations (customer IT, partner, managed service provider)?

## Commercial & GTM Alignment

- Is there a phased roadmap (departmental → enterprise-wide AI expansion)?
- Do you have a multi-year pro forma projection (TCO, ARR, margin uplift)?
- Are licensing uplift scenarios clear (vector DB, ISV software, AI ops tools)?
- Have you explored partner incentives (distributor margin, OEM/Intel co-funding)?
- Is the purchase aligned to budget cycles (CapEx vs OpEx, consumption models)?
- Do you have a services partner (Arrow, WWT, Presidio, etc.) for sizing & deployment?

## Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

**LIMITED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.