Enterprise Computational Fluid Dynamics (CFD) Readiness Checklist

Sensitivity: Public

Company: DiGiCOR PTY LTD

Status: Published

Date Published: 2026-07

A Practical How-To + Enterprise Readiness Assessment

Infrastructure, Hardware, Networking, and Facility Provisioning Blueprint

Purpose: This guide helps organisations understand where to start, what information to prepare, and how to make a right decision when deploying CFD infrastructure.

Who this is for:

- Engineering managers, Principal / Lead Engineer
- Research & Development (R&D) Engineer
- IT infrastructure teams
- HPC administrators
- CFD simulation engineers, CFD Analyst, CFD Specialist / Consultant,
- Aerodynamicist
- Hydrodynamicist
- Thermal Management Engineer
- HVAC Engineer
- Mechanical Design Engineer
- Process / Chemical Engineer
- Turbomachinery Engineer
- Academic Researcher / Professor

Industries:

- Aerospace Engineering
- Automotive/transportation
- Biomedical and Healthcare maintaining a precise temperature range during storage
- Data Centers and Cooling the Cloud
- Energy Sector and Optimizing Power Generation
- Electronics Cooling and Keeping Critical Components Cool
- Heating, Ventilation, and Air Conditioning (HVAC)
- Manufacturing and Chemical Processing
- Industrial Applications
- HPC/Scientific Computing
- Turbomachinery

Contents

How-To: The Practical Path to Computational Fluid Dynamics (CFD)

Step 1: Define the Use Case and Success Criteria

Identify the primary workflow

- Accelerated Computing Tools & Techniques
- Simulation/modelling/design

Define success

- Model Efficiency (instant operation of CFD)
- Development Process Change (optimum use of CFD)
- Clearly Defined Operating Conditions
- Validation Confidence (reliable simulation results and test data)
- Performance Targets (simulation time, turnaround time and concurrency)
- Support for Complex Physics (such as multi-phase fluid, combustion and turbulence)

Step 2: Power Infrastructure for Data Centre

Data centre power infrastructure keeps IT equipment and support systems powered and protected for continuous operation. Power the data centre around:

- Rack power
- UPS and redundancy
- Cooling
- Equipment density
- Physical protection

Prepare the data centre for sustained, high-density CFD workloads:

- Size rack power for sustained CFD compute loads.
- Provide redundant power and UPS protection.
- Use hot/cold aisle containment or liquid cooling for high-density systems where needed
- Install clean-agent gaseous fire suppression systems (such as FM-200 or Novec 1230) to protect high-value computing equipment.

Step 3: Size the Compute Environment

Define the expected CFD workload, including model size, solver type, simulation frequency, concurrency, and future growth.

- CPU architecture and core count
- GPU acceleration requirement and solver compatibility
- Memory capacity and bandwidth
- NUMA topology
- Expected concurrent jobs
- Parallel scaling requirements

Configure each compute node for stable and efficient CFD performance:

- Populate memory channels symmetrically to maximise available memory bandwidth (e.g., 24 DIMMs total for dual 12-channel CPUs).
- Use ECC memory for long-running and business-critical CFD workloads.
- Select enterprise-grade GPUs that support precision and acceleration (such as NVIDIA H100/H200 or AMD Instinct lines).
- Review BIOS power setting, NUMA-aware configuration, and hyper-threading.

Step 4: Build the Cluster Networking

For CFD simulations, design the network around:

- Latency and bandwidth
- Inter-node communication
- MPI communication
- RDMA support
- User access network
- Redundancy and failover

Build the network to support tightly coupled CFD workloads and large data transfers:

- Use a high-speed interconnect (NVIDIA Quantum InfiniBand) where multi-node CFD scaling requires ultra-low-latency communication.
- Separate management traffic from simulation and storage traffic where appropriate.
- Ensure primary core data switches feature extensive hardware packet buffering can support multi-gigabyte engineering data transfers without packet loss.
- Validate MPI and RDMA performance with representative CFD workloads.

Step 5: Choose the Storage Architecture

Define how simulation inputs, outputs, temporary files, project data, and archives will be stored and accessed.

Plan storage around:

- Input and output file sizes
- Temporary simulation data
- Shared project storage
- Backup and archive requirements
- Data retention policies
- Expected storage growth

Set up storage for processing, solving, post-processing, collaboration, and long-term retention:

- Provide high-performance scratch storage for temporary solver data.
- Use resilient shared storage for CAD, mesh files, and post-processed datasets.
- Use a parallel file system (such as Lustre, Weka, or Ceph) for low-latency multi-node file ingestion.
- Provide a lower-cost archive tier for long-term or regulatory project retention.

Step 6: Create the software environment

For fast and accurate CFD workloads, a secure and reliable software environment should be planned around:

- Pipeline & data security
- Workload scheduling
- Infrastructure security

Build a controlled and repeatable platform for CFD users and workloads:

- Use a workload scheduler such as Slurm or OpenPBS for job queues, user priority fairness, and hardware resource caps.
- Use Apptainer (formerly Singularity) or Docker environments to guarantee immutable, repeatable code runtimes.
- Integrate with central identity services such as Active Directory, LDAP, or FreeIPA.
- Use automated provisioning tools ((such as Ansible or SaltStack) to maintain consistent compute-node configurations.

Step 7: Select Your Deployment Model (Bare Metal vs Virtualised HPC vs Hyperconverged Infrastructure)

Use **bare metal** when predictable, direct performance is the priority.

Use **virtualised HPC** when you need a flexible enterprise environment and easier management of shared resources.

Use **hyperconverged** infrastructure when your workloads follow a specific operating pattern or require a hybrid design.

Step 8: Observability and Operation

Once the CFD-based infrastructure is ready, the organisation can begin implementation. Treat CFD infrastructure as a performance-sensitive HPC deployment with defined risks to manage:

- Inconsistent hardware and software configurations
- Poor CPU, NUMA, and memory placement
- Network latency and inefficient MPI communication
- Storage throughput and capacity bottlenecks
- Resource retention
- Insufficient testing with representative CFD workloads

Enterprise CFD Readiness Checklist

Optional scoring:

Score each item as:

- 0 = Not started
- 1 = In progress
- 2 = Implemented

Total possible score:

A) Strategy, Scope and Success Metrics

Score	Statement
	☐ CFD Use cases clearly defined
	☐ Stakeholders and owners identified (IT, business)
	☐ Success metrics defined (quality, safety, user experience)
	☐ The components, products or systems to be simulated identified
	☐ Operating conditions, boundary conditions, and required materials identified
	☐ Data sensitivity and business impact assessed
	☐ Test data validated and recognised benchmarks defined
	☐ Deployment model decided (bare metal, virtualised HPC, Hyperconverged infrastructure)
	☐ Performance targets defined (runtime, turnaround time, model size and concurrency)
	☐ User training and enablement plan created
	Total

B) Data Centre and Power Infrastructure

Score	Statement
	☐ Power capacity has been assessed for CFD workload
	☐ Redundant power and UPS requirements defined
	☐ Cooling capacity has been assessed for the compute density
	☐ Rack space and future infrastructure expansion have been considered
	Total

C) Compute

Score	Statement
	☐ CFD workloads selected for infrastructure sizing
	☐ CPU architecture, core count, and clock-speed requirement assessed
	☐ Memory capacity and memory bandwidth assessed
	☐ GPU acceleration and CFD solver compatibility assessed
	☐ Expected numbers of simultaneous users and job concurrency defined
	☐ Single-node and multi-node scalability evaluated
	☐ NUMA topology, CPU affinity, memory locality, and MPI process placement
	☐ Redundant power supplies sized for peak sustained load
	☐ Thermal design validated for sustained utilisation
	☐ Hardware lifecycle and warranty plan defined
	Total

D) Networking

Score	Statement
	☐ The communication of pattern of the CFD workloads assessed
	☐ Network bandwidth and latency assessed
	☐ MPI and RDMA requirements assessed
	☐ The network topology between compute nodes and storage has been designed
	☐ Management, user-access, storage, and compute traffic defined
	☐ Network redundancy, resilience, and future expansion considered
	☐ Network monitoring configured
	Total

E) Storage

Score	Statement
	☐ CFD input, output, checkpoint data, and post-processing data volumes estimated
	☐ Shared project storage defined
	☐ Required storage throughput and data performance assessed
	☐ Storage capacity and expected data growth estimated
	☐ Backup, archive, data retention workflows implemented

Score	Statement
	☐ Data access, permissions, and recovery defined
	☐ Storage monitoring and alerting configured
	Total

F) Enterprise Software Environment & Security

Score	Statement
	☐ The required CFD software stack defined CFD input
	☐ A workload scheduler configured
	☐ Software licensing requirements confirmed
	☐ User authentication and access management planned
	☐ Monitoring, logging, and operational management tools defined
	Total

G) Model Strategy (Bare Metal vs Virtualised HPC vs Hyperconverged Infrastructure)

Score	Statement
	☐ Bare-metal HPC has been evaluated based on CFD performance
	☐ Virtualised HPC has been evaluated according to flexibility, isolation, and management
	☐ Hyperconverged infrastructure has been evaluated aligned with performance, scalability, and operational requirements
	☐ The effect of virtualisation on CPU, memory, networking, and storage performance assessed
	☐ The preferred deployment model has been selected with a clear rationale
	Total

H) Observability and Operation

Score	Statement
	☐ Monitoring covers CPU platforms
	☐ Monitoring covers GPU utilisation, memory, and queue depth
	☐ Monitoring covers result collection and checkpointing
	☐ Alert thresholds defined and tested
	☐ Incident runbooks created (model outage, retrieval failure, latency spikes)

Score	Statement
	☐ On-call and escalation defined
	☐ Capacity planning process established
	☐ Post-incident review process defined
	☐ Documentation and architecture diagrams maintained
	Total

Readiness Score Interpretation

Score	Statement
	A) Strategy, Scope and Success Metrics
	B) Data Centre Power Infrastructure
	C) Compute
	D) Networking
	E) Storage
	F) Software Environment & Security
	G) Model Strategy (Bare Metal vs Virtualised HPC vs Hyperconverged Infrastructure)
	H) Observability and Operations
	Total

Score bands (out of 114):

0–39 (Pilot Only)

Significant gaps across security, governance, or platform fundamentals. High risk in CFD workload.

40–69 (Scaling Risk)

Core platform exists, but governance, security, or operations need strengthening before rollout.

70–89 (Production Capable)

Strong baseline; validate via load tests, security tests, and operational drills.

90–114 (Enterprise Ready)

Mature readiness; focus on continuous improvement and evolving controls.

We Can Get You Started

DiGiCOR supports enterprise CFD infrastructure from design to deployment:

- The infrastructure platform that underpins a CFD environment
- End-to-end infrastructure planning across compute, storage, and networking, and licensing/software
- Robust, scalable and suitable infrastructure layer properly matched to the workload

Why Choose DiGiCOR?

- Local expertise across enterprise infrastructure, HPC, AI, and CFD workloads
- Vendor-agnostic solution design matched to your performance, budget, and growth requirements
- Purpose-built systems configured, integrated, and validated for real-world workloads
- Australian and New Zealand support from planning and procurement through to deployment and lifecycle management

Assess Your CFD Readiness

Ready to deploy CFD infrastructure?

Fill out the CFD Readiness Assessment

- Validate your CFD use case, performance targets, and rollout plan
- Identify infrastructure gaps across compute, networking, storage and deployment models
- Choose the right deployment model aligned to your specific need
- Design a CFD infrastructure based on current requirement and future growth and contact us to do the configuration for you.

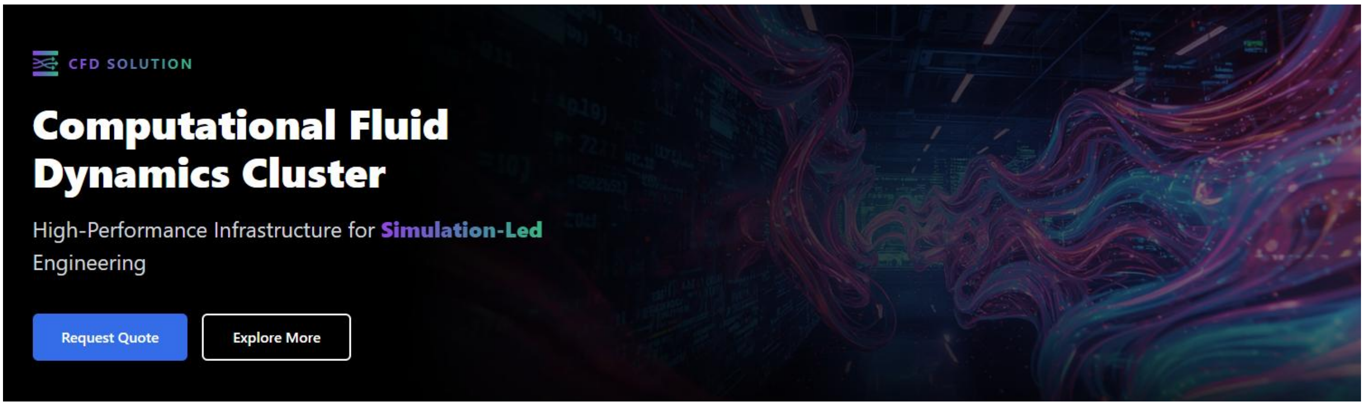

Disclaimer and Terms of Use

This document is provided for **general informational and educational purposes only** and does not constitute technical, legal, financial, or professional advice.

Implementation outcomes vary by workload, data, architecture, and operational constraints. Users should perform independent assessment and seek appropriate professional advice prior to acting on this guidance.