

Data Centre Power Management: A Practical Guide

Balancing Capacity, Redundancy, and Energy Efficiency

A professional reference guide for data centre managers, IT infrastructure professionals, facility engineers, and GSDC Certified Data Centre Professional candidates. This guide provides technically accurate, practical guidance on managing power requirements, evaluating infrastructure, selecting appropriate redundancy configurations, and improving energy performance without compromising operational reliability.

GSDC CDCP REFERENCE GUIDE

Understanding Data Centre Power Management

Data centre power management is the discipline of ensuring that electrical power is delivered reliably, efficiently, and in sufficient quantity to meet the demands of IT equipment and supporting infrastructure. It encompasses everything from the initial design of electrical systems to the ongoing monitoring of energy consumption and the planning of capacity for future workloads.

Reliable power is not simply a technical preference — it is a fundamental operational requirement. IT equipment depends on a continuous, stable supply of electricity to maintain data integrity, support business applications, and meet contractual availability obligations. Any interruption, however brief, can result in data corruption, service disruption, and significant financial or reputational consequences. For this reason, power management must be treated as a core operational discipline rather than a background concern.

The three central priorities of data centre power management — **capacity**, **redundancy**, and **efficiency** — are deeply interconnected. Increasing power capacity may introduce additional infrastructure that affects operational costs. Adding redundancy provides protection against failure but increases capital expenditure and can reduce overall energy efficiency. Improving efficiency, if poorly implemented, may inadvertently reduce headroom for redundancy or future growth. Effective power management requires balancing all three priorities simultaneously, with decisions in one area carefully considered against their impact on the others.

The challenge of managing data centre power requirements is growing. The adoption of cloud computing, artificial intelligence, and high-performance computing workloads is driving significant increases in rack power density. Many facilities designed for traditional enterprise IT are now facing demands that exceed their original design parameters. Understanding the fundamentals of power management is therefore essential for any professional responsible for data centre infrastructure.

Understanding Data Centre Power Infrastructure

Electricity reaches IT equipment through a series of infrastructure components, each performing a specific function in the power delivery chain. Understanding each component and its role is essential for assessing capacity, planning redundancy, and identifying efficiency improvements.

Utility Power Supply

The incoming electrical supply from the national or regional grid, typically delivered at high or medium voltage. Many data centres receive supply via two independent utility feeds to reduce the risk of a single point of failure at the utility level.

Transformers and Switchgear

Transformers step down incoming high-voltage supply to the operational voltage levels used within the facility. Switchgear controls, protects, and isolates electrical circuits, enabling safe switching between power sources and protecting equipment from fault conditions.

Uninterruptible Power Supply (UPS)

The UPS provides a continuous, conditioned power output to IT equipment, protecting against utility supply interruptions, voltage fluctuations, and power quality issues. It bridges the gap between a utility failure and generator start-up, typically using battery energy storage.

Backup Generators

Diesel or gas-powered generators provide sustained electrical power during prolonged utility outages. Once started and synchronised, they assume the facility load, allowing UPS batteries to recharge and normal operations to continue.

Power Distribution Units (PDUs)

PDUs distribute power from the main electrical infrastructure to individual server cabinets. Floor-standing PDUs serve rows of racks, whilst rack-mounted PDUs distribute power to individual servers, storage devices, and networking equipment within each cabinet.

Power Monitoring Systems

Intelligent power monitoring provides real-time visibility of power consumption at the facility, row, rack, and device level. This data supports capacity planning, efficiency analysis, and early detection of abnormal power conditions.

Power Capacity Planning

Power capacity planning is the process of determining how much electrical capacity a data centre requires to support its current IT load, provide sufficient headroom for growth, and accommodate the demands of future workloads. Poor capacity planning results either in infrastructure that cannot support operational requirements or in costly over-provisioning that wastes capital expenditure.

Key Capacity Planning Inputs

Current IT Power Load: The measured or estimated electrical demand of all active IT equipment, including servers, storage, and networking devices. Measured in kilowatts (kW), this forms the baseline for all capacity calculations.

Peak Power Demand: The maximum simultaneous power draw observed or anticipated during periods of highest workload. Infrastructure must be sized to accommodate peak demand, not average consumption.

Rack Power Density: The average and maximum power consumed per rack, typically expressed in kW per rack. Traditional enterprise racks may consume 3–10 kW, whilst high-density and AI-focused racks may require 20 kW or considerably more.

Future Workload Growth: An estimate of how IT power requirements will change over a defined planning horizon, typically three to five years. This must account for new services, increased utilisation, and hardware refresh cycles.

AI and High-Performance Computing: AI training and inference workloads, powered by graphics processing units (GPUs) and accelerator hardware, introduce power requirements significantly higher than those of traditional server infrastructure. Facilities planning to support these workloads must assess whether existing electrical and cooling infrastructure can accommodate the resulting density increases.

Practical Capacity Example

A data centre operates 200 racks with an average measured load of 6 kW per rack.

Current IT Load: 200 × 6 kW = 1,200 kW

Peak Demand Allowance (+15%): 1,200 × 1.15 = 1,380 kW

Three-Year Growth Forecast (+25%): 1,380 × 1.25 = 1,725 kW

Planned AI Rack Addition: 20 racks × 25 kW = 500 kW

Total Estimated Requirement: ~2,225 kW

This figure informs decisions about UPS capacity, generator sizing, and incoming supply requirements — and must be compared against available electrical infrastructure capacity.

Available electrical capacity must always exceed the total estimated power requirement, with headroom retained for redundancy systems and future expansion. Capacity that appears sufficient today may become a constraint within a short planning horizon as workload density increases.

Power Redundancy and Availability

Power redundancy is the practice of providing additional capacity, duplicate components, or alternative power paths beyond the minimum required to operate under normal conditions. Its purpose is to ensure that the failure of a single component, power source, or distribution path does not interrupt the supply of electricity to IT equipment.

It is important to distinguish between two forms of redundancy. **Component redundancy** involves duplicating individual items of equipment — for example, providing two UPS modules where one would be sufficient — so that the failure of one unit does not affect the overall system. **Path redundancy** involves providing entirely independent electrical paths from the power source through to IT equipment, ensuring that a fault in one complete path does not affect power delivery through an alternative route. The highest levels of availability typically require both component and path redundancy.

The relationship between redundancy, availability, and infrastructure cost is a critical consideration in data centre design. Greater redundancy generally provides higher tolerance for equipment failures and reduces the probability of a power-related outage reaching IT equipment. However, additional redundancy requires additional capital investment in electrical infrastructure, physical space, and ongoing maintenance. It also typically results in lower average utilisation of installed capacity, since equipment is intentionally oversized relative to normal operating loads.

Redundancy configurations are described using standardised notation. The letter **N** represents the capacity required to support the full IT load under normal operating conditions. Additional capacity above N is expressed as follows:

N Configuration

Exactly the capacity required. No redundancy beyond the minimum needed to operate.

N+1 Configuration

One additional component or unit beyond what is strictly required to carry the full load.

2N Configuration

Full duplication — two complete, independent systems each capable of supporting the entire IT load.

2N+1 Configuration

Two complete systems plus one additional unit — the highest level of redundancy in common practice.

Comparing Power Redundancy Configurations

Selecting the appropriate redundancy configuration requires careful evaluation of workload criticality, available budget, infrastructure complexity, and operational requirements. The table below provides a detailed comparison of the four principal configurations used in data centre power design.

Criteria	N	N+1	2N	2N+1
Architecture	Single path, no spare capacity	Single path, one spare component	Two fully independent paths	Two independent paths plus one spare unit
Backup Capacity	None	One additional unit per system	Full second system (100% redundant)	Full second system plus one additional unit
Single Failure Impact	May cause service disruption	System remains operational; redundancy is consumed	One path takes full load; operations continue	One path takes full load; spare unit remains available
Availability Level	Basic — suitable for non-critical workloads	Moderate — suitable for important workloads with some tolerance for planned maintenance outages	High — supports concurrent maintainability	Very high — supports concurrent maintainability with additional protection
Infrastructure Cost	Lowest	Moderate	High — requires duplication of all systems	Highest — full duplication plus additional units
Complexity	Simple	Moderate	Complex — dual bus, independent feeds to each rack	Most complex — requires rigorous change management
Typical Use Case	Development environments, low-criticality workloads	Standard enterprise data centres	Mission-critical operations, financial services, tier-sensitive facilities	Highest criticality environments where any service interruption is unacceptable

It is important to note that the redundancy configuration of the power system must be considered alongside the dual power supply capabilities of IT equipment itself. Servers equipped with dual power supply units can connect to two independent power paths, fully exploiting the protection offered by 2N or 2N+1 designs. Single-corded equipment connected to a 2N system retains only the protection offered to that single connection point.

The selection of a redundancy configuration should be driven primarily by workload criticality and recovery time requirements, not solely by budget constraints. Oversizing redundancy beyond operational requirements wastes capital; undersizing creates unacceptable availability risk.

Improving Data Centre Energy Efficiency

Energy efficiency in data centre operations refers to the degree to which electrical power consumed by the facility is used to perform productive IT work, rather than being consumed by supporting infrastructure such as cooling, power conversion, and lighting. Improving efficiency reduces operating costs and environmental impact without requiring a reduction in IT capability or service reliability.



UPS Efficiency

Modern UPS systems operating in eco or high-efficiency modes can significantly reduce conversion losses compared to older double-conversion designs. UPS efficiency should be evaluated at typical operating loads, not only at full rated capacity, as efficiency often varies substantially across the load range.



Transformer and Distribution Efficiency

Transformer losses contribute to overall facility energy consumption. Specifying transformers with low no-load and full-load losses, and operating distribution systems at appropriate load levels, reduces wasted energy in the power delivery chain.



Cooling Optimisation

Cooling typically represents the largest source of non-IT energy consumption in a data centre. Hot aisle and cold aisle containment, economiser modes, and correct airflow management reduce the energy required to maintain IT equipment within acceptable operating temperature ranges.



Server Consolidation and Virtualisation

Underutilised physical servers consume power disproportionate to their productive output. Virtualisation consolidates workloads onto fewer, better-utilised physical hosts, reducing the number of powered-on devices and the associated power and cooling requirements.



Continuous Energy Monitoring

Real-time power monitoring at the facility, row, rack, and device level enables the identification of inefficiencies, abnormal consumption, and opportunities for improvement. Without accurate measurement, energy management decisions are based on assumptions rather than evidence.

Understanding Power Usage Effectiveness (PUE)

Power Usage Effectiveness (PUE) is the most widely used metric for measuring the energy efficiency of a data centre's supporting infrastructure. It expresses the relationship between the total energy consumed by the entire facility and the energy consumed specifically by IT equipment.

The PUE Formula

$$\text{PUE} = \frac{\text{Total Data Centre Energy (kW)}}{\text{IT Equipment Energy (kW)}}$$

Worked Example

Total data centre energy consumption: **1,500 kW**

IT equipment energy consumption: **1,000 kW**

$$\text{PUE} = \frac{1{,}500}{1{,}000} = 1.5$$

This result indicates that for every 1 kW consumed by IT equipment, an additional 0.5 kW is consumed by supporting infrastructure — cooling, power conversion, lighting, and other ancillary systems.

Interpreting PUE Values

A PUE of **1.0** would represent a theoretically perfect facility where all consumed energy is used directly by IT equipment — in practice, this is not achievable. A PUE of **1.2 or below** is generally considered indicative of a highly efficient facility. Many established data centres operate with PUE values between 1.5 and 2.0, with older facilities sometimes exceeding 2.0.

A lower PUE generally indicates that a smaller proportion of total facility energy is being consumed by supporting infrastructure relative to IT equipment. However, it does not indicate that the IT equipment itself is operating efficiently — a facility with highly inefficient servers may still report a low PUE if its cooling and power conversion systems are well-optimised.

Limitations of PUE

PUE should not be used in isolation. It does not capture the reliability or availability of power delivery, the utilisation of installed IT capacity, or the efficiency of IT equipment itself. A data centre that achieves a low PUE by reducing redundancy may expose operations to unacceptable availability risk. PUE is most useful when considered alongside capacity utilisation metrics and availability performance data.

1/5

Excellent Efficiency

PUE of 1.2 or below — highly optimised facilities with advanced cooling and power systems.

2/5

Good Efficiency

PUE between 1.2 and 1.5 — well-managed facilities with effective cooling optimisation.

3/5

Average Efficiency

PUE between 1.5 and 2.0 — typical of many established enterprise data centres.

5/5

Poor Efficiency

PUE above 2.0 — older facilities with significant opportunities for improvement.

Balancing Capacity, Redundancy, and Efficiency

The central challenge of data centre power management is that capacity, redundancy, and efficiency do not operate independently — decisions in any one area create consequences in the other two. Understanding these interdependencies is essential for making infrastructure decisions that are technically sound and operationally sustainable.

How Capacity Decisions Affect Redundancy and Efficiency

Increasing IT power capacity requires additional electrical infrastructure — larger UPS systems, additional generator capacity, upgraded distribution equipment, and enhanced cooling. If redundancy configurations are maintained at their existing levels, the absolute cost of that redundancy increases proportionally. Facilities operating near the limits of their electrical capacity may also find their options for adding redundant capacity constrained by physical space, incoming supply, or structural limitations.

Equipment oversizing — installing electrical infrastructure significantly beyond current and near-term requirements — reduces average utilisation. Many electrical components, particularly UPS systems and transformers, operate less efficiently at very low load levels. A facility that installs capacity well in excess of its requirements may find its PUE adversely affected until utilisation levels increase sufficiently to bring operating loads into the efficient range.

How Redundancy Affects Costs and Utilisation

Moving from N+1 to 2N redundancy approximately doubles the installed electrical infrastructure. During normal operations, each redundant path typically carries only half the total load, since both paths are active simultaneously. This lower utilisation level affects efficiency and increases the capital cost per unit of productive IT capacity supported. The decision to invest in higher redundancy must therefore be justified by the criticality and recovery requirements of the workloads being supported.

Selecting Redundancy Based on Workload Criticality

Not all workloads within a data centre require the same level of power protection. A tiered approach — applying higher redundancy to mission-critical systems and lower redundancy to development or test environments — allows facilities to allocate infrastructure investment where it delivers the greatest operational value. This approach requires careful zone planning within the facility and clear classification of workload criticality.

Planning for Future Growth Without Overspending

Maintaining sufficient capacity headroom for future growth is a legitimate and necessary planning objective, but excessive headroom introduces both capital cost and efficiency penalties. A structured capacity planning process, refreshed on a regular cycle and informed by accurate consumption monitoring, enables facilities to maintain an appropriate balance between current operational requirements, near-term growth, and long-term strategic flexibility. Modular infrastructure designs allow capacity to be added incrementally as demand increases, reducing the need for large speculative investments.

Power Distribution and Infrastructure Design

Designing a power distribution system that reliably serves current workloads whilst accommodating future expansion requires a systematic approach to each stage of the electrical delivery chain. Each component must be sized, protected, and maintained to support both normal operating conditions and foreseeable failure scenarios.

1 Incoming Electrical Supply

The incoming supply from the utility network should be received at an appropriate voltage for the facility scale and stepped down through dedicated transformers. Where the criticality of operations justifies it, two independent utility feeds entering the facility via separate routes and terminating on separate switchgear boards provide protection against a single point of failure at the utility interface.

2 UPS Systems and Generator Integration

UPS systems must be sized to support the full IT load during the interval between utility failure and generator assumption of the load. Generator capacity must be sufficient to supply the entire facility load — including UPS recharge requirements — for the anticipated duration of a utility outage. Transfer switching arrangements must be tested to confirm correct and timely operation. UPS battery autonomy should be validated against realistic load conditions rather than manufacturer-rated values at idealised conditions.

3 Busways, PDUs, and Rack Distribution

Overhead busway systems provide a flexible and scalable method for distributing power to rows of server cabinets, allowing tap-off points to be added or relocated as the floor layout evolves. Floor-standing PDUs aggregate power for rack rows and typically include metering capabilities. Rack-mounted PDUs distribute power to individual devices and, in intelligent variants, provide outlet-level metering and remote switching. For 2N facilities, each rack must receive two independent feeds from separate PDUs connected to separate power paths.

4 Circuit Load Management and Maintenance Paths

Circuits should be loaded to no more than 80% of rated capacity under normal operating conditions, retaining headroom for load increases and ensuring that protective devices do not operate at their limits during transient peaks. The design must identify maintenance paths — the electrical configurations that allow individual components to be isolated for maintenance whilst keeping the IT load energised through alternative routes. Maintenance paths must be documented, tested, and kept clear of unauthorised modifications.

AI Workloads and Changing Power Requirements

The adoption of artificial intelligence infrastructure is fundamentally changing the power management landscape in data centres. AI training and inference workloads, powered by high-performance GPU clusters and dedicated accelerator hardware, introduce power and cooling requirements that are qualitatively different from those of traditional enterprise IT — not simply greater in degree, but different in kind, requiring a reassessment of infrastructure design assumptions.

High-Density AI Rack Requirements

Traditional enterprise server racks typically operate within a moderate power density range that most established data centres can accommodate with standard air-cooling infrastructure and conventional power distribution. AI server racks housing GPU-accelerated hardware may require substantially greater power per rack — in some configurations, an order of magnitude more than a traditional rack. This density concentration creates challenges for power distribution, cooling delivery, and structural loading that must be addressed in the facility design before AI equipment is deployed.

GPU and accelerator components have high and relatively stable power draw characteristics. Unlike traditional CPU-based servers whose power consumption varies significantly with computational load, GPU workloads during training phases can sustain near-peak power consumption for extended periods. This affects the sizing of electrical infrastructure, since the peak-to-average ratio that engineers might apply for traditional workloads may not be appropriate for sustained AI compute jobs.

Cooling and Infrastructure Implications

The concentrated heat output of high-density AI racks often exceeds the capacity of conventional raised-floor air cooling to manage effectively. Direct liquid cooling, rear-door heat exchangers, or immersion cooling solutions are increasingly being evaluated and deployed alongside or in place of air-based cooling for the highest-density installations. Each cooling technology has implications for power consumption, water usage, infrastructure complexity, and maintenance requirements.

Power demand forecasting for AI workloads requires a different approach to that used for traditional IT. The scheduling of AI training jobs can create significant demand spikes across a cluster, and the aggregated power draw of a large GPU installation may impose rapid load changes on electrical infrastructure. Facilities planning to support AI workloads should work closely with their electrical engineers to validate that UPS systems, generators, and distribution equipment can respond appropriately to these dynamic load profiles.

AI-Supported Operations

Beyond serving as a subject of infrastructure planning, AI-based tools are increasingly being applied to data centre operations themselves — supporting predictive maintenance of power and cooling equipment, optimising cooling setpoints in response to real-time load conditions, and identifying anomalous power consumption patterns that may indicate equipment degradation or configuration issues.

Practical Case Study: Planning for AI Infrastructure Expansion

> The following scenario is entirely hypothetical and is presented for illustrative purposes only. It is intended to demonstrate a structured decision-making process rather than to represent the outcomes of any specific facility.

A mid-sized enterprise data centre currently supports conventional server and storage workloads across a fully equipped computer room. Its management team has been asked to plan for the introduction of a GPU-accelerated AI computing cluster, whilst maintaining existing service levels and controlling operating costs.

01

Step 1: Assess Existing Power Consumption and Capacity

The facility team conducts a metered audit of current IT power consumption at the rack level, using intelligent PDU data and facility management system records. They identify the available headroom between current load and the rated capacity of the installed UPS, generator, and distribution infrastructure. They also identify any circuits or distribution paths already operating near their rated limits.

02

Step 2: Forecast Additional Power Requirements

Working with the IT team, the facility team determines the power draw characteristics of the planned AI hardware. They calculate the total additional load that the GPU cluster will impose, accounting for sustained near-peak consumption during training workloads, and assess whether the existing incoming supply and infrastructure can accommodate this addition without exceeding safe loading thresholds.

03

Step 3: Evaluate Redundancy Options

The team evaluates whether the AI cluster should be served by the existing N+1 power infrastructure or whether the criticality of the new workload justifies upgrading a portion of the facility to 2N. They consider the cost and complexity of each option, the recovery time requirements of the AI workloads, and the physical feasibility of routing independent power paths to the designated AI zone.

04

Step 4: Review UPS, Generator, and Distribution Capacity

UPS battery autonomy is validated against the new combined load. Generator fuel capacity and transfer switching arrangements are reviewed to confirm that sustained outage coverage remains adequate. PDU and busway capacity in the AI zone is assessed, and high-current rack PDUs with appropriate circuit protection are specified for the GPU hardware.

05

Step 5: Identify Cooling and Efficiency Improvements

The team identifies that standard air cooling is insufficient for the planned GPU rack densities. They evaluate rear-door heat exchanger options for the AI zone and assess the impact on facility PUE. Opportunities to improve cooling efficiency in the existing areas of the facility are also identified, partially offsetting the efficiency impact of the new high-density zone.

06

Steps 6 & 7: Implement Monitoring and Review Performance

Intelligent power monitoring is extended to the new AI zone, providing outlet-level visibility of GPU power consumption. Following expansion, the team reviews actual consumption against forecast, validates that redundancy configurations are performing as designed, and updates the facility capacity model to reflect the new baseline. PUE is recalculated to establish a revised efficiency benchmark for ongoing management.

Data Centre Power Management Assessment Checklist

The following checklist provides a practical framework for evaluating power management practices within a data centre facility. It is intended as a structured self-assessment tool for data centre managers, facility engineers, and infrastructure professionals reviewing their operational readiness.

Capacity Assessment

- [ ] Have current IT power loads been measured at the rack level using metered PDU or monitoring system data?
- [ ] Has peak power demand been recorded and accounted for in infrastructure sizing?
- [ ] Is rack power density being monitored and compared against infrastructure design limits?
- [ ] Have future workload requirements been estimated for a defined planning horizon?
- [ ] Is sufficient electrical headroom available within existing UPS, generator, and distribution capacity to accommodate planned growth?
- [ ] Have the power requirements of any planned AI or high-density workloads been specifically assessed?

Efficiency Assessment

- [ ] Is PUE being calculated and monitored on a regular basis using metered data?
- [ ] Are UPS systems operating at load levels within their efficient operating range?
- [ ] Are cooling systems optimised with hot/cold aisle containment and appropriate setpoints?
- [ ] Have underutilised servers been identified and considered for consolidation or decommissioning?
- [ ] Is energy consumption reviewed regularly to identify trends, anomalies, and improvement opportunities?
- [ ] Are transformer and distribution losses understood and factored into efficiency assessments?

Redundancy Assessment

- [ ] Is the redundancy configuration of power infrastructure clearly documented and understood by the operations team?
- [ ] Are UPS system capacity and battery autonomy validated against current IT load conditions?
- [ ] Are generator capacity and fuel reserves sufficient for the anticipated duration of a utility outage?
- [ ] Are independent power paths maintained to dual-corded IT equipment throughout the facility?
- [ ] Are backup systems tested under realistic load conditions on a defined schedule?
- [ ] Are maintenance switching procedures documented, reviewed, and practised?

Monitoring and Maintenance

- [ ] Are power monitoring systems in place at the facility, row, rack, and device level?
- [ ] Are equipment alarms and threshold alerts configured, reviewed, and acted upon promptly?
- [ ] Are maintenance activities documented with pre- and post-maintenance capacity and load verification?
- [ ] Are infrastructure upgrades followed by a formal review of capacity, redundancy, and efficiency status?
- [ ] Is there a defined process for updating the facility capacity model as workloads change?

Conclusion and Professional Development

Effective data centre power management is not a single technical discipline — it is the continuous balancing of three interdependent operational priorities. The decisions made in each area shape the reliability, cost, and sustainability of the facility as a whole, and must be revisited regularly as workloads, technologies, and business requirements evolve.

Capacity Planning

Ensures that sufficient electrical power is available to support current IT loads, peak demand, and foreseeable future workload growth — including the demands of high-density and AI-accelerated infrastructure.

Redundancy

Protects critical operations against the consequences of power equipment failures and supply interruptions, with the level of redundancy selected according to the criticality and recovery requirements of the workloads being supported.

Energy Efficiency

Helps control electricity consumption and operating costs through optimised UPS operation, cooling system management, server consolidation, and continuous energy monitoring — without compromising service reliability.

Continuous Monitoring

Provides the operational visibility needed to maintain the right balance as workloads change, identify inefficiencies, detect abnormal conditions, and inform informed infrastructure decisions.

Build Your Data Centre Infrastructure Knowledge With GSDC

The **GSDC Certified Data Centre Professional (CDCP)** certification is designed for professionals who wish to develop their knowledge of data centre infrastructure, operations, power management, and facility practices. The programme provides a structured framework for understanding how data centres are designed, operated, and optimised — covering the technical disciplines addressed throughout this guide, including power systems, cooling infrastructure, availability tiers, and operational best practices.

Whether you are preparing for a career in data centre operations, seeking to validate existing professional knowledge, or working towards a recognised industry qualification, the GSDC CDCP provides a credible and practical foundation for data centre infrastructure professionals.

Explore the GSDC CDCP Certification

Website: **https://www.gsdcouncil.org/certified-data-center-professional**

What the GSDC CDCP Covers

- Data centre infrastructure design and components
- Power systems, UPS, generators, and distribution
- Cooling systems and thermal management
- Redundancy configurations and availability tiers
- Energy efficiency and PUE measurement
- Capacity planning and workload management
- Data centre operations and maintenance practices
- Physical security and environmental controls

GSDC
Global Skill Development Council

Globally Recognised Certification Body
190+ countries · 2,50,000+ Certified Professionals

IAS
ACCREDITED
Personnel Certification Body

Certified Data Centre Professional



ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org